On Fat Partioning, Fat Covering and the Union Size of Polygons

M. van Kreveld

RUU-CS-93-36
November 1993

Utrecht University

Department of Computer Science

Padualaan 14, P.O. Box 80.089,
3508 TB Utrecht, The Netherlands,
Tel. : ... + 31 - 30 - 531454

On Fat Partioning, Fat Covering and the Union Size of Polygons

M. van Kreveld

Technical Report RUU-CS-93-36
November 1993

Department of Computer Science
Utrecht University
P.O.Box 80.089
3508 TB Utrecht
The Netherlands

ISSN: 0924-3275

On Fat Partitioning, Fat Covering and the Union Size of Polygons*

Marc van Kreveld

Abstract

The complexity of the contour of the union of simple polygons with n vertices in total can be $O(n^2)$ in general. A notion of fatness for simple polygons is introduced, which extends most of the existing fatness definitions. It is proved that a set of fat polygons with n vertices in total has union complexity is $O(n \log\log n)$, which is a generalization of a similar result for fat triangles [19]. Applications to several basic problems in computational geometry are given, such as efficient hidden surface removal, motion planning, injection molding, etc. The result is based on a new method to partition a fat simple polygon P with n vertices into $O(n)$ fat convex quadrilaterals, and a method to cover (but not partition) a fat convex quadrilateral with $O(1)$ fat triangles. The maximum overlap of the triangles at any point is two, which is optimal for any covering of a fat simple polygon by a linear number of fat triangles.

1 Introduction

The primary motivation of this research is to determine for what sets of geometric objects (regions bounded by Jordan arcs), the contour of the union has small complexity. When the union size is small, many geometric problems can be solved more efficiently and with simpler algorithms than in the general case.

Upper bounds on the union size have been found for several types of objects. Kedem et al.[13] show that the contour of the union of a set of n pseudo-discs in the plane has linear description size (a set of pseudo-discs is a set of simply connected regions of which any two boundaries intersect at most twice). It is easy to see that the contour of the union of a set of n isothetic rectangles can have $\Omega(n^2)$ connected components, and therefore quadratic description size, by placing them in a grid-like pattern. Since two isothetic rectangles intersect at most four times, the question arises what the maximum union size is of sets of unbounded regions of which every two boundaries intersect at most three

*The research of the author was supported by an NSERC international fellowship. Current address: Department of Computer Science, Utrecht University, P.O.Box 80.089, 3508 TB Utrecht, the Netherlands.

times. This case was settled by Edelsbrunner et al.[11], who show that the contour size is $O(n\alpha(n))$, and there are $O(n\alpha(n))$ connected components in the contour (where $\alpha(n)$ is the extremely slowly growing functional inverse of Ackermann's function). These bounds are tight in the worst case.

Recently, computational geometers have become interested in so-called *fat* objects. Well-known geometric problems can be reconsidered for cases where the given objects or subdivision satisfy a certain fatness condition, and more efficient, simpler algorithms can often be obtained [3, 6, 10, 12, 21, 25]. Fat objects are important in practice, since generally one does not deal with objects that are very thin. With respect to the contour size, Matoušek et al.[19] observed that for triangles, a quadratic lower bound example can only be constructed if the triangles have sharp angles. They proved that for a set of triangles of which any angle is at least δ, for some constant $\delta > 0$, the union determines only $O(n)$ holes, and the contour size is $O(n \log\log n)$. Notice that two such triangles can intersect six times.

In this paper we extend the results from [19] to the case of simple polygons. The fatness condition that each angle is bounded from below by a constant clearly is not good enough, because the lower bound example with rectangles still holds. To obtain a necessary and sufficient condition to bound the union size of simple polygons, we make the following definitions (see Figure 1):

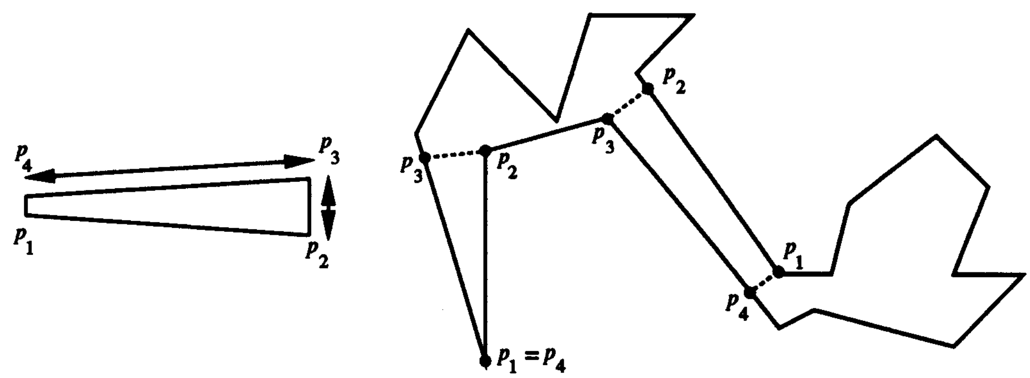

Figure 1: Example of a δ-corridor and a δ-wide simple polygon with two corridors indicated.

Definition 1 *For any* $0 < \delta \leq 1$, *a* δ-corridor *is a convex quadrilateral* Q *with vertices* p_1, p_2, p_3, p_4 *such that* $\angle p_1p_2p_3 = \angle p_2p_3p_4$ *and* $\angle p_3p_4p_1 = \angle p_4p_1p_2$, *and* $|\overline{p_1p_2}| = |\overline{p_3p_4}| = \frac{1}{\delta} \cdot \max\{|\overline{p_2p_3}|, |\overline{p_1p_4}|\}$.

For any $0 < \delta \leq 1$, *a simple polygon* P *(or any set of edges) is* δ-wide *if for any two edges* e *and* e' *of* P, *and any four points* $p_1, p_2 \in e$ *and* $p_3, p_4 \in e'$ *that are the vertices of a* γ*-corridor* Q *such that* interior$(Q) \subseteq$ interior(P), *it follows that* $\gamma \geq \delta$.

If P contains four vertices as in the definition, then we state that P *contains a γ-corridor.* If δ is a constant, we refer to P as a *wide polygon.* If P is δ-wide, then the minimum angle at any vertex of P is at least $2\arcsin(\delta/2) \geq \delta$ radians.

Several definitions of fatness have been used in other papers, but most of them do not apply to simple polygons. We mention three that are equivalent up to a constant factor to our definition of wideness, and the objects to which they apply: (i) for triangles, if every angle is at least a constant [19], (ii) for convex polygons, if the ratio of the radii of the maximum inscribed circle and the minimum enclosing circle is at least a constant [6], (iii) for convex polygons, if the ratio of the width and diameter is at least a constant. A fourth definition, for fatness of simple polygons and not equivalent to ours, is given in [21, 25]. This paper deals with fatness according to Definition 1 above, and the results also hold for fatness according to the three equivalent definitions just mentioned. To avoid confusion, we use the term *wide* for our definition.

Theorem 1 *Let S be a set of δ-wide simple polygons with n vertices in total. The contour of the union of the polygons in S has complexity $O((n\log\log n)/\delta^3)$.* [1]

The bound generalizes the similar bound for δ-fat (or δ-wide) triangles proved in [19]. The best known lower bound on the contour size is $\Omega(n/\delta + n\alpha(n))$, so for constant δ the upper bound is close to optimal ($\alpha(n)$ is the extremely slowly growing functional inverse of Ackermann's function).

The method we use to obtain the bound on the union size is interesting in its own right. If we partition a wide simple polygon into $O(n)$ wide triangles, then the result would follow immediately from the work of Matoušek et al.[19] mentioned before. Unfortunately, such a partitioning does not always exist. However, we can show instead that a wide simple polygon can be partitioned into $O(n)$ wide quadrilaterals, and also that a wide quadrilateral can be covered using $O(1)$ wide triangles. We make this more precise.

Definition 2 *A set S of quadrilaterals and triangles (only triangles) is a* weak Steiner quadrilateralization *(*weak Steiner triangulation*) of a simple polygon P if S is a partitioning of P. The set S is a* strong Steiner quadrilateralization *(*strong Steiner triangulation*) if additionally, no edge of any quadrilateral or triangle contains Steiner points in its interior. A set S of triangles is a* k-covering *of P if any point in the interior of P lies in the interior of at most k triangles of S, and the union of the triangles in S is P.*

We consider what simple polygons admit partitionings and coverings using wide quadrilaterals and triangles, using a set S of small size. It is easy to see that a rectangle with edge lengths 1 and m cannot be triangulated or k-covered with a set S of $O(1)$ wide triangles when m is large, see Figure 2. The rectangle is not wide, and therefore, we consider partitionings and coverings of δ-wide polygons in this paper. Also, a *wide* convex quadrilateral exists that cannot be partitioned into $O(1)$ wide triangles, see Figure 2 (we will show that $\Omega(\log m)$ wide triangles are needed).

[1] A previous version of this paper [14] claimed erroneously that the bound is $O((n\log\log n)/\delta)$.

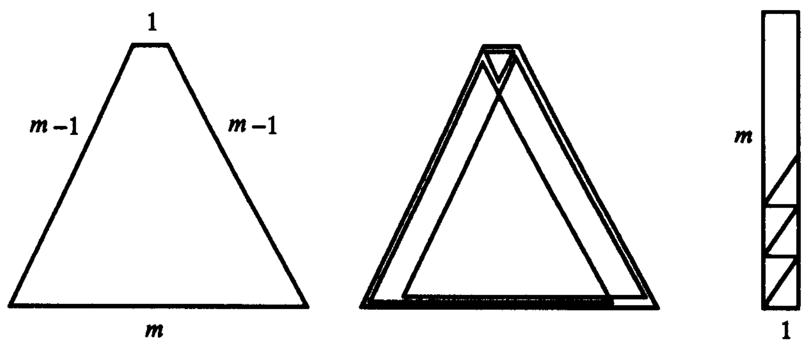

Figure 2: The leftmost wide quadrilateral cannot be partitioned into $O(1)$ wide triangles, but it can be covered by 3 wide triangles. The rightmost non-wide rectangle cannot be partitioned or covered by $O(1)$ wide triangles.

Let P be any δ-wide simple polygon, and let $\gamma = \min\{\delta, 1 - \frac{1}{2}\sqrt{3}\}$. (here we do not assume that δ is a constant). We show that P admits a weak Steiner quadrilateralization using $O(n)$ γ-wide quadrilaterals and triangles. Furthermore, we show that any δ-wide quadrilateral can be 2-covered using $O(1)$ γ'-wide triangles, where $\gamma' = \min\{\delta/4, \arctan(\pi/10)\}$. Consequently, P can be 2-covered using $O(n)$ γ''-wide triangles, where $\gamma'' = \min\{\gamma, \gamma'\} = \min\{\delta/4, 1 - \frac{1}{2}\sqrt{3}\}$. Thus, wide polygons do not admit a wide triangulation of linear size: bounds must depend on the ratio of edge lengths of the polygon as well. However, wide quadrilateralizations and wide triangle 2-coverings of linear size do exist for every wide polygon.

In the remainder of the paper, any polygon is simple by default. Furthermore, we let a quadrilateral be a *convex* 4-gon, and all angles are given in radians.

This paper is organized as follows. Section 2 shows that not every wide quadrilateral can be triangulated using a bounded number of wide triangles. In Section 3 we show that any simple polygon P with n vertices can be partitioned in $O(n \log^2 n)$ time into $O(n)$ quadrilaterals and triangles, such that no δ-corridor with $\delta < 1 - \frac{1}{2}\sqrt{3}$ is created. Consequently, a δ-wide polygon can be partitioned into $O(n)$ γ-wide quadrilaterals and triangles, where $\gamma = \min\{\delta, 1 - \frac{1}{2}\sqrt{3}\}$. The main construction tool is the *Edge Voronoi Diagram.* We show in Section 4 that a δ-wide quadrilateral can be 2-covered with $O(1)$ γ-wide triangles, where $1/\gamma = \Theta(1/\delta)$. Section 5 shows that the contour size of the union of δ-wide simple polygons with n vertices in total is $O((n \log \log n)/\delta^3)$. Applications are also given in this section. We close with the conclusions and open problems in Section 6.

2 Wide polygons cannot be partitioned into a bounded number of wide triangles

In this section we give an example of a wide quadrilateral that cannot be triangulated using a constant number of wide triangles. This immediately shows that wide quadrilateralizations are sometimes possible when wide triangulations of bounded size are not. A question that arises is whether wide partitionings into e.g. pentagons or hexagons are possible for a more general class of polygons than wide quadrilateralizations. In the next section it appears that this is not the case.

Intuitively, wide triangulations of bounded size are not always possible for the following reason. A wide triangle that has one short edge must necessarily have three short edges, and thus have a small area. This is not true for wide quadrilaterals, pentagons or hexagons, which provides some intuition why wide quadrilateralizations are possible whenever any partitioning into wide polygons of constant size is possible.

Theorem 2 *Not every wide convex quadrilateral can be partitioned into $O(1)$ wide triangles.*

Proof: Consider an equilateral triangle with edge length m, and truncate it by removing an equilateral triangle with edge length 1 from one of its vertices. Let Q be the resulting trapezoid, with edge lengths m, $m-1$, 1, $m-1$, see Figure 3. If $m \geq 2$, then Q is clearly wide.

Let $\mathcal{T}$ be any triangulation of Q into δ-wide triangles, for some constant $0 < \delta \leq 1$. We show that $\mathcal{T}$ consists of $\Omega(\log m)$ triangles. Position Q so that the edge denoted *base* with length m and the edge denoted *top* with length 1 are horizontal, with *top* above *base*. The other two edges of Q are called the *left side* and *right side*. We construct a sequence of small triangles that must be present in $\mathcal{T}$ by using an *active boundary*, which is a polygonal chain that connects *left side* to *right side* and has all currently chosen small triangles above it. "above" refers to being in the polygon bounded by the active boundary and *top*.

Initially, we let *top* be the *active boundary*, and we repeatedly find a new triangle that is small, has an edge adjacent to an edge of the active boundary and lies below it. Let $p_1, \ldots, p_j$ be the vertices of the present active boundary, where p_1 lies on the *left side* and p_j lies on the *right side* (see Figure 3). We omit all vertices where the active boundary makes an angle π. Let L be the length of the active boundary. A triangle of $\mathcal{T}$ is small with respect to the active boundary if and only if it has an edge no longer than L. Consider the triangle $t_1 \in \mathcal{T}$ incident to p_1 and the edge $\overline{p_1p_2}$, or a part of it, and lying below the active boundary. If this triangle is small w.r.t. the active boundary, we add it to the sequence of small triangles and change the active boundary accordingly. Otherwise, t_1 extends past the edge $\overline{p_1p_2}$ and therefore the vertex p_2 is reflex on the active boundary. Consider the triangle $t_2 \in \mathcal{T}$ incident to p_2 and the edge $\overline{p_2p_3}$. Again, if t_2 is small we add it, otherwise we continue at the reflex vertex p_3. Since p_j necessarily is a convex vertex, the triangle $t_{j-1} \in \mathcal{T}$ incident to p_{j-1} and the edge $\overline{p_{j-1}p_j}$ is small if the

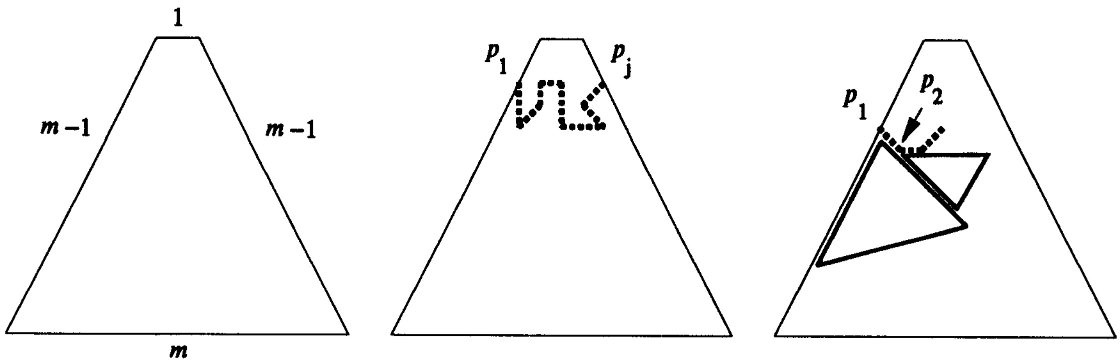

Figure 3: Left: the quadrilateral Q. Middle: an active boundary. Right: the argument of the proof illustrated.

triangle t_{j-2} is not. Therefore, we can always add a small triangle to the sequence (see Figure 3).

As long as the active boundary has length $L < \frac{1}{2}\sqrt{3}m$, the quadrilateral cannot contain a point on *base*, and have exhausted the triangles of $\mathcal{T}$. Initially, $L = 1$ (length of *top*), and any triangle added cannot add more than $L + 2L/\delta$ to L by wideness and by choice. (If a δ-fat triangle has an edge of length 1, then it cannot have an edge longer than $1/\delta$.) Hence, if $L(r)$ denotes the maximum value of L when r triangles have been added, we have $L(0) = 1$ and

$$L(r) \leq L(r-1) \cdot (2 + 2/\delta).$$

The maximum value of r such that $L(r) \leq \frac{1}{2}\sqrt{3}m$ is $\Omega(\log m)$, since δ is a constant. □

3 Partitioning simple polygons preserving wideness

In this section we show how to construct a weak Steiner quadrilateralization of any wide polygon, while preserving the wideness in any resulting subpolygon. Since the subdivision is weak, we may omit all vertices with angle π from the polygon. We give some elementary properties of δ-corridors and segments that are interior to a polygon P, which are straightforward to verify.

Lemma 1 *Two segments of which the supporting lines make an angle α cannot form a δ-corridor with $\delta < 2\sin(\alpha/2)$.*

Lemma 2 *Let C be a circle whose interior lies completely inside a polygon P. Any segment $\overline{pq}$ between two points p and q on C such that $\overline{pq}$ makes an angle α with the tangents to C at p and q cannot create a δ-corridor with $\delta < \frac{1-\cos\alpha}{2\sin\alpha}$.*

Lemma 3 *Let P be a polygon, let $e = \overline{uv}$ be an edge of P and let w be a vertex of P. If the interior of Δuvw does not intersect P, then the addition of the segment that bisects $\alpha = \angle uwv$ to partition P cannot create δ-corridors with $\delta < 2\sin(\alpha/4)$ in the two subpolygons.*

3.1 Partitioning a simple polygon into 6-gons

The *Edge Voronoi Diagram* inside P of the edges of P, denoted EVD, is a subdivision of the interior of P into regions where one edge is closest. See Figure 4 for an example. The EVD is also called the *medial axis* or *internal skeleton.*

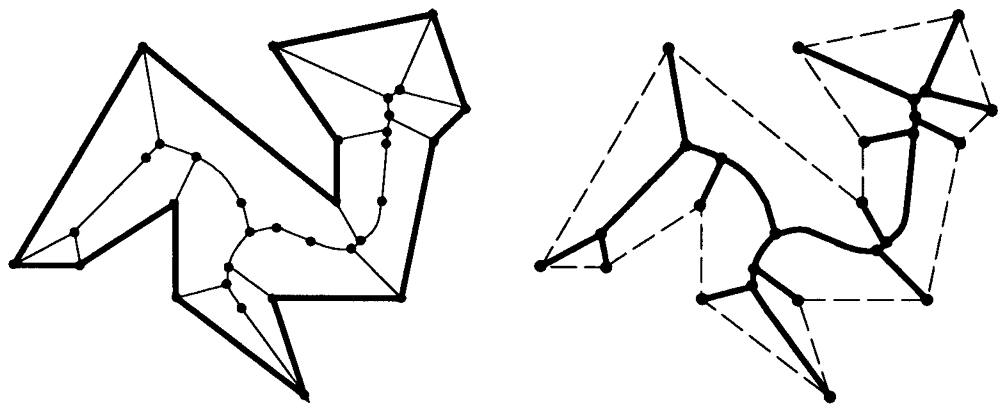

Figure 4: Left: the EVD of a simple polygon. Right: the tree T of the EVD with only nodes of degree 1 or 3.

The arcs and nodes of the EVD form a tree with n leaves, one for each vertex of P. Any arc is a line segment or paraboloid arc. The arc of the EVD incident to a leaf v (a vertex of P) is the bisector of the angle of the two edges of P that are incident to v. In non-degenerate cases, no node of the EVD has degree greater than 3, and there are exactly $n-2$ nodes with degree 3. Each node is the center of a circle that is at equal distance to 3 edges of P, 2 edges and 1 vertex of P, 1 edge and 1 vertex of P, or 2 vertices of P. No edge or vertex of P is nearer to that vertex of the EVD. It follows that the interior of the circle is empty.

We are only interested in the n leaves and the $n-2$ vertices of degree 3, which form a tree T when we identify the two arcs incident to each node of degree 2 (see Figure 4). T has the property that there exists a vertex v of degree 3 whose removal partitions the tree into 3 subtrees, each with at most $\lfloor n/2 \rfloor$ leaves of T. We use the largest circle with empty interior centered at v to partition polygon P into subpolygons, without creating δ-corridors with small δ.

Lemma 4 *Let P be a polygon and C a circle such that* $\text{interior}(C) \subset \text{interior}(P)$ *and which intersects P in 2 points. The two points can be connected with 1 or 2 segments, such that in either of the two subpolygons no δ-corridors are created with $\delta < \frac{1}{2}\sqrt{2} - \frac{1}{2}$. If C intersects P in 3 points, then these points can be connected with at most 4 segments, such that each of the 3 subpolygons that arise use at most 2 of these segments, and no δ-corridor is created with $\delta < 1 - \frac{1}{2}\sqrt{3}$. Furthermore, a $\frac{1}{2}$-wide quadrilateral may be created.*

Proof: (See Figure 5.) Let c be the center of C. Assume first that C intersects P in two points p_1 and p_2. If $\angle p_1 c p_2 \geq \pi/2$ then choose the segment $\overline{p_1 p_2}$. Otherwise, let q be the point inside C such that $p_1 p_2 q$ forms an isosceles right triangle with the right angle at q. Choose the segments $\overline{p_1 q}$ and $\overline{p_2 q}$. By Lemma 2, the segments do not create δ-corridors with $\delta < \frac{1}{2}\sqrt{2} - \frac{1}{2}$.

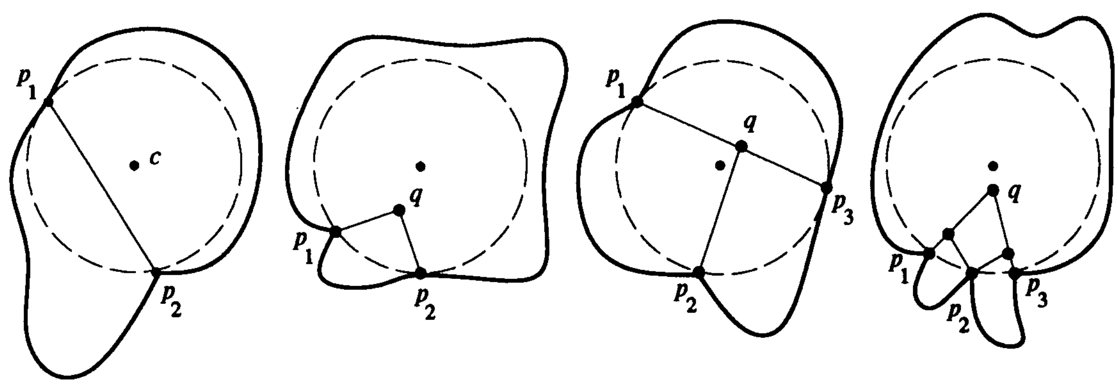

Figure 5: Splitting a polygon without creating narrow corridors using empty circles.

Next, assume that C intersects P in three points p_1, p_2, p_3. Assume w.l.o.g. that p_1 and p_3 are furthest apart. If $\angle p_1 c p_3 \geq \pi/3$ then choose the segment $\overline{p_1 p_3}$. Connect p_2 to this segment by a segment $\overline{p_2 q}$ which makes an angle $\pi/6$ or $\pi/3$ or $\pi/2$ (five possibilities in total) with the tangent to C at p_2. Choose the possibility for which $q \in \overline{p_1 p_3}$ and the angle $\overline{p_2 q}$ and $\overline{p_1 p_3}$ make is at least $\pi/6$ (this is always possible because $\angle p_1 p_2 p_3 \geq \pi/3$). Otherwise, if $\angle p_1 c p_3 < \pi/3$, then let q be the point inside C such that $p_1 p_3 q$ is an equilateral triangle. Choose the segments $\overline{p_1 q}$ and $\overline{p_3 q}$. Connect p_2 to these two segments with two new segments that make an angle of $\pi/6$ with the tangent to C at p_2 (both possibilities). The four new segments form a $\frac{1}{2}$-wide quadrilateral, and all three subpolygons of P use at most two of the new segments (recall that we omit angles of π radians). By Lemmas 1 and 2, the new segments do not create δ-corridors with $\delta < 1 - \frac{1}{2}\sqrt{3}$. □

The following algorithm is used to partition any polygon P with n vertices into 6-gons:

1. Compute the EVD of P using the algorithm of Lee [16] or Yap [26]. Consider it as a graph, and remove each node of degree 2 by identifying the arcs incident to it to obtain the tree T. Select a vertex v in T whose removal gives 3 subtrees which have at most $\lfloor n/2 \rfloor$ leaves each.

2. Compute the largest empty interior circle centered at v. It intersects P in 2 or 3 points. Add the segments as used in the proof of Lemma 4, to obtain 2 or 3 subpolygons of P (and, possibly, a wide quadrilateral).

3. For each subpolygon that has at least 7 vertices, recursively subdivide it.

Theorem 3 *Any δ-wide simple polygon P with n vertices can be partitioned in $O(n \log^2 n)$ time into $O(n)$ γ-wide polygons with at most 6 vertices and $\gamma = \min\{\delta, 1 - \frac{1}{2}\sqrt{3}\}$.*

Proof: Notice that splitting P as in Lemma 4 yields subpolygons with at most 3 new vertices. Therefore, the above algorithm gives subpolygons with at most $\lfloor n/2 \rfloor + 3$ vertices. As long as $n > 6$, a polygon is partitioned into subpolygons with fewer vertices. The wideness guarantee of the resulting polygons follows from Lemma 4. The time bound of the algorithm follows easily from the $O(n \log n)$ time algorithms for computing an EVD [16, 26]. □

3.2 Partitioning 6-gons, 5-gons and 4-gons into convex pieces

Let P be a 6-gon with vertices $v_1, \ldots, v_6$. If P is non-convex, we show how to partition P in $O(1)$ convex pieces. Assume w.l.o.g. that v_1 is a reflex vertex (see Figure 6, left). Since $\angle v_6 v_1 v_5 + \angle v_5 v_1 v_4 + \angle v_4 v_1 v_3 + \angle v_3 v_1 v_2 = \angle v_6 v_1 v_2 > \pi$, at least one of the four angles is $> \pi/4$. (If any of v_3, v_4, v_5 is not visible from v_1, it can simply be removed from consideration; the argument still holds with better constants.) We draw a segment from v_1 to bisect this angle. The two new angles in the two subpolygons at v_1 are at least $\pi/8$, and therefore the largest new angle at v_1 is at least $\pi/8$ less than $\angle v_6 v_1 v_2$. We continue this procedure until every subpolygon of P is convex. The same idea can be used for non-convex 5-gons and 4-gons. We obtain:

Lemma 5 *Any 6-gon can be partitioned into $O(1)$ convex polygons with at most 6 vertices each, without creating δ-corridors with $\delta < 2\sin(\pi/16)$. Any 5-gon can be partitioned into $O(1)$ convex polygons with at most 5 vertices each, without creating δ-corridors with $\delta < 2\sin(\pi/12)$. Any 4-gon can be partitioned into $O(1)$ convex polygons with at most 4 vertices each, without creating δ-corridors with $\delta < 2\sin(\pi/8)$.*

3.3 From convex 6-gons to 5-gons

Let P be a convex 6-gon with vertices $v_1, \ldots, v_6$. Consider the triangle $\Delta v_1 v_3 v_5$, and assume w.l.o.g. that v_1 has angle $\angle v_5 v_1 v_3 \geq \pi/3$ (see Figure 6, middle). Therefore,

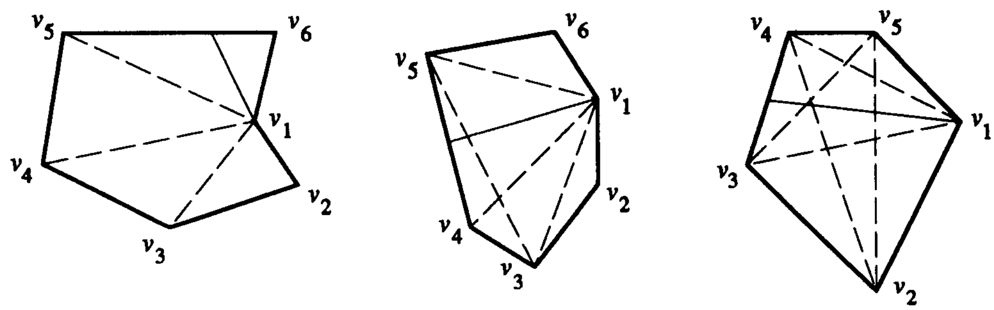

Figure 6: Left: one step of the partitioning of a 6-gon into convex pieces. Middle: partitioning a convex 6-gon. Right: partitioning a convex 5-gon.

$\angle v_3v_1v_4 \geq \pi/6$ or $\angle v_4v_1v_5 \geq \pi/6$. If $\angle v_3v_1v_4 \geq \pi/6$, then we bisect this angle with a segment to obtain a 5-gon and a 4-gon without creating δ-corridors with $\delta < 2\sin(\pi/24)$ (by Lemma 3). The other case is the same.

3.4 From convex 5-gons to 4-gons

Let P be a convex 5-gon with vertices $v_1, \ldots, v_5$. Consider the sequence $v_1v_3v_5v_2v_4$, which forms the star inscribed in P (see Figure 6, right). The sum of the five angles at the five points of the star is π, and hence, at least one the five angles is at least $\pi/5$. Assume w.l.o.g. that it is at v_1. We bisect the angle $\angle v_3v_1v_4$ with a segment to obtain two 4-gons without creating δ-corridors with $\delta < 2\sin(\pi/20)$ (by Lemma 3).

The above partitionings lead to the following result:

Theorem 4 *A δ-wide simple polygon P with n vertices can be partitioned in $O(n\log^2 n)$ time into $O(n)$ γ-wide quadrilaterals and triangles, where $\gamma = \min\{\delta, 1 - \frac{1}{2}\sqrt{3}\}$.*

Remark: The partitioning methods of this section do not introduce angles smaller than $\pi/12$ radians.

4 Covering quadrilaterals by triangles

We show that a δ-wide quadrilateral Q can be covered by $O(1)$ γ-wide triangles, where $\gamma = \min\{\delta/4, \arctan(\pi/10)\}$. No three triangles in the covering overlap in a positive area region.

The first stage of the covering is to separate angles that are smaller than $\pi/5$. Let Q be a convex quadrilateral, and assume that v is a vertex with angle $\alpha < \pi/5$. Let e and

e' be the edges incident to v with e the shorter one. Let w be the other endpoint of e, and let w' be the other endpoint of e'. Choose point u on e' such that $\angle w'uw$ is $\pi/5 + \alpha$ radians. Partition Q into the triangle Δuvw and a quadrilateral Q' which is Q with v replaced by u. The triangle Δuvw has angles $4\pi/5 - \alpha$, $\pi/5$ and α. Since $|e| \leq |e'|$, the angle $\angle vww' \geq \frac{\pi}{2} - \frac{\alpha}{2}$, and therefore, $\angle uww' = \angle vww' - \angle vwu = \frac{\pi}{2} - \frac{\alpha}{2} - \frac{\pi}{5} > \frac{\pi}{5}$. The quadrilateral Q' has a vertex with angle α replaced by an angle $\pi/5 + \alpha$ without creating an angle less than $\pi/5$.

Next we separate all angles greater than $4\pi/5$, without creating angles less than $\pi/5$ or greater than $4\pi/5$. Let v be a vertex such that the incident edges make an angle greater than $4\pi/5$. Let w_1, w_2, w_3 be the other vertices of Q in clockwise order. At least one of $\angle w_1vw_2$ and $\angle w_2vw_3$ must be greater than $2\pi/5$; partition this angle by a segment that splits Q into a triangle and a quadrilateral. It is easy to see that no angles smaller than $\pi/5$ or greater than $4\pi/5$ are created.

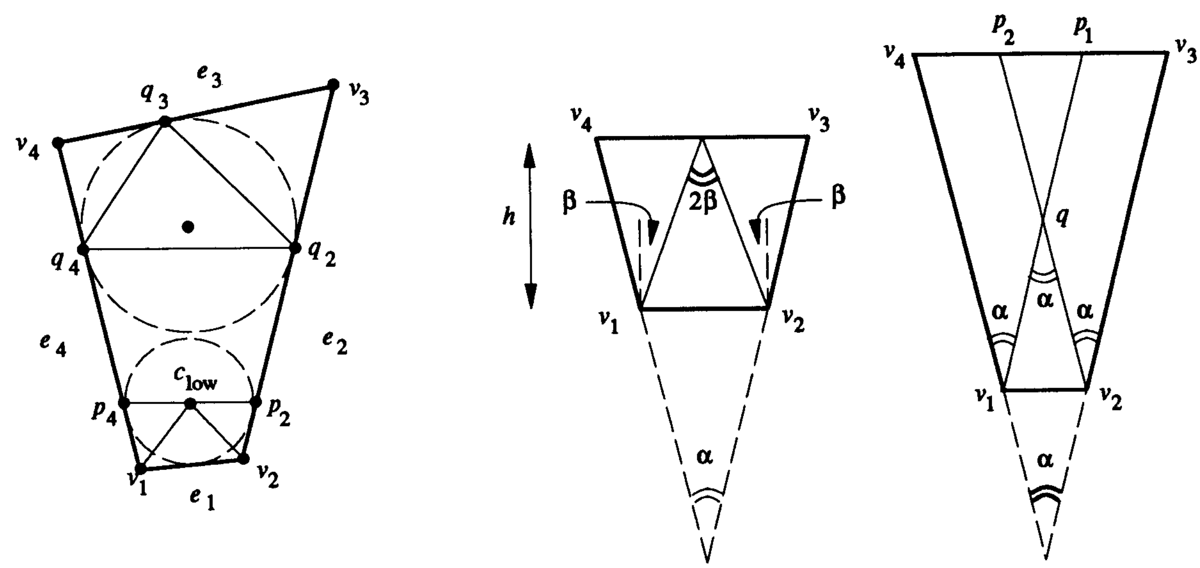

Figure 7: Left: the quadrilateral Q. Right: illustrations of the proof of Lemma 10.

Let Q be a quadrilateral with all angles at least $\pi/5$ and at most $4\pi/5$. We again consider the Edge Voronoi Diagram of the edges of Q. The EVD contains two nodes of degree 3, which are the centers of circles that touch Q in three edges. Orient Q such that the center of the larger circle is vertically above the center of the smaller circle (see Figure 7). Label the edge touching only the smaller circle with e_1, and label the other edges with e_2, e_3 and e_4 counterclockwise. Label the vertex incident to e_4 and e_1 with v_1, and the other vertices are v_2, v_3 and v_4. Let the circle tangent to e_2, e_3 and e_4 be C_{high} and the circle tangent to e_4, e_1 and e_2 is C_{low}. The center of C_{high} is c_{high}, the center of

C_{low} is c_{low}. Furthermore, $q_2 = C_{high} \cap e_2$, $q_3 = C_{high} \cap e_3$ and $q_4 = C_{high} \cap e_4$. Finally, ℓ_p is the horizontal line through c_{low}, $p_2 = \ell_p \cap e_2$ and $p_4 = \ell_p \cap e_4$.

Lemma 6 *ℓ_p indeed intersects e_2 and e_4.*

Proof: C_{low} is a circle that has all three points incident to Q on the lower semi-circle. Since ℓ_p lies above $C_{low} \cap e_2$ and $C_{low} \cap e_4$ and thus above v_1 and v_2, ℓ_p cannot intersect e_1. If e_3 and e_2 intersect ℓ, then, since e_3 does not intersect C_{low}, $\angle v_1 + \angle v_2 + \angle v_4 \geq 2\pi$, a contradiction. Similar if e_3 and e_4 intersect ℓ_p. Hence, e_2 and e_4 intersect ℓ_p. □

Lemma 7 *The following triangles have all angles at least $\pi/10$: $\triangle q_4 q_3 v_4$, $\triangle q_2 q_3 q_4$, $\triangle q_2 v_3 q_3$, $\triangle v_1 c_{low} p_4$, $\triangle v_2 c_{low} v_1$, $\triangle v_2 p_2 c_{low}$.*

Proof: Since $\pi/5 \leq \angle v_3 \leq 4\pi/5$, it follows that $\angle v_3 q_2 q_3 = \angle v_3 q_3 q_2 \geq \pi/10$ and they are $\leq 2\pi/5$. Symmetrically, $\angle v_4 q_3 q_4$ and $\angle v_4 q_4 q_3$ are bounded by $\pi/10$ and $2\pi/5$. It follows that $\angle q_2 q_3 q_4 \geq \pi/5$. Furthermore, since $\angle v_1$ and $\angle v_2$ are at most $4\pi/5$, we get $\angle q_2 q_4 v_1 = \angle q_4 q_2 v_2 \geq \pi/5$. Also, these angles are at most $\pi/2$. It follows that $\angle q_2 q_4 q_3$ and $\angle q_4 q_2 q_3$ are at least $\pi/10$. This proves the angles of the first three triangles. For the other three triangles, observe that $\overline{c_{low} v_1}$ bisects $\angle v_1$ and $\overline{c_{low} v_2}$ bisects $\angle v_2$. The bounds on the angles now follow in a similar fashion. □

The line ℓ_q through q_2 and q_4 and the line ℓ_p through p_2 and p_4 are horizontal. If ℓ_q lies below ℓ_p, then the triangles of the above lemma cover Q with maximum overlap 2, and we are done.

So assume ℓ_q lies above ℓ_p. Then $p_4 p_2 q_2 q_4$ is a trapezoid R. It remains to partition or cover R using fat triangles. R can be not wide in one of of two ways. If the distance between ℓ_q and ℓ_p is large compared to $|\overline{q_4 q_2}|$, then R is not wide and this is due to the fact that also the quadrilateral Q is not wide. On the other hand, if the distance between ℓ_q and ℓ_p is small, then also R is not wide, but in this case Q is wide. The latter case gives potential problems, since Q is wide and should be covered by $O(1)$ fat triangles, however, T is not wide and may require many fat triangles. To overcome this problem, we take a new horizontal line ℓ'_p below ℓ_p, and we *partition* the trapezoid T' between ℓ_q and ℓ'_p (we must partition T in order to obtain a 2-cover of Q, since T partially overlaps with three of the triangles chosen in Lemma 7). If the distance between ℓ_q and ℓ_p is less than $\frac{\pi}{20}|\overline{p_2 p_4}|$, then we define a new horizontal line ℓ'_p below ℓ_p such that the distance is $\frac{\pi}{20}|\overline{p'_2 p'_4}|$, where $p'_2 = \ell'_p \cap e_2$ and $p'_4 = \ell'_p \cap e_4$.

Lemma 8 *ℓ'_p is well-defined, i.e., there exists a horizontal line ℓ'_p below ℓ_q which intersects e_2 and e_4 such that the distance between ℓ_q and ℓ'_p is $\frac{\pi}{20}|\overline{p'_2 p'_4}|$.*

Proof: Assume w.l.o.g. that v_1 lies at least as high as v_2. Let ℓ_v be the horizontal line through v_1. Let $v'_2 = \ell_v \cap e_2$. Since $\angle p_4 c_{low} v_1 \geq \pi/10$, it follows that the distance between ℓ_v and ℓ_p is at least $\frac{1}{2}|\overline{v_1v_2'}| \cdot \tan\frac{\pi}{10} \geq \frac{\pi}{20}|\overline{v_1v_2'}|$. Therefore, it is possible to choose ℓ_p' between ℓ_p and ℓ_v such that the distance to ℓ_q is exactly $\frac{\pi}{20}|\overline{p_2'p_4'}|$. □

Let T be a trapezoid with edges e_1, e_2, e_3 and e_4 listed counterclockwise. The vertices are denoted v_1, v_2, v_3 and v_4, counterclockwise, where v_1 is incident to e_4 and e_1. Assume without loss of generality that e_1 and e_3 are horizontal, e_1 lies below e_3, and e_3 is at least as long as e_1. Denote by h the thickness of T, which is the distance from e_1 to e_3. Denote by α the angle of the lines supporting e_2 and e_4, and $\alpha = 0$ if they are parallel. Then the angle of T at v_1 and v_2 is $\pi/2 + \alpha/2$, and the angle at v_3 and v_4 is $\pi/2 - \alpha/2$. We have $0 \leq \alpha \leq 4\pi/5$ by the way T is obtained from Q, and also $h \geq \frac{\pi}{20}|e_1|$.

To solve the 2-covering problem for quadrilateral Q, it remains to *partition* T if $h = \frac{\pi}{20}|e_1|$, and to *cover* (or partition) T if $h \geq \frac{\pi}{20}|e_1|$.

Lemma 9 *The trapezoid T with $h = \frac{\pi}{20}|e_1|$ can be partitioned into 3 triangles with all angles at least* $\arctan(\pi/10)$.

Proof: Let p be the point in the middle of e_3. Then the triangles Δv_1pv_4, Δv_2pv_1 and Δv_2v_3p partition T and every angle is at least $\arctan(\pi/10)$ radians. □

Lemma 10 *The δ-wide trapezoid T with $h \geq \frac{\pi}{20}|e_1|$ can be 2-covered by 3 triangles with all angles at least* $\min\{\arctan(\pi/10), \delta/4\}$ *radians.*

Proof: (See Figure 7.) If $|e_3| \leq 2|e_1|$, then partition T as in the previous lemma. If we set $|e_2| = |e_4| = 1$, then $|e_3| = \delta$. Furthermore, $h = (\delta - |e_1|)/(2\tan(\alpha/2))$ and $\sin(\alpha/2) = (\delta - |e_1|)/2$, so

$$\tan(\beta) = \frac{|e_1|}{2h} = \frac{|e_1| \cdot \tan(\alpha/2)}{\delta - |e_1|} = \frac{|e_1| \tan(\arcsin((\delta - |e_1|)/2))}{\delta - |e_1|} \geq |e_1|/2 \geq \delta/4 .$$

Hence, all angles are at least $\min\{\arctan(\pi/10), \delta/4\}$ radians.

Otherwise, let p_1 and p_2 be the points on e_3 such that $\overline{v_1p_1}$ and e_2 are parallel, resp. $\overline{v_2p_2}$ and e_4 are parallel. Then $\overline{p_1v_1}$ and $\overline{p_2v_2}$ intersect in a point q, and the triangles $\Delta v_1p_1v_4$, $\Delta v_2v_3p_2$ and Δv_1v_2q cover T and only the first two overlap. Since $|e_3| > 2|e_1|$, we have $\alpha > \delta/2$, otherwise, T is not δ-wide. It follows easily that all angles of the three triangles are at least $\min\{\pi/5, \delta/2\}$ radians. □

The above lemmas provide all ingredients for the 2-covering of a quadrilateral by triangles with a guarantee on minimum angles.

Theorem 5 *A δ-wide quadrilateral Q can be 2-covered by $O(1)$ triangles with all angles at least* $\min\{\arctan(\pi/10), \delta/4\}$ *radians.*

Theorems 4 and 5 yield:

Theorem 6 *A δ-wide polygon P with n vertices can be 2-covered by $O(n)$ γ-wide triangles, where* $\gamma = 2\sin(\min\{\arctan(\pi/10), \delta/4\}/2)$.

Remark: The algorithms for partitioning into quadrilaterals and covering by triangles can also be used for the regions of a planar straight line graph or a point set. Firstly, the Euclidean minimum spanning tree is constructed on the connected components of the PSLG (see [17, 26]) or point set (see [9, 22]), and its edges are part of the partitioning. It can be shown that no narrow corridors are added. Secondly, every face of the subdivision is partitioned as before into quadrilaterals.

5 The contour of the union of simple polygons

The previous sections showed that a δ-wide simple polygon can be 2-covered by $O(n)$ γ-wide triangles, where $1/\gamma = \Theta(\delta)$. This result allows us to show a bound on the union size of a set of δ-wide simple polygons. We state the results below:

Theorem 7 *Let S be a set of δ-wide polygons with n vertices in total. The maximum complexity of the contour of the union for S is* $O((n \log\log n)/\delta^3)$.

Proof: For any polygon P_i in S, let C_{P_i} be a linear size 2-covering by γ-wide triangles, where $1/\gamma = \Theta(1/\delta)$. Then the union $\bigcup\{P_i \mid P_i \in S\}$ is the union of the 2-coverings by triangles $\bigcup\{t \mid t \in C_{P_i}, P_i \in S\}$. By [19], the complexity of the union of the triangles is $O((n \log\log n)/\gamma^3)$, which proves the theorem. □

Theorem 8 *Let S be a set of δ-wide polygons with n vertices in total. The maximum complexity of the boundaries of all cells covered by at most k polygons of S is* $O((nk \log\log(n/k))/\delta^3)$.

In the remainder of this section, we briefly discuss six applications in which a bound on the maximum complexity of the contour of the union leads to more efficient and simpler algorithms. Let δ be any positive constant.

Hidden surface removal. Katz, Overmars and Sharir [12] presented an efficient hidden surface removal algorithm for objects with small union size. They prove that for a set S of n non-intersecting objects in 3-space and a viewing point v, the visibility map of S as seen from v can be computed in $O((U(n) + k)\log^2 n)$ time, where $U(n)$ is the maximum complexity of the union in the projection, and k is the size of the resulting visibility map. It is assumed that the objects in S are ordered by depth from the viewing point. In [12] three cases are identified in which the union size is guaranteed to be small: (i) S is a set of balls, (ii) S is a set of triangles which are wide in the projection, and (iii) S consists of the set of triangles that form a polyhedral terrain. Applying Theorem 7, we obtain:

Theorem 9 *Let S be a set of polygons in 3-space with n vertices in total, let v be a viewing point and let a depth order for S exist and be given. If the projections of the polygons in S are δ-wide, then the visibility map of S can be computed in $O((n \log\log n + k)\log^2 n)$ time, where k is the complexity of the visibility map.*

Motion planning. The general motion planning problem is to find a sequence of motions that will take a robot from one position to another, without colliding with any of a set of obstacles. Often both the robot and the obstacles are modeled by simple or convex polygons. An important concept in motion planning is the configuration space, and the subspace of all free placements inside it (i.e., all placements of the robot for which it does not intersect any obstacle). Assume that the robot R is modeled by a simple k-gon, the obstacles by a set S of simple polygons $P_1, \ldots, P_m$, with n vertices in total, and that the robot is only allowed to translate. It is well-known that the free placement space has complexity $O((nk)^2)$ in this case, see e.g. [15, 23]. They also give examples to show that this bound is the best possible. If the obstacle polygons and the robot polygon are δ-wide, then it is easy to prove that the free placement space has complexity $O(nk \log\log nk)$. See also [25] for related results.

Injection molding. In [6], Bose, van Kreveld and Toussaint study the problem of injection molding under the optimization criterion of minimizing the number of venting holes needed to ensure a complete fill of the mold. In geometric terms, the problem is to find an orientation of a polyhedron in 3-space which minimizes the number of local maxima (in the vertically upward direction). They show that the problem for a polyhedron with n vertices can be transformed to a covering problem with at most n convex polygons in the plane, and with $O(n)$ vertices in total, which can be solved in $O(n^2)$ time. For polyhedra that satisfy a regularity condition, the convex polygons in the plane are wide, and the authors give a more efficient algorithm than for general polyhedra. The second algorithm runs in time $O(U(n,k)\log^2 n)$, where $U(n,k)$ is the total complexity of all regions of the plane covered by k or fewer polygons. By the wideness and by Theorem 8, this is $O(nk \log^2 n \log\log(n/k))$ time, where k is the number of local maxima in the optimal orientation. In practice, k will be much smaller than n, often a small constant.

Point stabbing. Let S be a set of simple polygons with n vertices in total. The point containment query problem for S is to preprocess S for the following type of queries: Given a point q, report all polygons of S that contain it. If S is a set of triangles, then the problem is related to the simplex range query problem, and complicated solutions that require $O(n \log^{O(1)} n)$ storage and $O(\sqrt{n} \log^{O(1)} n + k)$ query time are known (k is the output size). Other solutions require $O(n^2 \log^{O(1)} n)$ storage and $O(\log^{O(1)} n + k)$ query time. See e.g. Chazelle [7] and Matoušek [18]. A simple and more efficient solution for objects with small union size has been obtained by Sharir [24]. Using this result, we obtain:

Theorem 10 *Let S be a set of δ-wide polygons with n vertices in total. S can be preprocessed for point location queries in $O((n \log^2 n \log\log n)/\delta)$ expected time into a data* *structure of size* $O(n \log n \log \log n)$*, such that all* k *polygons of* S *that contain a given query point can be reported in* $O((k+1)\log n)$ *time.*

Ray shooting. Let S be a set of δ-wide polygons that lie in horizontal planes in 3-dimensional space, and let the polygons have n vertices in total. The vertical ray shooting problem for S is the following: Given a query point q, which polygon of S is hit first if q is translated vertically downward? An approach similar to the ones of Cole and Sharir [8] and de Berg and Overmars [5] leads to:

Theorem 11 *Let* S *be a set of* δ*-wide polygons with* n *vertices in total, which lie in horizontal planes in 3-dimensional space.* S *can be preprocessed in* $O(n \log^2 n \log \log n)$ *time into a data structure of size* $O(n \log n \log \log n)$*, such that vertical ray shooting queries can be answered in* $O(\log^2 n)$ *time.*

Red-blue intersection detection. Given a set R of red polygons and a set B of blue polygons with n vertices in total, we wish to decide whether there are a red and a blue polygon which intersect. The interior is also considered part of a polygon. If R is a set of lines and B a set of points, then the problem is called Hopcroft's problem, and an $O(n^{4/3}2^{O(\log^* n)})$ time solution is given by Matoušek [18], see also Chazelle [7] and Agarwal and Sharir [1].

If the sets R and B contain δ-wide polygons, then the following relatively simple solution may be more efficient. Compute the union $U(R)$ of all red polygons. Using divide-and-conquer and plane sweep for the merge, this requires $O(n \log^2 n \log \log n)$ time. Preprocess $U(R)$ for efficient point location in $O(n \log n \log \log n)$ time (see e.g. [9, 22]). For every vertex v of every blue polygon, test if $v \in U(R)$ by a point location query. If the answer is yes for any vertex, then a red-blue intersection is detected. Otherwise, we compute and preprocess $U(B)$, the union of the blue polygons, in the same way, and we query with every vertex w of every red polygon to test if $w \in U(B)$. Finally, if all queries are answered to the negative, then let S_R and S_B be the sets of $O(n \log \log n)$ segments in the contour of $U(R)$ and $U(B)$, respectively. Notice that the segments in S_R are disjoint, except at the endpoints, and the same holds for S_B. We test if any segment of S_R intersects any segments of S_B by a standard plane sweep in $O(n \log n \log \log n)$ time. If there is no intersection involving a segment of S_R and a segment of S_B, then we conclude that there is no red-blue intersection among R and B.

Theorem 12 *Let* R *be a set of* δ*-wide red polygons, let* B *be a set of* δ*-wide blue polygons, and let* n *be the total number of vertices of polygons in* R *and* B*. In* $O(n \log^2 n \log \log n)$ *time, one can determine if any red polygon intersects any blue polygon.*

6 Conclusions and open problems

In this paper we studied the complexity of the contour of the union of a set of simple polygons. The notion of δ-wide polygons was introduced, where the value of δ influences the contour size. An upper bound on the maximum union contour size was given which generalizes the results of Matoušek et al.[19] on the union size of fat (or wide) triangles. We also showed that a *partitioning* of a polygon into wide triangles cannot give the desired bounds, because too many wide triangles will needed.

The partitioning and covering algorithms presented in this paper require $O(n \log^2 n)$ and $O(n \log^2 n)$ time, respectively. It may be possible to improve upon this bound. We remark, however, that for most applications it is not necessary to perform the actual partitioning or covering, but instead regard the techniques as a proof that the union size is not large. To compute the actual union of a set of δ-wide polygons, a straightforward $O(n \log^2 n \log \log n)$ algorithm exists, see e.g. Kedem et al.[13]. A slightly more efficient, but randomized, algorithm is given by Miller and Sharir [20], see also [19]. A further speedup may be possible. A third open problem we take from [19]: The maximum union contour size of n fat triangles is $O(n \log \log n)$ and $\Omega(n\alpha(n))$. There is a gap to be closed. Also, the dependency of the union size on δ is not close to optimal. (Partial progress on these issues for wedges is made recently by Efrat, Rote and Sharir in [10].) More tight bounds would immediately give more tight bounds for the union contour size of δ-wide polygons.

References

[1] Agarwal, P.K., and M. Sharir, Red-blue intersection detection algorithms, with applications to motion planning and collision detection. *SIAM J. Comp.* **19** (1990), pp. 297–321.

[2] Agarwal, P.K., and M. Sharir, Applications of a new partitioning scheme. *Discr. & Comp. Geom.* **9** (1993), pp. 11–38.

[3] Alt, H., R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Shirra, and C. Uhrig, Approximate motion planning and the complexity of the boundary of the union of simple geometric figures. *Proc. 6th ACM Symp. Comp. Geom.* (1990), pp. 281–289.

[4] Bentley, J., and T. Ottmann, Algorithms for reporting and counting intersections. *IEEE Trans. Comp.* **28** (1979), pp. 643–647.

[5] de Berg, M., and M. Overmars, Hidden surface removal for axis-parallel polyhedra. *Proc. 31st IEEE Symp. Found. Comp. Science* (1990), pp. 252–261.

[6] Bose, J., M. van Kreveld, and G. Toussaint, Filling polyhedral molds. *Proc. 3rd Workshop on Algorithms and Data Structures* (1993), Lect. Notes in Comp. Science 709, Springer-Verlag, pp. 210–221.

[7] Chazelle, B., Cutting hyperplanes for divide-and-conquer. *Discr. & Comp. Geom.* **9** (1993), pp. 145–158.

[8] Cole, R., and M. Sharir, Visibility problems for polyhedral terrains. *J. Symb. Comp.* **7** (1989), pp. 11–30.

[9] Edelsbrunner, H., *Algorithms in Combinatorial Geometry.* Springer-Verlag, Berlin, 1987.

[10] Efrat, A., G. Rote, and M. Sharir, On the union of fat wedges and separating a collection of segments by a line. *Proc. 5th Canad. Conf. on Comp. Geometry*, pp. 115–120.

[11] Edelsbrunner, H., 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. *Discr. & Comp. Geom.* **4** (1989), pp. 523–539.

[12] Katz, M.J., M.H. Overmars, and M. Sharir, Efficient hidden surface removal for objects with small union size. *Computational Geometry–Theory and Applications* **2** (1992), pp. 223–234.

[13] Kedem, K., R. Livne, J. Pach, and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discr. & Comp. Geom.* **1** (1986), pp. 59–71.

[14] van Kreveld, M., On fat covering, fat partitioning and the union size of polygons (extended abstract). *Proc. 3rd Workshop on Algorithms and Data Structures* (1993), Lect. Notes in Comp. Science 709, Springer-Verlag, pp. 452–463. Also: Tech. Rep. SOCS-93-2, School of Computer Science, McGill University, 1993.

[15] Latombe, J.-C., *Robot motion planning.* Kluwer Academic Publishers, Boston, 1991.

[16] Lee, D.T., Medial axis transformation of a planar shape. *IEEE Trans. Pattern Anal. Mach. Intel.* **4** (1982), pp. 363–369.

[17] Lee, D.T., and R.L. Drysdale III, Generalization of Voronoi diagrams in the plane. *SIAM J. Comp.* **10** (1981), pp. 73–87.

[18] Matoušek, J., Range searching with efficient hierarchical cuttings. *Proc. 8th ACM Symp. Comp. Geom.* (1992), pp. 276–285.

[19] Matoušek, J., N. Miller, J. Pach, M. Sharir, S. Sifony, and E. Welzl, Fat triangles determine linearly many holes. *Proc. 32nd IEEE Symp. Found. Comp. Science* (1991), pp. 49–58. To appear in: *SIAM J. Comput.*

[20] Miller, N., and M. Sharir, *Efficient randomized algorithms for constructing the union of fat triangles and of pseudodiscs.* Manuscript, 1993.

[21] Overmars, M., *Point location in fat subdivisions.* Tech. Rep. RUU-CS-91-40, Dept. of Comp. Science, Utrecht University, 1991.

[22] Preparata, F.P., and M.I. Shamos, *Computational geometry – an introduction.* Springer-Verlag, New York, 1985.

[23] Sharir, M., Efficient algorithms for planning purely translational collision-free motion in two and three dimensions. *Proc. IEEE Int. Conf. on Robotics and Automation* (1987), pp. 1326–1331.

[24] Sharir, M., On k-sets in arrangements of curves and surfaces. *Discr. & Comp. Geom.* **6** (1991), pp. 593–613.

[25] van der Stappen, A. F., D. Halperin, and M. Overmars, *The complexity of free space for a robot moving amidst fat obstacles.* Tech. Rep. CS-RUU-92-5, Dept. of Comp. Science, Utrecht University, 1992.

[26] Yap, C.K., An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments. *Discr. & Comp. Geom.* **2** (1987), pp. 365–393.