On the boundary complexity of the union of fat triangles*

János Pach†
City College, CUNY and Rényi Institute, Hungarian Academy
Courant Institute, NYU, 251 Mercer Street, New York, NY 10012
pach@cims.nyu.edu

Gábor Tardos‡
Rényi Institute, Hungarian Academy
POB 127, H-1364 Budapest, Hungary
tardos@renyi.hu

MSC: 52C45, 52B55
Keywords: holes, boundary complexity, fat objects

Abstract

*A triangle is said to be δ-*fat *if its smallest angle is at least $\delta > 0$. A connected component of the complement of the union of a family of triangles is called* hole*. It is shown that any family of n δ-fat triangles in the plane determines at most $O\left(\frac{n}{\delta}\log\frac{2}{\delta}\right)$ holes. This improves on some earlier bounds of Efrat, Rote, Sharir and Matoušek et al. Solving a problem of Agarwal and Bern, we also give a general upper bound for the number of holes determined by n triangles in the plane with given angles. As a corollary, we obtain improved upper bounds for the boundary complexity of the union of fat polygons in the plane, which, in turn, leads to better upper bounds for the running times of some known algorithms for motion planning, for finding a separator line for a set of segments, etc.*

1. Introduction, main results

Many basic problems in computational geometry related to motion planning [SS83, SS89, SS90], range searching [K97, GJ97], computer graphics [AK94], and geographic information systems [BK97] lead to questions about the complexity of the boundary of the union of certain geometric objects. When the boundary is simple, these problems can usually be solved more efficiently [GS93]. This was the motivation behind a lot of research during the past fifteen years, establishing upper bounds for the complexity (or, equivalently, for the description size) of the union of various objects.

Perhaps the first results of this kind were the following. Given n simply connected regions in the plane, any two of which share at most 2 (resp., at most 3) boundary points, the boundary of their union consists of at most $6n-12$ (resp., at most $n\alpha(n)$) *simple arcs*, i.e., connected pieces whose interior belongs to the boundary of a single region [KL86] (resp., [EG89]; here $\alpha(n)$ denotes the extremely slowly growing inverse of Ackermann's function). In

* An extended abstract of this paper appeared in [PT00]
† Supported by NSF grant CCR-97-321018, OTKA T-020914 and AKP 2000-78 2.1
‡ Supported by OTKA grants T030059, T029255, FKFP 0607/1999, and AKP 2000-78 2.1

some sense, this result is best possible: if two regions are allowed to cross at 4 boundary points, then the boundary of their union may consist of $\Omega(n^2)$ simple arcs. Indeed, consider n very "skinny" pairwise crossing triangles, no three of which have a point in common.

However, it was discovered by Matoušek et al. [MP94] that if we restrict how skinny our triangles can be, we can still establish a nearly linear upper bound on the complexity of their union. For any $\delta > 0$, a triangle is said to be *δ-fat* if each of its angles is at least δ. (The reciprocal of the smallest angle of a triangle is often called its *aspect ratio*.) It turned out that for any fixed $\delta > 0$, the boundary of the union of n δ-fat triangles in the plane consists of at most $n \log\log n/\delta^3$ simple arcs.

The *boundary complexity* of the union of a family $\mathcal{T}$ of triangles (or simply connected regions) is defined as the number of simple arcs along $\mathrm{Bd}(\cup\mathcal{T})$, the boundary of the union of $\mathcal{T}$. A connected component of the complement of $\cup\mathcal{T}$ is called a *hole*. The heart of the argument in [MP94] was the following statement.

Theorem 0. (Matoušek et al.) *Any family of n δ-fat triangles in the plane determines $O\left(n/\delta^3\right)$ holes.*

The concept of δ-fatness, as well as the above theorem, has been extended to arbitrary polygons by van Kreveld [K98]. For other extensions and generalizations, see [SH93],[S94],[ES97], [EK98], and [E99].

For *wedges* (i.e., cones) in place of triangles, a somewhat better upper bound was found by Efrat, Rote, and Sharir [ER93]. They proved that the number of holes determined by n wedges in the plane (and the boundary complexity of their union) is $O\left(\frac{n}{\delta^2}\log\frac{2}{\delta}\right)$.

Our first theorem generalizes and strengthens this result.

Theorem 1. *Any family of n δ-fat triangles in the plane determines $O\left(\frac{n}{\delta}\log\frac{2}{\delta}\right)$ holes. This bound is tight up to the logarithmic factor.*

Theorem 1 can be used to establish a more general upper bound for the number of holes determined by a family of triangles with given angles.

Theorem 2. *Let $\mathcal{T} = \{T_1, \ldots, T_n\}$ be a family of $n > 1$ triangles in the plane, and let α_i denote the smallest angle of T_i $(1 \le i \le n)$. Suppose $0 < \alpha_1 \le \alpha_2 \le \cdots \le \alpha_n$, and let $k \le n$ be the largest integer satisfying $\sum_{i=1}^{k}\alpha_i < \pi$.*

Then $\mathcal{T}$ determines $O(nk\log k)$ holes. Furthermore, there exists a family $\mathcal{T}' = \{T'_1 \ldots, T'_n\}$, where T'_i and T_i are congruent for all i, and $\mathcal{T}'$ determines $\Omega(nk)$ holes.

Of course, the same result applies to wedges, provided that their angles are separated from π. Moreover, in this case an almost identical upper bound holds for the *boundary complexity* of the union.

Theorem 3. *Let $\mathcal{T}$ be a family of n wedges in the plane with angles $0 < \alpha_1 \le \cdots \le \alpha_n < \pi$. Let $k \le n$ be the largest integer satisfying $\sum_{i=1}^{k}\alpha_i < \pi$.*

If $k \ge 2$, then the boundary complexity of $\cup\mathcal{T}$ is $O(nk\log k)$. Furthermore, there exists a family of n wedges with angles $\alpha_1, \ldots, \alpha_n$, which determines $\Omega\left((\pi - \alpha_n)nk\right)$ holes.

Notice that Theorem 3 bounds the boundary complexity instead of the number of holes. The bound in Theorem 2 for the number of holes in families of triangles cannot be extended to boundary complexity as there are families of equilateral triangles (for which $k = 2$) with superlinear boundary complexities (at least $\Omega(n\alpha(n))$, cf. [WS88]).

In some applications, e.g., the overlay of triangulated environments in geographic information systems (GIS), we cannot assume that *all* of the participating triangles are fat (have bounded aspect ratios). However, very often *most* of them satisfy this condition [ZS99]. To deal with these situations, Agarwal and Bern [AB99] asked whether Theorem 0 can be generalized as follows. Let $\mathcal{T}$ be a family of n triangles in the plane, whose *average* aspect ratio is bounded by a constant. That is, $\sum_{i=1}^{n}\frac{1}{\alpha_i} = O(n)$, where α_i denotes the smallest angle of the i-th triangle. Is it then true that $\mathcal{T}$ determines only $O(n)$ (or nearly a linear number of) holes? Theorems 2 and 3 answer this question in the negative. Indeed, let

$$\alpha_i = \begin{cases} \frac{1}{\sqrt{n}} & \text{if } 1 \le i \le \sqrt{n} \\ 1 & \text{if } \sqrt{n} < i \le n. \end{cases}$$

Then we have $\sum_{i=1}^{n} \frac{1}{\alpha_i} < 2n$, but, according to the second statement of Theorem 2, the number of holes can be as large as $\Omega(n^{3/2})$. This is true even for wedges (Theorem 3). The first statement of Theorem 2 shows that this bound is tight, apart from the logarithmic factor.

In the case when some of the wedges have angles very close to π, Theorem 3 is not sufficiently tight. A more careful analysis can account more precisely for the contribution of the convex wedges with angles close to π:

Theorem 4. *Let $\mathcal{T}$ be a family of n wedges in the plane with angles $0 < \alpha_1 \leq \cdots \leq \alpha_n < 2\pi$. Let $l \in [0, n]$ be the largest integer satisfying $\alpha_l < 3\pi/4$, and let $m \in [l, n]$ be the largest integer satisfying $\alpha_m < \pi$. (We set $l = 0$ if $\alpha_1 \geq 3\pi/4$, and $l = m = 0$ if $\alpha_1 \geq \pi$.) Let k be the largest integer with $\sum_{j=1}^{k} \alpha_j < \pi$, and, for any $1 \leq i \leq m$, let $k_i \in [0, i)$ be the largest integer such that $\sum_{j=1}^{k_i} \alpha_j < \pi - \alpha_i$.*

Then the boundary complexity of $\cup\mathcal{T}$ is $O(n + \sum_{i=1}^{l} k_i \log k_i + \sum_{i=l+1}^{m} k_i) = O(n + lk \log k + \sum_{i=l+1}^{m} k_i)$, where the sum is taken over all i with $k_i \neq 0$. Furthermore, there is a family of n wedges with angles $\alpha_1, \ldots, \alpha_n$, which determines $\sum_{i=1}^{m'} k_i + m' + 1$ holes, where $m' = \min\{m, n-1\}$.

The rest of the paper is organized as follows. The proofs of Theorems 1, 2-3, and 4 are presented in Sections 2-3, 4, and 5, respectively. The last section contains some combinatorial and algorithmic consequences of the main results.

2. Reduction to rhombs and assignment of holes

By a *polygon* we mean a simply connected (bounded or unbounded) region in the plane, whose boundary consists of a finite number of straight-line segments and possibly two half-lines. A family of polygons is said to be in *general position*, if no three lines supporting different sides of the polygons pass through the same point. We say that a *point* is *incident* to a hole H, if it lies on the boundary of H.

Given a family $\mathcal{P}$ of polygons in the plane, let $h(\mathcal{P})$ and $H(\mathcal{P})$ denote the number of holes determined by $\mathcal{P}$ and the minimum number of non-overlapping convex polygons the union of these holes can be partitioned into, respectively. Furthermore, let $c(\mathcal{P})$ stand for the number of concave angles of $\Re^2 \setminus \cup\mathcal{P}$, the union of the holes.

Lemma 2.1. *For any family $\mathcal{P}$ of polygons in the plane, we have*

$$h(\mathcal{P}) \leq H(\mathcal{P}) \leq h(\mathcal{P}) + c(\mathcal{P}).$$

Proof: The lower bound on $H(\mathcal{P})$ follows from the fact that to cover each hole we need at least one convex set. To establish the upper bound, we show that every hole with k concave vertices can be partitioned into $k+1$ convex sets. In the case $k = 0$, the hole itself is convex. For $k > 0$, it is enough to observe that the total number of concave vertices decreases by cutting the hole into two along the angular bisector at a concave vertex. □

Lemma 2.2. *Let $\mathcal{P}$ and $\mathcal{P}'$ be two families of polygons in the plane such that $\cup\mathcal{P}' \subseteq \cup\mathcal{P}$ and any segment connecting two points of $\Re^2 \setminus \cup\mathcal{P}$ which intersects $\cup\mathcal{P}$ also intersects $\cup\mathcal{P}'$.*

Then $H(\mathcal{P}) \leq H(\mathcal{P}')$.

Proof: For any partition of $\Re^2 \setminus \cup\mathcal{P}'$ into a family $\mathcal{C}$ of convex sets, $\{C \setminus \cup\mathcal{P} \mid C \in \mathcal{C}\}$ forms a partition of $\Re^2 \setminus \cup\mathcal{P}$ into convex sets. Note that we needed the condition in the lemma to guarantee that each member of the latter family is convex. □

For the rest of this section we fix an angle $\delta \leq \pi/3$ and set $k := \lceil 2\pi/\delta \rceil$. Clearly, we have $k \geq 6$. Fix any straight line ℓ_0 in the plane. We say that a *segment* (line, half-line) is *canonical* if the angle between its supporting line and ℓ_0 is an integer multiple of π/k. A *rhomb* is called *canonical* if (a) all of its sides are canonical, and (ii) two of its angles are equal to π/k.

Next we show that Theorem 1 can be reduced to

Theorem 2.3. *Any family $\mathcal{R}$ of n canonical rhombs in the plane determines $O(nk \log k)$ holes.*

Proof of Theorem 1 (using Theorem 2.3): Let $\mathcal{T}$ be a family of n δ-fat triangles in the plane. Consider a vertex A of a triangle $ABC \in \mathcal{T}$. By the choice of k, there are at least two canonical half-lines emanating from A, whose initial segments belong to ABC. Therefore, we can pick a point A' on the segment BC and two other points, D and D', in ABC such that $R_A = ADA'D'$ is a canonical rhomb whose angles at A and A' are equal to π/k. Let R_A denote such a rhomb. Similarly, we can define two other canonical rhombs, R_B and R_C, within the triangle ABC. (See Figure 1.)

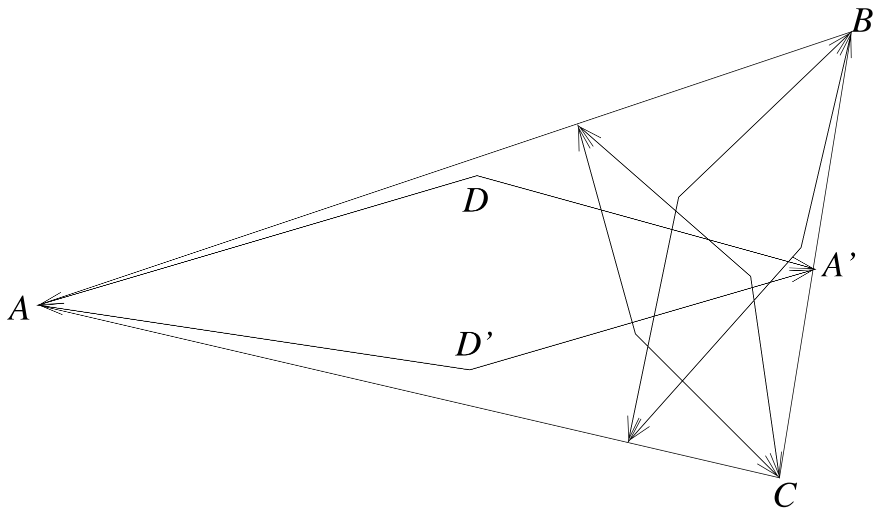

Substituting three canonical rhombs for a triangle

Figure 1

Let $\mathcal{T}'$ denote the family obtained from $\mathcal{T}$ by replacing every triangle $ABC \in \mathcal{T}$ by the three corresponding canonical rhombs, R_A, R_B, and R_C. By Theorem 2.3, $\mathcal{T}'$ determines $O(nk \log k)$ holes, i.e., $h(\mathcal{T}') = O(nk \log k)$. Now Lemmas 2.1 and 2.2 imply that

$$h(\mathcal{T}) \le H(\mathcal{T}) \le H(\mathcal{T}') \le h(\mathcal{T}') + c(\mathcal{T}').$$

In other words, the number of holes determined by $\mathcal{T}$ may be larger than the number of holes determined by $\mathcal{T}'$, but the difference cannot exceed the total number of concave corners (vertices) in all holes determined by $\mathcal{T}'$. However, every such corner corresponds to a vertex of one of the $3n$ rhombs defined above, so the difference is at most $12n$. $\Box$

In the rest of this and the next section, we establish Theorem 2.3. We may and will assume without loss of generality that the rhombs in $\mathcal{R}$ are in general position. We use the term *edge* only for the sides of the rhombs in $\mathcal{R}$. We call two edges *homothetic* if they are corresponding sides of two homothetic rhombs. We orient every edge e of a rhomb toward its vertex of angle π/k. This vertex is called the *apex* of e. The subsegments of e inherit the orientation of e.

Let e be an edge and let S be a homothecy class of edges not containing e. There may be several holes whose boundaries contain a piece of e next to a piece of some element of S. The first and last such holes along e are said to be *extreme*. Since there are $4n$ edges in at most $4k$ homothecy classes, the number of extreme holes is at most $32kn$. We call the non-extreme holes *intermediate*. Clearly, it is sufficient to bound the number of intermediate holes. As every hole incident to a vertex of a rhomb is extreme, the intermediate holes are convex.

Let H be an intermediate hole. Consider two consecutive segments, a and b, on the boundary of H, belonging to the edges e and f, respectively, and denote their common endpoint by X. Let P and Q denote the apices of e

and f, respectively. Suppose that b is oriented away from X. We say that H is a *hole assigned to* e if we have $|X - P| \le |X - Q|$. The distance $|X - P|$ is called the *depth of* H *along* e or, if it leads to no confusion, simply the *depth of* H. We say that H is an *in-hole* or an *out-hole* assigned to e, depending on whether a is oriented toward X or away from X. See Figure 2.

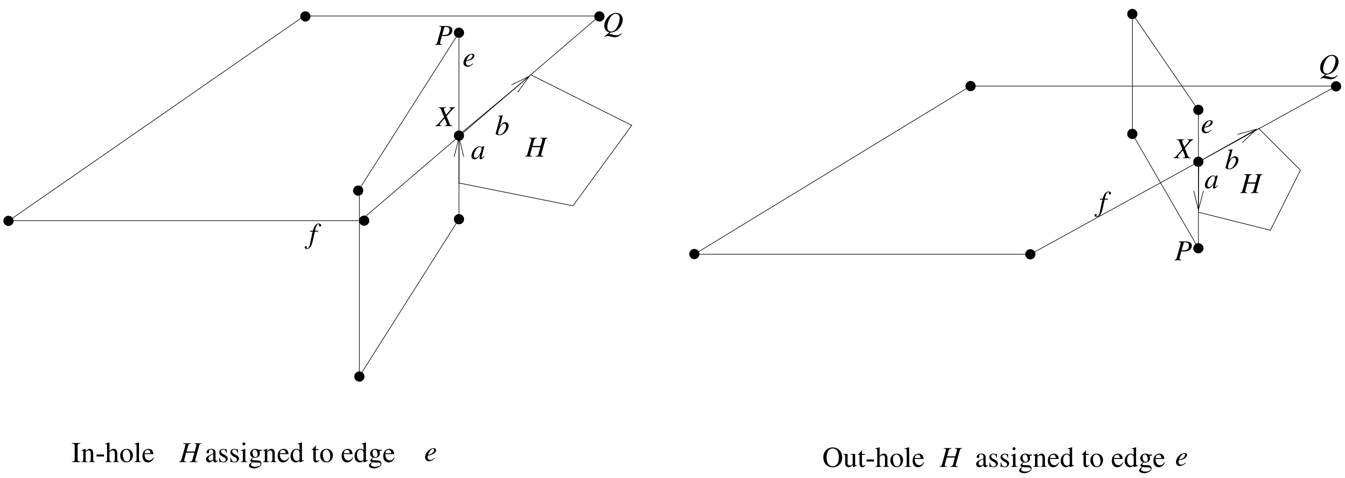

Figure 2

Lemma 2.4. *Every intermediate hole is assigned to at least one edge.*

Proof: Let H be an intermediate hole. If the segments bounding H are not cyclically oriented, we find two consecutive segments, both oriented away from their common endpoint. In this case, H is an out-hole assigned to one of the edges containing these two segments.

Suppose that the segments forming the boundary of H are cyclically oriented. Let $X_1, X_2, \ldots, X_k = X_0$ denote the vertices of H in this cyclic order, and let P_i denote the apex of the edge containing $X_{i-1}X_i$ $(1 \le i \le k)$. Set $P_{k+1} := P_1$. For every $1 \le i \le k$, if H is not an in-hole assigned to the edge containing $X_{i-1}X_i$, then we have

$$|X_i - P_i| > |X_i - P_{i+1}| = |X_i - X_{i+1}| + |X_{i+1} - P_{i+1}|.$$

Summing up these inequalities, we obtain

$$\sum_{i=1}^{k} |X_i - P_i| > \sum_{i=1}^{k} |X_i - P_i| + \mathrm{Per}(H),$$

where $\mathrm{Per}(H)$ stands for the perimeter of H. This contradiction proves that H is an in-hole assigned to an edge supporting one of its sides. $\Box$

Next we show that the depths of the in-holes along an edge are $\Omega(1/k)$ apart in a logarithmic scale, and the same is true for the depths of the out-holes. (The depth of an in-hole can be arbitrarily close to the depth of an out-hole though.) More precisely, we have:

Lemma 2.5. *Let H be an in-hole (out-hole) assigned to an edge e, whose depth is d. Then the depth of no other hole assigned to e is between d and $\left(1 - \frac{1}{k}\right) d$ (respectively. between d and $\left(1 + \frac{1}{k}\right) d$).*

Proof: Let a and b be two consecutive segments of the boundary of H causing H to be assigned to e. Let X be their common endpoint, f the edge containing b, and let P and Q be the apices of e and f, respectively. Clearly, a is on e oriented towards X and b is oriented away from X. The depth of H is $d = |X - P|$. Since H is not an extremal hole, e must cut through the rhomb belonging to the edge f. The length of the piece of XP covered by this rhomb is at least

$$|X - Q| \sin \frac{\pi}{k} \geq |X - P| \sin \frac{\pi}{k} > \frac{d}{k}.$$

Thus, the depth of no hole along e can belong to the interval $((1 - 1/k)d, d)$. The corresponding statement for out-holes can be proved similarly. □

3. Base points

Theorem 2.3 would immediately follow from Lemma 2.5, if we could show that the ratio of the largest and smallest depths of an in-hole (and out-hole) along the same edge is bounded by a polynomial of k. It is not hard to see that this holds for wedges rather than rhombs (and this can be used to give a direct proof of Theorem 3), but the general statement is false. We prove instead that the depths of the intermediate holes (in-holes and out-holes) assigned to a given edge fall into a small number of short intervals.

To formulate our result precisely, we need some preparation. We assign at most *three* so-called *base points* to each edge, according to the following definition.

Definition 3.1. (*Base points*) Let $e = NP$ be an edge of a rhomb $\Delta = NPSU \in \mathcal{R}$, and let e be oriented towards P. Let Y denote the point of e closest to N that does not belong to the interior of any member of $\mathcal{R}$. Let both P and Y be assigned to e as *base points*. Further, let h denote the oriented half-line (ray) starting at N and passing through U. Consider all edges homothetic to e that either intersect h beyond U, or intersect both h and the edge SP. If there is no such edge, then no other base point is assigned to e. Otherwise, let e' denote the edge with this property that intersects h closest to N. If e' intersects h beyond U, then we say that e' is *far from* e, and set $Y' := e' \cap h$. If e' intersects the segment PS, then we say that e' is *close to* e, and set $Y' := e' \cap PS$. In both cases, Y' is the third base point assigned to e, which will be referred to as the *base point forced by* e. Note that this third point is *not* on e. (See Figure 3.)

The *depth of a base point* Y sitting on an edge e is defined as the distance between Y and the apex of e.

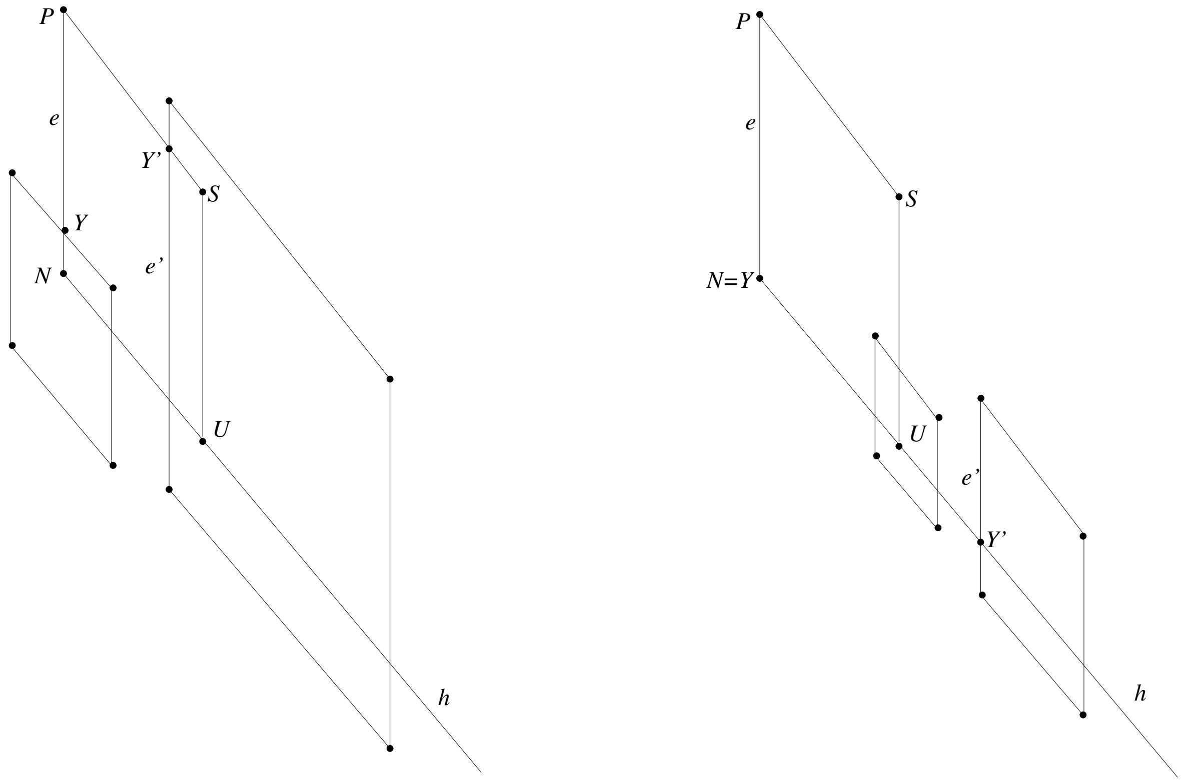

Y on e and
Y' on e' forced by e
e' close to e

Y on e and
Y' on e' forced by e
e' far from e

Figure 3

Using the above notation, no point of the open interval YN can be incident to a hole. Thus, the depth of a hole assigned to e cannot exceed the depth of the base point Y on e. Consequently, the depth of each hole assigned to e is between the depths of some pair of consecutive base points on e. Although there may be many base points along the same edge, the total number of base points is at most $12n$.

Lemma 3.2. *If there are two consecutive base points on an edge e with depths d_1 and $d_2 > k^2 d_1$, then there is at most one intermediate hole assigned to e with depth belonging to the interval $(d_1, d_2/k^2)$.*

Proof: Let $e = NP$ be an edge oriented toward P. Choose two consecutive base points on e with depths $d_1 < d_2$, respectively. Let H be an (intermediate) hole assigned to e with depth d satisfying $d > d_1$. We prove that for every other hole assigned to e, whose depth d' is larger than d, $d' \geq d_2/k^2$ holds.

As before, let a and b be the two consecutive segments on the boundary of H causing H to be assigned to e. Let the segment b belong to the edge f with apex Q. Let X be the common endpoint of a and b. Clearly, a is on e and b is oriented away from X.

The proof is based on the following claim:

Claim: *There is a base point Y on e of depth at least the depth of H such that the angle PYQ is at least $\frac{\pi}{2k}$.*

Proof: We assume without loss of generality that e is vertical and that its upper endpoint is P. As b is oriented away from X, and H is not an extreme hole, we find another edge e_0, homothetic to e, which intersects XQ. Consider the sequence of edges $e_0, e_1, e_2, \ldots, e_t$, where e_i is the edge homothetic to e, containing the base point forced by e_{i-1} $(1 \leq i \leq t)$. The last edge of this sequence, e_t, forces no base point. (See Figure 4.) Let e_i be the edge of the rhomb $\Delta_i = N_i P_i S_i U_i \in \mathcal{R}$ with $e_i = N_i P_i$ oriented toward P_i, and let h_i be the half-line starting at N_i and passing through U_i $(0 \leq i \leq t)$. Notice that h_i intersects e_{i+1} for $i < t$, and denote this point of intersection by X_{i+1}. The intersection of f and e_0 is denoted by X_0. Let Π be the directed polygonal path $QX_0N_0X_1N_1\dots X_tN_t$ followed by the half-line h_t. Since all segments of Π (except maybe QX_0) are pointing downward and are at least as steep as f, X must lie above Π.

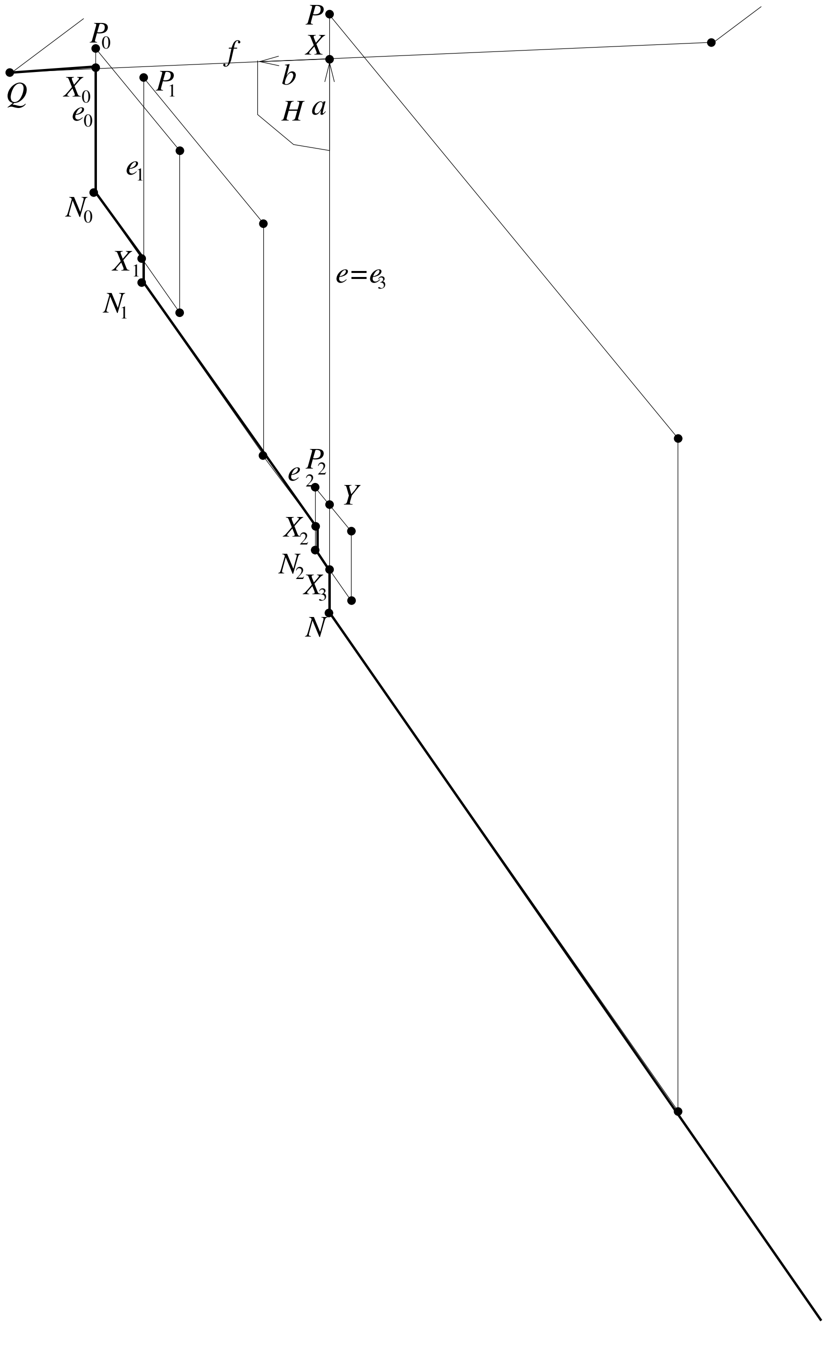

The construction of the path Π (the bold path)

Figure 4

We distinguish two cases.

Case I: $e = e_{i_0}$ for some $1 \le i_0 \le t$.

The point $X_{i_0} \in \Pi$ lies below X on the same vertical edge e. Let Y denote the base point on e forced by e_{i_0-1}. If e is far from e_{i_0-1}, the point Y coincides with X_{i_0}, so it is below X. If e is close to e_{i_0-1}, the open interval

$X_{i_0}Y$ belongs to the interior of Δ_{i_0-1}, and the point X incident to H cannot lie on this interval. Therefore, in this case Y cannot be above X either.

For $0 \le i \le t$, denote by σ_i the strip between the parallel lines containing P_iS_i and N_iU_i. Notice that if e_i is close to e_{i-1} ($1 \le i \le t$), then σ_i contains σ_{i-1}.

Let $1 \le j_1 < j_2 < \cdots < j_s$ be the sequence of indices $1 \le j \le i_0$, for which e_j is far from e_{j-1}, and let $j_0 = 0$. For $1 \le i \le s$, the portion of Π between $X_{j_{i-1}}$ and X_{j_i} is contained in σ_{j_i-1}. Thus, $X_{j_{i-1}}$ is below the line $P_{j_i-1}U_{j_i-1}$, while X_{j_i} is above the same line. This implies that the angle $X_{j_{i-1}}X_{j_i}P_{j_i}$ is larger than $\pi/(2k)$. If $j_s = i_0$, we have $X_{j_s} = X_{i_0} = Y$. Otherwise, $X_{j_s} \in \sigma_{i_0-1}$, while Y is on the upper boundary of this strip, so the angle $X_{j_s}YP$ is at least π/k. Since we have bounded the slope of each portion of Π, combining these bounds, it follows that the angle PYQ is at least $\pi/(2k)$. This completes the proof of the Claim in Case I.

Case II: $e \notin \{e_i : 1 \le i \le t\}$.

The path Π now cannot cross e. Otherwise, let X' be this point of intersection and suppose X' lies on the half line h_i ($i \le t$). When determining the base point forced by e_i we must consider the edge e, since the interval $X'P$ contains the point X incident to the hole H, and thus cannot be entirely covered by Δ_i. So we either have $e = e_{i+1}$ or X_{i+1} lies between N_i and X'. Both are contradictory to our assumptions, the letter since X' is not on Π in this case.

So e must lie entirely above Π. Let l denote the line containing e, and let Z be the intersection point $l \cap \Pi$. Suppose that Z belongs to the non-vertical segment of Π starting at the point N_{i_0}, for some $i_0 \le t$.

Let j be the largest index between (and including) 1 and i_0, for which e_j is far from e_{j-1}. If there is no such index, set $j = 0$. It follows in exactly the same way as in the previous case that the angle QX_jP_j is at least $\pi/(2k)$. Let Y denote the point of e closest to N that does not belong to the interior of any rhomb in $\mathcal{R}$. Recall that, according to Definition 3.1, Y is a base point, and notice that it does not lie above X.

If l intersects Δ_{i_0}, the point X_j is below the line $P_{i_0}S_{i_0}$, while Y' is on this line or above it. Therefore, in this case the angle X_jYP is at least π/k. If l does not intersect Δ_{i_0}, consider the portion of Π between X_j to Z, and notice that it lies in the strip between the lines $P_{i_0}S_{i_0}$ and $N_{i_0}U_{i_0}$. Thus, X_j is below the line $P_{i_0}U_{i_0}$ but Z and thus Y is above the same line. This implies that in this case the angle X_jYP is also at least $\pi/(2k)$. Combining this with the same lower bound for the angle QX_jP_j, we obtain that the angle PYQ is at least $\pi/(2k)$. This completes the proof of the Claim. □

Now we finish the proof of Lemma 3.2:

The depth of the base point Y is larger than d_1 thus we have $|P-Y| \ge d_2$. Using the law of sines for the triangle QYX, we obtain $|Y-X| \le |Q-X|/\sin\frac{\pi}{2k}$. Since H was assigned to e, we have that $|P-X| \le |Q-X|$. Therefore,

$$d_2 \le |P-Y| = |P-X| + |X-Y| \le$$

$$\le \left(1 + \frac{1}{\sin\frac{\pi}{2k}}\right)|Q-X| \le k|Q-X|.$$

Note that the rhomb in $\mathcal{R}$ belonging to the edge f covers an interval of e, whose length is at least $|Q-X|\sin(\pi/k) \ge d_2/k^2$. As X is an endpoint of this interval, any hole assigned to e, whose depth d' is larger than the depth $d = |P-X|$ of H, must satisfy the inequality $d' > d_2/k^2$, hence the lemma is true. □

Proof of Theorem 2.3: According to Lemma 2.5, the depths of the intermediate holes assigned to an edge are separated from each other by $\Omega(1/k)$ in a logarithmic scale. By Lemma 3.2, if an edge has $c+1$ base points, then the depths of all holes assigned to this edge, except for at most c of them fit in c intervals, each of length $2\log k$ in logarithmic scale. Thus this dege has $O(ck\log k)$ holes assigned to it. Thus by Lemma 2.4 the total number of intermediate holes is $O(nk\log k)$. Since the number of extreme holes is $O(nk)$, Theorem 2.3 follows. □

4. Generalizations—Proofs of Theorems 2 and 3

Proof of Theorem 2: Obviously, the triangles $T_{k+1}, \ldots, T_n$ are $\pi/(k+1)$-fat. By Theorem 1, they determine $O(nk \log k)$ holes. By adding the first k triangles, we increase the number of intersection points and, therefore, the number of holes by $O(nk)$.

As for the construction, let $l = \lfloor k/2 \rfloor$ and arrange $T'_1, \ldots, T'_l$ so that any two intersect in a single point v, and all of them are contained in a right angle wedge W with apex v. This is possible, because we have $\sum_{i=1}^{l} \alpha_i < \pi/2$. Let ε be the distance of v from the nearest other vertex of $T'_1, \ldots, T'_l$. We place T'_i for $i > l$ so that

- T'_i meets both rays bounding W but does not contain v;
- $T'_i \cap W$ is contained in the ball of radius ε around v;
- the sets $T'_i \cap W$ are pairwise disjoint for $i > l$.

All of these conditions can be satisfied by placing the triangles one by one so that we put a vertex of the next triangle, corresponding to an acute angle, sufficiently close to v. The obtained configuration has $(l-1)(n-l+1)$ holes inside W. This quantity is $\Omega(nk)$, as required, unless $l = 1$. In the latter case, we arrange the triangles T'_i $(i > 1)$ so that all of them have a point in common outside T'_1, and their intersections with T'_1 are distinct single points. These triangles determine at least $n - 1$ holes. □

We continue with Theorem 3. (Notice that it is a simple special case of Theorem 4.)

Proof of Theorem 3: First we prove $O(nk \log k)$ bound for the *number of holes* determined by the convex wedges in $\mathcal{T}$. As we have indicated before, the direct proof of this bound along the lines described in Section 2 is simpler than the proof of Theorem 1. However, at this point it is more convenient to deduce it from Theorem 2. Let h be the number of holes determined by the wedges. First we split each wedge with an obtuse angle into two congruent wedges. Then we replace each wedge by a triangle, intersecting it with a half-plane that contains all intersection points between the boundaries of the original wedges. We make sure that all new angles introduced exceed $\pi/4$. Following this procedure, we obtain a family of at most $2n$ triangles that determine at least $h - n + 1$ holes. The value k for this new family (as defined in Theorem 2) is at most 3 larger than the corresponding value for the original family of wedges. Applying Theorem 2 to the triangles, we obtain the desired bound for wedges.

To prove the same upper bound for the *boundary complexity*, notice that $\alpha_i \geq \pi/(k+1)$ for $i > k$. As in the proof of Theorem 2, we can disregard the first k wedges, because their contribution to the boundary complexity is at most $4kn$. We proceed as in [MP94]. We partition the remaining wedges into $2k+2$ classes so that all wedges belonging to the ith class contain in their interior a half-line having a $2\pi/(2k+2)$ positive angle from a reference direction. Now every vertex of the union of all wedges in a given class is either the apex of a wedge or is the last vertex along one of the open half-lines bounding the wedges. Thus, the boundary complexity of this union is linear in the number of wedges in the class. Using the Combination Lemma of [EG90] (see also Lemma 2.1 in [MP94]) to merge the classes in a binary tree-like fashion, it follows that the boundary complexity of the union of all wedges in all classes is $O(nk \log^2 k)$. To get rid of the extra $\log k$ factor, we consider the family of wedges $\mathcal{T}'$ we obtain at an intermediate step through the combination process. It is the union of some j of the original $2k+2$ families. We can make sure, that these are j consecutive families. Applying an affine transformation, if necessary, we can achieve that the angle of every wedge belonging to these families is $\Omega(1/j)$. Since such a transformation does not change the number of holes, we obtain the better bound $O(mj \log j)$ for the number of holes determined by $\mathcal{T}'$, where m is the number of wedges in $\mathcal{T}'$. Using the combination lemma with this better bound, we conclude that the boundary complexity of $\mathcal{T}$ is $O(nk \log k)$.

To verify the last statement of Theorem 3, we use almost the same construction as in the proof of Theorem 2. The only difference is that now we have to start with a wedge W whose angle is smaller than $\pi - \alpha_n$, otherwise no wedge of angle α_n could intersect it in the required manner. Let $l \in [0, n/2]$ be the largest integer satisfying $\sum_{i=1}^{l} \alpha_i < \pi - \alpha_n$. Clearly, we have $l \geq \min\{\lfloor(1-\alpha_n/\pi)k\rfloor, n/2\}$. Select l wedges in W with angles $\alpha_1, \ldots, \alpha_l$ such that the intersection of any two is the apex v of W. Then choose $n-l$ wedges of angles $\alpha_{l+1}, \ldots, \alpha_n$ such that

- each of them intersects both boundary half-lines of W;
- none of them contains v;
- their intersections with W are pairwise disjoint bounded sets.

The resulting family determines $(l-1)(n-l+1)$ holes in W. This is $\Omega((\pi - \alpha_n)nk)$, unless $l \leq 1$. In the latter case, we use a trivial construction similar to the one described at the end of the proof of Theorem 2: we pick $n-1$ wedges that intersect the remaining wedge in distinct single points. This family determines at least $n-1$ holes. $\Box$

5. Wedges of angles close to π—Theorem 4

This section is devoted to the proof of Theorem 4.

First we establish the upper bound $O(B)$ for the *number of holes* in $\mathcal{T}$, where $B = m + lk\log k + \sum_{i=l+1}^{m} k_i$. The same bound on the *boundary complexity* then follows directly from the Combination Lemma of [EG90]: by Theorem 3, the boundary complexity of the subfamily consisting of the l smallest wedges in $\mathcal{T}$ is $O(lk\log k)$, the boundary complexity of the subfamily consisting of the next $m-l$ wedges of $\mathcal{T}$ is $O(m)$, while the boundary complexity of the family of the $n-m$ largest wedges of $\mathcal{T}$ (whose angles are at least π) is clearly $O(m)$, as it determines a single convex hole.

We proceed as in the proof of Theorem 1, but now we have to deal with different angles.

Fix a reference direction and say that a wedge is *small* if its angle is $\pi/2^s$ for some integer $s \geq 2$ and if the angles between its boundary rays and the reference direction are integer multiples of $\pi/2^s$. A wedge is *large* if its angle is $\pi - 2\pi/2^s$ for some integer $s \geq 2$ and if the angles between its boundary rays and the reference direction are integer multiples of $\pi/2^s$.

Let W_i be the wedge in $\mathcal{T}$ of angle α_i and let P_i its apex. For $i \leq l$ define W'_i to be the maximal small wedge with apex P_i that is contained in W_i. For $l < i \leq m$ let W'_i be the maximal large wedge with apex P_i that is contained in W_i. Let $\mathcal{T}' = \{W'i | 1 \leq i \leq m\}$. By Lemmas 2.1 and 2.2, we have

$$h(\mathcal{T}) \leq H(\mathcal{T}) \leq H(\mathcal{T}') \leq h(\mathcal{T}') + m.$$

Thus, when passing from $\mathcal{T}$ to $\mathcal{T}'$, the number of holes cannot decrease by more than m.

Let α'_i denote the angle of the wedge W'_i ($i \leq m$), and let k' and k'_i be defined for $\mathcal{T}'$ in exactly the same way, as k and k_i were defined for $\mathcal{T}$ ($1 \leq i \leq m$). Notice that $\alpha_i/4 < \alpha'_i \leq \alpha_i$, so $k' \leq 4k+3$. Furthermore, we have $\pi - \alpha'_i < 4(\pi - \alpha_i)$, which implies that $k'_i \leq 16k_i + 15$. Thus, it is sufficient to prove the desired upper bound on the number of holes determined by the modified family $\mathcal{T}'$. For notational convenience, from now on we assume that $\mathcal{T}' = \mathcal{T}$, i.e., that $\mathcal{T}$ consists of l small and $m-l$ large wedges.

We use the terms *wedge, apex, ray,* and *hole* only for the wedges in $\mathcal{T}$, their apices, their rays, and for the holes determined by them. Assume without loss of generality that no three rays have a point in common. We say that two rays are *homothetic* if they are corresponding sides of two wedges that are translates of each other.

A triplet (X, r_1, r_2) is called a *vertex* if X is the intersection point of two distinct rays, r_1 and r_2, and it lies on the boundary of a hole. Our plan is to define several special types of vertices, to bound the number of vertices of each type separately, and finally to bound the number of holes by showing that each hole has a point on its boundary that appears in a vertex of some special type. A vertex (X, r_1, r_2) is said to be *small* (*large*), if both of the wedges supporting r_1 and r_2 are small (resp., large) wedges. The number of small and large vertices can be bounded by the boundary complexity of the family consisting of all small or all large wedges in $\mathcal{T}$, and, by Theorem 3, we have:

Claim 5.1. *There are $O(lk \log k)$ small and $O(m)$ large vertices.*

In what follows, when we consider a vertex (X, r_1, r_2), the wedges supporting r_1 and r_2 will be denoted by W_1 and W_2, respectively. The apices of W_1 and W_2, will be denoted by P_1 and P_2, respectively. If a vertex is neither small nor large, then it is said to be *hybrid*. For a hybrid vertex (X, r_1, r_2), one of W_1 and W_2 is small and the other is large. In this case, let i and j denote the indices of the corresponding small and large elements of $\mathcal{T}$, resp., with angles $\alpha_i = \pi/2^s$ and $\alpha_j = \pi - 2\pi/2^t$, resp., where $1 \leq i \leq l$, $l+1 \leq j \leq m$, $s \geq 2$ and $t \geq 2$.

A vertex (X, r_1, r_2) is called *strongly extremal* if it is the first or last vertex along r_1 or r_2. Obviously, we have

Claim 5.2. *There are at most $4m$ strongly extremal vertices.*

A hybrid vertex (X, r_1, r_2) is said to be *blunt* if $s < t$. Considering the canonical properties of the wedges in $\mathcal{T}$, we see that if a blunt vertex (X, r_1, r_2) is not strongly extremal, then the boundaries of the wedges meet in three vertices, and X is the middle one (along either boundary). Furthermore, if W_1 is large we find that r_2 is perpendicular to the angular bisector of W_1 (see Figure 5).

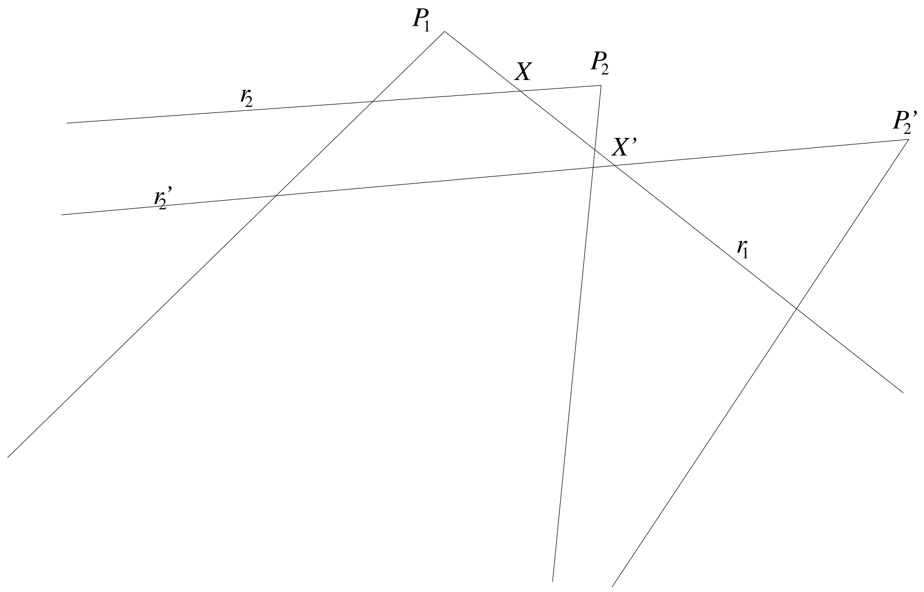

Blunt vertices
Figure 5

We claim that along a ray bounding a large vertex there is at most one blunt vertex that is not strongly extremal. Suppose, for contradiction, that (X, r_1, r_2) and (X, r_1, r_2') are both blunt and not strongly extremal and W_1 is a large wedge with apex P_1. We may also assume that X belongs to the interval $X'P_1$. Here the ray r_2' is parallel to r_2, so the wedges W_1 and W_2 together cover the part of r_2' on one side of X' (see Figure 5). Therefore, (X', r_1, r_2') must be strongly extremal, a contradiction. Therefore, the number of blunt vertices that are not strongly extremal is at most $2m$, which implies

Claim 5.3. *The total number of blunt vertices is $O(m)$.*

A hybrid vertex (X, r_1, r_2) is said to be *sharp* if $s > t + \log k$.

Claim 5.4. *The number of sharp vertices is $O(\sum_{j=l+1}^{m} k_j)$.*

Proof: Let (X, r_1, r_2) be sharp. We have $i \leq k_j$, for otherwise none of the angles of the first $k_j + 1$ wedges in $\mathcal{T}$ would exceed $\pi/2^s < \pi/2^{t+\log k}$, and their sum would be less than $\pi - \alpha_j = 2\pi/2^t$, contradicting the definition of k_j. Thus, there are at most $O(k_j)$ sharp vertices involving the jth wedge in $\mathcal{T}$ ($j > l$), and Claim 5.4 follows. □

A vertex (X, r_1, r_2) is called *extremal*, if it is the first or the last vertex along r_1 among all vertices of the form (X', r_1, r_2') with r_2' being homothetic to r_2.

Claim 5.5. *The number of extremal vertices is $O(B)$.*

Proof: It is sufficient to establish this bound for those hybrid extremal vertices which are neither strongly extremal, nor blunt, nor sharp. Fix such a vertex (X, r_1, r_2). We distinguish two cases.

Case i: W_1 is small and W_2 is large. For a given small wedge W_1 (which determines s), there are at most $\log k + 1$ possible values for t, without the vertex being sharp or blunt. For a given t, there are at most four homothecy classes possible for W_2 without the vertex being strongly extremal. Thus, each ray of a small wedge is involved in at most $8 \log k + 8$ extremal vertices satisfying the condition of Case i. Therefore, the total number of extremal vertices of this type is $O(l \log k)$.

Case ii: W_1 is large and W_2 is small. Fix W_1 to be the jth wedge in $\mathcal{T}$. This determines the value of t. For a given value $s > t$, there are less than 2^{s-t+2} possible homothecy classes for W_2 without the vertex being strongly extremal. Notice that if z among these classes for all $s > t$ are nonempty then the sum of $\lceil z/2 \rceil$ of the smallest angles is less than $\pi - \alpha_j = 2\pi/2^t$. This implies $z = O(k_j)$. The total number of extremal vertices involving r_1 that satisfy the condition in Case ii is $O(k_j)$, and the total number of all extremal vertices of type 2 is $O(\sum_{i=l+1}^{m} k_i)$. □

A vertex (X, r_1, r_2) is called *covered*, if the interval XP_1 is at least as long as the interval XP_2 and P_1 and W_2 lie on different sides of r_2.

Claim 5.6. *The number of covered vertices is $O(B)$.*

Proof: It is enough to consider those hybrid covered vertices that are neither extremal, nor sharp, nor blunt. Let (X, r_1, r_2) be such a covered vertex. As (X, r_1, r_2) is not extremal, there exists another vertex (X', r_1, r_2') for which X' belongs to the interval XP_1 and the ray r_2' is homothetic to r_2. Here the wedge W_2', whose boundary ray is r_2', covers all of r_2 except an interval P_2Q. Our goal is to show that W_1 covers a sufficiantly large subinterval XY of P_2Q. Considering these intervals XY for different covered vertices involving r_2, we find that they cover any point at most twice. Therefore, any lower bound on the lengths of these intervals yields an upper bound on their number. As in the proof of the previos claim, we distinguish two cases.

Case i: W_1 is small and W_2 is large. By the relative position of the wedges W_1, W_2, and W_2', we have $|P_2Q| = O(|XY|/2^{s-t})$. This implies that r_2 appears in $O(k_j + 1)$ covered vertices that satisfy the condition of Case i. Therefore, the total number of covered vertices of this type is $O(m + \sum_{j=l+1}^{m} k_j)$.

Case ii: W_2 is small and W_1 is large. Here we have to use that (X, r_1, r_2) is not sharp, so that $s \leq t + \log k$. As in Case i, one can show that $|P_2Q| = O(k|XY|)$. Hence, r_2 appears in $O(k)$ covered vertices satisfying the condition of Case ii, and the total number of covered vertices of this type is $O(lk)$. □

The proof of the upper bound in Theorem 4 can now be completed by showing that each hole determined by $\mathcal{T}$ has a point on its boundary that appears as the leading term of a strongly extremal or a covered vertex. Indeed, if no such strongly extremal vertex exists, then the hole must be a (bounded) convex polygon. Consider the orientation of the edges inherited from the rays oriented toward their apices. If it is not *cyclic*, we find a vertex X with two outgoing edges. Obviously, either (X, r_1, r_2) or (X, r_2, r_1) is covered, where r_1 and r_2 are the two rays containing X. We deal with the cyclically oriented case in exactly the same way as in the proof of Theorem 1: assuming that no vertex around the hole is covered, we obtain several inequalities, whose sum gives a contradiction. This concludes the proof of the upper bound in Theorem 4 on the number of holes as well as on the boundary complexity of $\mathcal{T}$.

It remains to describe a construction for the lower bound. Choose an $\epsilon > 0$ such that $\epsilon < (\pi - \alpha_i - \sum_{j=1}^{k_i} \alpha_j)/n$ for all i, and let ℓ_0 be a fixed horizontal line. By the *direction* of a half-line h from a point P on ℓ_0 we mean the angle between the half-lines h and the part of ℓ_0 to the right of P. We place the wedges W_i with angles α_i $(i = 1, \ldots, m')$ one by one, according to the following rules.

1. The apex of W_i is on ℓ_0 sufficiently to the right, so that all intersection points of boundaries of wedges already placed are outside W_i and to the left of it;
2. W_i is above ℓ_0;
3. the direction of the right ray of W_i is ϵ larger than the direction of the left ray of W_{k_i}, or it is simply ϵ if $k_i = 0$.

By the choice of k_i and ϵ, all the above requirements can be satisfied. The value $H = H(\{W_1, \ldots, W_{m'}\})$ of the resulting family (see the beginning of Section 2 for the definition) is the bound in Theorem 4 for the number of holes. For $i > m'$, we place a wedge of angle α_i such that it does not contain any intersection point between boundaries of W_j for $j \le m'$, but one of its boundary rays is on ℓ_0 and contains the apices of all W_j $(j \le m')$.

6. Concluding remarks, applications

As in [MP94], Theorems 1 and 2 yield the following upper bounds for the boundary complexity of a family of triangles.

Corollary 6.1. *The boundary complexity of any family of n δ-fat triangles in the plane is $O\left(\frac{n}{\delta}(\log\log n \log\frac{2}{\delta} + \log^2\frac{2}{\delta})\right)$. Moreover, the boundary of the union can be computed in time $O\left(\frac{n\log n}{\delta}(\log\log n \log\frac{2}{\delta} + \log^2\frac{2}{\delta})\right)$.*

Corollary 6.2. *Let $\mathcal{T} = \{T_1, \ldots, T_n\}$ be a family of $n > 1$ triangles in the plane, and let α_i denote the smallest angle of T_i $(1 \le i \le n)$. Suppose $0 < \alpha_1 \le \alpha_2 \le \cdots \le \alpha_n$, and let $k \le n$ be the largest integer such that $\sum_{i=1}^{k} \alpha_i < \pi$.*

Then the boundary complexity of $\mathcal{T}$ cannot exceed $O(nk(\log\log n \log k + \log^2 k))$.

Plugging Theorem 3 into the analysis of the running time of the algorithm described in [ER93], we obtain the following two results.

Corollary 6.3. *The union of a family of n δ-fat wedges in the plane can be computed in time $O\left(\frac{n}{\delta}\log^2\frac{2}{\delta} + n\log n\right)$.*

A line ℓ is called a *separator* for a family $\mathcal{S}$ of pairwise disjoint segments in the plane, if ℓ avoids all members of $\mathcal{S}$, and there is at least one member of $\mathcal{S}$ on both of its sides.

Corollary 6.4. *Given a family $\mathcal{S}$ of n line segments in the plane such that the ratio between the length of the shortest segment in $\mathcal{S}$ and the diameter of $\cup\mathcal{S}$ is at least $\delta > 0$, there is an algorithm which determines whether $\mathcal{S}$ admits a separator, and finds one if it exists, in $O\left(\frac{n}{\delta}\log\frac{n}{\delta}\log^2\frac{1}{\delta}\right)$ time and $O\left(\frac{n}{\delta}\log^2\frac{1}{\delta}\right)$ space.*

M. van Kreveld [K98] extended the definition of fatness to (not necessarily convex) simple polygon. He proved that every δ-fat simple polygon of k vertices can be covered by $O(k)$ δ-fat triangles, and such a covering can be constructed in $O(k\log k)$ time. Therefore, Theorem 1 generalizes to δ-fat simple polygons whose total number of sides is n.

A. Efrat [E99] introduced another generalization of the notion of fatness to compact connected regions of 'constant description complexity' depending on two real parameters. He established an upper bound on the boundary complexity of a system of 'fat' objects according to this definition. The dependence of his bounds on the parameters can be improved by using Theorem 1 instead of the results in [MP94].

References

[AB99] P. K. Agarwal and M. Bern: Open problem raised at International Conference on Discrete and Computational Geometry (Centro Stefano Franscini, ETH Zürich), Monte Verita, Switzerland, July 1999.

[AK94] P. K. Agarwal, M. J. Katz, and M. Sharir: Computing depth orders and related problems, in: *Algorithm Theory – SWAT '94, Lecture Notes in Computer Science* **824**, Springer-Verlag, Berlin, 1994, 1–12.

[AF92] H. Alt, 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, *Algorithmica* **8** (1992), 391–406.

[BK97] M. de Berg, M. Katz, F. van der Stappen, and J. Vleugels: Realistic input models for geometric algorithms, in: *Proceedings of 13th Annual Symposium on Computational Geometry,* ACM Press, 1997, 294–303.

[EG89] 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 and Computational Geometry* **4** (1989), 523–539.

[EG90] H. Edelsbrunner, L. Guibas, and M. Sharir: The complexity and the construction of many faces in arrangements of lines and of segments, *Discrete and Computational Geometry* **5** (1990), 161–196.

[E99] A. Efrat: The complexity of the union of (α, β)-covered objects, *Proceedings of the 15th Annual Symposium on Computational Geometry,* ACM Press, 1999, 134–142.

[EK98] A. Efrat and M. Katz: On the union of κ-curved objects, *Computational Geometry: Theory and Applications* **14** (1999), 241–254.

[EK97] A. Efrat, M. Katz, F. Nielsen, and M. Sharir: Dynamic data structures for fat objects and their applications, *Computational Geometry: Theory and Applications* **15** (2000), 215–227.

[ER93] 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* **3** (1993), 277–288.

[ES97] A. Efrat and M. Sharir: On the complexity of the union of fat objects in the plane, *Proceedings of the 13th Annual Symposium on Computational Geometry,* ACM Press, 1997, 104–172.

[GS93] L. Guibas and M. Sharir: Combinatorics and algorithms of arrangements, Chapter 1 in: *New Trends in Discrete and Computational Geometry* (J. Pach, ed.), Springer-Verlag, Berlin, 1993, 9–36.

[GJ97] P. Gupta, R. Janardan, and M. Smid: A technique for adding range restrictions to generalized searching problems, *Information Processing Letters* **64** (1997), 263–269.

[K97] M. J. Katz: 3-D vertical ray shooting and 2-D point enclosure, range searching, and arc shooting amidst convex fat objects, *Computational Geometry: Theory and Applications* **8** (1997), 299–316.

[KL86] K. Kedem, R. Livne, J. Pach, and M. Sharir: On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete and Computational Geometry* **1** (1986), 59–71.

[K98] M. van Kreveld: On fat partitioning, fat covering and the union size of polygons, *Computational Geometry: Theory and Applications* **9** (1998), 197–210.

[MP94] J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl: Fat triangles determine linearly many holes, *SIAM Journal of Computing* **23** (1994), 154–169.

[PT00] J. Pach and G. Tardos: On the boundary complexity of the union of fat triangles, in: *Proceedings of 41st Annual Symposium on Foundations of Computer Science (FOCS'00)*, 2000, 423–431.

[SS83] J. T. Schwartz and M. Sharir: On the "piano movers" problem I,II, *Communications in Pure and Applied Mathematics* **36** (1983), 345–398 and *Advances of Applied Mathematics* **4** (1983), 298–351.

[SS89] J. T. Schwartz and M. Sharir: A survey of motion planning and related geometric algorithms, in: *Geometric Reasoning* (D. Kapur and J. Mundy, eds.), MIT Press, Cambridge, MA, 1989, 157–169.

[SS90] J. T. Schwartz and M. Sharir: Algorithmic motion planning in robotics, in: *Handbook of Theoretical Computer Science* (J. van Leeuwen, ed.), Elsevier, Amsterdam, 1990, 391–430.

[S94] F. van der Stappen: *Motion Planning amidst Fat Obstacles (Ph. D. Thesis,* Faculteit Wiskunde & Informatica, Universiteit Utrecht, 1994.

[SH93] F. van der Stappen, D. Halperin, and M. Overmars: The complexity of the free space for a robot moving amidst fat obstacles, *Computational Geometry: Theory and Applications* **3** (1993), 353–373.

[WS88] A. Wiernik and M. Sharir: Planar realization of non-linear Davenport-Schinzel sequences by segments, *Discrete and Computational Geometry* **3** (1988), 15–47.

[ZS99] Y. Zhou and S. Suri: Analysis of a bounding box heuristic for object intersection, *Journal of the ACM,* **46** (1999), 833–858.