The Common Exterior of Convex Polygons in the Plane*

Boris Aronov[†] Micha Sharir[‡]

Abstract

We establish several combinatorial bounds on the complexity (number of vertices and edges) of the complement of the union (also known as the common exterior) of k convex polygons in the plane, with a total of n edges. We show:

1. The maximum complexity of the entire common exterior is $\Theta(n\alpha(k) + k^2)$.[1]
2. The maximum complexity of a single cell of the common exterior is $\Theta(n\alpha(k))$.
3. The complexity of m distinct cells in the common exterior is $O(m^{2/3}k^{2/3}\log^{1/3}(\frac{k^2}{m}) + n\log k)$ and can be $\Omega(m^{2/3}k^{2/3} + n\alpha(k))$ in the worst case.

1 Introduction

In this paper we establish several combinatorial bounds on the complexity of the *common exterior* (namely, the complement of the union) of a collection of k convex polygons in the plane, with a total of n edges. The arrangement of such a collection of polygons can be viewed as a special case of an arrangement of n segments, but we prefer to regard it as a generalization of an arrangement of k segments, where each segment is replaced by a convex polygon. Arrangements of segments have been studied extensively in [3, 9, 11, 19] (see also [18]). It is shown in these papers that

- the maximum combinatorial complexity (i.e., number of edges and vertices) of an arrangement of n segments is $\Theta(n^2)$;
- the maximum complexity of a single face of such an arrangement is $\Theta(n\alpha(n))$ [11, 19]; and
- the maximum complexity of any m distinct faces in such an arrangement is $O(m^{2/3}n^{2/3} + n\alpha(n) + n\log m)$ and can be $\Omega(m^{2/3}n^{2/3} + n\alpha(n))$ [3, 9].

Let $\mathcal{P}$ be a set of k convex polygons with a total of n edges. We denote the common exterior of $\mathcal{P}$ by $\mathcal{E} = \mathcal{E}(\mathcal{P})$. If k is proportional to n or, equivalently, if the average size of

*Work by Boris Aronov has been supported by NSF Grant CCR-92-11541 and a Sloan Research Fellowship. Work by Micha Sharir has been supported by NSF Grants CCR-91-22103, CCR-93-11127, and CCR-94-24398, by a Max-Planck Research Award, and by grants from the U.S.-Israeli Binational Science Foundation, the Israel Science Fund administered by the Israeli Academy of Sciences, and the G.I.F., the German-Israeli Foundation for Scientific Research and Development.

†Department of Computer Science, Polytechnic University, Brooklyn, NY 11201 USA.

‡School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 USA.

[1] Here and thereafter $\alpha(n)$ denotes the functional inverse of the Ackermann's function.

a polygon of $\mathcal{P}$ is bounded by a constant, then such an arrangement of k convex polygons is not much different from an arrangement of n segments, and no bounds better than those just cited can be obtained. However, the situation changes drastically if the average size of a polygon of $\mathcal{P}$ is increased, that is, if $k \ll n$. For instance, the maximum complexity of the entire arrangement, which can be as high as $\Theta(n^2)$ in an arrangement of n arbitrary segments, now reduces to $\Theta(kn)$—this follows trivially from the observation that an edge of one polygon of $\mathcal{P}$ can intersect the boundary of another convex polygon in at most two points, and the lower bound is equally trivial to establish. Below, the maximum total complexity of $\mathcal{E}(\mathcal{P})$ is shown to be $\Theta(k^2 + n\alpha(k))$, which is asymptotically much smaller than $\Theta(kn)$ (see Theorem 2.1). The maximum complexity of a single face of $\mathcal{E}(\mathcal{P})$ is shown to be $\Theta(n\alpha(k))$, slightly improving the general upper bound cited above. In fact, this bound applies to any face of the arrangement of $\mathcal{P}$. Finally, we analyze the complexity of any m distinct faces of $\mathcal{E}(\mathcal{P})$, and establish an upper bound of $O(m^{2/3}k^{2/3}\log^{1/3}\frac{k^2}{m} + n\log k)$, improving considerably the general upper bound cited above.

Intuitively, these bounds indicate that the complexity measures under consideration (of the common exterior, of any single face of the exterior, and of several faces of the exterior) depend mainly on the number of polygons k, with the parameter n contributing only an almost-linear additive term to these bounds. In fact, these additive terms are all linear in n, with small multiplicative factors, which are at most logarithmic in k and independent of n.

This paper, whose original preparation started about 1989, has found applications in [2]. Another motivation for studying the problem comes from the study in [15] of *translational motion planning* in the plane. It is shown there that if $\mathcal{P}$ is a collection of *Minkowski sums* of the form $A_i \oplus B$, where $A_1, \ldots, A_k$ and B are convex polygons, and the A_i's have pairwise-disjoint interiors, then the complexity of the common exterior $\mathcal{E}(\mathcal{P})$ is $O(n)$, where n is the total number of edges of the polygons of $\mathcal{P}$. In fact, the result of [15] is more general. It states that if the polygons of $\mathcal{P}$ have the property that the boundaries of any pair of them intersect in at most two points (a property that holds in the above case of Minkowski sums), then the complexity of their common exterior is $O(n)$. However, if pairs of polygon boundaries can intersect in four (or more) points, the complexity of the common exterior can increase to $\Omega(k^2)$, as is easily seen. The results of this paper give a tight calibration of the complexity of $\mathcal{E}(\mathcal{P})$ (and a fairly accurate calibration of appropriate portions of $\mathcal{E}(\mathcal{P})$) in this more general case. We also mention a recent extension of the results of this paper to three dimensions: we have shown in [4, 5] that the complexity of the union of k convex polyhedra in 3-space, with a total of n faces, is $O(k^3 + kn\log k)$, and can be $\Omega(k^3 + kn\alpha(k))$ in the worst case.

In the remainder of the paper we assume that the polygons of $\mathcal{P}$ are in *general position*, i.e., no three polygon boundaries meet at a common point and no polygon vertex lies on the boundary of another polygon. It is easy to show that appropriate slight expansions of the given polygons put them in general position, without decreasing the complexity measures under consideration. Thus there is no loss of generality in assuming general position.

2 The Complexity of the Common Exterior

Theorem 2.1 *The maximum number of edges bounding the common exterior of k convex polygons with a total of n edges is $\Theta(n\alpha(k) + k^2)$.*

Proof: We first give a brief high-level overview of the proof. We form the arrangement $\mathcal{A}$ of k 'sentinel lines', each passing through the leftmost and rightmost points of some polygon, and of the $2k$ vertical lines passing through these leftmost and rightmost points. We next decompose the plane into $O(k^2)$ 'boxes', by splitting each face of $\mathcal{A}$ further by vertical lines passing through the vertices on its lower boundary. We then show that the analysis of the complexity of the common exterior can be reduced to the analysis of the complexity of the lower envelopes, one in each box, of the polygon boundaries that cross the box. Within a particular box C, each polygon boundary edge that crosses C is either 'short' (terminates within C) or 'long' (crosses the boundary of C twice.) We show that the overall number of long edges, summed over all boxes, is $O(k^2)$, and the overall number of short edges is $O(n)$. Using standard results on lower envelopes, the upper bound follows readily. The second term in the lower bound is based on the simple observation that any lower bound that holds for faces in an arrangement of k line segments can be made to hold for the common exterior of k convex polygons, by replacing each segment by a sufficiently thin rectangle. The first term requires a more complex construction, built upon the $\Omega(k\alpha(k))$ lower bound construction for lower envelopes of k segments, as given in [19].

We now present the proof in full detail: Let $\mathcal{P} = \{P_1, \ldots, P_k\}$ be a collection of k convex polygons in the plane. Let n_i denote the number of vertices of P_i, for $i = 1, \ldots, k$, so that $n = \sum_{i=1}^{k} n_i$. In what follows, we will denote the common exterior $\mathcal{E}(\mathcal{P})$ simply as $\mathcal{E}$.

For each i, define the *sentinel segment* s_i of P_i to be the segment connecting the leftmost vertex of P_i to its rightmost vertex. (Without loss of generality, we can assume that each P_i has a unique leftmost vertex and a unique rightmost vertex.) The *sentinel line* ℓ_i of P_i is the line containing s_i. Put $S = \{s_1, \ldots, s_k\}$ and $\mathcal{L} = \{\ell_1, \ldots, \ell_k\}$. Draw a vertical line through each endpoint of each segment in S, and let $\mathcal{L}^*$ denote the union of $\mathcal{L}$ with the set of these $2k$ vertical lines.

Consider the arrangement $\mathcal{A}(\mathcal{L}^*)$ of $\mathcal{L}^*$. Its combinatorial complexity is $O(k^2)$. Moreover, the number of intersections between the boundaries of the P_i's and the lines of $\mathcal{L}^*$ is also $O(k^2)$, since a line can intersect the boundary of a convex polygon in at most two points.

Let C be one of the (necessarily convex) cells of $\mathcal{A}(\mathcal{L}^*)$, and let P_i be a polygon in $\mathcal{P}$. If the x-projection of s_i is disjoint from that of C, then clearly P_i cannot contribute any feature to $\partial\mathcal{E} \cap C$. Otherwise, s_i can be classified as lying either (possibly, partially) above C or (partially) below C. Let $\mathcal{P}_C^+$ (resp. $\mathcal{P}_C^-$) denote the subcollection of all the P_i's for which s_i lies above (resp. below) C.

It is easily seen that $\mathcal{E} \cap C$ is the set of all points within C that lie between the lower envelope $E_{\mathcal{P}_C^+}$ of the polygons in $\mathcal{P}_C^+$ and the upper envelope $E_{\mathcal{P}_C^-}$ of the polygons in $\mathcal{P}_C^-$. Hence, using a standard argument, the complexity of $\mathcal{E} \cap C$ is proportional to the sum of the complexities of these two envelopes within C.

We estimate separately the complexity of the lower envelope portions $C \cap E_{\mathcal{P}_C^+}$ and of the upper envelope portions $C \cap E_{\mathcal{P}_C^-}$, over all $C \in \mathcal{A}(\mathcal{L}^*)$. It suffices to consider the case of lower envelopes, since the case of upper envelopes is fully symmetric.

We subdivide each face C of $\mathcal{A}(\mathcal{L}^*)$ into subfaces, by drawing a maximal vertical segment within C from each vertex on the lower boundary of C. Note that each of the left, right, and bottom sides of each subface τ consists of a single segment (the left and right sides may degenerate to a single point, and, in case of unbounded subfaces, some of the sides may be absent altogether), whereas the top boundary of τ is a concave polygonal chain; see Figure 1. For each resulting subface τ, let $\mathcal{P}_\tau^+$ denote the set of polygons in $\mathcal{P}_C^+$ whose lower boundaries intersect τ, where C is the face of $\mathcal{A}(\mathcal{L}^*)$ containing τ; clearly, only the lower boundaries of these polygons can intersect τ. (If some polygon P_i covers τ completely, it does not contribute

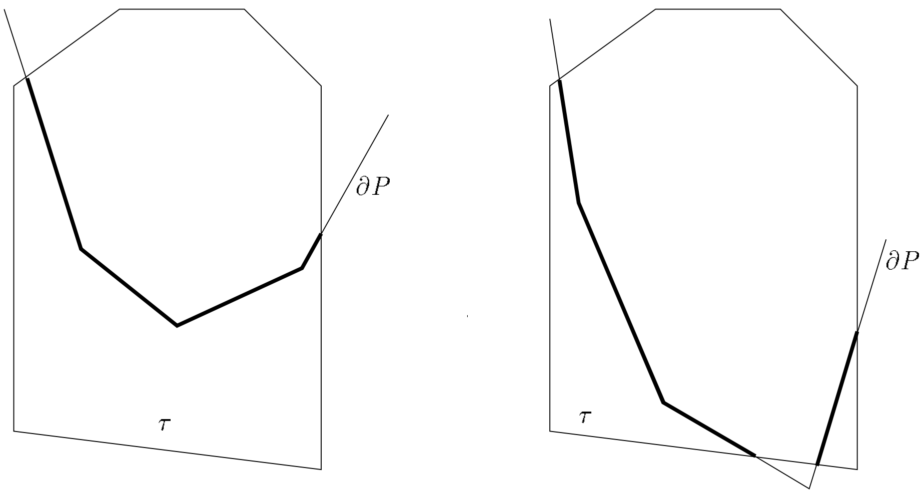

Figure 1: Possible interaction of a subface τ of $\mathcal{A}(\mathcal{L}^*)$ and a polygon P in $\mathcal{P}_\tau^+$. Thick lines indicate $\partial P \cap \tau$.

to $\mathcal{E}(\mathcal{P})$ and τ need not be considered at all, so we will assume hereafter that no polygon of $\mathcal{P}_C^+$ that meets τ covers it completely.)

Let B_τ^+ denote the set of segments that constitute the intersections $\partial P_i \cap \tau$, for $P_i \in \mathcal{P}_\tau^+$. Let S_τ^+ denote the set of those segments $e \in B_\tau^+$ such that at least one endpoint of e is an original vertex of the polygon on whose boundary e lies, and let R_τ^+ denote the set of all other segments in B_τ^+ (both endpoints of each of these segments lie on $\partial\tau$). Clearly, $\sum_\tau |S_\tau^+| \le 2n$.

We claim that $\sum_\tau |R_\tau^+| = O(k^2)$. To see this, fix a polygon $P_i \in \mathcal{P}$, and consider the *outer zone* of P_i in $\mathcal{A}(\mathcal{L}^*)$. This is defined as the collection of those portions of the faces of $\mathcal{A}(\mathcal{L}^*)$ crossed by the boundary of P_i, which lie outside P_i. If we erase, for each line $\ell \in \mathcal{L}^*$, the segment $\ell \cap P_i$, we obtain a collection $\mathcal{L}'$ of at most $6k$ lines and rays, and the outer zone is a portion of a single face of $\mathcal{A}(\mathcal{L}')$. As is well known [1], the complexity of such a face is $O(k)$.

To complete the estimate of $\sum_\tau |R_\tau^+|$, we charge each segment $e \in R_\tau^+$, for any subface τ, as follows. If e has an endpoint that lies on a line of $\mathcal{L}^*$, we charge e to that endpoint; the overall number of such points, as noted above, is $O(k^2)$ and each will be charged at most twice. Otherwise, the right endpoint of e lies on a vertical side g of τ, which is a newly added vertical segment, erected upwards from some vertex v of $\mathcal{A}(\mathcal{L}^*)$. In this case, we charge e to v, and observe that v is a vertex of the outer zone of the polygon P_i on whose boundary e lies, and that v can be charged at most once by segments e lying on the lower boundary of the same P_i. It follows that the number of such segments e, over all subfaces τ, that lie on the boundary of any fixed P_i is $O(k)$, so the overall number of such segments is $O(k^2)$.

Consider next the lower envelope $E_{\mathcal{P}_\tau^+}$, within some subface τ. Any vertex of $E_{\mathcal{P}_\tau^+}$ that lies in the interior of τ must also be a vertex of the lower envelope $E_{B_\tau^+}$ of the segments in B_τ^+ (note that the converse statement may fail to hold). Any other vertex of $\tau \cap E_{\mathcal{P}_\tau^+}$ lies on $\partial\tau$ and is an endpoint of a segment in B_τ^+, which also appears as a vertex of $\tau \cap E_{B_\tau^+}$. We consider separately the lower envelopes $E_{S_\tau^+}$ and $E_{R_\tau^+}$ of the subcollections S_τ^+ and R_τ^+, respectively, and observe that $E_{B_\tau^+}$ is the pointwise minimum of $E_{S_\tau^+}$ and $E_{R_\tau^+}$, so its complexity is proportional to the sum of the complexities of these two 'sub-envelopes'.

The complexity of $E_{S_\tau^+}$ is $O(|S_\tau^+|\alpha(k))$. This follows from the results of [13], exploiting the fact that any vertical line intersects at most k segments of S_τ^+. The complexity of $E_{R_\tau^+}$ is $O(|R_\tau^+|)$. To see this, we take each segment $e \in R_\tau^+$ and replace it by a line or ray, as follows. If e does not intersect the bottom side of τ, we replace e by the full line containing e. If e intersects the bottom side of τ at some (unique) point u, we replace e by the ray emanating from u and containing e. Let R_τ^* denote the resulting collection of lines and rays. It is easy to verify that the lower envelope $E_{R_\tau^*}$ of R_τ^* coincides, within τ, with the envelope $E_{R_\tau^+}$ (the extended portions of the lines and rays in R_τ^* all lie outside τ, and none of them passes below τ). Using, once again, the results of [1], we conclude that the complexity of $E_{R_\tau^*}$, and thus also that of $E_{R_\tau^+}$, is $O(|R_\tau^+|)$.

We have thus shown that the sum of the complexities of the envelope portions $\tau \cap E_{\mathcal{P}_\tau^+}$, over all subfaces τ, is proportional to

$$\sum_\tau O(|S_\tau^+|\alpha(k) + |R_\tau^+|) = O(n\alpha(k) + k^2)\,.$$

Applying a symmetric analysis to the upper envelopes $E_{\mathcal{P}_\tau^-}$, and combining the bounds, we obtain the upper bound asserted in the theorem.

The lower bound is established by two constructions. The first construction yields a collection of k convex polygons with a total of n edges, whose common exterior $\mathcal{E}$ consists of a single face of complexity $\Omega(n\alpha(k))$. This is achieved as follows. Construct a collection $S = \{s_1, \ldots, s_k\}$ of k segments, whose upper envelope has complexity $\Omega(k\alpha(k))$ [19]. As is easily verified, one can extend each segment in S to a total continuous piecewise-linear function, by adding a steeply ascending (resp. descending) half-line immediately to the left (resp. to the right) of the segment, so that, if the absolute value of the slopes of all these half-lines is greater than some threshold value, then the combinatorial complexity of the upper envelope of these functions is at least as large as that of the upper envelope of S.

Let $q = \lfloor \frac{n}{2k} \rfloor$, and fix $0 < \epsilon < \frac{\pi}{q}$. Without loss of generality, we may assume that $q \geq 3$. Scale S so that it fits into a unit disc, and shrink it further vertically, so that any line forming an angle of less than $\frac{\pi}{2} - \frac{\pi}{q} + \epsilon$ with the vertical direction is "sufficiently steep" in the above sense (in particular, the slope of such a line has larger absolute value than the slope of any segment in S). Let Q be a regular q-gon whose side has length $\frac{10}{\epsilon}$. Place a rotated copy of S at each vertex v of Q, aligning its vertical direction with the radial direction of the ray from the center of Q through v. Denote this rotated copy by $S^{(v)} = \{s_1^{(v)}, \ldots, s_k^{(v)}\}$, where $s_i^{(v)}$ is the image of s_i in S_v, for each $i = 1, \ldots, k$ and each vertex v of Q. Construct k $2q$-gons, $Q_1, \ldots, Q_k$, where Q_i, for $i = 1, \ldots, k$, is the convex hull of $\bigcup_v s_i^{(v)}$, where the union is taken over all vertices v of Q. The boundary of each Q_i consists of a sequence of edges, alternating between copies of s_i and long 'connecting' edges between endpoints of pairs of successive copies of s_i. This follows from the way in which S was shrunk, as argued in more detail in the following paragraph; see also Figure 2.

Let pq be an edge of one of these polygons Q_i, so that p is an endpoint of $s_i^{(v)}$ and q is an endpoint of $s_i^{(w)}$, for a pair of adjacent vertices v, w of Q. By construction, p (resp. q) is in a unit disc centered at v (resp. w), and v and w are $\frac{10}{\epsilon}$ apart. In particular, this implies that the angle between vw and pq cannot exceed ϵ. As a result, in any single copy of S, the added connecting segments lie on lines forming an angle at most $\frac{\pi}{2} - \frac{\pi}{q} + \epsilon$ with the local "vertical" direction, and thus do not decrease the combinatorial complexity of the upper envelope of each copy of S (notice, by the way, that this also implies that the polygons Q_i do indeed

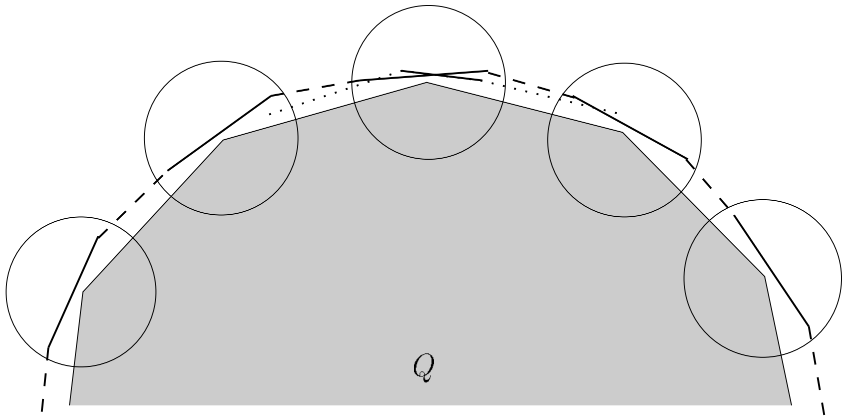

Figure 2: The lower bound construction in Theorem 2.1, not to scale. One of the polygons Q_i is shown dashed and part of the boundary of another polygon Q_j is shown dotted.

have the alternating structure claimed above). We have thus obtained a collection $\mathcal{P}$ of k convex polygons with a total of $2qk \leq n$ edges, where the unbounded face (which is the only component of the common exterior $\mathcal{E}$) has complexity $q \cdot \Omega(k\alpha(k)) = \Omega(n\alpha(k))$, as desired.

The second term, $\Omega(k^2)$, in the asserted lower bound is trivial to obtain, e.g., by a collection of k long and thin rectangles, half of which have their long edge horizontal, and the other half have their long edge vertical. Combining these two constructions, the asserted lower bound follows.

This completes the proof of the theorem. $\Box$

3 Multiple Exterior Faces

The goal of this section is to establish sharp bounds on the overall complexity of m distinct faces of the common exterior of a collection of convex polygons. We begin by stating the standard 'combination lemma' of [9]:

Lemma 3.1 *Let $\mathcal{A}_1, \mathcal{A}_2$ be two arrangements of a total of n segments in the plane, and let $\mathcal{A}$ denote the arrangement obtained by superimposing $\mathcal{A}_1$ and $\mathcal{A}_2$. Let M be a set of m points, none lying on any segment. Then the overall complexity of the faces of $\mathcal{A}$ that contain points of M is at most $C_1 + C_2 + O(m+n)$, where C_i is the total complexity of the faces of $\mathcal{A}_i$ that contain points of M, for $i = 1, 2$.*

The main tool that we use is the so-called 'multi-color combination lemma', which deals with the complexity of several faces in an overlay of many arrangements. There are two known variants of this lemma:

Lemma 3.2 (Combination Lemma I) *Let $\mathcal{A}_1, \ldots, \mathcal{A}_t$ be t arrangements of a total of n segments in the plane, and let M be a set of points, none of which lies on any segment. Let $\mathcal{T}$ be a binary tree of height $O(\log t)$, whose leaves correspond to the individual arrangements $\mathcal{A}_i$, and let us associate with each internal node ν of $\mathcal{T}$ the arrangement $\mathcal{A}_\nu$ obtained by the* *superposition of all the arrangements corresponding to the leaves of the subtree rooted at ν; the root is thus associated with the superposition $\mathcal{A}$ of all the given arrangements. For any node ν of $\mathcal{T}$, let $\mathcal{F}_\nu$ denote the collection of faces of $\mathcal{A}_\nu$ that contain points of M, and let $\mathcal{F}$ denote this collection at the root of $\mathcal{T}$. Let C_i denote the total combinatorial complexity of the faces of the original arrangement $\mathcal{A}_i$ containing points of M, for $i = 1, \ldots, t$, and let $C = \sum_i C_i$. Then the total combinatorial complexity of the faces of $\mathcal{F}$ is at most*

$$C + O(n \log t) + O\left(\sum_{\nu \text{ a non-leaf}} |\mathcal{F}_\nu| \right).$$

Lemma 3.2 is easily proved by applying Lemma 3.1 to the nodes of $\mathcal{T}$ in a bottom-up fashion; see, for example, the proof of Lemma 2.3 in [3].

Lemma 3.3 (Combination Lemma II [12]) *Let $\mathcal{A}_1, \ldots, \mathcal{A}_t$ be t arrangements of segments in the plane, and let p be a point not lying on any of these segments. Let F_i denote the face of $\mathcal{A}_i$ containing p, for $i = 1, \ldots, t$, and let F denote the face containing p in the arrangement $\mathcal{A}$ obtained by the superposition of all the arrangements $\mathcal{A}_i$. Let C_i denote the combinatorial complexity of F_i, for $i = 1, \ldots, t$, and let $C = \sum_i C_i$. Then the combinatorial complexity of F is $O(C\alpha(t))$.*

We begin this section with a tight bound on the maximum complexity of a single face of the common exterior of a collection of convex polygons. As a matter of fact, our upper bound holds for *any* face of the arrangement of such polygons:

Theorem 3.4 *The maximum number of edges bounding any single face of the arrangement of k convex polygons with a total of n edges is $\Theta(n\alpha(k))$. The same bound holds for any face of the arrangement.*

Proof: The lower bound is an immediate consequence of the first lower bound construction given in the proof of Theorem 2.1. (Enclosing the entire construction in a large triangle formed by three elongated rectangles produces an arrangement formed by the boundaries of k convex polygons with a total of n edges, in which there is a non-exterior face with complexity $\Omega(n\alpha(k))$.) The upper bound is an easy consequence of Combination Lemma II: Let $\mathcal{P} = \{P_1, \ldots, P_k\}$ be a collection of k convex polygons with a total of n edges, and let F be a face of the arrangement of $\mathcal{P}$. Let $\mathcal{A}_i$, for $i = 1, \ldots, k$, be the arrangement formed by the edges of P_i. Then F_i (the face of $\mathcal{A}_i$ containing F) is either the exterior or the interior of P_i, so, in the notation of Lemma 3.3, $C_i = O(n_i)$, where n_i is the number of edges of P_i, and $C = O(n)$. The asserted upper bound is now an immediate consequence of Lemma 3.3. □

We next analyze the complexity of many faces in the common exterior of a collection of convex polygons, as above. First of all, the number of components of the common exterior is only $O(k^2)$, as easily follows from the results of Katona [14] and Kovalev [16]:

Theorem 3.5 *The common exterior of any collection of k compact convex bodies in the plane has at most $\binom{k-1}{2} + 1 = O(k^2)$ components.*

Before proving the main theorem of this section, we introduce the following technical tool. Let $\mathcal{L}$ be a collection of m lines in the plane, and let $1 \leq r \leq m$ be an integer. As is well known [7, 17], there exists a tiling of the plane by $O(r^2)$ triangles, so that the closure of no triangle meets more than $\frac{cm}{r}$ lines of $\mathcal{L}$, where c is some absolute constant. Moreover, one can construct a binary tree $\mathcal{T}$ of depth $O(\log r)$, whose leaves correspond to those triangles, such that, for every node ν of $\mathcal{T}$, any line of $\mathcal{L}$ that misses all triangles stored at the leaves of the subtree $\mathcal{T}_\nu$ rooted at ν, lies either above all these triangles or below all of them (see [3, Lemma 2.2]); in other words, such a line cannot separate the triangles of T_ν. We refer to such a tiling as an $(\mathcal{L}, r)$-*tiling*.

Theorem 3.6 *The maximum number $C(m,k,n)$ of edges bounding any $m \le \binom{k-1}{2}+1$ distinct faces of the common exterior of k convex polygons with a total of n edges is*

$$O(m^{2/3}k^{2/3}\log^{1/3}\left(\frac{k^2}{m}\right) + n\log k) \quad \textit{and} \quad \Omega(m^{2/3}k^{2/3} + n\alpha(k)) .$$

Proof: Let $\mathcal{P}$ be such a collection of polygons, and let M be a set of m 'marking' points, one point in each of m given faces of $\mathcal{E}(\mathcal{P})$.

(a) We first prove the weaker bound $C(m,k,n) = O(k\sqrt{m} + n\log k)$. Since $m \le \binom{k-1}{2} + 1$, we can partition the polygons into $t = \left\lceil \frac{k}{\sqrt{m}} \right\rceil \ge 2$ groups, each group containing $O(\sqrt{m})$ polygons. Let $\mathcal{P}_j$ denote the set of polygons in the j-th group, and apply Theorem 2.1 to each arrangement $\mathcal{A}(\mathcal{P}_j)$, to conclude that the total complexity of the faces of $\mathcal{A}(\mathcal{P}_j)$ containing points of M (all these faces belong to the common exterior of $\mathcal{P}_j$) is $O(m + N_j\alpha(m))$, where N_j is the total number of edges of the polygons in $\mathcal{P}_j$. We apply Combination Lemma I to the t arrangements $\mathcal{A}(\mathcal{P}_j)$, and use the trivial estimate $\sum_{\nu \text{ a non-leaf}} |\mathcal{F}_\nu| = O(mt)$. Hence we obtain

$$\begin{aligned} C(m,k,n) &= \sum_{j=1}^{t} O(m + N_j\alpha(m)) + O(mt + n\log t) \\ &= O(mt + n\alpha(m) + n\log t) = O(k\sqrt{m} + n\log k), \end{aligned}$$

as asserted.

(b) We next establish the sharper upper bound asserted in the theorem. Our proof is similar to an argument used in [3] for arrangements of segments. Let $\mathcal{P}$ and M be as above. Using the terminology introduced in Section 2, let L be the set of the k non-vertical sentinel lines of the polygons in $\mathcal{P}$.

Since we have considerable freedom in choosing the exact position of the points in the marking set M, we may assume that they lie in general position, i.e., that no three points lie on a common line. We consider the dual transformation that maps any point $q = (a,b)$ to the non-vertical line $q^* : y = -ax + b$, and maps any non-vertical line $\ell : y = cx + d$ to the point $\ell^* = (c,d)$; this transformation preserves the above-below relation between points and lines [8]. Pass to the dual plane, and consider the set M^* of lines dual to the marking points, and the set L^* of points dual to the sentinel lines. Let r be a fixed integer between 1 and m, to be determined below. We construct an (M^*, r)-tiling of the dual plane, as defined just before the theorem; let $\mathcal{T}$ be the associated binary tree built upon the triangles of the tiling.

Consider a triangle τ, corresponding to some leaf of $\mathcal{T}$, that we also denote by τ. Let $\mathcal{P}_\tau$ be the set of polygons P_i whose dual points ℓ_i^* lie in τ, let N_τ denote the total number of edges of the polygons in $\mathcal{P}_\tau$, and let M_τ be the set of marking points whose dual lines intersect the closure of τ. As above, there is no loss of generality in assuming that each point dual to a sentinel line lies in the interior of some triangle. Therefore the sets $\mathcal{P}_\tau$, as above, form a partition of $\mathcal{P}$, so that $\sum_\tau |\mathcal{P}_\tau| = k$ and $\sum_\tau N_\tau = n$. We now apply Combination Lemma I to the family of arrangements $\mathcal{A}(\mathcal{P}_\tau)$ and to the marking set M, to deduce that the total complexity of the marked faces of $\mathcal{E}(\mathcal{P})$ is at most

$$\sum_{\tau \text{ a leaf}} C_\tau + O\left(\sum_{\nu \text{ a non-leaf}} |\mathcal{F}_\nu|\right) + O(n \log r) \ , \qquad (1)$$

where C_τ is the complexity of the marked faces in $\mathcal{A}(\mathcal{P}_\tau)$, and $\mathcal{F}_\nu$ is the collection of marked faces in the overlay of the arrangements $\mathcal{A}(\mathcal{P}_\tau)$, over all leaves τ of the subtree $\mathcal{T}_\nu$ of T rooted at ν.

The first step in simplifying this expression is to bound the quantities $|\mathcal{F}_\nu|$. Let τ be a triangle in the above partitioning of the dual plane. If a marking point p does not belong to M_τ then τ is fully contained in one of the halfplanes, say in the upper halfplane, bounded by the dual line p^*. In the primal plane, this means that p lies below all the sentinel lines of the polygons in $\mathcal{P}_\tau$. This is easily seen to imply that p lies in the unbounded face of the common exterior $\mathcal{E}(\mathcal{P}_\tau)$. In other words, we have just argued that the number $|\mathcal{F}_\tau|$ of marked faces in $\mathcal{A}(\mathcal{P}_\tau)$ is at most $\frac{cm}{r}+1$. The same argument, using the non-separability property of subtrees of $\mathcal{T}$ stated above, implies that, for any non-leaf node ν of $\mathcal{T}$, we have $|\mathcal{F}_\nu| \leq \frac{cm}{r} \cdot |\nu| + 1$, where $|\nu|$ is the number of leaves of $\mathcal{T}_\nu$. Indeed, if p is any marking point p whose dual line p^* misses all the triangles τ stored at the leaves of $\mathcal{T}_\nu$, then p^* passes either above all these triangles or below all of them. In both cases, the point p must lie in the unbounded face of the overlay of the arrangements $\mathcal{A}(\mathcal{P}_\tau)$, over all leaves τ of the subtree $\mathcal{T}_\nu$. Since the number of marking points whose dual lines do intersect one of these triangles τ is at most $\frac{cm}{r} \cdot |\nu|$, the claim follows.

Using the weaker bound derived in step (a) above, we have

$$C_\tau = O(|\mathcal{P}_\tau|\sqrt{\frac{cm}{r}+1} + N_\tau \log k) \ .$$

Thus, the bound in (1) is

$$\sum_{\tau \text{ a leaf}} O(|\mathcal{P}_\tau|\sqrt{\frac{cm}{r}+1} + N_\tau \log k) + O\left(\sum_{\nu \text{ a non-leaf}} (\frac{cm}{r} \cdot |\nu| + 1)\right) + O(n \log r) \ .$$

The first sum in this expression is $O(k\sqrt{\frac{cm}{r}+1} + n \log k)$. To bound the second sum, note that $\sum_\nu |\nu|$, over all nodes ν at a fixed level of $\mathcal{T}$, is $O(r^2)$, so the sum $\sum_\nu |\nu|$, over all nodes of $\mathcal{T}$, is $O(r^2 \log r)$. To summarize,

$$C(m,k,n) = O(k\sqrt{\frac{m}{r}} + n \log k + mr \log r). \qquad (2)$$

Now choose

$$r = \left\lceil \frac{k^{2/3}}{m^{1/3} \log^{2/3}(k^2/m)} \right\rceil .$$

It is easily checked that if $m \geq \sqrt{k}$ then $1 \leq r \leq m$ (the inequality $r \geq 1$ follows from the fact that $\left\lceil \frac{k}{\sqrt{m}} \right\rceil \geq 2$). Thus, if $m \geq \sqrt{k}$, this choice of r implies

$$C(m,k,n) = O(k^{2/3} m^{2/3} \log^{1/3}\left(\frac{k^2}{m}\right) + n \log k).$$

If $m = O(\sqrt{k})$, it is sufficient to show that $C(m, k, n) = O(n \log k)$, which can be obtained by putting $r = m$ in (2).

The first term in the lower bound of the theorem follows from the fact that there exists an arrangement of k segments that contains m faces whose overall complexity is $\Omega(m^{2/3}k^{2/3})$ [6, 10], and from the observation, already made above, that in this construction the segments can be replaced by appropriate elongated rectangles. The second term in the lower bound follows from the lower bound construction of Theorem 2.1. □

Note: The estimate obtained in part (a) of the above upper bound proof can be written in the following slightly sharpened form:

$$C(m, k, n) = O(k\sqrt{m} + n\alpha(m) + n \log \frac{k^2}{m}).$$

Using this estimate in the remainder of the upper bound proof produces the following expression:

$$C(m, k, n) = O\left(k^{2/3}m^{2/3}\log^{1/3}\left(\frac{k^2}{m}\right) + n\alpha\left(\left\lceil\frac{m^2}{k}\right\rceil\right) + n \log \frac{k^2}{m}\right),$$

which is microscopically better than the bound quoted in the theorem, for m close to $\Theta(k^2)$. It gives $O(k^2 + n\alpha(k))$ for the largest possible value of m, matching the lower bound. However, a tight bound over the entire range of values of m, k, and n, remains an open problem.

References

[1] P. Alevizos, J. D. Boissonnat, and F. P. Preparata, An optimal algorithm for the boundary of a cell in a union of rays, *Algorithmica* 5 (1990), 573–590.

[2] B. Aronov, M. Bern, and D. Eppstein, Arrangements of polytopes with applications, manuscript, 1992.

[3] B. Aronov, H. Edelsbrunner, L. Guibas and M. Sharir, Improved bounds on the number of edges of many faces in arrangements of line segments, *Combinatorica* 12 (1992), 261–274.

[4] B. Aronov and M. Sharir, The union of convex polyhedra in three dimensions, *Proc. 34th IEEE Symp. on Foundations of Computer Science* (1993), pp. 518–527.

[5] B. Aronov, M. Sharir, and B. Tagansky, The union of convex polyhedra in three dimensions, *SIAM J. Comput.*, to appear.

[6] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir, and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* 5 (1990), 99–160.

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

[8] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.

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

[10] H. Edelsbrunner and E. Welzl, On the maximal number of edges of many faces in arrangements, *J. Comb. Theory, Ser. A* 41 (1986), 159–166.

[11] L. Guibas, M. Sharir, and S. Sifrony, On the general motion planning problem with two degrees of freedom, *Discrete Comput. Geom.* 4 (1989), 491–521.

[12] S. Har Peled, *The complexity of many faces in the overlay of arrangements,* M.Sc. Thesis, Department of Computer Sciences, Tel Aviv University, 1995.

[13] D. P. Huttenlocher, K. Kedem, and J. Kleinberg, Voronoi diagrams of rigidly moving sets of points, *Inform. Process. Lett.*, 43 (1992), 217–223.

[14] G. O. Katona, On a problem of L. Fejes Tóth, *Stud. Sci. Math. Hung.* 12 (1977), 77–80.

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

[16] М. Д. Ковалев, Свойство выпуклых множеств и его приложение, *Математические Заметки*, **44** (1988), 89–99.

[17] J. Matoušek, Construction of ε-nets, *Discrete Comput. Geom.* 5 (1990), 427–448.

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

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