On the external boundary of a union of rays

Jean-Daniel Boissonnat, Franco Preparata

UNITÉ DE RECHERCHE
INRIA-SOPHIA ANTIPOLIS

Institut National
de Recherche
en Informatique
et en Automatique

Domaine de Voluceau
Rocquencourt
B.P.105
78153 Le Chesnay Cedex
France
Tél.:(1) 39 63 55 11

Rapports de Recherche

N° 705

ON THE EXTERNAL BOUNDARY OF A UNION OF RAYS*

Jean-Daniel BOISSONNAT
Franco P. PREPARATA

JUILLET 1987

ON THE EXTERNAL BOUNDARY OF A UNION OF RAYS*

Jean-Daniel Boissonnat[†] and Franco P. Preparata[‡]

Abstract

In this paper, we consider the external contour of a union of n rays, that is, the boundary of the unbounded region in the complement of this union. We show that the external contour of union of rays has $O(n)$ edges (differently from the contour of a union of segments). We also show that, if all ray termini are known to belong to the contour, the contour can be computed in optimal time $\Theta(n \log n)$.

*This work was partly supported by the CEE ESPRIT Project P-940, by the Ecole Normale Supérieure, Paris, France, and by the NSF grant ECS-84-10902.

†INRIA Avenue Emile Hugues, 06565 Valbonne, France

‡University of Illinois at Urbana-Champain, IL 61801, USA. This work was done while this author was on sabbatical leave at the Ecole Normale Supérieure, Paris, France.

1 Introduction

Given is a set $R = \{r_1 ..., r_n\}$ of n rays in the plane (i.e. half lines, with the property that all intersect a line l and that all ray termini lie on the same side of l). To avoid insignificant degeneracies, we assume that the rays are in general position, i.e., all termini are distinct, no three rays intersect in the same finite point, and no two rays intersect on l. The union $F = \bigcup_{j=1}^{n} r_j$ partitions the plane into a collection of regions, some of them being unbounded. We call external boundary E of F, the boundary of the unbounded region passing through the ray terminus which is the most distant from l. Without loss of generality, we will assume, throughout this paper, that l is the x axis and that the ray termini have positive ordinates.

We show, in this paper, that, while the boundary of F may consist of $O(n^2)$ edges, its external boundary consists of, at most, $O(n)$ edges. This problem is closely related to the problem of finding the upper envelope (i.e. pointwise maximum) of n functions. This problem was first studied by Atallah [1] and then by Sharir, Hart and Wiernick [4,8].

For the case of n line-segments in the plane, none of which is vertical, Hart and Sharir have shown that the upper envelope consists of, at most, $O(n\alpha(n))$ edges where $\alpha(n)$ is the extremely slowly growing inverse Ackermann's function. Their proof extends immediately to the case of rays. Very recently, Wiernick and Sharir [8] proved that this bound is tight for line-segments.

However, as a consequence of our result here, the bound is not tight for sufficiently long line-segments and, in particular, for rays.

Moreover, we present an algorithm that computes the external boundary of a set of n rays whose termini are on the contour in optimal $O(n \log n)$ time. Our algorithm makes use of a result by Alevizos, Boissonnat, and Yvinec [2] who showed that the order of the ray termini along the external boundary can be obtained in time $O(n \log n)$. This order is based on the following binary relation: let x_i and x_j be the abscissae of the intersections of r_i and r_j with l, respectively. Ray r_i precedes ray r_j if either r_i and r_j do not intersect in the half-plane $y > 0$ and $x_i < x_j$ or r_i and r_j intersect in the half-plane $y > 0$ and $x_i > x_j$. It is shown in [2] that this relation is a total order under our hypotheses (referred to as the ABY-order). In the following, the attributes "consecutive", "adjacent", etc..., as applied to rays, refer to their ABY-order.

Throughout the paper, we will make the two following assumptions :

1. The external boundary E is connected,
2. E passes through all ray termini.

Assumption 1 is trivially not restrictive. In Section 5, we will justify that Assumption 2 is not restrictive for the bound, while the algorithm explicitly

A PROPOS DU CONTOUR EXTERIEUR D'UNE UNION DE RAYONS

Jean-Daniel Boissonnat et Franco P. Preparata

Juin 1987

Résumé

Soient n demi-droites dans le plan (appelées *rayons*) qui coupent toutes une droite donnée l et dont toutes les extrémités sont du même côté de l. On montre que le contour extérieur de l'union de ces rayons, c'est à dire le bord de la région non-bornée du complémentaire de l'union des rayons, a $O(n)$ arêtes, ce qui n'est pas vrai pour le contour extérieur d'une union de segments de droite. On fournit également, dans le cas où toutes les extrémités des rayons appartiennent au contour extérieur, un algorithme optimal de complexité $\Theta(n \log n)$ qui construit ce contour.

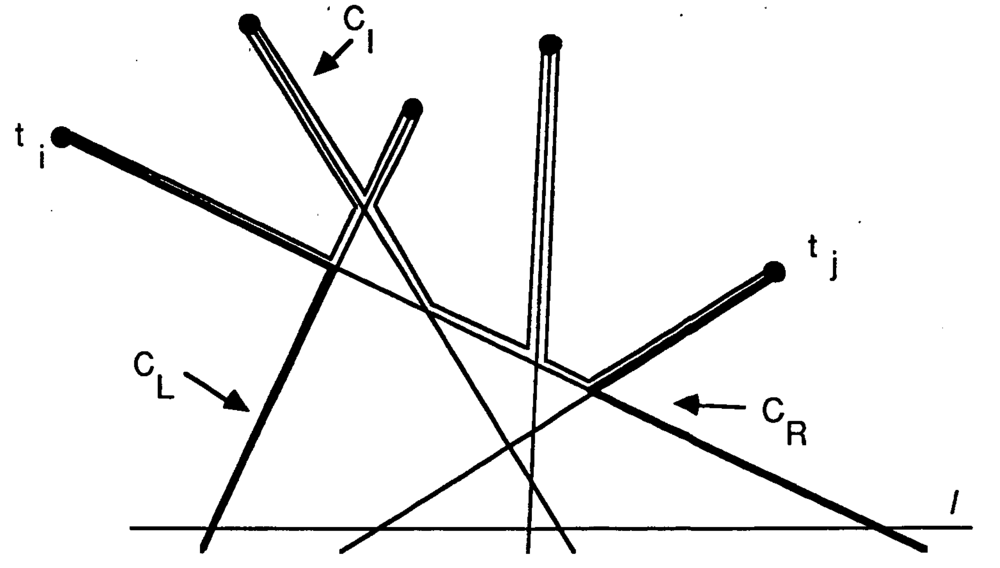

Figure 1: Illustration of the notion of $b(r'_i, \dots, r'_j)$

rests on it.

2 Bundles and mergeable bundles

In their ABY-order, the rays form a string $\rho = r'_1 \dots r'_n$. Crucial to our techniques is the notion of *bundle*. Given a substring $\alpha = r'_i \dots r'_j$ of ρ, the external boundary of $\bigcup_{k=i}^{j} r'_k$ consists of three portions : a polygonal chain C_I, called the *intermediate chain*, between the termini t_i of r'_i and t_j of r'_j, a polygonal chain C_L, called the *left chain*, between infinity and t_i, and a polygonal chain C_R, called the *right chain*, between infinity and t_j. It is immediate to realize that C_L and C_R are convex, and oppose their convex profiles. The bundle pertaining to α, denoted $b(\alpha)$, consists of chains C_L, C_I and C_R (see Figure 1).

We note that a single ray r with terminus t is itself a bundle, pertaining to a singleton string, with C_I degenerating to t and C_L and C_R both coinciding with the ray itself. We also note the following useful property.

Lemma 1 : Chains C_L and C_R are monotone with respect to the line orthogonal to l.

Proof : By contradiction. Let $\alpha = r'_i \dots r'_j$. Suppose that one convex chain of $b(\alpha)$, say C_L, is not monotone. Then there is a vertex p of C_L which is closest to l, and let r' be the first ray whose segment in C_L violates monotonicity (see Figure 2).

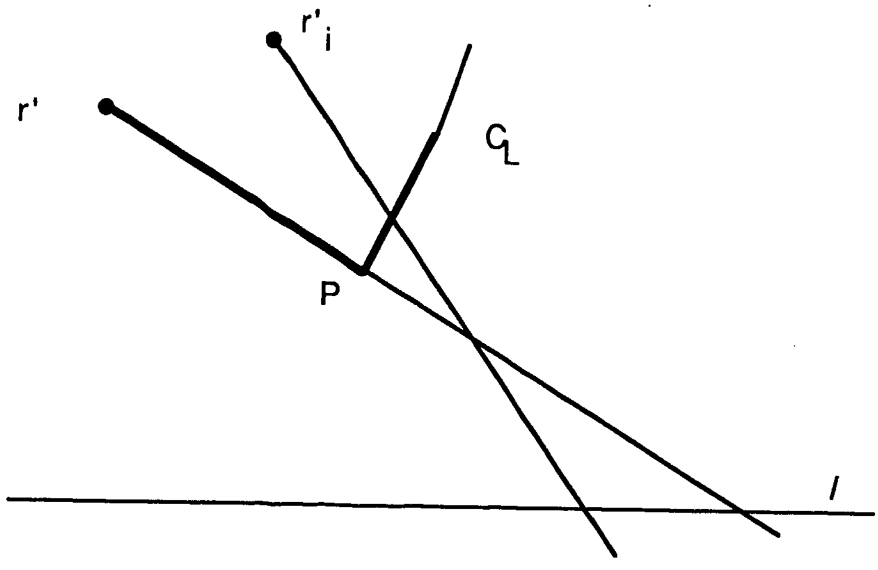

Figure 2: Chain C_L is monotone

Due to the convexity of C_L, r' precedes r'_i in the ABY-order, contrary to the hypothesis that r'_i is the leftmost term of string α. □

Two bundles are disjoint when they pertain to disjoint substrings. We now establish the following crucial property :

Lemma 2 : Any two chains among the right and left chains of two disjoint bundles intersect in at most one point.

Proof: Let (C^i_L, C^i_I, C^i_R) be the chain of $b(\alpha_i)(i = 1, 2)$, with $\alpha_1 \bigcap \alpha_2 = \emptyset$. Referring to their ABY-order, let α_1 be the left bundle.

First we observe that C^1_L and C^2_R cannot intersect. Indeed, suppose they do. Then, due to their opposing convexities, they intersect in two points, p_1 and p_2 with $y(p_1) > y(p_2)$ (see Figure 3). Let r_1 and r_2 be the rays intersecting in p_2, with $r_1 \in \alpha_1$ and $r_2 \in \alpha_2$; according to the ABY-sorting rule given in the introduction, r_2 appears before r_1 in the ABY-order, contradicting the hypothesis.

Consider now the pair (C^1_R, C^2_L). They are both monotone with respect to the vertical (Lemma 1) and with opposing convexities; again they may intersect in at most two points p_1 and p_2 with $y(p_1) > y(p_2)$(see Figure 4). Applying the same argument to the rays intersecting in p_1, we reach an analogous contradiction. Thus C^R_1 and C^L_2 intersect in at most one point (hereafter denoted with the letter "u").

Finally, consider the pair (C^1_L, C^2_L) (the pair (C^1_R, C^2_R) is treated analo-

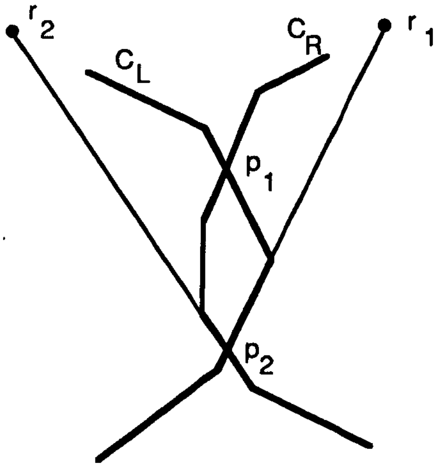

Figure 3: C_L^1 and C_R^2 have at most one point of intersection

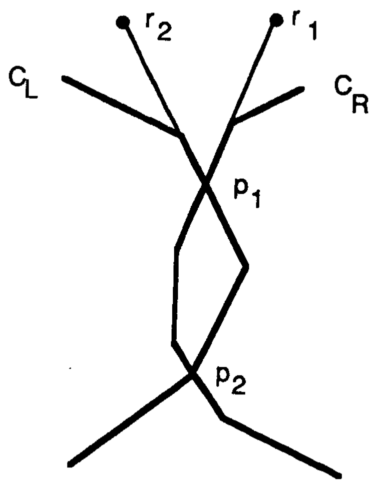

Figure 4: C_R^1 and C_L^2 have at most one point of intersection

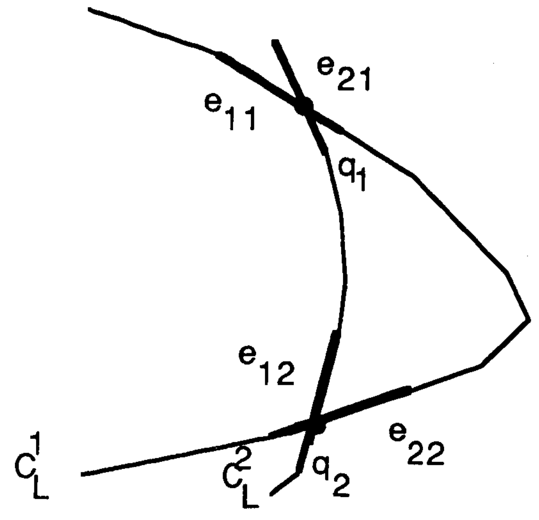

Figure 5: C_L^1 and C_L^2 have at most one point of intersection

gously). Suppose that they intersect in more than one point. This means there are at least two intersections, and we consider two consecutive ones, q_1 and q_2. There are therefore four rays r_{11}, r_{12}, r_{21}, and r_{22}, such that r_{i1} and r_{i2} belong to $C_L^i (i = 1, 2)$, and r_{1j} intersects $r_{2j} (j = 1, 2)$(see Figure 5).

It is immediate to realize that these rays are interlaced in order (r_{11}, r_{21}, r_{12}, r_{22}), contrary to the disjointness hypothesis. Thus C_L^1 and C_L^2 intersect in at most one point (hereafter denoted with the letter "v_L"). □

We can now define the operation of merging two bundles. Given two disjoint bundles $b(\alpha_1)$ and $b(\alpha_2)$, they are classified as "left" and "right" on the basis of the order of their respective substrings in ρ. Two disjoint bundles $b(\alpha_1)$ and $b(\alpha_2)$ are said to be *mergeable* if the right chain of $b(\alpha_1)$ intersects the left chain of $b(\alpha_2)$.

Let $b(\alpha_1)$ and $b(\alpha_2)$ be two mergeable bundles (refer to Figure 6), with $b(\alpha_j) = (C_L^j, C_R^j)(j = 1, 2)$.

Among the chains we have at least one intersection point (u between C_R^1 and C_L^2), and at most three intersection points (u, v_L between C_L^1 and C_L^2, and v_R between C_R^1 and C_R^2). If v_L does not exist, it is conventionally assumed as lying at infinity on both C_L^1 and C_L^2 ; and similary for v_R. Thus, in general, C_L^1 is split by v_L into two branches. The *initial* branch from a ray terminus to v_L and the *terminal* branch from v_L to infinity (the latter may be empty) ; analogously, C_R^2 is split into two branches. Chain C_R^1, instead, is split into three branches : the initial branch, from a ray terminus to u, the intermediate branch, from u to v_R, and the terminal branch from v_R to

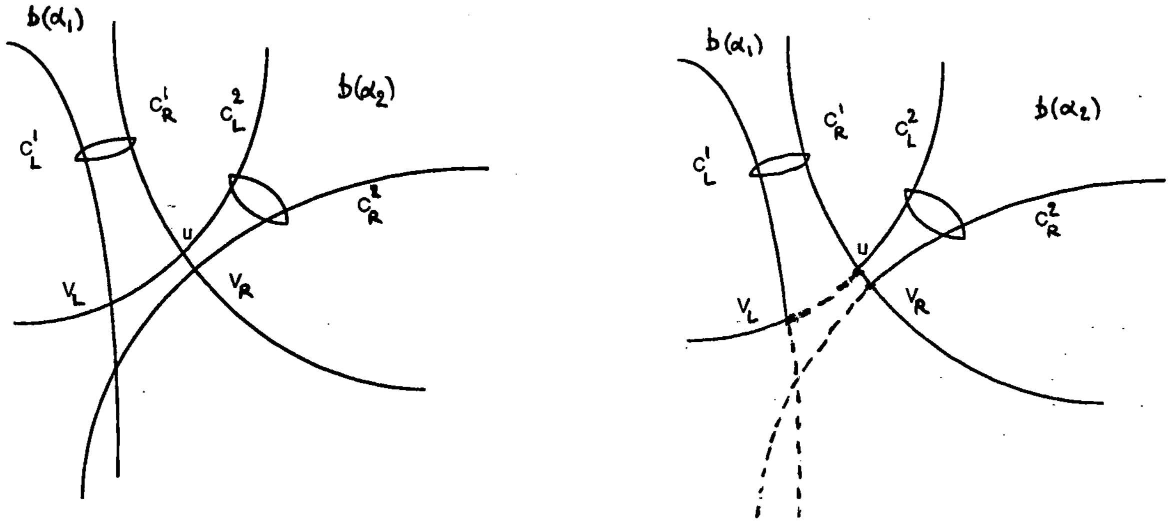

Figure 6: Illustration of the merging operation

infinity (again, the last one may be empty); analogously C_L^2 is split into three branches.

With this nomenclature, the merging of two mergeable bundles gives rise to two structures (see Figure 6) :

1. A *pocket*, concatenating the initial branches of C_R^1 and C_L^2.

2. A *resultant*, consisting in turn of two chains : a chain C_L concatenating the initial branch of C_L^1 and the terminal branch of C_L^2; a chain C_R, concatenating the initial branch of C_R^2 and the terminal branch of C_R^1.

Note that (with the convention that l is the x-axis and that the ray termini have positive ordinates) u, a vertex of the pocket, has minimal ordinate in the pocket (this follows from Lemma 1 and the fact that C_R^1 and C_L^2 have opposing convexities).

In addition, if the two mergeable bundles $b(\alpha_1)$ and $b(\alpha_2)$ are also adjacent, i.e. they pertain to contiguous substrings ($\alpha_1\alpha_2$ is itself a substring of ρ), then the resultant of their merging is itself a legal bundle.

3 Number of edges of E

As previously noted, in their ABY-order, the rays form a string $\rho = r'_1...r'_n$. The ABY-order on the rays induces an order on the edges of the external boundary E. The attributes "before", "after" etc., as applied to edges of E, refer to this order.

Ray r'_n is referred to as the *last* ray of set R, and is also denoted r^*. Ray r^* (with terminus t^*) contributes $k \geq 1$ edges of E (i.e. contains k edges of E) before t^* (in the ABY-order). We denote them $e_1, ..., e_k$, in the order

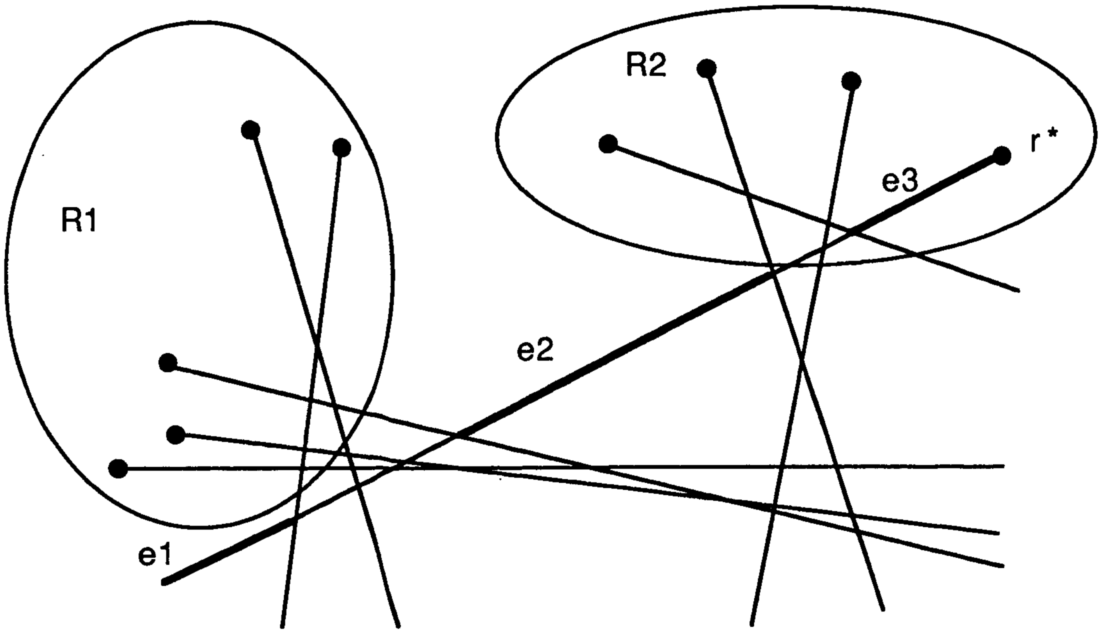

Figure 7: Subsets R_1 and R_2

above.

If e_1 is itself a ray, we define R_1 as the subset of R consisting of all rays containing the edges of E preceding e_1 (see Figure 7).

Clearly a substring $\alpha_1 = r'_1...r'_m$ of ρ is associated to R_1 and the complementary substring $\alpha_2 = r'_{m+1}...r'_n$ is associated to the subset $R_2 \triangleq R - R_1$.

Thus the bundles $b(\alpha_1)$ and $b(\alpha_2)$ pertaining to α_1 and to α_2 are disjoint and mergeable.

With this nomenclature, we are ready to prove the main theorem of this section :

Theorem 1 : The external boundary of the union of n rays has at most $4n - 2$ edges.

Proof : The proof is by induction. The theorem obviously holds for $n = 1$; we assume that it holds for any positive integer $t \leq n - 1$.

Let us denote E_i the external boundary of the union of the rays of set R_i, and $|E_i|$ the number of edges of $E_i (i = 1, 2)$. According to the inductive hypothesis, we have $|E_1| \leq 4m - 2$ and $|E_2| \leq 4(n - m) - 2$.

Let us denote (C^i_L, C^i_I, C^i_R) the chains of $b(\alpha_i)(i = 1, 2)$. Because C^1_I appears between e_1 and e_2, it cannot intersect R_2 and similarly, because C^2_I appears between e_2 and e_3, it cannot intersect R_1. Thus the only edges of E_1 and E_2 which intersect belong to right or left chains of the associated bundles.

Due to Lemma 2, these chains have at most three points of intersection.

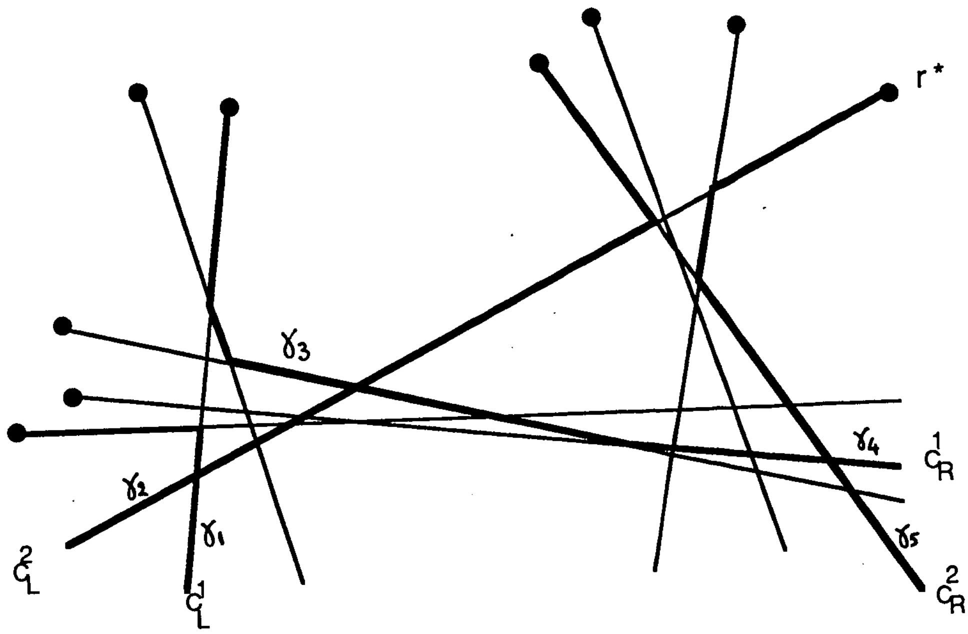

Figure 8: Intersections between E_1 and E_2

Specifically, we have at most (refer to Figure 8):

- an intersection between an edge of C_L^1, say γ_1, and the first edge of C_L^2, say γ_2, contained in r^*.
- an intersection between an edge of C_R^1, say γ_3, and edge γ_2 of C_L^2.
- an intersection between an edge of C_R^1, say γ_4, and an edge of C_R^2, say γ_5.

The external boundary E of the union of all the rays $(R_1 \bigcup R_2)$ is the external boundary of the union of the external boundaries E_1 and E_2. It is clear, from the above discussion, that at most one edge of E_1 (γ_3) and at most one edge of E_2 (γ_2) may be cut into three pieces when merging the two bundles, two of them appearing as edges of E. Thus we have

$$|E| \leq |E_1| + |E_2| + 2 \leq 4n - 2$$

which completes the proof. □

4 Construction of E

Before describing an algorithm for computing the external boundary E of F, we observe that $\Omega(n \log n)$ is a lower bound for the time complexity of any algorithm computing E. Indeed let $x_1, ..., x_n$ be n real numbers. For each $x_i (i = 1, ..., n)$, we construct the vertical ray with terminus $(x_i, 1)$ and extending to $y = -\infty$. We add the horizontal ray with terminus $(\max x_i + 1, 0)$ and extending to $x = -\infty$. The external contour of this set gives the sorting of the given numbers. Thus sorting is $O(n)$ transformable to computing the external boundary of a set of rays.

In this section, we explicitly assume that E passes through all ray termini. As previously noted, the external boundary E of F consists of a sequence of convex polygonal chains, each delimited by two consecutive ray termini (where the points at infinity of the first and last edges of E are conventionally treated as termini). Now, if each of these chains is viewed as a pocket generated by the merging of two adjacent bundles, then there is a family of merging schedules where pockets are exactly those found on the external contour of F. Any such schedule must have the property that, when merging two adjacent mergeable bundles, the ensuing pocket is not intersected by any ray external to the substring of the resulting bundle.

The strategy we are about to describe exhibits the above property.

The intersection of two adjacent bundles (which exists only if the two bundles are also mergeable) is the vertex u obtained in their merging (see Section 2).

Initially, each ray is a bundle, and all rays are in their ABY-order. For each pair of adjacent bundles we compute its intersection ; the existing intersections are ordered by decreasing ordinate.

The general step of the algorithm selects the bundle intersection with maximum ordinate ; let $b(\alpha_1)$ and $b(\alpha_2)$ (in this order), be the two bundles whose intersection is u. We assign their pocket to the contour and replace the pair $b(\alpha_1)$ and $b(\alpha_2)$ with their resultant $b(\alpha_1\alpha_2)$. We also update the intersections with the bundles respectively adjacent to $b(\alpha_1)$ to the left and to $b(\alpha_2)$ to the right.

The process terminates when there is just one bundle.

To establish the correctness on the outlined approach, consider the pocket generated by the current merging at $b(\alpha_1)$ and $b(\alpha_2)$ and assume for a contradiction, that there is a ray r, not belonging to $\alpha_1\alpha_2$ which intersects its interior (refer to Figure 9).

Let $p \neq u$ be the intersection of r with the pocket having largest ordinate. Since u has minimal ordinate in the pocket, $y(p) > y(u)$. Assume, without loss of generality, that r precedes the substring $\alpha_1\alpha_2$ in ρ : then ρ intersects the left chain of $b(\alpha_1)$ in a point q such that $y(q) \geq y(p)$. Since the intersection of two bundles is the maximum of the pairwise intersections of their respective members, $y(q)$ is not larger than the ordinate of the intersection u' of $b(\alpha_1)$ with the bundles to which r belongs. But this shows $y(u') \geq y(q) \geq y(p) > y(u)$, contrary to the hypothesis that $y(u)$ is maximum. Therefore :

Theorem 2 : The outlined procedure correctly generates the external boundary of F.

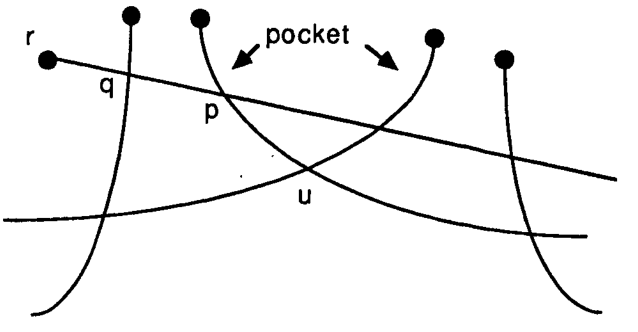

Figure 9: For the proof of correctness

We now turn our attention to the performance of the above procedure :
We observe :

1. The ABY-order of the rays is obtained in time $O(n \log n)$ [2].

2. The initial computation of the intersections of the $(n-1)$ adjacent pairs of bundles and their sorting by ordinate is also accomplished in time $O(n \log n)$.

3. The general step of the algorithm involves intersecting three pairs of convex monotone chains (each pair intersecting in at most one point). A chain is supposed to be stored in a height-balanced (mergeable) tree. We shall distinguish two cases depending upon whether the two chains have opposing or concordant convexities. Let C' and C'' be the two chains in question, where C' belongs to the left bundle and C'' to the right bundle.

a) C' and C'' have opposing convexities. In this case, we apply a bisection technique. Let e' be an edge of C' and e'' an edge of C'', and let lines l' and l'' contain e' and e'' respectively, with p denoting the intersection of l' and l''. Edge e', oriented by decreasing ordinate, is classified as "before", "crossing", or "after" depending upon whether its extremes both precede, both follows or contain point p. Thus, e' and e'' give rise to nine cases, illustrated in Figure 10 (symmetric cases being omitted).

In the case (crossing, crossing) the task is completed ; in all other cases we can eliminate a subchain (shown by a wiggly line) because it cannot contain the intersection. An additional comment is needed for the case (before, before), where we eliminate the subchain whose terminus has larger ordinate. In this manner, at each step we eliminate a bounded fraction of one chain, which ensures termination in time $O(\log s)$, where s is the larger of the two

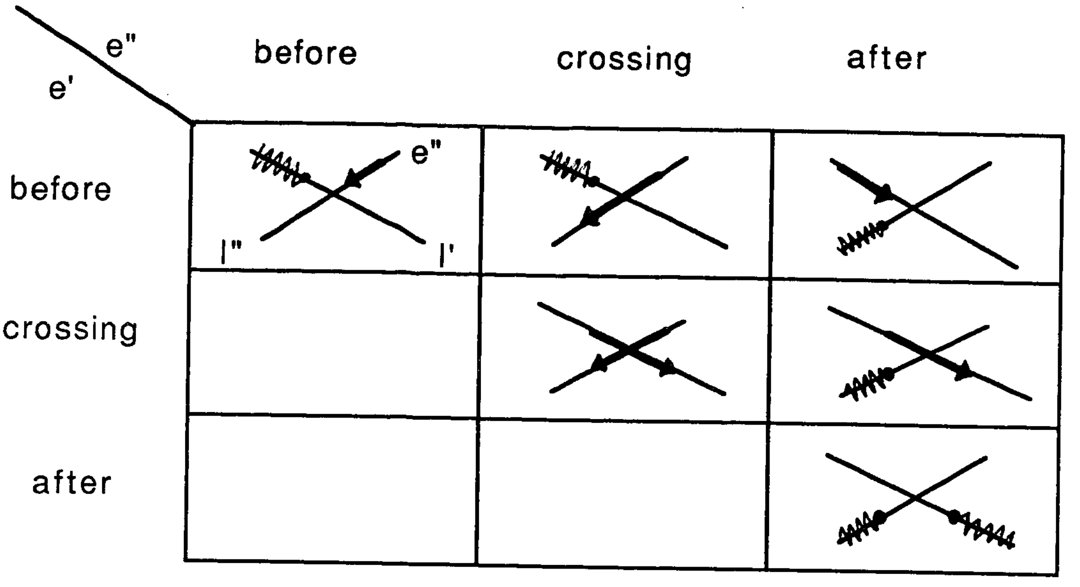

Figure 10: Intersection of two chains with opposing convexities

chain sizes.

b) C' and C'' have concordant convexities. Referring to Figure 11, each chain terminates with a ray (r' for C', and r'' for C'').

A necessary and sufficient condition for C' and C'' to have one (and only one) intersection is that the slope of r'' be smaller than the slope of r': indeed, in the opposite case, C'' lies entirely to the right of C', since C' and C'' have at most one intersection. If the slope test indicates the existence of an intersection, the latter is determined as follows. We "march" on C' by increasing ordinate and we "jump" on C'' by a bisection strategy.

The idea -analogous to that presented in [7]- is not to advance on the chain whose current edge may contain a yet to be found intersection. Initially the two edges are the chain rays r' and r''. We advance on C' until the current edge e_1 crosses the line l_0 containing r'' ; at this point, we stop the march on C' and, by an additional step of binary search, determine an edge e_2 of C'' which either intersects e_1 (thus completing the task) or lies entirely to the left of the (detected) line l_1 containing e_1 ; and so on.

This process terminates in time $O(\log r + s)$, where r is the size of C' and s is the size of the terminal branch of C''. Note that this terminal branch is traversed only once.

Moreover, the intersection with adjacent bundles have to be updated. From the discussion of case b) above, each intersection is computed in $O(\log n)$ time. If we store the ordinates of the intersections in a priority queue, each general step updated involves three deletions and two insertions, each executable in time $O(\log n)$.

In summary, the work due to the $(n-1)$ executions of the general step is

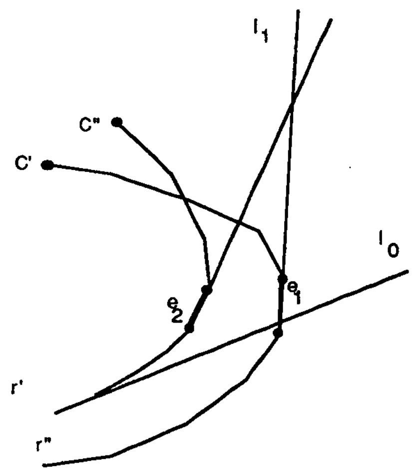

Figure 11: Intersection of two chains with concordant convexities

$O(n \log n) + \sum_{j=1}^{n-1} s_j$, where s_j is the size of the terminal branches eliminated at the j-th step. Since the external contour of a set of n rays has $O(n)$ edges, it follows that $\sum_{j=1}^{n-1} s_j$ is also $O(n)$. Therefore we conclude :

Theorem 3 : The external contour of the union of n rays can be constructed in time $O(n \log n)$.

5 Extensions and variants

5.1 Justification of assumptions

Throughout the previous sections, we have assumed that E passes through all ray termini. This assumption is not restrictive.

Let us suppose that a ray terminus is not on the external boundary; clearly, if we extend the ray so that it crosses the external boundary and its terminus appears on the external boundary, the number of edges of E increases. Thus the maximum number of edges of the external boundary is reached when all the ray termini are on the external boundary.

When it is not known whether all ray termini lie on E, we are not aware of any $O(n \log n)$ time algorithm to compute E. However, the following divide-and-conquer technique, running in time $O(n \log^2 n)$ can be used to solve the problem. The set of rays is arbitrarily subdivided into two equal-size subsets, for which the external contours are recursively computed. Viewing each such contour as a simple polygon closed through the line at infinity of the plane, we obtain two simple intersecting polygons. According to an observation of Pollack, Sharir and Sifrony [5], the contour of the union of these two polygons can be obtained by applying the ray-shooting technique of Chazelle and Guibas [3] in time $O(\log n)$ per contour edge, i.e. in total time $O(n \log n)$. Since this is the time of the merge step, we obtain -as noted- an $O(n \log^2 n)$-time algorithm.

5.2 The case of wedges

For ease of presentation, we have considered rays in the previous sections. However our results extend in a straight forward manner to the case of wedges. Rays are simply particular cases of wedges, whose branches are coincident.

5.3 The case of line segments

Wiernick and Sharir [8] have proved that the upper-envelope of the union of n line segments may have $O(n\alpha(n))$ edges where $\alpha(n)$ is the inverse Ackermann's function. However Theorem 1 proves that this bound is not tight when the line segments are long enough.

The following proposition is a direct consequence of Theorem 1.

Proposition 1 : The external boundary of a set of n line segments each intersecting a given line l has $O(n)$ edges.

Proof : Suppose without loss of generality that l is the x axis. Each line segment e_i has an end-point p_i with positive ordinate and the other q_i with negative ordinate. Segment e_i can be considered as the intersection of two rays r_P^i with terminus p_i and r_Q^i with terminus q_i. Due to Theorem 1, both the external boundary E_P of the union of the $r_P^i (i = 1, n)$ and the external boundary E_Q of the union of the $r_Q^i (i = 1, n)$ have $O(n)$ edges. The total number of edges of the external boundary of the segments is not greater than the sum of the numbers of edges of E_P and E_Q which is $O(n)$. The total number of edges of the external boundary of the segments is not greater than the sum of the numbers of edges of E_P and E_Q which is $O(n)$. □

Remark : The algorithms given in Sections 4 and 5.1 can each be easily adapted to compute the external boundary of n such line segments within the same time bounds.

Let us suppose now that each one of the n line segments intersects at least one of k given lines $l_1, ..., l_k$. We partition the set S of line segments into k classes $S_1, ..., S_k$. Each element of S_i intersects $l_i (i = 1, ..., k)$. If a line segment intersects more than one l_i we arbitrarily put it in one of the corresponding classes. The external boundaries E_i of the union of the line segments of class S_i is a simple polygonal line with, at most, $O(n_i)$ edges if n_i is the number of elements of $S_i (i = 1, k)$. Clearly the external boundary E of the union of the line segments is the external boundary of the union of the $E_i (i = 1, ..., k)$.

In order to prove that the number of edges of E is $O(n)$, we will prove that the external boundary of the union of k simple polygons, for fixed k, has at most $O(n)$ edges, if n is the total number of edges of the polygons.

Proposition 2 :The external boundary of the union of k simple polygons with a total number of n edges has at most kn edges.

Proof: 1. $k = 2$: Let P and Q be two polygons, E the external boundary of their union. An edge of E is either contained in an edge of P or in an edge of Q. If an edge of E is contained in an edge e of P or Q, we label it with e. Thus we associate to E a circular sequence L of labels.

An edge of P or Q will be called of type C_i if it contains i edges of E. Let γ_i be the number of such edges and m the maximum of i.

The following observation is due to Pollack, Sharir and Sifrony [5] : L cannot contain a subsequence of the form $...e...f...e...f...e...$

From the above observation, we deduce that among the edges of P (resp.Q) intersecting a given edge e of Q (resp. P), the only ones which may be of type $C_i, i > 1$, are the first one and the last one (when marching along E). Moreover these "extreme" edges cannot intersect e twice.

As a consequence, if $e \in P$ (resp. Q) and is of type $C_j, j \geq 3$, Q (resp.P) has at most two edges of type $C_i, i > 1$, and at least $2j - 4$ edges of type C_1. Thus

$$\gamma_i \geq \sum_{j=3}^{m} (2j-4)\gamma_j.$$

Besides we have

$$\sum_{i=1}^{m} \gamma_i = n$$

and

$$\sum_{i=1}^{m} i\gamma_j = |E|.$$

Hence we conclude that

$$|E| \leq 2n - \frac{\gamma_1}{2} \leq 2n.$$

2. $k > 2$: By recursively splitting the set of k polygons into two subsets with no more than $\lceil \frac{k}{2} \rceil$ polygons, we get $|E_k| \leq 2|E_{\frac{k}{2}}| \leq kn$. □

Applying this result to our initial problem, proves the following theorem:

Theorem 4 : Let $l_1, ..., l_k$ be a given set of lines, and let S be a set of n line segments $e_1, ..., e_n$ such that

$$\forall i, \exists j : e_i \bigcap l_j \neq \emptyset.$$

Then the external boundary of the union of the line segment $e_1, ..., e_n$ has at most $O(n)$ edges.

Remark : This external boundary can also be efficiently computed using the algorithms of Sections 4 and 5.1. First we calculate all the E_i's and subsequently we merge the boundaries $E_1, ..., E_k$ -according to an arbitrary merge schedule- using the already mentionned algorithm of Pollack, Sharir and Sifrony [5]. It is immediate to verify that the time bounds of Sections 4 and 5.1 are applicable to this problem.

5.4 Other connected components of the boundary of F

Theorem 1 extends to the case of any connected component (or cycle) of the boundary of F. This is obvious for those cycles which do not contain a terminus. Let us consider a cycle I containing $k \geq 1$ ray termini $r'_1, ..., r'_k$. We denote by C the smallest convex polygon bounded by rays which encloses the k ray termini. The cycle I consists of edges of C and of edges contained in rays $r'_1 ... r'_k$.

We partition the set of rays $R = \{r'_i, ..., r'_k\}$ into disjoint subsets consisting of all the rays intersecting C between two successive occurrences of edges of I contained in edges of C. Let m be the number of such subsets and n_i the number of rays pertaining to subset $i (i = 1, ..., m)$. Disjoint mergeable bundles are associated to the subsets and we have

$$|I| \leq \sum_{i=1}^{m}((4n_i - 2) + |C| + m).$$

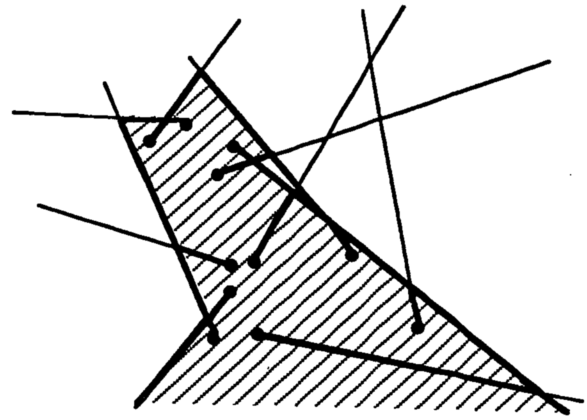

Figure 12: Illustration of the smallest polygon enclosing the object

Indeed Theorem 1 holds for each bundle and merging C and the bundles adds at most m edges to I. With $|C| < n$ and $\sum_{i=1}^{m} n_i < n$, we conclude $|I| < 5n - m$. Thus we can reformulate Theorem 1 as

Theorem 5 : Any connected component of the union of n rays has at most $O(n)$ edges.

We leave as an open problem the question of whether or not an $O(n \log n)$ algorithm can be designed to compute a given connected component.

5.5 Application to contour reconstruction from rays

In [2], the following problem was presented : a robot moves around an unknown object; the robot is equipped with an optical device, such as a laser range-finder, which measures the coordinates of points on the contour of a planar cross-section of the object. Thus the information about the object consists of a set of points and of a set of rays issuing from these points that are known not to pass through the object. In [2], an $\Theta(n \log n)$ algorithm is described which computes the unique simple polygon with the measured points as vertices and which is not intersected by the rays. If we don't content ourselves with a polygonal approximation of the object, but if we want to know the smallest polygon C which surely contains the object (see Figure 12), we can apply the results of this paper and solve this problem in $\Theta(n \log n)$ time.

Indeed polygon C is the connected component of the union of the rays which passes through all the measured points. Let $e = [ab]$ be an edge of CH, the convex hull of the measured points, a and b its endpoints. The portion of C between a and b is the external boundary of the rays intersecting e. C is simply the concatenation of the different portions corresponding to the different edges of CH. Because a ray can only intersect one edge of CH, the total time complexity is clearly $O(n \log n)$.

6 Concluding remarks

The main result of this paper establishes a linear upper bound for the number of edges of the external boundary of the union of n rays. This result improves the previous $O(n\alpha(n))$ upper bound which is known to be tight if we consider line segments instead of rays. Thus an interesting distinction is introduced between half-lines and line segments. Moreover, for some special configurations of line segments, our result applies, yielding an $O(n)$ upper bound for the number of edges of their union. Such is the case if there exists a bounded number of lines intersecting all the line segments.

It would be interesting to find other special configurations admitting such linear bound. In particular, does the bound apply to the external boundary of n line segments whose directions belong to a finite set ? This is known to be true for the case of line segments restricted to be either horizontal or vertical, i.e. to belong to just two directions [6].

Several other questions remain open. In particular, does the bound hold for the nontrivial boundary of the union of n rays ? The nontrivial boundary is defined as the union of all connected components of the boundary, each of which contains a ray terminus.

From the algorithmic point of view, our main result is an $\Theta(n \log n)$ algorithm which constructs the external boundary of the union of n rays when the ray termini are known to lie on the boundary. Several open questions remain : is it possible to construct the boundary of a general collection of rays in time $\Theta(n \log n)$? Is it possible to construct any given connected component of the boundary also in time $O(n \log n)$?

References

[1] M. Atallah : Dynamic Computational Geometry, IEEE Symposium on Theory of Computing, 1983.

[2] P. Alevizos, J.D. Boissonnat, M. Yvinec : An optimal $O(n \log n)$ Algorithm for Contour Reconstruction from Rays, ACM Symposium on

Computational Geometry, Waterloo, June 1987.

[3] B. Chazelle, L. Guibas : Visibility and Intersection Problems in Plane Geometry, ACM Symposium on Computational Geometry, Baltimore, June 1985.

[4] S. Hart, M. Sharir : Non Linearity of Davenport-Schinzel Sequences and of Generalized Path Compression Schemes, Combinatorica 6(2), 1986, pp 151-177.

[5] R. Pollack, M. Sharir, S. Sifrony : Separating Two Simple Polygons by a Sequence of Translations, Technical Report No 215, April 1986, NYU Courant Institute.

[6] W. Lipski, Jr., F.P. Preparata, Finding the contour of a union of iso-oriented rectangles, Journal of Algorithms 1,235-246 (1980).

[7] J. O'Rourke, C.-B Chien, T. Olson, D. Naddor : A New Linear Algorithm for Intersecting Convex Polygons, Computer Graphics and Image Processing, 19, 384-391 (1982).

[8] A. Wiernik, M. Sharir : Planar Realizations of Non-linear Davenport - Schinzel sequences by segments, Technical Report No 224, 1986, NYU Courant Institute.

Imprimé en France

par

l'Institut National de Recherche en Informatique et en Automatique