Arrangements of Segments that Share Endpoints: Single Face Results

Esther M. Arkin(1) * Dan Halperin(2) Klara Kedem(2) †
Joseph S. B. Mitchell(1) ‡ Nir Naor(2)

(1)Department of Applied Mathematics and Statistics
State University of New York, Stony Brook, NY 11794-3600
(2)Department of Computer Science
School of Mathematical Sciences, Tel Aviv University

February 6, 1994

Abstract

We provide new combinatorial bounds on the complexity of a face in an arrangement of segments in the plane. In particular, we show that the complexity of a single face in an arrangement of n line segments determined by h endpoints is $O(h \log h)$. While the previous upper bound, $O(n\alpha(n))$, is tight for segments with distinct endpoints, it is far from being optimal when $n = \Omega(h^2)$. Our result shows that, in a sense, the fundamental combinatorial complexity of a face arises not as a result of the number of *segments*, but rather as a result of the number of *endpoints*.

1 Introduction

Let $S = \{s_1, \ldots, s_n\}$ be a finite set of n line segments in the plane. Then S induces a partition of the plane, called the *arrangement* $A(S)$ of S, into $O(n^2)$ faces, edges and vertices. Refer to Figure 1. Arrangements of segments play a fundamental role in computational geometry (see, e.g., [EGS]).

The problem studied here is that in which the segments S are allowed to share endpoints, so that the total number of endpoints, h, may be substantially less than $2n$. In this context, we provide improved combinatorial bounds on the complexity of a face in an arrangement of segments or pseudo-segments in the plane. In particular, we show that the complexity of a single face in an arrangement of n line segments determined by h endpoints is $O(h \log h)$. The previously known upper bound ([EGS],[PSS]) was $O(n\alpha(n))$. When $n = O(h)$, this bound is tight, since there are arrangements of h segments such that the complexity of a single face is $\Omega(h\alpha(h))$ [Sh, WS]. However, when n is superlinear in h, the known upper bound is far from being optimal. In fact,

*Partially supported by NSF Grants ECSE-8857642 and CCR-9204585.

†Partially supported by the Pikkovski-Valazzi Fund and by the Eshkol grant 04601-90 from the Israeli Ministry of Science and Technology.

‡Partially supported by a grant from Hughes Research Laboratories, Malibu, CA, and by NSF Grants ECSE-8857642 and CCR-9204585.

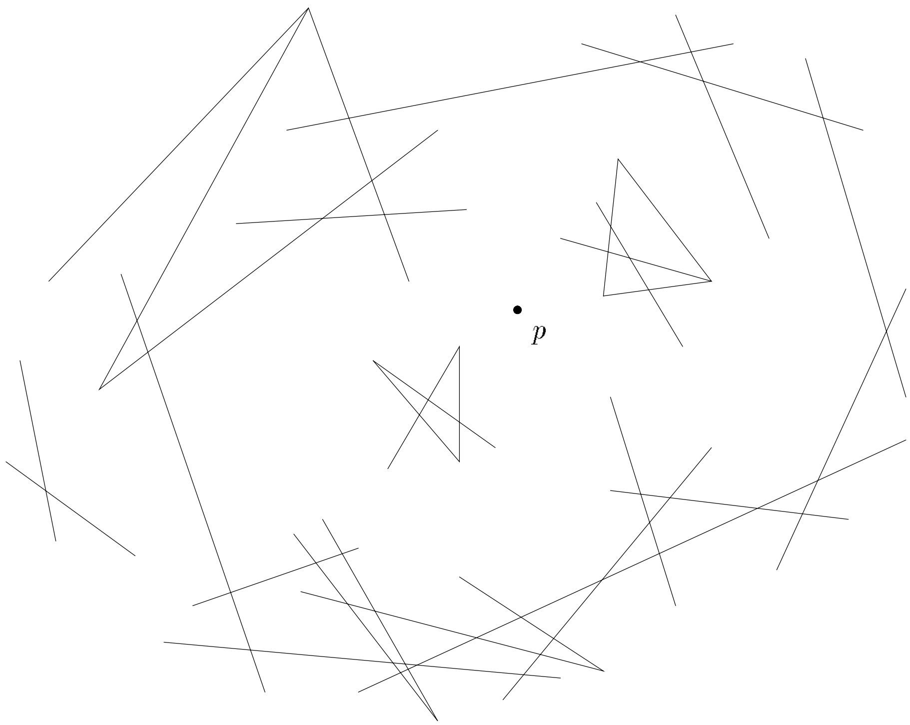

Figure 1: An arrangement of segments.

n may be quadratic in h, in which case the upper bound of $O(h^2\alpha(h))$ is an order of magnitude greater than the lower bound. Our result shows, in a sense, that the inherent complexity of a face is due not to the *segments* themselves, but to their *endpoints*. We expect that this same phenomenon is manifested in a broad class of combinatorial problems on arrangements.

Another way to state this result is to say that the complexity of any single face in a drawing of a graph with h nodes and $n = O(h^2)$ straight edges is $O(h \log h)$, improving the previously known bound of $O(n\alpha(n))$ implied by [EGS, PSS] for the cases in which n is significantly larger than h.

Our result is established by a new technique, which we call the method of "favorable stabbers", that partitions the set of segments into $O(\log h)$ classes in a manner very similar to that of interval trees, except that instead of *straight, vertical* "stabbing" lines determining which segments belong to which class, we define a set of "parallel" *pseudo-lines* that determine the classes. The pseudo-lines are chosen in such a way that the segments stabbed by any one pseudo-line are "favorable" in that the complexity of the face containing a fixed point p in the arrangement of the stabbed segments together with the pseudo-line has complexity at most $O(h)$ (rather than $O(n\alpha(n))$).

This paper is organized as follows. In the next section, we begin with some "warm-up" exercises that demonstrate the appearance of h in various combinatorial bounds on arrangements. In Section 3, we give our main result on the complexity of a single face. Finally, we conclude with some open problems discussed in Section 4.

Remark. This paper corrects an error in an earlier draft ([AHKMN]), where it was claimed that only $O(h)$ line segments can be "exposed" on any single face of $A(S)$ (which then immediately implies a face complexity of $O(h\alpha(h))$). There is an error in Lemma 11 of [AHKMN], as has been pointed out to us by A. Lubiw and J. Spinrad. It remains open how many line segments are exposed on any one face. Our results here show an upper bound of $O(h \log h)$. In the case of pseudo-segments, A. Lubiw and S. Suri have communicated to us an example showing that in fact there is a $\Omega(h \log h)$ lower bound on the complexity of a single face.

2 Warm-Up Results

In order to get a feel for the nature of the results we prove, we begin with some simple problems.

First, we show that the complexity of a single face in an arrangement of n rays with a total of h termini is $\Theta(h)$, generalizing the bound of $O(n)$ given by [ABP].

Theorem 1 *The complexity of a single face in an arrangement of n rays with a total of h termini is $\Theta(h)$.*

Proof. Let the face f of interest be designated by a point p. Consider a single terminus v, which may have many rays emanating from it. The rays emanating from v partition the plane into cones. Exactly one of these cones contains the point p. Thus, the only rays emanating from v that can possibly participate in the boundary of the face f containing p are the two rays bounding the cone that contains p. Thus, there are only $2h$ rays that can participate in the face f, so we are done, by the results of [ABP]. □

Next, we show that the *lower envelope* of n segments that have h endpoints is only $O(h \log h)$. We begin with the following simple result, which will form a motivating basis for our method of "favorable stabbers" when we prove the general single-face result in the next section.

Lemma 2 *Let S be a set of n line segments with a total of h endpoints, such that all segments are stabbed by a vertical line ℓ. Then the complexity of the lower envelope of S is $O(h)$.*

Proof. It suffices to show that the complexity of the lower envelope to the left of ℓ is $O(h)$. First, notice that, for a given endpoint v to the left of ℓ, we can discard all segments incident on v except for the bottommost one. We are left with $O(h)$ segments. Now, considering only contributions to the lower envelope to the left of ℓ, it is easy to see that there can be no pattern of the form "a...b...a...b" in the listing of contributions to the lower envelope, where "a" and "b" denote labels of the $O(h)$ bottommost segments. Refer to Figure 2. Then, by the theory of Davenport-Schinzel sequences ([HS]), we know that the lower envelope complexity to the left of ℓ (and, similarly, also to the right of ℓ) is $O(h)$. □

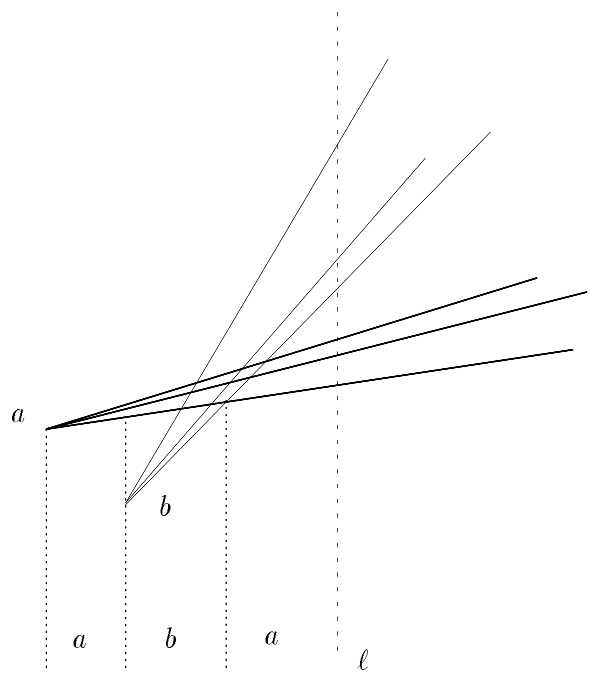

Figure 2: No "a...b...a...b" contribution to the lower envelope.

It is now easy to get the following theorem:

Theorem 3 *Let S be a set of n line segments with a total of h endpoints. The complexity of the lower envelope of S is $O(h \log h)$.*

Proof. Consider a vertical line ℓ that bisects the set of endpoints of S and the corresponding partitioning of the segments S into three groups: S_L (those with both endpoints left of ℓ), S_R (those with both endpoints right of ℓ), and S_ℓ (those stabbed by ℓ). By the previous lemma, we know that the lower envelope of the set S_ℓ has complexity $O(h)$.

Let $g(h, n)$ be the worst-case complexity of the lower envelope of n segments with h endpoints ($h/2 \leq n \leq \binom{h}{2}$). Then the lower envelope of S_L (resp., S_R) has complexity bounded by $g(h/2, |S_L|)$

(resp., $g(h/2, |S_R|)$). The lower envelope of S is obtained by combining (merging) the lower envelopes of S_L, S_ℓ, and S_R. By applying the Combination Lemma of [EGS] (specialized to the case of lower envelopes), we get that the complexity of the lower envelope of S is at most the sums of the complexities of the lower envelopes of S_L, S_ℓ, and S_R, *plus* $O(h)$ (since only the h endpoints can be "reflex" vertices with respect to the polygon below the envelopes). We therefore get the following recursion:

$$g(h, n) = g(h/2, |S_L|) + g(h/2, |S_R|) + O(h),$$

with $g(2, 1) = O(1)$. The solution is $g(h, n) = O(h \log h)$. □

We now turn to the more general question of bounding the complexity of a single face in an arrangement of segments. We obtain a bound of $O(h^{1.5}\alpha(h))$, using a result from extremal graph theory. It was the discovery of this result that first suggested to us that subquadratic (in h) bounds were possible for this problem.

Theorem 4 *Let S be a set of n line segments with a total of h endpoints. Then at most $O(h^{1.5})$ segments can participate in the boundary of any single face of $A(S)$, which implies that the complexity of any single face of $A(S)$ is at most $O(h^{1.5}\alpha(h))$.*

Proof. We orient each segment from left to right, and we look only at those segments that have a contribution to the face of interest on their *top*. The proof is based on the following claim: *The directed graph that we draw in this way can have no subgraph that is a directed $K_{2,2}$ with 2 nodes of in-degree two and 2 nodes of out-degree two.* This claim is checked by examining all possible ways of embedding a directed $K_{2,2}$ in the plane with straight edges. Refer to Figure 3. There are two cases: (1). The four segments defining edges of a $K_{2,2}$ define a simple polygon (as shown on the right in the figure); and (2). Two of the segments cross (as shown on the left). In case (1), we get a contradiction to the fact that all four segments should have a top contribution, since two of the segments have their top exposed on one side of the quadrilateral, while the other two have their tops exposed on the other side of the quadrilateral. In case (2), one of the two triangular faces in the arrangement of four segments must contain the top contribution of one segment, while the other three segments have their top contributions outside this triangle.

A result in extremal graph theory (see [Lo], problem 10.36(a)) shows that the number of edges in a directed graph with n nodes having no such directed $K_{2,2}$ subgraph can be at most $O(n^{1.5})$, completing the proof. □

Remark. The above result will be improved to show that the worst-case complexity of a single face is $O(h \log h)$, in Theorem 8 in the next section.

As a final opening remark, note that when $n = \binom{h}{2}$ (the case of a complete graph on h vertices), it is clear that the complexity of a single face is at most h, since all bounded cells are convex, and at most one of the edges incident on any vertex can contribute on its "right side" to any one face. It is possible, however, for $n = \Omega(h^2)$, while there exist faces of complexity $\Omega(h\alpha(h))$. (Construct a complete graph on $h/2$ vertices, and, independently, construct a configuration of $h/4$ segments with lower envelope of size $\Omega(h\alpha(h))$.) This implies that an arrangement that is "nearly" a complete graph does not necessarily give rise to faces of complexity $O(h)$. (See, however, the discussion of open problem (2) in the conclusion of the paper.)

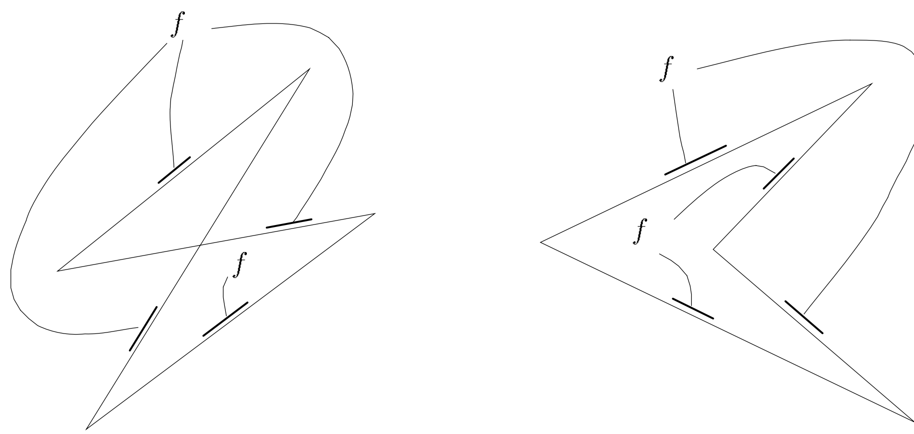

Figure 3: Embedding a directed $K_{2,2}$ separates the face f.

3 The Complexity of a Single Face

We assume, without loss of generality, that the face of interest in $A(S)$ is the face at infinity, denoted by $\phi(S)$; we consider $\phi(S)$ to be an *open* set, and denote its boundary by $\partial(\phi(S))$. Our goal in this section is to show that the complexity of $\phi(S)$ is $O(h \log h)$, where h is the number of endpoints of S.

In general, $\phi(S)$ is multiply connected, and its boundary consists of several connected components. It suffices for us to consider a single connected component of the boundary of $\phi(S)$ and to show that its complexity is at most $O(h' \log h')$, where h' is the number of endpoints of segments of S taking part in one such component; summing over all components, then, gives the desired bound on the total complexity of $\phi(S)$.

Thus, let γ denote a connected component of the boundary of $\phi(S)$. Then, γ is a (*possibly degenerate*) simple polygon whose edges are subsegments of those segments of S that are surrounded by $\phi(S)$. The degeneracy that is potentially present in γ is that two of its edges may coincide — one edge corresponding to a subsegment on one "side" of a segment $s_i \in S$, and another edge corresponding to the same subsegment on the other "side" of s_i. We refer to an edge of γ that is a subsegment of $s_i \in S$ as a *contribution* of s_i to the face at infinity.

Our proof that the complexity of γ is $O(h \log h)$ is strongly motivated by the proof of Theorem 3: We determine an endpoint bisecting "stabber" ℓ such that the "stabbed" segments' contribution to γ is only *linear* in the number, h_ℓ, of their endpoints. We then apply the Combination Lemma of [EGS] to obtain a recursion whose solution gives the claimed result. In order to make these statements precise, we need several definitions.

Let ℓ be a simple, unbounded (directed) curve in the plane; thus, ℓ divides the plane into two closed *halfspaces* — that component of $\Re^2$ to the left of ℓ (denoted L) and that component to the right of ℓ (denoted R). Let $int(L)$ ($int(R)$) denote the interior of L (R). We say that a curve ℓ is a *pseudo-line with respect to S* if, for every segment $s \in S$ that intersects ℓ, $s \cap \ell$ is connected.

In the following definitions, assume that ℓ is a pseudo-line with respect to S. Let $S_\ell \subseteq S$ denote the segments that intersect ℓ. Let h_ℓ denote the number of endpoints of S_ℓ. If ℓ is such that $\ell \cap \phi^c(S)$

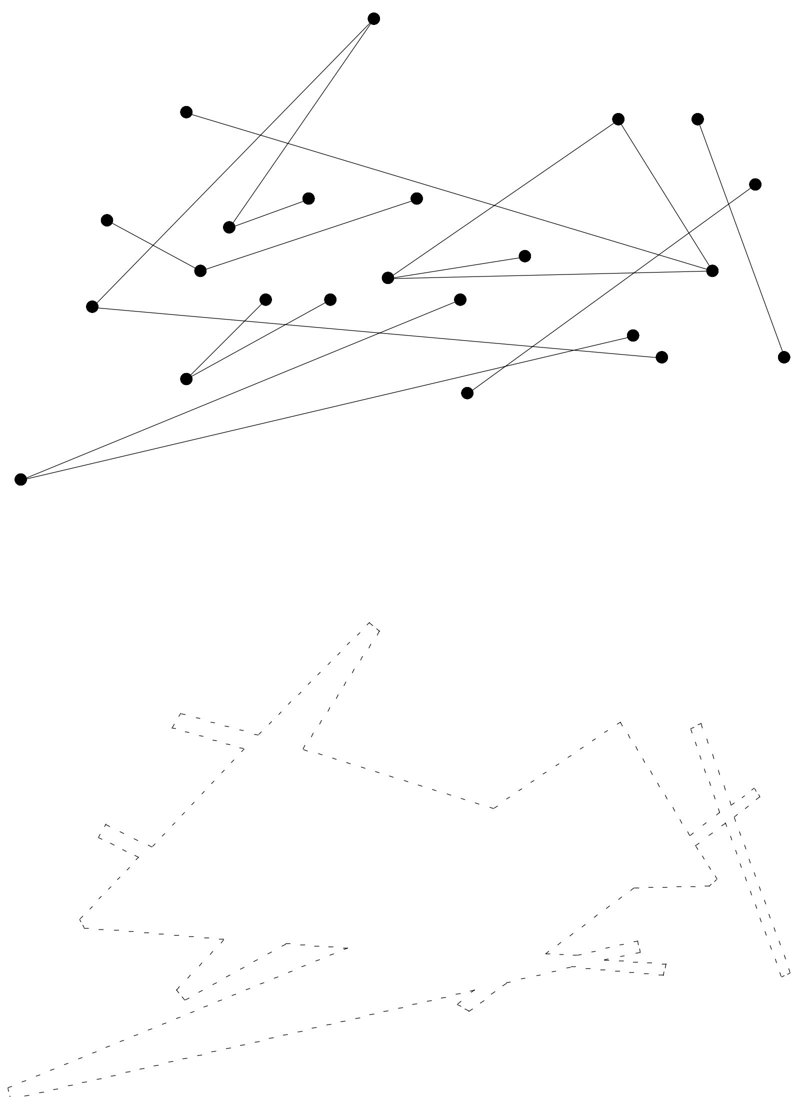

Figure 4: Disposing of degeneracies in γ.

is a connected subcurve, ℓ', of ℓ, then we say that ℓ *favorably stabs* S_ℓ *with respect to* S. If ℓ is such that there are at most $\lfloor h/2 \rfloor$ endpoints strictly to its left (i.e., in $int(L)$) and at most $\lfloor h/2 \rfloor$ endpoints strictly to its right (i.e., in $int(R)$), then we say that ℓ is a *bisector* of the endpoints of S. If, further, ℓ favorably stabs S_ℓ, then we say that ℓ is a *favorable bisector* of the endpoints of S.

Favorable bisectors play a crucial role in our proof of the $O(h \log h)$ bound on the complexity of $\phi(S)$. First, we argue that a favorable bisector always exists, using a method we call "sweeping geodesics":

Lemma 5 *For any set S of line segments in the plane, there exists a favorable bisector of the endpoints of S.*

Proof. Without loss of generality we assume that the boundary γ of $\phi(S)$ is connected, so that γ is the boundary of a simple polygon, P, containing all of the segments S; P is possibly degenerate, in the sense mentioned above. In order to assist in visualizing P as a simple polygon, we can imagine "thickening" the segments of S by a very small $\varepsilon > 0$, so that now each segment is represented by two parallel and very close subsegments; refer to Figure 4. Our discussion and arguments are made, however, with respect to the original (unperturbed) γ and P.

Let β^+ (resp., β^-) denote the point at $y = +\infty$ (resp., $y = -\infty$). Let p_0 denote the leftmost point of γ. (If this is not unique, just pick p_0 to be the topmost leftmost point of γ.) Let $p^+(t)$ (resp., $p^-(t)$) denote the point on γ that is at distance t from p_0, measuring distance along the curve γ going clockwise (resp., counterclockwise) from p_0. Here, $t \in [0, T]$ is a real number between 0 and T, the half-perimeter of $\phi(S)$.

Let ρ denote the leftward ray out of point p_0. Then, $Q = \phi(S) \setminus \rho$ is simply connected and can be treated as a simple polygon (which is unbounded and contains both β^+ and β^-).

Now, for $t \in [0, T]$, define the curve $\ell(t)$ as follows: $\ell(t)$ goes from β^+ to $p^+(t)$ along the (unique) geodesic path within Q, then from $p^+(t)$ to $p^-(t)$ along the (unique) geodesic path within P, and then from $p^-(t)$ to β^- along the (unique) geodesic path within Q. As t varies from 0 to T, the curves $\ell(t)$ sweep across the plane.

Note that each $\ell(t)$ is a simple curve. Further, each $\ell(t)$ is a pseudo-line with respect to S, since $s \cap \ell(t)$ must either be empty or be a subsegment of s. (The geodesic path from $p^+(t)$ to $p^-(t)$ within P cannot intersect s in two or more pieces, since it is a shortest path within P, $s \subset P$, and the shortest path between two points of s must necessarily be a subsegment of s.)

Claim: The set of values of t for which $\ell(t)$ passes through a vertex v (of some segment of S) is an interval, I_v.

Proof. Assume to the contrary that there are 3 values of t — call them t_1, t_2, t_3 — such that $\ell(t_1)$ and $\ell(t_3)$ go through v, but $\ell(t_2)$ does not. Then, the geodesic path within P from $p^+(t_1)$ to $p^-(t_1)$ goes through endpoint v, as does the geodesic path from $p^+(t_3)$ to $p^-(t_3)$, but the geodesic path from $p^+(t_2)$ to $p^-(t_2)$ does not go through v. Now, consider the ("funnel") region C^- enclosed by the Jordan curve that goes from $p^-(t_1)$ to v (along the geodesic path in P), from v to $p^-(t_3)$ (along the geodesic path in P), and from $p^-(t_3)$ to $p^-(t_1)$ (along γ). Similarly, define the region C^+ enclosed by the Jordan curve that goes from $p^+(t_1)$ to v (along the geodesic path in P), from v to $p^+(t_3)$ (along the geodesic path in P), and from $p^+(t_3)$ to $p^+(t_1)$ (along γ). The curve $\ell(t_2)$ enters C^- (resp., C^+) at point $p^-(t_2)$ (resp., $p^+(t_2)$) and must exit it at some point w^- (resp., w^+). Because the geodesic path within P from w^- to w^+ is unique, we must have that $w^- = w^+ = v$, in contradiction to the assumption that $\ell(t_2)$ does not go through v. □

Now, a favorable bisector is given by $\ell(t^*)$, where t^* is a median among the $2h$ left/right endpoints of the intervals I_v. □

Next, we give a precise statement of the sense in which ℓ is "favorable":

Lemma 6 *Let ℓ be a pseudo-line with respect to S. Let $S_\ell \subseteq S$ be the segments intersected by ℓ, and assume that ℓ is a favorable stabber of S_ℓ with respect to S. Let h_ℓ denote the number of endpoints of S_ℓ, and let $\ell' = \ell \cap \phi^c(S)$. Then the segments S_ℓ contribute at most $O(h_\ell)$ pieces to the boundary complexity of $\phi(S_\ell \cup \{\ell'\})$; i.e., the combinatorial complexity of $Q_\ell = S_\ell \cap \partial(\phi(S_\ell \cup \{\ell'\}))$ is $O(h_\ell)$.*

Proof. We apply a Davenport-Schinzel argument as follows. It suffices to show that the number of contributions by segments of S_ℓ to $\phi(S_\ell \cup \{\ell'\})$ in L (to the left of ℓ) is at most $O(h_\ell)$; a symmetric argument applies to contributions within R.

Let v be an endpoint of one or more segments of S_ℓ, with $v \in int(L)$. Let $S_\ell(v)$ denote the segments of S_ℓ that have v as one endpoint. The segments $S_\ell(v)$ form a "fan" at apex v, so there is a well-defined leftmost and rightmost segment of $S_\ell(v)$ when looking from v towards ℓ. We refer to these (at most two) segments of $S_\ell(v)$ as the *extremal* segments with respect to v.

First, we claim that the only segments of $S_\ell(v)$ that have contributions to $\phi(S_\ell \cup \{\ell'\})$ within L are the (at most two) extremal segments with respect to v. This follows from the fact that ℓ is a favorable stabber: Any non-extremal segment of $S_\ell(v)$ is "trapped" from contributing within L by the extremal segments and ℓ'.

Next, we distinguish between two kinds of contributions to $\phi(S_\ell \cup \{\ell'\}) \cap L$: those on the "left" of a segment, and those on the "right" of a segment, where left/right is defined according to when we look along a segment from its endpoint in L towards the pseudo-line ℓ. It suffices to show a bound of $O(h_\ell)$ on the total number of left contributions (the right contributions can be handled symmetrically).

Consider an ordered listing of the left contributions of S_ℓ to $\phi(S_\ell \cup \{\ell'\}) \cap L$, with each contribution labeled by the containing segment. For definiteness, assume that the ordering is counterclockwise about the boundary of $\phi^c(S_\ell \cup \{\ell'\})$. We claim that, in this sequence:

1. *There cannot be two consecutive identical symbols (e.g., s_i, s_i) corresponding to subsegments $[p,p']$ and $[q,q']$ of $s_i = [u_i, w_i]$.*

   Proof. Let us assume that $u_i \in L$, that p (resp., q) is closer to u_i than p' (resp., q'), and that $[p,p']$ is encountered before $[q,q']$ in a counterclockwise traversal of the boundary of $\phi^c(S_\ell \cup \{\ell'\})$. Refer to Figure 5.

   Let λ denote a point interior to $L \cap \phi(S_\ell)$. Now, the set $L \cap \phi(S_\ell)$ is simply connected (since every segment of S_ℓ must intersect $\ell = \partial L$). Thus, there exist paths π_p, from λ to p, and $\pi_{q'}$, from λ to q', such that (1) $\pi_p, \pi_{q'} \subset L \cap \phi(S_\ell)$; (2) $\pi_p \cap \pi_{q'} = \lambda$; and (3) π_p (resp., $\pi_{q'}$) are incident to p (resp., q') *from the left side* of s_i. Thus, the following (oriented) cycle is a closed Jordan curve that surrounds a bounded region, $\mathcal{R}$, on its left: go from λ to q' along $\pi_{q'}$, then from q' to p along s_i, then back to λ along π_p.

   Each of the points p, p', q, q' is an intersection point of s_i with some other segment of S_ℓ. Let $s_j = [u_j, w_j]$ denote the other segment containing p, so that $p = s_i \cap s_j$, and $u_j \in L$. Since s_j is crossed once by ℓ, the endpoint w_j cannot lie within region $\mathcal{R} \subset L$. Thus, $u_j \in \mathcal{R}$,

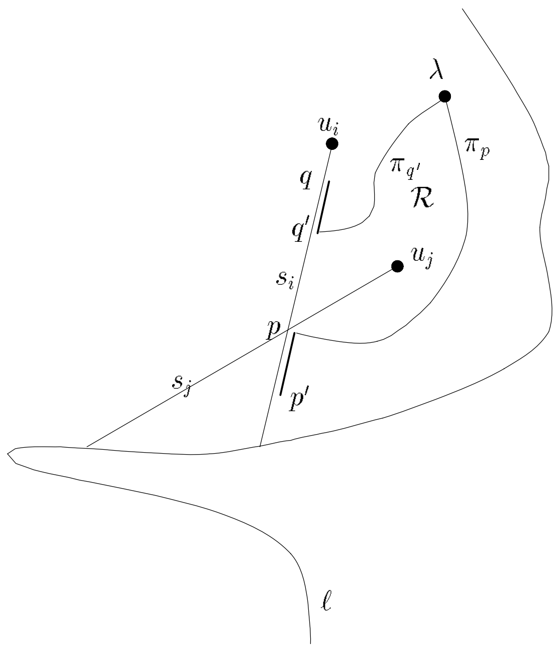

Figure 5: Proof of lemma: There cannot be two consecutive identical symbols "s_i", "s_i".

implying that π_p is incident to s_j at point p *from the left side* of s_j. But this means that there is a symbol "s_j" in the sequence immediately following the symbol "s_i" that corresponds to $[p, p']$. □

2. *There cannot be a subsequence of the form* $s_i, \ldots, s_j, \ldots, s_i, \ldots, s_j$.

   Proof. This follows from the fact that ℓ is a favorable stabber, since the segments s_i and s_j, together with ℓ', trap the left portion of s_j between the crossing point $s_i \cap s_j$ and the point where s_j crosses ℓ. Refer to Figure 6. □

Thus, we conclude that there are at most $2h_\ell$ segments of S_ℓ that contribute to the face $\phi(S_\ell \cup \{\ell'\})$, and that the listing of left contributions made by these segments form a Davenport-Schinzel sequence of order 2. This implies that S_ℓ makes at most $4h_\ell - 1$ left contributions to $\phi(S_\ell \cup \{\ell'\}) \cap L$. □

Now let us see how to put our results together, using the Combination Lemma of [EGS]. Let ℓ be a favorable bisector of the endpoints of S. Let $\ell' = \ell \cap \phi^c(S)$, and let S_L (resp., S_R) denote those segments of S whose endpoints both lie to the left (resp., right) of ℓ. Let $Q_\ell = S_\ell \cap \partial(\phi(S_\ell \cup \{\ell'\}))$ be the portion of the boundary of $\phi(S_\ell \cup \{\ell'\})$ contributed by the segments S_ℓ, and let $\phi(Q_\ell)$ denote the corresponding "face at infinity".

Our goal is to bound the boundary complexity of the region $\phi(S)$. We can think of $\phi(S)$ as a *polygon*, since it is an open face in an arrangement of line segments — such is the definition used in [EGS]. Similarly, the sets $\phi(S_L)$, $\phi(S_R)$, and $\phi(Q_\ell)$ are also polygons. We now claim that

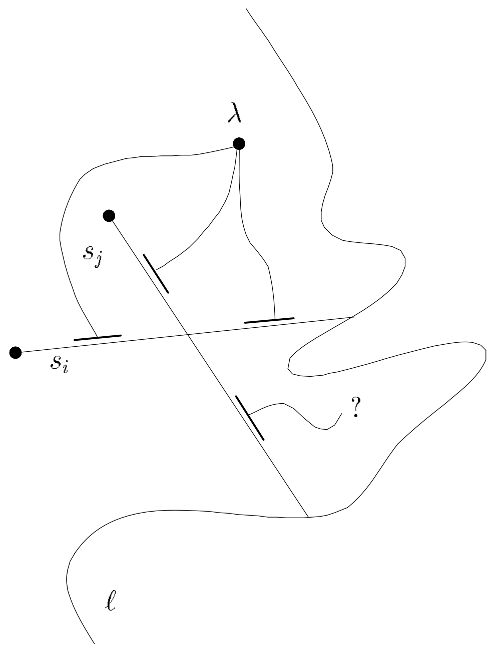

Figure 6: Proof of lemma: There cannot be a subsequence of the form $s_i, \ldots, s_j, \ldots, s_i, \ldots, s_j$.

Lemma 7 *$\phi(S)$ is the face at infinity in the arrangement of polygons $\phi(S_L)$, $\phi(S_R)$, and $\phi(Q_\ell)$.*

Proof. It is clear that $\phi(S)$ is the face at infinity in the arrangement of $\phi(S_L)$, $\phi(S_R)$, and $\phi(S_\ell)$. What we must show is that we are justified in not considering the portions of the boundary of $\phi(S_\ell)$ that are "concealed" by the curve ℓ'. Let x be a point on the boundary of $\phi(S_\ell)$ that is *not* on the boundary of $\phi(Q_\ell)$ (i.e., not in the set Q_ℓ). Then, there is a path, π_x, connecting x to the point at infinity, such that π_x crosses no segment of S_ℓ but π_x does cross ℓ'. Let y denote the first point where π_x crosses ℓ', when going from x to the point at infinity. Since $\ell' \subset \phi^c(S)$, we see that $y \in \phi^c(S)$, so the subpath of π_x that joins y to the point at infinity must cross the boundary of $\phi(S)$ an odd number of times. This implies that π crosses the boundary of $\phi(S)$ an odd number of times, and thus that x is not exposed on the face $\phi(S)$ (i.e., $x \notin \partial(\phi(S))$). □

Now let us apply the Combination Lemma of [EGS]. First, let us look only at the two polygons $\phi(S_L)$ and $\phi(S_R)$, which we think of as being a (single) "red" polygon, R_1, and a (single) "blue" polygon, B_1. (Each polygon is multiply-connected, possibly with many holes.) There is a single point p_1, which is the point at infinity. Following the notation of [EGS], then, $k = s = t = 1$. The Combination Lemma implies that the number of edges of the polygon $R_1 \cap B_1$ is at most $|\phi(S_L)| + |\phi(S_R)| + O(1) + O(r)$, where r is the number of reflex vertices of $\phi(S_L)$ and $\phi(S_R)$. But we know that $r \leq h$, since the only vertices of $\phi(S_L)$ and $\phi(S_R)$ that can be reflex are the endpoints of segments S. Thus, we get

$$|\phi(S_L) \cap \phi(S_R)| \leq |\phi(S_L)| + |\phi(S_R)| + O(h).$$

Now we apply the Combination Lemma again, to combine the polygon $\phi(S_L) \cap \phi(S_R)$ and the polygon $\phi(Q_\ell)$. Again, the total number of reflex vertices is $O(h)$, since Lemma 6 implies that the total number of vertices of $\phi(Q_\ell)$ is only $O(h)$. Thus,

$$|(\phi(S_L) \cap \phi(S_R)) \cap \phi(Q_\ell)| \leq |\phi(S_L) \cap \phi(S_R)| + |\phi(Q_\ell)| + O(1) + O(h) \leq |\phi(S_L)| + |\phi(S_R)| + O(h).$$

This gives us the following recursion for the worst-case complexity, $g(h, n)$, of the complexity of $\phi(S)$:

$$g(h, n) \leq g(h/2, n_L) + g(h/2, n_R) + O(h),$$

for $n_L + n_R \leq n$, and $g(2, 1) = O(1)$.

Solving this recursion gives the desired bound, $|\phi(S)| = O(h \log h)$, and allows us to state our main combinatorial result:

Theorem 8 *The worst-case complexity of a single face in an arrangement of n line segments with a total of h endpoints is $O(h \log h)$.*

4 Conclusion

The general theme underlying our line of research is that the number of endpoints is "more important" than the number of segments when it comes to complexity issues in planar arrangements. Our main result is that a single face in an arrangement of segments in the plane has worst-case complexity $O(h \log h)$, where h is the number of endpoints. Several open questions remain:

(1) The only lower bound we know for the complexity of a single face among line segments is $\Omega(h\alpha(h))$, as reported in [WS]. It would be interesting to see if the $\Omega(h \log h)$ lower bound construction for pseudo-segments (provided to us by A. Lubiw and S. Suri) can be "straightened" to show a similar bound for line segments, thereby closing the gap with our upper bound.

(2) One may be able to get improved *lower* bounds on the complexity of a single face by writing the bounds in terms of parameters other than n or h. For example, can we obtain a non-trivial lower bound on the complexity of a single face in terms of h and k, where $k = n - \binom{h}{2}$ is the number of pairs of endpoints that are *not* joined by an edge (e.g., k is a measure of how much the graph differs from the complete graph)? When $k = 0$, we have seen that the complexity of a single face is bounded by h, which is significantly less than the $O(h \log h)$ upper bound we have shown in general.

(3) We have concentrated on *single-face* results. A natural extension would be to look at the complexity of m faces in an arrangement of n segments with a total of h endpoints.

(4) What can be said in higher dimensions? For example, what is the complexity, $f(n, m, k)$, of a single cell in an arrangement of n triangles in space, where m is the number of edges and k is the number of vertices?

Acknowledgement

Christine Piatko has played a vital role in this research, at all stages. We particiularly thank her for pointing out several errors in earlier drafts. We are also grateful to A. Lubiw and J. Spinrad for pointing out an error in an earlier draft. S. Suri helped in providing an example of a set of pseudo-segments whose lower envelope complexity is $\Omega(h \log h)$.

References

[ABP] 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), pp. 573-590.

[AHKMN] E.M. Arkin, D. Halperin, K. Kedem, J.S.B. Mitchell, and N. Naor, "Arrangements of Segments that Share Endpoints: Single Face Results", *Proc. Seventh Annual ACM Symposium on Computational Geometry*, North Conway, NH, June 10–12, 1991, pp. 324-333.

[MN] J.S.B. Mitchell, and N. Naor, "Combinatorial Bounds for Single Faces in Arrangements of Pseduo-Segments and Chords in Polygons", Manuscript, 1991.

[At] M. Atallah, "Some Dynamic Computational Geometry Problems", *Comput. Math. Appl.* **11** (12) (1985), pp. 1171-1181.

[CEGSS] B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir, and J. Snoeyink, "Computing a Face in an Arrangement of Line Segments", *Proc. Second Annual ACM-SIAM Symposium on Discrete Algorithms*, San Francisco, January 28-30, 1990, pp. 441-448.

[EGS] H. Edelsbrunner, L.J. Guibas, and M. Sharir, "The Complexity and Construction of Many Faces in Arrangements of Lines and of Segments", *Discrete and Computational Geometry* **5** (1990), pp. 161-196.

[HS] S. Hart and M. Sharir, "Nonlinearity of Davenport-Schinzel Sequences and of Generalized Path Compression Schemes", *Combinatorica* **6** (1986), pp. 151-177.

[He] J. Hershberger, "Finding the Upper Envelope of n Line Segments in $O(n \log n)$ Time", *Information Processing Letters* **33** (1989), pp. 169-174.

[Lo] L. Lovász, *Combinatorial Problems and Exercises*, North-Holland, Amsterdam, New York, Oxford, 1979.

[Mi] J.S.B. Mitchell, "On Computing a Single Face in an Arrangement of Line Segments", Manuscript, July, 1990.

[PSS] R. Pollack, M. Sharir, and S. Sifrony, "Separating Two Simple Polygons by a Sequence of Translations", *Discrete and Computational Geometry* **3** (1988), pp. 123-136.

[PS] F.P. Preparata and M.I. Shamos, *Computational Geometry*, Springer-Verlag, New York, Berlin, Heidelberg, Tokyo, 1985.

[Sh] P.W. Shor, "Geometric Realizations of Superlinear Davenport-Schinzel Sequences (Part 1: Line Segments)", Manuscript, AT& T Bell Laboratories, 1990.

[WS] A. Wiernik and M. Sharir, "Planar Realization of Nonlinear Davenport-Schinzel Sequences by Segments", *Discrete and Computational Geometry* **3** (1988), pp. 15-47.