Determining Sector Visibility of a Polygon

Binay K. Bhattacharya
School of Computing Science,
Simon Fraser University

David G. Kirkpatrick
Department of Computer Science,
University of British Columbia

Godfried T. Toussaint
School of Computer Science,
McGill University

We consider a generalization of several notions of external visibility of simple polygons, namely weak external visibility, weak external visibility from a line and monotonicity, that we call *sector* visibility. Informally, sector visibility addresses the question of external visibility along rays (or sight lines) whose angles are restricted to a sector (wedge) of specified width σ. This provides an interesting measure of the degree of external visibility of a polygon. Our framework also permits a unification and extension of a number of previously unrelated results. Finally, our results uncover a curious complexity discontinuity in this family of problems: algorithms are $\Theta(n)$ when $\sigma \leq \pi$ or $\sigma = 2\pi$, but require $\Omega(n \log n)$ time (at least), when $\pi < \sigma < 2\pi$.

1. Introduction

Any sequence of n points $p_1,...,p_n$ in the Euclidean plane E^2 defines a *polygonal chain* $C[p_1,...,p_n]$ whose *vertices* are the points $p_1,...,p_n$ and whose *edges* are the finite line segments $[p_i,p_{i+1}]$, i = 1,...,n-1. A polygonal chain $C[p_1,...,p_{n+1}]$ with $p_1 = p_{n+1}$ is called a *polygon* (or *n-gon*).

Semi-infinite line segments are referred to as *rays*. We denote by *ray*(x, ψ) the ray with endpoint x and direction ψ. The ray r = *ray*(x, ψ) is said to *support* polygon P at x if $r \cap P = \{x\}$.

A polygonal chain is *simple* if no non-consecutive pair of its edges intersect. A simple polygon P has a well defined (bounded) interior (denoted by *int*(P)) and (unbounded) exterior (denoted by *ext*(P)). We denote by P^* the union of P and *int*(P). We assume that the point sequence defining a given simple polygon P satisfies the property that each directed line segment $[p_i,p_{i+1}]$ has the interior of P to its left. Hereafter, polygonal chains (including polygons) will be assumed to be simple.

Two points x and y are said to be *visible* (with respect to a polygon P) if the interior of the line segment $[x,y]$ lies either completely in *int*(P) or completely in *ext*(P). If *int*($[x,y]$) $\subseteq$ *int*(P) (respectively, *int*($[x,y]$) $\subseteq$ *ext*(P)) then x and y are said to be *internally* (respectively *externally*) *visible* (with respect to P).

A point set T is said to be *weakly visible* from a point set S if, for each point $p \in T$, there exists a point $q \in S$ such that p and q are visible. The notion of weak visibility has received attention in both the mathematics and computer science literature. Horn and Valentine [HV] have characterized L-sets in terms of weak visibility properties while such characterizations for convex, star-shaped and other sets have been presented by Bezdek, Bezdek and Bisztriczy [BBB] and Shermer and Toussaint [ST]. A polygon P is said to be an L-set provided that for every pair of points $x, y \in P$, there exists a third point $z \in P$ (possibly dependent on x and y) such that both x and y are visible from x. Avis and Toussaint [AT] showed that given a polygon P and a specified e of P, whether P^* is weakly visible from e can be determined in O(n) time. A more difficult problem is to determine whether there exists an edge of P from which P is weakly visible. Clearly, by applying the algorithm in [AT] to each edge in turn the latter problem can be solved in $O(n^2)$ time. Subsequently Sack and Suri [SS] discovered a linear-time algorithm for determining all (if any) such edges of a given polygon. Recently, Yan Ke [Ke] considered the problem of detecting the weak visibility of a polygon from an internal line segment. He presents an $O(n \log n)$ time algorithm that tests if a polygon is weakly visible from some internal line segment and reports such a line segment if it exists. He also shows that the shortest such segment can be found in $O(n \log n)$ time. Finally he addresses the query version of this problem: given a query line segment S in P, is P weakly visible from S? he shows that this question can be answered in $O(\log n)$ time after the polygon is preprocessed in $O(n \log n)$ time using $O(n)$ space.

In this paper we focus on weak *external* visibility of a polygon. This topic is as yet quite unexplored compared to its internal counterpart. Toussaint and Avis [TA] considered the problem of determining if a polygon is weakly externally visible. (Since we will restrict ourselves hereafter to notions of visibility that are both weak and external we will drop these adjectives for the sake of less cumbersome terminology). A polygon P is *edge-visible* if for each point $x \in P$ there exists a ray that supports P at x. This is equivalent to saying that P is visible from a circle at infinity (or, in fact, any circle that properly encloses P). Toussaint and Avis [TA], using related results of [AT], show that edge-visibility of polygons can be recognized in $O(n)$ time. This result is proved by showing that the edge-visibility problem is equivalent to the somewhat less constrained *vertex-visibility problem*: determine, for each *vertex* $v \in P$, if there exists a ray that supports P at v.

The notion of monotonicity, which enjoys numerous applications [PSh], [TE], [CRS] can also be cast as a kind of external visibility problem. A polygonal chain C is said to be *monotone* with respect to a line L if every line orthogonal to L intersects C in at most one point. Equivalently, C is weakly visible from one point on the circle at infinity (defined by the family of sight-lines orthogonal to L). A polygon is monotone with respect to a line L if it can be decomposed into two chains each of which is monotone with respect to L. Preparata and Supowit [PSu] show that monotonicity of a polygon, in fact a description of *all* directions of monotonicity, can be determined in $O(n)$ time.

Intermediate to the notions of edge-visibility and monotonicity is the notion of edge-visibility from a line, the study of which was the starting point for the research presented here. A polygon P is *edge-visible from a line* if there exists a line L in $ext(P)$ such that P is edge-visible from L. (Equivalently, P is edge-visible from a semicircle at infinity, whose points correspond to sight lines in an interval bounded by the two orientations of L). Fig. 1 illustrates a polygon P edge-visible from a

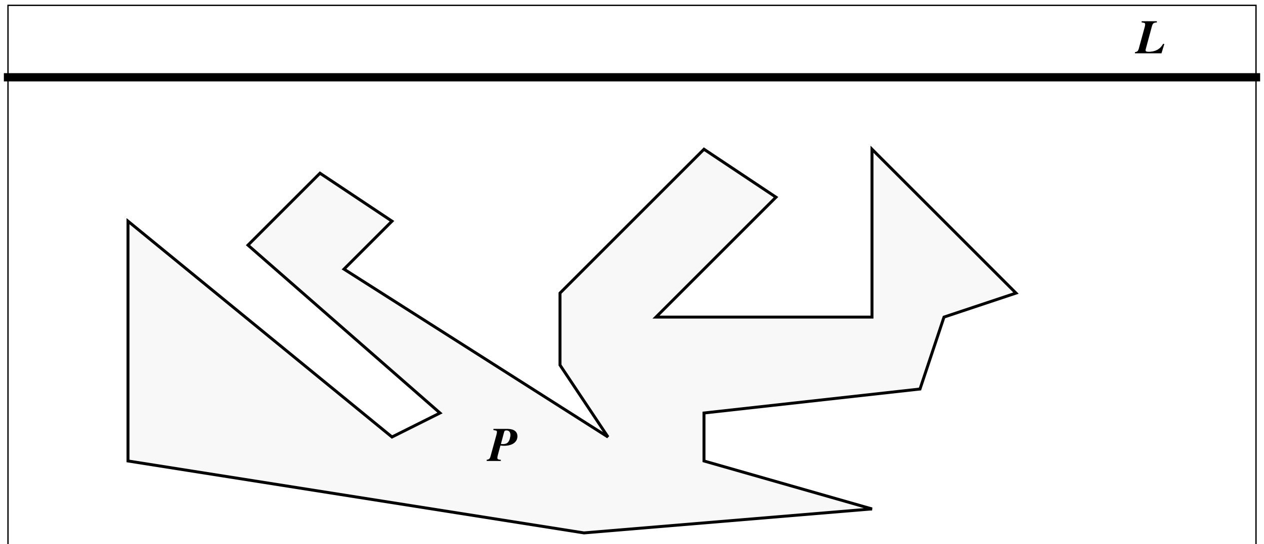

Fig. 1: A polygon *P* (weak external) *edge-visible* from a line *L*.

line *L*.

The above notions of external visibility have a natural unification and generalization. We refer to arbitrary angles as *sight-lines*. An open fixed angular interval *W* of sight-lines ψ satisfying $\psi^B < \psi < \psi^F$ and denoted by (ψ^B,ψ^F) is referred to as a (visibility) *wedge*; closed wedges are defined similarly. We denote by $|W|$ the (angular) *width* of W, namely $\psi^F - \psi^B$. A polygonal chain *C* is said to be *W-edge-visible* (respectively, *W-vertex-visible*) for a given *W,* if for every point (respectively, vertex) *x* of *C* there exists a $\psi \in W$ such that *ray*(x,ψ) supports *C* at *x*. Furthermore, *C* is said to be σ–*sector-(edge/vertex)-visible* if there exists a wedge *W* of width σ such that *C* is *W*-(edge/vertex)-visible. It should be clear from the discussion above that (weak external) edge-visibility corresponds to 2π-sector-edge-visibility, edge-visibility from a line corresponds to π–sector-edge-visibility, and monotonicity (of chains) corresponds to ε-sector-edge-visibility, for all sufficiently small $\varepsilon > 0$.

We should add here that since the completion of the work presented here, the problem of determining the shortest line segment from which a polygon is weakly externally visible has been solved in O(n) time by Bhattacharya, Mukhopadhyay and Toussaint [BMT]. The algorithm in [BMT] can also be used to solve the 2π-sector-edge-visibility problem considered here.

Given the framework described above, a number of natural questions arise:

1. [Wedge (edge/vertex)-visibility problem]
   Given a polygon *P* and a wedge *W* of sight lines, determine whether *P* is *W*-(edge/vertex)-visible.

2. [Sector (edge/vertex)-visibility problem]

Given a polygon P and an angle σ, $0 \leq \sigma < 2\pi$, determine whether P is σ-sector-(edge/vertex)-visible, and if so describe all wedges W that realize this sector visibility.

3. [Minimum sector (edge/vertex)-visibility problem]

Given a polygon P, determine the minimum σ for which P is σ-sector-(edge/vertex)-visible.

We show that the inherent (worst case) complexity of answering the sector (edge/vertex)-visibility problem exhibits a curious discontinuity. When $\sigma \leq \pi$ or $\sigma = 2\pi$ the complexity is $\Theta(n)$, yet when $\pi < \sigma < 2\pi$ it has an $\Omega(n \log n)$ lower bound. Furthermore, when $\sigma \leq \pi$, in at most $O(n)$ additional time a linear size description of *all* wedges realizing the specified sector visibility can be constructed, that permits wedge visibility queries to be answered in $O(\log n)$ time per query. The minimum sector (edge/vertex)-visibility problem inherits the same complexity bounds; it has a $\Theta(n)$ solution when the minimum is at most π and an $\Omega(n \log n)$ lower bound otherwise.

2. Determining wedges of support

If $P = C[p_1,...,p_{n+1}]$ is a polygon and x is any point of P we define *angle*(x) to be the external angle of P at point x. (In particular, *angle*(x) = π for all points of P which are not vertices.). The *local wedge of support of P at point x*, denoted $W_p^*(x)$, is given by $(\theta^B(x), \theta^F(x))$ where $\theta^B(p_1)$ is the angle in $[0,2\pi]$ formed by the ray with endpoint p_1 passing through p_{n-1}, $\theta^B(p_i) = \theta^B(p_{i-1}) + angle(p_{i-1}) - \pi$, for $i > 1$, $\theta^B(x) = \theta^B(p_i)$, if $x \in int([p_{i-1},p_i])$, and $\theta^F(x) = \theta^B(x) + angle(x)$, for all x in P.

Note that by this definition the local wedge of support of a point is dependent on the choice of initial vertex p_1. The redundancy evident in the representation of angles, though hard to motivate here, is exploited in subsequent algorithms. This redundancy is limited, however, by the fact that polygons that are vertex-visible cannot spiral too much. The following lemma quantifies this redundancy.

Lemma 1. If P is vertex-visible then $W_p^*(v) \subseteq (-5\pi, 5\pi)$, for all vertices v of P.

Proof. Suppose $\theta^B(p_i) < 5\pi$. Then it is straightforward to show that P does not admit a supporting ray at either p_1 or p_{i+1}. The argument when $\theta^F(p_i) > 5\pi$ is identical.

Q.E.D.

If P is a polygon and x is any point of P, the *global wedge of support of P at x*, denoted $W_p(x)$ (or simply $W(x)$ when P is understood), is the set of all angles $\psi \in W_p^*(x)$ such that *ray*(x,ψ) supports

P at x. Let $W(x) = (\psi^B(x),\psi^F(x))$. By the maximality of $W(x)$ it follows that both $ray(x,\psi^B(x))$ and $ray(x,\psi^F(x))$ intersect P - {x}. Let $t^B(x)$ (respectively, $t^F(x)$) be referred to as the *back* (respectively, *forward*) *tangent point* from x (see Figure 2). With this notation, it is clear that P is W-edge-visible if and only if for every point $x \in P$ there is a $\psi \in W_p(x)$ and a $\psi' \in W$ such that $\psi \equiv \psi'$ (mod 2π). If this is the case, we say that W *spans* the collection $\{W_p(x) \mid x \in P\}$.

In this section, we consider the efficient computation of $W_p(x)$ for all points x of a given polygon P. We assume without loss of generality that P is edge-visible; in fact, if this is not the case, it will be detected as part of the algorithm.

It will suffice to show how to determine $t^B(x)$ for all points x of P; a symmetric algorithm can be used to construct $t^F(x)$, and hence complete the determination of $W_p(x)$. The algorithm proceeds by refining P, through the addition of new vertices on some of its edges, and determining $t^B(x)$ for each vertex v of this refined polygon. The new vertices are chosen in such a way that for an arbitrary non-vertex point x on P, $t^B(x) = t^B(v)$, where v is the vertex following x on the refined chain.

The algorithm is most easily described as a simple modification of the on-line convex hull algorithm for simple polygonal chains due to Melkman [M]. Let $P = C[p_1,...,p_{n+1}]$ be a simple polygon. Suppose, without loss of generality, that p_1 is a vertex of the convex hull of P. If x_1,x_2 and x_3 ar three points then the function $side(x_1,x_2,x_3)$ takes the value 1, 0, or -1 depending on whether x_3 is to the right of, collinear with, or to the left of the line through x_1 and x_2 and directed from x_1 to x_2. The algorithm maintains a stack S of points (initially empty). The operations *push* and *pop* modify S in the obvious way. The variable *top* refers to the top element of S.

Algorithm back-tangents

begin
push p_1
$t^B(p_2) \leftarrow p_1$
$j \leftarrow 3$
while $j \leq n$ *do*
    if side (top, p_{j-1}, p_j) ≤ 0
        then if side (p_{j-2}, p_{j-1}, p_j) ≥ 0
                then HALT {p_{j-1} is not weakly externally visible}
                else $t^B(p_j) \leftarrow p_{j-1}$
                    push p_{j-1}
                    $j \leftarrow j + 1$

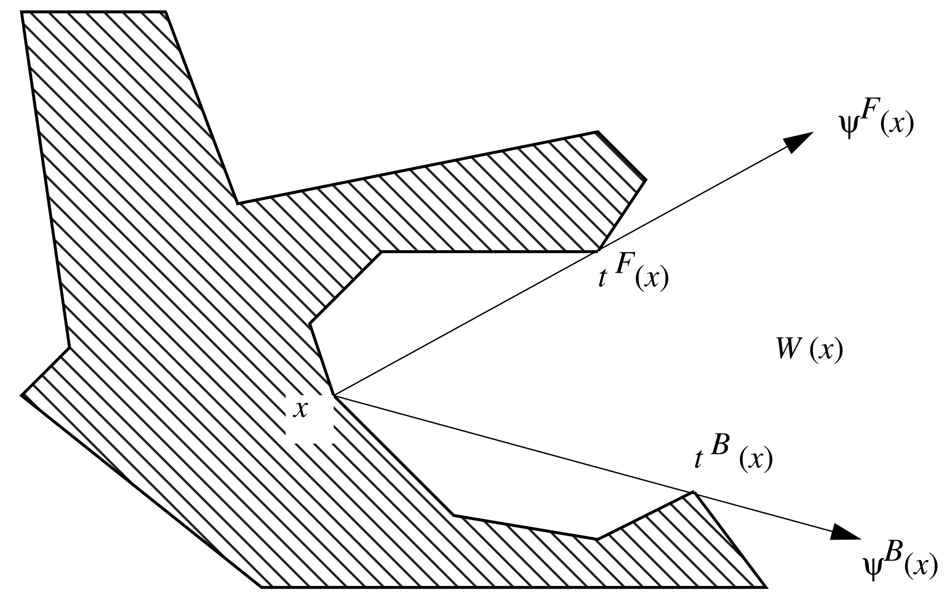

Figure 2

else $z \leftarrow$ top
    pop S
    while $S \neq \phi$ and side (top, z, p_j) > 0
        $w \leftarrow$ intersect (line(top, z); line($p_{j\text{-}1}$,p_j))
        insert w on $[p_{j\text{-}1},p_j]$
        $t^B(w) \leftarrow z$
        $z \leftarrow$ top
        pop S
    $t^B(p_j) \leftarrow z$
    push z
    $j \leftarrow j + 1$
end

The correctness of Algorithm Back-Tangents follows from a straightforward case analysis similar to that of Melkman [M], together with the invariant that the elements of S followed by $p_{j\text{-}1}$ describe the convex hull of the polygonal chain $C[p_1,...,p_{n\text{-}1}]$. It is also straightforward to confirm that the algorithm runs in O(n) steps and inserts O(n) new vertices into the edges of C. We summarize the result of this section with the following theorem.

Theorem 1: Given a polygon P, linear time suffices to construct a refinement $P\,'$ of P with the property that given any point x of $P\,'$ and its associated edge, the wedge of support of P at x can be determined in O(1) time.

It may be suspected that to determine the sector edge-visibility of a given polygon it suffices to determine its sector vertex-visibility. In fact, this is the case for both $\varepsilon-$ and 2π-sector visibility [TA], [PSu]. This is not true in general, however. Nevertheless, when $\sigma \le \pi$, we can reduce the sector edge-visibility problem to a closely related sector vertex-visibility problem.

Lemma 2: Let P be any polygon and let W be any wedge of sight lines of width at most π, then P is W-edge-visible if and only if the polygon P'', formed from P by subdividing each of its edges, is W-vertex-visible.

Proof: It suffices to observe that if both of the endpoints of some edge of P'' are W-visible then so is the entire edge. We know, from [AT], that if the endpoints of any edge e of P'' are W-visible then e is edge-visible.

If supporting rays from W at each of the endpoints of e diverge then at least one of these must support all of the points of e (here we use the fact that one of the two end points of e must be a subdivision point). Alternatively, it is easy to see that e is W'-edge-visible for some wedge W' bounded by two rays of W and satisfying $|W'| < \pi$. Since the width of W is at most π, it follows that $W' \subset W$, and hence e is W-edge-visible.

Q.E.D.

Note that lemma 2 does not hold for wedges of width greater than π. Figure 3 illustrates a polygon P and a visibility wedge W such that each vertex of P, including v, is W-visible, and yet the shaded edge is not W-visible.

3. The sector edge-visibility problem

In section 1 we introduced the sector edge- and vertex-visibility problems. In section 2 we showed that, when $\sigma \le \pi$, sector edge-visibility reduces to sector vertex-visibility. In this section we focus on the problem of sector vertex-visibility. Before addressing the general case, it is instructive to review the case where $\sigma = \pi$, what we originally called weak visibility from a line. The problem of sector vertex-visibility in this case can be interpreted as a transversal problem; specifically does the collection $\{W_p(v) \mid v \in V\}$ (where V denotes the set of vertices of P and each $W_p(v)$ is now viewed as a sector of the plane) admit a common transversal. In general, a family F of subsets of the plane is said to admit a *common transversal* if there exists a straight line L which intersects every member of F.

Common transversals for families of convex sets have been investigated for some time in both the mathematics [Gr], [Le] and computer science [AB], [AW1], [AW2], [Ed], literatures. In the latter, the more aggressive term *stabber* is more often used for *transversal*. Transversals in the plane find application in several areas including line-fitting [O´R] and updating triangulations [ET]. Edelsbrunner, Overmars and Wood [EOW] develop a method for planar visibility problems that yields a procedure for computing transversals for F, a family of *simple objects*, in $O(n^2 \log n)$ time, where n is the cardinality of F. By simple objects, it is meant those objects that have an $O(1)$ storage description each, and which are such that, for every pair of such objects, constant time suffices to compute basic properties such as their intersection and common tangents. O(n log n) time is sufficient for the special cases of vertical line segments [O´R], for line segments with arbitrary directions [EMPRWW], and for a set of *n translates* of a simple object in the plane [Ed]. Finally, for a set of *isothetic* rectangles O(n) time suffices via linear programming [Ed].

Even more closely related to the topic of this paper is the problem of computing *shortest* transversals of sets, when a transversal exists. In [BCETSU] O(n log n) time algorithms are given for computing a shortest transversal for a family of n lines, a family of n *l*ine segments and a family of convex polygons with a total of n vertices. The algorithms are optimal for the latter two families of objects.

Given a family F of n convex cones, as in the visibility problem considered in this paper, determining whether F admits a common transversal could certainly be accomplished in O(n^2 log n) time with the procedure of [EOW] or in O(n log n $\alpha(n)$) time, where $\alpha(n)$ is the extremely slowly growing inverse Ackermann's function, with the more recent technique of Atallah and Bajaj [AB]. We now show that the structure in our family F, namely the fact that our convex cones are not arbitrary and independent but with their apexes anchored on the vertices of a simple polygon, allows us to solve this transversal problem in O(n log n) time. In the next section, this result is improved to O(n) time.

In the remainder of this section, we show that the problem of finding a wedge, of width at most $\sigma < 2\pi$, that spans a collection of n wedges, has inherent worst-case complexity $\Theta(n \log n)$. We introduce a dual wedge cover problem that simplifies some of the arguments. Let $W = \{W_1,..., W_n\}$ be a set of wedges. The set W is said to *cover* the plane if for every angle ψ, $0 \le \psi < 2\pi$, there exists a wedge $W_i \in W$ such that $\psi \in W_i$. If $W = (\psi^B, \psi^F)$ is an (open) wedge then W^σ denotes the (closed) wedge $[\psi^F - 2\pi + \sigma,\ \psi^B]$. W^σ, which we call the σ-dual of wedge W, can be viewed as a generalized complement of wedge W. Furthermore,

Lemma 3: The set $W = \{W_1,..., W_n\}$ admits a spanning wedge of width σ if and only if the set $W^\sigma = \{W_1^\sigma,..., W_n^\sigma\}$ does not cover the plane.

Proof: This follows immediately from the observation that the wedge $(\psi, \psi + \sigma)$ spans W if and only if $\psi \notin \cup W_i^\sigma$.

Q.E.D.

Lemma 4. The plane covering problem for wedges has worst case time complexity $\Theta(n \log n)$.

Proof. An O(n log n) solution follows by simply lexicographically sorting the wedges (viewed as ordered pairs) and scanning the resulting list. The $\Omega(n \log n)$ lower bound, which says, in effect, that this sorting step is unavoidable in general, holds for arbitrary fixed order algebraic decision trees [B-O]. Ben-Or [B-O] shows that determining if a set $A = \{a_1,..., a_n\}$ is identical to the set $B = \{b_1,..., b_n\}$ requires $\Omega(n \log n)$ time on this model. This set equality problem can be reduced to

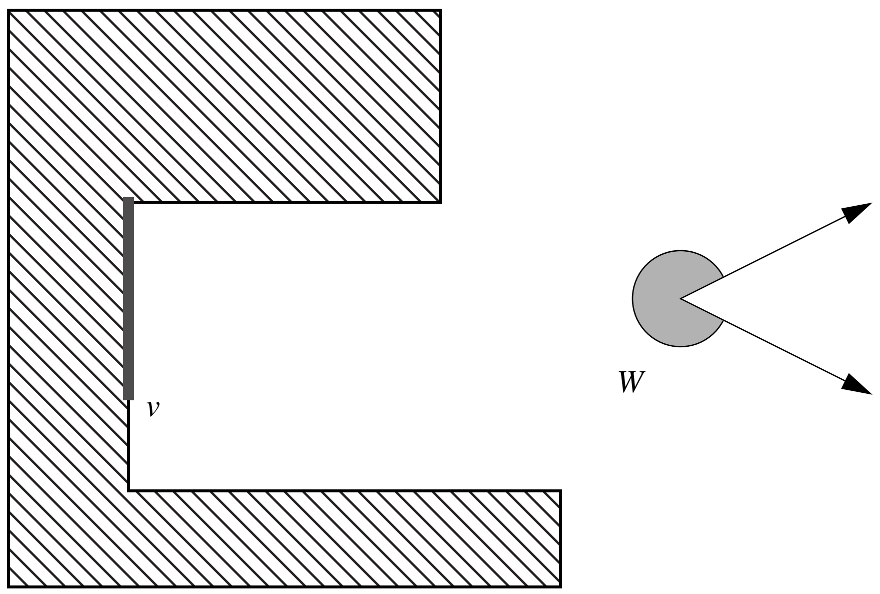

Figure 3

the plane covering problem by setting $W_i = [2\pi a_i/n, 2\pi(a_i+1)/n]$, for $1 \le i \le n$. Note that this is closely related to the so-called measure problem also discussed by Ben-Or.

Q.E.D.

The following theorem which summarizes the main result of this section is an immediate conesquence of the above lemmas.

Theorem 2: The sector vertex-visibility problem has complexity O(n log n).

In fact lemma 4 tells us something stronger than Theorem 2. In particular, if we have any hope of achieving an O(n log n) bound on the complexity of sector visibility problems, it must come as a result of exploiting the structure imposed by the underlying polygon on the set of visibility wedges associated with its vertices. Of course, this is precisely what makes problems for polygonal chains less complex than their unstructured counterparts in general. We pursue this idea in the next section.

4. Sector visibility when $\sigma \le \pi$

We have seen that sector edge-visibility reduces to sector vertex-visibility when $\sigma \le \pi$. Curiously, it is in precisely this situation that sector vertex-visibility itself exhibits a demonstrable complexity discontinuity. The purpose of this section is to substantiate this claim.

Superficial examination of the global wedges of support associated with the vertices of arbitrary simple polygon, reveals little apparent structure. For example, the wedges of adjacent vertices can intersect in an arbitrary fashion. It turns out that the useful structure is most easily seen by examining the dual wedges.

Recall that the global wedge of support of polygon P at vertex p_i, $W(p_i)$, is the interval $(\psi^B(p_i), \psi^F(p_i))$. Its σ-dual $W^\sigma(p_i)$ is the interval $[\psi^F(p_i) - 2\pi + \sigma,\ \psi^B(p_i)]$.

Lemma 5. If $j > i$ then $\psi^F(p_i) - \pi < \psi^B(p_j)$.

Proof. Let ς denote the direction of the ray from p_i through p_j. Then the existence of a chain from p_i to p_j in P ensures that $\psi^F(p_i) < \varsigma$ and $\psi^B(p_j) > \varsigma - \pi$.

Q.E.D.

Corollary. If $j > 1$ then $W^\sigma(p_j)$ either intersects $W^\sigma(p_i)$ or it contains angles strictly larger than those of $W^\sigma(p_i)$.

It follows from the corollary above that we can maintain $\cup\, W^\sigma(p_i)$ as a stack S of disjoint wedges where the angles of successive wedges strictly increase. this construction is made precise in the following algorithm.

Algorithm combine-dual-wedges

begin
$[\phi^B, \phi^F] \leftarrow [\psi^F(p_1) - 2\pi + \delta,\ \psi^B(p_1)]$
$j \leftarrow 2$
while $j \leq n$ *do*
    $[\varsigma^B, \varsigma^F] \leftarrow [\psi^F(p_j) - 2\pi + \delta,\ \psi^B(p_j)]$
    if $[\phi^B, \phi^F] \cap [\varsigma^B, \varsigma^F] \neq 0$
        then $\phi^B \leftarrow \mathbf{min}\ \{\phi^B, \varsigma^B\}$
            $\phi^F \leftarrow \mathbf{max}\ \{\phi^F, \varsigma^F\}$
            while $S \neq \phi$ and top $\cap\ [\phi^B, \phi^F] \neq 0$
                $[\varsigma^B, \varsigma^F] \leftarrow$ *top*
                $\phi^B \leftarrow \mathbf{min}\ \{\phi^B, \varsigma^B\}$
                $\phi^F \leftarrow \mathbf{max}\ \{\phi^F, \varsigma^F\}$

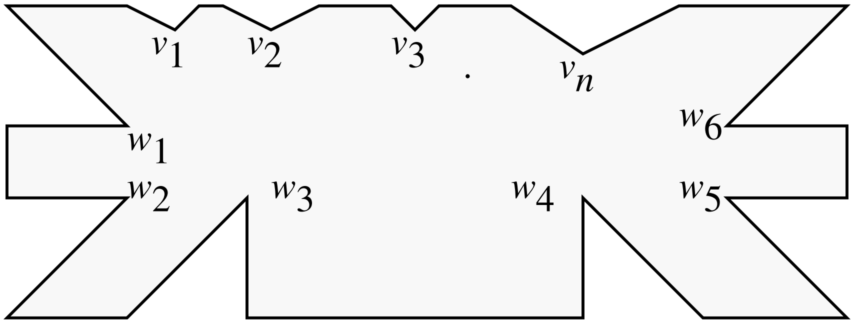

Figure 4

pop S

else push $[\phi^B, \phi^F]$

$\phi^B \leftarrow \varsigma^B$

$\phi^F \leftarrow \varsigma^F$

end

We know from lemma 1 that the angles in the wedges remaining on the stack containing S at the completion of algorithm Combine-Dual-Wedges lie in the interval $[-5\pi, 5\pi]$. It is now straightforward to complete the union of the dual wedges with all angles now reduced mod 2π. With this reduction, S partitions into O(1) ordered lists of intervals, which can be merged in O(n) time. Together with lemma 2 and lemma 3, this completes the proof of the following:

Theorem 3: The sector edge-visibility and the sector vertex-visibility problems both have complexity O(n), when $\sigma \leq \pi$.

It is immediate from their definition that dual wedges decrease linearly in width with increasing σ. This permits us to solve the minimum sector vertex-visibility problem by first using the above algorithm to check if P is π-sector-vertex-visible. If this is so, then the union of the dual wedges on the stack at completion fail to cover the plane. If the maximal uncovered wedge (which can be constructed in O(n) time by a simple scan) has width γ, then it is easy to see that P has minimum sector vertex-visibility $\pi - \gamma$. Hence we have,

Theorem 4: The minimum sector edge-visibility and the minimum sector vertex-visibility problems both have complexity $O(n)$, when the minimum σ is at most π.

It remains now to show that the condition $\sigma \leq \pi$ in the two preceding theorems cannot be relaxed. This is an immediate consequence of the following,

Theorem 5: The sector vertex-visibility problem requires $\Omega(n \log n)$ time, when $\pi < \sigma < 2\pi$.

Proof: As in lemma 4, we prove the lower bound by reduction from the set equality problem. Hence, our lower bound holds for arbitrary fixed order algebraic decision trees.

For simplicity, we will describe the reduction when $\sigma = 3\pi/2$; the generalization should be clear. Let $a_i \in \{1,..., n\}$. Consider the polygon P with $3n + 18$ vertices illustrated schematically in Figure 4. Vertex v_i, $1 \leq i \leq n$ by construction, has $W(v_i) = \frac{a_i\pi}{2n}, \frac{(a_i - 1)\pi}{2n} + \frac{\pi}{2}$.

Note that each such wedge $W(v_i) \subset (0, \pi)$ and thus constitutes a notch in the upper edge of polygon P. The dual wedges associated with vertices $w_1,..., w_6$ cover the entire plane except for the wedge $[0, \pi/2]$. Thus P is σ–sector-vertex-visible if and only if the dual wedges associated with vertices $v_1,..., v_n$ do not cover the wedge $[0, \pi/2]$. But, by construction, this holds precisely when $\{a_1,..., a_n\} = \{1,..., n\}$.

Q.E.D.

5. Concluding remarks

Sector visibility problems constitute what may be considered the easiest external visibility problems. Though we have characterized the asymptotic complexity of many of these exactly, some questions from within this family remain incompletely resolved. For example, what is the complexity of sector edge-visibility when $\pi < \sigma < 2\pi$?

Acknowledgment

The research reported here was supported in part by the National Science and Engineering Research Council. It results from collaboration during the third author's tenure as a British Columbia Advanced Systems Institute Visiting Fellow at Simon Fraser University. A preliminary version of this paper was presented at the *Fifth ACM Symposium on Computational Geometry* in Saarbrucken in 1989.

6. References

[AB] Atallah, M. and Bajaj, C., "Efficient algorithms for common transversals,", *Information Processing Letters*, Vol. 25, May 1987, pp. 87-91.

[AW1] Avis, D. and Wenger, R., "Algorithms for line stabbers in space,", *Proc. 3rd ACM Symposium on Computational Geometry*, 1987, pp. 300-307.

[AW2] Avis, D. and Wenger, R., "Polyhedral line transversals in space,", *Discrete and Computational Geometry*, to appear.

[AT] Avis, A. and Toussaint, G.T., "An optimal algorithm for determining the visibility of a polygon from an edge,", *IEEE Transactions on Computers*, Vol. C-30, No. 12, December 1981, pp. 910-914.

[BBB] Bezdek, A., Bezdek, K. and Bisztriczy, T., "On illumination in the plane by line segments," *Geometria Dedicata*, vol. 41, 1992, pp. 39-50.

[BCESTU] Bhattacharya, B. Czyzowicz, J., Egyed, P., Stojmenovic, I., Toussaint, G. T. and Urrutia, J., "Computing shortest transversals of sets," *International Journal of Computational Geometry and Applications*, vol. 2, No. 4, December 1992, pp. 417-43.

[BMT] Bhattacharya, B. K., Mukhopadhyay, A. and Toussaint, G. T., "A linear-time algorithm for computing the shortest line segment from which a polygon is weakly externally visible," in *Proc. Workshop on Algorithms and Data Structures*, F. Dehne, J. R. Sack and N. Santoro, (Eds.), Ottawa, Canada, August 14-16, 1991, pp. 412-42

[B-O] Ben-Or, M., "Lower bounds for algebraic computation trees,", *Proc. 15th Annual ACM Symposium on Theory of Computing*, 1983, pp. 80-86.

[BT] Bhattacharya, B. K. and Toussaint, G. T., "Computing shortest transversals," *Proc. ICALP 91*, Madrid, Spain, July 8-12, 1991, pp. 649-660.

[CRS] Chandru, V., Rajan, V. T. and Swaminathan, R., "Monotone pieces of chains," *ORSA Journal of Computing*, vol. 4, No. 4, Fall 1992, pp. 439-446.

[Ed] Edelsbrunner, H., "Finding transversals for sets of simple geometric figures,", *Theoretical Computer Science*, Vol 35, 1985, pp. 55-69.

[EOW] Edelsbrunner, H., Overmars, M.H., and Wood, D., "Graphics in flatland: a case study,", *Tech. Report F79*, Technical University of Graz, 1981.

[EMPRWW] Edelsbrunner, H., Maurer, H.A., Preparata, F.P., Rosenberg, A.L., Welzl, E. and Wood, D., "Stabbing line segments,", *BIT*, Vol. 22, 1982, pp. 274-281.

[ET] ElGindy, H., and Toussaint, G.T., "Efficient algorithms for inserting and deleting edges from triangulations," *Proc. Intl. Conf. on Foundations of Data Organization*, Kyoto, Japan, May 1985.

[Gr] Grunbaum, B., “On common transversals,”, *Arch. Math.*, Vol. 9, 1958, pp. 465-469.

[HV] Horn, A., and Valentine, F.A., “Some properties of L-sets in the plane,”, *Duke Mathematics Journal*, Vol. 16, 1949, pp. 131-140.

[Ke] Ke, Y., “Detecting the weak visibility of a simple polygon and related problems,” *The Johns Hopkins University*, manuscript, March 1988.

[Le] Lewis, T., “Two counter examples concerning transversals for convex subsets of the plane,” *Geometriae Dedicata*, vol. 9, 1980, pp. 461-465.

[Me] Melkman, A., “On-line construction of the convex hull of a simple polyline,” *Information Processing Letters*, vol. 25, April 1987, pp. 11-12.

[O’R] O’Rourke, J., “An on-line algorithm for fitting straight lines between data ranges,” *Communications of the ACM*, vol. 24, No. 9, September 1981, pp. 574-578.

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

[PSu] Preparata, F. P. and Supowit, K. J., “Testing a simple polygon for monotonicity,” *Information Processing Letters*, vol. 12, No. 4, August 1981, pp. 161-164.

[SS] Sack, J.-R. and Suri, S., “An optimal algorithm for detecting weak visibility of a polygon,” *IEEE Transactions on Computers*, vol. 39, No. 10, October 1990. pp. 1213-1219.

[ST] Shermer, T. and Toussaint, G.T., “Characterizations of convex and star-shaped polygons,” in *Snapshots of Computational and Discrete Geometry*, G.T. Toussaint, editor, Tech. Report SOCS-88.11, School of Computer Science, McGill University, June 1988.

[TA] Toussaint, G.T. and Avis, D., “On a convex hull algorithm for polygons and its application to triangulation problems,” *Pattern Recognition*, vol. 15, No. 1, 1982, pp. 23-29.

[TE] Toussaint, G. T. and ElGindy, H. A., "Separation of two monotone polygons in linear time," *Robotica*, vol. 2, 1984, pp. 215-220.