Minimal Witness Sets For Art Gallery Problems

Eyüp Serdar Ayaz* Alper Üngör*

Abstract

We study the problem of finding witness sets for polygons which can be used as a first step to solve the problem of guarding art galleries. For a polygon P, a set $W \subseteq P$ is called a *witness* set if every set G that guards W, is guaranteed to guard P. Previous study exists for computing a minimal witness set of points for polygons. However, very few polygons admit witness sets of points. Here we propose an algorithm for computing witness sets of points, line segments and, if necessary, regions in $O(n^4)$ time. The output witness set is shown to be minimal if the input polygon has one, otherwise is shown to be near-minimal (as defined later in the paper). This algorithm also determines whether guarding the boundary of a polygon is sufficient to guard the entire polygon.

1 Introduction

The Art Gallery Problem (AGP) was proposed by Klee to Chvatal in 1973 as a challenge to find the point locations of a minimum number of guards such that each point on the walls of an art gallery is seen by at least one guard [3]. In general, gallery interior needs to be guarded as well. Both the original [4] and the generic [6] versions of AGP are proven to be NP-hard. It is clear that a solution for the generic version also guards the wall, but the reverse proposition is not the case. Over the years many other versions of AGP have also been studied as surveyed in [7, 9].

A common approach to solve AGP for a polygon P employs integer programming and uses a formulation with two parameters: AGP(X, Y), where $X, Y \subseteq P$, X is the set of possible guard locations and Y is the set of points to be guarded [8]. Note that AGP(X, P) and AGP(P, Y) are upper and lower bounds respectively on the minimum number of guards. Various heuristics are used to initialize X and Y and then iteratively insert elements into them until lower and upper bounds converge.

The original witnessability problem formulation is credited to Joseph Mitchell in a paper by Chwa *et al.* [2]. A *witness set* W of a polygon P is defined as a set such that if any set G that guards W also guards P. The straight-forward use of witnessability concept is checking the visibility of a subset of a polygon instead of the whole polygon from a guard set. If a witness set consists of lower dimensional elements there can be further algorithmic advantages. Another use is within the initialization step of the integer programming approach discussed above. Using appropriate witness points as Y results in early convergence. Unfortunately, not all polygons admit a finite witness set of points (See Figure 1).

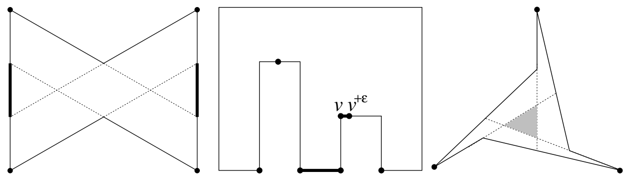

Figure 1: Three polygons admitting no witness set of points only. (Left) polygon has a minimal witness set of points and segments. (Middle) polygon has no minimal witness set but a near-minimal one, as it is necessary to include $\overline{vv^{+\epsilon}}$. (Right) polygon has a minimal witness set of three points and an interior region.

Here we extend the concept of witnessability using sets with points, line segments and if necessary regions. We propose an algorithm that finds a (near)-minimal witness set and, as a consequence, determines whether a solution for AGP(P, ∂P) guarantees a solution for AGP(P, P) i.e., ∂P is a witness set for P.

2 Preliminaries

The input for the witnessability problem is a simple (non-convex) polygon P with n vertices where $int(P)$, and ∂P denote the interior and the boundary of P, respectively. For a reflex vertex v of P, let $v^{-\epsilon}$, $v^{+\epsilon} \in \partial P$ be two infinitesimally close points to v on clockwise and counter-clockwise traversals, respectively. Next, we review some concepts and results from [2].

Two points $p, q \in P$ *see* each other if the whole line segment $\overline{pq}$ is in P. If a point p in P sees a reflex vertex v of P and the ray $\overrightarrow{pv}$ continues in P after hitting v then we say p *sees past* v. If the exterior part of the polygon is on the left side of $\overrightarrow{pv}$ in the immediate neighborhood of v then we say p *sees past left* v. Similarly, if the exterior is on the right, p *sees past right* v. (See Figure 2). For an edge e of P in counter-clockwise orientation, the half-plane induced by e is the set of points on the left side of the line

*CISE Department, University of Florida, Gainesville [ayaz,ungor]@cise.ufl.edu

through e and is denoted as $l^+(e)$. For two points $p, v \in P$ such that p sees past left (right) v, let $l^+(p, v)$ denote the half-plane that is right (left) of the line $\overrightarrow{pv}$. The closures of the half-planes $l^+(e)$ and $l^+(p, v)$ are denoted as $l^c(e)$ and $l^c(p, v)$.

The set of points in P that can be seen from a point $p \in P$ is called the *visibility polygon* of p [5], denoted as $\mathcal{V}(p)$. The set of points that can see every point in $\mathcal{V}(p)$ is called the *visibility kernel* of p, denoted as $\mathcal{VK}(p)$. For a set of points S, we use $\mathcal{VK}(S)$ as the union of the visibility kernels of the points in S.

A *witness set* can also be defined as a finite set $W \subseteq P$ such that, for any set of point guards G in P, $W \subseteq \bigcup_{g \in G} \mathcal{V}(g)$ implies $\bigcup_{g \in G} \mathcal{V}(g) = P$. If W is a witness set for P, then W is said to *witness* P. The same verb, to *witness*, can also be used with other geometric entities like points, if seeing the subject guarantees the object to be seen. A witness set W for P is said to be *minimal* if there exists no proper subset of W that witnesses P. If there exists a minimal witness set for P, then P is a *minimalizable* polygon. Otherwise, P is a *non-minimalizable* polygon.

Theorem 1 [2]. *A point set W is a witness set for a polygon P if and only if $\mathcal{VK}(W) = P$. Also the following statements are equivalent for $p, q \in P$:*
(i) p witnesses q; (ii) $q \in \mathcal{VK}(p)$; (iii) $\mathcal{VK}(q) \subseteq \mathcal{VK}(p)$.

Let p be a point on the boundary of P. Let $E(p)$ denotes the set of edges of P of which p sees at least one interior point. We define $RM(p)$, as the short form of rightmost vertex of p, the first vertex that p sees past left in counter-clockwise order from the viewpoint of p as the *rightmost* vertex with respect to p and denote as. The $leftmost$ vertex of p, $LM(p)$, is defined symmetrically. If p does not see past left (right) any vertex, then we set $RM(p)(LM(p))$ as the next vertex on ∂P in (counter-)clockwise order. Then, as proven by Chwa *et al.* [2], we have:

$$\mathcal{VK}(p) = l^c(p, RM(p)) \cap l^c(p, LM(p)) \cap \bigcap_{f \in E(p)} l^c(f) \quad (1)$$

3 Minimal and near-minimal witness sets

In this section, we define the lemmas and theorems to be used in our algorithm that finds a (near)-minimal witness set for a simple polygon.

3.1 Witnesses on the boundary of the polygon

We define *anchor points* to subdivide the boundary of the input polygon. Anchor points consist of three types: 1. The vertices of the polygon. 2. For each reflex vertex, the boundary points where the extension of each of the two edges incident to it hit first. 3. For every pair of reflex vertices that see past each other, the boundary points where the two extensions of the line segment between them hit first (See Figure 2). The line segments between two consecutive anchor points are *anchor edges*.

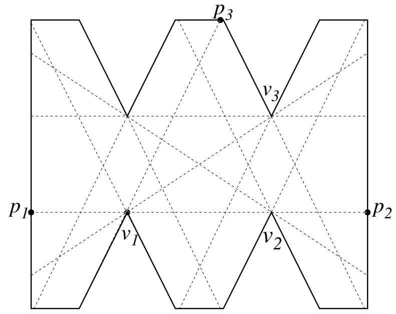

Figure 2: Partition of a polygon boundary and inducing line segments and their extensions. v_1(Type 1), p_3(Type 2) and p_1(Type 3) are three of the 32 anchor points. The line segment $\overline{v_1v_3}$ is a cross line but $\overline{v_1v_2}$ is not. Here, p_1 sees past right v_1 while p_2 sees past left v_1.

Observation 1 *Every point on an anchor edge sees past left and right the same set of reflex vertices (due to Type 2 anchor points). Also these points (partially) sees the same set of edges (due to Type 3 anchor points). Moreover the leftmost and the rightmost reflex vertices a point sees past is the same for all points within an anchor edge.*

We classify the boundary points of P exclusively into five types according to the vertices they see past left and/or right. If a boundary point p doesn't see past any vertex except the ones incident to the edge(s) p belongs to, it is of *Type Z* (See Figure 3). If there exists a vertex p sees past left but there are no vertices p sees past right, then p is of *Type L*. The symmetric version of Type L is *Type R*. If there exist a vertex p sees past left and another vertex p sees past right, there are two possibilities: If $\mathcal{VK}(p) = \{p\}$, then we say that p is of *Type D*. Otherwise, p is of *Type N*. Since the type of points within an anchor edge is the same, we also use these types for anchor edges.

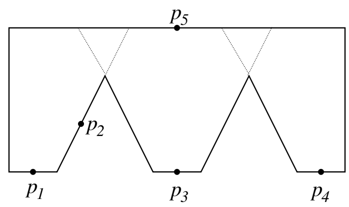

Figure 3: p_1 is of Type R, p_2 is of Type Z, p_3 is of Type N, p_4 is of Type L, p_5 is of Type D

Lemma 2 *Let e be an anchor edge, and p and p' be any two interior points of e in counter-clockwise order.*

If e is of Type L, then p witnesses p'. Similarly if e is of Type R, then p' witnesses p.

Proof. Consider the case e is of Type L. Notice that $E(p) = E(p')$ from Observation 1. Hence, $p' \in \bigcap_{f \in E(p)} l_c(f) = \bigcap_{f \in E(p')} l_c(f)$. Due to the orientation of p and p' along the boundary $p' \in l^c(p, RM(p))$. From (1), we have $p' \in \mathcal{VK}(p)$ and from Theorem 1 p witnesses p'. The symmetric case can be proven similarly. ◻

Theorem 3 *Let e be an anchor edge, v_1 and v_2 be the endpoints of e in counter-clockwise order, and p be an interior point of e. Let $A = \bigcap_{f \in E(p)} l^c(f)$, $B = l^+(v_1, RM(p)) \cup RM(p)$, and $C = l^+(v_2, LM(p)) \cup LM(p)$. $\mathcal{VK}(e)$ can be calculated using finitely many half-plane intersections:*
(a) If e is of Type D, then $\mathcal{VK}(e) = e$
(b) If e is of Type Z, then $\mathcal{VK}(e) = A$
(c) If e is of Type L, then $\mathcal{VK}(e) = A \cap B$
(d) If e is of Type R, then $\mathcal{VK}(e) = A \cap C$
(e) If e is of Type N, then $\mathcal{VK}(e) = A \cap B \cap C$

Proof. (a) It follows from the definition of Type D.
(b) From (1) we have:

$$\mathcal{VK}(e) = \bigcup_{q \in e} \mathcal{VK}(q) = \bigcup_{q \in e} \bigcap_{f \in E(q)} l^c(f) = A$$

(c) From (1), Theorem 1 and Lemma 2, we have:

$$\mathcal{VK}(e) = \bigcup_{q \in e} \mathcal{VK}(q) = \bigcup_{q \in e} (l^c(q, RM(q)) \cap \bigcap_{f \in E(q)} l^c(f))$$
$$= A \cap \bigcup_{q \in e} l^c(q, RM(q)) = A \cap B$$

(d) The symmetric case of part (c).
(e) From (1), we know that $\mathcal{VK}(p) \subset A$, hence $\mathcal{VK}(e) \subseteq A$. Also, due to the orientation of the points on e, $A \cap l^c(p, RM(p)) \subset A \cap B$. With this and the symmetrical equivalent, we can see that $\mathcal{VK}(e) \subseteq A \cap B \cap C$.

To prove the equivalence, we need to show that every point in $A \cap B \cap C$ needs to be witnessed by a point on e. When p approaches to v_1, $\mathcal{VK}(p)$ converges to $A \cap B \cap l^c(v_1, LM(p))$. With this and the symmetric version, we have $(A \cap ((B \cap l^c(v_1, LM(p)) \cup (C \cap l^c(v_2, RM(p))) \subset \mathcal{VK}(e)$. For the rest of the points, let r be a point in $(A \cap B \cap C) \setminus (A \cap ((B \cap l^c(v_1, LM(p)) \cup (C \cap l^c(v_2, RM(p)))$. Let r' be the first point the ray $\overrightarrow{v_1 r}$ hit on ∂P. Observing that $r' \in e$ implies $r \in \mathcal{VK}(r')$. ◻

3.2 Witnesses in the interior of the polygon

For every pair of reflex vertices that see past left each other or see past right each other, line segment between them is called a *cross line* (See Figure 2). A point $p \in P \setminus \mathcal{VK}(\partial P)$ is called a *ordinary point* if it is not on a cross line.

Proofs of Lemmas 4-8 are omitted due to space limitations. They can be found in the complete version.

Lemma 4 *Let p is a point on the cross line $\overline{vu}$ that is not in $\mathcal{VK}(\partial P)$. If p is not on the closure of $\mathcal{VK}(\partial P)$, then $\mathcal{VK}(p) = \{p\}$. Otherwise $\mathcal{VK}(p) \subseteq \overline{vu}$.*

Lemma 5 *A ordinary point p can only witness itself, i.e., $\mathcal{VK}(p) = \{p\}$. Moreover p is present in every witness set of P.*

3.3 Minimal witness sets

Following lemmas establish the groundwork of our results on minimal witness sets.

Lemma 6 *Let p and q be distinct points. If $\mathcal{VK}(p) \subset \mathcal{VK}(q)$, then p cannot be in any minimal witness set.*

Lemma 7 *If a point p is in $\mathcal{VK}(\partial P) \setminus \partial P$ then p cannot be in any minimal witness set.*

Lemma 8 *If a point p on ∂P is not witnessed by another point on ∂P, then p has to be in every minimal witness set.*

Theorem 9 *A minimal witness set consist of a set of boundary elements, all ordinary points, and at least one point on each cross line that is in $P \setminus \mathcal{VK}(\partial P)$.*

Proof. Let $p \in P$. Then p has to be one of these three disjoint sets: $\mathcal{VK}(\partial P)$, ordinary points and points on cross lines that are not in $\mathcal{VK}(\partial P)$. Lemma 5 states that ordinary points are in any witness set including a minimal one. Lemma 7 states that the points that are in $\mathcal{VK}(\partial P) \setminus \partial P$ cannot be in any minimal witness set. Points on cross lines that are not in $\mathcal{VK}(\partial P)$ needs to have at least one point in a minimal witness set according to Lemma 4. ◻

It is clear that not all polygons admit unique minimal witness sets as we can choose any point on a Type Z anchor edge and, for some cases, any point on a cross line. Other than those, the ordinary points are proven to be necessary in any witness set via Lemma 5. Also the points that are witnessed by ∂P proven not to be in any minimal witness set via Lemma 7. Therefore miminal witness sets for P differ by finitely many points and the cardinalities of them are equal.

3.4 Near-minimal witness sets

Suppose we have an edge segment with points Type L (R) and the left (right) endpoint of the segment is a reflex vertex. Using Lemma 2 we can show that for any point p, there is another point on the left (right) of p that witnesses p. However the line segment is open and we cannot reach on the left (right) endpoint of the line segment but we need to keep at least one point arbitrarily close to the reflex vertex. For these cases, we use the infinitesimally short line segments $\overline{pp^{-\epsilon}}$ or $\overline{pp^{+\epsilon}}$ in the witness set. (See Figure 1)

Definition 1 *(Near-minimality) Let W be a witness set for a non-minimalizable polygon P. W is* near-minimal *if it can be divided into two disjoint sets, W_{min} and W_ϵ such that W_ϵ consist of finitely many infinitesimally short line segment and removal of any element from W_{min} or W_ϵ makes W to violate the witnessing condition of W. We also call each element of W_ϵ as an* ϵ-witness.

Analogous results to Lemmas 7, 8, and Theorem 9 for near-minimal witness sets are rather immediate, however omitted here for the sake of brevity.

4 Algorithm

We use a subdivision of the polygon based on an arrangement of line segments of following three types: 1. Extensions of edges that are incident to reflex vertices until they hit other boundary points. 2. For each anchor point p, the line segments from p to $LM(p)$ and $RM(p)$. 3. For each anchor edge $\overline{v_1v_2}$, let p be an interior point of $\overline{v_1v_2}$; the line segments from v_1 to $RM(p)$ and from v_2 to $LM(p)$. We denote this subdivision as $\mathcal{A}(P)$ (stored as a doubly connected edge list) and the contiguous regions in the interior of the arrangement as *cells*.

Let W be a candidate (near-)minimal witness set initialized to be empty. For each cell c, we record the number of elements in W that witnesses c as $count(c)$. For each directed line segment s of $\mathcal{A}(P)$, incident to cells c_l and c_r, $\Delta count(s)$ stores $count(c_l) - count(c_r)$. Each line segment is marked once it is witnessed. The algorithm consists of five steps:

1. Find the anchor points and the edges they (partially) see: For this purpose, we simply employ a standard linear time visibility algorithm [5] per vertex resulting in a total cost of $O(n^2)$ time.

2. Compute $\mathcal{A}(P)$: There can be at most $O(n^2)$ line segments, therefore it takes $O(n^2 \log n + k)$ time where k is the number of intersecting points [1], which is in $O(n^4)$. Hence, the total time for this step is $O(n^4)$.

3. Find the elements of a (near-)minimal witness set on ∂P: The visibility kernel of each anchor point or edge is contiguous and inclusive. Therefore we can traverse the boundary of the visibility kernel starting from the anchor point or edge. When we insert an element b to W, we traverse the boundary of $\mathcal{VK}(b)$ in counter-clockwise order and we increment the $\Delta count$ value of each line segment by one and decrement the $\Delta count$ of the opposite direction. When we remove an element from W, which happens when another boundary point witnesses an element on W, we backtrack this increment/decrement. For Type N and D, the whole line segments are inserted to W. To keep the minimality of W, for Type D, we chose the middle point to be inserted to W. For Type L (R), we insert the left (right) endpoint of the line segment to W if it is not a reflex vertex. If the endpoint is a reflex vertex, we insert an ϵ-witness incident to the corresponding end point. Note that the visibility kernels of both anchor points and anchor edges are convex. Therefore, each of $O(n^2)$ line segments of $\mathcal{A}(P)$ can intersect a visibility kernel twice. There can be at most $O(n^2)$ visibility kernels. Therefore this step costs $O(n^4)$ time.

4. Find the cells of a (near-)minimal witness set in $int(P)$: The *count* values can be retrieved starting from a boundary cell and using the $\Delta count$ values of incident line segments of $\mathcal{A}(P)$. At the end, we insert the cells that have witness *count* 0 and the unwitnessed line segments of $\mathcal{A}(P)$ that are not in the closure of $\mathcal{VK}(\partial P)$ to W. As a last step, we traverse each line segment s on the boundary of $\mathcal{VK}(\partial P)$. If $s \notin \mathcal{VK}(\partial P)$ we insert either s or a single point of s depending if s is on a cross line. This step is done in $O(n^4)$ time.

If there exists no ϵ-witness element in W then W is minimal. Otherwise the polygon is not minimalizable and W is near-minimal. If step 4 does not find an element in $int(P)$, then ∂P witnesses P. As we can see from the steps, the algorithm works in $O(n^4)$ time.

References

[1] B. Chazelle, H. Edelsbrunner, An optimal algorithm for intersecting line segments in the plane, *Proc. 29th IEEE Symp. on Found. of Comp. Sci.*, 217-225, 1988.

[2] K.-Y. Chwa, B.-C. Jo, C. Knauer, E. Moet, R. van Oostrum, and C.-S. Shin, Guarding Art Galleries by Guarding Witnesses. *Int. J. Comp. Geometry Appl.*, 16(2-3): 205-226, 2006.

[3] R. Honsberger, *Mathematical Gems II*, Mathematical Association of America, 1976.

[4] A. Laurentini. Guarding the walls of an art gallery, *The Visual Computer*, 15(6):265-278, 1999.

[5] D.T. Lee, Visibility of a simple polygon, *Comp. Vision, Graphics, and Image Processing* 207-221, 1983.

[6] D.T. Lee and A. K. Lin, Computational complexity of art gallery problems. *IEEE Trans. Inf. Theor.*, 32(2):276-282, 1986.

[7] J. O'Rourke, Art Gallery Theorems and Algorithms. Oxford University Press, 1987.

[8] P. J. de Rezende, C. C. de Souza, S. Friedrichs,M. Hemmer, A. Krller and D. C. Tozoni, Engineering Art Galleries. arXiv preprint arXiv:1410.8720.

[9] J. Urrutia, Art gallery and illumination problems. *Handbook on Comp. Geo.*, North-Holland. 973-1027, 2000.