Guarding Art Galleries by Guarding Witnesses

Kyung-Yong Chwa, Byung-Cheol Jo,

Christian Knauer, Esther Moet,

René van Oostrum & Chan-Su Shin

institute of information and computing sciences, utrecht university

technical report UU-CS-2003-044

www.cs.uu.nl

Guarding Art Galleries by Guarding Witnesses

Kyung-Yong Chwa* Byung-Cheol Jo† Christian Knauer‡ Esther Moet§
René van Oostrum§ Chan-Su Shin¶

December 2003

Abstract

Let P be a simple polygon. We define a *witness set* W to be a set of points such that if any (prospective) guard set G guards W, then it is guaranteed that G guards P. We show that not all polygons admit a finite witness set. If a finite minimal witness set exists, then it cannot contain any witness in the interior of P; all witnesses must lie on the boundary of P, and there can be at most one witness per edge. We give an algorithm to compute a minimal witness set for P in $O(n^2 \log n)$ time, if such a set exists, or to report the non-existence within the same time bounds. We also outline two algorithms that use a witness set for P to test whether a (prospective) guard set sees all points in P.

1 Introduction

Visibility problems have been studied extensively in the Computational Geometry literature, and the so-called *Art Gallery Problems* form an important subcategory within this field. The problem of how many points or *guards* are always sufficient to guard any polygon with n vertices was posed by Victor Klee in 1976. Chvátal [4] showed soon thereafter that $\lfloor \frac{n}{3} \rfloor$ guards are always sufficient, and sometimes necessary. Since then, a lot of research in this field has been carried out, and in 1976 O'Rourke published a whole book on the topic [12], while many new results that had been achieved after the publication of O'Rourke's book have appeared in surveys by Shermer [13] and Urrutia [14].

Most of the papers related to the Art Gallery Problem research consider the computation of the location of the guards under various restrictions on the shape of the polygon (i.e., orthogonal polygons, polygons with holes, etc.) or on the placement of the guards (i.e., edge guards, vertex guards, mobile guards, etc.). Except for the case of a single guard, only few papers consider the computation of the guarded region for a *given* set of guards [3, 5].

Approximately seven years ago, Joseph Mitchell posed the *Witness Problem* to Tae-Cheon Yang during a research visit of the latter: "Given a polygon P, does it admit a *witness set*, i.e., a set of objects in P such that any (prospective) guard set that guards the witnesses is

*Department of Computer Science, KAIST, Daejon, Korea. Email: kychwa@tclab.kaist.ac.kr
†Taff System, Co. Ltd., Seoul, Korea. Email: mrjo@taff.co.kr
‡Freie Universität Berlin, Takustraße 9, D-14195 Berlin, Germany. Email: knauer@inf.fu-berlin.de
§Institute of Information & Computing Sciences, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands. Email: {esther,rene}@cs.uu.nl
¶School of Electronics and Information Engineering, Hankuk University of Foreign Studies, Yongin, Korea. Email: cssin@hufs.ac.kr

guaranteed to guard to whole polygon?" The simplest kind of witnesses are points, possibly constrained to lie on the vertices or on the boundary of P. The idea, of course, is that in the case of a set of moving guards, or of a guard set that permits the addition or removal of guards, it is easier to check point-to-point (i.e., guard-to-witness) visibility, than to update the complete visibility region of all guards. Other possible types of witnesses are edges of the polygon. Yang showed this problem to a student, and they made an initial effort to classify polygons that can be witnessed by guards placed at vertices at the polygons, and polygons that can be witnessed by (partial) edges of the polygon; see [15] (in Korean).

In this paper we consider point witnesses that are allowed to lie anywhere in the interior or on the boundary of the polygon. We want to determine for a given polygon P whether a finite witness set exists, and if this is the case, to compute a *minimal* witness set. The main contribution of our paper is the combinatorial/geometrical result that a finite minimal witness set for a polygon (if it exists) contains only witnesses on the boundary, and there are at most n of them, where n is the number of vertices of the polygon. Furthermore, we show that for any $n \geq 4$ there is a polygon that can be witnessed by no less than $n-2$ witnesses. The results are not trivial: there are polygons that are not witnessable by a finite witness set, but that need an infinite number of witnesses on their boundary or in their interior. There are also polygons for which there are witness sets that witness the boundary, but not the whole interior of the polygon.

For completeness, we also give an algorithm to compute a minimal witness set for a polygon if such a set exists, or to report the non-existence otherwise. The running time of this algorithm is $O(n^2 \log n)$. Also, we show how to use witness sets for testing if a set of (prospective) guards actually sees the entire polygon.

The paper is organized as follows: in the next section we introduce the formal definition of *witness sets*, and we prove some interesting basic properties of them. In Section 3 we present a number of examples to give the reader a better understanding of the problem. In Section 4 we study the properties of the so-called *visibility kernel*. These properties are used in the subsequent section to establish our main results on finite witness sets. In section 6 we give algorithms for computing and using minimal witness sets, and we wrap up in Section 7 with a brief discussion on our results and on problems that remain unsolved yet.

2 Preliminaries

Throughout this paper, P denotes a simple polygon with n vertices $V(P) = \{v_0, v_1, \ldots, v_{n-1}\}$; we assume that the vertices are ordered in counterclockwise direction. The edges of P are denoted with $E(P) = \{e_0, e_1, \ldots, e_{n-1}\}$, with $e_i = (v_i, v_{i+1 \bmod n})$. Geometrically, we consider the edge e_i to be the closed line segment between its incident vertices, i.e., an edge includes its endpoints. Edges e_i are directed from v_i to v_{i+1}, so that the interior of P lies locally to the left of e_i. We say that a point p lies in P if p lies in the interior of P (denoted with $int(P)$) or on its boundary (denoted with ∂P), i.e., we consider P to be a closed subset of $\mathbb{E}^2$.

A point p in P *sees* a point q in P if the line segment pq is contained in P. Since polygons are closed regions, the line-of-sight pq is not blocked by grazing contact with the boundary of P; this definition of visibility is commonly used in the Art Gallery literature [12].

We say that a point p in P *sees past* a reflex vertex v of P if p sees v, and the edges incident to v do not lie on different sides of the line through p and v (i.e., one of the edges may lie on this line).

Let p be a point in P. The *visibility polygon* of p is the set of points in P that are visible to p. We denote the visibility polygon by $\mathrm{VP}(p)$. The *visibility kernel* of a point p is the kernel of its visibility polygon and is denoted by $\mathrm{VK}(p)$.

Definition 1 *A* witness set *for a polygon P is a point set W in P for which the following holds: if, for any arbitrary set of points G in P, each element of W is visible from at least one point in G, then* every *point in P is visible from at least one point in G.*

The elements of G in the above definition are commonly referred to as *guards*, and we call the elements of a witness set *witnesses*.

The following theorem states the necessary and sufficient conditions on witness sets:

Theorem 1 *A point set W is a witness set for a polygon P if and only if the union of the visibility kernels of the elements of W covers P completely.*

Proof. Let $W = \{w_1, \ldots, w_k\}$ be a set of points in P such that $\cup_{w \in W} \mathrm{VK}(w)$ covers P completely. Let G be an arbitrary set of points in P such that every element $w_i \in W$ is visible from at least one $g_j \in G$. Since g_j lies in $\mathrm{VP}(w_i)$, we have that all points in $\mathrm{VK}(w_i)$ are visible from g_j. Since there is such a g_j for every w_i, and $\cup_{w \in W} \mathrm{VK}(w)$ covers P completely, it follows that every point in P is visible from at least one $g_j \in G$, and therefore W is a witness set for P.

Conversely, let W be an arbitrary witness set for P. Let us assume for the sake of contradiction that $\cup_{w \in W} \mathrm{VK}(w)$ does *not* cover P completely.

Let us first consider the case where the union of the visibility *polygons* (as opposed to the visibility *kernels*) of all elements of W do not cover P completely. In this case, a contradiction is easily derived: place a guard g_i at every witness w_i. Now every witness is seen by at least one guard, but the guards do *not* see the whole polygon, so W is not a witness set for P.

It remains to consider the more interesting case where the union of the visibility polygons of all elements of W *does* cover P completely. Then we pick an arbitrary point p in the region of P that is *not* covered by any of the visibility kernels of the witnesses. For any $w_i \in W$, p may or may not lie in $\mathrm{VP}(w_i)$, but in either case, since p does not lie in $\mathrm{VK}(w_i)$, p cannot see *all* points in $\mathrm{VP}(w_i)$. This means that for each $w_i \in W$ we can place a guard g_i in $\mathrm{VP}(w_i)$ such that g_i does *not* see p. So every witness w_i is seen by at least one guard (namely, g_i), but the guards do *not* see every point in P (for none of the guards sees p). This means that W is not a witness set for P, and we have a contradiction again. $\square$

We also apply the concept of witnesses to individual points. For two points p and q in a polygon P, we say that p is a witness for q (or alternatively, that p witnesses q), if any point that sees p also sees q.

The following lemma is analogous to Theorem 1, and we omit the proof (which is much simpler than the proof of the theorem):

Lemma 1 *If p and q are points in a polygon P, then p witnesses q if and only if q lies in $\mathrm{VK}(p)$.*

The following two lemmas show that witnessing is transitive:

Lemma 2 *Let P be a polygon. A point p in P witnesses a point q in P if and only if $\mathrm{VP}(p) \subseteq \mathrm{VP}(q)$.*

PROOF. If a point p witnesses a point q, then $q \in \mathrm{VK}(p)$ by the preceding lemma. This implies that everything that is visible from p is also visible from q, which means that $\mathrm{VP}(p) \subseteq \mathrm{VP}(q)$.

Conversely, if $\mathrm{VP}(p) \subseteq \mathrm{VP}(q)$, then any point that sees p, and therefore lies in $\mathrm{VP}(p)$, also lies in $\mathrm{VP}(q)$, and consequently sees q. ◻

Lemma 3 *Let P be a polygon, and let p, q, and r be points in P. If p witnesses q and q witnesses r, then p witnesses r.*

PROOF. If p witnesses q and q witnesses r, then by the preceding lemma, $\mathrm{VP}(p) \subseteq \mathrm{VP}(q)$ and $\mathrm{VP}(q) \subseteq \mathrm{VP}(r)$. This means that $\mathrm{VP}(p) \subseteq \mathrm{VP}(r)$ and thus that p witnesses r. ◻

This leads to the notion of *minimal witness sets*:

Definition 2 *Let P be a polygon and let W be a witness set for P. W is called a* minimal witness set *for P if, for any $w \in W$, $W \setminus \{w\}$ is* not *a witness set for P.*

The proofs of the following two lemmas are straightforward, and we omit them here.

Lemma 4 *Let P be a polygon, and let W be a witness set for P. W is a minimal witness set for P if and only if for any $w \in W$, w does not lie in $\mathrm{VK}(w')$ for any $w' \in W, w' \neq w$.*

Lemma 5 *Let P be a polygon. If W is a witness set for P, then (i) there exists a subset $W' \subseteq W$ such that W' is a minimal witness set for P, and (ii) for any superset $W'' \supseteq W$, W'' is a witness set for P.*

As for the second claim: note that, although adding witnesses to a witness set W may seem to place additional constraints on the placement of the guards, this is actually not the case. If the guards see all witnesses in W, then they already see the whole polygon, including any additional witnesses.

3 Examples

Although the definition of a witness set and the properties thereof that we discussed in the previous section are not hard to understand, we observed while explaining the problem to colleagues and while discussing it amongst the authors that it is very easy to confuse the constraints on the witness set with those on the guards. Furthermore, we think that it is helpful to provide the reader some intuition for the problem, and therefore we give a few examples in this section. We introduce the examples informally: we do not prove any claims that we make about them. If the reader wishes to check the claims, she may use the lemmas and theorems that we give in Sections 4 and 5.

Our examples are given in Figure 1; all four polygons are simple polygons (note that the white regions in the lower two polygons are not holes, but unwitnessed regions).

In Figure 1(a) we see a star-shaped polygon. Clearly, a point p in the kernel guards the whole polygon. One may mistakenly believe that a single point in the kernel is a witness set for the polygon as well, since p sees the witness, and p also guards the polygon. This, however, is incorrect: according to Definition 1, *any* set of points that sees the witnesses must see the whole polygon. Now let p play the role of witness. A single point in one of the arms of the

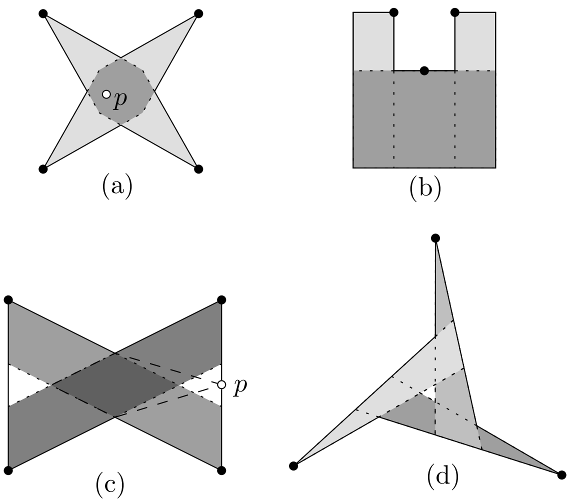

Figure 1: Examples of the regions in polygons witnessed by vertices and points on edges. Note that the white regions in the lower two polygons are not holes, but unwitnessed regions.

polygon (i.e., outside the kernel of the polygon) sees p, but it does not see the whole polygon, and therefore, $\{p\}$ is *not* a valid witness set.

The polygon of Figure 1(a) can be witnessed by the points shown in black. It is not true that a polygon has to be star-shaped to be witnessable by a finite witness set; see Figure 1(b). Neither is it true that *all* star-shaped polygons can be witnessed by a finite witness set; see Figure 1(c).

In Figure 1(b) we see a polygon that is witnessable by a set of three witnesses. One of the witnesses is placed in the interior of an edge. In fact, this polygon is not witnessable by a witness set that has only witnesses placed at vertices.

Figure 1(c) shows a star-shaped polygon, and a set of four points placed on vertices of the polygon. Their visibility kernels do not cover the whole polygon; the white regions are not witnessed. Adding witnesses at the remaining two vertices does not help, because these are already witnessed and witnessing is transitive. To witness the white regions, we need to place witnesses that lie on the boundary of the polygon and in the white regions. Consider such a witness p. Its visibility kernel is the convex hull of the kernel of the polygon (darkly shaded) and p itself, so the only point on the boundary that is witnessed by p is p itself. It turns out that we need to cover both unwitnessed segments on the boundary of the polygon completely with witnesses to get an (infinite) witness set for this polygon.

Finally, Figure 1(d) shows an example of a polygon that cannot be witnessed with a set of points on the boundary; even if we would cover the whole boundary of this polygon with witnesses, the white region in the middle would remain unwitnessed. A witness set for the polygon would have to include all points in this region (and each point in the region only witnesses itself).

4 Visibility kernels

In this section we study several properties of visibility kernels, that are used in the next section to establish our main results on finite witness sets.

Let P be a polygon with n vertices and edges, as defined in Section 2. Let $\ell(e)$ be the directed line through an edge e of P such that $\ell(e)$ has the same orientation as e. The positive halfspace induced by $\ell(e)$ is the region of points in $\mathbb{E}^2$ to the left of $\ell(e)$, and we denote it by $\ell^+(e)$. The negative halfspace $\ell^-(e)$ is defined analogously. The *closure* of a (possibly open) region of points $R \subset \mathbb{E}^2$ is the union of R and its boundary ∂R; we denote it with $cl(R)$.

It is well-known that the kernel of a polygon P is the intersection of halfspaces $\bigcap_{e\in E(P)} cl(\ell^+(e))$; when this kernel is non-empty, the polygon is said to be *star-shaped*. The visibility polygon $\mathrm{VP}(p)$ of a point p in P is star-shaped by definition (the kernel contains at least p), and its kernel $\mathrm{VK}(p)$ can be computed as $\bigcap_{e\in E(\mathrm{VP}(p))} cl(\ell^+(e))$. However, there is an alternative way of describing $\mathrm{VK}(p)$ that turns out to be useful.

The edges of the visibility polygon $\mathrm{VP}(p)$ can be classified into two groups:

1. An edge e of $\mathrm{VP}(p)$ coincides with the part of an edge e' of P that is visible from p. In this case, $\ell^+(e) = \ell^+(e')$. We denote the set of edges of P that are (partially) visible from p with $S_1(p)$.

2. An edge e of $\mathrm{VP}(p)$ is induced by the directed line $\ell(p, v)$ through p and a reflex vertex v of P such that p sees past v. The line $\ell(p, v)$ coincides with $\ell(e)$, and the orientation of the former line is defined to be the same as the orientation of the latter line. We denote the set of reflex vertices v_i of P such that p sees past v_i with $S_2(p)$.

The visibility kernel $\mathrm{VK}(p)$ of p can now alternatively be described as $\bigcap_{e\in S_1(p)} cl(\ell^+(e)) \cap \bigcap_{v\in S_2(p)} cl(\ell^+(p, v))$.

The reader may wonder why we introduce this seemingly complicated alternative representation of $\mathrm{VK}(p)$. The reason is that for any p, at most two vertices in $S_2(p)$ suffice to define $\bigcap_{v\in S_2(p)} cl(\ell^+(p, v))$. This, in turn, helps us to reduce the complexity of the data structures involved in computing the union of a set of visibility kernels; see Section 6.

Lemma 6 *Let $S_2(p)$ be the subset of the reflex vertices of P as defined above. Then either $\bigcap_{v\in S_2(p)} cl(\ell^+(p, v))$ consists of only the point p itself, or there exists a subset $S'_2(p) \subset S_2(p)$ of cardinality at most two such that $\bigcap_{v\in S'_2(p)} cl(\ell^+(p, v)) = \bigcap_{v\in S_2(p)} cl(\ell^+(p, v))$.*

Proof. We argue that we can construct the intersection of halfspaces incrementally, and we prove the lemma by induction. Let the vertices in $S_2(p)$ be labeled $v_1, v_2, \ldots, v_k$, and let $S_2^i(p) \subseteq S_2(p)$ be $\{v_1, v_2, \ldots, v_i\}$, for $1 \le i \le k$.

Note that $S_2^k(p)$ is equivalent to $S_2(p)$, and observe that all lines in $\{\ell(p, v) | v \in S_2(p)\}$ pass through p. Let R_i be a shorthand for $\bigcap_{v\in S_2^i(p)} cl(\ell^+(p, v))$.

The lemma holds trivially for $S_2^1(p)$ and $S_2^2(p)$. Suppose that the lemma holds for $S_2^i(p)$, for some $2 \le i < k$. Then we can distinguish two cases: either R_i consists of only the point p, or it is equivalent to $cl(\ell^+(p, v')) \cap cl(\ell^+(p, v''))$, for $v', v'' \in S_2^i(p)$.

1. if R_i consists of only the point p, then $R_{i+1} = R_i \cap cl(\ell^+(p, v_{i+1}))$ also consists of only the point p, since p is contained in $cl(\ell^+(p, v_{i+1}))$ as well.

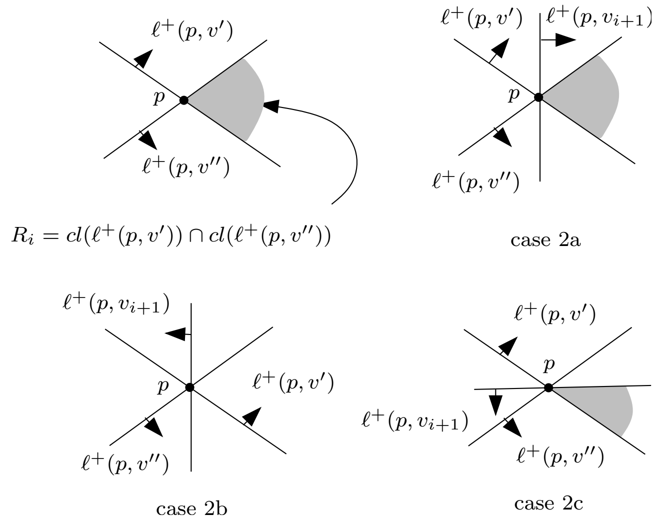

Figure 2: The intersection of closed halfspaces defined by lines through p is either p, or it is defined by two of the lines.

2. if R_i is equivalent to $cl(\ell^+(p, v')) \cap cl(\ell^+(p, v''))$, for $v', v'' \in S_2^i$, then we distinguish three sub-cases (see Figure 2):

   (a) R_i lies completely in $cl(\ell^+(p, v_{i+1}))$. In that case, $R_{i+1} = R_i$, and therefore R_{i+1} is also equivalent to $cl(\ell^+(p, v')) \cap cl(\ell^+(p, v''))$.

   (b) R_i lies completely in $cl(\ell^-(p, v_{i+1}))$. In that case, R_{i+1} consists of only the point p.

   (c) R_i lies partially in $cl(\ell^+(p, v_{i+1}))$. In that case, R_{i+1} is either $cl(\ell^+(p, v')) \cap cl(\ell^+(p, v_{i+1}))$ or
   $cl(\ell^+(p, v_{i+1})) \cap cl(\ell^+(p, v''))$

In either case and sub-case, if the lemma holds for $S_2^i(p)$, for some $2 \leq i < k$, it also holds for $S_2^{i+1}(p)$, and since it holds for $S_2^1(p)$ and $S_2^2(p)$, we conclude that the lemma holds for $S_2^k = S_2(p)$. □

Corollary 1 *Let P be a polygon, and let W be a point set, located in P, with cardinality m. For any point p, $S_1(p)$ and $S_2'(p)$ are defined as above. The union $\bigcup_{w \in W} S_1(w)$ has at most n elements, namely all edges of P. The union $\bigcup_{w \in W} S_2'(w)$ has cardinality at most $2m$.*

We conclude this section with some properties of visibility kernels that are of use in the remainder of this paper.

Lemma 7 *If a point p in a polygon P sees past a reflex vertex $v \in V(P)$, then p lies on the boundary of* $\mathrm{VK}(p)$.

Proof. The lemma follows easily from the facts that (i) p lies on $\ell(p, v)$, (ii) $\mathrm{VK}(p)$ is contained in $\ell^+(p, v)$, and (iii) p lies in $\mathrm{VK}(p)$. □

Lemma 8 *For a star-shaped polygon P and any point p in P, $\mathrm{VK}(p)$ contains the kernel of P.*

PROOF. Any point in P sees all points q in the kernel of P, and this holds no less for any point that can see p. Therefore, p is a witness for all points q in the kernel of P, and by Lemma 1 it follows that all the points q in the kernel of P lie in $\mathrm{VK}(p)$. ◽

Lemma 9 *Let p and q be points in a polygon P. If q lies outside $\mathrm{VK}(p)$, then q sees past at least one reflex vertex $v \in V(P)$.*

PROOF. If P is star-shaped, then q cannot lie in the kernel of P, since by the preceding lemma, q would otherwise lie in $\mathrm{VK}(p)$, which contradicts the premise of this lemma. So in this case, there must be a point q' in P that is not visible from q.

Otherwise, if P is not star-shaped, there is also a point q' that is not visible from q.

In both cases, the shortest geodesic path from q to q' has consecutive vertices $u_0 = q, u_1, \ldots, u_k, u_{k+1} = q'$, where for each $1 \leq i \leq k$ the vertex u_i is a reflex vertex of P [10], and $k \geq 1$. This means that q sees past u_1. ◽

Lemma 10 *Let P be a polygon, and let p be a point in its interior. Then P is convex if and only if p witnesses all points in P.*

PROOF. If P is convex, then any non-empty guard set guards the polygon. So in this case, the empty set is a witness set for P, and by Lemma 5 $\{p\}$ is also a witness set.

Conversely, if p witnesses all points in P, then p must also *see* all points in P; otherwise, a single guard placed at p would not see the whole polygon. This means that (i) p sees all edges of P, so $S_1(p) = E(P)$, and (ii) p does not see past any reflex vertex v of P (otherwise, one of the two edges incident to v would not be visible from p). So $S_2(p) = \emptyset$.

We now have that $\mathrm{VK}(p) = \bigcap_{e \in E(P)} \ell^+(e)$. But this intersection of half-spaces only covers P if P is convex. ◽

5 Finite witness sets

We would like to determine for a given simple polygon P whether a finite witness set for P exists, and if so, to compute such a set. Suppose that such a set W' exists. Then according to Lemma 5, there exists also a minimal subset $W \subseteq W'$. We denote its cardinality with k, so that $W = \{w_1, w_2, \ldots, w_k\}$.

We have seen that a point p witnesses a point q if q lies in $\mathrm{VK}(p)$. This means that for a witness set W, the union $\bigcup_{w \in W} \mathrm{VK}(w)$ must cover the whole polygon p. Lemma 6 suggests the following approach for testing whether a set W is indeed a witness set:

For an edge $e \in E(P)$, $\ell(e) \cap P$ consists of one or more connected parts (segments) that lie completely in P or on its boundary. The segment that contains e is denoted by $s(e)$. Similarly, for a point p in P and a reflex vertex $v \in S_2'(p)$ (as defined in Lemma 6), $\ell(p, v) \cap P$ consists of one or more connected parts (segments) that lie completely in P or on its boundary. The segment that contains p and v is denoted by $s(p, v)$. These two types of line segments are illustrated in Figure 3.

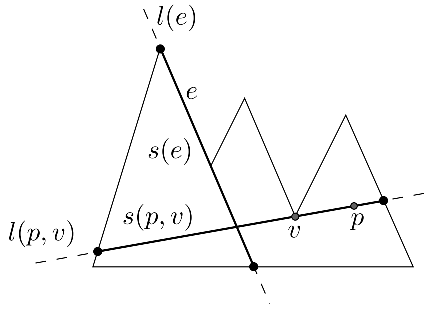

Figure 3: The two types of line segments that are used to form the arrangement $\mathcal{A}(P,W)$.

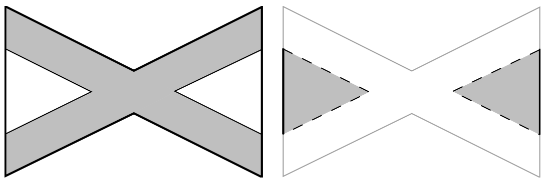

Figure 4: Since the union of visibility kernels (left) is closed, the resulting unwitnessed regions (right) may be neither open nor closed.

For a polygon P and a set W of points in P, let $\mathcal{A}(W)$ be the arrangement induced by the set of segments $\mathcal{S}(W) = \{s(e)|e \in E(P)\} \cup \{s(w,v)|w \in W \wedge v \in S'_2(w)\}$, with $S'_2(w)$ as defined in Lemma 6. For any cell c of $\mathcal{A}(W)$ and any point $w \in W$, c lies either completely inside or completely outside $\mathrm{VK}(w)$. This means that W is a witness set for P if and only if every cell of $\mathcal{A}(W)$ is contained in $\mathrm{VK}(w)$ for at least one $w \in W$.

If we denote the cardinality of W by m, then, by Corollary 1, $\mathcal{S}(W)$ contains at most $n + 2m$ segments, and the complexity of $\mathcal{A}(W)$ is $O((n+m)^2)$. We discuss how to test the cells of $\mathcal{A}(W)$ on containment in visibility kernels in Section 6.

In the remainder of this section we bound the number of witnesses of a finite minimal witness set W for a polygon P from above. We show first that the elements of W can only lie on the boundary of P (Lemma 11), and next, that any edge of P has at most one element of W on it (Lemma 13).

We denote the regions of points in P witnessed by any of the elements of a set S of points in P by $\mathcal{R}(S)$. If S is finite, then $\mathcal{R}(S)$ consists of one or more *closed* polygonal regions, since it is the union of a finite set of kernels of visibility polygons, which are closed. Note that in degenerate cases, a region in $\mathcal{R}(S)$ may be a single point or a line segment. The regions of points that are *not* witnessed by any of the elements of S are denoted by $\mathcal{Q}(S) = P \setminus \mathcal{R}(S)$, and these are called the *unwitnessed regions*. $\mathcal{Q}(S)$ consists of one or more connected (but not necessarily simply connected) polygonal regions that are either open (if the region lies in $int(P)$) or neither open nor closed (if the region contains part of ∂P; see Figure 4). The important fact is that no part of $\partial\mathcal{Q}(S)$ that lies in $int(P)$ belongs to $\mathcal{Q}(S)$ itself.

Lemma 11 *Let P be a simple polygon. If W is a finite minimal witness set for P, then no element of W lies in $int(P)$.*

Proof. If W is the empty set, then the lemma holds trivially.

W cannot have only one element. Otherwise, by Lemma 10, P would be convex, and since the empty set is also a valid witness set for convex polygons, this would contradict the minimality of W.

It remains to show that the lemma holds in case W has more than one element. In this case, consider any witness $w \in W$. Note that $\mathcal{Q}(W \setminus \{w\})$ cannot be a point or a one-dimensional region; otherwise, $\mathcal{R}(W \setminus \{w\})$ would not be closed, and this is only possible is W is an infinite set. From Lemma 4 we deduce that w lies in $\mathcal{Q}(W \setminus \{w\})$. Since no part of of $\partial\mathcal{Q}(W \setminus \{w\})$ that lies in $int(P)$ belongs to $\mathcal{Q}(W \setminus \{w\})$ itself, w can only lie in $int(\mathcal{Q}(W \setminus \{w\}))$ or on $\mathcal{Q}(W \setminus \{w\}) \cap \partial P$. Since W has at least two elements, there is another witness w', and $w \notin \mathrm{VK}(w')$ (otherwise, by Lemma 4, W would not be minimal). Then by Lemma 9 and Lemma 7, w lies on the boundary of $\mathrm{VK}(w)$. Since w cannot lie simultaneously on the boundary of $\mathrm{VK}(w)$ and in $int(\mathcal{Q}(W \setminus \{w\}))$, the conclusion is that w can only lie on $\mathcal{Q}(W \setminus \{w\}) \cap \partial P$. ◻

We established that witnesses of a finite minimal witness set lie on the boundary of P. Using the following lemma, we show in Lemma 13 that every edge of P has at most one element of W on it.

Lemma 12 *Let P be a simple polygon, let W be a finite minimal witness set for P, and let w be an element of W. If w lies on an edge e of P, then w does not see past any reflex vertex v of P that does not lie on $\ell(e)$.*

Proof. Suppose, for the sake of contradiction, that w *does* see past one or more reflex vertices v of P that do not lie on $\ell(e)$. Then $S_2(w)$ (see Section 4) is non-empty, and $\mathrm{VK}(w)$ is contained in $cl(\ell^+(w, v))$ for some $v \in S_2(w)$.

We now regard only the boundary of P. The intersection of any visibility kernel (which is convex) with ∂P consists of one or more parts that are homeomorphic to a point or a to a closed line segment. Similarly, since W is finite, $\mathcal{R}(W \setminus \{w\}) \cap \partial P$ consists of one or more parts that are homeomorphic to a point or a closed line segment, and $\mathcal{Q}(W \setminus \{w\}) \cap \partial P$ consists of one or more parts that are homeomorphic to an *open* line segment.

The remainder of the proof is essentially a one-dimensional version of the proof of Lemma 11. Since w lies on the intersection of e and $\ell(w, v)$, and $\mathrm{VK}(w)$ is contained in $cl(\ell^+(w, v))$ for some $v \in S_2(w)$, w must be the endpoint of one of the parts of $\mathrm{VK}(w) \cap \partial P$. Since $\mathcal{Q}(W \setminus \{w\}) \subset \mathrm{VK}(w) \cap \partial P$, and w must lie in $\mathcal{Q}(W \setminus \{w\})$ (otherwise, by Lemma 4, W would not be minimal), w must also lie on endpoint of one of the parts in $\mathcal{Q}(W \setminus \{w\}) \cap \partial P$. Since w cannot lie simultaneously in an open subset and on the boundary thereof, we conclude that w does not see past any reflex vertex that does not lie on $\ell(e)$. ◻

Note that the above lemma does not contradict Lemma 9. A witness w on an edge e sees the vertices of P that lie on $s(e)$, and combining the two lemmas implies that w must see past at least one of these vertices.

Lemma 13 *Let P be a simple polygon, and let W be a finite minimal witness set for P. No edge e of P contains more than one witness $w \in W$.*

Proof. Let e be an edge of P that contains a witness $w \in W$. Let e' be the closed segment $e \cap \mathrm{VK}(w)$, and let p and p' be the endpoints of e'. No other witness $w' \in W$ can lie on e in between p and p' (otherwise, w' would lie in $\mathrm{VK}(w)$, and by Lemma 4 this contradicts the

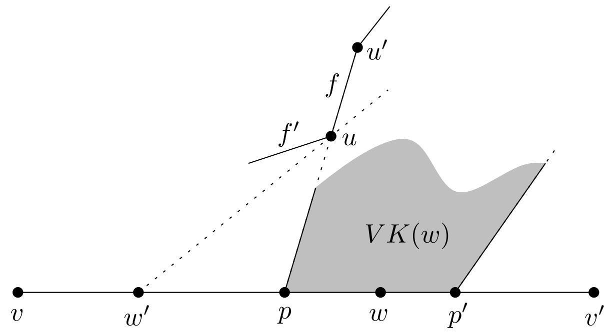

Figure 5: Any point w' in between v and p sees past u or (not shown) past another reflex vertex in the triangle formed by w', u, and p.

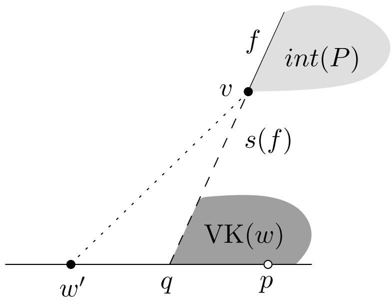

Figure 6: If an edge e is (partially) witnessed by a witness w not on e, then there cannot be a witness w' on e.

minimality of W). Therefore, if p and p' coincide with the vertices v and v' incident to e, then the lemma holds.

So let us consider the situation that not both of p and p' coincide with v and v', respectively, and let us assume without loss of generality that p does not coincide with v; see Figure 5. Since p lies on the boundary of VK(w), and, by Lemma 12, there is no reflex vertex u of P in $\ell^+(e)$ such that w sees past u, p must be the intersection of e and $s(f)$ (the extension of an edge f, as defined in Section 4) for some edge f of P, visible from w. Let u and u' be the vertices incident to f, with u the one closest to e. Let f' be the other edge incident to u. Note that f' must also be visible from w, otherwise, w would see past u, and that would contradict Lemma 12. Now suppose that there is a witness $w' \in W$ that lies on the half-open segment from v to p (including v, but not including p). If w' sees u, then w' also sees past u. Otherwise, if w' doesn't see u, there must be a reflex vertex u'' on the shortest geodesic path from w' to u such that w' sees past u'' (note that in the latter case, u'' must lie in the interior of the triangle formed by w', u, and p). In either case we derive a contradiction with Lemma 12, and the conclusion is that the half-open segment from v to p cannot contain any witness. Analogously, if p' doesn't coincide with v', then the half-open segment from v' to p' cannot contain a witness either. ◻

Lemma 14 *Let P be a simple polygon, and let W be a minimal finite witness set for P. If a witness $w \in W$ that does* not *lie on an edge e of P witnesses a point in the interior of e, then there cannot be a witness $w' \in W$ that lies on e.*

Proof. Suppose that there is a witness w, which does not lie on an edge e, that witnesses a point p in the interior of e. Since p lies in the interior of e, w must lie in $cl(\ell^+(e))$. Now suppose, for the sake of contradiction, that there is a witness w' on e, and without loss of generality, assume that w' lies to the left of p. See Figure 6. We have that p lies in $\mathrm{VK}(w)$, but w' does not lie in $\mathrm{VK}(w)$ (otherwise, by Lemma 4, the minimality of W would be contradicted). So the boundary of $\mathrm{VK}(w)$ must intersect e in between w' and p; let the intersection point be q. This point q is actually the intersection of e and $s(f)$, where f is an edge visible from w. Let the vertex incident to f that is closest to q be v. Let f' be the other edge incident to v. Note that the boundary of P cannot intersect the interior of the segment qv; otherwise, $s(f)$ would not extend all the way to q. We assume that f' does not lie on $s(f)$; otherwise, f' would also be visible from w, and we restart our argumentation with f replaced by f'.

We now consider two cases:

- If f' lies in $\ell^+(f)$ then $s(f)$ would not extend all the way to q, so this case is impossible.
- If f' lies in $\ell^-(f)$, then we consider two sub-cases: (i) if the boundary of P does not intersect the interior of the triangle formed by w', q, and v, then w' sees past the reflex vertex v; (ii) otherwise, if the boundary of P does intersect the interior of the triangle formed by w', q, and v, then there is a reflex vertex v' on the shortest geodesic path from w' to v (which lies completely in the aforementioned triangle), and w' sees past this vertex v'. Both sub-cases contradict Lemma 12.

Since both cases are not possible, we conclude that there cannot be a witness w' on e. ◻

Lemma 15 *Let P be a simple polygon, and let W be a finite minimal witness set for P. If $w \in W$ lies on an edge e of P, then w witnesses the whole edge e.*

Proof. If w does not witness e completely, there must be at least one other witness w' that witnesses a part of e. By Lemma 13, w' cannot lie on e. We know that w' witnesses at least one point on the interior of e, since the part of e that does *not* lie in $\mathrm{VK}(w)$ cannot consist of only one or both vertices of e (because visibility kernels are closed regions). But then by Lemma 14, there cannot be a witness that lies on e – and this contradicts the fact that w lies on e. ◻

Note that the situation depicted in Figure 5, where the points p and p' lie in between the vertices v and v', cannot occur if W is a finite witness set. In fact, a witness w can be placed anywhere on e, except on its reflex vertices, if any. A witness on a reflex vertex does not witness the edges incident to the vertex, so by the above lemma, a finite minimal witness set does not contain any witnesses that lie on reflex vertices.

Lemma 16 *For any $n \geq 4$ there exists a simple polygon P with n vertices, such that P is witnessable by a finite witness set, and such that the* minimum *size witness set has $n - 2$ elements.*

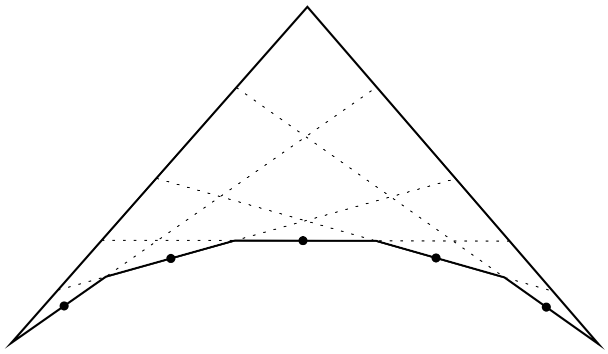

Figure 7: For any n there is a polygon with n vertices that is witnessable with no less than $n-2$ witnesses. Witnesses in the example are indicated with black dots.

Proof. The lowerbound construction is given in Figure 7. It is straightforward to verify that any edge on the bottom chain can only be witnessed by a witness in its interior (the leftmost and rightmost witnesses may also be moved to the leftmost and rightmost incident vertices of their edges, respectively). ◻

Combining Lemmas 11, 13, 15, and 16 we derive the following theorem:

Theorem 2 *Let P be a simple polygon with n edges. If a finite minimal witness set W for P exists, then all witnesses $w \in W$ lie on the boundary of P, and there is at most one witness $w \in W$ on each edge. Furthermore, a witness $w \in W$ on an edge e of P must witness e completely, as no other witness $w' \in W$ can witness any point on e. Finally, for each $n \geq 4$ there is a polygon that needs no less than $n-2$ witnesses to be witnessed.*

6 Algorithms

In this section we outline an algorithm that computes a minimal finite witness set W for a simple polygon P, if such a set exists, or reports the non-existence of such a set otherwise. We also outline two algorithms that use a witness set W for P to test whether a set of points G in P guards the whole polygon.

The algorithm to compute a minimal witness set W for a given simple polygon P with n vertices works as follows:

- First, we place witnesses at candidate positions. From Theorem 2, we know that a minimal witness set for P (if it exists) has at most one witness on every edge of P (but there is not necessarily a witness on *every* edge of P). We also know that a witness can be placed anywhere on an edge, except on a reflex vertex. Lemma 5 tells us that a superset W' of a witness set W is also a witness set (but unless $W' = W$, the superset is not a minimal witness set).

  We compute a candidate set W' as follows: we place a witness at every convex vertex of P, and one halfway each edge of P. If W' is a witness set for P, it is certainly not a minimal one. Any witness halfway an edge that has an incident convex vertex, is, by Lemma 15, already witnessed by that vertex. However, if a minimal witness set W exists, then W' is a superset of W. Note that there are smarter ways to create a smaller candidate set W', but these are more complicated to describe, do not improve the overall running time, and do not guarantee smaller minimal witness sets.

This step runs in $O(n)$ time.

- Next, we test whether W' is a witness set for P. If this is not the case, then no witness set for P exists. Otherwise, we extract a minimal witness set W from W' in the next step.

  Determining whether W' is a witness set for P is done as follows: first, we compute the visibility polygons $\mathrm{VP}(w)$ for each witness $w \in W'$. This can be done in $O(n)$ time per witness [7], and takes $O(n^2)$ time overall. The resulting polygons have linear size each. Next, from each visibility polygon $\mathrm{VP}(w)$, we compute its kernel $\mathrm{VK}(w)$. Again, this takes linear time per polygon [9] and $O(n^2)$ time overall. The resulting kernels have linear size each. Then we compute the arrangement induced by all visibility kernels and by P itself, using a sweepline algorithm [2]. During the sweep, we maintain for every cell in the arrangement whether it is covered by at least one visibility kernel, or only by P. In the latter case, we abort, and report that W' is not a witness set. Otherwise, if the algorithm is not aborted, W' is a witness set for P, and we proceed with the next step.

  A crude analysis of the last step tells us that since we have $O(n)$ polygons with $O(n)$ edges each, the total number of segments that form the arrangement is $O(n^2)$, and the complexity of the arrangement is therefore $O(n^4)$. However, recall from the discussion in Section 4 that a visibility kernel of a witness w is induced by $S_1(w)$ (the edges of P visible from w), and by $S_2'(w)$ (at most two reflex vertices of P visible from w). By Corollary 1, $\cup_{w \in W'} S_1(w)$ contains at most n edges, and there are also $O(n)$ segments derived from all sets $S_2'(w)$. So the arrangement has in fact quadratic complexity, and this step takes $O(n^2 \log n)$ time.

- If we didn't abort in the previous step, then we have a witness set W' for P at this point. We extract a minimal witness set W by repeatedly removing an unnecessary witness as follows: for each element w of W, we test if it lies in $\mathrm{VK}(w')$ for some $w' \in W$ that has not yet been removed. If it does, we remove w, for otherwise the resulting set would not be a minimal witness set by Lemma 4. Note that if w' would also be removed in a later iteration, then this would be because w' lies in some visibility kernel $\mathrm{VK}(w'')$, and by the transitivity property (Lemma 3), w would then also be contained in $\mathrm{VK}(w'')$. The resulting set is a minimal witness set W.

  After $O(n \log n)$ preprocessing for each visibility kernel, which amounts to $O(n^2 \log n)$ in total, each containment query can be done in $O(\log n)$ time [11]. Since we test a linear number of witnesses against a linear number of visibility kernels, the total query time is $O(n^2 \log n)$.

This leads to the following theorem:

Theorem 3 *Let P be a simple polygon with n vertices. If a finite witness set for P exists, a finite minimal such set W can be computed in $O(n^2 \log n)$ time. Otherwise, if no finite witness set for P exists, than this can be reported in the same running time.*

Next, we need an algorithm for testing whether a set G of g guards in a polygon P together see the whole polygon.

A straightforward way, without using witnesses, would be to compute the visibility polygons of the g guards in $O(gn)$ total time [7]. Their union has complexity $O(gn+n^2)$, and can be computed in $O((g^2+gn\log g)\log(g+n))$ time [3, 5]. Next, we test whether this union covers P; this can be done easily within the time bounds of the previous step, e.g., by comparing the areas of the polygons.

Can we do better if a witness set W of size m for P is given? One approach would be to test for each witness whether it can be seen by a guard by performing (at most) g ray shooting queries, or $O(gm)$ queries in total. P can be preprocessed for ray shooting in $O(n)$, after which a query takes $O(\log n)$ time [6]. So the total preprocessing time, including the computation of W, becomes $O(n^2\log n)$ time, and the query time is $O(gm\log n)$. Note that in the worst case $m=\theta(n)$. This query time is faster then the straightforward approach described above, but not necessarily very much (how much precisely depends on the parameters g and m). If m is small and g is large, then the gain is big.

Alternatively, we can compute the arrangement induced by the visibility polygons of the witnesses (note: *not* the visibility *kernels*), and preprocess it for point location. This arrangement has $O((mn)^2)$ complexity (which is $O(n^4)$ in the worst case), and it can be computed and preprocessed for point location in $O((mn)^2)\log mn)$ time [2, 11], after which point location takes $O(\log(n+m))=O(\log n)$ time per guard. If we store in each cell which visibility polygons cover the cell, then querying with g guards gets us g sets of size $O(m)$ in $O(g\log n)$ time. Then we must check whether these sets cover the set of witnesses; this can be done trivially in $O(gm)$ time. The total query time of this approach is slightly better than that of the ray shooting approach, but preprocessing time and memory requirements become much larger.

We conclude that the approaches that use witnesses are worthwhile if the polygon can be witnessed by a small number of witnesses, or if the number of guards is large.

7 Concluding remarks

We showed that if a polygon P admits a finite witness set, then any minimal witness set W for P has no witnesses in the interior of P, and there is at most one witness per edge of P. It follows that any minimal witness set for P has at most n elements. Furthermore, for any $n\geq 4$, there is a polygon for which the *minimum* size witness set has $n-2$ witnesses. A minimal finite witness set for P can be computed in $O(n^2\log n)$ time, if it exists.

It remains open whether the problem of finding a *minimum* size witness set for a given polygon is computable. It is well-known that the problem of finding a minimum size guard set is NP-complete [1, 8]. We conjecture, however, that finding a minimum size witness set is computable in polynomial time, and we are currently working towards turning our conjecture into a theorem.

Another interesting direction for further research is to consider other types of witnesses, such as (a subset of) the edges of the polygon. We believe that we can extend our current lemmas, theorems, and algorithms to test whether a polygon is witnessable by an minimal infinite witness set, where all the witnesses lie on the boundary of the polygon.

References

[1] Alok Aggarwal. *The art gallery problem: Its variations, applications, and algorithmic aspects*. PhD thesis, Dept. of Comput. Sci., Johns Hopkins University, Baltimore, MD, 1984.

[2] J. L. Bentley and T. A. Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Trans. Comput.*, C-28(9):643–647, September 1979.

[3] O. Cheong and R. van Oostrum. The visibility region of points in a simple polygon. In *Proc. 11th Canad. Comf. Comp. Geom.*, pages 87–90, 1999.

[4] V. Chvátal. A combinatorial theorem in plane geometry. *J. Combin. Theory Ser. B*, 18:39–41, 1975.

[5] L. Gewali, A. Meng, Joseph S. B. Mitchell, and S. Ntafos. Path planning in $0/1/\infty$ weighted regions with applications. *ORSA J. Comput.*, 2(3):253–272, 1990.

[6] J. Hershberger and Subhash Suri. A pedestrian approach to ray shooting: Shoot a ray, take a walk. *J. Algorithms*, 18:403–431, 1995.

[7] B. Joe and R. B. Simpson. Correction to Lee's visibility polygon algorithm. *BIT*, 27:458–473, 1987.

[8] D. Lee and A. Lin. Computational complexity of art gallery problems. *IEEE Trans. Inform. Theory*, 32(2):276–282, 1986.

[9] D. T. Lee and F. P. Preparata. An optimal algorithm for finding the kernel of a polygon. *J. ACM*, 26(3):415–421, July 1979.

[10] D. T. Lee and F. P. Preparata. Euclidean shortest paths in the presence of rectilinear barriers. *Networks*, 14:393–410, 1984.

[11] K. Mulmuley. A fast planar partition algorithm, I. *J. Symbolic Comput.*, 10(3-4):253–280, 1990.

[12] J. O'Rourke. *Art Gallery Theorems and Algorithms*. The International Series of Monographs on Computer Science. Oxford University Press, New York, NY, 1987.

[13] T. C. Shermer. Recent results in art galleries. *Proc. IEEE*, 80(9):1384–1399, September 1992.

[14] Jorge Urrutia. Art gallery and illumination problems. In Jörg-Rüdiger Sack and Jorge Urrutia, editors, *Handbook of Computational Geometry*, pages 973–1027. North-Holland, 2000.

[15] T.-C. Yang and C.-S. Shin. Guard sufficiency set for polygons. *Journal of Korean Information Science and Technology*, 28(1–2):73–79, 2001.