Guarding Art Galleries by Guarding Witnesses (Extended Abstract)

Kyung-Yong Chwa[1], Byung-Cheol Jo[2], Christian Knauer[3], Esther Moet[4], René van Oostrum[4], and Chan-Su Shin[5]

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

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. 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 in the interior of every edge. We give an algorithm to compute a minimum 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.

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 [1] 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 1987 O'Rourke published a whole book on the topic [2], while many new results that had been achieved after the publication of O'Rourke's book have appeared in surveys by Shermer [3] and Urrutia [4].

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.).

R. Fleischer and G. Trippen (Eds.): ISAAC 2004, LNCS 3341, pp. 352–363, 2004.

Except for the case of a single guard, only few papers consider the computation of the guarded region for a *given* set of guards [5, 6].

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 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 [7] (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 *minimum* witness set. The main contribution of our paper is the combinatorial/geometrical result that a finite minimum size 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. Examples of such polygons can be found in the full version of this paper, and in a preliminary version [8]. Note that in this preliminary version we only give results concerning minimal (but not necessarily minimum size) witness sets.

For completeness, we also outline an algorithm to compute a minimum size witness set for a polygon if such a set exists, or to report the non-existence otherwise. The running time of this algorithm is $O(n^2 \log n)$.

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. Using the results from Section 2, we study properties of *finite* witness sets in Section 3. In Section 4, we discuss the cardinality of minimum size witness sets. In Section 5 we give an algorithm for computing minimum size witness sets, and we wrap up in Section 6 with a brief discussion on our results and on open problems.

2 Preliminaries

Throughout this paper, P denotes a simple polygon with n vertices. We assume that P's vertices $V(P) = \{v_0, v_1, \ldots, v_{n-1}\}$ are ordered in counterclockwise direction. We also assume that the vertices of P are in general position, that is, no three vertices of P are collinear[1]. 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 [2]. 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 e be an edge e of P. Let $\ell(e)$ be the directed line through e 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 $cl(R)$.

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 vertex v of P, $\ell(p, v)$ is defined as the line through p and v, and $\ell(p, v) \cap P$ 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 $s(p, v)$.

The *kernel* of a polygon P is the set of points from which every point in P is visible. If the kernel is nonempty, we call P *star-shaped*. It is well-known that the kernel of a polygon P can be described as $\bigcap_{e \in E(P)} cl(\ell^+(e))$.

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 $\mathrm{VP}(p)$. We define the *visibility kernel* of a point p to be the kernel of its visibility polygon and we denoted it by $\mathrm{VK}(p)$. The visibility polygon $\mathrm{VP}(p)$ of a point p in P is star-shaped by definition (the kernel contains at least 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 non-empty 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*. Note that a guard set always is nonempty, but that a witness set is allowed to be empty (namely, the empty set is a witness set for any convex polygon).

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

[1] This assumption is only used in the proof of Lemma 17.

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. ◻

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*. A set W is called a *minimal witness set* for P if W is a witness set for P and, for any $w \in W$, $W \setminus \{w\}$ is *not* a witness set for P. The proofs of the lemmas in the rest of this section are straightforward, and we omit them here. The interested reader can find them in a preliminary version of this paper [8], as well as in the full version.

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.*

Lemma 6. *For a star-shaped polygon P and any point p in P,* $\mathrm{VK}(p)$ *contains the kernel of P.*

Lemma 7. *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 $v \in V(P)$.*

Lemma 8. *If a point p in a polygon P sees past a reflex vertex $v \in V(P)$, then p lies on the boundary of* $\mathrm{VK}(p)$*.*

Lemma 9. *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.*

By combining Lemmas 4, 7 and 8, we get the following lemma:

Lemma 10. *Let P be a polygon. If W is a minimal witness set for P, with $|W| > 1$, and w is an element of W, then w lies on the boundary of* $\mathrm{VK}(w)$*.*

3 Finite Witness Sets

In 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 in its interior (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). 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 9, 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 $\mathcal{Q}(W \setminus \{w\}) \cap \partial P$. By Lemma 10, w lies on the boundary of $\mathrm{VK}(w)$. Since w cannot lie simultaneously on the boundary of $\mathrm{VK}(w)$ (which is closed) and in the open set $int(\mathcal{Q}(W \setminus \{w\}))$, 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 W 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 in its interior.

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 sees past any reflex vertex v of P, then w must lie on an edge e of P such that v lies on $\ell(e)$.*

Proof. Since W is a finite witness minimal witness set for P, any $w \in W$ lies on some edge e of P by Lemma 11. Note that w lies on two such edges if it lies on a vertex. Now 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 there is at least one reflex vertex v such that w sees past v and $\mathrm{VK}(w)$ is contained in $cl(\ell^+(w,v))$.

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 (i) w lies on the intersection of e and $\ell(w,v)$, (ii) $\mathrm{VK}(w)$ is contained in $cl(\ell^+(w,v))$, and (iii) w sees past v, w must be the endpoint of one of the parts of $\mathrm{VK}(w) \cap \partial P$. Since W is a minimal witness set, by Lemma 4, w must also lie in $\mathcal{Q}(W \setminus \{w\})$. But w cannot lie simultaneously in an open subset and on the boundary thereof, so 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 7. A witness w on an edge e sees the vertices of P that lie on $\ell(e)$, and combining Lemmas 7 and 12 implies that w must see past at least one of these vertices.

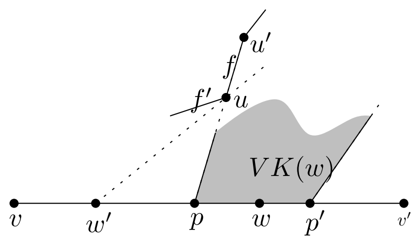

Fig. 1. Illustration to the proof of Lemma 13. Any point w' in between v and p sees past u or (not shown) past another reflex vertex in the triangle formed by w', u, and p

Lemma 13. *Let P be a simple polygon, and let W be a finite minimal witness set for P. If an edge e of P contains a witness $w \in W$ in its interior, then no other witness $w' \in W$ can lie on e.*

Proof. Let e be an edge of P that contains a witness $w \in W$ in its interior. 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 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 1. Since p lies on the boundary of $\mathrm{VK}(w)$, p must be the intersection of e and $s(f)$ (the extension of an edge f, as defined in Section 2) 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, for the sake of contradiction, suppose that there is a witness $w' \in W$ that lies in the half-open segment from (and including) v to (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. □

We omit the proof of the following lemma; it can be found in a preliminary version of this paper [8] and in the full version. The main arguments use

Lemma 12. The lemma itself is used to prove Lemma 15, which in turn is applied in Section 4.3, where we investigate the possible locations of the witnesses on the boundary of P.

Lemma 14. *Let P be a simple polygon, and let W be a finite minimal 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.*

Lemma 15. *Let P be a simple polygon, and let W be a finite minimal witness set for P. If $w \in W$ lies in the interior of 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 1, where the points p and p' lie in between the vertices v and v', cannot occur if W is a finite witness set.

Lemma 16. *Let P be a simple polygon and let W be a finite minimal witness set for P. Then no element of W lies on a reflex vertex of P.*

Proof. Suppose, for the sake of contradiction, that there is an element w of W that lies on a reflex vertex v of P. Observe that $\mathrm{VK}(w)$ does not contain any points on the edges e and e' incident to v, except v itself. This means that all points on e and e' (except v) must be witnessed by the elements of $W \setminus \{w\}$. However, the union of the visibility kernels of these witnesses is a closed region, since W is finite, and therefore v is also covered by these visibility kernels. This means that w is witnessed by another witness $w' \in W$, and this contradicts the minimality of W. □

4 Minimum Size Witness Sets

In this section we classify the edges of P into three categories, depending on whether their incident vertices are convex or not: we distinguish reflex-reflex edges, convex-reflex (or reflex-convex) edges, and convex-convex edges. In Section 4.1 we show that each reflex-reflex and convex-reflex edge of P contains exactly one witness, and in the following section we show that no witness lies on a convex-convex edge, except possibly at its vertices that are also incident to a convex-reflex edge. In Section 4.3 finally, we show *where* to place witnesses on reflex-reflex and convex-reflex edges, namely: for each reflex-reflex edge we can place a witness anywhere in its interior, and for each convex-reflex edge we can place a witness on its convex edge. This placement strategy establishes a minimum size witness set for P—if a finite witness set for P exists. Testing whether the candidate witness set is indeed a witness set is the subject of Section 5.

4.1 Reflex-Reflex and Convex-Reflex Edges

Lemma 17. *Let P be a simple polygon, let W be a finite minimal witness set for P and let e be an edge of P that is incident to at least one reflex vertex. Then there is exactly one witness $w \in W$ located on e.*

Proof. Let v be a reflex vertex of e. There must be a witness w of W that witnesses a point p in the interior of e, as well as v itself, since $\mathrm{VK}(w)$ is a closed region, and we have only finitely many witnesses.

We first show that w must lie on $s(e)$ (recall from Section 2 that $s(e)$ is the single piece of $\ell(e)$ that lies in $int(P)$ and that contains e). Suppose, for the sake of contradiction, w does not lie on $s(e)$. Obviously, both v and p are visible from w, so w must lie in $\ell^+(e)$, and by Lemma 11 w lies on the boundary of P. We have two cases to consider:

- w sees past v. However, according to 12, this is only allowed if v lies on $\ell(f)$, where f is the edge on which w lies. However, f cannot be incident to v (otherwise v would not be reflex), so this violates the assumption made in Section 2 that no three vertices of P are collinear.
- w does not see past v. Since w sees v, we have that w sees both edges e and e' incident to v and $\mathrm{VK}(w)$ is completely contained in $cl(\ell^+(e)) \cap cl(\ell^+(e'))$. But since v is reflex, this means that the only point on e that is contained in $\mathrm{VK}(w)$ is v, contradicting the fact that w also witnesses p.

Since both cases lead to a contradiction, we conclude that w lies on $s(e)$.

Since w lies on the boundary of P by Lemma 11, w lies either on e or on one the endpoints of $s(e)$. Let q be an endpoint of $s(e)$. q is either a convex vertex of e, or it does not lie on e. In the latter case, q sees a reflex vertex v' of e as well as both edges e and e' incident to v' (v' may or may not be v). Were w located on q in that case, then we argue again that the only point on e that is contained in $\mathrm{VK}(w)$ is v', contradicting the fact that w also witnesses p. We conclude that w lies on e.

By Lemma 16 w does not lie on a reflex vertex of e, and therefore we consider the following two cases:

- w lies in the interior of e. By Lemma 13, no other witness besides w can lie on e;
- w lies on the convex vertex of e (this case is only applicable if e is a convex-reflex edge). By Lemma 13, no other witness can lie in $int(e)$. By Lemma 16, the remaining reflex vertex of e also cannot contain a witness, so w is the only witness on e.

In both cases, w is the only witness on e and thus the proof is completed. □

4.2 Convex-Convex Edges

No witness of a minimal witness set is located on a convex-convex edge, except possible at a vertex that is also incident to a convex-reflex edge:

Lemma 18. *Let P be a simple polygon and let W be a finite minimal witness set for P. Then no element of W is located in the interior of a convex-convex edge or on a vertex incident to two convex-convex edges.*

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

W cannot have only one element. Otherwise, by Lemma 9, 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, suppose for the sake of contradiction, that there is a witness $w \in W$ that lies either in the interior of a convex-convex edge e, or on a convex vertex that is shared by two convex-convex edges e and e'. There is at least one other witness $w' \in W$, and w lies outside $\mathrm{VK}(w')$; otherwise, w' would witness w, and this contradicts the minimality of W. By Lemma 7 w sees past at least one reflex vertex v of P, and by Lemma 12 v lies on $\ell(e)$ (or on $\ell(e')$). However, all vertices incident to e (and e') are convex, so this is impossible. ⊓⊔

4.3 Placement of Witnesses

By Lemma 17, every finite minimal witness set W for P contains exactly one element w located on a convex-reflex or reflex-reflex edge. The following lemma shows to be useful in constructing a minimum size witness set for a polygon.

Lemma 19. *Let P be a simple polygon, and let W be a finite minimal witness set for P. Furthermore, let e be a convex-reflex or reflex-reflex vertex of P, and let w be the witness located on e. If w lies in $int(e)$, then $\{w'\} \cup W \setminus \{w\}$ is also a witness set for P, where $w' \notin W$ is an replacement witness placed anywhere in $int(e)$ or on the convex vertex of e (if any).*

Proof. If w' and w coincide, the lemma follows trivially. If on the other hand w and w' lie on different points on e, then we argue that w' witnesses w, and therefore, w' can replace w, as by Lemma 3, any point that is witnessed by w is also witnessed by w'.

Suppose, for the sake of contradiction, that w' does not witness w, or in other words, that w does not lie in $\mathrm{VK}(w')$. This means that the boundary of $\mathrm{VK}(w')$ intersects e in a point p that lies in between w and w'. This point p is the endpoint of some segment $s(f)$, (partially) visible from w', but not visible from w. Let u be the vertex of f closest to p. We have that u is a reflex vertex, and moreover, u lies in $\ell^+(e)$. Now, if u is visible from w, then w sees past u. Otherwise, if u is not visible from w, 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 triangleformed by w, u, and p). In either case, we derive a contradiction with lemma 12, and the conclusion is that w' indeed witnesses w. ⊓⊔

Combining the lemmas in Sections 3 and 4 we derive our main theorem:

Theorem 2. *Let P be a simple polygon. If P is witnessable by a finite witness set, then a* minimum size *witness set for P can be constructed by placing a witness anywhere in in the interior of every reflex-reflex edge and by placing a witness on the convex vertex of every convex-reflex edge.*

Furthermore, we have that for every $n \geq 4$ there is a polygon that is witnessable with no less than $n - 2$ witnesses, as is illustrated in Figure 2.

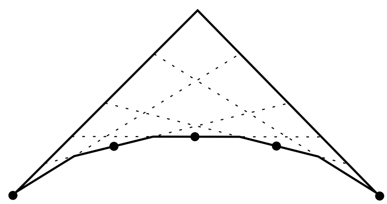

Fig. 2. For any $n \geq 4$ there is a polygon with n vertices that is witnessable with no less than $n - 2$ witnesses

5 Algorithm

In this section, we outline an algorithm that computes a minimum size witness set W for a simple polygon P with n vertices, if such a set exists, or reports the non-existence of such a set otherwise. The idea is straightforward: we place witnesses according to Theorem 2, giving W, and check whether the union of the visibility kernels of the witnesses covers the whole polygon. This can be done by computing the arrangement induced by the visibility kernels and by P itself, using a sweepline algorithm [9]. 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 report that P is not witnessable by a finite witness set. Otherwise, we report that W is a witnessset for P.

Since each visibility kernel has $O(n)$ edges and we have $O(n)$ witnesses, we have $O(n^2)$ edges in total, and the arrangement has $O(n^4)$ complexity. However, we can show that every visibility kernel has at most two edges that are *not* (part of the) edges of P. This leads to an arrangement of $O(n^2)$ complexity, and gives a running time of $O(n^2 \log n)$ for our algorithm. Details can be found in the full version of this paper, as well as in a preliminary version [8].

6 Concluding Remarks

We showed that if a polygon P admits a finite witness set, then no minimal witness set W for P has any witnesses in the interior of P, and a minimum size witness set for P can be constructed in linear time. If it is not known in advance whether P is witnessable with a finite witness set, then checking whether the constructed (candidate) witness set witnesses the polygon can be done in $O(n^2 \log n)$ time.

An 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. Chvátal, V.: A combinatorial theorem in plane geometry. J. Combin. Theory Ser. B **18** (1975) 39–41
2. O'Rourke, J.: Art Gallery Theorems and Algorithms. The International Series of Monographs on Computer Science. Oxford University Press, New York, NY (1987)
3. Shermer, T.C.: Recent results in art galleries. Proc. IEEE **80** (1992) 1384–1399
4. Urrutia, J.: Art gallery and illumination problems. In Sack, J.R., Urrutia, J., eds.: Handbook of Computational Geometry. North-Holland (2000) 973–1027
5. Cheong, O., van Oostrum, R.: The visibility region of points in a simple polygon. In: Proc. 11th Canad. Comf. Comp. Geom. (1999) 87–90
6. Gewali, L., Meng, A., Mitchell, J.S.B., Ntafos, S.: Path planning in $0/1/\infty$ weighted regions with applications. ORSA J. Comput. **2** (1990) 253–272
7. Yang, T.C., Shin, C.S.: Guard sufficiency set for polygons. Journal of Korean Information Science and Technology **28** (2001) 73–79
8. Chwa, K.Y., Jo, B.C., Knauer, C., Moet, E., van Oostrum, R., Shin, C.S.: Guarding art galleries by guarding witnesses. Report UU-CS-2003-044, Institute for Information & Computing Sciences, Utrecht University, Utrecht, Netherlands (2003)
9. Bentley, J.L., Ottmann, T.A.: Algorithms for reporting and counting geometric intersections. IEEE Trans. Comput. **C-28** (1979) 643–647
10. Joe, B., Simpson, R.B.: Correction to Lee's visibility polygon algorithm. BIT **27** (1987) 458–473
11. Hershberger, J., Suri, S.: A pedestrian approach to ray shooting: Shoot a ray, take a walk. J. Algorithms **18** (1995) 403–431
12. Mulmuley, K.: A fast planar partition algorithm, I. J. Symbolic Comput. **10** (1990) 253–280