Guarding Art Galleries by Guarding Witnesses

Kyung-Yong Chwa [a] Byung-Cheol Jo [b] Christian Knauer [c] Esther Moet [d]
René van Oostrum [d] Chan-Su Shin [e]

[a] *Department of Computer Science, KAIST, Daejon, Korea. Email: kychwa@tclab.kaist.ac.kr*
[b] *Taff System, Co. Ltd., Seoul, Korea. Email: mrjo@taff.co.kr*
[c] *Freie Universität Berlin, Takustraße 9, D-14195 Berlin, Germany. Email: knauer@inf.fu-berlin.de*
[d] *Institute of Information & Computing Sciences, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands. Email: {esther,rene}@cs.uu.nl*
[e] *School of Electronics and Information Engineering, Hankuk University of Foreign Studies, Yongin, Korea. Email: 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. 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 in the interior of any 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 an algorithm that uses a witness set for P to test whether a (prospective) guard set sees all points in P.

1. Introduction

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 the whole polygon?"

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.

A preliminary full version of this paper is available as technical report [3]. Due to space limitations, we omitted several lemmas of minor importance and the proofs of the remaining lemmas in this abstract.

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 $e_i = (v_i, v_{i+1 \bmod n})$. We consider an edge e_i to be the closed line segment between its incident vertices, and 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 [7].

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

Theorem 2 *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.*

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 2:

Lemma 3 *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 lemma shows that witnessing is transitive:

Lemma 4 *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.*

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

Definition 5 *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.*

Lemma 6 *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 7 *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.*

Observe that not all polygons are witnessable with a finite witness set; see Figure 1. The polygon on the left is witnessable by three witnesses (the black dots), but the polygon on the right needs an infinite number of witnesses. The visibility kernels of the witnesses indicated at four of the vertices of the polygon do not cover the complete polygon.

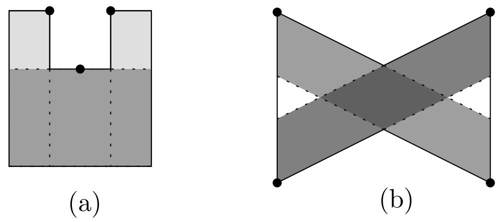

Fig. 1. The polygon on the left is witnessable with three witnesses, while the polygon on the right needs an infinite number of witnesses.

Adding witnesses to the remaining vertices does not help, as these vertices are already witnessed and witnessing is transitive. It turns out that we would need to cover both unwitnessed segments on the boundary of the polygon completely with witnesses to get an (infinite) witness set for this polygon.

3. 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. It is well-known that the kernel of a polygon P is the intersection of the positive halfspaces of its edges; 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). 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 (see Figure 2):

(i) An edge e of $\mathrm{VP}(p)$ coincides with the part of an edge e' of P that is visible from p.

(ii) 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 visibility kernel $\mathrm{VK}(p)$ is the intersection of the closure of the positive half-spaces induced by the lines through all edges of $\mathrm{VP}(p)$. The reader may wonder why we introduce this seemingly complicated alternative representation of the edges of $\mathrm{VP}(p)$. The reason is that for any p, there may be many edges in group (ii), but at most two of these contribute to $\mathrm{VK}(p)$. This helps us to reduce the complexity of the data structures involved in com-

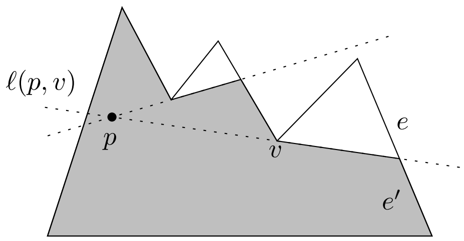

Fig. 2. Two types of edges of VP(P).

puting the union of a set of visibility kernels; see Section 5.

We conclude this section with a property of visibility kernels that is of use in the remainder of this paper.

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* VK(p).

4. 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.

We have seen that a point p witnesses a point q if q lies in VK(p). This means that for a witness set W, the union $\bigcup_{w \in W} \text{VK}(w)$ must cover the whole polygon p.

For a polygon P and a set W of points in P, let $\mathcal{A}(W)$ be the arrangement in P induced by the supporting lines of the line segments of type (i) and of those of type (ii) that contribute to VK(w), for every witness $w \in W$.

For any cell c of $\mathcal{A}(W)$ and any point $w \in W$, c lies either completely inside or completely outside 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 VK(w) for at least one $w \in W$.

We denote the cardinality of W by m. Because for every $w \in W$, there are at most two vertices in group (ii) that contribute to VK(w), there are in total at most $n + 2m$ line segments that define $\mathcal{A}(W)$, and therefore 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 5.

Via several lemmas that are derived from Lemmas 4 and 8, we arrive at the following lemma:

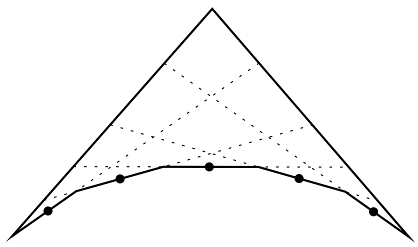

Fig. 3. For any n there is a polygon with n vertices that is witnessable with no less than $n - 2$ witnesses. Witnesses in the example are indicated with black dots.

Lemma 9 *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).*

Note that a convex polygon can be witnessed by a single point in its interior. However, such a one-element witness set is not minimal, as the empty set is also a witness set for any convex polygon.

Given the above lemma, we only need to concentrate on witnesses that lie on the boundary of P. Analyzing the possible configurations of witness sets, we arrive at the following theorem:

Theorem 10 *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. Each edge has at most one witness $w \in W$ in its interior. If an edge has one or two witnesses on its incident vertices, then there cannot be any witness in the interior of the edge. Finally, for each $n \geq 4$ there is a polygon that needs no less than $n - 2$ witnesses to be witnessed.*

The lowerbound construction is given in Figure 3.

5. 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 an algorithm that uses 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. We place a witness at every vertex of P, and one halfway each edge of P. This step runs in $O(n)$ time.
- Next, using $\mathcal{A}(W)$, we test whether W' is a witness set for P with a sweepline approach. If it is, we extract a minimal witness set W from W' in the next step; otherwise, we report that no finite witness set for P exists. This step takes $O(n^2 \log n)$ time.
- We extract a minimal witness set W by repeatedly removing an unnecessary witness, i.e. witnesses that are witnessed by another witness in W'. This step takes $O(n^2 \log n)$ time.

This leads to the following theorem:

Theorem 11 *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 check, without using witnesses, can be performed in $O((g^2 + gn \log g) \log(g + n))$ time [2,4].

Can we do better if a witness set W of size m for P is given? We 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 [5]. 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 (the number of witnesses) is small and g (the number of guards) is large, then the gain is big.

6. 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, there is at most one witness in the interior of each edge of P. If an edge has one or two witnesses on its incident vertices, then there cannot be a witness in the interior of this edge. 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 [1,6]. 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] 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.

[3] Kyung-Yong Chwa, Byung-Cheol Jo, Christian Knauer, Esther Moet, René van Oostrum, and Chan-Su Shin. Guarding art galleries by guarding witnesses. Report UU-CS-2003-044, Institute for Information & Computing Sciences, Utrecht University, Utrecht, Netherlands, 2003.

[4] 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.

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

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

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