

Allocating Vertex π -Guards in Simple Polygons via Pseudo-Triangulations*

Bettina Speckmann¹ and Csaba D. Tóth²

¹Department of Mathematics and Computer Science, TU Eindhoven,
P.O. Box 513, Eindhoven 5600 MB, The Netherlands
speckman@win.tue.nl

²Department of Computer Science, University of California at Santa Barbara,
Santa Barbara, CA 93106, USA
toth@cs.ucsb.edu

Abstract. We use the concept of pointed pseudo-triangulations to establish new upper and lower bounds on a well known problem from the area of art galleries: What is the worst case optimal number of vertex π -guards that collectively monitor a simple polygon with n vertices? Our results are as follows:

1. Any simple polygon with n vertices can be monitored by at most $\lfloor n/2 \rfloor$ *general* vertex π -guards. This bound is tight up to an additive constant of 1.
2. Any simple polygon with n vertices, k of which are convex, can be monitored by at most $\lfloor (2n - k)/3 \rfloor$ *edge-aligned* vertex π -guards. This is the first non-trivial upper bound for this problem and it is tight for the worst case families of polygons known so far.

1. Introduction

Pseudo-triangulations, also called geodesic triangulations, have received considerable attention in the last few years. They were originally introduced in the context of visibility complexes [13], [14] and ray shooting [4], [8]. However, only their recent applications to robot arm motion planning [19] and kinetic collision detection [1], [11], as well as the identification of *pointed pseudo-triangulations* by Streinu [19], initiated growing interest in their combinatorial and geometric properties (see also [18]).

* Research on this paper was supported by the Berlin–Zürich European Graduate Program “Combinatorics, Geometry, and Computation”, financed by the German Science Foundation (DFG) and ETH Zürich.

Art Galleries and *illumination* problems represent an old and popular field [15] of computational geometry intimately related to planar subdivisions. A typical art gallery theorem provides combinatorial bounds on the number of guards needed to cover a space visually [17], [20], [24]. Usually, a partition of the space into convex regions, each of which is visible by one guard, allows one to use a combinatorial argument on a finite set of possible guard locations.

In contrast, our allocation of guards is based on a decomposition of a simple polygon into *pseudo-triangles*—planar polygons with exactly three convex vertices. Given the inherently reflex nature of pseudo-triangles it seems somewhat surprising that they effectively support the placement of guards. In fact, our results confirm observations made by previous applications: pseudo-triangulations provide sparse tessellations of space that uniquely capture its geometric properties. We regard it as the main contribution of this paper to establish pseudo-triangulations as an efficient technique for attacking not only geometric but also combinatorial problems.

Art Galleries. The first theorem on *Art Galleries* is due to Chvátal [5] who showed that any simple polygon with n vertices can be monitored by $\lfloor n/3 \rfloor$ point guards; this bound is tight. The famous proof of Chvátal's theorem by Fisk [7] places point guards at vertices of the polygon (i.e., uses *vertex guards*). Recently it has been shown in [21] that $\lfloor n/3 \rfloor$ guards are always sufficient, even if the range of vision of each point guard is restricted to π (such guards are called *point π -guards*). Point π -guards may be placed at any point in the polygon, even allowing two π -guards to be placed at the same point.

Urrutia [6] asked the following question (which was restated in [17], [24], and [12]): What is the minimal number of vertex π -guards that can collectively monitor any simple polygon P with n vertices? A vertex π -guard is given by a pair (v, H^v) where v is a vertex of P and H^v is a closed half-plane such that v is on the boundary of H^v . There may be at most one π -guard at each vertex of P . A π -guard (v, H^v) monitors $a \in P$ if and only if the closed line segment va is in $P \cap H^v$.

The notion of a vertex α -guard can be defined for any angle $0 < \alpha < 2\pi$ as a cone of aperture at most α with the apex at a vertex of polygon P . Under the condition that there may be at most one vertex guard at each vertex, it is known [6] that for any angle $\alpha < \pi$, there exists a constant N_α such that for any $n, n \geq N_\alpha$, there is a convex polygon $P_{n,\alpha}$ with n vertices such that n α -guards cannot monitor $P_{n,\alpha}$.

A π -guard at a reflex vertex cannot monitor the complete angular domain bounded by the two incident sides. Restrictions on the orientation of the half-plane H^v , where v is a reflex vertex, lead to three variants of the problem. We may require that at every reflex vertex v , the possible guard (v, H^v) be

- *inward-facing*: the two sides of P incident to v are disjoint from the interior of H^v (originally this variant was proposed in [6]);
- *edge-aligned*: inward-facing and the boundary of H^v is collinear with one of the sides incident to v ;
- *general*: no restriction on the orientation of H^v . (See Fig. 1 for illustrations.)

Previous Work and Results. Urrutia [25] found a family of polygons that require $\lfloor 5(n-1)/8 \rfloor$ inward-facing or edge-aligned vertex π -guards, but showed only the suf-




Fig. 1. A polygon with an inward-facing (left), an edge-aligned (middle), and an outward-facing (right) vertex π -guard.

iciency of $n - 2$ vertex π -guards [6]. In the restricted models of *inward-facing* or *edge-aligned* π -guards, this is still the best previously known upper bound [17]. For *general* π -guards the upper bound was improved to $\lfloor (3n - 5)/4 \rfloor$ using a so-called *dense decomposition* of the simple polygon [22]; later Brumberg et al. [2] claimed a simple proof for an upper bound of $\lfloor 5n/6 \rfloor$. The best known lower bound, $\lfloor (n - 1)/2 \rfloor$, for the general case follows from a family of polygons called *monotone mountains* [16] (polygons with two chains of reflex vertices).

In this paper we improve the upper bounds for all three models as well as the lower bound in the edge-aligned case. In particular we present the following results: First, $\lfloor n/2 \rfloor$ vertex π -guards can always monitor a simple polygon with n vertices in the *general* model—this bound is tight for n odd, and tight up to an additive constant of one for n even. Second, $\lfloor 2n/3 \rfloor - 1$ edge-aligned vertex π -guards can always monitor a simple polygon with n vertices. Third, we construct a new family of polygons that requires $9n/14 - O(1)$ inward-facing vertex π -guards (where $9/14 \approx 0.643$), generalizing the ideas of Santos and Urrutia and using the concept of pseudo-triangles.

In fact, we prove a stronger upper bound on the number of edge-aligned vertex π -guards for simple polygons with more than three convex vertices: $\lfloor (2n - k)/3 \rfloor$ vertex π -guards can always monitor a simple polygon with n vertices, k of which are convex. This bound is tight for the family of polygons of Urrutia where the number of vertices is $n = 8k - 15$. Our lower bound construction has $n = 14k - 27$ vertices and it requires $9k - 18$ inward-facing vertex π -guards which implies that our upper bound including k is tight for this family as well. This suggest that k is a natural parameter of the problem.

Proof Techniques. As we pointed out before, our allocation of guards is based on a pseudo-triangulation of the simple polygon P . A pseudo-triangulation decomposes P along non-crossing diagonals into pseudo-triangles. It is easy to see (see Section 3.1) that a pseudo-triangle with ℓ vertices can always be monitored with at most $\lfloor (2\ell - 3)/3 \rfloor$ vertex π -guards. Note though, that this does not immediately imply an upper bound of $2n/3 + O(1)$ for the total number of guards required by P : adjacent pseudo-triangles of the decomposition share two vertices and guard allocations in two pseudo-triangles

sharing a vertex v can prescribe contradicting orientations for a π -guard at v . As we will see in Section 3.2, this kind of conflict can be resolved by using a pointed pseudo-triangulation, where every vertex is either a convex vertex of P or a reflex vertex of one of its adjacent pseudo-triangles.

Our proofs are constructive and we can find a guard allocation within our bounds in linear time if we are given the pseudo-triangulation of the simple polygon. A pointed pseudo-triangulation (specifically, the *pulling* pseudo-triangulation that we use in Section 4) is provided by a shortest path tree of the convex vertices and can be constructed in linear time [9], [10] if the triangulation of the simple polygon [3] is given.

Organization. The next section presents some basic definitions and notation concerning pseudo-triangulations of simple polygons. In Section 3 we prove our upper bound on the number of edge-aligned vertex guards, followed by the proof of the bound on the number of general vertex guards in Section 4. We conclude the paper by describing our lower bound construction in Section 5.

2. Pseudo-Triangulations

A *pseudo-triangle* T is a simple polygon with exactly three convex vertices, called *corners*. If a , b , and c are the corners of T in counter-clockwise order, then we denote by \overline{ab} (and similarly by \overline{bc} and \overline{ca}) the chain of reflex vertices between a and b in counter-clockwise direction on the boundary of T . For a simple polygon P , a *pseudo-triangulation* is a set D of non-overlapping pseudo-triangles such that $P = \bigcup D$ and vertices of pseudo-triangles are vertices of P . Since triangles are also pseudo-triangles, any triangulation of a simple polygon is also a pseudo-triangulation.

In this paper we consider exclusively a special kind of pseudo-triangulation, namely *pointed pseudo-triangulations*. A pointed pseudo-triangulation is a pseudo-triangulation with the minimum number of pseudo-triangular faces, which is always the number of convex vertices of P minus two. Equivalently, pointed pseudo-triangulations have the property that every vertex of the polygon P is either convex in P or reflex in a pseudo-triangle $T \in D$. This property is also called *pointedness*. For the remainder of the paper we use the term pseudo-triangulation to refer to a pointed pseudo-triangulation.

Every simple polygon has at least one pseudo-triangulation, but this pseudo-triangulation is not necessarily unique. One possible method to obtain a pseudo-triangulation of a simple polygon is to add non-crossing diagonals as long as they do not violate the pointedness property. Another approach is to dissect P along non-crossing *geodesic paths* connecting convex vertices [23], [4]. Pseudo-triangulations are also referred to as *geodesic triangulations*, since every pseudo-triangulation can be obtained in this way. A pseudo-triangle in a pseudo-triangulation is uniquely determined by its three corners; hence we use the notation abc for a pseudo-triangle with corners a , b , and c .

The *dual graph* $G(D)$ of a pseudo-triangulation D is defined by choosing D as the node set of $G(D)$ and connecting two nodes by an edge iff the corresponding pseudo-triangles share a side (a diagonal of P). If D is a pseudo-triangulation of a simple polygon, then $G(D)$ is a tree (see Fig. 2). Note that in our terminology, polygons have *vertices* and *sides*; graphs have *nodes* and *edges*.




Fig. 2. A pseudo-triangulation and its dual graph for a simple polygon.

Let $x \in V(P)$ be a convex vertex of P and let $y \in V(P)$ be a vertex adjacent to x . We denote by $G(D)_{xy}$ the rooted tree $G(D)$ where the root R corresponds to the pseudo-triangle containing side xy . We write $S < T$ for two pseudo-triangles $S, T \in D_{xy}$ if T is the ancestor of S in $G(D)_{xy}$. For example, $S < R$ for all pseudo-triangles $S \neq R$. As a shorthand, we denote by D_{xy} the pseudo-triangulation D together with the partial order of ancestor relationships in the tree $G(D)_{xy}$.

For a reflex vertex v of a pseudo-triangle we consider only vertex π -guards which are *inward-facing* (as defined above) or *outward-facing*, that is, the sides of P incident to v are in the closed half-plane H^v (see also Fig. 1). For a corner v , we consider only two possibilities: either (v, H^v) completely covers the angular domain, or it does not cover it at all, that is, we ignore partial coverage.

3. Edge-Aligned π -Guards

To prove our upper bound on the number of edge-aligned π -vertex guards we first show how to guard a single pseudo-triangle with $\lfloor (2n - 3)/3 \rfloor$ vertex π -guards and then generalize this approach to a pseudo-triangulation of a simple polygon.

3.1. Guarding a Pseudo-Triangle

Proposition 1. *A pseudo-triangle T with ℓ vertices can be monitored by at most $\lfloor (2\ell - 3)/3 \rfloor$ vertex π -guards: one at a corner and at most $\lfloor 2(\ell - 3)/3 \rfloor$ outward-facing guards at reflex vertices along two chains.*

Proof. Observe that one guard at a corner and outward-facing guards along the two adjacent chains can monitor the complete pseudo-triangle. For the three corners, this gives three possible guard allocations jointly using a total of $3 + 2(\ell - 3) = 2\ell - 3$ guards. This implies that at least one guard allocation uses at most $\lfloor (2\ell - 3)/3 \rfloor$ guards (see Fig. 3). \square




Fig. 3. Guards at a corner and along two adjacent chains monitor the complete pseudo-triangle.

Using the same idea, we can deduce an upper bound of $\lfloor (\ell - 1)/2 \rfloor$ for so-called *monotone mountains*—defined by O’Rourke [16].

Proposition 2. *A pseudo-triangle abc with ℓ vertices and $\overline{bc} = \emptyset$ can be monitored by at most $\lfloor (\ell - 1)/2 \rfloor$ vertex π -guards: one at the corner b or c (but not at both) and at most $\lfloor (\ell - 3)/2 \rfloor$ outward-facing guards at reflex vertices along a chain different from \overline{bc} .*

Proof. Consider the two guard allocations which put one π -guard at corner b (resp., c) and outward-facing π -guards at vertices along the adjacent chain \overline{ba} (resp., \overline{ca}). Both monitor T and jointly use a total of $\ell - 1$ guards (see Fig. 3 (right)). \square

3.2. Guarding an Art Gallery

Theorem 3. $\lfloor (2n - k)/3 \rfloor$ edge-aligned vertex π -guards can always monitor a simple polygon with n vertices, k of which are convex.

Proof. Following, and generalizing the idea of Proposition 1, we describe three different guard allocations for P that jointly use a total of $2n - k$ edge-aligned vertex π -guards. We define these allocations with respect to an arbitrary but fixed pseudo-triangulation D_{xy} . In all three guard allocations and for each pseudo-triangle $T \in D_{xy}$, there will be vertex π -guards at one corner of T and along the two chains adjacent to that corner such that these general guards collectively monitor T . The three guard allocation schemes will jointly use all three corners (and thus different pairs of chains) for each $T \in D_{xy}$.

Consider the pseudo-triangle T_1 at the root of $G(D)_{xy}$. We can guard T_1 in three different ways by choosing any one of its three corners and placing guards at that corner (the *guarded corner* $gc(T_1)$ of T_1) and along its two adjacent chains (see Fig. 3). Every choice of guard allocation for T_1 induces a direction on its chains: the sides on the chains adjacent to $gc(T_1)$ are directed *away* from $gc(T_1)$, the sides on the remaining




Fig. 4. Propagating the side directions along $G(D)_{xy}$.

chain have *no* direction. Note that we do indeed distinguish between three directions for each side—the two standard directions and *no* direction. Fixing the direction for one side of a pseudo-triangle uniquely determines the directions for the remaining ones. Each side of the pseudo-triangulation can only have one direction, i.e., a pseudo-triangle inherits its directions from its parent via their joined side. Therefore, choosing a guarded corner $gc(T_1)$ in T_1 not only induces a unique direction on the sides of T_1 , but actually determines the direction of all sides of D_{xy} (see Fig. 4). Since every pseudo-triangle T now has exactly one corner with two outgoing directed sides, we define this corner to be the guarded corner $gc(T)$. Note that for each corner v of every $T \in D_{xy}$ there is a choice of the guarded corner of T_1 , such that the resulting orientation of D_{xy} implies that v is the guarded corner of T .

After choosing $gc(T_1)$ and propagating the directions along $G(D)_{xy}$, we place vertex π -guards along the directed sides of D_{xy} . More specifically, in each $T \in D_{xy}$ we place a vertex π -guard at the corner $gc(T)$ and *outward-facing* π -guards along the two directed chains adjacent to $gc(T)$ (see Fig. 5). By Proposition 1, these general guards collectively monitor T . Note that it is always sufficient to place at most one π -guard at each vertex since every vertex is a reflex vertex of at most one pseudo-triangle of D_{xy} and π -guards at reflex vertices are outward-facing. (This already gives an upper bound of $\lfloor (2n - k)/3 \rfloor$ in the *general* model.)

Next we align every vertex π -guard (v, H^v) at every reflex vertex v with one of the sides of P incident to v . Recall that v is a reflex vertex of exactly one $T_v \in D_{xy}$. We placed a guard at v only if v is incident to a side s_{in} of T_v directed to v and a side s_{out} of T_v directed away from v (i.e., v lies on a chain adjacent to $gc(T_v)$). Let D_{xy}^v be the set of all other pseudo-triangles of D_{xy} incident to v , every $T \in D_{xy}^v$ has a corner at v . Note that $gc(T) = v$ for every pseudo-triangle $T \in D_{xy}^v$ separated from T_v by s_{out} , because if one side of T incident to v is directed away from v , then both sides are directed that way. Similarly, $gc(T) \neq v$ for every pseudo-triangle $T \in D_{xy}^v$ separated from T_v by s_{in} . Now for every reflex vertex v align the π -guard at v with the incident side of P directed away from v . The resulting edge-aligned π -guard at v monitors at least as much area of T_v as before, and it monitors the full angular domain of every $T \in D_{xy}^v$ whose guarded corner




Fig. 5. Guard allocation according to the vertex assignment in Fig. 4.

is v (see Fig. 6). So after alignment, the resulting set of vertex π -guards still monitors every $T \in D_{xy}$.

Finally, we show that the total number of guards is $2n - k$. We use $2\ell_1 - 3$ guards on the boundary of the pseudo-triangle T_1 corresponding to the root of D_{xy} , where ℓ_1 is the number of vertices of T_1 . For all other $T_i \in D_{xy}$ with ℓ_i vertices, two guards out of the $2\ell_i - 3$ are already counted on the boundary of the ancestor of T_i . (This can be easily checked for all possible mutual positions of two adjacent pseudo-triangles: i.e., the common vertices are both corners, both reflex, or one is a corner and the other one is reflex.) Hence all three guard allocations use a total of $2 + \sum_{i=1}^p (2\ell_i - 5) = 2 - 5p + 2 \sum_{i=0}^p \ell_i$ guards where p denotes the number of pseudo-triangles, i.e., $p = k - 2$. Since $n - 2 = \sum_{i=1}^p (\ell_i - 2) = (\sum_{i=1}^p \ell_i) - 2p$, we obtain $2n - 2 - p = 2n - k$ guards in total. \square




Fig. 6. Guard allocation according to the vertex assignment in Figure 4 after alignment.

4. General π -Guards

Following the structure of the previous section we first show how to guard a single pseudo-triangle with $\lfloor n/2 \rfloor$ vertex π -guards and then generalize this approach to a special pseudo-triangulation of a simple polygon.

4.1. Guarding a Pseudo-Triangle

Lemma 4. *A pseudo-triangle with ℓ vertices can be monitored by $\lfloor \ell/2 \rfloor$ vertex π -guards such that there is at most one inward-facing π -guard, which lies on a prescribed chain, and all other π -guards are either located at a convex vertex or are outward-facing.*

Proof. We propose two possible guard allocations for a pseudo-triangle abc with a total of at most ℓ guards, such that all π -guards at reflex vertices are outward-facing, except for at most one guard on the prescribed chain \overline{bc} . We consider two cases:

Case 1: There is a vertex $v \in \overline{bc}$ visible from a . One guard allocation uses a guard at the corner v of the mountain abv and outward-facing guards on the subchain \overline{bv} , combined with one guard at a and outward-facing guards along \overline{ac} monitoring avc . The other allocation uses a guard at the corner v of the mountain avc and outward-facing guards on the subchain \overline{vc} , combined with one guard at a and outward-facing guards along \overline{ab} monitoring abv (see Fig. 7). The two allocations jointly use a total of ℓ guards.

Case 2: There is no vertex on \overline{bc} visible from a . Part of the relative interior of a side uv , $u, v \in \overline{bc} \cup \{b, c\}$, is visible from a . Denote by $g(u)$ (resp., $g(v)$) the first vertex on \overline{ba} (resp., \overline{ca}) visited by the geodesic path connecting u to a (resp., v to a). The ray $\overrightarrow{ug(u)}$ (resp., $\overrightarrow{vg(v)}$) hits the polygonal chain connecting a and c at point $h(u)$ (resp., $h(v)$) on the polygonal chain connecting a and b). The segment $uh(u)$ (resp., $vh(v)$) dissects abc into three mountain pseudo-triangles $T_u(a) = ag(u)h(u)$, $T_u(b) = bug(u)$, and $T_u(c) = ch(u)u$ containing the vertices a , b , and c , respectively. Note that $T_u(b)$ (resp., $T_v(c)$) can be empty if $b = u$ (resp., $c = v$). One guard allocation uses an inward-facing guard $(u, (uv)^+)$, outward-facing guards along \overline{uc} monitoring $T_u(c)$, and outward-facing




Fig. 7. The two guard allocations in Case 1.




Fig. 8. The two guard allocations in Case 2.

guards along \overline{ab} monitoring $T_u(a)$ and $T_u(b)$. The other guard allocation uses an inward-facing guard $(v, (uv)^+)$, outward-facing guards along \overline{vb} monitoring $T_v(b) = bvh(v)$, and outward-facing guards along \overline{ac} monitoring $T_v(a) = ah(v)g(v)$ and $T_v(c) = cg(v)v$ (see Fig. 8). The two allocations jointly use a total of $\ell - 1$ guards. \square

Note the differences between Cases 1 and 2. In Case 1 there is always a π -guard at corner a , and no guard at b nor at c ; one allocation uses at most $\lfloor \ell/2 \rfloor$ guards. In Case 2 there is never a guard at a , and there is a guard at b or at c only if $b = u$ or $c = v$; one guard allocation uses at most $\lfloor (\ell - 1)/2 \rfloor$ guards. Case 2 also contains the proof of Proposition 2 with $b = u$ and $c = v$.

4.2. Guarding an Art Gallery

To establish the upper bound in the general model, we rely on a special pseudo-triangulation of P , which corresponds to the so-called *pulling triangulation* of convex polygons. A *pulling pseudo-triangulation* of a simple polygon is obtained by choosing (i.e., *pulling*) a convex vertex x and connecting it to all other convex vertices of P via geodesic paths (see Fig. 9). The pulling pseudo-triangulation has the special property that every pseudo-triangle $T \in D_{xy}$, with the exception of the root, has at least one corner adjacent to its parent, i.e., at least one of the two vertices which T shares with its parent is a corner of T .

Theorem 5. $\lfloor n/2 \rfloor$ vertex π -guards can always monitor a simple polygon with n vertices in the general model.

Proof. The basic idea of the proof is to construct a partition \mathcal{S} of the vertices V_P of P and to place vertex π -guards at at most $\frac{1}{2}|\mathcal{S}|$ vertices for each $S \in \mathcal{S}$. It follows that the total number of vertex π -guards will be at most $\lfloor n/2 \rfloor$.

In order to define the partition \mathcal{S} , consider an ordered pulling pseudo-triangulation D_{xy} . We assign a subset $S_T \subset V_T$ to each pseudo-triangle $T \in D_{xy}$ such that $\mathcal{S} := \{S_T : T \in D_{xy}\}$ is a partition of V_P . Then we place vertex π -guards at some vertices $R_T \subset S_T$.




Fig. 9. A pulling pseudo-triangulation with respect to x .

The pseudo-triangles of D_{xy} are processed from bottom to top, that is, $T \in D_{xy}$ is processed when all its descendants have already been processed. Processing $T \in D_{xy}$ involves the following three tasks:

1. Defining $S_T \subset V_T$.
2. Specifying $R_T \subset S_T \subset V_T$ with $|R_T| \leq \frac{1}{2}|S_T|$.
3. Verifying that
 - (a) the vertex π -guards placed on the boundary of T (not only those in R_T) can be oriented such that they collectively monitor T ,
 - (b) for each vertex v , the orientation of a π -guard at v is consistent in every $T' \in D_{xy}$ adjacent to v .

The consistency of the orientation relies heavily on the pointedness of our pseudo-triangulation: a π -guard at a vertex v can monitor all convex angular domains adjacent to v and at the same time it can be an outward-facing π -guard with several possible orientations in a pseudo-triangle where v is a reflex vertex. This means that if we include at one step in the construction a vertex v in R_T where v is a corner in T , then we are free to choose the orientation of the π -guard at v whenever we are processing an ancestor T' of T in which v is a reflex vertex as long as the guard remains outward-facing in T' .

Now let $T = abc \in D_{xy}$ and suppose that all descendants $T' < T$ of T have already been processed. Let $K_T = \{v \in V_T : \nexists T' < T, v \in S_{T'}\}$ be the set of vertices of T which are not assigned to any descendant of T . Since D_{xy} is a pulling pseudo-triangulation, we can assume that one of the corners of T , say a , is adjacent to its parent (or $a = x$ if T is the root of D_{xy}). Furthermore, we assume that $s \in \overline{ab} \cup \{b\}$ is the other vertex that T shares with its parent. Let $\tilde{K}_T = K_T \setminus \{a, s\}$.

Since none of the elements of \tilde{K}_T is a vertex of any ancestor of T it follows that $\tilde{K}_T \subseteq S_T$ has to hold. Furthermore, S_T cannot contain vertices of T that have already been assigned to any of the previously processed pseudo-triangles, that is, $S_T \subseteq K_T$. This implies that the set $S_T \in \mathcal{S}$ of vertices assigned to T has to be one of \tilde{K}_T , $\tilde{K}_T \cup \{a\}$, $\tilde{K}_T \cup \{s\}$, and K_T .

In order to ensure that T is guarded by the guards placed on its boundary, we establish the following invariant:

Invariant 6. *For every $v \in V_T \setminus K_T$ (i.e., for every vertex $v \in V_T$ that has previously been assigned to one of its descendants), if $v \in R_{T'}$, $T' < T$ (i.e., there is guard placed at v), then the guard at v*

- *is outward-facing in T if v is a reflex vertex of T ;*
- *covers the angular domain of v in T if v is a corner of T .*

Invariant 6 clearly holds as long as no pseudo-triangle has been processed, i.e., as long as no guard has been placed. In order to maintain this invariant while processing a pseudo-triangle T it is sufficient to respect the following two rules: (i) if a or s is in S_T , then it is also in R_T , and (ii) if $s \in S_T$ (which implies $s \in R_T$) and s is reflex in T , then the guard at s is outward-facing.

The following lemma guarantees that we can process each pseudo-triangle $T \in D_{xy}$ while maintaining Invariant 6:

Lemma 7. *Assume we are given a pseudo-triangle $T = abc$ and a partition of its vertex set $V_T = K_T \cup^* L_T$. Denote by s the vertex next to a in $\overline{ab} \cup \{b\} \subset V_T$. Then T can be monitored by vertex π -guards at a set of vertices $U_T \subset V_T$ such that*

- *π -guards at reflex vertices of $L_T \cup \{s\}$ are outward-facing,*
- *there is at most one inward-facing guard and it is located on $\overline{bc} \cap K_T$,*
- *at most $\lfloor |K_T|/2 \rfloor$ guards are placed at vertices of K_T , in particular if $|K_T|/2$ guards are placed, then there is a guard placed at both a and s , and if $(|K_T| - 1)/2$ guards are placed, then there is a guard placed at least one of a and s .*

In order to simplify the proof of Theorem 5 we relegate the many but technically not difficult details of the proof of Lemma 7 to the next subsection.

Now we are ready to process the pseudo-triangle T . Recall that at this point all descendants of T have already been processed and Invariant 6 holds. We apply Lemma 7 with $L_T = V_T \setminus K_T$. This yields a guard allocation which places guards at the set of vertices $U_T \subset V_T$. Let $S_T = \tilde{K}_T \cup (\{a, s\} \cap U_T)$ and $R_T = K_T \cap U_T$. Lemma 7 and Invariant 6 imply that T is now indeed guarded by consistently oriented π -guards placed on its boundary. Furthermore, the orientations for the guards specified in Lemma 7 assure that Invariant 6 is maintained.

It remains to show that $|R_T| \leq \frac{1}{2}|S_T|$. According to Lemma 7 there are at most $\lfloor |K_T|/2 \rfloor$ guards placed at vertices of K_T , that is, $|R_T| = |U_T \cap K_T| \leq |K_T|/2$. If $|R_T| = |K_T|/2$, then $\{a, s\} \subset U_T$ and therefore $S_T = \tilde{K}_T \cup (\{a, s\} \cap U_T) = K_T$. If $|R_T| = (|K_T| - 1)/2$, then $a \in U_T$ or $s \in U_T$ and therefore $S_T = \tilde{K}_T \cup \{a\}$, $S_T = \tilde{K}_T \cup \{s\}$, or $S_T = \tilde{K}_T \cup \{a, s\}$, i.e., $|S_T| \geq |K_T| - 1$. Finally, if $|R_T| \leq (|K_T| - 2)/2$, then $|S_T| \geq |\tilde{K}_T| = |K_T| - 2$. \square

4.3. Proof of Lemma 7

Before we begin the actual proof of Lemma 7 we formulate an additional lemma which is a generalization of Proposition 2 as well, but slightly weaker than Lemma 7 itself. It will be used as a building block during the rest of this section.

Lemma 8. *Assume we are given a mountain $T = xyz$ with $\overline{yz} = \emptyset$ and a partition of its vertex set $V_T = K_T \cup^* L_T$.*

- (1) *If $\{x\} \cup \overline{xy} \cup \overline{xz} \subset L_T$, then T can be monitored by vertex π -guards at a subset of V_T such that*
 - π -guards at reflex vertices are outward-facing,
 - we place no guard at vertices of K_T .
- (2) *If $x \in L_T$ and at least one of $\overline{xy} \cap K_T$ and $\overline{xz} \cap K_T$ is non-empty, then T can be monitored by either one of two guard allocations $A_2(y) \subset V_T$ and $A_2(z) \subset V_T$ such that*
 - π -guards at reflex vertices are outward-facing,
 - $A_2(y)$ places a guard at y but no guard at z , while $A_2(z)$ places a guard at z but no guard at y ,
 - $A_2(y)$ and $A_2(z)$ jointly place $|K_T| - 1$ guards at vertices of K_T .
- (3) *If $x \in K_T$, then T can be monitored by either one of two guard allocations $A_3(y) \subset V_T$ and $A_3(z) \subset V_T$ such that*
 - π -guards at reflex vertices are outward-facing,
 - $A_3(y)$ places a guard at y but no guard at z , while $A_3(z)$ places a guard at z but no guard at y ,
 - $A_3(y)$ and $A_3(z)$ jointly place $|K_T| - 1$ guards at vertices of K_T .

Proof. (1) If $\overline{xy} \subset L_T$ and $\overline{xz} \subset L_T$ (i.e., K_T is one of $\{y\}, \{z\}$, and $\{y, z\}$), then one guard at x and outward-facing guards along the chains \overline{xy} and \overline{xz} monitor T using no guard at K_T .

(2) If $\overline{xy} \cap K_T$ or $\overline{xz} \cap K_T$ is non-empty, then let $y_0 = y, y_1, y_2, \dots, y_p$ and $z_0 = z, z_1, z_2, \dots, z_q$ be the vertices on the chains $\{y\} \cup \overline{yx}$ and on $\{z\} \cup \overline{zx}$. If $\overline{yx} \cap K_T \neq \emptyset$, then let y_i be the closest vertex to x in $\overline{yx} \cap K_T$ and, equivalently, if $\overline{zx} \cap K_T \neq \emptyset$, then let z_j be the closest vertex to x in $\overline{zx} \cap K_T$. If y_i is defined, then let e denote the intersection of the ray $\overrightarrow{y_{i-1}y_i}$ and the polygonal chain between z and x . Similarly, if z_j is defined let f denote the intersection of the ray $\overrightarrow{z_{j-1}z_j}$ and the polygonal chain between y and x . Note that at least one of e and f is defined.

We may suppose without loss of generality that f is defined. If e is defined as well, then we may also suppose that f is between y_i and x . The segment $z_j f$ dissects T into two pseudo-triangles xfz_j and fyz as depicted in Fig. 10. (In Fig. 10 and those that follow we depict the vertices of L_T as hollow circles and the vertices of K_T as filled circles. Similarly, guards placed at vertices of L_T are drawn with black lines while guards placed at vertices of K_T are drawn, as before, with grey lines.)

Both guard allocations $A_2(y)$ and $A_2(z)$ monitor the pseudo-triangle xfz_j by one π -guard at x and outward-facing π -guards along the two adjacent chains \overline{xf} and $\overline{xz_j}$, that is, using no guards at vertices of K_T . $A_2(y)$ monitors the pseudo-triangle fyz by one guard at




Fig. 10. Two possible partitions $K_T \cup^* L_T$ (left, right) and one possible guard allocation for the partition depicted on the left (middle).

corner y and outward-facing π -guards at the adjacent chain \overline{yf} . Analogously $A_2(z)$ monitors fyz by one guard at corner z and outward-facing π -guards at the adjacent chain $\overline{zz_j}$.

Note that $z_j \in K_T$ is not a vertex of fyz and so neither guard allocation places a guard at z_j . Since the two allocations use all other vertices of K_T exactly once, they jointly use $|K_T| - 1$ guards at vertices of K_T .

(3) Following the allocation scheme described in Proposition 2, $A_3(y)$ places a guard at y and outward-facing π -guards along the chain \overline{yx} . Analogously $A_3(z)$ places a guard at z and outward-facing π -guards along the chain \overline{zx} . \square

Proof of Lemma 7. We distinguish three cases. Note that in all of the cases we assume that if $a \in L_T$ or $s \in L_T$, then there is a guard placed at a or s , respectively.

Case A: T is a mountain (this includes the case of $b = s$). Since T is a mountain there is at least one empty chain. Let d denote the corner opposite such a chain. To guard T we invoke Lemma 8 with $x = d$.

First, let us note that all allocations of Lemma 8 use only outward-facing π -guards at reflex vertices.

Second, Lemma 7 clearly holds in the case of Lemma 8(1). Otherwise one of the allocations uses at most $\lfloor (|K_T| - 1)/2 \rfloor$ guards at vertices of K_T . If both allocations use exactly $(|K_T| - 1)/2$ guards, and both $a \in K_T$ and $s \in K_T$, then it remains to show that one allocation places a guard at a or s : If $d = a$, then $A_3(y)$ places a guard at s . If $d \neq a$ then $a = y$ or $a = z$ and one of the allocations from Lemma 8(2) or (3) places a guard at a .

Case B: T is not a mountain and there is a vertex $v \in K_T \cap \overline{bc}$ visible from a . Note that this case corresponds to Case 1 of the proof of Lemma 4.

The segment av dissects the pseudo-triangle $T = abc$ into two mountains abv and avc . We guard abv and avc by invoking Lemma 8 with $x = b$ and $x = c$, respectively.

If abv or avc are handled by Lemma 8(1), i.e., $\{b\} \cup \overline{bv} \cup \overline{ba} \subset L_T$ or $\{c\} \cup \overline{ca} \cup \overline{cv} \subset L_T$, then no guard from K_T is necessary to monitor abv or avc , respectively, and the guard allocation for $T \setminus abv$ or $T \setminus avc$ is therefore subsumed by Case A.

Otherwise we obtain the two guard allocations for T by combining the two pairs of guard allocations for abv and avc described by Lemma 8(2) or (3). Specifically,

we combine allocation $A_i(a)$ for abv and $A_j(v)$ for avc , $i, j \in \{2, 3\}$, and similarly allocation $A_i(v)$ for abv and $A_j(a)$ for avc , $i, j \in \{2, 3\}$.

Both guard allocations place a π -guard at a and at v . Lemma 8(2) and (3) place only outward-facing π -guards at reflex vertices. Therefore, only the guard at v is not outward-facing with respect to T but it can always be oriented to be inward-facing.

The two guard allocations jointly use at most $|K_T \cap V_{abv}| - 1$ guards at vertices of $K_T \cap V_{abv}$ and at most $|K_T \cap V_{avc}| - 1$ guards at vertices of $K_T \cap V_{avc}$. Therefore, they use at most $|K_T \cap V_T| = |K_T|$ guards at vertices of K_T . It follows that one of them necessarily uses at most $|K_T|/2$ guards.

If both guard allocations for T use exactly $|K_T|/2$ guards at vertices of K_T , then it remains to show that there always is a guard at a and s . Recall that both allocations place a guard at a . Since T is not a mountain, $s \neq b$ which implies that $s \in \overline{ab}$. If $b \in K_T$, then allocation $A_3(a)$ places a guard at s .

If, however, $b \in L_T$, then it is possible that s corresponds to y_i or z_j in the proof of Lemma 8(2) and so neither allocation places a guard at s . In this case $K_T \cap \overline{ab} = \{s\}$. We can therefore replace $A_2(v)$ in our combination by an allocation that guards abv by also using the guard at a placed by the allocation $A_j(a)$ for avc , $j \in \{2, 3\}$, and outward-facing π -guards at \overline{ab} . This allocation uses at most as many guards as $A_2(v)$, since it does not place a guard at $v \in K_T$ but places a guard at s instead.

Case C: T is not a mountain and there is no vertex in $K_T \cap \overline{bc}$ visible from a . Note that this case corresponds to Case 2 of the proof of Lemma 4.

Let \overline{uv} be the largest subchain of \overline{bc} such that a sees all vertices of \overline{uv} . Necessarily $\overline{uv} \subseteq L_T$. Following the notation of Lemma 4, denote by $g(u)$ (resp., $g(v)$) the first vertex on \overline{ba} (resp., \overline{ca}) visited by the geodesic path connecting u to a (resp., v to a). Note that $s \in \overline{g(u)a} \cup \{g(u)\}$. The segments $ug(u)$ and $vg(v)$ dissect abc into two mountains $bug(u)$ and $cg(v)v$, and a pseudo-triangle auv (see Fig. 11). Note that $bug(u)$ (resp., $cg(v)v$) can be empty if $b = u$ (resp., $c = v$).

If $bug(u)$ (resp., $cg(v)v$) is empty, then we can simply remove it from all further consideration. If they are non-empty, then we guard $bug(u)$ and $cg(v)v$ by invoking Lemma 8 with $x = b$ and $x = c$, respectively.




Fig. 11. The partition of T .

If $\text{bug}(u)$ or $\text{cg}(v)v$ are handled by Lemma 8(1), i.e., $\{b\} \cup \overline{bu} \cup \overline{bg(u)} \subset L_T$ or $\{c\} \cup \overline{cg(v)} \cup \overline{cv} \subset L_T$, then no guard from K_T is necessary to monitor $\text{bug}(u)$ or $\text{cg}(v)v$, respectively. In this case we can treat $T \setminus \text{bug}(u)$ (resp. $T \setminus \text{vcg}(v)$) as if $\text{bug}(u)$ (resp. $\text{cg}(v)v$) were empty.

Otherwise we define two guard allocations for T that are a combination of the guard allocations for its parts. The first allocation guards $\text{bug}(u)$ by $A_i(u)$ and $\text{vcg}(v)$ by $A_j(g(v))$, $i, j \in \{2, 3\}$, and it guards auv by one guard at corner v and outward facing π -guards at $\overline{vu} \subset L_T$ and \overline{va} . The second allocation guards $\text{bug}(u)$ by $A_i(g(u))$ and $\text{vcg}(v)$ by $A_j(v)$, $i, j \in \{2, 3\}$, and it guards auv by one guard at corner u and outward facing π -guards at $\overline{uv} \subset L_T$ and \overline{ua} .

Either allocation places guards at every vertex of $\overline{uv} \cup \{u, v\}$ and neither places a guard at a . Lemma 8(2) and (3) place only outward-facing π -guards at reflex vertices. Observe that only the guard at either v or u (at one corner of auv) is not outward-facing with respect to T , but it can always be oriented to be inward-facing. The guards at $g(u)$ and $g(v)$ are always outward-facing with respect to T .

The two guard allocations jointly use at most $|K_T \cap V_{\text{bug}(u)}| - 1$ guards at vertices of $K_T \cap V_{\text{bug}(u)}$ if $\text{bug}(u)$ is non-empty, and at most $|K_T \cap V_{\text{cg}(v)v}| - 1$ guards at vertices of $K_T \cap V_{\text{cg}(v)v}$ if $\text{cg}(v)v$ is non-empty. They jointly use $|(K_T \setminus \{a\}) \cap V_{\text{auv}}|$ guards at vertices of $(K_T \setminus \{a\}) \cap V_{\text{auv}}$. Therefore, they use at most $|(K_T \setminus \{a\}) \cap V_T| = |K_T \setminus \{a\}|$ guards at vertices of $K_T \setminus \{a\}$. It follows that one of them necessarily uses at most $|K_T|/2$ guards at vertices of K_T .

If both guard allocations for T use exactly $|K_T \setminus \{a\}|/2$ guards at vertices of $K_T \setminus \{a\}$, then one of them places a π -guard at $s \in \overline{ag(u)} \cup \{g(u)\}$. So if both allocations use exactly $|K_T|/2$ guards at vertices of K_T , then necessarily $a \in L_T$, and we may suppose that there is a guard at a . \square

5. Lower Bound Construction

We describe a family of simple polygons where no member P (with n vertices, k of which are convex) can be monitored by less than $(2n - k)/3$ inward-facing vertex π -guards, attaining the bound of Theorem 3. The parameters n and k are not independent in our construction: for any $\alpha, \beta \in \mathbb{Z}$, $0 \leq \alpha \leq \beta$, and $k \in \mathbb{N}$, $k \geq 3$, we build polygons with $n = (2k - 3) + 3(k - 2)(\alpha + \beta)$ vertices.

Our construction is based on a rotational symmetric pseudo-triangle with $3(1 + \alpha + \beta)$ vertices, which contains exactly $\alpha + \beta$ reflex vertices on every reflex chain. We then glue $k - 2$ copies of this pseudo-triangle into one simple polygon with $n = (k - 2)(5 + 3\alpha + 3\beta) + 1$ vertices, k of which are convex. This generalizes the families of polygons of both Santos and Urrutia. Optimizing the parameters α and β , we obtain simple polygons which require more vertex π -guards in terms of n than any previously known examples. Note that the upper bound of Theorem 3, in terms of n and k , is tight for every choice of α and β , $0 \leq \alpha \leq \beta$.

An (α, β) -mountain for $\alpha, \beta \in \mathbb{Z}$, $0 \leq \alpha \leq \beta$, is a mountain with 0 , α , and β reflex vertices on its three chains, respectively. Denote the vertices on the two non-empty chains of an (α, β) -mountain axb (including the endpoints) by $(x = a_0, a_1, a_2, \dots, a_\alpha, a_{\alpha+1} = a)$ and $(x = b_0, b_1, b_2, \dots, b_\beta, b_{\beta+1} = b)$. The mountain is constructed such that the




Fig. 12. A (2, 4)-mountain.

ray $\overrightarrow{a_{i+1}a_i}$ hits the relative interior of the side $b_{i+1}b_i$ and the ray $\overrightarrow{b_{i+1}b_i}$ hits $\text{relint}(a_{i+1}a_i)$ for every $i = 1, 2, \dots, \alpha$; then the ray $\overrightarrow{b_{i+1}b_i}$ hits $\text{relint}(a_\alpha a)$ for every $i = \alpha + 1, \alpha + 2, \dots, \beta$. Thus axb is constructed recursively by locating the vertices $b_1, a_1, b_2, \dots, a_\alpha, b_\alpha, a_{\alpha+1} = a$ and then the vertices $b_{\alpha+1}, b_{\alpha+2}, \dots, b_\beta, b_{\beta+1} = b$. See Fig. 12 for an example with $\alpha = 2$ and $\beta = 4$.

Proposition 9. *The minimum number of inward-facing vertex π -guards required to monitor an (α, β) -mountain axb is at least*

- $\alpha + \beta + 1$ if we cannot use guards at a or b ,
- $\beta + 1$ if we cannot use a guard at a ,
- $\alpha + 1$ otherwise.

The proof of this proposition extends an argument by O'Rourke [16].

Proof. Every (α, β) -mountain axb has a unique triangulation consisting of $\alpha + \beta + 1$ triangles (see Fig. 12). If a triangle is jointly monitored by several inward-facing vertex π -guards of the mountain, then one of them already monitors it fully. Observe that every inward-facing π -guard can fully monitor at most two triangles, with the exception of a guard placed at a which can fully monitor $2 + \beta - \alpha$ triangles. Therefore axb cannot be monitored by less than $\alpha + 1$ inward-facing π -guards.

If we cannot use a guard at a , then each of the $\beta - \alpha + 1$ triangles $\Delta ab_\beta b_{\beta+1}, \dots, \Delta ab_\alpha b_{\alpha+1}$ requires one guard. The remaining 2α triangles require at least α more guards.

If we cannot use guards at a or b , then every triangle requires a separate guard. This is the case, since the triangle $\Delta ab_\beta b$ can only be monitored by the guard $(b_\beta, (b_\beta b)^+)$. Once the guard at b_β is placed the triangle $\Delta ab_{\beta-1} b_\beta$ can only be monitored by the guard $(b_{\beta-1}, (b_{\beta-1} b_\beta)^+)$. This argument can be iterated for every triangle in the triangulation. \square

Lemma 10. *For any $\alpha, \beta \in \mathbb{Z}$, $0 \leq \alpha \leq \beta$, there is a (rotational symmetric) pseudo-triangle with $6 + 3\alpha + 3\beta$ vertices which cannot be monitored by less than $3 + \beta + \min\{1 + 3\alpha, 2\beta\}$ inward-facing vertex π -guards.*

Proof. We construct the pseudo-triangle $P = xyz$ as the union of an equilateral triangle abc and three congruent (α, β) -mountains axb , byc , and cza . Let o denote the center of abc . The three mountains are rotational symmetric with respect to o . By symmetry, it suffices to describe one (α, β) -mountain, say, axb . Neither x nor any reflex vertex of axb can see the center o or any points in the relative interior of bc . See Fig. 13 for an example with $\alpha = 1$ and $\beta = 2$.

The pseudo-triangle xyz has a unique triangulation that consists of the triangle Δabc and the triangulations of the mountains (see Fig. 12). If a triangle in the triangulation of an (α, β) -mountain is jointly monitored by several inward-facing vertex π -guards, then




Fig. 13. Schematic picture of the pseudo-triangle in Lemma 10.

one of them at a vertex of the mountain already monitors it fully. Furthermore, only a vertex π -guard at a , b , or c can monitor the center o .

If a vertex π -guard at a (resp., b or c) monitors o , then the mountain axb (resp., byc , or cza) cannot use this vertex guard. Suppose that two (α, β) -mountains, axb and byc , use $(\alpha + 1)$ guards each and therefore use a guard at a and b , respectively. The guards at a and b do not monitor o , hence an inward-facing π -guard at c must monitor o . Consequently, cza cannot use guards at c nor at a . So we need at least $2(\alpha + 1) + 1 + (\alpha + \beta + 1) = 4 + 3\alpha + \beta$ guards. Finally, if all three (α, β) -mountains use $(\beta + 1)$ guards, then we need $3 + 3\beta$ guards, three of which monitor o as well. \square

Now suppose that we are given two pseudo-triangles P_1 and P_2 of the type described in Lemma 10 with n_1 and n_2 vertices requiring at least g_1 and g_2 inward-facing vertex π -guards, respectively. We can glue the two pseudo-triangles into one polygon with $n_1 + n_2 - 1$ vertices requiring at least $g_1 + g_2$ inward-facing vertex π -guards: Let v_1 and v_2 be two corners of P_1 and P_2 . Identify v_1 and v_2 such that two of their incident sides in the polygons are collinear.

We then construct a simple polygon with k convex vertices by recursively gluing $k - 2$ copies of the pseudo-triangle described in Lemma 10. At each of the $k - 3$ gluing steps we lose one vertex while the number of guards required remains the sum of the guards of all pseudo-triangles. This leads to the following lemma:

Lemma 11. *There exists a family of polygons with $n = (k - 2)(5 + 3\alpha + 3\beta) + 1$ vertices, k of which are convex, such that it cannot be monitored by less than $(k - 2)(3 + \beta + \min\{1 + 3\alpha, 2\beta\})$ inward-facing vertex π -guards.*

In our notation, Santos's polygons have parameters $\alpha = 0$ and $\beta = 0$, and Urrutia's polygons have $\alpha = 0$ and $\beta = 1$. The ratio of necessary inward-facing vertex π -guards and the total number of vertices is maximized by the choice $\alpha = 1$ and $\beta = 2$, which proves the following theorem.

Theorem 12. *For every $n \geq 3$, there is a simple polygon with n vertices that cannot be monitored by less than $9n/14 - O(1)$ inward-facing vertex π -guards, where $9/14 \approx 0.643$.*

References

1. P. K. Agarwal, J. Basch, L. J. Guibas, J. Hershberger, and L. Zhang, Deformable free space tilings for kinetic collision detection, in: *Algorithmic and Computational Robotics: New Directions, Proc. 5th Workshop Algorithmic Found. Robotics (Hanover, NH, 2001)*, pp. 83–96.
2. V. Brumberg, S. Ramaswami, and D. Souvaine, Experimental results on upper bounds for vertex π -lights, in: *Abstr. 11th Ann. Fall Workshop Comput. Geom. (Brooklyn, NY, 2001)*.
3. B. Chazelle, Triangulating a simple polygon in linear time, *Discrete Comput. Geom.* **6**(5) (1991), 485–524.
4. B. Chazelle, H. Edelsbrunner, M. Gringó, L. J. Guibas, J. Hershberger, M. Sharir, and J. Snoeyink, Ray shooting in polygons using geodesic triangulations, *Algorithmica* **12** (1994), 54–68.
5. V. Chvátal, A combinatorial theorem in plane geometry, *J. Combin. Theory Ser. B* **18** (1975), 39–41.

6. V. Estivill-Castro, J. O'Rourke, J. Urrutia, and D. Xu, Illumination of polygons with vertex guards, *Inform. Process. Lett.* **56** (1995), 9–13.
7. S. Fisk, A short proof of Chvátal's watchman theorem, *J. Combin. Theory Ser. B* **24** (1978), 374.
8. M. Goodrich and R. Tamassia, Dynamic ray shooting and shortest paths in planar subdivision via balanced geodesic triangulations, *J. Algorithms* **23** (1997), 51–73.
9. L. J. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. E. Tarjan, Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons, *Algorithmica* **2** (1987), 209–233.
10. J. Hershberger and J. Snoeyink, Computing minimum length paths of a given homotopy class, *Comput. Geom. Theory Appl.* **4** (1994), 63–97.
11. D. Kirkpatrick, J. Snoeyink, and B. Speckmann, Kinetic collision detection for simple polygons, *Internat. J. Comput. Geom. Appl.* **12** (2002), 3–27.
12. J. S. B. Mitchell and J. O'Rourke, Computational geometry column 42, *SIGACT News* **32**(3) (2001), 63–72.
13. M. Pocchiola and G. Vegter, Pseudo-triangulations: theory and applications, in: *Proc. 12th ACM Symp. Comput. Geom. (Philadelphia, PA, 1996)*, ACM Press, New York 1996, pp. 291–300.
14. M. Pocchiola and G. Vegter, Topologically sweeping visibility complexes via pseudo-triangulations, *Discrete Comput. Geom.* **16** (1996), 419–453.
15. J. O'Rourke, *Art Gallery Theorems and Algorithms*, The International Series of Monographs on Computer Science, Oxford University Press, New York, 1987.
16. J. O'Rourke, Vertex π -lights for monotone mountains, in: *Proc. 9th Canadian Conf. Comput. Geom. (Kingston, ON, 1997)*, pp. 1–5.
17. J. O'Rourke, Open problems in the combinatorics of visibility and illumination, in: *Advances in Discrete and Computational Geometry* (B. Chazelle, J. E. Goodman, and R. Pollack, eds.), AMS, Providence, RI, 1998, 237–243.
18. J. O'Rourke, Computational geometry column 43, *SIGACT News* **33**(1) (2002), 58–60.
19. I. Streinu, A combinatorial approach to planar non-colliding robot arm motion planning, in: *Proc. 41st Symp. Found. Comput. Sci. (Redondo Beach, CA, 2000)*, IEEE Computer Society Press, Los Alamitos, CA 2000, pp. 443–453.
20. L. Szabó, Recent results on illumination problems, in: *Intuitive Geometry*, Bolyai Society Mathematical Studies 6, Bolyai Mathematical Society, Budapest, 1997, pp. 207–221.
21. Cs. D. Tóth, Art gallery problem with guards whose range of vision is 180° , *Comput. Geom. Theory Appl.* **17** (2000), 121–134.
22. Cs. D. Tóth, Illuminating polygons with π -floodlights, in: *Proc. Internat. Conf. Comput. Sci. (San Francisco, CA, 2001), Part I*, vol. 2073 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, 2001, pp. 772–781.
23. G. T. Toussaint, Computing geodesic properties inside a simple polygon, *Rev. Intell. Artific.* **3** (1989), 9–42.
24. J. Urrutia, Art gallery and illumination problems, in: *Handbook on Computational Geometry* (J. R. Sack and J. Urrutia, eds.), Elsevier, Amsterdam, 2000, pp. 973–1027.
25. J. Urrutia, Personal communication, 2001.

Received March 31, 2003, and in revised form November 7, 2003. Online publication June 4, 2004.