Improved Bounds for Wireless Localization

Tobias Christ[1], Michael Hoffmann[1], Yoshio Okamoto[2], and Takeaki Uno[3]

[1] Institute for Theoretical Computer Science, ETH Zürich, Switzerland
[2] Tokyo Institute of Technology, Japan
[3] National Institute of Informatics, Tokyo, Japan

Abstract. We consider a novel class of art gallery problems inspired by wireless localization. Given a simple polygon P, place and orient guards each of which broadcasts a unique key within a fixed angular range. Broadcasts are not blocked by the edges of P. The interior of the polygon must be described by a monotone Boolean formula composed from the keys. We improve both upper and lower bounds for the general setting by showing that the maximum number of guards to describe any simple polygon on n vertices is between roughly $\frac{3}{5}n$ and $\frac{4}{5}n$. For the natural setting where guards may be placed aligned to one edge or two consecutive edges of P only, we prove that $n-2$ guards are always sufficient and sometimes necessary.

1 Introduction

Art gallery problems are a classic topic in discrete and computational geometry, dating back to the question posed by Victor Klee in 1973: "How many guards are necessary, and how many are sufficient to patrol the paintings and works of art in an art gallery with n walls?" Chvátal [2] was the first to show that $\lfloor n/3 \rfloor$ guards are always sufficient and sometimes necessary, while the beautiful proof of Fisk [6] made it into "the book" [1]. Nowadays there is a vast literature [12, 14, 16] about variations of this problem, ranging from optimization questions (minimizing the number of guards [10] or maximizing the guarded boundary [7]) over special types of guards (mobile guards [11] or vertex pi-guards [15]) to special types of galleries (orthogonal polygons [8] or curvilinear polygons [9]).

A completely different direction has recently been introduced by Eppstein, Goodrich, and Sitchinava [5]. They propose to modify the concept of visibility by not considering the edges of the polygon/gallery as blocking. The motivation for this model stems from communication in wireless networks where the signals are not blocked by walls, either. For illustration, suppose you run a café (modeled, say, as a simple polygon P) and you want to provide wireless Internet access to your customers. But you do not want the whole neighborhood to use your infrastructure. Instead, Internet access should be limited to those people who are located within the café. To achieve this, you can install a certain number of devices, let us call them guards, each of which broadcasts a unique (secret) key in an arbitrary but fixed angular range. The goal is to place guards and adjust their angles in such a way that everybody who is inside the café can prove this fact just by naming the keys received and nobody who is outside the café can provide such a proof. Formally this means that P can be described by a monotone Boolean formula over the keys, that is, a formula using the operators And and Or only, negation is not allowed. It is convenient to model a guard as a subset of the plane, namely the area where the broadcast from this guard can be received. This area can be described as an intersection or union of at most two halfplanes. Using this notation, the polygon P is to be described by a combination of the operations union and intersection over the guards. For example, the first polygon to the right can be described by $(a \cup b) \cap c \cap d$.

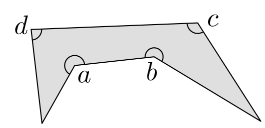

Natural guards. Natural locations for guards are the vertices and edges of the polygon. A guard which is placed at a vertex of P is called a *vertex guard*. A vertex guard is *natural* if it covers exactly the interior angle of its vertex. But natural vertex guards alone do not always suffice [5], as the second polygon P shown to the right illustrates: No natural vertex guard can distinguish the point p inside P from the point q outside of P. A guard placed anywhere on the line given by an edge of P and broadcasting within an angle of π to the inner side of the edge is called a *natural edge guard*. Dobkin, Guibas, Hershberger, and Snoeyink [4] showed that n natural edge guards are sufficient for any simple polygon with n edges.

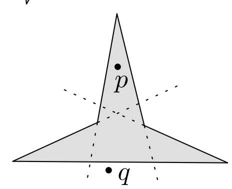

Vertex guards. Eppstein et al. [5] proved that any simple polygon with n edges can be guarded using at most $n-2$ (general, that is, not necessarily natural) vertex guards. More generally, they show that $n+2(h-1)$ vertex guards are sufficient for any simple polygon with n edges and h holes. This bound is not known to be tight. Damian, Flatland, O'Rourke, and Ramaswami [3] describe simple polygons with n edges which require at least $\lfloor 2n/3 \rfloor - 1$ vertex guards.

General guards. In the most general setting, we do not have any restriction on the placement and the angles of guards. So far the best upper bound known has been the same as for vertex guards, that is, $n-2$. On the other hand, if the polygon does not have collinear edges then at least $\lceil n/2 \rceil$ guards are always necessary [5]. The lower bound construction of Damian et al. [3] for vertex guards does not provide an improvement in the general case, where these polygons can be guarded using at most $\lceil n/2 \rceil + 1$ guards. As O'Rourke wrote [13]: "The considerable gap between the $\lceil n/2 \rceil$ and $n-2$ bounds remains to be closed."

Results. We provide a significant step in bringing the two bounds for general guards closer together by improving both on the upper and on the lower side. On one hand we show that for any simple polygon with n edges $\lfloor (4n-2)/5 \rfloor$ guards are sufficient. The result generalizes to polygons combined in some way by the operations intersection and/or union. Any simple polygon with h holes can be guarded using at most $\lfloor (4n-2h-2)/5 \rfloor$ guards. On the other hand we describe a family of polygons which require at least $\lceil (3n-4)/5 \rceil$ guards. Furthermore, we extend the result of Dobkin et al. [4] to show that $n-2$ natural (vertex or edge) guards are always sufficient. It turns out that this bound is tight.

Table 1. Number of guards needed for a simple polygon on n vertices.
The mark $*$ indicates the results of this paper.

	natural		general	
	vertex guards	guards	vertex guards	guards
upper bound	does not exist [5]	$n-2$ [$*$]	$n-2$ [5]	$\lfloor (4n-2)/5 \rfloor$ [$*$]
lower bound	does not exist [5]	$n-2$ [$*$]	$\lfloor 2n/3 \rfloor - 1$ [3]	$\lceil (3n-4)/5 \rceil$ [$*$]

2 Notation and Basic Properties

We are given a simple polygon $P \subset \mathbb{R}^2$. A *guard* g is a closed subset of the plane, whose boundary ∂g is described by a vertex v and two rays emanating from v. The ray that has the interior of the guard to its right is called the *left ray*, the other one is called the *right ray*. The *angle* of a guard is the interior angle formed by its rays. For a guard with angle π, the vertex is not unique.

A guard g *covers* an edge e of P *completely* if $e \subseteq \partial g$ and their orientations match, that is, the inner side of e is on the inner side of g. We say e is covered *partly* by g if their orientations match and $e \cap \partial g$ is a proper sub-segment of e that is not just a single point. We call a guard a k-guard if it covers exactly k edges completely. As P is simple, a guard can cover at most one edge partly. If a guard covers an edge partly and k edges completely, we call it a k'-guard. Assuming there are no collinear edges, a guard can cover at most two edges; then a natural vertex guard is a 2-guard and a natural edge guard is a 1-guard. A *guarding* $\mathcal{G}(P)$ for P is a formula composed of a set of guards and the operators union and intersection that defines P. The *wireless localization problem* is to find a guarding with as few guards as possible. The same problem is sometimes referred to as *guard placement for point-in-polygon proofs* or the *sculpture garden problem* [5]. The following basic properties are restated without proof.

Observation 1. *For any guarding $\mathcal{G}(P)$ and for any two points $p \in P$ and $q \notin P$ there is a $g \in \mathcal{G}(P)$ which* distinguishes *p and q, that is, $p \in g$ and $q \notin g$.*

Lemma 1. [4] *Every edge of P must be covered by at least one guard or it must be covered partly by at least two guards.*

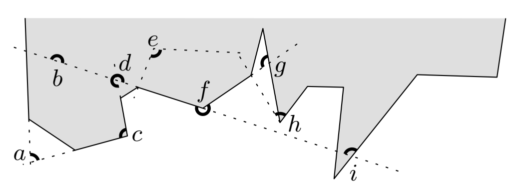

Fig. 1. (a) a 2-guard, (b) a 1-guard (and a natural edge guard), (c) a 2-guard (and a natural vertex guard), (d) a 2-guard, (e) a 0-guard, (f) a 0-guard, (g) a 1-guard (not a $1'$-guard), (h) a 1-guard (a non-natural vertex guard), (i) a $1'$-guard.

3 Upper Bounds

Following Dobkin et al. [4] we use the notion of a *polygonal halfplane* which is a topological halfplane bounded by a *simple bi-infinite polygonal chain* with edges $(e_1, \ldots, e_n)$, for $n \in \mathbb{N}$. For $n = 1$, the only edge e_1 is a line and the polygonal halfplane is a halfplane. For $n = 2$, e_1 and e_2 are rays which share a common source but are not collinear. For $n \geq 3$, e_1 and e_n are rays, e_i is a line segment, for $1 < i < n$, and e_i and e_j, for $1 \leq i < j \leq n$, do not intersect unless $j = i + 1$ in which case they share an endpoint. For brevity we use the term *chain* in place of simple bi-infinite polygonal chain in the following. For a polygonal halfplane H define $\gamma(H)$ to be the minimum integer k such that there exists a guarding $\mathcal{G}(H)$ for H using k guards. Similarly, for a natural number n, denote by $\gamma(n)$ the maximum number $\gamma(H)$ over all polygonal halfplanes H that are bounded by a chain with n edges. Obviously $\gamma(1) = \gamma(2) = 1$. Dobkin et al. [4] show that $\gamma(n) \leq n$.

Lemma 2. *Any simple polygon P on $n \geq 4$ vertices is an intersection of two polygonal halfplanes each of which consists of at least two edges.*

Proof. Let p_- and p_+ be the vertices of P with minimal and maximal x-coordinate, respectively. If they are not adjacent along P, split the circular sequence of edges of P at both p_- and p_+ to obtain two sequences of at least two segments each. Transform each sequence into a chain by linearly extending the first and the last segment beyond p_- or p_+ to obtain a ray. As p_- and p_+ are opposite extremal vertices of P, the two chains intersect exactly at these two points. Thus, the polygon P can be expressed as an intersection of two polygonal halfplanes bounded by these chains. Now consider the case that p_- and p_+ are adjacent along P. Without loss of generality assume that P lies above the edge from p_- to p_+. Rotate clockwise until another point q has x-coordinate larger than p_+. If q and p_- are not adjacent along P, then split P at these points. Otherwise the convex hull of P is the triangle qp_-p_+. In particular, q and p_+ are opposite non-adjacent extremal vertices and we can split as described above. □

Theorem 3. *Any simple polygon P with $n \geq 4$ edges can be guarded using at most $n - 2$ natural (vertex or edge) guards.*

Proof. Dobkin et al. [4] showed that for any chain there is a Peterson-style formula, that is, a guarding using natural edge guards only in which each guard appears exactly once and guards appear in the same order as the corresponding edges appear along the chain. Looking at the expression tree of this formula there is at least one vertex both of whose children are leaves. In other words, there is an operation (either union or intersection) that involves only two guards. As these two guards belong to two consecutive edges of P, we can replace this operation in the formula by the natural vertex guard of the common vertex, thereby saving one guard. Doing this for both chains as provided by Lemma 2 yields a guarding for P using $n - 4$ natural edge guards and two natural vertex guards. □

The (closure of) the complement of a polygonal halfplane H, call it $\overline{H}$, is a polygonal halfplane as well.

Observation 2. *Any guarding for H can be transformed into a guarding for $\overline{H}$ using the same number of guards.*

Proof. Use de Morgan's rules and invert all guards (keep their location but flip the angle to the complement with respect to 2π). Note that the resulting formula is monotone. Only guards complementary to the original ones appear (in SAT terminology: only negated literals); a formula is not monotone only if both a guard **and** its complementary guard appear in it.

Corollary 4. *Let $P_1, \ldots, P_m$ be a collection of $m \geq 1$ simple polygons t of which are triangles, for $0 \leq t \leq m$. Let R be a region that can be described as a formula composed of the operations intersection, union, and complement over the variables $\{P_1, \ldots, P_m\}$ in which each P_i appears exactly once. Then R can be guarded using at most $n - 2m + t$ natural (vertex or edge) guards, where n is the total number of edges of the polygons P_i, for $1 \leq i \leq m$.* □

Corollary 5. *Any simple polygon with $n \geq 4$ edges and h non-triangular holes can be guarded using at most $n - 2(h + 1)$ natural (vertex or edge) guards.* □

Our guarding scheme for chains is based on a recursive decomposition in which at each step the current chain is split into two or more subchains. At each split some segments are extended to rays and we have to carefully control the way these rays interact with the remaining chain(s). This is particularly easy if the split vertex lies on the convex hull because then the ray resulting from the segment extension cannot intersect the remainder of the chain at all. However, we have to be careful what we mean by convex hull. Instead of looking at the convex hull of a polygonal halfplane H we work with the convex hull of its bounding chain C. The convex hull $\mathrm{h}(C)$ of a chain $C = (e_1, \ldots, e_n)$, for $n \geq 2$, is either the convex hull of H or the convex hull of $\overline{H}$, whichever of these two is not the whole plane which solely depends on the direction of the two rays of C. The boundary of $\mathrm{h}(C)$ is denoted by $\partial\mathrm{h}(C)$. There is one degenerate case, when the two rays defining C are parallel and all vertices are contained in the strip between them; in this case, $\mathrm{h}(C)$ is a strip bounded by the two parallel lines through the rays and thus $\partial\mathrm{h}(C)$ is disconnected.

Theorem 6. *Any polygonal halfplane bounded by a simple bi-infinite polygonal chain with $n \geq 2$ edges can be guarded using at most $\lfloor (4n - 1)/5 \rfloor$ guards.*

Proof. The statement is easily checked for $2 \leq n \leq 3$. We proceed by induction on n. Let C be any chain with $n \geq 4$ edges. Denote the sequence of edges along C by $(e_1, \ldots, e_n)$ and let v_i, for $1 \leq i < n$, denote the vertex of C incident to e_i and e_{i+1}. The underlying (oriented) line of e_i, for $1 \leq i \leq n$, is denoted by ℓ_i. For $2 \leq i \leq n - 1$, let e_i^+ be the ray obtained from e_i by extending the segment linearly beyond v_i. Similarly e_i^- refers to the ray obtained from e_i by extending the segment linearly beyond v_{i-1}. For convenience, let $e_1^+ = \ell_1$ and $e_n^- = \ell_n$.

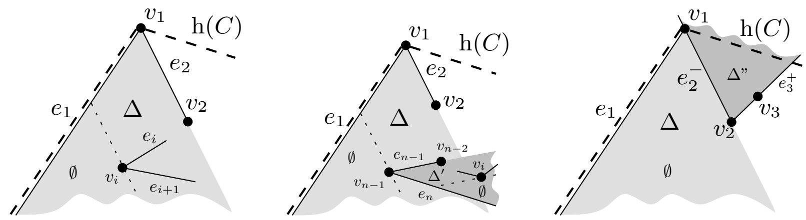

(a) Case 1. (b) Case 2. (c) Case 3.

Fig. 2. The chain C can interact with the shaded region Δ in three possible ways. The label $\emptyset$ marks an area which does not contain any vertex from C.

Without loss of generality (cf. Observation 2) suppose that v_1 is reflex, that is, the interior of the region bounded by C lies in the angle of C incident to v_1 which is larger than π. If there is any vertex v_i on $\partial\mathrm{h}(C)$, for some $1 < i < n-1$, then split C into two chains $C_1 = (e_1, \ldots, e_i^+)$ and $C_2 = (e_{i+1}^-, \ldots, e_n)$. We obtain a guarding for C as $\mathcal{G}(C_1) \cup \mathcal{G}(C_2)$ and thus $\gamma(C) \le \gamma(i) + \gamma(n-i)$, for some $2 \le i \le n-2$. As both $i \ge 2$ and $n - i \ge 2$, we can bound by the inductive hypothesis $\gamma(C) \le \lfloor (4i-1)/5 \rfloor + \lfloor (4n-4i-1)/5 \rfloor \le \lfloor (4i-1)/5 + (4n-4i-1)/5 \rfloor \le \lfloor (4n-1)/5 \rfloor$. Else, if both e_1 and e_n are part of $\partial\mathrm{h}(C)$ and ℓ_1 intersects ℓ_n then we place a guard g that covers both rays at the intersection of ℓ_1 and ℓ_n to obtain a guarding $g \cup \mathcal{G}(e_2^-, \ldots, e_{n-1}^+)$ for C. Therefore, in this case $\gamma(C) \le 1 + \gamma(n-2)$. Observe that this is subsumed by the inequality from the first case with $i = 2$. Otherwise, either ℓ_1 does not intersect ℓ_n and thus v_1 and v_{n-1} are the only vertices of $\partial\mathrm{h}(C)$ (the degenerate case where $\partial\mathrm{h}(C)$ is disconnected) or without loss of generality (reflect C if necessary) v_1 is the only vertex of $\partial\mathrm{h}(C)$. Let Δ denote the open wedge bounded by e_1 and e_2^+. We distinguish three cases.

Case 1. There is a vertex of C in Δ and among these, a vertex furthest from ℓ_2 is v_i, for some $3 \le i \le n-2$ (Fig. 2(a)). Split C into three chains, $C_1 = (\ell_1)$, $C_2 = (e_2^-, \ldots, e_i^+)$, and $C_3 = (e_{i+1}^-, \ldots, e_n)$. By the choice of v_i there is no intersection between C_2 and C_3 other than at v_i. A guarding for C can be obtained as $\mathcal{G}(C_1) \cup (\mathcal{G}(C_2) \cap \mathcal{G}(C_3))$. In this case $\gamma(C) \le 1 + \gamma(j) + \gamma(n-j-1)$, for some $2 \le j \le n-3$. Since $j \ge 2$ and $n-j-1 \ge n-(n-3)-1 = 2$, we can apply the inductive hypothesis to bound $\gamma(C) \le 1 + \lfloor (4j-1)/5 \rfloor + \lfloor (4n-4j-5)/5 \rfloor \le \lfloor (4n-1)/5 \rfloor$.

Case 2. There is a vertex of C in Δ and among these, the unique one furthest from ℓ_2 is v_{n-1} (Fig. 2(b)). We may suppose that ℓ_1 intersects ℓ_n; otherwise (in the degenerate case where $\partial\mathrm{h}(C)$ is disconnected), exchange the roles of v_1 and v_{n-1}. We cannot end up in Case 2 both ways. Let Δ' denote the open (convex) wedge bounded by e_n and e_{n-1}^-. If there is any vertex of C in Δ', let v_i be such a vertex which is furthest from ℓ_{n-1}. Let $C_1 = (e_1, \ldots, e_i^+)$ and $C_2 = (e_{i+1}^-, \ldots, e_{n-1}^+)$.

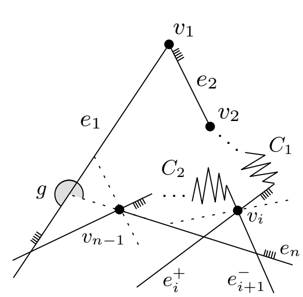

Both C_1 and C_2 are simple, except that their first and their last ray may intersect (in that case split the resulting polygon into two chains). Put a guard g at the intersection of ℓ_n with e_1 such that g covers e_n completely and e_1 partially (see figure, the small stripes indicate the side to be guarded). A guarding for C can be obtained as $g \cap (\mathcal{G}(C_1) \cup \mathcal{G}(C_2))$. Again this yields $\gamma(C) \leq 1+\gamma(i)+\gamma(n-i-1)$, for some $2 \leq i \leq n-3$, and thus $\gamma(C) \leq \lfloor (4n-1)/5 \rfloor$ as above in Case 1.

Otherwise there is no vertex of C in Δ'. We distinguish two sub-cases. If e^+_{n-1} intersects e_1 then put two guards (see figure): a first guard g_1 at the intersection of ℓ_n with e_1 such that g_1 covers e_n completely and e_1 partially, and a second guard g_2 at the intersection of ℓ_{n-1} with e_1 such that g_2 covers e_{n-1} completely and e_1 partially. Together g_1 and g_2 cover e_1 and $g_1 \cap (g_2 \cup \mathcal{G}(C'))$ provides a guarding for C, with $C' = (e^-_2, \ldots, e^+_{n-2})$. In this case we obtain $\gamma(C) \leq 2+\gamma(n-3)$ and thus by the inductive hypothesis $\gamma(C) \leq 2 + \lfloor (4n-13)/5 \rfloor \leq \lfloor (4n-1)/5 \rfloor$.

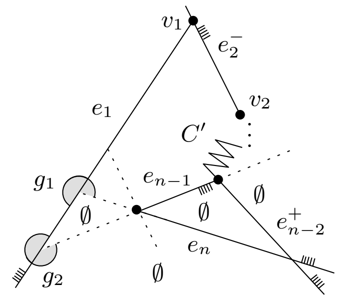

Finally, suppose that e^+_{n-1} does not intersect e_1. Then for the chain $C' = (e_1, \ldots, e^+_{n-1})$ there is some vertex other than v_1 on the convex hull boundary $\mathrm{h}(C')$. Thus we can obtain a guarding for C' as described above for the case that there is more than one vertex on the convex hull. Put a guard g at the intersection of ℓ_n with e_1 such that g covers e_n completely and e_1 partially (see figure). This yields a guarding $g \cap \mathcal{G}(C')$ for C with $\gamma(C) \leq 1 + \gamma(C') \leq 1 + \gamma(i) + \gamma(n-i-1)$, for some $2 \leq i \leq n-3$. As in Case 1 we conclude that $\gamma(C) \leq \lfloor (4n-1)/5 \rfloor$.

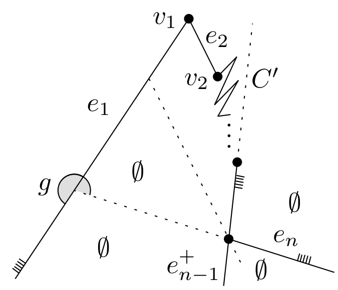

Case 3. There is no vertex of C in Δ (Fig. 2(c)). Let Δ'' denote the open (convex) wedge bounded by e^-_2 and e^+_3. If e^-_3 does not intersect e_1 then put a natural vertex guard g at v_1 to obtain a guarding $g \cap \mathcal{G}(C')$ for C, where $C' = (e^-_3, \ldots, e_n)$. This yields $\gamma(C) \leq 1 + \gamma(n-2)$ and thus by the inductive hypothesis $\gamma(C) \leq 1 + \lfloor (4n-9)/5 \rfloor \leq \lfloor (4n-1)/5 \rfloor$.

Now suppose that e^-_3 intersects e_1. We distinguish two sub-cases. If there is no vertex of C in Δ'', then place two guards: a natural vertex guard g_1 at v_1 and a guard g_2 at the intersection of e^-_3 with e_1 such that g_1 covers e_3 completely and e_1 partially. A guarding for C is provided by $g_1 \cap (g_2 \cup \mathcal{G}(C'))$, with $C' = (e^-_4, \ldots, e_n)$. In this case we obtain $\gamma(C) \leq 2 + \gamma(n-3)$ and thus in the same way as shown above $\gamma(C) \leq \lfloor (4n-1)/5 \rfloor$.

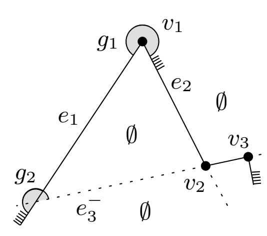

Otherwise there is a vertex of C in Δ''. Let v_i, for some $4 \leq i \leq n-1$, be a vertex of C in Δ'' which is furthest from ℓ_3. First suppose e^-_{i+1} does not intersect e_2. Then neither does e_i^+ and hence we can split at v_i in the same way as if v_i would be on $\partial\mathrm{h}(C)$. If $i = n-1$, e_n^- must intersect e_2 (otherwise, e_n would be on $\partial\mathrm{h}(C)$). Thus we have $i < n-1$ and both chains consist of at least two segments/rays.

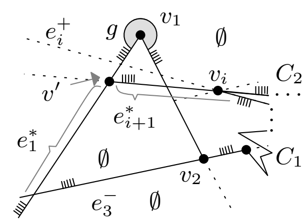

Now suppose that e_{i+1}^- intersects e_2 and thus e_1, and denote the point of intersection between e_{i+1}^- and e_1 by v'. Let e_1^* be the ray originating from v' in direction e_1, and let e_{i+1}^* denote the segment or ray (for $i = n-1$) originating from v' in direction e_{i+1}^-. Place a natural vertex guard g at v_1. Regardless of whether or not e_i^+ intersects e_2 and e_1, a guarding for C is provided by $g \cap (\mathcal{G}(C_1) \cup \mathcal{G}(C_2))$, with $C_1 = (e_3^-, \ldots, e_i^+)$ and $C_2 = (e_1^*, e_{i+1}^* \ldots, e_n)$ (if $i = n-1$ then $C_2 = (e_1^*, e_n^*)$). Observe that by the choice of v_i both C_1 and C_2 are simple and $\gamma(C) \leq 1 + \gamma(j) + \gamma(n-j-1)$, for some $2 \leq j \leq n-3$. As above, this yields $\gamma(C) \leq \lfloor (4n-1)/5 \rfloor$.

We have shown that in every case $\gamma(C) \leq \lfloor (4n-1)/5 \rfloor$ and as C was arbitrary it follows that $\gamma(n) \leq \lfloor (4n-1)/5 \rfloor$. □

Corollary 7. *Any simple polygon P with n edges can be guarded using at most $\lfloor (4n-2)/5 \rfloor$ guards.*

Corollary 8. *Let $P_1, \ldots, P_m$ be a collection of $m \geq 1$ simple polygons with n edges in total, and let R be a region that can be described as a formula composed of the operations intersection, union, and complement over the variables $\{P_1, \ldots, P_m\}$ in which each P_i appears exactly once. Then R can be guarded using at most $\lfloor (4n-2m)/5 \rfloor$ guards.*

Corollary 9. *Let P be any simple polygon P with h holes such that P is bounded by n edges in total. Then P can be guarded using at most $\lfloor (4n-2h-2)/5 \rfloor$ guards.*

4 Lower Bounds

For any natural number m we construct a polygon P_m with $2m$ edges which requires "many" guards. The polygon consists of spikes $S_1, S_2, ..., S_m$ arranged in such a way that the lines through both edges of a spike cut into every spike to the left (see Fig. 3). Denote the apex of S_i by w_i and its left vertex by v_i. The edge from v_i to w_i is denoted by e_i, the edge from w_i to v_{i+1} by f_i. We can construct P_m as follows: Consider the two hyperbolas $\left\{(x,y) \in \mathbb{R}^2 \mid x \geq 1, y = \frac{1}{x}\right\}$ and $\left\{(x,y) \in \mathbb{R}^2 \mid x \geq 1, y = -\frac{1}{x}\right\}$. Let $v_1 := (1,1)$ and $w_1 := (1,-1)$. Then choose f_1 tangential to the lower hyperbola. Let v_2 be the point where the tangent of the lower hyperbola intersects the upper hyperbola, that is, $v_2 = (1+\sqrt{2}, \frac{1}{1+\sqrt{2}})$. Choose w_2 to be the point where the tangent of the upper hyperbola in v_2 intersects the lower hyperbola, and proceed in this way. When reaching w_m, draw the last edge f_m from w_m to v_1 to close the polygon. Due to the convexity of the hyperbolas, P_m has the claimed property.

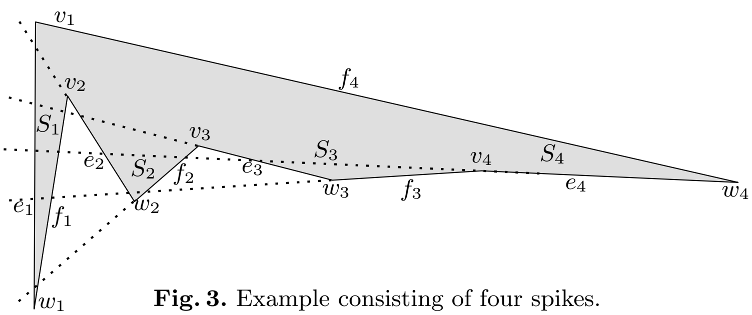

Fig. 3. Example consisting of four spikes.

No two edges of P_m are collinear. Consider the line arrangement defined by the edges of P_m. No two lines intersect outside P_m, unless one of them is the line through f_m. This leads to the following observation.

Observation 3. *In any guarding for P_m every 2-guard that does not cover f_m is a natural vertex guard.*

Theorem 10. *For any even natural number n there exists a simple polygon with n edges which requires at least $n-2$ natural guards.*

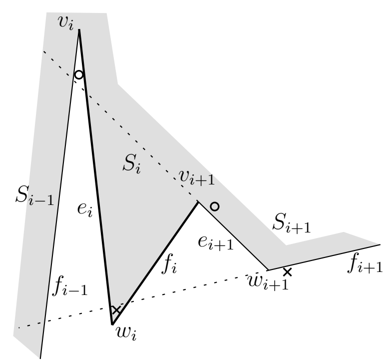

We say a guard *belongs* to a spike S_i if it is a natural edge guard on e_i or f_i or if it is a natural vertex guard on v_i or w_i. As only natural guards are allowed, every guard belongs to exactly one spike. The basic idea is that most spikes must have at least two guards. Obviously every spike S_i has at least one guard, since e_i must be covered (Lemma 1).

Lemma 11. *Consider a guarding $\mathcal{G}(P_m)$ using natural guards only, and let $i \in \{1, ..., m-1\}$. If only one guard belongs to S_i, then this guard must be on v_i or on e_i. If neither the guard at w_i nor the guard of f_i appear in $\mathcal{G}(P_m)$, then both the guard at v_{i+1} and the guard of e_{i+1} are in $\mathcal{G}(P_m)$.*

Proof. Assume only one guard from $\mathcal{G}(P_m)$ belongs to S_i. It cannot be the natural edge guard of f_i, because this would leave e_i uncovered (Lemma 1). If we had a guard on w_i only, there would be no guard to distinguish a point near v_i outside P_m from a point near v_{i+1} located inside P_m and below the line through f_i (see the two circles in the figure). Now assume there are no guards at w_i nor on f_i. Then to cover the edge f_i there must be a vertex guard on v_{i+1}. Furthermore, the edge guard on e_{i+1} is the only remaining natural guard to distinguish a point at the apex of S_i near w_i from a point located to the right of the apex of S_{i+1} near w_{i+1} and above the line through e_{i+1} (depicted by two crosses). □

This lemma immediately implies Theorem 10. Proceed through the spikes from left to right. As long as a spike has at least two guards which belong to it, we are fine. Whenever there appears a spike S_i with only one guard, we know that there must be at least two guards in S_{i+1} namely at v_{i+1} and on e_{i+1}. Either there is a third guard that belongs to S_{i+1}, and thus both spikes together have at least four guards; or again we know already two guards in S_{i+2}. In this way, we can go on until we either find a spike which at least three guards belong to or we have gone through the whole polygon. So whenever there is a spike with only one guard either there is a spike with at least three guards that makes up for it, or every spike till the end has two guards. Hence there can be at most one spike guarded by one guard only that is not made up for later. For the last spike S_m the lemma does not hold and we only know that it has at least one guard. So all in all there are at least $2(m-2)+1+1=n-2$ guards.

If we allow general (vertex) guards, it is possible to find guardings for P_m using roughly $2n/3$ guards, which is in accord with the lower bound in [3].

Theorem 12. *For any even natural number n there exists a simple polygon with n edges which requires at least* $\lceil(3n-4)/5\rceil$ *guards.*

Proof. Consider a polygon P_m as defined above, and let $\mathcal{G}(P_m)$ be a guarding for P_m. Define a to be the number of 2-guards in $\mathcal{G}(P_m)$, and let b be the number of other guards. All the n edges of P have to be covered somehow. An edge can be covered completely by a 2-guard, a 1-guard, or a $1'$-guard. If no guard covers it completely, then the edge must be covered by at least two guards partly (Lemma 1). Moreover, at least one of these guards, namely the one covering the section towards the right end of the edge, is a $0'$-guard, because the orientation can not be correct to cover a second edge. So if an edge e is not covered by a 2-guard, then there is at least one guard that does not cover any edge other than e. Therefore $2a+b\geq n$.

For any $i\in\{1,...,m-2\}$ let h_i be the directed line segment from the intersection of the lines through e_{i+1} and e_{i+2} to v_{i+2} (see Fig. 4). Similarly, let h'_i be the line segment from w_{i+1} to the intersection of the lines through f_i and f_{i+1}. As in Lemma 11, consider pairs $(p_1,q_1),...,(p_{m-2},q_{m-2})$ and $(p'_1,q'_1),...,(p'_{m-2},q'_{m-2})$ of points infinitesimally close to the starting point or the endpoint of the corresponding line segment, located as follows: $p_i,p'_i\in P_m$ for all i, $q_i,q'_i\notin P_m$ for all i, p_i is outside the natural vertex guard at w_{i+1}, whereas q_i is inside the natural vertex guard at w_{i+2}, and similarly, p'_i is outside the natural vertex guard at v_{i+2}, whereas q'_i is inside the natural vertex guard at v_{i+1}. There are $n-4$ such pairs, and they need to be distinguished somehow (Observation 1). Any natural vertex guard can distinguish at most one pair, and the same is true for any (non-natural) 2-guard located along the line through f_m. Thus any 2-guard in $\mathcal{G}(P_m)$ distinguishes at most one of the pairs (Observation 3).

We claim that every guard g in $\mathcal{G}(P_m)$ can distinguish at most three of these pairs. Denote the vertex of g by v_g, and let ℓ_g and r_g denote the left and right ray of g, respectively. Assume g distinguishes p_i from q_i. If v_g is to the left of h_i, then—in order to distinguish p_i from q_i—the ray r_g must intersect h_i. Symmetrically, if v_g is to the right of h_i, then ℓ_g must intersect h_i. Finally, if v_g is on the line through h_i then it must be on the line segment h_i itself. To

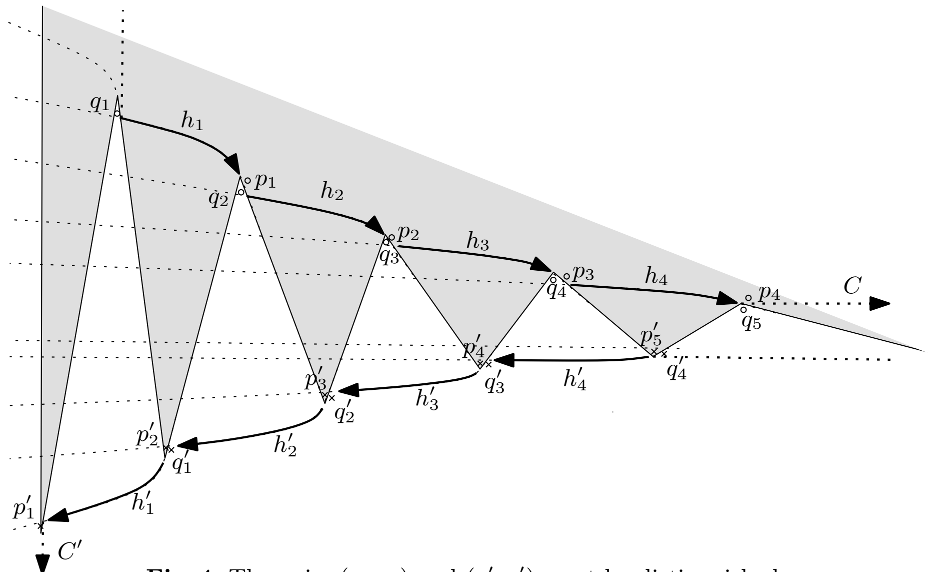

Fig. 4. The pairs (p_i, q_i) and (p'_i, q'_i) must be distinguished.

distinguish p_i from q_i, the endpoint of h_i (i.e. v_{i+2}) must be inside g (possibly on the boundary of g), hence ℓ_g must point to the left side of h_i or in the same direction as h_i, and r_g must point to the right side of h_i or in the same direction. Now assume g distinguishes p'_i and q'_i. If v_g is to the right of h'_i, then ℓ_g must intersect it, if it is to the left r_g must intersect it. If v_g lies on h'_i, ℓ_g leaves to the left and r_g to the right, or either or both rays lie on h'_i. In any case either ℓ_g intersects h_i (h'_i, respectively) coming from the right side of h_i (h'_i) and leaving to the left side, or r_g intersects h_i (h'_i) coming from the left side and leaving to the right, or ℓ_g starts on h_i (h'_i) itself leaving to the left or r_g starts on the line segment itself leaving to the right (see Fig. 5). If r_g leaves an oriented line segment to the right side of the segment or if ℓ_g leaves an oriented line segment to the left side, we say the ray *crosses* the line segment *with correct orientation*. So whenever a pair (p_i, q_i) or (p'_i, q'_i) is distinguished by g, then at least one of the rays ℓ_g or r_g has a correctly oriented crossing with h_i (h'_i, respectively). The line segments $h_1, ..., h_{m-2}$ lie on a oriented convex curve C, which we obtain by prolonging every line segment until reaching the starting point of the next one. Extend the first and last line segment to infinity vertically and horizontally, respectively. In the same way define a curve C' for $h'_1, ..., h'_{m-2}$ (see Fig. 4). Any ray can cross a convex curve at most twice. Because of the way C and C' are situated with respect to each other (a line that crosses C twice must have negative slope, to cross C' twice positive slope) a ray can intersect $C \cup C'$ at most three times. But we are only interested in crossings with correct orientation. If a ray crosses a curve twice, exactly one of the crossings has the correct orientation. If a ray crosses both C and C' once, exactly one of the crossings has the correct orientation. Therefore any ray can have at most two correctly oriented crossings. If one of the rays has two correctly oriented crossings, the other ray has at most

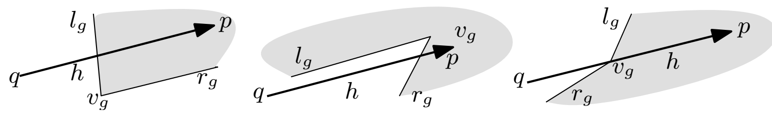

Fig. 5. Different ways g can distinguish p and q. In every case ℓ_g intersects h leaving to the left side or r_g intersects h leaving to the right.

one. Thus both rays together can have at most three correctly oriented crossings, and therefore g can distinguish at most three pairs. This leads to the second inequality $a + 3b \geq n - 4$. Both inequalities together imply $a + b \geq \frac{3n-4}{5}$. □

References

1. Aigner, M., Ziegler, G.M.: Proofs from THE BOOK. Springer, Berlin, 3rd edition (2003)
2. Chvátal, V.: A Combinatorial Theorem in Plane Geometry. J. Combin. Theory Ser. B 18, 39–41 (1975)
3. Damian, M., Flatland, R., O'Rourke, J., Ramaswami, S.: A New Lower Bound on Guard Placement for Wireless Localization. 17th Annual Fall Workshop on Computational Geometry (2007)
4. Dobkin, D.P., Guibas, L., Hershberger, J., Snoeyink, J.: An Efficient Algorithm for Finding the CSG Representation of a Simple Polygon. Algorithmica 10,1–23 (1993)
5. Eppstein, D., Goodrich, M.T., Sitchinava, N.: Guard Placement for Efficient Point-in-Polygon Proofs. In: Proc. 23rd Annu. Sympos. Comput. Geom., pp. 27–36 (2007)
6. Fisk, S.: A Short Proof of Chvátal's Watchman Theorem. J. Combin. Theory Ser. B 24, 374 (1978)
7. Fragoudakis, C., Markou, E, Zachos, S.: Maximizing the Guarded Boundary of an Art Gallery is APX-complete. Comput. Geom. Theory Appl. 38(3), 170–180 (2007)
8. Kahn, J., Klawe, M.M., Kleitman, D.J.: Traditional Galleries Require Fewer Watchmen. SIAM J. Algebraic Discrete Methods 4, 194–206 (1983)
9. Karavelas, M.I., Tsigaridas, E.P.: Guarding Curvilinear Art Galleries with Vertex or Point Guards. Rapport de recherche 6132, INRIA (2007)
10. Lee, D.T., Lin, A.K.: Computational Complexity of Art Gallery Problems. IEEE Trans. Inform. Theory 32(2), 276–282 (1986)
11. O'Rourke, J.: Galleries Need Fewer Mobile Guards: A Variation on Chvátal's Theorem. Geom. Dedicata 14, 273–283, 1983.
12. O'Rourke, J.: Visibility. In: J. E. Goodman and J. O'Rourke, editors, Handbook of Discrete and Computational Geometry, chapter 28, pp. 643–663. CRC Press LLC, Boca Raton, FL (2004)
13. O'Rourke, J.: Computational Geometry Column 48. ACM SIGACT News 37(3), 55–57 (2006)
14. Shermer, T.C.: Recent Results in Art Galleries. Proc. IEEE 80(9), 1384–1399 (1992)
15. Speckmann, B., Tóth, C.D.: Allocating Vertex Pi-guards in Simple Polygons via Pseudo-triangulations. Discrete Comput. Geom. 33(2), 345–364 (2005)
16. Urrutia, J.: Art Gallery and Illumination Problems. In: J.-R. Sack and J. Urrutia, editors, Handbook of Computational Geometry, pp. 973–1027. Elsevier Science Publishers B.V. North-Holland, Amsterdam (2000)