Guarding Art Galleries: The Extra Cost for Sculptures Is Linear*

Louigi Addario-Berry[1], Omid Amini[2], Jean-Sébastien Sereni[3], and Stéphan Thomassé[4]

[1] Department of Statistics, Oxford University, Oxford, UK
louigi@gmail.com
[2] Max-Planck-Institut für Informatik, Saarbrücken, Germany
amini@mpi-inf.mpg.de
[3] Institute for Theoretical Computer Science and Department of Applied Mathematics, Charles University, Prague, Czech Republic
sereni@kam.mff.cuni.cz
[4] LIRMM-Université Montpellier II, Montpellier, France
thomasse@lirmm.fr

Abstract. Art gallery problems have been extensively studied over the last decade and have found different type of applications. Normally the number of sides of a polygon or the general shape of the polygon is used as a measure of the complexity of the problem. In this paper we explore another measure of complexity, namely, the number of guards required to guard the boundary, or the walls, of the gallery. We prove that if n guards are necessary to guard the walls of an art gallery, then an additional team of at most $4n-6$ will guard the whole gallery. This result improves a previously known quadratic bound, and is a step towards a possibly optimal value of $n-2$ additional guards. The proof is algorithmic, uses ideas from graph theory, and is mainly based on the definition of a new reduction operator which recursively eliminates the simple parts of the polygon. We also prove that every gallery with c convex vertices can be guarded by at most $2c-4$ guards, which is optimal.

Keywords: Art Gallery, Pseudo-triangulation.

1 Introduction

Art gallery problems are, broadly speaking, the study of the relation between the shapes of regions in the plane and the number of points needed to *guard* them. The problem of determining how many guards are sufficient to see every point in the interior of an n-wall art gallery room was first posed by Klee [11]. Conceptually, the room is a simple polygon P with n vertices, and the guards are stationary points in P that can see any point of P connected to them by a straight line segment lying entirely within P. The first "art gallery theorem"

* This work was supported by the European project IST FET AEOLUS.

J. Gudmundsson (Ed.): SWAT 2008, LNCS 5124, pp. 41–52, 2008.

was obtained by Chvátal [3], who demonstrated that given any simple polygon with n sides, the interior of the polygon can be guarded with at most $\lfloor \frac{n}{3} \rfloor$ guards and that this number of guards is sometimes necessary. Fisk [8] later found a simpler proof which lends itself to an $O(n \log n)$-time algorithm developed by Avis and Toussaint [2] for locating these $\lfloor \frac{n}{3} \rfloor$ stationary guards. With some restriction on the shape of the polygon, for example if the polygon is rectilinear, that is, the edges of the polygon are either horizontal or vertical, Kahn *et al.* [12] have shown that $\lceil \frac{n}{4} \rceil$ guards are sufficient and sometimes necessary. Sack [20] and Edelsbrunner *et al.* [6] have devised an $O(n \log n)$-time algorithm to locate these $\lfloor \frac{n}{4} \rfloor$ guards. These classical results in the theory of art galleries have spawned a plethora of research (see the monograph by O'Rourke [18], and the surveys [21,23,25] for overviews of previous work). In particular, since then the art gallery problems have emerged as a research area that stresses complexity and algorithmic aspects of visibility and illumination in configurations comprising obstacles and guards.

In most of the reseach papers in the field, the number of sides of a polygon or restriction on the shape of the polygon is used as a very natural measure of the "complexity" of the polygon. The aim of this paper is to explore another measure of complexity, namely the number of guards required to guard the boundary, or the walls, of the gallery. As we will see in the next sections, this new complexity measure can be regarded as a mixture of the two named ones: the shape and the number of sides, but remains different and has its own characteristics. As shown in Figure 1, a team of guards inside a gallery can see the walls (where paintings are hung), without necessarily guarding the whole gallery (where sculptures are displayed), showing that these two notions of complexity are in general different. More precisely, the question we investigate in this paper is the following: given that the interior walls of a polygon can be guarded with at most n guards, how many *additional* guards may be needed to guard the whole interior? This question has been first explored by Aloupis *et all.* in [1] in their study of fat polygons. They proved that an additional number of at most $3n^2/2$ guards can guard the whole gallery.

Main Results. We prove the following linear bound.

Theorem 1. *Let M be a polygonal gallery. If the walls of M can be guarded by at most $n > 1$ guards, an additional set of $4n - 6$ guards is sufficient to guard the interior of M.*

Observe that when $n = 1$, the unique guard sees all the walls, hence sees the whole gallery. Most likely, the previous bound is not sharp. We offer the following conjecture.

Conjecture 1. *If the walls of a gallery can be guarded by $n > 1$ guards, then $n - 2$ additional guards are sufficient to guard the whole gallery.*

If Conjecture 1 is true then the given value would be optimal, as is shown by the example in Figure 2. In this example, there are $n-2$ "small rooms" attached by narrow entrances to a main room. Guarding the walls requires at most n guards

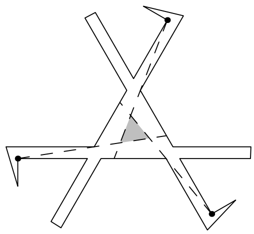

Fig. 1. Three guards are enough the guard the paintings (on the walls), but not the sculpture in the shaded area. Dashed lines are lines of sight of the guards.

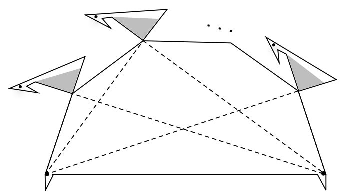

Fig. 2. Black dots indicate guards. The shaded areas indicate parts not seen by any of the current n guards, and dashed lines are lines of sight of the guards.

(as shown): one guard in each of the "small rooms" off the main room, and one guard each of the two far corners of the main room. These latter two guards each have a line of sight along one wall of each small room. However, with such a set of guards the parts of the gallery's interior shaded in dark grey are left unguarded. To guard the whole gallery requires *two* guards in each of the "small rooms", and an additional two guards in the main room.

The proof of Theorem 1 uses the fact that every gallery with c convex vertices can be guarded by at most $2c - 4$ guards. This latter result is optimal. In order to apply induction to bound the number of additional guards required to guard M, we first *reduce* our gallery to another gallery with certain guaranteed structural properties that make it easier to analyse. We do so by means of a new *transformation operator* $T(\cdot, \cdot)$, which takes as an argument a gallery N and a set of guards G that guards the walls of N, and returns another gallery N'. The operator T captures the complexity of the polygon by successfully deleting the parts which do not contribute to the main complexity. It has a nice definition and the general idea behind it may hopefully be applied to other contexts. On the complexity side, using the reduction operator T and earlier results [16,7], one can infer that the general problem of calculating the number of extra guards needed is NP-complete and does not admit a PTAS, i.e. is APX-hard.

Related Work. As we mentioned before, the literature on the art gallery problems is huge and different type of strategies and situations have been considered. Let us briefly review some of the works related to this paper. Laurentini [14] investigated the problem of covering the sides of the polygon and not necessarily the interior—related complexity questions being studied in another paper [9]. Efrat *et al.* [5] introduced the *link diagram* of a polygon. As we will see later, the last step in the proof of Theorem 1 is based on a certain kind of link diagram between the guards. It is interesting to explore the relations between the two notions. As the graph we use is based on the connectivity between guards, another related subject is that of the *guarded guard art gallery* problem [15,17]. In particular, one could investigate the guarded guard version of our problem.

Notations and Basic Definitions. Let us give some formal definitions. We let $\overline{S}$ be the closure of the set $S \subset \mathbf{R}^2$. A simply connected, compact set $M \subset \mathbf{R}^2$ is *polygonal* if its boundary ∂M is a simple closed polygon with finitely many vertices. The set M is *nearly* polygonal if M can be written as the union of polygonal galleries $M_1, \ldots, M_k$ such that

(i) for distinct $i, j \in \{1, \ldots, k\}$, letting $E_{i,j} = M_i \cap M_j$, either $E_{i,j} = \emptyset$ or $E_{i,j}$ contains a single point $e_{i,j}$; and
(ii) the connectivity graph with node set $\{v_1, \ldots, v_k\}$, where v_i represents the polygonal gallery M_i and nodes v_i and v_j are adjacent whenever $M_i \cap M_j \neq \emptyset$, is a tree.

We sometimes refer to the set ∂M as *the walls of* M, and to $M_1, \ldots, M_k$ as the *rooms* of M. A point $p \in M$ is a *cut-vertex* of M if $M \setminus \{p\}$ is not connected—so the cut-vertices of M are precisely the points $e_{i,j}$ defined above.

If M is a polygonal gallery, then we may describe M by simply listing the vertices of the polygon ∂M in their cyclic order, which we always assume is given in the "clockwise direction". Similarly, we may describe a nearly polygonal gallery M by listing the vertices of ∂M in cyclic order (again, in this paper always clockwise). If M is the nearly polygonal gallery described by $P = (p_1, \ldots, p_k, p_{k+1} = p_1)$, then M is polygonal precisely if P has no repeated points. Given points x and y of ∂M, by $\partial M[x, y]$ we mean the subset of ∂M starting at x and ending at y and following the cyclic order. These straightforward definitions and facts are depicted in Figures 3 and 4. We will also often abuse notation and write P or $P[x, y]$ in place of ∂M or $\partial M[x, y]$, respectively.

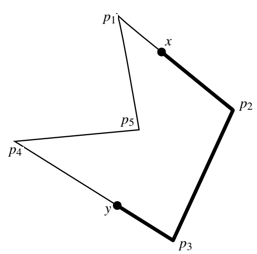

Fig. 3. (a) A polygonal gallery defined by the sequence $(p_1, p_2, p_3, p_4, p_5, p_1)$. The set $\partial M[x, y]$ is shown in bold.

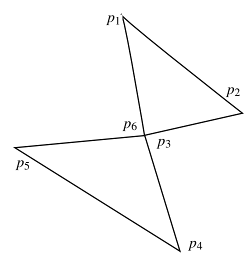

Fig. 4. (b) A nearly polygonal gallery defined by the sequence $(p_1, p_2, p_3, p_4, p_5, p_6, p_1)$, with $p_6 = p_3$

A *guard* is a point of M. A guard g *sees* a point p of M if the line segment $[g, p]$ is included in M. Unless otherwise stated, G is always a set of guards in M. We say that G *guards* M if every point of M is seen by a guard of G. Similarly, G *guards* ∂M (or G *guards the walls* of M) if every point of ∂M is seen by a guard of G. The *guarding number* of M is the minimum number of guards necessary to guard M.

2 Guards Versus Convex Vertices

Let $P = (p_1, \ldots, p_k, p_{k+1} = p_1)$ describe a nearly polygonal gallery M. The goal of this section is to prove that the guarding number of M is at most $2c - 4$, where c is the number of convex vertices of M. This bound is sharp: an example is given in Figure 5. The polygon shown in Figure 5 contains five convex vertices. To bound the number of guards required, consider the grey shaded region of the polygon. Regardless of how guards are placed outside the shaded region, the dark grey area remains unguarded, and no single guard can see all the dark grey area. Thus, the grey shaded region of the polygon must contain two guards. Similarly, the two other "concave triangular" areas must each contain two guards, for a total of six guards. This example can easily be generalised to show that for every $c \geq 5$, there is a polygon with c convex vertices requiring $2c - 4$ guards.

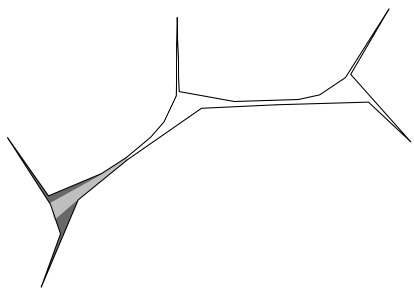

Fig. 5. A gallery with five convex vertices and guarding number six

We will use *pseudo-triangulations* of polygons to obtain our bound. A *pseudo-triangle* is a simple polygon with exactly three convex vertices. Given a simple polygon P, a *pseudo-triangulation* of P is a partition of the interior of P into non-overlapping pseudo-triangles whose vertices are all among vertices of P. We refer to the survey of Rote *et al.* [19] for further exposition about pseudo-triangulations. In our considerations, we need the following result.

Theorem 2. *Every simple polygon with k convex vertices admits a pseudo-triangulation consisting of $k - 2$ pseudo-triangles.*

It is easy to see that

Lemma 1. *The guarding number of a pseudo-triangle is at most 2. Moreover, it is one if the pseudo-triangle contains two consecutive convex vertices.*

The next theorem is a direct consequence of Theorem 2 and Lemma 1.

Theorem 3. *Let M be a polygonal gallery with c convex vertices for some integer $c \geq 3$. Then the guarding number of M is at most $2c - 4 - s$, where $s = 1$ if M contains two consecutive convex vertices, and $s = 0$ otherwise.*

More generally, using pseudo-triangulations, one can show [1]

Theorem 4. *Let M be a polygonal gallery with c convex vertices for some integer $c \geq 3$ such that these vertices appear in t chains of consecutive convex vertices. Then the guarding number of M is at most $c + t - 4$.*

As we don't use Theorem 4 in this full generality, we leave its proof to the full version of this paper.

3 Sculpture Galleries

We now turn our attention to the proof of Theorem 1, which we will prove inductively. For the purposes of our induction, we will in fact prove the following, stronger result.

Theorem 5. *Let M be a nearly polygonal gallery. If ∂M can be guarded with at most n guards, an additional set of $4n - 6$ guards is sufficient to guard M.*

In order to apply induction to bound the number of additional guards required to guard M, we first "reduce" M to another gallery M' with certain guaranteed structural properties that make it easier to analyse. We do so by means of a transformation operator $T(\cdot, \cdot)$, which takes as an argument a nearly polygonal gallery N and a set of guards G that guards the walls of N, and returns another nearly polygonal gallery N'.

Roughly speaking, the effect of T is to "trim off" a section of the polygon N that is unimportant to any of the lines of sight of the guards. Before defining T, then, we first formalise this notion of "importance". Let $U = U(N, G)$ be the set of points of N not seen by any guard $g \in G$. We say that a point p of N is *important* (with respect to N and G) if $p \in G$ or if $p \in \overline{U}$ or if p is a cut-vertex of N.

When there is no risk of confusion, we will write $T(N)$ instead of $T(N, G)$. We also remark that the operator T will be such that $T(N, G)$ is nearly polygonal, $U \subset T(N, G) \subset N$, $G \subset T(N, G)$, and G guards the walls of $T(N, G)$. We ask the reader to keep these properties in mind while reading the definition of $T(N, G)$, to which we now proceed.

3.1 The Definition of the Operator $T(N, G)$

Let N have rooms $N_1, \ldots, N_k$, and suppose that N is described by $P = (p_1, \ldots, p_m, p_1)$. We say N_i is a *leaf* if there is at most one $j \neq i$ such that $N_j \cap N_i \neq \emptyset$. By N_i^- we mean the set $N_i \setminus \cup_{j \neq i} N_j$. N_i is *empty* if $N_i^- \cap G = \emptyset$.

(A) If there is $1 \leq i \leq k$ such that N_i is an empty leaf then set $T(N) = \overline{N \setminus N_i^-} = \overline{N \setminus N_i}$.

[1] We are very gratefull to the first referee for pointing out that the proof of Theorem 3 we presented in the first version of this paper, which didn't use the pseudo-triangulations, could be simplified and generalised to derive theorem 4 without using the pseudo-triangulations.

(B) Otherwise, if N contains a cut-vertex p such that $p \notin G$ and such that for each simply connected component N^* of $N \setminus \{p\}$, $|G \cap \overline{N^*}| < |G|$ and $G \cap \overline{N^*}$ guards $\overline{N^*}$, then set $T(N) = N$.

(C) Otherwise, if every convex vertex of ∂N is important, set $T(N) = N$.

If none of (A),(B), or (C) occur, then ∂N contains a convex vertex p_i that is not important (and in particular is not a cut-vertex). Choose $x \in [p_{i-1}, p_i[$ and $y \in]p_i, p_{i+1}]$ such that $\triangle_{xp_iy}$ contains no important points except perhaps p_{i-1} and or p_{i+1}, with x as close to p_{i-1} as possible subject to this, and with y as close to p_{i+1} as possible subject to the previous constraints. We note that if $[x, y] \cap \partial N$ contains an interval of positive length, then $[x, y]$ must contain at least one guard; for, letting $[a, b]$ be some interval in $[x, y] \cap \partial N$, no finite set of guards lying outside $\triangle_{xp_iy}$ can see all of $\triangle_{ap_ib}$. (This situation is depicted in Figure 6.)

(D) If $G \cap]x, y[$ is non-empty choose points $x' \in]x, p_i[$ and $y' \in]p_i, y[$ arbitrarily. Let g be some guard of $G \cap]x, y[$ and let $T(N) = \overline{(N \setminus \triangle_{xp_iy}) \cup \triangle_{xx'g} \cup \triangle_{y'yg}}$. (This case is shown in Figure 7)

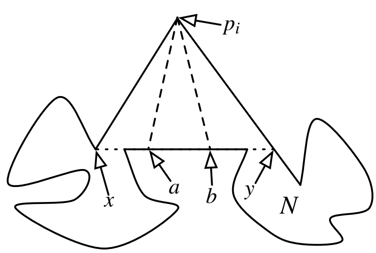

Fig. 6. No matter how guards are placed outside of $\triangle_{xp_iy}$, some part of $\triangle_{ap_ib}$ close to $[a, b]$ will not be seen by any guard

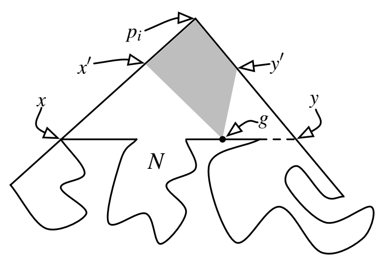

Fig. 7. The situation in case (D). The dark shaded region belongs to N but not to $T(N)$.

(E) Otherwise, suppose $x \in G$ or $y \in G$ – without loss of generality, we presume $x \in G$. Choose $z \in [x, y] \cap \partial N$ such that $[x, z]$ is not contained in ∂N, and as close to x as possible subject to this. Choose $z' \in \partial N$ very close to z and after z in the cyclic order. Finally, choose x' in $[x, y] \cap \partial N$ as far from x as possible such that $[x, x'] \subset \partial N$ (possibly $x' = x$), and set $T(N) = \overline{N \setminus \triangle_{x'zz'}}$. (This case is shown in Figure 8.)

If none of (A)-(E) occur then $[x, y] \cap G = \emptyset$, so $[x, y] \cap \partial N$ contains no interval of positive length.

(F) Otherwise, if $x = p_{i-1}$, $y = p_{i+1}$, or if $]x, y[\cap \partial N \neq \emptyset$, then set $T(N) = \overline{(N \setminus \triangle_{xp_iy})}$.

If (F) does not occur then we may assume without loss of generality that $x \neq p_{i-1}$. Since $[x, y] \cap G = \emptyset$ and $]x, y[\cap \partial M = \emptyset$, by our choice of x and y there must be $z \in [x, y[\cap \overline{U}$.

(G) Otherwise, if $x \in \overline{U}$, then set $T(N) = \overline{(N \setminus \triangle_{xp_iy})}$.

(H) Otherwise, let z be the point of $]x, y[$ that is closest to x such that $z \in \overline{U}$. Pick a point $z' \notin \triangle xp_iy$ chosen close to z in order to guarantee that $\triangle_{xzz'}$ does not contain a guard and is disjoint from U and from ∂N. Finally, set $T(N) = \overline{N \setminus (\triangle_{xp_iy} \cup \triangle_{x'zz'})}$. (This case is shown in Figure 9.)

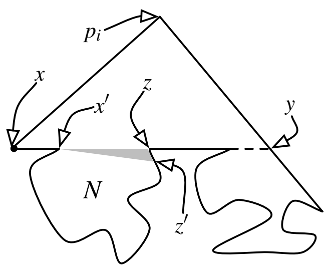

Fig. 8. The situation in case (E). The grey shaded region belongs to N but not to $T(N)$.

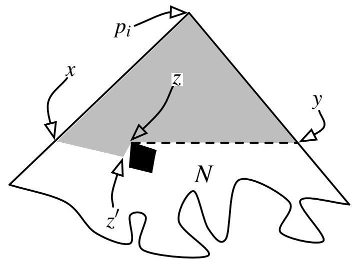

Fig. 9. The situation in case (H). The grey shaded region belongs to N but not to $T(N)$, and the black shaded region belongs to U.

When applying T repeatedly, we will usually write $T^2(N)$ in place of $T(T(N))$. As mentioned at the start of Section 3.1, a key property of this (rather cumbersome) transformation is that if N and G satisfy the hypotheses of Theorem 5 then $T(N)$ and G also satisfy the hypotheses of Theorem 5. Another important property of T is that its repeated application is guaranteed to increase the value of a certain bounded invariant that can be associated to gallery-guard set pairs N, G, and so by applying T to any such pair N, G enough times, we are guaranteed to reach a fixed point of the transformation T.

To define this invariant, we first introduce one additional piece of notation. Let C denote the set of cut-vertices of N, and, for $c \in C$, let $\kappa(c)$ be the number of simply connected components of $N \setminus \{c\}$. The invariant, which we denote $\Phi(N, G)$ (or $\Phi(N)$ when there is no risk of confusion), is equal to the number of convex vertices which are guards, plus the number of convex vertices in $\overline{U}$, plus $\sum_{c \in C \cap G}(\kappa(c) - 1)$.

We hereafter refer to vertices that are also guards as *occupied*, and vertices that are in $\overline{U}$ as *critical*. We observe that every occupied (resp. critical) convex vertex of N is an occupied (resp. critical) convex vertex of $T(N)$. It thus follows from the definition of Φ that $\Phi(T(N), G) \geq \Phi(N, G)$. The main property that makes this invariant useful is captured by the following lemma.

Lemma 2. *For any nearly polygonal gallery M with no set of empty leaves, and any finite set of guards G that see all of ∂M, if $T^2(M) \neq T(M) \neq M$ then either $\Phi(T(M)) > \Phi(M)$ or $T(M)$ has strictly fewer vertices than M.*

Let us postpone the proof of Lemma 2 to the end of the paper.

3.2 The Proof of Theorem 5

Let M and G be as in the statement of Theorem 5 and let $g = |G|$. If $g = 1$, the unique guard sees the whole gallery M. Next suppose that $g > 1$ and that the statement of Theorem 5 holds for all values $n < g$. As previously, we let $U = U(M, G)$ be the (open) set of points of the gallery M that are not seen by any guard of G.

Let $M_0 = M$; for $i \geq 1$ set $M_i = T(M_{i-1}, G)$, It turns out that the number of critical convex vertices in *all* of the galleries M_i can be bounded uniformly in terms of g; this is the substance of the following lemma.

Lemma 3. *For all $i \geq 0$, there are at most $g - 1$ critical convex vertices in M_i.*

We will prove this lemma along with Lemma 2, at the end of the paper.

We observe that M contains precisely $1 + \sum_{\{c \in C : \kappa(c) > 2\}} (\kappa(c) - 1)$ leaves (this can be seen by a straightforward induction). Since M contains no empty leaves, it follows that

$$g - 1 \geq \sum_{\{c \in C : \kappa(c) > 2\}} (\kappa(c) - 1) \geq \sum_{\{c \in C \cap G : \kappa(c) > 2\}} (\kappa(C) - 1)\,,$$

so

$$\sum_{c \in C \cap G} (\kappa(C) - 1) \leq g + \sum_{\{c \in C \cap G : \kappa(c) > 2\}} (\kappa(C) - 1) \leq 2g - 1\,.$$

Furthermore, there are at most g occupied convex vertices. It follows by Lemma 3, the above inequalities and the definition of Φ that $\Phi(M_i) \leq 4g - 1$ for all i. By the observation immediately preceding Lemma 2, $\Phi(M_{i+1}) \geq \Phi(M_i)$ for all i, and Lemma 2 then implies that there exists an integer j such that $T(M_j) = M_j$. In this case, by the definition of the operator $T(\cdot)$, one of (B) or (C) occurs for M_j.

We now show that an additional set of at most $4g-6$ guards suffices to guard M_j. Since $U \subset M_j$, these guards also guard U in M_j; since $M_j \subseteq M$, these guards also guard U in M; so together with G, they guard all of M, as claimed. We now assume, purely for the ease of exposition, that $j = 0$, i.e., that $P_j = P$ and $M_j = M$; this eases the notational burden without otherwise changing the proof.

If (B) occurs then we let p be a cut-vertex as described in (B). Let $N_1^-, \ldots, N_r^-$ be the simply connected components of $M - \{p\}$, and for $i \in \{1, \ldots, r\}$ let $N_i^+ = \overline{N_i^-}$, let $G_i = N_i^+ \cap G$ and let $g_i = |G_i|$. Since $p \notin G$, we have $\sum_{i=1}^r g_i = g$. Furthermore, since G_i guards N_i^+ and $g_i < g$ for all $i \in \{1, \ldots, r\}$, the induction hypothesis implies the existence of $H_i \subset N_i^+$ such that $|H_i| \leq 4g_i - 6$ and $G_i \cup H_i$ guards N_i^+. In this case $\bigcup_{i=1}^r (G_i \cup H_i)$ guards M, and

$$\left| \bigcup_{i=1}^r H_i \right| \leq \sum_{i=1}^r (4g_i - 6) \leq 4 \sum_{i=1}^r g_i - 6r \leq 4r - 8 < 4r - 4\,.$$

If (C) occurs then since the number of critical convex vertices is at most $g-1$ by Lemma 3, and the number of occupied vertices is at most g, the total number of convex vertices of our gallery M is at most $2g-1$. It follows from Theorem 3 that there is a set H of at most $2(2g-1)-4=4g-6$ additional guards that guard M, and hence guard U. □

3.3 Proofs of Lemmas 2 and 3

Having established Theorem 1 assuming that Lemmas 2 and 3 hold, we now turn to the proofs of these lemmas.

Proof (of Lemma 2). Let M and G be as in the statement of the lemma, and suppose that $T(M) \neq M$. Let $M_1 = T(M)$ and let $M_2 = T^2(M)$. For any of the conditions (A)-(H) in the description of T – say (E), for example – we will use the shorthand "(E) holds for M (resp. M_1, M_2)" if, letting $N = M$ (resp. M_1, M_2) in the definition of $T(\cdot)$, the condition described in (E) holds and none of the earlier conditions hold.

By the definition of $T(\cdot)$, if $M_1 \neq M$, then one of (A) or (D)-(H) must hold for M. We now show that in each case, either M_1 has strictly fewer vertices than P, $\Phi(M_1) > \Phi(M)$, or $M_2 = M_1$.

– If (A) holds for M then $T(N)$ has strictly fewer vertices than N.
– If (D) holds for M then g is a cut-vertex in M_1 and, more strongly, $M_1 \setminus \{g\}$ has strictly more connected components than $M \setminus \{g\}$. Thus $\Phi(M_1) > \Phi(M)$.
– If (E) holds for M then since $\mathrm{int}(\triangle_{xp_iy}) \cap G = \emptyset$, x' is a cutvertex in M_1 and, more strongly, (B) holds for M_1 (with $p = x'$). Thus $M_2 = M_1$.
– If (F) holds for M and $x = p_{i-1}$, $y = p_{i+1}$, then M_1 contains strictly fewer vertices than M. If (F) holds for M and $]x, y[\cap \partial M \neq \emptyset$, we first observe that since $[x, y] \cap G = \emptyset$, no guard on the line through x and y can see $\mathrm{int}(\triangle_{xp_iy})$. As every point in $\mathrm{int}(\triangle_{xp_iy})$ is guarded, it follows that every point in $[x, y]$ is seen by some guard g *not* on the line through $[x, y]$. Now let z be a point in $]x, y[\cap \partial M$; then z is a cut-vertex in M_1. Furthermore, by the above comments it must be the case that (B) holds for M_1 (with $p = z$). Thus $M_2 = M_1$.
– If (G) holds for M then x is a critical convex vertex in M_1 but not in M, so $\Phi(M_1) > \Phi(M)$.
– Finally, if (H) holds for M then z is a critical convex vertex in M_1 but not in M, so $\Phi(M_1) > \Phi(M)$.

This completes the proof of Lemma 2.

Proof (of Lemma 3). Let $P = (p_1, \ldots, p_k, p_1)$ describe M. We form a graph $\mathcal{G}$ whose vertex set is the set G of guards of M. For every critical convex vertex p_j, we choose one guard g_i that sees some non-empty interval $]x_j, p_j[$ of $[p_{j-1}, p_j]$, and one guard g'_j that sees some non-empty interval $]p_j, y_j[$ of $[p_j, p_{j+1}]$, and add the edge $g_j g'_j$ to $\mathcal{G}$. By construction, the number of critical convex vertices of P

is at most the number of edges of $\mathcal{G}$. We shall show that $\mathcal{G}$ contains no cycles, from which the conclusion immediately follows.

We first observe that if p_j is a critical convex vertex of P, then the angle $p_{j-1}p_jp_{j+1}$ is strictly positive. Therefore, g_j is different from g_{j+1}, or else the quadrilateral $g_jx_jp_jy_j$ would be entirely seen by g_j, contradicting the fact that p_j is in the closure of U. It follows that $\mathcal{G}$ contains no loops (cycles of length 1).

Next, suppose that $\mathcal{G}$ contains a cycle $g_1, g_2, \ldots, g_k, g_{k+1} = g_1$ with $k \geq 2$ (in which case $g_{i+1} = g'_i$ and g_ig_{i+1} is the edge corresponding to some critical convex vertex p_i, for $i \in \{1, \ldots, k\}$). In this case, the polygonal line

$$\mathrm{PL} = (g_1, p_1, g_2, p_2, \ldots, g_k, p_k, g_1)\,,$$

which is not necessarily simple or even uncrossing, contains some simple, closed polygonal line $\mathrm{PL}_1 = (x, g_i, p_i, \ldots, x)$ or $\mathrm{PL}_2 = (x, p_i, g_{i+1}, \ldots, x)$. We emphasise that though a line segment $[g_i, g_{i+1}]$ may not be contained within M, PL is fully contained within M.

Given a critical convex vertex p_j, as the angle at p_j is convex, p_j can only appear in PL as the endpoint of a line segment. Furthermore, by definition there is no guard, so no vertex of $\mathcal{G}$, at position p_j. It follows that PL contains each of $p_1, \ldots, p_k$ exactly once, so the point x is not the point p_i of PL_1 or PL_2. Suppose first that PL contains a closed circuit such as PL_1. Since x is not p_i, p_i is proceeded by $g_{i+1} = g'_i$ in PL_1. Since PL_1 is simple, its interior (the bounded component of $\mathbf{R}^2 \setminus \mathrm{PL}_1$) lies entirely within M. Since p_i is convex, the guard g_i sees a non-empty interval $]p_i, y[$ for some $y \in]p_i, g'_i[$. This means that the triangle g_ip_iy is entirely seen by g_i, which contradicts the fact that p_i is in the closure of U. A similar contradiction occurs when considering PL_2 instead of PL_1. Therefore, $\mathcal{G}$ contains no cycles of length at least 2, so no cycles at all, and hence has at most $g - 1$ edges.

4 Conclusion

It would be interesting to consider the approxamibility of the problem. In particular, we do not know if the problem admits a constant factor approximation (the best approximation algorithm [10,4] for the general art gallery problems has ratio log(OPT)). The generalisation of the problem to three dimensions is also another natural question to investigate.

References

1. Aloupis, G., Bose, P., Dujmović, V., Gray, C., Langerman, S., Speckmann, B.: Guarding Fat Polygons and Triangulating Guarded Polygons, September (preprint, 2007)
2. Avis, D., Toussaint, G.T.: An efficient algorithm for decomposing a polygon into star-shaped polygons. Pattern Recognition 13(6), 395–398 (1981)
3. Chvatal, V.: A combinatorial theorem in plane geometry. J. Combin. Theory Ser. B 18, 39–41 (1975)

4. Deshpande, A., Taejung, K., Demaine, E.D., Sarma, S.E.: A Pseudopolynomial Time O(log $c_{opt)}$-Approximation Algorithm for Art Gallery Problems. In: Dehne, F., Sack, J.-R., Zeh, N. (eds.) WADS 2007. LNCS, vol. 4619, pp. 163–174. Springer, Heidelberg (2007)
5. Efrat, A., Guibas, L.J., Har-Peled, S., Lin, D.C., Mitchell, J.S.B., Murali, T.M.: Sweeping simple polygons with a chain of guards. In: Proc. 11th annual ACM-SIAM symposium on Discrete algorithms, pp. 927–936 (2000)
6. Edelsbrunner, H., O'Rourke, J., Welzl, E.: Stationing guards in rectilinear art galleries. Computer vision, graphics, and image processing 27(2), 167–176 (1984)
7. Eidenbenz, S., Stamm, C., Widmayer, P.: Inapproximability Results for Guarding Polygons and Terrains. Algorithmica 31(1), 79–113 (2001)
8. Fisk, S.: A short proof of Chvatals watchman theorem. J. Combin. Theory Ser. B 24(3), 374 (1978)
9. Fragoudakis, C., Markou, E., Zachos, S.: Maximizing the guarded boundary of an Art Gallery is APX-complete. Comput. Geom. 38(3), 170–180 (2007)
10. Ghosh, S.: Approximation algorithms for art gallery problems. In: Proc. Canadian Inform. Process. Soc. Congress, pp. 429–434 (1987)
11. Honsberger, R.: Mathematical Gems II. Dolciani Mathematical Expositions No. 2. Mathematical Association of America, Washington, pp. 104–110 (1976)
12. Kahn, J., Klawe, M., Kleitman, D.: Traditional Galleries Require Fewer Watchmen. SIAM J. Algebraic Discrete Methods 4, 194–206 (1983)
13. Kranakis, E., Pocchiola, M.: Camera placement in integer lattices. Discrete and Comput. Geom. 12(1), 91–104 (1994)
14. Laurentini, A.: Guarding the walls of an art gallery. The Visual Computer 15(6), 265–278 (1999)
15. Liaw, B.C., Huang, N.F., Lee, R.C.T.: The minimum cooperative guards problem on k-spiral polygons. In: Proc. of 5th Canadian Conference on Computational Geometry, pp. 97–101 (1993)
16. Lee, D., Lin, A.: Computational complexity of art gallery problems. IEEE Trans. Inform. Theory 32(2), 276–282 (1986)
17. Michael, T., Pinciu, V.: Art gallery theorems for guarded guards. Comput. Geom. 26(3), 247–258 (2003)
18. O'Rourke, J.: Art gallery theorems and algorithms. Oxford University Press, New York (1987)
19. Rote, G., Santos, F., Streinu, I.: Pseudo-triangulations: A survey, Arxiv preprint math/0612672v1 (December 2006)
20. Sack, J.R.: An $O(n \log n)$ algorithm for decomposing simple rectilinear polygons into convex quadrilaterals. In: Proc. 20th Conference on Communications, Control, and Computing, pp. 64–74 (1982)
21. Shermer, T.C.: Recent results in art galleries [geometry]. Proc. IEEE 80(9), 1384–1399 (1992)
22. Speckmann, B., Tóth, C.D.: Allocating Vertex p-Guards in Simple Polygons via Pseudo-Triangulations. Discrete Comput. Geom. 33(2), 345–364 (2005)
23. Szabo, L.: Recent results on illumination problems. Bolyai Soc. Math. Stud. 6, 207–221 (1997)
24. Tóth, L.F.: Illumination of convex discs. Acta Math. Hungar. 29(3-4), 355–360 (1977)
25. Urrutia, J.: Art gallery and illumination problems. Handbook of Computational Geometry, 973–1027 (2000)