Guarding Art Galleries: The Extra Cost for Sculptures is Linear

Louigi Addario-Berry, Omid Amini, Jean-Sébastien Sereni, Stéphan Thomassé

▶ To cite this version:

Guarding Art Galleries: The Extra Cost for Sculptures is Linear[*]

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

[1] Department of Statistics, Oxford University, Oxford, UK
[2] Max-Planck-Institut für Informatik, Saarbrücken, Germany
[3] Institute for Theoretical Computer Science (ITI) and Department of Applied Mathematics (KAM), Charles University, Prague, Czech Republic
[4] LIRMM-Université Montpellier II, Montpellier, France

Abstract. 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. These problems have been extensively studied over the last decade and have found different type of applications in practical situation. 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. The aim of this paper is to present and 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 (visibility graph induced on the already placed guards), and is mainly based on the definition of a new reduction operator which recursively eliminates the simple parts of the polygon. We also use the fact that every gallery with c right-turn angles can be guarded by at most $2c - 4$ guards. This latter result is optimal.

1 Painting Galleries

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 Victor Klee [Hon76]. 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 that lies entirely within P. The first "art gallery theorem" was proved by Chvátal [Chv75], 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 [Fis78] later found a simpler proof which lends itself to an $\mathcal{O}(n \log n)$ algorithm developed by Avis and Toussaint [AT81] 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. [KKK83] have shown that $\lceil \frac{n}{4} \rceil$ guards are sufficient and sometimes necessary. Sack [Sac82] and Edelsbrunner, et al. [EOW84] have, based on the results of [KKK83] and [O'R83], respectively, devised an $\mathcal{O}(n \log n)$ algorithm for locating 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 [O'R87], and the surveys [Urr00,Sza97,She92] for overviews of previous work). In particular, since then the art gallery problems have emerged as a research area that stress complexity and algorithmic aspects of visibility and illumination in configurations comprising obstacles and guards. In fact by creating rather idealised situations the theory succeeds in abstracting the algorithmic essence
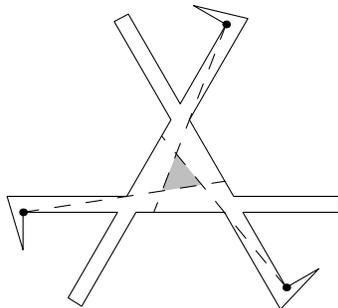

Fig. 1. Three guards are enough the guard the paintings (on the walls), but not the sculpture in the shaded area. Dashed lines are lines of sight of the guards.

of many visibility problems (like in partitioning theorems, mobile guard configurations, visibility graphs, . . .) thus significantly facilitating the study of their computational complexity, see [She92] and [KP94]. 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 present and 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 quite 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 Dujmović and Bose [DB07], who proved that an additional number of at most $3n^2/2$ guards can guard the whole gallery.

Main Results. The main result of this paper provides the following linear bound:

Theorem 1. *Let M be a polygonal gallery. If the walls of M can be guarded with at most n 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. The previous bound is then sharp for n equal to one, but is most likely not for larger values of n. We offer the following conjecture.

Conjecture 1 *Let n be an integer greater than one. If the interior walls of a gallery can be guarded by n guards, then $n - 2$ additional guards are sufficient to guard the whole gallery.*

As $n - 2$ is also the number of triangles in a triangulation of an n-sided polygon, it is tempting to try to prove the above theorem by finding a relationship between the original polygon P and some auxiliary polygon whose vertices are the set of guards that guard the interior boundary of P. We were not able to prove the conjecture in this way.

If Conjecture 1 is true then the given value would be optimal, as is shown by following example.
In Figure 2, $n - 2$ "small rooms" are attached by narrow entrances to a main room. Guarding the walls requires at most n 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.
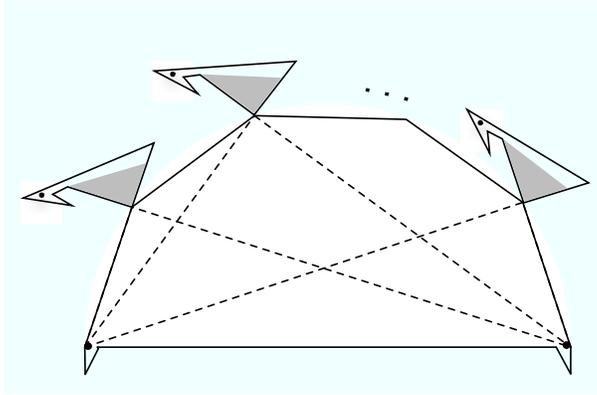
Fig. 2. Black dots indicate guards. The shaded areas indicate parts not seen by any of the current n guards, and dashed lines are lines of sight of the guards.

The proof of Theorem 1 uses the fact that every gallery with c *right-turn* angles 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 could hopfully be applied to other contexts.

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 by different authors. The books by O'Rourke [O'R87] and Urrutia [Urr00] provide a nice introduction to the topic. Here we cite some of the works related to this paper. In [Lau99], the author introduces and investigates the problem of covering the sides of the polygon and not necessarily the interior, see also [FMZ07] for complexity results. In [EGHP$^+$00], the notion of *link diagram* of a polygon has been introduced. 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 two notions. As the graph we used is based on the connectivity between guards, another related subject is that of the *guarded guard art gallery* problem [LHL93,MP03]. In particular, it is interesting to investigate the guarded guard version of our problem. The recent work of Speckmann and Tóth [ST05] suggests the use of pseudo-triangulations (see [PV96] for a survey) in the study of art-gallery problems. The result of Section 2 can be rewritten in the modern language of pseudo-triangulations, but due to the lack of space, our aim to simplify the presentation, and the fact that we were not able to apply the techniques to prove our main Theorem, we have decided to postpone these connections to the full version of this paper.

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

(i) for all $1 \le i \ne j \le k$, letting $E_{i,j} = M_i \cap M_j$, either $E_{i,j} = \emptyset$ or $E_{i,j}$ contains a single point, say $E_{i,j} = \{e_{i,j}\}$; and
(ii) the graph with vertex set $\{v_1, \ldots, v_k\}$ and edge set $\{v_i v_j : E_{i,j} \ne \emptyset\}$ 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. We say $p \in M$ is a *cut-vertex* of M if $M \setminus \{p\}$ is not connected – so the cut-vertices of M are precisely the points $e_{i,j}$, 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 the 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.
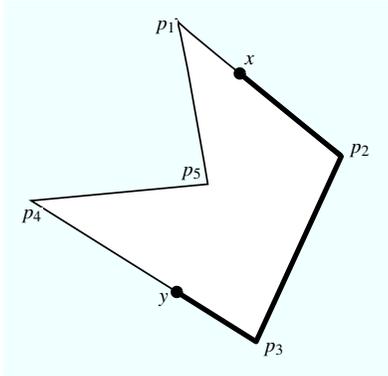

Fig. 3. (a) A polygonal gallery defined by the sequence (p_1, p_2, p_3, p_1). The set $\partial M[x, y]$ is shown in bold.
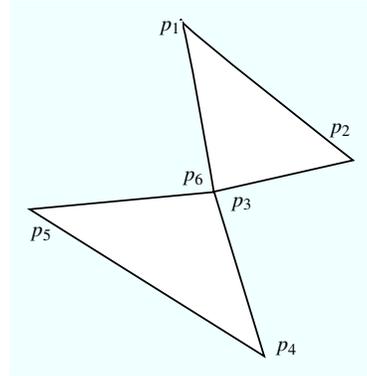
Fig. 4. (b) A nearly polygonal gallery defined by the sequence $(p_1, p_2, p_3, p_4, p_5, p_6, p_1)$, with $p_6 = p_3$.

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 Right-Turn 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 right-turn vertices of M. This bound is sharp: an example is given in Figure 5. The polygon shown in Figure 5 contains five right-turn 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 right-turn vertices requiring $2c - 4$ guards.

This bound on the guarding number in terms of the number of right-turn vertices follows from a stronger bound that applies only to polygonal galleries. We say that M is *special* (or that P is special)
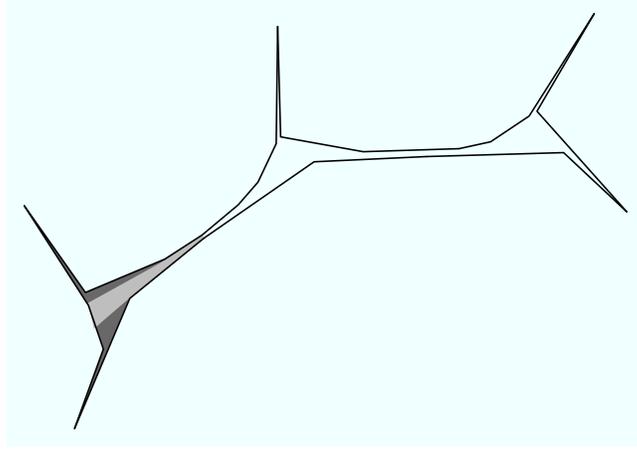
Fig. 5. A gallery with five right-turn vertices and guarding number six.

if it contains two consecutive right-turn vertices. We prove that if M is polygonal, then the guarding number of M is at most $2c - 4 - s$, where $s = 1$ if it is special, and $s = 0$ otherwise. The proof of the following (easy) lemma illustrates a technique we will widely use.

Lemma 1. *If P has exactly three right-turn vertices and they are consecutive, then p_i guards M.*

Proof. It is easily verified that if P has three right-turn vertices then M is polygonal. Suppose that the right-turn vertices of P are precisely p_{i-1}, p_i and p_{i+1}. Let x be the last point of $P[p_{i+1}, p_{i-1}]$ with respect to the cyclic order such that $P[p_{i+1}, x]$ is seen by p_i. We consider two cases.

  The polygon P is disjoint from $]p_i, x[$. In this case, the point x is not a right-turn vertex in $P' = P[x, p_i] \cup [p_i, x]$ (otherwise p_i would see a point further than x). This is impossible since P' would only have two right-turn vertices.
  The polygon P intersects $]p_i, x[$. Let $y \in P \cap]p_i, x[$. The point y is not a right-turn vertex in $P' = P[p_i, y] \cup [y, p_i]$. Since P' cannot have only two right-turn vertices, we infer that $x = p_{i-1}$.

This means that p_i sees all of P, and hence guards M. ∎

Lemma 2. *If P has three right-turn vertices and is special, then M has guarding number one.*

Proof. If the right-turn vertices are consecutive, the result follows from Lemma 1. Otherwise, we let p_1 and p_2 be the consecutive right-turn vertices, and p_i be the third one. Let p be a point of $P]p_{i+1}, p_{i-1}[$ that is seen by p_i. Observe that both $P_1 = P[p_i, p] \cup [p, p_i]$ and $P_2 = P[p, p_i] \cup [p_i, p]$ have three consecutive right-turn vertices in which p is the middle vertex. Thus, by Lemma 1, p guards both P_1 and P_2. ∎

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

Proof. We prove the result by induction on $c - s$. The case $c - s = 2$ follows from Lemma 2. Our approach to the inductive step depends on whether or not M is special.

  Case 1: *the gallery M is not special.* We pick a right-turn vertex p_i. Since both p_{i+1} and p_{i-1} are not right-turn vertices, there is a point p of $P[p_{i+2}, p_{i-2}]$ which is seen by p_i. Let c_1 and c_2 be the number of right-turn vertices of $P_1 = P[p, p_i] \cup [p_i, p]$ and $P_2 = P[p_i, p] \cup [p, p_i]$, respectively. Note that both

P_1 and P_2 are special as c_1 and c_2 are right-turn vertices in P_1 and P_2, respectively. Observe also that $c = c_1 + c_2 - 3$, unless p is a right-turn vertex of P in which case $c = c_1 + c_2 - 2$. Thus $c_1 \leq c$ and $c_2 \leq c$. Hence, by the induction hypothesis applied to P_1 and P_2, we infer that the guarding number of P_i is at most $2c_i - 5$, for each $i \in \{1, 2\}$. Therefore, the guarding number of M is at most $2c_1 - 5 + 2c_2 - 5 \leq 2(c + 3) - 10 = 2c - 4$.

Case 2: *The gallery M is special.* We consider two consecutive right-turn vertices p_i and p_{i+1}. Let p_j be the first right-turn vertex before p_i and p_k be the first right-turn vertex after p_{i+1}. If p_j sees p_k we form two polygons $P_1 = P[p_k, p_j] \cup [p_j, p_k]$ and $P_2 = P[p_j, p_k] \cup [p_k, p_j]$ with number of right-turn vertices c_1 and c_2, respectively. Since $c_1 + c_2 = c + 2$ and both P_1 and P_2 are special, it follows from the induction hypothesis that the guarding number of M is at most $2c - 6$.

Assume now that p_j does not see p_k. Let $x \in P]p_j, p_i]$ and $y \in P[p_{i+1}p_k[$ be chosen such that

(i)   the point x sees y;

(ii) the point x sees no point of $P]y, p_k]$; and

(iii) the point y sees no point of $P[p_j, x[$.

We observe that $]x, y[$ is disjoint from $P[x, y]$, and consider two cases.

The curve $P[y, x]$ is disjoint from $]x, y[$. We form two polygons $P_1 = P[x, y] \cup [y, x]$ and $P_2 = P[y, x] \cup [x, y]$. Observe that both x and y are right-turn vertices in P_2. In particular P_1 and P_2 together have $c + 2$ right-turn vertices, and since they are special, the induction hypothesis yields the result.

The curve $P[y, x]$ intersects $]x, y[$. We construct a sequence of points z_i, x_i as follows. Let z_1 be the point of $P[y, x] \cap]x, y[$ closest to x. Let x_1 be the point of $P[p_j, x]$ seen by z_1 that is closest to p_j. If either $x_1 = p_j$ or $]z_1, x_1[$ is disjoint from P, we stop the sequence. Otherwise, we let z_2 be the point of $P \cap]x_1, z_1[$ closest to x_1, and x_2 be the point of $P[p_j, x_1]$ seen by z_2 that is closest to p_j. We iterate this construction, which terminates at some pair x_i, z_i such that either $x_i = p_j$ or $]x_i, z_i[$ is disjoint from P. We partition P into $P_1 = P[x_i, z_i] \cup [z_i, x_i]$ and $P_2 = P[z_i, x_i] \cup [x_i, z_i]$. Let c_1 and c_2 be the number of right-turn vertices of P_1 and P_2, respectively. We observe that the vertex z_i is not a right-turn vertex of P_1, as the angle at z_i in P is greater than the angle $x_{i-1}z_iz_{i-1}$ (where if $i = 1$ then $x_{i-1} = x$ and $z_{i-1} = y$), and this latter angle is precisely π. If the sequence terminates at $x_i = p_j$, the vertex p_j is now a right-turn vertex in both P_1 and P_2, thus $c_1 + c_2 = c + 2$. Furthermore, P_1 is special as the angle formed by z_i in P_1 is also a right-turn vertex, and P_2 is special as both p_i and p_{i+1} are vertices of P_2. Therefore, by induction, the guarding number of M is at most $2c_1 - 5 + 2c_2 - 5 < 2c - 5$. If the sequence terminates with $]x_i, z_i[$ disjoint from P, since x_i is not a right-turn vertex in P_2, we deduce that $c_1 + c_2 = c + 2$. The polygon P_1 is still special, and $c_2 \leq c - 2$ as none of p_i, p_{i-1}, and p_k are contained in P_2. It follows that we may apply induction to both P_1 and P_2, so the total number of guards needed is at most $2c_1 - 5 + 2c_2 - 4 \leq 2c - 5$.

∎

It follows easily that a similar bound also holds for nearly polygonal galleries.

Corollary 1. *Let M be a nearly polygonal gallery with c right-turn vertices that are not cut-vertices. Then the guarding number of M is at most $2c - 4$.*

Proof. Let M be a nearly polygonal gallery. We write M as the union of polygonal galleries $M_1, \ldots, M_k$ and denote by c_i the number of right-turn vertices of M_i, for $i = 1, \ldots, k$.

Letting $e_{i,j}$ be the unique element of $M_i \cap M_j$ for pairs $i \neq j$ with $M_i \cap M_j \neq \emptyset$, we observe that if $M_i \cap M_j \neq \emptyset$ then e_i may be a right-turn vertex in M_i and/or in M_j, but is certainly not a right-turn vertex in any room to which it does not belong. Thus,

$$\sum_{i=1}^{k} c_i \leq c + 2|\{\{i, j\} : 1 \leq i \neq j \leq k, M_i \cap M_j \neq \emptyset\}|.$$

Furthermore, by condition (ii) in the definition of nearly polygonal galleries, it follows that at most $k-1$ of the pairs $\{\{i,j\} : 1 \leq i \neq j \leq k\}$ have $M_i \cap M_j \neq \emptyset$, so $\sum_{i=1}^{k} c_i \leq c + 2(k-1)$. By Theorem 2, for all $1 \leq i \leq k$, $2c_i - 4$ guards suffice to guard M_i. Thus, we can guard M with at most $\sum_{i=1}^{k}(2c_i - 4)$ guards, and

$$\sum_{i=1}^{k}(2c_i - 4) \leq 2\left(\sum_{i=1}^{k} c_i\right) - 4k \leq 2\left(c + 2(k-1)\right) - 4k = 2c - 4$$

<div align="right">∎</div>

It would be of independent interest to characterise the maximum guarding number of a gallery in terms of the sequence of the size of blocks of consecutive right-turn vertices. For instance, if P has four right-turn vertices, the situation is as follows.

- The maximum guarding number is four if the block sequence is $1, 1, 1, 1$;
- the maximum guarding number is three if the block sequence is $2, 1, 1$;
- the maximum guarding number is two if the block sequence is $3, 1$ or $2, 2$ or 4.

Using pseudo-triangulations, we have partial answers to these questions, that due to the lack of space we postpone 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 3. *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   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 $T(N) = N \setminus N_i^- = \overline{N \setminus N_i}$.
(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 right-turn vertex of ∂N is important, set $T(N) = N$.

If none of (A),(B), or (C) occur, then ∂N contains a right-turn 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 remark that if 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 shown 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)
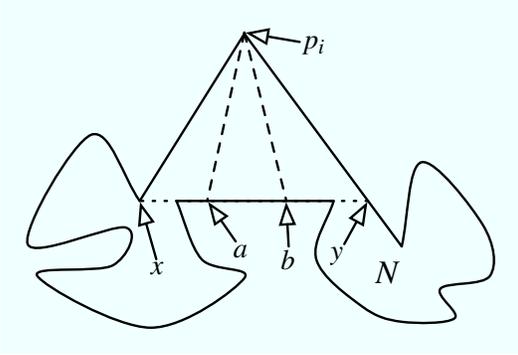

Fig. 6. No matter how guards are placed outside of $\triangle_{xp_iy}$, some part of $\triangle_{ap_ib}$ close to $[a, b]$ will not be seen by any guard.
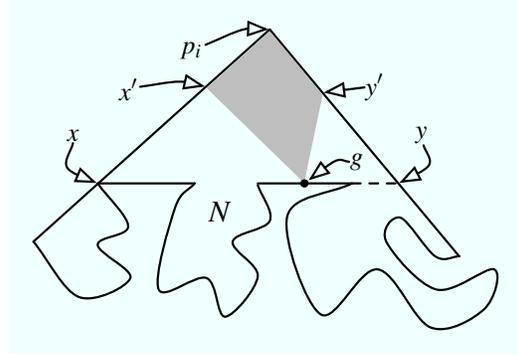
Fig. 7. The situation in case (D). The dark shaded region belongs to N but not to $T(N)$.

(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.)

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 3 then $T(N)$ and G also satisfy the hypotheses of Theorem 3. Another important property of T is that its repeated application is guaranteed to increase the value of a certain
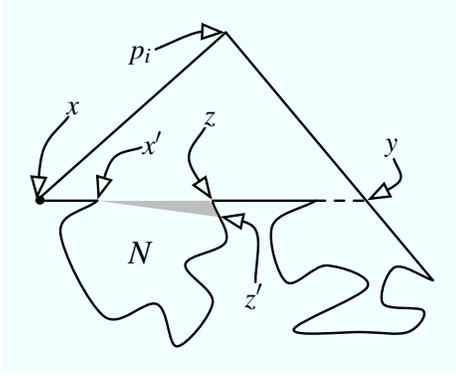
Fig. 8. The situation in case (E). The grey shaded region belongs to N but not to $T(N)$.
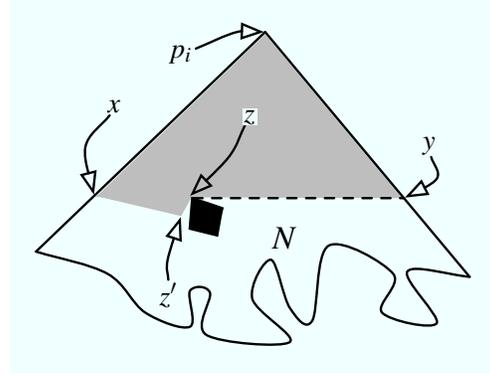
Fig. 9. The situation in case (H). The grey shaded region belongs to N but not to $T(N)$, and the black shaded region belongs to U.

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 right-turn vertices which are guards, plus the number of right-turn 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) right-turn vertex of N is an occupied (resp. critical) right-turn 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 claim.

Claim. For any nearly polygonal gallery M with no set of 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.

We postpone the proof of this claim to the end of the paper, first using it to prove Theorem 3.

3.2 Proof of Theorem 3

Let M and G be as in the statement of Theorem 3 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 3 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 right-turn vertices in *all* of the galleries M_i can be bounded uniformly in terms of g; this is the substance of the following claim.

Claim. For all $i \geq 0$, there are at most $g - 1$ critical right-turn vertices in M_i.

We will prove this claim along with Claim 3.1, after completing the proof of Theorem 3.

We observe that M contains precisely $1 + \sum_{\{c \in C : \kappa(c) > 2\}}(\kappa(c) - 1)$ leaves (this can easily be seen by 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

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

Furthermore, there are clearly at most g occupied right-turn vertices. It follows by Claim 3.2, the above inequalities and the definition of Φ that for all i, $\Phi(M_i)$ is at most $4g - 1$. By the observation immediately preceding Claim 3.1, $\Phi(M_{i+1}) \ge \Phi(M_i)$ for all i, and Claim 3.1 then implies that there exists an integer $j \ge 0$ 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 sake of the 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 $1 \le i \le r$ let $N_i^+ = \overline{N_i^-}$, let $G_i = N_i^+ \cap G$ and let $g_i = |G_i|$. Since $p \notin G$, $\sum_{i=1}^r g_i = g$. Furthermore, since G_i guards N_i^+ and $g_i < g$ for all $1 \le i \le r$, by induction there is $H_i \subset N_i^+$ such that $|H_i| \le 4g_i - 4$ and $G_i \cup H_i$ guards N_i^+. In this case $\bigcup_{i=1}^r (G_i \cup H_i)$ guards M, and

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

If (C) occurs then since the number of critical right-turn vertices is at most $g - 1$ by Claim 3.2, and the number of occupied vertices is at most g, the total number of right-turn vertices of our gallery M is at most $2g - 1$. It follows from Corollary 1 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 Claims 3.1 and 3.2

Having established Theorem 1 assuming that Claims 3.1 and 3.2 hold, we now turn to the proofs of these claims.

Proof (Proof of Claim 3.1). Let M and G be as in the statement of the claim, and suppose that $T(M) \ne M$. 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 $M_1 \ne M$, 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 $\text{int}(\triangle_{xp_iy}) \cap G = \emptyset$, 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 $]x, y[\cap \partial M \ne \emptyset$, we first observe that since $[x, y] \cap G = \emptyset$, no guard on the line through x and y can see $\text{int}(\triangle_{xp_iy})$. As every point in $\text{int}(\triangle_{xp_iy})$ 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 right-turn 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 right-turn vertex in M_1 but not in M, so $\Phi(M_1) > \Phi(M)$.

This completes the proof of Claim 3.1.

Proof (Proof of Claim 3.2). 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 right-turn 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 right-turn 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 claim immediately follows.

We first remark that if p_j is a critical right-turn vertex of P then the angle $p_{j-1} p_j p_{j+1}$ is strictly positive. Therefore, g_j is different from g_{j+1}, or else the quadrilateral $g_j x_j p_j y_j$ 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, for $i = 1, \ldots, k$, $g_{i+1} = g_i'$ and $g_i g_{i+1}$ is the edge corresponding to some critical right-turn vertex p_i). In this case, the polygonal line

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

which is not necessarily simple or even uncrossing, contains some simple, closed polygonal line $PL_1 = (x, g_i, p_i, \ldots, x)$ or $PL_2 = (x, p_i, g_{i+1}, \ldots, x)$. 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 right-turn vertex p_j, as the angle at p_j is right-turn, 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 $\mathbb{R}^2 \setminus PL_1$) lies entirely within M. Since p_i is right-turn, 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 $g_i p_i y$ 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. This completes the proof of Claim 3.2.

4   Conclusion

In this paper, we introduce a new measure of complexity for the polygonal shapes in the plane. We consider the problem of guarding the interior of an art gallery, assuming that the walls are already guarded. We present a linear bound on the size of extra guards, and present a precise conjecture on the optimal value. On the complexity side, based on the reduction operator that we define, and by using the results of [LL86,ESW01], it is not difficult to prove that the general problem of calculating the number of extra guards needed is NP-complet and does not admit PTAS, i.e. is APX-hard. It is however interesting to consider the approximability of the problem with good performance ratios. In particular we do not know if the problem admits constant factor approximation (we note that the best approximation algorithm for the general art gallery problems has ratio $\log(OPT)$ [Gho87,DKDS07], and it is not known whether or not there is a constant factor approximation algorithm). The generalisation of the problem to three dimension is also another natural problem to investigate.

References

[AT81]   D. Avis and GT Toussaint. An efficient algorithm for decomposing a polygon into star-shaped polygons. *Pattern Recognition*, 13(6):395–398, 1981.

[Chv75]   V. Chvatal. A combinatorial theorem in plane geometry. *J. Combin. Theory Ser. B*, 18:39–41, 1975.

[DB07]   V. Dujmović and P. Bose. 2007. manuscript.

[DKDS07]   Ajay Deshpande, Taejung Kim, Erik D. Demaine, and Sanjay E. Sarma. A Pseudopolynomial Time $O(\log c_{opt})$-Approximation Algorithm for Art Gallery Problems. In *WADS*, pages 163–174, 2007.

[EGHP$^+$00] A. Efrat, L.J. Guibas, S. Har-Peled, D.C. Lin, J.S.B. Mitchell, and TM Murali. Sweeping simple polygons with a chain of guards. *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, pages 927–936, 2000.

[EOW84]   H. Edelsbrunner, J. O'Rourke, and E. Welzl. Stationing guards in rectilinear art galleries. *Computer vision, graphics, and image processing*, 27(2):167–176, 1984.

[ESW01]   S. Eidenbenz, C. Stamm, and P. Widmayer. Inapproximability Results for Guarding Polygons and Terrains. *Algorithmica*, 31(1):79–113, 2001.

[Fis78]   S. Fisk. A short proof of Chvatals watchman theorem. *J. Combin. Theory Ser. B*, 24(3):374, 1978.

[FMZ07]   C. Fragoudakis, E. Markou, and S. Zachos. Maximizing the guarded boundary of an Art Gallery is APX-complete. *Computational Geometry: Theory and Applications*, 38(3):170–180, 2007.

[Gho87]   S. Ghosh. Approximation algorithms for art gallery problems. *Proc. Canadian Inform. Process. Soc. Congress*, pages 429–434, 1987.

[Hon76]   R. Honsberger. Mathematical Gems II. *Mathematical Association of America, Washington*, pages 104–110, 1976.

[KKK83]   J. Kahn, M. Klawe, and D. Kleitman. Traditional Galleries Require Fewer Watchmen. *SIAM Journal on Algebraic and Discrete Methods*, 4:194, 1983.

[KP94]   E. Kranakis and M. Pocchiola. Camera placement in integer lattices. *Discrete and Computational Geometry*, 12(1):91–104, 1994.

[Lau99]   A. Laurentini. Guarding the walls of an art gallery. *The Visual Computer*, 15(6):265–278, 1999.

[LHL93]   B.C. Liaw, NF Huang, and R.C.T. Lee. The minimum cooperative guards problem on k-spiral polygons. *Proc. of Fifth Canadian Conference on Computational Geometry*, pages 97–101, 1993.

[LL86]   D. Lee and A. Lin. Computational complexity of art gallery problems. *Information Theory, IEEE Transactions on*, 32(2):276–282, 1986.

[MP03]   T. Michael and V. Pinciu. Art gallery theorems for guarded guards. *Computational Geometry*, 26(3):247–258, 2003.

[O'R83]   J. O'Rourke. An alternate proof of the rectilinear art gallery theorem. *Journal of Geometry*, 21(1):118–130, 1983.

[O'R87]   J. O'Rourke. *Art gallery theorems and algorithms*. Oxford University Press, Inc. New York, NY, USA, 1987.

[PV96]   M. Pocchiola and G. Vegter. Pseudo-triangulations: theory and applications. *Proceedings of the twelfth annual symposium on Computational geometry*, pages 291–300, 1996.

[Sac82]   J.R. Sack. An O (n log n) algorithm for decomposing simple rectilinear polygons into convex quadrilaterals. *Proceedings 20th Conference on Communications, Control, and Computing*, pages 64–74, 1982.

[She92]   TC Shermer. Recent results in art galleries [geometry]. *Proceedings of the IEEE*, 80(9):1384–1399, 1992.

[ST05]   B. Speckmann and C.D. Tóth. Allocating Vertex p-Guards in Simple Polygons via Pseudo-Triangulations. *Discrete and Computational Geometry*, 33(2):345–364, 2005.

[Sza97]   L. Szabo. Recent results on illumination problems. *Bolyai Society Mathematical Studies*, 6:207–221, 1997.

[Tót77]   L.F. Tóth. Illumination of convex discs. *Acta Mathematica Hungarica*, 29(3):355–360, 1977.

[Urr00]   J. Urrutia. Art gallery and illumination problems. *Handbook of Computational Geometry*, pages 973–1027, 2000.