

Opposing Half Guards*

Erik Krohn[†]Bengt J. Nilsson[‡]Christiane Schmidt[§]

Abstract

We study the art gallery problem for opposing half guards: guards that can either see to their left or to their right only. We present art gallery theorems, show that the problem is NP-hard in monotone polygons, present approximation algorithms for spiral and staircase polygons, and show that the location of half guards in 2-guardable polygons is not restricted to extensions.

1 Introduction

The Art Gallery Problem (AGP), based on a question by Victor Klee, is one of the classical problems in Computational Geometry. Klee asked for the minimum number of stationary guards with 360° vision that we need to place to achieve complete visibility coverage of a polygon \mathbf{P} . Such a guard $g \in \mathbf{P}$ can see a point $p \in \mathbf{P}$ iff \overline{gp} is fully contained in \mathbf{P} . Typical results can be classified in two categories:

1. Computational complexity and algorithmic results for the minimization of the number of star-shaped polygons (the *visibility polygons* (VPs) of guards) that cover a polygon—computation of $G(\mathbf{P})$.
2. “Art Gallery Theorems”: Worst-case, combinatorial bounds on the number of VPs that are sometimes necessary and always sufficient to cover a class of polygons—bounds on the maximum value of $G(\mathbf{P})$ over all polygons of n vertices, $g(n)$.

Results on (2) are presented in, e.g., [4, 7, 11, 16]. Results settling the computational complexity, (1), are given in, e.g., [17, 13, 16, 6, 1], approximation algorithms are given in, e.g., [12, 3].

Here, the guards do not have 360° vision, but every guard can either see to the left or the right: imagine a spotlight as shown in Figure 1(a), for which the upper bow is fixedly mounted (and cannot rotate/yaw), so, the only degree of freedom for the spotlight is analogue to




Figure 1: (a) A spotlight (picture from flickr.com user Toby Simkin under licence CC BY-NC-SA 2.0.), (b) a polygon (black) with a left-looking half guard and its visibility polygon (blue) and a right-looking half guard and its visibility polygon (red).

an aircraft pitch. If we fully utilize this movement, the spotlight illuminates one of the two halfplanes defined by the line that contains the upper bow. Formally, a left-looking (right-looking) half guard $g \in \mathbf{P}$ sees a point $p \in \mathbf{P}$ iff \overline{gp} is fully contained in \mathbf{P} and if p does not have larger (smaller) x -coordinate than g , see Figure 1(b) for an illustration. We call such guards *opposing* in contrast to half guards that can all only see in one direction. Half guards that all see in only one direction have been considered by Gibson et al. [9].

Of course, if we found a feasible solution for the AGP with “ordinary” guards, placing a left- and a right-looking half guard for each guard yields a feasible solution also for the AGP with half guards—usually this will not be optimal, for monotone mountains (where we can compute an optimal solution for the AGP in polynomial time [5]) it would directly yield a 2-approximation. In this paper, we study questions of both types (1) and (2) for half guards.

2 Notation and Preliminaries

We let \mathbf{P} denote a simple polygon, $|\mathbf{P}| = n$. We let r denote the number of reflex vertices of P . A simple polygon \mathbf{P} is *x-monotone* if the intersection $\ell \cap \mathbf{P}$ of \mathbf{P} with any vertical line ℓ is a connected set. Any *x-monotone* polygon decomposes into two *x-monotone* polygonal chains between the rightmost and leftmost point of \mathbf{P} . An *x-monotone* polygon is a *monotone mountain* or *uni-monotone*, if one of its two chains—w.l.o.g. the upper chain—is a single horizontal segment, \mathcal{H} . A polygon \mathbf{P} is *orthogonal* (or *rectilinear*) if all of its edges are axis-parallel, that is, either horizontal or vertical. An orthogonal polygon \mathbf{P} is a *staircase polygon* if it is *x-* and *y-monotone*. We assume a leftmost (rightmost) point of a staircase polygon to also be its

*B. J. N. and C. S. were supported by grants 2018-04001 (Nya paradigmer för autonom obemannad flygledning) and 2021-03810 (Illuminate: bevisbart goda algoritmer för bevakningsproblem) from the Swedish Research Council (Vetenskapsrådet). E. K. was supported by a sabbatical grant from the University of Wisconsin - Oshkosh.

[†]Department of Computer Science, University of Wisconsin - Oshkosh, Oshkosh, WI, USA, krohne@uwosh.edu

[‡]Department of Computer Science and Media Technology, Malmö University, Sweden, bengt.nilsson.TS@mau.se

[§]Department of Science and Technology, Linköping University, Sweden, christiane.schmidt@liu.se

lowest (highest) point and denote this point by p_ℓ (p_u). We have two polygonal chains connecting p_ℓ and p_u , the chain for which there exists a point on the other chain with the same x -coordinate but larger y -coordinate is denoted as the lower chain, the other chain is the upper chain.

Our half guards can either look to their left or their right, formally, we define left-looking and right-looking half guards. Let p_x and p_y define the x - and y -coordinate of a point p , respectively.

- A *left-looking half guard* $g \in \mathbf{P}$ can see a point $q \in \mathbf{P}$ iff (\overline{gq}) does not intersect \mathbf{P} 's boundary AND $g_x \geq p_x$; we say that g *half sees* q .
- A *right-looking half guard* $g \in \mathbf{P}$ can see a point $q \in \mathbf{P}$ iff (\overline{gq}) does not intersect \mathbf{P} 's boundary AND $g_x \leq p_x$; we say that g *half sees* q .

Because we sometimes compare with “normal” visibility, we also define when a (full) guard $g \in \mathbf{P}$ *sees* a point $p \in \mathbf{P}$: when the line segment \overline{gp} does not intersect \mathbf{P} 's boundary. For a point p , we let $\mathbf{V}(p)$ denote the half-visibility polygon of p and $\mathcal{V}(p)$ denote the “normal” visibility polygon of p . In a polygon \mathbf{P} , a set of witnesses W is a set of points in \mathbf{P} , such that $\forall w_1, w_2 \in W : \mathcal{V}(w_1) \cap \mathcal{V}(w_2) = \emptyset$.

3 Art Gallery Theorems for Opposing Half Guards

In this section, we give Art Gallery Theorems, that is, statements of the type “ $x(n)$ guards are always sufficient and sometimes necessary for polygons with n vertices”, for different polygon classes.

Theorem 1 *In simple polygons with n vertices:*

- For $r > \frac{n}{2}$: $2\lfloor \frac{n}{3} \rfloor$ half guards are always sufficient and sometimes necessary.
- For $r \leq \frac{n}{2}$: $r + 1$ half guards are always sufficient and sometimes necessary.

Proof. For $r > \frac{n}{2}$, the upper bound follows trivially from the $\lfloor \frac{n}{3} \rfloor$ upper bound by Fisk [7] for “normal” guards (triangulating the polygon, three-coloring the vertices and using the least-frequently used color yields at most $\lfloor \frac{n}{3} \rfloor$): placing one right- and one left-looking half guard at each position of a “normal” guard results in $2\lfloor \frac{n}{3} \rfloor$ half guards.

For the lower bound, we construct a family of polygons \mathbf{P}_n , see Figure 2(a), that needs $2\lfloor \frac{n}{3} \rfloor$ half guards. \mathbf{P}_n is an x -monotone polygon: the upper polygonal chain has reflex vertices only (except for the rightmost and leftmost vertex, which are convex), the lower chain has alternating reflex and convex vertices. The reflex vertices of the upper chain have the same x -coordinate as the convex vertices of the lower chain. The lower-chain vertices incident to the rightmost and leftmost vertex of \mathbf{P}_n are either a reflex or convex vertex, such that we can define \mathbf{P}_n also for $(n \bmod 3) \neq 0$. For each




Figure 2: (a) Lower bound construction for simple (and monotone) polygons with $r > \frac{n}{2}$; (b) zoomed in on (a) showing \mathbf{P}_i (yellow) and \mathbf{P}_{i+1} (green). (c) Lower bound construction for simple (and monotone) polygons with $r \leq \frac{n}{2}$.

convex vertex c_i on the lower chain, we define a sub-polygon $\mathbf{P}_i \subset \mathbf{P}_n$. We extend the two edges incident to c_i ; let the two points where these extensions intersect with the upper chain be v_i and w_i . Let the reflex vertex of the upper chain with the same x -coordinate as c_i be u_i . The polygon \mathbf{P}_i is defined by c_i, v_i, u_i and w_i , see Figure 2(b). Note that $\mathbf{P}_i \cap \mathbf{P}_{i+1} \neq \emptyset$. We claim that we need two half guards per \mathbf{P}_i . Let p_i and q_i be two points on the edges incident to c_i within distance ε from the reflex vertices (marked in red in Figure 2(b)). The point p_i can be seen from a right-looking half guard g only if $g_x \leq p_{i,x}$, however—as indicated by the red line segment—such a half guard cannot see q_i . A similar argument holds for a left-looking half guard seeing q_i . Both p_i and q_i can be seen from points in \mathbf{P}_i only. However, we saw that \mathbf{P}_i and \mathbf{P}_{i+1} overlap: assume that we place a left-looking half guard g at w_i , it can see q_i and p_{i+1} . We still need two (more) half guards in \mathbf{P}_{i+1} : a right-looking half guard at w_i cannot see q_{i+1} , but a left-looking half guard that sees w_i cannot see all of \mathbf{P}_{i+1} . Hence, we need $2\lceil \frac{n}{3} \rceil$ half guards.

For the upper bound for $r \leq \frac{n}{2}$, we recursively partition the polygon into $r + 1$ convex pieces: we pick any reflex vertex and extend one of its incident edges until we hit the boundary. Then, in both subpolygons we created, this vertex is no longer a reflex vertex. Because we end up with convex pieces, we can cover each piece with either a left-looking half guard at its rightmost vertex or a right-looking guard at its leftmost vertex. This yields in total $r + 1$ half guards. For the lower bound, we construct a family of polygons \mathbf{P}_n , see Figure 2(c), that needs $r + 1$ half guards. No two of the vertices marked with a point can be seen by a single half guard. Hence, we need $r + 1$ half guards. \square

The lower bound constructions for simple polygons are in fact both also monotone polygons, we yield:

Corollary 2 *In monotone polygons with n vertices:*

- For $r > \frac{n}{2}$: $2\lfloor \frac{n}{3} \rfloor$ half guards are always sufficient and sometimes necessary.
- For $r \leq \frac{n}{2}$: $r + 1$ half guards are always sufficient and sometimes necessary.




Figure 3: (a) Guarding L-shaped pieces with a single half guard; right-looking half guards are shown in red, left-looking half guards are shown in blue. (b) Lower bound construction for orthogonal polygons.

Theorem 3 *In simple orthogonal polygons with n vertices, $\lfloor \frac{n}{4} \rfloor$ half guards are always sufficient and sometimes necessary.*

Proof. An orthogonal polygon can be partitioned in $\lfloor \frac{n}{4} \rfloor$ L-shaped pieces (in linear time) [15, 16]. L-shaped pieces are orthogonal hexagons. Any L-shaped piece can be guarded by a single half guard placed at the only convex vertex, v , that can see all other vertices of the L-shaped piece (when the piece is considered as a simple polygon). Depending on whether the interior of the L-shaped piece lies in the left or right half plane of the vertical line through v , we use a left-looking or right-looking half guard, respectively. See Figure 3(a) for examples of guarding L-shaped pieces of the four possible orientations. This establishes that $\lfloor \frac{n}{4} \rfloor$ half guards are always sufficient.

Figure 3(b) shows a family of polygons \mathbf{P}_n that needs $\lfloor \frac{n}{4} \rfloor$ half guards: no two of the vertices marked with a point can be seen by a single half guard. Hence, we need one half guard per four edges. \square

Theorem 4 *In monotone mountains with n vertices:*

- For $r < \frac{n}{2}$: $r + 1$ half guards are always sufficient and sometimes necessary.
- For $\frac{n}{2} \leq r \leq \frac{3n}{4}$: $\lfloor \frac{n}{2} \rfloor$ half guards are always sufficient and sometimes necessary.
- For $r > \frac{3n}{4}$: $2 \cdot (n - r - 2) \leq \frac{n}{2}$ half guards are sometimes necessary.

Proof. For $r \leq \frac{n}{2}$, the upper and lower bound follow as in the proof of Theorem 1 (i.e., the lower bound is shown in Figure 2(c) (where the lower horizontal chain is a horizontal segment)).

For the upper bound for $\frac{n}{2} \leq r \leq \frac{3n}{4}$, we consider the lower polygonal chain (the upper polygonal chain is a single horizontal segment \mathcal{H}). For guarding monotone mountains with “normal” guards, all guards can be placed on \mathcal{H} and it is sufficient to guard all points of the lower polygonal chain to guard the complete polygon [5], the arguments used there also hold for half guards.

In between any pair of consecutive convex vertices, we have a reflex chain. One or two of the vertices in a reflex chain have a larger y -coordinate than the other vertices of that reflex chain. We split the reflex chain at one of these two vertices. Now, any convex vertex v is




Figure 4: Upper bound construction for monotone mountains. The vertex v_1 and $C_R(v_1)$ and $C_L(v_1)$ are shown in dark red; v_2 , $C_R(v_2)$, and $C_L(v_2)$ are shown in dark green; and v_3 , $C_R(v_3)$, and $C_L(v_3)$ are shown in dark blue. Right-looking half guards are shown in red, left-looking half guards are shown in blue.




Figure 5: Lower bound construction for monotone mountains and $\frac{n}{2} \leq r \leq \frac{3n}{4}$.

adjacent to a split reflex chain both on its right, $C_R(v)$, and on its left, $C_L(v)$. Let ℓ be the vertex with maximal y -coordinate in $C_L(v)$. We distinguish three cases:

1. Both $C_R(v)$ and $C_L(v)$ have more than two edges: Let $h \in \mathcal{H}$ be the point with $h_x = v_x$, we place a right-looking and a left-looking guard at h . These half guards monitor $C_R(v)$ and $C_L(v)$, and we use two half guards for at least four edges. The vertex v_1 in Figure 4 is an example for this case.
2. One of the two split reflex chains, w.l.o.g. $C_R(v)$, has more than two edges: let h be the point on \mathcal{H} with $h_x = \ell_x$. We place a right-looking half guard at h . This half guard sees $C_R(v)$ and $C_L(v)$ ($C_L(v)$ has only one edge), and we use one half guard for at least three edges. The vertex v_2 in Figure 4 is an example for this case.
3. Both $C_R(v)$ and $C_L(v)$ have only one edge: let h be the point on \mathcal{H} with $h_x = \ell_x$. We place a right-looking half guard at h . It monitors $C_R(v)$ and $C_L(v)$, and we use one half guard for two edges. The vertex v_3 in Figure 4 is an example for this case.

In all cases, each half guard monitors on average at least two edges, and the claim follows.

For the lower bound for $\frac{n}{2} \leq r \leq \frac{3n}{4}$, we use a similar construction as in Figure 2(c), however, in between two consecutive convex vertices, we now include one, two or three reflex vertices, see Figure 5 for an example. If $C_R(v)$ and $C_L(v)$ have both one edge, we can guard both of these edges with a single (left- or right-looking) half guard. If $C_R(v)$ and $C_L(v)$ have both two edges, these four edges cannot be seen by a single half guard, and no half guard from neighboring reflex chains can fully monitor these edges. Hence, we use one half guard per two edges.

For $r > \frac{3n}{4}$, we only use case 1 (that is, only what is depicted in dark red in Figure 4): we place

a right-looking and a left-looking guard on \mathcal{H} at the x -coordinates of convex vertices—except for the leftmost and rightmost convex vertex that are located on \mathcal{H} . These guards see the complete lower polygonal chain. We have $c = n - r$ convex vertices, thus, we place $2(c-2) = 2(n-r-2)$ half guards. For the lower bound, we insert reflex chains as the dark-red reflex chains from Figure 4 in the construction from Figure 5. \square

4 Hardness Results for Opposing Half Guards

NP-hardness for point guarding a monotone polygon with half guards that only see to the right was claimed in [10]. The same reduction can be used for opposing half guards. In [10], the authors show an NP-hardness reduction from 3SAT. They show that certain vertices on the boundary represent truth values for variables in the original 3SAT instance. Clauses are represented by specific points on the boundary of the polygon. For example, if a clause $c = x_2 \vee x_5 \vee \bar{x}_7$ were in the original instance, then a vertex would exist on the boundary that would be seen by three vertices, namely the ones representing x_2, x_5 and \bar{x}_7 . We briefly look at each pattern and show that if the polygon is guardable with k guards, then all k guards must be right-looking half guards.




Figure 6: A starting pattern and a variable pattern for the NP-hardness reduction.

Starting Pattern: No left-looking half guard can see both v_3 and v_6 , see Figure 6. No left-looking half guard placed outside of the starting pattern can see v_3 or v_6 . Placing a left-looking half guard in the starting pattern for x_i would require at least two guards to be placed for the x_i starting pattern when one right-looking half guard is sufficient.

Variable Pattern: Distinguished vertices of a variable pattern that can be seen by guards outside of the

variable pattern are vertices v_{10} and v_{13} , see Figure 6. Neither of these vertices, and none of the other distinguished vertices in this pattern, can be seen by a left-looking half guard outside of the variable pattern. Any left-looking half guard placed inside the variable pattern will require too many guards to be placed to guard the entire variable pattern. In the original reduction, two right-looking guards are required to guard the variable pattern. Even with left-looking guards being allowed, two guards are still required.

We note that only one of v_{10} or v_{13} will be seen by a guard to the left of the variable pattern. An incorrectly placed guard that sees both v_{10} and v_{13} will be an additional guard and will not reduce the number of guards needed in the current variable pattern. No guard (left or right-looking) can see both v_2 and v_5 . Therefore, at least two guards are required to see v_2 and v_5 in each variable pattern. First, assume v_{10} is seen by a right-looking half guard to the left of the variable pattern. This leaves the following vertices to be guarded: $\{v_2, v_5, v_8, v_{12}, v_{13}\}$. No left-looking half guard can see more than one distinguished vertex in that list. If a left-looking half guard sees v_2 , then the only location to see both v_5 and v_{12} is a right-looking half guard at location v_{11} . If a left-looking guard is placed to see v_5 , then the only location that sees v_2 and v_{12} is a right-looking guard at location v_1 . In both cases, neither v_1 nor v_{11} sees v_8 and a third guard would be required.

Next, assume v_{13} is seen by a guard outside of the variable pattern. In this instance, no left-looking guard can see more than one of $\{v_2, v_5, v_8, v_{10}, v_{12}\}$. The same argument as above also applies here. Therefore, if a left-looking half guard is placed in this variable pattern, three guards are required when two right-looking half guards are sufficient. Those guard locations are $\{x_i, v_{11}\}$ or $\{\bar{x}_i, v_1\}$.

No extra guards are required to see any of the clause distinguished points. Any left-looking half guard that is placed to see a clause distinguished point will only see that particular clause distinguished point. No other distinguished points in any starting or variable pattern will be seen by such a guard. If k guards are sufficient to guard the entire polygon, then if a single left-looking half guard is placed in any pattern, an additional k guards are required to see the entire polygon. Thus, the reduction from [10] holds even if left-looking half guards are allowed to be placed inside the monotone polygon.

5 2-Guardable Polygons

In his Master thesis, Belleville [2] showed that if a polygon \mathbf{P} is two-guardable (two guards can fully monitor \mathbf{P}), \mathbf{P} is two-guardable by two guards that are located on edge extensions (including the edges themselves). We show that this statement does not hold for half guards:




Figure 7: Polygon construction for Theorem 5.

Theorem 5 *Let \mathbf{P} be a polygon for which the minimum half-guard cover has cardinality two, let these two guards be denoted as g_1 and g_2 . Then neither g_1 nor g_2 must be located on edge extensions. Moreover, neither g_1 nor g_2 must be located on polygon diagonals.*

Proof. We construct a polygon \mathbf{P} that can be covered by two half guards, but these guards do not lie on any edge extensions (or edges), that is, if we would restrict half-guard locations to edges and edge extensions, \mathbf{P} cannot be covered by two half guards. The polygon \mathbf{P} is shown in Figure 7. The shown half guards (one left-looking and one right-looking half guard) monitor \mathbf{P} , let the right-looking half guard be denoted as g_r and the left-looking half guard be denoted as g_l . The lines of sight meeting in the niches on the top and bottom of \mathbf{P} are shown in red and blue for g_r and g_l , respectively. The dots in the polygon edges indicate very long edges.

Clearly, as \mathbf{P} is an orthogonal polygon, neither g_r nor g_l lie on any edge extensions, nor on any diagonals. It remains to show that no other pair of half guards monitors \mathbf{P} .

First, note that we cannot use two half guards that look into the same direction to monitor \mathbf{P} with two half guards only: Then half guards cannot “share” seeing the niches, and we need one half guard per pair of mirrored niches. Hence, any minimum half-guard cover of \mathbf{P} must contain a right-looking and a left-looking half guard.

Now assume that we try to move g_r and g_l . Assume first that we only alter the y -coordinate of g_r . W.l.o.g.—the polygon is symmetric—we increase the y -coordinate of g_r and obtain guard g'_r . Then, the first point that g'_r sees on the top-right edge of \mathbf{P} (the horizontal edge ending in the upper right vertex of \mathbf{P}) has a larger x -coordinate than p_1 . Hence, we need to increase the x -coordinate of g_l . If we only alter the x -coordinate of g_l to obtain g'_l , points at distance ε from p_3 and p_4 on the same edges are not visible to g'_l , and we do no longer have a half-guard cover. Increasing also the y -coordinate of g_l leaves part on the top unseen. Thus, assume that we increase the x -coordinate and decrease the y -coordinate of g_l . Then, a point at distance ε from p_3 on its edge is not visible.

Similar arguments yield that changing the x -coordinate of g_r does not allow us to find a position for g_l such that \mathbf{P} is covered. \square




Figure 8: (a) Windows of a spiral polygon. (b) Illustrating the four operations.

Note that for “normal” guards, we may move the two guards to the x -coordinates of the left-most and right-most vertical edges of the niches.

6 An Approximation Algorithm for Spiral Polygons

A simple polygon \mathbf{P} is *spiral* if it has two convex vertices u and u' such that a clockwise boundary walk from u to u' encounters only convex vertices and a counterclockwise boundary walk from u to u' encounters only reflex vertices. Nilsson and Wood [14] show a linear time greedy algorithm to compute the minimum number of “normal” guards for spiral polygons. We show a $3/2$ -approximation for half guards based on dynamic programming.

For a half visibility polygon $\mathbf{V}(p)$ of a point p in polygon \mathbf{P} , we call a *window* a boundary edge of $\mathbf{V}(p)$ that does not coincide with the boundary edges of \mathbf{P} ; see Figure 8(a). We make the following claims without proof.

Lemma 6 *For a point p in a spiral polygon \mathbf{P} , the half visibility polygon $\mathbf{V}(p)$ has at most three windows.*

Lemma 7 *In a spiral polygon \mathbf{P} , there is an optimal set of half guards that all lie on the convex chain of \mathbf{P} .*

We let $n = n_c + n_r$ be the total number of vertices where n_c is the number of convex vertices including u and u' and n_r is the number of reflex vertices. We order the vertices from u to u' so that $u = u_1, \dots, u_{n_c} = u'$ are the convex vertices and v_1, \dots, v_{n_r} are the reflex vertices in counterclockwise order starting from the vertex after u and ending at the vertex before u' . To simplify, we let $v_0 = u$. We also denote by $e_i = [v_{i-1}, v_i]$ the edge of the reflex boundary connecting v_{i-1} and v_i , $1 \leq i \leq n_r$.

We identify special vertices that we call *corners*. A convex vertex u_j in \mathbf{P} is an (*outer*) *corner*, if the two incident edges to u_j both lie on the same side of a vertical line through u_j and assume for simplicity that \mathbf{P} has no vertical edges. We say that a corner is a *left corner* if the incident edges lie to the right of the vertical line through the corner, otherwise it is a *right corner*. Similarly, we can define the *inner corners* as the vertices of the reflex boundary chain for which the adjacent edges lie on the same side of the vertical line through the vertex. As we follow the convex chain in clockwise order from a left corner to a right corner, we say that a left looking half guard is a *backward guard* and a right looking half guard

is a *forward guard*. Similarly following the convex chain in clockwise order from a right corner to a left corner, a left looking half guard is a *forward guard* and a right looking half guard is a *backward guard*.

In addition, we define some useful operations as follows. For an edge e of the reflex chain, the point $\rightarrow(e)$ is the point on the convex chain intersected by a ray exuded in the counterclockwise direction along e .

For a point p on the convex chain, the point $\downarrow(p)$ is the point on the reflex chain intersected by a ray exuded in the vertical direction from p towards the interior of the polygon, if it exists, otherwise $\downarrow(p)$ is undefined.

For a point p on the convex chain, the point $\searrow(p)$ is the point on the convex chain intersected by a ray exuded towards the last vertex on the reflex chain seen by a forward half guard at p .

For a point p on the reflex chain, the point $\uparrow(p)$ is the point on the convex chain intersected by a ray exuded in the vertical direction from p towards the interior of the polygon. If p is an inner corner, the ray is exuded in the vertical direction making $\uparrow(p)$ the furthest of the two possible points along the convex chain; see Figure 8(b).

Consider a set of half guards \mathcal{G} on the convex chain and a half guard g placed at point p of the convex chain. Assume that the half guards in \mathcal{G} completely see the edges e_1, \dots, e_{i-1} and g completely sees the edges $e_{i'}, \dots, e_{i-1}$, for indices $i' \leq i$, then operation $\text{ix}_{\mathcal{G}}(p)$ evaluates to the index i . Our algorithm will always place guards so that they see a contiguous portion of the reflex chain starting at $v_0 = u$ and ending at some point p on e_i , $i \geq 1$. We can therefore assume that \mathcal{G} is any set of half guards that see the edges before $e_{i'}$ and define $\text{ix}(p) = i$ only based on the half guard g at p .

We specify our algorithm as a dynamic programming algorithm based on the following recursion. Let i be the index of the furthest reflex edge e_i not completely seen by the currently placed half guards, let p be the last point on e_i seen by the currently placed half guards, and let q be the last point on the convex chain seen by the currently placed half guards. The first call is $G(0, u, u)$.

$$G(i, p, q) = \min \begin{cases} G(\text{ix}(u_j), v_{\text{ix}(u_j)-1}, \searrow(u_j)) + 1, & \text{half guard at next} \\ & \text{corner } u_j, \text{ if } u_j \text{ sees } p \text{ and } q \\ G(\text{ix}(g), \downarrow(g), g) + 1, & \text{backward guard at } g = \rightarrow(e_i), \\ & \text{if } g \text{ sees } q, \downarrow(g) \text{ is defined, and } g \text{ is before next corner} \\ G(\text{ix}(u_j), v_{\text{ix}(u_j)-1}, \searrow(u_j)) + 2, & \text{backward guard at} \\ & g = \rightarrow(e_i) \text{ and half guard at next corner } u_j, \\ & \text{if } g \text{ sees } q, \downarrow(g) \text{ is undefined, and } g \text{ is before } u_j \\ G(\text{ix}(q), v_{\text{ix}(q)-1}, \searrow(q)) + 1, & \text{backward guard at } g = \searrow(q) \\ & \text{if } v_i \text{ does not see } q, \downarrow(g) \text{ is defined,} \\ & \text{and } g \text{ is before next corner} \\ G(\text{ix}(u_j), v_{\text{ix}(u_j)-1}, \searrow(u_j)) + 2, & \text{backward guard at} \\ & g = \searrow(q) \text{ and half guard at next corner } u_j, \text{ if } v_i \text{ does not} \\ & \text{see } q, \downarrow(g) \text{ is undefined, and } g \text{ is before } u_j \\ G(\text{ix}(q), v_{\text{ix}(q)-1}, \searrow(q)) + 1, & \text{forward guard at } q, \\ & \text{if } q \text{ lies before } \uparrow(p) \text{ on convex chain} \\ G(\text{ix}(g), v_{\text{ix}(g)-1}, \searrow(g)) + 1, & \text{forward guard at } g = \uparrow(p), \\ & \text{if } \uparrow(p) \text{ lies before } q \text{ on convex chain} \end{cases}$$




Figure 9: Illustrating the seven cases in the recursion.

We ignore the bottom of the recursion, when $i > n_r + 1$ as it follows the general description above without the recursive calls.

Each possible half guard position q can be pre-computed as the (outer) corners, the intersection points $\rightarrow(e)$ on the convex chain issuing from each edge e of the reflex chain, and the intersection points $\uparrow(v)$ on the convex chain issuing vertically from each vertex v on the reflex chain, giving a linear number of possible positions. For each of these, we identify the linear sized sequence of continued intersection points $\searrow(q)$ of the supporting segments with the convex boundary, giving at most a quadratic number of positions. Each position can be acquired in constant amortized time by a traversal of the two boundary chains taking quadratic time in total. The dynamic programming thus fills out a table of size $O(n) \times O(n^2)$, each position in constant time.

We next prove the correctness and approximation ratio of the algorithm.

Lemma 8 *If \mathcal{G}^* is a minimal set of half guards for a spiral polygon, then the algorithm covers the polygon and places at most $3|\mathcal{G}^*|/2 + 1$ half guards.*

Proof. The correctness of the algorithm follows by construction, since it ensures that each window created by a half guard is seen by the next half guard placed.

It remains to show the approximation ratio for the algorithm. To do so, let \mathcal{B} be the set of prespecified and computed points on the convex boundary chain of \mathbf{P} where the algorithm can place a half guard. We have $|\mathcal{B}| \in O(n^2)$ from the discussion above. Consider an optimum set of half guards \mathcal{G}^* that we can assume by Lemma 7 all lie on the convex chain. Follow the convex chain from u to u' until a half guard g_1 in \mathcal{G}^* is reached that does not lie on a point in \mathcal{B} . Let p be the last point on the reflex chain seen by the guards in \mathcal{G}^* that lie before g_1 in \mathcal{G}^* . If g_1 is the first half guard, then $p = u$. We will show that we can exchange g_1 and the subsequent half guard g_2 in \mathcal{G}^* for three half guards g , g' , and g'' that indeed lie on the corresponding points in \mathcal{B} . Repeating the argument as the process follows the convex chain in clockwise order, proves our claim.

If g_1 can be moved to the subsequent point in \mathcal{B} , without losing visibility, we do so. This will not increase the

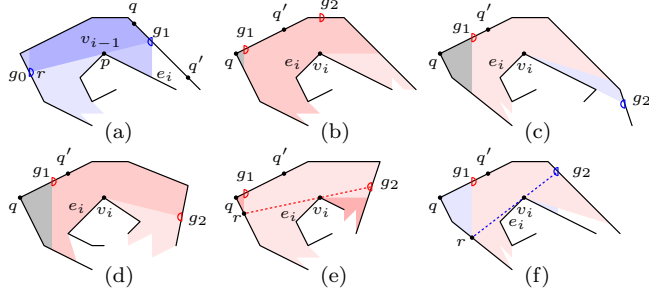


Figure 10: Illustrating the proof of Lemma 8. Grey regions are seen by previous half guards in \mathcal{G}^* .

number of guards.

Assume that g_1 is a backward guard between two subsequent points q and q' in \mathcal{B} and that g_1 cannot be moved to q' without losing visibility. Let v_{i-1} be the first vertex on the reflex chain that g_1 sees and thus g_1 sees e_i . So, p is some point on e_i (except v_i). Since g_1 cannot be moved to q' and the reflex chain is seen by the previous guards until p on e_i , there must be points on the convex chain that are not seen if g_1 is moved towards q' . Let r be such a point on the convex chain and let g_0 be the guard in \mathcal{G}^* that sees r . The guard g_0 must be a backward guard at r and r is by assumption in \mathcal{B} , hence, g_1 lies on the intersection of the convex chain and the supporting line from r through the reflex chain and this point lies in \mathcal{B} , a contradiction; see Figure 10(a).

Assume next that g_1 is a forward guard between two subsequent points q and q' in \mathcal{B} , that g_1 cannot be moved to q' without losing visibility, and that the region behind g_1 is seen by the previous guards in \mathcal{G}^* . Let v_i be the last vertex on the reflex chain that g_1 sees and assume furthermore that g_2 is a forward guard (or a backward guard with a corner between g_1 and g_2 along the convex chain). Without loss of generality, we can assume that g_2 lies on $\searrow(g_1)$ or on $\uparrow(v_i)$ depending on whether g_1 sees the next corner or not; see Figures 10(b), (c), and (d). If g_2 does not lie on any of these points, we can move it there without losing visibility. We can place two forward guards g and g' at q and q' . The half guard g' sees a vertex $v_{i'}$ with $i \leq i'$, hence we can replace g_2 by a half guard at $\searrow(q')$, if g' sees the next corner, or at $\uparrow(v_{i'})$. Note that if g_1 sees the next corner, so must g' . In both cases, g_1 and g_2 are replaced with three guards at positions in \mathcal{B} .

Finally, assume that g_1 is a forward guard, that the subsequent guard g_2 in \mathcal{G}^* is backward (or forward with a corner between them) and the region behind g_1 is not seen by previous guards in \mathcal{G}^* . We can argue for g_2 as we did in the first case to obtain a point r on the convex chain not seen by g_2 if g_2 is moved forward; see Figure 10(e) and (f). The forward guard g_1 lies on the intersection between vertical line segment intersection through r and the convex chain (the vertical segment spans between r and g_1). Thus, the convex region de-




Figure 11: (a) A staircase polygon with CWS points w_1 and w_2 , visibility polygons are shown in gray. (b) Placement of guards in G_{cw} . (c) Placement of guard in G_s (darkblue) and moving the guard to a feasible vertex (dash-dotted).

finied by the convex chain from r to g_1 and limited by the vertical segment through these points is seen by g_2 and it contains a corner. We can replace g_1 and g_2 by a half guard g at the corner, a forward guard g' at $\uparrow(v_i)$, where v_i is the last reflex vertex seen by g , and either a backward guard at $\rightarrow(e_{i'})$, where $e_{i'}$ is the first unseen edge of the reflex boundary, or a half guard at the subsequent corner, whichever point comes first along a traversal of the convex chain; see Figures 10(e) and (f).

In each case, a pair of half guards from the optimum is replaced by a triple of half guards at a subset of our designated positions. Hence, there exists a feasible placement of half guards having size at most $\lceil 3|\mathcal{G}^*|/2 \rceil$. Since the dynamic programming algorithm computes the minimum such placement, the lemma follows. \square

We have the following result.

Theorem 9 *The algorithm described computes a 3/2-approximate set of opposing half guards for spiral polygons in $O(n^3)$ time.*

7 2-Approximation for Staircase Polygons

Gibson et al. [8] show that staircase polygons allow for a 2-approximation for “normal” guards. Our algorithm is inspired by their CCCG algorithm (where CCCG stands for canonical convex corner guard).

Let G^* be an optimal opposing-half-guard set for a staircase polygon. We place a guard set G that is composed of two sets, that is, we have $G = G_{cw} \cup G_s$.

To construct G_{cw} , we place a set of witnesses on convex vertices, the *convex witness set* (CWS): We place witnesses alternately on (some) convex vertices of the lower and the upper chain. We place the first witness, w_1 , on the first convex vertex of the upper chain that does not lie on P 's lowest edge. We then define $\mathbf{P}_i = \mathbf{P} \setminus \mathcal{V}(w_{i-1})$, where $\mathcal{V}(p)$ denotes the visibility polygon of p under “normal” visibility. In \mathbf{P}_i we place a witness on the first convex on the lower chain vertex that does not lie on \mathbf{P}_i 's lowest edge, see Figure 11(a) for an exemplary witness placement. We yield $W = \{w_1, w_2, \dots\}$.

Lemma 10 *W is a set of witness points (and a CWS) and, hence, $|W| \leq |G^*|$.*

Proof. The topmost edge of $\mathcal{V}(w_i)$, for $i = 2k + 1, k = 0, \dots$, is a horizontal edge, e_i^h , and w_i and e_i^h have the same y -coordinate. By construction, the lowest edge of $\mathcal{V}(w_{i+1})$, e_{i+1}^h , is also horizontal, and its y -coordinate is larger than that of e_i^h . An analogous argument holds for the vertical edges limiting $\mathcal{V}(w_i)$ and $\mathcal{V}(w_{i+1})$ for $i = 2k, k = 0, \dots$. Thus, the visibility polygons of the w_i are pairwise disjoint. \square

We now place G_{cw} as follows: place a right-looking half guard on each w_i for which $i = 2k + 1, k = 0, \dots$, and a left-looking half guard on each w_i for which $i = 2k + 1, k = 1, \dots$, see Figure 11(b). Because the convex vertices are incident to vertical edges that limit their visibility to one half plane, the visibility polygons of the half guards coincide with the “normal” visibility polygons of our witnesses. We have $|G_{cw}| = |W| \leq |G^*|$.

For the construction of G_s , we consider the still unseen parts of \mathbf{P} : we can at most have $|G_{cw}|$ such polygon pieces (between each pair of witness visibility polygons and possibly one that includes either \mathbf{P} 's topmost or \mathbf{P} 's rightmost edge). We show that each such region is a staircase polygon for which either the upper or the lower chain has exactly two edges—a *stair*. Consider the placement of w_{2k+1} and $w_{2(k+1)}$: We place $w_{2(k+1)}$ on the first convex lower chain vertex that does not have the same y -coordinate as w_{2k+1} . Hence, the lower chain of the polygonal region between $\mathbf{V}(w_{2k+1})$ and $\mathbf{V}(w_{2(k+1)})$ consists of one horizontal edge (defined by the upper edge of w_{2k+1} 's visibility polygon) and one vertical edge (the upper end point of this edge has the same y -coordinate as $w_{2(k+1)}$). Thus, the polygonal region is a stair. We can place a left-looking half guard at its lowest-rightmost point that covers it completely, see Figure 11(c). Analogously, the polygonal region between $\mathbf{V}(w_{2k})$ and $\mathbf{V}(w_{2k+1})$ is a stair for which the upper chain has two edges, and we can guard it with a right-looking half guard at its highest-leftmost point. Thus, we obtain:

Theorem 11 *The set G covers all of \mathbf{P} and $|G| = |G_{cw}| + |G_s| \leq 2 \cdot |G_{cw}| \leq 2 \cdot |G^*|$.*

In fact, the result holds for vertex half guards: all half guards in G_{cw} are already placed on vertices; we observe that we can slide each guard $g \in G_s$ between w_{2k+1} and $w_{2(k+1)}$ down along the vertical edge it resides on without losing coverage of its stair, see Figure 11(c); analogously, a guard between w_{2k} and w_{2k+1} can be slid left along the horizontal edge it resides on.

References

- [1] M. Abrahamsen, A. Adamaszek, and T. Miltzow. The art gallery problem is $\exists\mathbb{R}$ -complete. *Journal of the ACM*, 69(1), 2021.
- [2] P. Belleville. Computing two-covers of simple polygons. Master's thesis, McGill University, 1991.

- [3] É. Bonnet and T. Miltzow. An Approximation Algorithm for the Art Gallery Problem. In *33rd International Symposium on Computational Geometry, SoCG'2017*, volume 77 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:15, 2017.
- [4] V. Chvátal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory B*, 18:39–41, 1974.
- [5] O. Daescu, S. Friedrichs, H. Malik, V. Polishchuk, and C. Schmidt. Altitude terrain guarding and guarding uni-monotone polygons. *Computational Geometry*, 84:22–35, 2019.
- [6] S. Eidenbenz. Inapproximability results for guarding polygons without holes. In *Algorithms and Computation, ISAAC'1998*, pages 427–437, 1998.
- [7] S. Fisk. A short proof of Chvátal's watchman theorem. *Journal of Combinatorial Theory B*, 24:374–375, 1978.
- [8] M. Gibson, E. Krohn, B. J. Nilsson, M. Rayford, and P. Zylinski. A note on guarding staircase polygons. In *31st Canadian Conference on Computational Geometry, CCCG'2019*, pages 105–109, 2019.
- [9] M. Gibson, E. Krohn, and M. Rayford. Guarding monotone polygons with half-guards. In *29th Canadian Conference on Computational Geometry, CCCG'2017*, pages 168–173, 2017.
- [10] H. M. Hillberg, E. Krohn, and A. Pahlow. On the complexity of half-guarding monotone polygons. <https://arxiv.org/abs/2204.13143>, 2022.
- [11] J. Kahn, M. Klawe, and D. Kleitman. Traditional art galleries require fewer watchmen. *SIAM Journal on Algebraic and Discrete Methods*, 4(2):194–206, 1983.
- [12] J. King and D. Kirkpatrick. Improved approximation for guarding simple galleries from the perimeter. *Discrete & Computational Geometry*, 46(2):252–269, 2011.
- [13] D.-T. Lee and A. K. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.
- [14] B. J. Nilsson and D. Wood. Watchmen routes in spiral polygons. In *2nd Canadian Conference on Computational Geometry, CCCG'1990*, pages 269–272, 1990.
- [15] J. O'Rourke. An alternate proof of the rectilinear art gallery theorem. *Journal of Geometry*, 21(1):118–130, 1983.
- [16] J. O'Rourke. *Art Gallery Theorems and Algorithms*. International Series of Monographs on Computer Science. Oxford University Press, New York, 1987.
- [17] J. O'Rourke and K. Supowit. Some NP-hard polygon decomposition problems. *Information Theory, IEEE Transactions on*, 29(2):181–190, 1983.