A 3-Approximation Algorithm for Guarding Orthogonal Art Galleries with Sliding Cameras

Stephane Durocher[1] and Saeed Mehrabi[1]

[1]Department of Computer Science, University of Manitoba, Winnipeg, Canada
{durocher, mehrabi}@cs.umanitoba.ca

Abstract. A *sliding camera* travelling along a line segment s in a polygon P can see a point p in P if and only if p lies on a line segment contained in P that intersects s at a right angle. The objective of the *minimum sliding cameras (MSC)* problem is to guard P with the fewest sliding cameras possible, each of which is a horizontal or vertical line segment. In this paper, we give an $O(n^3)$-time 3-approximation algorithm for the MSC problem on any simple orthogonal polygon with n vertices. Our algorithm involves establishing a connection between the MSC problem and the problem of guarding simple grids with *periscope guards*.

1 Introduction

Given a polygon P with n vertices in the plane, the art gallery problem is to find a minimum-cardinality set of guards such that every point in P is visible to at least one guard, where each guard g is a point in the plane that sees a point p if the line segment from g to p is contained in P. In the *orthogonal art gallery problem*, the input polygon P is orthogonal; that is, every edge of P is vertical or horizontal. The art gallery problem is NP-hard for both arbitrary [13] and orthogonal polygons [16]. Eidenbenz [4] proved that the art gallery problem is APX-hard on simple polygons, and that no polynomial-time algorithm can guarantee to find a solution with $o(\log n)$ times the minimum number of guards on polygons with holes, unless P=NP [5]. Ghosh [7] gave an $O(\log n)$-approximation algorithm for the art gallery problem that runs in $O(n^4)$ time on simple polygons and $O(n^5)$ time on polygons with holes. Krohn and Nilsson [12] gave a polynomial-time $O(OPT^2)$-approximation algorithm for the orthogonal art gallery problem, where OPT is the cardinality of the optimal solution. Many variants of the art gallery problem have been studied based on different types of visibility [14, 19], different polygonal domains (e.g., orthogonal polygons [8], or polyominoes [1]) and different types of guards (e.g., points or line segments). See the surveys by O'Rourke [15] or Urrutia [18] for a history of the art gallery problem.

Recently, Katz and Morgenstern [9] introduced a variant of the art gallery problem in which *sliding cameras* are used to guard an orthogonal polygon. Given an orthogonal polygon P with n vertices, a sliding camera is a point guard that travels back and forth along a horizontal or vertical line segment s inside P. The camera can see a point $p \in P$ if there is a point $q \in s$ such that the line segment $\overline{pq}$ is horizontal or vertical, and is contained in P. In the *minimum sliding cameras (MSC)* problem, the objective is to guard P using the minimum number of sliding cameras.

A grid D is a connected union of vertical and horizontal line segments; each maximal line segment in the grid is called a *grid segment*. We denote the set of grid segments of D by T_D. Moreover, a simple grid is defined as follows:

Definition 1 (Kosowski et al. [10]). *A grid D is* simple *if there exists $\delta > 0$ such that for every $\epsilon \in (0, \delta)$ and every grid segment d in D, both endpoints of d_ϵ lie in the outer face, where d_ϵ is the extension of d by ϵ units in both directions.*

A *periscope guard* x located on a grid segment s in a grid D is a point on s that sees a point y in D if some path from x to y in D has at most one bend. In other words, points x and y are mutually visible if and only if they lie on respective segments s_x and s_y in D such that $s_x \cap s_y \neq \emptyset$ (it could be that $s_x = s_y$). Periscope guards were introduced by Gewali and Ntafos [6] in their examination of the complexity of the orthogonal art gallery problem (the orthogonal art gallery problem was shown to be NP-hard three years later by Schuchardt and Hecker [16]). In the *minimum periscope guards (MPG)* problem on a grid, the objective is to guard the grid with the minimum number of periscope guards. The MPG problem can be defined on an orthogonal polygon P similarly: the goal is to locate the minimum number of periscope guards in P such that every point in P is guarded by at least one periscope guard.

Related Work. Katz and Morgenstern [9] first considered a restricted version of the MSC problem in which only vertical cameras are allowed; by reducing the problem to the minimum clique cover problem on chordal graphs, they solved the problem exactly in polynomial time. For the generalized case, where both vertical and horizontal cameras are allowed, they gave a 2-approximation algorithm for the MSC problem under the assumption that the polygon P is x-monotone. Durocher and Mehrabi [3] showed that the MSC problem is NP-hard when the polygon P is allowed to have holes (i.e., polygon P is not simple). They also gave an exact polynomial-time algorithm that solves a variant of the MSC problem, called the minimum-length sliding cameras (MLSC) problem, in which the objective is to minimize the sum of the lengths of line segments along which cameras travel. Seddighin [17] considered the MLSC problem under k-visibility, where a camera's line of sight can pass through k edges of the polygon, and proved that the MLSC problem is NP-hard under k-visibility for any fixed $k \geq 2$. Durocher et al. [2] gave an $O(n^{2.5})$-time (3.5)-approximation algorithm for the MSC problem on a simple orthogonal polygon with n vertices. Their algorithm uses different techniques from those used in this paper; specifically, it applies solutions to the minimum edge cover problem in graphs and the guarded mobile guard problem on grids (where each guard must be seen by at least one other guard). The complexity of the MSC problem on simple orthogonal polygons remains unknown.

Gewali and Ntafos [6] showed that the MPG problem is NP-hard on general three-dimensional grids and that it is polynomial-time tractable on simple two-dimensional grids (see Theorem 1). Moreover, Kosowski et al. [11] showed that the problem of guarding a two-dimensional grid with the minimum number of k-periscope guards is NP-hard (a point p on the grid is visible to a k-periscope guard g if there exists a path of at most k bends in the grid from p to g). Our results refer to the following theorem by Gewali and Ntafos [6].

Theorem 1 (Gewali and Ntafos [6]). *Given a simple two-dimensional grid G with n segments, the MPG problem can be solved exactly on G in $O(n^3)$ time.*

Our Result. In this paper, we give an $O(n^3)$-time 3-approximation algorithm for the MSC problem on any simple orthogonal polygon P. To this end, we describe a connection between the MSC problem on simple orthogonal polygons and the MPG problem on simple grids. We first construct a simple grid G_P associated with polygon P and then show that a reduction from the MSC problem on P to the MPG problem on grid G_P gives a set of sliding cameras whose cardinality is at most twice the cardinality of the solution to MPG problem on G_P. However, some new potentially unguarded regions are introduced. We show that the number of such unguarded regions is bounded from above by the cardinality of the optimal solution to the MPG problem, each of which can be guarded with a single sliding camera. Finally, we show that the cardinality of the optimal solution to the MPG problem is a lower bound for any feasible solution for the MSC problem. This results in an approximation factor of 3 (2 for each periscope guard in the solution of the MPG problem and 1 for guarding each unguarded region), improving the previous best approximation factor of 3.5 [2].

2 Preliminaries

Throughout the paper, let P denote a simple orthogonal closed polygonal with n vertices (including the polygon's interior). Observe that every simple orthogonal polygon with at most six vertices can be guarded by a single sliding camera; therefore, we assume throughout the paper that $n > 6$. Let OPT_P and OPT_{PG} denote optimal solutions for the MSC problem on P and the MPG problem on a simple grid, respectively. Let $V(P)$ denote the set of reflex vertices of P and let H_u and V_u be the maximum-length horizontal and vertical line segments, respectively, inside P through a vertex $u \in V(P)$. Let $L(P) = \{H_u \mid u \in V(P)\} \cup \{V_u \mid u \in V(P)\}$. Let L and L' be two orthogonal line segments (with respect to P) inside P; the *visibility region* of L is the union of points in P that are seen by a sliding camera that travels along L. We say L dominates L' if the visibility region of L' is a subset of that of L.

Let r be a reflex vertex of P. The lines through H_r and V_r partition the plane into four quadrants, exactly one of which contains the exterior of P in an ϵ-neighbourhood around r, for some $\epsilon > 0$; we call the quadrant that is opposite to this quadrant the *essential quadrant* of r, denoted by $Q(r)$. See Figure 1(a) for an example. Let u be a convex vertex of P such that both the vertices v and w of P that are adjacent to u are also convex. Let p and q denote the next vertices of P that are adjacent to v and w, respectively. We call vertex u a *pocket vertex* of

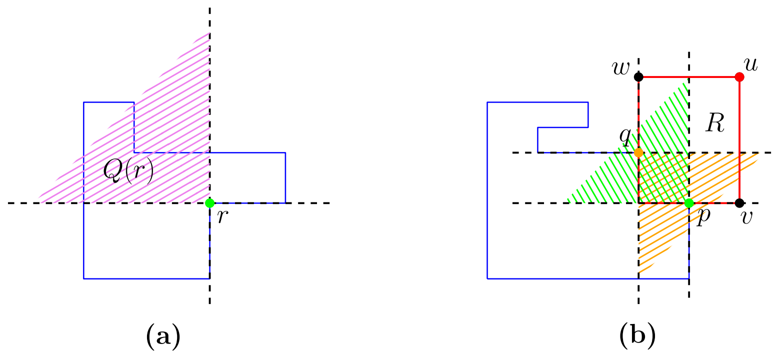

Fig. 1: (a) An example of a reflex vertex r with the essential quadrant $Q(r)$ (i.e., the open hatched quadrant) shown in pink. (b) An example of a pocket vertex u; both vertices p and q are reflex and $Q(p) \cap Q(q) \neq \emptyset$. The edges uv and uw are the pocket edges of the convex pocket R (red rectangle).

P if and only if (i) both the vertices p and q are reflex, and (ii) $Q(p) \cap Q(q) \neq \emptyset$. Moreover, we refer to the edges of P that are incident to a pocket vertex as the *pocket edges* of P and to the rectangular subregion of P whose sides are two of the pocket edges of P as a *convex pocket* of P. See Figure 1(b).

3 A 3-Approximation Algorithm

In this section, we describe the 3-approximation algorithm for the minimum sliding cameras (MSC) problem on simple orthogonal polygons. Given any simple orthogonal polygon P, we first construct a grid G_P associated with P as follows. Initially, let G_P be the set of all line segments in $L(P)$. Now, for any pair of reflex vertices u and v where H_u dominates H_v (resp., V_u dominates V_v) in P, we remove H_v (resp., V_v) from G_P; if two segments mutually dominate each other, remove one of the two arbitrarily. Next, for each convex pocket R of P, we add a segment into G_P for every pocket edge of R. We call a grid segment in G_P corresponding to a pocket edge of P a *pocket segment* of G_P. Let G_P denote the resulting grid. Each of the pocket segments remains in G_P even if it is dominated by another segment in G_P. See Figure 2 for an example.

Observe that the number of grid segments in G_P (i.e., $|T_G|$) is at most n, where n is the number of vertices of P. Moreover, G_P is simple because the construction preserves the property that the endpoints of each grid segment in T_G lie on the boundary of the polygon and, therefore, the endpoints of every grid segment in T_G lie on the outer face of G_P. To see that G_P is connected, it suffices to note that (i) the grid induced by the line segments in $L(P)$ is connected, and (ii) for each grid segment $s \in L(P)$ that is removed (due to domination), the set of grid segments that are intersected by s are also intersected by $s' \in T_G$, where s' is the grid segment that dominates s. Therefore, G_P is also connected and we have the following result.

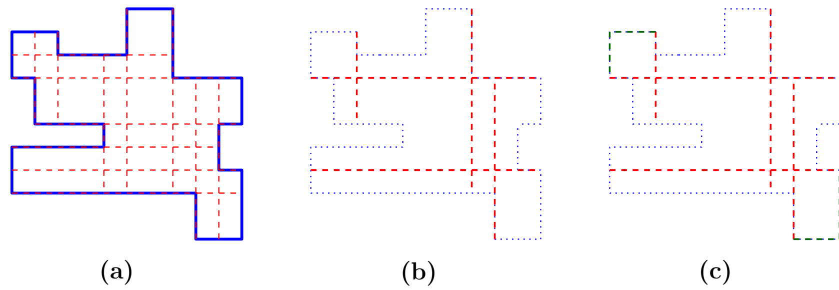

Fig. 2: (a) Polygon P with the initial grid G_P that consists of all line segments in $L(P)$. (b) Grid G_P after removing the dominated grid segments. (c) The final grid G_P after adding the segments corresponding to the pocket edges of the convex pockets of P; these grid segments are shown in green.

Lemma 1. *Grid G_P is a simple and connected grid.*

As described in Section 1, we reduce the MSC problem on P to the MPG problem on G_P. In general the visibility region of a periscope guard g cannot be guarded entirely by a single sliding camera; see Figure 3 for an example. Two sliding cameras suffice to guard the visibility region of a periscope guard.

Observation 1 *The visibility region of any periscope guard g in a polygon P can be guarded by the maximal vertical and horizontal line segments through g in P.*

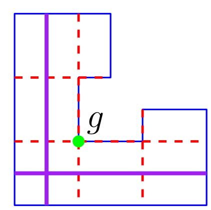

Fig. 3: Although the grid G_P induced by P can be guarded by a single periscope guard g, two sliding cameras (shown in purple) are needed to guard P.

By Theorem 1 we can obtain a set of periscope guards by solving the MPG problem on G_P in $O(n^3)$ time. Let M denote the set of sliding cameras obtained by placing a pair of sliding cameras on each periscope guard. Since the algorithm of Gewali and Ntafos [6] positions periscope guards only at the intersections of grid segments, this ensures that the sliding cameras located in P are all aligned with line segments in $L(P)$.

The procedure described above can result in a set M of sliding cameras whose cardinality exceeds three times that of an optimal solution. See Figure 4: the vertical segment in G_P that corresponds to the vertical pocket edge of convex pocket R cannot be guarded by periscope guard g_1, forcing the algorithm to add a second periscope guard, while a single sliding camera suffices to guard polygon P entirely. We now describe how to modify the grid G_P to bound $|M|$.

3.1 Pocket Segments and Desert Regions

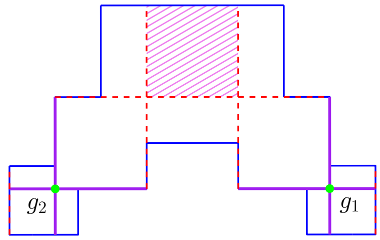

Fig. 5: A simple orthogonal polygon P and its corresponding grid G_P (dashed red). The set $\{g_1, g_2\}$ of periscope guards guard G_P. However, the sliding cameras located by the algorithm (solid purple) do not guard P entirely. The hatched pink region, called a *desert*, is not guarded by any sliding camera.

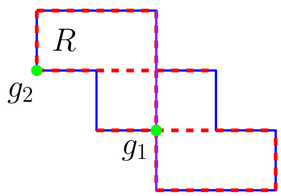

Fig. 4: By adding the two pocket segments into grid G_P corresponding to the pocket edges of every convex pocket of P, an optimal solution to the MPG problem uses two guards (g_1 and g_2). The algorithm for the MSC problem uses four sliding cameras (red), which is four times the size of the optimal solution (purple).

As illustrated in Figure 4, the cardinality of solution M may not be bounded by three times the cardinality of an optimal solution for the MSC problem. To resolve this problem, we add into G_P exactly *one* of the pocket segments corresponding to the pocket edges of every convex pocket of P as follows: let R be a convex pocket of P and let s_1 and s_2 be, respectively, the vertical and horizontal grid segments in G_P whose corresponding maximal line segments in P enter R. Observe that the vertical pocket edge (resp., horizontal pocket edge) of R intersects s_2 (resp., s_1). If the number of grid segments intersected by s_1 is greater than the number of grid segments intersected by s_2, then we remove from G_P the pocket segment that corresponds to the vertical pocket edge of R; otherwise, we remove the pocket segment that corresponds to the horizontal pocket edge of R. Note that we now have exactly one pocket segment in G_P for both the pocket edges of every convex pocket of P. We show later that by this modification the cardinality of M obtained by solving the MPG problem on G_P is at most three times that of OPT_P.

The set M might not still be a feasible solution for the MSC problem. See Figure 5 for an example. In the following, we characterize such unguarded regions, called the *desert regions*, and show that the number of desert regions is bounded from above by $|M|$. To characterize desert regions, take any unguarded point p in P and let R_p be a maximal axis-aligned rectangle contained in P that covers p and is also not guarded by the line segments in M. Observe that (i) rectangle R_p is visible to some line segments in G_P, and that (ii) no such line segments are in M because R_p is unguarded. Consider the maximal regions in P that lie immediately above, below, left, and right of R_p; we denote the union of these

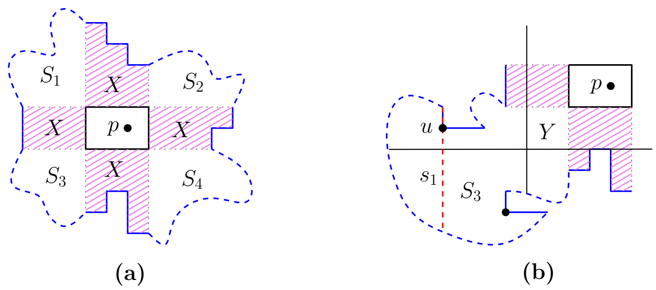

Fig. 6: (a) An unguarded point p inside a polygon P with maximal unguarded rectangle R_p. The hatched pink region of P indicates the region X; the four regions S_1, S_2, S_3 and S_4 are labelled accordingly. (b) An illustration in support of the proof of Lemma 2.

regions by X. See the hatched region in Figure 6(a) for an example. Any sliding camera that sees any part of R_p must intersect some region of X. Since R_p is unguarded, region X cannot contain any sliding camera in M; therefore, no periscope guard lies in X. Moreover, region X partitions the polygon into five subregions (see Figure 6(a)): the union of X and rectangle R_p, the subregion on the upper-left side of X (denoted by S_1), the subregion on the upper-right side of X (denoted by S_2), the subregion on the lower-left side of X (denoted by S_3) and the subregion on the lower-right side of X (denoted by S_4). Note that the periscope guards in S can only lie in regions S_1, S_2, S_3 and S_4. We first show the following results.

Lemma 2. *If for some $1 \leq i \leq 4$, the subregion S_i contains no periscope guards of S, then all reflex vertices of P in S_i face the unguarded rectangle R_p.*

Proof. Without loss of generality, assume that there is no periscope guard in S_3. Suppose, to the contrary of the lemma statement, that there exists a reflex vertex u of P in S_3 that is not faced towards rectangle R_p. Observe that there are only two possibilities for such reflex vertex as shown in Figure 6(b). We now continue the proof for the upper vertex u shown in Figure 6(b); the proof for the other vertex is similar. Consider the maximal vertical line segment s_1 that passes through u and let Y be the set of line segments in G_P that enter region S_3 from the other regions. First, note that either $s_1 \in G_P$ or otherwise $s_j \in G_P$, for some s_j that dominates s_1. Without loss of generality, assume that $s_1 \in G_P$ (otherwise, the proof will be similar by replacing s_1 with s_j). Since R_p is unguarded, there is no sliding camera and, therefore, no periscope guard located on L, for all $L \in Y$. Since there is no periscope guard in S_3, we conclude that s_1 is not guarded by any periscope guard, which is a contradiction to the fact that S is a feasible solution to the MPG problem on G_P. Therefore, all the reflex vertices of P inside S_3 must face towards rectangle R_p. □

By Lemma 2, we conclude that if there is no periscope guard in S_i, for some $1 \leq i \leq 4$, then the region S_i must be bounded by at most two staircases with their reflex vertices all facing towards the unguarded rectangle R_p. There are two possibilities for the staircases to lie in S_i depending on the orientation of the staircases: they can be either both horizontal or both vertical; see Figure 8 for an illustration in which $S_i = S_1$. Moreover, Lemma 2 implies that region S_i is *orthogonally convex*, because otherwise there must be a reflex vertex in S_i that is not faced toward rectangle R_p and, therefore, there will be a grid segment in S_i that is not guarded by any periscope guard.

Lemma 3. *If the subregion S_i, for some $1 \leq i \leq 4$, contains no periscope guards of S, then the subregion S_i has no convex pockets.*

Proof. Without loss of generality, assume that there is no periscope guard in S_3. Suppose, to the contrary of the lemma statement, that there exists a convex pocket R inside S_3 and let Y be the set of grid segments in G_P that intersect at least one of the pocket edges s_1 and s_2 of R; see Figure 7. Without loss of generality, assume that $s_1 \in G_P$. We first show that there is no periscope guard on L, for all $L \in Y$. Take any grid segment L in Y. Note that if L is entirely contained in S_3, then there is no periscope guard on L by the assumption. If L enters S_3 from another region, then it must intersect region X and, therefore, rectangle R_p is visible to L. Since R_p is unguarded, there is no sliding camera (and therefore no periscope guard) on L. This means that s_1 is not guarded by any periscope guard, which is a contradiction to the fact that S is a feasible solution to the MPG problem. This completes the proof of the lemma. □

By Lemma 3, we conclude that the staircases of S_i are joined with each other in such a ways that they do not create any convex pocket in S_i.

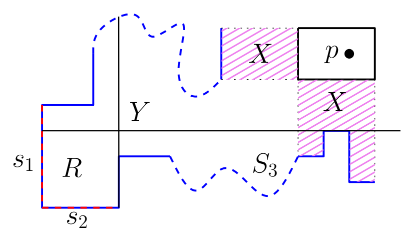

Fig. 7: An illustration in support of the proof of Lemma 3.

3.2 Characterizing Desert Regions

Recall that the periscope guards in S can only lie in $S_1 \cup S_2 \cup S_3 \cup S_4$. The structure of a desert region depends on how many of the four regions S_1, S_2, S_3 and S_4 contain at least one periscope guard. In the following, we consider all the four cases and show that the desert region in each case can be guarded entirely by a single sliding camera. Let $Z \subseteq \{S_1, S_2, S_3, S_4\}$ such that $S_i \in Z$, for all $1 \leq i \leq 4$, if and only if S_i contains at least one periscope guard.

Case 1: $|\mathbf{Z}| = \mathbf{4}$. In this case, there is at least one periscope guard in S_i, for all $1 \leq i \leq 4$. Since (i) the grid segments in G_P are all guarded by at least one periscope guard, and (ii) each part of region X (i.e., the parts that are immediately above, below, to the left and to the right of rectangle R_p) is intersected by at least one grid segment in G_P, we conclude that the region S_i is guarded by sliding cameras in M, for all $1 \leq i \leq 4$. Therefore, the desert region in this case is just the rectangle R_p and can be guarded by a single sliding camera. See Figure 5 for an example.

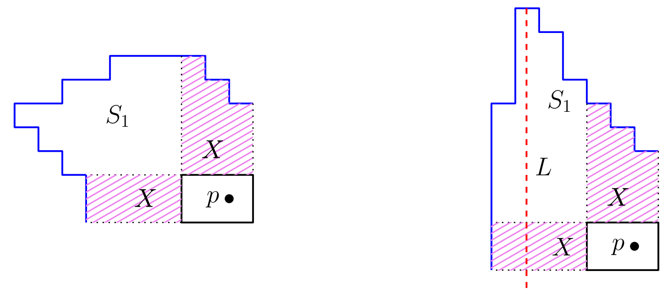

(a) **(b)**

Fig. 8: An example of a region S_1 such that it contains no periscope guards of S. The staircases in S_1 must both be either (a) horizontal, or (b) vertical. Note that S_1 is an orthogonally convex region.

Case 2: $|\mathbf{Z}| = \mathbf{3}$. Without loss of generality, assume that there is no periscope guard in region S_1. Note that the desert region in this case is the union of S_1 and rectangle R_p. Recall that the staircases in S_1 must both be horizontal or both vertical. Assume without loss of generality that the staircases are lied vertically in S_1 (i.e., Figure 8(b)). Since S_1 is an orthogonally convex region the maximal vertical line segment L that crosses the top most horizontal edge of S_1 guards the union of S_1 and rectangle R_p; we call L the *neighbour camera* associated with S_1. Therefore, one sliding camera located on L can guard the desert region.

Case 3: $|\mathbf{Z}| = \mathbf{2}$. Let S_i and S_j, for some $1 \leq i, j \leq 4$ and $i \neq j$, be the regions that contain no periscope guard. We observe that in this case, the desert region is the union of S_i, S_j and rectangle R_p. There are two cases depending on the positions of S_i and S_j:

(a) Suppose regions S_i and S_j are neighbours to each other. Without loss of generality, assume that $S_i = S_1$ and $S_j = S_2$ and that the staircases in S_1 are lied vertically; see Figure 9(a). Let L_1 and L_2 be the neighbour cameras of S_1 and S_2, respectively. If the staircases in S_2 are also lied vertically, then it is straightforward to see that there exists a maximal horizontal line segment inside P that guards the union S_1, S_2 and rectangle R_p (see Figure 9(a)). If the staircases in S_2 are lied horizontally, we show that L_1 guards the union of S_1, S_2 and rectangle R_p. First, note that L_1 guards the union of S_1 and rectangle R_p. Now, suppose to the contrary, that there exists a point $q \in S_2$ that is not visible to L_1; see Figure 9(b). Since S_2 consists of only two staircases and such staircases in S_2 are lied horizontally, they must be joined with each other in S_2 such that they form a convex pocket, which is a contradiction to Lemma 3. Therefore, S_2 is entirely visible to L_1.

(b) Suppose regions S_i and S_j are opposite to each other. Note that (i) each of the regions S_i and S_j consists of at most two staircases, and (ii) the staircases in S_i (or in S_j) are either both vertical or both horizontal. By an argument

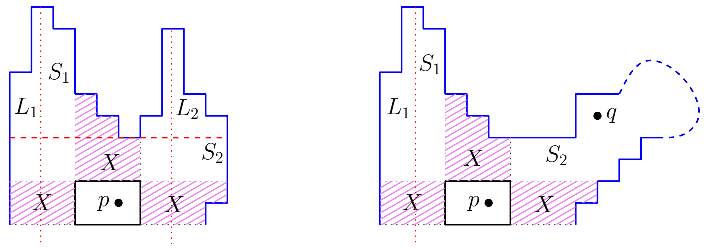

(a) **(b)**

Fig. 9: An illustration in support of Case 3.

analogous to that given in Case (a), we can conclude that the union of regions S_i, S_j and rectangle R_p can be guarded by one sliding camera.

By the two cases described above, we conclude that the desert region can be guarded by one sliding camera.

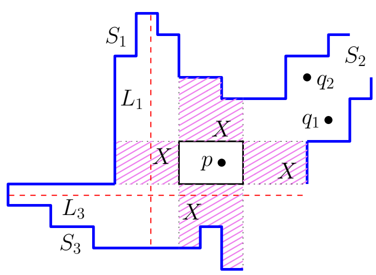

Fig. 10: An illustration in support of Case 4.

Case 4: $|\mathbf{Z}| = 1$. Without loss of generality, assume that all the periscope guards lie in S_4. We show that the subregion $P \setminus \{S_4\}$, which forms the desert region, can be guarded by a single sliding camera. Consider the neighbour camera L_3 associated with region S_3 and assume without loss of generality that the staircases in S_3 lie horizontally; see Figure 10. It is straightforward to see that L_3 guards the union of S_3 and rectangle R_p. We now check to see if L_3 can also guard the union of S_1 and S_2. If L_3 guards the union of S_1 and S_2, then the subregion $P \setminus \{S_4\}$ can be guarded by one sliding camera located on L_3. Otherwise, there are two possibilities:

(a) Suppose exactly one of the regions S_1 or S_2 is guarded by L_3. Without loss of generality, assume that S_2 is not guarded by L_3 entirely. Thus, there is a point $q_1 \in S_2$ that is not visible to L_3; see Figure 10. Since L_3 guards S_1 the staircases in S_1 must be lied vertically. Therefore, the neighbour camera L_1 (associated with region S_1) is vertical and guards the union of S_1, S_3 and rectangle R_p. Note that L_1 also guards S_2 because otherwise there must be a point $q_2 \in S_2$ that is not visible to L_1 (see Figure 10). But, the existence of points q_1 and q_2 in S_2 implies that the staircases in S_2 must be joined with each other in such a way that they form a convex pocket in S_2, which is a contradiction to Lemma 3. Therefore, in this case, L_1 guards the desert region entirely.

(b) Suppose neither S_1 nor S_2 is guarded entirely by L_3. Since L_3 is horizontal, the staircases in regions S_1 and S_2 must all have lain horizontally and, therefore, all the staircases in subregion $P \setminus \{S_4\}$ lie horizontally. It is now easy to observe that in this case there exists a maximal vertical line segment inside P that guards the subregion $P \setminus \{S_4\}$.

By the two possibilities above, we conclude that the desert region can be guarded by one sliding camera.

We observe that in each of the Cases 1 through 4, at least one periscope guard is required in characterizing the desert region. Therefore, by Cases 1 through 4 described above, we have the following theorem.

Theorem 2. *Every point in P that is not inside a desert region is guarded by at least one sliding camera in M. Each desert region of P consists of a set of staircases and it can be guarded entirely by a single sliding camera. Moreover, the number of desert regions is at most the number of periscope guards in S.*

To summarize the algorithm, we first solve the MPG problem on G_P and compute the set S of optimal periscope guards. By Observation 1, we locate two sliding cameras inside P for each periscope guard to obtain the set M. By Theorem 2, we then guard each desert region by a single sliding camera; let M' denote the set of sliding cameras that guard the desert regions. By Theorem 2, the set $M \cup M'$ of sliding cameras guards P entirely.

3.3 Analyzing the Algorithm

In this section, we analyze the running time and the approximation factor of the algorithm. To this end, we first give a lower bound on any feasible solution for the MSC problem on P. Recall OPT_P, an optimal solution to the MSC problem, and recall OPT_{PG}, an optimal solution for the MPG problem on G_P. We show the following result whose proof is omitted due to space constraints.

Lemma 4. $|OPT_P| \geq |OPT_{PG}|$.

We know that $|M| \leq 2 \cdot |S|$. By Theorem 2, we have that $|M'| \leq |S|$ and so $|M \cup M'| \leq 3 \cdot |S|$. Therefore, by Lemma 4 we have that $|M \cup M'| \leq 3 \cdot |OPT_P|$. To analyze the running time of the algorithm, we note that the construction of grid G_P can be completed in $O(n^2)$ time [2]. Since $|T_G| = O(n)$, where n is the number of vertices of P, the MPG problem can be solved on G_P in $O(n^3)$ time. Next, the desert regions of P can be detected in $O(n^2)$ time by detecting the visibility region of cameras in M and comparing their union with P. Finally, the desert regions can be guarded in $O(n)$ time by locating a sliding camera inside P, for each desert region. Therefore, we have the main result of this paper.

Theorem 3. *There exists an $O(n^3)$-time 3-approximation algorithm for the MSC problem on any simple orthogonal polygon with n vertices.*

4 Conclusion

In this paper, we gave an $O(n^3)$-time 3-approximation algorithm for the problem of guarding a simple orthogonal polygon P with n vertices using the minimum number of sliding cameras (i.e., the MSC problem). The complexity of the MSC problem is still unknown and remains the main direction for future work. Also, giving algorithms with better approximation factor or showing a hardness of approximation remains open as another direction for future work.

References

1. T. C. Biedl, M. T. Irfan, J. Iwerks, J. Kim, and J. S. B. Mitchell. The art gallery theorem for polyominoes. *Disc. & Comp. Geom.*, 48(3):711–720, 2012.
2. S. Durocher, O. Filtser, R. Fraser, A. D. Mehrabi, and S. Mehrabi. A (7/2)-approximation algorithm for guarding orthogonal art galleries with sliding cameras. In *Proc. LATIN*, pages 294–305, LNCS 8392, 2014.
3. S. Durocher and S. Mehrabi. Guarding orthogonal art galleries using sliding cameras: algorithmic and hardness results. In *Proc. MFCS*, pages 314–324, LNCS 8087, 2013.
4. S. Eidenbenz. Inapproximability results for guarding polygons without holes. In *Proc. ISAAC*, volume 1533 of *LNCS*, pages 427–436, 1998.
5. S. Eidenbenz. *Inapproximability of visibility problems on polygons and terrains.* PhD thesis, ETH Zurich, 2000.
6. L. Gewali and S. C. Ntafos. Covering grids and orthogonal polygons with periscope guards. *Comput. Geom.*, 2:309–334, 1992.
7. S. K. Ghosh. Approximation algorithms for art gallery problems in polygons. *Disc. App. Math.*, 158(6):718–722, 2010.
8. F. Hoffmann. On the rectilinear art gallery problem. In *Proc. ICALP*, pages 717–728, 1990.
9. M. J. Katz and G. Morgenstern. Guarding orthogonal art galleries with sliding cameras. *Int. J. of Comp. Geom. & App.*, 21(2):241–250, 2011.
10. A. Kosowski, M. Malafiejski, and P. Zylinski. An efficient algorithm for mobile guarded guards in simple grids. In *Proc. ICCSA (1)*, pages 141–150, 2006.
11. A. Kosowski, M. Malafiejski, and P. Zylinski. Cooperative mobile guards in grids. *Comp. Geom.*, 37(2):59–71, 2007.
12. E. Krohn and B. J. Nilsson. Approximate guarding of monotone and rectilinear polygons. *Algorithmica*, 66(3):564–594, 2013.
13. D. T. Lee and A. K. Lin. Computational complexity of art gallery problems. *IEEE Trans. on Info. Theory*, 32(2):276–282, 1986.
14. R. Motwani, A. Raghunathan, and H. Saran. Covering orthogonal polygons with star polygons: the perfect graph approach. In *Proc. ACM SoCG*, pages 211–223, 1988.
15. J. O'Rourke. *Art gallery theorems and algorithms.* Oxford University Press, Inc., New York, NY, USA, 1987.
16. D. Schuchardt and H. Hecker. Two NP-hard art-gallery problems for ortho-polygons. *Math. Logic Quarterly*, 41(2):261–267, 1995.
17. S. Seddighin. Guarding polygons with sliding cameras. Master's thesis, Sharif University of Technology, 2014.
18. J. Urrutia. Art gallery and illumination problems. In *Handbook of Comp. Geom.*, pages 973–1027. North-Holland, 2000.
19. C. Worman and J. M. Keil. Polygon decomposition and the orthogonal art gallery problem. *Int. J. of Comp. Geom. & App.*, 17(2):105–138, 2007.