Half-Guarding Weakly-Visible Polygons and Terrains

Nandhana Duraisamy
PSG College of Technology, Coimbatore, India

Hannah Miller Hillberg
University of Wisconsin, Oshkosh, WI, USA

Ramesh K. Jallu[1]
Indian Institute of Information Technology Raichur, India

Erik Krohn
University of Wisconsin, Oshkosh, WI, USA

Anil Maheshwari
Carleton University, Ottawa, Canada

Subhas C. Nandy
Indian Statistical Institute, Kolkata, India

Alex Pahlow
University of Wisconsin, Oshkosh, WI, USA

Abstract

We consider a variant of the art gallery problem where all guards are limited to seeing 180°. Guards that can only see in one direction are called half-guards. We give a polynomial time approximation scheme for vertex guarding the vertices of a weakly-visible polygon with half-guards. We extend this to vertex guarding the boundary of a weakly-visible polygon with half-guards. We also show NP-hardness for vertex guarding a weakly-visible polygon with half-guards. Lastly, we show that the orientation of half-guards is critical in terrain guarding. Depending on the orientation of the half-guards, the problem is either very easy (polynomial time solvable) or very hard (NP-hard).

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases Art Gallery Problem, Approximation Algorithm, NP-Hardness, Monotone Polygons, Half-Guards

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2022.18

Related Version This paper combines the results of [7] and [12], which were obtained independently. The first gave a PTAS for vertex guarding the boundary of a weakly-visible polygon. The second gave a PTAS for vertex guarding the vertices of a weakly-visible polygon and provided hardness results for half-guarding terrains and weakly-visible polygons. A preliminary version of [12] is presented in the Iranian Conference on Computational Geometry 2022 (an informal meeting).
Previous Version: `https://iccg2022.ce.sharif.edu/files/Papers/4-Guarding-Weakly-Visible-Polygons-with-Half-Guards.pdf`

1 Introduction

An instance of the original *art gallery problem* takes as input a simple polygon P having a sequence of vertices $V = \{v_1, v_2, \ldots, v_n\}$, and reports a guarding set of points $G \subseteq P$ such that every point $p \in P$ is seen by a point in G. Here, by the term *a point q sees a point p*

[1] The work was partially done while the author was affiliated with Indian Statistical Institute, Kolkata, India.

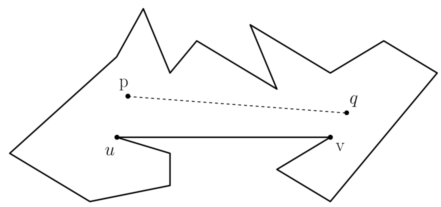

Figure 1 A WV-polygon where p sees q and every point in the polygon is seen by a point on $e = (u, v)$ or sees a point on e.

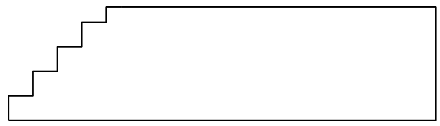

Figure 2 A WV-polygon that requires $\Omega(n)$ half-guards that see to the right. Half-guards need to be placed at every step like structure.

where $p, q \in P$, we mean that the line segment $[p, q]$ completely lies inside the polygon P. In this paper, we study the vertex guarding problem which says that guards are only allowed to be placed at the vertices of V, and the objective is to find the smallest such G.

1.1 Definitions

We use $p.x$ to denote the x-coordinate for point p. A weakly-visible polygon (WV-polygon) P contains a visibility edge $e = (u, v)$ such that every point in P sees at least one point on the edge e. When referencing a half-guard p sees to the right (i.e., guards see 180°towards the right), then a point q is visible to p if $p.x \leq q.x$ (see Fig. 1). Without loss of generality, we assume that the visibility edge e is aligned with the x-axis.

A (strictly) x-monotone terrain T with x-axis as the base is defined by a set of points $V = \{v_1, v_2, \ldots, v_n\}$ satisfying $v_i.x < v_{i+1}.x$ for $i = 1, 2, \ldots, n-1$, and $[v_1.x, v_n.x]$ is aligned with the x-axis.

1.2 Previous Works

Extensive studies are done on algorithmic aspects of *art galleries*. Several results related to hardness and approximations are available in [1, 8, 11, 17]. Due to the inherent difficulty in fully understanding the art gallery problem for simple polygons, work has been done for guarding polygons with some special structures (see [2, 4, 15]). In this paper we consider terrains and weakly-visible polygons.

Motivation of considering these two special structure comes from the fact that many cameras/sensors cannot generally sense in 360°, referred to as *full-guards* in this paper. In [2], the authors proposed a constant-factor approximation algorithm to guard interior of a WV-polygon. We consider the *half-guards* that can sense in 180°. For weakly-visible polygons, we restrict the problem even further by only allowing these half-guards to see to the right. Even with these restrictions, the vertex guarding problem is difficult to solve in weakly-visible polygons. Using half-guarding the authors in [9, 14] propose a 4-approximation algorithm for terrain guarding with full-guards. In Fig. 2, we demonstrate that there are WV-polygons P that can be completely guarded with one full-guard, but requires $\Omega(n)$ half-guards.

The difficulty with guarding terrains with half-guards is less straightforward. As shown in [5], it is possible to optimally guard a terrain in polynomial time using half-guards that only see to the right. If the visibility of half-guards is modified such that they see down in the terrain setting, as we will show, the problem becomes NP-hard.

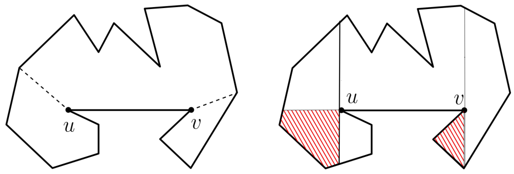

Figure 3 On the left, a full-guard placed at u and v cuts off the polygon below the WV-edge. On the right, the shaded regions below $\overline{uv}$ are still unseen after half-guards are placed at u and v.

1.3 Our Contribution

Ashur et al. [3] gives a local search based polynomial time approximation scheme (PTAS) for minimum dominating set in terrain-like graphs. They then show that several families of polygons have a visibility graph that is terrain-like. One such family is WV-polygons. However, their analysis uses the fact that vertices can see 360° and that the visibility graph of the vertices contains undirected edges. Since edges in a visibility graph for a WV-polygon that use half-guards are directed, it does not imply that the visibility graph is terrain-like. We provide additional observations in this paper which show a PTAS for vertex guarding the vertices of a WV-polygon with half-guards. We extend this to show that guarding the boundary of a WV-polygon with half-guards placed at vertices also admits a PTAS.

NP-hardness has been shown for many variants of the art gallery problem [4, 8, 17]. In many of those reductions, guards are allowed to see in all directions. If the problem is restricted enough, it can become polynomially time solvable, for example, see [6, 16]. If the polygon is restricted to be a WV-polygon, guards are restricted to be at the vertices and guards are only allowed to see to the right, even with these many restrictions, the problem is still NP-hard. To obtain this result, we first show that a variant of half-guarding a terrain is NP-hard.

In Section 2, we provide a PTAS for vertex guarding the vertices and boundary of a WV-polygon using half-guards that see right. In Section 3, we show that the half-guarding of terrains is NP-hard when the half-guards see down. In Section 4, the NP-hardness is shown for vertex guarding a WV-polygon using half-guards that see right.

2 PTAS for Vertex Guarding a WV-Polygon with Half-Guards

We use local search to obtain a PTAS for half-guarding a WV-polygon. We adapt the idea of the proof provided in [3] to this problem with some notable exceptions. The divergent details of the algorithm and proof along with a few identical details, included for readability, are provided below.

Much of the proof from [3] assumes that the WV-polygon lies above the WV-edge by placing guards at u and v and cutting off the portion of the polygon beneath the WV-edge, see Fig. 3(left). When every vertex under consideration lies above this WV-edge, the so-called order claim holds and the visibility graph is terrain-like. However, consider doing the same thing for a WV-polygon using half-guards that see to the right. In this case, the remaining unseen portion of the polygon cannot be assumed to be above the WV-edge. The shaded region in Fig. 3(right) demonstrates the invisible portions to the left-of-and-below the guards at u and v. A guard g *dominates* another guard g' if for every vertex v that g' sees, g also sees v. Unlike with full-guards, a half-guard placed at u will not dominate a half-guard placed in the shaded region to the left-of-and-below u. Thus, the proof of the order claim for WV-polygons with full-guards does not imply that the order claim is true for WV-polygons with half-guards.

Algorithm 1 Local Search Algorithm.

We set $k = \frac{\alpha}{\epsilon^2}$, for appropriate constants $\alpha, \epsilon > 0$. Let Q be the set of guards.
1. Initialize Q to the vertex set V.
2. Determine if there exists a subset $S \subseteq Q$ of size at most k and a subset $S' \subseteq V \setminus Q$ of size at most $|S| - 1$ such that $(Q \setminus S) \cup S'$ guards V.
3. If such S and S' exists then, set $Q \leftarrow (Q \setminus S) \cup S'$ and go to step 2. Else, return the set Q.

2.1 Algorithm Overview

In this section, we will assume that $e = (u, v)$ is a convex edge, that is, the interior angles at u and v are convex. Without loss of generality, we assume that $P \setminus e$ is contained in the open half-plane above the x-axis. We later show how to solve for the general case, where the convexity assumption of the angles at u and v is removed. The local search algorithm, described below, is identical to the one described in [3].

2.2 Local Search Analysis

For any two points a and b on the boundary of the polygon P, we use $a \prec b$ (or $b \succ a$) to say that a precedes b (or b succeeds a), i.e., when one reaches a before b, while traversing the boundary of P, clockwise from u. Claim 1 was proved to be true for weakly-visible polygons with full guards. The order claim proof from [3] shows that a and d see each other. However, they do not claim that d must be to the right of a. With full-guards, it does not matter since if a sees d, then d sees a. The following proof of the order claim for *half-guards* shows why a sees d. An important note to remember is that this claim is only true if the angles at u and v are convex.

▷ Claim 1. Let a, b, c, d be four points on the boundary of a WV-polygon such that $a \prec b \prec c \prec d$. If a sees c and b sees d, then a must see d (see Fig. 4).

Proof. Since a sees c and b sees d, $a.x \leq c.x$ and $b.x \leq d.x$. If $a.x \leq b.x$, then $a.x \leq b.x \leq d.x$. Using the proof from [3], a must see d. If $b.x < a.x$, then consider the line segment connecting b to d. Since $a \prec b \prec c \prec d$, it must be the case that the $\overline{bd}$ line segment crosses the $\overline{ac}$ line segment, see Fig. 4. Let o be the intersection point of $\overline{ac}$ and $\overline{bd}$. By definition, $a.x \leq o.x$ and $o.x \leq d.x$. Therefore, $a.x \leq d.x$ and a sees d. ◁

It should be noted that the orientation of the WV-edge does not matter with Claim 1. In other words, if the polygon is rotated in any direction, the original order claim from [3] still holds, i.e. a and d must see each other, see Fig. 4(right).

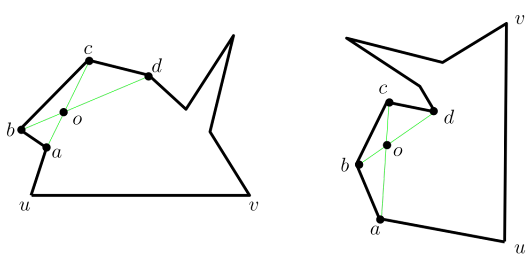

Figure 4 Order claim example. The orientation of the $\overline{uv}$ edge does not matter. If $a < b < c < d$, a sees c and b sees d, then a sees d.

Let the red set R be an optimal solution (minimum cardinality guard set), and the blue set B be the solution returned by Algorithm 1. For each vertex $x \in (R \cup B)$, we define color(x) to be the color of the vertex x, which is red if $x \in R$, and blue if $x \in B$. We can assume that $R \cap B = \emptyset$. Any vertices that are in both R and B are removed, i.e. the local search algorithm found an optimal guard and one can ignore vertices which are seen by such a guard. We prove that $|B| \leq (1+\epsilon) \cdot |R|$. We construct a bipartite graph $G = (R \cup B, E)$. As in [3], we can prove that (i) G is planar and (ii) G satisfies the locality condition as defined below.

▶ **Definition 2.** *A graph $G = (R \cup B, E)$ satisfies* ***locality condition*** *if for each vertex $w \in P$, there exist vertices $r \in R$ and $b \in B$, such that (i) r sees w, (ii) b sees w, & (iii) $(r, b) \in E$.*

Next, by using the proof scheme of Mustafa and Ray [19], if G satisfies these two conditions, then we have $|B| \leq (1+\epsilon) \cdot |R|$. As in [3], we now define $\lambda(\cdot)$ and $\rho(\cdot)$ for each vertex of P. While traversing the boundary clockwise from u, for a vertex $w \in V$, if there exists a vertex in $R \cup B$ that sees w and it precedes w on the path of traversal from u to w, then $\lambda(w)$ is set to the first such vertex. Similarly, $\rho(w)$ is set to the first vertex in $R \cup B$ that sees w while traversing the boundary counterclockwise from v. Since $R \cap B = \emptyset$, for every vertex $w \in V$, $\lambda(w)$ and/or $\rho(w)$ must be defined in $R \cup B$.

2.3 Constructing G

▷ Claim 3. Let $A_1 = \{(\lambda(w), w) \mid w \in V \text{ for which } \lambda(w) \text{ is defined}\}$. Then the segments of A_1 are non-crossing.

Proof. Consider two segments $(\lambda(x), x), (\lambda(y), y) \in A_1$, such that $\lambda(x) \neq \lambda(y)$. Without loss of generality assume that $\lambda(x) \prec \lambda(y)$. By contradiction, suppose the segments $(\lambda(x), x)$ and $(\lambda(y), y)$ intersect, then it must be the case where $\lambda(x) \prec \lambda(y) \prec x \prec y$. But, by Claim 1, this would imply that $\lambda(x)$ sees y, which is impossible by the definition of $\lambda(y)$. Thus, the segments are non-crossing. ◁

The edges of G are similar to the edges of G described in [3]. If it can be shown that the order claim holds and that the segments of A_1 are non-crossing, the details of how to create the graph are essentially the same. However, the description of those edges are provided here for completeness. For each vertex $x \in R \cup B$, if $\lambda(x)$ is defined and color($\lambda(x)$) $\neq$ color(x), add the edge $(\lambda(x), x)$ to E_1. Otherwise, if there exists a segment $(\lambda(w), w) \in A_1$, such that (i) $\lambda(w) \prec x \prec w$, (ii) $(\lambda(w), w)$ can be reached from x without going outside of P and without intersecting any other segment in A_1 (except possibly at x), and (iii) color($\lambda(w)$) $\neq$ color(x), add the edge $(\lambda(w), x)$ to E_1.

Analogously, we define the sets A_2 and E_2 by replacing λ with ρ. This implies, $A_2 = (w, \rho(w)) \mid w \in V$ for which $\rho(w)$ is defined}, and E_2 is defined with respect to A_2. For each vertex $w \in R \cup B$, if both $\lambda(x)$ and $\rho(x)$ are defined and color($\lambda(x)$) $\neq$ color($\rho(x)$), then add the edge $(\lambda(x), \rho(x))$ to E_3. Let $G = (R \cup B, E)$, where $E = E_1 \cup E_2 \cup E_3$. Now, we have the following lemmas.

▶ **Lemma 4.** *The graph $G = (R \cup B, E)$ is (a) planar, and (b) satisfies locality condition.*

Proof. We prove that G is planar by finding a suitable embedding of it in the plane. Let C be a circle of radius $r > 0$ in the plane. We map the vertices of P to equally-spaced points on C, and the edges in E are drawn between pairs of points inside and outside of

C. If $A_1 \cap E_1 \neq \emptyset$, then the edges in $A_1 \cap E_1$ are drawn inside C as straight line segments, and are non-crossing due to Claim 3. Note that for each edge $(\lambda(w), x) \in E_1 \setminus A_1$, there exists a segment $(\lambda(w), w) \in A_1$ such that $\lambda(w) \prec x \prec w$ and x can reach $(\lambda(w), w)$ without exiting from P and without intersecting any segment in A_1. We can argue that the edge $(\lambda(w), x)$ can be drawn as a straight line segment inside C without intersecting any edge that was already drawn. Suppose $(\lambda(w), x)$ intersects an edge (segment) $(\lambda(a), a) \in A_1 \cap E_1$. Now, if $\lambda(a) \prec \lambda(w) \prec a \prec x$, then the segments $(\lambda(a), a)$ and $(\lambda(w), w)$ are crossing, a contradiction to Claim 3. If $\lambda(w) \prec \lambda(a) \prec x \prec a$, there arise two cases: (i) $a \prec w$ or $a = w$, in this case x crosses the segment $(\lambda(a), a)$ to reach $(\lambda(w), w)$, a contradiction; (ii) $w \prec a$, in which case the segments $(\lambda(w), w)$ and $(\lambda(a), a)$ are crossing, a contradiction to Claim 3. If, suppose, $(\lambda(w), x)$ intersects an edge $(\lambda(w'), y) \in E_1 \setminus A_1$, then $\lambda(w) \prec \lambda(w') \prec x \prec y$. Implies $(\lambda(w), x)$ crosses the segment $(\lambda(w'), w') \in A_1$. If $(\lambda(w'), w') \in E_1$, then it is similar to the previous case. On the other hand, if $(\lambda(w'), w') \in E_1$, there arise two cases: either $w' \prec w$ or $w' = w$, in which the contradiction is that x crosses $(\lambda(w'), w')$ to reach $(\lambda(w), w)$; or $w \prec w'$, in which $\lambda(w)$ sees w' (order claim applied to $\lambda(w), \lambda(w'), w, w'$), contradiction to the definition of $\lambda(w')$. Thus the edges E_1 are properly embedded inside C. We embed the edges of E_2 outside C as curved line segments. The edges in $A_2 \cap E_2$, if any, can be drawn without intersecting due to the fact that the segments in A_2 are non-crossing. We can argue that the edges of $E_2 \setminus A_2$ can be embedded without crossing as in the previous case. Observe that the edges in E_3 correspond to a vertex $x / \in R \cup B$, having $\lambda(x)$ and $\rho(x)$ defined, such that $\text{color}(\lambda(x)) \neq \text{color}(\rho(x))$. Also, observe that neither $(\lambda(x), x)$ nor $(x, \rho(x))$ belong to $E_1 \cup E_2$ by the way E_1 and E_2 are defined. Hence, each edge $(\lambda(x), \rho(x)) \in E_3$ can be drawn as the union of the segments $(\lambda(x), x) \in A_1$ and $(x, \rho(x)) \in A_2$. Therefore, G is planar as all the edges in E are drawn without any intersection.

(b) Let x be any vertex of P. We will show that there exists $r \in R$ and $b \in B$, such that r sees x, b sees x, and $(r, b) \in E$. We argue by distinguishing between the following two cases. Case (i) $x \notin R \cup B$: If both $\lambda(x)$ and $\rho(x)$ are defined and $\text{color}(\lambda(x)) \neq \text{color}(\rho(x))$, then $(\lambda(x), \rho(x)) \in E_3$, and hence the claim is true. If both $\lambda(x)$ and $\rho(x)$ are defined and $\text{color}(\lambda(x)) = \text{color}(\rho(x))$, then there must exists $w \in R \cup B$ such that w sees x, and $\text{color}(w)$ is different from $\text{color}(\lambda(x))$ and $\text{color}(\rho(x))$. Without loss of generality we assume that w be the first such vertex while traversing P's boundary clockwise from u, and $\lambda(x) \prec w \prec x$ (a similar argument works if $x \prec w \prec \rho(x)$ assuming w be the last such vertex while traversing P's boundary clockwise from u). Let $(\lambda(y), y) \in A_1$ associate with w, implies $\lambda(x) \prec \lambda(y) \prec w \prec y \prec x$. If $y \neq x$, then $\lambda(y)$ sees x due to order claim on the vertices $\lambda(y), w, y$, and x. If $y = x$, then $\lambda(x) = \lambda(y)$, so $\lambda(y)$ sees x. Since w is the first vertex that sees x and satisfies $\lambda(x) \prec w \prec \rho(x)$, it must be that $\text{color}(\lambda(y)) \neq \text{color}(w)$. So the edge $(\lambda(y), w) \in E_1$. Hence, the claim is true. The above argument works even if only $\lambda(x)$ (or $\rho(x)$) is defined.

Case (ii) $x \in R \cup B$: If $\lambda(x)$ is defined and if $\text{color}(\lambda(x)) \neq \text{color}(x)$, then $(\lambda(x), x) \in E_1$, and hence the claim is true. Similarly, if $\rho(x)$ is defined and if $\text{color}(\rho(x)) \neq \text{color}(x)$, then $(x, \rho(x)) \in E_2$, and hence the claim is true. Suppose $\lambda(x)$ is defined, but $\text{color}(\lambda(x)) = \text{color}(x)$ (the argument is similar if $\rho(x)$ is defined and $\text{color}(x) = \text{color}(\rho(x))$). There must exist $w \in R \cup B$ such that w sees $x, \lambda(x) \prec w \prec x$, and $\text{color}(\lambda(x)) \neq \text{color}(w)$. Without loss of generality, let w itself be the first such vertex (while traversing P's boundary clockwise from u). Now, we can prove the claim by proceeding as in Case (i). ◀

▶ **Lemma 5.** *There exists a PTAS for vertex guarding a weakly-visible polygon, with visibility edge $e = (u, v)$ where the angles of u and v are convex, with half-guards where half-guards can only see to the right.*

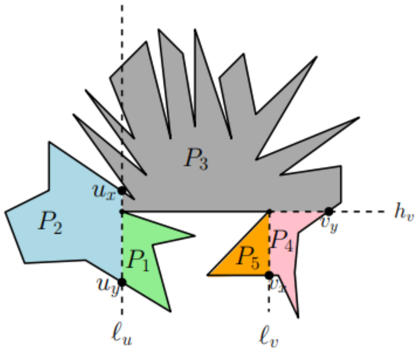

Figure 5 Removing the convexity assumption.

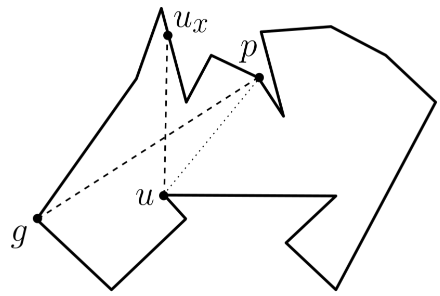

Figure 6 Any point $p \in P_3$ that is seen by a point $g \in P_2$ is also seen by u.

2.4 Removing the Convexity Assumption

In [3], the convexity assumption is handled by placing guards at u and v. It is proved that every vertex "below" the edge $e = (u, v)$ is seen by either u or v. For half-guards, if guards are similarly placed at u and v the analysis of [3] does not extend to half-guarding. A portion of the polygon below the $e = (u, v)$ edge is unseen, see, for example, P_5 and P_2 from Figure 5. We now show how to remove the convexity assumption of u and v. Let l_u be the vertical line going through u; u_x and u_y be the points of intersection of l_u with the boundary of P. Let h_v be the horizontal line through v that hits the boundary of P at v_y to the right of v. Let l_v be the vertical line through v that hits the boundary of P at v_x below v. The segments $\overline{u_x u_y}$, $\overline{v v_x}$, and $\overline{v v_y}$ split the polygon into at most five sub-polygons, say P_1, P_2, P_3, P_4, P_5; each weakly-visible from some edge (see Fig. 5). Thus, vertex guarding the vertices of each sub-polygon ensures vertex guarding the vertices of P. As shown in [3], if the WV-polygon can be guarded using only a constant number (say c) of vertices, then we try all such possibilities and obtain an optimal solution in $O(n^c)$ time. Assuming the polygon cannot be guarded with a constant number of guards, we place guards at u and v and show that the final guarding set, including these two guards, is a $(1 + \epsilon)-$approximation.

The entire P_1 region must be seen by the point u on the edge $e = (u, v)$. All the vertices in P_4 are seen by v. Thus, vertex guarding the vertices of P_1 and P_4 can be achieved by placing 1 half-guard at u and v, respectively. Let U be the set of vertices in P_3 that are neither guarded by u nor v. We apply Algorithm 1 on P_3 with the WV-edge $\overline{uv}$ to vertex guard U. Let S_3 be the solution obtained.

Observe that the polygons P_2 and P_5 are weakly-visible from $\overline{u_x u_y}$, and $\overline{v v_x}$, respectively. Hence, Algorithm 1 can be applied individually since u_x, u_y, v, and v_x are convex vertices in their respective sub-polygons. Since Claim 1 (half-guard order claim) holds for these polygons, the A_1 segments in those polygons are non-crossing. Thus, one can construct a graph G for each sub-polygon where G is planar and G satisfies the locality condition. Let S_2 and S_5 be the solutions obtained for P_2 and P_5, respectively.

It remains to show that the solutions S_2, S_3 and S_5 are disjoint. In other words, the local search algorithm can be run independently on each sub-polygon. No guard in S_2 will see any of the U vertices in P_3 nor P_5. Similarly, no guard in S_3 will see vertices in P_2 nor P_5. Lastly, no guard in S_5 will see vertices in P_2 nor the U vertices in P_3.

It is immediately obvious that no guard in S_3 nor S_5 can see any vertex in P_2. Any such vertex would need to see left to cross the $\overline{u_x u_y}$ line segment and no guard can see left. In a similar fashion, no guard of S_2 nor S_3 can see any vertex in P_5. In order to see into P_5, a guard must see to the left of the $\overline{vv_x}$ line segment and no guard can see left. When guarding P_3, we assumed that guards were already placed at u and v and therefore, we only need to run the local search algorithm on the remaining vertex set $U \in P_3$ such that no vertex in U is seen by either u or v.

▷ Claim 6. Any vertex $p \in P_3$ that is seen by any boundary point $g \in P_2$ is seen by u.

Proof. By definition, u sees u_x since u_x is directly above u. It must be the case that $u \prec g \preceq u_x$. Since u sees u_x and g sees p, the polygon cannot pierce the $\overline{uu_x}$ nor $\overline{gp}$ line segments. The only way to block u from seeing p is to come from the "other side," see Fig. 6. However, if blocking occurs in this manner, then p no longer has a line segment connecting it to the $\overline{uv}$ edge, contradicting the fact that the polygon is weakly-visible. Therefore, u must see p. ◁

▷ Claim 7. Any vertex $p \in P_3$ that is seen by a boundary point $g \in P_5$ is seen by v.

The set S_2, S_3 or S_5 may not be a feasible solution of P as S_2, S_3 and S_5 may contain guards at u_x, u_y, v_x, or v_y, which are not vertices of P. Assuming v_y is not a vertex, it is not considered by the algorithm for P_3. The vertex that precedes v_y is considered by the local search algorithm for P_3. No guard will be placed at v_x when running the local search algorithm on P_5. A guard at v_x only sees itself and v. A guard was already placed at v making any guard placed at v_x redundant.

If the local search algorithm places a guard at u_y or u_x when creating a solution for P_2, then that guard was only placed to guard u_y and/or u_x. Since the rest of the P_2 polygon is to the left of u_x and u_y, guards placed here will not see any other vertex of P_2. If this happens, then we sweep a vertical line leftward starting at $\overline{u_x u_y}$ until it hits a vertex of P_2. The guard is moved to this vertex. That vertex will see both u_x and u_y. As shown in Claim 6, neither u_x nor u_y will be responsible for guarding a vertex in P_3 and it is obvious they will not contribute to the other sub-polygons. Therefore, no guards will be placed at u_x, u_y, v_x nor v_y.

▶ **Theorem 8.** *For any weakly-visible polygon, there exists a PTAS for vertex guarding the vertices of the polygon with half-guards.*

Proof. Let P be a polygon that is weakly-visible from an edge $e = (u, v)$. We divide the polygon P into at most five sub-polygons as discussed above. Let S_2, S_5 be the set of half-guards obtained by our algorithm for P_2 and P_5. The obtained set S_3 guards the set of vertices U of P_3 that are not seen by the half-guards placed at u and v. Let OPT be an optimal half-guard set of vertices of P. Let $S = S_2 \cup S_3 \cup S_5 \cup \{u, v\}$.

It can be observed that, the vertices in P_2 (resp. P_5) can only be guarded by half-guards placed at the vertices of P_2 (resp. P_5). Let O_i be an optimal solution for vertex guarding the sub-polygon P_i, and $OPT_i \subset OPT$ be the set of half-guards (in the optimum solution of P) lying in P_i. We get, $|S_i| \leq (1+\epsilon) \cdot |O_i| \leq (1+\epsilon) \cdot |OPT_i|$ as $|O_i| \leq |OPT_i|$.

Consider the sub-polygon P_3. Let O_3 be an optimal solution for vertex guarding the set U in P_3. Mote that, the subset of vertices U is not guarded by $OPT_1 \cup OPT_2 \cup OPT_4 \cup OPT_5$. Let $OPT_3 \subset OPT$ be the set of guards lying in P_3 to guard U. We have, $|O_3| \leq |OPT_3|$, implying $|S_3| \leq (1+\epsilon) \cdot |O_3| \leq (1+\epsilon) \cdot |OPT_3|$. Combining the above inequalities, we have $|S| \leq (1+\epsilon) \cdot |O_2| + (1+\epsilon) \cdot |O_3| + (1+\epsilon) \cdot |O_5| + 2/$ Thus, $|S| \leq (1+\epsilon) \cdot OPT + 2$. We can get rid of 2 in the inequality by adjusting k in the algorithm. ◀

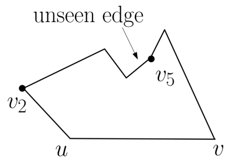

Figure 7 The vertices are half-guarded by v_2 and v_5 but an edge remains unseen.

2.5 Guarding the boundary of a weakly-visible polygon

In the previous section, we vertex guarded the vertices of a WV-polygon. However, even though all vertices are seen, it is possible that part of the boundary is unseen (see Fig. 7). Let $AGP(G, W)$ denote the problem of guarding the point set W with minimum number of entries in G. Let ∂P denote the boundary of P. We consider the following problem. Given a weakly-visible polygon P, construct a set $W \subset \partial P$ (witness points) such that an optimal solution for $AGP(V, W)$ is also an optimal solution for $AGP(V, \partial P)$. This discretization is based on Friedrichs et al. [10].

Computation of a witness set

We have a finite guard set V, which are the vertices of P. In order to guard ∂P, we find a discrete set of witness points $W \subset \partial P$. Let $g \in V$ be one of the half-guard candidates and $\mathcal{V}(g)$ be the visibility region of g on P's boundary ($\mathcal{V}(g) \subseteq \partial P$). $\mathcal{V}(g)$ creates $O(n)$ closed intervals (may be degenerate to a point) along the boundary ∂P. Considering the intervals generated by all the members in the guard set G, we split the entire boundary ∂P into maximal intervals, called *features* such that every point in a feature f is seen by the same set of guards $G(f) = \{g \in V \mid f \subseteq \mathcal{V}(g)\}$. Note that, (i) the union of the features is ∂P, and (ii) if any point in a feature is guarded by some guard in G, then the entire feature is guarded by that half-guard.

We pick any arbitrary point of each feature to the witness set W. So if we guard these $O(n|V|)$ witnesses in W, then the entire feature set, and as a result, ∂P is guarded. We can further reduce the number of witnesses by only using those features f with inclusion-minimal[2] $G(f)$. Let F be the feature set and G be the guard set (here $G = V$). Let w_f be an arbitrary point in a feature $f \in F$, then the witness of G is defined as $W = \{w_f \mid f \in F, G(f) \text{ is inclusion minimal}\}$.

▶ **Lemma 9.** *For a weakly-visible polygon P, with vertex set V and guard set G (here $G = V$), any feasible solution of $AGP(G, W)$ is also a feasible solution of $AGP(G, \partial P)$.*

Proof. Let $C \subset G$ be a feasible guard set of the witness set W. Suppose C does not guard some point $p \in \partial P$. Since every point on ∂P is visible to at least one vertex, some half-guard $g \in G$ should guard p. Therefore, p must be part of some feature $f \in F$. The set W either contains some witness $w_f \in f$ or a witness $w_{f'}$ such that $G(f') \subseteq G(f)$ (inclusion-minimality). In the first case, p must be guarded by C, otherwise w_f will not be guarded, and hence C is not a feasible solution for $AGP(G, W)$, leading to a contradiction. In the second case, some half-guard $g \in C$ guards $w_{f'}$. That half-guard should also guard f, thus it guards p; thus arriving at a contradiction of the statement that p is not guarded. ◀

[2] There does not exist another feature f' such that $G(f') \subseteq G(f)$

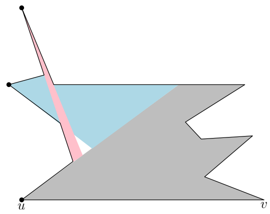

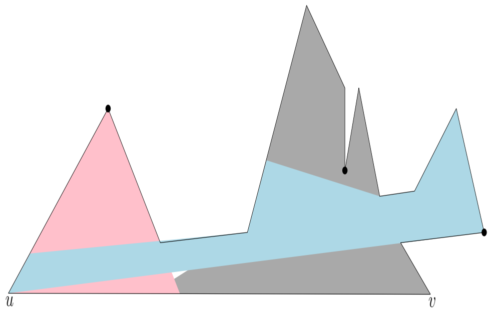

Figure 8 Guards that see the boundary of P but do not see the entire interior of P. (a) Example with half-guards. (b) Example with full-guards.

We apply the algorithm discussed above to guard the witness set W with the guard set V. Thus, we have the following result:

▶ **Theorem 10.** *For any weakly-visible polygon P, there exists a PTAS for finding the minimum number of half-guards at the vertices of P, that are required to guard the entire boundary of P.*

▶ Remark. Vertex guarding the entire boundary may not necessarily imply vertex guarding the entire polygon. In other words, even if the boundary is fully guarded there may be some pockets (invisible regions) inside the polygon. In Fig. 8(a), any solution (even an optimal one) guarding the boundary must have guards at the tip of the spike for guarding the spikes in the polygon as the guards can see only to the right. If the local search algorithm encounters the guards shown in the figure it will report that guard set. A similar observation can be observed with full-guards too. Note that for the weakly-visible polygon shown in Fig. 8(b), at least three full-guards are necessary to cover the boundary. If the local search algorithm encounters the guards shown in the figure, it will return them and exit as it cannot improve the solution further. In this case, there is a pocket. Indeed, there is another set of three points that cover the entire boundary as well as the interior of the polygon. However, the local search algorithm might not pick them.

3 NP-hardness for Vertex Guarding a Terrain with Half-Guards

Abusing notation, only in this section, we will assume that half-guards can only see "down." Let $p.y$ be the y-coordinate of a point p. We define seeing *down* as: a point p sees a point q if $p.y \geq q.y$. We begin this section by briefly sketching how the NP-hardness reduction works for terrain guarding with full-guards. The interested reader is encouraged to read [13] to see the full details of the original reduction. The reduction modifications begin in Section 3.1.

The terrain guarding reduction is from PLANAR 3SAT [18] where an instance has n variables and m clauses. The reduction works by assigning vertices on the terrain to truth values of variables from the PLANAR 3SAT instance. For each variable in the PLANAR 3SAT instance, variable gadgets are created such that the gadget contains a vertex representing the literal x_i and a vertex representing the literal $\overline{x_i}$, see Figure 9(middle). These variable gadgets are grouped together in chunks on the terrain. Figure 9(left) shows an example of one such chunk that contains one variable gadget for each variable in $\{x_1, x_2, \ldots, x_n\}$. These chunks are replicated on the terrain such that a guard placed at the vertex representing the literal x_i in chunk C_j would require a guard to be placed at the vertex representing the literal x_i in a different chunk C_k. There are points on the terrain that correspond to clauses in the PLANAR 3SAT instance. For example, if clause $c_i = (x_{i+1} \vee \overline{x_{i+3}} \vee x_{i+4})$ were in the original PLANAR 3SAT instance, then a point on the terrain would exist that is seen by three vertices corresponding to the literals x_i, $\overline{x_{i+3}}$ and x_{i+4}. If one of those vertices has a guard placed on it, then the clause would be satisfied. If none of those vertices has a guard placed on it, then an extra guard would be required to guard the terrain. If some minimum number of guards were placed and the entire terrain was seen, then the original instance was satisfied. If not, then the original instance was unsatisfiable.

The terrain guarding reduction only works if guards are full-guards. If the guards are half-guards that only see down, part of the terrain would go unseen. For example, in Figure 9, the portion of the terrain near the top could potentially be unseen. In the original reduction, a guard placed in the highest variable gadget of chunk will see this portion. Other parts of the terrain may also go unseen with the original reduction. Several tweaks and modifications to the reduction need to happen to ensure all of the terrain is guarded.

3.1 Terrain Hardness Modifications

If we restrict guards to be half-guards that see down in the terrain, then the terrain guarding problem with these half-guards is NP-hard. In regular terrain guarding, a point p *sees* another point q if the line segment connecting p and q does not go below the terrain. In this half-guard variant, the point p sees q only if the y-coordinate of p is greater than or equal to the y-coordinate of q and the line segment connecting p and q does not go below the terrain.

We describe the changes to the gadgets of [13] that must be made in order for the reduction to hold. Each part will explain why the original gadget doesn't work for this half-guard variant and what changes must be made in order to have the reduction hold.

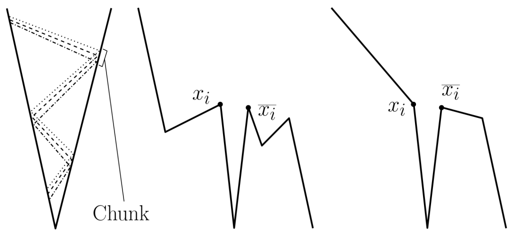

Figure 9 The left shows an overview of the NP-hardness reduction for terrain guarding. The middle is a variable gadget. The right is a starting gadget.

Between Chunks: Let us order the chunks from top to bottom $(C_1, C_2, \ldots, C_m)$. We will assume that all variables gadgets in chunk C_i are below all variable gadgets in chunk C_{i-1}. In the original reduction, the terrain between C_i the C_{i+2} is seen by any guard placed in the C_{i+1} chunk. Since guards can only see down in this variant, a portion of the terrain, which we will call u, below the lowest variable gadget in C_i and above the highest variable gadget in C_{i+1} will be unseen, see Figure 10(left). To fix this, we place a new gadget on the other side of the terrain from C_i. This gadget is at the same y-coordinate of the lowest variable gadget in chunk C_i. This ensures that the new guard will not affect the mirroring. The vertices f and g can be seen from vertex e and by no other point outside of this small region. Placing a guard at e will see f and g and also see the unseen region u below chunk C_i.

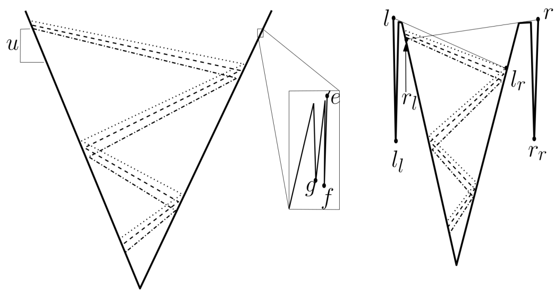

Figure 10 (Left) To ensure the entire terrain is seen, a guard is placed at e to see vertices f and g along with the u region. (Right) The overview of the terrain reduction where guards see down.

Variable Gadget: The following minor modifications need to be done in the variable gadgets to work for proving the hardness for half-guards.

Assume the $\overline{x_i}$ vertex in the variable gadget in chunk C_j has a guard placed on it and it sees b. In this case, a guard is placed at $\overline{x_i}$ in the variable gadget in chunk C_{j+1} to see a and c. The entire x_i variable gadget in chunk C_{j+1} is seen. Likewise, if a guard placed at x_i in chunk C_j sees a, then a guard is placed at x_i in chunk C_{j+1} sees b and c. A small portion of the terrain below $\overline{x_i}$ in chunk C_{j+1} may have been missed, see Figure 11. However, if the $\overline{x_i}$ vertex in chunk C_j is lowered just slightly, then the guard placed at x_i in C_j will see this region and an additional guard is not required. One must ensure the previously placed x_i does not see b but it can see anything in this gadget above b. Therefore, we ensure that $\overline{x_i}$ in the previous variable gadget is placed in such a way that it blocks x_i from seeing just above b in the subsequent variable gadget.

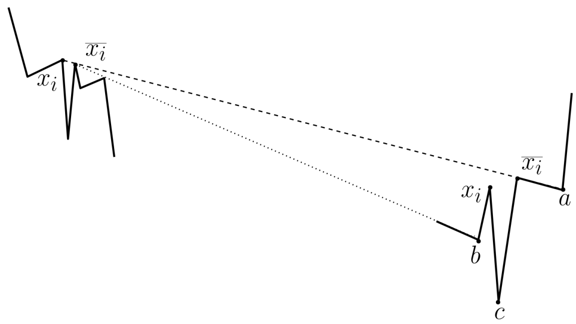

Figure 11 The variable gadget remains mostly unchanged from [13].

Removing a Variable: We will assume we are removing a variable from chunk C_i. When a variable gadget is removed going upwards, the gadget is modified slightly to remove the a and b vertices. Such a gadget is also called a starting gadget. When a guard is placed at the "lower" vertex in this gadget, a small portion of the terrain below the "higher" vertex remains unseen. In Figure 9(right), a starting gadget is shown. A guard placed at $\overline{x_i}$ would not see the small portion of the terrain below the x_i guard. To ensure this region is seen, we look at the variable patterns placed in the chunk above it, chunk C_{i-1}. The guard placed in the lowest variable pattern will see all of these potentially unseen portions. For example, in Figure 10(right), the guard placed in the variable gadget, x_k, directly above the r_l point will see all of these unseen regions in the variable patterns between l_r and the variable gadget for x_k in chunk C_2. Therefore, no modification is needed and no additional guard is needed.

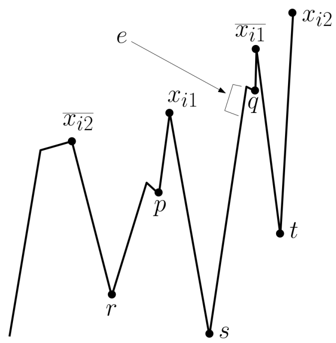

Figure 12 The inversion gadget for variable x_i in chunk C_{j+1}. This gadget is not tweaked but the variable gadget x_i in chunk C_j is modified slightly.

When removing a variable going down, the variable gadget is simply removed. The guard placed in the previous chunk's lowest variable gadget will see this region. If the lowest variable was being removed, then the e guard in the gadget that sees between chunks will see this region (see Figure 10(left)).

Above chunk C_1 and C_2: Since guards cannot see "up," a guard must be placed that guards the terrain above the first set of variable gadgets in chunk C_1 and above the variable gadgets in C_2. Assume without loss of generality that C_1 is on the left side of the terrain. Two guards are placed at points l and r as shown in Figure 10(right). These guards are required to see their own set of distinguished points, namely l_l, l_r, r_r and r_l. They see the "top" of the terrain above chunks C_1 and C_2. Since all gadgets in chunk C_1 are starting gadgets, r_l is placed below chunk C_1 to ensure the relevant part of the terrain in the starting gadgets and the terrain above the chunk are seen.

Inversion Gadgets: The updating required for the inversion gadgets are stated below:

See Figure 12 for a sample inversion gadget placed in chunk C_{j+1}. If a guard from the variable gadget x_i in chunk C_j sees point p, then when guards are placed at $\overline{x_{i1}}$ and $\overline{x_{i2}}$, the entire gadget is seen. If the guard from chunk C_j guard sees point q, then guards are placed at x_{i1} and x_{i2}. This leaves a small portion of the terrain unseen, the line segment e in Figure 12. To fix this, similar to the tweak of the variable gadget, the previously placed guard in chunk C_j is tweaked such that it sees just over the x_{i1} guard to see e. In this example, the previously placed guard must see q and not see p. As long as it is blocked from p, this is all that matters.

Clause Gadgets: No change is needed for the clause gadgets.

Putting it all together: As seen above, certain tweaks and updates are made to ensure that the entire terrain is seen. Making these changes will cause the minimum number of guards that must be placed, k, to increase by $m + 1$. An additional $m - 1$ guards are needed to see the unseen regions between chunks. Two additional guards are added at l and r to see the "top" of the terrain. This gives us a total of $m - 1 + 2 = m + 1$ additional required guards. None of these additional required guards see any of the original distinguished points of the terrain. They see their own set of distinguished points and also see the portion of the terrain that would have been unseen. As shown in [13], if k guards can guard the entire terrain, the instance is satisfiable. If more than k are needed, then the instance is not satisfiable.

▶ **Theorem 11.** *Finding the smallest vertex guard cover for guarding a terrain using half-guards that see down is NP-hard.*

4 NP-hardness for Vertex Guarding a WV-Polygon with Half-Guards

In this section, we give an overview of how to tweak the reduction from the previous section to show that vertex guarding a WV-polygon with half-guards is NP-hard. We will show how to use the terrain guarding hardness result for guards that only see down to show that vertex guarding a WV-polygon with half-guards that see to the right is NP-hard. One can take the modified terrain reduction, rotate it counterclockwise 90° and connect vertex l to vertex r to create a WV-polygon that is visible from the edge $e = (l, r)$, see Figure 13(left). The reduction holds the same way that it does for vertex guarding a terrain with half-guards that only see down.

One will notice that if the WV-edge is rotated slightly counterclockwise, the reduction still holds, see Figure 13(right). The key visibilities from the guards remain and the polygon is still weakly-visible. To help with seeing this, consider the variable patterns of Figure 14. The WV-polygon can be rotated, for example, 10° counterclockwise and the x_i and $\overline{x_i}$ guards still sees "right" to the distinguished points that they needs to see. The reduction begins to fail whenever a guard visibility from the original reduction starts to have a negative slope. In other words, if some point that guard x_i is supposed to see is to the left of x_i. If this happens, the guard no longer sees the distinguished point(s) to its right. To account for this, the original polygon is "stretched" further up-and-to-the-right such that none of the guards visibility have a negative slope. We now give a sketch of how to stretch the terrain/polygon while still maintaining the hardness reduction.

4.1 Stretching Hardness

For the terrain guarding hardness reduction, if half-guards see down and the terrain is rotated 1° in a counterclockwise direction, the terrain does not need to be modified much in order to maintain visibilities. The e guards placed to see between chunks C_i and C_{i+2} need to be shifted "up" to see the correct portion of the terrain between the chunks, see Figure 10. All gadgets are stretched but their visibilities do not change. A similar stretching idea is done for WV-polygons.

An example of stretching a WV-polygon is seen in Figure 15. In the bottom part of this Figure, the WV-edge is parallel to the x-axis and visibility to distinguished points are still to the right. As seen in the example, the original right-most vertex of the original WV-polygon is pulled "up-and-to-the-right" to ensure the polygon remains weakly-visible and all important lines of sight look to the right. The l and r vertices are tweaked slightly to ensure they also see their respective distinguished vertices to the right. Each gadget is also stretched "up-and-to-the-right" to maintain visibilities. An example of such a stretching

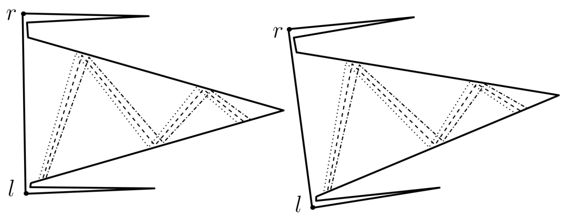

Figure 13 An overview of NP-hardness for WV-polygons.

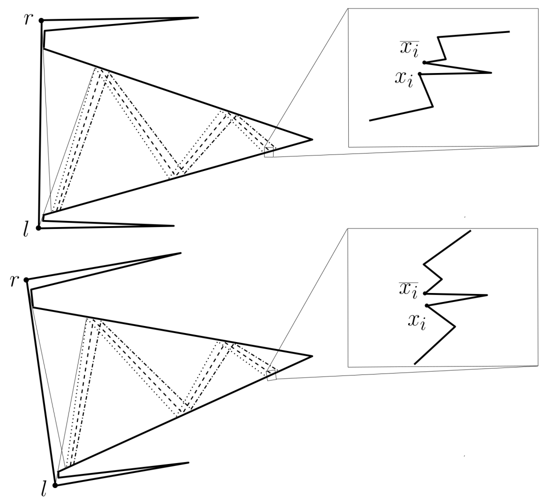

Figure 14 The top shows an original variable gadget. The bottom shows a rotated variable gadget.

is shown in Figure 15(bottom). In the original reduction, Figure 15(top), the x_i and $\overline{x_i}$ variables see right but they also see "up-and-to-the-right" to mirror its value to the next variable gadget. In the stretched WV-polygon, the variable gadget is stretched but the visibility of x_i and $\overline{x_i}$ that see "up-and-to-the-right" still exist. Those lines of sight are simply steeper than in the original reduction.

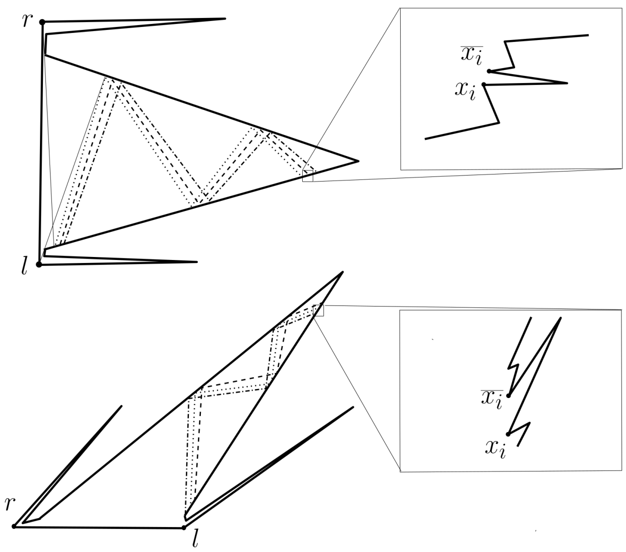

Figure 15 A stretched WV-polygon with a rotated variable gadget.

This stretching of the polygon works even if the WV-edge has a positive slope. The polygon can continue to be stretched as far "up-and-right" as necessary. However, the tweak fails when the WV-edge is a vertical edge and the inside of the polygon is to the left of the edge. We leave this as an open problem.

5 Conclusion and Future Work

In this paper, we present a PTAS for vertex guarding the vertices and boundary of a WV-polygon with half-guards that see to the right. This algorithm works regardless of the orientation of the WV-edge. We present an NP-hardness proof for vertex guarding a terrain with half-guards that see down. In contrast, if the half-guards see to the right for vertex guarding a terrain, the problem is polynomial time solvable. We also present an NP-hardness proof for vertex guarding a WV-polygon with half-guards that see to the right. Such a proof works for all instances except when the WV-edge is parallel to the y-axis and the "inside" of the polygon is to the left of the WV-edge. Whether or not this problem is NP-hard is left as an open problem. Future work might include finding a better approximation for the point guarding version of this problem. Insights provided in this paper may help with guarding polygons where the guard can choose to see either left or right, or in other natural directions. One may also be able to use these ideas when allowing guards to see 180° but guards can choose their own direction, i.e. 180°-floodlights.

References

1 Alok Aggarwal. *The art gallery theorem: its variations, applications and algorithmic aspects*. PhD thesis, The Johns Hopkins University, 1984.

2 Stav Ashur, Omrit Filtser, and Matthew Katz. A constant-factor approximation algorithm for vertex guarding a WV-polygon. *Journal of Computational Geometry*, 12(1), December 2021. doi:10.20382/jocg.v12i1a6.

3 Stav Ashur, Omrit Filtser, Matthew J. Katz, and Rachel Saban. Terrain-like graphs: PTASs for guarding weakly-visible polygons and terrains. *Computational Geometry*, 101:101832, 2022. doi:10.1016/j.comgeo.2021.101832.

4 Pritam Bhattacharya, Subir Ghosh, and Bodhayan Roy. Approximability of guarding weak visibility polygons. *Discrete Applied Mathematics*, 228, February 2017. doi:10.1016/j.dam.2016.12.015.

5 Danny Z. Chen, Vladimir Estivill-Castro, and Jorge Urrutia. Optimal guarding of polygons and monotone chains. In *Proceedings of the 7th Canadian Conference on Computational Geometry, Quebec City, Quebec, Canada, August 1995*, pages 133–138. Carleton University, Ottawa, Canada, 1995. URL: http://www.cccg.ca/proceedings/1995/cccg1995_0022.pdf.

6 Marcelo C. Couto, Pedro J. de Rezende, and Cid C. de Souza. An exact algorithm for minimizing vertex guards on art galleries. *Int. Trans. Oper. Res.*, 18(4):425–448, 2011. doi:10.1111/j.1475-3995.2011.00804.x.

7 Nandhana Duraisamy, Ramesh K. Jallu, Anil Maheshwari, and Subhas C. Nandy. A PTAS for vertex guarding weakly-visible polygons using half-guards. Unpublished Manuscript, 2019.

8 Stephan Eidenbenz. Inapproximability results for guarding polygons without holes. In *ISAAC*, pages 427–436, 1998. doi:10.1007/3-540-49381-6_45.

9 Khaled M. Elbassioni, Erik Krohn, Domagoj Matijevic, Julián Mestre, and Domagoj Severdija. Improved approximations for guarding 1.5-dimensional terrains. *Algorithmica*, 60(2):451–463, 2011. doi:10.1007/s00453-009-9358-4.

10 Stephan Friedrichs, Michael Hemmer, James King, and Christiane Schmidt. The continuous 1.5D terrain guarding problem: Discretization, optimal solutions, and PTAS. *J. Comput. Geom.*, 7(1):256–284, 2016. doi:10.20382/jocg.v7i1a13.

11 Subir Kumar Ghosh. On recognizing and characterizing visibility graphs of simple polygons. *Discrete & Computational Geometry*, 17(2):143–162, 1997. doi:10.1007/BF02770871.

12 Hannah Miller Hillberg, Erik Krohn, and Alex Pahlow. Guarding weakly-visible polygons with half-guards. *Iranian Conference on Computational Geometry*, 2022. URL: https://iccg2022.ce.sharif.edu/files/Papers/4-Guarding-Weakly-Visible-Polygons-with-Half-Guards.pdf.

13 James King and Erik Krohn. Terrain guarding is NP-hard. *SIAM J. Comput.*, 40(5):1316–1339, 2011. doi:10.1137/100791506.

14 Erik Krohn. Survey of terrain guarding and art gallery problems. Comprehensive Exam, University of Iowa, November 2007. URL: http://faculty.cs.uwosh.edu/faculty/krohn/papers/comprehensive.pdf.

15 Erik Krohn and Bengt J. Nilsson. Approximate guarding of monotone and rectilinear polygons. *Algorithmica*, 66(3):564–594, 2013. doi:10.1007/s00453-012-9653-3.

16 Alexander Kröller, Tobias Baumgartner, Sándor P. Fekete, and Christiane Schmidt. Exact solutions and bounds for general art gallery problems. *ACM J. Exp. Algorithmics*, 17, May 2012. doi:10.1145/2133803.2184449.

17 D. T. Lee and A. K. Lin. Computational complexity of art gallery problems. *IEEE Trans. Inform. Theory*, 32(2):276–282, March 1986.

18 David Lichtenstein. Planar formulae and their uses. *SIAM J. Comput.*, 11:329–343, 1982.

19 Nabil H. Mustafa, Rajiv Raman, and Saurabh Ray. QPTAS for geometric set-cover problems via optimal separators. *CoRR*, abs/1403.0835, 2014. arXiv:1403.0835.