Guarding Weakly-Visible Polygons with Half-Guards

Hannah Miller Hillberg* Erik Krohn† Alex Pahlow‡

Abstract

We consider a variant of the art gallery problem where all guards are limited to seeing to the right inside a weakly-visible polygon. Guards that can only see in one direction are called half-guards. In this paper, we give a polynomial time approximation scheme for vertex guarding a weakly-visible polygon with half-guards. We then show NP-hardness for vertex guarding a weakly-visible polygon with half-guards.

1 Introduction

An instance of the original *art gallery problem* takes as input a simple polygon P. A polygon P is defined by a set of points $V = \{v_1, v_2, \ldots, v_n\}$. There are edges connecting (v_i, v_{i+1}) where $i = 1, 2, \ldots, n-1$. There is also an edge connecting (v_1, v_n). If these edges do not intersect other than at adjacent points in V (or at v_1 and v_n), then P is called a simple polygon. The edges of a simple polygon give us two regions: inside the polygon and outside the polygon. For any two points $p, q \in P$, we say that p sees q if the line segment $\overline{pq}$ does not go outside of P. The art gallery problem seeks to find a guarding set of points $G \subseteq P$ such that every point $p \in P$ is seen by a point in G. In this paper, we study the vertex guarding problem which says that guards are only allowed to be placed at the vertices V. The optimization problem is defined as finding the smallest such G.

1.1 Previous Work

There are many results about guarding art galleries. Several results related to hardness and approximations can be found in [1, 5, 6, 10].

Additional Structure. Due to the inherent difficulty in fully understanding the art gallery problem for simple polygons, there has been some work done guarding polygons with additional structure, see [3, 8] for example. In this paper we consider *weakly-visible polygons* that we will describe below. Motivated by the fact that many cameras/sensors cannot sense in 360°, referred to as *full-guards* in this paper, we study guards that can sense in 180°, referred to as *half-guards*. We restrict the problem even further by only allowing these half-guards to see to the right. Even with these restrictions, the problem is difficult to solve.

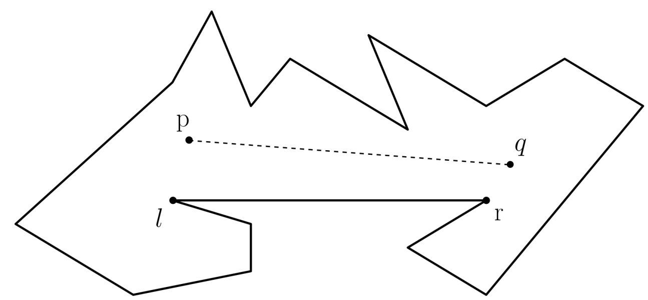

Figure 1: A WV-polygon where p sees q and every point in the polygon is seen by a point on $e = [l, r]$ or sees a point on e.

1.2 Definitions

A weakly-visible polygon (WV-polygon) P contains an edge $e = (l, r)$ such that every point in P sees at least one point on edge e. The usual definition of *sees* says that for any two points p and q inside of the polygon, if the line segment connecting p and q does not go outside of the polygon, then *p sees q*. Let $p.x$ be the x-coordinate for point p. In this paper, for a point p to see a point q, it must be the case that $p.x \leq q.x$, see Figure 1. The definition of a WV-polygon is slightly modified to say that it contains an edge $e = (l, r)$ such that every point in P sees (or is seen by) at least one point on edge e.

1.3 Our Contribution

NP-hardness has been shown for many variants of the art gallery problem. 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 [4, 9]. If the polygon is restricted to be a WV-polygon, restrict guards to be at the vertices and only allow them to see to the right, we show that even with these many restrictions, the problem is still NP-hard.

*Department of Computer Science, University of Wisconsin - Oshkosh, Oshkosh, Wisconsin, USA hillbergh@uwosh.edu

†Department of Computer Science, University of Wisconsin - Oshkosh, Oshkosh, Wisconsin, USA krohne@uwosh.edu
Supported by the Faculty Development Sabbatical Program at the University of Wisconsin - Oshkosh.
‡Department of Computer Science, University of Wisconsin - Oshkosh, Oshkosh, Wisconsin, USA pahloa45@uwosh.edu

Figure 2: A WV-polygon that requires $\Omega(n)$ half-guards that see to the right.

Ashur et al. [2] give a polynomial time approximation scheme (PTAS) for minimum dominating set in terrain-like graphs, which we will describe later. 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 does not imply that the visibility polygon of vertex guarding a WV-polygon *with half-guards* is terrain-like. We provide additional observations in this paper that show the visibility polygon is terrain-like. There are WV-polygons P that can be completely guarded with one full-guard but require $\Omega(n)$ half-guards considered in this paper, see Figure 2.

The remainder of the paper is organized as follows. Section 2 provides a PTAS for vertex guarding a WV-polygon using half-guards. Section 3 shows NP-hardness for vertex guarding a WV-polygon using half-guards. Finally, Section 4 gives a conclusion and future work.

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

In this section, we show that the visibility graph of a WV-polygon with half-guards is terrain-like. The visibility graph is a graph $G = (V, E)$ such that V corresponds to vertices in the WV-polygon and the edges E correspond to vertices that can be seen by other vertices. Then we use Theorem 1 from [2], see Appendix, to show that a PTAS exists for vertex guarding a WV-polygon with half-guards.

Using the definition from [2], a graph $G = (V, E)$ is *terrain-like* if one can assign a unique integer from the range $[1, |V|]$, where $|V|$ is the number of vertices in the polygon, to each vertex in V, such that, if both (i, k) and (j, l) are in E, for any $i < j < k < l$, then so is (i, l).

Much of the proof from [2] assumes that the WV-polygon lies above the WV-edge by placing guards at l and r and cutting off the portion of the polygon beneath the WV-edge, see Figure 3(left). When every vertex under consideration lies above this WV-edge, the 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 portion of the polygon cannot be assumed to be above the WV-edge. As shown in the shaded parts of Figure 3(right), the regions to the left of the placed guards are not seen. More so, a guard placed at l will

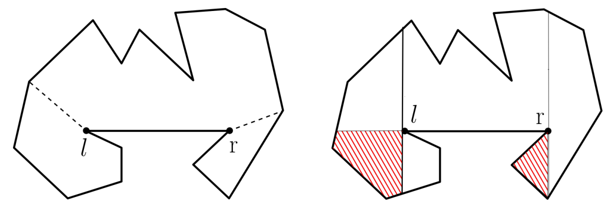

Figure 3: On the left, a full-guard placed at l and r cuts off the polygon below the WV-edge. On the right, the shaded regions are still unseen after half-guards are placed at l and r.

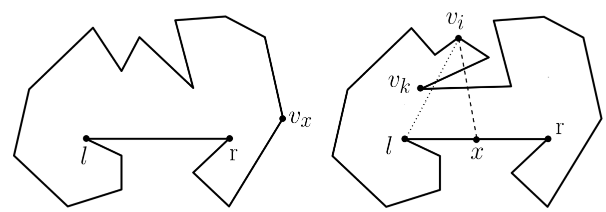

Figure 4: Example of v_x and Lemma 1.

not dominate a guard placed in the shaded region to the left of l. If the guards were full-guards, these regions would be seen and the l guard would dominate any optimal guard placed in this region. We show that even though these portions of the polygon are unseen and optimal guards can lie in these regions, the visibility graph connecting vertices to the guards that see them is still terrain-like.

2.1 Visibility Polygon is Terrain-Like

We will prove that the visibility polygon of the vertices is terrain-like by using the following modified order claim in WV-polygons that applies to full-guards as well as half-guards. Order the vertices walking in clockwise order starting from l: $(v_1 = l, v_2, v_2, \ldots, v_n = r)$.

Modified Order Claim: Assume that guards are placed at l and r and then consider the remaining unseen vertices. If 4 vertices are in order such that $a < b < c < d$, then if a sees c and b sees d, then a sees d.

As shown in [2], if all of a, b, c and d lie above the WV-edge, then the order claim holds and a sees d. If both a and d are below the WV-edge and both see l or both see r, then the same arguments from [2] hold for why a must see d.

The following lemmas are given for WV-polygons that apply to full-guards. For simplicity of the arguments, we assume that the WV-edge is parallel to the x-axis. Let v_x be the vertex such that every vertex $[v_{x+1}, v_{n-1}]$ is below the WV-edge. If no vertex meets this requirement, then $v_x = v_n = r$. In other words, v_x is the last vertex that is above the WV-edge when walking clockwise from l to r, see Figure 4(left).

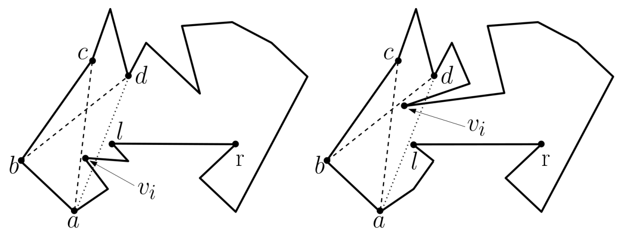

Figure 5: If a is blocked from d, then the blocker must be either l or r.

Lemma 1 *If $v_i \in [v_1, v_x]$, then no vertex $v_k \in [v_{i+1}, v_n]$ can block v_i from l.*

Proof. Assume that v_k blocks v_i from seeing l. If this happens, then v_k is to-the-left (when looking from l) of the $\overline{lv_i}$ line segment, see Figure 4(right). Let x be any point on the WV-edge that v_i sees. The $\overline{xv_i}$ line segment is to the right (when looking from l) of the $\overline{lv_i}$ line segment. This line segment is also blocked by v_k which means v_i does not see any point on the WV-edge. The polygon is not weakly-visible and we have a contradiction. □

Corollary 2 *Let v_w be the vertex such that every vertex $[v_2, v_{w-1}]$ is below the WV-edge. If no vertex meets this requirement, then $v_w = v_1 = l$. If $v_i \in [v_w, v_n]$, then no vertex $v_k \in [v_1, v_{i-1}]$ can block v_i from r.*

Lemma 3 *Consider 4 vertices in a WV-polygon such that $a < b < c < d$, a sees c, b sees d and a does not see d. If a is below the WV-edge and d is above the WV-edge, then a and d must both see l.*

Proof. Since a is below the WV-edge, a must see l. If a does not see l, then the polygon is not weakly-visible.

If d does not see l, then by Lemma 1, the v_i blocker for d must lie in the (l, d) range. If $v_i \in (b, d)$, then v_i would block b from d. If $v_i = (a, b]$, then a would not see c. It is not possible for $v_i = a$ since a is below the WV-edge. Lastly, if that vertex $v_i \in (l, a)$, then it would block a from seeing l, see Figure 5(left). Since there is no way to block d from l, d must see l. □

Corollary 4 *Consider 4 vertices in a WV-polygon such that $a < b < c < d$, a sees c, b sees d and a does not see d. If a is above the WV-edge and d is below the WV-edge, then a and d must both see r.*

This brings us to our final Lemma with full-guards:

Lemma 5 *If the order claim is broken, then a sees either l or r and also, d sees either l or r.*

Proof. If $a < b < c < d$, a sees c, b sees d, a does not see d and a and d are both below the WV-edge, then a sees l and d sees r. This, along with Lemma 3 and Corollary 4, cover the remaining cases. □

Returning to the discussion with respect to half-guards. Lemma 5 applies to full-guards. A modification of Lemma 5 is given to apply to half-guards:

Lemma 6 *If the order claim is broken, then at least one of a, b, c or d is seen by either l or r.*

The complete proof of Lemma 6 is given in the appendix. In short, if the order claim is broken, at least one of a, b, c or d must lie to the right of l or r and will necessarily be seen by one of l or r.

The analysis from [2] is now used to show a PTAS exists. First, one checks to see if the polygon can be guarded with a constant number of guards of some appropriate size. If an optimal guarding set of this size does not exist, then the first step of the algorithm is to place guards at l and r and remove the vertices that l or r see from consideration as they are already guarded. By Lemma 6, an order claim violation is not possible since at least one of the vertices involved in breaking the order claim has already been guarded by l or r and will not be in the modified problem instance. In other words, the visibility graph connecting vertices to the guards that see them are vertices that are not seen by l nor r. Since the order claim cannot be broken, when looking at the visibility graph of vertices connected to the leftmost and rightmost guards that see them, if $i < j < k < l$, $(i, k) \in E$ and $(j, l) \in E$, then $(i, l) \in E$. With this claim, the visibility graph for vertex guarding WV-polygons with half-guards that see to the right is terrain-like.

It should also be noted that the orientation of the polygon does not matter. For example, consider a polygon where the WV-edge is parallel to the y-axis and the "main" part of the polygon is to the right of the WV-edge, see Figure 6(left). In this instance, guards placed at l and r still cause the order claim to not be broken in the unguarded vertices that remain. Figure 6 shows polygons where the order claim will not be broken. No matter the orientation of the WV-edge, Lemma 6 holds for half-guards as well as full-guards.

Since the visibility graph is terrain-like, we use Theorem 1 from [2] to state the following:

Theorem 7 *There exists a PTAS for vertex guarding a weakly-visible polygon with half-guards where half-guards can only see to the right.*

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

In this section, we provide a sketch for showing that vertex guarding a WV-polygon with half-guards is NP-hard. NP-hardness for terrain guarding with full-guards was shown in [7], however, the entire terrain is not seen if guards are only allowed to look down. In the appendix,

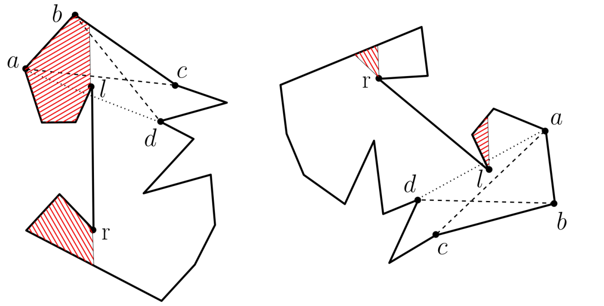

Figure 6: Shows the portion of the polygon "below" the WV-edge that is not seen by l or r.

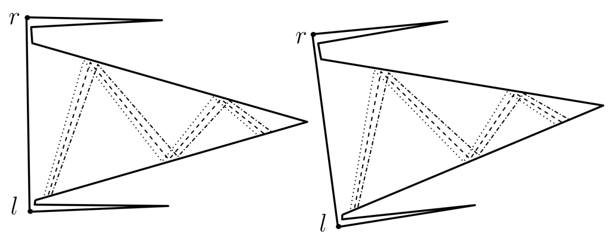

Figure 7: An overview of NP-hardness for WV-polygons.

we provide several new and updated gadgets to the reduction from [7] to show that vertex guarding a terrain with half-guards that only look down 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 7(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 when the WV-edge is rotated slightly counterclockwise, the reduction still holds, see Figure 7(right). The key visibilities from the guards remain and the polygon is still weakly-visible. The reduction begins to fail whenever a guard visibility from the original reduction starts to have a negative slope. If this happens, the guard no longer sees the distinguished point(s) to its right. To account for this, the original terrain is "stretched" such that none of the guard visibilities have a negative slope. Details of this stretching are in the appendix.

4 Conclusion and Future Work

In this paper, we present a PTAS for vertex guarding WV-polygons with half-guards that see to the right. This algorithm works regardless of the orientation of the WV-edge. 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] A. Aggarwal. *The art gallery theorem: its variations, applications and algorithmic aspects*. PhD thesis, The Johns Hopkins University, 1984.

[2] S. Ashur, O. Filtser, M. J. Katz, and R. Saban. Terrain-like graphs: Ptass for guarding weakly-visible polygons and terrains. *Computational Geometry*, 101:101832, 2022.

[3] P. Bhattacharya, S. Ghosh, and B. Roy. Approximability of guarding weak visibility polygons. *Discrete Applied Mathematics*, 228, 02 2017.

[4] M. C. Couto, P. J. de Rezende, and C. C. de Souza. An exact algorithm for minimizing vertex guards on art galleries. *Int. Trans. Oper. Res.*, 18(4):425–448, 2011.

[5] S. Eidenbenz. Inapproximability results for guarding polygons without holes. In *ISAAC*, pages 427–436, 1998.

[6] S. K. Ghosh. On recognizing and characterizing visibility graphs of simple polygons. *Discrete & Computational Geometry*, 17(2):143–162, 1997.

[7] J. King and E. Krohn. Terrain guarding is np-hard. *SIAM J. Comput.*, 40(5):1316–1339, 2011.

[8] E. Krohn and B. J. Nilsson. Approximate guarding of monotone and rectilinear polygons. *Algorithmica*, 66(3):564–594, 2013.

[9] A. Kröller, T. Baumgartner, S. P. Fekete, and C. Schmidt. Exact solutions and bounds for general art gallery problems. *ACM J. Exp. Algorithmics*, 17, may 2012.

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

[11] D. Lichtenstein. Planar formulae and their uses. *SIAM J. Comput.*, 11:329–343, 1982.

Appendix

Theorem 1 from [2]

Theorem *There exists a PTAS for (general) minimum dominating set in terrain-like graphs. That is, for any $\epsilon > 0$, there is a polynomial-time algorithm which, given a terrain-like graph $G = (V, E)$ and two sets $C, W \subseteq V$, returns $Q \subseteq C$ such that Q dominates W and $|Q| \leq (1+\epsilon) \cdot OPT$; here OPT is the size of a minimum subset of C that dominates W.*

Proof of Lemma 6

Proof. We break up this proof into several cases.

1. If a and d are both above the WV-edge, then the order claim cannot be broken as shown in [2].

2. If a is below the WV-edge and d is above the WV-edge, then by Lemma 3, a and d see l. It also must be the case that a and d are to the left of l (otherwise l would see one of them). Since a sees l, d sees l, a sees (or is seen by) c and b sees (or is seen by) d, the entire $\overline{ad}$ line segment is surrounded by visibility lines that cannot be pierced. Therefore, a must see (or be seen by) d, a contradiction that the order claim was broken. Since a and d cannot both be to the left of l, it must be that l sees one of them.

3. If a is above the WV-edge and d is below the WV-edge, then by Corollary 4, r sees (or is seen by) a and d. In order for r to not see a or d (and rather be seen by both a and d), both must be to the left of r. If b or c are below the WV-edge, then r will see them. Therefore, b and c must be above the WV-edge. Since d sees b, it must be the case that b is to the right of r. Since b is above and to the right of r, there must be a vertex that blocks r from seeing b. By Corollary 2, the blocker for r to b must be in the (b, r) range. If the blocker is in the (b, d) range, then d would not see b. If the blocker is in the (d, r) range, then d would not see r, a contradiction that the polygon is weakly-visible. It must be the case that r sees b.

4. If a and d are both below the WV-edge, a is to the left of l and d is to the left of r, then in order for d to see b, it must be the case that the $\overline{db}$ line segment goes below r forcing b to be to the right of r. Similar to case 3, r sees b. Although not necessary for the proof, using similar arguments, one can show that l sees c.

$\square$

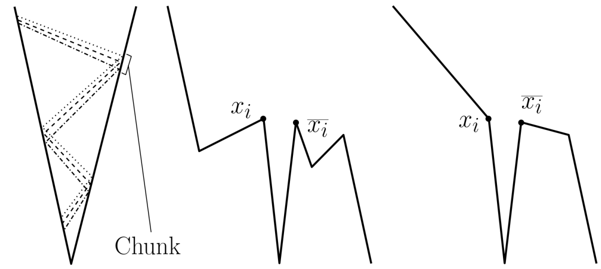

Figure 8: The left shows an overview of the NP-hardness reduction for terrain guarding. The middle is a variable gadget. The right is a starting gadget.

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." If we restrict guards to be half-guards that see "down" in the terrain, then the terrain guarding problem is still 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.

Sketch of Reduction

The terrain guarding reduction is from PLANAR 3SAT [11] 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 x_i and vertex representing $\overline{x_i}$, see Figure 8(middle). These variable gadgets are grouped together in chunks on the terrain. Figure 8(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 x_i vertex in the variable gadget of x_i in chunk C_j would require a guard to be placed at the x_i vertex in another variable gadget in a different chunk C_k, see Figure 10 for an example of such mirroring of data. There are points on the terrain that correspond to clauses of 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 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

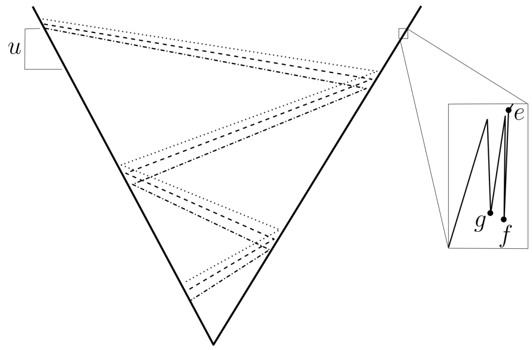

Figure 9: To ensure the entire terrain is seen, a guard is placed at e to see vertices f and g along with the u region.

terrain was seen, then the original instance was satisfied. If not, then the original instance was unsatisfiable.

The interested reader is encouraged to read [7] to see the full details of the original reduction. We modify the gadgets from [7] and show how these modified gadgets work for guards that can only see down.

Terrain Hardness Modifications

We describe the changes to the gadgets of [7] 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.

Between Chunks: Let us order the chunks from top to bottom $(C_1, C_2, \dots, C_m)$. 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 9. 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 only be seen from vertex e. Placing a guard at e will see f and g and also see the unseen region u below chunk C_i.

Variable Gadget: The variable gadgets, shown in Figure 10, do not need to be tweaked much to work. 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

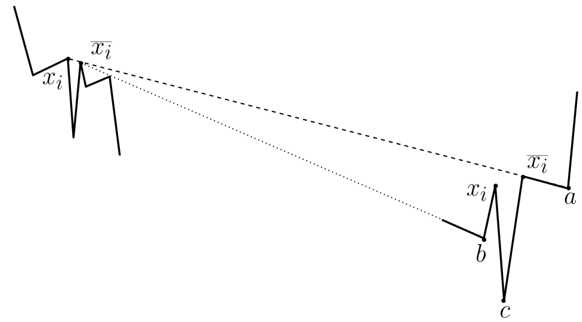

Figure 10: The variable gadget remains unchanged from [7].

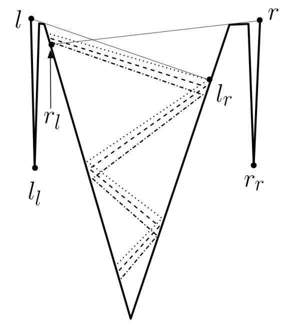

Figure 11: The overview of the terrain reduction where guards see down.

C_{j+1} may have been missed, see Figure 10. 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.

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 8(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 11, 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

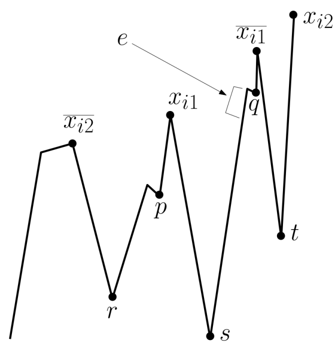

Figure 12: The inversion gadget for variable x_i in chunk C_{j+1}. This gadget is not tweaked but the variable gadget x_i in chunk C_j is modified slightly.

guard is needed.

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 9).

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 WLOG that C_1 is on the left side of the terrain. Two guards are placed at points l and r as shown in Figure 11. 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.

Clause Gadgets: No change is needed for the clause gadgets. All points in the clause gadgets are seen by the appropriate vertices above it.

Inversion Gadgets: A small update is needed for the inversion gadgets. 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. 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. Therefore, it can be tweaked to see the e line segment.

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 [7], if k guards can guard the entire terrain, the instance is satisfiable. If more than k are needed, then the instance is not satisfiable.

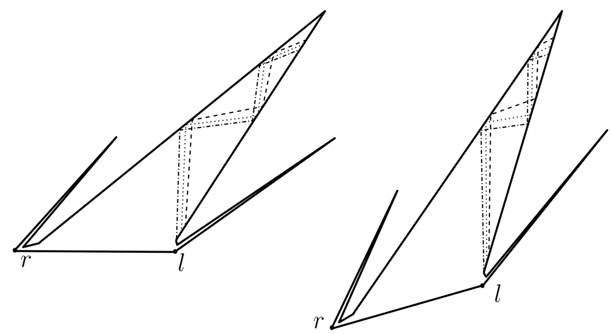

Figure 13: A stretched WV-polygon such that the NP-hardness result holds.

Theorem 8 *Finding the smallest vertex guard cover for guarding a terrain using half-guards that see down is NP-hard.*

Stretching WV-Polygon Hardness

An example of this stretching is seen in Figure 13(left). As seen in the example, the original bottom of the terrain 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.

This stretching of the polygon works even if the WV-edge has a positive slope, see Figure 13(right). 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.