Vertex Guarding in Weak Visibility Polygons

Pritam Bhattacharya†, Subir Kumar Ghosh & Bodhayan Roy

†Email: pritam.bhattacharya@cse.iitkgp.ernet.in

Abstract

The art gallery problem enquires about the least number of guards that are sufficient to ensure that an art gallery, represented by a polygon P, is fully guarded. In 1998, the problems of finding the minimum number of point guards, vertex guards, and edge guards required to guard P were shown to be APX-hard by Eidenbenz, Widmayer and Stamm. In 1987, Ghosh presented approximation algorithms for vertex guards and edge guards that achieved a ratio of $\mathcal{O}(\log n)$, which was improved to $\mathcal{O}(\log\log OPT)$ by King and Kirkpatrick in 2011. Ghosh also conjectured that constant-factor approximation algorithms exist for these problems. We settle the conjecture for the special class of polygons that are weakly visible from an edge and contain no holes by presenting a 6-approximation algorithm for finding the minimum number of vertex guards that runs in $\mathcal{O}(n^2)$ time. Contrastingly, for weak visibility polygons with holes, we present a reduction from Set Cover to show that there cannot exist a polynomial time algorithm for the vertex guard problem with an approximation ratio better than $\frac{1-\epsilon}{12}\ln n$ for any $\epsilon > 0$, unless NP = P .

Introduction

The art gallery problem enquires about the least number of guards that are sufficient to ensure that an art gallery (represented by a polygon P) is fully guarded, assuming that a guards field of view covers 360° as well as an unbounded distance. This problem was first posed by Victor Klee in 1973, and in the course of time, it has become one of the most well-studied problems in computational geometry.

A *polygon* P is defined to be a closed region in the plane bounded by a finite set of line segments, called edges of P, such that, between any two points of P, there exists a path which does not intersect any edge of P. If the boundary of a polygon P consists of two or more cycles, then P is called a *polygon with holes*. Otherwise, P is called a *simple polygon* or a *polygon without holes*. An art gallery can be viewed as an n-sided polygon P (with or without holes) and guards as points inside P. A point $z \in P$ is said to be *visible* from a guard g if the line segment zg does not intersect the exterior of P. In general, guards may be placed anywhere inside P. If the guards are allowed to be placed only on vertices of P, they are called *vertex guards*. If there is no such restriction, guards are called *point guards*. Point and vertex guards together are also referred to as *stationary guards*. If guards are allowed to patrol along a line segment inside P, they are called *mobile guards*. If they are allowed to patrol only along the edges of P, they are called *edge guards*. In 1975, Chvtal showed that $\lfloor \frac{n}{3} \rfloor$ stationary guards are sufficient and sometimes necessary for guarding a simple polygon.

The decision version of the art gallery problem is to determine, given a polygon P and a number k as input, whether the polygon P can be guarded with k or fewer guards. The problem was first proved to be NP-complete for polygons with holes by ORourke and Supowit. For simple polygons, it was proved to be NP-complete for vertex guards by Lee and Lin, and their proof was generalized to work for point guards by Aggarwal. The problem was shown to be NP-hard even for simple orthogonal polygons.

In 1987, Ghosh provided an $\mathcal{O}(\log n)$-approximation algorithm for the case of vertex and edge guards by discretizing the input polygon and treating it as an instance of Set Cover. For guarding simple polygons using vertex guards and perimeter guards, King and Kirkpatrick obtained an approximation ratio of $\mathcal{O}(\log\log OPT)$ in 2011. In 1998, Eidenbenz, Stamm and Widmayer proved that the problem is APX-complete, implying that an approximation ratio better than a fixed constant cannot be achieved unless P=NP. They also proved that if the input polygon is allowed to contain holes, then there cannot exist a polynomial time algorithm for the problem with an approximation ratio better than $((1-\epsilon)/12)\ln n$ for any $\epsilon > 0$, unless NP $\subseteq$ TIME($n^{\mathcal{O}(\log\log n)}$). Contrastingly, in the case of simple polygons without holes, the existence of a constant-factor approximation algorithm for vertex guards and edge guards has been conjectured by Ghosh since 1987. However, this conjecture has not been settled till date.

Our Contributions

A polygon P is said to be a *weak visibility polygon* if every point in P is visible from some point of an edge. We present a 6-approximation algorithm, having a running time of $\mathcal{O}(n^2)$, for vertex guarding polygons that are weakly visible from an edge and contain no holes. This result settles Ghosh's conjecture for a special class of polygons. By presenting a reduction from Set Cover, we also show that, for the special class of polygons containing holes that are weakly visible from an edge, there cannot exist a polynomial time algorithm for the vertex guard problem with an approximation ratio better than $((1-\epsilon)/12)\ln n$ for any $\epsilon > 0$, unless NP = P.

Placement of vertex guards in weak visibility polygons

Let P be a simple polygon. If there exists an edge uv in P (where u is the next clockwise vertex of v) such that P is weakly visible from uv, then it can be located in $\mathcal{O}(n^2)$ time. Henceforth, we assume that such an edge uv has been located. Let $bd_c(p,q)$ (or, $bd_{cc}(p,q)$) denote the clockwise (respectively, counterclockwise) boundary of P from a vertex p to another vertex q. Note that, by definition, $bd_c(p,q) = bd_{cc}(q,p)$. The *visibility polygon* of P from a point z, denoted by $VP(z)$, is defined to be the set of all points in P that are visible from z, i.e. $VP(z) = \{q \in P : q \text{ is visible from } z\}$.

The *shortest path tree* of P rooted at any point s of P, denoted by $SPT(s)$, is the union of Euclidean shortest paths from s to all the vertices of P (see Figure 1). This union of paths is a planar tree, rooted at r, which has n nodes, namely the vertices of P. For every vertex x of P, let $p_u(x)$ and $p_v(x)$ denote the parent of x in $SPT(u)$ and $SPT(v)$ respectively.

Suppose a guard is placed on every non-leaf vertex of $SPT(u)$ and $SPT(v)$. It is obvious that these guards see all points of P. However, the number of guards required may be very large compared to the size of an optimal guarding set. In order to reduce the number of guards, placing guards on every non-leaf vertex should be avoided. Let A be a subset of vertices of P. Let S_A denote the set which consists of the parents $p_u(z)$ and $p_v(z)$ of every vertex $z \in A$. Then, A should be chosen such that all vertices of P are visible from guards placed at vertices of S_A. We present below a naive algorithm for choosing A and S_A.

Algorithm 1 An $\mathcal{O}(n^2)$-algorithm for computing a guard set S_A for all vertices of P

```
Compute SPT(u) and SPT(v)
Initialize all the vertices of P as unmarked
Initialize A ← ∅, S_A ← ∅ and z ← u
while z ≠ v do
    z ← the vertex next to z clockwise on bd_c(u, v)
    if z is unmarked then
        A ← A ∪ {z} and S_A ← S_A ∪ {p_u(z), p_v(z)}
        Place guards on p_u(z) and p_v(z)
        Mark vertices of P visible from p_u(z) or p_v(z)
    end if
end while
return the guard set S_A
```

Lemma 1. *Let S_{opt} denote an optimal set of vertex guards. If, for all $z \in A$, every vertex of $bd_c(p_u(z), p_v(z))$ is visible from $p_u(z)$ or $p_v(z)$, then $|A| \leq |S_{opt}|$.*

Theorem 2. *If every vertex $z \in A$ is such that all vertices of $bd_c(p_u(z), p_v(z))$ are visible from $p_u(z)$ or $p_v(z)$, then $|S_A| \leq 2|S_{opt}|$.*

Proof. It is easy to see that $|S_A| = 2|A|$. By Lemma 1, $|A| \leq |S_{opt}|$. So, $|S_A| = 2|A| \leq 2|S_{opt}|$. □

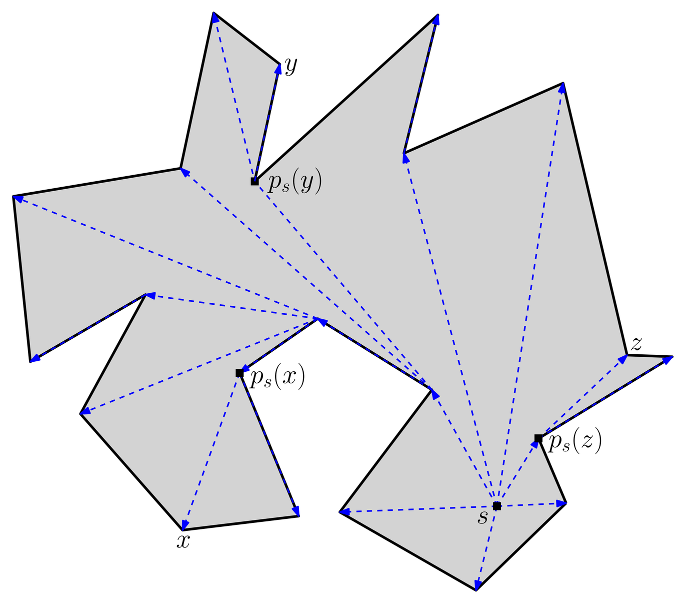

Figure 1: Euclidean shortest path tree rooted at s. The parents of vertices x, y and z in $SPT(s)$ are marked as $p_s(x)$, $p_s(y)$ and $p_s(z)$ respectively.

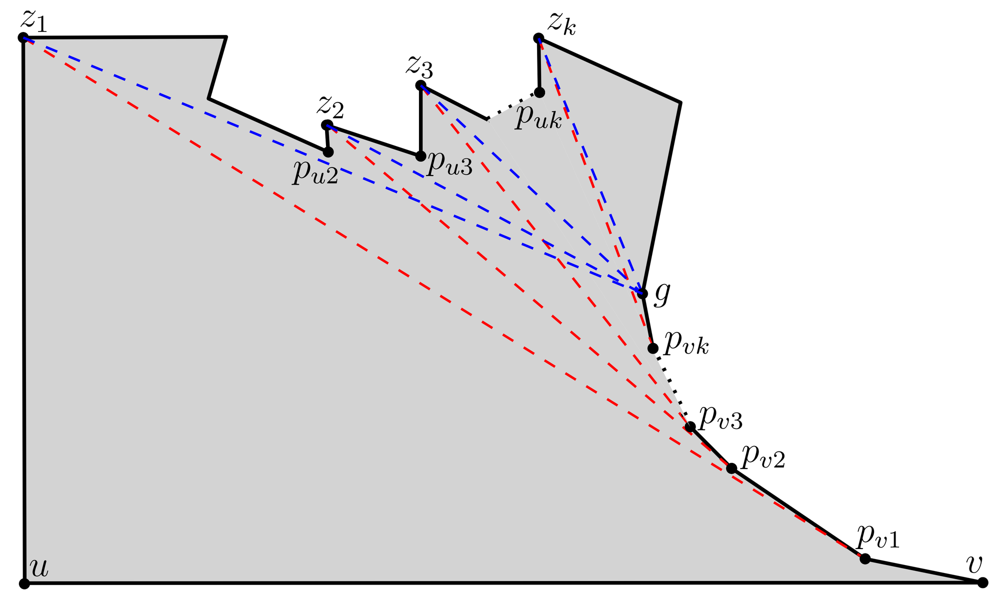

Figure 2: Instance where guard set S_A computed by Algorithm 1 is arbitrarily large compared to S_{opt}.

However, the above bound does not hold if there exists $z \in A$ such that some vertices of $bd_c(p_u(z), p_v(z))$ are not visible from $p_u(z)$ or $p_v(z)$. Consider Figure 2. For each $i \in \{1, 2, \ldots, k-1\}$, z_{i+1} is not visible from $p_u(z_i)$ or $p_v(z_i)$, which forces Algorithm 1 to place guards at $p_u(z_{i+1})$ and $p_v(z_{i+1})$. Therefore, Algorithm 1 includes $z_1, z_2, z_3, \ldots, z_k$ in A and ends up placing a total of $2k$ guards at vertices $u, p_{v1}, p_{u2}, p_{v2}, \ldots, p_{uk}, p_{vk}$. However, all vertices of P are visible from just two guards placed at u and g. Hence, $|S_A| = 2k$ whereas $|S_{opt}| = 2$. Since the construction in Figure 2 can be extended for any arbitrary integer k, $|S_A|$ can be arbitrarily large compared to $|S_{opt}|$. So we now present a new algorithm which gives us a 4-approximation.

In the new algorithm, described in pseudocode as Algorithm 2, $bd_c(u,v)$ is scanned to identify a set of unmarked vertices, denoted as B, such that all vertices of P are visible from guards in $S_B = \{p_u(z) | z \in B\} \cup \{p_v(z) | z \in B\}$. However, unlike the previous algorithm (see Algorithm 1), the new algorithm does not blindly include in B every next unmarked vertex that it encounters during the scan. During the scan, if z denotes the current unmarked vertex under consideration, then it may either choose to include z in B or skip ahead to the next unmarked vertex along the scan depending on certain properties of z.

Algorithm 2 An $\mathcal{O}(n^2)$-algorithm for computing a guard set S for all vertices of P

```
Compute SPT(u) and SPT(v)
Initialize all the vertices of P as unmarked
Initialize B ← ∅, S_B ← ∅ and z ← u
while there exists an unmarked vertex in P do
    z ← the first unmarked vertex on bd_c(u, v) from z
    if every unmarked vertex of bd_c(z, p_v(z)) is visible from p_u(z) or p_v(z) then
        B ← B ∪ {z} & S_B ← S_B ∪ {p_u(z), p_v(z)}
        Mark all vertices visible from p_u(z) or p_v(z)
        z ← p_v(z)
    else
        z' ← the first unmarked vertex on bd_c(z, v)
        while every unmarked vertex of bd_c(p_u(z'), z') is visible from p_u(z') or p_v(z') do
            z ← z'
            z' ← the first unmarked vertex on bd_c(z', v)
        end while
        B ← B ∪ {z} & S_B ← S_B ∪ {p_u(z), p_v(z)}
        Mark all vertices visible from p_u(z) or p_v(z)
        y ← z
        while ∃ an unmarked vertex on bd_c(u, z) do
            y ← first unmarked vertex on bd_cc(p_u(y), u)
            B ← B ∪ {y} & S_B ← S_B ∪ {p_u(y), p_v(y)}
            Mark all vertices visible from p_u(y) or p_v(y)
        end while
    end if
end while
return the guard set S_B
```

Lemma 3. *If S_{opt} denotes an optimal set, $|B| \leq 2|S_{opt}|$.*

Theorem 4. $|S_B| \leq 4|S_{opt}|$.

Proof. It is easy to see that $|S_B| = 2|B|$. By Lemma 3, $|B| \leq 2|S_{opt}|$. So, $|S_B| = 2|B| \leq 4|S_{opt}|$. □

It is not guaranteed that all interior points of P are visible from guards in S_B. Consider the scenario shown in Figure 3. While scanning $bd_c(u,v)$, Algorithm 2 places guards at $p_u(z)$ and $p_v(z)$ as all vertices of $bd_c(p_u(z), p_v(z))$ become visible from $p_u(z)$ or $p_v(z)$. Observe that in fact all vertices of P become visible from these two guards. But, $VP(p_u(z))$ has several left pockets and $VP(p_v(z))$ has several right pockets which intersect pairwise to create multiple invisible cells. In order to guard these invisible cells, a set S' of additional guards need to be placed.

Theorem 5. *There exists an algorithm with running time $\mathcal{O}(n^2)$ that returns a guard set S' for guarding all interior points of P such that $|S'| \leq 6|S_{opt}|$.*

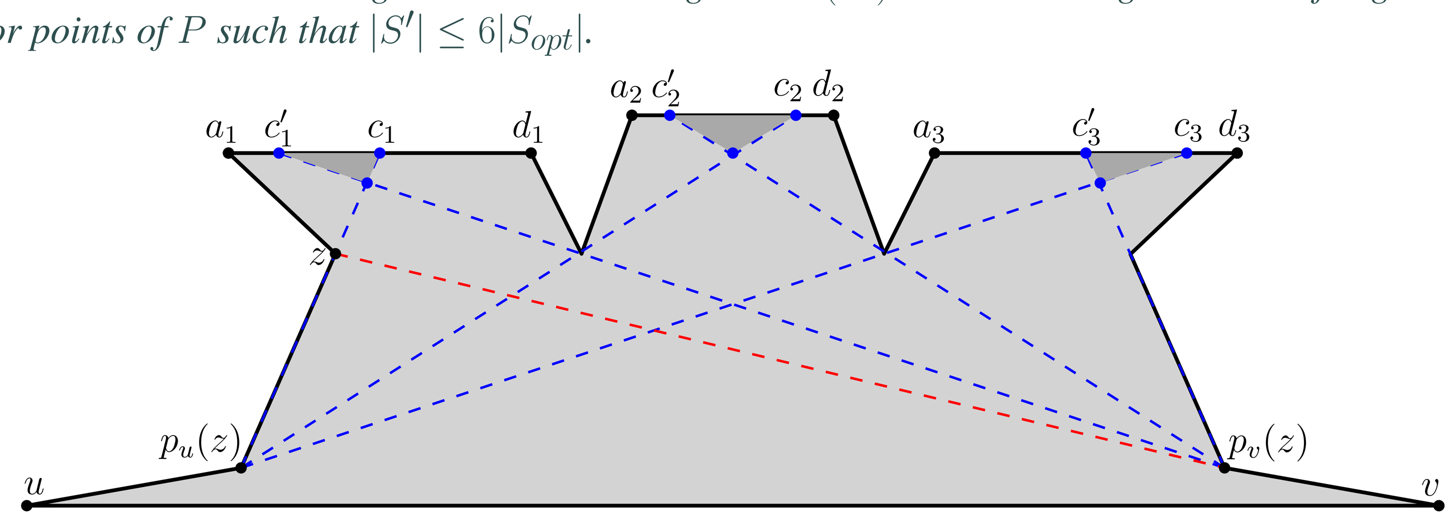

Figure 3: Multiple invisible cells exist that are not visible from the guards placed at $p_u(z)$ and $p_v(z)$.

An Inapproximability Result

For polygons with holes, Eidenbenz, Stamm and Widmayer proved in 1998 that there cannot exist a polynomial time algorithm for the art gallery problem with an approximation ratio better than $((1-\epsilon)/12)\ln n$ for any $\epsilon > 0$, unless NP $\subseteq$ TIME($n^{\mathcal{O}(\log\log n)}$). Modifying their technique, and taking into consideration a recent result by Dinur and Steurer, we obtained the following result.

Theorem 6. *For weak visibility polygons with holes, there cannot exist a polynomial time algorithm for Vertex Guard with an approximation ratio better than $((1-\epsilon)/12)\ln n$ for every $\epsilon > 0$, unless NP = P.*

Acknowledgements

The first author would like to acknowledge TCS for their financial support since July 1, 2015.