January 2001

VC-Dimension of Exterior Visibility of Polyhedra

Volkan Isler
University of Pennsylvania

Sampath Kannan
University of Pennsylvania, kannan@cis.upenn.edu

Kostas Daniilidis
University of Pennsylvania, kostas@cis.upenn.edu

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-01-34.

VC-Dimension of Exterior Visibility of Polyhedra

Abstract
In this paper, we address the problem of finding the minimal number of viewpoints outside a polyhedron in two or three dimensions such that every point on the exterior of the polyhedron is visible from at least one of the chosen viewpoints. This problem which we call the *minimum fortress guard problem (MFGP)* is the optimization version of a variant of the art-gallery problem (sometimes called the fortress problem with point guards) and has practical importance in surveillance and image-based rendering. Solutions in the vision and graphics literature are based on image quality constraints and are not concerned with the number of viewpoints needed. The corresponding question for art galleries (minimum number of viewpoints in the interior of a polygon to see the interior of the polygon) which we call the *minimum art-gallery guard problem (MAGP)* has been shown to be NP-complete. A simple reduction from this problem shows the NP-completeness of MFGP. Instead of relying on heuristic searches, we address the approximability of the camera placement problem. It is well known (and easy to see) that this problem can be cast as a hitting set problem. While the approximability of generic instances of the hitting set problem is well understood, Brönnimann and Goodrich[3] presented improved approximation algorithms for the problem in the case that the input instances have bounded Vapnik-Chervonenkis (VC) dimension.

In this paper we explore the VC-dimension of set systems associated with the camera placement problem described above. We show a constant bound for the VC dimension in the two dimensional case but a tight logarithmic bound in the three dimensional case. In the two dimensional case we are also able to present an algorithm that uses at most one more viewpoint than the optimal in the case that the viewpoints are restricted to be on a circumscribing circle - a restriction that is justified in practice.

Comments
University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-01-34.

VC-Dimension of Exterior Visibility of Polyhedra

Technical Report MS-CIS-01-34
University of Pennsylvania

Volkan Isler
Department of Computer and Information Science
University of Pennsylvania
isleri@grasp.cis.upenn.edu

Sampath Kannan
Department of Computer and Information Science
University of Pennsylvania
kannan@cis.upenn.edu

Kostas Daniilidis
Department of Computer and Information Science
University of Pennsylvania
kostas@cis.upenn.edu

ABSTRACT

In this paper, we address the problem of finding the minimal number of viewpoints outside a polyhedron in two or three dimensions such that every point on the exterior of the polyhedron is visible from at least one of the chosen viewpoints. This problem which we call the *minimum fortress guard problem (MFGP)* is the optimization version of a variant of the art-gallery problem (sometimes called the fortress problem with point guards) and has practical importance in surveillance and image-based rendering. Solutions in the vision and graphics literature are based on image quality constraints and are not concerned with the number of viewpoints needed. The corresponding question for art galleries (minimum number of viewpoints in the interior of a polygon to see the interior of the polygon) which we call the *minimum art-gallery guard problem (MAGP)* has been shown to be NP-complete. A simple reduction from this problem shows the NP-completeness of MFGP. Instead of relying on heuristic searches, we address the approximability of the camera placement problem. It is well known (and easy to see) that this problem can be cast as a hitting set problem. While the approximability of generic instances of the hitting set problem is well understood, Brönnimann and Goodrich[3] presented improved approximation algorithms for the problem in the case that the input instances have bounded Vapnik-Chervonenkis (VC) dimension.

In this paper we explore the VC-dimension of set systems associated with the camera placement problem described above. We show a constant bound for the VC dimension in the two dimensional case but a tight logarithmic bound in the three dimensional case. In the two dimensional case we are also able to present an algorithm that uses at most one more viewpoint than the optimal in the case that the viewpoints are restricted to be on a circumscribing circle — a restriction that is justified in practice.

1. INTRODUCTION

The problem of placing guards in the exterior of a polygon so as to see the entire boundary of the polygon has been called the fortress problem [11]. An important question is to determine the minimum number of guards needed in a given instance of the fortress problem.

Because of its close relationship to the art gallery problem this question can be shown to be NP-hard. Indeed, this is shown by transforming instances of MAGP to instances of MFGP by simply bounding the polygon by a bounding box and then "eviscerating" the polygon, so that what was the interior of the polygon becomes part of the exterior, and so that the boundary of the polygon is not visible from the rest of the exterior. A more precise description of this reduction is given in Section 2.3.

In this paper we examine the problem of finding an approximation algorithm for this problem in two and three dimensions. Bounds on approximation ratios have been derived for several variations of the art-gallery problem in the past. For a survey see [4].

In an instance of the minimum fortress guard problem (MFGP) in two dimensions (resp. 3D) we are given a polygon P (resp. polyhedron P) and a set of possible viewpoint (or camera) positions. Some possible sets of camera positions we consider are the following: (1) cameras permitted anywhere in the exterior of P; (2) cameras restricted to be on a circumscribing circle (resp. sphere) around P; (3) cameras restricted to be outside the convex hull of P. In all cases we assume that the boundary of P is visible from at least one of the allowable camera positions since otherwise the problem is insoluble.

In the decision version of MFGP we are also given an integer k and asked whether it is possible to locate k cameras so that they can see all of the exterior of P. In the optimization version we want to find the minimum number (and location) of cameras needed to see the exterior of P.

The decision version of the well-known *hitting set problem* is the following: We are given a set X and a collection of sets $\mathcal{R}$ where each $R \in \mathcal{R}$ is a subset of X. We are also given a number k. The question is whether there is a subset

$H \subset X$ such that $|H| \leq k$ and for each $R \in \mathcal{R}$ $R \cap H \neq \emptyset$. This problem is known to be NP-complete. The hitting set problem is the dual of the even better known *set cover problem* which is also NP-complete. Under the assumption that P $\neq$ NP it is known that both hitting set and set cover can be approximated to within a log factor of the maximum set sizes (in either the primal or the dual system) and not much better [15].

MFGP can be seen to be a particular case of the hitting set problem. The set X is the set of possible camera locations. For each point p on the boundary of the polyhedron P, there is a set R_p consisting of all camera locations that can see p. The hitting set problem assumes a finite set X and we have to implicitly deal with this issue when we attempt to pose MFGP as such a problem.

As stated earlier the general hitting set problem cannot be approximated to better than a log factor. However [3] shows that if the VC dimension can be bounded by d and the optimal hitting set has size c, then we can produce an $O(d \log cd)$ approximation. Thus we need to examine the VC-dimension of hitting set instances that can be produced from MFGP.

We are able to determine the VC-dimension both in 2D and in 3D. Surprisingly while in 2D the VC-dimension (for all three restrictions on camera placements) is bounded by a constant, in 3D the VC-dimension is $\Theta(\log n)$ where n is the number of vertices in the input polyhedron to MFGP. This means that the algorithm in [3] does not provide an improved approximation over the greedy algorithm in 3D. On the other hand, in 2D, we are able to get much better approximations in the case where camera placements are restricted to a circumscribing circle. In this case, we produce a solution that uses at most one more camera than the optimal solution!

The particular scenarios we are addressing are surveillance, object inspection, and image based rendering. In the case of surveillance, we need a complete coverage at any time so that no event will be missed. This is the reason why coverage with one mobile guard is not applicable. In case of object inspection, we know the prior geometry of an object, and we need the minimal number of views so that the object will be checked regarding defects. In this scenario, the object might be placed on a turntable and we ask then for the minimal number of rotations. The objects might be medical organs which have to be imaged from very few viewpoints of an endoscope guided by a robot manipulator. In the case of image based rendering, we have a prior map of the environment but we need to obtain a detailed reconstruction with a range sensor or we need just the appearance of the environment for visualization. This is very important for telepresence and immersive environments: After the environment has been captured and an immersed user changes her viewpoint any hole would cause a break in the sense of presence.

2. PRELIMINARIES

A set system is a pair $(X, \mathcal{R})$ where X is a set and $\mathcal{R}$ is a collection of subsets $R \subseteq X$. Given a set system, the minimum set cover problem asks for a minimum cardinality set $\mathcal{S} \subseteq \mathcal{R}$ such that $\bigcup_{R \in \mathcal{S}} R = X$.

As is well known the greedy algorithm finds an approximate solution to set cover with approximation factor equal to the log of the maximum cardinality of a set in $\mathcal{R}$. Dually, an algorithm due to Hochbaum[8] achieves an approximation factor which is the logarithm of the maximum number of sets that any element in X occurs in. These are essentially the best achievable approximations under reasonable complexity theoretic assumptions.

Given a set system $(X, \mathcal{R})$ one can define a dual set system $(Y, \mathcal{S})$ where $Y = \{R | R \in \mathcal{R}\}$ and $\mathcal{S}$ consists of a set S_x for each $x \in X$ where $S_x = \{R | x \in R\}$. The set cover problem in $(Y, \mathcal{S})$ is the *hitting set* problem in $(X, \mathcal{R})$. The hitting set problem seeks the smallest subset H of X such that every set $R \in \mathcal{R}$ contains an element from H. From the symmetry of the approximation bounds given above, it is clear that the hitting set problem and the set cover problem are approximable to the same factor.

The camera placement problem has several variations but the general form of the problem is the following: We are given an object P and a set of allowable camera positions $\mathcal{C}$. We would like to choose a subset of camera positions such that every point in P is visible from one of the chosen camera positions. Except for the possibility that infinite set systems arise in this formulation, it is easy to see that the camera placement problem can be modeled as a set cover problem. Each camera position represents a subset of P that is covered and we want to pick as few camera positions as possible to cover all of P. Dually, this problem can also be viewed as a hitting set problem. each point $p \in P$ defines a set of camera positions C_p that can see this point. The task is to pick a set of camera positions that hits the sets C_p $\forall p$.

In this paper the object P will always be the boundary of a polygon in two dimensions or a polyhedron in three dimensions. The set of allowable camera positions will be either a circumscribing circle/sphere around P or the entire exterior outside of the convex hull of P. Note that unless arguments can be made that only finitely many points on P need to be considered and that only finitely many camera positions need to be considered, the set systems that arise are typically infinite set systems.

Although one cannot expect a better approximation factor than $O(log\ n)$ for the general set-cover problem[15, 5], the camera placement problem is far from being general due to the underlying geometric constraints. One such constraint is the Vapnik-Chervonenkis(VC) Dimension[17].

2.1 VC-Dimension and Camera Placement

The concept of Vapnik-Chervonenkis Dimension, VC Dimension for short, was introduced by Vapnik and Chervonenkis in [17] and has applications in statistics, learning theory, computational geometry and complexity theory (see [10] for references). The VC-Dimension of a set system is defined as follows:

Definition 2.1. *Given a set system $(X, \mathcal{R})$, let A be a subset of X. We say A is shattered by $\mathcal{R}$ if $\forall Y \subseteq A$ $\exists R \in \mathcal{R}$ such that $R \cap A = Y$. The VC-dimension of $(X, \mathcal{R})$ is the* *cardinality of the largest set that can be shattered by $\mathcal{R}$.*

The VC-Dimension of a set system reveals a lot of information about properties of the set system. For example, if the set system $(X, \mathcal{R})$ has a constant VC-Dimension d, a small number, $O(\frac{d}{\epsilon} log \frac{1}{\epsilon})$, points sampled from X intersects all the subsets in $\mathcal{R}$ whose sizes are greater than $\epsilon \cdot |X|$ with high probability. Another useful property is that if $(X, \mathcal{R})$ has a constant VC-Dimension d, then the number of subsets in $\mathcal{R}$ is bounded by n^d as opposed to 2^n where $n = |X|$. A nice presentation of these results can be found in [9].

Perhaps the most crucial result for constant VC-Dimension systems from a camera placement perspective is the algorithm presented by Bronnimann and Goodrich in [3], that returns $O(log\ d)$ solutions to the set-cover of systems with bounded VC-Dimension where d is the optimal set-cover for the system. This is a significant improvement on the previous *log n*-approximation, when n is large but the optimal is small.

The VC-Dimension of set systems plays a very important role in randomized and geometric algorithms. The reader is referred to the surveys [10, 1] for further information.

In this paper we will address the problem of covering the boundary of a polygon P with as few cameras as possible. An *instance* of the camera placement problem is: Given P and a specification of possible camera locations find a minimum set cover of the system $(P, \{V(c_i)\})$, where $V(c_i)$ is the set of points visible on P from camera c_i and the index i varies over all possible camera locations. The definition of $V(c_i)$ can capture any optical constraints on what a camera can see. We will refer to the specification of possible camera locations as a *setup*. We say a set S of cameras *cover* P if $\bigcup_{c_i \in S} V(c_i) = P$.

Throughout this paper we will represent cameras with their projection centers c_i and say that c_i sees the point $p \in P$ if the only intersection of the line segment $\overline{pc_i}$ with P is p. We extend the notion of visibility to sets as follows: We say that a camera sees a set of points ω if it can see all the points in ω. The following notation will be useful for VC-Dimension proofs. Let $P_m = \{p_1, \ldots, p_m\}$ be m points. We say that camera c sees the subset $\omega \subseteq P_m$ if c can see all points in ω but no point in $P_m \setminus \omega$.

By the *VC-Dimension of a setup*, we will refer to the VC-Dimension of the maximum number of points that can be shattered over all instances of the camera placement problem for a specific setup. For example, if there are no restrictions on cameras and we want to cover simple polygons, we would like to find the VC-Dimension of the set system $(P, \{V(c_i)\})$ as P varies over the set of all simple polygons. Therefore, in order to give a lower bound m on the VC-Dimension of a setup it suffices to present one instance where m points are shattered, but for an upper bound one needs to show that there exists no instance such that m points can be shattered.

2.2 Previous work on visibility and VC-Dimension

There is quite a rich literature on the Art Gallery Theorem which states: $\lfloor \frac{n}{3} \rfloor$ guards are occasionally necessary and always sufficient to cover a polygon of n vertices[11, 14, 12]. A different version of the art gallery problem, known as Minimal Art Gallery Guarding Problem (MAGP) is: *Given a particular polygon, what is the minimum number of guards necessary to cover the polygon?* It has been shown that MAGP is NP-hard. All these results can be found in [12].

Approximation algorithms for Minimum Guard Coverage have been considered [6, 4, 7] for different versions of the problem, however there is still a gap between the inapproximability results and existing algorithms. For a survey of the approximation results see [4].

The VC-Dimension of 2D visibility systems have also appeared in the literature. For example in [16], Valtr proved that the VC-Dimension of the system $(P, \{V(x) \mid \forall x \in P\})$, where P is a simple polygon and $V(x)$ is the visibility polygon of point x in P, is somewhere between 6 and 23. He also established a $O(log(h))$ bound for polygons with holes where h is the number of holes. Recently, in [7] Banos et. al. showed a loose bound of the $log(n + h)$ bound for the dual of a system similar to the one considered in [16] where n is the number of vertices of the polygon and h is the number of holes. However, to the best of our knowledge, there are no results for the VC-Dimension of visibility systems in 3D.

2.3 NP-completeness of MFGP

The NP-completeness of MFGP (minimal fortress guard problem) follows immediately by a reduction from MAGP by an "evisceration" technique whereby the interior of a polygon is more or less transformed into its exterior. The following figure illustrates the transformation.

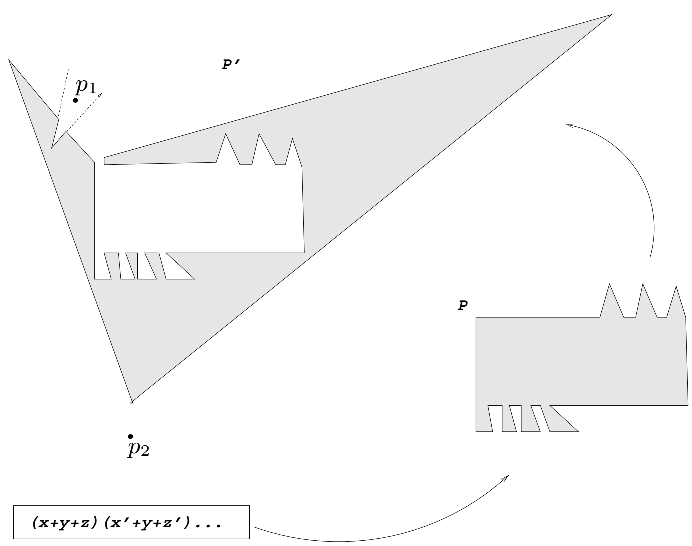

Figure 1: MFGP is NP-hard

Instances of MAGP that arise in the reduction from 3-SAT that proves the NP-completeness of MAGP have the form of the polygon P. Clearly if it takes k guards to guard the interior of P, it will take exactly $k + 2$ exterior guards to guard the exterior of the eviscerated polygon. This proves the NP-completeness of MFGP.

3. 3D SETTING

In this section, we consider the following setup which arises in typical tele-immersive applications:

DEFINITION 3.1. *We define* $3DSPHERE$ *as a setup where we are given a polyhedron* $\mathcal{P}$*, and a viewing sphere* $\mathcal{S}$ *such that* $\mathcal{P}$ *is totally contained in* $\mathcal{S}$*.*

We show that even under these restrictions there are polyhedra with n vertices such that $\Theta(log\ n)$ points can be shattered from the viewing sphere. Namely, we prove the following theorem:

THEOREM 3.2. *The VC-Dimension of 3DSPHERE is* $\Theta(log\ n)$ *where* n *is the number of vertices of the polyhedron* $\mathcal{P}$ *for the set system* $(\mathcal{P}, \{V(c_i)\})$ *such that the centers of cameras* c_i *are restricted to lie on a viewing sphere* $\mathcal{S}$ *that contains* $\mathcal{P}$*.*

In the next two subsections, we present the upper and lower bounds for the VC-Dimension of 3DSPHERE, in lemmata 3.3 and 3.4 respectively.

The implication of theorem 3.2 is that it is not likely that the algorithm in [3] helps for the camera placement problem.

3.1 Upper Bound

In this section we present an upper bound on the VC-Dimenion of 3DSPHERE.

LEMMA 3.3. *Let* d *be the VC-Dimension of 3DSPHERE.* $d = O(log\ n)$ *where* n *is the number of vertices of the polyhedron we want to cover.*

PROOF. In [13], Platinga and Dyer define *aspects* as changes in the topology of the image of a polyhedron. After presenting a catalogue of events that can change the aspects, they construct the viewing space partition, VSP, which is a partition of the viewpoint space into maximal regions of constant aspect and they present tight bounds for the number of regions in VSP. They show that the size of the VSP for a general(i.e. non-convex) polyhedron under orthographic projection is $\Theta(n^6)$ and their model for the orthographic projection is exactly the same as 3DSPHERE with $\mathcal{S}$ at infinity.

Let $P_m = \{p_1, \ldots, p_m\}$ be any m points to be shattered on a polyhedron. If we define an aspect as appereance/disappereance of $p_i, i = 1, \ldots, m$, and restrict the camera locations to a sphere that contains the polyhedron, we can use the catalogue of events in [13] to show that the size of the VSP for this new notion of aspects is still $\Theta(n^6)$. However, in order to shatter m points, one needs 2^m distinct partitions. Since we must have $n^6 \geq 2^m$, we have $m = O(log\ n)$ which gives us the desired upper bound. □

3.2 Lower Bound

In this section we show that the upper bound $log\ n$ on the VC-Dimension of 3DSPHERE is indeed tight, our main result as stated in the following lemma for theorem 3.2.

LEMMA 3.4. *For any given* m*, there exists a polyhedron* $\mathcal{P}$ *with* $\Theta(2^m)$ *vertices that is contained in a cone with base* B *and height proportional to* m *such that there are* m *marked points on* $\mathcal{P}$ *that can be shattered from* 2^m *cameras lying in* 2^m *disjoint connected viewing regions on the viewing sphere* $\mathcal{S}$*.*

PROOF. By induction on m. For the base case when $m = 1$ we start with the polyhedron in figure 2 which contains a point p_1, at height 1 from the base B. p_1 is connected to the base using a pedestal, which is in fact an infinitesimal pyramid whose shadow on $\mathcal{S}$ can be ignored. Then, we create a region on the sphere that cannot see p_1 by using a rectangular block O and we connect O to B using another pedestal (see figure 3). It is easy to see that the polyhedron satisfies the properties in theorem 3.4 with p_1 marked for $m = 1$.

We will maintain the following inductive hypothesis: There exists a polyhedron $\mathcal{P}$ contained in a cone with base B and height km, where k is a constant. There are m points $P_m = \{p_1, \ldots, p_m\}$ marked on $\mathcal{P}$ at heights $k/2, 3k/2, 5k/2, \ldots$ respectively, such that for any subset $\omega \subseteq P_m$, there exists a connected region V_ω from which all points in ω are visible, but no point in $P_m \setminus \omega$ is.

First, we introduce a new point p_{m+1} at height $k/2$ above the tip of the surrounding cone. Note that all the viewing regions V_ω can see p_{m+1} because they lie in the (say) northern hemisphere. We split each V_ω into two connected regions, V_ω^+ and V_ω^- (see figure 4) as follows: We consider the largest inscribed rectangle in V_ω and ensure that V_ω^+ and V_ω^- each contain half of this rectangle. (We ignore the rest of V_ω.) Next, for each V_ω^-, we put a rectangular obstacle, that will block the visibility of p_{m+1} from V_ω^-. This obstacle will be placed at a small distance ϵ from p_{m+1} and will have an area that is $\frac{\epsilon^2}{2^{\Theta(m)}}$. This will ensure that V_ω^- has an area of $\frac{R^2}{2^{\Theta(m)}}$ as will V_ω^+. (The constant of proportionality in the Θ here comes from the fact that our areas are getting smaller because we discard portions of V_ω at each stage. Thus this constant is greater than 1.)

Again, we place the obstacles so that they lie inside the cone with base B and height $k(m+1)$ and connect them using small pedestals to $\mathcal{P}$. While the obstacles were added to block p_{m+1} from various regions on $\mathcal{S}$, they could have the unintended effect of blocking the visibility of other marked points from the viewing sphere. Later we will choose k and argue that the small size of the obstacles and their distance from other marked points makes this effect negligible.

We now have $2 \cdot 2^m = 2^{m+1}$ viewing regions that can shatter the set $P_{m+1} = \{p_1, \ldots, p_m, p_{m+1}\}$ marked on the polyhedron $\mathcal{P}'$ which contains $\mathcal{P}$, p_{m+1}, the obstacles and pedestals used during the construction. Also, during the construction we add a constant number c_1 of vertices for each 2^m viewing regions. By the inductive hypothesis suppose that $\mathcal{P}$ had at most $c_2 \cdot 2^m$ vertices, for a constant $c_2 \geq c_1$. Therefore, $\mathcal{P}'$ has at most $c_1 \cdot 2^m + c_2 \cdot 2^m \leq 2c_2 2^m = c_2 2^{m+1}$ vertices proving the inductive hypothesis.

We return to the assumption that an obstacle does not significantly block marked points that it was not intended to block. Let A be the total area of all obstacle placed at stage $m+1$ and A' be the area of their shadow with respect to point p_i for $i \in [1..m]$. Note that the distance from p_i to these obstacles is $m+1-i$ and that the distance from p_i to the viewing region is at most $2R$. (See figure 5). Using similarity, we obtain:

$$\frac{A'}{A} \leq \left(\frac{2R}{m+1-i}\right)^2 \tag{1}$$

$$A' \leq A\left(\frac{2R}{m+1-i}\right)^2 \tag{2}$$

Let Δ be the total area of the shadows the obstacle casts for all points p_i. Then:

$$\Delta \leq A\left(\frac{2R}{k}\right)^2\left(1+\frac{1}{2^2}+\cdots+\frac{1}{m^2}\right) \leq A\left(\frac{2R}{k}\right)^2\frac{\pi^2}{6} \tag{3}$$

The second inequality follows from the fact that $\sum_{i=1}^{\infty}\frac{1}{i^2} = \frac{\pi^2}{6}$.

Using the fact that $A = O(\epsilon^2)$, we get the total size of the shadow is $O((\frac{\epsilon R}{k})^2)$. By choosing k big enough and letting $\epsilon = 1/R$ we ensure that these shadows are indeed negligible as claimed. $\square$

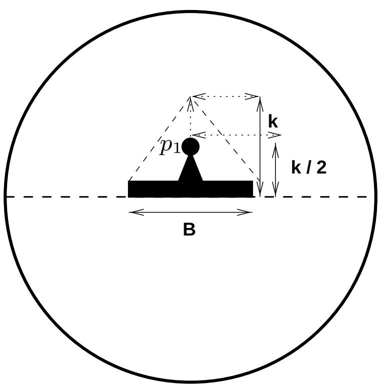

Figure 2: The base for the construction

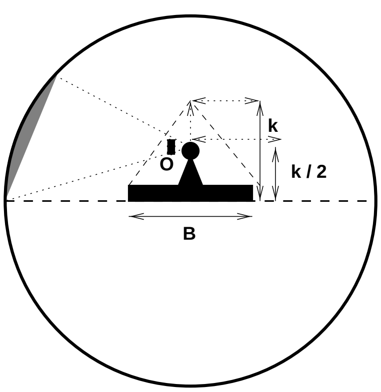

Figure 3: Introducing obstacles to block visibility

Proof of theorem 3.2. Theorem 3.2 is a direct consequence of lemmata 3.3 and 3.4. $\square$

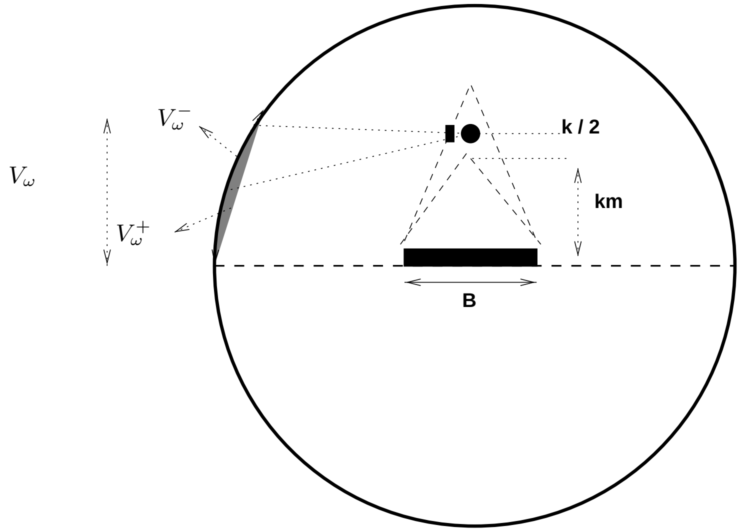

Figure 4: Inductive step

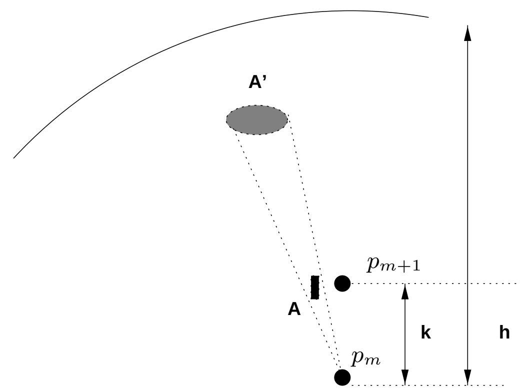

Figure 5: The obstacle introduced at step m may block the visibility of other points

4. RESULTS ON PLANAR CONFIGURATIONS

In this section we consider various settings where the camera locations are restricted and we decrease the existing upper bound of 23 on VC-Dimension under such restricted settings. First we restrict the camera locations to a circle around the polygon and obtain the exact VC-Dimension of 2. Then, we relax the restriction on camera locations and we allow cameras anywhere outside the convex hull of the polygon. In this case the VC dimension is bounded between 4 and 6.

4.1 Cameras restricted to a circle around a simple polygon

Consider a setup in which we want to cover the polygon P using cameras restricted to a circle C around P.

Definition 4.1. *We define 2DCIRCLE as a setup where a set of cameras whose locations are restricted to a circle C are to cover a simple polygon that is contained in C.*

We need the following technical lemma before proving our main theorem.

LEMMA 4.2. *Each point p on P is visible along a continuous arc on the circle C and nowhere else.*

PROOF. Let p be a point on P and let c be a camera located on the visibility circle C that can see p. As c moves from p clockwise along C, there will come a time when the ray R starting from p and passing from c intersects P. Let b be the location of c at this time. Similarly, if c moves counterclockwise from its initial location, the ray L starting from p and passing from c will intersect P. Call this location a. Clearly all the points on the arc $\widehat{ab}$ can see p.

Now suppose that p is also visible from points outside of the arc $\widehat{ab}$. Let c' be the first such point counterclockwise of a and c'' be the first such point clockwise of b. This means that in the regions $\widehat{c'a}$ and $\widehat{bc''}$ p is not visible. There must be points q and r in the boundary of P that lie in the sectors formed by p with $\widehat{c'a}$ and $\widehat{bc''}$ respectively. Traversing the boundary of P in order to go through p, q, r we find that the boundary must cut through one of the lines of visibility of p, a contradiction. □

We now prove the following theorem.

THEOREM 4.3. *The VC-Dimension of 2DCIRCLE is exactly 2.*

PROOF. Let c_i be a camera that sees a point p on P. If we move c_i clockwise along the circle, p will be visible to c_i until it reaches the intersection of the half-line L_p^R with the circle and it will remain invisible afterwards until c_i reaches L_p^L. Therefore for each point p, there exists an arc A on C such that p is visible on the arc and invisible otherwise.

Now, consider m points $P_m = \{p_1, \ldots, p_m\}$ that can be shattered. We must identify 2^m points c_ω on the circle corresponding to each subset ω of P_m such that c_ω can see all points in ω but no point in $P_m \setminus \omega$. As we discussed above, for each p_i there is an arc A_i from which p_i is visible. In general, some of these arcs will be overlapping. The $2m$ endpoints of $A_i,\ i = 1, \ldots, m$ divide C into $2m$ arcs $A'_j,\ j = 1, \ldots, 2m$ such that A'_j are disjoint and within an A'_j only a fixed subset of P_m is visible. Therefore each camera $c_{w_i},\ i = 1, \ldots, 2^m;\ \omega_i \in P_m$ must lie on a different $A'_j,\ j = 1, \ldots, 2m$. But this implies that $2m$ must be greater than or equal to 2^m which is only true for m less than 3. Thus the VC dimension is upper bounded by 2.

The lower bound is proved by the example in Figure 6 where the points p_1 and p_2 are shattered by the 4 cameras shown. □

4.2 Cameras restricted to remain outside the convex hull of a simple polygon

Let us now relax the restriction on camera locations so that we allow cameras anywhere outside the convex hull of the polygon.

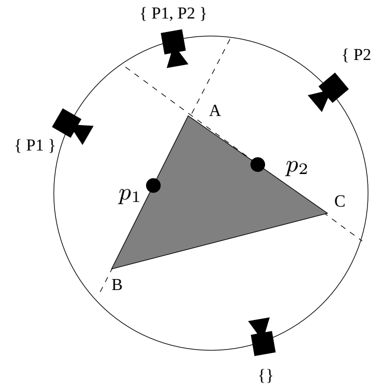

Figure 6: p_1 and p_2 can be shattered by four cameras. Each camera is labeled with the subset it can see. In this figure the polygon P is $\triangle ABC$.

DEFINITION 4.4. *We define 2DCONVEX as a setup where a set of cameras located outside the convex hull of a simple polygon P are to cover P.*

The upper bound on the VC-Dimension of 2DCONVEX slightly increases but it is still a small constant significantly less than the upper bound for the general case, 23.

PROPOSITION 4.5. *The VC-Dimension of 2DCONVEX is less than or equal to 6.*

PROOF. If $P_m = \{p_1, \ldots, p_m\}$ are m points to be shattered on the polygon P, then for each p_i there are two lines $L_{p_i}^L$ and $L_{p_i}^R$ such that p_i is visible from the region lying between $L_{p_i}^L$ and $L_{p_i}^R$ in a clockwise scan. Note that we restricted the cameras to lie outside the convex hull of P, otherwise there could be many lines L_i such that the visibility of a point p changes as we cross L_i. This is illustrated in figure 7.

Consider the arrangement of the $2m$ lines $\{L_{p_i}^L, L_{p_i}^R \mid i = 1, \ldots, m\}$ as in the proof of lemma 4.2. For contradiction's sake assume that there are two cameras c_ω and $c_{\omega'}$ located in the same face [1] of the arrangement, such that c_ω can see all points in $\omega \subseteq P_m$ but no point in $P_m \setminus \omega$ and similarly for $c_{\omega'}$. If ω and ω' are two different subsets, as we move from c_ω to $c_{\omega'}$ on the line defined by them either a point $q \in \omega \setminus \omega'$ disappears or a point $q \in \omega' \setminus \omega$ becomes visible. But the only reason of this visual event can be crossing L_q^L or L_q^R, which contradicts with the fact that c_ω and $c_{\omega'}$ are in the same face of the arrangement. Therefore, each of the 2^m cameras $c_\omega,\ \omega \subseteq P_m$ must lie in a different face of the arrangement. It is well known that an arrangement of $2m$ lines has at most $\binom{2m+1}{2}+1$ faces therefore we must have $\binom{2m+1}{2}+1 \geq 2^m$ which is only true for m up to 7. Hence the VC-Dimension of this system is at most 6. □

[1]The faces of an arrangement of lines are the connected regions on the plane remaining after the removal of the lines from the plane.

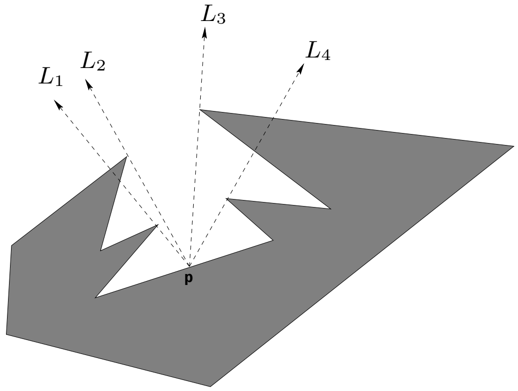

Figure 7: The effect of restricting the cameras to lie outside the convex hull for the setup 2DCONVEX: Only lines L_2 and L_3 are relevant to the visibility of point P.

Note that relaxing the camera locations from 2DCIRCLE to 2DCONVEX indeed increases the VC-Dimension, as we see in the following proposition.

PROPOSITION 4.6. *The VC-Dimension of 2DCONVEX is greater than or equal to 4.*

PROOF. Again, we present an instance where four points can be shattered from 16 cameras. In figure 8, the points $\{A, B, C, D\}$ can be by sixteen cameras lying outside the convex hull. See caption for details. □

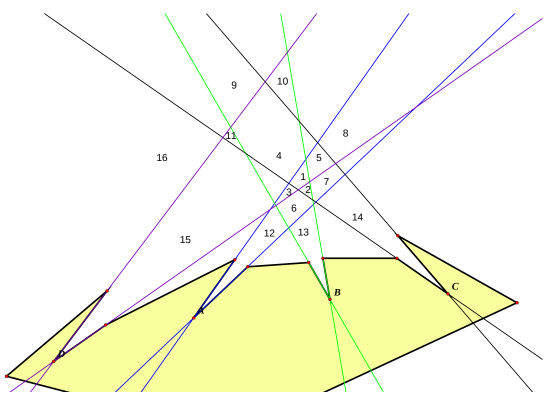

Figure 8: In the setup 2DCONVEX one can shatter a bigger subset than in 2DCIRCLE. The set $\{A, B, C, D\}$ can be shattered by the 16 cameras. Numbers in the figure correspond to the following subsets: 1.$\{A, B, C, D\}$, 2.$\{A, B, C\}$, 3.$\{A, B, D\}$, 4.$\{B, C, D\}$, 5.$\{A, C, D\}$, 6.$\{A, B\}$, 7.$\{A, C\}$, 8.$\{A, D\}$, 9.$\{B, C\}$, 10.$\{B, D\}$, 11.$\{C, D\}$, 12.$\{A\}$, 13.$\{B\}$, 14.$\{C\}$, 15.$\{D\}$, 16.$\{\}$.

REMARK 4.7. *If we further remove the restriction that the cameras are outside the convex hull, then the best known bound is 23 and the reader is referred to [16].*

4.3 Discussion: Bounded Number of Obstacles

What happens if the number of obstacles we would like to cover is more than one? This is related to art gallery theorems for polygons with holes and there are sufficiency results available for such settings [11]. In addition, the number of views of such systems have also been studied [2]. Unfortunately, the VC-Dimension of such systems is no longer constant [16, 7] but bounded by $2log(h) + 4log(log(h)) + o(1)$, where h is the number of holes (or number of objects for external visibility) of the polygon. This suggests that if the number of polygons in the scene is bounded, it is possible to come up with scenarios where the VC-Dimension is bounded for the value $2log(h) + 4log(log(h)) + o(1)$ will evaluate to a constant.

For example, consider k polygons on the plane and cameras restricted to lie on a circle that contains all the polygons. We can extend the discussion for 2DCIRCLE as follows. Consider a point $p_i \in P_m$, where $P_m = \{p_1, \ldots, p_m\}$ is the set to be shattered, located on one of the polygons. As we discussed in the case of 2DCIRCLE, if there were no other objects there would be an arc from which p_i is visible. Let us introduce the k-1 polygons one by one. We can have either,

1. the polygon introduced at step i totally blocks an arc, or,
2. the size of an arc is reduced, or,
3. an arc is split into 2, or,
4. an arc is not affected at all.

In addition, if there exists an arc that is split into 2, there can be no other arcs associated with p_i that splits. Therefore, at each step the number of visibility arcs that can see p_i increases by at most 1. This implies that for each p_i there will be at most k arcs on the circle such that visibility of p_i changes as you move in and out of these arcs. Again, using the endpoints we can divide the viewing circle into $2mk$ arcs, such that any camera c_ω that sees $\omega \subseteq P_m$ must lie in a different arc. Therefore we must have $2mk \geq 2^m$, which is true only up to a constant that depends on k.

Therefore when we know that there are bounded number of objects in the scene, the algorithm for systems with bounded VC-Dimension is still beneficial to find an optimal placement of cameras.

5. CIRCLE GUARDS

DEFINITION 5.1. *Let P be a simple polygon totally contained in a circle C such that every point on the boundary of P is visible from some point on C. The Guard Placement problem is to find a minimum cardinality set G of points on C such that G covers P.*

As shown in Lemma 4.2 the region of C from which any particular point p on P is a continuous arc.

DEFINITION 5.2. *Let q be a point on P. The visibility arc of q, A_q, is the set of points on C that can see q. A_q is called* *a minimal arc if there is no point p on P such that the arc A_p is properly contained in A_q.*

Alternatively, the guard placement problem can be viewed as the problem of hitting the set of minimal arcs.

The following subroutine for finding an optimal hitting set for intervals on a line will be useful.

IntervalHittingSet(S)

Input: A set S of intervals on a line.

Output: A minimum cardinality set of points, P, such that for each interval $s \in S$, there exists a point in P contained in s.

$G \leftarrow \emptyset$
Sort the intervals in S according to their right ends
while S is not empty
    Let $[a, b]$ be the interval that
        has the leftmost right end
    $P \leftarrow P \cup \{b\}$
    Remove all intervals that intersect b from S

LEMMA 5.3. *IntervalHittingSet returns a minimum cardinality set in polynomial time.*

PROOF. The lemma is proved trivially by induction. The first guard we place is not to the left of the first guard in any feasible solution and hence covers at least the set of intervals covered by the first guard in any solution. □

5.1 Overview of the Algorithm

Start with an arbitrary point on c which moves clockwise on C. We will maintain two lists:

- **Covered:** Contains all the portion of the boundary of P that is currently covered. Initially empty.
- **Active:** The parts of the boundary that have been seen by c but not currently covered. Initially contains the segments observed from the initial position of c.

For now, assume continuous motion of c. The algorithm proceeds as follows:

1. When c moves to the point c' on V
2. If new boundary becomes visible from c', add it to Active.
3. If a point in Active is not visible from c',
   - Place guard g at c'.
   - Add the part of the boundary visible from g to Covered.
   - Remove the part of the boundary visible from g from Active.
4. $c \leftarrow c'$
5. Repeat until Covered $= P$.

Obviously, we can not move c to all the points on C. The main challenge in this algorithm is to find a small finite set of "events" of interest along C. Let us define appearance/disappearance of a point p on P as c moves in clockwise direction as a *visual event.* In the next section we show that some visual events necessarily occur before others. Therefore we can consider finitely many locations of c.

5.2 Visual Events

By Lemma 4.2 for each point p there is an arc $A_p = \widehat{ab}$ on which p is visible. a is called the *entrance* of the arc and b is called the exit.

DEFINITION 5.4. *Let e be a line segment (either an edge or a portion of an edge) on the boundary of P. Then the weak visibility region of e is defined as $\bigcup_{p \in e} A_p$.*

Let e be a line segment on the boundary of P. We can rotate P so that e is horizontal and the exterior of P is locally above e. We will call this the canonical orientation of P with respect to e. It is clear that in this orientation, the interior of e can only be seen above the horizon defined by e. For two points a, b on the viewing circle C and above this horizon we say that $a \leq b$ if a appears clockwise of or coincides with b.

LEMMA 5.5. *Let P be oriented canonically with respect to line segment e on the boundary. Let p, p' be points in the interior of e with p to the left of p'. Let $A_p = \widehat{ab}$ and $A_{p'} = \widehat{cd}$. Then $c \leq a$ and $d \leq b$.*

PROOF. Suppose for contradiction that $a < c$. Then there is a point a' between a and c that can see p but cannot see p'. Clearly the boundary of the polygon must intersect the segment $\overline{a'p'}$. By definition the boundary of the polygon cannot intersect the wedge formed by c, d, and p' which lies to the right of this segment. Thus the boundary of P must intersect $\overline{a'p'}$ from the left. This will force it to first intersect the segment $\overline{a'p}$ contradicting the assumption that p is visible from a'. Thus we have a contradiction. A symmetric argument shows that $d \leq b$. □

A useful corollary of Lemma 5.5 is that if $[p, q]$ is a segment on the boundary of P, then in any clockwise traversal starting from some point that does not see $[a, b]$, b appears before any point in (a, b) and a disappears after all points in (a, b).

LEMMA 5.6. *The weak visibility region of a line segment on P is a continuous arc.*

PROOF. Let p and p' be two points on a line segment on the boundary of P. Assume that we have canonical orientation and p lies to the left of p'. We show the lemma by showing that one of the following must hold: Either A_p and $A_{p'}$ intersect or, for any point x on C which lies strictly between $A_{p'}$ and A_p x can see some point on e between p and p'.

Suppose A_p and $A_{p'}$ do not intersect. Then by Lemma 5.5 $A_{p'} = \widehat{cd}$ is to the left of A_p. Let the segments $\overline{dp'}$ and $\overline{ap}$ intersect at point o_1 and the segments $\overline{cp'}$ and $\overline{bp}$ intersect at point o_2. Suppose x is point between d and a and x is completely blocked from seeing the segment $\overline{pp'}$. Then either there is a portion of the boundary of P in the sector defined by d, a, and o_1 or there is a portion of the boundary of P in the sector defined by p, p', and o_2. Either of these possibilities leads to a contradiction with respect to the known visibilities. □

6. THE ALGORITHM

In this section we present the algorithm *PlaceGuards* that mimics the algorithm outlined in the previous section.

PlaceGuards(C, P)

Input: A circle C and a polygon P

Output: a set of points, G, on C such that G covers P.

```
G ← ∅
pick an arbitrary point p on C
Active ← {chains on P visible from p}
Covered ← ∅
Q_a ← ∅
For each vertex v of P
    Find the visibility arc A_v of v
    Q_a ← Q_a ∪ {A_v}
Sort Q_a according to their exits
(*from now on we keep Q_a always sorted*)
While Covered ≠ P
    For each chain s ∈ Active
        Let a, b be the endpoints of s
        Find the visibility arcsA_a and A_b
        Q_a ← Q_a ∪ {A_a, A_b}
    p ← Head(Q_a).exit
    G ← G ∪ {p}
    remove all the arcs that intersect p form Q_a
    add the portions of boundary visible from p
        to Covered.
```

THEOREM 6.1. *PlaceGuards uses at most one more guard than optimal.*

PROOF. After PlaceGuards has placed the first guard at point a on C we remove all arcs that pass through a. At this point the C can be "opened" at a and the remaining arcs can be viewed as intervals on a line. PlaceGuards now mimics the IntervalHittingSet algorithm which is optimal for this problem. □

6.1 Analysis

In this section, we show that the running time of PlaceGuards is polynomial by showing that the number of line segments in the lists *Covered* and *Active* is bounded by a polynomial.

DEFINITION 6.2. *We say a guard g splits a line segment s if s splits into two disjoint line segments s_1 and s_2 after the removal of the portion of s visible from g. s_1 and s_2 are called the children of s.*

LEMMA 6.3. *Let e be an edge of P. During the course of the algorithm PlaceGuards e may be split at most once and none of the children of e is split afterwards.*

PROOF. Initially an edge may be split if the locations of the current set of guards do not intersect the visibility arcs of its endpoints. However, as the algorithm proceeds clockwise, none of the children will be split due to lemma 5.5. □

COROLLARY 6.4. *At any instance of the algorithm PlaceGuards size of the list that contains visible line segments is at most $2n$.*

7. REFERENCES

[1] P. K. Agarwal and S. Sen. Randomized algorithms for geometric optimization problems. *Handbook of Randomization (P. Pardalos, S. Rajasekaran, J. Reif, and J. Rolim, eds.), Kluwer Academic Publishers*, to appear.

[2] B. Aronov, H. Brönnimann, D. Halperin, and R. Schiffenbauer. On the number of views of polyhedral scenes. *Proceedings of the Japan Conference on Discrete and Computational Geometry*, 2000.

[3] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite VC-dimension. *GEOMETRY: Discrete & Computational Geometry*, 14, 1995.

[4] S. Eidenbenz, C. Stamm, and P. Widmayer. Inapproximability results for guarding polygons and terrains. *Algorithmica*, 31:79–113, 2001.

[5] U. Feige. A threshold of ln n for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.

[6] S. Ghosh. Approximation algorithms for art gallery problems. *Proc. of the Canadian Information Processing Society Congress*, 1987.

[7] H. Gonzalez-Banos and J. Latombe. A randomized art-gallery algorithm for sensor placement. *Proc. ACM Symp. on Computational Geometry (SoCG'01)*, 2001.

[8] D. Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on Computing*, 11(3):555–556, 1982.

[9] M. J. Kearns and U. V. Vazirani. *An Introduction to Computational Learning Theory.* The MIT Press, 1994.

[10] J. Matousek. Geometric computation and the art of sampling. In *IEEE Symposium on Foundations of Computer Science*, 1998.

[11] J. O'Rourke. *Art Gallery Theorems And Algorithms*. Oxford University Press, 1987.

[12] J. O'Rourke. Visibility. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 25, pages 467–480. CRC Press LLC, 1997.

[13] W. Plantinga and C. Dyer. Visibility, occlusion, and the aspect graph. *International Journal of Computer Vision*, 5(2):137–160, 1990.

[14] T. Shermer. Recent results in art galleries. *Proceedings of the IEEE*, 80(9):1384–1399, 1992.

[15] P. Slavik. A tight analysis of the greedy algorithm for set cover. *Journal of Algorithms*, 25(2):237–254, November 1997.

[16] P. Valtr. Guarding galleries where no point sees a small area. *Israel Journal of Mathematics, 104:1-16*, 104:1–16, 1998.

[17] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.