

Approximation Algorithms to Minimum Vertex Cover Problems on Polygons and Terrains

Ana Paula Tomás^{*1}, António Leslie Bajuelos², and Fábio Marques³

¹ DCC-FC & LIACC, University of Porto, Portugal

apt@ncc.up.pt

² Department of Mathematics / R&D Unit “Mathematics and Applications”
University of Aveiro, Portugal

leslie@mat.ua.pt

³ School of Technology and Management, University of Aveiro, Portugal
fabio@estga.ua.pt

Abstract. We propose an anytime algorithm to compute successively better approximations of the optimum of MINIMUM VERTEX GUARD. Though the presentation is focused on polygons, the work may be directly extended to *terrains* along the lines of [4]. A major idea in our approach is to explore dominance of visibility regions to first detect pieces that are more difficult to guard.

1 Introduction, Related Works, and Problem Definition

The *Art Gallery Problem* [17] was posed in 1973 by V. Klee: “How many guards are always sufficient to guard any simple polygon P with n vertices?”. Each guard is considered a fixed point with 2π range visibility, that is, it can see in every direction. An equivalent formulation of this problem, is to ask how many point lights are needed to fully illuminate P , so that it is also known as an *Illumination Problem*. In our work, we focus on problems where guards may be only placed at vertices of P – VERTEX GUARDS. By simple polygon we mean a region of the plane enclosed by a simple cycle of straight line segments. In a simple cycle, non-adjacent segments do not intersect and two adjacent segments intersect only in their common endpoint. A simple polygon P is a polygon without holes, that is, its interior is topologically equivalent to a disk. We say that point x sees point y (or y is visible to x) iff the closed segment \overline{xy} does not intersect the exterior of P . The set of all points of P visible to $v \in P$ is called the visibility region of v and denoted $V(v)$. If $V(v)$ is the union of a polygon $Q \subseteq P$ and some line segments, we restrict $V(v)$ to Q , so that $V(v)$ will fit our definition of polygon. Chvátal has shown that $\lfloor \frac{n}{3} \rfloor$ guards are always sufficient and occasionally necessary to guard a simple polygon with n vertices [2]. In 1978, Fisk [5] gave a simpler proof based on the fact that the graph obtained by triangulating a simple polygon is

* Work partially supported by funds granted to LIACC through *Programa de Financiamento Plurianual, Fundação para a Ciência e Tecnologia* and *Programa POSI*.

three-colorable. This idea led Avis and Toussaint [1] to develop an $O(n \log n)$ algorithm for locating $\lfloor \frac{n}{3} \rfloor$ stationary guards in a simple polygon.

Orthogonal (rectilinear) polygons, that is, polygons whose edges are parallel to the axes, are of particular interest. Indeed, most real life buildings are, after all, orthogonal. Moreover, these kind of polygons arise naturally in certain applications, such as VLSI design and computer graphics. Interesting results on this subclass of art gallery problems include the proof by Kahn et al. [9] that $\lfloor \frac{n}{4} \rfloor$ vertex guards are sufficient and sometimes necessary to guard n -vertex orthogonal polygons. Edelsbrunner et al. [3] developed an $O(n \log n)$ algorithm for stationing $\lfloor \frac{n}{4} \rfloor$ guards in a rectilinear polygon, by performing an L-shaped partition of the polygon. Later, Sack and Toussaint [15] showed that guard placement in *monotone* rectilinear polygons can be done in linear time, while the problem can be solved for arbitrary n -vertex rectilinear polygons in $O(n \log \log n)$.

In contrast, MINIMUM VERTEX GUARD, that is the problem of finding the minimum number of vertex guards needed to cover P , is much harder. Lee and Lin [10] proved that it is NP-hard for polygons and Shuchardt and Hecker [16] have shown that it is also NP-hard for orthogonal polygons. In [17], Urrutia asserts that one approach that has been neglected in the study of art gallery problems is that of finding algorithms that obtain approximate solutions. The most well-known result on this subject is an algorithm by Ghosh [7], that finds in $O(n^5 \log n)$ time a vertex guard set that is at most $O(\log n)$ times the minimum number of vertex guards needed. Further results are given by Eidenbenz in [4], who also designed approximation algorithms and heuristics for several variations of terrain guarding problems, by finding the right extension of Ghosh's work.

Our contribution. We propose an anytime algorithm to compute successively better approximations of the optimum to MINIMUM VERTEX GUARD. Anytime means that the algorithm returns the best answer found even if it is not allowed to run to completion, and may improve it if it is allowed to run longer. Although the presentation is focused on polygons, our work may be directly extended to terrains along the lines of [4]. A major idea in our approach is to successively refine an initial partition of polygon P and to explore dominance of visibility regions to try to identify, as soon as possible, pieces that are more difficult to guard. An experimental evaluation of the algorithm is being done, in which standard Computational Geometry algorithms and Constraint Programming techniques will be used in an hybrid way.

The paper is structured as follows. In the next section we formalize the problem and state some useful results that support our method. In section 3, we describe the algorithm and present an example, and then conclude.

2 From Minimum Vertex Guard to Minimum Set Cover

MINIMUM VERTEX GUARD is usually modeled as a discrete combinatorial problem – MINIMUM SET COVER – using decompositions of polygon P . It is often required that each piece in such a decomposition is either totally visible or totally

invisible from each vertex v , which guarantees that the approximation is exact. This feature, which we shall term **PIECE RESTRICTION**, may render decompositions too grained [7]. Nevertheless, for some real-world applications, the solution to the discrete model seems to give a fairly good approximation [4], even if the decomposition does not verify **PIECE RESTRICTION**. As we mentioned already, our idea is to take successively finer partitions to improve an initial approximation. We now introduce some definitions about visibility and some results that we need to present the algorithm.

A vertex of a polygon P is called *convex* if the interior angle between its two incident edges is at most π ; otherwise it is called *reflex*. We use symbol r to represent the number of reflex vertices of P . It has been shown by O'Rourke [14] that $n = 2r + 4$ for every n -vertex orthogonal polygon. Thus the upper bound $\lfloor \frac{n}{4} \rfloor$ given in [9] is identical to the bound $\lfloor \frac{r}{2} \rfloor + 1$.

Definition 1. A partition Π of polygon P is a division of P into sub-polygons (named pieces) that do not overlap except on their boundaries.

Definition 2. A *rectilinear cut* (r -cut) of an orthogonal polygon P is obtained by extending each edge incident to a reflex vertex of P towards its interior until it hits the boundary of P . By drawing all r -cuts, we partition P into rectangles (called r -pieces, c.f. Fig. 1(a)). This partition will be denoted by Π_{r-cut} .




Fig. 1. Construction and visibility of the r -pieces.

Following the terminology of [3,14], we say that to draw an r -cut, we draw an horizontal and a vertical cut, but keep the resulting parts together until all r -cuts are done. Thus, Π_{r-cut} does not depend on the order we do cuts, which is useful to show Proposition 4 in section 2.2. By noting that each drawn horizontal (vertical) line intersects at most n previously drawn vertical (horizontal) lines, we may show that Π_{r-cut} has at most $1 + r + r^2$ pieces.

2.1 Notions of Visibility

Definition 3. A piece R is totally visible from vertex v if $R \subseteq V(v)$ and is partially visible from v if only part of the interior of R is visible from v . If R is partially visible from v then $V(v) \cap R$ is called a visibility section and is denoted by $V_R^s(v)$. A piece R is visible by sections if it is union of visibility sections.

We represent the sets of all pieces that are totally visible from v and partially visible from v , by G_v^t and G_v^s , respectively. Similarly, G_R^t and G_R^s represent the sets of all vertices that are visible from every point of R and that are visible from only part of the interior points of R . Thus, a vertex $v \in G_R^t$ iff $R \subseteq V(v)$ and $v \in G_R^s$ iff there are points p and q in the interior of R such that p is visible from v but q is not. (We note that G stands for guard.)

The following result gives a simple condition that we may check to decide whether a convex piece is totally visible from a vertex v of P . It may be used with Proposition 2, to efficiently find G_v^t and G_v^s , provided we have a good representation of the grid defined by Π_{r-cut} (or another partition Π).

Proposition 1. Let v be a vertex of an orthogonal simple polygon P and R an r -piece inside P . If e_1, e_2, e_3 and e_4 are the vertices of R and v and all e_i 's are visible, then for all $p \in R$, v and p are visible (i.e., $R \subseteq V(v)$).

Proof. We first observe that if $R \subseteq P$ and v and the four e_i 's are visible, then the line segments $\overline{ve_i}$ do not intersect the exterior of P . We now prove that the convex hull $\mathcal{H}(S)$ of the set of points $S = \{v, e_1, e_2, e_3, e_4\}$ is in P (see Fig. 1(b)). If it were not, we could find a point $q \in \mathcal{H}(S)$ with $q \notin P$. Hence, P would not define a closed simply-connected set and so it would not be a simple polygon. ■

Note 1. We actually have proved that, under the assumptions of this proposition, $\mathcal{H}(S) \subseteq V(v)$, for $S = \{v\} \cup \{\text{vertices of } R\}$. Therefore, Proposition 1 still holds for any simple polygon P , if we replace r -piece by convex piece.

Proposition 2. If a piece $R \subset P$ is partially visible from a vertex v , the visibility section $V_R^s(v)$ defines a closed simply-connected set.

Proof. If the proposition were false, $V_R^s(v)$ could be divided into polygons and there would be two disjoint polygons in such partition, say $V_{R,i}^s(v)$ and $V_{R,j}^s(v)$, and points $p, q \in R$ such that $p \in V_{R,i}^s(v)$, $q \in V_{R,j}^s(v)$ and $\overline{pq} \notin V_{R,i}^s(v) \cup V_{R,j}^s(v)$. Thus, we could find $w \in R$, $w \in \overline{pq}$, $w \notin V_{R,i}^s(v) \cup V_{R,j}^s(v)$ such that w and v were not visible ($\overline{vw} \notin P$). Therefore, there would exist t such that $t \notin P$ and $t \in \Delta pvq \subset P$ (see Fig. 1(c)). Hence, P would not be a simple polygon because it would not define a closed simply-connected set. ■

2.2 Different Discrete Models of the Problem

Let Γ_0^Π be the set of pieces in partition Π of P that are partially visible from at most one vertex. Let Γ_∞^Π be the set of pieces that are not visible by sections.

Lemma 1. *Every piece in Γ_0^{Π} is not visible by sections, that is $\Gamma_0^{\Pi} \subseteq \Gamma_{\infty}^{\Pi}$.*

Definition 4. MINIMUM VERTEX GUARD WITH PIECE RESTRICTION is the problem MINIMUM VERTEX GUARD with the additional restriction that guards are not allowed to cooperatively see a piece in the partition (i.e., visibility by sections is not allowed).

Definition 5. Given a region \mathcal{F} (i.e., a set of pieces), the optimum of MINIMUM VERTEX GUARD and the optimum of MINIMUM VERTEX GUARD WITH PIECE RESTRICTION w.r.t. \mathcal{F} are denoted by $OPT(\mathcal{F})$ and $OPT_{\square}(\mathcal{F})$, respectively, with guards placed in vertices of P .

Remark. In the rest of the paper, we assume that each piece R in a partition Π is totally visible from at least one vertex of P . Every partition may be refined to meet this condition. If a piece R is only visible by sections (i.e., when $G_R^t = \emptyset$), we replace R by the partition of R induced by its visibility sections, that we call \mathcal{Z}_R . The pieces in \mathcal{Z}_R are the cells of the arrangement of segments that define the boundary of R and the boundary of the visibility sections $V_R^s(v)$, for $v \in G_R^s$. Each piece in \mathcal{Z}_R is either totally visible or totally invisible to each vertex of P . In this way, we clearly have $OPT(\mathcal{F}) \leq OPT_{\square}(\mathcal{F})$, for all \mathcal{F} .

Now, for every $\Gamma \subseteq \Gamma_{\infty}^{\Pi}$, the optimum of MINIMUM VERTEX GUARD is the optimum of MINIMUM VERTEX GUARD WITH PIECE RESTRICTION, as stated by Lemma 2. It follows from $\Gamma_0^{\Pi} \subseteq \Gamma_{\infty}^{\Pi} \subseteq \Pi$ and the definitions of Γ and Γ_{∞}^{Π} .

Lemma 2. *For every subset Γ of Γ_{∞}^{Π} , we have $OPT_{\square}(\Gamma) = OPT(\Gamma)$. Furthermore, $OPT(\Gamma_0^{\Pi}) \leq OPT(\Gamma_{\infty}^{\Pi}) \leq OPT(P) \leq OPT_{\square}(\Pi)$, for every polygon P and partition Π of P , such $G_R^t \neq \emptyset$, for all $R \in \Pi$.*

A triangulation \mathcal{T} of polygon P is a decomposition of P into triangles by a maximal set of non-intersecting diagonals. Every triangle is totally visible from each of its vertices, that are also vertices of P . Proposition 3 is a direct corollary of Lemma 2 and of Fisk’s proof [5] of the upper bound $\lfloor \frac{n}{3} \rfloor$ for MINIMUM VERTEX GUARD on an n -vertex polygon.

Proposition 3. *For every triangulation \mathcal{T} of an n -vertex simple polygon P , $OPT_{\square}(\mathcal{T}) \leq \lfloor \frac{n}{3} \rfloor$.*

A similar result holds for orthogonal polygons, for which the partition Π_{r-cut} defined by r -cuts (c.f. Definition 2) is a natural decomposition.

Proposition 4. *For the partition Π_{r-cut} of an n -vertex orthogonal polygon P into r -pieces, $OPT_{\square}(\Pi_{r-cut}) \leq \lfloor \frac{n}{4} \rfloor = \lfloor \frac{r}{2} \rfloor + 1$.*

A key idea for proving it is the observation that Π_{r-cut} is a refinement of the decomposition proposed by Edelsbrunner et al. in [3]. This allows us to adapt the proof that their algorithm achieves this same upper bound. Ultimately both are based on O’Rourke’s proof of the Rectilinear Art Gallery Theorem [14] and proceed by induction on the number of reflex vertices. Each L-shaped polygon in their decomposition may appear as the union of three or more r -pieces in

Π_{r-cut} , but it is still true that all are totally visible from the single reflex vertex of the L-shaped polygon.

The works of Ghosh [7] and of Eidenbenz [4] have in common the approximation of $OPT_{\square}(\Pi)$ using a GREEDY algorithm [8] for MINIMUM SET COVER. It recursively adds to the solution the vertex that totally sees a maximum number of pieces not yet covered by the current solution. We argue that this heuristic does not take enough advantage of the polygon’s topology, and in particular, of the fact that some pieces are dominant over others: if they are visible so are the dominated ones. Next we formalize this notion of dominance and explore it to design an approximation algorithm.

3 An Anytime Approximation Algorithm

The main idea is to find a sequence of successively shorter intervals enclosing $OPT(P)$. To approximate $OPT(P)$ from below we consider an increasing powerful subset of dominant pieces that are not visible by sections. To find a sequence of decreasing upper bounds, we successively refine the given partition of P . We introduce two notions of dominance both for vertices and for pieces.

Definition 6. A vertex v_i is strictly dominant over a vertex v_j if $V(v_j) \subseteq V(v_i)$. A vertex v_i is \square -dominant over a vertex v_j if $G_{v_j}^t \subseteq G_{v_i}^t$.

Definition 7. A piece R_i is \square -dominant over a piece R_j if $G_{R_i}^t \subseteq G_{R_j}^t$, with $R_i, R_j \in \Pi$. A piece R in $\Gamma \subseteq \Gamma_{\infty}^{\Pi}$ is strictly dominant over a different piece $R' \in \Pi$ if $G_R^t \subseteq G_{R'}^t$, meaning that R' is necessarily guarded whenever R is guarded. We recall that Γ_{∞}^{Π} is the set of pieces that are not visible by sections.




Fig. 2. Example

Example 1. Though the polygon in Fig. 2 is quite simple, it illustrates the impact of finding dominant pieces and vertices. We may see that $G_{R_1}^t = \{0, 1, 2, 15\}$ and $G_{R_1}^s = \{4, 5, 6, 8, 9, 10, 12, 13, 14\}$ but R_1 is not visible by sections. Hence, all

other r -pieces R such that $G_R^t \supseteq G_{R_1}^t$ are \square -dominated and strictly dominated by R_1 . In this case, we may discard $R_6, R_{10}, R_{15}, R_{18}, R_{22}$ and R_{24} . Piece R_2 is not visible by sections either, so that we can throw away R_3, R_4 and R_5 . Eventually we would keep just three pieces R_1, R_{19} and R_{11} . The latter is visible by sections from the non-dominated vertices 2 and 13, which totally guard R_1 and R_{19} . Here, $\Gamma_0^{\Pi} = \{R_3, R_4, R_5, R_6, R_7, R_8, R_{10}, R_{14}, R_{15}, R_{17}, R_{18}, R_{21}, R_{22}\}$ but R_7 is strictly dominant over them. Also, $\Gamma_{\infty}^{\Pi} = \Pi \setminus \{R_{11}\}$ and Π is Π_{r-cut} .

Let $\mathcal{D}_{\square}(\mathcal{F})$ and $\mathcal{D}(\mathcal{F})$ be maximal sets of all \square -dominant pieces and of strictly non-dominated pieces in the family/region \mathcal{F} . We wrote maximal instead of maximum because R_i is \square -dominant over R_j and R_j is \square -dominant over R_i if $G_{R_i}^t = G_{R_j}^t$. To break ties, either R_i or R_j may be selected as dominant, but not both. Proposition 5 follows directly from these definitions of dominance.

Proposition 5. *For all partitions Π of P , we have $OPT_{\square}(\Pi) = OPT_{\square}(\mathcal{D}_{\square}(\Pi))$ and $OPT(\Gamma) = OPT(\mathcal{D}(\Gamma)) = OPT_{\square}(\Gamma) = OPT_{\square}(\mathcal{D}_{\square}(\Gamma))$, for every $\Gamma \subseteq \Gamma_{\infty}^{\Pi}$.*

Refining Partitions. As before, let Γ_0^{Π} be the set of pieces in partition Π of P that are partially visible from at most one vertex. Let us construct a sequence of pairs (Γ_i, Π_i) , for $i \geq 0$, being Π_i a partition of P and Γ_i a subset of pieces of Π_i that are not visible by sections, with $\Pi_0 = \Pi$ and $\Gamma_0 = \Gamma_0^{\Pi}$. We consider

$$\Delta_{i+1} = \{R \mid R \in \Pi_i \setminus \Gamma_i \text{ and } R \text{ is not } \square\text{-dominated by a piece in } \mathcal{D}(\Gamma_i)\}$$

and if $\Delta_{i+1} \neq \emptyset$, we select a piece $R \in \Delta_{i+1}$ and find whether R is visible by sections. For that, we compute the partition \mathcal{Z}_R of R determined by the visibility sections. As we noted in section 2.2, each piece in \mathcal{Z}_R satisfies PIECE RESTRICTION, i.e., is either totally visible or totally invisible to each vertex of P . So, R is not visible by sections iff there is a piece in \mathcal{Z}_R that is not visible to any vertex in G_R^s . We define $\Pi_{i+1} = (\Pi_i \setminus \{R\}) \cup \mathcal{Z}_R$ and $\Gamma_{i+1} = \Gamma_i \cup \mathcal{Z}_R$, if R is visible by sections, and $\Pi_{i+1} = \Pi_i$ and $\Gamma_{i+1} = \Gamma_i \cup \{R\}$, otherwise. This completion procedure stops when $\Delta_{i+1} = \emptyset$, because then $\mathcal{D}(\Gamma_i) = \mathcal{D}(\Pi_i)$ and $OPT_{\square}(\mathcal{D}(\Pi_i)) = OPT(\mathcal{D}(\Pi_i)) = OPT(P)$. Non-increasing intervals enclosing $OPT(P)$ are determined by this sequence $\{(\Gamma_i, \Pi_i)\}_{i \geq 0}$, as formally stated by Proposition 6. It follows from $\Gamma_i \subseteq \Gamma_{i+1}$, and the fact that Π_{i+1} refines Π_i .

Proposition 6. *For all $i \geq 0$,*

$$OPT_{\square}(\Gamma_i) \leq OPT_{\square}(\Gamma_{i+1}) \leq OPT(P) \leq OPT_{\square}(\Pi_{i+1}) \leq OPT_{\square}(\Pi_i)$$

We note that if $R \in \Pi_i$ is \square -dominated by some piece in Γ_i then this piece strictly dominates R , because pieces in Γ_i satisfy PIECE RESTRICTION. This implies that $\Delta_{i+1} \subset \Delta_i$, for all i , so that Δ_{i+1} can be obtained from Δ_i . Moreover, we may use a maximal dominant set $\mathcal{D}(\Gamma_i)$ instead of Γ_i .

In contrast to previous approaches [4,7], we do not start from a partition Π that satisfies PIECE RESTRICTION. Neither would we like to assume that $OPT_{\square}(\Pi)$ is a good approximation of $OPT(P)$, in practice, as argued by Eidenbenz for a problem of location of antennas [4]. Since MINIMUM SET COVER

is NP-hard, though for small instances it can naturally be solved to optimality, we would like to keep the number of elements in the partition small. This is why we decompose pieces of the initial partition Π that are visible by sections in a lazy evaluation fashion, and introduce the notion of dominance. A possible heuristic to select the piece $R \in \Delta_{i+1}$ that is the next one to be decomposed is to proceed from the boundary pieces to the most interior ones, to hopefully and quickly improve the bounds of the enclosing interval or reduce the number of dominant pieces. In particular, we may try to find pieces in *ears* of the polygon and may take more than one piece at each step. The duality between vertices and pieces may be further explored to reduce the sets of relevant pieces and vertices. This is the main idea behind the first step in the approximation algorithm for $OPT_{\square}(\mathcal{F})$, we describe next.

3.1 Approximation of $OPT_{\square}(\mathcal{F})$

The idea is to tighten the relevant sets $\mathcal{D}_{\square}(\mathcal{F})$ and $\mathcal{D}_{\square}(V_{\mathcal{F}})$ by applying jointly \square -dominance for vertices and for pieces. The algorithm is in Fig. 3: $V_{\mathcal{F}}$ is the set of vertices that totally see pieces of \mathcal{F} , and $G_{R,\mathcal{F}}^t = G_R^t \cap V_{\mathcal{F}}$ and $G_{v,\mathcal{F}}^t = G_v^t \cap \mathcal{F}$, for $v \in V_{\mathcal{F}}$ and $R \in \mathcal{F}$. Our intuition is that the number of relevant pieces is

```

Compute a maximal set  $\mathcal{D}_{\square}(V_{\mathcal{F}})$  of  $\square$ -dominant vertices in  $V_{\mathcal{F}}$ .
DO
  Replace  $V_{\mathcal{F}}$  by  $\mathcal{D}_{\square}(V_{\mathcal{F}})$ .
  Replace  $G_{R,\mathcal{F}}^t$  by  $G_{R,\mathcal{F}}^t \cap V_{\mathcal{F}}$ , for all  $R \in \mathcal{F}$ .
  Compute  $\mathcal{D}_{\square}(\mathcal{F})$ , a maximal set of  $\square$ -dominant pieces in  $\mathcal{F}$ .
  IF  $\mathcal{F} \neq \mathcal{D}_{\square}(\mathcal{F})$  THEN
    Replace  $\mathcal{F}$  by  $\mathcal{D}_{\square}(\mathcal{F})$ .
    Replace  $G_{v,\mathcal{F}}^t$  by  $G_{v,\mathcal{F}}^t \cap \mathcal{F}$ , for all  $v \in V_{\mathcal{F}}$ .
    Compute  $\mathcal{D}_{\square}(V_{\mathcal{F}})$ .
  WHILE  $\mathcal{D}_{\square}(V_{\mathcal{F}}) \neq V_{\mathcal{F}}$ .
    
```

Fig. 3. A polynomial algorithm for finding \square -dominance w.r.t. \mathcal{F} .

much smaller than the total number of pieces. The number of iterations is at most $\min(|V_{\mathcal{F}}|, |\mathcal{F}|)$. The algorithm preserves $OPT_{\square}(\mathcal{F})$, though the reduction of $V_{\mathcal{F}}$ may not preserve the set of optimal solutions. If we apply the GREEDY algorithm to the resulting \mathcal{F} , we find an approximation of $OPT_{\square}(\mathcal{F})$ with a ratio bound of $O(\log f)$ [12], where f is the number of pieces in the reduced \mathcal{F} . This approximation may be used to try to solve some instances to optimality, that is to find $OPT_{\square}(\mathcal{F})$.

3.2 Using Constraint Programming to Find $OPT_{\square}(\mathcal{F})$

The reformulation of MINIMUM VERTEX GUARD WITH PIECE RESTRICTION as MINIMUM SET COVER allows us to model it as a Constraint Satisfaction

Problem (CSP) and to solve it using Constraint Programming. The decision variables are: $X_v \in \{0, 1\}$, for $v \in V_{\mathcal{F}}$ (1 iff a guard is placed at vertex v) and $Y_R \in \{0, 1\}$, for $R \in \mathcal{F}$ (1 iff piece R is guarded). The constraints are (1)–(3)

$$\sum_{R \in \mathcal{F}} Y_R = |\mathcal{F}| \quad (1)$$

$$\sum_{v \in G_R^t} X_v \geq Y_R, \text{ for all } R \in \mathcal{F} \quad (2)$$

$$\sum_{R \in G_v^t} Y_R \geq X_v |G_v^t|, \text{ for all } v \in V_{\mathcal{F}} \quad (3)$$

and state that all pieces must be visible, that piece R is visible only if there exists $v \in G_R^t$ with $X_v = 1$, and that a guard at v sees all pieces in G_v^t , respectively. The goal is to minimize $\sum_{v \in V_{\mathcal{F}}} X_v$.

The concept of Constraint Programming (CP) was introduced in artificial intelligence and graphics in the 1960s and 1970s. For an introduction to CP see for instance [13]. Frühwirth et al [6] describe an application of CP to solve optimal antenna placement. CP languages and solvers help reduce the implementation effort and offer significant pruning of the search space through the reduction of the domains of variables by constraint propagation.

Additional constraints may be imposed here to reduce the search space. For example, for orthogonal polygons, we know that $\sum_{v \in V_{\mathcal{F}}} X_v \leq \lfloor \frac{n}{4} \rfloor$. More interesting are the constraints that result from Proposition 6. If $[m_i, M_i]$ is the interval found up to approximation step i , we may check whether it is possible to improve the lower bound m_i by imposing that $m_i < OPT_{\square}(\Gamma_{i+1}) \leq M_i$. If this condition, together with the dominance relations, renders the model inconsistent, we define $m_{i+1} = m_i$. If either $OPT_{\square}(\Gamma_{i+1}) = M_i$ or all pieces in $\mathcal{D}(\Pi_{i+1}) \setminus \Gamma_{i+1}$ are totally visible for the resulting solution, $OPT(P)$ was found. Otherwise, we may check whether the constraints (1)–(3) defined by Π_{i+1} are consistent with $OPT_{\square}(\Pi_{i+1}) \leq m_{i+1}$. If they are, $OPT(P)$ was found. If not, we compute $OPT_{\square}(\Pi_{i+1})$, but, to try to reduce the search, we impose the redundant constraint that $OPT_{\square}(\Pi_{i+1}) > m_{i+1}$. As before, we try to improve M_i by requiring also that $OPT_{\square}(\Pi_{i+1}) < M_i$, and consistently update M_{i+1} in a similar way. The approximation procedure may continue until $OPT_{\square}(\Pi_{i+1}) = m_{i+1}$, in which case we know that $OPT(P)$ was found.

3.3 The Main Steps of Our Approximation Algorithm

We summarize the main steps of the general algorithm. Suppose the n vertices of polygon P are given in counterclockwise order. Perform a decomposition Π of P (e.g., Π_{r-cut} for orthogonal polygons), such that each piece is totally visible from at least a vertex. Compute the visibility region of each vertex of P , by adapting Lee's algorithm [11,17] and keep only a maximal set of (strictly) non-dominated vertices. Determine G_v^t, G_v^s , for all $v \in \mathcal{D}(\{v_1, \dots, v_n\})$ and collect relevant information about G_R^t and G_R^s , but delay decomposition of pieces visible

by sections. Compute Γ_0^H and apply the algorithm described in the previous subsections. We note that to compute $OPT_{\square}(\mathcal{F})$ we first apply the algorithm in Fig. 3 to reduce \mathcal{F} and $V_{\mathcal{F}}$.

4 Conclusions

We give an anytime algorithm to compute the optimal solution of MINIMUM VERTEX GUARD by successive approximations. Though the presentation is focused on polygons, it may be directly extended to terrains following [4]. An experimental evaluation of the algorithm is being done, which will combine the application of standard Computational Geometry algorithms and Constraint Programming techniques.

References

1. Avis, D., Toussaint, G.: An efficient algorithm to decompose a polygon into star-shaped pieces. *Pattern Recognition* **13** (1981) 295–298.
2. Chvátal, V.: A combinatorial theorem in plane geometry. *J. of Combinatorial Theory (Series B)* **18** (1975) 39–41.
3. Edelsbrunner, H., O’Rourke, J., Welzl, E.: Stationing guards in rectilinear art galleries. *Computer Vision, Graphics, and Image Processing* **27** (1984) 167–176.
4. Eidenbenz, S.: Approximation algorithms for terrain guarding. *Information Processing Letters* **82** (2002) 99–105.
5. Fisk, S.: A short proof of Chvátal’s watchman theorem. *J. of Combinatorial Theory (Series B)* **24** (1978) 374.
6. Frühwirth, T., Brisset, P.: Optimal Placement of Base Stations in Wireless Indoor Communication Networks. *IEEE Intelligent Systems Magazine* **15(1)** (2000).
7. Ghosh, S. K.: Approximation algorithms for art gallery problems. *Proc. Canadian Information Processing Society Congress.* (1987) 429–434.
8. Johnson, D.: Approximation algorithms for combinatorial problems. *J. Comput. System Sci.* **9** (1974) 256–278.
9. Kahn, J., Klawe, M., Kleitman, D.: Traditional galleries require fewer watchmen. *SIAM J. Algebraic and Discrete Methods* **4** (1983) 194–206.
10. Lee, D. T., Lin, A. K.: Computational complexity of art gallery problems. *IEEE Transaction on Information Theory* **IT-32** (1986) 276–282.
11. Lee, D. T.: Visibility of a simple polygon. *Computer Vision, Graphics, and Image Processing* **22** (1983) 207–221.
12. Cormen, T. H., Leiserson, C. E., Rivest, R. L.: *Introduction to Algorithms*, 11st Ed., MIT Press (1994) 974–978.
13. Marriott, K., Stuckey, P.: *Programming with Constraints – An Introduction*, MIT Press (1998).
14. O’Rourke, J.: An alternate proof of the rectilinear art gallery theorem. *J. of Geometry* **21** (1983) 118–130.
15. Sack, J. R., Toussaint, G.: Guard placement in rectilinear polygons. *Computational Morphology*. G. T. Toussaint, ed., Elsevier Science Publishers (1988) 153–175.
16. Schuchardt, D., Hecker, H.: Two NP-hard problems for ortho-polygons. *Math. Logiv Quart.* **41** (1995) 261–267.
17. Urrutia, J.: Art gallery and illumination problems. In J.-R. Sack and J. Urrutia, editors, *Handbook on Computational Geometry*. Elsevier (2000).