The Visual Computer

Guarding the walls of an art gallery

Aldo Laurentini

Dipartimento di Automatica ed Informatica,
Politecnico di Torino, Corso Duca degli Abruzzi 24,
10129 Torino, Italy
E-mail: laurentini@polito.it

The standard "art gallery" problem consists in stationing a minimum set of guards in a polygon P so that each point of P is seen by at least one guard. We introduce and explore the edge-covering problem; the guards are required to observe the edges of P; metaphorically the paintings on the walls of the art gallery, and not necessarily every interior point. We compare minimum edge and interior covers for a given polygon and analyze the bounds and complexity for the edge-covering problem. We also introduce and analyze a restricted edge covering problem, where full visibility of each edge from at least one guard is required. For this problem we present an algorithm that computes a set of regions where a minimum set of guards must be located. The algorithm can also deal with the external visibility of a set of polygons.

Key words: Computational geometry – Visibility – Illumination – Visual sensor placement – Art gallery problem and theorems – Edge covering

The Visual Computer (1999) 15:265–278
© Springer-Verlag 1999

1 Introduction

Problems of visual sensor placement arise in many computer vision areas such as object recognition or reconstruction, robotic tasks, and area surveillance. Although these are generally 3D problems, in some cases, such as the surveillance of buildings, we can restrict our considerations to 2D areas.

The research in this area was triggered by Chvátal's "art gallery" Theorem (1975). He proved that at most $g(n) = \lfloor n/3 \rfloor$ point guards are required for covering a simple polygon P with n edges. The word polygon specifies a closed set including interior and boundary points. Two points of a polygon P are *visible* from each other if the line segment connecting the points lies entirely in P. A polygon P is *covered* by a set of viewpoints, or *guards*, lying in P, if each point in P is visible from at least one guard. Guards can lie in restricted positions (the vertices for instance); if not, they are *point* guards.

The upper tight bound $g(n,h) = \lfloor (n+h)/3 \rfloor$ for covering polygons with h polygonal holes and n edges, a conjecture for about 10 years, has been independently proven by Bjorling-Sachs and Suvaine (1995), and Hoffman et al. (1991).

Many variations or restricted versions of the problem have been considered. The reader is referred to O'Rourke (1987), Shermer (1992) and Urrutia (1996) for comprehensive surveys. Most of the problems in this area can also be interpreted as illumination problems, where point lights replace the guards.

In spite of the success in finding the worst-case numbers of guards, the *minimum cover problem*, that is, the practical problem of stationing a minimum set of $G(P)$ guards in a given polygon P, is still open. We use the terms ICH and IC to refer to the minimum cover problems for the interior of polygons with and without holes.

O'Rourke and Supowit (1983) have shown that the *decision* problem corresponding to ICH (given P, is there a cover of P with k or fewer guards?) is NP-hard. The same result for the IC has been obtained by Lee and Lin (1986).

Algorithms have been constructed for many non-polynomial problems. They work sufficiently well with the usual inputs, even though some unlikely worst case might occur. This is not the case for IC and ICH: at present, no exact algorithm for finding and locating a minimum set of guards in a given polygon has been found. Another possible approach is to construct approximate algorithms that can supply solutions close to the optimal. Also this approach seems unable to cope with the elusive nature of

IC and ICH, since the current approximate guard-placing algorithms have no guaranteed performances int the worst case (Laurentini 1996).

The problem addressed in this paper is different from the classic IC and ICH. We call it the *edge-covering problem*, since our minimum set of $GE(P)$ guards must only cover the *edges* of the polygon P, neglecting possible uncovered internal regions. This problem is in keeping with the name "art gallery", since the surveillance of paintings on the walls is the main task of the guards. In addition, if our problem deals with positioning point lights, the main concern is to illuminate the walls.

Even though covering the edges appears a sensible variation of the classic art gallery problem, it has received very little attention. To the author's knowledge, only O'Rourke (1987) mentions edge covering, when he acknowledges the distinction between the two covers.

The content of this paper is as follows. Worst cases, relations with interior covering, and complexity of the edge-covering problems are analyzed in Sect. 2 for polygons with and without holes. In Sect. 3 we introduce the restriction that each edge be *entirely visible* from at least one guard, analyze the worst cases of this subproblem, and compare it with unrestricted edge covering. In Sect. 4 we present an algorithm for entire visibility edge covering and its complexity analysis. In Sect. 5 we modify the algorithm for external visibility (or illumination) of polygons. In Sect. 6 we summarize the results obtained and discuss further lines of research.

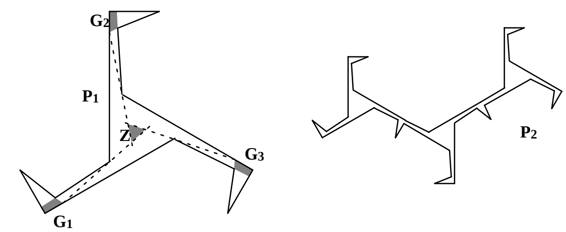

Fig. 1. Three guards are sufficient for covering the edges of polygon P_1, but four are required for covering its interior. For P_2, respectively 7 and 10 guards are required

2 Edge cover: comparison with interior cover, bounds, and complexity

In this section we compare minimum edges and interior covers for a given polygon and discuss bounds and complexity for the edge covering problem. Let EC and ECH indicate the edge-covering problem for polygons without and with holes and the term *wall guards* indicate the guards required for this purpose (we cannot use the term "edge guards", since in "art gallery" parlance it means guards restricted to the edges).

2.1 Minimum edge and interior covers for a given polygon

Let $GE(P)$ be the minimum number of wall guards necessary for covering (the edges of) a polygon P. Obviously $GE(P) \leq G(P)$, since observing the interior implies observing the edges. It is not difficult to find examples for which the minimum covers are different. For polygon P_1 (Fig. 1), three guards G_1, G_2, and G_3, located anywhere in the areas highlighted, form the minimum edge cover, but another guard is necessary for covering the central area Z. The difference $G(P) - GE(P)$ can assume any value (see polygon P_2 in Fig. 5, where the difference is three).

For the ratio $G(P)/GE(P)$ we have found the following result:

For polygons with holes, $G(P)/GE(P)$ can assume any value. This statement is proved by the example in Fig. 2. The polygon in the figure belongs to a family where $GE(P)$ is two and $G(P)$ is $O(n)$.

For polygons without holes, the situation seems to be different. Although we are not able to produce a proof, we conjecture that $G(P)/GE(P) \leq 2$. A family of polygons such that the ratio $G(P)/GE(P)$ is arbitrarily near 2 is shown in Fig. 3.

Since the interior-covering problem is still open, an important question is: if we succeed in finding a solution to the edge-covering problem for a given polygon, can we use the wall guards as a basis for a minimum interior cover?

In some cases, a minimum set of wall guards also covers the interior of P, and obviously it is a minimum set for interior covering as well. For the case in which the wall guards do not cover the interior, the examples of Figs. 1–3 may suggest that, when a solution for the EC is known, a solution for

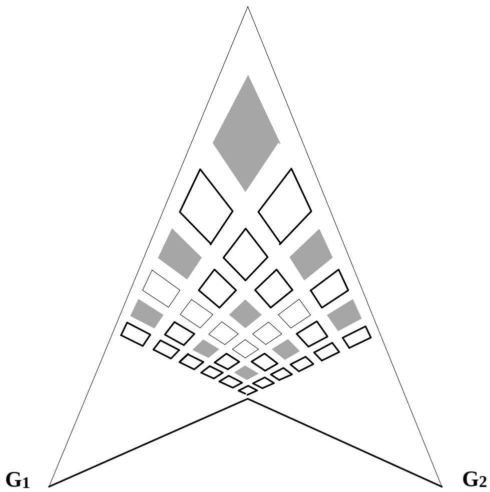

2

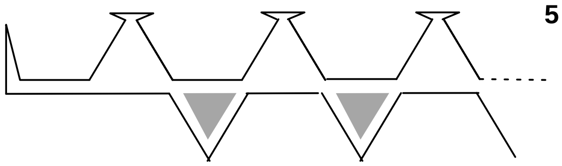

3

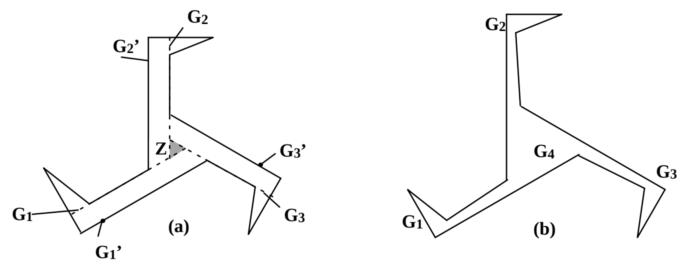

4

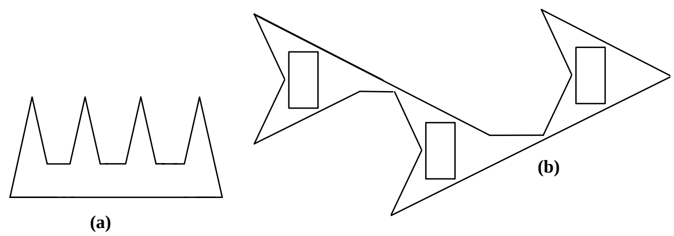

5

Fig. 2. Two guards G_1 and G_2 are sufficient for edge cover, but covering the interior requires one additional guard for each area highlighted

Fig. 3. A family of polygons such that the ratio $G(P)/GE(P)$ can be arbitrarily near to 2

Fig. 4a,b. These examples show that, in general, a solution for the IC cannot be obtained from a solution for the EC (**a**), or vice versa (**b**), by deleting or adding some guards

Fig. 5a,b. Worst-case polygons **a** Chvátal's; **b** Shermer's. The guards in the figure are minimum sets for both interior and edge covering

the IC could be obtained by adding some guards, or an EC solution could be obtained by deleting some guards of an IC solution. However, in general, it is impossible to establish such simple relations between a minimum edge and interior covers.

In fact, let us consider the polygon in Fig. 4a. In this case, G_1, G_2, and G_3 are a solution for the EC, but not for the IC, since region Z is not covered. However, a minimum set of guards for the IC is not obtained by adding a fourth guard for covering Z: only three guards such as G_1', G_2' and G_3' are sufficient. Conversely, let us consider the polygon in Fig. 4b. It requires four guards for the IC and three for the EC. The four guards of Fig. 4b are a solution for the IC, but no subset of three guards covers the edges.

Even though interior and edge covers appear to be close relatives, general simple ways for transforming a solution for one problem into a solution for the other do not seem to exist.

The situation could be different if we were content with an approximate solution with guaranteed performance. If $G(P)/GE(P) \leq 2$ is true for polygons without holes, a minimum set of guards for the IC is also a solution for the EC, which contains at most twice as many guards as are strictly necessary.

2.2 "Art gallery" theorems for edge cover

Let $ge(n)$ and $ge(n, h)$ be the worst-case minimum numbers of guards necessary to cover the edges of polygons with n edges without holes and with h holes, respectively. We have:

$-ge(n) = g(n) = \lfloor n/3 \rfloor$
$-ge(n,h) = g(n,h) = \lfloor (n+h)/3 \rfloor$.

In fact, $ge(n) \leq g(n)$, otherwise there would be some polygons (the worst-case polygons for the EC) that require more guards for edge covering than for interior covering. Chvatal's (1975) comb polygon (Fig. 5a), a worst case for the IC, is also a worst case for the EC, and thus the equal sign holds. The equation $ge(n,h) \leq g(n,h)$ must also hold, otherwise a contradiction follows. It is easy to see (Fig. 5b) that the polygons used as worst cases for the ICH (O'Rourke 1987) also requires $\lfloor (n+h)/3 \rfloor$ guards for the ECH, and thus the bound is tight.

2.3 Computational complexity of edge covering

Many covering and decomposition problems have been shown to be NP-hard: see O'Rourke (1987), Shermer (1992) and Culberson and Reckhow (1994). We have found that *the EC and the ECH are both NP-hard*. This can be shown by verifying that the proofs given by Lee and Linn (1986) for the IC and by O'Rourke and Supowit (1983) for the ICH also apply to the EC and the ECH. Both proofs are based on a reduction from the NP-complete 3-SAT problem. Given an instance of 3-SAT (a boolean expression formed by the product (AND) of sums (OR) of exactly 3 variables), the proofs construct in polynomial time polygons that can be covered by k guards if and only if the boolean expression is satisfiable (that is, there is an assignment of TRUE and FALSE values for the variables such that the whole expression is TRUE). Although the constructions cannot be reported here in full detail, we sketch the general ideas and show that they also apply to the edge-covering problems.

2.4 Polygons with holes

Let us consider the construction for polygons with holes and verify that it is also a reduction from 3-SAT to the ECH. For each boolean variable, the polygon P has a region shaped as a staircase corridor such that a guard located near a corner can fully observe the interior (or the walls) of two adjacent straight segments of the corridor (Fig. 6a). Let us label the corners with consecutive integers. Two minimal arrangements of guards are possible to cover the interior (or the walls) of the whole corridor. One arrangement, let us say that with guards located at the *odd* corners, corresponds to the TRUE assignment (Fig. 6b); the other, with guards located at the *even* corners, corresponds to the assignment FALSE (Fig. 6c).
It can be shown that the corridors can be bent and crossed without affecting the interior coverage properties. It is easy to see that this is also true for edge covering. The corridor is bent to form a loop for each variable. We must now construct, for each sum of three terms, a new part of P that is observed by the guards in the corridors if and only if the sum is TRUE.
Let us consider the sum $\bar{x}_1 + x_2 + x_3$. We construct a triangle (Fig. 6e) that shares three small areas of a corner with one corner of each corridor corresponding to the variables x_1, x_2, and x_3. In this case, the corner of the corridor corresponding to x_1 must be an even corner, which thus has a guard in it if and only if x_1 is FALSE, and the corners corresponding to x_2 and x_3 must be odd. With this construction, one or more corridor guards cover the interior (and the walls) of the triangle if and only if the sum $\bar{x}_1 + x_2 + x_3$ is TRUE. By repeating this construction for all the sums, we obtain the complete polygon whose edges are covered by k guards (half of the corners of the corridors) if and only if the boolean expression is satisfiable. O'Rourke and Supowit have shown that the construction can be performed in $O(mn)$ time with m sums and n variables.

2.5 Polygons without holes

Let us consider the proof by Lee and Linn for polygons without holes. First, they reduce 3-SAT to the IC with guards restricted to the vertices. Let us verify that this is also a reduction to the EC with guards restricted to the vertices. The upper part of the polygon they construct has m large concavities, one for each sum in the boolean expression (Fig. 7a). Three vertex guards for each concavity are required for observing the interior (and the walls) of the three spikes. There are two possible positions for each guard. Positions t_1, t_2, t_3 and f_1, f_2, f_3 correspond to TRUE and FALSE values of the first, second, and third literal of the sum. For instance, if the sum is $\bar{x}_1 + x_2 +$

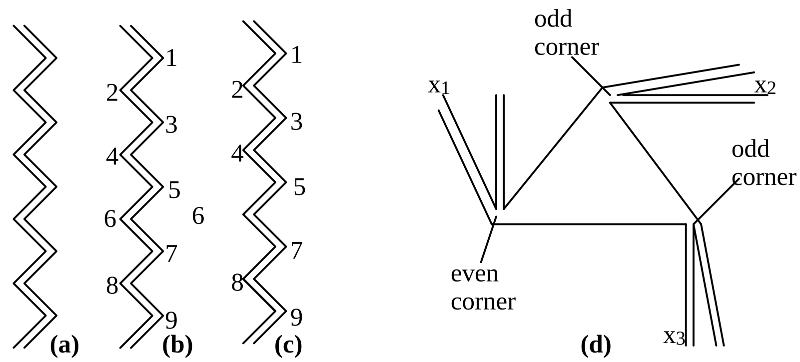

6

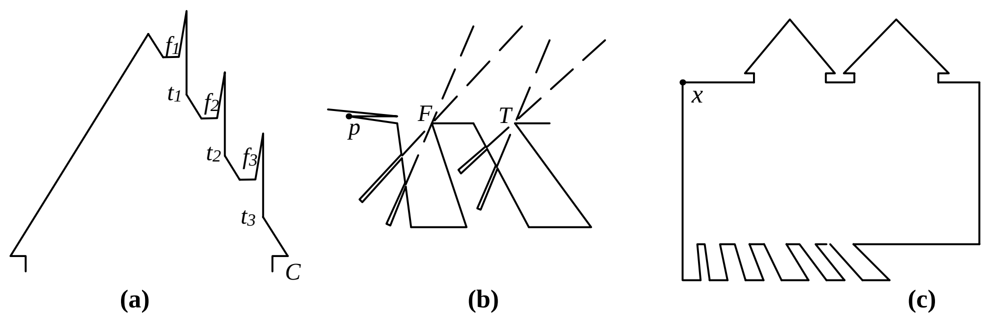

7

Fig. 6a. The staircase corridor; **b, c** The TRUE and FALSE assignments; **d** The walls of the triangle are visible from the corridor guards if and only if $\bar{x}_1 + x_2 + x_3$ is TRUE

Fig. 7a. The triangular concavity emulating a sum; **b** The pattern that forces the consistency in the sums of the truth assignment of one variable; **c** The final polygon for two sums

x_3 and $x_1 = x_2 = 0$ and $x_3 = 1$, the guards are located at t_1, f_2, and t_3. At least one guard must be in a t_i position, otherwise parts of the concavity (and of its walls), as the corner C, cannot be observed by these guards. Thus the cover of this concavity is minimal if and only if the sum is TRUE.

The lower part of the polygon forces the coverage to be minimal if and only if the truth assignment of the variables are the same in the various sums (Fig. 7b). For each variable there are two "wells", each containing a set of spikes, and another spike, containing point p, requiring one guard located at point F or point T to be observed. The first well contains a set of spikes that are covered if and only if the variable is FALSE in all sums. In fact, the dotted lines join F to all the positions of the guards in the upper concavities corresponding to the assignment FALSE to the variable. The second well has a similar purpose for the assignment TRUE. Figure 7b shows the wells in the case of a variable present in two sums. From this arrangement it follows that, if and only if a variable is TRUE (FALSE) in all sums, the upper guards cover the interior and walls of the spikes in the right (left) well. Further, only one other guard located at F (T) is required for covering the interior and walls of the spikes of the other well and the spike containing p.

Therefore, if and only if an assignment of the variables exists that makes each sum TRUE, the minimal vertex cover of both the interior and walls consists of $K = 3m$ upper guards plus n guards for the bottom spikes plus one other guard at vertex x for the interior and walls of the uncovered wells. A sketch of the whole polygon for two sums is shown in Fig. 7c (the spikes are omitted for simplicity). It is easy to see that the construction takes polynomial time.

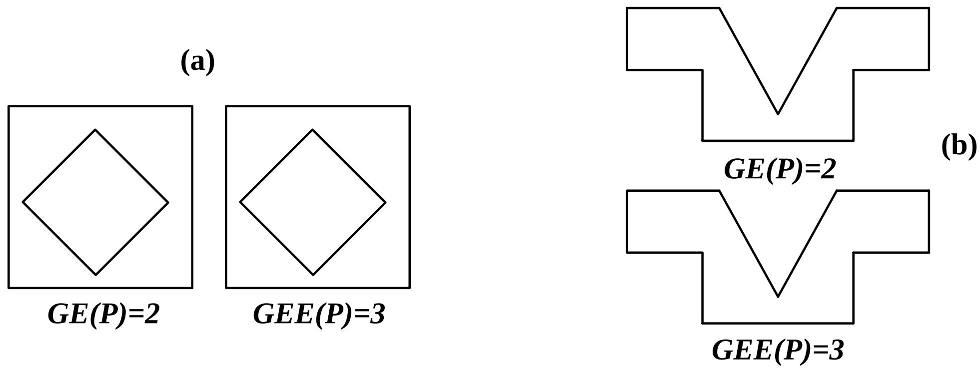

8

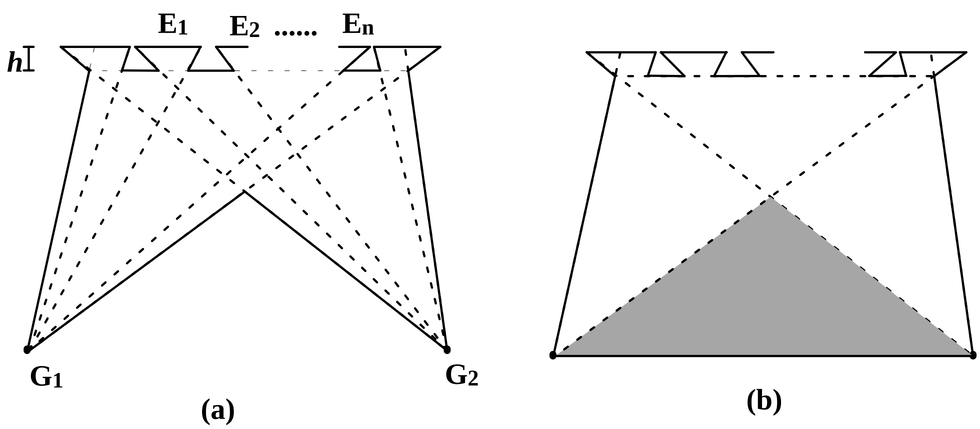

9

Fig. 8a,b. Two examples for which the entire edge cover requires one more guard. The points indicate possible locations of the guards

Fig. 9a,b. The examples show that $GEE(P)/GE(P)$ is $O(n)$ for polygons without and with holes

All the arguments of Lee and Linn are based on points lying on the edges, and they also apply to any EC with vertex guards, which thus is NP-complete. Finally, the EC with point guards can be shown to be NP-hard by Aggarwal's argument (Lee and Linn 1986).

3 The entire visibility edge-covering problem

Now we introduce a restriction that could make practical sense for the edge-covering problem. We require that each edge must be seen in its entirety by at least one guard. Let the EEC and the EECH be the problems of finding a minimum cover of this kind for polygons without and with holes. In this section we compare these problems with unrestricted edge cover, and we discuss their worst cases and computational complexity.

3.1 Entire edge cover versus edge cover for a given polygon

Let $GEE(P)$ be the minimum number of guards required for the entire edge cover of a polygon P. Obviously $GEE(P) \geq GE(P)$. In Fig. 8 we show a polygon with holes and another without holes where these minimum numbers are different. Connecting polygons such as those shown in the figure, we can easily obtain cases where $GEE(P) - GE(P)$ is $O(n)$.

What about the ratio $GEE(P)/GE(P)$? Were it bounded, a solution of the restricted problem could

be an approximate solution with guaranteed performance for edge cover. Unfortunately, this is not the case. The ratio is unbounded for polygons both with and without holes as shown by the examples of Fig. 9. Two unrestricted wall guards G_1 and G_2 are sufficient in both cases, but each upper edge E_1, $E_2, \ldots$ E_n requires a different guard to be observed entirely. If the height h is reduced, the example can be extended to any number of concavities, which shows that $GEE(P)/GE(P)$ can assume any value. The modified polygon in Fig. 9b, with tiny triangular holes conveniently located in the shaded area, shows that $GEE(P)/GE(P)$ could be $O(n)$ for polygons with holes also.
Observe that $GEE(P)$ could be equal (Fig. 5a), greater (Fig. 8) or smaller (Fig. 2) than $G(P)$.
Finally, in general one solution for the EEC cannot be obtained from a solution for the EC, or vice versa, by deleting or adding guards, as Fig. 8b may suggest. This is shown by the example in Fig. 8a. No minimum set containing $GEE(P) = 3$ guards can be obtained by adding one guard to the $GE(P) = 2$ wall guards shown, and no two guards subset of the $GEE(P) = 3$ guards shown covers the edges.

3.2 Art gallery theorems for entire edge cover

Let $gee(n)$ and $gee(n,h)$ be the worst-case minimum numbers of guards required by the EEC and EECH. For polygons without holes we have found that

$$gee(n) = ge(n) = g(n) = \lfloor n/3 \rfloor .$$

Proof. First, $gee(n) \leq \lfloor n/3 \rfloor$. This is shown by the classic proof of the art gallery theorem given by Fisk (1978). Fisk triangulates P without adding new vertices and exploits the existence of a three-coloring of the vertices. At least one color affects a set of $m \leq \lfloor n/3 \rfloor$ vertices. These vertices can be selected as interior guards, since each triangle of the triangulation has a vertex belonging to this set. Since each edge belongs entirely to one triangle, it is entirely observed by at least one of these guards, which therefore are also entire wall guards. The Chvátal comb polygon is also a worst case for the EEC, thus the bound is tight.

Finding tight bounds for polygons with holes is more difficult. For the case of one hole, we have found that

$$gee(n,1) = \lfloor (n+2)/3 \rfloor .$$

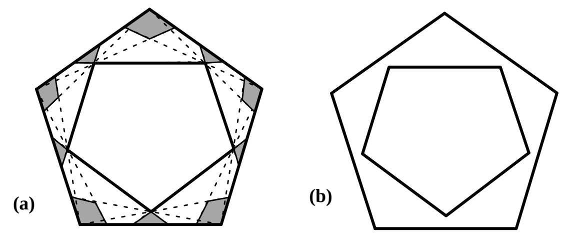

Fig. 10a. Three edges at most can be entirely observed by guards located in the highlighted regions. **b** At least four guards are required to observe the edges

Proof. It is easily seen that

$$\lfloor (n+h)/3 \rfloor = ge(n,h) \leq gee(n,h)$$
$$gee(n,h) \leq \lfloor (n+2h)/3 \rfloor .$$

In fact, $ge(n,h) \leq gee(n,h)$, otherwise there would be polygons (the worst case polygons for the EC) that require more guards for unrestricted edge covering than for restricted edge covering. The bound $\lfloor (n+2h)/3 \rfloor$, first established by O'Rourke (1987) for the ICH and combinatorial guards, is easily seen to apply to our case. By inserting h zero width "bridges" between vertices belonging to unconnected parts of the boundary, a polygon P with h holes is reduced to a simply connected polygon containing all the original edges. This allows us to use the $ge(n)$ bound for a polygon with $n+2h$ vertices.

For the case of one hole, $\lfloor (n+2)/3 \rfloor$ guards are sometimes necessary, as shown by the example in Fig. 10. No guard in this polygon can entirely observe four edges or more. It is possible to entirely observe three edges from the regions highlighted in Fig. 10a. It follows that at least four guards are required for observing the 10 edges (Fig. 10b). Observe that the bound $\lfloor (n+h)/3 \rfloor = 3$ would be incorrect in this case. The polygon is a member of a family composed of a regular polygonal hole inside a larger regular polygon with the same number of edges. If the polygons are sufficiently close and $n = 2m = 3p+1$, with m and p integers, $\lfloor (n+2)/3 \rfloor$ guards are required. The only triangulation possible for these polygons is an example of Shermer's *though triangulation* (O'Rourke 1987).
Is this result in agreement with the bound $\lfloor (n+h)/3 \rfloor$ for both the ICH and the ECH? The proofs

given by Bjorling-Sachs and Suvaine (1995) and Hoffman et al. (1991) cannot be extended to the EECH. Hoffmann et al. station each guard at the intersection of at most three convex polygons. In general, each convex polygon covers only a part of an edge, thus Hoffmann's guards could be insufficient for entire edge covering. The proof by Bjorling-Sachs and Suvain reduces the original polygon to a polygon without holes and $n+h$ vertices. Each step of the proof deletes a hole using one of three possible constructions: two of them split an edge in two, which in the new polygon is not guaranteed to be observed entirely by the same guard. The polygon presented in Fig. 10 allows only the edge-splitting constructions.

For $h \geq 2$, we have not been able to find polygons requiring more than $\lfloor (n+h)/3 \rfloor$ guards, even if they admit only the edge-splitting construction, and we conjecture that this bound is tight.

3.3 Computational complexity of the EEC and the EECH

The reduction of 3-SAT to the EC and the ECH also holds for the EEC and the EECH. In fact, both proofs construct polygons where each edge is observed entirely by at least one guard.

4 An algorithm for entire edge cover

The entire visibility restriction allows us to discretize the edge-guarding problem. In this section we describe an algorithm for both the EEC and the EECH. Given a polygon P, the algorithm computes a set of polygonal regions. A minimum set of $GEE(P)$ guards can be obtained, independently locating one guard anywhere in each region. A minor addition makes the algorithm suitable for the "fortress" or external guarding problem, where the edges of the holes are observed from an unbounded region.

The algorithm consists of the following steps:

Step 1. Compute a partition Π of P into regions Z_i such that:

- The same set $\mathbf{E}_i = (E_p, E_q, \dots, E_t)$ of edges is completely visible from all points of $Z_i \forall i$.
- The Z_i are maximum regions, i.e., $\mathbf{E}_i \neq \mathbf{E}_j$ for contiguous regions.

Step 2. Select the *dominant regions*. A region Z_i is defined to be dominant if there is no other region Z_j of the partition such that $\mathbf{E}_i \subset \mathbf{E}_j$.

Step 3. Select an optimal (or minimum) solution. A minimum solution consists of a set of dominant regions $\mathbf{S}_j = (Z_{j1}, Z_{j2}, \dots, Z_{jk})$ that covers $\mathbf{E} = \cup \mathbf{E}_i$ with the minimum number of members.

Observe that there could be a minimal solution containing nondominant regions. For instance, in Fig. 10a the dominant regions are those highlighted. The upper guard in Fig. 10b does not lie in a dominant region. We consider only sets of dominant regions for two reasons. First, a nondominant region can be replaced by a dominant region covering the same edges and some others. Multiple coverage of edges is preferable, for instance in the case of sensor failure. Second, we are looking for one optimal solution, not for all optimal solutions. Considering only the dominant regions reduces the computations in step 3, exponentially in the worst case.

It is worth noting that partition Π could be obtained by intersecting the *complete visibility polygons* of each edge (the term *strong visibility* is used by Shermer (1992) with the same meaning). Complete visibility polygons can be constructed in linear time for polygons without holes (O'Rourke 1986), but unfortunately, to the author's knowledge, no algorithm has yet been presented for polygons with holes.

We now present the details of the algorithm and the complexity analysis.

4.1 Step 1. Computing partition Π

We divide P into maximal regions Z_i from which the same set of edges $\mathbf{E}_i$ is entirely visible, and label each region Z_i with the set $\mathbf{E}_i$, using a visiting algorithm.

Let us discuss which lines are relevant to Π. Obviously, the lines supporting the edges are necessary. When a guard crosses line L_a as shown in Fig. 11a, the supported edge E_i becomes entirely visible. Taking into account occlusions requires other categories of lines. For instance, if a guard crosses line L_b in Fig. 11b, vertex v_k becomes visible. If there are no other occlusions, as in Fig. 11b, the entire edge E_i to which the vertex v_k belongs becomes visible. Otherwise, as in the case of Fig. 11c, the edge becomes only partially visible. Thus line L_b is potentially relevant to Π.

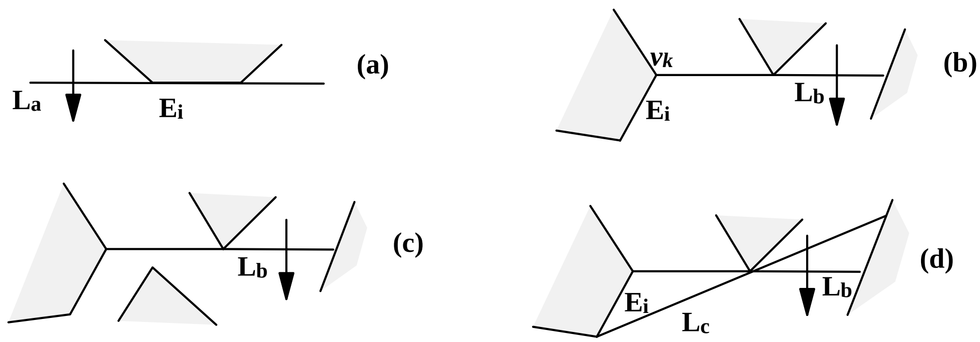

11

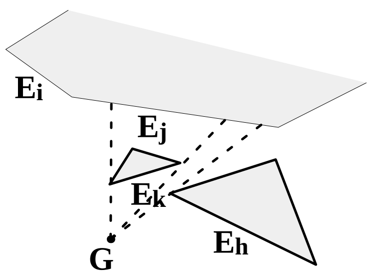

12

Fig. 11a–d. Examples of lines that change the visibility condition of an edge

Fig. 12. Edge E_j has three occlusions in the aspect relative to G

For dealing with such cases, we compute Π as a refinement of a more detailed partition Π', which also contains potentially relevant lines such as L_b and auxiliary lines. These lines are not relevant to Π, but change the state of occlusion of an edge, as for instance L_c in Fig. 11d.

Before defining partition Π', let us first define the *aspect* $\mathbf{A}(G)$ of a point G belonging to P.

$$\mathbf{A}(G) = ((E_h, n_h), (E_k, n_k), \ldots, (E_q, n_q)),$$

where $E_h, E_k, \ldots, E_q$ are the edges fully or partially visible from G, and $n_h, n_k, \ldots, n_q$ are the numbers of occlusions of these edges. For instance, the aspect relative to point G in Fig. 12 is:

$$\mathbf{A}(G) = ((E_i, 0), (E_j, 3), (E_k, 0), (E_h, 0), \ldots)$$

Π' is defined as the partition that divides P into regions Z'_i such that:

- All points of Z'_i have the same aspect $\mathbf{A}_i$.
- Z'_i are maximum regions, i.e., $\mathbf{A}_i \neq \mathbf{A}_j$ for contiguous regions.

Clearly, to belong to the same region of Π' is a necessary, but insufficient, condition for two points to belong to the same region of Π.

Computing Π' is strictly related to the computation of the *aspect graph* of a set of polygons, a problem not discussed in the literature. The vertices of the aspect graph represent all the topologically distinct views, or aspects, of an object. The graph structure of the aspect graph is the dual graph of the partition of the viewing space into regions whose points share the same aspect. The edges of the graph represent *visual events*, that is, qualitative or topological changes in the aspects of adjacent regions. Gualtieri et al. (1989) give algorithms for computing aspect graphs in two dimensions for one convex polygon and in three dimensions for polyhedra, solids of revolution, and other curved objects. For further details, the interested reader is referred to Gigus et al. (1991) and Plantinga and Dyer (1990).

We now present the catalogue of lines that form partition Π' and the associated visual events or changes in the aspects. A line is said to be *active* if it contains one or more *active segments*. An active segment is the boundary between points whose aspects are different. It is not difficult to verify that all active lines are those that join two vertices in the cases given in Table 1. The active segments are highlighted with a thick line. The arrows mark the *positive crossing directions*. The *positive visual event* is the change

	Active lines	Positive visual event	Negative visual event
a	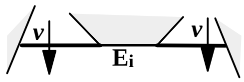	Add (E_i, 0)	Delete (E_i, 0)
b	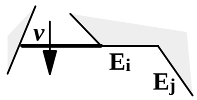	$n_j \leftarrow n_j - 1$ Add (E_i, 0)	$n_j \leftarrow n_j + 1$ Delete (E_i, 0)
c	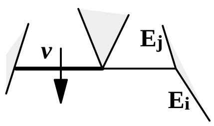	$n_i \leftarrow n_i - 1$ if E_j is already in the aspect: $n_j \leftarrow n_j + 1$; otherwise, add (E_j, 1)	$n_i \leftarrow n_i + 1$ if $n_j = 1$: delete (E_j, 1); otherwise, $n_j \leftarrow n_j - 1$
d	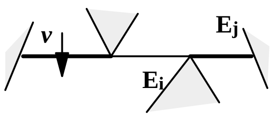	$n_i \leftarrow n_i - 1$ if E_j is already in the aspect: $n_j \leftarrow n_j + 2$; otherwise, add (E_j, 2)	$n_i \leftarrow n_i + 1$ if $n_j = 2$: delete (E_j, 2); otherwise, $n_j \leftarrow n_j - 2$
e	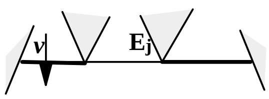	if E_j is already in the aspect: $n_j \leftarrow n_j + 1$; otherwise, add (E_j, 1)	if $n_j = 1$: delete (E_j, 1); otherwise, $n_j \leftarrow n_j - 1$
A←B means that B replaces A in the aspect.			

Table 1. Catalogue of lines for partition Π'

of aspect of a point which crosses the active segment along the positive direction; a similar definition holds for the *negative visual event*. For simplicity, in cases d and e, each with two active segments, we only indicate the visual event of the left segment, since the other is obtained by changing the subscripts. Observe that:

– The positive visual events of cases a and b always produce full visibility of a new edge.
– Positive visual events of cases c and d produce full visibility of a new edge only if the previous number of occlusions of this edge is one.
– Visual events of case e only affect the occlusion numbers.

Computing Π requires the following substeps:

1a. Computing the active segments of Π'
1b. Constructing Π'
1c. Refining Π' into Π with a visiting algorithm.

1a. *Computing the active segments.* To find the active lines we must consider $O(n^2)$ pairs of vertices (v_i, v_j). As is easily seen from Table 1, two conditions must be satisfied for a line to be active:

1. The line must not cross the boundary of P at both v_i and v_j (examples of vertices that do not produce active lines are shown in Fig. 13).
2. Segment $v_i v_j$ must belong entirely to P.

Condition 1 can be checked in constant time for each pair (v_i, v_j). Condition 2 can be verified and the active segments can be found in $O(n^2)$ time with Welzl's algorithm for constructing the *visibility graph* of a set of line segments (O'Rourke 1987), i.e., the graph whose edges represent visibility among the vertices of the segments. The algorithm requires $O(n^2)$ time for a topological sort of the $(n(n-1))/2$ slopes of the oriented lines joining each pair of vertices (one line for each pair), and an angular sweep that finds the first segment crossed (if any) in constant time for each oriented line. For our purposes we need only perform a second angular sweep with reversed directions of the lines. We mark each active segment with the corresponding visual event, which

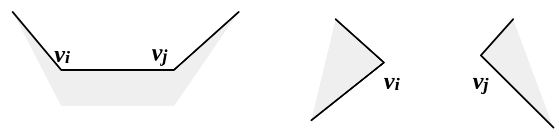

Fig. 13. Examples of vertices visible to each other that do not produce active lines

can be computed in constant time. Thus, $O(n^2)$ active segments are obtained in $O(n^2)$ time.

1b. *Constructing partition* Π'. A data structure containing vertices, edges, and regions of partition P' and pointers representing inclusion could be constructed by Edelsbrunner's (1987) classic algorithm in $O(n^4)$ time. The partition can also be constructed by a plane sweep algorithm in $O(p \log p)$ time, where p is the number of vertices of the partition (regions and edges also are $O(p)$). This approach, which makes the computation time of the algorithm sensitive to the size p of the output, has been applied by Gigus et al. (1991) for constructing a partition of the plane with segments of straight lines and conics. We refer the reader to the original paper for further detail. For each edge of Π' lying on an active segment, we store the positive direction and the visual event in the data structure.

The total time for computing partition Π' is $O(n^2 + p \log p)$.

1c. *Computing* Π' *from* Π. There can be adjacent regions in Π' with the same set of completely visible edges. These regions are separated by: (1) lines of type e, which only affect the occlusion numbers and (2) some lines of types c and d. To obtain Π, these construction lines must be removed and the regions separated by these lines must be merged. While traversing Π', partition Π is computed as follows.

We start at an arbitrary region, compute its aspect, and traverse Π' with a depth-first search on the dual graph. The aspect of each new region traversed is obtained by updating the aspect of the previous region with the positive or negative visual events at each boundary, according to the crossing direction. Thus, at each boundary of the partition, we are able to verify whether the visual event of an edge actually creates (or deletes) a fully visible edge. If this is not the case, we update the structure of the partition by deleting the edge and merging the two adjacent regions. Each region is marked with the set of fully visible edges, and each edge with the positive crossing directions, which indicate the region with a larger set of fully visible edges.

Computing the aspect of the starting region is equivalent to solving a 2D hidden line problem, and takes $O(n^2)$ time (Preparata and Shamos 1985). The time required for traversing the partition is $O(p)$ (Baase 1988). The dimension of each aspect is $O(n)$; therefore we can update the aspect and store the fully visible edges at each boundary crossing in $O(n)$ time. Updating the structure of the partition takes a constant time at each merging. The total time for visiting Π', computing Π, and storing all aspects is $O(n^2 + pn)$.

The overall time bound of step 1 is $O(n^2 + p \log p + pn)$.

4.2 Step 2. Computing the dominant regions

To find d dominant zones, we must compare the sets of fully visible edges $\mathbf{E}_i$ of each region Z_i. This process can be shortened if we observe that:

1. A necessary condition for a region Z_i of Π to be dominant is that $\mathbf{E}_j \subset \mathbf{E}_i$ for all the regions Z_j adjacent to Z_i, or, in other words, all the positive crossing directions of the edges of Z_i lead to the interior of the region (except for the edges of P).
2. All the edges of Z_i (except for the edges of P) being due to cases a and b of Table 1, i.e., they lie on lines supporting edges of P, condition 1 is also sufficient.

The first statement is obvious. For the second, it is sufficient to observe that, if the edges of Z_i lie on the lines supporting the edges $E_p, E_q, \ldots, E_k$, no other region is able to observe this set of edges. Thus, to find the dominant regions:

- First visit all regions of Π and check condition 1 for selecting c candidate dominant regions in $O(p)$ time. Regions also satisfying condition 2 are immediately recognized as dominant.
- Perform $O(c^2)$ comparisons in $O(nc^2)$ time to select d dominant zones.

Steps 1 and 2 of the algorithm require $O(n^2 + p \log p + pn + nc^2)$ time. In Fig. 14 we show both partition Π' and Π for a polygon with nine edges and one hole and the five resulting dominant regions.

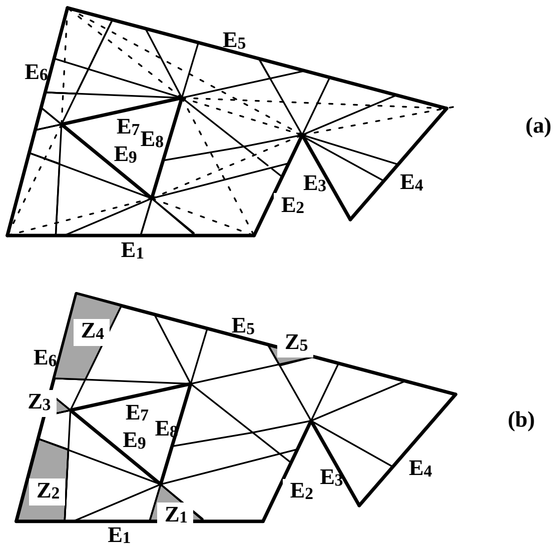

Fig. 14a,b. Partition $\mathbf{\Pi}'$ (**a**) and partition $\mathbf{\Pi}$ (**b**) of a polygon. The active segments of $\mathbf{\Pi}'$ are *solid*; the remaining parts of the active lines are *dotted*. Five dominant zones result. Three of them (Z_1, Z_3, Z_5) are immediately identified by the positive crossing directions of the edges

4.3 Step 3. Finding an optimal solution

An optimal (or minimum) solution is a set of dominant regions that covers **E** with the minimum number of members. This is an instance of the well-known set-covering problem. In general, given a set **S** and a number of subsets, an optimal cover is a set of subsets with a union **S** that minimizes the sum of the costs of the subsets. In our case, all costs are equal. The corresponding decision problem (is there a cover with k subsets or less?) is NP-complete (Parker and Rardin 1988).

Numerous practical situations have been modeled as set-covering problems, and a number of algorithms for set covering have been presented (see for instance Salkin and Mathur 1989). When, as in our case, only one minimal solution is required, much can be pruned. This is, for instance, the case of an algorithm developed for the minimization of switching functions; details can be found in Muroga (1979), and the complexity $(O(n2^d))$ has been determined by Laurentini (1996).

A nearly optimal solution obtained with a polynomial selection algorithm could be an interesting alternative. Such a solution can be obtained with a greedy heuristic, which selects the region covering the largest number of uncovered edges each time. A straightforward implementation of this algorithm is $O(np^2)$. Although its performance cannot be guaranteed to be data independent, it does not depends on the number of edges n. Let $GEE^G(P)$ be the number of regions obtained by the greedy algorithm and let r be the largest number of edges observed by a dominant region of P. It can be shown that (Nemhauser and Wolsey (1988):

$$GEE^G(P)/GEE(P) \leq 1 + \lg(r)$$

Observe that r could be small even for very large values of n.

5 The case of exterior visibility

A small addition makes Table 1 also fit for the case where only the polygonal holes are left and the region where the guards can be located is unbounded. This has been called the fortress problem when only one polygon is present (O'Rourke 1987). In this case, the active segments could be unbounded. Inspecting Table 1, we can verify that this only affects case d. If the right active segment is unbounded, the positive and negative visual events diminish to the parts affecting E_i. Thus, an enhanced table including the new entry d' shown in Fig. 15 also holds for the case of exterior visibility. The rest of the algorithm is unaffected.

6 Summary and discussion

We have proposed and explored the problem of covering the edges of a polygon with a minimum set of guards as an alternative to the classic art gallery problem, which requires complete interior cover. The minimum edge and interior covers have been compared for a given polygon. Even though a minimum set of wall guards could also cover the interior, for polygons with holes, the interior guards could be $O(n)$ times the wall guards. For polygons without holes, we conjecture that interior guards are at most twice the number of the edge guards. No simple rule seems to exist for obtaining a minimum set of interior guards starting from a minimum set of edge guards or vice versa.

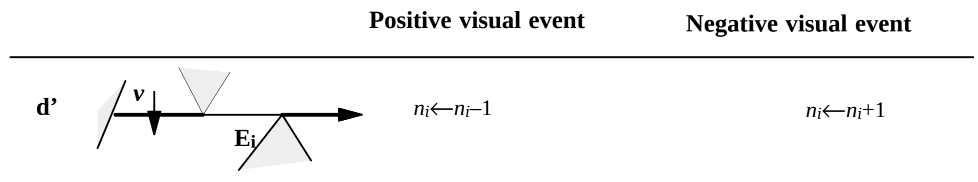

15

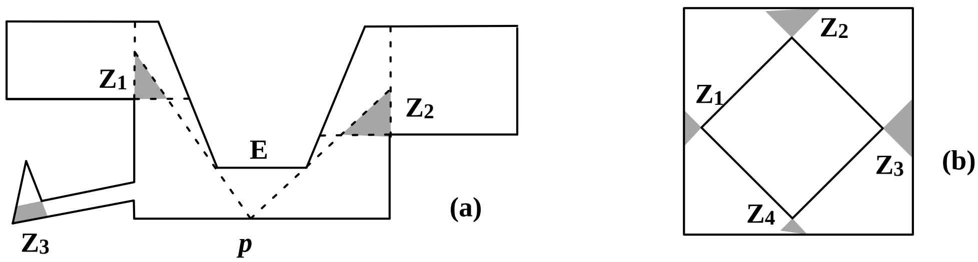

16

Fig. 15. Table 1 fits the case of exterior visibility of polygons with the addition of this entry

Fig. 16a. There are infinite pairs of maximum regions Z_1 and Z_2, depending on point *p*. Region Z_3 is necessary for covering *E*; **b** a solution for the ECH is obtained by locating two guards in the maximum regions Z_1 and Z_2 (or Z_2 and Z_4)

The worst-case numbers of guards have been found to be equal for edge and interior cover, and the edge-covering problem to be NP-hard for polygons with and without holes.

For the edge-covering problem, a restriction has been proposed that makes practical sense. The entire edge-covering problem requires each edge to be entirely visible from at least one guard. This problem is also NP-hard. For polygons without holes, the worst-case number of guards is $\lfloor n/3 \rfloor$, as for the unrestricted problem. For polygons with one hole, we have found that at most $\lfloor (n+2)/3 \rfloor$ guards are always sufficient and sometimes required. For more than one hole, we conjecture that is $\lfloor (n+h)/3 \rfloor$ is the tight bound.

We have described an algorithm for the entire edge covering problem. It computes a set of polygonal regions where the guards of a minimum set can be independently located. The algorithm is also suitable for the fortress or external guarding problem. The last step of the algorithm is an instance of the set-covering problem, exponential in the worst case. A greedy selection supplies nearly optimal solutions in polynomial time within a factor dependent only on the logarithm of the largest number of edges observed by a guard, and this factor is independent of n.

In practice, placing visual sensors is likely to have to satisfy additional constraints. A feature of the algorithm described is that it can easily be modified to take geometrical restrictions into account. For instance, a maximum distance and a minimum distance from each point observed could be required. This constrains the guards required for observing each edge into a region whose boundary lines can be inserted into Π for obtaining a modified partition Π^* of P.

We have observed that simple ways for deriving internal guards from wall guards do not seem to exist. This suggests some hope for solving for edge covering some of the problems unsolved for interior covering. Many lines of research could be explored for edge covering, such as:

– Attempting to construct exact algorithms, even if they become exponential in the worst case.
– Attempting to construct approximate algorithms with guaranteed performance.
– Looking for lower bounds for the edge guards of a given polygon. This would enable us to evaluate the quality of an approximate solution for a given polygon. It could be satisfactory even if it is supplied by approximate algorithms with no guaranteed performance.

– Looking for solutions for particular classes of polygons. For instance, in unrestricted edge covering, there are generally infinite sets of mutually dependent regions where minimum sets of guards can be located (Fig. 16a). However, there is a finite number of sets of independent maximum regions for some polygons, as for the case of entire visibility (Fig. 16b). Constructing algorithms for such polygons could be a much easier task.

References

1. Baase S (1988) Computer algorithms. Addison-Wesley, New York
2. Bjorling-Sachs I, Souvaine DL (1995) An efficient algorithm for guard placement in polygons with holes. Discrete Comput Geom 13:77–109
3. Chvátal V (1975) A combinatorial theorem in plane geometry. J Combinat Theor, Series B 18:39–41
4. Culberson JC, Reckhow RA (1994) Covering polygons is hard. J Algorithms 17:2–44
5. Edelsbrunner H (1987) Algorithms in combinatorial geometry. Springer, Berlin, Heidelberg
6. Fisk S (1978) Short proof of Chvatal's Watchman Theorem. J Combinat Theor, Series B 24:374
7. Gigus Z, Canny J, Seidel R (1991) Efficiently computing and representing aspect graphs of polyhedral objects. IEEE Trans Patt Anal Machine Intell 13:542–551
8. Gualtieri JA, Baugher S, Werman M (1989) The visual potential: one convex polygon. Comput Vision Graph Image Processing 46:96–130
9. Hoffmann F, Kaufmann M, Kriegel K (1991) The art gallery theorem for polygons with holes. Proceedings of the 32nd Symposium of Foundations of Computer Science, pp. 39–48
10. Laurentini A (1996) Guarding the walls of an art gallery. Internal Report, Politecnico di Torino, Italy
11. Lee D, Lin A (1986) Computational complexity of art gallery problems. IEEE Trans Inform Theor IT–32:276–282
12. Muroga S (1979) Logic design and switching theory. John Wiley, New York, Chapt. 4
13. Nemhauser G, Wolsey L (1988) Integer and combinatorial optimization. John Wiley, New York, pp. 466
14. O'Rourke J, Supowit K (1983) Some NP-hard polygon decomposition problems. IEEE Trans Inform Theor IT–29:181–190
15. O'Rourke J (1987) Art gallery theorems and algorithms. Oxford University Press, New York
16. Parker RG, Rardin RL (1988) Discrete optimization. Academic Press, San Diego, p. 37
17. Plantinga H, Dyer CR (1990) Visibility, occlusion and the aspect graph. Int J Comput Vision 5:137–160
18. Preparata F, Shamos M (1985) Computational geometry: an introduction. Springer, Berlin, Heidelberg, New York
19. Salkin HM, Mathur K (1989) Foundations of integer programming. Elsevier, New York, Chapt. 13
20. Shermer T (1992) Recent results in art galleries. IEEE Proc 80:1384–1399
21. Urrutia J (1996) Art gallery and illumination problems. In: Sack JR, Urrutia J (eds) Handbook of computational geometry, Internal Report, Ottawa University, Canada

ALDO LAURENTINI was born in Genova, Italy in 1940. He received the Laurea in Ingegneria Elettronica from the Politecnico di Milano, in 1963. In 1965 he joined the Politecnico di Torino, where he is Professor of Computer Science. His current interests include: application areas as physically correct simulation of the draping of clothes, and performing non-intrusive motion capture; theoretical areas in computer vision and computational geometry, such as exploiting the visual hull idea for object recognition or reconstruction from their silhouettes or shadows, and art gallery problems. Aldo Laurentini is member of IEEE and ACM, and author of more than 50 papers in leading journal and conferences.