

A New Lower Bound for Evaluating the Performances of Sensor Location Algorithms

Original

A New Lower Bound for Evaluating the Performances of Sensor Location Algorithms / Bottino, ANDREA GIUSEPPE; Laurentini, Aldo; Rosano, L.. - In: PATTERN RECOGNITION LETTERS. - ISSN 0167-8655. - STAMPA. - 30 (13):(2009), pp. 1175-1180. [10.1016/j.patrec.2009.05.020]

Availability:

This version is available at: 11583/2278664 since:

Publisher:

Elsevier

Published

DOI:10.1016/j.patrec.2009.05.020

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

A New Lower Bound for Evaluating the Performances of Sensor Location Algorithms

Andrea Bottino^{a,1}, Aldo Laurentini^a, Luisa Rosano^a

^a*DAUIN, Politecnico di Torino, Corso Duca degli Abruzzi 24, Torino, 10129, Italy*

Abstract

Locating sensors in 2D can be modelled as an Art Gallery problem. Tasks such as surveillance require observing or “covering” the interior of a polygon with a minimum number of sensors (IC, Interior Covering). Edge Covering (EC) is sufficient for tasks such as inspection or image based rendering. As IC, also EC is NP-hard, and no finite algorithm is known for its exact solution. A number of heuristics have been proposed for EC, but their performances with respect to optimality are unknown. Recently, a lower bound for the cardinality of the optimal EC solution, specific of a given polygon, has been proposed. It allows assessing the performances of approximate EC sensor location algorithms. In this paper, we propose a new lower bound. It is always greater than, or equal to the previous, and can be computed in reasonable time for environments with up to a few hundreds of edges. Tests over hundreds of polygons using a recent incremental EC algorithm show that the gap between the cardinality of the solution provided by the algorithm and the new lower bound is substantially reduced, and then the new lower bound outperforms the previous one.

Key words: Art Gallery, visual sensor positioning, visibility, edge covering, inspection, surveillance

PACS: 42.30.Sy

Email addresses: andrea.bottino@polito.it (Andrea Bottino)

URL: <http://www2.polito.it/ricerca/cgvv/> (Andrea Bottino)

¹Tel +39 011 5647175 / Fax +39 011 5647099

1. Introduction

Several computer vision and robotics tasks, as surveillance, inspection, image based rendering, constructing environment models, require multiple sensor locations, or the displacement of a sensor in multiple positions for fully exploring an environment or an object. Sensor placement, or planning, or location, is an important area of research, addressed in several surveys [17, 22, 23, 24]. Sensor location problems require considering a number of constraints, such as image resolution, field of view of the sensors, feature visibility, lighting, etc. Visibility is clearly the fundamental constraint. An omni directional or rotating sensor is usually modelled as a point. A feature of an object is visible from the sensor if any segment joining a point of the feature and the viewpoint does not intersect the environment or the object itself.

Although the general problem is three-dimensional, in several cases it can be restricted to 2D. This is for instance the case of buildings. The 2D visibility constraint is modelled by the classic Art Gallery problem, which asks to position a minimum set of “guards” able to see, or “cover” a polygonal environment. Tight upper bounds for the cardinality of the set of guards have been found in several cases. The famous Art Gallery Theorem by Chvátal states that at most $\lfloor n/3 \rfloor$ guards are required for covering any simple polygon with n edges, metaphorically the interior of an art gallery. The upper tight bound $\lfloor (n+h)/3 \rfloor$ holds for polygons with n edges and h polygonal holes. Many variations of the problem have been considered, as particular kind of polygons, restricted positions for the guards, additional constraints. For further details, the reader is referred to the monograph by O’Rourke [19] and to the surveys by Shermer [21] and Urrutia [25].

Unfortunately, the practical problem, that is locating a minimum set of guards in a given polygon, is NP-hard, and no finite exact algorithm is known for locating a minimum cover. Approximate algorithms with guaranteed per-

formance and polynomial in the worst case are unlikely to exist [10]. Anyway, several approximate polynomial algorithms have been proposed, reported by Shermer [21] or presented later, for instance by Bjorling-Sachs and Souvaine [3], Elnagar and Lulu [11], [12], Efrat and Har-Peled [9], Amit, Mitchell and Packer [1].

Observe that the sensors are required to cover the interior of a polygon for tasks such as *surveillance*. Other tasks, such as *inspection*, a main application of sensor planning according to the survey in Ref. [23], and *image based rendering*, only require observing the *boundary*. In this paper, we deal with the latter problem, referred to as the *Edge Covering (EC)* problem, while the classic problem is called the *Interior Covering (IC)* problem. The EC problem and its relations with IC have been analyzed in Ref. [16]. The Chvátal bound also holds for EC, but, although any interior cover is also an edge cover, in general an optimum set of IC guards is not an optimum set of EC guards and vice-versa. Examples show that the number of IC guards may be two times, for simple polygons, or $\Theta(n)$ times, for polygons with holes, the EC guards [16]. Then, even if efficient IC algorithms were found, using for EC the solution provided could result in a large sensor waste. EC is different from IC, but not easier. Actually, also EC is NP-hard [16], and no finite exact algorithm is known for locating a minimum set of EC guards in a given polygon.

Also for the EC problem, approximate sensor positioning algorithms have been presented. For instance, Kazazakis and Argyros [15] have proposed and implemented a polynomial heuristic that also takes into account the range constraint. The randomized approach (Danner and Kavraki [8], Gonzales-Banos and Latombe [13, 14]) attempts to approach the optimal solution by locating at random many sensors. Anyway, no experimental or theoretical results are supplied for evaluating the quality of the solution provided by these algorithms (cardinality of the set of guards compared with that of the optimal solution).

Recently, an EC incremental sensor location algorithm has been presented

[4, 5, 6]. This algorithm converges toward the optimal solution in an undefined number of steps, and makes use of a lower bound, *specific of the polygon considered*, for the minimum, or optimum, number of guards. The lower bound allows evaluating the quality of the solution obtained at each step, and halting the algorithm if the solution is satisfactory. Experimental results show that on the average the algorithm supplies solutions close to the lower bound, and then to the optimal cover. The idea of a lower bound for the IC problem has been presented in Ref. [1].

Clearly, since no known algorithm is able to compute the cardinality of a minimum set of EC guards, a tight lower bound is of great importance for evaluating the quality of sensor positioning algorithms. In this paper, we present and discuss a new, polygon specific, lower bound algorithm. The lower bound computed with this algorithm is equal or larger than that computed with the algorithm described in Refs. [4, 6]. The algorithm has been implemented and tested for many random polygons of different categories and different number of edge, and compared with the results supplied by the previous lower bound algorithm. The tests show that the new lower bound is significantly larger than that provided by the previous algorithm, and, most of all, that the gap between the solution provided by the incremental algorithm and the lower bound is substantially reduced. The algorithm is not polynomial, but its running time allows dealing with polygons with up to a few hundred of edges, and it will be shown that computing the new lower bound takes less time than the old one.

The paper is organized as follows. In Section 2, we describe the new lower bound algorithm. Section 3 provides the experimental results and comparisons. Concluding remarks are reported in Section 4.




Figure 1: Two weak visibility polygons. Each of these polygons must contain at least one guard.

2. The Lower Bound Algorithm for EC

2.1. The previous Lower Bound and its shortcomings

Let us first recall the lower bound algorithm described in Ref. [5]. It is based on the concept of *weak visibility polygon of an edge*. Two points of a polygon P are visible, or see each other, if the segment joining the points lies completely in P . According to Avis and Toussaint [2]:

Definition 1 *a polygon W is weakly visible from an edge e if for each point $w \in W$ there exists at least a point $z \in e$ such that w is visible from z .*

In other words, the weak visibility polygon $W(e_i)$ of an edge e_i is the polygon whose points see at least a point of e_i . Observe that points seeing *only one vertex* of e_i do not belong to $W(e_i)$. Examples of weak visibility polygons are shown in Fig. 1. Polynomial algorithms for computing weak visibility polygons of an edge are described in the literature [20]. In our case, however, weak visibility polygons are computed as a by-product of the sensor location algorithm described in Refs. [4, 6].

Weak visibility polygons allow us to determine a lower bound for the number of sensors needed. In fact, each weak visibility polygon must contain *at least one sensor*, otherwise no points of the edge are seen by any sensor. Therefore, a lower

bound $LB_W(P)$ for a polygon P is obtained by computing *the cardinality of the maximal subset of disjoint (not intersecting) weak visibility polygons $W(e_i)$ of P .*

A simple example is shown in Fig. 1. It is easy to verify by inspection that no more than two disjoint weak visibility polygons can be found, for instance $W(e_1)$ and $W(e_2)$, and thus $LB_W(P) = 2$, which is also the cardinality of the minimum set of sensors.

Computing LB_W requires solving the *maximum independent set problem* for a graph G where each node represents the weak visibility polygon of an edge of P , and each edge of G connects nodes corresponding to intersecting weak visibility polygons. The problem is equivalent to the *maximum clique problem* for the complement graph G' . Although this is an NP-complete problem, exact branch-and-bound algorithms for these problems have been presented and extensively tested [25, 18, 26], showing more than acceptable performances for graphs with hundreds of nodes.

The tests reported in Refs. [4, 6] also show that on the average the difference between the $LB_W(P)$ and the cardinality of the solution provided by the sensor location algorithm is small, and both are close to the optimum cardinality that lies in between.

However, the algorithm for computing LB_W fails to produce good results in some simple cases. Consider for example the case in Fig. 2, showing a comb polygon of a family used for showing that the Chvátal upper bound is tight for both IC and EC. Only two not intersecting weak visibility polygons can be found, for instance those shown in Fig. 2(a), and then $LB_W(P) = 2$. However, three EC guards are clearly required, one for each spike. The reason of the bad behaviour of the algorithm in this case can be appreciated from Fig. 2(b), where the weak visibility polygon $W(e_3)$ of one of the edges forming the central spike is shown. $W(e_3)$ intersects $W(e_1)$, and likewise $W(e_4)$ intersects $W(e_2)$.

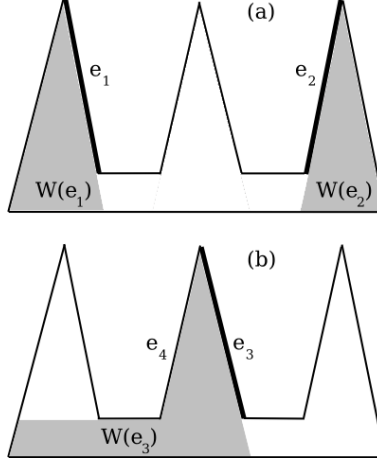


Figure 2: $LB_1(P)$ is two, but EC requires three guards.

Let us observe that similar arguments show that the lower bound LB_W is 2 for *all polygons* of the comb family: while the cardinality of the minimum set of guards increases with the number of spikes.

2.2. The new Lower Bound Algorithm

The previous example suggests considering visibility polygons of parts of the boundary *smaller* than an edge. Given a polygon P , let us recall the definition given by O'Rourke [19]:

Definition 2 *the point visibility polygon $VP(x)$ of a point x is the set of points $p \in P$ visible from x .*

In particular, we focus our attention on convex vertices of the polygon and thus consider $VP(v_i)$ of all convex vertices v_i of P . We only consider convex vertices, because they produce visibility polygons smaller than those of the edges converging at the vertices.

Consider the cardinality of the maximal subset of not intersecting VP s of convex vertices. It is clear that this cardinality is another lower bound, since

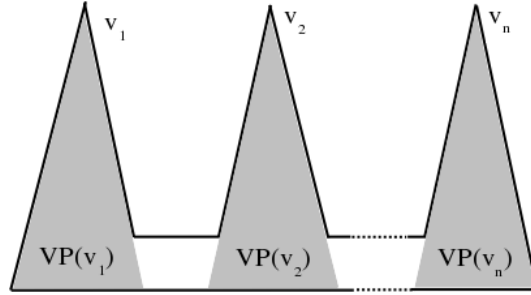


Figure 3: The non intersecting $VP(v_i)$ are as many as the guards.

each VP must contain at least one guard. If we use this new lower bound, the problem with the comb polygon family is solved, as shown in Fig. 3.

However, choosing as lower bound the cardinality of the larger set of VP s of convex vertices could be not satisfactory even in relatively simple cases. Consider for instance the polygon in Fig. 4. It can be easily verified that no more than four VP s of convex vertices exist, and precisely those of the vertices v_1, v_2, v_3, v_4 (Fig. 4(a)). However, five EC sensors are required, located for instance as shown in Fig. 4(b).

The examples discussed suggest to take into account both weak visibility polygons of edges and point visibility polygons of convex vertices.

Then we assume the following new definition of lower bound:

Definition 3 *The lower bound $LB_{W\&VP}(P)$ is the cardinality of the maximal subset (or subsets) of not intersecting weak visibility polygons $W(e_i)$ of edges e_i of P , and visibility polygons $VP(v_i)$ of convex vertices v_i of P .*

With this definition we solve the problems highlighted both for the comb polygons and for the polygon of Fig. 4. In the latter case the new definition supplies five and not four as lower bound. A maximum set of non intersecting visibility polygons is shown in Fig. 5. One of them is the weak visibility polygon of the edge e ; the other polygons can be interpreted either as visibility polygons

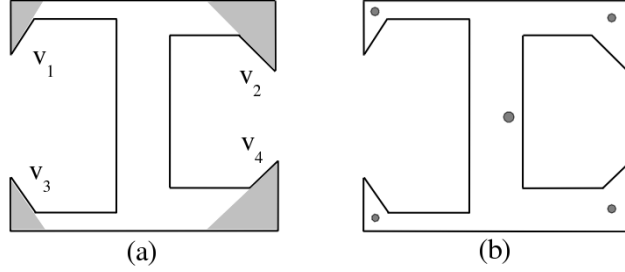


Figure 4: At most, four non intersecting VP of convex vertices can be found (a), but five guards are required (b).

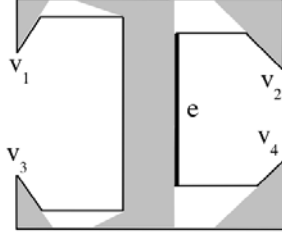


Figure 5: Five non intersecting visibility polygons are found.

of convex vertices, or as weak visibility polygons of edges converging in these vertices. Combining polygons as those shown in Fig. 2 and 4, we can easily produce examples where the new lower bound is better than those provided by weak visibility polygons and convex vertex visibility polygons separately.

In general, it is clear that $LB_{W\&VP}(P) \geq LB_W(P)$ for any P , and then $LB_{W\&VP}(P)$ is a better or equal lower bound. A block diagram of the new algorithm is shown in Fig. 6.

Polynomial algorithms for computing VP s of polygons with and without holes can be found in O'Rourke [19]. In addition, for polygons without holes it is possible to compute the point visibility polygon of a convex vertex v_i as the intersection of the weak visibility polygons of the edges converging into v_i . In our case, vertex visibility polygons are computed again as a by-product of the

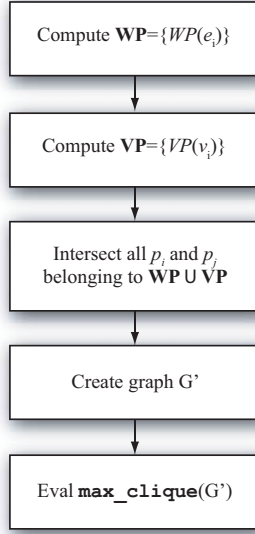


Figure 6: New algorithm.

sensor location algorithm described in Refs. [4, 6].

At a first glance, we could expect a heavier computational burden for the non polynomial part of the algorithm, that is the selection of the maximum independent set of vertices in the associated graph. However, as confirmed by the experimental section, this is not the case since an important reduction of the number of nodes of the graph can be performed. It is clear in fact that the nodes corresponding to weak visibility polygons of the edges converging at the convex vertices can be deleted from the graph, since they enclose the weak visibility polygon of the vertex. Adding up, considering c convex vertices reduces to $n - c$ the number of nodes of the graph.

2.3. Taking into account range and incidence

The new lower bound can be easily extended to take into account other geometrical constraints. EC algorithms usually also consider: a) minimal and maximal distances between the sensors and the observed boundary points; b)




Figure 7: $C(e_i)$ for range (a) and for incidence (b) constraints; $C(v_i)$ for range constraint.

minimal angle of incidence between an edge and the viewline [8, 13, 14, 15]. For each edge e_i each constraint defines a region $C(e_i)$ of P where the viewpoint can be located. These regions can be easily computed (we omit the obvious details). An example of these regions for range and incidence constraints can be seen in Fig. 7 (a) and (b). Further details can be found in Ref. [6]. The range constraint for convex vertices simply defines a region $C(v_i)$ as that shown in Fig. 7(c).

For computing the new lower bound $LB_{W \& VP}$, it is sufficient to consider reduced visibility polygons, obtained as intersections of $W(e_i)$ and $C(e_i)$ or of $VP(v_i)$ and $C(v_i)$.

Finally, observe that same care must be observed in selecting incidence and range parameters, otherwise the resulting regions could be empty, or lie outside the polygon.

3. Experimental results

In this section, we present experimental results showing that, on the average, the new lower bound significantly outperforms the previous.

In order to evaluate the performance of $LB_{W \& VP}$ compared to LB_W , we implemented it within the EC algorithm described in Ref. [5]. Thus, two versions of the EC algorithm are considered: one is the original version (described in Ref. [5]) computing the lower bound LB_W , while the second computes the $LB_{W \& VP}$ proposed in this article. In the following, results from the original

version of the EC algorithm are subscripted with W , while results from the new version are subscripted with $W\&VP$. Comparing the old and the new LB is not sufficient for a full evaluation. A better insight is provided by the reduction of the *absolute gap* between the lower bound and the solution provided by the EC algorithm, as well as by the reduction of the *relative gap*, that is the gap divided by the cardinality. Finally, computation times for the new and old lower bounds are compared.

Both versions of the EC algorithm were tested over several hundreds of polygons belonging to the following five categories:

- (A) generic random polygons, with edges oriented in generic directions;
- (B) generic random polygons with one to three holes;
- (C) orthogonal random polygons with no holes;
- (D) orthogonal random polygons with one to three holes;

Four different sets of polygons, with 30, 40, 50 and 60 edges, were constructed for each of the first four categories. Test results for each category are illustrated through Table 3 to Table 6.

Each line of the tables refers to a set of *no.* polygons with *nedges* edges used for tests. Data reported in these tables provide the following information averaged over the total number of polygons for each set:

- LB , the lower bound computed;
- C , the cardinality of the final EC solution. For polygons of the categories (A)-(D) the cardinality is given by the solution of the EC algorithm presented in Ref. [5], with four iterations without improvements and a time limit for the execution of 2400s. For polygons of category (E), the cardinality is given by the greedy solution of the EC algorithm;

- G , the gap between the lower bound and the cardinality of the EC solution. Precisely, $G_W = C - LB_W$ is the gap estimated for each polygon tested under the original version of the EC algorithm and $G_{W\&VP} = C - LB_{W\&VP}$ is the gap estimated under the new version of the algorithm. The smaller is the gap, the better a solution is. Clearly, in the optimal case, the gap is null;
- G/C , the *relative gap*; that is, respectively, G_W/C and $G_{W\&VP}/C$.
- $LBtime$, the total time, in seconds, spent to compute the lower bound computation (see below for further details);
- $G\ reduction$, the gap reduction percentage, defined as $100 \cdot (1 - G_{W\&VP}/G_W)$;
- $LBtime\ reduction$, the percentage of time saved computing the lower bound as $LB_{W\&VP}$ instead of LB_W (negative values stand for extra time spent).

The experiments show a substantial reduction of the gap between LB and the final EC solution and of the relative gap. These improvements are summarized per polygon category in Table 1, where we can see that the gap reduction ranges between 28% and 46%.

Table 1: Total gap reduction per polygon category.

Polygon category	G reduction
Random	42.59%
Random with holes	30.38%
Orthogonal	46.43%
Orthogonal with holes	28.93%
Total	35.80%

As a whole, considering all the experiments, the mean gap reduction is 35.80%. These results assert that the new lower bound presented in this article provides a substantially tighter approximation of the optimum.

Regarding the processing time, the total time reported includes:

- the *data structure time*, that is the time spent to construct the required data structure (weak visibility polygons in the case of LB_W , weak and point visibility polygons in the case of $LB_{W\&VP}$)
- the *max clique time*, which is the time taken to construct the dual graph from the set of visibility polygons and to solve the max clique problem.

Processing times required for computing the data structure and solving the max clique problem were individually recorded for each polygon tested and then averaged per each polygon category. Per cent time reductions for: 1) constructing the data structure, 2) solving the max clique problem and 3) computing the lower bound as a whole are summarized in Table 2.

Table 2: Lower bound reduction times.			
Average reduction times			
Polygon category	Data struct	Max Clique	Total
Random	-1.24%	31.40%	10.70%
Random with holes	-3.19%	30.69%	10.72%
Orthogonal	-1.21%	21.39%	5.96%
Orthogonal with holes	-15.46%	24.27%	-0.15%
Total	-3.60%	28.28%	8.33%

Positive values stand for time savings while negative values stand for extra time consumed. Table 2 shows that, as expected, the time spent in creating the data structure increases, but the time spent in evaluating the max clique decreases. This is particularly evident for polygons with a very high number of edges. On the average, a time saving results, particularly for large sizes of the associated graph.

Concluding, the new LB definitely outperforms the previous one.

4. Conclusions

We have studied, implemented and experimented a new lower bound for the minimum number of guards required for solving the EC problem, a variation of the Art Gallery problem important for tasks such as inspection.

The new lower bound has been experimentally compared with a previous lower bound, and with the cardinality of the coverage provided by an efficient incremental EC algorithm.

The results collected from a wide range of polygons, with and without holes, show that the new lower bound outperforms the previous one, since the gap between the lower bound and the solution provided by the EC algorithm is reduced on average of about one third. Furthermore, the new lower bound requires less time for its evaluation, since the bottleneck of the algorithm, computing a maximum independent set of vertices of an associated graph, is performed on a graph with less vertices.

Concluding, the new lower bound has been shown to be closer to the optimum, and than more effective for evaluating the performances of approximate EC sensor location algorithms.

Table 3: Random polygons - (A)

nedges	no.	C	LB		G		G/C		LBtime		G reduction	LBtime reduction
			W&VP	W	W&VP	W	W&VP	W	W&VP	W		
30	20	4.30	3.95	3.80	0.35	0.50	0.078	0.108	0.319	0.292	30.00%	-9.09%
40	20	5.40	5.20	4.90	0.20	0.50	0.035	0.087	0.643	0.598	60.00%	-7.58%
50	20	6.70	6.25	5.90	0.45	0.80	0.066	0.114	1.561	1.731	41.88%	9.83%
60	20	8.30	7.75	7.40	0.55	0.90	0.063	0.102	2.687	3.212	38.89%	16.37%

Table 4: Random polygons with 1-3 holes - (B)

nedges	no.	C	LB		G		G/C		LBtime		G	LBtime
			W&VP	W	W&VP	W	W&VP	W	W&VP	W	reduction	reduction
30	20	5.20	4.90	4.60	0.30	0.60	0.046	0.099	0.460	0.448	50.00%	-2.83%
40	20	6.45	5.70	5.40	0.75	1.05	0.115	0.159	1.674	1.728	28.57%	3.12%
50	20	7.45	6.75	6.40	0.70	1.05	0.106	0.137	2.106	2.345	33.33%	10.19%
60	20	8.40	7.40	7.15	1.00	1.25	0.115	0.146	2.982	3.569	20.00%	16.44%

Table 5: Orthogonal polygons - (C)

nedges	no.	C	LB		G		G/C		LBtime		G	LBtime
			W&VP	W	W&VP	W	W&VP	W	W&VP	W	reduction	reduction
30	20	4.50	4.25	3.95	0.25	0.55	0.053	0.128	0.195	0.173	54.55%	-13.06%
40	20	6.15	5.70	5.35	0.45	0.80	0.073	0.127	0.333	0.329	43.75%	-1.16%
50	20	7.30	6.60	6.15	0.70	1.15	0.095	0.159	2.236	1.055	39.13%	-111.85%
60	20	8.85	8.00	7.15	0.85	1.70	0.091	0.187	1.814	3.311	50.00%	45.22%

Table 6: Orthogonal polygons with 1-3 holes - (D)

nedges	no.	C	LB		G		G/C		LBtime		G	LBtime
			W&VP	W	W&VP	W	W&VP	W	W&VP	W	reduction	reduction
30	20	6.00	5.45	5.25	0.55	0.75	0.464	0.647	0.202	0.170	26.67%	-18.85%
40	20	8.30	7.70	7.55	0.60	0.75	0.069	0.086	0.248	0.182	20.00%	-36.38%
50	20	9.10	7.85	7.40	1.25	1.70	0.142	0.198	0.707	0.733	26.47%	3.53%
60	20	11.35	9.45	8.50	1.90	2.85	0.165	0.243	1.402	1.471	33.33%	4.67%

References

- [1] Amit, J., et al., 2007. Locating Guards for Visibility Coverage of Polygons. Workshop on Algorithm Engineering and Experiments, ALANEX.
- [2] Avis, D., Toussaint, G.T., 1981. An optimal algorithm for determining the visibility polygon from an edge. IEEE Trans. Comput., C-30(12), 910–914.
- [3] Bjorling-Sachs, I., Souvaine, D., 1995. An efficient algorithm for guard placement in polygons with holes. Discrete and Computational Geometry, 77–109.
- [4] Bottino, A., Laurentini A., 2005. A new art gallery algorithm for sensor location. Proc. ICINCO, 242–249.

- [5] Bottino, A., Laurentini A., 2006. Experimental results show near-optimality of a sensor location algorithm. Proc. of IEEE Int. Conf. on Robotics and Biomimetics.
- [6] Bottino, A., Laurentini, A., 2008. A Nearly Optimal Sensor Placement Algorithm for Boundary Coverage. Pattern Recognition, 41(11), 3343–3355.
- [7] Carraghan, R., Parados, P., 1990. An exact algorithm for the maximum clique problem. Operations Research Letters, 9, 375–382.
- [8] Danner, T., Kavraki, L. E., 2002. Randomized planning for short inspections paths. IEEE Int. Conf. On Robotics and Automation, 971-976.
- [9] Efrat, A., Har-Peled, S., 2006. Guarding galleries and terrains. Info. Proc. Letters, 100, 138–245.
- [10] Eidenbenz, S., et. al., 2001. Inapproximability results for guarding polygons and terrains. Algorithmica, 31, 79–113.
- [11] Elnagar, A., Lulu, L., 2005. An Art Gallery based approach to autonomous robot motion planning in global environments. Proc. IEEE Int. Conf. on Intelligent Robotics and Systems, 1367–1372.
- [12] A. Elnagar, A., Lulu, L., 2005. A comparative study between visibility-based roadmap path planning algorithms. Proc. IEEE Int. Conf. on Intelligent Robotics and Systems, 3700–3705.
- [13] Gonzales-Banos, H., Latombe, J-C., 2001. A randomised art gallery algorithm for sensor placement. ACM SCG’01.
- [14] Gonzales-Banos, H., Latombe, J-C., 1998. Planning robot motion for range-image acquisition and automatic 3D model construction. Proc. AAAI Fall Symposium Series.

- [15] Kazazakis, G. D., Argyros, A., 2002. Fast positioning of limited visibility guards for inspection of 2D workspaces. *Proc. 2002 IEEE/RSJ Int. Conf. On Intelligent Robots and Systems*, 2843–2848.
- [16] Laurentini, A., 1999. Guarding the walls of an art gallery. *The Visual Computer*, 15, 265–278.
- [17] Newman, T. S., Jain, A. K., 1995. A survey of automated visual inspection. *Comput. Vis. Image Understand.*, 61(2), 231–262.
- [18] Oestergard, P., 2002. A fast algorithm for the maximum clique problem. *Discrete Applied Mathematics*, 120, 197–207.
- [19] O’Rourke, J., 1987. *Art gallery theorems and algorithms*. Oxford University Press, New York.
- [20] Sack, J. R., Suri, S., 1990. An optimal algorithm for detecting weak visibility of a polygon. *IEEE Trans. on Computers*, 39, 1213–1219.
- [21] Shermer, T., 1992. Recent results in art galleries. *IEEE Proc.*, 80, 1384–1399.
- [22] Scott, W. R., Roth, G., 2003. View planning for Automated Three-Dimensional Object Reconstruction and Inspection. *ACM Comput. Surveys*, 35(1), 64–96.
- [23] Tarabanis, K. A., et al., 1995. A survey of sensor planning in computer vision. *IEEE Trans. Robot. and Automat.*, 11(1), 86–104.
- [24] Tarbox, G., Gottschlich, S., 1995. Panning for complete sensor coverage in inspection. *Comp.Vis. and Image Und.*, 61, 84–111.
- [25] Urrutia, J., 2000. Art gallery and illumination problems. *Handbook on Computational Geometry*, Elsevier Science Publishers, J.R. Sack and J. Urrutia eds., 973–1026.

- [26] Woods, D.R., An algorithm for finding a maximum clique in a graph, 1997.
Operations Research Letters, 21, 211–217.