

Strategies for Optimal Placement of Surveillance Cameras in Art Galleries

Marcelo C. Couto, Cid C. de Souza,* Pedro J. de Rezende†
Instituto de Computação
Universidade Estadual de Campinas
Campinas, Brazil

Abstract

The Art Gallery problem (AGP) consists of minimizing the number of cameras required to guard an art gallery whose boundary is an n -vertex polygon P . In this paper, we report our ongoing work in exploring an exact algorithm for a few variants of AGP, which iteratively computes optimal solutions to Set Cover problems (SCPs) corresponding to discretizations of P . Besides having proven in [Couto et al. 2007] that this procedure always converges to an exact solution of the original continuous problem, we have evidence that, in practice, convergence is achieved after only a few iterations, even for random polygons of hundreds of vertices. Nonetheless, we observe that the number of iterations required is highly dependent on the way P is initially discretized. As each iteration involves the solution of an SCP, the strategy for discretizing P is of paramount importance. We present here some of the discretization strategies we have been working with and new ones that will be studied in the near future. In comparison to the current literature, our results show a significant improvement in the size of the instances that can be solved to optimality while maintaining low execution times: no more than 65 seconds for random polygons of up to one thousand vertices.

Keywords: Camera Placement, Art Gallery, Surveillance, Visibility

1 INTRODUCTION

In 1973, Victor Klee posed the *Art Gallery Problem*, (AGP), which consists in determining the minimum number of cameras sufficient to guard the interior of an n -wall art gallery [Honsberger 1976]. Chvátal showed, that $\lfloor n/3 \rfloor$ cameras are occasionally necessary and always sufficient to guard a simple polygon with n vertices [Chvátal 1975].

We focus here on the specific variation of AGP in which the placement of cameras is restricted to the vertices of the polygon and they have a 360° field of vision. Since the corresponding minimization problem has been proven NP-hard [Lee and Lin 1986] even for orthogonal polygons [Schuchardt and Hecker 1995], placement of sub-optimal numbers of cameras has been studied [Avis and Toussaint 1981; Sack and Toussaint 1988], as well as a few approximation algorithms [Amit et al. 2007]. Moreover, modeling AGP as a discrete combinatorial problem followed by the solution of the corresponding optimization problem has been attempted by Erdem and Sclaroff [2006] who discretize the interior of the polygon using a fixed grid, yielding an approximation algorithm.

The approach that we have been studying consists of iteratively modeling the problem as a classical *Set Cover problem* (SCP). Among the results obtained so far, we have shown that the number of iterations executed by our method is polynomially bounded, and in practice, very small. Furthermore, the experimental analysis show that the number of iterations and the total time depends much on how the polygon is discretized. This becomes clearer when one realizes that at each iteration an instance of SCP, an NP-hard problem, has to be solved to optimality, by an Integer Programming (IP)

solver, in our case.

So far, we have completed a thorough experimental investigation concerning the trade-off between the number and nature of the discretizing points and the number of iterations, analyzing the practical viability of each approach. Our test data, available at [Couto et al. 2008b], includes multiple instances for each size of the vertex set, for various classes of polygons with up to a *thousand* vertices. The experimental results obtained thus far surpass, by more than five-fold, those reported elsewhere in [Tomás et al. 2006] and in [Couto et al. 2007].

In the next sections, we present the method, we summarize the experimental results, we describe our testing environment and we indicate venues of further investigation and a conclusion.

2 ALGORITHM

In an instance of AGP we are given a simple polygon P that bounds an art gallery and we are asked to determine the minimum number and an optimal placement of vertex cameras in order to keep the whole gallery under surveillance. Vertex cameras are assumed to have a range of vision of 360° .

The approach used by the algorithm described below transforms the continuous AGP into a discrete problem which, in turn, can be easily modeled as an instance of SCP.

Due to limited space, we will assume the reader's familiarity with the standard terminology from the literature, such as [O'Rourke 1987]. However, the following notation is best made explicit: a *vertex surveillance set* S is any subset of vertices of P such that $\bigcup_{s \in S} V(s) = P$, where $V(s)$, the *visibility region of s* , is the set of all points visible from vertex s . In other words, a vertex surveillance set for P gives the positions of stationary cameras which can oversee an entire art gallery of boundary P . Thus, AGP amounts to finding the smallest subset S of vertices that is a vertex surveillance set for P .

We now describe how the solutions to successively refined discrete instances of SCP are guaranteed to converge to an optimal solution to the original continuous problem. To this end, consider an arbitrary discretization of P into a finite set of points $D(P)$. An IP formulation of the corresponding SCP instance is:

$$\begin{aligned} z &= \min \sum_{j \in V} x_j \\ \text{s.t.} \quad &\sum_{j \in V} a_{ij} x_j \geq 1, \text{ for all } p_i \in D(P) \\ &x_j \in \{0, 1\}, \text{ for all } j \in V \end{aligned} \quad (1)$$

where the binary variable x_j is set to 1 if and only if vertex j from P is chosen to be in the surveillance set. Moreover, given a point p_i in $D(P)$ and a vertex j of P , a_{ij} is a binary value which is 1 if and only if $p_i \in V(j)$.

Given a feasible solution x for the IP above, let $Z(x) = \{j \in V \mid x_j = 1\}$. Constraint (1) states that each point $p_i \in D(P)$ is visible from at least one selected camera position in the solution and the objective function minimizes the cardinality z of $Z(x)$. Clearly, as the set $D(P)$ is finite, it may happen that $Z(x)$ does not form a vertex surveillance set for P . In this case, we must add a new point inside each uncovered region and include these points in $D(P)$. A new SCP instance is then created and the IP is solved again.

We are now able to describe the algorithm proposed in [Couto et al. 2007]: in the *preprocessing phase*, two procedures are exe-

*Partially supported by CNPq – Conselho Nacional de Desenvolvimento Científico e Tecnológico – Grants # 307773/2004-3, 472504/2007-0 and FAPESP – Fundação de Amparo à Pesquisa do Estado de São Paulo – Grant # 107/97

†Partially Supported by CNPq – Grants # 201205/2005-0, 472504/2007-0, Fapesp – Grant # 07/52015-0 and a Grant from FAEPEX/UNICAMP.

cuted. The first one computes the visibility polygons for the points in V . The second one computes the initial discretization $D(P)$ and builds the corresponding IP model. In the *solution phase*, the algorithm iterates as described above, solving SCP instances for the current discretization, until no regions remain uncovered.

We had shown in [Couto et al. 2007] that an upper bound on the number of iterations is $O(n^3)$ which derives from the fact that the edges of the visibility regions induce a subdivision of P which is comprised of $\Theta(n^3)$ faces or *Atomic Visibility Polygons (AVPs)* – see [Bose et al. 2002]. One point inside an AVP is enough to guarantee that this entire AVP will be covered by the solution to the discretized problem. Whence follows the upper bound on the number of iterations.

Moreover, the actual number of iterations that is required depends on how many uncovered regions can be successively generated. As the cost of each iteration is related to the number of constraints in (1), an interesting trade-off naturally sprouts and leads one to attempt multiple choices of discretization schemes. On the other hand, any method of cleverly choosing the initial points of the discretization will have a corresponding cost in preprocessing time, opening another intriguing time exchange consideration. These questions are precisely what we address next.

In Section 3 we briefly describe several discretization schemes leading to the various performance analysis summarized in Section 4.

3 DISCRETIZATION STRATEGIES

The key point in the IP approach is to set up instances of SCP that can rapidly be solved while minimizing the number of iterations required to attain an optimal solution to the original art gallery problem, within the least amount of time. However, one must take into account that sophisticated geometric properties used to build more efficient discretizations will generate a corresponding cost in preprocessing, possibly outweighing the benefits.

Regular Grid. The first discretization strategy considered here is based on the generation of a dense grid inside the polygon in the assumption that few iterations might result. This approach, however, leads to very large instances of SCPs which increase the time needed to run the IP solver. A summary of the outcome of the use of regular grids for von Koch polygons can be seen in Table 1.

Induced Grid. Given the perception that reflex vertices are partly responsible for the hardness of the problem, a natural discretization strategy is the grid induced by the edge extensions that intersect in the polygon. Here, we generate fewer constraints than in the previous strategy while capturing more of the intrinsic visibility information of the polygon. One can expect this to decrease the time to solve the instances.

Just Vertices. As an extreme case, consider the rather sparse grid consisting only of the vertices of P . Surprisingly, this strategy leads to an overall faster method than the two aforementioned ones, due to the fact that at each new iteration all added constraints correspond to areas harder to supervise and all SCP instances are small.

Reduced Atomic Visibility Polygons. It is easy to see that, by definition, if a camera surveillance set S covers the centroid of an AVP, then it must cover the entire AVP. Therefore, if S covers the centroids of every AVP of P , then S must be a surveillance set for P . As an initial discretization comprised of the centroids of all AVPs would lead to an impractically large instance of $O(n^3)$ constraints, we have devised a way to reduce it to an equivalent subset. We refer the reader to [Couto et al. 2008a] for the details of this reduction.

4 EXPERIMENTAL RESULTS

In this section, we summarize our experimental evaluation of the discretization strategies discussed above. A description of the testing environment is presented in section 5.

4.1 INSTANCES

We conducted the tests on a large number of instances, downloadable from [Couto et al. 2008b], of polygons from three classes. The first two of these are composed of n -vertex orthogonal (random and von Koch) polygons and the last one is comprised of random simple (non-orthogonal) polygons.




Figure: Sample polygons: Simple Random and Orthogonal von Koch (with 100 vertices).

(1) **Orthogonal Random:** These are n -vertex random orthogonal polygons generated as described in [Tomás and Bajuelos 2004].

(2) **Complete von Koch (CvK):** These polygons are based on a modified version of the fractal von Koch curve, which is generated, starting from a square, by recursively replacing each edge as shown in Figure 1, where $\overline{ar} = \overline{st} = \overline{ub}$ and $\overline{sr} = \overline{tu} = \frac{3}{4}\overline{ar}$.




Figure 1: Levels of modified von Koch polygons.

(3) **Simple Random:** This class consists of randomized simple (non orthogonal) polygons generated using [Auer and Held 1996]. The random instances were generated for the number of vertices n in the ranges: [20, 200] with step 20, (200, 500] with step 50 and (500, 1000] with step 100. The CvK class contains, by construction, only 3 instances with $n \in \{20, 100, 500\}$.

To have statistical significance, we chose the number of instances generated to be between 10 and 30 for each value of n . Thus, in total, our data set is composed of 643 instances, having between 20 and 1000 vertices.

4.2 RESULTS

We now discuss the actual experimental evaluation of the strategies described in Section 3. All values reported here are average results for all instances of each size, or multiple runs of the same instance of CvK polygons.

Table 1: Results for Complete von Koch polygons.

# vertices	Final $D(P)$			Total Time		
	20	100	500	20	100	500
Reg. Grid	45	500	6905	0.05s	1.57s	92.37s
Ind. Grid	24	205	1665	0.03s	1.41s	70.94s
Red. AVPs	28	324	5437	0.07s	3.14s	143.93s
Just Vert.	20	107	564	0.04s	0.97s	29.35s

The usage of discretization strategies based on dense grids becomes discouraging when we analyze the results in Table 1 which displays the execution time and the size of the discretization of the strategies proposed in Section 3 for the CvK polygons. One can see that for these instances, the *Induced Grid* strategy has a better performance than the *Regular Grid* strategy. The size of the discretization produced by *Regular Grid* grows quadratically in the number of vertices and inflates the number of constraints in the IP formulation increasing considerably the time necessary to optimally solve the SCP instances. The *Reduced AVP* strategy has a poor behavior for CvK polygons since the number of shadow AVPs increases fast in this case. The *Just vertices* strategy is the one that spends less time.

Figure 2 shows the amount of discretized points necessary for each strategy to achieve the optimal solution of AGP for random

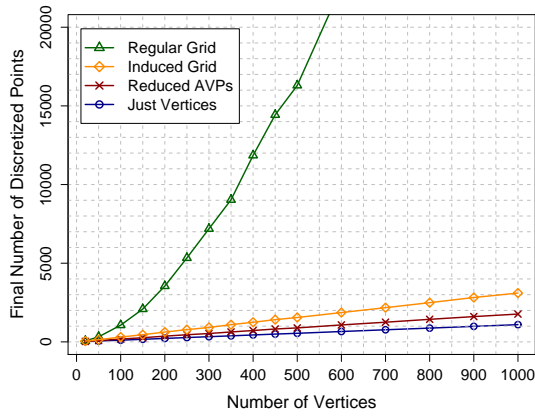


Figure 2: Final discretization size: Random ortho-polygons.

ortho-polygons. One can see that the *Regular Grid* strategy rapidly becomes impractical due to the huge size of the discretization and, therefore, will no longer be analyzed. On the other hand, the *Reduced AVP* strategy is very well-suited for random ortho-polygons.

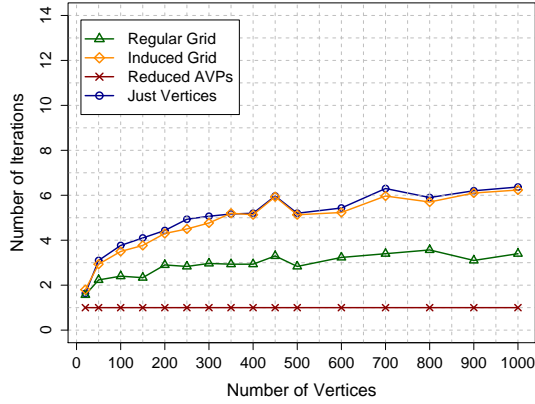


Figure 3: Number of iterations: Random ortho-polygons.

The curves corresponding to the *Just Vertices* strategy suggest that the set of vertices of the polygon is a good guess for the initial discretization since few new points are added to it to achieve the optimal solution of an AGP instance for this class of instances.

Figure 3 shows the number of iterations each strategy needs to achieve the optimal solution for random ortho-polygons. The chart displays the expected behavior with the number of iterations increasing as the size of the discretizations decrease. Now, relative to the size of the input polygon, the number of iterations remains negligible when compared to the theoretical bound of $\Theta(n^3)$.

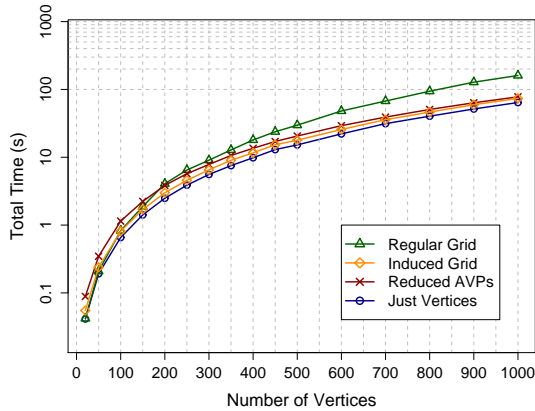


Figure 4: Total time: Random ortho-polygons.

Figure 4 shows the total time, including the preprocessing and processing phases, to solve instances from the random ortho-polygons. The curves are plotted in log \times linear format and both charts are in the same scale. One can see that for Random polygons, all the strategies behave similarly except, the *Regular Grid*.

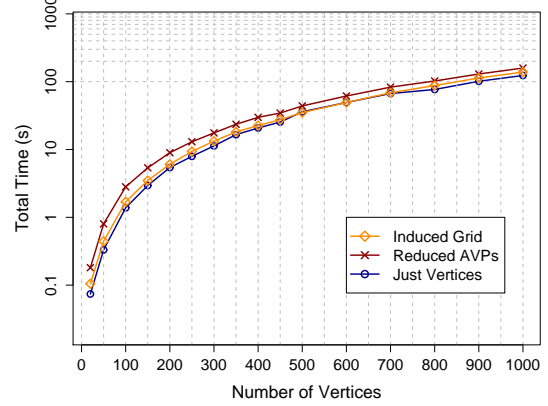


Figure 5: Total time: Random simple polygons.

It is interesting to notice that the corresponding graph for simple random polygons in Figure 5 is similar in shape, except that the times are doubled, due to this being a more difficult problem.

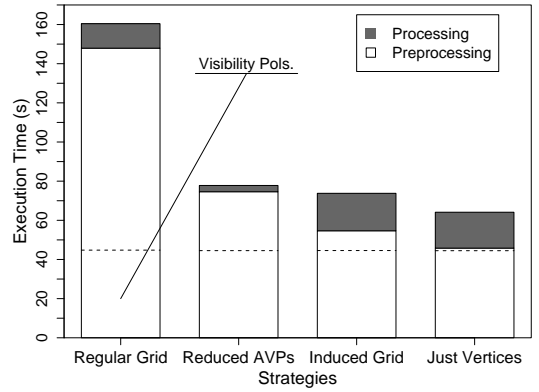


Figure 6: Execution time for polygons of 1000 vertices: Random ortho-polygons.

In Figure 6 one can see that even though we are solving NP-hard problems (SCPs) in the solution phase, most of the computing time is spent in the preprocessing phase, which is polynomial both while constructing the visibility polygons and while assembling the IP model after setting up the initial discretization. Note that the *Just Vertices* strategy requires almost no computation after the visibility polygons are built.

5 TESTING ENVIRONMENT

To evaluate our algorithm and the discretization strategies, we implemented, in C++, a testing environment designed in three layers, each one with a specific function. The first layer is a combination of the geometric library CGAL 3.2.1, the commercial IP solver Xpress v17.01.02 and literature standard algorithms to construct visibility polygons and to generate random polygons. This layer also contains a strategy depot. The system's architecture is depicted in Figure 7.

The second layer consists of the core of the application, i.e., the implementation of the method and uses several different interfaces for communication with the other layers and thus manages the algorithms that optimally solve the problem.

Finally, the third layer comprises the graphical interface of the application (see Figure 8) containing several toolbars that allows




Figure 7: Application Architecture.

the user to: change polygon type and the strategy used; run the method step by step with visual identification of the uncovered and covered regions; display the visibility region of any point inside the polygon; and zoom in to investigate details. It also displays information related to the discretization strategy and from the solution in view.

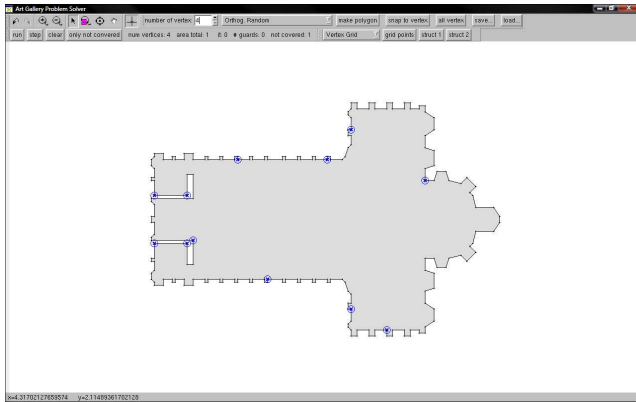


Figure 8: Application Interface showing an optimal solution for a polygon of 232 vertices corresponding to a simplified floor plan of the Basilica of St. Saturnin in Toulouse, France.

As for hardware, we used a desktop PC featuring a Pentium IV at 3.4 GHz and 1 GB of RAM running GNU/Linux 2.6.17.

6 CONCLUSIONS AND REMARKS

We conducted an experimental investigation of an exact algorithm for the NP-hard Art Gallery problem which relies on the discretization of the interior of the input polygon P and on the modeling of this simplified discrete problem as a Set Cover problem (SCP). The resulting SCP instance is solved to optimality by an IP solver and, if uncovered regions remain, additional constraints are included and the process is repeated. Clearly, the performance of the algorithm depends on the number of such iterations.

This work focused on different strategies to implement the discretization of P . Thorough experimentation was carried out to assess the trade-off between the number of iterations and time spent by the many variants of the algorithm that arise from the alternative discretization methods.

Our conclusion is that this exact algorithm is a viable choice to tackle instances of AGP, since the largest ones we solved were five times larger than those reported earlier in the literature.

Additional strategies for the initial discretization that are still under consideration include starting off with: (i) a single vertex, or a single internal point; (ii) only the reflex vertices, or only the convex vertices; (iii) only the centroids of AVPs that intersect edges of P .

Furthermore, our algorithm can solve (with minor modifications) a few related problems which we are currently working on:

- alternate sets of discrete potential camera spots (instead of just vertices), such as: (i) just reflex vertices, or just convex vertices; (ii) midpoints of all edges; (iii) a dense internal

grid, leading to an approximate solution to the continuous placement problem.

- redundant coverage (e.g., double coverage of P).

On the other hand, a promising venue of further investigation lies in trying to identify inexpensive geometric properties that might lead to a set of constraints that capture the essence of the hardness of the problem, such as a significant reduction on the number of AVPs.

ABOUT THE AUTHORS

Marcelo C. Couto is a graduate student in Computer Science at UNICAMP. His e-mail is: couto.marcelo@gmail.com

Cid. C. de Souza is a professor at UNICAMP. His e-mail is: cid@ic.unicamp.br

Pedro J. de Rezende is a professor at UNICAMP. His e-mail is: rezende@ic.unicamp.br

Correspondence to the authors can be addressed to:
 Instituto de Computação, C.P. 6176
 Universidade Estadual de Campinas, UNICAMP
 Campinas, SP, Brazil, 13083-970.

References

- AMIT, Y., MITCHELL, J. S. B., AND PACKER, E. 2007. Locating guards for visibility coverage of polygons. In *Proc. Workshop on Algorithm Eng. and Experiments*, 1–15.
- AUER, T., AND HELD, M., 1996. Heuristics for the generation of random polygons.
- AVIS, D., AND TOUSSAINT, G. T. 1981. An efficient algorithm for decomposing a polygon into star-shaped polygons. *Pattern Recogn. 13*, 395–398.
- BOSE, P., LUBIW, A., AND MUNRO, J. I. 2002. Efficient visibility queries in simple polygons. *Comput. Geom. 23*, 3, 313–335.
- CHVÁTAL, V. 1975. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory Series B 18*, 39–41.
- COUTO, M. C., DE SOUZA, C. C., AND DE REZENDE, P. J. 2007. An exact and efficient algorithm for the orthogonal art gallery problem. In *Proc. of the XX Brazilian Symp. on Comp. Graphics and Image Processing*, IEEE Computer Society, 87–94.
- COUTO, M. C., DE SOUZA, C. C., AND DE REZENDE, P. J. 2008. Experimental evaluation of an exact algorithm for the orthogonal art gallery problem. In *Proc. of the 7th International Workshop on Experimental Algorithms*, WEA 2008, Springer-Verlag, vol. 5038, 101–113.
- COUTO, M. C., DE SOUZA, C. C., AND DE REZENDE, P. J., 2008. OAGPLIB - orthogonal art gallery problem library. <http://www.ic.unicamp.br/~cid/Problem-instances/Art-Gallery/>.
- ERDEM, U. M., AND SCLAROFF, S. 2006. Automated camera layout to satisfy task-specific and floor plan-specific coverage requirements. *Comput. Vis. Image Underst. 103*, 3, 156–169.
- HONSBERGER, R. 1976. *Mathematical Gems II*. No. 2 in The Dolciani Mathematical Expositions. Math. Assoc. of America.
- LEE, D. T., AND LIN, A. K. 1986. Computational complexity of art gallery problems. *IEEE Trans. Inf. Theor. 32*, 2, 276–282.
- O’ROURKE, J. 1987. *Art Gallery Theorems and Algorithms*. Oxford University Press.
- SACK, J.-R., AND TOUSSAINT, G. T. 1988. Guard placement in rectilinear polygons. In *Computational Morphology*, G. T. Toussaint, Ed. North-Holland, 153–175.
- SCHUCHARDT, D., AND HECKER, H.-D. 1995. Two NP-hard art-gallery problems for ortho-polygons. *Mathematical Logic Quarterly 41*, 261–267.
- TOMÁS, A. P., AND BAJUELOS, A. L. 2004. Generating random orthogonal polygons. In *Current Topics in Artificial Intelligence*, Springer, vol. 3040 of LNCS, 364–373.
- TOMÁS, A. P., BAJUELOS, A. L., AND MARQUES, F. 2006. On visibility problems in the plane - solving minimum vertex guard problems by successive approximations. In *Proc. of the 9th Int. Symp. on Artificial Intelligence and Mathematics*.