Guarding Galleries with No Nooks (Extended Abstract)

David Kirkpatrick
kirk@cs.ubc.ca
Department of Computer Science
University of British Columbia
Vancouver BC, Canada V6T 1Z4

July 4, 2000

Abstract

We consider the problem of guarding galleries that have no small nooks (regions that are visible from only a small fraction of the entire gallery). Intuitively, such galleries (of which convex galleries are a special but uninteresting case) should need fewer guards.

We show that in any simply connected gallery in which every corner sees at least a fraction ϵ of the other corners there exist a set of at most $64\frac{1}{\epsilon}\log\log\frac{1}{\epsilon}$ corners (guard stations) from which every other corner is visible. Similar results also hold for guarding the entire boundary with guards stationed at points on the boundary.

These results are similar but slightly stronger than analogous results known for guarding the interior of a gallery. We provide a direct and simple guard assignment scheme that realizes the size bounds.

1 Introduction

We consider a variant of the well-studied problem of guarding art galleries. As usual, a gallery is modeled as a (bounded) simply connected polygonal region P in the plane. (The more general situation of galleries with interior walls (holes) is discussed in the concluding section.) We denote by $\hat{P}$ the simple polygon that describes the boundary of P and $\ddot{P}$ the set of vertices (corners) of P. Every point x of P has an associated *visibility polygon*, denoted $V(x, P)$, consisting of the set of all points y of P that are "visible" from x (that is, the line segment $\overline{xy}$ lies entirely within P). A set of points X *guards* a subset $Y \subseteq P$ if $Y \subseteq \bigcup_{x \in X} V(x, P)$. (In the special cases where $X \subseteq Y$ and $Y = P$ (respectively, $Y = \hat{P}$, $Y = \ddot{P}$) X is called an *interior* (respectively, *boundary*, *vertex*) *guard set*.)

It is well known that a guard set of size $\lfloor n/3 \rfloor$ is both necessary and sufficient to guard a gallery with n vertices, in the worst case. This is true even if it is only necessary to guard the boundary (or even just the vertices) of P, and even if guard points are restricted to vertices. Unfortunately, the problem of finding a minimum size guard set for a specific gallery is NP-hard, for all natural variations of the problem. (See [5, 6] for comprehensive surveys of basic results on art gallery guarding.)

In this paper we continue a line of research based on the intuition that galleries all of whose points are highly visible (which we make precise later but describe informally as "having no (small) nooks") must be easier to guard. The question was first raised by Kavraki *et al.* [3] in the context of randomized algorithms for robot motion planning. They asked if every compact simply connected region in Euclidean d-space that is ϵ-good (meaning every point is seen by a fraction ϵ of the entire region) can be guarded by a set of $g(\frac{1}{\epsilon})$ points, where g is some fixed polynomial. (In fact Kavraki *et al.* posed a somewhat more general question concerning regions with holes, which we will return to in the concluding section.)

Kalai and Matoušek [2] gave an affirmative answer to the question of Kavraki *et al.*, when $d = 2$. Specifically, they showed that $g(\frac{1}{\epsilon}) = C\frac{1}{\epsilon}\log\frac{1}{\epsilon}$ suffices for some (rather large) constant C. Their proof uses the concepts of VC-dimensionality (cf. [9]) and ϵ-nets (cf. [1, 4]). The essential result is that the set system formed by the visibility polygons for points in any compact simply-connected region in 2-space has bounded VC-dimension. Following the same approach as Kalai and Matoušek, Valtr [7] provided both a significantly tighter constant C as well as an affirmative (and asymptotically optimal) answer to the more general question concerning regions with holes. In a subsequent paper, Valtr [8] demonstrated that the question of Kavraki *et al.* does not have an affirmative answer in dimensions three and higher. Specifically, he showed how to construct $\frac{5}{9}$-good simply connected polyhedral regions in Euclidean d-space, for $d \geq 3$, that require more than any fixed constant number of guards.

We are concerned in this paper with the vertex and boundary guarding variants of the question of Kavraki *et al.* in the Euclidean plane. In particular, we address the question(s): if P is a simply connected polygonal region and every vertex (respectively, boundary point) of P is seen by at least a fraction ϵ of the vertices (respectively, total boundary) of P then how many vertices (respectively, boundary points), expressed as a function of $\frac{1}{\epsilon}$, suffice to guard $\ddot{P}$ (respectively, $\hat{P}$) in the worst case? These variant questions may inherit some of the practical motivation associated with the original question of Kavraki *et al.* However, their primary motivation here is that they provide a simpler setting in which to explore the apparent duality between degree of visibility and guard set size (expressed informally by the "no nooks" intuition).

While our questions are not directly addressed by the results of Kalai and Matoušek [2] and Valtr [7], their approach can be adapted to this different setting, exploiting the (easily confirmed) fact that the set system formed by the boundary visibility components associated with the boundary points of any simple polygon has constant VC-dimension. This implies an $O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ upper bound on the number of guards needed in the worst case.

In the remainder of this paper we show that the structure of boundary visibilty components can be exploited to give a explicit guard assignment scheme that uses $O(\frac{1}{\epsilon}\log\log\frac{1}{\epsilon})$ guards in the worst case. We develop this result for the more restricted question (finding small vertex guard sets) and discuss its generalization in the concluding section.

2 Preliminaries

Suppose that the polygonal region P has the property that each of its vertices is seen by (equivalently, sees) at least some fixed fraction ϵ of the the other vertices of P. The reader might well suspect that a greedy approach (iteratively select a vertex that guards the largest number of currently unguarded vertices) would lead to an optimal or near-optimal size vertex guard set for P. Unfortunately this intuition is incorrect; it is straightforward to modify the example presented by Kalai and Matoušek ([2], section 5) which shows that the analogous strategy applied to the interior guarding problem fails to provide a guard set whose size depends only on ϵ.

Our approach to guard placement is based on a hierarchical boundary subdivision scheme. At any time we consider the boundary of P to be partitioned into a sequence of *(boundary) fragments* each of which consists of a contiguous sequence of vertices of P. The simplicity of $\hat{P}$ implies that the cyclic (say, clockwise) order of vertices visible from any point $p \in P$ respects the cyclic (clockwise) order of the vertices on $\hat{P}$. In particular, the points of any fixed fragment A visible from p occupy a contiguous interval of angles centered at p, and we can speak of the "leftmost" and "rightmost" visible point (or vertex) of A from p.

Fragment A is visible from fragment B if there exist points $a \in A$ and $b \in B$ such that a and b are visible in P. We are interested in identifying *extreme points* of visibility: the leftmost and rightmost vertices of A visible from B, and vice versa. Obviously, every pair of fragments has associated with it at most four such extreme vertices.

Our strategy can now be summarized as follows: We start with a polygonal region P with n vertices and assume that every vertex of P is visible from at least $m = \epsilon n$ of the other vertices of P We view the entire boundary of P split at some arbitrary point as an initial fragment. Thereafter we proceed through some number $t(\epsilon)$ phases, where in each phase we partition all of the fragments from the preceeding phase into new smaller fragments, the degree of fragmentation depending on ϵ and the phase number. All fragments with a common parent are called *sibling fragments.* Associated with each phase is a allocation of guards to the extreme visible vertices associated with every pair of sibling fragments generated in that phase. Our main result follows by demonstrating that for a suitable choice of t and phase fragmentation factors, the resulting set of guards has size $O(\frac{1}{\epsilon}\log\log\frac{1}{\epsilon})$ and forms a vertex guard set for P.

3 Hierarchical Fragmentation and Guard Placement

Suppose that polygon P has n vertices and that every vertex of P is visible from at least $m = \epsilon n$ other vertices of P. Let $t = \lceil \log\log\frac{1}{\epsilon}\rceil$. The hierarchical fragmentation process involves t phases. Let f_i denote the total number of fragments after phase i. Then $f_0 = 1$ and, for all i, $1 \leq i \leq t$, we choose

$$f_i = 4t \cdot 2^{2^t - 2^{t-i} - t + i + 1}.$$

Thus, in phase i, every fragment from phase $i-1$ is split into s_i pieces each of size at most $\lceil n/f_i \rceil$, where $s_i (= f_i/f_{i-1})$ satisfies

$$s_1 = 2^{2^{t-1}+1} \cdot t/2^{t-2}$$

and, for $1 < i \leq t$,

$$s_i = 2^{2^{t-i}+1}.$$

In phase i guards are placed at the extreme visible vertices associated with each pair of sibling fragments introduced in that phase. By our observations above the number of guards introduced in phase i is less than $4f_{i-1} \cdot s_i^2/2$. Provided $t \geq 4$ this is at most $32t2^{2^t-t+i}$, and hence the total number of guards assigned over all phases is less than $64t2^{2^t}$.

It remains to show that the guards selected in this fashion form a vertex guard set for P. We will prove this by demonstrating that if any vertex x of P is not seen by any of the guards then x sees fewer than m vertices of P, contrary to our assumption. The key observation is the following:

Lemma 1 *If a vertex x is not seen by any of the guards then, for all i, $1 \leq i \leq t$, x sees at most $4i$ fragments at level i.*

The contradiction follows by taking $i = t$ and recalling that level t fragments have size at most $\lceil n/(4t \cdot 2^{2^t}) \rceil$.

To prove Lemma 1 we need to establish some adtional properties of the portions of boundary fragments visible from a fixed point on the boundary of P. A point x on the boundary of P is said to have a *tangent* to fragment A if there exists a chord $\overline{xy}$ in P such that A touches the interior of $\overline{xy}$ but neither of its endpoints (see Figure 1). A tangent $\overline{xy}$ is a left (respectively, right) tangent to A if x sees A to the right (respectively, left) of y.

Figure 1: A left tangent from x to fragment A

Note that for each successive pair A,B (in clockwise order) of fragments visible from x (ignoring the fragment containing x), either x has a right tangent to A or x has a left tangent to B.

Lemma 2 *Suppose vertex x has a left tangent to fragment A and the next visible fragment B to the right of A is a sibling of A that does not contain x. Then x sees an extreme point of B's visibility on A.*

Proof: There are four cases to consider depending on whether A is visible in a sector of size less than or greater than or equal to π at x, and whether A occludes B or B occludes A at the transition in their visibility. These are illustrated in Figures 2 through 5. In each case it is easy to see that the leftmost point a of B's visibility on A lies somewhere in the interval between points a_L and a_R, the left and rightmost points if x's visibility on A. Furthermore if a is not visible from x then either a_L is not visible from x or a is not the leftmost point on A visible from B, either of which contradict our assumptions. ■

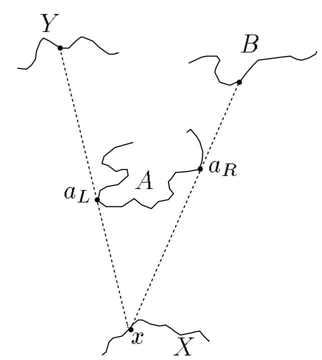

Figure 2: Case 1.

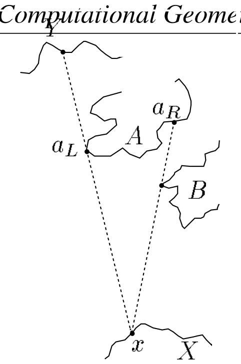

Figure 3: Case 2.

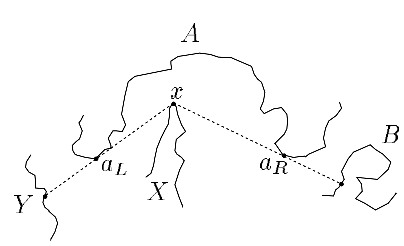

Figure 4: Case 3.

Of course, Lemma 2 remains true if left and right are interchanged throughout. It follows then that if x has a tangent to a fragment A at some level i then either x

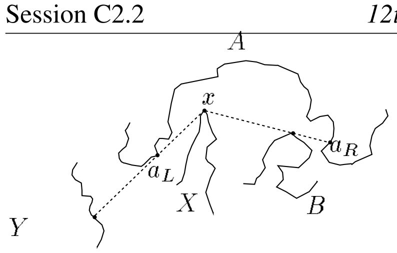

Figure 5: Case 4.

sees only one of the child fragments of A at level $i+1$ or x sees a guard associated with one of these child fragments. Since at any level there are at most two fragments seen by x without having a tangent from x, it follows that each level increases the total number of fragments visible from x by at most four. Lemma 1 follows by induction on the level index.

We summarize our main result in the following:

Theorem 3 *If the simply connected polygonal region P has the property that each of its vertices sees at least a fraction ϵ of the other vertices of P then there exists a set of at most $64\frac{1}{\epsilon}\log\log\frac{1}{\epsilon}$ vertices of P that guard the entire vertex set of P.*

4 Related Questions

If we guide our boundary fragmentation process by boundary length (measured as Euclidean distance or in terms of an arbitrary non-negative weight function defined on the boundary) rather than number of vertices then the results of the preceding section generalize to the following:

Theorem 4 *If the simply connected polygonal region P has the property that each of its boundary points sees a portion of the boundary whose measure is at least a fraction ϵ of the total measure of the boundary of P, then there exists a set of at most $64\frac{1}{\epsilon}\log\log\frac{1}{\epsilon}$ points on the boundary of P that guard the entire boundary of P.*

An obvious open question is whether Theorems 3 and 4 can be strengthened by making the size of the guard set $O(\frac{1}{\epsilon})$. (The reader will of course appreciate that the size of the guard set must be at least $\frac{1}{\epsilon}$ in the worst case.)

Another natural direction to extend this work is to consider polygonal regions with holes. While our proof can be modified to accomodate polygonal regions with h holes at a cost of a multiplicative factor of h^2 on the number of guards, this is far removed from the (asymptotically optimal) $\log h$ dependency demonstrated by Valtr [7] for the analogous (and possibly harder) interior guarding problem.

Finally, it seems tempting to move from guarding the boundary of simply connected polygonal regions to the analogous question for polyhedral regions. Unfortunately, just as Valtr [8] showed for the corresponding interior guarding problem, for any k there exist simply connected polyhedral regions P, each of whose vertices see a fixed fraction (say $1/4$) of the other vertices of P, for which no subset of k vertices suffice to guard all of the vertices of P.

The details on these and other related results (including algorithmic issues) will appear in the full version of this paper.

Acknowledgements

I would like to express my appreciation to both Tom Shermer and Jiří Matoušek for helpful discussions.

References

[1] D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete and Computational Geometry*, 2:127-151, 1987.

[2] G. Kalai and J. Matoušek. Guarding galleries where every point sees a large area. *Israel Journal of Mathematics*, 101:125-139, 1997.

[3] L. Kavraki, J-C Latombe, R. Motwani and P. Raghavan. Randomized query processing in robot motion planning. *Proc. 27th ACM Symposium on Theory of Computing*, pp. 353-362, 1995.

[4] J. Komlós, J. Pach and G. Woeginger. Almost tight bounds for ϵ-nets. *Discrete and Computational Geometry*, 7:163-173, 1992.

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

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

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

[8] P. Valtr. On galleries with no bad points. *Discrete and Computational Geometry*, 21:193-200, 1999.

[9] V. N. Vapnik and A. Ya. Chernonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16:264-280, 1971.