Recent Results in Art Galleries

THOMAS C. SHERMER

Invited Paper

Two points in a polygon are called visible if the straight line segment between them lies entirely inside the polygon. The ***art gallery problem*** *for a polygon P is to find a minimum set of points G in P such that every point of P is visible from some point of G. This problem has been shown to be NP-hard by Lee and Lin [71]. However, Chvátal showed that the number of points of G will never exceed $\lfloor n/3 \rfloor$ for a simple polygon of n sides [21]. This latter result is referred to as the art gallery theorem.*

Many variations on the art gallery problem have been studied, and work in this area has accelerated after the publication of the monograph of O'Rourke [92], which deals exclusively with this topic.

This paper provides an introduction to art gallery theorems, and surveys the recent results of the field. The emphasis is on the results rather than the techniques. In addition, this paper examines several new problems that have the same geometric flavor as art gallery problems.

I. Introduction

A. Definitions

This section contains necessary definitions, some background on art galleries, and a discussion of the scope of this paper. We begin with the definitions, following O'Rourke [92].

A *polygon* is generally defined as an ordered sequence of at least three points $v_1, v_2, \cdots, v_n$ in the plane, called vertices, and the n line segments $\overline{v_1v_2}, \overline{v_2v_3}, \cdots, \overline{v_{n-1}v_n}$, and $\overline{v_nv_1}$, called edges. A *simple* polygon is then a polygon with the constraint that nonconsecutive edges do not intersect. A simple polygon is a Jordan curve, and thus divides the plane into three subsets: the polygon itself, the (bounded) interior, and the (unbounded) exterior. However, we will henceforth use the term "polygon" to refer to "simple polygon plus interior." Polygons are thus closed, bounded sets in the plane.

A polygon P is said to be covered by a collection of subsets of P if the union of these subsets is exactly P. The collections of subsets is called a cover of P. A cover of P is called a decomposition if the intersection of each pair of subsets in the cover has zero area. A *triangulation* of a polygon is a decomposition of the polygon into triangles without adding vertices. This is done by chopping the polygon with *diagonals* (line segments between nonadjacent vertices). A polygon and one of its triangulations are shown in Fig. 1. A *triangulation graph* of a polygon P is the graph on the vertices of P where two vertices are joined if they share an edge, or are the endpoints of a diagonal in a fixed triangulation. The class of polygon triangulation graphs is the same as the class of *maximal outerplanar graphs.*

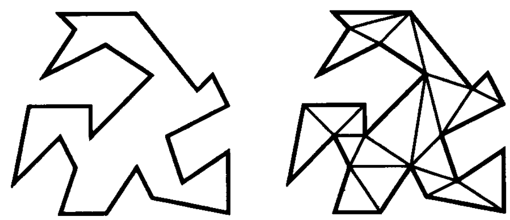

Fig. 1. A polygon and one of its triangulations.

Many algorithms in computational geometry incorporate a polygon triangulation step. Recently, Chazelle presented an algorithm that will compute a triangulation of a polygon in $O(n)$ time [15]; the algorithmic complexity results presented here have been reanalyzed in the light of this result. Consequently, the running times quoted in this paper often do not agree with that presented in the paper to which the algorithm is attributed. In particular, whenever a $\log\log n$ term is not present in a complexity result in this paper, but is in the original source, it is because Chazelle's algorithm has been substituted for the previously-best $O(n \log\log n)$ algorithm of Tarjan and van Wyk [117].

Let x and y be two points in a polygon P. We will say that x and y are visible if the line segment $\overline{xy}$ does not intersect the exterior of P. In Fig. 2, the point a is visible to b and c, but not d. Visible points are said to see one another. The set of all points of P visible from x is a polygon, called the visibility polygon of x, and is denoted $V(x, P)$.

We will distinguish some sets of points in a polygon by calling them guard sets. The individual elements of a guard

Manuscript received October 30, 1990; revised March 20, 1992. This work was supported by the Natural Sciences and Engineering Research Council of Canada and the Simon Fraser University Office of the President.

The author is with the School of Computing Science, Simon Fraser University, Burnaby, BC, Canada V5A 1S6.

IEEE Log Number 9204143.

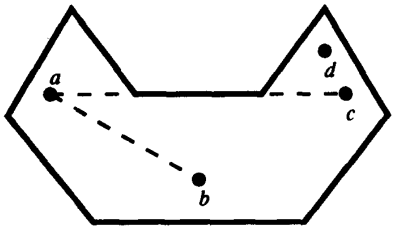

Fig. 2. Point a can see b and c, but not d.

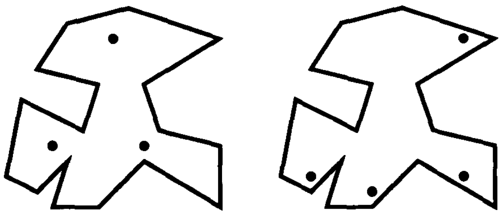

Fig. 3. A covering guard set; a hidden set.

set are called guards. If all of the points in a guard set are vertices of P, then G is called a vertex guard set, and the elements of G are called vertex guards. Otherwise G is called a point guard set, and its elements point guards. Other types of guards will be discussed later.

A guard set G is said to cover a polygon P if the collection of sets $\{V(g, P) \mid g \in G\}$ covers P. The points in the polyon on the left in Fig. 3 are a covering guard set. We will later see that the definition of covering for guard sets is a simple generalization of the usual definition of covering given above. The *art gallery problem* for a polygon P is to find a minimum-cardinality covering guard set G for P. It is so called because one envisions the polygon P as the floor plan of an art gallery, and the points of G as locations to place guards, so that every part of the art gallery is seen by at least one guard. We use $g(P)$ to denote the cardinality of a minimum covering guard set for the polygon P.

A concept similar to that of covering guard sets is *hidden sets*. A hidden set is a set of points H in a polygon such that no pair of points of H is visible. The points in the polygon on the right in Fig. 3 are a hidden set. Hidden sets are known in the mathematics literature as visually independent sets and are related to "property P_m" [65].

An *orthogonal polygon* is a polygon with edges that alternate between horizontal (zero slope) and vertical (infinite slope). Orthogonal polygons have also been called isothetic and rectilinear. Orthogonal polygons are an important subclass of polygons which arise in many computing applications, owing to the ease with which they are represented and manipulated, and to the design of many machines (such as image scanners and plotting devices) that are used in these applications. Restricting the polygons considered in the art gallery problem to orthogonal polygons creates an interesting subclass of problems, and has led to many results.

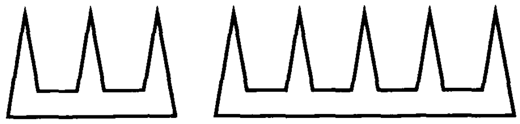

Fig. 4. Comb polygons.

B. Art Gallery History

The original art gallery problem, posed in conversation by Klee to Chvátal, is to find the smallest number of point guards necessary to cover any polygon of n vertices; this number will be denoted $g(n)$ (not to be confused with $g(P)$ as defined above). In terms of galleries, $g(n)$ is the minimum number of guards necessary to supervise any gallery with n walls.

Chvátal quickly proved that $g(n) = \lfloor n/3 \rfloor$, a result which has come to be known as the art gallery theorem [21]. First, he showed that $g(n) \geq \lfloor n/3 \rfloor$, by exhibiting the class of polygons now known as comb polygons; examples of comb polygons are shown in Fig. 4 for $n = 9$ and $n = 15$. Comb polygons exist for any n that is a multiple of 3. Each comb polygon requires $n/3$ guards, as no one guard can see into any two "teeth" (upward triangular regions) of the comb, and there are $n/3$ such teeth.

Next, Chvátal showed that $g(n) \leq \lfloor n/3 \rfloor$, by a relatively complex inductive argument on triangulation graphs of polygons. Fisk later gave the following more concise proof of this inequality [47]: First, triangulate the polygon. Next, three-color the vertices of the triangulation graph: assign each vertex one of three different colors, so that no two vertices which are adjacent in the graph have the same color. Each triangle of the graph, which corresponds to a triangle of the triangulation, will have one vertex of each color. Furthermore, every point of a triangle is visible to each vertex of that triangle. Therefore, choosing any of the three color classes will result in a set of vertices from which every point of every triangle, and thus every point of the polygon, is visible (i.e., each color class is a covering vertex guard set). By the pigeonhole principle, the smallest of these color classes will contain at most $\lfloor n/3 \rfloor$ vertices.

Later, Lee and Lin showed that the art gallery problem for polygons (given a polygon P, find the minimum number of guards necessary to cover P) is NP-hard [70], by reduction from Boolean three-satisfiability. Their result is for vertex guards, and this was extended to point guards by Aggarwal [2]. The reader unfamiliar with complexity theory is referred to the book of Garey and Johnson [50].

Although Lee and Lin's result implies that it is impractical to find a minimum set of guards (i.e., to find $g(P)$ guards) for a given polygon, Avis and Toussaint showed that it is possible to find a set of $g(n)$ guards for a polygon in polynomial time [7]. Algorithms for finding such guard sets are called guard placement algorithms. Most guard placement algorithms work by imitating upper-bound art gallery proofs, and Avis and Toussaint's algorithm is no exception, being an algorithmic imitation of Fisk's proof.

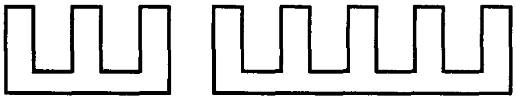

Fig. 5. Orthogonal comb polygons.

C. Orthogonal Art Galleries

Kahn, Klawe, and Kleitman investigated the art gallery problem restricted to orthogonal polygons [63]. They exhibit the orthogonal comb polygons of Fig. 5, establishing that $orth(n) \geq \lfloor n/4 \rfloor$ ($orth(n)$ denotes the maximum number of guards necessary for any orthogonal polygon of n vertices). They prove a matching upper bound, $orth(n) \leq \lfloor n/4 \rfloor$, in the same manner as Fisk proved the original art gallery theorem, but they decompose the polygon into *convex quadrilaterals* rather than triangles, and then four-color the quadrilateralization graph, so that each quadrilateral has one vertex of each of the four colors. The bulk of their paper is devoted to proving that every orthogonal polygon has a decomposition into convex quadrilaterals.

Edelsbrunner, O'Rourke, and Welzl gave an $O(n)$ point guard placement algorithm for orthogonal polygons, based on L-shaped partitioning [37]. Lubiw and Sack and Toussaint have presented other linear placement algorithms based on quadrilateralization [76], [100], [103].

D. Importance of Art Galleries

Art gallery problems are studied by computing scientists because they are fundamental visibility problems, and visibility is a central issue in many computing applications. Application areas for visibility include robotics [69], [123], motion planning [75], [80], vision [113], [124], graphics [79], [17], CAD/CAM [12], [38], computer-aided architecture [34], [99], and pattern recognition [5], [118]. Other reasons that art gallery problems are studied are that they are a continuous form of classical facility-location problems, have a simple formulation, and require an interesting interplay of graph theory, geometry, and computing science in their solution.

The monograph of O'Rourke [91] is devoted to art gallery problems and contains well-written detailed expositions of the results mentioned above, and the techniques used in their proofs. Since the publication of this book, activity in art gallery problems has rapidly increased, yielding many new theorems and algorithms. This paper is an attempt to collect these recent results into one place. However, we do not intend this paper to be a tutorial on the proof techniques used, and hence provide only a few details about the methods.

E. Organization of Paper

The remainder of this paper is organized into six sections. Section II contains results about different types of guards. Section III contains covering results, and Section IV is about covering the outsides of polygons. Section V contains results on visibility graphs, and Section VI contains problems which are not strictly art gallery problems, but have the same geometric feel. Conclusions are drawn in Section VII.

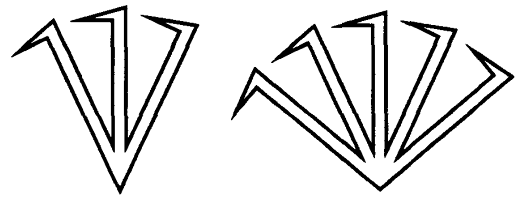

Fig. 6. Polygons requiring $\lfloor n/4 \rfloor$ edge guards.

II. Generalized Guards

In this section, we consider some variations of the art gallery problem that arise when specified subsets of the polygon, rather than just points, are allowed as elements of guard sets. A point will be called visible to such a subset if it is visible to some point in that subset. This notion of visibility from a subset is known as weak visibility, in contrast to *strong visibility*, where a point is called visible to a subset if it is visible to every point of the subset [6].

More formally, if R is a subset inside a polygon P, we let $V(R, P) = \{p \in P \mid \exists r \in R$ such that p and r are visible$\}$. We are still interested in finding minimum-cardinality covering guard sets, but now the individual guards will be various types of subsets. We are concerned only with the number of guards, and not with the sizes of the individual guards. Typical types of subsets used as guards are convex sets or polygon edges.

This branch of variations on the art gallery problem was started by Toussaint in 1981, when he asked how the art gallery theorem would change if guards were allowed to patrol individual edges of a polygon rather than continually standing at the same point. He wanted to know the guarding function $g^E(n)$, the minimum number of *edge guards* necessary to cover any polygon of n vertices.

Toussaint's conjecture is that if a small number of polygons are excluded, $g^E(n) = \lfloor n/4 \rfloor$. To lend weight to this conjecture, he exhibited the polygon class illustrated in Fig. 6, which establishes that $g^E(n) \geq \lfloor n/4 \rfloor$. Two types of polygons are known that require more than $\lfloor n/4 \rfloor$ edge guards. These polygons, proposed by Paige and Shermer, require $\lfloor (n+1)/4 \rfloor$ guards, and are shown in Fig. 7. However, these polygons are thought to be isolated exceptions, hence the qualification in Toussaint's conjecture.

O'Rourke was the first to make progress on Toussaint's conjecture. Although he was unable to establish an upper bound on $g^E(n)$, he was able to prove an upper bound on $g^M(n)$, the minimum number of *mobile guards* necessary for any polygon of n vertices [90]. Mobile guards are a slightly more general version of edge guards; each mobile guard can patrol either an edge or a diagonal of the polygon. Thus, every edge guard is a mobile guard, and $g^M(n) \leq g^E(n)$.