Art Gallery and Illumination Problems

Jorge Urrutia[1]
Instituto de Matemáticas
Univresidad Nacional Autónoma de México
Mexico D.F. MEXICO
E-mail: urrutia@matem.unam.mx

April 28, 2004

[1]The author was supported by CONACYT Grant 37540-A, and PAPIIT (UNAM) Proyecto IN110802.

Preface

In 1973, Victor Klee posed the following question:

> How many guards are necessary, and how many are sufficient to patrol the paintings and works of art in an art gallery with n walls?

This wonderfully naïve question of combinatorial geometry has, since its formulation, stimulated an increasing number of of papers and surveys. In 1987, J. O'Rourke published his book *Art Gallery Theorems and Algorithms* which has further fueled this area of research.

The present book is being written almost 10 years since the publication of O'Rourke's book, and the need for an up-to-date manuscript on Art Gallery or Illumination Problems is evident. Some important open problems stated in O'Rourke's book, such as ... have been solved. New directions of research have since been investigated, including: watchman routes, floodlight illumination problems, guards with limited visibility or mobility, illumination of families of convex sets on the plane, guarding of rectilinear polygons, and others. In this book, we study these results and try to give a complete overview of all the results known to us. We hope that this book will provide a renewed source of inspiration towards the study of Art Gallery Problems.

Jorge Urrutia, September 1996.
Ottawa, Canada

Contents

Chapter 1

Art Gallery Problems

1.1 Introduction

Illumination problems have always been a popular topic of study in mathematics. For example, it is well known that the boundary of any smooth compact convex S set on the plane can always be illuminated using three light sources located in the complement of S. An easy proof of this can be obtained by enclosing a smooth convex set within a triangle, then placing a light at each vertex of T; see Figure 1.1.

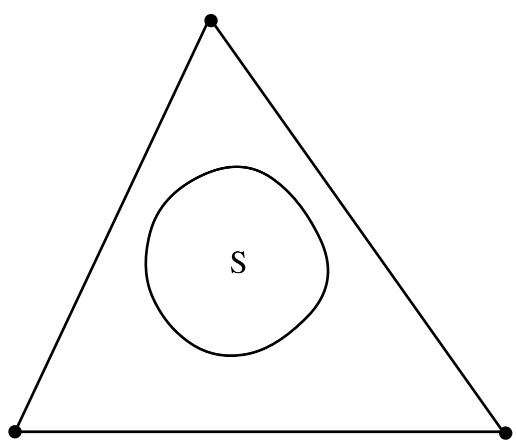

Figure 1.1: Three lights suffice to illuminate a compact convex set.

One famous—and until recently open—problem on illumination is attributed to Ernst Strauss (see E. G. Strauss and V. Klee [83]), who, in the early fifties posed the following problem:

Suppose that we live in a two-dimensional room whose walls form a simple closed polygon P and each wall is a mirror.

1. Is it true that if we place a light at any point of P, all of P will be illuminated using reflected rays as well as direct rays?

2. Is there necessarily a point from which a single light source will illuminate the entire room using reflected rays as well as direct rays?

The first part of Strauss's problem was recently proved to be false by G. W. Tokarsky [127]. Tokarsky's proof is surprisingly simple, using basic concepts of geometry that are easily understandable. We refer the interested reader to Tokarsky's original manuscript which is clearly written and a pleasure to read. The second part of Strauss's conjecture, though, remains open. It would be nice if a "simple" proof for it could be obtained.

More closely related to our topic of interest here is a question posed by V. Klee during a conference in Stanford in August 1976. Klee's question was:

> How many guards are always sufficient to guard any polygon with n vertices?

Soon after, V. Chvátal established what has become known as *Chvátal's Art Gallery Theorem*, namely: that $\lfloor \frac{n}{3} \rfloor$ guards are always sufficient and occasionally necessary to cover a simple polygon with n vertices [28].

Since the publication of this original result, a tremendous amount of research on illumination or Art Gallery problems has been carried out by mathematicians and computer scientists. In 1987, J. O'Rourke [104] published *Art Gallery Theorems and Algorithms*, the first book dedicated solely to the study of illumination problems of polygons on the plane. The publication of O'Rourke's book further fueled the study of Art Gallery type problems, and many variations to the original Art Gallery Theorem have since been studied. Two thorough survey papers have also been written on this subject, one in 1992 by T. Shermer [115], and a second one in 1996 by J. Urrutia [131]. In this book, we present most of the material known to us on Art Gallery or illumination problems. Following J. O'Rourke's lead, we have tried to write this book in a manner that makes the material presented here accessible to a wide audience. The only background required for this book is basic knowledge of algorithms, graph theory and data structures. This should make this book both accessible to senior undergraduate students and suitable for courses at the master's level. Numerous open problems are presented, making the book also useful to researchers working in this area.

1.2 Basic terminology

A *polygon* P is an ordered sequence of points $p_1, \ldots, p_n$, $n \geq 3$, called the vertices of P together with the set of line segments joining p_i to p_{i+1}, $i = 1, \ldots, n-1$ and p_n to p_1, called the edges of P. P is called *simple* if any two non-consecutive edges do not intersect. A simple polygon divides the plane into two regions, an unbounded one called the exterior region and a bounded one, the interior. Henceforth to simplify our presentation, the term polygon will be used to denote simple polygons together with their interior.

A polygon is called *orthogonal* if all its edges are parallel to either the x-axis or the y-axis.

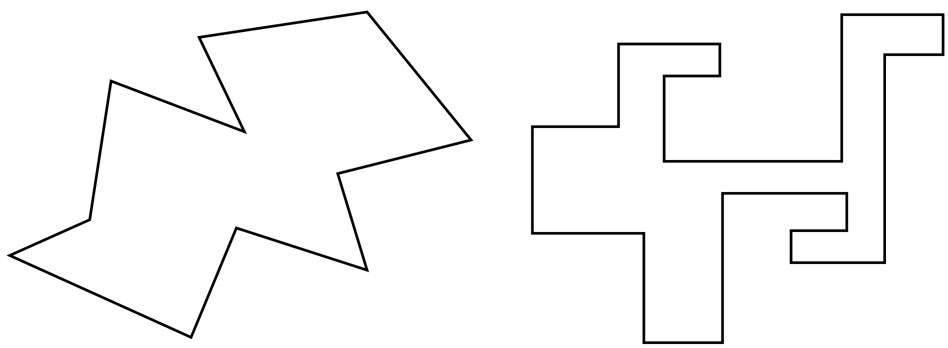

Figure 1.2: A simple and an orthogonal polygon.

A *graph* G consists of a set of elements V called the vertices of G, together with a set E of pairs of vertices of G called the edges of G. Two vertices u and v of G are called *adjacent* if the pair $\{u, v\}$ is an element of E. The *degree* of a vertex v of G is the number of vertices of G adjacent to v. A graph G is planar if it can be drawn on the plane in such a way that its vertices are represented by points on the plane, and each edge $\{u, v\}$ of G is represented by a simple curve joining points representing u and v. Moreover, two edges of G may only intersect at their endpoints. In Figure 1.3 we show two graphs; the first one is planar, and the second one is not.

A *path* of a graph G is a sequence of distinct vertices $v_1, \ldots, v_k$ such that v_i and v_{i+1} are adjacent in G, $i = 1, \ldots, k-1$, $k \geq 2$. A *cycle* of G is a path $v_1, \ldots, v_k$ together with the edge $\{v_k, v_1\}$, $k \geq 3$. A graph $G(V, E)$ is called connected if for every pair of vertices u and v of $G(V, E)$, there exists a path $u = v_1, \ldots, v_k = v$ starting at u and ending at v; otherwise $G(V, E)$ is called disconnected. A graph G is called a *tree* if it is connected and contains no

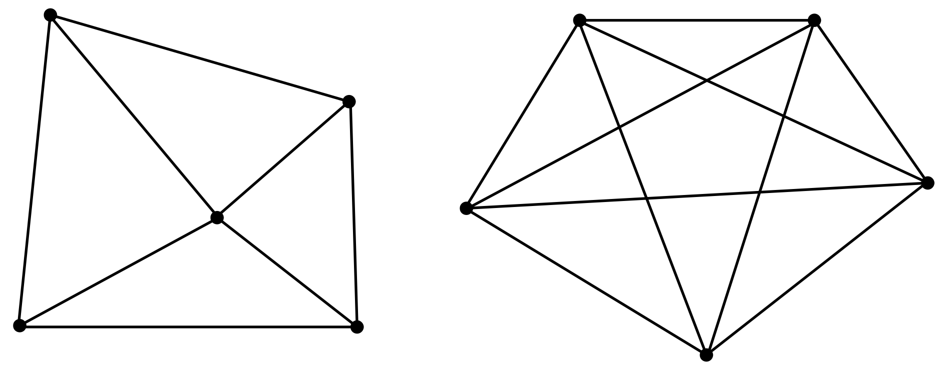

Figure 1.3: Two graphs with five vertices.

cycles.

A *triangulation* T of a polygon P is a partitioning of P into a set of triangles with pairwise disjoint interiors in such a way that the edges of those triangles are either edges or diagonals of P joining pairs of vertices.

The following result is easy to prove:

Theorem 1.2.1 *Any polygon can be triangulated; moreover any triangulation of a polygon P with n vertices contains exactly $n-2$ triangles.*

Triangulations of polygons play a central role in the study of art gallery problems. The problem of finding efficient algorithms to triangulate polygons has received much attention in computational geometry. In 1978, Garey, Johnson, Preparata and Tarjan [60] obtained the first $O(n \ln n)$ time triangulation algorithm. This result was improved in 1988 by Tarjan and van Wyk to $O(n \ln \ln n)$ [126]. Finally Chazelle [20] proved in 1990:

Theorem 1.2.2 *There is a linear time algorithm to triangulate simple polygons.*

The algorithmic details of Chazelle's triangulation algorithm are beyond the scope of this book, and will not be studied here.

Given a triangulation T of a polygon P, we can define a graph $GT(P)$ such that the vertices of $GT(P)$ are the vertices of P, and two vertices of $GT(P)$ are adjacent if they are connected by an edge of T.

The dual of a triangulation T of a polygon P is obtained by placing a vertex in the interior of each triangle of T, and connecting two vertices if their corresponding triangles share a common edge in T. It is easy to see that the dual graph of a triangulation of a polygon P is always a tree. In Figure 1.4 we show a polygon together with a triangulation T and its dual.

A *coloring* of the vertices of a graph G is an assignment of colors to its vertices in such a way that adjacent vertices receive different colors. The chromatic number of a graph G is the smallest integer k such that a coloring of G exists. In this case, we say that G has chromatic number k. The following result can be easily proved:

Theorem 1.2.3 *Let P be a polygon and T a triangulation of P. Then $GT(P)$ has chromatic number* 3.

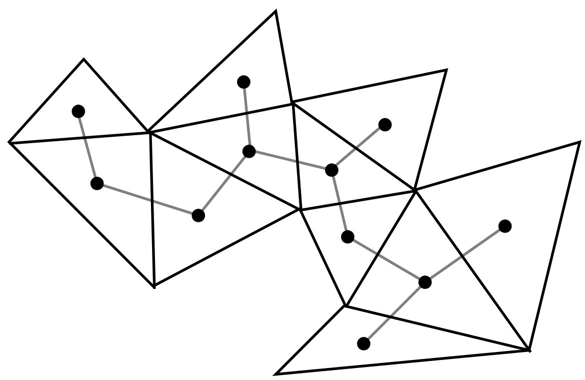

Figure 1.4: A triangulation of a polygon together with its dual.

Given two points p and q of a polygon P, we say that p is *visible* from q if the line segment joining p to q is totally contained in P.

A collection H of points of P *illuminates* or *guards* P if every point u of P is visible from a point p in H. The term *illuminates* follows the notion that if we place a light source that emits light *in all directions* on each element of H, P is totally illuminated. The use of the term *guard* follows the notion that if we station a guard at each element of H, all of P is guarded. The reader may easily verify that to illuminate the orthogonal polygon in Figure 1.2 we need four lights. We remark at this point that the terms *illumination* and *guarding* will be used interchangeably in this manuscript. Our choice of whether to "guard" or "illuminate" an object depends mainly on the term used in the original paper in which a particular result was proved.

The visibility graph $VG(P)$ of P is the graph whose vertex set is the set of vertices of P, two vertices u and v being adjacent in $VG(P)$ if they are visible in P. Visibility graphs were first introduced by Avis and ElGindy [6].

Avis and ElGindy gave an optimal (worst case) $O(n^2)$ time algorithm to compute the visibility graph of a polygon. Their result was later improved by Hershberger [72], who gave an $O(\mid E \mid)$ time algorithm, where $\mid E \mid$ is the number of edges of the visibility graph of P. One of the most important open problems in computational geometry is that of characterizing and recognizing visibility graphs.

For the case of planar graphs, we will use the following result by T. Nishizeki extensively:

Theorem 1.2.4 *(Nishizeki [97]): Any planar* 2*-connected graph* G *with* $n \geq 14$ *vertices and minimum vertex degree greater than or equal to* 3 *has a maximum matching with at least* $\lfloor \frac{n+4}{3} \rfloor$ *edges. When* $n \leq 14$*, the number of edges in such a matching is* $\lfloor \frac{n}{2} \rfloor$*.*

The proofs of these results are not of concern in this work; the interested reader can find them in many books on graph theory.

1.3 Exercises

Ex. 1.1 Prove that any tree with n vertices contains exactly $n-1$ edges.

Ex. 1.2 Prove Theorem 1.2.3.

Ex. 1.3 Prove that the graph $GT(P)$ obtained from a triangulation of a polygon P has at least two vertices of degree two.

Ex. 1.4 Prove that any polygon can be triangulated, and that any triangulation of a polygon with n vertices has exactly $n-2$ triangles.

Ex. 1.5 Let H be a tree such that the degree of all its vertices is at most 3. Show that there is a polygon P and a triangulation T of P such that the dual of $GT(P)$ is H. Moreover, show that there is a polygon P such that P has a unique triangulation T whose dual is H.

Ex. 1.6 Show that every convex polyhedron in R^3 can be decomposed into a set of convex tetrahedrons.

Ex. 1.7 Find a convex polyhedron with n vertices that can be decomposed into tetrahedrons in two different ways such that the first decomposition has a linear number of tetrahedrons, and the second partition has a quadratic number.

Chapter 2

Guarding Polygons

In this chapter, we study the problem of illuminating simple polygons on the plane. In section 2.1 we study illumination problems for simple polygons. In section 2.2 we study illumination of orthogonal polygons.

2.1 Chvátal's Art Gallery Theorem

We now give a proof of Chvátal's Art Gallery Theorem due to S. Fisk [56].

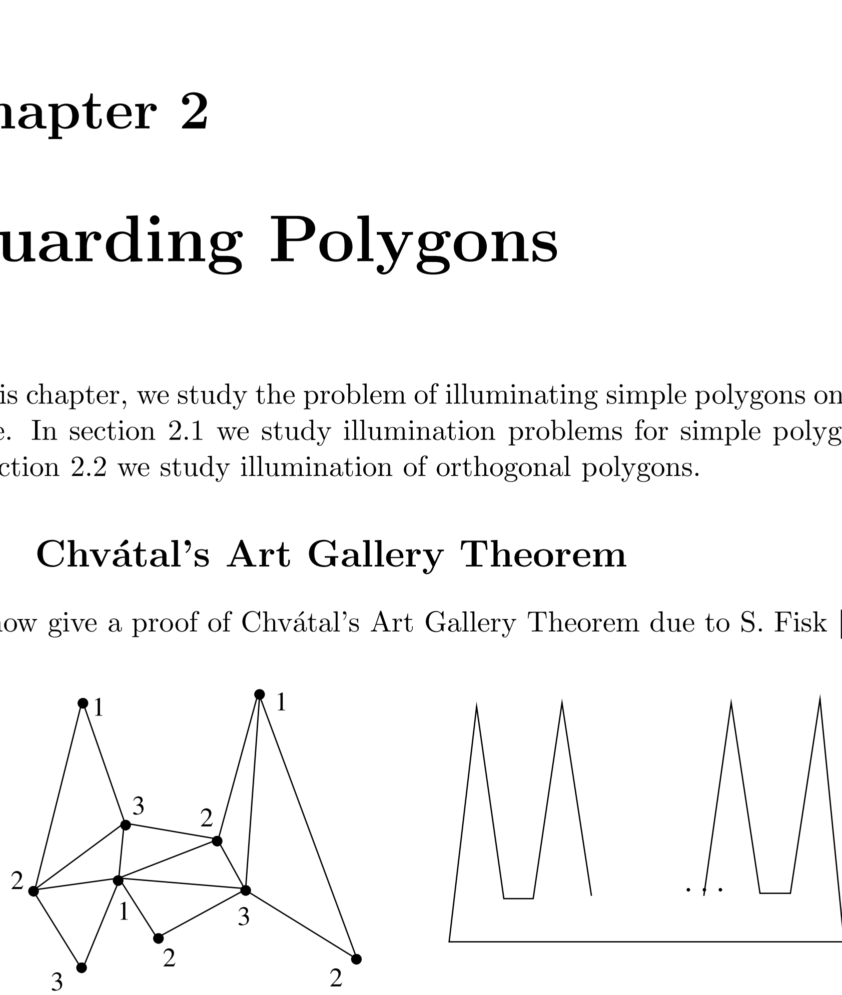

Figure 2.1: Illustration of proof of Chvátal's Art Gallery Theorem.

Theorem 2.1.1 $\lfloor \frac{n}{3} \rfloor$ *stationary guards are always sufficient and occasionally necessary to illuminate a polygonal art gallery with* n *vertices.*

Proof: Let P be a simple polygon with n vertices. Triangulate P by adding $n-2$ interior diagonals. (See Figure 2.1.)

By Theorem 1.2.3 we can color the vertices of P using three colors $\{1, 2, 3\}$ such that any two vertices joined by an edge of P or a diagonal in our triangulation receive different colors. This partitions the vertex set of P into three chromatic classes C_1, C_2 and C_3. Clearly one of our chromatic classes, say C_1 has at most $\lfloor \frac{n}{3} \rfloor$ vertices. Since the vertices of each triangle receive different colors, each of them has a vertex with color 1. It now follows that if we place a guard at all vertices with color 1, they guard all of P.

To see that $\lfloor \frac{n}{3} \rfloor$ guards are sometimes needed, consider the comb polygon $Comb_m$ with $n = 3m$ vertices presented in Figure 2.1. It is easy to see that to guard P_m we need at least m guards. ∎

Chvátal's proof of Theorem 2.1.1 more complicated than Fisk's. Nevertheless it is intersting to study it, as it exposes some intrinsic properties of triangulations of polygons that are not evident in Fisk's proof. Chvatal's proof also starts by triangulating a polygon P, and then selects a subset S of its vertices with at most $\lfloor \frac{n-3}{3} \rfloor$ elements such that each triangle in our triangulation has a vertex in S. Chvátal's proof is based in the following following result:

Lemma 2.1.1 *Any triangulation of a polygon with n vertices has a diagonal d that splits P into two polygons P_1 and P_2, such that P_1 has 5, 6, or 7 vertices, $n \geq 6$.*

The proof of this Lemma is left as an excercise. We now proceed by induction on the number of vertices of P. Observe first that if P has 3, 4 or 5 vertices, it can always be guarded with a single vertex. Suppose then that P has at least six vertices. Let T be any triangulation of P.

By Lemma we can find a diagonal d in T that splits P into two subpolygons P_1 and P_2 such that P_1 has at least 5 and at most 7 vertices. Several cases arise:

Case 1: We can choose d such that P_1 has five vertices. In this case, P_1 can be guarded with a single vertex, and since P_2 has $n-3$ vertices it can be guarded with with $\lfloor \frac{n-3}{3} \rfloor$ guards. This produces a total of $\lfloor \frac{n}{3} \rfloor$ guards that guard P.

Case 2 Suppose next that P_1 has six vertices labelled u, v, w, x, y, z in the clockwise direction such that u and v are the endvertices of d, see Figure **??**. If any vertex of P_1 guards all of it, then this vertex together with any guarding set of P_2 with at most $\lfloor \frac{n-4}{3} \rfloor$ vertices produces a guarding set of P with at most $\lfloor \frac{n-3}{3} \rfloor$ guards.

Let t be the triangle of T contained in P_1 that contains d. Notice that the third vertex of t cannot be w (or z), for in this case the diagonal uw cuts P into two subpoligons, one of which has five vertices $u,\ w,\ x,\ y,\ z$ (resp. $v,\ w,\ x,\ y, z$).

Suppose then that the third vertex of t is x. Observe that in this case, x is not adjacent to z, for otherwise x would guard all of P. It now follows that u guards the quadrilateral with vertices $u,\ x,\ y,\ z$.

Observe that $P_2 \cup t$ is a polygon with $n-3$ vertices, and by induction it has a subset S with at most $n-3$ of its vertices that guards it. By induction,we also have that one of the vertices of t is in S. Observe that if either of v or x is in S, then S together with u guards P. If u belongs to S, then $S \cup \{x\}$ also guard P. Our result follows.

The case when P_1 has seven vertices follows in a similar way and is left as an excercise. ■

The cases when P_1 has 6 and 7 vertices can be solved in a similar way, and are left as excercises.

The proof of this result will be studied later. This result gives rise to the study of algorithms that somehow take advantage of the structure of a polygon to obtain a better bound on the number of vertices required to guard them. For example, convex polygons can be guarded with one guard.

A vertex of a polygon P is called *reflex* if the internal angle of P at v is greater than π. We now prove:

Theorem 2.1.2 *Le P be a polygon with r reflex vertices, $r \geq 1$. Then r guards are always sufficient and occasionally necessary to guard P.*

The following result will be useful to prove Theorem 2.1.2.

Lemma 2.1.2 *Any polygon with r reflex vertices can be partitioned into at most $r+1$ convex polygons with disjoint interiors.*

Proof: Let $\{r_1, \ldots, r_k\}$ be the reflex vertices of P. For $i = 1, \ldots, k$ extend a ray starting at r_i that bisects the internal angle of P at r_i until it hits the boundary of P or a previously drawn ray; see Figure 2.2. An easy inductive argument can now be used to show that this set of rays divides P into $r+1$ convex polygons as desired. ■

We can now prove our result:

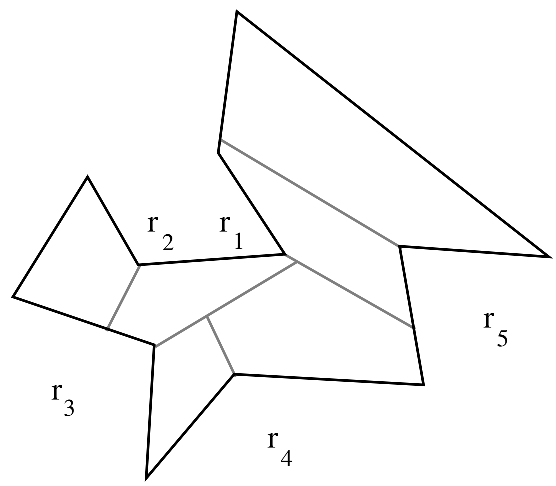

Figure 2.2: Partitioning a polygon into $r + 1$ convex polygons.

Proof of Theorem 2.1.2: We first show that r guards always suffice. Partition P as in the previous lemma. Notice that each of the convex polygons obtained is such that at least one of its vertices is a reflex vertex of P. Our result now follows.

A family of polygons for which r guards are necessary is shown in Figure 2.3. ∎

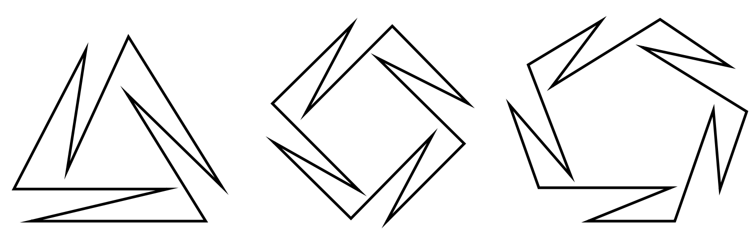

Figure 2.3: A family of polygons that require r guards.

2.2 Orthogonal Polygons

Of particular interest is the study of guarding problems for orthogonal polygons, that is; polygons whose edges are all parallel to the x- or y-axis.

Perhaps one of the motivations for the study of these polygons is that most "real life" buildings are, after all, "orthogonal". From a mathematical point of view, their inherent structure allows us to obtain very interesting and aesthetic results, and I believe this is an even stronger motivation to study them.

The first major result here is due to Kahn, Klawa and Kleiman [82]. They proved:

Theorem 2.2.1 *Any orthogonal polygon with n vertices can be guarded with at most $\lfloor \frac{n}{4} \rfloor$ guards.*

We will present three proofs for Theorem 2.2.1 The first one is new, and short. We then give two more proofs of the same theorem that although longer, are interesting in their own right, and demonstrate some interesting properties of partitioning of polygons into convex quadrilaterals and L-shaped polygons.

We start by proving the following result due to O'Rourke:

Lemma 2.2.1 *Let P be an orthogonal polygon with n vertices, r of which are reflex. Then $r = \frac{n-4}{2}$.*

Proof: Since P has n vertices, the sum of the internal angles of P is exactly $(n-2)\pi$. Notice that all the internal angles of P are of size $\frac{\pi}{2}$ or $\frac{3\pi}{2}$, depending on whether they are generated by a convex or reflex vertex respectively.

Thus if we have r reflex and $n-r$ convex vertices, we have:

$$(n-r)(\frac{\pi}{2}) + r(\frac{3\pi}{2}) = (n-2)\pi$$

Solving for r gives the desired result. ∎

A horizontal or vertical *cut* of P is the extension of a horizontal or vertical edge of P at a reflex vertex towards the interior of P until it hits the boundary of P. A horizontal or vertical cut is called an *odd cut* if it splits P into two non-empty polygons P_1 and P_2 with n_1 and n_2 vertices respectively such that at least one of n_1 or n_2 equals $4k+2$ for some k.

O'Rourke noted in his proof of the orthogonal art gallery theorem that if we can prove that any orthogonal polygon has an odd cut, then by an inductive argument, Theorem 2.2.1 would follow. Indeed, we can easily verify that if $n_1 + n_2 \leq n + 2$, and one of n_1 or n_2 equals $n_1 = 4k + 2$ then $\lfloor \frac{n_1}{4} \rfloor + \lfloor \frac{n_2}{4} \rfloor \leq \lfloor \frac{n}{4} \rfloor$.

Given an orthogonal polygon P, we label its edges as *top*, *right*, *bottom*, and *left* in the natural way, and then call a reflex vertex of P a *top-left* reflex vertex if the edges of P incident to it are a top and a left edge. *Top-right*, *bottom-right*, and *bottom-left* edges are defined in a similar way; see Figure 2.4.

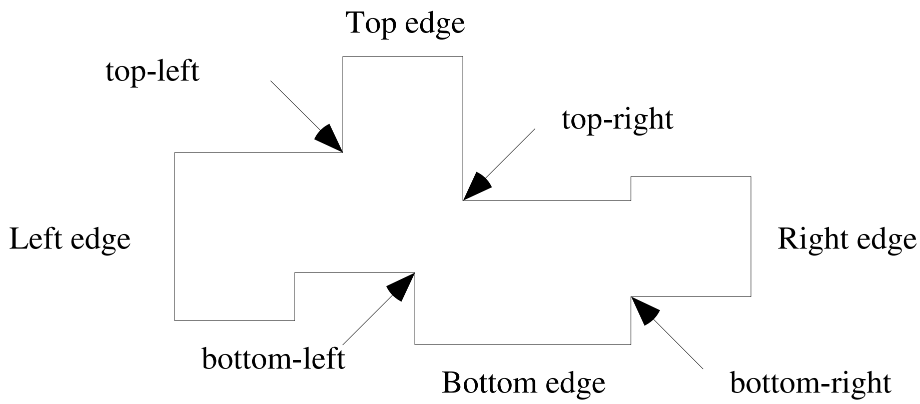

Figure 2.4: Classifying the edges and vertices of an orthogonal polygon.

We are ready to prove Theorem 2.2.1.

Proof: Split the set of reflex vertices of P into two sets, S_1 containing all the top-right and bottom-left reflex vertices of P, and S_2 containing all the top-left and bottom-right vertices. Since P has $\frac{n-4}{2}$ reflex vertices, one of S_1 or S_2 has at most $\lfloor \frac{n}{4} \rfloor$ vertices. Suppose it is S_1. If placing a light at every vertex of S_1 illuminates all of P, we are done. Suppose then that there is a point p in P not illuminated by S_1. Consider the longest horizontal line segment c containing p, and contained in P. Let e and f be the edges of P containing the endpoints of c; see Figure 2.5.

Consider the largest rectangle containing c and contained in P. Let e' and f' be edges of P that intersect the top and bottom edges of R respectively. Since p is not visible from any point in S_1, it follows that e and e' meet at the top-left corner point of R. Similarly f and f' meet at the right-bottom corner point of R.

Let q and q' be the top-right and bottom-left vertices of R. If they are vertices of P, it follows that P is R and there is nothing to prove.

Two cases arise:

1. Neither of q and q' are vertices of P.

2. Exactly one of them, say q, is a vertex of P.

If neither of q and q' are vertices of P, it is easy to see that we can generate two horizontal cuts that generate a rectangle contained in R and containing c as shown in Figure 2.5. It now follows that one of these two horizontal cuts is an odd cut of P.

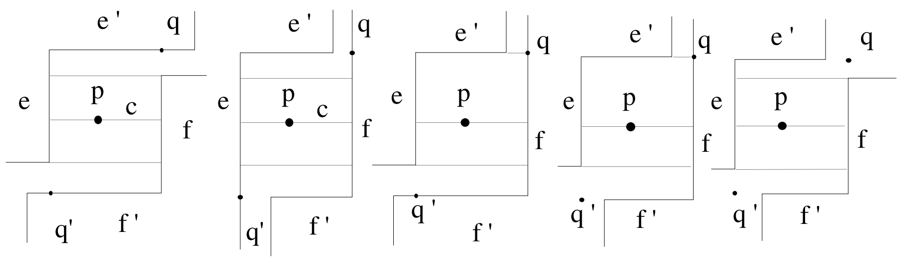

Figure 2.5: Up to symmetry, these are all the cases arising when none of q and q' are vertices of P.

Suppose then that only q is a vertex of P. If e and f' are properly contained in the left and bottom edge of R, then by extending the horizontal edge of P, incident to the bottom vertex of e, and the vertical edge of P incident to the left endpoint of f', we obtain a polygon with $n-4$ vertices that by induction can be guarded with $\lfloor \frac{n-4}{4} \rfloor$ guards. Since R can be guarded with a single guard, our result follows; see Figure 2.6.

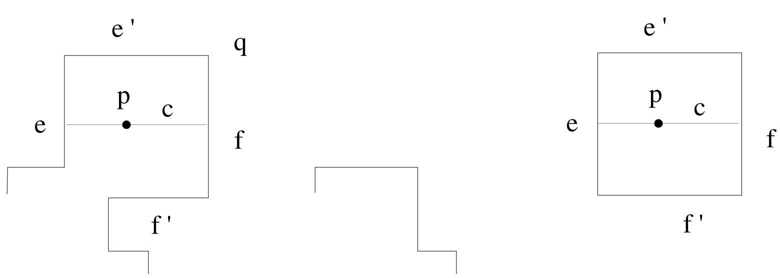

Figure 2.6: The case when e and f' are properly contained on the left and bottom edges of R.

Suppose then without loss of generality that f' properly contains the bottom edge of R and let w be the bottom vertex of e. Consider the horizontal line segment l joining w to a point in the base of R. Slide this segment until it hits a vertical edge of P or it reaches the left endpoint of f' or the leftmost endvertex of horizontal edge g incident to the bottom vertex of e. In the second case, if we reach the left endpoint of f', this point generates a vertical odd cut of P that leaves a polygon with six edges to its right, and another with $n-4$ on the left. In this case our result follows again by induction. The case when we reach the leftmost point of g follows in a similar way.

Suppose then that we hit a vertical edge of P. Let x be the highest vertex of P contained in l. Then we can generate two cuts of P at x, a horizontal cut h and a vertical cut h'. Let P' be the orthogonal subpolygon of P to the left of h' obtained when we cut P along h', and P'', the subpolygon on top of h obtained by cutting P along h. If P' has m vertices, P'' contains $m+2$ vertices, and thus either h or h' is an odd cut; see Figure 2.7. Our result follows. ∎

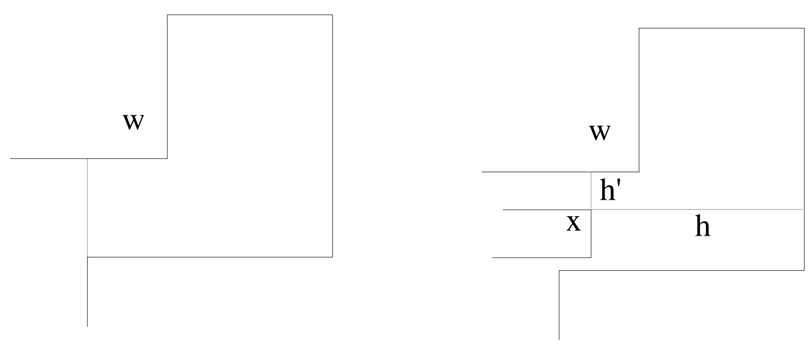

Figure 2.7: The case when f' properly contains the bottom edge of R.

2.3 Partitioning orthogonal polygons into convex pieces

A convex quadrilaterization of an orthogonal polygon P is a partition of P into a set of convex quadrilaterals with disjoint interiors such that the edges of these quadrilaterals are either edges of P or diagonals joining pairs of vertices of P; see Figure 2.8.

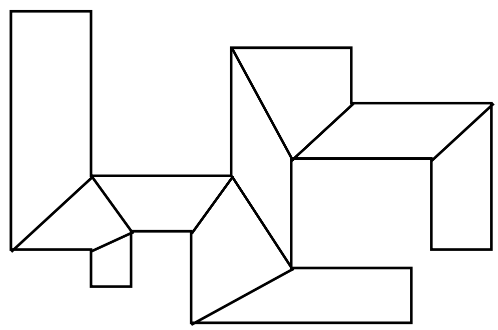

Figure 2.8: A quadrilaterization of an orthogonal polygon.

The proof of Theorem 2.2.1 is based on the following result:

Theorem 2.3.1 *Any orthogonal polygon is convex quadrilaterizable.*

We will return to the proof of this result in the following section; we now prove Theorem 2.2.1.

Proof of Theorem 2.2.1: In Figure 2.9, we show a family of orthogonal polygons that require $\lfloor \frac{n}{4} \rfloor$ guards.

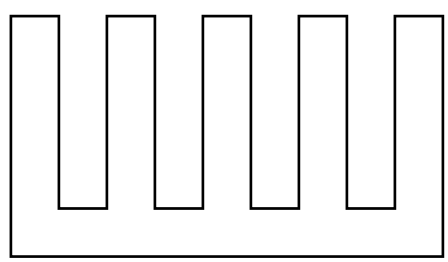

Figure 2.9: An orthogonal polygon that requires $\lfloor \frac{n}{4} \rfloor$ guards.

To show that $\lfloor \frac{n}{4} \rfloor$ guards are always sufficient, we first obtain a quadrilaterization Q of P. We next obtain a graph H by adding two diagonals connecting opposite vertices of every quadrilateral of Q, as shown in Figure 2.10.

We now show that this graph is 4-colorable. To this end, consider the dual graph Q^* of Q. The vertices of Q^* are the quadrilaterals of Q, two of which are adjacent if they share an edge of Q. Clearly Q^* is a tree. Remove from Q a quadrilateral corresponding to a leaf of Q^*. This produces a

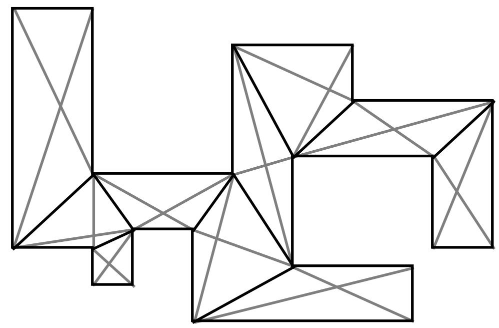

Figure 2.10: A 4-chromatic graph obtained from a quadrilaterization.

subgraph H' of H which, by an inductive argument, may be assumed to be 4-colorable. It is easy to see that this coloring can be extended to a 4-coloring of H. Notice that in the resulting 4-coloring, all the vertices of any quadrilateral of Q receive different colors, and thus the vertices of each chromatic class guard P. Place a guard at each vertex of the smallest chromatic class, and our result follows. ∎

We now show that every orthogonal polygon is convex quadrilaterizable. The proof we give here is due to Lubiw [91]. To this end, we define a *1-orthogonal polygon* P to be a polygon that satisfies the following conditions:

1. All the edges of P, with the possible exception of one distinguished edge denoted by e and called the *slanted edge*, are parallel to the x- or y axis.

2. With the possible exception of e, the edges alternate between horizontal and vertical.

3. All internal angles are less than or equal to $\frac{3\pi}{4}$.

4. The interior of the nose of the slanted edge contains no vertex of P.

5. P has an even number of edges.

Notice that orthogonal polygons are 1-orthogonal. In our definition, we do not explicitly forbid e from being parallel to the x- or y-axis. If P is orthogonal, then e may be chosen to be any of the edges of P. Thus if we show that any 1-orthogonal polygon is convexly quadrilaterizable, Theorem 2.2.1 will be proved.

The *nose* of the slanted edge e of P is the right triangle T towards the interior of P such that:

1. The hypotenuse of T is e.

2. The other two edges of T are parallel to the x- and y-axis.

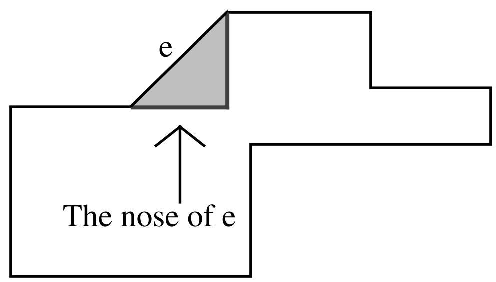

Figure 2.11: A 1-orthogonal polygon.

We now prove:

Theorem 2.3.2 *Any* 1*-orthogonal polygon is convexly quadrilateralizable.*

Proof: If P has four edges, it must be convex, and our result follows. We now show that if P has more than four edges, there is always a quadrilateral Q' whose removal splits P into smaller 1-orthogonal polygons. Using induction our result follows.

Assume without loss of generality that the slope of e lies between 0 and $\frac{\pi}{2}$, and that the vertices of e are labelled u and v such that u is the vertex of e with smallest y-coordinate. Let w be the third vertex of the nose T of e.

Since P has an even number of edges, and all internal angles are of size at most $\frac{3\pi}{2}$, both edges adjacent to e are horizontal. Let e' be the horizontal edge of P that precedes e in the counterclockwise order along the boundary of P. We now proceed to find the four vertices of Q'. The first two vertices of Q' are u and v. Let e" be the semi-open line segment joining v to w, open at w. Slide e" to the right until it hits a vertex of P. If e'' hits a whole edge of P, then the endpoints of this edge, together with u and v are the vertices of Q'. Suppose then that e'' hits a single vertex x of P; x is the third vertex of Q'.

Consider now the semi-open horizontal line segment h joining u to the vertical line through x; h does not contain u, but it contains its right end-point. Slide it down until it reaches a vertex of P, or it hits a horizontal edge f of P in the middle. In the first case, let y be the rightmost point that h hits. The fourth vertex of Q' is y. This case is illustrated in Figure 2.12.

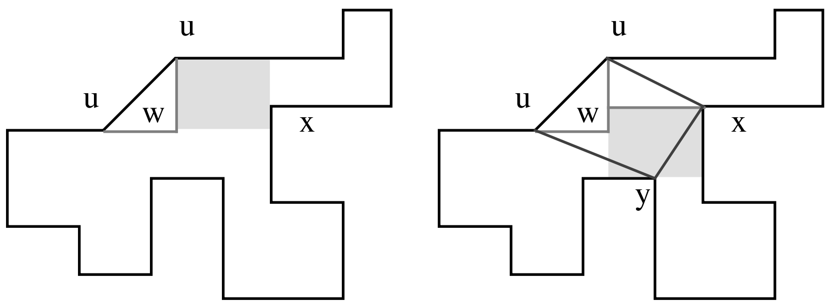

Figure 2.12: Finding Q', case 1.

We now deal with the case when h hits an edge f of P. This can only happen if x is the second end-vertex of e'. This case is illustrated in Figure 2.13. In this case, consider the semi-open vertical line segment h' joining x to f, closed at the bottom and open at the top. Slide h' to the right until it hits a vertex of P. If it hits more than one vertex, choose y, the lowest of these points, to be the last vertex of Q'.

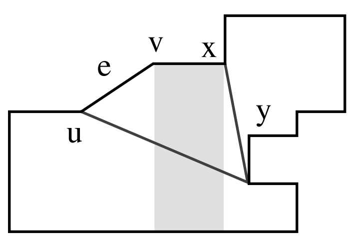

Figure 2.13: Finding Q', case 2.

It is now easy to verify that we can split $P - Q'$ into two or three 1-orthogonal polygons. ∎

2.4 Cutting orthogonal polygons into L-shaped pieces

A different proof of Theorem 2.2.1 was obtained by J. O'Rourke [102]. An L-shaped polygon is an orthogonal polygon with six vertices. Notice that any L-shaped polygon can be guarded by a single guard. O'Rourke's main idea is to divide an orthogonal polygon into L-shaped pieces, each of which can be guarded with one guard.

In view of Lemma 2.2.1 O'Rourke rephrases his result in terms of r, the number of reflex vertices of orthogonal polygons. He proves:

Theorem 2.4.1 $\lfloor \frac{r}{2} \rfloor + 1$ *guards are always sufficient and occasionally necessary to guard an orthogonal polygon with* r *reflex vertices.*

Proof: The proof of Theorem 2.4.1 is by induction on r. If $r = 0$ our result is clearly true. Two cases arise:

1. There is a horizontal or vertical line segment l connecting two reflex vertices of P such that the interior of l is totally contained in the interior of P.

2. No such horizontal or vertical line segment exists. In this case, we say that P is in general position.

Let l be a horizontal line segment as in (1). Then l splits P into two orthogonal polygons Q and R with s and t reflex vertices respectively, such that $s+t = r-2$. By induction, we can guard them with $\lfloor \frac{s}{2} \rfloor + 1$ and $\lfloor \frac{t}{2} \rfloor + 1$ guards respectively. But $(\lfloor \frac{s}{2} \rfloor + 1) + (\lfloor \frac{t}{2} \rfloor + 1) \leq \lfloor \frac{r}{2} \rfloor + 1$ and our result follows.

Some terminology and results will be neede to prove (2). A *cut* of an orthogonal polygon is an extension of one of the edges incident to a reflex vertex of P towards the interior of P until it hits the boundary of P. Notice that a cut through a reflex vertex u "resolves" or eliminates u since it is no longer a reflex vertex in either of the two polygons, say Q and R, into which P is partitioned by the cut. A cut is called an *odd cut* if at least one of Q and R has an odd number of reflex vertices.

To finish the proof of our result, we will use the following result, which will be proved after we finish our proof: *Every orthogonal polygon in general position has an odd-cut.*

Let P be an orthogonal polygon in general position, and consider an odd-cut that splits it into two subpolygons Q and R. Assume without loss of generality that Q has an odd number of reflex vertices. Let s and t be the number of reflex vertices of Q and R respectively. Then by induction, we can guard Q and R with $(\lfloor \frac{s}{2} \rfloor + 1)$ and $(\lfloor \frac{t}{2} \rfloor + 1)$ vertex guards respectively. Notice that $s + t = r - 1$, and since s is odd it follows by a simple case analysis that $(\lfloor \frac{s}{2} \rfloor + 1) + (\lfloor \frac{t}{2} \rfloor + 1) \leq \lfloor \frac{r}{2} \rfloor + 1$. Our result now follows. ■

To finish the proof of Theorem 2.4.1 we now prove that every orthogonal polygon in general position has an odd-cut.

We notice first that the general position condition is essential. The polygon in Figure 2.14(a) has no odd-cuts! Moreover a horizontal odd-cut does not always exist. The polygon in Figure 2.14(b) has no horizontal odd-cut. This complicates things a bit. We divide cuts into horizontal and vertical cuts, which we denote as *H-cuts* and *V-cuts* respectively.

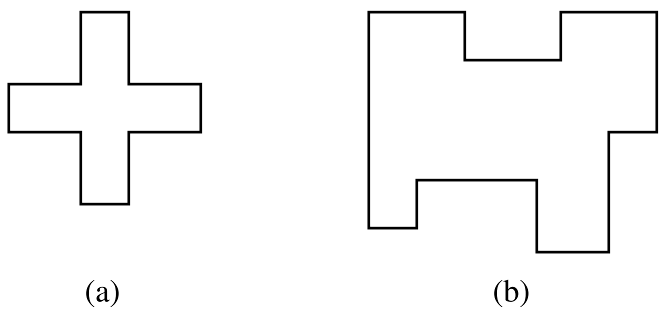

Figure 2.14: An orthogonal polygon with no odd-cuts, and a polygon with no horizontal cuts.

Notice first that if the number of reflex vertices of P is even, then the number of reflex vertices of Q plus those in R equals $r-1$, which is odd, and thus one of Q or R must have an odd number of reflex vertices. Suppose then that P has an odd number of reflex vertices.

Two reflex vertices of P are called a *horizontal pair* if they are the end vertices of a horizontal edge of P. A reflex vertex v of P is called *H-isolated* if the second end vertex of the horizontal edge incident at v is convex.

Obtain a partitioning of P by cutting P along an H-cut at each reflex vertex of P that belongs to a horizontal pair; see Figure 2.15. We call this the *H-partitioning* π of P. We now define a H-graph of this partitioning as the oriented graph whose vertex set is the set of regions of this partitioning. Two adjacent regions, X and Y are connected by an edge oriented from X

to Y if the H-cut that separates them is an extension of an edge of P on the boundary of X; see Figure 2.15.

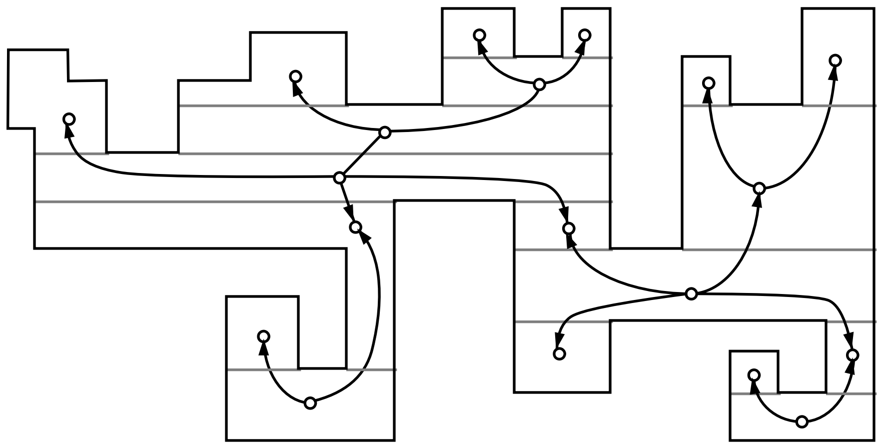

Figure 2.15: The H-partitioning of P and its H-graph.

We can now classify the vertices of the H-graph of P as follows:

Leaf nodes Vertices with in-degree 1.

Pit nodes Vertices with in-degree 2, and out-degree 0.

Source nodes Vertices with out-degree 2 or 4.

Branch nodes Vertices with out-degree 2 and in-degree 1.

Lemma 2.4.1 *If the H-graph of P contains a pit, then it has an odd-cut.*

Proof: Let K be a pit of P, and let h and h' be the top and lower H-cuts separating it from its neighbours. If K contains no H-isolated vertices, one of h or h' must be an odd-cut. If K has H-isolated vertices and h is an even cut, the highest H-isolated vertex u of K generates an odd-cut, and our result follows; see Figure 2.16. ∎

We now prove:

Lemma 2.4.2 *If P contains no odd-cuts, its H-graph contains a single source, P has exactly one H-isolated vertex, and it is located at the source of its H-graph.*

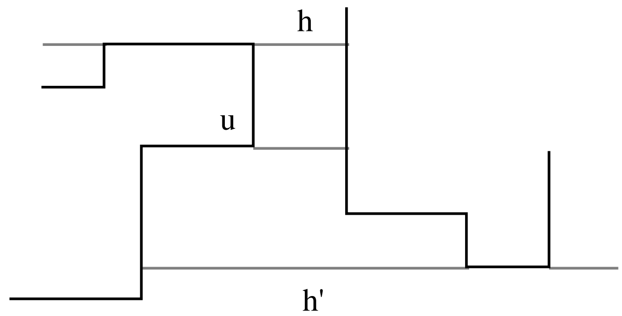

Figure 2.16: Finding an odd-cut in a pit.

Proof: Observe that if the H-graph of P contains two sources, then the path joining them must necessarily contain a pit. However since P admits no odd-cuts, its H-graph has no pits. It follows the H-graph of P contains a unique source. Since the H-graph of P is a tree with a unique source, each vertex of it, except for its source, has in-degree 1. Let K be one of these regions of the H-partitioning of P. Then since P admits no odd-cuts, the cut corresponding to the incoming edge of K is an even-cut. Suppose without loss of generality that this edge is a bottom edge of K. Then if K contains H-isolated vertices, the lowest of them generates an odd-cut, which is a contradiction.

However since P contains an odd number of reflex vertices (or any H-cut would be an odd-cut) it must contain at least one H-isolated vertex, and it must be a vertex of the source region of the H-partitioning of P. Moreover, notice that if the source region contains more than one H-isolated vertex, one of them would generate an odd-cut, which is a contradiction. Our result follows. ∎

We notice that all the results proved for H-cuts hold for V-cuts, and in particular we have that if P admits no vertical odd-cuts, it has exactly one V-isolated vertex located at the source face of the V-graph of P.

We now prove:

Lemma 2.4.3 *Any orthogonal polygon P in general position with an odd number r of vertices admits an odd-cut, $3 \leq r$.*

Proof: Since P admits no odd-cuts, it has exactly one H-isolated vertex u and one V-isolated vertex v. Moreover these vertices are in the sources of the H-graph and the V-graph of P. Any other reflex vertex of P belongs to an H- and a V-pair. Therefore all the reflex vertices of P are located in a chain of reflex vertices that starts, say, in u and ends in v; see Figure 2.17. This implies that both u and v are located in leaf regions of the H- and V-partitions of P, which is a contradiction. ∎

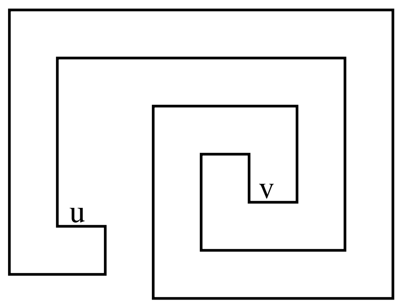

Figure 2.17: A spiral orthogonal polygon.

This concludes our proof of Theorem 2.4.1. ∎

2.5 Algorithms for L-shaped partitions

In this section we develop an $O(n \ln n)$ algorithm to partition an orthogonal polygon into L-shaped pieces. Using Chazelle's linear time triangulation algorithm, a linear time implementation of our algorithm can be obtained, however this is outside the scope of our book.

An orthogonal polygon Q is called a vertical *histogram* if it has a horizontal edge, called the *base* of Q, such that every point of Q is visible from some point in its base; see Figure 2.18.

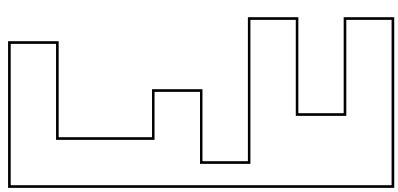

Figure 2.18: A histogram.

Lemma 2.5.1 *If an orthogonal polygon has no horizontal odd-cuts, it can be partitioned into* $\lfloor \frac{r}{2} \rfloor + 1$ *L-shaped pieces.*

Proof : Observe that by Lemma 2.4.2 the H-graph of any such polygon has exactly one source, and one H-isolated vertex located in this face. It is easy to see that this cut partitions P into two histograms. Our result now follows by partitioning each of these histograms by cutting them along a vertical line through every second reflex vertex. ∎

We now show how to obtain a partitioning of P into L-shaped pieces. To start, assume that P has no horizontal cuts that join two reflex vertices. This case can be easily dealt with. Furthermore we will assume that P contains an odd number of reflex vertices. The case when P has an even number of reflex vertices will be dealt with at the end of our proof for the odd case. Our approach is justified by the following observation: an odd-cut of a polygon with an odd number of reflex vertices splits it into two subpolygons, each with an odd number of reflex vertices. Moreover we observe that if we cut P along an odd-cut h, and h' is an odd-cut of P different from h, then h' is also an odd-cut of the subpolygon of P that contains it.

We now present our algorithm to find an L-partitioning of a polygon P with an odd number of reflex vertices:

Algorithm 2.5.1

1. *Find all the horizontal cuts of* P.
2. *Find all the horizontal odd-cuts.*
3. *Cut* P *at each horizontal odd-cut.*
4. *Each resulting piece has no horizontal odd-cut, i.e it is a histogram. Partition it by cutting it along a vertical through every second reflex vertex.*

We notice now that each subpolygon P_i of P obtained in Step 3 of our algorithm has no odd-cuts, and thus by Lemma 2.5.1 can be partitioned into at most $\lfloor \frac{r_i}{2} \rfloor - 1$ L-shaped pieces, where r_i is the number of reflex vertices of P_i. It now follows that our algorithm indeed produces a partitioning of P as desired.

Finding the odd-cuts of P can be done in $O(n \ln n)$ time by a horizontal line sweep. See Appendix.

Determining which cuts are odd-cuts can be accomplished as follows: Pick any reflex vertex of P and label it 1. Set a reflex counter of vertices of P to 1. We now *walk* along the boundary of P. Each time we encounter a reflex vertex of P we increase our counter by one, and label the vertex with the new value of our counter. Each time we encounter the end point of a horizontal cut that is not a vertex of P we simply label it with the value of our counter. A horizontal cut is even if the absolute value of the difference of the labels of its end points is odd. This can clearly be done in linear time.

Step 3 can also be done in linear time. It is obvious that this partition yields at most $\lfloor \frac{r}{2} \rfloor + 1$ L-shaped polygons. Thus we have proved:

Theorem 2.5.1 *An orthogonal polygon with an odd number of reflex vertices can be partitioned into at most $\lfloor \frac{r}{2} \rfloor + 1$ L-shaped polygons in $O(n \ln n)$ time.*

If P has an even number r of reflex vertices, we first create an extra reflex vertex of P as follows: let e be the lowest horizontal edge of P and u its leftmost vertex. Since we are assuming that P is in general position, all its vertices have different x- and y-coordinares. Let ϵ be such that the horizontal and vertical distance between any two vertices of P is at least ϵ. Cut away from P a square with sides parallel to the coordinate axes, such that the length of its edges is ϵ and u is its lowest leftmost vertex. Let P' be the resulting orthogonal polygon. P' has exactly $r+1$ reflex vertices. The following lemma, left as an exercise to the reader, can now be proved:

Lemma 2.5.2 *Any L-shaped partitioning of P' into at most $\lfloor \frac{r+1}{2} \rfloor + 1$ pieces can be modified to obtain a similar partitioning of P into $\lfloor \frac{r}{2} \rfloor + 1$ pieces.*

However since r is even, $\lfloor \frac{r+1}{2} \rfloor + 1 = \lfloor \frac{r}{2} \rfloor + 1$ and we have proved:

Theorem 2.5.2 *An orthogonal polygon with r reflex vertices can be partitioned into at most $\lfloor \frac{r}{2} \rfloor + 1$ L-shaped polygons in $O(n \ln n)$ time.*

2.6 Algorithms

Using Fisk's proof, together with Chazelle's linear time triangulation algorithm, we can develop a linear time algorithm to find $\lfloor \frac{n}{3} \rfloor$ guards to protect a polygon as follows:

Algorithm 2.6.1

1. *Triangulate p.*
2. *Three color the vertices of $GT(P)$.*
3. *Place the guards at the smallest chromatic class.*

To show that Algorithm 2.6.1 runs in linear time, all we have to show is that Step 2 of Algorithm 2.6.1 can be implemented in linear time. To show this, we notice that by Excercise 3 of Section 1.3, $GT(P)$ always has a vertex v of degree 2. We now notice that any 3-coloring of $GT(P) - v$ can be extended to a 3-coloring of $GT(P)$. It is now easy to see that this procedure leads to a linear time algorithm.

2.6.1 Input sensitive algorithms

Although Theorem 2.1.1 provides a general upper bound on the number of guards required to guard any polygon with n vertices, most polygons can be guarded with fewer than $\lfloor \frac{n}{3} \rfloor$ guards. It is natural then to ask for the existence of an efficient algorithm to find the minimun number of vertex guards to protect a polygon.

Unfortunately, the existence of such an algorithm is highly unlikely, due to the following result of D. T. Lee and A.K.Lin [86]:

Theorem 2.6.1 *The minimum vertex guard problem for polygons is NP-complete.*

2.7 Exercises

Ex. 2.1 Provide all the details of a linear time implementation of Algorithm 2.6.1.

Ex. 2.2 Prove Lemma 2.1. Hint: use the fact that the dual graph of a triangulation of a polygon is a tree in which the maximum degree of each vertex is 3.

Ex. 2.3 Finish Chvátal's proof of Theorem 2.1.1.

Ex. 2.4 Verify that $P - Q'$ can indeed be split into two or three 1-orthogonal polygons; i.e. verify that the resulting polygons have an even number of edges and their noses are empty.

Ex. 2.5 Prove Lemma 2.5.2.

Chapter 3

Guarding traditional art galleries

In the classical Art Gallery Theorem, an art gallery is a simple polygon on the plane. In a more realistic setting, a *traditional art gallery* is housed in a rectangular building subdivided into rectangular rooms. Assume that any two adjacent rooms have a door connecting them. (See Figure 3.1.)

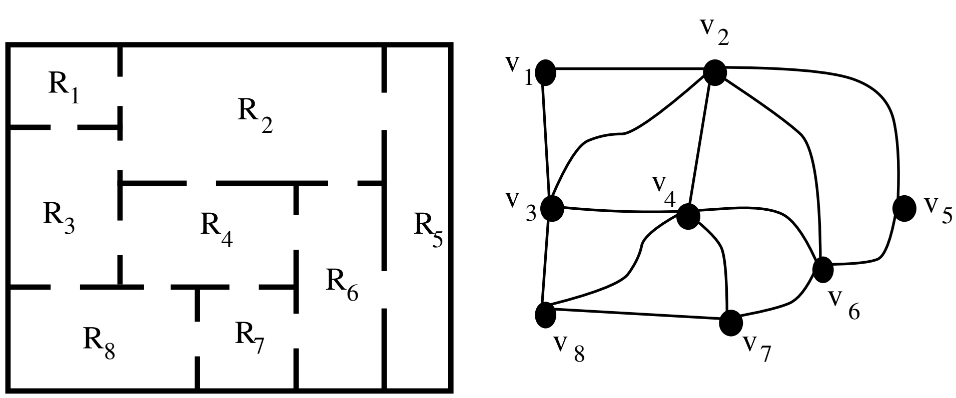

Figure 3.1: Traditional art gallery and its dual graph.

How many guards need to be stationed in the gallery so as to guard all its rooms? Notice that a guard who is stationed at a door connecting two rooms will be able to guard both rooms at once, and since no guard can guard three rooms, it follows that if the art gallery has n rooms we need at least $\lceil \frac{n}{2} \rceil$ guards. One of our objectives in this chapter is to prove:

Theorem 3.0.1 *[41] Any rectangular art gallery with n rooms can be guarded with exactly $\lceil \frac{n}{2} \rceil$ guards.*

3.1 A taste of matching theory

Some basic results in matching theory have been widely used in several results on illumination problems. To help the reader better understand the results about to be proved, a brief revisiting of these results is in order.

Given a graph G, a *matching* M of G is a subset of edges of G such that no two edges of M have a common vertex. A matching M of G is called *perfect* if each vertex of G is incident to an edge in M; see Figure 3.2.

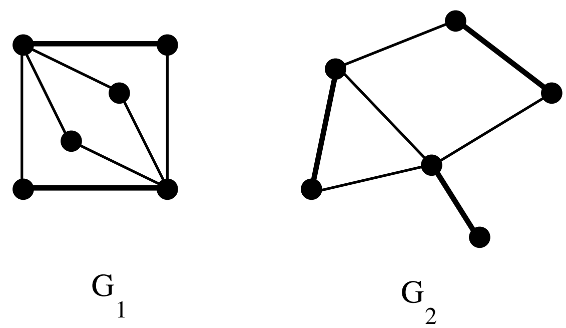

Figure 3.2: Graph G_1 has a matching with two edges, but no perfect matching. G_2 has a perfect matching.

Given a graph G and a subset S of the vertices of G, we define $Odd(S)$ to be the number of components of $G - S$ with an odd number of vertices. We now prove the following result due to W. T. Tutte which provides necessary and sufficient conditions for the existence of perfect matchings:

Theorem 3.1.1 *(Tutte [128]): A graph G has a perfect matching iff for every subset S of $V(G)$, $Odd(G - S) \leq |S|$.*

We first observe that if G has a subset of vertices S such that $Odd(G - S) > |S|$, then G cannot have a perfect matching. To see this we simply observe that in any perfect matching of G at least one element of an odd component of $G - S$ must be matched to a vertex in S. Thus if we have more odd components in $G - S$ than elements in S, a perfect matching cannot exist; see Figure 3.3.

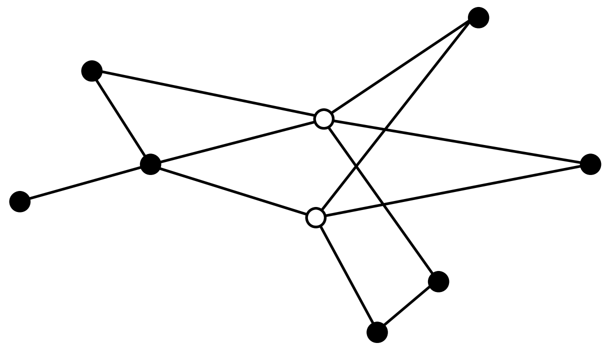

Figure 3.3: A graph with no perfect matching. The vertices of S are represented by small empty disks.

A graph G is called *saturated* if it has no perfect matching, but adding any other edge to G produces a graph that has a perfect matching.

Let S be the set of vertices of a saturated graph G adjacent to *all* the vertices of G, and let $T = V(G) - S$. We now prove:

Lemma 3.1.1 *Let u, v, w be elements of T such that u and w are adjacent to v. Then u and w are also adjacent.*

Proof: Assume that u and w are not adjacent. Since v is in T, there is a vertex x of G not adjacent to v. Let G' and G'' be the graphs obtained by adding the edges $u - w$ and $v - x$ to G respectively. Since G is saturated, G' and G'' have perfect matchings. Let M' and M'' be these matchings respectively. Since G has no perfect matchings, $u - w$ and $v - x$ are in M' and M'' respectively.

In the graph obtained by joining M' and M'' each vertex has degree two, and this graph consists of a set of even cycles with edges alternating between M' and M''. Two cases arise:

1. $u - w$ and $v - x$ belong to different cycles of $M' \cup M''$.
2. $u - w$ and $v - x$ belong to the same cycle.

Suppose then that $u - v$ and $v - x$ belong to different cycles in $M' \cup M''$. Let C be the cycle containing $u - w$. Then by deleting from M' all the edges of C in M' and replacing them by those edges of C in M'', we obtain a perfect matching of G' that does not contain $v - w$, which is a contradiction.

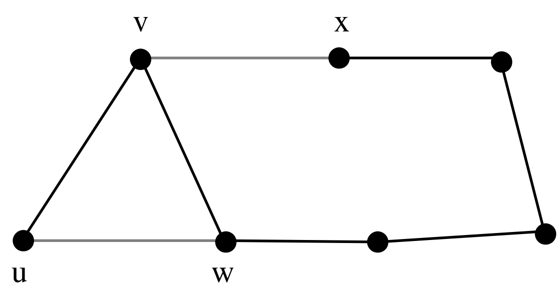

Figure 3.4:

If both edges belong to the same cycle C of $M \cup M''$, relabel the vertices of C $v_0 = v, v_1 = x, \ldots,$ according to the order in which they appear in C, and suppose without loss of generality that w is relabelled with a smaller label than v. Then since the edges of C alternate between M' and M'', $w = v_{2k+1}$ for some $k \geq 0$. However since v and w are adjacent, the cycle with vertices $v_0 = v, v_1 = x, \ldots, v_{2k+1}, v_0 = v$ is an even cycle C' such that every second edge of it belongs to M'; see Figure 3.4. By removingthe edges of C' in M' and replacing them by the edges of C' not in M' we obtain again a matching of G' that does not contain $v - x$, a contradiction. Therefore u and w are adjacent in G. ∎

We are now ready to prove Tutte's Theorem:

Proof: As mentioned before, if $Odd(G-S) > |S|$ then G cannot have a perfect matching. Suppose then that G has no perfect matching, and that G has an even number of vertices, otherwise if G has an odd number of vertices, take S to be the empty set, and our result follows.

We now show that G must contain a subset S of vertices such that $Odd(G - S) > |S|$. Let us add to G as many edges as possible to obtain a saturated graph G'. Since G is not a complete graph, G' is not complete either. As before, let S be the set of vertices of G' adjacent to all the vertices of G' and H_0 be the complete subgraph of G' induced by S, and $H_1, \ldots, H_k$ be the components of $G' - S$. By Lemma 3.1.1 $H_1, \ldots, H_k$ are also complete subgraphs of G'.

If at most S components of $G' - S$ are odd, then a perfect matching of G' can be easily found contradicting that G' is saturated. Thus $G' - S$ has at least $S + 1$ odd components, in fact by a parity argument, it must have at least $S + 2$ odd components. If $G' - S$ contains more than $S + 2$ odd components, we can check that we can add to G' an extra edge connecting two odd components and the resulting graph still has no perfect matching, a contradiction.

Similarly, if $G' - S$ contains an even component, then any edge connecting it to another component of $G' - S$ can be added without creating a perfect matching, which contradicts that G' is saturated.

To finish our proof, we delete from G' the edges we added to G, and observe that in doing so, each odd component of $G' - S$ generates at least an odd component of $G - S$. Thus $Odd(G - S) > |S|$ and our result follows. ∎

3.2 The proof of Theorem 3.0.1

We now prove Theorem 3.0.1 using Tutte's Theorem.

Proof: Given a rectangular art gallery T with n rooms $R_1, \ldots, R_n$, we can associate to it a dual graph $G(T)$ by representing each room R_i of T by a vertex v_i in $G(T)$, two vertices being adjacent if their corresponding rectangles share a line segment in their common boundary. See Figure 3.1.

Notice that the boundary of the union of the rectangles corresponding to the vertices of any connected subgraph of $G(T)$ forms an orthogonal polygon, possibly with some orthogonal holes.

We now show that if $G(T)$ has an even number of vertices, $G(T)$ has a perfect matching M. This suffices to prove our result since our $\lceil \frac{n}{2} \rceil$ points can now be chosen using M as follows: for every edge $\{v_i, v_j\}$ of $G(T)$ in M station a guard at the door connecting R_i and R_j. Clearly these guards will cover all the subrectangles of T. The case when $G(T)$ has an odd number of vertices follows by subdividing any room of T into two.

Suppose then that $G(T)$ has an even number of vertices. We now show that $G(T)$ satisfies Tutte's Theorem. Let S be any subset of vertices of $G(T)$, and k be the number of connected components of $G(T) - S$. Each component C_i of $G(T) - S$ is represented by an orthogonal subpolygon P_i of T. Each such polygon has at least four corner points and thus the total number of corner points generated by the k components in $G(T) - S$ is at least $4k$; see Figure 3.5.

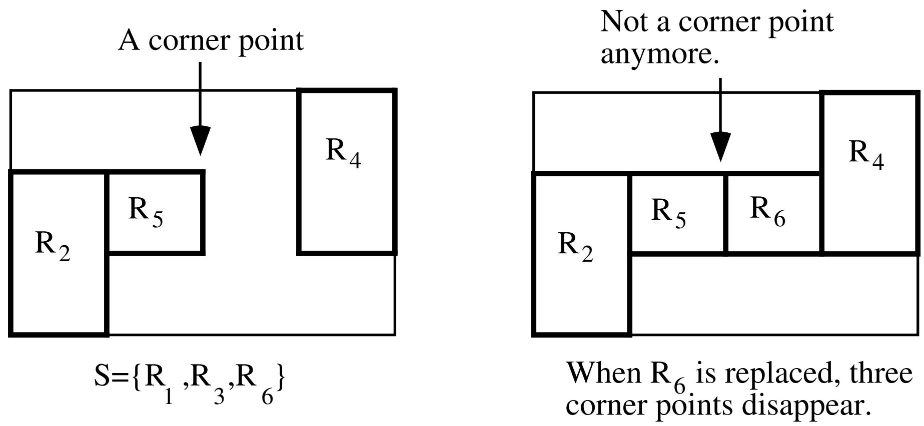

Figure 3.5: The components of $G - S$ are themselves orthogonal polygons.

The next observation is essential to our proof: When a rectangle represented by a point in S is now replaced, at most four corner points will disappear.

Once all rectangles in S are replaced, all the corner points generated by the components of $G(T) - S$ will disappear, except for the four corner points of T. It follows that $k \leq\mid S \mid +1$. The reader may verify that if $k =\mid S \mid +1$, then at least one of the components of $G(T) - S$ is even. Our result now follows. ■

A different proof of Theorem 3.0.1 can be obtained by using another well known result on planar graphs. It is common practice in most museums to indicate to visitors, using arrow signs, a sequence in which the rooms of the museum may be visited without either missing or repeating a room. Our second proof of Theorem 3.0.1 shows that in fact, any traditional Art Gallery always has such a path. A graph G is called *hamiltonean* if there is a cycle that contains all the vertices of G. A graph has a *hamiltonean path* if there are two vertices u and v and a path starting at u and ending in v that contains all the vertices of the graph. We now prove:

Theorem 3.2.1 *Every rectangular art gallery has a hamiltonean path. Moreover, this path starts and ends at external rooms.*

To prove our result, we will use the following result due again to Tutte [129]:

Theorem 3.2.2 *Every planar 4-connected planar graph is hamiltonean.*

A graph is called k-connected if we need to delete at least k vertices of it to obtain a disconnected graph. We now prove:

Theorem 3.2.3 *Any rectangular art gallery with n rooms has a hamiltonean path.*

To prove our result, we will show that $G(T)$ has a hamiltonean path. We notice first that $G(T)$ is planar, but not necessarily 4-connected. For example the graph arising from Figure 3.1 is 3-connected. To sove this problem we first *frame* our rectangular art gallery T using four rectangles as shown in Figure 3.6.

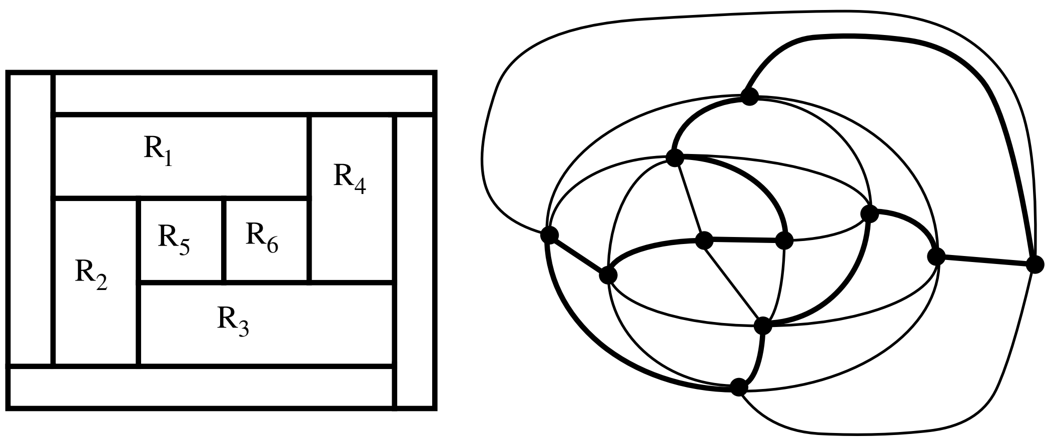

Figure 3.6: Framing a rectangular art gallery.

We now notice that the dual $G(T^*)$ of the resulting rectangular art gallery T^* in which we add a vertex to the unbounded external face is indeed 4-connected. Thus by Theorem 3.2.3 it is hamiltonean. This however is not enough, as the hamiltonean path of $G(T^*)$ guaranteed by Theorem 3.2.3 may induce two paths in $G(T)$. Referring again to Figure 3.6 the hamiltonean cycle of $G(T^*)$ shown in thick lines first visits R_2, R_5, R_6 and R_1, exits T and reenters it to visit R_4 and R_3, thus failing to produce a hamiltonean path for T. This can however be fixed as follows: Take three copies of T, and frame them as in Figure 3.7. The dual graph $G(T^{**})$ of the resulting configuration is again hamiltonean. Hovever, it is now easy to verify that the hamiltonean cycle of $G(T^{**})$ induces a hamiltonean path in at least one of the three copies of T used in our construction. That this path begins and ends at external rooms of T is obvious. Our result follows. ∎

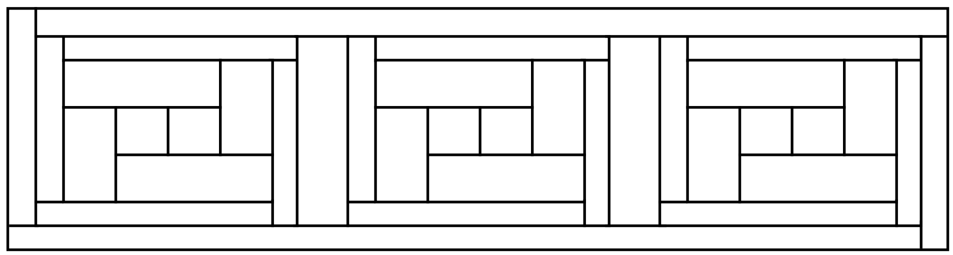

Figure 3.7: Framing three copies of a rectangular art gallery.

To show that Theorem 3.2.1 implies Theorem 3.0.1, we simply follow our hamiltonean path for T, and station a guard at the door connecting the first room we visit to the second, the door connecting the third to the fourth, etc. our result follows.

3.2.1 Algorithms

In this section we show that the placement of the $\lfloor \frac{n}{2} \rfloor$ guards to protect a rectangular art gallery can be done in linear time. Our result follows fron the following result due to N. Chiba and T. Nishizeki [27]

Theorem 3.2.4 *Finding hamiltonean cycles in* 4*-connected planar graphs can be done in linear time.*

By using Chiva and Nishizeki's algorithm on $G(T^{**})$, which is 4-connected, we have:

Theorem 3.2.5 *Guarding a rectangular art gallery with* $\lceil \frac{n}{2} \rceil$ *can be done in linear time.*

3.3 Non-rectangular Art Galleries

In the previous section, we studied the problem of guarding art galleries housed in rectangular buildings. In this section we study the problem of guarding art galleries housed in convex buildings that are subdivided into

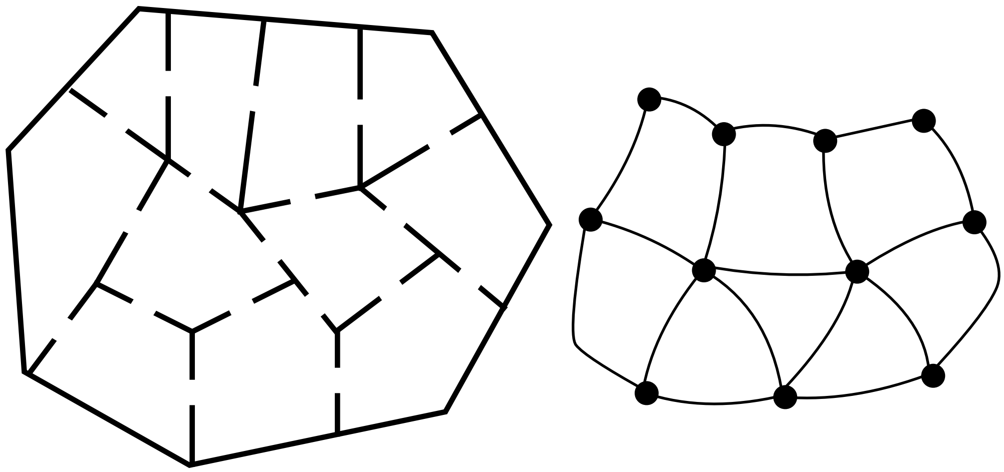

Figure 3.8: A convex art gallery subdivided into convex rooms.

convex rooms. As before, we will assume that whenever two rooms share a wall, there is a door connecting them. (See Figure 3.8.)

Our main result in this subsection is the following:

Theorem 3.3.1 *Any convex art gallery T with n convex rooms can be guarded with at most $\lfloor \frac{2n}{3} \rfloor$ guards.*

To prove our result, we use the following result due to Nishizeki [96] (see also [97]).

Theorem 3.3.2 *Let G be a 2-connected planar graph such that the degree of every vertex in G is at least 3. Then for all $n \geq 14$, G has a matching of size at least $\lceil \frac{n+4}{3} \rceil$. For $n < 14$, G has a matching of size $\lceil \frac{n}{2} \rceil$.*

We will not prove Nishizeki's theorem here. The interested reader may find the proof of this result in [97].

We proceed now to prove Theorem 3.3.1.

Proof: Let $G(T)$ be the dual graph of T in which the vertices of $G(T)$ are the rooms of T, two of which are adjacent if they share a common wall. Suppose that $G(T)$ contains no cut vertex v with degree 2. This case can be handled easily by an inductive argument. It may be the case that $G(T)$ has vertices of degree 2. Since our rooms are convex, any internal room, i.e. a room contained in the interior of T has degree greater than or equal to 3. Thus vertices of degree 2 in $G(T)$ correspond to *exposed* rooms of T, i.e. rooms having a sector of thier boundaries on the boundary of T. It now follows that if we add an artificial vertex to $G(T)$ adjacent to all the exposed roms of T to obtain a 2-connected graph such that all its vertices have degree at least 3.

We can now apply Nishizeki's Theorem to our augmented graph. This produces a matching M of size at least $\lceil \frac{n+5}{3} \rceil$. Each edge of M corresponds to two adjacent rooms that can be guarded with a single guard. This takes care of $2\lceil \frac{n+5}{3} \rceil$ rooms of T We must place a guard for each of the remaining rooms. Thus we can guard T with at most

$$\lceil \frac{n+5}{3} \rceil + ((n+1) - 2\lceil \frac{n+5}{3} \rceil) = \lfloor \frac{2n}{3} \rfloor$$

guards. An example in which $\lfloor \frac{2n}{3} \rfloor$ guards are necessary is shown in Figure 3.9. ∎

A point p covers or guards a convex set S_i if p is in the interior or the boundary of S_i. Let P be a partitioning of the plane into n convex sets. A point set G is called a *cover* of P if every element of P is covered by at least a point in G. We observe that using the same technique as that one used in the proof of our previous result we can prove:

Theorem 3.3.3 *Any partitioning of the palne into n convex sets has a covering with at most $\lfloor \frac{2n}{3} \rfloor$ points. The bound is tight.*

The proof of this result is left as an excercise.

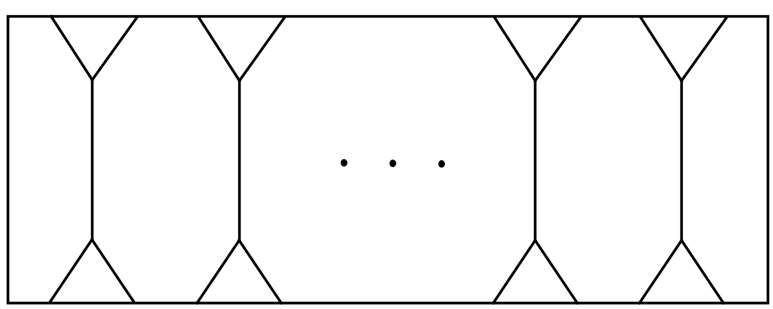

Figure 3.9: A convex art gallery with $3m + 1$ rooms that needs $2m$ guards.

3.4 Excercises

Ex. 3.1 Prove that if an orthogonal polygon with k reflex vertices is subdivided into n rectangular rooms as in Theorem 3.0.1 it can be guarded with at most $\lfloor \frac{n+k}{2} \rfloor$ guards. Show that this bound is tight.

Ex. 3.2 Show that Theorem 3.0.1 extends to decompositions of boxes in R^3 into boxes, such that no two of these boxes meet at a single point.

Ex. 3.3 Suppose that we choose to station a set of guards in a rectangular art gallery according to the following greedy algorithm:

> While not all the rectangles of T are guarded, choose, if possible, a pair of adjacent rectangles, guard them with a guard, and delete them them from $G(T)$. Once this can no longer be done, the collection of subrectangles left form an independent set in the dual graph $G(T)$. Guard each of these rectangles with a guard.

Prove that using this algorithm will produce a guarding set with at most $\lfloor \frac{3n}{4} \rfloor$ guards. Show that in some cases, our algorithm will actually choose $\lfloor \frac{3n}{4} \rfloor$ guards.

Ex. 3.4 Prove Theorem 3.3.3.

Chapter 4

Families of Line Segments

Let $F = \{S_1, \ldots, S_n\}$ be a collection of n line segments on the plane. We say that a set of points $Q = \{p_1, \ldots, p_k\}$ protects F if every element S_i of F contains a point visible from some point p_j in Q.

A natural motivation for this concept is that a guard is protecting an art work S if it can see at least a point of it. In this case, a robber cannot remove S without our guard noticing it! In Figure 4.1 we show a family of 7 line segments that need two points to protect them. Notice that the smaller line segments cannot be protected by the same guard.

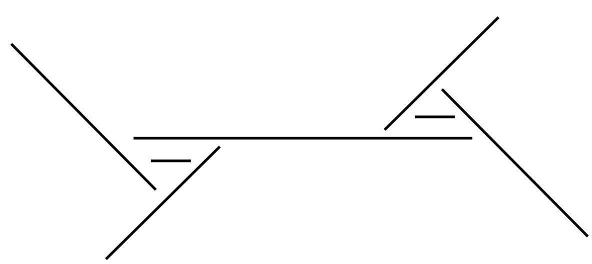

Figure 4.1: Two guards suffice.

Our main objective in this section is to prove the following result.

Theorem 4.0.1 *Any collection of n line segments can always be protected using at most $\lceil \frac{n}{2} \rceil$ points; $\lfloor \frac{2n-3}{5} \rfloor$ points are occasionally necessary.*

Some preliminary results and definitions will be needed before we can prove our result.

Consider a collection $F = \{S_1, \ldots, S_n\}$ of n disjoint line segments on the plane. Construct a graph $G(F)$ with n vertices $v_1, \ldots, v_n$ such that v_i is adjacent to v_j if and only if there is a point x on the plane that sees at least a point in the boundary of each of $S_i$and S_j, i.e. x can protect both line segments. In Figure 4.2 a collection F of five segments and its corresponding graph $G(F)$ is shown.

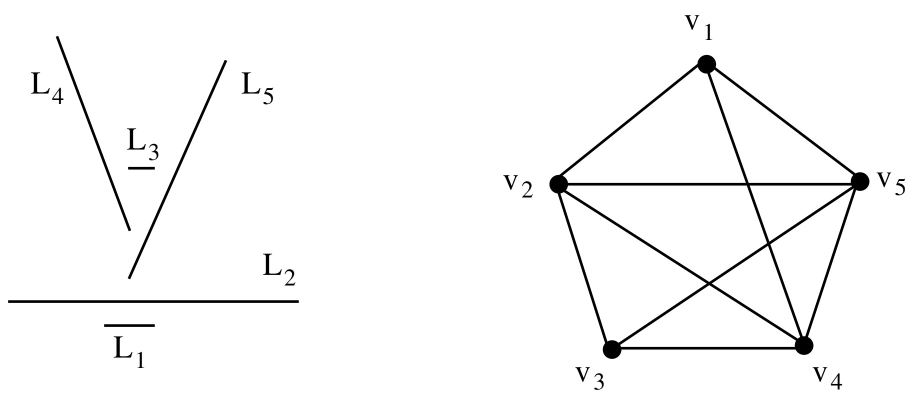

Figure 4.2: A family F of five line segments, and its corresponding graph $G(F)$.

The main idea in our proof of Theorem 4.0.1 is to show that for any family F with an even number n of disjoint segments, the graph $G(F)$ has a perfect matching.

The following lemma, given without proof, will be used to prove our main result.

Lemma 4.0.1 *Let Q be any convex polygon, $F = \{S_1, \ldots, S_n\}$ a family of n disjoint segments and let H be the subset of elements of F which intersect Q. The subgraph of $G(F)$ induced by the vertices of $G(F)$ representing elements in H is connected.*

We proceed now to prove Theorem 4.0.1.

Proof: Let $F = \{L_1, \ldots, L_n\}$ be a collection of n disjoint segments and let $G(F)$ be its associated graph. Assume that n is even, otherwise add another segment to F. We show that $G(F)$ satisfies Tutte's condition and thus has a perfect matching .

Consider any subset H of F and let S be the set of vertices of $G(F)$ representing the elements of H. We show that the number of connected components of $G(F) - S$ is at most $\mid S \mid = \mid H \mid$.

To start, delete from the plane all the segments *not* in H. One at a time, extend the elements of H until they meet another element of F, meet a previously extended element of F or become lines or semilines. Let π be the plane partition induced by the extended elements of H. It is easy to verify that π contains exactly $\mid H \mid +1$ polygonal faces. Replace the elements of F not in H; see Figure 4.3.

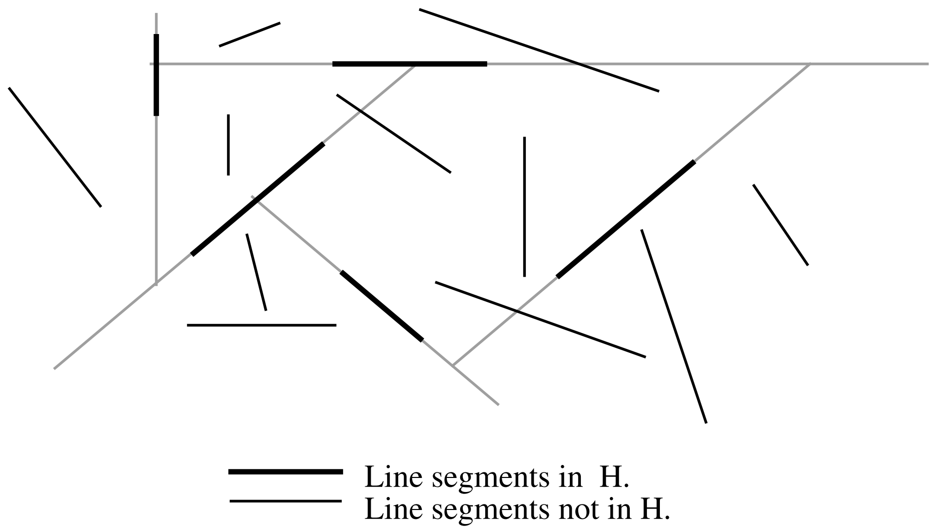

Figure 4.3: The partitioning generated by the elements in H.

By Lemma 4.0.1, the number of components of $G(F) - S$ is at most the number of faces of π, which is $\mid H \mid +1$. It can be easily verified that there are at least two adjacent faces in π such that the segments that intersect them are in the same component in $G(F) - S$. Thus the number of components in $G(F) - S$ is at most $\mid S \mid$, implying that $G(F)$ has a perfect matching M. It follows that F can be protected by a set, consisting of at most $\frac{n}{2}$ points, one for each edge of M.

To get the lower bound, construct an example of a family F with n segments in which $\lfloor \frac{2n-3}{5} \rfloor$ points are required to guard F, as follows.

Let H be a cubic plane graph with triangular outer face in which all the vertices, except the outer ones, are such that the three vectors emanating from the vertex along the edges positively span the plane. Let H have k vertices; it has $\frac{3k}{2}$ edges.

Substitute the edges of H by segments such that at each of the $k-3$ inner vertices we obtain a triangular face in which we insert a small segment; see Figure 4.4.

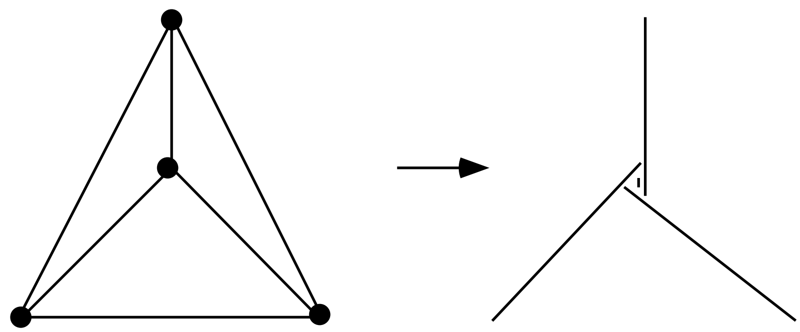

Figure 4.4: Generating a set of n line segments that require $\lfloor\frac{2n-3}{5}\rfloor$ points to protect them.

Discard the three edges of the outer face of H and disconnect each edge in a small neighborhood of its end vertices to form a collection of $n = (\frac{3k}{2}) - 3 + k - 3$ segments. No two of our $k-3$ small segments are visible from a single point, hence $k-3$ points are needed to guard the collection of segments. It is easy to verify that $k-3$ points are also sufficient. $k-3 = \lfloor\frac{2n-3}{5}\rfloor$, completing the proof of Theorem 4.0.1. ■

We close this section by noticing that the proof of Theorem 4.0.1 is also valid if we allow some of our line segments to be semilines, or even lines. Moreover in this case, the $\lceil\frac{n}{2}\rceil$ bound is tight, the set of semilines, and line shown in Figure 4.5 achieves this bound.

4.1 More on families of line segments

An attractive variation to guarding problems involving families of line segments was studied by J. O'Rourke [104].

Suppose we have a set of disjoint line segments $F = \{L_1, \ldots, L_n\}$ representing walls built anywhere on the plane, and we want to guard the whole plane. How many guards are always sufficient to achieve our goal?

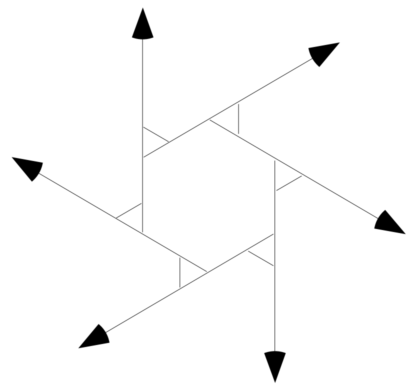

Figure 4.5: A family containing m semilines and m line segments that requires $\lceil \frac{n}{2} \rceil$ guards to protect it.

We must first clarify the conditions set by O'Rourke in his result. A guard at a point p sees a point q on the plane if the line segment joining p to q does not *cross* any element of F. For this purpose, it is better to consider the elements of F as open line segments. Moreover, O'Rourke allows for the line segment joining p to q to be collinear with any element of F, and also to locate guards on the line segments themselves.

Under these restrictions, O'Rourke proved:

Theorem 4.1.1 $\lfloor \frac{2n}{3} \rfloor$ *guards are always sufficient and occasionally necessary to guard the plane in the presence of* n *line segment obstacles.*

We will assume that our line segments are in *general position*, i.e. we assume that no two of them are collinear, and that the lines generated by the elements of F are such that no three meet at a point. These conditions can easily be taken care of.

We can now prove Theorem 4.1.1.

Proof: Consider a family $F = \{L_1, \ldots, L_n\}$ of n line segments in general position. As we did in the proof of Theorem 4.0.1, one by one extend each element of F until it hits another element of F, or an extension of it, or becomes a semiline or a whole line.

This creates a partitioning of the plane into exactly $n + 1$ convex polygons. Consider the dual of the obtained partitioning. If $n \geq 3$ this graph is 2-connected, however it may have vertices of degree 2; see Figure 4.6. It is easy to see that all vertices of degree 2 lie on infinite regions of our partitioning, and thus by adding an *artificial* vertex adjacent to all these vertices, we obtain a planar graph that satisfies Nishizeki's Theorem, and thus has a matching of size at least $\lceil\frac{((n+1)+1)+4}{3}\rceil = \lceil\frac{n+6}{2}\rceil$.

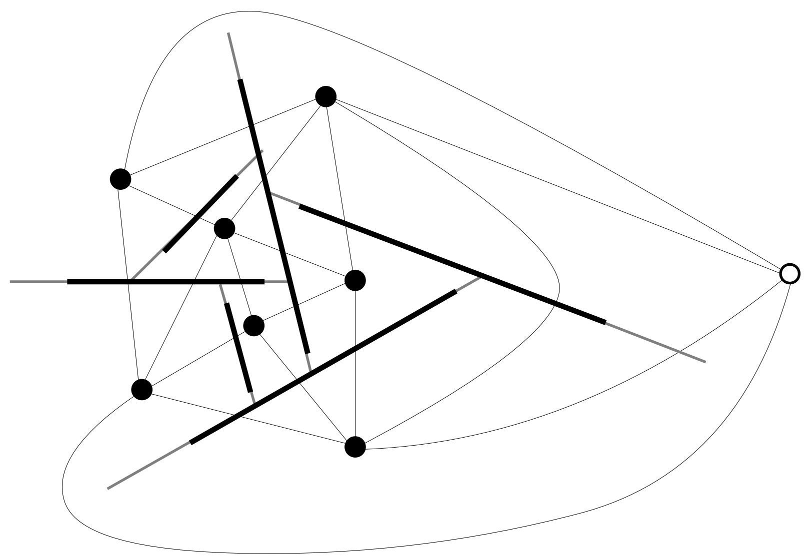

Figure 4.6: A convex partitioning of the plane generated by a set of line segments.

Notice that since two adjacent regions can be guarded with a single point, located *possibly on a line segment* of F, using $\lceil\frac{n+6}{2}\rceil$ points we can guard at least $2\lceil\frac{n+6}{2}\rceil - 1$ regions. The last -1 in the equation is to take into consideration the case when our artificial vertex is the endvertex of an edge in our matching. The remaining ones can be guarded with a point each. This gives a total of:

$$(n+1) - 2\lceil\frac{n+6}{2}\rceil - 1 + \lceil\frac{n+6}{2}\rceil = \lfloor\frac{2n}{3}\rfloor$$

guards. We now show that there are families of line segments for which $\lfloor\frac{2n}{3}\rfloor$ points are required. We start again with a cubic graph G with m vertices drawn on the plane such that its external face is a triangle, and for all internal vertices, the edges incident to them span the plane. We again discard the edges in the outer face and modify the remaining edges of G in such a way that at each of the internal vertices we obtain a *triangular* face. Our configuration has exactly three line segments, say a, b and c that can be extended to infinity.

We now shorten each of the resulting line segments by a sufficiently small ammount, and add three more lines l_a, l_b and l_c close to the end points of a, b, and c respectively as shown in Figure 4.7.

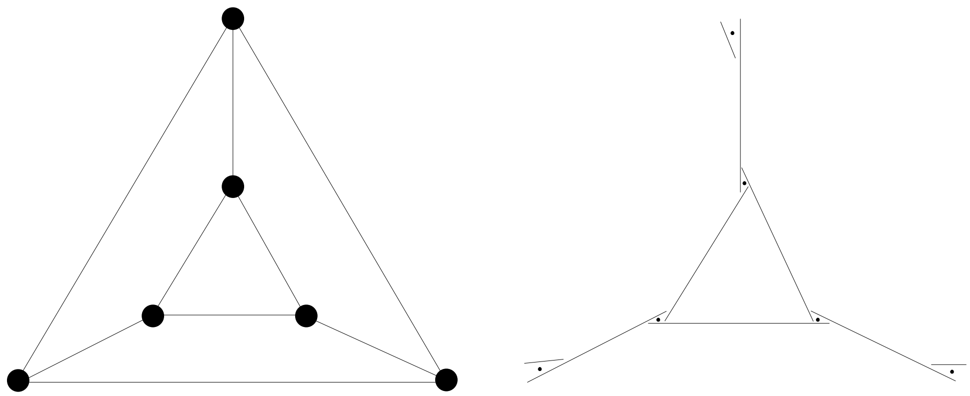

Figure 4.7: A family of n line segments for which $\lfloor \frac{2n}{3} \rfloor$ guards are necessary.

Since G has exactly $\frac{3m}{2}$ edges, our construction yields a total of $\frac{3m}{2}$ line segments. If we place a point in the middle of each triangle in our previous construction, plus a point between each of a and l_a, b and l_b, and c and l_c, we obtain a set of m points such that no two of them can be guarded by a single guard. Our result now follows. ∎

An interesting open problem arises:

Problem 4.1.1 *Verify if the bounds of Theorem 4.1.1 remain valid if we don't allow guards to be placed on line segments.*

At this point, it is worth noticing that the proof of Theorem 4.1.1 extends to families containing line segments, semilines, and lines. However if we do not allow guards to be located on our our obstacles, the bound proved in Theorem 4.1.1 is no longer valid. If our obstacls are n parallel lines, and no guard is allowed on them, we need $n+1$ guards to guard the plane. If we use semilines, it is easy to construct examples for which n guards are required.

4.2 Illuminating line segments

We now turn our attention to the following problem studied by Czyzowicz, Rivera-Campo, Urrutia, and Zaks [40]: Instead of viewing our line segments as obstacles, let us try to illuminate them. Notice however that a point on a line can be illuminated by a ray hitting the line from either side of it. How many lights are needed to illuminate all the line segments? We prove:

Theorem 4.2.1 *Any collection F of n disjoint line segments can be illuminated with at most* $\lceil \frac{2n}{3} \rceil - 3$ *light sources.*

Proof: Let $F = \{L_1, \ldots, L_n\}$ be a family of n disjoint line segments. Choose a triangle T containing all of the elements of F in its interior and let F' be the family containing all of the elements of F together with three line segments L_{n+1}, L_{n+2} L_{n+3} obtained from T by shortening the three sides of T by an $\epsilon > 0$, ϵ sufficiently small; see Figure 4.8.

Construct a family $H = \{S_1, \ldots, S_n, S_{n+1}, S_{n+2}, S_{n+3}\}$ of $n+3$ *strictly convex* compact sets with mutually disjoint interiors (our sets are allowed to touch each other at a single tangency point located in their boundaries) satisfying the following properties:

1. L_i is contained in S_i, $i = 1, \ldots, n+3$.
2. The number of points at which pairs of elements of H are tangent is maximized.

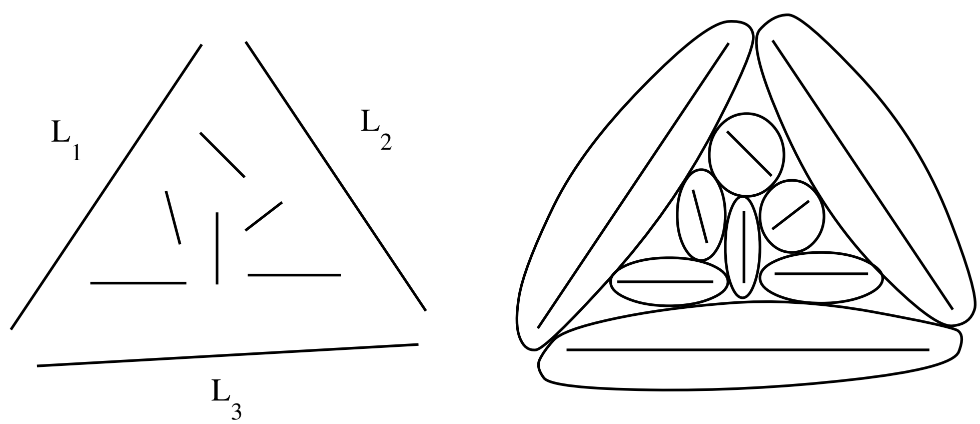

Figure 4.8: Illuminating line segments.

We can now see easily that every element in H is tangent to at least three elements in H. Construct a graph G as follows: For each element of H

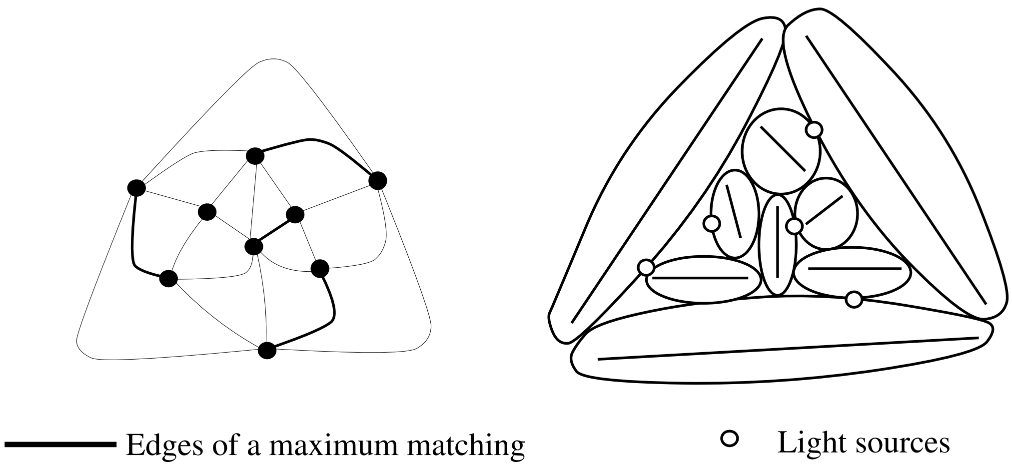

Figure 4.9: Placement of lights.

insert a vertex in G. Two vertices are adjacent if their corresponding sets in G are tangent; see Figure 4.9.

It is now easy to see that G is planar and 2-connected. In addition, since each set in H is tangent to at least three other elements in H the degree of each vertex in G is at least three. Then by Theorem 3.3.2, G has a matching M of size at least $\lceil\frac{n+3+4}{3}\rceil = \lceil\frac{n+1}{3}\rceil + 2$. For each pair of elements S_i, S_j matched in M by an edge of G, place a light source at the point in which they intersect. This light source will completely illuminate the line segments L_i and L_j contained in S_i and S_j respectively. Since M has at least $\lceil\frac{n+1}{3}\rceil+2$ elements, $2(\lceil\frac{n+1}{3}\rceil + 2)$ elements of F will be illuminated using $\lceil\frac{n+1}{3}\rceil + 2$ lights. For the remaining elements of F, an extra light source per element is needed. Then the total number of lights required with our technique is:

$$(\lceil\tfrac{n+1}{3}\rceil + 2) + ((n+3) - 2(\lceil\tfrac{n+1}{3}\rceil)) = n + 1 - \lceil\tfrac{n+1}{3}\rceil \leq \lceil\tfrac{2n}{3}\rceil - 3. \qquad \blacksquare$$

In terms of lower bounds, we notice that the set of n lines obtained in Theorem 4.0.1 also requires $\lfloor\frac{2n-3}{5}\rfloor$ to illuminate them. If we also consider illumination of sets of line segments and semilines, we can obtain sets that require $\lceil\frac{n}{2}\rceil$ points to illuminate them as follows: Consider a regular polygon P_n with n vertices. In the clockwise direction, extend all the sides of P_n to obtain a set of n semilines. Shorten them a bit at their finite endpoint to make them disjoint.

The gap between the lower and upper bounds for the illumination of line segments (or semilines for that matter) is large. A challenging problem is that of closing those gaps. For line segments, I believe that the correct bound is somewhere around $\lfloor \frac{2n-3}{5} \rfloor$. However even a proof that any family of n line segments can always be illuminated with $\lceil \frac{n}{2} \rceil + c$ lights has proved elusive, c a constant.

The case when our lines are *orthogonal*, i.e. parallel to the x- or y-axis is also open. In this case we can prove that any set of n orthogonal line segments can always be illuminated with $\lceil \frac{n+1}{2} \rceil$ lights, see Exercise 1. For the lower bound, we can only produce families that require $\lfloor \frac{n}{3} \rfloor$ lights; see Figure 4.10.

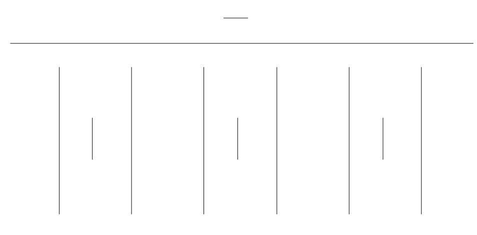

Figure 4.10: A set of orthogonal lines that requires $\lfloor \frac{n}{3} \rfloor$ lights.

4.3 Hiding behind walls

In this section, we study the following problem introduced by Hurtado, Serra and Urrutia [78]: Consider a set of disjoint line segments F. A collection of points P is called *hidden* with respect to F if any line segment joining two points in P intersects an element of F.

We now prove:

Theorem 4.3.1 *Any family of n disjoint line segments has a hidden point set with at least $\sqrt{n}$ elements.*

To prove our result, we will use the following result of Erdös and Szekeres [48]:

Theorem 4.3.2 *Any sequence of n numbers has an increasing or a decreasing subsequence with at least $\sqrt{n}$ elements. The bound is tight.*

For example, consider the following sequence of numbers:

$\pi = \{4, 6, 3, 7, 2, 6, 1, 9, 8\}$

The numbers $4, 6, 7, 9$ form an increasing subsequence of π. Notice that the elements of the subsequence do not have to appear consecutively in π. Similarly $7, 2, 1$ form a decreasing subsequence of π. The proof of this result is left as an exercise to the reader. See exercise 3. Consider a set $F = \{L_1, ..., L_n\}$ of line segments such that their projections on the x axis form a disjoint collection of intervals. Let s_i be the slope of L_i, $i = 1, \ldots, n$. We now prove:

Lemma 4.3.1 *If $a_1 <, \ldots, < a_n$ then F admits a hidden set of size n.*

Proof: For each L_i let p_i be its midpoint. Choose a point q_1 below p_i at distance ϵ from L_i, $i = 1, \ldots, n$. We now claim that if we choose ϵ small enough, $q_1, \ldots, q_n$ form a hidden set. Consider two integers $i < j$, and the line segment $L_{i,j}$ joining p_i to p_j. If a_i is smaller than the slope $a_{i,j}$ of $L_{i,j}$, then by choosing ϵ small enough, we can guarantee that the line segment joining q_i to q_j intersects L_i. If a_i is greater than or equal to $a_{i,j}$ the segment joining q_i to q_j will intersect L_j; see Figure 4.11. ■

We observe that Lemma 4.3.1 is also true when $a_1 >, \ldots, > a_n$, placing $q_1, \ldots, q_n$ *above* p_i, $i = 1, \ldots, n$.

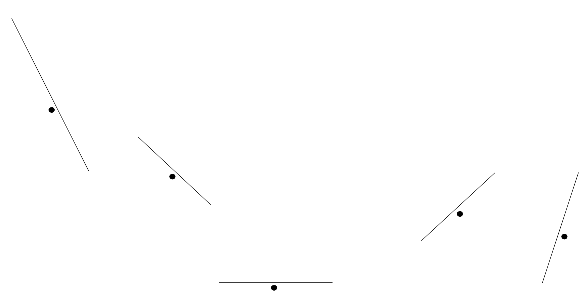

Figure 4.11: A family of 5 line segments with increasing slopes that has a hidden set of size 5.

We now prove Theorem 4.3.1:

Proof: Consider a family of n disjoint line segments. Pick a point p_i in the interior of each L_i in such a way that the x-coordinates $x_1, \ldots, x_n$ of $p_i, \ldots, p_n$ are all different. Assume without loss of generality that $x_1 <, \ldots, < x_n$. For each L_i, choose a segment L'_i contained in it and centered at p_i in such a way that $L'_1, \ldots, L'_n$ have disjoint projections on the x-axis. Let us consider now the sequence of slopes of $L'_1, \ldots, L'_n$. Then by Theorem 4.3.2 this sequence contains an increasing or decreasing subsequence with at least $\sqrt{n}$ elements. Our result now follows from Lemma 4.3.1. ∎

4.4 Excercises

Ex. 4.1 Prove that any family of n isothetic line segments, i.e. parallel to the x- or y-axis, can be illuminated with at most $\lfloor \frac{n+1}{2} \rfloor$ guards. *Hint: Enclose the segments within a rectangle and then partition it into $n+1$ rectangles. Now use Tutte's Theorem. Be careful, in this case we cannot assume that all adjacent faces are connected in the dual graph!*

Ex. 4.2 (Jennings and Lenhart [80] Prove that any family F of n disjoint line segments always has a subset S with at most $\lfloor \frac{n}{2} \rfloor$ elements such that every element in F is in S or is visible from at least one point on an element in S. The bound is tight.

Ex. 4.3 Prove Theorem 4.3.2. *Hint: Let $\pi = \{p_1, \ldots, p_n\}$ be a sequence of numbers. Associate to each p_i a point on the plane with integer coordinates (a_i, b_i) such that a_i is the length of the longest increasing subsequence of elements of π that ends in p_i, and b_i is the lenght of the longest decreasing subsequence of π that* starts *at p_i. Show that if $i \neq j$ then $(a_i, b_i) \neq (a_j, b_j)$. Conclude from this that for some i, a_i or b_i is at least $\sqrt{n}$.*

Ex. 4.4 (Boenke and Shermer [104])Show that if we place n disjoint line segments on the plane, and we restrict our guards to be located only within our line segments, then n guards are always sufficient, and ocassionally necessary to guard the plane.

Chapter 5

Floodlight illumination problems

In the previous chapters, we have assumed that the light sources emit light in all directions, or that the guards can patrol around themselves in all directions. We now present an illumination problem in which the light sources have a restricted angle of illumination. We call such light sources *floodlights*. Thus for the rest of this paper, a floodlight f_i is a source of light located at a point p of the plane, called its apex; f_i illuminates only within a positive angle of illumination α_i, and can be rotated around its apex. We start by study the following problem due to J. Urrutia:

Problem 5.0.1 The 3-floodlight illumination problem*: Let $\alpha_1 + \alpha_2 + \alpha_3 = \pi$ and consider any convex polygon P. Can we place three floodlights of sizes at most $\alpha_1, \alpha_2, \alpha_3$, no more than one per vertex, in such a way that P is completely illuminated?*

We now show that the Three Floodlight Illumination Problem always has a positive solution. Clearly our result is true if P has 3 vertices. Consider any convex polygon P with at least four vertices and suppose that $\alpha_1 \leq \alpha_2 \leq \alpha_3$. Notice first that $\alpha_2 < \frac{\pi}{2}$ and that since P has at least four vertices, the interior angle at one vertex v of P is at least $\frac{\pi}{2}$. Find a triangle T with internal angles α_1, α_2, and α_3 such that the vertex of T with angle α_2 lies on v, and the other vertices of T lie on two points x and y on the boundary of P. Suppose that x and y lie on different edges, say e_1 and e_3 of P. (The case when they lie on the same edge will be left to the reader.) See Figure 5.1.

Place a floodlight f_2 with angle of illumination α_2 at v, illuminating T. Consider the circle C passing through the vertices of T. It is easy to see

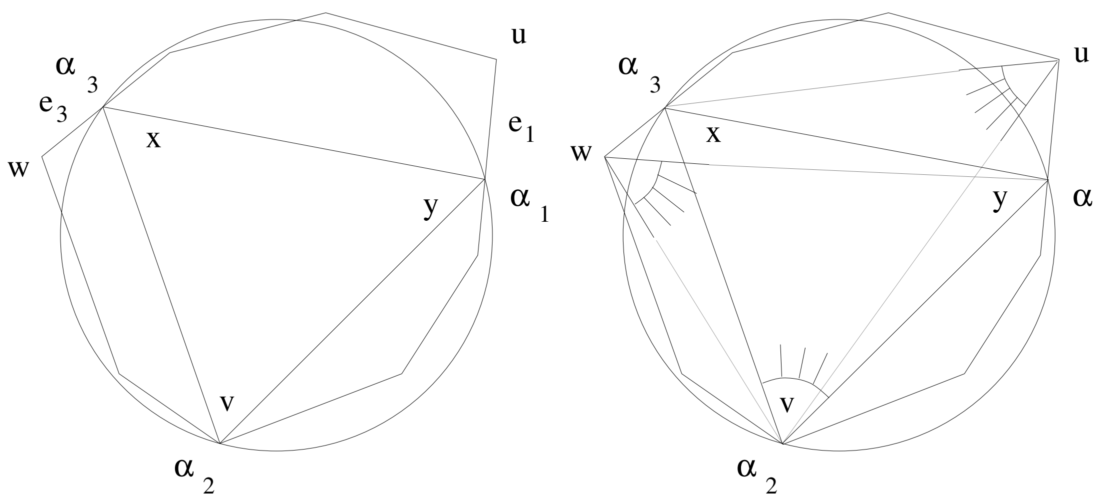

Figure 5.1: Illustration of proof of the Three Floodlight Problem.

that at least one vertex of each of e_1 and e_3 is not contained in the interior of C. Let u and w be these end points. Two cases arise:

1. $u \neq w$. Place floodlights f_1 and f_3 at u and w illuminating the angular region determined by v, u, x and v, w, y. Since u and w are not contained in the interior of C, the angles of illumination of f_1 and f_3 are at most α_1 and α_3 respectively. Since f_1, f_2, and f_3 illuminate P, our result follows.

2. $u = w$. By considering the tangents to C at x and y it is easy to verify that the angle α of P at u is at most $\pi - 2\alpha_2$ which is less than or equal to $\alpha_3 = \pi - (\alpha_1 + \alpha_2)$. Place a floodlight of size α at u. This illuminates P. ∎

5.1 Illuminating the plane

We now study the following floodlight illumination problem of the plane studied in [14]. Suppose we have four points $p_1, \ldots, p_4$ on the plane and four $\frac{\pi}{2}$ floodlights, one at each p_i. Can we orient them in such a way that all of the plane is illuminated? (We assume here that our lights cast no shadows.)

To see that under these conditions the plane can always be illuminated, consider a line that leaves two points on one side of it. Assume without loss of generality that the line is parallel to the x-axis. It is easy to see now that using the floodlights on top of l, we can illuminate the region below it, and using the lights below l, we can illuminate the region above it, as in Figure 5.2. It follows that the number of solutions to our plane illumination problem is actually infinite!

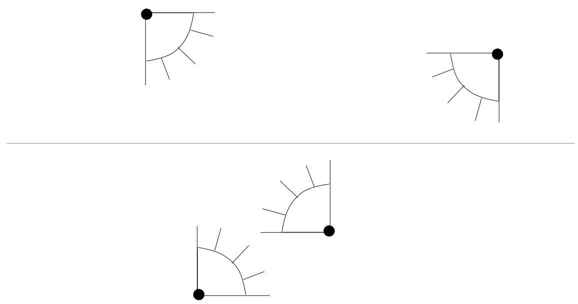

Figure 5.2: Four $\frac{\pi}{2}$ floodlights suffice to illuminate the plane.

We now prove an interesting generalization of the previous result. Let $\{\alpha_1, \ldots, \alpha_n\}$ be a set of angles such that each of them is at most π and $\alpha_1 + \ldots \alpha_n = 2\pi$. Consider a set of floodlights $f_1, \ldots, f_n$ such that the size of f_i is α_i, $i = 1, \ldots, n$.

Theorem 5.1.1 *Let P_n be a collection of n points on the plane, and $f_1, \ldots, f_n$ a set of floodlights of sizes $\{\alpha_1, \ldots, \alpha_n\}$ such that $\alpha_1 + \ldots + \alpha_n = 2\pi$. Then we can locate all our floodlights, one per element of P_n and point them in such a way that the plane is completely illuminated.*

Some preliminary results will be needed to prove our result. They are of interest in their own right. Suppose that we have three angles $\alpha_1, \alpha_2, \alpha_3$ such that their sum is 2π, and each of them is at most π. Consider any point set P_n with n points, and three integers n_1, n_2, n_3 such that $n_1+n_2+n_3 = n$. We now prove:

Theorem 5.1.2 *There is a point p such that the plane can be split into three infinite angular regions by three rays emanating from p such that the angles between the rays are $\alpha_1, \alpha_2, \alpha_3$ and the region with angle α_i contains exactly n_i elements of P_n in its interior.*

Proof: Choose the three directions 0, α_1, and $\alpha_1 + \alpha_2$, and assume without loss of generality that the angle of inclination of any line connecting two elements of P is not 0, α_1, or $\alpha_1 + \alpha_2$. If an element q of P is such that the three rays emanating from it at angles 0, α_1, and $\alpha_1+\alpha_2$ partition the plane in three regions that contain i, j,and k elements such that two of these values are equal to n_1, n_2, or n_3 in this order, say $i = n_1$, $j = n_2$, then by choosing p slightly above q our result follows. Assume then that this is not the case.

Assign to each point with coordinates (x, y) the vector (i, j, k) in which i is the number of elements q of P_n with the property that the direction d of the line connecting p to (x, y) is $0 \leq d < \alpha_1$. The integer j is the number of elements of P_n with $\alpha_1 \leq d < \alpha_1+\alpha_2$, and $k = n-(i+j)$. If (x, y) is in P_n, it will be considered to belong to the region bounded by the rays emanating from it with directions $\alpha_1 + \alpha_2$ and 2π.

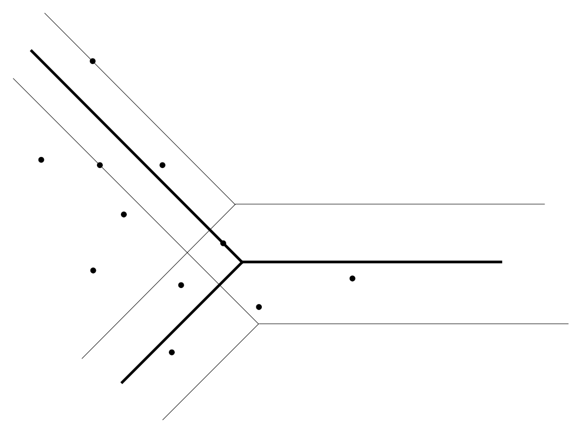

Figure 5.3: Finding a partitioning of the plane for $n = 11$, $n_1 = 2, n_2 = 6$, and $n_3 = 3$, $\alpha_1 = \frac{3\pi}{4}$, and $alpha_2 = \frac{\pi}{2}$ The solution appears in thick lines.

We observe now that for each y we can choose x large enough that $(i, j, k) = (0, n, 0)$. As x decreases in value and y remains constant, the second coordinate of (i, j, k) decreases one by one until it becomes n_2. Let $f(y)$ be the value such that if $x < f(y)$, then $j < n_2$. It is easy to verify that the function defined by $(f(y), y)$ is a piecewise continuous linear function. Moreover if we choose y_0 small enough, then for the point $(f(y_0), y_0)$ we have that $(i, j, k) = (n - n_2, n_2, 0)$ and that for y_1 large enough, we have $(i, j, k) = (0, n_2, n - n_2)$. It now follows that by increasing y_0 continuously, the value of the first coordinate of (i, j, k) decreases one by one, until it reaches the value n_3; see Figure 5.3. Our result follows. ∎

Let $\alpha \leq \pi$. An α-*wedge* of the plane is a region W bounded by two rays emanating from a point p, called the apex of W, such that the angle between the rays generating W is α; α is called the size of W. We define W^{-1} to be the wedge with apex at p but generated by rays emanating from p in opposite directions from those that generated W; see Figure 5.4.

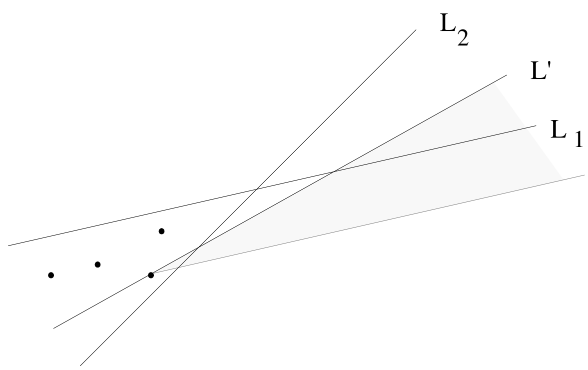

Figure 5.4: Illuminating a wedge.

We now prove:

Lemma 5.1.1 *Consider an α wedge W, and set of floodlights with sizes $\{\alpha_1, \ldots, \alpha_m\}$ such that $\alpha_1 + \ldots + \alpha_m = \alpha$, and a set of points P_m contained in W^{-1}. Then we can always illuminate W by placing a floodlight at each element of P_m.*

Proof: Our result is clearly true if $m = 1$. Suppose then that $m > 1$. Suppose that W and W^{-1} are opposite regions bounded by two lines L_1 and L_2 that intersect at a point p, and that W is above L_1 as shown in Figure 5.4. Consider a line L' that intersects L_1 at an angle α_1, contains exactly one point q of P_m, and all the remaining elements of P_m lie on or above L'. Notice that L' and L_1 define a wedge W_1 of size α_1 such that q belongs to W_1'. W_1 can now be illuminated from q using a floodlight of size α_1. Notice now that L' also defines with l_2, a wedge W" of size $\alpha - \alpha_1$, and all the elements of $P_m - q$ lie on W''^{-1}. By induction, we can now illuminate W''^{-1} with the floodlights of sizes $\alpha_2, \ldots, \alpha_m$, and our result follows. ∎

We are ready to prove Theorem 5.1.1. Partition our set of floodlights into three sets $F_1 = \{f_1, \ldots, f_i\}$, $F_2 = \{f_{i+1}, \ldots, f_j\}$ and $F_3 = \{f_{j+1}, \ldots, f_n\}$ such that $\beta_1 = \alpha_1 + \ldots + \alpha_i$, $\beta_2 = \alpha_{i+1} + \ldots + \alpha_j$ and $\beta_3 = \alpha_{j+1} + \ldots + \alpha_n$ are all less than π. Let $n_1 = i$, $n_2 = j - i$, and $n_3 = n - j$. By Theorem 5.1.2, we can partition the plane into 3 wedges W_1, W_2 and W_3 of sizes β_1, β_2 and β_3 such that each of them contains exactly n_1, n_2 and n_3 points. By Lemma 5.1.1, we can now illuminate W_1^{-1}, W_2^{-1} and W_3^{-1} by using the points in W_i, and using the floodlights at F_i, we can illuminate W_i^{-1}, $i = 1, \ldots, 3$. However $W_1^{-1} \cup W_2^{-1} \cup W_3^{-1}$ covers the plane! ∎

5.2 Illumination of stages

We now study the following problem introduced by J. Urrutia [14] in 1990. Suppose that we are testing the set of stage lights of a new theater, with a number of stage lights. The *Stage illumination problem*, consists in determining if the set of stage lights of the theater are enough to illuminate all of the stage at once. In mathematical terms we cast this problem as follows:

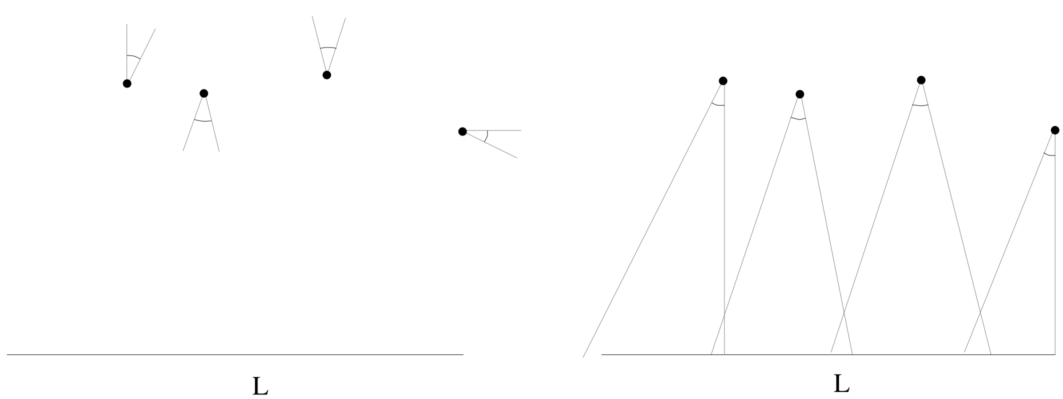

Figure 5.5: The stage illumination problem.

Suppose we have a line segment, representing the stage of a theater, and a set of stage lights, modeled by floodlights. Is it possible to point our floodlights in such a way that the line segment is completely illuminated?

The existence of an efficient algorithm to solve the stage illumination problem remains open. We now proceed to solve a problem that at first appears harder:

Given a set $F = \{f_1, \ldots, f_n\}$ of floodlights each with angle $\alpha_i, i = 1, \ldots, n$ we can associate to F an angular cost equal to $\alpha_1 + \ldots + \alpha_n$. We now solve the following problem:

Optimal floodlight illumination of stages Consider a stage, represented by a line segment S and a set $P = \{p_1, \ldots, p_n\}$ of n points. Determine a set of floodlights F that illuminates S such that the angular cost of F is minimized and the apex of each floodlight $f_i \in F$ is located at some point $p_j \in P$.

We point out here that we will allow for more than one floodlight to be located at each point of P. Moreover, we assume that each floodlight has size strictly greater than 0. Interestingly enough, our solution is such that only at one point in P we may place two lights; at the other points we place one or no light at all.

5.2.1 Illuminating the real line

Our first step consists in solving the problem of illuminating the real line using floodlights placed on the elements of a point set $P = \{p_1, \ldots, p_n\}$. To make our presentation easier, we will assume that all of our points have different y-coordinates. This restriction can be easily taken care of.

We study now the problem of illuminating the real line using only two points p_i and p_j. Assume that p_i is lower than p_j. Consider the two circles tangent to the real line and containing p_i and p_j, and let their intersection points with the real line be labelled $x_{i,j}$ and $y_{i,j}$ in such a way that $x_{i,j} < yi, j$; see Figure 5.6(a).

We now prove:

Lemma 5.2.1 *In the optimal floodlight illumination of the real line L from $\{p_i, p_j\}$ all points in the interval $[x_{i,j}, y_{i,j}]$ are illuminated from p_j, and all points in the intervals $(\infty, x_{i,j}]$ and $[y_{i,j}, \infty)$ are illuminated from p_i.*

Proof: Suppose that in an optimal illumination of the real line with a set of floodlights F, an interval with endpoints $a < b$ contained to the left of $x_{i,j}$ is illuminated by a small floodlight f at p_j. Consider the circle passing through a, b and p_j. It is easy to see that this circle leaves p_i outside, and thus the angle a, p_i, b is smaller than the angle a, p_j, b. Thus if we substitute f by a floodlight f' at p_i we obtain a set of lights that illuminates the real line

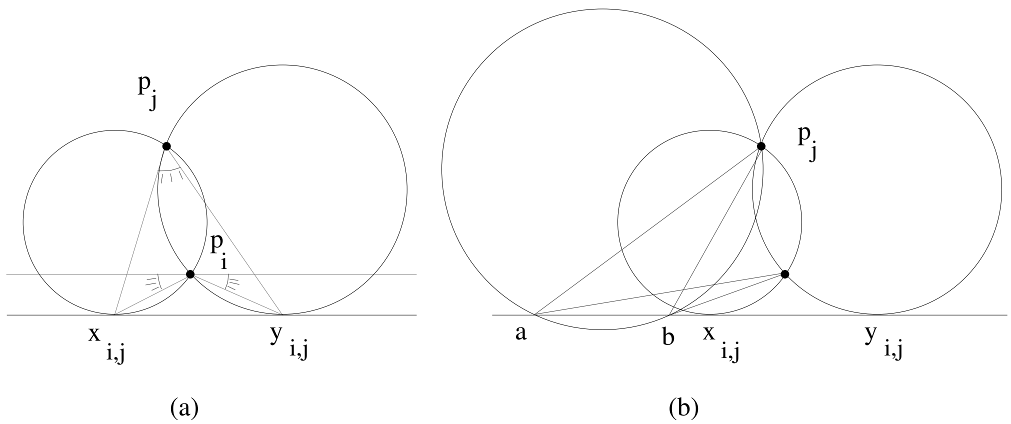

Figure 5.6: Illuminating the real line from two points.

with smaller weight than that of F, which is a contradiction. In now follows that all the points to the left of $x_{i,j}$ are illuminated from p_j. Similarly we can conclude that the interval $[x_{i,j}, y_{i,j}]$ is illuminated from p_j and $[y_{i,j}, \infty)$ is illuminated from p_i, see Figure 5.6. ∎

Next we prove:

Lemma 5.2.2 *Let $P = \{p_1, \ldots, p_n\}$ be a collection of points, and p_i a point in the interior of the convex hull of P. Then in any optimal illumination of the real line with floodlights at points of P, there is no floodlight located at p_i.*

Proof: Suppose that p_i is an interior point of the convex hull of P, and that there is an optimal illumination of the real line in which a floodlight f_i placed at p_i illuminates an interval, say $[x, y]$. Consider the smallest disk D containing x, y and all of the elements of P, see Figure 5.7.

Let p_j be a point of P located in t he boundary of D. Since p_i is in the interior of the convex hull of P, $p_i \neq p_j$; moreover p_i belongs to the interior of D. Therefore the angle x, p_i, y is greater than angle x, pj, y. Thus we could substitute the floodlight at p_i that illuminates the interval $[x, y]$ by a smaller one placed at p_j that illuminates the same interval . This contradicts our assumption on the optimality of F. ∎

The following corollary now follows:

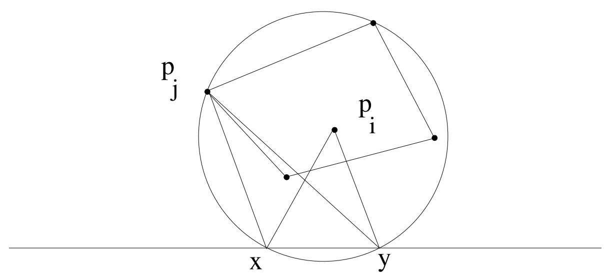

Figure 5.7: Only points in the convex hull are useful.

Corollary 5.2.1 *Consider an optimal floodlight illumination of the real line wit a set of floodlights F on a set of points P. Then if a point x is illuminated by a floodlight of F located at a point p_i of P, the disk circle C tangent to the real line and containing p_i contains all of the elements of P.*

It now follows that the optimal solution to the floodlight illumination problem of the real line can be obtained as follows: Consider the leftmost point $y_{i,j}$ generated by all pairs of points in the convex hull of P, and obtain the circle tangent to the real line and passing throuh p_i and p_j. Slide to the left a point p initially located at $y_{i,j}$. While we move p to the left, maintain a circle C tangent to the real line at p, and tangent to the convex hull of P. For any position of p, the vertex p_k on the convex hull of P and C, is the point were the floodlight that will illuminate p must be located, see Figure 5.8. Notice that at a finite set of positions of p, C will contain two points in the convex hull of P. These points correspond to the beginning and ending of the intervals to be illuminated by the elements of P, this situation is illustrated in Figure 5.8.

we now proceed to develop an algorithm to solve our floodlight illumination problem:

Algorithm 5.2.1 *FLIP*

Input *: A set $P = \{p_1, \ldots, p_n\}$ of n points on the plane with positive y-coordinates, and the real line L.*

Output *: A partitioning of the lines into a sequence of at most $n + 1$ intervals, each assigned to a point of P from where that interval is to be illuminated.*

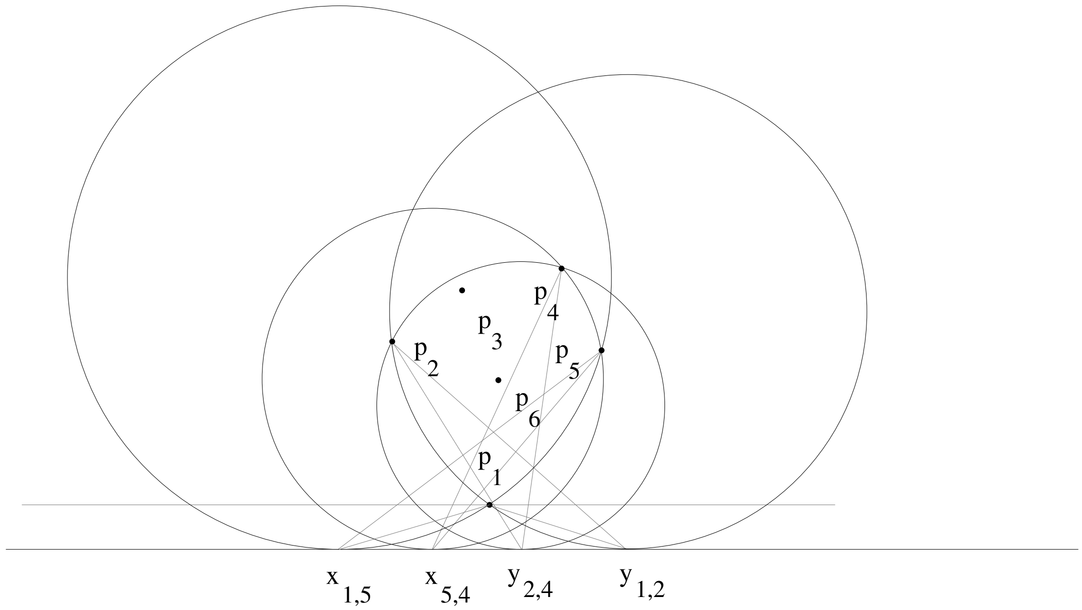

Figure 5.8: Optimal illumination of the real line.

1. *Calculate the convex hull of P. Relabel the vertices of the convex hull of P in the clockwise direction by $\{p_1, \ldots, p_k\}$ where p_1 is the point of P closest to L and k is the number of vertices in the convex hull of P.*

2. *Determine the point $y = y_{1,i}$ which is the rightmost point in $\{y_{1,j}; j = 2, ..., n\}$. Illuminate all of the points in the interval $[y, \infty)$ from p_1.*

   While i1 do:

   (a) *Find the smallest index $j > i$ (or take $j = 1$ if no such j exists) such that the disk D defined by the circle tangent to L containing p_i and p_j contains all the elements of P. Let x be the point in which C is tangent to L.*

   (b) *Illuminate the interval $[x, y]$ from p_i.*

   (c) *$i \leftarrow j$, $y \leftarrow x$*

   EndWhile

3. *: Illuminate the interval $(-\infty, y]$ from p_1.* **Stop**

For example in Figure 4, y initially takes the value $y_{1,2}$, and i the value 2. In the next iteration, y changes to $y_{2,4}$ and i to 4. Notice that even though p_3 is a vertex in the convex hull of $P = \{p_1, \ldots, p_6\}$, no interval is illuminated by p_3. This happens because the circle tangent to L through p_2 and p_3 does not contain p_4, and in the execution of our While loop for $i = 2$, j skips the value 3. The subsequent values for y are $x_{5,4}$ and $x_{1,5}$ and the values for i are 5 and 1 respectively. Thus all the points in the intervals $(-\infty, x_{1,5}]$ and $[y_{1,2}, \infty)$ are illuminated from p_1, and $[y_{2,4}, y_{1,2}]$, $[x_{5,4}, y_{2,4}]$ and $[x_{1,5}, x_{5,4}]$ are illuminated from p_2, p_4 and p_5 respectively.

5.3 Floodlight illumination of polygons

Let us consider a polygon P with n vertices. Observe that if we are allowed to place $n-2$ floodlights of size $\frac{\pi}{3}$ at the vertices of P, we can always do so in such a way that P is completely illuminated. This follows from the observation that any triangulation of P has $n-2$ triangles, each of which in turn has an internal angle of size at most $\frac{\pi}{3}$. Observe, however, that using this rule to illuminate P, we may place more than one floodlight at some vertices of P. If we allow a floodlight of size at most π at each vertex, we can indeed illuminate all of P. If P is a triangle, we can illuminate it by placing a floodlight at any of its vertices of size at most π. Suppose then that P has more than 3 vertices. Take any triangulation T of P, choose an ear e of T. Let v be the vertex of e of degree 2 in T. Place a floodlight of size equal to the internal angle of P at v. Delete v and by an inductive argument we can illuminate $P - v$. Our result follows.

A natural question now arises: is it possible to illuminate any polygon by placing at each vertex of P at most one floodlight of size at most α for some $\alpha < \pi$? We will now prove that for arbitrary polygons, the answer to this question is, surprisingly, no. We now prove:

Theorem 5.3.1 *For any $\epsilon > 0$ there is a polygon P that cannot be illuminated by placing at most one floodlight of size $\pi - \epsilon$ at each of its vertices.*

Proof: We will prove our result only for orthogonal floodlights. The proof presented here is due to O'Rourke and Xu [105]. The full result follows easily from the ideas presented here, and is left as an exercise. The interested reader can find the complete proof in [49]. Consider the symmetric polygon P shown in Figure 5.9. Divide P into a left and right section by cutting it along a vertical through vertex v. Notice that all the convex internal angles at the vertices of P are slightly bigger than $\frac{\pi}{2}$. Label the vertices of P as shown in Figure 5.9 and notice that the extensions l_2, l_1 of the edges of P connecting v_2 to r_2, and v_1 to r_1, respectively are such that l_2 hits the base of P to the left of the point at which l_1 hits it.

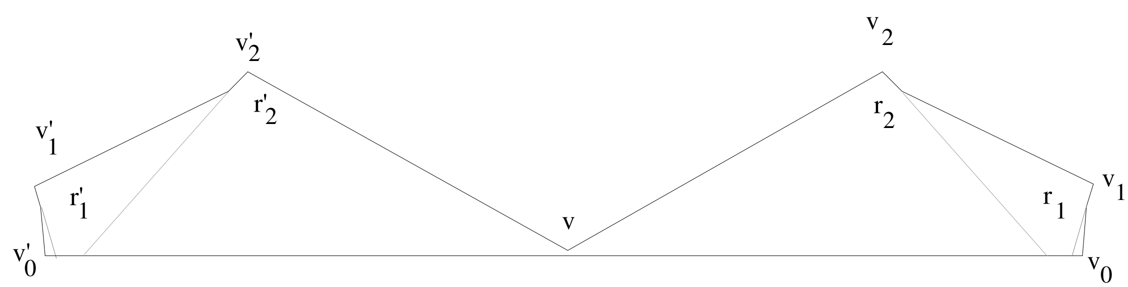

Figure 5.9: A polygon that cannot be illuminated by placing an orthogonal floodlight at each vertex.

We now show that P cannot be illuminated by placing an orthogonal floodlight at each of its vertices. The orthogonal floodlight at v illuminates vertices in L or R, but not both. Assume without loss of generality that this light illuminates vertices in L. Since the internal angle at v_2 is bigger than $\frac{\pi}{2}$ the only way we can illuminate all the points in a small neighborhood of v_2 is by using the floodlights at v_2 and r_2. Similarly to illuminate all the points in a small neighborhood of v_1, we need the lights at v_1 and r_1. Notice, however, that this will prevent the floodlights at r_1 and r_2 from pointing towards v_0. It now follows that all the points in a small neighborhood of v_0 have to be illuminated using only the floodlight at v_0. Notice that this is impossible, since the internal angle of P at v_0 is greater than $\frac{\pi}{2}$. Our result now follows. ■

The following problem arises naturally from the above result:

> *How many π-floodlights are always sufficient to illuminate any polygon with n vertices?*

F. Santos has produced a family of polygons with $5n+1$ vertices that require $3n$ vertex π floodlights to illuminate them; see Figure 5.10. We now conjecture:

Conjecture 5.3.1 $\lfloor \frac{3n-3}{5} \rfloor$ *π vertex floodlights are always sufficient and occasionally necessary to illuminate any polygon with n vertices.*

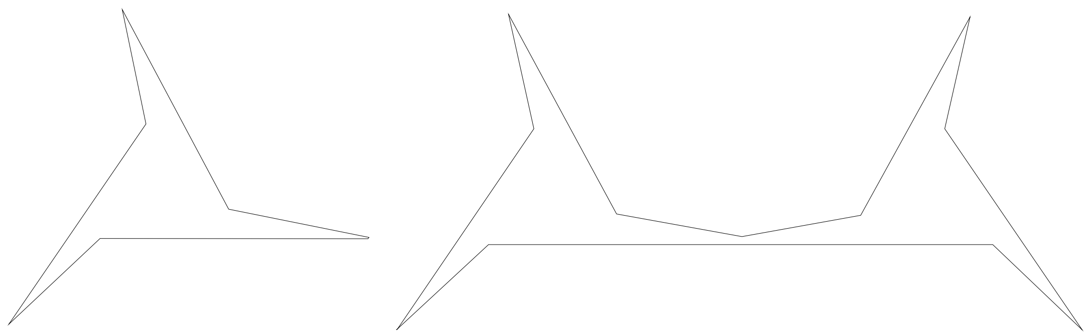

Figure 5.10: A family of polygons with $5n+1$ vertices that require $3n$ π vertex floodlights.

5.4 Orthogonal polygons

We now turn our attention to the study of the illumination of orthogonal polygons with vertex floodlights. A floodlight of size $\frac{\pi}{4}$ will be called an orthogonal floodlight. We now prove:

Theorem 5.4.1 *Any orthogonal polygon P with n vertices can be illuminated with at most $\lfloor \frac{3n-4}{8} \rfloor$ orthogonal vertex floodlights.*

Proof: We start by proving that P can be illuminated by choosing a set of orthogonal floodlights according to the following rule, which we call the *top-left* illumination rule:

Top-left illumination rule:

1. At the top vertex of every left edge of P place an orthogonal floodlight illuminating the set of directions between $\frac{3\pi}{4}$ and 2π.

2. At the left vertex of every top edge of P again place an orthogonal floodlight illuminating the set of directions between $\frac{3\pi}{4}$ and 2π.

To see that the set of floodlights chosen by the top-left illumination rule illuminates P, choose any point q in P. Consider the longest horizontal line segment h passing by q, totally contained in P. Slide h upwards until it hits a top edge of P or it reaches the top vertex of the edge of P containing the left end point of h. In the former case, there is a floodlight at this vertex that illuminates q. Suppose then that h hits a top edge e of P. Notice that

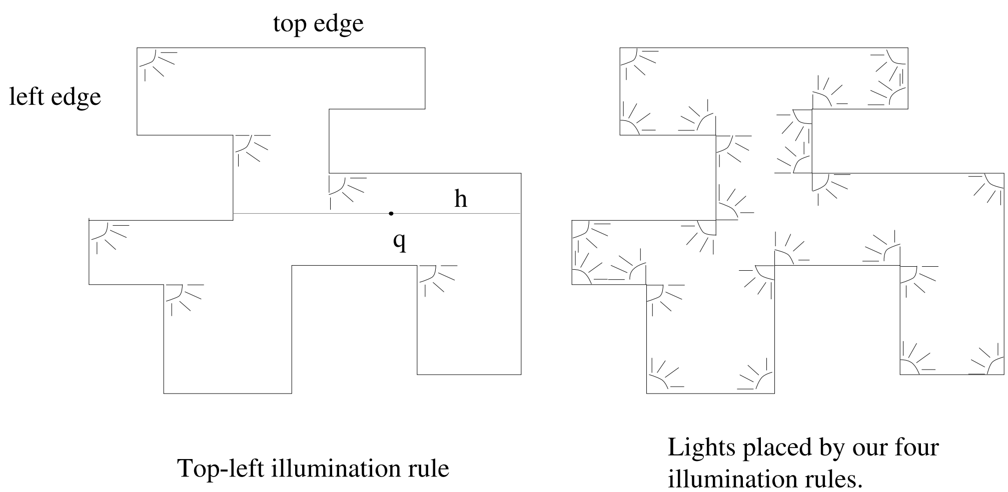

Figure 5.11: Illuminating an orthogonal polygon with orthogonal floodlights.

by our illumination rule, we have a floodlight at the left vertex of e, and this floodlight illuminates q; see Figure 5.11.

In a similar way, we can define the *top-right*, *bottom-right*, and *bottom-left* illumination rules, each of which illuminates P. Moreover, if we use the four illumination rules simultaneously, we will place exactly two floodlights at each reflex vertex, and one at each convex vertex. Thus if P has r reflex vertices, the total number m of floodlights used by our four rules is $2r + (n - r)$. By Lemma 2.2.1 $r = \frac{n-4}{2}$, and thus we have that $m = 2\frac{n-4}{2} + \frac{n+4}{2} = \frac{3n-4}{2}$. It now follows that one of our illumination rules has at most $\lfloor \frac{3n-4}{8} \rfloor$ floodlights.

The polygon P_{12} shown in Figure 5.12(a) has 12 vertices and requires $4 = \frac{3(12)-4}{8}$ orthogonal vertex floodlights. If we now *paste* copies of this polygon as shown in the same figure, we can generate a family of polygons with $12 + 8k$ vertices, each of which requires $4 + 3k$ vertex floodlights. Our result now follows. ∎

We now show that $\frac{\pi}{2}$ is essential in our previous result, i.e. we prove:

Theorem 5.4.2 *For any $\epsilon > 0$ there is an orthogonal polygon that cannot be illuminated with $\frac{\pi}{2} - \epsilon$ vertex floodlights.*

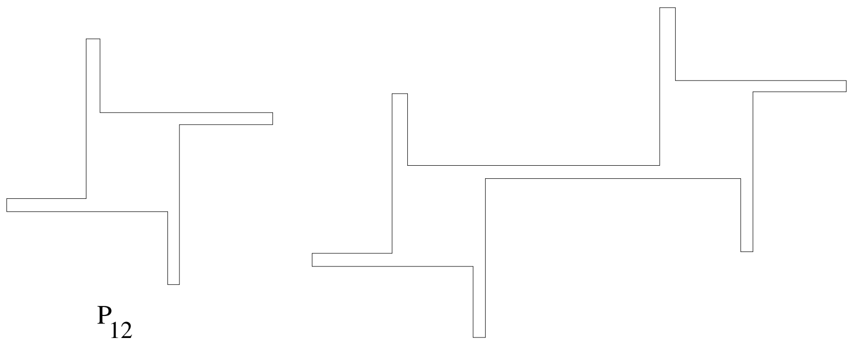

Figure 5.12: A family of polygons that requires $\lfloor \frac{3n-4}{8} \rfloor$ orthogonal vertex floodlights.

Proof: For any $\epsilon > 0$ construct an orthogonal polygon with 12 vertices consisting of a square S with four long rectangles $R_1, \ldots, R_4$ attached to it as shown in Figure 5.13. Label the vertices of P as shown in the same figure. It is easy to see that P can be constructed in such a way that the following conditions hold:

1. The angle $u_i, v - i, r_i$ is greater than $\frac{\pi}{2} - \epsilon$, $i = 0, \ldots, 3$.

2. The angle p, r_i, r_{i+1} is greater than $\frac{\pi}{2} - \epsilon$, $i = 0, \ldots, 3$ addition taken mod4.

Under these conditions, we can verify that the $\frac{\pi}{2} - \epsilon$ floodlights at u_i and v_i cannot illuminate all of $R_i, i = 0, \ldots, 3$ To complete the illumination of each R_i, we must then use a floodlight placed at a reflex vertex of P, $i = 0, \ldots, 3$. It now follows that the point p at the center of S is not illuminated. ■

5.5 The two floodlight illumination problem

In this section we study the following problem: Given a convex polygon P with n vertices, find two floodlights of sizes α_1 and α_2 that illuminate P and their locations such that $\alpha_1 + \alpha_2$ is minimized.

Consider two points a and b on the boundary of a polygon P, not necessarily vertices of P. The vertex interval (a, b) is defined to be the set of

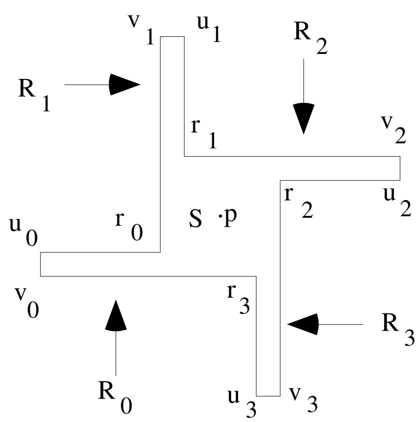

Figure 5.13: An orthogonal polygon that cannot be illuminated with orthogonal vertex floodlights.

vertices of P that we meet when we move in the clockwise direction on the boundary of P from a to b. For short, we will refer to the interval (a, b) rather than to the vertex interval (a, b). Notice that (a, b)
$neq(b, a)$. In particular, when a and b are interior points of edges of p, $(a, b) \cup (b, a)$ is the set of vertices of P. A pair of floodlights F_1 and F_2 that illuminates a polygon P will be called a floodlight illuminating pair and will be called a FLIP of P. If F_1 and F_2 are such that the sum of their apertures is minimized, we call them an optimal FLIP.

We call F_1, F_2 an opposite FLIP if the intersection of the regions illuminated by F_1 and F_2 is a quadrilateral with all of its vertices on the boundary of P, see Figure 5.14(a). If the interior of the regions illuminated by a FLIP F_1 andF_2 are disjoint, we call F_1, F_2 a dividing FLIP, see Figure 5.14(b). Observe that an optimal FLIP must be either an opposite FLIP or a dividing FLIP.

We now prove:

Lemma 5.5.1 *Let F_1, F_2 be an optimal FLIP of a polygon P. Then the apexes of F_1 and F_2 are located at vertices of P.*

Proof: Suppose that F_1, F_2 form an optimal FLIP. Two cases arise:

(a) F_1 and F_2 form an opposite FLIP with apexes at q and r such that r is not a vertex of P. Let x, r, y, q be the vertices of the intersection of the

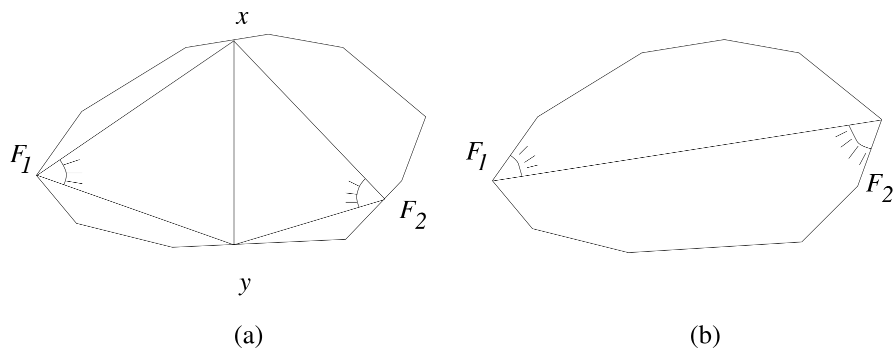

Figure 5.14: An opposite and a dividing FLIP.

areas illuminated by F_1 and F_2. Consider the circle $C(x, y, p)$ that passes through x, y and a vertex p of P such that $C(x, y, p)$ contains all the vertices of the interval (x, y) of P. Notice that the angle spanned by p and the line segment joining x to y is smaller than that spanned by r and the same segment. Therefore if we replace F_2 by a floodlight F_3 with apex in p and illuminating the angular sector determined by x, p, and y, F_1 and F_3 also illuminate P and the sum of their sizes is smaller than that of F_1 and F_2, see Figure 5.15(a).

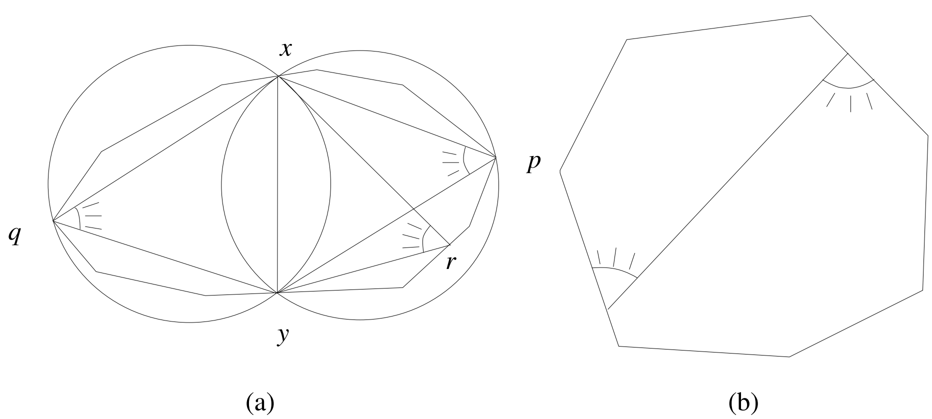

Figure 5.15: The vertices of an optimal FLIP must lie on vertices of P.

b) $F_1$and F_2 form a partitioning pair Figure 5.15(b). We can easily verify that in this case, we can slide the apex of F_1 and F_2 towards two vertices of P such that the sum of the sizes of F_1 andF_2 decrease. ∎

Suppose now that we want to find a FLIP F_1 and F_2 of a polygon P in such a way that their apexes are located on two fixed vertices p and q of P, and the sum of the sizes of F_1 and F_2 is minimized. Call such a FLIP an optimal $(p,q) - FLIP$. We now show:

Lemma 5.5.2 *Given two vertices p and q of P, we can find an opposite optimal $(p,q) - FLIP$ in linear time.*

Proof: Our objective here is to find two points, x and y, on the boundary of p such that the sum of the angles $\alpha = ypx$ and $\beta = xqy$ is minimized. Consider the angles γ and δ formed by pxq and qyp as in Figure 5.16(a).

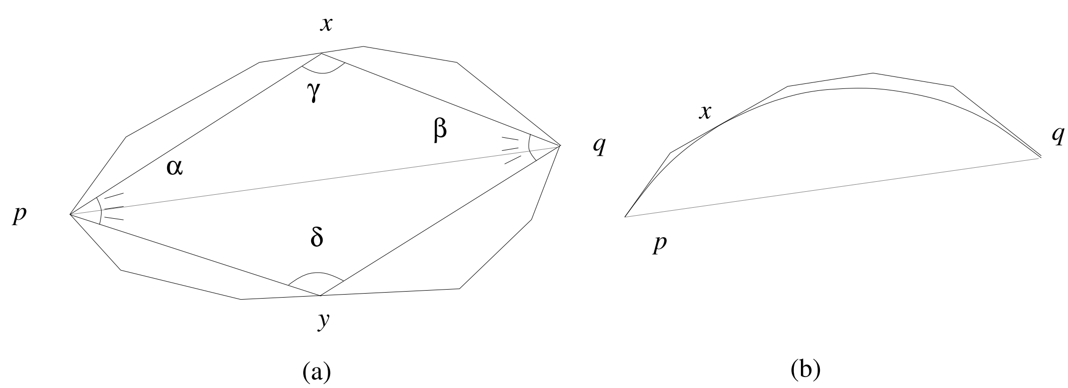

Figure 5.16: Finding an optimal $(p,q) - FLIP$.

We now observe that minimizing $\alpha + \beta$ is equivalent to maximizing $\gamma + \delta$. However, the maximization of $\gamma + \delta$ can be achieved by maximizing independently γ and δ! Thus all we need to do is to locate for each edge e of the chain of edges of P from p to q the point r_e on e that maximizes the angle pr_eq, and keep the point r_e wich maximizes γ. Clearly this can be done in constant time per edge, and x can be found in linear time. The same procedure is applied to find y. ∎

The previous Lemma, provides an $O(n^3)$ algorithm to find the optimal FLIP. To reduce the complexity of our algorithm to $O(n^2)$, we need to provide a criterion to reduce the number of pairs of vertices of P that can be apexes of an optimal opposite FLIP. Consider the set D of all diagonals of P joining all pairs of vertices of P with the property that there is a circle through their endpoints that contains P. Under our general circular position on the vertices of P, it is easy to see that the elements of D induce a triangulation $T(P)$ of P, see Figure 5.17(a). Using standard techniques for calculating Voronoi Diagrams of convex polygons, D and $T(P)$ can be found in linear time [4].

A subset of three vertices $\{u, v, w\}$ of the vertex set of p is called a $c-triple$ if $\{u, v, w\}$ is the set of vertices of a triangle of $T(P)$. By the definition of $T(P)$, it follows that the circle $C(u, v, w)$ determined by $\{u, v, w\}$ contains P. Under our general position assumption, the number of $c-triples$ of P is exactly $n-2$. We say that two $c-triples$ $\{u, v, w\}$ and $\{x, y, z\}$ are adjacent if they share two common elements. If $\{u, v, w\}$ and $\{u, v, x\}$ are two adjacent $c-triples$, The vertices w and x will be called the antipodal vertices of $\{u, v, w\}$ and $\{u, v, x\}$.

Given two vertices p and q of P, we say that a $c-triple$ $\{r, s, t\}$ *separates* them if p and q are both different from r, s, and t and when we traverse the boundary of p from p to q (and from q to p) in the clockwise order, we meet either one of r, s or t before we meet q (resp. p), see Figure 5.17(b).

Lemma 5.5.3 *If two vertices p and q of a polygon P are apexes of an optimal FLIP, then there is no $c-triple$ that separates them.*

Proof: Suppose there is a $c-triple$ that separates p and q, and that there is an optimal opposing FLIP F_1, F_2 of P with apexes at p and q, see Figure 5.17(b). Suppose further that F_1 and F_2 meet at two points x and y in the interior of two edges of P. As in the proof of Lemma 5.5.1, the circle through x, y and q contains in its interior all of the vertices of P between x and y in the clockwise direction, except for q itself, see Figure 5.17(b). Similarly for the circle trough y, p and x. However, this will force one of these circles has to intersect the circle through r, s and t at least four times, which is impossible. ∎

Corollary 5.5.1 *Let p and q, be the apexes of an optimal FLIP. Then p and q form an antipodal pair of vertices of two adjacent $c-triples$, or there is a circle through p and q containing$P$.*

We are now ready to give our algorithm O-FLIP to obtain an optimal floodlight illumination pair of a polygon.

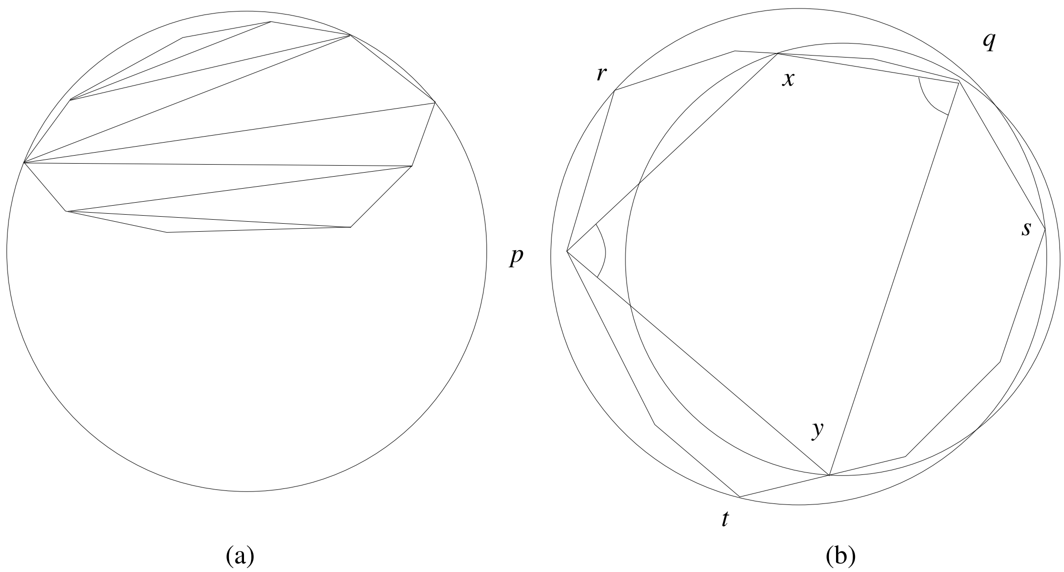

Figure 5.17: A polygon, and its corresponding $T(P)$. The $c - triple$ $\{r, s, t\}$ separates $p$and q .

Algorithm 5.5.1 *O-FLIP Input a polygon P with n vertices.*

1. *Find $T(P)$, and all $c - triples$ of P.*

2. *Find the set S of all pairs of vertices of P that are antipodal pairs, or adjecent in $T(P)$.*

3. *For each pair of p and q in S, find an optimal opposite $(p, q) - FLIP$.*

4. *For each pair of vertices p and q of P find in constant time the partitioning pair of floodlights F_1, F_2 with apexes at p and q that illuminate P and minimizes the sum of the sizes of F_1 and F_2.*

Output the FLIP of minimum weight identified in Steps 3 and 4.

We now have:

Theorem 5.5.1 *Given a convex polygon P with n vertices in general position, O-FLIP finds an optimal FLIP in $O(n^2)$ time.*

Proof: Steps 1 and 2 can be carried out in linear time. Step 3 requires $O(n^2)$ time, since there are at most $n-2$ antipodal pairs and $2n-3$ edges in $T(P)$. Solving each of them requires $O(n)$ time. In Step 4, we need to test all pairs of vertices p and q of P. However, for every of these pairs, there are exactly two partitioning FLIP's, and the complexity of Step 4 is $O(n^2)$. ∎

5.6 Exercises

1. Suppose that we associate to each vertex v of P a weight equal to the size of the internal angle of P at v. Show that we can always choose a set of vertices of P such that they illuminate P and the sum of their weights is at most $\frac{(n-2)\pi}{3}$. Show that this bound is tight.

2. Show that Theorem 5.1.2 still holds if we remove the condition that α_1, α_2 and α_3 are less than or equal to π.

3. Prove that Theorem 5.1.1 is not true if we remove the condition that the floodlights are of size at most π!

4. A polygon P is called co-circular if all the vertices of P lie on a circle. Show that the two-floodlight illumination problem for cocircular polygons can be solved in linear time.

Chapter 6

Moving guards

In the previous chapters, we studied problems daling with stationary guards, i.e. guards stationed at some fixed location. In this chapter we will study guarding problems in which our guards are represented by points that are allowed to move within a polygon.

We start with the following problem introduced by Suzuki and Yamashita, called by T. Shermer *The Hunter's Problem*. In this problem, we have a set of hunters and a prey that are allowed to move freely within a polygon. The hunters move with a bounded velocity, and the prey moves with unbounded velocity.

Our objective is to decide if there is a strategy that the hunters can follow to catch the prey. The prey is considered caught if it comes within sight of a hunter. A polygon P is k-searchable if k hunters are always sufficient to catch any prey in P. It follows easily from Theorem 2.1.1 that every polygon is $\lfloor \frac{n}{3} \rfloor$-searchable. This bound however is not tight, Urrutia proved:

Theorem 6.0.1 *$O(\ln n)$ hunters are always sufficient, and occasionally necessary to catch a prey in any polygon with n vertices.*

Proof: Let $f(n)$ be the minimum number of hunters needed to catch a prey in an n vertex polygon P. We now show that $f(n) \leq f(\frac{2n}{3}) + 1$. This will prove our result. We start by first finding a diagonal joining two vertices of P that cuts it into two subpolygons P_1 and P_2, each of size at most $\frac{2n}{3}$. It is well known that such a diagonal always exists. Station a hunter at an endpoint of this diagonal. This will ensure that no prey can go from P_1 to P_2. Next scan P_1 first, and then P_2. This can be accomplished with at most $f(\frac{2n}{3})$ hunters. Our result now follows. A family of polygons that require $O(\ln n)$ hunters can be obtained from a binary search tree T as follows: first draw T on the plane without crossing edges, and then substitute a sequence of three narrow corridors for each edge of T as in Figure 6.1. ∎

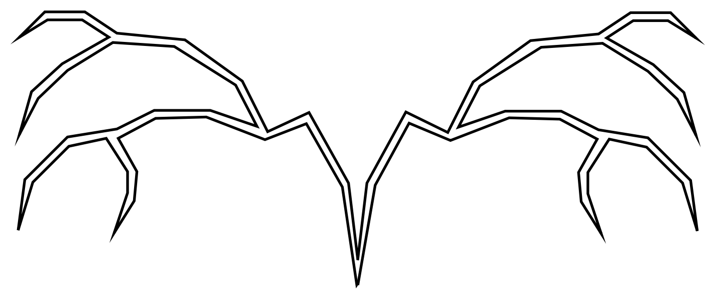

Figure 6.1: A binary polygon.

The problem of deciding if a polygon is k-searchable seems to be difficult, even for small values of k. Suzuki and Yamashita studied the problem of deciding if a simple polygon is 1-searchable. They were able to give some sufficient or necessary conditions for a polygon to be 1-searchable, but failed to provide a full characterization of these polygons. Three points x_1, x_2, and x_3 of P are called an asteroidal triple if they are such that the shortest path between x_i and x_j is not visible from any point visible from x_k; $i, j, k \in \{1, 2, 3\}$, $i \neq j \neq k$. For example in the polygon in Figure 6.2, x_1, x_2, and x_3 form an asteroidal triple. Moreover, it is easy to verify that this polygon is not 1-searchable. Suzuki and Yamashita [123] proved:

Theorem 6.0.2 *A polygon P that contains an asteroidal triple is not* 1-*searchable.*

The proof of this result is left as an exercise.

6.0.1 Lazy guards

The following problem involving moving guards was introduced by Colley, Meijer and Rappaport:

Problem 6.0.1 The lazy guard problem: *Given a simple polygon P, choose the minimum number of stations (points) in P such that a moving guard that visits all the stations guards P, i.e. while traveling to visit all the stations, every point in P will, at some point in time, be visible to our guard.*

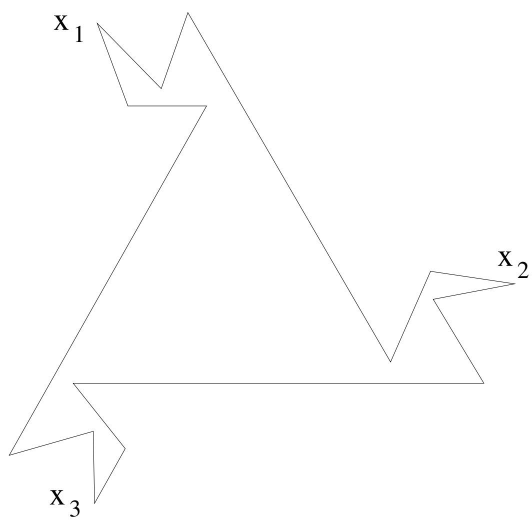

Figure 6.2: A 2-searchable polygon.

Since our guard wants to minimize the distance walked, we will assume that while traveling from one station to the next, he will choose the shortest path connecting them. The motivation for introducing stations is that if we want to make sure a guard patrols a polygon, the easiest way to do so is to place a few *check-in* stations such that if the guard visits them, P is guarded. In Figure 6.3 we show a polygon that needs three stations.

We now prove:

Theorem 6.0.3 *The lazy guard problem can be solved in linear time.*

We begin by proving the following result:

Lemma 6.0.1 *Let S be a set of k stations such that if a guard visits them in a given order, it guards P. Then if we visit the elements of S in any order, P is also guarded.*

Proof: Suppose that P is guarded while visiting the elements of S in a specific order. Let p and q be two elements of S that are visited in consecutive order, and r any point in P visible from some point x in the shortest path from p to q. Take the longest line segment L containing r and x contained in P. This will cut P into two subpolygons, one containing p

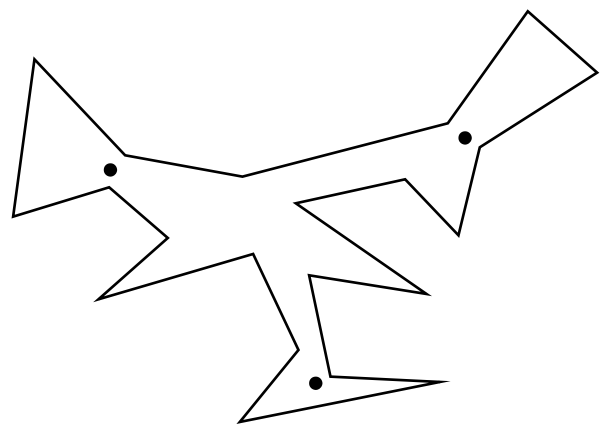

Figure 6.3: Three stations suffice to guard this polygon.

and the other containing q. Then if we visit the elements of S in any other order, say first p and then q, at some point we must cross L, and thus r is guarded. ∎

It is easy to see now that we can without loss of generality assume that the stations are visited starting and ending at the same station along the shortest circuit connecting them. This circuit defines a perhaps degenerate polygon P_S that may consist of the union of several simple polygons and line segments connecting them. We refer to this as *the guards polygon.*

We now prove:

Lemma 6.0.2 *If a polygon is guardable using k stations, it can also be guarded with k stations located on the boundary of P.*

Proof: The case $k = 2$ is easy. Let us assume then that $k \geq 3$, and suppose that we have a set S of k stations such that when we visit them, we guard P. Let s be a station in S not located on the boundary of P. Suppose that s is adjacent to two vertices s_1 and s_2 in P_S. Consider the bisector of the external angle of P_S defined by s_1, s, and s_2, and extend it until it hits the boundary of P at a point s'. Let $S' = S - \{s\} \cup \{s'\}$. We now prove that if we visit all the stations in S' we guard P. We notice first that $P_{S'}$ contains P_S. It now follows that if a point is visible from some point in P_S it is also visible from a point in $P_{S'}$. To finish our proof, we continue replacing the elements of S until all stations are on the boundary of P. ∎

Consider a reflex vertex v of P, and let v^- and v^+ be the clockwise and counter-clockwise vertices of P adjacent ot v. Let $r(v^+)$ and $r(v^-)$ be the points at which the rays starting at v^+ and v^- which pass through v hit the boundary of P; see Figure 6.4. Let $C(v, r^-)$ be the segment of the boundary of P traversed when moving from v to $r(v^+)$ in the clockwise direction along the boundary of P. Similarly we define $C(v^-, v)$.

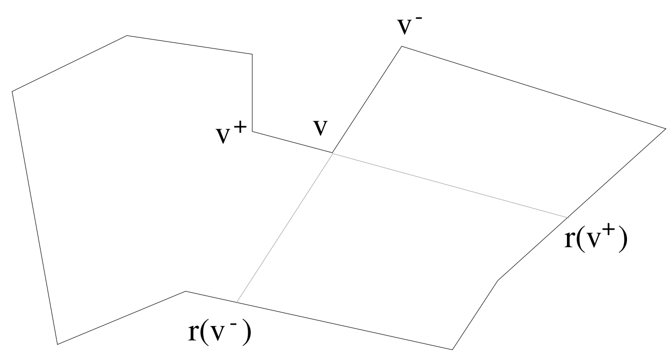

Figure 6.4: Finding $C(v, v^+)$ and $C(v^-, v)$.

We now prove:

Theorem 6.0.4 *A set of stations on the boundary of P suffices to lazy guard it if for every reflex vertex v of P, $C(v, v^+)$ and $C(v^-, v)$ contain an element of S.*

Proof: It is clear that if for some reflex vertex v of P no element of S is located on $C(v, v^+)$ then v^+ is not visible from P_S, which is a contradiction. Similarly for $C(v^-, v)$. The converse is obvious. ■

From Theorem 6.0.4 we can obtain the following algorithm to solve the lazy guard problem:

Algorithm 6.0.1

1. *For each reflex vertex of P compute $C(v, v^+)$ and $C(v^-, v)$.*

2. *Determine the smallest set of points S on the boundary of P such that each $C(v, v^+)$ and $C(v^-, v)$ contains at least a point in S.*

Step 1 of our algorithm can be done in linear time; see [71]. Step 2 can also be solved in linear time by reducing it to the problem of finding a minimum vertex cover of a circular arc graph, which is known to be solvable in linear time [85], as follows:

Map the boundary of P to a circle C. Each $C(v, v^+)$ and $C(v^-, v)$ is in turn mapped to an arc of our circle. Our problem is now reduced to that of finding a minimum number of points on C that covers the arcs defined by $C(v, v^+)$ and $C(v^-, v)$; v a reflex vertex of P.

Thus we have proved:

Theorem 6.0.5 *An optimal placing of stations to lazy guard a polygon can be found in linear time.*

For polygons with holes, the situation changes drastically. The order in which a set of stations is visited becomes relevant, and finding a minimum set of stations to lazy guard a polygon becomes NP-complete.

We now show that lazy guarding a polygon with holes is NP-complete. We use the following problem that is known to be NP-complete:

Problem 6.0.2 k-vertex cover of planar 3-regular graphs

INSTANCE: *A planar graph G with maximum degree 3, and an integer k.*

QUESTION: *Is there a set of vertices with k elements such that every edge of G has an endpoint in S?*

We now prove that the following problem is also NP-complete:

Problem 6.0.3 k-guarding of polygons with holes

INSTANCE: *A polygon P possible with holes.*

QUESTION: *Is there a set of points in p of cardinality k that guards P?*

Lemma 6.0.3 **k-guarding of polygons with holes** *is NP-complete.*

Proof: Take a planar graph G with maximum degree 3, and a plane embedding of it in which its edges are represented by straight lines. We can assume that no two edges of G are collinear. Slightly extend the line segment representing each edge in both sides by a small amount, and substitute this line segment by a thin rectangle $r(e)$ with its longer side parallel to e. Finally

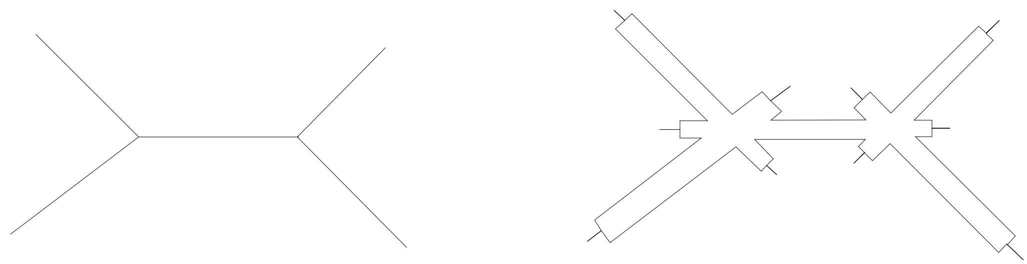

Figure 6.5: Substituing edges in a planar graph.

add two small spikes aligned with e and of width ϵ. The objective of these spikes is to ensure that any point that sees all of them belongs to $r(e)$ or to any of its the two spikes; see Figure 6.5. Let P be the union of all $r(e)$ and their spikes.

It is now easy to verify that a k-vertex cover of G generates a set of points that guard the resulting polygon. On the other hand, let S be a set of k points that guards P. If a point s in S guards points in more than one rectangle $r(e)$, these rectangles correspond to edges in G with a common end-vertex v. Move s to the location of v. It now follows that the vertices in which we place a guard correspond to a vertex cover of G. ∎

We now show:

Theorem 6.0.6 *Lazy guarding a polygon with k stations is NP-complete.*

Proof: Take a planar graph with maximum degree 3, and an embedding of it in which no edge is horizontal or vertical. Consider a set of $n+1$ horizontal line segments and two vertical ones such that the endpoints of each horizontal line segment are on the vertical segments, and each horizontal line segment, except for the bottom one, contain a vertex of G. Obtain from G the polygon P as in the previous lemma, and substitute each horizontal line segment by a narrow rectangle with a *hook* at each end, and the vertical lines by two narrow rectangles as shown in Figure 6.6. Let Q be the union of these polygons. We now prove that Q can be lazy guarded with $2n+k+2$ stations iff G has a vertex covering of size k.

To cover the spikes of the horizontal rectangles, we need a station in each of them, i.e. we need $2n+2$ stations. Suppose then that G has a vertex cover of size k. Then by placing k stations at the positions corresponding to the vertices of the vertex cover of G, we get a set of $2n+2+k$ stations that lazy guards Q.

Suppose now that Q is $2n+2+k$ lazy guardable, and that S is a set of $2n+2+k$ stations that lazy guard P. To guard the spikes of the horizontal rectangles of Q, we need $2n+2$ stations. Thus the remaining k stations must be used to ensure that P is lazy guarded. Notice first that if a corridor $r(e)$ of Q is such that no station is located in $r(e)$ or its spikes, then due to the way we constructed Q, we can find a path connecting all the stations of S that does not cover $r(e)$, which is a contradiction. As in the previous lemma, we can now move the k stations in S not used to guard hooks to locations corresponding to vertices of G to obtain a k vertex cover of G. It now follows that deciding if Q is lazy guardable with $2n+2+k$ stations is equivalent to deciding if P is k guardable, which is NP-complete. ∎

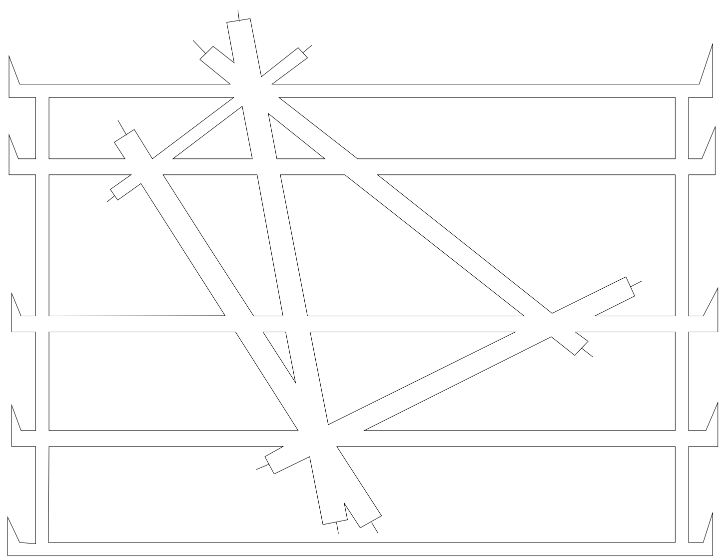

Figure 6.6: Constructing Q.

6.1 Exercises

1. Prove Theorem 6.0.2.
2. (Research problem) Characterize i-searchable polygons.

Chapter 7

Miscellaneous

In this chapter, we study variations of guarding or illumination problems studied by several authors.

7.1 Guarding triangulations of point sets

Let P_n be a set of n points on the plane. A *triangulation* T of P_n is a partitioning of the convex hull of P_n into a set of triangles with disjoint interiors. The vertices of these triangles are elements of P_n, and each element of P_n is a vertex of at least one triangle in T. The edges of T are the edges of its triangles.

A set S of edges of T guards it if every triangle of T has a vertex that is the end-vertex of an edge in S. A set H of vertices of T guards it if every face of T has a vertex in H. We will prove the following results, proved by Everett and Rivera-Campo [53], and Bose, Shermer, Toussaint and Zhu [16] respectively:

Theorem 7.1.1 *Any triangulation of a set of n points can be guarded with at most $\lfloor \frac{n}{3} \rfloor$ edges.*

Theorem 7.1.2 *Any triangulation of a point set can be guarded with at most $\lfloor \frac{n}{2} \rfloor$ vertices. This bound is tight.*

We notice first that the upper bound for Theorem 7.1.2 follows easily from the Four Color Theorem; simply take a four-coloring of the vertices of our triangulation, and by placing a guard at all the vertices of the two smaller chromatic classes our triangulation is guarded. In fact the first proof of Theorem 7.1.1 also used the Four Color Theorem. We now present a proof of both of these results that does not use the Four Color Theorem. This proof is due to Bose, Kirkpatrick and Li [15]. Some preliminary results will be needed. We first observe that our problems of guarding triangulations of point sets can be reduced to that of guarding maximal planar graphs. If the convex hull of our point set is a triangle, then any triangulation of our point set yields a maximal planar graph, otherwise by adding extra edges to the external face of the graph defined by our triangulation, we obtain a maximal planar graph.

A graph is called k-regular if all of its vertices have degree k. We start by proving the following well known result in Graph Theory:

Theorem 7.1.3 *Every 2-connected planar 3-regular graph has a perfect matching.*

Proof: To prove our result, we will show that any 3-regular 2-connected graph G satisfies Tutte's Theorem 3.1.1. Let S be a set of vertices of G and H an odd component of $G - S$. Since G is 3-regular and H has an odd number of vertices, the number of edges joining vertices of H to vertices in S has to be odd. However since G is 2-connected, there are at least 3 edges connecting S to H. It now follows that $Odd(G - S) \leq \mid S \mid$. ■

A 2-coloring of a graph is called K_3-free if it contains no triangle that is monochromatic. We now prove:

Theorem 7.1.4 *Every planar graph admits a K_3-free 2-coloring.*

Proof: We can assume that G is a maximal planar graph, otherwise insert extra edges until it becomes a maximal planar graph. Let G^* be the dual graph of G, i.e. for each face of G, insert a vertex in G^*. Two vertices in G^* are adjacent if their corresponding triangular faces share an edge of G. Notice that there is a one-to-one correspondence between the edges of G and those of G^*. Since G is maximal, G^* is a 3-regular and 2-connected graph, and thus by Theorem 7.1.3 it has a perfect matching M. Delete from G all the edges corresponding the the edges of G^* in M. It is now easy to see that all the faces of the remaining graph, call it H, are quadrilaterals, and H is therefore bipartite, i.e. the vertices of H can be 2-colored in such a way that no adjacent vertices receive the same color. It now follows that this 2-coloring of H induces a K_3-free coloring of G; see Figure 7.1. ■

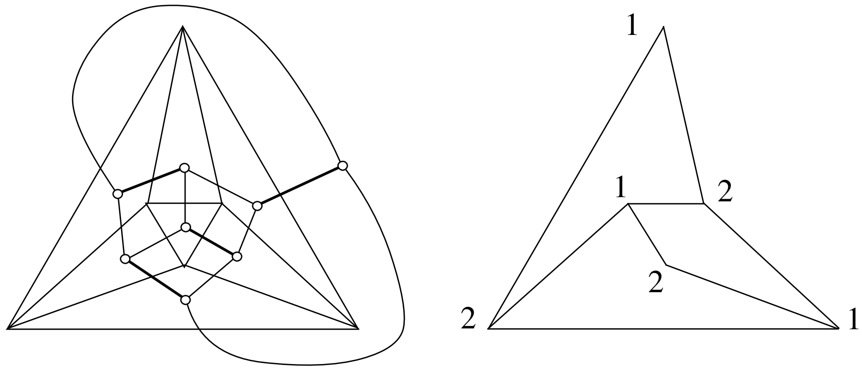

Figure 7.1: Obtaining a K_3-free coloring of a maximal planar graph.

We can now prove Theorem 7.1.2:

Proof: Take any K_3 2-coloring of T. Clearly each chromatic class in any such coloring guards T. Place a guard at each vertex in the smallest chromatic class, and our result follows. To show that our bound is tight, we generate a family of triangulations that require $\lfloor \frac{n}{2} \rfloor$ vertex guards as follows: Let S_1 be the seven vertex triangulation shown in Figure 7.2. $S_2, \ldots, S_n$ are now obtained by recursively replacing triangle x, y, z by a copy of S_1. We leave it to the reader to verify that S_k has $4k+3$ vertices, and that it requires $2k+1$ guards; see Excercise 1. ∎

The following will be needed to prove Theorem 7.1.1. Let G_1 and G_2 be the subgraphs of G induced by the chromatic classes of our 2-coloring of G, and let M_1 and M_2 be maximal matchings of G_1 and G_2 respectively.

Let us now choose two different sets S_1 and S_2 of edges of G that guard G as follows:

S_1 contains all the edges of M_1 as well as an extra edge incident to each unmatched vertex of G_1. S_2 is chosen in a similar way.

Next we show:

Lemma 7.1.1 *$M_1 \cup M_2$ guard G.*

Proof: Suppose that there is a triangle F of G that is not guarded by $M_1 \cup M_2$. Then none of the vertices of this face is incident to any edge in M_1 or M_2. Since we have a 2-coloring of G, two of the vertices of F receive the same color, say color 1. This contradicts the assumption that M_1 is maximal. ∎

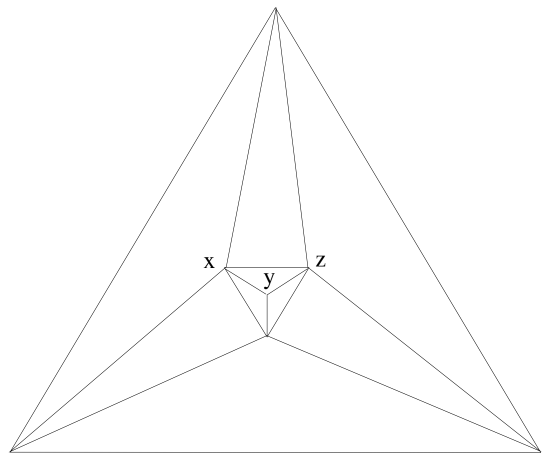

Figure 7.2: S_1 needs 3 guards to be guarded.

We are now ready to prove our main result:

Proof of Theorem 7.1.1: Notice that S_1, S_2 and $M_1 \cup M_2$ guard G. If $M_1 \cup M_2$ has at most $\lfloor \frac{n}{3} \rfloor$ edges, we are done. Suppose then that $\mid M_1 \cup M_2 \mid$ is greater that $\lfloor \frac{n}{3} \rfloor$. Notice that:

$\mid S_1 \mid = \mid V(G_1) \mid - \mid M_1 \mid$

and

$\mid S_2 \mid = \mid V(G_2) \mid - \mid M_2 \mid$

Moreover $\mid V(G_1) \mid + \mid V(G_2) \mid = \mid V(G) \mid = n$.

Then $\mid S_1 \mid + \mid S_2 \mid = \mid V(G_1) \mid + \mid V(G_2) \mid - (\mid M_1 \mid + \mid M_2 \mid)$

$\leq n - \lfloor \frac{n}{3} \rfloor \leq \lceil \frac{2n}{3} \rceil$.

It now follows that one of S_1 or S_2 has at most $\lfloor \frac{n}{3} \rfloor$elements. ∎

An example of a triangulation with $13n - 12$ vertices that needs $4n - 8$ edge guards can be obtained as follows: Consider the graph H shown in Figure 7.3. It is clear that for this triangulation we need two edges. Let T be any maximal planar graph with n vertices. By Euler's formula, T has exactly $2n - 4$ triangular faces. Place a copy of H in the middle of each face of T and retriangulate the resulting graph. The total number of vertices of the resulting graph is $n + 6(2n - 4)$. Since each copy of H needs two edge guards, and no two copies of H can share an edge guard, we need at least $2(2n - 4)$ edge guards.

A plane graph H is a planar graph G together with an embedding of it on the plane. An edge e and a face F of H are called incident if at least

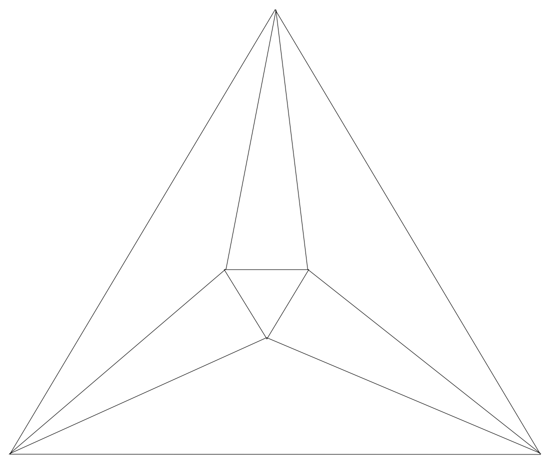

Figure 7.3: To guard H we need 2 edge guards.

one vertex of e lies on the boundary of F. A set S of edges of H guards it if each face of H is incident to an edge of S. Notice that the proof of the previous theorem does not apply to plane graphs; see exercise.

In a similar way, but now using the Four Color Theorem, we have [53]:

Theorem 7.1.5 *Every plane graph H on n vertices can be guarded with $\lfloor \frac{2n}{5} \rfloor$ edges.*

Proof: Triangulate H and four-color the vertices of H. Let V_1, V_2, V_3 and V_4 be the chromatic clases of our four coloring of H. For $1 \leq i < j \leq 4$ let $M_{i,j}$ be a maximal matching of the subgraph of H induced by $V_i \cup V_j$. Let $A_{i,j}$ be a set of edges defined as follows: $A_{i,j}$ contains all the edges of $M_{i,j}$ plus choose an extra edge for each vertex of $V_i \cup V_j$ not covered by $M_{i,j}$. Since each face of H must contain three vertices of different colors, it follows that each $A_{i,j}$ guards H. Moreover notice that $M_{1,2} \cup M_{1,3} \cup M_{2,4} \cup M_{3,4}$ is also a guarding set of H. However:

$$\mid A_{1,2} \mid + \mid A_{1,3} \mid + \mid A_{2,4} \mid + \mid A_{3,4} \mid + \mid M_{1,2} \cup M_{1,3} \cup M_{2,4} \cup M_{3,4} \mid = 2n.$$

It follows that at least one of these sets has at most $\lfloor \frac{2n}{5} \rfloor$ edges. ∎

From an algorithmic point of view, the proofs presented here for Theorems 7.1.1 and 7.1.2 yield algorithms with complexity $O(n^{\frac{3n}{2}})$. To prove this, we notice that calculating the dual graph of a triangulation can be done in linear time, finding its matching in $O(n^{\frac{3n}{2}})$; see [93]. Finding M_1, M_2 again takes $O(n^{\frac{3n}{2}})$ time, and calculating and S_1 and S_2 can be done in linear time. Thus we have:

Theorem 7.1.6 *Finding an edge and a vertex guard sets of sises* $\lfloor \frac{n}{3} \rfloor$ *and* $\lfloor \frac{n}{2} \rfloor$ *respectively can be done in* $O(n^{\frac{3n}{2}})$

An interesting open problem here is to find a proof of Theorem 7.1.5 that does not depend on the Four Color Theorem. It is easy to see that a proof based on K_3-free 2-colorings fails since $M_1 \cup M_2$ does not necessarily guard a plane graph. On the other hand if all the faces of a plane graph are odd, i.e. have an odd number of edges on their boundaries, the elements of $M_1 \cup M_2$ guard it. The proof of this is left as an excercise.

7.2 Exercises

1. Verify that $S_1, \dots S_k, \dots$ as defined in Theorem 7.1.2 require $2K + 1$ guards.

2. Prove that there are 2-connected plane graphs for which $\lfloor \frac{n}{3} \rfloor$ edges are necessary to guard them.

3. Prove that any plane graph such that all its faces are odd can be guarded with $\lceil \frac{n}{3} \rfloor$ edges.

Bibliography

[1] J. Abello, V. Estivill-Castro, T. Shermer and J. Urrutia. *Illumination of orthogonal polygons with orthogonal floodlights.* To appear in International Journal on Computational Geometry.

[2] M. Abellanas, J García-López, and F. Hurtado. *Consecutive Guards* Proc. of the Sixth Canadian Conference on Computational Geometry, pp. 393-398 (1994).

[3] A. Aggarwal, S. Ghosh, and R. Shyamasundar. *Computational complexity of restricted polygon decomposition.* In Computational Morphology, G. Toussaint, ed, North Holland, 1-11 (1988).

[4] A. Aggarwal, L.J. Guibas, J. Saxon and P.W. Shor, emph A linear time algorithm for computing the Voronoi diagram of a convex polygon, Discrete and Computational Geometry **4** (1989) 591-604 .

[5] A. Aggarwal. *The Art Gallery Theorem: its variations, applications and algorithmic aspects.* Ph.D. Thesis, Johns Hopkins Univ. 1984.

[6] D. Avis, and H. ElGindy. *A combinatorial approach to polygon similarity.* IEEE Transactions on Information Theory, **IT-29**, pp. 148-150, (1983).

[7] D. Avis, and G. Toussaint. *An optimum algorithm to determine the visibility of a polygon from an edge.* IEEE Trans. Comput. **C-30**, pp. 910-914, (1981).

[8] D. Avis and G. Toussaint. *An efficient algorithm to decompose a polygon into star-shaped pieces.* Pattern Recognition **13**, pp. 295-298, (1981).

[9] J. Bagga, L. Gewali, and D. Glasser. *The complexity of illuminating polygons by alpha-floodlights* Proceedings of the Eight Canadian Conference on Computational Geometry, pp. 337-342, (1996).

[10] P. Belleville. *On restricted boundary covers and convex three-covers.* Technical Report, Simon Fraser University. In preparation.

[11] P. Belleville, P. Bose, J. Czyzowicz, J. Urrutia, and J. Zaks. *K-guarding polygons on the plane.* Proceedings of the Sixth Canadian Conference on Computational Geometry, pp. 381-386, (1994).

[12] I. Bjorling-Sachs, and D. Souvaine. *A tight bound for guarding polygons with holes.* Rutgers University Technical Report LCSR-TR-165, (1991).

[13] I. Bjorling-Sachs, and D. Souvaine. *An efficient algorithm for guard placement in polygons with holes.* Discrete and Computational Geometry **13**, pp. 77-109, (1995).

[14] P. Bose, L. Guibas, A. Lubiw, M. Overmars, D. Souvaine, J. Urrutia. *The floodlight illumination problem.* To appear in International J. on Computational Geometry.

[15] P. Bose, D. Kirkpatrick, and Z. Li. *Efficient algorithms for guarding or illuminating the surface of a polyhedral terrain.* To appear Proc. of the Eight Canadian Conference on Computational Geometry, (1996).

[16] P. Bose, T. Shermer, G. Toussaint, and B. Zhu. *Guarding polyhedral terrains.* To appear in Computational Geometry, Theory and Applications.

[17] H. Bunting, and D. Larman. Personal Communication (1992).

[18] S. Carlsson, and H. Jonsson. *Guarding a treasury.* Proc. Fifth Canadian Conference on Computational Geometry, pp. 85-90, (1993).

[19] S. Carlsson, B.J. Nilsson, and S. Ntafos. *Optimum guard covers and m-watchmen routes for restricted polygons.* Proc. 2nd Workshop on Algorithms and Data Structures, Lecture Notes in Computer Science **519**, pp. 367-378, (1991).

[20] B. Chazelle. *Triangulating a simple polygon in linear time.* Discrete and Computational Geometry **6**, pp. 485-524, (1991).

[21] B. Chazelle, and D. Dobkin. *Optimal convex decompositions.* In Computational Geometry, G. Toussaint, ed., North Holland, (1985).

[22] D.Z. Chen. *An optimal parallel algorithm for detecting weak visibility of a simple polygon.* Proc. Eight Annual Computational Geometry Symposium, pp. 63-72, (1992).

[23] D.Z. Chen, V. Estivill-Castro, and J. Urrutia. *Optimal guarding of polygons and monotone chains.* Proc. of the Seventh Canadian Conference on Computational Geometry, pp. 133-138, (1995).

[24] W.P. Chin and S. Ntafos. *Optimum watchman routes.* Information Processing Letters, **28**, pp. 39-44, (1988).

[25] W.P. Chin and S. Ntafos. *Watchman routes in simple polygons.* Discrete and Computational Geometry, **6**, pp. 9-31, (1991).

[26] W.P. Chin, and S. Ntafos. *The zoo-keeper route problem.* Information Sciences **63**, pp 245-259, (1992).

[27] N. Chiba, and T. Nishizeki *The Hamiltonian cycle problem is linear-time solvable for 4-connected planar graphs.* J. Algorithms **10**, pp. 187-212, (1989).

[28] V. Chvátal. *A combinatorial theorem in plane geometry.* Journal of Combinatorial Theory Ser B **18**, pp. 39-41, (1975).

[29] P. Colley, H Meijer, and D. Rappaport. *Motivating lazy guards.* Proc. Seventh Canadian Conference on Computational Geometry, pp. 121-126, (1995).

[30] H. Conn, and J. O'Rourke. *Some restricted rectangle covering problems.* Proceedings of the 1987 Allerton Conference.

[31] C. Coullard, B. Gamble, W. Lenhart, W. Pulleyblank, and G Toussaint. *On illuminating a set of disks.* Manuscript, (1989).

[32] J. Culberson, and R. Reckhow. *Covering polygons is hard.* Proc. 29th Symposium on Foundations of Computer Science, pp. 601-611, (1988).

[33] J. Culberson, and R. Reckhow. *A unified approach to orthogonal polygon covering problems via Dent Diagrams.* University of Alberta TR. 89-6.

[34] G. Csizmadia, and G. Tóth. *Note on an art gallery problem.* Manuscript, (1996).

[35] J. Czyzowicz, P. Egyed, H. Everett, W. Lenhart, K. Lyons, D. Rappaport, T. Shermer, D. Souvaine, G. Toussaint, J. Urrutia and S. Whitesides. *The aquarium keeper's problem.* Proc. ACM-SIAM Symposium on Data Structures and Algorithms, pp. 459-464, (1991).

[36] J. Czyzowicz, I. Rival, and J. Urrutia. *Galleries, light matchings and visibility graphs.* WADS 1989, Ottawa, in: F. Dehne, J.-R. Sack and N. Santoro, (eds.), Lecture Notes in Computer Science 382, pp. 316-324, (1989).

[37] J. Czyzowicz, E. Rivera-Campo, and J. Urrutia. *Illuminating rectangles and triangles in the plane.* Journal of Combinatorial Theory Ser. B, pp. 1-17, (1993).

[38] J. Czyzowicz, E. Rivera-Campo, J. Urrutia and J. Zaks. *Illuminating lines and circles on the plane.* University of Ottawa, Computer Science Department TR-89-49, (1989).

[39] J. Czyzowicz, E. Rivera-Campo, J. Urrutia and J. Zaks. *Protecting convex sets.* Graphs and Combinatorics **10**, pp. 311-312, (1994).

[40] J. Czyzowicz E. Rivera-Campo, J. Urrutia and J. Zaks. *On illuminating line segments in the plane.* Discrete Mathematics **137**, pp. 147-153, (1995).

[41] J. Czyzowicz, E. Rivera-Campo, N. Santoro, J. Urrutia and J. Zaks. *Guarding rectangular art galleries.* Discrete Mathematics **50**, pp. 149-157, (1994).

[42] J. Czyzowicz E. Rivera-Campo, and J. Urrutia. *Optimal floodlight illumination of stages.* Proc. 5th Canadian Conference on Computational Geometry, pp. 393-398, (1993).

[43] J. Czyzowicz, B. Gaujal, E. Rivera-Campo, J. Urrutia and J. Zaks. *Illuminating high-dimensional convex sets.* Geometriae Dedicata, **56**, pp. 115-120 (1995).

[44] J. Czyzowicz, and J. Urrutia. Personal communication, (1996).

[45] L.L. Deneen, and S. Joshi. *Treasures in an art gallery.* Proc. Fourth Canadian Conference on Computational Geometry, pp. 17-22, (1992).

[46] H. Edelsbrunner, J. O'Rourke, and E. Welzl. *Stationing guards in rectilinear art galleries* Computer Vision, Graphics, and Image Processing **27**, pp. 167-176, (1984).

[47] H. ElGindy, and D. Avis. *A linear time algorithm for computing the visibility polygon from a point.* Journal of Algorithms, **2**, pp. 186-197, (1981).

[48] P. Erdös, and G. Szekeres. *A combinatorial problem in Geometry.* Compositio Mathematica **2**, pp. 463-470, (1935).

[49] V. Estivill-Castro, J. O'Rourke, J. Urrutia and D. Xu. *Illumination of polygons with vertex floodlights.* Information Processing Letters, **56**, Vol. 1, pp. 62-73, (1995).

[50] V. Estivill-Castro, and J. Urrutia. *Optimal floodlight illumination of orthogonal polygons.* Proc. Sixth Canadian Conference on Computational Geometry, pp. 81-86, (1994).

[51] V. Estivill-Castro, and J. Urrutia. *The two floodlight illumination problem.* Proc. Fourth Workshop on Algorithms and Data Structures, Lecture Notes in Computer Science 955, Springer Verlag, pp. 62-73, (1995).

[52] H. Everett, K. Lyons, B. Reed, and D. Souvaine. *Illuminating squares on a traversal.* Proc. Third Canadian Conference on Computational Geometry, pp. 118-121, (1991).

[53] H. Everett, and E. Rivera-Campo. $\lfloor \frac{n}{3} \rfloor$ *edge guards are sufficient to guard a polyhedral terrain.* To appear in Computational Geometry, Theory and Applications.

[54] H. Everett, and G. Toussaint. *On illuminating isothetic rectangles in the plane.* Manuscript, (1990).

[55] L. Fejes Tóth. *Illumination of convex disks.* Acta Math. Acad. Sci. Hung. **29**, pp. 355-360, (1977).

[56] S. Fisk. *A short proof of Chvátal's watchman theorem.* Journal of Combinatorial Theory Ser B **24**, p. 374, (1978).

[57] Z. Füredi, and D. Kleitman. *The prison yard problem.* Combinatorica **14**, pp. 287-300, (1994).

[58] J. García-López. *Problemas Algorítmico-Combinatorios de Visibilidad.* Ph.D. Thesis, Universidad Politécnica de Madrid, (1995).

[59] R. Garey and D. Johnson. *Computers and Intractability.* W.H. Freeman and Company, New York, 1979.

[60] M. Garey, D.S. Johnson, F.P. Preparata, and R.E. Tarjan. *Triangulating a simple polygon.* Information Processing Letters **7**, pp. 175-179, (1978).

[61] L. Gewali. *Placing guards inside orthogonal polygons.* Information Sciences **88**, pp. 1-14, (1996).

[62] L. Gewali, and R. Lombardo. *Watchman routes for a pair of convex polygons.* Lecture Notes in Pure and Applied Mathematics, **144**, (1993).

[63] L. Gewali, A.C. Meng, J.S.B. Mitchell, and S. Ntafos, *Path planning in 0/1/inf weighted regions with applications.* ORSA Journal on Computing, Vol. 2, **3**, pp. 253-272, (1990).

[64] L.P. Gewali and S. Ntafos. *Covering grids and orthogonal polygons with periscope guards.* Computational Geometry: Theory and Applications, **2**, pp. 309-334, (1993).

[65] L.P. Gewali and S. Ntafos. *Watchman routes in the presence of a pair of convex polygons.* In Proc. Seventh Canadian Conference on Computational Geometry, pp. 127-132, (1995).

[66] L.P. Gewali and I. Stojmenovic. *Computing external watchman routes on PRAM, BSR and interconnection network models of parallel computation.* Parallel Processing Letters, **4**, 182, pp. 83-93, (1994).

[67] S. Ghosh. *Approximation algorithms for Art Gallery Problems.* Proceedings of the Canadian Information Processing Society Congress, (1987).

[68] E. Györy. *A short proof of the rectilinear art gallery theorem.* SIAM J. Alg. Disc. Meth. **7** pp. 452-454, (1986).

[69] E. Györy, F. Hoffmann, K. Kriegel, and T. Shermer. *Generalized guarding and partitioning for rectilinear polygons.* Proceedings of the Sixth Canadian Conference on Computational Geometry, pp. 302-307, (1994).

[70] H.-D. Hecker, and D. Herwing. *Some NP-hard polygon cover problems.* J. Inf. Process Cybern. **25**, pp. 101-108, (1989).

[71] P.J. Heffernan. *An optimal algorithm for the two-guard problem.* International J. of Computational Geometry and Applications **6**, pp. 14-44, (1996).

[72] J. Hershberger. *Finding the visibility graph of a polygon in time proportional to its size.* Proceedings of the Third ACM Symposiuum on Computational Geometry, pp. 11-20, (1987).

[73] J. Hershberger, and J. Snoeyink. *An efficient solution to the Zookeeper's Problem.* Proc. Sixth Canadian Conference on Computational Geometry, pp. 104-109, (1994).

[74] G. Hernandez-Peñalver. *Controlling guards.* Proc. of the Sixth Canadian Conference on Computational Geometry, pp. 387-392, (1994).

[75] F. Hoffmann. *On the rectilinear Art Gallery Problem.* in Proc. ICALP, Springer Verlag Lecture Notes in Computer Science **90** pp. 717-728, (1990).

[76] F. Hoffman, M. Kaufman, and K. Kriegel. *The art gallery problem for polygons with holes.* Proceedings of the 32nd. Symposium on Foundations of Computer Science, pp. 39-48 (1991).

[77] F. Hoffman, and K. Kriegel. *A graph coloring result and its consequences for the polygon guarding problems.* Technical report B 93-08, Freie Universität Berlin, (1993).

[78] F. Hurtado, O. Serra, and J. Urrutia. *Hiding people in arrangements of segments.* To appear in Discrete Applied Mathematics.

[79] C. Icking, and R. Klein. *The two guards problem.* Inter. J. on Computational Geometry and Applications **2**, pp. 257-285, (1992).

[80] G.F. Jennings, and W. Lenhart. *An art gallery theorem for line segments in the plane.* Pattern Recognition Letters **14**, pp. 727-732 (1993).

[81] J.M. Keil. *Decomposing a polygon into simpler components.* SIAM Journal on Computing, Vol. 14, No. 4, pp. 799-817, (1985).

[82] J. Kahn, M. Klawe, and D. Kleitman. *Traditional galleries require fewer watchmen.* SIAM J. Algebraic and Discrete Methods **4**, pp. 194-206, (1983).

[83] V. Klee. *Is every polygonal region illuminated from some point?.* Amer. Math. Monthly **76**, p. 180, (1969).

[84] P.P. Kumar, and C.E.V. Madhavan. *Shortest Watchman tours in weak visibility polygons.* Proc. Fifth Canadian Conference on Computational Geometry, pp. 91-96, (1995).

[85] C.C. Lee, and D.T. Lee. *On a circle-cover minimization problem* Information Processing Letters, **18**, pp. 109-115, (1984).

[86] D.T. Lee and A.K. Lin. *Computational complexity of art gallery problems.* IEEE Trans. Info. Th. **IT-32**, pp. 415-421, (1979).

[87] D.T. Lee, and F.P. Preparata. *An optimal algorithm for finding the kernel of a polygon.* Journal of the ACM, **26**, pp. 276-282, (1979).

[88] B-C. Liaw, N.F. Huang, and R.C.T. Lee. *The minimum cooperative guards problem on k-spiral polygons.* Proc. of the Fifth Canadian Conference on Computational Geometry, pp. 97-102, (1993).

[89] A. Lingas. *The power of non rectilinear holes.* Proc. 9th Colloquium on Automata, Languages, and Programming, Springer-Verlag, pp. 369-383, (1982).

[90] R. Liu, and S. Ntafos. *On partitioning rectilinear polygons into star-shaped polygons.* Algorithmica, **6**, pp. 771-800, (1991).

[91] A. Lubiw. *Decomposing polygonal regions into convex quadrilaterals.* Proc. First ACM Computational Geometry Symposium, Baltimore, pp. 97-106, (1985).

[92] C.S. Mata and J. Mitchell. *Approximation algorithms for geometric tours and network design problems.* Proc. 11th ACM Symposium on Computational Geometry, pp. 360-370, (1995).

[93] S. Micali, and V. Vazirani *An $O(\sqrt{V}E)$ algorithm for finding a matching in general graphs.* Proc. 21st Annual IEEE Symposium on Foundations of Computer Science, pp. 17-27, 1980.

[94] B.J. Nilsson. *Guarding art galleries, methods for mobile guards.* Ph.D. Thesis, Department of Computer Science, Lund University (1994).

[95] B.J. Nilsson, and D. Wood. *Optimum watchmen routes in spiral polygons.* Proceedings of the Second Canadian Conference in Computational Geometry, pp. 269-272, (1990).

[96] T. Nishizeki. *Precise proofs of some lemmas in the paper "Lower bounds on the cardinality of the maximum matchings of planar graphs,"* Carnegie-Mellon tech. report (1977).

[97] T. Nishizeki and I. Baybars. *Lower bounds on cardinality of maximum matchings of planar graphs.* Discrete Mathematics **28**, pp. 255-267, (1979).

[98] S. Ntafos. *The robber problem.* Information Processing Letters **34**, pp. 59-63, (1990).

[99] S. Ntafos. *Watchman routes under limited visibility.* Computational Geometry, Theory and Applications, Vol. 1, No. 3, pp. 149-170, (1992).

[100] S. Ntafos, and L. Gewali. *External watchman routes.* The Visual Computer **10**, pp. 474-483, (1994).

[101] S. Ntafos, and M. Tsoukalas. *Optimal placement of guards.* Proceedings of the Third Canadian Conference on Computational Geometry, (1991) pp. 122-125. To appear in Information Sciences.

[102] J. O'Rourke. *An alternate proof of the rectitlinear art gallery theorem.* J. of Geometry **21**, pp. 118-130, (1983).

[103] J. O'Rourke. *Galleries need fewer mobile guards: a variation to Chvátal's Theorem.* Geometriae Dedicata **14**, pp. 273-283, (1983).

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

[105] J. O'Rourke, and D. Xu. *Illumination of polygons with* 90^o *vertex lights.* in Snapshots in Computational Geometry, Vol. 3, pp. 109-123, (1992).

[106] J. O'Rourke, T. Shermer, and I. Streinu. *Illuminating convex polygons with vertex floodlights.* in Proceedings of the Seventh Canadian Conference on Computational Geometry, pp. 151-156, (1995).

[107] G. Rote. *Personal Communication.* (1993).

[108] J.R. Sack. *Rectilinear computational geometry.* Carleton University, TR-SCS-54-1984.

[109] J.R. Sack, and G. Toussaint. *A linear time algorithm for decomposing rectilinear star-shaped polygons into convex quadrilaterals.* Proceedings of the 19th Allerton Conference on Communication, Control and Computing, Monticello, pp. 21-30, (1981).

[110] J.R. Sack, and G. Toussaint. *Guard placement in rectilinear polygons.* Computational morphology, G.T. Toussaint (editor) Elsevier Science Publishers, pp. 153-175, (1988).

[111] F. Santos. *Personal Communication.* (1995).

[112] D. Schuchardt, and H. Hecker. *Two NP-hard problems for ortho-polygons.* Math. Logiv Quart. **41**, pp. 261-267, (1995).

[113] T. Shermer. *Hiding people in polygons.* Computing **42**, pp. 109-131, (1989).

[114] T. Shermer. *Visibility properties of polygons.* Ph.D. Thesis, McGill University, Montreal, Canada, (1989).

[115] T. Shermer. *Recent results in Art Galleries.* Proceedings of the IEEE, pp. 1384-1399, (1992).

[116] T. Shermer. *On recognizing unions of two convex polygons.* Pattern Recognition Letters, **14**, pp. 737-745, (1993).

[117] T. Shermer. *A tight bound on the combinatorial edge guard problem.* In preparation.

[118] I. Streinu. *Positive and negative results in Computational Geometry.* Ph.D. Thesis, Rutgers University, (1994).

[119] W. Steiger, and I. Streinu. *Positive and negative results on the Floodlight Illumination Problem.* Proc. Sixth Canadian Conference on Computational Geometry, pp. 87-92, (1994).

[120] W. Steiger, and I. Streinu. *Illumination by floodlights.* Manuscript, (1996).

[121] K. Sugihara, I. Suzuki, and M. Yamashita. *The searchlight scheduling problem.* SIAM J. Computing, **19**, pp. 1024-1040, (1990).

[122] K. Sugihara, I. Suzuki, and M. Yamashita. *The searchlight scheduling problem.* Technical Report, Department of Electrical Engineering and Computer Science University of Milwaukee, October 1988.

[123] I. Suzuki, and M. Yamashita. *Searching for a mobile intruder in a polygonal region.* SIAM J. Computing, **21**, pp. 863-888, (1992).

[124] X.-H. Tan, T. Hirata, and Y. Inagaki. *An incremental algorithm for constructing the shortest watchman route.* Proc. ISAC'91, Lecture Notes in Computer Science, **557**, pp. 163-175, (1991).

[125] X.-H. Tan, T. Hirata, and Y. Inagaki. *Constructing the shortest watchman routes by divide and conquer.* Proc. 4th International Symposium on algorithms and Computation, Lecture Notes in Computer Science, **762**, pp. 68-77, (1993).

[126] R. Tarjan and C. van Wyk. *An* $O(n\,log \log n)$ *time algorithm for triangulating simple polygons.* SIAM Journal of Computing **17**, pp. 143-178, (1988).

[127] G.T. Tokarsky. *Polygonal rooms not illuminable from every point.* The American Mathematical Monthly, Vol. 102, No. 10, pp. 867-879, (1995).

[128] W.T. Tutte. *The factorization of a planar graph.* J. London Math. **22**, pp. 107-111, (1947).

[129] W.T. Tutte. *A theorem on planar graphs.*Trans. Amer. Math. Soc. **82**, pp. 99-116, (1956).

[130] J. Urrutia. *Iluminando polígonos con reflectores.* Actas VI Encuentro de Geometría Computacional, Barcelona, July 1995, F. Hurtado, M. Noy, and J. Trias (eds.), Universitat Politécnica de Catalunya, pp. 59-73, (1995).

[131] J. Urrutia. *Art Gallery and illumination problems.* To appear in *Handbook on Computational Geometry.* J.R. Sack and J. Urrutia, eds. Elsevier Science Publishers (1997).

[132] J. Urrutia, and J. Zaks. *Illuminating convex sets.* University of Ottawa TR-89-31, (1989).

[133] S. Viswanathan. *The edge guard problem for spiral polygons.* Proc. Fifth Canadian Conference on Computational Geometry, pp. 103-108, (1993).

[134] S.N. Yiu. *A generalized fortress problem using k-consecutive vertex guards.* Proc. Seventh Canadian Conference on Computational Geometry, pp. 139-144, (1995).

[135] S.M. Yiu, and A. Choi. *Edge guards on a fortress.* Proceedings of the Sixth Canadian Conference on Computational Geometry, pp. 296-301, (1994).