VISIBILITY PROPERTIES OF POLYGONS

by
Thomas C. Shermer

School of Computer Science
McGill University
Montréal

June 1989

A THESIS SUBMITTED TO THE FACULTY OF GRADUATE STUDIES AND RESEARCH
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Abstract

In this thesis, we establish tight bounds on the maximum size of maximum hidden sets, minimum guard sets, and minimum partitions and covers of polygons, using link-visibility. These results unify and generalize the guard set results of Chvátal and O'Rourke. Our method also provides tight bounds on independent and dominating sets in triangulation graphs, and almost-tight bounds on the size of hidden sets, guard sets, covers, and partitions of polygon exteriors. In addition, we prove that, using link-visibility, the optimization problems of finding maximum hidden sets, minimum guard sets, or minimum covers are NP-hard.

Link-visibility is an extended notion of visibility arising from robotics and motion planning problems. Hidden sets are sets of points in a polygon such that no two points of the set are visible, and guard sets are sets such that each point of the polygon is visible to some point in the guard set. Both maximum hidden set sizes and minimum guard set sizes can be used as polygon shape complexity measures.

Résumé

Grâce à la vue-liée, nous bornons de manière optimal la grandeur maximum des ensembles cloisonés maximums, des ensembles-sentinelles minimums et celles de partitions et de couvertures minimums de polygones.

Ces résultats unifient et généralisent ceux de Chvàtal et d'O'Rourke sur les ensembles-sentinelles.

En outre, des bornes optimums sur la grandeur d'ensembles indépendants et dominants, et des bornes quasi-optimums sur la grandeur d'ensembles cloisonés, d'ensembles-sentinelles, de couvertures et de partitions de la région externe d'un polygone sont dérivés par le même biais. De plus, nous démontrons à l'aide de la vue-liée que les problèmes d'optimisation de la recherche d'ensembles cloisonés maximums, d'ensembles sentinelles minimum et de couvertures minimums sont NP-durs.

La vue-liée est une généralisation de la notion de vue, provenant de la robotique et de la planification de trajectoire.

N'importe quels deux points membres d'un ensemble cloisoné sont mutuellement non-visibles, alors que n'importe quel point du polygone est vu par au moins un point d'un ensemble-sentinelle. Les concepts de grandeur maximum d'ensembles cloisonés et de grandeur minimum d'ensembles-sentinelles peuvent servir de mesure de la complexité de la forme d'un polygone.

Originality

This entire thesis, with the exception of many definitions in the introduction, and the entire review chapter, should be considered an original contribution to knowledge.

In the invention and preparation of this material, the assistance that I recieved from others was limited to clarifying discussions, proofreading, and the translation of the abstract.

Acknowledgements

I have many people to thank for their contributions to my being in a position to present this thesis. First, I would like to thank my parents, for raisin' me up real good.

I am endebted to my teachers David Meier, Fred Giffin, and Lou Montrose, from whom I not only learned mathematics but the love of mathematics. Their dedication to their profession and zeal for the subject has been a constant inspiration.

I must also thank Joseph O'Rourke, my mentor at Johns Hopkins, for introducing me to algorithms and computational geometry. It was his interest in and enthusiasm for the art gallery problems that rekindled my flagging interest in mathematics and geometry.

At the New York Institute of Technology, I was greatly influenced by Robert McDermott and Pat Hanrahan. Robert impressed upon me the importance of having a *lifework*, and Pat provided me with a sound example of a true scholar and good friend.

At McGill, I am deeply indebted to David Avis and Godfried Toussaint, for consistently good advice, good instruction, and inspirational ideas. Others here have also contributed to my education, most memorably: Sue Whitesides, for teaching me complexity theory; Rafe Wenger, for getting me a good office and being a good office-mate; Hossam ElGindy, Mike Houle, and David Samuel, for an uncountable number of stimulating discussions, and incessant proofreading; and finally, Luc Devroye, for being funny. Collectively, the computional geometry community members at McGill also warrant mention, for their tolerance with me and my work, and their expert criticisms and commentary.

Naji Mouawad gets a thousand praises and his own paragraph in these acknowledgements for translating the abstract.

I would also like to thank the Friends of McGill, for providing the fellowship that made my stay in Montréal possible.

My Guru, Swami Chidvilasananda, deserves more mention here than is possible. Without her steadying, uplifting, and loving presence in my life, the work that I have done here in Montréal would have taken at least four times as long, and would not have been nearly as pleasant.

Above all, I would like to thank my family: Lakshmi, Sylvain, and Toby. My wife Lakshmi has done an excellent job of taking care of the administrative details of my studies, and has provided me with all forms of necessary support, including, but not limited to, good meals, friendship, and lots of laughs.

Dedication

This thesis is dedicated to
and offered at the lotus feet of
my guru,
Swami Gurumayi Chidvilasananda,
by whose joy all are made joyous.

Contents

List of Figures

Chapter 1

Introduction

In this chapter we introduce the major ideas that this thesis is concerned with.

1.1 Notation and Terminology

We use the operator $\setminus$ to denote the usual set-theoretic difference.

If the intersection of two sets has zero measure (no area), then we shall say that the sets are nonoverlapping. Note that two sets may be intersecting but not overlapping.

We assume that the reader is familiar with elementary graph theory, and we use the usual graph theory notation (see, e.g., [H69] or [BM76]). We let the class of graphs include the *null graph* (the graph on zero vertices), and the class of trees include the *null tree*.

The vertex set and diameter of a graph G are denoted by $vert(G)$ and $diam(G)$, respectively. $d(x,y)$ is the (graph-theoretic) distance between vertices x and y. By the *induced graph difference* $G \ominus S$, where S is a subgraph of G, we mean the subgraph of G that is induced by $vert(G) \setminus vert(S)$.

In a rooted tree R, $depth(R)$ denotes the depth of R, and $st(R,w)$ is the subtree of R rooted at w. We define the *remaining tree* $rt(R,U)$ to be $R \ominus \bigcup_{u \in U} st(R,u)$.

A *contraction* of two vertices v and w in a graph G replaces G by a graph G^* which is G with v and w (and their edges) removed, and a new vertex v^* added, which is adjacent to all of the vertices that v and w were adjacent to. If H is some subgraph

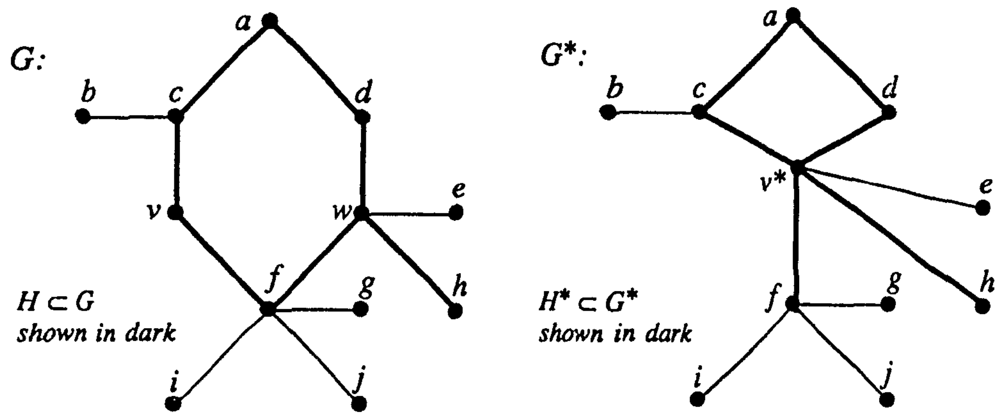

Figure 1.1: Contractions on graphs and subgraphs

of G, then H^* is the subgraph of G^* which results from contracting v and w (if they both exist) in H (see figure 1.1). For brevity, we refer to $(st(R,w)))^*$ as $st^*(R,w)$, and similarly define $rt^*(R,U)$.

We use the prefix D_k- to indicate that a graph-theoretic object has diameter at most k. For instance, a D_k-tree is a tree of diameter at most k.

1.2 Visibility

The major category that the work in this thesis falls under is called *visibility*; this is a well-studied notion in mathematics and computer science. Given some set of points R in E^d, we say that two points $x, y \in R$ are *visible* if the closed line segment from x to y lies entirely in R (see figure 1.2). Visibility is therefore a symmetric and reflexive relation on the points of R. Two points which are visible are said to *see* each other.

Given this definition of visibility, we can define two types of point sets (regions) R: *convex* regions, for which $\forall x, y \in R$, x sees y, and *star-shaped* regions, for which $\exists x \in R\ \forall y \in R$, x sees y. Examples of these types of regions are shown in figure 1.3.

We can extend the concept of visibility from points to regions: we say that a region $U \subset R$ is visible from a region $T \subset R$ if $\forall u \in U, \exists t \in T$ such that t sees u. In

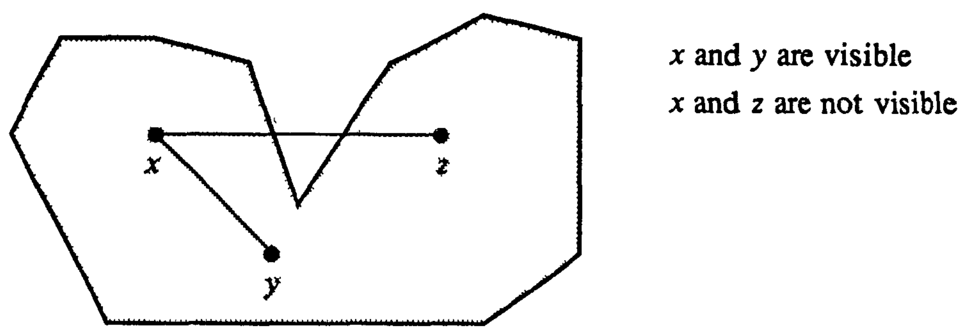

Figure 1.2: Illustrating visibility

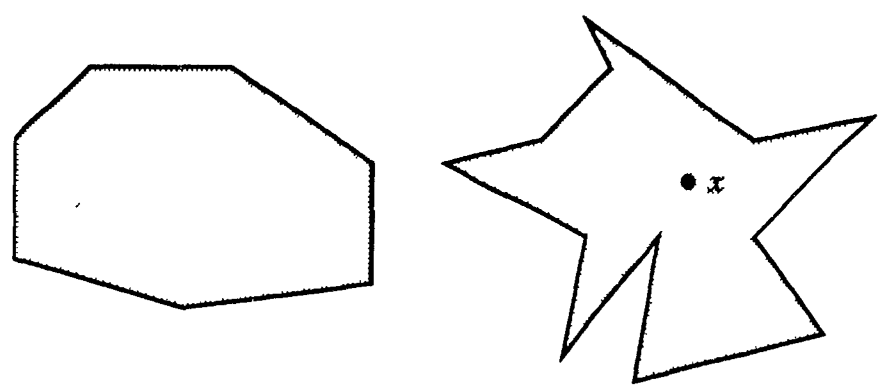

Figure 1.3: Convex and star-shaped regions

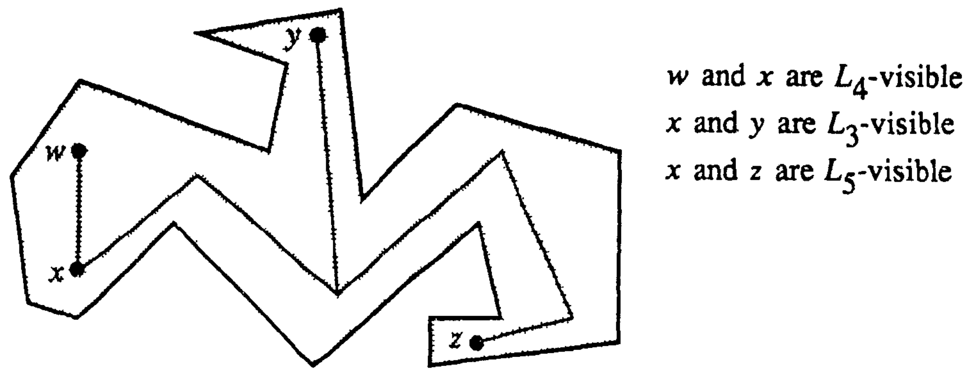

Figure 1.4: L_j-visible

such an instance we will also say that T *sees* U. This is the notion of *weak* visibility introduced in [AT81b]. We note that (weak) visibility is not a symmetric relation.

1.3 Generalized Visibility

One of the major contributions of this thesis is extension of known results about visibility to a more general visibility, which is called *link-j-visibility.* We will use the notation L_j as shorthand for "link-j".

We say that two points $x, y \in R$ are *L_j-visible* if there is some path $P \subseteq R$ joining x and y which consists of j or fewer straight line segments ("links"). Some examples are shown in figure 1.4. The smallest j such that x and y are L_j-visible is called the *link-distance* between x and y [S86a]. We note that the usual notion of visibility introduced above is exactly L_1-visibility.

We can define *L_j-convex* and *L_j-star-shaped* regions in a manner analogous to our definitions of convex and star-shaped regions: L_j-convex regions are those for which $\forall x, y \in R$, x and y are L_j-visible, and L_j-star-shaped regions are those for which $\exists x \in R\ \forall y \in R$, x and y are L_j-visible. Examples of these types of regions are shown in figure 1.5.

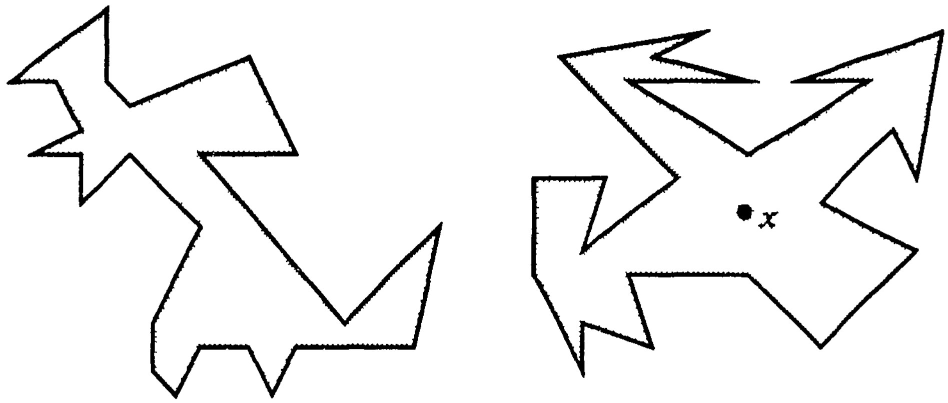

Figure 1.5: A L_3-convex region and a L_2-star-shaped region

We define L_j-visibility for regions in the same manner as we defined L_1-visibility for regions.

1.4 Polygon

The major type of regions that we will deal with in this thesis is *simple, singly-connected polygons.* A polygon is a finite figure in the plane that is bounded by a finite number of straight line segments. A singly-connected polygon is bounded by n points $v_1, v_2, \ldots, v_n$ (called *vertices*) and the n line segments $[v_1, v_2], [v_2, v_3], \ldots, [v_{n-1}, v_n]$, and $[v_n, v_1]$ (called *edges*). Such a polygon is called *simple* if no point of the plane belongs to more than two edges of the polygon and the only points which belong to precisely two edges are the vertices.

A simple polygon divides the plane into a bounded region, called the *interior*, and an unbounded region, called the *exterior.* We henceforth will use the term *polygon* to refer to the boundary and interior of a simple, singly-connected polygon. Several regions which we do not consider polygons are shown in figure 1.6.

A *chord* of a polygon is a line segment which is contained in the polygon, and has both endpoints on the boundary of the polygon. A *diagonal* of a polygon is a chord

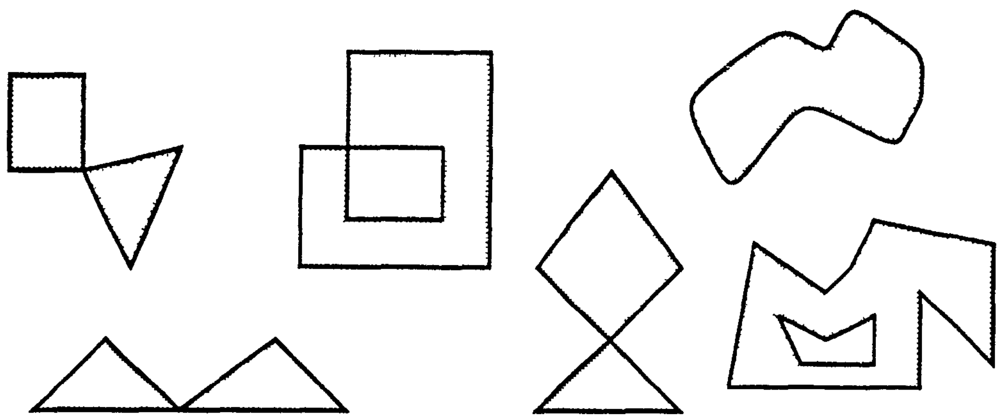

Figure 1.6: Regions which are not polygons

with both endpoints on vertices of the polygon. An *ear* of a polygon is a vertex whose adjacent vertices can be connected by a diagonal.

A polygon Q is called a *subpolygon* of a polygon P if every point of Q is also in P.

1.5 Polygon Covering

The polygon covering problem is: given a polygon P and some property π which is true of some polygons, find a minimum-cardinality set $Q = \{Q_i\}$ of subpolygons of P, each with property π, such that their union is P (i. e., $\bigcup_i Q_i = P$). Typical properties that are used as π are star-shaped, monotone, or convex. The collection Q is called a *cover* for P. A polygon with a cover by convex sets is illustrated in figure 1.7.

Given a specific property π, we call the polygon covering problem that uses property π the π *cover problem.* For instance, the polygon cover problem with π being "convex" is known as the *convex cover problem.*

A *partition* is a cover where no two of the covering objects overlap. A minimum partition of a polygon will therefore always have the same number of or more pieces than a cover of that polygon. However, we will show that over all polygons with n

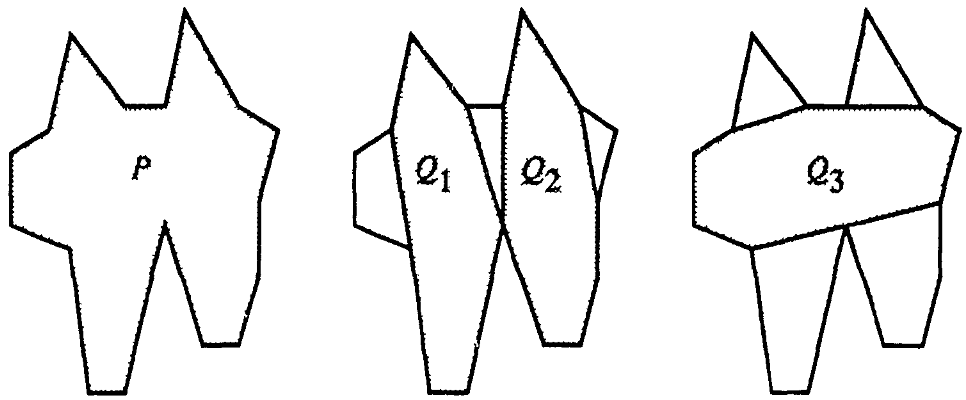

Figure 1.7: A polygon with a convex cover

vertices, the maximum size of a minimum π cover is the same as the maximum size of a minimum π partition.

It is well known that the maximum size of a convex cover of a polygon is $n-2$, and the same bound holds for convex partition. For star-shaped cover and partition, the bound is $\lfloor n/3 \rfloor$ [C75][F78]. The problem of computing the minimum convex cover of an input polygon has recently been shown to be NP-hard [CR88] [S88b]. The same is true of minimum star-shaped covers [LL86] [A84]. However, the minimum convex *partition* problem can be solved in polynomial time [CD85].

1.6 The Art Gallery Problem

A *polygon-guard class* is a collection of regions in a polygon. A *guard class* $\mathcal{C}$ is a function which for every polygon P maps to a polygon-guard class $\mathcal{C}(P)$. Typical guard classes are the ones where $\mathcal{C}(P)$ consists of the vertices of P, the points of P, all line segments in P, all diagonals of P, or all star-shaped regions in P.

A guard class $\mathcal{C}$ is said to *contain* another guard class $\mathcal{D}$ (written $\mathcal{C} \supseteq \mathcal{D}$ or $\mathcal{D} \subseteq \mathcal{C}$) if, for every polygon P, every member of $\mathcal{D}(P)$ is a subset of some member of $\mathcal{C}(P)$. For example, if we let $\mathcal{C}(P)$ be all of the chords of P, $\mathcal{P}(P)$ be all of the points of P, and $\mathcal{V}(P)$ be all of the vertices of P, then $\mathcal{V} \subseteq \mathcal{P} \subseteq \mathcal{C}$, as each vertex of a polygon is a point of the polygon, and each point in the polygon is contained in some chord.

The *art gallery problem* (see [O87]) is: given a polygon P and a guard class $\mathcal{C}$, what is the minimum cardinality g of a subset G of $\mathcal{C}(P)$ such that every point in P is in the visibility polygon of at least one element of G? For L_1-visibility, this problem has been shown to be NP-complete for vertex, point, and edge guards [LL86][A84].

However, exact bounds on $g(n)$, the maximum size of g over all polygons of n vertices, have been shown for vertex and point guards ($g(n) = \lfloor n/3 \rfloor$) [C75] [F78], and diagonal, line segment, and convex guards ($g(n) = \lfloor n/4 \rfloor$) [O83a]. For orthogonal polygons, bounds for vertex and point guards ($g(n) = \lfloor n/4 \rfloor$) [KKK83], and line segment and convex guards ($g(n) = \lfloor (3n+4)/16 \rfloor$) [A84] have also been shown.

Given a property π, we let $\mathcal{C}^{\pi}(P)$ be the collection of all subregions of P with property π. Then, the π cover problem can be viewed as the art gallery problem using $\mathcal{C}^{\pi}$ and L_0-visibility. Also, we can view the art gallery problem as a special case of the covering problem in that we are required to cover the polygon with visibility polygons.

In this thesis, we will be dealing mostly with the following two guard classes: $\mathcal{L}_k = \mathcal{C}^{L_{k-convex}}$, and $\mathcal{T}_k = \mathcal{C}^{\pi}$, where π is the property of being the vertices of a D_k-tree in some triangulation (of the polygon P).

1.7 Hidden Sets

A *hidden set* is a set of points in a polygon such that no two points in the set are visible to each other [S87]. A *hidden vertex set* is a hidden set which contains only vertices of the polygon. The maximum size of the hidden set (or hidden vertex set), over all polygons with n vertices, using L_j-visibility, is denoted $h_j(n)$. A (L_1-visibility) hidden set is shown in figure 1.8a.

Hidden sets are known in the mathematics literature as *visually independent sets* [KG70] [B76]. Another related concept from the math literature is that of k-convexity (or property P_k). A region is said to be k-convex if it has no hidden sets of size k or greater. Thus, convex sets are the same as 2-convex sets.

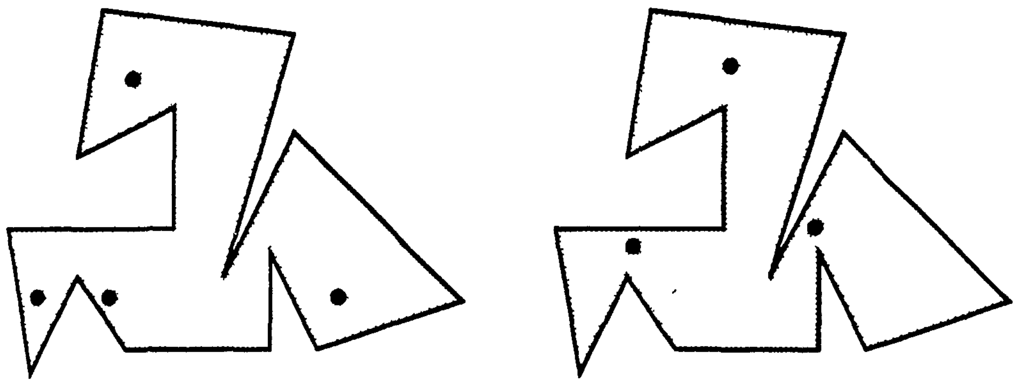

Figure 1.8: Hidden set and hidden guard set

A *hidden guard set* is a hidden set which is also a guard set. A hidden guard set is shown in figure 1.8b. Whereas every polygon admits a hidden guard set, we will see that not every polygon admits a *hidden vertex guard set*, which is a hidden vertex set that is also a guard set.

1.8 Point Visibility Graphs

We now introduce a structure which lends insight into the relationship between the guard set, hidden set, hidden guard set, and "link-distance" problems: the point visibility graph of a polygon.

Given a polygon P, we define the *point visibility graph* of P, denoted $PVG(P)$, as:

$$\begin{aligned} PVG(P) &= (V_P, E_P), \text{ where} \\ V_P &= \{p \mid p \in P\} \\ E_P &= \{[p,q] \mid p,q \in P \text{ and } p \text{ sees } q\} \end{aligned}$$

Note that this is an *infinite* graph, as the number of points in a polygon is infinite. (This graph may also be called a *continuous* graph, in the sense defined in [N-W73]). We hope that the reader will not mind the abuse of notation inherent in the above definition (we use p to refer to both a point in P and a vertex of $PVG(P)$, and similarly abuse $[p, q]$).

Without explicit reference to point visibility graphs, many of the properties of these graphs have been studied [BB64] [LPS87] [S86a]. For example, the distance of two vertices p and q of $PVG(P)$ is the link-distance of p and q in P. This means that any graph properties derived from distances have their link counterparts: the diameter of $PVG(P)$ is the link-diameter of P, the radius of $PVG(P)$ is the link-radius of P, the eccentricity of vertex p in $PVG(P)$ is the link-eccentricity of p in P (also known as the *covering radius*), and the center of $PVG(P)$ is the link-center of P. Also, a vertex-dominating set of $PVG(P)$ is a point guard set of P, an independent set of $PVG(P)$ is a hidden set of P, an independent vertex-dominating set of $PVG(P)$ is a hidden guard set of P, and a maximal clique of $PVG(P)$ is a maximal convex subset of P.

Another structure of interest is the *vertex visibility graph* of P, denoted $VVG(P)$, also known simply as the *visibility graph* of P. This is the subgraph of $PVG(P)$ which is induced by the vertices of P. Several papers have appeared on characterizing, recognizing, and computing visibility graphs [E85] [E89] [G86] [H87] [W85] [OW88] [KM88], and on applications of these graphs [AE83] [SH79]. A similar structure is the *edge visibility graph*, which has a vertex for every edge of the polygon, and an edge between two vertices if there are points on the corresponding edges which are visible. A survey of the use of these different types of visibility graphs can be found in [T88].

1.9 Triangulation

One of the major tools that we will use in this thesis is *polygon triangulation*. A polygon triangulation is a division of a polygon P into triangles such that there is no vertex in any triangle that is not a vertex of P. The edges of the triangles thus formed are either polygon edges or diagonals.

It is well known that every polygon can be triangulated (see, for instance [O87]), and some polygons may in fact have several triangulations. An example of a polygon with a triangulation is shown in figure 1.9a.

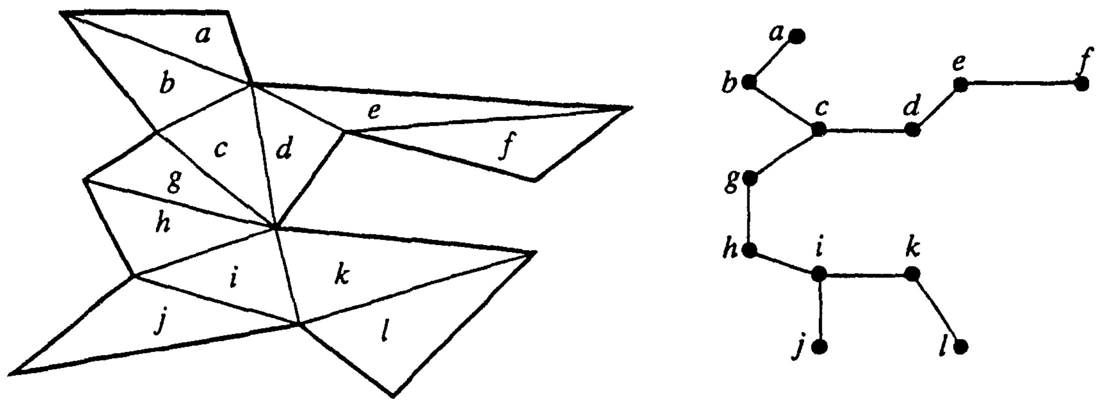

Figure 1.9: A triangulation and its dual tree

Given a polygon P with a triangulation, a *triangulation graph* is the plane graph formed by letting the vertices of the graph be the vertices of the polygon, and connecting two graph vertices if their corresponding polygon vertices share a triangulation diagonal or a polygon edge. The triangulation graph is the graph whose drawing results from drawing the polygon and the diagonals of a triangulation. For example, the drawing of the triangulation in figure 1.9a is also a drawing of its triangulation graph. The class of triangulation graphs is known to graph theorists as the class of *maximal outerplane graphs.*

Each triangulation graph has a unique Hamiltonian cycle, which corresponds to the edges of the polygon. We use the term *cycle edge* to refer to an edge in this hamiltonian cycle.

A triangulation graph is said to be *dominated* by a subset V' of its vertices if every triangle of the graph has at least one of its vertices in V'. We also say that a triangulation graph G is dominated by a subgraph (or a collection of subgraphs) of G if the vertices of the subgraph (or collection) dominates G. We will show that dominating a triangulation graph of a polygon is closely related to L_1-guarding that polygon.

The *dual tree* of a triangulation T is the graph which has one vertex for each triangle in T, and has an edge between two vertices if the two corresponding triangles share an edge. A sample dual tree is shown in figure 1.9b.

1.10 Organization of the Thesis

Throughout the text, we will use the variable j for the number of links in the visibility that we are using, and the variable k for the link-diameter of regions (generally subpolygons) under consideration.

The remainder of the thesis is organized as follows:

The following chapter (the second) is a review of the relevant mathematics and computer science literature.

The third chapter contains lower bounds on the maximum size of maximum hidden sets, minimum guard sets, and minimum polygon covers. Bounds are given for both polygon interiors and polygon exteriors.

The fourth chapter contains proofs of matching (for the interior) and almost-matching (for the exterior) upper bounds for the lower bounds presented in the third chapter.

In the fifth chapter, We show that most of the optimization problems associated with polygon covers, guard sets, and hidden sets are NP-hard.

The sixth and final chapter is the conclusion.

Chapter 2

Review

This thesis has its root in, and was inspired by, recent work in *art gallery* theorems. Art gallery theorems are so called because of the metaphor where one considers a polygon the floor plan of an art gallery in which all of the walls (and floor!) are covered with valuable artwork. One then asks the question: what is the fewest number of guards necessary to place in the art gallery such that there is no piece of artwork (point of the polygon) that is not seen by at least one guard?

This metaphor of polygon as room (or as art gallery) can be traced back at least 40 years. We find, for example, the following quote from [YB61]:

> Imagine a painting gallery consisting of several rooms connected with one another whose walls are completely hung with pictures. Krasnosel'skiĭ's Theorem states that if for each three paintings of the gallery there is a point from which all three can be seen, then there exists a point from which *all* the paintings of the gallery can be seen.

Art gallery problems, and visibility in general, have been studied by both mathematicians and computer scientists. Unless otherwise stated, the visibility used in the results discussed in this chapter is L_1-visibility.

2.1 Math Visibility

The recent mathematical interest in visibility was started by Valentine in his 1953 paper "Minimal Sets of Visibility" [V53]. In this paper, he characterizes those sets whose minimal *connected* guard sets are unique. In 1957, Valentine introduced a new generalization of convexity which inspired many mathematicians to consider visibility [V57]. Given our terminology, this generalization can be understood as follows: convex sets have a hidden set of size at most *one*; Valentine considered (and characterized) sets which have a hidden set of size at most *two*. These sets were further studied by Juul [J77]. Later investigators examined sets with hidden set of size at most m [KG70] [GK71] [B76] [BK76], and even further generalizations (which we will not detail here) [B73].

Horn and Valentine started mathematicians working on L_j-convex sets, in a paper which characterized planar L_2-convex sets [HV49]. Later work includes a paper by Breen about L_2-convex sets which are visible from a convex set [B77], a paper by Bruckner and Bruckner on the L_j-kernel (link center) of a set [BB64], a paper by Valentine on reflex points and L_j-convexity [V65b], and a paper by Sparks on intersections of maximal L_j-convex subsets of other sets [S70].

Covering has also been considered in the mathematics literature. There are many papers on sets which can be covered by two convex sets [SM63] [M66] [J77]. Also, covering and partitioning has been studied for sets with a bounded maximum hidden set size [B76] [BK76]. Other work has been done on covering with star-shaped sets [HK68] [KM66].

2.2 Art Gallery Results

It was a mathematician, Victor Klee, who finally got a computer scientist/graph theorist, Vašek Chvátal, interested in guarding problems. In 1973, Klee posed the art gallery question to Chvátal, who solved the problem, finding tight bounds on $g(n)$ for point and vertex guards. As a large portion of this thesis is devoted to generalizing this result, we review both the result and its proof.

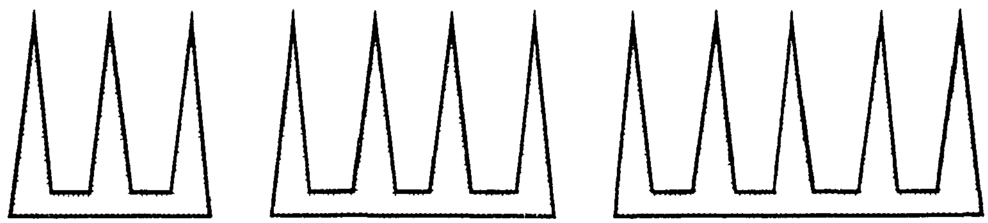

Figure 2.1: (Chvátal) Comb polygons

Theorem 2.1 (Chvátal 1975) *For any integer $n \geq 3$, $\lfloor n/3 \rfloor$ point or vertex guards are sometimes necessary and always sufficient to guard a polygon with n vertices.*

PROOF The "comb" polygons, as illustrated in figure 2.1, are polygons requiring $\lfloor n/3 \rfloor$ such guards; each upward spike on the comb requires its own guard. One can easily generalize the polygons shown to polygons of arbitrarily high n with $\lfloor n/3 \rfloor$ spikes. Thus, $\lfloor n/3 \rfloor$ point or vertex guards are necessary for some polygons.

To prove that this many are sufficient, we assume that we are given a polygon P with n vertices. First, triangulate P to get a triangulation graph T. Next, dominate T (by vertices) and place guards at each vertex in the dominating set. As each point of P is in some triangle of T, and each triangle has a vertex in the dominating set (and, therefore, in the guard set), each point will be visible to some guard. We now need only show that the triangulation graph T can be dominated by $\lfloor n/3 \rfloor$ vertices.

The proof is by induction. The induction basis is $n \leq 5$; any triangulation graph with 5 or fewer vertices (3 or fewer triangles) has one vertex which is incident on all triangles. The induction hypothesis is that $\lfloor n'/3 \rfloor$ vertices suffice to dominate any triangulation graph of $n' < n$ vertices. We wish to show that $\lfloor n/3 \rfloor$ suffices for $n > 5$.

First, find a diagonal $D = [a, b]$ that cuts the triangulation graph into two pieces, one of which has between 3 and 5 triangles, inclusive (such a diagonal always exists,

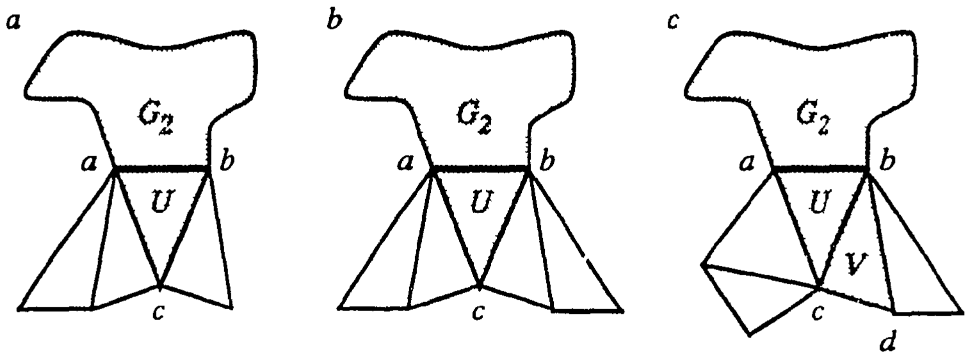

Figure 2.2: Cases in proof of Theorem 2.1

and we will prove a generalization of this statement later in the thesis). Let G_1 be the part of 3-5 triangles, and G_2 be the other part.

case 1: G_1 has all triangles incident on some vertex v.

In this case, we dominate G_2 by induction (with at most $\lfloor (n-3)/3 \rfloor = \lfloor n/3 \rfloor - 1$ vertices), and place v in the dominating set. This gives a total of at most $\lfloor n/3 \rfloor$ vertices.

case 2: G_1 has four triangles, as pictured in figure 2.2a.

Let U be the triangle of G_1 containing D, and c be the vertex of U not on D. Dominate $G_2 + U$ with at most $\lfloor n/3 \rfloor - 1$ vertices by induction. Either a, b, or c must be in the dominating set of G_2, else U is not dominated. If a is in the dominating set, then place b in the set as well. Similarly, if b or c is in the set, place a in the set. In either case we have added only one vertex to the dominating set, and G_1 is dominated. Hence, we have dominated G with at most $\lfloor n/3 \rfloor$ vertices.

case 3: G_1 has five triangles, as pictured in figure 2.2b.

Again, let $U = \Delta abc$ be the triangle of G_1 containing D, and dominate G_2+U with at most $\lfloor n/3 \rfloor - 1$ vertices by induction. If a is in the dominating set, place b in the set. If b is in the dominating set, place a in the set. If c is in the dominating set, remove it and place a and b in the set. In any case, we have added one vertex and dominated G_1. Also, removing c and replacing it with a and b does not affect any of the triangles in G_2. Therefore, we have dominated G with at most $\lfloor n/3 \rfloor$ vertices.

case 4: G_1 has five triangles, as pictured in figure 2.2c.

Let $U = \Delta abc$ be the triangle of G_1 containing D, $V = \Delta bcd$ be the adjacent triangle as pictured, and dominate G_2+U+V with at most $\lfloor n/3 \rfloor - 1$ vertices by induction. One of b, c, and d will be in the dominating set. If c is in the dominating set, then place b in the set. If b or d is in the set, place c in the set. In either case we have added only one vertex to the dominating set, and G_1 is dominated; we have again dominated G with at most $\lfloor n/3 \rfloor$ vertices.

As the above are the only possible cases, we have shown that $\lfloor n/3 \rfloor$ vertices suffice to dominate a triangulation graph, and hence $\lfloor n/3 \rfloor$ vertices suffice to guard a polygon. □

Fisk later found a more elegant proof of this theorem, by 3-coloring the triangulation graph, and placing guards on vertices which were colored with the least frequently used color [F78].

Later, Toussaint considered the problem of finding bounds on the number of guards when the guards are allowed to patrol fixed line segments or edges, and showed that $\lfloor n/4 \rfloor$ edge guards were sometimes necessary, and also conjectured the sufficiency of this number. O'Rourke proved sufficiency for line segments [O83a]; his proof is an extension of the method of Chvátal.

Theorem 2.2 (O'Rourke 1983) *For any integer $n \geq 4$, $\lfloor n/4 \rfloor$ vertex-pair, diagonal, or line segment guards are sometimes necessary and always sufficient to guard a polygon with n vertices.*

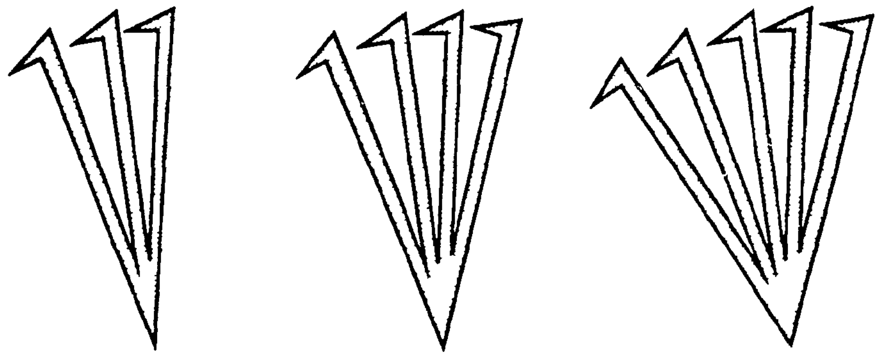

Figure 2.3: (Toussaint) Necessity for mobile guards

A *vertex-pair guard* is a pair of L_1-visible vertices.

PROOF The polygon class illustrated in figure 2.3 are polygons requiring $\lfloor n/4 \rfloor$ such guards; each arm on the polygon requires its own guard. One can easily generalize the pictured polygons to polygons of arbitrarily high n with $\lfloor n/4 \rfloor$ arms. Thus, $\lfloor n/4 \rfloor$ such guards are necessary for some polygons.

We only sketch the sufficiency proof. Assume that we are given a polygon P with n vertices. First, triangulate P to get a triangulation graph T. Next, dominate T by edges (more precisely, by units of two vertices connected by an edge), and place a guard at the geometric location for each unit in the dominating set. As each point of P is in some triangle of T, and each triangle has a vertex in a unit in the dominating set, each point will be visible to some guard. We now need only prove that the triangulation graph T can be dominated by $\lfloor n/4 \rfloor$ such units.

The proof is by induction. The induction basis is $n \leq 7$; any triangulation graph with 7 or fewer vertices (5 or fewer triangles) has one edge which is incident on all triangles (O'Rourke proves this by a lengthy case analysis). The induction hypothesis is that $\lfloor n'/3 \rfloor$ vertices suffice to dominate any triangulation graph of $n' < n$ vertices. We wish to show that $\lfloor n/3 \rfloor$ suffices for $n > 7$.

First, find a diagonal that cuts the triangulation graph into two pieces, one of which (called G_1) has between 4 and 7 triangles, inclusive; let G_2 be the other piece. Such a diagonal always exists.

If G_1 has 4 or 5 triangles, then induction can be applied to both G_1 and G_2, to get at most $1 + \lfloor (n-4)/4 \rfloor = \lfloor n/4 \rfloor$ guards.

If G_1 has 6 or 7 triangles, then the result is proved by a long case analysis based on the structure of G_1, similar to the analysis given in the proof of Theorem 2.1. □

2.3 Art Gallery Variants

Many variations on the art gallery problem have been investigated; O'Rourke has written a book which covers most of them [O87]. The most notable variations are those in which the class of regions investigated has been restricted, expanded, or changed. Typical work concentrates on star-shaped polygons, monotone polygons, spiral polygons, orthogonal polygons, polygons with holes, or polygon exteriors.

Star-shaped, monotone, and spiral polygons are restricted polygon classes which often arise in practice, and consideration of these classes has led to some interesting theorems (see [O87]). However, the study of visibility in these classes of polygons is not in the scope of this thesis.

Orthogonal polygons also often arise in practice. An orthogonal polygon is a polygon in which the edges alternate between horizontal and vertical. In [KKK83], Kahn, Klawe, and Kleitman proved that $\lfloor n/4 \rfloor$ vertex or point guards are necessary and sufficient for orthogonal polygons; in [A84], Aggarwal proved that $\lfloor (3n+4)/16 \rfloor$ line segment guards are necessary and sufficient. In this thesis, we do not consider orthogonal polygons, but we note that our method applies to orthogonal polygons when the diameter of the guarding object is even (such as 0 in the vertex guard case), but not when it is odd. Thus, we can generalize Kahn, Klawe, and Kleitman's result, but not Aggarwal's.

Another variant that has been considered are polygons with holes. In this problem, we attempt to guard regions which are polygons with subpolygons subtracted. We let h denote the number of holes, and consider point or vertex guards. Shermer has shown that there are polygons requiring $\lfloor (n+h)/3 \rfloor$, and O'Rourke proved that no polygon could require more than $\lfloor (n+2h)/3 \rfloor$. For $h = 1$, it has been shown that $\lfloor (n+1)/3 \rfloor$ suffice [S84] [S85]. Aggarwal, O'Rourke, and Shermer have also done work on orthogonal polygons with orthogonal holes; this and the other work on polygons with holes is summarized in [O87].

Work has also been done on guarding polygon exteriors. Here, O'Rourke and Wood have shown that $\lceil n/2 \rceil$ vertex guards are necessary and sufficient. Also, Aggarwal and O'Rourke proved that $\lceil n/3 \rceil$ point guards are necessary and sufficient, and Shermer has a simpler proof of this result showing that at most two of these point guards need to be located at points which are not vertices of the polygon (see [O87]). In this thesis we will not generalize these results, but we will generalize O'Rourke's slightly weaker result that $\lceil (n+1)/3 \rceil$ point guards are sufficient for a polygon exterior.

2.4 Computational Complexity

Lee and Lin have shown that determining the minimum number of vertex (or edge) guards necessary to guard a given simple polygon is NP-hard [LL86]; Aggarwal has generalized this proof to point guards (or star-shaped cover) [A84]. Also, Culberson and Reckhow have shown that determining the minimum number of convex subsets necessary to cover a polygon is NP-hard [CR88] (see [S88b] for an independent proof of this result). Previously, O'Rourke and Supowit had shown that these problems are NP-hard for polygons with holes [OS83].

In [S87], Shermer showed that many of the problems associated with hidden sets in polygons are difficult: computing the size of the maximum hidden set or hidden vertex set is NP-hard; computing the size of the minimum hidden guard set is NP-hard; determining if a polygon has a hidden vertex guard set is NP-complete; and computing the size of the minimum hidden vertex guard set is NP-hard, even if it is known that the polygon has a hidden vertex guard set.

2.5 Algorithms

A fair amount of work has been done on computing link-distance properties. Suri has given an algorithm for computing the link-distance between two points in a polygon; this algorithm runs in $O(n)$ time, given a triangulation of the polygon [S86a]. Suri has also given an $O(n \log n)$ algorithm for computing the link-diameter of a polygon (the smallest j such that the polygon is L_j-convex) [S86b]. The problems of computing the link-center and link-radius of a polygon were considered in [LPS87], where an $O(n^2)$ algorithm is given for both problems. This time was improved to $O(n \log n)$ (for both problems) while this thesis was in preparation [K89].

Although the minimum convex cover problem for a polygon is NP-hard, Chazelle and Dobkin have shown that the minimum convex *partition* problem for a polygon can be solved in $O(n^3)$ [CD85]. Shermer gives $O(n)$ algorithms for recognizing polygons which can be covered by two convex polygons and polygons which have a maximum hidden set of size two [S88c].

Several papers have appeared on the *guard placement problem*: find a set of guards, with the number of guards not exceeding the worst-case bound, for a given polygon. Avis and Toussaint first showed that the (point or) vertex guard placement problem for simple polygons can be solved in $O(n \log n)$ time [AT81a]. This can now be done in $O(n \log \log n)$ time using the trapezoidization/triangulation algorithm of Tarjan and Van Wyk [TV88], and imitating the art gallery proof of Fisk. Two papers exist which give $O(n \log \log n)$ algorithms for vertex guard placement in orthogonal polygons [EOW84] (as modified in [O87]) [ST88]. The quadrilateralization algorithm of Lubiw [L85] leads to an $O(n \log n)$ algorithm for this problem.

Chapter 3

Lower Bounds and Existence

3.1 Hidden Sets, Covering, and Guarding

We devote the first part of this section to the proof and corollaries of a lower-bound theorem for hidden sets in polygons; these results will show the close relationship between hiding, covering, and guarding. The section ends with some special-case bounds for L_1-visibility.

Theorem 3.1 *For any integers $j \geq 0$ and $n \geq j+1$, there exist polygons with n vertices that*

(a) have a L_j-hidden vertex set of size $\lfloor n/(j+1) \rfloor$, and

(b) require at least $\lfloor n/(j+1) \rfloor$ regions in any covering or partition by L_j-convex regions.

PROOF The polygon class illustrated in figure 3.1 consists of such polygons. Figure 3.1 shows representatives for each j, for j between 0 and 9, inclusive. The hidden set for the $j = 0$ example is the entire vertex set, and for the rest of the polygons, the hidden set is the set of vertices with acute angles at the end of the spiral "arms." Thus the hidden vertex set is of size $\lfloor n/(j+1) \rfloor$.

To get representatives for higher j, simply increase the number of turns on the spiral by an appropriate amount. To get representatives for other n, change the

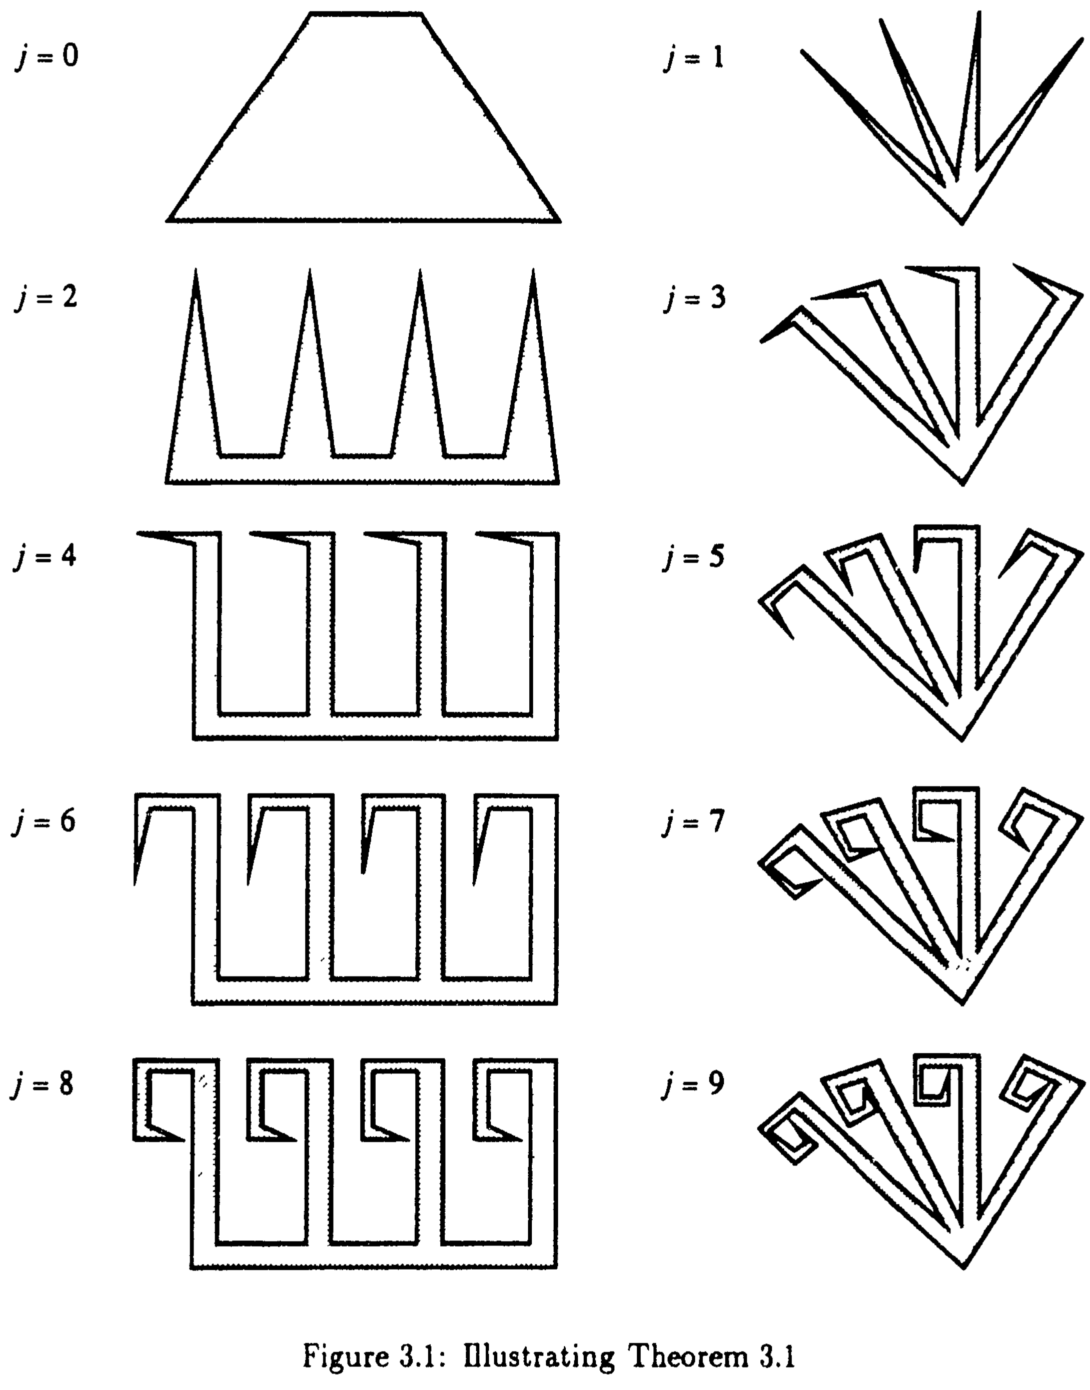

Figure 3.1: Illustrating Theorem 3.1

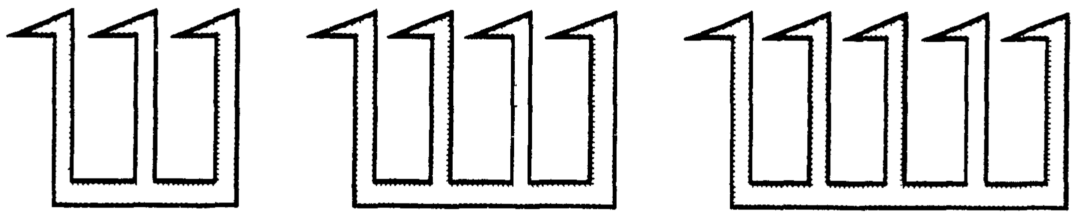

Figure 3.2: Increasing n

number of spirals on the polygon. For example, figure 3.2 shows representatives for $n = 15$, 20, and 25, for $j = 4$. Note than for values of n which are not multiples of $j + 1$, we can construct polygons by simply subdividing the appropriate number of edges of the polygon for the greatest multiple of $j + 1$ less than n.

Let P be a polygon with a L_j-hidden set of size $\lfloor n/(j+1) \rfloor$. If P were coverable with fewer than $\lfloor n/(j+1) \rfloor$ L_j-convex regions, then some region would contain two members of the hidden set, implying that these two members were link distance (at most) j apart. This means that these two members of the hidden set are L_j-visible, which is a contradiction. Therefore, P requires at least $\lfloor n/(j+1) \rfloor$ regions in a covering by L_j-convex regions. The bound holds for partitions as well, as every partition is also a covering. □

Corollary 3.1a *For any integers $j \geq 0$, $k \geq 0$, and $n \geq k + 2j + 1$, and any guard class C such that $C \subseteq \mathcal{L}_k$, using L_j-visibility, there exist polygons with n vertices requiring $\lfloor n/(k+2j+1) \rfloor$ guards.*

PROOF We claim that $VP_j(R)$, where $R \in C(P)$, is L_{k+2j}-convex. Take any two points x and y in $VP_j(R)$. The point x is L_j-visible to some point p in R, as it is in

$VP_j(R)$. Similarly, y is L_j-visible to some point q in R. Since $R \in \mathcal{C}(P)$, and $\mathcal{C} \subseteq \mathcal{L}_k$, the link distance between p and q is at most k. Therefore, there is a path from x to p to q to y that consists of at most $j + k + j = 2j + k$ links Since x and y were chosen arbitrarily in $VP_j(R)$, the claim follows. By the claim and Theorem 3.1(b), we have that there exist polygons requiring $\lfloor n/(k + 2j + 1) \rfloor$ guards, for $j \geq 0$ and $k \geq 0$ □

Corollary 3.1b *For any integers $j \geq 0$ and $n \geq j + 3$, there exist polygon triangulation graphs with n vertices having a distance-j independent set of size $\lfloor n/(j + 1) \rfloor$.*

PROOF Theorem 3.1 states that some polygon P exists with a L_j-hidden vertex set H of size $\lfloor n/(j + 1) \rfloor$. Let T be the triangulation graph of any triangulation of P. We claim that H is a distance-j independent set in T. If this were not the case, then there would exist some vertices h_1 and h_2 in H such that there is a D_j-path from h_1 to h_2 in T. As the geometric embeddings of each of the edges of a triangulation graph of a polygon is contained in the polygon, the D_j-path defines a link-j path from h_1 to h_2 in P. This contradicts the definition of H, therefore the claim holds. □

Corollary 3.1c *For any integers $k \geq 0$ and $n \geq k + 3$, there exist polygon triangulation graphs with n vertices that cannot be dominated by fewer than $\lfloor n/(k + 3) \rfloor$ D_k-subgraphs.*

PROOF By Corollary 3.1b, there exists some triangulation graph T with a distance-$(k + 2)$ independent set H of size $\lfloor n/(k + 3) \rfloor$. For each element h of H, we let $S(h)$ be some triangle of T containing h, and $S = \bigcup_{h \in H} S(h)$. S therefore has $\lfloor n/(k + 3) \rfloor$ members.

We claim that no D_k-subgraph of T can dominate more than one member of S If this were not the case, then there would be two members $S(h_1)$ and $S(h_2)$ of S such that there is a path of distance at most k between a vertex v_1 of $S(h_1)$ and a vertex v_2 of $S(h_2)$. As the distance from any h to any vertex of $S(h)$ is at most 1, there is some path from h_1 to v_1 to v_2 to h_2 of at most $1 + k + 1 = k + 2$ edges Since this contradicts the definition of H, the claim holds.

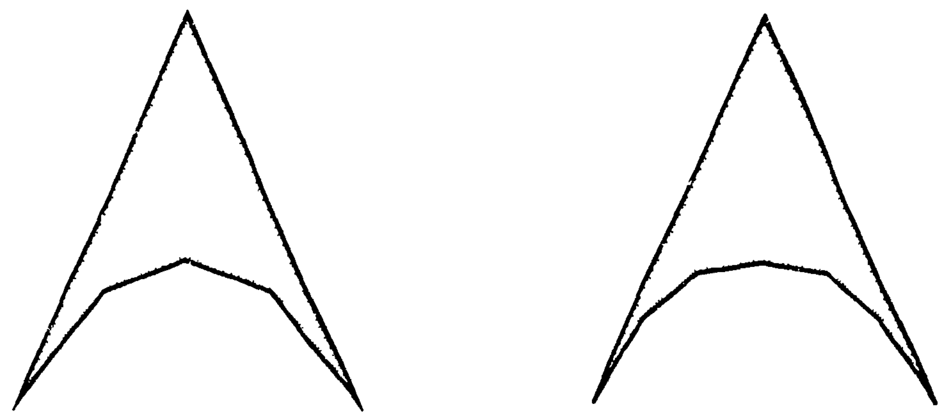

Figure 3.3: Spiral polygons with 3 convex vertices

As there are $\lfloor n/(k+3) \rfloor$ members of S, and no two members can be dominated by one D_k-subgraph, T cannot be dominated by fewer than $\lfloor n/(k+3) \rfloor$ such subgraphs. □

The following two theorems are bounds for hidden sets using L_1-visibility.

Theorem 3.2 *For any integer $n > 3$, there exist polygons with n vertices having a L_1-hidden set of size $n-2$. Furthermore, there do not exist polygons of n vertices having a L_1-hidden set of size greater than $n-2$.*

PROOF The spiral polygons with three convex vertices, as shown in figure 3.3, are such polygons. The hidden set is the set of midpoints of the edges on the reflex chain.

No polygon could have a hidden set of size greater than $n-2$, as every polygon can be triangulated (divided into $n-2$ triangles), and each triangle can only contain one member of a hidden set. □

The previous theorem strengthens the result of Theorem 3.1 for L_1-visibility.

3.2 Hidden Guard Sets

We now consider hidden guard sets, where the guards may be on any point in the polygon. In the next section we will see that the hidden vertex guard set problem is much more complicated. We start by showing that every polygon has a hidden guard set.

Theorem 3.3 *For any integer $j > 0$, polygon P, and L_j-hidden set $H \subseteq P$, there is a L_j-hidden guard set S for P which contains H.*

PROOF The following procedure generates such an S. First, let $S = H$. Repeatedly add points to S: at each step, add any point of P that is not seen (using L_j-visibility) from some point of S. Continue this until there are no such points left (i.e., S is a guard set for P). At each step in the construction, S is also a hidden set. This implies that S is finite, and thus the given procedure terminates. Therefore the final S is a hidden guard set. □

Theorem 3.4 *For any integers $j > 0$ and $n \geq 6j$, there exist polygons with n vertices with a minimum L_j-hidden guard set of size $\lfloor n/2j \rfloor - 1$.*

PROOF The *L_j-spur polygons* (shown in figure 3.4 for $j = 1, n = 8$ and 12, and $j = 4, n = 24$ and 40) are such polygons. In each polygon, there are $n/2j$ spiral arms, thus there are $n/2j$ vertices at the ends of spiral arms. Only guards placed in the the central region can be L_j-visible to more than one such end vertices. The spurs are constructed so that at most one (hidden) guard can be in this central region, and this guard will see at most two end vertices. Each other end vertex will require one guard; therefore these polygons require $n/2j - 1$ hidden guards. □

3.3 Hidden Vertex Guard Sets

In this section, we investigate and find bounds on hidden vertex guard sets in polygons. The first question that must be addressed is whether or not a given polygon

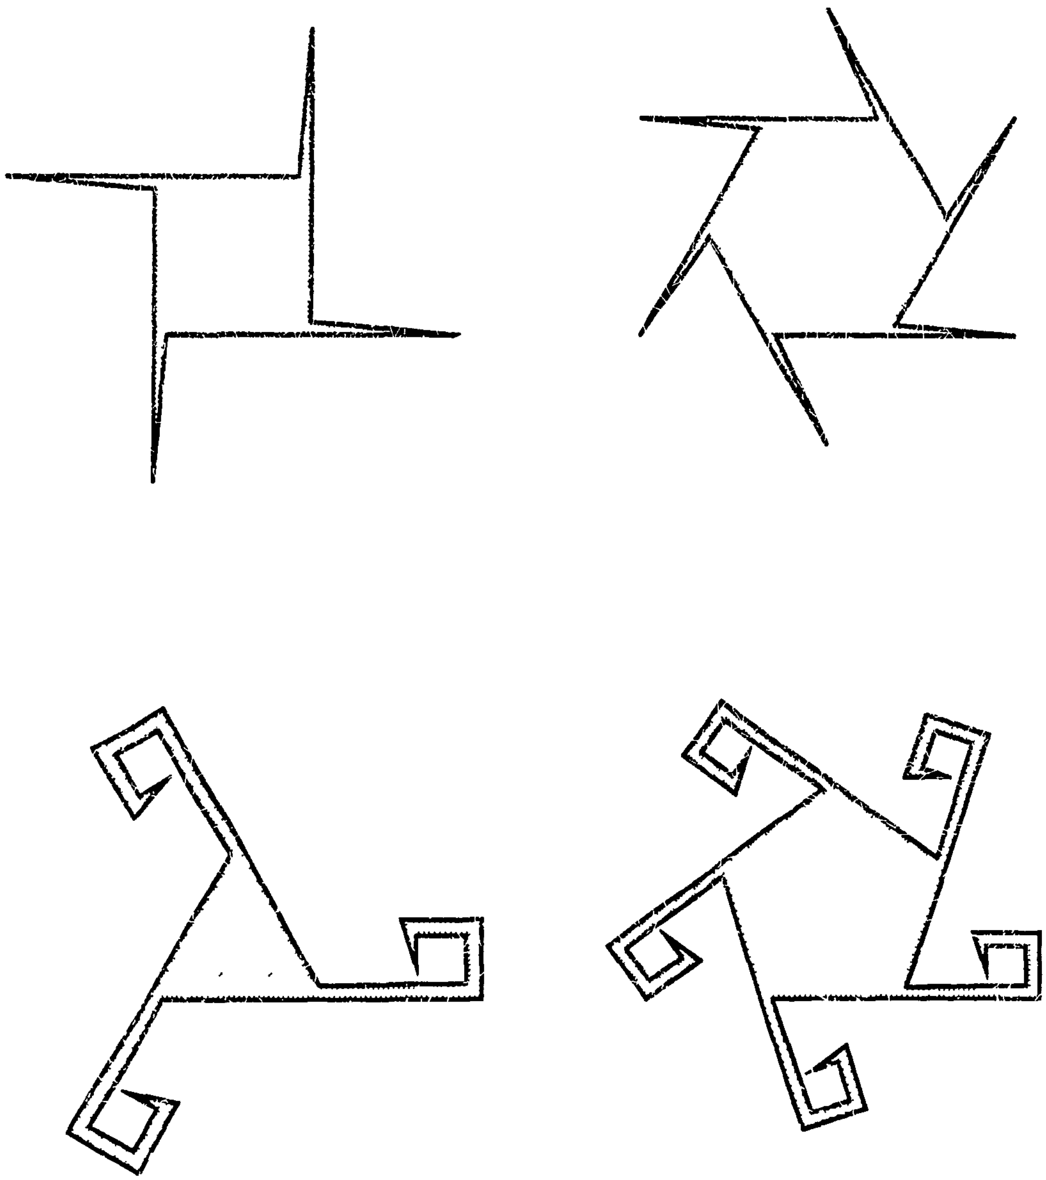

Figure 3.4: L_j-Spur polygons

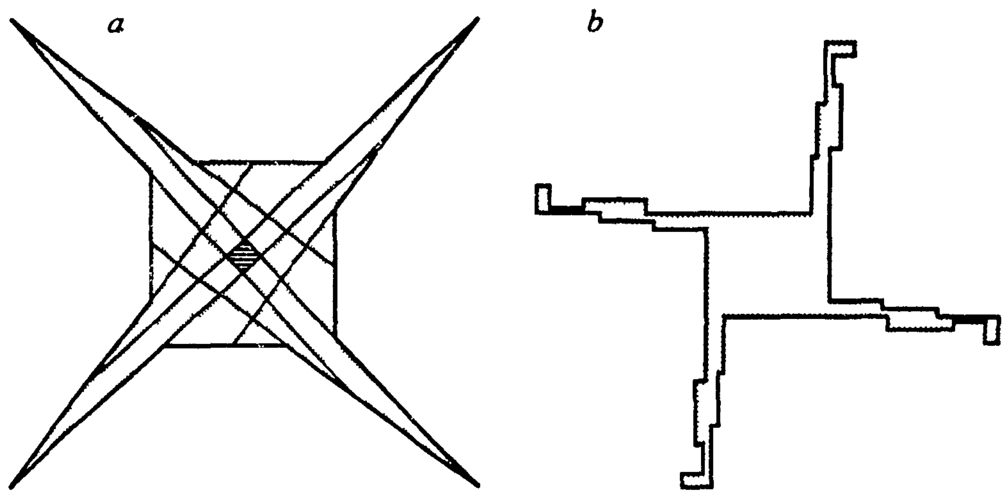

Figure 3.5: Polygons with no hidden vertex guard set

has a hidden vertex guard set. Surprisingly, there are polygons for which no hidden vertex guard set exists.

For example, neither of the polygons in figure 3.5 has a hidden vertex guard set. Consider figure 3.5a. Since guarding all of its extreme vertices does not cover the entire region, one of the interior vertices must be guarded (if this polygon is to admit a hidden vertex guard set). No more than one interior vertex can be guarded, however, as all interior vertices see one another. Guarding any interior vertex will leave two opposing triangles (and possibly some other region, whose guarding will not affect these triangles) as shown in figure 3.6: Since guarding neither v_a nor v_b covers both triangles, and v_a and v_b cannot both be guarded at the same time (they see one another), the polygon in figure 3.5a does not admit a hidden vertex guard set.

A similar argument holds for the polygon in figure 3.5b: to see the center point, one of the central four vertices must be guarded, but this leaves a triangle-pair as in figure 3.6.

We note that figure 3.5a is both star-shaped and monotone, and that figure 3.5b is orthogonal; therefore these classes of polygons (star-shaped, monotone, or orthogonal) do not always admit hidden vertex guard sets.

Spiral polygons, however, always admit hidden vertex guard sets; every other

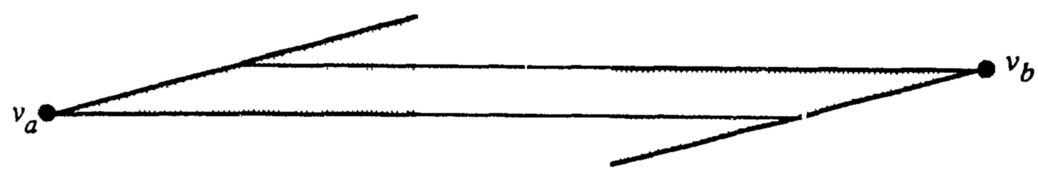

Figure 3.6: Opposing triangles

vertex of the reflex chain is such a set (see figure 3.7).

Since we know that some polygons do not admit a hidden vertex guard set, and some do, it is natural to ask whether or not there exists a good algorithm to determine whether a given polygon admits a hidden vertex guard set or not. This problem has been shown NP-complete in [S87].

Note that Theorem 3.3 implies that any maximal L_j-hidden set for P is also a L_j-guard set for P. We can also show the following interesting analog of this statement for L_j-hidden vertex sets:

Theorem 3.5 *For any integer $j \geq 0$, and any polygon P, any maximal L_j-hidden vertex set for P is also a L_{j+1}-guard set for P.*

PROOF Let H be a maximal L_j-hidden vertex set for P, and p be a point in P. Then, p is L_1-visible to some vertex v of P, as it is contained in some triangle of some triangulation of P. If v is in H, then p is L_1-visible, hence L_{j+1}-visible, to a member of H. If v is not in H, then there must be some element w of H that is L_j-visible to v, else v could be added to H, implying that it is not a maximal L_j-hidden vertex set. The j-link path from w to v followed by the segment from v to p is a $j+1$-link path from w to p. Thus, p is L_{j+1}-visible to an element of H.

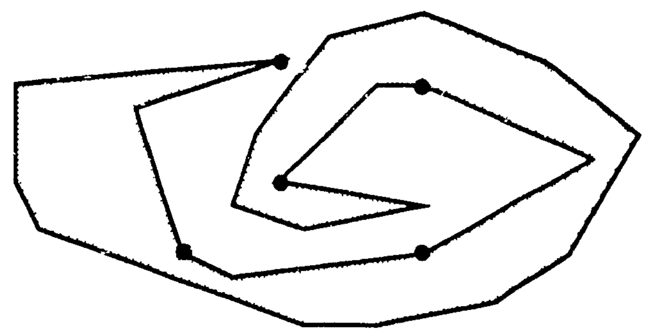

Figure 3.7: A hidden vertex guard set in a spiral polygon

As p was chosen arbitrarily, any such H will be a L_{j+1}-guard set for P. □

3.4 Polygon Exteriors

In this section, we obtain lower bounds on the maximum size of hidden sets, guard sets, and covers, for the exteriors of polygons.

We let $\mathcal{L}_k^e$ denote the guard class defined by letting $\mathcal{L}_k^e(P)$ be the collection of all L_k-convex sets in the exterior of P.

Theorem 3.6 *For any integers $j \geq 0$ and $n \geq j+1$, there exist polygons with n vertices having an exterior L_j-hidden vertex set of size $\lfloor n/(j+1) \rfloor$.*

PROOF Convex polygons are such polygons; the hidden vertex set consists of every $(j+1)$-th vertex. □

Theorem 3.7 *For any integer $n \geq 3$, there exist polygons with n vertices that*

(a) have an exterior L_1-hidden set of size n, and

(b) *require at least n regions in any covering or partition of the exterior by L_1-convex regions.*

Furthermore, no polygons exist with larger L_1-hidden sets or larger minimum covers by L_1-convex regions.

PROOF Convex polygons again provide the example; the hidden set consists of the midpoints of the edges of the polygon. As no two edge midpoints can be in the same L_1-convex region, a partition or cover by such regions must have at least n regions.

Also, the exterior of any polygon can be partitioned into n convex regions, using the naïve partitioning algorithm of [C80]. This provides the matching upper bound for partitions and covers. As no two members of any hidden set can be in the same convex region, this provides the upper bound on hidden sets as well. □

Part (b) of the above theorem is a well-known result (see, e.g. [O87]).

Theorem 3.8 *For any integers $j \geq 2$ and $n \geq j$, there exist polygons with n vertices that*

(a) *have an exterior L_j-hidden set of size $\lfloor (n+1)/(j+1) \rfloor$, and*

(b) *require at least $\lfloor (n+1)/(j+1) \rfloor$ regions in any covering or partition of the exterior by L_j-convex regions.*

PROOF The polygon class shown in figure 3.8 consists of such polygons. These polygons are derived by taking the polygons of Theorem 3.1 and turning them "inside-out." (This technique is due to O'Rourke and Aggarwal.)

The hidden sets are shown in the figure, and the covering result is proved in the same manner as the covering result of Theorem 3.1 was proved. □

Corollary 3.8a *For any integers $j \geq 0$, $k \geq 0$, and $n \geq k + 2j \geq 2$, and any guard class C such that $C \subseteq \mathcal{L}_k^e$, using L_j-visibility, there exist polygons with n vertices requiring $\lfloor (n+1)/(k+2j+1) \rfloor$ exterior guards.*

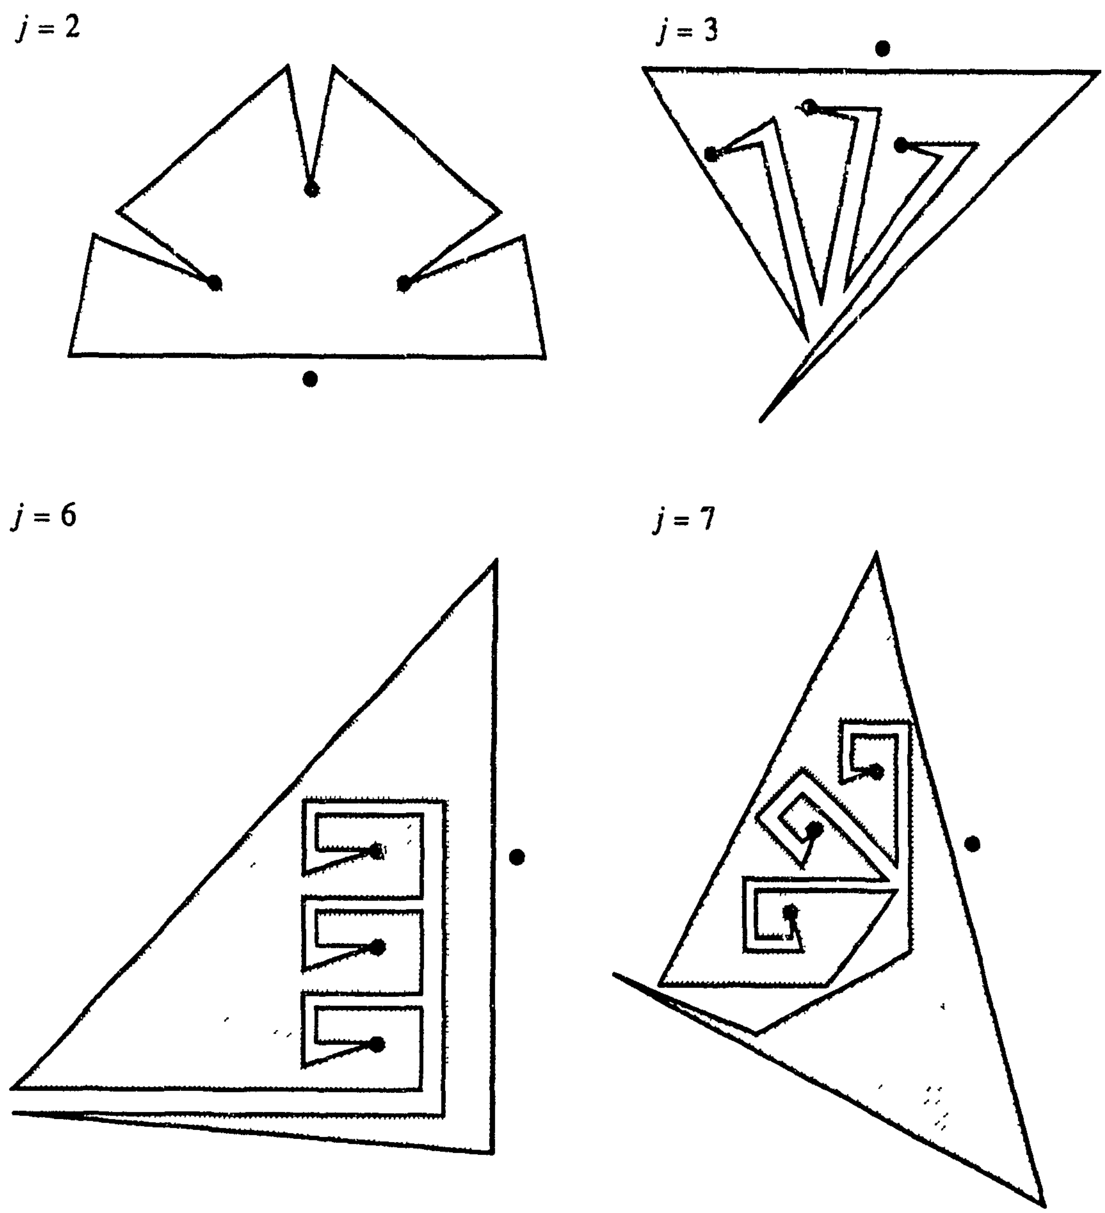

Figure 3.8: Illustrating Theorem 3.8

PROOF This is proved from Theorem 3.8 in the same manner as Corollary 3.1a was proved from Theorem 3.1. □

We note that this bound does not match the tight bound of $\lfloor (n+2)/(k+2j+1) \rfloor = \lfloor (n+2)/3 \rfloor$ for the $j = 1, k = 0$ case. However, we feel that $j = 1, k = 0$ is a special case in that it is the only case for which the exterior of a convex polygon (or the exterior of the hull of a non-convex polygon) requires two guards. We expect that the bound of Corollary 3.8a will be tight in all other cases.

Chapter 4

Upper Bounds

In this chapter we obtain upper bounds on hidden set, guard set, and polygon cover sizes, in polygon interiors, that are the same as the lower bounds presented in the previous chapter. In particular, we will prove the following general covering/guarding theorem:

Theorem 4.1 *For any guard class $\mathcal{C}$ such that $\mathcal{T}_k \subseteq \mathcal{C} \subseteq \mathcal{L}_k$, with $k \geq 0$ and $j > 0$, $\lfloor n/(k+2j+1) \rfloor$ guards of $\mathcal{C}$ are necessary (for some polygon) and sufficient (for all polygons) to guard polygons with n vertices, using L_j-visibility.*

This theorem unifies and generalizes the known guarding results. We also use it to obtain similar *almost-tight* bounds for polygon exteriors. We prove our theorem by generalizing the known art gallery proofs for simple polygons; these proofs were reviewed in chapter 2. Our generalization is not completely straightforward as there are several complexities introduced by the generalized dominating objects (D_k-trees) that we use.

These complexities necessitate two major differences between the known proofs and our proof. The first of these is that the induction strategy is altered. Previously, the proofs proceded in the following manner:

(1) Find a cutting diagonal D, dividing the triangulation graph G into a main piece G_1 and a small piece G_2.

(2) Based on the number of triangles in G_2, either:

(2a) Simply combine inductive dominating sets for G_1 and G_2 to get a dominating set for G, or

(2b) (2b1) Add some of the triangles in G_2 to G_1 to get G_1'.

(2b2) Inductively generate a dominating tree set for G_1'.

(2b3) Based on the inductive dominating tree set of G_1', find a dominating set for G_2 (and G).

In our proof, we change step (2b1) above to:

Perform an edge-contraction of D in G_1 to get G_1'.

This modification is not trivial, as the new induction does not always use subgraphs of our original graph, hence the inductive dominating tree sets may not be tree sets in our original graph. However, the changed induction simplifies the subsequent analysis; the original style of induction does not lead to a clean proof of our theorem.

The second major change is required because of the difference in the complexity of general trees as opposed to vertices or edges. At a critical point in the proof, it is desirable to have zero or one (rather than many) trees incident on any vertex. For $k = 0$ and $k = 1$, this is a triviality to enforce: for $k = 0$, we may throw away any duplicate trees (vertices); for $k = 1$, we may shorten one of any pair of intersecting trees (edges) to a vertex. For $k > 1$, however, no such simple strategy exists, and we are instead forced to complicate our proof by establishing and using a theorem about finding nonintersecting tree sets which cover the same vertices as a given intersecting tree set (Theorem 4.6).

We open the chapter with a section containing a simple proof of a tight upper bound on hidden vertex sets. Following that, we present the generalization of the proofs of Chvátal and O'Rourke. This comes in four sections: the first contains a general polygon cutting theorem, the second contains the theorem on finding nonintersecting tree sets in a graph, the third establishes the main theorem (an upper bound on the size of dominating tree sets), and the fourth contains important corollaries of the main theorem. We close the chapter with a section of results for polygon exteriors which parallel the corollaries of our main theorem.

4.1 Hidden Vertex Sets

Theorem 4.2 *For any integer $j > 0$, there are no polygons with a L_j-hidden vertex set of size larger than $\lfloor n/(j+1) \rfloor$.*

PROOF Suppose there was a polygon with a hidden vertex set $H = \{w_1, \ldots, w_h\}$, where h is larger than $\lfloor n/(j+1) \rfloor$ (this implies $n < h(j+1)$). Furthermore, we assume that the w_i's appear in counterclockwise order around the polygon, and we use the convention that $w_{h+1} = w_1$.

We label each edge with an integer, between 1 and h, such that an edge has label l if w_l is the first member of H clockwise around the polygon from the middle of the edge. By the pigeonhole principle, we see that there is some label l such that at most $\lfloor n/h \rfloor$ edges have label l. This means that between w_l and w_{l+1} there are at most $\lfloor n/h \rfloor$ edges.

Therefore, the link distance between these two members of H is at most $\lfloor n/h \rfloor$. However, $\lfloor n/h \rfloor < \lfloor (h(j+1))/h \rfloor = j+1$, hence $\lfloor n/h \rfloor$ is at most j. This means that the two elements w_l and w_{l+1} of the hidden set are visible (link-j), which is a contradiction. Therefore, there is no such polygon, and the theorem is proved. □

This theorem, combined with Theorem 3.1, establishes a tight bound of $\lfloor n/(j+1) \rfloor$ on the maximum size of a L_j-hidden vertex set inside (or outside) a simple polygon.

As the same argument applies to polygon exteriors, we have the following Corollary:

Corollary 4.2a *For any integer $j > 0$, there are no polygons with an exterior L_j-hidden vertex set of size larger than $\lfloor n/(j+1) \rfloor$.*

4.2 Cutting Diagonals in Polygons

In this section we present a result which we call the Cutting Diagonal Theorem. Given some t, this theorem guarantees the existance, in any triangulation graph of sufficient size, of at least one diagonal which cuts off between t and $2t-1$ triangles. The theorem is a generalization of lemmas due to Chvátal [C75] and O'Rourke [O83a], and our proof mimics their proofs. This theorem finds many uses in recursive algorithms and inductive proofs which deal with triangulations.

After our proof, we present several corollaries, many of which are known results.

Theorem 4.3 (The Cutting Diagonal Theorem) *Given a polygon triangulation graph G of n vertices, a cycle edge e of the graph and some positive integer $t \leq n-2$, there exists an edge D of G which separates G into two pieces G_1 and G_2 (with D in both pieces) such that:*

(a) G_1 has between t and $2t-1$ triangles, inclusive, and

(b) G_2 contains e.

The degenerate case $G_2 = e$ is allowed.

PROOF An edge E divides a triangulation graph G into two pieces $G_1(E)$ and $G_2(E)$, both containing the edge. We use the phrase "piece cut off by edge E" to indicate whichever piece ($G_1(E)$ or $G_2(E)$) does not contain e.

Let t' be the minimum number, greater than or equal to t, of triangles in any piece cut off by an edge, and let D be an edge which cuts off a piece with t' triangles. Such a D exists, as any cycle edge cuts off $n-2$ triangles, and $t \leq n-2$. Of the t' triangles cut off, let U be the one containing D (see figure 4.1). We note that t' is the sum of the triangles cut off by the other edges of U, plus one (for U). Each of the other edges of U must cut off less than or equal to $t-1$ triangles (else t' is not minimum). Therefore, $t' \leq 2(t-1)+1$, or $t' \leq 2t-1$. □

We will use the notation $CD(G, e, t)$ to denote the diagonal D guaranteed by this theorem, using G, e, and t as the graph, cycle edge, and integer in the hypothesis.

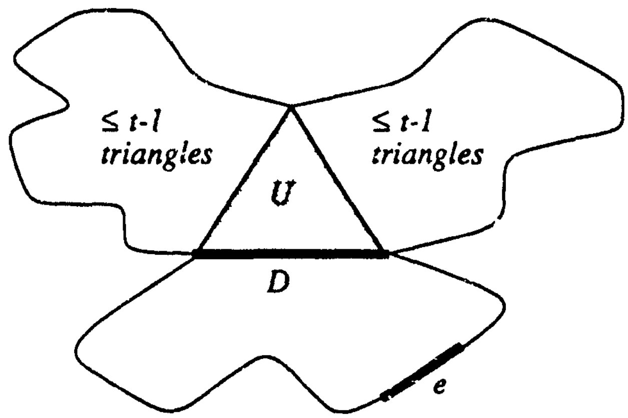

Figure 4.1: Illustrating Theorem 4.3

We are now ready to state some interesting corollaries of the Cutting Diagonal Theorem.

Corollary 4.3a *Given a polygon triangulation graph G of n vertices, and some positive integer $t \leq n-2$, there exists an edge D of G which separates G into two pieces G_1 and G_2 (with D in both pieces) such that G_1 has between t and $2t-1$ triangles, inclusive. The degenerate case $G_1 = G$ is allowed.*

PROOF Follows from Theorem 4.3, by choosing any cycle edge edge as e. □

Corollary 4.3b *Given a polygon triangulation graph G of n vertices, and some positive integer $t \leq \lfloor (n-1)/3 \rfloor$, there exists edges D_1 and D_2 of G which separate G into three pieces G_1, G_2, and G_3 such that:*

(a) G_1 and G_3 both have between t and $2t-1$ triangles, inclusive, and

(b) G_2 contains both D_1 and D_2.

The degenerate case $G_2 = D_1 = D_2$ is allowed.

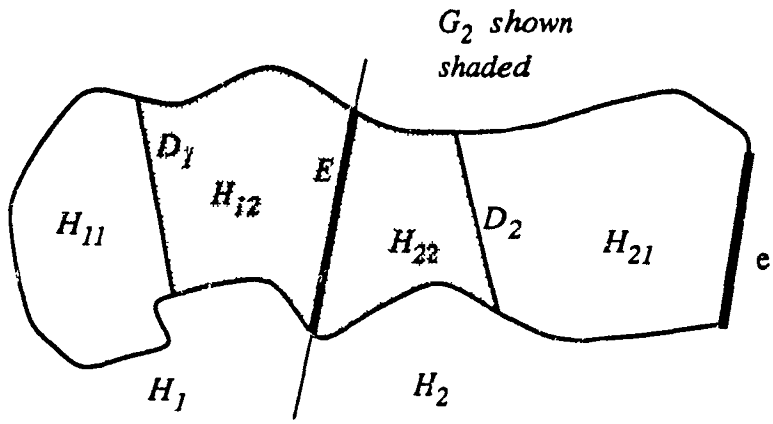

Figure 4.2: Illustrating Corollary 4.3b

PROOF Let e be any cycle edge of G, and find $E = CD(G, e, \lfloor (n-1)/3 \rfloor)$. E divides G into two parts, H_1 and H_2. Note that E is a cycle edge of both H_1 and H_2.

We now consider two cases:

case 1: There are at least t triangles in each of H_1 and H_2.

This implies that H_1 and H_2 have at least $t+2$ vertices. Thus, we may apply Theorem 4.3 to find $D_1 = CD(H_1, E, t)$ and $D_2 = CD(H_2, E, t)$. D_1 divides H_1 into parts H_{11} and H_{12}, with H_{11} containing t to $2t-1$ triangles, and H_{12} containing E. Similary, D_2 divides H_2 into parts H_{21} and H_{22}.

Finally, we let $G_1 = H_{11}$, $G_2 = H_{12} \cup H_{22}$, and $G_3 = H_{21}$. We note that G_2 is a single connected piece, as both H_{12} and H_{22} contain E. Therefore, we have D_1, D_2, G_1, G_2, and G_3 satisfying the theorem. The situation is illustrated in figure 4.2.

case 2: Either H_1 or H_2 has less than t triangles (without loss of generality, assume H_2 has less than t triangles).

By our choice of E as $CD(G, e, \lfloor (n-1)/3 \rfloor)$, the number of triangles in H_1 must be $\leq 2\lfloor (n-1)/3 \rfloor - 1$. The total number of triangles is therefore less than $t+2\lfloor(n-1)/3\rfloor-1 \leq 3\lfloor(n-1)/3\rfloor-1$. However, $3\lfloor(n-1)/3\rfloor \leq n-1$, so the total number of triangles is (strictly) less than $(n-1)-1=n-2$. This is a contradiction, as the number of triangles must be exactly $n-2$. Therefore this case cannot happen.

□

To state our next corollary, we must first generalize the notion of an ear of a polygon to that of a *k-ear*. Recall that an ear of a polygon is a vertex v_i such that the diagonal $[v_{i-1}, v_{i+1}]$ intersects the polygon boundary only at its endpoints. A k-ear is a collection of k to $2k-1$ consecutive vertices $v_i \dots v_j$ such that the diagonal $[v_{i-1}, v_{j+1}]$ intersects the polygon boundary only at its endpoints. An ear is then simply a 1-ear.

Corollary 4.3c (The Two k-Ears Theorem) *Every polygon has at least two nonoverlapping k-ears, for any positive integer* $k \leq \lfloor(n-1)/3\rfloor$.

PROOF This follows directly from the geometric interpretation of Corollary 4.3b and the definition of a k ear. □

The utility of the Cutting Diagonal Theorem and the aforementioned corollaries is illustrated by the many places in the literature in which we can find special cases and weaker versions of it. The following five corollaries are all lemmas and theorems from published papers.

Corollary 4.3d (Meisters' Two Ear Theorem [M75]) *Every polygon that is not a triangle has at least two nonoverlapping ears.*

This is simply the special case of Corollary 4.3c when $k=1$.

Corollary 4.3e (Chvátal [C75]) *Every polygon triangulation graph of at least 6 vertices has a diagonal which cuts off a piece with 3-5 triangles.*

This is the special case of Corollary 4.3a when $t=3$.

Corollary 4.3f (O'Rourke [O83a]) *Every polygon triangulation graph of at least 10 vertices has a diagonal which cuts off a piece with 4-7 triangles.*

This is the speical case of Corollary 4.3a when $t = 4$.

Corollary 4.3g (Avis-Toussaint [AT81a]) *Every polygon triangulation graph has a diagonal which cuts off a piece of between* $\lfloor n/4 \rfloor$ *and* $\lfloor 3n/4 \rfloor$ *vertices.*

This is implied by Corollary 4.3a when $\lfloor n/3 \rfloor \geq t \geq \lfloor n/4 \rfloor$.

Corollary 4.3h (Chazelle[C82]) *Every polygon triangulation graph has a diagonal which cuts off a piece of between* $\lfloor (n-2)/3 \rfloor$ *and* $2\lfloor (n-2)/3 \rfloor$ *triangles, inclusive.*

This corollary is actually only a special case of Chazelle's theorem; his theorem allows weights of 0 or 1 on each triangle, and finds a diagonal which makes the weight on each side between 1/3 and 2/3 of the total. Theorem 4.3 can be generalized in this manner, making an even more general cutting theorem. However, this generalization is beyond the purpose and scope of this thesis.

4.3 Nonintersecting Tree Sets

In this section we present a theorem on finding certain sets of nonintersecting subtrees in graphs. We first give two lemmas needed in the proof of this theorem.

Lemma 4.4 *Given a tree T which is a subgraph of a graph G, perform any finite series of contractions on G and T to give G^* and T^* respectively. Then, there is a tree S in G^* such that* $vert(S) = vert(T^*)$ *and* $diam(S) \leq diam(T)$.

PROOF It suffices to prove that the lemma holds for a single contraction, rather than a finite series of them. Repeated application of this proof then yields the lemma as stated.

If the two vertices being contracted are not both in T, or if they are adjacent in T, then $S = T^*$ satisfies the lemma.

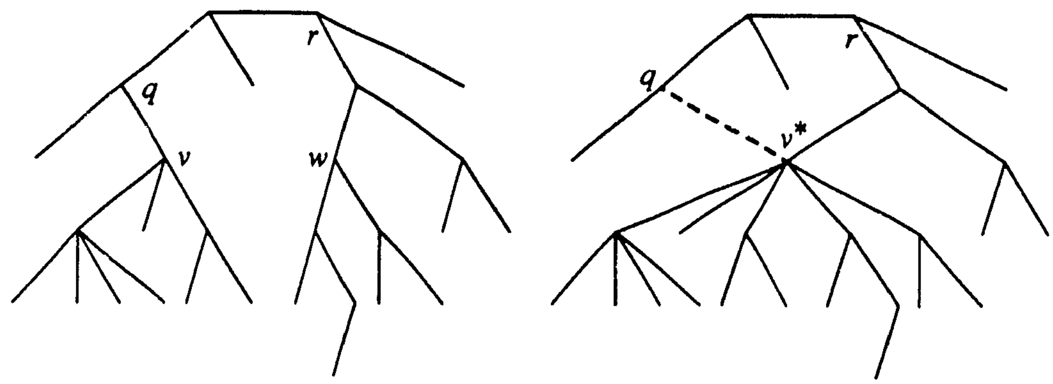

Figure 4.3: Illustrating Lemma 4.4

Otherwise, let v and w be the two vertices being contracted, and let v^* be the new vertex. Also, let r be some vertex on the path from v to w in T, and henceforth consider T and T^* to be rooted at r. Without loss of generality, assume that the depth of the subtree of T starting at w is not less than the depth of the subtree starting at v: $depth(st(T, w)) \geq depth(st(T, v))$. Let q be the first vertex (perhaps r) on the path from v to r in T (see figure 4.3).

Then, let S be T^* with the edge $[v^*, q]$ removed. Note that S is a tree, as the only cycle of T^* is formed by the contraction of the two ends of the path from v to w, and the removal of $[v^*, q]$ breaks this cycle. Note also that this definition of S implies that $vert(S) = vert(T)$.

We now check that $diam(S) \leq diam(T)$. Let x and y be any two vertices in S. Note that $vert(S) = vert(st^*(T, v) \cup st^*(T, w) \cup rt^*(T, [v, w]))$, therefore x and y must each be in one of these three components. We consider all cases. If x and y are both in $st^*(T, v)$, or both in $st^*(T, w)$, or both in $rt^*(T, \{v, w\})$, then their distance in T and in $T^*$are the same, and thus $d(x, y) \leq diam(T)$. This is also the case when one of them is in $st^*(T, w)$ and the other is in $rt^*(T, [v, w])$. If one is in $st^*(T, w)$ and one is in $st^*(T, v)$, then their distance has decreased (by the distance from v to w). Finally if one of them (without loss of generality assume that it is x) is in $rt^*(T, [v, w])$, and the other (y) is in $st^*(T,v)$, then (in tree S), $d(x,y) = d(x,v^*) + d(v^*,y) \leq d(x,v^*) + depth(st(T,v)) \leq d(x,w) + depth(st(T,w)) \leq diam(T)$. Therefore, in every case, $d(x,y) \leq diam(T)$, so $diam(S) \leq diam(T)$. Thus S is a tree satisfying the lemma, and the lemma is proved. □

Lemma 4.5 *Given a set of m nonintersecting trees $T = \{T_i\}$ in a graph G, perform any finite series of contractions on G and T to give G^* and $T^* = \{T_i^*\}$, respectively. Then, there is a set of m nonintersecting trees $S = \{S_i\}$ in G^* such that $vert(\bigcup_i S_i) = vert(\bigcup_i T_i^*)$, and $diam(S_i) \leq diam(T_i)$ for $1 \leq i \leq m$.*

PROOF Once again we need only consider the single-contraction case. Let v and w be the two vertices being contracted, and v^* be the new vertex. If v and w are not both in trees in T, then $S_i = T_i^*$ satisfies the lemma. If v and w are on the same tree (wlog, assume they are both on T_1), then let $S_i = T_i^*$ for all $1 < i \leq m$, and let S_1 be the tree S guaranteed by Lemma 4.4 (where the T of Lemma 4.4 is T_1 here). This choice of S_i clearly satisfies the lemma.

Otherwise, v and w are (wlog) on two trees T_1 and T_2 respectively, and we will let $S_i = T_i^*$ for all $2 < i \leq m$. Assume T_1 and T_2 are rooted at v and w.

Let $p_1, p_2, \ldots, p_P$ be the vertices of T_1 adjacent to v in T_1, and $q_1, q_2, \ldots, q_Q$ be the vertices of T_2 adjacent to w in T_2. Furthermore, choose p_P such that $depth(st(T_1,p_P)) \leq depth(st(T_1,p_i))$ for all $1 \leq i < P$, and similarly choose q_Q. Let bt_1 be $st(T_1,p_P)$ and lt_1 be $\bigcup_{i<P}(st(T_1,p_i) \cup [v,p_i])$. Similarly define bt_2 and lt_2. Note that $depth(lt_1) \leq diam(T_1)/2$, and $depth(lt_2) \leq diam(T_2)/2$. Without loss of generality assume $depth(lt_2) \geq depth(lt_1)$. Then, let $S_1 = bt_1$, and $S_2 = bt_2 \cup [v^*, q_Q] \cup lt_1^* \cup lt_2^*$. Figure 4.4 illustrates these definitions.

We claim that the S_i, as defined, have diameter less than the corresponding T_i. Certainly S_1, and S_i for $2 < i \leq m$, do. The only possible problem is with S_2. Arbitrarily choose two vertices x and y of S_2. Since $vert(S_2) = vert(bt_2 \cup lt_1^* \cup lt_2^*)$, each of x and y must be in one of those three components. We consider all cases. If x and y are both in $bt_2 \cup lt_2^*$, then $d(x,y)$ is the same in T_2 and in S_2, hence

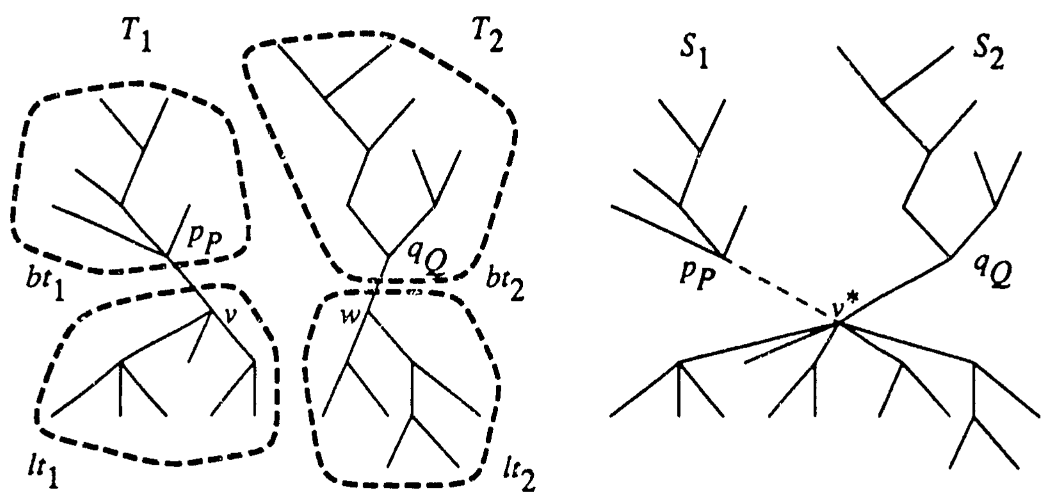

Figure 4.4: Illustrating Lemma 4.5

$d(x,y) \leq diam(T_2)$. If x and y are both in lt_1^*, then $d(x,y) \leq 2 * depth(lt_1) \leq 2 * depth(lt_2) \leq diam(T_2)$. If one is in lt_1^* and the other in $bt_2 \cup lt_2^*$, then $d(x,y) \leq depth(lt_1) + depth(bt_2) + 1 \leq depth(lt_2) + depth(bt_2) + 1 \leq diam(T_2)$. Therefore, in all cases, $d(x,y) \leq diam(T_2)$, so $diam(S_2) \leq diam(T_2)$.

Note also that there is no intersections among the trees S_i, as the only intersecting T_i's were T_1 and T_2, and S_1 and S_2 do not intersect, and contain no vertices other than those in T_1 and T_2 (hence could not possibly intersect another S_i). Therefore, the S_i, as defined, satisfy the lemma, and so the single-contraction version of the lemma holds. Hence, the lemma as stated holds. □

Theorem 4.6 *Given a set of m possibly intersecting trees $T = \{T_i\}$ in a graph G, there is a set $S = \{S_i\}$ of nonintersecting trees in G such that $vert(\bigcup_i S_i) = vert(\bigcup_i T_i)$, and (for $1 \leq i \leq m$), $diam(S_i) \leq diam(T_i)$.*

PROOF We construct a graph G^* as the union of m copies $G_1, G_2, \ldots, G_m$ of G, and a set of trees $T^* = \{T_i^*\}$ where T_i^* is the tree T_i as a subgraph of the graph G_i. We then perform, for each vertex v in G, a series of contractions in G^* that bring all m of the copies of v together into one vertex. Lemma 4.5 then gives the desired result. □

4.4 Dominating Tree Sets

In this section we will prove the major result of this chapter, concerning the number of D_k-trees sufficient to dominate a triangulation graph. We first present a lemma which establishes the induction basis for the proof.

Lemma 4.7 *One D_k-tree is sufficient to dominate any triangulation graph of up to $2k+5$ vertices ($2k+3$ triangles).*

PROOF By induction on k.

The induction basis, $k = 0$, is easily shown: any triangulation graph of 3, 4, or 5 vertices has a vertex which all triangles are incident on, and therefore can be dominated by one vertex (a vertex is a D_0-tree).

The induction hypothesis is that the lemma is true for all $k' < k$. We wish to show that the lemma is then true for k.

Let m be the number of vertices in the triangulation graph G under consideration. Then, $m \leq 2k+5$. We consider 2 cases:

case 1: $m \leq 2k+3$.

> In this case, the induction hypothesis states that G can be dominated by a D_{k-1}-tree. Since any D_{k-1}-tree is also a D_k-tree, the lemma holds.

case 2: $m = 2k+4$ or $2k+5$.

> By Corollary 4.3a (with $k = 2$), there is a diagonal D of G which cuts off 2 or 3 triangles. Use one that cuts off 2 if such a diagonal exists. Let G_1 and G_2 be the pieces, as in the corollary. Note that G_2 is a triangulation graph of $2k+1$ to $2k+3$ vertices.
>
> By the induction hypothesis, G_2 can be dominated by a D_{k-1}-tree. Let T be such a tree.

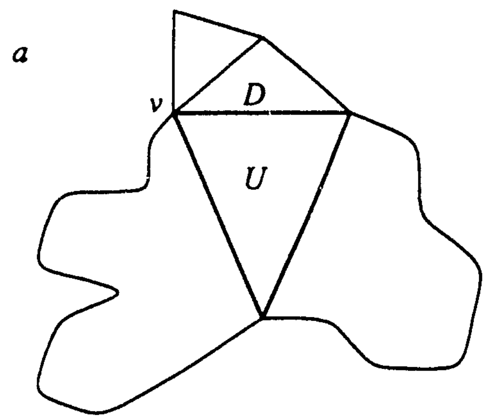

Figure 4.5: Illustrating Lemma 4.7

We now consider 2 subcases, depending on how many triangles are in G_1:

case 2a: G_1 contains 2 triangles (figure 4.5a).

These 2 triangles share some vertex v with D. Let U be the triangle of G_2 which has D as an edge. At least one of the vertices of U is a vertex of T. If v is such a vertex, then T is a tree (of diameter at most $k-1$) satisfying the lemma. Otherwise, let x be such a vertex, and join the edge $[x, v]$ to T, giving a tree of diameter at most $(k-1)+1 = k$ dominating G, and therefore satisfying the lemma.

case 2b: G_1 contains 3 triangles (figure 4.5b).

Note that figure 4.5b is the only possible configuration of 3 triangles which does not admit a diagonal which cuts off 2 triangles. These 3 triangles are dominated by the endpoints v and w of D, as shown. If both v and w are vertices of T, then T satisfies the lemma. If only one of v and w is in T, then add the diagonal D to T, giving a tree of diameter at most $(k-1)+1 = k$ satisfying the lemma.

Otherwise, we examine two cases. First, if $k = 1$, then G_2 is a single triangle (as we would otherwise have cut off two triangles). Therefore D is a dominating D_k-tree. Next, if $k > 1$, then we let x be the third vertex of U, and add the diagonals $[x, v]$ and $[x, w]$ to T. again yielding a D_k-tree satisfying the lemma.

Thus, in all cases, we have exhibited that the lemma holds for k; by induction it therefore holds for all finite k. □

Theorem 4.8 *For all $k \geq 0$, $\lfloor n/(k+3) \rfloor$ nonintersecting D_k-trees are sufficient to dominate any triangulation graph of $n \geq k+3$ vertices.*

PROOF Lemma 4.7 does the induction on k to provide us with the basis for the induction on n: one (obviously nonintersecting) D_k-tree suffices for $k+3 \leq n \leq 2k+5$. Therefore, for some fixed k, we assume that $\lfloor n'/(k+3) \rfloor$ D_k-trees suffice for all triangulation graphs of $n' < n$ vertices, where $n > 2k+5$. We will show that $\lfloor n/(k+3) \rfloor$ D_k-trees suffice for any triangulation graph of n vertices.

Let G be an arbitrary triangulation graph of n vertices. By Corollary 4.3a, we can find a diagonal D in G that cuts off a piece G_1 with between $k+2$ and $2k+3$ triangles, inclusive. We consider the case where D cuts off $k+3$ to $2k+3$ and the case where D cuts off $k+2$ separately.

If G_1 has between $k+3$ and $2k+3$ triangles, then G_2 (the remaining piece) has between $n-k-5$ and $n-2k-5$ triangles. We dominate on each piece by induction. G_1 has between $k+5$ and $2k+5$ vertices, which, by induction (or Lemma 4.7), requires 1 D_k-tree. G_2 has between $n-k-3$ and $n-2k-3$ vertices, which by induction requires at most $\lfloor (n-(k+3))/(k+3) \rfloor = \lfloor n/(k+3) \rfloor - 1$ D_k-trees. Combining these dominating D_k-tree sets for G_1 and G_2 gives a total of at most $\lfloor n/(k+3) \rfloor$ D_k-trees. However, since G_1 and G_2 share the diagonal D, these trees may intersect. If this is the case, then Theorem 4.6 may then be applied to give a nonintersecting tree set.

We now consider the case where G_1 has $k+2$ triangles. This means that G_1 contains $k+3$ cycle edges of G. Consider the triangle $U = (v, w, x)$ of G_2 that has

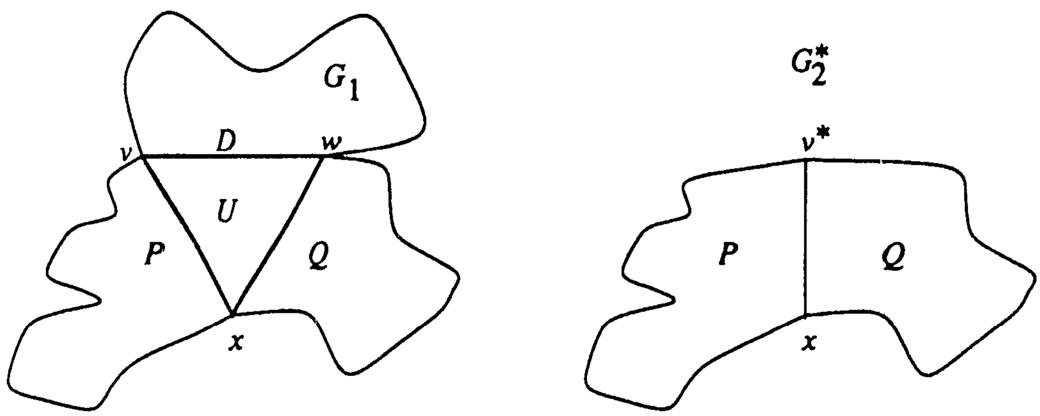

Figure 4.6: Illustrating Theorem 4.8

$D = [v, w]$ as an edge. U divides G_2 into two parts P and Q. We will perform induction of the graph G_2^*, which is obtained from G_2 by contracting v and w (see figure 4.6). G_2^* has $n - (k + 3)$ vertices, hence by induction can be dominated by $\lfloor (n - (k+3))/(k+3) \rfloor = \lfloor n/(k+3) \rfloor - 1$ D_k-trees. We now consider two cases based on the inductive dominating D_k-tree set Γ of G_2^*.

case 1: v^* has no tree incident on it.

Then, Γ is also a dominating D_k-tree set for $G_2 - D$ (G_2 with edge D removed). We inductively dominate the $k + 5$ triangle graph $G_1 + U$ with 1 D_k-tree, and note that the triangles of the pieces $G_2 - D$ and $G_1 + U$ are exactly the triangles of G. Therefore, we can combine the dominating D_k-tree sets for these two pieces to get a dominating D_k-tree set for G using at most $\lfloor n/(k+3) \rfloor - 1 + 1 = \lfloor n/(k+3) \rfloor$ (possibly intersecting) D_k-trees. By Theorem 4.6 we can make these nonintersecting trees.

case 2: v^* has a D_k-tree $T \in \Gamma$ incident on it.

Let T be rooted at v^*, and let V_p be the vertices of $T \cap P$ that are adjacent (in T) to v^*, and define V_q similarly (let x be in P and not in

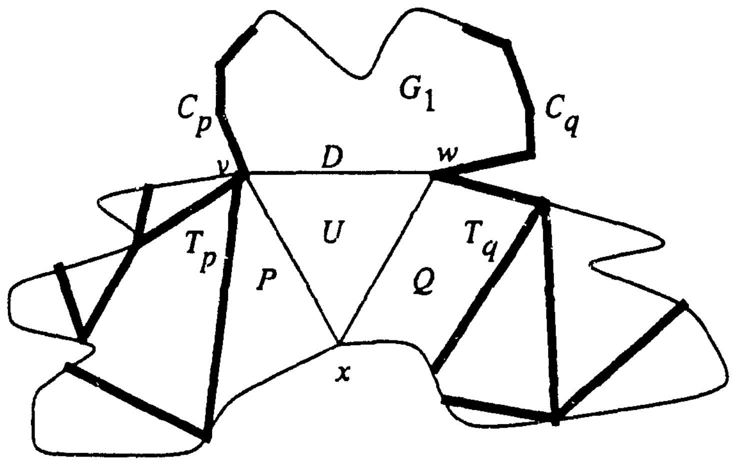

Figure 4.7: Case 2 in Theorem 4.8

Q for these definitions). Then, let T_p be $\bigcup_{p \in V_p} (st(T, p) \cup [v, p])$ and T_q be $\bigcup_{q \in V_q} (st(T, q) \cup [v, q])$.

We note that $depth(T_p) + depth(T_q) \leq k$, else T has diameter greater than k. We let C_p be a chain of $(k - depth(T_p))$ cycle edges of G starting at v and proceeding into G_1, and C_q be a chain of $(k - depth(T_q))$ cycle edges of G starting at w and proceeding into G_1. Let S_p be $C_p \cup T_p$, and S_q be $C_q \cup T_q$ (See figure 4.7).

We claim that $\Gamma - T + S_p + S_q$ is a dominating D_k-tree set for G. We note that this tree set certainly contains all of the vertices that the elements of Γ did, hence $G_2 - U$ is dominated. Also, both S_p and S_q dominate U, therefore all of G_2 is dominated. If S_p and S_q cover all of the vertices of G_1, then G_1 is dominated; else S_p and S_q have no common vertices. In that case, S_p covers $k - depth(T_p) + 1$ vertices and S_q covers $k - depth(T_q) + 1$ vertices of G_1. Thus a total of $2k + 2 - (depth(T_p) + depth(T_q))$ vertices of G_1 are covered. Since $depth(T_p) + depth(T_q) \leq k$, the total number of covered vertices of G_1 is at least $k + 2$. Since G_1 has $k + 2$ triangles, it has $k+4$ vertices, hence at most 2 vertices not covered by S_p and S_q. This means that each triangle must have at least one covered vertex, hence G_1 is dominated. Therefore G is dominated.

Note that the number of trees in this dominating tree set is $\lfloor n/(k+3) \rfloor$, and that S_p and S_q are of diameter at most k. If any of the trees of this dominating set intersect, we may apply Theorem 4.6 to make them nonintersecting.

In both cases we have exhibited a dominating D_k-tree set satisfying the theorem for n. Also, we need only consider these two cases (as the trees are or can be made nonintersecting, by Theorem 4.6). Hence, the theorem is proved for all finite k and n. □

4.5 Corollaries

Corollary 4.8a *For any guard class C such that $T_k \subseteq C$, $k \geq 0$, $\lfloor n/(k+3) \rfloor$ guards of C are sufficient to guard (using L_1-visibility) any polygon P of n vertices.*

PROOF Theorem 4.8 states that there is a set of $\lfloor n/(k+3) \rfloor$ dominating D_k-trees in any triangulation graph of P; the embeddings of these guards (each of which is a T_k-guard) will see the entire polygon, because each point of the polygon is in some triangle, and each triangle has a guard on some vertex, and all points in each triangle see one another. Since every T_k-guard is contained in a member of C, $\lfloor n/(k+3) \rfloor$ guards of C suffice. □

Corollary 4.8b *For any guard class C such that $T_k \subseteq C$, with $k \geq 0$ and $j > 0$, $\lfloor n/(k+2j+1) \rfloor$ guards of C are sufficient to guard (using L_j-visibility) any polygon P of n vertices.*

PROOF We let $k' = k + 2(j-1)$, and apply Corollary 4.8a (with the k of Corollary 4.8a equal to k') to show that some set Γ of $\lfloor n/(k'+3) \rfloor = \lfloor n/(k+2j+1) \rfloor$ $T_{k'}$-guards are sufficient to guard P, using L_1-visibility.

We claim that $VP_1(T)$, where T is any $T_{k'}$-guard, is contained in some $VP_j(S)$, where S is a $T_{k'-2(j-1)}$-guard (i.e., a T_k-guard). We prove the claim by induction on j, with k' fixed. If $j = 1$, then $k' = k$, and the claim is obvious. Therefore we assume that the claim holds for all $j' < j$, and show that it holds for j (i.e., that there is an S such that $VP_1(T) \subseteq VP_j(S)$). By the induction hypothesis, there is some S' which is a $T_{k'-2(j-2)}$-guard (i.e., a T_{k+2}-guard) such that $VP_1(T) \subseteq VP_{j-1}(S')$, hence the desired result follows if we show that there is some S such that $VP_{j-1}(S') \subseteq VP_j(S)$, for any T_{k+2}-guard S'.

Let $D(S')$ be a D_{k+2}-tree in a triangulation graph of P such that the embedding of the vertices of $D(S')$ is S'. Then, let $D' = (D(S')$ with all of its leaves removed); D' has diameter at most k, hence is a D_k-tree. If we let S be the embedding of the vertices of D', then S is a T_k-guard. Because $k \geq 0$, the vertices adjacent to leaves in $D(S')$ are in D', implying that $S' \subset VP_1(S)$. This in turn implies that $VP_{j-1}(S') \subseteq VP_j(S)$, proving the claim.

The claim implies that for every $T \in \Gamma$, we can find some S which is a T_k-guard which sees everything that T sees. Hence, $\lfloor n/(k+2j+1) \rfloor$ T_k-guards, using L_j-visibility, suffice. Thus $\lfloor n/(k+2j+1) \rfloor$ guards of $\mathcal{C}$ suffice. □

Corollary 4.8c *For any integer $k > 1$, $\lfloor n/(k+1) \rfloor$ guards of $\mathcal{L}_k$ are sufficient to partition any polygon P of n vertices.*

Before giving the proof of this corollary, we remark on two methods of proof which do not yield satisfactory results. The first, most obvious, method would be to triangulate, dominate the triangulation, and assign each triangle to any one of the trees which it is incident on. This method is incorrect, as the region assigned to a tree may not be connected. The second unsatisfactory method is a modification of the first; we assign triangles as before, but we also assign each edge of each tree to that tree's region. This method does yield connected regions, but the regions will be groups of triangles connected by line segments. We therefore also reject this method, as we can show that it is possible to find a collection of *polygons* that partition the given polygon.

PROOF First, let T be any triangulation of P. Next, let $k' = k - 2$, and apply Theorem 4.8 (with the k of Theorem 4.8 equal to k') to show that some set $\Gamma = \{I_1, I_2, \ldots, I_i\}$ of $i = \lfloor n/(k'+3) \rfloor = \lfloor n/(k+1) \rfloor$ $T_{k'}$-guards are sufficient to dominate T.

We will construct a region R_I for each $I \in \Gamma$. Initially, let each R_I be empty. For each triangle in T, we do the following:

Assume that the triangle has vertices a, b, and c. We let m_{ab}, m_{ac}, and m_{bc} be the midpoints of edges $\overline{ab}$, $\overline{ac}$, and $\overline{bc}$, respectively. Also, let m_{abc} be the center of gravity of m_{ab}, m_{ac}, and m_{bc}. We use the notation $T(v)$ to indicate which tree (element of Γ) is incident on vertex v, and $R_I \uplus Q$ to indicate that the current R_I is to be replaced by $R_I \cup Q$.

We examine three cases, based on the number of trees incident on the triangle.

case 1: There is one tree incident on $\triangle abc$. Without loss of generality, assume it is incident on vertex a, and let $R_{T(a)} \uplus \triangle abc$.

case 2: There are two trees incident on $\triangle abc$. We divide into two subcases:

case 2a: Both incident trees contain only one vertex of $\triangle abc$; without loss of generality let these vertices be a and b. Then, let $R_{T(a)} \uplus \triangle am_{ab}c$, and $R_{T(b)} \uplus \triangle bm_{ab}c$ (see figure 4.8a).

case 2b: One tree contains two vertices of $\triangle abc$; without loss of generality let these be a and b. Then, let $R_{T(a)} \uplus \square abm_{bc}m_{ac}$, and $R_{T(b)} \uplus \triangle cm_{bc}m_{ac}$ (see figure 4.8b).

case 3: There are three trees incident on $\triangle abc$. In this case, we let $R_{T(a)} \uplus \square am_{ab}m_{abc}m_{ac}$, $R_{T(b)} \uplus \square bm_{bc}m_{abc}m_{ab}$, and $R_{T(c)} \uplus \square cm_{ac}m_{abc}m_{bc}$ (see figure 4.8c).

The R_I's now partition P, as each part of each triangle has been placed in an R_I. An example of this is shown in figure 4.9.

Note that for any $T_{k'}$-guard I in Γ, if two vertices a, b are adjacent in I, then the segment $\overline{ab}$ is in R_I. Thus, between any two vertices u and v of I, there is a

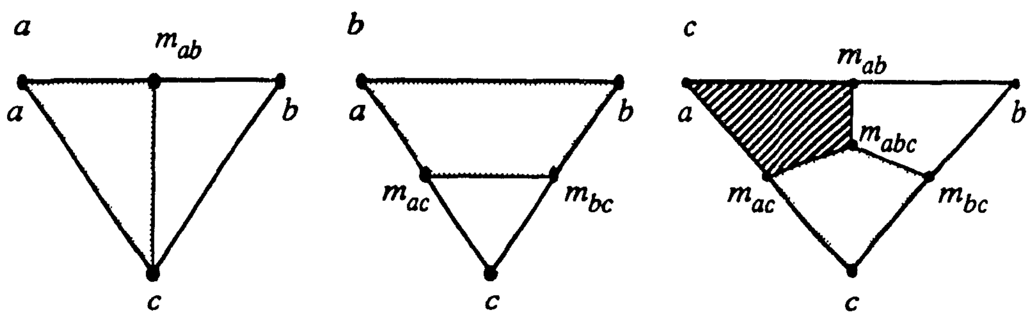

Figure 4.8: Partitioning Δabc

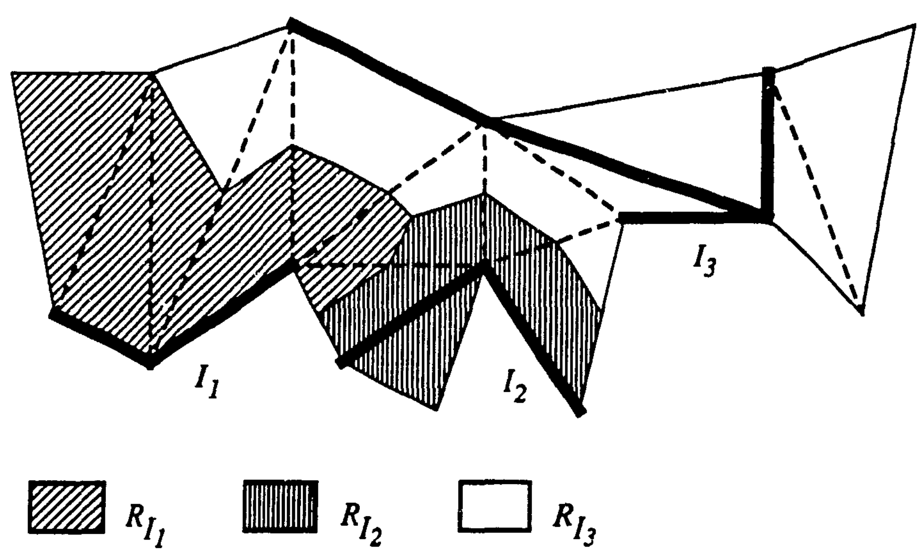

Figure 4.9: A sample partiton

link-k' path in R_I (just follow the edges of the $D_{k'}$-tree underlying I). Also, every point $x \in R_I$ is L_1-visible to a vertex $v(x)$ of I, as x is in some triangle U of the triangulation, and $R_I \cap U$ is convex and includes a vertex of I.

This means that between any two points $x, y \in R_I$, there is a link-k (recall that $k = k' + 2$) path in R_I (namely, the one from x to $v(x)$ to $v(y)$ to y), hence R_I is L_k-convex. Thus, the R_I's, which are a set of $\lfloor n/(k+1) \rfloor$ regions that partition P, are in guard class $\mathcal{L}_k$. □

Corollary 4.8d *For any guard class $\mathcal{C}$ such that $\mathcal{L}_k \subseteq \mathcal{C}$, with $k > 1$, $\lfloor n/(k+1) \rfloor$ guards of $\mathcal{C}$ are sufficient to cover any polygon P of n vertices.*

PROOF By Corollary 4.8c, there is a set of $\lfloor n/(k+1) \rfloor$ guards of $\mathcal{L}_k$ which partition P. As each guard in class $\mathcal{L}_k$ is contained in a guard of class $\mathcal{C}$, there must be a set of $\lfloor n/(k+1) \rfloor$ guards in $\mathcal{C}$ which cover P as desired. □

Note that Theorem 3.1 and Corollary 4.8c together imply that $\lfloor n/(k+1) \rfloor$ L_k-convex regions are sometimes necessary and always sufficient to partition or cover a polygon of n vertices, for $k > 1$.

Corollary 4.8e *There are no polygons of n vertices with a L_j-hidden set of size larger than $\lfloor n/(j+1) \rfloor$, for all $j > 1$ and $n \geq j + 1$.*

PROOF By Corollary 4.8c, any such polygon has a cover by L_j-convex regions of size $\lfloor n/(j+1) \rfloor$. Since no two elements of a L_j-hidden set can lie in a single L_j-convex region, the maximum hidden set is of size at most $\lfloor n/(j+1) \rfloor$. □

Note that Theorem 3.1 and Corollary 4.8e show that, for points, $h_j(n) = \lfloor n/(j+1) \rfloor$, for $j > 1$.

We have also now proved Theorem 4.1; it is a direct combination of Corollaries 3.1a and 4.8b. The following is a table of some of the consequences of this theorem for the art gallery problem. Note that the results for $j = 1$ and $k = 0$ and 1 are the known art gallery results for simple polygons [C75] [F78] [O83a] [O87].

k	j (L_j-visibility)	$g(n)$	$\mathcal{T}_k(P)$	$\mathcal{L}_k(P)$	other interesting classes of guards between $\mathcal{T}_k$ and $\mathcal{L}_k$
0	1	$\lfloor n/3 \rfloor$	vertices	points	
1	1	$\lfloor n/4 \rfloor$	vertex pair	convex	diagonals, line segments
2	1	$\lfloor n/5 \rfloor$	graph star vertices	L-convex	graph star, fan, star-shaped
k	1	$\lfloor n/(k+3) \rfloor$		L_k-convex	
0	j	$\lfloor n/(2j+1) \rfloor$	vertices	points	
1	j	$\lfloor n/(2j+2) \rfloor$	vertex pair	convex	diagonals, line segments

A *vertex pair guard* is a pair of vertices which are connected by a diagonal. A *graph star vertex guard* is the vertex set of a graph-theoretic star (tree with one non-leaf node). A *fan* is a star-shaped region with a vertex in the kernel.

4.6 Properties of Polygon Exteriors

We can use results of the previous two sections to get almost-tight bounds on visibility properties of polygon exteriors.

Given a polygon P, we can rotate P so that there is one uniquely highest vertex a. We can then place two points l and r to the left and right of P, below P's lowest vertex, and distant enough from P so that they both are L_1-visible to v. Let P^+ denote the set $P \cup l \cup r$, and $CH(P^+)$ denote its convex hull We define an *exterior triangulation* of P as a triangulation of the region interior to $CH(P^+)$ but exterior to P, for any such placement of l and r (see figure 4.10).Note that an exterior triangulation graph is not a triangulation graph.

Exterior T_k-guards are then defined as the geometric embedding of the vertices of some D_k-subtree in an exterior triangulation. We let $\mathcal{T}_k^e$ represent the guard class of exterior T_k-guards.

Theorem 4.9 *For any guard class C such that $\mathcal{T}_k^e \subseteq C$, $k \geq 0$, $\lfloor (n+3)/(k+3) \rfloor$ guards of C are sufficient to guard (using L_1-visibility) the exterior of any polygon P of n vertices.*

PROOF The following proof is a modification of the (special case $k = 0$) proof presented by Aggarwal and O'Rourke [O87].

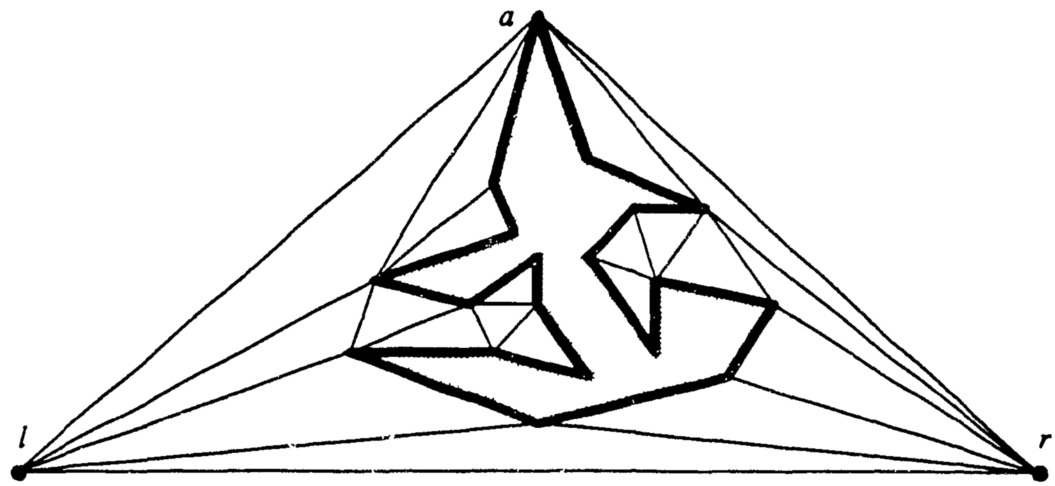

Figure 4.10: Exterior triangulation

Let T be an exterior triangulation of P as defined above and split the vertex a into two vertices a and a' so that the resulting graph T' is a (interior) triangulation graph (as shown in figure 4.11).

Let E_1, E_2, and E_3 be the edges of $CH(P^+)$. Each E_i is contained in some triangle U_i of T'. Let r_a, r_l, and r_r be the rays that bisect the exterior angles of $CH(P^+) \setminus \bigcup_{i=1}^{3}$ interior(U_i) at a, l, and r. These rays divide the exterior of $CH(P^+)$ into three regions S_1, S_2, and S_3, such that for any i, $S_i \cup U_i$ is convex. This construction is illustrated in figure 4.12.

We now dominate T', which has $n + 3$ vertices, with a set Γ' of $\lfloor (n + 3)/(k + 3) \rfloor$ D_k-trees, by Theorem 4.8. We claim that the entire exterior of P is seen by the set Γ of the exterior T_k-guards which are the embeddings of the elements of Γ'. We examine an arbitrary exterior point p:

If $p \in CH(P^+)$, then it is in some triangle of T'. Since each such triangle has an element of Γ on at least one vertex, p is seen by some guard.

If $p \notin CH(P^+)$, then it lies in some region S_i. $U_i \cup S_i$ is not only convex but also empty, as E is not a polygon edge. Therefore, p is seen by some element of Γ, as U_i has such an element incident on at least one vertex.

Therefore, Γ, a set of $\lfloor (n + 3)/(k + 3) \rfloor$ guards in class T_k^e, sees the entire exterior.

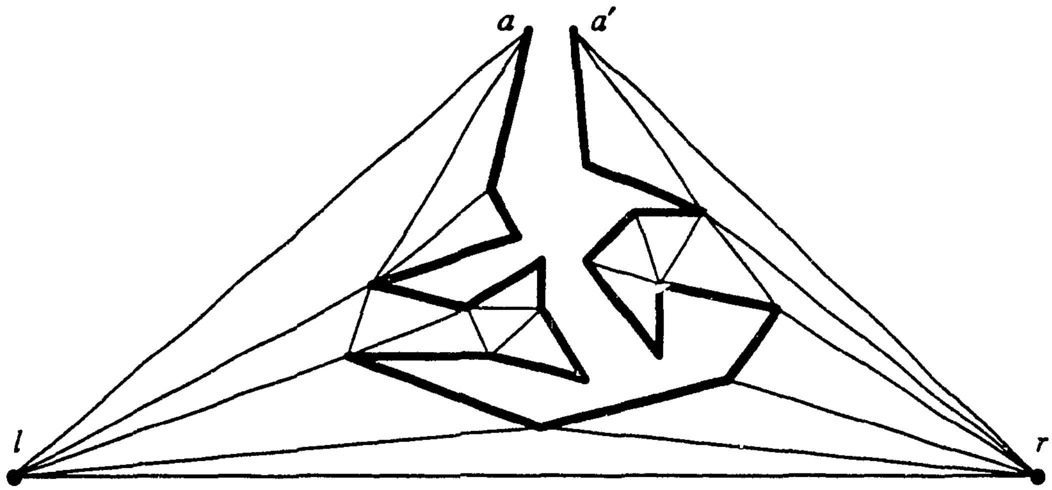

Figure 4.11: Splitting a vertex

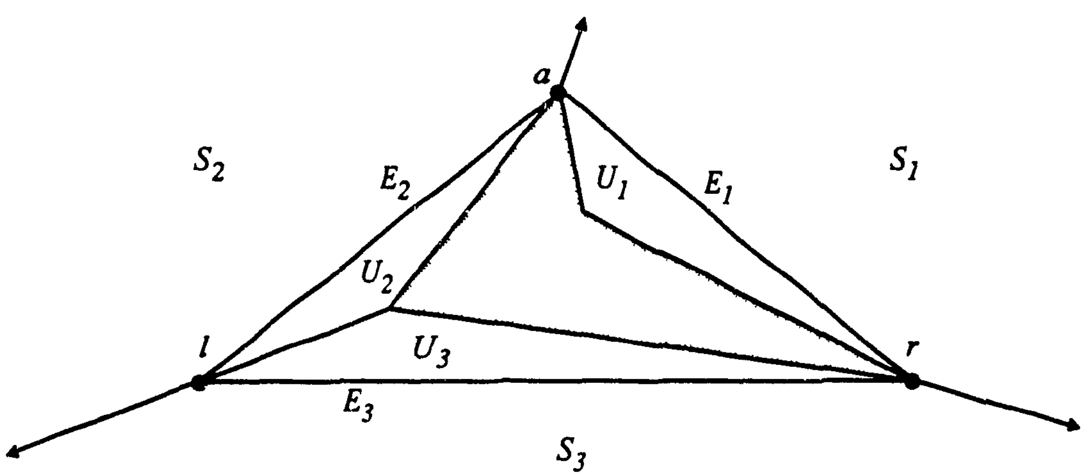

Figure 4.12: Regions exterior to the hull

As each guard of class $\mathcal{T}_k^e$ is contained in some guard of class $\mathcal{C}$, there is a set of $\lfloor (n+3)/(k+3) \rfloor$ guards of class $\mathcal{C}$ which sees the entire exterior. □

Corollary 4.9a *For any guard class $\mathcal{C}$ such that $\mathcal{T}_k^e \subseteq \mathcal{C}$, with $k \geq 0$ and $j > 0$, $\lfloor (n+3)/(k+2j+1) \rfloor$ guards of $\mathcal{C}$ are sufficient to guard (using L_j-visibility) the exterior of any polygon P of n vertices.*

PROOF The argument is identical to that of Corollary 4.8b, except that we start with Theorem 4.9 rather than Corollary 4.8a. □

Corollary 4.9b *For any integer $k > 1$, $\lfloor (n+3)/(k+1) \rfloor$ guards of $\mathcal{L}_k^e$ are sufficient to partition the exterior of any polygon P of n vertices.*

PROOF The proof is essentially the same as that of Corollary 4.8c, except that we must place the points of the regions S_1, S_2, and S_3 in the R_I's. First, we divide the region inside $CH(P^+) \setminus P$ as in Corollary 4.8c.

Next, for each S_i, let U_i be $\triangle abc$, with a and b the vertices of E_i. If the entire edge $\overline{ab}$ is contained in some R_V, then let $R_V \uplus S_i$. Otherwise, the edge $\overline{ab}$ is split between two regions R_V and R_W, with a in R_V and b in R_W. By the construction, $R_V \cap U_i$ and $R_W \cap U_i$ will share some edge F (either $\overline{cm_{ab}}$, $\overline{m_{ac}m_{ab}}$, $\overline{m_{bc}m_{ab}}$, or $\overline{m_{abc}m_{ab}}$). Let r_{ab} be a ray with vertex m_{ab} which in colinear with F and extends into S_i (see figure 4.13 for an example). r_{ab} divides S_i into two pieces S_a and S_b, with $a \in S_a$ and $b \in S_b$. At least one, and possibly both, of S_a and S_b are unbounded. Let $R_V \uplus S_a$, and $R_W \uplus S_b$.

Application of this procedure to each S_i yields a set of R_I's which partition the entire exterior of P, and are in class $\mathcal{L}_k^e$. □

Corollary 4.9c *There are no polygons with an exterior L_j-hidden set of size larger than $\lfloor (n+3)/(j+1) \rfloor$, for all $j > 1$.*

PROOF The argument is identical to that of Corollary 4.8e, except that we start with Corollary 4.9b rather than Corollary 4.8c. □

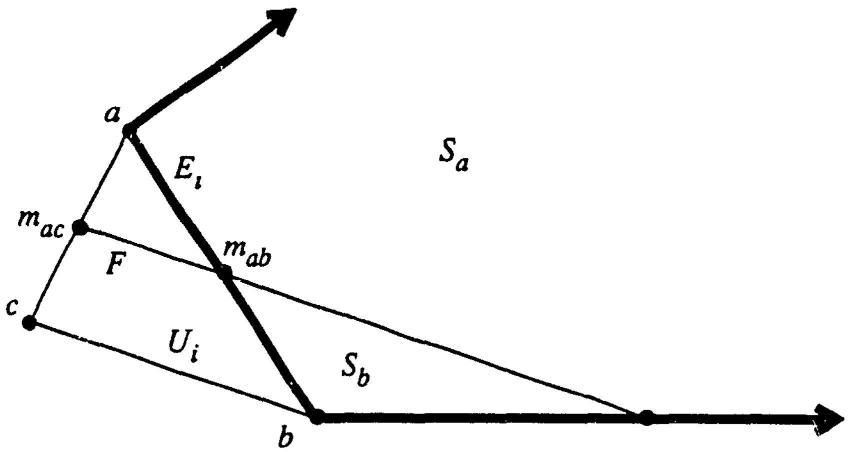

Figure 4.13: Partioning S_i

The above theorem and its corollaries are *almost-tight*: the lower bounds presented in the previous chapter are the same except that the numerator of the fraction in the floor is $n+1$ rather than $n+3$; this causes the resulting integers to differ by at most 1. Although the $k=0, j=1$ bound has an $n+2$ numerator, we do not expect this to generalize to larger k or j, as for $k > 0$ or $j > 1$ only one guard is needed to guard the exterior of a convex polygon (as compared to two for $k=0, j=1$). We therefore conjecture the following:

Conjecture 4.10 *For any guard class $\mathcal{C}$ such that $\mathcal{T}_k^e \subseteq \mathcal{C}$, $k > 0$, $\lfloor (n+1)/(k+3) \rfloor$ guards of $\mathcal{C}$ are sufficient to guard (using L_1-visibility) the exterior of any polygon P of n vertices.*

Conjecture 4.11 *For any guard class $\mathcal{C}$ such that $\mathcal{T}_k^e \subseteq \mathcal{C}$, with $k > 0$ and $j > 0$, $\lfloor (n+1)/(k+2j+1) \rfloor$ guards of $\mathcal{C}$ are sufficient to guard (using L_j-visibility) the exterior of any polygon P of n vertices.*

Conjecture 4.12 *For any integer $k > 1$, $\lfloor (n+1)/(k+1) \rfloor$ guards of $\mathcal{L}_k^e$ are sufficient to partition the exterior of any polygon P of n vertices.*

Conjecture 4.13 *There are no polygons with an exterior L_j-hidden set of size larger than $\lfloor (n+1)/(j+1) \rfloor$, for all $j > 1$.*

The simple proof of Aggarwal, O'Rourke, and Shermer [O87] for the $k = 0$ tight bound does not easily generalize to arbitrary k, for two reasons: First, their proof uses three-coloring in a manner similar to Fisk's proof of the original art gallery theorem. Second, their proof uses a restructuring of an exterior triangulation, by "flipping a diagonal" in a convex quadrilateral; this restructuring would need to be much more complex for higher k. However, a generalization of Fisk's proof (and of 3-coloring) has been found for L_j-visibility; so some hope of generalizing their proof remains.

Chapter 5

Computational Complexity

In this chapter, we will show that the optimization and decision problems for covers, guardings, and hidden sets are NP-hard. We present two fundamentally different constructions to obtain these results; one is a tranformation from Boolean 3-Satisfiability, and the other is a transformation from Exact Cover by 3-Sets.

We begin this chapter with a section on the formal definitions of the problems that we consider, and a section of remarks applying to all proofs. The sections following that are the constructions and proofs for our problems.

5.1 Problem Definitions

The first problem that we will be dealing with is the problem of determining if a polygon admits an L_j-convex cover of a given size. This is called the L_j-Convex-Cover problem:

L_j-CONVEX COVER (L_jCC)

INSTANCE: A polygon P, and an integer m.

QUESTION: Can P be covered by m or fewer L_j-convex sets?

We may also ask the minimization problem:

MINIMUM L_j-CONVEX COVER (MLjCC)

INSTANCE: A polygon P.

QUESTION: What is the smallest m such that P can be covered by m L_j-convex sets?

If a polynomial algorithm existed to solve **LjCC**, we could solve **MLjCC** in polynomial time as well: we would simply solve **LjCC** for values of m from 1 to $\lfloor n/(m+1) \rfloor$ (or to $n-2$ for $m=1$). The lowest value of m for which the **LjCC** problem has a yes answer would be the answer to the **MLjCC** instance (by Corollary 4.8c). Also, a polynomial algorithm for **MLjCC** would trivially provide a polynomial algorithm for **LjCC**. We therefore restrict our attention to the decision problem.

The situation for the other problems that we consider in this chapter is similar: there are equally powerful decision and optimization versions of the problem. In all instances we will consider only the decision problem.

It is often the case that even and odd link-diameters must be handled by separate cases. Our proof for **LjCC** is no exception; we must prove our result in two parts: one for the odd j's, and one for the even j's. For each of these two cases, we present a base case ($j=1$ or $j=2$), and a modification to the base case for larger j.

After our proof for **LjCC**, we consider covering polygons with L_j-star-shaped polygons, giving rise to the following problem:

L_j-STAR COVER (LjSC)

INSTANCE: A polygon P, and an integer m.

QUESTION: Can P be covered by m or fewer L_j-star-shaped sets?

L1SC is also known as Star Cover (or Point Guard), which was proved NP-hard by Lee, Lin, and Aggarwal ([LL86] [A84]) We will prove that **LjSC** is NP-hard by a modification of our proof of **LjCC** for even j.

Note that **LjCC** and **LjSC** are the two extremes of the general link-guarding problem:

LINK$_{j,k}$-GUARDING ($L_{j,k}G$)

INSTANCE: A polygon P, and an integer m.

QUESTION: Is there a collection C of m or fewer L_k-convex subpolygons of P such that P is covered by the link-j visibility polygons of the elements of C?

L_jCC is the same as **$L_{0,j}$G**, and **L_jSC** is the same as **$L_{j,0}$G**. We show that a modification of our proof for **L_jCC** will prove that **$L_{j,k}$G** is NP-hard (although it will not be necessary to modify our *construction*).

We will prove our NP-hardness results by transformation from two well-known NP-complete problems, Exact Cover by 3-Sets, and Boolean 3-Satisfiability (see [K72] or [GJ79]).

EXACT COVER BY 3-SETS (X3C)

INSTANCE: A finite set $X = \{X_1, X_2, \ldots, X_{3q}\}$, and a collection $C = \{C_1, C_2, \ldots, C_n\}$ of 3-element subsets of X.

QUESTION: Does C contain an *exact cover* for X: A subcollection $C' \subseteq C$ such that every X_i appears in exactly one member of C'?

BOOLEAN 3-SATISFYABILITY (3SAT)

INSTANCE: A finite set $U = \{U_1, U_2, \ldots, U_q\}$ of boolean variables and a collection $C = \{C_1, C_2, \ldots, C_n\}$ of 3-literal clauses on U.

QUESTION: Is there a truth assignment for U that satisfies all of the clauses in C?

Following our covering and guarding proofs, we note that our methods can be applied to the existing proofs for the NP-hardness of several problems relating to hidden sets, establishing these problems, using L_j-visibility, as NP-hard for odd j. The definitions of the hidden set problems considered are given in that section.

We end the chapter with a discussion of the comparative complexity of graph-theoretic problems and polygon visibility problems.

5.2 General Remarks

In this section, we give some general discussion on the complexity of the problems and transformations that we consider.

We will prove our NP-hardness results by using *component-design* transformations from **X3C** and **3SAT**. This means that we will construct geometric components (portions of polygons) which correspond to the elements of the **X3C** or **3SAT** problem.

The first matter which we wish to address is whether or not our problem transformations can be accomplished in polynomial time (polynomial in the size of the input **X3C** or **3SAT** instance). Our constructions all use a central rectangle, with many vertices located at integer coordinates on this rectangle. Each of these coordinates will take at most $O(\log n)$ bits to store. The remaining vertices of the transformation image polygon will be computable with a constant number of the following operations:

(1) Calculate the line between two points.

(2) Calculate the intersection of two lines.

(3) Find the midpoint of the line segment between two points.

Using rational computations, any one of these operations will result in a point location or a line equation which requires storage of at most two more than twice the number of bits of the input points or lines. Therefore, if the maximum height of a tree of these operations required to compute any vertex is c, and the points on the rectangle are expressible with b bits, then the resultant number of bits required to store any vertex will be $f(c)$, where

$$\begin{aligned} f(0) &= b, \text{ and} \\ f(x) &= 2 * f(x-1) + 2. \end{aligned}$$

The solution of this recurrence is:

$$f(x) = (b+2)2^x - 2$$

Thus, as c is constant, and b is $O(\log n)$, 2^c is a constant, and $f(c)$ is $O(\log n)$. Therefore, the number of bits required to store any vertex will be $O(\log n)$. Furthermore, each of the above operations can be accomplished in polynomial time in the number of input bits. Therefore each vertex, and the entire image polygon, can be computed in polynomial time. Thus, our problem transformations will take polynomial time.

The other concern that we want to address is the upper bound on the complexity of the problems we consider. These problems are all decidable; O'Rourke has shown this for the **L_1CC** problem [O82c], and methods similar to his can be used on any of our problems. It is difficult to determine whether or not our problems are in NP; it is suspected that there are polygon classes such that the height of a tree of line intersection/line determination calculations necessary to compute a vertex of the minimum cover increases with the size of the polygon [O82a] [O82b]. The recurrence discussed above, if the upper bound on storage that it represents is tight, indicates that a linear increase in the height of a tree of such calculations required to find a vertex of the minimum cover would reflect itself exponentially in the storage and time required to compute the cover. Thus, the existence of a polygon class exhibiting linear increase in the calculation tree height would suggest that the cover problem is not in NP. For a discussion of this and other related questions regarding the complexity of covering problems, the reader is referred to [O82a].

5.3 L_jCC: Odd j

We start our NP-hardness proofs with a proof that **L_1CC** (also known simply as **Convex Cover**) is NP-hard.

5.3.1 Problem Transformation

We will prove this result by transformation from **X3C**. Given an instance $I = (X, C)$ of **X3C**, we construct an instance $\psi_1(I) = (P, m)$ of **L_1CC** as follows:

First, we let $m = 2q + n + 1$ (q and n are from the definition of **X3C**).

We construct P as follows: we start with a rectangle (called the central rectangle),

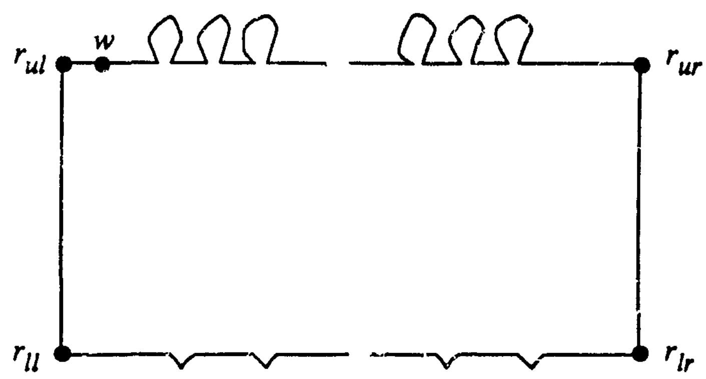

Figure 5.1: $\mathbf{L_1CC}$ Construction overview

to which we will connect structures corresponding to the X_u's (which we call *X-units*) along the bottom edge, and structures corresponding to the C_v's (which we call *C-units*) along the top edge (see figure 5.1). XU_u denotes the X-unit corresponding to X_u, and CU_v denotes the C-unit corresponding to C_v. Both X-units and C-units will be convex sets, and P will be the union of the X-units, C-units, and central rectangle.

Let r_{ul}, r_{ll}, r_{ur}, and r_{lr} be the upper-left, lower-left, upper-right, and lower-right vertices of the central rectangle. Also place a vertex w somewhere (anywhere) on the rectangle between r_{ul} and the leftmost C-unit.

Each X-unit is a 2-edge triangular notch, with its left edge colinear with w and right edge colinear with r_{ur}. (see figure 5.2). The X-units are evenly placed along the bottom of the central rectangle in order of increasing index. The three vertices of XU_u are called (from left to right) l_u, m_u, and r_u.

Let $C_v = \{X_A, X_B, X_C\}$ with $A < B < C$. The C-units for the C_v's are evenly placed on the top edge of the central rectangle in order of increasing index. A C-unit has seven vertices $(p_1, p_2, \ldots, p_7)$, attaching to the central rectangle at p_1 and p_7. p_3 is placed at the intersections of the lines $\overline{l_B p_1}$ and $\overline{r_C p_7}$, and p_5 is placed at the intersections of $\overline{r_B p_7}$ and $\overline{l_A p_1}$. p_2 and p_6 are placed colinear with $\overline{l_C p_1}$ and $\overline{r_A p_7}$, respectively. p_2, p_4, and p_6 are all placed so that $(p_1, p_2, \ldots, p_7)$ is convex, and such

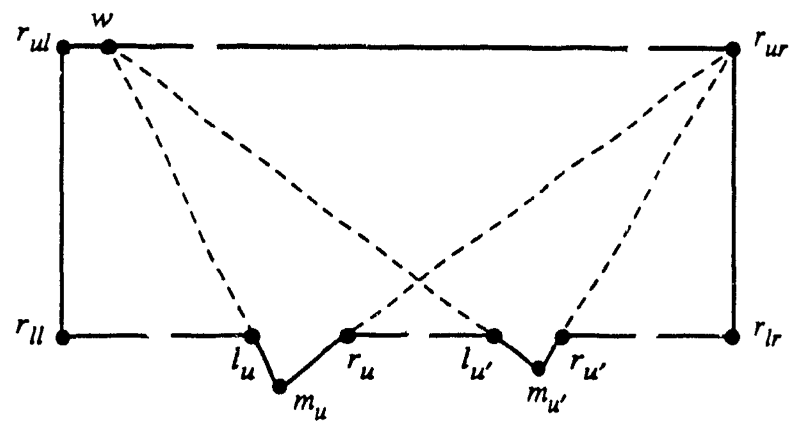

Figure 5.2: X-unit construction

that p_4 is not visible to m_A or m_C. This construction is illustrated in figure 5.3.

Given this basic structure, we now need to ensure that the segments where the X-units and C-units attach are small enough that we do not encounter either of the following two problems: (1) two C-units overlap, or (2) some convex set covering an m_u can cover significantly more of some C-unit if it includes only m_u rather than l_u, m_u, and r_u.

The first problem is handled by making the "gap distance" (distance between p_1 and p_7) for each C-unit very small, which will make the C-units themselves smaller. By considering the worst case that could happen ($C_v = \{X_1, X_2, X_u\}$ or $C_v = \{X_u, X_{3q-1}, X_{3q}\}$), and computing a gap distance small enough to keep the X-units from overlapping in these instances, we can guarantee that none of the X-units will overlap.

The second problem is handled by making the gap distance for the X-units (distance from l_u to r_u) smaller. We can do this by examining each C-unit in turn, and insuring that the gap distances for the concerned X-units are small enough that the following properties are satisfied:

1. Let q_1 be the intersection of $\overline{m_B p_1}$ and $\overline{p_2 p_3}$. Then q_1 must not be seen by m_{C+1}, if it exists.

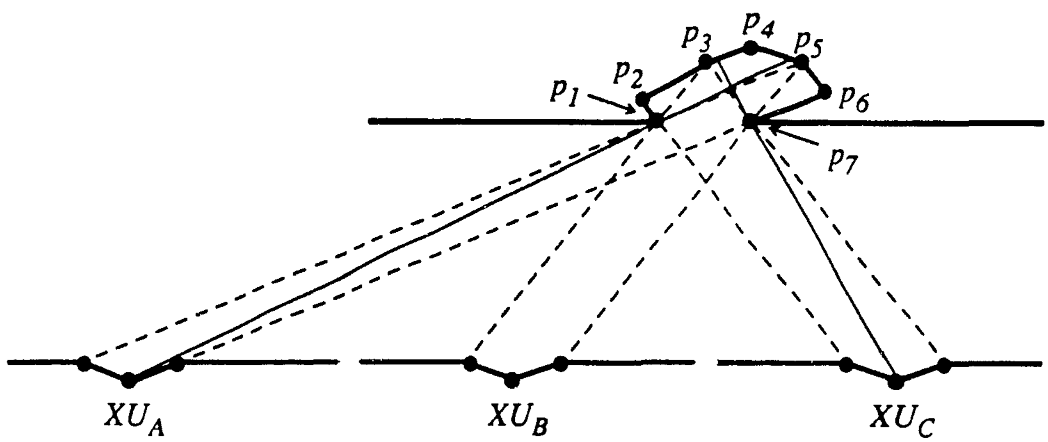

Figure 5.3: C-unit construction I

2. Let q_2 be the intersection of $\overline{m_C p_7}$ and $\overline{p_3 p_4}$. Then q_2 must not be seen by m_{B-1}.

3. Let q_3 be the intersection of $\overline{m_A p_1}$ and $\overline{p_4 p_5}$. Then q_3 must not be seen by m_{B+1}.

4. Let q_4 be the intersection of $\overline{m_B p_7}$ and $\overline{p_5 p_6}$. Then q_4 must not be seen by m_{A-1}, if it exists.

It is clear that, as the gap distances for XU_A, XU_B, and XU_C decrease, the points q_1, q_2, q_3, and q_4 draw closer to p_3, p_3, p_5, and p_5, respectively, and hence will be nonvisible as required.

We let q_1' be a point counterclockwise of and in the neighborhood of q_1 that cannot be seen by m_{C+1}. Such a point will exist, as m_{C+1} does not see q_1, and visibility polygons are closed regions. Similarly, let q_2' be a point clockwise of and in the neighborhood of q_2 that cannot be seen by m_{B-1}. Let q_3' and q_4' be defined symmetrically to q_2' and q_1'. Figure 5.4 illustrates this construction.

Figure 5.5 shows the full construction of P for the instance of **X3C** with $C = \{\{X_1, X_2, X_3\}, \{X_4, X_5, X_7\}, \{X_3, X_4, X_9\}, \{X_2, X_5, X_8\}, \{X_1, X_6, X_7\}, \{X_2, X_6, X_8\}, \{X_5, X_6, X_9\}\}$.

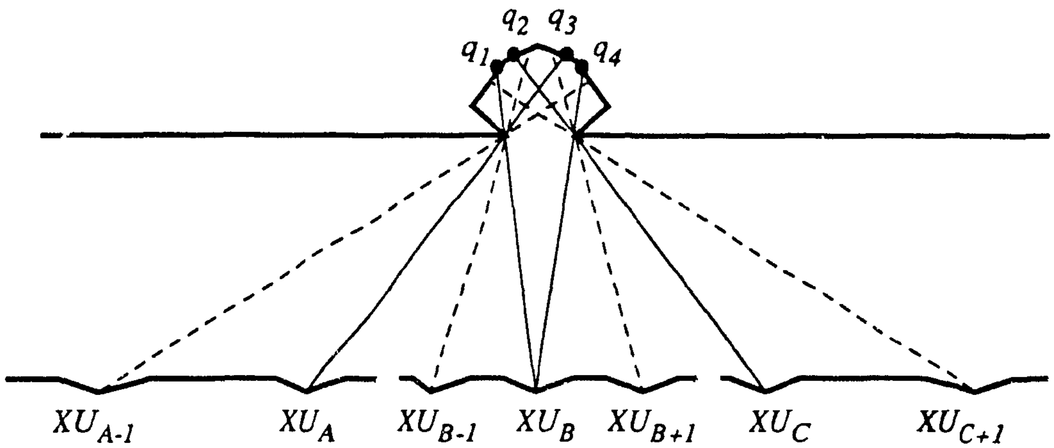

Figure 5.4: C-unit construction II

5.3.2 Properties of the Construction

Given a convex cover of P, we let S_u be a convex set of the cover which covers the vertex m_u. Such S_u's are called S-sets.

Our construction has the following important properties:

P1 The central rectangle can be covered by one convex set.

P2 Each X-unit XU_u can be covered by one convex set.

P3 The set of all m_u's plus r_{ul} form a hidden set.

P4 No convex set can help cover two C-units.

P5 No convex set containing r_{ul} can help cover any C-unit.

P6 Each C-unit CU_v will be coverable in three ways: either (a) by one convex set, (b) by four or more S-sets, or (c) by three S-sets, when the three S-sets are S_A, S_B, and S_C ($C_v = \{X_A, X_B, X_C\}$). Each of these S-sets are capable of covering the whole X-unit to which it corresponds in addition to the portion of CU_v which it covers.

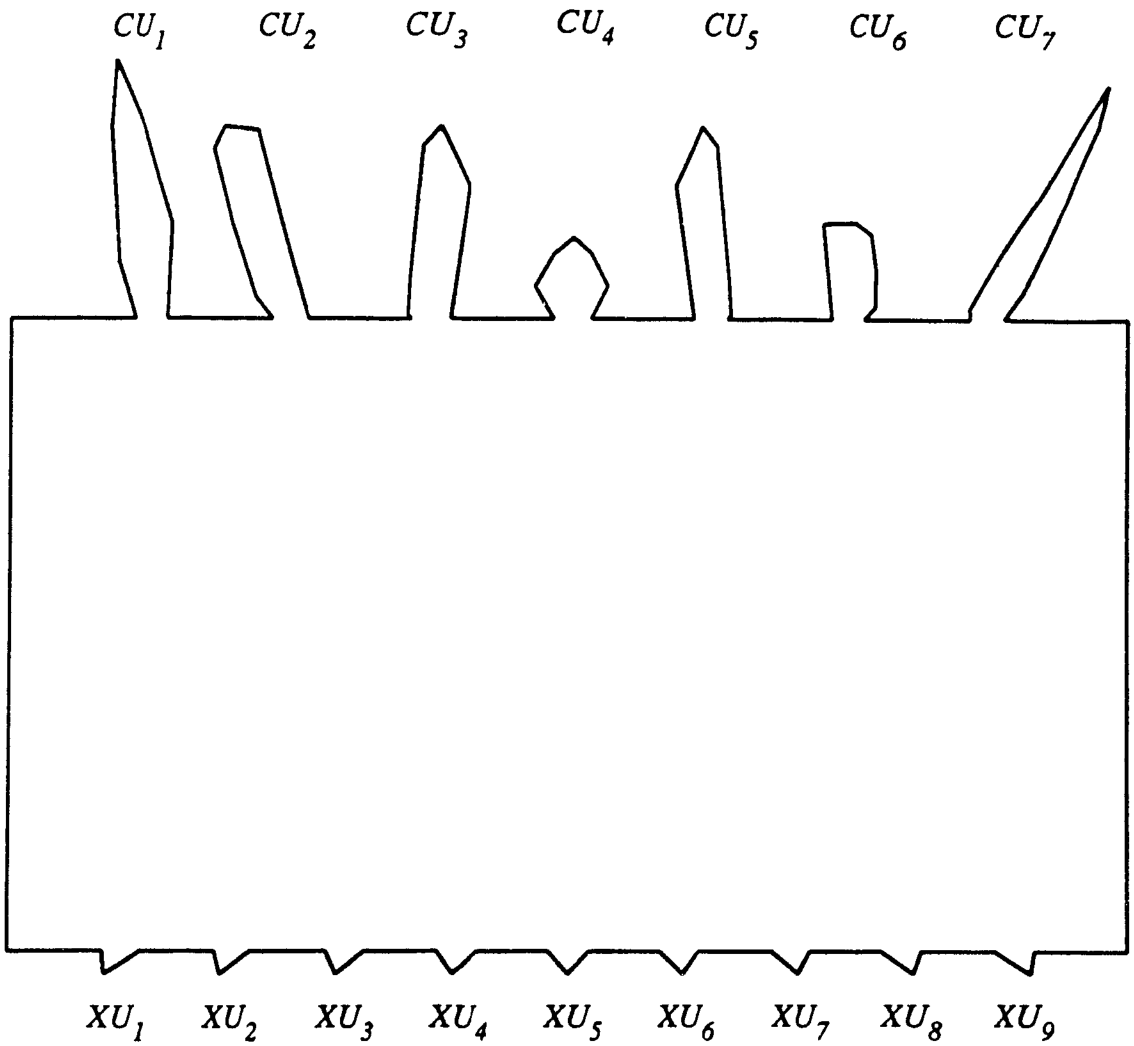

Figure 5.5: A Sample Construction for $\mathbf{L_1CC}$

Property **P1** will be satisfied, as we have cut no pieces off of our central rectangle.

As every triangle is convex, property **P2** is satisfied. Also, the set consisting of all m_u's along r_{ul} is a hidden set (property **P3** is satisfied).

As each C-unit is a convex set attached to the top of the central rectangle, no one convex subset of P can contain points from two C-units. Thus, property **P4** is satisfied.

Because r_{ul} lies along the upper edge of the rectangle, no convex set can contain both r_{ul} and any point of any CU_v (property **P5** is satisfied).

The following two lemmas help us establish property **P6**.

Lemma 5.1 *No C-unit can be covered by any two or fewer S-sets.*

PROOF Assume the contrary: some CU_v is covered by two S-sets, S_g and S_h. Without loss of generality, assume that S_g covers vertex p_6 of the C-unit. Then, by construction, $g \leq A$; this means that S_g can cover neither vertex p_4 nor vertex p_2. So S_h must cover p_2; then $h \geq C$, and S_h cannot cover p_4. Therefore, p_4 is not covered, which is a contradiction. Thus, the lemma holds. □

Lemma 5.2 *A C-unit CU_v can be covered by three S-sets iff the S-sets are S_A, S_B, and S_C.*

PROOF Assume that we have three S-sets S_a, S_b, and S_c covering CU_v. By the argument given in the proof of Lemma 5 1 we must have $a \leq A$ and $c \geq C$.

Assume that we have $c > C$. Then, q_1' is not covered by S_c. To cover q_1', we must have $b > B$. But then S_b would not cover q_3'; furthermore this point is not covered by S_a. We are thus not covering the C-unit. Therefore we must have $c = C$, and, symmetrically, $a = A$.

Furthermore, if $b \neq B$, then either q_2' or q_3' is not covered. Therefore, $b = B$, so that the only three S-sets which can cover the C-unit are S_A, S_B, and S_C. □

Lemmas 5.1 and 5.2 together with the convexity of the C-units imply that the construction has property **P6**.

Thus, the construction has all of the given properties, and we now procede with the proof of our theorem.

5.3.3 L_1CC

Theorem 5.3 L_1CC *is NP-hard.*

PROOF We show that the instance I of **X3C** will have a yes answer iff the instance $\psi_1(I)$ of $\mathbf{L_1}$CC has a yes answer (i e., P can be covered by $m = 2q + n + 1$ (L_1-) convex sets).

If the instance I of **X3C** has a yes answer, then we use the following cover for P: Let R cover the central rectangle. We choose S_u ($1 \leq u \leq 3q$) corresponding to the exact cover (via property **P6c**); each S_u covers XU_u and part of a C-unit, and q of the C-units are thus covered We have so far used only $3q + 1$ convex sets. For each of the remaining $n - q$ C-units, we cover each with its own convex set (by property **P6a**). Thus, we have a covering with $3q + 1 + (n - q) = 2q + n + 1$ convex sets.

We now assume that the instance of $\mathbf{L_1}$CC has a yes answer (we have covered P with $2q + n + 1$ convex sets).

Each m_u ($1 \leq u \leq 3q$), and the vertex r_{ul}, must be covered by at least one convex set. Let S_u be any of the sets covering m_u, and R be a set covering r_{ul}. By property **P3**, these sets must be distinct. Thus, in our covering, we have R, the S_u's and only $n - q$ other sets. Therefore, by properties **P6**, **P4**, and **P5**, at least q of the C-units were covered by the S-sets. Since no C-unit is coverable by 2 or fewer such sets, the only way we can cover this many C-units with S-sets is to have exactly q C-units covered with exactly 3 S-sets each.

However, the only covering for a C-unit by exactly 3 convex sets is by the convex sets contributed by the X-units corresponding to that C-units' members (property **P6c**). As no X-unit can contribute its set to more than one C-unit (property **P4**), the q covered C-units correspond to an exact cover for X. Therefore, the **X3C** instance has a yes answer. □

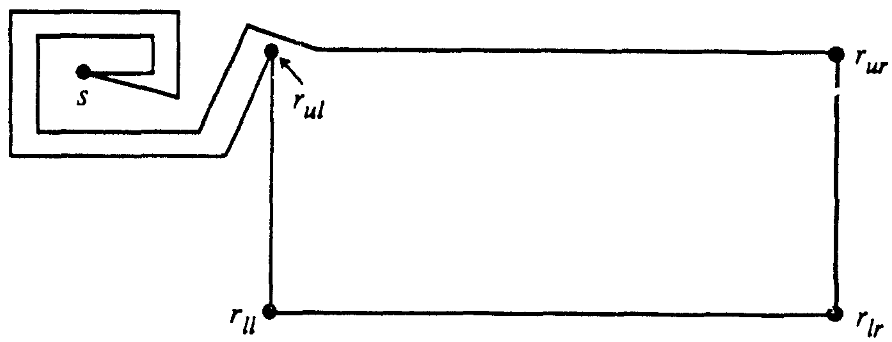

Figure 5.6: Central unit for $j = 7$

5.3.4 Extension to Higher Odd j

Theorem 5.4 *For any odd integer $j \geq 1$,* L$_j$CC *is NP-hard.*

PROOF The proof is similar to that of Theorem 5.3, with the units and the central rectangle slightly modified.

The necessary modifications to the units are as follows:

The central rectangle is changed to a "central unit," which is a rectangle with a spiral of $j - 1$ arms added at r_{ul}. We let s be the vertex at the end of the spiral, and the spiral arm connects to the rectangle so that $VP_{j-1}(s)$ intersects the rectangle only at r_{ul}. Figure 5.6 illustrates these definitions for $j = 7$.

We change each X-unit by adding a spiral of $(j - 1)/2$ arms at m_u. We let the m_u^* be the vertex at the end of the spiral. This is illustrated in figure 5.7 for $j = 7$.

We change each C-unit by adding several spirals of $(j - 1)/2$ arms: one each at p_2, p_4, p_6, q_1', q_2', q_3', and q_4'; the vertices at the ends of these spiral arms are p_2^*, p_4^*, p_6^*, q_1^*, q_2^*, q_3^*, and q_4^*, respectively. We let V_v^* be the set of these vertices at the end of the spirals on CU_v. This is illustrated in figure 5.8 for $j = 7$.

The spirals are shown schematically in figure 5.8. The actual geometry of these schematic representations are shown in figure 5.9 for the spirals on vertices (e.g., p_2),

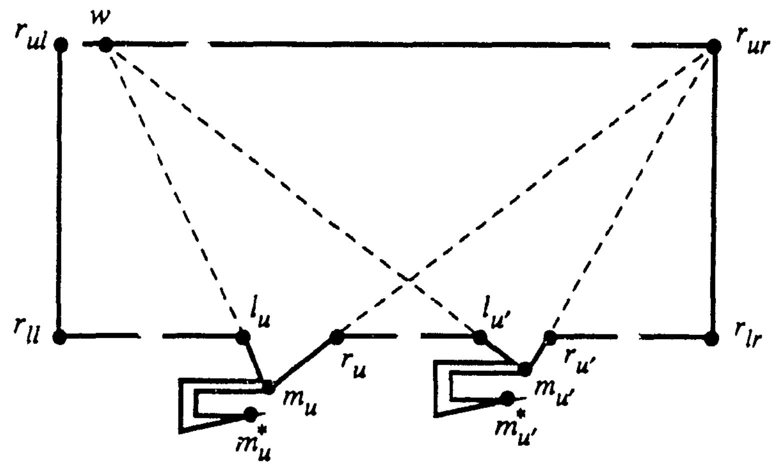

Figure 5.7: X-unit for $j = 7$

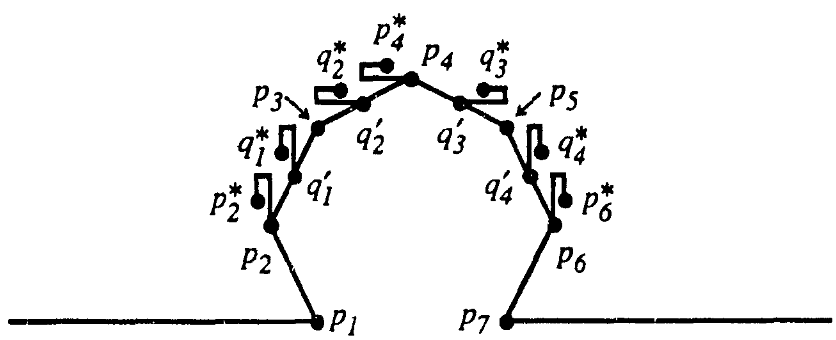

Figure 5.8. C-unit for $j = 7$

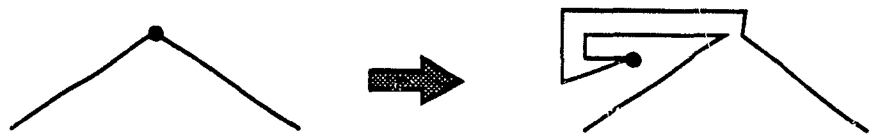

Figure 5.9: A spiral on a vertex

and in figure 5.10 for the spirals on intersection points (e.g., q_3).

The attachment of a spiral to a C-unit is made so that the link-$((j-3)/2)$ visibility polygon of the vertex at the end of the spiral is a small region containing the attachment point.

Given a L_j-convex cover of P, we let S_u be a L_j-convex set of the cover which covers the vertex m_u^*, and call such S_u's S-sets. We can then show the following properties:

P1′ The central unit can be covered by one L_j-convex set

P2′ Each X-unit XU_u can be covered by one L_j-convex set.

P3′ The set of all m_u^*'s plus s form a hidden set.

P4′ No L_j-convex set can cover elements of both V_g^* and V_h^* for $g \neq h$.

P5′ No L_j-convex set containing s can cover any element of V_h^* for any h.

P6′ Each C-unit CU_v will be coverable in three ways: either (a) by one L_j-convex set, (b) by four or more S-sets, or (c) by three S-sets, when the three S-sets are S_A, S_B, and S_C ($C_v = \{X_A, X_B, X_C\}$). Each of these S-sets are capable of

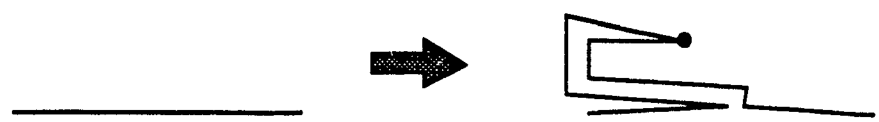

Figure 5.10: A spiral on an edge

covering the whole X-unit to which it corresponds in addition to the portion of CU_v which it covers.

These properties can be proved in a manner similar to that given in section 5.3.2, and then the theorem (NP-hardness) follows from the same proof as given for Theorem 5.3. □

5.4 $\mathbf{L_jCC}$: Even j

In this section we prove that $\mathbf{L_jCC}$ is hard for even j. We begin with the proof for L_2CC, which will then be generalized to the desired result.

5.4.1 Problem Transformation

We will prove that $\mathbf{L_2CC}$ is NP-hard by transformation from **3SAT**. Given an instance $I = (U, C)$ of **3SAT**, we construct an instance $\psi_2(I) = (P, m)$ of $\mathbf{L_2CC}$ as follows:

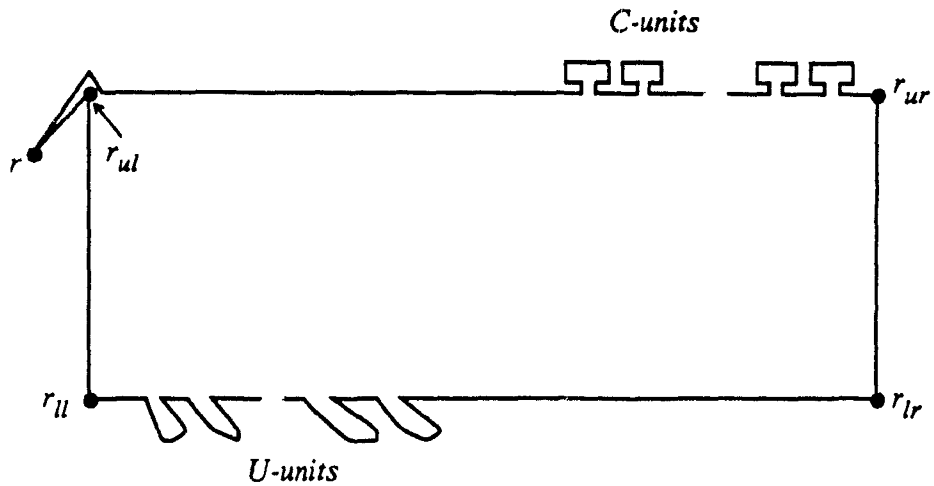

Figure 5.11: $\mathbf{L}_2$CC Construction overview

First, we let $m = 2q + 2n + 1$ (q and n are from the definition of **3SAT**). We assume that the literals in each clause appear in order of increasing index.

We construct P as follows: we start with a rectangle with an arm on the upper left corner, as in figure 5.11; we call this (rectangle and arm) the *central unit.* We let r_{ul}, r_{ll}, r_{ur}, and r_{lr} be the upper-left, lower-left, upper-right, and lower-right vertices of the rectangle, and r be the vertex at the end of the arm.

We will connect structures corresponding to the U_u's (called *U-units*) to the bottom of the central unit, and structures corresponding to the C_u's (called *C-units*) to the top of the central unit (see figure 5.11). Furthermore, all C-units are to the right of all U-units. We let UU_u denote the U-unit corresponding to U_u, and CU_v denote the C-unit corresponding to C_v. P will be the union of the U-units, C-units, and central unit.

Each C-unit is the union of four rectangles, as shown in figure 5.12 for CU_v. The vertices p_0 and p_1 will be on the upper edge of the central unit's rectangle; $b_{v,1}$ is the point p_0, and $b_{v,2}$ and $b_{v,3}$ are points one-third and two-thirds of the way from p_0 to p_1, respectively. We define CL_v as the union of the two rectangles shown shaded in figure 5.12, and the vertices $w_{v,1}$ and $w_{v,2}$ as shown The C-units are placed evenly along the right half of the upper edge of the central unit's rectangle in order of increasing

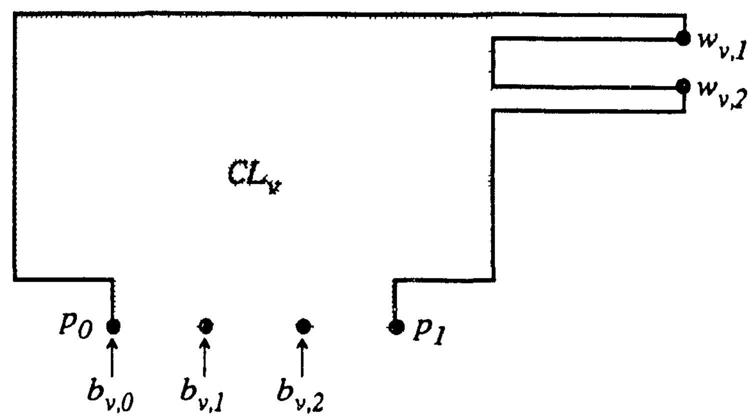

Figure 5.12: $\mathbf{L_2}$CC C-unit construction

index.

The U-units are a more complex structure which are placed evenly along the left half of the lower edge of the central unit's rectangle, also in order of increasing index. In the U-unit construction, we will be using *spikes*: these are very thin triangular notches, which we approximate by line segments sticking out from our polygon (as was done for the spiral arms in the C-unit construction for the proof of Theorem 5.4). We will show only the spikes in our description of the construction; keep in mind that these spikes will actually be replaced by thin triangles. The correct thinness for the spikes can easily be computed in polynomial time: for each spike, we find the radially closest (in both the clockwise and counterclockwise directions) sets that must be avoided by the spikes, and choose bounding edges for the triangles which replace the spikes so that these sets are not seen from the vertex at the end of the spike. This is a standard method (see [LL86] and [A84] for similar arguments).

The first stage of the U-unit construction for UU_u is illustrated in figure 5.13. The vertices p_0 and p_7 will be on the lower edge of the central unit's rectangle. The lines $\overline{p_0p_1}$, $\overline{p_2p_3}$, $\overline{p_4p_5}$, and $\overline{p_6p_7}$ are each colinear with r_{ul}. The lines $\overline{p_3p_4}$, $\overline{p_1p_2}$, and $\overline{p_5p_6}$ are horizontal, with $\overline{p_3p_4}$ high enough that p_4 can see r_{ur}, and $\overline{p_1p_2}$ and $\overline{p_5p_6}$ low enough that neither p_1 nor p_5 can see any vertex of any C-unit.

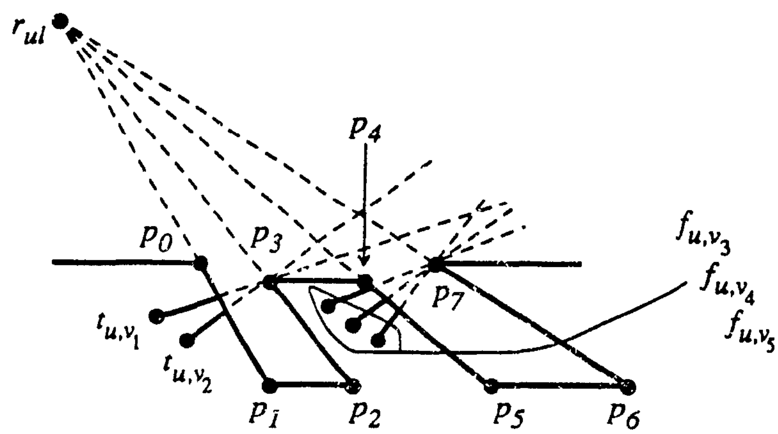

Figure 5.13: **L$_2$CC** U-unit construction I

For each clause C_v that U_u appears in, we will create a spike with vertex $t_{u,v}$. If U_u is the a^{th} literal of C_v, then we create the spike in edge $\overline{p_0p_1}$ colinear with $\overline{p_3b_{v,a}}$. Similarly, for each clause C_v that $\overline{U_u}$ appears in, we will create a spike with vertex $f_{u,v}$. This spike is in edge $\overline{p_4p_5}$, colinear with $\overline{p_7b_{v,a}}$, where $\overline{U_u}$ is the a^{th} literal of C_v. This part of the unit is similar to the construction for a variable pattern given in [LL86].

The second stage of construction for UU_u is illustrated in figure 5.14. Here we have added four spikes, and a small indentation on the edge $\overline{p_3p_4}$. First, a horizontal spike with vertex z_u is added in edge $\overline{p_0p_1}$, one-third of the vertical distance from $\overline{p_3p_4}$ to $\overline{p_0p_7}$ above $\overline{p_3p_4}$. Next, we let q_1 and q_4 be the points one-third and two-thirds of the way from p_3 to p_4, respectively. We create two new spikes, with vertices f_u and t_u, which intersect $\overline{p_0p_1}$ and $\overline{p_6p_7}$ (respectively) two-thirds of the vertical distance from $\overline{p_3p_4}$ to $\overline{p_0p_7}$. The spike with vertex f_u is made colinear with q_1, and the spike with vertex t_u is made colinear with q_4. Next, we place q_2 and q_3 such that $\overline{q_1q_2}$ and $\overline{q_3q_4}$ are colinear with r_{ul}, and $\overline{q_2q_3}$ is high enough that $VP_1(f_u)$ and $VP_1(t_u)$ do not intersect above $\overline{q_2q_3}$. Finally, we place a horizontal spike with vertex x_u at the vertex q_3.

To describe the final construction step, and in the subsequent proof, the following

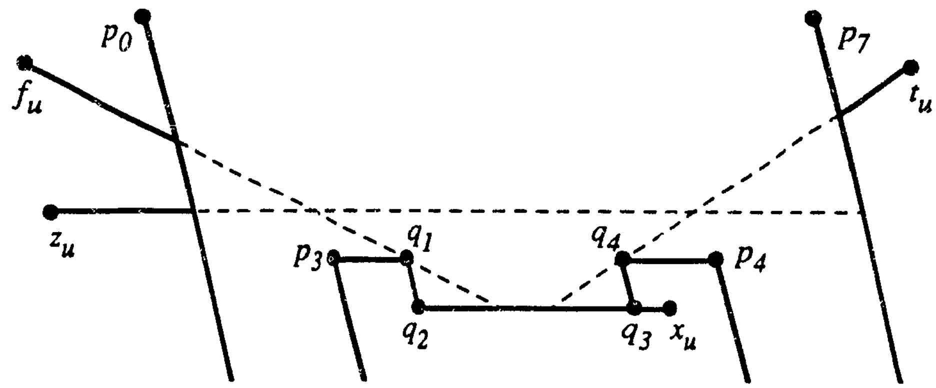

Figure 5.14: $\mathbf{L_2CC}$ U-unit construction II

alternative notation for the spikes $t_{u,v}$ and $f_{u,v}$ in clause unit CU_v will be useful: for every v and $a = 1$, 2, or 3, let

$$l_{v,a} = \begin{cases} t_{u,v} & \text{if the } a^{\text{th}} \text{ literal of } C_v \text{ is } U_u \\ f_{u,v} & \text{if the } a^{\text{th}} \text{ literal of } C_v \text{ is } \overline{U_u} \end{cases}$$

As a final step we must flatten out the C-units so that any two vertices of the form $l_{v,a}$ and $l_{v,a'}$ are not L_2-visible. This is done by computing the intersections of the three lines of the form $\overline{l_{v,a}b_{v,a}}$ for each v. These intersections will all be above the top of the central unit's rectangle, as the literals of a clause appear in sorted order, as do the U-units. We then place the horizontal edges of CU_v low enough that all of these line intersections are above the top edge and none of the lines intersect any other edge of the C-unit. We then have the situation illustrated in figure 5.15. This completes the construction.

5.4.2 Properties of the Construction

P1 $VP_2(r)$ is L_2-convex and covers all of P except the C-units and the spikes on the U-units.

P2 For all u, both $VP_1(x_u) \cup VP_1(t_u)$ and $VP_1(x_u) \cup VP_1(f_u)$ are L_2-convex.

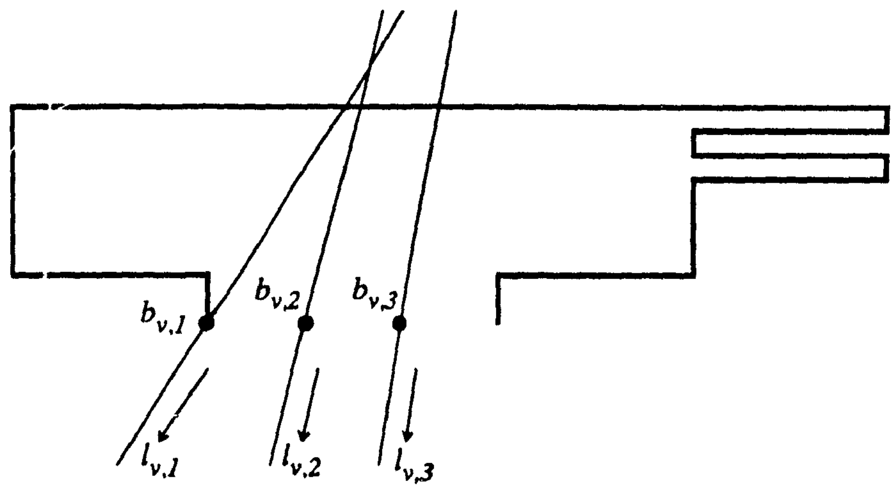

Figure 5.15: Flattening CU_v

P3 For all v, $a = 1, 2, 3$, and $b = 1$ or 2, $VP_1(w_{v,b}) \cup VP_1(l_{v,a}) \cup CL_v$ is L_2-convex.

P4 For all u, both $VP_1(z_u) \cup VP_1(t_u) \cup \bigcup_v VP_1(f_{u,v})$ and $VP_1(z_u) \cup VP_1(f_u) \cup \bigcup_v VP_1(t_{u,v})$ are L_2-convex.

P5 $H = \{r\} \cup \bigcup_u \{t_u, f_u\} \cup \bigcup_v \{w_{v,1}, w_{v,2}\}$ is a link-2 hidden set.

P6 For all u, $(H \cup \{x_u, z_u\}) \setminus \{t_u, f_u\}$ is a link-2 hidden set.

P7 For all u and v, $(H \cup \{x_u, t_{u,v}\}) \setminus \{f_u, w_{v,1}, w_{v,2}\}$ is a link-2 hidden set.
For all u and v, $(H \cup \{x_u, f_{u,v}\}) \setminus \{t_u, w_{v,1}, w_{v,2}\}$ is a link-2 hidden set.

P8 For all v, $\{l_{v,1}, l_{v,2}, l_{v,3}\}$ is a link-2 hidden set.

These properties are all easily verified from the construction.

5.4.3 L_2CC

Theorem 5.5 **L_2CC** *is NP-hard.*

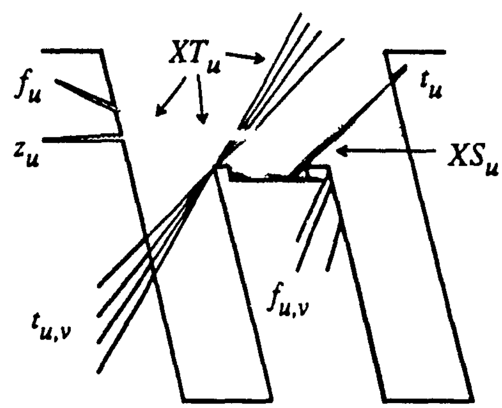

Figure 5.16: XS_u and XT_u when $\phi(U_u) = $ true

PROOF We show that the instance I of **3SAT** will have a yes answer iff the instance $\psi_2(I)$ of $\mathbf{L_2CC}$ has a yes answer (i.e., P can be covered by $m = 2q+2n+1$ L_2-convex sets).

If the instance I of **3SAT** has a yes answer, then there is some satisfying truth assignment $\phi : C \mapsto \{\text{true, false}\}$ for C. We will use the following cover for P.

First, we let $XQ = VP_2(r)$ be in the cover. XQ is L_2-convex, by property **P1** Also by property **P1**, we now need only cover the C-units and the spikes on the U-units. Then, for each U-unit UU_u, we let XS_u and XT_u be defined as follows:

$$XS_u = \begin{cases} VP_1(x_u) \cup VP_1(t_u) & \text{if } \phi(U_u) = \text{true} \\ VP_1(x_u) \cup VP_1(f_u) & \text{if } \phi(U_u) = \text{false} \end{cases}$$

$$XT_u = \begin{cases} VP_1(z_u) \cup VP_1(f_u) \cup \bigcup_v VP_1(t_{u,v}) & \text{if } \phi(U_u) = \text{true} \\ VP_1(z_u) \cup VP_1(t_u) \cup \bigcup_v VP_1(f_{u,v}) & \text{if } \phi(U_u) = \text{false} \end{cases}$$

These definitions are illustrated in figures 5.16 and 5.17 for $\phi(U_u) = $ true and $\phi(U_u) = $ false. The XS_u's and XT_u's are L_2-convex, by properties **P2** and **P4**. We place all XS_u's and XT_u's is the cover. We have thus used $2q+1$ sets, and have yet to cover only the C-units, and one set of spikes for each U-unit (either the spikes with vertices of the form $t_{u,v}$, or those with vertices of the form $f_{u,v}$).

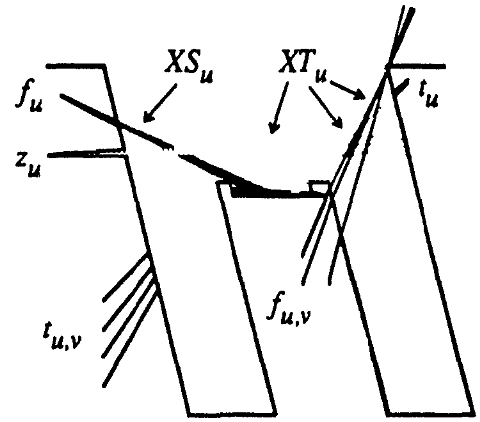

Figure 5.17: XS_u and XT_u when $\phi(U_u)$ = false

Next, for each C-unit CU_v, we let $a = 1, 2$, or 3 such that the a^{th} literal of C_v is true. We define $XR_{v,1}$ and $XR_{v,2}$ so that they will cover all of CU_v and two of the three spikes containing the vertices $l_{v,a}$; the uncovered spike will correspond to the a^{th} literal, which is known to be true:

$$XR_{v,1} = CL_v \cup VP_1(w_{v,1}) \cup \begin{cases} VP_1(l_{v,2}) & \text{if } a = 1 \\ VP_1(l_{v,1}) & \text{if } a = 2 \\ VP_1(l_{v,1}) & \text{if } a = 3 \end{cases}$$

$$XR_{v,2} = CL_v \cup VP_1(w_{v,1}) \cup \begin{cases} VP_1(l_{v,3}) & \text{if } a = 1 \\ VP_1(l_{v,3}) & \text{if } a = 2 \\ VP_1(l_{v,2}) & \text{if } a = 3 \end{cases}$$

The XR's are L_2-convex, by property **P3**. We place the XR's in the cover; we now have $2q + 2n + 1 (= m)$ sets.

We claim that the C-units and X-unit spikes are now covered. The XR's certainly cover the C-units, by the invariant part of their definitions.

Suppose there were some spike of some U-unit which were not covered. Assume that this spike contains $t_{u,v}$ (the case where the spike contains $f_{u,v}$ is similar). Since $t_{u,v}$ is not covered, it is in particular not covered by XT_u, implying that $\phi(U_u)$ = false.

The existance of $t_{u,v}$ implies that $t_{u,v} = l_{v,a}$ for $a = 1$, 2, or 3. Since neither $XR_{v,1}$ nor $XR_{v,2}$ covers $t_{u,v}$, we must have that the a^{th} literal of C_u satisfies C_u. However, this a^{th} literal must be U_u, by the definition of $l_{v,a}$; this implies that $\phi(U_u) =$ true. This is a contradiction.

Therefore the $2q + 2n + 1$ L_2-convex sets cover all of the spikes, and in fact cover the entire polygon. Hence the instance $\psi_2(I)$ of $\mathbf{L_2CC}$ has a yes answer.

If the instance of $\mathbf{L_2CC}$ has a yes answer, then there is a collection S of $2q + 2n + 1$ L_2-convex sets which cover P. Since, by property **P5**, $H = \{r\} \cup \bigcup_u \{t_u, f_u\} \cup \bigcup_v \{w_{v,1}, w_{v,2}\}$ is a link-2 hidden set (with size $2q+2n+1$), each member of S contains exactly one member of H. If $h \in H$, we let S_h be the member of S containing h.

We will use the following truth assignment ϕ foi our instance I of **3SAT**:

$$\phi(U_u) = \begin{cases} \text{true} & \text{if } S_{t_u} \text{ contains } x_u \\ \text{false} & \text{if } S_{f_u} \text{ contains } x_u \end{cases}$$

Note that property **P6** implies that x_u and z_u cannot be in $S_{h'}$ for any $h' \in H \backslash \{t_u, f_u\}$. This means that x_u and z_u must lie in $S_{t_u} \cup S_{f_u}$. As property **P6** also implies that x_u and z_u cannot be in the same S_h, exactly one of S_{t_u} and S_{f_u} contains x_u (and the other contains z_u).

We clain that ϕ is a satisfying truth assignment. We examine an arbitrary clause C_v: $S_{w_{v,1}}$ and $S_{w_{v,2}}$ can each cover at most one of $l_{v,1}$, $l_{v,2}$, and $l_{v,3}$, by property **P8**. Let a be such that $l_{v,a}$ is not covered by $S_{w_{v,1}}$ and $S_{w_{v,2}}$. We examine two cases, based on whether $l_{v,a} = t_{u,v}$ or $l_{v,a} = f_{u,v}$ for some u.

In the first case ($l_{v,a} = t_{u,v}$ for some u), $l_{v,a}$ must be covered by S_{f_u}, $S_{w_{v,1}}$, or $S_{w_{v,2}}$, by property **P7**. But by definition $l_{v,a}$ is not covered by $S_{w_{v,1}}$ or $S_{w_{v,2}}$. Therefore $l_{v,a}$ must be covered by S_{f_u}. Property **P7** then also implies that S_{f_u} can not contain x_u. Therefore, S_{t_u} must contain x_u, implying that $\phi(U_u) =$ true, by our definition of ϕ. Since U_u is the a^{th} literal of C_v, this means that C_v is satisfied.

If $l_{v,a} = f_{u,v}$, then a similar analysis holds: $l_{v,a}$ must be covered by S_{t_u}, which does not contain x_u. Therefore, S_{f_u} contains x_u, implying $\phi(U_u) =$ false. As the a^{th} literal of C_v is $\overline{U_u}$, C_v is satisfied.

Since in both cases C_v is satisfied, and C_v was chosen arbitrarily, all clauses are satisfied. Therefore ϕ is a satisfying truth assignment, and instance I of **3SAT** has a yes answer. □

We note that in this proof we have been using *multiply-connected* polygons as part of our cover (the XT_u's). However, changing the question of **L_jCC** from "can P be covered by m or fewer L_j-convex *sets*" to "can P be covered by m or fewer L_j-convex *polygons*" does not change our approach; for every set that we have placed in our cover that is not singly connected, we simply instead place the smallest simply-connected superset of that set in our cover. The following result (a corollary of [S70], theorem 4.5) shows that this will not affect our proof:

Theorem 5.6 *Let A be a compact L_j-convex subset of P. Then the smallest compact, simply-connected set in P containing A is also L_j-convex.*

5.4.4 Extension to Higher Even j

Theorem 5.7 *For any even integer $j \geq 2$,* **L_jCC** *is NP-hard.*

PROOF We modify the units and properties of the **L_2CC** construction.

The modifications to the units are as follows: the arm on the central unit is replaced by a spiral arm of $j-1$ arms, a spiral of $(j-2)/2$ arms is added at each $w_{v,1}$ and $w_{v,2}$, and a spiral of $(j-2)/2$ arms is added to the vertex of each spike on each U-unit. Sample modified units, for $j = 8$, are shown in figures 5.18, 5.19, and 5.20.

The properties of this construction are the same as the properties of the **L_2CC** construction, with "L_2-convex" replaced by "L_j-convex" "$VP_1(x)$" replaced by "$VP_{j/2}(x)$," "$VP_2(r)$" replaced by "$VP_j(r)$," and "link-2 hidden set" replaced by "link-j hidden set."

The proof of this theorem is then identical to that of Theorem 5.5. □

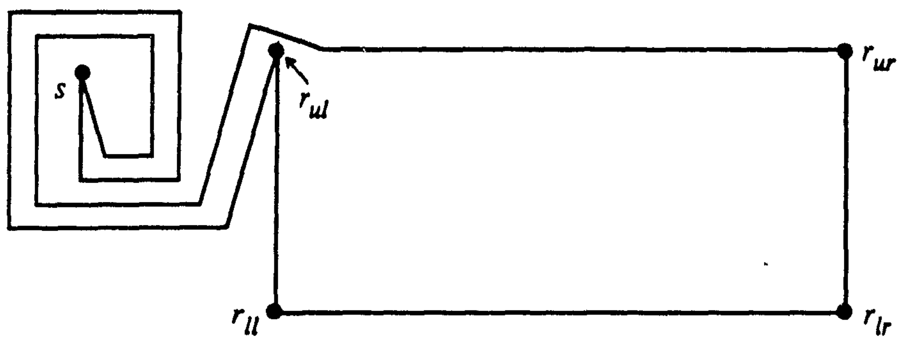

Figure 5.18: Central unit for $j = 8$

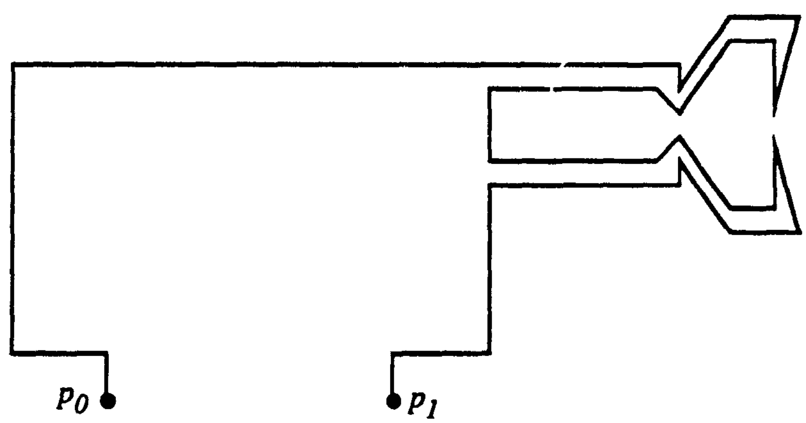

Figure 5.19: C-unit for $j = 8$

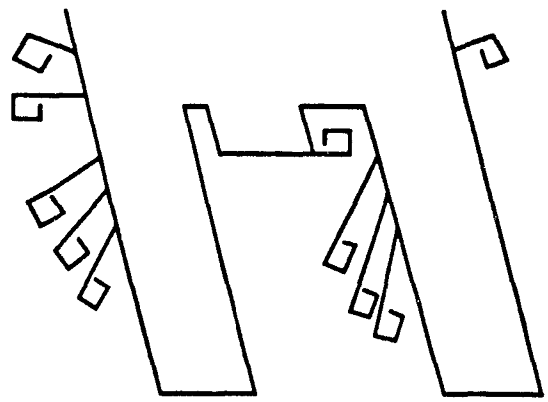

Figure 5.20: U-unit for $j = 8$

5.5 $\mathbf{L_jSC}$ and $\mathbf{L_{j,k}G}$

In this section we show that $\mathbf{L_jSC}$ and $\mathbf{L_{j,k}G}$ are NP-hard. We prove this for $\mathbf{L_jSC}$ by modification of our construction and proof for even-j $\mathbf{L_jCC}$. The result for $\mathbf{L_{j,k}G}$ is a combination and modification of the results for $\mathbf{L_jCC}$ and $\mathbf{L_jSC}$.

Theorem 5.8 *For any integer $j \geq 1$,* $\mathbf{L_jSC}$ *is NP-hard.*

PROOF We modify the U-units of the construction of Theorem 5.7. First, we let $z_{t,u}$ be a point in the intersection of $VP_1(z_u)$ and $VP_1(t_u)$, and $z_{f,u}$ be a point in the intersection of $VP_1(z_u)$ and $VP_1(f_u)$. We change the orientation and location of the spikes with vertices z_u, f_u, and t_u so that $z_{t,u}$ and $z_{f,u}$ both see r_{ur}. We must also change the height of the edges $\overline{p_1p_2}$, $\overline{p_3p_4}$, and $\overline{p_5p_6}$, so that the lines $\overline{p_1z_{f,u}}$ and $\overline{p_5z_{t,u}}$ intersect the top edge of the central unit to the left of all of the C-units.

We also change the spikes with vertices $t_{u,v}$ and $f_{u,v}$. We construct them colinear with $\overline{z_{f,u}b_{v,a}}$ and $\overline{z_{t,u}b_{v,a}}$, respectively (rather than colinear with $\overline{p_3b_{v,a}}$ and $\overline{p_7z_{f,u}b_{v,a}}$). A sample U-unit for $j = 1$ is shown in figure 5.21.

The interesting properties of this construction are the same as the properties of the construction for Theorem 5.7, except that we replace "L_2-convex" with "star-shaped." The changes to the U-unit were to make the sets considered in property **P4**

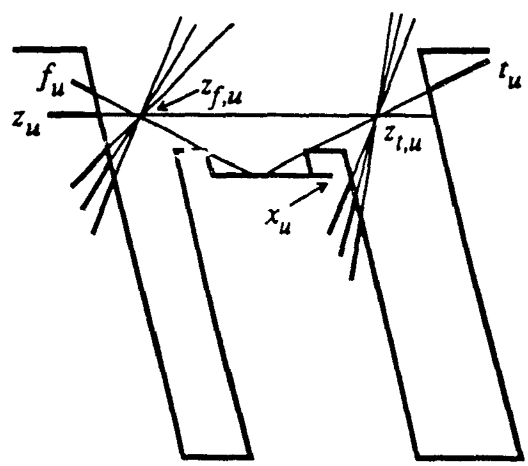

Figure 5.21: $\mathbf{L_jSC}$: U-unit

star-shaped; the other sets (properties **P1** - **P3**) were already star-shaped.

This theorem then is obtained by following the proof of Theorem 5.7, with "L_2-convex" replaced by "star-shaped." □

Theorem 5.9 *For any nonnegative integers j and k, at least one of which is positive,* $\mathbf{L_{j,k}G}$ *is NP-hard.*

PROOF If $k = 0$, then the problem is the $\mathbf{L_jSC}$ problem. If $j = 0$, then the problem is the $\mathbf{L_kCC}$ problem.

Otherwise, we use the construction for $\mathbf{L_{k+2j}CC}$.

If $k + 2j$ is odd, then k is odd. We note that the covers given in the proof of Theorem 5.4 consist entirely of subsets of sets of the form $VP_{(k-1)/2+j}(S)$, where S is L_1-convex. Thus those sets are also of the form $VP_j(S')$, where S' is L_k-convex; we let $S' = VP_{(k-1)/2}(S)$.

Similarly, if $k + 2j$ is even, k is even. The covers given in the proof of Theorem 5.7 consist entirely of subset of sets of the form $VP_{(k-2)/2+j}(S)$, where S is L_2-convex. Thus those sets are also of the form $VP_j(S')$, where S' is L_k-convex; we let $S' = VP_{(k-2)/2}(S)$.

Thus, in all cases, we can show that the polygons from the **L_jCC** constructions can be covered by the appropriate number of $VP_j(S)$'s when the instance I of **X3C** or **3SAT** has a yes answer.

The proof of the other implication (that the instance of **L_jCC** has a yes answer implies that the original problem instance has a yes answer) is unmodified from the **L_jCC** proofs. This is due to the observation that if the instance of **$L_{j,k}$G** has a yes answer, then the instance of **L_{k+2j}CC** must have a yes answer (which by the previous proofs give the desired result). □

5.6 Hidden Set Results

We can also use the methods of this chapter to show that the following problems, which were proved NP-hard for $j = 1$ in [S87], are NP-hard for any odd j.

L_j-HIDDEN SET (L_jHS)

INSTANCE: A polygon P, and an integer m.

QUESTION: Does P have a link-j hidden set with m or more members?

L_j-HIDDEN VERTEX SET (L_jHVS)

INSTANCE: A polygon P, and an integer m.

QUESTION: Does P have a link-j hidden vertex set with m or more members?

L_j-HIDDEN GUARD SET (L_jHGS)

INSTANCE: A polygon P, and an integer m.

QUESTION: Does P have a link-j hidden guard set with m or fewer members?

L_j**-HIDDEN VERTEX GUARD ADMISSABILITY (LjHVGA)**

INSTANCE: A polygon P.

QUESTION: Does P admit a link-j hidden vertex guard set?

L_j**-HIDDEN VERTEX GUARD SET (LjHVGS)**

INSTANCE: A polygon P, and an integer m.

QUESTION: Does P have a link-j hidden vertex guard set with m or fewer members?

We prove that these problems are NP-hard by a "link-j modification" of the $j = 1$ proofs. This modification is similar to the ones presented in this chapter for even and odd j **LjCC**. We omit the technical details.

Theorem 5.10 *For any odd integer* $j \geq 1$, **LjHS** *is NP-hard.*

Theorem 5.11 *For any odd integer* $j \geq 1$, **LjHVS** *is NP-complete.*

Theorem 5.12 *For any odd integer* $j \geq 1$, **LjHGS** *is NP-hard.*

Theorem 5.13 *For any odd integer* $j \geq 1$, **LjHVGA** *is NP-complete.*

Theorem 5.14 *For any odd integer* $j \geq 1$, **LjHVGS** *is NP-complete.*

These NP-hardness proofs all use similar constructions, so it is probably the case that one could prove the even-j variants with only one more construction We are thus led to conjecture:

Conjecture 5.15 *For all integers* $j \geq 1$, *the problems* **LjHS**, **LjHGS**, **LjHVS**, **LjHVGA**, *and* **LjHVGS** *are NP-hard.*

5.7 Graph and Polygon Complexity

It is one of the major contentions of this thesis that geometric visibility problems in polygons can be viewed as graph-theoretic problems on either the vertex-visibility graph or point-visibility graph of the polygon. In connection with this, it is interesting to point out the parallels in known visibility and known graph-theory complexity results: In *every* known instance, the complexity (either polynomial computability or NP-hardness) of the pure graph theoretic problem is the same as the associated polygon visibility problem.

For example, the independent set problem in a graph is NP-hard [K72], as is the hidden vertex set problem (independent set in $VVG(P)$), as is the hidden set problem (indpendent set in $PVG(P)$). A similar statement can be made about independent dominating sets (hidden vertex guard sets, hidden guard sets), and dominating sets (vertex guard sets, guard sets). Also, the k-colorability (chromatic number) problem for a graph is NP-hard [K72], as is the convex cover problem for polygons (k-colorability of the complement of $PVG(P)$).

Examples of polynomially-computable properties include the distance between two vertices (link-distance between two points), the center of a graph (link-center of a polygon), and the diameter and radius of a graph (link-diameter and link-radius of a polygon).

One must be careful with this relationship, though. For instance, a *maximal* clique in $PVG(P)$ corresponds to a *maximal* convex set in P, but a clique in $PVG(P)$ does not necessarily correspond to a convex set in P. Without this distinction one may become perplexed that there is a polynomial algorithm to find a minimum convex partition of a polygon, whereas it is NP-hard to find a minimum partition of a graph into cliques [K72]. Also, there are many NP-hard or NP-complete graph-theory problems that have no meaningful PVG counterpart; examples of these problems are finding a Hamiltonian circuit, finding a minimum maximal matching, and partitioning a graph into forests. These problems lose their substance on the infinite-vertex, infinite-degree graphs that we consider.

We have shown here many results using link-j visibility. This corresponds to solving problems in the jth power of a graph (see [H69] for definitions of powers of graphs). For example, solving a visiblity problem using L_2-visibility corresponds to solving a graph theoretic problem on the square of a PVG. This leads us to conjecture that the problems that we have proved NP-hard for polygons are also NP-hard on graphs, where we restrict our attention to graphs which are the jth power of some graph. The only result of this type of which I am aware is that finding a Hamiltonian circuit in the square of a graph is NP-complete [C76]; unfortunately, Hamiltonian circuit is a problem which is meaningless on PVGs.

Chapter 6

Conclusion

6.1 Method and Results

This thesis presents an extension and modification of the combinatorial method of Chvátal and O'Rourke, used for finding bounds on the value of many visibility properties of polygons. The bounds that we have obtained, which generalize and unify the previously-known bounds, are shown in figure 6.1.

The method presented can be used to obtain bounds for restricted polygon classes or other guard classes as well; it has been applied with success to the problems of finding bounds in orthogonal polygons (for even link-diameter covering/guarding) and finding bounds on the number of *edge guards* required for simple polygons.

This thesis also introduces the notion that visibility problems should be viewed as graph-theory problems on point-visibility graphs, and begins exploration of the comparative problem complexity of ordinary graphs and point-visibility graphs. We showed that the **LINK$_{j,k}$-GUARDING** problem, and all of its subproblems (including **L_j-CONVEX COVER** and **L_j-STAR COVER**) are NP-hard, using two constructions, and a method of extending the constructions to higher link-visibility or link-diameter. This extension method can be applied to NP-hardness proofs for any visibility property, and this was done for the hidden set problems proved NP-hard in [S87].

<table>
<tr><th rowspan="2">Object</th><th rowspan="2">Problem</th><th rowspan="2">j</th><th rowspan="2">k</th><th colspan="2">Bounds</th><th rowspan="2">Theorems, Corollaries</th></tr>
<tr><th>Lower</th><th>Upper</th></tr>
<tr><td rowspan="8">Polygon</td><td>Hidden Vertex Set</td><td></td><td>–</td><td colspan="2">$\lfloor \frac{n}{j+1} \rfloor$</td><td>3.1, 4.2</td></tr>
<tr><td rowspan="2">Hidden Set</td><td>1</td><td>–</td><td colspan="2">$n-2$</td><td>3.2</td></tr>
<tr><td>> 1</td><td>–</td><td colspan="2">$\lfloor \frac{n}{j+1} \rfloor$</td><td>3.1, 4.8e</td></tr>
<tr><td>L_k-Convex Cover</td><td>–</td><td>> 1</td><td colspan="2">$\lfloor \frac{n}{k+1} \rfloor$</td><td>3.1, 4.8d</td></tr>
<tr><td>L_k-Convex Partition</td><td>–</td><td>> 1</td><td colspan="2">$\lfloor \frac{n}{k+1} \rfloor$</td><td>3.1, 4.8c</td></tr>
<tr><td>Guarding $\mathcal{T}_k \subseteq \mathcal{C} \subseteq \mathcal{L}_k$</td><td></td><td></td><td colspan="2">$\lfloor \frac{n}{k+2j+1} \rfloor$</td><td>3.1a, 4.8b</td></tr>
<tr><td rowspan="2">Hidden Guard Set</td><td>1</td><td>–</td><td>$\lfloor \frac{n}{2} \rfloor - 1$</td><td>$n-2$</td><td>3.4, 3.2</td></tr>
<tr><td>> 1</td><td>–</td><td>$\lfloor \frac{n}{2j} \rfloor - 1$</td><td>$\lfloor \frac{n}{j+1} \rfloor$</td><td>3.4, 4.8e</td></tr>
<tr><td rowspan="2">Triangulation Graph</td><td>Independent Set</td><td></td><td>–</td><td colspan="2">$\lfloor \frac{n}{j+1} \rfloor$</td><td>3.1b, 4.8e</td></tr>
<tr><td>Dominating Set</td><td>–</td><td></td><td colspan="2">$\lfloor \frac{n}{k+3} \rfloor$</td><td>3.1c, 4.8</td></tr>
<tr><td rowspan="8">Polygon Exterior</td><td>Hidden Vertex Set</td><td></td><td>–</td><td colspan="2">$\lfloor \frac{n}{j+1} \rfloor$</td><td>3.6, 4.2a</td></tr>
<tr><td rowspan="2">Hidden Set</td><td>1</td><td>–</td><td colspan="2">n</td><td>3.7</td></tr>
<tr><td>> 1</td><td>–</td><td>$\lfloor \frac{n+1}{j+1} \rfloor$</td><td>$\lfloor \frac{n+3}{j+1} \rfloor$</td><td>3.8, 4.9c</td></tr>
<tr><td rowspan="2">L_k-Convex Cover</td><td>–</td><td>1</td><td colspan="2">n</td><td>3.7</td></tr>
<tr><td>–</td><td>> 1</td><td>$\lfloor \frac{n+1}{k+1} \rfloor$</td><td>$\lfloor \frac{n+3}{k+1} \rfloor$</td><td>3.8, 4.9b</td></tr>
<tr><td rowspan="2">L_k-Convex Partition</td><td>–</td><td>1</td><td colspan="2">n</td><td>3.7</td></tr>
<tr><td>–</td><td>> 1</td><td>$\lfloor \frac{n+1}{k+1} \rfloor$</td><td>$\lfloor \frac{n+3}{k+1} \rfloor$</td><td>3.8, 4.9b</td></tr>
<tr><td>Guarding $\mathcal{T}^e_k \subseteq \mathcal{C} \subseteq \mathcal{L}^e_k$</td><td colspan="2">$k + 2j > 1$</td><td>$\lfloor \frac{n+1}{k+2j+1} \rfloor$</td><td>$\lfloor \frac{n+3}{k+2j+1} \rfloor$</td><td>3.8a, 4.9a</td></tr>
</table>

Figure 6.1: Table of results

6.2 Open Problems

We have raised three major questions in the thesis to which we do not yet have answers:

- What is the exact tight bound for exterior visibility properties of polygons? The current bounds are *almost-tight*, but it is unsatisfying to not have exact bounds. Two methods have been used to get tight exterior bounds for point guarding (namely, that of Aggarwal and O'Rourke, and that of Aggarwal, O'Rourke, and Shermer [O87]), but neither of these methods seems easy to generalize.

- Are the visibility-property decision problems examined in the text in NP? This seems a hard question to answer, even for the simplest problem, Convex Cover [O82a].

- Can a construction be found for even j for the hidden set decision problems? This seems to be the easiest of these three questions.

There are also many questions which we did not explicitly raise, but which are nevertheless relevant. A sampling of these are:

- Linear algorithms exist to determine if a polygon has a hidden set of size two, and to determine if a polygon is the union of two convex sets [S88c]. Does there exist a good algorithm to determine if a polygon is the union of two star-shaped sets?

- The combinatorial method of this thesis can be applied to orthogonal polygons, when covering with sets of *even* link-diameter. What bounds can be found for covering orthogonal polygons with sets of *odd* link-diameter?

- Our combinatorial method is a generalization of Chvátal's art gallery proof. Preliminary research indicates that Fisk's proof can also be generalized; in particular, we can find a *k-thicket* in any triangulation graph. A k-thicket is a set of n unique D_k-trees such that:

  (1) Each tree is colored one of $k+3$ colors.

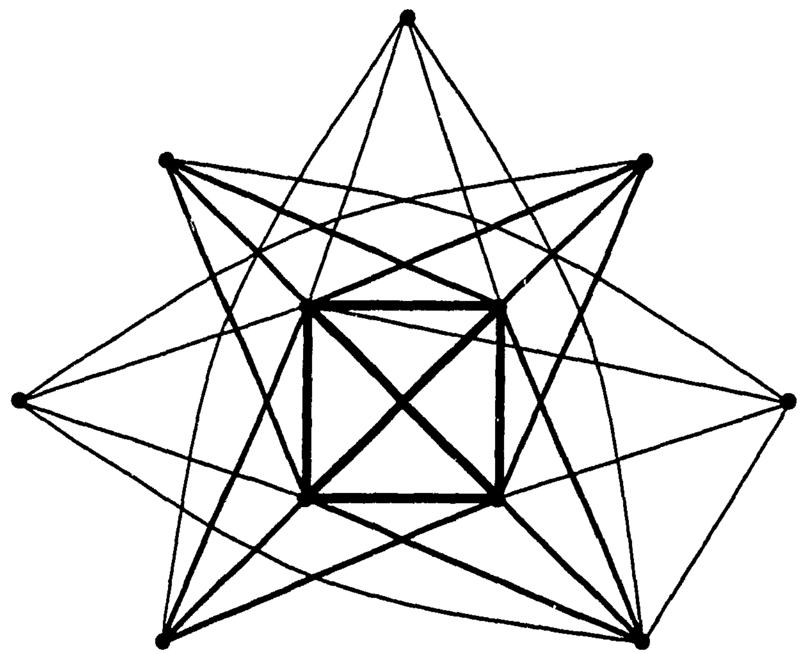

Figure 6.2: A forbidden induced subgraph

(2) Each triangle in the triangulation graph has at least one tree of each color incident on it.

A 0-thicket is exactly a 3-coloring, and k-thickets provide us with high link-diameter/link-visibility guard sets in the same manner that 3-coloring does for point or vertex guard sets. Can k-thickets be used to get tight exterior bounds? Are there any applications of k-thickets in graph theory?

- Can *PVG*s be characterized? Some progress has been made in this direction; there are examples of graphs which cannot be induced graphs of any *PVG* (see figure 6.2 for an example). Can all such forbidden induced subgraphs be characterized?

- Considering *PVGs* as graphs raises many questions. For example, when are two polygons *isomorphic* with respect to visibility? Is this problem decidable? It is known that all convex polygons are isomorphic, as are all polygons with one reflex vertex. Polygons with two reflex vertices are not all isomorphic; but nothing further is known. An interesting question is: how many different nonisomorphic polygons are there with two reflex vertices? It is suspected that there are infinitely many.

- Consider guarding and covering polygons with holes using the guard classes and visibility discussed here. For point guards, the leading conjecture is that $\lfloor (n+h)/(k+3) \rfloor$ guards are necessary and sufficient. However, no examples have been found for higher k which require more than $\lfloor n/(k+3) \rfloor$ guards. Is this the tight bound? This problem is very closely related to the exterior guarding problem (a polygon exterior can be treated as a hole without a polygon around it), and the remarks about the difference between $k = 0$ and $k > 1$ for that problem apply here as well.

- Are there any good approximation algorithms for the problems that we have shown to be NP-hard?

- Naïve implementation of the constructive proof for link-guards yields an $O(n^2)$ algorithm for guard placement. Can this time be improved?

6.3 Conclusion

Visibility problems are central to several applied subfields of computer science, including computer graphics, pattern recognition, robotics, computer-aided design, computer-aided architecture, and VLSI. The generalization of visibility that we have studied finds application mostly in robotics, but the generalized guard classes and covering objects are likely to be useful in many fields.

We have given tight combinatorial bounds on the size of hidden sets, guard sets, and covering sets, and have shown the close relationship between these properties. Although these bounds are more interesting to the geometer or graph theorist than the computer scientist, the proof method can be mimicked to get an $O(n^2)$ algorithm for guard placement (for any of the guard classes we use and any link-visibility). We have also shown that the optimization and decision problems relating to computing these properties are NP-hard.

Bibliography

[A84] A. Aggarwal, *The Art Gallery Theorem: its Variations, Applications, and Algorithmic Aspects*, PhD Thesis, The Johns Hopkins University, Baltimore, 1984.

[AE83] D. Avis and H. ElGindy, "A Combinatorial Approach to Polygon Similarity," *IEEE Transactions on Information Theory*, 1983, 148-150.

[AT81a] D. Avis and G. Toussaint, "An Efficient Algorithm for Decomposing a Polygon into Star-Shaped Pieces," *Pattern Recognition* **13**, 1981, 295-298.

[AT81b] D. Avis and G. Toussaint, "An Optimal Algorithm for Determining the Visibility of a Polygon from an Edge," *IEEE Trans. Computing* **C-30**, 1981, 910-914.

[BM76] J. Bondy and U. Murty, *Graph Theory with Applications*, North-Holland, New York, 1976.

[B73] M. Breen, "The Combinatorial Structure of (m,n)-Convex Sets," *Israel J. of Mathematics* **15**, 1973, 367-374.

[B76] M. Breen, "A Decomposition Theorem for m-Convex Sets," *Israel J. of Mathematics* **24**, 1976, 211-216.

[B77] M. Breen, "L_2 Sets Which are Almost Starshaped," *Geometriae Dedicata* **6**, 1977, 485-494.

[BK76] M. Breen and D. Kay, "General Decomposition Theorems for m-Convex Sets in the Plane," *Israel J. of Mathematics* **24**, 1976, 217-233.

[BB64] A. Bruckner and J. Bruckner, "Generalized Convex Kernels," *Israel J. of Mathematics* **2**, 1964, 27-32.

[C80] B. Chazelle, *Computational Geometry and Convexity*, PhD Thesis, Yale University, New Haven, 1980.

[C82] B. Chazelle, "A Theorem on Polygon Cutting with Applications," *Proceedings of the 23th Annual ACM Conference on the Foundations of Computer Science*, Chicago, 1982, 339-349.

[CD85] B. Chazelle and D. Dobkin, "Optimal Convex Decompositions," in *Computational Geometry*, G. Toussaint, ed., North-Holland, 1985.

[C75] V. Chvátal, "A Combinatorial Theorem in Plane Geometry," *J. Combinatorial Theory Series B* **18**, 1975, 39-41.

[C76] V. Chvátal, "Finding a Hamiltonian Circuit in the Square of a Graph is NP-complete," manuscript, 1976.

[CR88] J. Culberson and R. Reckhow, "Covering Polygons is NP-Hard," *Proceedings of the 29th Annual ACM Conference on the Foundations of Computer Science*, White Plains, 1988, 601-611.

[EOW84] H. Edelsbrunner, J. O'Rourke, and E. Welzl, "Stationing Guards in Rectilinear Art Galleries," *Computer Vision, Graphics, and Image Processing* **27**, 1984, 167-176.

[E85] H. ElGindy, *Hierarchical Decomposition of Polygons with Applications*, PhD Thesis, McGill University, Montréal, 1984.

[E89] H. Everett, "Visibility Graphs of Spiral Polygons," presented at *the First Canadian Conference on Computational Geometry*, Montréal, 1989.

[F78] S. Fisk, "A Short Proof of Chvátal's Watchman Theorem," *J. Combinatorial Theory Series B* **24**, 1978, 374.

[FK84] D. Franzblau and D. Kleitman, "An Algorithm for Covering Polygons with Rectangles," *Information and Control* **63**, 1984, 164-189.

[GJ79] R. Garey, and D. Johnson, *Computers and Intractability*, W. H. Freeman and Company, New York, 1979.

[G86] S. Ghosh, "On Recognizing and Characterizing Visibility Graphs of Simple Polygons," Johns Hopkins University Department of Computer Science, Report JHU/EECS-86/14, 1986.

[GK71] M. Guay and D. Kay, "On Sets Having Finitely Many Points of Local Nonconvexity and Property P_m," *Israel J. of Mathematics* **10**, 1971, 196-209.

[GHL86] L. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. Tarjan, "Linear Time Algorithms for Visibility and Shortest Path Problems Inside Simple Polygons," *Proceedings of the 2nd ACM Symposium on Computational Geometry*, Yorktown Heights, 1986, 1-13.

[H69] F. Harary, *Graph Theory*, Addison-Wesley, Reading, 1969.

[HK68] W. Hare, Jr., and J. Kenelly, "Intersection of Maximal Starshaped Sets," *Proceedings of the American Mathematical Society* **19**, 1968, 1299-1302.

[H87] J. Hershberger, "Finding the Visibility Graph of a Simple Polygon in Time Proportional to its Size," *Proceedings of the 3rd ACM Symposium on Computational Geometry*, Waterloo, 1987, 11-20.

[HV49] A. Horn and F. Valentine, "Some Properties of L-Sets in the Plane," *Duke Mathematics J.* **16**, 1949, 131-140.

[J77] K. Juul, "A Three-Point Convexity Property and the Union of Two Convex Sets," *Geometriae Dedicata* **6**, 1977, 181-192.

[KKK83] J. Kahn, M. Klawe, and D. Kleitman, "Traditional Galleries Require Fewer Watchmen," *SIAM J. Algebraic and Discrete Methods* **4**, 1983, 194-206.

[KM88] S. Kapoor and S. Maheshwari, "Efficient Algorithms for Euclidean Shortest Path and Visibility Problems with Polygonal Obstacles," *Proceedings of the 4th ACM Symposium on Computational Geometry*, Urbana-Champaign, 1988, 164-171.

[K72] R. Karp, "Reducibility Among Combinatorial Problems," in *Complexity of Computer Computations*, R. Miller and J. Thatcher, eds., Plenum Press, New York, 1972, 85-103.

[KG70] D. Kay and M. Guay, "Convexity and a Certain Property P_m," *Israel Journal of Mathematics* **8**, 1970, 39-52.

[K89] Y. Ke, "An Efficient Algorithm for Link Distance Problems," *Proceedings of the 5th ACM Symposium on Computational Geometry*, Saarbrücken, 1989, 69-78.

[KM66] C. Koch and J. Marr, "A Characterization of Unions of Two Star-Shaped Sets," *Proceedings of the American Mathematical Society* **17**, 1966, 1341-1343.

[LL86] D. Lee and A. Lin, "Computational Complexity of Art Gallery Problems," *IEEE Trans. Information Theory* **IT-32**, 1986, 276-282.

[LPS87] W. Lenhardt, R. Pollack, J.-R. Sack, R. Seidel, M. Sharir, S. Suri, G. Toussaint, S. Whitesides, and C. Yap, "Computing the Link Center of a Simple Polygon," *Proceedings of the 3rd ACM Symposium on Computational Geometry*, Waterloo, 1987, 1-10.

[L85] A. Lubiw, "Decomposing Polygonal Regions into Convex Quadrilaterals," *Proceedings of the ACM Symposium on Computational Geometry*, Baltimore, 1985, 97-106.

[MW84] H. Mannila and D. Wood, "A Simple Proof of the Rectilinear Art Gallery Theorem," University of Helsinki Computer Science Technical Report C-1984-16, 1984.

[M66] R. McKinney, "On Unions of Two Convex Sets," *Canadian Journal of Mathematics* **18**, 1966, 883-886.

[M75] G. Meisters, "Polygons Have Ears," *American Mathematics Monthly* **82**, 1975, 648-651.

[M87] D. Mount, personal communication, 1987.

[N-W73] C. Nash-Williams, "Unexplored and Semi-Explored Territories in Graph Theory," in *New Directions in Graph Theory*, F. Harary, Ed., Academic Press, New York, 1973.

[N86] S. Ntafos, "On Gallery Watchmen in Grids," *Information Processing Letters*, **23**, 99-102.

[O82a] J. O'Rourke, "Minimum Convex Covers for Polygons: Some Counterexamples," Johns Hopkins University Department of Computer Science, Report JHU/EECS-82/1, 1982.

[O82b] J. O'Rourke, "The Complexity of Computing Minimum Convex Covers for Polygons," *Proceedings of the 20th Allerton Conference*, Monticello, 1982, 75-84.

[O82c] J. O'Rourke, "The Decidability of Covering by Convex Polygons," Johns Hopkins University Department of Computer Science, Report JHU/EECS-82/4, 1982.

[O83a] J. O'Rourke, "Galleries Need Fewer Mobile Guards: A Variation on Chvátal's theorem," *Geometriae Dedicata* **14**, 1983, 273-283.

[O83b] J. O'Rourke, "An Alternate Proof of the Rectilinear Art Gallery Theorem," *J. of Geometry*, **21**, 1983, 118-130.

[O87] J. O'Rourke, *Art Gallery Theorems and Algorithms*, Oxford University Press, Oxford, 1987.

[OS83] J. O'Rourke and K. Supowit, "Some NP-hard Decomposition Problems," *IEEE Transactions on Information Theory* **IT-29**, 1983, 181-190.

[OW88] M. Overmars and E. Welzl, "New Methods for Computing Visibility Graphs," *Proceedings of the 4th ACM Symposium on Computational Geometry*, Urbana-Champaign, 1988, 164-171.

[ST82] J.-R. Sack and G. Toussaint, "A Linear-Time Algorithm for Decomposing Rectilinear Star-Shaped Polygons into Convex Quadrilaterals," *Proceedings of the 20th Annual Conference on Communication, Control, and Computing*, 1982, 64-74.

[ST88] J.-R. Sack and G. Toussaint, "Guard Placement in Rectilinear Polygons," in *Computational Morphology*, G. Toussaint, ed., North-Holland, 1988, 153-175.

[SH79] L. Shapiro and R. Haralick, "Decomposition of Two-dimensional Shapes by Graph-theoretic Clustering," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1979, 10-20.

[S84] T. Shermer, "Triangulation Graphs that Require Extra Guards," New York Institute of Technology Computer Graphics Laboratory TR 3D-13, 1984.

[S85] T. Shermer, "Polygon Guarding II: Efficient Reduction of Triangulation Fragments," New York Institute of Technology Computer Graphics Laboratory TR 3D-16, 1985.

[S87] T. Shermer, "Hiding People in Polygons," McGill University School of Computer Science TR SOCS.87-18, 1987 (to appear in *Computing*).

[S88a] T. Shermer, "Link Guarding Simple Polygons," McGill University School of Computer Science TR SOCS.88-12, 1988.

[S88b] T. Shermer, "Convex Cover is NP-hard," Manuscript, McGill University, 1988.

[S88c] T. Shermer, "On Recognizing Unions of Two Convex Sets," Manuscript, McGill University, 1988.

[S70] A. Sparks, "Intersections of Maximal L_n Sets," *Proceedings of the American Mathematical Society* **24**, 1970, 245-250.

[SM63] W. Stamey and J. Marr, "Unions of Two Convex Sets," *Canadian Journal of Mathematics* **15**, 1963, 152-156.

[S86a] S. Suri, "A Linear Time Algorithm for Minimum Link Paths Inside a Simple Polygon," *Computer Graphics, Vision, and Image Processing* **35**, 1986, 99-110.

[S86b] S. Suri, "Computing all Geodesic Furthest Neighbors of a Simple Polygon," Manuscript, The Johns Hopkins University, 1986.

[TV88] R. Tarjan and C. Van Wyk, "An $O(n \log\log n)$ Algorithm for Triangulating Simple Polygons," *SIAM Journal on Computing*, **17**, 1988, 143-178.

[T86] G. Toussaint, "Shortest Path Solves Edge-to-edge Visiblity in a Polygon," *Pattern Recognition Letters* **4**, 1986, 165-170.

[T87] G. Toussaint, "A Linear-time Algorithm for Solving the Strong Hidden-line Problem in a Simple Polygon," *Pattern Recognition Letters* **5**, 1987.

[T88] G. Toussaint, "Computing Visibility Properties of Polygons," in *Pattern Recognition and Artificial Intelligence*, E. Gelsema and L. Kanal, Eds., North-Holland, 1988, 130-122.

[V53] F. Valentine, "Minimal Sets of Visibility," *Proceedings of the American Mathematical Society* **4**, 1953, 917-921.

[V57] F. Valentine, "A Three Point Convexity Property," *Pacific Journal of Mathematics* **7**, 1957, 1227-1235.

[V64] F. Valentine, *Convex Sets*, McGraw-Hill, New York, 1964.

[V65a] F. Valentine, "Local Convexity and Starshaped Sets," *Israel Journal of Mathematics* **3**, 1965, 39-42.

[V65b] F. Valentine, "Local Convexity and L_n Sets," *Proceedings of the American Mathematical Society* **16**, 1965, 1305-1310.

[W85] E. Welzl, "Constructing the Visibility Graph for n Line Segments in $O(n^2)$ Time," *Information Processing Letters* **20**, 1985, 167-171.

[YB61] I. Yaglom and V. Boltyanskiĭ, *Convex Figures* (English translation by P. Kelly and L. Walton), Holt, Rinehart, and Winston, New York, 1961 (Original Russian published in 1951).

Index