Point Visibility Graphs and Restricted-Orientation Polygon Covering

by

David Dylan Bremner

B.Sc. (Hons.), University of Calgary, 1990

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE
in the School
of
Computing Science

SIMON FRASER UNIVERSITY
April 1993

APPROVAL

Name: David Dylan Bremner

Degree: Master of Science

Title of thesis: Point Visibility Graphs and Restricted-Orientation Polygon Covering

Examining Committee: Dr. Lou Hafer
Chair

Dr. Thomas Shermer
Senior Supervisor
Assistant Professor of Computing Science

Dr. Binay Bhattacharya
Professor of Computing Science

Dr. Pavol Hell
Professor of Computing Science

Dr. Robin Dawes
External Examiner
Assistant Professor of Computer Science
Queen's University

Date Approved: April 13, 1993

Abstract

A visibility relation can be viewed as a graph: the uncountable graph of a visibility relationship between points in a polygon P is called the *point visibility graph* (PVG) of P. In this thesis we explore the use of *perfect graphs* to characterize tractable subproblems of visibility problems. Our main result is a characterization of which polygons are guaranteed to have weakly triangulated PVGs, under a generalized notion of visibility called *$\mathcal{O}$-visibility*.

Let $\mathcal{O}$ denote a set of line orientations. Rawlins and Wood call a set P of points *$\mathcal{O}$-convex* if the intersection of P with any line whose orientation belongs to $\mathcal{O}$ is either empty or connected; they call a set of points $\mathcal{O}$-concave if it is not $\mathcal{O}$-convex. Two points are said to be $\mathcal{O}$-visible if there is an $\mathcal{O}$-convex path between them. A polygon is $\mathcal{O}$-starshaped if there a point from which the entire polygon is $\mathcal{O}$-visible.

Let $\mathcal{O}'$ be the set of orientations of minimal $\mathcal{O}$-concave portions of the boundary of P. Our characterization of which polygons have weakly-triangulated PVGs is based on restricting the cardinality and span of $\mathcal{O}'$. This characterization allows us to exhibit a class of polygons admitting an $O(n^8)$ algorithm for $\mathcal{O}$-convex cover. We also show that for any finite cardinality $\mathcal{O}$, $\mathcal{O}$-convex cover and $\mathcal{O}$-star cover are in NP, and have polynomial time algorithms for any fixed covering number. Our results imply previous results for the special case of $\mathcal{O} = \{0, 90\}$ of Culberson and Reckhow, and Motwani, Raghunathan, and Saran.

Two points are said to be *link-2 visible* if there is a third point that they both see. We consider the relationship between link-2 $\mathcal{O}$-convexity and $\mathcal{O}$-starshapedness, and exhibit a class of polygon/orientation set pairs for which link-2 $\mathcal{O}$-convexity implies $\mathcal{O}$-starshapedness.

Just because some of us can read and write and do a little math, that doesn't mean we deserve to conquer the Universe.

—Kurt Vonnegut, *Hocus Pocus*

Acknowledgments

Like most theses, this one is the product of several years of interaction with an enormous number of people. Of these people, seven in particular deserve special mention. I would like first and foremost to thank my supervisor Tom Shermer for what he modestly calls "just general supervisory stuff". I would also like to thank Binay Bhattacharya and Pavol Hell for serving on my supervisory committee, and Robin Dawes, for graciously consenting to be my external examiner. Finally, I would like to thank Patrice Belleville for his answers, Dave Peters for his questions, and Cheryl for putting up with being a thesis widow.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Definitions

This chapter contains some required definitions, a review of the some of the relevant previous work on visibility, polygon covering, and perfect graphs, and an overview of the organization of this thesis. We start with some standard geometric and graph theoretic definitions.

In this thesis all geometric objects under consideration are sets of points in the Euclidean plane. A *polygonal curve* denotes an ordered set of points $\langle v_1, v_2 \ldots v_n \rangle$ called *vertices* and a set of line segments $\{ \overline{v_1 v_2}, \overline{v_2 v_3} \ldots \overline{v_{n-1} v_n} \}$ called *edges*. A polygonal curve is called *closed* if it has at least 3 vertices and its first and last vertices are identical. A polygonal curve S is called *simple* if no point in the plane belongs to more than two edges of S, and the only points that belong to two edges are vertices of S. We use the term *path* to denote a simple polygonal curve. We define the *interior* of a path or line segment to mean the path or segment exclusive of its endpoints. A simple closed polygonal curve divides the plane into three regions: the bounded interior, the unbounded exterior, and the curve itself. We use the term *simple polygon* to denote a simple closed polygonal curve along with the interior of the curve. The polygonal curve bounding a simple polygon P is called the *boundary* of P. Unless otherwise specified, we take "polygon" to mean "simple polygon"; furthermore, we use "polygon" to refer interchangeably to the representation of a polygon P as a

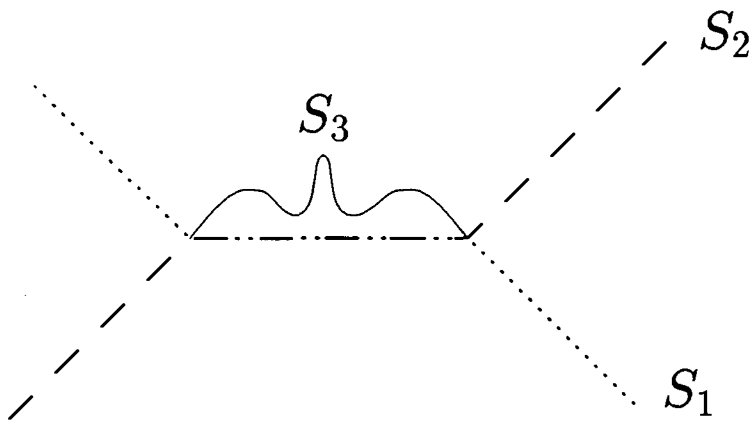

Figure 1.1: A proper crossing S_3 of two curves S_1 and S_2.

polygonal curve and to the point set consisting of the union of the boundary of P and the interior of P. A polygon is called *orthogonal* if it has only vertical and horizontal edges.

A maximal line segment in a polygon P that does not intersect the exterior of P is called a *chord* of P. The shortest path between two points x and y in a polygon P that does not intersect the exterior of the polygon is called the *geodesic* between x and y.

The neighbourhood of any point on the interior of a curve has two well defined sides. We use this notion of side to distinguish two kinds of local curve intersection. Let $\mathrm{int}(S)$ denote the interior of a curve S. A *tangency* of two curves S_1 and S_2 denotes a curve $S_3 \subseteq \mathrm{int}(S_1) \cap \mathrm{int}(S_2)$ such that as we traverse S_1 from one endpoint to the other, S_2 is on the same side of S_1 in the neighbourhood of the first endpoint encountered as in the neighbourhood of the second endpoint of S_3. If S_3 is a tangency for S_1 and S_2, we say that S_1 is tangent to S_2 at S_3. A *proper crossing* of two curves S_1 and S_2 denotes a curve $S_3 \subseteq \int S_1 \cap \int S_2$ such that as we traverse S_1 from one endpoint to another, S_2 is one side of S_1 in the neighbourhood of the first endpoint of S_3 encountered, and on the other side of S_1 in the neighbourhood the second endpoint of S_3 encountered (see Figure 1.1).

A closed polygonal curve S is called *weakly simple* if any pair of distinct points in S

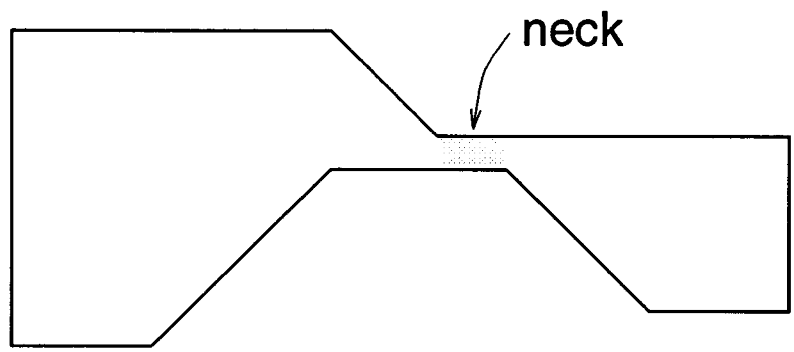

Figure 1.2: A polygon with a narrow neck.

divides S into two polygonal curves that have no proper crossings and the total angle traversed when S is traversed from any point on S is equal to 360 degrees. Like simple polygons, weakly simple polygonal curves have a well defined interior and exterior. A *weakly simple polygon* is defined to be a weakly simple closed polygonal curve, along with the interior of the curve. If a weakly simple polygon P is not simple then some pair of points on the boundary must divide the boundary of P into two polygonal curves that have at least one tangency; these intersections may be considered the limiting case of polygons with narrow "necks" (see Figure 1.2).

In this thesis we shall in particular be interested in the weakly simple subpolygons defined by a chord of a simple polygon. These *half polygons* consist of a single base edge $\overline{lr}$, a set of (possibly zero length) segments $\left\langle \overline{lr_0}, \overline{l_1r_1}, \ldots \overline{l_kr} \right\rangle$ collinear with $\overline{lr}$ and a set of non-intersecting simple polygonal chains $\langle \beta_1 \ldots \beta_k \rangle$ where β_i joins l_i to r_{i-1} (see Figure 1.3). If a half polygon Q contains only one polygonal chain β_i, then Q is called a *hat polygon* (see Figure 1.4). The zero width regions of Q between $\overline{lr_0}$ and the base edge and between $\overline{l_1r}$ and the base edge are called the *brim segments* of Q.

The notion of orientation will play a crucial role in this thesis. We will be interested in the orientation of lines, line-segments, and structures within a polygon. The orientation of a line denotes the smallest angle that the line makes with the positive x-axis. Since lines are undirected, we assume that all line orientations are in the range

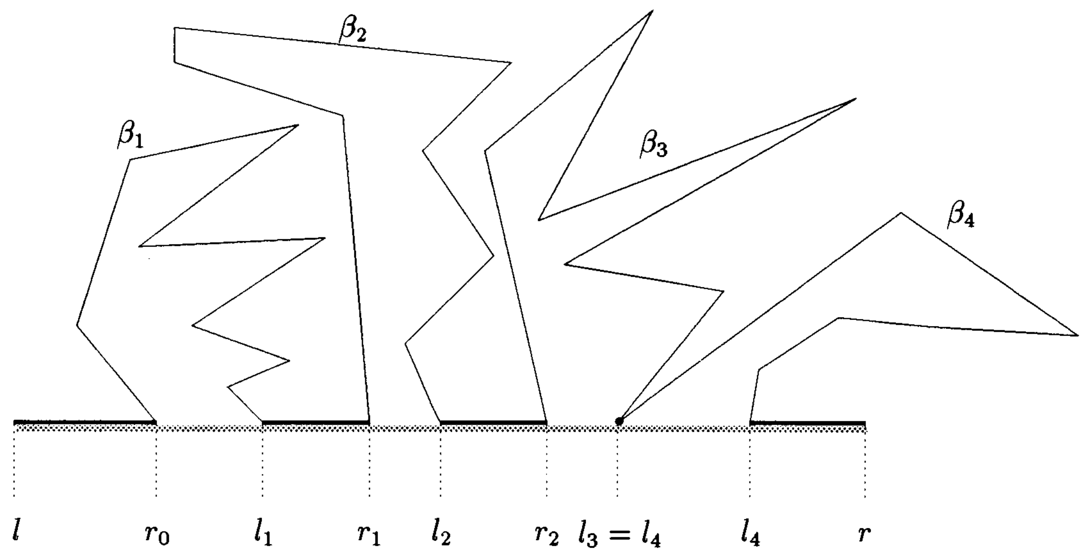

Figure 1.3: A half polygon.

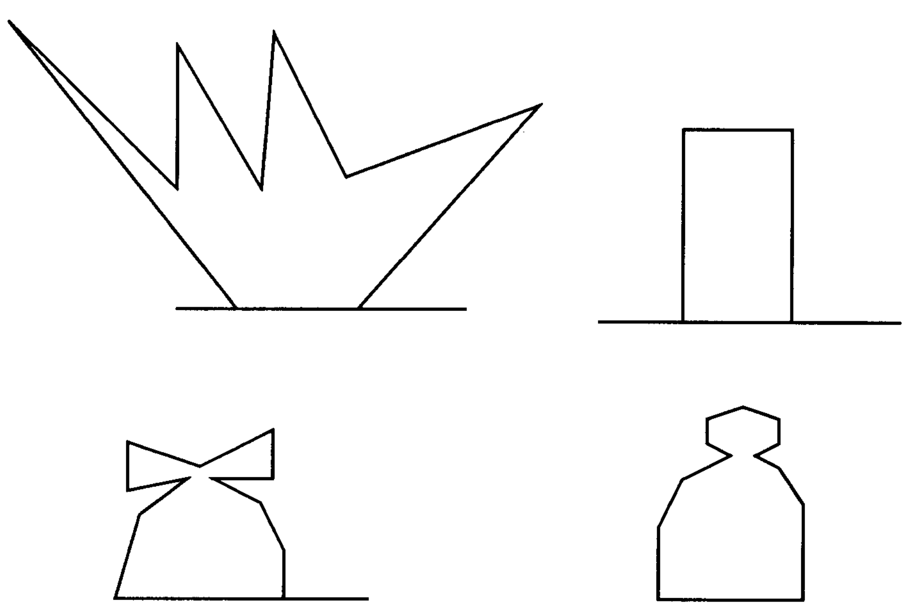

Figure 1.4: Some example hat polygons.

$[0°, 180°)$. All other objects will have orientations in the range $[0°, 360°)$. We assume that normal conventions of modulo arithmetic on orientations; that is we take the range $[a°, b°]$ to mean the range $[(a \bmod 360)°, (b \bmod 360)°]$. If $a > b$, then we take $[a, b]$ to mean $[0°, 360°) \setminus [b, a]$. Analogous rules apply to ranges open on one or both ends. Given a set of orientations $\mathcal{O}'$, we define the *span* of $\mathcal{O}'$ to be the smallest angle α such that

$$(\exists\theta)\, \mathcal{O}' \subseteq [\theta, \theta + \alpha].$$

We now define some necessary graph theoretic terms. See [8] for any omitted graph theoretic definitions. A *directed graph* (*digraph*) G is defined to be a pair (V, A) where

1. V is a set called the *vertices* of G, and

2. A is a subset of $V \times V$ called the *arcs* of G.

If (x, y) is an arc of of a digraph G, we write $x \overset{G}{\rightarrow} y$ or just $x \rightarrow y$ where the digraph G is understood from context.

A *subgraph* of a digraph G denotes a pair (V', A') such that $V' \subseteq V$ and $A' \subseteq A$. If $H = (V', A')$ is a subgraph of G and for any pair of vertices x and y of H,

$$(x \overset{G}{\rightarrow} y) \Rightarrow (x \overset{H}{\rightarrow} y)$$

then H is called an *induced subgraph* of G.

An ordered set of arcs $S = \langle\, (v_0, v_1), (v_1, v_2), \ldots (v_{k-1}, v_k)\,\rangle$ is called a *path* (k-path) from v_0 to v_k. If $v_0 = v_k$ then S is called a *cycle* (k-cycle). If every vertex occurs in at most two arcs, then S is called simple. If

$$(i + 1 \not\equiv j \bmod k) \Rightarrow v_i \not\rightarrow v_j$$

then S is called a *chordless* k-cycle. A digraph that contains no cycles is called an *directed acyclic graph* (DAG). A DAG $G = (V, A)$ that contains $|V| - 1$ edges and a vertex r such that there exists a path from r to every other vertex in G is called a *tree*; r is called the *root* of G.

A digraph $G = (V, A)$ is called a *undirected graph* (*graph*) if

$$(x \overset{G}{\rightarrow} y) \Rightarrow (y \overset{G}{\rightarrow} x)$$

Given an undirected graph G, if (v, v') and (v', v) are arcs of G, we say that $\{v, v'\}$ is an *edge* of G, and denote the edge as vv' or $v'v$ interchangeably. If vv' is an edge of G we say that v and v' are neighbours in G and write $v \overset{G}{\sim} v'$ or just $v \sim v'$ where the graph G is understood. The set of vertices of a graph G will be denoted $V(G)$ and the set of edges will be denoted $E(G)$. An undirected graph G is equivalently defined by the pair (V, E) where $V = V(G)$ and $E = E(G)$. Unless otherwise specified, all graphs in this thesis are undirected.

The *square* of a graph $G = (V, E)$ is a graph $G' = (V, E')$ where

$$E' = \left\{ v_1 v_3 \mid (\exists v_2)\, v_1 \overset{G}{\sim} v_2 \wedge v_2 \overset{G}{\sim} v_3 \right\}.$$

Given a graph $G = (V, E)$, we say that vertices v_0 and v_1 are *equivalent* and write $v_0 \overset{\mathcal{N}}{\equiv} v_1$ if v_0 and v_1 have the same neighbours in G. The maximal equivalence classes of the relation "$\overset{\mathcal{N}}{\equiv}$" are called the *vertex equivalence classes* of G. Two vertex equivalence classes V_0 and V_1 are said to be *adjacent* if

$$((v_0 \in V_0) \wedge (v_1 \in V_1)) \Rightarrow (v_0 \sim v_1).$$

A graph H is called the *quotient graph* of a graph G if each vertex of H corresponds to a vertex equivalence class of G and there is an edge between two vertices v_0 and v_1 of H if and only if the corresponding vertex equivalence classes are adjacent. If each vertex of H is labeled with the cardinality of the corresponding vertex equivalence class, H is called the *labeled quotient graph* of G.

Let G be a graph. Let $\omega(G)$, the *clique number* of G, denote the size of the largest complete subgraph of G. Let $\chi(G)$, the *chromatic number* of G, denote the minimum number of colours needed to colour G. Let $\alpha(G)$, the *independence number* of G, denote the size of the largest independent set of G. Let $k(G)$, the *clique cover number* of G, denote the minimum number of cliques needed to cover G.

A graph G is called *reflexive* if for any vertex v of G, $v \overset{G}{\sim} v$. Let G be a reflexive graph. Let H be the quotient graph of G. Let $\bar{G}$ be the graph theoretic complement of G. We shall make use of the following facts:

$$\alpha(G) = \alpha(H) \tag{1.1}$$

$$k(G) = k(H) \tag{1.2}$$

$$\alpha(G) = \omega(\bar{G}) \tag{1.3}$$

$$k(G) = \chi(\bar{G}). \tag{1.4}$$

In this thesis we are concerned with both finite and infinite graphs. For notational convenience, let $\aleph_0$ denote the cardinality of the integers and $\aleph_1$ the cardinality of the reals (in this thesis we do not assume the continuum hypothesis). A graph will be called *finite* if $|V(G) \cup E(G)|$ is finite, *countable* if $|V(G) \cup E(G)| = \aleph_0$ and *uncountable* if the cardinality of $|V(G) \cup E(G)| \geq \aleph_1$. For a treatment of infinite graphs, see [14, 33, 34]. Unless otherwise specified, graphs in this thesis are finite.

1.2 Perfect Graphs

Many graphs theoretic problems are "universal" in the sense that many problems can be reduced to them. For example, Roberts [39] gives the following problems that can be reduced to finding the chromatic number of a graph: meeting scheduling, register allocation in compilers, television and radio frequency assignments, and garbage truck routing. Unfortunately these same universal graph theoretic problems are often NP-Hard on general graphs. The most common solution is to show that the graphs that arise from a particular problem belong to some "nice" subclass for which the graph theoretic problem is tractable. Examples of such classes of graphs include trees, bipartite graphs, triangulated graphs, tree decomposable graphs, planar graphs, interval graphs, circular arc graphs, and maximal outerplanar graphs. A class of graphs that contains many of these classes is the class of perfect graphs. Let G_A denote the subgraph of G induced by the vertex set A. A graph G is called χ-perfect if

$$(\forall A \subseteq V(G)) \;\; \omega(G_A) = \chi(G_A).$$

A graph G is called α-perfect if

$$(\forall A \subseteq V(G)) \;\; \alpha(G_A) = k(G_A).$$

A graph is called *perfect* if it is both α-perfect and χ-perfect. The notions of α-perfection and χ-perfection were introduced by Berge [5, 6], who conjectured that a graph was α-perfect if and only if it was χ-perfect. Lovász proved the following:

Theorem 1.1 (Lovász[28]) *For a finite undirected graph $G = (V, E)$, the following statements are equivalent:*

(Γ_1) *G is χ-perfect,*
(Γ_2) *G is α-perfect,*
(Γ_3) $(\forall A \subseteq V)\, \omega(G_A)\alpha(G_A) \leq |A|$.

For uncountable graphs, the equivalence of α-perfection and χ-perfection does not hold. Furthermore, classes of uncountable graphs that are generalizations of classes of finite perfect graphs to the uncountable case are not necessarily α-perfect *or* χ-perfect [48].

Grötschel, Lovász and Schrijver [21] provided a polynomial algorithm based on the ellipsoid method for the problems of maximum clique, chromatic number, independent set, and clique cover on perfect graphs. Furthermore, many efficient and practical polynomial algorithms are known for subclasses of perfect graphs [7, 20].

A *hole* is a chordless cycle in a graph of size greater than 3. An *antihole* is a hole in the complement graph, or the graph theoretic complement of a hole. A graph is called triangulated if it contains no hole of size greater than 3. A graph is called *weakly triangulated* if it contains no hole or antihole of size greater than 4. Hayward [24] has shown that finite weakly triangulated graphs are perfect. Since a 5-antihole is a 5-hole and any antihole of size greater than 5 contains a 4-hole, triangulated graphs are weakly triangulated. Raghunathan [36] found an algorithm that finds a maximum clique and a minimum colouring of a weakly triangulated graph $G = (V, E)$ in $O(ev^2)$ time where $e = |E|$ and $v = |V|$.

1.3 Visibility

Visibility is a central notion in computational geometry. Informally, visibility problems are concerned with whether or not pairs of geometric objects within a set of obstacles can "see" one and other. Recent research has considered generalized visibility or "reachability" problems, where the notion of straight line visibility is generalized to reachability by some sort of constrained path. In this section we review the traditional notion of visibility, and some of the generalizations of visibility that researchers have investigated. In particular we describe the notion of *restricted orientation visibility* that this thesis investigates.

Two points x and y in a polygon P are said to be *visible* if the line segment between them does not intersect the exterior of P. A set of points P is said to be *convex* if every pair of points in P is visible. A set of points P is said to be *starshaped* if P contains some point k from which all of P is visible. A set H of points in P is called a *hidden set* if no pair of points in H is visible.

Several authors have considered alternative notions of visibility, that is relationships between points in a polygon that preserve some of the essential features of visibility. Munro, Overmars, and Wood [32] investigated *region visibility* in point sets where two points are considered visible if there is a region having some property (e.g. a square) containing only those two points. Shermer and Toussaint consider [46] *geodesic visibility* where two points x and y are considered geodesically visible with respect to a a subpolygon Q of a polygon P if the geodesic (in P) between them is contained in Q. Schuirer, Rawlins, and Wood [40] consider visibility in abstract convexity spaces.

Orthogonal polygons are commonly studied in computational geometry both because many problems intractable or open on general polygons become tractable on orthogonal polygons, and because they arise in many important applications (e.g., in VLSI and image processing). Keil [26], Culberson and Reckhow[13], and Motwani, Raghunathan, and Saran [30, 31] investigated the notion of *orthogonal visibility* in orthogonal polygons. In many applications not only the boundary but also internal paths (e.g. wires in a chip) are constrained to be chains of orthogonal line segments;

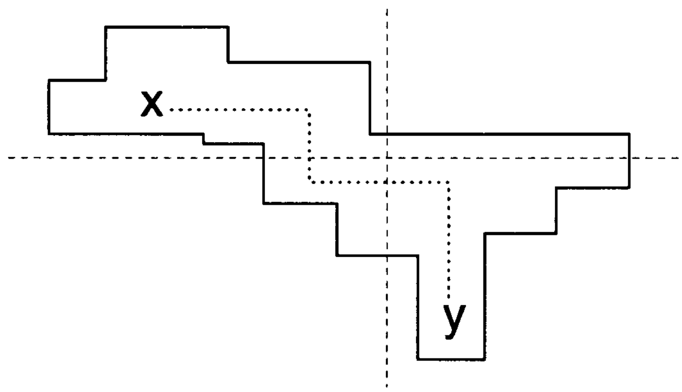

Figure 1.5: An orthogonally convex polygon

in this case it becomes natural to say that two points x and y are *orthogonally visible* if there is a path between them that is monotone with respect to both axes (see Figure 1.5), since no shorter path is realizable under the constrained geometry. A polygon P is called *orthogonally convex* if every pair of points in P is orthogonally visible (see Figure 1.5); this is equivalent to requiring that the intersection of P with any horizontal or vertical line be empty or connected. A polygon P is called an *orthogonal star* if there exists some set of points K in P such that every point in P is orthogonally visible from each point in K.

Recent VLSI designs, motivated by the need for increased device density, allow line segments to have additional orientations [9, 29]. Widmayer et al. [49] and Souvaine and Bjorling-Sachs [47] find efficient algorithms for several polygon union and intersection problems where the edges of the input polygons are restricted to have some bounded number of orientations. Rawlins and Wood [37, 38] generalized orthogonal convexity to the notion of *restricted orientation convexity*, or *$\mathcal{O}$-convexity.* Let $\mathcal{O}$ denote a fixed, but unspecified set of line orientations. A line is called an *$\mathcal{O}$-line* if it has an orientation in $\mathcal{O}$. A set P of points is called *$\mathcal{O}$-convex* if the intersection of P with any $\mathcal{O}$-line is either empty or connected (see Figure 1.6). A set P of points is called *$\mathcal{O}$-concave* if it is not $\mathcal{O}$-convex. A finite $\mathcal{O}$-convex path is called a *staircase*. Restricted orientation convexity is a generalization of both orthogonal convexity ($\mathcal{O} = \{\, 0°, 90° \,\}$) and the standard notion of convexity ($\mathcal{O} = [0°, 180°)$).

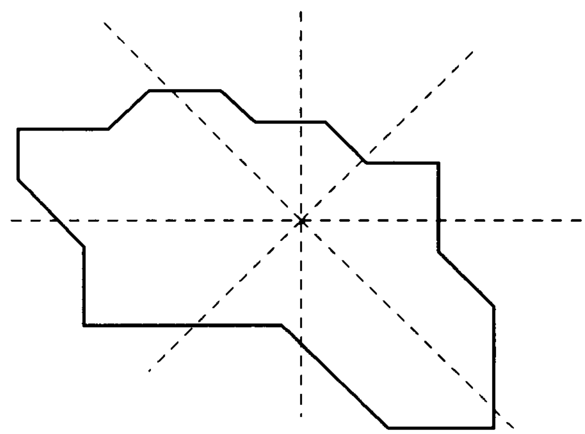

Figure 1.6: A polygon that is $\mathcal{O}$-convex for $\mathcal{O} = \{\, 0°, 45°, 90°, 135° \,\}$

We may assume without loss of generality that any set $\mathcal{O}$ of orientations contains a pair of orthogonal orientations. Following Rawlins [37], for any line orientation θ we define an *orthogonalizing transformation* T_θ follows:

$$T_\theta = \begin{bmatrix} \sin\theta & 0 \\ -\cos\theta & 1 \end{bmatrix}.$$

The transformation T_θ maps horizontal lines to horizontal lines and lines with orientation θ to vertical lines. Let $T_\theta(\phi)$ denote the orientation in the transformed space of line with orientation ϕ in the untransformed space. Let $\tilde{\mathcal{O}}_\theta$ denote the set $\{\, T_\theta(\phi) \mid \phi \in \mathcal{O} \,\}$. Since an orthogonalizing transformation is *affine* (see [17]), collinearity and the order of points along a line is preserved in the transformed space.

Observation 1.1 *A set of points P is $\mathcal{O}$-convex if and only if $T_\theta(P)$ is $\tilde{\mathcal{O}}_\theta$-convex for any θ.*

We may assume without loss of generality that some orientation ϕ in any set $\mathcal{O}$ of orientations is $0°$. By Observation 1.1, we may assume that any other single orientation in $\mathcal{O}$ is orthogonal to ϕ.

For paths, the following alternative characterization of $\mathcal{O}$-convexity will prove useful:

Observation 1.2 *A path S is $\mathcal{O}$-convex if and only if there is no $\mathcal{O}$-line tangent to the interior of S.*

Given two points x and y in a polygon, we say that x is *$\mathcal{O}$-visible* to y (x sees y) and write $x{\sim}y$ if there is a staircase between x and y that does not intersect the exterior of the polygon. The following lemma establishes the standard relationship between visibility and convexity:

Lemma 1.3 (Rawlins and Wood [38]) *If a point set P is connected, then P is $\mathcal{O}$-convex if and only if for any pair of points p and q in P, p sees q.*

In this thesis we are concerned with $\mathcal{O}$-visibility inside polygons. Our main results are for finite $\mathcal{O}$, but we investigate within as general a framework as possible, and some of our results provide insight into the combinatorial structure of $\mathcal{O}$-visibility for infinite $\mathcal{O}$, including traditional straight line visibility. Our definition of $\mathcal{O}$-visibility is slightly at variance with that given in [37, 38] since we require paths to be polygonal and Rawlins and Wood do not. Within a polygon, however, the two definitions are equivalent.

A set of points K contained in a polygon P is called the *$\mathcal{O}$-kernel* of P if every point in P is $\mathcal{O}$-visible from each point in K. A polygon is called *$\mathcal{O}$-starshaped* if it contains a non-empty $\mathcal{O}$-kernel. The smallest $\mathcal{O}$-convex polygon Q that contains a polygon P is called the $\mathcal{O}$-hull of P. Rawlins [37] provides efficient algorithms for computing the $\mathcal{O}$-kernel and the $\mathcal{O}$-hull of a polygon when $\mathcal{O}$ consists of a finite number of closed ranges.

Research on visibility has to a great extent taken a combinatorial approach— the general technique has been to reduce these geometric problems to problems on graphs and then apply algorithmic results from graph theory. In addition to providing insight into the computational complexity of visibility problems, this approach separates the combinatorial aspects of a problem from the geometric aspects. This is desirable, because geometric algorithms are prone to representational difficulty and numerical instability.

Shermer [43] introduced the notion of a *point visibility graph* as a unifying combinatorial framework for visibility problems. The point visibility graph (PVG) of a polygon P denotes the uncountable graph $G = (V, E)$ whose vertices are the points of P and whose edges are exactly those pairs (x, y) for which x is visible to y in P. Many visibility problems can be viewed as combinatorial problems on the PVG but the infinite size of these graphs prevents a straightforward application of graph theoretic algorithms. The combinatorial approach to visibility problems in the literature can be viewed as embodying two methods for dealing with the continuous nature of PVGs. The first method is to consider particular finite subgraphs of the PVG, most commonly those induced by the vertices of the polygon [16, 19]. Although results on specific problems have been obtained with this method, in general the finite substructures involved fail to capture many important visibility properties. The second method is to find classes of polygons, and appropriate definitions of visibility, for which there is a finite combinatorial representation for the entire PVG [13, 30, 43, 45]. Herein we are concerned with this second method.

1.4 Polygon Covering

Polygon covering is a class of problem closely associated with visibility problems. For any property Π of a set of points, the Π-covering problem is defined as follows:

Instance: A polygon P and an integer k.

Question: Does there exist a family $\mathcal{Q}$ of subsets of P such that

1. Each element of $\mathcal{Q}$ has property Π,
2. $P = \bigcup_{Q \in \mathcal{Q}} Q$, and
3. $|\mathcal{Q}| \leq k$.

If k is the smallest integer for which the answer to the Π-covering problem for a polygon P is yes, then we say that k is the Π-*covering number* of P. Covering problems are well studied in computational geometry, particularly those defined by the visibility properties convexity and starshapedness. Many versions of the problem are known to be NP-Hard (for a treatment of the theory of NP-Completeness, see

[18]). Star cover (i.e where $\Pi =$ "starshapedness") was shown to be NP-Hard by Lee and Lin [27] and Aggarwal [1]. Convex cover was shown to be NP-Hard by Culberson and Reckhow [12] and independently by Shermer [41]. Culberson and Reckhow further showed that even the restricted case of covering orthogonal polygons with rectangles is NP-Hard. Aupperle et al. [2] show that covering orthogonal polygons with squares is NP-Complete if the polygons to be covered may contain holes. In most cases, the obvious technique of guessing the sets of a minimal cover fails to establish that polygon covering problems are in NP, because it is not known if the locations of the vertices of the covering subpolygons are representable in a polynomial number of bits [35].

Since convex cover is special case of $\mathcal{O}$-convex cover and star cover is a special case of $\mathcal{O}$-star cover, we know that both of $\mathcal{O}$-convex cover and $\mathcal{O}$-star cover are NP-Hard. In this thesis we investigate the application of graph theoretic techniques to characterize tractable subproblems of $\mathcal{O}$-convex cover and $\mathcal{O}$-starshaped cover. Because we are considering a whole spectrum of different kinds of visibility, it does not suffice to consider restricted classes of polygons, since the visibility properties of a polygon change with type of visibility under consideration. We therefore introduce the notion of a *visibility instance*, defined to be a pair $(P, \mathcal{O})$ where P is a polygon and $\mathcal{O}$ is a set of orientations. The PVG of a visibility instance $(P, \mathcal{O})$ denotes the $\mathcal{O}$-visibility PVG ($\mathcal{O}$-PVG) of P.

Culberson and Reckhow [13] introduced the term *dent* to denote an edge of an orthogonal polygon with two reflex endpoints. Previous authors [13, 30] showed that orthogonal polygons with at most three dent orientations have weakly triangulated orthogonal visibility PVGs, and used this to give polynomial algorithms for covering orthogonal polygons having at most 3 dent orientations with orthogonally convex polygons. Culberson and Reckhow also give a polynomial algorithm for a subset of orthogonal polygons with all four possible dent orientations. Motwani et al. [31] have shown that the graph theoretic square of any orthogonal visibility PVG is weakly triangulated, and used this result to give a polynomial algorithm for orthogonal star cover. This thesis considers to what extent these results generalize to $\mathcal{O}$-visibility for more general sets of orientations $\mathcal{O}$. We give a natural extension of Culberson and Reckhow's notion of dent to more general sets of orientations and show that to guarantee a weakly triangulated $\mathcal{O}$-PVG, not only must a visibility instance have a maximum of 3 dent orientations, but if it does have 3 dent orientations, the span of these orientations must be at least 180°. We also consider under what conditions $\mathcal{O}$-star cover can be reduced to clique cover on the graph theoretic square of the PVG.

The rest of this thesis is organized as follows. In Chapter 2 we generalize the PVG discretization techniques of Culberson and Reckhow to more general sets of orientations and show how to compute the labeled quotient graph of the PVG of P when $|\mathcal{O}|$ is finite. We also show that the existence of this finite combinatorial representation of the $\mathcal{O}$-PVG (for finite $\mathcal{O}$) implies that $\mathcal{O}$-convex cover and $\mathcal{O}$-star cover are in NP, and provide simple brute force algorithms for these problems that are polynomial for any fixed covering number. In Chapter 3 we characterize when the PVG of a visibility instance is necessarily weakly triangulated. We give a polynomial algorithm for $\mathcal{O}$-convex cover on visibility instances that meet these conditions. In Chapter 4 we consider other types of structure present in PVGs. In Chapter 5 we investigate the relationship between clique cover on the square of the PVG and star cover. In Chapter 6 we present some conclusions and directions for future work.

Chapter 2

Cell Visibility Graphs

2.1 Dent Decompositions

Point visibility graphs provide a unifying combinatorial framework for visibility problems, but without a finite combinatorial representation for the point visibility graphs in question, such a framework is of little computational interest. Shermer [44] defined a *pure visibility problem* to be one solvable by solving a graph theoretic problem on the point visibility graph and mapping the solution back into a geometric object. Problems satisfying this definition include convex cover, star cover, and hidden set. In this chapter we show that for finite cardinality $\mathcal{O}$, there exists a polynomial size combinatorial representation for the $\mathcal{O}$-PVG of any polygon. We can solve pure visibility problems by computing this representation and then applying combinatorial algorithms. Our general approach will be to consider the decomposition of a polygon into visibility equivalence classes called *dent cells* and to consider the visibility graph of these cells (the *cell visibility graph*). We then show how to compute the cell visibility graph and consider some straightforward applications of the resulting data structure. We start by introducing a decomposition of the polygon called the *dent decomposition* and show that its cells are visibility equivalence classes.

An *oriented chord* is defined to be a pair $\delta = (\gamma, \theta)$ where γ is a chord of P, and θ is one of the two orientations perpendicular to γ. Each line of orientation ϕ can give rise to two chord orientations, $\phi + 90°$ and $\phi - 90°$, hence chord orientations are in

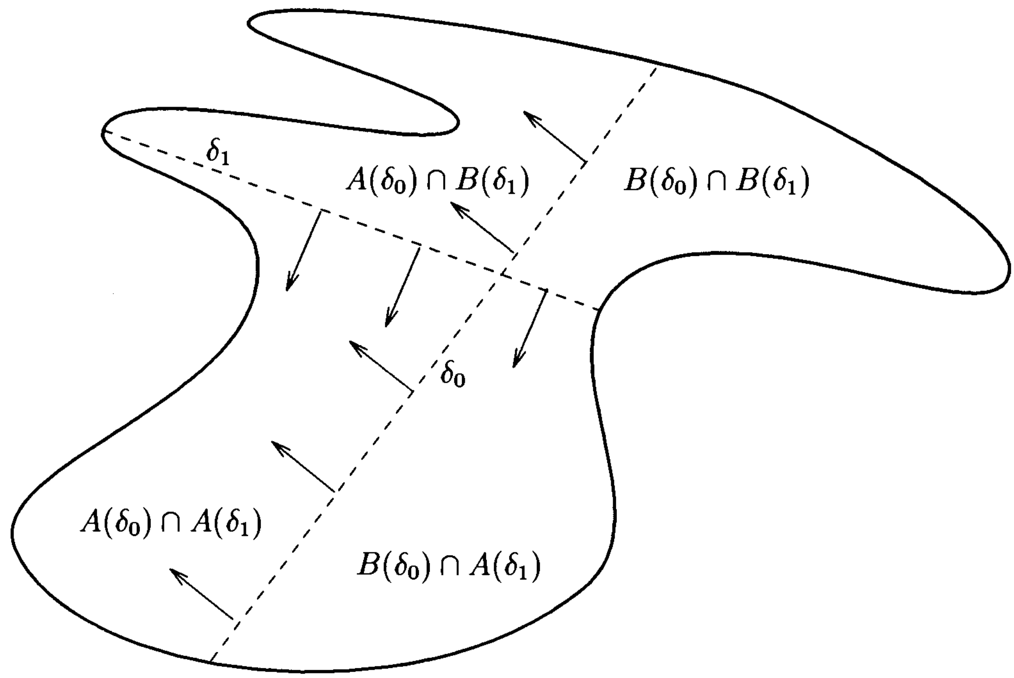

Figure 2.1: The partition of a polygon induced by a pair of oriented chords

the range $[0°, 360°)$. If $\delta = (\gamma, \theta)$ is an oriented chord, $\gamma(\delta)$ denotes γ and $\theta(\delta)$, called the orientation of δ, denotes θ. When there is no ambiguity, we sometimes use δ to mean $\gamma(\delta)$. We call an oriented chord $\delta = (\gamma, \theta)$ an *$\mathcal{O}$-chord* if the orientation of γ (which is distinct from θ, orientation of δ) belongs to $\mathcal{O}$.

Let R be a weakly simple polygonal region of P, and let $\delta = (\gamma, \theta)$ be an oriented chord such that a segment of γ is an edge of R. We say that δ *faces into* (respectively *faces out of*) R if some ray from γ with orientation θ (respectively $\theta + 180°$) is in R in the neighbourhood of γ. Each oriented chord divides the polygon into two weakly simple subpolygons; $A(\delta)$ denotes the one that δ faces into and $B(\delta)$ denotes the one that δ faces out of (see Figure 2.1). By convention δ is included in $A(\delta)$ but not in $B(\delta)$. If $x \in A(\delta)$ then we say that x is *$\mathcal{O}$-above* δ; conversely, if $x \in B(\delta)$, we say that x is *$\mathcal{O}$-below* δ. Where there is no ambiguity, we take "above" to mean $\mathcal{O}$-above and "below" to mean $\mathcal{O}$-below.

If a path S from x to y goes from $\mathcal{O}$-below an oriented chord δ to $\mathcal{O}$-above it, then we say that S crosses δ *upward*. Analogously, if a path S goes from $\mathcal{O}$-above δ to $\mathcal{O}$-below it, we say that S crosses δ *downward*. Since $\gamma(\delta)$ is $\mathcal{O}$-above δ, a curve S that is tangent to $\gamma(\delta)$ and $\mathcal{O}$-below δ in the neighbourhood of $\gamma(\delta)$ crosses δ both upwards and downwards.

Lemma 2.1 *If a path S crosses an $\mathcal{O}$-chord δ upward, and crosses an $\mathcal{O}$-chord δ' of the same orientation downward, then S is not $\mathcal{O}$-convex.*

Proof. Let θ be the orientation of the two oriented chords. Let $\mathcal{L}$ be the farthest line in the direction θ that is parallel to δ and intersects S. The path S crosses some chord of orientation θ $\mathcal{O}$-below $\mathcal{L}$ downward, and some chord of orientation θ $\mathcal{O}$-below $\mathcal{L}$ upward, so in particular S crosses upward and downward a line $\mathcal{L}'$ parallel to $\mathcal{L}$, arbitrarily close to $\mathcal{L}$ in direction $180 + \theta$. Since it intersects an $\mathcal{O}$-line in a disconnected set, S is not $\mathcal{O}$-convex. ∎

A vertex v of a polygon P is called *reflex* if the angle formed by the two edges meeting at v, inside the polygon, is greater than $180°$. An edge of a polygon is called reflex if both of its endpoints are reflex. Let τ be a reflex vertex or edge such that

1. There exists some $\mathcal{O}$-chord $\delta = (\gamma, \theta)$ such that γ is tangent to τ,

2. Any ray from τ in the direction θ is inside P in the neighbourhood of τ.

In this situation we call the ordered pair $D = (\tau, \theta)$ a *dent*, and call δ the *dent chord* of D, written $\vec{D}$. A given reflex vertex may be part of more than one dent, but a given reflex edge may be part of at most one. If an edge or vertex τ participates in some dent, we say that there is a dent *at* τ. Given a dent $D = (\tau, \theta)$, $\tau(D)$ denotes τ and $\theta(D)$ denotes θ. We sometimes use the term dent and the notation D to refer to $\tau(D)$. We use the *orientation of* D to mean the orientation of $\vec{D}$, $A(D)$ to denote $A(\vec{D})$, and $B(D)$ to denote $B(\vec{D})$. Given a set of dents $\mathcal{D}$, $D \in \mathcal{D}$ is called a *maximal* element of $\mathcal{D}$ if

$$(\nexists D' \in \mathcal{D})\ B(D) \subset B(D').$$

We may further subdivide $B(D)$. The dent chord $\vec{D}$ can be thought of as two disjoint collinear line segments from $\tau(D)$ to the polygon boundary. We define $B_l(D)$

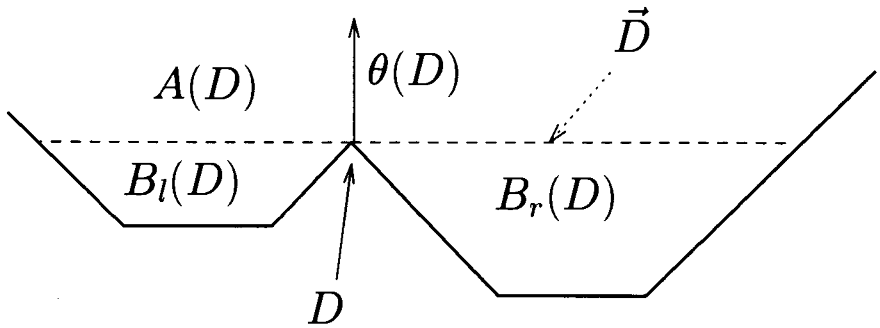

Figure 2.2: A dent, and the three subpolygons induced by it.

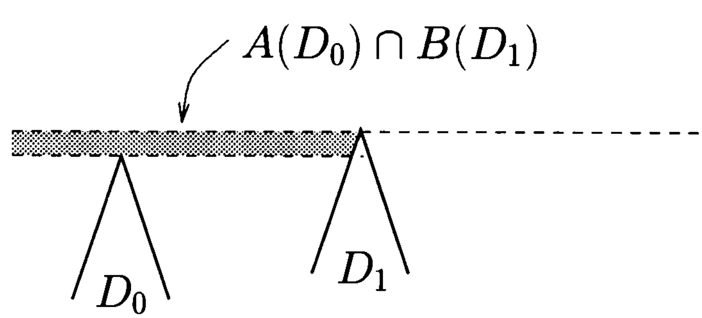

Figure 2.3: The zero width region between two coincident dent chords with the same orientation

(respectively $B_r(D)$) to be the weakly simple subpolygon induced by D containing the polygon edge clockwise (respectively counterclockwise) from $\tau(D)$ (see Figure 2.2). We call $B_l(D)$ and $B_r(D)$ the two *sides* of $B(D)$. For any point $p \in B(D)$, we use $B_p(D)$ to denote the side of $B(D)$ containing p and $B_{\bar{p}}(D)$ denote the other side of $B(D)$.

In the degenerate case, two or more dent chords may be coincident (i.e. have the same endpoints). We assume that the dents in the boundary of the polygon are ranked in some arbitrary but fixed way.

- Suppose two coincident dent chords $\vec{D}_0$ and $\vec{D}_1$ have the same orientation. Without loss of generality suppose D_1 is the higher ranked of the two dents. In this case we assume that $\vec{D}_0$ stops at $\tau(D_1)$ and consider there to be a zero width

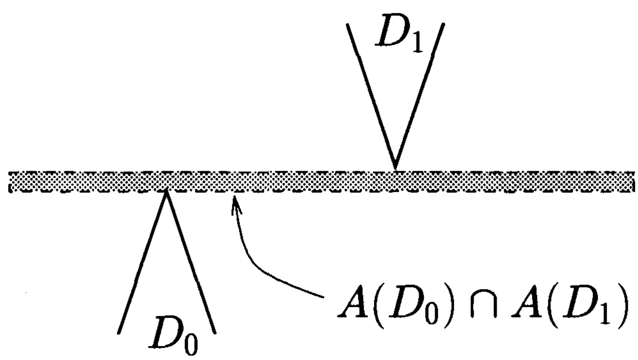

Figure 2.4: The zero width region between two opposite facing coincident dent chords

region between $\vec{D}_0$ and $\vec{D}_1$ that is $\mathcal{O}$-above D_0 but $\mathcal{O}$-below D_1 (see Figure 2.3).

- Suppose two coincident dent chords $\vec{D}_0$ and $\vec{D}_1$ have orientations that differ by exactly 180°. In this case we consider there to be a zero width region between $\vec{D}_0$ and $\vec{D}_1$ that is $\mathcal{O}$-above both D_0 and D_1 (see Figure 2.4).

We define a relation "$\prec$" (read "below") between oriented chords and points in a polygon as follows:

$$\begin{aligned} p \prec \delta \quad &\text{if} \quad p \in B(\delta) \\ \delta \prec p \quad &\text{if} \quad p \in A(\delta) \end{aligned}$$

We shall take $a \succ b$ (read "above") to be equivalent notation for $b \prec a$. We use $p \prec D$ to mean $p \prec \vec{D}$.

Let $\mathcal{D}$ be the set of all dents in the polygon boundary. We say that points p and p' are *dent equivalent* and write $p \overset{\mathcal{D}}{\equiv} p'$ if for any $D \in \mathcal{D}$

$$\begin{aligned} p \in A(D) \quad &\Leftrightarrow \quad p' \in A(D) \\ p \in B_l(D) \quad &\Leftrightarrow \quad p' \in B_l(D) \\ p \in B_r(D) \quad &\Leftrightarrow \quad p' \in B_r(D) \end{aligned}$$

We call the maximal equivalence classes of the relation "$\overset{\mathcal{D}}{\equiv}$" *dent cells*. For uniformity of terminology, we consider a degenerate cells between coincident dent chords

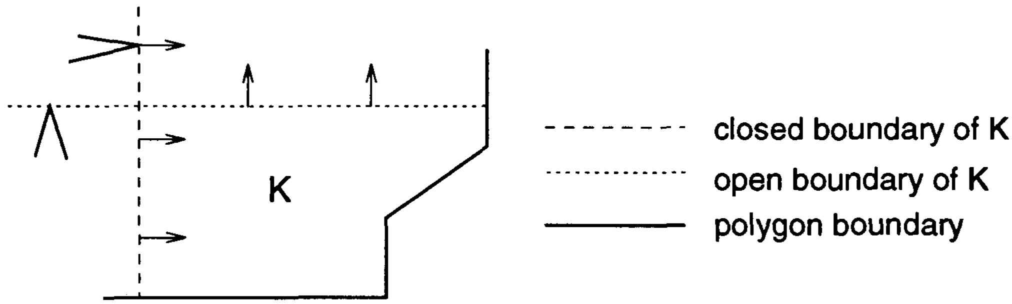

Figure 2.5: A dent cell K.

to be bounded by a polygon with two or more zero length edges, so that every pair of neighbouring dent cells meets at a well defined edge. If an edge of a dent cell boundary is a segment of a dent chord that faces out of the cell, we consider the cell to be open along that edge (see Figure 2.5). If an edge of a dent cell is a segment of a dent chord that faces into the cell, or a segment of a polygon edge, then we consider the cell to be closed along that edge (see Figure 2.5).

We call the decomposition of the polygon into dent cells the *dent decomposition* of the polygon.

A *separating dent* for two points x and y is a dent D such that $x \in B_l(D)$ and $y \in B_r(D)$ or vice versa. A dent D is called a *supporting dent* for a curve S if S is tangent to $\tau(D)$.

Lemma 2.2 *If x does not see y then the geodesic from x to y is supported by some separating dent for x and y.*

Proof. Suppose x does not see y. Consider a geodesic S from x to y. No path from x to y inside P is $\mathcal{O}$-convex, so from Observation 1.2, there must be some point t on S where an $\mathcal{O}$-line $\mathcal{L}$ is tangent to the curve. Suppose S is not supported by a portion of the polygon boundary at t; then we can find an $\mathcal{O}$-line $\mathcal{L}'$ parallel to $\mathcal{L}$ intersecting S twice in the neighbourhood of t interior to P. Let x' and y' be the intersection points of S and $\mathcal{L}'$ (see Figure 2.7). We can create a shorter path from x to y by replacing the subpath of S between $\mathcal{L}$ and $\mathcal{L}'$ with the segment $\overline{x'y'}$; but this contradicts our

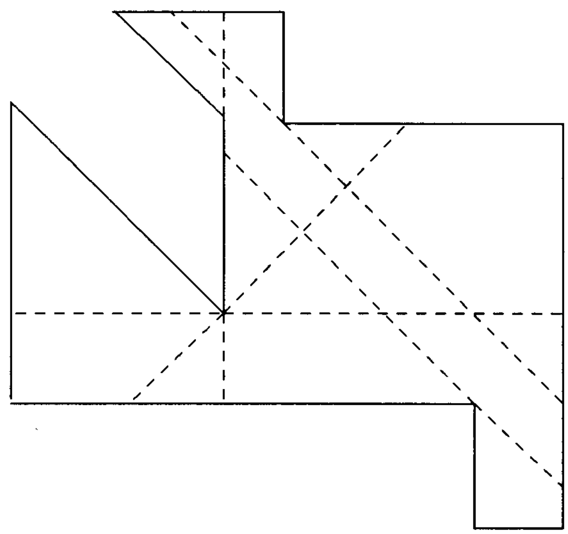

Figure 2.6: A dent decomposition of a polygon; $\mathcal{O} = \{ 0°, 45°, 90°, 135° \}$.

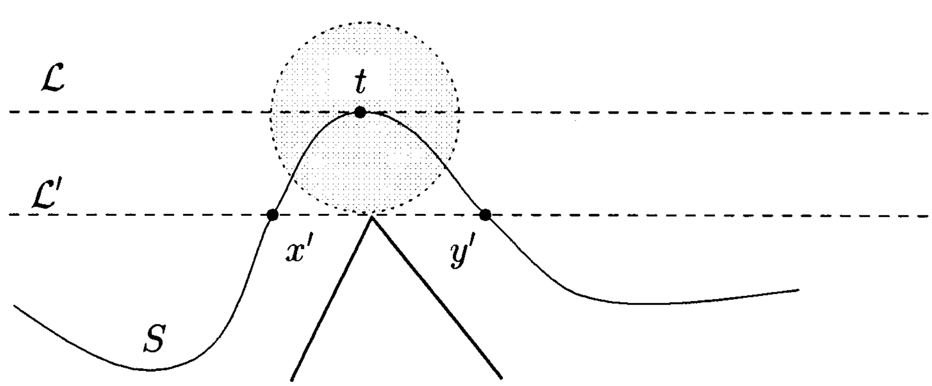

Figure 2.7: If a tangent point of a path is not supported by the polygon boundary, we can find a shorter path.

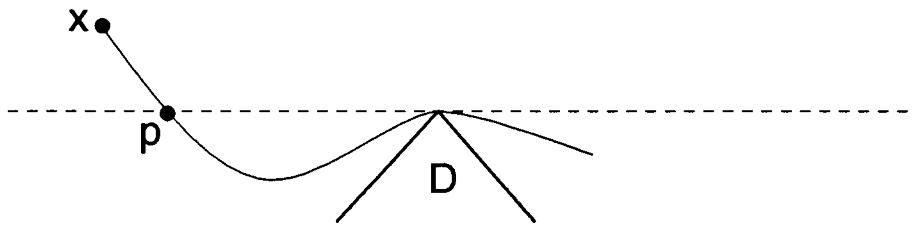

Figure 2.8: x is above the supporting dent.

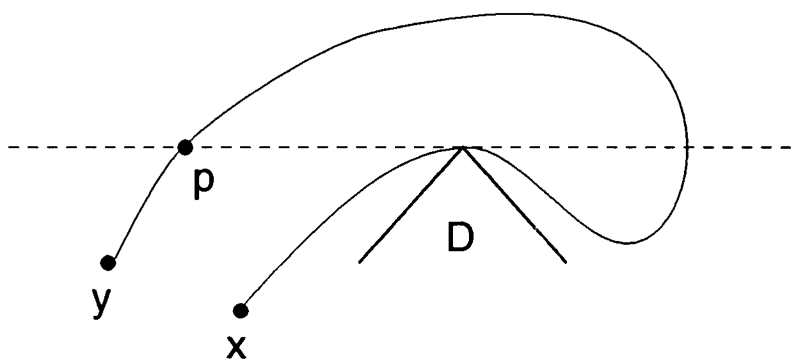

Figure 2.9: x is below the supporting dent on the same side as y.

assumption that S is a geodesic, so S must be supported by the polygon boundary at t. It follows that the $\mathcal{O}$-line $\mathcal{L}$ is also tangent to the polygon boundary at t; hence $D = (t, \theta)$ is a dent for some θ. Since $\tau(D)$ is tangent to S, D is a supporting dent for S.

We now argue that D must be a separating dent for x and y. We first argue that both x and y must be $\mathcal{O}$-below D. Let $B^x(D)$ be the side of $B(D)$ intersected last before x as S is traversed from y to x. Since $\vec{D}$ is tangent to S at D, S is in $B^x(D)$ in the neighbourhood of D. Suppose $x \succ D$; then S must exit $B^x(D)$ by crossing $\vec{D}$ at some point p such that the segment from p to D (where S entered $B^x(D)$) does not intersect the boundary of P (see Figure 2.8). Since joining the entry and exit points by a line segment would produce a shorter path, this is a contradiction, and $x \prec D$. By a symmetric argument, $y \prec D$.

We have established that both x and y are $\mathcal{O}$-below D. Suppose x and y are on the same side of D. Since S is in $B_{\bar{x}}(D)$ in the neighbourhood of D, S must go from $B_x(D)$ to $A(D)$ by crossing $\vec{D}$ at a point p such that the segment from p to D does not intersect the boundary of P. We can again create a shorter path by joining the entry and exit points with a line segment, this time cutting off the portion of S in $A(D)$ and the portion of S in $B_{\bar{x}}(D)$; but S is a geodesic so this is contradiction. It follows that $y \in B_{\bar{x}}(D)$, and D is a separating dent for x and y. ∎

Lemma 2.3 *Two points x and y in a polygon P are $\mathcal{O}$-visible if and only if there is no separating dent for x and y.*

Proof. We prove the contrapositives.

(If) Suppose x does not see y. Let S be the geodesic from x to y. From Lemma 2.2 there must be some supporting dent D for S that separates x from y.

(Only If) Suppose there exists a separating dent D for x and y. Any path in P from x to y must cross $\vec{D}$ both upward and downward; from Lemma 2.1, there is no $\mathcal{O}$-convex path from x to y. From the definition of $\mathcal{O}$-visibility, x does not see y. ∎

Corollary 2.4 *Two points x and y are $\mathcal{O}$-visible if and only if the geodesic between them is $\mathcal{O}$-convex.*

Proof.

(If) Suppose the geodesic from x to y is $\mathcal{O}$-convex. The point x sees the point y from the definition of $\mathcal{O}$-visibility.

(Only If) Suppose x sees y but the geodesic S between them is not $\mathcal{O}$-convex. By Observation 1.2 there is some $\mathcal{O}$-line tangent to S. From the proof of Lemma 2.2 it follows that there is a separating dent for x and y. By Lemma 2.3 x does not see y. This is a contradiction, so S must be $\mathcal{O}$-convex. ∎

Corollary 2.5 *If a point x sees points y and z, x must be above any separating dent for y and z.*

We have now arrived at a generalization of Tietze's convexity theorem to $\mathcal{O}$-convex sets. A similar characterization appears as Theorem 5.7 in [38].

Corollary 2.6 *A polygon P is $\mathcal{O}$-convex if and only if there are no dents in the boundary of P.*

Lemma 2.7 *Dent cells are $\mathcal{O}$-convex.*

Proof. Let K be a dent cell. We argue that there cannot be a dent in the boundary of K. Let v be a vertex in the boundary of K. Let e and e' be the two edges of the boundary of K adjacent to v.

1. Suppose e and e' are both segments of dent chords; then v cannot be reflex. Since v is not reflex, there cannot be a dent at v or at either edge containing v.

2. Suppose e is a segment of a dent chord $\delta = (\gamma, \theta)$ and e' is a segment of the polygon boundary. Again v cannot be reflex, since γ is a chord of P. Since v is not reflex, there cannot be a dent at v or at either edge containing v.

3. Suppose e and e' are both segments of the polygon boundary. There cannot be a dent at v or at either edge containing v since there would have to be some dent chord tangent to v (or tangent to the edge containing v), but there are no dent chords in the interior of a dent cell.

Since there are no dents in the boundary, from Corollary 2.6 K is $\mathcal{O}$-convex. ∎

Lemma 2.8 *Two points in a polygon are in the same dent cell if and only if they see the same set of points.*

Proof. We prove the contrapositives.

(If) Suppose x and y are in the same dent cell but do not see the same set of points. From Lemma 2.7 x sees y. Let z be a "distinguishing point" visible to one of these two points, but not the other. Without loss of generality, suppose y sees z and x does not see z. Let D be some separating dent for x and z. The point y sees both x and

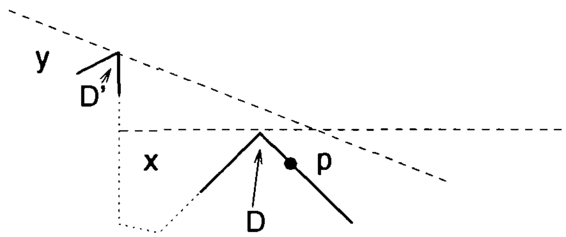

Figure 2.10: $p \in B_r(D')$

z, so from Corollary 2.5, y must be $\mathcal{O}$-above D. Since $x \prec D \prec y$, the points x and y cannot be in the same dent cell.

(Only If) Let x and y be two points in a polygon not in the same dent cell. We show that there is some some distinguishing point for x and y.

Suppose x does not see y; then trivially either x or y is a distinguishing point.

Suppose x sees y. From the fact x and y are in different dent cells, along with the assumption that x sees y, there must be some dent D such either $x \prec D \prec y$ or $y \prec D \prec x$. Let

$$\mathcal{D} = \{ D_i \mid (x \prec D_i \prec y) \vee (y \prec D_i \prec x) \}.$$

Without loss of generality, we assume that

- D maximizes, over $\mathcal{D}$, the area of the side of $B(D)$ containing neither x nor y,
- That $x \in B_l(D)$ and $y \in A(D)$, and that
- The orientation of D is 90°.

Let e be the polygon edge of the boundary of $B_r(D)$ incident on D. Let p be a point of e, arbitrarily close to $\tau(D)$. We now argue that p must be a distinguishing point for x and y. By Lemma 2.3 x does not see p. Suppose y does not see p; by Lemma 2.3, we know that there exists some dent D' such that $y \in B(D')$ and $p \in B_{\bar{y}}(D')$. Since p is arbitrarily close to $\tau(D)$, $\vec{D}'$ cannot intersect e between p and $\tau(D)$; furthermore $\tau(D')$ cannot be between p and $\tau(D)$. Since y is above $\vec{D}$ and $\vec{D}$ does not intersect the interior of e, $\tau(D')$ must also be above $\vec{D}$. We now consider two cases.

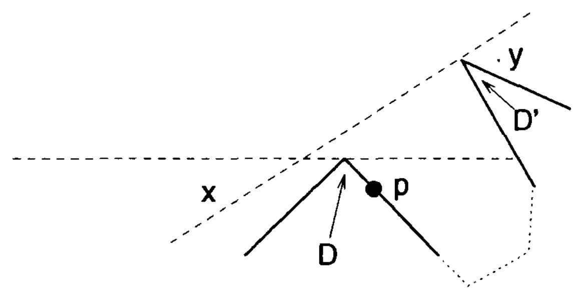

Figure 2.11: $p \in B_l(D')$

1. Suppose that $p \in B_r(D')$ (see Figure 2.10). In order for p to be $\mathcal{O}$-below $\vec{D}'$, $\tau(D)$ must also be in $B_r(D')$. It follows that $B_x(D) \subseteq B_{\bar{y}}(D')$, hence x and y are on opposite sides of D'. This contradicts our assumption that x sees y.

2. Suppose $p \in B_l(D')$ (see Figure 2.11). In order for p to be $\mathcal{O}$-below $\vec{D}'$, $\vec{D}'$ must intersect $\vec{D}$ at or to the left of $\tau(D)$.

   (a) Suppose x is below $\vec{D}'$. Since y and $\tau(D')$ are both above $\vec{D}$, x does not see y; this is a contradiction.

   (b) Suppose x is above $\vec{D}'$; then $\vec{D}'$ meets $\vec{D}$ at some non-zero angle, so $B_{\bar{x}}(D) \subset B_{\bar{y}}(D')$; but this contradicts our assumption that D maximizes over all $D_i \in \mathcal{D}$ the area of the side of $B(D_i)$ containing neither x nor y.

Since the assumption that y does not see p leads to a contradiction, and we know that x does not see p, p is a distinguishing point for x and y. ■

We have shown that all of the points in a given dent cell see the same subset of the polygon. We use these equivalence classes to define the labeled quotient graph of the PVG. Given two dent cells K_0 and K_1, we say that K_0 sees K_1 if the points in K_0 see the points in K_1. We define the *cell visibility graph* (CVG) of a polygon as follows: the vertices of a CVG are the cells of the dent decomposition, labeled with cardinalities of those cells, and there is an edge between two vertices if and only if the corresponding dent cells see one another. If the cardinality of $\mathcal{O}$ is finite, then the CVG is a finite combinatorial representation of the point visibility graph. Each vertex of a polygon can participate in at most $|\mathcal{O}|$ dents. If a polygon P has n vertices, the dent decomposition of P is an arrangement of at most $|\mathcal{O}|n + n$ line segments, so we have the following observation:

Observation 2.9 *Let P be a polygon with n vertices. The cell visibility graph of P contains $O((|\mathcal{O}|n)^2)$ vertices.*

2.2 Calculating the CVG

In this section we give a worst case optimal algorithm for calculating the cell visibility graph when $|\mathcal{O}|$ is finite. To simplify analysis, we assume that $|\mathcal{O}|$ is a fixed constant. We describe the algorithm from the bottom up, starting with the subproblems of computing the dent decomposition of a polygon and computing the subpolygon visible from a point p, called the $\mathcal{O}$-visibility polygon of p. Throughout this section, n denotes the number of vertices in the polygon whose CVG is being computed. All of the algorithms in this section presume that the polygon has been triangulated. Chazelle [10] has shown that a polygon can be triangulated in linear time.

We first give an algorithm to compute the dent decomposition. Given a dent chord $\vec{D} = (\gamma, \theta)$, if γ' is a maximal segment of γ that intersects $\tau(D)$ only at its endpoints, we call the pair $\sigma = (\gamma', \theta)$ a *half dent chord* of $\vec{D}$. Let $\sigma = (\gamma, \theta)$ be a half dent chord. We call θ the *orientation of* σ. Analogously to oriented chords, we call the region of the polygon that σ faces into $A(\sigma)$ and the region that σ faces out of $B(\sigma)$. Let σ be a half chord of $\vec{D}$. Let $B_\sigma(D)$ (respectively $B_{\bar{\sigma}}(D)$) denote the side of D identical with (respectively distinct from) $B(\sigma)$. There are exactly two half dent chords contained in any dent chord; if σ is a half dent chord contained in $\vec{D}$, we call the other half dent chord contained in $\vec{D}$ the *twin* of σ. The dent decomposition will be represented using standard subdivision representation. Each edge in the data structure for the dent decomposition will be augmented by a pointer to the half chord that contains it. A *ray shooting* query is defined as follows: given a point $p \in P$ and a direction r, find the the first intersection of a ray from p in direction r with boundary of the polygon. Guibas et al. [22] have shown how to preprocess a polygon in $O(n)$ time to allow $O(\log n)$ time ray shooting queries. We allow any number of dent chords to intersect at a point, but perturb coincident dent chords so as to compute the dent cells defined to be between them. We assume the existence sequence of an infinite sequence $\langle \epsilon_0, \epsilon_1, \epsilon_2, \ldots \rangle$ such that ϵ_i is arbitrarily small but non-zero, and $\epsilon_i \gg \epsilon_{i+1}$.

Algorithm 1: ComputeDentDecomposition(P:polygon)

1. Find the set $\mathcal{D}$ of $\mathcal{O}$-dents in the boundary of P by checking each reflex vertex or edge.
2. Preprocess P for ray shooting.
3. $k \leftarrow 0$
4. For each $D \in \mathcal{D}$
5.   For each $(\lambda, \phi) \in \{ (\text{LEFT}, \theta(D) + 90°), (\text{RIGHT}, \theta(D) - 90°) \}$
6.   Let v be the ϕ-most vertex of $\tau(D)$.
7.   Let q the first intersection of a ray from v with orientation ϕ.
8.   If q is a vertex of P then
9.   Perturb $\tau(D)$ and q by ϵ_k in direction $\theta(D) + 180°$.
10. $k \leftarrow k + 1$
11.   End if
12.   Set the λ half dent chord of D to $(\overline{vq}, \theta(D))$
13.   End For
14. End For
15. Compute the arrangement of the half dent chords of $\mathcal{D}$

End ComputeDentDecomposition

Finding the dents can be done in linear time by walking around the boundary of the polygon. Each ray shooting query takes $O(\log n)$ and there are $O(n)$ dents, so a total of $O(n \log n)$ time is spent in step 7. Let i denote the number of intersection points of dent chords in the dent decomposition. We can build the arrangement of half dent chords in $O(n \log n + i \log n)$ time using a modification of the algorithm of Bentley and Ottmann [4] described in [3]. Let k denote the number of cells in the dent decomposition. A dent decomposition is a connected planar graph, so from Euler's formula, $k \geq 2i - 4$ (for a proof of Euler's formula see [8]). It follows that the total time to calculate the dent decomposition is in $O(n \log n + k \log n)$.

We next consider the problem of calculating the $\mathcal{O}$-visibility polygon of an arbitrary point in a polygon. For any point $p \in P$, the *shortest path tree* of p is the union of the geodesics from p to each vertex of P. Guibas et al. [22] have shown how to

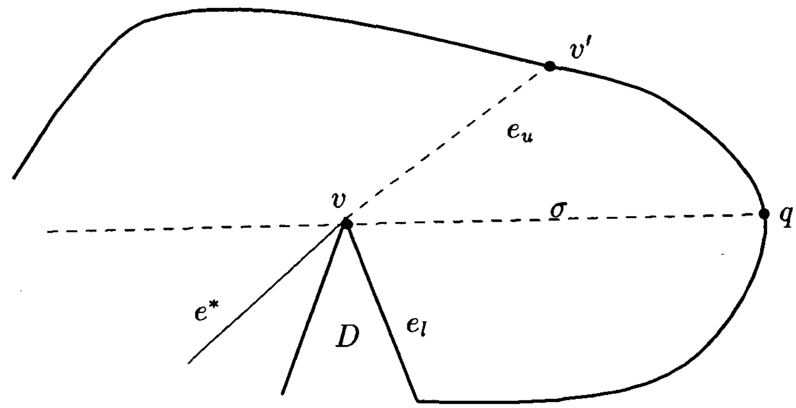

Figure 2.12: Finding the shadow chords adjacent to a given vertex.

compute the shortest path tree of a point in a triangulated polygon in linear time. We use SPT(p) to denote the shortest path tree of a point p. We are actually interested in the geodesics between p and every cell in the dent decomposition, but since this would take $\Omega(k)$ space to store and we want to use it in an $O(n)$ algorithm , we make the following observation.

Observation 2.10 *If S is the geodesic from p to q, then all but the last edge of S (i.e. the edge of S containing q) are edges of* SPT(p).

Let the *parent edge* of a vertex v in an embedded tree denote the edge between the parent of v and v. Let $e = (p, c)$ be an edge of an embedded tree where p is the parent of c. The parent edge of e is the parent edge of p. The forward extension of e denotes the maximal line segment collinear with e that does not intersect the exterior of P and that intersects e only at c. A half dent chord γ is called a *shadow chord* for a point p if $B(\gamma)$ is not $\mathcal{O}$-visible from p.

Observation 2.11 *Let σ be a half dent chord of $\vec{D}$. The half dent chord σ is a shadow chord for p if and only if $p \in B_{\bar{\sigma}}(D)$.*

Each half chord σ in the dent decomposition is associated with a flag ShadowFlag[σ] that marks whether or not σ is currently considered a shadow chord. We assume in the next algorithm that the dent decomposition of the polygon has been computed in a preprocessing step.

Algorithm 2: MarkShadowChords(p:point)

1. Compute SPT(p)
2. For each vertex v of P
3. Let e^* be the parent edge of v in SPT(p). Let $e_u = (v, v')$ be the forward extension of e. Let e_l be the edge of P adjacent to v that forms an acute angle with e_u (see Figure 2.12).
4. Let H be the set of half dent chords in the angle between e_l and e_u.
5. If $H \neq \emptyset$ then
6. Let $\sigma = (\overline{vq}, \phi)$ be the half dent chord in H that minimizes the angle $\angle v'vq$.
7. ShadowFlag[σ] $\leftarrow 1$
8. End if.
9. End for.

End MarkShadowChords

Recall that the shortest path tree of p can be computed in linear time. Since there is a constant number of dent chords incident on each dent we can find the shadow chords incident on a given vertex (if any) in $O(1)$ time. It follows that that Algorithm 2 terminates in $O(n)$ time after preprocessing.

Lemma 2.12 *Let Σ_p be the set of half dent chords marked as shadow chords by Algorithm 2 given a point p as a parameter. A point q is $\mathcal{O}$-visible from p if and only there is a path from p to q that does not intersect $B(\sigma)$ for any $\sigma \in \Sigma_p$.*

Proof.

(If) Suppose that p sees q. Further suppose that every path from p to q intersects $B(\sigma)$ for some $\sigma \in \Sigma_p$. Let S be a geodesic from p to q. Let $\sigma = (\gamma, \theta)$ be an element of Σ_p such that S intersects $B(\sigma)$. Let $\vec{D}$ be the dent chord containing σ. Since σ is marked as a shadow chord by Algorithm 2, there must exist some vertex τ of p such that there exists a parent edge e^* of v in SPT(p) (i.e. S is not a line segment) and σ forms an acute angle with the forward extension of e^* (see Figure 2.12). Since S intersects $B(\sigma)$, S must be in $B_\sigma(D)$ in the neighbourhood of σ (see Figure 2.13). It

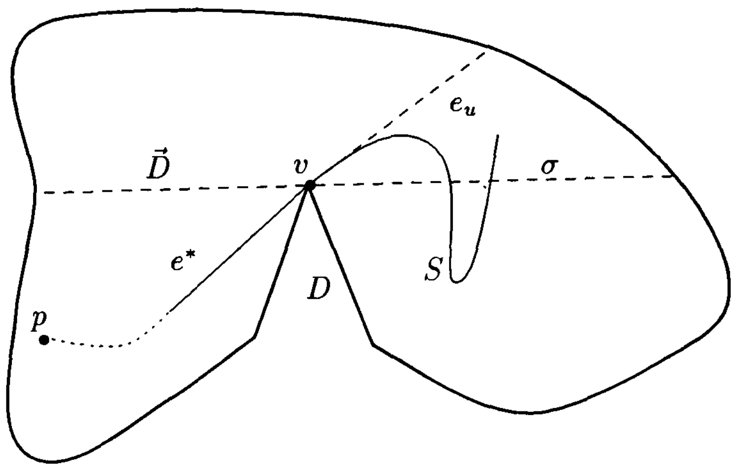

Figure 2.13: If S intersects $B(\sigma)$, then S is not $\mathcal{O}$-convex.

follows that S crosses $\vec{D}$ both upwards and downwards, hence by Lemma 2.1 is not $\mathcal{O}$-convex. By Corollary 2.4, x does not see y. This is a contradiction, so there must be some path from x to y that does not intersect $B(\sigma)$ for any $\sigma \in \Sigma_p$.

(Only If) Suppose p does not see q. Let S be a geodesic from p to q. By Lemma 2.2 there must be some dent $D = (\tau, \theta)$ that separates p from q and supports S. Since τ contains a vertex of P, there must be some edge of SPT(p) incident on τ. Let e^* be the parent edge of τ in SPT(p). Since D is a supporting dent of S, $\vec{D}$ must be tangent to S at τ. Let e_l be the edge of S after τ on a traversal from p to q. Let e_u be the forward extension of e^* (see Figure 2.14). Since S is tangent to $\vec{D}$ at τ, there must be a half chord σ of $\vec{D}$ between e_l and e_u. It follows that either σ was marked as a shadow chord or some half chord incident on τ and between σ and e_u was marked as a shadow chord. Since q is in $B_{\bar{p}}(D)$, it follows that q is in $B(\sigma)$. ∎

We could retrieve the explicit $O(n)$ sized representation of the visibility polygon in $O(n)$ time by starting at a vertex of P and walking around the boundary, taking short cuts across any shadow chords encountered. Since we are interested not in the visibility polygon *per se* but rather in the cells of the dent decomposition (i.e. vertices of the CVG) visible from a given cell, we instead present a modified depth first search subroutine that finds all of the cells reachable by a path from a given cell that

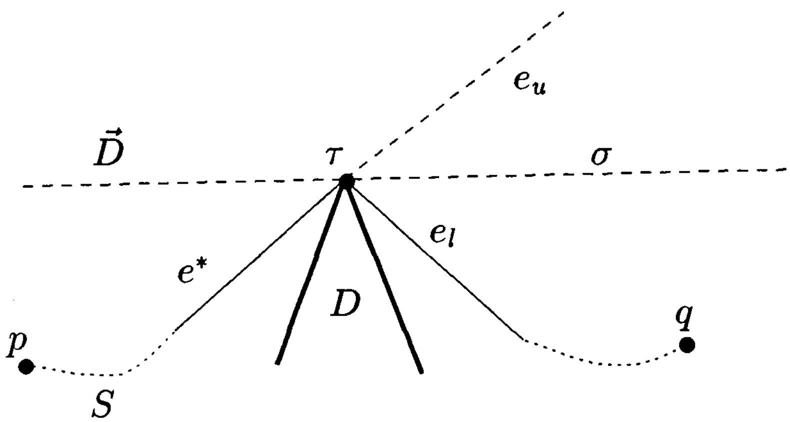

Figure 2.14: If $p \not\leadsto q$, then q is below some shadow chord of p.

does not intersect $B(\sigma)$ for any σ marked as a shadow chord in the dent decomposition. We assume that each cell K in the dent decomposition has an associated flag ReachedFlag$[K]$, initially cleared, that marks whether or not a cell has been reached by the following subroutine. We call an edge of a dent cell that is not an edge of P an *interior edge*.

```
Subroutine : ReachableFrom(K: cell)
    R ← ∅
    For each interior edge e of K
        Let σ be the half dent chord containing e.
        If ¬ShadowFlag[σ] then
            Let K_e be the cell that shares e with K.
            If ¬ReachedFlag[K_e] then
                ReachedFlag[K_e] ← 1
                R ← R ∪ ReachableFrom(K_e)
            End If
        End If
    End For
    Return R
End ReachableFrom
```

The time taken by Subroutine ReachableFrom is proportional to the total number of edges of the cells returned. Since the dual graph of the dent decomposition is a

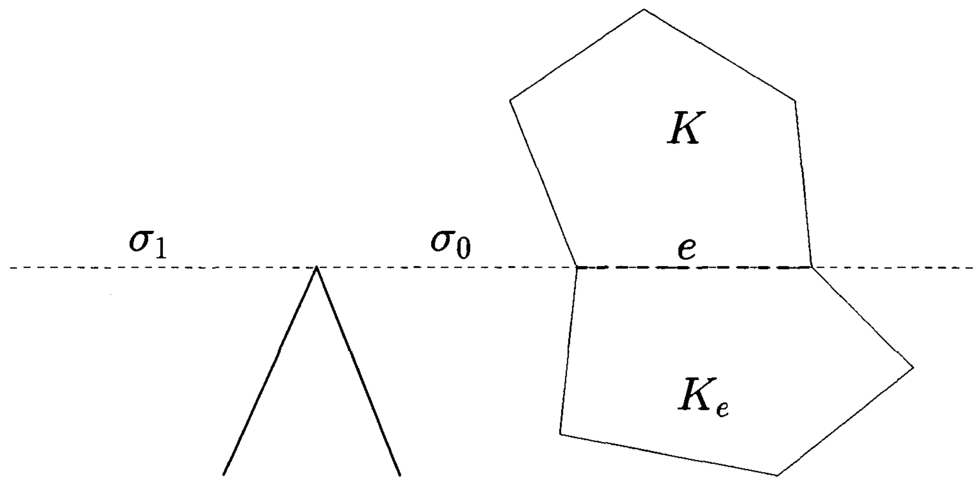

Figure 2.15: Crossing a shadow chord downwards.

planar graph, the number of edges is linear in the number of cells returned. It follows that the total time taken is in $O(|R|)$. By Lemma 2.12, if the half dent chords in the dent decomposition have been marked by Algorithm 2 with some $p_k \in K$ as a parameter, then the set of cells returned by ReachableFrom(K) is precisely the set of cells $\mathcal{O}$-visible from K.

We now present the algorithm to compute the cell visibility graph, another modified depth first search. The algorithm is similar to the Hershberger's optimal algorithm for computing the vertex visibility graph of a polygon [25]. It first calculates the $\mathcal{O}$-visibility polygon from a single cell, then incrementally modifies it in constant time to become the $\mathcal{O}$-visibility polygon of a neighbouring cell. Every cell is visible to itself, so we do not compute or store these edges of the CVG. Algorithm 3 also assumes that the dent decomposition has been computed as a preprocessing step. We associate with each cell K of the dent decomposition a flag EnumFlag[K] that marks whether or not K has been visited by the top level depth first search.

Algorithm 3: FindCVG($P, \mathcal{O}$)

1. $V \leftarrow \emptyset$; $E \leftarrow \emptyset$
2. For each half dent chord σ
   ShadowFlag[σ] $= 0$
3. For each cell K in the dent decomposition
   EnumFlag[K] $= 0$; ReachedFlag[K] $= 0$.
4. Choose some cell K as a starting cell.
5. let v_k be a point in K.
6. MarkShadowChords(v_k, P).
7. NextVertex(K)
8. Output (V, E).

End FindCVG

Subroutine : NextVertex(K: cell)

1. $V \leftarrow V \cup \{ K \}$
2. $V_r \leftarrow$ ReachableFrom(K)
3. For each $K_r \in V_r$
4.   ReachedFlag[K_r] $\leftarrow 0$
5. $E \leftarrow E \cup \{ \{ K, K_r \} \}$
6. End For
7. For each interior edge e of K
8.   Let K_e be the cell sharing e with K
9.   If $\neg$EnumFlag[K_e] then
10.   Let σ_0 be the half dent chord containing e.
11.   Let σ_1 be the twin half dent chord of σ_0.
12. $t \leftarrow$ ShadowFlag[σ_1]
13.   If $K \prec \sigma_0 \prec K_e$ then
                ShadowFlag[σ_1] $\leftarrow 0$
            else
                ShadowFlag[σ_1] $\leftarrow 1$ (see Figure 2.15).
            End If
14.   EnumFlag[K_e] $\leftarrow 1$
15.   NextVertex(K_e)
16.   ShadowFlag[σ_1] $\leftarrow t$
17.   End If
18. End For

End NextVertex

Lemma 2.13 *At any invocation of NextVertex(K) in Algorithm 3, the half dent chords of the dent decomposition marked as shadow chords are exactly the shadow chords for K.*

Proof.

(Basis) By Lemma 2.12, at the first invocation the correct half dent chords are marked as shadow chords.

(Induction) : Suppose that the correct half dent chords were marked for the parent invocation of the current one. Let the cell explored by the parent invocation be K_p. Let the cell being explored by the current invocation of NextVertex be K_c. Since K_c and K_p share an edge e of the dent decomposition, there must be exactly one distinguishing dent D for K_p and K_c. Furthermore, the edge must e must be contained in $\vec{D}$. Let σ_p be some shadow chord for K_p.

Suppose $\sigma_p \subseteq \vec{D}$; then step 13 correctly determines whether or not σ_p is also a shadow chord for K_c.

Suppose $\sigma_p \not\subseteq \vec{D}$. Let $\vec{D}_p$ be dent chord containing σ_p. The dent D_p is not the distinguishing dent for K_p and K_c and K_p is $\mathcal{O}$-below D_p, so it follows that K_c is also below D_p. Since K_p and K_c share an edge, $K_c \subseteq B_{\bar{\sigma}}(D)$. By Observation 2.11 σ_p is also a shadow dent for K_c. ∎

The correctness of Algorithm 3 follows from Lemma 2.12, Lemma 2.13, and the fact that after a given neighbour of a cell is visited, step 16 restores the shadow chords to again be correct for the current cell.

We now consider the time complexity of Algorithm 3. Recall that k denotes the number of cells in the dent decomposition of P. Preprocessing time is dominated by the $O(n \log n + k \log n)$ time it takes to build the dent decomposition. NextVertex is a modified depth first search on a planar graph (the dual graph of the dent decomposition is searched implicitly), so the time taken in NextVertex exclusive of steps 2 through 6 is in $O(k)$. Let j denote the number of edges in the CVG. Since each cell returned by ReachableFrom yields an edge of the CVG and each edge is yielded only twice, it follows that a total of $O(j)$ time is spent in steps 2 through 6 of NextVertex. Since the time for initialization in Algorithm 3 is in $O(n)$, it follows that total time complexity of Algorithm 3 is in

$$O(j + k \log n + n \log n).$$

We know the following about the relationship between j, k and n:

$$\begin{array}{ccccc} 1 & \leq & k & \leq & c_1 n^2 \\ 0 & \leq & j & \leq & c_2 k^2 \end{array}$$

A CVG can have $\Omega(n^2)$ vertices and $\Omega(n^4)$ edges (see Chapter 4), so Algorithm 3 is worst case optimal. If $j \in O(k \log n)$ then the time complexity of Algorithm 3 is dominated by the time to compute the dent decomposition. The algorithm of Bentley and Ottmann used to calculate the dent decomposition is not optimal. There is an optimal algorithm for intersecting line segments due to Chazelle and Edelsbrunner [11], but this algorithm requires that the line segments be in general position, a condition that dent decompositions do not satisfy. On the other hand, it should be possible to take advantage of the fact that the line segments whose arrangement we are computing are all chords of a polygon and have a bounded number of orientations.

2.3 Fixed Cover Numbers

In this section we consider some straightforward applications of the CVG. In particular we show that if $\mathcal{O}$ is finite, $\mathcal{O}$-convex cover and $\mathcal{O}$-star cover are both solvable in polynomial time if the number of covering polygons is fixed. We consider first the problem of $\mathcal{O}$-convex cover.

We call an $\mathcal{O}$-convex subpolygon Q of a polygon P maximal if there there is no $\mathcal{O}$-convex Q' such that $Q \subset Q' \subseteq P$. A polygon P can be covered with k maximal $\mathcal{O}$-convex subpolygons if and only if it can be covered with k $\mathcal{O}$-convex subpolygons. It follows that the complexity of the problem of maximal $\mathcal{O}$-convex cover provides an upper bound on the complexity of $\mathcal{O}$-convex cover.

A vertex or edge τ of a polygon P is said to be an $\mathcal{O}$-extremity if there exists some $\mathcal{O}$-line $\mathcal{L}$ tangent to τ and $\mathcal{L}$ is not in the interior of P in the neighbourhood of τ (see Figure 2.16). We consider $\mathcal{O}$-extremities to be oriented in the direction of a ray from τ perpendicular to $\mathcal{L}$ out of the the polygon.

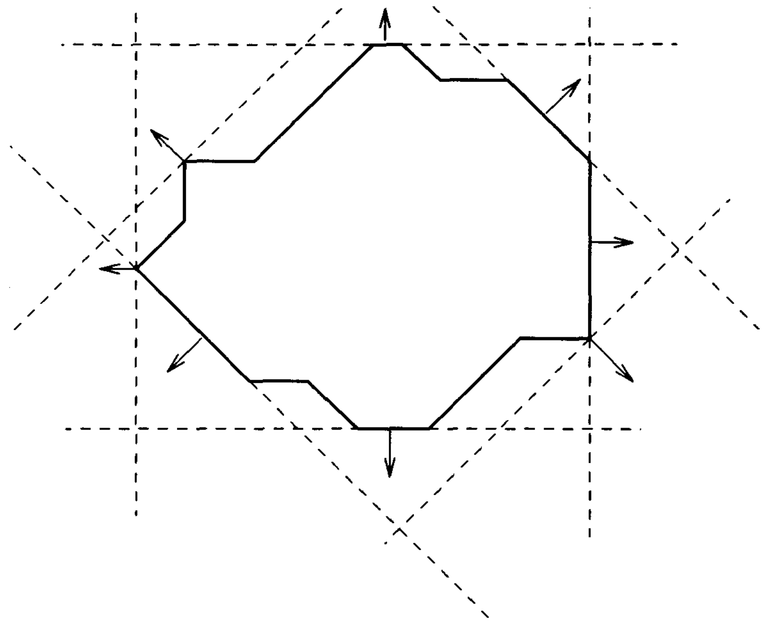

Figure 2.16: The $\{0°, 45°, 90°, 135°\}$-extremities of a polygon

Lemma 2.14 *An $\mathcal{O}$-convex polygon P has exactly $2|\mathcal{O}|$ $\mathcal{O}$-extremities.*

Proof. A polygon must have at least one $\mathcal{O}$-extremity of each possible orientation. Each orientation in $\mathcal{O}$ can have two $\mathcal{O}$-extremity orientations perpendicular to it. Suppose an $\mathcal{O}$-convex polygon has two distinct $\mathcal{O}$-extremities τ and τ' of the same orientation θ. Let δ and δ' be oriented chords of P of orientation θ arbitrarily close (inside P) to τ and τ' respectively. δ and δ', along with the polygon boundary, enclose two distinct subpolygons. Any path from τ to τ' must cross δ upwards and δ' downwards. Hence from Lemma 2.1 τ does not see τ', but this contradicts our assumption that P is $\mathcal{O}$-convex. ■

Each $\mathcal{O}$-extremity of a maximal $\mathcal{O}$-convex subpolygon Q of a polygon P must fall on an edge of P. We call these edges of P the *bounding edges* of Q.

Lemma 2.15 *A maximal $\mathcal{O}$-convex subpolygon of a polygon P is uniquely specified by its bounding edges.*

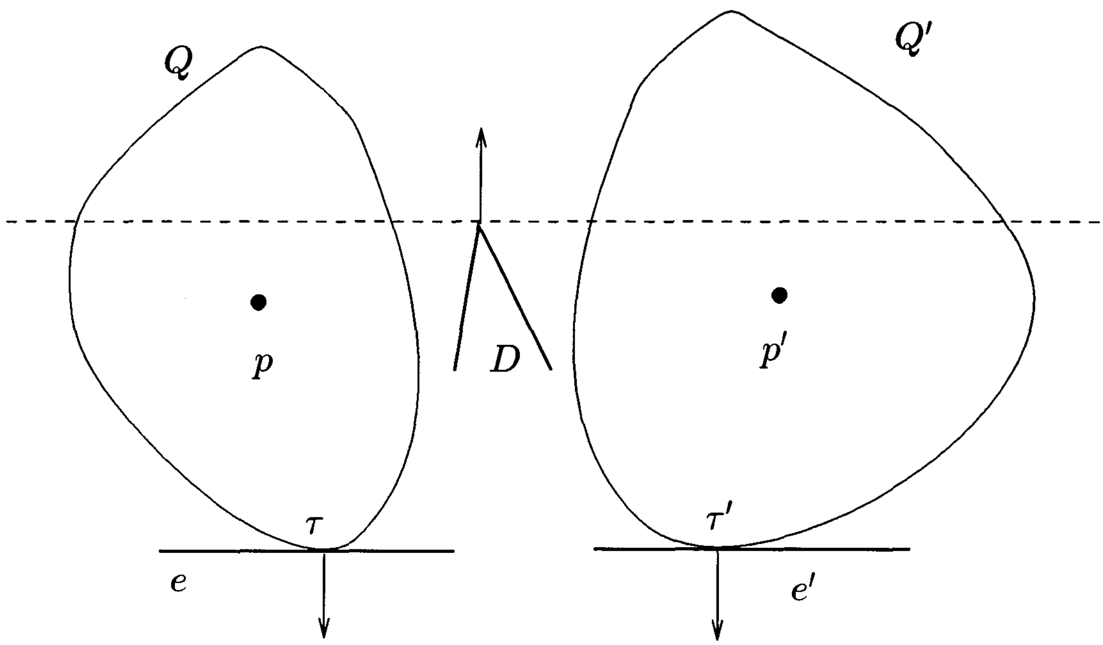

Figure 2.17: Illustration of the proof of Lemma 2.15

Proof. Let Q and Q' be maximal $\mathcal{O}$-convex subpolygons of P with the same bounding edges. Let p' be some point in Q' but not in Q. Since p' is not in Q, there must be some point $p \in Q$ such that p does not see p'. By Lemma 2.3, there must be some separating dent D for p and p'. Since there is a point of Q in $B_p(D)$ (the point p), there must also be some $\mathcal{O}$-extremity τ of Q with orientation $\theta(D) + 180°$ as D in $B_p(D)$. Similarly, there must be some $\mathcal{O}$-extremity τ' of Q' with orientation $\theta(D) + 180°$ in $B_{\bar{p}}(D)$. The $\mathcal{O}$-extremity τ must fall on some bounding edge e. The $\mathcal{O}$-extremity τ' must fall on some bounding edge e' (see Figure 2.17). No edge of P can be in both $B_p(D)$ and $B_{\bar{p}}(D)$, so $e \neq e'$. This contradicts our assumption that Q and Q' have the same bounding edges; it follows that there cannot be a point p' in $Q' \setminus Q$. ■

Lemma 2.16 *Let n be the number of vertices in a polygon P. There are at most $n^{2|\mathcal{O}|}$ maximal cliques in the CVG of P.*

Proof. Let G be the CVG of a polygon P. Each maximal clique in G corresponds to a maximal $\mathcal{O}$-convex subpolygon of P. From Lemma 2.15 and Lemma 2.14, each maximal $\mathcal{O}$-convex polygon is uniquely specified by choosing $2|\mathcal{O}|$ bounding edges. There are at most n choices for each bounding edge. ■

We now show that computing $\mathcal{O}$-convex cover is polynomial for any fixed covering number. Let n be the number of vertices in a polygon P. Let k be the number of $\mathcal{O}$-convex subsets to be used to cover P.

From Lemma 2.16 there are

$$\binom{n^{2|\mathcal{O}|}}{k} \in O(n^{2|\mathcal{O}|k})$$

distinct subsets of k maximal cliques in the CVG. Since from Observation 2.9 each of the cliques has $O((|\mathcal{O}|n)^2)$ vertices and $O((|\mathcal{O}|n)^4)$ edges, and since we can check whether a set of subgraphs covers a graph in linear time in the total size of the subgraphs and graph to be covered, it follows that each of the subsets of k cliques can be checked in time

$$O(k(|\mathcal{O}|n)^4).$$

Given the CVG of P, we can answer the question "Can P be covered by k or fewer $\mathcal{O}$-convex subpolygons?" in time

$$\begin{aligned} T(n) &\in O(k(|\mathcal{O}|n)^4 n^{2|\mathcal{O}|k}) \\ &\in O(|\mathcal{O}|^4 k n^{2|\mathcal{O}|k+4}). \end{aligned} \tag{2.1}$$

We now consider the problem of $\mathcal{O}$-star cover. Let n again be the number of vertices in a polygon P, and let k be the number of maximal stars to be used to cover P. Since each vertex of a graph defines some some maximal star, and there are $O((|\mathcal{O}|n)^2)$ vertices in the CVG, there are $O((|\mathcal{O}|n)^2)$ maximal $\mathcal{O}$-stars in P. It follows that there are $O((|\mathcal{O}|n)^{2k})$ subsets of k maximal stars in the the CVG. Since each maximal star has size $O((|\mathcal{O}|n)^4)$, each set of k maximal stars can be checked in time $O(k(|\mathcal{O}|n)^4)$. Given the CVG of P, we can answer the question "Can P be covered by k or fewer $\mathcal{O}$-starshaped subpolygons?" in time

$$\begin{aligned} T(n) &\in O(k(|\mathcal{O}|n)^4(|\mathcal{O}|n)^{2k}) \\ &\in O((|\mathcal{O}|n)^{2k+4}). \end{aligned}$$

This is a better bound than (2.1) since the time is not exponential in the number of orientations.

Neither of the algorithms presented in this section is polynomial for unbounded k, but they are in contrast with the case of non-finite $\mathcal{O}$, where it is not known whether or not convex cover is solvable in polynomial time for $k = 4$ or whether or not star cover is solvable in polynomial time for $k = 2$ (see [44] for a summary of results). In the next chapter we exhibit a class of visibility instances for which $\mathcal{O}$-convex cover is solvable in time polynomial in the input size and independent of k.

Chapter 3

Weakly Triangulated PVGs

Many problems that are NP-Hard on general graphs become tractable on suitable classes of graphs. Perhaps the most well-known such class is the class of perfect graphs.

In this chapter we show that if the set of dent orientations is restricted sufficiently, the resulting PVGs are weakly triangulated. Since CVGs are induced subgraphs of PVGs and the property of being weakly triangulated is a hereditary property of graphs, this will show that the corresponding CVGs are perfect. This will provide a duality between $\mathcal{O}$-hidden set and $\mathcal{O}$-convex cover on a restricted class of visibility instances, and a polynomial algorithm for both problems.

Culberson and Reckhow [13] have shown that orthogonal visibility CVGs of orthogonal polygons with up to 2 dent orientations are comparability graphs and Motwani, Raghunathan, and Saran [30] extended this work to show that orthogonal visibility CVGs of orthogonal polygons with up to 3 dent orientations are weakly triangulated. Following the terminology of [13, 30], let a *class k visibility instance* be one with at most k dent orientations. Our set of class 3 visibility instances includes the class 3 polygons of the previous authors as a special case. We show that not all class 3 visibility instances have weakly triangulated PVGs, but that class 3 visibility instances whose dent orientations span 180° have weakly triangulated point visibility graphs. This is more general than the previous results in two ways:

1. It places no restrictions on the orientation of polygon edges, emphasizing the

importance of dents in the visibility structure of a polygon.

2. The three dent orientations have any value, as long as their span is at least 180°.

If two dents D and D' have the same orientation, we write $D \parallel D'$; conversely if they have different orientations we write $D \not\parallel D'$. Note that this is stronger than the usual usage of the symbol "$\|$": if the orientations of D and D' differ by exactly 180° then $\vec{D}$ is parallel to $\vec{D}'$ but $D \not\parallel D'$.

Lemma 3.1 *Let D_0 and D_1 be two dents. If there exist two points x and y such that x sees y, $x \prec D_0 \prec y$, and $y \prec D_1 \prec x$, then $D_0 \not\parallel D_1$.*

Proof. Suppose $D_0 \parallel D_1$; call this orientation θ. Consider a path from x to y. This path must go from below $\vec{D}_0$ to above it, crossing a dent chord of direction θ upward, and from above $\vec{D}_1$ to below it, crossing a dent chord of direction θ downward . From Lemma 2.1, this path is not $\mathcal{O}$-convex, hence $x \not\sim y$. ■

Suppose that D is a dent, and a, b, and c are sets of points such that $a \subseteq B(D)$, $b \subseteq B_{\bar{a}}(D)$, and $c \subseteq A(D)$. We denote this situation by

$$D = \frac{c}{a \,\big|\, b}.$$

Let $\mathcal{D}$ be a collection of not necessarily distinct dents. The *incompatibility graph* G of $\mathcal{D}$ is a graph where the vertices of G are the elements of $\mathcal{D}$, and $\{D_i, D_j\}$ is an edge of G only if D_i is constrained not to be the same orientation as D_j. The chromatic number of an incompatibility graph is a lower bound on the the minimum number of distinct dent orientations present in $\mathcal{D}$, and therefore on the minimum number of distinct dents in $\mathcal{D}$.

Theorem 3.1 *A polygon whose PVG contains an antihole of size k, $k > 4$, must have at least $\lceil k/2 \rceil$ dent directions.*

Proof. Let the vertices of a k-antihole C be labeled cyclically $v_0 \ldots v_{k-1}$. In the following, we take all vertex indices modulo k. Furthermore, where $a < b$ we take the ordered set $\langle v_b \dots v_a \rangle$ to mean $\langle v_0 \dots v_{k-1} \rangle \setminus \langle v_{a+1} \dots v_{b-1} \rangle$. Let $|i-j|$ denote the "modular difference" between i and j, that is the smallest j' such that,

$$(i + j' \equiv j \bmod k) \vee (j + j' \equiv i \bmod k).$$

Two vertices v_i and v_j are $\mathcal{O}$-visible if and only if $|i-j| \neq 1$. Given two vertices, v_i and v_{i+1}, let β_i denote the set of vertices that see both v_i and v_{i+1}. Let $\bar{\beta}_i$ denote the set of vertices adjacent to either v_i or v_{i+1} in the complement graph. It follows that:

$$\begin{aligned} \bar{\beta}_i &= \langle v_{i-1}, v_i, v_{i+1}, v_{i+2} \rangle \\ \beta_i &= \langle v_0 \dots v_{k-1} \rangle \setminus \bar{\beta}_i \\ &= \langle v_{i+3} \dots v_{i-2} \rangle \end{aligned}$$

We know that there must be a separating dent D_i for each pair (v_i, v_{i+1}). Since β_i is the set of all vertices that see both v_i and v_{i+1}, from Corollary 2.5 every vertex in β_i must be $\mathcal{O}$-above this dent; i.e.:

$$D_i = \frac{\beta_i}{v_i \;\big|\; v_{i+1}} \tag{3.1}$$

We now consider the pairwise compatibility of the separating dents $\{ D_i \mid 0 \leq i < k \}$. Let i and j be vertex indices such that $|i-j| > 1$.

- Suppose $v_j \prec D_i$. It follows that

$$\begin{aligned} v_j &\in \langle v_{i+2} \dots v_{i-2} \rangle \setminus \beta_i \\ &\in \langle v_{i+2} \dots v_{i-2} \rangle \setminus \langle v_{i+3} \dots v_{i-2} \rangle \\ &= v_{i+2} \\ \beta_j &= \langle v_{j+3} \dots v_{j-2} \rangle \\ &= \langle v_{i+5} \dots v_i \rangle \end{aligned}$$

  Since there are at least 5 vertices in the antihole, $v_i \in \beta_j$ (see Figure 3.1), hence $v_i \succ D_j$. Furthermore, $j+1 = i+3$, so $v_{j+1} \in \beta_i$, and hence $v_{j+1} \succ D_i$. We

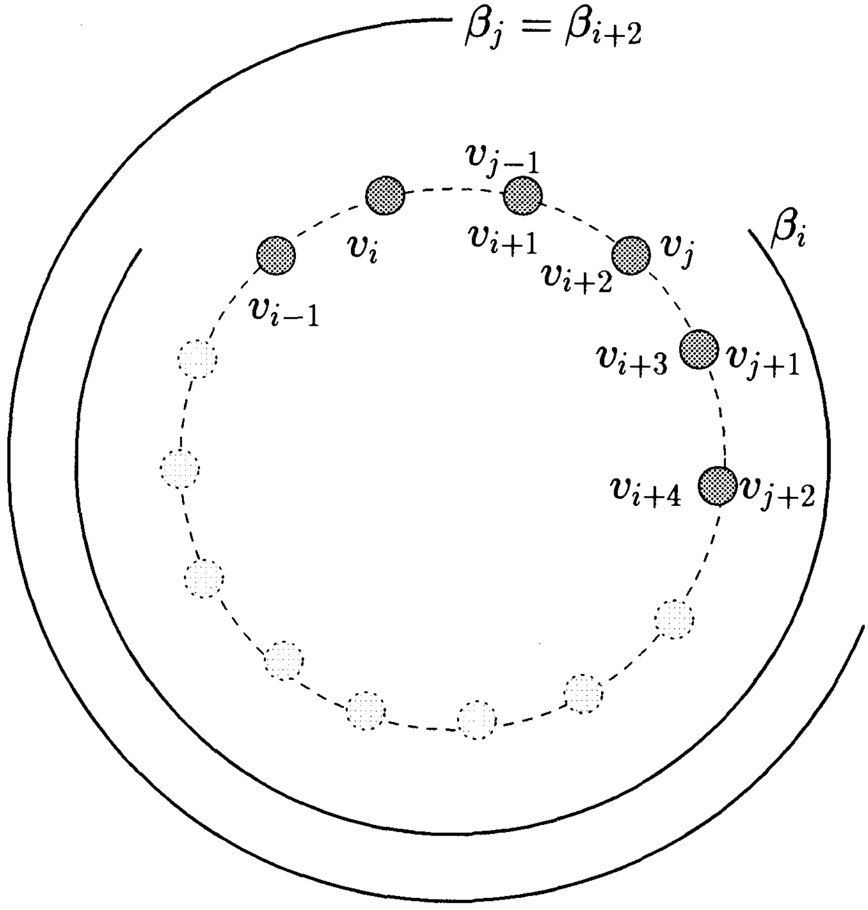

Figure 3.1: If $|i - j| = 2$, then $v_i \in \beta_j$.

know from (3.1) that $v_{j+1} \prec D_j$ and $v_i \prec D_i$, so have the following relationships:

$$\begin{aligned} v_{j+1} \prec \; & D_j \; \prec v_i \\ v_i \prec \; & D_i \; \prec v_{j+1}. \end{aligned}$$

It follows from Lemma 3.1 that $D_i \not\parallel D_j$.

- Otherwise $v_j \succ D_i$. From (3.1),

$$v_{i+1} \prec D_i \prec v_j. \tag{3.2}$$

Suppose that $v_{i+1} \prec D_j$; it follows from our assumption that $|i-j| > 1$ that

$$\begin{aligned} v_{i+1} &\in \overline{\beta}_j \setminus \{ v_j \mid 1 \geq |i-j| \} \\ &\in \langle v_{j-1}, v_j, v_{j+1}, v_{j+2} \rangle \setminus \langle v_{j-1}, v_j, v_{j+1} \rangle \\ &= v_{j+2}. \end{aligned}$$

But this contradicts our assumption that $|i-j| > 1$, so $v_{i+1} \succ D_j$, hence

$$v_j \prec D_j \prec v_{i+1}. \tag{3.3}$$

It follows from (3.2), (3.3) and Lemma 3.1 that $D_i \not\parallel D_j$.

We have established the following implication:

$$|i-j| > 1 \Rightarrow D_i \not\parallel D_j \tag{3.4}$$

From (3.4), we can see that the incompatibility graph of the set of dents generating a k-antihole contains a (not necessarily induced) k-anticycle; i.e. there exist k nodes in the complement graph with at most a cycle connecting them (see Figure 3.2). It follows that the complement graph (a k-cycle) requires at least $\lceil k/2 \rceil$ cliques to cover it; hence the chromatic number of a graph containing a k-anticycle is no less than $\lceil k/2 \rceil$. ∎

An embedding of a graph G in a polygon P such that two vertices are adjacent in G if and only if the corresponding embedded vertices are $\mathcal{O}$-visible is called an *instantiation* of G. An instantiation is called *planar* if the interiors of edge instantiations

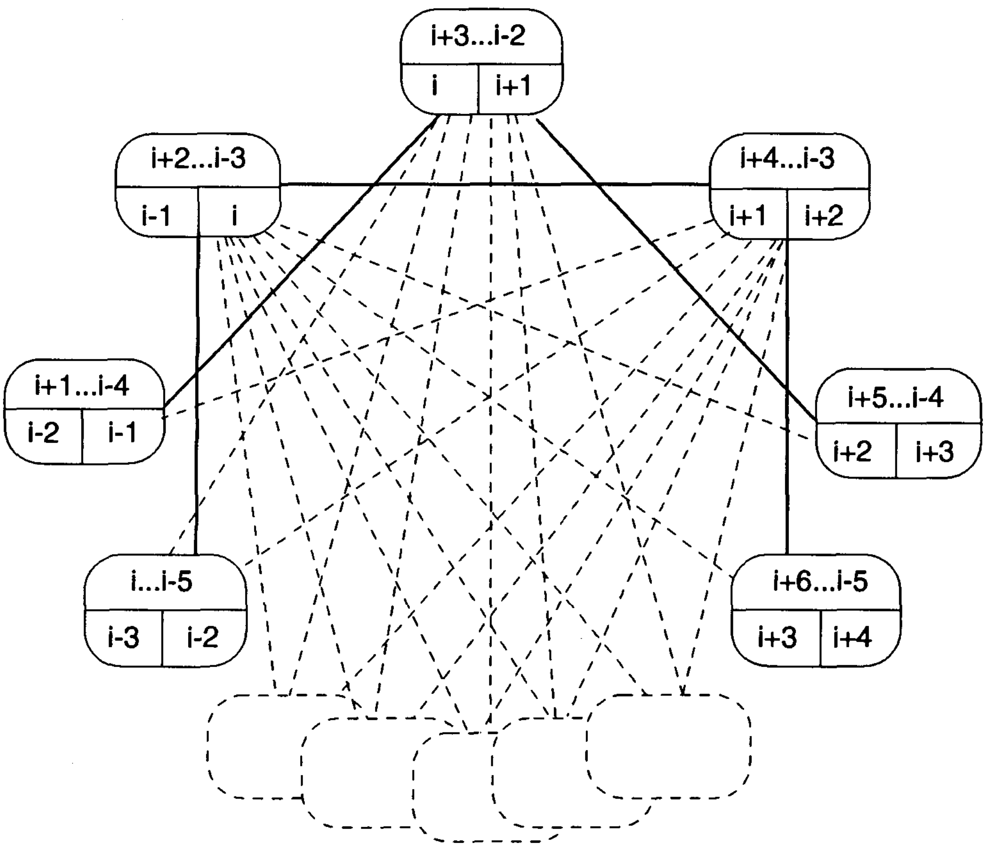

Figure 3.2: Dents incompatible with D_{i-1}, D_i, and D_{i+1} in a PVG k-antihole; only vertex indices are shown.

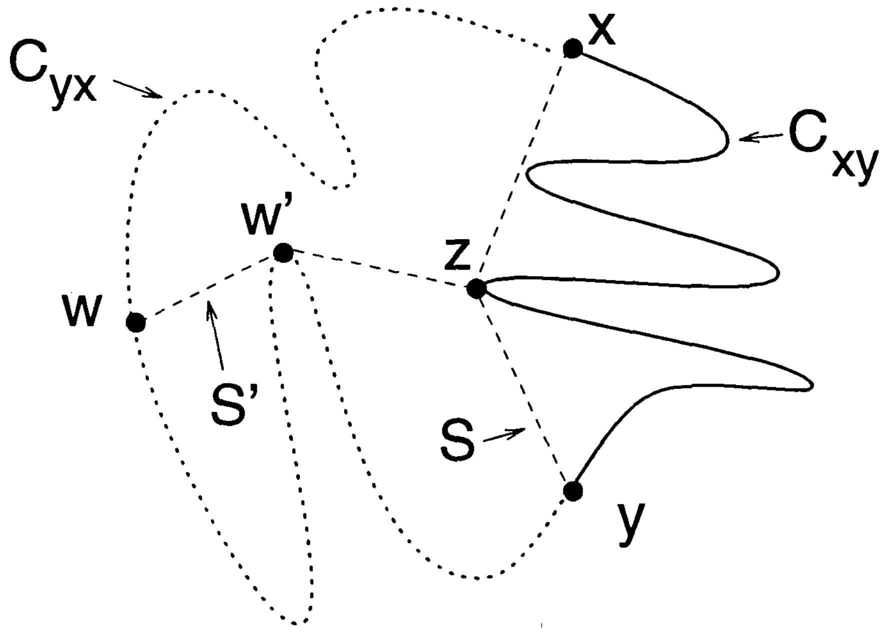

Figure 3.3: Illustration of proof of the crossing lemma.

do not intersect. In the following discussion, if x is a vertex or edge of G we write x to mean the point or staircase in the instantiation of G corresponding to x. The following is a generalization of a lemma contained in [30]:

Lemma 3.2 (The Crossing Lemma) *Let C be a hole, $|C| \geq 4$. In any instantiation of C, there exists a pair of non-adjacent edges that cross.*

Proof. Suppose there were a planar instantiation of C. A planar instantiation of a chordless cycle bounds a simple polygon. Let Q be the simple polygon bounded by a planar instantiation of C. Since the boundary of Q between vertices of C consists of staircases, dents in the boundary of Q can only occur at vertices of C.

Consider two non-adjacent vertices of C, x and y. Let the non-empty chain of C clockwise from x to y be C_{xy}. Let the non-empty chain of C clockwise from y to x be C_{yx}. Since C is chordless, $x \not\sim y$. Let S be a geodesic between x and y. Recall from the proof of Lemma 2.3 that a geodesic between two non-visible points must be supported by a dent in the polygon boundary. Let z be some vertex or edge of C

that is a supporting dent for S. Without loss of generality, assume $z \in C_{xy}$. Let w be some vertex in C_{yx}. Suppose w sees z; then the lemma holds. Otherwise, let S' be a geodesic between w and z, and let w' be the closest (along S') supporting dent of S' to z (see Figure 3.3). All of C_{xy} is on the same side of S as z, hence w' must be the instantiation of some vertex in C_{yx}. There is no dent on the subpath of S' between w' and z, hence w' must be $\mathcal{O}$-visible from z, but this is a contradiction of the chordlessness of C. It follows that there cannot be a planar instantiation of C. ∎

From Lemma 2.1, we have the following observation:

Observation 3.3 *Let x and y be the two endpoints of an $\mathcal{O}$-convex path through a third point q. Let D be a dent. Then:*

$$x \prec D \prec q \;\Rightarrow\; D \prec y$$
$$x \succ D \succ q \;\Rightarrow\; y \in B_q(D).$$

Let a *pushed chord* of a dent D through a point $p \succ D$ denote an oriented chord through p with the same orientation as $\vec{D}$.

Lemma 3.4 *Let x and y be points and D a dent such that $y \prec D \prec x$ and x is $\mathcal{O}$-visible from some point on $\tau(D)$; y is $\mathcal{O}$-below the pushed chord of D through x.*

Proof. Let S be an $\mathcal{O}$-convex path from some point on $\tau(D)$ to x. Since S is $\mathcal{O}$-convex, in particular it is monotone with respect to the orientation of D (recall that the orientation of D is perpendicular to the tangent $\mathcal{O}$-line). It follows that if δ and δ' are pushed chords of D, and δ' is farther along S from y to x, then $B(\delta) \subset B(\delta')$. ∎

Lemma 3.5 *Let D_0 be a separating dent for x and y. Let D_1 be a separating dent for x and z. If $D_0 \parallel D_1$ and $y \sim z$ then $D_0 = D_1$.*

Proof. Suppose $D_0 \neq D_1$ There are two cases.

1. Suppose $B(D_0) \subseteq A(D_1)$ and $B(D_1) \subseteq A(D_0)$. Any path from y to z must cross $\vec{D}_0$ upward and $\vec{D}_1$ downward, hence from Lemma 2.1 cannot be $\mathcal{O}$-convex and this is a contradiction.

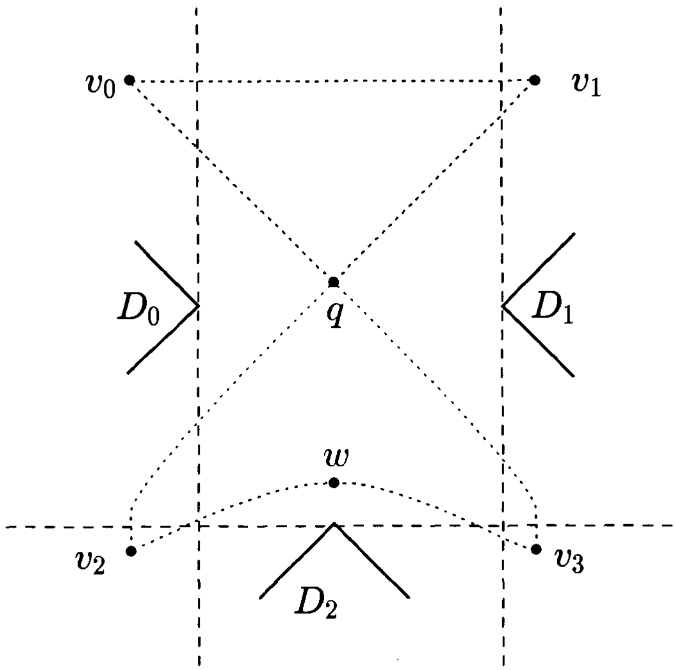

Figure 3.4: Vertex layout of a k-cycle, $k \geq 5$ in a class 3 visibility instance.

2. Suppose that $B(D_0) \subset B(D_1)$ or $B(D_1) \subset B(D_0)$, without loss of generality $B(D_0) \subset B(D_1)$. D_1 is a separating dent for y and z, and this is also a contradiction.

■

Lemma 3.6 *If the PVG of a a polygon P contains a k-hole, $k \geq 5$, then there are at least three distinct dent orientations in P.*

Proof.

Let C be a k-hole, $k \geq 5$, in the PVG of a class 3 visibility instance P. From the crossing lemma there are two non-adjacent edge staircases that cross in any instantiation of C. Let v_0v_3 and v_1v_2 be two edges whose instantiations cross; let q be the point where the two edges cross; q is visible from v_0, v_1, v_2, and v_3. Without loss of generality let the cyclic order of v_0, v_1, v_2, v_3 around C, possibly with intervening

Label	Vertex relations
D_0	$\dfrac{v_1, v_3, q}{v_0 \mid v_2}$
D_1	$\dfrac{v_0, v_2, q}{v_1 \mid v_3}$
D_2	$\dfrac{v_0, v_1, q}{v_2 \mid v_3}$

Table 3.1: Dents implied by a k-hole in a PVG.

vertices, be $\langle v_0, v_1, v_2, v_3 \rangle$. There must be at least one more vertex w in C, without loss of generality between v_2 and v_3 (see Figure 3.4). From Corollary 2.5 and 3.3, and the chordlessness of C, the dents shown in Table 3.1 must exist. From Lemma 3.1, the orientations of these three dents are pairwise incompatible, so the polygon must contain dents of at least three distinct orientations. ■

Observation 3.7 *Let S be an $\mathcal{O}$-convex path from x to y through a third point q. Let δ be a pushed chord through q whose orientation is perpendicular to some orientation in $\mathcal{O}$. The points x and y must be on opposite sides of δ.*

Lemma 3.8 *If the PVG of a class 3 visibility instance $(P, \mathcal{O})$ contains a k-hole, $k \geq 5$, then the span of the three dent orientations is strictly less than* $180°$.

Proof. Let C be a k-hole, $k \geq 5$ in the PVG of a polygon P. Let v_0v_3 and v_1v_2 be two edges whose instantiations cross; let q be the point where the two edges cross. Without loss of generality let the order of v_0, v_1, v_2, v_3 around C, possibly with intervening vertices, be $\langle v_0, v_1, v_2, v_3 \rangle$. There must be at least one more vertex in C, without loss of generality between v_2 and v_3. Let the chain of vertices from v_2 to v_3 be labeled $\langle w_1 \dots w_{k-4} \rangle$. From the proof of Lemma 3.6, the separating dents in Table 3.1 must exist and have distinct orientations.

Without loss of generality, suppose C is a minimum counterexample to the lemma; that is, suppose that k is the smallest integer greater than or equal to 5 for which a

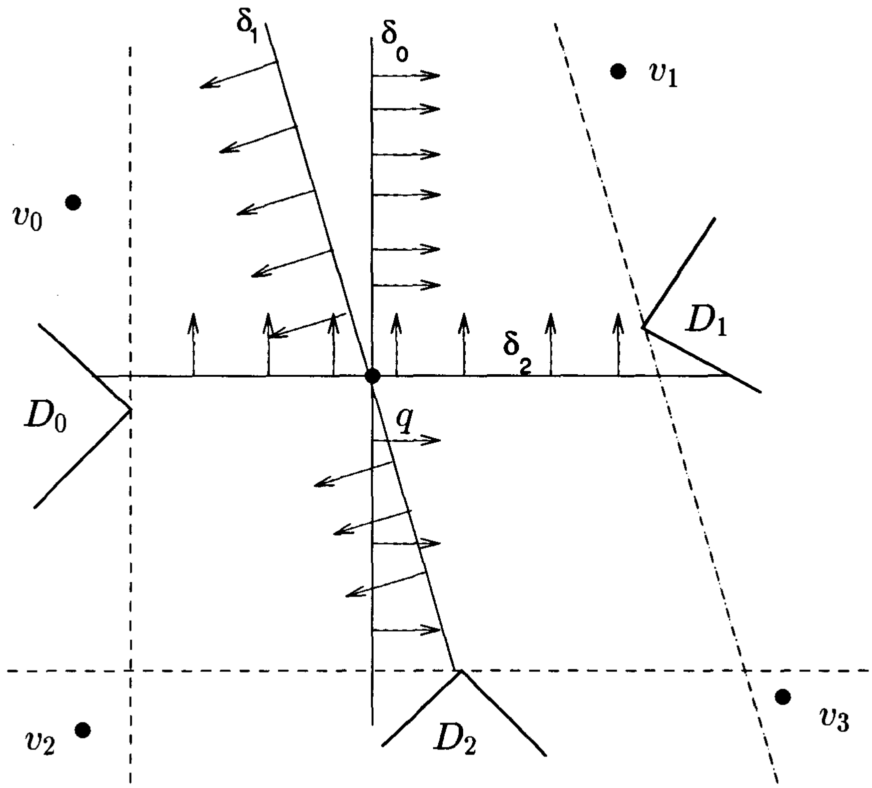

Figure 3.5: Pushed chords of D_0, D_1, and D_2.

k-hole is generated by a class 3 visibility instance with dent orientation span greater than or equal to 180°.

Suppose $v_0 \not\sim v_1$; then the following separating dent would exist:

$$D_z = \frac{v_2, v_3, q}{v_0 \mid v_1}$$

From Lemma 3.1, D_z cannot be the same orientation as any one of D_0, D_1, and D_2; but this is a contradiction, so $v_0 \sim v_1$.

Without loss of generality, suppose that the orientation of D_2 is 90°. By Observation 1.1, we may assume the orientation of D_0 is either 0° or 180°. By symmetry, suppose it is 0°. Since the span of dent orientations is at least 180°, the orientation of D_1 must be $180° + \phi$ where $0° \leq \phi \leq 90°$. Since the case of $\phi = 90$ is equivalent by

	relation to		
name	δ_0	δ_1	δ_2
κ_0	B	A	A
κ_1	A	B	A
κ_2	B	A	B
κ_3	A	B	B
κ_4	B	B	A
κ_5	A	A	B
κ_6	B	B	B
κ_7	A	A	A

Table 3.2: Equivalence classes for the relation "$\prec$" and the chords δ_0, δ_1, and δ_2.

	relation to		
vertex	δ_0	δ_1	δ_2
v_0	B	A	A
v_1	A	B	A
v_2	B	A	B
v_3	A	B	B

Table 3.3: Relations between vertices of a 5-hole and pushed chords.

relabeling to the case of $\phi = 0$, we may assume that $0° \le \phi < 90°$. Let δ_i be a pushed chord through q for dent D_i (see Figure 3.5). Every point in a polygon is either above or below a given oriented chord, but not both, so any set of all possible equivalence classes for "$\prec$" with respect to a given set of oriented chords forms a partition of the polygon. We label the equivalence classes with respect to $\{\delta_0, \delta_1, \delta_2\}$ as in Table 3.2.

Both κ_6 and κ_7 are empty so the set of weakly simple subpolygons $\{\kappa_0 \ldots \kappa_5\}$ partitions P (See Figure 3.6).

If $p \prec \delta$, we say that the relationship between p and δ is "B", otherwise we say it is "A". From Lemma 3.4, the "B" relations in Table 3.3 must hold. From

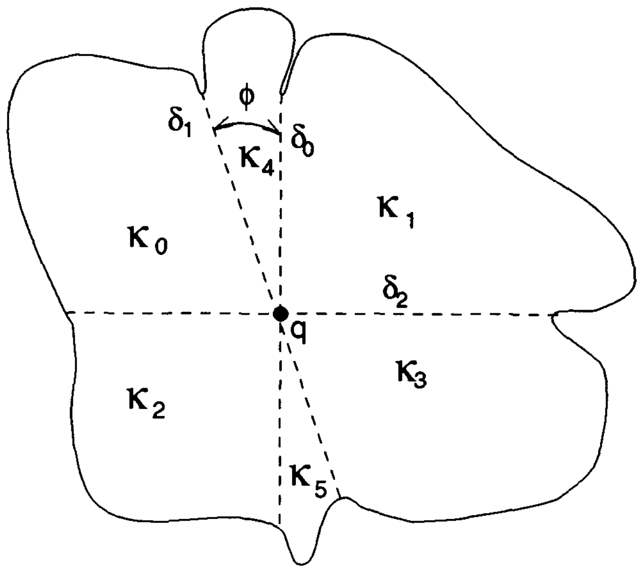

Figure 3.6: The partition of P induced by pushed chords δ_0, δ_1, and δ_2.

Observation 3.7,

$$\begin{aligned}
v_0 \prec \delta_0 &\Rightarrow v_3 \succ \delta_0 \\
v_1 \prec \delta_1 &\Rightarrow v_2 \succ \delta_1 \\
v_2 \prec \delta_0 &\Rightarrow v_1 \succ \delta_0 \\
v_2 \prec \delta_2 &\Rightarrow v_1 \succ \delta_2 \\
v_3 \prec \delta_1 &\Rightarrow v_0 \succ \delta_1 \\
v_3 \prec \delta_2 &\Rightarrow v_0 \succ \delta_2
\end{aligned}$$

hence the "A" relationships in Table 3.3 must hold. It follows that

$$i \in \{0, 1, 2, 3\} \Rightarrow v_i \in \kappa_i. \tag{3.5}$$

Let w' denote the vertex after w_1 in a cyclic traversal of C that reaches the vertices $\{v_0, v_1, v_2, v_3\}$ in the order $\langle v_0, v_1, v_2, v_3 \rangle$.

We first establish the following:

$$w_1 \sim q \vee w' \sim q. \tag{3.6}$$

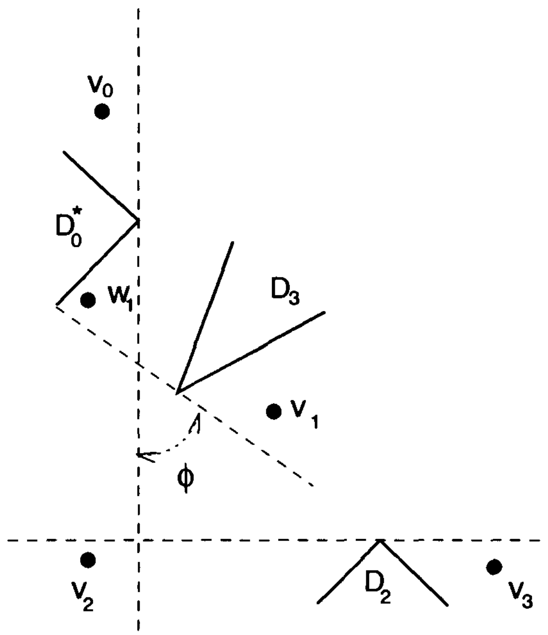

Figure 3.7: Dent layout of a k-hole when $w_1 \sim q$ and w_1 is on the top side of D_3.

To see that (3.6) must hold, we consider two possible cases.

1. Suppose $k = 5$; then $w' = v_3$ and v_3 sees q.

2. Suppose $k > 5$. Let w_j be the first of the tuple $\langle\, w_1, \ldots w_{k-4}, v_3 \,\rangle$ that sees q. If neither w_1 nor w_2 sees q then $C' = \langle\, q, v_2, w_1, w_2, \ldots w_j \,\rangle$ forms a hole of size k', $5 \leq k' < k$, generated by the same set of dent orientations, and this contradicts our assumption that C was the minimal counterexample.

We now consider the two possible cases of (3.6):

(Suppose $w_1 \sim q$) From the chordlessness of C and Observation 3.3 we know that the following dents must exist:

$$D_3 = \frac{q, v_2}{w_1 \mid v_1} \tag{3.7}$$

$$D_4 = \frac{q, v_3}{w_1 \mid v_0} \tag{3.8}$$

Recall that each of D_0, D_1, and D_2 has a unique orientation. The dent D_3 is incompatible with D_0 and D_2, so it must have the same orientation as D_1; similarly D_4 is incompatible with D_1 and D_2, hence must have the same orientation as D_0.

The pushed chords of dents with the same orientation through the same point must be identical, so $\delta_3 = \delta_1$ and $\delta_4 = \delta_0$. From Lemma 3.4, (3.7), and (3.8), w_1 must be below both δ_0 and δ_1. Since κ_4 is the only κ_i below both δ_0 and δ_1,

$$w_1 \in \kappa_4. \tag{3.9}$$

Since there is a point in κ_4 and both dent chord edges of κ_4 are open, $\phi > 0$ (see Figure 3.6).

From Lemma 3.5, Corollary 2.5, and Observation 3.3 the following dent must exist:

$$D_0^* = \frac{q, v_3, v_1}{v_0 \mid w_1, v_2}$$

1. Suppose $\vec{D}_0^*$ does not intersect $\vec{D}_3$; then $B(D_0^*)$ must be entirely $\mathcal{O}$-above or entirely $\mathcal{O}$-below and on a single side of D_3. We know that $w_1 \in B(D_0^*) \cap B(D_3)$, so $B(D_0^*)$ must be $\mathcal{O}$-below D_3 on the same side as w_1. This means that D_3 is a separating dent for v_0 and v_1, but this is a contradiction.

2. Suppose $\vec{D}_0^*$ intersects $\vec{D}_3$. We define the *top* (*bottom*) side of a dent to be the side whose boundary chord (i.e. half of $\vec{D}$ from $\tau(D)$ to the polygon boundary) projects highest (lowest) onto the y-axis. From (3.5) we know $v_2 \prec \delta_2 \prec v_0$. It follows that v_0 is in the top side of D_0^* and v_2 is in the bottom side of D_0^*. Since w_1 is on the same side of D_0^* as v_2, w_1 is on the bottom side of D_0^*. It follows that $\vec{D}_3$ must intersect $\vec{D}_0^*$ below $\tau(D_0^*)$ since otherwise w_1 would be $\mathcal{O}$-above D_3. Consider the possible positions of w_1.

   (a) Suppose w_1 is on the bottom side of D_3. From (3.9), $w_1 \succ \delta_2$, and from (3.5) $v_3 \prec \delta_2$. Since δ_2 is horizontal, it follows that v_3 must either be on the bottom side of D_3 or $\mathcal{O}$-above it. Since some point on the bottom side of D_3 (w_1) is in the bottom side of D_0^* and $0° < \phi < 90°$, $\vec{D}_3$ must be entirely contained in the bottom side of D_0^*. It follows that both the

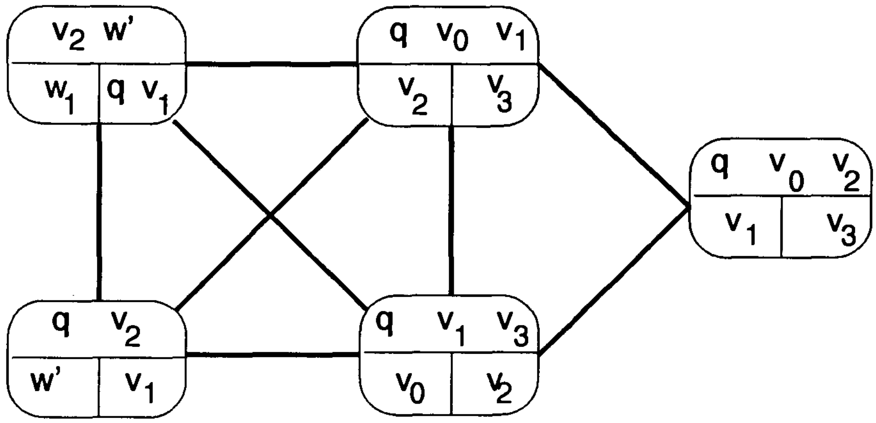

Figure 3.8: Incompatibility graph of the dents inducing a k-hole if $w_1 \not\sim q$.

bottom side of D_3 and $A(D_3)$ are contained in the bottom side of D_0^*, hence v_3 is contained in the bottom side of D_0^*. Since v_0 is in the top side of D_0^*, D_0^* is a separating dent for v_0 and v_3 but this is a contradiction because $v_0 \sim v_3$.

(b) Otherwise w_1 is on the top side of D_3 (see Figure 3.7). Since the angle ϕ between $\vec{D}_3$ and $\vec{D}_0^*$ is strictly between 0 and 90 degrees, v_0 is below D_3 on the same side as w_1. It follows that D_3 is a separating dent for v_0 and v_1 but this is a contradiction because $v_0 \sim v_1$.

We have shown that w_1 cannot see q. We now consider the other possible case of (3.6).

(Otherwise $w_1 \not\sim q \wedge w' \sim q$) Consider a separating dent D_q for w_1 and q. The points v_2 and w' see both w_1 and q, so from Corollary 2.5 must be $\mathcal{O}$-above D_q. From Observation 3.3, v_1 must be $\mathcal{O}$-below and on the same side of D_q as q.

$$D_q = \cfrac{v_2, w'}{w_1 \;\big|\; v_1, q}$$

From the chordlessness of C, there must be some separating dent D_w for w' and v_1 (see Figure 3.4). From Corollary 2.5 q must be $\mathcal{O}$-above D_w; it follows from Observation 3.3

that v_2 must also be $\mathcal{O}$-above D_w.

$$D_w = \begin{array}{c|c} \multicolumn{2}{c}{q, v_2} \\ \hline w' & v_1 \end{array}$$

Consider the incompatibility graph of D_0, D_1, D_2, D_w and D_q; as illustrated in Figure 3.8, this graph contains a clique of size 4, so no three colouring is possible, and this also is a contradiction. ∎

We call visibility instances with less than 3 dent orientations, or 3 dent orientations of span at least 180° *class 3a* visibility instances. We call visibility instances with three dent directions with span less than 180° *class 3b* visibility instances. We call a set of orientations a class 3a (3b) set of orientations if it has cardinality 3 and span at least (less than) 180°.

A 5-cycle is not perfect, hence any graph containing a 5-cycle as an induced subgraph is not perfect. The following lemma shows that Lemma 3.8 is as strong as possible.

Lemma 3.9 *For any class 3b set of orientations, there exists a polygon with dent orientations only in that set and a 5-hole in its PVG.*

Proof. Suppose we have a set of three dent orientations $\{\theta_0, \theta_1, \theta_2\}$ contained in a half-plane. Without loss of generality suppose that $\theta_0 < \theta_1 < \theta_2$ and that $\theta_0 = 0° + \phi$ and $\theta_2 = 360 - \phi$ for some $0 < \phi < 90$. By appropriate horizontal scaling of Figure 3.9 we can create a polygon with same PVG, but with the two non vertical dent orientations of Figure 3.9 being θ_0 and θ_2. Furthermore, since θ_1 is between θ_0 and θ_2 we can tilt the bottom reflex edge so that the the third dent orientation of our figure is θ_1, also without changing the PVG. ∎

Corollary 3.10 *Let $\mathcal{O}$ be a set of orientations. If $|\mathcal{O}| \geq 4$, then there exists a polygon with dent orientations only in that set and a 5-hole in its PVG.*

Proof. Let $\mathcal{O}$ be a set of orientations, $|\mathcal{O}| \geq 4$.

1. Suppose $|\mathcal{O}| = 4$.

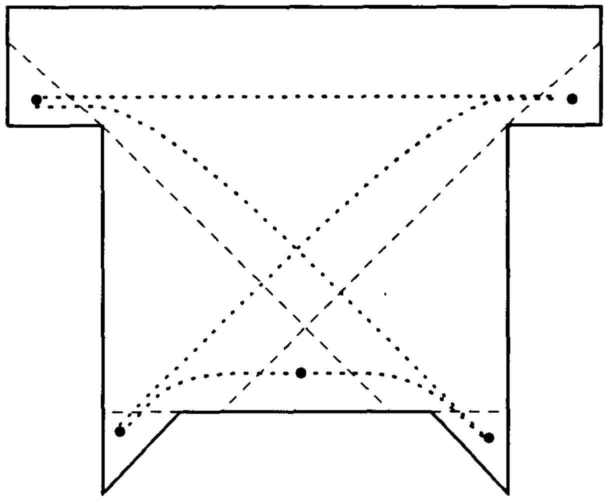

Figure 3.9: A polygon with 3 dent orientations ($\mathcal{O} = \{0°, 45°, 90°, 135°\}$) and a chordless 5-cycle in its PVG.

(a) If the orientations in $\mathcal{O}$ form two antipodal pairs, then we can apply Observation 1.1 and the construction for class 4 orthogonal polygons from [30].

(b) Otherwise, consider a line through the origin, along some orientation $\phi \in \mathcal{O}$ such that $\phi + 180° \notin \mathcal{O}$. Two of the orientations in $\mathcal{O}$ must be on one side of this line, and one of the orientations on the other. The two orientations on one side of the line, along with ϕ, form a class 3b set of orientations, and we can apply the construction of Lemma 3.9.

2. Otherwise $|\mathcal{O}| \geq 5$. Consider a line through the origin that is not along any orientation in $\mathcal{O}$. By the pigeonhole principle, one of the half planes induced by this line must contain a class 3b set of orientations so we can again apply the construction of Lemma 3.9.

■

Lemma 3.11 *A polygon with only three dent directions cannot have a 6-antihole in its PVG.*

Proof. Consider a six antihole in the PVG of a class 3 visibility instance. As in the proof of Theorem 3.1 let the separating dent for vertices v_i and v_{i+1} be D_i. From the proof of Theorem 3.1, we know that the dent incompatibility graph itself contains a six anticycle. Since we assume that the six antihole is generated by a visibility instance with only three dent orientations, it follows that the dent incompatibility graph must be 3-colourable. The only way to three colour this graph is to assign the same colour to successive pairs of vertices along the anticycle. Without loss of generality, assume that D_0 is assigned the same colour (i.e. is the same orientation as) D_1, D_2 is assigned the same colour as D_3, and so on. It follows that the separating dents for the the 6-antihole are divided into equivalence classes as follows

Orientation	Dents	
θ_1	$\dfrac{v_3, v_4}{v_0 \mid v_1}$	$\dfrac{v_4, v_5}{v_1 \mid v_2}$
θ_2	$\dfrac{v_5, v_0}{v_2 \mid v_3}$	$\dfrac{v_0, v_1}{v_3 \mid v_4}$
θ_3	$\dfrac{v_1, v_2}{v_4 \mid v_5}$	$\dfrac{v_2, v_3}{v_5 \mid v_0}$

From Lemma 3.5 we can deduce the existence of the following three dents:

$$\begin{aligned} D_1' &= \frac{v_3, v_4, v_5}{v_1 \mid v_0, v_2} \\ D_3' &= \frac{v_5, v_0, v_1}{v_3 \mid v_2, v_4} \\ D_5' &= \frac{v_1, v_2, v_3}{v_5 \mid v_4, v_0} \end{aligned} \tag{3.10}$$

Applying Observation 1.1 and symmetry, we may assume that the orientation of D_1' is 90° and the orientation of D_3' is 180°.

Since v_0 is $\mathcal{O}$-below D_1' on the the same side as v_2 and $\mathcal{O}$-above D_3', $\tau(D_1')$ must be above D_3'. Similarly since v_4 is $\mathcal{O}$-below D_3' on the same side as v_2 and $\mathcal{O}$-above D_1', $\tau(D_3')$ must be $\mathcal{O}$-above D_1'.

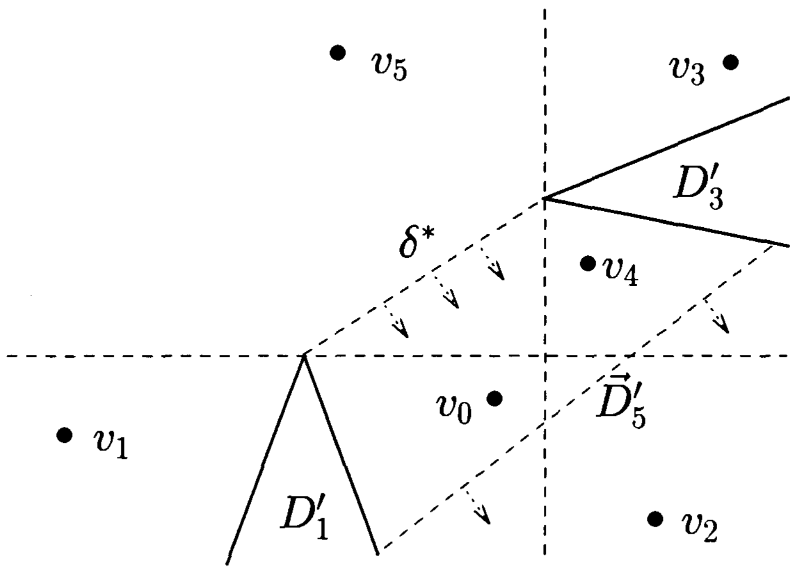

Figure 3.10: Dent layout for a 6 antihole with 3 dent orientations.

Let δ^* be a line segment from the rightmost point of $\vec{D}_1'$ to the bottommost point of $\vec{D}_3'$ (see Figure 3.10). Any point on the opposite side of δ^* from v_2 must be $\mathcal{O}$-below D_5' since $\vec{D}_5'$ separates $\{v_0, v_4\}$ from v_2, $B(\delta_*) \subset B(D_5')$. It follows that v_1 and v_3 are $\mathcal{O}$-below D_5' but this contradicts the definition of D_5'. ∎

We have now found enough forbidden subgraphs for class 3a visibility instances to prove the main result of the thesis:

Theorem 3.2 *If $(P, \mathcal{O})$ is a class 3a visibility instance then the $\mathcal{O}$-PVG of P is a weakly triangulated graph.*

Proof. Let G be the PVG of a class 3a visibility instance. From Lemma 3.8, G has no 5-antihole, or k-hole for $k > 5$. From Lemma 3.11, G has no 6-antihole. From Theorem 3.1, G has no k-antihole, $k > 6$. ∎

Theorem 3.2 provides both an interesting structural duality about class 3a visibility instances and an algorithmic result. Motwani et al. [30] noted that since the CVGs of class 3 orthogonal polygons orthogonal polygons are perfect, the maximum hidden set in a class 3 orthogonal polygon is the same size as the minimum convex cover. This result generalizes to the following:

Corollary 3.12 *The maximum $\mathcal{O}$-hidden set of a class 3a visibility instance is the same size as the minimum $\mathcal{O}$-convex cover.*

Proof. Let $(P, \mathcal{O})$ be a class 3a visibility instance. By Theorem 3.2, the $\mathcal{O}$-PVG G_p of P is weakly triangulated. Let G_c be the $\mathcal{O}$-CVG of P. Since G_c is an induced subgraph of G_p, G_c is also weakly triangulated. Finite weakly triangulated graphs are perfect, so $\alpha(G_c) = k(G_c)$. From the fact that G_c is the quotient graph of G_p and the fact that PVGs are reflexive, it follows that

$$\begin{aligned} \alpha(G_p) &= \alpha(G_c) \\ &= k(G_c) \\ &= k(G_p) \end{aligned}$$

■

Raghunathan [36] has given an algorithm that finds a maximum clique and a minimum colouring of a weakly triangulated graph $G = (V, E)$ in $O(ev^2)$ time where $e = |E|$ and $v = |V|$.

Corollary 3.13 *Let $(P, \mathcal{O})$ be a class 3a visibility instance with P having n vertices. $\mathcal{O}$-convex cover and $\mathcal{O}$-hidden set can be computed on P $O(n^8)$ time.*

Proof. We can compute the $\mathcal{O}$-CVG of P in $O(n^4)$ time using Algorithm 3. The complement of a weakly triangulated graph is also weakly triangulated, so we can use Raghunathan's algorithm on the graph theoretic complement of the CVG to find a maximum independent set and a minimum clique cover of the CVG. If a polygon P has n edges, then the CVG of P has $O(n^2)$ vertices and $O(n^4)$ edges. It follows that in $O(n^8)$ time we can find a maximum independent set and minimum clique cover in the CVG. We can map an independent set in the CVG back into a hidden set in O(n) time by simply choosing one point from each cell. We can merge the sets of cells corresponding to cliques in the CVG into $\mathcal{O}$-convex polygons in time proportional to the number of cells by marking each cell of the dent dent decomposition as to which $\mathcal{O}$-convex polygon is belongs to, and then doing a depth first walk of the dent decomposition from a cell in each clique to find the boundary of the covering subpolygons. ■

Chapter 4

Source Cells

From Corollary 3.10, we know that only a relatively small class of visibility instances have weakly triangulated PVGs. In this chapter, we examine other kinds of structure present in PVGs, in particular structure in the subgraph of the CVG induced by the "source cells" of the dent decomposition, that is the cells of the dent decomposition that are not $\mathcal{O}$-above any dent. First we show that this subgraph, called the source visibility graph, can be used to give a lower bound on the size of the largest independent set in an $\mathcal{O}$-PVG. Second, we generalize a result of Culberson and Reckhow that shows that to cover the CVG with cliques, it suffices to cover the source visibility graph with maximal cliques. Finally, we consider the question of whether or not the use of source visibility graphs can be used to enlarge the class of visibility instances for which $\mathcal{O}$-convex cover is tractable, or to improve the bounds already known.

Following Culberson and Reckhow [13], we define the *cell DAG* of a dent decomposition as follows: the nodes of the cell DAG are the cells of the dent decomposition, and there is an arc from K_1 to K_2 if they share a common edge and K_1 is $\mathcal{O}$-below the dent chord between them (see Figure 4.1). If there is an arc from K_1 to K_2 in the cell DAG, we write $K_1 \rightarrow K_2$. A *source vertex* is a vertex in the cell DAG with in degree 0. We refer to the cell of the dent decomposition corresponding to a source vertex as a *source cell*.

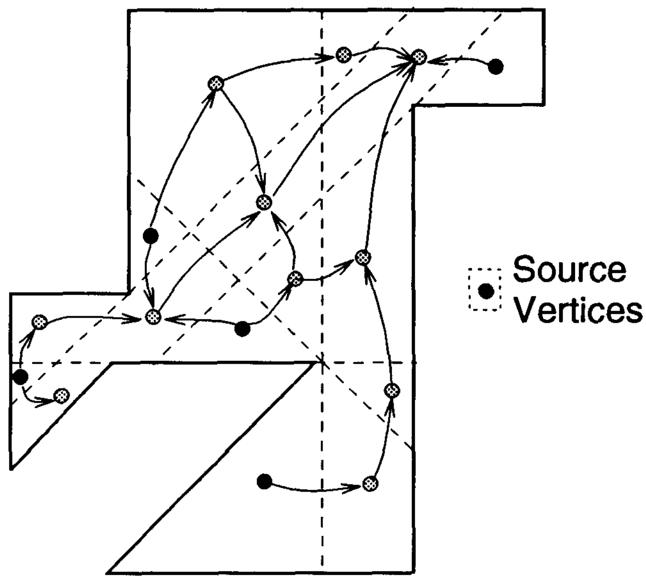

Figure 4.1: A cell DAG

Lemma 4.1 *If the boundary of a simple polygon P contains d dents of the same orientation then P contains a hidden set of size at least $d+1$.*

Proof. Let $\mathcal{D}$ be a subset of dents with orientation ϕ in the boundary of a polygon P. Let Δ be the dent decomposition induced by $\mathcal{D}$. Let $\mathcal{K}$ be the set of source cells of Δ. We first show by induction on the size of $\mathcal{D}$ that

$$|\mathcal{K}| = |\mathcal{D}| + 1. \tag{4.1}$$

(Basis) Suppose $|\mathcal{D}| = 0$; then $\mathcal{K}$ has exactly one element, the cell corresponding to the whole polygon.

(Induction) Suppose that (4.1) holds for $|\mathcal{D}| < n$. Suppose $|\mathcal{D}| = n$. Let D^* be a maximal element of $\mathcal{D}$. Let

$$\begin{aligned} \mathcal{D}_l &= \{\, D_i \in \mathcal{D} \mid \tau(D_i) \in B_l(D^*)\,\} \\ \mathcal{D}_r &= \{\, D_i \in \mathcal{D} \mid \tau(D_i) \in B_r(D^*)\,\}. \end{aligned}$$

Since D^* is maximal, and every element of $\mathcal{D}$ has the same orientation:

$$|\mathcal{D}_l| + |\mathcal{D}_r| = n - 1. \tag{4.2}$$

Consider $B_l(D^*)$ and $B_r(D^*)$ as weakly simple polygons. By the inductive hypothesis there are $|\mathcal{D}_l| + 1$ source cells in $B_l(D^*)$ and $|\mathcal{D}_r| + 1$ source cells in $B_r(D^*)$. When we add another dent chord that does not intersect the interior of any of the dent cells in $B(D^*)$, all of the source cells in $B(D^*)$ are also source cells in the new polygon. Since there are no source cells in $A(D^*)$ there are

$$(|\mathcal{D}_l| + 1) + (|\mathcal{D}_r| + 1) = n + 1 \tag{4.3}$$

source cells in Δ.

Let K_0 and K_1 be two elements of $\mathcal{K}$. Since K_0 and K_1 are source cells in a dent decomposition with only one dent chord orientation, each one must have exactly one edge that is not a segment of a polygon edge. It follows that any path from K_0 to K_1 must cross a chord of orientation ϕ both upwards and downwards, hence from Lemma 2.1 cannot be $\mathcal{O}$-convex. It follows that we can find a hidden set of size $|\mathcal{D}| + 1$ by taking one point from each element of $\mathcal{K}$. ∎

Theorem 4.1 *Let $G = (V, E)$ be an $\mathcal{O}$-CVG of a polygon P for some finite cardinality $\mathcal{O}$. Let p be the number of edges in P.*

$$\alpha(G) \geq \left\lceil \frac{-1 + \sqrt{8|V| - 7} - 2p}{4|\mathcal{O}|} \right\rceil + 1.$$

Proof. Edelsbrunner [15] has shown that an arrangement of n lines in the plane has

$$k = \frac{n^2 + n + 2}{2}$$

regions. Inverting, we get that if an arrangement of n lines has k regions,

$$n = \frac{-1 + \sqrt{8k - 7}}{2}.$$

Since we know that $n \geq 1$ and $k \geq 1$, we know that the discriminant is real and positive.

Now suppose we have an arrangement of p polygon edges and d dent chords having k' cells. Since extending each line segment to a line does not decrease the number of cells in the arrangement,

$$\begin{aligned} p+d &\geq \frac{-1+\sqrt{8k-7}}{2} \\ d &\geq \frac{-1+\sqrt{8k-7}}{2} - p. \end{aligned}$$

There are at most $2|\mathcal{O}|$ dent orientations, so by the pigeonhole principle, there must be

$$d' \geq \left\lceil \frac{-1+\sqrt{8k-7}-2p}{4|\mathcal{O}|} \right\rceil$$

dents having the same orientation. By Lemma 4.1, $\alpha(G) \geq d'+1$. ■

Theorem 4.2 (Shermer [42]) *Let $G=(V,E)$ be an PVG of a polygon P. Let p be the number of edges in P.*

$$\alpha(G) \leq p-2$$

The *source visibility graph* (SVG) of a polygon P is an undirected graph of the visibility relation between source cells of the dent decomposition of P. Culberson and Reckhow [13] show that clique cover of the the orthogonal visibility SVG of an orthogonal polygon P is polynomial time equivalent to clique cover of the orthogonal visibility PVG of P. We show that this result holds for general $\mathcal{O}$, although the SVG is only necessarily finite for finite $\mathcal{O}$.

Since there is a (possibly zero width) region between any two dent chords, exactly one dent "$\prec$" relationship must change in crossing a dent chord between neighbouring cells of the dent decomposition. We define the *distinguishing dent* for two neighbouring dent cells K_a and K_b, as the dent whose relationship (under the relation "$\prec$") to K_a is different from its relationship to K_b.

The following lemma is a generalization of Lemma 2.6 in [13].

Lemma 4.2 *If $K_1 \rightarrow K_2$ then every maximal $\mathcal{O}$-convex subpolygon that includes K_1 also covers K_2.*

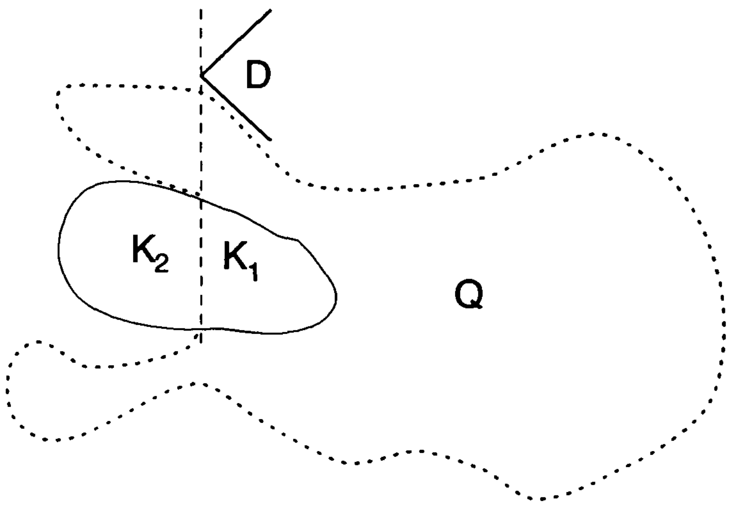

Figure 4.2: Illustration of the proof of Lemma 4.2

Proof. Let K_1 and K_2 be dent cells such that $K_1 \rightarrow K_2$. Let Q be a maximal $\mathcal{O}$-convex subpolygon that contains K_1. Suppose that $K_2 \not\subseteq Q$; since Q is maximal, $Q \cup K_2$ must be $\mathcal{O}$-concave. It follows that there must exist points x and y such that

$$(x \in K_2) \wedge (y \in Q) \wedge (x \not\sim y). \tag{4.4}$$

From Lemma 2.3, there exists some separating dent D such that

$$x \prec D \wedge y \in B_{\bar{x}}(D). \tag{4.5}$$

From (4.4), (4.5), and Lemma 2.8,

$$K_2 \subseteq B_{\bar{y}}(D). \tag{4.6}$$

There are two possible cases:

1. Suppose $K_1 \prec D$ (see Figure 4.2). Since $y \sim K_1$, $K_1 \subseteq B_y(D)$. But this, along with (4.6) contradicts the assumption that K_1 and K_2 share an edge.

2. Otherwise, $K_1 \succ D$. From (4.6) D is the distinguishing dent for K_1 and K_2. Since $K_1 \rightarrow K_2$ it follows that K_1 is below D and K_2 is above D. This contradicts (4.6). ■

Theorem 4.3 *If C is a set of maximal $\mathcal{O}$-convex polygons that includes every source cell of some polygon P, then C covers P.*

Proof. Since every node in the cell DAG is either a source vertex, or reachable by a directed path from a source vertex, this theorem follows by inductively applying Lemma 4.2. ∎

We have shown that clique cover of the $\mathcal{O}$-SVG of a polygon P is polynomial time equivalent to clique cover of the $\mathcal{O}$-PVG of P. The polynomial algorithms for class 3 orthogonal polygons presented by Culberson and Reckhow and Motwani et al. are both based on showing that the SVGs of class 3 orthogonal visibility instances are perfect. Since SVGs are induced subgraphs of PVGs, our previous results imply that the SVGs of class 3a visibility instances are perfect. Thus our results for SVGs are as strong as those for CVGs.

We now consider the question of whether the notion of source visibility graphs can be used to provide a polynomial algorithm for $\mathcal{O}$-convex cover for some non-trivial class of visibility instances beyond those in class 3a, or to improve the upper bounds in Chapter 2 and Chapter 3. Our conclusions are for the most part negative.

Each vertex in an independent set of a graph G must be in separate element of a clique cover of G, so Theorem 4.1 provides a lower bound on the $\mathcal{O}$-convex cover number of polygon that can be computed in $O(n^4)$ time if $|\mathcal{O}|$ is finite.

Since both the construction of [30] and the construction of Lemma 3.9 give induced 5-cycles in the source visibility graph, we have the following corollary of Lemma 3.9.

Corollary 4.3 *Let $\mathcal{O}$ be a set of orientations. If $|\mathcal{O}| \geq 4$, then there exists a polygon with dent orientations only in that set and a 5-hole in its SVG.*

Thus the use of SVGs does not provide a larger class of polygons for which clique cover algorithms for perfect graphs can be applied.

Culberson and Reckhow [13] argue that for orthogonal visibility class 3 polygons, the SVG will have $O(n)$ edges and $O(n^2)$ vertices while the cell visibility graph may have $\Omega(n^2)$ vertices and $\Omega(n^2)$ edges. They further note that orthogonal visibility class 4 polygons may have $\Omega(n^2)$ vertices and $\Omega(n^4)$ edges in the source visibility

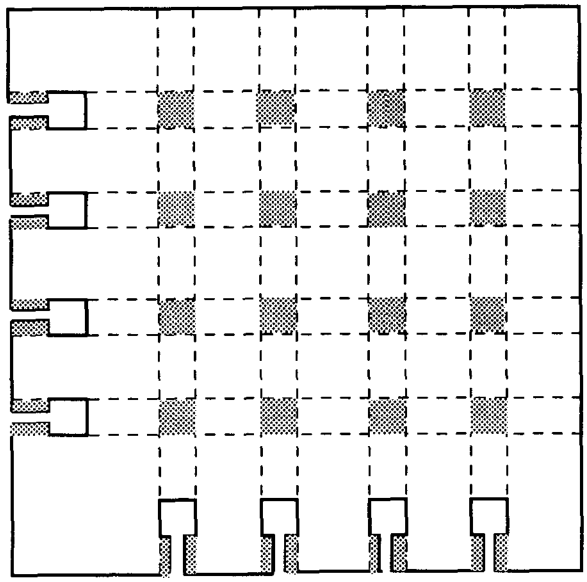

Figure 4.3: The SVGs of orthogonal visibility instances may contain $\Omega(n^2)$ vertices and $\Omega(n^4)$ edges.

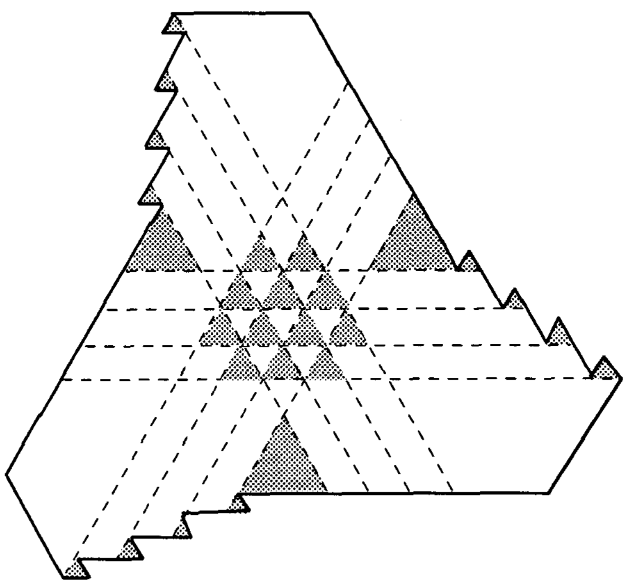

Figure 4.4: The SVGs of class 3 visibility instances may contain $\Omega(n^2)$ vertices and $\Omega(n^4)$ edges.

graph (see Figure 4.3). Given three dent orientations whose span is greater than 180° we can generalize the grid construction of Figure 4.3 to form a triangular grid of $\Omega(n^2)$ source regions with $\Omega(n^4)$ visibility edges between them (see Figure 4.4).

Chapter 5

Link 2 PVGs and Star Cover

A point x is called *link-k* $\mathcal{O}$-visible (or just link-k visible) from a point y if there is some path between x and y (not intersecting the exterior of the polygon) consisting of at most k staircases joined at their endpoints. A set of points P is called link-k $\mathcal{O}$-convex (or just link-k convex) if every pair of points in P is link-k $\mathcal{O}$-visible. Motwani et al. [31] prove that for orthogonal visibility, star cover reduces to link-2 convex cover. Motwani et al. further show that the link-2 orthogonal point visibility graph (the square of the orthogonal point visibility graph) is weakly triangulated. In this chapter we show that for $|\mathcal{O}| > 2$, neither of these results holds. We exhibit a class of visibility instances for which the equivalence of starshapedness and link-2 convexity does hold. This class of visibility instances does not contain all orthogonal visibility instances, but does have an interesting Helly like characterization.

An equivalent definition of link-2 convexity is that a set of points P is link-2 convex if any pair of points x and y in P both see some third point z in P. To see that any starshaped polygon is link-2 $\mathcal{O}$-convex, we note that any two points in a starshaped polygon see some point in the kernel. A link-2 convex polygon is not necessarily starshaped because every pair of points does not necessarily see the *same* point z. Motwani et. al showed that for $\mathcal{O} = \{ 0°, 90° \}$ any link-2 $\mathcal{O}$-convex polygon is $\mathcal{O}$-starshaped. By Observation 1.1, for $|\mathcal{O}| = 2$, any link-2 convex polygon is $\mathcal{O}$-starshaped. Figure 5.1 shows a polygon that is link-2 $\mathcal{O}$-convex for any $\mathcal{O}$ but is not $\mathcal{O}$-starshaped for the set of three orientations shown. It follows that for $|\mathcal{O}| > 2$,

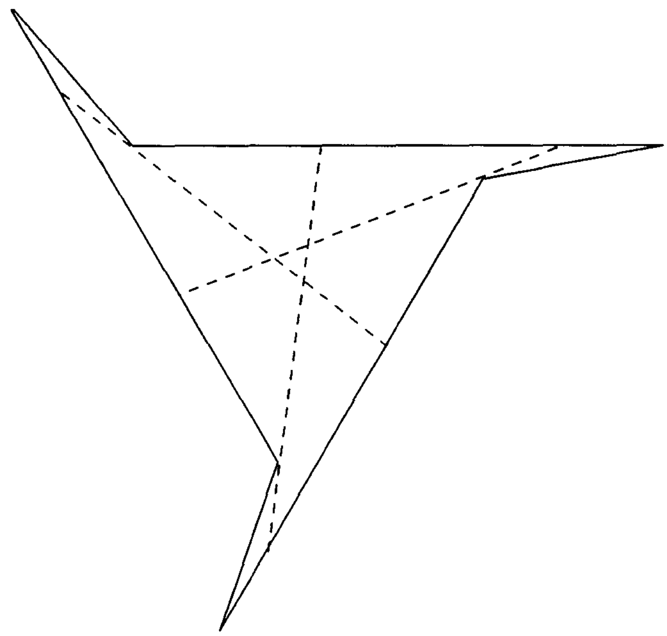

Figure 5.1: A class 3 visibility instance that is link-2 $\mathcal{O}$-convex, but not $\mathcal{O}$-starshaped.

link-2 $\mathcal{O}$-convexity is not necessarily equivalent to $\mathcal{O}$-starshapedness.

We say that a set of dents $\mathcal{D}$ *covers* a polygon P (or $\mathcal{D}$ is a *covering set* for P) if $P = \bigcup_{D \in \mathcal{D}} B(D)$. If $|\mathcal{D}| = 2$, we say that $\mathcal{D}$ is a *covering pair* for P. We say that a dent D covers a point p if $p \prec D$.

Lemma 5.1 *A polygon P is $\mathcal{O}$-starshaped if and only if it contains no covering set of dents.*

Proof.

(If) Let R be a region $\mathcal{O}$-above every dent in P. Suppose a point x in R did not see some other point y in P; then there must be a separating dent D for x and y. The point x cannot be below D, so this is a contradiction and R must be a kernel for P.

(Only If) Suppose P contains a covering set of dents. It follows that every point in P is below some dent. Let x be an arbitrary point in P. The point x must be below some dent; let D be such a dent. The point x cannot see any point in $B_{\bar{x}}(D)$, so x cannot be in a kernel for P. ∎

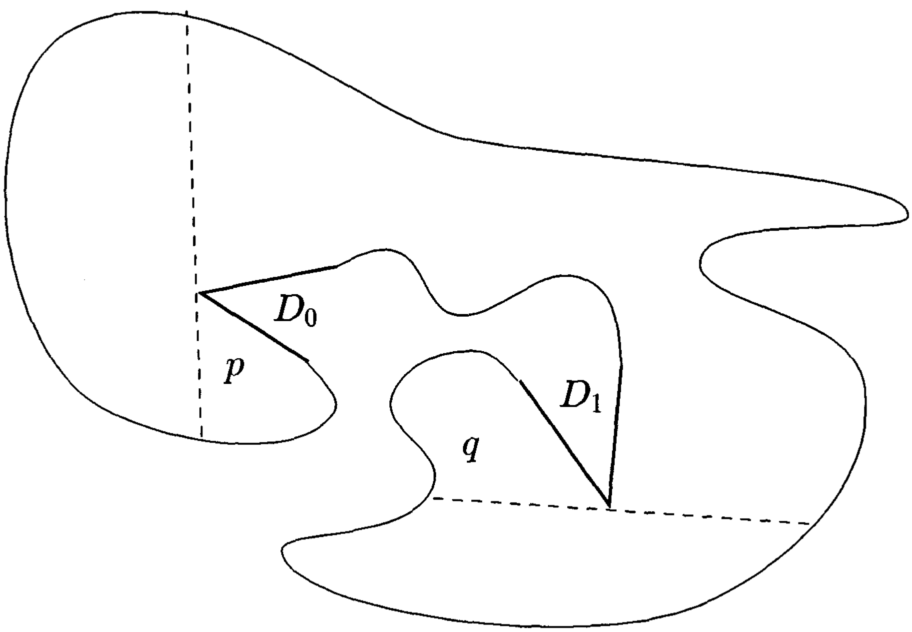

Figure 5.2: Illustration of the proof of Lemma 5.3

We define a *dent* in an open path analogously to a dent in the boundary of a polygon. A vertex of a polygonal path S is called *interior* if it is not one of the two endpoints of S. An edge is called interior if both endpoints are interior. Let S be a polygonal path. Let τ be an interior vertex or edge of S. If there exists some $\mathcal{O}$-chord $\delta = (\gamma, \theta)$ such that γ is tangent to τ, then we call the ordered pair $D = (\tau, \theta)$ a *dent* in S.

Observation 5.2 *If a path S has k dents, S is not link-k convex.*

Lemma 5.3 *If P is link-2 convex, then it contains no covering pair of dents.*

Proof. We prove the contrapositive. Let D_0 and D_1 be a covering pair of dents for P. Any point in P must be below one of D_0 or D_1. It follows that $A(D_0) \subset B(D_1)$ and $A(D_1) \subset B(D_0)$. Furthermore, for any dent D' no connected subset of P can intersect both $B_l(D')$ and $B_r(D')$ without also intersecting $A(D')$. Since $A(D_0)$ does

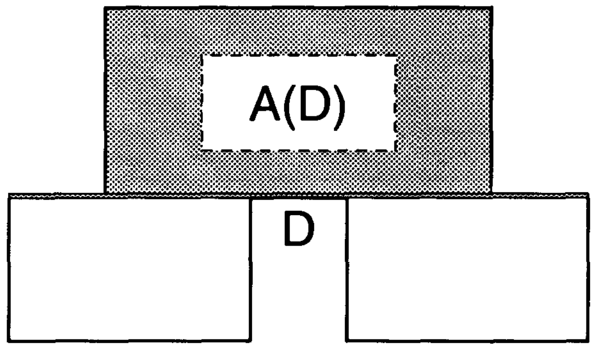

Figure 5.3: A dent D such that $A(D)$ is a hat polygon

not intersect $A(D_1)$, $A(D_0)$ must be entirely contained on one side of $B(D_1)$ and $A(D_1)$ must be entirely contained on one side of $B(D_0)$. Let p be a point $\mathcal{O}$-below D_0 on the opposite side from $A(D_1)$. Let q be a point $\mathcal{O}$-below D_1 on the opposite side from $A(D_0)$ (see Figure 5.2. Let S be a path from p to q. The point q is in $B_{\bar{p}}(D_0)$ so S must cross $\vec{D}_0$ upwards at some point u_0 and downward at some point d_0. It follows that there must be some dent t_0 on S between u_0 and d_0 such that $\tau(t_0) \in A(D_0)$. Similarly $p \in B_{\bar{q}}(D_1)$ so there must be some dent $t_1 \in A(D_1)$ on S. Since $A(D_0) \cap A(D_1) = \emptyset$, $t_0 \neq t_1$. Since S has two distinct dents, it cannot be link-2 convex. Since p is not link-2 visible from q, P is not link-2 convex. ∎

Let $\mathcal{D}$ be a set of dents. $\Theta(\mathcal{D})$ denotes the set $\{\, \theta(D_i) \mid D_i \in \mathcal{D} \,\}$ and the *span of* $\mathcal{D}$ denotes the span of $\Theta(\mathcal{D})$.

Lemma 5.4 *Let $\mathcal{D}$ be the set of dents in the boundary of a simple polygon P. If the span of $\mathcal{D}$ is at most $180°$ then for any dent $D \in \mathcal{D}$, If $(\theta(D) + 180°) \notin \Theta(\mathcal{D})$ then $A(D)$ is a hat polygon.*

Proof. Let $\mathcal{D}$ be the set of orientations in the boundary of P. Suppose the orientations in $\Theta(\mathcal{D})$ are contained in a closed half plane, without loss of generality the upper one one induced by a horizontal line. Let D be an element of $\mathcal{D}$ such that $(\theta(D) + 180°) \notin \Theta(\mathcal{D})$. It follows that $0° < \theta(D) < 180°$. Suppose $A(D)$ were not a hat polygon; since $\vec{D}$ is a chord of P, $A(D)$ is weakly simple. Since P is simple there

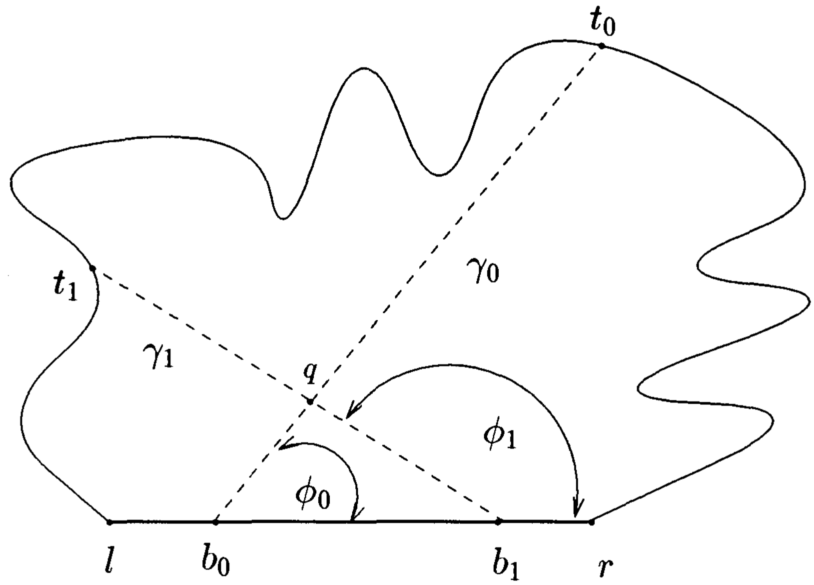

Figure 5.4: The chord γ_0 intersects the chord γ_1.

must be some point in $A(D)$ where the boundary of P is tangent to $\vec{D}$. This would imply the existence of a downward facing dent in $\mathcal{D}$, but this is a contradiction. It follows that $A(D)$ is a hat polygon (see Figure 5.3). ∎

Lemma 5.5 *Let P be a hat polygon. Let $e = (l, r)$ be an edge of P where r is after l in the counterclockwise traversal of the boundary of P. Let $\gamma_0 = (b_0, t_0)$ and $\gamma_1 = (b_1, t_1)$ be two chords of P such that*

1. *The points b_0 and b_1 are contained in the interior of e, and*

2. *Either*

   (a) *γ_0 intersects γ_1, or*

   (b) *The extensions of γ_0 and γ_1 to lines intersect in the half plane defined by e not containing γ_0 and γ_1.*

If $\angle t_0 b_0 r < \angle t_1 b_1 r$, then t_0 is encountered before t_1 on a counterclockwise walk of the boundary of P from r.

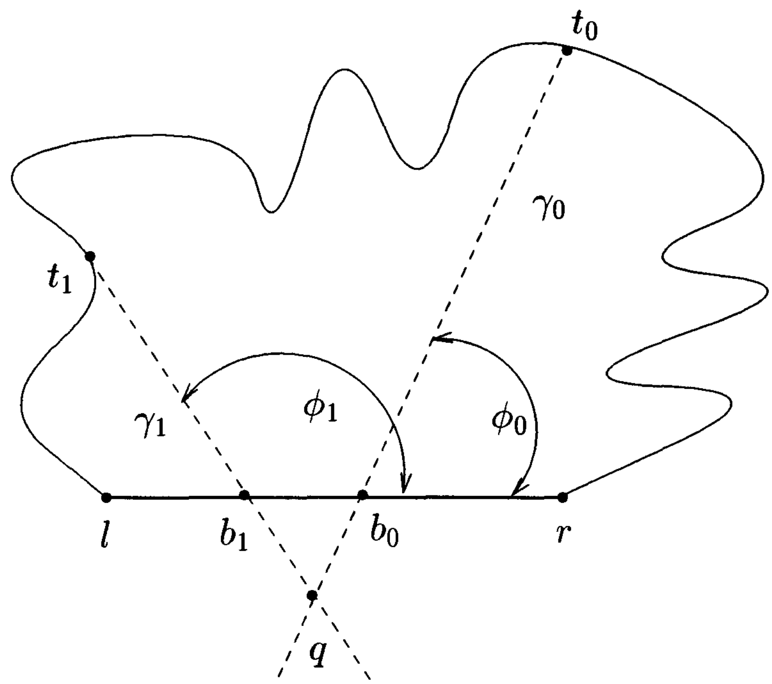

Figure 5.5: The extensions of γ_0 and γ_1 to lines intersect below e.

Proof. Without loss of generality, suppose e is horizontal and a ray upward from e is inside P in the neighbourhood of e. Let ϕ_0 (ϕ_1) denote the magnitude of the angle $\angle t_0 b_0 r$ ($\angle t_1 b_1 r$).

If b_0 or b_1 is in a brim segment of P, then γ_0 cannot cross γ_1 so it must be the case that b_0 is to the right of b_1 on e. Let e_l (e_r) denote the edge of P before (after) e on the counterclockwise traversal of the boundary of P. Suppose b_0 is in a brim segment of P; then since $b_0 = t_0$ and $b_1 = t_1$, t_1 must either be before t_0 on the clockwise traversal of e_l or on another edge. In either case t_0 is reached before t_1 on a counterclockwise walk of the boundary of P from r. The case of b_1 in the brim of P.

Otherwise if P is simple (i.e. has no brim segments) or neither b_0 or b_1 is in the brim of P, it suffices to consider the simple polygon P' consisting of the closure of the interior of P. The chord γ_1 divides P' into two weakly simple polygons: P_l on the left and P_r on the right. We show that t_0 must be in P_r. There are two cases:

1. Suppose γ_0 intersects γ_1 (see Figure 5.4). Let q be the intersection point. Since both γ_0 and γ_1 are chords of P and intersect the interior of e, $0° < \phi_0 < \phi_1 < 180°$, the line segment $(q, t_0]$ must be entirely contained in P_r.

2. Otherwise the extensions of γ_0 and γ_1 to lines intersect in the half plane below e (see Figure 5.5). Let q again be the intersection point. Since the orientation of a ray from q through t_0 is smaller than that of a ray from q through t_1, b_1 must be to the left of b_0. It follows that the the line segment $[b_0, t_0]$ is entirely contained in P_r.

Since t_0 is between r and t_1 on a walk (clockwise or counterclockwise) of the boundary of P_r from r, t_0 must be between r and t_1 on a counterclockwise walk of the boundary of P from r. ■

Given two oriented chords $\delta_0 = (\gamma_0, \phi_0)$ and $\delta_1 = (\gamma_1, \phi_1)$, we say that δ_0 *crosses* δ_1 and write $\delta_0 \bowtie \delta_1$ if the intersection of the chords γ_0 and γ_1 is a proper crossing. We use $D_0 \bowtie D_1$ as equivalent notation for $\vec{D}_0 \bowtie \vec{D}_1$. Let $l(D)$ (respectively $r(D)$) denote the endpoint of $\vec{D}$ incident on $B_l(D)$ (respectively $B_r(D)$).

Lemma 5.6 *Let D_0 and D_1 be dents such that $\{\,\theta(D_0), \theta(D_1)\,\} \subseteq [0°, 180°]$ and $D_0 \bowtie D_1$.*

$$l(D_0) \prec D_1 \quad \Leftrightarrow \quad \theta(D_1) < \theta(D_0)$$
$$r(D_0) \prec D_1 \quad \Leftrightarrow \quad \theta(D_1) > \theta(D_0)$$

Proof. We first prove the forward direction of the implications.

1. Suppose $\theta(D_0) = 0°$; any dent which crosses D_0 must have orientation distinct from, hence greater than $0°$.

2. Suppose $\theta(D_1) = 180°$; any dent which crosses D_0 must have orientation distinct from, hence less than $180°$.

3. Otherwise $0° < \theta(D_0) < 180°$, so $l(D_0)$ is to the left of $r(D_0)$. Let q be the intersection point of $\vec{D}_0$ and $\vec{D}_1$.

(a) Suppose $l(D_0) \prec D_1$. Let δ be an oriented chord through q and $l(D_0)$ with the same orientation as $\vec{D}_0$. In order to rotate δ about q so that $l(D_0)$ is $\mathcal{O}$-below δ, we must rotate in a clockwise direction, i.e. decrease the orientation of δ.

(b) Suppose $r(D_0) \prec D_1$. Let δ be an oriented chord through q and $r(D_0)$ with the same orientation as $\vec{D}_0$. In order to rotate δ about q so that $r(D_0)$ is $\mathcal{O}$-below δ, we must rotate in a counterclockwise direction, i.e. increase the orientation of δ.

We now show the the reverse directions of the implications hold. Observe that $\vec{D}_0$ and $\vec{D}_1$ are both line segments, so if $\vec{D}_0 \bowtie \vec{D}_1$ then exactly one endpoint of D_0 is on either side of $\vec{D}_1$. Suppose $\theta(D_0) < \theta(D_1)$ and $l(D_0) \succ D_1$. Since exactly one endpoint of $\vec{D}_0$ is $\mathcal{O}$-above D_1, $r(D_0) \prec D_0$. From the forward direction of the lemma, $\theta(D_0) > \theta(D_1)$, but this a contradiction, so $l(D_0) \prec D_1$. Similarly, if $\theta(D_0) > \theta(D_1)$ then $r(D_0)$ must be be $\mathcal{O}$-below D_1 ■

Lemma 5.7 *Let $\mathcal{D}$ be the set of dents in the boundary of a polygon P with the span of $\mathcal{D}$ at most* 180°. *If $\mathcal{D}$ contains no covering pair, then $\mathcal{D}$ contains no covering set of dents.*

Proof. Let $\mathcal{D}$ be the set of dents in the boundary of a polygon P. Suppose that the span of $\mathcal{D}$ is at most 180° and $\mathcal{D}$ contains no covering pair. Without loss of generality, suppose the orientations of dents in $\mathcal{D}$ are contained in the (closed) upper half plane induced by a horizontal line. Let $\mathcal{D}'$ be the set of maximal elements of $\mathcal{D}$. If there is a covering set of dents in $\mathcal{D}$ then there is a covering set in $\mathcal{D}'$. Suppose there were a covering set of dents in $\mathcal{D}'$.

We first show the following.

$$(\{ D_0, D_1 \} \subseteq \mathcal{D}') \Rightarrow (\vec{D}_0 \bowtie \vec{D}_1) \tag{5.1}$$

Let D_0 and D_1 be two elements of $\mathcal{D}'$. Since D_0 and D_1 are both maximal, it follows that

$$B(D_0) \not\subset B(D_1)$$

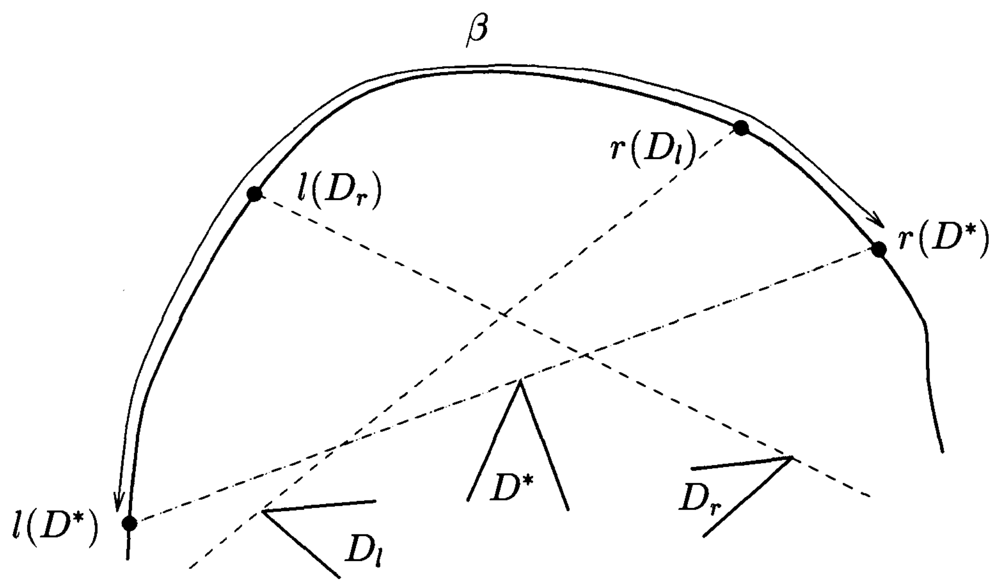

Figure 5.6: Dents covering the endpoints of a maximal dent chord.

$$B(D_1) \not\subset B(D_0).$$

Since $\{ D_0, D_1 \}$ is not a covering pair,

$$A(D_0) \cap A(D_1) \neq \emptyset.$$

It follows that $D_0 \bowtie D_1$.

Suppose the orientation of some $D \in \mathcal{D}'$ were 0° or 180°; it follows the associated dent chord $\vec{D}$ would be vertical. Let p be the highest of $\{ l(D), r(D) \}$. Let D_p be a dent that covers p. From (5.1), $D \bowtie D_p$. But this would imply that D had an orientation between 180° and 360°, which is a contradiction. Thus if there exists such a D_p, there is no covering set for D. We can now assume that

$$(D \in \mathcal{D}') \Rightarrow (0^\circ < \theta(D) < 180^\circ) \tag{5.2}$$

Let D^* be some element of $\mathcal{D}'$. A dent D is called a *right* dent if $0 \leq \theta(D) < \theta(D^*)$, and a *left* dent if $\theta(D^*) < \theta(D) \leq 180$. From (5.1) and Lemma 5.6 we know that $l(D^*)$ must be covered be a right dent and $r(D^*)$ must be covered by a left dent. Let

β be the path along the polygon boundary from $l(D^*)$ to $r(D^*)$ $\mathcal{O}$-above D^*. Let D_l be the left dent in $\mathcal{D}'$ whose chord intersects β closest to $l(D^*)$. Let D_r be the right dent in $\mathcal{D}'$ whose chord intersects β closest to $r(D^*)$ (see Figure 5.6).

From (5.1) D_l and D_r must both cross D^*. It follows from Lemma 5.6 that

$$\begin{aligned} \vec{D}_l \cap \beta &= r(D_l) \\ \vec{D}_r \cap \beta &= l(D_r) \end{aligned}$$

From Lemma 5.4, $A(D^*)$ is a hat polygon. It follows from Lemma 5.5 that $r(D_l)$ must be closer to $r(D^*)$ on β than $l(D_r)$ is. Let D_c be some dent that covers $r(D_l)$. From (5.1) $\vec{D}_c$ must intersect $\vec{D}_l$. Since D_c covers $r(D_l)$ and $\theta(D_c)$ is contained in the upper half-plane, it follows from Lemma 5.6 that $\theta(D_l) < \theta(D_c)$. Since D_l is a left dent, and

$$0° < \theta(D_l) < \theta(D_c) < 180°$$

it follows that D_c is a left dent. Since $D_c \bowtie D_l$ and $\theta(D_l) < \theta(D_c)$ it follows from Lemma 5.5 that $\vec{D}_c$ intersects β closer to $l(D^*)$ than $\vec{D}_l$ does. This contradicts our definition of D_l, so there is no covering set of dents in $\mathcal{D}'$, hence no covering set of dents in $\mathcal{D}$. ∎

We can restate the previous lemma in a manner analogous to Helly's theorem for planar convex sets.

Corollary 5.8 *Let $\mathcal{D}$ be the set of dents in the boundary of a polygon P with the span of $\mathcal{D}$ at most* 180°. *Let $\mathcal{A}$ be the set $\{ A(D) \mid D \in \mathcal{D} \}$. If every pair of elements of $\mathcal{A}$ has a point in common, then $\bigcap_{Q \in \mathcal{A}} Q \neq \emptyset$.*

Theorem 5.1 *Let $\mathcal{D}$ be the set of dents in the boundary of a polygon P with the span of $\mathcal{D}$ at most* 180°. *If P is link-2 $\mathcal{O}$-convex then P is $\mathcal{O}$-starshaped.*

Proof. Let P be a polygon whose dent orientations are contained in a half plane. Further suppose that P is link-2 $\mathcal{O}$-convex. From Lemma 5.3, P contains no covering pair of dents. From Lemma 5.7, P contains no covering set of dents. From Lemma 5.1, P is $\mathcal{O}$-starshaped. ∎

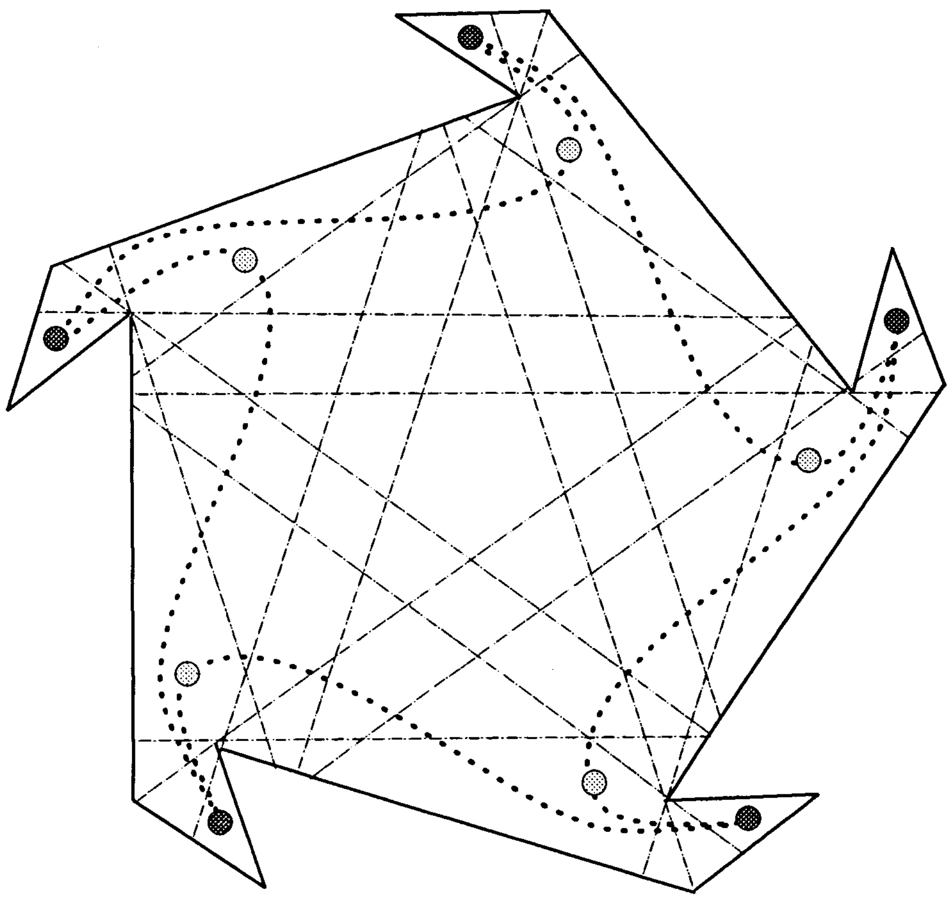

Figure 5.7: A polygon with an induced five cycle in the link-2 point visibility graph. $\mathcal{O} = \{\,0°, 36°, 72°, 108°, 144°\,\}$.

Motwani et al. [31] show that the link-2 orthogonal cell visibility graph is weakly triangulated. They use this fact, along with the fact that if a polygon is link-2 $\{0°, 90°\}$-convex, then it is $\{0°, 90°\}$-starshaped, to give a polynomial time algorithm for $\{0°, 90°\}$-star cover. We have argued above that while the reduction from star cover to link-2 convex cover does not hold in general for $|\mathcal{O}| \geq 3$, it does hold if the span of the dents in a visibility instance is at most 180°. The question of whether link-2 convex cover is tractable for these visibility instances remains open. We do know that not all link-2 $\mathcal{O}$-PVGs are perfect; Figure 5.7 shows a polygon with a chordless 5 cycle in the link-2 PVG.

Chapter 6

Conclusions

This thesis follows [13, 30] in using the notion of dent orientation to characterize which visibility instances have weakly triangulated PVGs. We show however, that for more general kinds of visibility, the number of dent orientations is not sufficient for this characterization. We introduce the notion of the span of a set of dent orientations and show that to guarantee a weakly triangulated PVG, not only must a polygon have a maximum of 3 dent orientations, but if it does have 3 dent orientations, the span of these orientations must be at least 180°.

From Corollary 3.10 we know that any class of visibility instances with more than 3 dent directions does not have perfect $\mathcal{O}$-PVGs. On the other hand, as investigated in Chapter 4, there is considerable structure in an arbitrary PVG.

Question 1 *Does there exist some nontrivial class of visibility instances distinct from the set of 3a visibility instances for which clique cover of the $\mathcal{O}$-PVGs is tractable?*

In Chapter 5 we investigate when $\mathcal{O}$-star cover is reducible to link-2 $\mathcal{O}$-convex cover. We show that if the span of the orientations of dents in the boundary of P is at most 180° then this reduction holds. Unlike the results of Chapter 3, this is not a generalization of previous results. Let $\mathcal{P}_2$ be defined to be the class of visibility instances $(P, \mathcal{O})$ such that $|\mathcal{O}| = 2$. Motwani et al. [31], along with Observation 1.1, show that for $(P, \mathcal{O}) \in \mathcal{P}_2$,

1. P is $\mathcal{O}$-starshaped if and only if P is link-2 $\mathcal{O}$-convex, and

2. The link-2 $\mathcal{O}$-PVG of P is weakly triangulated.

This leads us to ask the following two questions:

Question 2 *Is there a class $\mathcal{P}_\pi$ of visibility instances such that*

1. *$\mathcal{P}_2 \subset \mathcal{P}_\pi$, and*
2. *$\forall (P, \mathcal{O}) \in \mathcal{P}_\pi$ P is $\mathcal{O}$-starshaped if and only if P is link-2 $\mathcal{O}$-convex?*

Question 3 *Let $\mathcal{D}$ be the set of dents in the boundary of a polygon P. If the span of $\mathcal{D}$ is at most* $180°$*, is the link-2 PVG of P necessarily perfect?*

For finite cardinality $\mathcal{O}$, the distinction between perfect PVGs and perfect CVGs is practically unimportant. Rawlins and Wood show that several visibility problems are solvable not only for finite cardinality $\mathcal{O}$ but for $\mathcal{O}$ consisting of a finite set of closed ranges; in this case the CVG would not necessarily be finite. In Chapter 3 we argued that the PVGs of class 3a polygons are weakly triangulated. Since the statement "All triangulated graphs of cardinality $\aleph_1$ are perfect" is known to be independent of ZFC (Zermelo-Fraenkel set theory with the axiom of choice) [48], the perfection of uncountable weakly triangulated graphs is unlikely to be provable in general. However, it is also known that triangulated graphs that contain no infinite independent set and those that contain no infinite clique are perfect [23, 48]. Since no restricted-orientation PVG contains an infinite independent set, we ask the following:

Question 4 *If G is an uncountable weakly triangulated graph and G has no infinite independent set, is G necessarily perfect?*

Bibliography

[1] A. Aggarwal. *The Art Gallery Theorem: Its Variations, Applications and Algorithmic Aspects.* PhD thesis, Dept. of Elec. Eng. and Computer Science, Johns Hopkins University, 1984.

[2] L. J. Aupperle, H. E. Conn, J. M. Keil, and J. O'Rourke. Covering orthogonal polygons with squares. In *Proc. 26th Allerton Conf. Commun. Control Comput.*, pages 97–106, Oct. 1988.

[3] P. Belleville. Computing two-covers of simple polygons. Master's thesis, McGill University, Montréal, September 1991.

[4] J. Bentley and T. Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Trans. Comput.*, pages 643–647, 1979.

[5] C. Berge. *Farbung von Graphen sämtlich bzw. deren ungerade Kreise starr sind*, pages 114–115. Halle-Wittenberg Math-Natur, 1961.

[6] C. Berge. Sur un conjecture relative au problème des codes optimaux. In *comm. 13ème assemblée général de l'URSI Tokyo, 1962*, 1962.

[7] C. Berge. *Graphs and Hypergraphs.* North-Holland, second edition, 1973.

[8] J. Bondy and U. Murty. *Graph Theory with Applications.* North Holland, 1976.

[9] K. Chaudhary and P. Robinson. Channel routing by sorting. *IEEE Trans. Comput.-Aided Design Integrated Circuits*, 10(6), June 1991.

[10] B. Chazelle. Triangulating a Simple Polygon in Linear Time. In *Proc. 31st Symp. Foundations of Comput. Sci.*, pages 220–230, 1990.

[11] B. Chazelle and H. Edelsbrunner. An Optimal Algorithm for Intersecting Line Segments in the Plane. In *Proc. 20th Symp. Theory of Comput.*, pages 590–600, 1988.

[12] J. Culberson and R. A. Reckhow. Covering polygons is hard (preliminary abstract). In *Proc. 29th Symp. Foundations of Comput. Sci.*, pages 601–611. ACM, 1988.

[13] J. Culberson and R. A. Reckhow. Orthogonally convex coverings of orthogonal polygons without holes. *J. Comput. Syst. Sci.*, 39:166–204, 1989.

[14] R. Diestel. *Graph Decompositions: a study in infinite graph theory.* Oxford, 1990.

[15] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, 1987.

[16] H. Everett and D. Corneil. Recognizing visibility graphs of spiral polygons. *J. Algorithms*, 11:1–26, 1990.

[17] D. Gans. *Transformations and Geometries.* Appleton-Century-Crofts, 1969.

[18] M. R. Garey and D. S. Johnson. *Computers and Intractability.* W.H. Freeman and Company, 1979.

[19] S. Ghosh. On recognizing and characterizing visibility graphs of simple polygons. In R. Karlsson and A. Lingas, editors, *Lecture Notes in Computer Science 318*, pages 96–104. Springer-Verlag, New York, 1988.

[20] M. C. Golumbic. *Algorithmic Graph Theory and Perfect Graphs.* Academic Press, 1980.

[21] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences on combinatorial optimization. *Combinatorica*, 1:169–197, 1981.

[22] L. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. E. Tarjan. Linear time algorithms for visibility and shortest path problems insde simple polygons. In *Proc. 2nd Symp. on Comp. Geom.*, pages 1–12, 1986.

[23] R. Halin. A note on infinite triangulated graphs. *Discrete Mathematics*, 31:325–326, 1980.

[24] R. Hayward. Weakly triangulated graphs. *J. Comb. Theory B*, 39:200–209, 1985.

[25] J. Hershberger. An optimal visibility graph algorithm of triangulated simple polygons. *Algorithmica*, 4:141–155, 1989.

[26] J. M. Keil. Minimally covering a horizontally convex orthogonal polygon. In *Proc. 2nd Symp. on Comp. Geom.*, pages 43–51, 1986.

[27] D. Lee and A. Lin. Computational complexity of art gallery problems. *IEEE Trans. Inform. Theory*, 32:276–282, 1986.

[28] L. Lovász. A characterization of perfect graphs. *J. Comb. Theory B*, 13:95–98, 1972.

[29] D. Marple, M. Smulders, and H. Hegen. Tailor: A layout system based on trapezoidal corner stitching. *IEEE Trans. Comput.-Aided Design Integrated Circuits*, 9(1):66–90, January 1990.

[30] R. Motwani, A. Ragunathan, and H. Saran. Perfect graphs and orthogonally convex covers. *SIAM J. Disc. Math.*, 2:371–392, 1989.

[31] R. Motwani, A. Ragunathan, and H. Saran. Covering orthogonal polygons with star polygons: The perfect graph approach. *J. Comput. Syst. Sci.*, 40:19–48, 1990.

[32] J. I. Munro, M. H. Overmars, and D. Wood. Variations on visibility. In *Proc. 3rd Symp. on Comp. Geom.*, pages 291–299, 1987.

[33] C. Nash-Williams. Infinite graphs— a survey. *J. Comb. Theory*, 3:286–301, 1967.

[34] C. Nash-Williams. Unexplored and semi-explored territories in graph theory. In F. Harary, editor, *New Directions in the Theory of Graphs*, pages 149–186. Academic Press, 1973.

[35] J. O'Rourke. On the complexity of computing minimum covers for polygons. In *Proc. 20th Allerton Conf. Commun. Control Comput.*, pages 75–84, 1982.

[36] A. Raghunathan. *Polygon Decomposition and Perfect Graphs.* PhD thesis, University of California at Berkeley, 1988.

[37] G. J. Rawlins. *Explorations in Restricted Oriented Geometry.* PhD thesis, University of Waterloo, 1987.

[38] G. J. Rawlins and D. Wood. Restricted-oriented convex sets. *Inf. Sci.*, 54:263–281, 1991.

[39] F. S. Roberts. *Applied Combinatorics.* Prentice-Hall, 1984.

[40] S. Schuierer, G. J. Rawlins, and D. Wood. Towards a general theory of visibility. In J. Urrutia, editor, *Proc. 2nd Canadian Conf. on Comp. Geom.*, pages 354–357, 1990.

[41] T. C. Shermer. Convex cover is NP-hard. note, 1988.

[42] T. C. Shermer. Hiding People in Polygons. *Computing*, 42:109–131, 1989.

[43] T. C. Shermer. *Visibility Properties of Polygons.* PhD thesis, McGill University, Montréal, June 1989.

[44] T. C. Shermer. Recent results in art galleries. *Proc. IEEE*, 80(9):1384–1399, 1992.

[45] T. C. Shermer and G. Mac Donald. Isomorphism of spiral polygons. manuscript, 1992.

[46] T. C. Shermer and G. Toussaint. Characterizations of star-shaped polygons. Technical Report 92-11, Simon Fraser University School of Computing Science, 1992.

[47] D. L. Souvaine and I. Bjorling-Sachs. The contour problem for restricted-orientation polygons. *Proc. IEEE*, 80(9):1449–1470, 1992.

[48] S. Wagon. Infinite triangulated graphs. *Discrete Mathematics*, 22:183–189, 1978.

[49] P. Widmayer, Y. F. Wu, M. D. F. Schlag, and C. Wong. On some union and intersection problems for polygons with fixed orientations. *Computing*, 36:183–197, 1986.