Dear Sir,

I would greatly appreciate a reprint of your paper

"Explorations in restricted orientation"

(RR 87/57)

sent
AUG 11 1988

~~from~~

and related papers, if you have copies for distribution.

Thanking you in advance, Yours sincerely

K. Unger

377 Druckerei Aßmann Bm 915/85 5,0 214

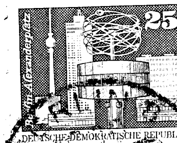

17.06.88-17

Postkarte

Kanada -
Canada

Prof G. J. E Rawlins
University of Waterloo
Dep. of Comp. Sci.

Waterloo, Ontario

N2L 3G1

Akademie der Wissenschaften der DDR
Karl-Weierstraß-Institut
für Mathematik
Kristel Unger
DDR 1086 Berlin
Mohrenstraße 39 · PF 1304

Explorations in Restricted Orientation Geometry

Gregory J.E. Rawlins

Research Report CS-87-57

Department of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
N2L 3G1

Faculty

of

Mathematics

University of Waterloo
Waterloo, Ontario, Canada
N2L 3G1

Explorations in Restricted Orientation Geometry

Gregory J.E. Rawlins

Research Report CS-87-57

Department of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
N2L 3G1

Explorations in Restricted Orientation Geometry

by

Gregory J. E. Rawlins

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Computer Science

Waterloo, Ontario, 1987

I hereby declare that I am the sole author of this thesis.

I authorize the University of Waterloo to lend this thesis to other institutions or individuals for the purpose of scholarly research.

Gregory Rawlins

I further authorize the University of Waterloo to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Gregory Rawlins

The University of Waterloo requires the signatures of all persons using or photocopying this thesis. Please sign below, and give address and date.

Abstract

Computational geometry is concerned with algorithms on geometric objects. For various theoretical and algorithmic reasons, computational geometry has singled out a special class of planar objects called *orthogonal objects* to investigate. Orthogonal objects are planar objects whose edge orientations are restricted to be either horizontal or vertical. This thesis demonstrates, through some example problems, that there is a much larger class of objects of restricted orientation that are worth investigating. *Restricted Orientation Geometry* is the study of the properties of planar geometric objects whose edge orientations are restricted. More generally, it is the study of the interaction of subclasses of such restricted objects and unrestricted planar objects.

In this thesis several results about restricted orientation convexity and visibility are proved. It is shown how the two apparently different notions of convex sets and orthogonally convex sets may be unified under restricted orientation convex sets. Further, several otherwise unsupported observations in the literature are validated.

This thesis also makes contributions to the sub-area of modern algebra called *abstract convexity theory*. A fundamental hull decomposition result follows from these contributions and also a new proof of a decomposition result on locked transaction systems.

Finally, several algorithms for restricted objects are developed and analysed. In particular, algorithms are given to construct the convex hull, to construct restricted versions of the kernel, and to place guards in restricted Art Galleries.

Acknowledgements

I wish to acknowledge the enormous debt of gratitude i owe to my supervisor, Dr. Derick Wood. This thesis would not have been possible without Derick's untiring direction, support and inspiration. I thank the other members of my thesis committee—Dr. Charles Colbourn, Dr. Ian Munro, Dr. Bruce Richmond, and my external examiner Dr. David Kirkpatrick—for being so careful in their reading and for suggesting many extensions to the work here presented.

I also wish to thank my readers Ehab El-Mallah and Darrell Raymond for catching many obscurities and infelicities in an earlier draft of this thesis. Finally, a special vote of thanks to Terence Jones who made the printing of the final version possible.

It is impossible to name here all the people who have contributed to my peace of mind during my doctoral studies so i content myself with thanking my parents for encouraging me and my friends for bearing with me. Thank you all.

To My Parents,

who made it possible,

And Especially To My Mother,

who couldn't stay.

Contents

List of Figures

Chapter 1

Introduction

The existence of analogies between central features of various theories implies the existence of a general theory which underlies the particular theories and unifies them with respect to those central features.

—E. H. Moore. *Introduction to a Form of General Analysis,* New Haven Mathematics Colloquium, Yale University Press, 1910.

1.1 Computational Geometry and Orthogonal Objects

Computational geometry, which might be defined as the study of algorithmic solutions to geometric problems, has found application in computer graphics, VLSI design, robotics, geographic databases, pattern recognition and computer aided design ([47]).

Within computational geometry the study of *orthogonal objects* has achieved a special significance ([64]). An orthogonal object is a planar figure whose edges are parallel to either the x or the y axis. Orthogonal objects are also known as rectilinear, isothetic, iso-oriented, x-y or aligned objects. Not only are orthogonal objects simpler—that is, more restricted—than arbitrary geometric objects but algorithms for geometric objects are often more efficient, descriptively simpler, or both, when the objects are restricted to being orthogonal.

Thus, orthogonal objects are of special interest to computational geometry both from a theoretical and an algorithmic point of view. Indeed, a case could be made that planar computational geometry has to some extent split into two camps—results for orthogonal objects and everything else. As evidence, the reader should note that both of the two textbooks in computational geometry ([47,38]) devote about one-sixth of their space to these special objects. The ratio is about the same for the most recent ACM computational geometry conference (*Third Annual ACM Symposium on Computational Geometry,* ACM Press, 1987).

This thesis examines to what extent these theoretical and algorithmic advantages are restricted solely to orthogonal objects and whether such advantages can be extended to more general objects. To accomplish this we have generalized the usual notions of convexity and visibility and have established connections between the special case of orthogonal objects and that of general objects. We call this sub-area *Restricted Orientation Geometry.*

1.2 Restricted Orientation Geometry

As mentioned previously computational geometry has been concerned, almost exclusively, with orthogonal objects and arbitrary objects. Recently ([21,22,23,63,62]) there have been attempts to unify these apparently disparate classes of objects by considering the geometry of planar objects of *restricted orientation*, that is, planar objects whose edges have orientations that belong only to some given set of orientations, $\mathcal{O}$. (The *orientation* of a line in the plane is the smaller of the two angles the line makes with the positive x axis.)

The notion of $\mathcal{O}$-orientation, for finite $\mathcal{O}$, was introduced by Güting in 1983, in [21], and followed up in [22,23]. Güting was concerned mainly with designing fast algorithms to find intersections of $\mathcal{O}$-oriented objects. The metric properties of $\mathcal{O}$-oriented objects, (again for finite $\mathcal{O}$), were investigated, contemporaneously with the work reported in this thesis, by Widmayer *et al.* [63,62]. Edelsbrunner [14] and also Culberson and Rawlins [9] considered two related classes of restricted objects.

Restricted Orientation Geometry is the study of the geometric properties of planar objects whose edge orientations belong to some set of orientations. More generally, it is the study of the interaction of some subclass of such restricted objects and general objects in the plane. In this thesis we will only be concerned with closed sets of orientations.

The reader should not be misled by the adjective "restricted" into thinking that restricted orientation geometry is somehow a sub-area of Euclidean geometry, indeed we allow *any* closed subset of the set of all possible orientations to be the given set of orientations. Thus restricted orientation geometry contains Euclidean geometry *as a special case*, namely when the set of orientations is the set of all orientations. Of course, it also contains orthogonal geometry as another special case.

1.3 Origins of the Dichotomy

Possibly the main reason for the distinction made between orthogonal objects and objects of arbitrary orientation is that the major application areas of computational geometry, namely: VLSI design; computer aided design; digital picture processing; computer vision and computer graphics, have traditionally placed emphasis on orthogonally-oriented objects.

This traditional dependence on orthogonal objects is due, mainly, to technical restrictions; for example, most input/output devices and layout schemes have traditionally been orthogonal. It may be also partly attributable to a form of mental laziness; we have been conditioned by cartesian coordinates into thinking in terms of horizontal and vertical axes.

More recently, however, VLSI designs have allowed objects to have more than two orientations and, as a result, designers are now concerned with horizontal, vertical, 45° and 135° orientations (see Widmayer *et al.* [63,62]). Some companies even allow any finite number of orientations. This suggests that we should investigate the geometry of objects whose edges belong to more than the usual two, but less than all possible, orientations.

Another justification for the special study of orthogonal objects is that algorithms for orthogonal objects are very "tight" since we can often examine all possible cases when designing, and/or proving correct, our algorithms (see Sack [54]). It is natural to speculate on whether we can increase the number of allowed orientations and still have fast and/or easy to prove correct algorithms.

As in the digitization of images we ask whether efficient algorithms can be devised for *approximations* of arbitrary objects where the original edges have been approximated by edges chosen from some given set of orientations. The question arises: What can be said about objects whose edges have orientations in some given set *vis a vis* the two extreme cases of orthogonal objects on the one hand and arbitrary objects on the other?

The dichotomy between orthogonal and arbitrary objects is not specific to computational geometry but can be seen in many areas of mathematics. For example, planar geometry can be divided into the study of the (simple, well-defined, highly restricted) regular polygons and arbitrary polygons. Once the more elementary properties of these two extremes are exhausted (in any area) some suitable generalization is attempted that will throw light on the objects "in between." This thesis examines classes of objects in between orthogonal objects and arbitrary objects with respect to the twin notions of visibility and convexity.

One of the more important questions to be answered is: is it possible to easily extend algorithms designed for orthogonal objects to restricted orientation objects? For some intersection problems Güting [21] has shown that they can be solved in $O(|\mathcal{O}|^2 f(n))$ time, where $\mathcal{O}$ is a *finite* set of orientations and $f(n)$ is the worst case time bound for orthogonal objects. This result should be contrasted with the optimal $\Theta(f(n) + r)$ time solution for the convex hull problem in chapter 7 (where $\mathcal{O}$ is an *arbitrary* closed set of orientations, r is the number of disjoint ranges in $\mathcal{O}$, and $f(n)$ is the worst case time bound for computing the orthogonal convex hull).

1.4 Convex Sets

A *convex set* is a set of points whose intersection with any line is either empty or connected. Convex sets are a comparatively recent but fruitful concept in geometry having applications in optimization, statistics, geometric number theory, functional analysis and combinatorics (see Klee [32] and Preparata and Shamos [47]).

Computational geometry is directly concerned with such sets for several reasons. The main reason is that the smallest convex set containing an object (the *convex hull*) is, typically, simpler than the object itself but is still representative of the object and so is much used in testing for intersections among objects (see for example Preparata and Shamos [47] and Toussaint [58]). Computer graphics and computer vision frequently use the "bounding box" of an object for precisely the same reason.

Also, algorithms are often more efficient when restricted to convex polygons and so polygons are often decomposed into convex subpolygons in order to answer various queries efficiently (see Preparata and Shamos [48] and Toussaint [58]). Finally, the convex hull was one of the first concepts studied in computational geometry (see Shamos [55]) and so deserves especial attention.

1.4.1 Orthogonal and Restricted Orientation Convexity

A set is said to be *orthogonally convex* if its intersection with any horizontal or vertical line is either empty or connected. Orthogonally convex sets were defined and investigated in computational geometry as a byproduct of the emphasis placed on orthogonal objects. The motivation for such a definition is that in many instances (for example, in VLSI design) the possible lines are restricted to horizontal and vertical ones and intersection of these lines and planar objects is the main concern. Orthogonal convexity has been defined not only in computational geometry (see Wood [64] for further references) but also in digital picture processing (see [53,29,6]) and in combinatorics (for example, see Bender [2]).

It was felt that there should be a class of orthogonal objects analogous to normal convex sets. The hope was that this class of objects would possess the same advantages as their more general cousins (that is, normal convex sets). In particular, the notion of the *hull* of an orthogonal object should be well defined and simpler, in general, than the original object.

Several researchers attempted to define orthogonally convex sets and their associated hulls in such a way as to retain all of the advantages of normal convex sets (see Nicholl *et al.* [41], Montuno and Fournier [39] and Ottmann *et al.* [43]). However, in [44], it was realized that there were several problems with the alternative definitions chosen and that there needed to be more investigation to determine a better one. In this thesis we construct a theory of *$\mathcal{O}$-convex sets* which leads to a reasonable definition of orthogonal convexity that, in a sense to be made precise later, maximizes the number of analogous properties that can hold between convex and orthogonally convex sets.

A set is said to be *$\mathcal{O}$-convex* if its intersection with any line whose orientation belongs to the set of orientations $\mathcal{O}$ is either empty or connected. Earlier researchers worked under the self-imposed (and unnecessary) restriction of only dealing with orthogonal sets. By applying our definition to *arbitrary* sets of points we generalize all previous results and, at the same time, verify some otherwise unsupported observations in the literature.

As it turns out, this is an advantageous generalization since the new notion completely encompasses the old and there is no additional complexity. Indeed, the hull algorithms are slightly simplified. The reason for the simplification is that the generalization allows the elimination of inessential details that were specific only to orthogonal polygons.

Curiously, our investigation demonstrates that in the plane we may treat restricted-orientation convexity just as if it were orthogonal convexity. In other words, it seems that *for convexity problems* it is always possible to construct a case analysis which is only concerned with at most two orientations at a time.

1.4.2 Abstract Convexity and Abstract Visibility

Abstract convexity theory is a branch of Topology and Modern Algebra that is concerned with classes of sets over some ground set which obey two weak axioms (closure under intersection and possession of both the ground set and the empty set). Abstract convexity theory could be said to have started with Levi's ground-breaking paper ([35]) but has only received real interest in the past four years or so (see Jamison [24] for further references).

Abstract convexity theory is a useful starting point to the work in this thesis because it serves as useful point of vantage from which to view basic properties of convex sets. It is quite easy to show that restricted-orientation convex sets form an abstract convexity space and all algebraic results derivable from purely formal manipulation in abstract convexity theory have analogues in the plane.

Unfortunately, researchers in abstract convexity have mainly limited themselves to the study of "Helly-type" ([10]) properties (that is, mainly combinatorial properties) of such topological spaces and not the more interesting (for computational geometry) problems of visibility and connectedness in such topological spaces. As a result, abstract convexity theory does not help us with questions related to visibility. In an attempt to remedy this defect we add two simple axioms to the usual definition of a convexity space and lay the groundwork for an abstract theory of visibility as well as convexity.

1.5 Description of this Thesis

This thesis is organized along the following lines:

Chapter 1 (this chapter) gives a statement of the general aims of the thesis and it points to related work in this area.

Chapter 2 contains some of the definitions and conventions that are used throughout the thesis. For the reader's convenience, following chapter 8, there is a brief list of the conventions followed and a glossary of terms used in this thesis.

Chapter 3 defines the area of *abstract convexity*. It contains several useful lemmas and theorems that find application in chapters 5 and 6. In particular, the *Refinement Theorem* and (a weak) *Decomposition Theorem* are proved. This chapter ends with an application of the results derived so far to locked transaction systems.

Chapter 4 defines the important concept of a *stairline* and considers some elementary separation results. Stairlines are generalizations of lines that maintain the relevant properties that we need for our more generalized convex sets. Because they generalize lines they are also fundamental to generalized visibility and separation. The chapter ends with algorithms to find the intersection of two polygonal stairlines.

Chapter 5 develops the theory of restricted-orientation convex sets and contains the proof of the *Decomposition Theorem* and related results. These results are used to characterize the generalized convex hull and to reduce the time complexity of the hull algorithms (given in chapter 7). Parts of this chapter, restricted to special cases, have appeared in [51]. A more comprehensive treatment has appeared as an internal report in [50].

Chapter 6 investigates a generalization of *visibility* based on the notion of a stairline introduced in chapter 4. It considers several variant definitions of visibility and shows how to unify them using some simple algebraic properties. For one of the variants of visibility we investigate the *kernel* and *starshaped sets*. This chapter contains algorithms to find the kernel and to place Art Gallery guards for special kinds of polygons and under special kinds of visibility. It concludes with two generalizations of the well-known *Art Gallery Theorem* of Chvátal ([7]).

Chapter 7 gives *convex hull algorithms* for polygons arising from the theoretical considerations of chapters 4 and 5. It also gives algorithms to find the hulls of point sets. An early version of the work reported in chapter 7 has appeared in [49].

Finally, chapter 8 concludes the thesis with a list of open problems.

Chapter 2

Definitions and Conventions

In this thesis, all objects defined and results stated are in the plane ($\Re^2$).

In this chapter we give short, intuitive definitions of several basic topological, geometric, and computational concepts. The interested reader is referred to [4] for the topological, [20,17] for the geometrical, and [47] for the computational definitions.

2.1 Topological Background

If we treat the plane as a rubber sheet then any deformation of the plane without tearing or folding is called a *topological transformation.* A set of points in the plane is *simply connected* if it is the image of a disc under a topological transformation. A point set is *path connected* if for every two points in the point set there is a continuous path in the point set joining the two points. A *neighbourhood* of a point is the set of all points within some disc centered at the point. The *boundary* of a point set is the set of points in the point set all of whose neighbourhoods contain at least one point in the point set and one point not in the point set. Notice that a point may be on the boundary of a set and not actually be in the set. The *interior* of a point set is the set of all points in the point set which are not boundary points. A set is said to be *closed* if it contains all of its boundary points.

In this thesis *connected* is used to mean *path connected*. Also, to avoid special cases, we define the empty set and all singleton sets to be (vacuously) connected.

2.2 Sets of Points and Polygons

We shall, without further comment, denote subsets of $\Re^2$ by **bold face capital letters** (e.g., **P** and **Q**) and elements of such sets by *lower case italic letters* (e.g., p and q). We treat a subset of $\Re^2$ as a set of interior points together with its boundary. That is, unless specified otherwise, all subsets of $\Re^2$ we consider are assumed to be *closed*.

By a *curve* we mean the image of $[0, 1]$ under any continuous function (see [4]). We assume that all curves are *simple* (that is, they do not cross themselves, see figure 2.1). A *polygonal curve* is a finite sequence of line segments such that each segment endpoint is shared by at most two line segments. Moreover, we allow the first and/or last line segment to be *rays* (that is, to extend to infinity). We assume that all polygonal curves are simple (see figure 2.2).

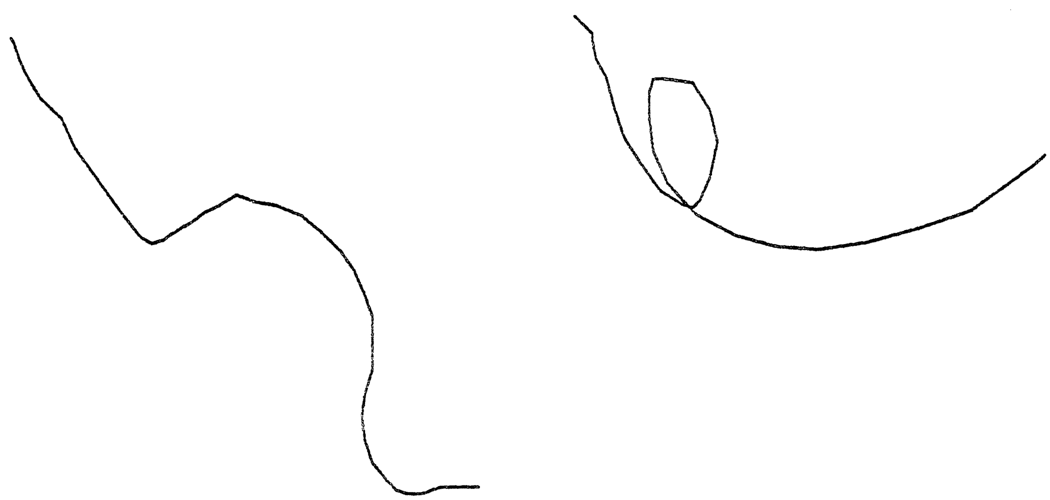

Figure 2.1: Simple and Non-Simple Curves

A *closed polygonal curve* is a finite sequence of line segments such that each segment

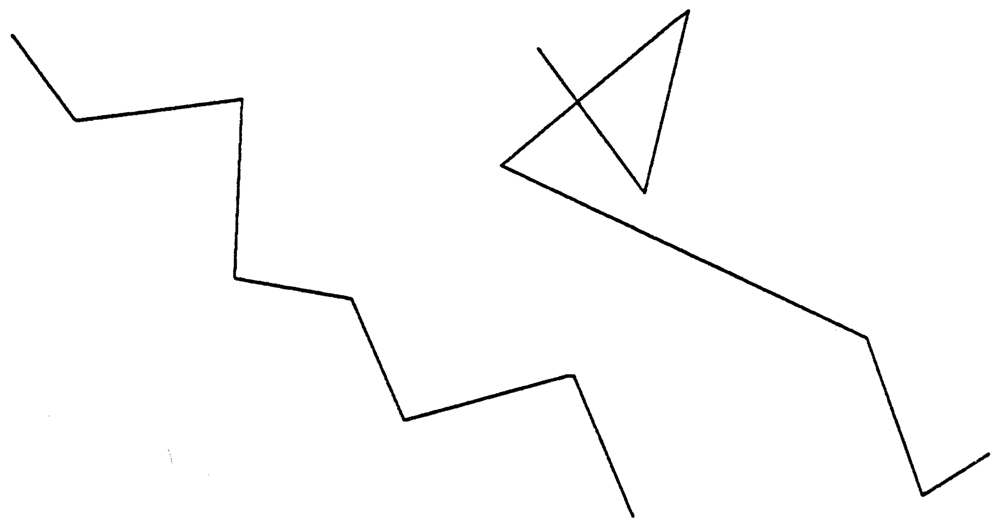

Figure 2.2: Simple and Non-Simple Polygonal Curves

endpoint is shared by exactly two line segments. A *polygon* is a connected region of the plane which is bounded by a closed polygonal curve. Note that a polygon is not necessarily simply connected (that is, the topological image of the unit disc) since line segments making up its boundary may cross each other at points other than their endpoints and the closed polygonal curve may be disconnected—that is, the polygon may contain *holes* (see figure 2.3). A polygon is *simple* if it is simply connected. All polygons and polygonal curves discussed in this thesis are simple unless explicitly stated otherwise. We assume that polygons are described as a list of vertices in clockwise order where no three consecutive vertices are collinear. It is common to treat polygons both as regions of the plane and as polygonal curves (that is, the boundary of such a region) ([47]). We will follow this convention wherever it does not lead to confusion.

2.3 Ranges

We use the convention that a square bracket ("[" or "]") bracketing a range on one endpoint implies that the endpoint is included in the range and a parenthesis ("(" or ")") implies that the endpoint is excluded. So, the range "$[\theta_1, \theta_2)$" is half-open.

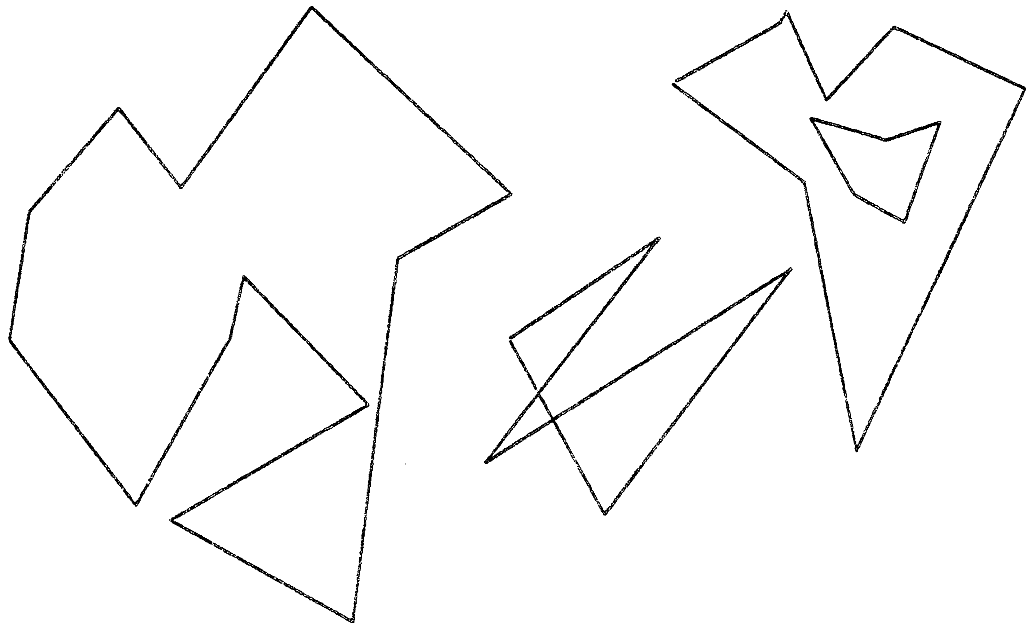

Figure 2.3: Simple and Non-Simple Polygons

We use angle brackets ("$\langle$" or "$\rangle$") when we wish to make statements that apply whether one or both of the endpoints are included. Thus, for example, any statement involving the range "$[\theta_1, \theta_2\rangle$" applies whether θ_2 is or is not included in the range, that is, whether the range is closed or half-open.

2.4 Sets of Orientations

The *orientation* of a line is the smaller of the two possible angles it makes with the positive x-axis (see figure 2.4). For the purposes of this thesis, all orientations are assumed to be mapped into the range $[0°, 180°)$ and we assume that $0°$ is identified with $180°$. This identification has the technical corollary that the set of orientations $[0°, \theta_1] \cup [\theta_2, 180°)$ *is connected.*

The symbol $\mathcal{O}$, with or without subscripts, is used to refer to a (possibly empty) *closed* set of orientations, that is:

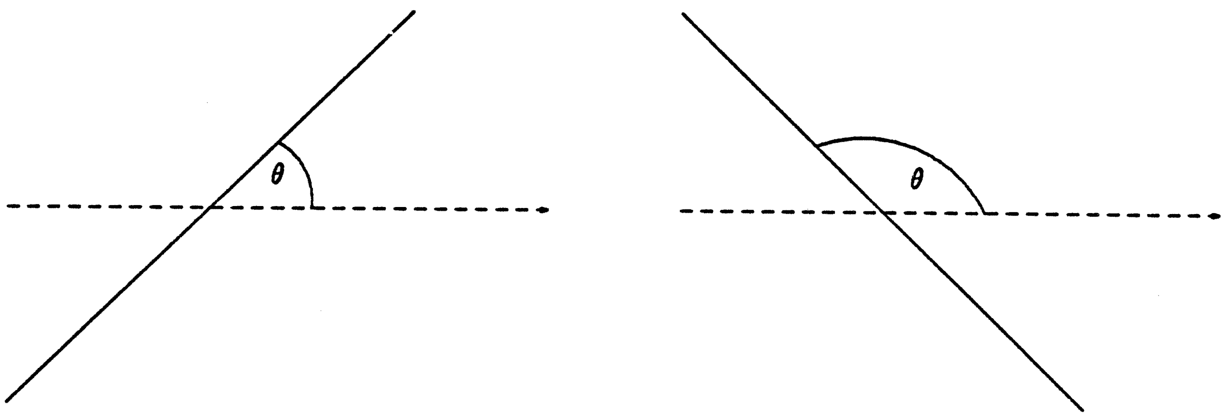

Figure 2.4: The Orientation of A Line

- $\mathcal{O} \subseteq [0°, 180°)$ and
- $\mathcal{O}$ consists of a set of closed disjoint ranges.

For example, $\mathcal{O}$ may be the set $\mathcal{O} = \{0°, \theta_5, \theta_6\} \cup \{\theta \mid \theta_1 \leq \theta \leq \theta_2\} \cup \{\theta \mid \theta_3 \leq \theta \leq \theta_4\}$ where each set is disjoint and $0° \leq \theta_i < 180°$ for all $1 \leq i \leq 4$.

For conciseness, we abuse set notation and abbreviate the description of $\mathcal{O}$ by listing its (maximal) ranges within braces unless $\mathcal{O}$ consists of one range in which case we write the range with no surrounding braces. For example, we describe the above set of orientations as $\{0°, [\theta_1, \theta_2], [\theta_3, \theta_4], \theta_5, \theta_6\}$, where it is understood that $0° < \theta_1 < \theta_2 < \theta_3 < \theta_4 < \theta_5 < \theta_6 < 180°$, but the set of orientations $\{\theta \mid \theta_1 \leq \theta \leq \theta_2\}$ is written as $[\theta_1, \theta_2]$.

Since $\mathcal{O}$ is a subset of an ordered set and each range is disjoint, the *next* range in $\mathcal{O}$ after a given range is well-defined (the successor of the last range is the first range). In the above example, $[\theta_1, \theta_2]$ is the next range after the range $0°(= [0°, 0°])$.

The range $\langle\theta_1, \theta_2\rangle$ is said to be *$\mathcal{O}$-free* if $\mathcal{O} \cap \langle\theta_1, \theta_2\rangle = \emptyset$. Observe that, if a range is *empty*, then it is $\mathcal{O}$-free. The range $\langle\theta_1, \theta_2\rangle$ is said to be a *maximal* $\mathcal{O}$-free range if it is $\mathcal{O}$-free and it is not a proper subset of any other $\mathcal{O}$-free range. In the above example the maximal $\mathcal{O}$-free ranges are $(0°, \theta_1), (\theta_2, \theta_3), (\theta_4, \theta_5), (\theta_5, \theta_6)$ and $(\theta_6, 180°)$.

A collection of lines, line segments or rays is said to be *$\mathcal{O}$-oriented* if the set of orientations of the elements of the collection is a subset of $\mathcal{O}$. Thus we speak of "$\mathcal{O}$-lines," "$\mathcal{O}$-line segments," "$\mathcal{O}$-rays," "$\mathcal{O}$-polygonal curves" and "$\mathcal{O}$-polygons" to mean $\mathcal{O}$-oriented lines, line segments, rays, polygonal curves and polygons.

2.5 The Orthogonal Transformation

For any collection of lines some of which have one of two orientations θ_1 and θ_2 $(0° \leq \theta_1 < \theta_2 < 180°)$ we may rotate the collection until the θ_1 lines are horizontal. Thus assume without loss of generality that we are given a collection of lines each of which is either horizontal or has orientation θ $(0° < \theta < 180°)$.

The transformation

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

maps all horizontal lines to horizontal lines and maps all θ-lines $(0° < \theta < 180°)$ to vertical lines. Since the transformation is *affine* (see Gans [17]) all incidences are preserved and none are introduced.

T is said to be *the orthogonal transformation.* The existence of T implies that any geometric problem for lines of two orientations which is not distance related can be transformed into *a similar problem for orthogonal lines.* Thus, when faced with a collection of lines, line segments and rays, we may, *without loss of generality*, assume that any two distinct orientations are orthogonal.

2.6 Lines, Line Segments and Halfplanes

For two distinct points p and q, we use the notation $\mathcal{L}[p,q]$ to mean the line passing through them and, similarly, the notation $\mathcal{LS}\langle p,q\rangle$ to mean the (open, half-open or closed) line segment with endpoints p and q. As with endpoints of ranges we use the convention that a square bracket implies closure at that endpoint, a parenthesis implies non-closure at that endpoint, and an angle bracket implies either. We use the notation $\mathcal{L}[\theta, p]$ to mean the line of orientation θ passing through the point p.

We use the notation $\mathcal{H}\langle\mathcal{L}, p]$ to mean the (open or closed) halfplane bounded by the line $\mathcal{L}$ that includes the point p. Again, we use $\mathcal{H}(\mathcal{L}, p]$ (respectively, $\mathcal{H}[\mathcal{L}, p]$) if the halfplane is open (closed).

We also use the notation $\Theta(\mathcal{L})$ (where $\mathcal{L}$ is a line, line segment or ray) to mean the orientation of $\mathcal{L}$.[1] If $\Theta(\mathcal{L}) \notin \mathcal{O}$, then by *the maximal $\mathcal{O}$-free range of $\mathcal{L}$* we mean the unique maximal $\mathcal{O}$-free range which contains $\Theta(\mathcal{L})$. If $\Theta(\mathcal{L}) \in \mathcal{O}$, then $\mathcal{L}$'s maximal $\mathcal{O}$-free range is empty.

As noted in the previous section any collection of lines, line segments or rays of exactly two orientations can be mapped bijectively to another collection having the same incidence structure as the first but with two completely different orientations.

When considering a particular $\mathcal{L}$ and $\mathcal{O}$ where $\Theta(\mathcal{L}) \notin \mathcal{O}$ and $\mathcal{O}$ has two or more orientations, we assume, for ease of exposition, that $(0°, 90°)$ is $\mathcal{L}$'s maximal $\mathcal{O}$-free range.

2.7 Monotone, Convex and Starshaped Sets

$\mathbf{P}$ is said to be *monotone* with respect to a line $\mathcal{L}$ if for every line $\mathcal{L}_1$ orthogonal to $\mathcal{L}$ the intersection of $\mathbf{P}$ and $\mathcal{L}$ is either empty, a point or a line segment (the figures in figure 2.5 are monotone with respect to the x axis).

$\mathbf{P}$ is said to be *convex* if for every line $\mathcal{L}$ the intersection of $\mathbf{P}$ and $\mathcal{L}$ is either empty,

[1] This notation conflicts with the usual convention used to state optimality of algorithms (see the section on order notation), however, this should cause no confusion as it should be clear from the context which meaning of "Θ" is being used.

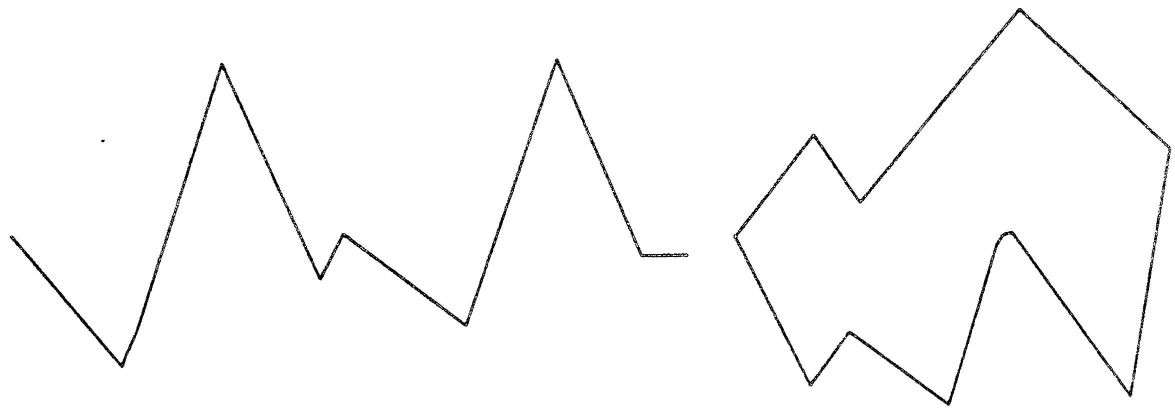

Figure 2.5: A Monotone Polygonal Curve and a Monotone Polygon

a point or a line segment. As we have defined the empty set and all singleton sets to be connected, we simplify this definition to: $\mathbf{P}$ is said to be convex if $\forall \mathcal{L}$, $\mathbf{P} \cap \mathcal{L}$ is connected. Alternately, $\mathbf{P}$ is convex if $\forall p, q \in \mathbf{P}$, $\mathcal{LS}[p,q] \subseteq \mathbf{P}$. If $\mathbf{P}$ is convex then $\mathbf{P}$ is monotone with respect to every line $\mathcal{L}$.

The smallest convex set containing $\mathbf{P}$ is said to be the *convex hull* of $\mathbf{P}$ and is denoted $hull(\mathbf{P})$ (figure 2.6 shows the convex hull of the dashed line polygon).

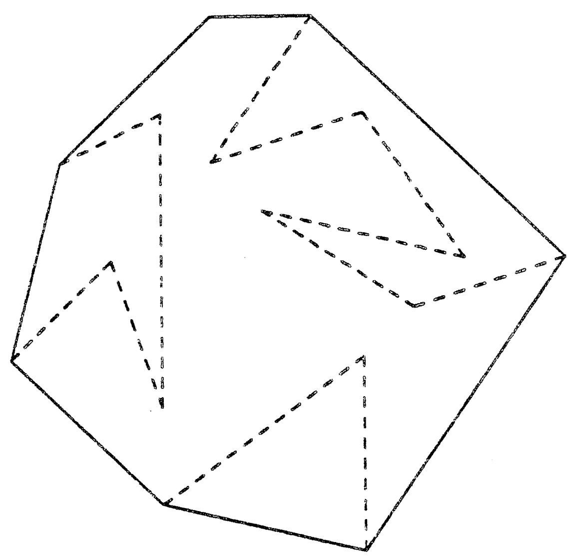

Figure 2.6: The Convex Hull of A Polygon

A line $\mathcal{L}$ is said to *separate* two sets $\mathbf{P}$ and $\mathbf{Q}$ if $\mathbf{P}$ and $\mathbf{Q}$ lie wholly in separate halfplanes bounded by $\mathcal{L}$. $\mathcal{L}$ is a *line of support* of $\mathbf{P}$ if $\mathcal{L}$ intersects $\mathbf{P}$ and $\mathbf{P}$ lies wholly in one of the two halfplanes defined by $\mathcal{L}$. A point p in $\mathbf{P}$ is an *extremal point* of $\mathbf{P}$ if there is a line of support of $\mathbf{P}$ through p. If $\mathbf{P}$ is convex then every point on $\mathbf{P}$'s boundary is an extremal point of $\mathbf{P}$.

A point p in $\mathbf{P}$ is said to *see* a point q in $\mathbf{P}$ if $\mathcal{LS}[p,q] \subseteq \mathbf{P}$. $\mathbf{P}$ is said to be *starshaped* if there is a point in $\mathbf{P}$ which sees every point in the set, that is, $\mathbf{P}$ is starshaped if $\exists p \in \mathbf{P}$ such that $\forall q \in \mathbf{P}\ \ \mathcal{LS}[p,q] \subseteq \mathbf{P}$. The set of all points that see every point in $\mathbf{P}$ is called the *kernel* of $\mathbf{P}$. If $\mathbf{P}$ is convex then $\mathbf{P}$ is starshaped and the kernel of $\mathbf{P}$ is $\mathbf{P}$ itself.

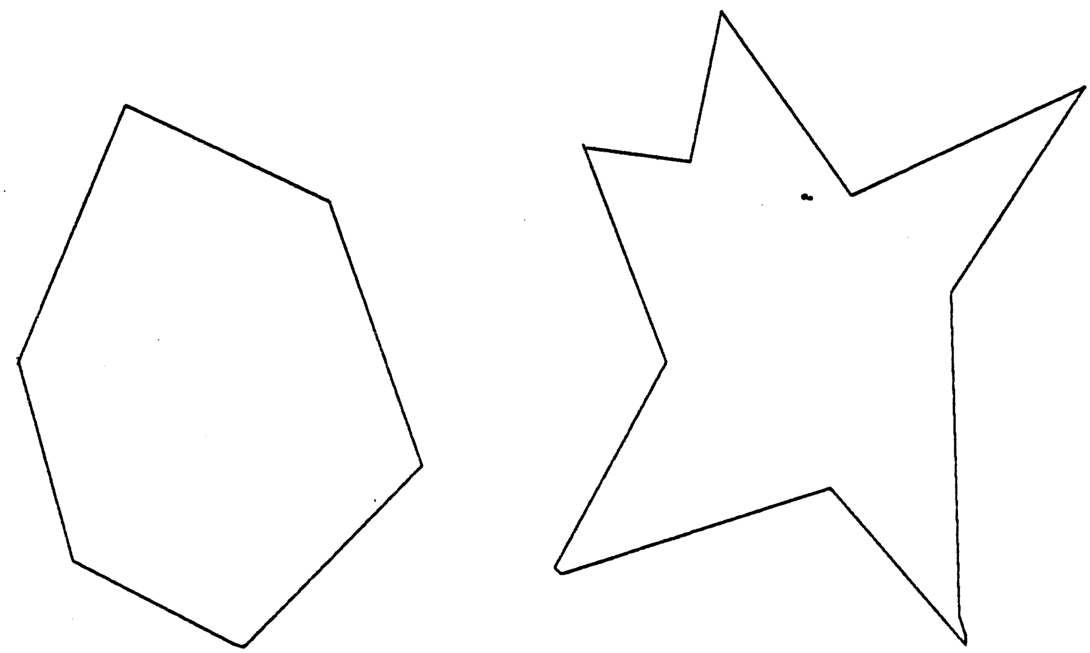

Figure 2.7: Convex and Starshaped Polygons

2.8 Order Notation

For two positive functions f and g

- $g(n) = O(f(n))$ if there exists positive constants c and n_0 with $g(n) \leq cf(n)\ ,\quad \forall n \geq n_0$.

- $g(n) = \Omega(f(n))$ if there exists positive constants c and n_0 with $g(n) \geq cf(n)$, $\forall n \geq n_0$.
- $g(n) = \Theta(f(n))$ if there exists positive constants c_1, c_2 and n_0 with $c_1 \leq g(n) \leq c_2 f(n)$, $\forall n \geq n_0$.
- $g(n) = o(f(n))$ if $\lim_{n\to\infty} g(n)/f(n) = 0$.

2.9 Models of Computation

In this thesis we adopt the following two models to demonstrate upper and lower bounds on various problems.

For upper bounds we use a modification of the standard *random-access machine* (RAM) model of computation (see [1]) called the *real RAM model*. Besides the standard attributes of a RAM the real RAM has the following extra capabilities:

1. Each memory location is capable of holding a single real number.
2. A single arithmetic operation $(+, -, \times, /)$; a comparison between two reals; a single k^{th} root, k^{th} power, or trigonometric operation; and a single indirect memory address each cost unit time.

For lower bounds we use the *algebraic decision tree* model of computation (see [11]) which may be described as follows:

An algebraic decision tree on a set of variables $\{x_1, \ldots, x_n\}$ is a finite, loop-free binary tree program with statements $L_1, \ldots, L_p$ of the form:

1. Compute $y := f(x_1, \ldots, x_n)$; if $y = 0$ then goto L_i else goto L_j. (Here, f is an algebraic function).
2. Halt and output 'YES'.

3. Halt and output 'NO'.

It has been shown that algebraic decision trees and real RAMs behave identically for a wide class of problems ([11]). This class includes the problems discussed in this thesis.

Within the algebraic decision tree model the following problem has been shown ([11]) to take $\Omega(n \lg n)$ time on inputs of size n:

The Element Uniqueness Problem: Given n real numbers decide whether any two are equal.

We will use this problem and its lower bound to establish lower bounds on two problems in this thesis.

Chapter 3

Abstract Convexity

A *convex set* is a set of points whose intersection with any line is connected. An *orthogonally convex set* is a set of points whose intersection with any *horizontal or vertical* line is connected. Given a set of orientations $\mathcal{O}$ a set of points is said to be $\mathcal{O}$-*convex* if the intersection of the set with any $\mathcal{O}$-line is connected. Clearly this definition includes both of the above definitions as special cases. In this chapter we investigate an abstraction of convex sets (known as an *abstract convexity space*) which includes $\mathcal{O}$-convex sets as a special case. We show that classes of abstract convexity spaces obey a (weak) decomposition result (theorem 3.4.3) and we apply this result to a problem in locked transaction systems. A strengthened form of this decomposition result also finds application in chapter 5.

3.1 Convexity Spaces

Abstract convexity theory is concerned with collections of subsets of a set which obey two weak axioms. Such a collection of sets has been variously called a *convexity space* ([35]), *convexity structure* ([56]), *alignment* ([24]), or *algebraic closure system* ([8]). We shall use the term "convexity space" since we are more interested in the geometry, as opposed to the algebra, of such a collection. A convexity space, in the sense we use it here, is intended to be an abstraction of the more essential properties of convex sets in $\Re^n$.

Definition: Given a set, $\mathcal{S}$, and a family, $\mathcal{C}$, of subsets of $\mathcal{S}$ the structure $(\mathcal{S}, \mathcal{C})$ is said to be a *convexity space* if

1. $\emptyset, \mathcal{S} \in \mathcal{C}$ and
2. $\forall C \subseteq \mathcal{C}\ , \quad \bigcap C \in \mathcal{C}$,

   where $\bigcap C = \bigcap_{X \in C} X$.

$\mathcal{S}$ is called the *groundset* of the convexity space and any element of the family $\mathcal{C}$ is said to be $\mathcal{C}$*-convex*. The interpretation of the family $\mathcal{C}$ is that it is the set of all convex sets over some space, where we have deferred the operational question of what we *mean* by convexity. The dominant characteristic of convex sets is taken to be closure under intersection.

Definition: Given a convexity space $(\mathcal{S}, \mathcal{C})$ we define the associated *hull operator*, $\mathcal{C}$*-hull*, as follows:

$$\forall \mathbf{P} \subseteq \mathcal{S}\ , \quad \mathcal{C}\text{-}hull(\mathbf{P}) = \bigcap \{\mathbf{Q} \mid \mathbf{P} \subseteq \mathbf{Q} \ \wedge \ \mathbf{Q} \in \mathcal{C}\}$$

It is straightforward to show that $\mathcal{C}\text{-}hull(\mathbf{P})$ exists, is unique and is the smallest $\mathcal{C}$-convex set which contains $\mathbf{P}$.

3.2 Example Convexity Spaces

Here are some example convexity spaces:

1. *The trivial convexity space* $(\mathcal{S}, \{\emptyset, \mathcal{S}\})$. In this convexity space there are only two convex sets—$\emptyset$ and $\mathcal{S}$—thus the hull of any non-empty subset of $\mathcal{S}$ is $\mathcal{S}$ itself.

2. *The complete convexity space* $(\mathcal{S}, \mathcal{P}(\mathcal{S}))$ where $\mathcal{P}(\mathcal{S})$ is the powerset of $\mathcal{S}$. In this convexity space every set is its own hull, or to put it another way, all subsets of $\mathcal{S}$ are convex since $\forall \mathbf{P} \subseteq \mathcal{S}$, $\mathbf{P} \in \mathcal{P}(\mathcal{S})$ and so

$$\mathcal{P}(\mathcal{S})\text{-}hull(\mathbf{P}) = \bigcap\{\mathbf{Q} \mid \mathbf{P} \subseteq \mathbf{Q} \ \wedge \ \mathbf{Q} \in \mathcal{P}(\mathcal{S})\} = \mathbf{P}.$$

3. *The partition convexity space* $(\mathcal{S}, \{\emptyset, \mathcal{S}\} \cup \{\mathcal{S}_i \mid i \in I\})$ where the $\mathcal{S}_i$'s partition $\mathcal{S}$. That is, $\cup_{i \in I} \mathcal{S}_i = \mathcal{S}$ and $\forall i, j \in I$, $i \neq j \Longrightarrow \mathcal{S}_i \cap \mathcal{S}_j = \emptyset$.

   As a special case we have the *point convexity space* $(\mathcal{S}, \{\emptyset, \mathcal{S}\} \cup \{\{s\} \mid s \in \mathcal{S}\})$.

   As this example shows the *points* (elements) of $\mathcal{S}$ are *not necessarily* $\mathcal{C}$-convex.

4. $(\mathcal{S}, \mathcal{T}(\mathcal{S}))$ where $\mathcal{S}$ is the complete graph on n nodes and $\mathcal{T}(\mathcal{S})$ is the set of all transitively (or reflexively or both) closed subgraphs of $\mathcal{S}$.

   This convexity space is usually expressed in terms of relations. For example, if $\mathcal{R}$ is a relation then the *transitive closure* of $\mathcal{R}$ is the smallest relation $\mathcal{R}^+$ such that:

   - $\mathcal{R} \subseteq \mathcal{R}^+$ and
   - if $(a, b) \in \mathcal{R}^+$ and $(b, c) \in \mathcal{R}$ then $(a, c) \in \mathcal{R}^+$.

We now show that, for any set of orientations $\mathcal{O}$, the set of all $\mathcal{O}$-convex sets over the groundset $\Re^2$ forms an abstract convexity space.

Lemma 3.2.1 *For any set of orientations $\mathcal{O}$, if $\mathbf{P}$ is convex then $\mathbf{P}$ is $\mathcal{O}$-convex.*

Proof: If $\mathcal{O}$ is empty then, vacuously, all sets are $\mathcal{O}$-convex. Suppose that $\mathcal{O}$ is non-empty. Since convex sets are by definition sets whose intersection with any line is connected then they are *a fortiori* $\mathcal{O}$-convex for any $\mathcal{O}$. ∎

The following sets are convex, and hence $\mathcal{O}$-convex for any $\mathcal{O}$: the empty set, $\Re^2$, and, any point, line, line segment, ray or halfplane in $\Re^2$.

Lemma 3.2.2 *For any set of orientations $\mathcal{O}$, if $\mathcal{C}$ is a non-empty collection of $\mathcal{O}$-convex sets, then $\bigcap \mathcal{C}$ is $\mathcal{O}$-convex.*

Proof: The result is vacuously true if $\mathcal{O}$ is empty since all sets are $\emptyset$-convex. Assume that $\mathcal{O}$ is non-empty. If there are no two points in $\bigcap \mathcal{C}$ which lie on an $\mathcal{O}$-line then the intersection of any $\mathcal{O}$-line and $\bigcap \mathcal{C}$ is either empty or a single point and such sets are defined to be connected. On the other hand, if there are two points in $\bigcap \mathcal{C}$ which lie on an $\mathcal{O}$-line then they belong to each member of $\mathcal{C}$. Since each member is $\mathcal{O}$-convex, the segment joining such a pair is in each member of $\mathcal{C}$ and so is in $\bigcap \mathcal{C}$. Hence $\bigcap \mathcal{C}$ is $\mathcal{O}$-convex. ∎

THEOREM **3.2.1** *For any set of orientations $\mathcal{O}$, let $\mathcal{C}_{\mathcal{O}}$ be the set of all $\mathcal{O}$-convex sets in $\Re^2$.*

Then $(\Re^2, \mathcal{C}_{\mathcal{O}})$ is a convexity space.

Proof: (1) $\emptyset$ and $\Re^2$ are both convex sets and every convex set is $\mathcal{O}$-convex for each $\mathcal{O}$ (lemma 3.2.1), thus $\emptyset, \Re^2 \in \mathcal{C}_{\mathcal{O}}$. (2) $\mathcal{C}_{\mathcal{O}}$ is closed under intersection (lemma 3.2.2). Thus, every choice of $\mathcal{O}$ gives rise to a convexity space. ∎

3.3 Basic Results

The results stated in the following theorem are well-known in abstract convexity theory (see Kay and Womble [27]).

THEOREM **3.3.1** *Given a convexity space $(\mathcal{S}, \mathcal{C})$; then $\forall \mathbf{P}, \mathbf{Q} \subseteq \mathcal{S}$*

1. *$\mathcal{C}$-hull$(\mathbf{P}) \in \mathcal{C}$;*
2. *$\mathbf{P} \subseteq \mathcal{C}$-hull$(\mathbf{P})$;*

3. $\mathcal{C}$-*hull*$(\mathbf{P}) = \mathbf{P} \iff \mathbf{P} \in \mathcal{C}$;

4. $\mathbf{P} \subseteq \mathbf{Q} \Longrightarrow \mathcal{C}$-*hull*$(\mathbf{P}) \subseteq \mathcal{C}$-*hull*$(\mathbf{Q})$;

5. $\mathcal{C}$-*hull*$(\mathcal{C}$-*hull*$(\mathbf{P})) = \mathcal{C}$-*hull*$(\mathbf{P})$.

In topology any operator which has properties (2), (4) and (5) is known as a *closure operator* ([13]). Although we do not do it here it is possible to show that a closure operator *induces* a convexity space. That is, given a set $\mathcal{S}$ and a closure operator $\mathcal{C}$, the set of all $\mathcal{C}$-closed sets over $\mathcal{S}$ is a convexity space.

3.4 The Refinement Theorem

Part of the aim of this thesis is to elucidate the geometry of families of convex sets each convex with respect to a particular set of orientations $\mathcal{O}$. Each choice of $\mathcal{O}$ generates a different set of "convex" sets. Thus it is natural for us to consider *families* of convexity spaces over a fixed groundset $\mathcal{S}$. In the *Refinement Theorem* (theorem 3.4.2) we show that the union of several hulls of a set, where each hull is formed in a distinct convexity space, is a subset of a *"composed"* hull formed from the separate hulls, and, the composed hull is in turn a subset of the hull with respect to the intersection of the different convexity spaces. In the *Decomposition Theorem* (theorem 3.4.3) we specialize the refinement theorem to convexity spaces which act as if they were independent of each other (so called, *invariant* convexity spaces). The decomposition theorem is then applied to a problem on locked transaction systems.

Definition: Let $(\mathcal{S}, \mathcal{C}_1)$ and $(\mathcal{S}, \mathcal{C}_2)$ be two convexity spaces defined on the groundset $\mathcal{S}$. Then the space $(\mathcal{S}, \mathcal{C}_2)$ is said to be a *refinement* of the space $(\mathcal{S}, \mathcal{C}_1)$ if $\mathcal{C}_1 \subseteq \mathcal{C}_2$. Alternatively, $(\mathcal{S}, \mathcal{C}_1)$ is said to be *coarser* than $(\mathcal{S}, \mathcal{C}_2)$. Intuitively, any $\mathcal{C}_1$-convex set is $\mathcal{C}_2$-convex.

Notice that in the example convexity spaces given previously the complete convexity space is a refinement of the trivial convexity space. Indeed, the complete convexity space is a refinement of all convexity spaces over S and the trivial convexity space is coarser than any other convexity space over S.

Definition: Let (S, C_1) and (S, C_2) be two convexity spaces defined on the groundset S. We use the notation $C_1 \wedge C_2$ to represent the subset of C_1 produced by composing the two hull operators. That is,

$$C_1 \wedge C_2 = \{C_1\text{-}hull(C_2\text{-}hull(\mathbf{P})) \mid \mathbf{P} \subseteq S\}$$

Observe that we may simplify this definition to

$$C_1 \wedge C_2 = \{C_1\text{-}hull(\mathbf{P}) \mid \mathbf{P} \in C_2\}$$

The following example shows that the composition of two convexity spaces is not necessarily a convexity space:

Let $S = \{a, b, c\}$; $C_1 = \{\emptyset, S, \{a, c\}, \{b, c\}, \{c\}\}$; $C_2 = \{\emptyset, S, \{a\}, \{b\}\}$; then $C_1 \wedge C_2 = \{\emptyset, S, \{a, c\}, \{b, c\}\}$. But $C_1 \wedge C_2$ is not a convexity space since $\{a, c\} \cap \{b, c\} = \{c\} \notin C_1 \wedge C_2$.

Although the composition of two convexity spaces is not necessarily a convexity space we can extend the notion of a hull operator to such families of sets as follows: Given a family of subsets C of S and a set $\mathbf{P} \subseteq S$, the *C-hull* of $\mathbf{P}$ is the intersection of all sets in C which contain $\mathbf{P}$. Note that if (S, C) is not a convexity space, the hull of a set may not be in C. Furthermore, the *C-hull* of $\mathbf{P}$ is not defined if $\mathbf{P}$ is not in at least one set in C.

Definition: Let (S, C_1) and (S, C_2) be two convexity spaces defined on the groundset S. We use the notation $(C_1 \wedge C_2)$-*hull* to represent the composition of the two *C-hull* operators. That is, $(C_1 \wedge C_2)\text{-}hull(\mathbf{P}) \equiv C_1\text{-}hull(C_2\text{-}hull(\mathbf{P}))$. Note that for all $\mathbf{P} \subseteq S$, $(C_1 \wedge C_2)\text{-}hull(\mathbf{P})$ is well-defined.

The following example shows that hull operators do not commute under composition:

Let $S = \{a, b, c\}$; $\mathcal{C}_1 = \{\emptyset, S, \{a\}\}$; $\mathcal{C}_2 = \{\emptyset, S, \{a, b\}\}$. Then $(\mathcal{C}_1 \wedge \mathcal{C}_2)$-$hull(\{a\}) = S$ but $(\mathcal{C}_2 \wedge \mathcal{C}_1)$-$hull(\{a\}) = \{a, b\}$.

The following theorem establishes a weak form of commutativity of composed hull operators. We show here that two hull operators commute if one of the convexity spaces is a subset of the other. In fact, in this case the composed hulls are both equal to the hull formed by the coarser of the two convexity spaces.

THEOREM **3.4.1** *Let $(S, \mathcal{C}_1)$ and $(S, \mathcal{C}_2)$ be two convexity spaces defined on the groundset S then*

$$
\begin{aligned}
\mathcal{C}_1 \subseteq \mathcal{C}_2 \quad &\Longleftrightarrow \quad (1) \quad \forall \mathbf{P} \subseteq S \quad \mathcal{C}_2\text{-}hull(\mathbf{P}) \subseteq \mathcal{C}_1\text{-}hull(\mathbf{P}) \\
&\Longleftrightarrow \quad (2) \quad \forall \mathbf{P} \subseteq S \quad (\mathcal{C}_2 \wedge \mathcal{C}_1)\text{-}hull(\mathbf{P}) = \mathcal{C}_1\text{-}hull(\mathbf{P}) \\
&\Longleftrightarrow \quad (3) \quad \forall \mathbf{P} \subseteq S \quad (\mathcal{C}_1 \wedge \mathcal{C}_2)\text{-}hull(\mathbf{P}) = \mathcal{C}_1\text{-}hull(\mathbf{P})
\end{aligned}
$$

Proof:

1. If $\mathcal{C}_1 \subseteq \mathcal{C}_2$ then $\forall \mathbf{P} \subseteq S$ $\mathcal{C}_1$-$hull(\mathbf{P}) \in \mathcal{C}_2$. Hence, $\forall \mathbf{P} \subseteq S$ $\mathcal{C}_2$-$hull(\mathbf{P}) \subseteq \mathcal{C}_1$-$hull(\mathbf{P})$.

   Conversely, suppose that $\forall \mathbf{P} \subseteq S$ $\mathcal{C}_2$-$hull(\mathbf{P}) \subseteq \mathcal{C}_1$-$hull(\mathbf{P})$. If $\mathbf{P} \in \mathcal{C}_1$ then $\mathbf{P} \subseteq \mathcal{C}_2$-$hull(\mathbf{P}) \subseteq \mathcal{C}_1$-$hull(\mathbf{P}) = \mathbf{P}$. Therefore $\mathbf{P} = \mathcal{C}_2$-$hull(\mathbf{P})$ and hence $\mathbf{P} \in \mathcal{C}_2$. Hence $\mathcal{C}_1 \subseteq \mathcal{C}_2$.

2. If $\mathcal{C}_1 \subseteq \mathcal{C}_2$ then $\forall \mathbf{P} \subseteq S$ $\mathcal{C}_1$-$hull(\mathbf{P}) \in \mathcal{C}_2$. Hence, $\forall \mathbf{P} \subseteq S$ $\mathcal{C}_2$-$hull(\mathcal{C}_1$-$hull(\mathbf{P})) = \mathcal{C}_1$-$hull(\mathbf{P})$.

   Conversely, suppose that $\forall \mathbf{P} \subseteq S$ $(\mathcal{C}_2 \wedge \mathcal{C}_1)$-$hull(\mathbf{P}) = \mathcal{C}_1$-$hull(\mathbf{P})$. If $\mathbf{P} \in \mathcal{C}_1$ then $\mathcal{C}_1$-$hull(\mathbf{P}) = \mathbf{P}$ and hence $\mathcal{C}_2$-$hull(\mathbf{P}) = \mathbf{P}$, hence $\mathbf{P} \in \mathcal{C}_2$. Hence $\mathcal{C}_1 \subseteq \mathcal{C}_2$.

3. If $\mathcal{C}_1 \subseteq \mathcal{C}_2$ then from (1) we have that $\forall \mathbf{P} \subseteq S$ $\mathcal{C}_2$-$hull(\mathbf{P}) \subseteq \mathcal{C}_1$-$hull(\mathbf{P})$.

   This implies that $\forall \mathbf{P} \subseteq S$ $\mathcal{C}_1$-$hull(\mathcal{C}_2$-$hull(\mathbf{P})) \subseteq \mathcal{C}_1$-$hull(\mathcal{C}_1$-$hull(\mathbf{P})) = \mathcal{C}_1$-$hull(\mathbf{P})$.

   But $\mathcal{C}_1$-$hull(\mathbf{P}) \subseteq \mathcal{C}_1$-$hull(\mathcal{C}_2$-$hull(\mathbf{P}))$. Therefore, $\mathcal{C}_1$-$hull(\mathcal{C}_2$-$hull(\mathbf{P})) = \mathcal{C}_1$-$hull(\mathbf{P})$

We prove the converse by proving its contrapositive.

Suppose that $\forall \mathbf{P} \subseteq S \quad (\mathcal{C}_1 \wedge \mathcal{C}_2)\text{-}hull(\mathbf{P}) = \mathcal{C}_1\text{-}hull(\mathbf{P})$. If $\mathbf{P} \notin \mathcal{C}_2$ then $\mathbf{P} \subset \mathcal{C}_2\text{-}hull(\mathbf{P})$. But this implies that $\mathbf{P} \subset \mathcal{C}_1\text{-}hull(\mathcal{C}_2\text{-}hull(\mathbf{P})) = \mathcal{C}_1\text{-}hull(\mathbf{P})$. Therefore, $\mathbf{P} \notin \mathcal{C}_1$. Hence $\mathcal{C}_1 \subseteq \mathcal{C}_2$. ∎

Result (1) has been previously proved by Sierksma ([56]). We now prove that convexity spaces over the same groundset are closed under intersection but not union. This is used in the proof of the refinement theorem below.

Lemma 3.4.1 *Let $\{(S, \mathcal{C}_i) \mid i \in I\}$ be a non-empty family of convexity spaces defined on S; then the structure $(S, \bigcap_{i \in I} \mathcal{C}_i)$ is a convexity space on S.*

Proof: $\emptyset, S \in \mathcal{C}_i, \quad \forall i \in I \Longrightarrow \emptyset, S \in \bigcap_{i \in I} \mathcal{C}_i$.

$$\mathcal{C} \subseteq \bigcap_{i \in I} \mathcal{C}_i \Longrightarrow \mathcal{C} \subseteq \mathcal{C}_i, \quad \forall i \in I \Longrightarrow \bigcap \mathcal{C} \in \mathcal{C}_i \quad \forall i \in I \Longrightarrow \bigcap \mathcal{C} \in \bigcap_{i \in I} \mathcal{C}_i \ \blacksquare$$

The following example shows that a similar result does not hold for the union of convexity spaces:

Let $S = \{a, b, c\}$; $\mathcal{C}_1 = \{\emptyset, S, \{a, b\}\}$; $\mathcal{C}_2 = \{\emptyset, S, \{a, c\}\}$; and $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 = \{\emptyset, S, \{a, b\}, \{a, c\}\}$. But $(S, \mathcal{C})$ is not a convexity space since $\{a, b\} \cap \{a, c\} = \{a\} \notin \mathcal{C}$.

THEOREM **3.4.2 (The Refinement Theorem)** *Given $n \geq 1$ convexity spaces $(S, \mathcal{C}_i)$, $1 \leq i \leq n$, then, $\forall \mathbf{P} \subseteq S$,*

$$\bigcup_{i=1}^{n} (\mathcal{C}_i\text{-}hull(\mathbf{P})) \subseteq (\bigwedge_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathbf{P}) \subseteq (\bigcap_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathbf{P})$$

Proof: The proof is by induction on n.

The theorem is vacuously true for $n = 1$.

Let $n = 2$. From theorem 3.3.1(2) we have $\mathrm{P} \subseteq \mathcal{C}_2\text{-}hull(\mathrm{P})$ and, hence, from theorem 3.3.1(4) we have $\mathcal{C}_1\text{-}hull(\mathrm{P}) \subseteq \mathcal{C}_1\text{-}hull(\mathcal{C}_2\text{-}hull(\mathrm{P}))$. Also, from theorem 3.3.1(2) we have $\mathcal{C}_2\text{-}hull(\mathrm{P}) \subseteq \mathcal{C}_1\text{-}hull(\mathcal{C}_2\text{-}hull(\mathrm{P}))$. Hence, $(\mathcal{C}_1\text{-}hull(\mathrm{P}) \cup \mathcal{C}_2\text{-}hull(\mathrm{P})) \subseteq \mathcal{C}_1\text{-}hull(\mathcal{C}_2\text{-}hull(\mathrm{P}))$.

Since $\mathcal{C}_1 \cap \mathcal{C}_2 \subseteq \mathcal{C}_2$ theorem 3.4.1(1) implies $\mathcal{C}_2\text{-}hull(\mathrm{P}) \subseteq (\mathcal{C}_1 \cap \mathcal{C}_2)\text{-}hull(\mathrm{P})$. Hence, $\mathcal{C}_1\text{-}hull(\mathcal{C}_2\text{-}hull(\mathrm{P})) \subseteq \mathcal{C}_1\text{-}hull((\mathcal{C}_1 \cap \mathcal{C}_2)\text{-}hull(\mathrm{P})) = (\mathcal{C}_1 \cap \mathcal{C}_2)\text{-}hull(\mathrm{P})$.

Thus, $(\mathcal{C}_1\text{-}hull(\mathrm{P}) \cup \mathcal{C}_2\text{-}hull(\mathrm{P})) \subseteq \mathcal{C}_1\text{-}hull(\mathcal{C}_2\text{-}hull(\mathrm{P})) \subseteq (\mathcal{C}_1 \cap \mathcal{C}_2)\text{-}hull(\mathrm{P})$. Hence the theorem holds for $n = 2$.

Suppose now that the inductive hypothesis holds for all $3 \leq k < n$. Consider the two innermost operators of the composite operator:

$$(\bigwedge_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathrm{P}) = \mathcal{C}_1\text{-}hull(\ldots \mathcal{C}_{n-1}\text{-}hull(\mathcal{C}_n\text{-}hull(\mathrm{P}))\ldots)$$

From the theorem for $n = 2$ we know that

$$\mathcal{C}_{n-1}\text{-}hull(\mathcal{C}_n\text{-}hull(\mathrm{P})) \subseteq (\mathcal{C}_{n-1} \cap \mathcal{C}_n)\text{-}hull(\mathrm{P})$$

Since there are no more than $n - 2$ operators surrounding these inner operators we can apply the inductive hypothesis to conclude that

$$\begin{aligned}
(\bigwedge_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathrm{P}) &= (\bigwedge_{i=1}^{n-2} \mathcal{C}_i)\text{-}hull(\mathcal{C}_{n-1}\text{-}hull(\mathcal{C}_n\text{-}hull(\mathrm{P}))) \\
&\subseteq (\bigcap_{i=1}^{n-2} \mathcal{C}_i)\text{-}hull(\mathcal{C}_{n-1}\text{-}hull(\mathcal{C}_n\text{-}hull(\mathrm{P})))
\end{aligned}$$

But both $(\bigcap_{i=1}^{n-2} \mathcal{C}_i)\text{-}hull$ and $(\mathcal{C}_{n-1} \cap \mathcal{C}_n)\text{-}hull$ are hull operators (lemma 3.4.1). Therefore from the theorem for $n = 2$ we have that

$$\begin{aligned}
(\bigcap_{i=1}^{n-2} \mathcal{C}_i)\text{-}hull(\mathcal{C}_{n-1}\text{-}hull(\mathcal{C}_n\text{-}hull(\mathrm{P}))) &\subseteq (\bigcap_{i=1}^{n-2} \mathcal{C}_i)\text{-}hull((\mathcal{C}_{n-1} \cap \mathcal{C}_n)\text{-}hull(\mathrm{P})) \\
&\subseteq ((\bigcap_{i=1}^{n-2} \mathcal{C}_i) \cap (\mathcal{C}_{n-1} \cap \mathcal{C}_n))\text{-}hull(\mathrm{P}) \\
&= (\bigcap_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathrm{P})
\end{aligned}$$

Thus

$$(\bigwedge_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathbf{P}) \subseteq (\bigcap_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathbf{P})$$

All that remains is to show that

$$\bigcup_{i=1}^{n} (\mathcal{C}_i\text{-}hull(\mathbf{P})) \subseteq (\bigwedge_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathbf{P})$$

Consider the j^{th} hull operator in the composition:

$$(\bigwedge_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathbf{P}) = \mathcal{C}_1\text{-}hull(\ldots(\mathcal{C}_j\text{-}hull(\ldots(\mathbf{P})\ldots))\ldots)$$

Consider the set on which $\mathcal{C}_j$-*hull* acts. That is, the set $(\ldots(\mathbf{P})\ldots)$ produced by the hull operators inside the inner pair of brackets. This set must contain $\mathbf{P}$ since each hull operator is expansive (theorem 3.3.1(2)).

Thus from theorem 3.3.1(4) we have, for any $1 \leq j \leq n$,

$$\mathcal{C}_j\text{-}hull(\mathbf{P}) \subseteq \ldots(\mathcal{C}_j\text{-}hull(\ldots(\mathbf{P})\ldots))\ldots$$

Thus

$$\bigcup_{i=1}^{n} (\mathcal{C}_i\text{-}hull(\mathbf{P})) \subseteq (\bigwedge_{i=1}^{n} \mathcal{C}_i)\text{-}hull(\mathbf{P})$$

■

Note that this result is true independent of the order in which the hulls are composed in the middle term.

The following example shows that, in general, we cannot replace any of the containments by equality:

Let $S = \{a, b, c, d\}$; $\mathcal{C}_1 = \{\emptyset, S, \{a\}, \{a, b, c\}\}$; $\mathcal{C}_2 = \{\emptyset, S, \{a, b\}\}$; and $\mathcal{C} = \mathcal{C}_1 \cap \mathcal{C}_2 = \{\emptyset, S\}$. Then $\mathcal{C}_1\text{-}hull(\{a\}) \cup \mathcal{C}_2\text{-}hull(\{a\}) = \{a, b\}$; $(\mathcal{C}_1 \wedge \mathcal{C}_2)\text{-}hull(\{a\}) = \{a, b, c\}$; and $\mathcal{C}\text{-}hull(\{a\}) = S$.

However, if the hull operators are *invariant* in the sense defined below then the second two terms are equal (that is, the composition hull is equal to the intersection hull). In chapter 5 we show, with respect to $\mathcal{O}$-convex sets, that all three terms are equal for the special class of connected point sets.

Definition: Given two convexity spaces $(\mathcal{S}, \mathcal{C}_1)$ and $(\mathcal{S}, \mathcal{C}_2)$, $\mathcal{C}_2$ is said to be $\mathcal{C}_1$-*invariant* if $\forall \mathbf{P} \in \mathcal{C}_2$ $\mathcal{C}_1$-$hull(\mathbf{P}) \in \mathcal{C}_2$.

The idea is that $\mathcal{C}_2$ is $\mathcal{C}_1$-invariant if forming the $\mathcal{C}_1$-*hull* of any set which is already $\mathcal{C}_2$-convex does not destroy the $\mathcal{C}_2$-convexity of the set.

THEOREM 3.4.3 **(The (Weak) Decomposition Theorem)** *Let $(\mathcal{S}, \mathcal{C}_1)$ and $(\mathcal{S}, \mathcal{C}_2)$ be two convexity spaces defined on the groundset $\mathcal{S}$.*

If $\mathcal{C}_2$ is $\mathcal{C}_1$-invariant then $\forall \mathbf{P} \subseteq \mathcal{S}$ $(\mathcal{C}_1 \wedge \mathcal{C}_2)$-$hull(\mathbf{P}) = (\mathcal{C}_1 \cap \mathcal{C}_2)$-$hull(\mathbf{P})$

Proof: Let $\mathbf{P}$ be a subset of $\mathcal{S}$. By definition, $\mathcal{C}_2$-$hull(\mathbf{P}) \in \mathcal{C}_2$. If $\mathcal{C}_2$ is $\mathcal{C}_1$-invariant then $\mathcal{C}_1$-$hull(\mathcal{C}_2$-$hull(\mathbf{P})) \in \mathcal{C}_2$. But $\mathcal{C}_1$-$hull(\mathcal{C}_2$-$hull(\mathbf{P})) \in \mathcal{C}_1$. Therefore $\mathcal{C}_1$-$hull(\mathcal{C}_2$-$hull(\mathbf{P})) \in \mathcal{C}_1 \cap \mathcal{C}_2$. But $(\mathcal{C}_1 \cap \mathcal{C}_2)$-$hull(\mathbf{P})$ is a subset of all $(\mathcal{C}_1 \cap \mathcal{C}_2)$-convex sets which contain $\mathbf{P}$. Therefore $(\mathcal{C}_1 \cap \mathcal{C}_2)$-$hull(\mathbf{P}) \subseteq (\mathcal{C}_1 \wedge \mathcal{C}_2)$-$hull(\mathbf{P})$.

But from the refinement theorem we know that $\forall \mathbf{P} \subseteq \mathcal{S}$ $(\mathcal{C}_1 \wedge \mathcal{C}_2)$-$hull(\mathbf{P}) \subseteq (\mathcal{C}_1 \cap \mathcal{C}_2)$-$hull(\mathbf{P})$. Hence the result follows. ∎

The following example shows that the union of two hulls is not necessarily equal to their composition (and their intersection) even if the convexity spaces are invariant.

Let $\mathcal{S} = \{a, b, c, d\}$; $\mathcal{C}_1 = \{\emptyset, \mathcal{S}, \{a, b\}\}$; $\mathcal{C}_2 = \{\emptyset, \mathcal{S}, \{a, c\}\}$; and $\mathcal{C} = \mathcal{C}_1 \cap \mathcal{C}_2 = \{\emptyset, \mathcal{S}\}$. Then $\mathcal{C}_2$ is $\mathcal{C}_1$-invariant but $\mathcal{C}_1$-$hull(\{a\}) \cup \mathcal{C}_2$-$hull(\{a\}) = \{a, b, c\}$; $(\mathcal{C}_1 \wedge \mathcal{C}_2)$-$hull(\{a\}) = \mathcal{S}$ and $\mathcal{C}$-$hull(\{a\}) = \mathcal{S}$.

3.5 Convexity Spaces and Locked Transaction Systems

As an illustration of the generality of the previous results we discuss how they apply to a problem on locked transactions in a database.

In attempting to solve a database concurrency problem Yannakakis *et al.* ([66]) found a correspondence between the safety of a locked transaction system and what was later to be called the *NESW-closure* of a collection of orthogonal rectangles. Lipski and Papadimitriou ([36]) then found an $O(n \lg n \lg\lg n)$ time $O(n \lg n)$ space algorithm to find the *NESW-closure* of a set of orthogonal rectangles. Later, Soisalon-Soininen and Wood ([57]) proved that the *NESW-closure* can be decomposed into the composition of two simpler closures. With this decomposition result in hand they were able to derive a simple—and optimal—$O(n \lg n)$ time and $O(n)$ space algorithm to find the *NESW-closure*. We reprove their decomposition result for connected sets using the results proved in the previous sections.

Definition: Let $p = (x_1, y_1)$ and $q = (x_2, y_2)$ be two points in the plane. p and q are said to be *comparable* if p and q lie on a line of positive slope (that is, if $(x_1 - x_2)(y_1 - y_2) > 0$). Similarly, p and q are said to be *incomparable* if p and q lie on a line of negative slope.

Definition: Let $p = (x_1, y_1)$ and $q = (x_2, y_2)$ be two incomparable points and assume that p is to the left of q. The point (x_2, y_1) $((x_1, y_2))$ is said to be the *NE-conjugate* (*SW-conjugate*) of p and q. If p and q are comparable we have the NW- and SE-conjugates respectively.

Definition: $\mathbf{P}$ is said to be *NE-closed* if for every incomparable pair of points p and q in $\mathbf{P}$, *NE-conjugate*$(p, q) \in \mathbf{P}$. Similarly, we define NW-, SE- and SW-closed sets. If $\mathbf{P}$ is both NE-closed and SW-closed $\mathbf{P}$ is said to be $NESW$-closed and similarly we obtain $NWSE$-closed sets. If $\mathbf{P}$ is both $NESW$-closed and $NWSE$-closed $\mathbf{P}$ is said to be R-closed.

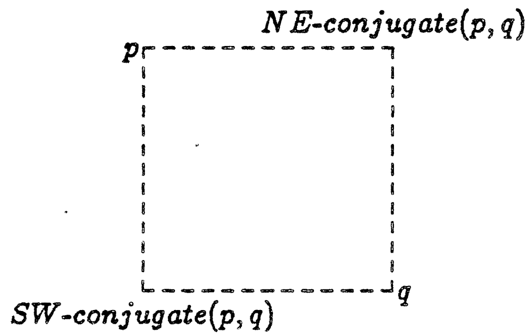

Figure 3.1: NE- and SW-Conjugates

This is a relaxation of the definitions used by Soisalon-Soininen and Wood since they required that the two points be *connected* in the given set.

Definition: For any $X \in \{NE, SW, NW, SE, NESW, NWSE, R\}$ $\mathbf{Q}$ is said to be the ***X-closure*** of $\mathbf{P}$ if $\mathbf{Q}$ is the smallest X-closed region containing $\mathbf{P}$.

It is easy to prove that, for any $X \in \{NE, SW, NW, SE, NESW, NWSE, R\}$, ***X-closure*** is a hull operator. That is, the set of all X-closed sets forms a convexity space. From this it follows that all of the results given in this chapter apply to X-closed sets.

The fundamental decomposition result Soisalon-Soininen and Wood needed was the following theorem (compare the Orientation Decomposition Theorem—theorem 5.4.1):

THEOREM **3.5.1 (The NESW Decomposition Theorem)** $\forall \mathbf{P}$

$$NE\text{-}closure(SW\text{-}closure(\mathbf{P})) = NW\text{-}closure(SE\text{-}closure(\mathbf{P})) = NESW\text{-}closure(\mathbf{P})$$

Soisalon-Soininen and Wood proved their results by examining the boundaries of the various closures. We re-prove their results for the special case of connected regions by showing that NE-closure is SW-invariant (invariance of NW-closure with respect to SE-closure is similar). The above decomposition theorem restricted to connected regions then follows from theorem 3.4.3.

In theorem 3.5.2 we show that ***NE-closure*** is SW-invariant (the other invariance proof is similar). First however we need the following simple lemma.

Lemma 3.5.1 *If $p \in SW\text{-}closure(\mathbf{P}) \setminus \mathbf{P}$ then the vertical line through p intersects $\mathbf{P}$ above p and the horizontal line through p intersects $\mathbf{P}$ to the right of p.*

Proof: If $p \in SW\text{-}closure(\mathbf{P}) \setminus \mathbf{P}$ and the vertical line through p does not intersect $\mathbf{P}$ above p then we may delete p and all points in $SW\text{-}closure(\mathbf{P})$ above p and still have a SW-closed set. Since no points in $\mathbf{P}$ were deleted this set still contains $\mathbf{P}$. But $SW\text{-}closure(\mathbf{P})$ is the smallest set which both contains $\mathbf{P}$ and is SW-closed, a contradiction. Thus the vertical line through p must intersect $\mathbf{P}$ above p. A similar argument applies to the other condition. ∎

THEOREM **3.5.2** *NE-closure is SW-invariant.*

Proof: We show that if $\mathbf{P}$ is NE-closed then

$$\forall p, q \in SW\text{-}closure(\mathbf{P}) \;, \quad NE\text{-}conjugate(p,q) \in SW\text{-}closure(\mathbf{P}).$$

This establishes that if $\mathbf{P}$ is NE-closed then $SW\text{-}closure(\mathbf{P})$ is NE-closed.

Let p and q be in $SW\text{-}closure(\mathbf{P})$ such that p and q are incomparable. If no such points exist we are done. Assume, without loss of generality, that p is north-west of q (see figure 3.2).

Case 1: $p, q \in \mathbf{P}$. In this case $NE\text{-}conjugate(p,q) \in \mathbf{P}$ since $\mathbf{P}$ is NE-closed.

Case 2: $p \in \mathbf{P}$, $q \notin \mathbf{P}$. Since $q \notin \mathbf{P}$, from lemma 3.5.1 we know that the horizontal line through q must intersect $\mathbf{P}$ at some point α to the right of q and the vertical line through q must intersect $\mathbf{P}$ at some point β above q.

Case 2a: β is on the horizontal line through p. In this case $NE\text{-}conjugate(p,q) = \beta \in \mathbf{P}$.

Case 2b: β is below the horizontal line through p. In this case $NE\text{-}conjugate(p,q) = NE\text{-}conjugate(p,\beta) \in \mathbf{P}$.

Case 2c: β is above the horizontal line through p. In this case $NE\text{-}conjugate(p,q) = SW\text{-}conjugate(NE\text{-}conjugate(p,\alpha),\beta) \in SW\text{-}closure(\mathbf{P})$.

Case 3: $p \notin \mathbf{P}$, $q \in \mathbf{P}$. This case is a mirror reflection of case 2.

Case 4: $p, q \notin \mathbf{P}$. Since $p, q \notin \mathbf{P}$, from lemma 3.5.1 we know that the horizontal line through p must intersect $\mathbf{P}$ at some point α to the right of p and the vertical line through q must intersect $\mathbf{P}$ at some point β above q.

Case 4a: α is left of q and β is below p. In this case

$$NE\text{-}conjugate(p,q) = NE\text{-}conjugate(\alpha,\beta) \in \mathbf{P}.$$

Case 4b: α is right of q and β is above p. In this case

$$NE\text{-}conjugate(p,q) = SW\text{-}conjugate(\beta,\alpha) \in SW\text{-}closure(\mathbf{P}).$$

Case 4c: α is left of q and β is above p. Again from lemma 3.5.1 we know that the horizontal line through q must intersect $\mathbf{P}$ at some point γ to the right of q. In this case

$$NE\text{-}conjugate(p,q) = SW\text{-}conjugate(\beta, NE\text{-}conjugate(\alpha,\gamma)) \in SW\text{-}closure(\mathbf{P}).$$

Case 4d: α is right of q and β is below p. This case is a mirror reflection of case 4c.

Hence in all cases if $\mathbf{P}$ is NE-closed $SW\text{-}closure(\mathbf{P})$ is NE-closed. ∎

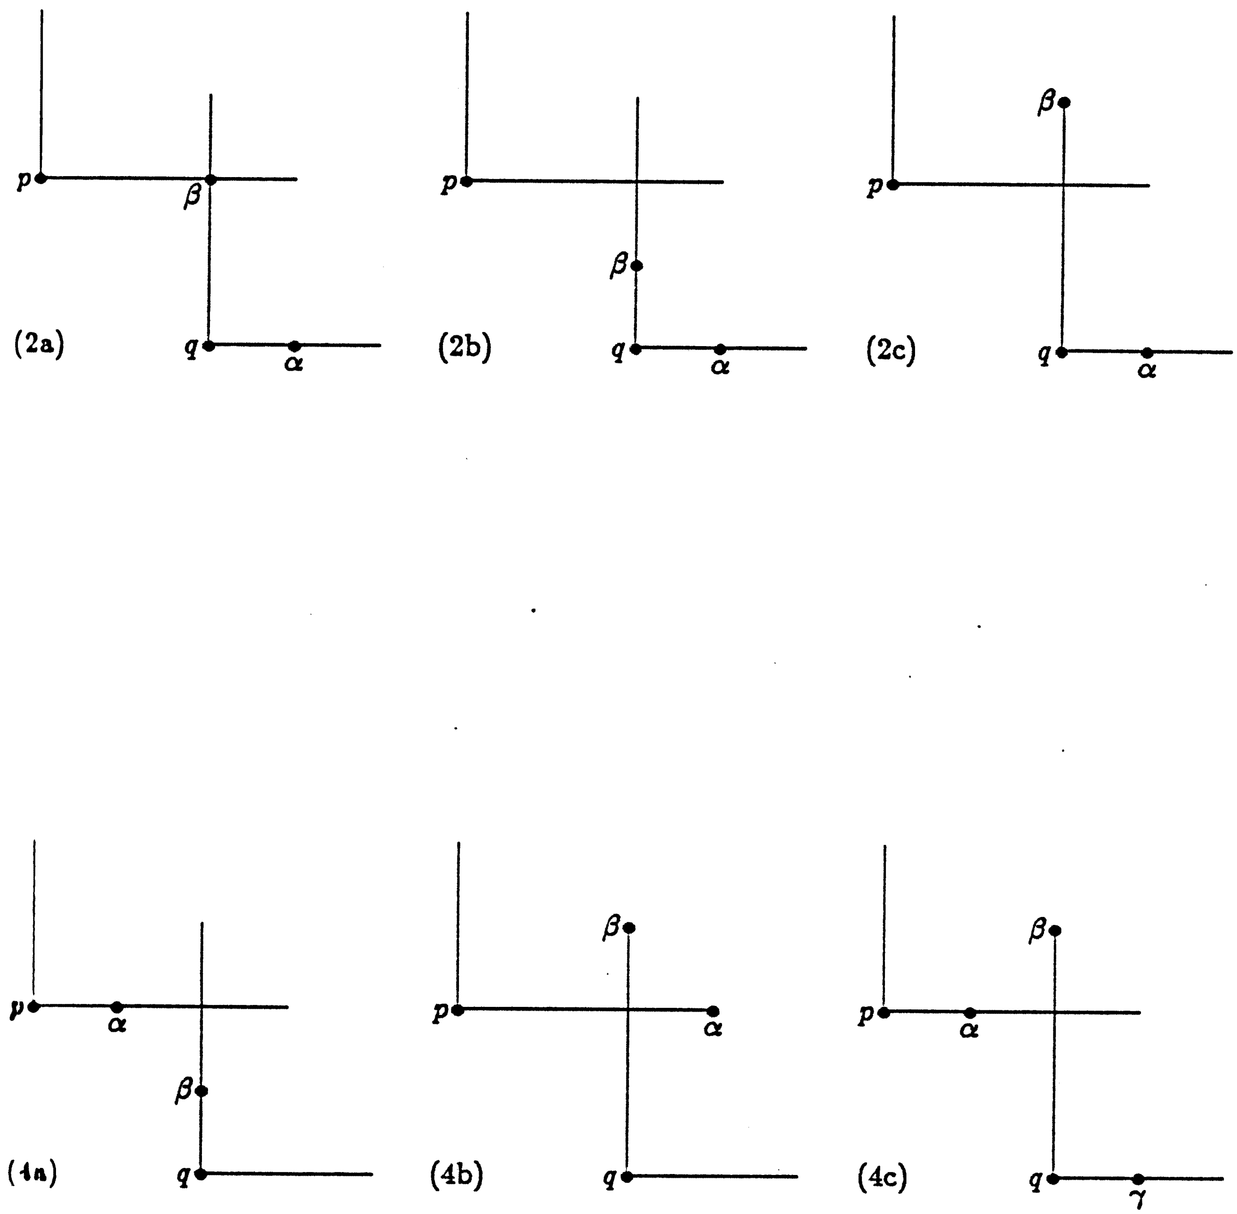

Figure 3.2: The Various Possible Cases

Chapter 4

Stairlines

4.1 Lines and Staircases

In the study of orthogonally convex sets a *staircase* ([64]) is defined as a polygonal curve made up of orthogonal line segments. The line segments are connected such that they are alternately horizontal and vertical. As we shall see in chapter 5, with respect to orthogonally convex sets, staircases play the same role as *lines* do for convex sets. In this chapter we investigate a class of curves which generalize them both. We call these curves "$\mathcal{O}$-stairlines," deriving the term as a portmanteau of (orthogonal) staircases and (straight) lines. We present a simple characterization of $\mathcal{O}$-stairlines in terms of the orientations they cover (their *spans*) and we completely characterize the union of the set of all $\mathcal{O}$-stairline segments between any two points (their *$\mathcal{O}$-parallelogram*). These characterizations will be of fundamental importance in the theory of $\mathcal{O}$-convex sets developed in chapters 5 and 6. We then briefly examine one possible way to define the notion of *parallelism* for $\mathcal{O}$-stairlines and sketch algorithms to find the intersection point(s) of two parallel or non-parallel polygonal $\mathcal{O}$-stairlines.

Recall that a *convex* set is a set whose intersection with any line is connected; an *orthogonally convex* set is a set whose intersection with any horizontal or vertical line connected; an $\mathcal{O}$-*convex* set is a set whose intersection with any $\mathcal{O}$-line is connected.

4.2 $\mathcal{O}$-Stairlines

Definition: A curve S is said to be an $\mathcal{O}$-*stairline* if S is $\mathcal{O}$-convex.

By analogy with lines we also use the terms "$\mathcal{O}$-stairsegment" and "$\mathcal{O}$-stairray" with the obvious meanings. We derive an equivalent but more convenient definition of an $\mathcal{O}$-stairline in lemma 4.2.1 below.

Definition: A set $\mathbf{P}$ is said to have *span* $\langle\theta_1, \theta_2\rangle$ (where $\theta_1 \le \theta_2$) if $\forall p, q \in \mathbf{P}$, where $p \neq q$, $\Theta(\mathcal{L}[p,q]) \in \langle\theta_1, \theta_2\rangle$ and this is true for no smaller range. The span of the empty set or an isolated point is empty.

At this point the reader should recall that 0° and 180° are identified. Thus, any span which crosses the horizontal is still connected *even though it is described as the union of two (apparently) disjoint ranges.*

Observe that $\mathbf{P}$ has a single orientation in its span if and only if $\mathbf{P}$ is a subset of a line. Also note that the span of a set is not necessarily a closed range, it may be open or half-open.

If $\mathbf{P}$'s span is $\mathcal{O}$-free then there are no two distinct points in $\mathbf{P}$ lying on an $\mathcal{O}$-line and thus $\mathbf{P}$ is $\mathcal{O}$-convex. The converse is false in general as $\Re^2$ is always $\mathcal{O}$-convex but its span is never $\mathcal{O}$-free (unless $\mathcal{O}$ is empty). However, as we show, if $\mathbf{P}$ is a *curve* then a slightly weaker converse holds.

Lemma 4.2.1 *A curve with span $\langle\theta_1, \theta_2\rangle$ is an $\mathcal{O}$-stairline if and only if the open range (θ_1, θ_2) is $\mathcal{O}$-free.*

Proof: Suppose that S is a curve with span $[\theta_1, \theta_2]$ (the other cases are handled similarly).

If S is $\mathcal{O}$-convex then the open range (θ_1, θ_2) must be $\mathcal{O}$-free, hence conversely, suppose that the open range (θ_1, θ_2) is $\mathcal{O}$-free.

If $\mathcal{O}$ is empty then every set is $\mathcal{O}$-convex and, in particular, S is $\mathcal{O}$-convex, so suppose that $\mathcal{O}$ is non-empty. If $\theta_1 = \theta_2$ then S is a line, line segment or ray and hence is $\mathcal{O}$-convex from lemma 3.2.1. Suppose then that $\theta_1 \neq \theta_2$. Without loss of generality assume that $[\theta_1, \theta_2] = [0°, 90°]$. (Recall that the orthogonal transformation given in chapter 2 allows us to assume this and it also preserves incidence.)

Suppose that there exists an $\mathcal{O}$-line $\mathcal{L}$ which cuts S at two distinct points p and q. Since $(0°, 90°)$ is $\mathcal{O}$-free, $\mathcal{L}$ can only be horizontal or vertical (S's span is $[0°, 90°]$). Suppose that $\mathcal{L}$ is horizontal and that p is to the left of q. Consider any point r on S in between p and q. Suppose that $r \notin \mathcal{LS}(p, q)$ and, without loss of generality, that r is above $\mathcal{LS}(p, q)$.

Either $\mathcal{L}[q, r]$ is vertical or r is to the right of q otherwise $\Theta(\mathcal{L}[q, r]) \notin [0°, 90°]$. Consider a line $\mathcal{L}_1$ through q of orientation θ where $\theta > 90°$ (see figure 4.1). Since S is connected, p and r are connected and so $\mathcal{L}_1$ must cut S at some point other than q. But this means that there exist two distinct points in S whose line segment has slope outside of S's span, a contradiction. Hence r cannot be above (or similarly, below) $\mathcal{LS}(p, q)$. Therefore $r \in \mathcal{LS}(p, q)$ for all r in S in between p and q. That is, between p and q, S is a line segment. Hence, even if θ_1 (or θ_2) is an orientation in $\mathcal{O}$ then S is $\mathcal{O}$-convex. ∎

Observe that if $\theta_1 = \theta_2$ then (θ_1, θ_2) is empty and hence is (vacuously) $\mathcal{O}$-free and so any line (line segment or ray) is an $\mathcal{O}$-stairline ($\mathcal{O}$-stairsegment or $\mathcal{O}$-stairray).

Definition: If an $\mathcal{O}$-convex curve divides the plane into two halfspaces the two halfspaces are said to be *$\mathcal{O}$-stairhalfplanes*.

The following lemma is straightforward:

Lemma 4.2.2 *If* $\mathbf{P}$ *is an $\mathcal{O}$-stairhalfplane then* $\mathbf{P}$ *is $\mathcal{O}$-convex.*

Observe that any closed curve divides the plane into two $\emptyset$-stairhalfplanes.

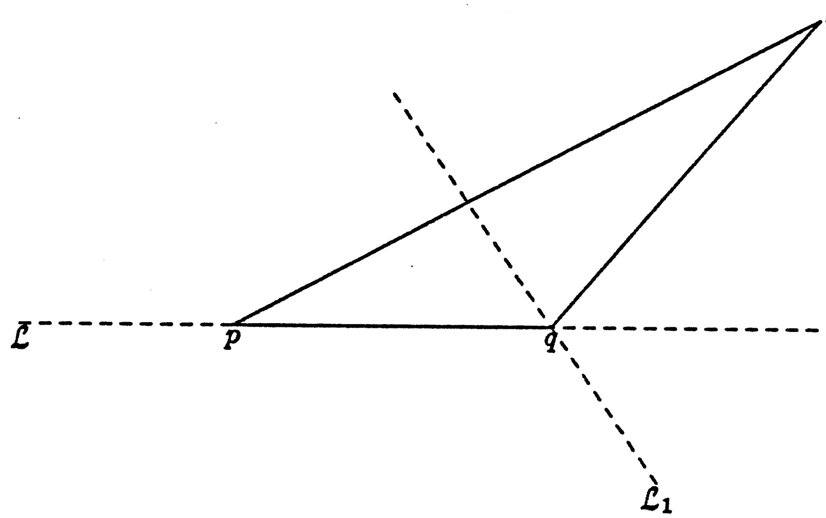

Figure 4.1: r must be on $\mathcal{LS}[p,q]$

Definition: An $\mathcal{O}$-convex polygonal curve is said to be a *polygonal* $\mathcal{O}$-stairline. If S is a polygonal $\mathcal{O}$-stairline made up of n lines, line segments or rays then S is said to be of *length* n.

It is easy to show that if a polygonal curve made up of lines, line segments or rays $\mathcal{L}_1, \mathcal{L}_2, \ldots, \mathcal{L}_n$ $(n \geq 1)$ forms a polygonal $\mathcal{O}$-stairline, $\mathcal{O}$-stairsegment or $\mathcal{O}$-stairray with span $\langle\theta_1, \theta_2\rangle$ then

1. It is monotone with respect to the orientations $\theta_1 + 90°$ and $\theta_2 + 90°$.[1]
2. $\forall 1 \leq i \leq n\ , \quad \Theta(\mathcal{L}_i) \in \langle\theta_1, \theta_2\rangle$.

Definition: An $\mathcal{O}$-oriented polygonal $\mathcal{O}$-stairsegment is called an *$\mathcal{O}$-staircase*.

[1] Recall that all orientations are taken modulo 180°.

See figure 4.2 for examples of an $\mathcal{O}$-stairsegment, a polygonal $\mathcal{O}$-stairsegment, and an $\mathcal{O}$-staircase for $\mathcal{O} = \{0°, 90°\}$.

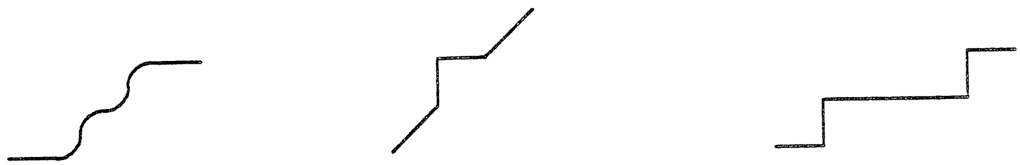

Figure 4.2: Example $\mathcal{O}$-Stairsegments

4.3 $\mathcal{O}$-Parallelograms

We now wish to characterize the set of *all* $\mathcal{O}$-stairsegments joining any two points in the plane.

Definition: We call the union of all $\mathcal{O}$-stairsegments joining p and q the ***$\mathcal{O}$-parallelogram*** of p and q and write $\mathcal{O}$-$\|[p,q]$.

Observe that $\mathcal{LS}[p,q]$ is always an $\mathcal{O}$-stairsegment and thus $\mathcal{LS}[p,q] \subseteq \mathcal{O}$-$\|[p,q]$ for all $\mathcal{O}$.

See figure 4.3 for examples of $\mathcal{O}$-$\|[p,q]$ for $\mathcal{O} = \{0°, 90°\}$.

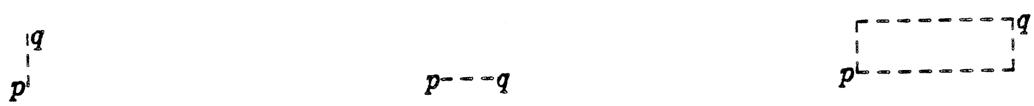

Figure 4.3: Example $\mathcal{O}$-Parallelograms

Our next three results show the term "parallelogram" to be a misnomer! A parallelogram is only the non-degenerate case; there are three possible degeneracies.

Lemma 4.3.1 *$\forall p, q$ if $\Theta(\mathcal{LS}[p,q]) \in \mathcal{O}$ then $\mathcal{O}$-$\|[p,q] = \mathcal{LS}[p,q]$.*

Proof: Any curve that connects p and q which departs from $\mathcal{LS}[p,q]$ cannot be $\Theta(\mathcal{LS}[p,q])$-convex. The proof of this claim is similar to the proof of lemma 4.2.1 and is omitted. ∎

Corollary 4.3.1 *If* $\mathcal{O} = [0°, 180°)$ *then,* $\forall p, q$ $\mathcal{O}\text{-}\|[p,q] = \mathcal{LS}[p,q]$.

For the next result recall that $\mathcal{H}$ denotes a halfplane; $\mathcal{H}[\mathcal{L}, p]$ denotes the closed halfplane with boundary line $\mathcal{L}$ containing the point p; and $\mathcal{L}[\theta, p]$ denotes the line through p of orientation θ.

THEOREM **4.3.1 ($\mathcal{O}$-Parallelogram Characterization)** $\forall p, q$

1. *If* $\mathcal{O} = \emptyset$ *then* $\mathcal{O}\text{-}\|[p,q] = \Re^2$
2. *If* $\mathcal{O} = \theta$ *then* $\mathcal{O}\text{-}\|[p,q] = \mathcal{H}[\mathcal{L}[\theta,p],q] \;\cap\; \mathcal{H}[\mathcal{L}[\theta,q],p]$
3. *If* $|\mathcal{O}| \geq 2$ *then*

$$\mathcal{O}\text{-}\|[p,q] = \mathcal{H}[\mathcal{L}[\theta_1,p],q] \;\cap\; \mathcal{H}[\mathcal{L}[\theta_1,q],p] \;\cap\; \mathcal{H}[\mathcal{L}[\theta_2,p],q] \;\cap\; \mathcal{H}[\mathcal{L}[\theta_2,q],p]$$

(where (θ_1, θ_2) *is* $\mathcal{L}[p,q]$*'s maximal* $\mathcal{O}$*-free range).*

Proof: If $\Theta(\mathcal{L}[p,q]) \in \mathcal{O}$ then $\mathcal{O}\text{-}\|[p,q] = \mathcal{LS}[p,q]$ and the theorem follows. Suppose then that $\Theta(\mathcal{L}[p,q]) \notin \mathcal{O}$.

Case 1: $|\mathcal{O}| = 0$. If $\mathcal{O}$ is empty then every range is $\mathcal{O}$-free and so *any* curve connecting p and q is an $\mathcal{O}$-stairsegment, hence $\emptyset\text{-}\|[p,q] = \Re^2$.

Case 2: $|\mathcal{O}| = 1$. If $\mathcal{O} = \theta$ then any curve connecting p and q which leaves the infinite slab bounded by the two θ-lines passing through p and q cannot be θ-convex. Further, if r is any point in the infinite slab then the curve $\mathcal{LS}[p,r] \cup \mathcal{LS}[r,q]$ is θ-convex and so is a θ-stairsegment. Thus, any point in the infinite slab is in $\theta\text{-}\|[p,q]$.

Case 3: $|\mathcal{O}| \geq 2$. Without loss of generality, let $(0°, 90°)$ be $\mathcal{L}[p,q]$'s maximal $\mathcal{O}$-free range. If any curve connecting p and q leaves the orthogonal rectangle with diagonal endpoints p and q then it can only be monotone in either the horizontal or vertical direction but not both and so cannot be an $\mathcal{O}$-stairsegment. Hence when $\mathcal{O}$ contains two or more orientations then all $\mathcal{O}$-stairsegments must lie within the parallelogram with diagonal endpoints p and q.

If r is any point in the orthogonal rectangle then, as in case 2, the curve $\mathcal{LS}[p,r] \cup \mathcal{LS}[r,q]$ is an $\mathcal{O}$-stairsegment connecting p and q passing through r. In fact we can just as easily construct an $\mathcal{O}$-*staircase* connecting p and q passing through r (see figure 4.4 for a simple example $\mathcal{O}$-staircase for $\mathcal{O} = \{0°, 90°\}$). ■

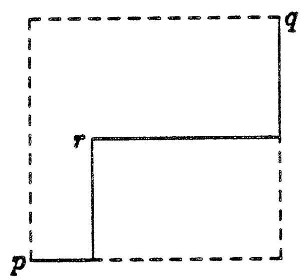

Figure 4.4: Any Point in $\mathcal{O}$-$\|[p,q]$ lies on an $\mathcal{O}$-Stairsegment Connecting p and q

Definition: If $|\mathcal{O}| \geq 2$ and $\Theta(\mathcal{L}[p,q]) \notin \mathcal{O}$ then the two extremal $\mathcal{O}$-staircases bounding $\mathcal{O}$-$\|[p,q]$ are said to be the *arms* of $\mathcal{O}$-$\|[p,q]$.

For example, in figure 4.4 the two dashed line segments connecting p, *NW-conjugate*(p,q), and q form the *upper arm* of $\mathcal{O}$-$\|[p,q]$. Similarly, the *lower arm* is the length 2 $\mathcal{O}$-staircase joining p, *SE-conjugate*(p,q), and q.

Remark: As can be seen in the above theorem, the special case $|\mathcal{O}| \leq 1$ gives odd results, but proof techniques for $|\mathcal{O}| \leq 1$ are similar to those for $|\mathcal{O}| \geq 2$.

In the remainder of this thesis we will often only discuss the case $|\mathcal{O}| \geq 2$ since results for the two special cases can be proved by similar techniques to those given for the general case.

4.4 The Analogy Between Lines and $\mathcal{O}$-Stairlines

With respect to $\mathcal{O}$-convex sets, $\mathcal{O}$-stairlines are perhaps the most natural analogues of straight lines with respect to convex sets, in that: the intersection of an $\mathcal{O}$-line and an $\mathcal{O}$-stairline is connected (as we saw in the proof of lemma 4.2.1), and, two $\mathcal{O}$-stairlines with disjoint spans intersect in at most one point.

Unfortunately, the intersection of two $\mathcal{O}$-stairlines whose spans are non-disjoint is either empty, connected or disconnected—unlike the simpler case for straight lines. Further, $\mathcal{O}$-stairlines can be non-intersecting without being parallel (in the Euclidean sense). Finally, two points no longer necessarily define a unique $\mathcal{O}$-stairline—so, in general, this is not an incidence geometry.

Ideally, we would like to say that two $\mathcal{O}$-stairlines are non-parallel if their spans are disjoint, for then their intersection is guaranteed to be connected. However, there is another condition which guarantees a connected intersection and so we use the following scheme:

Definition: Two $\mathcal{O}$-stairlines are said to be *non-parallel* if they intersect and either

1. the spans of the two $\mathcal{O}$-stairlines are disjoint or
2. at least one of the the $\mathcal{O}$-stairlines is an $\mathcal{O}$-line.

If neither condition holds the two $\mathcal{O}$-stairlines are said to be *parallel.*

If a collection of $\mathcal{O}$-stairlines are pairwise non-parallel then they have most of the properties of lines that are relevant to convex sets (for example, they intersect exactly once, are 1-separable continua etc.).

Note that under this definition of parallelism Euclid's fifth postulate does not hold (unless $\mathcal{O}$ is the set of all orientations); lines, having only one orientation in their spans, must always have a connected intersection; and finally when $\mathcal{O}$ is the set of all orientations then an $\mathcal{O}$-stairline is a line and two lines are parallel if and only if they do not intersect.

4.5 Computing $\mathcal{O}$-Stairline Intersections

We now sketch an algorithm to find the intersection of two polygonal $\mathcal{O}$-stairlines assuming that we are told whether or not they are parallel (in the sense defined in the previous section). For simplicity of description assume that two parallel $\mathcal{O}$-stairlines share the common span $[0°, 90°]$ and two non-parallel $\mathcal{O}$-stairlines have disjoint spans $[0°, 90°]$ and $(90°, 180°)$. We assume that the $\mathcal{O}$-stairlines are given as a sorted list of their vertices and that we can randomly access this list.

4.5.1 Non-Parallel Polygonal $\mathcal{O}$-Stairlines

First observe that it is sufficient to solve this problem for doubly infinite $\mathcal{O}$-stairlines (that is $\mathcal{O}$-stairlines that go to infinity at both ends) since if one or both are not doubly infinite then we adjoin up to four rays and find their common intersection. After finding the intersection we can decide, in constant time, whether the intersection was caused by the adjoined rays.

If either of the $\mathcal{O}$-stairlines are lines then we perform a binary search to find their intersection point. If both are lines then we find their intersection point in constant time. Thus we have reduced the problem to:

Problem: Given two doubly infinite $\mathcal{O}$-stairlines, over disjoint spans, and each with two or more turns, determine their common intersection.

Let the $\mathcal{O}$-stairlines be S_1 and S_2. Assume that S_1's span is $[0°, 90°]$ and that S_2's span is $(90°, 180°)$. Let their vertices be $p_1, p_2, \ldots, p_n$ and $q_1, q_2, \ldots, q_m$ respectively. Assume that p_1 is the lowest vertex in S_1 and that q_1 is the highest vertex in S_2.

Consider $\mathcal{O}$-$\|[p_{\lfloor n/2 \rfloor}, q_{\lfloor m/2 \rfloor}]$. There are four possible cases and in each case we: find the intersection; discard half of the vertices of S_1; discard half of the vertices of S_2; or discard half of the vertices of both.

1. $p_{\lfloor n/2 \rfloor}$ and $q_{\lfloor m/2 \rfloor}$ share neither an x nor a y coordinate. Suppose that $p_{\lfloor n/2 \rfloor}$ is below and to the left of $q_{\lfloor m/2 \rfloor}$ (the other three possibilities are similar). In this case we

discard the lower half of S_1 since it cannot contain the intersection point (see figure 4.5).

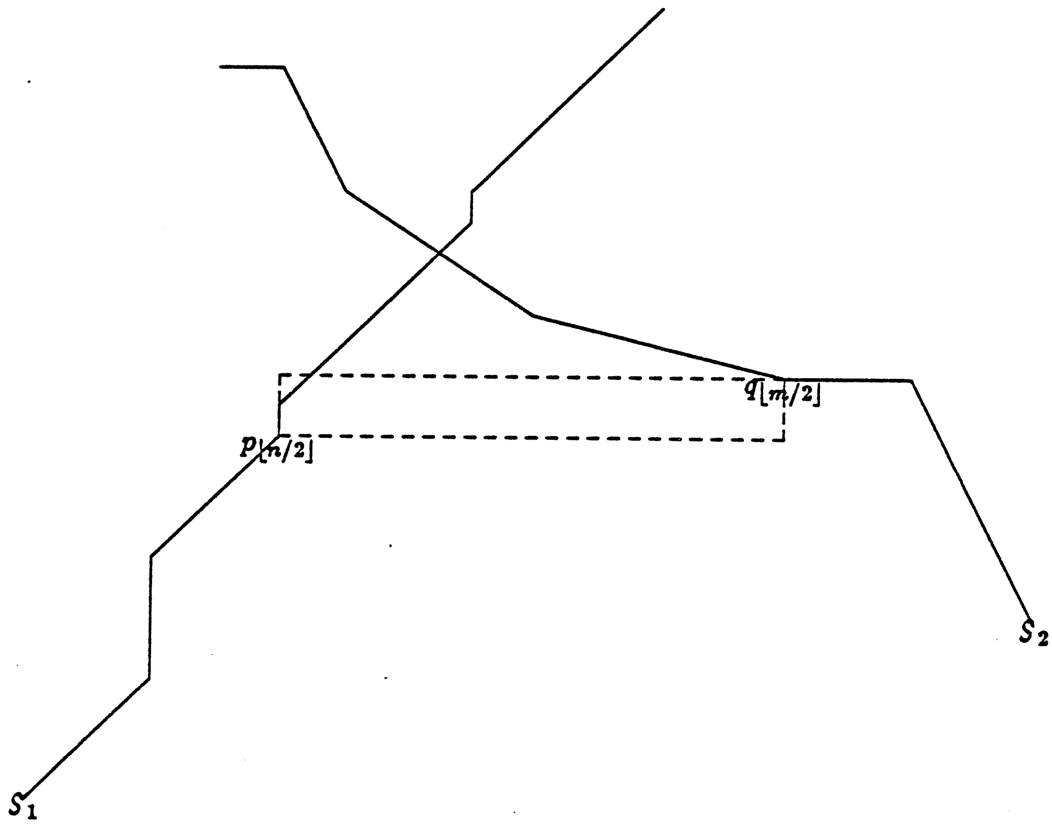

Figure 4.5: Illustrating Case 1

2. $p_{\lfloor n/2 \rfloor}$ and $q_{\lfloor m/2 \rfloor}$ share both x and y coordinates. Halt; intersection found.

3. $p_{\lfloor n/2 \rfloor}$ and $q_{\lfloor m/2 \rfloor}$ share an x but not a y coordinate. In this case we discard half of S_1 *and* half of S_2.

4. $p_{\lfloor n/2 \rfloor}$ and $q_{\lfloor m/2 \rfloor}$ share a y but not an x coordinate. This case is similar to the last.

If after discarding some number of vertices we have not found the intersection point and we have reduced S_1 or S_2 to one edge then we perform a binary search for the single edge in the remaining list of vertices.

The recurrence governing the number of comparisons that must be done in the worst case is:

$$f(n,m) \le \begin{cases} \lg n & m = 1 \\ \lg m & n = 1 \\ \max\{f(n, \lfloor m/2 \rfloor), f(\lfloor n/2 \rfloor, m)\} + c & \text{otherwise} \end{cases}$$

This[2] implies that

$$f(n,m) = O(\lg n + \lg m) = O(\max\{\lg n, \lg m\})$$

The $\mathcal{O}$-stairline intersection finding algorithm sketched above can be seen as a generalization of the $O(\lg n)$ algorithm to discriminate a point against a monotone chain given in [47] pg. 49. (Discriminating a point against a curve means determining which side of the curve the point lies on.)

We can establish a lower bound of $\Omega(\max\{\lg n, \lg m\})$ on the number of comparisons needed to find the intersection point by reducing the problem of searching for an element in a list of either n or m elements (whichever is larger) to finding the intersection of a line and a polygonal $\mathcal{O}$-stairline. Hence the above algorithm is optimal in the worst case in the algebraic decision tree model.

4.5.2 Parallel Polygonal $\mathcal{O}$-Stairlines

When S_1 and S_2 are parallel $\mathcal{O}$-stairlines then there can be $O(\min\{n,m\})$ intersection points (for example, see figure 4.6). Thus any algorithm to solve this problem cannot do better than a linear number of comparisons in the worst case. It is simple to find the intersection points in linear time by stepping along the vertices of S_1 and S_2 in sequence that is, perform essentially a linear merge of the two sorted lists.

[2]This is a well-known recurrence, it is the dual of a recursive algorithm to find the boundary of the convex hull of n points (see [30]).

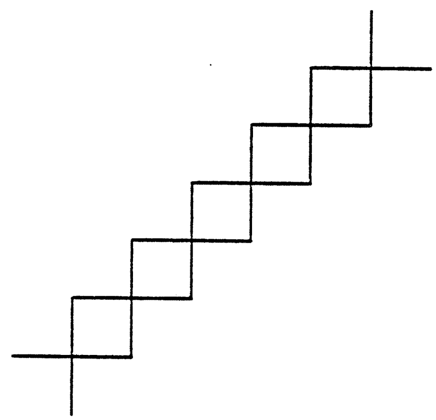

Figure 4.6: Parallel $\mathcal{O}$-Stairlines Can Have $O(n)$ Intersections

Chapter 5

Restricted Orientation Convexity

In this chapter we see how the notions of convexity spaces and $\mathcal{O}$-stairlines can be com bined to construct a theory of $\mathcal{O}$-convex sets. $\mathcal{O}$-convexity serves as a useful vantage poin to survey and unify scattered results and observations in the literature of computationa geometry. In this chapter we show that many of the more important properties of conve sets in the plane are special cases of the more general properties of $\mathcal{O}$-convex sets whe restricted only to *connected sets.* Thus $\mathcal{O}$-convexity is a reasonable generalization of bot normal convexity and orthogonal convexity. Further, we show that the $\mathcal{O}$-convex hull (de fined in an analogous manner to the hull defined in chapter 3) obeys a strong *Decompositio Theorem* (theorem 5.4.1) when restricted to connected sets.

5.1 Properties of Convex Sets

To provide an appropriate generalization of convex sets we need to know what properti of convex sets distinguish them from other sets. What makes them useful, elegant, important? We list below some of the most salient properties of planar convex sets (s Grünbaum [20]), where P is a planar convex set:

1. P is simply connected.

2. The intersection of $\mathbf{P}$ and any line is connected.

3. $\mathbf{P}$ is the intersection of all convex sets which contain it.

4. If $p \notin \mathbf{P}$ then there exists a line separating p and $\mathbf{P}$.

5. $\mathbf{P}$ is the intersection of all halfplanes which contain it.

6. If $p, q \in \mathbf{P}$ then $\mathcal{LS}[p,q] \subseteq \mathbf{P}$.

Except for property (1), all of these are defining characteristics of convex sets. In the concluding section of this chapter we list analogous properties of $\mathcal{O}$-convex sets which include these as special cases.

5.2 $\mathcal{O}$-Convex Sets

Figure 5.1 contains some example figures which are $\mathcal{O}$-convex for various $\mathcal{O}$. Figure 5.1 (a) is not $\mathcal{O}$-convex for any non-empty $\mathcal{O}$, but is $\mathcal{O}$-convex if $\mathcal{O} = \emptyset$, as are all the other figures. Figures (b) and (c) are convex with respect to any horizontal line, as are (d), (e) and (f), so they are all $0°$-convex besides being $\emptyset$-convex. Note that (b) and (c) are not convex for any other orientation. Figures (d), (e) and (f) are convex with respect to any vertical line as well and so they are also $\{0°, 90°\}$-convex. Note that (d) is not convex for any other orientation. Figures (e) and (f) are convex with respect to any line with orientation in the range $[90°, 180°]$ and so they are also $[90°, 180°]$-convex. Note that (e) is not convex for any other orientation. Figure (f) is $\mathcal{O}$-convex for any $\mathcal{O}$.

In lemma 3.2.1 we showed that every convex set is $\mathcal{O}$-convex. In the following lemma we completely characterize convex sets as a subclass of $\mathcal{O}$-convex sets:

Lemma 5.2.1 $\forall \mathbf{P}$, $\mathbf{P}$ *is convex if and only if* $\mathbf{P}$ *is* $\mathcal{O}$*-convex, for* $\mathcal{O} = [0°, 180°)$. *Moreover, this is not true for any* $\mathcal{O} \subset [0°, 180°)$.

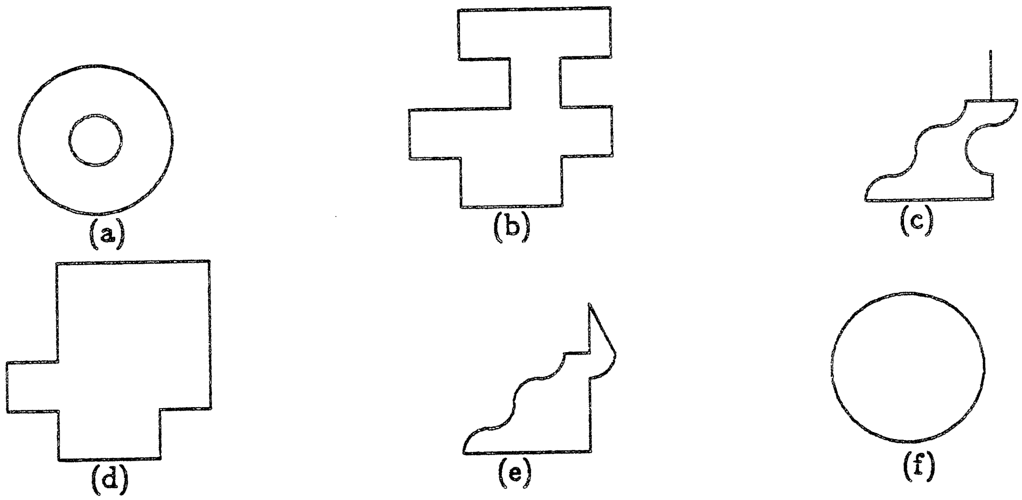

Figure 5.1: Some Examples of $\mathcal{O}$-convex Sets

Proof: From property 2 above **P** is convex if and only if **P** is $\mathcal{O}$-convex, for $\mathcal{O} = [0°, 180°)$. It is easy to construct examples to show that this is true for no smaller set of orientations. For example, if we delete just *one* orientation (say θ_1) then any set consisting of just two distinct points on a θ_1-line is θ-convex for all $\theta \neq \theta_1$ but is, of course, not convex. Indeed, such examples demonstrate that "for all **P**, **P** is connected if **P** is $\mathcal{O}$-convex" holds if and only if $\mathcal{O} = [0°, 180°)$. ∎

Definition: The intersection of all $\mathcal{O}$-convex sets containing **P** is said to be the *$\mathcal{O}$-hull* of **P** and is denoted *$\mathcal{O}$-hull*(**P**) (cf. the definition of the *C-hull* in chapter 3).

Observe that if $\mathcal{O}$ or **P** is empty then *$\mathcal{O}$-hull*(**P**) = **P**. When $\mathcal{O} = \theta$ and **P** is a polygon then the *$\mathcal{O}$-hull* of **P** has been called the "θ-visibility hull" of **P** ([54,60]).

As examples of hulls observe that in figure 5.1, (f) is the *$\mathcal{O}$-hull* of (a) for any non-empty $\mathcal{O}$ and (d) and (e) are the 90°-*hulls* of (b) and (c) respectively.

Theorem 3.2.1 established that every choice of $\mathcal{O}$ gave rise to a convexity space thus we may employ the results of chapter 3. From theorem 3.3.1 we have:

Corollary 5.2.1 $\forall \mathcal{O}, \mathbf{P}, \mathbf{Q}$

$$\mathbf{P} \subseteq \mathcal{O}\text{-}hull(\mathbf{P})$$

$$\mathcal{O}\text{-}hull(\mathbf{P}) = \mathbf{P} \iff \mathbf{P} \text{ is } \mathcal{O}\text{-convex}$$

$$\mathbf{P} \subseteq \mathbf{Q} \Longrightarrow \mathcal{O}\text{-}hull(\mathbf{P}) \subseteq \mathcal{O}\text{-}hull(\mathbf{Q})$$

Note that the second relation is equivalent to property (3).

Observe that if $\mathcal{O}_1 \subseteq \mathcal{O}_2$ then, for all $\mathbf{P}$, $\mathcal{O}_1$-*hull*($\mathbf{P}$) $\subseteq$ $\mathcal{O}_2$-*hull*($\mathbf{P}$) as was proved in theorem 3.4.1, although the reader should note that the inclusion is *reversed* since if $\mathcal{O}_1 \subseteq \mathcal{O}_2$ then the set of $\mathcal{O}_2$-convex sets is a refinement of the set of $\mathcal{O}_1$-convex sets. That is, if $\mathcal{O}_1 \subseteq \mathcal{O}_2$ then every $\mathcal{O}_2$-convex set is $\mathcal{O}_1$-convex.

In some sense as a set of orientations $\mathcal{O}$ "grows" to include all possible orientations then $\mathcal{O}$-*hull*($\mathbf{P}$) "grows" to be the convex hull of $\mathbf{P}$. Indeed we have the following theorem:

THEOREM **5.2.1** *The set of all $\mathcal{O}$-convex sets over $\Re^2$ for all $\mathcal{O}$ form a lattice under refinement. The set of $[0°, 180°)$-convex sets is the supremum of this lattice and the set of $\emptyset$-convex sets is the infimum.*

Proof: Grätzer ([19] exercise 9.(i), page 7) shows that the set of all subsets of a set form a lattice under inclusion. We need merely observe that refinement is equivalent to inclusion for $\mathcal{O}$-convex sets. ■

Figure 5.2 shows a small sublattice of this lattice. Each bullet is meant to indicate the set of all $\mathcal{O}$-convex sets for that $\mathcal{O}$; as usual the lines indicate dominance.

Observe that if $\mathcal{O}_1 \subset \mathcal{O}_2$ we can always construct sets which are $\mathcal{O}_1$-convex but not $\mathcal{O}_2$-convex, we call such sets *strictly $\mathcal{O}_1$-convex sets.* For example, the dashed polygons in figure 5.3 are $\mathcal{O}$-convex only for special sets of orientations. For example, (a) is $\mathcal{O}$-convex if and only if $\mathcal{O} \subseteq \{0°, 90°\}$, (b) is $\mathcal{O}$-convex if and only if $\mathcal{O} \subseteq \{0°, 45°, 90°\}$ and (c) is $\mathcal{O}$-convex if and only if $\mathcal{O} \subseteq \{0°, 45°, 90°, 135°\}$.

Now we wish to establish conditions on when an $\mathcal{O}$-convex set is connected or simply connected to match property (1).

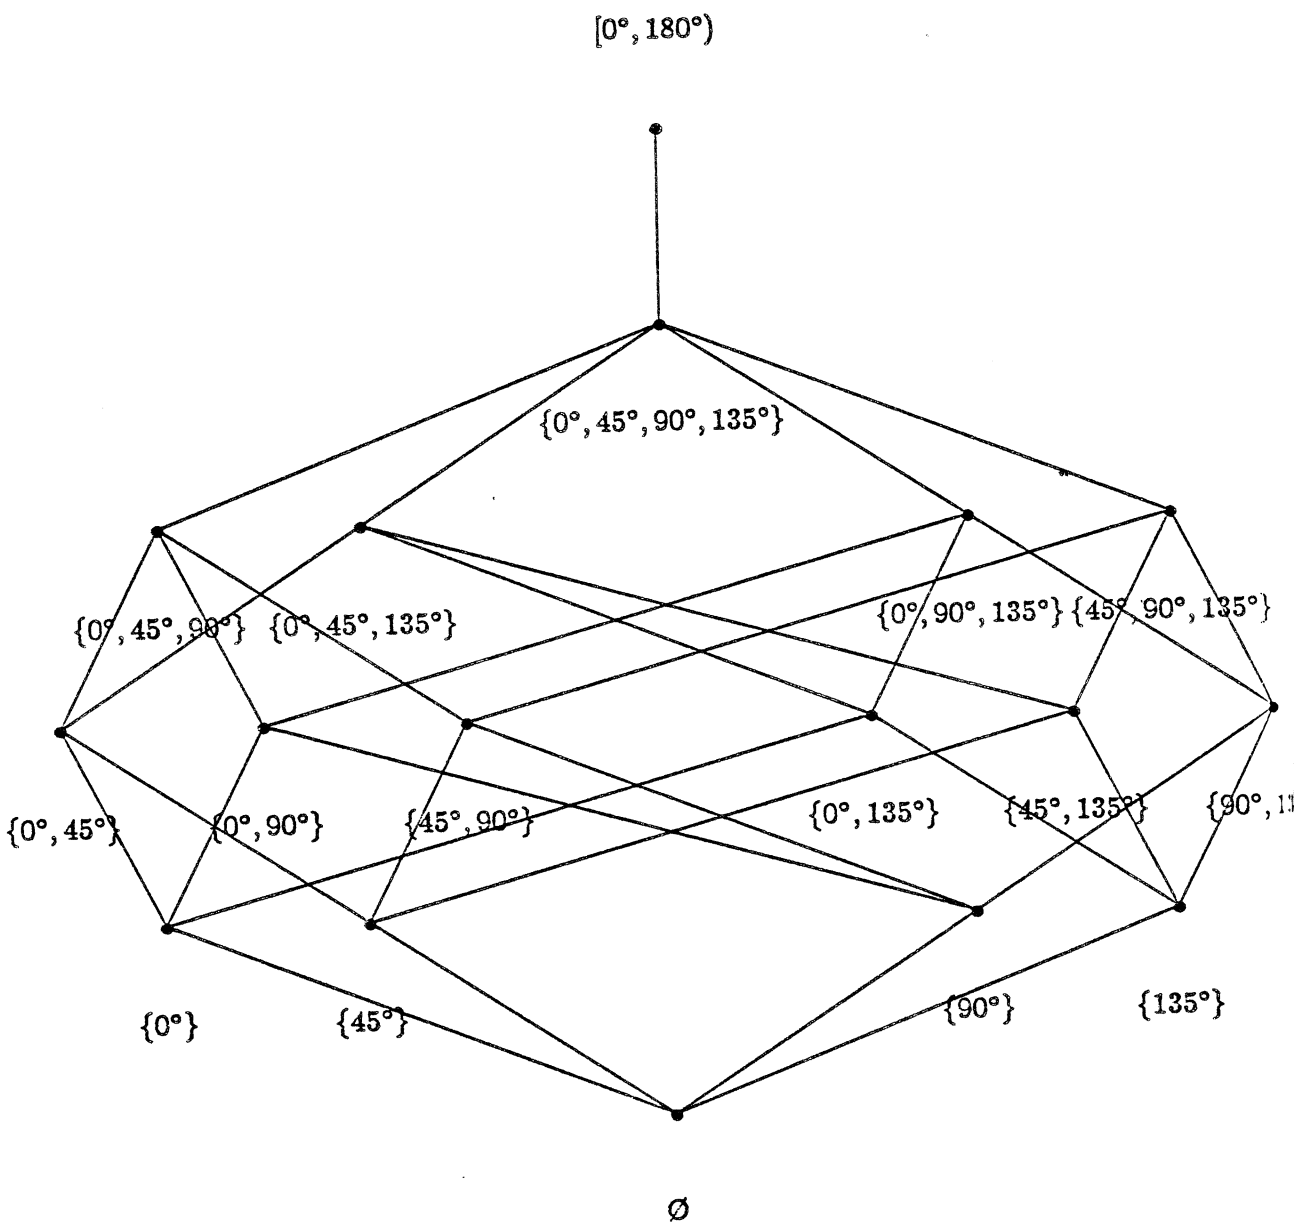

Figure 5.2: A Sublattice of the Lattice of $\mathcal{O}$-Convex Sets

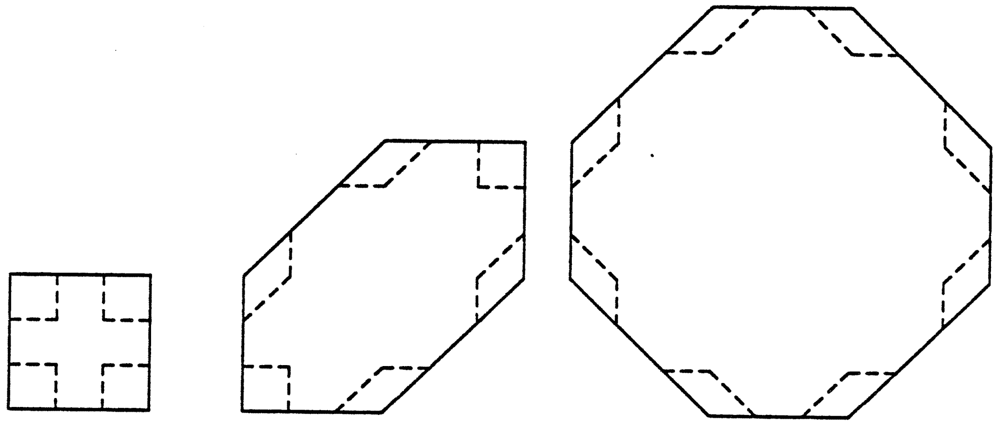

Figure 5.3: Strictly $\mathcal{O}$-Convex Sets

Lemma 5.2.2 *If $\mathcal{O}$ is non-empty and* $\mathbf{P}$ *is connected, then $\mathcal{O}$-hull*$(\mathbf{P})$ *is simply connected.*

Proof: If $\mathbf{P}$ is empty we have nothing to prove, so suppose $\mathbf{P}$ is non-empty.

Suppose that $\mathcal{O}$-*hull*$(\mathbf{P})$ is disconnected. Since $\mathbf{P}$ is connected it can only belong to one of the connected components of $\mathcal{O}$-*hull*$(\mathbf{P})$ (it must belong to at least one otherwise $\mathcal{O}$-*hull*$(\mathbf{P})$ does not contain $\mathbf{P}$). This component must be $\mathcal{O}$-convex, otherwise the entire hull is not $\mathcal{O}$-convex. Hence we may discard all of the other components of $\mathcal{O}$-*hull*$(\mathbf{P})$ and have a smaller $\mathcal{O}$-convex set which contains $\mathbf{P}$. But $\mathcal{O}$-*hull*$(\mathbf{P})$ is the smallest such set. Therefore $\mathcal{O}$-*hull*$(\mathbf{P})$ must be connected if $\mathbf{P}$ is connected.

Suppose that $\mathcal{O}$-*hull*$(\mathbf{P})$ is connected but contains a hole. Since $\mathcal{O}$ is non-empty there exists at least one $\mathcal{O}$-line which cuts this hole. Hence there exists an $\mathcal{O}$-line whose intersection with $\mathcal{O}$-*hull*$(\mathbf{P})$ is not connected. But this implies that $\mathcal{O}$-*hull*$(\mathbf{P})$ is not $\mathcal{O}$-convex hence $\mathcal{O}$-*hull*$(\mathbf{P})$ must be simply connected. ∎

Corollary 5.2.2 *If $\mathcal{O}$ is non-empty and* $\mathbf{P}$ *is connected and $\mathcal{O}$-convex, then* $\mathbf{P}$ *is simply connected.*

Proof: If $\mathbf{P}$ is $\mathcal{O}$-convex then it is its own hull. ∎

Since convex sets are $[0°, 180°)$-convex sets (lemma 5.2.1) and $\mathcal{O} = [0°, 180°)$ is non-empty, by corollary 5.2.2 every convex set is simply connected as stated in property (1) in the first section.

Lemma 5.2.3 *A set is $\mathcal{O}$-convex if and only if it consists of a set of disjoint connected components such that each component is $\mathcal{O}$-convex and no $\mathcal{O}$-line intersects any pair of components.*

Proof: Let $\mathbf{P}$ consist of a set of disjoint connected components such that each component is a connected $\mathcal{O}$-convex set and no $\mathcal{O}$-line intersects any pair of components. Since no $\mathcal{O}$-line can intersect any two of them simultaneously and each component is separately $\mathcal{O}$-convex, the entire collection is $\mathcal{O}$-convex.

Conversely, let $\mathbf{P}$ be disconnected and $\mathcal{O}$-convex. If any one of its components is not $\mathcal{O}$-convex then $\mathbf{P}$ cannot be $\mathcal{O}$-convex. Similarly, if there exists an $\mathcal{O}$-line which intersects any two components then $\mathbf{P}$ cannot be $\mathcal{O}$-convex. ∎

Observe that if $\mathcal{O}$ is the set of all orientations then for each pair of connected components there exists at least one $\mathcal{O}$-line which intersects them. Hence, all $[0°, 180°)$-convex sets are connected. Thus, lemma 5.2.3 together with corollary 5.2.2 implies property (1) when $\mathcal{O} = [0°, 180°)$.

5.3 The Separation Theorem

Now we establish a strong result (the *Separation Theorem*) from which will follow the generalization of the separation property (property 4) of convex sets. Further, we use the separation theorem to prove the *Decomposition Theorem* in next section.

Lemma 5.3.1 *If $\mathbf{P}$ is connected and $p \in \mathcal{O}$-hull$(\mathbf{P})$, then each $\mathcal{O}$-line through p intersects* $\mathbf{P}$.

Proof: If $\mathcal{O}$ is empty, $\mathbf{P}$ is empty, or both, then the lemma is vacuously true since in each case $\mathcal{O}$-*hull*$(\mathbf{P}) = \mathbf{P}$. Further, if $p \in \mathbf{P}$ we have nothing to prove. So suppose that both $\mathcal{O}$ and $\mathbf{P}$ are non-empty and that $p \notin \mathbf{P}$.

Suppose that there exists a $\theta \in \mathcal{O}$ such that the θ-line through p does not intersect $\mathbf{P}$. Then there exists a convex set (and hence an $\mathcal{O}$-convex set) which contains $\mathbf{P}$ and does not contain p, namely, any halfplane bounded by a θ-line separating p and $\mathbf{P}$. Hence p cannot be in the intersection of all $\mathcal{O}$-convex sets which contain $\mathbf{P}$ and so cannot be in its $\mathcal{O}$-*hull*. Thus, each $\mathcal{O}$-line through p intersects $\mathbf{P}$. ∎

This lemma is false if $\mathbf{P}$ is disconnected as the following example shows. In figure 5.4 $\mathbf{P}$ is the pair of points indicated by the bullets. The point p is not in $\mathbf{P}$ yet it is in $\mathcal{O}$-*hull*$(\mathbf{P})$ whenever $0° \in \mathcal{O}$. However, only one $\mathcal{O}$-line through p is guaranteed to intersect $\mathbf{P}$.

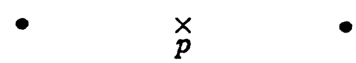

Figure 5.4: Disconnected Point Sets Do Not Force Intersection

THEOREM **5.3.1 (The Separation Theorem)** *Let* $\mathbf{P}$ *be connected and* $p \notin \mathbf{P}$. $p \in \mathcal{O}$-*hull*$(\mathbf{P})$ *if and only if there exists a* $\theta \in \mathcal{O}$ *such that the* θ*-line through* p *intersects* $\mathbf{P}$ *in, at least, two points on either side of* p.

Proof: If either $\mathcal{O}$ or $\mathbf{P}$ is empty then the theorem is vacuously true since $\mathcal{O}$-*hull*$(\mathbf{P}) = \mathbf{P}$. So suppose that both $\mathcal{O}$ and $\mathbf{P}$ are non-empty.

If $p \notin \mathbf{P}$ and there exists an $\mathcal{O}$-line which intersects $\mathbf{P}$ at two points which bracket p then p must be in the $\mathcal{O}$-*hull* of $\mathbf{P}$ (otherwise the $\mathcal{O}$-*hull* would not be $\mathcal{O}$-convex).

Conversely, if $\mathbf{P}$ is connected and $p \in \mathcal{O}\text{-}hull(\mathbf{P}) \setminus \mathbf{P}$ then all $\mathcal{O}$-lines through p must intersect $\mathbf{P}$ (lemma 5.3.1).

We prove the claim for the three cases in which we have exactly one orientation in $\mathcal{O}$, two or more with at least one $\mathcal{O}$-free range and, finally, if all orientations are in $\mathcal{O}$ (that is, there are no $\mathcal{O}$-free ranges).

Case 1: $\mathcal{O} = \theta$.

The θ-line through p must cut $\mathbf{P}$. Suppose that it only cuts it on one side of p (say to the right of p). Then we may delete p and all other points in $\theta\text{-}hull(\mathbf{P})$ on the left θ-ray from p and so obtain a smaller θ-convex set which contains $\mathbf{P}$. But $\theta\text{-}hull(\mathbf{P})$ is the smallest such set. Hence p cannot be in $\theta\text{-}hull(\mathbf{P})$. Hence the θ-line through p must cut $\mathbf{P}$ on both sides of p.

Case 2: $\mathcal{O}$ contains two or more orientations but not all.

Every $\mathcal{O}$-line through p must cut $\mathbf{P}$. Suppose that none of them cut $\mathbf{P}$ both to the left and to the right of p. We show that if $\mathbf{P}$ is connected that this implies that there is an $\mathcal{O}$-stairhalfplane containing $\mathbf{P}$ and not p.

Assume, without loss of generality, that $0° \in \mathcal{O}$. p divides the horizontal line through p into two rays, arbitrarily call one ray the "left ray" and the other the "right ray." The horizontal line through p is an $\mathcal{O}$-line and so either the left or the right ray cannot intersect $\mathbf{P}$. By assumption, at least one must intersect $\mathbf{P}$. Assume, without loss of generality, that the right ray intersects $\mathbf{P}$ (and, consequently, that the left ray does not).

Let θ_1 be the largest orientation in $\mathcal{O}$ such that the θ_1-line through p intersects $\mathbf{P}$ on the right. Similarly, let θ_2 be the smallest orientation in $\mathcal{O}$ such that the θ_2-line through p does not intersect $\mathbf{P}$ on the right. Since $\mathcal{O}$ is closed θ_1 and θ_2 always exist. By construction, the range (θ_1, θ_2) is $\mathcal{O}$-free, for if there exists a $\theta \in (\theta_1, \theta_2)$ such that $\theta \in \mathcal{O}$ then either θ_1 is not the largest or θ_2 is not the smallest such orientation.

Assume, without loss of generality, that $[\theta_1, \theta_2] = [0°, 90°]$. From the above argument we know that $(0°, 90°)$ is a maximal $\mathcal{O}$-free range; that the 0°-line through p intersects $\mathbf{P}$ to the right of p (and does not intersect $\mathbf{P}$ to the left of p); and that the 90°-line through p intersects $\mathbf{P}$ below p (and does not intersect $\mathbf{P}$ above p). Thus, the curve made up of the horizontal ray to the left of p and the vertical ray above p does not intersect $\mathbf{P}$. Further, this curve has span $[0°, 90°]$ and the open range $(0°, 90°)$ is $\mathcal{O}$-free thus, from lemma 4.2.1, this curve is an $\mathcal{O}$-stairline. Since $\mathbf{P}$ is connected it lies wholly in one of the two $\mathcal{O}$-stairhalfplanes this $\mathcal{O}$-stairline defines (in this case the complement of the "north-west" quadrant). Thus there is an $\mathcal{O}$-convex set containing $\mathbf{P}$ and not p. Hence p is not in the intersection of all $\mathcal{O}$-convex sets containing $\mathbf{P}$, and so, $p \notin \mathcal{O}\text{-}hull(\mathbf{P})$, a contradiction.

Hence at least one of the $\mathcal{O}$-lines through p must cut $\mathbf{P}$ to the left and to the right of p. (See figure 5.5 for a simple example with $\mathcal{O} = \{0°, 90°, 135°\}$).

Case 3: $\mathcal{O} = [0°, 180°)$.

Here $\mathcal{O}\text{-}hull(\mathbf{P})$ is the convex hull of $\mathbf{P}$. The result then follows from the well-known separation theorem of Fenchel ([16]). ∎

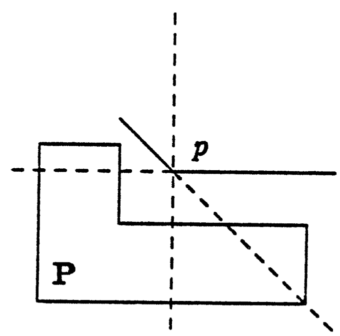

Figure 5.5: A Simple Example of the Separation Theorem

This theorem is false if $\mathbf{P}$ is allowed to be disconnected as the following example shows. In figure 5.6, $\mathbf{P}$ is the set of points indicated by the bullets. The point p is not in $\mathbf{P}$ yet it is in $\mathcal{O}\text{-}hull(\mathbf{P})$ whenever $\{0°, 90°\} \subseteq \mathcal{O}$. However, there does not exist an $\mathcal{O}$-line through

p which cuts $\mathbf{P}$ on both sides of p.

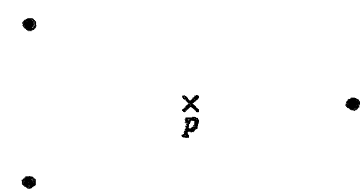

Figure 5.6: Disconnected Point Sets Do Not Support Separation

Corollary 5.3.1 *Let* $\mathbf{P}$ *be connected. If* $p \notin \mathcal{O}$*-hull*$(\mathbf{P})$ *then there exists an* $\mathcal{O}$*-stairlin separating* p *and* $\mathbf{P}$.

Proof: Let $\mathbf{P}$ be connected. If $p \notin \mathcal{O}$-*hull*$(\mathbf{P})$ then from theorem 5.3.1 there exists n $\theta \in \mathcal{O}$ such that $\mathcal{L}[\theta, p]$ intersects $\mathbf{P}$ on both sides of p. Thus we may construct an $\mathcal{O}$-stairline separating p and $\mathbf{P}$ as in the proof of theorem 5.3.1. ∎

When $\mathcal{O} = [0°, 180°)$ all $\mathcal{O}$-stairlines are lines and this corollary becomes property (5).

Corollary 5.3.2 *If* $\mathbf{P}$ *is connected then* $\mathcal{O}$*-hull*$(\mathbf{P})$ *is the intersection of all* $\mathcal{O}$*-stairhalfplan s containing* $\mathbf{P}$.

Proof: Since $\mathcal{O}$-*hull*$(\mathbf{P})$ is a subset of every $\mathcal{O}$-convex set which contains $\mathbf{P}$ it then follow that $\mathcal{O}$-*hull*$(\mathbf{P})$ is a subset of the intersection of all $\mathcal{O}$-stairhalfplanes which contain $\mathbf{P}$.

If $p \notin \mathcal{O}$-*hull*$(\mathbf{P})$ then from corollary 5.3.1 there exists an $\mathcal{O}$-stairline separating p an $\mathbf{P}$. Thus, there exists an $\mathcal{O}$-stairhalfplane containing $\mathbf{P}$ and not p. Thus, p is not in th intersection of all $\mathcal{O}$-stairhalfplanes which contain $\mathbf{P}$. Therefore, if $p \notin \mathcal{O}$-*hull*$(\mathbf{P})$ then p not in the intersection of all $\mathcal{O}$-stairhalfplanes which contain $\mathbf{P}$. Thus, contrapositively, th intersection of all $\mathcal{O}$-stairhalfplanes which contain $\mathbf{P}$ is a subset of $\mathcal{O}$-*hull*$(\mathbf{P})$. Hence th result follows. ∎

When $\mathcal{O} = [0°, 180°)$ all $\mathcal{O}$-stairhalfplanes are halfplanes and this corollary becomes property (5).

The following example shows that this corollary is false if **P** is disconnected. In figure 5.7(a) **P** is the set made up of the three points indicated by the bullets. If $\mathcal{O} = \{0°, 90°\}$ then $\mathcal{O}$-*hull*(**P**) = **P** but the intersection of all $\mathcal{O}$-stairhalfplanes containing **P** is figure 5.7(b).

Figure 5.7: Intersection of $\mathcal{O}$-Stairhalfplanes Do Not Form the Hull of Disconnected Sets

5.4 The Decomposition Theorem

Intuitively, we may think of forming the $\mathcal{O}$-*hull* of a set **P** by sweeping a line of each orientation in $\mathcal{O}$ across **P** and adding suitable line segments to the hull formed so far so that it is convex in each orientation in $\mathcal{O}$ (if $\mathcal{O}$ is empty then we do not add anything to **P**). Thinking of it this way it appears reasonable that the hull we eventually produce is unchanged if we choose a different sweeping order. As we prove in theorem 5.4.1 this is, in fact, the case but only for *connected* sets. For disconnected sets lemma 5.4.1 is the strongest possible result.

As a by-product of the following theorems we establish the validity of two assumptions made in the literature for orthogonal hulls. Sack ([54]) showed, in the orthogonal case, that the horizontal hull of the vertical hull of an orthogonal polygon (or alternately the vertical hull of the horizontal hull) is equivalent to the union of both hulls. It was taken as self-evident that their union is the smallest horizontally and vertically convex polygon enclosing the orthogonal polygon. Theorem 5.4.1 validates this assumption. Toussaint and Sack ([60]) made the observation that the convex hull is the union of the "visibility hulls" over all directions of visibility. Theorem 5.4.1 supplies the proof for this observation.

The following lemma follows from the refinement theorem (theorem 3.4.2). Observe that the hull with respect to the intersection of the set of convexity spaces has been replaced by the hull with respect to the union of the sets of orientations. This is because when we form the union of two sets of orientations the hull with respect to their union is convex with respect to *both* sets of orientations.

Lemma 5.4.1 *Given $n \geq 1$ sets of orientations $\mathcal{O}_i$, $1 \leq i \leq n$; then, $\forall \mathbf{P}$,*

$$\bigcup_{i=1}^{n}(\mathcal{O}_i\text{-}hull(\mathbf{P})) \subseteq (\bigwedge_{i=1}^{n} \mathcal{O}_i)\text{-}hull(\mathbf{P}) \subseteq (\bigcup_{i=1}^{n} \mathcal{O}_i)\text{-}hull(\mathbf{P})$$

Simple counterexamples (see figure 5.8) show that this result is the best possible, in that there exists sets for which the respective converses are false. However, we can strengthen lemma 5.4.1 considerably by restricting $\mathbf{P}$ to be connected. As we prove below when $\mathbf{P}$ is connected the containment relations in lemma 5.4.1 become equalities. Thus, in the language of chapter 3, *with respect to connected subsets of $\Re^2$ any two orientation convexity spaces are mutually invariant.*

THEOREM **5.4.1 (The Orientation Decomposition Theorem)** *Given $n \geq 1$ sets of orientations $\mathcal{O}_i$, $1 \leq i \leq n$; if $\mathbf{P}$ is connected,*

$$\bigcup_{i=1}^{n}(\mathcal{O}_i\text{-}hull(\mathbf{P})) = (\bigwedge_{i=1}^{n} \mathcal{O}_i)\text{-}hull(\mathbf{P}) = (\bigcup_{i=1}^{n} \mathcal{O}_i)\text{-}hull(\mathbf{P})$$

Proof: Because of lemma 5.4.1 we need establish only that if $\mathbf{P}$ is connected then $(\bigcup_{i=1}^{n} \mathcal{O}_i)\text{-}hull(\mathbf{P}) \subseteq \bigcup_{i=1}^{n}(\mathcal{O}_i\text{-}hull(\mathbf{P}))$. If $\mathbf{P}$ or $(\bigcup_{i=1}^{n} \mathcal{O}_i)$ is empty then this holds, so assume that both are non-empty.

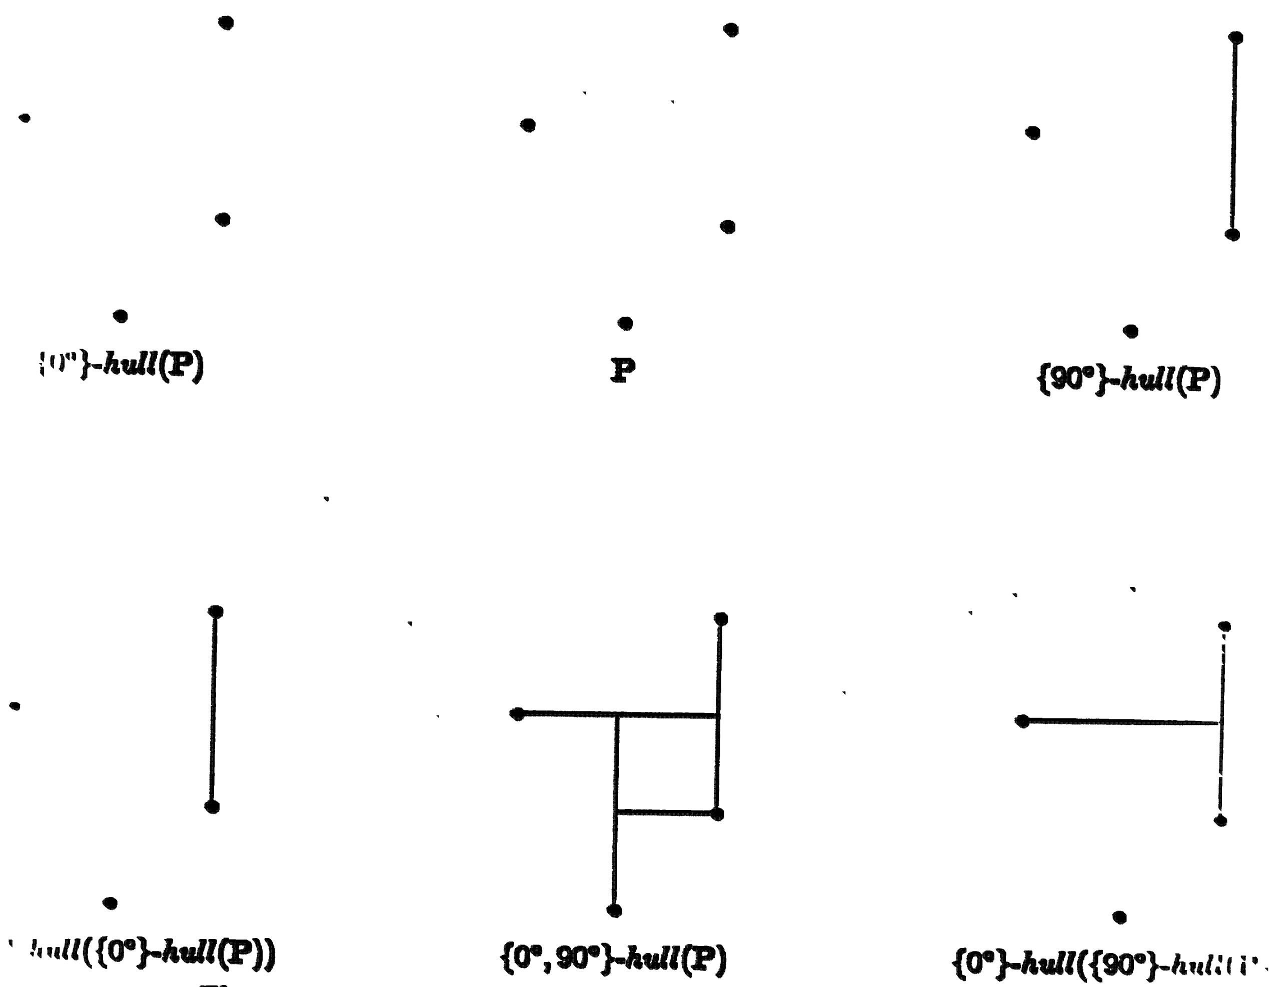

Figure 5.8: Hulls of Disconnected Sets Are not Decomposable

Let $p \in (\bigcup_{i=1}^{n} \mathcal{O}_i)$-*hull*(P). If $p \in \mathbf{P}$ then $p \in \bigcup_{i=1}^{n}(\mathcal{O}_i\text{-}hull(\mathbf{P}))$. So suppose that $p \in (\bigcup_{i=1}^{n} \mathcal{O}_i)$-*hull*$(\mathbf{P}) \setminus \mathbf{P}$. From theorem 5.3.1, we know that there must exist a $\theta \in (\bigcup_{i=1}^{n} \mathcal{O}_i)$ such that the θ-line through p cuts $\mathbf{P}$ to the left and right of p. This means that p must be in $\mathcal{O}_i$-*hull*(P) for some i such that $\theta \in \mathcal{O}_i$. Hence, $p \in \bigcup_{i=1}^{n}(\mathcal{O}_i\text{-}hull(\mathbf{P}))$ and the result follows. ∎

This decomposition result immediately yields an algorithm to find the hull of any connected set given that we can find the hull in one orientation and that we can find the union of two or more hulls. As it turns out, however, connected $\mathcal{O}$-convex sets have considerably more structure than this and we can exploit this structure to construct optimal algorithms to find the hull of any connected set. We return to this topic in chapter 7.

5.5 Characterizing Connected $\mathcal{O}$-Convex Sets

In this section we characterize *connected* $\mathcal{O}$-convex sets by a necessary and sufficient condition on the shape of their boundary and we prove a generalized version of property (6). Recall that a point is *extremal* in $\mathbf{P}$ if it is a point of support of $\mathbf{P}$ with respect to some line.

Definition: p is said to be *$\mathcal{O}$-extremal* in $\mathbf{P}$ if p is a point of support of $\mathbf{P}$ with respect to an $\mathcal{O}$-line.

Definition: A subset of a curve is said to be a *maximal $\mathcal{O}$-stairsegment* in the curve if it is an $\mathcal{O}$-stairsegment and it is not a proper subset of any other $\mathcal{O}$-stairsegment in the curve.

We now show that the boundary of a simply connected $\mathcal{O}$-convex set may be completely characterized in terms of maximal $\mathcal{O}$-stairsegments.

THEOREM **5.5.1 (The Boundary Theorem)** *If $|\mathcal{O}| \geq 2$ and $\mathbf{P}$ is a simply connected set then $\mathbf{P}$ is $\mathcal{O}$-convex if and only if the portions of its boundary between any two consecutive $\mathcal{O}$-extremal points are maximal $\mathcal{O}$-stairsegments.*

Proof: If $\mathbf{P}$ is simply connected and its boundary is made up only of $\mathcal{O}$-stairsegments meeting at $\mathcal{O}$-extremal points in $\mathbf{P}$ then the only way in which $\mathbf{P}$ could fail to be $\mathcal{O}$-convex is if some $\mathcal{O}$-line intersects one of the $\mathcal{O}$-stairsegments more than once, since no $\mathcal{O}$-line can intersect the boundary of $\mathbf{P}$ more than twice. But this is impossible, since the intersection of an $\mathcal{O}$-line and an $\mathcal{O}$-stairline is connected. Hence $\mathbf{P}$ must be $\mathcal{O}$-convex.

Conversely, suppose that $\mathbf{P}$ is $\mathcal{O}$-convex. Consider any pair of distinct consecutive $\mathcal{O}$-extremal points p and q of $\mathbf{P}$. Let $S[p,q]$ be the portion of $\mathbf{P}$'s boundary connecting p and q. If $\Theta(\mathcal{LS}[p,q]) \in \mathcal{O}$ then $S[p,q] = \mathcal{LS}[p,q]$ (since p and q are consecutive $\mathcal{O}$-extremal points of $\mathbf{P}$) and $\mathcal{LS}[p,q]$ is an $\mathcal{O}$-stairsegment. Hence assume that $\Theta(\mathcal{LS}[p,q]) \notin \mathcal{O}$. Without loss of generality, assume that $\mathcal{LS}[p,q]$'s maximal $\mathcal{O}$-free range is $(0°, 90°)$ and that p is below and to the left of q (see figure 5.9).

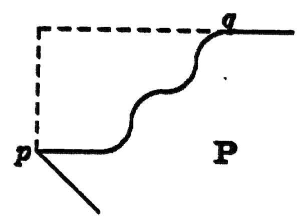

Figure 5.9: $S[p,q]$ is a Maximal $\mathcal{O}$-Stairsegment

Now sweep a horizontal line, that is, a 0°-line, from q down to p. If at any time in this sweep this line intersects $S[p,q]$ more than once then $\mathbf{P}$ cannot be $\mathcal{O}$-convex, and similarly for a vertical line sweeping from p to q. Hence $S[p,q]$ is an $\mathcal{O}$-stairsegment connecting p and q. Trivially, it is maximal since its endpoints are $\mathcal{O}$-extremal in $\mathbf{P}$. ■

Observe that for the convex hull (that is, $\mathcal{O} = [0°, 180°)$) all boundary points are $\mathcal{O}$-extremal and so the maximal $\mathcal{O}$-stairsegments in the boundary can only be points, lines, line segments, or rays.

In the theory of convex sets two points are said to be mutually *visible* in a given set if the line segment joining them lies wholly in the set. Thinking of $\mathcal{O}$-stairlines as the analogues of straight lines we are led to define a generalized form of visibility in which two points in a set are mutually visible if there exists *at least one $\mathcal{O}$-stairsegment* joining them which lies wholly in the set. This leads to the next characterization of $\mathcal{O}$-convex sets and again it only applies to *connected* $\mathcal{O}$-convex sets. In particular, when $\mathcal{O} = [0°, 180°)$ this characterization corresponds to property (6), since convex sets are connected.

THEOREM **5.5.2 (The Visibility Theorem)** *If $|\mathcal{O}| \geq 2$ and $\mathbf{P}$ is a simply connected set then $\mathbf{P}$ is $\mathcal{O}$-convex if and only if for all p and q in $\mathbf{P}$ there exists an $\mathcal{O}$-stairsegment in $\mathbf{P}$ with endpoints p and q.*

Proof: Suppose that $\mathbf{P}$ is simply connected and for all $p, q \in \mathbf{P}$ there exists an $\mathcal{O}$-stairsegment joining them lying in $\mathbf{P}$. If $\Theta(\mathcal{LS}[p,q]) \in \mathcal{O}$ then $\mathcal{O}$-$\|[p,q] = \mathcal{LS}[p,q]$. Hence there is only one $\mathcal{O}$-stairsegment joining p and q and so it must lie in $\mathbf{P}$. Hence $\mathbf{P}$ is $\mathcal{O}$-convex.

Conversely, suppose that $\mathbf{P}$ is connected and $\mathcal{O}$-convex. If $p, q \in \mathbf{P}$ and $\Theta(\mathcal{LS}[p,q]) \in \mathcal{O}$ then there exists an $\mathcal{O}$-stairsegment lying in $\mathbf{P}$ joining p and q—namely, $\mathcal{LS}[p,q]$ (otherwise $\mathbf{P}$ is not $\mathcal{O}$-convex). Suppose then that $\Theta(\mathcal{LS}[p,q]) \notin \mathcal{O}$. Consider $\mathcal{O}$-$\|[p,q]$. If either arm of $\mathcal{O}$-$\|[p,q]$ lies in $\mathbf{P}$ then there exists an $\mathcal{O}$-stairsegment lying in $\mathbf{P}$ joining p and q since either arm of $\mathcal{O}$-$\|[p,q]$ is an $\mathcal{O}$-stairsegment. Assume then that neither arm lies wholly in $\mathbf{P}$. Since the lower arm (say) consists of two $\mathcal{O}$-segments and it does not lie wholly in $\mathbf{P}$ then it must intersect the boundary of $\mathbf{P}$ exactly twice (otherwise $\mathbf{P}$ would not be $\mathcal{O}$-convex).

Both of these intersection points must belong to one maximal $\mathcal{O}$-stairsegment since if they belonged to separate maximal $\mathcal{O}$-stairsegments then there must be at least one $\mathcal{O}$-extremal point on $\mathbf{P}$'s boundary between the two intersection points. This means that there must exist at least one $\mathcal{O}$-orientation in $\mathcal{LS}[p,q]$'s $\mathcal{O}$-free range. But this is impossible.

Therefore we can construct an $\mathcal{O}$-stairsegment lying in $\mathbf{P}$ connecting p and q by starting at p and following the lower arm of $\mathcal{O}$-$\|[p,q]$ until we encounter $\mathbf{P}$'s boundary then follow the boundary until we intersect the arm again, and then follow the arm to q. ∎

With $\mathcal{O}$-stairlines replacing lines we can generalize convexity in ways other than the one we investigate in this chapter. For example, a set $\mathbf{P}$ is said to be *strongly $\mathcal{O}$-convex* if for every pair of points p and q in $\mathbf{P}$ *all* $\mathcal{O}$-stairsegments with endpoints p and q lie in $\mathbf{P}$. In the next chapter we show that this definition of convexity always produces *convex $\mathcal{O}$-oriented* sets. Indeed, when $\mathcal{O} = \{0°, 90°\}$, the strong $\mathcal{O}$-convex hull of $\mathbf{P}$ is just the *bounding box* of $\mathbf{P}$. We shall see more of this and other variants of $\mathcal{O}$-convexity in the succeeding chapters.

5.6 $\mathcal{O}$-Convexity as a Generalization of Convexity

We have shown that $\mathcal{O}$-convex sets are a generalization of both convex sets and orthogonally convex sets and that the properties of both are special cases of the properties of $\mathcal{O}$-convex sets. The main characteristic of convex sets that we have lost in generalizing to $\mathcal{O}$-convex sets is *connectivity*. A convex set is always connected.

Connected $\mathcal{O}$-convex sets have all of the properties of convex sets listed at the beginning of the chapter if we replace "line" by "$\mathcal{O}$-stairline" and "line segment" by "$\mathcal{O}$-stairsegment." However, note that there is no longer a unique "line segment" joining two points. In the following we assume that $\mathbf{P}$ is a *connected* $\mathcal{O}$-convex set.

1. If $\mathcal{O}$ is non-empty then $\mathbf{P}$ is simply connected (corollary 5.2.2). Indeed, the connected components of any $\mathcal{O}$-convex set are simply connected if $\mathcal{O}$ is non-empty (lemma 5.2.3).

2. The intersection of $\mathbf{P}$ and any $\mathcal{O}$-line is connected (by definition). This holds even if $\mathbf{P}$ is disconnected.

3. $\mathbf{P}$ is the intersection of all $\mathcal{O}$-convex sets which contain it (corollary 5.2.1). This holds even if $\mathbf{P}$ is disconnected.

4. If $p \notin \mathbf{P}$ then there exists an $\mathcal{O}$-stairline separating p and $\mathbf{P}$ (corollary 5.3.1).

5. **P** is the intersection of all $\mathcal{O}$-stairhalfplanes which contain it (corollary 5.3.2).

6. If $p, q \in \mathbf{P}$ then there exists an $\mathcal{O}$-stairsegment in **P** connecting p and q (theorem 5.5.2).

Chapter 6

Restricted Orientation Visibility

Closure under intersection provides one of the main characterizations of convex sets. However, there is an equally important characterization of convex sets, namely, a set is convex if *every pair of points in the set see each other*. In this chapter we investigate alternative ways in which points may be said to "see" each other when we have a restricted set of orientations. We also add the concept of *visibility* to abstract convexity spaces and show how the new definitions of visibility may be unified. Finally, we give algorithms to find a generalized kernel and we generalize Chvátal's Art Gallery Theorem.

6.1 Weak and Strong $\mathcal{O}$-Convexity

Definition: Given two sets $\mathbf{P}$ and $\mathbf{Q}$, we say that *p sees q in $\mathbf{P}$ via $\mathbf{Q}$* if $p, q \in \mathbf{Q}$, $\mathbf{Q}$ is connected, and $\mathbf{Q} \subseteq \mathbf{P}$.

Definition: $\mathbf{P}$ is said to be *weakly $\mathcal{O}$-convex* (*strongly $\mathcal{O}$-convex*) if $\forall p, q \in \mathbf{P}$ p sees q in P via an $\mathcal{O}$-stairsegment (via $\mathcal{O}$-$\|[p, q]$).

Observe that the set of weakly $\emptyset$-convex sets is just the set of all connected sets and that there are only two strongly $\emptyset$-convex sets, namely, the empty set, $\emptyset$, and the whole plane,

$\Re^2$. Also, observe that the set of weakly $\mathcal{O}$-convex sets is not closed under intersection and so does not form a convexity space whereas the set of strongly $\mathcal{O}$-convex sets does. In fact, as we show, the set of convex sets forms a convexity space which is a refinement of the set of strongly $\mathcal{O}$-convex sets.

A strongly $\mathcal{O}$-convex set is the usual Minkowski convex set in an arbitrary metric space. In a metric space two points do not necessarily have a unique line segment joining them and a set is said to be (metrically) convex if for every pair of points in the set every line segment they define is also in the set. It is straightforward to show that:

THEOREM **6.1.1** $\mathbf{P}$ *is strongly $\mathcal{O}$-convex if and only if the intersection of* $\mathbf{P}$ *and every $\mathcal{O}$-stairline is connected.*

Observe that convex sets are a subset of weakly $\mathcal{O}$-convex sets and that strongly $\mathcal{O}$-convex sets are a subset of convex sets. Further, both containments are proper. The three definitions of convex sets are identical only when $\mathcal{O} = [0°, 180°)$.

The notion of weak $\mathcal{O}$-convexity coincides exactly with that of $\mathcal{O}$-convexity when restricted to *connected sets*. The following result is a restatement of the Visibility Theorem (theorem 5.5.2).

THEOREM **6.1.2** $\mathbf{P}$ *is weakly $\mathcal{O}$-convex if and only if it is a connected $\mathcal{O}$-convex set.*

We now completely characterize strongly $\mathcal{O}$-convex sets.

Lemma 6.1.1 *If* $\mathbf{P}$ *is strongly $\mathcal{O}$-convex then* $\mathbf{P}$ *is convex.*

Proof: If $\mathbf{P}$ is strongly $\mathcal{O}$-convex then $\forall p, q \in \mathbf{P}$ $\mathcal{O}$-$\|[p,q] \subseteq \mathbf{P}$. Since $\mathcal{LS}[p,q] \subseteq \mathcal{O}$-$\|[p,q]$, $\mathbf{P}$ is convex. ∎

THEOREM **6.1.3** $\mathbf{P}$ *is strongly $\mathcal{O}$-convex if and only if* $\mathbf{P}$ *is a convex $\mathcal{O}$-object.*

Proof: Suppose that **P** is a convex $\mathcal{O}$-object and let p and q be two points in **P**. If $\Theta(\mathcal{LS}[p,q]) \in \mathcal{O}$ then $\mathcal{O}\text{-}\|[p,q] = \mathcal{LS}[p,q]$. Thus, since **P** is convex, $\mathcal{O}\text{-}\|[p,q] \subseteq \mathbf{P}$. Otherwise, if $\Theta(\mathcal{LS}[p,q]) \notin \mathcal{O}$ suppose, without loss of generality, that $(0°, 90°)$ is $\mathcal{L}$'s maximal $\mathcal{O}$-free range. Let p be below and to the left of q. If $\mathcal{O}\text{-}\|[p,q] \not\subseteq \mathbf{P}$ then there exists some point $r \in \mathcal{O}\text{-}\|[p,q]$ such that $r \notin \mathbf{P}$. Since **P** is convex there must exist a line $\mathcal{L}$ separating **P** and r. Since $p, q \in \mathbf{P}$ and **P** is convex, $\Theta(\mathcal{L}) \in (0°, 90°)$ (otherwise $\mathcal{L}$ separates p and q). Thus $\mathcal{L}$ cuts off one of the corners of $\mathcal{O}\text{-}\|[p,q]$, say *NW-conjugate*$(p,q)$ (see figure 6.1). Translate $\mathcal{L}$ until it becomes a line of support for **P**. Since **P** is $\mathcal{O}$-oriented and $\mathcal{L}$ is not, no edge in **P**'s boundary can be collinear with $\mathcal{L}$. Thus its boundary nearest r must contain a reflex angle. But then **P** is not convex, a contradiction. Thus for all $p, q \in \mathbf{P}$ $\mathcal{O}\text{-}\|[p,q] \subseteq \mathbf{P}$.

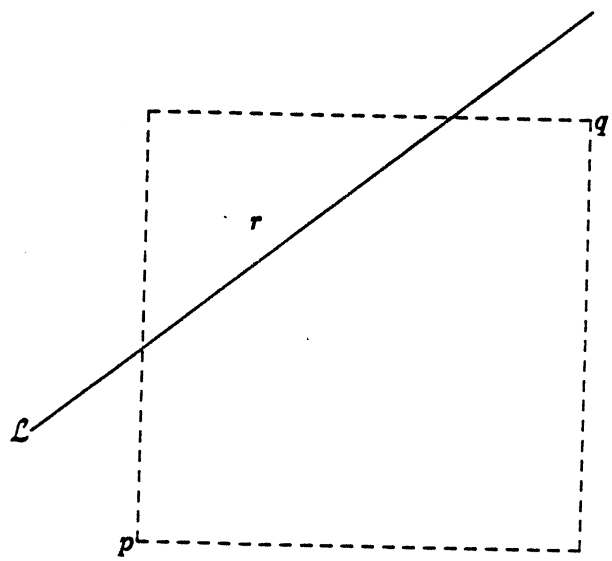

Figure 6.1: The Strong $\mathcal{O}$-*hull* Is A Convex $\mathcal{O}$-Object

Conversely, suppose that **P** is strongly $\mathcal{O}$-convex but is not an $\mathcal{O}$-object. This means that some section of its boundary, say between the two points p and q, is not an $\mathcal{O}$-line segment, that is, $\Theta(\mathcal{LS}[p,q]) \notin \mathcal{O}$. Since **P** is convex (lemma 6.1.1) we may assume that p and q are extremal in **P**. Without loss of generality, let $(0°, 90°)$ be $\mathcal{LS}[p,q]$'s maximal $\mathcal{O}$-free range and assume that p is below and to the left of q. Now *NW-conjugate*(p,q) cannot be in **P** since p and q are extremal in **P**. But this point must be in **P** since it is in

$\mathcal{O}$-$||[p, q]$ and P is strongly $\mathcal{O}$-convex, a contradiction. Thus P must be an $\mathcal{O}$-object. ∎

6.2 $\mathcal{O}$-Starshapedness

In [54] Sack characterized starshaped sets for orthogonal polygons. However his definition of visibility was the usual one of *straight line visibility*. When talking about restricted orientation sets it seems more reasonable to define visibility in terms of *$\mathcal{O}$-stairlines* rather than lines. For example, Sack (see [54]) showed that under line visibility a starshaped orthogonal polygon is orthogonally convex. This is contrary to the expected analogy from convex sets since a convex set is starshaped but a starshaped set is not necessarily convex.

Definition: p and q are said to be *weakly $\mathcal{O}$-visible (strongly $\mathcal{O}$-visible)* in P if p sees q in P via an $\mathcal{O}$-stairsegment (via $\mathcal{O}$-$||[p, q]$).

Definition: The set of points in P which weakly (strongly) see all points in P is said to be the *weak (strong) $\mathcal{O}$-kernel* of P. P is said to be weakly (strongly) *$\mathcal{O}$-starshaped* if P empty or if P is non-empty and the weak (strong) *$\mathcal{O}$-kernel*(P) is non-empty.

Keil [28] investigates strong $\mathcal{O}$-visibility in $\mathcal{O}$-polygons for $\mathcal{O} = \{0°, 90°\}$. That is, two points see each other in P if the orthogonal rectangle they define is in P. Overmars and Wood ([46]), and Munro *et al.* [40] examine a variant of this version of visibility.

The proofs of the following results follow directly from the definitions and are omitted.

THEOREM **6.2.1** $\forall$**P** ,

1. *If* P *is starshaped then* P *is weakly $\mathcal{O}$-starshaped.*
2. *If* P *is strongly $\mathcal{O}$-starshaped then* P *is starshaped.*
3. *Weak (strong) $\mathcal{O}$-kernel*(P) *is weakly (strongly) $\mathcal{O}$-convex.*

4. **P** *is weakly (strongly) $\mathcal{O}$-convex if and only if weak (strong) $\mathcal{O}$-kernel*$(\mathbf{P}) = \mathbf{P}$.

Because of the similarity of the above results, in the next section we search for an algebraic characterization of visibility which includes all three versions of visibility.

6.3 Abstract Visibility

Most results in abstract convexity theory are generalizations of combinatorial results such as Helly's Theorem, Carathéodory's Theorem and Radon's Theorem ([24]). Abstract convexity theory needs a definition of visibility capturing the essence of visibility in the Euclidean plane to make it more useful to geometers. The definition should be general enough so that its results apply to most convexity spaces. In particular, they should be applicable to the $\mathcal{O}$-visibility spaces we have just defined. Towards this end we define the following *visibility relation:*

Definition: Given a convexity space $(\mathcal{S}, \mathcal{C})$ and a set $\mathbf{P} \subseteq \mathcal{S}$, p is said to *see* q in $\mathbf{P}$ if $\exists \mathbf{Q} \in \mathcal{C}$ such that $p, q \in \mathbf{Q}$ and $\mathbf{Q} \subseteq \mathbf{P}$.

This generalizes the notion of visibility in a set in $\Re^2$, where instead of requiring that the line segment joining the two points be in $\mathbf{P}$ we merely require that the two points be in some convex set which is contained in $\mathbf{P}$. Thus we avoid defining what, if anything, we mean by a "line segment" since we have not given any metric to our space. The above visibility relation includes line visibility since in normal convexity a line segment is convex. It also includes weak and strong $\mathcal{O}$-visibility.

Observe that we cannot remove the condition "$\mathbf{Q} \subseteq \mathbf{P}$" since otherwise *all* points would see each other in every set since the groundset of a convexity space is always convex and contains all sets in the space.

An alternative way to define visibility is to require that p sees q in $\mathbf{P}$ if $\mathcal{C}$-*hull*$(\{p, q\}) \subseteq \mathbf{P}$. However, this definition fails to give geometric insight if the hull of two disconnected points may be disconnected (as happens with weak $\mathcal{O}$-visibility).

Definition: Given a convexity space $(\mathcal{S}, \mathcal{C})$ and a subset $\mathbf{P}$ of $\mathcal{S}$, the *$\mathcal{C}$-kernel* of $\mathbf{P}$ is the set of all points in $\mathbf{P}$ which see every point in $\mathbf{P}$. $\mathbf{P}$ is said to be *$\mathcal{C}$-starshaped* if either $\mathbf{P} = \emptyset$ or $\mathcal{C}\text{-}kernel(\mathbf{P}) \neq \emptyset$.

Lemma 6.3.1 *Given a convexity space $(\mathcal{S}, \mathcal{C})$, if $\mathbf{P} \in \mathcal{C}$ then $\mathcal{C}\text{-}kernel(\mathbf{P}) = \mathbf{P}$.*

Proof: We need merely take $\mathbf{Q} = \mathbf{P}$ in the definition of the visibility relation. ∎

This result says that any two points in a given convex set "see" each other, which is just what we want. It also establishes that every two points in a convexity space can see each other in the groundset since the groundset is always convex.

Corollary 6.3.1 *Given a convexity space $(\mathcal{S}, \mathcal{C})$, if $\mathbf{P}$ is $\mathcal{C}$-convex then $\mathbf{P}$ is $\mathcal{C}$-starshaped.*

As a first step in constructing a theory of abstract visibility we would like to at least prove that the *$\mathcal{C}$-kernel* of any set is $\mathcal{C}$-convex. To prove this we add two simple axioms to those of a convexity space to give what we call a visibility space.

Definition: Given a set, $\mathcal{S}$, and a family of subsets, $\mathcal{C}$, of $\mathcal{S}$, the structure $(\mathcal{S}, \mathcal{C})$ is said to be a ***visibility space*** if

1. $\emptyset, \mathcal{S} \in \mathcal{C}$;
2. $\forall C \subseteq \mathcal{C}, \quad \bigcap C \in \mathcal{C}$;
3. $\forall p \in \mathcal{S}, \quad \{p\} \in \mathcal{C}$;
4. $\forall p \in \mathcal{S}, \forall C \in \mathcal{C}, \quad \bigcup_{q \in C} \mathcal{C}\text{-}hull(\{p, q\}) \in \mathcal{C}$.

Thus we require that every point be $\mathcal{C}$-convex and that the cone created by a point and any $\mathcal{C}$-convex set be $\mathcal{C}$-convex.

We can now prove a useful result which characterizes $\mathcal{C}$-*kernel*(P) in visibility spaces as the intersection of the "largest" convex subsets of $\mathbf{P}$. From this result it will follow that $\mathcal{C}$-*kernel*(P) is $\mathcal{C}$-convex since it is the intersection of a non-empty collection of convex sets.

Definition: We say that the $\mathcal{C}$-convex set $\mathbf{Q}$ is *maximally $\mathcal{C}$-convex in* $\mathbf{P}$ if $\mathbf{Q} \subseteq \mathbf{P}$ and there does not exist a $\mathcal{C}$-convex set $\mathbf{R}$ such that $\mathbf{Q} \subset \mathbf{R} \subseteq \mathbf{P}$. We denote the set of all maximally $\mathcal{C}$-convex sets in $\mathbf{P}$ by $\mathcal{C}_{\mathbf{P}}$.

THEOREM **6.3.1** *Given a visibility space* $(\mathcal{S}, \mathcal{C})$,

$$\forall \mathbf{P} \subseteq \mathcal{S}\ ,\quad \mathcal{C}\text{-}kernel(\mathbf{P}) = \bigcap \mathcal{C}_{\mathbf{P}}$$

Hence, $\forall \mathbf{P} \subseteq \mathcal{S}$, $\mathcal{C}$-*kernel*(P) *is* $\mathcal{C}$-*convex.*

Proof: We first prove that $\mathcal{C}\text{-}kernel(\mathbf{P}) \subseteq \bigcap \mathcal{C}_{\mathbf{P}}$. If $\mathcal{C}\text{-}kernel(\mathbf{P}) = \emptyset$ then $\mathcal{C}\text{-}kernel(\mathbf{P}) \subseteq \bigcap \mathcal{C}_{\mathbf{P}}$ so suppose that $\mathcal{C}\text{-}kernel(\mathbf{P}) \neq \emptyset$.

Let p be in $\mathcal{C}$-*kernel*(P) and suppose that C_i is a maximally convex subset of $\mathbf{P}$ such that $p \notin C_i$. Since p is in $\mathcal{C}$-*kernel*(P), $\forall q \in C_i$ there exists a $C_j \in \mathcal{C}$ such that $p, q \in C_j$ and $C_j \subseteq \mathbf{P}$ (otherwise p cannot see q). Since $\mathcal{C}\text{-}hull(\{p,q\})$ is a subset of every $\mathcal{C}$-convex set which contains p and q then $\mathcal{C}\text{-}hull(\{p,q\}) \subseteq C_j \subseteq \mathbf{P}$. Therefore, $\bigcup_{q\in C_i} \mathcal{C}\text{-}hull(\{p,q\}) \subseteq \mathbf{P}$. But from axiom 4, $\bigcup_{q\in C_i} \mathcal{C}\text{-}hull(\{p,q\})$ is $\mathcal{C}$-convex. Also, $C_i \subset \bigcup_{q\in C_i} \mathcal{C}\text{-}hull(\{p,q\})$ since $p \notin C_i$. But C_i is maximal in $\mathbf{P}$, a contradiction. Therefore, no such C_i can exist and so, $p \in \mathcal{C}_{\mathbf{P}}$.

We now prove that $\bigcap \mathcal{C}_{\mathbf{P}} \subseteq \mathcal{C}\text{-}kernel(\mathbf{P})$. If $\bigcap \mathcal{C}_{\mathbf{P}} = \emptyset$ then $\bigcap \mathcal{C}_{\mathbf{P}} \subseteq \mathcal{C}\text{-}kernel(\mathbf{P})$ so suppose that $\bigcap \mathcal{C}_{\mathbf{P}} \neq \emptyset$.

Suppose that $p \in \mathcal{C}_{\mathbf{P}}$. Let q be in $\mathbf{P}$. From axiom 3 we know that $\{q\}$ is $\mathcal{C}$-convex, thus q is in at least one $\mathcal{C}$-convex set which is in $\mathbf{P}$. This implies that there is a maximally convex subset C_i in $\mathbf{P}$ such that $q \in C_i$. But $p \in C_i$ since p is in the intersection. Therefore, both p and q are in C_i and C_i is $\mathcal{C}$-convex and is in $\mathbf{P}$. Thus p sees q in $\mathbf{P}$ for all q in $\mathbf{P}$. Hence, $p \in \mathcal{C}\text{-}kernel(\mathbf{P})$. ∎

Corollary 6.3.2 $\forall \mathbf{P} \subseteq \mathcal{S}\ ,\quad \mathbf{P} = \mathcal{C}\text{-}kernel(\mathbf{P}) \iff \mathbf{P} = \mathcal{C}\text{-}hull(\mathbf{P}) \iff \mathbf{P} \in \mathcal{C}$

Observe that the strong $\mathcal{O}$-visibility space is a visibility space, as is the normal convexity space, however weak $\mathcal{O}$-visibility does not even give rise to a convexity space since the intersection of two weakly $\mathcal{O}$-visible sets is not necessarily weakly $\mathcal{O}$-visible.

Observe that our abstract definition of visibility does not imply that two points which see each other are in fact connected (in the conventional sense) unless convex sets (however they happen to be defined) are connected sets.

6.4 The Weak $\mathcal{O}$-Kernel

In the rest of this chapter we investigate two computational consequences of weak $\mathcal{O}$-visibility. From now on we restrict attention to weak $\mathcal{O}$-visibility and so by the *$\mathcal{O}$-kernel* we mean the weak $\mathcal{O}$-kernel and by *$\mathcal{O}$-visibility* we mean weak $\mathcal{O}$-visibility. We develop an algorithm to construct the $\mathcal{O}$-kernel of a polygon in this section and we generalize Chvátal's Art Gallery Theorem in the next section.

6.4.1 The 0°-*kernel* of a Polygon

We begin by giving an algorithm to compute the *$\mathcal{O}$-kernel* when $|\mathcal{O}| = 1$. We assume without loss of generality that $\mathcal{O} = \{0°\}$. In the next subsection we consider the case $|\mathcal{O}| \geq 2$.

Given a polygon $\mathbf{P}$, we wish to determine $0°\text{-}kernel(\mathbf{P})$. That is, the set of all points which see every point in $\mathbf{P}$ via 0°-stairsegments which lie in $\mathbf{P}$ (figure 6.3(a) gives a example $0°$-*kernel*). Recall that a $0°$-stairsegment a $0°$-convex curve, that is, a curve which is monotone with respect to any vertical line. We develop an algorithm to find $0°$-*kernel*(P) in linear time and space.

Consider any reflex angle in **P** at a vertex v_i where v_{i-1} and v_{i+1} are either both above or both below v_i (see figure 6.2). These vertices create two types of "walls"—*up-walls* and *down-walls* (in the figure the first wall is an up-wall).

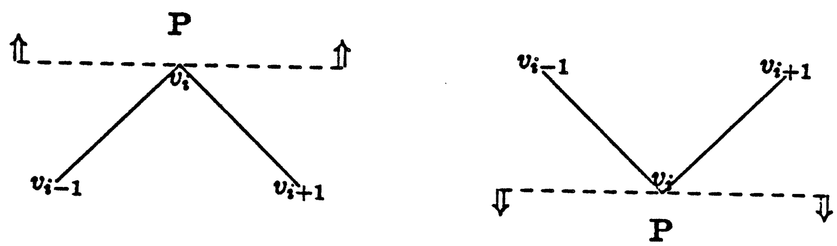

Figure 6.2: Walls Created By Interior Angles

Lemma 6.4.1 *Given a polygon* **P**, $0°$-*kernel*(**P**) *lies above any up-wall and below any down-wall.*

Proof: Consider an up-wall $\mathcal{L}$. If $0°$-*kernel*(**P**) lies below $\mathcal{L}$ then either v_{i-1} or v_{i+1} cannot be seen from $0°$-*kernel*(**P**) since any curve crossing $\mathcal{L}$ twice cannot be $0°$-convex. ∎

To find $0°$-*kernel*(**P**) we proceed as follows. In linear time we find y_{max} and y_{min}—the maximum and minimum y-values achieved by the vertices of **P**. y_{max} and y_{min} are the initial boundaries of $0°$-*kernel*(**P**). Scan the vertices of **P** starting from any vertex and create the walls (if any). As each new vertex is examined it may create an up-wall or a down-wall. If it creates an up-wall (say), check whether it is higher than y_{min}. If it is, update y_{min} and check that y_{min} is still lower than y_{max}. If not then $0°$-*kernel*(**P**) is empty otherwise we will end up with a range of y-values $[y_{min}, y_{max}]$ in which $0°$-*kernel*(**P**) must lie. It is easy to show that since $[y_{min}, y_{max}]$ is above every up-wall and below every down-wall (by construction), any point in **P** in this range is in $0°$-*kernel*(**P**). Thus to find the

$0°$-*kernel* we need merely scan $\mathbf{P}$ testing each vertex against $[y_{min}, y_{max}]$. Since a polygonal description of the $0°$-*kernel* is recoverable from $[y_{min}, y_{max}]$ in linear time we can assume that all $0°$-*kernels* are described by ranges. Since we must examine every vertex in $\mathbf{P}$ we see that:

THEOREM **6.4.1** *The* $0°$*-kernel of an* n*-vertex polygon can be found in* $\Theta(n)$ *time and* $\Theta(n)$ *space.*

6.4.2 The $\mathcal{O}$-*kernel* of a Polygon

In this section we show that the $0°$-*kernel* finding algorithm outlined in the previous section can be used to find the $\mathcal{O}$-*kernel* for any $\mathcal{O}$. To demonstrate this we need the following result:

THEOREM **6.4.2** *Let* $\mathbf{P}$ *be a simple polygon.*

Then

$$(\bigcup_i \mathcal{O}_i)\text{-}kernel(\mathbf{P}) = \bigcap_i \mathcal{O}_i\text{-}kernel(\mathbf{P})$$

Proof: Let p be in $(\bigcup_i \mathcal{O}_i)$-*kernel*$(\mathbf{P})$. For each q in $\mathbf{P}$, p sees q via a $(\bigcup_i \mathcal{O}_i)$-stairsegment lying in $\mathbf{P}$. Each such curve is $(\bigcup_i \mathcal{O}_i)$-convex and hence $\mathcal{O}_i$-convex for each i. Hence, for each q, p sees q via an $\mathcal{O}_i$-stairsegment lying in $\mathbf{P}$. Thus, for all i, $p \in \mathcal{O}_i$-*kernel*$(\mathbf{P})$. Therefore, $(\bigcup_i \mathcal{O}_i)$-*kernel*$(\mathbf{P}) \subseteq \bigcap_i \mathcal{O}_i$-*kernel*$(\mathbf{P})$.

Conversely, let p be in $\bigcap_i \mathcal{O}_i$-*kernel*$(\mathbf{P})$. For each q in $\mathbf{P}$, and for each i, p sees q via an $\mathcal{O}_i$-stairsegment lying in $\mathbf{P}$. We wish to show that we can construct a $(\bigcup_i \mathcal{O}_i)$-stairsegment lying wholly in $\mathbf{P}$ connecting p and q. Observe that for each $\theta \in \bigcup_i \mathcal{O}_i$ p must be able to see every point in $\mathbf{P}$ θ-convexly (that is, for each q in $\mathbf{P}$ there must exist some θ-convex curve lying wholly in $\mathbf{P}$ connecting p and q).

If $\Theta(\mathcal{LS}[p,q]) \in \bigcup_i \mathcal{O}_i$ then from lemma 4.3.1 $(\bigcup_i \mathcal{O}_i)$-$\|[p,q] = \mathcal{LS}[p,q]$. Suppose that $\Theta(\mathcal{LS}[p,q]) \in \mathcal{O}_j$. $\mathcal{LS}[p,q]$ is the only $\mathcal{O}_j$-convex curve connecting p and q. Thus, since

p sees q via an O_j-stairsegment lying in **P**, p sees q via a $(\bigcup_i O_i)$-stairsegment lying in **P**, namely $\mathcal{LS}[p,q]$. Therefore, suppose that $\Theta(\mathcal{LS}[p,q]) \notin \bigcup_i O_i$ and let (θ_1, θ_2) be $\mathcal{LS}[p,q]$'s maximal $(\bigcup_i O_i)$-free range. Since (θ_1, θ_2) is $(\bigcup_i O_i)$-free then, from lemma 4.2.1, to prove the claim it is sufficient to construct a $\{\theta_1, \theta_2\}$-stairsegment lying wholly in **P** connecting p and q. If one of the arms of $\{\theta_1, \theta_2\}$-$||[p,q]$ does not intersect the boundary of **P** then that arm is a $\{\theta_1, \theta_2\}$-stairsegment lying wholly in **P** connecting p and q. Therefore, suppose that both of the arms of $\{\theta_1, \theta_2\}$-$||[p,q]$ intersect the boundary of **P**.

Assume, without loss of generality, that $[\theta_1, \theta_2] = [0°, 90°]$. The following arguments hinge on the fact that p must be able to see every point in **P** 0°-convexly and 90°-convexly (otherwise $p \notin \bigcap_i$ O-*kernel*(**P**)). Call the four line segments making up $\{0°, 90°\}$-$||[p,q]$ *top, bottom, left and right.* Suppose that *bottom* intersects **P**'s boundary (if this isn't the case then *right* must intersect **P**'s boundary and a similar argument to that given below applies).

Let p_b be the first (leftmost) point on *bottom* which is on **P**'s boundary. Let S_1 be the boundary of **P** starting from p_b that lies in $\{0°, 90°\}$-$||[p,q]$. Since **P** is simple and p and q are in **P** then S_1 must exit $\{0°, 90°\}$-$||[p,q]$ at some point. If S_1 intersects *top* then p cannot see q 0°-convexly. If S_1 intersects *left* then either p cannot see q 0°-convexly or p cannot see q 90°-convexly. If S_1 intersects *bottom* at some point $r \neq p_b$ then p cannot see r 0°-convexly. Thus S_1 must exit $\{0°, 90°\}$-$||[p,q]$ through *right.* Further, S_1 must be 0°-convex otherwise there is a point on S_1 which p cannot see 0°-convexly.

Construct the curve S_2 from S_1 by replacing each segment of S_1 which is non-90°-convex by a vertical line segment (that is, form 90°-*hull*(S_1), here however we are only interested in the portion of the hull on the same side as p). We claim that S_2 lies wholly in **P**, that is, S_2 does not intersect **P**'s boundary at any point not on S_1. Let S' be any boundary segment of **P** in $\{0°, 90°\}$-$||[p,q]$ not contained in S_1. Suppose that S' intersects S_2 at some point r. Since **P** is simple r cannot be on S_1, therefore r lies on one of the vertical line segments introduced to form the "hull" of S_1. If S' enters and exits $\{0°, 90°\}$-$||[p,q]$ through *top* then there is a point that p cannot see 90°-convexly (either the entry point or the exit point). S' cannot enter or exit $\{0°, 90°\}$-$\|[p, q]$ through *right* for either S' is a subset of S_1 or p cannot see q 90°-convexly. Thus S' enters through *left*. Further, S' must be 90°-convex otherwise there is a point which p cannot see 90°-convexly. Since S' is be 90°-convex and r is on S', S' intersects S_1. But $\mathbf{P}$ is simple, a contradiction. Therefore, no S' can intersect S_2 and so S_2 lies wholly in $\mathbf{P}$.

Therefore, the curve S_2 is connected, it lies wholly in $\mathbf{P}$, it connects p and q and it is $\{0°, 90°\}$-convex. Since $(0°, 90°)$ is $(\bigcup_i O_i)$-free S_1 is a $(\bigcup_i O_i)$-stairsegment connecting p and q in $\mathbf{P}$. Therefore, $p \in (\bigcup_i O_i)$-*kernel*$(\mathbf{P})$. ∎

Thus, we may generalize the previous algorithm to find the O-*kernel* of an n-vertex polygon in $O(nr + r \lg r)$ time and $O(n)$ space where O consists of r ranges. We require $O(nr)$ time to find the $2r$ separate kernels and $O(r \lg r)$ time to find their intersection. This time bound reduces to $O(n)$ when r is constant, and this compares favourably with Lee and Preparata's linear time kernel algorithm ([34]). However, if r is an input parameter then we would prefer to have an $O(n \lg r + r)$ algorithm and we conjecture that such an algorithm exists.

In figure 6.3 we provide some example O-kernels: (a) is the 0°-*kernel*; (b) is the 90°-*kernel*; (c) is the $\{0°, 90°\}$-*kernel*; (d) is the $[0°, 90°]$-*kernel*; (e) is the $[90°, 180°]$-*kernel* and (f) is the $[0°, 180°)$-*kernel*.

6.4.3 The 0°-*kernel* of a Polygon with Holes

In lemma 6.4.3 below we examine the possible effects of holes on the 0°-*kernel*. First though we need the following result:

Lemma 6.4.2 *Let* $\mathbf{P}$ *be a polygon with at least one hole and let* p *be in* $\mathbf{P}$.

If $\mathcal{L}[0°, p]$ *intersects a hole then* $p \notin 0°$-*kernel*$(\mathbf{P})$.

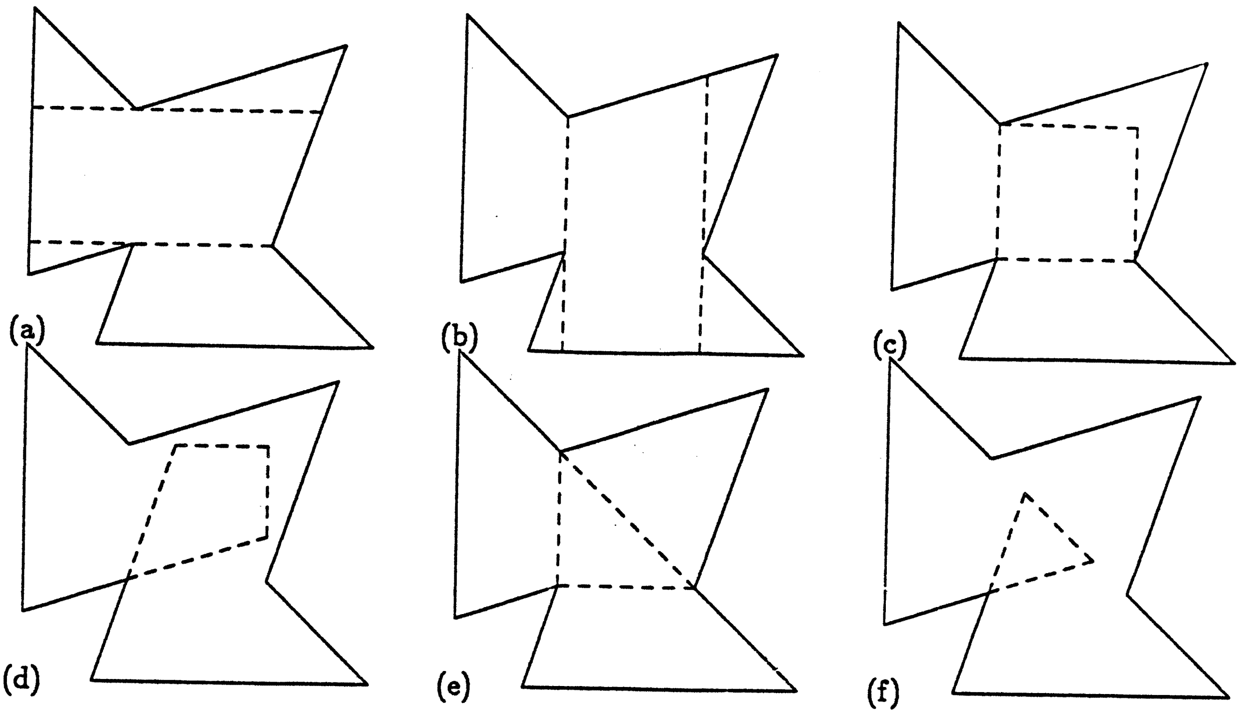

Figure 6.3: Some Example $\mathcal{O}$-Kernels of a Polygon

Proof: Suppose that $p \in \mathbf{P}$ and $\mathcal{L}[0°, p]$ intersects a hole. $\mathcal{L}[0°, p]$ must cut $\mathbf{P}$ at some point q on the opposite side of the hole to p. Since p and q lie on a horizontal line, from lemma 4.3.1 $0°\text{-}\|[p, q] = \mathcal{LS}[p, q]$, and thus, no 0°-convex curve lying wholly in $\mathbf{P}$ can connect p and q. Thus, p cannot be in 0°-*kernel*($\mathbf{P}$). ∎

Holes can affect the 0°-*kernel* beyond deleting a subrange from the range defining the 0°-*kernel* of the "hole-less" version of the polygon. Call a hole 0°-*convex* if it is 0°-convex when viewed as a polygon in its own right. In figure 6.4 we give three examples to illustrate what happens when the hole is and is not 0°-convex. The hole in the first polygon is non-0°-convex and we think of it as generating an up-wall and a down-wall. However, the up-wall is above the down-wall and so the 0°-*kernel* of the polygon as a whole is empty (no single point can see both indentations 0°-convexly). The hole in the second polygon is also non-0°-convex, however in this case the hole only generates an up-wall. Thus the 0°-*kernel* may lie above this up-wall. Finally, the hole in the third polygon is 0°-convex and it generates neither an up-wall nor a down-wall. Of course, all three holes always eliminate (at least) the subrange they define from the range defining the 0°-*kernel* of the hole-less polygon.

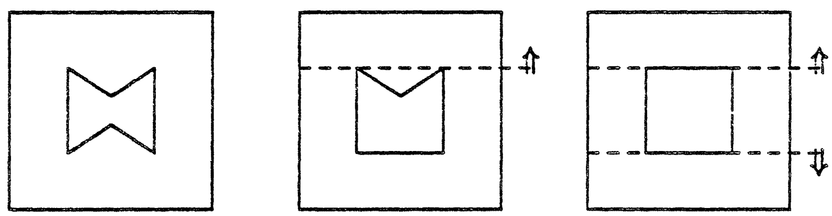

Figure 6.4: The Effect of Non-0°-Convex and 0°-Convex Holes

Lemma 6.4.3 *If* P *is a polygon with at least one hole then* 0°-*kernel*(P) *may be disconnected.*

Proof: If the 0°-*kernel* of the hole-less version of P is empty then 0°-*kernel*(P) is empty, so assume that the 0°-*kernel* of the hole-less version of P is non-empty and let its range be $[y_{min}, y_{max}]$. Suppose that P has a hole lying in the range $[y_1, y_2]$.

From lemma 6.4.2 we know that no point in P in the range $[y_1, y_2]$ can lie in 0°-*kernel*(P). There are, essentially, four cases to examine with respect to $[y_1, y_2]$ and $[y_{min}, y_{max}]$ (we illustrate the cases in figure 6.5):

Case 1: $y_{min} < y_1 < y_2 < y_{max}$

In this case 0°-*kernel*(P) can only lie in the ranges $[y_{min}, y_1)$ and $(y_2, y_{max}]$. If the hole is 0°-convex then this is its only effect on 0°-*kernel*(P). If the hole is non-0°-convex then it may eliminate either $[y_{min}, y_1)$, $(y_2, y_{max}]$, or both, depending on whether it generates an up-wall, a down-wall, or both.

Case 2: $y_1 < y_{min} < y_{max} < y_2$.

In this case 0°-*kernel*(P) is empty.

Case 3: $y_1 < y_{min} < y_2 < y_{max}$ (and similarly, $y_{min} < y_1 < y_{max} < y_2$).

In this case $0°$-*kernel*(P) can only lie in the range $(y_2, y_{max}]$ (respectively, $[y_{min}, y_1)$). If the hole is $0°$-convex then this is its only effect on $0°$-*kernel*(P). If the hole is non-$0°$-convex then it may eliminate $(y_2, y_{max}]$ (respectively, $[y_{min}, y_1)$), if it generates a down-wall (respectively, an up-wall).

Case 4: $y_1 < y_2 < y_{min} < y_{max}$ (and similarly, $y_{min} < y_{max} < y_1 < y_2$).

In this case $0°$-*kernel*(P) can still lie in the range $[y_{min}, y_{max}]$. If the hole is $0°$-convex then it has no effect on $0°$-*kernel*(P). If the hole does not create a down-wall (respectively, an up-wall) then it also has no effect on $0°$-*kernel*(P). If the hole creates a down-wall (respectively, an up-wall) then $0°$-*kernel*(P) is empty. ∎

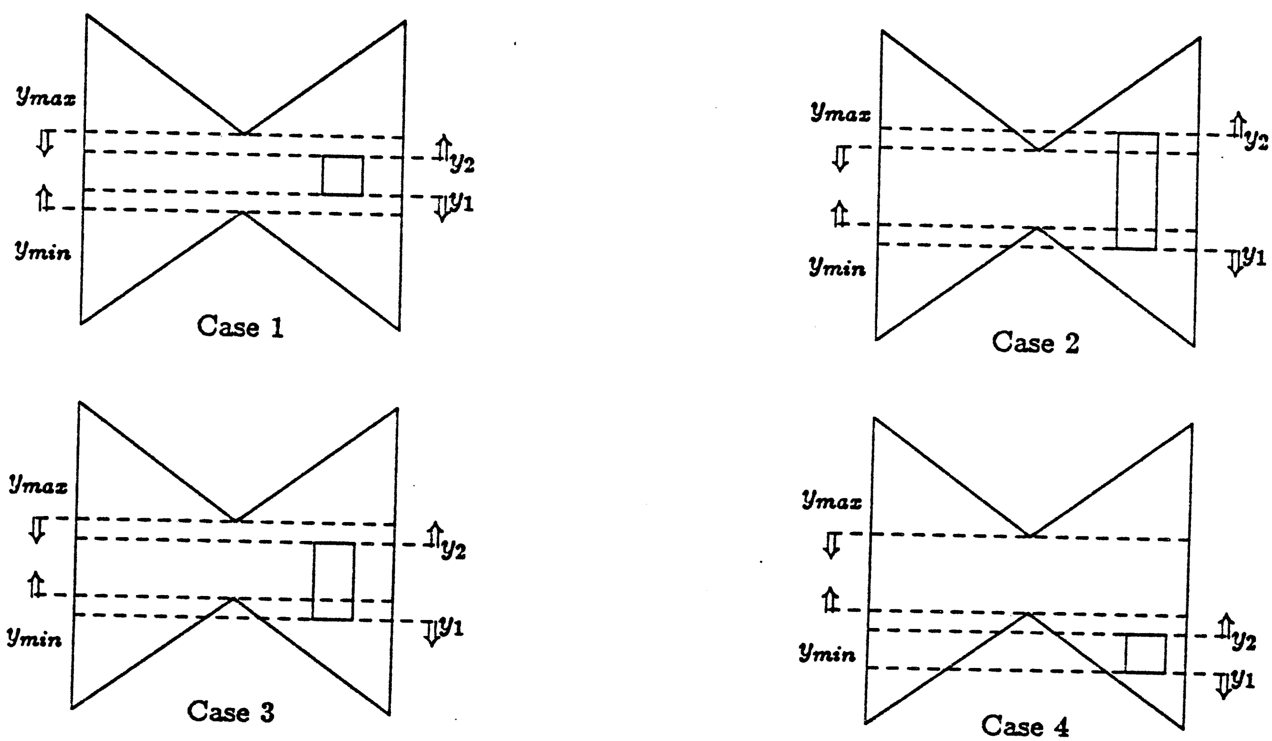

Figure 6.5: The $0°$-*kernel* of a Polygon With Holes May Be Empty, Connected or Disconnected

We can find the $0°$-*kernel* of a polygon with holes by first finding the $0°$-*kernel* of the hole-less version. This gives us a single range of y values over which the $0°$-*kernel* lies (if it is empty we are done). Now we process the holes to determine the ranges they eliminate from the allowed range. This takes linear time. Each hole eliminates some subrange of y values over which the $0°$-*kernel* may lie. Each such elimination can be accomplished in $O(\lg n)$ time by performing a binary search into the $O(n)$ remaining subranges. We now process each hole to determine if any of them is non-$0°$-convex. In chapter 7 we show that this can be accomplished in linear time. If none of the holes create any walls we are done. If the walls they create move past each other then the $0°$-*kernel* is empty, otherwise the $0°$-*kernel* is unchanged. Thus the set of y ranges over which the $0°$-*kernel* is defined can be found in $O(n \lg n)$ time.

Lemma 6.4.4 *Finding the $0°$-kernel of an n-vertex polygon with holes requires $\Omega(n \lg n)$ time in the algebraic decision tree model of computation.*

Proof: Suppose that $\Omega(n \lg n)$ is not a lower bound on finding the $0°$-*kernel* of a polygon with holes. We show how to solve the Element Uniqueness problem in $o(n \lg n)$ time, a contradiction ([11]).

Given n values $y_1, y_2, \ldots, y_n$, we construct a polygon with $4(n+1)$ vertices containing n *point holes* as follows. The i^{th} hole will be the point (i, y_i). Find the minimum and maximum values in y and construct an appropriate surrounding box (see figure 6.6 in which the point holes have been expanded to squares for clarity). All of these operations take linear time. Run the $0°$-*kernel* algorithm and count, in linear time, how many components compose the $0°$-*kernel*. If the number of components is less than $n+1$ then some pair of y coordinates are equal otherwise they are all distinct. ∎

THEOREM **6.4.3** *The $0°$-kernel of an n-vertex polygon with holes can be found in $\Theta(n \lg n)$ time and $\Theta(n)$ space.*

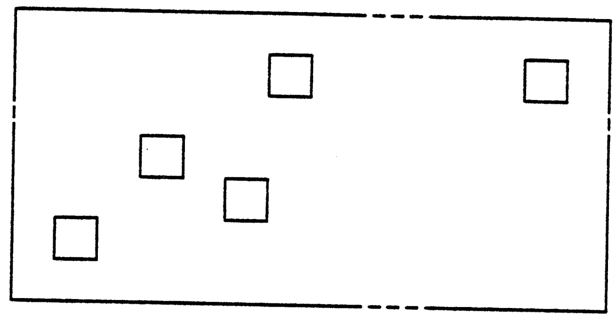

Figure 6.6: Finding The 0°-*kernel* of a Polygon with Holes Requires $O(n \lg n)$ Time

6.4.4 The $\mathcal{O}$-*kernel* of a Polygon with Holes

The $\{0°, 90°\}$-*kernel* of a polygon with holes may have $\Omega(n^2)$ components if $|\mathcal{O}| \geq 2$ and $\mathcal{O}$ is not the set of all orientations (for example see figure 6.7 where $\mathcal{O} = \{0°, 90°\}$). Thus, just to report each component of the $\mathcal{O}$-*kernel* requires $\Omega(n^2)$ time in the worst case. The kernel with respect to each range of orientations in $\mathcal{O}$ can be found in $O(n \lg n)$ time since we need only find the intersection of the kernels with respect to the range boundaries. If there are r ranges in $\mathcal{O}$ then the intersection of their kernels can be found in $O(r^2 n^2)$ time, thus the $\mathcal{O}$-*kernel* of a polygon with holes can be found in $O(r^2 n^2)$ time in the worst case.

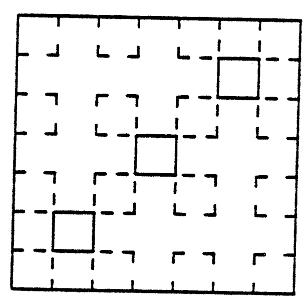

Figure 6.7: The $\mathcal{O}$-*kernel* of a Polygon With Holes May Have $\Omega(n^2)$ Components

6.5 The Art Gallery Problem

Given a polygon **P**, whose edges represent an art gallery, and a set of points (representing art gallery guards) in **P**, **P** is said to be *guarded* if for each point in **P** there is at least one guard which sees that point via a line. The Art Gallery Problem is: How many guards are necessary and sufficient to guard any n-vertex polygon?

In [7] Chvátal proved that $\lfloor n/3 \rfloor$ guards are necessary and sufficient to guard any simple n-vertex polygon and Kahn *et al.* ([26]) showed that $\lfloor n/4 \rfloor$ guards are necessary and sufficient to guard any *orthogonal* polygon. The latter result holds for any class of $\mathcal{O}$-polygons for which $|\mathcal{O}| = 2$. If $|\mathcal{O}| \geq 3$ then there exist $\mathcal{O}$-polygons which require $\lfloor n/3 \rfloor$ guards (see figure 6.8 in which $\{0°, 45°, 135°\} \subseteq \mathcal{O}$).

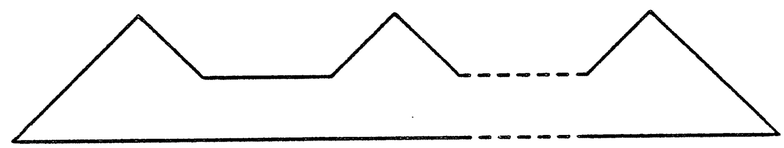

Figure 6.8: If $|\mathcal{O}| \geq 3$ $\mathcal{O}$-Polygons Require $\lfloor n/3 \rfloor$ Guards

In the above mentioned results the way in which the guards "see" is kept constant while the class of polygons to be guarded is changed. We now examine the effect of modifying the way the guards see as well as allowing the class of polygons to vary.

Definition: A guard which can see along any $\mathcal{O}$-stairline is said to be an *$\mathcal{O}$-guard*.

Of course, when $\mathcal{O}$ is the set of all orientations $\mathcal{O}$-guards are guards in the usual sense. Reckhow and Culberson ([52]) have, independently, considered the special case of $\mathcal{O}$-guards in $\mathcal{O}$-polygons when $\mathcal{O} = \{0°, 90°\}$. We call such $\mathcal{O}$-guards, *orthogonal guards*. A number of other variants of the art gallery problem have been defined: for example, O'Rourke [42] considers mobile guards, and Keil [28] considers guards which see via rectangles instead of lines. [1]

[1]As for $\mathcal{O}$-visibility, we may define *strong $\mathcal{O}$-guards* to be guards which see along any $\mathcal{O}$-parallelogram.

6.5.1 Orthogonal Guards in Orthogonal Polygons

Definition: Given an orthogonal polygon **P**, an edge of **P** is said to be a *north edge* of **P** if the interior of **P** lies locally below the edge. Similarly, we define *south, east* and *west* edges of **P**. A convex vertex of **P** is said to be a *NE vertex* of **P** if it is made by a north and east edge of **P**. Similarly, we define *NW*, *SE* and *SW* vertices of **P**.

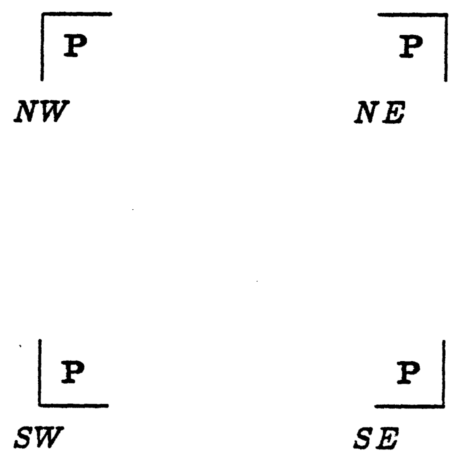

Figure 6.9: Kinds of Convex Interior Angles for Orthogonal Polygons

Lemma 6.5.1 *Every point in an orthogonal polygon sees at least one NE, NW, SE and SW vertex via* $\{0°, 90°\}$*-stairsegments.*

Proof: Let **P** be an orthogonal polygon and let p be a point in **P**. We shall prove the result only for NE vertices, a similar construction works for any of the other three convex vertices.

If p is at a NE vertex we are done. If p is on an east edge of **P** then travel up the edge until the next vertex is reached. If this vertex is a NE vertex we are done, otherwise we are at a reflex vertex. This brings us to the general case: p is somewhere in the interior of **P**.

(In this framework, the "O-guards" defined above would be *weak O-guards*.) Observe that Keil's ([28]) "rectangle guards" would be the same as strong orthogonal guards. It is possible to show that there exist polygons which require infinitely many strong O-guards once O is not the set of all orientations.

Given that p is in the interior of $\mathbf{P}$ construct the following polygonal curve. p can see only one east edge, call it e_1. Take the horizontal line segment from p to e_1 as the first edge in the curve. If this terminates at a NE vertex then the curve is complete, otherwise this new point on e_1 sees one north edge, call it n_1. Add the vertical line segment to n_1 to the curve.

The intersection point on n_1 can see only one east edge. Again add the horizontal line segment to this new edge to the curve and so on inductively. Since $\mathbf{P}$ is a polygon it has only a finite number of edges and we will eventually exhaust all of them so this process must stop. At this point we are either facing east and are blocked by a north edge or we are facing north and are blocked by an east edge. In both cases we are at a NE vertex. The curve we have constructed consists of alternating horizontal and vertical line segments and is horizontally and vertically monotone thus it is a $\{0°, 90°\}$-stairsegment. Thus any point in an orthogonal polygon can see a NE vertex via a $\{0°, 90°\}$-stairsegment. ■

THEOREM **6.5.1** *$\lfloor (n+4)/8 \rfloor$ 0°- or $\{0°, 90°\}$-guards are necessary and sufficient to guard an orthogonal art gallery.*

Proof: From the previous lemma it suffices to place guards at all of the NE, NW, SE or SW convex vertices. We first determine which assignment gives the minimum number of guards and assign guards to those convex vertices. By counting interior angles it is easy to show that in an orthogonal polygon the number of convex vertices is $(n+4)/2$. Thus the number of $\{0°, 90°\}$-guards placed is at most $\lfloor (n+4)/8 \rfloor$.

Since a $\{0°, 90°\}$-stairsegment is a 0°-stairsegment this number of 0°-guards suffices as well.

The necessity of this many guards in the worst case follows from the class of polygons illustrated in figure 6.10(a). ■

Thus it suffices to place the 0°-guards or $\{0°, 90°\}$-guards only at convex vertices. This result motivates us to ask whether a polygon can be guarded by placing $\mathcal{O}$-guards only at

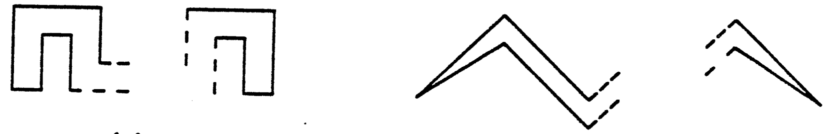

(a) (b)

Figure 6.10: Orthogonal And Arbitrary Snakes

convex vertices? The counterexample in figure 6.11(a) shows that this is not possible if $|\mathcal{O}| \geq 2$ (in the figure $\{0°, 90°\} \subseteq \mathcal{O}$ and no $\mathcal{O}$-guard placed at a convex vertex can see p via a $\{0°, 90°\}$-stairsegment). Even if we restrict ourselves to $\mathcal{O}$-polygons the counterexample in figure 6.11(b) shows that this is not possible if $|\mathcal{O}| \geq 3$ (in the figure $\{0°, 45°, 90°\} \subseteq \mathcal{O}$ and no $\mathcal{O}$-guard placed at a convex vertex can see q via a $\{0°, 45°, 90°\}$-stairsegment).

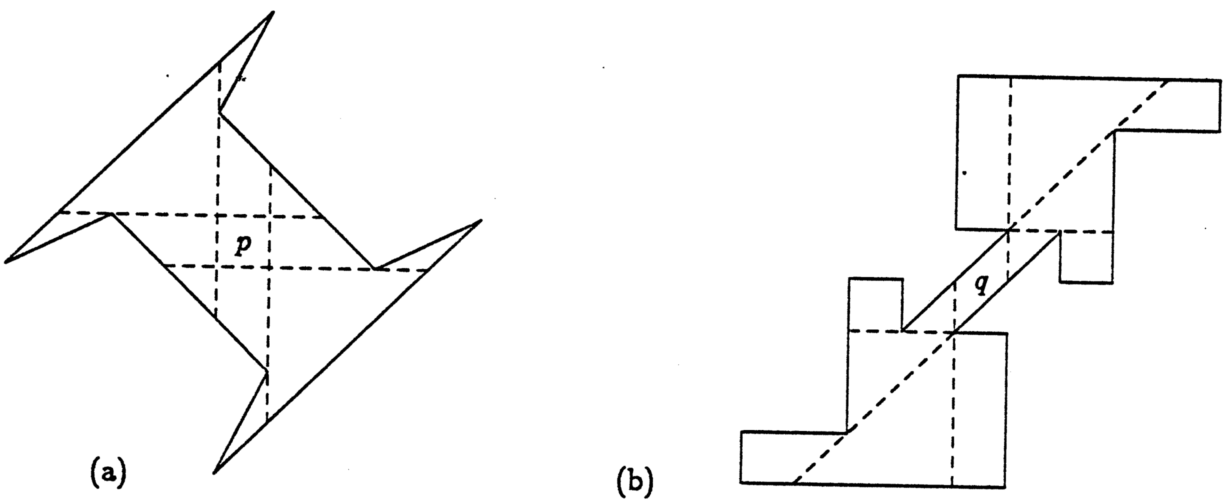

(a) (b)

Figure 6.11: Counterexamples to Placing $\mathcal{O}$-Guards at Convex Vertices

6.5.2 The Art Gallery Theorems

THEOREM 6.5.2 **(The Orthogonal Art Gallery Theorem)** *If $|\mathcal{O}| = 0$ then 1 $\mathcal{O}$-guard is necessary and sufficient to guard an orthogonal art gallery.*

If $|\mathcal{O}| \geq 1$ then $\lfloor n/4 \rfloor$ $\mathcal{O}$-guards are necessary and sufficient to guard an orthogonal art gallery.

Proof: $|\mathcal{O}| = 0$ is trivial.

$|\mathcal{O}| \geq 1$: Since line-visible points are $\mathcal{O}$-visible for all $\mathcal{O}$ sufficiency follows from the quadrilateralization argument of Kahn *et al.* ([26]). Necessity follows from the class of polygons illustrated in figure 6.12 (here $45° \in \mathcal{O}$). ∎

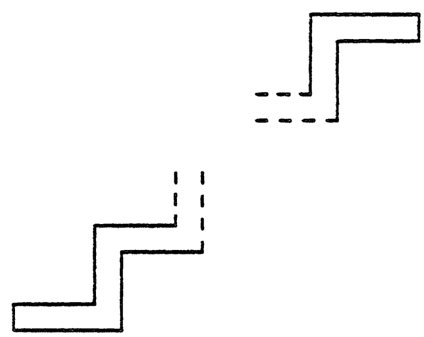

Figure 6.12: If $|\mathcal{O}| \geq 1$ Orthogonal Art Galleries Require $\lfloor n/4 \rfloor$ $\mathcal{O}$-Guards

THEOREM **6.5.3 (The Art Gallery Theorem)** *If $|\mathcal{O}| = 0$ then 1 $\mathcal{O}$-guard is necessary and sufficient to guard an art gallery.*

If $|\mathcal{O}| = 1$ then $\lfloor (n+2)/4 \rfloor$ $\mathcal{O}$-guards are necessary and sufficient to guard an art gallery.

If $|\mathcal{O}| \geq 2$ then $\lfloor n/3 \rfloor$ $\mathcal{O}$-guards are necessary and sufficient to guard an art gallery.

Proof: $|\mathcal{O}| = 0$ is trivial.

$|\mathcal{O}| = 1$: Sufficiency follows from an argument similar to placing 0°-guards at convex vertices in orthogonal polygons. There are a total of at most $(n+2)/2$ vertices which are convex with respect to any fixed direction. The $\mathcal{O}$-guards are placed at the convex vertices which give the smaller number. Necessity follows from the class of polygons illustrated in figure 6.10(b) (here $0° \in \mathcal{O}$).

$|\mathcal{O}| \geq 2$: Since line-visible points are $\mathcal{O}$-visible for all $\mathcal{O}$ sufficiency follows from the triangulation argument of Chvátal ([7]). Necessity follows from the class of polygons illustrated in figure 6.13 (here $\{0°, 90°\} \subseteq \mathcal{O}$). ∎

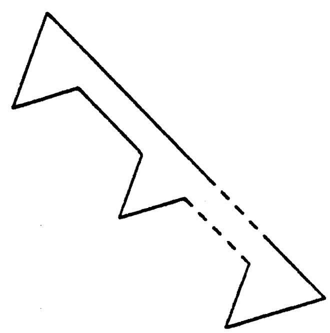

Figure 6.13: If $|\mathcal{O}| \geq 2$ Art Galleries Require $\lfloor n/3 \rfloor$ $\mathcal{O}$-Guards

Chapter 7

Computing the Hulls

In this chapter we give optimal algorithms to find the $\mathcal{O}$-*hull* and the strong $\mathcal{O}$-*hull* of a simple polygon. The algorithms have worst case time complexity $\Theta(n+r)$ and use $\Theta(n)$ space where n is the number of edges of the polygon and r is the number of ranges in $\mathcal{O}$.

7.1 The Three $\mathcal{O}$-*hulls*

So far we have seen *hull*(P) (the convex hull of P), $\mathcal{O}$-*hull*(P), the weak $\mathcal{O}$-*hull* of P and the strong $\mathcal{O}$-*hull* of P. In [62] Widmayer *et al.* investigated $\mathcal{O}$-*oriented* $\mathcal{O}$-convex polygons (for finite $\mathcal{O}$). They were interested in this class of polygons since, for their application area (VLSI design), not only must all the polygons under consideration be $\mathcal{O}$-oriented but also, the only reasonable definition of "convexity" is $\mathcal{O}$-convexity. The reason is that in VLSI design not only are the polygons constrained to be $\mathcal{O}$-oriented but so are the wires connecting them. It is worth noting that all previous work on orthogonal convexity ([41,39,43]) also assumed that an orthogonally convex set should be orthogonal.

If we restrict ourselves to considering only $\mathcal{O}$-oriented polygons then the following definition of the $\mathcal{O}$-*hull* is a natural one:

Definition: Given two polygons P and Q, Q is said to be the *$\mathcal{O}$-oriented $\mathcal{O}$-hull* of P ($\mathcal{O}\mathcal{O}$-*hull*(P)) if Q is the smallest $\mathcal{O}$-convex $\mathcal{O}$-polygon which contains P.

As we shall see, there is no reason to define $\mathcal{O}\mathcal{O}$-*hull*(P) since in most cases it is equivalent to the less restrictive $\mathcal{O}$-*hull*(P); further unlike $\mathcal{O}$-*hull*(P), it does not always exist. Notice that since polygons are connected sets, from theorem 6.1.2 the weak $\mathcal{O}$-*hull* of P is just $\mathcal{O}$-*hull*(P).

Recall that in theorem 6.1.3 we showed that the strong $\mathcal{O}$-*hull* is convex and $\mathcal{O}$-oriented. Thus it is a natural generalization of the bounding box of a polygon, that is, the smallest orthogonal rectangle which encloses the polygon. Bounding boxes are currently in use in probabilistic tests of intersection: to test whether two objects intersect find their bounding boxes and test whether they intersect, if they don't then the objects don't, otherwise test the two objects by brute force. Of course, there is nothing sacred about the x and y directions, we may expand the bounding box test to first pick some number of random orientations and find the bounding boxes, that is, the strong $\mathcal{O}$-*hull*s, of the two objects for these orientations then test them for intersection.

Figures 7.1 and 7.2 show the various hulls for two example polygons where $\mathcal{O} = \{0°, 45°, 90°\}$. In figure 7.1 the polygon is $\mathcal{O}$-oriented, whereas in figure 7.2 it is not. Observe that the $\mathcal{O}\mathcal{O}$-*hull* of the non-$\mathcal{O}$-oriented polygon (P_2) does not exist.

7.2 Characterizing $\mathcal{O}$-Convex Polygons

From the boundary characterization result of chapter 5 (theorem 5.5.1) we may immediately state the following special cases:

THEOREM 7.2.1 *A polygon is $\mathcal{O}$-convex if and only if its boundary consists of a sequence of maximal polygonal $\mathcal{O}$-stairsegments meeting at $\mathcal{O}$-extremal points.*

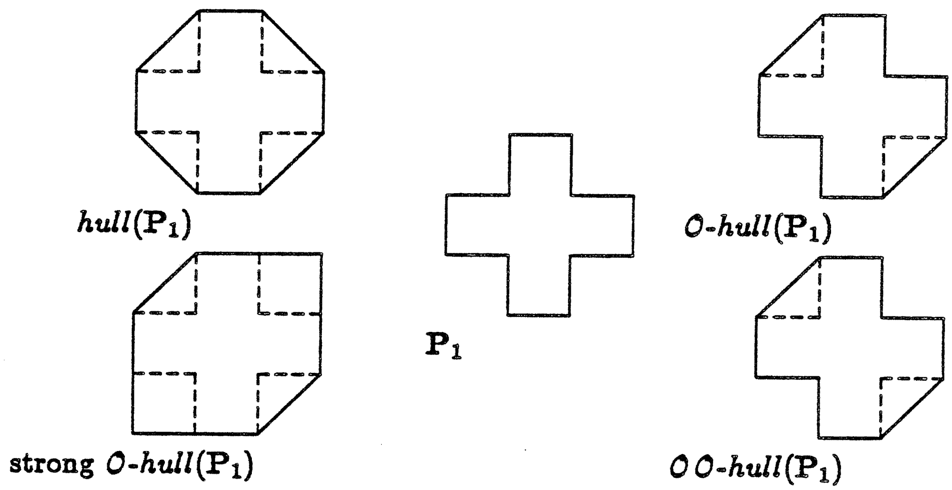

Figure 7.1: An $\mathcal{O}$-Polygon and its Various Hulls

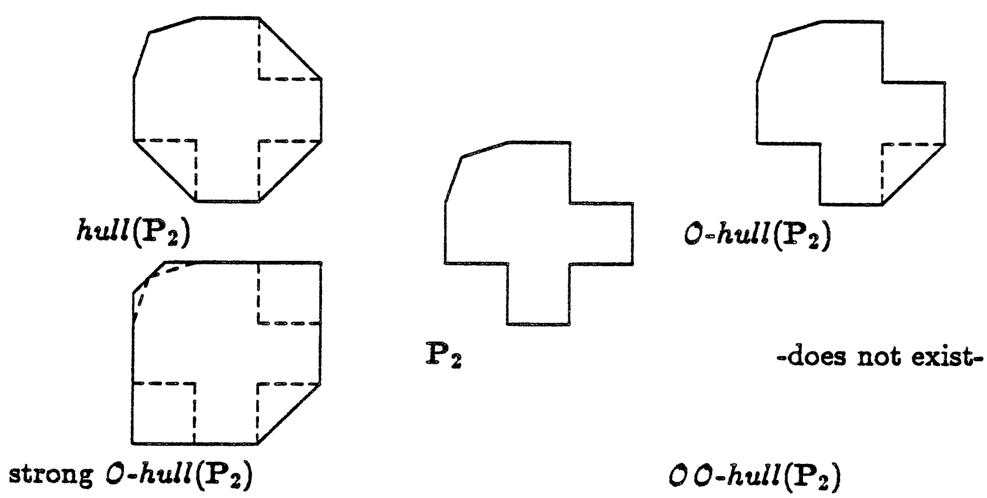

Figure 7.2: An Arbitrary Polygon and its Various Hulls

Corollary 7.2.1 *An $\mathcal{O}$-polygon is $\mathcal{O}$-convex if and only if its boundary consists of a sequence of maximal $\mathcal{O}$-staircases meeting at $\mathcal{O}$-extremal points.*

For finite $\mathcal{O}$ corollary 7.2.1 has been stated in [62]. The characterization of the boundary of $\mathcal{O}$-convex polygons as a sequence of maximal polygonal $\mathcal{O}$-stairsegments is a direct generalization of the case for orthogonal polygons (see for example Wood [64]).

Theorem 7.2.1 enables us to prove:

THEOREM **7.2.2** *If the ranges in $\mathcal{O}$ are given in sorted order, then testing a polygon* $\mathbf{P}$ *for $\mathcal{O}$-convexity and strong $\mathcal{O}$-convexity can be accomplished in $\Theta(n+r)$ time (worst case) and $\Theta(n)$ space, where n is the number of edges of* $\mathbf{P}$ *and r is the number of ranges in $\mathcal{O}$.*

Proof: We first consider testing for strong $\mathcal{O}$-convexity. From theorem 6.1.3 we know that a strong $\mathcal{O}$-convex set is convex and $\mathcal{O}$-oriented. Thus to test $\mathbf{P}$ for strong $\mathcal{O}$-convexity we check whether $\mathbf{P}$ is convex then check whether it is $\mathcal{O}$-oriented. Both tests may be done in linear time as follows. First find the range in $\mathcal{O}$ to which one of $\mathbf{P}$'s edges belong, this takes $O(\lg r)$ time if $\mathcal{O}$ has r ranges. Then from the next edge in $\mathbf{P}$ determine the order in which to march along $\mathcal{O}$ (there are only two possible orders). Having found this examine each new edge and check that it lies either in the current range or in the next one. If it does not then $\mathbf{P}$ is not $\mathcal{O}$-oriented. When we find that it lies in the next range the next range becomes the current range and the process is repeated until we have exhausted the edges. Since we either discard a range in $\mathcal{O}$ or an edge of $\mathbf{P}$ in constant time and we only examine each range in $\mathcal{O}$ at most twice then this algorithm takes $O(n+r)$ time.

To see that this is optimal consider a convex $\mathcal{O}$-polygon. Such a polygon can be constructed for any $\mathcal{O}$ (for example, see the solid line polygon in figure 7.3, here $\mathcal{O} = \{36°, 72°, 108°, 140°\}$). Such a polygon will be non-$\mathcal{O}$-oriented if one of its edges is not in $\mathcal{O}$. Since each range in $\mathcal{O}$ is represented in the polygon then any algorithm is forced to examine each range in $\mathcal{O}$ at least once. Also any algorithm must examine each edge of the polygon at least once thus it requires $\Omega(n+r)$ time to test any polygon for $\mathcal{O}$-convexity.

We may test a polygon for $\mathcal{O}$-convexity either using a slightly more complicated algorithm than the one outlined above or we may use the monotone testing algorithm of Preparata and Supowit ([48]). First find all directions of monotonicity using the algorithm of Preparata and Supowit. This algorithm takes linear time and it outputs an ordered doubly-linked list of orientations of size $O(n)$. Transform the list in linear time to the set of corresponding orthogonal orientations. That is, if θ is in the monotonicity list, replace it by $\theta + 90°$. Now, find the first $\mathcal{O}$-orientation in the list (linear time) then check each range in $\mathcal{O}$ in order. This takes $O(n+r)$ time overall.

To see that this is optimal consider an $\mathcal{O}$-convex $\mathcal{O}$-polygon which consists of alternating convex and reflex angles. Such a polygon can be constructed for any $\mathcal{O}$ (for example, see the dashed line polygon in figure 7.3, here $\mathcal{O} = \{36°, 72°, 108°, 140°\}$).

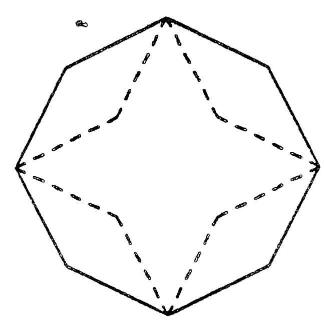

Figure 7.3: Polygons Realising the Lower Bound

Such a polygon will be non-$\mathcal{O}$-convex if any of the pairs of $\mathcal{O}$-line segments making up a reflex angle are non-consecutive in $\mathcal{O}$, and so any algorithm is forced to examine each range in $\mathcal{O}$ at least once. Also any algorithm must examine each edge of the polygon at least once thus it must take $\Omega(n+r)$ time to test any polygon for $\mathcal{O}$-convexity. ■

7.3 Relationships Between the Hulls for $\mathcal{O}$-Polygons

If **P** is a $\mathcal{O}$-polygon then all of the hulls exist and two of the hulls are always equal as we now demonstrate.

Lemma 7.3.1 *If* $\mathbf{P}$ *is an* $\mathcal{O}$*-polygon then* $\mathcal{O}\mathcal{O}$*-hull*$(\mathbf{P}) = \mathcal{O}$*-hull*$(\mathbf{P})$

Proof: If $\mathcal{O}$-*hull*$(\mathbf{P})$ is not $\mathcal{O}$-oriented then it has some non-$\mathcal{O}$-oriented edge $\mathcal{L}$. $\mathcal{L}$ cannot be a subset of an edge of $\mathbf{P}$ since $\mathbf{P}$ is an $\mathcal{O}$-polygon. Hence there is at least one point in $\mathcal{O}$-*hull*$(\mathbf{P})$ between $\mathcal{L}$ and the polygon, and this point can be removed by replacing a part of $\mathcal{L}$ by a length 2 $\mathcal{O}$-staircase made up of line segments of orientations θ_1 and θ_2 where (θ_1, θ_2) is $\mathcal{L}$'s maximal $\mathcal{O}$-free range.

Such a replacement preserves the $\mathcal{O}$-convexity of the hull because no $\mathcal{O}$-line can cut both of the new segments simultaneously and no $\mathcal{O}$-line can cut $\mathcal{L}$ and either one of the segments since (θ_1, θ_2) is $\mathcal{O}$-free. Therefore we have obtained an $\mathcal{O}$-convex polygon which contains $\mathbf{P}$ and is properly contained in $\mathcal{O}$-*hull*$(\mathbf{P})$, a contradiction. ∎

The above proof fails if $\mathbf{P}$ is not $\mathcal{O}$-oriented.

Lemma 7.3.2 *If* $\mathbf{P}$ *is an* $\mathcal{O}$*-polygon, then*

$$\mathbf{P} \subseteq \mathcal{O}\text{-}hull(\mathbf{P}) = \mathcal{O}\mathcal{O}\text{-}hull(\mathbf{P}) \subseteq hull(\mathbf{P}) \subseteq \ strong\ \mathcal{O}\text{-}hull(\mathbf{P})$$

Proof: These containment results follow from the observations of the previous chapter that strongly $\mathcal{O}$-convex sets are convex, and convex sets are $\mathcal{O}$-convex. ∎

These relationships hold even when $\mathbf{P}$ is not $\mathcal{O}$-oriented except that, as is easy to show, $\mathcal{O}\mathcal{O}$-*hull*$(\mathbf{P})$ may not exist (see figure 7.2).

7.4 Computation of the Hulls

Using the algorithms already developed for the orthogonal convex hull as our starting point gives $O(nr)$ time algorithms to find both the $\mathcal{O}$-*hull* and the strong $\mathcal{O}$-*hull*, respectively, of an n-vertex polygon. This is because the algorithms presented in [41,39,43,54] essentially decomposed the problem into finding the hull in first one then the other orthogonal direction.

This approach also leads to a straightforward $O(nr)$ time algorithm when we have r ranges. That this can be improved upon is due to the high degree of structure possessed by both the *O-hull* and the strong *O-hull*.

From theorem 7.2.1 we know that the *O-hull* of a polygon $\mathbf{P}$ must be the smallest polygon which both contains $\mathbf{P}$ and is composed of a sequence of polygonal $\mathcal{O}$-stairsegments meeting at convex angles. If we were to find the *convex hull* of $\mathbf{P}$, how would we have to modify it to obtain *O-hull*($\mathbf{P}$)? (From [37,3,18,33] we know that *hull*($\mathbf{P}$) can be found in $\Theta(n)$ time.) It is easy to see that the boundary of *O-hull*($\mathbf{P}$) must lie between the boundaries of $\mathbf{P}$ and *hull*($\mathbf{P}$). In fact, we can prove—

Lemma 7.4.1 *If* $\mathbf{P}$ *is a polygon then the extremal vertices of* $\mathbf{P}$*, that is, the vertices of hull*($\mathbf{P}$)*, are vertices of O-hull*($\mathbf{P}$) *and the edges of O-hull*($\mathbf{P}$) *between any two consecutive* $\mathcal{O}$*-extremal vertices of* $\mathbf{P}$ *form a polygonal* $\mathcal{O}$*-stairsegment.*

Proof: The proof that the extremal vertices are vertices of the *O-hull* is similar to the proof of case 2 of theorem 5.5.1 and is omitted. The second assertion is a slight refinement of theorem 7.2.1. ∎

There is an analogous result for the strong *O-hull* of $\mathbf{P}$.

We adapt the "gift wrapping" method used by Jarvis ([25]) to find *O-hull*($\mathbf{P}$) in $O(n+r)$ time. Consider how to construct the portion of the hull lying between two consecutive extremal vertices of $\mathbf{P}$. Let the two vertices be v_i and v_j where $i < j$.

If $\Theta(\mathcal{LS}[v_i, v_j]) \in \mathcal{O}$ then all we need do to process these two vertices is to add them to the list of vertices that we have already found to be in the *O-hull*. Also, if $j = i+1$ then we will always add the two vertices to the list since v_i and v_j are extremal. Otherwise assume without loss of generality that $(0°, 90°)$ is $\mathcal{LS}[v_i, v_j]$'s maximal $\mathcal{O}$-free range. Consider the boundary of $\mathbf{P}$ from v_i to v_j (see figure 7.4).

If any vertex is obstructed in the x-direction from v_i or in the y-direction from v_j then it cannot be on the list of vertices belonging to *O-hull*($\mathbf{P}$). It is now easy to construct an

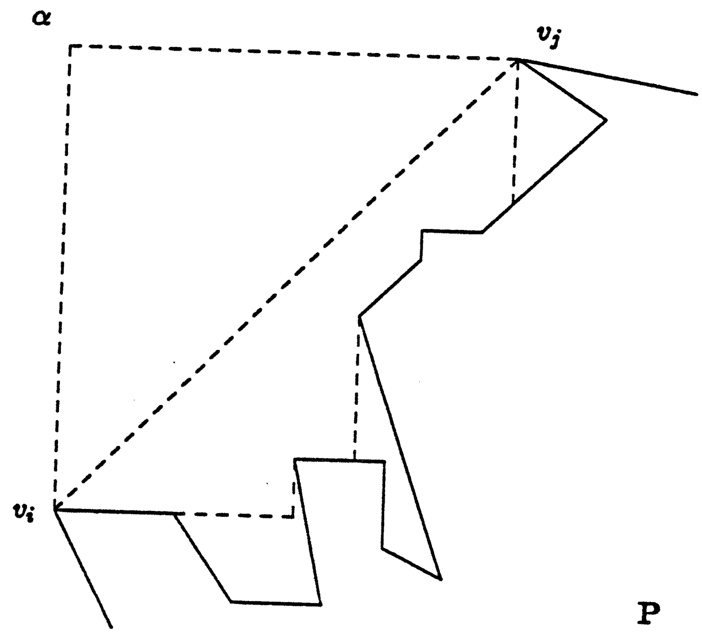

Figure 7.4: Forming the $\mathcal{O}$-*hull*

algorithm to step along the vertices in between each pair of consecutive extremal vertices of **P** and find the vertices of **P** which belong in $\mathcal{O}$-*hull*(**P**). As in the proof of theorem 7.2.2 we need only to find the range that contains the orientation of any one of the line segments defined by the consecutive extremal vertices of **P** and from that point on march along the extremal vertices of **P** and the ranges in $\mathcal{O}$ in lock-step. Since we are examining the vertices of a convex polygon (that is, the convex hull of **P**) *in order* we will only encounter any one pair of consecutive $\mathcal{O}$-range boundaries at most twice. Also, we will never see any two consecutive extremal vertices more than once. So at each step we either discard one of the vertices of the polygon or we discard one of the possible pairs of consecutive $\mathcal{O}$-range boundaries in constant time per step.

The strong $\mathcal{O}$-*hull* of **P** is even simpler to compute, since the only vertex we need add is one of the corners of $\mathcal{O}$-$\|[v_i, v_j]$ when v_i and v_j are extremal. For example, in figure 7.4 we need only add the new vertex α and delete v_i if this would make it a collinear vertex.

Although both hulls may require the introduction of new vertices, the number of vertices is not increased.

Hence we have proved the following theorem:

THEOREM 7.4.1 *If* P *is a polygon then both* $\mathcal{O}$*-hull*(P) *and the strong* $\mathcal{O}$*-hull of* P *can be found optimally in* $\Theta(n+r)$ *time in the worst case where* n *is the number of vertices of* P *and* r *is the number of ranges in* $\mathcal{O}$.

7.5 The $\mathcal{O}$-*hulls* for Finite Point Sets

Case 1: $\mathcal{O} = \emptyset$. The strong $\emptyset$-*hull* of any set containing at least two points is the entire plane. The $\emptyset$-*hull* of any set is the set itself. Thus both $\mathcal{O}$-*hulls* may be found in $O(1)$ time.

Case 2: $\mathcal{O} = 0°$. The strong $0°$-*hull* is the infinite section of the plane bounded by two horizontal lines through the points with minimum and maximum y coordinate. Thus the strong $0°$-*hull* of a set of n points can be computed in $O(n)$ time. Since finding the minimum and maximum of n values requires $\Omega(n)$ time this is an optimal procedure.

To find the $0°$-*hull* sort the points by y coordinate and if any two points have the same value then the horizontal line segment joining them is added to the hull. For a set of n points this takes $O(n \lg n)$ time. From the Element Uniqueness problem, this is an optimal procedure since whether a set of n values is unique can be answered in the following fashion: Construct a set of points whose y coordinates are the n values and whose x coordinates are the numbers $1, 2, 3, \ldots, n$ in succession. This can be done in linear time. Find the $0°$-*hull* of these points and test whether the points form their own $0°$-*hull*. This can also be done in linear time. At least one line segment is added to the set of points if and only if at least two y coordinates are the same.

Case 3: $|\mathcal{O}| \geq 2$. The strong $\mathcal{O}$-*hull* is the bounding box in at most r orientations. It can be computed in $O(n \lg n + r)$ time by finding the convex hull then finding its strong $\mathcal{O}$-*hull*. Observe that if $\mathcal{O}$ is finite then there is an alternate algorithm with time bound

$O(nr)$ ($|\mathcal{O}| = r$). This algorithm finds the maximum and minimum of the points in each orientation in order (following the order of the orientations in $\mathcal{O}$). We do not have a non-trivial lower bound for this problem.

The $\mathcal{O}$-*hull* may be computed in $O(rn \lg n)$ time by adapting the orthogonal hull algorithm of Ottmann *et al.* ([43]). $\Omega(n \lg n)$ time is a lower bound since it is a lower bound when $|\mathcal{O}| = 1$.

We shall follow Ottmann *et al.* by first deciding which points belong in a connected component of the $\mathcal{O}$-*hull* and then find the $\mathcal{O}$-*hull* of each separate component.

Suppose that we are given a set of at most n points which belong to a connected component of the $\mathcal{O}$-*hull*. For each pair of adjacent range boundaries in $\mathcal{O}$ we run the orthogonal hull algorithm to produce in $O(rn \lg n)$ time $2r$ partial hulls. The union of the set of partial hulls is the $\mathcal{O}$-*hull* of the set of points since, from the decomposition theorem, if a set is connected then the union of the partial hulls is the $\mathcal{O}$-hull.

Thus we have reduced the problem to identifying which points belong to separate connected components in the $\mathcal{O}$-*hull*. To decide this we proceed as follows: For each pair of consecutive range boundaries in $\mathcal{O}$ run the connected components version of the orthogonal hull algorithm on the set of points. This takes $O(rn \lg n)$ time and produces at most n connected components. From lemma 5.2.3 we know that the $\mathcal{O}$-*hull* must be a collection of disjoint connected components such that no $\mathcal{O}$-line intersects any pair of components. That is, each component must be *separated* from every other component by a θ-line for every $\theta \in \mathcal{O}$. We adapt the orthogonal hull algorithm to scan the set of partial hulls produced and amalgamate any pair of components which share a θ-range for some θ in $\mathcal{O}$. Note that it is sufficient to examine only the *range boundaries* in $\mathcal{O}$. This procedure takes $O(rn \lg n)$ time and produces at most n disjoint connected components.

The set of disjoint $\mathcal{O}$-*hulls* produced is then the $\mathcal{O}$-*hull* of the set of n points since each connected component is $\mathcal{O}$-convex and they are all $\mathcal{O}$-separable from each other.

It may be that this algorithm can be improved to $O(n \lg r)$ time using techniques similar to those of Kirkpatrick and Seidel (see [31]). This is quite possible if we know beforehand that the points form a *connected* $\mathcal{O}$*-hull*, however, it is not immediately obvious how to find the connected components of the $\mathcal{O}$*-hull* in $O(n \lg r)$ time.

Chapter 8

Conclusions and Open Problems

8.1 Open Problems

8.1.1 Decomposability

Probably the single most important observation arising from this work is that for convexity and visibility problems it is no more difficult to work with $\mathcal{O}$-objects than the more restrictive orthogonal objects. Indeed, in a sense, problems involving $\mathcal{O}$-convexity are basically *orthogonal* problems, that is, they can be decomposed so that at any point in time we need only consider *two* orientations (the two nearest). The first open problem is to characterize *a priori* those problems which are decomposable.

As a first problem to examine we suggest the following: Widmayer *et al.* [62] give an $O(|\mathcal{O}|n)$ (finite $\mathcal{O}$) algorithm to find the minimum distance between two $\mathcal{O}$-convex polygons. Can this be improved to $O(n+r)$ as for our convex hull algorithms?

A related problem is to find the *diameter* (least upper bound of distances between points of the set) of an $\mathcal{O}$-convex set.

Finally, how fast can we determine if a polygon is $\mathcal{O}$-oriented? It is easy to construct an algorithm with time bound $O(n \lg r)$ (perform a binary search for each edge in $\mathcal{O}$). Can this be improved?

8.1.2 Mathematical Underpinnings

Restricted orientation geometry gives rise to many interesting mathematical questions. The most obvious one is "What are the purely metric and purely topological properties of the metric space induced by $\mathcal{O}$?" (By the induced metric space we mean, following Widmayer *et al.* [62], the metric space with distance function between two points defined as the Euclidean length of the shortest $\mathcal{O}$-staircase connecting the two points.)

On the algebraic side there is a powerful characterization of convex sets which we have not been able to generalize:

P is said to be *convex* if $\forall p, q \in \mathbf{P}$ and $\forall 0 \leq \alpha \leq 1$, $\alpha p + (1 - \alpha) q \in \mathbf{P}$.

This algebraic characterization allows the (simple and easily generalized) algebraic proof of many of the other characterizations of convex sets. Further it makes it easy to discover new applications of convex sets (usually to the theory of linear programming).

The open question is to characterize $\mathcal{O}$-convex sets in a similar algebraic manner.

8.1.3 Algorithms on $\mathcal{O}$-Objects

$\mathcal{O}$-Boolean Masking: (see [45] for definitions) Given r sets of possibly overlapping $\mathcal{O}$-polygons find their Boolean combination. Widmayer and Wood ([61]) show that the $O((n + k)(r + \lg n))$ time bound for the general case can be reduced to $O(n \lg n + k)$ by restricting attention to orthogonal polygons. Can the time bound for this problem be improved if the polygons are $\mathcal{O}$-oriented?

$\mathcal{O}$-Line Segment Intersection: Can this be done any faster than testing for intersection in each of $|\mathcal{O}|$ directions? (Assuming finite $\mathcal{O}$.)

$\mathcal{O}$-Nearest Neighbour: Widmayer *et al.* ([62]) present several algorithms to answer closest point problems by building a Voronoi Diagram using the induced $\mathcal{O}$-metric (see previous section). These algorithms were all for finite $\mathcal{O}$. Can they be generalized to arbitrary $\mathcal{O}$ and still yield efficient solutions?

$\mathcal{O}$-Convex Skull: How fast can the largest $\mathcal{O}$-convex subpolygon of a polygon be found? Wood and Yap [65] show that the time bound of $O(n^7)$ for the general case, that is the convex skull of an arbitrary polygon ([5]), can be improved to $O(n^2)$ for the orthogonally convex skull of an orthogonal polygon. Are time bounds improved if the polygon is $\mathcal{O}$-oriented and we ask for the $\mathcal{O}$-convex skull? As a special case consider finding the largest empty $\mathcal{O}$-parallelogram in a polygon.

8.1.4 Abstract Convexity

The Refinement and Decomposition Theorems (theorems 3.4.2 and 3.4.3) for abstract convexity spaces were discovered as a result of research on $\mathcal{O}$-convex sets. Because geometric relationships are usually easy to see, $\mathcal{O}$-convex sets can serve as an ideal testbed for various conjectures on convex sets in an abstract convexity setting.

Also, abstract convexity theory has been severely criticized by many geometers for being too sterile.[1] One possible reason for this is that abstract convexity has no notion of *visibility*. The work in chapter 6 on abstract visibility provides a basis for a theory of visibility but much remains to be done. As a first step we ask for the simplest (in some sense) new axiom to add to the two convexity space axioms such that we are able to prove that for all of our definitions of visibility the kernel under each of these definitions is convex.

[1]Peter McMullen, private communication.

8.1.5 Properties of Convex Sets

Unlike convex sets, $\mathcal{O}$-convex sets are not closed under affine transformations. But as we have seen other properties of convex sets carry over directly—for example:

THEOREM **8.1.1** *The intersection of two $\mathcal{O}$-convex sets is $\mathcal{O}$-convex.*

Clearly not all theorems on convex sets carry over directly to $\mathcal{O}$-convex sets, the open question being to determine which ones do. For example, for which notions of $\mathcal{O}$-convexity is the following result true?

THEOREM **8.1.2** *Given an $\mathcal{O}$-convex set* $\mathbf{P}$ *and a point* $p \notin \mathbf{P}$ *the set consisting of all $\mathcal{O}$-stairsegments from* p *to each point in* $\mathbf{P}$ *is again $\mathcal{O}$-convex.*

8.1.6 Other Definitions of $\mathcal{O}$-Convexity

There are many other interesting ways to define $\mathcal{O}$-convex sets, in particular the following variants should be investigated, as they may be of some relevance to robotic path problems.

Definition: $\mathbf{P}$ is said to be *weakly (strongly) arm $\mathcal{O}$-convex* if $\forall p, q \in \mathbf{P}$ p sees q in $\mathbf{P}$ via at least one (both) of the arms of $\mathcal{O}$-$\|[p, q]$.

It is easy to show that strong arm $\mathcal{O}$-convexity is equivalent to strong $\mathcal{O}$-convexity, what though can be said about weak arm $\mathcal{O}$-convexity?

Definition: $\mathbf{P}$ is said to be *n-$\mathcal{O}$-staircase $\mathcal{O}$-convex* if $\forall p, q \in \mathbf{P}$ p sees q in $\mathbf{P}$ via an $\mathcal{O}$-staircase of length n.

This last definition is too restrictive when $n = 1$, since it requires that $\mathbf{P}$ be an $\mathcal{O}$-line, $\mathcal{O}$-line segment or $\mathcal{O}$-stairray. But the following result should hold:

CONJECTURE **8.1.1** $\mathbf{P}$ *is weakly arm $\mathcal{O}$-convex if and only if* $\mathbf{P}$ *is 2-$\mathcal{O}$-staircase $\mathcal{O}$-convex.*

n-$\mathcal{O}$-staircase convexity is a special case of *L_m-convexity* introduced by ElGindy *et al* in [15]. An L_m-convex set is a set for which every two points in the set can be connected by a polygonal curve of length m. They present a linear time algorithm to triangulate any L_2-convex polygon and a quadratic time algorithm to check whether a polygon is L_2-convex. Since n-$\mathcal{O}$-staircase convexity is a special case is there a linear time algorithm to determine whether a polygon is n-$\mathcal{O}$-staircase convex?

We can extend n-staircase convexity to allow arbitrarily many turns in each staircase:

Definition: $\mathbf{P}$ is said to be *$\mathcal{O}$-staircase $\mathcal{O}$-convex* if $\forall p, q \in \mathbf{P}$ p sees q in $\mathbf{P}$ via an $\mathcal{O}$-staircase.

The reader should observe that there are serious compatibility problems with $\mathcal{O}$-staircase $\mathcal{O}$-convexity since an arbitrary *line* is *not* $\mathcal{O}$-staircase $\mathcal{O}$-convex *unless* the line happens to be an $\mathcal{O}$-line.

Finally, we may define convexity "componentwise."

Definition: $\mathbf{P}$ is said to be *component weakly $\mathcal{O}$-convex* (*component strongly $\mathcal{O}$-convex*) if $\forall p, q \in \mathbf{P}$, whenever p and q are connected in $\mathbf{P}$, p sees q in $\mathbf{P}$ via an $\mathcal{O}$-stairsegment (via $\mathcal{O}$-$\|[p,q]$).

Component strong $\mathcal{O}$-convexity is the same as Soisalon-Soininen and Wood's "R-closure" (see chapter 3), when $\mathcal{O} = \{0°, 90°\}$. Note that if $\mathbf{P}$ is weakly $\mathcal{O}$-convex then $\mathbf{P}$ is $\mathcal{O}$-convex and if $\mathbf{P}$ is $\mathcal{O}$-convex then $\mathbf{P}$ is component weakly $\mathcal{O}$-convex and both containments are proper. It is possible to prove that this last definition yields convexity spaces, that is:

1. *The component strong $\mathcal{O}$-convexity space* ($\Re^2$, set of all component strongly $\mathcal{O}$-convex sets). This convexity space is a refinement of the strong $\mathcal{O}$-convexity space.
2. *The component weak $\mathcal{O}$-convexity space* ($\Re^2$, set of all component weakly $\mathcal{O}$-convex sets). This convexity space is a refinement of both the normal convexity space and the component strong $\mathcal{O}$-convexity space.

8.1.7 Orientations as Lines

Rather than considering $\mathcal{O}$ to be a set of orientations we may consider it to be a pencil of lines. This suggests that we generalize $\mathcal{O}$ to be a pencil of arbitrary curves. A set is said to be $\mathcal{O}$-convex if its intersection with any translate of an $\mathcal{O}$-curve is connected. For example, call a polygonal curve a *broken line* if it is made up of two rays. We may define $\mathcal{O}$ to be a collection of broken lines. In particular, if the broken lines are made up of two orthogonal rays and we add the extra condition that the intersection must be connected *and contain the corner point of the broken line* then we have the NE- and SW-closures of Soisalon-Soininen and Wood ([57]).

Drandell ([12]) has investigated some of the topological properties of such a pencil of curves, however he requires that for every pair of points there is a *unique* curve in the pencil passing through them. Thus this generalization fails to include weakly $\mathcal{O}$-convex sets. Note though that any collection of $\mathcal{O}$-stairlines *with disjoint spans* satisfies this requirement.

8.1.8 The Art Gallery Theorem

In chapter 6 it was shown that if $\mathcal{O}$ contains two or more orientations then $\lfloor n/3 \rfloor$ $\mathcal{O}$-guards are necessary and sufficient to guard any simple n vertex polygon. This result was shown by appealing to Chvátal's Art Gallery Theorem. Since Chvátal's theorem is really only a special case of the result we ask whether the more general result can be proved more directly—perhaps by showing that a particular decomposition with respect to some set of orientations in $\mathcal{O}$ always exists.

8.1.9 Decomposing Polygons

Many algorithms in computational geometry depend on decomposing an arbitrary polygon into a collection of covering or partitioning subpolygons of some special type. For example, polygons are often decomposed into monotone, starshaped, or convex subpolygons to answer point location problems more efficiently ([58]). In this thesis we have shown that the following containments amongst families of different kinds of "convex" sets hold:

$$\text{strongly } \mathcal{O}\text{-convex} \subseteq \text{convex} \subseteq \mathcal{O}\text{-convex} \subseteq \text{monotone}$$

$$\text{strongly } \mathcal{O}\text{-starshaped} \subseteq \text{starshaped} \subseteq \text{weakly } \mathcal{O}\text{-starshaped}$$

We ask whether efficient decomposition algorithms can be designed for monotone polygons into $\mathcal{O}$-convex polygons, $\mathcal{O}$-convex polygons into convex polygons or for convex polygons into strongly $\mathcal{O}$-convex polygons and similarly for starshaped polygons.

Observe that starshaped polygons are not necessarily $\mathcal{O}$-convex and that $\mathcal{O}$-convex polygons are not necessarily starshaped. Thus, although convex sets are a subset of starshaped sets, starshaped sets are incomparable with $\mathcal{O}$-convex sets.

8.1.10 Allowing Open Ranges and Directed Lines

Can the notion of $\mathcal{O}$-orientation be generalized so that $\mathcal{O}$ may be *any* subset of the set of all orientations? That is, can a reasonable theory be built up if we allow open or half-open ranges in $\mathcal{O}$? The problem with allowing open ranges in $\mathcal{O}$ is that there is no longer any notion of a *nearest* orientation in $\mathcal{O}$ to a given orientation not in $\mathcal{O}$. Many of our results depend on identifying the two nearest orientations to a given orientation.

We ask the same question for $\mathcal{O}$ being a collection of *directions* as opposed to orientations. That is, we may distinguish between the two possible directions of a line.

8.1.11 Generalizing to Higher Dimensions

In attempting to generalize the notions of $\mathcal{O}$-convexity and $\mathcal{O}$-objects to higher dimensions we would like to retain most, if not all, of the advantages of the two dimensional case. Thus, we would like definitions which are easy to use and which are intuitive; which allow space-sweep algorithms which are easy to design; and which are—at least—sensible in three dimensions. In the following we explore one such generalization.

There is a straightforward extension of $\mathcal{O}$-orientation geometry to $\Re^n$ which singles out a subclass, $\mathcal{O}$, of $(n-1)$-dimensional hyperplanes and in which a set is defined to be $\mathcal{O}$-oriented if its boundary hyperplanes are translates of those in $\mathcal{O}$. However, the "natural" definition of $\mathcal{O}$-convexity under this definition soon leads to confusion.

The problem is that the natural definition of convexity is line convexity. That is, $\mathcal{O}$ should be a collection of lines and a set is said to be $\mathcal{O}$-convex if its intersection with any translate of a line in $\mathcal{O}$ is connected. Thus, for example, the natural generalization of orthogonal convexity to $\Re^3$ entails that an orthogonally convex set have a connected intersection with any line parallel to one of the three axes. As we show there is some difficulty in having both "$\mathcal{O}$-oriented objects" and "$\mathcal{O}$-convex sets" for dimensions higher than 2.

Suppose that we define a set to be $\mathcal{O}$-convex if its intersection with any translate of an $\mathcal{O}$-hyperplane is connected. When $n = 2$ this definition is precisely the definition given before since 1-dimensional hyperplanes are lines. Any k-flat in $\Re^n$ is $\mathcal{O}$-convex in this sense. In particular, for $n = 2$, the flats are: any point (a 0-dimensional hyperplane), any line (a 1-dimensional hyperplane) and the whole plane (a 2-dimensional hyperplane). Moreover, these are all "$\mathcal{O}$-convex."

However there is a problem with this definition once $n \geq 3$. Consider an empty box in $\Re^3$. Suppose that $\mathcal{O}$ is a set of three mutually orthogonal planes and that the box is $\mathcal{O}$-oriented. This box, *although it is empty,* is "$\mathcal{O}$-convex" by our definition because we only require that the intersection with any $\mathcal{O}$-plane be connected *and not* simply connected.

If we try to remedy this by strengthening the definition to simple connectedness instead of just connectedness then the following set will have to be accepted as "$\mathcal{O}$-convex." Consider a (solid) box with a notch in it. The notch is in the shape of a smaller box cut out of the large box such that two faces of the inner box are exposed. This object is "$\mathcal{O}$-convex" even under the stronger "simply connected" definition of $\mathcal{O}$-convexity.

Thus we might define $\mathcal{O}$-convexity in the following way:

Definition: If $\mathcal{H}$ is a hyperplane in $\mathcal{O}$ then the set $\mathbf{P}$ is $\mathcal{O}_{\mathcal{H}}$*-convex* if the intersection of $\mathbf{P}$ and any translate of $\mathcal{H}$ is convex with respect to the family of lines generated on $\mathcal{H}$ by every other hyperplane in $\mathcal{O}$.

Definition: A set $\mathbf{P}$ is $\mathcal{O}$*-convex* if for each hyperplane $\mathcal{H} \in \mathcal{O}$ the intersection of $\mathbf{P}$ and any translate of $\mathcal{H}$ is $\mathcal{O}_{\mathcal{H}}$-convex.

This definition corresponds to our intuition when $n = 3$, and the open problem is to determine the "best" such generalization.

8.2 Conclusion

It may be argued that since computational geometry concerns itself with figures in $\Re^n$ that it is not necessary to develop the theory of $\mathcal{O}$-oriented sets in as general a setting as is possible. There are two obvious rejoinders to this point of view. First, the history of algorithms for finding the convex hull of a polygon illustrates that unaided *geometric intuition* is not sufficiently powerful to avoid incorrect algorithms as there have been several algorithms proposed over time (and accepted as correct) which were later shown to be incorrect (see, for example, [3,59,37]). Any theoretical machinery that may aid insight seems desirable. Second, there is a well-demonstrated synergism between theoretical investigations and practical problems, in that practice suggests new areas for theory and in turn a developing theory suggests a broadening and sharpening of practique.

Further, I believe that this material may have at least one important application area—specifically, restricted orientation VLSI design—and that it will be of continuing theoretical interest. Also, this thesis has shown that the investigation of convexity and visibility topics in restricted orientation geometry benefits from and can enrich subjects as seemingly far afield as abstract convexity theory. Finally, if any further justification were needed, I submit that the study of restricted orientation geometry is of interest in its own right.

While the practical concerns from which computational geometry grew will continue to change and expand, the broad outlines of computational geometry that serve to delineate it from classical geometry and combinatorial geometry are, in my opinion, sufficiently well defined that it can now, in its turn, give impetus to the development of new directions of geometry. I submit that restricted orientation geometry is one such new direction.

In sum, it is my opinion that the field of computational geometry will benefit from a full investigation of the notion of restricted orientation geometry in that the investigation should continue to unearth hitherto unsuspected and theoretically pleasing generalizations.

Notation

$\mathcal{O}$	A closed set of orientations.
$\Re^2$	The plane.
p, q	Points in the plane.
$\mathbf{P}, \mathbf{Q}$	Sets of points in the plane.
$\mathcal{H}$	A halfplane.
$\mathcal{L}$	A line.
$\langle a, b\rangle$	An indeterminate range. The range may be open, half-open or closed.
$\mathcal{H}\langle \mathcal{L}, p]$	The (open or closed) halfplane bounded by the line $\mathcal{L}$ and containing p.
$\mathcal{L}[\theta, p]$	The line of orientation θ through p.
$\mathcal{LS}\langle p, q\rangle$	The (open, half-open or closed) line segment through p and q.
$\Theta(\mathcal{L})$	The orientation of the line (or line segment) $\mathcal{L}$.

Glossary

$\mathcal{O}$-range A contiguous set of orientations in $\mathcal{O}$. All sets of orientations are described by listing their (disjoint) ranges of orientations in order.

$\mathcal{O}$-free range A contiguous set of orientations not in $\mathcal{O}$.

maximal $\mathcal{O}$-free range An $\mathcal{O}$-free range which is not a proper subset of of any other $\mathcal{O}$-free range.

$\mathcal{O}$-line A line whose orientation is in $\mathcal{O}$.

$\mathcal{O}$-extremal A point of support of some set with respect to an $\mathcal{O}$-line.

$\mathcal{O}$-convex set A set which has a connected intersection with every $\mathcal{O}$-line.

$\mathcal{O}$-stairline (Also, $\mathcal{O}$-stairsegment, $\mathcal{O}$-ray.) An $\mathcal{O}$-convex curve.

$\mathcal{O}$-staircase An $\mathcal{O}$-convex $\mathcal{O}$-oriented polygonal curve.

$\mathcal{O}$-parallelogram (Denoted $\mathcal{O}$-$\|[p, q]$.) The union of all $\mathcal{O}$-stairsegments connecting p and q.

maximal $\mathcal{O}$-stairsegment A subset of a curve which is an $\mathcal{O}$-stairsegment and which is not a proper subset of any other $\mathcal{O}$-stairsegment in the curve.

span The span of a set is the smallest range of orientations which contains every orientation defined by each pair of points in the set.

$\mathcal{O}$-*hull*(**P**) The smallest $\mathcal{O}$-convex set containing **P**.

Weakly $\mathcal{O}$-convex set A set such that there is an $\mathcal{O}$-stairsegment joining each pair of points in the set that is wholly in the set.

Strongly $\mathcal{O}$-convex set A set such that every $\mathcal{O}$-stairsegment joining each pair of points in the set is wholly in the set.

Weak $\mathcal{O}$-guard A point in a given set which can only see points connected to it by an $\mathcal{O}$-stairsegment that joins them and is wholly in the set.

Strong $\mathcal{O}$-guard A point in a given set which can only see points connected to it by every $\mathcal{O}$-stairsegment that joins them and these must all lie wholly in the set.

Weak $\mathcal{O}$-*kernel*(**P**) The set of points p which for each $q \in \mathbf{P}$ there is at least one $\mathcal{O}$-stairsegment connecting p and q in **P**.

Strong $\mathcal{O}$-*kernel*(**P**) The set of points p which for each $q \in \mathbf{P}$ every $\mathcal{O}$-stairsegment connecting p and q is in **P**.

Bibliography

[1] Aho, A. V., Hopcroft, J. E. and Ullman, J. D.; *The Design and Analysis of Computer Algorithms,* Addison-Wesley, Reading Mass., 1974.

[2] Bender, E. A.; "Convex n-ominoes," *Discrete Mathematics,* **8**, 219-226, 1974.

[3] Bhattacharya, B. K. and El Gindy, H.; "A New Linear Convex Hull Algorithm for Simple Polygons," *IEEE Transactions on Information Theory,* **IT-30**, 85-88, 1984.

[4] Binmore, K. G.; *Foundations of Analysis: A Straightforward Introduction. Book 2: Topological Ideas,* Cambridge University Press, 1981.

[5] Chang, J. S. and Yap, C. K.; "A Polynomial Solution for the Potato-Peeling Problem," *Discrete and Computational Geometry,* **1**, 155-182, 1986.

[6] Chassery, J.-M.; "Discrete Convexity: Definition, Parametization, and Compatibility with Continuous Convexity," *Computer Vision, Graphics, and Image Processing* **21**, 326-344, 1983.

[7] Chvátal, V.; "A Combinatorial Theorem in Plane Geometry," *Journal of Combinatorial Theory B,* **18**, 39-41, 1975.

[8] Cohn, P. M.; *Universal Algebra,* Harper & Row, New York, 1965.

[9] Culberson, J. and Rawlins, G. J. E.; "Turtlegons: Generating Simple Polygons from Sequences of Angles," *Proceedings of the* 1^{st} *Annual ACM Symposium on Computational*

Geometry, 305-310, 1985.

[10] Danzer, L., Grünbaum, B. and Klee, V.; "Helly's Theorem and Its Relatives," *Convexity, Proceedings of Symposia in Pure Mathematics,* **VII**, 101-180, 1963. American Mathematical Society, Providence, Rhode Island.

[11] Dobkin, D. and Lipton, R.; "On the Complexity of Computations Under Varying Sets of Primitives," *Journal of Computer and Systems Science,* **18**, 86-91, 1979.

[12] Drandell, M.; "Generalized Convex Sets in the Plane," *Duke Mathematics Journal,* **19**, 537-547, 1952.

[13] Dwinger, Ph.; "On the Lattice of the Closure Operators of a Complete Lattice," *Nieuw Archief voor Wiskunde,* **VI**, 112-117, 1956.

[14] Edelsbrunner, H.; "Intersection Problems in Computational Geometry," Technical University of Graz and Austrian Computer Society, Graz, Austria, Institute für Informationsverarbeitung, Report F 93, 1982.

[15] El Gindy, H., Avis, D. and Toussaint, G.; "Applications of a Two-Dimensional Hidden-Line Algorithm to Other Geometric Problems," *Computing,* **31**, 191-202, 1983.

[16] Fenchel, W.; "Über Krümmung und Windung geschlossener Raumkurven," (in German) *Math. Ann.,* **101**, 238-252, 1929.

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

[18] Graham, R. L. and Yao, F. F.; "Finding the Convex Hull of a Simple Polygon," *Journal of Algorithms,* **4**, 324-331, 1983.

[19] Grätzer, G.; *General Lattice Theory,* Birkhäuser, 1978.

[20] Grünbaum, B.; *Convex Polytopes,* Wiley-Interscience, New York, 1967.

[21] Güting, R. H.; *"Conquering Contours: Efficient Algorithms for Computational Geometry,"* Doctoral Dissertation, Universität Dortmund, 1983.

[22] Güting, R. H.; "Stabbing C-Oriented Polygons," *Information Processing Letters,* **16**, 35-40, 1983.

[23] Güting, R. H.; "Dynamic C-Oriented Polygonal Intersection Searching," *Information and Control,* **3**, 143-163, 1984.

[24] Jamison-Waldner, R. E.; "A Perspective on Abstract Convexity: Classifying Alignments by Varieties," *Convexity and Related Combinatorial Geometry, Proceedings of the* 2^{nd} *University of Oklahoma Conference,* (Kay, D. C. and Breen, M., eds.). Lecture Notes in Pure and Applied Mathematics, **76**, 113-150, 1982. Marcel Dekker, Inc., New York and Basel.

[25] Jarvis, R. A.; "On the Identification of the Convex Hull of a Finite Set of Points in the Plane," *Information Processing Letters,* **2**, 18-21, 1973.

[26] Kahn, J., Klawe, M. and Kleitman, D.; "Traditional Galleries Require Fewer Guards," *SIAM Journal on Algebraic and Discrete Methods,* **7**, 452-454, 1986.

[27] Kay, D. C. and Womble, E. W.; "Axiomatic Convexity Theory and Relationships between the Carathéodory, Helly and Radon Numbers," *Pacific Journal of Mathematics,* **38**, 471-485, 1971.

[28] Keil, J. M.; "Minimally Covering a Horizontally Convex Orthogonal Polygon," *Proceedings of the* 2^{nd} *Annual ACM Symposium on Computational Geometry,* 43-51, 1986.

[29] Kim, C. E. and Rosenfeld, A.; "On the Convexity of Digital Regions," *Proceedings of the* 5^{th} *Joint Conference on Pattern Recognition,* 1010-1015, 1980.

[30] Kirkpatrick, D. G. and Seidel, R.; "Output-Size Sensitive Algorithms for Finding Maximal Vectors," *Proceedings of the* 1^{st} *Annual ACM Symposium on Computational Geometry,* 89-96, 1985.

[31] Kirkpatrick, D. G. and Seidel, R.; "The Ultimate Planar Convex Hull Algorithm," *SIAM Journal of Computing,* **15**, 287-299, 1986.

[32] Klee, V.; "What is a Convex Set?," *American Mathematical Monthly,* **78**, 616-631, 1971.

[33] Lee, D. T.; "On Finding the Convex Hull of a Simple Polygon," *International Journal of Computer and Information Sciences,* **12**, 87-99, 1983.

[34] Lee, D. T. and Preparata, F. P.; "An Optimal Algorithm for Finding the Kernel of a Polygon," *Journal of the ACM,* **26**, 415-421, 1979.

[35] Levi, F. W. "On Helly's Theorem and The Axioms of Convexity," *Journal of the Indian Mathematical Society,* **15**, 65-76, 1951.

[36] Lipski, W. and Papadimitriou, C. H.; "A Fast Algorithm for Testing for Safety and Detecting Deadlocks in Locked Transaction Systems," Technical Report MIT/LCS/TM-181, Laboratory for Computer Science, MIT, 1980.

[37] McCallum, D. and Avis, D.; "A Linear Algorithm for Finding the Convex Hull of a Simple Polygon," *Information Processing Letters,* **9**, 201-206, 1979.

[38] Mehlhorn, K.; *Data Structures and Algorithms 3: Multi-dimensional Searching and Computational Geometry,* Springer-Verlag, 1984.

[39] Montuno, D. Y. and Fournier, A.; "Finding the X-Y Convex Hull of a Set of X-Y Polygons," Technical Report CSRG-148, University of Toronto, 1982.

[40] Munro, J. I., Overmars, M. and Wood, D.; "Variations on Visibility," *Proceedings of the* 3^{rd} *Annual ACM Symposium on Computational Geometry,* 291-299, 1987.

[41] Nicholl, T. M., Lee, D. T., Liao, Y. Z. and Wong, C. K.; "Constructing the X-Y Convex Hull of a Set of X-Y Polygons," *BIT,* **23**, 456-471, 1983.

[42] O'Rourke, J.; "Galleries Need Fewer Mobile Guards: A Variation of Chvatal's Theorem," *Geometriae Dedicata,* **14**, 273-283, 1983.

[43] Ottmann, Th., Soisalon-Soininen, E. and Wood, D.; "On The Definition and Computation of Rectilinear Convex Hulls," *Information Sciences,* **33**, 157-171, 1984.

[44] Ottmann, Th., Soisalon-Soininen, E. and Wood, D.; "Partitioning and Separating Sets of Orthogonal Polygons," *Information Sciences,* **40**, to appear, 1987.

[45] Ottmann, Th., Widmayer, P. and Wood, D.; "A Fast Algorithm for the Boolean Masking Problem," *Computer Vision, Graphics, and Image Processing,* **30**, 249-268, 1985.

[46] Overmars, M. and Wood, D.; "On Rectangular Visibility," *Journal of Algorithms,* to appear, 1987.

[47] Preparata, F. P. and Shamos, M. I.; *Computational Geometry,* Springer-Verlag, New York, 1986.

[48] Preparata, F. P. and Supowit, K. J.; "Testing a Simple Polygon for Monotonicity," *Information Processing Letters,* **12**, 161-164, 1981.

[49] Rawlins, G. J. E. and Wood, D.; "Optimal Computation of Finitely Oriented Convex Hulls," *Information and Computation,* **72**, 150-166, 1987.

[50] Rawlins, G. J. E. and Wood, D.; "Restricted-Orientation Convexity," Technical Report No. 86-36, University of Waterloo, 1986.

[51] Rawlins, G. J. E. and Wood, D.; "Orthoconvexity," *Computational Morphology,* to appear (Toussaint, G., ed.). North Holland, Amsterdam, 1987.

[52] Reckhow, R. A. and Culberson, J.; "Covering a Simple Orthogonal Polygon with a Minimum Number of Orthogonally Convex Polygons," *Proceedings of the* 3^{rd} *Annual ACM Symposium on Computational Geometry,* 268-277, 1987.

[53] Rosenfeld, A; "Digital Straight Line Segments," *IEEE Transactions on Computers,* c-23, 1264-1269, 1974.

[54] Sack, J.-R.; *"Rectilinear Computational Geometry,"* Doctoral Dissertation, Carleton University, School of Computer Science, 1984.

[55] Shamos, M. I.; *"Problems in Computational Geometry,"* Doctoral Dissertation, Yale University, Computer Science Department, 1978.

[56] Sierksma, G.; "Extending a Convexity Space to an Aligned Space," *Indagationes Mathematicae,* 46, 429-435, 1984.

[57] Soisalon-Soininen, E. and Wood, D.; "Optimal Algorithms to Compute the Closure of a Set of Iso-Rectangles," *Journal of Algorithms,* 5, 199-214, 1984.

[58] Toussaint, G. T.; "Pattern Recognition and Geometrical Complexity," *Proceedings of the International Conference on Pattern Recognition,* 2, 1324-1347, 1980.

[59] Toussaint, G. T. and Avis, D.; "On the Convex Hull Algorithm and its Application to Triangulation Problems," *Pattern Recognition,* 15, 23-29, 1982.

[60] Toussaint, G. T. and Sack, J.-R.; "Some New Results on Moving Polygons in the Plane," *Proceedings of the Robotic Intelligence and Productivity Conference,* 158-164, 1983.

[61] Widmayer, P. and Wood, D.; "A Time- and Space-Optimal Algorithm for Boolean Mask Operations for Orthogonal Polygons," *Computer Vision, Graphics, and Image Processing,* to appear, 1987.

[62] Widmayer, P., Wu, Y. F. and Wong, C. K.; "Distance Problems in Computational Geometry for Fixed Orientations," *Proceedings of the* 1^{st} *Annual ACM Symposium on Computational Geometry,* 186-195, 1985.

[63] Widmayer, P., Wu, Y. F., Schlag, M. D. F. and Wong, C. K.; "On Some Union and Intersection Problems for Polygons with Fixed Orientations," *Computing,* **36**, 183-197, 1986.

[64] Wood, D.; "An Isothetic View of Computational Geometry," *Computational Geometry,* (Toussaint, G., ed.), 429-459. North Holland, Amsterdam, 1985.

[65] Wood, D. and Yap, C. K.; "The Orthogonal Convex Skull Problem," *Discrete and Computational Geometry,* to appear, 1987.

[66] Yannakakis, M., Papadimitriou, C. H. and Kung, H. T.; "Locking Policies: Safety and Freedom from Deadlock," *Proceedings of the* 20^{th} *Annual IEEE Symposium on Foundations of Computer Science,* 286-297, 1979.