Restricted Orientation Visibility*

Sven Schuierer[†] Derick Wood [‡]

Dezember 1991

Abstract

Let $\mathcal{O}$ be some set of orientations, i.e., $\mathcal{O} \subseteq [0°, 360°)$. In this paper we look at the consequences of defining visibility based on curves that are monotone w.r.t. to the orientations in $\mathcal{O}$. We call such curves *$\mathcal{O}$-staircases*. Two points p and q in a polygon **P** are said to *$\mathcal{O}^s$-see* each other if there exists an $\mathcal{O}$-staircase from p to q that is completely contained in **P**. We investigate some of the structural properties of $\mathcal{O}^s$-visibility and then turn to the computation of the $\mathcal{O}^s$-kernel of a polygon.

The *$\mathcal{O}^s$-kernel* of a polygon **P** is then the set of all points which $\mathcal{O}^s$-see all other points. We show that the $\mathcal{O}^s$-kernel of a simple polygon can be obtained as the intersection of all $\{\theta\}^s$-kernels, with $\theta \in \mathcal{O}$. With the help of this observation we are able to develop an $O(n \log |\mathcal{O}|)$ algorithm to compute the $\mathcal{O}^s$-kernel in a simple polygon, for finite $\mathcal{O}$. We also show how to compute the *external $\mathcal{O}^s$-kernel* of a polygon in time $O(n + |\mathcal{O}|)$. Both algorithms can be combined to compute the $\mathcal{O}^s$-kernel of a polygon with holes in time $O(n^2 + n|\mathcal{O}|)$.

*This work was supported by the Deutsche Forschungsgemeinschaft under Grant No. Ot 64/8-1.

[†]Institut für Informatik, Universität Freiburg, Am Flughafen 17, D-79110 Freiburg, Fed. Rep. of Germany; email: schuiere@informatik.uni-freiburg.de

[‡]Department of Computer Science, University of Waterloo, Waterloo, Ontario, N2L3G1, Canada. email: dwood@dragon.waterloo.edu

Contents

1 Introduction

Restricted orientation visibility is based on the concept of *restricted orientation convexity.* Restricted orientation convexity tries to bridge the gap between Euclidean convexity and $\{0°, 90°\}$-convexity. A set **S** is called $\{0°, 90°\}$*-convex* or *orthogonally convex* if the intersection of **S** with any axis-parallel line is connected, i.e., a point, line segment or the empty set. $\{0°, 90°\}$-convexity is a well-studied area; see [CR89, Kei86, MRS88b, NLLW83, RC87, WY88]. In [Raw87, RW88a, RW87, RW88b] Rawlins and Wood take the idea of $\{0°, 90°\}$-convexity one step further by developing the theory of restricted orientation convexity or $\mathcal{O}$*-convexity.* Instead of considering only axis-parallel lines they allow lines with orientation in some fixed set $\mathcal{O}$. A set **S** is called $\mathcal{O}$*-convex* if the intersection of **S** with any line whose orientation is in $\mathcal{O}$ is connected. Note that restricted orientation convexity encompasses both $\{0°, 90°\}$-convexity—if $\mathcal{O} = \{0°, 90°\}$—and Euclidean convexity—if $\mathcal{O}$ is the set of all orientations—as special cases.

The framework of $\{0°, 90°\}$-convexity spawns a definition of visibility called *staircase visibility* which is based on $\{0°, 90°\}$-convex paths. Two points p and q in a set **S** are *staircase visible* from each other if there exists a $\{0°, 90°\}$-convex path from p to q that is completely contained in **S**. Staircase visibility has been considered by Reckhow and Culberson [CR89] and Motwani *et al.* [MRS88a]. Both papers deal with covering polygons with the minimum number of $\{0°, 90°\}$-starshaped sets, i.e., sets that contain one point p such that all other points are staircase visible from p. In the same way $\mathcal{O}$-convexity gives rise to a new definition of visibility we call $\mathcal{O}^s$*-visibility.* Two points p and q in a set **S** are $\mathcal{O}^s$*-visible* from each other or $\mathcal{O}^s$*-see* each other if there exists an $\mathcal{O}$-convex path from p to q that is completely contained in **S**. Note that $\mathcal{O}^s$-visibility again encompasses staircase visibility and Euclidean visibility as special cases.

In this paper we present an $O(n^2 + n|\mathcal{O}|)$ algorithm to compute the $\mathcal{O}^s$-kernel of a multiply connected polygon. The problem of computing the $\mathcal{O}^s$-kernel of a simple polygon has already been treated in [Raw87, Chapter 6] where a linear time algorithm to compute the $\{0°\}^s$-kernel of a simple polygon is presented.

In this paper we continue our investigations of generalized notions of visibility started in [SW91] in a much more specialized setting. While previously we were only concerned with the structural aspects of kernels under a very general notion of visibility, it is now our aim to give algorithms to compute kernels of polygons in the plane. Recall that the kernel of a set **S** is the set of points that see all other points in **S**. While for Euclidean visibility this is a well-studied problem where several optimal algorithms have been developed [CG88, LP79], the corresponding problem for $\{0°, 90°\}$-visibility has not been considered. We introduce a new definition of sight called $\mathcal{O}^s$*-visibility.* It is based on the framework of *restricted orientation convexity* which was first considered by G. Rawlins in [Raw87]. $\mathcal{O}^s$-visibility offers the opportunity to develop an algorithm that computes the Euclidean kernel as well as the $\{0°, 90°\}$-kernel depending on the input parameter $\mathcal{O}$ and that is competitive in both cases. The algorithm we present here is a first step in this direction.

An important requirement in the geometry of fixed orientations is to try to find algorithms that *do not* solve the problem separately, for each given orientation, and then put the obtained solutions together; instead one is interested in algorithms that exploit the "coherence" of the orientations in the solution of the problem. Our algorithms show that this can be done very efficiently in the computation of the $\mathcal{O}^s$-kernel.

The paper is organized as follows. We start off with a precise definition of $\mathcal{O}$-convexity and $\mathcal{O}^s$-visibility in the next section. In Section 3 we take a short glimpse at the topology of $\mathcal{O}$-convex sets in order to justify our assumption that the set of orientation $\mathcal{O}$ can be assumed to be closed. This is followed by the presentation of an algorithm for the computation of the $\mathcal{O}$-kernel of a polygon which is based on a second definition of $\mathcal{O}$-visibility which we did not mention above. In Section 5 we take a look at the structural properties of $\mathcal{O}^s$-visibility making use of the approach developed in [SW91]. Finally, in Section 6 we give an $O(n^2 + n|\mathcal{O}|)$ algorithm to compute the $\mathcal{O}^s$-kernel of a multiply connected polygon consisting of n vertices.

2 Definitions and Notation

In this section we give short definitions of the basic concepts we will make use of throughout the rest of this paper.

2.1 Topology

We start off with some notations of elementary topology. We are only concerned with the topology of the plane $\mathbb{E}^2$ which is based on the family of norms $\|\cdot\|_p$ where $p \in \mathbb{R}$, $p \geq 1$. It is defined as follows. If $x = (x_1, x_2)$ is a point in the plane, then the $\|\cdot\|_p$-norm of x is

$$\|x\|_p = (|x_1|^p + |x_2|^p)^{1/p}.$$

$\|\cdot\|_p$ induces the *L_p-metric (or d_p-metric)* which is defined by $L_p(x, y) = \|x - y\|_p$ where x and y are two points in the plane. Note that all norms in $\mathbb{E}^2$ induce the same topology. We will only consider the L_2-metric (d_2-metric)—which we also call the *Euclidean metric*—and the L_1-metric.

The interior of a set $\mathbf{S}$ is denoted by $int(\mathbf{S})$ and the closure*closure* of $\mathbf{S}$ by $\overline{\mathbf{S}}$. The *boundary* of $\mathbf{S}$ which is denoted by $\partial\mathbf{S}$ is the intersection of the closure of $\mathbf{S}$ with the closure of $\mathbb{E}^2 \setminus \mathbf{S}$.

Another important concept are *curves* or *paths*. A *curve* or *path* is an equivalence class $\mathcal{C}$ of continuous mappings $\varphi : I \longrightarrow \mathbb{E}^2$ where I is a real interval. $\varphi \in \mathcal{C}$ is called a *parametrization of* $\mathcal{C}$. We will mostly consider a curve as a point set but it should be kept in mind that a curve can always be endowed with a parametrization. Given two curves $\mathcal{P}$ and $\mathcal{Q}$ with two parametrizations $\varphi_{\mathcal{P}} : \langle a, b] \longrightarrow \mathbb{E}^2$ and $\varphi_{\mathcal{Q}} : [c, d\rangle \longrightarrow \mathbb{E}^2$ (the brackets $\langle$ and $\rangle$ are used to indicate that it doesn't matter if the left or right end point belongs to the interval or not) such that $\varphi_{\mathcal{P}}(b) = \varphi_{\mathcal{Q}}(c)$ we define the *concatenation of* $\mathcal{P}$ *and* $\mathcal{Q}$, denoted by $\mathcal{P} * \mathcal{Q}$, as the equivalence class of the mapping $\varphi_{\mathcal{P}*\mathcal{Q}} : \langle a, d-c+b\rangle \longrightarrow I\!E^2$ which is given by

$$\varphi_{\mathcal{P}*\mathcal{Q}}(x) = \begin{cases} \varphi_{\mathcal{P}}(x) & \text{if } x \in \langle a, b] \\ \varphi_{\mathcal{Q}}(x+c-b) & \text{if } x \in [b, d-c+b\rangle. \end{cases}$$

We say that a set $\mathbf{S}$ is *path-connected* if, for any two points x and y in $\mathbf{S}$, there exists a path $\mathcal{P}$ connecting x and y that is completely contained in $\mathbf{S}$. The sets we will consider in this paper are usually compact, path-connected sets. More precisely, they are the domain of a simple curve as defined below.

2.2 Jordan Curves

A curve $\mathcal{C}$ is called *simple* if there exists a parametrization $\varphi : \langle a, b\rangle \longrightarrow I\!E^2$ of $\mathcal{C}$ that is injective on $[a, b)$. We say $\mathcal{C}$ is closed if there exists a parametrization φ of $\mathcal{C}$ over $[0, 1]$ such that $\varphi(0) = \varphi(1)$. We call a simple closed curve according to [Rin75] a *Jordan curve.* By the famous Jordan Curve Theorem [Rin75, 40.1] a simple closed curve $\mathcal{J}$ partitions the plane into two parts, one bounded called *the interior of* $\mathcal{J}$ and denoted by $int(\mathcal{J})$, and the other unbounded called *the exterior of* $\mathcal{J}$ and denoted by $ext(\mathcal{J})$. We define the *domain* of $\mathcal{J}$ to be the union of $int(\mathcal{J})$ and $\mathcal{J}$ and denote it by $dom(\mathcal{J})$.

A simple, closed curve is a special case of a *multiply connected* curve[1] that we define as follows. A closed curve $\mathcal{C}$ is *multiply connected* if it is a finite collection $\{\mathcal{C}_1, \ldots, \mathcal{C}_m\}$ of simple curves that satisfy the following properties.

1. All $\mathcal{C}_i$ are pairwise disjoint, $1 \leq i \leq m$.
2. There is one curve, say $\mathcal{C}_1$, called the *enclosing cycle* such that $\mathcal{C}_i \subseteq int(\mathcal{C}_1)$, for $2 \leq i \leq m$. The remaining curves are called *holes.*
3. The interiors of the holes are pairwise disjoint.

We define the domain of $\mathcal{C}$ to be

$$dom(\mathcal{C}) = (int(\mathcal{C}_1) \setminus \bigcup_{2 \leq i \leq m} dom(\mathcal{C}_i)) \cup \mathcal{C}.$$

Clearly, the domain $\mathbf{D}$ of a multiply connected curve $\mathcal{C}$ is path-connected and compact with $\partial\mathbf{D} = \mathcal{C}$. The interior and exterior of $\mathcal{C}$ are defined in the natural way, $ext(\mathcal{C}) = I\!E^2 \setminus dom(\mathcal{C})$ and $int(\mathcal{C}) = int(dom(\mathcal{C}))$.

[1]We are slightly misusing the definition of curve here since a multiply connected curve is not the equivalence class of continuous mappings of an interval into the plane. Still this should not lead to misunderstandings.

2.3 Polygons

Polygons are a special case of multiply connected curves. Usually a *polygon* $\mathcal{P}$ is defined to be a closed curve that consists of a finite number of line segments, i.e., there is a parametrization of $\mathcal{P}$ that is piecewise linear. The *edges* of a polygon $\mathcal{P}$ are the maximal line segments it consists of. Many of the results in this paper are concerned with *simple polygons*, i.e., piecewise linear Jordan curves. The *vertices* of a simple polygon are defined to be the intersection points of the edges.

We call a polygon $\mathcal{P}$ *multiply connected* if $\mathcal{P}$ is a multiply connected curve and all cycles of $\mathcal{P}$ are simple polygons. If not mentioned explicitly otherwise, we use the term *polygon* to denote a simple polygon.

Since we are mostly interested in the properties of the domain of a polygon instead of the curve itself, we will also use the term polygon to denote the *domain* of a (multiply connected) polygon. In order to distinguish between curves and domains, we use capital calligraphic letters like $\mathcal{P}$ for the first and capital bold-face letters like $\mathbf{P}$ for the latter.[2] So whenever we want to refer to the curve that encloses a given domain $\mathbf{D}$ we will speak of the *boundary of* $\mathbf{D}$. If we speak of the vertices and edges of $\mathbf{P}$, we, of course, mean the vertices and edges of the boundary of $\mathbf{P}$. If $\mathbf{P}$ is a multiply connected polygon, we call the domain of the enclosing cycle $\mathcal{Q}$ the *enclosing polygon*. The term *holes* is also used to denote the the domains of the polygons interior to the enclosing polygon.

As for representing a polygon $\mathbf{P}$ we will assume that $\mathbf{P}$ is given as a list of its vertices sorted according to a counterclockwise traversal of $\partial\mathbf{P}$.

2.4 Halfplanes, Lines, and Line Segments

We now turn to the definitions of the geometric objects we will be dealing with. As mentioned above we are only concerned with the plane $\mathbb{E}^2$. Sometimes it is helpful to consider $\mathbb{E}^2$ as a vector space. So, for a point p in $\mathbb{E}^2$, we denote the corresponding vector by $\vec{p}$.

An oriented line l is a set of points in the plane of the form $\{p \in \mathbb{E}^2 \mid \vec{p} = \vec{p_1} + \lambda\vec{p_2}, \lambda \in \mathbb{R}\}$, with $p_1, p_2 \in \mathbb{E}^2$, together with a unit vector $\vec{u}$ that is contained in the linear hull of $\vec{p_2}$. We call $\vec{u}$ the *associated vector* of l. We say two lines are *parallel* if their associated vectors are collinear. Let $\vec{u}^\perp$ be the unit vector that is obtained if we turn $\vec{u}$ for $90°$ counterclockwise. We say a point p is (strictly) to the left of l if p has a non-negative (positive) second coordinate in the $(\vec{u}, \vec{u}^\perp)$-coordinate system with origin somewhere on l. Similarly, we say p is (strictly) to the right of l if p has a non-positive (negative) second coordinate. We denote the closed halfplane of all points that are to the left of l by $h^+(l)$ and the closed halfplane to the right of l by $h^-(l)$. Since $\vec{u}$ is completely specified if its angle with the x-axis is known, we will not distinguish between angles and unit vectors. We use the term *orientation* to denote either an unit vector or an angle.

[2]Note that calligraphic letters may also be used to denote a family of sets.

We denote the line segment between two distinct points p and q by $\overline{pq}$. In order to denote the oriented line from p to q we use the notation $\ell(p,q)$ and, similarly, for the ray starting in p that goes through q the notation $r(p,q)$. If we are given a point p and an orientation θ, $\ell(p,\theta)$ denotes the oriented line through p in direction θ and, in the same way, $r(p,\theta)$ denotes the ray starting in p in the direction of θ.

2.5 Computational Foundations

The algorithms developed in this thesis are analysed according to their *asymptotic worst-case behaviour*. The underlying model of computation we make use of here is the *real RAM (random-access machine) model*. A real RAM consists of an arithmetic unit and an unbounded number of storage cells each of which is capable of storing one real number. The arithmetic unit consists of a constant number of registers each of which can also store a real number. The commands to control a real RAM resemble the assembly code of existing processors. So we assume that the following operations can be carried out in one step on a real RAM.

(i) to store/load a real number in/from a storage location from/to a register;

(ii) to execute a conditional or unconditional jump;

(iii) to compare two reals;

(iv) to add, subtract, multiply or divide two reals;

(v) to compute the k^{th} root or the k^{th} power of a real or to perform a trigonometric operation.

Since the algorithms in this thesis are not formulated as real RAM-programs, the analysis of the running time (depending on the input size) can only be accurate up to a constant factor, i.e., we are only interested in the growth rate of an algorithm. This can be conveniently expressed in the O- and Ω-notation. We define for a function $f : \mathbb{N} \longrightarrow \mathbb{R}_+$ (see [OW90])

$$O(f) \quad = \quad \{g \mid \exists c_1, c_2 > 0 : \text{for all } \ n \in \mathbb{N} : g(n) \leq c_1 f(n) + c_2\}$$

and

$$\Omega(f) \quad = \quad \{h \mid \exists c > 0 : \text{for infinitely many } n \in \mathbb{N} : h(n) \geq c f(n)\}.$$

If a function g is in $O(f)$, we also write $g(n) = O(f(n))$ and in the same way if $h \in \Omega(f)$, then we write $h(n) = \Omega(f(n))$. $\Theta(f)$ is defined as the intersection of $O(f)$ and $\Omega(f)$. Again we write $g(n) = \Theta(f(n))$ if $g \in \Theta(f)$. Finally, we define

$$o(f) \quad = \quad \{g \mid \lim_{n \longrightarrow \infty} g(n)/f(n) = 0\}.$$

We use the *algebraic computation tree* model to prove lower bounds for the problems we consider. It can be specified as follows. The statements s_i, $1 \leq i \leq m$, of an algebraic computation tree program are allowed to be only of the following four types (we assume that the statement we currently consider has index i)

1. $z \leftarrow u \bullet v$ with $\bullet \in \{+, -, \times, /\}$; go to statement s_{i+1};
2. If $u : 0$ (where $: \in \{=, <, >\}$), then go to statement s_j else to statement s_k with j, $k > i$;
3. Halt and output "yes";
4. Halt and output "no".

The variables u and v above have to be either results of previous computations or input parameters. Note that the algebraic computation tree is a weaker model of computation than the real RAM. The time needed by an algebraic computation tree program is defined to be the longest possible executed sequence of statements starting with s_1. A problem is said to have complexity at least $T(n)$ if the time needed by any algebraic computation tree program is $\geq T(n)$.

It has been shown that the following problem takes time $\Omega(n \log n)$ in the algebraic computation tree model [Yao89].

The Integer Element Uniqueness Problem: Given n integers decide whether two of them are equal.

This will be used in Section 4.3 to prove a lower bound on the problem of finding the $\{0°\}$-kernel for polygons with holes.

2.6 Restricted Orientation Convexity

Restricted orientation convexity can be understood as a generalization of the usual notion of convexity in the plane. There are many ways to define the family of convex sets. The definition we give leads in a natural way to the family of restricted oriented convex sets. We say a set is *convex* if its intersection with any line in the plane is connected. Recall that we regard the empty set to be connected as well.

If we consider only the intersection with a certain class of lines and still require the intersection of a set to have a connected intersection with each line in the class, the resulting family of sets shares some of the properties of convex sets while others are lost. In the following we will investigate the properties of the family of sets if we restrict the *orientation* of the lines we require a set to have a connected intersection with.

If we are given an oriented line l in the plane by $l = \{p \in \mathbb{E}^2 \mid \vec{p} = \vec{p_1} + \lambda \vec{p_2}\}$ with associated vector $\vec{u}$, we define its *orientation* to be the angle that $\vec{u}$ forms with the x-axis and denote it by $\Theta(l)$. Of course, we can speak in the same way of the orientation of a line

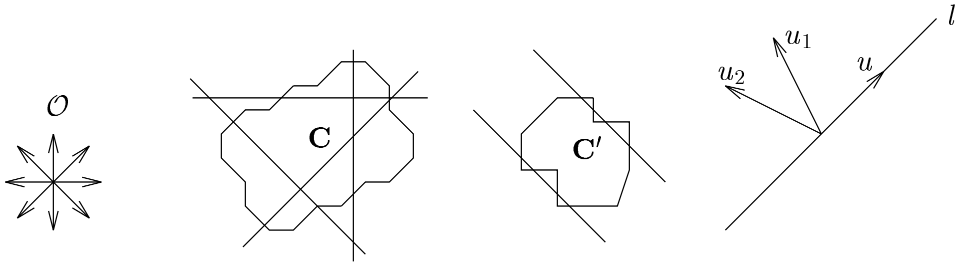

(a) (b)

Figure 1: **C** *is $\mathcal{O}$-convex while* **C′** *is not.*

segment or a ray. The orientation of a point p is defined to be the orientation of the line segment from the origin $\mathbf{0}$ to p.

In order not to have to deal with angles that are greater than 360°, we assume from now on that the addition and subtraction of two orientations is done modulo 360°. The restricted oriented convex sets can now be defined as follows.

Definition 2.1 *Let $\mathcal{O}$ be a subset of* $[0°, 360°)$. *A set* $\mathbf{C} \subseteq \mathbb{E}^2$ *is* $\mathcal{O}$-convex *if* $l \cap \mathbf{C}$ *is connected, for all lines* l *with* $\Theta(l) \in \mathcal{O}$.

An example is shown in Figure 1. The first thing to note about this definition is that we can assume that $\mathcal{O}$ is symmetric w.r.t. the origin, that is, if orientation θ is in $\mathcal{O}$, then a set is $\mathcal{O}$-convex if and only if it is $\mathcal{O} \cup \{\theta + 180°\}$-convex. We denote the orientation $\theta + 180°$ by θ^{-1}. So there is no loss in generality if we assume that, for all $\theta \in \mathcal{O}$, we also have $\theta^{-1} \in \mathcal{O}$. We also want to be able to compare two orientations and say which one comes (counterclockwise) before or after the other. In order to do so for two given orientations θ_1 and θ_2, we consider an oriented line l through the origin such that the normal vectors $\vec{u_1}$ and $\vec{u_2}$ that correspond to θ_1 and θ_2 are to the left of l. We now say that θ_1 is before θ_2 if the angle enclosed by the associated vector $\vec{u}$ of l with $\vec{u_1}$ is less than the angle between $\vec{u_2}$ and $\vec{u}$. More formally, we define

$$\theta_1 \leq \theta_2 \text{ if } \theta_2 - \theta_1 \leq 180°.$$

For an illustration refer to Figure 1b. Since $\mathcal{O}$ always contains either both orientations θ and θ^{-1} or none of them, we use the notation $|\mathcal{O}|$ to denote *half* the cardinality of $\mathcal{O}$. Furthermore, we will only specify the orientations in $[0°, 180°)$ to define a particular $\mathcal{O}$ though it should be kept in mind that $\mathcal{O}$ also contains the opposite orientations in $[180°, 360°)$.

Another consequence of the fact that $\mathcal{O}$ is symmetric w.r.t. to the origin is that given some orientation $\theta \in \mathcal{O}$ we can speak of the *next* orientation in $\mathcal{O}$ w.r.t. θ which is defined as the infimum of all orientations θ' in $\mathcal{O}$ with $\theta \leq \theta'$.[3] In the same way we can speak of the previous orientation in $\mathcal{O}$ w.r.t. θ. We say a range (θ_1, θ_2) is *$\mathcal{O}$-free* if $(\theta_1, \theta_2) \cap \mathcal{O} = \emptyset$. A range (θ_1, θ_2) is called a *maximal $\mathcal{O}$-free range* if (θ_1, θ_2) is $\mathcal{O}$-free and there is no other range (θ_1', θ_2') that is also $\mathcal{O}$-free and contains (θ_1, θ_2). If θ is some orientation in $[0°, 360°)$, the *maximal $\mathcal{O}$-free range of* θ is the maximal $\mathcal{O}$-free range that contains θ or if such a range does not exist, the emptyset. In the same way we will speak of the maximal $\mathcal{O}$-free range of a line, ray, line segment, or point which is, of course, meant to denote the maximal $\mathcal{O}$-free range of the orientation of the geometric object. Two orientations θ_1, θ_2 in $\mathcal{O}$ are called *adjacent* if (θ_1, θ_2) is $\mathcal{O}$-free.

Since $\mathcal{O}$-free ranges play an important role in the theory of $\mathcal{O}$-convex sets, it is important to note that we can map any two linearly independent unit vectors $\vec{u_1}$ and $\vec{u_2}$ that correspond to two orientations θ_1 and θ_2 with a bijective linear mapping f to the unit vectors $\vec{e_1}$ and $\vec{e_2}$ that have orientation $0°$ and $90°$, respectively. Since f is linear and bijective, all incidences of lines are preserved and none are introduced. Hence, we may assume in the following that if we consider a particular non-empty $\mathcal{O}$-free range (θ_1, θ_2), then $\theta_1 = 0°$ and $\theta_2 = 90°$. Another property of $\mathcal{O}$-free ranges is that they are only useful if $|\mathcal{O}| \geq 2$. Hence, it will be often necessary to treat the case $|\mathcal{O}| = 1$ separately.

It is easy to see that $\mathcal{O}$-convex sets are closed under intersection and, thus, form a convexity space. We define the *$\mathcal{O}$-hull* of a set $\mathbf{S}$ by

$$\mathcal{O}\text{-}hull(\mathbf{S}) = \bigcap \{\mathbf{C} \subseteq \mathbb{E}^2 \mid \mathbf{C} \text{ contains } \mathbf{S} \text{ and } \mathbf{C} \text{ is } \mathcal{O}\text{-convex}\}.$$

Obviously, we have that if $\mathcal{O}$ equals $[0°, 360°)$, then the $\mathcal{O}$-convex sets are exactly the convex sets in the usual sense.

2.6.1 Two Definitions of $\mathcal{O}$-Visibility

As we already mentioned before it is the aim of this paper to investigate some properties of visibility in the context of $\mathcal{O}$-convexity. So the next thing we need is a "natural" generalization of the usual line segment visibility. Our first approach is based on the $\mathcal{O}$-hull of two points.

Definition 2.2 *Let $\mathbf{P} \subseteq \mathbb{E}^2$ and $p, q \in \mathbf{P}$; we say p and q $\mathcal{O}$-*see *each other or are $\mathcal{O}$-*visible *from each other if $\mathcal{O}$-hull$(\{p, q\}) \subseteq \mathbf{P}$.*

Although this definition is straightforward once $\mathcal{O}$-convexity has been established, the resulting notion of visibility has some rather undesirable properties as the $\mathcal{O}$-hull of two points can be disconnected (see Figure 2). In general we can say that, for any two points $p, q \in \mathbb{E}^2$, we have that

$$\mathcal{O}\text{-}hull(\{p, q\}) = \begin{cases} \overline{pq} & \text{if } \Theta(\overline{pq}) \in \mathcal{O} \\ \{p, q\} & \text{if } \Theta(\overline{pq}) \notin \mathcal{O}. \end{cases}$$

[3]Note that since we have θ and θ^{-1} in $\mathcal{O}$, there always is one orientation in $\mathcal{O}$ that greater or equal to some other given orientation θ' and there is also one orientation that is less or equal to θ'.

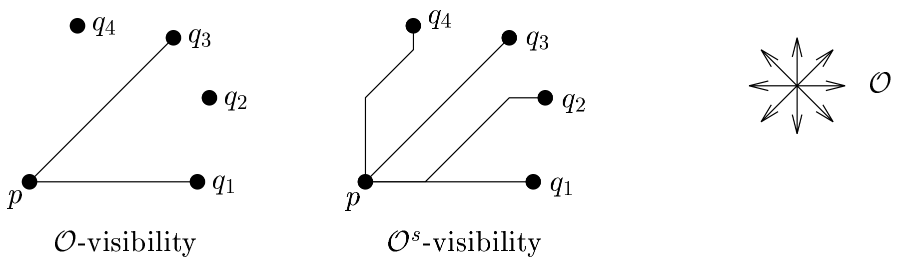

Figure 2: *Two definitions of visibility in the context of $\mathcal{O}$-convexity.*

Taking the above equation as a definition, the $\mathcal{O}$-convex sets can also be characterized as those sets $\mathbf{C}$ that contain $\mathcal{O}$-*hull*$(\{p, q\})$, for all $p, q \in \mathbf{C}$.

On the other hand, connectedness is a vital part of "normal" visibility and, therefore, generalizations should try to preserve this strong property. This leads to a second definition of visibility in the context of $\mathcal{O}$-convexity (see Figure 2).

Definition 2.3 *Let* $\mathbf{S}$ *be a subset of the plane and p, q be two points in* $\mathbf{S}$*; we say p and q* $\mathcal{O}^s$-see *each other or are* $\mathcal{O}^s$-visible *from each other if there is an* $\mathcal{O}$-convex path *from p to q that is completely contained in* $\mathbf{S}$.

A subset $\mathbf{S}$ of the plane is called $\mathcal{O}^s$*-convex* if any two points of $\mathbf{S}$ $\mathcal{O}^s$-see each other. In Theorem 5.5.2 of [Raw87] it is shown that $\mathbf{S}$ is $\mathcal{O}^s$-convex if and only if $\mathbf{S}$ is $\mathcal{O}$-convex and connected. $\mathcal{O}^s$-convexity is also a natural generalization of the notion of *ortho-convexity* (i.e., $\mathcal{O} = \{0°, 90°\}$) as defined in [RW88b].
From now on we will call an $\mathcal{O}$-convex path an $\mathcal{O}$*-stairsegment*. An $\mathcal{O}$-stairsegment that consists of a finite number of edges is called an $\mathcal{O}$*-staircase*.

There is an interesting and rather obvious connection to monotone paths. Let $\mathcal{O}^{\perp}$ denote the set of orientations which are orthogonal to those in $\mathcal{O}$. Then, a path $\mathcal{P}$ is $\mathcal{O}$-convex if and only if $\mathcal{P}$ is monotone with respect to all orientations in $\mathcal{O}^{\perp}$.

We are mostly interested in the two kernels of a planar set that arise from the above definitions. Recall that the kernel of a set $\mathbf{P}$ is the set of points that see all of $\mathbf{P}$. We will denote them by $\mathcal{O}$*-kernel* and $\mathcal{O}^s$*-kernel* to distinguish between the two kinds of visibility we are concerned with.

2.6.2 Previous Results

Since the work presented here continues the investigations of [Raw87], it will be sometimes necessary to refer to some of the results of [Raw87] as we did above with Theorem 5.5.2. We repeat all the results we reference for the convenience of the reader since we make use of them in several proofs. Since the terminology used in [Raw87] differs from the one developed here, we allow ourselves to adapt the formulation of the result where adequate.

Definition: A set $\mathbf{P}$ is said to have *span* $[\theta_1, \theta_2]$ (where $\theta_1 \leq \theta_2$) if, for all p, $q \in \mathbf{P}$, where $p \neq q$, $\Theta(\overline{pq}) \in [\theta_1, \theta_2]$ and this is true for no smaller range.

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

Let $h^+(p, \theta)$ denote the halfplane to the left of the θ-oriented line through p.

Lemma 4.3.1 *Let p and q be two points in the plane with $\Theta(\overline{pq}) \notin \mathcal{O}$. If $|\mathcal{O}| \geq 2$, then*

$$\mathcal{O}\text{-}\|[p, q] = h^+(p, \theta_1) \cap h^+(q, \theta_2) \cap h^+(q, \theta_1^{-1}) \cap h^+(p, \theta_2^{-1})$$

if (θ_1, θ_2) is the maximal $\mathcal{O}$-free range of $\overline{pq}$.

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

Theorem 5.3.1 *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.*

Theorem 5.5.1 *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 two consecutive $\mathcal{O}$-extremal points are maximal $\mathcal{O}$-stairsegments.*

Theorem 5.5.2 *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.*

Theorem 6.1.3 *$\mathbf{P}$ is strong-$\mathcal{O}$-convex if and only if $\mathbf{P}$ is convex and $\mathcal{O}$-oriented.*

Lemma 6.4.1 *Given a polygon $\mathbf{P}$, $\{0°\}^s$-kernel$(\mathbf{P})$ lies above any reflex maximum and below any reflex minimum.*

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

$$(\textstyle\bigcup_i \mathcal{O}_i)^s\text{-}kernel(\mathbf{P}) = \displaystyle\bigcap_i \mathcal{O}_i^s\text{-}kernel(\mathbf{P}).$$

Lemma 6.4.2 *Let $\mathbf{P}$ be a polygon with at least one hole and let p be in $\mathbf{P}$. If $\ell(p, 0°)$ intersects a hole, then $p \notin \{0°\}^s$-kernel$(\mathbf{P})$.*

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

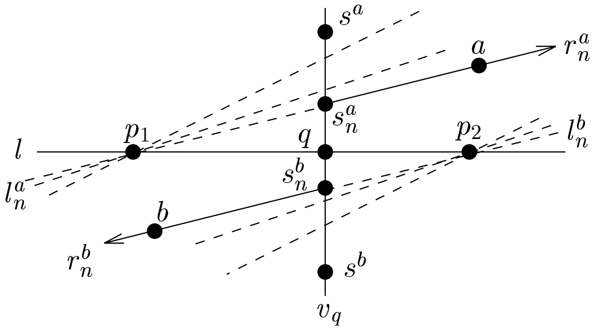

Figure 3: *Illustration to the proof of Theorem 3.1.*

3 The Topology of $\mathcal{O}$-convex Sets

Before we start to investigate the properties of $\mathcal{O}$- and $\mathcal{O}^s$-visibility, we want to take a look at the topological properties of the family of $\mathcal{O}$-convex sets. This aspect is not treated in [Raw87] though we feel that it also illustrates some of the more interesting properties of $\mathcal{O}$-convex sets. In this section we abandon our convention that we only consider path-connected sets. Hence, the term connected is used in its usual meaning.

In [Raw87] it is assumed that $\mathcal{O}$ is always a closed set. This leads to the question if a (useful) theory of $\mathcal{O}$-convexity can be developed if $\mathcal{O}$ is allowed to be any subset of $[0°, 360°)$ (see Section 8.1.10 of [Raw87]). In the following theorem we can show that as long as we are concerned with closed connected sets the notion of $\mathcal{O}$-convexity and $\overline{\mathcal{O}}$-convexity do not differ. Hence, there is no loss of generality in assuming $\mathcal{O}$ to be a closed set. Note, in particular, that our definition of $\mathcal{O}^s$-visibility is based on $\mathcal{O}$-stairsegments which are compact connected sets.

Theorem 3.1 *Let* $\mathbf{C}$ *be a closed connected set in the plane and* $\mathcal{O}$ *a set of orientations.* $\mathbf{C}$ *is* $\mathcal{O}$*-convex if and only if* $\mathbf{C}$ *is* $\overline{\mathcal{O}}$*-convex.*

Proof: Since $\mathcal{O} \subseteq \overline{\mathcal{O}}$, the if-part of the claim is immediate. So assume $\mathbf{C}$ is $\mathcal{O}$-convex and let l be a line with $\Theta(l) = \theta \in \overline{\mathcal{O}} \setminus \mathcal{O}$. Let p_1 and p_2 be two points of $l \cap \mathbf{C}$. Since $\theta \in \overline{\mathcal{O}}$, there is a sequence $(\theta_n)_{n \in \mathbb{N}}$ of orientations in $\mathcal{O}$ that converges to θ. W.l.o.g. assume that $\theta = 0°$ and that all θ_n are greater than $0°$. Further label p_1 and p_2 such that p_1 is to the left of p_2. Let q be some point in $\overline{p_1 p_2}$. We want to show that $q \in \mathbf{C}$ (see Figure 3).

Let v_q be the vertical line through q; furthermore, let s^a be the closest point to q in $\mathbf{C} \cap v_q$ above q and s^b the closest point below q. Since $\mathbf{C} \cap v_q$ is closed and $\mathbf{C}$ is connected, one of s^a and s^b clearly exist. If one of them does not exist, we assume it to be an arbitrary point above or below q, respectively. If q equals either s^a or s^b, then we are done. So assume $q \notin \mathbf{C}$.

Choose n_0 big enough such that, for all $n \geq n_0$, all the lines l_n^a through p_1 with orientations θ_n intersect v_q below s^a and all the lines l_n^b through p_2 with orientations θ_n

intersect v_q above s^b. We denote the intersection of v_q and l_n^a by s_n^a and the intersection of v_q and l_n^b by s_n^b. Let r_n^a be the ray starting in s_n^a with orientation θ_n, i.e., r_n^a does not contain p_1. r_n^b is defined analogously, i.e., r_n^b has orientation θ_n^{-1}. Let $\mathcal{P}$ be the biinfinite curve $r_n^a * \overline{s_n^a s_n^b} * r_n^b$. It divides the plane into two disjoint open sets one of which contains p_1 and the other contains p_2. Since $\mathbf{C}$ is connected, it intersects $\mathcal{P}$.

$\mathbf{C}$ cannot intersect $\overline{s_n^a s_n^b}$ because of the definition of s^a and s^b and the way we chose n. Therefore, $\mathbf{C}$ either intersects r_n^a or r_n^b. Assume $\mathbf{C}$ intersects r_n^a and let a be in the intersection of $\mathbf{C}$ and r_n^a. Because $a \in l_n^a$, we have $\Theta(\overline{p_1 a}) = \theta_n \in \mathcal{O}$ and the $\mathcal{O}$-convexity of $\mathbf{C}$ yields $\overline{p_1 a} \subseteq \mathbf{C}$. Furthermore, since $a \in r_n^a$, we have $s_n^a \in \overline{p_1 a}$ and, hence, $s_n^a \in \mathbf{C}$. This contradicts our choice of s^a as the closest point on v_q to q. Therefore, q belongs to $\mathbf{C}$.

Since q is arbitrary in $\overline{p_1 p_2}$, we have $\overline{p_1 p_2} \subseteq \mathbf{C}$. Hence, $l \cap \mathbf{C}$ is path-connected and this immediately yields the claim. □

Since the properties of $\mathcal{O}$-convexity we are interested in are usually related to closed connected sets, we will assume from now on that $\mathcal{O}$ is closed. The necessity of connectedness in the above claim is established by two points which lie on a horizontal line and which are $\mathcal{O}$-convex, for any set of orientations $\mathcal{O}$ that is contained in $[0°, 360°) \setminus \{0°, 180°\}$. If we assume the sets we are dealing with to be *compact*, i.e., closed and bounded as well, we get some interesting properties from the *Hausdorff Metric.*

Definition 3.1 (R. Schneider: Convex Geometry Section 1.8 [Sch89]) *Let $\mathbf{C}_1$ and $\mathbf{C}_2$ be two compact sets in the plane. We denote the* Hausdorff distance *of $\mathbf{C}_1$ and $\mathbf{C}_2$ by $d_H(\mathbf{C}_1, \mathbf{C}_2)$ and define it to be the maximum of $\sup_{p \in \mathbf{C}_1} \inf_{q \in \mathbf{C_2}} d_2(p, q)$ and $\sup_{q \in \mathbf{C}_2} \inf_{p \in \mathbf{C}_1} d_2(p, q)$.*

It can be shown that the Hausdorff distance defines a metric in the space $\mathcal{C}^2$ of compact sets in the plane and that the metric space $(\mathcal{C}^2, d_H)$ is complete. Our main intrest is in the properties of the subset of all connected $\mathcal{O}$-convex sets of $(\mathcal{C}^2, d_H)$ which we denote by $\mathcal{O}$-$\mathcal{C}^2$.

Theorem 3.2 *$\mathcal{O}$-$\mathcal{C}^2$ is a closed subset of $\mathcal{C}^2$.*

Proof: To show the claim let $(\mathbf{C}_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{O}$-$\mathcal{C}^2$ that converges to $\mathbf{C} \in \mathcal{C}^2$. We have to show that $\mathbf{C}$ is $\mathcal{O}$-convex and connected.

First we will prove that $\mathbf{C}$ is $\mathcal{O}$-convex. Let l be an $\mathcal{O}$-line that has a non-empty intersection with $\mathbf{C}$ and let p and q in $l \cap \mathbf{C}$. W.l.o.g. we assume that $\Theta(l) = 0°$ and that p is to the left of q. Let $(p_n)_{n \in \mathbb{N}}$ and $(q_n)_{n \in \mathbb{N}}$ be sequences with $p_n, q_n \in \mathbf{C}_n$ and $\lim_{n \to \infty} p_n = p$ and $\lim_{n \to \infty} q_n = q$. Let r be a point in $\overline{pq}$ and let v_r be the vertical line through r. We first want to construct a sequence $(r_n)_{n \in \mathbb{N}}$ with $r_n \in \mathbf{C}_n$ that converges to r (for illustration see Figure 4).

There is a $n_0 \in \mathbb{N}$ such that, for all $n \geq n_0$, we have

$$\max\{d_2(p_n, p), d_2(q_n, q)\} < \min\{d_2(p, r), d_2(q, r)\}$$

and, hence, p_n is on one side of v_r and q_n on the other. For $n < n_0$ we choose r_n to be p_n.

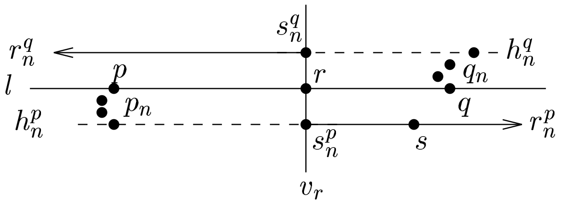

Figure 4: *The construction of a sequence the converges to r.*

Now let $n \geq n_0$ and consider the points p_n and q_n. W.l.o.g. assume that p_n is lower than q_n. We consider the two horizontal lines h_n^p through p_n and h_n^q through q_n. The intersection point of v_r and h_n^p is denoted by s_n^p and the intersection point of v_r and h_n^q is denoted by s_n^q, respectively. Let further r_n^p be the horizontal ray to the right that originates in s_n^p, i.e., $\Theta(r_n^p) = 0°$. Analogously let r_n^q be the horizontal ray to the left that originates in s_n^q, i.e., $\Theta(r_n^q) = 180°$. Now let $\mathcal{P}$ be the biinfinite curve $r_n^p * \overline{s_n^p s_n^q} * r_n^q$. It divides the plane into two disjoint open regions one of which contains p_n and the other contains q_n. Since $\mathbf{C}_n$ is connected, $\mathbf{C}_n$ intersects $\mathcal{P}$.

If $\mathbf{C}_n$ intersects $\overline{s_n^p s_n^q}$, then we can choose r_n to be a point in $\mathbf{C}_n \cap \overline{s_n^p s_n^q}$. If $\mathbf{C}$ intersects r_n^p in, say, s, then $\overline{sp_n}$ is contained in $\mathbf{C}_n$ since $\mathbf{C}_n$ is $\{0°\}$-convex. This yields $s_n^p \in \mathbf{C}_n$ and we can choose r_n to be s_n^p. If $\mathbf{C}$ intersects r_n^q, then we have $s_n^q \in \mathbf{C}_n$ and we can choose r_n to be s_n^q. In any case we have that $d_2(r_n, r) < \max\{d_2(p, p_n), d_2(q, q_n)\}$ and, hence, $r_n \to r$.

With the help of $(r_n)_{n \in \mathbb{N}}$ we can now contruct a sequence $(r'_k)_{k \in \mathbb{N}}$ in $\mathbf{C}$ that converges to r. For all $k \in \mathbb{N}$, there is a $n_k \in \mathbb{N}$ with $d_H(\mathbf{C}_{n_k}, \mathbf{C}) < 1/(2k)$ and $d_2(r_{n_k}, r) < 1/(2k)$. Hence, there is a point r'_k in $\mathbf{C}$ with $d_2(r_{n_k}, r'_k) < 1/(2k)$ and, therefore, $d_2(r, r'_k) \leq d_2(r, r_{n_k}) + d_2(r_{n_k}, r'_k) < 1/k$. So, the sequence $(r'_k)_{k \in \mathbb{N}}$ converges to r. Since $\mathbf{C}$ is closed, this implies $r \in \mathbf{C}$ and, thus, $\overline{pq} \subseteq \mathbf{C}$.

Now we proof that $\mathbf{C}$ is connected. For assume that this is not the case; then, there are two disjoint, non-empty sets $\mathbf{C}'$ and $\mathbf{C}''$ which are open and closed in $\mathbf{C}$. Since they are closed in $\mathbf{C}$, they are compact in $\mathbb{E}^2$ and, hence, the Hausdorff distance between $\mathbf{C}'$ and $\mathbf{C}''$ is positiv, say $\varepsilon > 0$. Therefore, the open sets $\mathbf{D}' = \bigcup\{N_{\varepsilon/2}(p) \mid p \in \mathbf{C}'\}$ and $\mathbf{D}'' = \bigcup\{N_{\varepsilon/2}(p) \mid p \in \mathbf{C}''\}$ are disjoint. Now choose n big enough such that $d_H(\mathbf{C}_n, \mathbf{C}) < \varepsilon/2$. $\mathbf{C}_n$ is then contained in $\mathbf{D}' \cup \mathbf{D}''$ and $\mathbf{D}'$ and $\mathbf{D}''$ both have a non-empty intersection with $\mathbf{C}_n$. This contradicts the connectedness of $\mathbf{C}_n$. $\square$

It should be noted that the family of all path-connected, compact, and $\mathcal{O}$-convex sets is not a closed subset of $\mathcal{C}^2$. To see this just consider the graph of the function $\sin(1/x)$. If we restrict the x-values to the interval $[1/n, 1]$ the graph is a compact, path-connected, and $\{90°\}$-convex set. But if $n \to \infty$, the limit consists of the line segment $\overline{(0,0)(0,1)}$ together with the graph of $\sin(1/x)$ over the interval $(0, 1]$ and it is well known that this is a connected but not path-connected set.

On the other hand, if we give up connectivity the claim is false as well. As a counterexample consider the sequence $\mathbf{S}_n$ of sets which is defined by $\mathbf{S}_n = \{(0, 1/n), (1, 1/(n+1))\}$.

Each $\mathbf{S}_n$ is $\{0°\}$-convex but the sequences converges to $\{(0,0),(1,0)\}$ which is not $\{0°\}$-convex.

Another interesting question is the topological nature of the $\mathcal{O}$*-hull*. We have the following result on this.

Theorem 3.3 *If $\mathcal{O}$ is closed and $\mathbf{A}$ is a connected, compact set, then $\mathcal{O}$-hull$(\mathbf{A})$ is also connected and compact.*

Proof: The connectedness is just Lemma 5.2.2 of [Raw87]. It can be easily seen, that $\mathcal{O}$-$hull(\mathbf{A})$ is bounded. So we just focus on the closedness of $\mathcal{O}$-$hull(\mathbf{A})$. Let $(p_n)_{n\in\mathbb{N}}$ be a converging sequence in $\mathcal{O}$-$hull(\mathbf{A})$ with $p_n \rightarrow p$. By Theorem 5.3.1 of [Raw87] there exist points q_n and r_n in $\mathbf{A}$ with $p_n \in \mathcal{O}$-$hull(\{q_n, r_n\})$ and $\Theta(\overline{q_n r_n}) \in \mathcal{O}$. Since $\mathbf{A}$ is compact, there exists a converging subsequence $(q_{n_{k_1}})_{k_1\in\mathbb{N}}$ of q_n. For the same reason there exists another converging subsequence $(r_{n_{k_2}})_{k_2\in\mathbb{N}}$ of r_n with $n_{k_2} \in \{n_{k_1} \mid k_1 \in \mathbb{N}\}$. Let $q \in \mathbf{A}$ be the limit of $q_{n_{k_2}}$ and $r \in \mathbf{A}$ the limit of $r_{n_{k_2}}$. Since $\mathcal{O}$ is closed, the sequence $\Theta(\overline{q_{n_{k_2}} r_{n_{k_2}}})$, $k_2 \in \mathbb{N}$, converges to a value in $\mathcal{O}$. And it is easy to see that $p \in \mathcal{O}$-$hull(\{q,r\}) = \overline{qr} \subseteq \mathcal{O}$-$hull(\mathbf{A})$. □

Because of the above results, we will assume throughout the rest of this chapter that $\mathcal{O}$ is a closed subset of $[0°, 360°)$.

4 Computing the $\mathcal{O}$-Kernel of a Polygon

In this section we present an algorithm to compute the $\mathcal{O}$-kernel of a polygon if $\mathcal{O}$ is finite. The approach we take is the following. First we show that the $\mathcal{O}$-kernel can be computed as the intersection of the $\{\theta\}$-kernels, taken over all $\theta \in \mathcal{O}$. Since these intersections can be computed efficiently, we concentrate on the computation of the $\mathcal{O}$-kernel if $\mathcal{O}$ consists of only one orientation. Clearly, we can assume w.l.o.g. that $\mathcal{O} = \{0°\}$. We give a linear time algorithm to compute the $\{0°\}$-kernel of a simple polygon and show how to adapt it to handle multiply connected polygons.

We first turn to proving the above mentioned intersection property of the $\mathcal{O}$-kernel.

Lemma 4.1 (Intersection Lemma) *Given a set $\mathbf{P} \subseteq \mathbb{E}^2$ and sets of orientations $\mathcal{O}_i$, $i \in I$[4], with $\mathcal{O} = \bigcup_{i\in I} \mathcal{O}_i$ we have that*

$$\mathcal{O}\text{-}kernel(\mathbf{P}) = \bigcap_{i\in I} \mathcal{O}_i\text{-}kernel(\mathbf{P}).$$

Proof: Let $p \in \mathcal{O}$-$kernel(\mathbf{P})$. We obtain $\mathcal{O}$-$hull(\{p,q\}) \subseteq \mathbf{P}$, for all $q \in \mathbf{P}$. Since $\mathcal{O}$-$hull(\{p,q\})$ is $\mathcal{O}_i$-convex, for all $i \in I$, we have that $p \in \mathcal{O}_i$-$kernel(\mathbf{P})$ and this implies that p is in $\bigcap_{i\in I} \mathcal{O}_i$-$kernel(\mathbf{P})$.

Now let $p \in \bigcap_{i\in I} \mathcal{O}_i$-$kernel(\mathbf{P})$ and $q \in \mathbf{P}$ arbitrary. Further, let ϑ be the orientation of the line defined by p and q. We distinguish two cases.

[4] I denotes an arbitrary index set throughout this paper.

First we assume that ϑ is in $\mathcal{O}$. Since $\mathcal{O} = \bigcup_{i\in I} \mathcal{O}_i$, there is an $i_0 \in I$ with $\vartheta \in \mathcal{O}_{i_0}$. Hence, $\mathcal{O}$-$hull(\{p,q\}) = \overline{pq} = \mathcal{O}_{i_0}$-$hull(\{p,q\}) \subseteq \mathbf{P}$, because p $\mathcal{O}_{i_0}$-sees q. But then p $\mathcal{O}$-sees q. On the other hand, if $\vartheta \notin \mathcal{O}$, then $\mathcal{O}$-$hull(\{p,q\}) = \{p,q\} \subseteq \mathbf{P}$, and p $\mathcal{O}$-sees q. □

As we already observed the above lemma enables us to compute the $\mathcal{O}$-kernel, for finite $\mathcal{O}$, once we are able to compute the $\{\theta\}$-kernel of a polygon. In the following we present a linear time algorithm to compute $\{0°\}$-$kernel(\mathbf{P})$[5], for a *simple* polygon $\mathbf{P}$.

4.1 The $\{0°\}$-Kernel of a Simple Polygon

In order to develop an algorithm that computes the $\{0°\}$-kernel of a polygon, we need to gain some insight in its structure. As a first step the next lemma describes the $\{0°\}$-kernel of a polygon $\mathbf{P}$ as the area of those points that intersect the boundary of $\mathbf{P}$ only once if they are moved infinitely far to the left or right.

Lemma 4.2 *Let* $\mathbf{P}$ *be a simple polygon. A point* p *in* $\mathbf{P}$ *belongs to* $\{0°\}$-$kernel(\mathbf{P})$ *if and only if the intersection of the horizontal line* h_p *through* p *with* $\mathbf{P}$ *is connected. Furthermore, if* $p \in \{0°\}$-$kernel(\mathbf{P})$*, then any point on* $h_p \cap \mathbf{P}$ *belongs to* $\{0°\}$-$kernel(\mathbf{P})$.

Proof: Let p be a point in $\mathbf{P}$, h_p the horizontal line through p, and $q \in \mathbf{P} \cap h_p$. If p belongs to $\{0°\}$-$kernel(\mathbf{P})$, then we have that $\{0°\}$-$hull(\{p,q\}) = \overline{pq} \subseteq \mathbf{P}$. This implies that $h_p \cap \mathbf{P}$ is connected since q is chosen arbitrarily. Now let $h_p \cap \mathbf{P}$ be connected. Hence, p sees all the points in $\mathbf{P}$ that have the same y-coordinate as p. Since $\{0°\}$-$hull(\{p,q\}) = \{p,q\} \subseteq \mathbf{P}$, for $q \notin h_p$, we have that p $\{0°\}$-sees all points in $\mathbf{P}$ and, hence, $p \in \{0°\}$-$kernel(\mathbf{P})$.

To see the second claim, just note that we have $h_p \cap \mathbf{P} = h_q \cap \mathbf{P}$, for any point q in $h_p \cap \mathbf{P}$. Therefore, we have that if $p \in \{0°\}$-$kernel(\mathbf{P})$, then $h_q \cap \mathbf{P}$ is connected and by the first claim q also belongs to $\{0°\}$-$kernel(\mathbf{P})$. □

We give an example of the $\{0°\}$-kernel of a polygon in Figure 5. Clearly, the above statement implies that the intersection of $\{0°\}$-$kernel(\mathbf{P})$ with any horizontal line is connected. We state this as a corollary.

Corollary 4.3 *If* $\mathbf{P}$ *is a polygon, then* $\{0°\}$-$kernel(\mathbf{P})$ *is* $\{0°\}$*-convex.*

The corollary yields that the $\{0°\}$-*kernel* of a polygon $\mathbf{P}$ consists of $\{0°\}$-convex components whose y-ranges do not overlap. There is an interesting relationship between the maximal $\mathcal{O}$-convex subsets and the $\mathcal{O}$-kernel of a polygon $\mathbf{P}$ if $|\mathcal{O}| = 1$. The maximal $\mathcal{O}$-convex subsets of $\mathbf{P}$ are called $\mathcal{O}$-*skulls* and are denoted by $\mathcal{O}$-$skulls(\mathbf{P})$. For the (rather simple) proof of the following lemma we need the notion of the $\mathcal{O}$-*join* of the $\mathcal{O}$-hull of two points with a third point. In general the $\mathcal{O}$-*join* of a set $\mathbf{C}$ and a point p is defined to be $\bigcup_{q\in\mathbf{C}} \mathcal{O}$-$hull(\{p,q\})$ and denoted by $\mathcal{O}$-$join(p,\mathbf{C})$. In Section 6.1 we will make use of a different join concept to prove a similar result in the case of $\mathcal{O}^s$-visibility.

[5] Remember that the $\{0°\}$-$kernel(\mathbf{P})$ actually is the $\{0°, 180°\}$-kernel of $\mathbf{P}$.

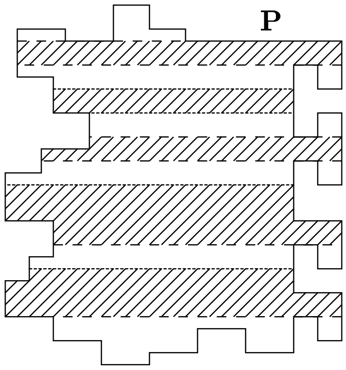

Figure 5: *An example of the* $\{0°\}$*-kernel of a polygon.*

Theorem 4.4 *If* $\mathbf{P}$ *is a polygon, then we have*

$$\{0°\}\text{-}kernel(\mathbf{P}) = \bigcap \{0°\}\text{-}skulls(\mathbf{P}).$$

Proof: Since the semi-convex space of $\mathcal{O}$-convex sets is simple and complete, we only have to prove that $\{0°\}$-$join(p, \{0°\}\text{-}hull(\{q, r\}))$ is $\{0°\}$-convex, for all points p, q, $r \in I\!E^2$ according to the Kernel Theorem of [SW91]. First we consider the case that $\{0°\}$-$hull(\{q, r\}) = \overline{qr}$. If p lies above or below $\overline{qr}$, then the join just consists of p and $\overline{qr}$ which, obviously, is $\{0°\}$-convex. If, on the other hand, p has the same height as q and r, then the join consists of one line segment, which, of course, is again $\{0°\}$-convex.

Now suppose that $\{0°\}$-$hull(\{q, r\}) = \{q, r\}$. If p does not lie on a common horizontal line with either q or r, then we have the three points (with different heights) as the join and the join is $\{0°\}$-convex. Should p have the same height as one of the points q or r, then the join consists again of a line segment and a point and is, thus, $\{0°\}$-convex. $\square$

It can easily be seen that this nice property does not carry over to the case $|\mathcal{O}| \geq 2$. An examples is shown in Figure 6.

Since the connected components of the $\{0°\}$-kernel of a polygon $\mathbf{P}$ are $\{0°\}$-convex, they contain two monotone chains w.r.t. the y-axis and at most two extremal horizontal edges. Clearly, the monotone chains are part of the boundary of $\mathbf{P}$. So the main task in computing $\{0°\}$-$kernel(\mathbf{P})$ consists of finding the maximal and minimal coordinates of the components.

Before we continue, we need to introduce some more notation which we will also make use of later on in this paper. Let θ be some orientation in $\mathcal{O}$. We consider the coordinate system that has a θ-oriented x-axis and a $(\theta + 90°)$-oriented y-axis. For a polygon $\mathbf{P}$ and a point p on the boundary of $\mathbf{P}$, we say that p is a (local) *θ-maximum of* $\mathbf{P}$ if there exists a neighbourhood N of the connected component of $\partial\mathbf{P} \cap \ell(p, \theta)$ such that there is no point in $\partial\mathbf{P} \cap N$ that is higher than p in the θ-coordinate system. For illustration refer to Figure 7.

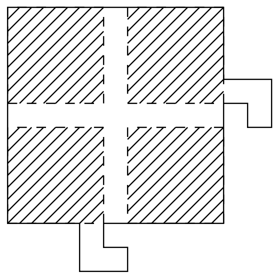

Figure 6: *The* $\{0°, 90°\}$*-kernel of a polygon is not* $\{0°, 90°\}$*-convex.*

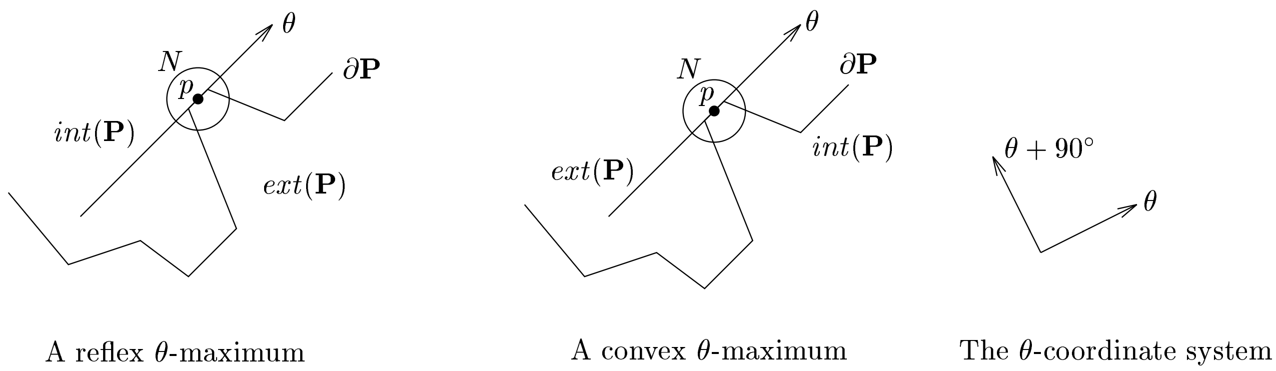

Figure 7: *The* θ*-maximum of a polygon* $\mathbf{P}$*.*

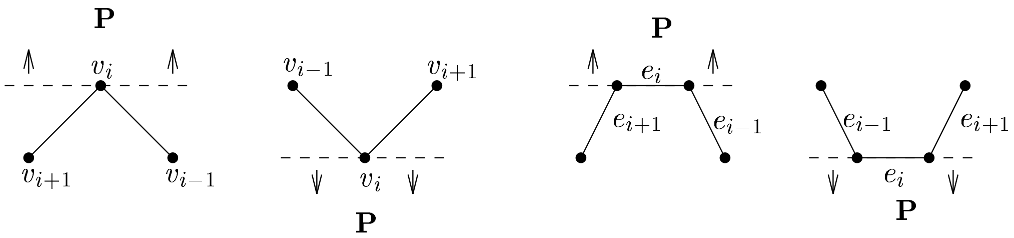

Figure 8: *Possible maxima and minima of a polygon.*

p is called a (local) *reflex θ-maximum* if p is a θ-maximum and there exists a neighbourhood N of p such that $ext(\mathbf{P}) \cap N$ is below the θ-oriented line through p in the θ-coordinate system, otherwise p is called a *convex θ-maximum*. A θ-maximum p is said to be *global* if it is the highest local θ-maximum of $\mathbf{P}$. The definition of θ-minima is analogous to the definition of θ-maxima. If $\theta = 0°$, we just speak of maxima and minima. Note that a maximum or minimum may consist of a vertex or an edge of $\mathbf{P}$ (see Figure 8). An (*oriented*) *tangent* t of a set $\mathbf{S}$ is an oriented line such that t intersects the boundary of $\mathbf{S}$ and such that $\mathbf{S}$ is completely contained in $h^+(t)$. A *horizontal* tangent is a tangent t such that $\Theta(t)$ is either equal to $0°$ or $180°$. With the above terminology, we have the following result about the $\{0°\}$-kernel of a polygon $\mathbf{P}$.

Lemma 4.5 *Let* $\mathbf{P}$ *be a polygon and let* $\mathbf{C}$ *be a component of* $\{0°\}$*-kernel*$(\mathbf{P})$*, then*

1. *Either* $\mathbf{C}$ *contains the highest vertex of* $\mathbf{P}$ *or the topmost horizontal tangent* t *to* $\mathbf{C}$ *contains a local minimum of* $\mathbf{P}$. t *intersects* $\mathbf{C}$ *if and only if* t *does not touch a convex minimum.*

2. *Either* $\mathbf{C}$ *contains the lowest vertex of* $\mathbf{P}$ *or the bottommost horizontal tangent* t *to* $\mathbf{C}$ *contains a local maximum.* t *intersects* $\mathbf{C}$ *if and only if* t *does not touch a convex maximum.*

Proof: We only prove the first claim since the proof of the second claim is completely symmetric to the first one. Suppose that $\mathbf{C}$ does not contain one of the highest vertices of $\mathbf{P}$.

Case 1 There is at least one point $p \in \mathbf{C}$ with highest y-coordinate.
Clearly, p is intersected by the topmost tangent to $\mathbf{C}$. Let h be the horizontal line through p. By Lemma 4.2 we have $h \cap \mathbf{P} \subseteq \mathbf{C}$. So there is a highest non-degenerate edge e in $\mathbf{C}$. Furthermore, there is some horizontal line in every neighbourhood above e that intersects $\mathbf{P}$ in at least two disconnected components. Hence, we obtain a sequence of horizontal lines h_n that converge to h and that intersect $\mathbf{P}$ in at least two components. So, there are at least four non-horizontal edges e_i in $\mathbf{P}$, $1 \leq i \leq 4$, that are intersected consecutively in that order from left to right by almost all h_n (see Figure 9a). We can assume that e_1

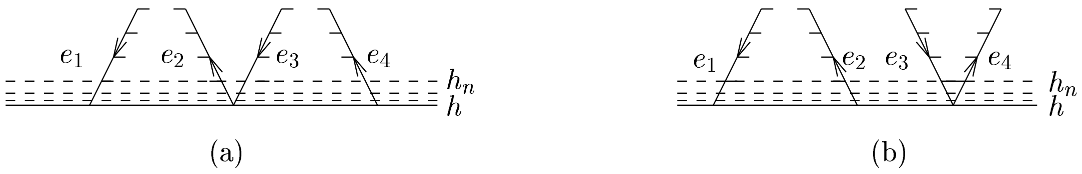

Figure 9: *The two cases that can occur for the highest segment of a* $\{0°\}$*-kernel component.*

and e_3 are directed downward and e_2 and e_4 are directed upward if we orient the boundary of $\mathbf{P}$ counterclockwise. Since the edges e_i are closed, we have that h also intersects them again in the same order from left to right. Therefore, e_2 and e_3 meet in one vertex or edge that is contained in h. Since e_2 is a upward directed edge and e_3 is a downward directed edge, they form a reflex minimum of $\mathbf{P}$.

Case 2 There is no point with highest y-coordinate in $\mathbf{C}$.
By the same reasoning as above we obtain a sequence of horizontal lines that converge and intersect $\mathbf{P}$ in two disconnected components (see Figure 9b); only now the limit h also intersects $\mathbf{P}$ in two disconnected components. Since every horizontal line in a sufficiently small neighbourhood below h has a connected intersection with $\mathbf{P}$, there is at least one vertex or edge on h before or after the one connected component of $h \cap \mathbf{P}$ that contains $\partial\mathbf{C}$; therefore, we get a local minimum with a convex vertex or edge that is touched by h.

□

4.2 The $\{0°\}$-Kernel Algorithm

In this section we present a linear time algorithm to compute the $\{0°\}$-kernel of a simple polygon $\mathbf{P}$. Later we will give a modification of this algorithm that needs time $O(n \log n)$ to compute the $\{0°\}$-kernel of a multiply connected polygon.

The algorithm follows a very simple idea. Starting with the lowest vertex we scan the boundary of $\mathbf{P}$ counterclockwise. During this scan we maintain a list $\mathcal{L}$ of y-intervals that correspond either to possible components of $\{0°\}$-$kernel(\mathbf{P})$ or to ranges that have already been recognized as not being part of $\{0°\}$-$kernel(\mathbf{P})$, i.e., ranges in which we have already detected a horizontal line that intersects $\mathbf{P}$ in at least two line segments. With a y-interval that corresponds to a (possible) component $\mathbf{C}$ of $\{0°\}$-$kernel(\mathbf{P})$ we also store those edges of $\mathbf{P}$ that belong to the boundary of $\mathbf{C}$ and that we have discovered so far. For each edge e we encounter in the scan, we have to update the list $\mathcal{L}$ of y-intervals. Furthermore, we may have to add e to some part of the boundary of the currently considered component. In the scan we distinguish three kinds of y-intervals.

1. The intervals that do not intersect the y-range of any of the so far processed edges. We call these *free intervals*. In the beginning, the whole y-axis is one free interval.

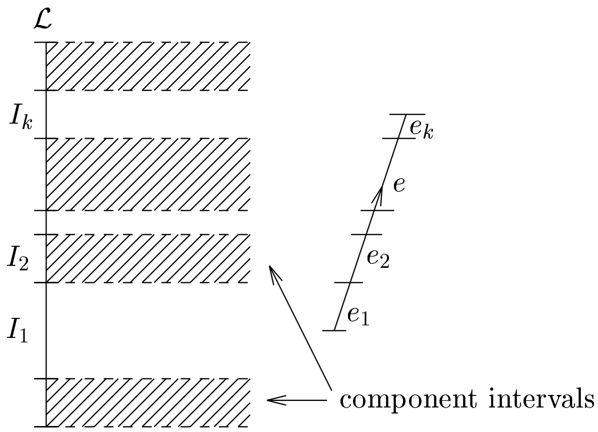

Figure 10: *Splitting up edge e in subedges* $e_1, \ldots, e_k$.

After having looked at at least one edge there are exactly two free intervals, one above the highest and one below the lowest so far encountered edge.

2. The intervals which may contain the y-range of a component of $\{0°\}$-$kernel(\mathbf{P})$. We call these *component intervals*. A component interval I has neither been completely covered more than twice by the edges of $\mathbf{P}$ we have looked at so far nor have we encountered an upward edge to the left of a downward edge in the horizontal slab spanned by I. Both of these conditions would exclude I from the intervals that may contain a component of $\{0°\}$-$kernel(\mathbf{P})$.

3. Finally, there are the intervals that have been crossed at least three times or that contain an upward edge to the left of a downward edge. These are called *deleted intervals* since such an interval can never contain a component of $\{0°\}$-$kernel(\mathbf{P})$.

Note that the intervals in $\mathcal{L}$ are alternatingly component and deleted intervals except for the first and last interval which are free intervals since any sequence of adjacent intervals of the same kind can, of course, be joined into one interval.

So suppose that we have processed some part of the boundary of $\mathbf{P}$ and now encounter edge e. e can either be directed upward or downward. Since these cases are symmetric, we will assume that e is an upward edge, i.e., the interior of $\mathbf{P}$ is locally to the left of e. The y-range of e may span several intervals $I_1, \ldots, I_k$ in $\mathcal{L}$. Therefore, we split e up into a number of subedges $e_1, \ldots, e_k$ such that the y-range of e_i is contained in I_i (see Figure 10). Now consider some subedge e_i. Let v_s be the start vertex of e_i and v_t the end vertex. We distinguish several cases depending on which kind of interval v_s belongs to. If v_s is on the boundary of two intervals, we just consider the upper one since e is directed upward. We denote the interval v_s belongs to by I_s and the interval below I_s by I_b.

Case 1 I_s is a free interval.
If I_b is a component interval, then we just extend I_b to the y-coordinate of v_t and add e_i to the right boundary of the component of $\{0°\}$-$kernel(\mathbf{P})$ that corresponds to I_b. If I_b is a deleted interval, then we create a new component interval I'_s between I_b and I_s that spans the y-range of e_i. Again the part of e_i that is above I_b is added to the right boundary of the component that corresponds to I'_s. In both cases the lower end point of I_s is moved up to the y-coordinate of v_t. In the very beginning we have that v_s belongs to the only interval of the list which is a free interval. In this case we create a new component interval I'_s that spans the y-range of e_i and split the free interval I_s up into two free intervals, I_1 and I_2, with I_1 below and I_2 above I'_s.

Case 2 I_s is a component interval.
Let $\mathbf{C}$ be the (possible) component of $\{0°\}$-$kernel(\mathbf{P})$ that corresponds to I_s.

Case 2.1 There is a downward edge to the right of e_i in $\mathbf{C}$ or there is another upward edge at the same height as e_i in $\mathbf{C}$.
In both of theses cases a horizontal line that has the same height as e_i intersects $\mathbf{P}$ in at least two disconnected components. We then delete the part I'_s of I_s from the y-coordinate of v_s to the y-coordinate of v_t. This divides I_s into at most three parts, the part I_1 that is below I'_s, I'_s itself, and the part I_2 that is above I'_s. If one of I_1 or I_2 is empty and there is a deleted interval below or above I_s, we can combine it with I'_s. Note that both of I_1 and I_2 are non-empty if and only if the y-range of edge e is completely contained in I_s. This will be important in the analysis of the algorithm.

Case 2.2 There is no downward edge to the right of e_i in $\mathbf{C}$ nor an upward edge at the same height.
The right boundary of $\mathbf{C}$ has not been traversed yet and the insertion of e_i does not lead to a disconnected intersection with a horizontal line. Hence, we add e_i to the right boundary of $\mathbf{C}$ and leave $\mathcal{L}$ unchanged.

Case 3 I_s is a deleted interval.
In this case there is nothing to do since I_s is already deleted and e_i cannot belong to any component of $\{0°\}$-$kernel(\mathbf{P})$.

The proof of correctness of the above algorithm is given by the following two lemmas.

Lemma 4.6 *If* $\mathbf{P}$ *is a polygon in the plane and* $p \in \{0°\}$-$kernel(\mathbf{P})$*, then there is a component interval* I *in the output of the* $\{0°\}$*-kernel algorithm for* $\mathbf{P}$ *that contains the* y*-coordinate of* p*. Furthermore, the component corresponding to* I *contains* p*.*

Proof: Let $p \in \{0°\}$-$kernel(\mathbf{P})$. Consider the horizontal line h through p and assume it intersects the edge $e_r \in \mathbf{P}$ on the right hand side. Clearly, this edge is directed upward. Since p belongs to $\{0°\}$-$kernel(\mathbf{P})$, h does not intersect any other edge of $\mathbf{P}$ to the right of it nor an upward edge to the left of it; hence, e_r was added either in Case 1 or 2.2. The same argument applies to the edge e_l that is intersected by h on the left hand side.

Therefore, the y-coordinate of p belongs to one of the component intervals that are created during the execution of the algorithm and the y-coordinate of p cannot be deleted from it since then there would be a third edge that intersects h. In the same way we obtain that the edges e_r and e_l form the right and left boundary of the component corresponding to the interval in $\mathcal{L}$ p belongs to. $\square$

Lemma 4.7 *If* $\mathbf{P}$ *is a polygon in the plane and* $p \notin \{0°\}$*-kernel*$(\mathbf{P})$*, then there is a deleted or free interval in the output of the* $\{0°\}$*-kernel algorithm for* $\mathbf{P}$ *that contains the* y*-coordinate of* p*.*

Proof: If p is above or below $\mathbf{P}$, the claim is obvious. Hence, suppose that p belongs to the horizontal slab spanned by $\mathbf{P}$. Again consider the horizontal line h through p. p is not contained in $\{0°\}$-*kernel*$(\mathbf{P})$, therefore, h intersects at least four edges $e_1, \ldots, e_m$ with $m \geq 4$. Let the numbering denote the sequence in which the algorithm encounters the edges. Assume that e_1 is an upward edge. By the topological properties of polygons e_2 is then directed downward. If e_2 is to the right of e_1, then we delete the y-coordinate of p from the component interval by Case 2.1. If e_2 is to the left of e_1, then the y-coordinate of p is deleted from the component interval when we encounter the upward edge e_3. $\square$

Analysis

So it just remains to be shown that the algorithm, indeed, needs linear time. Obviously, each of the cases we considered above needs only a constant amount of time. Hence, processing a subedge e_i of e can be done in a constant number of steps. So, if e is not split up, then we are within our time bound. On the other hand, if e is split up into a number of subedges $e_1, \ldots, e_k$, then note that the y-range of e includes at least $\lfloor (k-2)/2 \rfloor$ component intervals since the intervals in $\mathcal{L}$ are alternatingly deleted and component intervals with the possible exception of the first and last interval. If we denote the number of subedges edge e is split up into by k_e and the number of component intervals the y-range of e contains by c_e, we obtain the following estimate for the total number of subedges created during the algorithm

$$\sum_{e \in \mathbf{P}} k_e \leq \sum_{e \in \mathbf{P}} (2(c_e + 1) + 2).$$

Next we show that the sum on the right hand side is linear in the number of vertices of $\mathbf{P}$. Our first claim is that we create at most n component intervals during the scan if n is the size of $\mathbf{P}$. To see this we note that any edge leads to the creation of at most one component interval. For, consider the Cases 1 and 2.1 which are the only ones that may lead to a new component interval. In both cases we have that the end point of the edge we consider is contained in the interval I_s if the number of component intervals increases.

Now by Case 2.1 a component interval I is contained in at most two y-ranges that belong to some edges of $\mathbf{P}$ before it is combined with a deleted interval. Hence, if we sum the number of edges of $\mathbf{P}$ whose y-range contain I over all component intervals I that are created in the scan, we cannot obtain more than twice the number of intervals. In other words, if c_I is the number of edges that cover component interval I we have that $\sum_I c_I \leq 2n$. Clearly, we have that $\sum_{e \in \mathbf{P}} c_e = \sum_I c_I$ since both sums count how often the y-range of an edge covers a component interval. Hence, we get the following estimate for the total number of subedges created during the execution of the algorithm

$$\begin{aligned}
\sum_{e \in \mathbf{P}} k_e &\leq \sum_{e \in \mathbf{P}} (2(c_e + 1) + 2) \\
&\leq 2 \sum_I c_I + 4n \\
&= O(n).
\end{aligned}$$

This completes the analysis of our algorithm. We have shown the following theorem.

Theorem 4.8 *If* $\mathbf{P}$ *is a simple polygon in the plane, then* $\{0°\}$-*kernel*$(\mathbf{P})$ *can be computed in time linear in the number of vertices of* $\mathbf{P}$.

4.3 The $\mathcal{O}$-kernel for polygons with holes

If we allow the polygons to be multiply connected (but they still have to be simple), we have to account for the regions of the $\mathcal{O}$-kernel that are deleted by the presence of holes. Fortunately, the Intersection Lemma holds for any subset of the plane. So, again, we only have to show how to compute the $\{0°\}$-kernel, for a given polygon $\mathbf{P}$. The lemma stated below completely describes the relationship between the $\{0°\}$-kernel of $\mathbf{P}$ and $\{0°\}$-kernel of the enclosing polygon and the holes.

Lemma 4.9 *If* $\mathbf{P}$ *is a simple multiply connected polygon with enclosing polygon* $\mathbf{Q}$ *and holes* $\mathbf{H}_1, \ldots, \mathbf{H}_m$*, then we have that* $p \in \{0°\}$-*kernel*$(\mathbf{P})$ *if and only if*

1. $p \in \{0°\}$-*kernel*$(\mathbf{Q})$ *and*

2. *there is no hole that intersects the horizontal line* h_p *through* p.

Proof: Obviously, we have $\{0°\}$-*kernel*$(\mathbf{P}) \subseteq \{0°\}$-*kernel*$(\mathbf{Q})$. So let $p \in \{0°\}$-*kernel*$(\mathbf{Q})$ and assume there is no $\mathbf{H}_i$ that intersects the horizontal line h_p through p and let $q \in \mathbf{P}$. If $\Theta(\overline{pq}) = 0°$, then $q \in h_p$ and, hence, $\{0°\}$-*hull*$(\{p,q\}) = \overline{pq} \subseteq \mathbf{P}$. If $\Theta(\overline{pq}) \neq 0°$, then $\{0°\}$-*hull*$(\{p,q\}) = \{p,q\} \subseteq \mathbf{P}$.

On the other hand, if $p \in \{0°\}$-*kernel*$(\mathbf{P})$, then h_p does not intersect any of the $\mathbf{H}_i$ since otherwise p cannot see all points in $\mathbf{P} \cap h_p$. □

The above lemma leads to the following algorithm to compute the $\{0°\}$-kernel of a multiply connected polygon consisting of the enclosing polygon $\mathbf{Q}$ and holes $\mathbf{H}_1, \ldots, \mathbf{H}_m$ (the algorithm is very similar to the algorithm described in [Raw87, p. 82]). We first compute $\{0°\}$-*kernel*$(\mathbf{Q})$. The y-coordinates of the components of $\{0°\}$-*kernel*$(\mathbf{Q})$ are stored in a range tree T. We use T to answer stabbling queries with the end points of the y-intervals that are covered by the holes $\mathbf{H}_i$. Any component interval that lies between

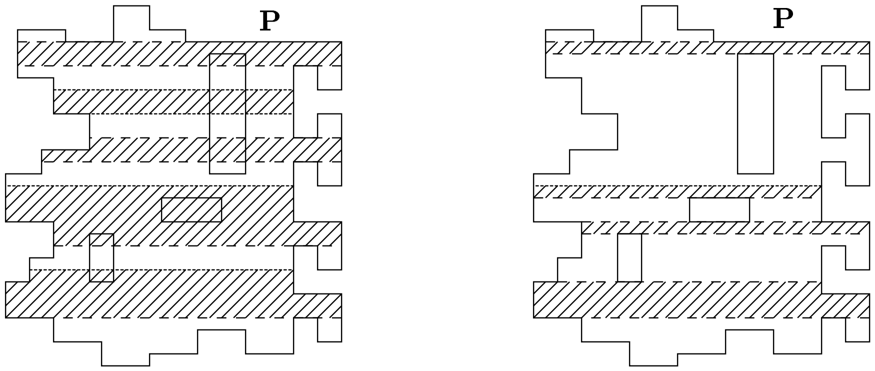

The $\{0°\}$-kernel of **P** without holes   The $\{0°\}$-kernel of **P** with holes

Figure 11: *An example of the $\mathcal{O}$-kernel of a polygon with holes.*

two such end points, is deleted since its y-range intersects a hole. Since the y-ranges of the components are disjoint, we can answer a stabbling query in time $O(\log(\#$ of components of $\{0°\}\text{-}kernel(\mathbf{Q}))) = O(\log n)$. The algorithm can now be described as follows (an example is shown in Figure 11).

Algorithm $\{0°\}$-kernel for polygons with holes
compute $\{0°\}\text{-}kernel(\mathbf{Q})$;
for each $\mathbf{H}_i$ do
    compute the maximal y_i^{max} and minimal y_i^{min} y-coordinate of $\mathbf{H}_i$;
    find the components of $\{0°\}\text{-}kernel(\mathbf{Q})$ that have a y-range in $[y_i^{min}, y_i^{max}]$.
    let $\mathbf{Z}_{min}$ be the component such that y_i^{min} is in the y-range of $\mathbf{Z}_{min}$;
    define $\mathbf{Z}_{max}$ analogously;
    delete the part of $\mathbf{Z}_{min}$ in the slab that is spanned by $[y_i^{min}, y_i^{max}]$;
    delete the part of $\mathbf{Z}_{max}$ in the slab that is spanned by $[y_i^{min}, y_i^{max}]$;
    delete all components between $\mathbf{Z}_{min}$ and $\mathbf{Z}_{max}$;
end for;

The maximum number of components we can get this way is the sum of the number of components and the number of holes since every hole can split at most one component into two. Finding $\mathbf{Z}_{min}$ and $\mathbf{Z}_{max}$ can thus always be done in time $O(\log n)$ if we use the above mentioned range tree. Deletion of a part of a component can be accomplished in time $O(\log n)$ plus $O($the number of edges deleted$)$ and the deletion of all components between $\mathbf{Z}_{min}$ and $\mathbf{Z}_{max}$ can also be achieved in time proportional to the number of components deleted. Since we can get at most a linear number of edges and components and since each edge and component can be deleted only once, the overall deletion time is linear. Hence,

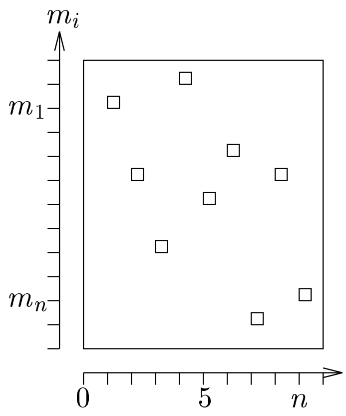

Figure 12: *The construction of a polygon for n integers.*

the running time of the $\{0°\}$-kernel algorithm for an n-vertex multiply connected simple polygon is $O(n \log n)$. This is indeed optimal, as the following lemma shows.[6]

Lemma 4.10 *Finding the $\{0°\}$-kernel of an n-vertex polygon with holes requires $\Omega(n \log n)$ time in the algebraic computation tree model.*

Proof: Suppose there is an algorithm that computes the $\{0°\}$-kernel of a polygon with holes in time $o(n \log n)$. We give a linear reduction to the integer element uniqueness problem. Recall that we introduced this problem in Section 2.5 and that it requires time $\Omega(n \log n)$ in the algebraic computation tree model.

Given n integers $m_1, m_2, \ldots, m_n \in I\!N$ we construct a polygon $\mathbf{P}$ with $4(n+1)$ vertices containing n holes as follows (for illustration refer to Figure 12). The i^{th} hole is the rectangle spanned by the points $(i, m_i + 1/2)$ and $(i + 1/2, m_i)$. We find the minimum and maximum values of the m_i and construct an appropriate surrounding box. All of these operations take linear time. After applying the $\{0°\}$-kernel algorithm to $\mathbf{P}$ we can count, in linear time, the number of components. If it is equal to $n+1$, all n_i are distinct, otherwise there is some pair of equal integers. □

5 $\mathcal{O}^s$-Convexity in Simple Polygons

In the rest of this paper we will look at the second definition of visibility we proposed in the context of $\mathcal{O}$-convexity. Recall that in Section 2.6.1 we defined that two points p and

[6]The proof idea for this lemma is taken from the proof of Lemma 6.4.4 of [Raw87] which makes a similar statement about the time needed to compute the $\{0°\}^s$-kernel of a polygon with holes. Though the construction used is exactly the same as presented here the proof in [Raw87] is not correct since a reduction to the element uniqueness problem over the *reals* is claimed which only works if *point holes* are considered. A lower bound for the Integer Element Uniqueness Problem was discovered only later.

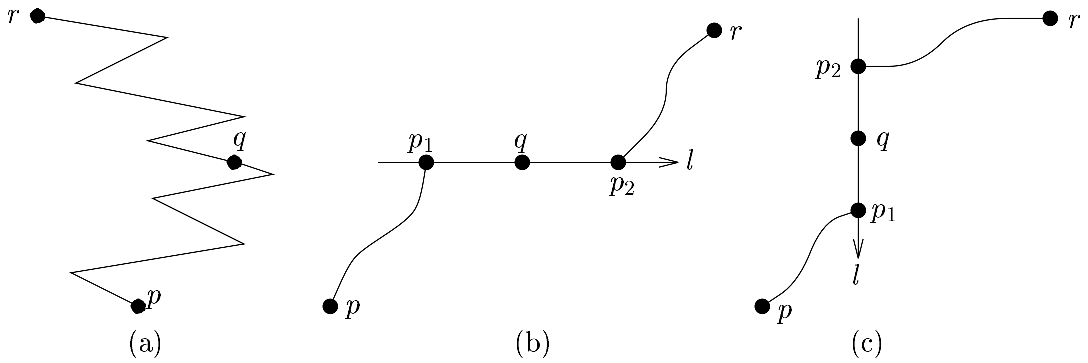

Figure 13: *How to join two staircases*

q in a polygon **P** $\mathcal{O}^s$-see each other if there exists an $\mathcal{O}$-stairsegment in **P** that connects p and q. In this section we will investigate $\mathcal{O}^s$-visibility in simple polygons. Our main interest will be to obtain a result similar to the intersection lemma (Lemma 4.1) which drastically simplifies the problem of computing the $\mathcal{O}^s$-kernel.

Since $\mathcal{O}^s$-visibility is based on $\mathcal{O}$-stairsegments, the following section is concerned with some basic properties of $\mathcal{O}$-stairsegments we will make use of throughout the rest of this section.

5.1 $\mathcal{O}$-Stairsegments

We need some technical statements on how to construct new $\mathcal{O}$-stairsegments from some given $\mathcal{O}$-convex curves. The first lemma shows how to glue two $\mathcal{O}$-stairsegments together.

Lemma 5.1 *Let p, q, and r be points in the plane and (α, β) an $\mathcal{O}$-free range with $\Theta(\overline{pq}) \in [\alpha, \beta]$ and $\Theta(\overline{qr}) \in [\alpha, \beta]$. If $\mathcal{S}_{pq}$ is an $\mathcal{O}$-stairsegment from p to q and $\mathcal{S}_{qr}$ is an $\mathcal{O}$-stairsegment from q to r, then $\mathcal{S}_{pq} * \mathcal{S}_{qr}$ is an $\mathcal{O}$-stairsegment from p to r.*

Proof: We use the same notation as in the claim. We distinguish two cases according to the cardinality of $\mathcal{O}$.

Case 1 $|\mathcal{O}| = 1$.
W.l.o.g. suppose that $\mathcal{O} = \{0°\}$. Hence, q is above p and r is above q (see Figure 13a).[7] So consider a horizontal line l that intersects $\mathcal{S}_{pq} * \mathcal{S}_{qr}$. If l is strictly below q, then l only intersects $\mathcal{S}_{pq}$ since $\mathcal{S}_{qr}$ is $\{0°\}$-convex and, therefore, stays above q. Hence, $l \cap (\mathcal{S}_{pq} * \mathcal{S}_{qr})$ is connected in this case. Similarly, if l is strictly above q, it only intersects $\mathcal{S}_{qr}$ and

[7]Note that we will use the terms "above" and "below" meaning *at least as high as* resp. *at least as low as*. If we want to indicate that one object is higher (lower) than some other object will use the terms *strictly above* (*strictly below*). The same remark applies to the terms "to the right of" and "to the left of".

$l \cap (\mathcal{S}_{pq} * \mathcal{S}_{qr})$ is again connected. If l has the same height as q, then $l \cap \mathcal{S}_{pq}$ consists of a horizontal line segment s_1 that contains q and $l \cap \mathcal{S}_{qr}$ also consists of a horizontal line segment s_2 that contains q. Thus, $l \cap (\mathcal{S}_{pq} * \mathcal{S}_{qr}) = s_1 \cup s_2$ is also a line segment.

Case 2 $|\mathcal{O}| \geq 2$.
We assume w.l.o.g. that the $\mathcal{O}$-free range (α, β) equals $(0°, 90°)$. Consider any $\mathcal{O}$-line l. Suppose that l intersects $\mathcal{S}_{pq} * \mathcal{S}_{qr}$ in at least two points p_1 and p_2. We will show that p_1 and p_2 are in a connected component of $l \cap (\mathcal{S}_{pq} * \mathcal{S}_{qr})$.

If p_1 and p_2 are both contained in the same $\mathcal{O}$-stairsegment, say $\mathcal{S}_{pq}$, then they belong to the same connected component of $l \cap (\mathcal{S}_{pq} * \mathcal{S}_{qr})$ since $\mathcal{S}_{pq}$ is $\mathcal{O}$-convex. Hence, assume p_1 and p_2 do not both belong to the same staircase. Let $p_1 \in \mathcal{S}_{pq}$ and $p_2 \in \mathcal{S}_{qr}$. Since $\Theta(l) \in \mathcal{O}$ and $(0°, 90°)$ is $\mathcal{O}$-free, we have that $\Theta(l) \in [90°, 180°] \cup [270°, 360°]$. W.l.o.g. we assume that $\Theta(l) \in [270°, 360°]$, i.e., l is either vertical or has a non-positive inclination. We distinguish two cases.

Case 2.1 p_1 is the first intersection point on l.
If l is not horizontal, then p_2 is strictly below p_1 and $\mathcal{S}_{qr}$ cannot contain p_2 since q is above p_1 and all points of $\mathcal{S}_{qr}$ are above q. Therefore, l is horizontal (see Figure 13b). Since q cannot be strictly above p_2 or strictly below p_1, l contains q. Since $\mathcal{S}_{pq}$ is $\mathcal{O}$-convex and p_1, $q \in \mathcal{S}_{pq}$ with $\Theta(\overline{p_1 q}) = 0°$, $\mathcal{S}_{pq}$ contains $\overline{p_1 q}$. The same holds for $\mathcal{S}_{qr}$ and $\overline{q p_2}$. Hence, $l \cap (\mathcal{S}_{pq} * \mathcal{S}_{qr}) \supseteq \overline{p_1 q} \cup \overline{q p_2}$; thus, p_1 and p_2 are connected in $l \cap (\mathcal{S}_{pq} * \mathcal{S}_{qr})$.

Case 2.2 p_2 is the first intersection point on l.
If l is not vertical, then p_2 is strictly to the left of p_1 and $\mathcal{S}_{qr}$ cannot contain p_2 since q is to the right of p_1 and $\mathcal{S}_{qr}$ is also to the right of q (see Figure 13c). Therefore, l is vertical and, as above, $\overline{p_1 q}$ and $\overline{q p_2}$ are contained in $\mathcal{S}_{pq} * \mathcal{S}_{qr}$; hence, p_1 and p_2 are connected in $l \cap (\mathcal{S}_{pq} * \mathcal{S}_{qr})$.

□

Now consider two points p and q and an $\mathcal{O}$-convex path $\mathcal{S}$ from p to q. If we choose any two points v, w on $\mathcal{S}$ with v before w, we claim that we can connect them by an $\mathcal{O}$-stairsegment $\mathcal{S}'$ and the part $\mathcal{S}_{pv}$ of $\mathcal{S}$ from p to v and the part $\mathcal{S}_{wq}$ of $\mathcal{S}$ from w to q together with $\mathcal{S}'$ form an $\mathcal{O}$-stairsegment, i.e., $\mathcal{S}_{pv} * \mathcal{S}' * \mathcal{S}_{wq}$ is $\mathcal{O}$-convex. To see this note that $\mathcal{S}_{pv}$ and $\mathcal{S}'$ satisfy the conditions of Lemma 5.1. Hence, $\mathcal{S}_{pv} * \mathcal{S}'$ is $\mathcal{O}$-convex. The same holds for $\mathcal{S}_{pv} * \mathcal{S}'$ and $\mathcal{S}_{wq}$. We state this as a corollary (see Figure 14).

Corollary 5.2 *Let p, v, w, and q be four points in the plane that occur in this order on the $\mathcal{O}$-stairsegment $\mathcal{S}$. If we replace the part of $\mathcal{S}$ between v and w by some other $\mathcal{O}$-stairsegment, the resulting path $\mathcal{S}'$ is $\mathcal{O}$-convex.*

Another corollary tells us how to construct $\mathcal{O}$-staircases from edges.

Lemma 5.3 *Let $\mathcal{S}$ be a polygonal curve consisting of the edges e_1, $e_2, \ldots, e_n$ in this sequence. $\mathcal{S}$ is an $\mathcal{O}$-stairsegment if and only there exists an $\mathcal{O}$-free range (α, β) such that $\Theta(e_i) \in [\alpha, \beta]$, for $i \in \{1, \ldots, n\}$.*

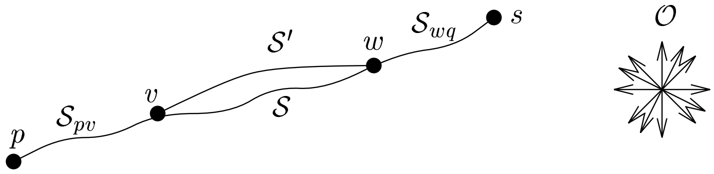

Figure 14: *The part of $\mathcal{S}$ from q to r can be replaced by any $\mathcal{O}$-stairsegment $\mathcal{S}'$.*

Proof: The "if"-part of the claim is an immediate consequence of Lemma 4.2.1 of [Raw87]. So we just have to show the "only if"-part. We prove it by induction on n. The induction hypothesis clearly holds for $n = 1$. Hence, assume it holds for all $k < n$ and n is greater than 1. Let $\mathcal{S}$ be a polygonal curve of n line segments e_i, $i \in \{1, \ldots, n\}$ such that $\Theta(e_i) \in [\alpha, \beta]$.

Let $\mathcal{S}'$ be the staircase consisting of the edges $e_1, \ldots, e_{n-1}$. $\mathcal{S}'$ is $\mathcal{O}$-convex by induction hypothesis. If p is the start point of $\mathcal{S}'$, q the start point of e_n, and r the end point of e_n, then we have the same situation as in Lemma 5.1. Hence, $\mathcal{S} = \mathcal{S}' * e_n$ is $\mathcal{O}$-convex. □

5.2 The $\mathcal{O}^s$-Hull of Two Points

In our further investigations of $\mathcal{O}^s$-visibility we want to make use of the results obtained in [SW91].[8] In analogy to $\mathcal{O}$-convexity one might want to introduce the following definition.

Definition 5.1 *Let $\mathcal{O}$ be some set of orientations. We say a set $\mathbf{C}$ is $\mathcal{O}^s$-*convex *if any two points in $\mathbf{C}$ $\mathcal{O}^s$-see each other.*

The problem we encounter with this definition is that the intersection of two $\mathcal{O}^s$-convex sets is not necessarily $\mathcal{O}^s$-convex again. Hence, the family of $\mathcal{O}^s$-convex sets does not form a semi-convex space[9]. There are two ways to remedy this situation. The first is to require that *all* $\mathcal{O}$-stairsegments between two points are contained in the set instead of just one. This leads to the notion of *strong-$\mathcal{O}$-convexity* which we will not pursue further. The second one makes use of a convexity concept based on "straight" lines w.r.t. some d. It is not very surprising that we actually can use a norm d such that $\mathcal{O}$-stairsegments become d-straight lines.

Definition 5.2 *Let p be a point in $\mathbb{E}^2$. We denote the $\mathcal{O}$-*norm *of p by $\|p\|_{\mathcal{O}}$ and define it by*

$$\|p\|_{\mathcal{O}} = \begin{cases} \|p\|_2 & \text{if } \Theta(p) \in \mathcal{O} \\ \min\{\|p_1\|_2 + \|p_2\|_2 \mid \Theta(p_i) \in \mathcal{O} \text{ and } \vec{p_1} + \vec{p_2} = \vec{p}\} & \text{if } \Theta(p) \notin \mathcal{O}. \end{cases}$$

Figure 15 illustrates the definition. Let $\overline{pq}$ be a line segment with $\Theta(\overline{pq}) \notin \mathcal{O}$ and (α, β)

[8] This section is heavily based on the terminology used in [SW91]. We assume the reader to be familiar with Sections 3.3, 3.4, 5.3, and the Kernel Theorem of [SW91].

[9] A semi-convex space over a groundset $\mathbf{X}$ is a family $\mathcal{C}_{\mathbf{X}}$ of subsets of $\mathbf{X}$ that is closed under intersection (see [SW91]).

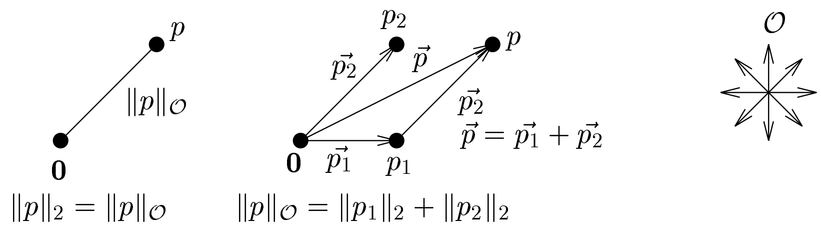

Figure 15: *The definition of* $\|\cdot\|_{\mathcal{O}}$.

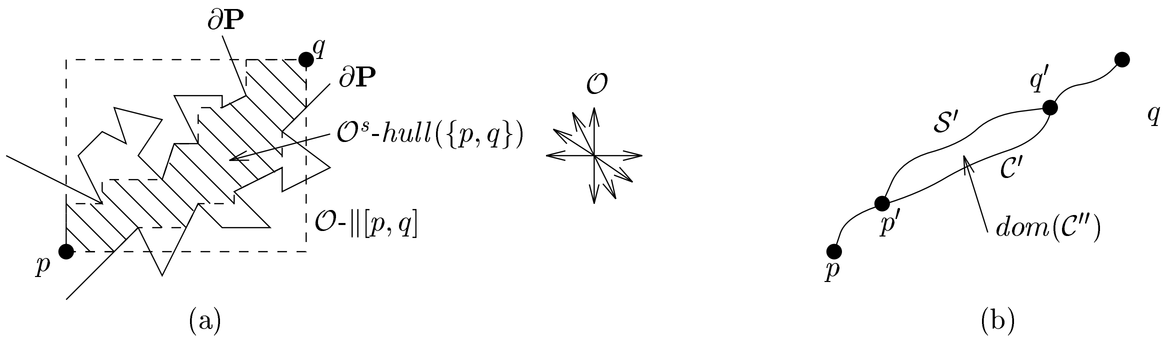

Figure 16: *The* $\mathcal{O}^s$*-hull of two points.* $dom(\mathcal{C}'')$ *can be added to* $\mathbf{C}$.

the maximal $\mathcal{O}$-free range of $\overline{pq}$. The distance from p to q is just half the perimeter of the parallelogram spanned by p and q with sides parallel to α and β. As in the case of $\{0°, 90°\}^s$-convexity (see [SW91]) it can be shown that $\|\cdot\|_{\mathcal{O}}$ is topological, extensible, and stiff (for a formal definition of these concepts we also refer to [SW91]). For further reference we state this as a lemma.

Lemma 5.4 *$\|\cdot\|_{\mathcal{O}}$ is a topological, extensible, and stiff norm in the plane.*

Since the results obtained in [SW91] only apply to simple domains, we will restrict our discussion to simple polygons. Hence, from now on throughout the rest of this section we use the term *polygon* to denote a simple polygon.

Definition 5.3 *If* $\mathbf{P}$ *is a polygon in the plane and* p *and* q *are two points in* $\mathbf{P}$*, we define the* $\mathcal{O}^s$-hull *of* p *and* q *to be the union of all* $\mathcal{O}$*-stairsegments in* $\mathbf{P}$ *that connect* p *and* q *and denote it by* $\mathcal{O}^s$*-hull*$(\{p,q\})$.

An example is shown in Figure 16a. Note that the definition of the $\mathcal{O}^s$-hull of two points depends on the polygon we consider. If the $\mathcal{O}^s$-hull of two points p and q consists of *all* $\mathcal{O}$-stairsegments from p to q, we also use the notation $\mathcal{O}$-$\|[p,q]$ (read: the $\mathcal{O}$-parallelogram of p and q). In Theorem 4.3.1 of [Raw87] it is shown that if $|\mathcal{O}| \geq 2$, then $\mathcal{O}$-$\|[p,q]$ is either the line segment between p and q if $\Theta(\overline{pq}) \in \mathcal{O}$ or the parallelogram spanned by p and q whose sides are parallel to the orientations of the maximal $\mathcal{O}$-free range of $\Theta(\overline{pq})$ if $\Theta(\overline{pq}) \notin \mathcal{O}$. Note that the definition of $\mathcal{O}$-$\|[p,q]$ is independent of the polygon we consider. Hence, if $\mathcal{O}^s$-$hull(\{p,q\}) = \mathcal{O}$-$\|[p,q]$, then the polygon does not present an obstacle for the $\mathcal{O}^s$-visibility between p and q. The $\mathcal{O}^s$-hull of two points has the following nice property.

Lemma 5.5 *Let* $\mathbf{P}$ *be a polygon and* p *and* q *be two points in* $\mathbf{P}$. *If* $\mathbf{C}$ *is a maximal, connected* $\mathcal{O}$*-convex subset of* $\mathcal{O}$-$\|[p,q] \cap \mathbf{P}$ *that contains* p *and* q*, then we have that* $\mathbf{C} = \mathcal{O}^s$-$hull(\{p,q\})$.

Proof: "$\subseteq$": Let r be a point of $\mathbf{C}$. Since $\mathbf{C}$ contains p and q, we have by Theorem 5.5.2 of [Raw87] that p sees r via an $\mathcal{O}$-stairsegment $\mathcal{S}_p$ in $\mathbf{C}$. Of course, the same holds for q which sees r via an $\mathcal{O}$-stairsegment $\mathcal{S}_q$. Note that we have, for a point r in $\mathcal{O}$-$\|[p,q]$ that $\Theta(\overline{pr}) \in [\alpha,\beta]$ and $\Theta(\overline{rq}) \in [\alpha,\beta]$ if (α,β) is the maximal $\mathcal{O}$-free range of $\overline{pq}$. Therefore, we can directly apply Lemma 5.1 to $\mathcal{S}_p$ and $\mathcal{S}_q$ which yields that $\mathcal{S}_p * \mathcal{S}_q$ is an $\mathcal{O}$-stairsegment. Hence, $r \in \mathcal{O}^s$-$hull(\{p,q\})$.
"$\supseteq$": Let $\mathcal{S}$ be an $\mathcal{O}$-stairsegment from p to q. Assume $\mathcal{S}$ is not contained in $\mathbf{C}$. Let p' be the first and q' the second cross point of $\partial\mathbf{C}$ and $\mathcal{S}$ (as seen from p). Let $\mathcal{S}'$ be the part of $\mathcal{S}$ and $\mathcal{C}'$ be the part of the boundary of $\mathbf{C}$ between p' and q' as shown in Figure 16b. Clearly, $\mathcal{S}' * \mathcal{C}'$ forms a simple closed curve $\mathcal{C}''$. $dom(\mathcal{C}'')$ belongs to $\mathbf{P}$ because $\mathbf{P}$ is simply connected; further, it is a part of $\mathcal{O}$-$\|[p,q]$ since $\mathcal{S}'$ and $\mathcal{C}'$ are contained in $\mathcal{O}$-$\|[p,q]$. Any $\mathcal{O}$-line l intersects $\mathcal{S}'$ at most once because $\mathcal{S}'$ is an $\mathcal{O}$-stairsegment. The same holds for $\mathcal{C}'$. Hence, $l \cap dom(\mathcal{C}'')$ is connected and since $l \cap \mathbf{C}$ is also connected, $\mathbf{C} \cup dom(\mathcal{C}'')$ is $\mathcal{O}$-convex and clearly path-connected. This contradicts the maximality of $\mathbf{C}$. $\square$

The $\mathcal{O}^s$-*hull* indeed satisfies the conditions associated with a hull operator, i.e., we have that

$$v, w \in \mathcal{O}^s\text{-}hull(\{p,q\}) \Rightarrow \mathcal{O}^s\text{-}hull(\{v,w\}) \subseteq \mathcal{O}^s\text{-}hull(\{p,q\}).$$

as is shown in the next lemma.

Lemma 5.6 *Let* $\mathbf{P}$ *be a polygon and* p *and* q *be two points in* $\mathbf{P}$. *If* $v, w \in \mathcal{O}^s$-$hull(\{p,q\})$, *then* $\mathcal{O}^s$-$hull(\{v,w\}) \subseteq \mathcal{O}^s$-$hull(\{p,q\})$.

Proof: W.l.o.g. we assume that the maximal $\mathcal{O}$-free range of $\overline{pq}$ is $(0°, 90°)$. Let $\mathcal{S}$ be one arm of $\mathcal{O}^s$-$hull(\{p,q\})$ and assume that $\mathcal{S}$ crosses some path $\mathcal{S}'$ in $\mathcal{O}^s$-$hull(\{v,w\})$. We distinguish two cases.

Case 1 $\Theta(\overline{vw}) \in [0°, 90°]$.
Assume that $\mathcal{S}$ is the upper arm of $\mathcal{O}^s$-$hull(\{p,q\})$. Let p' be the first intersection point of $\mathcal{S}$ and $\mathcal{S}'$ and q' the last. Consider the part $\mathcal{S}''$ of $\mathcal{S}'$ between p' and q'. Since $\mathcal{S}$ crosses $\mathcal{S}'$ and v and w are in $\mathcal{O}^s$-$hull(\{p,q\})$, there is a point on $\mathcal{S}''$ not in $\mathcal{O}^s$-$hull(\{p,q\})$. But by Corollary 5.2 $\mathcal{S}''$ is part of an $\mathcal{O}$-stairsegment from p to q. This contradicts Lemma 5.5 and, hence, $\mathcal{S}$ cannot cross $\mathcal{S}'$. The same reasoning applies to the lower arm of $\mathcal{O}^s$-$hull(\{p,q\})$ and, therefore, $\mathcal{S}'$ is contained in $\mathcal{O}^s$-$hull(\{p,q\})$. Again by Lemma 5.5 we then get that $\mathcal{O}^s$-$hull(\{v,w\}) \subseteq \mathcal{O}^s$-$hull(\{p,q\})$.

Case 2 $\Theta(\overline{vw}) \notin [0°, 90°]$.
Assume that $\mathcal{S}$ is the upper arm of $\mathcal{O}^s$-*hull*$(\{p, q\})$ and w.l.o.g. assume that v is to the left and above w. And consider an $\mathcal{O}$-stairsegment $\mathcal{S}'$ from p to q through v. Clearly, $\mathcal{S}$ has to stay above $\mathcal{S}'$. Now $\mathcal{S}'$ cannot leave the first and third quadrant around v. Since $\mathcal{O}^s$-*hull*$(\{v, w\})$ is contained in the forth quadrant around v, it is completely below $\mathcal{S}'$ and, hence, below $\mathcal{S}$. For similar reasons, $\mathcal{O}^s$-*hull*$(\{v, w\})$ is above the lower arm of $\mathcal{O}^s$-*hull*$(\{p, q\})$. Therefore, $\mathcal{O}^s$-*hull*$(\{v, w\})$ is contained in $\mathcal{O}^s$-*hull*$(\{p, q\})$.

□

In the special case that $\mathcal{O}$-$\|[v, w]$ is a subset of $\mathbf{P}$ we obtain the following result.

Corollary 5.7 *Let* $\mathbf{P}$ *be a polygon and* p *and* q *two points in* $\mathbf{P}$. *For every pair of points* $v, w \in \mathcal{O}^s$-*hull*$(\{p, q\})$ *with* $\mathcal{O}$-$\|[v, w] \subseteq \mathbf{P}$ *we have that* $\mathcal{O}$-$\|[v, w] \subseteq \mathcal{O}^s$-*hull*$(\{p, q\})$.

Proof: Just note that if $\mathcal{O}$-$\|[v, w] \subseteq \mathbf{P}$, then $\mathcal{O}^s$-*hull*$(\{v, w\}) = \mathcal{O}$-$\|[v, w]$. □

5.3 $\mathcal{O}^s$-visibility with Different Sets of Orientations

The first question we are interested in is how $\mathcal{O}^s$-visibility behaves if we change the set of orientations we consider. More precisely, given sets of orientations $\mathcal{O}_i$, $i \in I$, we want to make some statements about the semi-convex space induced by $\bigcup_i \mathcal{O}_i$.

In order to do so we consider $\mathcal{O}^s$-*hull*$(\{p, q\})$, for two points p and q in a polygon $\mathbf{P}$. The boundary of $\mathcal{O}^s$-*hull*$(\{p, q\})$ has a rather special structure. According to Theorem 5.5.1 of [Raw87] the part $\mathcal{S}_1$ of the boundary of $\mathcal{O}^s$-*hull*$(\{p, q\})$ from p to q is a maximal $\mathcal{O}$-stairsegment as is the part $\mathcal{S}_2$ of $\partial\mathcal{O}^s$-*hull*$(\{p, q\})$ from q to p if we assume that $\partial\mathcal{O}^s$-*hull*$(\{p, q\})$ is oriented counterclockwise. If the maximal $\mathcal{O}$-free range of $\overline{pq}$ is (α, β), then $\mathcal{S}_1$ as well as $\mathcal{S}_2$ are $\{\alpha, \beta\}$-stairsegments. Clearly, $\mathcal{S}_1$ and $\mathcal{S}_2$ are staircases and do not cross though they may touch each other. We call $\mathcal{S}_1$ the *lower arm* and $\mathcal{S}_2$ the *upper arm* of $\mathcal{O}^s$-*hull*$(\{p, q\})$.

The arms of $\mathcal{O}^s$-*hull*$(\{p, q\})$ will be the main tool in investigating the structure of the $\mathcal{O}^s$-*hull*. We now can state some properties of the upper and lower arm of $\mathcal{O}^s$-*hull*$(\{p, q\})$.

Lemma 5.8 *Let* $\mathbf{P}$ *be a polygon and* p *and* q *two points in* $\mathbf{P}$. *If* e *is an edge of the boundary of* $\mathcal{O}^s$-*hull*$(\{p, q\})$, *then we have either that* $e \subseteq \partial\mathbf{P}$ *or* $\Theta(e) \in \mathcal{O}$.

Proof: The proof is by contradiction. So suppose there is an edge e on the boundary of $\mathcal{O}^s$-*hull*$(\{p, q\})$ that is neither contained in $\partial\mathbf{P}$ nor do we have $\Theta(e) \in \mathcal{O}$ (see Figure 17). Since $\partial\mathbf{P}$ is a polygonal path, there is a subedge e' of e that does not intersect $\partial\mathbf{P}$. Since $\partial\mathbf{P}$ is closed, we can choose two points r and s on e' that are close enough to each other such that $\mathcal{O}$-$\|[r, s]$ is contained in $\mathbf{P}$. Furthermore, since r and s are on the boundary of $\mathcal{O}^s$-*hull*$(\{p, q\})$, there is one arm of $\mathcal{O}$-$\|[r, s]$ that contains points outside $\mathcal{O}^s$-*hull*$(\{p, q\})$.[10] Now r and s belong to one edge of $\partial\mathcal{O}^s$-*hull*$(\{p, q\})$ and, therefore, there is one arm $\mathcal{S}$ of $\mathcal{O}^s$-*hull*$(\{p, q\})$ that contains both r and s. By Corollary 5.2 we can replace the part of

[10] Note that since $\mathcal{O}$ is closed and $\Theta(e) \notin \mathcal{O}$, $\mathcal{O}$-$\|[r, s]$ does not equal $\overline{pq}$.

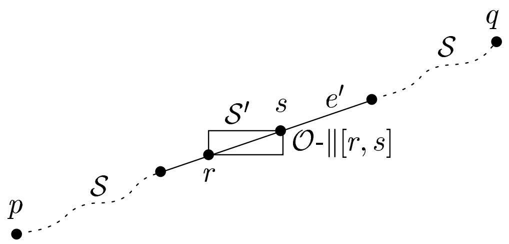

Figure 17: *An edge that is not contained in* $\partial\mathbf{P}$*.*

$\mathcal{S}$ from r to s with one arm of $\mathcal{O}$-$\|[r,s]$ and the resulting new staircase $\mathcal{S}'$ will remain $\mathcal{O}$-convex. But since one arm of $\mathcal{O}$-$\|[r,s]$ is not contained in $\mathcal{O}^s$-$hull(\{r,s\})$ we obtain an $\mathcal{O}$-staircase that is not contained in $\mathcal{O}^s$-$hull(\{p,q\})$ which contradicts the definition of $\mathcal{O}^s$-$hull(\{p,q\})$. Note that this proof also carries through if $|\mathcal{O}| = 1$ since then we just can take any $\mathcal{O}$-convex path from r to s that is not contained in $\mathcal{O}^s$-$hull(\{p,q\})$ which exists by our assumption on e. □

The main result of this section is stated in the following theorem which shows the strong analogy between $\mathcal{O}$-visibility and $\mathcal{O}^s$-visibility.

Theorem 5.9 *Let* $\mathbf{P}$ *be a polygon,* p *and* q *two points in* $\mathbf{P}$*, and* $\mathcal{O}$ *a set of orientations. We have that* p $\mathcal{O}^s$*-sees* q *if and only if* p $\{\theta\}^s$*-sees* q*, for all* $\theta \in \mathcal{O}$*.*

Proof: The "only if" direction of the claim is obvious since any $\mathcal{O}$-staircase from p to q is $\{\theta\}$-convex, for all $\theta \in \mathcal{O}$. So we turn immediately to the "if" direction. If $\Theta(\overline{pq}) \in \mathcal{O}$, then there is a $\theta \in \mathcal{O}$, with $\Theta(\overline{pq}) = \theta$. Since p $\{\theta\}^s$-sees q, we have that $\overline{pq} \subseteq \mathbf{P}$ and, hence, p $\mathcal{O}^s$-sees q.

So suppose that $\Theta(\overline{pq}) \notin \mathcal{O}$. W.l.o.g. we assume that the maximal $\mathcal{O}$-free range of $\Theta(\overline{pq})$ is $(0°, 90°)$ and that p is to the left and below q. By assumption we have that $\{0°\}^s$-$hull(\{p,q\})$ and $\{90°\}^s$-$hull(\{p,q\})$ are both non-empty. Consider the leftmost arm $\mathcal{S}_1$ of $\{0°\}^s$-$hull(\{p,q\})$ and the topmost arm $\mathcal{S}_2$ of $\{90°\}^s$-$hull(\{p,q\})$. By Theorem 5.5.1 of [Raw87] and Lemma 5.5 $\mathcal{S}_1$ is a $\{0°\}$-staircase and $\mathcal{S}_2$ is a $\{90°\}$-staircase from p to q. We claim that the lower enveloppe $\mathcal{S}$ of $\mathcal{S}_1 \cup \mathcal{S}_2$ is a $\{0°, 90°\}$-staircase from p to q. Note that $\mathcal{S}$ is the part of $\mathcal{S}_1 \cup \mathcal{S}_2$ that is contained in $\{0°\}^s$-$hull(\{p,q\}) \cap \{90°\}^s$-$hull(\{p,q\})$. For illustration see Figure 18a.

Let $\mathcal{Q}_1, \ldots, \mathcal{Q}_m$ be the connected components of $\mathcal{S}_1 \cap \mathcal{S}_2$ numbered according to the sequence in which they appear on $\mathcal{S}_1$. We denote the starting point of $\mathcal{Q}_i$ by p_i and the end point by q_i. Since $\mathcal{S}_1$ is horizontally convex and $\mathcal{S}_2$ is vertically convex, $\mathcal{Q}_i$ is $\{0°, 90°\}$-convex, for $1 \leq i \leq m$. More precisely, $\mathcal{Q}_i$ consists of edges with orientation in $[0°, 90°]$. In the following we will show that we can connect q_i to p_{i+1} with an edge of orientation $0°$ or $90°$. This immediately implies that the concatenation of the $\mathcal{Q}_i$ with the connecting edges is a $\{0°, 90°\}$-staircase from p to q by Lemma 5.3.

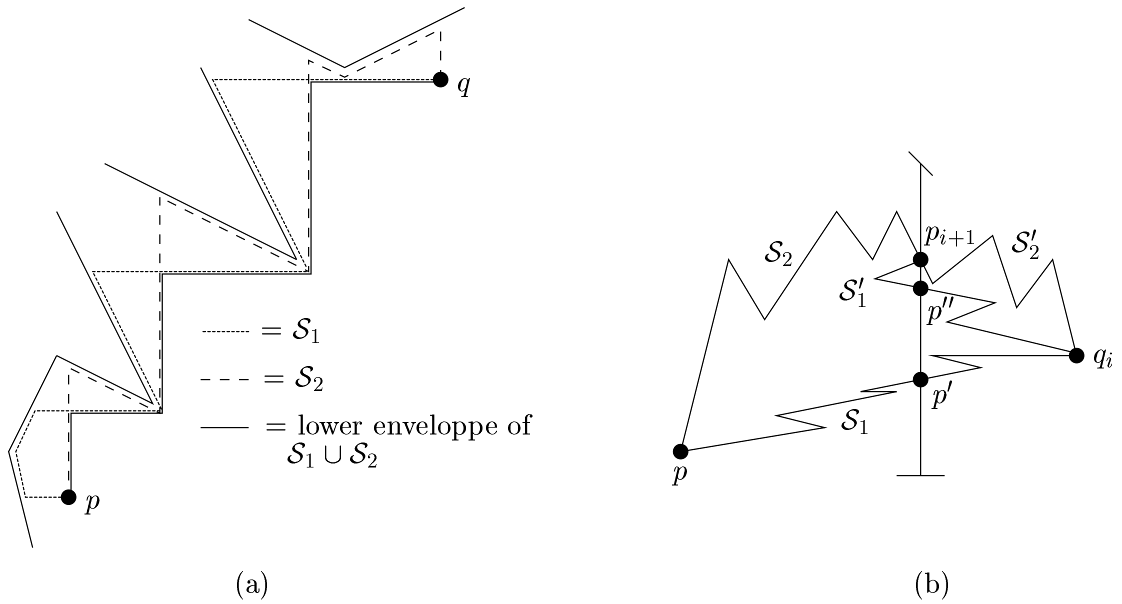

Figure 18: *The lower enveloppe of* $\mathcal{S}_1 \cup \mathcal{S}_2$ *is* $\{0°, 90°\}$*-convex.*

Let $\mathcal{S}_1'$ be the part of $\mathcal{S}_1$ from q_i to p_{i+1} and $\mathcal{S}_2'$ the part of $\mathcal{S}_2$ from q_i to p_{i+1}. Since q_i occurs before p_{i+1} on $\mathcal{S}_1$, q_i is below p_{i+1}. If q_i has the same height as p_{i+1}, then the only $\{0°\}$-staircase from q_i to p_{i+1} is a horizontal line segment and we are done. So assume that p_{i+1} is strictly above q_i. We claim that q_i is to the left of p_{i+1} (refer to Figure 18b). To see this consider the vertical chord $c_{p_{i+1}}$ through p_{i+1}. $c_{p_{i+1}}$ splits $\mathbf{P}$ into two subpolygons one to the left of $c_{p_{i+1}}$ which we denote by $\mathbf{P}_l$ and the other to the right of $c_{p_{i+1}}$ which we denote by $\mathbf{P}_r$. $\mathcal{S}_2$ is vertically convex and, hence, the part of $\mathcal{S}_2$ after p_{i+1} is contained in $\mathbf{P}_r$. Since we assume q_i to be to the right of p_{i+1}, q_i is also contained in $\mathbf{P}_r$. On the other hand, p is in $\mathbf{P}_l$ since p occurs before p_{i+1} on $\mathcal{S}_2$. This implies that $\mathcal{S}_1$ crosses $c_{p_{i+1}}$ in a point p' below p_{i+1} that occurs before q_i which contradicts the fact that $\mathcal{S}_1$ is the leftmost $\{0°\}$-staircase from p to q; for, we can substitute the part of $c_{p_{i+1}}$ from p' to the next intersection point p'' of $\mathcal{S}_1$ above p' for the part of $\mathcal{S}_1$ from p' to p'' and the resulting path is still $\{0°\}$-convex but contains points to the left of $\mathcal{S}_1$. Hence, q_i is to the left of p_{i+1}.

In the same way as above we can prove that $\mathcal{S}_1'$ does not cross the vertical chord $c_{p_{i+1}}$ through p_{i+1}. Therefore, $\mathcal{S}_1'$ and $\mathcal{S}_2'$ are contained in $\mathbf{P}_l$. Similarly, it can be shown that $\mathcal{S}_2'$ does not cross the horizontal chord c_{q_i} through q_i and which implies that $\mathcal{S}_1' \cup \mathcal{S}_2'$ is contained in the part $\mathbf{P}_a$ of $\mathbf{P}$ above c_{q_i}. Hence, $\mathbf{P}_i = \mathbf{P}_l \cap \mathbf{P}_a$ contains both $\mathcal{S}_1'$ and $\mathcal{S}_2'$. Let $\mathcal{P}$ be the part of the boundary of $\mathbf{P}_i$ from q_i to p_{i+1} if we orient $\partial\mathbf{P}_i$ counterclockwise.

We claim that $\mathcal{S}_1'$ does not intersect $\mathcal{P}$. To see this assume the contrary. We consider the horizontal chord c_x through a point x of $\mathcal{S}_1' \cap \mathcal{P}$ as illustrated in Figure 19a. . The part of $\mathcal{S}_1$ after x is contained in the part of $\mathbf{P}_i$ above c_x, and, hence, the left intersection point

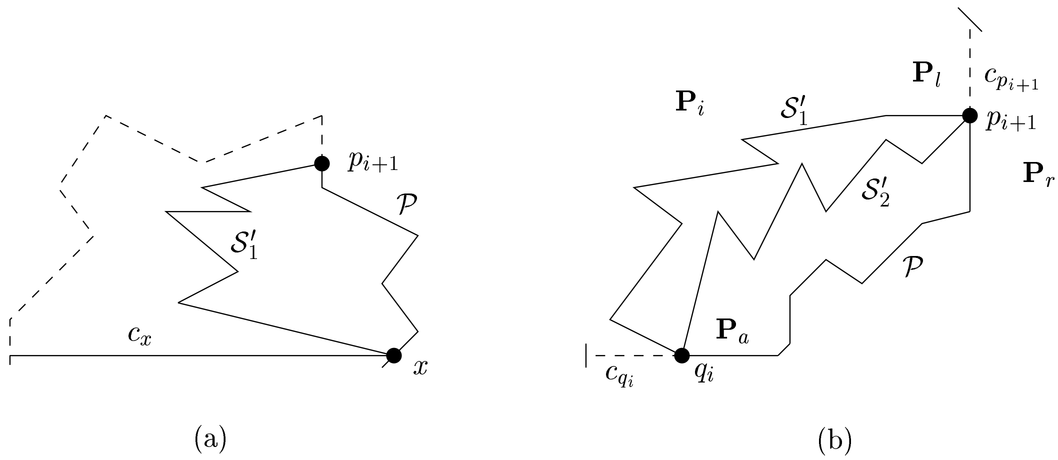

Figure 19: *$\mathcal{S}_1$ is to the left of $\mathcal{P}$.*

of c_x with $\partial\mathbf{P}_i$ does not belong to $\mathcal{P}$ since the part of $\mathbf{P}_i$ above c_x contains p_{i+1}. Similarly, we have that $\mathcal{S}'_2$ does not intersect $\mathcal{P}$.

Thus, $\mathcal{C}_1 = \mathcal{S}'_1 * \mathcal{P}$, $\mathcal{C}_2 = \mathcal{S}'_2 * \mathcal{P}$, and $\mathcal{C}_3 = \mathcal{S}'_1 * \mathcal{S}'_2$ are simple closed curves. Since $\mathcal{S}'_1$, $\mathcal{S}'_2$, and $\mathcal{P}$ start and end in the same points, either $\mathcal{C}_1$ contains $\mathcal{S}'_2$ or $\mathcal{C}_2$ contains $\mathcal{S}'_1$ or $\mathcal{C}_3$ contains $\mathcal{P}$. Since $\mathbf{P}$ is a simple polygon and $\mathcal{S}'_1$ and $\mathcal{S}'_2$ are contained in $\mathbf{P}$, $\mathcal{P}$ as part of the boundary of $\mathbf{P}$ cannot be contained in $\mathcal{S}'_1 * \mathcal{S}'_2$. So suppose w.l.o.g. that $\mathcal{S}_2$ is contained in $\mathcal{C}_1$ (see Figure 19b).

Since an edge of $\mathcal{S}'_2$ is non-vertical only if it is contained $\partial\mathbf{P}$, $\mathcal{S}'_2$ contains no non-vertical edges. Therefore, $\mathcal{S}'_2$ consists of a single vertical edge from q_i to p_{i+1}. This implies that $\mathcal{Q}_1 * \overline{q_1 p_2} * \mathcal{Q}_2 * \cdots * \mathcal{Q}_{m-1} * \overline{q_{m-1} p_m} * \mathcal{Q}_m$ is a $\{0°, 90°\}$-staircase from p to q that is contained in $\mathbf{P}$ and, therefore, p $\mathcal{O}^s$-sees q. □

Observe that the claim is false for polygons with holes. A counterexample is depicted in Figure 20. q can be reached from p with a $\{0°\}$- and a $\{90°\}$-staircase but there exists no $\{0°, 90°\}$-staircase from p to q that does not intersect the hole of $\mathbf{P}$.

The above theorem leads to following slightly more general corollary.

Corollary 5.10 *Let $\mathbf{P}$ be a polygon, p and q two points in $\mathbf{P}$ and $\mathcal{O}_i$, $i \in I$, a family of sets of orientations and $\mathcal{O}$ their union, i.e., $\mathcal{O} = \bigcup_{i\in I} \mathcal{O}_i$. p $\mathcal{O}^s_i$-sees q, for all $i \in I$, if and only if p $\mathcal{O}^s$-sees q.*

Proof: If p $\mathcal{O}^s_i$-sees q, for all $i \in I$, then we, in particular, have that p $\{\theta\}^s$-sees q, for all $\theta \in \bigcup \mathcal{O}_i = \mathcal{O}$. Theorem 5.9 now implies that p $\mathcal{O}^s$-sees q. The reverse implication is obvious. □

It is interesting to note that although the visibility structure of $\mathcal{O}$-convexity and $\mathcal{O}^s$-convexity is essentially the same (an analog to Corollary 5.10 is contained in the proof of Lemma 4.1), the hull properties differ considerably. Recall that we have for the $\mathcal{O}$-hull of

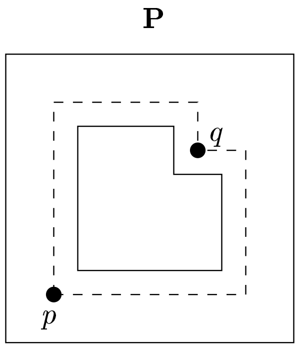

Figure 20: *p can not $\{0°, 90°\}^s$-see q.*

two points p and q that

$$\mathcal{O}\text{-}hull(\{p,q\}) = \bigcup_{i \in I} \mathcal{O}_i\text{-}hull(\{p,q\})$$

and contrast this with the following corollary.

Corollary 5.11 *Let $\mathbf{P}$ be a polygon and p and q two points in $\mathbf{P}$. If $\mathcal{O}_i$, $i \in I$, is a family of sets of orientations and $\mathcal{O}$ is their union, then we have that*

$$\mathcal{O}^s\text{-}hull(\{p,q\}) = \bigcap_{i \in I} \mathcal{O}_i^s\text{-}hull(\{p,q\}).$$

Proof: Let $\mathbf{I} := \bigcap_{i \in I} \mathcal{O}_i^s\text{-}hull(\{p,q\})$. Since every $\mathcal{O}$-stairsegment is an $\mathcal{O}_i$-stairsegment, for all $i \in I$, we obviously have $\mathcal{O}^s\text{-}hull(\{p,q\}) \subseteq \mathbf{I}$.

Now let r be a point $\mathbf{I}$. Since $r \in \mathbf{I}$, p $\mathcal{O}_i^s$-sees r (in $\mathbf{P}$), for all $i \in I$. The prerequisites of Corollary 5.10 are satisfied, and therefore p $\mathcal{O}^s$-sees r. The same holds for q. Hence, there is an $\mathcal{O}$-stairsegment $\mathcal{S}_{pr}$ from p to r and an $\mathcal{O}$-stairsegment $\mathcal{S}_{rq}$ from r to q. With Lemma 4.3.1 of [Raw87] it can be easily seen that $\mathcal{O}\text{-}\|[p,q] = \bigcap_i \mathcal{O}_i\text{-}\|[p,q]$. Since $\mathcal{O}_i^s\text{-}hull(\{p,q\}) \subseteq \mathcal{O}_i\text{-}\|[p,q]$, we have that $r \in \mathcal{O}\text{-}\|[p,q]$ and as in the proof of Lemma 5.5 Lemma 5.1 then implies that $\mathcal{S} = \mathcal{S}_{pr} * \mathcal{S}_{rq}$ is an $\mathcal{O}$-stairsegment from p to q containing r. By Definition 5.3 we have that $\mathcal{S} \subseteq \mathcal{O}^s\text{-}hull(\{p,q\})$ and, hence, $r \in \mathcal{O}^s\text{-}hull(\{p,q\})$ which proves the claim. □

The main implication of Theorem 5.9 is the following corollary which is a restatement of Theorem 6.4.2 in [Raw87]. Unfortunately, the proof presented in [Raw87] is not correct.

Corollary 5.12 *If $\mathbf{P}$ is a simple polygon, $\mathcal{O}_i$, $i \in I$, a family of sets of orientations and $\mathcal{O}$ their union, then*

$$\mathcal{O}^s\text{-}kernel(\mathbf{P}) = \textstyle\bigcap_i \mathcal{O}_i^s\text{-}kernel(\mathbf{P}).$$

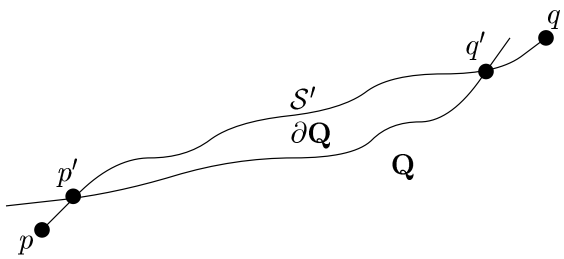

Figure 21: *Any $\mathcal{O}$-staircase from p to q belongs to* $\mathbf{Q}$.

Proof: Let q be a point in $\mathbf{P}$. If p is in $\mathcal{O}^s$-$kernel(\mathbf{P})$, then p $\mathcal{O}^s$-sees q; since every $\mathcal{O}$-stairsegment is also an $\mathcal{O}_i$-stairsegment, for all $i \in I$, we have $p \in \bigcap_i \mathcal{O}_i^s$-$kernel(\mathbf{P})$. On the other hand, if p is in $\bigcap_i \mathcal{O}_i^s$-$kernel(\mathbf{P})$, then p $\mathcal{O}_i^s$-sees q, for all $i \in I$ and Corollary 5.10 implies that p $\mathcal{O}^s$-sees q. Hence, $p \in \mathcal{O}^s$-$kernel(\mathbf{P})$. □

The above corollary will be very helpful in the next section where we are concerned with algorithms to compute $\mathcal{O}^s$-$kernel(\mathbf{P})$. But before we turn to this problem we take a look at $\mathcal{O}^s$-skulls which correlate nicely to maximal, connected $\mathcal{O}$-convex sets.

5.4 $\mathcal{O}^s$-Skulls

In order to use the notion of skull correctly we have to specify the semi-convex space we are looking at. So let $\mathbf{P}$ be a polygon. We define the family $\mathcal{O}$-$\mathcal{C}_{\mathbf{P}}$ of $\mathcal{O}^s$-convex sets as the family of all sets $\mathbf{C} \subseteq \mathbf{P}$ such that, for any two points p and q in $\mathbf{C}$, we have that $\mathcal{O}^s$-$hull(\{p,q\}) \subseteq \mathbf{C}$. An $\mathcal{O}^s$-*skull* $\mathbf{Q}$ of $\mathbf{P}$ is a maximal $\mathcal{O}^s$-convex subset of $\mathbf{P}$. We denote the set of $\mathcal{O}^s$-skulls of $\mathbf{P}$ by $\mathcal{O}^s$-$skulls(\mathbf{P})$.

Recall that in Theorem 5.5.2 of [Raw87] it is shown that in a connected $\mathcal{O}$-convex set any two points $\mathcal{O}^s$-see each other. Furthermore, note that an $\mathcal{O}^s$-convex set in a polygon $\mathbf{P}$ is necessarily $\mathcal{O}$-convex and connected but the reverse is not true. But if we consider $\mathcal{O}^s$-skulls and the *maximal* $\mathcal{O}$-convex connected subsets of a polygon $\mathbf{P}$, we can show that the two notions coincide.

Lemma 5.13 *If* $\mathbf{P}$ *is a polygon and* $\mathbf{Q}$ *a maximal, connected* $\mathcal{O}$*-convex subset of* $\mathbf{P}$*, then* $\mathbf{Q} \in \mathcal{O}^s$-$skulls(\mathbf{P})$.

Proof: The proof is by contradiction. Let $\mathbf{Q}$ be a maximal, connected $\mathcal{O}$-convex subset of $\mathbf{P}$. Suppose there are two points p and q in $\mathbf{Q}$ such that $\mathcal{O}^s$-$hull(\{p,q\})$ is not contained in $\mathbf{Q}$, i.e., there is an $\mathcal{O}$-staircase $\mathcal{S}$ from p to q that is not contained in $\mathbf{Q}$. Hence, $\mathcal{S}$ intersects $\partial\mathbf{Q}$ at least twice. W.l.o.g. we assume that the maximal $\mathcal{O}$-free range of $\overline{pq}$ is $(0°, 90°)$ and that p is to the left and below q. Therefore, $\mathcal{S}$ is a $\{0°, 90°\}$-staircase.

Let p' be the first cross point of $\mathcal{S}$ with $\partial\mathbf{Q}$ and q' be the second (see Figure 21). Since $\mathcal{S}$ is a $\{0°, 90°\}$-staircase, q' is to the right and above p'. By Theorem 5.5.1 of [Raw87]

$\partial\mathbf{Q}$ is also a $\{0°, 90°\}$-staircase from p' to q'. If we replace the part of $\partial\mathbf{Q}$ from p' to q' with the part $\mathcal{S}'$ of $\mathcal{S}$ from p' to q', then the new boundary still encloses an $\mathcal{O}$-convex set $\mathbf{Q}'$ by a similar argument as in the proof of Lemma 5.5. Since $\mathcal{S}'$ consists only of points outside $\mathbf{Q}$, we have that $\mathbf{Q}'$ properly contains $\mathbf{Q}$ which contradicts the maximality of $\mathbf{Q}$. This concludes the proof. □

If we denote the family of maximal connected $\mathcal{O}$-convex sets of $\mathbf{P}$ by $\mathcal{O}\text{-}\mathcal{S}(\mathbf{P})$, the above lemma states that $\mathcal{O}\text{-}\mathcal{S}(\mathbf{P}) \subseteq \mathcal{O}^s\text{-}skulls(\mathbf{P})$. We want to show that actually equality holds. So let $\mathbf{Q}$ be a set in $\mathcal{O}^s\text{-}skulls(\mathbf{P})$. Clearly, $\mathbf{Q}$ is connected and $\mathcal{O}$-convex. Our first claim is that there exists a set $\mathbf{Q}' \in \mathcal{O}\text{-}\mathcal{S}(\mathbf{P})$ containing $\mathbf{Q}$. To see this consider a nested chain $\mathcal{N}$ of connected $\mathcal{O}$-convex sets containing $\mathbf{Q}$ and let $\mathbf{C} = \bigcup\mathcal{N}$. $\mathbf{C}$ is clearly connected and $\mathcal{O}$-convex. Hence, by Zorn's Lemma there exists a maximal connected $\mathcal{O}$-convex set $\mathbf{Q}' \in \mathcal{O}\text{-}\mathcal{S}(\mathbf{P})$ containing $\mathbf{Q}$. By the above lemma we have $\mathbf{Q}' \in \mathcal{O}^s\text{-}skulls(\mathbf{P})$. The definition of a skull then implies that $\mathbf{Q}' = \mathbf{Q}$ and, hence, $\mathbf{Q} \in \mathcal{O}\text{-}\mathcal{S}(\mathbf{P})$.

Note that the semi-convex space $(\mathbf{P}, \mathcal{O}\text{-}\mathcal{C}_{\mathbf{P}})$ is induced by the norm $\|\cdot\|_{\mathcal{O}}$. By Lemma 5.4 and Theorem 5.3 of [SW91] we have that $(\mathbf{P}, \mathcal{O}\text{-}\mathcal{C}_{\mathbf{P}})$ satisfies the join condition which is stated in the following corollary for further reference.

Corollary 5.14 *If $\mathbf{P}$ is a simple polygon, then $(\mathbf{P}, \mathcal{O}\text{-}\mathcal{C}_{\mathbf{P}})$ satisfies the join-condition.*

This, in particular, implies that we can apply the Kernel Theorem of [SW91], i.e., we have

$$\mathcal{O}\text{-}\mathcal{C}_{\mathbf{P}}\text{-}kernel(\mathbf{P}) = \bigcap \mathcal{O}\text{-}\mathcal{C}_{\mathbf{P}}\text{-}skulls(\mathbf{P}).$$

Note that $\mathcal{O}\text{-}\mathcal{C}_{\mathbf{P}}\text{-}skulls(\mathbf{P})$ is just another name for $\mathcal{O}^s\text{-}skulls(\mathbf{P})$ by definition. Lemma 5.13 now yields the connection to $\mathcal{O}^s$-visibility. $\mathcal{O}^s\text{-}kernel(\mathbf{P})$ equals $\mathcal{O}\text{-}\mathcal{C}_{\mathbf{P}}\text{-}kernel(\mathbf{P})$ since there is an $\mathcal{O}$-stairsegment from p to q if and only if $\mathcal{O}^s\text{-}hull(\{p, q\}) \neq \emptyset$. On the other hand, we have that the set of all maximal, connected $\mathcal{O}$-convex subsets of $\mathbf{P}$ equals $\mathcal{O}^s\text{-}skulls(\mathbf{P})$. We conclude that $\mathcal{O}^s\text{-}kernel(\mathbf{P}) = \bigcap \mathcal{O}\text{-}\mathcal{S}(\mathbf{P})$. Note that on both sides of the equation we have terms which are defined just by making use of $\mathcal{O}^s$-visibility *independently* of the notion of $\mathcal{O}^s$-convexity. An immediate consequence of the above equation is the following corollary.

Corollary 5.15 *If $\mathbf{P}$ is a simple polygon, then the $\mathcal{O}^s$-kernel of $\mathbf{P}$ is $\mathcal{O}^s$-convex. In particular, $\mathcal{O}^s\text{-}kernel(\mathbf{P})$ is a connected, $\mathcal{O}$-convex set.*

The semi-convex space $(\mathbf{P}, \mathcal{O}\text{-}\mathcal{C}_{\mathbf{P}})$ obviously models the aspects of $\mathcal{O}^s$-visibility we are interested in quite adequately. Hence, we feel that other properties of this interesting example of a semi-convex space should be explored. As an example we present another corollary concerning $\mathcal{O}^s$-skulls which shows that $\mathcal{O}^s$-skulls can be obtained by intersection from $\mathcal{O}_i^s$-skulls.

Corollary 5.16 *Let $\mathbf{P}$ be a polygon in the plane, $\mathcal{O}_i$ set of orientations, for $i \in I$, and $\mathcal{O}$ their union. If $\mathbf{Q} \in \mathcal{O}^s\text{-}skulls(\mathbf{P})$ and $\mathbf{Q}_i \in \mathcal{O}_i^s\text{-}skulls(\mathbf{P})$ with $\mathbf{Q} \subseteq \mathbf{Q}_i$, for all $i \in I$, then $\mathbf{Q} = \bigcap_i \mathbf{Q}_i$.*

Proof: We clearly have that $\mathbf{Q} \subseteq \bigcap_i \mathbf{Q}_i$. So now let p be a point in $\bigcap_i \mathbf{Q}_i$. Since $p \in \bigcap_i \mathbf{Q}_i \supseteq \mathbf{Q}$, p $\mathcal{O}_i^s$-sees all points in $\mathbf{Q}$, for all $i \in I$. Corollary 5.10 implies that p $\mathcal{O}^s$-sees all points in $\mathbf{Q}$. Hence, the $\mathcal{O}^s$-join of p and $\mathbf{Q}$ is non-empty and contained in $\mathbf{P}$. Since we have that the $\mathcal{O}^s$-join of a point and an $\mathcal{O}^s$-convex set is $\mathcal{O}^s$-convex by Corollary 5.14, the maximality of $\mathbf{Q}$ implies that p is contained in $\mathbf{Q}$. □

6 Computing the $\mathcal{O}^s$-Kernel of a Polygon

In the following we turn to the computational aspects of $\mathcal{O}^s$-visibility. In particular, we are interested in finding the $\mathcal{O}^s$-kernel of a polygon. In the first part of this section we are concerned with simple polygons. With Corollary 5.12 as the analogon to the Intersection Lemma of Section 4 the computation of the $\mathcal{O}^s$-kernel is reduced to the task of computing the $\{\theta\}^s$-kernel, for all $\theta \in \mathcal{O}$, and of intersecting the resulting kernels. With the help of a simple observation we show how to obtain an algorithm that runs in time $O(n \log |\mathcal{O}|)$, for finite $\mathcal{O}$. The second part then deals with multiply connected polygons. Unfortunately, we lose many of the more pleasing properties that hold for the $\mathcal{O}^s$-kernel of simple polygons. This leads to the notion of the *external $\mathcal{O}^s$-kernel* of a polygon. The last part is then concerned with the relationship between the $\mathcal{O}$- and $\mathcal{O}^s$-kernels where we show that the $\mathcal{O}^s$-kernel is a component of the $\mathcal{O}$-kernel.

6.1 The $\mathcal{O}^s$-kernel for Simple Polygons

As we noted above Corollary 5.12 leaves us with the computation of the $\mathcal{O}^s$-kernel in the case $|\mathcal{O}| = 1$. Of course, we can again assume that $\mathcal{O} = \{0°\}$. In the following we give a characterization the $\{0°\}^s$-kernel of a polygon. Although the characterization turns out to be quite a bit simpler than the one for the $\{0°\}$-kernel of a polygon, the proofs get more involved. In particular, we need the following very basic lemma about Jordan curves which relates reflex and convex minima on the curve.

Lemma 6.1 *Let $\mathcal{J}$ be a Jordan curve. If p and q are two (different) convex minima (maxima) of $\mathcal{J}$, then there is at least one reflex maximum (minimum) of $\mathcal{J}$ above (below) them.*

Proof: We only prove the lemma in the case that p and q are two convex minima of $\mathcal{J}$. We assume that $\mathcal{J}$ is oriented in such a way that the interior of $\mathcal{J}$ is locally to the left of $\mathcal{J}$. W.l.o.g. suppose that p is below q (see Figure 22a for illustration). We prove the claim by contradiction. So from now on suppose that all maxima of $\mathcal{J}$ above q are convex. We denote the part of $\mathcal{J}$ from q to p by $\mathcal{Q}$. Let h_q be the horizontal line through q and c_q be the connected of component of $\mathcal{J} \cap h_q$ that contains q (see Figure 22b).

Our first claim is that the "first" intersection point of $\mathcal{Q}$ and h_q is to the left of c_q. By the first intersection point we mean the first point of $\mathcal{Q}$ that intersects h_q if we follow the orientation of $\mathcal{Q}$ from q to p. To see this note that p is below q and, hence, $\mathcal{Q}$ intersects h_q.

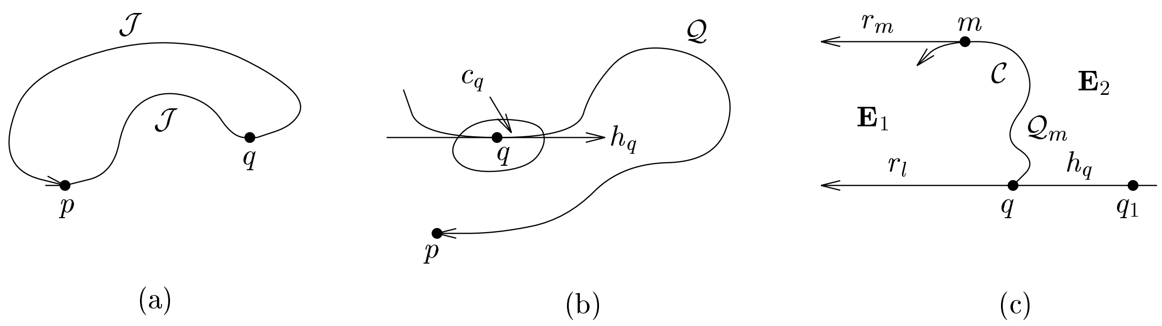

Figure 22: *There is a reflex maximum between p and q.*

Suppose the first intersection point q_1 of $\mathcal{Q}$ with h_q is to the right of c_q. Since c_q is a convex minimum, there are points of $\mathcal{Q}$ in any neighbourhood of c_q that are above h_q. Therefore, there is a maximum in the part $\mathcal{Q}_1$ of $\mathcal{Q}$ from q to q_1 which is above q and q_1. Let m be the leftmost of the highest maxima of $\mathcal{Q}_1$. By our assumption on $\mathcal{J}$, m is convex. Consider the biinfinite curve $\mathcal{C}$ that consists of the ray r_l towards the left starting in q, the part $\mathcal{Q}_m$ of $\mathcal{Q}_1$ from q to m, and the ray r_m towards the left in m as shown in Figure 22c. By our choice of m and the assumption that $\mathcal{Q}$ does not intersect r_l, $\mathcal{C}$ is simple and divides the plane into two connected components $\mathbf{E}_1$ and $\mathbf{E}_2$. Let $\mathbf{E}_1$ be the component that is locally above r_l and $\mathbf{E}_2$ the component that is locally below r_l. By definition q_1 does not belong to $\mathcal{Q}_m$ or r_m since m is above q_1 and $\mathcal{Q}_m$ is only a part of $\mathcal{Q}_1$. Furthermore, q_1 is to the right of q and, thus, not part of r_l. By our choice of q_1, $\mathcal{Q}_m$ does not intersect h_q. This implies that $\mathcal{Q}_m$ stays above h_q and it can be easily seen that there are paths that connect q_1 with points locally below r_l without crossing $\mathcal{C}$. On the other hand, the points locally above r_l can be connected with a path that does not cross $\mathcal{C}$ to the points locally below r_m and, thus, also to points locally below m. Since m is convex, there are points of $\mathcal{Q}_1$ in the part from m to q_1 that are locally below r_m, i.e., in $\mathbf{E}_1$. Hence, $\mathcal{Q}_1$ intersects points in $\mathbf{E}_1$ and $\mathbf{E}_2$ and, therefore, crosses $\mathcal{C}$. By our choice of m, $\mathcal{Q}_1$ does not cross r_m and, furthermore, since $\mathcal{J}$ is simple $\mathcal{Q}_1$ does not cross $\mathcal{Q}_m$. This implies that $\mathcal{Q}_1$ intersects r_l in contradiction to our choice of q_1.

So from now on we can assume that the first intersection point q_1 of $\mathcal{Q}$ with h_q is to the left of c_q. As above we denote the part of $\mathcal{Q}$ from q to q_1 by $\mathcal{Q}_1$. Note that $\mathcal{Q}_1$ stays above h_q. We consider the Jordan curve $\mathcal{J}_1$ which consists of $\mathcal{Q}_1$ concatenated with the line segment $\overline{q_1 q}$ from q_1 to q (see Figure 23a). Since $\mathcal{Q}_1$ stays above h_q, the interior of $\mathcal{J}_1$ is locally above $\overline{q_1 q}$ and, hence, also locally above c_q. Let $\mathcal{P}$ be the part of $\mathcal{J}$ q_1 to q. Since c_q is a convex minimum, there are points of $\mathcal{J}$ before q, i.e., points of $\mathcal{P}$, that are locally above c_q. Therefore, $\mathcal{P}$ intersects $int(\mathcal{J}_1)$. Let p_1 be the last intersection point of $\mathcal{P}$ with $\overline{q_1 q}$ and $\mathcal{P}_1$ the part of $\mathcal{P}$ from p_1 to q. Since the points of $\mathcal{P}_1$ that are locally before q belong to $int(\mathcal{J}_1)$, $\mathcal{P}_1$ is contained in $int(\mathcal{J}_1)$ and there is maximum of $\mathcal{P}_1$ contained in $int(\mathcal{J}_1)$. We choose m_1 to the rightmost of the highest maxima of $\mathcal{P}_1$. By our assumption

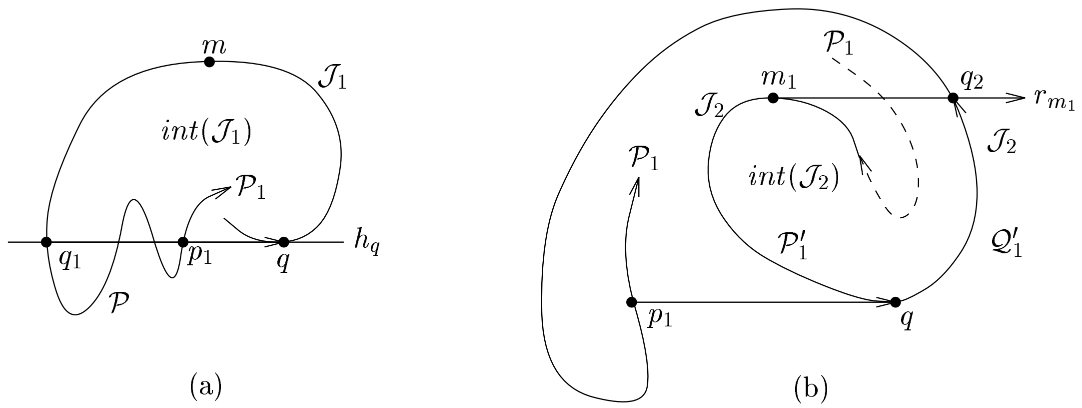

Figure 23: *m_1 is not convex.*

on $\mathcal{J}$, m_1 is convex. We will show that this leads to a contradiction.

We now consider the ray r_{m_1} towards the right starting in m_1. Since m_1 is contained in $int(\mathcal{J}_1)$, r_{m_1} intersects $\mathcal{J}_1$. Let q_2 be the first intersection point of r_{m_1} with $\mathcal{Q}_1$. We denote the part of $\mathcal{P}_1$ from m_1 to q by $\mathcal{P}_1'$ and the part of $\mathcal{Q}_1$ from q to q_2 by $\mathcal{Q}_1'$. $\mathcal{J}_2 = \mathcal{P}_1' * \mathcal{Q}_1' * \overline{q_2 m_1}$ is a Jordan curve whose interior is locally below $\overline{q_2 m_1}$ (see Figure 23b). To see this just note the interior of $\mathcal{J}_2$ is locally above q since $\mathcal{J}_2$ is above h_q and by following $\mathcal{P}_1'$ we can connect the points locally below $\overline{q_2 q}$ with the points locally above q and, hence, these points also belong to $int(\mathcal{J}_2)$.

Since m_1 is convex, the part of $\mathcal{P}_1$ before m_1 intersects $int(\mathcal{J}_2)$. Furthermore we have that p_1 does not belong to $int(\mathcal{J}_2)$ since otherwise there are at least two intersection components of $\mathcal{J}_2$ with h_q, one to the left and one to the right of p_1, which contradicts the fact that $\mathcal{P}_1'$ and $\mathcal{Q}_1'$ intersect h_q only in c_q. Hence, $\mathcal{P}_1$ intersects points in the interior and exterior of $\mathcal{J}_2$ and, thus, crosses $\mathcal{J}_2$. Furthermore, since $\mathcal{J}$ is simple, it intersects $\overline{q_2 m_1}$. But we chose m_1 to be one of the highest maxima of $\mathcal{P}_1$ in the interior of $\mathcal{J}_2$ which contradicts the fact that $\mathcal{P}_1$ crosses $\overline{q_2 m_1}$. This yields the desired contradiction. Hence, one of the maxima above p and q is reflex. □

Of course, we get the same result if we exchange convex and reflex by reverting the orientation of the curve. Now we turn to the characterization of the $\{0°\}^s$-kernel of a polygon. As we indicated above the conditions can be easily tested which results in a simple algorithm. This result was first stated in [Raw87, Lemma 6.4.1]. Unfortunately, the proof contains only the easy "only if" part which we also give for completeness.

Theorem 6.2 *If* $\mathbf{P}$ *is a simple polygon, then* $p \in \{0°\}^s$*-kernel*$(\mathbf{P})$ *if and only if*

1. *p is above every reflex maximum of* $\mathbf{P}$ *and*

2. *p is below every reflex minimum of* $\mathbf{P}$.

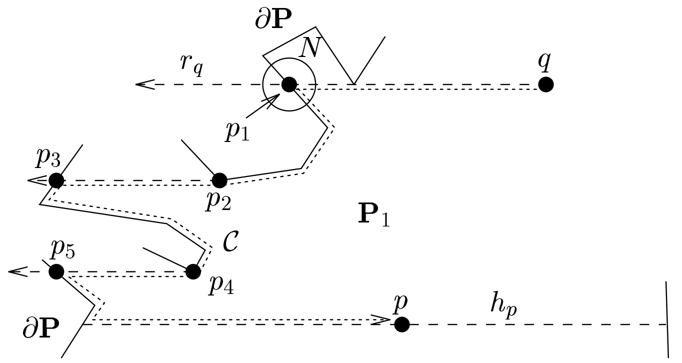

Figure 24: *The construction of a $\{0°\}$-convex path from p to q.*

Proof: As mentioned above, the only if part of the claim is very easy to see. W.l.o.g. assume that p is below a reflex maximum and let q be some point of the same height as p on the other side of the maximum. Any curve that connects p and q cannot be $\{0°\}$-convex.

Now let p be above every reflex maximum and below every reflex minimum and q some point in $\mathbf{P}$. W.l.o.g. assume that q is above p. We construct a $\{0°\}$-convex curve $\mathcal{C}$ that connects p and q (for illustration refer to Figure 24). To do this let r_q be the horizontal ray towards the left starting in q. We follow r_q until we reach the exterior of $\mathbf{P}$. Let the cross point with the boundary of $\mathbf{P}$ be p_1. Consider a sufficiently small neighbourhood N of p_1 (see Figure 24). Since we cross $\partial\mathbf{P}$, there is some part of $\partial\mathbf{P} \cap N$ below r_q. We follow the part of $\partial\mathbf{P}$ below r_q until we encounter a reflex minimum or intersect the maximal horizontal line segment h_p through p in p_2. In the latter is the case we just follow h_p to p which is always possible since h_p is a chord in $\mathbf{P}$. Otherwise, we continue from p_2 to the left until we again intersect $\partial\mathbf{P}$ and continue the process. We claim that the resulting curve $\mathcal{C}$ finally reaches p and that $\mathcal{C}$ is $\{0°\}$-convex.

To see this note that h_p splits the polygon $\mathbf{P}$ into a number of components. Let $\mathbf{P}_1$ be the component that contains q. By the construction of $\mathcal{C}$ we do not cross h_p and, hence, $\mathcal{C}$ is contained in $\mathbf{P}_1 \cup h_p$. We first show that $\mathcal{C}$ is $\{0°\}$-convex. For assume the contrary. Then, there is a horizontal line h that intersects $\mathcal{C}$ in two disconnected components. Clearly, both components belong to $\mathbf{P}_1$ since the part of $\mathcal{C}$ that is contained in h_p is connected by construction.

There is a minimum m in the part of $\mathcal{C}$ between the two disconnected components. By the construction of $\mathcal{C}$ m is a convex minimum that belongs to $\partial\mathbf{P}$. If m is on h_p, then $\mathcal{C}$ does not continue after p which contradicts the fact that h intersects $\mathcal{C}$ in two disconnected components. Hence, suppose that m is not on h_p.

Note that h_p contains a convex minimum of $\mathbf{P}_1$. In the following we show that there is no other convex minimum in $\mathbf{P}_1$. This immediately yields the desired contradiction since we already noted above that $\mathcal{C}$ is contained in $\mathbf{P}_1 \cup h_p$. So assume there is a convex minimum m in $\mathbf{P}_1$ that is not on h_p. Lemma 6.1 yields that there is a reflex maximum r

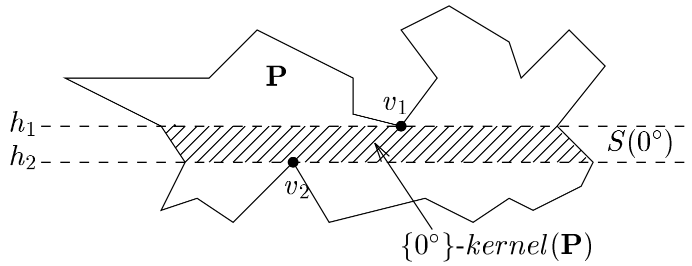

Figure 25: *$\mathcal{O}^s$-kernel*(**P**) *is the intersection of* $S(0°)$ *and* **P**.

between h_p and m and above both. Clearly, r belongs to $\partial\mathbf{P}$ and, therefore, this contradicts Conditions 1 of the claim. Hence, there can be no line that intersects $\mathcal{C}$ in two components and $\{0°\}$-convex.

Since $\mathcal{C}$ is $\{0°\}$-convex, i.e., monotone and we follow each edge to its end point, we encounter each edge of **P** at most once. Hence, the procedure has to stop after a finite number of steps. Since we stop the construction only if we reach p, $\mathcal{C}$ is a $\{0°\}$-convex path from q to p as claimed. □

By the above theorem we are immediately led to an algorithm to compute $\{0°\}^s$-*kernel*(**P**). We just have to find the lowest reflex minimum v_l and highest reflex maximum v_h and then output the left and right parts of the boundary of **P** with a y-range in the interval $[v_h, v_l]$. All of this can be done in linear time. For more details see p. 75/76 of [Raw87]. We summarize our results in the following theorem.

Theorem 6.3 *Let* **P** *be a simple polygon, then* $\{0°\}^s$-*kernel*(**P**) *can be computed in time linear in the number of vertices.*

6.2 A Fast Algorithm for the $\mathcal{O}^s$-Kernel of a Polygon

Above we presented an algorithm which computes the $\{\theta\}^s$-kernel of a polygon **P**, for each $\theta \in \mathcal{O}$, and then intersects the kernels obtained to yield $\mathcal{O}^s$-*kernel*(**P**), for finite $\mathcal{O}$. Since we have to apply the $\{\theta\}^s$-kernel subroutine $|\mathcal{O}|$-times, the algorithm needs at least time $\Omega(n \cdot |\mathcal{O}|)$ to compute $\mathcal{O}^s$-*kernel*(**P**) even if we do not take the time to intersect the $\{\theta\}^s$-kernels of **P** into account.

In the following we will give an algorithm to compute the $\mathcal{O}^s$-kernel of a simple polygon **P** that only takes $O(n \log |\mathcal{O}|)$ time if $\mathcal{O}$ is given as a sorted set. In order to develop an algorithm that obtains the above time bound, we have to view the $\{0°\}^s$-kernel a little differently. Let v_1 be the lowest reflex minimum and v_2 the highest reflex maximum of **P**. If we denote the horizontal line through v_i by h_i, for $i = 1, 2$, then Theorem 6.2 implies that $\{0°\}^s$-*kernel*(**P**) is the intersection of the slab $S(0°)$ between h_1 and h_2 with **P** (see Figure 25). Hence, if we are given the slabs $S(\theta)$ that lie between the θ-oriented lines through the lowest reflex θ-minimum and the highest reflex θ-maximum, for all $\theta \in \mathcal{O}$, then we have that

$$\begin{aligned}
\mathcal{O}^s\text{-}kernel(\mathbf{P}) &= \bigcap_{\theta\in\mathcal{O}} (S(\theta)\cap\mathbf{P}) \\
&= \left(\bigcap_{\theta\in\mathcal{O}} S(\theta)\right)\cap\mathbf{P}.
\end{aligned}$$

In the following we denote $\bigcap_{\theta\in\mathcal{O}} S(\theta)$ by $S_{\mathcal{O}}(\mathbf{P})$. Since $S(\theta)$ is the intersection of two θ-oriented halfplanes, $S_{\mathcal{O}}(\mathbf{P})$ is strong-$\mathcal{O}$-convex by Theorem 6.1.3 of [Raw87]. This, in particular, implies that it is convex and consists of at most $2|\mathcal{O}|$ edges. Hence, $S_{\mathcal{O}}(\mathbf{P})\cap\mathbf{P}$ can be computed in time $O(n\log|\mathcal{O}|)$ by testing, for each edge e of $\mathbf{P}$, if e intersects an edge of $\partial S_{\mathcal{O}}(\mathbf{P})$. This yields the intersection points of $\partial S_{\mathcal{O}}(\mathbf{P})$ with $\partial\mathbf{P}$ and enables us to compute the vertices of $S_{\mathcal{O}}(\mathbf{P})\cap\mathbf{P}$ in additional linear time.

So we only have to show how to compute $S_{\mathcal{O}}(\mathbf{P})$ in time $O(n\log|\mathcal{O}|)$. In order to do so we introduce the following notation. Given a point p in the plane we denote the (closed) halfplane to the left of the θ-oriented line through p by $h^+(p,\theta)$. For a given orientation θ, we denote the highest reflex θ-maximum of $\mathbf{P}$ by v_θ.

The first observation we make is that $S(\theta)$ is the is the intersection of $h^+(v_\theta,\theta)$ with $h^+(v_{\theta^{-1}},\theta^{-1})$ since the lowest reflex θ-minimum is, of course, the highest reflex θ^{-1}-maximum. Since we always have that θ^{-1} is in $\mathcal{O}$ if $\theta\in\mathcal{O}$, we only have to find the intersection of the θ-oriented halfplanes above the reflex θ-maxima, taken over all $\theta\in\mathcal{O}$. Therefore, if v_θ is the highest reflex θ-maximum of $\mathbf{P}$, for each $\theta\in\mathcal{O}$, we want to compute $\bigcap_{\theta\in\mathcal{O}} h^+(v_\theta,\theta)$.

The idea of the algorithm is to scan the boundary of $\mathbf{P}$ counterclockwise and, thus, successively process the vertices of $\mathbf{P}$. During the scan we keep a list $\mathcal{L}$ of the highest reflex θ-maxima we have encountered so far, for each $\theta\in\mathcal{O}$. When we process the next reflex vertex v, we check, for which orientations $\theta\in\mathcal{O}$, v is a reflex a θ-maximum. We replace those θ-maxima v_θ for which v is higher than v_θ (in the θ-coordinate system) by v. The problem is that there may be $\Omega(|\mathcal{O}|)$ of orientations θ which have v_θ replaced by v. Hence, the algorithm may still take time $\Omega(n|\mathcal{O}|)$ if there are $\Omega(n)$ reflex vertices, each being a reflex θ-maximum for $\Omega(|\mathcal{O}|)$ orientations θ.

The important observation to reduce the running time of the algorithm is that we are only interested in the intersection of the halfplanes $h^+(v_\theta,\theta)$, with $\theta\in\mathcal{O}$. If we have that the so far encountered highest reflex θ-maximum v_θ is lower than v, for a subset $\mathcal{O}'$ of orientations θ in $\mathcal{O}$, the following lemma shows that there exist two orientations θ_1 and θ_2 such that $h^+(v,\theta_1)\cap h^+(v,\theta_2)$ already equals $\bigcap_{\theta\in\mathcal{O}'} h^+(v,\theta)$. If we can find θ_1 and θ_2 in logarithmic time, we only have to replace v_{θ_1} and v_{θ_2} in $\mathcal{L}$ since only these may influence the final intersection. If $\mathcal{L}$ is stored as a dictionary, then replacing v_{θ_1} and v_{θ_2} can be achieved in time $O(\log|\mathcal{O}|)$. Hence, it only has to be shown that θ_1 and θ_2 exist and can be found in time $O(\log|\mathcal{O}|)$.

Lemma 6.4 *If $\mathbf{P}$ is a polygon, $\mathcal{O}$ a finite set of orientations, and v a reflex vertex of $\mathbf{P}$,*

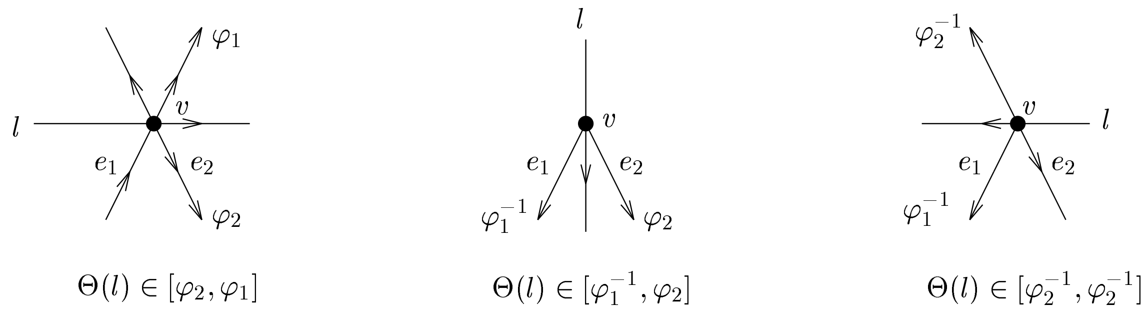

Figure 26: *v is a reflex θ-maximum only for the orientations $\theta \in [\varphi_2, \varphi_1]$.*

then there are two orientations θ_1 and θ_2 in $\mathcal{O}$ such that

$$h^+(v, \theta_1) \cap h^+(v, \theta_2) = \bigcap_{\theta} h^+(v, \theta)$$

if the intersection on the right side is taken over all $\theta \in \mathcal{O}$ for which v is a reflex θ-maximum. Furthermore, θ_1 and θ_2 can be found in time $O(\log |\mathcal{O}|)$.

Proof: Let e_1 be the edge that is (counterclockwise) before v and e_2 the edge after v. To see the claim we denote the orientation of e_i by φ_i. Note that v is a reflex θ-maximum only for the orientations θ that are in the interval $[\varphi_2, \varphi_1]$.[11] This can be shown as follows (see Figure 26). If $\theta \in [\varphi_1, \varphi_2^{-1}]$ or $\theta \in [\varphi_1^{-1}, \varphi_2]$, then the start point of e_1 is on the other side of the line $l(v, \theta)$ through v with orientation θ than the end point of e_2. Hence, v is no θ-extremum. If θ is in $[\varphi_2^{-1}, \varphi_1^{-1}]$, then e_1 and e_2 are above $l(v, \theta)$ and v is no θ-maximum. Only if $\theta \in (\varphi_2, \varphi_1)$, then both of e_1 and e_2 are completely below $l(v, \theta)$. There is a little problem if φ_2 is in $\mathcal{O}$, since then it depends on the edge after e_2 if v is a φ_2-maximum (refer to Figure 27). To see this suppose that $\varphi_2 \in \mathcal{O}$ and let e_3 be the edge after e_2. If $\Theta(e_3) < \varphi_2$, then e_2 is a φ_2-maximum since $\varphi_2 < \varphi_1$ and if $\Theta(e_3) > \varphi_2$, then e_2 merely contains some φ_2-saddlepoints. Since we only have to spend a constant amount of time to find the edge after e_2, we will not consider this special case anymore. Clearly, it can be checked in logarithmic time if φ_2 is in $\mathcal{O}$ if $\mathcal{O}$ is sorted.

So let θ_1 be the last orientation in $\mathcal{O}$ that is (counterclockwise) before φ_1 and θ_2 be the first orientation in $\mathcal{O}$ after φ_2 if they exist (see Figure 27). We denote the interval (φ_2, φ_1) by I. If there is one θ in $\mathcal{O}$ with $\theta \in I$, then, obviously, we have that θ_1 and θ_2 both exist. Hence, if they are not defined, then there is no θ in $\mathcal{O}$ for which v is a θ-maximum and we are done. So assume that $\theta_1, \theta_2 \in I$ and that they are different (if $\theta_1 = \theta_2$, then there is only one element in $\mathcal{O} \cap I$ and the lemma is trivially true.)

Consider the wedge w that is formed by $h^+(v, \theta_1) \cap h^+(v, \theta_2)$. It can easily be seen that w is contained in any halfplane $h^+(v, \theta)$ if $\theta \in [\theta_2, \theta_1]$. Since we have that $\mathcal{O} \cap I = \mathcal{O} \cap [\theta_2, \theta_1]$

[11] Since v is reflex, we have that $\Theta(e_1) > \Theta(e_2)$.

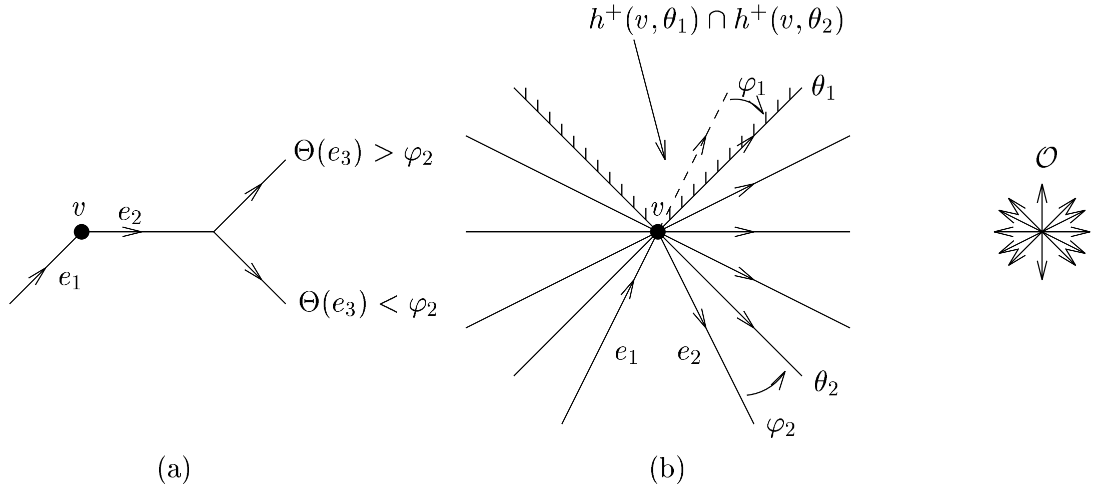

Figure 27: *How to find θ_1 and θ_2.*

by our choice of θ_1 and θ_2 and the orientations in $\mathcal{O} \cap I$ are the only ones for which v is a θ-maximum, we have that

$$h^+(v,\theta_1) \cap h^+(v,\theta_2) = \bigcap_{\theta \in \mathcal{O} \cap [\theta_1,\theta_2]} h^+(v,\theta) = \bigcap_{\theta \in \mathcal{O} \cap I} h^+(v,\theta).$$

Hence, it suffices to just look at v_{θ_1} and v_{θ_2} and possibly replace them with v. Note that if v_{θ_1} is higher than v in the θ_1-coordinate system, i.e., if we *do not* replace v_{θ_1} with v, then $h^+(v_{\theta_1},\theta_1) \subseteq h^+(v,\theta_1)$, and, hence, $h^+(v_{\theta_1},\theta_1) \cap h^+(v,\theta_2)$ is still contained in each $h^+(v,\theta)$, with $\theta \in [\theta_2,\theta_1]$.

If $\mathcal{O}$ is given as a sorted array, it is easy to see that θ_1 and θ_2 can be found in time $O(\log|\mathcal{O}|)$ given φ_1 and φ_2. $\square$

So the scanning algorithm computes a number of vertices v_θ such that

$$\bigcap_{\theta \in \mathcal{O}} h^+(v_\theta,\theta) \cap \mathbf{P}$$

yields the $\mathcal{O}^s$-kernel of $\mathbf{P}$. Above we have shown that we can compute the vertices v_θ in time $O(n\log|\mathcal{O}|)$. Observe that there are at most $|\mathcal{O}|$ such vertices, one for each orientation $\theta \in \mathcal{O}$. Furthermore, note that the above computation yields the vertices v_θ sorted according to θ. In Appendix 7 we show how to compute the intersection of halfplanes that are sorted according to slope in time linear in the number of halfspaces. Hence, the intersection $\bigcap_{\theta \in \mathcal{O}} h^+(v_\theta,\theta)$ can be computed in time $O(|\mathcal{O}|)$. This proves the following theorem.

Theorem 6.5 *The $\mathcal{O}^s$-kernel of a simple polygon with n vertices can be computed in time $O(n\log|\mathcal{O}| + |\mathcal{O}|)$, for finite $\mathcal{O}$, given $O(|\mathcal{O}|\log|\mathcal{O}|)$ preprocessing time to sort $\mathcal{O}$.*

There are two questions in connection with the above result.

(i) Is it possible to compute the $\mathcal{O}^s$-kernel of polygon within the same time bounds if $\mathcal{O}$ consists of a finite number of intervals, i.e., the algorithm should have a running time of $O(n \log r)$ if $\mathcal{O}$ consists of r intervals?

(ii) Is there an algorithm that computes the $\mathcal{O}^s$-kernel faster or can a lower bound be proven?

As to the first question we conjecture that this is possible using an approach that is similar to the kernel algorithm of Lee and Preparata [LP79], i.e., processing the vertices in a scan as in our algorithm and keeping track of a suitable representation of the $\mathcal{O}^s$-kernel.

For the second question the situation is less clear. On the one hand, a lower bound of $\Omega(n \log |\mathcal{O}|)$ on finding all vertices which are a θ-maximum, for some θ in $\mathcal{O}$, is easy to show (see Appendix 8); on the other hand, there is a kernel algorithm by Cole and Goodrich [CG88] which could possibly be adapted to yield an $O(n + |\mathcal{O}|)$ time algorithm for the computation of $\mathcal{O}^s\text{-}kernel(\mathbf{P})$.

6.3 The $\mathcal{O}^s$-Kernel in the General Setting

If we allow a polygon to have holes, the situation changes considerably. The $\mathcal{O}^s\text{-}kernel(\mathbf{P})$ is not necessarily connected or $\mathcal{O}$-convex anymore as shown in Figure 28. Furthermore, we lose the Intersection Lemma as Figure 29 illustrates. Since the algorithm to compute the $\mathcal{O}^s$-kernel of a simple polygon $\mathbf{P}$ and finite $\mathcal{O}$ is based on Corollary 5.12 for which there is no analog if $\mathbf{P}$ is multiply connected, we need a different approach to be able to compute the $\mathcal{O}^s$-kernel for multiply connected polygons. Holes can be viewed as polygons themselves. But instead of being able to see their inside we are now concerned with their external visibility. This gives rise to the definition of an *external kernel* of a polygon.

Definition 6.1 *Let* $\mathbf{P}$ *be a polygon in the plane and* $\mathcal{O}$ *some set of orientations. We define the* external kernel of $\mathbf{P}$ *as the set of all points that belong to* $\mathcal{O}^s\text{-}kernel(\mathbb{E}^2 \setminus int(\mathbf{P}))$ *and denote it by* $\mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P})$.

It is not very surprising that the $\mathcal{O}^s$-kernel of a multiply connected polygon $\mathbf{P}$ is the intersection of the $\mathcal{O}^s$-kernel of the enclosing polygon of $\mathbf{P}$ intersected with the external kernels of the holes of $\mathbf{P}$. This is shown in the following lemma.

Lemma 6.6 *If* $\mathcal{O}$ *is a set of orientations and* $\mathbf{P}$ *a multiply connected simple polygon with enclosing polygon* $\mathbf{Q}$ *and holes* $\mathbf{H}_1, \ldots, \mathbf{H}_m$*, then*

$$\mathcal{O}^s\text{-}kernel(\mathbf{P}) = \mathcal{O}^s\text{-}kernel(\mathbf{Q}) \cap \bigcap_{i=1}^{m} \mathcal{O}^s\text{-}kernel_{ext}(\mathbf{H}_i).$$

Proof: We denote the intersection of $\mathcal{O}^s\text{-}kernel(\mathbf{Q})$ with the external kernels of the holes $\mathbf{H}_i$, $1 \leq i \leq k$, by $\mathbf{K}$. It can be easily seen that $\mathcal{O}^s\text{-}kernel(\mathbf{P}) \subseteq \mathbf{K}$ so we concentrate on

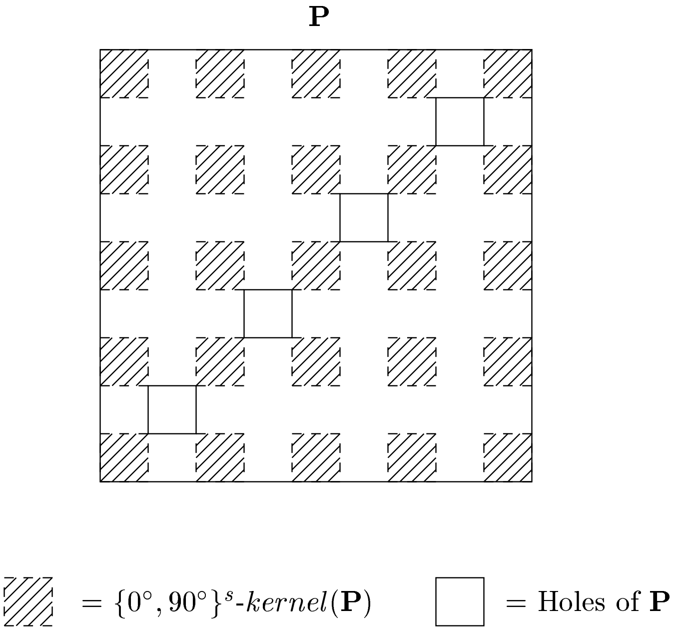

Figure 28: *The $\mathcal{O}^s$-kernel of* $\mathbf{P}$ *is disconnected and not $\mathcal{O}$-convex.*

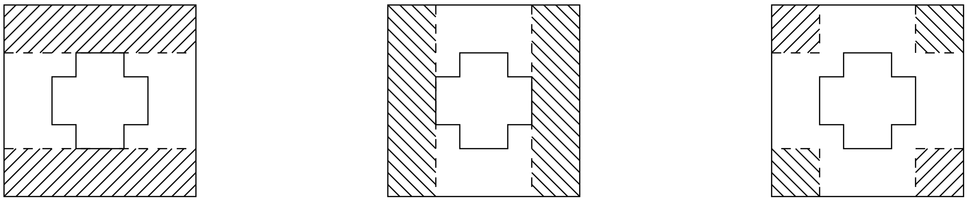

Figure 29: *The $\{0^\circ, 90^\circ\}^s$-kernel is not the intersection of the $\{0^\circ\}^s$-kernel and the $\{90^\circ\}^s$-kernel.*

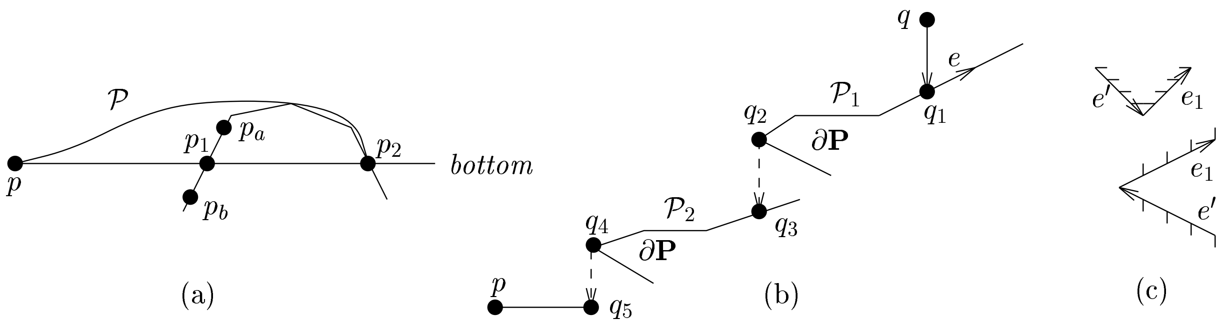

Figure 30: *The construction of an $\mathcal{O}$-convex path from q to p.*

proving the reverse inclusion. For this, let p be a point in $\mathbf{K}$ and q a point in the polygon $\mathbf{P}$. We have to show that there is an $\mathcal{O}$-convex path from p to q. If $\Theta(\overline{pq}) \in \mathcal{O}$, then neither the holes nor $\mathbf{Q}$ intersect $\overline{pq}$ and, hence, p $\mathcal{O}$-sees q in $\mathbf{P}$. Therefore, suppose that $\Theta(\overline{pq}) \notin \mathcal{O}$.

We first treat the case that $|\mathcal{O}| \geq 2$. So, w.l.o.g. we assume that the maximal $\mathcal{O}$-free range of p and q is $(0°, 90°)$. Hence, we have to construct a $\{0°, 90°\}$-convex stairsegment in $\mathbf{P}$ that connects p and q.

For brevity, we call the four line segments making up $\{0°, 90°\}$-$\|[p, q]$ *top, bottom, left,* and *right* with the obvious meaning. If some part of the boundary of $\mathbf{P}$ enters $\{0°, 90°\}$-$\|[p, q]$ through the lower arm of $\{0°, 90°\}$-$\|[p, q]$, it has to leave through it again since otherwise q would be cut off from p and p cannot belong to either the kernel of $\mathbf{Q}$ or the external kernel of some hole $\mathbf{H}_i$.

The first observation we make is that $\partial\mathbf{P}$ crosses *bottom* and *left* at most once. For suppose that there is a part of the boundary of $\mathbf{P}$ that crosses, say *bottom*, twice. Starting from p let p_1 resp. p_2 be two points that belong to the first resp. second component of $\partial\mathbf{P} \cap$ *bottom* as illustrated in Figure 30a. Note that we can assume that p_1 and p_2 both belong to either the boundary of the same hole or to $\partial\mathbf{Q}$ for the components of $ext(\mathbf{P})$ just correspond to $ext(\mathbf{Q})$ or $int(\mathbf{H}_i)$. W.l.o.g. suppose that p_1 and p_2 belong to $\partial\mathbf{H}_i$. Since $\partial\mathbf{H}_i$ crosses *bottom* in p_1, there are points p_a and $p_b \in \partial\mathbf{H}_i$ locally above and below p_1. Consider a path $\mathcal{P}$ from p to p_2. $\mathcal{P}$ passes either above p_a or below p_b. Since its start and end point lie on *bottom*, there is a line parallel to *bottom* that intersects $\mathcal{P}$ twice; hence, $\mathcal{P}$ is not $\{0°, 90°\}$-convex which contradicts the fact that p belongs to $\mathcal{O}^s$-$kernel_{ext}(\mathbf{H}_i)$.

We now construct a $\{0°, 90°\}$-convex path $\mathcal{S}$ from q to p. The intuitive idea is to follow a path like water would flow from q to p. Let q_1 be the first cross point of *right* and $\partial\mathbf{P}$ (as seen from q). If q_1 does not exist, we reach *bottom* and continue horizontally to p. Since there is at most one cross point of $\partial\mathbf{P}$ with *bottom*, this is always possible. So suppose q_1 exists and belongs to edge e (see Figure 30b). Since the interior of $\mathbf{P}$ is locally to the left of e, e is oriented from left to right. The slope of e is positive, i.e., $\Theta(e) \in (0°, 90°)$ since otherwise p does not $\{0°, 90°\}^s$-see q_1. Note that this is true for any edge of $\partial\mathbf{P}$ in

$\{0°, 90°\}$-$\|[p, q]$ that is oriented from left to right, i.e., which has the interior of $\mathbf{P}$ locally above it. We continue *downward*, i.e., contrary to the orientation of e, along $\partial\mathbf{P}$ until we reach a $\{90°\}$-reflex, locally leftmost vertex q_2. The first observation we make is that the part $\mathcal{P}_1$ of $\partial\mathbf{P}$ from q_1 to q_2 consists only of edges that oriented from left to right and that have an orientation in $[0°, 90°]$. Let e_1 be the last edge of $\mathcal{P}_1$ with $\Theta(e_1) \in [0°, 90°]$. Note that there is at least one such edge since $\Theta(e) \in [0°, 90°]$. Let e' be the edge before e_1 on $\partial\mathbf{P}$. If $\Theta(e') \in (180° + \Theta(e_1), 360°)$, then e' cannot be $\{0°, 90°\}$-seen by p and if $\Theta(e') \in (90°, 180° + \Theta(e_1))$ then the common vertex of e' and e_1 is a $\{90°\}$-reflex, locally leftmost vertex (see Figure 30c). Lemma 5.3 yields that $\mathcal{P}_1$ is $\{0°, 90°\}$-convex. Secondly, we note that $\mathcal{P}_1$ does not cross *left*. For assume the contrary. Since the interior of $\mathbf{P}$ is locally above $\mathcal{P}_1$ and p is below the cross point of $\mathcal{P}_1$ with *left*, there is a second intersection point of $\partial\mathbf{P}$ with *left* which contradicts our above observation.

From q_2 we again drop vertically until we reach another part of $\partial\mathbf{P}$ in q_3 where q_3 is the first cross point of the vertical ray downward through q_2 and $\partial\mathbf{P}$ in $\{0°, 90°\}$-$\|[p, q]$. We continue in this way until we finally reach *bottom*. Note that we finally have to arrive either at *bottom* or at *left* since we always move down or to the left and $\partial\mathbf{P}$ consists only of a finite number of edges. From there on we complete the path horizontally or vertically to p. This is always possible since $\partial\mathbf{P}$ crosses *bottom* and *left* only once.

With the same arguments as above we can show that the parts $\mathcal{P}_i$ of $\partial\mathbf{P}$ from q_{2i-1} to q_{2i}, for $i = 1, \ldots, m$—if we have repeated the procedure m times—are $\{0°, 90°\}$-stairsegments. By Lemma 5.1 we have that the concatenation of the $\mathcal{P}_i$ with the vertical line segments connecting them is an $\mathcal{O}$-stairsegment $\mathcal{S}$. By construction $\mathcal{S}$ is contained in $\mathbf{P}$. Hence, p $\mathcal{O}^s$-sees q.

If $|\mathcal{O}| = 1$, say $\mathcal{O} = \{90°\}$, a similar approach works. To see this let e_i be again the last edge of $\mathcal{P}_i$ with orientation in $[-90, +90]$ and e' the edge before e_1 ($\partial\mathbf{P}$). If $\Theta(e') \in (180° + \Theta(e_1), 270°)$, then e' is not $\{0°\}$-seen by p and if $\Theta(e') \in (90°, 180° + \Theta(e_1))$, then $e_1 \cap e'$ is a leftmost vertex. $\square$

The above lemma provides the tools to compute the $\mathcal{O}^s$-kernel of a multiply connected polygon. Since we know how to compute the $\mathcal{O}^s$-kernel of the enclosing polygon, we only have to show how to compute the external $\mathcal{O}^s$-kernel of a polygon which is done in the following section.

6.4 The External Kernel of a Polygon

We will now look at the problem of computing the external $\mathcal{O}^s$-kernel of a simple polygon. Recall that Figure 29 implies that the external kernel cannot be reduced to the calculation of the kernel for one direction. So we have to develop an algorithm that directly computes $\mathcal{O}^s$-$kernel_{ext}(\mathbf{P})$. This time we allow the set of orientations $\mathcal{O}$ to consist of a finite number r of closed intervals (or ranges) $[\alpha_1, \beta_1], \ldots, [\alpha_r, \beta_r]$. We call $\mathcal{O}$ *interval finite* in this case. Recall that an oriented *tangent* t of $\mathbf{P}$ is an oriented line such that t intersects $\partial\mathbf{P}$ and that $\mathbf{P}$ is completely contained in $h^+(t)$. Obviously, there is only one tangent to $\mathbf{P}$, for any given $\theta \in [0°, 360°)$. We denote it by $t(\theta)$. The halfplane to the left of $t(\theta)$ is denoted by

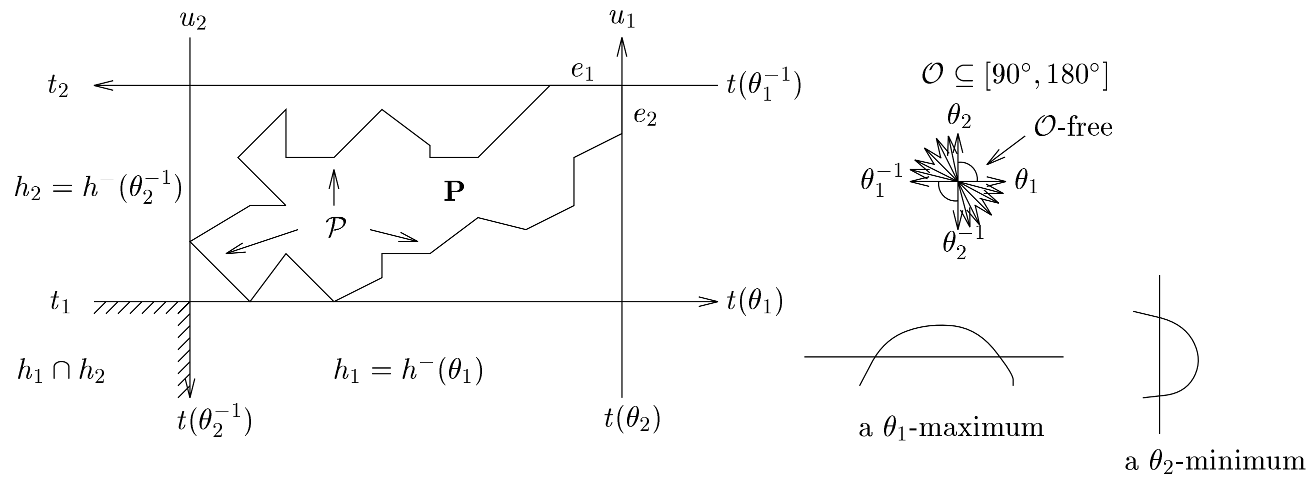

Figure 31: *The external kernel of a polygon*

$h^+(\theta)$ and, similarly, the halfplane to the right by $h^-(\theta)$.

We start with a simple observation (see also Lemma 6.4.2 of [Raw87]).

Observation 6.7 No point which lies between a pair of tangents to $\mathbf{P}$ that are parallel to some $\theta \in \mathcal{O}$ belongs to the external $\mathcal{O}^s$-kernel of $\mathbf{P}$.

This is an immediate consequence of Lemma 4.9 since $\mathcal{O}^s$-visibility implies $\mathcal{O}$-visibility. We now turn to characterizing the components of the external kernel of a simple polygon $\mathbf{P}$.

Lemma 6.8 *Let* $\mathbf{P}$ *be a polygon in the plane and* $\mathcal{O}$ *be a set of orientations with* $|\mathcal{O}| \geq 2$. *If* θ_1 *and* θ_2 *are two adjacent*[12] *orientations in* $\mathcal{O}$, *with* $\theta_1 \leq \theta_2$, *such that*

1. $e_1 = t(\theta_1^{-1}) \cap \mathbf{P}$ *and* $e_2 = t(\theta_2) \cap \mathbf{P}$ *are two (maybe degenerate) edges of* $\mathbf{P}$ *that meet in one vertex of* $\mathbf{P}$, *and*

2. $\mathcal{P} = \partial\mathbf{P} \setminus (e_1 \cup e_2)$ *contains neither a local convex* θ_1*-maximum nor a local convex* θ_2*-minimum,*[13]

then $h^-(\theta_1) \cap h^-(\theta_2^{-1})$ *belongs to* $\mathcal{O}^s$*-kernel*$_{ext}(\mathbf{P})$.

Proof: W.l.o.g. we assume that $\{\theta_1, \theta_2\} = \{0°, 90°\}$. The situation is illustrated in Figure 31. Let h_1 be the halfplane $h^-(\theta_1) = h^-(0°)$ and h_2 the halfplane $h^-(\theta_1^{-1}) = h^-(270°)$. For brevity, we denote $t(0°)$ by t_1, $t(180°)$ by t_2, $t(90°)$ by u_1, and $t(270°)$ by u_2. We have that t_1 is below t_2 and u_1 is to the right of u_2 (in the plane, not w.r.t. to their orientations). $h_1 \cap h_2$ is below and to the left of $\mathbf{P}$. First suppose there is some point $p \in h_1 \cap h_2$ that sees $\mathbb{E}^2 \setminus int(\mathbf{P})$. We have to show that Conditions 1 and 2 hold.

[12] Note that adjacency, in particular, implies that (θ_1, θ_2) is $\mathcal{O}$-free.

[13] Refer to page 16 for the definition of (local) θ-extremum.

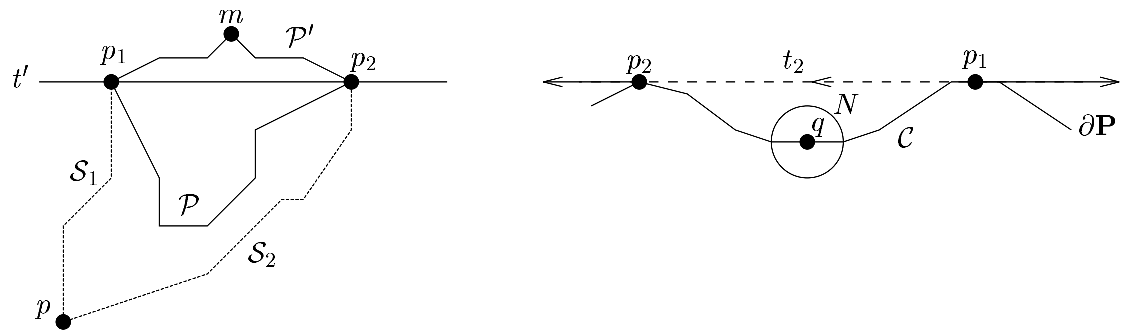

Figure 32: *If $p \in \{0°\}^s$-kernel$(\mathbf{P})$, then $\mathcal{P}$ contains no $\{0°\}$-maximum and $t(\theta_1^{-1}) \cap \partial\mathbf{P}$ is connected.*

1. Assume there is a convex $\{0°\}$-maximum m in $\mathcal{P}$. (The argument is analogous for a convex $\{90°\}$-minimum.) Hence, we can choose some line t' parallel to t_1 such that t' intersects $\mathcal{P}$ slightly below m in two points p_1 and p_2, p_1 to the left of p_2, and such that the part $\mathcal{P}'$ of $\mathcal{P}$ from p_1 to p_2 contains m as its only $\{0°\}$-maximum (see Figure 32). Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be two $\mathcal{O}$-stairsegments from p to p_1 and p_2, respectively which exist by our assumption. Since $\mathcal{S}_1$ and $\mathcal{S}_2$ are $\{0°, 90°\}$-convex, they do not cross t'. By our choice of t' the concatenation of $\mathcal{P}'$ with $\mathcal{S}_1$ and $\mathcal{S}_2$ forms a closed curve whose interior contains $\mathbf{P}$. This holds since $\mathcal{S}_1$ and $\mathcal{S}_2$ are completely contained in the exterior of $\mathbf{P}$ and $\partial\mathbf{P}$ is simple. Furthermore, the only part above t' is $\mathcal{P}'$. Hence, m is a global $\{0°\}$-maximum of $\partial\mathbf{P}$ and is intersected by t_2 which contradicts the assumption that m belongs to $\mathcal{P}$.

2. We now show that $t_2 \cap \partial\mathbf{P}$ is connected and contains $t_2 \cap u_1$ (The proof that $u_1 \cap \partial\mathbf{P}$ is connected is, of course, analogous.) So let p_1 and p_2 be two points in $t_2 \cap \partial\mathbf{P}$ such that p_1 occurs before p_2 on t_2. We will show that p_1 and p_2 belong to the same connected component of $t_2 \cap \partial\mathbf{P}$. For assume the contrary. Then, there is a point q in the part of $\partial\mathbf{P}$ from p_1 to p_2 that is below t_2. We consider the curve $\mathcal{C}$ consisting of the part of t_2 to the right of p_1 and the part of $\partial\mathbf{P}$ from p_1 to p_2 and the part of t_2 to the left of p_2. $\mathcal{C}$ clearly is a biinfinite curve. W.l.o.g. we assume that $\mathcal{C}$ is simple and, hence, divides the plane into two components one of which contains p and the other of which contains the points exterior to $\mathbf{P}$ in a sufficiently small neighbourhood N of q. Therefore, any path $\mathcal{P}$ that connects p to q and that is contained in the exterior of $\mathbf{P}$ intersects t_2 twice. Since p and q are below t_2, this implies that $\mathcal{P}$ is not $\{0°\}$-convex and p is not in $\mathcal{O}^s$-$kernel_{ext}(\mathbf{P})$.

   That $u_1 \cap t_2$ belongs to $\partial\mathbf{P}$ can be seen by employing an argument similar to the one above only that now p_1 is chosen in $u_1 \cap \partial\mathbf{P}$ and p_2 in $t_2 \cap \partial\mathbf{P}$. We then consider the biinfinite curve consisting of the part of u_1 below p_1, the part of $\partial\mathbf{P}$ from p_1 to p_2

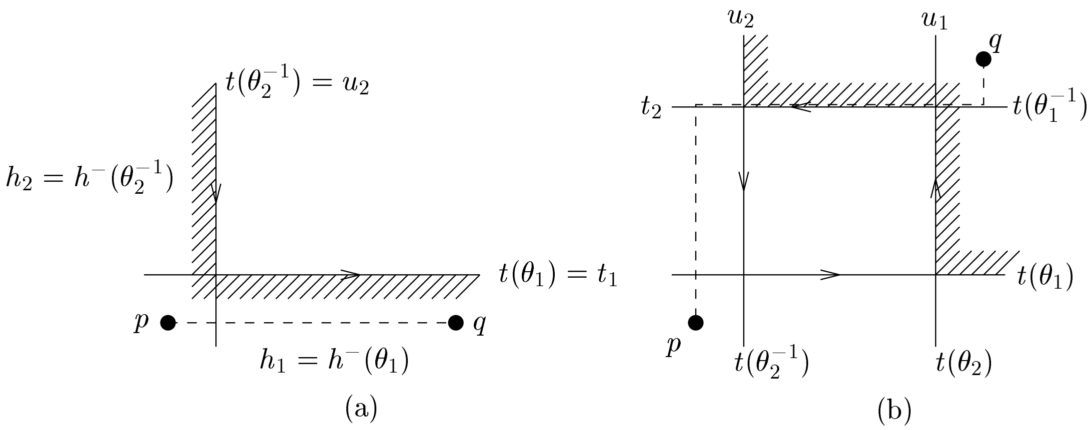

Figure 33: *q is in $h_1 \cup h_2$ or in $h^-(\theta_1^{-1}) \cup h^-(\theta_2)$.*

and the part of t_2 to the left of p_1.

This concludes the only if direction of the claim. Now suppose that Conditions 1 and 2 hold. We will show that then every point in $h_1 \cap h_2$ can see $\mathbb{E}^2 \setminus int(\mathbf{P})$. To this end let $p \in h_1 \cap h_2$ and $q \in \mathbb{E}^2 \setminus int(\mathbf{P})$. We distinguish several cases.

Case 1 q is in $h_1 \cup h_2$ (see Figure 33a).
Since p and q are on the same side of t_1 or u_1 and there are no obstacles in these half-planes, we have $\overline{pq} \subseteq \mathbb{E}^2 \setminus int(\mathbf{P})$.

We have treated the case that q is to the left of or below p. Hence, from now on we can assume that q is to the right and above p which implies that the maximal $\mathcal{O}$-free range of $\Theta(\overline{pq})$ is $(0°, 90°)$. Therefore, it suffices to construct a $\{0°, 90°\}$-convex stairsegment from p to q.

Case 2 q is in $h^-(\theta_1^{-1}) \cup h^-(\theta_2)$ and q is to the right and above p (see Figure 33b).
We just consider the case that q is in $h^-(\theta_1^{-1})$. Starting from q vertically downward there is no obstruction until we reach t_2. We follow t_2 to the left till we hit the x-coordinate of p. This point on t_2 must be to the left of u_2 (since p is to the left of u_2) and, hence we can connect it to p by another vertical segment. Obviously, the vertical segments do not share common y-coordinates (apart from the one point on t_2) and, hence, the path constructed is $\{0°, 90°\}$-convex. Note that this case also covers $q \in e_1 \cup e_2$ as a degeneracy since, then, the first vertical segment is unnecessary.

Case 3 q is in $\mathbf{R} = h^+(\theta_1) \cap h^+(\theta_2) \cap h^+(\theta_1^{-1}) \cap h^+(\theta_2^{-1})$ but not in the interior of $\mathbf{P}$.
Consider the vertical ray r downward from q. We assume that r intersects $\partial\mathbf{P}$ and denote the first edge that r crosses by e_1. We claim that e_1 is tilted downward (if we follow the counterclockwise orientation of e_1), i.e., $\Theta(e_1) \in [180°, 270°)$, as shown in Figure 34. To

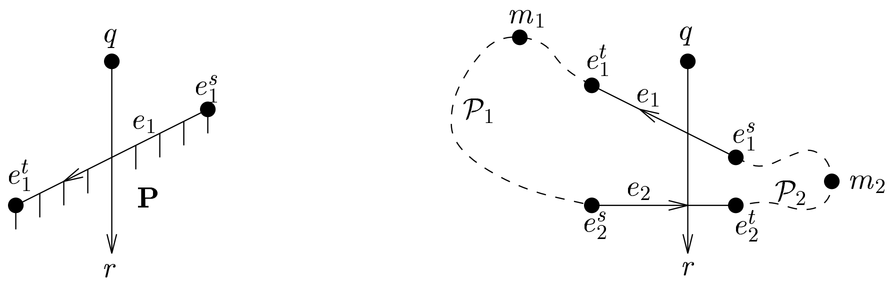

Figure 34: *q is in the bounding box of* $\mathbf{P}$.

see this suppose that the contrary holds and e_1 is tilted upward. Since q is on the outside of $\mathbf{P}$, it is above the line through e_1, i.e., the source vertex e_1^s of e_1 is to the right of q and the target vertex e_1^t is to the left of q. Furthermore, r crosses $\partial\mathbf{P}$ in a second edge e_2 with e_2^s to the left of r and e_2^t to the right of r. Let $\mathcal{P}_1$ be the part of $\partial\mathbf{P}$ from from e_1^s to e_2^s and $\mathcal{P}_2$ the part of $\partial\mathbf{P}$ from from e_2^s to e_1^s.

Since e_1^t is above e_1^s, there is a convex $\{0°\}$-maximum m_1 on $\mathcal{P}_1$. Also since e_2^t is to the right of e_2^s there is a rightmost point m_2 on $\mathcal{P}_2$. m_2 is a convex $\{90°\}$-minimum. By Condition 2 we have that m_1 belongs to $t_2 \cap \partial\mathbf{P}$ and m_2 belongs to $u_1 \cap \partial\mathbf{P}$. Condition 1 implies that the part of $\partial\mathbf{P}$ from m_2 to m_1 consists only of two edges, one having orientation 90° and the other having orientation 180° which contradicts the fact that $\Theta(e_1) \in (90°, 180°)$.

We are now in a position to construct a $\{0°, 90°\}$-convex stairsegment from p to q in very much the same fashion as in Lemma 6.6. The idea is again to follow a path like water would flow from p to q. Any edge of $\partial\mathbf{P}$ we encounter is tilted downward to the left (to see this just choose some point immediately above the edge and the previous argument applies). If we reach some locally leftmost edge of $\partial\mathbf{P}$, we continue vertically downward until we either intersect $\partial\mathbf{P}$ or t_1. Proceeding in this way we will reach t_1 in a point to the right of or on u_2 since the number of edges of $\mathbf{P}$ is finite. Since p is to the left of u_2 and below t_1, we clearly can continue the $\{0°, 90°\}$-convex path constructed so far to p.

□

With Observation 6.7 and the above lemma we are able to decide if a point p that is contained in a slab between two parallel $\mathcal{O}$-oriented lines or in a wedge enclosed by two tangents with adjacent orientations belongs to $\mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P})$ or not. Are there any other points in $\mathbb{E}^2 \setminus int(\mathbf{P})$ that do not fall into one of these two categories? The answer is no. To see this note that there are $2r$ $\mathcal{O}$-Oriented tangents to $\mathbf{P}$ (see Figure 35). The arrangement of the $2r$ tangents divides the plane into $O(r^2)$ components. It is easy to show that a point p in the plane either lies between two parallel tangents or p lies within a wedge formed by two tangents with adjacent orientations. More precisely, if $\mathbf{P}$ is a simple polygon and $p \in ext(\mathbf{P})$, then we have that either p lies between a pair of $\mathcal{O}$-oriented parallel tangents to $\mathbf{P}$ or there are two tangents t_1 and t_2 to $\mathbf{P}$ whose orientations are adjacent in $\mathcal{O}$ such that

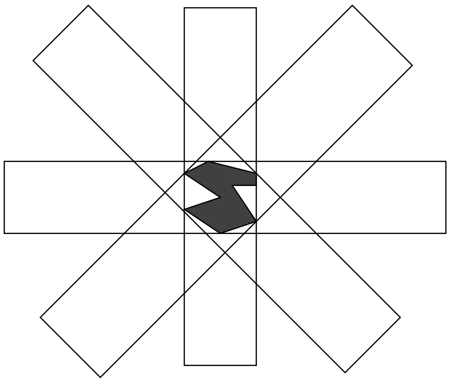

Figure 35: *The arrangement of* $2r$ *parallel tangents to* $\mathbf{P}$.

$h^-(t_1) \cap h^-(t_2)$ contains p. To see this we imagine the slabs $h^+(\theta) \cap h^+(\theta^{-1})$, with $\theta \in \mathcal{O}$, as being painted black (see Figure 35). Since these slabs have a common intersection—all of them contain $\mathbf{P}$ by definition—the arrangement of the slabs is equivalent to an arrangement of θ-oriented lines, with $\theta \in \mathcal{O}$, that intersect in one point (shrinking a slab does not create any new regions as long as the common intersection of all slabs remains simply connected). For the latter arrangement it is obvious that a point p either lies on a θ-oriented line or in a wedge that is formed by two lines whose orientations are adjacent in $\mathcal{O}$.

The above approach obviously only works if $|\mathcal{O}| \geq 2$. So it remains to look at $\mathcal{O}^s$-$kernel_{ext}(\mathbf{P})$ if $\mathcal{O} = \{\theta\}$. As expected the situation simplifies considerably and we get the following characterization. This result is contained on p. 79/80 of [Raw87] but not proven there.

Lemma 6.9 *If* $\mathbf{P}$ *is a simple polyon in the plane,* $\mathcal{O} = \{0°\}$, *and* p *is above the highest point of* $\mathbf{P}$, *then*

$$p \text{ belongs to } \mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P}) \text{ if and only if } \mathbf{P} \text{ contains no reflex maxima.}$$

Proof: First suppose that p belongs to $\mathcal{O}^s$-$kernel_{ext}(\mathbf{P})$ and assume that $\mathbf{P}$ has a reflex maximum m. Consider a horizontal line l slightly below m. Since m is a reflex maximum l intersects $\partial\mathbf{P}$ in at least two points p_1 and p_2 that belong to different components of $\partial\mathbf{P} \cap l$. W.l.o.g. we can assume that the part of $\partial\mathbf{P}$ from p_1 to p_2 does not intersect l in a third component. Suppose that p_1 is to the left of p_2. We choose a point q on $\partial\mathbf{P}$ above and to the right of p_1. The path which consists of l to the left of p_1 and to the right of p_2 and the part of the boundary of $\mathbf{P}$ between p_1 and p_2 divides the plane into two disconnected components one of which contains p and the other one contains all the points exterior to $\mathbf{P}$ in a sufficiently small neighbourhood N of q. Any path $\mathcal{P}$ in the exterior of $\mathbf{P}$ that connects p and q and that does not cross $\partial\mathbf{P}$ intersects l either to the left of p_1 or to the right of p_2. Since q is above l, $\mathcal{P}$ intersects l a second time and, hence, $\mathcal{P}$ is not $\{0°\}$-convex.

Now suppose we have no reflex maximum in $\partial\mathbf{P}$. Let q be some point in $\mathbb{E}^2 \setminus int(\mathbf{P})$. We have to show that p $\{0°\}$-sees q. If q is above the highest point or below the lowest point of $\mathbf{P}$, it is easy to construct a $\{0°\}$-convex stairsegment from p to q. So let q be between the two horizontal tangents to $\mathbf{P}$ and assume the horizontal ray to the right of q intersects $\mathbf{P}$. We follow the boundary of $\mathbf{P}$ upward until we reach a maximum m. Since this maximum is convex and q is to the left of $\partial\mathbf{P}$, the points immediately to the right of m belong to $\mathbb{E}^2 \setminus int(\mathbf{P})$ and, hence, we can continue our path along the horizontal ray to the right of m until we hit the boundary of $\mathbf{P}$ again or are to the right of the rightmost point of $\mathbf{P}$. Carrying on in this fashion leads us to a point either completely to the right of or completely above $\mathbf{P}$ from where we can easily complete the $\{0°\}$-convex path to p started in q. □

6.5 Computing the External Kernel

In order to develop an efficient algorithm to compute $\mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P})$ we need to bound the number of components that it can maximally consist of. The following lemma shows that the external $\mathcal{O}^s$-kernel of a polygon has only a constant number of components independent of the cardinality of $\mathcal{O}$ or the size of $\mathbf{P}$. In the subsequently presented algorithm we will make crucial use of this fact.

Lemma 6.10 *Let* $\mathbf{P}$ *be a simple polygon. The maximum number of components* $\mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P})$ *consists of is*

(i) 2 if $|\mathcal{O}| = 1$,

(ii) 4 if $|\mathcal{O}| = 2$,

(iii) 3 if $|\mathcal{O}| = 3$, and

(iv) 2 if $|\mathcal{O}| > 3$,

and these bounds are tight. Furthermore, if $|\mathcal{O}| = 3$ *and* $\mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P})$ *consists of three components, then* $\mathbf{P}$ *is a triangle.*

Proof: To see that the external kernel of a polygon can have as many components as claimed refer to Figure 36.

In the first case the claim is an immediate consequence of Lemma 6.9 since either all points in the halfplane above $t(180°)$ belong to $\mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P})$ or none of it. The same, of course, holds for the halfplane below $t(0°)$. The proof of the second claim is as easy to see since the two $\mathcal{O}$-oriented slabs that contain $\mathbf{P}$ divide the plane into at most four quaterplanes. By Lemma 6.8 a quaterplane belongs either completely to the external $\mathcal{O}^s$-kernel of $\mathbf{P}$ or not at all. Hence, $\mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P})$ consists of at most four components.

Note that Condition 1 of Lemma 6.8 implies that, for each component w of $\mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P})$, there is a vertex v of $conv(\mathbf{P})$ whose adjacent edges have an interior

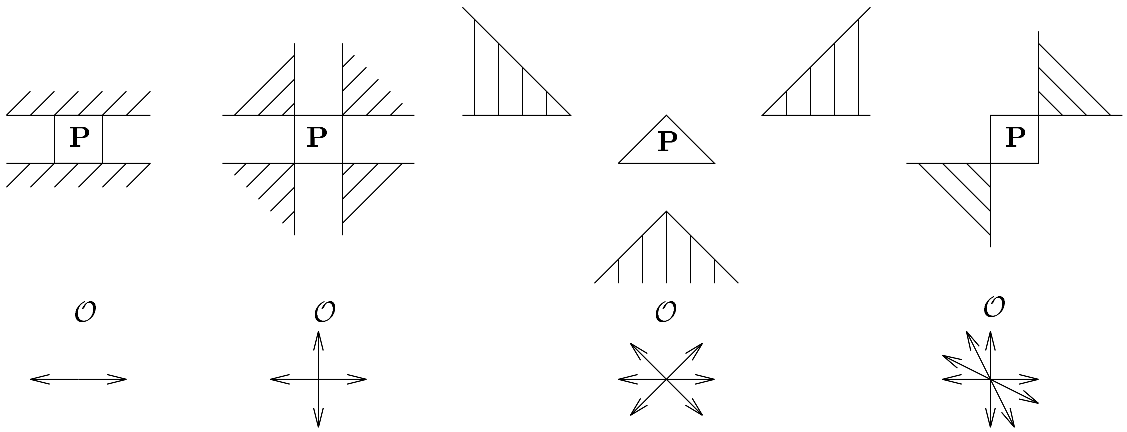

Figure 36: *Lower bounds on the number of components of the external kernel of a polygon.*

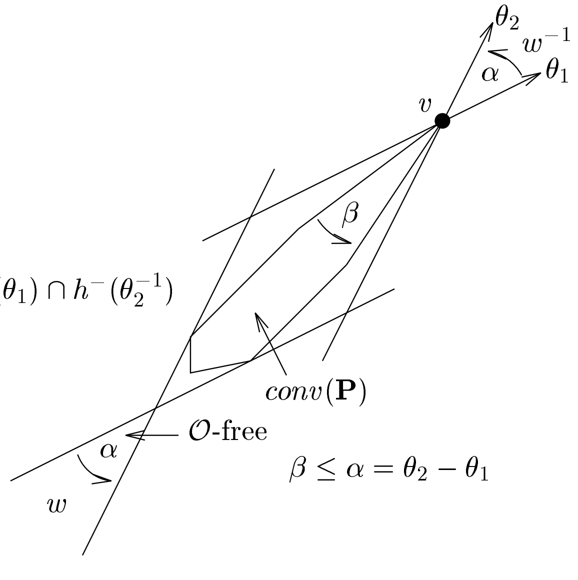

Figure 37: *There is a vertex in $conv(\mathbf{P})$ that has an interior angle of at most α.*

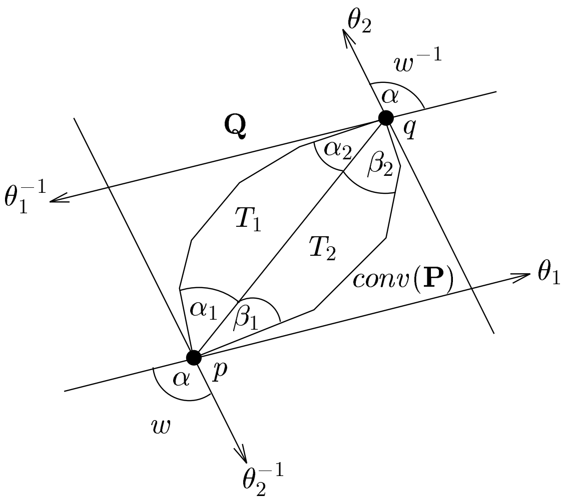

Figure 38: $\alpha_1 + \alpha_2 \leq \alpha$ *and* $\beta_1 + \beta_2 \leq \alpha$

angle which is less or equal to the $\mathcal{O}$-free range of the two tangents that enclose w (see Figure 37).

We now turn to proving the third claim. The proof is by contradiction. So suppose that $|\mathcal{O}| = 3$ and there are at least four components w_1, w_2, w_3, and w_4 of $\mathcal{O}^s$-$kernel_{ext}(\mathbf{P})$. Suppose the enclosing tangents of $w_1, \ldots, w_4$ span the $\mathcal{O}$-free ranges (α_1, β_1), (α_2, β_2), (α_3, β_3), and (α_4, β_4). As we noted above there are four vertices v_i of $conv(\mathbf{P})$, $1 \leq i \leq 4$, with an interior angle σ_i that is less or equal to $\beta_i - \alpha_i$. Since $|\mathcal{O}| = 3$, the closure of the four ranges cannot cover all of $[0°, 360°)$, i.e., $\bigcup_{i=1}^4 [\alpha_i, \beta_i] \neq [0°, 360°)$.

Now $conv(\mathbf{P})$ is a convex, say m-gon, whose sum of interior angles is $(m-2) \cdot 180°$. Since $\sum_{i=1}^4 \sigma_i \leq \sum_{i=1}^4 \beta_i - \alpha_i < 360°$ the remaining $(m-4)$ interior angles have to sum up to $(m-2) \cdot 180° - \sum_{i=1}^4 \sigma_i > (m-2) \cdot 180° - 360° = (m-4) \cdot 180°$ which contradicts the fact that all the interior angles of $conv(\mathbf{P})$ are less than $180°$. Hence, $\mathcal{O}^s$-$kernel_{ext}(\mathbf{P})$ has at most three components if $|\mathcal{O}| = 3$.

Before we treat the case $|\mathcal{O}| \geq 4$, we need a technical statement. We say a wedge w is a (θ_1, θ_2)-wedge if (θ_1, θ_2) is $\mathcal{O}$-free and w is the wedge between $h^-(\theta_1)$ and $h^-(\theta_2^{-1})$ (see again Figure 37). Let $\alpha = \theta_2 - \theta_1$. We claim that if a (θ_1, θ_2)-wedge w and its opposite, the $(\theta_1^{-1}, \theta_2^{-1})$-wedge w^{-1} belong to $\mathcal{O}^s$-$kernel_{ext}(\mathbf{P})$, then all the remaining interior angles of $conv(\mathbf{P})$ are greater than $180° - \alpha$.

To see this note that $conv(\mathbf{P})$ is contained in the parallelogram $\mathbf{Q} = h^+(\theta_1) \cap h^+(\theta_2) \cap h^+(\theta_1^{-1}) \cap h^+(\theta_2^{-1})$ whose sides that are parallel to θ_1 and θ_2 and that has the $(\theta_1, \theta_2^{-1})$-apex p and the $(\theta_1^{-1}, \theta_2)$-apex q in common with $conv(\mathbf{P})$ as shown in Figure 38. The line segment $\overline{pq}$ partitions $\mathbf{Q}$ into two triangles T_1 and T_2. Let α_1, α_2 be the angles of T_1 at $\overline{pq}$ and β_1, β_2 the angles of T_2 at $\overline{pq}$. Elementary geometry yields that $\alpha_1 + \alpha_2 \leq \alpha$ and $\beta_1 + \beta_2 \leq \alpha$. Hence, the remaining m_1 interior angles γ_i, $i = 1, \ldots, m_1$, of the part of

$conv(\mathbf{P})$ in T_1 satisfy the inequality

$$\sum_{i=1}^{m_1} \gamma_i + \alpha \geq m_1 \cdot 180^\circ$$

and the analog holds for the m_2 interior angles δ_i, $i = 1, \ldots, m_2$, of the part of $conv(\mathbf{P})$ in T_2. Since all γ_i and δ_i are less than 180°, we, in particular, have $\gamma_i > 180^\circ - \alpha$ and $\delta_i > 180^\circ - \alpha$.

Now if there are three components of $\mathcal{O}^s\text{-}kernel_{ext}(\mathbf{P})$, then we have three disjoint, $\mathcal{O}$-free ranges (α_i, β_i) such that there are three vertices of $conv(\mathbf{P})$ with interior angles $\sigma_i \leq \beta_i - \alpha_i$, $1 \leq i \leq 3$, by Lemma 6.8. If two of the ranges are opposites to each other, say $\alpha_2 = \alpha_1^{-1}$ and $\beta_2 = \beta_1^{-1}$, then $\sigma_3 \geq 180^\circ - \sigma_1$ by the above argument. This contradicts the fact that $[\alpha_1, \beta_1] \cup [\alpha_3, \beta_3]$ does not cover 180°. Hence, (α_1, β_1), (α_2, β_2), and (α_3, β_3) are not opposites of each other. Since $|\mathcal{O}| \geq 4$, we have that $\sigma_1 + \sigma_2 + \sigma_3 < 180^\circ$ and, hence, there are three vertices in $conv(\mathbf{P})$ whose sum of interior angles is less than 180°. If we denote the remaining interior angles of $conv(\mathbf{P})$ by γ_i, $i = 1, \ldots, m-3$, we have the following inqualities

$$\begin{aligned}(m-2) \cdot 180^\circ &= \sum_{i=1}^{3} \sigma_i + \sum_{i=1}^{m-3} \gamma_i \\ &< 180^\circ + \sum_{i=1}^{m-3} \gamma_i\end{aligned}$$

and, hence, $\sum_{i=1}^{m-3} \gamma_i > (m-3) \cdot 180^\circ$ which again contradicts the fact that $conv(\mathbf{P})$ is a convex polygon, i.e., that $\gamma_i \leq 180^\circ$.

If $|\mathcal{O}| = 3$ and there are three components, then we claim that $conv(\mathbf{P})$ is a triangle. To see this note that by the above argument all three ranges have to be different and, hence, add up to exactly 180°. This implies that the interior angle of any other vertex equals 180°, i.e., there are only trivial (collinear) other vertices. Clearly, $\mathbf{P}$ coincides with its convex hull in this case; otherwise it violates one of the conditions of Lemma 6.8. This proves the last claim of the lemma. □

To compute the external kernel we now proceed as follows. We assume that $\mathbf{P}$ is given as a list of n edges and that $\mathcal{O}$ consists of r ranges $[\alpha_1, \beta_1], \ldots, [\alpha_r, \beta_r]$ which are given as a sorted array. First we note that the vertices where the $\mathcal{O}$-tangents touch the boundary of $\mathbf{P}$ belong to the convex hull of $\mathbf{P}$. Hence, we compute the convex hull of $\mathbf{P}$ which can be done in linear time by [BG84, GY83, Lee83, MA79]. Suppose that $conv(\mathbf{P})$ consists of the (counterclockwise oriented) edges $e_1, \ldots, e_m$ and vertices $v_1, \ldots, v_m$ with $v_i \in e_i \cap e_{i+1}$ ($e_{m+1} = e_1$).

By Lemma 6.8 there is a vertex v_j in $conv(\mathbf{P})$, for each (θ_1, θ_2)-wedge, such that $\Theta(e_j)$ and $\Theta(e_{j+1})^{-1}$ are in $[\theta_1, \theta_2]$ (see Figure 39). Using the same argument as in the proof of Lemma 6.10 it can be seen that there are at most four vertices v_j on $conv(\mathbf{P})$ that satisfy $[\Theta(e_{j+1})^{-1}, \Theta(e_j)] \subseteq [\theta_1, \theta_2]$, for some $\mathcal{O}$-free range (θ_1, θ_2) and $\theta_1, \theta_2 \in \mathcal{O}$. We say v_j is a *candidate vertex of* (θ_1, θ_2) in this case.

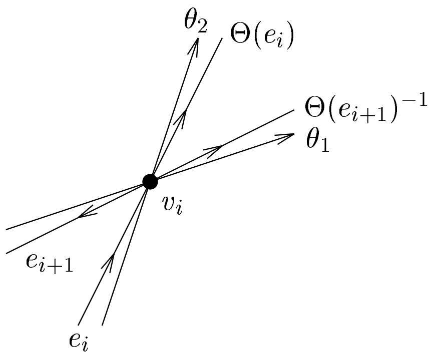

Figure 39: $[\Theta(e_{j+1})^{-1}, \Theta(e_i)]$ *is contained in* $[\theta_1, \theta_2]$.

The idea of the algorithm is to find these candidate vertices and then to check if the conditions of Lemma 6.8 are satisfied. Since we have to test for θ_1-maxima and θ_2-maxima at most four times and this can be done in time linear in the number of edges of $\mathbf{P}$, we need at most additional $O(n)$ steps once we have found the candidate vertices.

In the following we show how to compute the candidate vertices of $conv(\mathbf{P})$ in time $O(n + r)$. The idea is to step through the edges of $conv(\mathbf{P})$ and the ranges of $\mathcal{O}$ simultaneously. If we have a pointer to the ranges in $\mathcal{O}$ that keeps track of the orientation of the currently processed edge e_i, it can be checked in constant time whether $[\Theta(e_{j+1})^{-1}, \Theta(e_i)]$ is contained in a maximal $\mathcal{O}$-free range. We say a range $r = [\alpha, \beta]$ *corresponds to* orientation θ if (α, β) is $\mathcal{O}$-free and β is the least upper end point of a maximal $\mathcal{O}$-free range such that $\theta \leq \beta$. More precisely, the algorithm can now be described as follows.

Algorithm External $\mathcal{O}^s$-kernel
Let r_1 be the $\mathcal{O}$-free range that corresponds to $\Theta(e_1)$;
for each edge e_i of $\mathbf{P}$ **do**
  if $[\Theta(e_{i+1})^{-1}, \Theta(e_i)] \subseteq r_i$
    then output v_i as a candidate vertex for r_i
  end if
  if $\Theta(e_{i+1}) \leq \beta_i$
    then let r_{i+1} be r_i
    else step through $\mathcal{O}$ to the $\mathcal{O}$-free range r_{i+1} that corresponds
      to $\Theta(e_{i+1})$ $(*)$
  end if
end for;

Since we look at each range of $\mathcal{O}$ only once in Step $(*)$, the algorithm obviously needs time $O(n + r)$. In this way we compute the at most four candidate vertices of $conv(\mathbf{P})$ together with the $\mathcal{O}$-free ranges they belong to. As we already mentioned above, we need only additional linear time to check if the conditions of Lemma 6.8 are satisfied. This proofs the following theorem.

Theorem 6.11 *The external $\mathcal{O}^s$-kernel of a polygon with n vertices can be computed in time $O(n+r)$, for a set of orientations $\mathcal{O}$ consisting of r ranges, given $O(r \log r)$ preprocessing time to sort the ranges of $\mathcal{O}$.*

6.6 Computing the $\mathcal{O}^s$-Kernel of a Multiply Connected Polygon

With the above result, Lemma 6.6, and Theorem 6.5 we can compute the $\mathcal{O}^s$-kernel of a multiply connected polygon $\mathbf{P}$ with enclosing polygon $\mathbf{Q}$ and holes $\mathbf{H}_1, \ldots \mathbf{H}_m$ by computing $\mathcal{O}^s$-$kernel(\mathbf{Q})$ and intersecting it with $\mathcal{O}^s$-$kernel_{ext}(\mathbf{H}_i)$, for $1 \leq i \leq m$. We now want to obtain an estimate on the time needed for this procedure. Note that we have to restrict ourselves to finite $\mathcal{O}$. So suppose that $\mathbf{Q}$ consists of n_0 edges and each $\mathbf{H}_i$ of n_i edges. Let $n = n_0 + \sum_{i=1}^{m} n_i$ be the number of edges $\mathbf{P}$. The first step is to compute $\mathcal{O}^s$-$kernel(\mathbf{Q})$ and $\mathcal{O}^s$-$kernel_{ext}(\mathbf{H}_i)$, for $1 \leq i \leq m$, which requires time

$$O(n_0 \log |\mathcal{O}| + \sum_{i=1}^{m} n_i + m|\mathcal{O}|).$$

With the scan-line algorithm by Chazelle and Edelsbrunner [CE88] we can intersect m polygons with all together n edges in time $O(n \log n +$ (the number of intersections of the edges)). Since the external kernel consists only of at most eight edges and $\mathcal{O}^s$-$kernel(\mathbf{Q})$ consists of n_0 edges, this takes time $O((n_0 + m)\log(n_0 + m) +$ (the number of intersections)). Let k be the number of intersections. We want to obtain an estimate on k. k consists of two parts, k_1 and k_2, where k_1 counts the intersections of $\mathcal{O}^s$-$kernel(\mathbf{Q})$ with the external kernels of the holes and k_2 counts the intersections among the edges of the external kernels of the holes. Since an edge e of $\mathcal{O}^s$-$kernel_{ext}(\mathbf{H}_i)$ is $\mathcal{O}$-oriented and $\mathcal{O}^s$-$kernel(\mathbf{Q})$ is $\mathcal{O}$-convex by Corollary 5.15, $\partial\mathcal{O}^s$-$kernel(\mathbf{Q})$ intersects e at most twice. Hence, k_1 is at most $2 \cdot 8m = O(m)$. The $O(m)$ edges of the external kernels of the holes have at most $O(m^2)$ intersection points. Hence, $k = O(m^2)$. Note that Figure 28 shows that $\bigcap_{1 \leq i \leq m} \mathcal{O}^s$-$kernel_{ext}(\mathbf{H}_i)$ may, indeed, consist of $O(m^2)$ components. Therefore, these bounds are optimal. We have proven the following theorem.

Theorem 6.12 *The $\mathcal{O}^s$-kernel of a multiply connected polygon with n vertices and m holes can be computed in time $O(n(\log |\mathcal{O}| + \log n) + m(|\mathcal{O}| + m))$ time, for finite $\mathcal{O}$.*

Note that this result is optimal for small $\mathcal{O}$, i.e., if $|\mathcal{O}| = O(m)$ by our above observation on the number of components and Lemma 6.4.4 of [Raw87].

Before we turn to the relationship between the $\mathcal{O}$-kernel and the $\mathcal{O}^s$-kernel we want to take a look the components of $\mathcal{O}^s$-$kernel(\mathbf{P})$, for a multiply connected polygon $\mathbf{P}$. Let $\mathbf{H}$ denote the intersection of all external kernels of holes of $\mathbf{P}$, i.e.,

$$\mathbf{H} = \bigcap_{1 \leq i \leq m} \mathcal{O}^s\text{-}kernel_{ext}(\mathbf{H}_i).$$

In the following lemma we show that we do not introduce new components by intersecting $\mathbf{H}$ with $\mathcal{O}^s$-*kernel*$(\mathbf{Q})$.

Lemma 6.13 *Let* $\mathbf{P}$ *be a multiply connected polygon in the plane with enclosing polygon* $\mathbf{Q}$ *and holes* $\mathbf{H}_1, \ldots, \mathbf{H}_m$. *If* p *and* q *are in* $\mathcal{O}^s$-*kernel*$(\mathbf{P})$ *and* p *and* q *belong to the same component of* $\bigcap_{i=1}^m \mathcal{O}^s$-*kernel*$_{ext}(\mathbf{H}_i)$, *then there is an* $\mathcal{O}$-*convex path connecting* p *and* q *in* $\mathcal{O}^s$-*kernel*$(\mathbf{P})$.

Proof: Let $\mathbf{C}$ be the component of $\bigcap_{i=1}^m \mathcal{O}^s$-*kernel*$_{ext}(\mathbf{H}_i)$ that contains p and q. As we noted above $\mathbf{C}$ is strong-$\mathcal{O}$-convex; hence, $\mathbf{C}$ contains $\mathcal{O}$-$\|[p, q]$, i.e., $\mathbf{C}$ contains all $\mathcal{O}$-convex paths from p to q. By Lemma 5.15 $\mathcal{O}^s$-*kernel*$(\mathbf{Q})$ is $\mathcal{O}^s$-convex. Therefore, there is an $\mathcal{O}$-convex path $\mathcal{P}$ in $\mathcal{O}^s$-*kernel*$(\mathbf{Q})$ that connects p and q. Since $\mathbf{C}$ contains all $\mathcal{O}$-convex paths from p to q, we have $\mathcal{P} \subseteq \mathbf{C} \cap \mathcal{O}^s$-*kernel*$(\mathbf{Q}) \subseteq \mathcal{O}^s$-*kernel*$(\mathbf{P})$. □

This again implies that there are at most $O(m^2)$ components in $\mathcal{O}^s$-*kernel*$(\mathbf{P})$. An immediate consequence of the above lemma is the following corollary.

Corollary 6.14 *The components of* $\mathcal{O}^s$-*kernel*$(\mathbf{P})$ *are* $\mathcal{O}$-*convex.*

6.7 The Relationship between $\mathcal{O}^s$-*kernel*$(\mathbf{P})$ and $\mathcal{O}$-*kernel*$(\mathbf{P})$

In this subsection we restrict ourselves to simple polygons. We investigate the relationship between $\mathcal{O}^s$-*kernel*$(\mathbf{P})$ and $\mathcal{O}$-*kernel*$(\mathbf{P})$ which is straightforward and not very surprising. In the following lemma we show that $\mathcal{O}^s$-*kernel*$(\mathbf{P})$ is a component of $\mathcal{O}$-*kernel*$(\mathbf{P})$.

Theorem 6.15 *If* $\mathbf{P}$ *is a simple polygon, then we have that*

$$\mathcal{O}^s\text{-}kernel(\mathbf{P}) \text{ is a component of } \mathcal{O}\text{-}kernel(\mathbf{P}).$$

Proof: Quite obviously we have $\mathcal{O}^s$-*kernel*$(\mathbf{P}) \subseteq \mathcal{O}$-*kernel*$(\mathbf{P})$. Hence, there is a component $\mathbf{C} \subseteq \mathcal{O}$-*kernel*$(\mathbf{P})$ that contains $\mathcal{O}^s$-*kernel*$(\mathbf{P})$. We have to show that $\mathbf{C} = \mathcal{O}^s$-*kernel*$(\mathbf{P})$. We first prove the claim for $|\mathcal{O}| = 1$ and, then, extend it with the intersection lemmas to a general $\mathcal{O}$.

So for now assume w.l.o.g. that $\mathcal{O} = \{0°\}$. Clearly, it suffices to show that the edges of $\partial\mathbf{C}$ that do not completely belong to $\partial\mathbf{P}$ equal the edges in $\mathcal{O}^s$-*kernel*$(\mathbf{P}) \setminus \partial\mathbf{P}$. We know that every edge of $\mathbf{C}$ or $\mathcal{O}^s$-*kernel*$(\mathbf{P})$ that contains inner points of $\mathbf{P}$ is horizontal. Let e be an upper edge of $\mathcal{O}^s$-*kernel*$(\mathbf{P})$ that contains inner points.

Since e exists, we know by Theorem 6.2 that e intersects a reflex $\{0°\}$-minimum p. Clearly, no point q of $\mathbf{C}$ can be in a sufficiently small neighbourhood above e since otherwise the horizontal line through q crosses the edges incident on p. The same argument holds for the lower edge of $\mathcal{O}^s$-*kernel*$(\mathbf{P})$. This proves the claim for $|\mathcal{O}| = 1$.

Now, let $|\mathcal{O}| > 1$. For each $\theta \in \mathcal{O}$, we have $\{\theta\}^s$-*kernel*$(\mathbf{P})$ is a component of $\{\theta\}$-*kernel*$(\mathbf{P})$. Hence, $\mathcal{O}^s$-*kernel*$(\mathbf{P}) = \bigcap_{\theta\in\mathcal{O}} \{\theta\}^s$-*kernel*$(\mathbf{P})$ is one of the components of $\bigcap_{\theta\in\mathcal{O}} \{\theta\}$-*kernel*$(\mathbf{P}) = \mathcal{O}$-*kernel*$(\mathbf{P})$. □

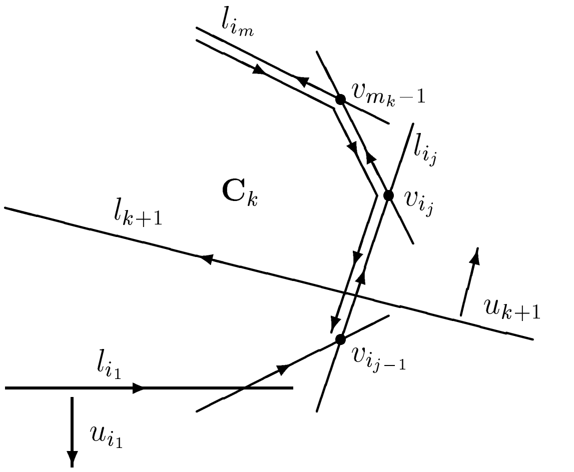

(a) The unbounded case.

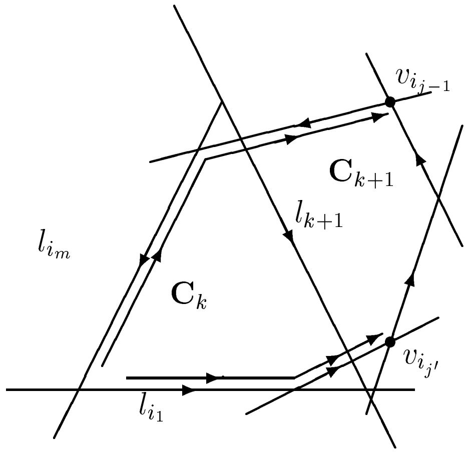

(b) The bounded case.

Figure 40: *Two cases can occur when halfplanes are intersected.*

7 Intersecting Presorted Halfplanes

In this section we present two linear time algorithms to compute the intersection of n halfplanes $h_1, \ldots, h_n$ which are sorted according to slope. For each $1 \leq i \leq n$, we denote the unique oriented line l such that $h^+(l) = h_i$ by l_i. We say that l_i is *associated to* h_i in this case. The intersection of the halfplanes $h_1, \ldots, h_i$ is denoted by $\mathbf{C}_i$ and, furthermore, the normal vector of line l_i that points away from h_i by u_i.

The first algorithm we present uses an incremental approach and computes the sequence $\mathbf{C}_1, \mathbf{C}_2, \ldots, \mathbf{C}_n$. In order to avoid distinguishing between the case that $\mathbf{C}_i$ is unbounded and that $\mathbf{C}_i$ is bounded we will not represent $\mathbf{C}_i$ by its actual boundary, i.e., as a list of rays, vertices and edges as illustrated in Figure 40, but rather as the sequence $l_{i_1}, \ldots, l_{i_m}$ of the "relevant" lines for $\mathbf{C}_i$ and their intersection points $v_{i_j} = l_{i_j} \cap l_{i_{j+1}}$. We say a line l_j is *relevant* for $\mathbf{C}_i$ if the intersection of the halfplanes h_k, with $1 \leq k \leq i$ and $k \neq j$ does not equal $\mathbf{C}_i$, i.e., l_i intersects $\partial \mathbf{C}_i$ in more than one point. Of course, the edges and, if necessary rays, that constitute $\partial \mathbf{C}_i$ can be easily recontructed from $l_{i_1}, \ldots, l_{i_m}$ and $v_{i_1}, \ldots, v_{i_{m-1}}$. If $\mathbf{C}_n$ is unbounded, then there is a pair of halfplanes whose slope differ by more than 180°. Obviously, there can be only one such pair since the slope of the halfplanes is contained in $[0°, 360°)$. W.l.o.g. we assume in the following that such a gap is between h_n and h_1 if it occurs at all. Therefore, we will not consider the case that there are two consecutive halfplanes h_i and h_{i+1} such that $\Theta(h_{i+1}) - \Theta(h_i) > 180°$ in the algorithm.

7.1 An Incremental Algorithm

Assume we are given a description of $\mathbf{C}_k$ as the list of relevant lines $L_k = \{l_{i_1}, \ldots, l_{i_m}\}$ together with the list of vertices $V_k = \{v_{i_1}, \ldots, v_{i_{m-1}}\}$. We will show how to compute the new list of relevant lines and vertices of $\mathbf{C}_{k+1}$ after the intersection of $\mathbf{C}_k$ with h_{k+1} in time proportional to the number of lines and vertices deleted in $L_k \cup V_k$. This immediately yields that the algorithm runs in linear time since line l_k is added at most once to the list of relevant lines L_k and maybe deleted once in a later step.

First assume that $\Theta(h_{k+1}) \leq \Theta(h_1) + 180°$. (See Figure 40 (a) for illustration). Since h_{k+1} differs by more than 180° from h_1 if we measure the angle clockwise, $\mathbf{C}_{k+1}$ is unbounded. Clearly, there is only one intersection point of $\partial\mathbf{C}_k$ and l_{k+1}. Note that $\Theta(h_i) < \Theta(h_{k+1})$, for $i \leq k$. To find it we proceed as follows. Starting with $v_{i_{m-1}}$ we walk clockwise along the boundary of $\mathbf{C}_k$ until we find the first vertex $v_{i_{j-1}}$ that is contained in h_{k+1} and delete all the lines $l_{i_m}, \ldots, l_{i_{j+1}}$ together with the vertices $v_{i_{m-1}}, \ldots, v_{i_j}$. The new list of lines L_{k+1} is given by $l_{i_1}, l_{i_2}, \ldots, l_{i_j}, l_{k+1}$. V_k is then given by $v_{i_1}, \ldots, v_{i_j}$ with $v_{i_j} = l_{i_j} \cap l_{k+1}$. The part of the boundary of $\mathbf{C}_k$ that has not been removed is obviously in h_{k+1} so that all the lines in L_{k+1} are still relevant. To show the correctness of our algorithm we need the following lemma.

Lemma 7.1 *Let l_1, l_2, and l_3 be any three oriented lines with $h_i = h^+(l_i)$ such that*

(i) the lines are sorted according to slope,

(ii) the angle between l_1 and l_2 and the angle between l_2 and l_3 is less than 180°, and

(iii) the intersection of l_1 with l_2 is after the intersection of l_2 with l_3 on l_2.

If $p = l_1 \cap l_3$ is contained in h_2, then $\bigcap h_i$ is non-empty and h_3 is irrelevant, that is $h_1 \cap h_2 \cap h_3$ is the same as $h_1 \cap h_2$.

On the other hand, if p is not contained in h_2, then $\bigcap h_i$ is empty.

Proof: W.l.o.g. assume that $\Theta(l_2) = 180°$, i.e., h_2 is *below* l_2, and that $s = l_1 \cap l_2$ is the origin. Since $t = l_2 \cap l_3$ is before s, t has a positive x-coordinate. Suppose the equations of the halfplanes are given by (cf. Figure 41)

$$h_1\colon\ \cos\theta_1 x_1 + \sin\theta_1 x_2 \leq 0 \text{ and} \tag{1}$$

$$h_3\colon\ \cos\theta_3 x_1 + \sin\theta_3 x_2 + c \leq 0. \tag{2}$$

Let $\mathbf{W}$ be the wedge enclosed by h_1 and h_3. Before we start to prove the claim we derive an inequality that holds for all points in $\mathbf{W}$. Note that θ_1 and θ_3 are the angles of the normal vectors of l_1 and l_3. Hence, we have $\theta_1 \in (-90°, 90°)$ and $\theta_3 \in (90°, 270°)$. This implies that $\cos\theta_1 > 0$ and $\cos\theta_3 < 0$. Let $q = (q_1, q_2)$ be a point in $\mathbf{W}$. From Inequality 1 we immediately get $q_1 \leq -(\sin\theta_1/\cos\theta_1)q_2$ since $\cos\theta_1 > 0$. Further, Inequality 2 yields

$$-\frac{\sin\theta_3}{\cos\theta_3}q_2 - \frac{c}{\cos\theta_3} \leq q_1$$

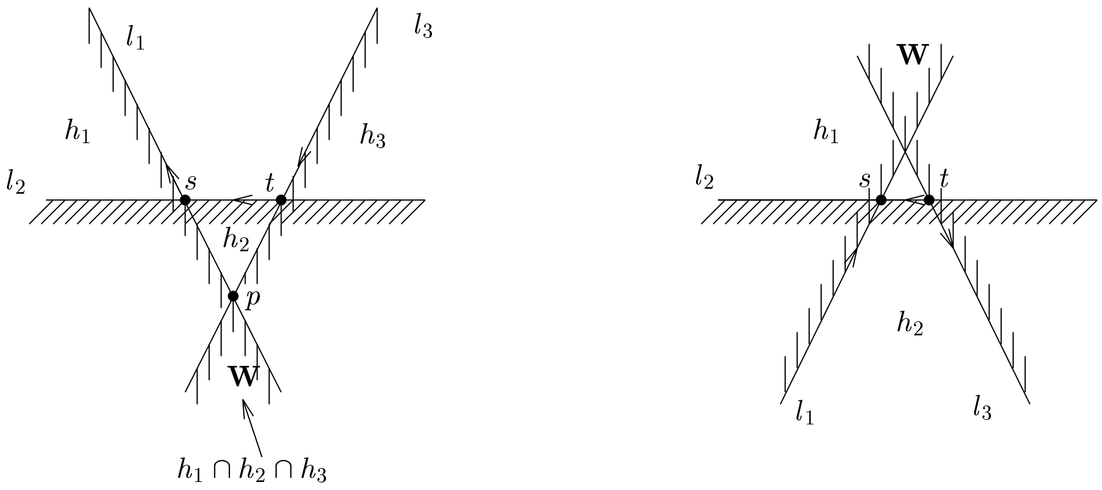

Figure 41: *(a) Either* $H_1 \cap H_1 = H_1 \cap H_1 \cap H_3$ *or* *(b)* $H_1 \cap H_2 \cap H_3$ *is empty.*

because $\cos\theta_3 < 0$. Combining these inequalities we get

$$\frac{\sin\theta_1}{\cos\theta_1}q_2 \quad \leq \quad \frac{\sin\theta_3}{\cos\theta_3}q_2 + \frac{c}{\cos\theta_3}. \tag{3}$$

It is our aim to show that $\mathbf{W}$ is either completely above or completely below l_2. This immediately implies the claimed result since if p is above l_2, then all of $\mathbf{W}$ is above l_2 and $\bigcap h_i$ is empty. On the other hand, if p is below l_2, then $\mathbf{W}$ is completely contained in h_2 and h_2 is irrelevant.

The proof is by contradiction. So assume there are points $q = (q_1, q_2)$ and $r = (r_1, r_2)$ in $\mathbf{W}$ with $q_2 < 0$ and $r_2 > 0$. We then get

$$\begin{aligned} \frac{\sin\theta_1}{\cos\theta_1}q_2 \quad &\leq \quad \frac{\sin\theta_3}{\cos\theta_3}q_2 + \frac{c}{\cos\theta_3} \\ &< \quad \frac{\sin\theta_3}{\cos\theta_3}q_2. \end{aligned}$$

Note that since $t = (t_1, t_2)$ is on l_3 and $t_2 > 0$, we have that $c = -t_2/\cos\theta_3 > 0$ and, hence, $-c/(\cos\theta_3) > 0$. Therefore, $\sin\theta_3/\cos\theta_3 < \sin\theta_1/\cos\theta_1$. Substituting r in Inequality 3 we have that

$$\left(\frac{\sin\theta_1}{\cos\theta_1} - \frac{\sin\theta_3}{\cos\theta_3}\right) r_2 \leq \frac{c}{\cos\theta_3}$$

which, of course, contradicts the fact that $\sin\theta_3/\cos\theta_3 < \sin\theta_1/\cos\theta_1$, $r_2 > 0$ and $c/\cos\theta_3 < 0$. So there cannot be points above and below l_2 in $\mathbf{W}$. $\square$

If we set the first line to be $l_{i_{s-1}}$, the second to be l_{i_s} and the third to be l_{k+1}, for $j+1 \leq s \leq m$, then we see that we can delete l_{i_s} if $v_{i_{s-1}}$ is not in h_{k+1} (see Figure 42).

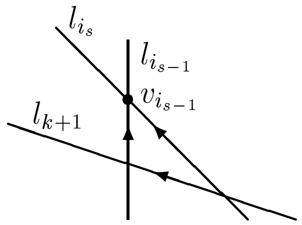

Figure 42: l_{i_s} *can be deleted.*

On the other hand, since there is only one intersection point of l_{k+1} with $\partial\mathbf{C}_k$ and this is on the edge $\overline{v_{i_{j-1}}v_{i_j}}$, the remaining part of $\partial\mathbf{C}_k$ is entirely contained in h_{k+1}.

Now suppose that $\Theta(h_{k+1}) > \Theta(h_1) + 180°$. We assume for the moment that $\mathbf{C}_k$ is bounded and define v as the intersection point of l_{i_m} and l_{i_1}. If v is contained in h_{k+1}, then h_{k+1} is irrelevant for $\mathbf{C}_{k+1}$. To see this we again apply Lemma 7.1 with the first line as l_{i_m}, the second line as l_{k+1} and the third line as l_{i_1}.

So now suppose that either $\mathbf{C}_k$ is unbounded or v is not contained in h_{k+1}. Then, we proceed as outlined above to find an index j such that $v_{i_{j-1}}$ contained in h_{k+1}. Furthermore, starting with v_{i_1} we walk counterclockwise around the boundary until we reach the first vertex $v_{i_{j'}}$ that is contained in h_{k+1}. Since l_{k+1} has at most two intersections with $\partial\mathbf{C}_k$, this suffices. L_{k+1} will be given as the list $l_{i_{j'}}, \ldots, l_{i_j}, l_{k+1}$ and $V_{k+1} = \{v_{i_{j'}}, \ldots, v_{i_{j-1}}, v_j\}$ with $v_{i_j} = l_{i_j} \cap l_{k+1}$.

We describe the algorithm in a PASCAL-like notation.

Algorithm Halfplane intersection
Input:
halfplanes $h_i : a_i x_1 + b_i x_2 + c_i \leq 0,\ i \in \{1, \ldots, n\}$ that are sorted according to slope and whose difference between two consecutive slopes is never greater than 180°.
Output: Lists L and V of lines and vertices that give the boundary of $\mathbf{C}_n$.
Method:
for $k := 1$ **to** n **do**
  if $\mathbf{C}_{k-1}$ is bounded **and** $l_{i_m} \cap l_{i_1}$ is in h_k
    then continue with $k+1$
    else find the first vertex $v_{i_{j-1}}$ in V starting at v_{i_m} that is in h_k;
      delete the lines in L after l_{i_j} and the vertices after v_{i_j} in V;
      add l_k to L and $v_{i_j} = l_k \cap l_{i_j}$ to V;
      if $(\Theta(l_k) > \Theta(l_{i_1}) + 180°)$
        then find the first vertex $v_{i_{j'}}$ in V starting at v_{i_1} that is in h_k;
          delete the lines in L before $l_{i_{j'}}$ and the vertices

before $v_{i_{j'}}$ in V;
end if
end if
end for;

if V and L are empty
then $\bigcap h_i$ is empty
else the vertices of $\mathbf{C}_n$ are given by V;
if $\mathbf{C}_n$ is unbounded
then the rays of $\mathbf{C}_n$ are given by l_{i_1} and l_{i_m}
else the remaining vertex of $\mathbf{C}_n$ is given by $l_{i_1} \cap l_{i_m}$;
end if
end if

7.2 Halfplane Intersection and the Convex Hull

In this section we show how to reduce the intersection of halfplanes that are sorted according to slope to the computation of the convex hull of certain point sets that are sorted on one coordinate. To this end we apply the concept of *geometric duality* which maps points onto lines and vice versa while keeping incidence relations. More precisely we apply the dual map $\mathcal{D}_0$ of H. Edelsbrunner's book [Ede87] which is defined by

$$\begin{aligned}\mathcal{D}_0(p) &= \{(x,y) \in I\!R^2 \mid y = p_1 x - p_2\} \text{ with } p = (p_1, p_2) \text{ and} \\ \mathcal{D}_0(l) &= (a_1, -a_2) \text{ with } l = \{(x,y) \in I\!R^2 \mid y = a_1 x + a_2\}\end{aligned}$$

$\mathcal{D}_0$ has the special property that is preserves order, i.e., if p is above the line l, then the point $\mathcal{D}_0(l)$ is above the line $\mathcal{D}_0(p)$. In order to compute the intersection of the halfplanes $\mathcal{H} := \{h_1, \ldots, h_n\}$ we now proceed as follows. We split $\mathcal{H}$ into two sets $\mathcal{H}_+$ and $\mathcal{H}_-$ such that all the halfplanes that have their interior above them are in $\mathcal{H}_+$ and all the halfplanes that have their interior below them are in $\mathcal{H}_-$. Further, let $\mathcal{L}_+$ be the set of lines that are associated to the halfplanes in $\mathcal{H}_+$ and let $\mathcal{L}_-$ be defined analogously for $\mathcal{H}_-$. Hence, $\bigcap \mathcal{H} = \bigcap \mathcal{H}_+ \cap \bigcap \mathcal{H}_-$. Now $p \in \bigcap \mathcal{H}_+$ if and only if p is above all lines l in $\mathcal{L}_+$; in the dual plane this corresponds to the fact that for all $l \in \mathcal{L}_+$, $\mathcal{D}_0(l)$ is above $\mathcal{D}_0(p)$, i.e., if the convex hull of $\mathcal{D}_0(\mathcal{L}_+)$ is above $\mathcal{D}_0(p)$. If we consider line l' through the vertices v_1 and v_2 on the lower convex hull of $\mathcal{D}_0(\mathcal{L}_+)$, we have that all the points of $\mathcal{D}_0(\mathcal{L}_+)$ are above l' except v_1 and v_2 which lie on l'. So l' corresponds to a point p' in the original plane that is above all lines in $\mathcal{L}_+$ and is the intersection point of $\mathcal{D}_0(v_1)$ and $\mathcal{D}_0(v_2)$. In other words, p' is a vertex of $\bigcap \mathcal{H}_+$. So we have just shown the following observation.

Observation 7.2 *The lines through the edges of the lower convex hull of $\mathcal{D}_0(\mathcal{L}_+)$ correspond to the vertices of $\bigcap \mathcal{H}_+$.*

In the same way we get that $p \in \bigcap \mathcal{H}_-$ if and only if, for all $l \in \mathcal{L}_-$, $\mathcal{D}_0(l)$ is below $\mathcal{D}_0(p)$, i.e., if the convex hull of $\mathcal{D}_0(\mathcal{L}_-)$ is below $\mathcal{D}_0(p)$. Since we obviously can find the intersection of $\bigcap \mathcal{H}_+$ and $\bigcap \mathcal{H}_-$ in linear time, we only have to look at the problem of computing the convex hull of the points $\mathcal{D}_0(\mathcal{L}_+)$ resp. $\mathcal{D}_0(\mathcal{L}_-)$. Clearly, the slopes of the lines $l \in \mathcal{L}_+$ are contained in the interval $[0°, 90°) \cup (270°, 360°)$. If the equation of the line bounding a halfplane with slope θ is $y = a_1 x + a_2$, we have the relation $a_1 = \tan\theta$. Since the halfplanes are sorted according to slope, we can sort the halfplanes in $\mathcal{H}_+$ according to a_1 in linear time if we append the ones with slope in $[0°, 90°)$ to the ones with slope in $(270°, 360°)$. But then $\mathcal{D}_0(\mathcal{L}_+)$ is sorted according to the first coordinate and we can compute the convex hull in time $O(n)$. This gives a second linear time algorithm to compute the intersection of presorted halfplanes.

It is interesting to note that we are able to obtain an "ultimate" halfspace intersection algorithm, i.e., an algorithm that runs in time $O(n \log h)$ if h is the number of vertices in the intersection, if we use a slightly different technique. By applying Meggido's technique for linear programming to the problem of finding a point in $\bigcap h_i$ we can determine in linear time whether $\bigcap h_i$ is empty or find a point p in the interior of $\bigcap h_i$. Translating the halfplanes so that p becomes the origin we can apply the classical dual transformation $\mathcal{D}$ which maps a point (p_1, p_2) to the line given by the equation $y = p_1 x + p_2$ and vice versa. If we apply $\mathcal{D}$ to the lines that correspond to the boundaries of the halfplanes, the vertices of the intersection of the halfplanes just correspond to the lines through the edges of the convex hull of the dual points. The convex hull of the points can be computed in time $O(n \log h)$ by the algorithm of Seidel and Kirkpatrick [KS86] if h is the number of edges on the boundary. By retransforming the convex hull in linear time we obtain the intersection of the halfplanes within the same time bound.

8 A Lower Bound for Reflex θ-Maxima

In this section we give a lower bound of $\Omega(n \log r)$ for the following problem.

> *Given a polygon* $\mathbf{P}$ *with* n *edges and a set of orientations* $\mathcal{O}$ *with* $|\mathcal{O}| = r$ *where* $r \leq n$*, compute all the vertices of* $\mathbf{P}$ *that are reflex* θ*-maxima, for some* $\theta \in \mathcal{O}$.

We also say that a vertex or edge is an *reflex* $\mathcal{O}$*-maximum* if it is a reflex θ-maximum, for some $\theta \in \mathcal{O}$. In order to show this lower bound we first have to look at a different problem.

8.1 A Lower Bound for a Subset-Containment Problem

So we consider the following situation. Given a multiset $A = \{a_1, \ldots, a_n\}$ and a set $B = \{b + (i-1)\varepsilon \mid b, \varepsilon \in \mathbb{R}_+ \text{ and } 1 \leq i \leq r\}$. How much time is needed to find out if every $a_i \in A$ is in B?

The solution space S to this problem is a subset of $\mathbb{R}^{n+2}$ and can be specified as follows.

$$S = \{(a_1, \ldots, a_n, b, \varepsilon) \in \mathbb{R}^{n+2} \mid \forall 1 \leq i \leq n \ \exists 0 \leq k \leq r-1 : a_i = b + k\varepsilon\}.$$

If we are given the function $\lfloor \cdot \rfloor$, then the membership problem for S of a point $p \in \mathbb{R}^{n+2}$ can be solved in time $O(n)$; but we assume the *algebraic computation tree model* which was introduced in Section 2.5 and which does not contain the floor-function. By the theorem of Ben-Or [Ben83] a lower bound for the membership problem for S is $\Omega(\log \#S - n)$ if S consists of $\#S$ connected components. So we will compute the number of components S consists of.

Fixing the last two coordinates of S corresponds to a projection S' of S onto $\mathbb{R}^n$. Since a projection clearly does not introduce any new connected components, the number of components of S' is at least as large as the number of components of S. S' is given by

$$S' = \{(a_1, \ldots, a_n) \in \mathbb{R}^n \mid \text{for all } \ 1 \leq i \leq n \quad \exists\, 0 \leq k \leq r-1 : a_i = b + k\varepsilon\}$$

where b and ε are arbitrary but fixed values. If $\varepsilon \neq 0$, then S' obviously consists of r^n distinct points. Hence, by Ben-Or's theorem a lower bound of the above problem is given by $\Omega(\log r^n - n) = \Omega(n \log r)$.

8.2 A Reduction to the θ-Maxima Problem

Above we have shown that the subset containment problem as we stated it has complexity $\Omega(n \log r)$. We want to use this fact to prove that the problem of finding the reflex $\mathcal{O}$-maxima of a n-vertex polygon has the same complexity. In order to do this we reduce the containment problem to the reflex $\mathcal{O}$-maxima problem in linear time. So assume we are given a multiset $A = (a_1, \ldots, a_n)$ and two values b and ε. W.l.o.g. we assume that $b = 0$, and $0 < \varepsilon < \pi/4$, and that $a_i \in (0, \pi/4 - \varepsilon)$. We now construct a polygon $\mathbf{P}$ as follows. For each a_i, we append two edges e_{2i-1} and e_{2i} to the boundary of $\mathbf{P}$, one with orientation

a_i and the other with orientation $a_i + \varepsilon$. By Lemma 5.3 we obtain a $\{0°, 90°\}$-convex curve $\mathcal{C}$ in this way. If we translate $\mathcal{C}$ to the left and connect the start and end point of $\mathcal{C}$ by two horizontal edges to its translation we obtain a simple polygon $\mathbf{P}$. Let the boundary of $\mathbf{P}$ be oriented counterclockwise and let the first vertex v_1 of $\mathbf{P}$ be at the beginning of edge e_1. Assume we are given an algorithm Q that computes all θ-reflex maxima, with $\theta \in \mathcal{O} = \{0, \varepsilon, 2\varepsilon, \ldots, (r-1)\varepsilon\}$. Let L be the output of Q. If an edge e_j is in L, i.e., e_j is a reflex $\mathcal{O}$-maximum, then $a_{\lceil j/2 \rceil}$ belongs to $\mathcal{O}$ since an edge e is a reflex $\mathcal{O}$-maximum only if $\Theta(e_i) \in \mathcal{O}$ and by construction we have $\Theta(e_{2i-1}) \in \mathcal{O}$ if and only if $\Theta(e_{2i}) \in \mathcal{O}$. If vertex v_{2i} is a reflex $\mathcal{O}$-maximum but edge e_{2i} is not, then there is an orientation $\theta \in \mathcal{O}$ that is between the orientation of e_{2i-1} and e_{2i} which implies that neither of them can belong to $\mathcal{O}$. Hence, we in particular, have that $a_i = \Theta(e_{2i-1}) \notin \mathcal{O}$. So we can solve the above stated subset containment problem in the same time bound that is used by Q. This implies that the complexity of finding the reflex $\mathcal{O}$-maxima is at least $\Omega(n \log |\mathcal{O}|)$ time and proves the following theorem.

Theorem 8.1 *The problem of computing all reflex $\mathcal{O}$-maxima of a n-vertex polygon requires $\Omega(n \log r)$ time in the algebraic computation tree model.*

9 Conclusions

We have introduced a new definition of sight that encompasses both the usual notion of visibility based on straight line segments as well as staircase visibility. It is based on the theory of restricted orientation convexity. We have been able to show that the $\mathcal{O}^s$-kernel of a simple polygon can be obtained as the intersection of the θ-kernels, taken over all $\theta \in \mathcal{O}$. With some additional observations this leads to an $O(n \log |\mathcal{O}|)$ algorithm to compute the $\mathcal{O}^s$-kernel of a simple polygon, for finite $\mathcal{O}$. In order to compute the kernel of a multiply connected polygon we introduced the external $\mathcal{O}^s$- kernel of a polygon and gave an $O(n+r)$ to compute it where r is the number of ranges $\mathcal{O}$ consists of. Combining the two algorithms the $\mathcal{O}^s$-kernel of a multiply connected polygon can be computed in time $O(n^2 + n|\mathcal{O}|)$.

There are several open problems in connection with the above results. As we already mentioned in Section 6.2, it is not known if the algorithm to compute the $\mathcal{O}^s$-kernel of a simple polygon is optimal or if it can be adapted to handle sets of orientations $\mathcal{O}$ that consist of a finite number of ranges. The algorithm to compute the external kernel of a polygon clearly is optimal but we conjecture that the algorithm for the $\mathcal{O}^s$-kernel of multiply connected polygon can be improved upon.

Above we have always assumed that we are given the set of orientations we use for our visibility considerations. In the light of Corollary 5.12 one might ask the following question. Given a polygon $\mathbf{P}$, for which orientations $\theta \in [0°, 360°)$ is $\mathbf{P}$ $\{\theta\}^s$-starshaped, i.e., $\{\theta\}^s$-*kernel*$(\mathbf{P})$ non-empty? This question can be seen as a natural generalization of the problem if a polygon is starshaped.

References

[Ben83] M. Ben-Or. Lower bounds for algebraic computation trees. In *Proceedings of the 15th ACM Symposium on Theory of Computing*, pages 80–86, 1983.

[BG84] B. K. Bhattachara and H. El Gindy. A new linear convex hull algorithm for simple polygons. *IEEE Transactions on Information Theory*, IT-30:85–88, 1984.

[CE88] Bernhard Chazelle and Herbert Edelsbrunner. An optimal algorithm to intersect line segments in the plane. In *Proceedings of the 29th Annual Syposium on Foundations of Computer Science*, pages 590–600, IEEE, Computer Society Press, 1988.

[CG88] Richard Cole and Michael T. Goodrich. Optimal parallel algorithms for polygon and point-set problems. In *Proceedings of the Fourth Annual Symposium on Computational Geometry*, pages 201–210, ACM, ACM Press, Champaign, Illinois, June 1988.

[CR89] J. Culberson and R. Reckhow. *A Unified Approach to Orthogonal Polygon Covering Problems via Dent Diagrams.* Technical Report TR 89-6, Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada, February 1989.

[Ede87] Herbert Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer Verlag, Berlin Heidelberg New York, eatcs monographs on theoretical computer science edition, 1987.

[GY83] R. L. Graham and F. F. Yao. Finding the convex hull of a simple polygon. *Journal of Algorithms*, 4:324–331, 1983.

[Kei86] J. M. Keil. Minimally covering a horizontally convex polygon. In *Proceedings of the Second Annual Symposium on Computational Geometry*, pages 43–51, ACM, Yorktown Heights, New York, June 1986.

[KS86] D. G. Kirkpatrick and R. Seidel. The ultimate planar convex hull algorithm. *SIAM Journal of Computing*, 12(1):287–299, 1986.

[Lee83] D. T. Lee. On finding the convex hull of a simple polygon. *International Journal of Computer and Information Sciences*, 12:87–99, 1983.

[LP79] D. T. Lee and F. P. Preparata. An optimal algorithm for finding the kernel of a polygon. *Journal of the ACM*, 26(3):415–421, July 1979.

[MA79] D. McCallum and D. Avis. A linear algorithm for finding the convex hull of a simple polygon. *Information Processing Letters*, 9:201–206, 1979.

[MRS88a] Rajeev Motwani, Arvind Raghunathan, and Huzur Saran. Covering orthogonal polygons with star polygons: the perfect graph approach. In *Proceedings of the Fourth Annual Symposium on Computational Geometry*, pages 211–223, ACM, Urbana, Illinois, June 1988.

[MRS88b] Rajeev Motwani, Arvind Raghunathan, and Huzur Saran. Perfect graphs and orthogonally convex covers. In *Fourth Annual SIAM Conference on Discrete Mothematics*, page , 1988.

[NLLW83] T. M. Nicholl, D. T. Lee, Y. Z. Liao, and C. K. Wong. Constructing the x-y convex hull of a set of x-y convex polygons. *BIT*, 23:456–471, 1983.

[OW90] Thomas Ottmann and Peter Widmayer. *Algorithmen und Datenstrukturen*. BI Wissenschaftsverlag, 1990.

[Raw87] Gregory J. E. Rawlins. *Explorations in Restricted-Orientation Geometry*. PhD thesis, University of Waterloo, 1987.

[RC87] Robert A. Reckhow and Joseph Culberson. Covering a simple orthogonal polygon with a minimum number of orthogonally convex polygons. In *Proceedings of the Third Annual Symposium on Computational Geometry*, pages 268–277, ACM, Waterloo, Ontario, Canada, June 1987.

[Rin75] W. Rinow. *Topologie*. Volume 79 of *Hochschulbücher für Mathematik*, VEB Deutscher Verlag der Wissenschaften, Berlin, 1975.

[RW87] Gregory J. E. Rawlins and Derick Wood. On the optimal computation of finitely-oriented convex hulls. *Information and Computation*, 72:150–166, 1987.

[RW88a] G. J. E. Rawlins and D. Wood. Computational geometry with restricted orientations. In *Proceedings of the 13th IFIP Conference on System Modelling and Optimization*, Springer Verlag, 1988. Lecture Notes in Computer Science.

[RW88b] Gregory J. E. Rawlins and Derick Wood. Ortho-convexity and its generalizations. In Godfried T. Toussaint, editor, *Computational Morphology*, pages 137–152, Elsevier Science Publishers B. V., (North-Holland), 1988.

[Sch89] R. Schneider. Convex geometry. SS 1989. Script to Lecture.

[SW91] S. Schuierer and D. Wood. *Visibility in Semi-convex Spaces*. Technical Report 37, Institut für Informatik, Universit"at Freiburg, 1991.

[WY88] Derick Wood and Chee K. Yap. The orthogonal convex skull problem. *Discrete and Computational Geometry*, 3(4):349–365, 1988.

[Yao89] Andrew Chi-Chih Yao. Lower bounds for algebraic computation trees with integer inputs. In IEEE, editor, *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, pages 308–313, Computer Society Press, October 1989.