Generalized Kernels of Polygons with Holes*

Sven Schuierer† Derick Wood ‡

Abstract

Let $\mathcal{O}$ be some set of orientations, that is, $\mathcal{O} \subseteq [0°, 360°)$. We consider the consequences of defining visibility based on curves that are monotone with respect 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 *$\mathcal{O}$-see* each other if there exists an $\mathcal{O}$-staircase from p to q that is completely contained in P. The *$\mathcal{O}$-kernel* of a polygon P is then the set of all points which $\mathcal{O}$-see all other points. It can be shown that the $\mathcal{O}$-kernel of a simple polygon can be computed in time $O(n \log |\mathcal{O}|)$. In this paper we show how to compute the *external $\mathcal{O}$-kernel* of a polygon in optimal time $O(n + |\mathcal{O}|)$ and how to combine the two algorithms to compute the $\mathcal{O}$-kernel of a polygon with holes in time $O(n^2 + n|\mathcal{O}|)$.

1 Introduction

Visibility problems play an important role in computational geometry. Apart from the usual line segment visibility several other notions of visibility have been investigated in the past years: *staircase visibility* [4,11], *rectangular visibility* [7,13,15] and *periscope visibility* [5]. In this paper we introduce a new definition of sight called *$\mathcal{O}$-visibility*. It is based on the framework of *restricted orientation convexity* which was first considered by Rawlins [17].

Restricted-orientation convexity tries to bridge the gap between Euclidean convexity and $\{0°, 90°\}$-convexity. Recall that a set S is *$\{0°, 90°\}$-convex* or *orthogonally convex* if the intersection of S with any axis-parallel line is connected. $\{0°, 90°\}$-convexity is a well-studied area [4,7,12,14,20,23]. Rawlins and Wood generalize the idea of $\{0°, 90°\}$-convexity and develop the theory of *restricted-orientation convexity* or *$\mathcal{O}$-convexity* [17,16,18,19]. 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 (when $\mathcal{O} = \{0°, 90°\}$) and Euclidean convexity (when $\mathcal{O}$ is the set of all orientations) as special cases.

The framework of $\{0°, 90°\}$-convexity spawns a new definition of visibility called *staircase visibility* or *$\{0°, 90°\}$-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 [4] and Motwani et al. [11]. Both papers deal with covering polygons with the minimum number of $\{0°, 90°\}$-starshaped sets; that is, 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 *$\mathcal{O}$-visibility* [21]. $\mathcal{O}$-visibility again encompasses $\{0°, 90°\}$-visibility and Euclidean visibility as special cases.

One of the central visibility problems is the computation of the kernel of a polygon. The *kernel* of a set S is the set of points that see all other points in S. For Euclidean visibility the computation of the kernel is a well-studied problem for which several optimal algorithms have been developed [3,9], whereas the corresponding problem for $\{0°, 90°\}$-visibility has not been considered. $\mathcal{O}$-visibility offers us 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 is a first step in this direction (see also [21]).

Visibility questions can also be seen as *reachability* questions. Consider the problem of "guarding" a polygon with one robot whose motion is restricted in such a way that the robot's path must be *monotone* in some set of orientations. In order to decide whether it is possible to guard a polygon and, if so, where to place the guard, we consider the *$\mathcal{O}$-kernel* of a polygon; that is, the set of points that $\mathcal{O}$-see all other points.

The rest of this paper is organized as follows. We start off with a precise definition of $\mathcal{O}$-convexity and $\mathcal{O}$-visibility in the next section. In Section 3 we turn to the computation of the $\mathcal{O}$-kernel of a polygon with holes. In contrast to Euclidean visibility, the $\mathcal{O}$-kernel is not necessarily empty if there are holes. Since holes are polygons themselves, there exist polygons which can be seen completely from their exterior. This observation gives rise to the definition of the *external $\mathcal{O}$-kernel* of a polygon which, in turn, can be employed to compute the $\mathcal{O}$-kernel of a polygon with holes. This is considered in Section 4.

2 Basic Definitions for Restricted Orientation Convexity

If we are given an oriented line l in the plane, we define its *orientation* to be the angle it 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 segment or a ray. As already stated in the introduction the $\mathcal{O}$-convex sets can now be defined as follows.

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

In order not to have to deal with orientations that are greater than 360°, we assume from now on that the addition and sub-

*This work was supported by the Deutsche Forschungsgemeinschaft under Grant No. Ot 64/5-4 and the Natural Sciences and Engineering Research Council of Canada and Information Technology Research Centre of Ontario.

†Institut für Informatik, Universität Freiburg, Rheinstr. 10-12, D-7800 Freiburg, Fed. Rep. of Germany; email: schuierer@informatik.uni-freiburg.de

‡Department of Computer Science, University of Western Ontario, London, Ontario, N6A 5B7, Canada. email: dwood@watdragon.uwaterloo.ca

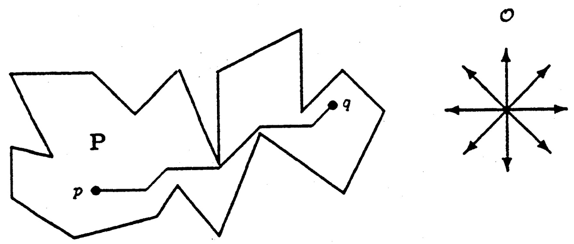

Figure 1: *The definition of $\mathcal{O}$-visibility.*

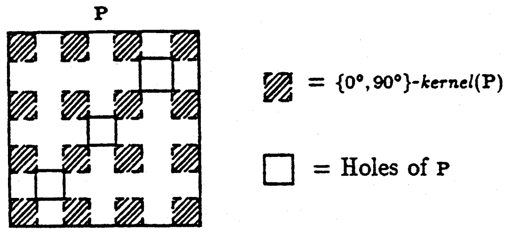

Figure 2: *The $\mathcal{O}$-kernel of* P *is disconnected and not $\mathcal{O}$-convex.*

traction of two orientations is done modulo 360°. The first thing to note about the definition of $\mathcal{O}$-convexity is that we can assume that $\mathcal{O}$ is symmetric with respect to 180°, that is, if the 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 $\theta^{-1} \in \mathcal{O}$. 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 specific $\mathcal{O}$ though it should be kept in mind that $\mathcal{O}$ also always contains the opposite orientations in $[180°, 360°)$.

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 empty set.

If p and q are two points in the plane, we denote the line segment between p and q by $\overline{pq}$. It can be shown that if $\Theta(\overline{pq}) \notin \mathcal{O}$ and (θ_1, θ_2) is the maximal $\mathcal{O}$-free range of $\overline{pq}$, then a curve $\mathcal{S}$ from p to q is $\mathcal{O}$-convex if and only if $\mathcal{S}$ is $\{\theta_1, \theta_2\}$-convex [17]. Note that if $\Theta(\overline{pq}) \in \mathcal{O}$, then the only $\mathcal{O}$-convex curve from p to q is $\overline{pq}$.

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*. Two points p and q in a set P are *$\mathcal{O}$-visible* from each other or *$\mathcal{O}$-see* each other if there exists an $\mathcal{O}$-stairsegment from p to q that is completely contained in P (see Figure 1).

As we already mentioned, we are interested in the *$\mathcal{O}$-kernel* of a set P which is the set of points in P that $\mathcal{O}$-see all the other points in P and which is denoted by $\mathcal{O}$-*kernel*(P). We are mainly concerned with *polygons*. A *simple polygon* is the union of a simple closed curve and its interior such that the simple closed curve consists of (a finite number of) line segments (called edges) and no two consecutive edges are collinear. If we want to refer to the curve that surrounds a simple polygon P, we speak of the *boundary of* P which is denoted by ∂P. The *exterior* of P is defined to be the set of points of $\mathbb{E}^2$ that do not belong to P and is denoted by *ext*(P). We define a *polygon* P *with holes* to be a simple polygon P_0 called the *enclosing polygon of* P that contains a number of disjoint simple polygons $\mathrm{H}_1, \ldots, \mathrm{H}_k$ in its interior called the *holes of* P. The exterior of P is defined as union of the exterior of P_0 and the interior of the holes of P. P is also called *multiply connected* in this case.

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

In the following it is our main aim to show how to compute the $\mathcal{O}$-kernel of a polygon. As far as simple polygons are concerned the following theorem holds [22].

Theorem 3.1 *The $\mathcal{O}$-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 $\mathcal{O}$.*

It is based on the following observation about the $\mathcal{O}$-kernel of a simple polygon [22].

Lemma 3.2 *If* P *is a simple polygon and $\mathcal{O}$ a set of orientations, then*

(i) $\mathcal{O}$-*kernel*(P) $= \bigcap_{\theta \in \mathcal{O}} \{\theta\}$-*kernel*(P),

(ii) $\mathcal{O}$-*kernel*(P) is $\mathcal{O}$-convex and connected.

3.1 The $\mathcal{O}$-Kernel of a Polygon with Holes

If we allow a polygon to have holes, the situation changes considerably. The $\mathcal{O}$-*kernel*(P) is neither necessarily connected nor $\mathcal{O}$-convex any more as shown in Figure 2. Furthermore, we lose the Intersection Lemma as Figure 3 illustrates; hence, we need a different approach to compute the $\mathcal{O}$-kernel of a multiply connected polygon. Since holes can be viewed as polygons themselves, we are now also concerned with *external visibility*. This notion gives rise to the definition of the *external kernel* of a polygon.

Definition 3.1 *Let* P *be a polygon in the plane and $\mathcal{O}$ some set of orientations. We define the* external kernel of P *as the set of*

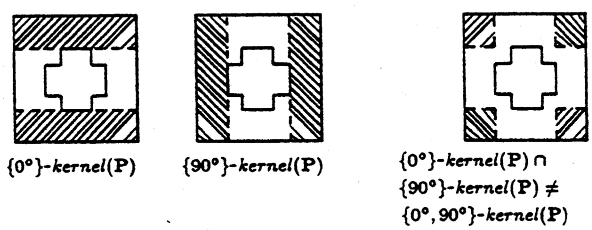

Figure 3: $\{0°, 90°\}$-*kernel*(P) $= \emptyset$ *is not the intersection of* $\{0°\}$-*kernel*(P) *and* $\{90°\}$-*kernel*(P).

all points that belong to $\mathcal{O}$-kernel($\mathbb{E}^2 \setminus int(\mathrm{P})$) and denote it by $\mathcal{O}$-kernel$_{ext}$(P).

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

Lemma 3.3 *If $\mathcal{O}$ is a set of orientations and P a multiply connected simple polygon with enclosing polygon Q and holes $\mathrm{H}_1, \ldots, \mathrm{H}_m$, then*

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

Proof: omitted. □

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

3.2 The External Kernel of a Polygon

We allow the set of orientations $\mathcal{O}$ to consist of a finite number r of closed intervals (or ranges) $[\alpha_1, \beta_1], \ldots, [\alpha_r, \beta_r]$. Before we continue, we need to introduce some more notation. 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 P and a point p on the boundary of P, we say that p is a (local) *θ-maximum of* P if there exists a neighbourhood N of the connected component of ∂P with the θ-oriented line through p such that there is no point in $\partial\mathrm{P} \cap N$ that has a larger $(\theta + 90°)$-coordinate than p.

Given an oriented line l we denote the halfplane to the left of l by $h^+(l)$ and the halfplane to the right of l by $h^-(l)$. An *oriented tangent* t of P is an oriented line such that t intersects ∂P and that P is completely contained in $h^+(t)$. Obviously, there is only one tangent to 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 $h^+(\theta)$ and, similarly, the halfplane to the right by $h^-(\theta)$.

We start with a simple observation (see also [17]).

Observation 3.4 No point which lies between a pair of tangents to P that are parallel to some $\theta \in \mathcal{O}$ belongs to the external $\mathcal{O}$-kernel of P.

We now turn to characterizing the components of the external kernel of a simple polygon P.

Lemma 3.5 *If P is a polygon in the plane, $\mathcal{O}$ a set of orientations with $|\mathcal{O}| \geq 2$, and θ_1 and θ_2 are two adjacent[1] orientations in $\mathcal{O}$, with $\theta_1 \leq \theta_2$, then $h^-(\theta_1) \cap h^-(\theta_2^{-1})$ belongs to $\mathcal{O}$-kernel$_{ext}$(P) if and only if*

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

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

[1] Note that adjacency, in particular, implies that (θ_1, θ_2) is $\mathcal{O}$-free.
[2] A local θ-extremum v is convex if v is a convex vertex of P.

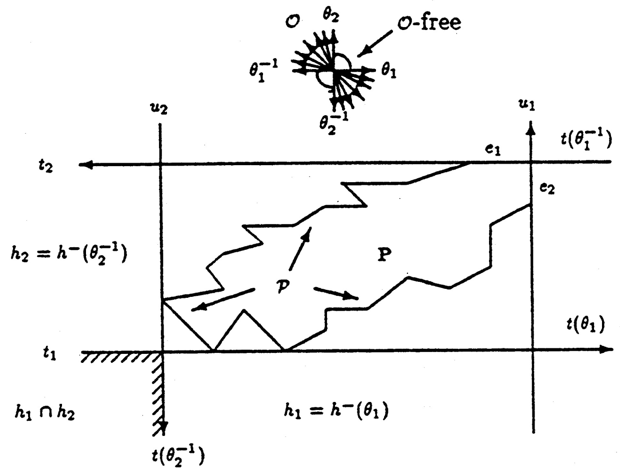

Figure 4: *The external kernel of a polygon*

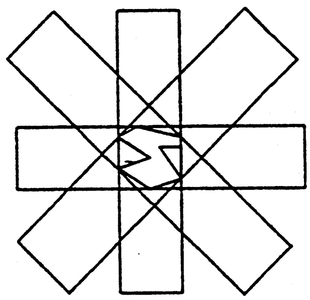

Figure 5: *The arrangement of 2r parallel tangents to* P.

Proof: omitted. □

With Observation 3.4 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 $\mathcal{O}$-kernel$_{ext}$(P) or not. There are no other points in $\mathbb{E}^2 \setminus int(\mathrm{P})$. To see this note that there are r $\mathcal{O}$-oriented pairs of tangents to P and the arrangement of the slabs between these pairs is topologically equivalent to an arrangement of r lines that meet in one point (see Figure 5).

The above approach obviously only works if $|\mathcal{O}| \geq 2$. So it remains to look at $\mathcal{O}$-kernel$_{ext}$(P) if $\mathcal{O} = \{\theta\}$. As expected the situation simplifies considerably and we get the following characterization.

Lemma 3.6 *If P is a simple polyon in the plane, $\mathcal{O} = \{0°\}$, and p is below the lowest point of P, then p belongs to $\mathcal{O}$-kernel$_{ext}$(P) if and only if P contains no reflex minima.*

3.3 Computing the External Kernel

In order to develop an efficient algorithm to compute $\mathcal{O}$-kernel$_{ext}$(P) we need to bound the number of components that it can maximally consist of. The following lemma shows that

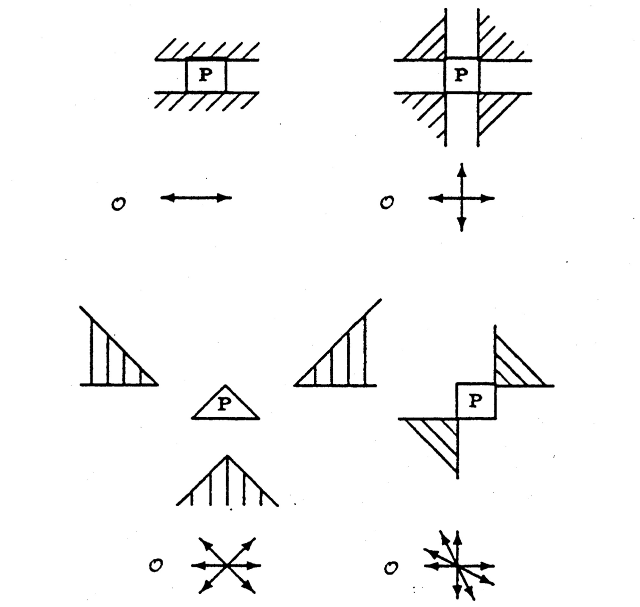

Figure 6: *Lower bounds on the number of components of the external kernel of a polygon.*

the external $\mathcal{O}$-kernel of a polygon has only a constant number of components independent of the cardinality of $\mathcal{O}$ or the size of P. In the subsequently presented algorithm we will make crucial use of this fact.

Lemma 3.7 *Let* P *be a simple polygon. The maximum number of components $\mathcal{O}$-kernel$_{ext}$(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.*

Proof: To see that the external kernel of a polygon can have as many components as claimed refer to Figure 6.

In the first case the claim is an immediate consequence of Lemma 3.6 since either all points in the halfplane below $t(0°)$ belong to $\mathcal{O}$-*kernel*$_{ext}$(P) or none of it. The same, of course, holds for the halfplane above $t(180°)$. The proof of the second claim is as easy to see since the two $\mathcal{O}$-oriented slabs that contain P divide the plane into at most four quarterplanes. By Lemma 3.5 a quarterplane belongs either completely to the external $\mathcal{O}$-kernel of P or not at all. Hence, $\mathcal{O}$-*kernel*$_{ext}$(P) consists of at most four components.

Note that Condition 1 of Lemma 3.5 implies that, for each component w of $\mathcal{O}$-*kernel*$_{ext}$(P), there is a vertex q of the convex hull *conv*(P) of P whose adjacent edges have an interior angle which is less or equal to the $\mathcal{O}$-free range of the two tangents that enclose w, that is, in Figure 7 $\alpha_2 + \beta_2 \leq \alpha = \theta_2 - \theta_1$.

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 $\mathcal{O}$-*kernel*$_{ext}$(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*(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 $[0°, 360°)$, that is, $\bigcup_{i=1}^{4}[\alpha_i, \beta_i] \neq [0°, 360°)$.

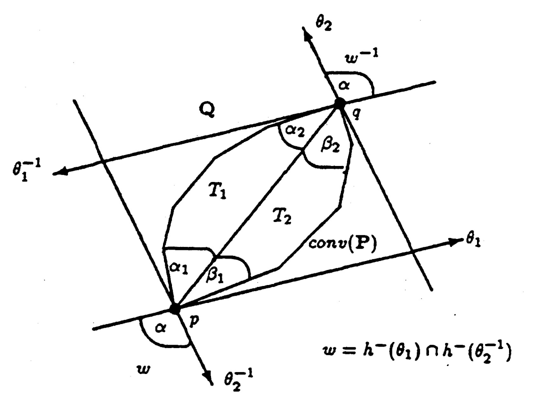

Figure 7: $\alpha_1 + \alpha_2 \leq \alpha$ *and* $\beta_1 + \beta_2 \leq \alpha$

Now *conv*(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*(P) are less than 180°. Hence, $\mathcal{O}$-*kernel*$_{ext}$(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 $h^-(\theta_2^{-1})$. 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 $\mathcal{O}$-*kernel*$_{ext}$(P), then all the remaining interior angles of *conv*(P) are greater than $180° - \alpha$.

To see this note that *conv*(P) is contained in the parallelogram $Q = h^+(\theta_1) \cap h^+(\theta_2) \cap h^+(\theta_1^{-1}) \cap h^+(\theta_2^{-1})$ with sides that are parallel to θ_1 and θ_2 and that Q has the $(\theta_1, \theta_2^{-1})$-apex p and the $(\theta_1^{-1}, \theta_2)$-apex q in common with *conv*(P) as shown in Figure 7. The line segment $\overline{pq}$ partitions Q into two triangles T_1 and T_2. Let α_1 (α_2) be the angle formed by the edge of *conv*(P) in T_1 incident to p (q) with $\overline{pq}$ and β_1 (β_2) defined analogously for T_2 (see Figure 7). 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*(P) in T_1 satisfy the inequality

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

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

Now if there are three components of $\mathcal{O}$-*kernel*$_{ext}$(P), then we have three disjoint, $\mathcal{O}$-free ranges (α_i, β_i) such that there are three vertices of *conv*(P) with interior angles $\sigma_i \leq \beta_i - \alpha_i$, $1 \leq i \leq 3$, by Lemma 3.5. If two of the ranges are opposites to each other, say $\alpha_2 = \alpha_1^{-1}$ and $\beta_2 = \beta_1^{-1}$, then $\sigma_3 \geq 180° - \sigma_1$ 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 $\sigma_1 + \sigma_2 + \sigma_3 < 180°$ and, hence, there are three vertices in *conv*(P) whose sum

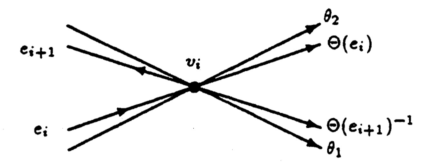

Figure 8: $[\Theta(e_{j+1})^{-1}, \Theta(e_i)]$ *is contained in* $[\theta_1, \theta_2]$.

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 inequalities

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

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, that is, that $\gamma_i \leq 180^\circ$.□

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 [1,6,8,10]. 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 3.5 there is a vertex v_j in $conv(\mathbf{P})$, for each component $h^-(\theta_1) \cap h^-(\theta_2^{-1})$ of $\mathcal{O}\text{-}kernel_{ext}(\mathbf{P})$, such that $\Theta(e_j)$ and $\Theta(e_{j+1})^{-1}$ are contained in $[\theta_1, \theta_2]$ (see Figure 8). As in the proof of Lemma 3.7 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.

The idea of the algorithm is to find these candidate vertices and then to check if the conditions of Lemma 3.5 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 $\mathcal{O}$-kernel
Let (α_1, β_1) be the $\mathcal{O}$-free range that corresponds to $\Theta(e_1)$;
for each edge e_i of $conv(\mathbf{P})$ **do**
    if $[\Theta(e_{i+1})^{-1}, \Theta(e_i)] \subseteq (\alpha_i, \beta_i)$
        then output v_i as a candidate vertex for (α_i, β_i)
    end if
    if $\Theta(e_{i+1}) \leq \beta_i$
        then let $(\alpha_{i+1}, \beta_{i+1}) := (\alpha_i, \beta_i)$
        else step through $\mathcal{O}$ to the $\mathcal{O}$-free range
            $(\alpha_{i+1}, \beta_{i+1})$ 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 3.5 are satisfied. This proves the following theorem.

Theorem 3.8 *The external $\mathcal{O}$-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 $\mathcal{O}$.*

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

With the above result, Lemma 3.3, and Theorem 3.1 we can compute the $\mathcal{O}$-kernel of a multiply connected polygon $\mathbf{P}$ with enclosing polygon $\mathbf{Q}$ and holes $\mathbf{H}_1, \ldots \mathbf{H}_m$ by computing $\mathcal{O}\text{-}kernel(\mathbf{Q})$ and intersecting it with $\mathcal{O}\text{-}kernel_{ext}(\mathbf{H}_i)$, 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 $\mathcal{O}\text{-}kernel(\mathbf{Q})$ and $\mathcal{O}\text{-}kernel_{ext}(\mathbf{H}_i)$, 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 [2] 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 $\mathcal{O}\text{-}kernel(\mathbf{Q})$ consists of at most 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 $\mathcal{O}\text{-}kernel(\mathbf{Q})$ 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 $\mathcal{O}\text{-}kernel_{ext}(\mathbf{H}_i)$ is $\mathcal{O}$-oriented and $\mathcal{O}\text{-}kernel(\mathbf{Q})$ is $\mathcal{O}$-convex by Lemma 3.2, $\partial\mathcal{O}\text{-}kernel(\mathbf{Q})$ intersects e at most twice. Hence, k_1 is at most $O(m)$. 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 2 shows that $\bigcap_{1 \leq i \leq m} \mathcal{O}\text{-}kernel_{ext}(\mathbf{H}_i)$ may, indeed, consist of $O(m^2)$ components. Therefore, these bounds are optimal. We have proven the following theorem.

Theorem 4.1 *The $\mathcal{O}$-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 $\mathcal{O}$, that is, if $|\mathcal{O}| = O(m)$ by our above observation on the number of components and the fact that the computation of the $\mathcal{O}$-kernel of a multiply connected polygon can be shown to have a lower bound of $\Omega(n \log n)$ [17].

5 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. The $\mathcal{O}$-kernel of a simple polygon can be computed in time $O(n \log |\mathcal{O}|)$. In order to compute the kernel of a multiply connected polygon, we introduce the external $\mathcal{O}$-kernel of a polygon and give an $O(n+r)$ time algorithm to compute it where r is the number of ranges in $\mathcal{O}$. Combining the two algorithms, the $\mathcal{O}$-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 these results. The algorithm to compute the external kernel of a polygon is clearly optimal, but we conjecture that the efficiency of the algorithm for the $\mathcal{O}$-kernel of a multiply connected polygon can be improved upon.

We have always assumed that the set of orientations that we use for our visibility considerations is given in advance. If we drop this assumption, the following question can be seen as a natural generalization of the problem if a polygon is starshaped: Given a polygon P, for which orientations $\theta \in [0°, 360°)$ is P $\{\theta\}$-starshaped or $\{\theta\}$-*kernel*(P) non-empty? Though it seems that the techniques developed so far should be helpful in the solution, it is an open problem how fast this question can be answered.

References

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

[2] B. Chazelle and H. Edelsbrunner. An optimal algorithm to intersect line segments in the plane. In *Proc. 29th IEEE Syposium on Foundations of Computer Science*, pages 590–600, 1988.

[3] R. Cole and M. Goodrich. Optimal parallel algorithms for polygon and point-set problems. In *Proc. 4th ACM Symp. on Computational Geometry*, pp. 201–210, 1988.

[4] 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, Canada, February 1989.

[5] L. Gewali and S. Ntafos. Minimum covers for grids and orthogonal polygons by periscope guards. In Jorge Urrutia, editor, *Proc. 2nd Canadian Conference in Computational Geometry*, pages 358–361, University of Ottawa, 1990.

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

[7] J. Keil. Minimally covering a horizontally convex polygon. In *Proc. 2nd ACM Symposium on Computational Geometry*, pages 43–51, 1986.

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

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

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

[11] R. Motwani, A. Raghunathan, and H. Saran. Covering orthogonal polygons with star polygons: the perfect graph approach. In *Proc. 4th ACM Symposium on Computational Geometry*, pages 211–223, 1988.

[12] R. Motwani, A. Raghunathan, and H. Saran. Perfect graphs and orthogonally convex covers. In *4th SIAM Conference on Discrete Mathematics*, 1988.

[13] I. Munro, M. Overmars, and D. Wood. Variations on visibility. In *Proc. 3rd ACM Symposium on Computational Geometry*, pages 291–300, 1987.

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

[15] M. Overmars and D. Wood. On rectangular visibility. *Journal of Algorithms*, 9:372–390, 1988.

[16] G. Rawlins and D. Wood. Computational geometry with restricted orientations. In *Proc. 13th IFIP Conference on System Modelling and Optimization*, Springer Verlag, 1988.

[17] G. Rawlins. *Explorations in Restricted-Orientation Geometry*. PhD thesis, University of Waterloo, 1987.

[18] G. Rawlins and D. Wood. On the optimal computation of finitely-oriented convex hulls. *Information and Computation*, 72:150–166, 1987.

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

[20] R. Reckhow and J. Culberson. Covering a simple orthogonal polygon with a minimum number of orthogonally convex polygons. In *Proc. 3rd ACM Symposium on Computational Geometry*, pages 268–277, 1987.

[21] S. Schuierer. *On Generalized Visibility*. PhD thesis, Universität Freiburg, 1991.

[22] S. Schuierer, G. Rawlins, and D. Wood. A generalization of staircase visibility. In T. Shermer, editor, *Proc. 3rd Canadian Conference on Computational Geometry*, pp. 96–99, Simon Fraser University, 1991.

[23] D. Wood and C. Yap. The orthogonal convex skull problem. *Discrete and Computational Geometry*, 3(4):349–365, 1988.