HKUST Theoretical Computer Science Center Research Report HKUST-TCSC-98-06

Multiple-Guard Kernels of Simple Polygons*

Sven Schuierer† Derick Wood‡

Abstract

We study a generalization of the kernel of a polygon. A polygon P is k guardable if there are k points in P such that, for all points p in P, there is at least one of the k points that sees p. We call the k points a k-guard set of P and the k-kernel of a polygon P is the union of all k-guard sets of P. The usual definition of the kernel of a polygon is now the one-kernel in this notation.

We show that the two-kernel of a simple polygon P with n edges has $O(n^4)$ components and that there are polygons with $\Omega(n)$ components. Moreover, we show that the components of the two-kernel of a simple polygon can be paired in a natural manner which implies that the two-kernel of a simple polygon has either one component or an even number of components. Finally, we consider the question of whether there is a non-starshaped simple polygon P such that two-$kernel(P) = P$. We show that when the two-kernel has only one component, it contains a hole; hence, the two-kernel of a simple polygon P is never P itself. For every $k \geq 1$, there are, however, polygons P_k with holes such that k-$kernel(P_k) = P_k$.

1 Introduction

The concept of visibility plays an important role in many fields of mathematics and computer science. In particular, in planar geometry many aspects of visibility have been studied. One of the central concepts in dealing with visibility is the **kernel** of a set. It is defined as the subset of all points in the given set that see all other points. It can be easily shown that the kernel of a set is convex. The dimension of the kernel and its relationship to various other properties of a set are well studied [4, 13].

A different approach than considering sets that can be seen by only one point is to consider the minimal number of points that together see a given set. One famous result obtained in this context is the so called Art Gallery theorem [6, 7]. It states that $\lfloor n/3 \rfloor$ points are always sufficient to guard an art gallery with n walls. If the art gallery contains h holes, then the number increases to $\lfloor (n+h)/3 \rfloor$ points [5, 9]. These bounds are tight since there are polygons that need this many guards. On the other hand, $\lfloor n/3 \rfloor$ is a coarse upper bound when a specific polygon is considered. It is an NP-hard problem to compute the minimum number of points necessary to guard a simple polygon [1, 10]. With this viewpoint, a polygon has a nonempty kernel if and only if it can be guarded by one guard.

In this paper, we examine a generalization of the kernel of a polygon for more than one guard. We call a polygon ***k*-guardable** if k points are enough to guard it. The ***k*-kernel** of a polygon is defined to be the union of all sets of k points that guard the given polygon. Whereas the kernel of a polygon has only one convex component, the k-kernel displays a much greater morphological variety even in the case $k = 2$. In particular, when a polygon has n edges, there can be $\Omega(n)$ components.

*This work was supported under grant No. Ot 64/8-2, "Diskrete Probleme" from the the Deutsche Forschungsgemeinschaft, grants from the Natural Sciences and Engineering Council of Canada, from the Information Technology Research Centre of Ontario, and from the Research Grants Council of Hong Kong.

†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, Hong Kong University of Science & Technology, Clear Water Bay, Kowloon, Hong Kong. E-mail address: dwood@cs.ust.hk.

One intriguing property of convex sets is that they are identical to their kernel. Clearly, this is the only class of one-guardable sets that can be guarded by one point that has this property since the kernel itself is convex. It is an open question whether there exist other classes of polygons that have the property that every point is a member of a minimal set that guards the polygon.

We are interested in two questions concerning the k-kernel of a polygon. First, what can be said about the number of components of $k\text{-}kernel(P)$? Second, are there simple polygons with $k\text{-}kernel(P) = P$ and $k = g(P) > 1$? In other words, is there a simple polygon P with $g(P) = k$ that has a trivial nonempty k-kernel?

We concentrate on the case $k = 2$ and we elucidate some properties of two-kernels. In particular, we prove the following three results. Let $|\text{two-}kernel(P)|$ be the number of components of two-$kernel(P)$.

1. If P is a simple polygon with n edges, then $|\text{two-}kernel(P)| \in O(n^4)$ and there is a simple polygon P with $|\text{two-}kernel(P)| \in \Omega(n)$, for every $n \geq 3$.

2. If P is a simple polygon, then two-$kernel(P)$ consists either of one component or of an even number of components.

3. If P is a simple polygon that is not starshaped, then two-$kernel(P) \neq P$.

Belleville [3] has given an $O(n^4)$-time and $O(n^2 \log n)$-space algorithm to decide if a simple polygon is two-guardable. For rectilinear polygons this is improved by Tan and Zhu to $O(n^3 \log n)$ time and $O(n)$ space [14]. Our results can be seen as extension of this work.

2 Definitions

We begin with some basic definitions for polygons and guards. A **polygonal chain** is a path that consists of (a finite number of) line segments called **edges** such that no two consecutive line segments are collinear. The end points of edges are called **vertices**. A **(simple) polygon** is a simple, closed polygonal chain. For convenience, we identify a polygon with the region it encloses. If we want to refer to the polygonal chain enclosing a polygon P, we use the notation ∂P for the boundary of P. A **multiply connected polygon** is a simple polygon that contains polygonal holes in its interior. We assume that the edges of a polygon P are oriented counterclockwise; that is, if e is an edge of P, then the interior of P is locally to the left of e.

If p and q are two points inside a polygon P, p is said to **see** q if the line segment $\overline{pq}$ does not intersect the exterior of P. The **visibility polygon $vis(p)$ of a point p** is defined as the set of points inside P that p sees. If Q is a subset of P, then the visibility polygon $vis(Q)$ of Q is defined as $\bigcup\{vis(p) \mid p \in Q\}$.[1] A maximal line segment of the boundary of $vis(Q)$ that does not belong to the boundary of P is called a **window** of $vis(Q)$. A window w splits P into two parts one of which contains $vis(Q)$. The part that does not contain $vis(Q)$ is called the **pocket of w**. If p and q are two different points in P, then the **visibility ray $r_p(q)$** of p w.r.t. q is the maximal line segment in P that has one end point at p and that contains $\overline{pq}$.

The **kernel** of polygon P is defined as the set of points p such that $vis(p) = P$. If the kernel of a polygon is nonempty, the polygon is **starshaped** (see Figure 1a).

Let P be a polygon and G be a subset of P. We say that **G sees (or guards) P** if $vis(G) = P$. We denote the cardinality of G by $|G|$. If $|G| = k$, then G is called a **k-guard set.** Furthermore, if there is no k'-guard set $G' \subseteq P$ with $k' < k$, then G is called a **minimal guard set** of P. The cardinality $g(P)$ of a minimal guard set G of P is called the **guarding number** of P. We say that P is **k-guardable** if $g(P) \leq k$ and **exactly k-guardable** if $g(P) = k$.

The art gallery theorem can now be stated as $g(P) \leq \lfloor n/3 \rfloor$, for all polygons P with n edges. As we noted in the introduction, given a simple polygon P, it is NP-hard to compute $g(P)$. We investigate polygons that have a fixed guarding number $g(P)$ that is independent of n. In

[1] If $\mathcal{F}$ is a family of sets, then $\bigcup \mathcal{F}$ denotes the set $\bigcup_{S \in \mathcal{F}} S$.

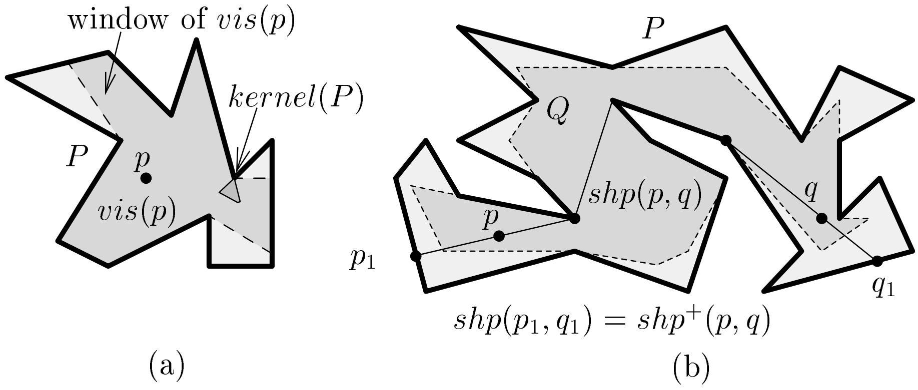

Figure 1: (a) The visibility polygon of the point p and the kernel of the starshaped polygon P. (b) Q is a geodesically convex set in P and the linear extension of the shortest path $shp(p, q)$ from p to q is given by $shp(p_1, q_1)$.

particular, we are interested in a generalization of the kernel of a polygon. Clearly, if $g(P) > 1$, then the kernel of P is empty.

There are numerous ways to generalize the notion of the kernel of a polygon for more than one guard. We base our investigations on the following definition.

Definition 2.1 *Let P be a simple polygon and $k \geq 1$. If $\mathcal{G}_k$ is the family of all k-guard sets of P, then we define the **k-kernel of P** to be the set $\bigcup \mathcal{G}_k$ which we denote by k-kernel(P).*

Sometimes we consider the k-kernel of a (not necessarily connected) subset Q of P where the visibility of two points in Q is determined by P and not by Q. In this case we denote the k-kernel of Q w.r.t. P by k-$kernel_P(Q)$.

It is also possible to view the k-kernel of a polygon P as a subset of the k-fold Cartesian product of P when we treat P as a set.

Definition 2.2 *Given a simple polygon P, we define the k^+-kernel of P to be the set of all k-tuples $(g_1, \ldots, g_k)$ such that $\{g_1, \ldots, g_k\}$ guards P.*

Note that the k^+-*kernel* of P is a subset of $\mathbb{R}^{2k}$ and the k-kernel, as defined in Definition 2.1, is the projection of the subset of $\mathbb{R}^{2k}$ onto $\mathbb{R}^2$ given by projecting, say, the first two components.

If G_k is a k-guard set of P and $g, g' \in G_k$, then g is called a **partner** of g'. For a point g in k-$kernel(P)$, we define the **co-kernel of g** to be the set $(k-1)$-$kernel_P(P \setminus vis(g))$ and denote it by co-$kernel_k(g, P)$.

There is a close connection between guarding a polygon and covering it with starshaped polygons as stated in the following observation.

Observation 2.1 *A polygon P is k-guardable if and only if P is the union of k starshaped sets.*

Obviously, k-$kernel(P)$ is empty, for $0 \leq k < g(P)$, and k-$kernel(P)$ is P, for $k > g(P)$. So the only interesting value of k is $g(P)$. We define the **generalized kernel $kernel^*(P)$ of P** to be k-$kernel(P)$ where $k = g(P)$. Hence, if $g(P) = 1$, then $kernel^*(P)$ is the usual kernel of P. We denote the number of components of k-$kernel(P)$ by $|k$-$kernel(P)|$.

Sometimes we need to speak about the distance between two sets A and B. If A and B are two compact sets, then the Hausdorff-metric $H(A, B)$ is defined as

$$\max(h(A, B), h(B, A))$$

where

$$h(A, B) = \min_{a \in A} \max_{b \in B} d(a, b).$$

It is easy to see that $H(A, B)$ is a metric on the family of compact sets in the plane [8].

An important concept that we make use of is **paths**. Strictly speaking a **path** or **curve** is an equivalence class of continuous functions from an interval to the plane, but we always assume the interval to be [0,1] and we sometimes identify a path π with the its image $\pi([0,1])$ in the plane. If π is a path from p_1 to p_2, then π^{-1} denotes the path with the reverse orientation from p_2 to p_1, that is, $\pi^{-1}(t) = \pi(1-t)$. If the end point of a path π_1 is the start point of a path π_2, then we denote the catenation of π_1 with π_2 by $\pi_1 * \pi_2$.

We also need to define the shortest path between two points p and q, which we denote by $shp(p, q)$. We say a set $Q \subset P$ is **geodesically convex** if, for all pairs p, q of points $\in Q$, the shortest path between p and q is contained in Q. Note that a geodesically convex set is connected and that the intersection of geodesically convex sets is again geodesically convex (see Figure 1b).

If p and q are two points in P, then $shp^+(p, q)$ denotes the **linear extension** of $shp(p, q)$, that is, $shp^+(p, q)$ is the path that we obtain by extending the first and last line segment of $shp^+(p, q)$ as much as possible without intersecting the exterior of P.

3 The Number of Components of the Two-Kernel

We establish upper and lower bounds on the number of components of the two-kernel. We start with an observation about the relationship between the shortest path between two points and visibility polygons.

Lemma 3.1 *Let P be a simple polygon. If p is a point in P, then $vis(p)$ is geodesically convex.*

Proof: The proof is by contradiction. So consider two points q_1 and q_2 in $vis(p)$ and assume that the shortest path from q_1 to q_2 is not contained in $vis(p)$. Then, there is a window w of $vis(p)$ that is crossed by $shp(q_1, q_2)$. Since both q_1 and q_2 are contained in $vis(p)$ and $vis(p)$ is simply connected, $shp(q_1, q_2)$ crosses w twice, once to enter the pocket of w and once to leave it again. But then the part of $shp(q_1, q_2)$ inside the pocket of w can be replaced by the part of w between the cross points. This yields a shorter path—a contradiction. □

Corollary 3.2 *Let P be a simple polygon and p and q two points in P. If x is a point on the shortest path between p and q, then*

$$vis(x) \supseteq vis(p) \cap vis(q).$$

Proof: Let t be a point in $vis(p) \cap vis(q)$, that is, t is seen by both p and q and, therefore, $p, q \in vis(t)$. By Lemma 3.1 $shp(p, q)$ belongs $vis(t)$. Hence, if x is a point on $shp(p, q)$, then $x \in vis(t)$ and x sees t as claimed. □

The upper bound on the number of components is based on the observation that every component intersects an extension of an edge.

Definition 3.1 *Let P be a simple polygon and e be one of its edges. The maximal line segment collinear with e inside P is called the* **extension of *e***.

The following lemma is shown by Belleville [3].

Lemma 3.3 *If P is a simple polygon such that $P = P_1 \cup P_2$, for two maximal starshaped sets P_1 and P_2, then $kernel(P_1)$ and $kernel(P_2)$ intersect an extension of an edge of P.*

Belleville shows that each edge extension can be partitioned into $O(n^3)$ line segments such that each line segment either is completely contained in the two-kernel or is completely disjoint from it [3]. In conjunction with Lemma 3.3 it implies the following result.

Lemma 3.4 *The two-kernel of a simple polygon with n edges has at most $O(n^4)$ components.*

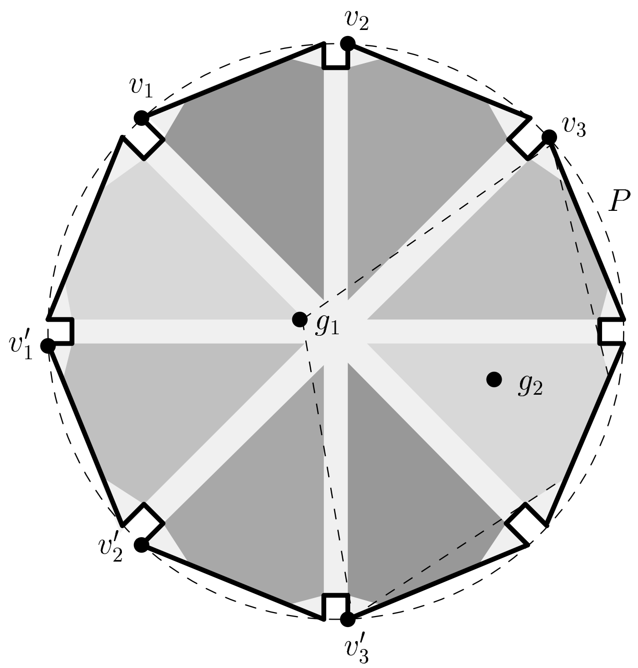

Figure 2: A polygon whose two-kernel has eight components. Components that correspond to each other have the same shading intensity. The dashed lines issuing from g_1 illustrate how the boundaries of the component of g_2 are computed. The point g_2 must see the points v_1, v_2, and v_3, and also see the points v_1', v_2', and v_3' that g_1 does not see.

Proof: If each edge extension is partitioned into $O(n^3)$ line segments such that, for each line segment l, either $l \subseteq \text{two-}kernel(P)$ or $l \cap \text{two-}kernel(P) = \emptyset$, then there are $O(n^4)$ such line segments. Each component C of two-$kernel(P)$ intersects at least one edge extension by Lemma 3.3; hence, it contains at least one line segment l. If component C' also intersects l, then $l \subseteq C \cap C'$ which implies that $C = C'$. Hence, two-$kernel(P)$ has $O(n^4)$ components. □

Figure 2 shows an instance of a family of polygons with n edges whose two-kernel has $\Omega(n)$ components.

4 Pairs of Components

One interesting aspect of the lower bound example in Figure 2 is that the components of the two-kernel of P always seem to occur in pairs. We show that this holds in general if there are at least two components. In particular, the number of components of the two-kernel is either one or even. In the following we investigate the relationship between two-guard sets $\{g_1, g_2\}$ and $\{h_1, h_2\}$.

Lemma 4.1 *Let $\{g_1, g_2\}$ and $\{h_1, h_2\}$ be guard sets of a simple polygon P such that $g_2 \in shp(g_1, h_1)$; then, $\{g_2, h_2\}$ guards P.*

Proof: Let $p \in P$ be chosen arbitrarily. We have to show that g_2 or h_2 sees p. Assume that h_2 does not see p. Since $\{h_1, h_2\}$ guards P, h_1 sees p. If g_1 does not see p, then g_2 sees p since $\{g_1, g_2\}$ guards P. If g_1 sees p, then since $g_2 \in shp(g_1, h_1)$, g_2 sees p by Corollary 3.2. □

Lemma 4.2 *If g_1, g_2, and g_3 are three points in a polygon P such that $\{g_1, g_2\}$ and $\{g_1, g_3\}$ guard P, then $shp(g_2, g_3)$ is contained in two-$kernel(P)$ and g_1 is a partner of all points on $shp(g_2, g_3)$.*

Proof: Let p be some point on $shp(g_2, g_3)$. We claim that $\{g_1, p\}$ sees P. Let q be a point in P and assume that g_1 does not see q. Since $\{g_1, g_2\}$ and $\{g_1, g_3\}$ guard P, both g_2 and g_3 see q, that is, $q \in vis(g_2) \cap vis(g_3)$. By Corollary 3.2 this implies that $q \in vis(p)$ and, hence, p sees q. Therefore, $shp(g_2, g_3)$ belongs to two-$kernel(P)$. $\square$

The above result, in particular, implies that g_2 and g_3 are connected by a path in two-$kernel(P)$. For further reference, we state this implication as a corollary.

Corollary 4.3 *If g_1, g_2, and g_3 are three points in a polygon P such that $\{g_1, g_2\}$ and $\{g_1, g_3\}$ guard P, then g_2 and g_3 are in the same component of two-$kernel(P)$.*

The above observations lead to the following result which implies that there is a natural pairing between the components of the two-kernel. If C is a component of two-$kernel(P)$, then we denote in the following the set of points of two-$kernel(P)$ that have a partner in C by $part(C)$.

Lemma 4.4 *If P is a simple polygon and C is a component of two-$kernel(P)$, then $part(C)$ is contained in one component of two-$kernel(P)$.*

Proof: The proof is by contradiction. So assume there is a component C of two-$kernel(P)$ such that $part(C)$ intersects at least two different components of two-$kernel(P)$.

It is easy to see that the set of partners is closed. Since $part(C)$ intersects at least two different components of two-$kernel(P)$ and every point in C belongs to at least one partner set, C is the union of a finite number of—at most $O(n^4)$—closed sets by Lemma 3.4. Since C is connected, the closed sets cannot be disjoint. But then there is a point p in C such that $p \in part(C') \cap part(C'')$ for two different components C' and C'' of two-$kernel(P)$. By Corollary 4.3 this implies that there is a path in two-$kernel(P)$ between C' and C''—a contradiction. $\square$

Lemma 4.4 implies that the components of two-$kernel(P)$ occur either as pairs or as singletons; we show that the latter is not possible if there is more than one component. To do so we have to investigate the structure of the components of two-$kernel(P)$ that contain their partners. An important observation about guarding a polygonal region R in a simple polygon is that it suffices to see the vertices of R.

Lemma 4.5 *Let P be a polygon and R be a polygonal region in P. If a point p sees all the vertices of R, then p sees all of R.*

Proof: Let q be a point in R and s be the line segment from p to q. We extend s until it intersects an edge e of R. If p sees e, then p also sees q. Hence, we have to show only that p sees all the edges of R.

Let e be an edge of R with vertices v and w. Since p sees v and w, Lemma 3.1 implies that p also sees $shp(v, w)$ which equals e. This proves the claim. $\square$

Lemma 4.6 *If P is a simple polygon and p is a point in P, then co-$kernel_2(p, P)$ is the intersection of the visibility polygons of the vertices of $P \setminus vis(p)$.*

Proof: By Lemma 4.5 q belongs to co-$kernel_2(p, P)$ if and only if q sees all the vertices of $P \setminus vis(p)$; that is, co-$kernel_2(p, P)$ equals the intersection of the visibility polygons of the vertices of $P \setminus vis(p)$. $\square$

Lemma 4.7 *If P is a simple polygon and p is a point in P, then co-$kernel_2(p, P)$ is a geodesically convex polygon.*

Proof: Let $\{v_1, \ldots, v_m\}$ be the vertices of $P \setminus vis(p)$. By Lemma 4.6 co-$kernel_2(p, P)$ equals the intersection $\bigcap\{vis(v_i) \mid 1 \le i \le m\}$. Since $vis(v)$ is geodesically convex by Lemma 3.1, co-$kernel_2(p, P)$ is a geodesically convex polygon as the intersection of geodesically convex polygons. $\square$

We now show that if one guard moves on a polygonal path π, then there is a polygonal path π' for its partner such that the polygon is guarded at all times t by $\{\pi(t), \pi'(t)\}$. This result can be viewed as a strengthening of Lemma 4.4.

We start off with an observation about the boundary of $co\text{-}kernel_2(\pi(t), P)$ if π is a path in two-$kernel(P)$.

Lemma 4.8 *If P is a simple two-guardable polygon and π is a polygonal path in two-kernel(P), then the boundary of co-kernel$_2(\pi(t), P)$ depends continuously on t if we use the topology induced by the Hausdorff-metric.*

Proof: Let $p_t = \pi(t)$. We have to show that, for all $t_0 \in [0, 1]$,

$$\lim_{t\to t_0} co\text{-}kernel_2(p_t, P) = co\text{-}kernel_2(p_{t_0}, P)$$

where the convergence of $co\text{-}kernel_2(p_t, P)$ is measured w.r.t. to the Hausdorff-metric. In fact, it suffices to show that, for all $t_0 \in [0, 1]$,

$$\lim_{t\to t_0-} co\text{-}kernel_2(p_t, P) \quad = \quad co\text{-}kernel_2(p_{t_0}, P) \tag{1}$$

since once we have shown this equality, we can consider the inverse path π^{-1} of π. If $p_t^{-1} = \pi^{-1}(t)$ and t_0 is in $[0, 1]$, then Equation 1 implies

$$\begin{aligned} \lim_{t\to t_0+} co\text{-}kernel_2(p_t, P) &= \lim_{t\to(1-t_0)-} co\text{-}kernel_2(p_t^{-1}, P) \\ &= co\text{-}kernel_2(p_{1-t_0}^{-1}, P) = co\text{-}kernel_2(p_{t_0}, P) \end{aligned}$$

which proves the continuity of $co\text{-}kernel_2(p_t, P)$ in t_0.

We assume in the following that the vertices of P are in general position, that is, no three vertices of P are collinear. The claim for the degenerate case follows easily from our proof because of the continuity of the functions we consider.

So let $t_0 \in [0, 1]$. As we observed above, $co\text{-}kernel_2(p_t, P)$ is the intersection of the visibility polygons of the vertices $V(t)$ of $P \setminus vis(p_t)$, for $t \in [0, 1]$. Since P is a polygon and π is a polygonal path, there is a $t' < t_0$ such that $V(t)$ is a continuous function of t, for all $t \in (t', t_0)$. Let $V(t_0^-) = \lim_{t\to t_0-} V(t)$ (in the Hausdorff-metric) and $vis\text{-}int(V(t_0^-))$ the intersection of the visibility polygons of the vertices in $V(t_0^-)$, that is,

$$vis\text{-}int(V(t_0^-)) = \bigcap_{v\in V(t_0^-)} vis(v) \qquad \text{where } V(t_0^-) = \lim_{t\to t_0-} V(t).$$

Note that the vertices in $V(t_0^-)$ may be different from the vertices in $V(t_0)$. We prove Equation 1 in two Steps. First we show that

$$\lim_{t\to t_0-} co\text{-}kernel_2(p_t, P) = vis\text{-}int(V(t_0^-))$$

and then we show that

$$vis\text{-}int(V(t_0^-)) = co\text{-}kernel_2(p_{t_0}, P).$$

Step 1 Let $V(t) = \{v_1(t), \ldots, v_m(t)\}$ and $v_i(t_0) = \lim_{t\to t_0-} v_i(t)$, for $1 \le i \le m$. There are two reasons why $\lim_{t\to t_0-} co\text{-}kernel_2(p_t, P)$ may be different from $vis\text{-}int(V(t_0^-))$.

1. $\lim_{t\to t_0-} vis(v_i(t))$ may be different from $vis(v_i(t_0))$ and
2. $\lim_{t\to t_0-} \bigcap_i vis(v_i(t))$ may be different from $\bigcap_i \lim_{t\to t_0-} vis(v_i(t))$.

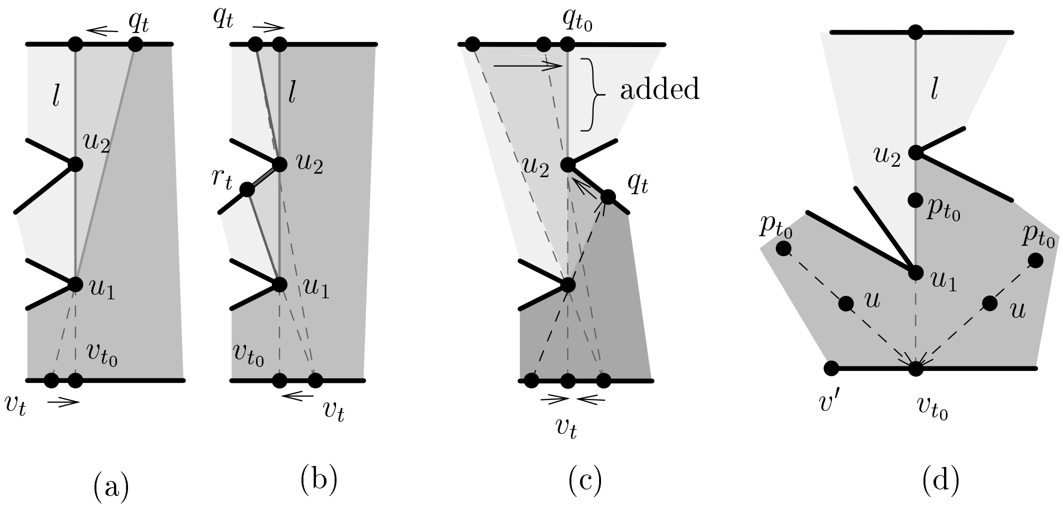

Figure 3: (a) The boundary of $vis(v_t)$ deforms continuously as v_t crosses l. (b) The line segment $\overline{u_2 q_{t_0}}$ is added. (c) There is a vertex of $V(t_0^-)$ that does not see $\overline{u_2 q_{t_0}}$

Since P is a closed set, it is easy to see that

$$\lim_{t \to t_0 -} vis(v_i(t)) \subseteq vis(v_i(t_0))$$

and

$$\lim_{t \to t_0 -} \bigcap_i vis(v_i(t)) \subseteq \bigcap_i \lim_{t \to t_0 -} vis(v_i(t)).$$

This implies that

$$\lim_{t \to t_0 -} co\text{-}kernel_2(p_t, P) = \lim_{t \to t_0 -} \bigcap_i vis(v_i(t)) \subseteq \bigcap_i vis(v_i(t_0)) = vis\text{-}int(V(t_0^-)).$$

So we only need to prove that $vis\text{-}int(V(t_0^-)) \subseteq \lim_{t \to t_0 -} co\text{-}kernel_2(p_t, P)$. Let q be a point in $vis\text{-}int(V(t_0^-))$. We show that q belongs to $\lim_{t \to t_0 -} co\text{-}kernel_2(p_t, P)$.

Step 1.1 As a first step we show that q belongs to $\lim_{t \to t_0 -} vis(v_i(t))$, for all $1 \leq i \leq m$. Let v_t be a vertex of $V(t)$ that depends continuously on t, for $t \in (t', t_0)$. Note that $\lim_{t \to t_0 -} vis(v_t) \neq vis(v_{t_0})$ only if v_{t_0} belongs a maximal line segment l in P that intersects two vertices u_1 and u_2 such that v_{t_0} does not belong to $\overline{u_1 u_2}$ and u_i is a reflex vertex such that both edges incident to u_i are on one side of l, for $i = 1, 2$ [11]. We assume that l is vertical, that u_1 is below u_2, and that the edges incident to u_1 are to the left of l. Depending on the location of the edges incident to u_2 we consider two cases.

Case 1 *Both edges incident to u_2 are also to the left of l.*

First assume that v_t crosses l from left to right at the time t_0 (see Figure 3a). Let q_t be the end point of the visibility ray $r_{v_t}(u_1)$, for $t \leq t_0$. If $t < t_0$, then q_t and u_1 are vertices of $vis(v_t)$. If $t = t_0$, then q_{t_0} and u_1 are also vertices of $vis(v_{t_0})$. Since the remaining vertices of $vis(v_t)$ depend continuously on t, this implies that $\lim_{t \to t_0 -} vis(v_t) = vis(v_{t_0})$.

Now assume that v_t crosses l from right to left at the time t_0 (see Figure 3b). Let q_t be the end point of the visibility ray $r_{v_t}(u_2)$ and r_t the end point of the visibility ray $r_{v_t}(u_1)$.

At $t = t_0$, u_2 and r_t become collinear with the window $\overline{u_1 q_{t_0}}$ and vanish as vertices of $vis(v_{t_0})$. Nevertheless, the boundary of $vis(v_t)$ depends continuously on t and again $\lim_{t \to t_0 -} vis(v_t) = vis(v_{t_0})$.

Case 2 *Both edges incident to u_2 are to the right of l.*

We first assume that v_t crosses l from left to right at the time t_0 (see Figure 3c). Let q_t be the end point of the visibility ray $r_{v_t}(u_1)$. At time t_0, the point q_t "jumps" from u_2 to q_{t_0}, that is,

$$\lim_{t \to t_0-} q_t = u_2 \neq q_{t_0}.$$

Hence, the line segment $\overline{u_2 q_{t_0}}$ is added to $vis(v_t)$ at time t_0. We claim that there are vertices of $\lim_{t \to t_0-} P \setminus vis(p_t)$ that do not see $\overline{u_2 q_{t_0}}$; hence, q does not belong to $\overline{u_2 q_{t_0}}$.

Since v_t is a moving vertex of $P \setminus vis(p_t)$, p_t sees v_t and there is a reflex vertex u that intersects $\overline{p_t v_t}$. Note that we can choose t' close enough to t_0 such that the reflex vertex u does not change, for all $t \in (t', t_0)$. We distinguish three cases according to the location of p_t (see Figure 3d). Note that since p_t and u are collinear, u belongs to the vertices $V(t)$ of $P \setminus vis(p_t)$ and, therefore, also to the vertices of $V(t_0^-)$.

If p_t is to the left of l, then u is below u_1 and u_1 blocks the view of u to $\overline{u_2 q_{t_0}}$. If u is to the right of l, then u is below u_2 and u_2 blocks the view of u to $\overline{u_2 q_{t_0}}$. In both cases $\overline{u_2 q_{t_0}}$ is not seen by a vertex of $\lim_{t \to t_0-} P \setminus vis(p_t)$.

Finally assume that p_t is located on l. In this case u equals either u_1 or u_2. Since u_1 is below u_2, the left end point v' of the edge that contains v_t is a vertex of $\lim_{t \to t_0-} P \setminus vis(p_t)$ and v' does not see $\overline{u_2 q_{t_0}}$. Hence, q does not belong to $\overline{u_2 q_{t_0}}$.

If v_t crosses l from right to left at the time t_0, then $\overline{u_2 q_{t_0}}$ belongs to $\lim_{t \to t_0-} vis(v_t)$ which implies that $\lim_{t \to t_0-} vis(v_t) = vis(v_{t_0})$.

Note that the argument does not change if v_{t_0} belongs to more than one such line segment l. Hence, q is contained in $\lim_{t \to t_0-} vis(v_i(t))$, for all $1 \leq i \leq m$.

Step 1.2 Now we show that q belongs to

$$\lim_{t \to t_0-} \textit{co-kernel}_2(p_t, P) = \lim_{t \to t_0-} \bigcap_{v \in V(t)} vis(v).$$

The proof is by contradiction. So assume that q does not belong to $\lim_{t \to t_0-} \textit{co-kernel}_2(p_t, P)$. Since $\lim_{t \to t_0-} \textit{co-kernel}_2(p_t, P)$ is a closed set, there is an $\varepsilon > 0$ such that the geodesic disc $B_\varepsilon(q)$ in P of radius ε does not intersect $\lim_{t \to t_0-} \textit{co-kernel}_2(p_t, P)$.

Since q belongs to $\bigcap_i \lim_{t \to t_0-} vis(v_i(t))$, all the vertices in $V(t_0^-)$ see q. There are some vertices $v_i(t_0)$ such that the line segment from $v_i(t_0)$ to q intersects a reflex vertex u_i that blocks the view of $v_i(t)$ to q, for $t \in (t_0 - \delta, t_0)$; otherwise, clearly, there is a $\delta > 0$ such that q belongs to $\bigcap_i vis(v_i(t))$, for $t \in (t_0 - \delta, t_0)$, and since the intersection $\bigcap_i vis(v_i(t))$ is a closed set, q also belongs to $\lim_{t \to t_0-} \bigcap_i vis(v_i(t))$.

Let I be the set of indices of vertices $v_i(t_0)$ such that $v_i(t_0)$ is collinear with a vertex u_i that blocks the view of $v_i(t)$ to q, for all $t \in (t_0 - \delta, t_0)$; moreover, let $w_i(t)$ be the window that is collinear with $\overline{v_i(t) u_i}$ and $R_i(t)$ the subpolygon of $w_i(t)$ that contains $v_i(t)$.

Since, for $i \in I$, $v_i(t_0)$ sees q but $v_i(t)$ does not see q, for $t \in (t_0 - \delta, t_0)$, the windows $w_i(t_0)$ intersect in q; moreover, since $w_i(t)$ rotates around u_i, there is a $0 < \delta' < \delta$, such that all the intersection points of the windows $w_i(t)$ are contained in $B_\varepsilon(q)$, for all $i \in I$ and all $t \in (t_0 - \delta', t_0)$. This implies that all the cells of the arrangement of the line segments $w_i(t)$ intersect $B_\varepsilon(q)$, and, therefore, all the cells of the arrangement of the regions $R_i(t)$ in P intersect $B_\varepsilon(q)$.

By an argument similar to Case 2 of Step 1.1 it is easy to see that there is at most one reflex vertex u_i that intersects the line segment $\overline{q v_i(t_0)}$. Hence, the part of $B_\varepsilon(q)$ that is visible to $v_i(t)$ increases monotonely as t approaches t_0. We now choose an $\varepsilon' > 0$ such that $B_{\varepsilon'}(q)$ intersects at most one window of $vis(v_i(t))$, for $t \in (t_0 - \delta', t_0)$ and change δ' such that it fulfills the same conditions for $B_{\varepsilon'}(q)$ as it does for $B_\varepsilon(q)$.

Since $B_{\varepsilon'}(q)$ intersects at most one window of $vis(v_i(t))$ for $t \in (t_0 - \delta', t_0)$, the intersection of $B_{\varepsilon'}(q)$ with $\bigcap_i R_i(t)$ equals the intersection of $B_{\varepsilon'}(q)$ with $\bigcap_i vis(v_i(t))$. As P is two-guardable,

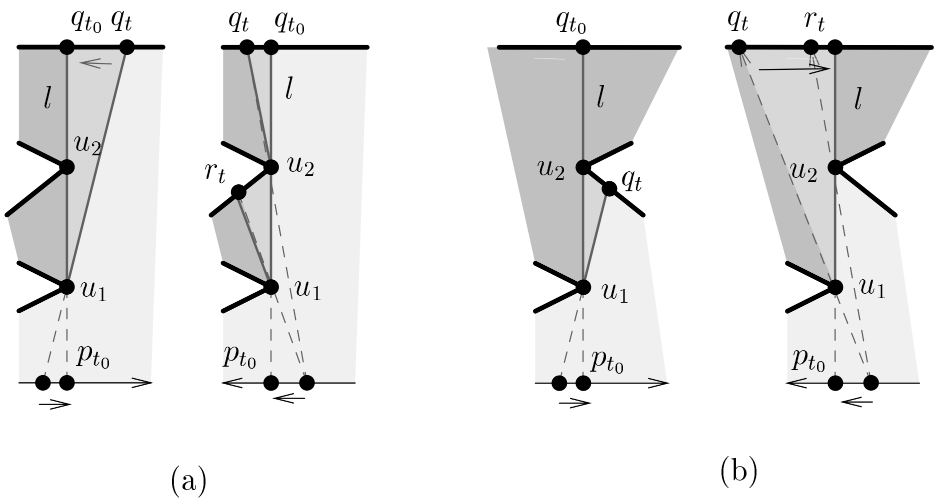

Figure 4: (a) The edges incident to u_2 are to the left of l. (b) The edges incident to u_2 are to the right of l.

this implies that $\bigcap_i R_i(t)$ is non-empty and there is a $t' \in (t_0 - \delta', t_0)$ such that the cell $C(t')$ of the arrangement that is contained in all regions $R_i(t')$ intersects $B_{\varepsilon'}(q)$. Since more and more points of $B_{\varepsilon'}(q)$ become visible to the vertices $v_i(t)$ as t approaches t_0, the intersection of $C(t')$ with $B_{\varepsilon'}(q)$ is a subset of $\bigcap_i R_i(t)$ and $\bigcap_i vis(v_i(t))$, for all $t \in (t_0 - \delta', t_0)$ and, therefore, also of $\bigcap_i vis(v_i(t_0))$—in contradiction to our choice of ε and ε'.

Hence, q belongs to $\lim_{t \to t_0 -}$ *co-kernel*$_2(p_t, P)$ and

$$vis\text{-}int(V(t_0^-)) \subseteq \lim_{t \to t_0 -} co\text{-}kernel_2(p_t, P).$$

Step 2 In the second step we need to show that

$$vis\text{-}int(V(t_0^-)) = co\text{-}kernel_2(p_{t_0}, P).$$

Consider $vis\text{-}int(V(t_0^-))$. As above we see that $\lim_{t \to t_0 -} V(t) \neq V(t_0)$ only if p_{t_0} belongs a maximal line segment l in P that intersects two vertices u_1 and u_2 such that p_{t_0} does not belong to $\overline{u_1 u_2}$ and both vertices u_1 and u_2 are reflex vertices such that the edges incident to v_i are on one side of l, for $i = 1, 2$. We again assume that l is vertical, that u_1 is below u_2, and that the edges incident to u_1 are to the left of l. Depending location of the edges incident to u_2 there are again two cases to be considered.

Case 1 *Both edges incident to u_2 are also to the left of l.*

First assume that p_t crosses l from left to right at the time t_0 (see Figure 4a). Let q_t be the end point of the visibility ray $r_{p_t}(u_1)$, for $t \leq t_0$. If $t < t_0$, then q_t, u_1, and u_2 are vertices of $P \setminus vis(p_t)$. If $t = t_0$, then q_{t_0}, u_1, u_2 are vertices of $V(t_0)$. Since the remaining vertices of $P \setminus vis(p_t)$ depend continuously on t, this implies that $V(t_0^-) = V(t_0)$.

Now assume that p_t crosses l from right to left at the time t_0. Let q_t be the end point of the visibility ray $r_{p_t}(u_2)$ and r_t the end point of the visibility ray $r_{p_t}(u_1)$.

At $t = t_0$, u_2 and r_t join into one vertex of $P \setminus vis(p_{t_0})$. Since the remaining vertices of $P \setminus vis(p_t)$ depend continuously on t, this implies again that $V(t_0^-) = V(t_0)$.

Case 2 *Both edges incident to u_2 are to the right of l.*

First assume that p_t crosses l from left to right at the time t_0 (see Figure 4b). Let q_t be the end point of the visibility ray $r_{p_t}(u_1)$. The points q_t, u_1, and u_2 are vertices of $P \setminus vis(p_t)$.

Note that $\lim_{t \to t_0 -} q_t = u_2 \neq q_{t_0}$, that is, at the time t_0 the point q_{t_0} becomes a new vertex of $P \setminus vis(p_{t_0})$. If q_{t_0} is a vertex of P, then q_{t_0} is also a vertex of $P \setminus vis(p_t)$, for $t < t_0$ and $\lim_{t \to t_0 -} V(t) = V(t_0)$. If q_{t_0} is not a vertex of P, then let e be the edge that contains q_{t_0}. The visibility polygon of q_{t_0} contains the intersection of the visibility polygons of the end points of e and, therefore, *co-kernel*$_2(p_{t_0}, P) \subseteq vis(q_{t_0})$. Hence, *co-kernel*$_2(p_{t_0}, P) = vis\text{-}int(V(t_0^-))$.

If p_t crosses l from right to left at the time t_0, then $\lim_{t \to t_0 -} q_t = q_{t_0}$ and again $V(t_0^-) = V(t_0)$.

Note that the arguments above do not change if v_{t_0} belongs to more than one such line segment l. This concludes the proof. □

In the above proof Step 1.2 may seem expendable. However, consider the following simple example. Let l_1 and l_2 be two line segments of the same length that intersect in one end point p. If we rotate l_1 around p until it equals l_2 at time t_0, then

$$\lim_{t \to t_0 -} l_1(t) \cap l_2(t) = p \neq l_1(t_0) = \lim_{t \to t_0 -} l_1(t) \cap \lim_{t \to t_0 -} l_2(t).$$

We need the following definition.

Definition 4.1 *Let P be a polygon. The* **geodesic center** *of P is the point in P which minimizes the maximum geodesic distance to the boundary of P.*

It can be shown that the geodesic center is unique [2]. Therefore, the geodesic center depends continuously on the boundary of P.

Lemma 4.9 *Let P be a continuous function from $[0,1]$ to the set of simple polygons. If, for every $t \in [0,1]$, $c(t)$ is the geodesic center of $P(t)$, then c is also a continuous function.*

Proof: Let $t_0 \in [0,1]$. We have to show that $\lim_{t \to t_0} c(t) = c(t_0)$. Let t_j a sequence with $\lim_{j \to \infty} t_j = t_0$. We show that $\lim_{j \to \infty} c(t_j) = c(t_0)$.

Since $P(t)$ depends continuously on t, $\bigcup_j P(t_j) \cup P(t_0)$ is a closed bounded set and there is a subsequence t_{j_l} of t_j such that $c(t_{j_l})$ converges to a point c. We first show that $c_{t_0} = c$.

Let $d(t)$ be the maximum distance of $c(t)$ to the boundary of $P(t)$ and $d = \lim_{l \to \infty} d(t_{j_l})$. If $d(t_0) = d$, then $c_{t_0} = c$ since the geodesic center is unique. If $d(t_0) < d$, say $d - d_{t_0} = \varepsilon$, then there is an $l > 0$ such that $H(P(t_{j_l}), P_{t_0}) \leq \varepsilon/4$. This observation implies that the distance of the closest point c'_{t_0} of $P(t_{j_l})$ to c_{t_0} to the boundary of $P(t_{j_l})$ is at most $d_{t_0} + \varepsilon/4$. On the other hand, the distance of $c(t_{j_l})$ to the boundary of $P(t_{j_l})$ is at least $d - \varepsilon/4$—a contradiction.

Hence, $c = c(t_0)$. If there is an $\varepsilon > 0$ such that there are infinitely many $c(t_j)$ outside $B_\varepsilon(c)$ in $\bigcup_j P(t_j) \cup P(t_0)$, then there is another subsequence of $c(t_j)$ that converges to a point c' different from c. But above we have shown that $c' = c(t_0)$—again a contradiction. Therefore, $\lim_{j \to \infty} c(t_j)$ exists and equals $c(t_0)$ as claimed. □

With the help of the geodesic center we are now able to construct a path π' for a given path π in two-*kernel*(P) such that $\{\pi(t), \pi'(t)\}$ guards P, for all $t \in [0,1]$.

Lemma 4.10 *Let P be a simple two-guardable polygon and $\{g_1, g_2\}$ and $\{h_1, h_2\}$ two two-guard sets of P. If there is a polygonal path π_1 from g_1 to h_1 in two-kernel(P), then there is a polygonal path π_2 from g_2 to h_2 such that $\{\pi_1(t), \pi_2(t)\}$ guards P, for all $t \in [0,1]$.*

Proof: We assume that π_1 is parameterized in such a way that, for $t \in [0, 1/4]$, $\pi_1(t) = g_1$ and, for $t \in [3/4, 1]$, $\pi_1(t) = g_2$.

For $t \in [1/4, 3/4]$ let $\pi_2(t)$ be the geodesic center of *co-kernel*$_2(\pi_1(t), P)$. Since the boundary of *co-kernel*$_2(\pi_1(t), P)$ is a continuous function of t by Lemma 4.8 and the geodesic center depends continuously on the boundary of *co-kernel*$_2(\pi_1(t), P)$, $\pi_2(t)$ is a path in P such that $\{\pi_1(t), \pi_2(t)\}$ guards P.

Clearly, if $\pi_1(t)$ moves on a polygonal path, then so do the vertices of the intersection of the visibility polygons of the vertices of $P \setminus vis(\pi_1(t))$, that is, the vertices of *co-kernel*$_2(\pi_1(t), P)$.

Since the geodesic center only depends on at most three vertices of $co\text{-}kernel_2(\pi_1(t), P)$ [12], it also moves on a polygonal path.

Let the geodesic center of $co\text{-}kernel_2(g_1, P)$ be $\overline{g}_2$ and the geodesic center of $co\text{-}kernel_2(h_1, P)$ be $\overline{h}_2$. Note that $\pi_2(1/4) = \overline{g}_2$ and $\pi_2(3/4) = \overline{h}_2$.

Consider the shortest path $shp(g_2, \overline{g}_2)$ from g_2 to $\overline{g}_2$. We assume that $shp(g_2, \overline{g}_2)$ is parameterized by length, that is, $shp(g_2, \overline{g}_2)(t)$ is t-times the length of $shp(g_2, \overline{g}_2)$, for $t \in [0, 1]$. Since $shp(g_2, \overline{g}_2)$ is contained in $co\text{-}kernel_2(g_1, P)$ by Lemma 4.7, we can define $\pi_2(t) = shp(g_2, \overline{g}_2)(4t)$, for $t \in [0, 1/4]$.

Similarly, we can define $\pi_2(t) = shp(h_2, \overline{h}_2)(4t - 3)$, for $t \in [3/4, 1]$ where we assume again $shp(h_2, \overline{h}_2)$ is parameterized by length. Hence, π_2 is a path from g_2 to h_2 such that, for all $t \in [0, 1]$, $\{\pi_1(t), \pi_2(t)\}$ guards P. □

If π_1 and π_2 are paths in a polygon P such that $\{\pi_1(t), \pi_2(t)\}$ guards P, then π_2 is called a **partner path** of π_1 and vice versa. As an interesting aside, Lemma 4.4 implies that the two$^+$-$kernel(P)$ and the two-kernel of a simple polygon P have the same number of components.

Theorem 4.11 *If P is a simple polygon, then $|two^+\text{-}kernel(P)| = |two\text{-}kernel(P)|$.*

Proof: If P is starshaped, then two$^+$-$kernel(P) = P \times P$ and two-$kernel(P) = P$; hence, the claim holds. Similarly, if P is not two-guardable, then two$^+$-$kernel(P)$ = two-$kernel(P) = \emptyset$. Therefore, we can assume that P is exactly two-guardable. Clearly, $|\text{two}^+\text{-}kernel(P)| \geq |\text{two-}kernel(P)|$. Let $\Pi_1 : I\!R^4 \longrightarrow I\!R^2$ be the projection onto the first two coordinates. As we observed before, $\Pi_1(\text{two}^+\text{-}kernel(P)) = \text{two-}kernel(P)$. Let C be a component of two-$kernel(P)$ and consider $C^{-1} = \Pi_1^{-1}(C) \cap \text{two}^+\text{-}kernel(P) = \{(g_1, g_2) \in P^2 \mid g_1 \in C, g_2 \in co\text{-}kernel_2(g_1, P)\}$. If two-$kernel(P)$ consists of the components $C_1, \ldots, C_m$, then, clearly, two$^+$-$kernel(P) = \bigcup_{i=1}^{m} C_i^{-1}$. We claim that C_i^{-1} consists of only one component, for $1 \leq i \leq m$. Consider two pairs (g_1, g_2) and (g'_1, g'_2) of points in C_i^{-1}. Since g_1 and g'_1 belong to the same component C_i of two-$kernel(P)$, there is a path π_1 from g_1 to g'_1 in C_i. By Lemma 4.10 there is a path π_2 in two-$kernel(P)$ from g_2 to g'_2 such that $\pi_2(t) \in co\text{-}kernel_2(\pi_1(t), P)$. By the definition of C_i^{-1}, we deduce that $(\pi_1(t), \pi_2(t)) \in C_i^{-1}$, for all $t \in [0, 1]$. Hence, (π_1, π_2) is a path in C_i^{-1} that connects (g_1,g_2) to (g'_1,g'_2) and both points belong to the same component of two$^+$-$kernel(P)$. □

One of the other applications of Lemma 4.10, which we investigate further, is to the situation when g_1 and g_2 belong to the same component of two-$kernel(P)$.

Corollary 4.12 *Let P be a simple two-guardable polygon and $\{g_1, g_2\}$ be a two-guard set of P. If there is a polygonal path π_1 from g_1 to g_2 in two-kernel(P), then there is a polygonal path π_2 from g_2 to g_1 such that $\{\pi_1(t), \pi_2(t)\}$ guards P, for all $t \in [0, 1]$.*

One useful property of a set of partnered paths that have opposite start and end points is that every point of P is intersected by the extension of the shortest path $shp^+(\pi_1(t), \pi_2(t))$, for some $0 \leq t \leq 1$. Before we prove this fact, we make a few simple observations about shortest paths between guards.

Lemma 4.13 *If P is a simple polygon and $\{g_1, g_2\}$ is a two-guard set of P, then $shp(g_1, g_2)$ intersects at most one reflex vertex.*

Proof: Assume that there are at least two reflex vertices v_1 and v_2 that are intersected by $shp(g_1, g_2)$ such that both edges incident to v_1 or to v_2 are to the same side of $shp(g_1, g_2)$. Clearly, neither g_1 nor g_2 can see the part of $shp(g_1, g_2)$ between the v_1 and v_2 which contradicts the fact that $\{g_1, g_2\}$ guards P (see Figure 5). □

An immediate consequence of Lemma 4.13 is that the shortest path between two guards has at most two line segments which we state as a corollary for further reference.

Corollary 4.14 *If $\{g_1, g_2\}$ is a guard set of P, then $shp(g_1, g_2)$ has at most two line segments.*

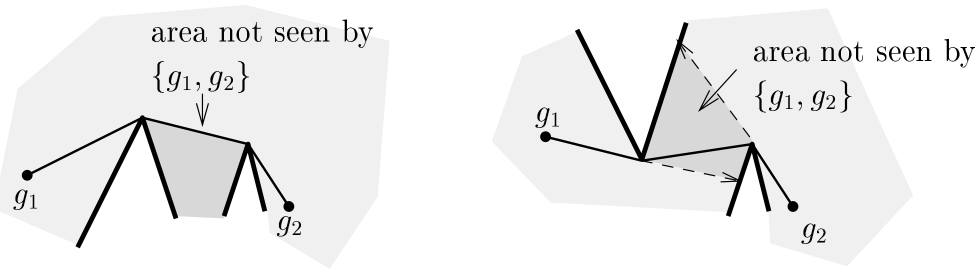

Figure 5: If $shp(g_1, g_2)$ intersects two reflex vertices, then $\{g_1, g_2\}$ does not guard P.

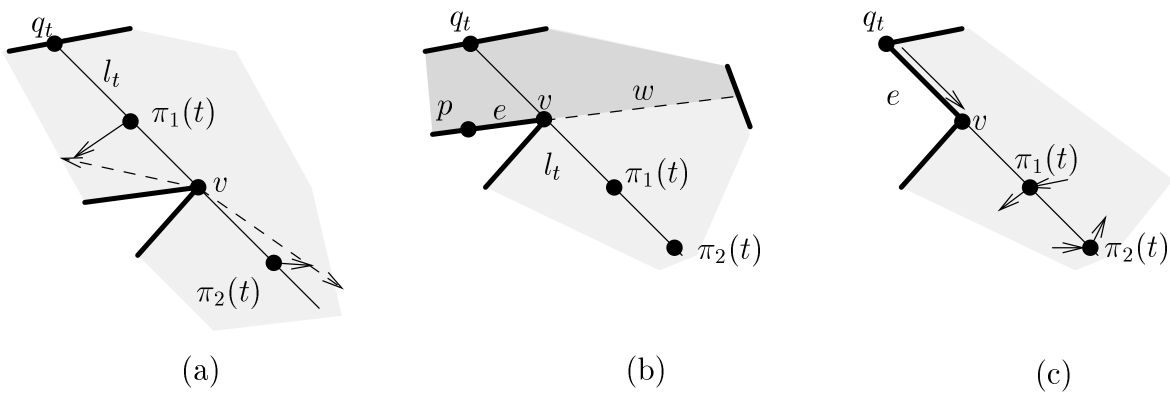

Figure 6: (a) l_t intersects a reflex vertex. Note that if we consider $\pi_1(1-t)$ and $\pi_2(1-t)$, that is, we revert the direction of π_1 and π_2, then l_t merges with the second line segment of $shp^+(\pi_1(t), \pi_2(t))$. (b) The part of l_t from q_t to $\pi_1(t)$ intersects a reflex vertex. (c) The edge e is collinear with l_t.

Lemma 4.15 *Let P be a simple two-guardable polygon, $\{g_1, g_2\}$ a two-guard set of P, and π_1 and π_2 a pair of polygonal paths that connects them in two-kernel(P) such that $\{\pi_1(t), \pi_2(t)\}$ guards P, for all $t \in [0, 1]$. If p is a point in P, then there is a $t \in [0, 1]$ such that $p \in shp^+(\pi_1(t), \pi_2(t))$.*

Proof: We first show that the end points of $shp^+(\pi_1(t), \pi_2(t))$ depend continuously on t. So let q_t be one end point of $shp^+(\pi_1(t), \pi_2(t))$, say the end point closer to $\pi_1(t)$. Since both $\pi_1(t)$ and $\pi_2(t)$ depend continuously on t, so do all points of $shp(\pi_1(t), \pi_2(t))$. Let l_t be the line segment of $shp^+(\pi_1(t), \pi_2(t))$ that q_t belongs to. Since l_t is just the extension of a line segment of $shp^+(\pi_1(t), \pi_2(t))$, q_t moves continuously unless an event happens that changes l_t discontinuously.

The events that change l_t may be either that l_t merges with the other line segment of $shp^+(\pi_1(t), \pi_2(t))$ into one line segment or that l_t intersects a reflex vertex v such that both edges incident to v are to one side of l_t (see Figure 6a). Clearly, in the first case the end point q_t still moves continuously. In the second case q_t also moves continuously if v intersects $shp(\pi_1(t), \pi_2(t))$ since in this case only the turn axis of l_t changes.

So assume that v does not belong to $shp(\pi_1(t), \pi_2(t))$. First assume that neither of the edges that are incident to v is collinear with $\pi_1(t)$ (see Figure 6b). We show that P is not completely seen by $\{\pi_1(t), \pi_2(t)\}$ in this case. Since neither edge that is incident to v is collinear with $\pi_1(t)$ and both edges are to one side of l_t, there is one edge e that is completely invisible to $\pi_1(t)$. We show that e is also invisible to $\pi_2(t)$. Let p be a point on the edge $e \setminus v$. Consider the visibility polygon $vis(p)$ of p. Clearly, there is one window w of $vis(p)$ that is collinear with e such that the pocket of w contains $\pi_1(t)$ (see Figure 6b). Since $shp^+(\pi_1(t), \pi_2(t))$ is a shortest path, it does not intersect w twice and, therefore the part of $shp^+(\pi_1(t), \pi_2(t))$ after $\pi_1(t)$ is also contained in the pocket of w and $\pi_2(t)$ does not see p. This contradicts the fact that $\{\pi_1(t), \pi_2(t)\}$ guards P.

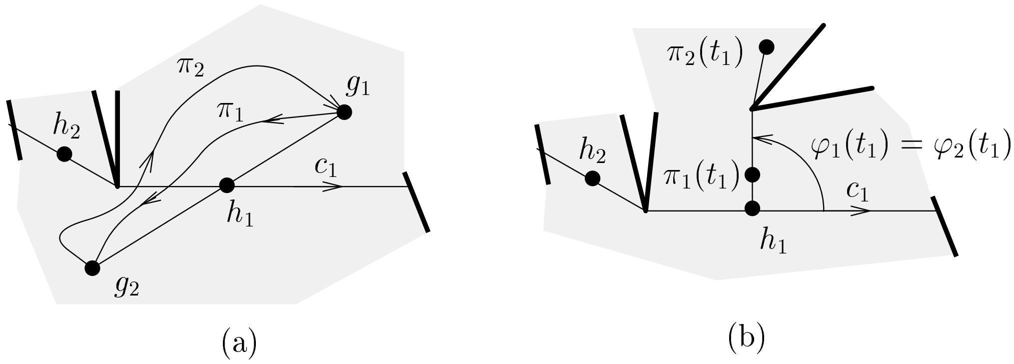

Figure 7: h_1 switches sides.

Hence, one edge e incident to v is collinear with $\pi_1(t)$ (see Figure 6c). In this case we assume that π_1 and π_2 are parameterized in such a way that the movement stops for some time $[t_1, t_2] \subseteq (0, 1)$ and q_t slides from one end point of e to the other in the interval $[t_1, t_2]$.

Hence, the mapping from t to q_t is continuous. Now consider the volume $P \times [0, 1]$. We parameterize $shp^+(\pi_1(t), \pi_2(t))$ such that, for $s \in [0, 1]$, the length of $shp^+(\pi_1(t), \pi_2(t))(s)$ is s times the length of $shp^+(\pi_1(t), \pi_2(t))$. Since both end points of $shp^+(\pi_1(t), \pi_2(t))$ depend continuously on t, the length and, therefore, $shp^+(\pi_1(t), \pi_2(t))(s)$ depends continuously on s and t. Hence, the mapping S which maps (s, t) to $shp^+(\pi_1(t), \pi_2(t))(s)$ is a continuous mapping, and S describes a surface in $\mathbb{R}^3$. The boundary of S is formed by the two curves $shp^+(\pi_1(0), \pi_2(0))$, $shp^+(\pi_1(1), \pi_2(1))$, and the two paths described by the end points of $shp^+(\pi_1(t), \pi_2(t))$ both of which are completely contained in the boundary of $P \times [0, 1]$. Therefore, S divides $P \times [0, 1]$ into a number of connected components and every path from one component to another intersects S. Note that all points in P that are to the left of $shp^+(\pi_1(0), \pi_2(0))$ are to the right of $shp^+(\pi_1(1), \pi_2(1))$. Hence, if p is a point in $P \setminus shp^+(q_1, q_2)$, then $(p, 0)$ is in one component of $(P \times [0, 1]) \setminus S$ and $(p, 1)$ in another. This implies that there is one $t \in [0, 1]$ such that (p, t) belongs to S. In other words, $p \in shp^+(\pi_1(1), \pi_2(1))$ as claimed. $\square$

We are now able to prove that when $|\text{two-}kernel(P)| > 1$, there is no component that contains its own partner.

Lemma 4.16 *If P is a simple polygon and $\{g_1, g_2\}$ and $\{h_1, h_2\}$ are two guard sets of P such that g_1 and g_2 belong to the same component C of two-kernel(P), then both h_1 and h_2 also belong to C.*

Proof: It suffices to show that one of h_1 or h_2 belongs to C since, by Lemma 4.4, all the points in C have their partners in the same component.

Since g_1 and g_2 belong to the same component of two-$kernel(P)$, there are two paths π_1 and π_2 connecting g_1 and g_2 in two-$kernel(P)$ such that $\{\pi_1(t), \pi_2(t)\}$ guards P, for $t \in [0, 1]$ by Lemma 4.10. By Lemma 4.15, there is a t_0 in $[0, 1]$ such that h_1 belongs to $shp^+(\pi_1(t_0), \pi_2(t_0))$. Now if h_1 does not belong to $shp(\pi_1(t_0), \pi_2(t_0))$, then either $\pi_1(t_0)$ belongs to $shp(h_1, \pi_2(t_0))$ or $\pi_2(t_0)$ belongs to $shp(h_1, \pi_1(t_0))$. Say, the latter case holds. Then, we have, by Lemma 4.1, that $\{\pi_2(t_0), h_2\}$ guards P. Since h_2 and $\pi_1(t_0)$ are both partners of $\pi_2(t_0)$, Corollary 4.3 implies that h_2 and $\pi_1(t_0)$ belong to the same component of two-$kernel(P)$; that is, $h_2 \in C$.

So assume that h_1 belongs to $shp(\pi_1(t_0), \pi_2(t_0))$. Without loss of generality, we may assume that $t_0 = 0$. Consider $shp^+(h_1, h_2)$ (see Figure 7a). By Corollary 4.14, it consists of at most two chords c_1 and c_2. We assume that c_1 contains h_1. Note that c_1 splits P into two parts one of which contains g_1 and the other contains g_2 since $h_1 \in shp(g_1, g_2)$. Hence, π_1 and π_2 cross c_1. Assume that there is some crossing point $\pi_1(t)$ of π_1 and c_1 that does not belong to $shp(h_1, h_2)$. Then, $h_1 \in shp(\pi_1(t), h_2)$ and, by Lemma 4.1, $\{h_1, \pi_2(t)\}$ guards P. Again since h_1 and $\pi_1(t)$ are both partners of $\pi_2(t)$, Corollary 4.3 implies that h_1 and $\pi_1(t)$ belong to C.

So from now on assume that all crossing points of π_1 and π_2 with c_1 belong to $shp(h_1, h_2)$. For ease of presentation we assume c_1 to be horizontal with $c_1 \cap \pi_1$ to the left of h_1 and g_1 above and g_2 below it (see Figure 7a).

Let $\varphi_i(t)$ be the cummulative angle of the first link of $shp(h_1, \pi_i(t))$ with c_1 where $\varphi_i(0)$ is defined as the angle of the first link of $shp(h_1, \pi_i(0))$ with c_1, for $i = 1, 2$ (see Figure 7b). Note that since neither π_1 nor π_2 intersects c_1 to the right of h_1, $\varphi_1([0,1]) \subseteq (0, 2\pi)$ and $\varphi_2([0,1]) \subseteq (0, 2\pi)$. In particular, $\varphi_1(0) = \varphi_2(1)$ and $\varphi_1(1) = \varphi_2(0)$. Hence, $\varphi_2(t) - \varphi_1(t)$ is a continuous function that assumes all values between $\varphi_2(0) - \varphi_1(0) > 0$ and $\varphi_2(1) - \varphi_1(1) < 0$. Therefore, there is a value t_1 for which $\varphi_1(t_1) = \varphi_2(t_1)$.

At the time t_1 either the shortest path from $\pi_1(t_1)$ to h_1 contains $\pi_2(t_1)$ or the shortest path from $\pi_2(t_1)$ to h_1 contains $\pi_1(t_1)$ (see Figure 7b). As we observed above, this condition implies that h_1 and h_2 belong to C. □

Theorem 4.17 *The number of components of the two-kernel of a simple polygon is either one or even.*

Proof: Let P be a simple polygon and assume that two-$kernel(P)$ consists of $m > 1$ components $C_1, C_2, \ldots, C_m$. By Lemma 4.4 there is one unique component C_j, for each component C_i, such that all partners of points in C_i are contained in C_j and vice versa. By Lemma 4.16, there is no component C_i with itself as a partnered component since $m > 1$. Hence, m is even and $\{C_1, C_2, \ldots, C_m\}$ can be partitioned into $m/2$ pairs. □

5 The Shape of the Two-Kernel

We now investigate various aspects of the shape of the two-kernel. In particular, we show that if $g(P) = 2$ and two-$kernel(P)$ has one component, then two-$kernel(P)$ contains a hole. But first we prove that two-$kernel(P)$ is a polygon.

Lemma 5.1 *Let P be a simple polygon. If l is a line segment contained in two-$kernel(P)$, then co-$kernel_2(l, P) = \bigcup_{p \in l}$ co-$kernel_2(p, P)$ is a polygon.*

Proof: Let p_0 be the one end point of l and p_1 the other and $p_t = (1-t)p_0 + tp_1$, with $0 \le t \le 1$.

Clearly, the vertices of $vis(p_t)$ describe piecewise linear curves on the boundary of P depending on t and, so do the vertices of the intersection of their visibility polygons which equals co-$kernel_2(p_t, P)$ as shown above.

As shown in Lemma 4.8 the topology of co-$kernel_2(p_t, P)$ only changes if either p_t or one of the vertices of $vis(p_t)$ intersects a line segment in P that is collinear with two vertices of P. Since the paths described by p_t, the vertices of $vis(p_t)$, and the vertices of co-$kernel_2(p_t, P)$ are line segments, there are only a finite number of these events. Hence, there is an $m \ge 1$ and numbers $0 = t_0 < t_1 < \cdots < t_m = 1$ such that no event occurs in (t_i, t_{i+1}), for $0 \le i \le m-1$, and the vertices of the polygon co-$kernel_2(p_t, P)$ move on straight line segments that may intersect but do not cross for $t \in [t_i, t_{i+1}]$, $0 \le i \le m-1$. Clearly, $\bigcup_{t \in [t_i, t_{i+1}]}$ co-$kernel_2(p_t, P)$ is a simple polygon P_i. Since the finite union of simple polygons is a polygon, this proves the claim. □

Now we can show that two-$kernel(P)$ is also a polygon.

Lemma 5.2 *If P is a simple polygon, then two-$kernel(P)$ is a polygon.*

Proof: To prove the claim it suffices to show that two-$kernel(P)$ is the union of finitely many co-kernels of line segments since co-kernels of line segments are polygons by Lemma 5.1. As mentioned earlier there is a set $\mathcal{L}$ of $O(n^4)$ line segments such that each component intersects a line segment in $\mathcal{L}$ [3]. We show that, for all $p \in$ two-$kernel(P)$, there is an $l \in \mathcal{L}$ such that $p \in$ co-$kernel_2(l, P)$.

So let p be a point in two-$kernel(P)$. There is some point q such that $co\text{-}kernel_2(q, P)$ contains p. Let V be some maximal starshaped polygon in P that contains $vis(q)$. Note that $kernel_P(P \backslash V)$ contains p. The kernel of V intersects an extension l of an edge by Lemma 3.3. Let l' be the maximal line segment in two-$kernel(P)$ that contains the intersection of l with $kernel(V)$. The point p obviously belongs to $co\text{-}kernel_2(l', P)$. Clearly, l' can be partitioned into line segments $l_1, \ldots, l_m$, with $l_i \in \mathcal{L}$. Since $co\text{-}kernel_2(l', P) = \bigcup_{i=1}^{m} co\text{-}kernel_2(l_i, P)$, there is one $1 \leq i \leq m$ such that $p \in co\text{-}kernel_2(l_i, P)$ which proves the claim. □

In order to show that two-$kernel(P)$ contains a hole, we need to investigate homotopies between partner paths.

Lemma 5.3 *Let P be a simple polygon and g_1 and g_2 be two partners in the same component of two-$kernel(P)$. If π_1 and π_2 are two partner paths of π, then there is a homotopy $\mathcal{H}$ from π_1 to π_2 such that, for all $(s,t) \in [0,1]^2$, $\mathcal{H}(s,t)$ is a partner of $\pi(t)$.*

Proof: Let π_1 and π_2 be two partner paths of π. We define the homotopy $\mathcal{H}$ from π_1 to π_2 as follows. For each $t \in [0,1]$ consider the shortest path $shp(\pi_1(t), \pi_2(t))$ from $\pi_1(t)$ to $\pi_2(t)$. We again parameterize $shp(\pi_1(t), \pi_2(t))$ by length. Since π_1 and π_2 are continuous and the parameterization of the shortest path by length depends continuously on the end points, it is easy to see that the above mapping of (s,t) to $shp(\pi_1(t), \pi_2(t))(s)$ is continuous. Moreover, since both $\pi_1(t)$ and $\pi_2(t)$ are partners of $\pi(t)$, Corollary 3.2 implies that, for all $t \in [0,1]$, all the points of $shp(\pi_1(t), \pi_2(t))$ are partners of $\pi(t)$ as claimed. □

Lemma 5.4 *Let P be a simple two-guardable polygon that is not starshaped. If two-$kernel(P)$ has only one component, then it contains a hole.*

Proof: We show that if two-$kernel(P)$ contains no hole, then P is starshaped. So assume that P is a 2-guardable polygon such that two-$kernel(P)$ consists of one component.

Let g_1 be a point in two-$kernel(P)$ and g_2 be a partner of g_1. Let π_1 be a polygonal path from g_1 to g_2 and π_2 a polygonal partner path of π_1 which exists by Lemma 4.10. Hence, for all $t \in [0,1]$, $\{\pi_1(t), \pi_2(t)\}$ guards P.

We first show that there are two homotopies $\mathcal{H}$ from π_1 to π_2 and $\mathcal{I}$ from π_2 to π_1 such that, for all $(s,t) \in [0,1] \times [0,1]$, $\{\mathcal{H}_s(t), \mathcal{I}_s(t)\}$ guards P.

First we observe that since two-$kernel(P)$ contains no hole, there exists a homotopy $\mathcal{H} : [0,1]^2 \longrightarrow I\!R^2$ from π_1 to π_2^{-1} such that $\mathcal{H}_0(t) = \pi_1(t)$ and $\mathcal{H}_1(t) = \pi_2^{-1}(t)$.[2] For reasons that will become apparent below we choose $\mathcal{H}$ in such a way that, for all $s \in [1/2, 1]$, $\mathcal{H}_s(t) = \pi_2^{-1}(t)$.

We now construct a homotopy $\mathcal{I}$ from π_2 to π_1^{-1} such that, for all $(s,t) \in [0,1]^2$, $\{\mathcal{H}_s(t), \mathcal{I}_s(t)\}$ guards P. Let π_s be the partner path for the path $\mathcal{H}_s$ as constructed in the proof of Lemma 4.10. First consider $t \in [1/4, 3/4]$. $\pi_s(t)$ is defined as the geodesic center of $co\text{-}kernel_2(\mathcal{H}_s(t), P)$. Since the boundary of $co\text{-}kernel_2(\mathcal{H}_s(t), P)$ depends continuously on s and t by Lemma 4.8 and the geodesic center depends continuously on the boundary of $co\text{-}kernel_2(\mathcal{H}_s(t), P)$, $\pi_s(t)$ is a continuous function of $(s,t) \in [0,1] \times [1/4, 3/4]$.

Now consider $t \in [0, 1/4]$. Let $\overline{g}_s = \pi_s(1/4)$. $\pi_s(t)$ is defined as $shp(g_2, \overline{g}_s)(4t)$. Since $\overline{g}_s$ depends continuously on s and, clearly, $shp(g_2, \overline{g}_s)$ depends continuously on $\overline{g}_s$, $\pi_s(t)$ is again a continuous function of s and t. A similar argument applies to $\pi_s(t)$ if $t \in [3/4, 1]$. Hence, if we define $\mathcal{I} : [0, 1/2] \times [0,1] \longrightarrow I\!R^2$ by $\mathcal{I}_s(t) = \pi_s(t)$, then $\mathcal{I}$ is also a homotopy from the partner path π_2 of π_1 to a partner path $\mathcal{I}_{1/2}$ of $\mathcal{H}_{1/2} = \pi_2^{-1}$. Note that $\mathcal{I}_{1/2}$ may be different from π_1^{-1}.

Since π_1^{-1} as well as $\mathcal{I}_{1/2}$ are both partner paths of π_2^{-1}, $\mathcal{I}$ can be extended as a homotopy for $1/2 \leq s \leq 1$ such that $\mathcal{I}_1 = \pi_1^{-1}$ by Lemma 5.3. Hence, $\mathcal{I}_0 = \pi_2$ and $\mathcal{I}_1 = \pi_1^{-1}$.

Summarizing $\mathcal{H}$ and $\mathcal{I}$ have the following properties.

(i) $\pi_1(t) = \mathcal{H}_0(t) = \mathcal{I}_1(1-t)$,

(ii) $\pi_2(t) = \mathcal{H}_1(1-t) = \mathcal{I}_0(t)$, and

[2] $\mathcal{H}_1(t) = \pi_2^{-1}(t)$ since the start point of π_2 is g_2 and the end point is g_1.

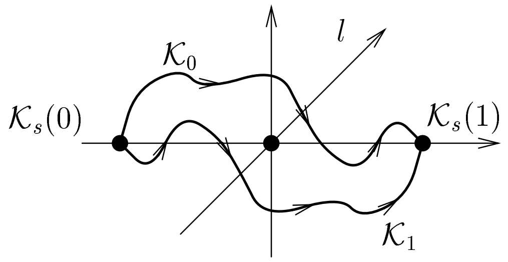

Figure 8: The line l intersects $\mathcal{K}_1([0,1])$ and $\mathcal{K}_0([0,1])$ on both sides of the origin an odd number of times.

(iii) $\{\mathcal{H}_s(t), \mathcal{I}_s(t)\}$ guards P, for all $(s,t) \in [0,1]^2$.

Property (iii) holds since $\mathcal{H}_s(t)$ and $\mathcal{I}_s(t)$ are partners by construction.

In the following we show that there is one $(s,t) \in [0,1]^2$ with $\mathcal{H}_s(t) = \mathcal{I}_s(t)$. By Property (iii) this implies that $p = \mathcal{H}_s(t)$ sees all of P and P is starshaped.

In order to see that there is one $(s,t) \in [0,1]^2$ with $\mathcal{H}_s(t) = \mathcal{I}_s(t)$ consider the continuous mapping $\mathcal{K}_s(t) = \mathcal{H}_s(t) - \mathcal{I}_s(t)$. We show that there is one (s,t) such that $\mathcal{K}_s(t) = (0,0)$.

We first observe that by Property (ii)

$$\mathcal{K}_0(t) = \mathcal{H}_0(t) - \mathcal{I}_0(t) = \mathcal{H}_0(t) - \mathcal{H}_1(1-t)$$

and

$$\begin{aligned}\mathcal{K}_1(t) &= \mathcal{H}_1(t) - \mathcal{I}_1(t) = \mathcal{H}_1(t) - \mathcal{H}_0(1-t) \\ &= -\left(\mathcal{H}_0(1-t) - \mathcal{H}_1(t)\right) = -\mathcal{K}_0(1-t).\end{aligned}$$

In particular, $\mathcal{K}_1([0,1]) = -\mathcal{K}_0([0,1])$ and

$$\begin{aligned}\mathcal{K}_s(0) &= \mathcal{H}_s(0) - \mathcal{I}_s(0) = g_1 - g_2 \\ \mathcal{K}_s(1) &= \mathcal{H}_s(1) - \mathcal{I}_s(1) = g_2 - g_1 = -\mathcal{K}_s(0).\end{aligned}$$

In the following we assume for convenience that $\mathcal{K}_s(0)$ and with it $\mathcal{K}_s(1)$ are contained in the x-axis. Consider a line l through the origin (see Figure 8). We claim that if l is non-horizontal, then l crosses $\mathcal{K}_0$ an odd number of times. To see this we note that the line l divides the plane into two halfplanes, one of which contains $\mathcal{K}_0(0)$ and the other contains $\mathcal{K}_0(1) = -\mathcal{K}_0(0)$. Since $\mathcal{K}_0$ is a path from $\mathcal{K}_0(0)$ to $\mathcal{K}_0(1)$, $\mathcal{K}_0$ crosses l an odd number of times.

The line l consists of two rays r_1 and r_2 that start at the origin. One of r_1 and r_2 crosses $\mathcal{K}_0$ an odd number of times and one crosses $\mathcal{K}_0$ an even number of times, say r_1 an odd and r_2 an even number. Since $\mathcal{K}_1([0,1]) = -\mathcal{K}_0([0,1])$, $\mathcal{K}_1([0,1])$ is obtained from $\mathcal{K}_0([0,1])$ by mirroring $\mathcal{K}_0([0,1])$ at the origin. Since r_2 is obtained by mirroring r_1 at the origin, r_2 intersects $\mathcal{K}_1([0,1])$ an odd number of times and r_1 an even number. Hence, every non-horizontal ray starting at the origin intersects the catenation of $\mathcal{K}_0([0,1])$ with $\mathcal{K}_1([0,1])$ an odd number of times.

This condition implies that the origin belongs to the area that is enclosed by $\mathcal{K}_0$ and $\mathcal{K}_1$. Since $\mathcal{K}$ is a homotopy from $\mathcal{K}_0$ to $\mathcal{K}_1$, there is a point $(s,t) \in [0,1]^2$ with $\mathcal{K}_s(t) = (0,0)$ which implies that P is starshaped. $\square$

Corollary 5.5 *If P is two-guardable, then two-kernel$(P) = P$ if and only if P is starshaped.*

Proof: The "if"-direction is obvious. On the other hand, if P is not starshaped, then by Lemma 5.4 two-$kernel(P)$ has at least one hole. Since P is a simple polygon, this implies that two-$kernel(P) \neq P$. $\square$

Another way of stating the preceding result is that the minimum guard-set of a two-guardable simple polygon P is never equal to P. It is interesting to note that there are multiply-connected

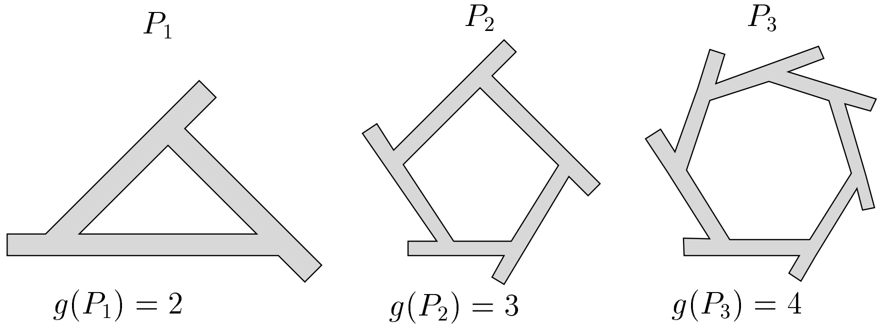

Figure 9: Examples of multiply-connected polygons with $kernel^*(P) = P$.

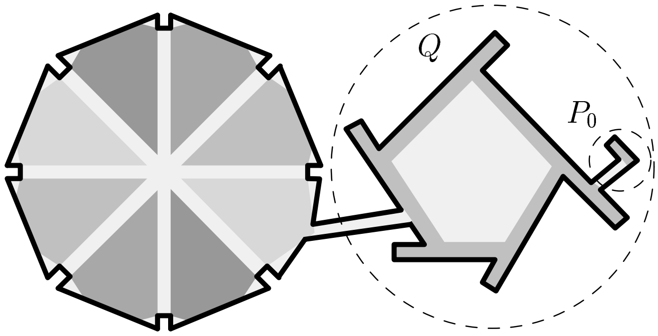

Figure 10: A 6-guardable polygon with ten components of $kernel^*(P)$.

polygons that coincide with their own guard set. In Figure 9 some examples are shown that are easily generalized, for any $k > 4$.

It should be noted that if the hole is removed from P_1, P_2, or P_3 in Figure 9, the two-kernel does not change. Indeed, for all $i \geq 1$, the polygons P_i can be transformed into simple polygons that are $(i+1)$-guardable and whose two-kernel has only one component.

6 More Than Two Guards

We now briefly discuss the number of components of the k-kernel when $k > 2$. We show that the techniques employed for the two-kernel do not generalize in a natural way. First we note that for $k \geq 4$ the k-kernel may have any number of components and may have $\Omega(n)$ components. Both of these claims follow from the lower bound examples of Figures 2 and 9. Let m be the number of components required. If m is even, then we take a starshaped polygon P_0 and combine it with the polygon P_{k-3} defined at the end of Section 5 (see Figure 10). Note that the resulting polygon Q is $(k-2)$-guardable and has a $(k-2)$-kernel with two components. Finally, we combine Q with the polygon displayed in Figure 2 that has a two-kernel with $m-2$ components. If m is odd, we directly combine P_{k-2} with the polygon displayed in Figure 2 whose two-kernel has $m-1$ components.

For $k = 3$, this construction yields polygons with an odd number of components. It is not known whether, for each $m > 0$, there is a three-guardable polygon such that the number of the components of its three-kernel is $2m$. For $m = 1$, this is easily seen to hold. For $m = 2$, a three-guardable polygon P whose kernel has four components is shown in Figure 11.

To see that the shaded regions in Figure 11(a) constitute the three-kernel of the polygon we make some observations about P. First note that to guard P there must be a guard in each one of the three regions R_1, R_2, and R_3 as well as in the slabs $S_1, \ldots, S_4$ as shown in Figure 11(b). Consider the dark shaded regions in Figure 11(c). If a guard is placed in one of these regions,

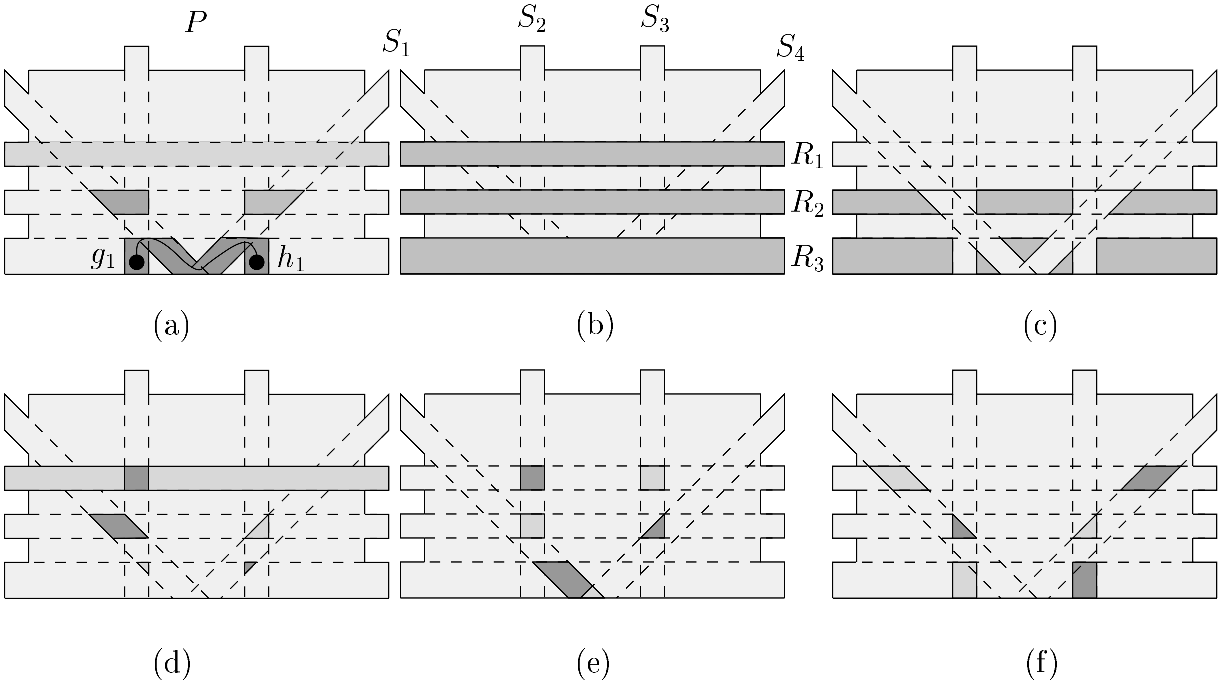

Figure 11: A three-guardable polygon P such that three-*kernel*(P) has four components.

it does not guard any of the slabs S_i. Note that a guard placed in R_1 can see at most one slab. Since no three slabs intersect in a common region in P, one additional guard cannot see all of the remaining three slabs. Hence, the three-kernel is no larger than the shaded regions in Figure 11(a).

To show that the shaded regions of Figure 11(a) do indeed belong to the three-kernel of P, we perform a detailed case analysis depending on where the first guard is placed. Since there has to be one guard in each of the three regions R_1, R_2, and R_3, we place g_1 in R_1, g_2 in R_2, and g_3 in R_3. In each region R_i the components of two-*kernel*(P) are considered from left to right.

Case 1 $g_1 \in R_1$ *(see lightly shaded regions in Figure 11d).* Place g_2 in $S_3 \cap S_4 \cap R_2$ and g_3 in $S_1 \cap S_2 \cap R_3$.

Case 2 $g_2 \in S_1 \cap R_2$ *(see dark shaded regions in Figure 11d).* Place g_1 in $S_2 \cap R_1$ and g_3 in $S_3 \cap S_4 \cap R_3$.

Case 3 $g_2 \in S_2 \cap R_2$ *(see lightly shaded regions in Figure 11e).* Place g_1 in $S_3 \cap R_1$ and g_3 in $S_1 \cap S_4 \cap R_3$.

Case 4 $g_2 \in S_3 \cap R_2$*.* This is symmetric to Case 3.

Case 5 $g_2 \in S_4 \cap R_2$*.* This is symmetric to Case 2.

Case 6 $g_3 \in S_1 \cap R_3$ *(see dark shaded regions in Figure 11e).* Place g_1 in $S_2 \cap R_1$ and g_2 in $S_3 \cap S_4 \cap R_2$.

Case 7 $g_3 \in S_2 \cap R_3$ *(see lightly shaded regions in Figure 11f).* Place g_1 in $S_1 \cap R_1$ and g_2 in $S_3 \cap S_4 \cap R_2$.

Case 8 $g_3 \in S_3 \cap R_3$ *(see dark shaded regions in Figure 11f).* This is symmetric to Case 7.

Case 9 $g_3 \in S_4 \cap R_3$ *.* This is symmetric to Case 6.

We have proved that the three-kernel of P has the regions displayed in Figure 11(a). The three-kernel of P also demonstrates that there is no immediate analog for Lemma 4.10, for $k \geq 3$.

For consider the path π in R_3 from g_1 to h_1. Note that as π crosses from S_1 into S_4, the partners of $\pi(t)$ suddenly have to "jump" from $S_2 \cap R_1$ to $S_3 \cap R_1$ and from $S_3 \cap R_2$ to $S_2 \cap R_2$ and there is no continuous path connecting them. This observation is the more surprising as $S_2 \cap R_1$ and $S_3 \cap R_1$ belong to the same component of three-$kernel(P)$. Note though that these point sets belong to different components of three$^+$-$kernel(P)$. Hence, it seems a new approach is needed to investigate k-kernels with $k \geq 3$.

7 Conclusions

We presented a generalization of the notion of a kernel of a starshaped polygon which we call the k-kernel of a polygon. The k-kernel is defined as the union of all sets of k points that guard the polygon. In particular, we investigated the properties of the two-kernel and showed that the number of its components is $O(n^4)$ if the polygon has n edges. Moreover, we gave a family of examples of polygons whose two-kernel has $\Omega(n)$ components. One of the interesting properties of the two-kernel is that its components can be paired in a natural way, which implies that the number of components of the two-kernel is either one or even, for all simple polygons.

Finally, we showed that if the two-kernel has only one component that is different from P, then it contains a hole. This fact implies that there is no simple nonstarshaped polygon P such that two-$kernel(P) = P$.

Numerous open questions remain. We showed that, for every $k > 3$ and every $i \in \mathbb{N}$, there is polygon P such that $|k\text{-}kernel(P)| = i$. As $\{|\text{two-}kernel(P)| : P \text{ is a simple polygon}\} = \{1\} \cup 2\mathbb{N}$, only the case $k = 3$ remains to be considered. It is easy to see that, for every odd number m, there is a polygon P with $|\text{three-}kernel(P)| = m$. Although we presented an example with $|\text{three-}kernel(P)| = 4$, it is open whether, for every even number m, there is a polygon P such that $|\text{three-}kernel(P)| = m$.

The main open question is whether there is a natural number $k > 2$ and a simple polygon P such that $g(P) = k$ and k-$kernel(P) = P$. We conjecture that such a pair does not exist. Also the gap between the upper and lower bound on the number of components of the two-kernel is still quite large and remains to be reduced. It should be noted that it is not known whether the problem of computing $g(P)$ is in NP. If it could be shown that there is a polynomial number of points in P that can be represented by a number of bits that is polynomial in the bits needed to represent P and that contain a k-guard set, then the problem is in NP. Of course, the questions considered here can also be asked for polygons with holes—which we have done partially—and for higher dimensional bodies.

References

[1] A. Aggarwal. *The Art Gallery Theorem: Its Variations, Applications, and Algorithmic Aspects.* PhD thesis, The Johns Hopkins University, Baltimore, Maryland, 1984.

[2] Te. Asano and G. Toussaint. Computing the geodesic center of a simple polygon. In D. S. Johnson, editor, *Discrete Algorithms and Complexity*, Perspectives in Computing, pages 65–74. Academic Press, 1987.

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

[4] Marilyn Breen. Clear visibility and the dimension of kernels of starshaped sets. *Proceedings of the American Mathematical Society*, 85:414–418, 1982.

[5] I. Bjorling-Sachs and D. Souvaine. An efficient algorithm for guard placement in polygons with holes. *Discrete & Computational Geometry*, 13(1):77–110, 1995.

[6] V. Chvàtal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory B*, 13(6):395–398, 1975.

[7] S. Fisk. A short proof of Chvàtal's watchman theorem. *Journal of Combinatorial Theory B*, 24:374, 1978.

[8] F. Hausdorff. *Grundzüge der Mengenlehre.* Verlag von Veit, Leizig, 1914.

[9] F. Hoffmann, M. Kaufmann, and Klaus Kriegel. The art gallery theorem for polygons with holes. In *Proc. 32nd IEEE Symposium on Foundations of Computer Science*, pages 39–48, 1991.

[10] D. T. Lee and A. K. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, IT-32:276–282, 1986.

[11] Joseph O'Rourke. *Art Gallery Theorems and Algorithms.* International series of monographs on computer science. Oxford University Press, 1987.

[12] R. Pollack, M. Sharir, and G. Rote. Computing the geodesic center of a simple polygon. *Discrete and Computational Geometry*, 4(6):611–626, 1989.

[13] F. A. Toranzos. Critical visibility and outward rays. *Journal of Geometry*, 33:155–167, 1988.

[14] X. Tan and B. Zhu. Two-guarding a rectilinear polygon. In *Proc. 2nd Conf. on Computing and Combinatorics*, number 1090 in LNCS, pages 391–400, 1996.