Searching for the Kernel of a Polygon: A Competitive Strategy Using Self-Approaching Curves*

Christian Icking** Rolf Klein** Elmar Langetepe**

Abstract

We present a competitive strategy for walking into the kernel of an initially unknown star-shaped polygon. From an arbitrary start point, s, within the polygon, our strategy finds a path to the closest kernel point, k, whose length does not exceed $5.3331\ldots$ times the distance from s to k. This is complemented by a general lower bound of $\sqrt{2}$. Our analysis relies on a result about a new and interesting class of curves which are *self-approaching* in the following sense. For any three consecutive points a, b, c on the curve the point b is closer to c than a to c. We show a tight upper bound of $5.3331\ldots$ for the length of a self-approaching curve over the distance between its endpoints.

Keywords. Competitive strategy, on-line strategy, simple polygon, kernel, curves with increasing chords, self-approaching curves, geometric optimization.

1 Introduction

Suppose that a mobile robot equipped with a 360 degree vision system "wakes up" in an unknown environment. Its task is to go to some location from which the whole environment is visible. This should be achieved as efficiently as possible.

In attacking this problem we are using the following model. The robot is a point, its environment is a simple polygon, P. Clearly, we have to assume that P is star-shaped, or no location with the required property would exist; see Section 2 for definitions. At each point u, the robot is provided with the visibility polygon $\mathrm{vis}(u)$. The cost of the robot's motion is measured by the arc length of its path; we ignore the cost of planning the path, which will in fact turn out to be negligible.

If the robot knew a map of the polygon and its start position beforehand, it could employ one of the classic algorithms for computing the kernel, see [2, 16]. Then it could determine the kernel point, k, closest to its start position, s, and go straight from s to k; note that by definition each point of the kernel can see any other point of P, so that the line segment from k to s cannot be blocked. This would result in a

*This work was partially supported by the Deutsche Forschungsgemeinschaft, grant Kl 655/8-2. A preliminary version has appeared in the proceedings of the 11th Annual ACM Symposium on Computational Geometry.

**FernUniversität Hagen, Praktische Informatik VI, Informatikzentrum, 58084 Hagen, Germany, {Christian.Icking,Rolf.Klein,Elmar.Langetepe}@fernuni-hagen.de

perfect solution at cost $d(k, s)$, where d denotes the Euclidean distance.

But since the polygon P is not known to the robot a priori, its actions are bound to contain elements of try and error which cause extra cost. For example, in the situations depicted in Figure 1 a detour cannot be avoided; see Lemma 5.1.

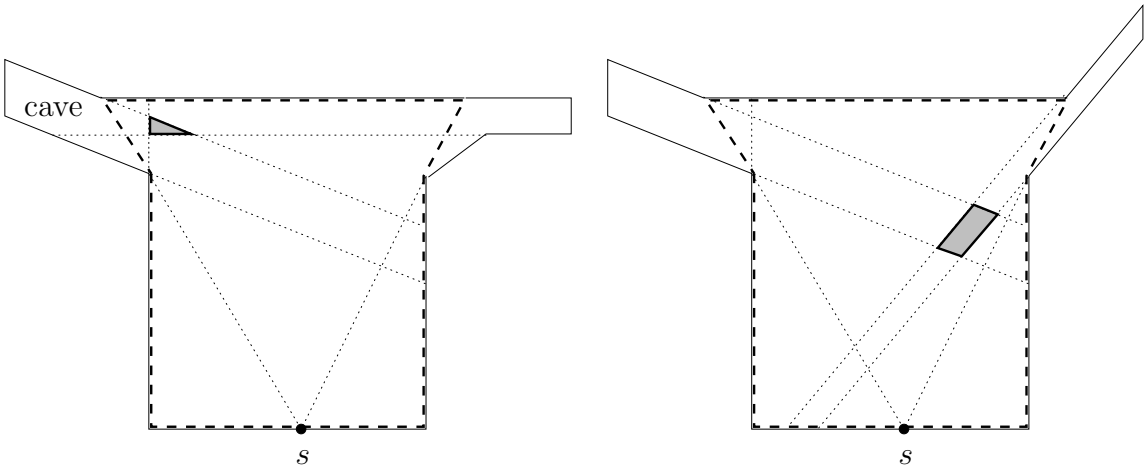

Figure 1: Identical visibility polygons but completely different kernels.

In this paper we present a strategy that guarantees a path length bounded by $5.3331\ldots$ times $d(k, s)$. In general, a strategy for a problem class PC is called *competitive* with competitive factor C if each instance P of PC can be solved at a cost not exceeding C times the cost of a perfect solution of P with full information in advance.

Competitive strategies for autonomous robots have recently received considerable interest in computational geometry. For example, the problem of finding a goal in an unknown environment has been studied in [5, 9, 11, 13, 15, 17]. In [6, 7, 10] the task of drawing a map is addressed. How to decide on the robot's current location if a map is available has been investigated in [8, 14]. A special case of the present problem, where P contains only one reflex vertex, has been optimally solved in [12].

Competitive strategies offer many advantages. They are superior to heuristic approaches because they come with a proven performance guarantee. Though, they can often be made as simple as heuristic rules, as the present paper demonstrates. Sometimes competitive strategies exist even for problems whose optimal solution would be NP-hard; see [7, 8]. However, analyzing a competitive strategy is often difficult so that the proven bound is considerably bigger than the real competitive factor.

This paper is organized as follows. In Section 2 we briefly recall some definitions and state a simple fact on visibility polygons. Then, in Section 3, we extract a criterion for approaching the kernel and describe our strategy CAB (Continuous Angular Bisector) for finding the closest kernel point. The strategy is very easy to implement; at each time an angular bisector between two vertices is chosen as direction. If the polygon P is not star-shaped, the robot notices this and ends its motion. Otherwise it will arrive at the closest kernel point, k.

In Section 4 we introduce the class of *self-approaching* curves, i. e. oriented planar curves that will always get closer to all their future points with respect to the orien-

tation. We characterize the self-approaching property through normals to the curve and show that paths generated by CAB are self-approaching. Then we prove that for self-approaching curves there is a tight upper bound of $5.3331\ldots$ for their *detour*, i. e. the ratio of their arc length divided by the distance between their endpoints. Therefore, we have the same value as an upper bound for the competitive factor of strategy CAB. Curves that are self-approaching in *both* directions have been called *curves with increasing chords*. In [18], Rote has solved an old open problem mentioned as problem G3 in [4] by proving a tight bound for this subclass of self-approaching curves. Unfortunately, it seems that his technique cannot be applied to our non-symmetric case. Our proof for the upper bound is based on the observation that the arc length of a self-approaching curve does not exceed the perimeter of its convex hull. Finally, in Section 5, we prove that there is a lower bound of $\sqrt{2}$ for the competitive factor of *any* possible strategy for finding the kernel.

2 Visibility in a polygon

In this section we briefly recall some elementary facts on visibility in simple polygons. As before, let P denote a simple planar polygon, interior area plus boundary.

Definition 2.1 Let u, v be two points in P. Then v is *visible* from u if the line segment $\overline{u\,v}$ is contained in P. The set of all points of P that are visible from u is called the *visibility polygon* of u, denoted by $\mathrm{vis}(u)$ (sometimes $\mathrm{vis}_P(u)$ to avoid ambiguity).

Definition 2.2 The *kernel* of P, $\ker(P)$, is the set of all points in P from which each point of P is visible. If $\ker(P)$ is not empty, then P is called *star-shaped*.

Clearly, visibility is a symmetric relation. Only for convex polygons does $P = \ker(P)$ hold. Otherwise the polygon has at least one *reflex vertex*, i. e. one whose internal angle is greater than $180°$, and the kernel is a proper subset of P which can be obtained in the following way. Each edge e of P defines two halfplanes, an inner one which locally contains points of the interior of P, and an outer one. The kernel of P is known to be the intersection of all inner halfplanes. In particular, the kernel is convex. In Figure 1, the kernels are represented by shaded areas.

Each visibility polygon $\mathrm{vis}(u)$ is star-shaped and contains u in its kernel. Its boundary consists of (segments of) original edges of P, and of edges that do not belong to P. These *spurious* edges are segments of lines emanating from u that touch a reflex vertex of P before they hit the boundary of P. A spurious edge separates the part of P that is visible from u from a part which is not, a so-called *cave*; see Figure 1, where the visibility polygons are depicted by dashed lines.

In the two pictures of Figure 1 the visibility polygons are identical, whereas the kernels are quite different. Thus, from point s the robot has only little information about the kernel of P. But from any point in P, one can at least identify the kernel of the actual visibility polygon! The following lemma shows how these sets are related: the kernel of the visibility polygon shrinks as visibility increases.

Lemma 2.3 *Let u, v be points in P such that* $\mathrm{vis}(v) \subseteq \mathrm{vis}(u)$. *Then we have the following inclusions.*

$$\ker(P) \subseteq \ker(\mathrm{vis}(u)) \subseteq \ker(\mathrm{vis}(v))$$

Proof. Let x be a point in $\ker(P)$, and let us set $P' = \mathrm{vis}(u)$.

For the first inclusion we want to show that $x \in \ker(\mathrm{vis}(u))$, i. e. x is in P' and $\mathrm{vis}_{P'}(x) = P'$. The point x can see all of P, in particular the point u, so $x \in \mathrm{vis}(u) = P'$ and for some arbitrary point $y \in P'$ we know that the line segment $\overline{x\,y}$ is contained in P. No edge of P' can intersect $\overline{x\,y}$, neither the spurious edges because otherwise one of x and y would lie in a cave of P invisible from u, nor the other edges, of course. So $\overline{x\,y} \subset P'$ for any point $y \in P'$, which means that x can see, from within the polygon P', all of P', in other words $x \in \ker(P') = \ker(\mathrm{vis}(u))$.

For the second inclusion we apply the first inclusion for P' and v. By assumption

$$v \in \mathrm{vis}(v) \subseteq \mathrm{vis}(u) = P'$$

holds, therefore the application is justified and we conclude $\ker(\mathrm{vis}(u)) \subseteq \ker(\mathrm{vis}_{P'}(v))$.

To complete the proof we show $\mathrm{vis}_{P'}(v) = \mathrm{vis}(v)$. The inclusion $\mathrm{vis}_{P'}(v) \subseteq \mathrm{vis}(v)$ holds since a line segment $\overline{v\,x}$ in P' lies inside $P \supseteq P'$. The reverse inclusion is also fulfilled: If $x \in \mathrm{vis}(v)$ holds then by assumption

$$\overline{v\,x} \subseteq \mathrm{vis}(v) \subseteq \mathrm{vis}(u) = P'$$

is true and also $x \in \mathrm{vis}_{P'}(v)$ holds. Altogether the second inclusion is true. □

In the next sections we will use the benefits of Lemma 2.3. We prefer strategies that increase the visibility polygon and under application of Lemma 2.3 we conclude that they approach the kernel successively.

3 A strategy for the kernel problem

Throughout this section we assume that P is a star-shaped simple polygon.

3.1 A criterion for approaching the kernel

At the start point, s, the robot may check whether $\mathrm{vis}(s)$ contains spurious edges. If not, it can conclude that s lies in the kernel of P, and stop. Otherwise, the robot should proceed from s on such a path that it gains insight into all the caves, and never lets out of sight points of P it has already seen, so that the visibility polygon grows monotonically.

The robot can only gain insight into a cave if it walks into the inner halfplane defined by the cave's spurious edge, causing the edge to rotate about its supporting reflex vertex *into* the cave. The robot can only keep an eye on what it has so far seen if, in addition, it never leaves the inner halfplane of a visible edge of P; in particular, it must stay within P. This leads to the following definition. We assume that the intersection of an empty set of halfplanes equals the full plane.

Definition 3.1 Let u be a point in P.

(i) Let $G(u)$ denote the *gaining wedge* that results from intersecting the inner halfplanes of the spurious edges of $\text{vis}(u)$. The reflex vertices of P associated with the two halflines bounding $G(u)$ are called the *maximum constraint vertices* at u.

(ii) Let $K(u)$ denote the *keeping wedge* that results from intersecting the inner halfplanes of all edges of P that contain u, or are visible from and collinear with u.

A point u belongs to the kernel of P if and only if $G(u)$ equals the full plane. For almost all points of P the keeping wedge $K(u)$ is equal to the full plane.

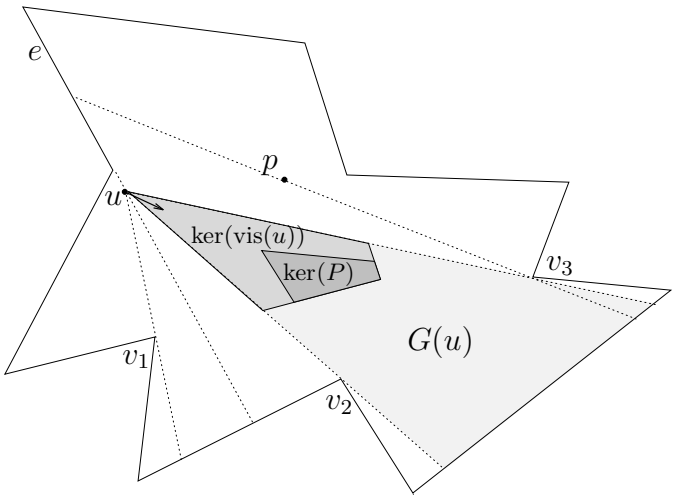

Figure 2: The set $G(u) \cap K(u)$ equals $\ker(\text{vis}(u))$ in the neighborhood of u and the kernel of P is contained in $\ker(\text{vis}(u))$.

In Figure 2 for point u there are two maximum constrained vertices, v_2 and v_3. For point p there is only one maximum constraint vertex, namely v_3, because this is the only visible reflex vertex that causes a cave with respect to p. At u the keeping wedge, $K(u)$, is the inner halfplane of the supporting line of edge e.

The next lemma shows that walking continuously into $G(u) \cap K(u)$ is always possible and leads to a strategy for increasing $\text{vis}(u)$.

Lemma 3.2 *Within a sufficiently small neighborhood of a point $u \in P$ the set $G(u) \cap K(u)$ equals $\ker(\text{vis}(u))$ and is nonempty. If the robot moves into this set then $\text{vis}(u)$ grows or it remains the same if u belongs to $\ker(P)$.*

Proof. The set $\ker(\text{vis}(u))$ is defined by the intersection of three types of halfplanes. First we have the inner halfplanes of spurious edges of P which also define the gaining wedge, $G(u)$. Second we have the inner halfplanes of all edges of P that contain u, or are visible from and collinear with u which are those that also define the keeping wedge, $K(u)$. The remaining halfplanes bounding $\ker(\text{vis}(u))$ stem from supporting lines of edges which do not intersect a sufficiently small neighborhood of $u \in P$. For an example see Figure 2.

Due to Lemma 2.3, $\ker(\text{vis}(u))$ contains $\ker(P)$, therefore it is nonempty. By construction of $G(u) \cap K(u)$, the set $\text{vis}(u)$ grows, unless u belongs to $\ker(P)$. □

From Lemma 2.3 and Lemma 3.2 we conclude that a process of stepping continuously into $G(u) \cap K(u)$ causes $\ker(\text{vis}(u))$ to shrink while $\ker(P)$ is always contained in it (see Figure 2), such that eventually $\ker(\text{vis}(u))$ shrinks to $\ker(P)$.

3.2 The CAB strategy

In the previous section we have seen that for increasing visibility it seems reasonable to move into the gaining wedge $G(u)$ while not leaving the keeping wedge $K(u)$. Our idea is to continuously follow the angular bisector of $G(u)$ or, if this is not contained in $K(u)$, the nearest possible direction in $K(u)$. The resulting strategy is called CAB (Continuous Angular Bisector) and is formally defined as follows.

Strategy CAB
  $u := s$;
  while $\text{vis}(u) \neq P$ **do**
    compute gaining wedge $G(u)$;
    $m :=$ angular bisector of $G(u)$;
    if m leaves the keeping wedge $K(u)$
      then walk in direction of the projection of m along the boundary of $K(u)$
      else walk along m
    end-if
  end-while;
  go straight to the point $k \in \ker(P)$ closest to s
end

This strategy is meant to be continuous, so the direction is adjusted very frequently as compared to the robot's speed; see Figure 3 for an example of a complete path generated by CAB.

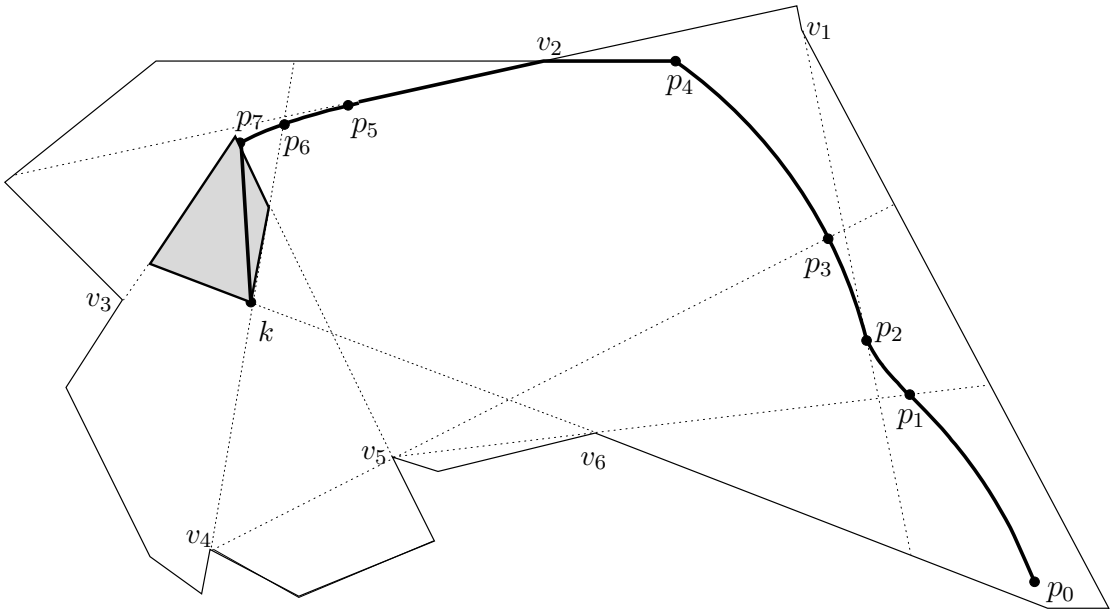

Figure 3: A path generated by strategy CAB.

Figure 4 illustrates a case where the **then**-branch applies. The following lemma shows why the projection of the angular bisector, m, is well-defined in this case.

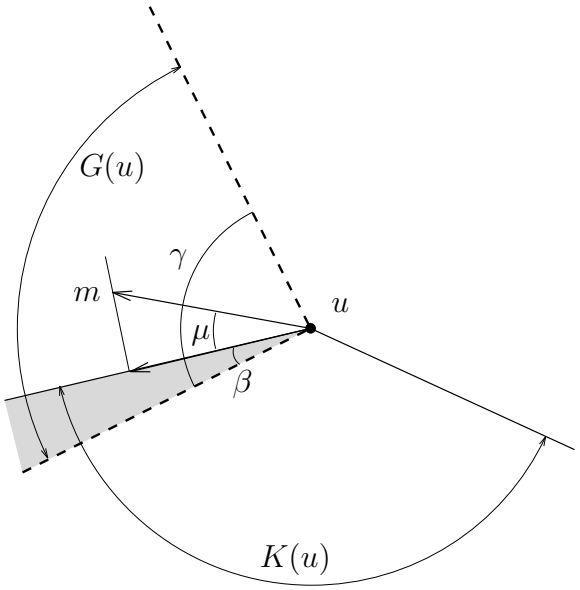

Figure 4: If the angular bisector m of the gaining wedge, $G(u)$, leaves the keeping wedge, $K(u)$, the robot walks along the boundary of $K(u)$.

Lemma 3.3 *If the angular bisector m of $G(u)$ leaves the wedge $K(u)$ then it forms an angle $\mu < 90°$ with exactly one of the two bounding halflines of $K(u)$. Therefore, the projection onto $K(u)$ is well-defined. Moreover, $G(u) \cap K(u)$ includes an angle $\beta \leq 90°$.*

Proof. Using the notations of Figure 4, we have $\mu = \gamma/2 - \beta \leq 90° - \beta \leq 90°$. Equality could hold for *both* halflines bounding $K(u)$ only if $G(u)$ and $K(u)$ are both $180°$ wedges with disjoint interiors. Since $\ker(P)$ is contained in their intersection, this means that u can already look into the cave(s) responsible for $G(u)$. This is a contradiction to Definition 3.1 because an edge is spurious only if it separates a non-empty part of P from $\text{vis}(u)$. Finally, we have $\beta \leq 90°$ because μ is bigger than 0, by assumption. □

Figure 3 shows an example where the robot walks along boundaries of keeping wedges. Between p_3 and p_4, the maximum constraint vertices are v_4 and v_5. At point p_4, it hits the extension of the upper edge incident to reflex vertex v_2. The robot follows this extension to v_2. Here, $K(u)$ is defined by both edge extensions. Now the robot follows the extension of the lower edge incident to v_2 until it arrives at p_5. Still, v_4 and v_5 are the maximum constraint vertices. From p_5 on, the robot resumes following the angular bisector which no longer crosses the edge extension.

As soon as the robot arrives in the kernel of P at p_7, it knows the whole polygon. According to CAB, it then determines which kernel point k would have been closest to its start point, s, and walks to k within $\ker(P)$. In practice, one might skip the last step because the robot's job can be considered done as soon as it reaches *some* point of the kernel. Since this does not lead to a smaller competitive factor, we can as well account for the last step, too.

Lemma 3.4 *The path M generated by strategy CAB consists of segments of hyperbolae and ellipses.*

Proof. For two maximum constraint vertices v_1 and v_2 the two cases depicted in Figure 5 may arise.

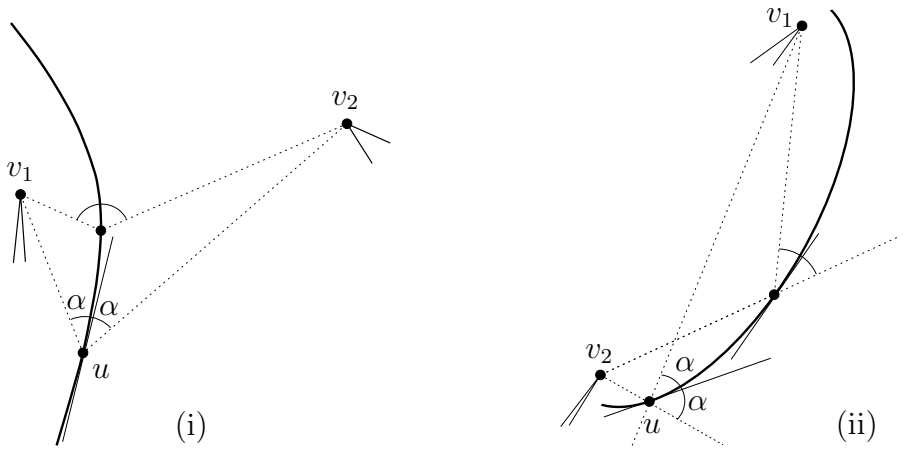

Figure 5: Following the angular bisector generates (i) hyperbolae or (ii) ellipses.

In (i), the boundary of the gaining wedge $G(u)$ of a point u consists of halflines originating from u that pass through vertices v_1 and v_2. In this case CAB describes a curve which is a unique solution of an ordinary differential equation which turns out to be a hyperbola with focus points v_1 and v_2. While following this hyperbola, the difference $d(p, v_1) - d(p, v_2)$ remains constant. If the difference equals 0 one obtains a line, as a special case (line segments are also generated while the robot follows the boundary of a keeping wedge).

In (ii), the gaining wedge $G(u)$ of a point u is formed by one halfline originating from u that runs through its reflex vertex, v_1, and another one, whose prolongation beyond u runs through v_2. In this case CAB describes a piece of an ellipse with focus points v_1 and v_2. Here, the sum $d(p, v_1) + d(p, v_2)$ remains constant.

A circle arises as a special case if there is only one constraint vertex. □

Lemma 3.5 *The number of hyperbolic and ellyptic segments the path M generated by CAB consists of is $O(n)$, where n denotes the number of vertices of P.*

Proof. There are three possibilities for a segment, as described in Figure 5, to end.

First, a maximum constrained vertex may be newly discovered or may get fully explored. For each vertex this happens at most once.

A segment also ends if the strategy begins to follow the boundary of the keeping wedge. In this case the robot follows an edge of P or the prolongation of an edge of P. The robot leaves this line either for the supporting line of another edge or for a piece of a hyperbola or ellipse. In the second case the path can only return to the same supporting line after a maximum constrained vertex has changed as described in the previous paragraph, due to the convex form of the segments and because its starting direction is flush with the line.

Finally a piece of the path ends if the kernel is reached; this occurs only once. Altogether the total number of segments in M is in $O(n)$. □

4 Self-approaching curves

In the first part of this section, we introduce the notion of self-approaching curves. Then we show that a path created by CAB belongs to this class. Finally, we prove a *tight* upper bound for the length of a self-approaching curve divided by the distance of its endpoints.

4.1 Definitions and properties

The curves considered here are assumed to be piecewise smooth curves in the plane, as is the case for all curves generated by strategy CAB, due to Lemma 3.4 and Lemma 3.5. For a curve C and a point a inside a smooth piece of C, the *tangent* to C at a and the *normal* to C at a, which is perpendicular to the tangent, are uniquely determined. Let a be a point of C such that two smooth pieces of C meet at a. The two normals N_1 and N_2 to the corresponding smooth pieces at a define a set of lines (see Figure 6), each line of this set is regarded as a normal to C at a.

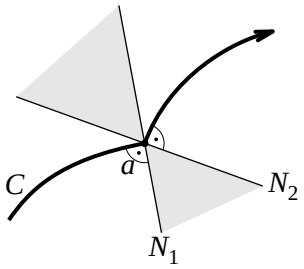

Figure 6: The bundle of lines defined as normals to C at a.

As before, let $d(a, b)$ denote the Euclidean distance between two points a and b. For two points $a \leq b$ ($a < b$) on a directed curve C, $C^{\geq a}$ ($C^{>a}$) denotes the part of C from a to the end (without a), $C[a, b]$ means the part of C between a and b, and $\text{length}(C[a, b])$ means its arc length.

Definition 4.1 An oriented curve is called *self-approaching* if the inequality

$$d(a, c) \geq d(b, c)$$

is fulfilled for any three consecutive points a, b, c on the curve.

Let C be an oriented curve from a to b. Then the quantity

$$\frac{\text{length}(C[a, b])}{d(a, b)}$$

is called the *detour* of a curve from a to b.

The following lemma shows that the self-approaching property is equivalent to the fact that for any point a on the curve the rest of the curve lies fully on one side of any normal to C at a, we call this the *normal-property*.

Lemma 4.2 *An oriented curve C is self-approaching iff any normal to C at any point a does not cross $C^{>a}$.*

Proof. The *normal-property* means that in point a we move closer or hold the distance to every point in $C^{>a}$. This property holds continuously, so for any three consecutive points the self-approaching property holds.

If the *normal-property* is not fulfilled then there exists a point a such that a normal to C at a crosses $C^{>a}$ in c'. Then in a we move away from some points in $C^{>c'}$. So there are points $b \in C^{>a}$ and $c \in C^{>c'}$ for which the self-approaching property is not true. □

Example. The logarithmic spiral, directed to the center, is an interesting example for such a curve. In polar coordinates it is the set of all points $(\varphi, e^{\varphi \cot \alpha})$ with constant $\alpha < 90°$, which is the angle between the tangent and the radius to each point on the curve, see Figure 7. It is self-approaching if α fulfills

$$\alpha \leq \arctan\left(\frac{3\pi}{2\,\mathrm{W}(\frac{3}{2}\pi)}\right) \approx 74.66°$$

in which W denotes Lambert's W function [3] defined by the functional equation $W(x)\,e^{W(x)} = x$. Figure 7 shows the limiting case where the normal at any point is tangent to the rest of the curve.

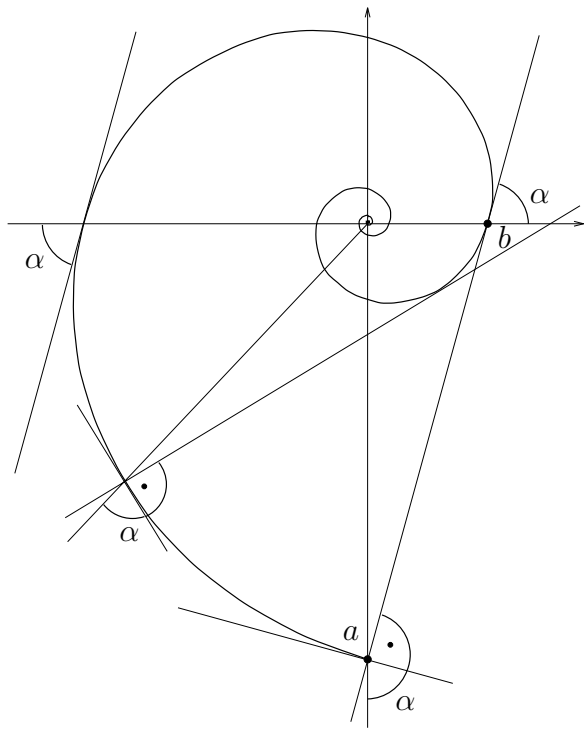

Figure 7: The narrowest self-approaching logarithmic spiral.

This special curve is in a sense the narrowest self-approaching logarithmic spiral. Let us suppose that we fix a string at point b of Figure 7 and attach a pencil at point a. Now we move the pencil clockwise holding the string taut. Then the pencil

draws the spiral while the string wraps around the inner circle of the spiral. Therefore this curve is its own involute.

One can show that its detour equals $1/\cos\alpha_{\max} \approx 3.78$, but despite its optimized form there are other self-approaching curves with a bigger detour as we will see in Section 4.3.

4.2 CAB curves are self-approaching

Let M be a path generated by CAB. We want to show that M is self-approaching. As we have shown in Section 3.2, M is an oriented planar curve that consists of a finite number of smooth pieces. So it is sufficient to prove the *normal-property*, due to Lemma 4.2.

Lemma 4.3 *For every point u on M, $M^{>u}$ is contained in* $\ker(\mathrm{vis}(u))$.

Proof. We only make use of the fact that at each point u CAB walks into $K(u) \cap G(u)$. From Lemma 3.2 we know that the actual visibility polygon grows. So for a point $v \in M^{>u}$ we have $\mathrm{vis}(u) \subseteq \mathrm{vis}(v)$ and from Lemma 2.3 we conclude $\ker(\mathrm{vis}(v)) \subseteq \ker(\mathrm{vis}(u))$. By definition $v \in \ker(\mathrm{vis}(v))$ and the claim follows. □

Theorem 4.4 *The path M generated by CAB is self-approaching.*

Proof. Let u be a point on M. From Lemma 3.2 we know that $\ker(\mathrm{vis}(u))$ is enclosed in the wedge $K(u) \cap G(u)$. We show that a normal to M at u cannot cross the boundary of $K(u) \cap G(u)$.

When the robot follows the angular bisector of $G(u)$ (CAB's **else**-branch) then $G(u)$ has a maximum angle of 180° and the wedge $K(u) \cap G(u)$ is enclosed by $G(u)$. So the normal to the bisector of $G(u)$ at u cannot cross the boundary of $K(u) \cap G(u)$. Otherwise (**then**-branch) the robot follows the boundary of $K(u)$. Due to Lemma 3.3 the wedge $G(u) \cap K(u)$ includes an angle $\beta \leq 90°$ and so the normal to a border of $G(u) \cap K(u)$ at u cannot cross $G(u) \cap K(u)$.

Summarizing, no normal to M at u can cross $\ker(\mathrm{vis}(u))$, which in turn contains $M^{>u}$, due to Lemma 4.3. Therefore M is self-approaching. □

4.3 Analysing self-approaching curves

In this section we analyse the detour of self-approaching curves. First we show that their length is bounded by the perimeter of their convex hull. Then we estimate the perimeter of their convex hull and prove, by giving an example, that the bound is tight.

Let $\mathrm{ch}(C)$ denote the convex hull of a curve C and $\mathrm{per}(C)$ the length of the perimeter of $\mathrm{ch}(C)$.

Theorem 4.5 *The arc length of a self-approaching curve C is less than or equal to the perimeter,* $\mathrm{per}(C)$*, of its convex hull.*

Proof. The arc length of a piecewise smooth curve C is, by definition, the supremum of the lengths of all polygonal chains with vertices on C in the same order as they appear on C. Therefore, an upper bound for the length of all such chains is also an upper bound for the length of C.

We take an arbitrary polygonal chain Q whose vertices lie on C in the same order. By induction on the number of vertices of Q, we will prove that Q is shorter than the perimeter, $\mathrm{per}(Q)$, of its convex hull, $\mathrm{ch}(Q)$, which in turn is bounded by $\mathrm{per}(C)$. Note that the vertices of $\mathrm{ch}(Q)$ are also vertices of Q and are therefore points on C.

The assertion is true for Q being a line segment, so let us assume that Q has at least three vertices, the first two are called a and b. The induction hypothesis is that $\mathrm{length}(Q^{\geq b}) \leq \mathrm{per}(Q^{\geq b})$.

We distinguish two cases depending on whether b lies on the boundary of $\mathrm{ch}(Q)$ or not.

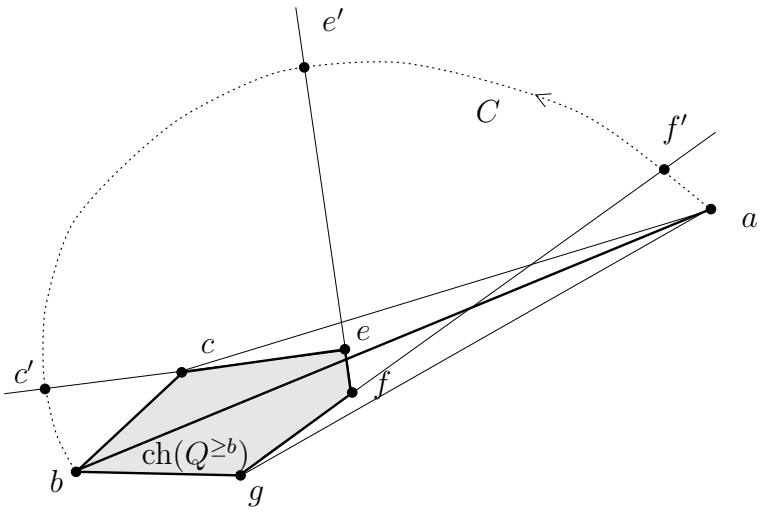

Figure 8: $d(c,a) + d(g,a) - d(a,b) \geq d(g,f) + d(f,e) + d(e,c)$.

Case 1. The point b is on the boundary of $\mathrm{ch}(Q)$. We have a situation as depicted in Figure 8. From $\mathrm{ch}(Q^{\geq b})$ to $\mathrm{ch}(Q)$, the convex hull changes as follows. The segments $\overline{gf}$, $\overline{fe}$ and $\overline{ec}$ which belong to $\mathrm{ch}(Q^{\geq b})$ are replaced by the segments $\overline{ca}$ and $\overline{ga}$. Since $\mathrm{length}(Q^{\geq b}) + d(a,b) = \mathrm{length}(Q)$ it suffices to prove that

$$d(c,a) + d(g,a) - d(a,b) \geq d(g,f) + d(f,e) + d(e,c).$$

We do not know which way C takes from a to b but there are either points $f' \in \ell(g,f)$, $e' \in \ell(f,e)$ and $c' \in \ell(e,c)$ in exactly this order on $C[a,b]$ or there are points $c' \in \ell(c,e)$, $e' \in \ell(e,f)$ and $f' \in \ell(f,g)$ in this order on $C[a,b]$. W.l.o.g. we assume the first case.

While the curve C moves from a to f' it gets closer to f. Therefore

$$d(g,a) \geq d(g,f') = d(g,f) + d(f,f')$$

By the same argument C gets closer to f while it runs from f' to e'. Therefore

$$d(f,f') \geq d(f,e') = d(f,e) + d(e,e')$$

Similarly we have $d(e,e') \geq d(e,c) + d(c,c')$ and also $d(c,c') \geq d(c,b)$. Altogether we conclude

$$d(c,a) + d(g,a) \geq d(g,f) + d(f,e) + d(e,c) + d(c,b) + d(c,a).$$

The fact $d(c,b) + d(c,a) \geq d(a,b)$ finishes the proof.

Notice that the arguments generalize to any number of vertices of $\mathrm{ch}(Q^{\geq b})$, instead of e, f and c and also for cases with $c = b$ or $g = b$.

Notice that this argument generalizes to any number of vertices of $\mathrm{ch}(Q^{\geq b})$, instead of c and e.

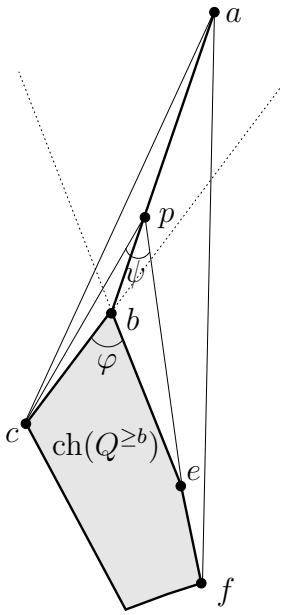

Figure 9: $d(c,p) + d(e,p) - d(b,c) - d(b,e) \geq d(b,p)$

Case 2. The point b is not on the boundary of $\mathrm{ch}(Q)$. Then a must lie in the wedge formed by the prolongations of the adjacent edges of $\mathrm{ch}(Q^{\geq b})$ at b, see Figure 9. The neighbouring vertices of b in $\mathrm{ch}(Q^{\geq b})$ are called c and e. W.l.o.g. we can assume that c appears before e on C. Then $d(b,e) \geq d(c,e)$ holds, and we have $\varphi \leq 90°$ since $\overline{ce}$ is not the longest edge of the triangle bce.

We assume a situation as in Figure 9 where the points b and e of $\mathrm{ch}(Q^{\geq b})$ are not on the boundary of $\mathrm{ch}(Q)$. Let p be the intersection point of $\overline{a\,b}$ and $\ell(e,f)$. Using the induction hypothesis and the fact $\varphi \leq 90°$ we show that $\mathrm{length}(Q^{\geq p})$ is not greater than $\mathrm{per}(Q^{\geq p})$. From $\mathrm{ch}(Q^{\geq b})$ to $\mathrm{ch}(Q^{\geq p})$, the convex hull changes as follows. The segments $\overline{b\,c}$ and $\overline{b\,e}$ which belong to $\mathrm{ch}(Q^{\geq b})$ are replaced by the segments $\overline{p\,c}$ and $\overline{p\,e}$ of $\mathrm{ch}(Q^{\geq p})$. Since $\mathrm{length}(Q^{\geq p})$ is equal to $\mathrm{length}(Q^{\geq b}) + d(b,p)$ it is sufficient to show that

$$d(c,p) + d(e,p) - d(b,c) - d(b,e) \geq d(b,p)$$

holds, this is exactly the conclusion of Lemma A.1 in the appendix on page 18.

Since we can use the assumption $\mathrm{length}(Q^{\geq p}) \leq \mathrm{per}(Q^{\geq p})$ and the fact $\psi \leq \varphi \leq 90°$ the same argument holds also for $Q^{\geq a} = Q$ and also for the case that more vertices of

$\mathrm{ch}(Q^{\geq b})$ than only b and e do not reappear as vertices in $\mathrm{ch}(Q)$. This concludes the proof. □

In the following, for two points p and q let $\mathrm{circ}_p(q)$ denote the circle with center p passing through q.

As an immediate consequence of Theorem 4.5, we have an upper bound of 2π for the detour of self-approaching curves, because any such curve from point a to point b must be contained in $\mathrm{circ}_b(a)$. The following theorem refines this argument to a better bound.

Theorem 4.6 *The detour of a self-approaching curve is not greater than*

$$c_{\max} := \max_{\beta \in [0..\frac{\pi}{2}]} \frac{2\beta + \pi + 2}{\sqrt{5 - 4\cos\beta}} \approx 5.3331\ldots$$

Proof. Let a, f denote the first resp. last point of a self-approaching curve C. The proof works as follows: We show that $\frac{\mathrm{length}(C)}{d(a,f)} \leq c_{\max}$ holds if the curve does *not* cross the line segment $\overline{a\,f}$. If it does, we apply this bound for each subcurve between two successive curve points on $\overline{a\,f}$ and add up the lengths. Due to the self-approaching property, the curve points on $\overline{a\,f}$ appear in the same order as on C.

So for the rest of this section we may assume that $\mathbb{M}C$ does not cross the line segment $\overline{a\,f}$. In the proof of Theorem 4.5, we have seen in the first paragraph of Case 2 that the whole of C lies inside a wedge at a of $\varphi = 90°$. We choose such a wedge consisting of two orthogonal halflines X and Y starting at a. W.l.o.g. we assume that the initial part of C lies on the left side of the edge directed from a to f, and, if necessary, we rotate such that the halfline on the other(right) side of $\overline{a\,f}$ touches C at a point e, see Figure 10. Let h and w be the height resp. width of the bounding box of a and f according to rectangle with sides parallel to the wedge and with diagonal $\overline{a\,f}$.

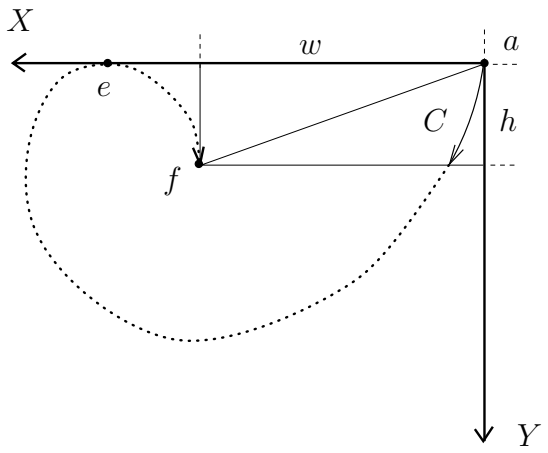

Figure 10: Any self-approaching curve is contained in a wedge of 90°.

We will construct a convex area, A, that contains C and we will show that the perimeter of this area divided by $d(a,f) = \sqrt{h^2 + w^2}$ is bounded by $c_{\max}$. Then from Theorem 4.5, it follows that $c_{\max}$ bounds the detour $\frac{\mathrm{length}(C)}{d(a,f)}$.

This construction goes as follows, refer to Figure 11. First we know that C is contained in the wedge at a, and $C[a, e]$ is contained in $\mathrm{circ}_e(a)$ and $C[e, f]$ is contained in $\mathrm{circ}_f(e)$, due to the self-approaching property.

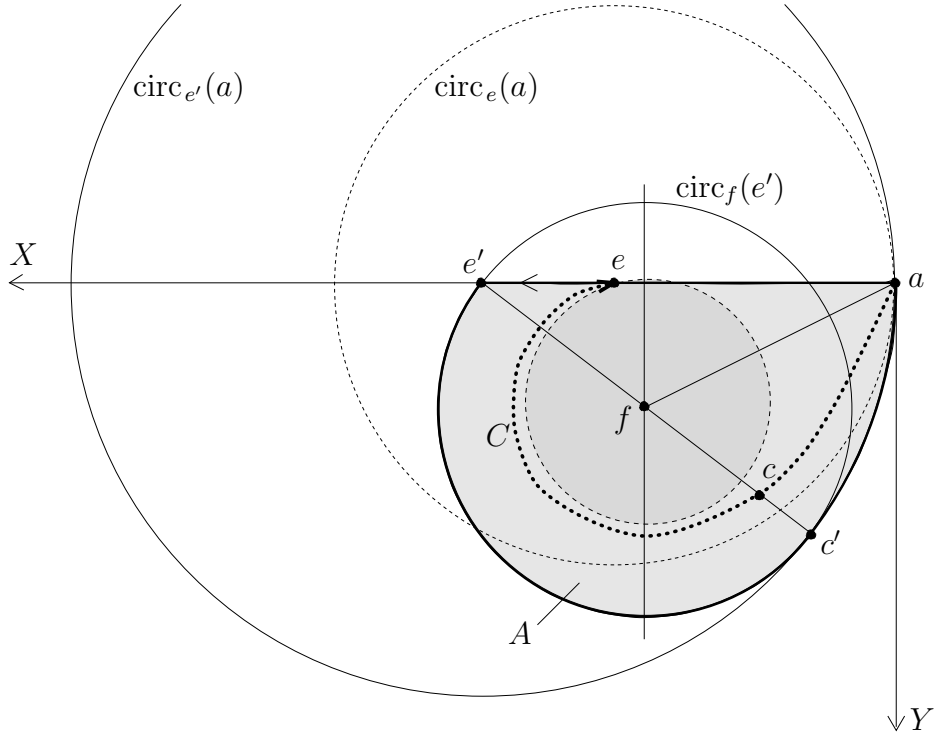

Figure 11: Curve C is contained in the constructed area A.

We also know that $C[a, e]$ must avoid $\mathrm{circ}_f(e)$, but pass around it to reach e because it must not cross $\overline{a\,f}$. We conclude that $\mathrm{circ}_f(e)$ is contained in $\mathrm{circ}_e(a)$.

Now, we will enlarge these circles to a certain extent. Instead of e, we use a point e' on X with $d(a, e') \geq d(a, e)$, such that $\mathrm{circ}_{e'}(a)$ still contains $\mathrm{circ}_f(e')$ and touches it in one point c'. This is possible because for every position of e' on X with $d(a, e') \geq d(a, e)$, the whole circle $\mathrm{circ}_{e'}(a)$ is always on one side of Y, while $\mathrm{circ}_f(e')$ must eventually cross l. Note that $d(e', a) = d(e', c') = d(c', f) + d((f, e') = 2d(f, e')$ holds, in other words the radius of $\mathrm{circ}_{e'}(a)$ equals the diameter of $\mathrm{circ}_f(e')$.

Now let c be the point at which curve C crosses $\overline{f\,c'}$ first. We know that $C[a, c]$ is contained in $\mathrm{circ}_e(a) \subseteq \mathrm{circ}_{e'}(a)$, $C[c, f]$ is contained in $\mathrm{circ}_f(c) \subseteq \mathrm{circ}_f(c') = \mathrm{circ}_f(e')$. So the curve C is included in the convex area A limited by $\overline{e'\,a}$, the circular arc from a to c' about e', and the halfcircle from c' to e' about f, see Figure 11.

We choose a scale such that $d(f, c') = 1$. Let β be the angle between $\overline{e'\,a}$ and $\overline{e'\,c'}$, see Figure 12. Then $d(e', c') = 2 = d(e', a)$, $w = 2 - \cos\beta$ and $h = \sin\beta$. The length of the arc from a to c' equals 2β while the halfcircle from c' to e' has length π.

Then the perimeter of the constructed convex area A divided by $d(a, f) = \sqrt{h^2 + w^2}$ is given by the following function f in β.

$$f(\beta) := \frac{2\beta + \pi + 2}{\sqrt{\sin^2\beta + (2 - \cos\beta)^2}} = \frac{2\beta + \pi + 2}{\sqrt{5 - 4\cos\beta}}$$

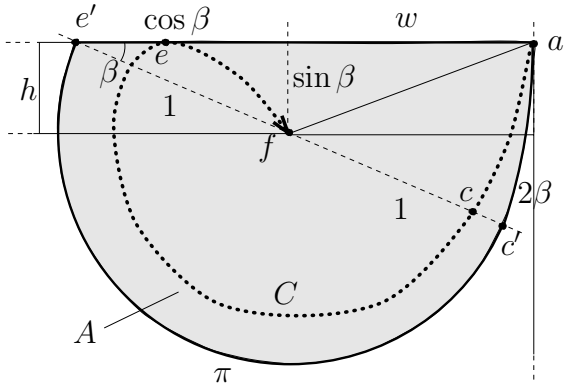

Figure 12: Measuring the perimeter of the convex area A.

The maximum value, c_{max}, of $f(\beta)$ equals aproximately $5.3331\ldots$ for $0 \le \beta \le \pi/2$ and is attained at $\beta_{opt} \approx 11.22°$ with $8\cos\beta_{opt} + 2\sin\beta_{opt}(2\beta_{opt} + \pi + 2) = 10$. □

Surprisingly, it turns out that there cannot be a smaller upper bound for self-approaching curves.

Theorem 4.7 *The upper bound c_{max} for the detour of self-approaching curves is tight.*

Proof. We construct a curve with a convex hull similar to the bounding area A in Theorem 4.6.

As a first attempt we consider the curve in Figure 13. From the start point, a, to the end, f, it consists of a circular arc of radius 2 and angle β, a half circle of radius 1, and a line segment of length 1. This curve is self-approaching, its length equals $2\beta + \pi + 1$ while $d(a, f) = \sqrt{5 - 4\cos\beta}$. The ratio takes on a maximum value of approximately 4.38.

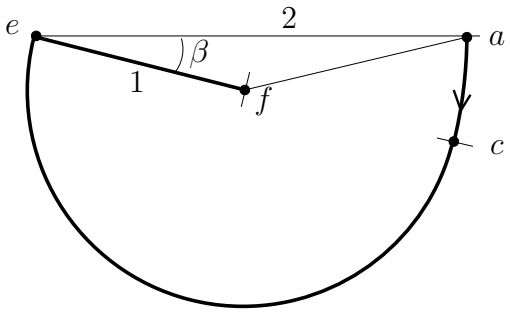

Figure 13: This self-approaching curve has a detour of 4.38.

But there is some room for improvements in the last step, i. e. the line segment from e to f. Instead of walking straight from e to f we use a sequence of pieces of small cycloids. (A cycloid is known to be the orbit of a point on the boundary of a rolling circle and it has another cycloid as its involute.) For an odd number $n \in \mathbb{N}$ we can fill a rectangle of height h and width $w = 2nh/\pi$ with n successive congruent pieces of cycloids such that they form a curve from the lower left to the upper right corner, see Figure 14. Each piece is a cycloid generated by a circle of radius h/π

rolling on a vertical line, and each one is the involute of its successor. The resulting curve is self-approaching, and its length is exactly $2w$ since the length of a piece is twice its width.

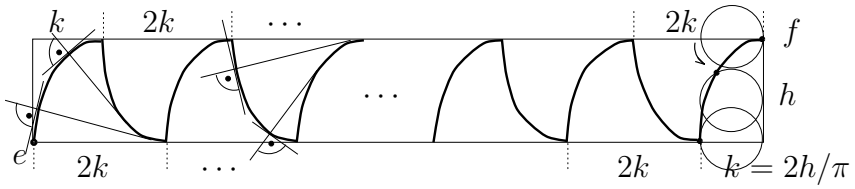

Figure 14: How to fill a rectangle of width w and height h with a self-approaching curve of length $2w$ using pieces of cycloids.

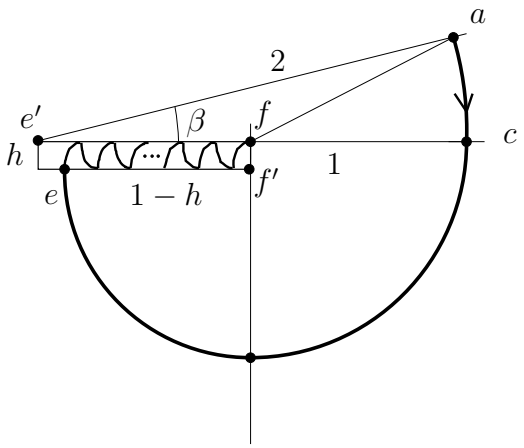

Figure 15: Replacing the line segment by a rectangle that gives room for improvement.

Now let us replace the line segment in our first attempt by such a construction in a rectangle of height $h = \frac{1}{(2n/\pi)+1}$ and width $w = 1 - h = 2nh/\pi$, see Figure 15.

The curve consists of

- a circular arc about e' of length 2β
- a circular arc about f of length $\pi/2$
- a circular arc about f' of length $(1-h)\pi/2$ and
- a sequence of cycloids from e to f of overall length $2(1-h)$.

We choose $\beta = \beta_{opt}$ from Theorem 4.6 and for $n = 1, 3, 5 \ldots$ we have a sequence of self-approaching curves with a detour of at least

$$\frac{2\beta_{opt} + \pi/2 + (1-h)(\pi/2+2)}{\sqrt{5 - 4\cos\beta_{opt}}}$$

which converges to c_{max} for $n \to \infty$ (i. e. $h \to 0$). □

Actually, as n tends to infinity, one could think of the curve of Figure 14 as a thick line segment of length 2, while its endpoints are only distance 1 apart. One might wonder if a factor bigger than 2 can be achieved by a different technique. Note that this is not possible as a direct consequence of Theorem 4.6.

5 A lower bound

A lower bound for any strategy for the kernel problem can be derived in the following way.

Lemma 5.1 *No strategy for the kernel problem can guarantee a detour of less than* $\sqrt{2}$.

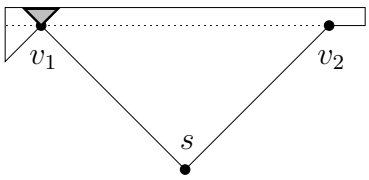

Figure 16: Any strategy can be forced to make a detour of $\sqrt{2}$.

Proof. The polygon P shown in Figure 16 is a rectangular, isosceles triangle with two small caves at the upper corners. We allow only two cases, the kernel may be near the left corner, as shown in the picture, or near the right corner, i. e. if we flip left and right corners in the picture.

The robot starts at the lower corner, s, and before reaching the line $\overline{v_1\,v_2}$ no strategy can determine which of the two cases occurs. If $\overline{v_1\,v_2}$ is reached to the left of its midpoint, then the kernel can be at the right corner and vice versa. But then a detour of $\sqrt{2}$ is unavoidable. □

By the way, the greatest detour of strategy CAB we know of is $\pi + 1$, as shown in Figure 17. However such a one cave polygon can be explored much more efficient at a cost of $1.21\ldots$, see [12].

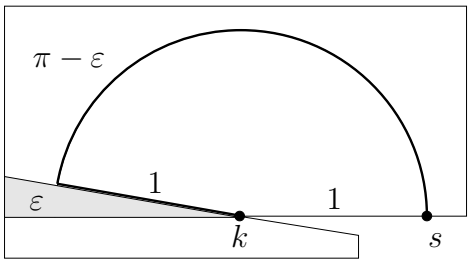

Figure 17: In this polygon, CAB achieves a detour of $\pi-\varepsilon+1$, and ε may be arbitrarily small.

A Appendix

Lemma A.1 *Let v be a point inside a triangle abc. We connect each vertex to v using segments $l_1 = \overline{b\,v}$, $r_1 = \overline{c\,v}$, and $z = \overline{a\,v}$. Let $l_2 = \overline{a\,b}$ and $r_2 = \overline{a\,c}$ be two edges of the triangle; see Figure 18. If the angle $\varphi \leq \pi$ between l_1 and r_1 is less than or equal to $\pi/2$ then for the lengths of the segments $l_1 + r_1 + z \leq l_2 + r_2$ holds.*

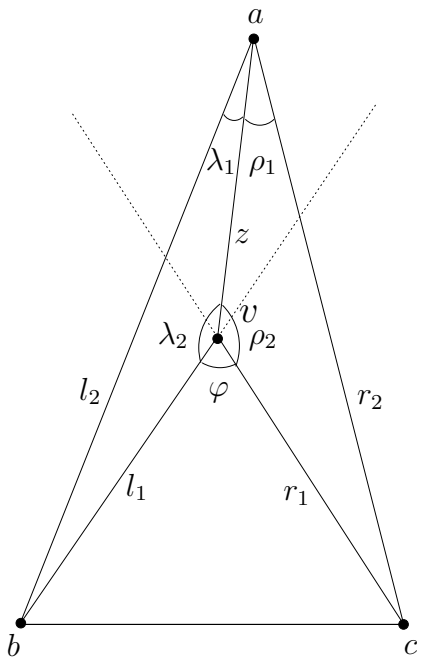

Figure 18: $l_1 + r_1 + z \leq l_2 + r_2$ holds for $\varphi \leq \pi/2$.

Proof. The assumption is obviously true for $z = 0$. Let $z \neq 0$. We have to prove the inequality $l_2 - l_1 + r_2 - r_1 \geq z$. Let $\lambda_1 \leq \pi$ be the angle between l_2 and z, $\lambda_2 \leq \pi$ be the angle between l_1 and z, $\rho_1 \leq \pi$ be the angle between r_2 and z and $\rho_2 \leq \pi$ be the angle between r_1 and z. Using the law of sines we substitute l_1 with $z\frac{\sin\lambda_1}{\sin(\lambda_1+\lambda_2)}$ and l_2 with $z\frac{\sin\lambda_2}{\sin(\lambda_1+\lambda_2)}$. We transform r_1 and r_2 analogously and divide the whole expression by z. So we have to show that

$$\frac{\sin\lambda_2 - \sin\lambda_1}{\sin(\lambda_1+\lambda_2)} + \frac{\sin\rho_2 - \sin\rho_1}{\sin(\rho_1+\rho_2)} \geq 1$$

is true. We consider some simple transformations:

$$\begin{aligned}
\frac{\sin\lambda_2 - \sin\lambda_1}{\sin(\lambda_1+\lambda_2)} &= \frac{\sin\lambda_2 - \sin((\lambda_1+\lambda_2)-\lambda_2)}{\sin(\lambda_1+\lambda_2)} \\
&= \frac{\sin\lambda_2 - \sin(\lambda_1+\lambda_2)\cos\lambda_2 + \cos(\lambda_1+\lambda_2)\sin\lambda_2}{\sin(\lambda_1+\lambda_2)} \\
&= -\cos\lambda_2 + \underbrace{\frac{(1+\cos(\lambda_1+\lambda_2))\sin\lambda_2}{\sin(\lambda_1+\lambda_2)}}_{(*)} \geq -\cos\lambda_2
\end{aligned}$$

Note that $(*) \geq 0$ holds because of $0 \leq \lambda_1 + \lambda_2 \leq \pi$. Similarly we have $\frac{\sin\rho_2-\sin\rho_1}{\sin(\rho_1+\rho_2)} \geq -\cos\rho_2$ so it is sufficient to show that $-\cos\lambda_2 - \cos\rho_2 \geq 1$ is true. We conclude $\frac{\lambda_2+\rho_2}{2} = \pi - \frac{\varphi}{2}$ from $\lambda_2 + \rho_2 + \varphi = 2\pi$. Since λ_2 and ρ_2 are inner angles of a triangle we know $\lambda_2, \rho_2 \leq \pi$ and from $\varphi \leq \pi/2$ we conclude $\lambda_2, \rho_2 \geq \pi/2$. Therefore $|\lambda_2 - \rho_2| \leq \pi/2$ is true. Now

$$-\cos\lambda_2 - \cos\rho_2 = -2\cos\left(\frac{\lambda_2+\rho_2}{2}\right)\cos\left(\frac{|\lambda_2-\rho_2|}{2}\right) \geq 1$$

holds since $-\cos\left(\frac{\lambda_2+\rho_2}{2}\right) = \cos\left(\frac{\varphi}{2}\right) \geq \cos\left(\frac{\pi}{4}\right) = \frac{1}{\sqrt{2}}$ and $\cos\left(\frac{|\lambda_2-\rho_2|}{2}\right) \geq \frac{1}{\sqrt{2}}$ are fulfilled. □

References

[1] R. C. Buck. *Advanced Calculus*. McGraw-Hill, New York, 1978.

[2] R. Cole and M. T. Goodrich. Optimal parallel algorithms for polygon and point-set problems. *Algorithmica*, 7:3–23, 1992.

[3] R. M. Corless, G. H. Gonnet, D. E. G. Hare, and D. J. Jeffrey. Lambert's W function in Maple. *The Maple Techical Newsletter*, Issue 9:12–22, 1993.

[4] H. T. Croft, K. J. Falconer, and R. K. Guy. *Unsolved Problems in Geometry*. Springer-Verlag, New York, 1991.

[5] A. Datta and C. Icking. Competitive searching in a generalized street. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 175–182, 1994.

[6] X. Deng, T. Kameda, and C. Papadimitriou. How to learn an unknown environment. In *Proc. 32nd Annu. IEEE Sympos. Found. Comput. Sci.*, pages 298–303, 1991.

[7] X. Deng, T. Kameda, and C. H. Papadimitriou. How to learn an unknown environment I: the rectilinear case. Technical Report CS-93-04, Department of Computer Science, York University, Canada, 1993.

[8] G. Dudek, K. Romanik, and S. Whitesides. Localizing a robot with minimum travel. In *Proc. 6th ACM-SIAM Sympos. Discrete Algorithms*, pages 437–446, 1995.

[9] K. Ghosh and S. Saluja. Optimal on-line algorithms for walking with minimum number of turns in unknown streets. Technical Report TCS-94-2, Tata Institute of Fundamental Research, Bombay, 1994.

[10] F. Hoffmann, C. Icking, R. Klein, and K. Kriegel. A competitive strategy for learning a polygon. In *Proc. 8th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1997.

[11] C. Icking. *Motion and visibility in simple polygons*. PhD thesis, Department of Computer Science, FernUniversität Hagen, Germany, 1994.

[12] C. Icking, R. Klein, and L. Ma. How to look around a corner. In *Proc. 5th Canad. Conf. Comput. Geom.*, pages 443–448, Waterloo, Canada, 1993.

[13] R. Klein. Walking an unknown street with bounded detour. *Comput. Geom. Theory Appl.*, 1:325–351, 1992.

[14] J. M. Kleinberg. The localization problem for mobile robots. In *Proc. 35th Annu. IEEE Sympos. Found. Comput. Sci.*, 1994.

[15] J. M. Kleinberg. On-line search in a simple polygon. In *Proc. 5th ACM-SIAM Sympos. Discrete Algorithms*, pages 8–15, 1994.

[16] D. T. Lee and F. P. Preparata. An optimal algorithm for finding the kernel of a polygon. *J. ACM*, 26:415–421, 1979.

[17] A. López-Ortiz and S. Schuierer. Going home through an unknown neighbourhood. Technical report, Department of Computer Science, University of Waterloo, Canada, 1994.

[18] G. Rote. Curves with increasing chords. *Mathematical Proceedings of the Cambridge Philosophical Society*, 115(1):1–12, 1994.