Reachability by paths of bounded curvature in convex polygons

Hee-kap Ahn* Otfried Cheong* Jiří Matoušek† Antoine Vigneron*

Abstract

Let B be a point robot moving in the plane, whose path is constrained to forward motions with a curvature at most 1, and let P be a convex polygon with n vertices. Given a starting *configuration* (a location and a direction of travel) for B inside P, we characterize the region of all points of P that can be reached by B, and show that it has linear complexity.

1. Introduction

The problem of planning the motion of a robot subject to non-holonomic constraints [10, 15] (for instance, bounds on velocity or acceleration [13, 5, 6], bounds on the turning angle) has received considerable attention in the robotics literature. Theoretical studies of non-holonomic motion planning are far sparser.

In this paper we consider a point robot in the plane whose turning radius is constrained to be at least 1 and that is not allowed to make reversals. This restriction corresponds naturally to constraints imposed by the steering mechanism found in car-like robots. We assume that the robot is located at a given position (and orientation) inside a convex polygon, and we are interested in the set of points in the polygon that can be reached by the robot. We put no restriction on the orientation with which the robot can reach a point.

*Dept. of Computer Science, Hong Kong University of Science & Technology, Clear Water Bay, Kowloon, Hong Kong. Email: {pastel|otfried|antoine}@cs.ust.hk. This research was supported by the Hong Kong Research Grants Council.

†Dept. of Applied Mathematics, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic. Email: matousek@kam.mff.cuni.cz. This research was supported by Czech Republic Grant GAČR 0194 and by Charles University grants No. 193 and 194.

The lack of such a restriction distinguishes our work from most of the previous theoretical work on curvature-constrained paths, which usually assumes that not a point, but a *configuration* (a location with orientation) is given. Dubins [7] was perhaps the first to study curvature-constrained shortest paths. He proved that a curvature-constrained shortest path from a given starting configuration to a given final configuration consists of at most 3 segments, each of which is either a straight line or an arc of a unit-radius circle. Reeds and Shepp [12] extended this characterization to robots that are allowed to make reversals. Using ideas from control theory, Boissonnat et al. [3] gave an alternative proof for both cases. Sussmann [16] extended the characterization to the 3-dimensional case. Recently, Agarwal et al. [1] presented an $O(n^2 \log n)$-time algorithm to compute a curvature-constrained shortest path between two given configurations inside a convex polygon.

In the presence of obstacles, Fortune and Wilfong [8] gave a decision procedure to verify if the source and target placement of a point robot may be joined by a curvature constrained path avoiding the polygonal obstacles. Jacobs and Canny [9] and Wang and Agarwal [18] gave algorithms computing an approximate curvature constrained path, and Wilfong [19] designed an exact algorithm for the case where the curvature constrained path is limited to some fixed straight "lanes" and circular arc turns between the lanes. Švestka and Overmars [17] applied the probabilistic approach to compute curvature constrained paths for car-like robots.

Agarwal et al. [2] considered the restricted case of pairwise disjoint moderate obstacles (convex obstacles whose boundary has curvature bounded by 1) and gave efficient approximation algorithms. For the same problem, Boissonnat and Lazard [4] gave an exact algorithm of $O(n^2 \log n)$-time. Reif and Wang [14] confirmed that the problem of deciding whether there exists a curvature-constrained path amidst general objects is NP-hard.

Our result is a characterization of the region of points reachable by paths under curvature constraints from a given starting configuration inside a convex polygon. We show that all points reachable from the starting configuration are also reachable by paths of a fairly special form. We also show that the reachable region has complexity linear in the complexity of the polygon. The characterization is constructive and could be used to compute the reachability region.

2. Terminology and some lemmas

Let P be a convex polygon in the plane. A *configuration* $\mathbf{s} = (s, \mathbf{d})$ is a point s together with a direction of travel $\mathbf{d}$ (a unit vector). Unless otherwise stated, we will have $s \in P$. By a *path* we mean a continuously differentiable curve (the image of a C^1-mapping of $[0, 1]$ to $\mathbf{R}^2$) consisting of finitely many line segments and finitely many circular arcs of radius at least 1. Unless stated otherwise, we also assume that a path is completely contained in P. A *configuration on the path* π is a configuration $\mathbf{s} = (s, \mathbf{d})$ with s on π such that $\mathbf{d}$ is the forward tangent to π in s. The *starting configuration* of π is the starting point of π with its forward tangent.

Given a configuration $\mathbf{s} = (s, \mathbf{d})$, the *left disc* $D_L(\mathbf{s})$ (*right disc* $D_R(\mathbf{s})$) is the unit disc touching s and completely contained in the left (right) half-plane defined by the directed line through $\mathbf{d}$. (All *unit discs* in this paper are *unit-radius* discs.)

The *left directly accessible region* $\mathrm{LDA}(\mathbf{s})$ is the set of all points in P that can be reached by a path with starting configuration $\mathbf{s}$ consisting of a single (possible zero-length) circular arc on the boundary of $D_L(\mathbf{s})$ followed by a single (possible zero-length) line segment. The *right directly accessible region* is defined analogously. The *directly accessible region* $\mathrm{DA}(\mathbf{s})$ is the union of the left and right directly accessible region.

Pestov-Ionin lemma.

The following lemma is perhaps the foundation for all our results. In a slightly less general form, it was proven by Pestov and Ionin [11].

Lemma 1 (Pestov-Ionin) *Let Γ be a simple closed arc, C^1 and piecewise C^2 everywhere except maybe at one point. Suppose its curvature is at least 1. Then a unit disc lies in the interior of Γ.*

We need a somewhat more general result. The proof uses the same idea as Pestov and Ionin.

Lemma 2 *Let D be a closed disc, and Γ be a simple arc with endpoints (a, b) such that $\Gamma \cap D = \{a, b\}$. Suppose that Γ is C_1 and piecewise C_2, with curvature at least 1. Then a unit disc lies within the region $\mathcal{R}$ bounded by Γ and the exterior arc of the boundary ∂D of D.*

Proof. We proceed by induction on the length ℓ of Γ (or, more precisely, by induction on $\lfloor \ell/\pi \rfloor$).

When $\ell < \pi$, we can prove by integration that a unit disc tangent to Γ does not cross Γ. We consider a unit disc D_0 tangent to Γ at $m \notin \{a, b\}$ on the interior side. Since $D_0 \setminus D$ has only one connected component and $m \in D_0 \setminus D$, clearly D_0 is contained in $\mathcal{R}$.

Otherwise, let m be the point halfway between a and b on Γ. Let D_1 be the largest disc contained in $\mathcal{R}$ that is tangent to Γ in m. If the radius of D_1 is larger or equal to 1, then we are done. If not, D_1 must touch $\partial\mathcal{R}$ in another point m'. Clearly m' lies on Γ, and the length of the arc Γ' of Γ between m and m' is at most $\ell/2$. By the induction hypothesis, a unit disc D_2 lies inside the region $\mathcal{R}_1$ bounded by Γ' and D_1. Since $\mathcal{R}_1 \subseteq \mathcal{R}$, the lemma follows. ◻

Lemma 1 follows from Lemma 2 by observing that it still holds when D degenerates to a point.

Filling.

By $\mathrm{FIL}(P)$ we denote the set of all unit-radius discs that are completely contained in P, and we let $\mathit{Fil}(P)$ be the union of all the discs in $\mathrm{FIL}(P)$. Both $\mathrm{FIL}(P)$ and $\mathit{Fil}(P)$ will be called the *filling* of P.

Pockets.

The connected components of $P \setminus \mathit{Fil}(P)$ are called the *pockets* of P. A pocket R of P is bounded by a single circular arc (lying on one disc of $\mathrm{FIL}(P)$) and a connected part of the boundary of P. The first and last edge on this connected chain are called the *mouth edges* of the pocket. Its extremities are called the *mouth points*. The mouth edges form an angle smaller than π (this is equivalent to observing that the mouth points lie on the same disc of $\mathrm{FIL}(P)$ and form an angle smaller than π) [1]. Agarwal et al. [1] proved the following lemma.

Lemma 3 (Pocket lemma) *A path entering a pocket from $\mathit{Fil}(P)$ cannot leave the pocket anymore.*

Reachability for a union of discs.

For a set $\mathcal{D}$ of unit discs, we let $\mathrm{conv}(\mathcal{D})$ denote the set of all unit discs contained in the convex hull of $\bigcup \mathcal{D}$. Equivalently, $\mathrm{conv}(\mathcal{D})$ consists of the unit discs centered at points of the convex hull of the set of all centers of the discs in $\mathcal{D}$.

Lemma 4 *Let $\mathcal{D}$ be a set of unit discs in the plane. Then $\bigcap \mathcal{D} = \bigcap \mathrm{conv}(\mathcal{D})$.*

Given a convex set Q, a *configuration on the boundary* of Q is a configuration $\mathbf{s} = (s, \mathbf{d})$ with s on the boundary of Q and such that $\mathbf{d}$ is tangent to the boundary of Q in s.

Lemma 5 *Let $\mathcal{D}$ be a set of unit discs, and let $\mathbf{s}$ be a starting configuration on the boundary of $\bigcup \mathrm{conv}(\mathcal{D})$. Then no point in the interior of $\bigcap \mathcal{D}$ can be reached by a path starting at $\mathbf{s}$ and contained in $\bigcup \mathrm{conv}(\mathcal{D})$, or even in any convex polygon P such that $\mathrm{FIL}(P) = \mathrm{conv}(\mathcal{D})$.*

Proof. Assume to the contrary that there is a path γ with starting configuration $\mathbf{s}$ on the boundary of $\bigcup \mathrm{conv}(\mathcal{D})$ and ending point in the interior of $\bigcap \mathcal{D}$. Assume for the moment that γ lies completely in $\bigcup \mathrm{conv}(\mathcal{D})$.

We extend γ to infinity using a straight ray, such that the extended path is still C^1. We then extend γ backwards, by attaching a single loop around the boundary of $\bigcup \mathrm{conv}(\mathcal{D})$ at $\mathbf{s}$. To summarize, the extended path γ' starts at $\mathbf{s}$, makes a single loop around the boundary of $\bigcup \mathrm{conv}(\mathcal{D})$, then follows the original path γ, and finally escapes to infinity along a straight line. We can now construct a simple, closed path γ'' as follows: starting at infinity, we follow γ' *backwards*, until we encounter the first intersection of γ' with the part that we have already seen. Such an intersection must exist since the two extensions intersect. Note that γ'' must contain the endpoint of γ in the interior of $\bigcap \mathcal{D}$.

By the Pestov-Ionin Lemma, γ'' contains a unit-disc D. Since γ'' is contained in $\bigcup \mathrm{conv}(\mathcal{D})$, we have $D \in \mathrm{conv}(\mathcal{D})$. Consequently, the interior of $\bigcap \mathcal{D}$ lies in the interior of D, a contradiction with the fact that γ'' contains a point in the interior of $\bigcap \mathcal{D}$.

The lemma still holds when we allow the path to lie inside a convex polygon P with $\mathrm{FIL}(P) = \mathrm{conv}(\mathcal{D})$. After all, by Lemma 3, the path cannot return to $Fil(P)$ after it has left it. ◻

The characterization.
Consider a starting configuration $\mathbf{s}$ on the boundary of the filling $Fil(P)$. It is easy to see that any point in P not in $\bigcap \mathrm{FIL}(P)$ can be reached by a path from $\mathbf{s}$. On the other hand, by Lemma 5, no point in $\bigcap \mathrm{FIL}(P)$ can be reached, and so we have a complete characterization of the region reachable from $\mathbf{s}$ as the complement of $\bigcap \mathrm{FIL}(P)$.

If the starting configuration lies on the boundary of an arbitrary unit disc contained in P, the same characterization holds. For arbitrary starting configurations, however, the situation becomes far more complicated. There is the possibility that the boundary of $Fil(P)$ is not reachable from $\mathbf{s}$, and the filling has no relation to the reachable region.

If the boundary of $Fil(P)$ is reachable from $\mathbf{s}$, all points outside $\bigcap \mathrm{FIL}(P)$ are reachable, but it is still possible that some points inside $\bigcap \mathrm{FIL}(P)$ are reachable, for instance because they lie in $\mathrm{DA}(\mathbf{s})$, or because $\mathbf{s}$ lies in a pocket with additional maneuvering space (that we would not have been able to exploit if starting inside $Fil(P)$ by Lemma 3).

In the rest of this paper, we give a complete characterization of the reachable region, for any starting configuration in P.

3. Paths starting along a side

We assume that the starting configuration $\mathbf{s}$ is on the boundary of P (recall that this means also that the direction is tangent to the boundary). Without loss of generality, we also assume that the direction of $\mathbf{s}$ is counterclockwise along the boundary, so that points of P are reached locally by a left turn from $\mathbf{s}$. For the ease of description, we will always consider $\mathbf{s}$ directed vertically upwards. Let's write just D_L for $D_L(\mathbf{s})$.

We will see in this section that for $\mathbf{s}$ on the boundary as described above, we can restrict ourselves to paths containing no right-turning arcs. We will call such a path a *left-turning path*.

If $\mathrm{FIL}(P) \neq \emptyset$, we define the point $c = c(P, \mathbf{s})$ on the boundary of P as the first point of $Fil(P)$ reached by traversing the boundary of P counterclockwise starting from s (see Fig. 1).

If the left disc D_L is not completely contained in P, we define

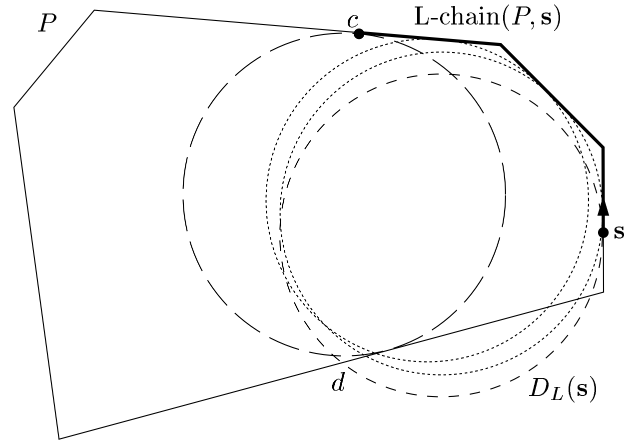

Figure 1: The left chain and the extreme discs in $\mathrm{LFIL}(P, \mathbf{s})$.

the point $d = d(P, \mathbf{s}) \in \partial P$ as the point where the boundary of P, when traversed from s counterclockwise, first enters the interior of D_L (more precisely, d is the last point outside of this interior). Note that at least one of the points c and d is always well-defined.

The *left chain* L-chain$(P, \mathbf{s})$ is the part of the boundary of P from s to c or d, whichever comes first (in the counterclockwise direction). The *left filling* $\mathrm{LFIL}(P, \mathbf{s})$ is the set of all unit-radius discs that touch the boundary of P from inside at some point of the left chain but whose interior does not intersect the left chain. Note that $\mathrm{LFIL}(P, \mathbf{s})$ always contains the left disc D_L.

The *extreme left-turning path* $\lambda(P, \mathbf{s})$ is a path from $\mathbf{s}$ going along the boundary of $\mathrm{conv}\left(\bigcup \mathrm{LFIL}(P, \mathbf{s})\right)$ (so, the extreme left-turning path is a left-turning path that turns left *as little as possible*). If the left chain ends in the point d, then $\lambda(P, \mathbf{s})$ ends there as well. If the left chain ends at c, then $\lambda(P, \mathbf{s})$ goes all the way to c, and then once around the boundary of $Fil(P)$ back to the point c (Fig. 2). Let's remark that

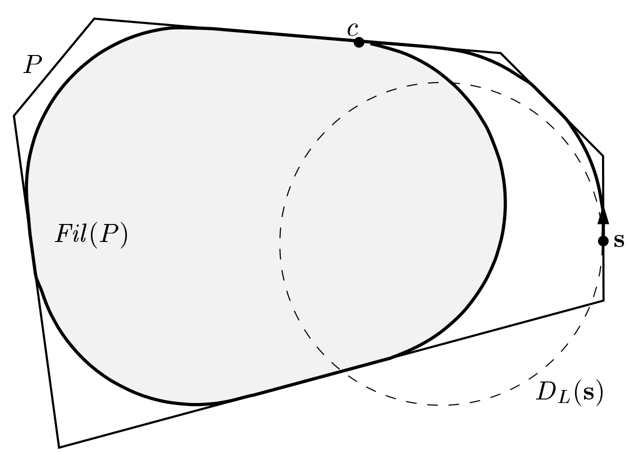

Figure 2: The extreme left-turning path.

the set of centers of the discs in $\mathrm{LFIL}(P, \mathbf{s})$ need not be connected or convex. For example, Fig. 3 shows a situation where the left filling consists of just 3 discs; their centers are marked by black dots (in general, any number of connected components is possible).

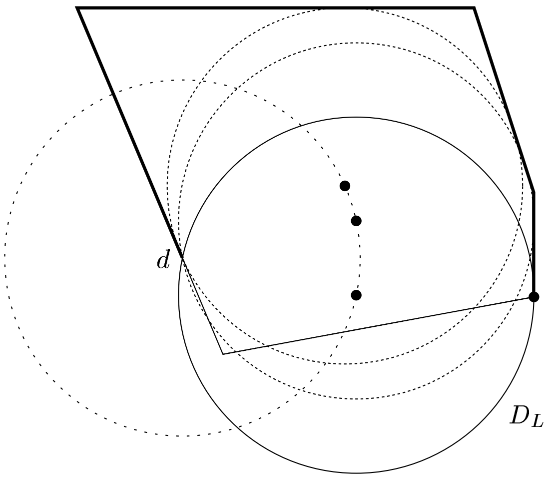

Figure 3: Example where the left filling consists of just 3 discs.

The following proposition describes the region reachable from $\mathbf{s}$. The essential information is that the reachable region is the union of the left directly accessible regions along the extreme left-turning path; in particular, all reachable points can be reached by left-turning paths.

Proposition 6 *Let P and $\mathbf{s}$ be as above. Then the set of points of P reachable from $\mathbf{s}$ equals $\bigcup_{\mathbf{r}} \mathrm{LDA}(P, \mathbf{r})$, where $\mathbf{r}$ runs over all the starting configurations on the extreme left-turning path.*

There are three cases, two of which have an alternative description of the reachable region:

(i) If the left chain ends in c, the reachable region is the complement of the interior of $\bigcap(\mathrm{FIL}(P) \cup \mathrm{LFIL}(P, \mathbf{s}))$.

(ii) If the left chain ends at d and turns at most π, then the reachable region is $\mathrm{LDA}(P, \mathbf{s})$.

(iii) If the left chain ends at d and turns more than π, then the reachable region is $\bigcup_{\mathbf{r}} \mathrm{LDA}(P, \mathbf{r})$, with $\mathbf{r}$ as above.

For a polygon with n sides, the complexity of the reachable region can be $\Omega(n)$ (for example, a regular polygon on n vertices with diameter between 2 and 4 has a reachable region of complexity n), and $\Omega(n)$ discs of $\mathrm{LFIL}(P)$ can contribute to its boundary. On the other hand, in case (i), if all points of P are reachable this is "witnessed" by at most three discs of $\mathrm{FIL}(P) \cup \mathrm{LFIL}(P, \mathbf{s})$, by Helly's theorem. In cases (ii) and (iii), there are always non-reachable points (for instance, a point directly "below" the starting point is unreachable in case (iii) by the Pestov-Ionin lemma).

Clearly, all the points of $\bigcup_{\mathbf{r}} \mathrm{LDA}(P, \mathbf{r})$ are reachable by left-turning paths (by first following the extreme left-turning path and then, from some $\mathbf{r}$ on, a suitable arc in the left directly accessible region of $\mathbf{r}$). The proposition states that no other points can be reached. We split its proof into three parts, corresponding to the three cases.

Case (i): Left chain ends in c.

Proof. We first show that the alternative description is equivalent to the general description. Since $\mathrm{FIL}(P) \cup \mathrm{LFIL}(P, \mathbf{s})$ contains the left discs of all the configurations $\mathbf{r}$ on the extreme left-turning path, no interior point of $\bigcap(\mathrm{FIL}(P) \cup \mathrm{LFIL}(P, \mathbf{s}))$ lies in $\mathrm{LDA}(P, \mathbf{r})$ for any $\mathbf{r}$.

Further, if a point $t \in P$ does not lie in the interior of some disc $F \in \mathrm{FIL}(P)$ then, as is easy to check, it is in the left directly accessible region of some $\mathbf{r}$ on the boundary of $Fil(P)$. Finally, assume that t is a point outside the interior of a disc $D \in \mathrm{LFIL}(P, \mathbf{s}) \setminus \mathrm{FIL}(P)$. Let r be the first point where D touches the left chain, and let $\mathbf{r}$ be configuration on the boundary of P at that point. Define the *forward arc* of D as the part of ∂D contained in the left directly accessible region $\mathrm{LDA}(P, \mathbf{r})$ (Fig. 4). Note that since D is not a part of the filling, it intersects the boundary of P at least twice. Thus, $P \setminus D$ has several components. We claim that those not contained in $\mathrm{LDA}(P, \mathbf{r})$ do not intersect the filling (and consequently they are reachable using some disc $F \in \mathrm{FIL}(P)$).

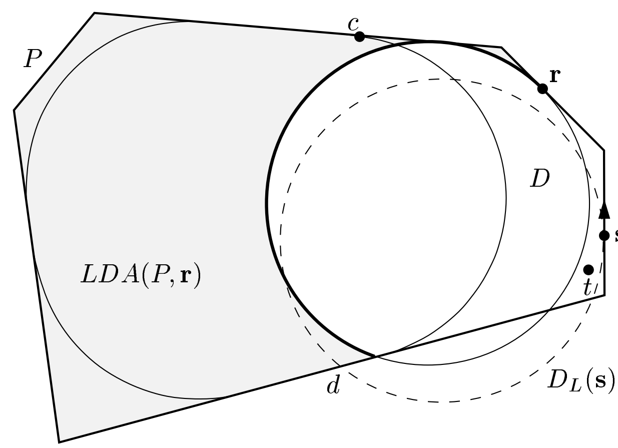

Figure 4: The forward arc.

Indeed a disc $F \in \mathrm{FIL}(P)$ can only intersect the boundary of D in the forward arc.

Let's now prove the general description. Since the left chain ends in c, we have $\mathrm{FIL}(P) \neq \emptyset$ and $c \in \mathrm{LDA}(P, \mathbf{s})$.

Let $\mathcal{D} = \mathrm{FIL}(P) \cup \mathrm{LFIL}(P, \mathbf{s})$, and let t lie in the interior of $\bigcap \mathcal{D}$. We need to show that t is not reachable from $\mathbf{s}$. If s is on the boundary of the filling of P, this is an immediate consequence of Lemma 5. In the sequel, we thus assume that s lies in a pocket of P. For contradiction, suppose that there is path γ from $\mathbf{s}$ to t. Let P' be the union of the pocket of P containing s and of $Fil(P)$. By the Pocket Lemma 3, $\gamma \subset P'$.

Let q be the semiline emanating from s downwards (see Fig. 5). For each inner point $x \in q$, consider the non-vertical tangent $\tau(x)$ to the left disc D_L. Let $x_0 \in q$ be the first point

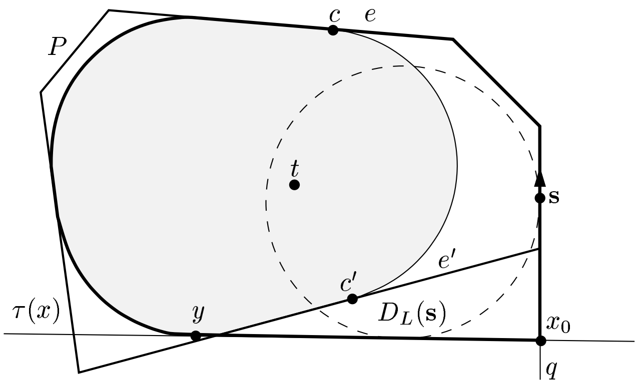

Figure 5: The set P''.

along q such that $\tau(x)$ touches the filling *Fil*(P). (Such an x_0 exist: if x is very close to s then the tangent avoids the filing, and for x receding down infinity, the tangent must intersect the filing, for otherwise we would get a pocket with parallel mouth edges.) Let y be the point of the tangent $\tau(x_0)$ where it first touches the boundary of the filling. Let P'' be the convex set bounded by the portion of the boundary of *Fil*(P) from c to y (counterclockwise), by the segments yx_0 and x_0s, and by the left chain (the portion of the boundary of P from s until c). We need the following two claims.

Claim A. $P' \subseteq P''$.

Claim B. $\mathrm{FIL}(P'') = \mathrm{conv}(\mathrm{FIL}(P) \cup \mathrm{LFIL}(P, \mathbf{s}))$.

First we finish the proof of case (i) of Proposition 6 assuming these claims. By Claim A, the hypothetical path $\gamma \subset P'$ is contained in P''. By Lemma 5, there is no path starting at $\mathbf{s}$ and reaching a point in the interior of $\bigcap \mathrm{FIL}(P'')$. By Claim B, the latter set equals $\bigcap(\mathrm{FIL}(P) \cup \mathrm{LFIL}(P, \mathbf{s}))$, and we get a contradiction to the existence of γ. ◻

Proof of Claim A. It is sufficient to show that the pocket of s in the polygon P is contained in the halfplane bounded by the tangent $\tau(x_0)$ and containing s. Let c and c' be the two mouth points of that pocket (where c is the endpoint of the left chain), let e and e' be the respective mouth edges of P, and let $F \in \mathrm{FIL}(P)$ be the disc touching e at c and e' at c' (see Fig. 5). We need to show that c' cannot precede the touching point y of $\tau(x_0)$ (counterclockwise) along the boundary of *Fil*(P). But if it did, it would be possible to move F slightly, translating its center towards the center of the left disc D_L. Locally, the moved F cannot collide with e', and it cannot collide with the other mouth edge e (containing c), because D_L does not intersect the line containing e. But the moved F would reach into the pocket of s. This contradiction proves Claim A. ◻

Proof of Claim B. It is easy to see that all the discs of $\mathrm{FIL}(P) \cup \mathrm{LFIL}(P)$ are completely contained in P''.

Let $F \in \mathrm{FIL}(P)$ be the disc touching the boundary of P'' at y. By the definition of the boundary of P'', any unit disc contained in P'' and touching its boundary must lie in $\mathrm{FIL}(P) \cup \mathrm{LFIL}(P)$, or have center on the segment connecting the centers of F and D_L. The center of any disc in $\mathrm{FIL}(P'')$ lies on a segment connecting the centers of two discs touching the boundary. This proves Claim B. ◻

Case (ii): Left chain ends in d turning at most π.

Proof. If the left chain turns by at most π, then the left directly accessible region $\mathrm{LDA}(P, \mathbf{s})$ can be enlarged to a pocket of the left disc D_L. Thus, no point outside $\mathrm{LDA}(P, \mathbf{s})$ is reachable by the Pocket Lemma 3. The general description is true since the left disc D_L is the only element of $\mathrm{LFIL}(P, \mathbf{s})$ and the extreme left-turning path goes along its boundary from s to d. (Any unit disc touching the left chain from inside, except for D_L, also intersects the left chain by the interior.) ◻

Case (iii): Left chain ends in d turning more than π.

Proof. In this case the filling must be empty, $\mathrm{FIL}(P) = \emptyset$, for otherwise there would be a disc of the filling touching the left chain; see Fig. 6.

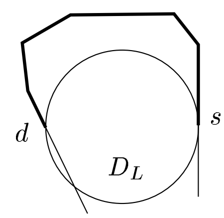

Figure 6: In case (iii) $\mathrm{FIL}(P) = \emptyset$.

Let t be a point reachable from $\mathbf{s}$ by a path γ. First, let's suppose that t belongs to the interior of D_L. Let t' be the first point of γ after s that belongs to the line st. Then t' must lie between s and t. (Otherwise γ would have to reach the segment st' from below later, we could get a loop by extending γ at this point along its tangent, and get a contradiction through the Pestov-Ionin lemma.) Let γ' be the section of γ between s and t'.

Let D_s be the smallest left disc tangent to P at s that intersects $\gamma' \setminus \{s\}$ (see Fig. 7). Then $D_s \subset D_L$. Let $t_0 \neq s$ be a point of $\gamma' \cap D_s$. Let φ be the arc of D_s going counterclockwise from t_0 to s. Let $\mathcal{A}$ be the region bounded by the arc of γ that goes from s to t_0 and by φ. Lemma 2 tells us that there is a unit disc D_0 that lies entirely within $\mathcal{A}$. Let c_0 be its center.

Since $\mathrm{FIL}(P) = \emptyset$, D_0 has to cross the segment st_0, as well as the segment st' (otherwise D_0 would be in the region delimited by the segment st' and γ', which is entirely within P). The center c_0 lies above the line st_0, since D_0 cannot cross φ. Moreover, c_0 lies above the line st, since t_0 lies above st. (See Fig. 8).

Let D_0' be the disc obtained by moving D_0 upwards until it touches the left chain. We have $D_0' \in \mathrm{LFIL}(P)$. Clearly, t is directly accessible through a path that goes around D_0'.

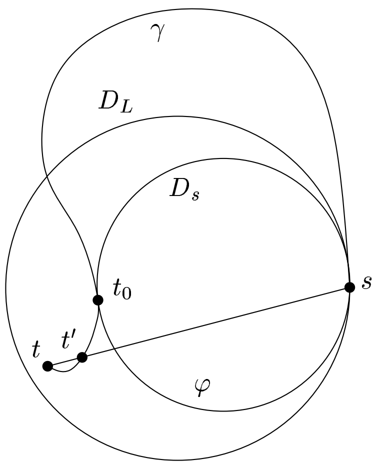

Figure 7: The disc D_s

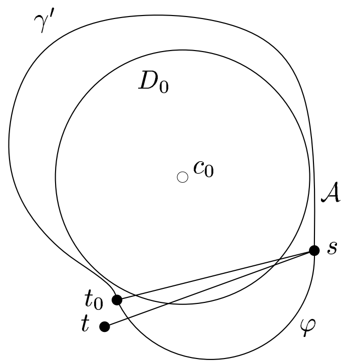

Figure 8: c_0 lies above st

Now let's consider the case where t is outside D_L. If it is in the upper connected component of $P \setminus D_L$, then t is obviously directly accessible. If not, γ has to cut through the interior of D_L, so we can construct a point t_1 in the interior of D_L as above, and use the same argument with t_1 playing the role of t. $\square$

4. Special left-right and right-left paths suffice

Here we show that for determining the reachability by paths in a convex polygon, it suffices to consider paths of a fairly special form.

Let $\mathbf{s}$ be a starting configuration. A *canonical SL-start* from $\mathbf{s}$ is a path from $\mathbf{s}$ to a configuration $\mathbf{r}$ on the boundary of P that begins with a straight segment and continues with a left-turning arc of unit radius with angle smaller than π ending at $\mathbf{r}$ (and tangent to the boundary of P there (Fig. 9). A *canonical RL-start* from $\mathbf{s}$ is defined analogously but the initial segment is replaced by a right-turning arc of unit radius and of angle smaller than π. Note that for each side of the polygon P and for a given $\mathbf{s}$, there is at most one canonical SL-start and at most one canonical RL-start ending on that side. A canonical SR-start and a canonical LR-start

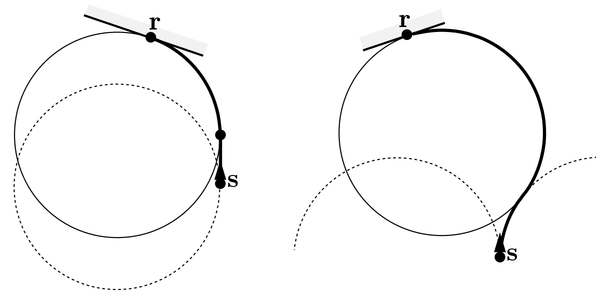

Figure 9: A canonical SL-start and a canonical RL-start.

are defined analogously.

Proposition 7 *Let P be a convex polygon and $\mathbf{s}$ a starting configuration, and let $t \in P$ be reachable from $\mathbf{s}$ by an arbitrary path. Then t lies in the directly accessible region $\mathrm{DA}(P, \mathbf{s})$, or it can be reached by a path of one of the following forms: a canonical SL-start or RL-start followed by a left-turning path (starting on a side of P), or a canonical SR-start or LR-start followed by a right-turning path (starting on a side of P).*

Proof. We use results about shortest paths of bounded curvature. In particular, we need the following facts [1]: Given two configurations $\mathbf{s}$ and $\mathbf{t}$ in a convex polygon P, such that $\mathbf{t}$ is reachable from $\mathbf{s}$ by a path (in our sense), then there is a shortest path from $\mathbf{s}$ to $\mathbf{t}$ consisting of finitely many segments and circular arcs of unit radius (in fact, there are no more than 8 portions [1, Theorem 3.1], but we don't need this fact here). The first and last portions are circular arcs, possibly of zero length. Each of the arcs on the path, except possibly for the first and last portions of the path, turns by an angle strictly greater than π [1, Lemma 2.3].

In our situation, given a point t reachable from $\mathbf{s}$, we consider a shortest path to some configuration $\mathbf{t}$ at the point t with these properties. By a series of transformations, we convert this path into a path of the form described in the proposition.

First, suppose that there is an arc on the path turning by at least π. Then we consider the first such arc α along the path. Let x be the initial point of α and let y be the point on α diametrically opposite to x. We translate the semicircle defined on α by x and y, in direction parallel to the tangents to α at x and y, until it touches the boundary of P; see Fig. 10. In this way, we get a new path that touches the boundary at some point r of the moved arc. Now by Proposition 6, the part of this path from r to t can be replaced by another path that is only turning in the same direction as the arc α.

If α was the first arc of the original path, we already have a path of the form as in the proposition. Let's suppose that α was the second arc, preceded by an arc α_0 and possibly by a segment. Let α_1 be the part of the translated semicircle of α from the initial point until the first point of contact with the boundary of P.

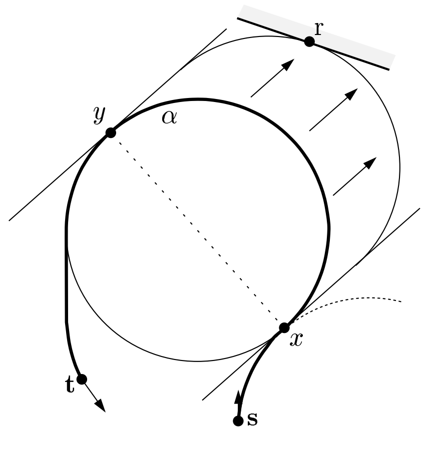

Figure 10: Moving the semicircle until it touches the boundary.

If α_1 and α_0 turn in the same direction, we may assume that their total turning angle is smaller than π, for otherwise we could translate the initial portion of the path that makes the turn by π until it touches the boundary, similar to the transformation made with α. Then we can execute the transformation indicated in Fig. 11, extending the initial segment of the path and shrinking the segment between α_0 and α_1. Either this segment disappears, or we reach a point where the translated arc α_0 touches the boundary of P. In both cases, we have a path as in the proposition.

If α_0 and α_1 turn in different directions, we make the transformation indicated in Fig. 12. Namely, we extend the initial arc α_0, translate the second arc α_1 along the boundary of P while extending it backwards, and rotate and shorten the segment between α_0 and α_1. In this way, we reach a position where the segment between the arcs disappears, or where one of the arcs attains length at least π. In the former case, we have a canonical LR-start or RL-start, and in the latter case, we may apply the transformation described above to the arc with length π.

It remains to consider the situation when the original path has no arc of angle π or greater. In this case, the path consists of at most two arcs and possibly a straight segment between them. Any point reachable by such a path lies in the directly accessible region $\mathrm{DA}(P, \mathbf{s})$. ◻

5. Complexity of the reachable region

The previous two sections enable us to construct the reachable region for any starting configuration in a convex polygon. For a starting configuration on the boundary, we can construct the extreme left-turning chain, and from it the reachable region as the exterior of the intersection of a set of discs. For a starting configuration in the interior, we can construct all possible canonical starts, and take the union of the reachable region for each of them, as well as the directly accessible region.

There are at most $4n$ canonical starts for a given configuration inside a convex polygon with n edges. However, the reachable region for each boundary configuration might involve a linear number of discs. We now prove that, nevertheless, the reachable region has linear complexity.

Proposition 8 *The complexity of the reachable region is $O(n)$, and this bound is tight.*

Proof. The reachable region consists of the directly accessible region of the starting configuration, and a potentially infinite number of LDA or RDA starting along the boundary of P. We claim that given an edge of ∂P, at most two configurations span the LDA of all the reachable configurations of this edge. That implies that the reachable region is the union of at most $4n + 1$ directly accessible regions, and hence has linear complexity.

We shall focus on an edge e that we assume to be vertical and on the right side of P without loss of generality. Let $\mathbf{s}_1$ be the lowest reachable configuration going upwards along e. Let D_1 be the left disc at $\mathbf{s}_1$. Let d_1 be the first intersection point between ∂D_1 and ∂P starting counterclockwise from s_1. $\mathcal{A}_1$ is the connected component of $P \setminus D_1$ directly ahead of $\mathbf{s}_1$.

First we assume that d_1 is in the upper hull (the edges of P locally above the interior of P). Then $\mathcal{A}_1$ is clearly the set of all the LDA starting along e. (Fig. 13)

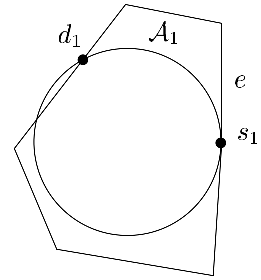

Figure 13: d_1 is in the upper hull

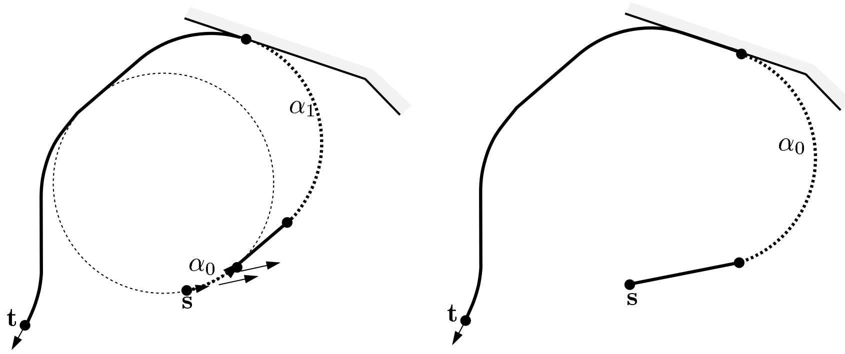

Figure 11: Reducing the initial arcs α_0 and α_1 turning in the same direction.

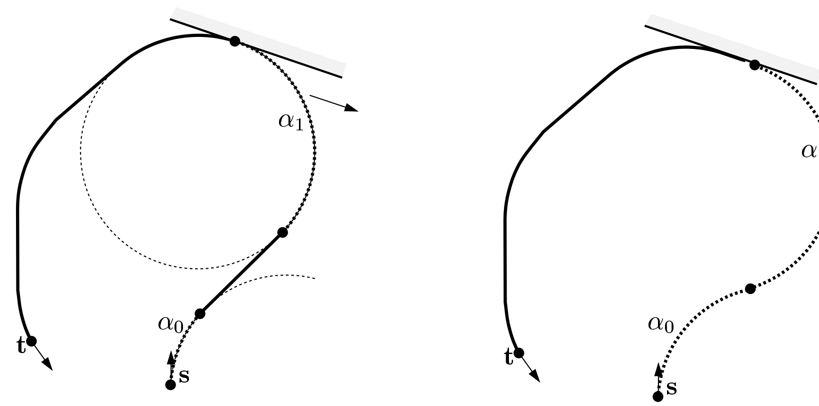

Figure 12: Reducing the initial arcs α_0 and α_1 turning in distinct directions.

Otherwise, for instance if D_1 is a filling disc, then we need another disc, namely D_2 which is obtained by translating D_1 upwards until it touches the upper hull. We denote by $\mathbf{s}_2$ the configuration associated with e and D_2.

The LDA of the configurations lying above s_2 are included in LDA($\mathbf{s}_2$). We claim that for any configuration $\mathbf{s}$ on the segment $s_1 s_2$, we have $\text{LDA}(\mathbf{s}) \subset \text{LDA}(\mathbf{s}_1) \cup LDA(\mathbf{s}_2)$. Let D be the left disc at $\mathbf{s}$ and assume that t is in LDA($\mathbf{s}$). If t is reached after an arc of less than π on D followed by a line segment rt then it is clearly in LDA($\mathbf{s}_1$) (Fig. 14). If t is reached by an arc of at least π on D, followed by a segment rt, we claim that t can be reached from D_2. Let r_2 be the tangent point to t on D_2. We have to argue that the arc from s_2 to r_2 on D_2 lies in P. This follows from the following three facts: D_2 is obtained by translating D upwards until it touches the upper hull of P; the arc from s to r on D is in P; the arc from s_2 to r_2 is shorter than the arc from s to r. ❑

In the full version of this paper we show how to extend this proof to an efficient algorithm for constructing the reachable region.

Acknowledgments

We thank Hazel Everett and Steve Wismath for posing the question discussed in this work, Ngoc-Minh Lê for finding the paper by Pestov and Ionin, and Sylvain Lazard for discussions on his interesting alternate proof of the Pestov-Ionin lemma.

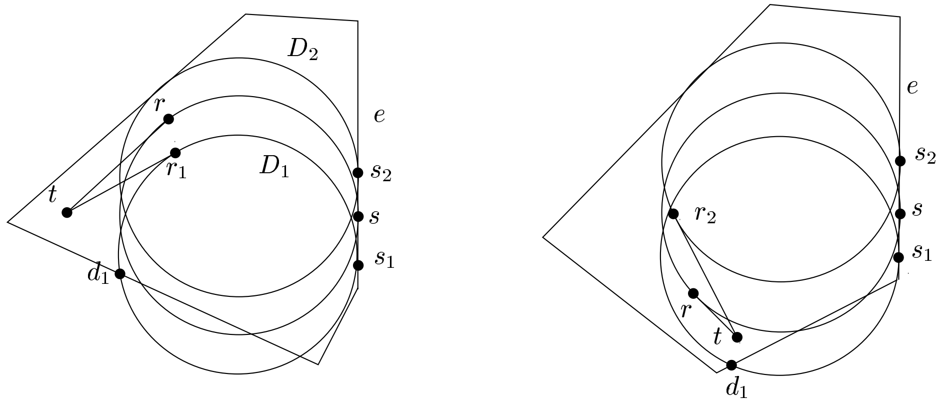

Figure 14: $\mathrm{LDA}(\mathbf{s}_1) \cup \mathrm{LDA}(\mathbf{s}_2)$ **contains all the** LDA **along** e

References

[1] P. K. Agarwal, T. Biedl, S. Lazard, S. e Robbins, S. Suri, and S. Whitesides. Curvature-constrained shortest paths in a convex polygon. In *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, pages 392–401, 1998.

[2] P. K. Agarwal, P. Raghavan, and H. Tamaki. Motion planning for a steering-constrained robot through moderate obstacles. In *Proc. 27th Annu. ACM Sympos. Theory Comput.*, pages 343–352, 1995.

[3] J.-D. Boissonnat, A. Cérézo, and J. Leblond. Shortest paths of bounded curvature in the plane. *Internat. J. Intell. Syst.*, 10:1–16, 1994.

[4] J.-D. Boissonnat and S. Lazard. A polynomial-time algorithm for computing a shortest path of bounded curvature amidst moderate obstacles. In *Proc. 12th Annu. ACM Sympos. Comput. Geom.*, pages 242–251, 1996.

[5] J. Canny, B. R. Donald, J. Reif, and P. Xavier. On the complexity of kinodynamic planning. In *Proc. 29th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 306–316, 1988.

[6] J. Canny, A. Rege, and J. Reif. An exact algorithm for kinodynamic planning in the plane. *Discrete Comput. Geom.*, 6:461–484, 1991.

[7] L. E. Dubins. On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents. *Amer. J. Math.*, 79:497–516, 1957.

[8] S. Fortune and G. Wilfong. Planning constrained motion. *Annals of Math. and AI*, 3:21–82, 1991.

[9] P. Jacobs and J. Canny. Planning smooth paths for mobile robots. In Z. Li and J. F. Canny, editors, *Nonholonomic Motion Planning*, pages 271–342. Kluwer Academic Pubishers, Norwell, MA, 1992.

[10] J.-C. Latombe. *Robot Motion Planning.* Kluwer Academic Publishers, Boston, 1991.

[11] G. Pestov and V. Ionin. On the largest possible circle imbedded in a given closed curve. *Dok. Akad. Nauk SSSR*, 127:1170–1172, 1959. (In Russian).

[12] J. A. Reeds and L. A. Shepp. Optimal paths for a car that goes both forwards and backwards. *Pacific Journal of Mathematics*, 145, 1990.

[13] J. Reif and M. Sharir. Motion planning in the presence of moving obstacles. *J. ACM*, 41:764–790, July 1994.

[14] J. Reif and H. Wang. The complexity of the two dimensional curvature-constrained shortest-path problem. In *Proc. 3rd Workshop Algorithmic Found. Robot.*, page to appear, 1998.

[15] J. T. Schwartz and M. Sharir. Algorithmic motion planning in robotics. In J. van Leeuwen, editor, *Algorithms and Complexity. Handbook of Theoretical Computer Science*, vol. A, pages 391–430. Elsevier, Amsterdam, 1990.

[16] H. J. Sussmann. Shortest 3-dimensional paths with a prescribed curvature bound. In *Proc. Conf. on Decision and control*, pages 3306–3311, 1995.

[17] P. Švestka and M. H. Overmars. Motion planning for car-like robots using a probabilistic learning approach. *Internat. J. Robot. Res.*, 16:119–143, April 1997.

[18] H. Wang and P. K. Agarwal. Approximation algorithms for curvature constrained shortest paths. In *Proc. 7th ACM-SIAM Sympos. Discrete Algorithms*, pages 409–418, 1996.

[19] G. Wilfong. Motion planning for an autonomous vehicle. In *Proc. IEEE Internat. Conf. Robot. Autom.*, pages 529–533, 1988.