Computing Minimum-Link Path in a Homotopy Class amidst Semi-Algebraic Obstacles in the Plane

D. Grigoriev[1♭]
Dept. of Computer Science,
Penn State University, USA
and
Steklov Inst. for Mathematics,
Acad. of Sci. of Russia,
St-Petersburg, Russia

A. Slissenko[2♯]
Dept. of Informatics,
Univ. Paris-12, France
and
Inst. for Informatics,
Acad. of Sci. of Russia,
St-Petersburg, Russia

Abstract. Given a set of semi-algebraic obstacles in the plane and two points in the same connected component of the complement, the problem is to construct a polygonal path between these points which has the minimum number of segments (links) and the minimum 'total turn', that is the sum of absolute values of angles of turns of the consecutive polygon links. We describe an algorithm that solves the problem spending polynomial time to construct one segment of the minimum-link and minimum-turn polygon if to use a modification of real RAMs which permits to handle the solutions of algebraic equations. It is known that the number of segments in such a minimum-link polygon can be exponential as function of the length of the input data or even of the degree of polynomials representing the semi-algebraic set. In fact, we describe how to construct a minimum-link-turn path for a given class of homotopy (whose shortest path has no self-intersections), and provide a rigorous and rather a universal way of reasoning about homotopy classes in contexts related to algorithms. It was previously shown by Heintz-Krick-Slissenko-Solernó that a shortest path in the situation under consideration is semi-algebraic, and an extended real RAM that is able to compute integrals of algebraic functions can find it in polytime.

1 Introduction.

We consider the problem of constructing minimum-link minimum-turn polygon amidst semi-algebraic obstacles in the plane either for a given homotopy class or globally. Usually this problem is motivated by robotics (see, e. g. [HS94, MPA92,

[1] *Address:* Dept. of Comput. Sci., Penn State University, University Park, PA 16802, USA. *E-mail:* grigorie@sol4.cse.psu.edu

[♭] *Supported by a grant of Ministère de l'Enseignement Supérieur et de la Recherche de la France (December 1995 – January 1996)* and by *NSF grant CCR-9424358.*

[2] *Address:* Dept. of Informatics, University Paris-12, 61 Av. du Gén. de Gaulle, 94010, Créteil, France. *E-mail:* slissenko@univ-paris12.fr

[♯] *Partially supported by INTAS-93-2530-ext project.*

MRW92]): if a robot is hard to turn then we try to minimize, firstly, the number of turns, and secondly, the total turn that is the sum of absolute values of all turns. We consider the case of semi-algebraic obstacles in the plane. It is known that one cannot reduce this case to the case of polygonal obstacles, approximating the initial ones, without exponential growth of the complexity of representation (e. g. see [HKSS94, HKSS93]).

1.1 Semi-algebraic obstacles.

A semi-algebraic [BCR87] set S in the plane is a set represented by a disjunctive normal form formula which atoms are polynomial equations and inequalities:

$$\bigvee_{1\leq i\leq N_1} \bigwedge_{1\leq j\leq N_2} f_{i,j}\omega_{i,j}0 \tag{1}$$

where $f_{i,j} \in \boldsymbol{Z}[x,y]$ and $\omega_{i,j} \in \{\leq,<,=\}$.
We treat the set as representing *obstacles*. It may consist of many connected components, and every its connected component will be also called an obstacle. Zero-dimensional componets, i. e. isolated points can hardly be obstacles, so we exclude them from the set. These points can be easily described by a formula of Tarsky algebra with small number of quantifiers and containing polynomials of small number of variables.
As we consider the complexity on a qualitative level, namely, polynomial, exponential, we do not need to go into details of descriptions of the sets under consideration. Our starting set is (1), denote it by S, and the complexity of its representations is measured by the following parameters: d, the maximum of degrees of the polynomials f_{ij}, and M, the least integer such that 2^M bounds the absolute values of the coefficients of f_{ij}. Related sets as the closure $\bar{S}$, the interior S°, the boundary S^β, the complement coS of S, the set of isolated points of S etc. can be represented by formulas of Tarski algebra which complexity (the number of quantifiers, the number of variables in polynomials involved) is bounded by a constant. We need a procedure which, given such a set, constructs its connected components and a procedure that recognizes whether two points are in the same connected component. Within the mentioned context such polytime procedures are known (for general case see, e. g. [HRS90, Gri88, GV92, Ren92], for 2-dimensional case see also [AM88]).
To assure the existence of a shortest path, the space admissible for trajectories is constituted by open space $co\bar{S}$, that is the *free space*, and by the boundary $\Delta =_{df} S^\beta$. The trajectories can go anywhere in the free space, but they are forbidden to go through boundaries though allowed to border them alongside.

1.2 Path Problems.

A *path* or *curve* is a continuous piecewise smooth image of a closed segment. A *simple* path or *quasi-segment* is a path without self-intersections.
Let $\boldsymbol{s}$ and $\boldsymbol{t}$ be two points in the free space. We consider paths between $\boldsymbol{s}$ and $\boldsymbol{t}$ not intersecting obstacles. When speaking about the order of points on a path ψ between $\boldsymbol{s}$ and $\boldsymbol{t}$ we mean *later* or *after* in the sense of a continous, 'length increasing' parametrisation $\psi = \{\psi(t)\}_{0\leq t\leq 1}$ of ψ such that $\boldsymbol{s} = \psi(0)$, $\boldsymbol{t} = \psi(1)$.

A polygonal path is *minimum-link* if it has minimum possible number of links among all polygonal paths between $\boldsymbol{s}$ and $\boldsymbol{t}$ in the same class of homotopy of the free space, and such a path is *globally minimum-link* if its number of links is minimum among all polygonal paths between $\boldsymbol{s}$ and $\boldsymbol{t}$.
Such a minimum-link (for a class of homotopy or globally) polygon always exists, is not unique and may have exponential number of links (for lower bound see [HKSS94, HKSS93], and for upper bound see [HRS94]). The latter fact implies that it is reasonable to estimate the complexity of construction of a minimum-link path in terms of the complexity of building one link.
Turn of two segments $[X, Y]$ and $[Y, Z]$ is the angle between vectors $\overrightarrow{XY}$ and $\overrightarrow{YZ}$. (The angle between these two vectors is the angle between the ray emanating from X in the direction $\overrightarrow{XY}$ and the ray emanating from X in the direction parallel to $\overrightarrow{YZ}$). *Total turn* of a polygonal line is the sum of absolute values of turns of all its consecutive links, say, from $\boldsymbol{s}$ to $\boldsymbol{t}$. Minimum-link path is minimum-turn if its total turn is minimum among all minimum-link paths.
Our construction of a minimum-link path uses as an initial information a shortest path in the same class of homotopy of the free space; to get a globally minimum-link path we start from the graph of 'shortest paths' constructed in [HKSS94]. A shortest path is a path between $\boldsymbol{s}$ and $\boldsymbol{t}$ having the minimum length in its class of homotopy. For the case of semi-algebraic obstacles in the plane it was shown in [HKSS94] that a shortest path can be constructed in polytime by an extended version of real RAM [BSS89] that is able to calculate in polytime integrals of algebraic functions. For the purposes of this paper as a computational model we use another extension which we call a real RAM with equation solver. It allows at a step of computation to compute a root of a polynomial with coefficients computed at previous steps. In the usual real RAM [BSS89] the admitted polynomials have degree 1, so rational operations can be performed. Since a real RAM is a generalization of usual RAM (e. g. [Pap94]), polytime algorithms for bitwise models can be implemented on the model we consider.
Slight development of the construction of [HKSS94] gives an algorithm (for a usual RAM) that outputs a path whose length is ε-close to the shortest paths length and which time complexity is polynomial in the complexity of representation of obstacles and in $\log \frac{1}{\varepsilon}$. The problem of constructing a shortest path in $\boldsymbol{R}^3$ even amidst polyhedral obstacles is known to be NP-hard [CR87].

1.3 Previous results.

The problem of minimizing the number of links of a polygonal path amidst polygonal obstacles and the total turn was studied in the case when obstacles are polygonal, see e. g. [MRW92], [HS94], [CGM$^+$95], and section 4.4.2 in [MS95]. A more complicated problem of minimizing the length of minimum-link paths is considered in [MPA92]. In [HS94] the authors consider also the problem of constructing a minimum-link path for a given class of homotopy, again for polygonal obstacles. However, no article known to us does contain rigorous reasoning about classes of homotopy.

1.4 Our result.

Our contribution to the minimum-link problem consists of, firstly, considering the problem for semi-algebraic obstacles and, secondly, in presenting rather a universal method of rigorous reasoning about paths in relation to classes of homotopy in algorithmic context. The method implies a subroutine for fast testing of homotopicity of two paths.

We prove that for any class of homotopy containing a shortest path without self-intersections and represented by such a path one can construct a minimum-link, minimum-turn path spending polytime to construct a single link of the path. (The requirement of being not self-intersecting is not very restrictive.) And the same is valid for constructing a globally (thus regardless of a class of homotopy) minimum-link, minimum-turn path. In contrast with usual shortest path, where an ϵ-approximation to the length can be found by a usual RAM in time polynomial in the length of input and that of $1/\varepsilon$ (or even of $\log(1/\varepsilon)$ for 2-dimensional case), the problem of approximation of minimum-link paths is more delicate and is not considered here.

2 Shortest and Minimum-Link Paths of a Given Homotopy Class

The basic observation of [HKSS94] (though not stated explicitly) is that the shortest paths between two algebraic points $\boldsymbol{s}$ and $\boldsymbol{t}$, of all classes of homotopy can be represented by a graph of polynomial size with edges weighted by the lengths of pieces of locally convex algebraic curves. The lengths of these pieces can be transcendental. But up to this point the graph can be found in polytime by a RAM (the weights can be computed on an extended real RAM which allows to compute integrals of algebraic functions). Every class of homotopy has exactly one shortest path. We will consider only shortest paths without self-intersections just to simplify the constructions.

2.1 Structure of minimum length paths.

Let S be a semi-algebraic set. We assume (without loss of generality) that the two end points $\boldsymbol{s}$ and $\boldsymbol{t}$ of our paths are inside a square which boundaries are obstacles, and that the points lie in the free space, but for technical reasons are viewed as infinitely small obstacles.

Denote by $B(X,r)$ the open *ball* centered at X and of radius r, and by $\bar{B}(X,r)$ its closure.

A path ψ *intersects the boundary* Δ at its point $X \in \Delta$ if for all small enough $\varepsilon > 0$ there is a closed quasi-segment $\sigma \subseteq \Delta \cap B(X,\varepsilon)$ such that $B(X,\frac{\varepsilon}{2}) \setminus \sigma$ consists of two connected components each containing points of ψ.

It is clear that a shortest path in the free space is a rectilinear segment. If its end meets an obstacle the segment must be locally supporting at the point of contact with the obstacle. Topologically it means that no small enough extension of the segment beyond the point of contact intersects the boundary. One can also define this property in Tarski algebra. A segment $\sigma = [\sigma^-,\sigma^+]$ is *(locally) supporting* to the boundary Δ at a point $X \in \Delta \cap \sigma$ if for every small enough $\varepsilon > 0$ either to the left of or to the right of $\sigma_\varepsilon = [X - \varepsilon(\sigma^+ - \sigma^-), X + \varepsilon(\sigma^+ - \sigma^-)]$ there are no points of the obstacles in $B(X, \varepsilon^2)$. "To the left" and "to the right" can be easiliy described in algebraic terms (e. g. in terms of the sign of an appropriate linear function).

A supporting segment $[X, Y]$ is *maximum to the right* (*to the left*) if any its extension to the right (respectively, to the left) intersects obstacles. By an extension to the right we mean a segment of the form $[X, Y + \varepsilon(Y - X)]$, an extension to the left is similar. A supporting segment is *maximum* if it is maximum in both directions.

A simple path is *locally convex* if the angle function of its tangent vector is monotone.

The *(tangent) angle function* of a smooth piece of the path is a function of, say, length parameter, giving for any point of this piece the (oriented) angle between the tangent vector at this point and some fixed direction.

For a junction point of two smooth pieces one can take as the value of the *angle function* the corresponding one-side limit of this function for any of these pieces.

A path is *globally convex* if it is a part of the boundary of its convex hull. Clearly, when a quasi-segment of a shortest path touches the obstacles and goes along the boundary, this piece of boundary must be locally convex, as well as the quasi-segment on the whole.

In the general case a shortest path is not globally convex, even its locally convex quasi-segment can be not convex because of a too big rotation (imagine a spiral corridor turning several times around some point). The path can change its convexity (i. e. the type of monotonicity of its angle function), but only via an *inflection segment*, i. e. a maximum rectilinear piece of the path such that small enough preceeding and subsequent pieces of the path are separated by the straight line determined by the segment.

The closure of pieces of the path between two consecutive inflection segments are *monotone*; we call a quasi-segment of a path *monotone* if for some its small extension (on the path) its angle func- tion is monotone.

It was shown in [HKSS94] that every shortest path φ consists of a polynomial number of semi-algebraic quasi-segments such that each of them is either an inflection segment of φ locally supporting to Δ at its both ends, or a semi-algebraic monotone quasi-segment which is constituted of pieces of Δ or of rectilinear segments between such pieces, the latter being locally supporting to the boundary at both its ends (imagine going around a circular saw blade as obstacle).

The first and the last segments of the shortest paths under consideration will be treated as inflection segments.

Let φ be a shortest path. Its *standard alternating representation* (or *decomposition*) is the following (finite) sequence D_φ of quasi-segments: the quasi-segments $D_\varphi(2k-1)$, $k \geq 1$, are the consecutive inflection segments of φ; each quasi-segment $D_\varphi(2k)$, $k \geq 1$, is the monotone quasi-segment of φ constituted by the right end of $D_\varphi(2k-1)$, left end of $D_\varphi(2k+1)$ and by the piece of φ between these ends (this piece may be empty).

2.2 Graph of shortest paths.

We can represent the shortest paths of all homotopic classes as a graph $G = G_S$, as it was done in [HKSS94] or in a 'dual' form as follows.
As vertices V of the graph we take $\boldsymbol{s}$, $\boldsymbol{t}$ and all points that are endpoints of segments, that are locally supporting to the obstacles at these endpoints, and which interior lies in the free space. Denote $S^e =_{df} \Delta \setminus V$. Two vertices X and Y of V are connected by an edge if they are two endpoints of a locally supporting segment mentioned above (and this segment is considered as the 'realization' of the edge) or if this is not the case, but X and Y consitute two endpoints of a locally convex connected component of S^e. Denote by E the set of just defined edges. As we do not use lengths, neither of the mentioned connected components nor of the segments, the graph can de found in polytime by a RAM.

2.3 Representation of Homotopy Classes.

We speak about paths between $\boldsymbol{s}$ and $\boldsymbol{t}$ that are *homotopic* in the free space.
Generators. The plane is supposed to be oriented.
As generators of the fundamental group of coS we take cuts that are in a way dual to usual generators as circles [ST80]. (We do not know whether this type of generators was explicitly mentioned elsewhere although it appears to be quite convenient for algorithmical needs.)
Choose in every connected component of the obstacles a point, and launch from it a curve homeomorphic to a ray (we will call these curves *cuts*) such that all the cuts are pairwise disjoint and go to infinity. Attribute to each cut a letter. The set of cuts constitutes a set of generators of a free group.
Now one can define the homotopic type of a path in the plane as follows. Consider the consecutive intersections of the path with the cuts. We assume, without loss of generality, that the path intersects the cuts in isolated points. This sequence of intersections defines the following word: if the ith intersection is with a cut α in the clockwise direction then the ith letter of the word is α, otherwise α^{-1}. Reduce the word as an element of the corresponding free group to the incontractible (irreducible) one. Thus for every path ψ in the plane and for every set of generators F we have defined the word $\Omega_F(\psi)$ that can be considered as a representation of the class of homotopy of ψ (thereby, there is a bijective correspondence between the classes of homotopy and the elements of the free group).
For technical reasons one can consider larger representations taking more points or launching more cuts from each point. The mentioned representation can then be obtained as a homomorphic image of such extended one.

3 Canonical Minimum-Link Polygon of a Homotopy Class.

Our construction of a minimum-link minimum-turn polygon for a given class of homotopy is described below. The construction is rather natural, the key problem is a rigorous proof of correctness of the construction; the situations that may appear are more diverse that one usually could imagine.

3.1 Canonical polygon.

Suppose that the shortest path φ of a homotopy class is given, and it has no self-intersections. We describe an algorithm that constructs a minimum-link minimum-turn polygon belonging to the same homotopy class as the path. This polygonal line will be called the *canonical polygon* of the homotopy class or of the shortest path φ.
The canonical polygon will contain some extension of every inflection segment of the shortest path and will envelop in a minimum way each its monotone quasi-segment up to the next inflection segment.
The canonical polygon is defined by iterations of procedure *NxtSeg*, described below, until reaching the point $\boldsymbol{t}$. The initial data are constituted by $\boldsymbol{s}$, by the first segment of the shortest path which is an inflection segment, and by the subsequent monotone quai-segment that controls the process of extension of the first segment.
The procedure *NxtSeg* constructs the next segment of the canonical polygon starting with its current argument which consists of a current segment $\sigma = [\sigma^-, \sigma^+]$, and of a monotone quasi-segment D of φ 'controlling' the construction of the polygon such that σ^+ lies on $D \cap \Delta$ (and σ is supporting to the obstacles at σ^+). The procedure produces an extension $\tilde{\sigma}$ of the segment σ which is appended to the canonical polygon, a segment σ_1 emanating from its last endpoint and supporting to the boundary, and a 'controlling' quasi-segment $D_\varphi(k_1)$ that may stay unchanged.

3.2 Algorithm for constructing the canonical polygon.

The algorithm *CanonPolyg* for constructing the canonical polygon has as input two points $\boldsymbol{s}$ and $\boldsymbol{t}$ and the shortest path φ of some class of homotopy going from $\boldsymbol{s}$ to $\boldsymbol{t}$. As its result the algorithm outputs the canonical polygon connecting the same two points and homotopic to φ.
The algorithm starts with the first segment emanating from $\boldsymbol{s}$ (that is an inflection one) and tries to advance it as far as possible not leaving the class of homotopy. This is done by the algorithm *NxtSeg*. The iteration of the latter until the point $\boldsymbol{t}$ is reached, constitutes the algorithm *CanonPolyg*. Having finished one application of *NxtSeg* the algorithm *CanonPolyg* gets from *NxtSeg* a segment $\tilde{\sigma}$ that it appends to the canonical polygon, and a segment σ_1 and a number k_1 that it uses as input for *NxtSeg* at the subsequent iteration.
Algorithm *NxtSeg* proceeds as follows, see Figure 1. Suppose that we have advanced up to some segment σ with the last point σ^+ belonging to Δ and to a quasi-segment $D = D_\varphi(k)$ which is not inflection one. Denote by R_σ the ray determined by vector $\overrightarrow{\sigma^-\sigma^+}$ and starting at σ^+. Consider points $X \in R_\sigma$. Such a point may determine the longest segment $[X, Y]$, where Y lies on D after σ^+, that is locally supporting to $D \cap \Delta$ at Y and that does not intersect obstacles. It is unique due to the local convexity of D and absence of self-intersections of φ.
Taking σ^+ as initial value of X we move point X along the the ray R_σ towards

∞ until one of the following (not disjoint) events happens:

(a) The ray R_σ intersects at X the straight line determined by the subsequent inflection segment $D_\varphi(k+1)$.

(b) The segment $[X, Y]$ meets an obstacle at a point $Z \in [X, Y)$.

The first event that happens determines X and Y. The interior of the triangle determined by the segments $[\sigma^+, X]$, $[X, Y]$ and the piece of D between σ^+ and Y is free of obstacles. We extend σ from σ^- to X and append the latter segment to our polygon, and the take $[X, Y]$ as the segment to play the role of σ for the next iteration. If it contains an inflection segment $D_\varphi(k+1)$ we set $D = D_\varphi(k+2)$.

We write this algorithm in a more rigorous form, see Figure 1, to obtain its modification needed to construct globally minimum-link paths in the last section.

NxtSeg(input: σ, k; output: $\tilde{\sigma}$, σ_1, k_1):

-- This procedure extends σ to $\tilde{\sigma}$ which will be appended to Q;

-- constructs the next segment σ_1 to be extended by the subsequent iteration;

-- finds the next controlling monotone quasi-segment $D_\varphi(k_1)$ that can be either $D_\varphi(k)$ or $D_\varphi(k+2)$,

-- the current controlling monotone quasi-segment being represented by $D_\varphi(k)$.

1: Let $D = D_\varphi(k)$; $\beta = D_\varphi(k+1)$.
Denote by R_σ the ray emanating from σ^+ in the direction defined by vector $\overrightarrow{\sigma^-\sigma^+}$; and by R_σ^+ its right 'infinite' end.

2: Construct the following three points $X_1, X_2, X_3 \in R_\sigma$:

a: If R_σ intersects the straight line determined by the inflection segment β, then X_1 is the point of this intersection, otherwise $X_1 = R_\sigma^+$.

b1: X_2 is the first point of intersection of R_σ with the obstacles. It always exists as the obstacles are inside a square.

b2: In the graph of shortest paths look for segments $\xi = [X', Y']$ such that ξ is supporting to D at $Y' \in D \cap \Delta$, ξ is supporting to Δ at X', ξ separates obstacles at X' and Y', the point X' is at the same side of R_σ as Y'. Choose ξ for which the angle between $\overrightarrow{X'Y'}$ and $\overleftarrow{\sigma^-\sigma^+}$ is maximum. If there are several such segments take the longest one. Extend it towards R_σ, then X_3 is the point of its intersection with R_σ.

3: Take as X the point in $\{X_1, X_2, X_3\}$ that is the closest to σ^+. Having found X take as Y the point on D to the right of σ^+ such that $Y \neq \sigma^+$, $[X, Y]$ is the longest segment locally supporting to D at Y and not intersecting obstacls;

4: **if** $[X, Y]$ contains β **then** $k_1 := k + 2$ **else** $k_1 := k$;

5: **return** $(\tilde{\sigma} := [\sigma^-, X]$; $\sigma_1 := [X, Y]$; $k_1)$

Fig. 1. *NxtSeg* algorithm

4 Link and turn optimality of the canonical polygon.

Here we sketch proofs of the main properties of the canonical polygon.

Theorem 1 (Link and turn optimality of the canonical polygon.)
The canonical polygon is minimum-link and minimum-turn.

This theorem is a direct consequence of Corollary 1 and Lemma 2 below.

4.1 Choice of Generators.

Assume that segments of any polygon P under consideration are numbered in the direction from $\boldsymbol{s}$ to $\boldsymbol{t}$ starting with 1. The kth segment of a polygon P will be denoted $P(k)$, and the number of segments of P by $\#P$.
A segment $P(k)$, $2 \leq k \leq \#P-1$, is called *monotone* if the segments $P(k-1)$ and $P(k+1)$ are situated at the same side of the straight line determined by $P(k)$, and is called *inflection* otherwise.
Without loss of generality we assume that the boundary $\varDelta$ everywhere bounds some non empty interior of the obstacles, maybe infinitely thin, thus a point of the boundary can be always moved 'slightly' inside obstacles.
Let $\sigma = [\sigma^-, \sigma^+]$ be a link of the canonical polygon Q of a shortest path φ. Its end σ^+ is a *blocking* point if any extension of σ beyond this point in the direction $\overrightarrow{\sigma^-\sigma^+}$ intersects the obstacles. A pair (α, β) of points of σ is a pair of its *points of rigidity* if $\alpha, \beta \in \varDelta$, β lies on φ, and either α is a *blocking* point of the preceding segment σ' of Q or the obstacles at α and β are separated by σ. The latter means that for some $\varepsilon > 0$ the (non empty) sets $(B(\alpha, \varepsilon) \cap S) \setminus L_\sigma$ and $(B(\beta, \varepsilon) \cap S) \setminus L_\sigma$, where L_σ is the straight line determined by σ, lie on opposite sides of L_σ.
Points of rigidity of a segment will be used in constructions below as starting points of some cuts, and we wish that at this points the cuts do not intersect the segment. For this reason they will be always presumed to be 'slightly' perturbated, that is, a pair (α, β) will be replaced by a pair (α', β') of points situated inside the obstacles close enough to (α, β) (respectively) and to the boundary (to leave them on the appropriate side of σ, the perturbated point of a blocking point must be on the side opposite to the side of β). The meaning of 'close enough' will be clear from the context.
To study the interaction of the canonical polygon Q with another polygon of the same class of homotopy we consider two consecutive segments σ_0 and σ_1 of Q and choose generators of the fundamental group as described below.
Denote by ζ the articulation point of σ_0 and σ_1, i. e. $\zeta = \sigma_0^+ = \sigma_1^-$, and by $\widehat{\zeta}$ the convex (not greater π) angle determined by σ_0 and σ_1 that are considered as emanating from ζ.
We distinguish 4 cases depending on the type of segments σ_0 and σ_1, see Figure 2:
(InfInf) σ_0 and σ_1 are both inflection segments;
(InfMon) σ_0 is an inflection segment and σ_1 is a monotone one;

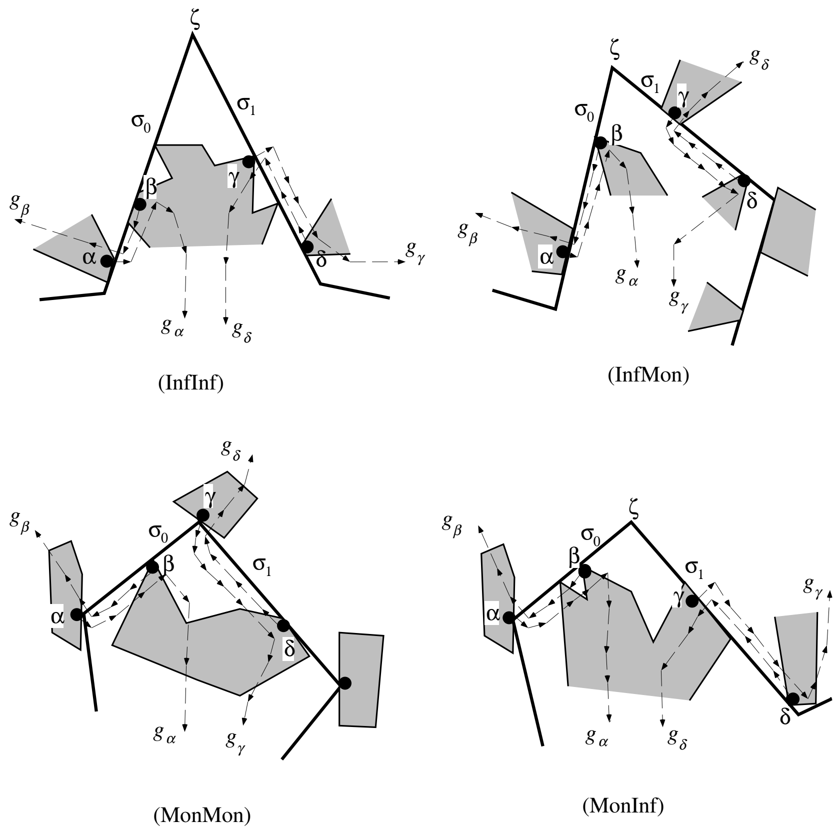

Fig. 2. Choice of generators.

(MonMon) σ_0 and σ_1 are both monotone segments;
(MonInf) σ_0 is an monotone segment and σ_1 is an inflection one.
To choose the points determining the tails of cuts we start from σ_0^- and go towards σ_1^+. Choose two points α and β of rigidity of σ_0, α being before β. If the previous segment has a blocking point, and σ_0 is monotone then α is such a point. Choose two points γ and δ of rigidity of σ_1. Again if σ_0 has a blocking point, and σ_1 is monotone then γ is such a point. We suppose that if γ is a blocking point then it is on the side of σ_0 opposite to β. 'Slightly' perturbate these points as it was described above, preserving the notations for them.
For the pairs (α, β) and (γ, δ) we define 2 cuts in the following way, see Figure 2.
Cut g_α immediately crosses σ_0 and goes along σ_0, infinitely close to it, up to β, then goes to infinity inside $\widehat{\zeta}$.
Cut g_β goes along σ_0, infinitely close to it (between σ_0 and g_α), in the direction

to α, then crosses σ_0 just before α, and after that goes to infinity outside $\widehat{\zeta}$.
If γ is outside $\widehat{\zeta}$ then *cut* g_γ immediately crosses σ_1, and goes along σ_1, infinitely close to it, up to δ, and after that goes to infinity inside $\widehat{\zeta}$.
If γ is inside $\widehat{\zeta}$ then *cut* g_γ immediately crosses σ_1 and goes along σ_1, infinitely close to it, up to δ, and after δ goes to infinity outside $\widehat{\zeta}$.
If δ is outside $\widehat{\zeta}$ then *cut* g_δ goes along σ_1, infinitely close to it (between σ_1 and g_γ) , in the direction to γ, then crosses σ_1 just before γ, and after that goes to infinity inside $\widehat{\zeta}$.
If δ is inside $\widehat{\zeta}$ then *cut* g_δ goes along σ_1, infinitely close to it (between σ_1 and g_γ) up to γ, crosses σ_1 just before γ, and goes to infinity outside $\widehat{\zeta}$.
All the *other cuts* are chosen so that they do not touch the piece of Q between α and β, the mentioned points included. Without loss of generality consider that the cut g_α is crossed by σ_0 *counter clockwise* (see Figure 2).
Denote by V the word describing the homotopy type, i. e. the intersection with the chosen cuts, of prefix Q_α of Q from $\boldsymbol{s}$ to α, the latter point excluded. And by W denote the word of intersection with the cuts of the suffix of Q from δ to $\boldsymbol{t}$, the former point excluded. Thus, the word describing the homotopy type of Q is of the form VUW, where U is constituted only by letters from

$$\Sigma =_{df} \{g_\alpha, g_\alpha^{-1}, g_\beta, g_\beta^{-1}, g_\gamma, g_\gamma^{-1}, g_\delta, g_\delta^{-1}\}.$$

Moreover, we consider that other cuts are chosen in such a way that the last letter V as well as the first letter of W are not from Σ.
For example, for the case (InfInf) the homotopy type of Q in terms of the chosen generators is of the form $V g_\alpha^{-1} g_\beta g_\gamma g_\delta^{-1} W$.

4.2 Proof of Theorem 1.

The link minimality of the canonical polygon follows from Lemma 1.

Lemma 1 *For any polygon P and the canonical polygon Q of the same class of homotopy, for every k, $1 \le k \le \#Q$ there exists m such that $k \le m$, and*
(a) $P(m)$ intersects $Q(k)$ in a point X lying between any two rigidity points of $Q(k)$,
(b) the prefix Q_X of Q starting at $\boldsymbol{s}$ and ending at X and the prefix P_X of P starting at $\boldsymbol{s}$ and ending at X have the same homotopic type,
(c) $P(m)$ arrives at crossing with $Q(k)$ from the side opposite to any point α that can constitute a pair (α, β) of points of rigidity of $Q(k)$.

Proof. We proceed by induction on the number k of segments of Q. The case $k = 1$ is evident.
Suppose the lemma is valid for the first k segments of Q.
Consider the non trivial case when the kth segment is not the last one, then it has a consecutive one. Denote the kth and $(k+1)$st segments of Q by σ_0 and σ_1 respectively. Choose the cuts to represent homotopies as described above.
Let $P(m_0)$ be the last segment of P which crosses Q at a point X lying between α and β and satisfies the induction hypothesis. It enters σ_0 from the side opposite to α, and Q_X and P_X have homotopy type represented by the word

$Vg_\alpha^{-1}g_\beta$. Within the choice of generators, the homotopy type of Q is of the form $Vg_\alpha^{-1}g_\beta GW$, where V, G and W are of the type described above. In particular, if σ_1 is monotone then $G = g_\gamma^{-1}g_\delta$, and if σ_1 is inflection then $G = g_\gamma g_\delta^{-1}$. Let $G = g'g''$, where g' and g'' are defined as has been just described depending on the case under consideration.
Now, to assure the demanded homotopy type the polygon P must realize g' (that is to cross g_γ from the appropriate side) and then realize g''.
Suppose that such a crossing is done outside of $[\gamma, \delta]$. Note that each pair of cuts, (g_α, g_β) and (g_γ, g_δ) determine a cut of the plane up to an infinitely narrow corridor between points α and β (and, respectively, γ and δ) with 'walls' constituted by the corresponding pieces of g_α and g_β (respectively, of g_γ and g_δ). And the segment $P(m_0)$ has entered X in a 'half' plane determined by (g_α, g_β). To enter the other 'half' plane without crossing $[\alpha, \beta]$ it must go through the corridor from the side of β, but this is an obstacle, so this is impossible. Because of a similar reason it is impossible to penetrate in the corridor from the side of α.
Now, the point X is in the intersection of two 'half' planes determined by (g_α, g_β) and (g_γ, g_δ). To respect the homotopy type the polygon P must realize an intersection with cuts that gives $g'g''$.
Suppose that σ_1 is an inflection segment (the case when it is monotone is completely similar). Is it possible to have $g_\gamma g_\delta^{-1}$ without crossing $[g_\gamma, g_\delta]$? Clearly not, because to do it the polygon P must penetrate the corridor between g_γ and g_δ that is impossible.
Thus, P must cross (g_γ, g_δ) from the side opposite to γ, and the last such crossing must have the homotopy type $Vg_\alpha^{-1}g_\beta G$. Clear, this crossing can not be realized by $P(m_0)$, but by a further segment of P.
□

Corollary 1 *Every canonical polygon is minimum-link.*

Corollary 2 *For any minimum-link polygon P and the canonical polygon Q of the same class of homotopy, for any k the kth segments of Q and P are either both monotone or inflection ones, and the kth segments of P and Q intersect each other.*
Corollary 1 is a straightforward consequence of Lemma 1, (a) and its condition on m.
Corollary 2 is implied by Lemma 1 and minimality of P. The latter imposes an intersection of $P(k)$ and $Q(k)$.

Lemma 2 *For any minimum-link polygon P and the canonical polygon Q of the same class of homotopy, and for any point X of intersection of an inflection segment of the canonical polygon Q with the corresponding (i.e. having the same number) inflection segment of P, the total turn of Q_X is not greater than the the total turn of P_X, where Q_X and P_X are the prefixes of respectively Q and P between s and X.*

Proof. It is sufficient to compare turns between two consecutive inflection segments of the canonical polygon Q and that of the corresponding segment of

another minimum-link polygon P, and to take into consideration Lemma 1, (c). Let σ_0 and σ_1 be two consecutive inflection segments of the canonical polygon Q. It is obvious, that the total turn of all monotone segments between σ_0 and σ_1 is equal to the total turn from σ_0 to σ_1 (it can be arbitrarily large). The basic observation is illustrated on the Figure 3: ξ_0 and ξ_1 are segments of P corresponding to Q (in the sense of Lemma 1). The angles to compare are φ (total turn modulo π from σ_0 to σ_1), and ψ (total turn modulo π from ξ_0 to ξ_1). Clearly, $\varphi \leq \psi$, because the triangle XYZ is inside of the triangle XYV.
□

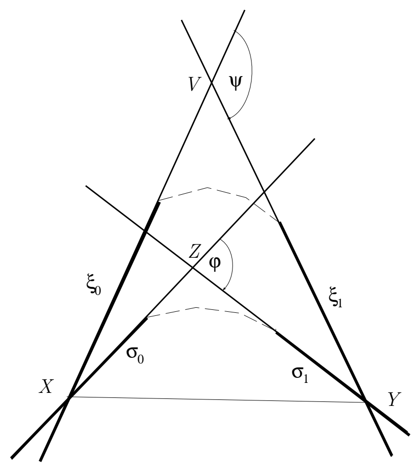

Fig. 3. Comparing total turns.

5 Algorithm of Global Search and its Complexity.

Now we describe a wave algorithm which constructs a minimum-link path between $\boldsymbol{s}$ and $\boldsymbol{t}$. It is based on the standard idea, as, for example, in Dijkstra's shortest path algorithm.
The algorithm starts from the graph of the shortest paths, denote it by $G = (V, E)$. It advances simultaneously along all the paths constructing the corresponding canonical polygons link by link. The number of links gives the current value of weight, and the algorithm builds only polygons of the minimum weight. As landmarks of this search it takes inflection segments which direct the wave

propagation and permit to determine a quasi-segment for applying slightly modified *NxtSeg* algorithm to calculate the link distance of the front of the current wave from $\boldsymbol{s}$.

To make this idea more precise denote by $G_{infl} = (V_{infl}, E_{infl})$ the following graph of inflection segments that will control the procedure. Let V_{infl} be the set of all pairs $[xy]$ of vertices of G corresponding to inflection segments $[x, y]$ (note that a segment cannot be a monotone segment for one path and be an inflection segment for another path). Two vertices $[ux]$ and $[yz]$ form an edge of E_{infl} iff the points x and y are connected by a path in G constituting a monotone quasi-segment. If two inflection segments are incident, i. e. have a common end point, then we consider them to be connected by an edge whose weight is 1.

This quasi-segment connecting the vertices will be called the *realisation* of the edge, and will be denoted by $D(\Gamma)$ for an edge Γ. One can remark that the canonical polygon is defined by consecutive inflection segments independently of its other parts as stated in Lemma 3.

Lemma 3 *For any two adjacent vertices $[uv]$ and $[xy]$ of G_{infl} and whatever be simple paths φ and ψ of G containing as subpath the path from u via v via $D([uv][xy])$ via x to y, the pieces of their canonical polygons that lie between u and y are equal.*

The piece of canonical polygon mentioned in Lemma 3 can be constructed by *NxtSeg* algorithm slightly modified in the following way.

Denote the modified version *NxtSegM*. This algorithm is applied each time when the global algorithm finds an edge (α, β) of G_{infl} to analyse. *NxtSegM* starts from α being directed by the quasi-segment corresponding to (α, β), i. e. by $D(\alpha, \beta)$. It has as arguments a segment σ and the quasi-segment $D = D(\alpha, \beta)$ that is supposed to be not an inflection segment. The quasi-segment D can be empty if α and β are incident. The algorithm returns a segment playing the role of σ for the next iteration. One can recognize when it has finished the piece of canonical polygon between α and β; this happens when σ contains β.

NxtSegM has an input σ, D, β and as output σ, where $D = D(\alpha, \beta)$ is a monotone quasi-segment directing the canonical polygon, and σ is a segment as in *NxtSeg*. The description of

$NxtSegM(input : \sigma, D, \beta; output : \sigma)$ can be obtained from the description of *NxtSeg* by expunging the initial comments in lines 3-5, the operators 1 and 4, and by deleting $\tilde{\sigma}$ and k in operator 5.

To uniformize the description of the algorithm of global search append to G_{infl} two vertices $[\boldsymbol{s}\,\boldsymbol{s}\,]$ and $[\boldsymbol{t}\,\boldsymbol{t}\,]$ and connect them by zero-weight edges with all vertices corresponding to inflection segments of G emanating from $\boldsymbol{s}$ and $\boldsymbol{t}$ respectively.

The algorithm *MnLnkPath* which construct a minimum-link polygon, and in fact all such polygons in a usual compressed form, proceeds as follows:

(a) All the edges of G_{infl} are classified as "with calculated weight" and "with not yet calculated weight". Initially, the only edges with calculated weight equal to 0, are those that emanate from $[\boldsymbol{s}\,\boldsymbol{s}\,]$ and $[\boldsymbol{t}\,\boldsymbol{t}\,]$.

(b) The algorithm advances as in Dijkstra's algorithm, starting from $\boldsymbol{s}$ but only

via edges of G_{infl} with calculated weights. Having reached $\boldsymbol{t}$ it stops.
(c) The algorithm keeps a set of reached vertices with calculated weights which constitute the front of the wave. For each edge outgoing from such a vertice α it calculates its weight by constructing one segment of the canonical polygon, and hence augments its weight by one for all edges with not yet calculated weights. When canonical polygon reaches the corresponding inflection segment, the weight is calculated. If after this adding of 1 the set of edges with calculated weights has augmented, the algorithm make one step of the Dijkstra algorithm, maybe advancing from front vertices one edge farther to not explored vertices. And iterates this procedure.
One can represent minimum-link paths by pointers which permit extract one such path in time polynomial in the number of its links.

Theorem 2 *Algorithm MnLnkPath constructs a minimum-link path in time polynomial in the number of links of the path.*

Theorem 3 *A minimum-link path with minimum turn can be found in time polynomial in the number of links of the path.*

Acknowledgements. The first author is thankful to the Department of Informatics of the University Paris-12 for the hospitability during his stay there for two months in 1995 and 1996.

References

[AM88] M. E. Alonso and Raimondo M. The computation of the topology of a planar semialgebraic set. *Rend. Sem. Mat. Univers. Politecn. Torino*, 46(3):327–342, 1988.

[BCR87] J. Bochnak, M. Coste, and M.-F. Roy. *Géométrie algébrique réelle.* Springer-Verlag, 1987.

[BSS89] L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over real numbers: NP-completeness, recursive functions and universal machines. *Bull. Amer. Math. Soc.*, 1:1–46, 1989.

[CGM+95] V. Chandru, S. K. Ghosh, A. Maheshwari, V. T. Rajan, and S. Saluja. NC-algorithms for minimum link path and related problems. *J. of Algorithms*, 19:173–203, 1995.

[CR87] J. Canny and J. Reif. New lower bound technique for robot motion planning problemms. In *Proc. 28th Annu. IEEE Symp. on Foundations of Comput. Sci.*, pages 49–60, 1987.

[Gri88] D. Yu. Grigoriev. Complexity of deciding Tarski algebra. *J. Symb. Comput.*, 5:65–108, 1988.

[GV92] D. Yu. Grigoriev and N. N. Vorobjov. Counting connected components of a semi-algebraic set in subexponential time. *Computational Complexity*, 2(2):133–184, 1992.

[HKSS93] J. Heintz, T. Krick, A. Slissenko, and P. Solernó. Une borne inférieure pour la construction de chemins polygonaux dans $\boldsymbol{R}^n$. In *Publications du département de mathématiques de l'Université de Limoges*, pages 94–100. Université de Limoges, France, 1993.

[HKSS94] J. Heintz, T. Krick, A. Slissenko, and P. Solernó. Search for shortest path around semialgebraic obstacles in the plane. *J. Math. Sciences*, 70(4):1944–1949, 1994. Translation into English of the paper published in *Zapiski Nauchn. Semin. LOMI*, vol. 192(1991), p. 163-173.

[HRS90] J. Heintz, M.-F. Roy, and P. Solernó. Sur la complexité du principe de Tarski-Seidenberg. *Bull. Soc. Math. de France*, 118:101–126, 1990.

[HRS94] J. Heintz, M.-F. Roy, and P. Solernó. Single exponential path finding in semi-algebraic sets. part 2: the general case. In Ch. L. Bajaj, editor, *Algebraic Geometry and Its Applications*, 1994.

[HS94] J. Hershberger and J. Snoeyink. Computing minimum length paths of a given homotopy class. *Computational Geometry*, 4:63–97, 1994.

[MPA92] J. S. B. Mitchell, C. Piatko, and E. M. Arkin. Computing a shortest k-link path in a polygon. In *Proc. 33rd IEEE FOCS*, pages 573–582, 1992.

[MRW92] J. S. B. Mitchell, G. Rote, and G. Woeginger. Minimum-link paths among obstacles in the plane. *Algorithmica*, 8:431–459, 1992.

[MS95] J. S. B. Mitchell and Suri S. *A survey on computational geometry*, volume 7 of *Hanbooks in Operations Research and Management Sciences*, chapter 7, pages 425–479. Elsevier Science B. V., 1995.

[Pap94] C. H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.

[Ren92] J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. parts 1-3. *J. Symb. Comput.*, 13(3):255–352, 1992.

[ST80] H. Seifert and W. Threlfall. *A Textbook of Topology*. Academic Press, 1980.