Minimizing the endpoint trace length of rod motions amidst polygonal obstacles is NP-hard

Extended Abstract

Tetsuo Asano[1], David Kirkpatrick[2]*, and Chee K. Yap[3]†
[1]School of Information Science, JAIST, Japan,
[2]University of British Columbia, Canada,
[3]Courant Institute, New York University, USA

Abstract

We continue, and in a sense complete, our study the motion of a rod (line segment) in the plane in the presence of polygonal obstacles, under an optimality criterion based on minimizing the trace length of a fixed but arbitrary point (called the *focus*) on the rod. In an earlier paper, we showed that this problem is *NP*-hard when the focus is in the relative interior of the rod. Our proof did not cover the case where the focus lies at one of the rod endpoints. Indeed, considerable evidence suggested that this special case might admit a polynomial time solution. In this paper we settle this open problem by proving, by means of a non-trivial adaptation of our earlier construction, that this remaining case is also *NP*-hard.

1 Introduction

Although the *feasibility* of motion planning is very well studied, comparatively little is known about *optimal* motion planning. The exceptions to this concern the cases where the robot body is a point (or, more generally, a sphere) or a *rod* (a directed line segment). The only known efficient general algorithms apply to the case where the robot body is a disc in the plane. Other cases are hard; specifically:

(A) In three dimensional space, the problem of determining the shortest path for a point robot moving amidst polyhedral obstacles is *NP*-hard[4];

(B) In the plane, the problem of minimizing the"d_1- distance" of a rod motion (defined to be the length of the trace of a specified point (the *focus*) on the rod) in the presence of polygonal obstacles is *NP*-hard, when the focus lies in the rod interior [1].

For rod motions, d_1-distance is certainly not the only reasonable notion of distance to try and optimize. One natural alternative is that of maximizing the clearance (*i.e.*, distance to the nearest obstacle). Here, an efficient algorithm based a generalization of the Voronoi diagram is known [11, 12]. Another interesting approach is to minimize the area swept by the motion of the rod. This problem attracted some interest in the past under the name "Kakeya's problem" (see [3]), but this turns out not to be a good idea. (For example, one can sweep an arbitrarily small area in a motion that moves a rod from any position to its complementary position, which occupies the same points in space. The obvious way to achieve this is to rotate the rod by 180° about its midpoint.)

Other notions of optimality can be formulated in terms of other concepts of distance or length. If X is any fixed point on the rod, the curve traced by X in any continuous motion μ of the rod is called the *trace* of X in μ. One natural choice here is to minimize the *average* lengths of the traces of the two endpoints of the rod. In the absence of obstacles, this has been called *Ulam's problem* [13]. Again it has an interesting history (see [5] and the references therein). The paper of Icking *et al.* [5] revisits this problem, introducing a simple tool based on Cauchy's surface area formula. They call the metric in Ulam's problem the d_2-metric. There is a natural generalization to d_n for any $n \geq 3$ or $n = \infty$ (corresponding to minimizing the average trace length of n evenly distributed points on the rod).

Other than [5], there are few previous papers on d_n-optimal motions. Papadimitriou and Silverberg [7] studied d_1-optimal motions with the focus F at one endpoint of the rod. However, they severely restricted the motions so that F travels only in straight lines between obstacle vertices. Their results were improved by Sharir [10]. O'Rourke [6] studied d_∞-optimal motions restricted to either pure translations or rotations by $\pm 90°$.

In contrast to these last cited papers, we are interested in *unrestricted* motions of the rod, except of course when the rod collides with obstacles. In this paper we settle the remaining open problem concerning d_1-optimal motion of a rod in the plane. Specifically, we show the following:

*Supported in part by NSERC grant #3583
†Supported in part by Supported by NSF/ITR Grant #CCR-0082056.

Theorem 1 *The problem of determining if their exists a collision free motion of a rod, from a specified initial configuration to a specified final configuration, whose d_1-length, with the focus at one of the rod endpoints, is at most some specified value D, is NP-hard.*

This result is somewhat surprising giving the considerable evidence that suggested that this problem (essentially that of finding optimal motions of a point with a line segment "tail") would admit a polynomial time solution.

2 The Construction

Our *NP*-hardness construction is an non-trivial modification of our earlier construction [1] that establishes the *NP*-hardness of determining d_1-optimal motions when the focus is in the rod interior. Like the previous construction our modified construction exploits the characterization of d_1-optimal motions developed in [1] (see the appendix of [2] for a proof). In particular, it makes essential use of the fact that polygonal obstacles induce a collection of curves that behave like mirrors in the sense that in optimal motions the trace of the focus point visits and "reflects" off these mirror curves (in accordance with Snell's law). Figure 1 illustrates a simple polygonal obstacle, its associated displaced features (mirrors), and two optimal motions from initial configuration S to target configurations T_1 and T_2.

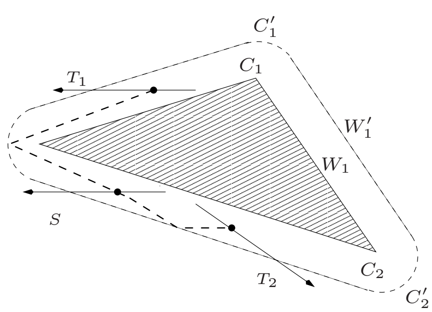

Figure 1: Displaced features (mirrors) and reflecting motions

If optimal motions never include more than some fixed constant number of reflections between stopovers (configurations where the focus lies at one of polynomially many locations determined by the obstacle set and the initial rod configuration), then it it is straightforward to demonstrate a polynomial time algorithm for constructing optimal motions, by reduction to a shortest path search in a "visibility graph" on the set of stopovers. (Indeed this is precisely the approach that gives a polynomial time algorithms for finding shortest paths for a disc, and for finding ε-approximation algorithms for d_1-optimal paths for a rod [2]. It is also what made the existence of a polynomial time algorithm for minimizing the trace of a rod endpoint quite plausible.)

Unfortunately, as in the case where the focus is in the rod interior, there exist polygonal obstacle sets, with a total of n vertices, with respect to which d_1-optimal motions may involve $\Theta(n)$ consecutive reflections. Indeed, our *NP*-hardness proofs are based on a obstacle set for which there exist placements with exponentially many distinct optimal connecting motions, all of which involve a sequence of $\Theta(n)$ reflections and *no* stopovers.

Our *NP*-hardness proofs involve a polynomial time reduction from 4CNF-satisfiability. Specifically, suppose Φ is a formula in 4CNF involving m clauses and the k variables $X_0, \ldots, X_{k-1}$. We show how to construct a polygonal environment E, whose description is bounded in size by some polynomial in k, together with free placements S and T, and a distance D, such that there exists a collision-free motion from S to T whose d_1-length is at most D if and only if Φ is satisfiable.

The overall structure of our reduction is similar to the Canny-Reif proof [4] that the shortest-path problem (for a *point* amidst polygonal obstacles) in 3-dimensions is *NP*-hard: A basic environment is designed that admits 2^k topologically distinct d_1-minimal motions between two specified placements; these paths are associated with distinct truth assignments to the variables $X_0, \ldots, X_{k-1}$; and finally, the environment is augmented with some additional obstacles that serve to block (filter) every path whose associated truth assignment does not satisfy Φ.

Our construction is necessarily different from that of Canny and Reif since our problem is set in two dimensions. As indicated above, the key to our construction is to exploit the mirror-like properties of reflection curves (displaced features). Our construction is modular in the sense that it consists of an assembly of certain pre-fabricated modules. In fact, our new construction borrows heavily from the collection of modules used in our earlier construction; we essentially need just one new module together with a slightly more involved (more global) analysis.

The reader is referred to [1] for details of the full construction. Here we concentrate on the modifications needed to make the modules work in the case where the focus is at one of the rod endpoints. Additional details will be presented in the full version. Each module is a collection of line segment barriers together with certain distinguished points, called *terminals*. Terminals play the dual role of attachment points for neighbouring modules and checkpoints on (potential) shortest paths. The trace of the rod focus

F, as the rod follows shortest paths between placements in our modules, is referred to as a *beam*. Beams that connect terminals are called *canonical* beams. There is just one basic module from which several others are fabricated:

Wide beam splitter $WBS(\lambda)$. This module has one input terminal a and two output terminals b_0 and b_1, with a vertical separation of λ units (the *separation factor*). Let μ_x denote an optimal motion between the horizontal rod placements H_a and H_x (with the focus at a and x respectively). Then, for all points x on the line $\overline{b_0 b_1}$ through b_0 and b_1, the d_1-distance of μ_x is minimized exactly when $x = b_0$ or $x = b_1$. We denote this minimum distance by σ. (If the maximum value of λ is fixed then σ can be fixed independent of λ). Figure 2 gives a schematic description of this module; the details of its construction are exactly where the current paper differs from [1].

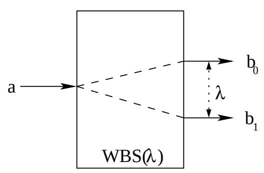

Figure 2: Wide Beam Splitter Schematic

In fact, we need a more general property to hold for our wide beam splitter that says that the same splitting property holds for a sufficiently narrow cluster of parallel beams (a *sheaf*). Specifically, if a' (respectively, b'_0, b'_1) is the point υ units above a (respectively, b_0, b_1), where $|\upsilon|$ is sufficiently small, and μ^{υ}_x denotes an optimal motion between the horizontal placements $H_{a'}$ and H_x then, for all points x on the line $\overline{b_0 b_1}$, the d_1-distance of μ^{υ}_x is minimized exactly when $x = b'_0$ or $x = b'_1$. Furthermore, this minimum distance is equal to σ (independent of υ).

The left-right mirror image of a wide beam splitter behaves like a wide beam combiner. It is not hard to imagine (see [1] for details) that wide beam splitters and combiners (with suitably chosen values of parameter λ) can be composed to form a module that splits a single beam into a sheaf of 2^k parallel beams all of which correspond to d_1-optimal motions from a single initial configuration. Composing this multi-beam splitter with its mirror image produces a module with exactly 2^k distinct minimum length beams joining two specified points a and a' (corresponding to d_1-optimal motions from initial placement H_a to final placement $H_{a'}$).

The reduction proceeds by interpreting the 2^k distinct paths as possible truth assignments to the variables of a given k-variable formula Φ. Additional modules inserted between the multi beam splitter and combiner are designed to filter out all but those beams whose associated truth assignments satisfy Φ. The result is a module which admits a beam traversal of a specified length joining two specified points (corresponding to d_1-optimal motions joining two specified rod configurations) if and only if the formula Φ is satisfiable.

3 Realizing the components

When the focus lies in the rod interior, a wide beam splitter module can be built out of two simpler components called turn "gadgets". The basic (optional turn) gadget (cf. Figure 3) behaves like a half-silvered mirror: a beam originating at point I can proceed straight through to point E_0 or can reflect and exit at point E_1.

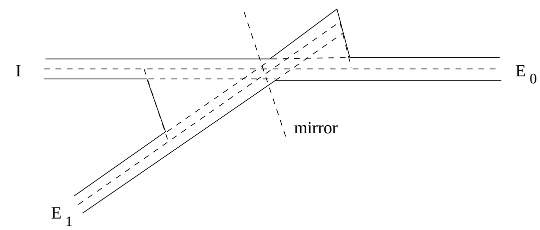

Figure 3: Optional Turn Gadget

A simple modification (blocking the exit to E_0) changes this into a forced turn gadget. Four of these turn gadgets (one optional and three forced) are combined into a wide beam splitter, as shown in Figure 4. Note that the length of the paths traced by the split beams can be easily adjusted (by modifying the length of the channels) to achieve some fixed value (σ) independent of λ.

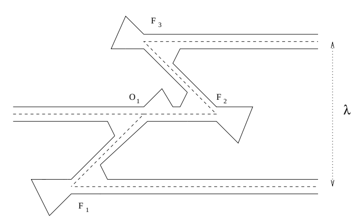

Figure 4: Mirror-based wide beam splitter

Unfortunately, when the focus lies at one of the rod endpoints this construction fails because there is only one endpoint that induces a mirror curve and the construction relies on (alternating) mirror reflections with respect to both rod endpoints. The essential new gadget permits a sequence of mirror reflections with respect to the same endpoint. This gadget, together with the rod (with focus at the tail end) in six successive positions in transit, is illustrated in Figure 5 below.

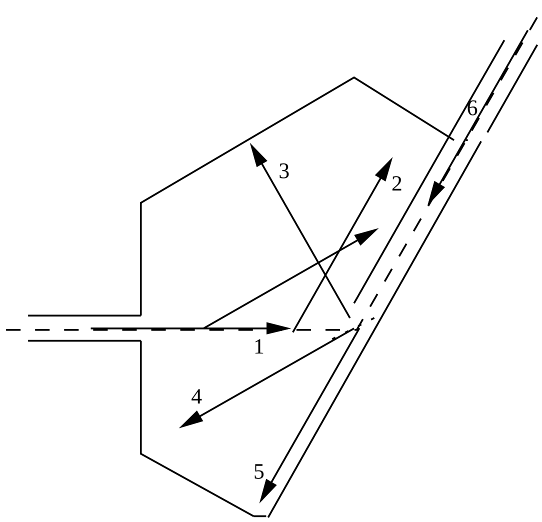

Figure 5: New Turn Gadget

The new turn gadget achieves not only a forced turn (by $\pi/3$ as illustrated) but also (importantly) a reorientation of the rod: the rod enters with the focus end behind, and exits with the focus end in front. This, alternating with our conventional $2\pi/3$ optional and forced turn gadgets (the alternation ensuring that the reflection always occurs with respect to the correct end of the rod) allows us to build a wide beam splitter that exploits reflection behaviour at only one end of the rod.

References

[1] T. Asano, D. Kirkpatrick, and C. Yap. d_1-Optimal motion of a rod. *12th ACM Symposium on Computational Geometry*, pages 252–263, 1996.

[2] T. Asano, D. Kirkpatrick, and C. Yap. Pseudo approximation algorithms with applications to optimal motion planning. *18th ACM Symposium on Computational Geometry*, pages 170–178, 2002.

[3] A. S. Besicovitch. On Kakeya's problem and a similar one. *Mathematische Zeitschrift*, 27:312–320, 1928.

[4] J. Canny and J. H. Reif. New lower bound techniques for robot motion planning problems. *IEEE Foundations of Computer Science*, 28:49–60, 1987.

[5] C. Icking, G. Rote, E. Welzl, and C. Yap. Shortest paths for line segments. *Algorithmica*, 10:182–200, 1993.

[6] J. O'Rourke. Finding a shortest ladder path: a special case. IMA Preprint Series 353, Institute for Mathematics and its Applications, University of Minnesota, 1987.

[7] C. H. Papadimitriou and E. B. Silverberg. Optimal piecewise linear motion of an object among obstacles. *Algorithmica*, 2:523–539, 1987.

[8] J. H. Reif, J. D. Tygar, and A. Yoshida. Computability and complexity of ray tracing. *Discrete Comput. Geom.*, 11:265-287, 1994.

[9] J. T. Schwartz and M. Sharir. On the piano movers' problem: I. the case of a two-dimensional rigid polygonal body moving amidst polygonal barriers. *Communications on Pure and Applied Mathematics*, 36:345–398, 1983.

[10] M. Sharir. A note on the Papadimitriou-Silverberg algorithm for planning optimal piecewise-linear motion of a ladder. NYU Robotics Report 188, Courant Institute, New York University, 1989.

[11] M. Sharir, C. O'D'únlaing, and C. Yap. Generalized Voronoi diagrams for moving a ladder I: topological analysis. *Communications in Pure and Applied Math.*, XXXIX:423–483, 1986.

[12] M. Sharir, C. O'D'únlaing, and C. Yap. Generalized Voronoi diagrams for moving a ladder II: efficient computation of the diagram. *Algorithmica*, 2:27–59, 1987.

[13] S. M. Ulam. *Problems of Modern Mathematics*. Science Editions, New York, 1964. Originally published as: *A Collection of Mathematical Problems*, Interscience Publishers, New York, 1960.

[14] Chee Yap. Algorithmic motion planning. In J. T. Schwartz and C. K. Yap, editors, *Advances in Robotics, Vol. 1: Algorithmic and geometric issues*, chapter 3. Lawrence Erlbaum Associates, 1987.