アルゴリズム 58－8
（1997. 9. 19）

線状ロボットの d_1－最適な移動問題 (2)

浅野哲夫[1]、D. G. Kirkpatrick[2]、C. K. Yap[3]

[1]北陸先端科学技術大学院大学（石川県辰口町）
[2]University of British Columbia（Vancouver, Canada）
[3]New York University（New York, U.S.A.）

本論文では、多角形の障害物を避けて線状のロボットを移動するのに、線上に任意ではあるが固定された点（参照点とよぶ）の軌跡の長さを最小化する問題の計算複雑度を扱う。本文では、このような線状ロボットに回転や平行移動などの任意の動作を許した場合について、参照点の軌跡を最小にする d_1-最適な動作を特徴づける定理に基づいてこの問題が NP-困難であることを示す。

d_1－Optimal Motion for a Rod (2)

T. Asano[1]、D. G. Kirkpatrick[2]、C. K. Yap[3]

[1]School of Computer Science, JAIST (Tatsunokuchi, Ishikawa, Japan)
[2]Dept. of Computer Science, University of British Columbia（Vancouver, Canada）
[3]Courant Institute, New York University（New York, U.S.A.）

We 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 orbit length of a fixed but arbitrary point (called the *focus*) on the rod. In this paper, we prove NP-hardness of the problem of finding such an optimal motion, based on the local characterization of a d_1−optimal motion of a rod that minimizes the orbit length of a reference point, allowing arbitrary kinds of motions including rotation and translation.

1 Introduction

Although the *feasibility* of motion planning is very well studied, little is known about *optimal* motion planning except for the case where the robot body is a disc. In this paper we address the problem of characterizing and computing optimal motions for a *rod* (a directed line segment) in the plane. Of course, a rod is the next simplest planar body to study.

A non-trivial issue arising in the study of "optimal" motion of a rod is the choice of a reasonable yet tractable motion of optimality. Turning to notions of optimality based on some distance or length concept, we may describe such optimal motions as "shortest". 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 *orbit* of X in μ. One natural choice here is to minimize the *average* lengths of the orbits of the two endpoints of the rod. In the absence of obstacles, this has been called *Ulam's problem* [5]. As Icking *et al.* [3] note, the "d_1-distance", based on minimizing the orbit length of the mid-point of the rod, is not a metric (rotation about the mid-point produces distinct rod placements with d_1-distance zero). It is nevertheless a rather natural measure of distance, capturing the idea of "charging" for translation but not for rotation about the mid-point. In this paper, we generalize this "d_1-distance" to refer to the family of distance functions based on minimizing the orbit length of a fixed but arbitrary point F (the "focus") in the relative interior of the rod.

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. Motion planning problems come in two variations depending on whether the two endpoints of the rod are distinguished (in which case the rod is said to be *directed*) or not (in which case the rod is said to be *undirected*). Clearly, motion planning problem for undirected rods can be reduced to two instances of the motion planning problem for directed rods. (The converse relationship is not known.) Unless otherwise stated we assume that rods are directed.

This paper proves NP-hardness of the problem of finding such an optimal motion, based on the local characterization of a d_1-optimal motion of a rod that minimizes the orbit length of a reference point, allowing arbitrary kinds of motions including rotation and translation.

1.1 Local Characterization of Optimal Motion

Intuitively, the trace of an optimal motion must travel in a straight line unless it must bend around a convex corner or it is constricted (in which case there is no choice but to trace the conchoid or elliptic curves). But there is one other possibility, namely, the trace can reflect off a displaced feature in accordance to Snell's law. For more details refer to the previous presentation [?]. Our characterization result is summarized as follows:

Theorem 1 *Any optimal motion μ can be transformed into a motion μ' such that $F\mu = F\mu'$ and μ' consists of a finite sequence of $O(n^4)$ submotions in which each submotion has one of the following forms:*
1. pure rotation around a pivot
2. pure translation along a straight line segment
3. dragging an endpoint along a wall in a straight trace
4. dragging the rod along a convex corner in a straight trace
5. sliding along two walls in an elliptical trace
6. sliding along a wall and a corner in a conchoidal trace.

2 *NP*-Hardness: the problem with mirrors

If optimal motions never included more than some fixed constant number of reflections between stopovers, then it would be straightforward to demonstrate a polynomial time algorithm for constructing optimal motions. Unfortunately, this is not so. In fact there exist obstacle sets with respect to which optimal motions may involve $\Theta(n)$ consecutive reflections. Indeed, the *NP*-hardness proof outlined in this section is 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.

For the remainder of this section, we fix rod AB to have unit length and a focus F at distance α from A. We assume that α is a rational satisfying $\alpha > 0$ and, without loss of generality, $\alpha \leq 1/2$. We formulate the decision version of our optimal motion problem as a quadruple $\langle E, s, t, D \rangle$, where E is a polygonal environment (obstacle set), specified by a listing of its walls; s and t are free placements of the rod with respect to E; and D is distance. All coordinates, distances and angles in the specification are assumed to be rational[1]. The problem is to decide if there is a free motion from s to t of d_1-distance at most D.

We show that this decision problem is *NP*-hard by describing 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 $\langle E, s, t, D \rangle$ is a yes-instance of our decision problem if and only if Φ is satisfiable.

The overall structure of our reduction is similar to the Canny-Reif proof [2] 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 distnict shortest motions between two specified placements; these paths are associated with distinct truth assignments to the

[1] An angle is said to be *rational* if both $\sin\phi$ and $\cos\phi$ are rational numbers.

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. Fo convenience we express distances in our modules as integer multiples of some rational *unit* distance. (It will suffice to to choose the unit to be $\alpha/2^{(2k+2)}$.) We describe modules in terms of our assumed orientation of the rod, but of course analogous modules exist for the rod in its opposite orientation. In fact, some of our modules, those we describe as "inverse" modules, are formed by mirror image of one of these opposite-orientation modules. 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:

(A) **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 *separating factor*). Let μ_x denote 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_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). Fig. 1 gives a schematic description of this module; the details of its construction are deferred to the next section.

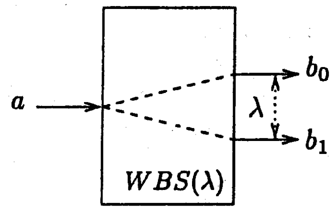

Fig. 1. Wide Beam Splitter.

The left-right mirror image of a wide beam splitter (for a rod with the opposite orientation) behaves like a beam combiner. In fact, our wide beam splitter satisfies a more general property: if a' (respectively, b'_0, b'_1) is the point v units above a (respectively, b_0, b_1), where $|v|$ is sufficiently small, and μ^v_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^v_x is minimized exactly when $x = b'_0$ or $x = b'_1$. Furthermore, this minimum distance is equal to σ (independent of v). Thus, our wide beam splitter (respectively, its mirror image) also behaves as a *wide sheaf splitter* (respectively, a *wide sheaf combiner*, where a *sheaf* is a cluster of horizontal beams. Fig. 2 illustrates this sheaf-splitting/combining property of $WBS(\lambda)$. As in all of our subsequent illustrations of modules, the figure includes (scaled-down) schematic descriptions of each module (which are used in describing other modules).

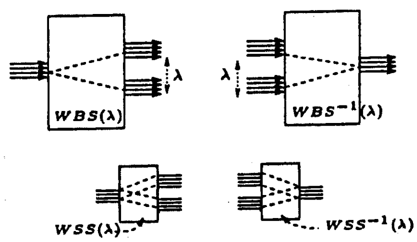

Fig. 2. Wide Sheaf Splitter and Combiner (with schematics).

We call a collection of 2^i horizontal beams, such that adjacent beams are separated by δ units, an (i, δ)-sheaf.

(B) **Narrow sheaf splitter** $SS(j, \delta)$. A wide sheaf splitter and a (slightly smaller) wide sheaf combiner can be composed to form a narrow sheaf splitter. Specifically, $SS(j,\delta)$ takes a (j,δ)-sheaf as input and produces a $(j+1,\delta)$-sheaf as output, by passing the sheaf in sequence through a $WSS(\lambda)$ and a $WSS^{-1}(\lambda - 2^j\delta)$-module (see Fig. 3).

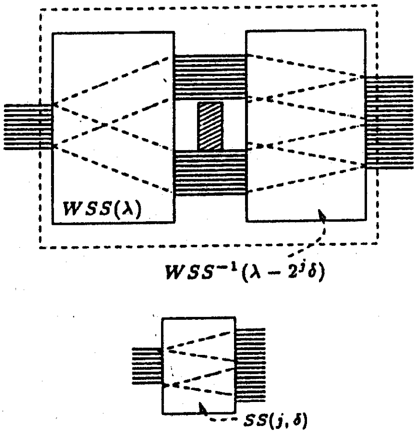

Fig. 3. Narrow Sheaf Splitter $SS(j,\delta)$.

The mirror image of this narrow sheaf splitter behaves as a narrow sheaf combiner (for a rod with the opposite orientation). Narrow sheaf splitters are also useful for manufacturing (i,δ)-sheaves:

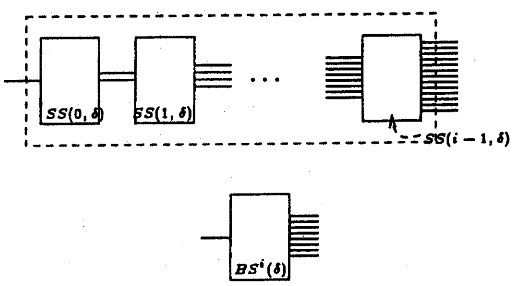

Fig. 4. Narrow Beam Splitter $BS^i(\delta)$.

(C) **General beam splitter** $BS^i(\delta)$. The basic idea, as illustrated in Fig. 4, is to compose a sequence of i narrow sheaf splitters. The j-th splitter has separating factor $\lambda_j = 2^{i-j}\delta$. The module $BS^i(\delta)$ has one input terminal a and 2^i δ-separated output terminals $b_0 \dots b_{2^i-1}$. It follows immediately from the properties of the constituent wide beam splitters that (i) the shortest motion from horizontal placement H_a to a horizontal placement H_x, with $x \in b_0\bar{b}_{2^i-1}$ has length greater than $i\sigma$, unless $x \in \{b_0, \dots, b_{2^i-1}\}$, and (ii) the shortest motion from H_a to any one of $H_{b_0}, H_{b_1}, \dots, H_{b_{2^i-1}}$ is exactly $i\sigma$.

We use a $BS^k(\delta)$ module to generate a (k,δ)-sheaf. In general the beams in a (k,δ)-sheaf will carry labels that we interpret as truth assignments. These label assignments (from the top to bottom beam in a sheaf) are always of the form:

$$\begin{array}{l}
X_iX_{i+1}\cdots X_{i+k-2}X_{i+k-1},\\
X_iX_{i+1}\cdots X_{i+k-2}\overline{X_{i+k-1}},\\
X_iX_{i+1}\cdots \overline{X_{i+k-2}}X_{i+k-1},\\
X_iX_{i+1}\cdots \overline{X_{i+k-2}X_{i+k-1}}\\
\quad\vdots\\
\overline{X_iX_{i+1}}\cdots \overline{X_{i+k-2}}X_{i+k-1}\\
\overline{X_iX_{i+1}}\cdots \overline{X_{i+k-2}X_{i+k-1}}
\end{array} \tag{1}$$

for some i, $0 \le i < k$ (where indices are understood to be reduced mod k). This is referred to as an $(X_i, X_{i+1}, \dots, X_{i+k-1})$-labelling. The output sheaf of $BS^k(\delta)$ is assigned an $(X_0, X_1, \dots, X_{k-1})$-labelling.

The mirror image of $BS^k(\delta)$, which we denote by $BS^{-k}(\delta)$, serves to combine the beams in a (k,δ)-sheaf into a single output beam. It has 2^k input terminals $b_0 \dots b_{2^k-1}$ and a single output terminal c. Furthermore, (i) the shortest motion from a horizontal placement H_x, with $x \in b_0\bar{b}_{2^i-1}$, to the horizontal placement H_a, has length greater that $i\sigma$, unless $x \in \{b_0, \dots, b_{2^i-1}\}$, and (ii) the shortest motion from any one of $H_{b_0}, H_{b_1}, \dots H_{b_{2^i-1}}$ to H_a is exactly $i\delta$.

Clearly, composing $BS^k(\delta)$ and $BS^{-k}(\delta)$, by identifying the identically labeled terminals, yields a module with 2^k distinct shortest motions from initial placement H_a to final placement H_c. The remaining modules are designed to separate $BS^k(\delta)$ and $BS^{-k}(\delta)$ in this trivial composition, blocking all but those beams whose associated truth assignments satisfy the formula Φ.

(D) **Elementary shuffle module** $SH(k,\delta)$. Consider the wide sheaf splitter illustrated in Fig. 2. If obstacles are placed to block the upper half of the upper sheaf and the lower half of the lower sheaf, and the resulting (half)-sheaves are input to a (slightly smaller) inverse sheaf splitter, the result is that these two (half)-sheaves are interleaved (shuffled) into a new full sheaf (of twice the density of the original); see Fig. 5. We refer to the resulting module as the elementary shuffle module $SH(k,\delta)$. This name derives from the fact that if the input sheaf to a an elementary shuffle module has an $(X_i, X_{i+1}, \ldots, X_{i+k-1})$-labelling, the output sheaf has an $(X_{i+1}, X_{i+2}, \ldots, X_{i+k})$-labelling.

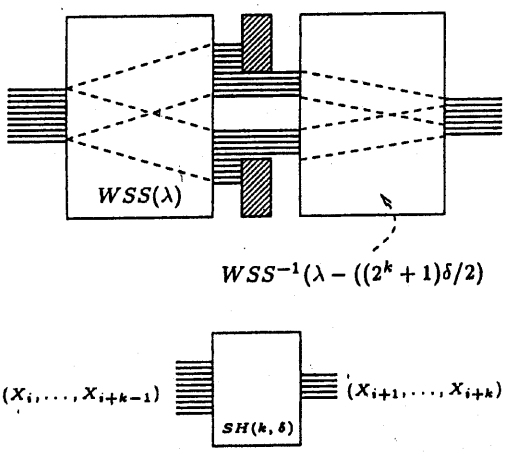

Fig. 5. Elementary Shuffle Module $SH(k,\delta)$

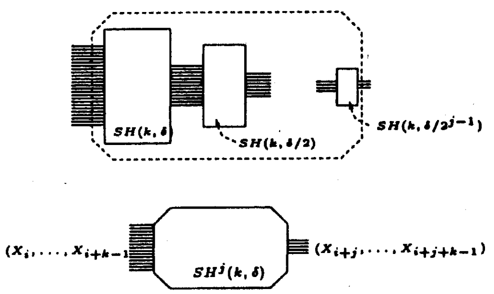

Fig. 6. General Shuffle Module $SH^j(k,\delta)$.

(E) **General shuffle module** $SH^j(k,\delta)$. Just as a narrow beam splitter is formed from narrow sheaf splitters by composition, a general shuffle module $SH^j(k,\delta)$ is formed by composing a sequence of j elementary shuffle modules, $SH(k,\delta), SH(k,\delta/2), \ldots, SH(k,\delta/2^{j-1})$. Note that $SH^j(k,\delta)$ takes a (k,δ)-sheaf with an $(X_i, X_{i+1}, \ldots, X_{i+k-1})$-labelling, and produces a $(k, \delta/2^j)$-sheaf with an $(X_{i+j}, X_{i+j+1}, \ldots, X_{i+j+k-1})$-labelling. In this output sheaf all beams whose labels contain X_{i+j} (respectively, $\bar{X}_{i+j}$) appear contiguously, permitting them all to be blocked with a single obstacle. The mirror image of $SH^j(k,\delta)$, which we denote by $SH^{-j}(k,\delta)$ will play the role of a sheaf "unshuffler".

(F) **Literal filter** $LF(k, X_j)$. The module $LF(k,X_j)$ has $(k,2^k)$-sheaves as input and output. It is formed by composing the modules $SH^j(k,2^k)$, $SH^{k-j}(k,2^{k-j})$, and $SH^{-k}(k,2^k)$. In addition, a vertical obstacle is placed in such a way that it blocks the lower half of the $(k,2^{k-j})$-sheaf joining the first two modules (see Fig. 7). The literal filter $LF(k,\bar{X}_j)$ is identical, except that the added obstacle blocks the upper half of the same sheaf. It should be clear that $LF(k,L)$ obstructs exactly those beams that correspond to truth assignments fail to satisfy the literal L.

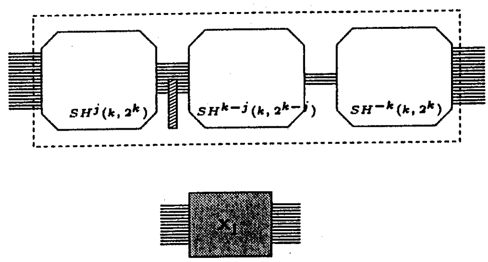

Fig. 7. Literal Filter.

(G) **Clause filter** $CF(k,C)$. A clause C of the formula Φ is the disjunction of four literals, say L_0, L_1, L_2, and L_3. The module $CF(k,C)$ has $(k,2^k)$-sheaves as input and output. It is formed by the composition of three sheaf-splitters (with sufficiently large splitting factors), the four literal filters $LF(k,L_0), \ldots, LF(k,L_3)$, and three inverse sheaf-splitters, as shown in Fig. 8. Filter $CF(k,C)$ obstructs exactly those beams that correspond to truth assignments that fail to satisfy the clause C.

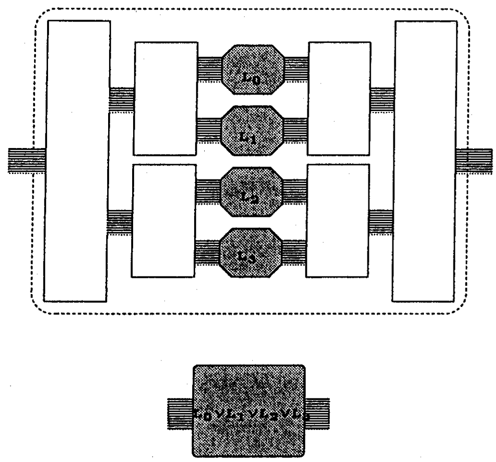

Fig. 8. Clause Filter.

(H) **Formula filter** $FF(k, \Phi)$. Suppose that formula Φ is the conjunction of clauses C_0, ..., C_{m-1}. The module $FF(k, \Phi)$ has $(k, 2^k)$-sheaves as input and output. It is formed by the composition of clause filters $CF(k, C_0), \ldots, CF(k, C_{m-1})$, as shown in Fig. 9. Filter $FF(k, \Phi)$ obstructs exactly those beams that correspond to truth assignments that fail to satisfy the CNF-formula Φ.

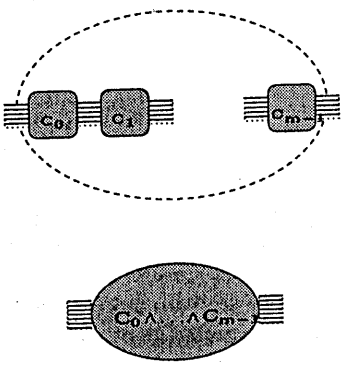

Fig. 9. Formula Filter.

If we compose the formula filter $FF(k, \Phi)$ with the beam splitter $BS^k(2^k)$ and inverse beam splitter $BS^{-k}(2^k)$, as shown in Fig. 10, we produce a module that admits a canonical beam from initial placement H_a to final placement H_b if and only if formula Φ is satisfiable. It is straightforward to confirm that all canonical beams in our full construction have length $D = (m(4k+4)+2k)\sigma$, and a motion of length at most D exists if and only if at least one canonical beam is unfiltered. It remains to prove that our elementary modules can be constructed, and that the bit-complexity for their specification is polynomial in the input size n (which can be assumed to be $\Theta(m)$). This is taken up in the next section.

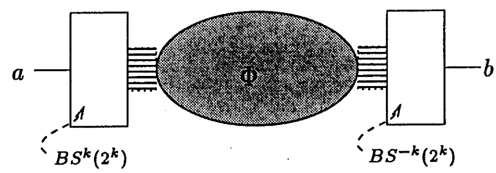

Fig. 10. Full Reduction.

3 Realizing the components

Our elementary modules are in turn built up of smaller components, that we will prefer to call "gadgets". There are basically two kinds of gadgets: a half-silvered mirror gadget and a fully-silvered mirror gadget. Let θ be a fixed rational angle less than $\tan^{-1}(1/3)$ (for example, $\theta = \tan^{-1}(7/24)$). This angle is used in our mirror constructions.

We first describe the *half-silvered mirror gadget*, $H(I, C, E^0, E^1, \Delta)$ where I, C, E^0, E^1 are points such that $|IC| = |CE^0| = |CE^1| = 2$ and $\angle(I, C, E) = 2\theta$. It is best to describe the gadget in terms of its "skeleton", shown in Fig. 11.

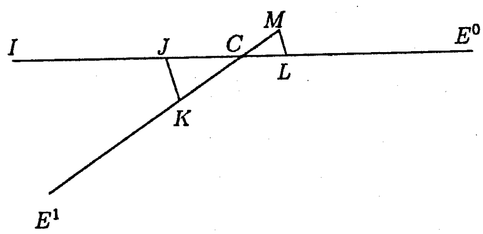

Fig. 11. Skeleton of Half-silvered Mirror $H(l, C, E^0, E^1, \Delta)$.

Note that points I, J, C, L and E^0 are all colinear, as are the points E^1,K, C, and M. Furthermore, $|JC| = |KC| = (1-\alpha)/\cos\theta$ and $|LC| = |MC| = \alpha/\cos\theta$, so the normal distance from C to $\bar{JK}$ (resp. $\bar{ML}$) is $(1-\alpha)$ (resp. α). For the actual gadget, we essentally thicken the line segments $[I, E^0]$ and $[E^1, M]$ into channels of width Δ. This gadget is shown in Fig. 12, where the skeleton (in dashed lines) is superimposed on the actual lines of the gadget.

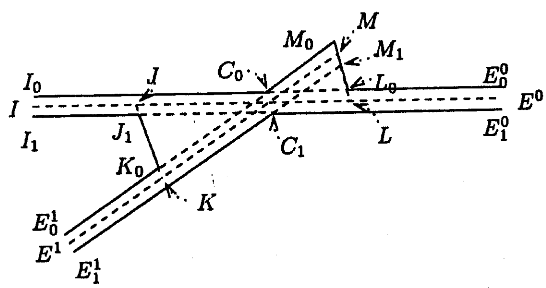

Fig. 12. Half-silvered Mirror $H(I, C, E^0, E^1, \Delta)$.

The segment $[C_0, C_1]$ is called the "neck" of the gadget. We choose $\Delta = \alpha \sin\theta$ so that the following condition holds: When $[C_0, C_1]$ is projected orthogonally onto $[J, K]$, it lies between J_1 and K_0.

This "neck property" of our construction guarantees that the B-displaced wall associated with segment $[J_1, K_0]$ spans the channel along line $\overline{C_0C_1}$.

Refering to Fig. 12, for $0 < s < 1$, let $C_s := sC_1 + (1-s)C_0$ denote an arbitrary point on the neck. Similarly, let $I_s := sI_1 + (1-s)I_0$, $E^0_s := sE^0_1 + (1-s)E^0_0$, and $E^1_s := sE^1_1 + (1-s)E^1_0$. Let Z_s, Z^0_s, Z^1_s be free placements such that

(a) $F[Z_s] = I_s$, $F[Z^0_s] = E^0_s$, $F[Z^1_s] = E^1_s$,

(b) $A[Z_s]$, $A[Z^1_s]$, and $B[Z^0_s]$ lie inside the enclosure[2] of the gadget.

Lemma 1 *(a) If $\mu : Z_s \longrightarrow Z^0_t$ has minimum d_1-distance among all motions that start at Z_s and end at a free placement Z^0_x, with $0 < x < 1$, then $t = s$ and the trace $F[\mu]$ consists of the line segment $[I_s, E^0_s]$.*

(b)If $\mu : Z_s \longrightarrow Z^1_t$ has minimum d_1-distance among all motions that start at Z_s and end at a free placement Z^1_x, with $0 < x < 1$, then $t = s$ and the trace $F[\mu]$ consists of the 2 line segments $[I_s, C_s]$, $[C_s, E^1_s]$.

Proof: Part (a) is immediate. For part (b), observe that our choice of θ precludes motions whose trace does not cross the neck. Hence, μ must go through a placement Z_1 where $AB[Z_1]$ is perpendicular to the line $\overline{JK}$, and $B[Z_1]$ touches $\overline{JK}$. Because of the neck property, it follows that $F[Z_1]$ lies on the neck. In fact, Z_1 is unique and $F[Z_1] = C_s$, by our characterization theorem. The lemma follows. Q.E.D.

Remarks:

(1) We call $H(I, C, E^0, E^1, \Delta)$ a "half-silvered" mirror since it is clear that an incoming beam through I_s can proceed to either (directly) to E^0_s or (by reflection) to E^1_s.

(2) There is an analogous gadget where the roles of A- and B-ends are interchanged. We denote this gadget by $H'(I, C, E^0, E^1, \Delta)$. Clearly we have $H'(I, C, E^0, E^1, \Delta) = H(I, C, E^0, E^1, \Delta)$ in case $\alpha = 1/2$.

(3) We can seal off the exit E^0 in Fig. 12, by introducing the segment $[E^0_0, E^0_1]$, thus obtaining a "fully-silvered mirror". We denote this gadget by $G(I, C, E^0, E^1, \Delta)$.

(4) Finally, note that any mirror images of the above gadgets obtained by reflecting about a vertical or horizontal line will also be denoted by the same notation $H(I, C, E^0, E^1, \Delta)$, etc.

Roughly speaking, our beam splitter is constructed out of one half-silvered mirror (h_0) and three fully-silvered mirrors (g'_1, g_2, and g'_3) connected as shown in Fig. 13.

Note that the lenth 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.

This completes the construction. We summarize our result in the following:

[2] The enclosure of this gadget is the region inside the polygon obtained by "sealing" off the terminals in Fig. 12, by introducing the line segments $[I_0, I_1]$, $[E^0_0, E^0_1]$ and $[E^1_0, E^1_1]$.

Theorem 2 *It is NP-hard to compute an optimal motion of a rod in the presence of polygonal obstacles when a focus is in the relative interior of the rod.*

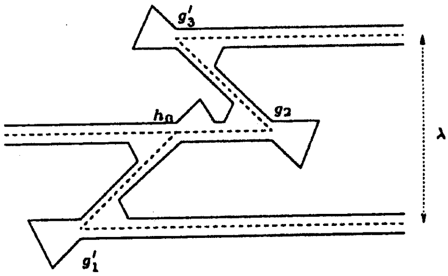

Fig. 13. Mirror-based beam splitter.

Remark It is interesting to note that our proof of Theorem 2 above shows that the ray-tracing question of determining whether a beam directed from a specified source, in a scene comprised of silvered and half-silvered linear mirrors, reaches a specified destination by a path of length bounded by some specified value, is *NP*-hard. The corresponding problem in three-dimensional space was shown to be PSPACE-hard by Reif *et al.* [4]. They left the two-dimensional version as an open problem.

References

[1] T. Asano, D.G. Kirkpatrick and C.K. Yap. "d_1-optimal Motion for a Rod," Technical Report of IPSJ, AL-57-4, pp.25-32, May, 1997.

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

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

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

[5] 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.