The Complexity of Fine Motion Planning

B. K. Natarajan

TR 86-734
February 1986

Department of Computer Science
Cornell University
Ithaca, NY 14853

THE COMPLEXITY OF FINE MOTION PLANNING

B.K. Natarajan
Dept. of Computer Science
Cornell University
Ithaca, NY 14850

ABSTRACT

We study the complexity of fine motion planning for robots with position measurement and damping. A reduction from fine motion planning with position measurement only to the "classical piano mover's problem" is developed, thereby showing it to be feasible in polynomial time. We then show that deciding the existence of fine motion plans for robots with damping in three dimensional scenes is PSPACE-hard and, with a view to finding the cause for the jump in complexity, we identify a restricted subclass of the PSPACE-hard problem that is PSPACE-complete. Finally, we show how to restrict this subclass to permit polynomial time algorithms for the problems in it.

1. INTRODUCTION

Motion planning for robots has received considerable attention in recent years. But the attention has been rather partial to gross motion planning at the cost of fine motion planning, which is a little studied although equally important problem. Loosely speaking, the term *gross motion* is applicable to situations in which the uncertainty in the relative position of the goal and the object to be moved is negligible, while the term *fine motion* is applicable to situations in which this uncertainty is significant. As an example, consider the typical problem of inserting a peg in a hole: Fig.1a shows

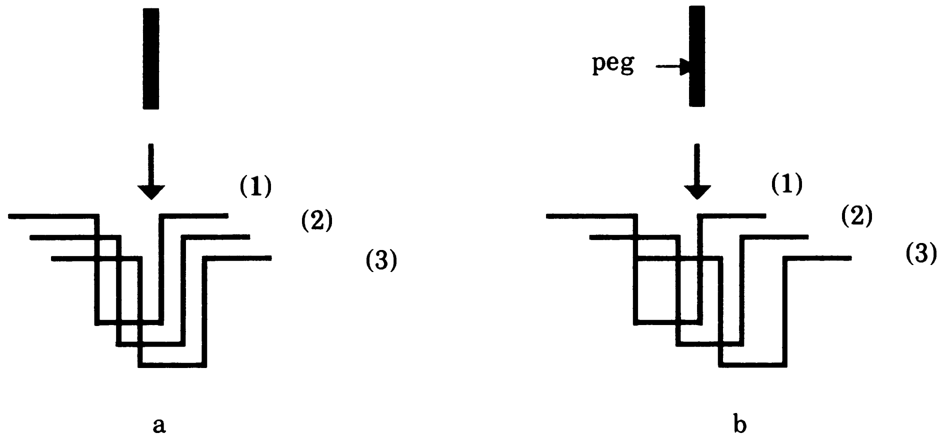

Fig.1. Motion planning with uncertainty; an example

three possible positions for the hole (1), (2) and (3), with the peg above the hole. We are required to plan a motion for the peg that would insert it in the hole. We see that a straight line motion downward will successfully insert the peg regardless of whether the hole is at position (1), (2) or (3). Therefore the uncertainty in the position of the hole is insignificant and this is an example of gross motion planning. Fig.1b on the other hand is not so straightforward. There is no *a priori* motion that will perform the insertion, and any attempt to insert the peg in the hole will require fine adjustments that depend on the exact location of the hole. Hence we conclude that the uncertainty in the position of the hole is significant here and that this is an example of fine motion planning. A slightly

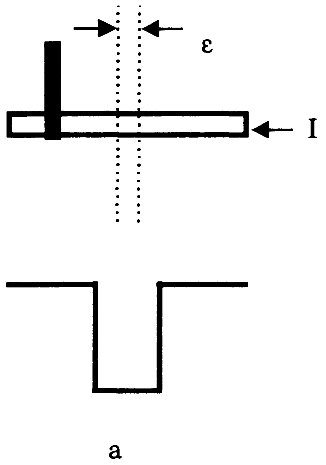

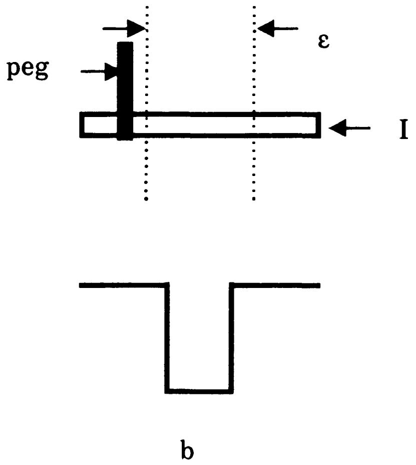

Fig.2. Motion planning with uncertainty.

different version of the same problem is shown in Fig.2. Here the position of the hole is fixed, but the initial position of the peg can be anywhere in the region marked I. The difficulty arises when the position of the peg is known only within some error ε. If ε is smaller than the clearance of the peg in the hole, Fig.2a, the task is possible with position feedback alone, since we could measure the position of the peg and utilize this information to plan a motion that will successfully insert the peg in the hole. Else, Fig.2b, an attempt to insert the peg in the hole will require additional information on the relative position of the peg with respect to the hole. This information could be obtained by any sensory mode – force sensing, generalized damping, vision etc. However, Fig.2a and Fig.2b are both examples of fine motion planning as no single motion will perform either task.

1.1 The Robot Model and A Formal Statement of the Problem

Robots With Position Measurement

We now describe the robot model on which we base our discussion. The model is that of Mason and Lozanó-Perez [1,2]. In the following, we shall use the term *effector* to refer to a fixed reference point on the object gripped by the robot. The location of any point on the gripped object with respect to this point is known precisely at all times.

The effector has up to six degrees of freedom - three translational and three rotational. The location of the effector in some fixed global frame is known at any time to within a fixed error bound

ε. In particular, the actual position of the effector is always within a sphere of radius ε centered at the observed position and vice versa. See Fig. 3a. Similarly, the orientation of the effector is known within a fixed error bound ε_θ at all times and the actual orientation of the effector is within a sphere of radius ε_θ centered at the observed orientation. We will also allow for velocity errors in the robot –

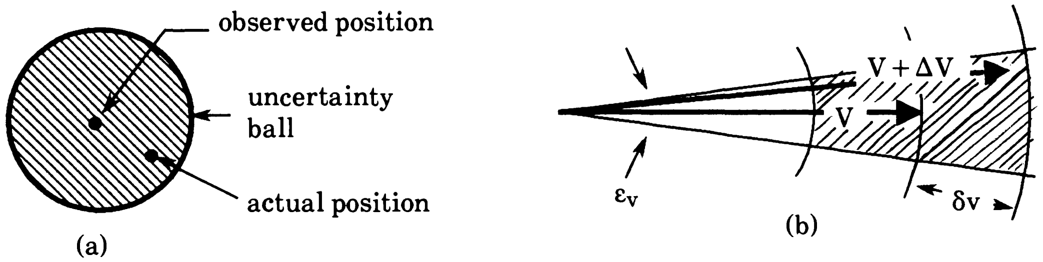

Fig.3. (a) The uncertainty ball in position measurement
(b) The uncertainty in velocity

differences between the commanded velocity and the veloctity actually attained by the effector. (We use the following convention for vectors: vector quantities are in upper case while the corresponding scalars are in lower case). The attained velocity $(V + \Delta V)$ is within an angle ε_v in direction and within δv in magnitude of the commanded velocity vector V. The assumption is illustrated in Fig.3b, where the shaded area is the set of possible locations for the arrowhead of the attained velocity vector. The magnitude of the commanded velocity, v, can be any real positive value. The error magnitude δv is assumed to vary linearly with v, while ε_v is assumed invariant. A motion plan M for the above effector consists of a sequence of moves $m_1, m_2, \ldots,$. Each move m_i is structured thus:

while $F_i(p)$ **do**

  move with velocity $f_i(p)$

  for time interval τ;

od

Here $F_i(p)$ is a boolean function of the observed effector position p and any program variables. The function f_i is the velocity function associatated with the i^{th} move and maps observed effector positions to command velocities. Every application of $f_i(p)$ to the effector is terminated after time interval τ, the *time-out* period, which can take any positive real value. Associated with the time-out period τ is an error $\delta\tau$, which varies linearly with τ. The computations involved in a motion plan are carried out to some fixed precision that is independent of the aforementioned error parameters.

The fine motion planning problem for robots with position measurement is as follows:

Input: A three dimensional *scene* S consisting of a finite set of planar walls in a closed and bounded three dimensional region D, a rigid polyhedral object O, a set of *initial positions* for the object I and a set of *goal positions* G.

Property: Is there a motion plan M for a given robot with a fixed error ε in position measurement that moves O from every position in I to some position in G without contacting any of the walls in S or leaving D? (We refer to such a plan M as a *fine* motion plan from I to G.)
We illustrate the problem with the following example in two dimensions.

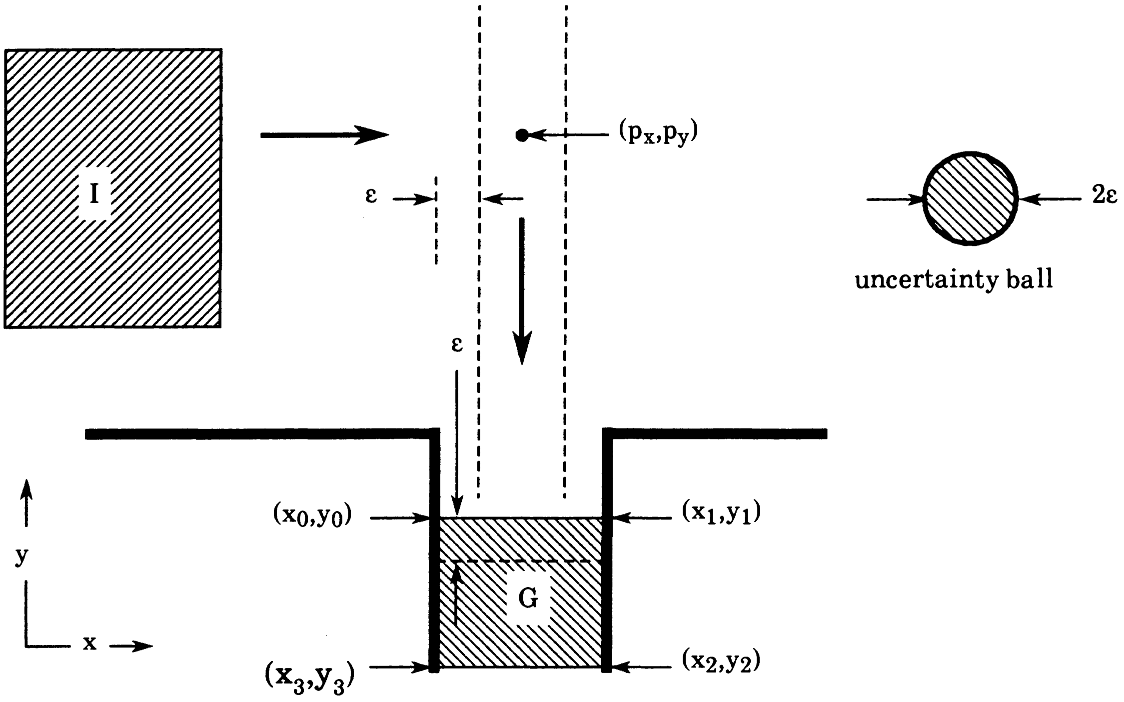

Fig.4. Fine motion planning with position measurement; an example.

Example 1: The problem as shown in Fig. 4 requires finding a motion plan to move a point object from the set of initial positions I to the set of goal positions G. The uncertainty ball associated with positon measurement is shown in the top right corner of Fig.4. We notice that a point directly above G can be moved downward into G. What constraints should the observed position (p_x,p_y) of the effector satisfy if its actual position is to be directly above G? Points with $p_x \geq x_0+\varepsilon$ and $p_x \leq x_1-\varepsilon$ are acceptable as they always correspond to actual positions that are directly above G. Our plan now is to move the effector to the right until $x_0+\varepsilon \leq p_x \leq x_1-\varepsilon$ is satisfied. Then, a downward motion until $y_3+\varepsilon \leq p_y \leq y_1-\varepsilon$ will put the effector in the goal. In particular, we give a two-move fine motion plan M as follows:

(p_x,p_y) := observed effector position;
pick τ such that $(\tau+\Delta\tau) \leq (1/v)\min(x_1-x_0-2\varepsilon, y_0-y_2-2\varepsilon)$;
while $p_x \leq x_0-\varepsilon$ **do**
    move in the $+x$ direction at speed v for time τ;

od

while $p_y \geq y_0 - \varepsilon$ **do**

move in the -y direction at speed v for time τ;

od

Here τ is picked to be small enough that overshoot of the goal and subgoal regions does not occur.

□

Robots with Damping and Position Measurement

Next we extend our robot model to include *damping*. Damping is a limited ability to conform to obstacles encountered during a motion. The concept is best explained with an example. Fig. 5a shows the path traversed by the effector of a robot with damping. The commanded velocity is V at position

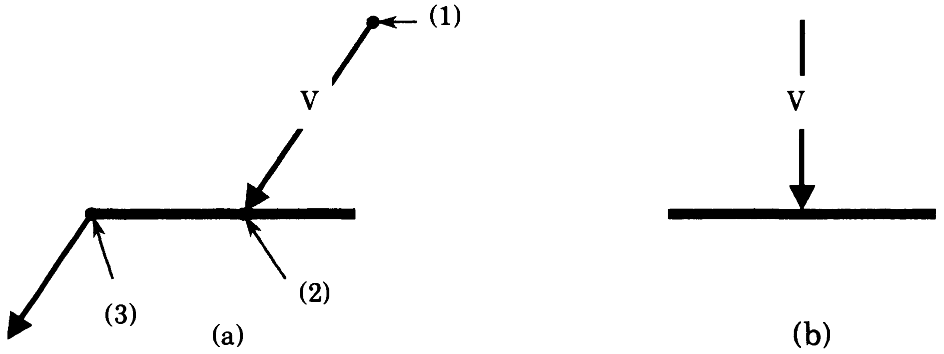

Fig. 5. Damped motion.

(1) on the path and is equal to the attained velocity. When the effector strikes the wall at position (2) it conforms to the wall by attaining a velocity equal to the component of the commanded velocity parallel to the wall. The effector slides along the wall until it reaches position (3) at which point it is no longer constrained by the wall and hence reattains the commanded velocity. In Fig.5b, the commanded velocity is perpendicular to the wall and hence the effector sticks to the wall as the component of the commanded velocity parallel to the wall is zero. In general, the effector will stick or slide depending on whether or not the commanded velocity is within the friction cone of the wall. The friction cone is a range of velocities about the normal to the wall that is determined by the coefficient of friction between the wall and the effector. For our purposes we need only know that incidence within the friction cone leads to sticking while outside the friction cone sliding occurs. See Fig.6. A fuller discussion of friction cones can be found in Mason [2]. In the presence of directional velocity errors, if there are to be commanded velocities that guarantee sticking on a wall, the friction cone should be larger than the directional velocity error. We will assume this to be the case in our discussion.

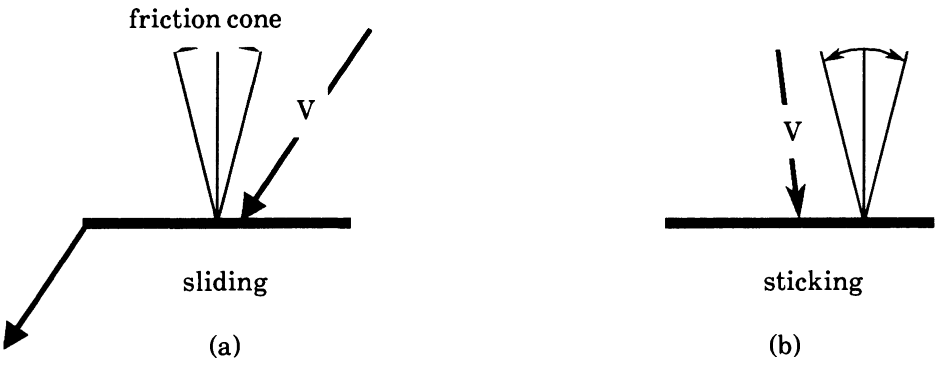

Fig. 6. Friction cones

A motion plan in this setting is almost identical to the one for robots without damping. The only difference lies in the range of values the time-out period τ can take. It is essential that the following restriction be placed on v, δv, τ, $\delta\tau$:

$$v\delta\tau + \tau\delta v + \delta v\delta\tau \geq \varepsilon.$$

Otherwise, it may be possible to reach any location in the scene with better than ε accuracy, making the uncertainty parameter ε meaningless. This is made clear in the following example.

Example 2: Consider the scene of Fig.7. Here, the corner vertex at the intersection of the two walls

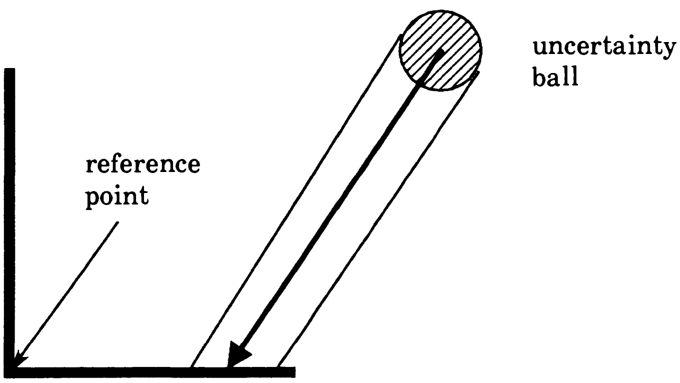

Fig. 7. Finding a reference point

can be used as a reference point. The effector moves towards the wall, slides along it and sticks at the corner (labeled reference point). Now the positon of the effector is known to the same precision as the location of the corner, which can be much better than the uncertainty ball. The effector can then be moved to any location with no more than $v\delta\tau+\tau\delta v$ error. Hence $v\delta\tau+\tau\delta v \geq \varepsilon$ should hold if the uncertainty parameter ε is to be meaningful. □

The statement of the problem in this setting is similar to the one given earlier with the exception that the robot now has the additional capability of damped motion.

Input: A three dimensional *scene* S consisting of a finite set of planar walls in a closed and bounded three dimensional region D, a rigid polyhedral object O, a set of *initial positions* for the object I and a set of *goal positions* G.

Property: Is there a motion plan M for a given robot with a fixed error ε in position measurement and damping that moves O from every position in I to some position in G without leaving D?

We illustrate the problem with the following example in two dimensions.

Example 3: The problem shown in Fig.8 is similar to the one of Fig.4 and requires moving a point

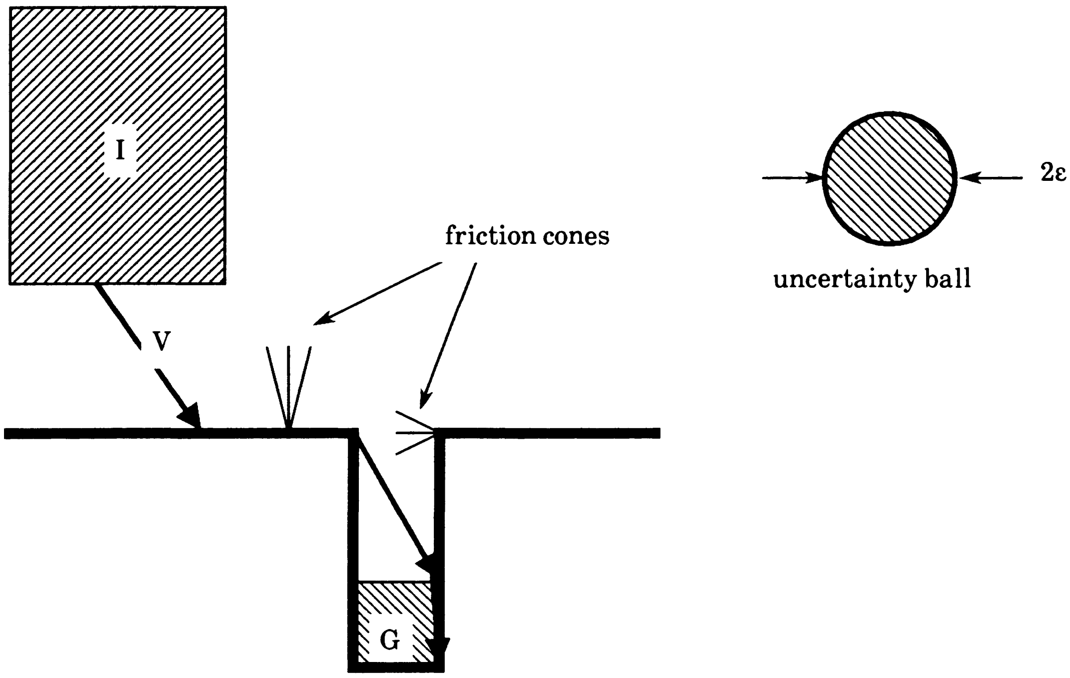

Fig.8. Fine motion planning with damping; an example.

object from the initial set I to the goal set G. Applying Theorem 1 to the scene tells us that there is no fine motion plan involving position measurement alone. In particular, there is point x in I such that the uncertainty ball around x is entirely in I while there is no such point in G. Hence there does not exist a path along which the uncertainty ball around x can be moved to a position where it is entirely in G. Assuming that the directional velocity error ε_v is zero, it is easy to verify that the velocity V shown in the figure is sufficient to move every point in I to some point in G. Notice the friction cones on the two walls and that V is outside both of them.

1.2 Overview of Results

In section 2 we exhibit a direct reduction from fine motion planning for robots with position measurement to the path planning problem for rigid objects – the"classical piano mover's" problem – which can be solved in polynomial time. In section 3 we show that fine motion planning for robots with damping is PSPACE-hard. This is the main result of the paper. Although there are algorithms in the literature [1,2] for this problem, this is the first attempt at the complexity of the problem. The result is significantly different from other PSPACE-hardness results [3,4,6] related to motion planning in that it concerns rigid objects as opposed to objects with unbounded degrees of freedom. In section 4 we show a restricted version of the problem to be PSPACE-complete. We then identify the key to the hardness of the problem and suggest a restriction that allows a polynomial time solution.

2. FINE MOTION PLANNING FOR ROBOTS WITH POSITION MEASUREMENT

2.1 Some Basic Results

We notice in our first example that a fine motion plan from I to G may exist even if I cannot be moved to G as a rigid body. In fact, a fine motion plan from I to G may be viewed as a translation of any continuous deformation of I to G. The limitation is that I cannot be deformed into an object smaller than an uncertainty ball. Before we make this more precise in the following lemma, we make some modifications to our robot model. In particular, in this section and the next, we shall consider robots that are capable of traversing arbitrary paths without any velocity error. (This is equivalent to piecewise linear motion with infinitesimal moves and can be realized by the model of Section 1.1 as a limiting case.) We also assume that the motion plan is executed on a machine capable of real arithmetic. These modifications are in the interest of clarity and in Section 2.3, we address the issue of approximate solutions for robots with piecewise linear motion plans and finite precision arithmetic.

Lemma 1: Let G be the goal set for a point object in a three dimensional scene S. Let x be any point in S such that B(x), the uncertainty ball centered at x, does not contact any of the walls. Then, there exists a fine motion plan from B(x) to G if and only if there exists a path along which B(x) can be translated as a rigid body to a position in which it is entirely in G.

Proof: Let x_G be a point in G such that $B(x_G) \subseteq G$. Let O be the rigid body defined by the boundary of B(x) and let R be the path connecting O centered at x to O centered at x_G. We now describe a fine motion plan M from B(x) to G. In particular, M moves the point object from B(x) to $B(x_G)$ by tracing out the path R. Clearly M maintains the point object within the volume swept by O along path R. Hence there can be no contact between the point object and the walls in S. Also, when M terminates, the point object is guaranteed to be in G as the uncertainty ball around x_G is completely contained in G. We conclude that M is a legal fine motion plan from B(x) to G.

Suppose there exists a fine motion plan M from B(x) to G. Let R be the path traced out by x when M is applied to x. At any point x' on R, B(x') cannot contact any wall in S. For otherwise, there exists a location of the effector contacting a wall in S with observed position x', implying that M is not a valid fine motion plan. Furthermore, let R terminate at a point x_G in G. Now, if $B(x_G)$ is not entirely in G, M is again invalid. It follows that R is a path along which B(x) can be translated to a position where it is entirely in G. □

In the light of the above lemma we have the following theorem. But first we need the following definition: A set I is *star shaped* about a point x with respect to the uncertainty ball B, if $B(x) \subseteq I$ and for any $y \in I$ and $z \in B(x)$, the line yz is in I. Fig. 9 is an example of a star shaped set in two dimensions.

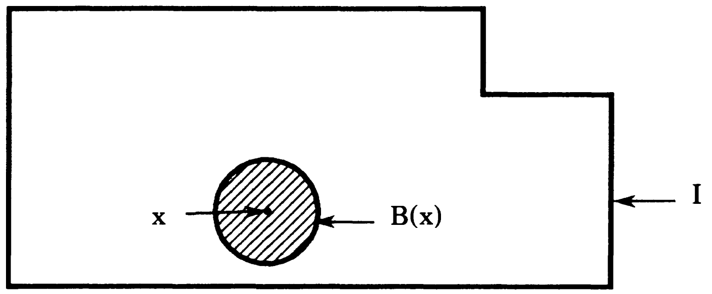

Fig. 9. A star shaped set

Proposition: If a set I is convex and contains a point a point x such that $B(x) \subseteq I$, then I is star shaped.

Proposition: A closed set I is star shaped with respect to B if there exists a point x in I such that $B(x) \subseteq I$ and for any boundary point y of I and any boundary point z of B(x), yz is in I.

Theorem 1: Let I and G be the initial and final sets for a point object in a three dimensional scene. If I is star shaped about x with respect to B, then there exists a fine motion plan from I to G iff there exists a path along which B(x) can be translated as a rigid body to a position where it is entirely in G.

Proof: Suppose there exists a fine motion plan M from I to G. Then M is also a fine motion plan from B(x) to G. It follows from Lemma 1 that there exists a path along which B(x) can be translated to a position in which it is entirely in G.

Suppose there exists a path R along which B(x) can be translated into G. Then by Lemma 1 there exists a fine motion plan M from B(x) to G. We now construct a fine motion plan $M' = m \circ M$ from I to G where m is the move given by

begin

p := observed position;

trace the line joining p and x;

end

Since $B(x) \subseteq I$, at the end of move m the effector is guaranteed to be in B(x) regardless of its inital position in I. Since I is star shaped with respect to B, the effector is always within I during m and hence cannot contact any of the walls in S. Since M is a fine motion plan from B(x) to G it follows that $m \circ M$ is a fine motion plan from I to G. □

2.2 A Polynomial Time Reduction to the Classical Piano Mover's Problem

Theorem 1 suggests a direct reduction from fine motion planning with position measurement to the "classical piano mover's problem" – path planning for rigid objects with *no* uncertainty – which is decidable in polynomial time [3,4,5]. We reproduce below the statement of the problem as presented by Reif [3].

"The classical piano mover's problem in d-space is:

Input: (R,S,p_I,p_F) where R is a set of polyhedral obstacles fixed in Euclidean d-space, and S (say, a sofa)is a rigid polyhedron with distinguished positions p_I and p_F.

Property: Can S be moved (by a sequence of translations and rotations in d-space) from position p_I to p_F without contacting any element of R?"

We now exhibit a polynomial time reduction from fine motion planning in our setting to the classical piano mover's problem. Consider the set

$$\mathrm{core}(G) = \{x \mid x \in G \text{ and } B(x) \subseteq G\}.$$

It has the property that if a point is observed to be in core(G) then it is definitely in G. Also, it follows from Theorem 1 that if $\mathrm{core}(G) = \phi$ then G is not attainable for any I such that $\mathrm{core}(I) \neq \phi$. We also observe that if a and b are two points in core(G) that are path connected in core(G) then for any point x in S, there exists a rigid translation of B(x) to a iff there exists a rigid translation of B(x) to b. We are now ready for the reduction to the piano mover's problem.

Problem : Given a set I of initial positions for a point object such that I is star shaped about a point x with respect to the uncertainty ball B and a convex polyhedral set G of goal positions in a three dimensional scene S, is there a fine motion plan from I to G ?

Reduction:

(1) Pick x_G in core(G). If no such exists, answer NO.

(2) Answer YES iff there exists a rigid translation of B(x) to $B(x_G)$.

Proof: Immediate from Theorem 1. □

Now, Step(1) is computable in polynomial time as G is convex and polyhedral. Since the piano mover's problem is decidable in polynomial time [3,4,5], we conclude that Step(2) and hence the fine motion planning problem as stated above is decidable in polynomial time.

2.3 An Approximate Solution

Next we consider the problem of fine motion planning for robots with non-zero velocity error, piecewise linear motion plans and fixed computational precision. Here, a tight reduction to rigid body motion planning is impossible as the set of possible locations for the object grows as it is moved around. Consequently, we present an approximate solution in the form of weaker versions of Lemma 1 and Theorem 1. First, we need some definitions. The *discretization parameter* δl is a measure of the distance moved by the effector in any iteration of a move. It absorbs the error in the velocity, the time out period as well as the piecewise linear approximation. The *computational error parameter* α allows for fixed precision execution of the motion plan and should be chosen accordingly. Extending our definition of the uncertainty ball B, we define B_δ to be the uncertainty ball B expanded by $\delta l + \alpha$.

Lemma 1′: Let G be the goal set for a point object in a three dimensional scene S. Let x be any point in S such that B(x), the uncertainty ball centered at x, does not contact any of the walls. Then,

(1) there exists a fine motion plan from B(x) to G if there exists a path along which $B_\delta(x)$ can be translated as a rigid body to a position entirely in G.

(2) there does not exist a fine motion plan from B(x) to G if B(x) cannot be translated as a rigid body to a position where it is entirely in G.

Proof: Let x_G be a point in G such that $B_\delta(x_G) \subseteq G$. Let O be the rigid body defined by the boundary of $B_\delta(x)$ and let R be the path that translates O from x to x_G. If R is not piecewise linear, break it up into segments at points $x_1, x_2, x_3, \ldots$, such that the line joining x_{i-1} and x_i deviates no more than $\delta l_i < \delta l$ from R. We now describe a fine motion plan $M = m_1 m_2 \ldots$ from B(x) to G. In particular, the i^{th} move m_i moves the point object from $B(x_{i-1})$ to $B(x_i)$ and is structured thus:

While observed position $p \neq x_i$ **do**

move along the line joining p and x_i towards x_i

at speed v for time τ:

od

The speed v and the time interval τ are to be chosen for each move to satisfy:

$(v\delta\tau + \tau\delta v + \delta v\delta\tau) + \delta l_i \leq \delta l.$

Clearly M maintains the point object within the volume swept by O along path R. Hence there can be no contact between the point object and the walls in S. Also, when M terminates, the point object is guaranteed to be in G as the uncertainty ball around x_G is completely contained in G. We conclude that M is a legal fine motion plan from B(x) to G.

The second claim in the lemma follows from Lemma 1. □

Theorem 1′: Let I and G be the initial and final sets for a point object in a three dimensional scene. If I satisfies both of the following conditions:

(1) I is star shaped about a point x with respect to B

(2) I grown by $\alpha+\delta l$ is a valid set of configurations for the point object

then

(1) there exists a fine motion plan from I to G if there exists a path along which $B_\delta(x)$ can be translated to a configuration in which it is entirely in G.

(2) there does not exist a fine motion plan from I to G if there does not exist a path along which B(x) can be translated to a configuration in which it is entirely in G.

Proof: An application of Lemma 1′ in much the same way as in Theorem 1. □

2.4 Objects Other than Point Objects

Finally we tackle the case of a three dimensional object O in a three dimensional scene S. Rather than work with the solid object in three dimensions, we transform the problem to moving a point object in six-dimensional configuration space as follows. For any three dimensional scene S we can construct a configuration scene S′ in six dimensional space such that a point $x=(x_1,x_2,x_3,\theta_1,\theta_2,\theta_3)$ in S′ denotes O at coordinates (x_1,x_2,x_3) and in orientation $(\theta_1,\theta_2,\theta_3)$. Obstacles are defined in S′ by the following condition: $x=(x_1,x_2,x_3,\theta_1,\theta_2,\theta_3)$ is an obstacle point if and only if O in the corresponding configuration intersects a wall in S. S′ is a subset of the product space $\mathbb{R}^3\times\mathbb{C}^3$, where $\mathbb{C}=\mathbb{R}$ mod2π. The initial and goal sets I and G are also subsets of $\mathbb{R}^3\times\mathbb{C}^3$. Since $\mathbb{R}^3\times\mathbb{C}^3$ is periodic and non-euclidean, we need to define a straight line and a star shaped set afresh. If $x=(x_1,x_2,x_3,\theta_1,\theta_2,\theta_3)$ and $y=(y_1,y_2,y_3,\phi_1,\phi_2,\phi_3)$ then a straight line xy joining x and y is the following parametric in t, $t\in[0,1]$:

$$x + t(y_1-x_1, y_2-x_2, y_3-x_3, \phi_1-\theta_1+n_1 2\pi, \phi_2-\theta_2+n_2 2\pi, \phi_3-\theta_3+n_3 2\pi)$$

for arbitrary integers n_1,n_2 and n_3. We say xy is generated by n_1,n_2 and n_3. A set I in $\mathbb{R}^3\times\mathbb{C}^3$ is star shaped about x with respect to B if both the following conditions hold:

(1) $B(x)\subseteq I$

(2) let u be any point in $\mathbb{R}^3\times\mathbb{C}^3$ and let $D=B(u)\cap I$. Then there exists integers n_1,n_2 and n_3 such that for every $y\in D$ and $z\in B(x)$, the line zy generated by n_1,n_2 and n_3 is in I.

With these definitions, it is not hard to verify that Lemma 1 and Theorem 1 stand in their original form. Although the above definition is not friendly, it says little more than its Euclidean counterpart that we saw earlier.

2.5 Remarks

Thus far we have only discussed cases where the initial set I completely contains an uncertainty ball. This can be a strong restriction. What can be said about the problem when this is not the case? We have yet another version Theorem 1 for some restricted shapes of the initial set I. In particular, if

I is *compressible* by the uncertainty ball B as defined in the following: A convex set $I \subseteq \mathbb{R}^n$ is *compressible* (to $B(x^*) \cap I$) by a convex set $B \subseteq \mathbb{R}^n$ if

$\exists x^* \in I : \forall x \in \mathbb{R}^n : B(x) \cap I$ can be *translated* to a position

where it is entirely in $B(x^*) \cap I$.

Fig. 10 illustrates the property. It shows two rectangles I and B and their intersections at two

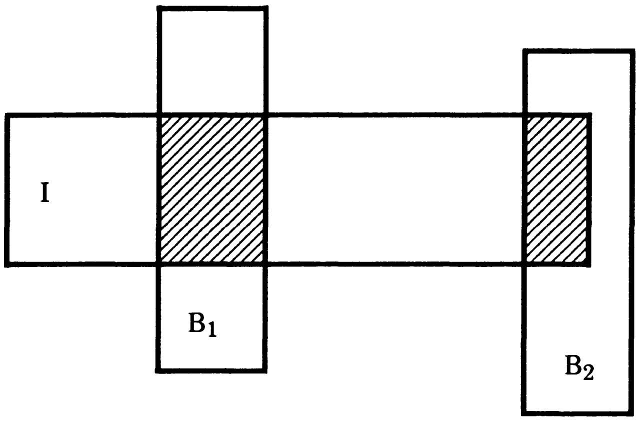

Fig 10. Compressibility. I is compressible by B.

different positions B_1 and B_2 of B. Verify that $B_2 \cap I$ can be fitted into $B_1 \cap I$ and so can any other intersection of I and B. Exactly what pairs of classes of objects satisfy the compressibility property is not known. However, the following merit mention:

(1) If B is spherical then it appears that only spherical regions I satisfy the property.

(2) If B is a cuboid then cuboidal regions I that are axially aligned with B satisfy the property.

(3) Line segments are compressible by all B.

Theorem 1″: Let I and G be the initial and final sets for a point object in a three dimensional scene. If I is convex and compressible by the uncertainty ball B to $B(x^*) \cap I$ for some x^* in I, then there exists a fine motion plan from I to G iff there exists a path along which $B(x^*) \cap I$ can be translated as a rigid body to a position where it is entirely in G. □

If the velocity error is non-zero, any region of possible positions that is smaller than an uncertainty ball will grow as it is moved around. This means that a straightforward reduction to rigid body motion planning is impossible. Hence , there is no straightforward extension of Theorem 1′ in this setting.

3. FINE MOTION PLANNING WITH DAMPING AND POSITION MEASUREMENT

3.1 A Complexity Result

In example 3, we saw that robots with damping can have feasible fine motion plans in situations where robots with just position measurement cannot. Unfortunately, attendant to this increased effectiveness is the increased complexity of deciding the existence of a motion plan.

Theorem 2: The fine motion planning problem for point objects in three dimensional scenes and robots with damping and position measurement is PSPACE-hard.

Proof: Given a Turing machine T with a binary tape alphabet that operates in a polynomial space bound S(n) and a binary string W, we construct a scene for which a fine motion plan exists if and only if T accepts W. The proof borrows from Reif [3] and Joseph et al. [6].

We assume that T has a set of states Q with a starting state q_0 and one accepting state q_f that is distinct from q_0 and has no transitions out of it. Also, T accepts by printing zeros on all the tape squares and entering q_f. The building blocks of our scene are conduits of $3S(|W|) + |Q|$ *channels* – three channels for each tape square of T and $|Q|$ channels to represent the state of T's finite control. See Fig.11. We say a channel is *active* at any point in a motion plan if executing the remainder of the

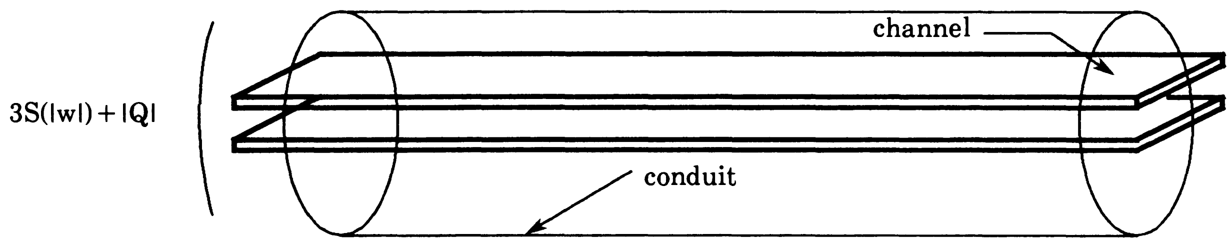

Fig.11. A conduit of channels

plan translates every point in the channel to the goal. We use the activity of the channels to encode T's configuration, which consists of its tape-contents, the state of its finite control and its head position. The three channels for each tape square are used thus: The first channel is active if the tape square contains a 1, the second channel is active if it contains a 0. The third channel is active if the tape square is currently scanned by T's head. Of the $|Q|$ state channels, the i^{th} one is active if T's finite control is in state q_i.

The overall layout of the scene is shown in Fig.12. There are two corridors linked by a series of gates G_1, G_2... One of the corridors is linked to a terminal area T_G containing the goal region G. T_G encodes the initial configuration of T by simply blanking out all the channels with walls and designating the walls in the channels that are to be active as the surfaces that make up the goal region G. See Fig.13. The other corridor is connected to a terminal area T_I containing the initial

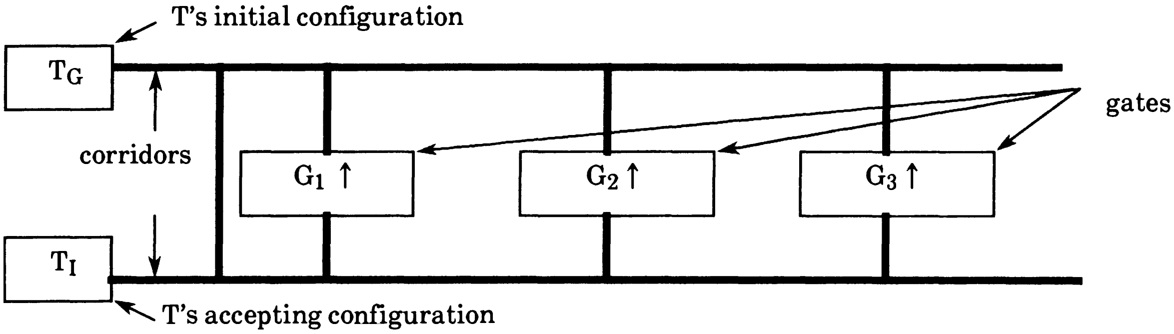

Fig.12. Schematic of the scene (heavy lines represent conduits).

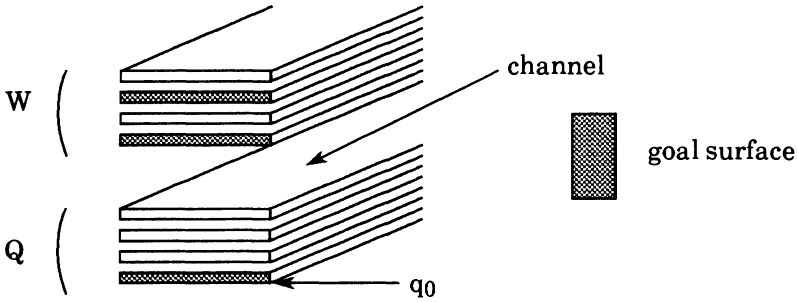

Fig.13. The terminal area TG

region I. T_I encodes the final configuration of T by blanking out all the channels and designating the blanking walls on those channels that are to be active as the surfaces that make up the initial region I.

A transition of the Turing machine T is a function from the set of configurations to the set of configurations and represents a move of T – reading a tape square, change of state of the finite control, writing on the tape square and moving the tape head. Each of the transition gates G_1, G_2... represent a legal transition of T running on a S(n) tape bound. If $S(|W|) = m$, then there are at most $2m|Q|$ transitions requiring at most $2m|Q|$ gates. G_i checks to see if the incoming set of active channels represents a aconfiguration valid for the i^{th} transition. If so, it sets the outgoing channels to reflect the transition. We now describe how these gates are to be built. To implement the transitions of T using these gates we need only be able to perform logical AND's on the activity of the channels. First we set ε, the error in position measurement, to be larger than any dimension in the scene. This ensures that position measurement is insufficient to infer to which channel the observed position corresponds. To compute $C = A \wedge B$, we connect A and B to a diffuser as shown in Fig.14. C is connected to a nozzle facing the diffuser. The dimensions of the diffuser are chosen with respect to the

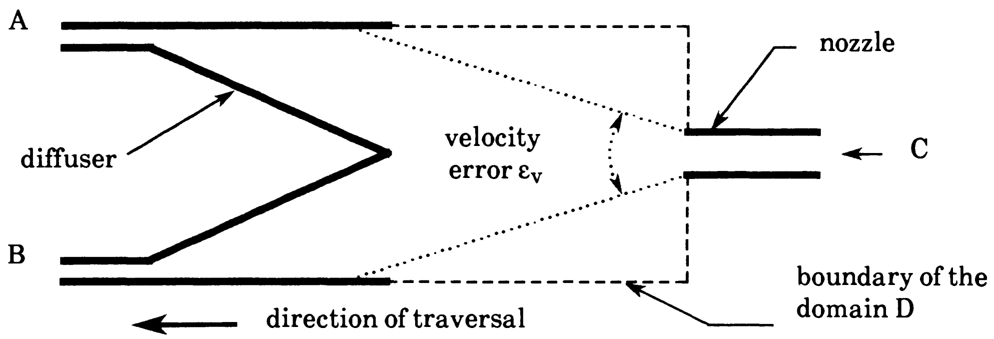

Fig.14. An AND gate

velocity error ε_v (assumed to be non-zero) so that a point at the mouth of the nozzle moved towards the diffuser cannot enter A or B preferentially. As paths outside D are not permitted, C can be translated to $A \wedge B$ or not at all, i.e., C can be made active iff both A and B are. Also, the dimension of the nozzle is chosen with respect to ε_v so that the gate canot be entered at A or B and exited through either of the other two ports. This makes the AND gate unidirectional so that $A \wedge B \Rightarrow C$ without side effects like $C \Rightarrow A$, $C \Rightarrow B$ etc.

If a transition gate G_i is to check to see if T is in state q and whether the head is at square k and reads 1, it performs AND's on the corresponding channels. To set the entry channels to reflect the transition, G_i sets the j^{th} entry channel to be the AND of the j^{th} exit channel and the result of the above validity check. Of course, it resets some channels to reflect the rewriting of tape square k and head movement by T. Fig.15 shows the schematic of a gate that implements the transition (q_1, 1 in tape square 1) to (q_2, print 0 in tape square 1, move to tape square 2). We note that a transition gate is unidirectional as the AND gates are. Hence the gates cannot be traversed the wrong way to realize illegal transitions.

Recall that the error in position measurement ε is set to be larger than any dimension in the scene. Hence position measurement cannot be used to selectively enter a transition gate. To make selective entry of gates posssible, the gates are arranged in a cascade. At the i^{th} level in the cascade, a motion plan has to choose between entering the i^{th} transition gate and continuing on to gates $i+1$, $i+2, \ldots$, as shown in Fig.16. Finally, we set the friction coefficient of the walls to be large enough that the friction cone angle is bigger than the directional velocity error ε_v.

Let S_I and S_G be two activity assignments to the channels encoding two valid configurations C_I and C_G of T.

Claim 1: There exists a fine motion plan translating S_I to S_G iff T started in C_G attains C_I.

Proof: In the following, we refer to a motion plan by the sequence of gates it traverses. The length of a plan is the length of the sequence. We proceed by induction on the plan-length l.

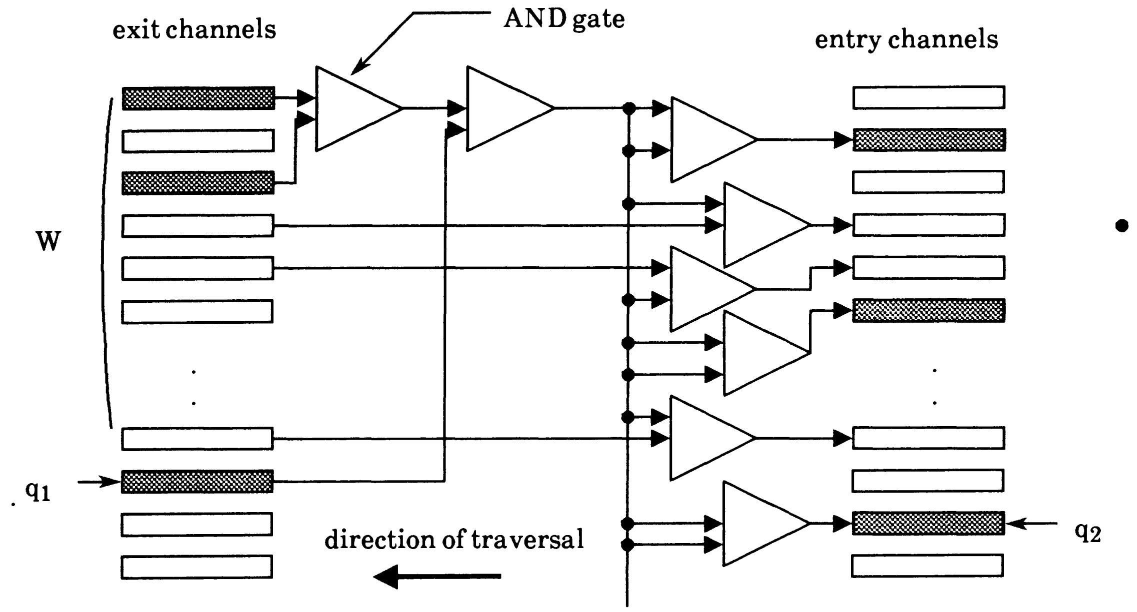

Fig.15. Schematic of a typical transition gate

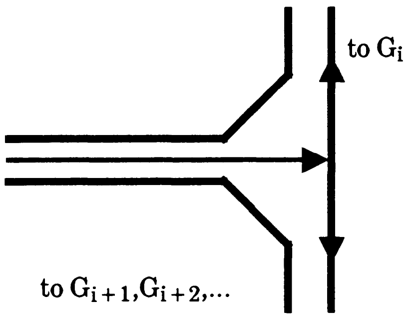

Fig.16. Cascading the gates.

Basis $l=0$; immediate.

Induction Assume true for $l=k$. Let $P = G_{P_1} \circ P'$ be a plan of length $k+1$, where P' is a plan of length k. Let $S_{I'}$ be the result of applying the plan G_{P_1} to S_I. By the correctness of the gates, S_I can be translated to $S_{I'}$ iff T started in $C_{I'}$ attains C_I. By the inductive hypothesis, $S_{I'}$ can be traslated to S_G by P' iff T started in C_G attains $C_{I'}$. It follows that S_I can be translated to S_G by a path of length $k+1$ iff T started in C_G attains C_I. □

From the above claim it follows that there exists a fine motion plan from I to G iff T accepts W. □

3.2 Remarks

Remark 1

In the proof of Theorem 2, a non-zero directional velocity error ε_v was key in the construction of the AND gates. Even if ε_v is zero, AND gates can be constructed provided the range of velocity directions available to the effector is discrete. In particular, we exhibit an AND gate for a cartesian robot - one that can command velocities only along three fixed and mutually orthogonal axes. See Fig. 17. The construction is similar to that of Fig.14. As shown, the nozzle is incident on a guide

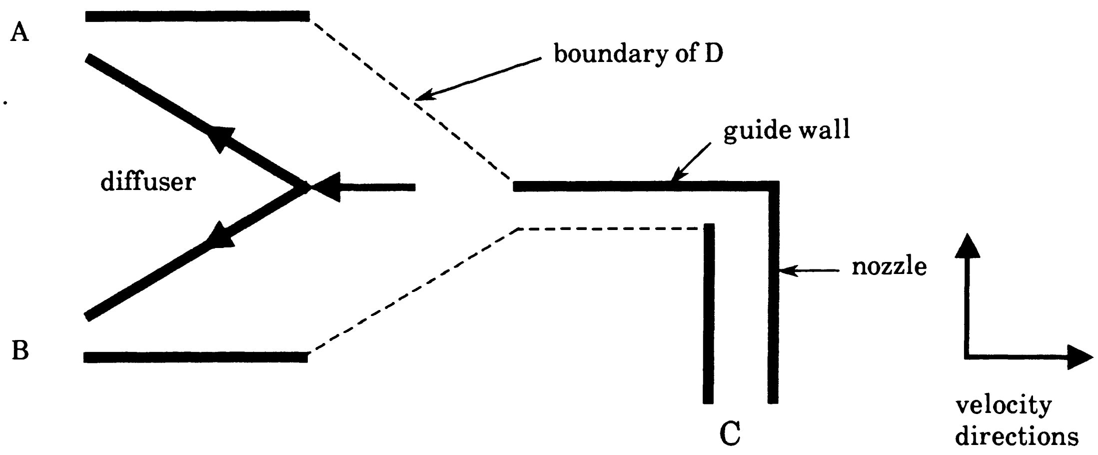

Fig.17. An AND gate for a cartesian robot.

wall. All the walls are frictionless so that the only way to exit the nozzle is to slide along the guide wall which is positioned directly opposite the apex of the diffuser. Assuming that the effector striking the apex does not stick but enters either A or B, we see that C can be active iff both A and B are, obtaining the desired relationship. It is easy to see that this version of the AND gate is unidirectional as well.

Since our proof of Theorem 2 only required motions in orthogonal directions, it holds for cartesian robots with zero directional velocity error as well, using the above AND gates.

Remark 2

Our construction of Theorem 2 utilized an initial set I that was made up of a number of disconnected components. This is not an essential ingredient of the proof. The proof goes through even for a single-point initial set, utilizing cascaded AND gates to generate the final configuration of the Turing machine T. Fig. 18 illustrates the idea.

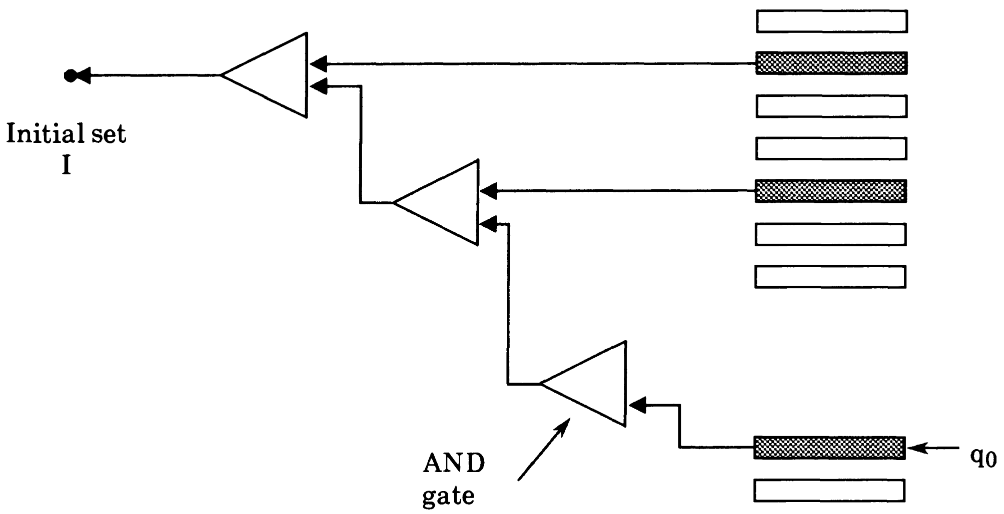

Fig.18. Generating configurations
from single-point initial sets.

3.3 More Complexity Results

Our next result is mainly of technical interest. Having shown the existence problem in Theorem 2 to be PSPACE-hard, we would like to show at least a restricted version of the problem to be PSPACE-complete. If we can eliminate much of the haze and show a pared down version of the problem to be PSPACE-complete, we can then inquire into the reason for the jump in complexity between the problems for robots with and without damping. We shall proceed with an algorithm for fine motion planning in rectilinear scenes. First, some definitions. A three dimensional scene is *rectilinear* if

(1) All the walls in the scene are rectangular planes of zero thickness.

(2) The scene has a fixed orthogonal reference frame and each wall has at least one of its edges parallel a reference axis.

(3) No two walls intersect, except on their boundaries.

It is easy to verify that our construction in Theorem 2 is a rectilinear scene. A scene S is *fine* with respect to a robot if there exists a point x in S such that $S \subseteq B(x)$, the uncertainty ball around x. In other words, S can be entirely contained by an uncertainty ball. Because of our lower bound condition on the time-out period τ for damped robots, a move in a fine scene S that commences at a point in S will terminate in S if and only if the effector sticks on a wall in S. Hence goal regions in fine scenes have to be wall surfaces to be attainable.

The 'Pacman' Problem

Input: Given are a damped cartesian robot with zero directional velocity error, a fine rectilinear scene in three dimensions with reference frame aligned with the robot axes, a single-point initial set I and a goal set G for a point object.

Property: Is there a fine motion plan with damping and position measurement from I to G in S?

Algorithm 1:

The following is a non-deterministic algorithm for the above problem.

start:

P := I; {set of possible positions of the point object}

If ($P \subseteq G$) **then** halt and report success;

d := pick a direction of the possible six of the cartesian robot;

if d was picked twice in a row **then** halt and report failure;

simulate move in direction d;

P' := new set of possible locations for the point object:

if $P' \not\subseteq S$ **then** halt and report failure;

else P := P'; goto *start;*

Claim 2: Algorithm 1 can be implemented in deterministic polynomial space.

Proof: In order to prove the above claim, we prove two subclaims:

(1) P can be never have more than polynomially (with respect to the number of vertices in the scene S) many points in it. Furthermore, these points need only be represented to the same precision as the vertices in S.

(2) Computing P' from P after each move takes polynomial space.

Proof of subclaim (1):

Let X, Y and Z be the reference axes of the rectilinear scene and let the input be specified as coordinates in this frame. Consider the set

$$S_X = \{x \mid \exists \text{ an edge in S parallel to Y or Z with X-coordinate } = x\} \cup \{x \mid (x,y,z) \in I\}$$
$$\cup \{x \mid \exists \text{ an edge along which two planes meet with X-coordinate } = x\}$$

and similarly S_Y and S_Z. Clearly there at most $O(n^2)$ points in each of the sets, where n is the number of vertices in the scene S. Whence it follows that

$$|S_X \times S_Y \times S_Z| = O(n^6).$$

We now show that $S_{XYZ} = S_X \times S_Y \times S_Z$ is closed under cartesian moves in the sense that applying a cartesian move to an element of S_{XYZ} will translate it to a set of points in S_{XYZ} or move it outside the scene S. Let (x_1,y_1,z_1) be a point (see Appendix A in this regard) in S_{XYZ} and let (x_2,y_2,z_2) be reached from (x_1,y_1,z_1) in a single cartesian move along direction d. Without loss of generality assume d is in the $+X$ direction. Then, (x_2,y_2,z_2) must lie on a wall normal to the x axis or on an edge along which two planes meet. In either case $x_2 \in S_X$. Let y be the coordinate of the last edge parallel to the Z direction that was traversed (made contact with) in the move.

Then $y_2 = y$ clearly. If no such edge was traversed, $y = y_1$. In either case $y_2 \in S_Y$. Similarly z_2 is in S_Z. Whence it follows that $(x_2,y_2,z_2) \in S_{XYZ}$. The argument is illustrated in Fig.20, where the point is 'split' and the Y-coordinate of both of the resultant points is left unaltered but the Z-coordinate of one is changed by deflection against the oblique wall in the path.

Proof of subclaim (2):

We must keep in mind here that a point p in P translated in direction d could 'split' by striking an edge as shown in Fig.19. At each split, we can identify one of the resultant points with p, so that we

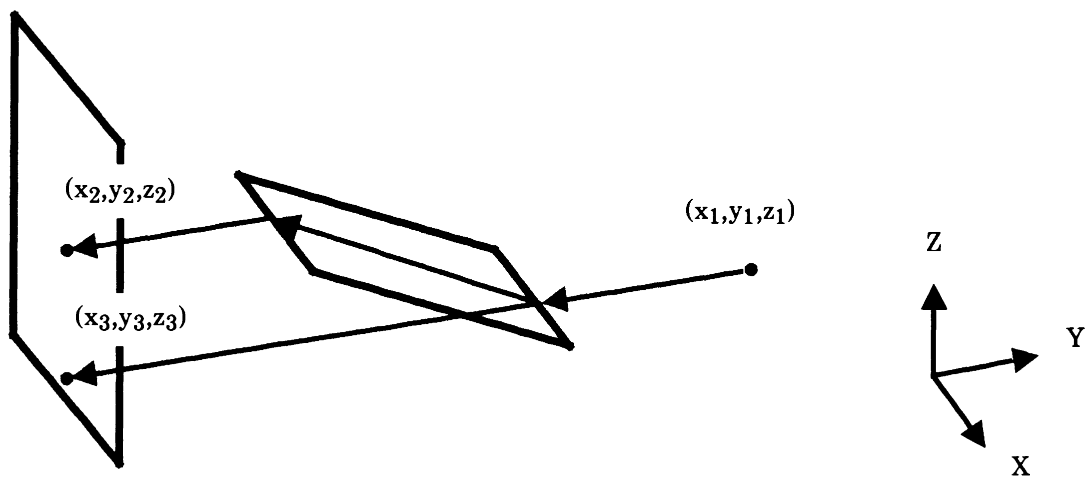

Fig. 19. Computing moves in rectilinear scenes.

can say that each $p \in P$ is split at most O(n) times - once for each wall in S - before it sticks and thereby becomes an element of P' or leaves the scene S. Consequently, we can recursively compute P' from P with the depth of the recursion being O(n) as follows.

for each $p \in P$ **do**

  Move(p,d);

od;

procedure Move(p:point;d:direction);

  {*simulate moving a point in direction d*}

(*) simulate moving p in direction d **until** it either

  (a)sticks

  (b)hits an edge

(c)goes outside S;

if (a) **then** P' := P'∪{p};

if (b) **then** let p split into p_1 and p_2;

Move(p_1,d);

Move(p_2,d);

if (c) **then** abort and report failure;

end;{Move}

Clearly computing P' from P as shown above takes polynomial space. In fact, the line marked (*) above takes polynomial time in a straightforward implementation.

From the two subclaims it follows that Algorithm 1 runs in polynomial space and hence can be implemented in non-deterministic polynomial space. Since NPSPACE = PSPACE by Savitch's thorem [7], it follows that Algorithm 1 can be implemented in deterministic polynomial space. □

On the basis of Claim 2 and Theorem 2 we can now state

Theorem 3: The 'Pacman' problem is PSPACE-complete.

Proof: By Theorem 2, Remarks 1 and 2 of Theorem 2, and Claim 2. □

We now have a very restricted version of the existence question for fine-motion plans that is PSPACE-complete. This will permit us to inquire into the reasons that contribute to the complexity of the problem. The key to the complexity is the number of 'degrees of freedom' of the object - the number of possible locations for the object at any point in a motion plan. If this were to be bounded by some constant, the problem can be solved in polynomial time. To illustrate this we look at a modified version of the problem of Algorithm 1 and exhibit a polynomial time algorithm for it.

<u>The Restricted "Pacman" Problem</u>

Input: Given are a damped cartesian robot with zero directional velocity error, a fine rectilinear scene in three dimensions with reference frame aligned with the robot axes, a single point initial set I and a goal set G for a point object.

Property: Is there a fine motion plan with damping and position measurement from I to G such that at any point in the plan the object has at most k (for fixed k) possible locations?

Algorithm 2:

Define a configuration to be the set of possible locations of the object at any instant. There are at most $\Pi^{k}_{i=1}\,{}^{(n^6)}C_i = O(n^{6k})$ configurations by our arguments in Claim 2. We now construct a graph with one vertex per configuration. We place a directed edge between two vertices u and v if one of the six directional moves of the robot translates u to v. For each vertex u we can find the outgoing edges using the following procedure.

```
for each direction d do
    P: = u;{P is the set of possible locations }
    v := Φ;
    while P ≠ Φ do
        simultaneously simulate moving all points in P in direction d until some point
        p in P does one of
                (a)sticks
                (b)hits an edge
                (c)goes outside S;
        if (a) then v := v∪{p};
        if (b) then let p split into p1 and p2;
                    P := (P-{p})∪{p1}∪{p2};
        if (c) then P := Φ; v := Φ;
    od
    draw an edge from u to v;
od
```

This procedure takes polynomial time as at worst P can run through all the configurations, since no configuration can be repeated in a simulation as the coordinates of the points in P are non-decreasing in direction d. We now label a vertex a 'goal vertex' if every point in the corresponding configuration is in the goal set G. We then search the graph for a path from the vertex representing I to any goal vertex. Clearly, there exists a motion plan from I to G that has no more than k possible locations for the object at any instant iff there is such a path.

Since the construction of the graph, the labelling and the searching of the graph are all feasible in polynomial time, Algorithm 2 can be implemented in polynomial time and space polynomial in klogn.

4. CONCLUSION

We began with fine motion planning for robots with position measurement and showed that if the initial region contains an uncertainty ball, we can reduce the problem in polynomial time to motion planning for a rigid object. Since the latter is decidable in polynomial time, the former is as well. We then looked at fine motion planning for a point object in three dimensions and robots with damping and showed a rather general class of problems to be PSPACE-hard. Discarding the unnecessary ingredients in the problem, we identified a seemingly innocuous class that is PSPACE-complete. Restricting this class to motion plans with a constant bound on the number of possible locations for the object reduced the complexity of the problem to polynomial time. From this we conclude that the number of disconnected regions in the set of possible locations for the object is the key contributor to the complexity of the problem. A practical algorithm for fine motion planning should only look for solutions that obey a bound in this regard.

5. APPENDIX A

It is strictly insufficient to refer to a point by its coordinates alone. If the point is on a plane, it is necessary to know on which side of the plane, for instance. In general, a point could lie on an edge along which many planes meet. We assign a unique number to each plane and a sign, say + and -, to refer to the two sides of a plane. Then, to locate a point we need its coordinates, the planes on whose surfaces it lies and the signs (+,-) for those surfaces. Since we need only three planes to uniquely identify a point, representing this information takes O(logn) space per point and hence requires no change in the argument.

The cardinality of the set of possible locations needs no adjustment on this account as we are considering the walls pairwise anyway.

6. ACKNOWLEDGEMENTS

Thanks are due to J.R. Gilbert, J.E. Hopcroft, D.Kozen, R.Palmer and A.Wechsler for their help and patience.

7. REFERENCES

[1] Lozano-Pérez, T., Mason, M.T., Taylor,R.T., The Int. J. of Robotics Research, Vol 3, No. 1, Spring 1984.

[2] Mason, M.T., Automatic Planning of Fine Motions: Correctness and Completeness, Carnegie-Mellon Univ., CMU-RI-TR-83-18.

[3] Reif, J.H., FOCS, 1979.

[4] Schwartz, J.T. and Sharir, M., On the 'Piano Movers' Problem II, Courant Institute of Math. Sci., Report No. 41, Feb. 1982.

[5] Lozano-Pérez, T. and Wesley, M.A., Comm. of the ACM, Oct. 1979, Vol 22, No. 10.

[6] Joseph, D.A. and Plantinga, H.W., ACM Symp. on Computational Geometry, 1985.

[7] Hopcroft, J.E. and Ullman, J.D., Intro. to Automata Theory, Languages and Computation, Addison-Wesley, 1979.