An Exact Algorithm for Kinodynamic Planning in the Plane

John Canny *   *Ashutosh Rege* *   *John Reif* †

Computer Science Division
University of California, Berkeley

Computer Science Department
Duke University, Durham, N.C.

1 Introduction

Consider the following problem : we are given an object in physical space, an initial point and a final point. We have to plan a motion for this object, through physical space, avoiding obstacles present therein. Additionally, the motion should respect certain other constraints, such as given bounds on the velocity and acceleration. Depending on the constraints involved, we can therefore define a broad class of problems. *Kinodynamic planning* deals with synthesizing robot motions subject to both kinematic constraints (such as avoiding obstacles) and dynamic constraints (such as bounds on the acceleration and velocity). A long-standing open problem in robotics has been that of devising algorithms for generating *time-optimal* motions under kinodynamic constraints. This problem has been considered previously in the literature and approximation algorithms have been provided for the two and three dimensional cases [CDRX] but with the exception of the one-dimensional case, no exact algorithms have been given. In this paper, we provide the first exact algorithm for time-optimal kinodynamic motion planning in the two-dimensional case.

2 The Problem Statement

Consider two-dimensional physical space, i.e. $\Re^2$, with polygonal obstacles. A point mass must be moved from a specified start position and velocity $S = (s, \dot{s})$ to an

end position and velocity $F = (f, \dot{f})$ avoiding the obstacles. The point mass is moved by the application of command accelerations (via command forces). Denote the velocity and acceleration of the point mass over time t by $v(t)$ and $a(t)$ respectively. The motion is subject to dynamic constraints in the form of upper bounds on the magnitude of the velocity and acceleration in some given norm. That is,

$$\| v(t) \| \leq v_{max} \quad \text{and} \quad \| a(t) \| \leq a_{max} \quad (1)$$

In this paper, we consider the L_∞ case, though we also provide partial results for arbitrary norms. The L_∞ case models Cartesian robots such as *RobotWorld*^{TM}. Let us denote by C, the physical space in which the object moves. Let $O \subseteq C$ be the *obstacle* space, i.e. the space occupied by the polygonal obstacles, and let $F \subseteq C$ be the free space. We assume that the obstacle space is specified as an arrangement of n *vertices* with rational coordinates which are joined together by *edges*. Let CP denote phase space, which is isomorphic to $\Re^4$. A point Q in CP is a pair $(q, \dot{q})$, corresponding to position and velocity. Similarly, let OP denote the phase obstacle space corresponding to forbidden positions and velocities, and FP denote phase free space. In general, a kinodynamic problem is given by a tuple $(O, a_{max}, v_{max}, S, F)$. Let $a : [0, a] \to \Re^2$ be some command acceleration where $[0, a]$ is an interval in time. Denote by $p : [0, a] \to C$, the *path* in physical space, and by $\Gamma : [0, a] \to CP$, the *trajectory* in phase space, corresponding to this acceleration. Then, a solution to the problem $(O, a_{max}, v_{max}, s, f)$ is a command acceleration $a : [0, a] \to \Re^2$ such that (1) is satisfied and $\Gamma([0, a]) \subseteq FP$, $\Gamma(0) = S$, $\Gamma(a) = F$. The *time* of solution is a. And a *time-optimal solution* to the given problem is a solution such that the time is minimized.

3 Previous results

The one-dimensional case is studied in [O]. The problem solved there is planning the motion of a particle moving on the real line such that given two "pursuit"

271

functions $f(t)$ and $g(t)$, the position function, say $x(t)$, satisfies $f(t) \leq x(t) \leq g(t)$ for all t. In addition, the motion must obey a bound on the acceleration. A polynomial time exact algorithm is provided for this problem. In the three-dimensional case, finding an exact solution is known to be NP-hard. Not much is known about lower bounds for finding an exact solution in the 2D case.

There is an extensive literature on time-optimal trajectory planning subject to different constraints. One focus has been to determine the time-optimal control for a manipulator moving along a given path [BDG, SM]. Other work attempts to characterize the time-optimal solutions analytically [Hol, Sch]. Sahar and Hollerbach [SH] and Shiller and Dubowsky [SD] provide approximation algorithms for robots with several degrees of freedom and full dynamics. These approaches use grid methods to compute solutions which take time that is approximately that taken by an exact time-optimal solution. However, these algorithms run in time exponential in the number of grid points. The first polynomial-time approximation algorithm for the two and three dimensional cases was provided in [CDRX]. More specifically, they provide an algorithm which finds a "safe" trajectory which runs in time $(1+\epsilon)$ times the time for an optimal trajectory. The trajectory is safe in the sense that there exists a "tube" of a certain size (the radius being a function of the particle velocity) around the trajectory which does not pass through through obstacle space. The algorithm runs in time which is polynomial in ϵ^{-1} and in n, the complexity of the input. Jacobs et al. [JHCP] give an algorithm for generating near time-optimal trajectories for an open-kinematic-chain manipulator with guaranteed bounds on the closeness of the approximation. Related problems include determining the shortest bounded curvature path in the presence of obstacles (without any dynamic constraints). Fortune and Wilfong [FW] address this problem in the two-dimensional case. They provide an exact algorithm which solves the reachability problem, i.e. it determines if there exists a path between two points which obeys the given constraints but it does not generate the path. Their algorithm runs in time and space exponential in the complexity of the polygonal arrangement. An approximation algorithm for generating a bounded-curvature path is provided in [JC].

4 Statement of Results and Overview of Approach

In this paper, we give the first algorithm to generate an *exact* time-optimal solution to the two dimensional problem $(O, a_{max}, v_{max}, \mathbf{S}, \mathbf{F})$ where $\mathbf{S} = (\mathbf{s}, \dot{\mathbf{s}}) = (\mathbf{s}, 0)$ and $\mathbf{F} = (\mathbf{f}, \dot{\mathbf{f}}) = (\mathbf{f}, 0)$ (i.e. the object starts and ends

at rest). The velocity and acceleration are assumed to be bounded in the L_∞ norm. We also provide partial results for arbitrary norms, in particular the L_2 norm. The algorithm requires space which is polynomial in the input (i.e. PSPACE) and runs in exponential time.

We can summarize our approach as follows :

- We first define and characterize, in the case of the L_∞ norm, a certain class of solutions which we call the *canonical* solutions, which satisfy certain homotopy properties.

- Then, we prove a tracking lemma which provides bounds on the time difference between a given trajectory and (a slower) one which tracks it. This lemma is applicable to arbitrary norms and is of independent interest.

- We then prove the Loop Elimination theorem which shows that any solution which generates a self-intersecting path can be replaced by one which does not give rise to a self-intersecting path. This theorem too applies to arbitrary norms.

- Next, we show that there exists a canonical solution which gives rise to a non-self-intersecting path and which is time-optimal.

- Finally, we show how to use the theory of the reals (with a bounded number of quantifier alternations) to obtain the canonical solution and hence the associated trajectory.

5 The canonical solutions

Consider first the one-dimensional problem : If we do not have a bound on the velocity, it can be easily seen that the fastest way to get from a point s with velocity $\dot{s}$ to another point f with velocity $\dot{f}$ is to apply a "bang-bang" command acceleration as shown in Fig. 1. Consider now the motion of the given object in two dimensional physical space without the presence of obstacles. Suppose the particle is to be moved from a point $\mathbf{S} = (\mathbf{s}, \dot{\mathbf{s}})$ to $\mathbf{F} = (\mathbf{f}, \dot{\mathbf{f}})$. Since the velocity and acceleration are bounded in the L_∞ norm, the bounds in the x and y directions are decoupled. We could then treat the two-dimensional problem as two one-dimensional problems and apply bang-bang accelerations in both directions. However, in general, the time-optimal motion in one of the directions will take more time than the other. We will refer to that direction as the *saturated direction*. Assume that the saturated direction is the x direction. Then, as shown in Fig. 2, we can apply a bang-bang acceleration in the x direction. Since the time-optimal motion in the y direction takes time at most equal to that in the x direction, it follows that

we can arrange the velocity profile in the y direction in a triple-bang fashion as shown in Fig. 2 such that the area under the velocity curve is exactly the distance to be traversed in the y direction. Since there is only one position of the middle segment of the velocity profile in the y direction which corresponds to the right y distance, the motion thus obtained is uniquely specified. Let the velocity profiles in the x and y directions have the general shape shown in Fig. 2. Let $\mathbf{s} = (s_x, s_y)$, $\dot{\mathbf{s}} = (\dot{s}_x, \dot{s}_y)$, $\mathbf{f} = (f_x, f_y)$ and $\dot{\mathbf{f}} = (\dot{f}_x, \dot{f}_y)$. For time $t \leq t_1$, the equations for the path described in the x and y directions are given by :

$$x = s_x + \dot{s}_x t + \frac{1}{2} a_{max} t^2 \qquad (2)$$

and

$$y = s_y + \dot{s}_y t + \frac{1}{2} a_{max} t^2 \qquad (3)$$

We can write similar expressions for the intervals $[t_1, t_2]$ and $[t_2, t_3]$. If we let $s_f - s_x = \Delta x \geq 0$, we get

$$t_1 = \frac{\dot{s}_x}{a_{max}} + \frac{1}{a_{max}} \sqrt{\frac{\dot{s}_x^2 + \dot{f}_x^2}{2} + a_{max} \Delta x} \qquad (4)$$

We can obtain similar expressions for t_2 and t_3. Thus, given the initial and final positions and velocities of a canonical segment we can write parametrized expressions (with the positions and velocities as indeterminates) describing the parabolic pieces which comprise the path. The parameter we use is time t. In section 8 we will show how we can use such expressions to write predicates which enable us to generate the optimal trajectory. The above analysis holds for the type of velocity profiles shown in Fig. 2. We can get different types of velocity profiles depending on the relative values of the velocities as shown in Fig. 3. For each case we can obtain the extremal time instances as above thereby giving us similar expressions which describe the corresponding path. The previous discussion assumed that there was no bound on the velocity of the particle. If we have a bound on the velocity then the motion of the particle would, in general, appear as shown in Fig. 4. We can then determine the extremal time instances in the same manner as before and obtain expressions for the path.

Now consider the motion of the object through physical space with obstacles. As the object moves from $\mathbf{S}$ to $\mathbf{F}$, the path might touch vertices or edges of the obstacle space. We will call such points on the path (or on the trajectory) *contact points* and will refer to the portion of a path (or trajectory) between two consecutive contact points as a *segment* of the path (or trajectory). For the sake of completeness we will also include the start and finish points of the trajectory as contact points. Consider a segment of a path where the command acceleration in both the x and y directions

is bang-bang and the motion is saturated in at least one direction. If the velocity profile is such that it is saturated in at least one direction and is arranged in the unsaturated direction as shown in Fig. 2, then we will call the segment a *canonical segment*. Finally, we have,

Definition 1 *Let* $\mathbf{a}$: $[0, b] \rightarrow \Re^2$ *be a solution to (i.e. the command acceleration for) a given problem* $(O, a_{max}, v_{max}, \mathbf{S}, \mathbf{F})$. *We say that this solution is* canonical *if it is made up of a finite number of canonical segments.*

Thus, equivalently a solution is canonical if it is made up of a finite number of constant saturated acceleration steps.

Note : The canonical solutions have an important property which we will need in Section 7 to prove the existence of time-optimal canonical solutions : Both the position and the velocity are continuous functions of the end-point position and velocity, i.e. if we vary the endpoint position and velocity continuously, we will obtain a continuous deformation of the trajectory we started out with. In other words, we obtain trajectories which are homotopic to the original one. This follows from the fact that t_1, t_2 and t_3 as obtained above are continuous functions of the end-point constraints. Therefore so are x, y and the velocities in the two directions. This is true even when the saturated direction changes since it does so in a continuous manner.

6 The loop elimination theorem

In this section, we show that any solution path which contains a loop, or more generally, intersects itself, can be replaced by one which does not intersect itself and which takes time which is less than or equal to that taken by the original trajectory. This enables us to restrict attention to loop-free paths in our search for an optimal trajectory using the theory of the reals. In order to prove this theorem, we prove a tracking lemma. The essence of this lemma is as follows : given any trajectory in free space, we show that there exists another trajectory starting at rest, which follows the same path in physical space, such that the time difference between the two trajectories at any given point on the path is bounded from above. This is done by establishing an invariant relating the two trajectories which holds at all points along the path. We prove this lemma for *any norm* and in particular for the L_∞ and L_2 norms. Though we use this lemma to prove the loop elimination theorem, it is of independent interest since it provides a simple and elegant characterization of the relationship between the original trajectory and the tracking trajectory.

Notation : Let $\mathbf{a}$: $[0, b] \rightarrow \Re^2$ be a solution to some given problem. Denote by $\bar{\mathbf{v}}(s)$ the arc-parametrized

velocity of the trajectory corresponding to this solution, where s denotes arc distance traversed along the path. Let $\tilde{a}(s)$ denote the arc-parametrized acceleration and $\tilde{p}(s)$ the arc-parametrized path. Let $t(s)$ be the time taken by the solution to traverse a distance of s along the path. Let $\| \cdot \|$ denote any norm.

Lemma 1 *(Tracking Lemma) Consider a command acceleration* $a_1 : [0, b_1] \to \Re^2$ *Then there exists a command acceleration* a_2 *such that*

1. $v_2(0) = 0$ *(the second trajectory starts at rest)*

2. $\vartheta(s)$ *is non-increasing, where* $\vartheta(s) = \| \Delta \tilde{v}(s) \| + \Delta t(s) a_{max}$, $\Delta \tilde{v}(s) = \tilde{v}_1(s) - \tilde{v}_2(s)$ *and* $\Delta t(s) = t_2(s) - t_1(s)$.

3. $\tilde{p}_1(s) = \tilde{p}_2(s)$ *(i.e. the paths described by the two trajectories are the same).*

Proof: We want to find a command acceleration which will satisfy conditions 1 to 3 in the statement of the lemma. Our strategy will be simple : we will follow the original path and at every point on the path we will apply the maximum possible acceleration in the tangential direction till we (possibly) encounter a point where the velocity in the second case equals that in the first case. From this point on, we will apply the same acceleration as that in the original trajectory. Consider Fig. 5 representing the accelerations and velocities at some point s on the two paths. The bounding closed curve represents the equinormal contour corresponding to a_{max} of the norm under consideration. *We note that from the properties of a norm the equinormal contour describes a convex region.* We have scaled the velocity $\tilde{v}_1(s)$ so that it has the same norm value as a_{max}. Both motions give rise to a centripetal acceleration of magnitude $\| \tilde{v}_1(s) \|_2^2 / \rho(s)$ in the original trajectory and $\| \tilde{v}_2(s) \|_2^2 / \rho(s)$ in the second case. Here $\rho(s)$ is the radius of curvature at s. Both accelerations are perpendicular to the velocity direction and point toward from the center of curvature. We show these accelerations $\tilde{a}_{c1}$ and $\tilde{a}_{c2}$ in Fig. 5. At every point, the velocity in the second case is at most that in the first case. Thus the centripetal acceleration in the tracking trajectory is at most that in the original trajectory. This means that as long as the velocity in the tracking trajectory does not exceed that in the original trajectory, we can apply the maximum possible acceleration in the tangential direction and yet follow the original path. Since the first motion is a legal one, so is the second one. In either case the amount of tangential acceleration that can be applied is limited by the centripetal acceleration. The tangential acceleration $\tilde{a}_{t2}(s)$ applied is the maximum possible in the tangential direction as shown in Fig. 5. The acceleration $\tilde{a}_2(s)$ is then obtained as a vector sum of the normal (centripetal) and tangential

accelerations. Since we are following the original path at each point and since we start with a zero velocity, conditions 1 and 3 in the lemma are satisfied.

We will now prove that condition 2 is also satisfied. We will show that for two points with arc distance s and $s + \delta s$, where δs is sufficiently small, the difference $\delta \vartheta(s) = \vartheta(s + \delta s) - \vartheta(s)$ is negative. We have

$$\delta \vartheta(s) = \delta \| \Delta \tilde{v}(s) \| + a_{max} \delta \Delta t(s)$$

The general relationship between $\tilde{v}_1(s)$, $\tilde{v}_2(s)$, $\tilde{v}_1' = \tilde{v}_1(s + \delta s)$ and $\tilde{v}_2' = \tilde{v}_2(s + \delta s)$ is as shown in Fig. 6. Since $\tilde{v}_1'$ and $\tilde{v}_2'$ are in the same direction,

$$\| \Delta \tilde{v}(s + \delta s) \| = \| \tilde{v}_1' \| - \| \tilde{v}_2' \|$$

Similarly,

$$\| \Delta \tilde{v}(s) \| = \| \tilde{v}_1 \| - \| \tilde{v}_2 \|$$

And,

$$
\begin{aligned}
\delta \Delta t(s) &= t_2(s + \delta s) - t_1(s + \delta s) - t_2(s) - t_1(s) \\
&= (t_2(s + \delta s) - t_2(s)) - (t_1(s + \delta s) - t_1(s)) \\
&= \delta t_2(s) - \delta t_1(s)
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\delta \vartheta(s) &= (\| \tilde{v}_1' \| - \| \tilde{v}_2' \|) - (\| \tilde{v}_1 \| - \| \tilde{v}_2 \|) \\
&\quad + a_{max}(\delta t_2(s) - \delta t_1(s)) \\
&= (\| \tilde{v}_1' \| - \| \tilde{v}_1 \|) - (\| \tilde{v}_2' \| - \| \tilde{v}_2 \|) \\
&\quad + a_{max}(\delta t_2(s) - \delta t_1(s)) \quad\quad (5)
\end{aligned}
$$

We can assume, *wlog*, the general picture shown in Fig. 6. As δs approaches 0, the difference $(\| \tilde{v}_1' \| - \| \tilde{v}_1 \|)$ in the norms of $\tilde{v}_1$ and $\tilde{v}_1'$ will be measured in the direction perpendicular to the norm contour at $\tilde{v}_1$. As shown in Fig. 7, the maximum change in $\| \tilde{v}_1 \|$ will occur in this direction. Since, in general, $\tilde{a}_1$ will be applied in another direction, we will not achieve the maximum change. The discrepancy between the two can be measured by taking the difference between the maximum change in velocity and the projection of the given change in velocity, $\tilde{a}_1 \delta t_1$ in the direction of the maximum change, as shown in Fig. 7. As $\delta s \to 0$, this measures the actual discrepancy. Similarly, we can get the discrepancy in the $\tilde{v}_2$ case. Since the direction of $\tilde{a}_1 \delta t_1$ and $\tilde{a}_1$ is the same (as is that of $\tilde{a}_2 \delta t_2$ and $\tilde{a}_2$), we can translate the discrepancies in velocities to those in accelerations as shown in Fig. 5. As can be seen from Fig. 5, the discrepancies, given by ϵ_1 and ϵ_2, are proportional to the centripetal accelerations. The actual discrepancy for $\tilde{a}_2$ has value at most ϵ_2 shown in Fig. 5 since due to the convexity of the norm, the point p lies above the point q. Therefore, we have

$$\frac{\epsilon_1}{\epsilon_2} = \frac{v_1^2}{v_2^2} \quad\quad (6)$$

where v_1 and v_2 denote the magnitude of the corresponding velocities in the L_2 norm. Now, in (5) we have $(\| \, \tilde{\mathbf{v}}_1' \, \| - \| \, \tilde{\mathbf{v}}_1 \, \|) = (a_{max} - \epsilon_1)\delta t_1$ and $(\| \, \tilde{\mathbf{v}}_2' \, \| - \| \, \tilde{\mathbf{v}}_2 \, \|) = (a_{max} - \epsilon_2)\delta t_2$. Combining with (6) and rearranging we get

$$
\begin{aligned}
\delta\vartheta(s) &= \epsilon_2\delta t_2 - \epsilon_1\delta t_1 \\
&= \epsilon_1\left(\frac{v_2{}^2\delta t_2 - v_1{}^2\delta t_1}{v_1{}^2}\right)
\end{aligned}
\tag{7}
$$

As $\delta s \to 0$, we have $v_1 = (\delta s/\delta t_1)$ and $v_2 = (\delta s/\delta t_2)$ and therefore $v_1\delta t_1 = v_2\delta t_2$. Therefore,

$$
\delta\vartheta(s) = \epsilon_1\delta t_1\left(\frac{v_2 - v_1}{v_1}\right)
\tag{8}
$$

Since $v_2 \le v_1$, $\delta\vartheta(s) \le 0$ as $\delta s \to 0$. In other words $\frac{\delta\vartheta(s)}{\delta s}$ is non-positive i.e. $\vartheta(s)$ is non-increasing. ∎

We are now ready to prove the loop elimination theorem in the case of arbitrary norms :

Theorem 1 *Consider a solution* $\mathbf{a_1} : [0, b_1] \to \Re^2$ *to some given problem* $(O, a_{max}, v_{max}, \mathbf{S}, \mathbf{F})$ *where* $\dot{s} = \dot{f} = 0$. *Let* $\tilde{\mathbf{p}}_1 : [0, l_1] \to C$ *be the arc-parametrized path corresponding to this solution. If there exist* c, $d \in [0, l_1]$, $c \neq d$, *such that* $\tilde{\mathbf{p}}_1(c) = \tilde{\mathbf{p}}_1(d)$ *(i.e. the path intersects itself) then there exists a solution* $\mathbf{a_2} : [0, b_2] \to \Re^2$ *with corresponding path* $\tilde{\mathbf{p}}_2 : [0, l_2] \to C$ *such that*

1. $b_2 \le b_1$.

2. there do not exist points c' and d' such that $c' \neq d'$ and $\tilde{\mathbf{p}}_2(c') = \tilde{\mathbf{p}}_2(d')$, (i.e. the path does not intersect itself).

Proof: The idea of the proof is as follows : We will construct the new trajectory by moving along the path of the original trajectory till we come to rest at the loop intersection point. At this point, we will ignore the loop and continue along the rest of the path. We will show that coming to a stop at the intersection does not penalize us and that the new trajectory takes time which is less than or equal to that taken by the original one.

Consider the portion of the (arc-parametrized) trajectory from d to l_1. By Lemma 1 above, there exists a solution $\tilde{a}_2$ defined on $[d, l_1]$ such that $\tilde{\mathbf{p}}_2(s) = \tilde{\mathbf{p}}_1(s)$ in this range, $\vartheta(s)$ is non-increasing in this interval and $\tilde{\mathbf{v}}_2(d) = 0$. Similarly, we can define $\tilde{a}_2$ on the interval $[0,c]$ by applying the above lemma, in the reverse direction, from c to 0. Now, $\tilde{\mathbf{v}}_2(c) = 0$ and $\tilde{\mathbf{v}}_2(d) = 0$. And c and d correspond to the same point in C. So, we can can now define the new acceleration $\tilde{a}_2$ as stated with $c = d$. If there are any more loops or self-intersections,

we can repeat the construction to eliminate them. (We note that this includes trivial loops of the kind where the particle just moves back and forth across the same portion of the path.) Since both pieces constitute a solution to and follow the path of the original trajectory, the combined path is a legal path and the combined solution is a solution to the original problem. We now show that for each loop elimination, we get a solution which takes time less than or equal to that of the trajectory in the previous iteration.

Consider the portions of the original and the new trajectories from d to l_1. (By abuse of notation, we will use the same arc-parametrization for this piece of both trajectories.) Treat these two trajectories as complete trajectories in themselves from d to l_1. Thus, e.g. $\Delta t(s)$ measures the time difference starting from d. But $\vartheta(d) = \| \, \Delta\mathbf{v}(d) \, \| + \Delta t(d)a_{max}$. Putting $\Delta t(d) = 0$ and $v_2(d) = 0$, we get $\vartheta(d) = \| \, \mathbf{v_1}(d) \, \|$. Similarly, we get $\vartheta(l_1) = \Delta t(l_1)a_{max}$. In like manner, consider the portions of the original and new trajectories from c to 0 in reverse arc-parametrization. We get $\vartheta(c) = \| \, \mathbf{v_1}(c) \, \|$. and $\vartheta(0) = \Delta t(0)a_{max}$ where we abuse our notation and denote the "time difference" between these two segments at 0 by $\Delta t(0)$. From Lemma 1 we have $\vartheta(d) \ge \vartheta(l_1)$ and $\vartheta(c) \ge \vartheta(0)$. Thus,

$$
\| \, \mathbf{v_1}(c) \, \| + \| \, \mathbf{v_2}(d) \, \| \ge (\Delta t(0) + \Delta t(l_1))a_{max}
\tag{9}
$$

Now, consider the loop. We can assume, *wlog*, the picture in Figure 8 (a). Here, c denotes the entry point of the loop and d the exit point. Fig. 8 (b) shows the velocity vectors $\mathbf{v_1}(c)$ and $\mathbf{v_1}(d)$ with the respective norm contours. Consider the component of $\mathbf{v_1}(c)$ in the direction perpendicular to the tangent to the contour at the point corresponding to $\mathbf{v_1}(c)$. By convexity, this component has norm at least as much at that of $\mathbf{v_1}$. When the velocity direction is along the tangent direction (point m in Fig. 8 (b)), this component becomes zero. The time required is therefore $\ge \frac{\|\mathbf{v_1}(c)\|}{a_{max}}$. We can use a similar argument for the velocity $\mathbf{v_1}(d)$, i.e. consider the component of $\mathbf{v_1}(d)$ in the direction perpendicular to the tangent to the contour at $\mathbf{v_1}(d)$. The time required for this component to go to zero by moving in the reverse (clockwise) direction is similarly $\frac{\|\mathbf{v_1}(d)\|}{a_{max}}$. However, in order to get a lower bound on the time taken in the loop as a sum of the times taken for the above two cases, we must ensure that the tangent direction at d has slope less than that at c as shown in Figure 8 (b) (the slope being defined as the angle made with the positive x axis). To see that this is true, we observe that from convexity of the norm contour the tangent direction has decreasing slope as we move clockwise along the contour of the norm. Thus the tangent direction at d has slope less than that at c. In the limiting case, the object leaves the loop in the direction exactly opposite to the direction of entrance.

In this case the tangent directions coincide. Thus, in all cases we get two distinct time intervals, one from c to m and the other from q to n with

$$t_1(m) - t(c) \geq \frac{\| \mathbf{v_1}(c) \|}{a_{max}}$$

$$t_1(d) - t(n) \geq \frac{\| \mathbf{v_1}(d) \|}{a_{max}}$$

Hence,

$$t_1(d) - t(c) \geq \frac{\| \mathbf{v_1}(c) \| + \| \mathbf{v_1}(d) \|}{a_{max}}$$

Combining this with (9) we get,

$$t_1(d) - t(c) \geq \Delta t(0) + \Delta t(l_1)$$

Thus, the time taken in the loop is at least as much as the time difference between the two trajectories. Ergo, we can eliminate the loop to get a new trajectory, without the loop, which takes time less than or equal to that taken by original trajectory. Repeating this procedure, we obtain a faster trajectory which has a path which does not intersect itself.

∎

7 Optimality of the canonical trajectory

In this section, we show that for any problem there exists a canonical solution, which is time-optimal.

Theorem 2 *Let* $\mathbf{a_1} : [0, b_1] \to \Re^2$ *be an optimal solution to some problem* $(O, a_{max}, v_{max}, \mathbf{S}, \mathbf{F})$. *Then there exists an optimal solution* $\mathbf{a_2}$ *to this problem which is canonical and which has* $O(n)$ *segments.*

Proof: We will show the existence of an optimal canonical trajectory by deforming the path corresponding to the given acceleration homotopically into one which is canonical. Starting at s, we move along the path $\mathbf{p_1}$ till we reach a point z which has the following property : assume that the trajectory from $(\mathbf{s}, \dot{\mathbf{s}})$ to $(\mathbf{z}, \dot{\mathbf{z}}) = \mathbf{Z}$ is canonical. Thus we get a new trajectory from $\mathbf{S}$ to $\mathbf{F}$ such that the trajectory passes through $\mathbf{S}$ and $\mathbf{Z}$ but such that the new trajectory has a canonical portion from $\mathbf{S}$ to $\mathbf{Z}$. The point $\mathbf{Z}$ is extremal in the sense that the canonical path from s to z is tangent to a vertex or to a segment of the obstacle space but is a legal path through C. In other words, we "canonize" the trajectory as far as we can till it touches some contact point. (We will refer to points in CP such as $\mathbf{Z} = (\mathbf{z}, \dot{\mathbf{z}})$ where z is a contact point in C as "contact points".) Denote this new trajectory by $\mathbf{\Gamma_1}$. It is possible for the path corresponding to this trajectory to go through more than one contact point. Let $\mathbf{Z_1}$ be the furthest contact

point down the trajectory between $\mathbf{S}$ and $\mathbf{Z}$. Since the portion of the trajectory between $\mathbf{S}$ and $\mathbf{Z}$ is canonical, it follows that the portion between $\mathbf{S}$ and $\mathbf{Z_1}$ is a canonical segment. We now repeat the process with the new trajectory $\mathbf{\Gamma_1}$, starting at $\mathbf{Z_1}$ to obtain another trajectory $\mathbf{\Gamma_2}$ and so on. Thus, at the k^{th} step, we have a trajectory $\mathbf{\Gamma_k}$ and contact points $\mathbf{Z_1}$ to $\mathbf{Z_k}$ on it such that the segments $\overline{\mathbf{Z_i}, \mathbf{Z_{i+1}}}$ (for $1 \leq i \leq k - 1$) are canonical. Then, we canonize the trajectory from $\mathbf{Z_k}$ till we obtain an extremal point $\mathbf{Z'}$ such that the trajectory from $\mathbf{Z_k}$ to $\mathbf{Z'}$ has a contact point $\mathbf{Z_{k+1}}$ which is closest to $\mathbf{Z'}$.

We need to prove several things in order to complete the proof of the theorem. First, we need to show that for any k the deformation of the trajectory $\mathbf{\Gamma_k}$ into $\mathbf{\Gamma_{k+1}}$ is continuous so as to ensure that even if the saturated direction changes the trajectory obtained does not pass through obstacle space. I.e. we need to ensure that there exists a point $\mathbf{Z'}$ such that the canonized section between $\mathbf{Z_k}$ and $\mathbf{Z'}$ passes through a contact point $\mathbf{Z_{i+1}}$. This can be done only by ensuring that we deform the trajectory continuously. But this follows from the observation made in Section 5. We noted that the position and velocity were continuous functions of the end-point position and velocity. At the k^{th} step, we move along the path continuously from z_k. We are therefore varying the position and velocity of the end-point continuously. Accordingly, the canonical segment so obtained, starting at $\mathbf{Z_k}$, also varies continuously. Hence, we are assured that either we will reach the final position or the segment will touch a point z_{k+1} of obstacle space. We then repeat the process starting at $\mathbf{Z_{k+1}}$. This ensures the validity of the above process of obtaining a canonical trajectory from the given one.

Next, we show that we can obtain a trajectory which has $O(n)$ contact points or we can transform it into one which does. If the trajectory passes through the same contact point more than once then it has a loop and by the loop elimination theorem we can replace it with another trajectory which does not pass through the same point more than once. This clearly does not affect the fact that the trajectory is made of canonical segments. As shown in Fig. 9(a), it is possible for the path corresponding to this new trajectory to touch the same edge of the obstacle space at different points. Let $\mathbf{w}$ and $\mathbf{w'}$ denote the first and last points which touch the same edge as shown in the Fig. 9(a). We note that the velocity at $\mathbf{w'}$ either has the same or opposite direction as that at $\mathbf{w}$. We can therefore treat the section of the trajectory from $\mathbf{w}$ to $\mathbf{w'}$ as a one-dimensional problem. Hence we can replace the entire section with a path running along the edge from $\mathbf{w}$ to $\mathbf{w'}$ as shown in Fig. 9(b) - the command acceleration for this section being bang-bang and parallel to the edge. This transformation does not affect the canoni-

cal nature of the trajectory since in the one-dimensional case a single bang-bang switch gives the optimal trajectory. If we treat only the first and last points along such a path (i.e. one moving along an edge) as contact points, we see that each edge can contribute only two contact points to the path. Other configurations which may arise with contacts at a given segment are treated similarly. Thus, the path is made up of $O(n)$ canonical segments.

Finally, we see that the trajectory obtained is time-optimal : for all k, Γ_{k+1} takes time less than or equal to that taken by Γ_k. This follows from the fact that in Γ_k the section from Z_k to Z' is bang-bang in the constrained direction and therefore takes time less than or equal to that taken by Γ_k between the same points. The theorem therefore follows.

■

8 Generating the optimal canonical trajectory

In this section, we will sketch the process of generating the optimal canonical trajectory. From the previous section, we know that there exists an optimal canonical trajectory from $\mathbf{S}$ to $\mathbf{F}$ which passes through at most n distinct vertices and has at most $2n$ contact points which touch edges. In this section, we show how to find it using only a polynomial (in the input) amount of space. We will use the theory of the reals to that end. An expression in the first order theory of the reals has the form

$$Q_1 x_1 \ldots Q_n x_n P(x_1, \ldots, x_n)$$

where each Q_i is one of the quantifiers $\exists$ or $\forall$ and $P(x_1, \ldots, x_n)$ is a quantifier free boolean formula with atomic predicates of the form $f_j(x_1, \ldots, x_n) = 0$ or $g_k(x_1, \ldots, x_n) \geq 0$, f_j and g_k being real polynomials.

For any two (indeterminate) points in free space, with two (indeterminate) velocities, we can write a collection of polynomial clauses which describe the canonical segment between the two points, in terms of the end-point indeterminates, as was shown in Section 5. Combining these clauses with other clauses which check that the segment does not pass through obstacle space, we get a clause $Q(\mathbf{z_i}, \mathbf{z_{i+1}}, \mathbf{v_i}, \mathbf{v_{i+1}}, t_i, t_{i+1})$ which is true iff there exists a canonical segment joining $\mathbf{z_i}$ and $\mathbf{z_{i+1}}$ with velocities $\mathbf{v_i}$ and $\mathbf{v_{i+1}}$ and at times t_i and t_{i+1} respectively. Let M denote a sequence of vertices and edges, i.e. M describes a possible sequence of contact points for a canonical trajectory. For every vertex point in S, we have two indeterminates, the velocities in the x and y directions. Similarly for every edge, we have first of all two contact points, and for each contact point we have two indeterminates - one is the distance along the edge and the other is the

velocity magnitude (since we know the velocity direction is along the edge from the first contact point to the second one). For each point $\mathbf{z_i}$ (whether a vertex or edge point), we will represent its indeterminates by $\mathbf{p_i}$. Let $R_M(\mathbf{z_0}, \mathbf{z_{3n+1}}, t)$ (where $(\mathbf{z_0}, \mathbf{v_0}) = \mathbf{S}$ and $(\mathbf{z_{3n+1}}, \mathbf{v_{3n+1}}) = \mathbf{F}$) be true iff there exists a path of at most $3n + 1$ canonical segments from $\mathbf{S}$ to $\mathbf{F}$ following the sequence M. We can write R_M in terms of Q as follows,

$$\exists \mathbf{p_0} \, \exists \mathbf{p_1} \ldots \exists \mathbf{p_{3n+1}} \, \exists t_0 \, \exists t_1 \ldots \exists t_{3n+1}$$
$$\bigcap_{i=0,\ldots,n} Q(\mathbf{z_i}, \mathbf{z_{i+1}}, \mathbf{v_i}, \mathbf{v_{i+1}}, t_i, t_{i+1}) \cap (t = t_{3n+1})$$

where the $\mathbf{p_i}$'s denote the indeterminates corresponding to each contact point as mentioned above. We will try every possible sequence of vertices and edges. At each stage we will compare a new sequence M' with the previous best, say M, using the following formula,

$$\exists t' \, \forall t \, (R_{M'}(\mathbf{z_0}, \mathbf{z_{3n+1}}, t'))$$
$$\wedge (R_M(\mathbf{z_0}, \mathbf{z_{3n+1}}, t) \Rightarrow (t' < t))$$

We note that the number of alternations between $\forall$'s and $\exists$'s is a constant (the Q's each have a $\forall$ but all the Q's are placed together, so there is only one alternation). We therefore use Renegar's PSPACE algorithm [Re] for the theory of reals with a bounded number of quantifier alternations. At each call to the algorithm, we only need remember the previous best, thus we need only a polynomial amount of space all told. We can get a polynomial number of bits of the minimum time and the velocities at each contact by making a polynomial number of calls as above. Alternately, we can use Renegar's algorithm to give us the time and velocities as roots of some polynomials. Therefore, we have

Theorem 3 *Under the L_∞ norm, the two dimensional kinodynamic problem $(O, a_{max}, v_{max}, \mathbf{S}, \mathbf{F})$, where $\mathbf{S} = (s, 0)$ and $\mathbf{F} = (f, 0)$, is in PSPACE.*

9 Conclusions and Open Problems

In this paper we gave the first algorithm for generating an exact time-optimal trajectory for a kinodynamic planning problem in the plane. This algorithm runs in PSPACE and takes exponential time. We characterized the nature of a class of solutions which provides us with a time-optimal solution. We proved a tracking lemma which is applicable to arbitrary norms. This lemma should prove to be useful in dealing with the different variations to this problem and is also of general interest since it provides bounds on the time difference between a given trajectory and one which tracks it. We then showed how to eliminate loops from a given trajectory in order to limit the number of contact points a time-optimal trajectory might have. Finally, we drew upon

277

the theory of reals (with a bounded number of alternations between quantifiers) to obtain a time-optimal solution. Several questions arise almost immediately :

1. The complexity of this problem - is it NP-hard or can it be done in polynomial time ?

2. Can the above algorithm be modified to handle the case of non-zero velocities at the beginning and end of the trajectory ?

3. Can the above approach be extended to the case of arbitrary norms, especially the L_2 norm? (See [RT] for work in this direction.) We have shown how to eliminate loops for arbitrary norms; the main problem now is characterizing the time-optimal solutions.

4. Exact algorithms for the three-dimensional case.

There are several other open problems within the broad framework of kinodynamic planning. E.g., as mentioned earlier, Fortune and Wilfong [FW] and Jacobs and Canny [JC] address the problem of determining the shortest bounded curvature path in the plane in the presence of obstacles. However, at present, only an approximation algorithm is known for this problem and it would be of significant interest to determine the complexity of this problem and to provide an exact algorithm for it. Additionally, a considerable amount of work needs to be done to translate such results into practice.

10   References

[BDG]   Bobrow, J.E., Dubowsky, S. and Gibson, J.S., "On the Optimal Control of Robotic Manipulators with Actuator Constraints", *Proc. of the ACC*, San Francisco, (1983), pp. 782-787.

[CDRX]   Canny, J., Donald, B., Reif J. and Xavier, P., "On the Complexity of Kinodynamic Planning", *Proc. 29th IEEE Symposium on Foundations of Computer Science*, New York, (1988), pp. 306-318.

[CR]   J.F. Canny and J. Reif, "New lower bound techniques for robot motion planning," in *28th IEEE Symposium on Foundations of Computer Science*, (Los Angeles), 1987.

[FW]   Fortune, S. and Wilfong, G., "Planning Constrained Motion" *Proc. ACM Symposium on the Theory of Computing*, Chicago, (1988), pp. 445-459.

[Hol]   J. M. Hollerbach, "Dynamic scaling of manipulator trajectories," in *Proceedings of the ACC*, (San Francisco, CA), pp. 752-756, The American Automatic Control Council, June 1983.

[JC]   P. Jacobs and J. Canny, "Planning smooth paths for mobile robots," in *Proceedings of the 1989 International Conference on Robotics and Automation*, pp. 2-7, IEEE, May 1989.

[JHCP]   P. Jacobs, G. Heinzinger, J. Canny and B. Paden, "Planning Guaranteed Near-Time-Optimal Trajectories for a Manipulator in a Cluttered Workspace", Technical Report ESRC 89-20/RAMP 89-15, Engineering Systems Research Center, University of California, Berkeley, October 1989.

[O]   C. Ó'Dúnlaing, "Motion planning with inertial constraints", *Algorithmica*, vol. 2(4), (1987), pp. 431-475

[RT]   Reif, J. and Tate, S., "Approximate kinodynamic planning using L_2 norm dynamic bounds", Duke University Tech. Report (1989).

[Re]   Renegar, J., "On the computational complexity and geometry of the first order theory of the reals", Technical Report no. 853, School of O.R. and I.E., Cornell University, NY, July 1989.

[SH]   G. Sahar and J. M. Hollerbach, "Planning of minimum-time trajectories for robot arms", Tech. Rep. A.I. Memo No. 804, MIT, November 1984.

[Sch]   H.M. Schaettler, "On the Optimality of Bang-Bang Trajectories in $\Re^3$", *Bull. AMS*, vol. 18(1), pp. 113-6 (1987).

[SD]   Z. Shiller and S. Dubowsky, "Global Time-Optimal Motions of Robotic Manipulators in the Presence of Obstacles", *IEEE Int. Conf. on Robotics and Automation*, Philadelphia, (1988).

[SM]   K. G. Shin and N. D. McKay, "Minimum-time control of robotic manipulators with geometric path constraints", *IEEE Transactions on Automatic Control*, vol. AC-30, pp. 531-541, June 1985.
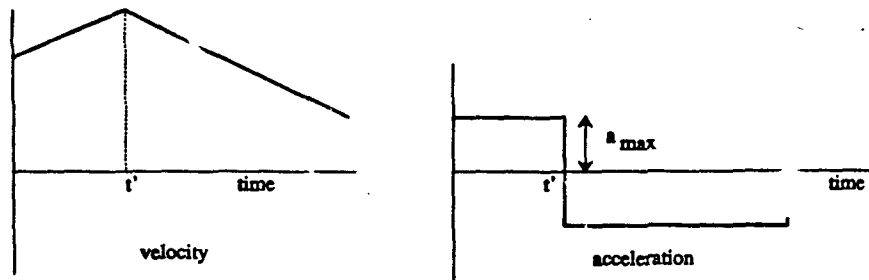
Figure 1. The time-optimal velocity and command acceleration in the one-dimensional case
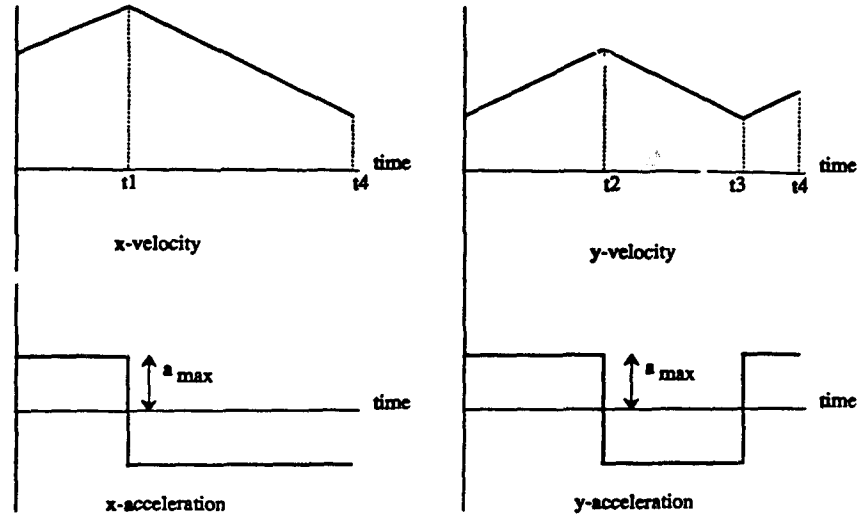

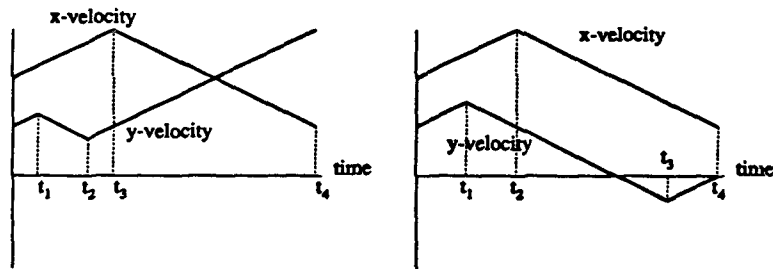
Figure 2. Canonical motion between two points

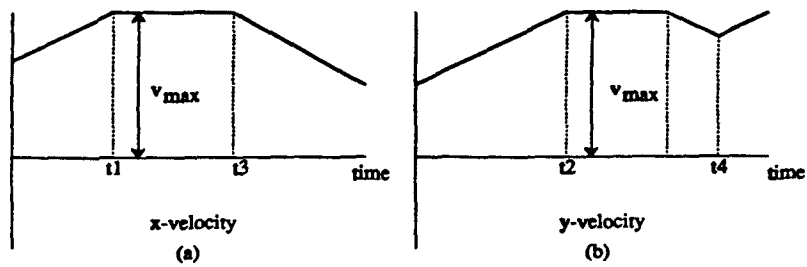Figure 3. Other types of velocity profiles

Figure 4

CONTOUR FOR q_{max}

Figure 5

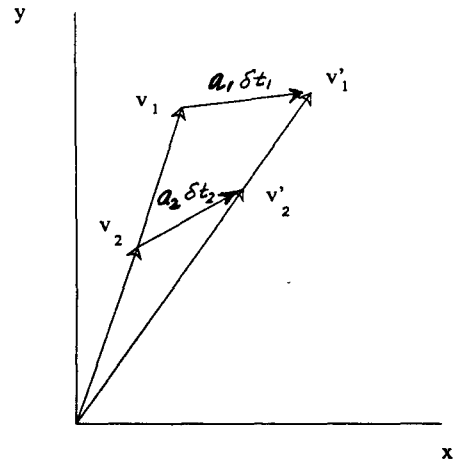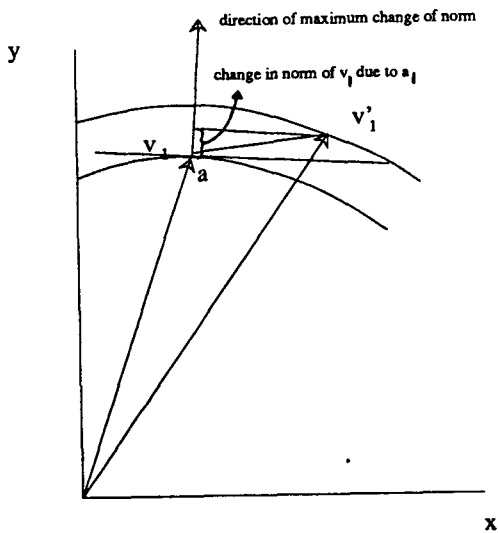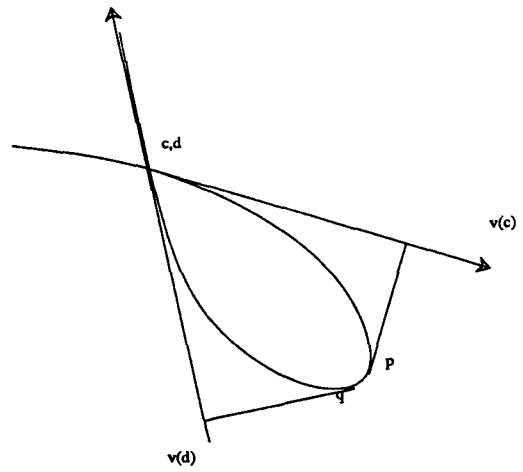Figure 6

direction of maximum change of norm
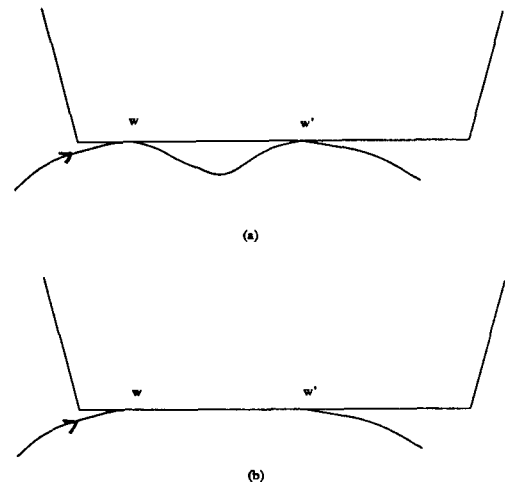
change in norm of v_1 due to a_1
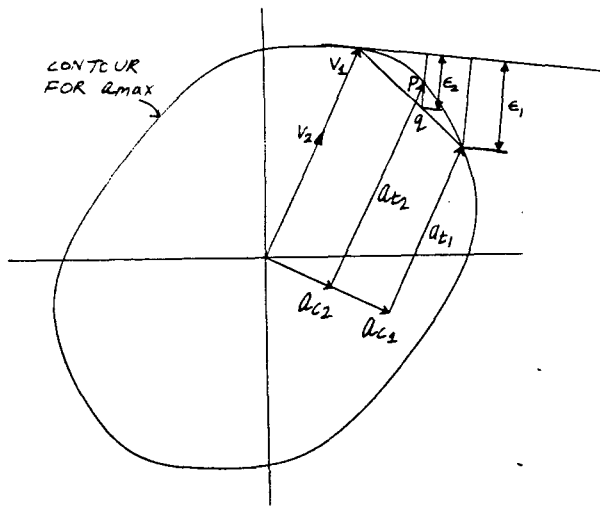
Figure 7

Figure 8 (a)

FIGURE 8 (b)

(a)

(b)

Figure 9