

Finite-Horizon Optimal Control and Stabilization of Time-Scalable Systems ¹

J. Alexander Fax

Richard M. Murray

Engineering and Applied Science
California Institute of Technology
Pasadena, CA 91125
{fax,murray}@cds.caltech.edu

Abstract

In this paper, we consider the optimal control of time-scalable systems. The time-scaling property is shown to convert the PDE associated with the Hamilton-Jacobi-Bellman (HJB) equation to a purely spatial PDE. Solution of this PDE yields the value function at a fixed time, and that solution can be scaled to find the value function at any point in time. Furthermore, in certain cases the unscaled control law stabilizes the system, and the unscaled value function acts as a Lyapunov function for that system. The PDE is solved for the well-known example of the nonholonomic integrator.

1 Introduction

In addressing optimal control problems, one frequently looks for structures in the dynamical system being controlled which influence the structure of optimal trajectories and which simplify the solution of the problem. In this paper, we will discuss one such structure, namely the optimal control of dynamical systems which are time-scalable, meaning the equations of motion remain the same when time is reparameterized.

Examples of time scalable dynamic systems include driftless systems and mechanical systems without potential energy. Application areas include locomotion of systems with nonholonomic constraints, control of underwater vehicles, and satellite reorientation [1, 7, 4]. These systems are not linearly controllable (when underactuated) and analysis of these systems has relied heavily on concepts and tools from differential geometry.

Optimal control of these classes of systems has been studied previously [8, 9], primarily in the context of simplifying the Euler-Lagrange equations using symme-

try reduction. In contrast, we present some preliminary results regarding the structure of the Hamilton-Jacobi-Bellman (HJB) equation using the notion of time scalability. Specifically, we show that if the cost function is compatible with the time scaling, one can deduce *a priori* the time-dependence of the value function associated with the HJB PDE. This results in conversion of that PDE, which depends on space and time, into a purely spatial PDE. The solution of this PDE, combined with the known time dependency, yields the control law which generates the optimal finite time trajectories. If the time variation of the control law is omitted, the same value function can produce a stabilizing (though not optimal) control law, with the value function playing the role of Lyapunov function.

The paper is structured as follows. In Section 2, we define time scalability and derive some properties of time-scalable dynamic systems. In Section 3, we discuss the application of time scalability to the optimal control problem. In Section 4, we solve the PDE for a simple example and show via simulation how it produces both finite-time and infinite-time stabilizing trajectories.

2 Time Scaling

2.1 Definitions

Let $x \in \mathbb{R}^p$, $u \in \mathbb{R}^s$. Let an n th order control system be a dynamical system of the form $x^{(n)} = f(x, \dot{x}, \dots, x^{(n-1)}, u)$, where the superscript in parentheses indicates the order of differentiation. Let $\bar{x} \in \mathbb{R}^{pn}$ be a point of the form $\bar{x} = (x, \dot{x}, \dots, x^{(n-1)})$.

Definition 1 An n th order control system is said to be *affinely time scalable* if, for any affine scaling of time $\tau = \alpha t + \beta$, there exists a scaling of u such that the scaled equations of motion are identical to the original equations. Let $K : \mathbb{R}^{pn} \times \mathbb{R}^s \times \mathbb{R} \rightarrow \mathbb{R}^{pn} \times \mathbb{R}^s \times \mathbb{R}$ denote the time scaling operator associated with a given time scaling. That is, $K(\bar{x}, u, t)$ maps each of the operands to its time-scaled counterpart.

¹Research supported in part by AFOSR Grant F49620-99-1-0190. First author also supported by an NSF Graduate Research Fellowship and an ARCS Foundation Fellowship

We can easily write how K acts on the various operands by applying the chain rule. We see that

$$\frac{dx}{dt} = \frac{dx}{d\tau} \frac{d\tau}{dt} = \alpha \frac{dx}{d\tau} \quad (1)$$

and thus we see that K maps \dot{x} to \dot{x}/α . More generally, we can see that K maps $x^{(j)}$ to $x^{(j)}/\alpha^j$. Because x is independent of time, it does not vary by time scaling.. To determine if a control system is time-scalable, one simply makes this substitution and then sees if u can be compatibly scaled such that α can be factored out of the equation.

Remark Because K acts on each operand independently, we will use $K(u)$, for example, to denote the image of u under the time scaling.

Example 1 First-order driftless systems, meaning systems of the form $\dot{x} = f(x)u$, are time scalable. If we scale time and substitute, we arrive at the equation

$$\frac{dx}{d\tau} = f(x) \frac{u}{\alpha}. \quad (2)$$

If we define $K(u) = u/\alpha$, and substitute, we recover the original equation. More generally, systems of the form $x^{(n)} = f(x)u$, are also time scalable, with the substitution $K(u) = u/\alpha^n$.

Example 2 Mechanical systems, including those with nonholonomic constraints, are time-scalable when the control vector fields are functions of the configuration space and the Lagrangian consist solely of a kinetic energy term. In coordinates, these systems take the form

$$\ddot{q}^i + \Gamma_{jk}^i \dot{q}^j \dot{q}^k = g^{ij}(q)u_j \quad (3)$$

where Γ are the Christoffel symbols associated with the connection which defines the system. See [1] for more information on nonholonomic mechanical control systems. When $u = 0$, the time-scalability corresponds with the well-known time scalability of geodesics [5]. When u is nonzero, it can still be scaled by the procedure noted above.

Definition 2 A point \bar{x} which exhibits the property $K(\bar{x}) = \bar{x}$ is said to be *scale invariant*. A scaling K for which $K(T) = T$, where $T \in \mathbb{R}$, is said to be *T-invariant*.

Remark The scale invariant points are clearly those for which all derivatives are zero. For a first-order control system, these consist of the entire space. A time scaling is T -invariant if $\beta = (1 - \alpha)T$. We see that the function $(T - t)^{-n}$ is compatible with any T -invariant time scaling, and the scale factor is $f(\alpha) = \alpha^{-n}$.

Definition 3 A function $L(\bar{x}, u, t)$ is said to be *compatible* with a time-scaling operation K if the following

holds there exists a function $s : \mathbb{R} \rightarrow \mathbb{R}$ such that:

$$L(K(\bar{x}, u, t)) = s(\alpha)L(\bar{x}, u, t) \quad (4)$$

Remark An obvious property of $s(\alpha)$ is that $s(1) = 1$. Also, it is easy to see that $s(\alpha^{-1}) = s(\alpha)^{-1}$.

We see that the class of linearly time-scalable systems is fairly rich, and includes systems for which significant work has been done in the areas of controllability and trajectory generation. To our knowledge, however, the time scaling property of these systems has never explicitly been exploited. In the development which follows, we discuss the implications time scalability has for the optimal control problem.

2.2 Properties of Time-Scalable Systems

The following proposition follows immediately:

Proposition 1 *Given a time-scalable control system, initial conditions \bar{x}_i at initial time $t = t_i$, a control time history $u(t)$ defined on the interval $t \in [t_i, t_f]$, and a resulting trajectory $\bar{x}(t)$ defined on the interval, the time scaled control $K(u)$ defined on the interval $K([t_i, t_f])$, when applied to the time-scaled initial condition $K(\bar{x}_i)$ at time $K(t_i)$, will produce the time-scaled trajectory $K(\bar{x}(K(t)))$.*

Remark In other words, the time scaling of an integral curve yields another integral curve generated by the time-scaled control history. In particular, if the initial and final points are scale invariant, then a control time history which connects the two points in a given time interval can be scaled to connect those two points in any interval of time.

Proposition 2 *Given a time-scalable control system, a time scaling K for which $\alpha \geq 0$, a function $L(\bar{x}, u)$ which is compatible with K and for which $f(\alpha) \geq 0$, and a control history $u^*(t), t \in [t_i, t_f]$ which drives the system between $m_i \in M$ at time $t = t_i$ and $m_f \in M$ at time $t = t_f$ and minimizes the quantity*

$$J = \int_{t_i}^{t_f} L(\bar{x}, u) dt. \quad (5)$$

Then the time-scaled control history $K(u^(K(t)))$ minimizes J between the time-scaled endpoints.*

Proof: To see the correspondence of J with its scaled counterpart, let us evaluate

$$\hat{J} = \int_{K(t_i)}^{K(t_f)} L(\bar{x}, v) d\tau \quad (6)$$

with endpoints $K(m_i)$ and $K(m_f)$. We can evaluate the integral in the unscaled interval by changing coordinates through K^{-1} . Letting $t = K^{-1}(\tau)$, we apply

the formula for change of coordinates of an integral to arrive at

$$\hat{J} = \int_{t_i}^{t_f} L(K(\bar{x}, v))\alpha dt \quad (7)$$

where the extra α term is the Jacobian of the coordinate transformation. Replacing v with $K(u)$ and applying the definition of compatibility yields

$$\hat{J} = \int_{t_i}^{t_f} \alpha s(\alpha) L(\bar{x}, u) dt = \alpha s(\alpha) J \quad (8)$$

which is clearly minimized by $u = u^*$, or $v = K(u^*)$, since the quantity $\alpha s(\alpha)$ is assumed positive. ■

Proposition 2 guarantees that $K(u^*)$ does in fact drive the system to the desired endpoint, and hence is an admissible control history. Furthermore, it guarantees that every control history which is admissible for the scaled function is also admissible for the original function. Hence no other control history could exist which yields a lower J than $v = K(u^*)$.

We now consider the properties of the integral function J for a special case. Suppose the final point m_f of the previous proposition is scale invariant, and let us fix the final time t_f . The function J can be thought of as a function of the initial point \bar{x} and the initial time $t_i \in (-\infty, t_f)$. It is not an explicit function of u , since u is chosen according to the initial position and time. From the proof above, the following proposition holds:

Proposition 3 *Given the above conditions, $J(\bar{x}, t_i)$ is compatible with the set of all t_f -invariant scalings. Specifically,*

$$J(K(\bar{x}, t_i)) = \alpha f(\alpha) J(\bar{x}, t_i). \quad (9)$$

The following proposition shows how the *temporal* partial derivative can be converted to a *spatial* quantity which includes time as a parameter.

Proposition 4 *Let $L(\bar{x}, u, t)$ be a function which is compatible with the set of all T -invariant time scalings. For simplicity, assume $K(u) = u/\alpha^n$. Then, for all $t_0 \neq T$,*

$$\frac{\partial L}{\partial t} \Big|_{t=t_0} = \frac{1}{t_0 - T} \left(\frac{\partial s}{\partial \alpha} \Big|_{\alpha=1} L + \sum_{i=1}^{n-1} i \frac{\partial L}{\partial x^{(i)}} x^{(i)} + n \frac{\partial L}{\partial u} u \right). \quad (10)$$

Proof: Recall the definition of the partial derivative:

$$\frac{\partial L}{\partial t} \Big|_{t=t_0} = \lim_{\Delta t \rightarrow 0} \frac{L(\bar{x}, u, t_0 + \Delta t) - L(\bar{x}, u, t_0)}{\Delta t}. \quad (11)$$

We can map the first term on the right hand side using the time scaling operator. If we set

$$\alpha = \frac{t_0 - T}{t_0 + \Delta t - T}, \quad \beta = \frac{T \Delta t}{t_0 + \Delta t - T} \quad (12)$$

we recover the unique T -invariant time-scaling for which $K(t_0 + \Delta t) = t_0$. Thus,

$$L(\bar{x}, u, t_0 + \Delta t) = \frac{1}{f(\alpha)} L \left(x, \frac{\dot{x}}{\alpha}, \dots, \frac{x^{(n-1)}}{\alpha^{n-1}}, \frac{u}{\alpha^n}, t_0 \right). \quad (13)$$

Now $\alpha^{-n} = 1 + n\Delta t/(t_0 - T) + \mathcal{O}(\Delta t^2)$. Using this, we can Taylor expand the right hand side to first order, to arrive at

$$L(\bar{x}, u, t_0 + \Delta t) = \frac{1}{f(\alpha)} \left[L + \left(\sum_{i=1}^{n-1} i \frac{\partial L}{\partial x^{(i)}} x^{(i)} + n \frac{\partial L}{\partial u} u \right) \frac{\Delta t}{t_0 - T} \right], \quad (14)$$

where the right-hand side is evaluated at (\bar{x}, u, t_0) . If we recall that at $\Delta t = 0$, we have $\alpha = 1$ and thus $s(\alpha)$ goes to 1 in the limit. Substituting into Equation (11) and taking the limit, we have

$$\frac{\partial L}{\partial t} \Big|_{t=t_0} = \lim_{\Delta t \rightarrow 0} \frac{1 - s(\alpha)}{s(\alpha)\Delta t} L + \frac{1}{t_0 - T} \left(\sum_{i=1}^{n-1} i \frac{\partial L}{\partial x^{(i)}} x^{(i)} + n \frac{\partial L}{\partial u} u \right).$$

The final term can be simplified by noting that $\alpha = 1 - \frac{\Delta t}{t_0 + \Delta t - T}$, Taylor expanding $s(\alpha)$ to first order, and taking the limit as before. This results in the following:

$$\frac{\partial L}{\partial t} \Big|_{t=t_0} = \frac{1}{t_0 - T} \left(\frac{\partial s}{\partial \alpha} \Big|_{\alpha=1} L + \sum_{i=1}^{n-1} i \frac{\partial L}{\partial x^{(i)}} x^{(i)} + n \frac{\partial L}{\partial u} u \right)$$

which proves our proposition. ■

3 Optimal Control of Time-Scalable Systems

3.1 Application of Time-Scalability to the HJB equation

We are now ready to consider the relationship between time scaling and the optimal control problem. Recall that the value function $V(\bar{x}, t)$ associated with the Hamilton-Jacobi-Bellman (HJB) equation is defined as the minimum cost necessary to drive a system from a point \bar{x} at time t to a desired final condition at time (not necessarily finite) T . To apply the above theory to this problem, we consider a finite time optimal control problem with fixed final state which is scale invariant, and a cost function $L(\bar{x}, u, t)$ which is compatible with all time scalings. Given this, we see that the value function is given by

$$V(\bar{x}, t) = \min_u \int_t^T L(\bar{x}, u, \tau) d\tau. \quad (15)$$

Applying Proposition 3, we see that $V(\bar{x}, t)$ is compatible with all T -invariant time scalings, and hence Proposition 4 applies as well.

The HJB equation for this optimal control problem is given by

$$\frac{\partial V}{\partial t} = - \min_u \left[L(\bar{x}, u, t) + \frac{\partial V}{\partial \bar{x}} g(\bar{x}, u) \right] \quad (16)$$

Proposition 4 allows us to replace the left-hand side with a term involving V and its spatial partial derivatives. Solving this spatial PDE yields the value function at one point in time, and the value function at any other time can be found using the time scaling.

The presence of the $\frac{1}{t-T}$ term in Equation (10) indicates that as t approaches T , the value function will approach infinity at all points except the desired endpoint. This is expected - as time runs out, the cost associated with driving the system to the endpoint increases. Ordinarily, the HJB equation is solved by defining $V(\bar{x}, T)$ using a terminal cost function and propagating that function backwards in time. In our case, the $\bar{x}(T)$ is fixed, so a terminal cost constraint is not meaningful. Nonetheless, the Principle of Optimality from which the HJB equation is derived is still applicable, so the use of that equation is still valid.

In the next sections, we consider the cases of first-order systems and mechanical control systems in greater detail. For the former, we shall show that solving the spatial PDE yields a stabilizing controller in addition to solving the finite-horizon optimal control problem.

3.2 First-Order Systems

Let us consider a system of the form

$$\dot{x} = g(x)u \quad (17)$$

where $g(x)$ is C^1 , and a value function

$$V(x, t) = \min_u \frac{1}{2} \int_{t_0}^T u^T u dt \quad (18)$$

where V is minimized over control histories with the property that they integrate Equation (17) to the value $x(T) = 0$. Evaluating the HJB equation for this system yields

$$V_t = \frac{1}{2} V_x^T g(x) g^T(x) V_x \quad (19)$$

and optimal control $u^*(x, t) = -g^T(x) V_x$, where the subscripts on V indicate partial derivatives. We now turn to replacing V_t with its spatial counterpart. The cost function $u^T u$ is compatible with any time scaling, and the scaling function is $s(\alpha) = \alpha^{-2}$. The scaling for V , following Proposition 4, is $s(\alpha) = \alpha^{-1}$. Applying Proposition 4, we have

$$V_t = \frac{1}{T-t} V. \quad (20)$$

Furthermore, applying the definition of time-scalability, we see that

$$V(x, t) = \frac{T-t_0}{T-t} V(x, t_0). \quad (21)$$

Letting $t_0 = 0$ and defining $\tilde{V}(x) = V(x, 0)$, we can substitute into Equation (19) and derive the following spatial PDE:

$$\frac{1}{T} \tilde{V} = \frac{1}{2} \tilde{V}_x^T g(x) g^T(x) \tilde{V}_x. \quad (22)$$

Solving this PDE yields the optimal finite-horizon control law by setting $\tilde{u}(x) = -g^T(x) \tilde{V}_x$ and scaling $\tilde{u}(x)$ appropriately in time to arrive at $u^*(x, t)$.

We now consider the possibility of using $\tilde{u}(x)$ as a control law without time scaling, and asking whether that control law stabilizes the system. When \tilde{V} is smooth, it is easy to show, using Equations (17),(22) and the definition of $\tilde{u}(x)$, that

$$\dot{\tilde{V}} = -\frac{2}{T} \tilde{V}. \quad (23)$$

Because \tilde{V} is a value function, it must be positive at all points away from the origin, and thus $\dot{\tilde{V}}$ must be negative. This implies that \tilde{V} is a Lyapunov function, and that $\tilde{u}(x)$ asymptotically stabilizes the system. Thus, solving a single spatial PDE yields both optimal finite-horizon trajectories and an asymptotically stabilizing controller whose gain can be set. Conceptually, a controller of this form can be thought of as applying the optimal control one would use if the final time were always $T - t_0$ away, reminiscent of receding-horizon control concepts used in model predictive control [6, 10]. Note further that by varying the value of t_0 used to define \tilde{V} , we can influence the rate of convergence of \tilde{V} , i.e. the gain of the controller.

3.3 Mechanical Control Systems

In this section, we consider systems of the form

$$\ddot{q}^i + \Gamma_{jk}^i \dot{q}^j \dot{q}^k = g^{ij}(q) u_j \quad (24)$$

as discussed in Section 1, with the cost function in Equation (18). In this case, the scaling function is $s(\alpha) = \alpha^{-4}$, and the scaling function for V is $\alpha s(\alpha) = \alpha^{-3}$. Thus, following Proposition 3, we have the following equation for V :

$$V(q, \dot{q}, t) = \frac{(T-t_0)^3}{(T-t)^3} V(q, \frac{T-t}{T-t_0} \dot{q}, t_0), \quad (25)$$

from which V_t is calculated and found to satisfy Proposition 4. Again, we set $t_0 = 0$ and denote the spatial PDE as $\tilde{V}(x, \dot{x}) = V(q, \dot{q}, 0)$, where $x = q$, $\dot{x} = \dot{q}(T-t)/T$. As in the first order case, Equation (16) can be evaluated for this system, and we can substitute

Equation (25) and its derivative to arrive at the PDE for \tilde{V} :

$$\frac{3}{T}\tilde{V} - \frac{1}{T}\tilde{V}_x\dot{x} = -\tilde{V}_x\dot{x} + \tilde{V}_x\Gamma_{jk}^i\dot{x}^j\dot{x}^k + \frac{1}{2}\tilde{V}_x^T g(x)g^T(x)\tilde{V}_x. \quad (26)$$

If we attempt to use this value function as a Lyapunov function as done earlier, we arrive at the following equation:

$$\dot{\tilde{V}} = -\frac{3}{T}\tilde{V} - \frac{1}{2}\tilde{V}_x^T g(x)g^T(x)\tilde{V}_x + \frac{1}{T}\tilde{V}_x\dot{x}. \quad (27)$$

The above equation is not necessarily negative due to the last term. Given a solution \tilde{V} , this equation can be evaluated to determine if \tilde{V} is in fact a Lyapunov function.

3.4 Example - The Nonholonomic Integrator

As an example of this theory, we consider a classic nonlinear optimal control problem, namely the first-order nonholonomic integrator [2]. The system, is defined by

$$\begin{aligned} \dot{x}_1 &= u_1 \\ \dot{x}_2 &= u_2 \\ \dot{x}_3 &= x_1 u_2 - x_2 u_1 \end{aligned} \quad (28)$$

with the cost function in Equation (18). The system is driftless, linearly uncontrollable, and time-scalable. The HJB equation for this system is

$$V_t = \frac{1}{2} [V_{x_1}^2 + V_{x_2}^2 + dV_{x_3}^2 + 2V_{x_3}(x_1V_{x_2} - x_2V_{x_1})] \quad (29)$$

where $d = x_1^2 + x_2^2$. Following the development in Section 3.1, we can replace $V(x, t)$ with $\tilde{V}(x)$ and V_t with \tilde{V} in the above equation. We now seek a functional form for \tilde{V} which is consistent with this equation. If we assume the form $\tilde{V} = \tilde{V}(d, x_3)$, we can substitute and simplify the equation to

$$\tilde{V} = 2d\tilde{V}_d^2 + \frac{d}{2}\tilde{V}_{x_3}^2. \quad (30)$$

If we further assume that $\tilde{V} = df(r)/2$, where $r = x_3/d$, we can reduce the equation to an ODE:

$$f = f^2 - 2rf f' + (r^2 + \frac{1}{4})f'^2 \quad (31)$$

This equation is quadratic in f' , and can be factored to form two ODEs. To solve either ODE, we need an initial condition. Note that when $r = 0$, this is equivalent initial conditions in the $x_1 - x_2$ plane. Reviewing the equations of motion, we see that this case can be solved directly using constant $u_1 = x_1/T_r, u_2 = x_2/T_r$. This corresponds to $V(d, 0) = \frac{d}{2T_r}$, meaning $f(0) = 1$. One of the ODEs has a nonzero solution at $f(0) = 1$ for $r \geq 0$, and the other has a nonzero solution for $r \leq 0$. Thus, by solving the initial value problem for both positive and negative r using the ODE which does not

give a trivial solution, we derive $f(r)$, and thus the value function as well.

At points away from the x_3 axis, we have the following control law:

$$\tilde{u}_1 = -x_1(f - f'r) + x_2 f'/2 \quad (32)$$

$$\tilde{u}_2 = -x_2(f - f'r) - x_1 f'/2 \quad (33)$$

While $\lim_{r \rightarrow \pm\infty} f(r) = \infty$, the value function remains finite, as well as the limit of $f'(r)$. This allows us to define V at points where $x_1, x_2 = 0$. Furthermore, $\lim_{r \rightarrow \infty} u_i$ is finite, but dependent on the path along which one takes the limit. It can be shown that along the x_3 axis a control law of the form

$$\tilde{u}_1 = (2\tilde{V})^{1/2} \sin \theta \quad (34)$$

$$\tilde{u}_2 = (2\tilde{V})^{1/2} \cos \theta \quad (35)$$

for arbitrary θ yields a control law which is continuous along trajectories.

Figures 1 and 2 show optimal finite-time trajectories and stabilizing trajectories for two initial conditions. The stabilizing trajectories were generated using the value function at $t_0 = 9$ to give faster convergence. The initial conditions in Figure 1 are at a point where the value function is smooth. Those for Figure 2 are at a nonsmooth point. Both show the trajectory reaching zero at the appropriate time for the finite-horizon case, and decaying to zero in the infinite horizon case.

This example is a classic one in nonlinear control theory, and the optimal trajectory for initial conditions lying on the x_3 axis can be found many places in the literature. The trajectories found here are consistent with those results, and also provide the globally minimizing trajectory for all points in the space.

4 Future Work

In the example shown above, we see that the solution to the spatial PDE is nonsmooth. This is unsurprising - the existence of a smooth solution would imply the existence of a continuous stabilizing feedback, which was shown by Brockett to be impossible [3]. The discontinuity, however, is relatively benign, in that trajectories never cross it. It would be desirable to show that the discontinuities inherent in the solution do not prevent the control law from stabilizing the system. Such a proof would involve tools from nonsmooth analysis, and will be the subject of future research. Additionally, we seek to refine the conditions under which mechanical control systems can be stabilized. Issues of smoothness will have to be addressed in this setting as well. Another area of future work is the development of numerical methods for solving the spatial PDE which




Figure 1: Finite-Time Optimal Trajectory, Stabilizing Trajectory, $x_0 = (2, -3, 4)$

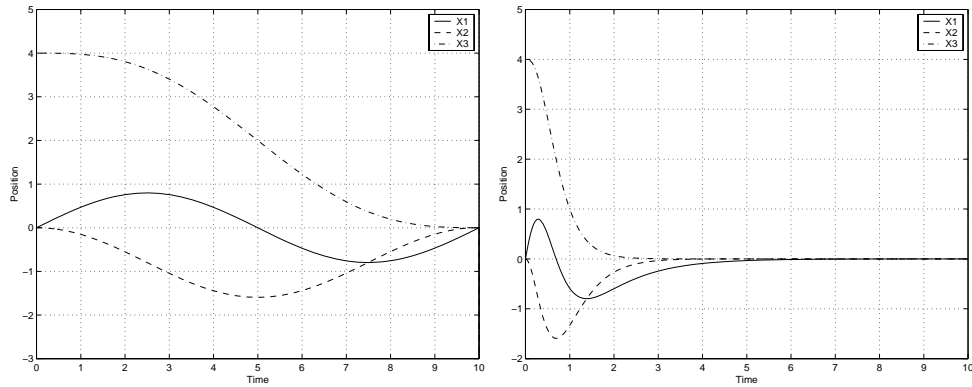


Figure 2: Finite-Time Optimal Trajectory, Stabilizing Trajectory, $x_0 = (0, 0, 4)$

exploit the structure of the underlying system. The literature regarding the geometric structure of the systems discussed in this paper is very rich. We expect to be able to exploit this structure in designing numerical methods.

References

- [1] A. Bloch, P. Krishnaprasad, J. Marsden, and R. Murray. Nonholonomic mechanical systems with symmetry. *Arch. Rat. Mech. Anal.*, 136:21–99, 1996.
- [2] R. W. Brockett. Control theory and singular Riemannian geometry. In P. Hilton and G. Young, editors, *New Directions in Applied Mathematics*. Springer-Verlag, 1981.
- [3] R. W. Brockett. Asymptotic stability and feedback stabilization. In R. W. Brockett, R. Millman, and H. Sussmann, editors, *Differential Geometric Control Theory*. Birkhauser, 1983.
- [4] F. Bullo, N. Leonard, and A. D. Lewis. Controllability and motion algorithms for underactuated Lagrangian systems on Lie groups. *To appear, IEEE Transactions on Automatic Control*, Nov. 97.
- [5] M. P. do Carmo. *Riemannian Geometry*. Birkhauser, Boston, 1992.
- [6] C. Garcia, D. Prett, and M. Morari. Model predictive control: Theory and practice. a survey. *Automatica*, 25(3):335–348, 1989.
- [7] S. Kelly and R. Murray. Geometric phases and locomotion. *Journal of Robotic Systems*, 12(6):417–431, 1995.
- [8] W.-S. Koon and J. Marsden. Optimal control for holonomic and nonholonomic mechanical systems with symmetry and Lagrangian reduction. *SIAM Journal on Control and Optimization*, 35(3):901–929, 1997.
- [9] J. Ostrowski, J. Desai, and V. Kumar. Optimal gait selection for nonholonomic locomotion systems. *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 786–791, 1997.
- [10] J. A. Primbs. *Nonlinear Optimal Control: A Receding Horizon Approach*. PhD thesis, California Institute of Technology, Pasadena CA, 1999.