Receding Horizon Control of Nonlinear Systems: A Control Lyapunov Function Approach

Thesis by

Ali Jadbabaie

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

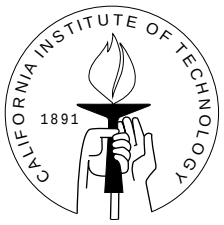

California Institute of Technology

Pasadena, California

2000

(Defended October 13, 2000)

To Nikroo

Acknowledgements

I would like to express my gratitude to the people who have made this thesis possible. First, I am obliged to my advisor, professor John Doyle, for taking the risk of accepting me as his student in the middle of the academic year. John has taught me to think independently, while creating a unique environment in which I could collaborate with different people. He has always given me inspiration, motivation, and support.

My deepest thanks go to professor John Hauser of the University of Colorado at Boulder. This research started when he came to Caltech as a visiting faculty. During his stay at Caltech, he was not only my coadvisor and mentor, but also a colleague and a great friend. He taught me mathematical rigor, while guiding my intuition to obtaining interesting results. His unique position as a control theoretician with a strong mathematical background as well as being a former pilot in the U.S. Air Force made him an invaluable resource for this project.

I am also thankful to professor Richard Murray, who has always made time for me in his extremely busy schedule and has given me valuable feedbacks at several occasions. Richard has been my role model for organization, professionalism, and hard work. I acknowledge professor Jerry Marsden for accepting to be a member of my thesis committee.

Aside from the faculty, there are many students and post doctoral scholars who have helped me through my studies and have made my experience at Caltech most gratifying. In particular, I would like to thank my friend Jim Primbs, for introducing me to receding horizon control. I am grateful to my officemate, Jie Yu, for her contributions to Section 5.2 and to Monica Giannelli for contributing to Section 2.4. I am also indebted to Mark Milam for building the Caltech ducted fan experiment. Additionally, I would like to acknowledge Mauricio Barahona, Yun Huang, Sanjay Lall, Pablo Parrilo, Marc Trotoux, Luz Vela and Yuji Yamada, with whom I have had many interesting discussions. My special thanks go to our department secretary,

Charmaine Boyd, for doing a wonderful job as our department's administrator.

Finally I would like to thank my parents, my brothers and my sister, for their unconditional love and support. Last but not least, I would like to thank my wonderful and loving wife Nikroo, to whom I have dedicated this thesis. Without her, none of this was possible.

Abstract

With the advent of faster and cheaper computers, optimization based control methodologies have become a viable candidate for control of nonlinear systems. Over the past twenty years, a group of such control schemes have been successfully used in the process control industry where the processes are either intrinsically stable or have very large time constants.

The purpose of this thesis is to provide a theoretical framework for synthesis of a class of optimization based control schemes, known as receding horizon control techniques for nonlinear systems such as unmanned aerial vehicles.

It is well known that unconstrained infinite horizon optimal control may be used to construct a stabilizing controller for a nonlinear system. In this thesis, we show that similar stabilization results may be achieved using unconstrained finite horizon optimal control. The key idea is to approximate the tail of the infinite horizon cost-to-go using, as terminal cost, an appropriate control Lyapunov function (CLF). A CLF can be thought of as generalization of the concept of a Lyapunov function to systems with inputs.

Roughly speaking, the terminal CLF should provide an (incremental) upper bound on the cost. In this fashion, important stability characteristics may be retained without the use of terminal constraints such as those employed by a number of other researchers. The absence of constraints allows a significant speedup in computation.

Furthermore, it is shown that in order to guarantee stability, it suffices to satisfy an improvement property, thereby relaxing the requirement that truly optimal trajectories be found.

We provide a complete analysis of the stability and region of attraction/operation properties of receding horizon control strategies that utilize finite horizon approximations in the proposed class. It is shown that the guaranteed region of operation contains that of the CLF controller and may be made as large as desired by increasing

the optimization horizon (restricted, of course, to the infinite horizon domain). Moreover, it is easily seen that both CLF and infinite horizon optimal control approaches are limiting cases of our receding horizon strategy. The key results are illustrated using a familiar example, the inverted pendulum, as well as models of the Caltech ducted fan at hover and forward flight, where significant improvements in guaranteed region of operation and cost are noted.

We also develop an optimization based scheme for generation of aggressive trajectories for hover and forward flight models of the Caltech ducted fan experiment, using a technique known as trajectory morphing. The main idea behind trajectory morphing is to develop a simplified model of the nonlinear system and solve the trajectory generation problem for that model. The resulting trajectory is then used as a reference in a receding horizon optimization scheme to generate trajectories of the original nonlinear system. Several aggressive trajectories are obtained in this fashion for the forward flight model of the Caltech ducted fan experiment.

Contents

List of Figures

List of Tables

Chapter 1 Introduction

1.1 Background

For quite some time now, optimization has played a crucial role in decision processes concerning physical or organizational systems. Every time there is a need to make a selection between a set of possible choices, one would like to pick the decision that costs the least, satisfies all constraints, and is practical to implement. Scarcity of resources often dictates the need to make decisions that distribute available resources in an optimum manner. Therefore, it is not surprising that *optimization* has become an integral part of any scientific and engineering discipline.

Despite the fact that the mathematical foundation of optimization theory goes way back to great mathematicians such as Gauss, Lagrange, Euler, the Bernoulis, etc., it was only 5 decades ago, with the advent of digital computers, that optimization became widespread and practical.

It was perhaps with the invention of Simplex method [Dan51], an efficient algorithm for solving linear programs, that both the theory and practice of optimization gained momentum. Linear programing was, and still is, the most natural way to formulate a variety of optimization problems. In addition to the elegance of the mathematical solution, perhaps the simplicity in the formulation phase has made linear programing a pure success.

Nevertheless, there are a huge class of optimization problems which can not be formulated as a linear program, for the simple reason that the constraints and/or the objective are not linear in variables. This brings us to *nonlinear* programs [GMW81]. Much of the early theory associated with this type of optimization has been focused on obtaining necessary and sufficient conditions for a solution point. Also, due to the non-convex nature of the problem, the solutions are at best local extrema, i.e., there is no guarantee, in general, that the nonlinear program is solved globally.

Although intrinsically more difficult than its linear counterpart, the past couple of decades has witnessed a tremendous amount of research in development of efficient algorithms for solving nonlinear programs. Several interior point methods have been developed for solution of quadratic programing problems which are a special class of nonlinear programs [Wri98]. General nonlinear programing problems can be solved by finding the solution to a sequence of quadratic programs using a method known as sequential quadratic programing [GMSW98].

Both linear and nonlinear programing, however, are static optimization problems. An interesting and useful class of optimization problems is dynamic optimization. In dynamic optimization, a functional is maximized/minimized subject to a set of differential and algebraic constraints. Such problems arise in planning, control, estimation, game theory, etc. This class of problems are essentially more difficult to solve, since the solutions are elements of *infinite dimensional* rather than *finite dimensional* vector spaces, as in the case of nonlinear programs. This simple difference introduces a whole new set of challenges to the problem. Namely, to solve the problem, one has to somehow convert the problem into a nonlinear programing problem, and then use off-the-shelf nonlinear programing solvers to solve the problem. In this thesis, our focus is on this type of optimization problems. Specifically, we focus our attention on the optimal control problem.

Although optimal control theory has its roots in calculus of variations and John Bernoulli's *brachistochrone* (Greek for "shortest time") problem, it was formally developed about 50 years ago in the seminal works of L.S. Pontryagin [Ponan] in the former Soviet Union and Richard Bellman [Bel52] in the United States. While Pontryagin introduced the *minimum principle*, which basically gave necessary conditions for the existence of optimal trajectories, Bellman introduced the concept of *dynamic programing*. The development of dynamic programing led to the notion of the celebrated *Hamilton-Jacobi Bellman (HJB)* partial differential equation, which had the *value function* as its solution. The value function provided the cost-to-go for any initial condition, resulting in an optimal *feedback* policy.

The idea behind dynamic programing was extremely simple and elegant. It was

mainly based on *principle of optimality* which merely suggested that if a trajectory is optimal, any end portion of it had to be optimal too. This resulted in a revolution in optimal control, and with pioneering work of Kalman [Kal60, KB61, Kal64] in 1960's, led to the theory of linear optimal control.

The decade 1960-1970 witnessed intensive research activities on the Linear Quadratic Gaussian (LQG) problem [Ath71], also known as the $\mathcal{H}_2$ optimal control problem. The HJB partial differential equation, prohibitive to solve in general, had now become two separate Riccati equations, which could be solved very efficiently.

At that point in time it was pretty much believed that the linear optimal control problem was solved. However, a simple counter example in [Doy78] proved otherwise. This simple counter example showed that LQG regulators can have arbitrarily small robustness margins. This result was quite contrary to the common belief at the time and ignited a new line of research that mixed control theory and pure mathematics for the most of 80's and led to the solution of the well-known $\mathcal{H}_\infty$ optimal control problem [DGKF89].

The story of the nonlinear optimal control, however, was unfortunately not as fruitful, for the very good reason that the problems were inherently more difficult in the nonlinear domain. The HJB equation was solved only in some low dimensional cases, and that was by brute force numerical methods that exploited gridding and meshing techniques.

Furthermore, the problem of finding an optimal state and control trajectory from a given initial condition was solved by solving a two point boundary value problem. This, however, was just a way to come up with the optimal control action, as a function of time and not as a function of state. Therefore, for obvious reasons, it was vulnerable to uncertainty, and not useful for control purposes. Besides, even in the linear case that the problem was much more mature, dealing with control and state constraints was not properly addressed in this context.

Several excellent textbooks emerged on the optimal control theory in 1960's and 1970's [AF66, LM67, BH75]. Among these books, however, [LM67] had an interesting paragraph that described a hypothetical method for obtaining a closed-loop

controller from open-loop trajectories [MRRS00]: "One technique for obtaining a feedback controller synthesis from knowledge of open-loop controllers is to measure the current control process state and then compute very rapidly for the open-loop control function. The first portion of this function is then used during a short time interval, after which a new measurement of the process state is made and a new open-loop control function is computed for this new measurement. The procedure is then repeated."

This very important remark was simply forgotten, probably due to the high computational cost of the algorithm. More than ten years later, this idea led to the Model Predictive Control (MPC), also known as Receding Horizon Control (RHC) technology.

Perhaps the earliest application of this idea was in Richalet, Rault, Testud and Papon [RRTP78]. These authors developed a technique known as IDCOM (IDentification and COMmand). Their method employed a finite horizon pulse response linear model, a quadratic cost function, and input and output constraints. It was, however, developed totally independent from the earlier results of optimal control, such as the one in [LM67], and was designed to address the needs and concerns of the industry, which required methods that can handle constraints and uncertainty. Several other methods were later developed along the same lines. Some of these were Dynamic Matrix Control (DMC) [CR80, PG79], and later Quadratic Dynamic Matrix control (QDMC) [GM86]. In QDMC, quadratic programing was employed to solve the constrained open-loop optimal control problem that results from having a linear system with linear constraints and quadratic cost in the optimization.

These methods were mainly used in the petro-chemical and process control industries. In these industries, the operating points were obtained by solving linear programs, and due to economic considerations, they were required to be on the boundary of feasibility. This made the use of optimization quite attractive.

Despite the fact that the industrial proponents of this approach did not address stability issues directly, they were well aware of its importance. In fact, an earlier observation by Kalman [Kal60] had suggested that even in the absence of constraints,

optimality did not necessarily imply stability, and that only under some conditions (stabilizability and detectability), and infinite length of the optimization horizon, optimal controllers were stabilizing. Later, standard Lyapunov stability techniques were used to establish stability and it turned out that the value function associated with the infinite horizon optimal control problem is a suitable Lyapunov function.

Since solving an infinite horizon optimal control problem in real time was a prohibitive task, except possibly for standard $\mathcal{H}_2$ and $\mathcal{H}_\infty$ problems, researchers came up with alternative schemes to use finite horizon approximations. According to Mayne *et al.* [MRRS00], early examples of results in this direction were Kleinman [Kle70], and Thomas [Tho75]. Kleinman considered the minimum energy control of a linear system over a finite horizon and picked the inverse of the controllability Grammian over the finite horizon as the Lyapunov matrix.

Further results by Kwon and Pearson [KP77] and Kwon, Brukstein, and Kailath [KBK83] solved a more general linear quadratic problem with a finite horizon. These results were important, but of limited application, because they were only appropriate for unconstrained linear systems.

Since then, MPC has become a standard technique for control of multivariable, constrained chemical processes. Over 2000 applications of this technology has made it a multimillion dollar industry [MRRS00, ML99, GPM89]. However, until recently, the attention of MPC proponents has been mainly restricted to open-loop stable processes or systems with large "settling times."

At first only linear systems were considered and the effect of horizon length and cost function parameters on stability and performance were analyzed. Since the employed machinery was linear analysis, hard constraints did not appear in these results.

Around the same time, however, a closely similar methodology was being developed in the adaptive control community and that was named Generalized Predictive Control (GPC) [CMT87a, CMT87b]. The main issue under consideration in GPC was stability of input-output linear discrete time systems where states were not available for measurement and noise was present. Stability was achieved by imposing

constraints on inputs and outputs over a finite interval.

The turning point in stability analysis of the MPC scheme is probably the paper by Chen and Shaw [CS82]. These authors show that the stability of the receding horizon scheme can be guaranteed by utilizing a terminal constraint and using the value function as a Lyapunov function candidate. Later, Keerthi and Gilbert [KG88] imposed a terminal state equality constraint $x(t+T) = 0$ and used the value function as the Lyapunov function to ensure the stability of constrained discrete-time nonlinear systems. Since then, the value function has been almost universally used as the Lyapunov function to ensure stability of the receding horizon scheme.

The 90's has witnessed a number of proposals in ensuring the stability of the MPC scheme. Some of these approaches follow and extend the last two references, and employ terminal equality constraints to guarantee stability. The work of Keerthi and Gilbert [KG88] resulted in a finite-horizon optimization problem which turned out to be computationally demanding and almost impossible to satisfy in some cases.

This result was later relaxed by Michalska and Mayne [MM93] which ensured closed-loop stability by requiring that the state at the end of the horizon to enter a suitable neighborhood of the origin rather than being at the origin. Once the state entered that neighborhood, the control was switched to a local linear controller which stabilized the system from that point. This approach was known as the dual mode control. A similar approach was also used in Sznaier and Damborg [SD90] in a different context.

Another important scheme which was developed by Bitmead *et al.* [BGW90] utilized a terminal cost to ensure closed-loop stability of unconstrained linear systems. There was no need to impose terminal constraints in this approach and the controller was computed off-line. In addition to these results, a very important result due to Rawlings and Muske [RM93] was developed in the context of linear stable systems with input constraints. They proposed the terminal cost to be the infinite horizon value function associated with zero control.

A different approach combined both of the aforementioned methods, namely, using a terminal constraint as well as a terminal cost. These were proposed by

Parisini and Zoppoli [PZ95] and later by Chen and Allgöwer [CA98] in the context of nonlinear systems. The idea was to use a quadratic endpoint penalty of the form $ax(t+T)^T Px(t+T)$ for some $a > 0$ and some positive definite matrix P, as well as a terminal inequality constraint.

In a more recent paper by Magni and Sepulchre [MS97] and later by De Nicolao *et al.* [NMS98], stability of the receding horizon scheme was proven (for continuous-time and discrete-time systems, respectively) by using a (possibly non-quadratic) end point penalty which is the cost incurred if a locally stabilizing linear control law is applied at the end of the time horizon T. The linear control law ensures local exponential stability of the equilibrium at $x = 0$, and it is assumed that the region of attraction of the linear controller is large enough that can be reached from the initial condition within the time interval $[0, T]$. Moreover, it is assumed that the optimization is performed over *admissible* control sequences, i.e., control sequences which guarantee that at the end of the horizon the state has reached a suitable neighborhood of the origin which is an exponential stability region for the linear controller. In other words, a state inequality constraint is implicitly imposed.

An approach for the receding horizon control of globally stabilizable nonlinear systems was developed by Primbs *et al.* [PND00, Pri98]. In this approach, first a globally stabilizing control law is achieved by finding a global control Lyapunov function (CLF) [Son89]. As it will be explained later in the thesis, a CLF is generalization of the Lyapunov function to systems with input. Once the global CLF is obtained, closed-loop stability is enforced by including additional state constraints that require the derivative of the CLF along the receding horizon trajectory to be negative and also that the decrease in the value of the CLF be greater than that obtained using the controller derived from the CLF. This approach is attractive in the sense that global stability is retained without, e.g., enforcing an equality endpoint constraint. (Unfortunately, it does require a global CLF.) This approach divides the problem into two phases. The first phase is to obtain a CLF and the second is to use the CLF in a receding horizon scheme.

Since the couple of lines of introduction in [LM67], MPC has come a long way.

However, there are several issues which have not been addressed properly.

On the theory side, despite the variety of approaches in stability analysis of the closed-loop system, there are a few problems still unsolved. Most importantly, the issue of stability is not completely resolved yet. Other issues such as regions of attractions of MPC or RHC have not been fully addressed either.

On the practical side, although various success stories have been reported from the process control industry, such as plants with as many as 50 states and multiple inputs and outputs, it has been mainly confined to slow or stable processes. Specifically no successful application of the MPC or RHC scheme or any of its variants which can be coined under the term *optimization based control*, has been reported in flight control. Part of the problem stems from the fact that control theory has been very developed in terms of regulation whereas the area of trajectory generation and tracking is not developed as much, and performing very exotic maneuvers can not be addressed properly in a regulation setting.

As it will be seen in the next section, the purpose of this thesis is to develop a mathematical framework for the MPC-RHC scheme that can be employed for the control of highly maneuverable aerial vehicles.

1.2 Contents of this thesis

In the previous section, we presented a brief overview of the status of model predictive control/receding horizon control methods, also known as optimization based control approaches. In the next few chapters, we develop a stabilizing scheme for receding horizon control of nonlinear systems. The proposed scheme utilizes a special class of control Lyapunov functions (CLFs), which can be obtained off-line, and used as a terminal cost in the receding horizon scheme to guarantee stability.

Specifically, we show that contrary to some earlier results described in the previous section, there is no need to impose terminal equality and /or inequality constraints or any other CLF based constraint to guarantee stability of the receding horizon scheme. This, as will be seen later throughout the thesis, will speed up the calculations

drastically.

This research has been performed at Caltech as part of the Software Enabled Control (SEC) at DARPA. The idea behind SEC program is to explore challenges that one faces when trying to use optimization-based control strategies for control of Unmanned Aerial Vehicles (UAVs). Until recently, use of RHC in control of UAVs seemed impossible, due to heavy cost of online computations required for stability and high performance of the vehicle. Thanks to faster and cheaper hardware, and more efficient software, there seems to be at last some light at the end of the tunnel. Proper use of RHC technology, in addition to proper software, undoubtedly makes this once thought farfetched goal as close as ever to reality.

The main contribution of this thesis is the development of an RHC strategy which is suitable for control of high performance Unmanned Aerial Vehicles (UAVs). In Chapter 2, we show that stabilization results may be achieved using an *unconstrained* finite horizon optimal control problem. The key idea is to approximate the tail of the infinite horizon cost-to-go using, as terminal cost, an appropriate control Lyapunov function. Roughly speaking, the terminal CLF should provide an (incremental) upper bound on the cost. In this fashion, important stability characteristics may be retained *without* the use of terminal constraints such as those employed by a number of other researchers mentioned in 1.1. The absence of constraints allows a significant speedup in computation.

We provide a complete analysis of the stability and region of attraction/operation properties of receding horizon control strategies that utilize finite horizon approximations in the proposed class. It is shown that the guaranteed region of operation contains that of the CLF controller and may be made as large as desired by increasing the optimization horizon (restricted, of course, to the infinite horizon domain). Moreover, it is easily seen that both CLF and infinite horizon optimal control approaches are limiting cases of our receding horizon strategy. The key results are illustrated using a familiar example, the inverted pendulum, where significant improvements in guaranteed region of operation and cost are noted.

A requirement implied in these results was being able to solve the optimizations

globally. This assumption is removed in Chapter 3 and it is shown that the optimality can be replaced by an *improvement* property. Specifically, it is shown that in order to guarantee stability, it suffices to satisfy an improvement property, thereby relaxing the requirement that truly *optimal* trajectories be found. A numerical example using the same inverted pendulum is presented to illustrate this point.

Furthermore, several approaches for obtaining CLF suitable for receding horizon purposes will be discussed in Chapter 4. The methods used range from Jacobian linearization of the nonlinear plant and designing an LQR controller, to using convex optimization techniques to come up with quadratic CLFs and finally Linear Parameter Varying (LPV) methods for quadratic and state dependent control Lyapunov functions.

The practical portion of the thesis deals with the application of receding horizon/CLF methods developed earlier in this thesis to models of a flight control experiment developed at Caltech. The experiment is a tethered flying wing with a fan and a duct, hence, it is called "ducted fan." The dynamics of the system are representative of a Vertical Landing and Take off (VTOL) aircraft, such as a Harrier around hover or a thrust vectored aircraft such as F18-HARV or X-31 in forward flight. Therefore, it serves as a perfect testbed for the purpose of this thesis. This section starts from Chapter 5. In this chapter, the Caltech ducted fan experiment is completely described and two models of the experiment, one in hover mode and one in forward flight, are developed and discussed. Several simulations are performed and comparisons are made between different results.

All the chapters so far deal with the problem of regulation. We briefly touch on the subject of trajectory generation in Chapter 6. Since trajectory generation is an inherently difficult problem, a simplified model of the ducted fan is developed for trajectory generation purposes. The generated trajectories are used as reference in a receding horizon scheme and trajectories for the full model are generated in this fashion. Finally, conclusions are presented in Chapter 7.

Chapter 2 Receding horizon control

2.1 Problem setting

The nonlinear system under consideration is

$$\dot{x} = f(x, u)$$

where the vector field $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ is C^2 and possesses a linearly controllable critical point at the origin, e.g., $f(0,0) = 0$ and $(A, B) := (D_1 f(0,0), D_2 f(0,0))$ is controllable. We require the set $f(x, \mathbb{R}^m) \subset \mathbb{R}^n$ to be convex for each $x \in \mathbb{R}^n$. Given an initial state x and a control trajectory $u(\cdot)$, the state trajectory $x^u(\cdot; x)$ is the (absolutely continuous) curve in $\mathbb{R}^n$ satisfying

$$x^u(t; x) = x + \int_0^t f(x^u(\tau; x), u(\tau))\, d\tau$$

for $t \geq 0$. We require that the trajectories of the system satisfy an *a priori* bound

$$\|x(t)\| \leq \beta(x, T, \|u(\cdot)\|_1) < \infty, \qquad t \in [0, T],$$

where β is continuous in all variables and monotone increasing in T and $\|u(\cdot)\|_1 = \|u(\cdot)\|_{L_1(0,T)}$. Most models of physical systems will satisfy a bound of this type.

The performance of the system will be measured by a given incremental cost $q : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ that is C^2 and fully penalizes both state and control according to

$$q(x, u) \geq c_q(\|x\|^2 + \|u\|^2), \qquad x \in \mathbb{R}^n, u \in \mathbb{R}^m$$

for some $c_q > 0$ and $q(0,0) = 0$. We further require that the function $u \mapsto q(x, u)$ be convex for each $x \in \mathbb{R}^n$. These conditions imply that the quadratic approximation of

q at the origin is positive definite, $D^2q(0,0) \geq c_q I > 0$.

We will also suppose that f and q are sufficiently compatible to uniquely define a C^2 Hamiltonian for the (optimized) system [LM67]. In particular, we will require that there is a C^2 function $\bar{u}^* : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^m : (x,p) \mapsto \bar{u}^*(x,p)$ providing a global minimum of the pre-Hamiltonian $K(x,p,u) := p^T f(x,u) + q(x,u)$ so that the Hamiltonian $H(x,p) := K(x,p,\bar{u}^*(x,p))$ is C^2. Such a $\bar{u}^*(\cdot,\cdot)$ is locally guaranteed by the implicit function theorem (though we would require $f, q \in C^3$). Note that this condition is trivially satisfied for control affine f and quadratic q (for then $u \mapsto K(x,p,u)$ is strictly convex and *real analytic*).

The cost of applying a control $u(\cdot)$ from an initial state x over the infinite time interval $[0,\infty)$ is given by

$$J_\infty(x, u(\cdot)) = \int_0^\infty q(x^u(\tau; x), u(\tau))\, d\tau \ .$$

The optimal cost (from x) is given by

$$J^*_\infty(x) = \inf_{u(\cdot)} J_\infty(x, u(\cdot))$$

where the control functions $u(\cdot)$ belong to some reasonable class of admissible controls (e.g., piecewise continuous). The function $x \mapsto J^*_\infty(x)$ is often called the *optimal value function* for the infinite horizon optimal control problem. For the class of f and q considered, we know that J^*_∞ is a positive definite C^2 function on a neighborhood of the origin. This follows from the *geometry* of the corresponding Hamiltonian system [vdS91, vdS94]. In particular, since $(x,p) = (0,0)$ is a hyperbolic critical point of the Hamiltonian vector field $X_H(x,p) := (D_2H(x,p), -D_1H(x,p))^T$, the local properties of J^*_∞ are determined by the linear-quadratic approximation to the problem and, moreover, $D^2J^*_\infty(0) = P > 0$ where P is the stabilizing solution of the appropriate algebraic Riccati equation.

For practical purposes, we are interested in approximating the infinite horizon optimization problem with one over a finite horizon. In particular, we would like to

somehow take the discarded tail of the cost into account. To do this, let V be a nonnegative C^2 function and define the finite horizon cost (from x using $u(\cdot)$) to be

$$J_T(x, u(\cdot)) = \int_0^T q(x^u(\tau; x), u(\tau))\, d\tau + V(x^u(T; x))$$

and denote the optimal cost (from x) as

$$J_T^*(x) = \inf_{u(\cdot)} J_T(x, u(\cdot)) \ .$$

As in the infinite horizon case, one can show, by geometric means, that J_T^* is locally smooth (C^2). Other properties, e.g., local positive definiteness, will depend on the choice of V and T. As we will see in the sequel, certain choices of V will allow us to keep the desirable features of the infinite horizon problem.

Let Γ^∞ denote the domain of J_∞^* (the subset of $\mathbb{R}^n$ on which J_∞^* is finite). It is not too difficult to show that the cost functions J_∞^* and J_T^*, $T \geq 0$ are continuous functions on Γ_∞ using the same arguments as in proposition 3.1 of [BCD97]. We make the following assumption.

Standing Assumption (SA): *The minimum value of cost functions J_∞^*, J_T^*, $T \geq 0$, is attained.*

The assumption (SA) guarantees the existence of an optimal trajectory

$$(x_T^*(t; x), u_T^*(t; x)), \qquad t \in [0, T],$$

such that

$$J_T(x, u_T^*(\cdot; x)) = J_T^*(x) \ .$$

Continuity of $u_T^*(.; x)$ follows directly from Pontryagin's Maximum Principle. This trajectory is not necessarily unique. In fact, in examples one finds two trajectories of equal (minimal) cost originating at points where J_T^* is only continuous (and not differentiable). Under assumptions of the sort given (convexity, boundedness, etc.), one can turn (SA) into a proposition. This involves the use of techniques from regularity

theory and the direct methods of the calculus of variations; see [Ces83] and [BMH98].

It is easy to see that J^*_∞ is proper on its domain so that the sub-level sets

$$\Gamma^\infty_r := \{x \in \Gamma^\infty : J^*_\infty(x) \leq r^2\}$$

are compact and path connected and moreover $\Gamma^\infty = \bigcup_{r\geq 0} \Gamma^\infty_r$. Note also that Γ^∞ may be a proper subset of $\mathbb{R}^n$ since there may be states that can not be driven to the origin. We use r^2 (rather than r) here to reflect the fact that our incremental cost is quadratically bounded from below. We refer to sub-level sets of J^*_T and V using

$$\Gamma^T_r := \text{path connected component of } \{x \in \Gamma^\infty : J^*_T(x) \leq r^2\} \text{ containing } 0,$$

and

$$\Omega_r := \text{path connected component of } \{x \in \mathbb{R}^n : V(x) \leq r^2\} \text{ containing } 0.$$

2.2 Infinite and finite horizon optimization

In this section, we explore some of the relationships between an infinite horizon optimal control problem and its finite horizon approximations. We will show that the use of an appropriate terminal cost allows us to retain desirable features of the infinite horizon problem.

2.2.1 Infinite horizon properties

What infinite horizon problem properties are interesting for finite horizon approximations and, in particular, are useful for receding horizon strategies? This is a question that we intend to answer in this section.

Let $(x^*_\infty, u^*_\infty)(\cdot; x)$ be *any* optimal trajectory originating at x. Then, for any $\delta > 0$, we have

$$J^*_\infty(x^*_\infty(\delta; x)) = J^*_\infty(x) - \int_0^\delta q(x^*_\infty(\tau; x), u^*_\infty(\tau; x))\, d\tau\ . \tag{2.1}$$

Since, by (SA), $\tau \mapsto u^*_\infty(\tau; x)$ is continuous for $\tau \geq 0$, we see that

$$\lim_{\delta \searrow 0} \frac{J^*_\infty(x^*_\infty(\delta; x)) - J^*_\infty(x)}{\delta} = -q(x, u^*_\infty(0; x)) \leq -c_q \|x\|^2$$

so that J^*_∞ possesses a well defined (negative definite) directional derivative in *each* optimal direction $f(x, u^*_\infty(0; x))$. In fact, we may write

$$\dot{J}^*_\infty(x, u^*_\infty(0; x)) + q(x, u^*_\infty(0; x)) = 0 \tag{2.2}$$

where $\dot{J}^*_\infty(x, u)$ is the directional derivative of J^*_∞ in the direction $f(x, u)$ (when it exists). (At points of differentiability, $\dot{J}^*_\infty(x, u) = DJ^*_\infty(x) \cdot f(x, u)$.)

We conclude that each sub-level set Γ^∞_r, $r > 0$, is positively invariant under optimal actions, both incremental ($\delta > 0$) and infinitesimal. Also, in all cases, these sets are attracted to the origin exponentially fast. In particular, the (not necessarily unique) feedback $u = k_\infty(x) := u^*_\infty(0; x)$ exponentially stabilizes the origin.

2.2.2 Finite horizon properties

As noted above, one may use optimal (infinite horizon) actions to provide a stabilizing feedback for a nonlinear system. It is natural to expect that a similar result would be possible using a finite horizon optimization. For instance, one could implement a *receding horizon* scheme as follows. From the current state $x(t)$, obtain an optimal trajectory $(x^*_T, u^*_T)(\tau; x(t))$, $\tau \in [0, T]$, and use as feedback $u(t) = u^*_T(0; x(t))$. (This feedback is not uniquely defined at points where more than one optimal trajectory is available.) This approach requires one to continuously re-solve the finite horizon optimization. An alternative scheme is to solve the finite horizon optimization every $\delta > 0$ seconds and use the control trajectory $u^*_T(\tau; x(t))$, $\tau \in [0, \delta]$, to drive the system from $x(t)$ at time t to $x^*_T(\delta, x(t))$ at time $t + \delta$. (Practically speaking, a better idea is to use a local tracking controller to regulate the system about the desired trajectory $(x^*_T, u^*_T)(\tau; x(t))$, $\tau \in [0, \delta]$.) We will denote this receding horizon scheme as $\mathcal{RH}(T, \delta)$. One might also consider using a variable δ_k, which will be denoted

as $\mathcal{RH}(T, \{\delta_k\})$. Note that the receding horizon strategy defined a (sampled data) *feedback* law in contrast with the one shot use of an open loop optimal trajectory. In defining (unconstrained) finite horizon approximations to the infinite horizon problem, the key design parameters are the terminal cost function V and the horizon length T (and, perhaps also, the increment δ). What choices will result in success?

It is well known (and easily demonstrated with linear examples), that simple truncation of the integral (i.e., $V \equiv 0$) may have disastrous effects if $T > 0$ is too small. Indeed, although the resulting value function may be nicely behaved, the "optimal" receding horizon closed loop system can be unstable!

A more considered approach is to make good use of a suitable terminal cost V. Evidently, the best choice for the terminal cost is $V(x) = J^*_\infty(x)$ since then the optimal finite and infinite horizon costs are the same. Of course, if *the* optimal value function were available there would be no need to solve a trajectory optimization problem. What properties of the optimal value function should be retained in the terminal cost? To be effective, the terminal cost must account for the discarded tail by ensuring that the origin can be reached from the terminal state $x^u(T; x)$ in an efficient manner (as measured by q). One way to do this is to use an appropriate control Lyapunov function (CLF).

A *control Lyapunov function* (CLF) is a C^1, proper, positive definite function $V : \mathbb{R}^n \to \mathbb{R}_+$ such that:

$$\inf_u \left[\dot{V}(x, u)\right] \leq 0 \tag{2.3}$$

If it is possible to make the derivative negative at every point by an appropriate choice of u, then we have achieved our goal and can stabilize the system with V a Lyapunov function for the closed loop that we choose. This is exactly the condition given in (2.3). It can be shown that the existence of a CLF is equivalent to the existence of an asymptotically stabilizing control law $u = k(x)$ which is smooth everywhere except possibly at $x = 0$ [Art83]. Moreover, one can calculate such a control law k when the system is affine in control, i.e., $f(x, u) = \tilde{f}(x) + g(x)u$, explicitly from $\tilde{f}$, g and V

[Son89]. To this end, suppose that V is a proper C^2 function satisfying $V(0) = 0$,

$$V(x) \geq c_v \|x\|^2, \qquad x \in \mathbb{R}^n,$$

and that is compatible with the incremental cost in the sense that

$$\min_u (\dot{V} + q)(x, u) \leq 0 \tag{2.4}$$

on a neighborhood of $x = 0$. Here $\dot{V}(x, u) := DV(x) \cdot f(x, u)$.

Indeed, the feedback

$$u = k_V(x) := \bar{u}^*(x, DV(x)^T) \tag{2.5}$$

stabilizes the origin. Note that V can be thought of as a control Lyapunov function which is also an upper bound on the cost-to-go. The maximum principle ensures that $V = J^*_\infty$ also satisfies (2.4) according to (2.2). Continuity and properness of V guarantee the existence of a continuous nondecreasing function $r \mapsto \bar{c}_v(r)$ such that $V(x) \leq \bar{c}_v(r)\|x\|^2$ for all $x \in \Omega_r$ so that $x \notin \Omega_{r_0}$ implies that $\|x\|^2 \geq r_0^2/\bar{c}_v(r_0)$. Also, let $r_v > 0$ be the largest r such that (2.4) is satisfied for all $x \in \Omega_r$.

The following result provides a basis for the use of finite horizon optimization in a receding horizon control strategy (cf. [JYH99b]).

Theorem 2.2.1 *Suppose that $x \in \mathbb{R}^n$ and $T > 0$ are such that*

$$x^*_T(T; x) \in \Omega_{r_v} \ . \tag{2.6}$$

Then, for each $\delta \in [0, T]$, the optimal cost from $x^*_T(\delta; x)$ satisfies

$$J^*_T(x^*_T(\delta; x)) \leq J^*_T(x) - \int_0^\delta q(x^*_T(\tau; x), u^*_T(\tau; x))\, d\tau \ . \tag{2.7}$$

Note that $(x^*_T, u^*_T)(\cdot; x)$ can be *any* optimal trajectory for the problem with horizon T.

Proof: Let $(\tilde{x}(t), \tilde{u}(t))$, $t \in [0, 2T]$, be the trajectory obtained by concatenating $(x_T^*, u_T^*)(t; x)$, $t \in [0, T]$, and $(x^k, u^k)(t-T; x_T^*(T; x))$, $t \in [T, 2T]$. Here, $(x^k, u^k)(s; x_0)$ is the closed loop trajectory starting from x_0 at time $s = 0$,

$$x^k(s; x_0) = x_0 + \int_0^s f(x^k(\tau; x_0), k(x^k(\tau; x_0))) \, d\tau \ ,$$

where $u = k(x)$ is any feedback law such that $(\dot{V} + q)(x, k(x)) \leq 0$ for $x \in \Omega_{r_v}$, e.g., that defined by (2.5). Consider now the cost of using $\tilde{u}(\cdot)$ for T seconds beginning at an initial state $x_T^*(\delta; x)$, $\delta \in [0, T]$; see Figure 2.1. We have

$$\begin{aligned} J_T(x_T^*(\delta; x), \tilde{u}(\cdot)) &= \int_\delta^{T+\delta} q(\tilde{x}(\tau), \tilde{u}(\tau)) \, d\tau + V(\tilde{x}(T+\delta)) \\ &= J_T^*(x) - \int_0^\delta q(x_T^*(\tau; x), u_T^*(\tau; x)) \, d\tau - V(x_T^*(T; x)) \\ &\quad + \int_T^{T+\delta} q(\tilde{x}(\tau), \tilde{u}(\tau)) \, d\tau + V(\tilde{x}(T+\delta)) \\ &\leq J_T^*(x) - \int_0^\delta q(x_T^*(\tau; x), u_T^*(\tau; x)) \, d\tau \end{aligned}$$

where we have used the fact that $q(\tilde{x}(\tau), \tilde{u}(\tau)) \leq -\dot{V}(\tilde{x}(\tau), \tilde{u}(\tau))$ for all $\tau \in [T, 2T]$. The result follows since the optimal cost satisfies $J_T^*(x_T^*(\delta; x)) \leq J_T(x_T^*(\delta; x), \tilde{u}(\cdot))$. □

The following corollary [JYH99b, NMS98, MS97] easily follows:

Corollary 2.2.1 *Suppose the CLF is replaced by the infinite horizon cost-to-go resulting from the application of an a priori obtained stabilizing controller, such as the one defined in (2.5). Then (2.7) still holds.*

Proof: The proof easily follows that of Theorem 2.2.1. □

At this point, one is tempted to conclude that our approach to approximating the infinite horizon problem using a CLF terminal cost has been successful. After all, (2.7) is an appropriate approximation to (2.1) for invariance purposes. In fact, Theorem 2.2.1 is sufficient to conclude the desired invariance and attractiveness properties

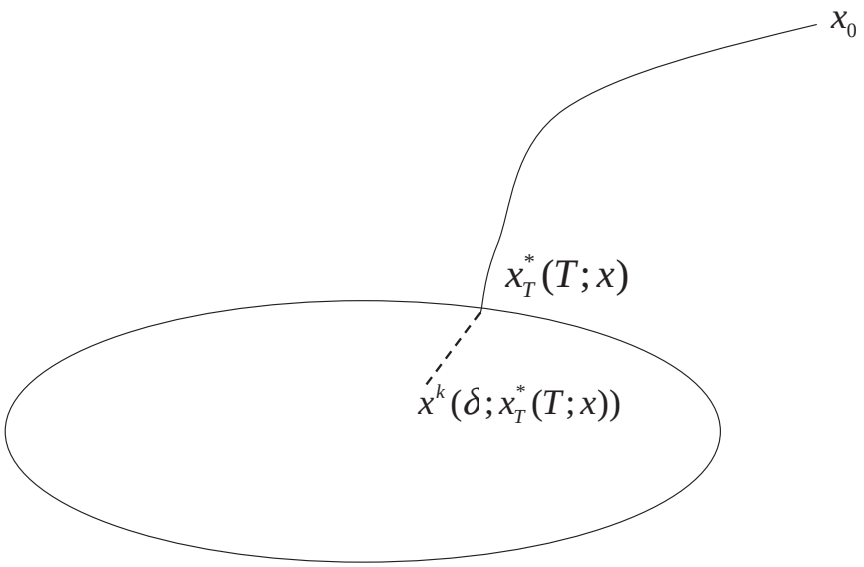

Figure 2.1: CLF sub-level set Ω_{r_v} as well as $x_T^*(\cdot)$ and $x^k(\cdot)$.

in the case that V is a *global* CLF for then that pesky "if" condition (2.6) will be trivially satisfied.

The situation when V is but a local CLF is much more delicate. Indeed, we must determine conditions under which (2.7) will hold under iteration of the receding horizon map, i.e., whether $x_T^*(T; x_T^*(\delta; x)) \in \Omega_{r_v}$ holds. One way to ensure success is to solve a *constrained* optimization that imposes such a condition; see, e.g., [NMS98, MM93]. We will show that such an approach is unnecessary.

We begin with a surprising lemma that helps us control the behavior of the terminal state of optimal trajectories.

Lemma 2.2.1 *Suppose that $x \in \Omega_r$, $r \leq r_v$. Then $x_T^*(T; x) \in \Omega_r$ for every $T \geq 0$.*

Proof: As before, let $(x^k, u^k)(t; x)$, $t \geq 0$, be the trajectory (starting at x) obtained using a feedback control $u = k(x)$ satisfying $(\dot{V} + q)(x, k(x)) \leq 0$ on Ω_{r_v}. The optimal cost with horizon $T \geq 0$ satisfies

$$\begin{aligned} J_T^*(x) &\leq \int_0^T q(x^k(\tau; x), u^k(\tau; x))\, d\tau + V(x^k(T; x)) \\ &\leq \int_0^T -\dot{V}(x^k(\tau; x), u^k(\tau; x))\, d\tau + V(x^k(T; x)) \\ &= V(x) \leq r^2 \ . \end{aligned}$$

Thus,

$$\begin{aligned} V(x_T^*(T; x)) &= J_T^*(x) - \int_0^T q(x_T^*(\tau; x), u_T^*(\tau; x))\, d\tau \\ &\leq J_T^*(x) \leq V(x) \leq r^2 \ . \end{aligned}$$

□

Note that Lemma 2.2.1 *does not* say that $x_T^*(t; x) \in \Omega_{r_v}$ for all $t \in [0, T]$ when $x \in \Omega_{r_v}$. This is false in general as simple examples show. Indeed, one might say that methods that attempt to maintain the invariance of Ω_r, $r \leq r_v$, are inefficient. (Moreover, adding constraints of that sort also drive up the computation cost.)

Figure 2.2 depicts the situation for a linear system with quadratic incremental cost. The (smaller) elliptical region is an invariant (sub-level) set for the chosen ter-

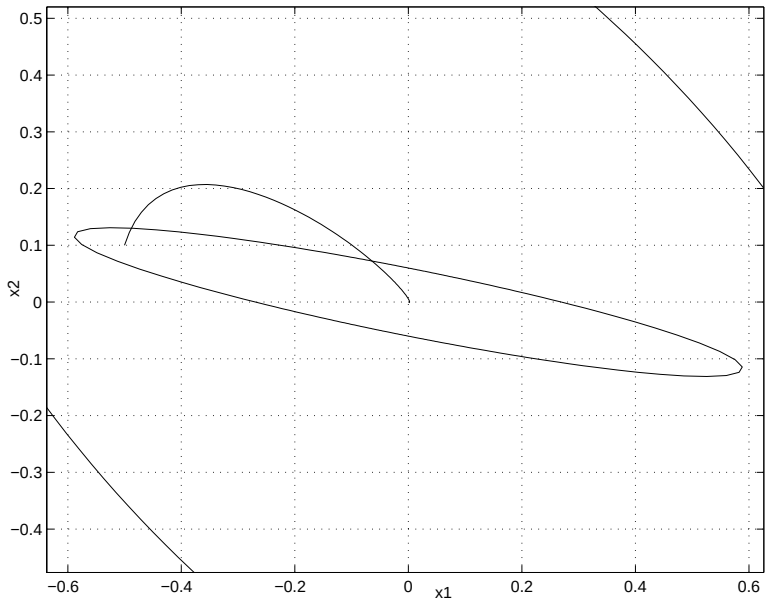

Figure 2.2: Optimal trajectories that enter Ω_{r_v} need not stay in Ω_{r_v}.

minal CLF (and corresponding controller). The trajectory originating at $(-0.5, 0.1)$ is optimal for $T > 5$. Note that sub-level sets of the terminal CLF are not invariant under optimal (incremental) flow. The shape of the sets that are invariant (sub-level set of J_T^*) is indicated by the larger (truncated) ellipse.

The elliptical region represents an invariant set defined by a quadratic Lyapunov function, whereas the larger set is the region of attraction of the infinite horizon control law. As it can be seen in this figure, the trajectories starting in Ω_{r_v} need not stay there for all $t \in [0, T]$.

One might say that methods that attempt to maintain the invariance of Ω_{r_v} are inefficient Since the available methods either require a terminal equality or inequality constraints which further add to the computational burden.

A key motivation for using on-line optimization is to *enlarge* the operating region for a controller. We are now in a position to show that the receding horizon controller does at least as good a job as the CLF controller, from the point of view of theoretical operating region predictions.

Proposition 2.2.1 *For all $T \geq 0$, $x \in \Gamma_{r_v}^T$ implies that $x_T^*(T; x) \in \Omega_{r_v}$. Moreover,*

$\Omega_{r_v} \subset \Gamma_{r_v}^T$ *for all* $T \geq 0$.

Proof: Let $T \geq 0$ and $x \in \Gamma_{r_v}^T$ and note that

$$V(x_T^*(T;x)) \;\leq\; r_v^2 - \int_0^T q(x_T^*(\tau;x), u_T^*(\tau;x))\, d\tau \;\leq\; r_v^2 \;.$$

The second statement was proved in the proof of Lemma 2.2.1. Figure 2.3 illustrates the set inclusions. □

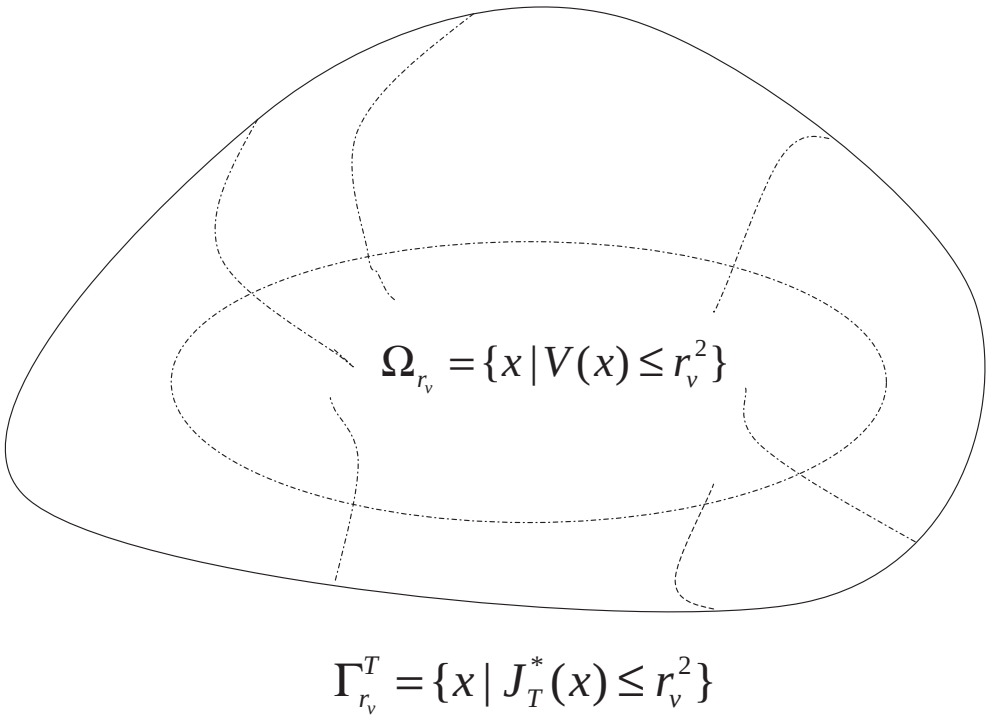

Figure 2.3: CLF sub-level set Ω_{r_v} and $\Gamma_{r_v}^T$. The optimal trajectories starting on $\Gamma_{r_v}^T$ end up inside Ω_{r_v} after T seconds.

Recall, however, that a key motivation for using on-line optimization is to *enlarge* the operating region for a controller. Moreover, in many cases, the best that one can hope for is a local CLF.

The following result provides a performance guarantee for our receding horizon control strategies.

Proposition 2.2.2 *Suppose that $T, r > 0$ are such that $x_T^*(T;x) \in \Omega_{r_v}$ for all $x \in \Gamma_r^T$. Let $x_0 \in \Gamma_r^T$ and consider a trajectory $(x_{\rm rh}(t), u_{\rm rh}(t))$, $t \geq 0$, resulting from the use of a receding horizon strategy $\mathcal{RH}(T, \{\delta_k\})$ (with $\delta_k > 0, \sum_{k=0}^{l} \delta_k \to \infty$ as $l \to \infty$).*

$$J_\infty(x_0, u_{\rm rh}(\cdot)) \leq J_T^*(x_0) \ .$$

Proof: Consider, at first, the use of $\mathcal{RH}(T, \delta)$ with constant $\delta > 0$. The receding horizon strategy defines a sequence of points $\{x_k\}_{k=0}^{\infty}$ according to

$$x_{k+1} = x_T^*(\delta, x_k)$$

starting with x_0 so that $x_k = x(k\delta)$. Now, by the principle of optimality, the cost of the arc from x_k to x_{k+1} is given by

$$\int_{k\delta}^{(k+1)\delta} q(x_{\rm rh}(\tau), u_{\rm rh}(\tau))\, d\tau = J_T^*(x_k) - J_{T-\delta}^*(x_{k+1}) \ .$$

Hence, the total cost of this strategy is

$$\begin{aligned} J_\infty(x_0, u(\cdot)) &= J_T^*(x_0) - J_{T-\delta}^*(x_1) + J_T^*(x_1) - J_{T-\delta}^*(x_2) + \cdots \\ &= J_T^*(x_0) + \sum_{k=1}^{\infty} \left\{ J_T^*(x_k) - J_{T-\delta}^*(x_k) \right\} \\ &\leq J_T^*(x_0) \end{aligned}$$

where the final inequality follows from the fact (shown in Proposition 2.2.3) that $J_T^*(x_k) \leq J_{T-\delta}^*(x_k)$ for all $\delta \geq 0$ and all $k \geq 0$. Clearly this result does not require $\delta > 0$ to be constant but merely that $\sum_{k=0}^{l} \delta_k \to \infty$ as $l \to \infty$.

The case of receding horizon control with continuous update follows by a limiting argument.

□

The above proposition generalizes the fact that

$$(\dot{V}+q)(x(t),u(t)) \le 0,\ t \ge 0 \quad \Longrightarrow \quad \int_0^\infty q(x(\tau),u(\tau))\,d\tau \le V(x(0))$$

when V is positive definite (implying $x(t) \to 0$). In both cases, we obtain an upper bound on the cost for a *family* of trajectories. We also point out that the cost of using a receding horizon control strategy approaches the infinite horizon cost as the horizon T is increased since $J_\infty^*(x_0) \le J_\infty(x_0, u_{\rm rh}(\cdot)) \le J_T^*(x_0)$ and $J_T^*(x_0) \to J_\infty^*(x_0)$ as $T \to \infty$.

We are now ready to show that application of the receding horizon strategy results in the exponential convergence of the trajectory to the origin:

Theorem 2.2.2 *Let $T > 0$ and consider the use of a receding horizon scheme $\mathcal{RH}(T,\{\delta_k\})$ with each $\delta_k \in (0,T]$ and $\sum_{j=0}^k \delta_j \to \infty$ as $k \to \infty$. Then, for each $x_0 \in \Gamma^T_{\bar{r}(T)}$, the resulting trajectory converges to the origin exponentially fast, where $\bar{r}(T)$ is the largest radius such that for each $x_0 \in \Gamma^T_{\bar{r}(T)}$, $x_T^*(T;x_0) \in \Omega_{r_v}$.*

Proof: Given $T > 0$, set $c_{\rm T}$ and m_∞ such that $J_T^*(x) \le c_{\rm T}||x||^2 \quad \forall x \in \Gamma^T_{\bar{r}(T)}$ and $J_\infty^*(x) \ge m_\infty||x||^2 \quad \forall x \in \Gamma^T_{\bar{r}(T)}$. Let $(x_{\rm RH}(t), u_{\rm RH}(t))$, $t \ge 0$, be the receding horizon trajectory originating from an arbitrary $x_0 \in \Gamma^T_{\bar{r}(T)}$ and define

$$W(t;x_0,u_{\rm RH}(\cdot)) := \int_t^\infty q(x_{\rm RH}(\tau),u_{\rm RH}(\tau))d\tau\,.$$

(The control trajectory $u_{\rm RH}(\cdot)$ is piecewise continuous since each optimal control trajectory $u_T^*(t;x)$, $t \ge 0$, is continuous as a function of time.) As shown in Proposition 2.2.2, $W(t;x_0,u_{\rm RH}(\cdot)) \le J_T^*(x_{\rm RH}(t))$. Also, since $\mathcal{RH}(T,\{\delta_k\})$ is a suboptimal strategy, $J_\infty^*(x_{\rm RH}(t)) \le W(t;x_0,u_{\rm RH}(\cdot))$. Now, since $q(x_{\rm RH}(t),u_{\rm RH}(t)) \to 0$ as $t \to \infty$, we have

$$\begin{aligned}\frac{\partial}{\partial t}W(t;x_0,u_{\rm RH}(\cdot)) &= -q(x_{\rm RH}(t),u_{\rm RH}(t)) \\ &\le -c_q||x_{\rm RH}(t)||^2\end{aligned}$$

$$\begin{aligned} &\leq -\frac{c_q}{c_{\mathrm{T}}} J_T^*(x_{\mathrm{RH}}(t)) \\ &\leq -\frac{c_q}{c_{\mathrm{T}}} W(t; x_0, u_{\mathrm{RH}}(\cdot)) \end{aligned}$$

so that $W(t; x_0, u_{\mathrm{RH}}(\cdot)) \leq e^{-\frac{c_q}{c_{\mathrm{T}}} t} W(0; x_0, u_{\mathrm{RH}}(\cdot))$. The result follows since

$$\begin{aligned} m_\infty ||x_{\mathrm{RH}}(t)||^2 &\leq J_\infty^*(x_{\mathrm{RH}}(t)) \leq W(t; x_0, u_{\mathrm{RH}}(\cdot)) \\ &\leq e^{-\frac{c_q}{c_{\mathrm{T}}} t} W(0; x_0, u_{\mathrm{RH}}(\cdot)) \leq e^{-\frac{c_q}{c_{\mathrm{T}}} t} J_T^*(x_0) \\ &\leq c_{\mathrm{T}} e^{-\frac{c_q}{c_{\mathrm{T}}} t} ||x_0||^2 . \end{aligned}$$

□

Note that the optimal control $u_T^*(\cdot; x)$ is uniquely defined in a neighborhood of the origin since $J_T^*(\cdot)$ is locally C^2 so that the locally defined *instantaneous* receding horizon control $u = k_T(x) := u_T^*(0; x)$ (i.e., $\delta = 0$) defines a feedback providing local exponential stability of the origin. Indeed, the resulting feedback law is identical to that obtained by solving the associated Hamilton-Jacobi-Bellman PDE. When there are states x possessing multiple optimal trajectories (as occurs in the example below), it is no longer clear that an instantaneous receding horizon control can be successfully employed. From a practical point of view, the restriction to $\delta > 0$ is quite sufficient as some computation time is always required.

Theorem 2.2.2 says that for every *fixed* $T > 0$, the receding horizon scheme using a T-horizon optimization is effective. What it does not say, in particular, is that we may vary T and expect a stable process, i.e., stability is not guaranteed (by our results) when the different horizon lengths are allowed at each receding horizon iteration. Note that the dual mode approach described in [MM93] uses (actually requires) a variable optimization horizon. In that case, stability is ensured by *switching* to a stabilizing (linear) controller when close to the equilibrium. In contrast, we note that one does not need to use a *fixed* δ when implementing a receding horizon scheme since (2.7) implies that $x_T^*(\delta; x) \in \Gamma_{r_v}^T$ for all $\delta \in (0, T]$. The stability results are thus independent of δ.

As was mentioned in the introduction, a major issue which to the best of our knowledge has not been addressed rigorously, is the region of attraction of receding horizon control strategies.

One expects that the region of effectiveness should grow as the optimization horizon T is increased, eventually covering all of Γ^∞. This can not be done without increasing r beyond r_v as the following result on inclusions shows.

Proposition 2.2.3 *Let $r > 0$ be given and suppose that $T > 0$ is such that*

$$x_T^*(T;x) \in \Omega_{r_v}$$

for all $x \in \Gamma_r^T$. Then

$$\Gamma_r^T \subset \Gamma_r^{T_1}$$

for all $T_1 \geq T$ so that, in particular, $\Gamma_r^T \subset \Gamma_r^\infty$.

Proof: Using (an extended version of) $\tilde{u}(\cdot)$ from the proof of Theorem 2.2.1, we see that

$$\begin{aligned} J_{T_1}(x, \tilde{u}(\cdot)) &= \int_0^T q(x_T^*(\tau;x), u_T^*(\tau;x))\, d\tau + \int_T^{T_1} q(\tilde{x}(\tau), \tilde{u}(\tau))\, d\tau + V(\tilde{x}(T_1)) \\ &\leq \int_0^T q(x_T^*(\tau;x), u_T^*(\tau;x))\, d\tau + V(x_T^*(T;x)) = J_T^*(x). \end{aligned}$$

It follows that $J_{T_1}^*(x) \leq J_T^*(x)$ for all $x \in \Gamma_r^T$ (cf. [MS97]). □

An important question is whether there exists a suitable horizon length for any desired radius r. The following result guarantees the existence of a suitable optimization horizon for a given (desired) radius r.

Proposition 2.2.4 *For any $r > 0$ there is a $T_v = T_v(r)$ such that*

$$x_T^*(T;x) \subset \Omega_{r_v}$$

for all $x \in \Gamma_r^\infty$ and all $T \geq T_v(r)$. In particular, $x_T^*(T;x) \subset \Omega_{r_v}$ for all $x \in \Gamma_r^T$.

Proof: First, note that $J_T^*(x)$ is bounded (hence well defined) on Γ_r^∞ for all $T \geq 0$ since

$$J_T^*(x) \;\leq\; \int_0^T q(x_\infty^*(\tau;x), u_\infty^*(\tau;x))\, d\tau + V(x_\infty^*(T;x)) \;\leq\; J_\infty^*(x) + b_v(r)$$

where $b_v(r) := \max_{x\in\Gamma_r^\infty} V(x)$. Next, we note that, regardless of the horizon length T, the trajectory $x_T^*(\cdot;x)$ must enter the set Ω_{r_v} within a bounded interval of time. Indeed, let $x \in \Gamma_r^\infty$ and $T > 0$ be arbitrary and suppose that $x_T^*(t;x) \notin \Omega_{r_v}$ on an interval $t \in [0, t_1)$. In this case, the optimal cost satisfies

$$J_T^*(x) \;=\; \int_0^T q(x_T^*(\tau;x), u_T^*(\tau;x))\, d\tau + V(x_T^*(T;x)) \;\geq\; \int_0^{t_1} c_q \|x_T^*(\tau;x)\|^2\, d\tau \;\geq\; \frac{c_q}{\bar{c}_v(r_v)} r_v^2 t_1 \;.$$

Combining the two inequalities, we see that, for $T > 0$ sufficiently large, $x_T^*(\cdot;x)$ must enter Ω_{r_v} with the first arrival time $t_1(x,T)$ satisfying

$$t_1(x,T) \leq \bar{t}_1(r) := \frac{\bar{c}_v(r_v)}{c_q} \frac{r^2 + b_v(r)}{r_v^2} \;.$$

In particular, we see that using $T_v = \bar{t}_1(r) + \epsilon$, $\epsilon > 0$, guarantees the existence of times $t_1(x) < T_v$, $x \in \Gamma_r^\infty$, such that $V(x_{T_v}^*(t_1(x);x)) \leq r_v^2$. The result $x_{T_v}^*(T_v;x) \in \Omega_{r_v}$ follows by Lemma 2.2.1 completing the proof. □

The following corollary follows immediately from the above Proposition.

Corollary 2.2.2 *Let $x_0 \in \Gamma^\infty$ be arbitrary. There exist $r, T < \infty$ such that*

1. *$x_0 \in int\, \Gamma_r^T$*
2. *$x_T^*(T;x) \in \Omega_{r_v}$ for all $x \in \Gamma_r^T$*

Proof: Use $r^2 = J_\infty^*(x_0) + r_v^2 + \epsilon$, $\epsilon > 0$, and $T = T_v(r)$. □

This also shows that Γ^∞ is an open set.

We are now prepared to present the following theorem:

Theorem 2.2.3 *Let Λ be a compact subset of Γ^∞. There is a $T < \infty$ such that Λ is contained in the exponential region of attraction for the receding horizon strategy $\mathcal{RH}(T, \delta)$ for every $\delta \in (0, T]$.*

Proof: For each $x \in \Lambda$, let $U(x) = \text{int } \Gamma_{r(x)}^{T(x)}$ where $T(x)$ and $r(x)$ are given by Corollary 2.2.2. The collection $\{U(x)\}_{x\in\Lambda}$ is an open cover of Λ. By compactness, there is a finite sub-cover $\{U(x_i)\}_{i\leq N}$. Setting $T_i = T(x_i)$ and $r_i = r(x_i)$ we see that

$$\Lambda \subset \bigcup_{i\leq N} \Gamma_{r_i}^{T_i} \subset \bigcup_{i\leq N} \Gamma_{r_i}^{T_m} \subset \Gamma_{r_m}^{T_m}$$

where $T_m = \max_i T_i$, $r_m = \max_i r_i$ and the last two inclusions follow from Proposition 2.2.3.

Setting $T = T_m$ (and $r = r_m$) we see that $x_T^*(T; x) \in \Omega_{r_v}$ for all $x \in \Gamma_r^T \supset \Lambda$. The result follows since (2.7) ensures that $x_T^*(\delta; x) \in \Gamma_r^T$ for all $\delta \in (0, T]$. □

Theorem 2.2.3 tells us that we may make the effective operating region of a receding horizon control strategy as large as we like (relative to the infinite horizon operating region). Of great importance is the fact that this result is obtained using finite horizon optimization without imposing *any* constraints on the terminal cost.

2.3 Unconstrained receding horizon control with no terminal cost

One would expect that as the horizon length grows, the effect of the terminal cost should diminish. In fact, it has been shown that when stability is enforced by terminal stability constraints, the effect of these constraints diminishes as the horizon length is increased [CM96, SR98]. Therefore, it would be reasonable to ask whether there exists a *finite* horizon length such that the receding horizon scheme would remain stabilizing *without* using a CLF as a terminal cost.

We know that in the infinite horizon case, the minimum cost qualifies as a Lyapunov function. Also, we know that as $T \to \infty$, $J_T^* \to J_\infty^*$. The remaining question

is whether there is a (large enough) finite horizon length for which the minimum cost qualifies as a suitable Lyapunov function. This question was answered fairly recently in the context of constrained discrete-time linear systems [PN00]. We will show that a similar result holds in the case of unconstrained nonlinear systems. Before proving this result, we will need the following definition and proposition [Roy88]:

Definition 2.3.1 *An extended real-valued function f is called upper semicontinuous at the point y if $f(y) \neq +\infty$ and $f(y) \geq \lim_{x \to y} f(x)$. We say that f is upper semicontinuous on a set if it is upper semicontinuous at each point of that set.*

Proposition 2.3.1 (Dini): *Let $\{f_n\}$ be a sequence of upper semicontinuous, real-valued functions on a countably compact space X, and suppose that for each $x \in X$, the sequence $\{f_n(x)\}$ decreases monotonically to zero. Then the convergence is uniform.*

Now we are ready to present the following theorem:

Theorem 2.3.1 *Consider the receding horizon scheme $\mathcal{RH}(T, \delta)$ with zero terminal cost. Then there exists a T^* such that for $T \geq T^*$, $\mathcal{RH}(T, \delta)$ is stabilizing with $\Gamma_{r_T}^{T-\delta}$ being the region of attraction, where for any given $r > 0$, $r_T > 0$ is the largest value such that $\Gamma_{r_T}^{T-\delta} \subset \Gamma_r^{\infty}$.*

Proof: Using the principle of optimality, we can write the following:

$$J_T^*(x) - J_{T-\delta}^*(x_T^*(\delta; x)) = \int_0^{\delta} q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau \quad \forall\ x \in \Gamma_r^{\infty} \tag{2.8}$$

where $J_T^*(x) := \int_0^T q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau$, i.e., the same as before except for zero terminal cost.

Adding $J_{T-\delta}^*(x)$ to both sides of (2.8), and taking $J_T^*(x)$ to the other side, we obtain the following:

$$J_{T-\delta}^*(x) - J_{T-\delta}^*(x_T^*(\delta; x)) = J_{T-\delta}^*(x) - J_T^*(x) + \int_0^{\delta} q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau. \tag{2.9}$$

Also, note that contrary to the case with a CLF terminal cost, the shorter the horizon length, the lower the cost value. To see this, suppose $T_1 < T_2$, then

$$J^*_{T_2}(x) = \int_0^{T_1} q(x^*_{T_2}(\tau;x), u^*_{T_2}(\tau;x))d\tau + \int_{T_1}^{T_2} q(x^*_{T_2}(\tau; x^*_{T_2}(T_1;x)), u^*_{T_2}(\tau))d\tau$$

which is clearly greater than $J^*_{T_1}(x)$. Therefore, we can write the following inequalities:

$$J^*_\delta(x) \leq \int_0^\delta q(x^*_T(\tau;x), u^*_T(\tau;x))d\tau \quad 0 < \delta \leq T. \tag{2.10}$$

The last inequality follows from the fact that u^*_T is not (necessarily) optimal over $[0, \delta]$. Hence, we can write (2.9) as the following inequality:

$$J^*_{T-\delta}(x) - J^*_{T-\delta}(x^*_T(\delta;x)) \geq J^*_{T-\delta}(x) - J^*_T(x) + J^*_\delta(x) \quad \forall\, 0 < \delta \leq T. \tag{2.11}$$

Also, (2.10) implies that the finite horizon sub level sets are indeed larger than the infinite horizon ones, i.e.,

$$\Gamma^T_r \subset \Gamma^\infty_r \quad \forall\, r > 0.$$

All we need in order to prove stability is to show that there exist a large enough but finite T^* such that the right-hand side of (2.11) is positive. To show this, we define the following function over the compact set Γ^∞_r:

$$\begin{aligned} \Phi_T(x) &:= \frac{J^*_\infty(x) - J^*_T(x)}{J^*_\delta(x)} \quad \forall\, x \neq 0 \\ \Phi_T(0) &:= \limsup_{x\to 0} \Phi_T(x) \end{aligned} \tag{2.12}$$

It is clear that for each $x \neq 0$, $\Phi_T(x)$ is continuous, because the numerator of $\Phi_T(x)$ is continuous and the denominator is continuous and non-zero. Hence it is also *upper semicontinuous* everywhere except possibly at the origin. In order to show that $\Phi_T(x)$ is upper semicontinuous at the origin as well, all we need to prove is that $\Phi_T(0) \neq \infty$. The rest is clear from definition 2.3.1 and the definition of $\Phi_T(0)$.

Also, note that for all non-zero values of $x \in \Gamma^\infty_r$, the sequence $\{\Phi_T(x)\}$ is a

monotone, decreasing sequence, tending to zero as $T \to \infty$. We must show that $\Phi_T(0)$ is bounded and converges to zero when $T \to \infty$ as well. This is true due to the fact that around the origin, $J_T^*(x) = x^T P_T x + r_T(x)$, $J_\infty^*(x) = x^T P_\infty x + r_\infty(x)$, and $J_\delta^*(x) = x^T P_\delta x + r_\delta(x)$ where P_T, P_δ, and P_∞ are the positive definite solutions of the corresponding differential and algebraic Riccati equations respectively [DAC95], and

$$\lim_{x \to 0} \frac{r_i(x)}{||x||^2} = 0 \quad i = \{T, \infty, \delta\}.$$

Therefore, in a small neighborhood of the origin we have:

$$\Phi_T(x) \approx \frac{x^T(P_\infty - P_T)x}{x^T P_\delta x} \leq \frac{\lambda_{\max}(P_\infty - P_T)}{\lambda_{\min}(P_\delta)} \to 0$$

as $T \to \infty$. This also proves that $\Phi_T(0)$ is bounded from above, showing that $\Phi_T(x)$ is upper semicontinuous at the origin. The same argument can be used for $\Phi_{T-\delta}(x)$, by merely replacing T with $T - \delta$.

We can now conclude that $\{\Phi_{T-\delta}(x)\}$ is a sequence of upper semicontinuous functions that converges monotonically and pointwise to zero; therefore by Proposition 2.3.1, the convergence is uniform.

Hence, by the definition of uniform convergence, there exists a T^* such that for any $T \geq T^*$, $\Phi_{T-\delta} < \frac{1}{2}$, implying that

$$J_\infty^*(x) - J_{T-\delta}^*(x) < \frac{1}{2} J_\delta^*(x).$$

Thus,

$$J_{T-\delta}^*(x_T^*(\delta; x)) - J_{T-\delta}^*(x) \leq -\frac{1}{2} J_\delta^*(x) \quad \forall\, x \in \Gamma_r^\infty$$

(note that $J_\infty^*(x) \geq J_T^*(x)$). This implies that $J_{T-\delta}^*(x)$ is a Lyapunov function proving asymptotic (and in fact exponential) stability of the closed-loop system with $\Gamma_{r_T}^T$ being the region of attraction, and r_T^2 being the value of the largest sub level set of $J_{T-\delta}^*(x)$ that lies inside Γ_r^∞, for any given $r > 0$. (See Figure 2.4 for an illustration of the set inclusions). $\square$

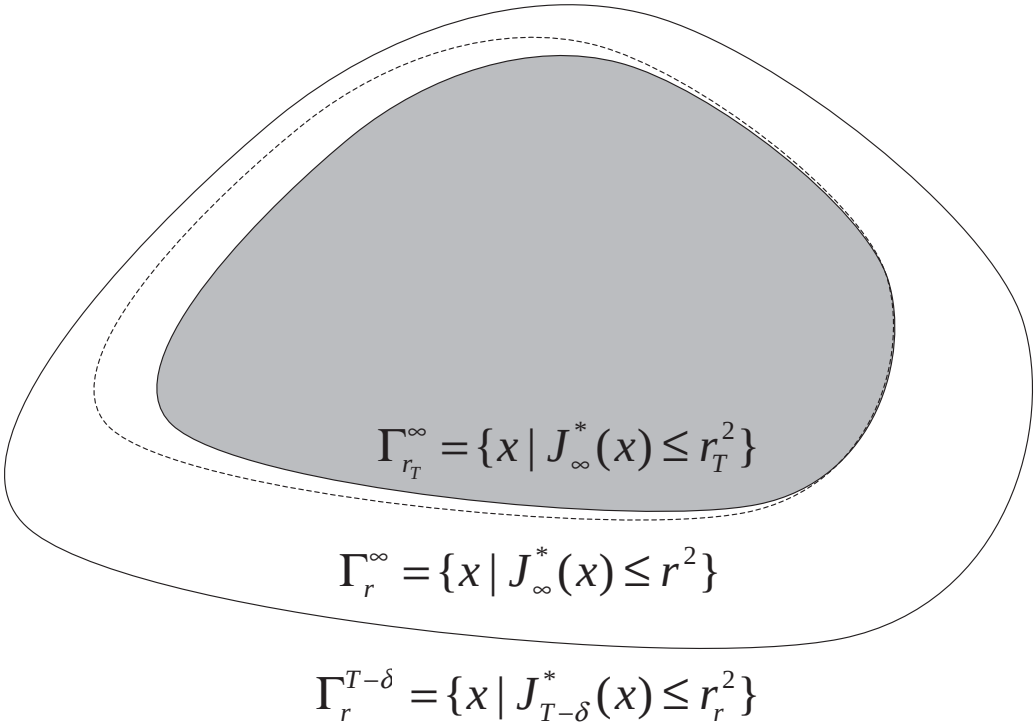

Figure 2.4: The region of attraction $\Gamma_{r_T}^{\infty}$ (in gray) as well as $\Gamma_{r}^{T-\delta}$ (dotted line) and Γ_{r}^{∞} (solid line).

The main consequence of this theorem is that when the horizon length is long enough, no off line calculation is required to ensure closed-loop stability. This opens up the possibility of using the receding horizon scheme in cases where the dynamics and/or the cost change in real-time (such as when a fault occurs). A practical application of these results will be demonstrated in Chapter 6.

2.4 Unconstrained receding horizon control with magnitude saturation

As it was mentioned in Chapter 1, one of the main attractions of the receding horizon control scheme is its ability to handle *input constraints*. Due to physical limits in actuation authority, any practical control strategy has to be able to cope with these types of constraints.

In this section we formulate a slightly different optimization scheme for inclusion of magnitude constraints on controls, i.e., instead of explicitly imposing the magnitude constraints in the optimization, we use a smooth saturating function such as a hyperbolic tangent function to implicitly impose the saturation constraint. The optimization problem to be solved can be written as follows:

$$J_T^*(x) = \inf_{v(\cdot)} \int_0^T q(x^u(\tau;x), \sigma(v(\tau))) + \frac{1}{2}||v-u||_R^2 \, d\tau + V(x^u(T;x)) \; .$$

subject to : $\dot{x} = f(x, \sigma(v)) \quad u = \sigma(v)$

where $\sigma(v) = \tanh(v)$ is a smooth saturating function. Note that the above optimization problem is effectively similar to the one with the magnitude constraint and with u as an input except for the fact that by eliminating the constraint we have kept the problem as unconstrained. The extra penalty term is added to force the control to stay in the linear part of the saturation curve.

In order to be able to utilize the stabilizing receding horizon approach developed

in this chapter in this case, we have to use a CLF that respects the input constraint. Since such CLFs are extremely hard to come up with in general, we will use the same unconstrained CLF. However, we have to reduce r_v, the radius of the region of attraction, such that the input constraints are satisfied. We would search for the largest level set of the CLF which satisfies $\min_{||u||\leq 1}(\dot{v} + q_e)(x, u) \leq 0$, where q_e is the modified cost functional including the extra term $\frac{1}{2}||v - u||_R^2$. Without loss of generality, we have assumed a unity bound on the magnitude of control action. This way we have implicitly taken the saturation into account without searching for a constrained CLF or imposing the constraints on the optimization.

2.5 Example

For the purpose of illustration, we consider the problem of balancing an inverted pendulum on a cart. We discard the states associated with the cart to allow two-dimensional visualization. (Please note that this is a highly unrealistic system as it allows equilibria where the cart is experiencing continuous acceleration—the system is *for visualization only*.) The pendulum is modeled as a thin rod of mass m and length $2l$ (the center of mass is at distance l from pivot) riding on a cart of mass M with applied (horizontal) force u. The dynamics of the pendulum are then given by (with θ measured from the vertical up position)

$$\ddot{\theta} \;=\; \frac{g/l\ \sin\theta - m_r\dot{\theta}^2/2\ \sin 2\theta - m_r/ml\ \cos\theta\ u}{4/3 - m_r\cos^2\theta}$$

where $m_r = m/(m+M)$ is the mass ratio and g is the acceleration of gravity. Specific values used are $m = 2$ kg, $M = 8$ kg, $l = 1/2$ m, and $g = 9.8$ m/s^2.

System performance is measured using the quadratic incremental cost $q(x, u) = 0.1x_1^2 + 0.05x_2^2 + 0.01u^2$ where as usual the state is $(x_1, x_2) = (\theta, \dot{\theta})$. To obtain an appropriate control Lyapunov function, we modeled the system locally as a Polytopic Linear Differential Inclusion (PLDI) [BGFB94]. This approach is quite satisfactory for this simple (planar) system over a large range of angles. (The details of obtaining

such CLFs will be discussed in detail in Chapter 4.) Working over a range of plus or minus 60 degrees, we obtained the quadratic CLF $V(x) = x^T P x$ with

$$P = \begin{bmatrix} 151.57 & 42.36 \\ 42.36 & 12.96 \end{bmatrix}.$$

Simple numerical calculations (in low dimensions!) show that $r_v \approx 6.34$, that is, $\min_u(\dot{V} + q)(x, u)$ is negative on solid P-ellipses Ω_r with a radius $r < 6.34$. An optimization technique that can be adapted to the problem of computing r_v in higher dimensions can be found in [LH93].

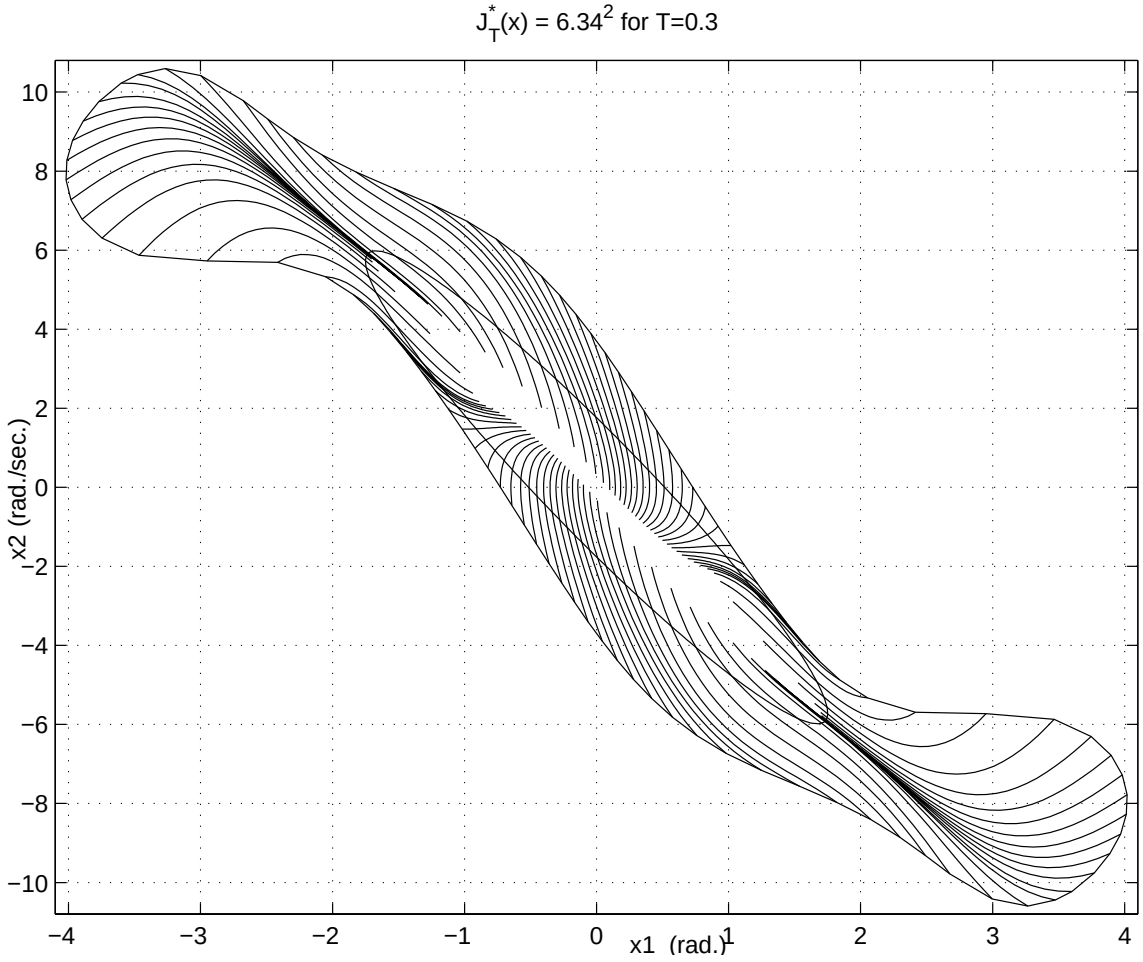

Figure 2.5: The sublevel set Γ_r^T for $T = 0.3$ and $r = r_v = 6.34$ together with Ω_{r_v}. Also depicted are the trajectories $x_T^*(\cdot; x)$ for x on the boundary of Γ_r^T.

By Theorem 2.2.2, we know that, for $T \geq 0$, $\Gamma_{r_v}^T$ is an invariant subset of the region

of attraction for the receding horizon controller $\mathcal{RH}(T, \delta)$ with $\delta \in (0, T]$. Figure 2.5 depicts the set $\Gamma^T_{r_v}$ for $T = 0.3$, $r_v = 6.34$ together with trajectories $x^*_T(\cdot; x)$ for x on the boundary. Also shown is the set Ω_{r_v}. The inclusion $\Omega_{r_v} \subset \Gamma^T_{r_v}$ (Proposition 2.2.3) is evident as is the fact that $x^*_T(T, x) \in \Omega_{r_v}$ for $x \in \Gamma^T_{r_v}$.

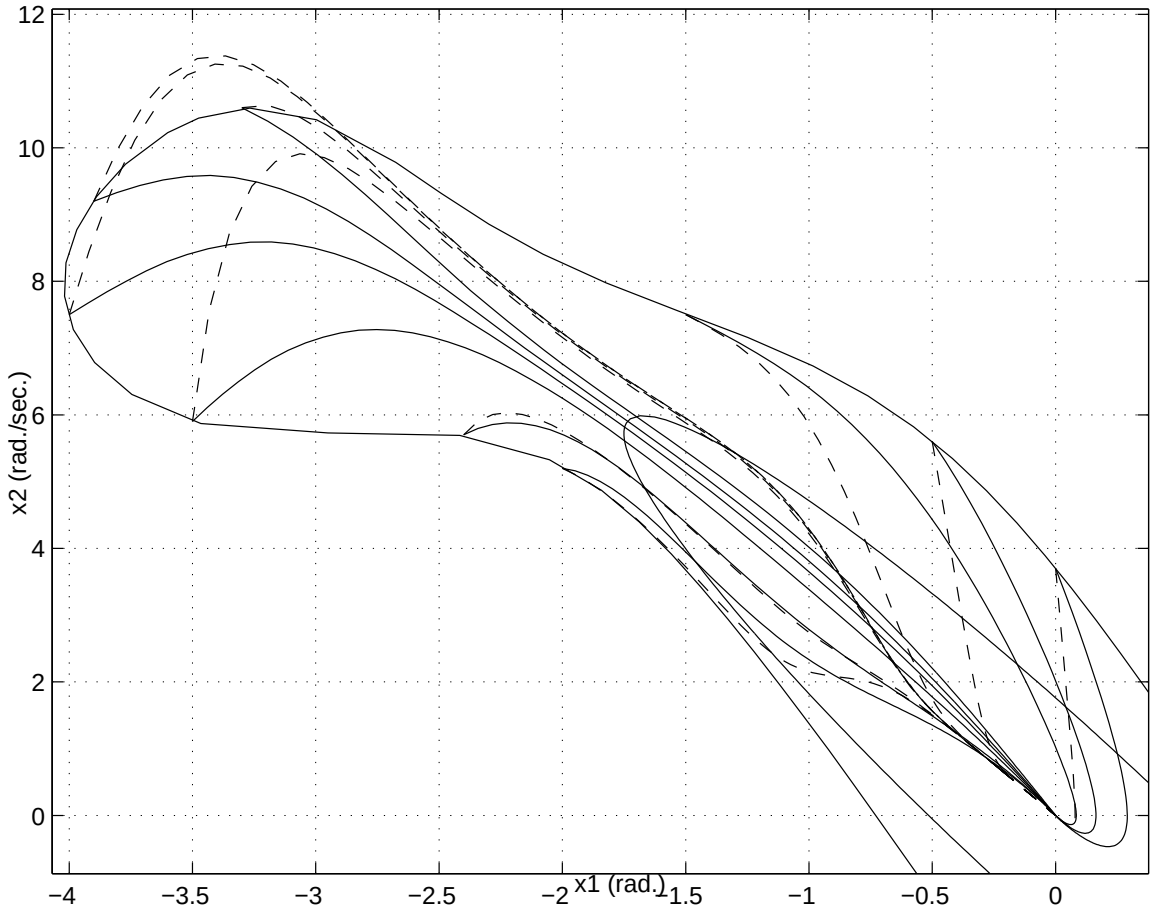

Figure 2.6: Receding horizon $\mathcal{RH}(0.3, 0.05)$ and CLF controller (dashed) trajectories.

Figure 2.6 provides a comparison of receding horizon trajectories (for $\mathcal{RH}(0.3, 0.05)$) with those obtained using the CLF controller $u = k(x) = \arg\min_w (\dot{V} + q)(x, w)$. Note that $\Gamma^T_{r_v}$ is not invariant under the CLF flow. As expected, the receding horizon trajectories do remain inside $\Gamma^T_{r_v}$.

We also note that the CLF controller often requires significantly more control authority. For example, as shown in Figure 2.7, the CLF controller demands almost 15 times as much authority when stabilizing from $x_0 = (-3.5, 5.9)$. (The pair orig-

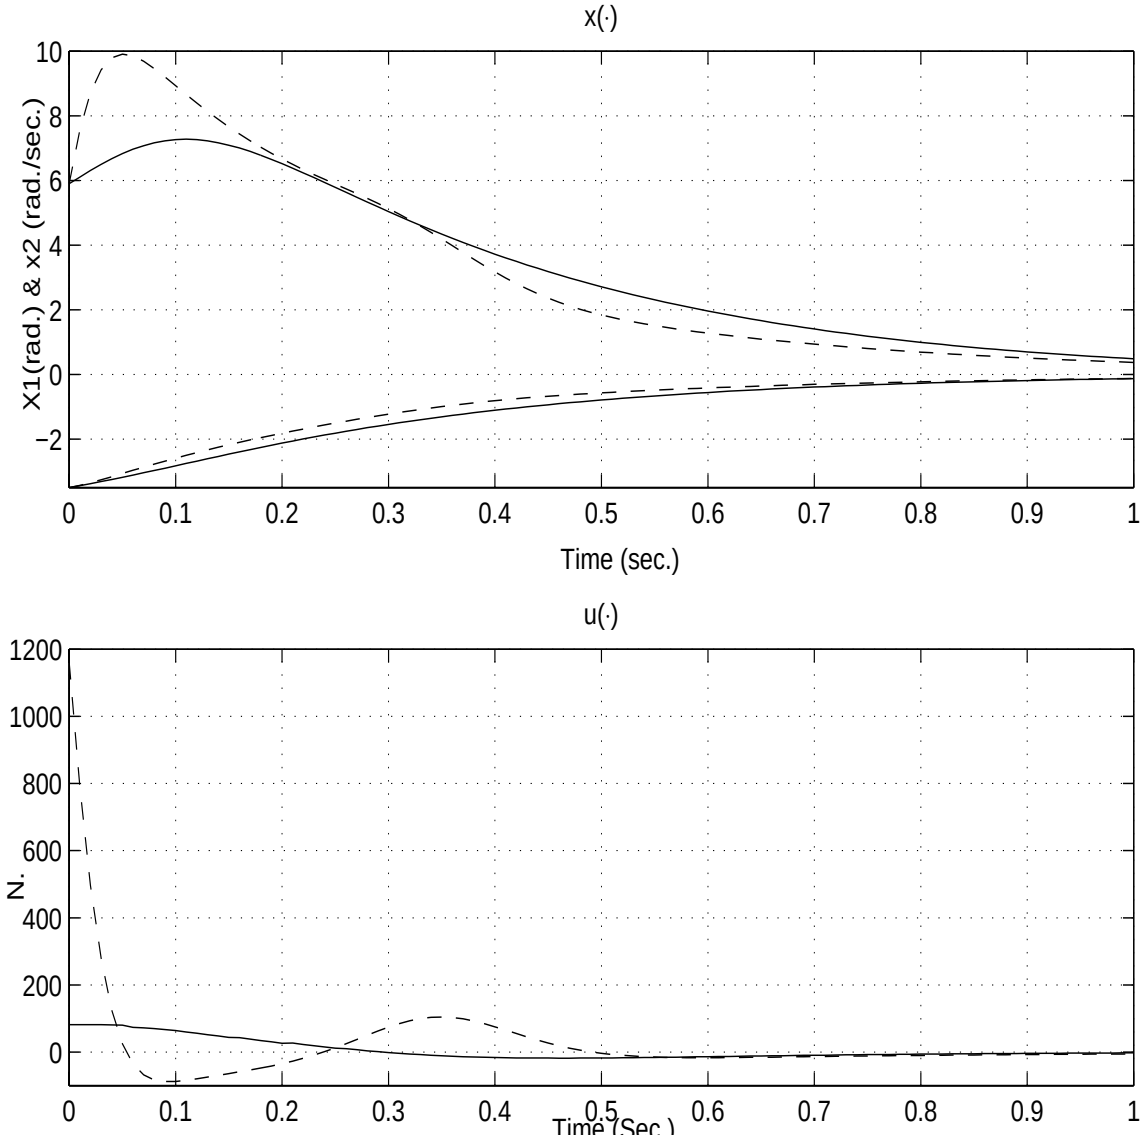

Figure 2.7: State and control trajectories (RH-solid and CLF-dashed) from $x_0 = (-3.5, 5.9)$.

inating at $x_0 = (-4.0, 7.5)$ has a ratio greater than 35.) This is not too surprising since the CLF controller was designed for angular deviations of perhaps 60 degrees and qualified on the set Ω_{r_v}. The chosen x_0 is well outside of the guaranteed CLF performance region. In contrast, a small optimization horizon ($T = 0.3$ compared with a convergence time of > 1.5) allows the receding horizon controller to exploit its knowledge of the nonlinear system dynamics in this region.

Some relevant quantities for the trajectories in Figure 2.6 are given in the following table (the square root of the cost is shown in parentheses):

x	$J_T^*(x)$	$J_\infty(x, u_{\rm rh}(\cdot))$	$J_\infty(x, u_{\rm clf}(\cdot))$
$(-3.5, 5.9)$	40.10	10.23 (3.20)	184.21 (13.57)
$(-0.5, 5.6)$	40.18	38.26 (6.19)	113.47 (10.65)
$(-3.9, 9.2)$	39.96	7.92 (2.81)	80.07 (8.95)
$(-2.0, 5.2)$	42.41	31.29 (5.59)	35.14 (5.93)
$(-4.0, 7.5)$	39.90	2.00 (1.41)	155.98 (12.49)
$(0.0, 3.7)$	40.18	38.59 (6.21)	101.34 (10.07)
$(-3.3, 10.6)$	40.25	17.64 (4.20)	27.74 (5.27)
$(-2.4, 5.7)$	39.69	28.26 (5.32)	30.80 (5.55)
$(-1.5, 7.5)$	39.76	31.88 (5.65)	80.30 (8.96)

Table 2.1: The cost of receding horizon control for various T values ($\sqrt{\text{cost}}$).

These values confirm the fact (Proposition 2.2.2) that $J_\infty(x, u_{\rm rh}(\cdot)) \leq J_T^*(x)$. In this case we see that significant performance improvements are obtained through the use of a relatively inexpensive receding horizon strategy.

The appropriate finite horizon optimization problems were solved numerically using RIOTS [Sch96] as well as some local codes that are under development.

Table 2.1 illustrates the effectiveness of using various horizon lengths T in an unconstrained receding horizon scheme (with $\delta = 0.05$). The table compares the cost (and its square root) of using receding horizon control with $T = 0.3$ and $T = 1.0$ with that resulting from the use of the CLF controller ($T = 0.0$). One observes a great reduction in cost for trajectories obtained using even modest amounts of online

optimization. We also see that near optimal (infinite horizon) performance may be obtained on a larger operating region by using a longer optimization horizon. See, especially, the results for $x(0) = (\pi/2, 2)$.

Further insights into the contrasting nature of these trajectories can be found by presenting the trajectories in phase coordinates as in Figure 2.8. The $r = 6.34$ level set contours for V and J_1^* are also shown in that figure. One notes how natural the receding horizon trajectories appear, especially by comparison to the corresponding CLF trajectories.

One notes that, away from the stable manifold of the uncontrolled system, the cost of using the CLF controller may be very large. (In fact, the phase portrait of the CLF controlled closed loop system indicates that the CLF controller uses a great amount of control energy to quickly force the system onto a slow manifold for transit to the origin. Moreover, this behavior is observed in the large (at great cost) with small eddies such as that observed near $(\pi/2, 0)$.)

We should point out that there is a curve of points where J_∞^* is not differentiable that roughly parallels the $r = 6.34$ contour in a northwest direction from the CLF eddy. From those points, one finds two very different trajectories with precisely the same cost. The presence of such a shock (in the "solution" to the Hamilton-Jacobi equation) somewhat complicates the story by allowing kinks in the level sets.

Surprisingly, our use of finite horizon approximations seems to improve the situation somewhat. In fact, the use of receding horizon control with a CLF terminal cost allows us to, in effect, find near optimal trajectories over a larger region by exploiting local CLF stability properties. This will be discussed in more detail in Chapter 3. We also mention that the optimal trajectories for this problem, even those that originate at points where the value function is not differentiable, appear to satisfy local second order sufficiency conditions for an isolated local minimum. In that case, reliable computations should be possible provided suitable initial trajectories may be obtained.

As it can be seen from Figure (2.7), when no saturation is in effect, a large control action is required to stabilize the system, even in the case of the receding horizon

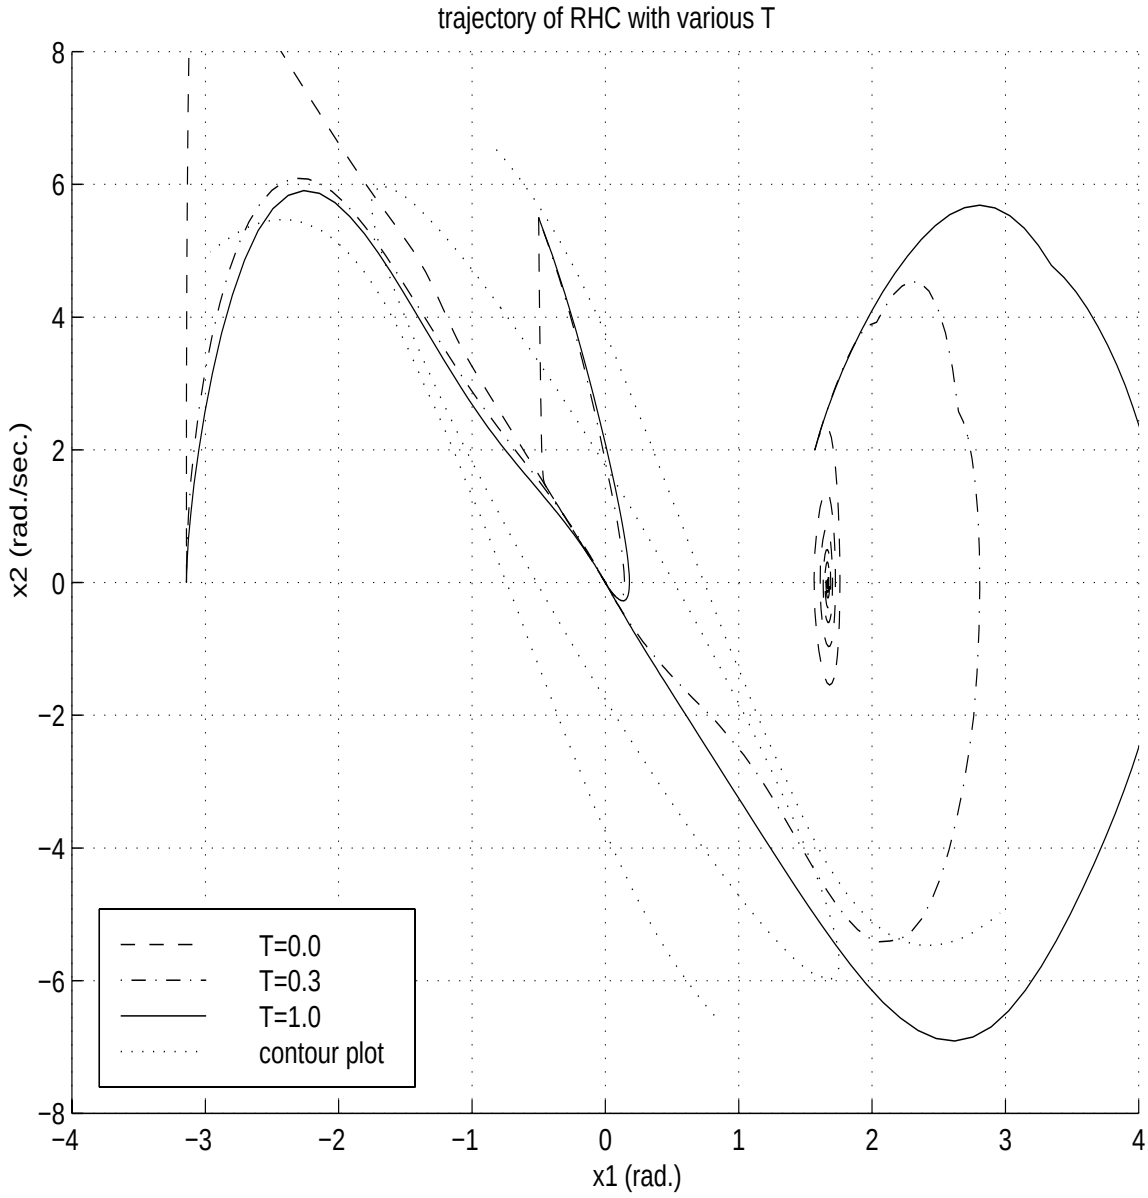

Figure 2.8: Comparison of trajectories for receding horizon control using various horizons T.

controller. Figure (2.9) illustrates what happens when the magnitude of the control action $||u||$ is required to be less than 60. In order for V to remain a valid CLF, r_v is reduced from 6.34 to 4.5. However, by choosing a larger horizon length of $T = 1$, we can compensate for the reduction in the size of the region of attraction. Control and state trajectories are depicted in Figure (2.10).

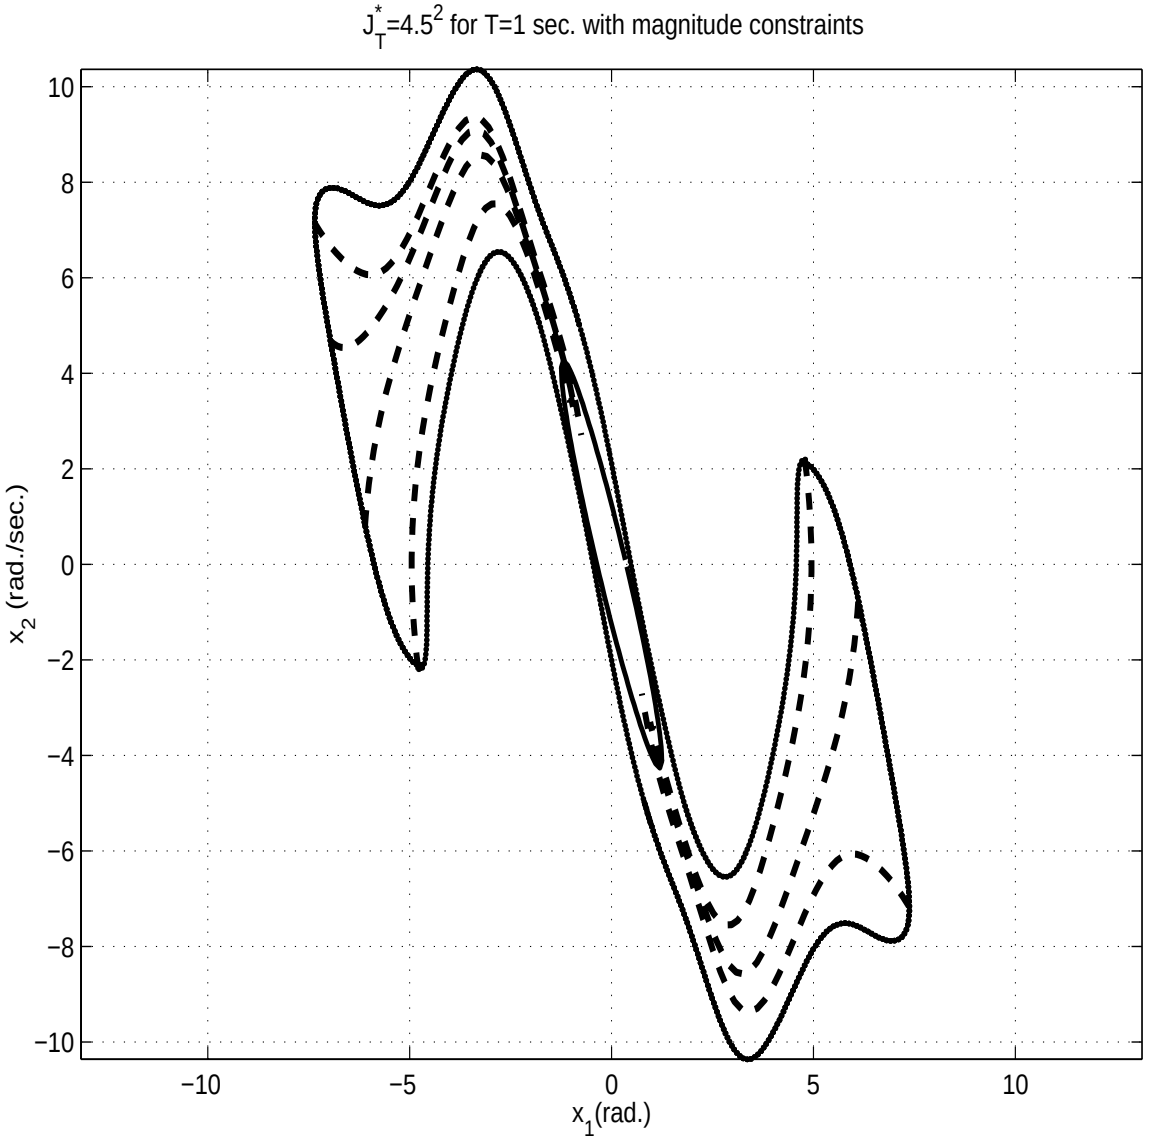

Figure 2.9: The sub level set Γ_r^T for $T = 1.0$ and $r = r_v = 4.5$ together with Ω_{r_v}. Also depicted are the trajectories $x_T^*(\cdot; x)$ for x on the boundary of Γ_r^T with magnitude constraints $|u| \leq 60$ from $x_0 = (-3.5, 5.9)$.

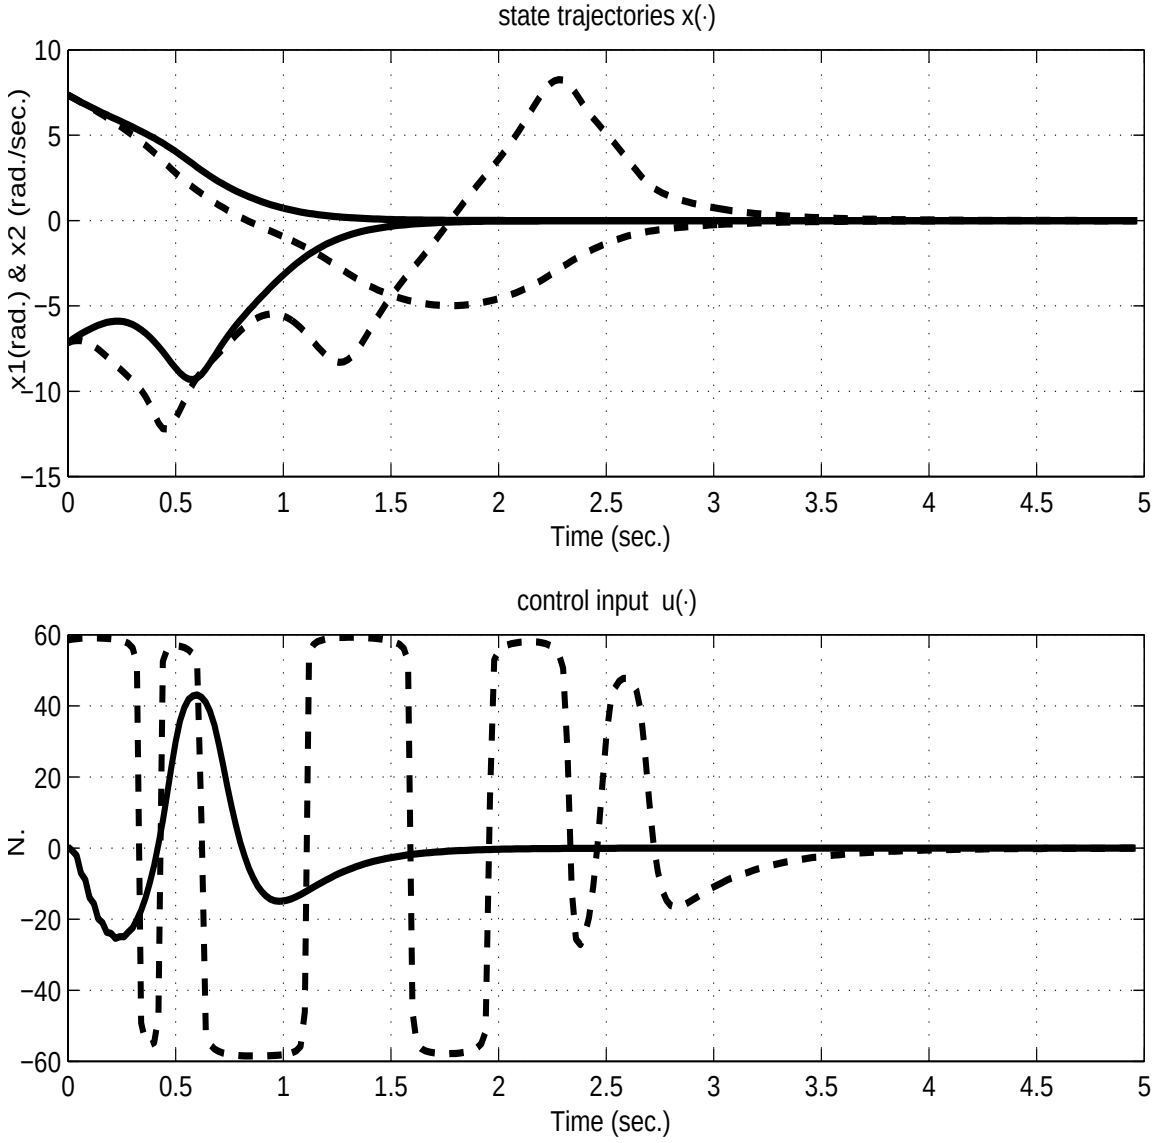

Figure 2.10: State and control trajectories (RH-solid and CLF-dashed) from $x_0 = (7.36, -7.15)$.

2.6 Summary

In this chapter, we have developed a family of receding horizon control strategies that obtain excellent stability and performance properties through the use of a control Lyapunov function as terminal cost. This approach is quite natural, providing a happy medium between the use of a CLF controller and an ideal infinite horizon controller. Of practical significance, we have shown that this approach does not require the introduction of terminal constraints for stability (such as the ones required by [Pri98]). As was discussed in Chapter 1, results of [Pri98] enforce stability by requiring that the derivative of CLF decreases sufficiently along the open-loop optimal trajectories. We have shown that such an approach is not necessary, thereby eliminating a key source of computational burden

In fact, it appears that these computations may be made fast enough to allow their use even in challenging areas such as flight control. An implicit assumption in this chapter is that (globally) optimal trajectories are calculated precisely. This restriction will be relaxed in Chapter 3. An interesting direction is the extension of these techniques to the case of the trajectory tracking for nonlinear systems. Of course, the situation is much more complicated since the problem of finding useful trajectories of a nonlinear system is itself a rather difficult problem. A first step in that direction is the use of *trajectory morphing* techniques [HM98b, HM98a]. This will be addressed in Chapter 6.

Chapter 3 Relaxing the optimality condition in receding horizon control

In the previous chapter, we have detailed the theoretical properties of *ideal* receding horizon strategies wherein a *global* minimum is computed at each step. Only in very special cases (e.g., linear dynamics, strictly convex cost, etc.) can one expect reliable (approximate) computation of a global minimum. It is the purpose of this chapter to illustrate one of the many ways in which this requirement may be relaxed. Several authors have addressed this issue in the context of *constrained* optimization. The purpose of this section, however, is to show that *without* requiring constraint optimization in each iteration, stability can be guaranteed with the same approach of Chapter 2; see Chen and Allgöwer [CA98], and Scokaert, Mayne, and Rawlings [SMR99] for results of this nature.

As was discussed earlier, receding horizon techniques produce a sequence of (state and control) trajectories with ever decreasing cost. Stabilization or, more precisely, convergence of the cost may be obtained by ensuring that there is sufficient *improvement* at each step. Thus we may replace the optimality test at *each* step by a test for improvement *between* steps.

3.1 Relaxing the requirement for optimality

The following result provides a sufficient condition to ensure convergence of the state to the origin.

Proposition 3.1.1 *Fix $T, \delta > 0$ and let $x_i, u_i(\cdot)$, $i \geq 0$, be such that $x_{i+1} = x^{u_i}(\delta; x_i)$ and*

$$J_T(x_{i+1}, u_{i+1}(\cdot)) \leq J_{T-\delta}(x_{i+1}, u_i(\cdot + \delta)) \; . \tag{3.1}$$

Then $x_i \to 0$ *as* $i \to \infty$.

Proof: Note that the sequence of costs $c_i := J_T(x_i, u_i(\cdot))$ is monotone decreasing and bounded from below. It follows that the incremental cost

$$\delta c_i = c_i - c_{i+1} \geq \int_0^\delta q(x^{u_i}(\tau; x_i), u_i(\tau))\, d\tau$$

must go to zero as $i \to \infty$, since the cost is lower bounded by zero and at each iteration there is *sufficient* amount of decrease. This implies that $x_i \to 0$ since there is a $\kappa > 0$ such that $\int_0^\delta q(x^u(\tau; x), u(\tau))\, d\tau \geq \kappa \cdot \min\{1, \|x\|^2\}$ for every $u(\cdot)$. □

How may we ensure, at each step, the existence of an improving control $u_{i+1}(\cdot)$? Similar to the controller $\tilde{u}(\cdot)$ constructed in the proof of Theorem 2.2.1, we can use a dual mode control. The difference here is the fact that the control need not be optimal over $[\delta, T]$:

Proposition 3.1.2 *Suppose that* x_0 *and* $u_0(\cdot)$ *are such that* $x^{u_0}(T; x_0) \in \Omega_{r_v}$. *Then, there exists a sequence of controls* $\{u_i(\cdot)\}_1^\infty$ *such that* $x^{u_i}(\delta; x_i) = x_{i+1} \to 0$ *as* $i \to \infty$.

Proof: Given x_i, $u_i(\cdot)$, choose $u_{i+1}(\cdot)$ such that $x^{u_{i+1}}(T; x_{i+1}) \in \Omega_{r_v}$ and the improvement property (3.1) is satisfied. One choice is the control obtained by using the remainder of $u_i(\cdot)$ in the interval $[\delta, T]$, followed by a CLF feedback control (as in the proof of Theorem 2.2.1) for the interval $[T, T + \delta]$. The rest of the proof follows that of Theorem 2.2.1. □

One may (and many have) use constrained optimization to solve, at each step, a feasibility problem of the sort indicated. In that regard, the above result shows that the problem will remain feasible if it is initially thus. Also, since feasible controls may be obtained *for free*, we may use any means whatsoever (including unconstrained optimization) in our search for better controls, accepting only those that satisfy both terminal and improvement conditions. Specifically, we can use the results of Theorem 2.2.2 to indicate that the sub level sets of the cost function $J_T(x^u, u(\cdot)$ is an estimate of the regions of attraction of the *unconstrained* receding horizon scheme. In other

words, one can infer all the results obtained in Chapter 2 in this case, by merely replacing the optimal trajectory u_T^* with any control trajectory that guarantees (3.1).

Of course, there is always a need for an *initial* feasible trajectory. As was shown earlier, inside Ω_{r_V}, one can use the CLF controller for warm starting the optimization.

3.2 Inverted pendulum revisited

The inverted pendulum example was used to illustrate the theoretical properties of our proposed receding horizon scheme in Chapter 2. We now use the same example to illustrate a very interesting phenomenon. It was mentioned in Section 2.5 that for long horizon lengths, the optimized trajectory behaves in a qualitatively different manner. See Figure (2.8).

It is the purpose of this example to elaborate more on this issue and provide a clear-cut picture. For small values of the horizon length T, the set $\Gamma_{r_v}^T$ looks very much like the ellipse Ω_{r_v} which is reasonably well lined up with the stable manifold of the pendulum. As the value of T is increased, the ends of $\Gamma_{r_v}^T$ begin to open up, eventually wrapping back around toward the inverted equilibrium, indicating that it can be efficient (from a cost standpoint) to allow the pendulum to swing down before bringing it back up to the vertical position. Figure 3.1 depicts the nature of this wrap-around for $T = 2.0$ and $r = r_v = 6.34$. The set $\Gamma_{r_v}^T$ is shown without an overlap by plotting half of the set boundary which, together with trajectories starting on the boundary, provide an unwrapped view of the set.

Figure 3.2 provides a close up view of the overlapping set $\Gamma_{r_v}^T$ together with the set Ω_{r_v}. At each point in the overlap region, there are (at least) two local minima. Strict use of the *global* optimum in a receding horizon strategy would indicate a preference for letting the pendulum fall in many situations where the pendulum can be brought back to the vertical quickly and for a reasonable, though suboptimal, cost.

Consider, for example, the use of a receding horizon strategy with $\delta = 0.1$ (and $T = 2.0$) starting at the initial condition $x(0) = x_0 = (-1.803, 8.413)$ with optimal cost $J_T^*(x_0) = r_v^2 = 40.1956$. The situation is depicted in Figure 3.3. After $\delta = 0.1$

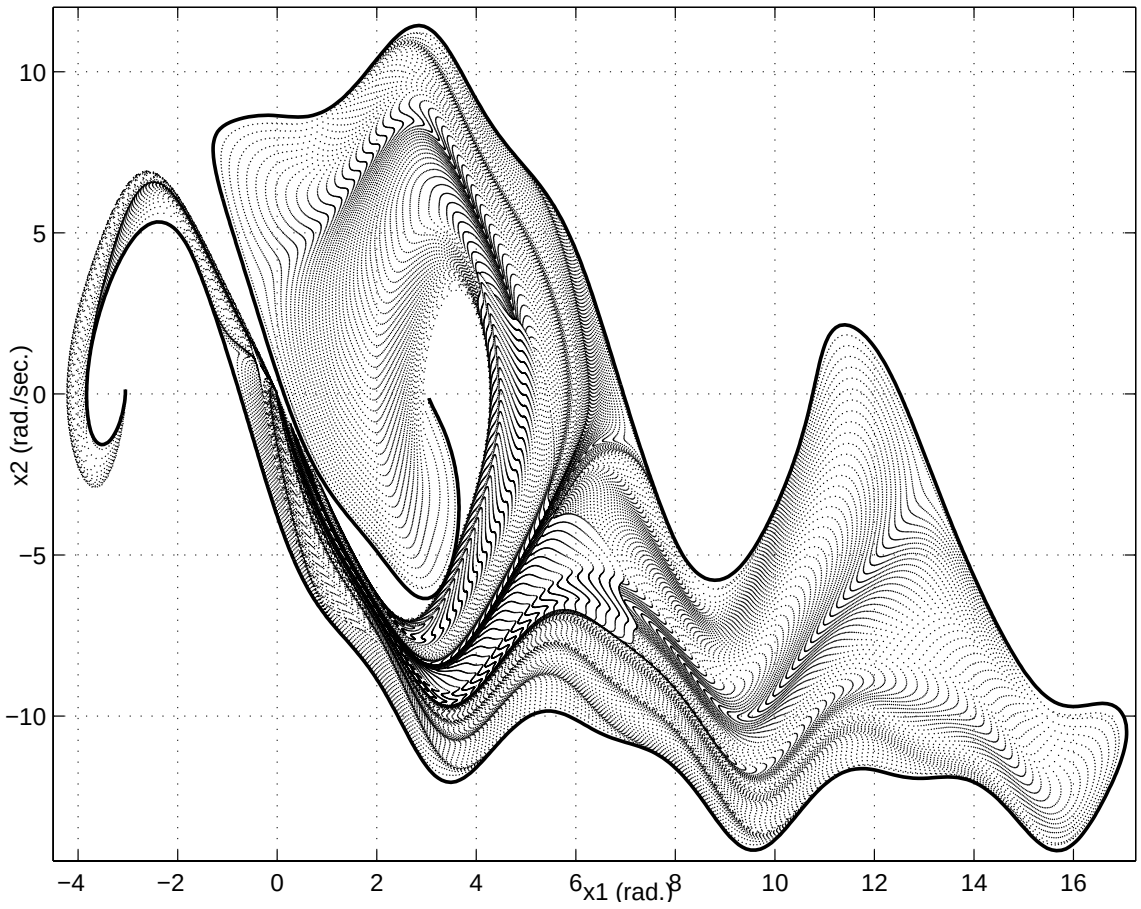

Figure 3.1: The sub-level set Γ_r^T for $T = 2.0$ and $r = r_v = 6.34$. Half of the boundary (together with trajectories) is shown in an *unwrapped* fashion to aid in understanding the overlapping nature of the set.

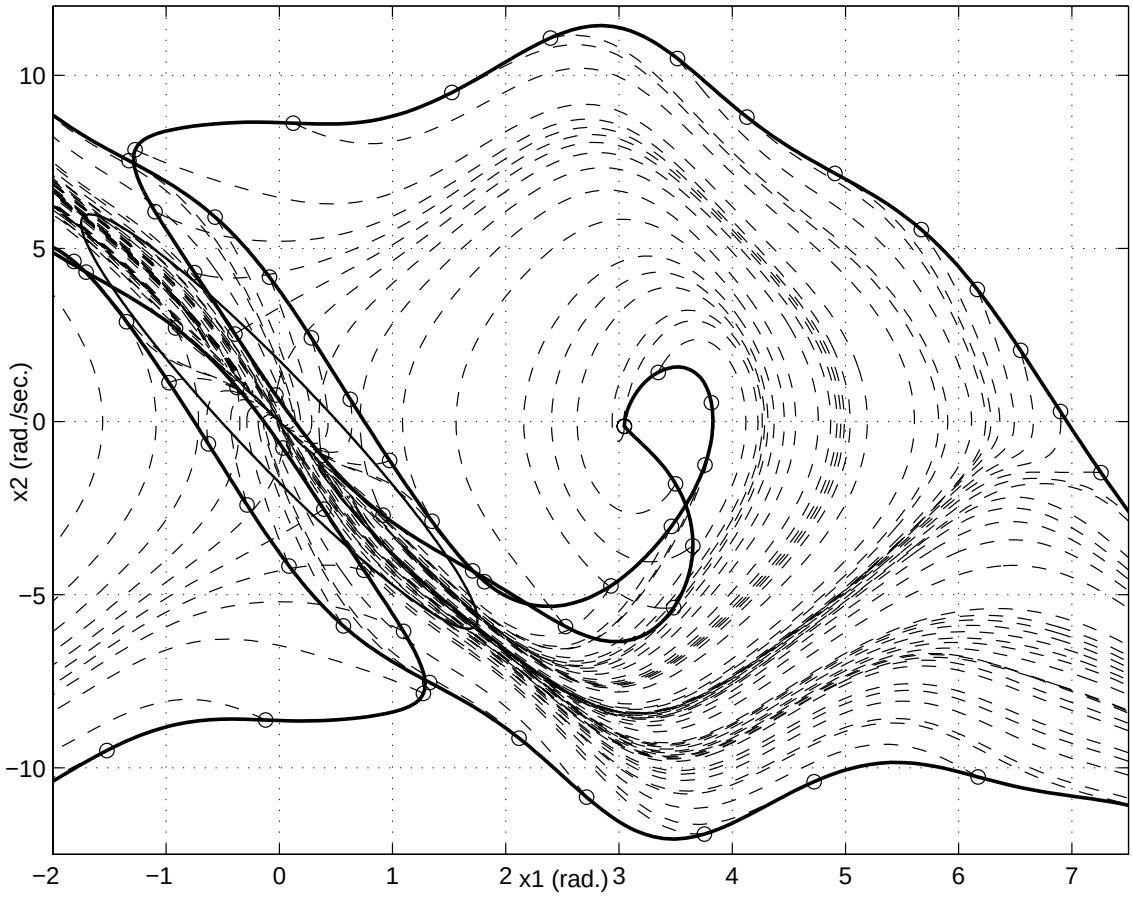

Figure 3.2: A closeup of the sub-level set Γ_r^T for $T = 2.0$ and $r = r_v = 6.34$ together with Ω_{r_v}. Also depicted are several locally optimal trajectories beginning on the boundary.

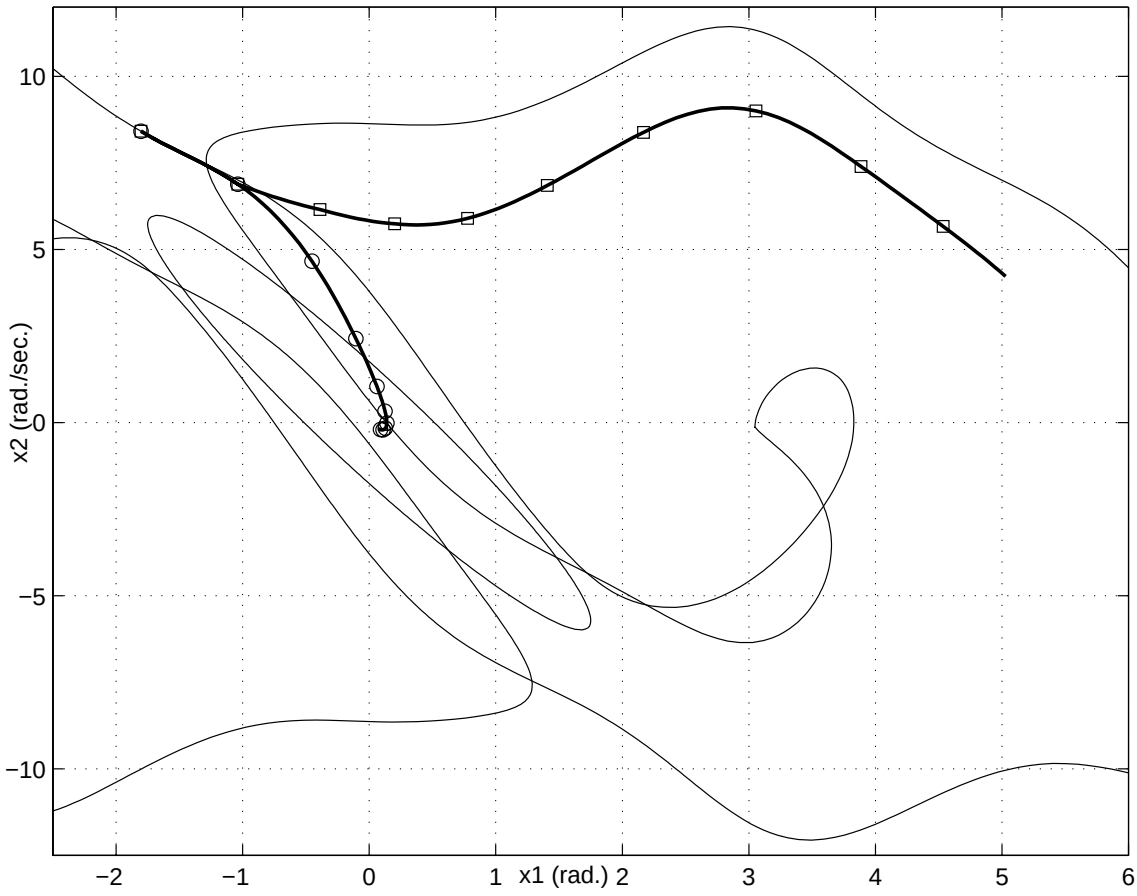

Figure 3.3: Receding horizon trajectories using $\delta = 0.1$ beginning at $x(0) = (-1.803, 8.413)$. At $x(\delta) = (-1.038, 6.887)$, local minima with costs of 36.973 (square) and 37.323 (circle) are found providing two different strategies.

seconds, we arrive at $x(\delta) = (-1.038, 6.887)$ and find two local minima with values 36.973 and 37.323, offering two potential strategies. As both costs are less than r_v^2, it is clear that either course will result in convergence to the inverted equilibrium. The resulting trajectories are shown in figure 3.3. The evolution of the costs is shown in figure 3.4 verifying its decreasing nature as well as the possibility of discrete jumps, indicating strict inequality in (3.1).

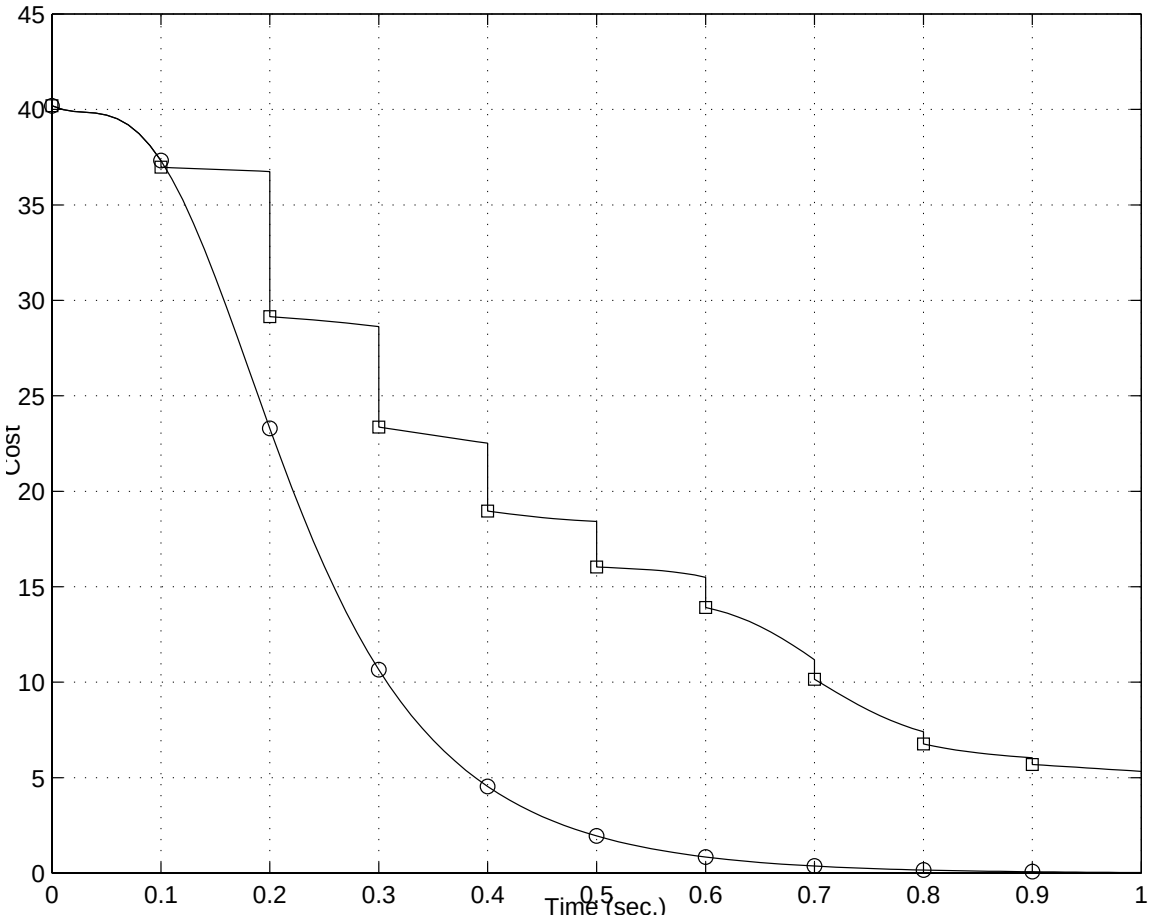

Figure 3.4: Evolution of the cost for the two strategies shown in figure 3.3.

This phenomenon is surprising, since in this special case, the *optimal* solution is not the *desired* one.

3.3 Summary

The purpose of this Chapter was to extend the results of Chapter 2 to the case were the optimizations are not solved exactly. Instead of requiring the receding horizon trajectories to be optimal in each iteration, a certain decrease in the value of the cost was required. Furthermore, it was shown that there always exists a controller which provides the proper decrease. The control trajectory consisted of two parts; the first part was the tail of the trajectory obtained from the previous iteration and the second was a feedback obtained from the CLF. A numerical example using the inverted pendulum compared the use of locally optimal and globally optimal trajectories. Simulations indicated that there are regions in which more than one locally optimal trajectory exist and both of them are stabilizing.

Chapter 4 Generating a suitable CLF

In this chapter, we address the problem of obtaining a CLF suitable for the receding horizon strategy outlined in Chapters 2 and 3. As was shown earlier, the required CLF has to be an upper bound on the cost-to-go, in the sense of (2.4). In the sequel, we discuss a variety of methods used to obtain such CLFs. Each method will be discussed with its advantages and disadvantages.

Despite their differences, these methods have one thing in common: They are all derived using *linear methods*. Other *nonlinear* methods such as differential flatness [FLMR95] and back-stepping [KKK95] can also be used to obtain suitable CLFs. Note that in the case of CLFs that are not an upper bound on the cost-to-go, i.e., the ones that do not satisfy (2.4), one can always use scaling by a large enough constant so that (2.4) is satisfied [JYH99b].

4.1 Jacobian linearization

Perhaps the simplest method for deriving a CLF for a nonlinear system is to use the Jacobian linearization of the system around the desired equilibrium point and generate a CLF by solving an LQR problem.

It is a well known result that the problem of minimizing the quadratic performance index:

$$\begin{aligned} \mathcal{J} &= \int_0^\infty (x^T(t)Qx(t) + u^T Ru(t))dt \\ \textbf{subject to:}\quad \dot{x} &= Ax + Bu \quad u = -Kx \end{aligned}$$

results in finding the positive definite solution of the following Riccati equation [DAC95]:

$$A^T P + PA - PBR^{-1}B^T P + Q = 0 \tag{4.1}$$

The optimal control action is given by

$$u = -R^{-1}B^T P x.$$

In the case of the nonlinear system $\dot{x} = f(x,u)$, A and B are assumed to be

$$A = \frac{\partial f(x,u)}{\partial x}|_{(0,0)} \quad B = \frac{\partial f(x,u)}{\partial x}|_{(0,0)}$$

where the pairs (A,B) and $(Q^{0.5},A)$ are assumed to be stabilizable and detectable respectively . Obviously the obtained CLF $V(x) = x^T P x$ will be valid only in a region around the equilibrium $(0,0)$. Therefore, if used as a terminal cost in the receding horizon scheme developed in Chapter 2, the terminal set Ω_{r_v} will be small in general, requiring longer horizon lengths to maintain stability without enforcement of terminal constraints. We will discuss this in more detail in the context of an example in Chapter 5.

4.2 Global linearization

The idea of global linearization has its roots in early works on the problem of absolute stability in the Soviet Union [LP44]. The basic idea behind this approach is to model a nonlinear system as a Polytopic Linear Differential Inclusion (PLDI) [BGFB94]. The dynamics of the nonlinear system are approximated to lie in the convex hull of a set of linear plants. The problem of quadratic stability of the obtained PLDI, i.e., stability provable by a quadratic Lyapunov function, is then recast as a convex optimization problem which can be solved very efficiently using interior point methods [NN94].

The PLDI describing the nonlinear model can be written as

$$\begin{aligned} \dot{x} &= \sum_{i=1}^{m} \alpha_i(t)(A_i x + B_i u) \\ u &= -Kx \\ \alpha_i(t) &\geq 0 \quad \sum_{i=1}^{m} \alpha_i(t) = 1. \end{aligned}$$

Using a quadratic cost function such as (4.1), the problem of minimizing an upper bound on the cost $\mathcal{J}$ can be written as the following convex optimization problem:

Minimize

$$tr(Z)$$

Subject to:

$$\begin{array}{rcl}
Y & > & 0 \\
\begin{bmatrix} YA_i^T + A_iY - B_iX - X^TB_i^T & YQ^{1/2} & X^TR^{1/2} \\ Q^{1/2}Y & -I & 0 \\ R^{1/2}X & 0 & -I \end{bmatrix} & < & 0 \\
\begin{bmatrix} Z & I \\ I & Y \end{bmatrix} & > & 0 \\
i & = & 1, \ldots, m
\end{array}$$

where Z is a slack variable, $Y = P^{-1}$, and $X = KY$ are the change of variables made to recast the matrix inequalities as LMIs [BGFB94]. Q and R are performance index matrices, with $Q \geq 0$ and $R > 0$. The initial conditions are assumed to be random vectors with zero mean and an identity covariance.

This is a very powerful method for obtaining suitable CLFs for nonlinear systems. However, it can be quite conservative, since there are many trajectories that are a trajectory of the PLDI, but are not a trajectory of the nonlinear system. Furthermore, the above LMI conditions verify quadratic stability, which is stability provable by a quadratic Lyapunov function. Another advantage of this method is that robustness can be addressed in this context. Both parametric uncertainties and unmodeled dynamics can be addressed in this framework [BGFB94]. A successful application of this strategy was illustrated in Section 2.5.

4.3 Frozen Riccati Equation (FRE) method

This method was first introduced by Cloutier *et al.* in [CDM96]. The basic idea behind this method, sometimes called *State Dependent Riccati Equations*, is to solve a Riccati equation pointwise, along the trajectory of the nonlinear system. A state dependent formulation of the nonlinear system is obtained, and nonlinearity of the representation at each point in state space is totally ignored. Although results are often promising, there are no rigorous justifications for even maintaining mere stability. Nevertheless, the simplicity of the implementation makes the FRE approach a plausible alternative in some applications. To apply this method, the nonlinear system is written in the following state dependent form:

$$\dot{x} = A(x)x + g(x)u \tag{4.2}$$

At each *frozen state* the Riccati equation is solved, and then the resulting state feedback controller is applied to the system, i.e., the state feedback nonlinear control law is obtained by solving the following:

$$\begin{aligned} 0 &= A(x)^T P(x) + P(x)A(x) - P(x)g(x)g^T(x)P(x) + Q \\ u &= -g^T(x)P(x)x \end{aligned} \tag{4.3}$$

The quantity $V(x) = x^T P(x)x$ generated by this technique is in general only a local CLF.

One of the major drawbacks of this method is the lack of a systematic procedure for selecting, among the infinite possibilities, a single parameterization for $f(x)$ (in the form of equation (4.2)) which achieves stability and acceptable performance [HJ99]. The CLF obtained using this approach has been used successfully in a receding horizon scheme [SCH+00].

4.4 Linear Parameter Varying (LPV) methods

This method was originally developed to design controllers for linear systems with *varying parameters*, hence the term LPV was coined [SA90, SA91, SA92]. It was assumed that the varying parameter $\rho(t)$ is available for measurement in real time, and also that it lied in a compact region with a bound on its variations known *a priori*. Although these conditions limit the application of this strategy, it has been successfully used in the aerospace industry. A variant of this approach, known as the *quasi* LPV scheme, was used to design controllers for *nonlinear* systems. The main ideas are identical, except for the fact that the varying parameters are a subset of the states, on which the nonlinearities in the state dependent representation depend. The so-called quasi-LPV representation of a nonlinear input-affine system can be written as follows:

$$\dot{x} = A(\rho(x))x + B(\rho(x))u \tag{4.4}$$

Assume the underlying parameter ρ varies in the allowable set

$$\mathcal{F}_{\mathcal{P}}^{\nu} := \{\rho \in \mathcal{C}^1(\mathbb{R}^+, \mathbb{R}^m) : \rho \in \mathcal{P},\ \underline{\nu}_i(\rho) \leq \dot{\rho}_i \leq \overline{\nu}_i(\rho),\ i = 1, \cdots, m\} \tag{4.5}$$

where $\mathcal{P} \subset \mathbb{R}^m$ is a compact set. If there exists a positive definite $X(\rho)$ such that the following inequality is satisfied

$$\begin{bmatrix} -\sum_{i=1}^{m} \overline{\nu}_i(\rho)\frac{\partial X}{\partial \rho_i} + A(\rho)X(\rho) + X(\rho)A^T(\rho) - B(\rho)R^{-1}B^T(\rho) & X(\rho)C^T(\rho) \\ C(\rho)X(\rho) & -I \end{bmatrix} < 0 \tag{4.6}$$

for all $\rho \in \mathcal{P}$ where $C(\rho) = Q^{\frac{1}{2}}(\rho(x))$, then the closed-loop system is stable with the state feedback

$$u(x) = -R^{-1}B^T(\rho(x))X^{-1}(\rho(x))x.$$

Moreover, an upper bound on the optimal value function $V^*(x)$ (which also serves as a CLF) is given by

$$V(x) = x^T X^{-1}(\rho(x))x \geq V^*(x).$$

The notation $\sum_{i=1}^{m} \underline{\overline{\nu}}_i(\rho)$ in (4.6) means that every combination of $\overline{\nu}_i(\rho)$ and $\underline{\nu}_i(\rho)$ should be included in the inequality. For instance, when $m = 2$, $\overline{\nu}_1(\rho) + \overline{\nu}_2(\rho)$, $\overline{\nu}_1(\rho) + \underline{\nu}_2(\rho)$, $\underline{\nu}_1(\rho) + \overline{\nu}_2(\rho)$ and $\underline{\nu}_1(\rho) + \underline{\nu}_2(\rho)$ should be checked individually. In other words, (4.6) actually represents 2^m inequalities. Additionally, solving (4.6) involves gridding the parameter space $\mathcal{P}$ and choosing a finite set of basis for $X(\rho)$ (see [WYPB96] for details). As can be easily seen, the number of LMIs grows exponentially with the number of varying parameters. Furthermore, the gridding technique can be quite costly. Due to these restrictions this method is usually suitable when there are only a few varying parameters (states) in the state dependent representation.

From a theoretical point of view, in order to guarantee stability, gridding has to be finer than a critical limit [WYPB96]. However, finding this limit requires knowledge of some information about the nonlinear system which is usually not available. Hence it is not practical to use. This makes the gridding phase an *ad hoc* procedure.

Similar to the FRE approach, the non-uniqueness of the state dependent representation of the nonlinear system can make the results conservative. However, due to the fact that bounds on the parameter variations are taken into account, stability arguments are more powerful than regular gain scheduling or the FRE method. The non-uniqueness issue of the state dependent representation can be properly addressed by including an additional degree of freedom in the optimizations [HJ99, Hua98].

The quasi-LPV control techniques have been quite successful in a variety of aerospace applications [BFL+97]. Since the main application area of this thesis is intended to be the aerospace industry, we choose the LPV approaches to be our method of choice for obtaining a CLF.

4.5 Summary

The purpose of this chapter was to discuss several methods of generating a CLF, suitable for receding horizon purposes. Following the results of [Pri98], it was shown that while these methods generate their own controllers, it is beneficial to use the CLF, rather than the controller, to be used in the receding horizon schemes developed earlier in this thesis. The methods discussed are only a few of the methods available in the literature. The discussed methods were all developed from a linear setting, and are suitable for generating CLFs that are useful in the context of the results of Chapters 2 and 3.

Chapter 5 The Caltech ducted fan

The Caltech ducted fan is an experimental testbed developed to facilitate the study of Uninhabited Combat Aerial Vehicles (UCAVs). These highly maneuverable aircrafts are designed to operate without having a pilot on board. The pilot will be in a ground or airborne based control room far from the actual vehicle. Using the state of the art virtual reality technology as well as proper control schemes is believed to make UCAVs a reality. The main advantage of using such vehicles is that it takes the pilots out of harm's way.

As was mentioned in the introduction, one of the tasks at DARPA's Software Enabled Control Program, is Online Control and Customization (OCC) for UCAVs. A viable candidate for OCC has been receding horizon control strategies. In this chapter, we deal with two models of the ducted fan experiment, one around hover and one at forward flight. The RHC method developed in this thesis is successfully applied to these models [JYH99a].

5.1 Ducted fan at hover

The Caltech Ducted Fan is a small flight control experiment whose dynamics are representative of a VTOL (Vertical Landing and Take off) aircraft such as Harrier in hover mode or a thrust vectored aircraft such as F18-HARV or X-31 in forward flight [Mur98]. This system has been used for a number of studies and papers. In particular, a comparison of several linear and nonlinear controllers was performed in [KBPM95, BBK96, NM96]. In this section we describe the simple planar model of the fan shown in Figure 5.1. This model is useful for initial controller design and would serve as a good testbed for purposes of this thesis.

Let (x, y, θ) denote the position and orientation of a point on the main axis of the fan. We assume that the forces acting on the fan consist of a force f_1 perpendicular

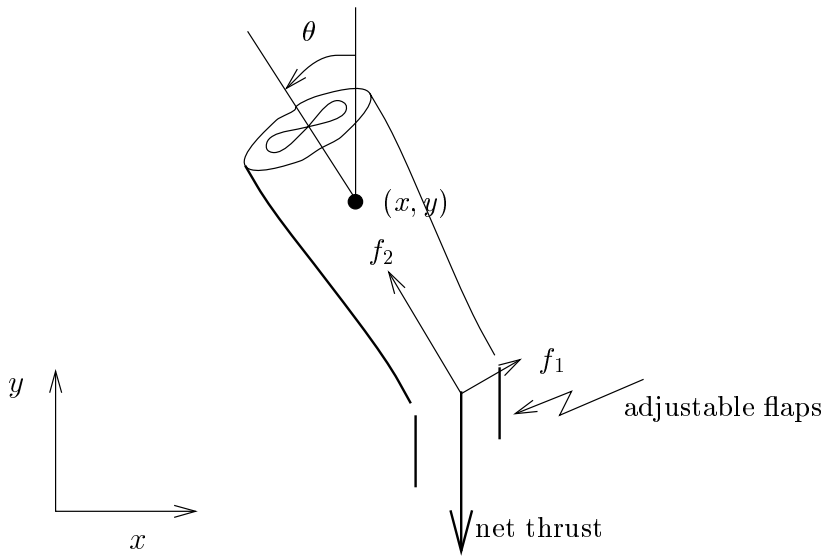

Figure 5.1: Schematics of the planar ducted fan.

to the axis of the fan acting at a distance r and a force f_2 parallel to the axis of the fan. Assuming m, J, and g to be the mass of the fan, the moment of inertia, and the gravitational constant respectively, the equations of motion can be written as follows [HSM92]:

$$\begin{aligned} m\ddot{x} &= -d\dot{x} + f_1 \cos\theta - f_2 \sin\theta \\ m\ddot{y} &= -d\dot{y} + f_1 \sin\theta + f_2 \cos\theta - mg \\ J\ddot{\theta} &= r f_1 \end{aligned} \tag{5.1}$$

where the drag terms are modeled as viscous friction with d being the viscous friction coefficient. The numerical data for the experiment are as follows: $m = 11.2\ kg$, $g = 0.28\ m/sec^2$, $J = 0.0462\ kg\ m^2$, $r = 0.156\ m$, $d = 0.1\ N\ sec$. Note that due to the use of a counterweight for balancing the fan, the gravitational constant is not $9.8\ m/s^2$. It is convenient to redefine the inputs so that the origin is an equilibrium point of the system with zero input. If we let $u_1 = f_1$ and $u_2 = f_2 - mg$, the equations of motion can be written as:

$$\begin{aligned} m\ddot{x} &= -mg\sin\theta - d\dot{x} + u_1\cos\theta - u_2\sin\theta \\ m\ddot{y} &= mg(\cos\theta - 1) - d\dot{y} + u_1\sin\theta + u_2\cos\theta \\ J\ddot{\theta} &= ru_1. \end{aligned} \tag{5.2}$$

These equations are referred to as the *planar ducted fan equations.* The following quadratic performance index was used in the quasi-LPV scheme:

$$\mathcal{J} = \int_0^\infty (\bar{x}^T(t)Q\bar{x}(t) + u^T(t)Ru(t))dt$$

where $\bar{x} = [x\ \dot{x}\ y\ \dot{y}\ \theta\ \dot{\theta}]^T$, R and Q is chosen to be diagonal matrices with the following diagonal terms:

$$Q = diag(\begin{bmatrix} 0.5 & 1 & 5 & 1 & 5 & 1 \end{bmatrix}) \tag{5.3}$$

$$R = diag(\begin{bmatrix} 2 & 1 \end{bmatrix}) \tag{5.4}$$

To obtain a CLF for the ducted fan using quasi LPV methods, we have to choose a set of varying parameters and bounds on parameter variation in case we are searching for parameter dependent Lyapunov functions.

For the ducted fan, $\rho = \theta$ is chosen as the varying parameter, and the operation range $\mathcal{P} = [-\frac{\pi}{2}, \frac{\pi}{2}]$. As a first attempt, we tried to find a quadratic CLF so that no *a priori* bound on parameter variation is needed. Unfortunately, the resulting convex program is not feasible; therefore, we switch to searching for parameter dependent Lyapunov functions with θ as the varying parameter. The bound on the rate variation on θ is set to be 10, i.e., $|\dot{\theta}| \leq 10$. Obvious parameterizations of f and g are given by

setting $A(\theta)$ and $B(\theta)$ as follows [JYH99b, JYH99a, YJPH99]:

$$A(\theta) = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & -\frac{d}{m} & 0 & 0 & -\frac{g\sin\theta}{\theta} & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -\frac{d}{m} & \frac{g(\cos\theta-1)}{\theta} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$B(\theta) = \begin{bmatrix} 0 & 0 \\ \frac{\cos\theta}{m} & -\frac{\sin\theta}{m} \\ 0 & 0 \\ \frac{\sin\theta}{m} & \frac{\cos\theta}{m} \\ 0 & 0 \\ \frac{r}{j} & 0 \end{bmatrix} \tag{5.5}$$

A set of polynomial basis functions are chosen to compute $X(\rho)$, i.e., $X(\rho) = \sum_{i=1}^{5} c_i(\rho) X_i$ where X_i's are symmetric coefficient matrices (LMI variables) and $\{c_i(\rho)\}$ are fifth order Legendre polynomials on $\mathcal{P}$:

$$\begin{aligned} \{c_i(\rho)\} = \{&1,\ \frac{2}{\pi}\theta,\ (3(\frac{2}{\pi}\theta)^2 - 1)/2,\ (5(\frac{2}{\pi}\theta)^3 - 3(\frac{2}{\pi}\theta))/2, \\ &(35(\frac{2}{\pi}\theta)^4 - 30(\frac{2}{\pi}\theta)^2 + 3)/8\} \end{aligned}$$

Once the CLF is obtained, it can be used as the terminal cost in a receding horizon optimization. Also, as it was mentioned earlier, if the horizon length is long enough, the local CLF obtained by solving an LQR problem for the linearized dynamics can also be used. In the next section, we discuss some simulation results for the planar ducted fan model, using CLFs from LQR and LPV as a terminal cost in the receding horizon scheme.

5.2 Simulation results

The following initial condition was chosen

$$\begin{bmatrix} x & \dot{x} & y & \dot{y} & \theta & \dot{\theta} \end{bmatrix} = \begin{bmatrix} -1 & 2.5 & 0 & 0 & -45^\circ & 0 \end{bmatrix}.$$

Simulations were performed for several different horizon lengths. Also, in order to have a measure for the performance, the "optimal" control trajectories were calculated by solving finite horizon optimization problems for long horizon lengths. It can be seen from the simulation plots that a simple LQR controller does not stabilize the system, whereas the LPV controller does. This is not surprising since the region of attraction of the LQR controller is much smaller.

As we include on-line optimization in the process, it can be seen that when a simple quadratic CLF from LQR is chosen as the terminal cost, the controller is still not stabilizing the system. Figure 5.2 depicts $x - y$ trajectories of the ducted fan for four receding horizon iterations with different horizon lengths of 0.05, 0.1, 0.15, and 0.2 seconds. As shown in Figure 5.2, the terminal point is far from the origin, hence the resulting receding horizon controller is not stabilizing. For long enough horizon lengths, however, the $x - y$ trajectories converge to the origin, and the resulting receding horizon controller stabilizes the closed-loop system, as shown in Figure 5.3. This is due to the fact that the set Ω_{r_V}, defined in Section 2.1, gets larger as the horizon length is increased. Note that the trajectories also converge to the optimal ones as T approaches infinity (see Figure 5.3).

Figure 5.3 shows that if the horizon length is long enough, the CLF form LQR is a suitable choice for the terminal cost. However, if the horizon length is short, the receding horizon with CLF from LQR method does not give a good performance and in fact the closed-loop system is unstable.

The $x-y$ trajectories for several receding horizon iterations with different horizon lengths with CLF from LPV is depicted in Figure 5.4. The initial condition chosen was the same as in (5.6). Note the difference in the orientation of the fan in the case of

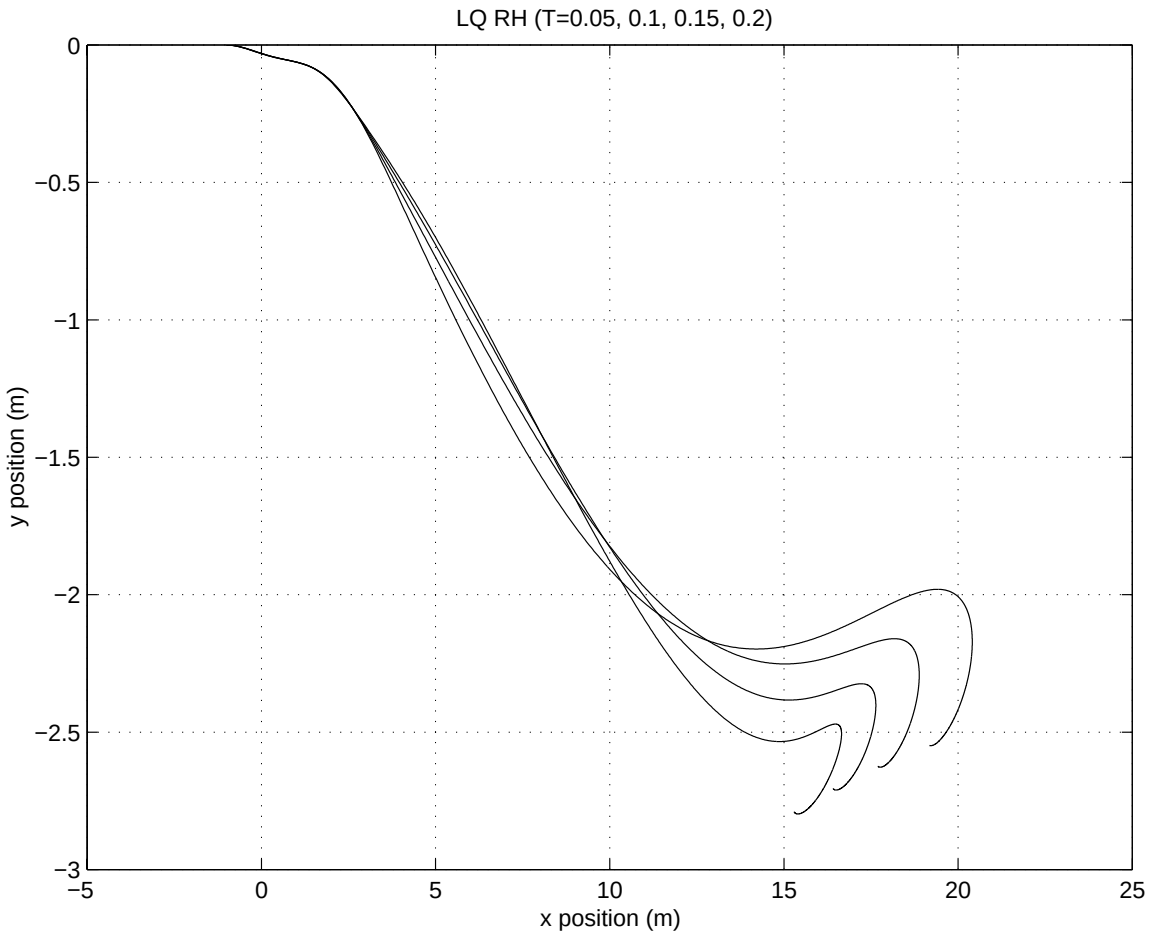

Figure 5.2: $x-y$ position of the fan for four one shot trajectories with different horizon lengths, from 0.05 *sec* to 0.2 second. The resulting receding horizon controllers are not stabilizing.

LQR and LPV terminal costs. Also, note that the CLF from LPV is stabilizing with a 0 horizon length which corresponds to no online computation at all, i.e., applying the LPV controller.

One might think that since stability is guaranteed by having a CLF as terminal cost, any positive definite function lower bounded by the CLF would also be a legitimate choice. Simulation results indicate that this is not the case, and merely putting a quadratic penalty which is greater than the value of the CLF (and therefore an upper bound on the cost-to-go) is not sufficient. Figure 5.5 depicts the simulation results for the receding horizon scheme with having a penalty of the form of $\gamma||x(t+T)||^2$ where $\gamma > 0$ is chosen such that $\gamma||x||^2 > V(x)\ \forall\ x$. Shown in Figure 5.5 are again

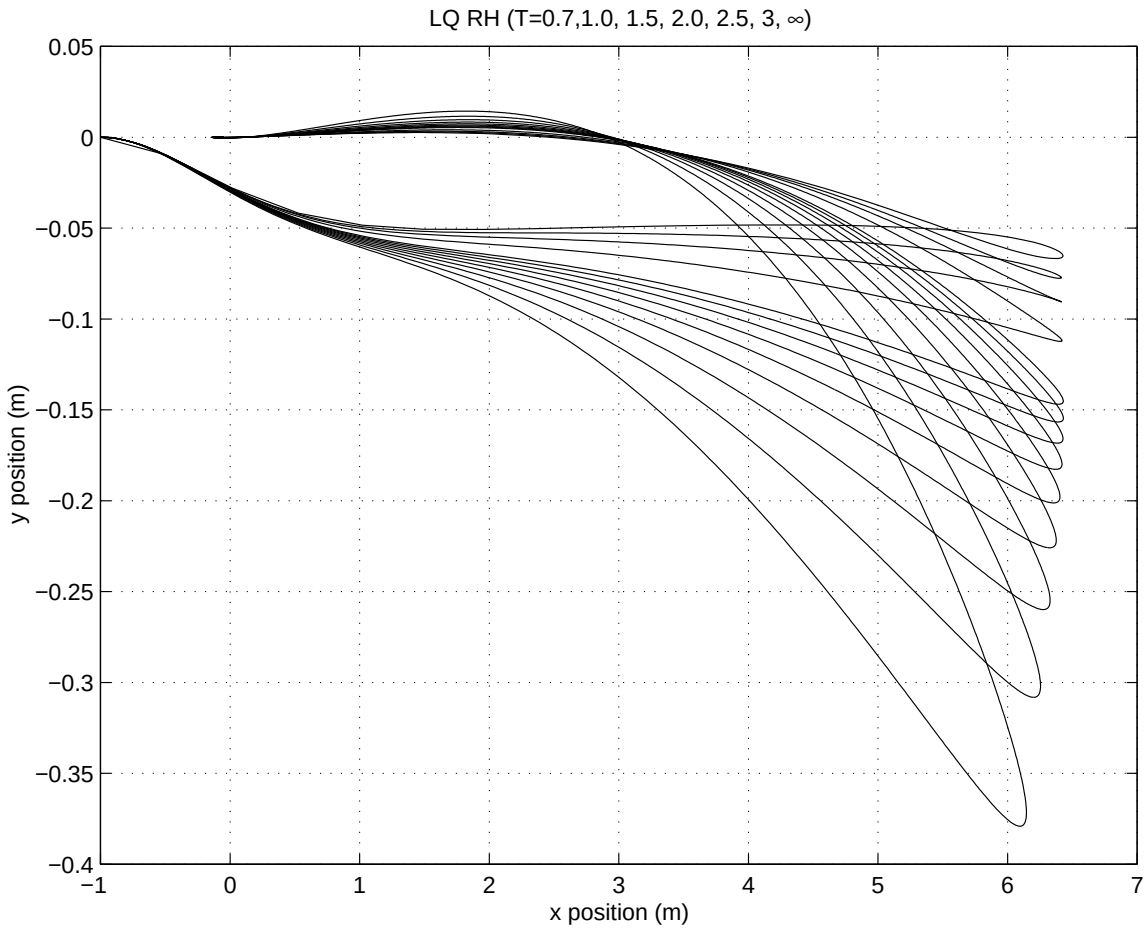

Figure 5.3: $x - y$ position of the fan for RH with quadratic CLF.

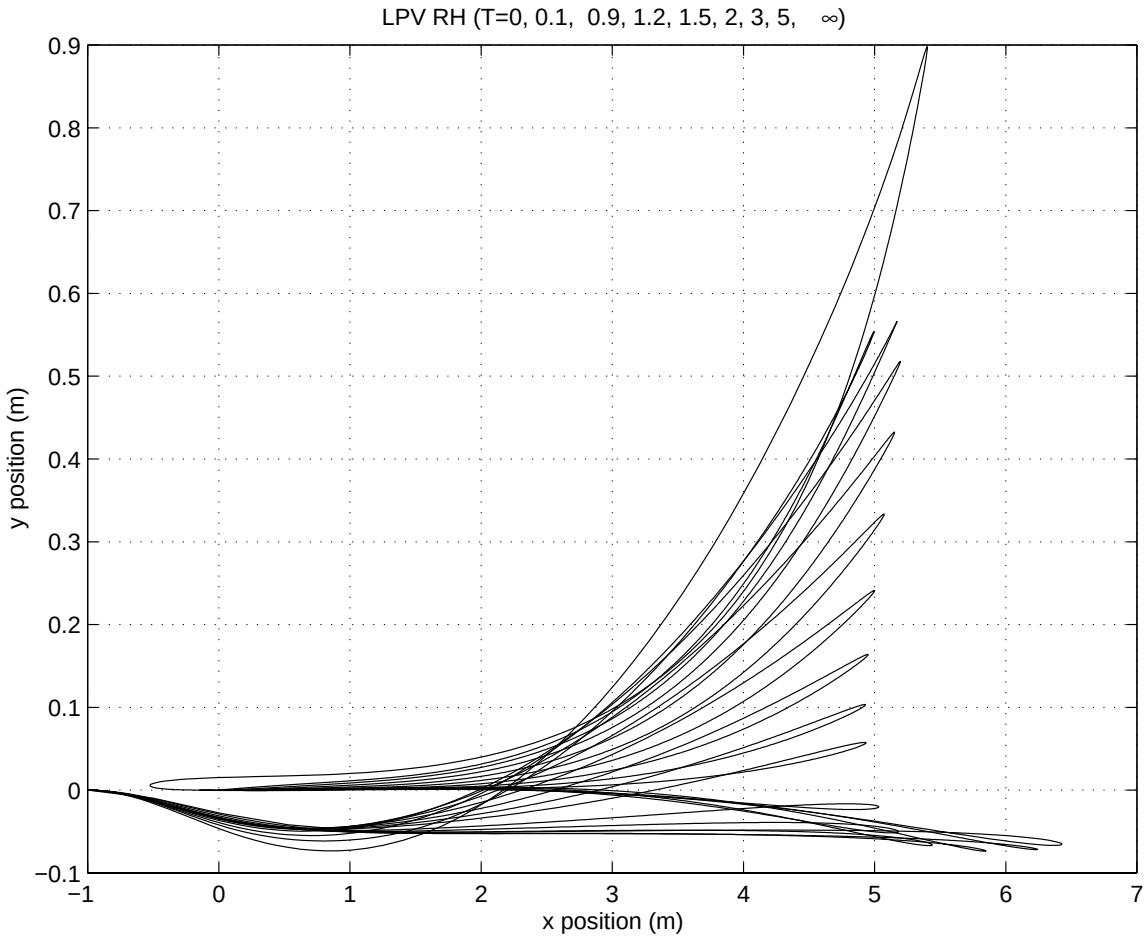

Figure 5.4: $x - y$ position of the fan for receding horizon with the CLF from LPV.

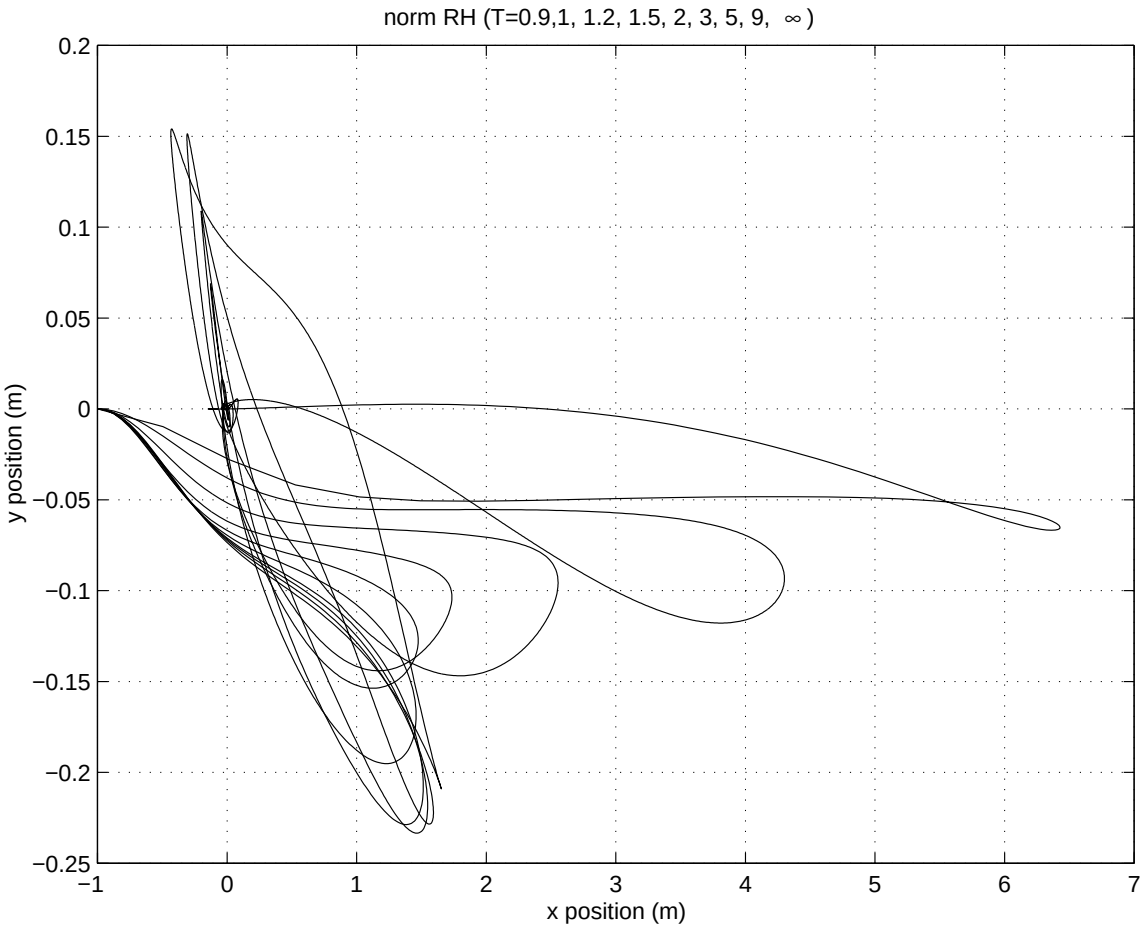

Figure 5.5: $x - y$ position of the fan for receding horizon with $\gamma||x||^2$ as penalty.

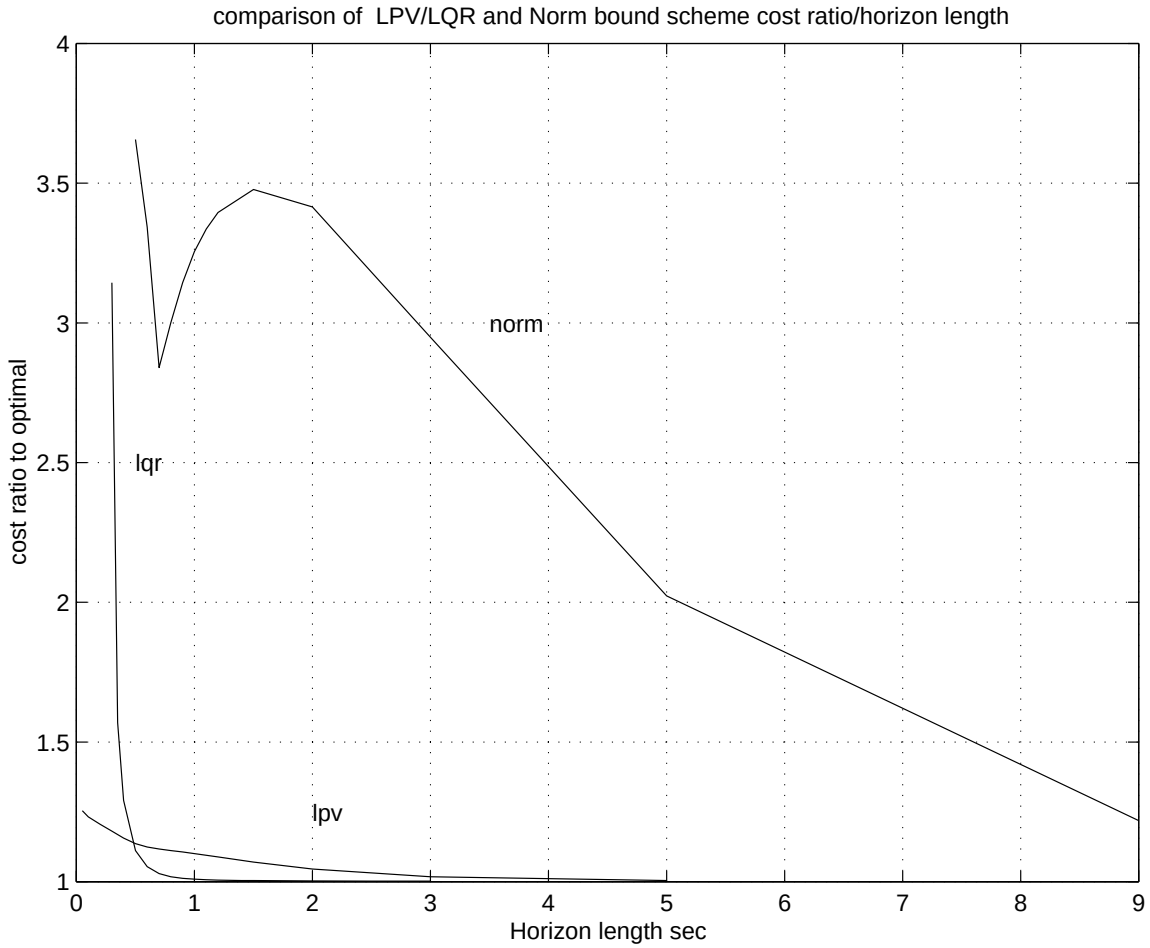

Figure 5.6: Ratio of the cost to the optimal vs. the horizon length.

$x - y$ trajectories for various choices of T. Note that even for long horizon lengths, the performance is not satisfactory at all.

An interesting comparison between the three discussed choices for the terminal cost is shown in Figure 5.6. The horizon length is plotted on the horizontal axis, and the ratio of the actual cost to the optimal cost is plotted on the vertical axis. Note the fast convergence of the cost to the optimal, when the CLF from LQR is chosen as the terminal cost. Interestingly, the choice of CLF from LPV as terminal cost starts at a much lower cost but longer horizon length is needed to converge to the optimal one.

In fact, for horizon lengths of $T \geq 0.6$, using the CLF from LQR results in a lower cost due to the fact that the optimal value function is well approximated by the CLF from LQR method, in a suitable neighborhood of the origin. We can therefore conclude that the CLF from LPV methods is suitable when long optimization horizons are not possible due to costly computation. Finally, having a terminal cost of the form of a quadratic upper bound on the cost-to-go does not work even for very long horizon lengths.

5.3 Ducted fan in forward flight

In the previous section, we dealt with a simplified model of the ducted fan around hover. The fan was modeled as a rigid body and the aerodynamic forces were totally ignored. In forward flight, however, one can no longer ignore the aerodynamic forces and moments, since they are the dominant forces that make the ducted fan fly.

An activity to build up aerodynamic models for the ducted fan was established at Caltech [MM99, Hau00]. However, obtaining steady state flight data has proved to be much more challenging than expected. This is mainly due to a periodic drag force resulting from the interaction of a wing tip vortex with the walls in the lab. (See Figures 5.7 and 5.8 for a view of the experimental setup.)

A first step in obtaining suitable models is to familiarize oneself with the flight of this vehicle by doing a significant amount of manual flight. The aggressive ca-

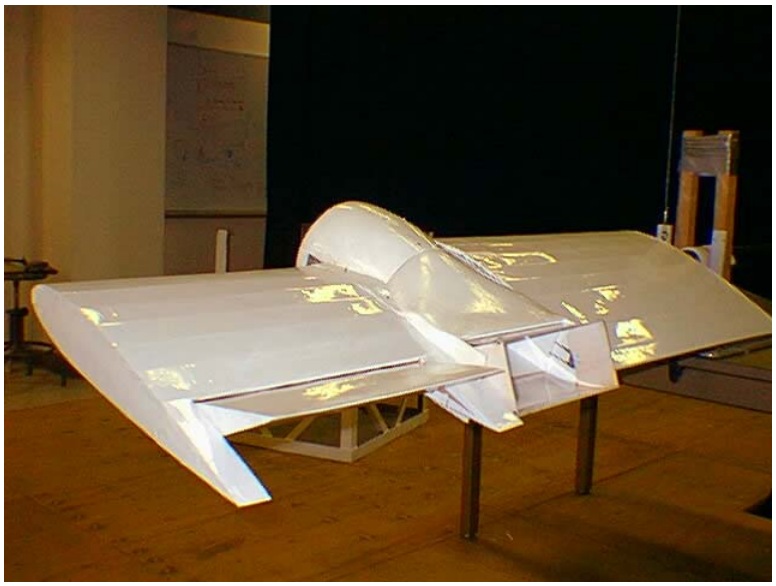

Figure 5.7: The wing has been re-designed to improve the thrust vectoring.

pabilities of the system have been explored through *manual* flip and turn around maneuvers [MM99, HJ00, Tro00]. Some of these *high angle of attack* maneuvers are depicted in Figure 5.9. Angle of attacks of close to 70 degrees are reported in these experiments [HJ00].

In spite of the problems arising from the periodic drag force, the results look quite promising, especially for the purpose of obtaining a model that captures the *essential features* of the system.

After getting a general picture of the capabilities of the system, an important step in modeling is to obtain the experimental *equilibrium manifold*. The equilibrium manifold is the surface resulting from setting the right-hand side of the equations of motion to zero. Experimentally, points on this surface can be found by performing steady flights at different velocities. Corresponding to each velocity, there is an angle of attack as well as the thrust force and thrust angle that would keep the vehicle in that equilibrium point. Figures 5.10 and 5.11 depict the experimental and the model equilibrium manifolds.

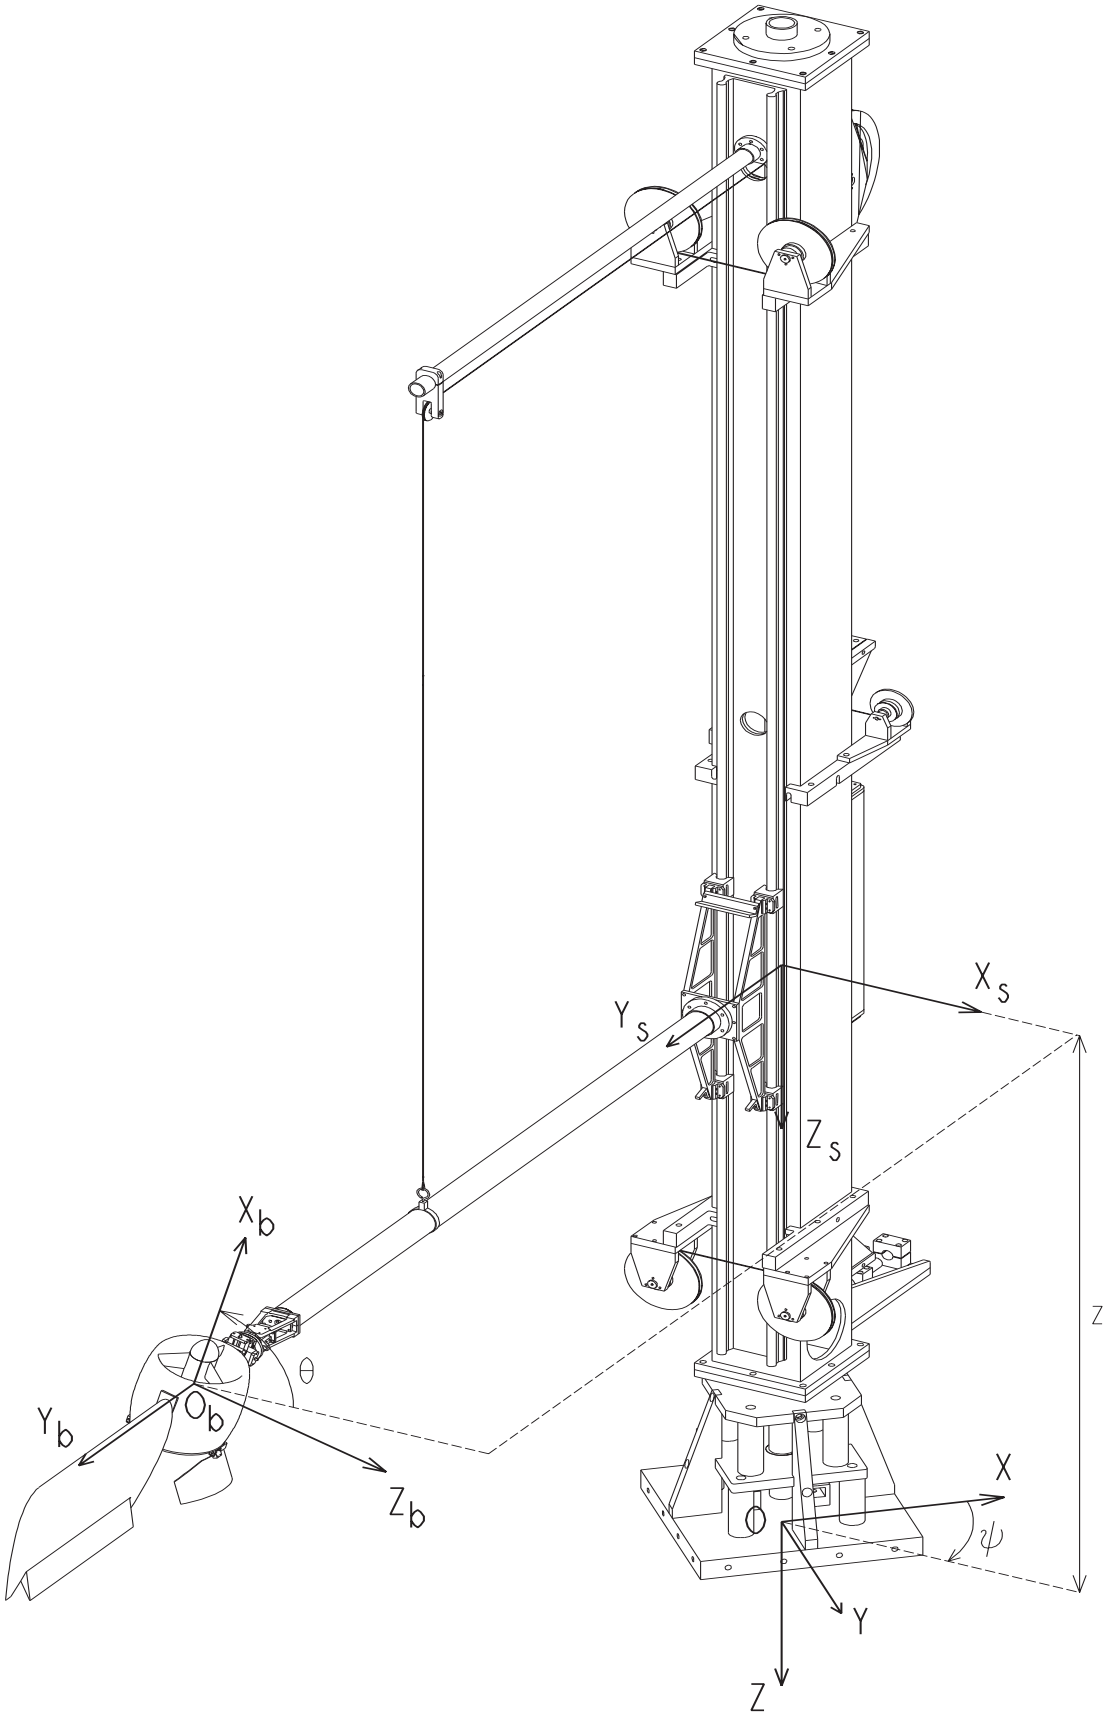

Figure 5.8: The Caltech ducted fan [MM99].

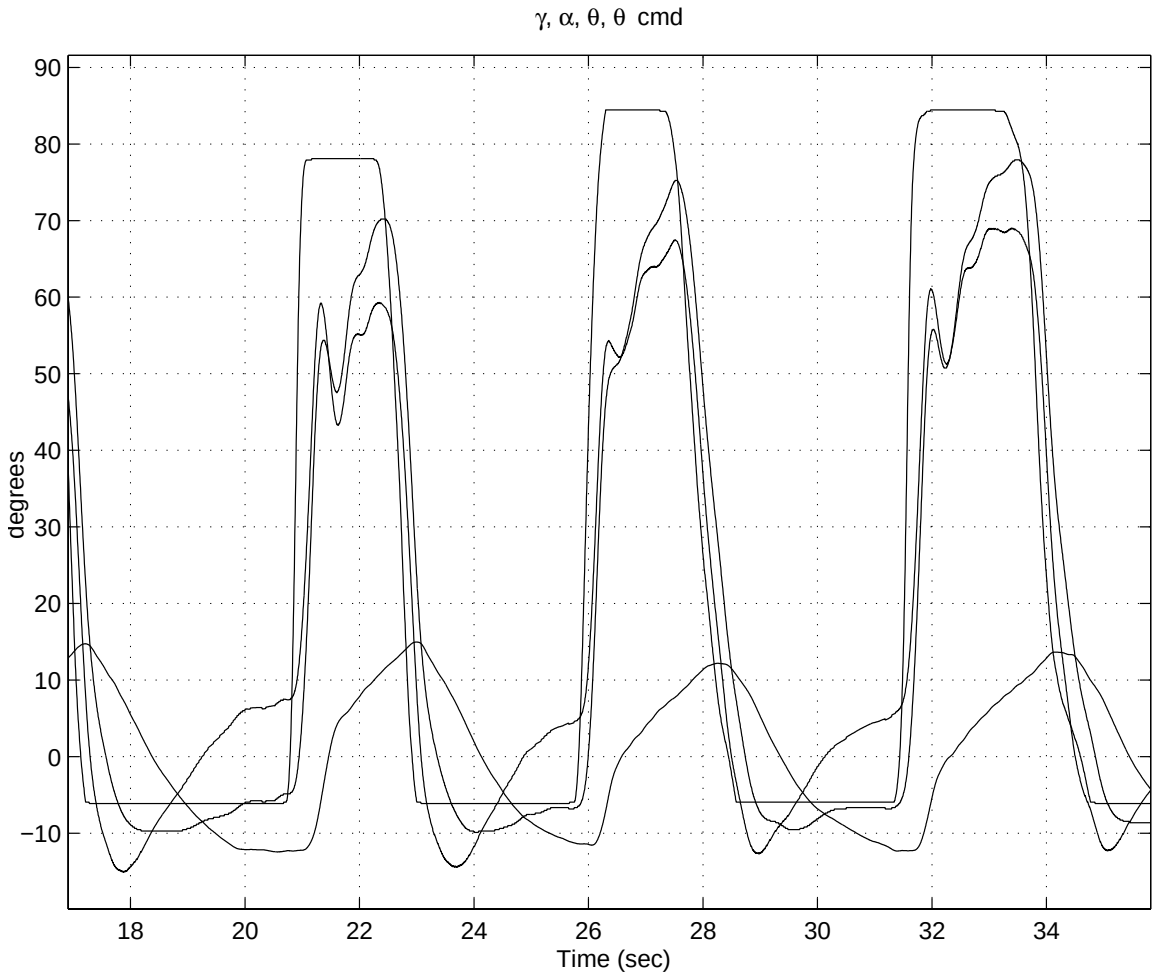

Figure 5.9: High angle of attack maneuvering of the Caltech ducted fan. Shown here are the commanded pitch θ_{cmd}, α, θ and γ.

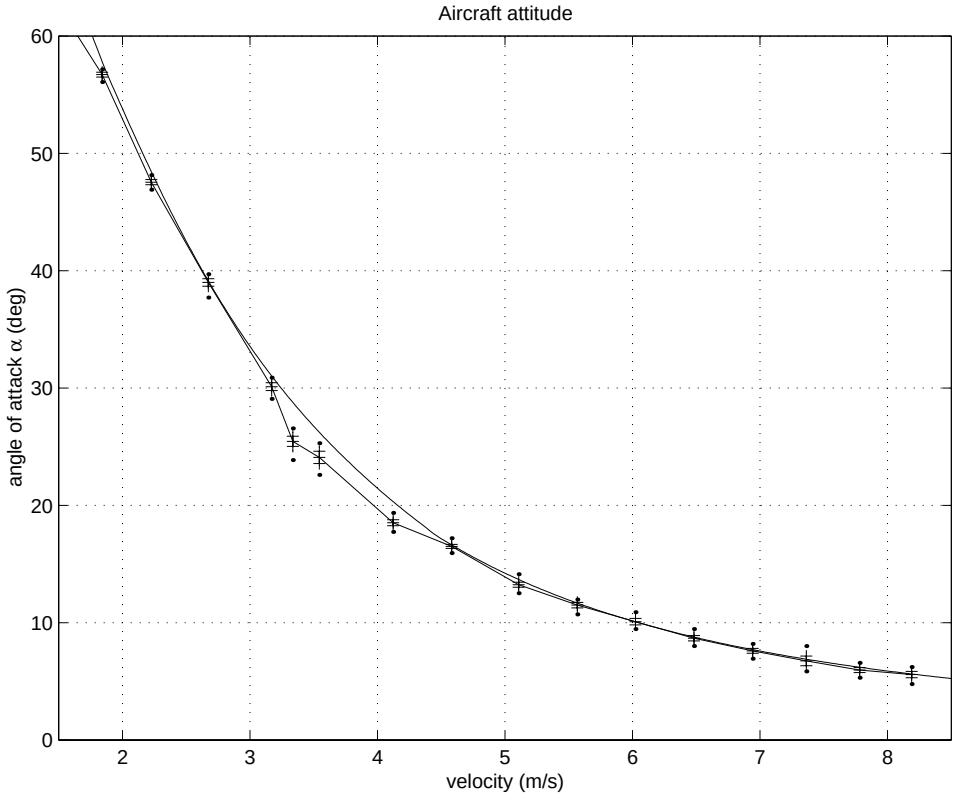

Figure 5.10: Required angle of attack vs. velocity for points on the model and experimental equilibrium manifold. For each flight condition, the average, standard deviation, and the range of values (minimum and maximum) are plotted [HJ00].

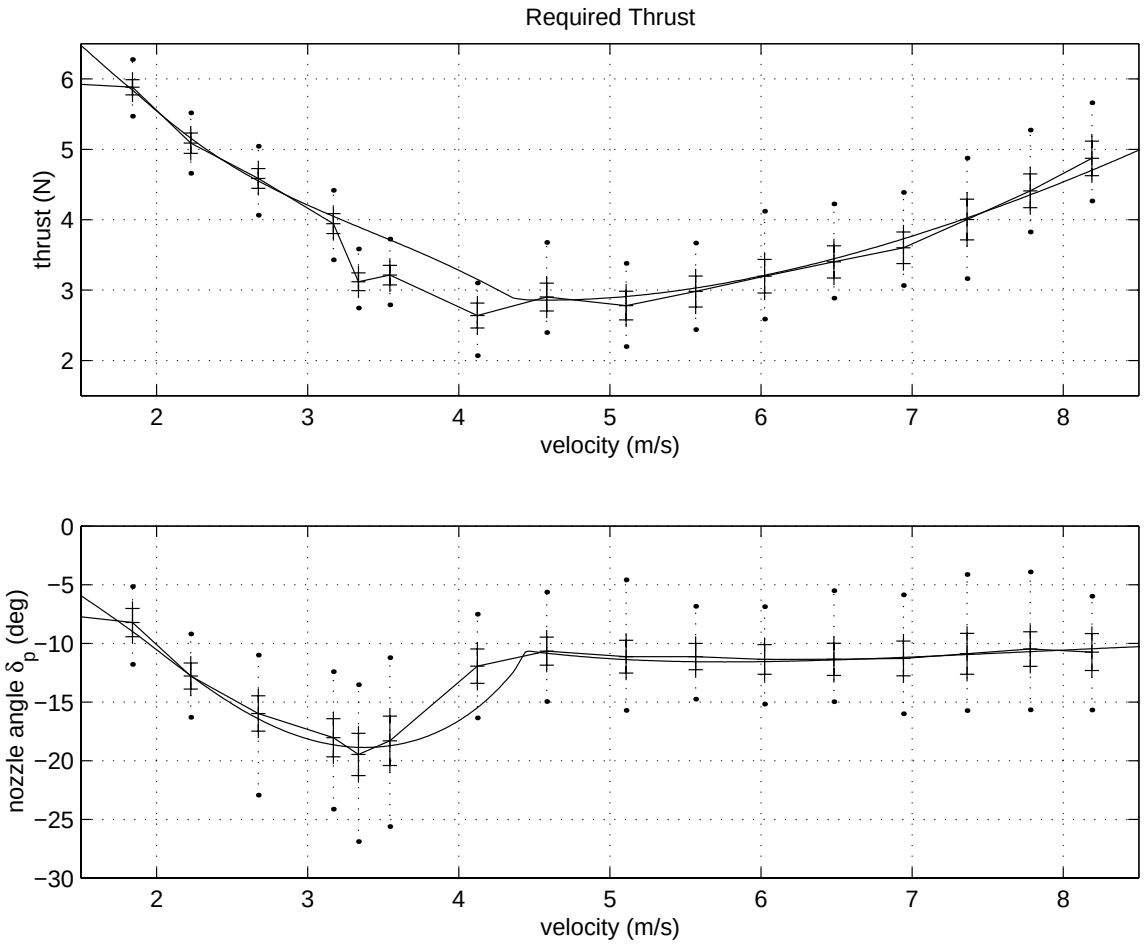

Figure 5.11: Required thrust T and paddle (nozzle) angle δ_p at different velocities for points on the equilibrium manifold. For each flight condition, the average, standard deviation, and the range of values (minimum and maximum) are plotted.

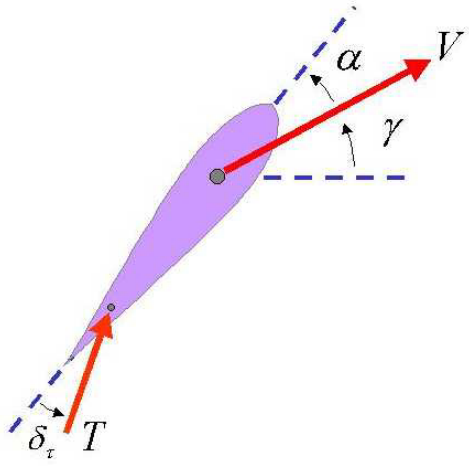

Figure 5.12: Schematics of the thrust vectored flying wing.

5.3.1 Modeling

The ducted fan is modeled as a flying wing with vectored thrust, as in Figure 5.12. The effect of the newly designed thrust vectoring nozzle is modeled as a vectored force applied at a fixed position on the flight vehicle. The parameters for this model were estimated using static force/moment measurements. Using the experimental equilibrium manifold, a plausible aerodynamic model was obtained.

The equations of motion for the thrust vectored flying wing of Figure 5.12 are as follows:

$$\begin{aligned} m\dot{V} &= -D(V,\alpha) - W\sin\gamma + T\cos(\alpha+\delta_T) \\ mV\dot{\gamma} &= L(V,\alpha) - W\cos\gamma + T\sin(\alpha+\delta_T) \\ \dot{\theta} &= q \\ J\dot{q} &= M(V,\alpha) - Tl_T\sin(\delta_T) \end{aligned} \tag{5.6}$$

where T is the thrust force, and δ_T is the angle at which thrust is applied. It will be natural to take as state and control variables $x = (V, \gamma, q, \alpha)$, where $\gamma = \theta - \alpha$ is the *flight path angle*, $q = \dot{\theta}$, and $u = (T, \delta_T)$.

The physical parameters for this model are $m = 12\ kg$, $g = 0.6\ m/s^2$, $S = 0.61\ m^2$, $\rho = 1.2\ kg/m^3$, $l_T = 0.31\ m$, $J = 0.24kg\ m^2$. The lift, drag and moment terms can

be written as follows:

$$\begin{aligned} L(V,\alpha) &= \frac{1}{2}\rho V^2 S C_L(\alpha) \\ D(V,\alpha) &= \frac{1}{2}\rho V^2 S C_D(\alpha) \\ M(V,\alpha) &= \frac{1}{2}\rho V^2 S \bar{c} C_M(\alpha). \end{aligned} \tag{5.7}$$

Where $\bar{c} = 0.5\ m$, and $C_L(\alpha),\ C_D(\alpha),\ C_M(\alpha)$ are the lift, drag and moment coefficients respectively. Using the experimental data from the equilibrium manifold as well as estimates of the thrust and thrust angle given in Figures 5.10 and 5.11, one can come up with plausible estimates for the lift, drag and moment coefficients as depicted in Figures 5.13 and 5.14.

From the discontinuity in the lift curve, one can notice that stall occurs between 16 and 20 degrees angle of attack. More noticeable is the change in pitching moment at these angle of attacks. Stall can also be noticed from Figure 5.11 which depicts the thrust angle vs. velocity for the equilibrium manifold.

An interesting point worth mentioning is that one would not see the usual drop in the lift curve after the wing stalls. This is probably due to the fact that the fuselage housing the fan starts to act as a lifting body at these angle of attacks. Worthwhile noting is the fact that the lift, drag, and moment curves fit nicely with the equilibrium manifold data. Locally (i.e., up to stall), lift, drag and moment curves are expressed by the following equations as functions of the angle of attack in radians.

$$\begin{aligned} C_L(\alpha) &= C_{l_\alpha}\alpha = 3.256\alpha \\ C_D(\alpha) &= C_{d_0} + C_{d_\alpha}\alpha^2 = 0.1716 + 2.395\alpha^2 \\ C_M(\alpha) &= C_{M_\alpha}\alpha = -0.0999\alpha \end{aligned} \tag{5.8}$$

Roughly speaking, the thrust can range from 0 to 13.5 N and can be vectored a little more than 25 degrees, i.e., $\delta_T \leq 0.45$ radians.

While the input variables in the model are the thrust force T and the thrust angle δ_T, the control inputs in the experiment are the motor voltage V_m and the commanded

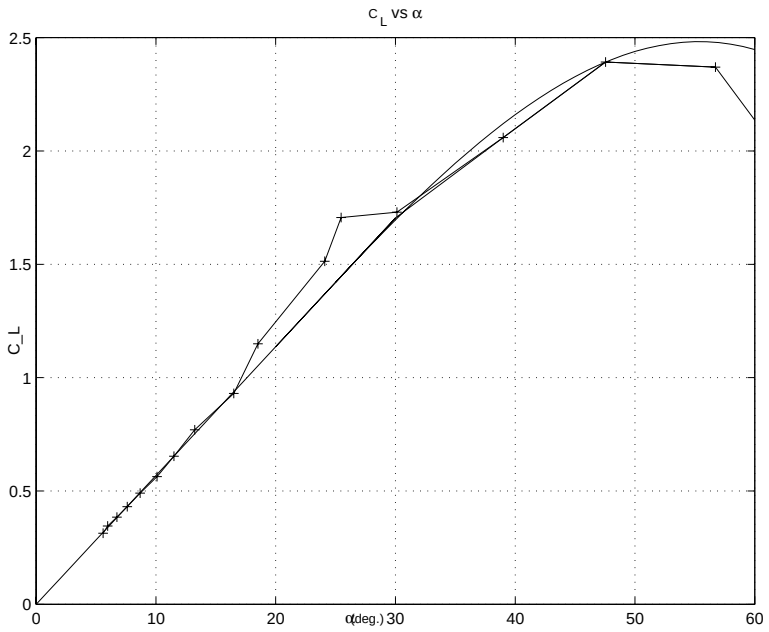

Figure 5.13: Experimental and model curves for the lift coefficient $C_L(\alpha)$.

paddle angle δ_p. We have the following relationship between the two pair of inputs:

$$\begin{aligned} \delta_T &= 0.684\delta_p \\ T &= 46.5V_m - 5. \end{aligned} \tag{5.9}$$

Overall, the above model is a plausible description for the Caltech ducted fan, although much work remains in the understanding of the nature of uncertainties, both external disturbances and model uncertainties.

5.3.2 Control of the ducted fan in forward flight

So far, we have discussed a model of the Caltech ducted fan which is suitable for control purposes.

Following the discussion in Chapter 4, we now develop a quasi LPV scheme to generate a control Lyapunov function suitable for receding horizon control. The objective of the controller is to regulate the ducted fan at a specific point on the equilibrium manifold.

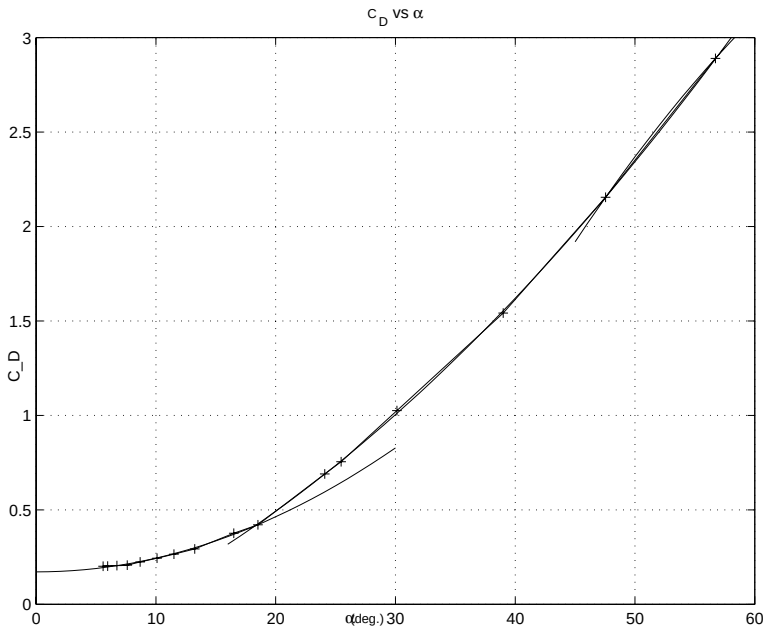

Figure 5.14: Experimental and model curves for the drag coefficient $C_D(\alpha)$.

Our objective is to regulate the ducted fan at $V = 6\ m/s$. Using Figures 5.10 and 5.11, the corresponding angle of attack, thrust, and thrust angle are as follows: $\alpha_0 = \theta_0 = 10.149°,\ \ T_0 = 3.2005\ N, \delta_{T_0} = -7.9°$.

In order to be able to use the quasi-LPV scheme to generate a CLF, we first have to do a coordinate translation so that the desired equilibrium is the origin of the translated system. The next step is to choose a set of varying parameters. Contrary to the planar model where only the pitch angle θ was the varying parameter, we need to choose three terms. We would need 2 out of the three angles $\alpha,\ \theta,\ \gamma$ as well as the velocity V.

For convenience, α, θ, and V are picked. Note that the new states are $\bar{V} = V - V_0$, $\bar{\alpha} = \alpha - \alpha_0$, and $\bar{\theta} = \theta - \theta_0$. Also, $\bar{q} = q = \dot{\theta}$. As can be seen from (5.7), the model is not *affine* in control, hence not suitable for use in the LPV scheme. In order to avoid this problem, the following change of variables are made, and forces in the x and z direction are chosen as inputs: $F_x := T\cos(\delta_T)\ \ F_z := T\sin(\delta_T)$. We can now describe the equations of motion in the new translated coordinates as follows:

$$\dot{\bar{x}} = A(\bar{V}, \bar{\alpha}, \bar{\theta})\bar{x} + B(\bar{V}, \bar{\alpha}, \bar{\theta})\bar{u} \tag{5.10}$$

where

$$\bar{u} := u - u_0 = \begin{bmatrix} \bar{F}_x \\ \bar{F}_z \end{bmatrix} = \begin{bmatrix} F_x - F_{x_0} \\ F_z - F_{z_0} \end{bmatrix}, \tag{5.11}$$

and

$$A(\bar{V},\bar{\alpha},\bar{\theta}) := \begin{bmatrix} A_{11} & A_{12} & A_{13} & 0 \\ A_{21} & A_{22} & A_{23} & 1 \\ 0 & 0 & 0 & 1 \\ A_{41} & A_{42} & 0 & 0 \end{bmatrix} \quad B(\bar{V},\bar{\alpha},\bar{\theta}) := \begin{bmatrix} \frac{\cos(\bar{\alpha}+\alpha_0)}{m} & \frac{\sin(\bar{\alpha}+\alpha_0)}{m} \\ -\frac{\sin(\bar{\alpha}+\alpha_0)}{m(\bar{V}+V_0)} & \frac{\cos(\bar{\alpha}+\alpha_0)}{m(\bar{V}+V_0)} \\ 0 & 0 \\ 0 & \frac{l_T}{J} \end{bmatrix}.$$

After a set of rather tedious calculations, the elements of the A matrix can be written as follows:

$$\begin{aligned}
A_{11} &= -\frac{1}{2}\frac{\rho\, S(\bar{V}+2\,V_0)(C_{d_0}+C_{d_\alpha}\,(\bar{\alpha}+\alpha_0)^2)}{m} \\
A_{12} &= -\frac{1}{2}\frac{\rho\, {V_0}^2 S C_{d_\alpha}\,(\bar{\alpha}+2\,\alpha_0)}{m} + g\sin(\frac{\bar{\theta}-\bar{\alpha}}{\pi})\pi(\bar{\theta}-\bar{\alpha})^{-1} + \\
&\quad \frac{1}{2}\,(\frac{\cos(\alpha_0+\frac{1}{2}\,\bar{\alpha})F_{z_0}}{m} - \frac{\sin(\alpha_0+\frac{1}{2}\,\bar{\alpha})F_{x_0}}{m})\sin(\frac{1}{2}\,\frac{\bar{\alpha}}{\pi})\alpha^{-1}\pi \\
A_{13} &= -g\sin(\frac{\bar{\theta}-\bar{\alpha}}{\pi})\pi(\bar{\theta}-\bar{\alpha})^{-1} \\
A_{21} &= -\frac{1}{2}\frac{\rho\, S C_{l_\alpha}\,(\bar{\alpha}+\alpha_0)}{m} - \frac{g}{V_0\,(\bar{V}+V_0)} + \frac{F_{x_0}\,\sin(\alpha_0) - F_{z_0}\,\cos(\alpha_0)}{m V_0\,(\bar{V}+V_0)} \\
A_{22} &= -\frac{g(\cos(\bar{\theta}-\bar{\alpha})-1)}{(\bar{\theta}-\bar{\alpha})(\bar{V}+V_0)} - \frac{1}{2}\,\frac{\rho\, S C_{l_\alpha} V_0}{m} \\
&\quad -2\,(\cos(\alpha_0+\frac{\bar{\alpha}}{2})F_{x_0} + \sin(\alpha_0+\frac{\bar{\alpha}}{2})F_{z_0})\sin(\frac{\alpha}{2\pi})\frac{\pi}{\alpha}(m(\bar{V}+V_0))^{-1} \\
A_{23} &= \frac{g(\cos(\bar{\theta}-\bar{\alpha})-1)}{(\bar{\theta}-\bar{\alpha})(\bar{V}+V_0)} \\
A_{41} &= \frac{1}{2}\,\frac{\rho\,(\bar{V}+2\,V_0) S\bar{c}\, C_{M_\alpha}\,(\bar{\alpha}+\alpha_0)}{J} \\
A_{42} &= \frac{1}{2}\,\frac{\rho\, {V_0}^2 S\bar{c}\, C_{M_\alpha}}{J}.
\end{aligned} \tag{5.12}$$

Once the equations of motion are represented in the appropriate state dependent form, the next step is to formulate a suitable performance index to be used in the

quasi-LPV scheme.

It is well known that choosing a suitable performance index such that all desired aspects of the problem are taken into account, can be quite difficult. We know that in order to regulate the ducted fan to a point on the equilibrium manifold, the flight path angle must be regulated to 0 therefore, it would be appropriate to put a large weight on the flight path angle in the optimization criterion. Since in our quasi-LPV modeling we did not choose γ as the varying parameter, it would be convenient at this point to choose the flight path angle γ over the pitch angle θ to be a state. Hence we perform the following change of coordinates:

$$\begin{bmatrix} V \\ \gamma \\ \alpha \\ q \end{bmatrix} = T_\gamma \begin{bmatrix} V \\ \alpha \\ \theta \\ q \end{bmatrix} \qquad T_\gamma = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{5.13}$$

The new set of coordinates is called the *wind coordinates*. Also, the following Q and R matrices are chosen in conjunction with the quasi LPV model to generate a CLF.

$$\begin{aligned} Q &= diag\,[3\ 36\ 2\ 1] \\ R &= diag\,[.46251\ \ 0.98851] \times 10^{-3} \end{aligned} \tag{5.14}$$

The reason for such a choice for R is the relationship between the forces in the x and z directions and the actual control inputs V_m and δ_p.

To use the quasi LPV scheme discussed in Chapter 4, we first have to pick a gridding region for the varying parameters (states).

Since our simple aerodynamic model is valid locally, the following region is picked for gridding: $V \in [4, 8]$, and $\alpha,\ \theta \in [-8.149°, 28.149°]$.

Note that the simulation model [Hau00, MM99] exceeds this local region in terms of the aerodynamic modeling. For simplicity of the LPV design, we have chosen to use the above limited gridding region for CLF design.

Also, note that since the CLF is merely a stability safeguard to be used as terminal

cost in the receding horizon scheme, we would be quite happy with a quadratic CLF, provided the resulting convex optimization problem is feasible. This would greatly reduce the number of LMI constraints in equation (4.6). It turns out that the optimization problem is indeed feasible, and a quadratic CLF does exist. The resulting CLF is $V(x) = x^T P x$ where

$$P = \begin{bmatrix} 2.0162 & -0.3103 & -0.4178 & -0.0818 \\ -0.3103 & 43.6804 & 10.3974 & 1.2838 \\ -0.4178 & 10.3974 & 4.8598 & 0.5232 \\ -0.0818 & 1.2838 & 0.5232 & 0.2064 \end{bmatrix}. \tag{5.15}$$

The above CLF is then used in a receding horizon scheme with a horizon length of $T = 0.5$ and sampling rate of $\delta = 0.05$ seconds. Also, to compare the results with "optimal" trajectories, the horizon length is increased to 3 seconds.

An important issue is magnitude constraints imposed on the control action. As was mentioned earlier, the thrust can not exceed 13.5 Newtons, and can be only vectored for about 25 degrees. This would translate to the following bounds on the actual control actions V_m and δ_p.

$$\begin{aligned} 0.13\ Volts &\leq V_m \leq 0.35\ Volts \\ -40^\circ &\leq \delta_p \leq 40^\circ \end{aligned} \tag{5.16}$$

In order to simulate the system in the forward flight mode, the following initial condition was chosen: $V = 6\ m/s \quad \gamma = -40^\circ \quad \alpha = 10.2\ ^\circ, \quad q = 0$. The above initial condition represents a pull-up maneuver. Simulation results are depicted in Figures 5.15 and 5.16. Also, Figures 5.17 and 5.18 depict the one shot "optimal" trajectories for the same initial condition and a longer horizon length of 3 seconds.

A more interesting way of plotting trajectories is to show the actual trajectory of the flying wing in the x_z space instead of showing the time trajectories. These plots are shown in Figures 5.19 and 5.20 for the receding horizon and one shot trajectories respectively. In both figures, the ducted fan is shown as a flying wing, with a red

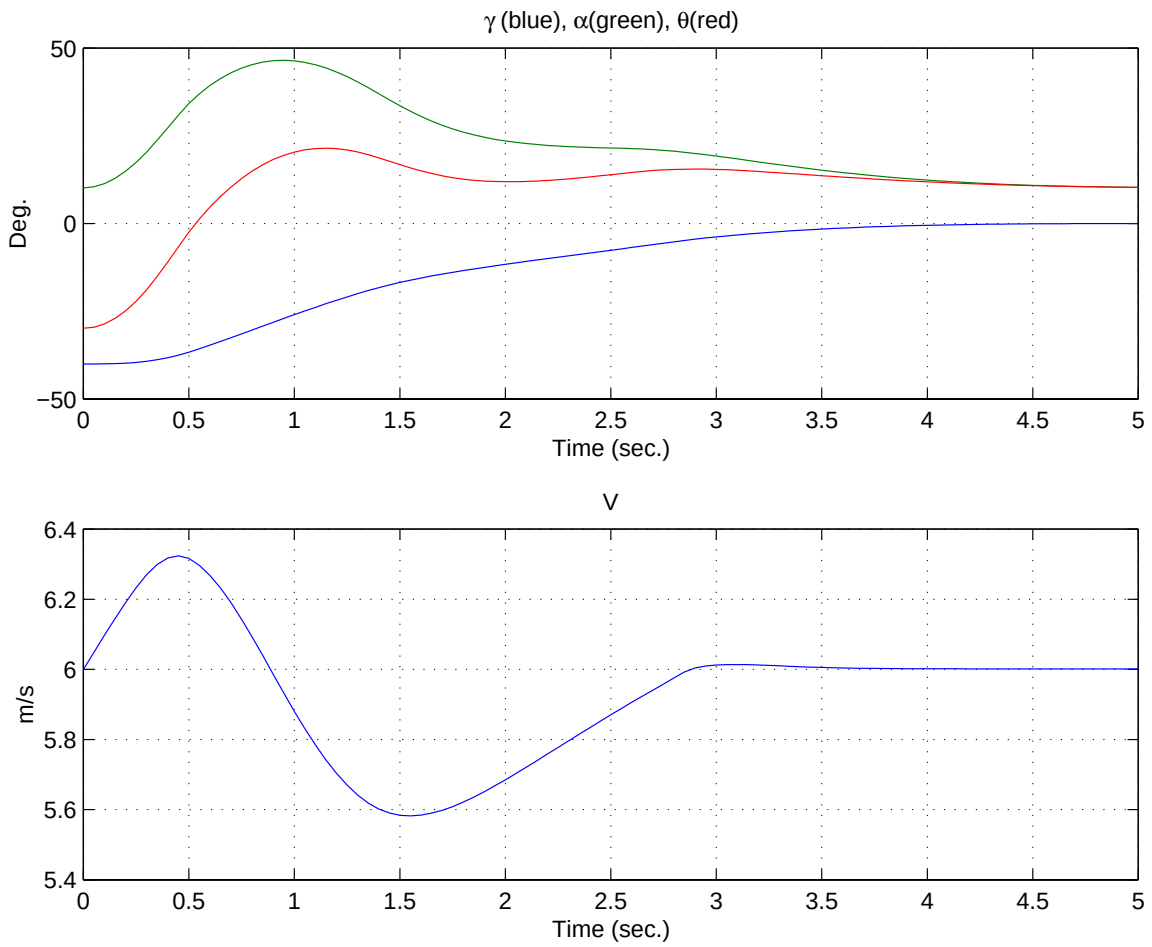

Figure 5.15: Simulation results for the ducted fan in forward flight with a receding horizon controller $\mathcal{RH}(0.5, 0.05)$. Shown here are the states $V,\ \gamma,\ \alpha,\ \theta$.

line representing the thrust. These figures indicate that with a horizon of only 0.5 seconds, we get very close to the optimal solution.

One can also include the z dynamics in the equations of motion by letting $\dot{z} = V \sin\gamma$. Using a similar LPV scheme and by penalizing z with a weight of 20, we obtain a quadratic Lyapunov function $V_z(x) = x^T P_z x$ with P_z given as follows:

$$P_z = \begin{bmatrix} 59.3330 & -1.8051 & -175.5690 & -39.2179 & -5.0851 \\ -1.8051 & 47.6105 & 11.1567 & -3.5507 & -0.6610 \\ -175.5690 & 11.1567 & 854.3814 & 176.9442 & 27.1237 \\ -39.2179 & -3.5507 & 176.9442 & 59.6604 & 7.3196 \\ -5.0851 & -0.6610 & 27.1237 & 7.3196 & 2.0231 \end{bmatrix}. \qquad (5.17)$$

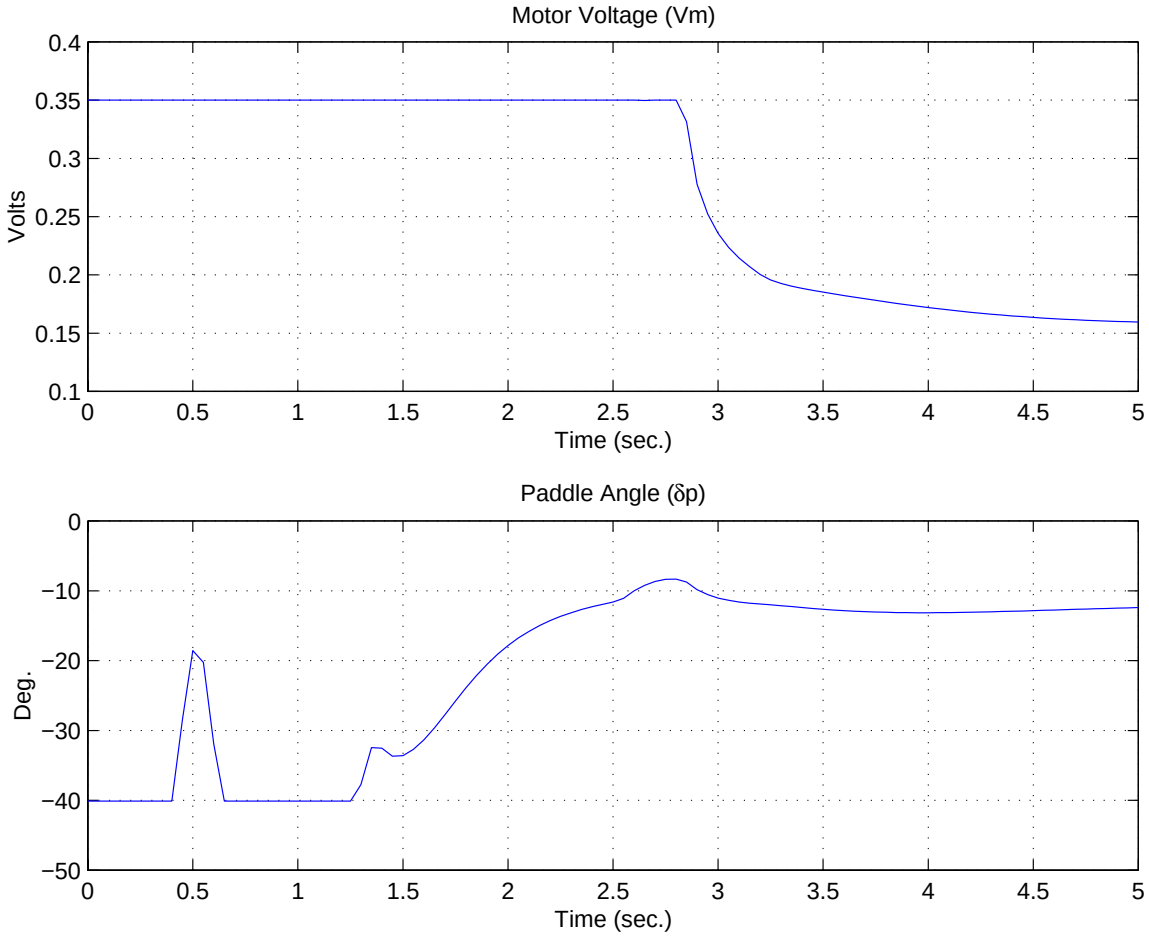

Figure 5.16: Simulation results for the ducted fan in forward flight with a receding horizon controller $\mathcal{RH}(0.5, 0.05)$. Shown here are the controls V_m and δ_p.

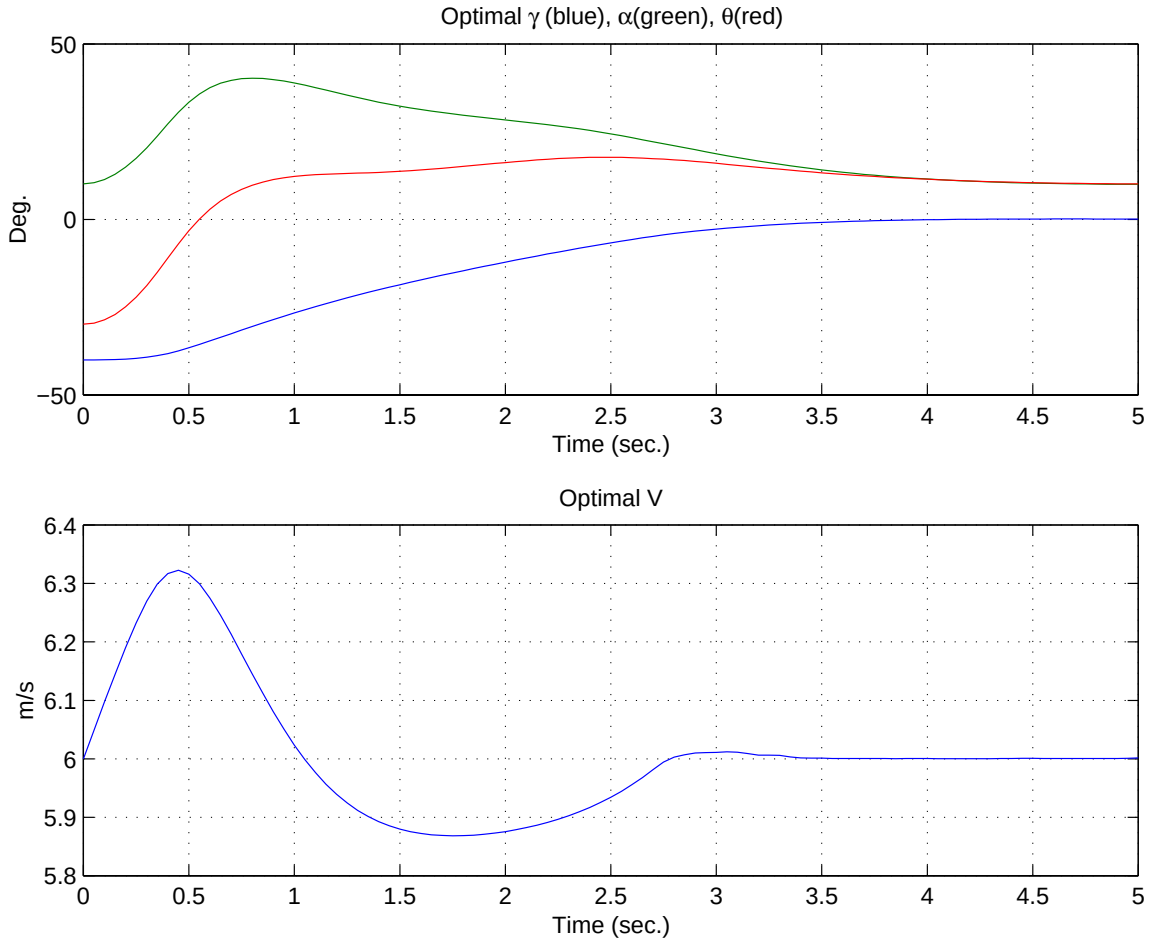

Figure 5.17: Simulation results for the ducted fan in forward flight with a one shot optimal controller $T = 3$ seconds. Shown here are the states V, γ, α, θ.

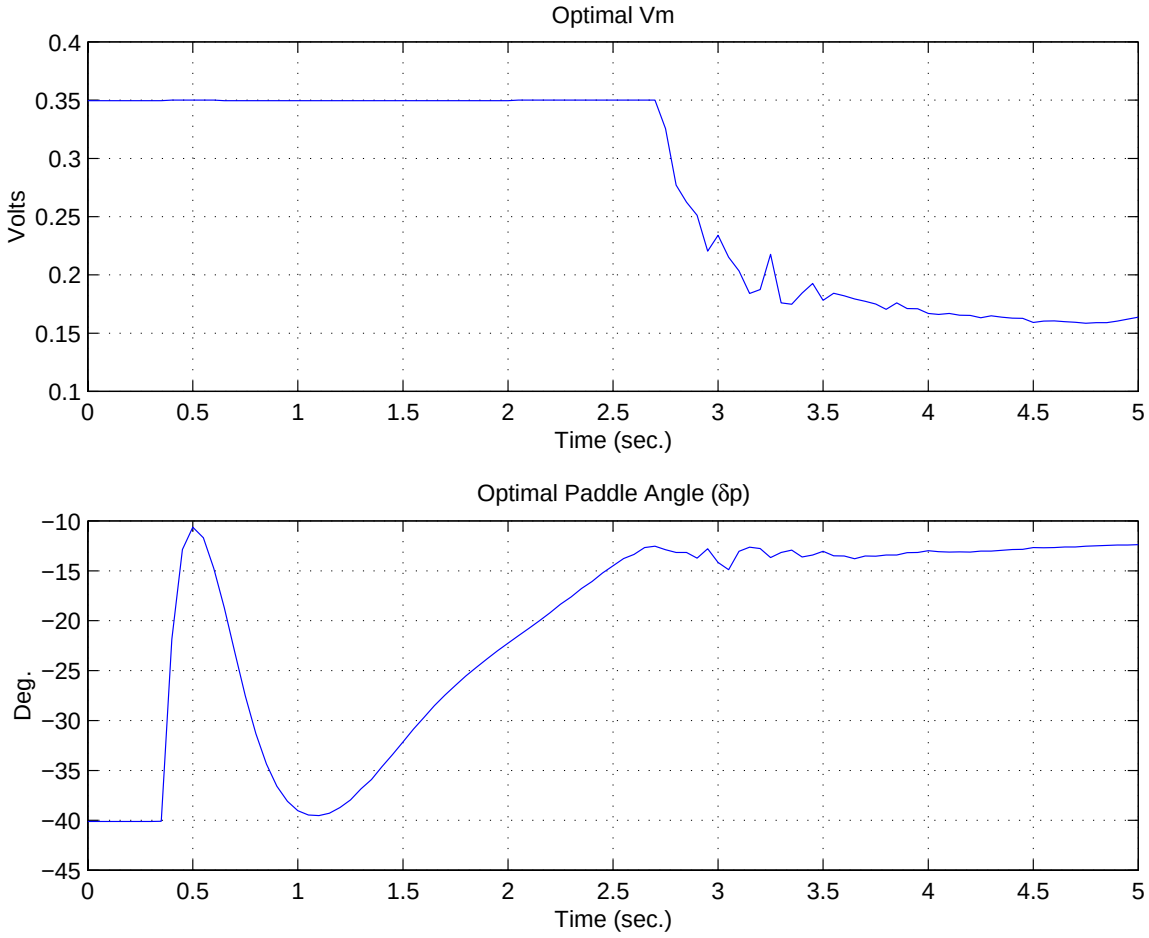

Figure 5.18: Simulation results for the ducted fan in forward flight with a one shot controller with $T = 3$ seconds. Shown here are the controls V_m and δ_p.

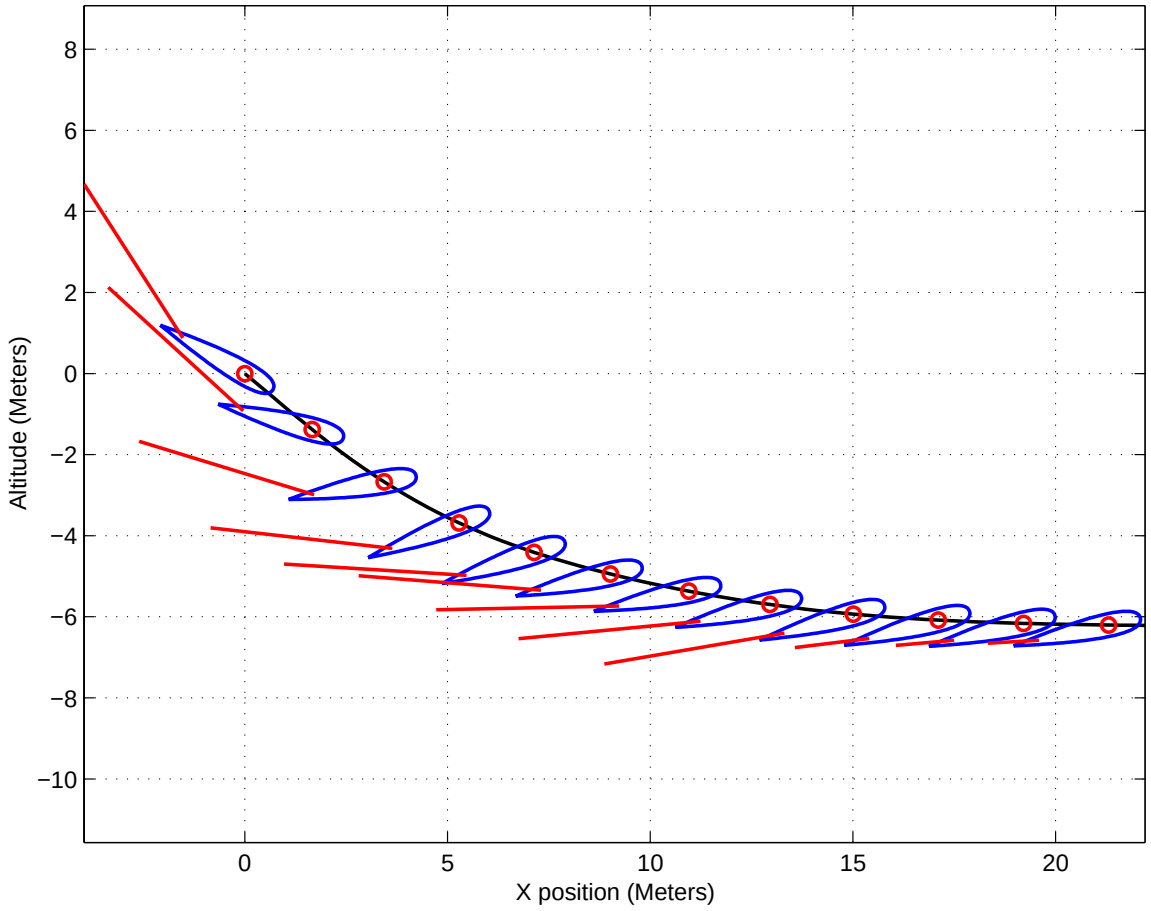

Figure 5.19: The $x-z$ trajectories for the ducted fan in forward flight with a receding horizon controller $\mathcal{RH}(0.5, 0.05)$. The red lines indicate the thrust, scaled according to the thrust value.

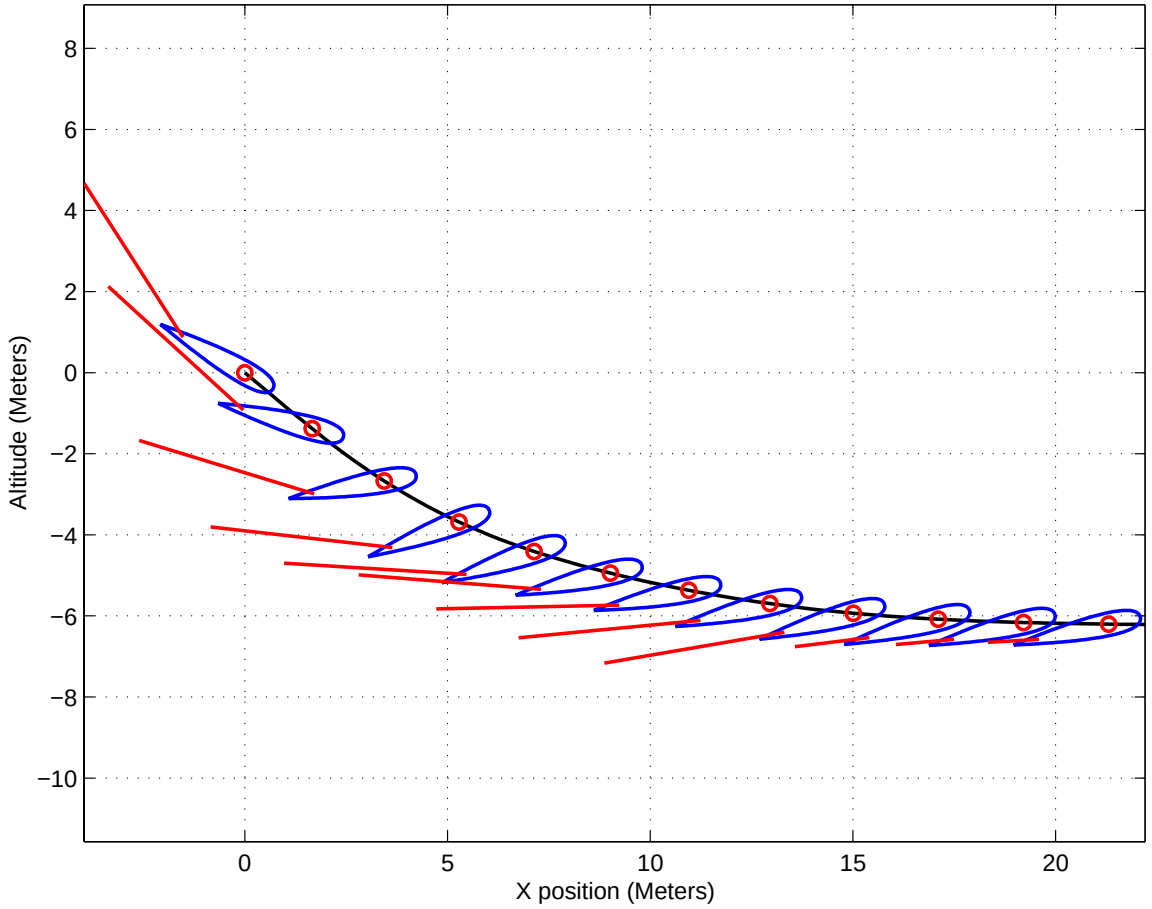

Figure 5.20: The $x - z$ trajectories for the ducted fan in forward flight with a one shot controller with $T = 3$ seconds. The red lines indicate the thrust, scaled in length according to the thrust value.

Figures 5.21 and 5.22 depict the simulation results when the desired flight path angle γ is a square wave, switching between -15° and 15°. Figure 5.21 depicts the time trajectories of the states and controls, whereas Figure 5.22 shows the trajectory in the $x - z$ space. As can be seen from both figures, the controls are saturated most of the time. Again, it should be noted that the simulation model extends the local range and includes the stall information. Also, the objective function in the receding horizon optimizations was the same as in the quasi-LPV scheme.

The resulting receding horizon optimizations were performed using RIOTS [Sch96] in the Matlab environment. Each receding horizon iteration (for a 0.5 second horizon length) took about 0.6 seconds on a Sun Ultra 30 machine. The next step is to verify the simulation results on the actual experiment using a much more efficient optimal trajectory solver recently developed at Caltech [MMM00].

5.4 Summary

The purpose of this chapter was to demonstrate the results of Chapters 2 and 4 on two distinct models of the Caltech ducted fan experiment. Equations of motion were obtained for the hover mode as well as the forward flight mode. Using quasi-LPV methods discussed in Chapter (4), two CLFS were obtained and used as a terminal cost in the receding horizon scheme. Several simulations were performed, and the results suggested that the proposed receding horizons scheme is suitable for regulation of the Caltech ducted fan and guaranteeing closed-loop stability.

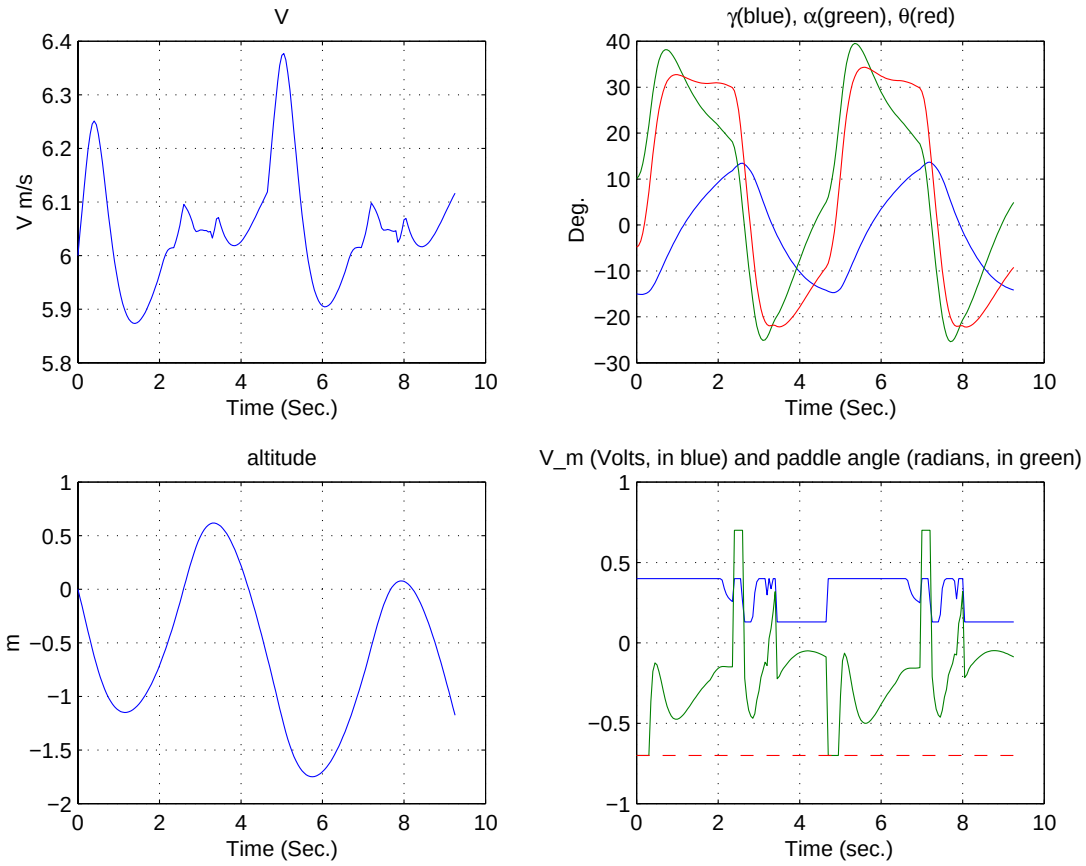

Figure 5.21: Simulation results for the ducted fan in forward flight with a receding horizon controller $\mathcal{RH}(0.5, 0.05)$. Shown here are the time trajectories for V, γ, α, and θ as well as controls V_m and δ_p. The commanded γ is switched between -15 and 15 degrees.

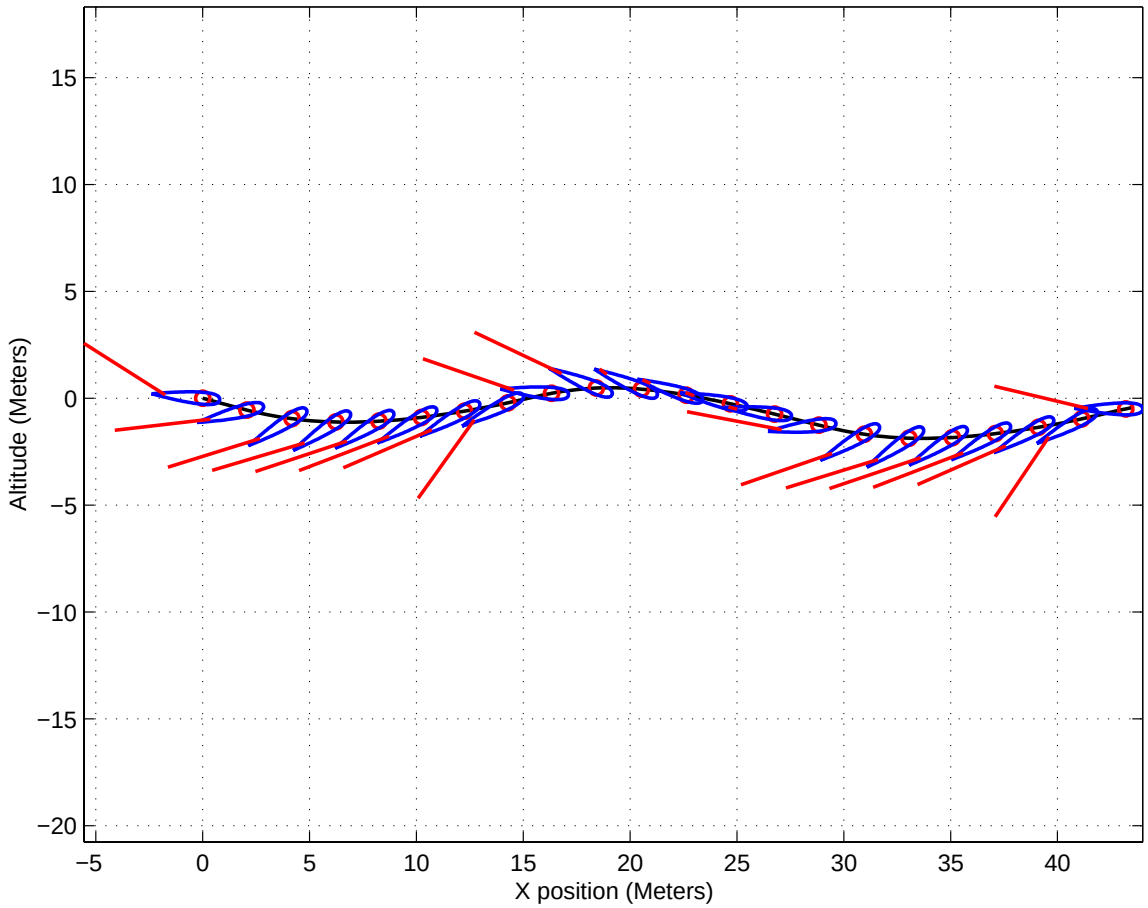

Figure 5.22: Simulation results for the ducted fan in forward flight with a receding horizon controller $\mathcal{RH}(0.5, 0.05)$. Shown here is the phase trajectory in the $x - z$ plane. The flying wing represents the ducted fan and the red line depicts the thrust, scaled in length according to the actual thrust value. The commanded γ is switched between -15 and 15 degrees.

Chapter 6 Receding horizon trajectory generation for the Caltech ducted fan

6.1 Trajectory generation and morphing

So far, we have outlined a receding horizon strategy which utilizes an *a priori* obtained terminal cost to guarantee closed-loop stability. However, it was assumed that the goal of the control strategy is to regulate the states to an equilibrium point, such as the origin.

Although this is an important task on its own, and there exist a plethora of results for addressing this problem, there are much fewer results that successfully address the issue of *trajectory generation* for a general nonlinear system. In this case, the main question being asked is not about finding a stabilizing controller, but about finding a pair of solutions (x, u) in the trajectory space of the system $\dot{x} = f(x, u)$, such that the pair (x, u) is "close enough" to a desired maneuver. In other words, instead of asking the system to "just get to the desired point," one would like to first provide a trajectory, and then regulate the system along that.

This has led to the *two degree of freedom* paradigm (see [NM96, MMM00] and the references therein). In this approach, the problem can be divided to two phases. The first phase deals with *generating* a state and control trajectory that respects the dynamics and other possible input and state constraints, and the second phase is to design a controller that would regulate the system around the trajectory.

Luckily, there exist a class of systems for which this problem can be easily tackled. It is known that if a system is *differentially flat* [FLMR95], the problem of trajectory generation reduces to solving a set of algebraic equations. Loosely speaking, a system is differentially flat, if there exist an output known as the *flat output*, such that states and the input can be written as a function of the flat output and a finite number of

its derivatives. It can be shown that all linear systems are differentially flat, as well as nonlinear systems that are *feedback equivalent* to a linear system. Unfortunately, it is extremely difficult, in general, to determine whether a given nonlinear system is differentially flat. Moreover, there are a lot of nonlinear systems that are *not* differentially flat.

The ducted fan in forward flight is an example of a nonlinear system which is not differentially flat. How should we go about generating trajectories for such a system? This is the question we intend to address in this chapter.

Our approach for trajectory generation and tracking is based on combining the phase one and phase two of the above mentioned two degree of freedom paradigm.

Using the idea of trajectory morphing [HM98b], we parameterize the trajectories of a complicated nonlinear system by those of a simpler one. With a homotopy connecting the simple system to the complicated one, one may *morph* simple trajectories to those of the complicated system. One thus seeks a simple system for which trajectory exploration and specification is tractable and that is sufficiently rich to capture the essential dynamic coupling of the target system. In an extreme case of simplification, one could use the system obtained by Jacobian linearization of the nonlinear dynamics as the simple system.

To illustrate this point, consider the following optimization problem:

$$\begin{aligned} \text{minimize} \quad & \frac{1}{2}\int_{t_0}^{t_f} \|(x(\tau)-x_d(\tau), u(\tau)-u_d(\tau))\|^2 d\tau \\ \text{subject to} \quad & \dot{x}(t) = f(x(t),u(t)), \qquad t \in [t_0, t_f] \\ & x(t_0) = x_d(t_0) \end{aligned}$$

where f describes the dynamics of the complicated system and the integrand may be a weighted square. Also, we might include a terminal cost to condition the problem.

The above nonlinear least square problem would result in a trajectory (x, u) that is closest in an $\mathcal{L}_2$ sense, to that of the simplified system.

Morphing makes use of the fact that we know that $(x_d, u_d)(\cdot)$ satisfies the equations of motion for the simplified system. Thus, after a suitable augmentation of the

simplified state and controls, we may define a homotopy connecting the simple system f_0 and the complicated system f, e.g., $f_\lambda = (1-\lambda)f_0 + \lambda f$. We know the solution of the above optimization problem when $f = f_\lambda$, it is simply (x_d, u_d). If f and f_0 are well chosen (e.g., by a clever engineer), then λ can be continued from 0 to 1 resulting in a trajectory of the complicated system that resembles the prototype trajectory of the simple system.

Since the simple system is often differentially flat, the desired trajectory can be easily generated. If the simplified system captures the essential dynamics and features of the more complicated model, one can also use a receding horizon scheme to morph the trajectories of the simplified model to those of the more complicated one, instead of using the above mentioned homotopy argument. Due to the fact that the receding horizon approach results in a sampled data feedback, (often times) a stabilizing controller around the trajectory is also generated, therefore removing the need for a tracking controller and combining both phases of the two degree of freedom paradigm.

Of course the issue of a suitable terminal cost becomes an important one, since the usual methods of obtaining a CLF are only valid for regulation around the trivial trajectory, i.e., the equilibrium at the origin. In order for our stability arguments to be valid, the terminal cost should be a CLF proving stability of the system around the trajectory of the full model. Such CLF is in general impossible to find, since we do not know the trajectory of the full model *a priori*.

However, we know from Theorem 2.3.1 that for a *long enough* horizon length, the terminal cost becomes unnecessary for guaranteeing stability, although it might be still useful to have one to condition the problem numerically.

Before moving to trajectory generation for the vectored thrust model (5.6), we consider the planar ducted fan equation (5.2).

6.2 Trajectory generation for the planar ducted fan

The planar ducted fan is a differentially flat system. However, to illustrate the receding horizon trajectory morphing scheme, it would be instructive to use it as an example of our trajectory morphing scheme. We would like to consider the following simplification of the planar equations in (5.2):

$$\begin{aligned} m\ddot{x} &= -f_2 \sin\theta \\ m\ddot{y} &= f_2 \cos\theta - mg \\ J\ddot{\theta} &= r f_1 \end{aligned} \tag{6.1}$$

As can be seen from the above equations, we have ignored the effect of f_1 on the x and y dynamics. With this simplification, we can choose x and y to be the flat outputs.

Our goal is to generate a trajectory for the simplified dynamics in (6.2), and morph that into a trajectory of (5.2).

Consider the following desired values for x and y:

$$x_d = 0.5\cos(0.6\pi t) \quad y_d = 0.$$

The above choice of x_d results in the following choice for θ_d:

$$\theta_d = -\arctan(\frac{\ddot{x}_d}{0.28}) = -\arctan(6.34\ \sin(0.6\pi t)).$$

Once θ_d is computed, one can compute f_{1d} and f_{2d} to be the following

$$f_{2d} = \frac{mg}{\cos\theta} \quad f_{1d} = \ddot{\theta}_d \frac{J}{r}.$$

Now that the desired trajectory is available, we can use it as a reference in the least square optimization discussed earlier. We choose a "long" horizon length of 3 seconds

to ensure stability of the receding horizon scheme.

Simulation results are depicted in Figure 6.1. Perfect tracking is obtained in x and θ, while there is a tracking error in y. This is perhaps due to the fact that the desired trajectory is not a trajectory of the full model.

In the next section, we discuss some flight maneuvers for the vectored thrust model of the ducted fan in forward flight.

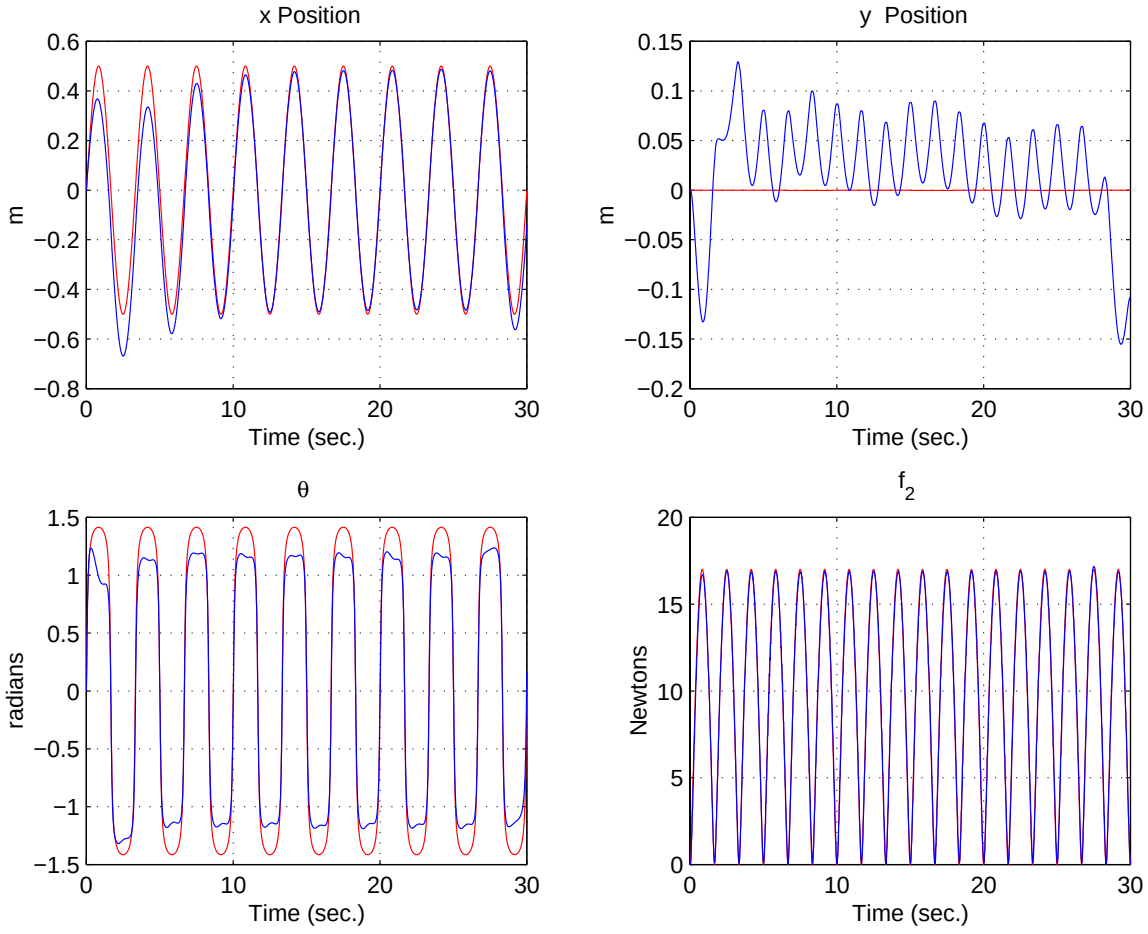

Figure 6.1: The desired trajectory (in red), and the actual trajectory (in black) for the planar ducted fan model. Shown here are the ducted fan positions x and y, as well as the pitch angle θ and the control action f_2.

6.3 Trajectory generation for the Caltech ducted fan in forward flight

As we have seen in Chapter 5, the Caltech Ducted Fan, modeled as a vectored thrust flying wing, is a highly nonlinear system with unique capabilities. From experience (and simple experiments), we know that operation near steady flight conditions can be easily achieved. Moreover, it is not too difficult to compute the associated equilibrium state and control values, as we did in Chapter 5. It is a much more difficult matter to specify feasible *non-equilibrium* trajectories, especially aggressive maneuvers that push the limits of performance.

The use of vectored thrust allows us to operate the ducted fan beyond the friendly *linear* aerodynamic regime, providing unique opportunities. For example, during a dive recovery, one may sacrifice a certain amount of energy (while creating high drag) for the sake of improved maneuverability (higher lift plus thrust vectoring).

How may we approach the specification of high performance, or aggressive, maneuvers? Although models of the sort described above are *nearly* differentially flat, we suspect that the presence of aerodynamic forces and moments breaks the necessary symmetries. Hence, trajectory generation for such a problem is by no means simple. To make things worse, the trajectory generation scheme has to handle input constraints as well.

In the case of the vectored thrust flying wing, one may obtain a simplified model by removing from consideration the (internal) pitch dynamics and then using the angle of attack α as a pseudo-control together with the thrust T. This results in a system with two states and two controls:

$$\begin{aligned} m\dot{V} &= -D(V,\alpha) - W\sin\gamma + T\cos(\alpha) \\ mV\dot{\gamma} &= L(V,\alpha) - W\cos\gamma + T\sin(\alpha). \end{aligned} \tag{6.2}$$

Where V and γ are the states, and T and α are the controls. Given a trajectory $(V_d(\cdot), \gamma_d(\cdot), T_d(\cdot), \alpha_d(\cdot))$ of the system (6.2), we build up a desired trajectory for the

full system (5.6) by defining $q_d(\cdot)$ as the (approximate, if necessary) derivative of $(\gamma_d + \alpha_d)(\cdot)$ and choosing $\delta_{\mathrm{T},d}$ so that

$$J\dot{q}_d = \quad M(V_d, \alpha_d) - T_d l_\tau \, \delta_{\mathrm{T},d}$$

if $\delta_{\mathrm{T},d}(\cdot)$ is sufficiently small (keeping the sin otherwise). The desired trajectory is then defined to be $x_d(\cdot) = (V_d, \gamma_d, q_d, \alpha_d)(\cdot)$ and $u_d(\cdot) = (T_d, \delta_{\tau,d})(\cdot)$.

6.4 Aggressive maneuvers

In this section we present some of the aggressive maneuvers obtained by implementing the morphing technique in a receding horizon fashion on maneuvers developed using the simplified model.

Note that as we mentioned earlier, in addition to actually generating trajectories of the more complicated system, the receding horizon approach also provides a (sampled data) feedback.

To get an idea of the maneuverability of the ducted fan, we first push the flying wing through a number of periodic climb/dive maneuvers similar to what we have flown of the real ducted fan. The period of the maneuvers was chosen to be 5 seconds. Figures 6.2 and 6.4 show the nature of these maneuvers as time trajectories. We see that the simplified system actually does a decent job of specifying approximate trajectories. The actual trajectories of the ducted fan for the above two cases are plotted in Figures 6.3 and 6.5 respectively.

A standard aerobatic maneuver for reversing direction is made up of a half loop, better known as an Immelman. For the ducted fan (in up and away flight sans floor and ceiling), the idea is to make the flight path angle γ go from 0 degrees to 180 degrees. The piloted maneuver would be completed with a 180 degree roll but we have no roll axis with the ducted fan and, moreover, the ducted fan can easily fly *inverted* since the system is more or less symmetric. Figures 6.6 and 6.7 show such a maneuver. Once again, the simplified system works quite well. In this figure, we

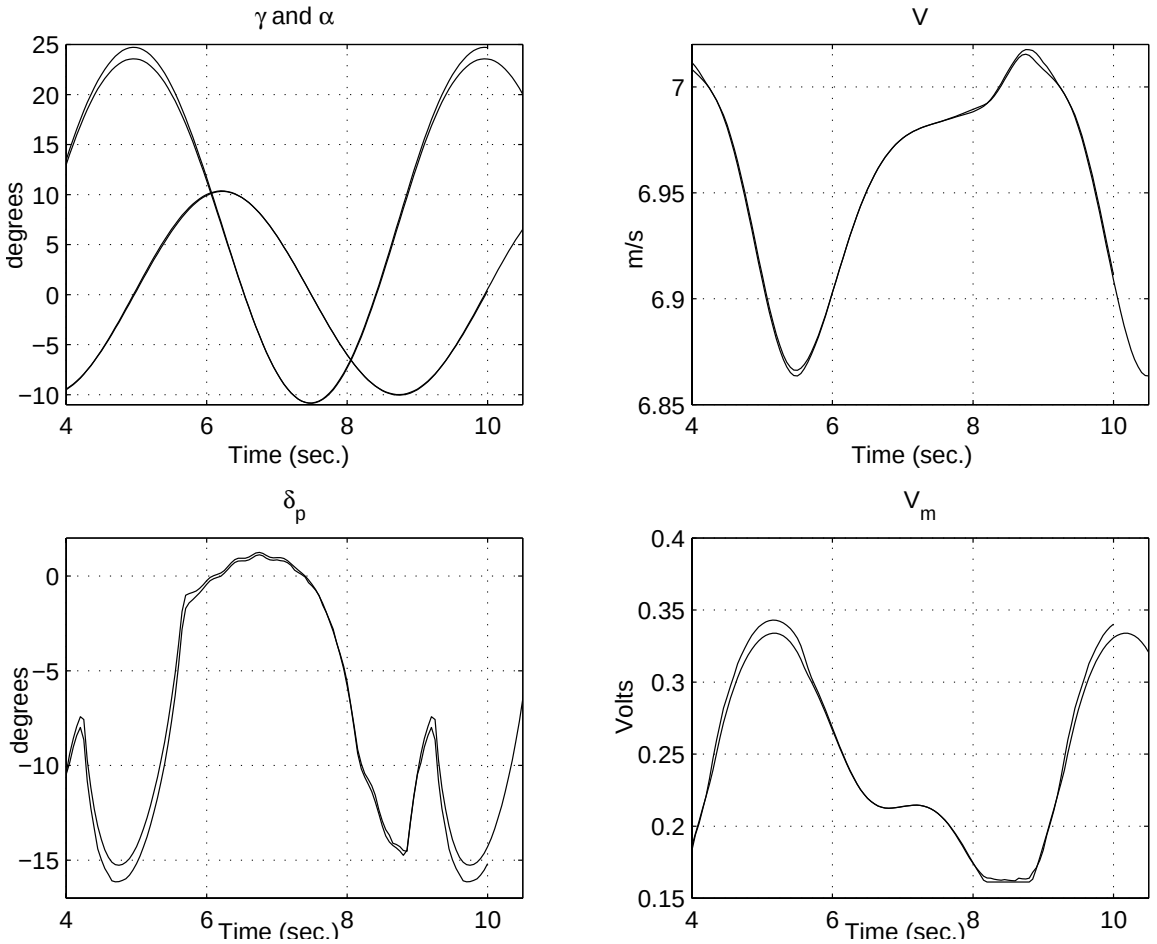

Figure 6.2: Periodic diving and climbing maneuver with approximately 20 degrees flight path angle γ change. Note that the angle of attack α approaches 25 degrees.

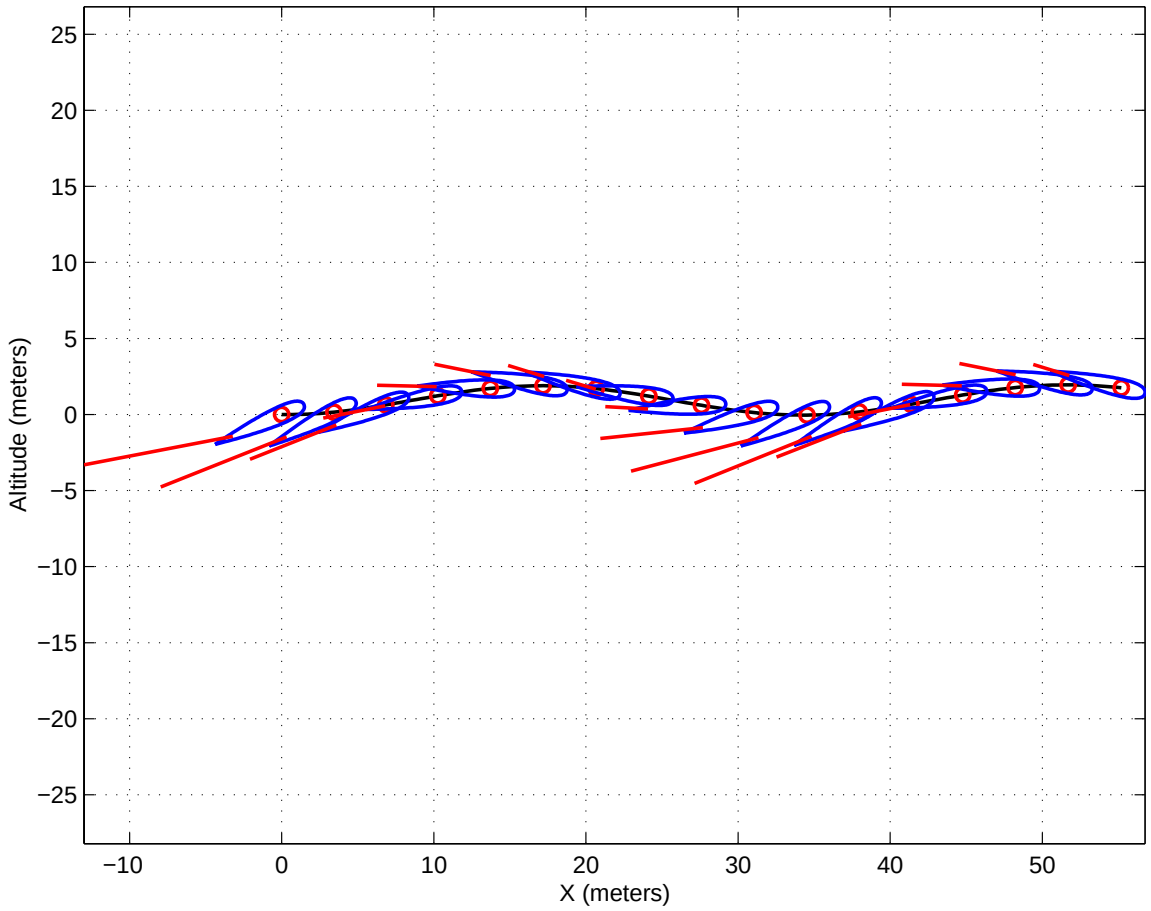

Figure 6.3: Periodic diving and climbing maneuver with approximately 20 degrees flight path angle γ change. Shown here is the actual trajectory of the ducted fan in the $x - z$ plane. The length of the red lines scales with the actual thrust value.

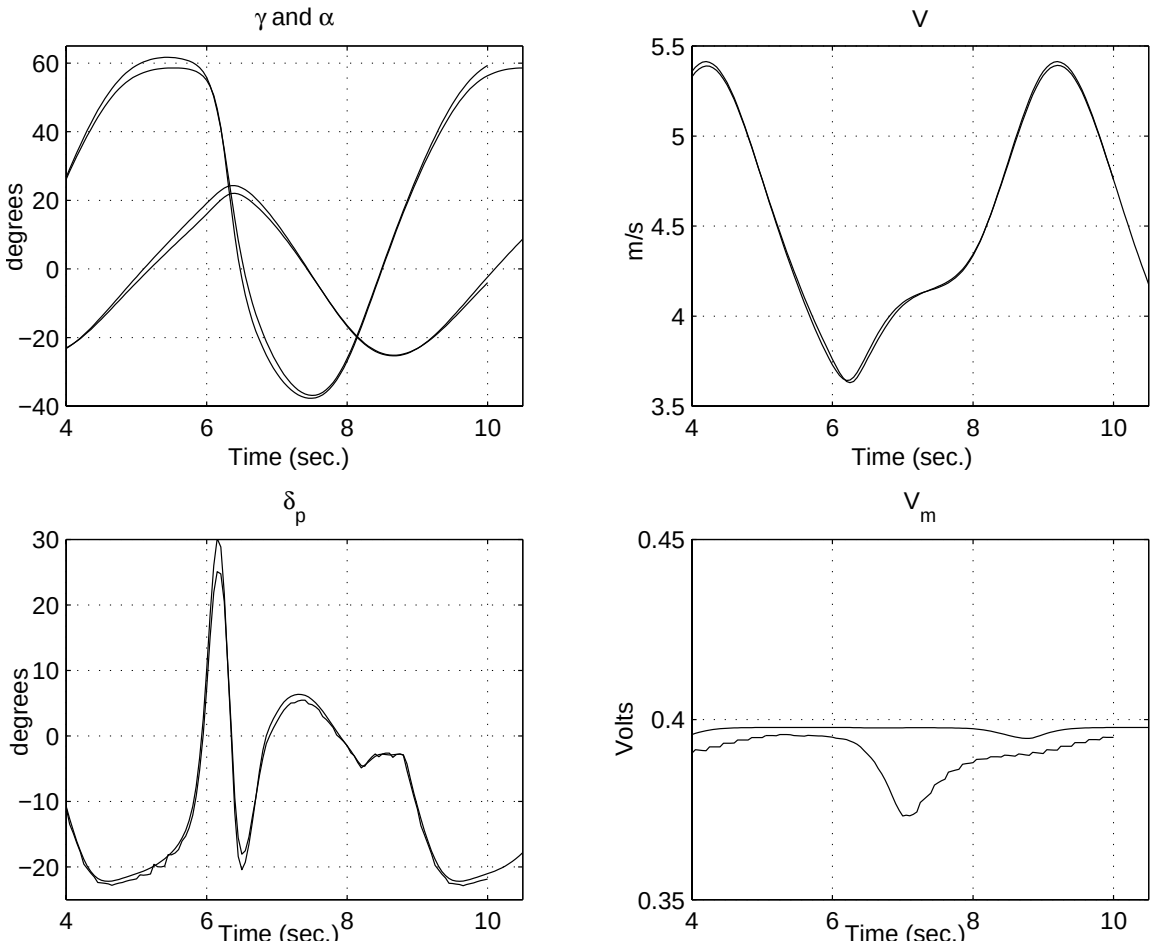

Figure 6.4: Periodic diving and climbing maneuver with approximately 50 degrees flight path angle γ change. Note that the angle of attack α goes to higher than 60 degrees. Also, we see that the thrust is almost continuously at the maximum.

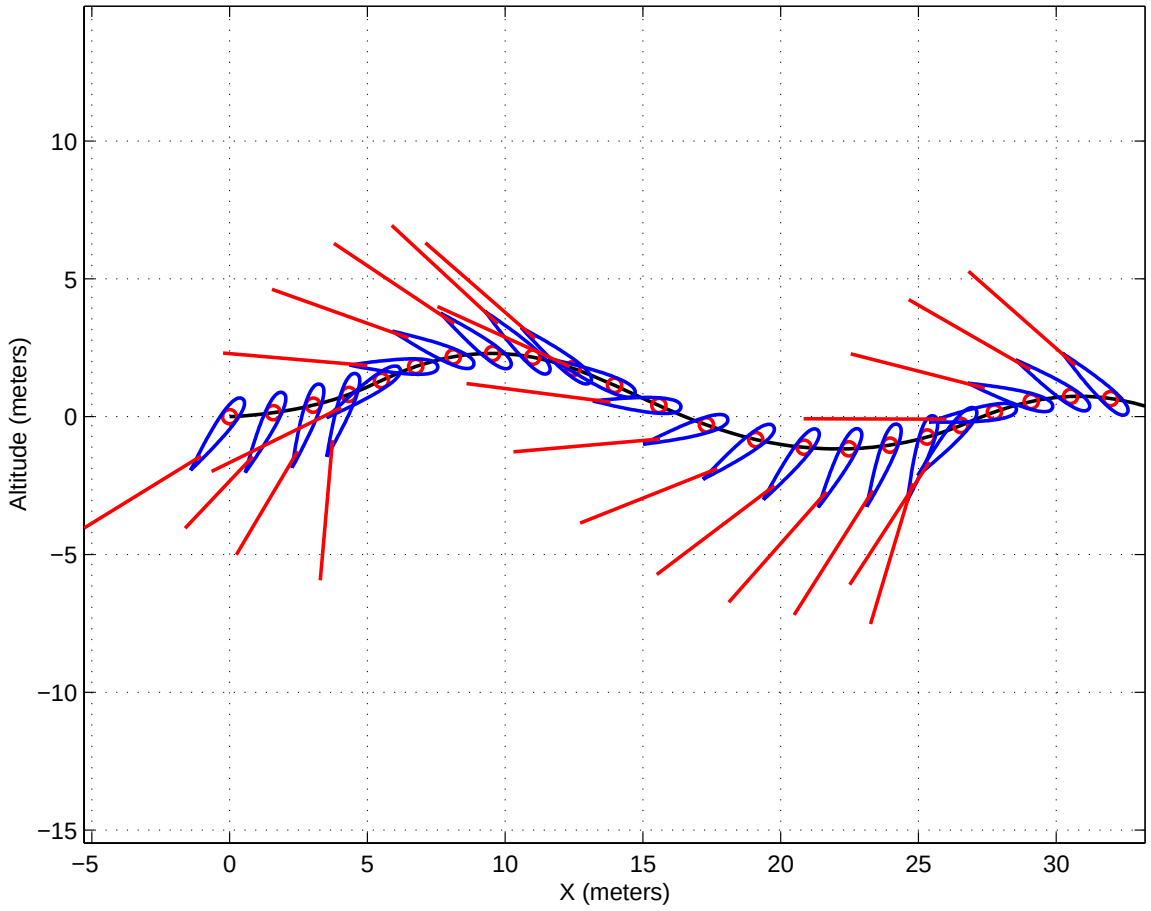

Figure 6.5: Periodic diving and climbing maneuver with approximately 50 degrees flight path angle γ change. Note that the angle of attack α goes to higher than 60 degrees. Shown here is the actual trajectory of the ducted fan in the $x - z$ plane. The length of the red lines scales with the actual thrust value.

have actually plotted the entire set of one second optimal trajectories (with artifacts between them).

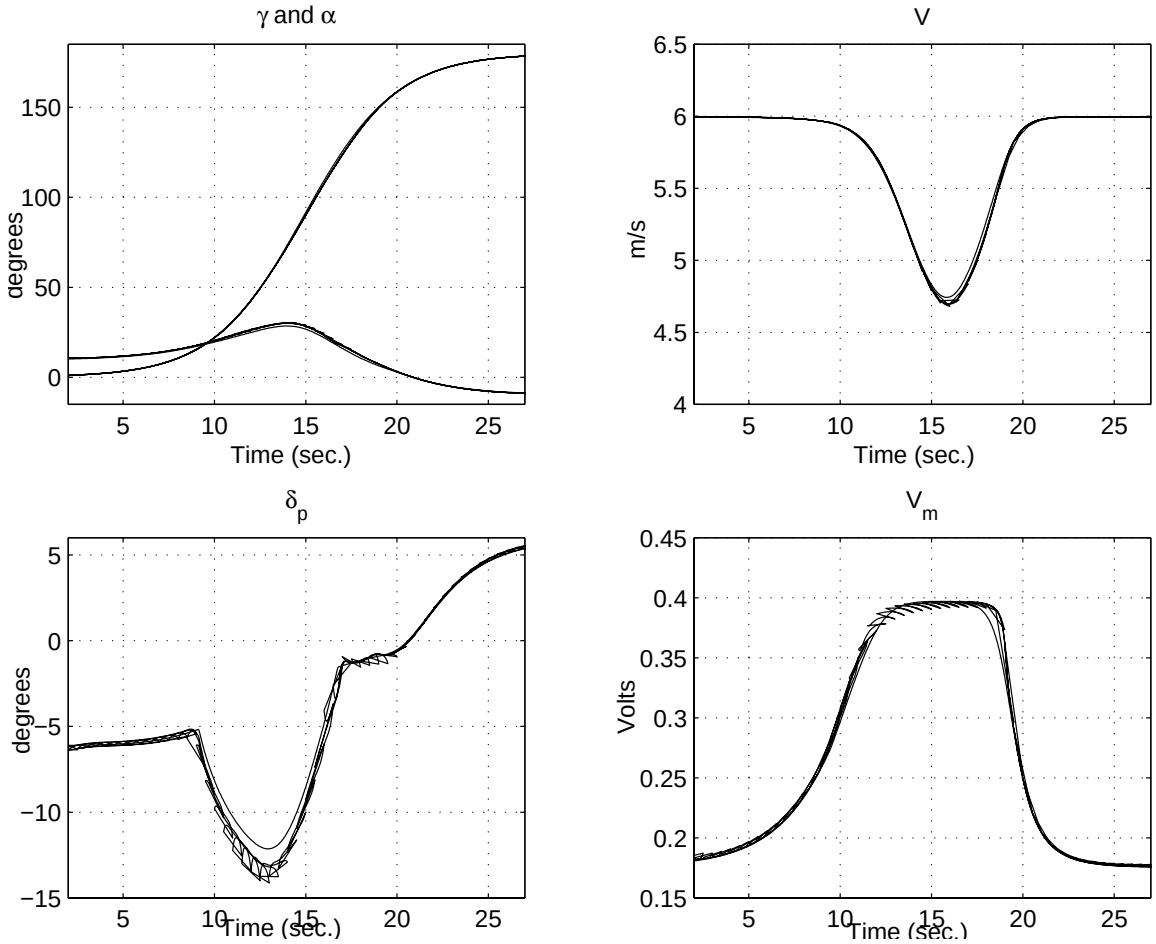

Figure 6.6: Immelman maneuver for reversing direction. Note the radical change in δ_p as stall is encountered.

6.5 Numerical considerations

The simulations presented in this chapter were performed using RIOTS [Sch96] in the Matlab environment. The software package RIOTS is a general purpose trajectory optimization solver. first, the problem is reduced to a nonlinear program and then the resulting nonlinear program is solved by NPSOL Version 5.0 [GMSW98] using a sequential quadratic programing (SQP) algorithm. RIOTS uses a method known

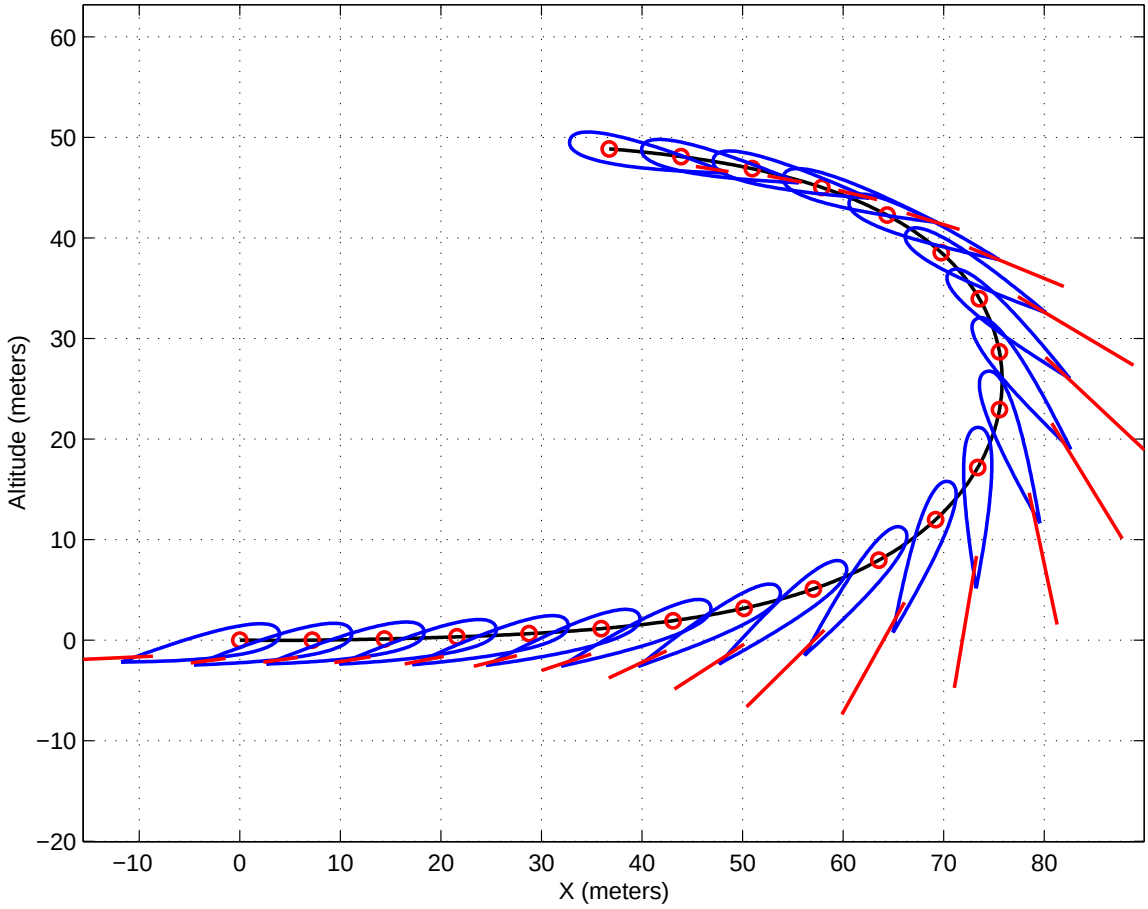

Figure 6.7: The actual trajectory of the ducted fan in the $x-z$ plane is shown during an Immelman maneuver. Thrust is represented by red lines, the length of which is scaled with the magnitude of thrust.

as consistent approximations [Pol93]. In this approach, a solution is obtained as an accumulation point of the solutions to a sequence of discrete-time optimal control problems that are consistent approximations to the original continuous-time optimal control problem. The discrete-time optimal control problems are constructed by discretizing the system dynamics with Runge-Kutta integration method and by parameterizing the control with finite-dimensional B-splines [Sch96]. Note that in this approach, only the control trajectory is parameterized. The control trajectory is fed into the system dynamics as a function of *time*, and the resulting system is solved using Runge-Kutta methods.

Since this is performed in an *open-loop* fashion, it suffers from *instability of shooting* [Hau00]. This makes it extremely difficult to solve optimization problems for horizon lengths of more than 2 seconds.

An alternative approach is to use collocation methods [MMM00], and parameterize both state and control trajectories. At certain collocation points, the dynamics are enforced, reducing the problem into a nonlinear program with equality constraints. Using ideas from differential flatness, this can be done in a very efficient manner, opening up the possibility of running the optimizations real-time [MMM00].

6.6 Summary

The purpose of this chapter was to demonstrate the use of receding horizon strategies for trajectory generation and tracking. Using an approach known as morphing, trajectories were obtained for a simplified model of the system under consideration. A least squares optimization problem is then formulated and solved in a receding horizon fashion, such that a sampled data receding horizon feedback controller is obtained at the same time as the trajectory of the full system.

These results were demonstrated in simulations on the planar model as well as the thrust vectored forward flight model of the Caltech ducted fan.

Chapter 7 Conclusions

7.1 Thesis summary

The purpose of this thesis was to develop optimization-based control schemes that are suitable for control of unmanned aerial vehicles. The developed scheme is based on a receding horizon control strategy, also known as model predictive control or moving horizon control.

The main idea behind receding horizon control is to solve a series of open-loop finite horizon optimal control problems in real time. The control trajectory resulting from this optimization is then applied to the system under consideration for a small fraction of the horizon length, after which a new state measurement is made and this process is repeated.

Although this idea and its variants have been used in industry and extensively studied, its use has only been limited to problems in process control industry, where the processes are usually open-loop stable and have large time constants.

With the advent of faster and cheaper computers, as well as state of the art numerical methods for solving nonlinear optimization problems faster than ever, it was suspected that this technology can be used beyond process control.

The main obstacle in applying the receding horizon technology in stability critical areas such as flight control has been the fact that existing methodologies were not tailored for such purposes. Moreover, even the few theoretical results that dealt with nonlinear systems had very high computational cost.

We developed a theoretical framework for stability analysis of receding horizon controllers, based on utilizing a control Lyapunov function as terminal cost. It was shown that if an appropriate control Lyapunov function (obtained off-line) is used as a terminal cost, the stability of the receding horizon scheme can be guaranteed *without* the need for imposing additional stability constraints. Specifically, it was shown that

the region for which the unconstrained receding horizon scheme is stabilizing can be computed and estimated. Furthermore, we proved that this region can be grown to contain any compact subset of the infinite horizon region of attraction.

We showed that the horizon length can be thought of as a homotopy parameter, which can have a CLF controller at one extreme (no optimization), and the optimal control (infinite horizon optimization) at the other. Also, in another important proposition, we showed that there always exists a finite horizon length for which the receding horizon scheme is stabilizing even without the presence of a terminal cost. Several propositions and theorems were stated and proven, illustrating the exponential convergence of the resulting receding horizon trajectories to the origin.

Also, it was shown that the requirement that truly optimal trajectories be computed can be relaxed by an improvement condition. This was illustrated by an instructive example, showing that sub-optimal controllers can sometimes be even more desirable. We also proposed a number of possible approaches that provide the CLF suitable for receding horizon purposes.

One of these approaches provided local CLFs based on the linearization of the nonlinear system at the origin, and using the linear quadratic regulator theory, while another, used the more sophisticated approach of quasi-Linear Parameter Varying (quasi-LPV) methods, generating state dependent Lyapunov functions resulting in larger regions of attractions. This was based on solving a convex optimization problem off-line.

These results were then applied to a model Unmanned Aerial Vehicle (UAV), known as the Caltech ducted fan. The Caltech ducted fan is a tethered flying wing developed as a testbed for flight control. Two modes of operation of the Caltech ducted fan, namely hover and forward flight modes, were considered. Several simulation studies and numerical experiments were performed to illustrate the developed theory.

A simplified model ignoring the aerodynamic forces was developed for the hover mode. Two different CLFs were chosen as terminal costs and comparisons were made to explain the theoretical contributions in more detail.

For the forward flight mode, the ducted fan was modeled as a thrust vectored flying wing, and aerodynamic forces and moments were taken into account. The model parameters were verified by performing a series of experiments to develop the experimental equilibrium manifold. It was shown that the model and experimental equilibrium manifolds follow each other closely.

In the next step, we chose an operating condition and developed a state dependent representation of the forward flight model, suitable for use in the quasi-LPV scheme. Simulation results indicated that the obtained results are indeed close to the optimal solution.

In Chapter 6, we touched on the issue of trajectory generation, and used the idea of trajectory morphing to generate trajectories for the planar ducted fan model as well as aggressive maneuvers for forward flight. The main idea behind the morphing concept was to use a simplified version of the model which captured the essential dynamics of the system, to generate aggressive maneuvers. Once these maneuvers were generated, they can be morphed to trajectories for the original complicated model by solving a least square optimization in a receding horizon fashion. The obtained results were quite promising.

In the next section we talk about possible future directions.

7.2 Future research directions

7.2.1 Running the experiment

The next obvious step in this research will be using the developed receding horizon methodology on the actual experiment. In fact, with availability of trajectory generation software developed at Caltech [MMM00], it would be possible to try the stabilization schemes on the experimental setup. Currently, this work is underway at Caltech to use the trajectory generation software in conjunction with the experiment. The removal of stability constraints has sped up the optimizations drastically. Therefore, we expect that using efficient software in conjunction with our results would

make it possible to run the receding horizon scheme on the experiment in real time.

7.2.2 Inclusion of state constraints

As was mentioned in the Introduction, one of the main advantages of optimization-based control schemes is their ability to handle constraints.

As far as the theoretical results are concerned, the method developed in this thesis can provably handle input and state constraints as long as the CLF can handle them as well. We showed how to use the same unconstrained CLF in the constrained case. However, as we saw in the thesis, the region of attraction gets smaller when the constraints are present. There are quite a few papers in the literature that propose methods for generating CLFs in the presence of input constraints. State constraints are, on the other hand, more difficult to handle. Perhaps the same methodology of using the unconstrained CLF, and getting a smaller region of attraction, can be a first step.

7.2.3 Multi vehicle extensions and decentralized control

The ultimate goal of this research project is to develop a decentralized, optimization based scheme for control of flocks of UAVs rather than just one. The single vehicle platform was just a starting point to establish the fact that receding horizon methodology is suitable for control of UAVs. A next step would be to use two vehicles, perhaps the actual experimental setup in addition to a virtual vehicle which can be in a computer.

The idea would be to try to tackle issues that do not appear in the single vehicle case, such as decentralization, coordination, communication and conflict resolution.

In the multi vehicle case, a challenge would be the design of decentralized controllers. In this setting, each vehicle would have only partial information about the state of other neighboring vehicles. This problem is well documented and well-known to be difficult in control literature. It would be interesting to pursue this in a receding horizon setting. Stability arguments will be undoubtedly more challenging, and

search for suitable terminal costs seem to be important issues.

7.2.4 Fault tolerance and online adaptation

A very important aspect of any flight control scheme is its ability to handle faults and adapt to changes in the environment and/or system. One of the possible advantages of online control schemes is their ability to adapt to changes. Since the control is not pre-programed, and is computed in real time, sudden changes in the cost function, dynamics and/or constraints can be handled. Of course, the choice of a terminal cost becomes a crucial one, since that usually requires off-line calculations. One possible direction to pursue would be generating the CLF in parallel to the optimization. In particular, when the simple CLF from LQR is chosen as terminal cost, one can solve Ricatti equations, in real-time, so that the CLF can be changed if necessary.

7.2.5 Uncertainty management

Throughout this thesis, the focus has been on developing sub-optimal control schemes that can handle nonlinearity. An aspect which has not been addressed is the issue of uncertainty. Although there are various approaches available for addressing uncertainty in linear systems, there are not nearly as many for nonlinear systems. There have been various attempts to address robustness by solving min-max optimization problems in a receding horizon setting, but so far none these have been amenable to real time calculations.

There are many possible directions that one can pursue for uncertainty management. Specifically, it would be interesting to find out whether use of CLFs that are obtained having robustness in mind, will do any better in terms of disturbance rejection. Due to the difficulty of solving min-max problems in real time, it is not possible to address the robustness issue in a direct fashion. However, a starting point would be to use a CLF derived using robust methods, such as those from $\mathcal{H}_\infty$ approaches. It is well known from the full information $\mathcal{H}_\infty$ control theory that the worst case disturbance as well as the optimal controller are both of the state feedback form.

These results have also been extended to the Linear Parameter Varying (LPV) and quasi-LPV cases. By augmenting the objective function in the robust receding horizon scheme with the *a priori* computed worst case disturbance, the min-max problem can be turned into a minimization problem, which can then be solved in a receding horizon framework having the CLF from $\mathcal{H}\infty$ methods as terminal cost.

7.2.6 Software development

Another important issue is the development of proper optimal control software packages. The newly developed software package [MMM00] has addressed a significant number of issues that were not properly addressed in RIOTS [Sch96]. However, there are still some problems that need to be studied further. One important issue is the choice of collocation points. How many of these points does one need to pick to guarantee satisfactory performance? How should the gridding be performed? These are some of the issues that need to be addressed in a more rigorous fashion.

7.2.7 Trajectory generation and morphing

As we saw in Chapter 6, trajectory morphing ideas can be used in conjunction with receding horizon schemes to generate trajectories for nonlinear systems such as the vectored thrust model of the Caltech ducted fan in forward flight. We mentioned that for *long enough* horizon lengths, the terminal cost is not necessary for stability guarantees, and we used that to justify the receding horizon morphing scheme. More theoretical work is needed to come up with rigorous conditions under which the receding horizon morphing technique is successful. It should be noted that when the trajectory ends up in an equilibrium point, the same machinery that was developed in this thesis can be used to guarantee stability of the receding horizon/morphing scheme. In order to develop a complete theory, one should answer questions such as: When can the morphing be successful without the homotopy argument? Are there any better ways to transition from trajectories of simplified models to the more complicated ones, other than forming convex combinations? Under what conditions can

we perform the morphing, in a receding horizon setting, and what would be a suitable terminal cost for shorter horizon lengths? The answers to these questions would help develop a nice framework for aggressive maneuvering and trajectory generation for nonlinear systems.

Bibliography

[AF66] M. Athans and P. L. Falb. *Optimal Control*. McGraw-Hill, New York, 1966.

[Art83] Z. Artstein. Stabilization with relaxed controls. *Nonlinear Anal.*, 7(11):1163–1173, 1983.

[Ath71] M. Athans. The role and use of the stochastic linear quadratic Gaussian problem in control system design. *IEEE Transactions on Automatic Control*, 16(12):529–552, December 1971.

[BBK96] B. Bodenheimer, P. Bendotti, and M. Kantner. Linear parameter-varying control of a ducted fan engine. *International Journal of Robust and Nonlinear Control*, 6:1023–1044, Nov 1996.

[BCD97] M. Bardi and I. Capuzzo-Dolcetta. *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Birkhauser, Boston, 1997.

[Bel52] R. Bellman. The theory of dynamic programming. In *Proc. Nat. Acad. Sci. USA*, number 38, 1952.

[BFL+97] G. Balas, I. Fialho, L. Lee, V. Nalabantoglu, A. Packard, W. Tan, E. Wemhoff, G. Wolodkin, and F. Wu. Theory and application of linear parameter varying control techniques. In *Tutorial presented at the 1997 Automatic Control Conference*, Albuquerque, NM, 1997.

[BGFB94] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*, volume 15 of *Studies in Applied Mathematics*. SIAM, Philadelphia, PA, June 1994.

[BGW90] R. R. Bitmead, M. Gevers, and V. Wertz. *Adaptive Optimal Control*. Prentice Hall, Englewood Cliffs, N.J., 1990.

[BH75] A. E. Bryson and Y. C. Ho. *Applied Optimal Control*. Taylor & Francis, Bristol, PA, 1975.

[BMH98] G. Buttazzo, G. Mariano, and S. Hildebrandt. *Direct Methods in the Calculus of Variations*. Oxford University Press, New York, 1998.

[CA98] H. Chen and F. Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Automatica*, 34:1205–1217, 1998.

[CDM96] J.R. Cloutier, C.N. D'Souza, and C.P. Mracek. Nonlinear regulation and nonlinear H_∞ control via the state-dependent riccati equation technique. In *Proc. 1st Internat. Conf. on Nonl. Problems in Aviation and Aerospace*, Daytona Beach, FL, 1996.

[Ces83] L. Cesari. *Optimization - Theory and Applications: Problems with Ordinary Differential Equations*. Springer-Verlag, New York, 1983.

[CM96] D. Chmielewski and V. Manousiouthakis. On constrained infinite-time linear quadratic optimal control. *Systems and Control Letters*, 29:121–129, 1996.

[CMT87a] D. W. Clarke, C. Mohtadi, and P. S. Tuffs. Generalized predictive control–I. The basic algorithm. *Automatica*, 23:137–148, 1987.

[CMT87b] D. W. Clarke, C. Mohtadi, and P. S. Tuffs. Generalized predictive control–II. Extensions and interpretations. *Automatica*, 23:149–160, 1987.

[CR80] C. R. Cutler and B. L. Ramaker. Dynamic matrix control– A computer control algorithm. In *Joint Automatic Control Conf.*, San Francisco, California, 1980.

[CS82] C. Chen and L. Shaw. On receding horizon feedback control. *Automatica*, pages 349–352, 1982.

[DAC95] P. Dorato, C. Abdallah, and V. Cerone. *Linear-Quadratic Control, An Introduction*. Prentice Hall, Englewood Cliffs, New Jersey, 1995.

[Dan51] G. B. Dantzig. Maximization of a linear function of variables subject to linear inequalities. In T. C. Koopmans, editor, *Activity analysis of production and allocation*, volume 1660 of *Cowles comission monograph*. John Wiley, New York, 1951.

[DGKF89] J.C. Doyle, K. Glover, P. Khargonekar, and B. Francis. State-space solutions to standard $\mathcal{H}_2$ and $\mathcal{H}_\infty$ control problems. *IEEE Transactions on Automatic Control*, 34(8):831–847, Aug 1989.

[Doy78] J. C. Doyle. Guaranteed margins for LQG regulators. *IEEE Transactions on Automatic Control*, 23(4):756–757, Aug 1978.

[FLMR95] M. Fliess, J. Levine, Ph. Martin, and P. Rouchon. Flatness and defect of non-linear systems: introductory theory and examples. *International Journal of Control*, 61:1327–1361, 1995.

[GM86] C. E. Garcia and A. M. Morshedi. Quadratic programming solution of dynamic matrix control (QDMC). *Chemical Engineering Communications*, 46:73–87, 1986.

[GMSW98] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. User's guide for NPSOL (version 5.0): a fortran package for nonlinear programming. Stanford University SOL 86-1, Stanford University, Revised, July 1998.

[GMW81] P. E. Gill, W. Murray, and M. H. Wright. *Practical Optimization*. Academic Press, London and New York, 1981.

[GPM89] C. E. García, D. M. Prett, and M. Morari. Model predictive control: Theory and practice – A survey. *Automatica*, 25(3):335–348, May 1989.

[Hau00] J. Hauser. *Personal Communications*. 2000.

[HJ99] Y. Huang and A. Jadbabaie. Nonlinear H_∞ control: an enhanced quasi-LPV approach. In *Proceedings of the IFAC World Congress, Vol. F*, pages 85–90, 1999.

[HJ00] J. Hauser and A. Jadbabaie. Aggressive maneuvering of a thrust vectored flying wing: A receding hoizon approach. In *IEEE Conference on Decision and Control*, Sydney, Australia, 2000.

[HM98a] J. Hauser and D. G. Meyer. The trajectory manifold of a nonlinear control system. In *IEEE Conference on Decision and Control*, pages 1034–1039, 1998.

[HM98b] J. Hauser and D. G. Meyer. Trajectory morphing for nonlinear systems. In *American Control Conference*, pages 2065–2070, 1998.

[HSM92] J. Hauser, S. Sastry, and G. Meyer. Nonlinear control design for slightly nonminimum phase systems - application to v/stol aircraft. *Automatica*, 28:665–679, 1992.

[Hua98] Y. Huang. *Nonlinear Optimal Control: An enhanced quasi-LPV approach.* PhD thesis, California Institute of Technology, Pasadena, CA, 1998.

[JYH99a] A. Jadbabaie, J. Yu, and J. Hauser. Receding horizon control of the Caltech ducted fan: A control Lyapunov function approach. In *IEEE Conference on Control Applications*, pages 51–56, 1999.

[JYH99b] A. Jadbabaie, J. Yu, and J. Hauser. Stabilizing receding horizon control of nonlinear systems: A control Lyapunov function approach. In *American Control Conference*, pages 1535–1539, 1999.

[Kal60] R. E. Kalman. Contributions to the theory of optimal control. *Bol. Soc. Mat. Mexicana*, 5:102–119, 1960.

[Kal64] R. E. Kalman. When is a linear control system optimal? *J. Basic Engrg. Trans. ASME Ser. D*, 86:51–60, 1964.

[KB61] R. E. Kalman and R. Bucy. New results in linear filtering and prediction theory. *J. Basic Eng., Trans. ASME, Ser. D*, 83:95–108, 1961.

[KBK83] W. H. Kwon, A. N. Bruckstein, and T. Kailath. Stabilizing state feedback design via the moving horizon method. *International Journal of Control*, 37(3):631–643, 1983.

[KBPM95] M. Kantner, B. Bodenheimer, P.Bendotti, and R. M. Murray. An experimental comparison of controllers for a vectored thrust ducted fan engine. In *American Control Conference*, pages 1956–1961, 1995.

[KG88] S. Keerthi and E. Gilbert. Optimal infinite-horizon feedback laws for a general class of constrained discrete-time systems: Stability and moving-horizon approximations. *Journal of Optimization Theory and Applications*, pages 265–293, 1988.

[KKK95] M. Kristić, I. Kanellakopoulos, and P. Kokotović. *Nonlinear and Adaptive Control Design*. John Wiley & Sons, New York, 1995.

[Kle70] B.L. Kleinman. An easy way to stabilize a linear constant system. *IEEE Transactions on Automatic Control*, 15:693, 1970.

[KP77] W.H. Kwon and A.E. Pearson. A modified quadratic cost problem and feedback stabilization of a linear system. *IEEE Transactions on Automatic Control*, 22(5):838–842, 1977.

[LH93] M. C. Lai and J. Hauser. Computing maximal stability region using a given Lyapunov function. In *American Control Conference*, pages 1500–1502, 1993.

[LM67] E. B. Lee and L. Markus. *Foundations of Optimal Control Theory*. Wiley, New York, 1967.

[LP44] A. I. Lur'e and V. N. Postnikov. On the theory of stability of control systems. *Applied Mathematics and Mechanics*, 8(3), 1944.

[ML99] M. Morari and J. H. Lee. Model predictive control: past, present, and future. *Computers and Chemical Engineering*, 23:667–682, 1999.

[MM93] H. Michalska and D.Q. Mayne. Robust receding horizon control of constrained nonlinear systems. *IEEE Transactions on Automatic Control*, 38(11):1623–1633, November 1993.

[MM99] M. Milam and R. M. Murray. A testbed for nonlinear flight control techniques: The Caltech Ducted Fan. In *IEEE Conference on Control Applications*, pages 345–351, Big Island, Hawaii, 1999.

[MMM00] M. Milam, K. Mushambi, and R. M. Murray. A computational approach to real-time trajectory generation for constrained mechanical systems. In *IEEE Conference on Decision and Control*, 2000.

[MRRS00] D. Q. Mayne, J. B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[MS97] L. Magni and R. Sepulchre. Stability margins of nonlinear receding horizon control via inverse optimality. *Systems and Control Letters*, 32:241–245, 1997.

[Mur98] R.M. Murray. Modeling of the Caltech ducted fan, class notes for cds 111. Technical report, California Institute of Technology, Control and Dynamical Systems 107–81, Pasadena, CA 91125, 1998.

[NM96] M.van Nieuwstadt and R. M. Murray. Real time tarjectory generation for differentially flat systems. In *Proceedings of the 1996 IFAC World Congress*, 1996.

[NMS98] G. De Nicolao, L. Magni, and R. Scattolini. Stabilizing receding-horizon control of nonlinear time-varying systems. *IEEE Transactions on Automatic Control*, 43(7):1030–1036, 1998.

[NN94] Yu. Nesterov and A. Nemirovsky. *Interior-point polynomial methods in convex programming*, volume 13 of *Studies in Applied Mathematics*. SIAM, Philadelphia, PA, 1994.

[PG79] D. M. Prett and R. D. Gillette. Optimization and constrained multivariable control of a catalytic cracking unit. In *AIChE Meeting*, Houston, TX, 1979.

[PN00] J.A. Primbs and V. Nevistic. Feasibility and stability of constrained finite receding horizon control. *Automatica*, 36(7):965–971, 2000.

[PND00] J. A. Primbs, V. Nevistić, and J. C. Doyle. A receding horizon generalization of pointwise min-norm controllers. *IEEE Transactions on Automatic Control*, 45:898–909, June 2000.

[Pol93] E. Polak. On the use of consistent approximations in the solution of semi-infinite optimization and optimal control problems. *Mathematical Programing*, 62:385–415, 1993.

[Ponan] L. S. Pontryagin. Optimal control processes. *Uspehi Mat. Nauk*, 14:3–20, 1959 (in Russian).

[Pri98] J. A. Primbs. *Nonlinear Optimal control: A receding horizon approach*. PhD thesis, California Institute of Technology, Pasadena, CA, 1998.

[PZ95] T. Parisini and R. Zoppoli. A receding horizon regulator for nonlinear systems and a neural approximation. *Automatica*, 31:1443–1451, 1995.

[RM93] J. B. Rawlings and K. R. Muske. The stability of constrained receding horizon control. *IEEE Transactions on Automatic Control*, 38(10):1512–1516, 1993.

[Roy88] H. L. Royden. *Real Analysis*. Macmilan, New York, 1988.

[RRTP78] J. Richalet, A. Rault, J. L. Testud, and J. Papon. Model predictive heuristic control: Applications to industrial processes. *Automatica*, 14(5):413–428, 1978.

[SA90] J. Shamma and M. Athans. Analysis of gain scheduled control for nonlinear plants. *IEEE Transactions on Automatic Control*, pages 898–907, 1990.

[SA91] J. Shamma and M. Athans. Guaranteed properties of gain scheduled control for linear parameter-varying plants. *Automatica*, pages 559–564, 1991.

[SA92] J. Shamma and M. Athans. Gain-scheduling: Potential hazards and possible remedies. *IEEE Control Systems Magazine*, 12(3):101–107, June 1992.

[Sch96] A. Schwartz. *Theory and Implementation of Numerical Methods Based on Runge-Kutta Integration for Optimal Control Problems*. PhD Dissertation, University of California, Berkeley, 1996.

[SCH+00] M. Sznaier, J. Cloutier, R. Hull, D. Jacques, and C. Mracek. Receding horizon control lyapunov function approach to suboptimal regulation of nonlinear systems. *Journal of Guidance, Control, and Dynamics*, 23(3):399–405, 2000.

[SD90] M. Sznaier and M. J. Damborg. Heuristically enhanced feedback control of constrained discrete-time linear systems. *Automatica*, 26:521–532, 1990.

[SMR99] P. Scokaert, D. Mayne, and J. Rawlings. Suboptimal model predictive control (feasibility implies stability). *IEEE Transactions on Automatic Control*, 44:648–654, March 1999.

[Son89] E.D. Sontag. A 'universal' construction of Artstein's theorem on nonlinear stabilization. *Systems and Control Letters*, 13(2):117–123, 1989.

[SR98] P. Scokaert and J. Rawlings. Constrained linear quadratic regulation. *IEEE Transactions on Automatic Control*, 43:1163–1169, August 1998.

[Tho75] Y.A. Thomas. Linear quadratic optimal estimation and control with receding horizon control. *Electronic Letters*, 11:19–21, 1975.

[Tro00] M. Trotoux. *Personal Communications*. 2000.

[vdS91] A. J. van der Schaft. On a state space approach to nonlinear H_∞ control. *Systems and Control Letters*, 116:1–8, 1991.

[vdS94] A. J. van der Schaft. *L_2-Gain and Passivity Techniques in Nonlinear Control*, volume 218 of *Lecture Notes in Control and Information Sciences*. Springer-Verlag, London, 1994.

[Wri98] M. H. Wright. The interior-point revolution in constrained optimization. In R. DeLeone, A. Murli, P. M. Pardalos, and G. Toraldo, editors, *High-Performance Algorithms and Software in Nonlinear Optimization*, pages 359–381. Kluwer Academic Publishers, 1998.

[WYPB96] F. Wu, H. X. Yang, A. Packard, and G. Becker. Induced L_2-Norm control for LPV systems with bounded parameter variation. *International Journal of Robust and Nonlinear Control*, vol. 6:983–998, 1996.

[YJPH99] J. Yu, A. Jadbabaie, J. Primbs, and Y. Huang. Comparison of nonlinear control designs for a ducted fan model. In *Proceedings of the IFAC World Congress, Vol. E*, pages 53–58, 1999.