

Nonlinear Optimal Control: A Receding Horizon Approach

Thesis by

James A. Primbs

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy



California Institute of Technology

Pasadena, California

1999

(Submitted January, 1999)

© 1999

James A. Primbs

All Rights Reserved

Acknowledgements

I have been particularly fortunate during my graduate studies to have had the opportunity to collaborate with a broad range of faculty, students, and visitors. First and foremost, I am obliged to my advisor, John Doyle, who provides an atmosphere that truly allows graduate students to transcend their traditional role and assume a position as an equal contributor to our discipline.

Additionally, I benefited from the guidance of a number of professors. Art Krener of UC Davis served as my undergraduate thesis advisor and introduced me to the subject of control theory. The CDS professors, Richard Murray and Jerry Marsden, generously made themselves available for consultation through the years. I would also like to thank Manfred Morari of ETH for inviting me to spend the summer of 1998 at the Automatic Control Laboratory in Zurich, Switzerland. Finally, I would like to acknowledge Mani Chandy, who along with John, Richard, Jerry, and Art, served on my thesis committee.

I have had the pleasure of sharing my graduate years with numerous visiting and fellow students. In particular, I would like to thank Yun Huang, who both as an officemate and through the organization of two workshops was always willing to help me sort through my numerous (and often random) questions. Furthermore, Yun, along with Ali Jadbabaie and Jie Yu, contributed greatly to Chapter 6 of this thesis. I am also grateful to Monica Giannelli, a visiting graduate student from Italy, whose knowledge of model predictive control led to informative and enjoyable collaborations resulting in Chapter 7. Additionally, numerous conversations with Pablo Parrilo, Xiaoyun Zhu, Sven Khatiri, Sonja Glavaski, Michael Kantner, Mark Milam, Muruhan Rathinam, Ben Shapiro, Hendra Adiwidjaja, Michael Gibson, Uri Keich, Jesse Brown, and others, contributed greatly to my entire educational experience.

Beyond any doubt, the person I am most indebted to is Vesna Nevistić, who was a frequent visitor to Caltech during her graduate studies in Switzerland. When I was

a new graduate student, she introduced me to her research topic and its community, inviting me to make them mine as well. Vesna was at times my greatest supporter, and at others my harshest critic, but always with the understanding, guidance, and care of a big sister. Technically, her influence and perspective helped shape a large portion of this thesis, including the core concepts in Chapters 2 through 5, but even more influential was her generous sense of sharing and teamwork that made our collaborations so productive and memorable.

Finally, I would like to thank my parents for their support throughout my entire education, not only my graduate studies. Their love and encouragement was worth far more than any degree.

Abstract

As advances in computing power forge ahead at an unparalleled rate, an increasingly compelling question that spans nearly every discipline is how best to exploit these advances. At one extreme, a tempting approach is to throw as much computational power at a problem as possible. Unfortunately, this is rarely a justifiable approach unless one has some theoretical guarantee of the efficacy of the computations. At the other extreme, not taking advantage of available computing power is unnecessarily limiting. In general, it is only through a careful inspection of the strengths and weaknesses of all available approaches that an optimal balance between analysis and computation is achieved. This thesis addresses the delicate interaction between theory and computation in the context of optimal control.

An exact solution to the nonlinear optimal control problem is known to be prohibitively difficult, both analytically and computationally. Nevertheless, a number of alternative (suboptimal) approaches have been developed. Many of these techniques approach the problem from an off-line, analytical point of view, designing a controller based on a detailed analysis of the system dynamics. A concept particularly amenable to this point of view is that of a control Lyapunov function. These techniques extend the Lyapunov methodology to control systems. In contrast, so-called receding horizon techniques rely purely on on-line computation to determine a control law. While offering an alternative method of attacking the optimal control problem, receding horizon implementations often lack solid theoretical stability guarantees.

In this thesis, we uncover a synergistic relationship that holds between control Lyapunov function based schemes and on-line receding horizon style computation. These connections derive from the classical Hamilton-Jacobi-Bellman and Euler-Lagrange approaches to optimal control. By returning to these roots, a broad class of control Lyapunov schemes are shown to admit natural extensions to receding horizon schemes, benefiting from the performance advantages of on-line computation. From

the receding horizon point of view, the use of a control Lyapunov function is a convenient solution to not only the theoretical properties that receding horizon control typically lacks, but also unexpectedly eases many of the difficult implementation requirements associated with on-line computation. After developing these schemes for the unconstrained nonlinear optimal control problem, the entire design methodology is illustrated on a simple model of a longitudinal flight control system. They are then extended to time-varying and input constrained nonlinear systems, offering a promising new paradigm for nonlinear optimal control design.

Contents

Acknowledgements	iii
Abstract	v
1 Introduction	1
1.1 Background	2
1.2 Thesis outline	4
2 Nonlinear Optimal Control	7
2.1 Introduction	7
2.2 Dynamic Programming: Hamilton-Jacobi-Bellman equations	8
2.3 Calculus of variations: Euler-Lagrange equations	12
2.4 Summary	16
3 Control Lyapunov Function Techniques	18
3.1 Introduction	18
3.2 HJB equations and CLF techniques	19
3.2.1 Control Lyapunov Functions (CLFs)	20
3.2.2 The value function as a CLF	21
3.2.3 CLF a substitute for the value function: Sontag's formula	22
3.2.4 Pointwise min-norm controllers	24
3.3 Example	27
3.4 Summary	30

4	Receding Horizon Control	32
4.1	Introduction	32
4.2	Receding Horizon Control (RHC)	32
4.2.1	Computational issues	34
4.2.2	Stability	36
4.2.3	Approaches to guaranteed stability	39
4.3	Summary	42
5	A Receding Horizon Extension of Pointwise Min-Norm Controllers	44
5.1	Introduction	44
5.2	Limits of receding horizon control	45
5.3	A receding horizon generalization of pointwise min-norm controllers	47
5.4	Implementation issues	53
5.5	Example	54
5.6	Summary	56
6	Control of a Ducted Fan Model	72
6.1	Introduction	72
6.2	An optimal control design paradigm	72
6.3	Caltech ducted fan model	73
6.4	Generation of CLFs	75
6.4.1	Jacobian linearization	76
6.4.2	Global linearization	77
6.4.3	Frozen Riccati Equation (FRE) method	78
6.4.4	Linear Parameter Varying (LPV) methods	80
6.5	CLF based control schemes	81
6.5.1	Sontag's formula	81
6.5.2	RHC extensions of CLF formulas	82
6.6	Comparisons	83
6.7	Implementation of on-line schemes	85

6.7.1	Time considerations	86
6.8	Summary	89
7	Extensions	95
7.1	Introduction	95
7.2	Time-varying optimal control	95
7.2.1	Optimal control for time-varying systems	96
7.2.2	CLF formulas for time-varying systems	97
7.2.3	A time-varying Sontag's formula	97
7.2.4	Receding horizon extensions of CLF schemes	99
7.2.5	CLFs for feedback linearizable systems	100
7.2.6	Example	103
7.3	Input constrained systems	109
7.3.1	Constrained nonlinear optimal control	110
7.3.2	A stabilizing bounded feedback control law	111
7.3.3	Receding horizon extensions	114
7.3.4	Example	115
7.4	Summary	116
8	Conclusions	119
8.1	Summary of main results	119
8.2	Future research	121
	Bibliography	123

List of Figures

2.1	Hamilton-Jacobi-Bellman vs. Euler-Lagrange Approach.	17
3.1	Contours of the value function (solid) and CLF (dashed).	28
3.2	Phase Portrait: Optimal (solid), Sontag's (dashed).	29
3.3	CLFs within the optimal control picture.	31
4.1	RHC for various horizon lengths.	38
4.2	Maximum eigenvalue versus horizon length for discretized linear system.	40
4.3	RHC within the optimal control picture.	43
5.1	Unified framework.	47
5.2	Performance constraint (5.7).	49
5.3	Phase portrait of receding horizon controllers.	56
5.4	The RHC+CLF scheme within the optimal control picture.	57
6.1	Planar ducted fan model.	74
6.2	Comparison of time versus horizon length.	88
6.3	Comparison of computation time for various constraints.	90
6.4	The frozen Riccati controller, initial condition $[5; 5; -0.9(\pi/2); 5; 0; 0]$	92
6.5	The LPV controller, initial condition $[5; 5; -0.9(\pi/2); 5; 0; 0]$	93
6.6	Optimal, initial condition $[5; 5; -0.9(\pi/2); 5; 0; 0]$	94
7.1	State and control trajectories from initial condition $[3, -2]$: Reference (dotted), CLF (dashed) and feedback linearized (dash-dot).	105
7.2	State and control trajectories from initial condition $[1, 6]$: Reference (dotted), CLF (dashed) and feedback linearized (dash-dot).	106
7.3	State and control trajectories from initial condition $[3, -2]$: Reference (dotted), RHC+CLF $T = 0.10$ (dashed) and $T = 0.25$ (dash-dot).	107

7.4	State and control trajectories from initial condition [1,6]: Reference (dotted), RHC+CLF $T = 0.10$ (dashed) and $T = 0.25$ (dash-dot). . .	108
7.5	Comparison of control trajectories from the pointwise min-norm and RHC+CLF controller from the initial condition [1,1].	117
7.6	Comparison of state trajectories from the pointwise min-norm and RHC+CLF controller from the initial condition [1,1].	118

List of Tables

3.1	Cost of Sontag's formula vs. the optimal from the initial condition [3, -2].	29
4.1	Comparison of controller performance from initial condition [3, -2].	39
5.1	Summary of controller costs from initial condition [3, -2].	57
6.1	Values of the cost function \mathcal{J} using different methods for the ducted fan example.	84
6.2	Horizon versus the number of time steps used in the integration scheme RK45.	87
7.1	Comparison of time-varying controller costs.	109
7.2	Cost of pointwise min-norm vs. RHC+CLF controller from initial condition [1, 1].	116

Notation

- \mathbf{C} The set of continuous functions.
- \mathbf{C}^n The set of functions n -times continuously differentiable.
- $\varphi(\cdot)$ Terminal weight function (used in optimal control objective function).
- $q(\cdot)$ Penalty on the state (used in optimal control problem).
- \mathbb{R} Real numbers.
- \mathbb{R}_+ Non-negative real numbers.
- σ Positive definite parameter in the pointwise min-norm problem.
- σ_s Pointwise min-norm parameter corresponding to Sontag's formula.
- T Horizon length in receding horizon schemes.
- T_s Sampling time in receding horizon schemes.
- u^* Optimal infinite horizon control trajectory corresponding to V^* .
- \hat{u}_T Optimal control trajectory from the RHC+CLF problem with horizon T .
- $\frac{\partial V}{\partial x}$ The row vector of partial derivatives of V , $[\frac{\partial V}{\partial x_1}, \frac{\partial V}{\partial x_2}, \dots, \frac{\partial V}{\partial x_n}]$.
- V^* The value function (minimum cost-to-go).
- V A control Lyapunov function.
- x^* Optimal infinite horizon state trajectory corresponding to V^* .
- \hat{x}_T Optimal state trajectory from the RHC+CLF problem with horizon T .

Chapter 1 Introduction

It is natural that when faced with a decision one would like, in some sense, to choose the “best” among the existing available alternatives. The process of finding the “best” has been formalized mathematically in the field of optimization. The standard approach is to rank the relative worth of each alternative strategy by a single real valued number known as the “performance index” or “objective function”. Optimization proceeds by selecting among the available strategies or decisions that which produces the greatest relative worth as characterized by the “performance index”.

While this approach suffers from obvious drawbacks, its mathematical simplicity and broad applicability have made it the dominant paradigm. Examples of the usefulness of this approach to decision making are abundant and span many fields of interest. To name a few, problems of allocation, planning, approximation, games, estimation, and control, spanning the interests of fields as diverse as economics, systems engineering, operations research, statistics, business, finance and others, are commonly posed in this framework.

While many problems can be properly formulated in the framework of modern optimization theory, this does not imply that the solution of these problems follows directly. In fact, it is the search for improved tools for the solution of these problems that comprises the current field of optimization. In recent decades, computers have played an increasing role in advancing the ability to solve optimization problems. This has also resulted in changing the face of optimization theory. Many old techniques have been rendered obsolete, being replaced by previously impractical methods which can now be efficiently performed with modern computing capabilities. Currently, the field has evolved into a blend of old and new, using numerical procedures, firmly rooted and justified in modern complexity theory but resting upon the fundamentals established by such greats as Gauss, Lagrange, Euler, the Bernoullis, Von Neumann and others. It is this synergy between the theoretical and numerical that holds great

promise for the future of optimization.

This thesis concentrates specifically on the nonlinear optimal control problem. This problem involves the optimization of a performance index associated with a system developing dynamically through time. Examples of such optimal control problems include: maximizing the range of a rocket, maximizing the profit produced by an economic enterprise, minimizing the error in the estimation of an object's position, minimizing the energy or cost required to achieve some specified terminal state, or any of a wide variety of similar tasks. The solution techniques for such problems, including the face of nonlinear optimal control in general, is changing in a manner similar to that of optimization, pushing its boundaries by reinventing techniques of the past and coupling them with the powerful computational tools of the present.

1.1 Background

For systems whose dynamics are linear and time-invariant, optimal control theory now has well developed tools for optimizing a number of performance indexes that embody desirable objectives. For instance, in addition to the classical \mathcal{H}_2 theory [AM89], there now exist not only theoretically elegant but computationally tractable solutions to the \mathcal{H}_∞ [DGKF89] and l_1 [DP87] problems.

In contrast, nonlinear optimal control (optimization constrained by a nonlinear dynamical system) is still a developing field. While its roots were laid down in the 1950s with the introduction of dynamic programming (leading to Hamilton-Jacobi-Bellman partial differential equations [Bel52]) and the Pontryagin maximum principle (a generalization of the Euler-Lagrange equations deriving from the calculus of variations [Pon59]), these were more theoretical contributions than practical design techniques. From these beginnings, numerous design methodologies for nonlinear optimal control have developed, often following different paths and techniques. Today, it appears as a fragmented field. The nonlinear optimal control problem is now attacked on many different fronts: by extending the linear theory, utilizing generalizations of the Lyapunov methodology, and brute force computation to name a few.

With less ambitious goals in mind than an exact solution to the nonlinear optimal control problem, classical Lyapunov theory (see, for example, [Kha92]) was extended to aid in the design of control laws. This led to the concept of a *control Lyapunov function* (CLF). Fueled by results establishing the equivalence of a control Lyapunov function and a continuous stabilizing control law [Art83], interest in control Lyapunov functions for design became active. An important contribution to this theory was the *explicit* construction of a stabilizing feedback control law given by Sontag [Son89]. Furthermore, systematic procedures emerged for deriving control Lyapunov functions for systems possessing special structure (e.g., feedback linearizable, strict feedback and feedforward systems [KKK95]).

More recently, the optimality properties of CLF based control laws have been analyzed. A concept referred to as *inverse optimality* [Kal64, MA73] was used to begin to bring optimality back into the picture of CLF based design. It was shown that every CLF is the value function solving a Hamilton-Jacobi-Bellman equation corresponding to a meaningful cost [FK95]. Furthermore, with the development of so-called pointwise min-norm controllers, an entire class of inverse optimal CLF control laws were introduced [FK95, FK96a].

At the same time, the advent of the microprocessor and the subsequent computer revolution opened up an entirely new possibility for optimal control: solution directly through numerical computations. While the solution of the Hamilton-Jacobi-Bellman equation remained intractable in all but the simplest cases, Euler-Lagrange type trajectory optimizations, deriving from the classical calculus of variations, provided an alternate more computationally feasible approach. By solving trajectory optimizations (which produce open-loop control trajectories as a function of time as opposed to a state-feedback law), computers were able to provide relatively efficient solutions. Feedback could then be incorporated by the repeated *on-line* solution of these trajectory optimizations, an approach known as receding or moving horizon. This spawned the technique of model predictive control [GPM89], which heavily exploited the receding horizon methodology. These techniques were first applied to plants with

slow dynamics where on-line intersample computation was feasible. Additionally, it was a natural approach to constrained systems because constraints could be directly incorporated into the optimizations. These techniques found success especially in industrial process applications [CR79, GPM89, Ric93, RRTP78].

Today, model predictive control or receding horizon control is gaining popularity as computers become increasingly faster. While connections with classical Euler-Lagrange type trajectory optimizations are often not mentioned explicitly, they implicitly provide the foundation for receding horizon techniques. While results from practical applications have been promising, these techniques have struggled to establish theoretical stability properties. Considerable success has been made for linear systems (see [GPM89] and references therein). While some of the same results hold for nonlinear systems (e.g., end constraints [MM90], infinite horizon [MS97, NMS98]) it is the structure of linear systems that general makes them computationally feasible.

1.2 Thesis outline

This thesis is an attempt to provide a more unified framework in which to understand the contributions of various design procedures toward nonlinear optimal control, and to exploit previously unrecognized connections to develop improved design methodologies. Our framework derives from the two classical approaches to the problem of optimal control. While only providing theoretical guidance, a deep understanding of their fundamental properties allows us to provide a better characterization and classification of state of the art techniques.

In this thesis we focus on CLF based control laws and the receding horizon methodology, but beginning from the premise that they inherit properties from the classical Hamilton-Jacobi-Bellman and Euler-Lagrange solutions to the optimal control problem. Hence, we begin by reviewing these two classical approaches to the optimal control problem in Chapter 2, highlighting their important properties and differences.

Next, in Chapter 3 we introduce control Lyapunov functions and their associated control laws in the context of Hamilton-Jacobi-Bellman equations. By viewing

a CLF as an approximation to the solution of the Hamilton-Jacobi-Bellman equation (commonly referred to as the *value function*), we derive a slight variation of Sontag’s formula [Son89] which has strong connections to an associated optimal control problem. Furthermore, it is shown that Sontag’s formula is actually a special case of pointwise min-norm controllers [FK95], which are known to possess inverse optimality properties. Nevertheless, we show that these CLF based control laws are similar in that they rely heavily on the shape of the level curves of the control Lyapunov function, and this can lead to poor performance when that shape does not resemble those of the value function.

Chapter 4 reviews the receding horizon methodology [GPM89]. This time, we relate the properties of receding horizon control to its roots in Euler-Lagrange type trajectory optimizations. Even though the receding horizon methodology produces a state feedback control law, it still inherits fundamental properties of open-loop Euler-Lagrange trajectory optimizations. We use this as a framework to review the stability properties of receding horizon control, and show how various stabilizing formulations have been developed to address these difficulties.

With the preceding chapters providing the foundation, in Chapter 5 we develop previously unrecognized connections that exist between pointwise min-norm controllers and the receding horizon methodology. By viewing pointwise min-norm controllers as a limiting case of a receding horizon scheme, this suggests that extensions of pointwise min-norm controllers to a receding horizon scheme should be possible. We develop such a scheme, extending pointwise min-norm controllers to incorporate on-line receding horizon style computation. This is presented within a new framework for nonlinear optimal control, in which optimal control and CLF based pointwise min-norm controllers are extreme cases of the new CLF based receding horizon scheme. Furthermore, philosophically this approach has a satisfying interpretation as a blend of the classical Hamilton-Jacobi-Bellman and Euler-Lagrange approaches to optimal control. While a CLF represents a global approximation to the value function (the solution of the Hamilton-Jacobi-Bellman equation), on-line trajectory optimizations represent local approximations. These two points of view are combined into a single

methodology.

In Chapter 6 the new design methodology is tested on a simple model of a longitudinal flight control system. We place existing techniques in a two stage design paradigm suggested by the framework developed in the previous chapter. The first stage involves the derivation of a CLF. For this task we consider techniques including Jacobian linearization, global linearization [LP44, BGFB94], frozen Riccati equations [CDM96], and quasi linear-parameter-varying methods [WYPB96]. The second stage requires the selection of a CLF based control law. For this we consider not only the standard implementation associated with each technique used in the first stage, but also Sontag's formula and its receding horizon extension. Simulation results indicate that our new control schemes, which fully utilize the contributions of existing techniques, can significantly outperform individual laws.

Chapter 7 extends the methodology to time-varying systems, which arise in problems of trajectory tracking, and input constrained systems. Simple examples are used to illustrate the methodology in these cases. Finally, Chapter 8 presents conclusions and future areas of research suggested by this thesis.

Chapter 2 Nonlinear Optimal Control

2.1 Introduction

Historically, the background of optimal control theory shows a long stream of scientific thought concerned with wave propagation and variational principles in physics, beginning with Huygens, continuing with Bernoulli, and finally achieving its maturity with the work of the great masters of the nineteenth century: Hamilton, Jacobi, and Lie. As argued by Sussmann [Sus96], perhaps the true birth of optimal control theory was in 1696 in the Netherlands, when Johann Bernoulli challenged his contemporaries with the *brachystochrone problem*. Given two points A and B in a vertical plane, find the orbit AMB of the movable point M which, starting from A and under the influence of its own weight, arrives at B in the shortest possible time. Bernoulli's brachystochrone problem was a true minimum-time problem, and the first to deal with a dynamical behavior and explicitly ask for the optimal selection of the path [Sus96].

Optimal control theory, in its modern sense, began in the 1950s with the formulation of two design optimization techniques: Dynamic Programming and the Pontryagin Maximum Principle. While the maximum principle, which represents a far-reaching generalization of the Euler-Lagrange equations from the classical calculus of variations, may be viewed as an outgrowth of the Hamiltonian approach to variational problems, the method of dynamic programming may be viewed as an outgrowth of the Hamilton-Jacobi approach to variational problems. In this chapter we explore the roots of these two modern approaches. This provides an important foundation for the following chapters, not in its technical detail, but rather in clarifying the fundamental differences between these two points of view. Later, we will see that many suboptimal approaches to nonlinear optimal control are aligned with one of these two approaches, leading to inherited advantages and disadvantages.

2.2 Dynamic Programming: Hamilton-Jacobi-Bellman equations

The nonlinear system under consideration will be of the form

$$\dot{x} = f(x) + g(x)u, \quad f(0) = 0 \quad (2.1)$$

with $x \in \mathbb{R}^n$ denoting the state, $u \in \mathbb{R}^m$ the control, and $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ continuously differentiable in all arguments.

Throughout this thesis, we will be concerned with the infinite horizon nonlinear optimal control problem stated below:

$$\begin{aligned} \min_{u(\cdot)} \quad & \int_0^\infty (q(x) + u^T u) dt \\ \text{s.t.} \quad & \dot{x} = f(x) + g(x)u \end{aligned} \quad (2.2)$$

for $q : \mathbb{R}^n \rightarrow \mathbb{R}$ positive semi-definite and \mathbf{C}^1 and the desired solution being a *state feedback* control law. We will also assume that the system $[f(x), q(x)]$ is *zero-state detectable*. (That is, for all $x \in \mathbb{R}^n$, $q(\phi(t, x)) = 0 \Rightarrow \phi(t, x) \rightarrow 0$ as $t \rightarrow \infty$ where $\phi(t, x)$ is the state transition function of the system $\dot{x} = f(x)$ from the initial condition $x(0) = x$.)

The dynamic programming solution

In this section we derive the Hamilton-Jacobi-Bellman partial differential equation solution to the nonlinear optimal control problem. The solution follows the technique known as *dynamic programming*, popularized by Bellman [Bel52]. We first explain the concept of dynamic programming, rooted in the so-called *principle of optimality*, then apply this concept to the optimal control problem in order to derive the Hamilton-Jacobi-Bellman partial differential equation.

The basis for the dynamic programming solution to the optimal control problem

is the so-called *principle of optimality*, formally stated as follows:

Definition 2.2.1 *Principle of Optimality:* *If $u^*(\tau)$ is optimal over the interval $[t, t_f]$, starting at state $x(t)$, then $u^*(\tau)$ is necessarily optimal over the subinterval $[t + \Delta t, t_f]$ for any Δt such that $t_f - t \geq \Delta t > 0$.*

The basic assumption underlying the principle of optimality is that the system can be characterized by its state $x(t)$ at time t , which completely summarizes the effect of all inputs $u(\tau)$ prior to time t . This allows for a local characterization of optimality as given in the principle of optimality. More details, as well as proof of the principle of optimality, can be found in many references [Sag68, AF66, AM89, DAC95].

Dynamic programming is the concept of using the principle of optimality to formulate an optimization problem as a recurrence relation, i.e., the remaining sub-problem has precisely the same structure as the previous sub-problem. In this way, a particular optimization problem is solved by studying a family of problems which contain the particular problem as a member.

For instance, in the optimal control problem, if one considers a function which associates to every point in state space the optimal cost starting from that point (such a function is often called a *value function*), then it is possible to write a recurrence relation in terms of the optimal value function which is valid for the entire state space. If this relation can be solved, the value function obtained is associated with an entire family of optimal control problems, each with a different initial point. While knowledge of the optimal value associated with a *single* initial condition provides no way of determining the minimizing trajectory itself, knowledge of the value function on the entire state space does allow one to determine the minimizing trajectory for any particular member of the family of problems. We demonstrate this idea more concretely by using it to solve the infinite horizon optimal control problem subject to time-invariant dynamics. More general problems (time-varying, finite horizon, etc.) are solved in a similar manner and can be found in the references [Nev97, Bel52, BH75, DAC95].

Define $V^*(x_0)$ to be the minimum of the performance index taken over all admissible trajectories $(x(t), u(t))$ where x starts at x_0 :

$$\begin{aligned} V^*(x_0) &= \min_{u(\cdot)} \int_0^\infty (q(x(t)) + u^T(t)u(t))dt \\ \text{s.t.} \quad \dot{x} &= f(x(t)) + g(x(t))u(t) \\ x(0) &= x_0. \end{aligned} \tag{2.3}$$

If no such trajectory exists, then $V^*(\cdot) = +\infty$. The function $V^* : \mathbb{R}^n \rightarrow \mathbb{R}_+ \cup \{\infty\}$, which determines the rule associating an optimal value with each initial point, is called the *value function* or *Bellman's function* of the optimal control problem. An *optimal pair* (often simply referred to as an “*optimal trajectory*”) is a pair $(x(t), u(t))$ that has a starting point x_0 and achieves the optimal cost $V^*(x_0)$.

Notice that $V^*(x_0)$ is independent of $u(\cdot)$, precisely because knowledge of the initial state abstractly determines the particular control by the requirement that the control minimizes the performance index. Rather than just searching for the control minimizing (2.1) and for the value of $V^*(x(t))$ for various x_0 , the problem is approached by considering the evaluation of $V^*(x(t))$ for all $x(t)$, and the associated optimal control.

Now let us apply the principle of optimality. Consider $V^*(x)$ given by (2.3), and let $u[t, \infty)$ be defined as the control signal over the interval $[t, \infty)$. Using the additive properties of integrals and the principle of optimality yields

$$V^*(x(t)) = \min_{u[t, t+\Delta t]} \left\{ \int_t^{t+\Delta t} [q(x(\tau)) + u^T(\tau)u(\tau)] d\tau + V^*(x(t + \Delta t)) \right\}. \tag{2.4}$$

That is, the optimal cost at state $x(t)$ is given by the minimum of the cost it takes to move to state $x(t + \Delta t)$ plus the optimal cost from $x(t + \Delta t)$. In essence, by using the principle of optimality the problem of finding an optimal control over the interval $[t, \infty)$ has been reduced to finding an optimal control over the reduced interval $[t, t + \Delta t]$.

Continuing further, when Δt is small, the integral in (2.4) can be approximated by

$[q(x(t)) + u^T(t)u(t)]\Delta t$. Applying a multivariable Taylor-series expansion of $V^*(x(t + \Delta t))$ about $x(t)$, with $x(t + \Delta t) - x(t)$ approximated by $[f(x(t)) + g(x(t))u(t)]\Delta t$, gives

$$V^*(x) = \min_u \left\{ [q(x) + u^T u]\Delta t + V^*(x) + \left(\frac{\partial V^*}{\partial x} \right) [f(x) + g(x)u]\Delta t + o(\Delta t) \right\}, \quad (2.5)$$

where $\frac{\partial V^*}{\partial x}$ denotes the gradient of V^* with respect to the vector x , and $o(\Delta t)$ denotes higher-order terms in Δt . Cancelling $V^*(x)$ on both sides and taking the limit as Δt goes to zero yields

$$\min_{u(t)} \left\{ [q(x(t)) + u^T(t)u(t)] + \left(\frac{\partial V^*}{\partial x} \right) [f(x(t)) + g(x(t))u(t)] \right\} = 0. \quad (2.6)$$

The boundary condition for this equation is given by $V^*(0) = 0$ where $V^*(x)$ must be positive for all x (since it corresponds to the optimal *cost* which must be positive).

Equation (2.6) is one form of the so-called Hamilton-Jacobi-Bellman equation. In many cases, this is not the final form of the equation. Two more steps can often be performed to reach a more convenient representation of the Hamilton-Jacobi-Bellman equation.

1. First, the indicated minimization is performed, leading to a control law of the form

$$u^* = -\frac{1}{2}g^T(x)\frac{\partial V^*}{\partial x}. \quad (2.7)$$

2. The second step is to substitute (2.7) back into (2.6), and solve the resulting nonlinear partial differential equation

$$\frac{\partial V^*}{\partial x} f(x) - \frac{\partial V^*}{\partial x} g(x)g^T(x)\frac{\partial V^*}{\partial x} + q(x) = 0 \quad (2.8)$$

for $V^*(x)$.

Equation (2.8) is what we will often refer to as the Hamilton-Jacobi-Bellman (HJB) equation. The actual calculation of the optimal control action proceeds in an

opposite fashion to the steps given above. First the HJB equation (2.8) is solved for V^* , then this is substituted into (2.7) where we obtain the optimal control action that achieves this minimal performance.

Properties of the HJB solution

There are some important aspects of the HJB solution that should be highlighted for clarity. We consider them below:

Closed Loop: The resulting solution is a *state feedback* control law as given in (2.7).

Global: The solution provides the optimal control trajectory from every initial condition. Hence, it solves the optimal control problem for every initial condition, all at once.

Sufficient: The solution of the HJB equation provides a sufficient condition for the solution to the corresponding optimal control problem.

Finally, perhaps the most important remark to make about the HJB equation (2.8) is that in general it is computationally intractable. This single fact is in large part the reason for the existence of the discipline of nonlinear optimal control. Hence, from one point of view, nonlinear optimal control can be thought of as the development of computationally tractable sub-optimal solutions to the optimal control problem. This explanation is attractive from a pedagogical viewpoint because it provides a natural justification for the tight connection between many popular approaches and the HJB equation.

2.3 Calculus of variations:

Euler-Lagrange equations

In this section we solve an optimal control problem by the techniques of classical variational calculus, leading to a derivation of the *Euler-Lagrange* equations. The optimal control problem solved in this section is *not* equivalent to that solved in the previous

section by dynamic programming techniques. Nevertheless, it will be important for two reasons. First, this problem will be used to motivate the introduction of receding horizon control. Secondly, it helps to illustrate some of the fundamental differences between the dynamic programming and calculus of variations approaches.

The *Euler-Lagrange* solution results by considering the optimal control problem in the framework of a constrained optimization:

$$\min_{u(\cdot)} \int_0^T (q(x) + u^T u) dt + \varphi(x(T)) \quad (2.9)$$

$$s.t. \quad \dot{x} = f(x) + g(x)u \quad (2.10)$$

$$x(0) = x_0 \quad (2.11)$$

Before proceeding, note the following two differences between this problem and that solved in the previous section. The objective function is based on a finite horizon length with a terminal weight $\varphi(\cdot)$ applied at the end of the horizon. (This cost is equivalent to an infinite horizon cost only when the terminal weight is chosen as the value function, i.e., $\varphi(\cdot) = V^*(\cdot)$, which can only be found from the solution to the HJB equation.) Secondly, in addition to viewing the dynamics as a constraint, a specific initial condition is imposed.

The calculus of variations solution can be thought of as a standard application of the necessary conditions for constrained optimization, the only twist being that the optimization is infinite dimensional. Hence, the first step is to use Lagrange multipliers to adjoin the constraints to the performance index. Since the constraints are determined by the system differential equation (2.10) and represent *equality* constraints that must hold at each instant in time, an associated multiplier $\lambda(t) \in \mathbb{R}^n$ is a function of time. Thus the augmented performance index is given by

$$\begin{aligned} & \int_0^T (q(x(t)) + u^T(t)u(t)) dt + \varphi(x(T)) + \int_0^T \lambda^T(t)(f(x(t)) + g(x(t))u(t) - \dot{x}) dt \\ & = \varphi(x(T)) + \int_0^T [q(x(t)) + u^T(t)u(t) + \lambda^T(t)(f(x(t)) + g(x(t))u(t) - \dot{x})] dt. \end{aligned} \quad (2.12)$$

Defining, for convenience, the following scalar function H , called *the Hamiltonian*,

$$H(x(t), u(t), \lambda(t)) = q(x(t)) + u^T(t)u(t) + \lambda^T(t)(f(x(t)) + g(x(t))u(t)) \quad (2.13)$$

and integrating the last term on the right side of (2.12) by parts yields

$$\varphi(x(T)) - \lambda^T(T)x(T) + \lambda^T(0)x(0) + \int_0^T [H(x(t), u(t), \lambda(t)) + \dot{\lambda}^T(t)x(t)]dt. \quad (2.14)$$

According to the theory of Lagrange multipliers, the problem of determining the control function $u(t)$ that minimizes the original performance index subject to the constraints (2.10) has been converted to the problem of finding stationary points of (2.14) without constraints.

Now consider the equation for variations of (2.14) with respect to $x(t)$ and $u(t)$

$$\left[\left(\frac{\partial \varphi}{\partial x} - \lambda^T \right) \delta x \right]_{t=T} + [\lambda^T \delta x]_{t=0} + \int_{t_0}^{t_f} \left[\left(\frac{\partial H}{\partial x} + \dot{\lambda}^T \right) \delta x + \frac{\partial H}{\partial u} \delta u \right] dt. \quad (2.15)$$

For a stationary point, it is required that this be equal to zero for all allowable variations. First, looking at the variation δx , in order to cause the coefficients of δx in (2.15) to vanish, the multiplier functions $\lambda(t)$ have to be chosen according to

$$\dot{\lambda}^T = -\frac{\partial H}{\partial x}, \quad 0 \leq t \leq T \quad (2.16)$$

with boundary condition

$$\lambda^T(T) = \frac{\partial \varphi}{\partial x} \Big|_{t=T}. \quad (2.17)$$

Equation (2.15) then becomes

$$\lambda^T(0)\delta x(0) + \int_0^T \left[\frac{\partial H}{\partial u} \delta u \right] dt. \quad (2.18)$$

Now, since in this problem the initial condition is given and fixed, this implies $\delta x(0) = 0$. Finally, since for a stationary point the variation must be zero for arbitrary $\delta u(t)$,

the following must be satisfied

$$\frac{\partial H}{\partial u} = 0, \quad 0 \leq t \leq T. \quad (2.19)$$

The above equations, (2.16), (2.17), and (2.19), plus the original dynamics and initial condition, represent necessary conditions for optimality known as the *Euler-Lagrange equations*. These equations are used to design the control $u(t)$ that minimizes the performance index, and can be summarized as follows:

$$\dot{x} = f(x) + g(x)u \quad (2.20)$$

$$\dot{\lambda} = - \left(\frac{\partial H}{\partial x} \right)^T \quad (2.21)$$

$$\frac{\partial H}{\partial u} = 0 \quad (2.22)$$

with boundary conditions

$$x(0) \quad \text{given} \quad (2.23)$$

$$\lambda(T) = \left(\frac{\partial \varphi}{\partial x} \right)^T \Big|_{t=T}. \quad (2.24)$$

The optimizing control action $u^*(t)$ is determined by

$$u^*(t) = \arg \min_u H(x^*(t), u, \lambda^*(t)) \quad (2.25)$$

where $x^*(t)$ and $\lambda^*(t)$ denote the solution corresponding to the optimal trajectory.

Note that the Lagrange multiplier $\lambda(t)$ is a dynamical variable that satisfies its own dynamical equation (2.21), the so-called *costate* or *adjoint* equation that evolves *backward* in time (by defining the backward time variable $\tau = T - t$ it follows that $d\tau = -dt$), with the final condition $\lambda(T)$ given by equation (2.24).

The Euler-Lagrange equations are coupled ordinary differential equations with two-point boundary conditions. That is, they are expressed by the state equation

(2.20) with initial condition (2.23) and the costate equation (2.21) with final condition (2.24). The optimal control $u(t)$ is then generally determined in terms of $x(t)$ and $\lambda(t)$ by using the *stationarity condition* given by (2.22). This condition guarantees a stationary point with respect to changes in $u(t)$. Finally, expression (2.25) does not yield an optimal control *feedback* law, but an optimal *open-loop* control (time function).

Properties of the EL solution

In contrast to the HJB solution to the infinite horizon optimal control problem, the Euler-Lagrange solution is characterized as follows:

Open-Loop: The resulting optimal trajectory is explicitly solved for as a function of time $u(t)$, not as a feedback law.

Local: The resulting solution is only valid for the specified initial condition $x(0)$. When a new initial condition is specified, the problem must be resolved.

Necessary: Since the Euler-Lagrange equations specify the conditions for the existence of a stationary point, they represent necessary conditions for an optimal trajectory.

2.4 Summary

We have outlined the two basic approaches to problems of optimal control, highlighting the differences in their basic approach and in the properties of their solutions. These differences are summarized in Figure 2.1.

A dynamic programming approach to the problem of optimal control leads to a derivation of the Hamilton-Jacobi-Bellman equation. It provides a global control law in the form of a state feedback controller. Unfortunately, it involves the solution of a partial differential equation, which is in general computationally intractable.

The calculus of variations solution, on the other hand, only requires the solution to a two-point boundary value ordinary differential equation, known as the Euler-Lagrange equations. While still presenting a challenge, this is tractable when com-

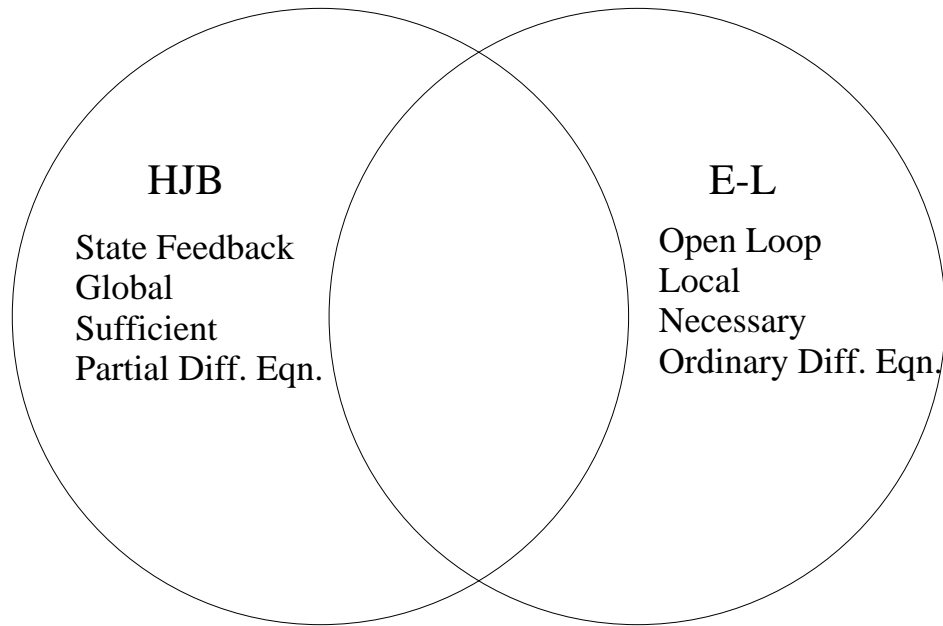


Figure 2.1: Hamilton-Jacobi-Bellman vs. Euler-Lagrange Approach.

pared to the HJB partial differential equation. But, this solution is not equivalent to that given by the Hamilton-Jacobi-Bellman equation. The Euler-Lagrange equations solve instead a trajectory optimization problem. That is, they provide an open-loop trajectory corresponding to a specific initial condition. Hence, computational tractability is traded for the lack of a global solution.

A deep understanding of these two viewpoints toward the optimal control problem provides the proper background and context in which to interpret a number of suboptimal strategies.

Chapter 3 Control Lyapunov Function Techniques

3.1 Introduction

The optimal control of nonlinear systems is one of the most challenging and difficult subjects in control theory. As detailed in the previous chapter, it is well known that the nonlinear optimal control problem can be reduced to the Hamilton-Jacobi-Bellman partial differential equation [BH75], but due to difficulties in its solution, this is not a practical approach. Instead, the search for nonlinear control schemes has generally been approached on less ambitious grounds than requiring the exact solution to the Hamilton-Jacobi-Bellman partial differential equation.

In fact, even the problem of stabilizing a nonlinear system remains a challenging task. Lyapunov theory, a successful and widely used tool, is a century old. Despite this, there still do not exist systematic methods for obtaining Lyapunov functions for general nonlinear systems. Nevertheless, the ideas put forth by Lyapunov nearly a century ago continue to be used and exploited extensively in the modern theory of control for nonlinear systems. One notably successful use of the Lyapunov methodology is the concept of a *control Lyapunov function* (CLF) [Son83, Son89, FK95, FP96, KKK95, FK96b, FK96a], the idea of which is to first choose a function which can be made into a Lyapunov function for the closed loop system by choosing appropriate control actions. The knowledge of such a function is then used to design control laws. Once again, there do not exist systematic techniques for finding CLFs for general nonlinear systems, but this approach has been applied successfully to many classes of systems for which CLFs can be found (feedback linearization, back-stepping, forwarding [KKK95, FK96b, FK95]).

In this chapter we focus on methods for producing a control law once a CLF

has been derived. Specifically, we explore the connection between Sontag’s formula [Son89], pointwise min-norm controllers [FK96a], and the Hamilton-Jacobi-Bellman point of view. Since both Sontag’s formula and pointwise min-norm controllers are suboptimal, they select a control policy by prioritizing and trading-off properties of the optimal state feedback controller that they seek to approximate. CLF based techniques first require stability. This is guaranteed completely by the fact that a CLF exists, and leaves extra degrees of freedom in the choice of the specific control policy. We show how different approaches tie these extra degrees of freedom to the HJB equation, clarifying both their strengths and limitations.

3.2 HJB equations and CLF techniques

Recall that the nonlinear system under consideration is given by

$$\dot{x} = f(x) + g(x)u \quad f(0) = 0 \quad (3.1)$$

with $x \in \mathbb{R}^n$ denoting the state, $u \in \mathbb{R}^m$ the control, and $f(x)$ and $g(x)$ are \mathbf{C}^1 . The objective function is

$$\begin{aligned} \min_{u(\cdot)} \quad & \int_0^\infty (q(x) + u^T u) dt \\ \text{s.t.} \quad & \dot{x} = f(x) + g(x)u \end{aligned} \quad (3.2)$$

for $q(x) \in \mathbf{C}^1$, positive semi-definite and the desired solution being a *state feedback* control law. We have also assumed that the system $[f(x), q(x)]$ is *zero-state detectable*.

The solution to this problem is

$$u^* = -\frac{1}{2}g^T \frac{\partial V^*}{\partial x}$$

where V^* solves the HJB equation

$$\frac{\partial V^*}{\partial x} f - \frac{1}{4} \frac{\partial V^*}{\partial x} g g^T \frac{\partial V^*}{\partial x} + q = 0 \quad (3.3)$$

and is the minimum “cost to go,” which is commonly referred to as the *value function*

$$V^*(x(0)) = \min_u \int_0^\infty (q(x) + u^T u) dt.$$

In what follows we develop connections between nonlinear control techniques based on a control Lyapunov functions, and the HJB approach to the optimal control problem. When a CLF is viewed beyond a mere Lyapunov stability framework, as an approximation to the value function V^* , many CLF approaches have natural derivations from the HJB framework. We pursue these connections here, focusing the majority of our attention on Sontag’s formula [Son89] and pointwise min-norm controllers [FK96a].

3.2.1 Control Lyapunov Functions (CLFs)

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

$$\inf_u \left[\frac{\partial V}{\partial x} f(x) + \frac{\partial V}{\partial x} g(x)u \right] < 0 \quad (3.4)$$

for all $x \neq 0$ [Art83, Son83, Son89]. This definition is motivated by the following consideration. Assume we are supplied with a positive definite function V and asked whether this function can be used as a Lyapunov function for a system we would like to stabilize. To determine if this is possible, we would calculate the time derivative of this function along trajectories of the system, i.e.

$$\dot{V}(x) = \frac{\partial V}{\partial x} [f(x) + g(x)u].$$

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 controlled system under the chosen control actions. This is exactly the condition given in (3.4).

Given a general system of the form (3.1), it may be difficult to find a CLF or even to determine whether one exists. Fortunately, there are significant classes of systems for which the systematic construction of a CLF is possible (back-stepping, feedback linearization, forwarding, LPV, etc.). This has been explored extensively in the literature ([KKK95, FK96b, FK95] and references therein). We will not concern ourselves with this question. Instead, we will pay particular attention to techniques for designing a stabilizing controller once a CLF has been found, and their relationship to the nonlinear optimal control problem.

3.2.2 The value function as a CLF

First, let us understand why it is reasonable to view the value function as a CLF. Rewriting the HJB equation (3.3) as

$$\frac{\partial V^*}{\partial x} \left[f - \frac{1}{2} g g^T \frac{\partial V^{*T}}{\partial x} \right] + \frac{1}{4} \frac{\partial V^*}{\partial x} g g^T \frac{\partial V^{*T}}{\partial x} + q = 0$$

and recalling that

$$u^* = -\frac{1}{2} g^T \frac{\partial V^{*T}}{\partial x}$$

allows us to reformulate (3.3) as

$$\frac{\partial V^*}{\partial x} [f + g u^*] = - \left(\frac{1}{4} \frac{\partial V^*}{\partial x} g g^T \frac{\partial V^{*T}}{\partial x} + q \right).$$

Note that now the left-hand side appears as in the definition of a control Lyapunov function (cf. (3.4)). Hence, if the right-hand side is negative, then V^* is a control Lyapunov function. Technically, the right-hand side need only be negative semi-definite¹ and hence the value function may only be a so-called weak CLF. Of course, for any positive definite cost parameter q , this equation shows that V^* is in fact a

¹This is why it is necessary to impose zero-state detectability to ensure stability of the optimal closed loop system.

strict CLF. It is important to keep this connection in mind as we proceed, because many CLF based techniques can be viewed as assuming that a CLF is an estimate of the value function, which is ideal for performance purposes.

3.2.3 CLF a substitute for the value function:

Sontag's formula

It can be shown that the existence of a CLF for the system (3.1) is equivalent to the existence of a globally asymptotically stabilizing control law $u = k(x)$ which is continuous everywhere except possibly at $x = 0$ [Art83]. Moreover, one can calculate such a control law k explicitly from f , g and V . Perhaps the most important formula for producing a stabilizing controller based on the existence of a CLF was introduced in [Son89] and has come to be known as Sontag's formula. We will consider a slight variation of Sontag's formula (which we will continue to refer to as Sontag's formula with slight abuse), originally introduced in [FP96]:

$$u_{\sigma_s} = \begin{cases} - \left[\frac{\frac{\partial V}{\partial x} f + \sqrt{\left(\frac{\partial V}{\partial x} f\right)^2 + q(x) \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V}{\partial x}^T\right)}}{\frac{\partial V}{\partial x} g g^T \frac{\partial V}{\partial x}^T} \right] g^T \frac{\partial V}{\partial x}^T & \frac{\partial V}{\partial x} g \neq 0 \\ 0 & \frac{\partial V}{\partial x} g = 0. \end{cases} \quad (3.5)$$

(The use of the notation u_{σ_s} will become clear later.) While this formula enjoys similar continuity properties to those for which Sontag's formula is known (i.e., for $q(x)$ positive definite it is continuous everywhere except possibly at $x = 0$ [Son89]), for us its importance lies in its connection with optimal control. At first glance, one might note that the cost parameter associated with the state, $q(x)$ (refer to eqn. (3.2)), appears explicitly in (3.5). In fact, the connection runs much deeper and our version of Sontag's formula has a strong interpretation in the context of Hamilton-Jacobi-Bellman equations.

Optimality, Sontag's formula and level curves

Below, we unravel some key connections between level curves of the value function V^* and Sontag's formula (3.5). It is shown that Sontag's formula, in essence, uses the directional information supplied by a CLF, V , and scales it properly to solve the HJB equation. In particular, if V has level curves that agree with those of the value function, then Sontag's formula produces the optimal controller [FP96].

Assume that V is a CLF for the system (3.1). For the sake of motivation, assume that V possesses the same shape level curves as those of the value function V^* . Even though in general V would not be the same as V^* , this does imply a relationship between their gradients. We may assert that there exists a scalar function $\lambda(x)$ such that $\frac{\partial V^*}{\partial x} = \lambda(x) \frac{\partial V}{\partial x}$ for every x (i.e., the gradients point in the same direction at every point). In this case, the optimal control can also be written in terms of the CLF V ,

$$u^* = -\frac{1}{2}g^T \frac{\partial V^{*T}}{\partial x} = -\frac{\lambda(x)}{2}g^T \frac{\partial V^T}{\partial x}. \quad (3.6)$$

Additionally, the HJB equation can be used to determine $\lambda(x)$ by substituting $\frac{\partial V^*}{\partial x} = \lambda(x) \frac{\partial V}{\partial x}$ into the HJB equation (3.3)

$$\lambda \frac{\partial V}{\partial x} f - \frac{\lambda^2}{4} \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x} \right) + q(x) = 0. \quad (3.7)$$

This is a quadratic equation in λ . Solving for λ and taking only the positive square root gives

$$\lambda = 2 \left(\frac{\frac{\partial V}{\partial x} f + \sqrt{[\frac{\partial V}{\partial x} f]^2 + q(x) [\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x}]}}{\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x}} \right). \quad (3.8)$$

Substituting this value into the control u^* given in (3.6) yields

$$u^* = \begin{cases} - \left[\frac{\frac{\partial V}{\partial x} f + \sqrt{(\frac{\partial V}{\partial x} f)^2 + q(x) (\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x})}}{\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x}} \right] g^T \frac{\partial V^T}{\partial x} & \frac{\partial V}{\partial x} g \neq 0 \\ 0 & \frac{\partial V}{\partial x} g = 0, \end{cases}$$

which is exactly Sontag's formula, u_{σ_s} (3.5). In this case, Sontag's formula will result in the optimal controller.

For an arbitrary CLF V , we may still follow the above procedure which results in Sontag's formula. Hence Sontag's formula may be thought of as using the direction given by the CLF (i.e., $\frac{\partial V}{\partial x}$), which, by the fact that it is a CLF will result in stability, but pointwise scaling it by λ so that it will satisfy the HJB equation as in (3.7). Then $\lambda \frac{\partial V}{\partial x}$ is used in place of $(\frac{\partial V^*}{\partial x})$ in the formula for the optimal control u^* , (3.6). Hence, we see that there is a strong connection between Sontag's formula and the HJB equation. In fact, Sontag's formula just uses the CLF V as a substitute for the value function in the HJB approach to optimal control.

Next, we introduce the notion of pointwise minimum norm controllers ([FK96a, FK95, FK96b]), and demonstrate that Sontag's formula is the solution to a specific pointwise minimum norm problem. It is from this framework that connections with optimal control have generally been emphasized.

3.2.4 Pointwise min-norm controllers

Given a CLF, $V > 0$, by definition there will exist a control action u such that $\dot{V} = \frac{\partial V}{\partial x}[f + gu] < 0$ for every x . In general there are many such u that will satisfy $\frac{\partial V}{\partial x}[f + gu] < 0$. One method of determining a specific u is to pose the following optimization problem [FK96a, FK95, FK96b]:

(Pointwise Min-Norm)

$$\text{minimize} \quad u^T u \quad (3.9)$$

$$\text{subject to} \quad \frac{\partial V}{\partial x}[f + gu] \leq -\sigma(x) \quad (3.10)$$

where $\sigma(x)$ is some continuous, positive definite function satisfying $\frac{\partial V}{\partial x} f(x) \leq -\sigma(x)$ whenever $\frac{\partial V}{\partial x} g(x) = 0$, and the optimization is solved pointwise (i.e., for each x). This formula pointwise minimizes the control energy used while requiring that V be a Lyapunov function for the closed loop system and decrease by at least $\sigma(x)$ at

every point. The resulting controller can be solved for off-line and in closed form (see [FK95] for details).

In [FK96a] it was shown that every CLF V is the value function for some meaningful cost functional. In other words, it solves the HJB equation associated with a meaningful cost. This property is commonly referred to as being “inverse optimal” [FK96a]. Note that a CLF V does not uniquely determine a control law because it may be the value function for many different cost functions, each of which may have a different optimal control. What is important is that the pointwise min-norm formulation always produces one of these inverse optimal control laws [FK96a].

To intuitively understand why pointwise min-norm controllers possess such strong connections to HJB equations, let us reconsider the optimization in (3.9), but this time use a Lagrange multiplier to deal with the constraint. Hence, we can write the Lagrangian for the problem as

$$L(u, \lambda) = u^T u + \lambda \left(\frac{\partial V}{\partial x} [f + gu] - \sigma \right)$$

where λ is the Lagrange multiplier (required to be positive, etc., in accordance with the Kuhn-Tucker conditions [KT61]). Lagrangian duality tells us that the optimizing u should minimize the Lagrangian. Furthermore, we can exploit the fact that adding or subtracting terms to the Lagrangian that do not contain u will not effect the solution. So, we will add the term $q(x)$ and subtract the $-\lambda\sigma$ term to obtain

$$\min_u \left\{ [q(x(t)) + u^T(t)u(t)] + \lambda(x(t)) \frac{\partial V}{\partial x} [f(x(t)) + g(x(t))u(t)] \right\} = 0,$$

which is identical to the HJB equation (2.6) except with $\left(\frac{\partial V^*}{\partial x}\right)$ replaced by $\lambda \frac{\partial V}{\partial x}$. Furthermore, by performing the minimization, we find that the resulting state feedback is of the form

$$u_\sigma = -\frac{\lambda}{2} g^T \frac{\partial V}{\partial x}.$$

This is identical to the relationship used to derive Sontag’s formula. Hence, we see that pointwise min-norm formulas are similar to Sontag’s formula in that they

substitute $\lambda \frac{\partial V}{\partial x}$ for the true gradient of the value function $(\frac{\partial V^*}{\partial x})$. The only difference is that pointwise min-norm controllers can use a different criterion to select the scaling λ . This degree of freedom is basically contained in the choice of σ . Therefore, we can view pointwise min-norm formulas as a generalization of Sontag's formula.

We now explicitly derive the parameter $\sigma(x)$ that generates Sontag's formula in the pointwise min-norm formulation. Let us assume that the solution to the above pointwise min-norm problem results in Sontag's formula. It should be clear that for $\frac{\partial V}{\partial x} g \neq 0$, the constraint will be active, since u will be reduced as much as possible. Knowing that u will turn out to be Sontag's formula results in the following value for σ [FP96]:

$$\begin{aligned} -\sigma &= \frac{\partial V}{\partial x}(f + gu_{\sigma_s}) \\ &= \frac{\partial V}{\partial x}f + \frac{\partial V}{\partial x}g \left(-\frac{\frac{\partial V}{\partial x}f + \sqrt{\left(\frac{\partial V}{\partial x}f\right)^2 + q(x)\left(\frac{\partial V}{\partial x}gg^T\frac{\partial V^T}{\partial x}\right)}}{\frac{\partial V}{\partial x}gg^T\frac{\partial V^T}{\partial x}} \right) g^T \frac{\partial V^T}{\partial x} \\ &= -\sqrt{\left(\frac{\partial V}{\partial x}f\right)^2 + q(x)\left(\frac{\partial V}{\partial x}gg^T\frac{\partial V^T}{\partial x}\right)}. \end{aligned}$$

Hence, the special choice of σ (which we denote by σ_s),

$$\sigma_s = \sqrt{\left(\frac{\partial V}{\partial x}f\right)^2 + q(x)\left(\frac{\partial V}{\partial x}gg^T\frac{\partial V^T}{\partial x}\right)} \quad (3.11)$$

in the pointwise min-norm scheme (3.9) results in Sontag's formula. This provides us with an important alternative method for viewing Sontag's formula. It is the solution to the above pointwise min-norm problem with parameter σ_s . Hence, our version of Sontag's formula enjoys all the properties of pointwise min-norm controllers.

We have seen that these CLF based techniques share much in common with the HJB approach to nonlinear optimal control. Nevertheless, the strong reliance on a CLF, while providing stability, can lead to suboptimal performance when applied naively, as demonstrated in the following example.

3.3 Example

Throughout this thesis we will call upon the following example to illustrate key points. Consider a two dimensional nonlinear oscillator

$$\begin{cases} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= -x_1 \left(\frac{\pi}{2} + \arctan(5x_1) \right) - \frac{5x_1^2}{2(1+25x_1^2)} + 4x_2 + 3u \end{cases}$$

with performance index

$$\int_0^{\infty} (x_2^2 + u^2) dt.$$

This example was created using the so-called converse HJB method [DPSN96] so that the optimal solution is known. For this problem, the value function is given by

$$V^* = x_1^2 \left(\frac{\pi}{2} + \arctan(5x_1) \right) + x_2^2$$

which results in the optimal control action

$$u^* = -3x_2.$$

A simple technique for obtaining a CLF for this system is to exploit the fact that it is feedback linearizable [Isi95]. In the feedback linearized coordinates, a quadratic function may be chosen as a CLF. In order to ensure that this CLF will at least produce a locally optimal controller, we choose the quadratic CLF to agree with the quadratic portion of the true value function.² This results in the following CLF

$$V = \frac{\pi}{2} x_1^2 + x_2^2.$$

(This function is actually not a CLF in the strict sense in that there exist points where \dot{V} may only be made equal to zero and not strictly less than zero. This is sometimes referred to as a *weak* CLF. Nevertheless, we will use this CLF since it is

²This can be done without knowledge of the true value function by performing Jacobian linearization and designing an LQR optimal controller for the linearized system.

the only quadratic function that locally agrees with our value function (which itself is not even a strict CLF for this system). Furthermore, asymptotic stability under Sontag's formula is guaranteed by LaSalle's invariance principle.)

We will compare Sontag's formula using this CLF to the performance of the optimal controller. Figure 3.1 is a plot of the level curves of the true value function V^* versus those of the CLF V . Clearly, these curves are far from the level curves of a quadratic function. Since Sontag's formula uses the directions provided by the CLF, one might suspect that Sontag's formula with the quadratic CLF given above will perform poorly on this system.




Figure 3.1: Contours of the value function (solid) and CLF (dashed).

This is indeed the case, as shown in Figure 3.2 where Sontag's formula (dotted) accumulates a cost of over 250 from the initial condition $[3, -2]$. The costs achieved by Sontag's formula and the optimal controller from the initial condition $[3, -2]$ are

summarized in Table 3.1.




Figure 3.2: Phase Portrait: Optimal (solid), Sontag's (dashed).

Table 3.1

Controller	Cost
Sontag	258
Optimal	31.7

Table 3.1: Cost of Sontag's formula vs. the optimal from the initial condition $[3, -2]$.

This example shows that CLF based designs can be particularly sensitive to differences between the CLF and the value function, even for a technique such as Sontag's formula that directly incorporates information from the optimal control problem into the controller design process.

One might note that we have naively utilized the CLF methodology without thought as to how to better craft a more suitable and sensible CLF for the this

problem. In this simple example, it is not too difficult to iterate on the selection of parameters and find a controller that performs admirably. Nevertheless, it exactly illustrates the more subtle issues involved in CLF design that often require experience and expertise to be able to modify the methodology on a problem by problem basis.

3.4 Summary

Control Lyapunov functions are best interpreted in the context of Hamilton-Jacobi-Bellman equations, especially a variation of Sontag's formula that naturally arises from HJB equations and furthermore is a special case of a more general class of CLF based controllers known as pointwise min-norm controllers. Even with strong ties to the optimal control problem, CLF based approaches err on the side of stability and can result in poor performance when the CLF does not closely resemble the value function.

In terms of the overall picture of nonlinear optimal control, control Lyapunov functions sit squarely on the side of the Hamilton-Jacobi-Bellman equation. Figure 3.3 shows this pictorially. In the following chapter we will fill in the Euler-Lagrange side with so-called receding horizon techniques.

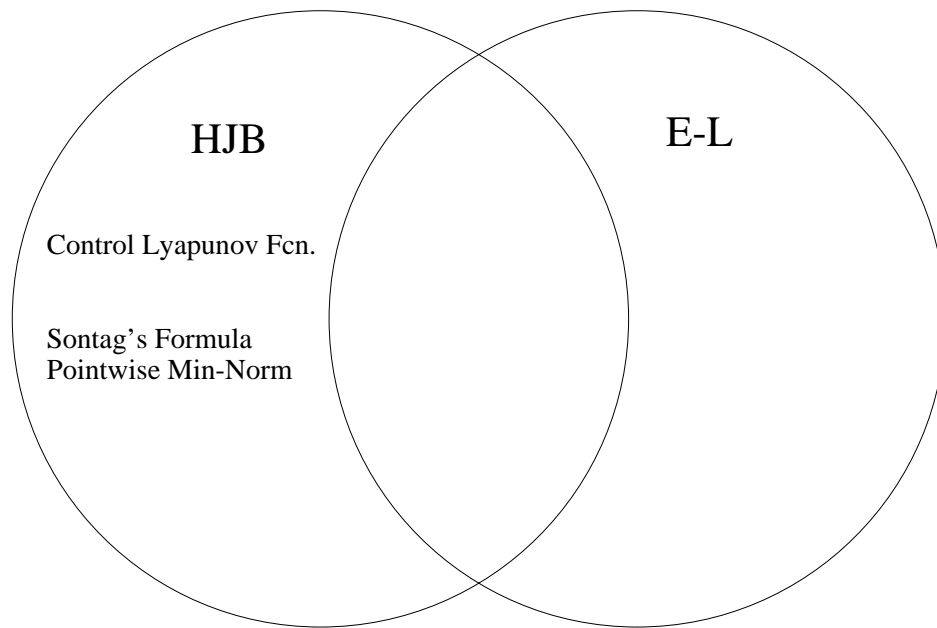


Figure 3.3: CLFs within the optimal control picture.

Chapter 4 Receding Horizon Control

4.1 Introduction

In contrast to the emphasis on guaranteed stability that is the primary goal of CLFs, another class of nonlinear control schemes that goes by the names receding horizon, moving horizon, or model predictive control places importance on optimal performance [KP77, KBK83, MM90, GPM89, KG88]. These techniques apply a receding horizon implementation in an attempt to approximately solve the optimal control problem through on-line computation. The receding horizon methodology is to solve a trajectory optimization emanating from the current state, and implement the resulting open-loop solution until a new state update is received and the process is repeated. For systems under which on-line computation is feasible, receding horizon control (RHC) has proven quite successful [RRTP78, Ric93], but guaranteed stability has remained a concern for some time.

In this chapter we argue that some of these difficulties are rooted in the fact that receding horizon control adopts an Euler-Lagrange framework for optimal control, but translates it to the desired state feedback solution by employing the receding horizon methodology. Yet, it still inherits the properties of the Euler-Lagrange solution to trajectory optimizations, namely that each solution only provides information for a specific initial condition and trajectory and hence this leads to difficulties when attempting to establish properties such as stability. We begin our exploration by defining exactly what we mean by receding horizon control.

4.2 Receding Horizon Control (RHC)

Receding horizon techniques (cf. [KP77, KBK83, GPM89, MM90, KG88, GPM89]) are based upon using on-line computation to repeatedly solve optimal control prob-

lems emanating from the current measured state [KP77, KBK83, GPM89, MM90, KG88, GPM89]. To be more specific, the current control at state x and time t is obtained by determining on-line the optimal control \hat{u} over the interval $[t, t + T]$ respecting the following objective:

(Receding Horizon Control)

$$\text{minimize} \quad \varphi(x(t + T)) + \int_t^{t+T} (q(x(\tau)) + u^T(\tau)u(\tau))d\tau \quad (4.1)$$

$$\text{subject to} \quad \dot{x} = f(x) + g(x)u \quad (4.2)$$

and implementing the optimizing solution $\hat{u}(\cdot)$ until a new state update is received. Note that this optimization uses a finite horizon T and a terminal weight $\varphi(\cdot)$, and hence is solved as an Euler-Lagrange type trajectory optimization with the current state measurement $x(t)$ serving as the initial condition. Repeating these calculations for every new measured state yields a state feedback control law. As is evident from this sort of control scheme, obtaining a reduced value of the performance index is of utmost importance.

The philosophy behind receding horizon control is to exploit the simplicity of the Euler-Lagrange approach to optimal control as compared to the HJB approach. Unfortunately, the Euler-Lagrange solution is valid only for a single initial condition and produces an open-loop trajectory. This is in contrast to a desired state feedback law. Hence, this is overcome by *resolving* an Euler-Lagrange type optimization at every encountered state, producing a state feedback. This is possible due to the local, open-loop nature of the Euler-Lagrange formulation which makes it computationally much simpler than the HJB equation. Furthermore, this methodology only requires that the optimal control problem be solved for the states encountered along the current trajectory, again circumventing the global nature of the HJB approach and its associated computational intractability.

In general, the solution to each receding horizon optimization provides an approximation to the value function at the current state, as well as an accompanying open-loop control trajectory, but this information is specific to the current state and

indicates that difficulties may arise when considering properties such as stability which are typically established for regions. In following sections we explore this and other associated difficulties encountered in the receding horizon framework.

4.2.1 Computational issues

Despite the computational advantages of an Euler-Lagrange approach over those of the HJB viewpoint, the on-line implementation of receding horizon control is still computationally demanding. In fact, the practical implementation of receding horizon control is often hindered by the computational burden of the on-line optimization which, in some theoretical settings, must be solved continuously [MM90]. In reality, the optimization is most commonly solved at discrete sampling times and the corresponding control moves are applied until they can be updated at the next sampling instance. The choice of both the sampling time and horizon are largely influenced by the ability to solve the required optimization within the allowed time interval. These considerations often limit the application of receding horizon control to systems with sufficiently slow dynamics to be able to accommodate such on-line inter-sample computation. Applications in the process industries represent the most prominent examples of the successes of the receding horizon methodology [RRTP78, Ric93, CR79].

For linear systems under quadratic objective functions, the on-line optimization is reduced to a tractable quadratic program, even in the presence of linear input and output constraints. This ability to incorporate constraints was the initial attraction of receding horizon control. For nonlinear systems the optimization is in general non-convex and hence has no efficient solution. There are a number of different approaches to the problem of implementing nonlinear receding horizon control. Below we expound on the most common and promising of these approaches.

Standard nonlinear receding horizon control

Nonlinear receding horizon control relies on standard techniques for solving trajectory optimization problems of the form in (4.1-4.2). These include both direct and indirect approaches relying typically on either shooting or collocation techniques [BH75]. Such

an optimization is generally non-convex. Beyond this, the major difficulty is that each evaluation of the performance index requires the simulation of nonlinear dynamics, which is computationally burdensome.

Stabilized continuation techniques

One approach to counter the computational difficulties of the optimizations in receding horizon control is the stabilized continuation method, where the boundary constraint from the Euler-Lagrange equations (2.24) at the end of the horizon (i.e., $\lambda(t + T) = \frac{\partial \varphi}{\partial x} \Big|_{x(t+T)}$) is treated as a function of the initial conditions ($x(t)$ and $\lambda(t)$) and dynamically stabilized rather than imposed. This allows the solution to the Euler-Lagrange equations to be propagated as a function of the initial condition $x(t)$ by differential equations. This is one approach to a continuous time implementation of receding horizon control, in contrast to the majority of other techniques that resolve each receding horizon optimization at discrete sampling times. Details can be found in [OF94b, OF94a, OF96, Oht96, RD83]. Most other approaches attempt to ease the computational burden of the optimization by simplifying the dynamics in some way.

Receding horizon control with feedback linearization

The idea of this approach is to transform the dynamics to those of a linear system through feedback linearization. Since linear systems can be integrated efficiently and accurately, this can dramatically improve the speed of the receding horizon optimizations [DRCN93, NM95, Nev97]. Note that this also involves a transformation of the cost and constraints, often resulting in state dependence. If desired, one may attempt to approximate the transformed cost and constraints by a quadratic cost and linear constraint, in which case the nonlinear receding horizon control problem is approximated completely by a linear problem [PN97b]. When the transformation to linear coordinates is well conditioned, the approach can be quite successful, but approximations can be inaccurate when the transformation is poorly conditioned. The idea of feedback linearization has even been found to be computationally beneficial when only a portion of the plant can be linearized [PN97b].

Gain scheduled receding horizon control

Gain scheduled receding horizon control simply uses linear approximations to the dynamics in the receding horizon optimization. With quadratic costs and linear constraints, this reduces each optimization to the standard optimization in linear receding horizon control (a quadratic programming problem). This approach can be found in Garcia [Gar84].

Receding horizon control with time-varying linear models

By using the solution to the current receding horizon optimization as a candidate trajectory for the receding horizon optimization at the next state measurement, linearization about this trajectory provides a time-varying linear model. In this approach, each receding horizon optimization is computed with respect to these time-varying linear dynamics instead of the true nonlinear dynamics. This fact can be used to reduce the optimization (under quadratic cost and linear constraints) to a large quadratic program which can be efficiently solved. This idea was introduced by Nevistić [Nev97].

While the numerical and practically oriented issues are compelling, there are fundamental issues related to the theoretical foundations of receding horizon control that deserve equal scrutiny. The most critical of these are well illustrated by considering the stability and performance properties of idealized receding horizon control.

4.2.2 Stability

While using a numerical optimization as an integral part of the control scheme allows great flexibility, especially concerning the incorporation of constraints, it complicates the analysis of stability and performance properties of receding horizon control immensely. Beyond limitations imposed by the Euler-Lagrange philosophy, additional difficulties arise as well. Since the control action is determined through a numerical on-line optimization at every sampling point, there is often no closed form expression for the controller or for the resulting closed loop system.

The lack of a complete theory for a rigorous analysis of receding horizon stability properties in nonlinear systems often leads to the use of intuition in the design process. Unfortunately, this intuition can be misleading. Consider, for example, the statement that horizon length provides a tradeoff between the issues of computation and of stability and performance. A longer horizon, while being computationally more intensive for the on-line optimization, will provide a better approximation to the infinite horizon problem and hence the controller will inherit the stability guarantees and performance properties enjoyed by the infinite horizon solution. While this intuition is correct in the limit as the horizon tends to infinity [PN97a], for horizon lengths applied in practice the relationship between horizon and stability is much more subtle and often contradicts such seemingly reasonable statements. This is best illustrated by the example used previously in Section 3.3. Recall that the system dynamics were given by

$$\begin{cases} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= -x_1 \left(\frac{\pi}{2} + \arctan(5x_1) \right) - \frac{5x_1^2}{2(1+25x_1^2)} + 4x_2 + 3u \end{cases}$$

with performance index

$$\int_0^{\infty} (x_2^2 + u^2) dt.$$

For simplicity we will consider receding horizon controllers with no terminal weight (i.e., $\varphi(x) = 0$) and use a sampling interval of 0.1. By investigating the relationship between horizon length and stability through simulations from the initial condition $[3, -2]$, a puzzling phenomena is uncovered. Beginning from the shortest horizon simulated, $T = 0.2$, the closed loop system is found to be unstable (see Figure 4.1). As the horizon is increased to $T = 0.3$, the results change dramatically and near optimal performance is achieved by the receding horizon controller. At this point, one might be tempted to assume that a sufficient horizon for stability has been reached and longer horizons would only improve the performance. In actuality the opposite happens and as the horizon is increased further, the performance deteriorates and returns to instability by a horizon of $T = 0.5$. This instability remains present even

past a horizon of $T = 1.0$. The simulation results are summarized in Table 4.1 and Figure 4.1.




Figure 4.1: RHC for various horizon lengths.

It is important to recognize and understand that the odd behavior we have encountered is not a nonlinear effect, nor the result of a cleverly chosen initial condition or sampling interval, but rather inherent to the receding horizon approach. In fact, the same phenomena takes place even for the linearized system

$$\dot{x} = \begin{bmatrix} 0 & 1 \\ -\frac{\pi}{2} & 4 \end{bmatrix} x + \begin{bmatrix} 0 \\ 3 \end{bmatrix} u. \quad (4.3)$$

In this case, a more detailed analysis of the closed loop system is possible due to the fact that the controller and closed loop system are linear and can be computed in closed form. Figure 4.2 shows the magnitude of the maximum eigenvalue of the

Table 4.1

Controller	Performance
$T = 0.2$: (dotted)	unstable
$T = 0.3$: (dash-dot)	33.5
$T = 0.5$: (dashed)	unstable
$T = 1.0$: (solid)	unstable

Table 4.1: Comparison of controller performance from initial condition $[3, -2]$.

closed loop versus the horizon length of the receding horizon controller.¹ This plot shows that stability is only achieved for a small range of horizons that include $T = 0.3$ and longer horizons lead to instability. It is not until a horizon of $T = 3.79$ that the controller becomes stabilizing once again.

4.2.3 Approaches to guaranteed stability

The stability problems demonstrated in the previous section are not new in the receding horizon control community and related phenomena have been noted before by Bitmead et al. in the context of Riccati difference equations [BGPK85]. This delicate relationship between horizon length and stability has been addressed by various means. For linear systems, the literature is well developed (see [GPM89] and references therein) and generally exploits computable properties of linear systems. Nonlinear systems lack exploitable structure in terms of computation, and hence lend themselves to fewer practical stabilizing formulations of receding horizon control.

Proving stability for nonlinear systems ultimately boils down to finding a Lyapunov function. In receding horizon control, the vast majority of stabilizing approaches use the optimal cost of the receding horizon optimization as a Lyapunov function. To make this work, it is necessary to use either a constraint, or terminal weight, or combination of the two that guarantee that each receding horizon optimization has a cost less than that computed at the previous measured state.

¹The receding horizon controller was computed by discretizing the continuous time system using a first order hold and time step of 0.001, and solving the Riccati difference equation. Hence the eigenvalues correspond to a discrete-time system with stability occurring when the maximum eigenvalue has modulus less than 1.




Figure 4.2: Maximum eigenvalue versus horizon length for discretized linear system.

Below, we review some approaches to guaranteeing stability in receding horizon formulations.

Finite horizon with zero end constraint

By requiring that $x(t + T) = 0$ directly in each receding horizon optimization, the control sequence $\hat{u}[t, t + T]$ from each receding horizon optimization will drive the system to the origin. In addition, this control sequence is feasible for the receding horizon optimization at time $t + \Delta t$ (i.e., $\hat{u}[t + \Delta t, t + T]$) and it achieves a cost less than the cost of the optimization at time t . Hence, this cost can be easily shown to be a Lyapunov function. This idea of an end constraint was first introduced by Kwon and Pearson [KP77, KP78] and has been used by others [CS82, MM90] to prove the stability of receding horizon control for nonlinear systems. There has been an attempt to remove such end constraints due to the fact that numerically they can be difficult

to satisfy, and appear artificial since they are not achieved in closed loop. In fact, ideally one would choose a terminal weight for each receding horizon optimization that is equal to the value function. Heuristically, an end constraint $x(t + T) = 0$ is akin to an infinite terminal weight, clearly a poor choice.

Infinite horizon for open-loop stable systems

The name infinite horizon is actually a bit misleading since the control variables are only optimized over a finite horizon $[t, t + T]$. But the name infinite horizon comes from the fact that the terminal weight is chosen as the open-loop infinite horizon cost, and hence each receding horizon optimization can be thought of as an infinite horizon optimization, but where the control actions can only be chosen over $[t, t + T]$. That is, the receding horizon objective is

$$\int_t^{t+T} (q(x) + u^T u) dt + \int_{t+T}^{\infty} q(x) dt.$$

Once again stability follows from the fact that the optimal control sequence from time t provides a feasible trajectory beginning at time $t + \Delta t$ and achieves a lower cost. This idea was first introduced by Rawlings and Muske in the context of linear systems [RM93]. For nonlinear systems, difficulty arises because the terminal weight $\varphi(x(t+T))$ must be the cost accumulated by the open-loop system. For linear systems, this is easily obtained from the Lyapunov equation, but for nonlinear systems no easily computable formula exists.

Hybrid (dual mode) receding horizon control

Hybrid receding horizon control, or dual mode MPC, is an attempt to relax the restrictiveness of end constraints. While end constraints require that by the end of the horizon the origin has been reached ($x(t + T) = 0$), hybrid receding horizon control only requires that $x(t + T)$ lie within a pre-specified set W around the origin. It is assumed that inside the set W , a stabilizing controller is known. Therefore, the receding horizon controller only attempts to bring the state into W , from which the controllers are switched and the local stabilizing controller takes over. In this

case, the cost used as a Lyapunov function is the cost to arrive in the set W . This Lyapunov function proves that the system enters W (in finite time). This idea was introduced by Michalska and Mayne [MM93] for nonlinear systems.

Quasi-infinite prediction horizon

This approach is both a generalization of infinite horizon receding horizon control, and another approach to the dual mode concept. Once again, it is assumed that a stabilizing controller is known locally within a set W . Instead of switching to this controller, as is done in hybrid receding horizon control, the infinite horizon cost of this controller (in the set W) is computed, off-line, and used as a terminal weight for the receding horizon problem (4.1–4.2). Furthermore, a constraint is added to the receding horizon problem that requires the final state $x(t + T)$ to lie within W . This is basically equivalent to pre-stabilizing the system and applying the infinite horizon results. In this case, no switching is required, regardless of whether the states are inside or outside of W . Chen and Allgöwer proposed this approach [CA96].

Finally, we mention one approach not based upon the idea of using the cost as a Lyapunov function.

Contractive receding horizon control

Contractive receding horizon control simply imposes a constraint to each receding horizon optimization that a norm of the state has contracted over the sampling interval. In essence, this is imposing a Lyapunov function on the closed loop system. This idea was introduced by De Oliveira and Morari [DOM96].

Other variations of these techniques continue to be developed, with both implementability and stability as the motivating factors.

4.3 Summary

Receding horizon control, which is based on the repeated on-line solution of open-loop trajectory optimal control problems, closely relates to an Euler-Lagrange framework.

The intractability of the HJB equations are overcome by solving for the optimal control only along the current trajectory through on-line computation. This approach chooses to err on the side of performance and in its purest form lacks guaranteed stability properties. Stability and performance concerns become even more critical when short horizons must be used to accommodate the extensive on-line computation required. This has led to the development of stabilizing receding horizon formulations. They typically involve the alteration of the receding horizon optimization to guarantee that its optimal value can be used as a Lyapunov function, imitating the fact that the value function is a Lyapunov function in the HJB framework.

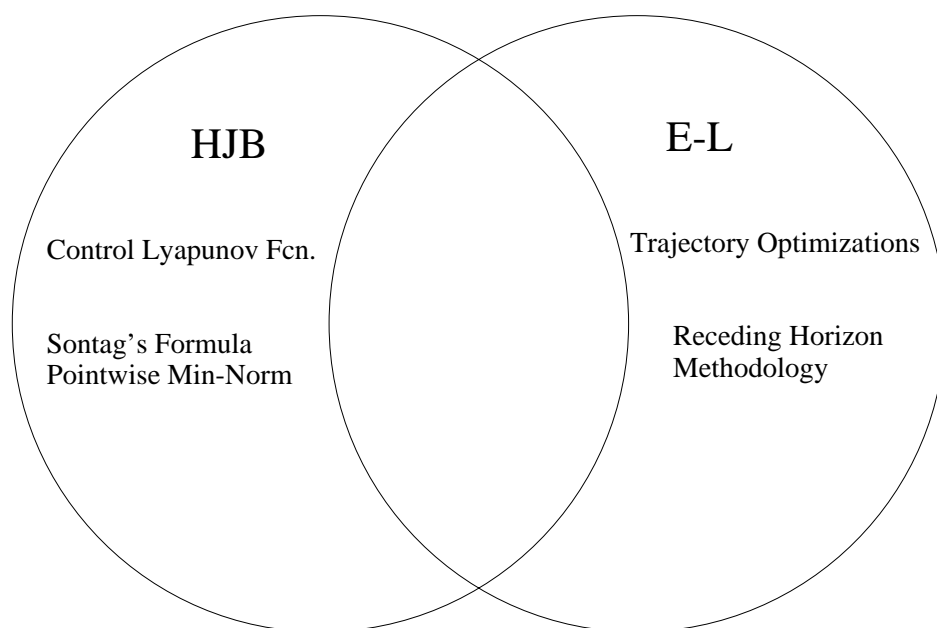


Figure 4.3: RHC within the optimal control picture.

In our optimal control framework, the receding horizon methodology comes under the Euler-Lagrange heading (Figure 4.3). It is based on trajectory optimizations, which only provide local open-loop information. It is the repeated solution of these optimizations, namely the receding horizon approach, that results in the desired state feedback solution.

Chapter 5 A Receding Horizon Extension of Pointwise Min-Norm Controllers

5.1 Introduction

In the previous two chapters, two popular approaches to the nonlinear optimal control problem were presented. Control Lyapunov function (CLF) based methodologies were considered first, where in particular they were discussed in relation to the Hamilton-Jacobi-Bellman (HJB) optimization equation. A variation of Sontag's famous CLF formula was highlighted as resulting from a special choice of parameters in the pointwise min-norm formulation and for possessing special optimality properties and interpretations in the context of the HJB partial differential equation. But as was clearly demonstrated in the example of Chapter 3, the performance of these controllers can be quite sensitive to the shape of the CLF and may result in poor performance when the CLF does not resemble the value function.

In stark contrast to the global and stability oriented philosophy which is the cornerstone of CLF techniques, the receding horizon methodology, which was reviewed in Chapter 4, aims for optimal performance through on-line computation. It proceeds by repeatedly solving finite horizon open-loop control problems emanating from the current state and applying the initial control actions until the next state measurement is available. This approach is more analogous to Euler-Lagrange based techniques for optimal control, which apply to finite horizon problems for a specified initial condition and result in open-loop control trajectories. While this approach aims for performance, guarantees on fundamental properties such as stability have generally been lacking or difficult to obtain. Moreover, successful implementation for nonlinear

systems can be troublesome due to the requirement of solving a generally non-convex optimization at each time step.

Based on their underlying connection with the optimal control problem, in this chapter we show that both control Lyapunov function based methods and receding horizon control can be cast in a single unifying framework where the advantages of both can be exploited. The strong stability properties of CLFs can be carried into a receding horizon scheme without sacrificing the excellent performance advantages of receding horizon control. With this flexible new approach, computation can be used to its fullest to approach optimality while stability is guaranteed by the presence of the CLF. This approach in essence combines and unites the best properties of CLFs and receding horizon control.

We begin by connecting the approaches reviewed in Chapters 3 and 4 by providing a unified framework in which to view them. From this common vantage point, we are able to introduce a new RHC+CLF scheme which represents a receding horizon extension of the pointwise min-norm controllers of Chapter 3. It is shown to possess various theoretical and implementation properties, including a special choice of parameters that corresponds to Sontag's formula. Finally, this approach is tested on our oscillator example of previous chapters.

5.2 Limits of receding horizon control

In Chapter 3 and 4, the philosophical underpinnings of two approaches (CLFs and RHC) were shown to lie in the two classical approaches (HJB and Euler-Lagrange) to the optimal control problem. A deeper look at the actual form of the underlying optimization involved in the following three schemes; optimal control, pointwise min-norm, and receding horizon; leads to an even more striking connection. In this section we develop a heuristic framework for viewing both optimal control (2.2) and pointwise min-norm control (3.9)–(3.10) as limiting cases of receding horizon control.

Our starting point will be to consider the standard open-loop optimization that is solved on-line at every time instance in receding horizon control, but without the

terminal weight $\varphi(\cdot)$

$$\int_t^{t+T} (q(x) + u^T u) d\tau. \quad (5.1)$$

First, we make the trivial observation that as the horizon T tends to infinity, the objective in the optimal control problem (2.2) is recovered,

$$\int_t^\infty (q(x) + u^T u) d\tau. \quad (5.2)$$

At the other extreme, consider what happens as the horizon T tends to zero. First, note that for any T an equivalent objective function is given by

$$\frac{1}{T} \int_t^{t+T} (q(x) + u^T u) d\tau \quad (5.3)$$

since dividing by T has no effect on the optimizing u . Now, letting $T \rightarrow 0$ yields

$$q(x(t)) + u^T(t)u(t). \quad (5.4)$$

Since $x(t)$ is known there is no need to include the term $q(x(t))$, leaving

$$u^T(t)u(t)$$

which is recognized as the objective function used in the pointwise min-norm formulation (3.9).

Hence, this indicates that we may heuristically view the pointwise min-norm problem as a receding horizon problem with a horizon length of zero. These considerations suggest the following interpretation: optimal control and pointwise min-norm formulations should represent extreme cases of a properly conceived receding horizon scheme. This is pictured in Figure 5.1.

Ideally, we would hope to incorporate the best properties of each approach into a single scheme parameterized by horizon length. These properties should include:

1. Stability for any horizon T .

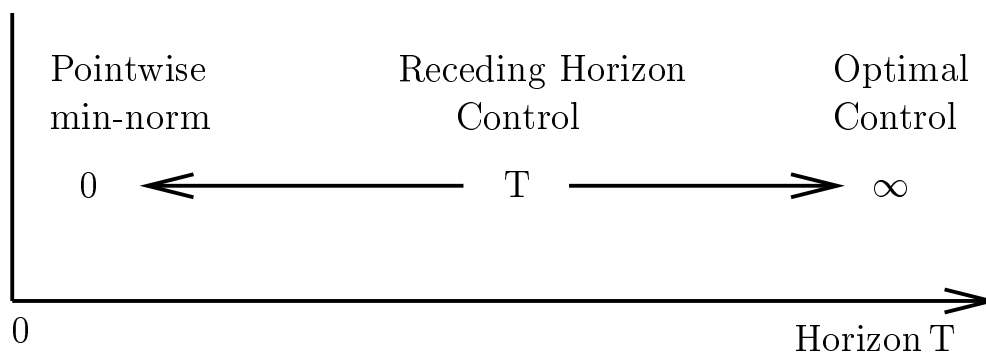


Figure 5.1: Unified framework.

2. Pointwise min-norm controllers for $T = 0$.
3. Optimality for $T = \infty$.

Additionally, there should exist an extension of Sontag's formula that will recover the optimal controller if the level curves of the CLF correspond to those of the value function, regardless of the horizon length T . With these goals in mind, we present a new class of control Lyapunov function based receding horizon control schemes.

5.3 A receding horizon generalization of pointwise min-norm controllers

In this section a new class of control schemes is introduced that retain the global stability properties of control Lyapunov function methods while taking advantage of the on-line optimization techniques employed in receding horizon control. In essence it represents a natural extension of the CLF based pointwise min-norm concept to a receding horizon methodology, including an appropriate interpretation as a conceptual blend of HJB and Euler-Lagrange philosophies. This interaction of approaches is found to inherit not only the theoretical advantages of each methodology, but unexpectedly results in practical and advantageous implementation properties.

Let V be a CLF and let u_σ and x_σ denote the control and state trajectories

obtained by solving the pointwise min-norm problem with parameter $\sigma(x)$ (cf. (3.9)–(3.10)). Consider the following receding horizon objective:

(RHC+CLF)

$$\text{minimize} \quad \int_t^{t+T} (q(x) + u^T u) d\tau \quad (5.5)$$

$$\text{subject to} \quad \dot{x} = f(x) + g(x)u$$

$$\frac{\partial V}{\partial x}[f + gu(t)] \leq -\epsilon\sigma(x(t)) \quad (5.6)$$

$$V(x(t+T)) \leq V(x_\sigma(t+T)) \quad (5.7)$$

with $0 < \epsilon \leq 1$ (preferably ϵ is small).

The preceding scheme is best interpreted in the following manner. It is a standard receding horizon formulation with two CLF constraints. The first constraint (5.6) is a direct stability constraint in the spirit of that which appears in the pointwise min-norm formulation (3.10). The parameter ϵ is merely used to relax this constraint as compared to its counterpart in the pointwise min-norm formulation. Note that this constraint need only apply to the implemented control actions, which in the ideal case of the optimization being continuously resolved is only the *initial* control action. In essence, this constraint requires V to be a Lyapunov function for the closed loop system. (In the actual implementation of receding horizon control, the constraint should apply at least over the entire sampling time interval in which the optimizing control solution is implemented. In fact, situations may exist where it is reasonable to apply a constraint of this form over the entire horizon T , in which case ϵ can even be chosen as a function of time on $[t, t+T]$. This is discussed further in Section 5.4.)

In contrast to the first constraint which is a direct stability constraint, the second constraint (5.7) is oriented toward performance and replaces the terminal weight used in the standard receding horizon formulation. As will be seen later, when the pointwise min-norm problem corresponding to Sontag's formula is used (i.e., $\sigma = \sigma_s$ (eqn. 3.11)), this constraint preserves the property that when the level curves of the CLF (V) correspond to those of the value function (V^*), the optimal controller

is recovered. It is obtained by first simulating the control from the solution to the pointwise min-norm problem for time T , which results in a state trajectory that ends at $x_\sigma(t+T)$, then evaluating the CLF at this point ($V(x_\sigma(t+T))$). The constraint then requires that all other potential sequences reach a final state that obtains a smaller value of V . A nice interpretation is in terms of level curves. The constraint (5.7) requires that the final state of all potential sequences lie inside the level curve of V that passes through $x_\sigma(t+T)$ (see Figure 5.2). The constraint (3.10) in the pointwise min-norm formulation can be thought of as a differential version of this constraint.




Figure 5.2: Performance constraint (5.7).

This combination of control Lyapunov functions and receding horizon control yields a number of theoretically appealing properties, as listed below:

1. *Stability is guaranteed for any horizon T .*

The constraint (5.6) requires that V is a Lyapunov function for the receding horizon controlled system and hence guarantees stability.

2. In the limit as the horizon goes to zero ($T \rightarrow 0$), the pointwise min-norm optimization problem is recovered.

It was already shown that as $T \rightarrow 0$, the limiting performance objective is given by $u^T u$. We only need to show that the constraints reduce to the pointwise min-norm constraint (3.10).

Subtracting $V(x(t))$ from both sides of the performance constraint (5.7) gives

$$V(x(t+T)) - V(x(t)) \leq V(x_\sigma(t+T)) - V(x(t)).$$

Dividing by T and taking the limit as $T \rightarrow 0$ yields

$$\begin{aligned} \frac{\partial V}{\partial x}[f(x) + g(x)u(t)] &\leq \frac{\partial V}{\partial x}[f(x) + g(x)u_\sigma(x(t))] \\ &\leq -\sigma(x(t)). \end{aligned}$$

In fact, it is simple to see that the constraints

$$\frac{\partial V}{\partial x}[f(x) + g(x)u(t)] \leq \frac{\partial V}{\partial x}[f(x) + g(x)u_\sigma(x(t))]$$

and

$$\frac{\partial V}{\partial x}[f(x) + g(x)u(t)] \leq -\sigma(x(t))$$

produce the same control actions in the pointwise min-norm formulation.

Since we require that $\epsilon \leq 1$ in the stability constraint (5.6) the above constraint supersedes the stability constraint in the limit. Hence, the receding horizon optimization problem is reduced to:

$$\begin{aligned} \text{minimize} \quad & u^T(t)u(t) \\ \text{s.t.} \quad & \frac{\partial V}{\partial x}[f(x) + g(x)u(t)] \leq -\sigma(x). \end{aligned}$$

3. *If V is a Lyapunov function for the closed loop system under the optimal control, u^* , which always satisfies the constraint (5.6), then an infinite horizon length will always recover the optimal controller.*

With an infinite horizon ($T = \infty$), the objective becomes an infinite horizon objective

$$\int_t^\infty (q(x) + u^T u) d\tau.$$

With no constraints the solution to this is the optimal control u^* . We only need to show that under the assumptions, the optimal control is feasible. By assumption, it is feasible for the first constraint (5.6). For an infinite horizon, the performance constraint (5.7) becomes that the state must approach zero as t approaches infinity. Clearly this is satisfied under the optimal control. Hence, the optimal unconstrained control is a feasible solution and therefore optimal.

While we have been rather informal about our justification of the above properties, in the appendix a rigorous treatment is given under stringent technical conditions. The argument above that the optimization problem reduces to the optimal infinite horizon problem or the pointwise min-norm formulation as the horizon tends to infinity or zero is strengthened to show that the receding horizon control action obtained from the RHC+CLF problem will converge to the optimal control action u^* or the pointwise min-norm controller u_σ as the horizon extends to infinity or shrinks to zero. Details are contained in the appendix.

Additionally, for the parameter choice $\sigma(x) = \sigma_s(x)$ corresponding to Sontag's formula in the pointwise min-norm problem (see eqn. 3.11), the optimality property of Sontag's formula is preserved.

Theorem 5.3.1 *Let $\sigma(x) = \sigma_s(x)$ (cf. eqn. 3.11). If V has the same level curves as the value function V^* , then the optimal control is recovered for any horizon length.*

Proof: Assume that V has the same level curves as the value function V^* . In this case, Sontag's formula results in an optimal state trajectory x_{σ_s} and control action

u_{σ_s} . Let us assume that x_{σ_s} and u_{σ_s} does not solve the optimization problem (5.5–5.7). Hence, there exist trajectories x and u such that

$$\int_t^{t+T} (q(x) + u^T u) d\tau < \int_t^{t+T} (q(x_{\sigma_s}) + u_{\sigma_s}^T u_{\sigma_s}) d\tau. \quad (5.8)$$

Furthermore, since x and u satisfy the constraint (5.7), we have that

$$V(x(t+T)) \leq V(x_{\sigma_s}(t+T))$$

or using the fact that V has the same level curves as V^* ,

$$V^*(x(t+T)) \leq V^*(x_{\sigma_s}(t+T)). \quad (5.9)$$

Combining (5.8) and (5.9) and the fact that Sontag's formula is optimal gives

$$\begin{aligned} \int_t^{t+T} (q(x) + u^T u) d\tau + V^*(x(t+T)) &< \int_t^{t+T} (q(x_{\sigma_s}) + u_{\sigma_s}^T u_{\sigma_s}) d\tau + V^*(x_{\sigma_s}(t+T)) \\ &= V^*(x(t)) \end{aligned}$$

which is a contradiction, since V^* is the minimum cost. ■

Before addressing some of the implementation issues faced in this new RHC+CLF scheme, let us summarize the key ideas behind this approach. From a practical viewpoint, it involves a mix of the guaranteed stability properties of control Lyapunov functions combined with the on-line optimization and performance properties of receding horizon control. Conceptually, it blends the philosophies behind the Hamilton-Jacobi-Bellman and Euler-Lagrange approaches to the nonlinear optimal control problem. The control Lyapunov function represents the best approximation to the value function in the HJB approach. The on-line optimization then proceeds in an Euler-Lagrange fashion, optimizing over trajectories emanating from the current state, improving the solution by using as much computation time as is available.

5.4 Implementation issues

In addition to the theoretical properties of the previous section, the RHC+CLF scheme possesses a number of desirable implementation properties.

1. *An initial feasible trajectory for the optimization is provided by the solution to the pointwise min-norm problem.*

For the performance constraint (5.7), it is necessary to simulate the solution to the pointwise min-norm problem over the horizon T to obtain $x_\sigma(t+T)$. Additionally, the control and state trajectory from this pointwise min-norm problem provide an initial feasible trajectory from which to begin the optimization.

2. *The optimization may be preempted without loss of stability.*

Since the constraint (5.6) ensures that V will be a Lyapunov function for the closed loop system, any control that satisfies this constraint will be stabilizing. In particular, if the optimization cannot be completed one may implement the current best solution and proceed without any loss of stability. Hence, there is no requirement of a global optimum to the non-convex optimization (5.5)–(5.7) to guarantee stability.

3. *The horizon may be varied on-line without loss of stability.*

This is again due to the stability constraint (5.6). Since stability is guaranteed by the constraint (5.6) and is independent of the objective function, it is clear that the horizon may be varied on-line without jeopardizing stability. In particular, one could imagine a situation where the amount of time available for on-line computation is not constant. When more time is available, the horizon can be extended on-line to take advantage of this. On the other hand, if at various times no on-line computation is available, the horizon can be drawn in to zero where the control is given by the pointwise min-norm solution. In essence, one may use the available computation time to its fullest by adjusting the horizon on-line, all without any concern of losing stability.

In practice, receding horizon control is typically not implemented in continuous time but rather at discrete sampling times. Over each sampling interval the receding horizon control problem is solved and the optimizing control solution is applied until a new state update is received at the next sampling time, in which the process repeats. To guarantee stability, the constraint (5.6) should apply over the entire sampling interval so that all control actions that are implemented conform to V being a Lyapunov function. There may even be cases in which it is convenient to impose the constraint (5.6) over the entire horizon $[t, t+T]$. For example, this situation may occur when the horizon length and/or sampling interval is allowed to vary dramatically, and hence cannot be determined a priori. In any case, the parameter ϵ need not be a fixed constant, but rather can be a function of time $\epsilon(\tau)$, $\tau \in [t, t+T]$ satisfying

1. $\epsilon(\tau) \leq 1$ for all $\tau \in [t, t+T]$
2. $\epsilon(\tau) > 0$ for all $\tau \in [t, t+T_s]$

where T_s is the sampling time. Beyond this, $\epsilon(\tau)$ is a free design parameter.

In the next section we demonstrate the RHC+CLF approach on our familiar two dimensional oscillator example.

5.5 Example

Once again we return to the two dimensional nonlinear oscillator used in Chapters 3 and 4, showing that now armed with both the stability properties of CLFs and the performance advantages of on-line receding horizon computation, the RHC+CLF approach provides an effective solution. Recall that the system dynamics are given by

$$\begin{cases} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= -x_1 \left(\frac{\pi}{2} + \arctan(5x_1) \right) - \frac{5x_1^2}{2(1+25x_1^2)} + 4x_2 + 3u \end{cases}$$

with performance index

$$\int_0^{\infty} (x_2^2 + u^2) dt.$$

For this problem, the value function is given by

$$V^* = x_1^2 \left(\frac{\pi}{2} + \arctan(5x_1) \right) + x_2^2$$

which results in the optimal control action

$$u^* = -3x_2$$

The same control Lyapunov function used in the example in Chapter 3 is also used here,

$$V = \frac{\pi}{2}x_1^2 + x_2^2.$$

Again, it is emphasized that the level curves of this CLF are far from those of the value function (see Fig. 3.1). As was explained in Chapter 3, this accounts for the poor performance of Sontag's formula, which accumulates a cost of over 250 from the initial condition $[3, -2]$.

Building upon this Sontag's formula approach (i.e., using σ_s in (3.11)), a horizon was introduced in accordance with the newly developed RHC+CLF scheme (as described in Section 5.3). In our implementation we used discrete time intervals of 0.1 over which the control inputs were held constant. Furthermore, the stability constraint (5.6) was applied over this entire 0.1 intersample time using $\epsilon = 0.01$. As shown in Figure 5.3, the erratic behavior demonstrated by the receding horizon controllers in Chapter 4 has been tamed and drastically improved performance is achieved for each of the tested horizons. Table 5.1 summarizes the costs accumulated for each of the horizons $T = 0.2, 0.3, 0.5$ and 1.0 . A surprising result is that even a short horizon dramatically reduces the cost over that of Sontag's formula alone, demonstrating the power of the combination of CLF techniques with even a minimal amount of on-line computation.

The fact that the cost does not decrease monotonically as a function of horizon length is attributable to the erratic behavior that receding horizon control by itself displays. It is interesting to observe that while alone both Sontag's formula and




Figure 5.3: Phase portrait of receding horizon controllers.

receding horizon control perform miserably, the proper combination of them results in consistent near optimal controllers.

5.6 Summary

The ideas behind CLF based pointwise min-norm controllers and receding horizon control were combined to create a new class of control schemes. These new results were facilitated by the development of a framework within which both optimal and pointwise min-norm controllers served as limiting cases of receding horizon control. This led us to propose a natural extension of the pointwise min-norm formulation to allow for on-line computation in a receding horizon implementation. In particular, this even provided a receding horizon “extension” of Sontag’s formula, and resulted in numerous theoretical and implementation advantages over present CLF and receding

Table 5.1

Controller	Cost
Sontag	258
RHC+CLF ($T = 0.2$)	35.3
RHC+CLF ($T = 0.3$)	37.9
RHC+CLF ($T = 0.5$)	33.6
RHC+CLF ($T = 1.0$)	36.8
Optimal	31.7

Table 5.1: Summary of controller costs from initial condition $[3, -2]$.

horizon methodologies. As summarized in our picture of optimal control (Figure 5.4), the RHC+CLF schemes complete the optimal control framework by combining both of the classical viewpoints and their offspring into a single unified approach.

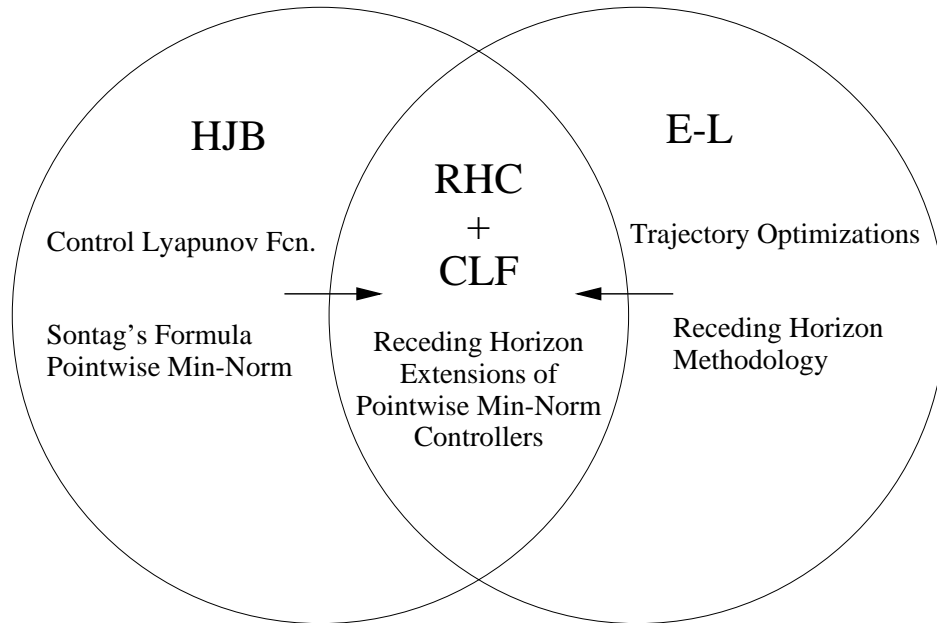


Figure 5.4: The RHC+CLF scheme within the optimal control picture.

Although we illustrated some of the advantages of receding horizon extensions on the common example used in both Chapters 3 and 4, a more thorough investigation is undertaken next. In the following chapter we work through the details of a more concrete and realistic example: the control of a planar ducted fan model. That example

more accurately illustrates the steps required to implement the design methodology presented in this chapter.

Appendix

In this appendix we show that the control actions from the RHC+CLF scheme converge to those of the pointwise min-norm controller and the optimal infinite horizon controller as the horizon is brought to zero and infinity, respectively. But first, we begin by establishing some required notation and assumptions.

Let $|\cdot|$ and $|\cdot|_\infty$ denote the standard Euclidean and infinity norms on \mathbb{R}^N . We will assume that both the CLF V and the value function V^* are \mathbf{C}^1 and proper. As before, $x_\sigma(\cdot)$ and $u_\sigma(\cdot)$ will denote the state and control corresponding to the pointwise min-norm problem, and $x^*(\cdot)$ and $u^*(\cdot)$ will represent the state and control of the optimal infinite horizon controller. For any optimization with a non-zero horizon, the positive semi-definite cost parameter $q(\cdot)$ will be at least \mathbf{C}^0 , the initial condition will be denoted $x(0)$, and the optimization will be taken over all piecewise \mathbf{C}^0 functions with the assumption that the infimum is achieved and is unique. The notation \hat{V}_T will be used to denote the optimal cost of the RHC+CLF problem with horizon T . The corresponding optimizing state and control trajectories will be denoted by $\hat{x}_T(\cdot)$ and $\hat{u}_T(\cdot)$. As before the dynamics are

$$\dot{x} = f(x) + g(x)u = f(x) + \sum_{i=1}^m g_i(x)u_i$$

with $x \in \mathbb{R}^n$ and $u = [u_1, u_2, \dots, u_m]^T \in \mathbb{R}^m$. We will assume that $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is globally Lipschitz with Lipschitz constant K_f and each $g_i : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is globally Lipschitz with common Lipschitz constant K_g .

For the pointwise min-norm problem (3.9) we will assume the parameter $\sigma(x)$ is continuous, locally Lipschitz, positive definite and satisfies

$$x \neq 0, \quad \frac{\partial V}{\partial x} g(x) = 0 \quad \Rightarrow \quad \frac{\partial V}{\partial x} f(x) < -\sigma(x).$$

Under these conditions, the pointwise min-norm controller $u_\sigma(x)$ is also continuous and locally Lipschitz everywhere except possibly at the origin [FK95]. Hence, for

small enough t , it satisfies

$$|u_\sigma(x(0)) - u_\sigma(x(t))| \leq Kt$$

for some K .

To prove connections between the pointwise min-norm problem, and the RHC+CLF problem, we will require a similar assumption on the control trajectories from the RHC+CLF problems, stated as follows:

(A1) Given a fixed initial condition $x(0)$, for all horizons T sufficiently small $\hat{u}_T(t)$ is \mathbf{C}^0 and satisfies the following Lipschitz condition

$$|\hat{u}_T(0) - \hat{u}_T(t)| \leq Kt, \quad \forall t \in [0, T] \quad (5.10)$$

for some K .

The assumption A1 also provides some regularity on the variation of the state trajectories $\hat{x}_T(\cdot)$. To see this consider the state trajectory $\hat{x}_T(\cdot)$ from the RHC+CLF problem beginning at state $x(0)$ and assume A1, then for small enough T :

$$\begin{aligned} |\hat{x}_T(t) - x(0)| &= \left| \int_0^t \left(f(\hat{x}_T(s)) + g(\hat{x}_T(s))\hat{u}_T(s) \right) ds \right| \\ &\leq \int_0^t \left(|f(\hat{x}_T(s))| + |g(\hat{x}_T(s))\hat{u}_T(s)| \right) ds \\ &\leq \int_0^t |f(\hat{x}_T(s))| ds \\ &\quad + \int_0^t |[g(\hat{x}_T(s)) - g(x(0)) + g(x(0))]\hat{u}_T(s)| ds \\ &\leq \int_0^t \left(|f(\hat{x}_T(s)) - f(x(0))| + |f(x(0))| \right) ds \\ &\quad + \int_0^t \left(\sum_{i=1}^m |[g_i(\hat{x}_T(s)) - g_i(x(0))]\hat{u}_{iT}(s)| \right) ds \\ &\quad + \int_0^t \left(\sum_{i=1}^m |g_i(x(0))\hat{u}_{iT}(s)| \right) ds \\ &\leq \int_0^t \left(|f(\hat{x}_T(s)) - f(x(0))| + |f(x(0))| \right) ds \end{aligned}$$

$$\begin{aligned}
& + \int_0^t \left(\sum_{i=1}^m |g_i(\hat{x}_T(s)) - g_i(x(0))| |\hat{u}_{iT}(s)| \right) ds \\
& + \int_0^t \left(\sum_{i=1}^m |g_i(x(0))| |\hat{u}_{iT}(s)| \right) ds \\
\leq & \int_0^t \left(|f(\hat{x}_T(s)) - f(x(0))| + |f(x(0))| \right) ds \\
& + \int_0^t \left(\sum_{i=1}^m |g_i(\hat{x}_T(s)) - g_i(x(0))| (|\hat{u}_T(0)| + Ks) \right) ds \\
& + \int_0^t \left(\sum_{i=1}^m |g_i(x(0))| (|\hat{u}_T(0)| + Ks) \right) ds \\
\leq & \int_0^t \left(K_f |\hat{x}_T(s) - x(0)| + |f(x(0))| \right) ds \\
& + \int_0^t \left(\sum_{i=1}^m K_g |\hat{x}_T(s) - x(0)| (|\hat{u}_T(0)| + Ks) \right) ds \\
& + \int_0^t \left(\sum_{i=1}^m |g_i(x(0))| (|\hat{u}_T(0)| + Ks) \right) ds \\
\leq & \int_0^t \left((K_f + mK_g (|\hat{u}_T(0)| + Ks)) |\hat{x}_T(s) - x(0)| \right) ds \\
& + \int_0^t \left(|f(x(0))| + (|\hat{u}_T(0)| + Ks) \sum_{i=1}^m |g_i(x(0))| \right) ds \\
\leq & \int_0^t \left((K_f + mK_g (|\hat{u}_T(0)| + Ks)) |\hat{x}_T(s) - x(0)| \right) ds \\
& + \left(|f(x(0))| + |\hat{u}_T(0)| \sum_{i=1}^m |g_i(x(0))| \right) t + \left(K \sum_{i=1}^m |g_i(x(0))| \right) \frac{t^2}{2}.
\end{aligned}$$

where we have used assumption A1 and that f and g are Lipschitz. If we let

$$\lambda(t) = \left(|f(x(0))| + |\hat{u}_T(0)| \sum_{i=1}^m |g_i(x(0))| \right) t + \left(K \sum_{i=1}^m |g_i(x(0))| \right) \frac{t^2}{2}$$

and

$$\mu(s) = (K_f + mK_g (|\hat{u}_T(0)| + Ks))$$

then we have that

$$|\hat{x}_T(t) - x(0)| \leq \int_0^t \mu(s) |\hat{x}_T(s) - x(0)| ds + \lambda(t).$$

An application of the Gronwall-Bellman Lemma [Kha92] gives,

$$|\hat{x}_T(t) - x(0)| \leq \lambda(t) + \int_0^t \lambda(s) \mu(s) \exp\left[\int_s^t \mu(\tau) d\tau\right] ds. \quad (5.11)$$

This provides an explicit bound for the amount by which \hat{x}_T is allowed to vary in time t . Finally, we will implicitly assume that all limits, when stated, exist.

A further justification for some of the above assumptions can be made as follows. Optimal control problems are typically solved by representing the control trajectory over a finite dimensional spline space. This involves the choice of a *knot* sequence (i.e., a nondecreasing sequence (t_i)) which the splines are defined with respect to. Most splines will allow discontinuities only on the knot sequence and can be chosen to be smooth in between. The optimization is carried out by using the coefficient of each spline basis function as a decision variable. If these coefficients are restricted to lie in some compact set, then assumption A1 will necessarily be satisfied. These considerations help to make the continuity and Lipschitz assumptions a bit more natural.

The first theorem shows that the control actions obtained from the RHC+CLF problem converge to the pointwise min-norm solution as the horizon is brought to zero.

Theorem 5.6.1 *Denote the initial condition for the RHC+CLF optimization problems by $x(0)$, and assume that $\lim_{T \rightarrow 0} \hat{u}_T(0) = \hat{u}_0$. Under the assumptions stated above, $\hat{u}_0 = u_\sigma(x(0))$ where $u_\sigma(x(0))$ solves the corresponding pointwise min-norm problem.*

Proof: First we show that \hat{u}_0 is feasible for the zero horizon problem (i.e., the pointwise min-norm problem with parameter $\sigma(x)$ as in (3.10)). For this purpose, it is sufficient to show that

$$\frac{\partial V}{\partial x}[f + g\hat{u}_0] \leq \frac{\partial V}{\partial x}[f + gu_\sigma(x(0))]. \quad (5.12)$$

Since it is known that each \hat{u}_T satisfies (5.7),

$$V(\hat{x}_T(T)) \leq V(x_\sigma(T)),$$

subtracting $V(x(0))$ and dividing by T gives:

$$\frac{1}{T} \left[V(\hat{x}_T(T)) - V(x(0)) \right] \leq \frac{1}{T} \left[V(x_\sigma(T)) - V(x(0)) \right].$$

By the definition of a derivative and the chain rule, taking the limit as $T \rightarrow 0$ gives (5.12). Hence \hat{u}_0 is feasible for the zero-horizon (pointwise min-norm) problem.

Now assume that $\hat{u}_0 \neq u_\sigma(x(0))$. Since \hat{u}_0 is feasible, we must have that $\hat{u}_0^T \hat{u}_0 > u_\sigma^T(x(0))u_\sigma(x(0))$ (otherwise this contradicts that $u_\sigma(x(0))$ is the unique solution to the zero horizon (pointwise min-norm) problem [FK95]). This means that for some $\epsilon > 0$ we can find a horizon T' small enough so that

$$q(x(0)) + u_\sigma^T(x(0))u_\sigma(x(0)) + \epsilon \leq q(x(0)) + \hat{u}_{T'}^T(0)\hat{u}_{T'}(0).$$

But, by the Lipschitz condition (5.10) on $\hat{u}_T(\cdot)$ and the bound (5.11) on the rate of variation of the state trajectory $\hat{x}_T(\cdot)$ a similar inequality must hold over a small enough horizon T' . (Note that equation (5.11) actually depends on $\hat{u}_T(0)$ through $\lambda(t)$ and $\mu(t)$. Furthermore, $\hat{u}_T(0)$ is different for each horizon T . Nevertheless, we know that $\hat{u}_T(0)$ converges to \hat{u}_0 and hence can still guarantee a bound on the rate of variation of \hat{x}_T which is independent of the horizon T .) Hence, there exists a T' sufficiently small so that,

$$q(x_\sigma(t)) + u_\sigma^T(x_\sigma(t))u_\sigma(x_\sigma(t)) < q(\hat{x}_{T'}(t)) + \hat{u}_{T'}^T(t)\hat{u}_{T'}(t), \quad t \in [0, T'].$$

Integration from zero to T' completes the contradiction since $\hat{u}_{T'}(t)$ was assumed optimal for this horizon. Hence $\hat{u}_0 = u_\sigma$. ■

Before exploring the solution to the RHC+CLF problem as the horizon is increased to ∞ , we remind the reader of the following definition.

Definition 5.6.1 *A function $W : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is said to belong to class \mathcal{K}_∞ if:*

1. *it is continuous.*
2. $W(0) = 0$.
3. *it is strictly increasing.*
4. $W(s) \rightarrow \infty$ when $s \rightarrow \infty$.

We will require the nonlinear system to satisfy an additional condition. Using notation from [KG88], we refer to the following as *property C*:

Definition 5.6.2 *The system $\dot{x} = f(x) + g(x)u$ is said to satisfy property C if there exists a time T^c , and a \mathcal{K}_∞ function W_c such that for any $x_0 \in \mathbb{R}^n$, there exist continuous state and control trajectories $(x^c(t), u^c(t))$ such that $x^c(0) = x_0$ and $x^c(T^c) = 0$ with*

$$\int_0^{T^c} |(x^c(t), u^c(t))| \leq W_c(|x_0|).$$

We will say that the system $\dot{x} = f(x) + g(x)u$ locally has property C if property C holds for some neighborhood of the origin. Note that for $q(\cdot)$ locally Lipschitz, local satisfaction of property C implies that

$$\int_0^{T^c} q(x^c(t) + |u^c(t)|)^2 \leq W_c'(|x_0|) \tag{5.13}$$

is also satisfied locally for some \mathcal{K}_∞ function W_c' .

Remark: Property C can be thought of as a weak controllability condition. Consider a linear system: $\dot{x} = Ax + Bu$ with (A, B) controllable. Then from any initial condition the state can be brought to the origin using the minimum energy control. It can be shown that this will satisfy property C [KG88].

Theorem 5.6.2 *Assume that $q(x)$ is continuous, locally Lipschitz and that $q(x) \geq \alpha(|x|)$ where α is \mathcal{K}_∞ . Additionally, assume that the optimal infinite horizon control u^* satisfies the CLF stability constraint (5.6). Furthermore, assume that the nonlinear system $\dot{x} = f(x) + g(x)u$ locally satisfies property C. Then over any compact set \mathcal{S}*

$$\hat{V}_T(x) \xrightarrow{T \rightarrow \infty} V^*(x) \text{ uniformly.}$$

Furthermore, if there exists an interval $[0, \beta]$ on which $\hat{u}_T(\tau)$ is continuous for each T and $\hat{u}_T(\tau) \rightarrow \hat{u}_\infty(\tau)$ uniformly, then $\hat{u}_\infty(\tau) = u^(\tau)$ for $\tau \in [0, \beta]$.*

Proof: To establish notation, recall that V^* is the value function corresponding to the optimal cost of the unconstrained infinite horizon optimal control problem with state and control trajectories x^* and u^* . Let V_T^* denote the cost of applying the infinite horizon optimal control action u^* , but only over a horizon of length T . Finally recall that \hat{V}_T is the optimal cost of the RHC+CLF problem with horizon T and state and control trajectories \hat{x}_T and \hat{u}_T .

Choose $\epsilon > 0$ and consider the set $\mathcal{N} = \{x : W'_c(|x|) \leq \epsilon\}$ (with $W'_c(\cdot)$ as in (5.13)) which contains a neighborhood of the origin. Furthermore, let $\hat{q} > 0$ be the infimum of $q(x)$ outside of \mathcal{N} . Now let \mathcal{S} be any compact set and denote the maximum of V^* over \mathcal{S} by v . Then for $T > T^* = v/\hat{q}$, there exists a $t \in [0, T]$ such that the state $x^*(t) \in \mathcal{N}$. That is, from any initial condition in \mathcal{S} , after T^* seconds it is guaranteed that the optimal trajectory $x^*(\cdot)$ has intersected \mathcal{N} . This is because if there does not exist a $t \in [0, T]$ with $x^*(t) \in \mathcal{N}$ then $q(x^*(t)) \geq \hat{q}$ for all $t \in [0, T]$ and hence

$$V^*(x) > V_T^*(x) = \int_0^T \left(q(x^*(t)) + u^{*T}(t)u^*(t) \right) dt \geq T\hat{q} > v$$

which is a contradiction.

Now, for the RHC+CLF problem with horizon $T > T^* + T^c$, consider the following feasible control actions. Apply u^* (this is feasible by assumption) until the state enters \mathcal{N} , then use u^c (cf., Definition 5.6.2) to drive the state to the origin. If $T^{\mathcal{N}} \leq T^*$

denotes the first time that $x^*(\cdot)$ enters \mathcal{N} , then the cost of this trajectory is less than or equal to $V_{T^*}^* + W'_c(|x^*(T^*)|)$ which is less than or equal to $V_T^* + \epsilon$. Furthermore, this trajectory ends at the origin, and hence also provides an upper bound for the optimal infinite horizon cost, V^* . From this we can assert the following: for every horizon $T > T^* + T^c$, we have

$$V_T^* + \epsilon \geq V^* \geq V_T^*$$

and

$$V_T^* + \epsilon \geq \hat{V}_T \geq V_T^*$$

which implies

$$|V^* - \hat{V}_T| \leq \epsilon$$

proving the first part of the theorem.

The second portion of the theorem follows in three steps:

1.) \hat{x}_∞ exists and is unique and continuous on $[0, \beta]$.

By assumption there exists an interval $[0, \beta]$ where $\hat{u}_T(\tau)$ is continuous and $\hat{u}_T(\tau) \rightarrow \hat{u}_\infty(\tau)$ uniformly. Hence, \hat{u}_∞ is continuous on $[0, \beta]$. Since $[0, \beta]$ is compact, $\hat{u}_\infty(t)$ is bounded. Let $\max_{t \in [0, \beta]} |\hat{u}_\infty(t)|_\infty = M$.

Now let \hat{x}_∞ be the state trajectory corresponding to the input \hat{u}_∞ over the interval $[0, \beta]$. If we define $\hat{f}(x, t) = f(x) + g(x)\hat{u}_\infty(t)$ on $t \in [0, \beta]$, then $\hat{f}(x, t)$ is Lipschitz since

$$\begin{aligned} |\hat{f}(x, t) - \hat{f}(y, t)| &= |f(x) - f(y) + [g(x) - g(y)]\hat{u}_\infty(t)| \\ &\leq |f(x) - f(y)| + |[g(x) - g(y)]\hat{u}_\infty(t)| \\ &\leq |f(x) - f(y)| + \sum_{i=1}^m |[g_i(x) - g_i(y)]\hat{u}_{i\infty}(t)| \\ &\leq |f(x) - f(y)| + \sum_{i=1}^m M|[g_i(x) - g_i(y)]| \end{aligned}$$

$$\begin{aligned}
&\leq |f(x) - f(y)| + \sum_{i=1}^m MK_g |x - y| \\
&\leq K_f |x - y| + mMK_g |x - y| \\
&= (K_f + mMK_g) |x - y|
\end{aligned}$$

where we have used that f and g are Lipschitz with Lipschitz constants K_f and K_g , and that $\hat{u}_\infty(\cdot)$ is bounded in infinity norm by M . Therefore, by standard existence and uniqueness theorems for differential equations (see [Kha92], pg. 81), the state trajectory \hat{x}_∞ exists and is unique and continuous on $[0, \beta]$.

2.) \hat{x}_T converges to \hat{x}_∞ on $[0, \beta]$.

Let us show that \hat{x}_T converges pointwise to \hat{x}_∞ on $[0, \beta]$. This is basically an exercise in using Lipschitz constants, and an application of the Gronwall-Bellman lemma ([Kha92], pg. 68).

$$\begin{aligned}
|\hat{x}_\infty(t) - \hat{x}_T(t)| &= \left| \int_0^t \left(f(\hat{x}_\infty(s)) - f(\hat{x}_T(s)) \right) ds \right. \\
&\quad \left. + \int_0^t \left(g(\hat{x}_\infty(s))\hat{u}_\infty(s) - g(\hat{x}_T(s))\hat{u}_T(s) \right) ds \right| \\
&\leq \int_0^t \left(|f(\hat{x}_\infty(s)) - f(\hat{x}_T(s))| \right) ds \\
&\quad + \int_0^t \left(|g(\hat{x}_\infty(s))\hat{u}_\infty(s) - g(\hat{x}_T(s))\hat{u}_T(s)| \right) ds \\
&= \int_0^t \left(|f(\hat{x}_\infty(s)) - f(\hat{x}_T(s))| \right) ds \\
&\quad + \int_0^t \left(|g(\hat{x}_\infty(s))\hat{u}_\infty(s) - g(\hat{x}_\infty(s))\hat{u}_T(s)| \right) ds \\
&\quad + \int_0^t \left(|g(\hat{x}_\infty(s))\hat{u}_T(s) - g(\hat{x}_T(s))\hat{u}_T(s)| \right) ds \\
&\leq \int_0^t \left(K_f |\hat{x}_\infty(s) - \hat{x}_T(s)| \right) ds \\
&\quad + \int_0^t \left(\sum_{i=1}^m |g_i(\hat{x}_\infty(s))| |\hat{u}_{i\infty}(s) - \hat{u}_{iT}(s)| \right) ds \\
&\quad + \int_0^t \left(\sum_{i=1}^m |g_i(\hat{x}_\infty(s)) - g_i(\hat{x}_T(s))| |\hat{u}_{iT}(s)| \right) ds
\end{aligned}$$

$$\begin{aligned}
&\leq \int_0^t \left(K_f |\hat{x}_\infty(s) - \hat{x}_T(s)| \right) ds \\
&\quad + \int_0^t \left(\sum_{i=1}^m |g_i(\hat{x}_\infty(s))| |[\hat{u}_{i\infty}(s) - \hat{u}_{iT}(s)]| \right) ds \\
&\quad + \int_0^t \left(\sum_{i=1}^m |[g_i(\hat{x}_\infty(s)) - g_i(\hat{x}_T(s))]| |\hat{u}_{iT}(s)| \right) ds
\end{aligned}$$

Now note that each $|g_i(\hat{x}_\infty(s))|$ is bounded from above on $[0, \beta]$ since it is a continuous function over a compact set. Hence, choose an M_g such that $\max_{t \in [0, \beta]} |g_i(\hat{x}_\infty(t))| \leq M_g$ for $i = 1 \dots m$. Furthermore, by the fact that \hat{u}_T converges *uniformly* to \hat{u}_∞ , by choosing T large enough we can bound $\max_{t \in [0, \beta]} |\hat{u}_T(\cdot)|_\infty$ by $M+1$ (recall that $\max_{t \in [0, \beta]} |\hat{u}_\infty(\cdot)|_\infty = M$). Hence, returning to our bound

$$\begin{aligned}
|\hat{x}_\infty(t) - \hat{x}_T(t)| &\leq \int_0^t \left(K_f |\hat{x}_\infty(s) - \hat{x}_T(s)| \right) ds \\
&\quad + \int_0^t \left(\sum_{i=1}^m |g_i(\hat{x}_\infty(s))| |[\hat{u}_{i\infty}(s) - \hat{u}_{iT}(s)]| \right) ds \\
&\quad + \int_0^t \left(\sum_{i=1}^m |[g_i(\hat{x}_\infty(s)) - g_i(\hat{x}_T(s))]| |\hat{u}_{iT}(s)| \right) ds \\
&\leq \int_0^t \left(K_f |\hat{x}_\infty(s) - \hat{x}_T(s)| \right) ds \\
&\quad + \int_0^t \left(\sum_{i=1}^m M_g |\hat{u}_{i\infty}(s) - \hat{u}_{iT}(s)| \right) ds \\
&\quad + \int_0^t \left(\sum_{i=1}^m (M+1) K_g |\hat{x}_\infty(s) - \hat{x}_T(s)| \right) ds \\
&\leq \int_0^t \left((K_f + m(M+1)K_g) |\hat{x}_\infty(s) - \hat{x}_T(s)| \right) ds \\
&\quad + \int_0^t \left(m M_g |\hat{u}_\infty(s) - \hat{u}_T(s)|_\infty \right) ds.
\end{aligned}$$

Now let $\epsilon = \max_{t \in [0, \beta]} |\hat{u}_\infty(t) - \hat{u}_T(t)|_\infty$. Since \hat{u}_T converges uniformly to \hat{u}_∞ , then $\epsilon \rightarrow 0$ as $T \rightarrow \infty$. So,

$$|\hat{x}_\infty(t) - \hat{x}_T(t)| \leq \int_0^t \left((K_f + m(M+1)K_g) |\hat{x}_\infty(s) - \hat{x}_T(s)| \right) ds + m M_g \epsilon t.$$

By an application of the Gronwall-Bellman lemma, we obtain

$$\begin{aligned} |\hat{x}_\infty(t) - \hat{x}_T(t)| &\leq mM_g\epsilon t + \int_0^t \left(mM_g\epsilon s(K_f + m(M+1)K_g)e^{(K_f+m(M+1)K_g)(t-s)} \right) ds \\ &= \epsilon \left(mM_g t + \int_0^t \left(mM_g s(K_f + m(M+1)K_g)e^{(K_f+m(M+1)K_g)(t-s)} \right) ds \right) \end{aligned}$$

which tends to zero as ϵ approaches zero. Hence, \hat{x}_T converges pointwise to \hat{x}_∞ on $[0, \beta]$ as $T \rightarrow \infty$ (in fact the convergence is uniform).

3.) $(\hat{x}_\infty, \hat{u}_\infty)$ satisfies the principle of optimality.

By definition, the cost $\hat{V}_T(x(0))$ can be written in terms of \hat{x}_T and \hat{u}_T as,

$$\hat{V}_T(x(0)) = \int_0^T \left(q(\hat{x}_T(t)) + \hat{u}_T^T(t)\hat{u}_T(t) \right) dt$$

where \hat{x}_T and \hat{u}_T satisfy the constraints (5.6) and (5.7). By the principle of optimality, $\hat{x}_T(\tau)$ and $\hat{u}_T(\tau)$ for $\tau \in [\beta, T]$ solves the optimization problem:

$$\begin{aligned} \underset{u_{[\beta, T]}}{\text{minimize}} \quad & \int_\beta^T (q(x) + u^T u) d\tau \\ \text{subject to} \quad & \dot{x} = f(x) + g(x)u \\ & x(\beta) = \hat{x}_T(\beta) \\ & V(x(T)) \leq V(x_\sigma(T)). \end{aligned}$$

(The only difference between this problem and the RHC+CLF problem is that the stability constraint (5.6) is absent since it applies only to the initial control action at time zero (i.e., $\hat{u}_T(0)$.) Let us denote the optimal cost of this problem by $\tilde{V}_{T-\beta}(\hat{x}_T(\beta))$. By an argument identical to that given for \hat{V}_T , we can also prove that \tilde{V}_T converges uniformly to V^* on any compact set. Furthermore, a restatement of the principle of optimality is that

$$\hat{V}_T(x(0)) = \int_0^\beta \left(q(\hat{x}_T(t)) + \hat{u}_T^T(t)\hat{u}_T(t) \right) dt + \tilde{V}_{T-\beta}(\hat{x}_T(\beta)). \quad (5.14)$$

Now take the limit as $T \rightarrow \infty$. On the left-hand side of (5.14), from the first part of this theorem we have that

$$\hat{V}_T(x(0)) \rightarrow V^*(x(0)).$$

Now consider the right-hand side of (5.14). We can show that the second term on the right-hand side converges to $V^*(\hat{x}_\infty(\beta))$ as follows

$$\begin{aligned} |V^*(\hat{x}_\infty(\beta)) - \tilde{V}_{T-\beta}(\hat{x}_T(\beta))| &\leq |V^*(\hat{x}_\infty(\beta)) - V^*(\hat{x}_T(\beta))| \\ &\quad + |V^*(\hat{x}_T(\beta)) - \tilde{V}_{T-\beta}(\hat{x}_T(\beta))|. \end{aligned}$$

The term

$$|V^*(\hat{x}_\infty(\beta)) - V^*(\hat{x}_T(\beta))|$$

tends to zero since V^* is continuous and $\hat{x}_T(\beta)$ converges to $\hat{x}_\infty(\beta)$. Additionally, the term

$$|V^*(\hat{x}_T(\beta)) - \tilde{V}_{T-\beta}(\hat{x}_T(\beta))|$$

tends to zero since by choosing T large enough we can assert that $\hat{x}_T(\beta)$ lies in a compact set (this is because $\hat{x}_T(\beta)$ is a convergent sequence). As mentioned earlier, by the same argument as for \hat{V}_T in the first portion of this theorem, we can assert that $\tilde{V}_{T-\beta}$ converges *uniformly* to V^* on any compact set. Therefore, this term also tends to zero. So, we conclude that

$$\tilde{V}_{T-\beta}(\hat{x}_T(\beta)) \rightarrow V^*(\hat{x}_\infty(\beta)).$$

Finally, we consider the limit of the first term on the right-hand side of (5.14),

$$\lim_{T \rightarrow \infty} \int_0^\beta \left(q(\hat{x}_T(t)) + \hat{u}_T^T(t) \hat{u}_T(t) \right) dt.$$

The dominated convergence theorem [Roy88] justifies an exchange of the limit and integral. By assumption $\hat{u}_T \rightarrow \hat{u}_\infty$ and by step 2.) $\hat{x}_T \rightarrow \hat{x}_\infty$. Hence, this term

converges to

$$\int_0^\beta \left(q(\hat{x}_\infty(t)) + \hat{u}_\infty^T(t) \hat{u}_\infty(t) \right) dt.$$

Therefore, taking the limit as $T \rightarrow \infty$ of equation (5.14) gives

$$V^*(x(0)) = \int_0^\beta \left(q(\hat{x}_\infty(t)) + \hat{u}_\infty^T(t) \hat{u}_\infty(t) \right) dt + V^*(\hat{x}_\infty(\beta))$$

which shows by the principle of optimality that \hat{u}_∞ is optimal for the infinite horizon problem over the interval $[0, \beta]$. ■

Chapter 6 Control of a Ducted Fan Model

6.1 Introduction

In the previous chapter we introduced a new paradigm for optimal control that combines ideas from control Lyapunov function based schemes with an on-line receding horizon approach. In this chapter we mesh the previous chapter's methodology with existing nonlinear control tools by designing and comparing controllers for a simple model of a longitudinal flight control system.

While the focus of this chapter is a specific example, we begin by clarifying the steps involved in nonlinear optimal control design. We offer a two stage design paradigm that clearly separates the controller selection process into an off-line or analysis portion, and an on-line or implementation stage. This allows us to understand the contribution of various existing techniques to the methodology proposed in the previous chapter. This approach is then validated on the ducted fan model.

6.2 An optimal control design paradigm

We will divide the controller design process into the following two distinct steps.

1. Generation of a CLF
2. Selection of a CLF based control scheme

While this distinction is somewhat artificial since most existing techniques span both steps, it helps to clarify the understanding that these techniques actually provide two separate contributions. Furthermore, a single technique does not have to be

used throughout the entire design process, but rather techniques can be “mixed and matched,” often resulting in improved controllers.

We will apply this methodology to the control of the ducted fan model by showing how existing control techniques provide either a CLF, a control law from a CLF, or both. We will compare the following methods: Jacobian Linearization, Frozen Riccati Equations (FRE), Linear Parameter Varying methods (LPV), Control using Global Linearization, and finally Receding Horizon Control (RHC), including hybrid approaches such as Receding Horizon Control combined with the CLF obtained using LPV.

On the generation of a CLF side, we explore Jacobian Linearization, Frozen Riccati Equations (FRE), Global Linearization, and Linear Parameter Varying methods (LPV). While each of these techniques also provides a specific control law, we first focus on the CLF that they produce. When deciding on the choice of a specific control law, we consider the standard implementation of each technique above, plus Sontag’s formula and its receding horizon extension as presented in the previous chapter.

6.3 Caltech ducted fan model

The Caltech ducted fan is a small flight control experiment whose dynamics are representative of either a Harrier in hover mode or a thrust vectored aircraft such as the 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 [KBBM95, BBK96, vNM96]. In this section we describe the simple planar model of the fan shown in Figure 6.3 which ignores the stand dynamics. This model is useful for initial controller design and serves as a good testbed for purposes of this chapter.

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




Figure 6.1: Planar ducted fan model.

mass of the fan, the moment of inertia, and the gravitational constant respectively, the equations of motion can be written as follows:

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

where the drag terms are modeled as viscous friction with d being the viscous friction coefficient. 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 $u_2 = f_2 - mg$, then

$$\begin{aligned}
 m\ddot{x} &= -mg \sin \theta - d\dot{x} + u_1 \cos \theta - u_2 \sin \theta \\
 m\ddot{y} &= mg(\cos \theta - 1) - d\dot{y} + u_1 \sin \theta + u_2 \cos \theta \\
 J\ddot{\theta} &= r u_1.
 \end{aligned} \tag{6.2}$$

These equations are referred to as the *planar ducted fan equations*. We chose the

parameter values:

$$m = 4.25kg, \quad r = 0.26cm, \quad J = 0.0475kg \, m^2, \quad d = 0.1kg/sec, \quad g = 9.8m/sec^2.$$

The following quadratic cost function was used for comparison of different design techniques:

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

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

$$Q = \text{diag}\left(\begin{bmatrix} 10 & 10 & 1 & 1 & 1 & 1 \end{bmatrix}\right) \quad (6.3)$$

Hence, the desired objective was to regulate the states to the origin, or the hover position for the fan. Associated with this optimal control problem is the corresponding *value function*, defined as:

$$V^*(x) = \min_{u(t); x_0=x} \mathcal{J} \quad (6.4)$$

which is also the solution to the Hamilton-Jacobi-Bellman (HJB) partial differential equation:

$$\frac{\partial V^*}{\partial x} f - \frac{1}{4} \frac{\partial V^*}{\partial x} g g^T \frac{\partial V^*}{\partial x} + x^T Q x = 0, \quad V^*(0) = 0$$

6.4 Generation of CLFs

A concept that underlies many nonlinear design methodologies is that of a control Lyapunov function. In simple terms, a control Lyapunov function is the natural extension of the Lyapunov methodology to control systems. To review from previous chapters, consider the following nonlinear system:

$$\dot{x} = f(x) + g(x)u \quad (6.5)$$

where $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$. A *control Lyapunov function* (CLF) is a \mathbf{C}^1 , proper, positive definite function $V : \mathbb{R}^n \rightarrow \mathbb{R}_+$ such that

$$\inf_u \left[\frac{\partial V}{\partial x} f(x) + \frac{\partial V}{\partial x} g(x)u \right] < 0 \quad (6.6)$$

for all $x \neq 0$ [Art83, Son83, Son89].

As was mentioned in the introduction, nonlinear control design can be thought of as having two stages. The first, and perhaps the most challenging stage, is to find a control Lyapunov function. In what follows, we present some of the widely used methods in nonlinear control design, and interpret each approach in the context of the search for a control Lyapunov function.

6.4.1 Jacobian linearization

Perhaps the simplest method of deriving a CLF is to use the Jacobian linearization of the system and generate a CLF by solving an LQR problem. It is a well known result that the problem of minimizing the quadratic performance index:

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

subject to

$$\dot{x} = Ax + Bu$$

is solved by finding the positive definite solution of the following Riccati equation:

$$A^T P + PA - PBB^T P + Q = 0. \quad (6.7)$$

The optimal control action is given by

$$u = -B^T P x$$

with corresponding quadratic CLF:

$$V(x) = x^T P x.$$

In the case of the nonlinear system

$$\dot{x} = f(x) + g(x)u$$

A and B are assumed to be

$$A = \left. \frac{\partial f(x)}{\partial x} \right|_{x=0} \quad B = \left. \frac{\partial g(x)}{\partial x} \right|_{x=0}.$$

Obviously the obtained CLF $V(x) = x^T P x$ will be valid only in a region around the equilibrium. Therefore, we should not expect good performance from initial conditions far from the origin. This is indeed the case as simulation results show that this method cannot stabilize the planar ducted fan model for large initial conditions.

6.4.2 Global linearization

The idea of global linearization has its roots in early works from the Soviet Union [LP44] on the problem of absolute stability. The basic idea behind this approach is to model a nonlinear system as a Polytopic Linear Differential Inclusion (PLDI) [BGF94]. The dynamics of the nonlinear system are approximated as a convex hull of a set of linear models. The problem of quadratic stability of the obtained PLDI, i.e., stability provable by a quadratic Lyapunov function, can be recast as an LMI feasibility problem which can be solved efficiently using interior point convex optimization methods. The PLDI describing the planar ducted fan model can be written as

$$\begin{aligned} \dot{x} &= \sum_{i=1}^2 \alpha_i(t) (A_i x + B_i u) \\ u &= -Kx. \end{aligned} \tag{6.8}$$

Using the same cost function \mathcal{J} as before, the problem of minimizing an upper bound on the cost \mathcal{J} can be written as the following convex optimization problem:

Minimize

$$\text{tr}(Z)$$

Subject to:

$$\begin{aligned} & Y > 0 \\ & \begin{bmatrix} Y A_i^T + A_i Y - B_i X - X^T B_i^T & Y Q^{1/2} & X^T \\ & Q^{1/2} Y & -I & 0 \\ & X & 0 & -I \end{bmatrix} < 0 \\ & \begin{bmatrix} Z & x_0^T \\ x_0 & Y \end{bmatrix} > 0 \\ & i = 1, 2 \end{aligned}$$

where $Y = P^{-1}$ and $X = KY$ are the change of variables made to recast the matrix inequalities as LMIs [BGFB94]. Q is chosen as before, and A_1, B_1 are obtained by linearization of the ducted fan model at the origin and A_2 and B_2 are chosen such that the dynamics lie in the convex hull described by (6.8). This method turns out to be conservative, since there are many trajectories that are a trajectory of the PLDI, but are not a trajectory of the nonlinear system. Using the LMI formulation of the LQR problem for PLDIs [BGFB94], we can find a CLF (given by $V(x) = x^T P x$) for the ducted fan model for positive values of θ . However, a global constant quadratic CLF does not exist. Simulation results for this method show that the closed loop system is stable, but may suffer from poor performance.

6.4.3 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

the Riccati equation online, at each time step. Although results are often promising, there is no rigorous justification for even maintaining mere stability. Nevertheless, the simplicity of the implementation makes the frozen Riccati equation approach a plausible alternative in some applications. To apply this method, the planar ducted fan model is written as

$$\dot{x} = A(x)x + g(x)u. \quad (6.9)$$

At each *frozen state* the Riccati equation is solved, and then the resulting control action is fed back to the system. That is, a state feedback nonlinear control law is obtained by solving the following:

$$\begin{aligned} 0 &= A(x)^T P(x) + P(x)A(x) - P(x)g(x)g^T(x)P(x) + Q \\ u &= -g^T(x)P(x)x. \end{aligned} \quad (6.10)$$

The quantity $V(x) = x^T P(x)x$ generated by this technique is in general only a local CLF. Furthermore, 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 (6.9)) which achieves stability and acceptable performance. However, in the case of the ducted fan model, the *obvious parameterization* of $f(x)$ appears to work in simulation studies. The dynamics of the fan are written as follows:

$$\dot{\bar{x}} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & -\frac{g \sin \theta}{\theta} & -\frac{d}{m} & 0 & 0 \\ 0 & 0 & \frac{g(\cos \theta - 1)}{\theta} & 0 & -\frac{d}{m} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \bar{x} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \frac{\cos \theta}{m} & -\frac{\sin \theta}{m} \\ \frac{\sin \theta}{m} & \frac{\cos \theta}{m} \\ \frac{r}{J} & 0 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} \quad (6.11)$$

Results of this approach are shown in Figure 6.4 at the end of the chapter.

6.4.4 Linear Parameter Varying (LPV) methods

In this technique, the following so-called quasi-LPV representation of a nonlinear input-affine system is used to design a state feedback controller

$$\dot{x} = A(\rho(x))x + B(\rho(x))u \quad (6.12)$$

where ρ is a parameter depending on the state. Hence, we have a linear parameterization of the dynamics through the parameter ρ . Further, it is assumed that the underlying parameter ρ varies in the allowable set

$$\mathcal{F}_{\mathcal{P}}^{\nu} := \{\rho \in \mathbf{C}^1(\mathbb{R}_+, \mathbb{R}^m) : \rho \in \mathcal{P}, \underline{\nu}_i(\rho) \leq \dot{\rho}_i \leq \bar{\nu}_i(\rho), i = 1, \dots, m\} \quad (6.13)$$

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

$$\begin{bmatrix} -\sum_{i=1}^m \bar{\nu}_i(\rho) \frac{\partial X}{\partial \rho_i} + A(\rho)X(\rho) + X(\rho)A^T(\rho) - B(\rho)B^T(\rho) & X(\rho)C^T(\rho) \\ C(\rho)X(\rho) & -I \end{bmatrix} < 0 \quad (6.14)$$

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

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

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 $\sum_{i=1}^m \bar{\nu}_i(\rho)$ in (6.14) means that every combination of $\bar{\nu}_i(\rho)$ and $\underline{\nu}_i(\rho)$ should be included in the inequality. For instance, when $m = 2$, $\bar{\nu}_1(\rho) + \bar{\nu}_2(\rho)$, $\bar{\nu}_1(\rho) + \underline{\nu}_2(\rho)$, $\underline{\nu}_1(\rho) + \bar{\nu}_2(\rho)$ and $\underline{\nu}_1(\rho) + \underline{\nu}_2(\rho)$ should be checked individually. In other words, (6.14) actually represents 2^m inequalities. Additionally, solving (6.14) involves gridding the parameter space \mathcal{P} and choosing a finite set of basis function for $X(\rho)$.

(See [WYPB96] for details.)

For the ducted fan, $\rho = \theta$ was chosen as the varying parameter, and the operating range as $\mathcal{P} = [-\frac{\pi}{2}, \frac{\pi}{2}]$. The bound on the rate variation on θ was set to 10, i.e., $|\dot{\theta}| \leq 10$. Both $A(\rho)$ and $B(\rho)$ were the same as in the model used for the frozen Riccati equation method (eqn. 6.11). A set of basis functions was chosen to compute $X(\rho)$, i.e., $X(\rho) = \sum_{i=1}^5 f_i(\rho) X_i$ where the X_i 's are symmetric coefficient matrices and the $\{f_i(\rho)\}$ are fifth order Legendre polynomials on \mathcal{P} :

$$\{f_i(\rho)\} = \left\{1, \frac{2}{\pi}\theta, \left(3\left(\frac{2}{\pi}\theta\right)^2 - 1\right)/2, \left(5\left(\frac{2}{\pi}\theta\right)^3 - 3\left(\frac{2}{\pi}\theta\right)\right)/2, \left(35\left(\frac{2}{\pi}\theta\right)^4 - 30\left(\frac{2}{\pi}\theta\right)^2 + 3\right)/2\right\}.$$

Simulation of the closed loop system is shown in Figure 6.5 at the end of the chapter.

6.5 CLF based control schemes

So far, we have discussed several methods for generating a CLF. Each of the above mentioned methods have their own technique for generating a controller. However, once a CLF is obtained there are a number of alternative methods that can be used to implement a controller purely from the knowledge of the CLF. We will briefly review some of the options available from previous chapters.

6.5.1 Sontag's formula

We have analyzed Sontag's formula extensively in previous chapters. For reference, we include it once more here:

$$u_{\sigma_s} = \begin{cases} - \left[\frac{\frac{\partial V}{\partial x} f + \sqrt{\left(\frac{\partial V}{\partial x} f\right)^2 + (x^T Q x) \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V}{\partial x} T\right)}}{\frac{\partial V}{\partial x} g g^T \frac{\partial V}{\partial x} T} \right] g^T \frac{\partial V}{\partial x} T & \frac{\partial V}{\partial x} g \neq 0 \\ 0 & \frac{\partial V}{\partial x} g = 0 \end{cases} \quad (6.15)$$

In Chapter 3 we learned that Sontag's formula, in essence, uses the directional information given by the CLF, V , and scales it properly to solve the Hamilton-Jacobi-

Bellman (HJB) equation. That is, Sontag’s formula can be “derived” by assuming the control action to be of the form:

$$u = -\frac{\lambda(x)}{2}g^T \frac{\partial V^T}{\partial x}$$

and determining λ by solving the HJB equation pointwise with $\lambda(x) \frac{\partial V}{\partial x}$ substituting for the gradient of the value function. In particular, if V has level curves that agree with those of the value function, then Sontag’s formula produces the optimal controller. On the other hand, when a CLF does not closely resemble the value function, poor performance may result [FP96]. In the comparison section, this motivates our use of the CLF from LPV in Sontag’s formula.

6.5.2 RHC extensions of CLF formulas

In Chapter 5 we introduced an extension of the class of pointwise min-norm controllers [FK95] to receding horizon schemes. Since Sontag’s formula was shown to be a special case of pointwise min-norm controllers, it also admitted an extension. Recall the RHC+CLF scheme presented in Chapter 5:

(RHC+CLF)

$$\text{minimize} \quad \int_t^{t+T} (x^T Q x + u^T u) d\tau \quad (6.16)$$

$$\text{s.t.} \quad \dot{x} = f(x) + g(x)u \quad (6.17)$$

$$\frac{\partial V}{\partial x}[f + gu(t)] \leq -\epsilon\sigma(x) \quad (6.18)$$

$$V(x(t+T)) \leq V(x_\sigma(t+T)) \quad (6.19)$$

where $1 \geq \epsilon > 0$ and x_σ represents the state trajectory from the pointwise min-norm controller with parameter $\sigma(x)$. We chose the parameter σ to correspond to Sontag’s formula. That is

$$\sigma_s = \sqrt{\left(\frac{\partial V}{\partial x} f\right)^2 + (x^T Q x) \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x}\right)}.$$

For implementation reasons, we replaced the constraint (6.18) with

$$\frac{\partial V}{\partial x}[f + gu(\tau)] \leq 0$$

and applied it over the entire horizon $\tau \in [t, t + T]$. While this approach does not require a fixed horizon length or even a completion of the optimization, again due to the software at our disposal, these properties were not taken advantage of. Results for various horizon lengths are compared in following sections, where we also detail the exact implementation procedures.

6.6 Comparisons

In this section we present a comparison of the approaches presented in the previous sections. By choosing a large time horizon we found the optimal cost for the quadratic cost \mathcal{J} from the chosen initial conditions by solving a single trajectory optimization. This allows us to see exactly how suboptimal techniques are. Values of the cost function for all of the methods described in this chapter are given in Table 1. These costs correspond to the following three initial conditions:

1. $[x, y, \theta, \dot{x}, \dot{y}, \dot{\theta}] = [5, 5, -\frac{0.9\pi}{2}, 5, 0, 0]$
2. $[x, y, \theta, \dot{x}, \dot{y}, \dot{\theta}] = [5, 5, \frac{0.9\pi}{2}, -5, 0, 0]$
3. $[x, y, \theta, \dot{x}, \dot{y}, \dot{\theta}] = [1, 1, \frac{\pi}{4}, 0, 0, 0]$.

The first initial condition is the most difficult of the three, and starts with the fan at a large initial condition and flying away from the origin. The second initial condition is slightly easier, still with a large initial condition but a simpler initial velocity. The third initial condition is close enough to the origin and mild enough that it should not present too difficult a challenge for any of the tested techniques.

A review of Table 6.1 leads to some interesting observations. At the top is Jacobian linearization plus LQR. Not surprisingly, this is found to be unstable for the first initial condition. This illustrates the true nonlinear nature of the problem and indicates that

Method	1	2	3
Jacobian Linearization+LQR	unstable	574	20.6
Frozen Riccati Equation	2801	800	22.0
Global Linearization+LQR using LMIs	2688	1128	29.2
Quasi-LPV	1805	617	27.2
CLF from LPV+Sontag	1761	506	25.4
CLF from LPV+RHC ($T = 0.1, T_s = 0.05$)	1564	468	25.2
CLF from LPV+RHC ($T = 0.3, T_s = 0.05$)	1463	449	22.9
CLF from LPV+RHC ($T = 0.5, T_s = 0.05$)	1421	446	20.5
CLF from LPV+RHC ($T = 1.0, T_s = 0.05$)	1382	434	19.3
Optimal	1368	431	19.3

Table 6.1: Values of the cost function \mathcal{J} using different methods for the ducted fan example.

nonlinear techniques are needed. On the other hand, from both the second and third initial condition, Jacobian linearization performs admirably, even out-doing some of the more sophisticated techniques.

Only slightly more sophisticated than Jacobian linearization plus LQR is the frozen Riccati equation technique. Even though it also lacks global stability guarantees, it is stabilizing from all three initial conditions, although with a rather poor cost. Simulation results for the frozen Riccati equation approach are supplied in Figure 6.4.

Next, we find that while global linearization techniques provide a guarantee of stability, on this example they suffer from very poor performance. To retain the guarantee of stability, but also aim for improved performance, more off-line computation must be thrown at the problem as in LPV techniques. The result of the LPV simulation from the first initial condition is given in Figure 6.5. We found that LPV provides reasonable levels of performance for all three initial conditions.¹ This, combined with the fact that it provides a global CLF, indicates that it might provide a fair representation of the true value function. Hence, the CLF from LPV is a reasonable choice for use in Sontag’s formula.

¹Note that in LPV and frozen Riccati equation techniques, a design choice is involved in the selection of a state dependent representation. Although no systematic procedure was used, the results obtained here were the best of the state dependent representations that were tested.

Applying Sontag’s formula with the aid of the CLF from LPV resulted in trajectories very similar to those obtained from the standard LPV implementation, although with slightly reduced costs from all three initial conditions. It was this controller that we decided to extend to an on-line receding horizon implementation.

Details of the implementations of the on-line RHC+CLF controllers are given later, but at first glance we observe that on-line computation is quite beneficial in terms of the cost. As the horizon was increased from $T = 0.1$ to eventually $T = 1.0$, the cost steadily decreased, providing the lowest cost observed in any of the simulations. Due to the similarity in results, only the optimal trajectory is supplied in Figure 6.6 for reference.

To summarize, in general the following trends were observed. While not uniformly true, the more detailed and sophisticated techniques, which generally involve extensive off-line analysis, tended to outperform the simpler, less theoretically sound techniques. Extensive computation was also found to be extremely beneficial, especially when employed in an on-line manner, but only when used under the guidance of a solid theoretical framework.

6.7 Implementation of on-line schemes

While the example provided in this chapter gives strong indication that on-line computation can be extremely beneficial, implementation issues can easily discourage its use. Therefore, we provide some of the details of the implementation procedure used in this chapter, pointing out potential pitfalls along the way.

The schemes involving on-line computation (RHC+CLF) were implemented with the use of the RIOTS² trajectory optimization software package. This package runs off of the nonlinear programming package NPSOL.³ RIOTS uses direct shooting methods and parameterizes input trajectories over a finite dimensional spline space to solve

²RIOTS stands for “Recursive Integration Optimal Trajectory Solver” and was written by Adam Lowell Schwartz as part of his Ph.D. thesis at UC Berkeley, 1996.

³NPSOL can be purchased from Stanford Business Software, Inc., 2680 Bayshore Parkway, Suite 304, Mountain View, CA 94043.

constrained trajectory optimization problems.

In our implementation, we first used the fact that Sontag's formula is a stabilizing state feedback controller to pre-stabilize the system. We then applied RIOTS to the pre-stabilized system. Shooting methods can have difficulty when applied to open-loop unstable systems, so by pre-compensating the system with Sontag's formula we removed this problem. In fact, when RIOTS was applied to the open-loop system (which is unstable), we encountered numerous numerical difficulties. Hence, the fact that a CLF also provides a stabilizing control law which can be used to pre-stabilize the dynamics before performing trajectory optimizations is yet another example of the synergies available between CLFs and receding horizon control. The resulting optimizations appeared to be very well conditioned for shooting techniques, and no further numerical problems were encountered.

For each of our trajectory optimizations we selected the RK45 integration option in RIOTS and fixed the time step size at $0.025s$. (The number of time steps used for each horizon length is as given in Table 6.2.) We also chose to use the *warm start* option available in RIOTS. This uses the Lagrange multipliers from the previous RIOTS solution as an initial guess at the multipliers for the new problem. Although the difference was not significant, it was generally perceived that this sped up the optimizations. All simulations were performed on a $450Mhz$ Pentium II processor.

6.7.1 Time considerations

The time required to solve on-line optimizations is perhaps the single most important factor limiting the application of receding horizon techniques. Realistically, many of the proposed receding horizon schemes, both in this thesis and elsewhere, are currently beyond present computing capabilities. Nevertheless, in the not so distant future they will be viable, indicating that they will represent a real alternate for control design.

There are two basic tradeoffs relating to computation time: time versus horizon length and time versus complexity of the optimization problem. We will present rough tradeoffs for both by comparing various implementations on the ducted fan model.

Time versus Horizon

We begin with a comparison of the time required for implementation of the same RHC+CLF scheme but under different horizons. In Table 6.2 we list the number of integration time steps used for each horizon, and the average time required to solve each receding horizon optimization for more than 100 initial conditions.

Horizon (T)	# of time steps	avg. time per opt.
$T = 0.1$	4	1.06
$T = 0.3$	12	6.15
$T = 0.5$	20	13.15
$T = 1.0$	40	58.42

Table 6.2: Horizon versus the number of time steps used in the integration scheme RK45.

A more accurate picture of computation times is presented in Figure 6.2 where we have plotted the time required by RIOTS to solve each receding horizon trajectory optimization along the trajectory beginning from the first initial condition $[5; 5; -0.9(\pi/2); 5; 0; 0]$.

As is fairly evident, computation times rise rather dramatically as a function of the horizon length. These results should be considered in light of the following fact. The constraint (6.18) ($\dot{V} < 0$) was imposed over the entire horizon T , not merely over the sampling time T_s . This was forced upon us by RIOTS. This means that the longer the horizon and the more sampling points chosen along that horizon, the more constraints were added to the problem. This fact alone makes the optimization numerically more difficult for longer horizons. Below, we will see more explicitly the effect that various constraints have on the overall computation time.

Time versus Constraint Complexity

Below we compare the amount of time each optimization takes for various constraints. In order to see the effect of the constraints on optimization time, we implemented three versions of receding horizon control.

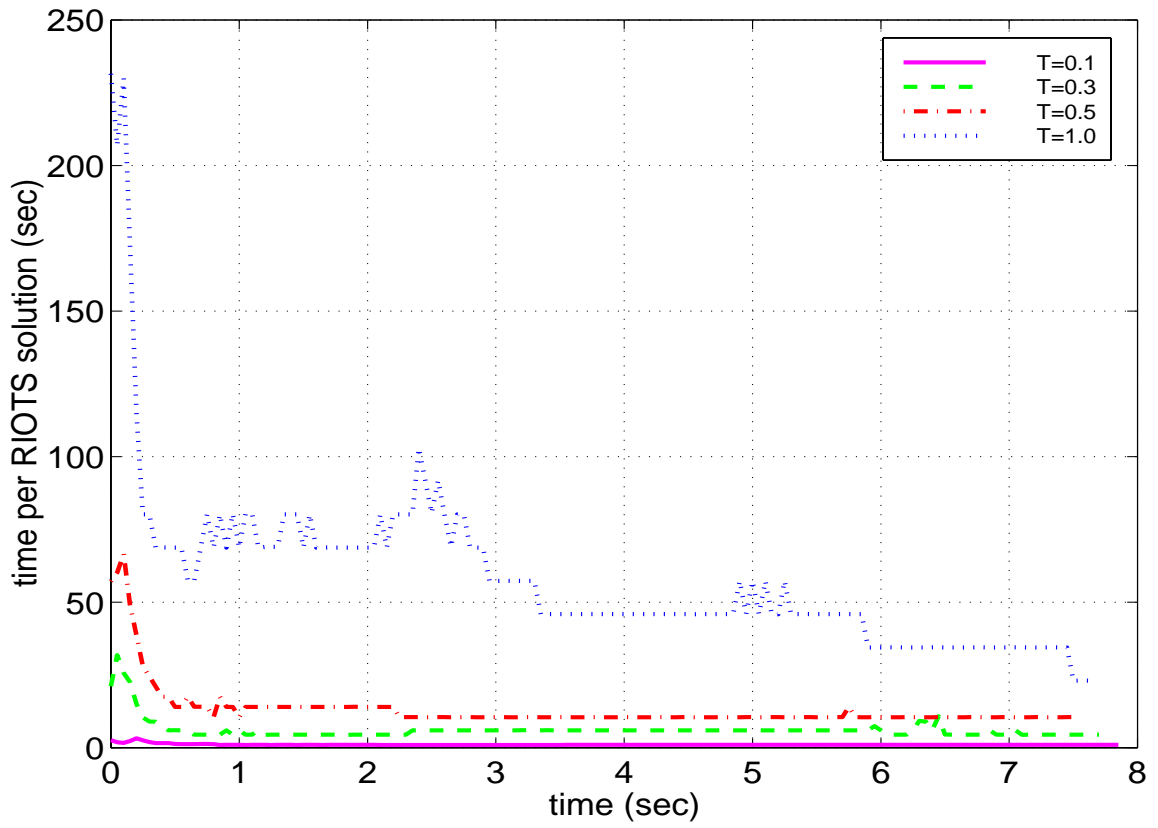


Figure 6.2: Comparison of time required for each RIOTS solution between RHC schemes with horizons of $T = 0.1s$ (solid), $T = 0.3s$ (dash), $T = 0.5s$ (dash-dot), $T = 1.0s$ (dotted).

1. Standard RHC+CLF

The standard RHC+CLF implementation involves two constraints (6.18) and (6.19). Recall that we were imposing the constraint (6.18) over the entire prediction horizon T .

2. RHC+CLF without constraint (6.18)

In this implementation we removed the constraint (6.18), but retained the end point constraint (6.19).

3. RHC with the CLF as a terminal weight

We applied receding horizon control with no constraints, but with the CLF as a terminal weight (i.e., $\varphi(\cdot) = V(\cdot)$).

All three implementations used the horizon $T = 0.3s$. The computation times associated with each of these implementations is plotted in Figure 6.3. It shows the time required for RIOTS to solve each on-line optimization for the simulation from the first initial condition $[5; 5; -0.9(\pi/2); 5; 0; 0]$.

From these plots we see that the constraint over the entire horizon adds a substantial amount of time to each optimization. On the other hand, there is not a large time difference between the implementation using an end-point constraint and the implementation using a terminal weight. The cost obtained by the unconstrained scheme was 1493, while the constrained approaches achieved a cost of 1463. These results indicate that trajectory constraints may be time consuming, but a single end point constraint does not add substantial difficulty over no constraints.

6.8 Summary

In this chapter we presented a concrete example of the framework for nonlinear optimal control developed in preceding chapters. We placed existing techniques in a two stage design procedure. The first is the derivation of a CLF. Potential techniques for

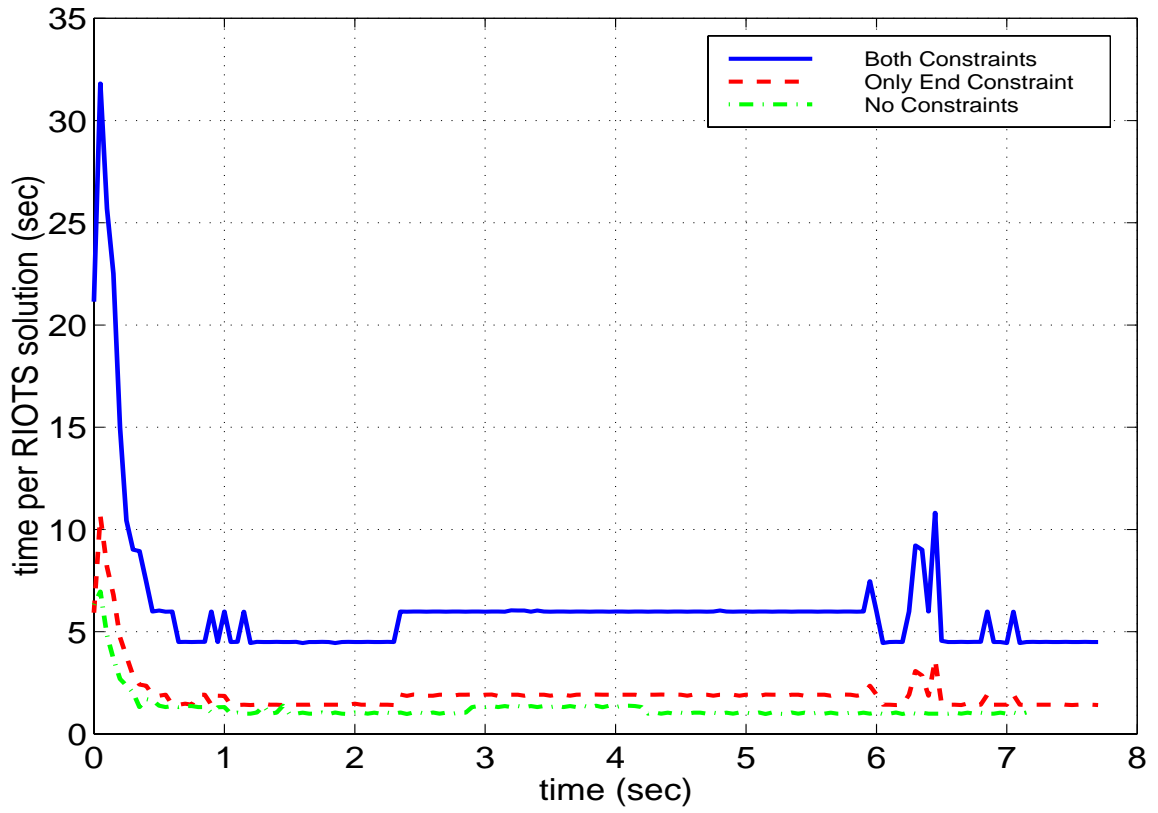


Figure 6.3: Comparison of time required for each RIOTS solution between RHC schemes with horizon $T = 0.3s$: RHC+CLF (solid), RHC+CLF w/o constraint (6.18) (dash), RHC+CLF w/ CLF as terminal weight (dash-dot).

this stage were: Jacobian linearization, global linearization, frozen Riccati equations, and linear parameter varying (LPV) techniques. The second stage involves using the CLF to produce a control scheme. In this step, one has additional choices including Sontag's formula, pointwise min-formulas, and their extensions to receding horizon schemes.

A ducted fan model was used as the test case for this design methodology, and a simulation study was used to test the results. It was found that a combination of off-line analysis in determining a CLF and on-line computation produced the best results. But, on-line results typically come at high implementation prices. Through simulation examples, we analyzed the fundamental issues facing the implementation of the proposed RHC+CLF schemes. While certain constraints and implementations do not appear to be numerically limiting or difficult, trajectory constraints were found to add considerable complexity, especially over long horizons.

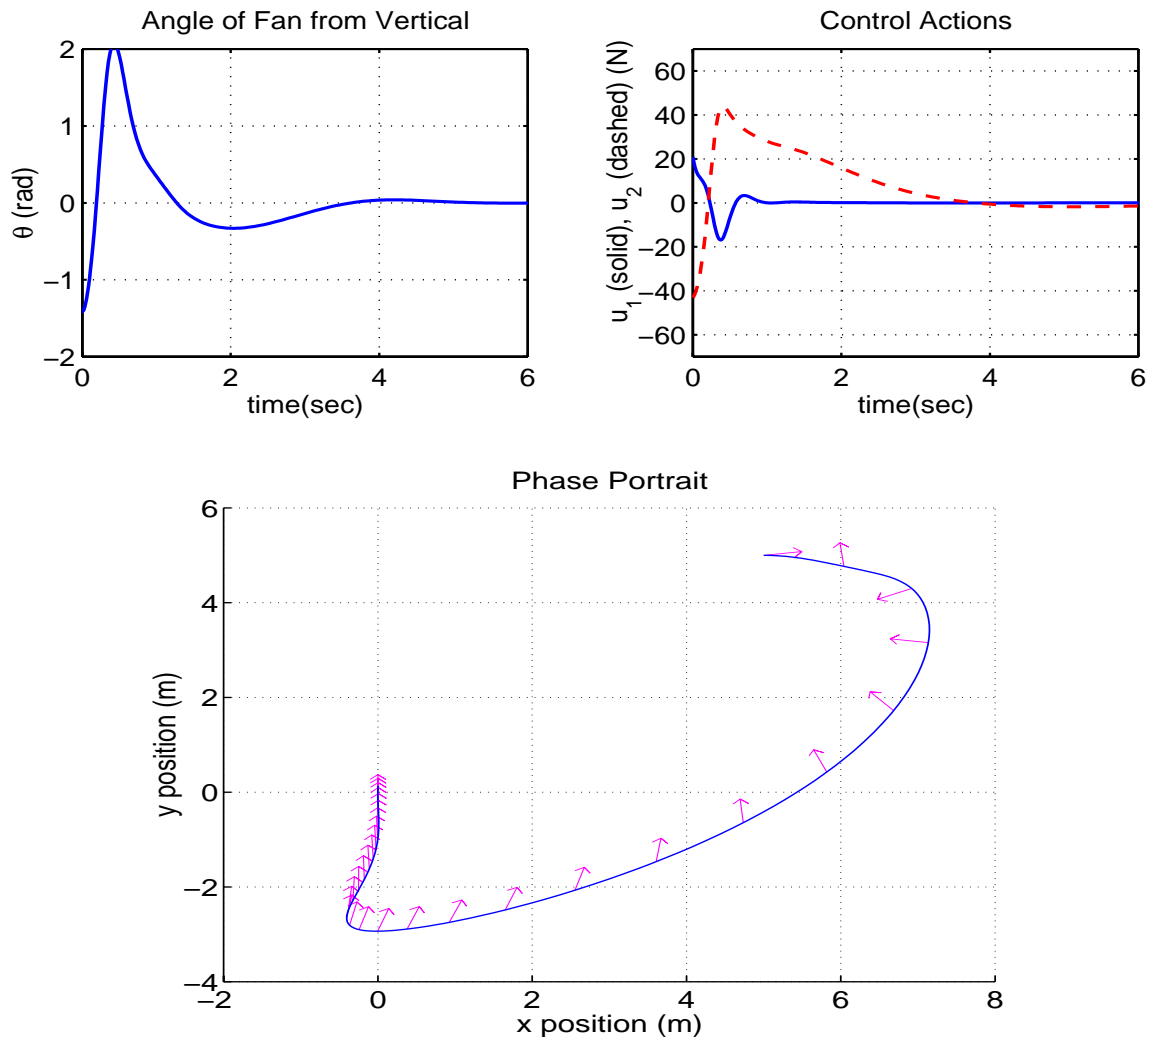


Figure 6.4: The frozen Riccati controller, initial condition $[5; 5; -0.9(\pi/2); 5; 0; 0]$.

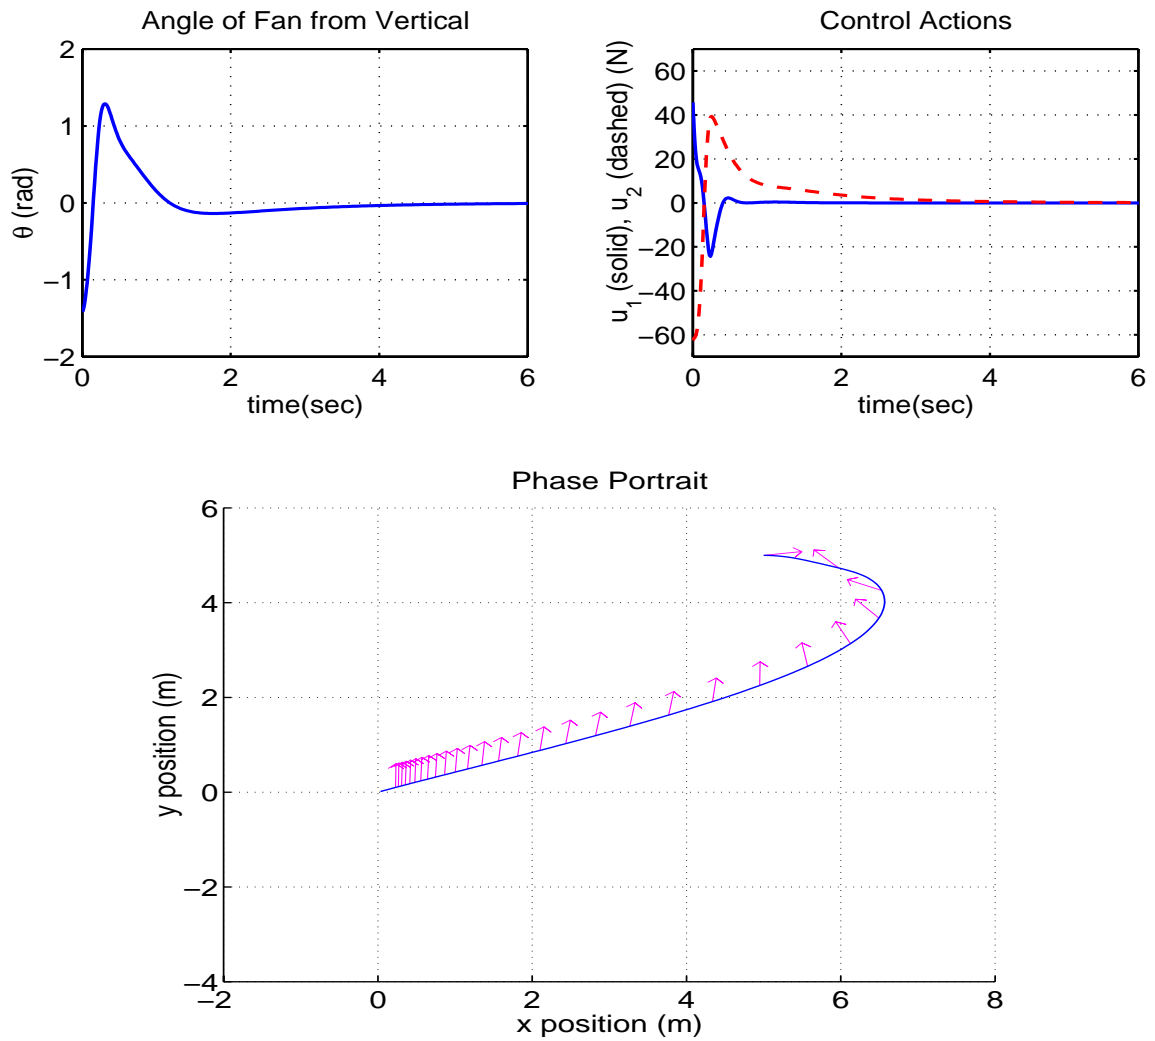


Figure 6.5: The LPV controller, initial condition $[5; 5; -0.9(\pi/2); 5; 0; 0]$.

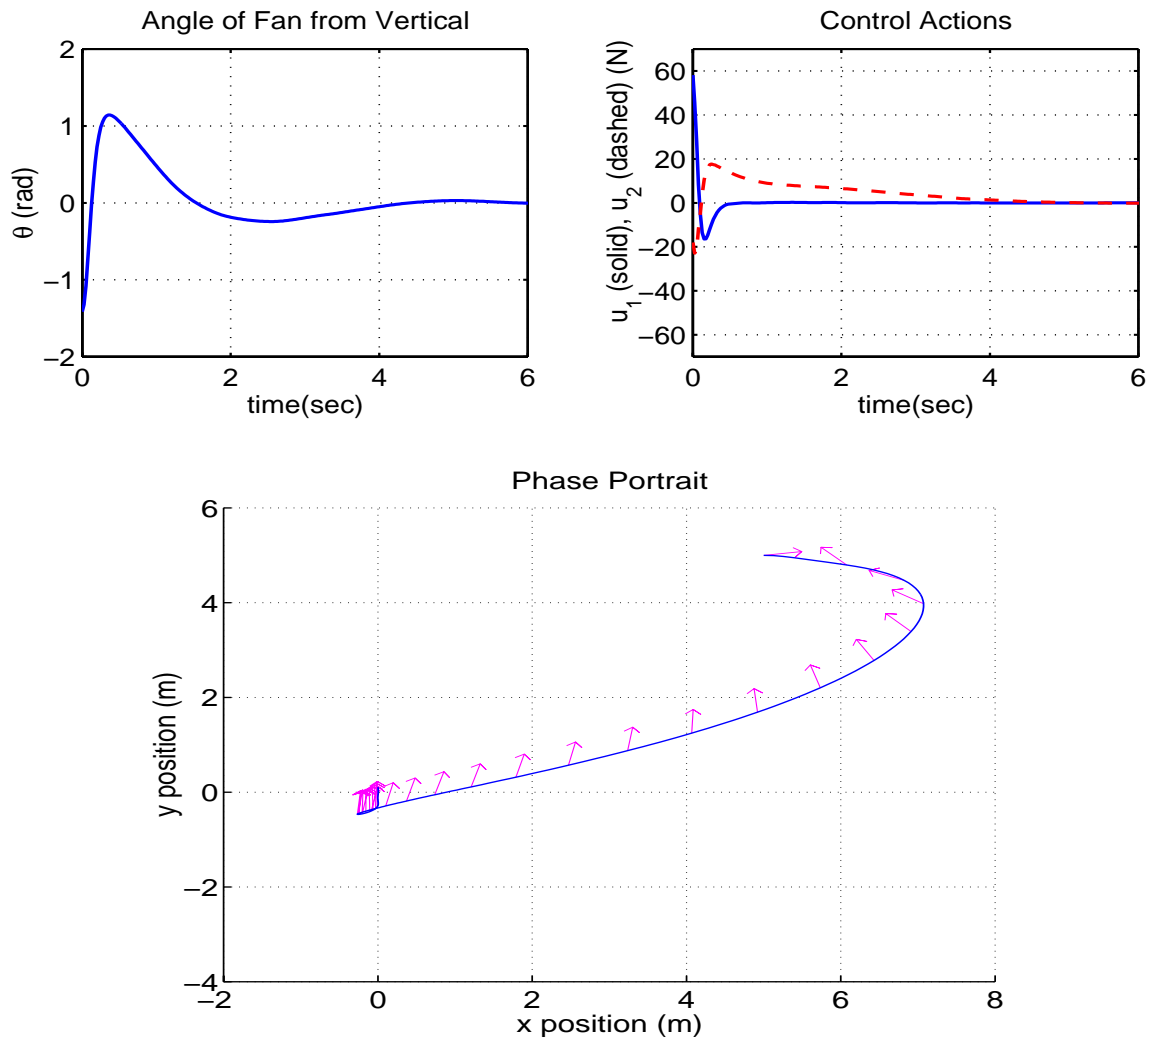


Figure 6.6: Optimal, initial condition $[5; 5; -0.9(\pi/2); 5; 0; 0]$.

Chapter 7 Extensions

7.1 Introduction

In this chapter we present two important extensions to the framework developed in Chapter 5. First we extend the methodology to handle time-varying dynamics. These results follow in a straightforward manner from those in Chapter 5. Next, we confront the issue of input constraints. While receding horizon control can naturally incorporate constraints directly into its on-line optimizations, pointwise min-norm controllers must be reformulated before their receding horizon extensions will carry through. In both cases, simple two dimensional examples are used to illustrate key points.

7.2 Time-varying optimal control

Currently the focus of nonlinear control research is directed toward time-invariant systems. Specifically, many modern approaches focus on the determination of a control Lyapunov function (CLF) [Kha96]. While the advantages of a CLF approach have been well documented for time-invariant nonlinear systems, the time-varying problem has received far less attention [AS97]. Yet, time-varying control problems naturally arise by considering the error dynamics in trajectory tracking problems. In this section we focus not on determining CLFs for time-varying dynamics, but on the selection of stabilizing control laws from a CLF. Following the new methodology presented in Chapter 5 for the incorporation of CLFs into on-line receding horizon schemes, in this section those results are extended to the time-varying case. Finally, the new schemes are tested on a trajectory tracking problem for a simple two-dimensional example.

7.2.1 Optimal control for time-varying systems

Consider a time-varying nonlinear control affine system

$$\begin{aligned} \dot{x} &= f(x, t) + g(x, t)u & f(0, t) = 0 & \quad \forall t \\ y &= h(x, t) \end{aligned} \tag{7.1}$$

with an infinite horizon objective

$$\begin{aligned} \min_{u(\cdot)} & \int_0^{\infty} (q(x) + u^T u) dt \\ \text{s.t.} & \quad \dot{x} = f(x, t) + g(x, t)u \end{aligned} \tag{7.2}$$

where $q(x)$ is continuously differentiable, positive semi-definite and $[f, q]$ is zero-state detectable.

Using a standard dynamic programming approach ([BH75] and Chapter 2), the above optimal control problem can be reduced to the time-varying Hamilton-Jacobi-Bellman optimization equation

$$-\frac{\partial V^*}{\partial t} = \min_{u(t)} \left\{ q(x) + u^T u + \frac{\partial V^*}{\partial x} [f + gu] \right\} \tag{7.3}$$

where once again

$$V^*(x, t) = \min_{u(\cdot)} \int_t^{\infty} (q(x) + u^T u) d\tau, \tag{7.4}$$

i.e., $V^*(x, t)$ is the *value function* and can be thought of as the minimum cost to go from the state $x(t)$. Performing the optimization in (7.3) leads to a control law of the form

$$u^* = -\frac{1}{2} g^T \frac{\partial V^{*T}}{\partial x}. \tag{7.5}$$

Substituting this in (7.3) results in the time-varying Hamilton-Jacobi-Bellman (HJB) partial differential equation

$$\frac{\partial V^*}{\partial t} + \frac{\partial V^*}{\partial x} f - \frac{1}{4} \frac{\partial V^*}{\partial x} g g^T \frac{\partial V^{*T}}{\partial x} + q(x) = 0 \tag{7.6}$$

whose solution is the value function V^* . Note that the only difference between the HJB equation in the time-varying case versus the time-invariant case (see eqn. (2.8)) is the term $\frac{\partial V}{\partial t}$.

7.2.2 CLF formulas for time-varying systems

While control Lyapunov function design is routinely applied to time-invariant nonlinear systems, a somewhat less established area is the use of control Lyapunov functions for *time-varying* nonlinear systems. To extend the concept of a control Lyapunov function to time-varying systems, first recall the following definition:

Definition 7.2.1 *A continuous function $\alpha : [0, a) \rightarrow [0, \infty)$ is said to belong to class \mathcal{K} if it is strictly increasing and $\alpha(0) = 0$. It is said to belong to class \mathcal{K}_∞ if $a = \infty$ and $\alpha(r) \rightarrow \infty$ as $r \rightarrow \infty$.*

We can then define a time-varying control Lyapunov function as follows:

Definition 7.2.2 *A function $V(x, t) : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^+$ is a global control Lyapunov function if:*

1. $V(x, t) \in C^1$,
2. There exist \mathcal{K}_∞ functions α_1, α_2 and a \mathcal{K} function α_3 , such that

$$\begin{aligned} \alpha_1(|x|) &\leq V(x, t) \leq \alpha_2(|x|), \quad \forall t \\ \inf_u \left\{ \frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} [f + gu] \right\} &\leq -\alpha_3(|x|), \quad \forall x, t \end{aligned}$$

7.2.3 A time-varying Sontag's formula

Let $V(x, t)$ be a CLF and for pedagogical purposes assume that the value function $V^*(x, t)$ is related to the CLF in the following manner.

$$V^*(x, t) = \Lambda(V(x, t)) \tag{7.7}$$

where Λ denotes a function from $\mathbb{R} \rightarrow \mathbb{R}$. Then we may write the HJB equation (7.6) in terms of the CLF as

$$\left(\frac{\partial \Lambda}{\partial V}\right) \frac{\partial V}{\partial t} + \left(\frac{\partial \Lambda}{\partial V}\right) \frac{\partial V}{\partial x} f - \left(\frac{\partial \Lambda}{\partial V}\right)^2 \frac{1}{4} \frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x} + q(x) = 0 \quad (7.8)$$

which upon solving (7.8) as a quadratic equation in terms of $\frac{\partial \Lambda}{\partial V}$ yields

$$\frac{\partial \Lambda}{\partial V} = \frac{\left(\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f\right) + \sqrt{\left(\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f\right)^2 + q(x) \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x}\right)}}{\frac{1}{2} \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x}\right)}. \quad (7.9)$$

Finally, recalling that the optimal control action is given by

$$u^* = -\frac{1}{2} g^T \frac{\partial V^{*T}}{\partial x} = -\frac{1}{2} \left(\frac{\partial \Lambda}{\partial V}\right) g^T \frac{\partial V^T}{\partial x} \quad (7.10)$$

leads to the following form for the optimal controller upon substitution of (7.9):

$$u^* = \begin{cases} - \left(\frac{\left(\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f\right) + \sqrt{\left(\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f\right)^2 + q(x) \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x}\right)}}{\left(\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x}\right)} \right) (g^T \frac{\partial V^T}{\partial x}) & , \frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x} \neq 0 \\ 0 & , \frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x} = 0. \end{cases} \quad (7.11)$$

This is an explicit formula for a control law as a function of the CLF. One might note that this derivation is essentially identical to that used to derive the time-invariant version in Chapter 3. Although it was derived by assuming the relationship (7.7), one can ask whether this is a valid stabilizing control law for an arbitrary CLF. This is simple to check by considering the time derivative of the CLF for the closed loop system:

$$\dot{V} = -\sqrt{\left(\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f\right)^2 + (q(x)) \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V^T}{\partial x}\right)}.$$

Hence, it is sufficient for stability to verify that there exists a \mathcal{K} function that bounds

$$\sqrt{\left(\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f\right)^2 + (q(x)) \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V}{\partial x}\right)}$$

from below, which is almost always the case (when this is not the case, one can often argue stability from LaSalle's invariance principle [Kha96]).

One might recognize that (7.11) is similar to Sontag's formula for time-invariant systems [Son89] and furthermore is the direct extension of the formula presented in Chapter 3. A simple modification of Sontag's original proof [Son89] shows that it enjoys the same continuity properties as in the time-invariant case, namely that when $q(x)$ is positive definite, it is as smooth as the data $(\frac{\partial V}{\partial x} f, \frac{\partial V}{\partial t}, \frac{\partial V}{\partial x} g$ and $q(x))$ except possibly at the origin. Furthermore, one can note that this scheme will produce an optimal controller for any CLF that actually satisfies the condition (7.7). Of course, realistically this condition cannot be expected to occur except in extremely rare cases.

7.2.4 Receding horizon extensions of CLF schemes

An alternate route to the formula in (7.11) is through the solution of the following pointwise min-norm problem:

$$\text{minimize} \quad u^T u \tag{7.12}$$

$$\text{subject to} \quad \frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} [f + gu] \leq -\sigma \tag{7.13}$$

where

$$\sigma = \sqrt{\left(\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f\right)^2 + (q(x)) \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V}{\partial x}\right)}.$$

A detailed discussion of the pointwise min-norm methodology can be found in Chapter 3. This alternate perspective provides the appropriate starting point for an extension of (7.11) to a receding horizon scheme. In Chapter 5, it was shown in the time-invariant case that pointwise min-norm schemes can be naturally extended to on-line receding horizon schemes that solve a *finite horizon optimal control problem* at every

encountered state. An analogous construction for time-varying systems leads to the following CLF based receding horizon scheme:

$$\text{minimize} \quad \int_t^{t+T} q(x(\tau)) + u^T(\tau)u(\tau)d\tau \quad (7.14)$$

$$\text{subject to} \quad \dot{x} = f(x, \tau) + g(x, \tau)u \quad (7.15)$$

$$\frac{\partial V}{\partial t} + \left(\frac{\partial V}{\partial x} \right) [f(x(t), t) + g(x(t), t)u(t)] \leq -\epsilon\sigma \quad (7.16)$$

$$V(x(t+T), t+T) \leq V(x_\sigma(t+T), t+T) \quad (7.17)$$

where x_σ represents the state trajectory produced by (7.12-7.13) and $0 < \epsilon \leq 1$ is a design parameter used to relax the constraint (7.16). This optimization is solved at each encountered state and the resulting solution is implemented in a receding horizon fashion.

As in the time-invariant case, this scheme possesses critical implementation properties that facilitate its efficient use of on-line computation:

1. Equation (7.11) provides a feasible control action.
2. Guaranteed stability for any horizon length (or variable horizons).
3. No requirement of an optimizing solution for stability.

The reader is referred to Chapter 5 for details.

Before applying these techniques to an example, in the next section we briefly outline one method for obtaining CLFs for trajectory tracking involving feedback linearizable systems.

7.2.5 CLFs for feedback linearizable systems

The subject of deriving CLFs is an active and vast research area in itself. Any attempt to cover all of the various available approaches, even at a superficial level, could occupy an entire thesis by itself. Hence, we will limit our scope to serve our

specific purpose here and briefly demonstrate how one may determine a CLF for trajectory tracking in feedback linearizable systems [Isi95]. There are two simplifying reasons that we choose to focus on this class:

1. Given a desired output trajectory it is possible to compute the corresponding state and inputs required to produce the output [Isi95].
2. There do not exist systematic techniques for finding CLFs for general nonlinear systems, but for feedback linearizable systems a quadratic function in the linearized coordinates may be used.

More specifically, consider the trajectory tracking problem for full-state feedback linearizable nonlinear control affine systems:

$$\dot{x} = \hat{f}(x) + \hat{g}(x)u \quad (7.18)$$

$$y = \hat{h}(x) \quad (7.19)$$

where $y_r(t)$ is a desired reference trajectory and satisfies

$$\dot{x}_r = \hat{f}(x_r) + \hat{g}(x_r)u_r \quad (7.20)$$

$$y_r = \hat{h}(x_r). \quad (7.21)$$

Since (7.18) is full state feedback linearizable, there exists a suitable change of coordinates and feedback transformation such that (7.18) is transformed into a linear and controllable system [Isi95]. In the new coordinates, $z = \beta(x)$, the system will be described by equations of the form

$$\begin{aligned} \dot{z}_1 &= z_2 \\ \dot{z}_2 &= z_3 \\ &\vdots \\ \dot{z}_{n-1} &= z_n \\ \dot{z}_n &= v \\ y &= z_1 \end{aligned} \quad (7.22)$$

where n is the order of the system and $v = b(z) + a(z)u$. More compactly, we may write the resulting linear system as

$$\dot{z} = Az + bv \quad (7.23)$$

$$y = cz. \quad (7.24)$$

Similarly, the reference trajectory in the z -coordinates is given by

$$\dot{z}_r = Az_r + bv_r \quad (7.25)$$

$$y_r = cz_r. \quad (7.26)$$

By defining error signals $\tilde{z} = z - z_r$, $\tilde{v} = v - v_r$, and $\tilde{y} = y - y_r$, it follows that the error dynamics are also linear:

$$\dot{\tilde{z}} = A\tilde{z} + b\tilde{v} \quad (7.27)$$

$$\tilde{y} = c\tilde{z}. \quad (7.28)$$

Stabilizing these dynamics is equivalent to tracking the desired reference signal and hence a CLF for these dynamics will also be one for the original trajectory tracking problem. A CLF can be easily determined by solving the Riccati equation corresponding to an LQR problem,

$$A^T P + PA - PBR^{-1}B^T P + Q = 0, \quad (7.29)$$

and using the resulting solution

$$\tilde{V}(\tilde{z}) = \tilde{z}^T P \tilde{z}. \quad (7.30)$$

Expressing this CLF in terms of the original error coordinates $\tilde{x} = x - x_r$ leads to a time-varying CLF:

$$V(\tilde{x}, t) = (\beta(\tilde{x} + x_r) - z_r)^T P (\beta(\tilde{x} + x_r) - z_r). \quad (7.31)$$

We can use (7.31) for the formula in (7.11), or to solve the receding horizon optimization problem as indicated in (7.14) - (7.17).

7.2.6 Example

Consider the following two-dimensional nonlinear oscillator:

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= -x_1 \left(\frac{\pi}{2} + \arctan(5x_1) \right) - \frac{5x_1^2}{2(1 + 25x_1^2)} + u \end{aligned}$$

with output $y = x_1$. The problem is to track the reference signal

$$y_r = \sin(t)$$

while minimizing the cost functional

$$\int_0^\infty (\tilde{x}_1^2 + 0.1\tilde{x}_2^2 + \tilde{u}^2) dt$$

where $\tilde{x} = x - x_r$ and $\tilde{u} = u - u_r$ are the state and control “error” signals. Since these dynamics are feedback linearizable with no coordinate change, it is easy to see that the quadratic function:

$$\tilde{x}^T \begin{bmatrix} 1.45 & 1 \\ 1 & 1.45 \end{bmatrix} \tilde{x}$$

is a CLF for the error system. This CLF results from solving the LQR problem with the given cost and linearized dynamics.

Both a feedback linearized controller and the CLF based controller as presented in Section 7.2.2 were tested on this system. The results indicate that a priori there is no advantage in one design technique over another. In fact, depending on the initial condition chosen, either controller can outperform the other considerably. Note that this is despite the strong connection that the CLF based formula has with the HJB equation.

Consider, for example, the initial condition $[3, -2]$. From this starting point the feedback linearized controller outperforms the CLF controller by a cost of 62 to 85. The corresponding trajectories are shown in Figure 7.1. On the other hand, from the initial condition $[1, 6]$, the CLF based controller achieves a cost of 59 compared to 103 for the feedback linearized controller. These results are given in Figure 7.2.

It is important to recognize that while the feedback linearized and CLF controller do not seem to possess inherent advantages over one another, the receding horizon scheme produced significantly improved performance over both. The results are presented in Figures 7.3 (initial condition $[3, -2]$) and 7.4 (initial condition $[1, 6]$) and show the improvement that is possible by utilizing on-line computation in accordance with the scheme presented in Section 7.2.4. For the same initial conditions, we tested the horizon lengths $T = 0.1$ and $T = 0.25$ (with a sampling time of $T_s = 0.05s$ and $\epsilon = 0.05$). From both initial conditions, a horizon of only $T = 0.1$ improved upon the CLF controller, but for the initial condition $[3, -2]$ it still did not achieve a performance better than the feedback linearized controller. By increasing the horizon to $T = 0.25$, a dramatic improvement over the horizon of $T = 0.1$ was apparent, and these controllers performed far better than either the CLF or feedback linearized controller. For the initial condition $[3, -2]$ it even transformed the poor performing CLF scheme into a controller that outperformed the others by a wide margin. A summary of the results is supplied in Table 7.1.




Figure 7.1: State and control trajectories from initial condition $[3, -2]$: Reference (dotted), CLF (dashed) and feedback linearized (dash-dot).




Figure 7.2: State and control trajectories from initial condition $[1, 6]$: Reference (dotted), CLF (dashed) and feedback linearized (dash-dot).




Figure 7.3: State and control trajectories from initial condition $[3, -2]$: Reference (dotted), RHC+CLF $T = 0.10$ (dashed) and $T = 0.25$ (dash-dot).




Figure 7.4: State and control trajectories from initial condition $[1, 6]$: Reference (dotted), RHC+CLF $T = 0.10$ (dashed) and $T = 0.25$ (dash-dot).

Controller	[3, -2]	[1, 6]
Feedback Lin.	61.7	103.5
CLF	85.3	59.3
RHC+CLF ($T = 0.10$)	79.7	55.8
RHC+CLF ($T = 0.25$)	42.9	34.8

Table 7.1: Comparison of time-varying controller costs.

7.3 Input constrained systems

Input saturations represent an inherent limitation on actuators and arise in virtually every problem of practical interest. Previously, constrained systems have only been studied explicitly by a small portion of the control community. However, in recent years there has been a renewed interest in the study of both linear and nonlinear systems subject to input saturations.

More recently, focus has shifted toward techniques for constrained nonlinear systems that employ a control Lyapunov function point of view. This approach consists of first deriving a control Lyapunov function for the constrained system, and then determining a constrained input law consistent with the control Lyapunov function. Research addressing the problem of determining a CLF for a constrained nonlinear system has been quite active lately. While similar to the unconstrained case in that no general systematic procedure exists for the derivation of a constrained CLF, procedures have emerged to construct CLFs for special classes of constrained systems. In [Lin94] conditions are obtained which ensure global asymptotic stability for control affine nonlinear systems such that their free dynamics are asymptotically stable. More general results may be found in [MP96] for systems in the so called “forwarding” form and in [FP98] for systems in the “backstepping” form. These references provide a method for the construction of a constrained CLF when the system possesses special structure.

The second stage in the control design procedure is the actual selection of a bounded control law from the knowledge of a CLF. While most techniques for the

determination of a constrained CLF also result in a constrained control scheme, there has additionally been work on the selection of a control law purely from the knowledge of a constrained CLF. For instance, Lin and Sontag [LS91] have derived a smooth control law in which the control actions take values in the unit ball, extending the well known results obtained in the unconstrained case [Son89].

In this section we extend the framework presented in previous chapters to handle input constraints. This is done by first presenting a new pointwise min-norm scheme for constrained nonlinear systems, then showing that it easily extends to a receding horizon problem. Finally, a simple example demonstrates the new methodology.

7.3.1 Constrained nonlinear optimal control

We will consider nonlinear systems of the form:

$$\dot{x} = f(x) + g(x)u, \quad f(0) = 0 \quad (7.32)$$

with $x \in \mathbb{R}^n$ denoting the state and $f(x), g(x) \in \mathbf{C}^1$. The input will be constrained to lie in a specified set:

$$u \in \Omega_u \subset \mathbb{R}^m$$

where it is assumed that Ω_u contains a neighborhood of the origin.

Our motivation will derive from the constrained infinite horizon nonlinear optimal control problem, stated as follows:

$$\begin{aligned} & \text{minimize} && \int_0^\infty (q(x) + u^T u) dt \\ & \text{subject to} && \dot{x} = f(x) + g(x)u \\ & && u(\cdot) \in \Omega_u. \end{aligned} \quad (7.33)$$

This is the standard nonlinear regulator problem, with the desired solution being a state feedback controller $u^* = k(x)$.

As has been extensively outlined in previous chapters, the problem (7.33) is in general

prohibitively difficult, even when there are no constraints. In the constrained case issues of feasibility complicate the problem even further, leading to questions of the mere existence of a stabilizing controller.

While model predictive control has always been primarily motivated by constraints, it is only recently that methodologies for the derivation of control Lyapunov functions have broached the subject [MP96, JSK96, FP98]. As this subject matures, it will become increasingly useful to develop a theory analogous to that presented in Chapter 5, but for the constrained optimal control problem. Our present aim is to extend those results to the case of saturated inputs.

7.3.2 A stabilizing bounded feedback control law

Consider the system (7.32). Our purpose is to find a stabilizing state-feedback

$$u = k(x)$$

such that

$$k(x) \in \Omega_u. \quad (7.34)$$

Suppose that a CLF V is given for (7.32) such that

$$\inf_{u \in \Omega_u} \left\{ \frac{\partial V}{\partial x} f + \frac{\partial V}{\partial x} g u \right\} < 0 \quad \forall x \neq 0. \quad (7.35)$$

Condition (7.35) implies that for each nonzero state x one can diminish the value of V by applying *some* control in the set Ω_u . Note that the problem of determining a CLF for a constrained system is, by itself, much more difficult than in the unconstrained case. This topic has become the focus of research as of late where considerable progress has been made [FP98, MP96, JSK96]. As a standing assumption, we will assume both the existence and knowledge of a constrained CLF.

As in the unconstrained case, we first introduce a pointwise min-norm problem based on a control Lyapunov function approach. Later, this approach will be extended to a corresponding receding horizon control problem. As a direct extension of

the unconstrained pointwise min-norm problem in Chapter 5, consider the following constrained formulation:

$$\text{minimize} \quad u^T u \quad (7.36)$$

$$\text{subject to} \quad \frac{\partial V}{\partial x}[f + gu] \leq -\hat{\sigma}(x(t)) \quad (7.37)$$

$$u \in \Omega_u \quad (7.38)$$

with $\hat{\sigma}(x(t)) > 0$. In the unconstrained case, the design parameter σ can be chosen almost without restriction. It is easy to see that σ is only required to satisfy $\frac{\partial V}{\partial x} f \leq -\sigma$ whenever $\frac{\partial V}{\partial x} g = 0$ to be an admissible choice. However, the situation is now more complicated. Since the input is bounded, the stability constraint (7.37) may make the problem infeasible for an arbitrary choice of σ . To signify that σ must be chosen with this in mind, we denote it by $\hat{\sigma}$ in the constrained problem.

In order to avoid infeasibility, $\hat{\sigma}$ must be properly chosen. We propose to accomplish this by solving the following optimization problem in u and μ :

$$\text{minimize} \quad u^T u + \rho \mu^2 \quad (7.39)$$

$$\text{subject to} \quad \frac{\partial V}{\partial x}[f + gu] \leq -\sigma(x(t)) + \mu \quad (7.40)$$

$$\mu \geq 0 \quad (7.41)$$

$$-\sigma(x(t)) + \mu \leq 0 \quad (7.42)$$

$$u \in \Omega_u \quad (7.43)$$

and setting

$$\hat{\sigma}(x(t)) = \sigma(x(t)) - \mu \quad (7.44)$$

with $\rho > 0$ a design knob to be properly chosen.

Note that when Ω_u describes linear constraints on u (e.g., magnitude saturation constraints), the optimization in (7.39)-(7.43) is pointwise a quadratic program, which can be efficiently solved.

The problem (7.39)-(7.43) can be viewed as a pointwise min-norm problem in which the objective function contains the penalty term $\rho\mu^2$. In fact, the stability constraint (7.38) in the standard pointwise min-norm problem may lead to infeasibility, and thus this term is used to “soften” that constraint. Hence, one may view σ as being the desired parameter for the pointwise min-norm problem, but due to the constraint (7.43) it must be compromised to $\hat{\sigma} = \sigma - \mu$. The parameter ρ measures one’s averseness to deviations from the desired σ . For each arbitrarily large but finite ρ , the problem (7.39)-(7.43) is always feasible due to the condition (7.35).

Even when condition (7.35) is not known to be satisfied, i.e., one is not sure whether the CLF V is valid for the constrained system, the above scheme is a reasonable approach to the design of a constrained control law. By removing the constraint (7.42) (which ensures that \dot{V} is negative), and using a large value of ρ , the above problem will select u in the constraint set Ω_u that makes \dot{V} as negative as possible whenever σ is not feasible for the standard problem (7.36)–(7.38). In this sense, the control law will attempt to provide a stabilizing control law in the bounded set Ω_u if such a law is possible.

We have the following important connection between the pointwise min-norm problems (7.39)-(7.43) and (7.36)-(7.38):

Lemma 7.3.1 *Let (u^*, μ^*) be the optimal solution of the problem (7.39)-(7.43) for any given state $x(t)$, then u^* is also the optimal solution of (7.36)-(7.38) with $\hat{\sigma}(x(t)) = \sigma(x(t)) - \mu^*$.*

Proof: Set $\mu = \mu^*$. The problem (7.39)-(7.43) is then an optimization with respect to u only. With $\hat{\sigma}(x(t)) = \sigma(x(t)) - \mu^*$, the constraints (7.41) and (7.42) are ineffective. Therefore, this problem reduces to the pointwise min-norm problem (7.36)-(7.38) with the parameter $\hat{\sigma}(x(t))$ in (7.44). ■

The importance of Lemma 7.3.1 lies in the fact that we do not need to solve two optimization problems; i.e., first (7.39)-(7.43) to solve for the optimal μ allowing the

computation of $\hat{\sigma}$, and then (7.36)-(7.38) to obtain the pointwise min-norm control input u using $\hat{\sigma}$. This allows us to always refer to the pointwise min-norm problem (7.36)-(7.38), even though the problem (7.39)-(7.43) is effectively solved in order to obtain a feasible solution.

By determining a feasible $\hat{\sigma}$ for the constrained pointwise min-norm problem, this allows the approach in Chapter 5 to be used to extend the pointwise min-norm controller to its natural receding horizon formulation.

7.3.3 Receding horizon extensions

The ability to extend pointwise min-norm controllers to receding horizon schemes is useful in a number of respects. First of all, the advantages of on-line computation have already been well established in techniques such as model predictive control, especially in the handling of constraints. With the development of new CLF based techniques for dealing with constraints, it is important to recognize that these new approaches complement the existing receding horizon based approach.

For the constrained problem, the extension of pointwise min-norm controllers to a receding horizon scheme takes the following form. Let $u_{\hat{\sigma}}$ and $x_{\hat{\sigma}}$ denote the control and state trajectories, respectively, obtained by solving the pointwise min-norm problem (7.36)-(7.38) with parameter $\hat{\sigma}$. Consider the following receding horizon optimal control problem:

$$\text{minimize} \quad \int_t^{t+T} (q(x(\tau)) + u^T(\tau)u(\tau))d\tau \quad (7.45)$$

$$\text{subject to} \quad \dot{x} = f(x) + g(x)u \quad (7.46)$$

$$\frac{\partial V}{\partial x}(f(x(t)) + g(x(t))u(t)) \leq -\epsilon\hat{\sigma}(x(t)) \quad (7.47)$$

$$V(x(t+T)) \leq V(x_{\hat{\sigma}}(t+T)) \quad (7.48)$$

$$u(\cdot) \in \Omega_u \quad (7.49)$$

where ϵ is chosen as in the unconstrained case. This optimization is solved on-line

and implemented in a receding horizon fashion. Except for the constraint (7.49), this receding horizon scheme is identical to that presented in Chapter 5. As a consequence, it inherits the same stability and implementation properties. Again, the reader is referred to Chapter 5 for details.

7.3.4 Example

In this section a constrained nonlinear example is presented to demonstrate the approach. Consider the following dynamics:

$$\dot{x}_1 = -\frac{(1 + 0.1x_1^2)}{(1 + 0.1x_2^2)}x_2 + (1 + 2x_1^2)u_1 \quad (7.50)$$

$$\dot{x}_2 = \frac{20(1 + 0.1x_1^2)}{(1 + 0.1x_2^2)}x_1 + u_2 \quad (7.51)$$

with performance objective

$$\int_0^\infty (x_1^2 + 5x_2^2 + u_1^2 + u_2^2)dt$$

and input constraints

$$|u_1| \leq 1, \quad |u_2| \leq 1.$$

For this system, a constrained CLF is given by

$$V = 10x_1^2 + 0.5x_2^2.$$

We will test both a constrained pointwise min-norm controller, and its receding horizon extension. In the pointwise min-norm scheme, we selected the parameters $\rho = 1 \times 10^6$ and $\sigma(x)$ corresponding to Sontag's unconstrained formula, i.e.,

$$\sigma(x) = \sqrt{\left(\frac{\partial V}{\partial x} f\right)^2 + (x_1^2 + 5x_2^2) \left(\frac{\partial V}{\partial x} g g^T \frac{\partial V}{\partial x}\right)}$$

where f and g correspond to the system dynamics in (7.50)–(7.51). For the receding horizon problem, we chose a horizon of 0.15 seconds (with a sampling time of 0.05s) and an almost negligible value of $\epsilon = 1 \times 10^{-6}$.

Simulation results from the initial condition $[1, 1]$ are shown in the figures. A summary of the cost achieved by each controller is given in Table 7.2.

Table 7.2

Controller	Cost
Pointwise Min-Norm	27.1
RHC+CLF (T=0.15)	22.9

Table 7.2: Cost of pointwise min-norm vs. RHC+CLF controller from initial condition $[1, 1]$.

The control trajectories of the pointwise min-norm and RHC+CLF controllers are contrasted in Figure 7.5. Note that the RHC+CLF controller is saturated for a much shorter time than the pointwise min-norm controller, contributing to its smaller cost. For reference, the state trajectories of the pointwise min-norm and RHC+CLF controllers are shown in Figure 7.6.

The results of this example follow the general trend of those given previously. While a pointwise min-norm controller typically displays reasonable performance, especially under proper tuning of the parameters, the addition of a receding horizon often leads to significant improvements, even with the application of relatively short horizons.

7.4 Summary

In this chapter we began by presenting a straightforward extension of the RHC+CLF control scheme derived in Chapter 5 to the time-varying problem. This involved deriving a time-varying version of Sontag’s CLF based formula, and then extending it to a receding horizon scheme. Simulation results indicate that even though the CLF formula was derived from the HJB equation, it does not possess any inherent per-

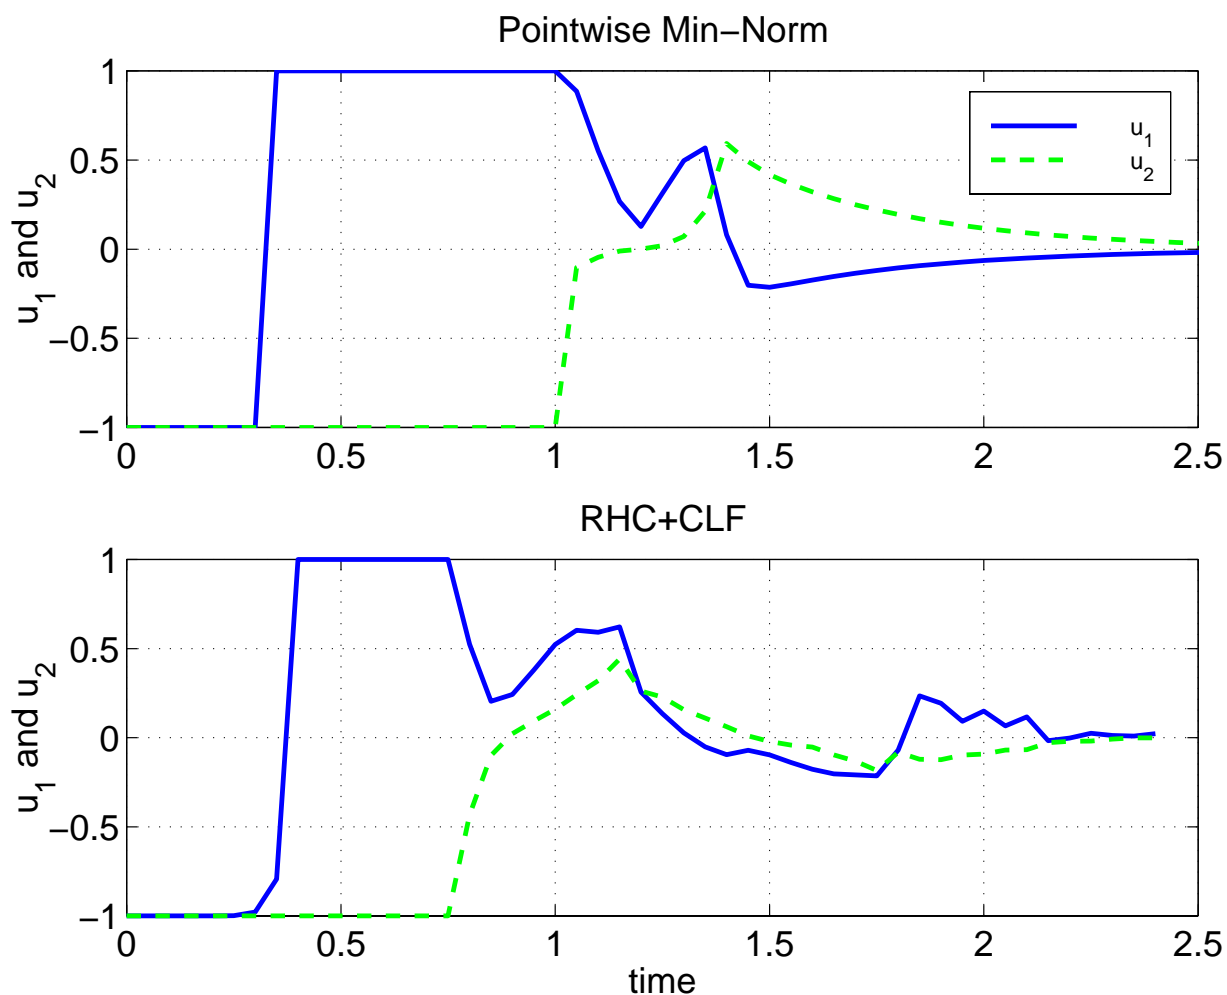


Figure 7.5: Comparison of control trajectories from the pointwise min-norm and RHC+CLF controller from the initial condition $[1, 1]$.

formance advantages over other schemes. On the other hand, the extended receding horizon schemes exhibit improved performance over other controllers, demonstrating the power of on-line computation coupled with the information provided by a CLF.

Next, we extended the framework introduced in Chapter 5 to include control constraints. This first involved the development of a constrained pointwise min-norm control scheme. This scheme is based on a modification of the unconstrained pointwise min-norm scheme, and as well as providing a controller for the constrained system, it generates the appropriate parameters required to establish a receding hori-

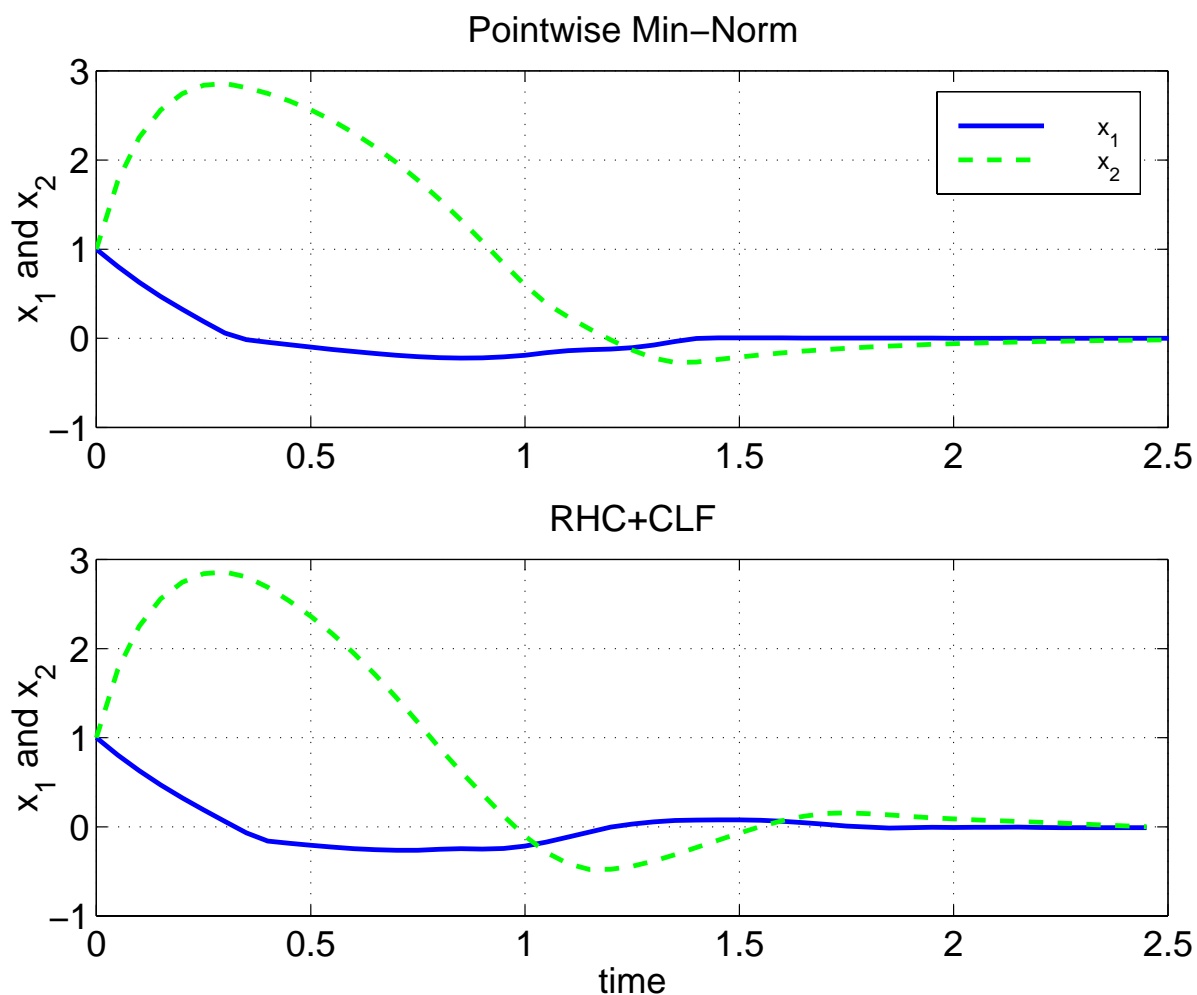


Figure 7.6: Comparison of state trajectories from the pointwise min-norm and RHC+CLF controller from the initial condition $[1, 1]$.

zon extension. As in the unconstrained case, in addition to providing a more flexible and implementable receding horizon scheme, it inherits the stability properties of the pointwise min-norm controller. Both of the frameworks presented in this chapter provide the foundation for new contributions in CLF and RHC theory to be effectively utilized.

Chapter 8 Conclusions

8.1 Summary of main results

We began this thesis with a review of the classical approaches to the problem of nonlinear optimal control: dynamic programming and calculus of variations. It was emphasized that these two solutions represent distinct points of view, and lead to Hamilton-Jacobi-Bellman partial differential equations, and the two point boundary value Euler-Lagrange ordinary differential equations, respectively. Furthermore, these two viewpoints acted as our guide for the rest of the thesis, providing a foundation for the interpretation of existing control approaches.

We focused on two popular approaches: those based on control Lyapunov functions, and the receding horizon methodology. While control Lyapunov functions can be thought of as generalizations of the Lyapunov methodology, the receding horizon methodology was made a practical reality by the computer revolution. In the context of optimal control, these techniques were shown to relate well to the two classical approaches to optimal control. Furthermore, this viewpoint was not only beneficial for understanding the contributions of existing techniques, but also led to the derivation of new control laws that exploit previously unrecognized connections.

First, we explicitly developed the connections between control Lyapunov function based schemes, specifically Sontag's formula and pointwise min-norm controllers, and the Hamilton-Jacobi-Bellman equation. For pointwise min-norm controllers, such relationships had been established previously, but our new variation of Sontag's formula was shown to couple even more tightly with the Hamilton-Jacobi-Bellman equation and furthermore be a special case of the pointwise min-norm formulation. This led to a deeper understanding of the pointwise min-norm controllers as well, and revealed both their strengths and weaknesses. In general, Sontag's formula and pointwise min-norm controllers rely on the information provided in the level curves of the CLF.

Despite the inverse optimality properties, if these level curves are far from those of the value function, these controllers are apt to lead to poor performance.

Next, receding horizon control was explored in the context of the Euler-Lagrange solution to trajectory optimizations. Again, this helped to explain both the advantages and disadvantages of receding horizon techniques, and it provided a clearer picture of the existing stabilizing formulations. In essence, receding horizon control exploits the computational simplicity of the Euler-Lagrange viewpoint to avoid the computational intractability associated with Hamilton-Jacobi-Bellman equations. The receding horizon methodology is merely a means to produce a state feedback control law from the repeated solution of trajectory optimizations.

Chapter 5 pieced the entire picture together by presenting a new framework in which optimal control and pointwise min-norm controllers could be interpreted as limits of a special receding horizon scheme. This even allowed us to present a receding horizon extension of Sontag's formula. In addition to leading to a clarification of the contributions of existing techniques, these new schemes demonstrated that both the Hamilton-Jacobi-Bellman and Euler-Lagrange points of view were complementary and could be combined in a beneficial manner. Theoretically, these schemes inherited both the stability properties of control Lyapunov functions and the performance advantages of on-line receding horizon style computation. Additionally, they were shown to possess desirable implementation properties, easing some of the difficulties associated with on-line intersample computation.

This new methodology was put into practice in Chapter 6 where it was applied to a simple model of a longitudinal flight control system. This example illustrated step-by-step the construction of control laws using a new two-stage design paradigm. The first stage involved the derivation of a CLF. It was shown how a number of standard and state-of-the-art techniques were natural candidates for this. The second stage required the selection of a CLF based control law. While the techniques used in the first stage offered their own implementations, it was recognized that Sontag's formula, pointwise min-norm controllers, and receding horizon extensions were also valid choices. Furthermore, simulations confirmed that this point of view was able

to utilize the contributions of existing approaches to produce improved control laws. Additionally, by having receding horizon implementations available, it naturally incorporated on-line computation, which will undoubtedly be a crucial advantage in the future. Finally, the framework was shown to extend to time-varying and input constrained systems, providing the foundation to include other advances in control theory as they become available.

8.2 Future research

Nonlinear optimal control is a vast subject and this thesis has only touched upon limited aspects of it. In particular, we have developed a framework to understand and utilize the contributions of existing techniques. While in one sense new control schemes were introduced, in another the ideas were already there, merely waiting to be formed into a coherent picture. This picture allowed us to leverage the contributions of existing techniques to design improved controllers. On the other hand, it also brought to the forefront those aspects of nonlinear optimal control that must be confronted in the future.

At its essence, the nonlinear control design process contains two stages: derivation of a CLF, and determination of a control law from the CLF. Let us outline some of the future challenges involved in each stage.

Derivation of a control Lyapunov function for nonlinear systems is a difficult task. No general procedures exist except for classes of systems that possess special structure. Nevertheless, exploiting special structure represents a promising approach to extending the ability to derive CLFs. In problems of trajectory tracking, for example, feedback linearizability can provide answers to both the problems of planning trajectories and determining CLFs around trajectories. Mechanical systems present another example of a class of systems that possess exploitable structure, and energy often provides a starting point in the derivation of a CLF.

When a CLF is desired for more than mere stability, to conform to constraints or robustness margins, the set of plants for which known techniques exist is limited

even further. Hence, these issues must be tackled if progress is to be made on the CLF based approach to nonlinear optimal control. The importance of this is widely recognized, and research on these subjects over the last couple of years has increased dramatically, especially in the area of constrained systems.

The second stage involves the determination of a control law. The use of more on-line computation in this step will undoubtedly occupy a large future area of research. Clearly, a deeper understanding of the properties of control schemes based solely upon on-line intersample optimization is needed. Currently, fundamental issues still remain to be sorted out, especially concerning constraints and robustness. While new analysis techniques have recently emerged for linear systems, extending results to the nonlinear problem will be challenging, but potentially extremely rewarding.

In the end, practical questions of implementation always have the final say. For control to take advantage of on-line computation, control designers must become more familiar with the computational tools available. This will involve an increased interaction with other communities, particularly computer science. As demonstrated in this thesis, an open mind to the offerings of different points of view can only serve to strengthen our ability to confront the problems of the future.

Bibliography

- [AF66] M. Athans and P. L. Falb. *Optimal Control*. McGraw–Hill, Inc., 1966.
- [AM89] B. D. O. Anderson and J. B. Moore. *Optimal Control; Linear Quadratic Methods*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1989.
- [Art83] Z. Artstein. Stabilization with relaxed controls. *Nonlinear Anal.*, 7(11):1163–1173, 1983.
- [AS97] F. Albertini and E. Sontag. Control Lyapunov Functions for time-varying set stabilization. In *Proceedings of 1997 European Control Conference*, Brussels, Belgium, 1997.
- [BBK96] B. Bodenheimer, P. Bendotti, and M. Kantner. Linear parameter-varying control of a ducted fan engine. *Int. Journal of Robust and Nonlinear Control*, 6:1023–1044, November 1996.
- [Bel52] R. Bellman. The theory of dynamic programming. In *Proc. Nat. Acad. Sci. USA*, number 38, 1952.
- [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, Pennsylvania, June 1994.
- [BGPK85] R. R. Bitmead, M. Gevers, I. R. Petersen, and R.J. Kaye. Monotonicity and stabilizability properties of solutions of the riccati difference equation: propositions, lemmas, theorems, fallacious conjectures and counterexamples. *System & Control Letters*, pages 309–315, 1985.
- [BH75] A. E. Bryson and Y. Ho. *Applied Optimal Control*. Hemisphere Publ. Corp., Washington D.C., 1975.

- [CA96] H. Chen and F. Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme for constrained nonlinear systems. In *Proceedings of the 16th Chinese Control Conference*, pages 309–316, Qindao, China, 1996.
- [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, Florida, 1996.
- [CR79] C. R. Cutler and B. L. Ramaker. Dynamic matrix control— A computer control algorithm. In *Proceedings of the AIChE 86th National Meeting*, Houston, Texas, April 1979.
- [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.
- [DGKF89] J. C. Doyle, K. Glover, P. Khargonekar, and B. Francis. State-space solutions to standard \mathcal{H}_2 and \mathcal{H}_∞ control problems. *IEEE Trans. Aut. Control*, 34(8):831–847, August 1989.
- [DOM96] S. L. De Oliveira and M. Morari. Contractive model predictive control for constrained nonlinear systems. *IEEE Trans. Aut. Control*, 1996. submitted.
- [DP87] M. A. Dahleh and J. B. Pearson. l_1 optimal feedback controllers for mimo discrete-time systems. *IEEE Trans. Aut. Control*, 32(4):314–322, 1987.
- [DPSN96] J. C. Doyle, J. Primbs, B. Shapiro, and V. Nevistić. Nonlinear games: examples and counterexamples. In *Proceedings of the 35th IEEE Conference on Decision and Control*, pages 3915–3920, Kobe, Japan, December 1996.

- [DRCN93] L. Del Re, J. Chapuis, and V. Nevistić. Predictive control with embedded feedback linearization for bilinear plants with input constraints. In *Proceedings of the 32nd IEEE Conference on Decision and Control*, pages 2984–2989, San Antonio, Texas, 1993.
- [FK95] R. Freeman and P. Kokotović. Optimal nonlinear controllers for feedback linearizable systems. In *Proceedings of the 1995 American Control Conference*, pages 2722–2726, Seattle, Washington, June 1995.
- [FK96a] R. Freeman and P. Kokotović. Inverse optimality in robust stabilization. *SIAM J. Contr. Optimiz.*, 34:1365–1391, July 1996.
- [FK96b] R. Freeman and P. Kokotović. *Robust Nonlinear Control Design*. Birkhäuser, Boston, 1996.
- [FP96] R. Freeman and J. Primbs. Control lyapunov functions: New ideas from an old source. In *Proceedings of the 35th IEEE Conference on Decision and Control*, pages 3926–3931, Kobe, Japan, December 1996.
- [FP98] R. Freeman and L. Praly. Integrator backstepping for bounded controls and control rates. *IEEE Trans. Aut. Control*, 43(2):258–262, 1998.
- [Gar84] C. E. García. Quadratic dynamic matrix control of nonlinear processes. An application to a batch reactor process. In *AIChE Annual Meeting*, San Francisco, California, 1984.
- [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.
- [Isi95] A. Isidori. *Nonlinear Control Systems*. Springer-Verlag, third edition, 1995.
- [JSK96] M. Jankovic, R. Sepulchre, and P. Kokotovic. Constructive lyapunov stabilization of nonlinear cascade systems. *IEEE Trans. Aut. Control*, 41(12):1723–1735, 1996.

- [Kal64] R. Kalman. When is a linear control system optimal. *Trans. ASME Ser. D: J. Basic Eng.*, 86:1–10, 1964.
- [KBBM95] M. Kantner, B. Bodenheimer, P. Bendotti, and R. M. Murray. An experimental comparison of controllers for a vectored thrust, ducted fan engine. In *Proceedings of the 1995 American Control Conference*, pages 1956–1961, Seattle, Washington, 1995.
- [KBK83] W. H. Kwon, A. N. Bruckstein, and T. Kailath. Stabilizing state feedback design via the moving horizon method. *Int. Journal of Control*, 37(3):631–643, 1983.
- [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.
- [Kha92] H. K. Khalil. *Nonlinear systems*. Macmillan Publishing Company, 1992.
- [Kha96] H. K. Khalil. *Nonlinear Systems*. Prentice Hall, Inc., Upper Saddle River, New Jersey, second edition, 1996.
- [KKK95] M. Krstić, I. Kanellakopoulos, and P. Kokotović. *Nonlinear and Adaptive Control Design*. John Wiley & Sons, New York, 1995.
- [KP77] W.H. Kwon and A.E. Pearson. A modified quadratic cost problem and feedback stabilization of a linear system. *IEEE Trans. Aut. Control*, 22(5):838–842, 1977.
- [KP78] W.H. Kwon and A.E. Pearson. On feedback stabilization of time-varying discrete linear systems. *IEEE Trans. Aut. Control*, 23:479–481, 1978.
- [KT61] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Proc. 2nd Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492, Berkeley, California, 1961.

- [Lin94] W. Lin. Input saturation and global stabilization by output feedback for affine systems. In *Proc. 33rd IEEE Conference on Decision and Control*, pages 1323–1328, Lake Buena Vista, Florida, December 1994.
- [LP44] A. I. Lur’e and V. N. Postnikov. On the theory of stability of control systems. *Applied Mathematics and Mechanics*, 8(3), 1944.
- [LS91] Y. Lin and E. Sontag. A universal formula for stabilization with bounded controls. *System & Control Letters*, 16:393–397, 1991.
- [MA73] P. J. Moylan and B. D. O. Anderson. Nonlinear regulator theory and an inverse optimal control problem. *IEEE Trans. Aut. Control*, 18:460–464, October 1973.
- [MM90] D. Q. Mayne and H. Michalska. Receding horizon control of nonlinear systems. *IEEE Trans. Aut. Control*, 35(7):814–824, 1990.
- [MM93] H. Michalska and D.Q. Mayne. Robust receding horizon control of constrained nonlinear systems. *IEEE Trans. Aut. Control*, 38(11):1623–1633, November 1993.
- [MP96] F. Mazenc and L. Praly. Adding integrations, saturated controls, and stabilization of feedforward systems. *IEEE Trans. Aut. Control*, 41(11):1559–1578, 1996.
- [MS97] L. Magni and R. Sepulchre. Stability margins of nonlinear receding-horizon control via inverse optimality. *System & 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, California 91125, 1998.
- [Nev97] V. Nevistić. Constrained control of nonlinear systems. Ph.D. thesis, ETH Zürich, 1997.

- [NM95] V. Nevistić and M. Morari. Constrained control of feedback-linearizable systems. In *Proceedings of the European Control Conference*, pages 1726–1731, Rome, Italy, 1995.
- [NMS98] G De Nicolao, L. Magni, and R. Scattolini. Stabilizing receding-horizon control of nonlinear time-varying systems. *IEEE Trans. Aut. Control*, 43(7):1030–1036, July 1998.
- [OF94a] T. Ohtsuka and H. Fujii. Nonlinear State Feedback Control of Space Vehicle Model Using a Real-Time Optimization Technique. In *AIAA Guidance Navigation and Control Conference*, Scottsdale, Arizona, 1994.
- [OF94b] T. Ohtsuka and H. Fujii. Stabilized Continuation Method for Solving Optimal Control Problems. *Journal of Guidance, Control, and Dynamics*, 17(5):950–957, 1994.
- [OF96] T. Ohtsuka and H. Fujii. Nonlinear Receding-Horizon State Estimation by Real-Time Optimization Technique. *Journal of Guidance, Control, and Dynamics*, 19(4):863–870, 1996.
- [Oht96] T. Ohtsuka. Time Variant Receding-Horizon Control of Nonlinear Systems. In *AIAA Guidance Navigation and Control Conference*, San Diego, California, 1996.
- [PN97a] J. A. Primbs and V Nevistić. Constrained finite receding horizon linear quadratic control. In *Proc. of the 36th IEEE Conference on Decision and Control*, pages 3196–3201, San Diego, California, 1997.
- [PN97b] J. A. Primbs and V Nevistić. MPC extensions to feedback linearizable systems. In *Proceedings of the American Control Conference*, Albuquerque, New Mexico, 1997.
- [Pon59] L. S. Pontryagin. Optimal control processes. *Uspehi Mat. Nauk*, 14:3–20, 1959. (in Russian).

- [RD83] S.L. Richter and R.A. DeCarlo. Continuation Methods: Theory and Applications. *IEEE Trans. Aut. Control*, AC-28(6):660–665, 1983.
- [Ric93] J. Richalet. Industrial applications of model based predictive control. *Automatica*, 29:1251–1274, 1993.
- [RM93] J. B. Rawlings and K. R. Muske. The stability of constrained receding horizon control. *IEEE Trans. Aut. Control*, 38(10):1512–1516, October 1993.
- [Roy88] H. L. Royden. *Real Analysis*. Macmillan, 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.
- [Sag68] A. P. Sage. *Optimum Systems Control*. Prentice Hall, Englewood Cliffs, New Jersey, 1968.
- [Son83] E.D. Sontag. A Lyapunov-like characterization of asymptotic controllability. *SIAM J. Contr. Optimiz.*, 21(3):462–471, 1983.
- [Son89] E.D. Sontag. A 'universal' construction of Artstein's theorem on nonlinear stabilization. *System & Control Letters*, 13(2):117–123, 1989.
- [Sus96] H. J. Sussmann. From the brachistochrone to the maximum principle and back. In *Proceedings of the 35th IEEE Conference on Decision and Control*, pages 1588–1593, Kobe, Japan, December 1996.
- [vNM96] M. van Nieuwstadt and R. M. Murray. Real time trajectory generation for differentially flat systems. In *Proceedings of the 1996 IFAC World Congress*, 1996.
- [WYPB96] Fen Wu, Hua Xin Yang, Andy Packard, and Greg Becker. Induced L_2 -Norm control for lpv system with bounded parameter variation. *Int. Journal of Nonlinear and Robust Control*, volume 6:983–998, 1996.