

On Receding Horizon Extensions and Control Lyapunov Functions

James A. Primbs*

Vesna Nevistić†

John C. Doyle*

*Control and Dynamical Systems 107-81
California Institute of Technology
Pasadena, California 91125
jprimbs, doyle@cds.caltech.edu

†Automatic Control Laboratory
Swiss Federal Inst. of Technology (ETH)
CH-8092 Zürich, Switzerland
vesna@aut.ee.ethz.ch

Abstract

Control Lyapunov functions (CLFs) are used in conjunction with receding horizon control (RHC) to develop a new class of control schemes. In the process, strong connections between the seemingly disparate approaches are revealed, leading to a unified picture that ties together the notions of pointwise min-norm, receding horizon, and optimal control. This framework is used to develop a control Lyapunov function based receding horizon scheme, of which a special case provides an appropriate extension of a variation on Sontag's formula. These schemes are shown to possess a number of desirable theoretical and implementation properties. An example is provided, demonstrating their application to a nonlinear control problem.

1. Introduction

The optimal control of nonlinear systems is one of the most challenging subjects in control theory. It is well known that the nonlinear optimal control problem can be reduced to the Hamilton-Jacobi-Bellman partial differential equation [2], but due to difficulties in its solution, this is not a practical approach. In fact, even the problem of stabilizing a nonlinear system remains a challenging task. Lyapunov theory, the most 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) [13, 3, 6, 5, 4]. The knowledge of such a function greatly facilitates the design of control schemes. 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.

In contrast to the emphasis on guaranteed stability that is the primary goal of CLFs, another class of nonlinear control schemes that go by the names receding horizon, moving horizon, or model predictive control place impor-

tance on optimal performance [10, 9, 11, 7, 8]. These techniques apply a receding horizon implementation in an attempt to approximately solve the optimal control problem through on-line computation. For systems under which on-line computation is feasible, receding horizon control (RHC) has proven quite successful. But guaranteed stability has remained a concern for some time.

Based on their underlying connection with the optimal control problem, we show that both CLF based methods and RHC 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 RHC. 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. In essence, it unites the best properties of CLFs and RHC.

2. Background

The reals will be denoted by \mathbb{R} , with \mathbb{R}_+ indicating the set of nonnegative real numbers. For notational convenience, the gradient of a function V with respect to x will be denoted by V_x (i.e. $V_x = \frac{\partial V}{\partial x}$).

We shall consider nonlinear systems of the form:

$$\dot{x} = f(x) + g(x)u \quad f(0) = 0 \quad (1)$$

with $x \in \mathbb{R}^n$ denoting the state, $u \in \mathbb{R}^m$ the control and $f(x), g(x) \in \mathcal{C}^1$.

2.1. Nonlinear Optimal Control

The standard nonlinear optimal control problem is formulated as follows:

Optimal Control Problem

$$\begin{aligned} \inf_{u(\cdot)} & \int_0^\infty (q(x) + u^T u) dt \\ \text{s.t.} & \quad \dot{x} = f(x) + g(x)u \end{aligned} \quad (2)$$

for $q(x) \in \mathcal{C}^1$, positive semi-definite.

A standard dynamic programming argument reduces the above optimal control problem to the Hamilton-Jacobi-Bellman partial differential equation (HJB) [2],

$$V_x^* f - \frac{1}{4}(V_x^* g g^T V_x^{*T}) + q(x) = 0 \quad (3)$$

where V^* is commonly referred to as the *value function* and can be thought of as the minimum cost to go, i.e.:

$$V^*(x(t)) = \inf_{u(\cdot)} \int_t^\infty (q(x(\tau)) + u^T(\tau)u(\tau))d\tau \quad (4)$$

If there exists a C^1 solution to the HJB equation (3), then the optimal control action is given by:

$$u^* = -\frac{1}{2}g^T V_x^{*T} \quad (5)$$

Unfortunately, the HJB pde (3) is extremely difficult to solve and in general precludes any hope of an exact solution to the optimal control problem.

A related optimal control problem is the finite horizon problem with terminal weight and specified initial condition:

$$\begin{aligned} \inf_{u(\cdot)} \quad & \int_0^T (q(x) + u^T u)dt + \varphi(x(T)) \\ \text{s.t.} \quad & \dot{x} = f(x) + g(x)u \\ & x(0) = x_0 \end{aligned}$$

This problem can be reduced to the *Euler-Lagrange* equations [2] which are much easier to solve than the HJB equation, but the solution is not equivalent to that of the infinite horizon problem unless the terminal weight is the value function V^* (which is found by solving the HJB equation). Furthermore, this problem is only solved for a *single initial condition* and produces an *open-loop* control law, in contrast to the *global, closed-loop* solution that the HJB approach provides.

The difficulties associated with an exact solution to the optimal control problem have motivated a number of alternative approaches, two of which are presented below.

2.2. Control Lyapunov Functions

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

$$\inf_u [V_x(x)f(x) + V_x(x)g(x)u] < 0 \quad (6)$$

for all $x \neq 0$ [1, 12, 13]. This definition is basically equivalent to saying that there always exists a choice of u such that the time derivative of V is decreasing. In this sense, a CLF will act as a Lyapunov function for the closed loop system.

A CLF may be used to design a stabilizing controller by posing the following optimization [3]:

Pointwise Min-Norm

$$\begin{aligned} \text{minimize} \quad & u^T u \\ \text{subject to} \quad & V_x(f + gu) \leq -\sigma(x) \end{aligned} \quad (7)$$

where $\sigma(x)$ is some positive definite function 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 σ at every point. When the following choice of $\sigma(x)$ is used:

$$\sigma_s = \sqrt{(V_x f)^2 + q(x)(V_x g g^T V_x^T)} \quad (9)$$

the solution reduces to a slight variation of Sontag's famous CLF formula [13]:

$$u_{\sigma_s} = \begin{cases} - \left[\frac{V_x f + \sqrt{(V_x f)^2 + q(x)(V_x g g^T V_x^T)}}{V_x g g^T V_x^T} \right] g^T V_x^T & V_x g \neq 0 \\ 0 & V_x g = 0 \end{cases} \quad (10)$$

In addition to possessing the continuity properties enjoyed by Sontag's formula, it also leads to an interpretation in terms of the optimal control problem. It can be thought of as using the gradient direction of the CLF, but scaled to satisfy the HJB equation. In particular, if a CLF has the same level curves as the value function, this formula recovers the optimal controller [6].

2.3. Receding Horizon Control

In receding horizon control (cf. [10, 9, 7, 11, 8]), the current control at state x and time t is obtained by determining on-line the optimal control $u^*(\cdot)$ over the interval $[t, t+T]$ with respect to the following objective:

Receding Horizon Control

$$\begin{aligned} \inf_{u(\cdot)} \quad & \int_t^{t+T} (q(x(\tau)) + u^T(\tau)u(\tau))d\tau + \varphi(x(t+T)) \\ \text{s.t.} \quad & \dot{x} = f(x) + g(x)u \end{aligned}$$

The above problem is solved in an open-loop/Euler-Lagrange fashion, exploiting the fact that a solution is only needed for the current state as an initial condition. The current control action is set to the initial optimizing solution $u^*(t)$. These computations are repeated at each time step, updating the initial condition with the current state and resolving the above optimization. In essence, this approach attempts to compute the value function and optimal control action on-line for those states encountered along its trajectory. In this way, receding horizon control approaches optimal control in a local fashion as opposed to the more global in nature control Lyapunov function methodologies. A substantial drawback is that the on-line computation required tends to be extensive and stability guarantees are often lacking.

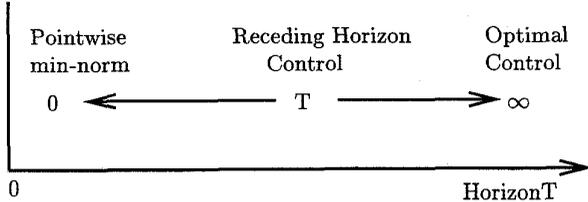


Figure 1: Unified framework

While appearing to represent different approaches to the optimal control problem, a deeper look at the actual form of the underlying optimization involved in each of our three problems (optimal control, pointwise min-norm, and receding horizon) leads to a striking connection.

3. Limits of receding horizon control

Consider the standard open loop optimization that is solved on-line at every time instance in receding horizon control, but without the terminal weight $\varphi(\cdot)$:

$$\int_t^{t+T} (q(x) + u^T u) d\tau \quad (11)$$

First, observe that as the horizon T tends to infinity, the objective of the optimal control problem (2) is recovered.

$$\int_t^\infty (q(x) + u^T u) d\tau \quad (12)$$

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:

$$\frac{1}{T} \int_t^{t+T} (q(x) + u^T u) d\tau \quad (13)$$

since dividing by T has no effect on the optimizing u . Now, letting $T \rightarrow 0$ yields: $q(x(t)) + u^T(t)u(t)$. Since $x(t)$ is known, there is no need to consider 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 (7). This indicates that we may heuristically view the pointwise min-norm problem as a receding horizon problem with a horizon length of zero (See Figure 1). Yet, it is clearly not the limit of the standard receding horizon formulation. This motivates one to consider if a scheme might exist that incorporates the best properties of CLFs and RHC, including:

- Stability for any horizon T .
- Pointwise min-norm controllers for $T = 0$.
- 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.

4. 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.

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 $\sigma(x)$. Consider the following receding horizon optimal control problem:

RHC+CLF

$$\text{minimize} \quad \int_t^{t+T} (q(x) + u^T u) dt \quad (14)$$

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

$$V_x(f + gu(t)) \leq -\epsilon\sigma(x) \quad (15)$$

$$V(x(t+T)) \leq V(x_\sigma(t+T)) \quad (16)$$

with $1 \geq \epsilon > 0$ (preferably ϵ is small). This optimization is solved on-line and implemented in receding horizon fashion. It is a standard receding horizon formulation with two CLF constraints. The first constraint (15) is a direct stability constraint in the spirit of that which appears in the pointwise min-norm formulation (8). Note that this constraint applies only to the *initial* control action. Since in the receding horizon approach only this first control action is implemented, V is required to be a Lyapunov function for the closed loop system. Ideally, ϵ is chosen small so that this constraint is not overly restrictive.

The second constraint (16) is oriented toward performance and replaces the terminal weight used in the standard RHC formulation. 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 . In terms of level curves this means that the final state of all potential sequences must lie inside the level curve of V that passes through $x_\sigma(t+T)$ (See Figure 2).

Below we explicitly list the following important properties possessed by the RHC+CLF scheme:

Theoretical Properties:

1. Stability is guaranteed for any horizon T .

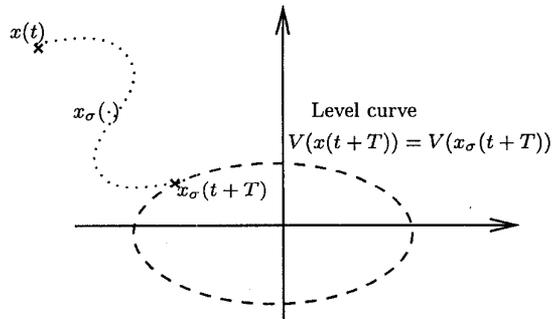


Figure 2: Performance constraint (16)

2. In the limit as the horizon goes to zero ($T \rightarrow 0$), the pointwise min-norm optimization problem is recovered.
3. If V is a Lyapunov function for the closed loop system under the optimal control, u^* , which always satisfies the constraint (15), then an infinite horizon length will always recover the optimal controller.

Additionally, for the parameter choice $\sigma(x) = \sigma_s(x)$ corresponding to Sontag's formula in the pointwise min-norm problem (see eqn. 9), the optimality property of Sontag's formula is preserved. That is, if the CLF V has the same level curves as the value function V^* , then the optimal controller is recovered for any horizon length.

In addition to the above theoretical properties, this receding horizon scheme possesses a number of desirable implementation properties.

Implementation Properties:

- Due to the Stability Constraint (15) the optimization may be preempted and/or the horizon may be varied on-line without loss of 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.

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

For the performance constraint (16), 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 initial feasible trajectories from which to begin the optimization.

Before presenting an example demonstrating 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. Example

Consider the 2d nonlinear oscillator:

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

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:

$$u^* = -3x_2$$

As a control Lyapunov function¹ we will use:

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

Both Sontag's formula (10) and receding horizon control alone are found to perform poorly on this system. Sontag's formula relies heavily on the "shape" of the level curves of the CLF (see [6]). Its poor performance in this example is attributable to the substantial difference between the level curves of the value function and the CLF. From the initial condition $[3, -2]$ Sontag's formula accumulates a cost of

¹This 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 the only quadratic function that locally agrees with our value function (which itself is not even a strict CLF for this system)

over 250 compared to the mere cost of 31.7 achieved by the optimal controller.

On the other hand, receding horizon control demonstrates an erratic behavior as a function of horizon length, wandering from instability to stability to instability once again as the horizon varies from 0.2 to longer than 1.0. This unpredictability is present even for receding horizon control applied to the linearization of this system, indicating that such behavior is inherent to receding horizon control and is not merely a nonlinear phenomena.

Building upon Sontag's formula, a horizon was introduced in accordance with the newly developed RHC+CLF scheme (as described in Section 4.). The erratic behavior demonstrated by the receding horizon controllers was tamed and drastically improved performance achieved for each of the tested horizons. Table 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.

Table 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 1: Cost from initial condition [3, -2]

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

6. Conclusions

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. This even provided a receding horizon "extension" of Sontag's formula and resulted in numerous theoretical and implementation advantages over present CLF and RHC methodologies. In

the end, we hope that these results will help to spawn new directions of research that reveal and exploit the synergistic relationships that exist between many of the current approaches to nonlinear control.

References

- [1] Z. Artstein. Stabilization with relaxed controls. *Nonlinear Anal.*, 7(11):1163–1173, 1983.
- [2] A. E. Bryson and Y. Ho. *Applied Optimal Control*. Hemisphere Publ. Corp., Washington D.C., 1975.
- [3] 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.
- [4] R. Freeman and P. Kokotović. Inverse optimality in robust stabilization. *SIAM J. Contr. Optimiz.*, 34:1365–1391, July 1996.
- [5] R. Freeman and P. Kokotović. *Robust Nonlinear Control Design*. Birkhäuser, Boston, 1996.
- [6] 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.
- [7] C. E. García, D. M. Prett, and M. Morari. Model predictive control: Theory and practice – A survey. *Automatica*, 25(3):335–348, May 1989.
- [8] 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.
- [9] W. H. Kwon, A. N. Bruckstein, and T. Kailath. Stabilizing state feedback design via the moving horizon method. *Int. J. Control*, 37(3):631–643, 1983.
- [10] 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.
- [11] D. Q. Mayne and H. Michalska. Receding horizon control of nonlinear systems. *IEEE Trans. Aut. Control*, 35(7):814–824, 1990.
- [12] E.D. Sontag. A lyapunov-like characterization of asymptotic controllability. *SIAM J. Contr. Optimiz.*, 21(3):462–471, 1983.
- [13] E.D. Sontag. A 'universal' construction of Artstein's theorem on nonlinear stabilization. *Syst. Contr. Lett.*, 13(2):117–123, 1989.