



12-4-2001

On the stability of unconstrained receding horizon control with a general terminal cost


Ali Jadbabaie

University of Pennsylvania, jadbabai@seas.upenn.edu

John Hauser

University of Colorado

Follow this and additional works at: http://repository.upenn.edu/ease_papers

 Part of the [Electrical and Computer Engineering Commons](#), and the [Systems Engineering Commons](#)

Recommended Citation

Ali Jadbabaie and John Hauser, "On the stability of unconstrained receding horizon control with a general terminal cost", . December 2001.

Copyright 2001 IEEE. Reprinted from *Proceedings of the 40th IEEE Conference on Decision and Control 2001*, Volume 5, pages 4826-4831.
Publisher URL: <http://dx.doi.org/10.1109/.2001.980971>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

NOTE: At the time of publication, author Ali Jadbabaie was affiliated with Yale University. Currently (March 2005), he is a faculty member in the Department of Electrical and Systems Engineering at the University of Pennsylvania.

On the stability of unconstrained receding horizon control with a general terminal cost

Abstract

This paper deals with unconstrained receding horizon control of nonlinear systems with a general, non-negative terminal cost. Earlier results have indicated that when the terminal cost is a suitable local control Lyapunov function, the receding horizon scheme is stabilizing for any horizon length. In a recent paper, the authors show that there always exist a *uniform* horizon length which guarantees stability of the receding horizon scheme over any sub-level set of the finite horizon cost when the terminal cost is identically zero. In this paper, we extend this result to the case where the terminal cost is a general non-negative function.

Keywords

nonlinear control systems, predictive control, stability, general terminal cost, model predictive control, nonlinear systems, optimal control, unconstrained receding horizon control, uniform horizon length

Disciplines

Electrical and Computer Engineering | Engineering | Systems Engineering

Comments

Copyright 2001 IEEE. Reprinted from *Proceedings of the 40th IEEE Conference on Decision and Control 2001*, Volume 5, pages 4826-4831.

Publisher URL: <http://dx.doi.org/10.1109/.2001.980971>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

NOTE: At the time of publication, author Ali Jadbabaie was affiliated with Yale University. Currently (March 2005), he is a faculty member in the Department of Electrical and Systems Engineering at the University of Pennsylvania.

On the stability of unconstrained receding horizon control with a general terminal cost*

Ali Jadbabaie[†] and John Hauser[‡]

ali.jadbabaie@yale.edu hauser@colorado.edu

[†]Electrical Engineering, PO Box 208257
Yale University, New Haven, CT 06520

[‡]Electrical and Computer Engineering
University of Colorado, Boulder, CO 80309-0425

Abstract

This paper deals with unconstrained receding horizon control of nonlinear systems with a general, non-negative terminal cost. Earlier results have indicated that when the terminal cost is a suitable local control Lyapunov function, the receding horizon scheme is stabilizing for any horizon length. In a recent paper, the authors show that there always exist a *uniform* horizon length which guarantees stability of the receding horizon scheme over any sub-level set of the finite horizon cost when the terminal cost is identically zero. In this paper, we extend this result to the case where the terminal cost is a general non-negative function.

Keywords: receding horizon control, nonlinear control design, model predictive control, optimal control.

Introduction

In receding horizon control, an open-loop finite horizon optimization is solved, generating an open-loop trajectory. The resulting open-loop control trajectory is applied to the system for a fraction of the horizon length. This process is then repeated, resulting in a sampled feedback law. Although receding horizon control has been successfully used in the process control industry, its application to stability critical areas has been more difficult. This is mainly due to two reasons. The first problem stems from the fact that the finite horizon optimizations have to be solved in a relatively short period of time. Second, it is well known and can be easily demonstrated using linear examples that a naive application of the receding horizon strategy can have disastrous effects and renders the system unstable. Several different approaches have been proposed to tackle this problem. (See [12] for an excellent review of this literature.)

These approaches used additional endpoint equality constraints [9] or endpoint inequality constraints [11, 13, 14, 15], in order to guarantee closed-loop stability. An alternative approach was presented in [16], where the authors use an *a priori* obtained control Lyapunov function (CLF) to guarantee stability. In this approach, stability is *enforced* by imposing additional constraints requiring that the derivative of the CLF along the optimal trajectory be negative definite and that the decrease in the value of the CLF at the end of the horizon be larger along the optimal trajectory than the CLF trajectory. This approach

assumes the existence of a global control Lyapunov function.

Recently, another method was developed by the authors in [8]. Similar to [16], this approach obtained stability guarantees through the use of an *a priori* obtained CLF. However, the CLF was used as terminal cost. By utilizing a suitable control Lyapunov function, the stability of the receding horizon scheme is guaranteed in a more efficient manner. In this setting, it was shown that the stability constraints are automatically satisfied, and can be eliminated from the optimization. This resulted in a speedup in the computations.

Furthermore, it was shown in [8] that the region of attraction of the unconstrained receding horizon control law is always larger than that of the CLF and it can be grown further to contain any compact subset of the infinite horizon region of attraction by a suitable choice of the horizon length. Other authors, such as [4, 18, 17] have shown (in the context of constrained linear systems) that for long enough horizon lengths, the terminal stability constraint is implicitly satisfied. The results of [4] are especially important since they also provide a bound on how large the horizon length should be.

In a recent paper [7], the authors show that there always exist a finite horizon length for which the unconstrained receding horizon scheme is stabilizing without the use of a terminal cost or constraint. While letting the horizon length be large is not practical due to the high computational cost, the result is of theoretical importance since they show a *uniform* horizon length exists.

The purpose of this paper is to develop similar results when a general terminal cost is employed. (For purely computational reasons, it is often beneficial to have a non-zero terminal cost rather than a zero one. See [6] for further details.) Specifically, it will be shown that there always exist a finite horizon length for which the unconstrained receding horizon scheme with a terminal cost which is an upper bound on the infinite horizon cost-to-go over a given level set of the infinite horizon cost, is stabilizing.

We will show that the case of a general terminal cost can be addressed in similar fashion by using the two cases mentioned above, namely, the zero and the upper bound cases.

This paper is organized as follows: In section 1., we formulate the problem and present our notation. Section 2. is a review of the unconstrained receding horizon scheme with a CLF terminal cost. A review of our results on the stability of the receding horizons scheme

*Research supported in part by DARPA and AFOSR. This research was performed when the first author was with Control and Dynamical Systems, Mail Code 107-81, Caltech, Pasadena, CA 91125.

with zero terminal cost [7] are presented in section 3.. In section 4., we discuss the case of a terminal cost which is an upper bound on the infinite horizon cost-to-go. In section 5., we present our main result, by combining the results of the two previous sections. Finally, our conclusions are presented in section 6..

1. Problem setting

The nonlinear system under consideration is

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

where the vector field $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is C^2 and possesses a linearly controllable critical point at the origin, e.g., $f(0, 0) = 0$ and $(A, B) := (D_1 f(0, 0), D_2 f(0, 0))$ is controllable.

Given an initial state x and a control trajectory $u(\cdot)$, the state trajectory $x^u(\cdot; x)$ is the (absolutely continuous) curve in \mathbb{R}^n satisfying

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

for $t \geq 0$.

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

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

for some $c_q > 0$ and $q(0, 0) = 0$. It follows that the quadratic approximation of q at the origin is positive definite, $D^2 q(0, 0) \geq c_q I > 0$.

To ensure that the solutions of the optimization problems of interest are nice, we impose some convexity conditions. We require the set $f(x, \mathbb{R}^m) \subset \mathbb{R}^n$ to be convex for each $x \in \mathbb{R}^n$. We also require that the pre-Hamiltonian function $u \mapsto p^T f(x, u) + q(x, u) =: K(x, u, p)$ be strictly convex for each $(x, p) \in \mathbb{R}^n \times \mathbb{R}^n$ and that there is a C^2 function $\bar{u}^* : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^m : (x, p) \mapsto \bar{u}^*(x, p)$ providing the global minimum of $K(x, u, p)$. The Hamiltonian $H(x, p) := K(x, \bar{u}^*(x, p), p)$ is then C^2 ensuring that extremal state, co-state, and control trajectories will all be somewhat smooth (C^1 or better). Note that these conditions are trivially satisfied for control affine f and quadratic q .

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

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

The optimal cost (from x) is given by

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

where the control functions $u(\cdot)$ belong to some reasonable class of admissible controls (e.g., piecewise continuous or

measurable). The function $x \mapsto J_\infty^*(x)$ is often called the *optimal value function* for the infinite horizon optimal control problem.

For the class of f and q considered, we know that $J_\infty^*(\cdot)$ is a positive definite C^2 function on a neighborhood of the origin. This follows from the *geometry* of the corresponding Hamiltonian system [19, 20, 6]. In particular, since $(x, p) = (0, 0)$ is a hyperbolic critical point of the C^1 Hamiltonian vector field $X_H(x, p) := (D_2 H(x, p), -D_1 H(x, p))^T$, the local properties of $J_\infty^*(\cdot)$ are determined by the linear-quadratic approximation to the problem and, moreover, $D^2 J_\infty^*(0) = P > 0$ where P is the stabilizing solution of the appropriate algebraic Riccati equation.

For practical purposes, we are interested in finite horizon approximations of the infinite horizon optimization problem. In particular, let $V(\cdot)$ be a nonnegative C^2 function with $V(0) = 0$ and define the finite horizon cost (from x using $u(\cdot)$) to be

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

and denote the optimal cost (from x) as

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

As in the infinite horizon case, one can show, by geometric means, that $J_T^*(\cdot)$ is locally smooth (C^2). Other properties will depend on the choice of V and T .

Let Γ^∞ denote the domain of $J_\infty^*(\cdot)$ (the subset of \mathbb{R}^n on which J_∞^* is finite). It is not too difficult to show that the cost functions $J_\infty^*(\cdot)$ and $J_T^*(\cdot)$, $T \geq 0$, are continuous functions on Γ^∞ using the same arguments as in proposition 3.1 of [1]. For simplicity, we will allow $J_\infty^*(\cdot)$ to take values in the extended real line so that, for instance, $J_\infty^*(x) = +\infty$ means that there is no control taking x to the origin.

We will assume that f and q are such that the minimum value of the cost functions $J_\infty^*(x)$, $J_T^*(x)$, $T \geq 0$, is attained for each (suitable) x . That is, given x and $T > 0$ (including $T = \infty$ when $x \in \Gamma^\infty$), there is a (C^1 in t) optimal trajectory $(x_T^*(t; x), u_T^*(t; x))$, $t \in [0, T]$, such that $J_T(x, u_T^*(\cdot; x)) = J_T^*(x)$. For instance, if f is such that its trajectories can be bounded on finite intervals as a function of its input size, e.g., there is a continuous function β such that $\|x^u(t; x_0)\| \leq \beta(\|x_0\|, \|u(\cdot)\|_{L_1[0, t]})$; then (together with the conditions above) there will be a minimizing control (cf. [10]). Many such conditions may be used to good effect, see [3] for a nearly exhaustive set of possibilities. In general, existence of minima can be guaranteed through the use of techniques from the direct methods of the calculus of variations—see [2] (and [5]) for an accessible introduction.

It is easy to see that $J_\infty^*(\cdot)$ is proper on its domain so that the sub-level sets

$$\Gamma_r^\infty := \{x \in \Gamma^\infty : J_\infty^*(x) \leq r^2\}$$

are compact and path connected and moreover $\Gamma^\infty = \bigcup_{r>0} \Gamma_r^\infty$. Note also that Γ^∞ may be a proper subset of \mathbb{R}^n since there may be states that cannot be driven to the origin. We use r^2 (rather than r) here to reflect the fact that our incremental cost is quadratically bounded from below. We refer to sub-level sets of $J_T^*(\cdot)$ and $V(\cdot)$ using $\Gamma_r^T :=$ path connected component of $\{x \in \Gamma^\infty : J_T^*(x) \leq r^2\}$ containing 0, and $\Omega_r :=$ path connected component of $\{x \in \mathbb{R}^n : V(x) \leq r^2\}$ containing 0.

2. Unconstrained receding horizon control with CLF terminal cost

Receding horizon control provides a practical strategy for the use of model information through on-line optimization. Every δ seconds, an optimal control problem is solved over a T second horizon, starting from the current state. The first δ seconds of the optimal control $u_T^*(\cdot; x(t))$ is then applied to the system, driving the system from $x(t)$ at current time t to $x_T^*(\delta, x(t))$ at the next sample time $t+\delta$. We denote this receding horizon scheme as $\mathcal{RH}(T, \delta)$.

In defining (unconstrained) finite horizon approximations to the infinite horizon problem, the key design parameters are the terminal cost function V and the horizon length T (and, perhaps also, the increment δ). What choices will result in success?

It is well known (and easily demonstrated with linear examples), that simple truncation of the integral (i.e., $V \equiv 0$) may have disastrous effects if $T > 0$ is too small. Indeed, although the resulting value function may be nicely behaved, the ‘‘optimal’’ receding horizon closed loop system can be unstable!

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

The following theorem states that the use of such a CLF guarantees exponential convergence of the receding horizon trajectories to the origin.

Theorem 1 [12, 8] *Let $T > 0$ be given, and consider the receding horizon scheme $\mathcal{RH}(T, \delta)$ for any $\delta \in (0, t]$, with a C^2 CLF terminal cost V satisfying $\min_u (\dot{V} + q)(x, u) \leq 0$ over the set Ω_r for any $r \leq r_v$. Then for each $x_0 \in \Gamma_r^T$, the resulting receding horizon trajectories go to zero exponentially.*

Theorem 1 shows that for *any* horizon length $T > 0$ and *any* sampling time $\delta \in (0, T]$, the receding horizon scheme is exponentially stabilizing over the set $\Gamma_{r_v}^T$. For a given T , the region of attraction can be enlarged by increasing r beyond r_v , as long as $x_T^*(T; x) \in \Omega_{r_v}$. An important feature of the above result is that there is no need for imposing additional stability constraints which make the optimizations harder to solve. Of course this method requires that a suitable CLF be generated off-line. If the resulting optimizations can be solved over longer horizons, it is desirable to obtain stability arguments that hold without the need of having a CLF as terminal cost. As we will show in the next section, there always exist a finite horizon length that would guarantee exponential stability of the receding horizon scheme with a zero terminal cost for fixed δ . Moreover, we will show that the same result holds when the terminal cost is merely an upper bound on the infinite horizon cost-to-go $J_\infty^*(x)$. These two cases can be thought of as limiting cases of a general terminal cost. We will show that the same result holds for a general positive terminal cost as long as it can be quadratically bounded over compact sets. The main tool in proving these results is to show that the finite horizon cost converges to the infinite horizon one, uniformly over Γ_r^∞ for any $r > 0$.

3. Receding horizon control with zero terminal cost

This problem was addressed fairly recently in the context of constrained discrete-time linear systems [17]. The case of continuous time nonlinear systems was treated in [7]. Unfortunately these results do not provide a bound on the horizon length, however, they guarantee that a finite and uniform horizon length always exists, although they might be too long for practical purposes. We present the main result of [7] for completeness:

Theorem 2 *Let $r > 0$ be given and suppose that $V(x) \equiv 0$. For each $\delta > 0$ there is a $T^* < \infty$ such that, for any $T \geq T^*$, the receding horizon scheme $\mathcal{RH}(T, \delta)$ is exponentially stabilizing. Moreover, the set $\Gamma_{r_1}^{T-\delta}$, with $\Gamma_{r_1}^{T-\delta} \subset \Gamma_r^\infty$, is contained in the region of attraction of $\mathcal{RH}(T, \delta)$.*

Proof: By the principle of optimality,

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

so that

$$\begin{aligned} J_{T-\delta}^*(x_T^*(\delta; x)) - J_{T-\delta}^*(x) &= J_T^*(x) - J_{T-\delta}^*(x) \\ &\quad - \int_0^\delta q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau \\ &\leq -J_{\delta,0}^*(x) + J_T^*(x) - J_{T-\delta}^*(x). \end{aligned}$$

Since $V(x) \equiv 0$, it is clear that $T_1 \leq T_2$ implies that $J_{T_1}^*(x) \leq J_{T_2}^*(x)$ for all x so that

$$J_{T-\delta}^*(x_T^*(\delta; x)) - J_{T-\delta}^*(x) \leq -J_{\delta,0}^*(x) + J_\infty^*(x) - J_{T-\delta}^*(x).$$

If we can show, for example, that there is a T^* such that $T \geq T^*$ yields

$$J_\infty^*(x) - J_{T-\delta}(x) \leq \frac{1}{2} J_{\delta,0}^*(x)$$

for $x \in \Gamma_r^\infty$, stability (in fact, exponential stability) over any sub-level set of $J_{T-\delta}^*(\cdot)$ contained in Γ_r^∞ will be assured. To that end, define, for $x \in \Gamma_r^\infty$,

$$\psi_T(x) := \begin{cases} \frac{J_\infty^*(x) - J_{T-\delta}^*(x)}{J_{\delta,0}^*(x)}, & x \neq 0 \\ \limsup_{x \rightarrow 0} \psi_T(x), & x = 0 \end{cases}$$

and note that $\psi_T(\cdot)$ is upper semicontinuous on Γ_r^∞ . This follows easily since $\psi_T(\cdot)$ is continuous at all $x \neq 0$ ($J_{\delta,0}^*(x) > 0$ for $x \neq 0$) and is finite at $x = 0$ with $\psi_T(0) = \max_{\|x\|=1} \frac{x^T(P_\infty - P_{T-\delta})x}{x^T P_\delta x}$ where $P_{T-\delta}$, P_δ , and P_∞ are the positive definite matrices defined above.

We see that $\{\psi_T(\cdot)\}_{T>0}$ is a monotonically decreasing family of upper semicontinuous functions defined over the compact set Γ_r^∞ . Hence, by Dini's Theorem, there is a $T^* < \infty$ such that $\psi_T(x) < \frac{1}{2}$ for all $x \in \Gamma_r^\infty$ and all $T \geq T^*$. The result follows since, for $r_1 > 0$ such that $\Gamma_{r_1}^{T-\delta} \subset \Gamma_r^\infty$, we have

$$J_{T-\delta}^*(x_T^*(\delta; x)) - J_{T-\delta}^*(x) \leq -\frac{1}{2} J_{\delta,0}^*(x)$$

for $x \in \Gamma_{r_1}^{T-\delta}$. \square

We see that when the calculation horizon is chosen to be sufficiently long, the trivial terminal cost $V(x) \equiv 0$ is fine. In a sense, if no off-line calculations are used to determine a suitable CLF, more online computations may be required (which might be impractical in some cases) to ensure closed loop stability of the receding horizon scheme. One might imagine that a suitably long horizon might also be adequate to ensure the stability of a receding horizon scheme when the dynamics and/or cost change in real-time such as when a fault occurs or a new objective is required.

4. Using an upper bound on the infinite horizon cost-to-go as a terminal cost

In the previous section (with $V(x) \equiv 0$), we exploited the fact that $J_T^*(x)$ increases monotonically with T to show that $J_{T-\delta}^*(\cdot)$, with T large, could be used as a Lyapunov function. A similar monotonicity property (actually reversed) is obtained when a CLF terminal cost providing an incremental upper bound on the infinite horizon cost-to-go is used [8]. In both of these cases monotonicity plays an important role in ensuring stability of the receding horizon scheme. Such a monotonicity result does not hold in the general case. Fortunately, uniform convergence of $J_T^*(\cdot)$ to $J_\infty^*(\cdot)$ on Γ_r^∞ , a key consequence of

monotonicity, is in fact sufficient for the task at hand. In this section, we take a different approach to show such uniform convergence when $V(\cdot)$ is merely an upper bound on $J_\infty^*(\cdot)$.

We begin by deriving a general upper bound of the difference between finite and infinite horizon costs.

Lemma 3 $J_T^*(x) - J_\infty^*(x) \leq V(x_\infty^*(T; x))$ for all $x \in \Gamma_\infty^\infty$.

Proof: The result follows easily by noting that

$$\begin{aligned} J_T^*(x) &\leq \int_0^T q(x_\infty^*(\tau; x), u_\infty^*(\tau; x)) d\tau + V(x_\infty^*(T; x)) \\ &\leq J_\infty^*(x) + V(x_\infty^*(T; x)). \end{aligned}$$

\square

In the case that the terminal cost is an upper bound on the infinite horizon cost-to-go, we get a lower bound on the difference between finite and infinite horizon costs. We will call a C^2 function $V(\cdot)$ with $V(0) = 0$ *strictly increasing* if it is proper and $\nabla V(x) \neq 0$ for all $x \in \mathbb{R}^n \setminus \{0\}$.

Lemma 4 Let $r > 0$ be given and suppose that the non-negative C^2 function $V(\cdot)$ is strictly increasing and such that $V(x) \geq J_\infty^*(x)$ for $x \in \Gamma_r^\infty$. Then, for any $T > 0$, $J_T^*(x) \geq J_\infty^*(x)$ for all $x \in \Gamma_r^\infty$.

Proof: Suppose, for the sake of contradiction, that this is not true. Then there is an $x_0 \in \Gamma_r^\infty$ such that $J_T^*(x_0) < J_\infty^*(x_0) =: r_0^2$. We have

$$\begin{aligned} \int_0^T q(x_T^*(\tau; x_0), u_T^*(\tau; x_0)) d\tau + V(x_T^*(T; x_0)) &< J_\infty^*(x_0) = r_0^2 \\ &\leq \int_0^T q(x_T^*(\tau; x_0), u_T^*(\tau; x_0)) d\tau + J_\infty^*(x_T^*(T; x_0)) \end{aligned}$$

so that $V(x_T^*(T; x_0)) < J_\infty^*(x_T^*(T; x_0))$ (with $J_\infty^*(x_T^*(T; x_0))$ possibly infinite) which implies that $x_T^*(T; x_0) \notin \Gamma_r^\infty$. On the other hand, $V(x_T^*(T; x_0)) < r_0^2 < r^2$, which clearly is a contradiction since $V(x) > r^2$ on $\mathbb{R}^n \setminus \Gamma_{r_1}^\infty$ (as $V(\cdot)$ is strictly increasing). \square

The above lemmas enable us to show that the difference between the finite and infinite horizon costs can be bounded according to

$$0 \leq J_T^*(x) - J_\infty^*(x) \leq V(x_\infty^*(T; x))$$

over the set Γ_r^∞ . If the mapping $x \mapsto V(x_\infty^*(T; x))$ was continuous and monotone (in fact, it's really a set-valued mapping since there may be multiple optimal trajectories!), we could apply Dini's theorem to complete our task. The next lemma shows that an appropriate continuity-like property actually holds.

Lemma 5 Let $r > 0$ and $\epsilon > 0$ be given. There is a $T^* < \infty$ such that for any $T \geq T^*$,

$$J_{\infty}^*(x_{\infty}^*(T; x)) \leq \epsilon J_{\infty}^*(x)$$

for all $x \in \Gamma_r^{\infty}$, where $x_{\infty}^*(\cdot; x)$ is any optimal trajectory.

Proof: Let $x \in \Gamma_r^{\infty}$ be arbitrary, and let $x_{\infty}^*(\cdot; x)$ be any optimal trajectory starting from x . Since the function $t \mapsto J_{\infty}^*(x_{\infty}^*(t; x))$ is monotonically decreasing (by the principle of optimality), once $x_{\infty}^*(\cdot; x)$ enters the set $\Gamma_{\epsilon J_{\infty}^*(x)}^{\infty}$, it remains there for all time. We will show that the first arrival time of $x_{\infty}^*(\cdot; x)$ to the set $\Gamma_{\epsilon J_{\infty}^*(x)}^{\infty}$ can be uniformly bounded over all $x \in \Gamma_r^{\infty}$ (and all optimal trajectories from such x). Indeed, let t_1 be the first arrival time of $x_{\infty}^*(\cdot; x)$ to the set $\Gamma_{\epsilon J_{\infty}^*(x)}^{\infty}$, so that $\|x_{\infty}^*(t; x)\|^2 \geq \frac{\epsilon}{b_r} J_{\infty}^*(x)$ for all $t \in [0, t_1]$ where b_r is such that $J_{\infty}^*(x) \leq b_r \|x\|^2$ for $x \in \Gamma_r^{\infty}$ (possible by compactness). It follows that

$$\begin{aligned} J_{\infty}^*(x) &\geq \int_0^{t_1} q(x_{\infty}^*(\tau; x), u_{\infty}^*(\tau; x)) d\tau \\ &\geq \int_0^{t_1} c_q \|x_{\infty}^*(\tau; x)\|^2 d\tau \geq t_1 \frac{\epsilon c_q}{b_r} J_{\infty}^*(x) \end{aligned}$$

which implies that $t_1 \leq \frac{b_r}{\epsilon c_q}$. The result follows by letting $T^* = \frac{b_r}{\epsilon c_q}$. \square

With these preliminary results in hand, we can show that upper bound type terminal costs also provide stabilization when the horizon is sufficiently long.

Theorem 6 Let $r > 0$ be given and suppose that the non-negative C^2 function $V(\cdot)$ is strictly increasing, locally quadratically bounded, and such that $V(x) \geq J_{\infty}^*(x)$ for $x \in \Gamma_r^{\infty}$. For each $\delta > 0$, there is a $T^* < \infty$ such that, for any $T \geq T^*$, the receding horizon scheme $\mathcal{RH}(T, \delta)$ is exponentially stabilizing. Moreover, the set $\Gamma_{r_1}^{T-\delta}$, with $\Gamma_{r_1}^{T-\delta} \subset \Gamma_r^{\infty}$, is contained in the region of attraction of $\mathcal{RH}(T, \delta)$.

Proof: As in the proof of Theorem 2, we will show that $J_{T-\delta}^*(\cdot)$ can be used as a Lyapunov function provided T is chosen sufficiently large. Once again, the fundamental relation is

$$J_{T-\delta}^*(x_{T-\delta}^*(\delta; x)) - J_{T-\delta}^*(x) \leq -J_{\delta,0}^*(x) + J_T^*(x) - J_{T-\delta}^*(x).$$

Our task is then to show that, over Γ_r^{∞} , the difference $J_T^*(x) - J_{T-\delta}^*(x)$ (with nonzero terminal costs) can be uniformly small relative to the (zero terminal cost) positive definite function $J_{\delta,0}^*(x)$.

Since $J_{\infty}^*(\cdot)$, $J_{\delta,0}^*(\cdot)$, and $V(\cdot)$ can each be quadratically bounded from above and below on the compact set Γ_r^{∞} , there exist $\epsilon_1, \epsilon_2 > 0$ such that $\epsilon_1 J_{\infty}^*(x) \leq \frac{1}{4} J_{\delta,0}^*(x)$ and $V(x) \leq \epsilon_2 J_{\infty}^*(x)$ for all $x \in \Gamma_r^{\infty}$. Now, using Lemma 5,

choose $T^* < 0$ so that $J_{\infty}^*(x_{\infty}^*(T; x)) \leq \epsilon_1/\epsilon_2 J_{\infty}^*(x)$ for all $T \geq T^* - \delta$ and all $x \in \Gamma_r^{\infty}$. Then, noting that

$$V(x_{\infty}^*(T; x)) \leq \epsilon_2 J_{\infty}^*(x_{\infty}^*(T; x)) \leq \epsilon_1 J_{\infty}^*(x) \leq \frac{1}{4} J_{\delta,0}^*(x),$$

and using the upper bound provided by Lemma 3, we see that

$$|J_T^*(x) - J_{T-\delta}^*(x)| \leq |J_T^*(x) - J_{\infty}^*(x)| + |J_{T-\delta}^*(x) - J_{\infty}^*(x)| \leq \frac{1}{2} J_{\delta,0}^*(x)$$

for all $T \geq T^*$ and all $x \in \Gamma_r^{\infty}$. Exponential stability of $\mathcal{RH}(T, \delta)$ over $\Gamma_{r_1}^{T-\delta}$ follows. \square

In what follows, by combining the results of this theorem together with Theorem 2, we will show that $\mathcal{RH}(T, \delta)$ with a general terminal cost is stable provided the horizon is sufficiently long.

5. Receding horizon control with a general terminal cost

We are now ready to present our main result.

Theorem 7 Let $r > 0$ be given and suppose that the non-negative C^2 function $V(\cdot)$ is locally quadratically bounded. For each $\delta > 0$, there is a $T^* < \infty$ such that, for any $T \geq T^*$, the receding horizon scheme $\mathcal{RH}(T, \delta)$ is exponentially stabilizing. Moreover, the set $\Gamma_{r_1}^{T-\delta}$, with $\Gamma_{r_1}^{T-\delta} \subset \Gamma_r^{\infty}$, is contained in the region of attraction of $\mathcal{RH}(T, \delta)$.

Proof: For $r > 0$, let $V_1(\cdot)$ be a locally quadratic, strictly increasing C^2 function that majorizes $V(\cdot)$ over \mathbb{R}^n and $J_{\infty}^*(\cdot)$ over Γ_r^{∞} and denote by $J_{T,1}^*(\cdot)$ the optimal cost with $V_1(\cdot)$ as terminal cost. It is then easy to show that

$$J_{T,0}^*(x) \leq J_T^*(x) \leq J_{T,1}^*(x)$$

and hence that

$$|J_T^*(x) - J_{\infty}^*(x)| \leq \max\{J_{\infty}^*(x) - J_{T,0}^*(x), J_{T,1}^*(x) - J_{\infty}^*(x)\}$$

for all $x \in \Gamma_r^{\infty}$ so that $J_T^*(\cdot)$ also converges uniformly to $J_{\infty}^*(\cdot)$ with respect to any locally quadratic positive definite function. The theorem follows directly using the results and techniques of Theorems 2 and 6. \square

6. Conclusion

The purpose of this paper was to demonstrate the stability of unconstrained nonlinear receding horizon control with a general terminal cost and without stability constraints. Previously, it was demonstrated that when the terminal cost is zero, Dini's theorem on uniform convergence of upper semicontinuous functions can be used to show that there exists a finite horizon length that guarantees stability of the receding horizon scheme for all points in an appropriate sub-level set of a finite horizon cost. This result was then extended to the case of a terminal

cost that is an upper bound on the infinite horizon cost-go. Finally, we showed that by combining these two results, the stability of the receding horizon scheme can be guaranteed when a general positive definite terminal cost is used.

References

- [1] M. Bardi and I. Capuzzo-Dolcetta. *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Birkhauser, Boston, 1997.
- [2] G. Buttazzo, G. Mariano, and S. Hildebrandt. *One-dimensional Variational Problems*. Oxford University Press, New York, 1998.
- [3] L. Cesari. *Optimization - Theory and Applications: Problems with Ordinary Differential Equations*. Springer-Verlag, New York, 1983.
- [4] D. Chmielewski and V. Manousiouthakis. On constrained infinite-time linear quadratic optimal control. *Systems and Control Letters*, 29:121–129, 1996.
- [5] B. Dacorogna. *Direct Methods in the Calculus of Variations*. Springer-Verlag, New York, 1989.
- [6] J. Hauser and H. Osinga. On the geometry of optimal control: the inverted pendulum example. In *American Control Conference*, 2001.
- [7] A. Jadbabaie, J. Primbs, and J. Hauser. Unconstrained receding horizon control with no terminal cost. In *American Control Conference*, 2001.
- [8] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 46, May 2001.
- [9] 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.
- [10] E. B. Lee and L. Markus. *Foundations of Optimal Control Theory*. John Wiley & Sons, New York, 1989.
- [11] L. Magni and R. Sepulchre. Stability margins of nonlinear receding horizon control via inverse optimality. *Systems and Control Letters*, 32:241–245, 1997.
- [12] D. Q. Mayne, J. B. Rawlings, C.V. Rao, and P.O.M. Sokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.
- [13] H. Michalska and D.Q. Mayne. Robust receding horizon control of constrained nonlinear systems. *IEEE Transactions on Automatic Control*, 38(11):1623–1633, November 1993.
- [14] G. De Nicolao, L. Magni, and R. Scattolini. Stabilizing receding-horizon control of nonlinear time-varying systems. *IEEE Transactions on Automatic Control*, 43(7):1030–1036, 1998.
- [15] T. Parisini and R Zoppoli. A receding horizon regulator for nonlinear systems and a neural approximation. *Automatica*, 31:1443–1451, 1995.
- [16] J. A. Primbs, V. Nevistić, and J. C. Doyle. A receding horizon generalization of pointwise min-norm controllers. *IEEE Transactions on Automatic Control*, 45:898–909, June 2000.
- [17] J.A. Primbs and V. Nevistic. Feasibility and stability of constrained finite receding horizon control. *Automatica*, 36(7):965–971, 2000.
- [18] P. Sokaert and J. Rawlings. Constrained linear quadratic regulation. *IEEE Transactions on Automatic Control*, 43:1163–1169, August 1999.
- [19] A. J. van der Schaft. On a state space approach to nonlinear H_∞ control. *Systems and Control Letters*, 116:1–8, 1991.
- [20] A. J. van der Schaft. *L₂-Gain and Passivity Techniques in Nonlinear Control*, volume 218 of *Lecture Notes in Control and Information Sciences*. Springer-Verlag, London, 1994.