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

Ali Jadbabaie[†] and John Hauser[‡]
alij@cds.caltech.edu hauser@colorado.edu

[†]Control and Dynamical Systems, Mail Code 107-81
California Institute of Technology, Pasadena, CA 91125

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

Abstract

We study the stability and region of attraction properties of a family of unconstrained receding horizon schemes for nonlinear systems. In a recent paper, using Dini's theorem on the uniform convergence of functions, we showed that there is always a *finite* horizon for which the corresponding receding horizon scheme is stabilizing *without* the use of a terminal cost or terminal constraints. In this paper, after showing that optimal infinite horizon trajectories possess a uniform convergence property, we show that exponential stability may also be obtained with a sufficient horizon when an upper bound on the infinite horizon cost is used as terminal cost. Combining these important cases together with a sandwiching argument, we are able to conclude that exponential stability is obtained for unconstrained receding horizon schemes with a general nonnegative terminal cost for sufficiently long horizons. Region of attraction estimates are also included in each of the results.

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

Introduction

In receding horizon control, a finite horizon optimal control problem is solved, generating an open-loop state-control trajectory. The resulting control trajectory is then applied to the system for a fraction of the horizon length. This process is then repeated, resulting in a sampled data feedback law. Although receding horizon control has been successfully used in the

*Research supported in part by DARPA and AFOSR.

process control industry, its application to fast, stability critical nonlinear systems has been more difficult. This is mainly due to two reasons. The first problem stems from the fact that the finite horizon optimizations must 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, often rendering a system unstable. Various approaches have been proposed to tackle this problem. See [17] for an excellent, up to date, review of this literature.

A number of approaches employ the use of terminal state equality [14] or inequality [18, 20, 5, 16, 19] constraints, often together with a terminal cost, to ensure closed loop stability. In [21], aspects of a stability guaranteeing global control Lyapunov function (CLF) were used, via state and control constraints, to develop a stabilizing receding horizon scheme with many of the nice characteristics of the CLF controller together with better cost performance. Unfortunately, a global control Lyapunov function is rarely available and often not possible.

Motivated by the difficulties involved in and cost of solving constrained optimal control problems, the authors developed an unconstrained receding horizon control strategy for the stabilization of nonlinear systems [11, 12, 13]. In this approach, closed loop stability is ensured through the use of a terminal cost consisting of a control Lyapunov function that is an incremental upper bound on the optimal cost to go. In the absence of explicit constraints, a dramatic speedup in computation is noted. Also, questions of existence and regularity of optimal solutions (very important for online optimization) can be dealt with in a rather straight forward manner.

Furthermore, it was shown in [12, 13] that region of attraction estimates of the unconstrained receding horizon control law are always larger than those of the CLF controller and can be grown to include any compact subset of the infinite horizon region of attraction by a suitable choice of the horizon length. Other authors, including [6, 24, 22] have shown (in the context of constrained linear systems) that, for sufficiently long horizons, the terminal stability constraints are implicitly satisfied. In a recent paper [22], it was shown that, in the case of constrained discrete-time linear systems, there always exists a finite horizon length for which the receding horizon scheme is stabilizing without the use of a terminal cost or constraint.

In a recent paper [10], we studied the stability and region of attraction properties of a family of unconstrained receding horizon schemes for nonlinear systems. Using Dini's theorem on the uniform convergence of functions, we showed that there is always a *finite* horizon for which the corresponding receding horizon scheme is stabilizing *without* the use of a terminal cost or terminal constraints. In this paper, after showing that optimal infinite horizon trajectories possess a uniform convergence property, we show that exponential stability may also be obtained with a sufficient horizon when an upper bound on the infinite horizon cost is used as terminal cost. Combining these important cases together with a sandwiching argument, we are able to conclude that exponential stability is obtained for unconstrained receding horizon schemes with a general nonnegative terminal cost for sufficiently long horizons. Region of attraction estimates are also included in each of the results.

This paper is organized as follows: After presenting the problem setting in 1, we review the unconstrained receding horizon control of nonlinear systems with stability resulting from the use of a CLF terminal cost. Our results on the stability of receding horizon schemes with zero terminal cost are reviewed and presented in section 3. In section 4, we discuss the case where the terminal cost 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.

1 Problem setting

The nonlinear system under consideration is

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

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

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

for some $c_q > 0$ and $q(0,0) = 0$. 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-Hamiltonion 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 $\bar{u}^* : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^m : (x,p) \mapsto \bar{u}^*(x,p)$ 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 [25, 26, 8]. 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 [2]. 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. [15]). Many such conditions may be used to good effect, see [4] for a nearly exhaustive set of possibilities. In general, the existence of minima can be guaranteed through the use of techniques from the direct methods of the calculus of variations—see [3] (and [7]) 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 $\Gamma^\infty = \bigcup_{r \geq 0} \Gamma_r^\infty$. 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

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

and

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

2 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 $V(\cdot)$ 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., $V(x) \equiv 0$) 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 $V(\cdot)$. 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 should 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 shows that the use of a particular type of CLF is in fact effective, providing rather strong and specific guarantees.

Theorem 1 [13] *Suppose that the terminal cost $V(\cdot)$ is a control Lyapunov function such that $\min_{u \in \mathbb{R}^m}(\dot{V} + q)(x, u) \leq 0$ for each $x \in \Omega_{r_v}$ for some $r_v > 0$. Then, for* every *$T > 0$ and $\delta \in (0, T]$, the receding horizon scheme $\mathcal{RH}(T, \delta)$ is exponentially stabilizing. For each $T > 0$, there is an $\bar{r}(T) \geq r_v$ such that $\Gamma^T_{\bar{r}(T)}$ is contained in the region of attraction of $\mathcal{RH}(T, \delta)$. Moreover, given any compact subset Λ of Γ^∞, there is a T^* such that $\Lambda \subset \Gamma^T_{\bar{r}(T)}$ for all $T \geq T^*$.*

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 $\Gamma^T_{r_v}$. For a given T, the region of attraction estimate is enlarged by increasing r beyond r_v to $\bar{r}(T)$ according to the requirement that $V(x_T^*(T; x)) \leq r_v^2$ on that set. An important feature of the above result is that, for operations with the set $\Gamma^T_{\bar{r}(T)}$, there is no need to impose stability ensuring constraints which would likely make the online optimizations more difficult and time consuming to solve. Of course, this method requires a suitable CLF. There are various techniques, requiring substantial offline computation, for the successful construction of such CLFs—see [9] for a detailed example using a quasi-LPV method.

Experience has shown that receding horizon strategies with terminal costs not satisfying the above condition are often effective provided that an optimization horizon of suitable length is used. It is therefore desirable to develop stability arguments that are valid for a more general class of terminal costs. As it was shown in [10] and reviewed in the next section, there is always a finite horizon length for which exponential stability of the receding horizon scheme with a zero terminal cost and fixed δ is guaranteed. Moreover, we will show that the same result holds when the terminal cost is a locally quadratic upper bound on the infinite horizon cost-to-go $J^*_\infty(\cdot)$. As these two cases are, in some sense, limiting cases of a general terminal cost, we will show that similar stability results hold in the general case. All of the results follow rather naturally once the uniform convergence (over compact sets) of the finite horizon costs to the infinite horizon cost is shown.

3 Receding horizon control with zero terminal cost

One would expect that as the horizon length grows, the effect of the terminal cost should diminish. Therefore it is reasonable to ask whether there is a *finite* horizon such that the receding horizon scheme would be stabilizing with a *zero* terminal cost, i.e., $V(x) \equiv 0$.

We know that, when the horizon is infinite, the minimum cost function $J^*_\infty(\cdot)$ qualifies as a Lyapunov function for proving the stability of corresponding optimal feedback system. Also, we know that, as $T \to \infty$, $J^*_T(\cdot) \to J^*_\infty(\cdot)$ in many ways (e.g., pointwise in x). An important question is whether there is a (sufficiently large, yet finite) horizon length T for which the minimum cost $J^*_T(\cdot)$ qualifies as a Lyapunov function for proving the stability of a corresponding receding horizon scheme, e.g., $\mathcal{RH}(T,\delta)$.

This question was answered fairly recently in the context of constrained discrete-time linear systems [22]. We showed in [10] that a similar result holds in the case of unconstrained nonlinear systems and zero terminal cost.

Recall that an extended real valued function $f(\cdot)$ is upper semicontinuous if $f^{-1}((-\infty,c)) := \{x \in \mathbb{R}^n : f(x) < c\}$ is open for each $c \in \mathbb{R}$. We will make use of the following well known result [23].

Theorem 2 (Dini) *Let $\{f_n\}$ be a sequence of upper semicontinuous, real-valued functions on a countably compact space X, and suppose that for each $x \in X$, the sequence $\{f_n(x)\}$ decreases monotonically to zero. Then the convergence is uniform.*

We begin with a rather simple result that will be used here and in the sequel. The proof is a simple exercise but is included for completeness.

Lemma 3 *For each $\delta > 0$, $J^*_{\delta,0}(\cdot)$ is continuous and positive definite on $\mathbb{R}^n$ and locally quadratic positive definite. That is, $J^*_{\delta,0}(x) > 0$ for all $x \in \mathbb{R}^n \setminus \{0\}$ and $J^*_{\delta,0}(x) \geq a\|x\|^2$ in a neighborhood of 0 for some $a > 0$. Moreover, for any $r > 0$, there is an $a > 0$ such that $J^*_{\delta,0}(x) \geq a\|x\|^2$ for all $x \in \Gamma_r^\infty$.*

The '0' in the subscript is used to indicate $J^*_{\delta,0}(x) = J^*_\delta(x)$ with *zero* terminal cost. This special notation is needed as this function will also be used in the discussion of receding horizon schemes with nonzero terminal cost.

Proof: Continuity of $J^*_{\delta,0}(\cdot)$ on $\mathbb{R}^n$ is easily shown using arguments of the sort used in proposition 3.1 of [2].

It is easy to show, e.g., by geometric methods [25, 26, 8], that $J^*_{\delta,0}(\cdot)$ is C^2 near 0 with

$$J^*_{\delta,0}(x) = \frac{1}{2} x^T P_\delta \, x + o(\|x\|^2)$$

where $P_\delta = P(-\delta)$ satisfies the Riccati equation

$$\begin{aligned}\dot{P}(t) + (A - BR^{-1}S^T)^T P(t) + P(t)(A - BR^{-1}S^T) \\ - P(t)BR^{-1}B^T P(t) + (Q - SR^{-1}S^T) = 0\end{aligned}$$

with $P(0) = 0$ where $Df(0,0) = \begin{bmatrix} A & B \end{bmatrix}$ is controllable and $D^2q(0,0) = \begin{bmatrix} Q & S \\ S^T & R \end{bmatrix} > \frac{c_q}{2} I > 0$. Clearly, P_δ is positive semi-definite since $\frac{1}{2}x^T P_\delta \, x$ is the optimal value of the corresponding linear quadratic optimal control problem. That P_δ is actually positive definite is easily shown by contradiction. Following [1], if there is an $x_0 \neq 0$ such that $x_0^T P_\delta x_0 = 0$ then, since the corresponding optimal control must be zero (as u is fully penalized), it must also be true that $e^{At}x_0 \equiv 0$ (as x is also fully penalized—an observability condition). Thus, $P_\delta > 0$ for each $\delta > 0$ and $J^*_{\delta,0}(\cdot)$ is locally quadratically positive definite. (One may also note the well known fact that $\delta_2 > \delta_1 > 0$ implies $P_{\delta_2} > P_{\delta_1} > 0$.)

Similarly, suppose that there is a nonzero x_0 such that $J^*_{\delta,0}(x_0) = 0$. Once again, since x is fully penalized, this would imply that the zero input nonlinear system trajectory beginning at x_0 would be identically zero, a clear contradiction.

The final claim follows easily from the continuity of $J^*_{\delta,0}(\cdot)$. $\square$

We have the following result (cf. [10]).

Theorem 4 *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 $\Gamma^{T-\delta}_{r_1}$, with $r_1 < r$ such that $\Gamma^{T-\delta}_{r_1} \subset \Gamma^\infty_r$, 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

$$\begin{aligned} J^*_{T-\delta}(x^*_T(\delta; x)) - J^*_{T-\delta}(x) &= J^*_T(x) - J^*_{T-\delta}(x) - \textstyle\int_0^\delta q(x^*_T(\tau; x), u^*_T(\tau; x))\, d\tau \\ &\leq -J^*_{\delta,0}(x) + J^*_T(x) - J^*_{T-\delta}(x)\ . \end{aligned}$$

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 $x \in \Gamma^\infty_r$, stability (and, in fact, exponential stability) over any sublevel set of $J^*_{T-\delta}(\cdot)$ contained in Γ^∞_r will be assured. To that end, define, for $x \in \Gamma^\infty_r$,

$$\psi_T(x) := \begin{cases} \dfrac{J^*_\infty(x) - J^*_{T-\delta}(x)}{J_{\delta,0}(x)}\ , & x \neq 0 \\ \limsup\limits_{x\to 0} \psi_T(x)\ , & x = 0 \end{cases}$$

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 as 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 $x \in \Gamma_{r_1}^{T-\delta}$. □

We see that when the optimization horizon is chosen to be sufficiently long, the trivial terminal cost $V(x) \equiv 0$ is fine. In a sense, if no offline calculations are used to determine a suitable CLF, more online computations may be required 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 [12, 13]. In both of these cases monotonicity plays an important role in the arguments that ensure 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 5 *$J_T^*(x) - J_\infty^*(x) \leq V(x_\infty^*(T; x))$ for all $T > 0$ and $x \in \Gamma^\infty$.*

Proof: The result follows easily by noting that

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

$\square$

In the case that the terminal cost is an upper bound on the infinite horizon cost-to-go, we can also get a lower bound on the difference between finite and infinite horizon costs.

We call a continuous function $W(\cdot)$ *strictly increasing* if it is proper and its sublevel sets are strictly increasing with respect to set inclusion, that is, $W^{-1}((-\infty, w_1]) \subset W^{-1}((-\infty, w_2)) \subset W^{-1}((-\infty, w_2])$ for all $w_1 < w_2$. Examples of strictly increasing functions include $J_\infty^*(\cdot)$ and differentiable proper functions $V(\cdot)$, $V(0) = 0$, with $\nabla V(x) \neq 0$ for all $x \in \mathbb{R}^n \setminus \{0\}$. Much like class $\mathcal{K}$ functions, strictly increasing functions provide a measure of the distance of a point x from the global minimum of the function, often the origin.

Lemma 6 *Let $r > 0$ be given and suppose that the nonnegative 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

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

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

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 out task.

It is clear that each infinite horizon trajectory must converge to the origin. The following result shows that the T parametrized family of set valued maps $x \mapsto x^*_\infty(T;x)$ (abusing notation) converges uniformly on compact subsets of Γ^∞ with respect to the strictly increasing function $J^*_\infty(\cdot)$. We will thus obtain the desired *uniform convergence* of, for example, the T parametrized family of functions $x \mapsto \sup_{\text{optimal } x^*_\infty(\cdot;x)} V(x^*_\infty(T;x))$.

Proposition 7 *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 $x \in \Gamma^\infty_r$, where $x^*_\infty(\cdot;x)$ is any *optimal trajectory.*

Proof: Let $x \in \Gamma^\infty_r$ 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 $\Gamma^\infty_{\epsilon J^*_\infty(x)}$, it remains there for all time. We will show that the first arrival time of $x^*_\infty(\cdot;x)$ to the set $\Gamma^\infty_{\epsilon J^*_\infty(x)}$ can be uniformly bounded over all $x \in \Gamma^\infty_r$ (and all optimal trajectories from such x). Indeed, let t_1 be the first arrival time of $x^*_\infty(\cdot;x)$ to the set $\Gamma^\infty_{\epsilon J^*_\infty(x)}$, 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 $x \in \Gamma^\infty_r$ (possible by compactness). It follows that

$$\begin{aligned} J^*_\infty(x) \;&\geq\; \int_0^{t_1} q(x^*_\infty(\tau;x), u^*_\infty(\tau;x))\; d\tau \\ &\geq\; \int_0^{t_1} c_q \|x^*_\infty(\tau;x)\|^2 \; d\tau \;\geq\; t_1 \tfrac{\epsilon c_q}{b_r} J^*_\infty(x) \end{aligned}$$

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

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

Theorem 8 *Let $r > 0$ be given and suppose that the nonnegative C^2 function $V(\cdot)$ is strictly increasing, locally quadratically bounded, and such that $V(x) \geq J^*_\infty(x)$ for $x \in \Gamma^\infty_r$. 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 $\Gamma^{T-\delta}_{r_1}$, $r_1 \geq r$ with $\Gamma^{T-\delta}_{r_1} \subset \Gamma^\infty_r$, is contained in the region of attraction of $\mathcal{RH}(T,\delta)$.*

Proof: As in the proof of theorem 4, 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

$$J^*_{T-\delta}(x^*_T(\delta;x)) - J^*_{T-\delta}(x) \leq -J^*_{\delta,0}(x) + J^*_T(x) - J^*_{T-\delta}(x)\ .$$

Our task is then to show that, over Γ_r^∞, the difference $J_T^*(x) - J_{T-\delta}^*(x)$ (with nonzero terminal cost) can be made 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 proposition 7, choose $T_1 < \infty$ so that $J_\infty^*(x_\infty^*(T;x)) \leq \epsilon_1/\epsilon_2\, J_\infty^*(x)$ for all $T \geq T_1$ 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 5, 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 $T \geq T^* := T_1 + \delta$ 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 4, 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 9 *Let $r > 0$ be given and suppose that the nonnegative C^2 terminal cost 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 4 and 8. $\square$

In each of the above theorems, the region of attraction is estimated by a set of the form $\Gamma_{r_1}^{T-\delta}$. Intuitively, we expect that this set can be made as large as we like by increasing the computation horizon T. Indeed, suppose that we would like the region of attraction to include the compact set $\Gamma_{r_2}^\infty$ (or any compact subset of Γ^∞). By the uniform convergence of $J_{T,0}^*(\cdot)$ and $J_{T,1}^*(\cdot)$ (hence $J_T^*(\cdot)$) to $J_\infty^*(\cdot)$, it is clear that, given $r > r_1 > r_2$, there is a $T_1 < \infty$ such that

$$\Gamma_{r_2}^\infty \subset \Gamma_{r_1}^{T,1} \subset \Gamma_{r_1}^\infty \subset \Gamma_{r_1}^{T,0} \subset \Gamma_r^\infty$$

for all $T \geq T_1$. Since $\Gamma_{r_1}^{T,1} \subset \Gamma_{r_1}^T \subset \Gamma_{r_1}^{T,0}$ for all $T > 0$, it is clear that the region of attraction of the general terminal cost receding horizon scheme can be made to include any compact subset of the infinite horizon region of attraction.

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. First, 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 to 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] Brian D.O. Anderson and John B. Moore. *Optimal Control: Linear Quadratic Methods*. Prentice-Hall, 1990.

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

[3] G. Buttazzo, G. Mariano, and S. Hildebrandt. *One-dimensional Variational Problems.* Oxford University Press, New York, 1998.

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

[5] H. Chen and F. Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. In *European Control Conference*, 1997.

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

[7] B. Dacorogna. *Direct Methods in the Calculus of Variations.* Springer-Verlag, New York, 1989.

[8] J. Hauser and H. Osinga. On the geometry of optimal control: the inverted pendulum example. In *American Control Conference*, 2001.

[9] A. Jadbabaie and J. Hauser. Control of the Caltech ducted fan in forward flight: A receding horizon—lpv approach. In *American Control Conference*, 2001.

[10] A. Jadbabaie, J. Primbs, and J. Hauser. Unconstrained receding horizon control with no terminal cost. In *American Control Conference*, 2001.

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

[12] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control: Stability and region of attraction results. In *IEEE Conference on Decision and Control*, 1999.

[13] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 46, May 2001.

[14] 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.

[15] E. B. Lee and L. Markus. *Foundations of Optimal Control Theory.* John Wiley & Sons, New York, 1989.

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

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

[18] 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.

[19] 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.

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

[21] 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.

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

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

[24] P. Scokaert and J. Rawlings. Constrained linear quadratic regulation. *IEEE Transactions on Automatic Control*, 43:1163–1169, August 1999.

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

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