Unconstrained receding horizon control with no terminal cost*

Ali Jadbabaie†, James Primbs†, and John Hauser‡

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

‡Electrical and Computer Engineering, and Aerospace Engineering
University of Colorado, Boulder, CO 80309-0425
{alij,jprimbs,hauser}@cds.caltech.edu

Abstract

In this paper, we discuss a stabilizing receding horizon scheme for unconstrained nonlinear systems. Using Dini's theorem on the uniform convergence of functions, we show that there always exist a *finite* horizon length for which the corresponding receding horizon scheme is stabilizing *without* using terminal costs and/or constraints.

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

1. Introduction

In receding horizon control, an open-loop finite horizon optimization is solved. 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 [10] or endpoint inequality constraints [4, 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, 9]. 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 dramatic speedup in the computations.

Furthermore, it was shown in [9] 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 [5, 19] have shown (in the context of constrained linear systems) that for long enough horizon lengths, the terminal stability constraint is implicitly satisfied.

In a recent paper [17], the authors show that in the case of constrained discrete-time linear systems, there always exist a finite horizon length for which

*Research supported in part by AFOSR and DARPA.

the receding horizon scheme is stabilizing without the use of a terminal cost or constraint.

The purpose of this paper is to develop similar results for unconstrained, continuous-time nonlinear systems. Specifically, it will be shown that there always exist a finite horizon length for which the unconstrained receding horizon scheme with no terminal cost is stabilizing.

This paper is organized as follows: In section 2., we formulate the problem and present our notation. Section 3. is a review of the unconstrained receding horizon scheme with a CLF terminal cost [9, 7]. Our main results are presented in section 4.. Finally, our conclusions are presented in section 5..

2. Problem statement

The nonlinear system under consideration is

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

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

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

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

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

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

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

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

for some $c_q > 0$ and $q(0,0) = 0$. We further require that the function $u \mapsto q(x, u)$ be convex for each $x \in \mathbb{R}^n$. These conditions imply that the quadratic approximation of q at the origin is positive definite, $D^2 q(0,0) \geq c_q I > 0$.

We will require that there is a C^2 function $\bar{u}^* : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^m : (x,p) \mapsto \bar{u}^*(x,p)$ providing a global minimum of the pre-Hamiltonian $K(x,p,u) := p^T f(x,u) + q(x,u)$ so that the Hamiltonian $H(x,p) := K(x,p,\bar{u}^*(x,p))$ is C^2. Note that this condition is trivially satisfied for control affine f and quadratic q (for then $u \mapsto K(x,p,u)$ is strictly convex and *real analytic*).

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

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

The optimal cost (from x) is given by

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

where the control functions $u(\cdot)$ belong to some reasonable class of admissible controls (e.g., piecewise continuous). The function $x \mapsto J_\infty^*(x)$ is often called the *optimal value function* for the infinite horizon optimal control problem. For the class of f and q considered, we know that J_∞^* is a positive definite C^2 function on a neighborhood of the origin. This follows from the *geometry* of the corresponding Hamiltonian system [20, 21]. In particular, since $(x,p) = (0,0)$ is a hyperbolic critical point of the Hamiltonian vector field $X_H(x,p) := (D_2 H(x,p), -D_1 H(x,p))^T$, the local properties of J_∞^* 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 approximating the infinite horizon optimization problem with one over a finite horizon. In particular, let V be a nonnegative C^2 function and define the finite horizon cost (from x using $u(\cdot)$) to be

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

and denote the optimal cost (from x) as

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

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

Let Γ^∞ denote the domain of J_∞^* (the subset of $\mathbb{R}^n$ on which J_∞^* is finite). It is not too difficult to show that the cost functions J_∞^* and J_T^*, $T \geq 0$ are continuous functions on Γ_∞ using the same arguments as in proposition 3.1 of [1]. We make the following assumption.

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

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

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

such that

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

Continuity of $u_T^*(.;x)$ follows directly from Pontryagin's Maximum Principle.

One can turn (SA) into a proposition. This involves the use of techniques from regularity theory and the direct methods of the calculus of variations, (see [3] and [2]).

It is easy to see that J_∞^* 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^* and V using $\Gamma_r^T :=$ path connected component of

$$\{x \in \Gamma^\infty : J_T^*(x) \leq r^2\} \text{ containing } 0,$$

and $\Omega_r :=$ path connected component of

$$\{x \in \mathbb{R}^n : V(x) \leq r^2\} \text{ containing } 0.$$

3. Unconstrained receding horizon control with CLF terminal cost

A practical receding horizon strategy is based on solving the finite horizon optimization every $\delta > 0$ seconds and using the control trajectory $u_T^*(\tau; x(t))$, $\tau \in [0, \delta]$, to drive the system from $x(t)$ at time t to $x_T^*(\delta, x(t))$ at time $t+\delta$. We will 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. As we show in the following theorem, the use of such a CLF guarantees exponential convergence of the receding horizon trajectories to the origin. (cf. [9]).

Theorem 1 *Let $T > 0$ and consider the use of a receding horizon scheme $\mathcal{RH}(T,\delta)$ with each $\delta \in (0,T]$. Then, for each $x_0 \in \Gamma_{\bar{r}(T)}^T$, the resulting trajectory converges to the origin exponentially fast.*

Proof: Given $T > 0$, set c_{T} and m_∞ such that $J_T^*(x) \leq c_{\mathrm{T}} ||x||^2 \quad \forall x \in \Gamma_{\bar{r}(T)}^T$ and $J_\infty^*(x) \geq m_\infty ||x||^2 \quad \forall x \in \Gamma_{\bar{r}(T)}^T$.

Let $(x_{\mathrm{RH}}(t), u_{\mathrm{RH}}(t))$, $t \geq 0$, be the receding horizon trajectory originating from an arbitrary $x_0 \in \Gamma_{\bar{r}(T)}^T$ and define

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

(The control trajectory $u_{\mathrm{RH}}(\cdot)$ is piecewise continuous since each optimal control trajectory $u_T^*(t;x)$,

$t \geq 0$, is continuous as a function of time.) As shown in Proposition 9 of [9], $W(t;x_0,u_{\rm RH}(\cdot)) \leq J^*_T(x_{\rm RH}(t))$. Also, since $\mathcal{RH}(T,\delta)$ is a suboptimal strategy, $J^*_\infty(x_{\rm RH}(t)) \leq W(t;x_0,u_{\rm RH}(\cdot))$. Now, we have

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

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

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

□

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

4. Unconstrained receding horizon control with no terminal cost

One would expect that as the horizon length grows, the effect of the terminal cost should diminish. Therefore it would be reasonable to ask whether there exists a *finite* horizon length such that the receding horizon scheme would remain stabilizing *without* using a CLF as a terminal cost.

We know that in the infinite horizon case, the minimum cost qualifies as a Lyapunov function. Also, we know that as $T \to \infty$, $J^*_T \to J^*_\infty$. The remaining question is whether there is a (large enough) finite horizon length for which the minimum cost qualifies as a suitable Lyapunov function. This question was answered fairly recently in the context of constrained discrete-time linear systems [17]. We will show that a similar result holds in the case of unconstrained nonlinear systems. Before proving this result, we will need the following definition and proposition [18]:

Definition 4..1 *An extended real-valued function f is called upper semicontinuous at the point y if $f(y) \neq +\infty$ and $f(y) \geq \limsup_{x\to y} f(x)$. We say that f is upper semicontinuous on a set if it is upper semicontinuous at each point of that set.*

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

Now we are ready to present the following theorem:

Theorem 3 *Consider the receding horizon scheme $\mathcal{RH}(T,\delta)$ with zero terminal cost. Then there exists a T^* such that for $T \geq T^*$, $\mathcal{RH}(T,\delta)$ is stabilizing with $\Gamma^{T-\delta}_{r_T}$ being the region of attraction, where for any given $r > 0$, $r_T > 0$ is the largest value such that $\Gamma^{T-\delta}_{r_T} \subset \Gamma^\infty_r$.*

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

$$J^*_T(x) - J^*_{T-\delta}(x^*_T(\delta;x)) = \int_0^\delta q(x^*_T(\tau;x),u^*_T(\tau;x))d\tau \quad \forall\, x \in \Gamma^\infty_r \tag{1}$$

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

Adding $J^*_{T-\delta}(x)$ to both sides of (1), and taking $J^*_T(x)$ to the other side, we obtain the following:

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

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

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

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

$$J^*_\delta(x) \le \int_0^\delta q(x^*_T(\tau;x), u^*_T(\tau;x))d\tau \quad 0 < \delta \le T. \tag{2}$$

The last inequality follows from the fact that u^*_T is not (necessarily) optimal over $[0,\delta]$. Hence, we can write (2) as the following inequality:

$$J^*_{T-\delta}(x) - J^*_{T-\delta}(x^*_T(\delta;x)) \ge J^*_{T-\delta}(x) - J^*_T(x) + J^*_\delta(x) \tag{3}$$

where $0 < \delta \le T$. Also, (2) implies that the finite horizon sub level sets are indeed larger than the infinite horizon ones, i.e.,

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

All we need in order to prove stability is to show that there exist a large enough but finite T^* such that the right hand side of (3) is positive. To show this, we define the following function over the compact set Γ^∞_r:

$$\begin{aligned} \Phi_T(x) &:= \frac{J^*_\infty(x) - J^*_T(x)}{J^*_\delta(x)} \ \forall\, x \ne 0 \\ \Phi_T(0) &:= \limsup_{x\to 0} \Phi_T(x) \end{aligned} \tag{4}$$

It is clear that for each $x \ne 0$, $\Phi_T(x)$ is continuous, because the numerator of $\Phi_T(x)$ is continuous and the denominator is continuous and non-zero. Hence it is also *upper semicontinuous* everywhere except possibly at the origin. In order to show that $\Phi_T(x)$ is upper semicontinuous at the origin as well, all we need to prove is that $\Phi_T(0) \ne \infty$. The rest is clear from definition 4..1 and the definition of $\Phi_T(0)$.

Also, note that for all non-zero values of $x \in \Gamma^\infty_r$, the sequence $\{\Phi_T(x)\}$ is a monotone, decreasing sequence, tending to zero as $T \to \infty$. We must show that $\Phi_T(0)$ is bounded and converges to zero when $T \to \infty$ as well. This is true due to the fact that around the origin, $J^*_T(x) = x^T P_T x + R_T(x)$, $J^*_\infty(x) = x^T P_\infty x + R_\infty(x)$, and $J^*_\delta(x) = x^T P_\delta x + R_\delta(x)$ where P_T, P_δ, and P_∞ are the positive definite solutions of the corresponding differential and algebraic Riccati equations respectively [6], and

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

Therefore, in a small neighborhood of the origin we have

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

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

We can now conclude that $\{\Phi_{T-\delta}(x)\}$ is a sequence of upper semicontinuous functions that converges monotonically and pointwise to zero, therefore by Proposition 2, the convergence is uniform.

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

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

Thus,

$$J^*_T(x) - J^*_{T-\delta}(x) < \frac{1}{2} J^*_\delta(x)$$

(note that $J^*_\infty(x) \ge J^*_T(x)$). Combining the above equation with 3

$$J^*_{T-\delta}(x^*_T(\delta;x)) - J^*_{T-\delta}(x) \le -\frac{1}{2} J^*_\delta(x) \quad \forall\, x \in \Gamma^\infty_r$$

This implies that $J^*_{T-\delta}(x)$ is a Lyapunov function proving asymptotic (and in fact exponential) stability of the closed-loop system with $\Gamma^T_{r_T}$ being the region of attraction, and r^2_T being the value of the largest sub level set of $J^*_{T-\delta}(x)$ that lies inside Γ^∞_r, for any given $r > 0$. □

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

5. Conclusion

The purpose of this paper was to extend some of the authors' previous results on unconstrained receding horizon control. It is shown that when a CLF is used as terminal cost, the receding horizon trajectories converge to the origin exponentially fast. Also, when the horizon length is long enough, similar exponential stability results can be obtained without using a terminal cost by using Dini's theorem on the uniform convergence of semicontinuos functions.

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. *Direct Methods in the Calculus of Variations*. Oxford University Press, New York, 1998.

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

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

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

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

[7] A. Jadbabaie and J. Hauser. Relaxing the optimality condition in receding hor izon control. In *IEEE Conference on Decision and Control*, Sydney, Australia, Dec. 2000.

[8] 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.

[9] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control of nonlinear systems. In *IEEE Conference on Decision and Control*, 1999.

[10] 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.

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

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

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

[21] 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.