JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 23, No. 3, May–June 2000

Receding Horizon Control Lyapunov Function Approach to Suboptimal Regulation of Nonlinear Systems

M. Sznaier*
Pennsylvania State University, University Park, Pennsylvania 16802
J. Cloutier†
U.S. Air Force Research Laboratory, Eglin Air Force Base, Florida 32542
R. Hull‡
Coleman Research Corporation, Orlando, Florida 32819
and
D. Jacques§ and C. Mracek¶
U.S. Air Force Research Laboratory, Eglin Air Force Base, Florida 32542

The problem of rendering the origin an asymptotically stable equilibrium point of a nonlinear system while optimizing some measure of performance has been the object of much attention in the past few years. In contrast to the case of linear systems where several optimal synthesis techniques (such as $\mathcal{H}_\infty$, $\mathcal{H}_2$, and ℓ^1) are well established, their nonlinear counterparts are just starting to emerge. Moreover, in most cases these tools lead to partial differential equations that are difficult to solve. In this paper we propose a suboptimal regulator for nonlinear affine systems based upon the combination of receding horizon and control Lyapunov function techniques. The main result of the paper shows that this controller is nearly optimal provided that a certain finite horizon problem can be solved on-line. Additional results include 1) sufficient conditions guaranteeing closed-loop stability even in cases where there is not enough computational power available to solve this optimization on-line and 2) an analysis of the suboptimality level of the proposed method. These results are illustrated with two simple examples comparing the performance of the suboptimal controller against that achieved by some other popular nonlinear control methods.

I. Introduction

A LARGE number of control problems involve designing a controller capable of rendering some point an asymptotically stable equilibrium point of a given time-invariant system while optimizing some performance index. This problem is relevant for disturbance rejection, tracking, and robustness to model uncertainty.[1] In the case of linear dynamics, this problem has been thoroughly explored during the past decade, leading to powerful formalisms such as μ synthesis and ℓ^1 optimal control theory that have been successfully employed to solve some hard practical problems. More recently, these techniques have been extended to handle multiple, perhaps conflicting, performance specifications (see, for example, Refs. 2 and 3 and references therein).

In the case of nonlinear dynamics, popular design techniques include Jacobian linearization (JL),[4] feedback linearization (FL),[4] the use of control Lyapunov functions (CLF),[5–7] recursive backstepping,[4] and recursive interlacing.[8] Although these methods provide powerful tools for designing globally (or semiglobally) stabilizing controllers, performance of the resulting closed-loop systems can vary widely, as we illustrate in the sequel with the problem of controlling a thrust-vectored aircraft. A simplified planar model of the system is shown in Fig. 1, with the corresponding dynamics given by (see Refs. 9 and 10 for details):

$$\begin{bmatrix} \ddot{x} \\ \ddot{y} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} -g\sin\theta \\ g(\cos\theta - 1) \\ 0 \end{bmatrix} + \begin{bmatrix} \cos\theta/m & -\sin\theta/m \\ \sin\theta/m & \cos\theta/m \\ r/J & 0 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} \tag{1}$$

where x, y, and θ denote horizontal, vertical, and angular position, respectively, and where u_1 and $u_2 = \tilde{u}_2 + mg$ (Following Ref. 9, the control u_2 has been shifted to compensate for gravity.) are the control inputs.

Assume that the goal is to drive the system to the origin while minimizing a performance index of the form

$$J(x_0, u) = \int_0^\infty [\xi' Q \xi + u' R u]\,\mathrm{d}t \tag{2}$$

$$\xi(0) = \begin{bmatrix} 0 & 0 & 0 & 12.5 & 0 & 0 \end{bmatrix} \tag{3}$$

$$Q = \mathrm{diag}\begin{bmatrix} 5 & 5 & 1 & 1 & 1 & 5 \end{bmatrix}, \qquad R = I_{2\times 2} \tag{4}$$

corresponding to the following choice of state variables: $\xi = [x\ y\ \theta\ \dot{x}\ \dot{y}\ \dot{\theta}]$.

Table 1 compares the performance achieved by several commonly used nonlinear control design methods. As shown there, in this case the performance of the CLF and JL controllers is worse, by an order of magnitude, than the optimal cost (obtained by offline optimization using a conjugate gradient algorithm[11]). Indeed, an invited session at the 1996 Conference on Decision and Control (CDC)[12] and a recent workshop on nonlinear control[10] have shown that while the methods just mentioned can recover the optimal under certain conditions; in general there are no guarantees on the performance of the resulting system.

As an alternative, during the past few years nonlinear counterparts of $\mathcal{H}_\infty$ (Refs. 13–16) and ℓ^1 (Ref. 17) have started to emerge. Although these theories are appealing because they guarantee optimal performance (at least in a given sense), from a practical standpoint they suffer from the fact that they lead to Hamilton–Jacobi–Isaacs type partial-differential equations that are hard to solve, except in some restrictive, low-dimensional cases.

Finally, there are two classes of methods that do not guarantee, in general, global stability, but work well in practice. The first is the class of state-dependent Riccati equation methods (SDRE) developed by Cloutier and coworkers,[18] based upon recasting the

Received 21 April 1998; revision received 2 December 1998; accepted for publication 23 October 1999.

*Associate Professor, Department of Electrical Engineering; msznaier@frodo.ee.psu.edu.
†Principal Research Scientist, Navigation and Control Branch.
‡Senior Systems Engineer, 7675 Municipal Drive.
§Assistant Professor, Navigation and Control Branch; currently with Air Force Institute of Technology, Wright–Patterson AFB, OH 45433.
¶Senior Principal Engineer, Navigation and Control Branch; currently with Raytheon Systems Company, Tucson, AZ 85734.

Table 1 Comparison of different methods for the thrust-vectored aircraft example

Method	Cost
Exact	1115
CLF[10]	2.53×10^4
JL + LQR[a 10]	1.1×10^5

[a]Linear quadratic regulator.

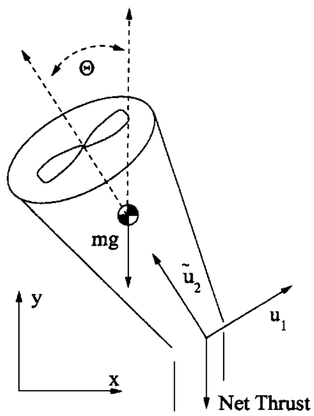

Fig. 1 Simplified model of a thrust-vectored aircraft.

nonlinear plant into a linear-like, state-dependent form and solving at each point a Riccati equation. Although at the present time only local stability of the resulting system has been formally established, consistent experience indicates that the method works well, often outperforming other techniques.[10] For instance, for the thrust-vectored aircraft example just mentioned, SDRE regulation yields a cost $J = 1646$, significantly outperforming the JL and CLF methods.

The second class, receding horizon (RH) control methods, is based upon the on-line solution of a finite-horizon optimal control problem at each sampling instant kT_s. The first element of the resulting (open-loop) control sequence is implemented over the interval $[kT_s, (k+1)T_s]$. At time $(k+1)T_s$ the process is repeated with the new initial condition x_{k+1}. These methods are appealing because they allow for explicitly handling constraints and guarantee optimality in some sense. Moreover, because the optimization is carried only along the present trajectory of the system (i.e, locally), the resulting computational complexity is far less than that associated with finding the true global optimal control (a task that entails solving a Hamilton–Jacobi type equation). However, in contrast with the linear case (where global stability has been established[19,20]), for nonlinear plants only local stability results are available.[21] For instance, for the thrust-vectored aircraft example a standard receding horizon controller fails to stabilize the plant until the horizon is increased to $T_H = 2$ s. In this case such an horizon effectively amounts to solving on-line the complete trajectory, an approach clearly impractical, even for moderately sized problems. Modified nonlinear RH formulations addressing this problem have been proposed in Refs. 22 and 10, but these approaches achieve stability at the expense of performance.

In this paper we propose an alternative controller for suboptimal regulation of nonlinear affine systems, based upon the combination of receding horizon and control Lyapunov functions ideas. This approach follows in the spirit of a similar controller successfully used in the case of constrained linear systems,[19,23] where a receding horizon was used to drive the system to an invariant neighborhood of the origin where a stabilizing controller is available. In the first part of the paper, we show that combining these ideas with a suitable finite-horizon approximation of the performance index leads to globally stabilizing, nearly optimal controllers, provided that enough computational power is available to solve on-line an optimization problem. As a comparison, the nonlinear RH controller proposed in Ref. 22, also based upon the use of dual-mode control, although guaranteed to stabilize the plant under comparable computational-time constraints, does not approximate the optimal cost, leading to a mismatch between the predicted and actual trajectories and their respective associated costs.

In the second part of the paper, we show how to modify the proposed controller to guarantee global stability in the face of computational time constraints, and we establish a connection with the well known CLF methodology. Additional results include an analysis of the suboptimality of the proposed method and show that if an approximate solution to the problem is known in a set containing the origin then our controller yields an extension of this solution with the same performance level.

The paper is organized as follows: in Sec. II we introduce the notation to be used and some preliminary results. In Sec. III we show that the infinite horizon regulation problem is approximately equivalent to a finite horizon problem, and we analyze the properties of this approximation. In Sec. IV we present the proposed controller, and we analyze the properties of the resulting closed loop. Section V briefly addresses the issue of selecting appropriate CLFs. Section VI illustrates the synthesis method with two simple examples. Finally, in Sec. VII we summarize our results, and we indicate directions for future research.

II. Preliminaries

A. Notation and Definitions

Consider the control-affine nonlinear system

$$\dot{x} = f(x) + g(x)u \tag{5}$$

where $x \in R^n$, $u \in R^m$, the vector fields $f(\cdot)$ and $g(\cdot)$ are known $\mathcal{C}^1$ functions and where $f(0) = 0$. Given a function $V: R^n \longrightarrow R$, its Lie derivative along f is defined as

$$L_f V(x) = \frac{\partial V}{\partial x} f(x)$$

Definition 1: A positive definite and radially unbounded $\mathcal{C}^1$ function $V: R^n \longrightarrow R_+$ is a CLF for the system (5) if

$$\inf_u [L_f V(x) + L_g V(x) u] < 0 \qquad \forall x \neq 0 \tag{6}$$

It is well known (see, for example, Ref. 4, p. 26) that existence of a CLF is equivalent to the existence of a globally asymptotically stabilizing feedback control law $u(x)$. (If, in addition, the so-called small control property holds then the stabilizing control law is continuous.)

B. Nonlinear Regulator Problem

Consider the nonlinear system (5). Our goal is to find a feedback control law $u(x)$ that minimizes the following performance index:

$$J(x_0, u) = \int_0^\infty [x' Q(x) x + u' R(x) u] \, \mathrm{d}t, \qquad x(0) = x_0 \tag{7}$$

where $Q(\cdot)$ and $R(\cdot)$ are $\mathcal{C}^1$, positive definite matrices. [This condition can be relaxed to $Q(x) \geq 0$.] It is well known (Ref. 24, Sec. 8.5) that this problem is equivalent to solving the following Hamilton–Jacobi–Bellman (HJB) partial differential equation:

$$0 = \frac{\partial V}{\partial x} f - \frac{1}{4} \frac{\partial V}{\partial x} g R^{-1} g' \frac{\partial V'}{\partial x} + x' Q(x) x, \qquad V(0) = 0 \tag{8}$$

If this equation admits a $\mathcal{C}^1$ nonnegative solution V, then the optimal control is given by $u = -\frac{1}{2} R^{-1} g' (\partial V' / \partial x)$ and $V(x)$ is the corresponding optimal cost (or storage function), i.e.,

$$V(x_0) = \min_u J(x_0, u)$$

III. Equivalent Finite Horizon Regulation Problem

Unfortunately, the complexity of Eq. (8) prevents its solution, except in some very simple, low-dimensional cases. To solve this difficulty, in this section we introduce a finite horizon approximation to the nonlinear regulation problem, and we analyze its properties. This approximation forms the basis of the proposed method.

Lemma 1: Consider a compact set $\mathcal{S}$ containing the origin in its interior and assume that the optimal storage function $V(x)$ is known for all $x \in \mathcal{S}$. Let $v = \min_{x \in \partial S} V(x)$, where ∂S denotes the boundary of S. Finally, define the set $\tilde{S} = \{x: V(x) \leq v\}$. Consider the following two optimization problems:

$$\min_u \left\{ J(x_0, u) = \int_0^\infty [x' Q(x) x + u' R u] \, dt : x(0) = x_0 \right\} \quad (9)$$

$$\min_u \left\{ J_T(x_0, u) = \int_0^T [x' Q(x) x + u' R u] \, dt + V[x(T)]: x(0) = x_0 \right\} \quad (10)$$

subject to Eq. (5). Then the following facts hold:

1) Let $u^0(\cdot), x^0(\cdot)$ denote an optimal solution of problem (10) and the corresponding trajectory, respectively. If $x^0(T) \in \tilde{S}$, then the control law $u^0(\cdot)$ is also optimal for Eq. (9) in the interval $[0, T]$.

2) Consider now $T_1 > T$. If $x^0(T) \in \tilde{S}$, then a controller that optimizes J_T is also optimal with respect to J_{T_1} in $[0, T]$.

Proof: Consider first the following free terminal time problem:

$$J^0(x_0, t) = \min_u \left\{ V[x(t_f)] + \int_t^{t_f} [x' Q x + u' R u] \, dt : x(0) = x_0 \right\}$$

subject to:

$$x(t_f) \in \tilde{S} \quad (11)$$

The optimal return function satisfies[24]

$$0 = \frac{\partial J}{\partial t} + \frac{\partial J}{\partial x} f - \frac{1}{4} \frac{\partial J}{\partial x} g R^{-1} g' \frac{\partial J'}{\partial x} + x' Q(x) x \quad (12)$$

with boundary condition $J(x, t) = V(x)$ for $x \in \tilde{S}$. Clearly this equation admits as solution $J(x, t) \equiv V(x)$. Thus problems (9) and (11) are equivalent. To establish the first claim, we will show that an optimal solution u^0 of Eq. (10) is also optimal for Eq. (11) [and thus Eq. (9)], provided that $x^0(T) \in \tilde{S}$. To this effect, note that the Euler–Lagrange optimality conditions for problems (10) and (11) are identical, except for the additional transversality condition $H[u^0, x^0(t_f), \lambda^0(t_f)] = 0$ that appears in the latter, where $\lambda(t)$ denotes the costates. The boundary condition for λ in problem (10) is given by

$$\lambda^0(T) = \left. \frac{\partial V}{\partial x} \right|_{x(T)} \quad (13)$$

Because $x^0(T) \in \tilde{S}$, it follows that $x^0(T), u^0(T), \lambda^0(T)$ satisfy the HJB equation (8), or equivalently $H[u^0, x^0(T), \lambda^0(T)] = 0$. Thus, an optimal solution of Eq. (10) is also optimal for Eq. (11).

To establish the second claim, note that the set $\tilde{S}$ is positively invariant with respect to the optimal control law. Consider now the following control law:

$$u^* = \begin{cases} u^0(t) & \text{for} \quad t \in [0, T] \\ -\frac{1}{2} R^{-1} g' \dfrac{\partial V}{\partial x} & \quad t > T \end{cases} \quad (14)$$

Because $V(x)$ is the optimal return function in $\tilde{S}$ and this set is invariant, it follows that

$$V[x(T)] = V[x(T_1)] + \int_T^{T_1} (x^{*\prime} Q x^* + u^{*\prime} R u^*) \, dt$$

Thus $J_{T_1}(x, u^*) = J_T$. From the equivalence of Eq. (10) and the free terminal time problem (11), it follows that $J_{T_1} \geq J_T$. Hence $\min_u J_{T_1} = J_T$, and u^* is the corresponding optimal control law. □

This lemma shows that if a solution to the HJB equation (8) is known in a neighborhood of the origin, then it can be extended via an explicit finite horizon optimization, well suited for an on-line implementation. This suggest a receding-horizon-type control combining an on-line optimization with an off-line phase to find a local solution to Eq. (8). However, finding (and storing) this local solution can be very computationally demanding in cases where the dimension of the problem is not low. Thus, it is of interest to consider the case where an approximation $\Psi(x)$ rather than the true storage function $V(x)$ is used in Eq. (10). The next result shows that in this case the approximation error does not grow (to the first order). In other words, the difference between the true optimal $V[x(0)]$ and $J_T[x(0)]$ is approximately equal to the difference between $V[x(T)]$ and $\Psi[x(T)]$.

Theorem 1: Let Ψ: $R^n \to R_+$ be a positive definite function and consider the following optimization problem:

$$J_\Psi(x, t) = \min_u \int_t^T \left(x' Q x + u' R u \right) d\tau + \Psi[x(T)] \quad (15)$$

subject to Eq. (5). Then

$$J_\Psi(x, t) - V[x(t)] = \Psi[x(T)] - V[x(T)] + \mathcal{O}\left[\left\| \left. \frac{\partial e}{\partial x} \right|_{x(T)} \right\|^2 \right] + \mathcal{O}(dt^2) \quad (16)$$

where $e(x, t) \doteq J_\Psi(x, t) - V(x)$ denotes the approximation error.

Proof: By considering the Hamilton Jacobi equations for J_Ψ and V, it can be easily shown that $e(t, x)$ satisfies the following equation:

$$0 = \frac{\partial e}{\partial t} + \frac{\partial e}{\partial x} \left(f - \frac{1}{2} g R^{-1} g' \frac{\partial J_\Psi}{\partial x} \right) + \frac{1}{4} \frac{\partial e}{\partial x} g R^{-1} g' \frac{\partial e'}{\partial x} \quad (17)$$

By exploiting the fact that the optimal control law for Eq. (15) is given by

$$u_\Psi = -\frac{1}{2} R^{-1} g' \frac{\partial J_\Psi}{\partial x} \quad (18)$$

Eq. (17) can be rewritten as

$$0 = \frac{\partial e}{\partial t} + \frac{\partial e}{\partial x}(f + g u_\Psi) + \frac{1}{4} \frac{\partial e}{\partial x} g R^{-1} g' \frac{\partial e'}{\partial x} = \frac{\partial e}{\partial t} + \frac{\partial e}{\partial x} \dot{x} + \frac{1}{4} \frac{\partial e}{\partial x} g R^{-1} g' \frac{\partial e'}{\partial x} = \dot{e} + \frac{1}{4} \frac{\partial e}{\partial x} g R^{-1} g' \frac{\partial e'}{\partial x} \quad (19)$$

From this last equation it follows that

$$\dot{e}(T) = -\frac{1}{4} \frac{\partial e}{\partial x} g R^{-1} g' \left. \frac{\partial e'}{\partial x} \right|_{x(T)} \quad (20)$$

Expanding $e(t)$ in a Taylor series around $t = T$ yields

$$e(t) = e(T) + \dot{e}(T) \, dt + \mathcal{O}(dt^2) = e(T) - \frac{1}{4} \frac{\partial e}{\partial x} g R^{-1} g' \left. \frac{\partial e'}{\partial x} \right|_{x(T)} dt + \mathcal{O}(dt^2) \quad (21)$$

□

Corollary 1: Assume that Ψ is selected so that $\|\partial e[x(T), T]/\partial x\| \cong 0$. Then $J_\Psi(x, t) - V[x(t)] \cong \Psi[x(T)] - V[x(T)]$ (to the first order in dt) along the trajectories of the system.

From Lemma 1 it follows that, given an initial condition $x(0)$, problem (9) can be solved by solving a sequence of problems of the form [Eq. (10)] with increasing T until a solution is found such that $x(T) \in \tilde{S}$. Moreover, once such a solution is obtained, no further improvement of the cost can be achieved by increasing the horizon T. These results suggest the following receding-horizon-type control law. Let $x(t)$ denote the current state of system (5). Then

0) Data: $V(x)$, $\tilde{S}$, a sampling interval δT.

1) If $x(t) \in \tilde{S}$, $u = -\frac{1}{2} R^{-1} g' (\partial V' / \partial x)$.

2) If $x(t) \notin \tilde{S}$, then solve a sequence of optimization problems of the form [Eq. (10)] until a solution such that $x(T) \in \tilde{S}$ is found. Use the corresponding control law $u(t)$ in the interval $[t_0, t_0 + \delta T]$.

From the preceding results it is clear that the resulting control law is globally optimal and thus globally stabilizing. However, as we indicated before, the computational complexity associated with finding $V(x)$ (even only in the region $\tilde{S}$) may preclude the use of this control law in many practical cases. Thus, it is of interest to consider a modified control law where an approximation $\Psi(x)$ [rather than $V(x)$] is used. To this effect, consider a compact set S containing the origin in its interior and let Ψ: $S \to R_+$, $\Psi \in \mathcal{C}^1$ be a CLF for system (5). Denote by u_Ψ the corresponding control law. Finally, let

$c = \min_{x\in\partial S}\Psi(x)$ and define the set $S_\Psi \subseteq S = \{x : \Psi(x) \le c\}$. Then we propose the following modified control law:

0) Data: $\Psi(x)$, S_ψ, δT, a positive definite function $\sigma(x)$.

1) If $x \in S_\Psi$, $u_\Psi(x) \doteq \mathrm{argmin}_u\{\|u\|: L_f\Psi + L_g\Psi u \le -\sigma(x) < 0\}$.

2) If $x \notin S_\Psi$, then consider an increasing sequence T_i. Let

$$u_T^* = \mathrm{argmin}\left\{\int_0^T \left(x'Qx + u'Ru\right)\mathrm{d}t + \Psi[x(T)]\right\}$$

Denote by $x^*(\cdot)$ the corresponding optimal trajectory and define $T(x) = \inf\{T : x^*(T) \in S_\Psi\}$. [From Barbalat's Lemma[4] (p. 491) we have $x'Qx + u'Ru \to 0$ as $T \to \infty$. Hence for every x_0, $T(x_0)$ is finite.] Then $u_\Psi(x) \doteq u^*_{T(x)}(t)$, $t \in [t_0, t_0 + \delta T]$.

From Theorem 1 it follows that the suboptimality associated with the modified algorithm is approximately given by $e_\Psi = \sup_{x\in S_\Psi}|\Psi(x) - V(x)|$.

Theorem 2: Assume that $\underline{\sigma}[Q(x)] > \sigma_m > 0$ for all x, where $\underline{\sigma}(\cdot)$ denotes the minimum singular value. Then the control law u_Ψ globally stabilizes Eq. (5).

Proof: Consider first an initial condition $x(0) = x_0 \notin S_\Psi$. Let

$$J(x_0) = \int_0^{T(x_0)} (x^{*\prime}Qx^* + u^{*\prime}Ru^*)\,\mathrm{d}t + \Psi\left\{x^*[T(x_0)]\right\}$$

where $u^*(\cdot)$, $x^*(\cdot)$ denote the optimal control and the corresponding trajectory. Let $x_1 \doteq x^*(\mathrm{d}t)$ [with $\mathrm{d}t > 0$ small enough so that $x^*(\mathrm{d}t) \notin S_\Psi$]. Because $x^*(t)$, $t \in [\mathrm{d}t, T]$ is also a feasible trajectory starting from x_1, we have that

$$J(x_1) = \inf_u\left(\int_{\mathrm{d}t}^{T(x_1)} (x'Qx + u'Ru)\,\mathrm{d}\tau + \Psi\{x[T(x_1)]\}\right)$$
$$\le \int_{\mathrm{d}t}^{T(x_0)} (x^{*\prime}Qx^* + u^{*\prime}Ru^*)\,\mathrm{d}\tau + \Psi\{x^*[T(x_0)]\}$$
$$= J(x_0) - \int_0^{\mathrm{d}t} (x^{*\prime}Qx^* + u^{*\prime}Ru^*)\,\mathrm{d}\tau \tag{22}$$

Thus,

$$J(x_1) - J(x_0) \le -\int_0^{\mathrm{d}t} (x^{*\prime}Qx^* + u^{*\prime}Ru^*)\,\mathrm{d}\tau$$
$$\le -\int_0^{\mathrm{d}t} (x^{*\prime}Qx^*)\,\mathrm{d}\tau \tag{23}$$

It follows that

$$\dot{J} = \lim_{\mathrm{d}t\to 0}\frac{J\{x(t+\mathrm{d}t) - J[x(t)]\}}{\mathrm{d}t} \le -x(t)'Qx(t) \le -\sigma_m\cdot\|x\| < 0$$

Because $J(x) > 0$ and $\dot{J}(x) < 0$ for all $x \notin S_\Psi$, it follows that trajectories originating outside this set reach it in a finite time. Asymptotic stability now follows from the facts that S_Ψ is invariant with respect to u_Ψ (i.e., trajectories starting in the set never leave it) and that $\Psi(x)$ is a CLF there. □

IV. Modified Receding Horizon Controller

In the last section we outlined a receding-horizon-type law that under certain conditions is nearly optimal and globally stabilizes system (5). Although most of these conditions are rather mild (essentially equivalent to the existence of a CLF), the requirement that T should be large enough so that $x(T) \in S_\Psi$ could pose a problem, especially in cases where the system has fast dynamics. In this section we propose a modified control law that is guaranteed to stabilize the system, even when this condition fails, and that takes into account computational time constraints.

Consider the following receding horizon control law:

Algorithm 1:

0) Data: a CLF $\Psi(x)$, an invariant region S_Ψ such that $0 \in \mathrm{int}\{S_\Psi\}$, a horizon T.

1) If $x(t) \in S_\Psi$, $u_\Psi(x) = \mathrm{argmin}_u\{\|u\|: L_f\Psi + L_g\Psi u \le -\sigma(x) < 0\}$.

2) If $x(t) \notin S_\Psi$, then $u_\Psi(x) = u(t)$ where $u(\lambda)$, $\lambda \in [t, t+T]$ is given by

$$u = \mathrm{argmin}_u\left\{\int_t^{T+t} (x'Qx + u'Ru)\,\mathrm{d}t + \Psi[x(T+t)]\right\} \tag{24}$$

Subject to:

$$-\sigma[x(t+T)] \ge x(t+T)'Qx(t+T) + L_f\Psi|_{x(t+T)} + \min_v\left\{v'Rv + L_g\Psi|_{x(t+T)}v\right\} - x(t)'Qx(t) - u'(t)Ru(t) \tag{25}$$

where $\sigma(x)$ is some positive definite function such that $\sigma(x) \le x'Qx$.

Theorem 3: The control law u_Ψ generated by Algorithm 1 has the following properties:

1) It renders the origin a globally asymptotically stable equilibrium point of Eq. (5).

2) It coincides with the globally optimal control law when $\Psi(x) = V(x)$.

3) It is nearly optimal (in the sense of Theorem 1) when $x(T) \in S_\Psi$.

Proof: To prove stability, proceeding as in Theorem 2, consider first an initial condition $x_0 \notin S_\Psi$. Denote by u^*, x^* the optimal control and associated trajectory respectively. Then

$$J[x(t+\mathrm{d}t)] = \min_u\left\{\int_{t+\mathrm{d}t}^{T+t+\mathrm{d}t} (x'Qx + u'Ru)\,\mathrm{d}t + \Psi[x(T+t+\mathrm{d}t)]\right\} \le \int_{t+\mathrm{d}t}^{T+t} (x^{*\prime}Qx^* + v^{*\prime}Rv^*)\,\mathrm{d}t + \Psi[x^*(T+t)]$$
$$+ \min_v\{x^{*\prime}(T+t)Qx^*(T+t) + v'Rv + \dot{\Psi}[x^*(T+t)]\}\mathrm{d}t$$
$$= J[x(t)] - [x^*(t)'Qx^*(t) + u^{*\prime}(t)Ru^*(t)]\,\mathrm{d}t$$
$$+ \min_v\{x^{*\prime}(T+t)Qx^*(T+t) + v'Rv + \dot{\Psi}[x^*(T+t)]\}\mathrm{d}t \tag{26}$$

Therefore, if Eq. (25) holds, then we have

$$\dot{J} = \lim_{\mathrm{d}t\to 0}\frac{J[x(t+\mathrm{d}t)] - J[x(t)]}{\mathrm{d}t} \le -\sigma(x) < 0 \tag{27}$$

Hence the trajectories starting outside S_Ψ reach this set in a finite time. As in the proof of Theorem 2, once there asymptotic stability is guaranteed by the control Lyapunov function Ψ.

To prove item 2, note that when $\Psi(x) = V(x)$ then from the Hamilton Jacobi equation (8) we have

$$x(t+T)'Qx(t+T) + L_f\Psi|_{x(t+T)} + \min_v\left\{v'Rv + L_g\Psi|_{x(t+T)}v\right\} \le 0 \tag{28}$$

Thus in this case the constraints (25) are redundant. The proof follows now immediately from Lemma 1. Finally, the proof of item 3 follows from Theorem 1.

Remark 1: If $T \to 0$ in Eq. (24), then u_Ψ is given by the solution to the following optimization problem:

$$u_\Psi = \mathrm{argmin}_u\{u'Ru + L_g\Psi u\} \tag{29}$$

Subject to:

$$\sigma[x(t)] + L_f\Psi[x(t)] + L_g\Psi[x(t)]u + \alpha[x(t)] = 0 \tag{30}$$

where $-\alpha(x)$ is the desired negativity margin.[25] The solution to this optimization problem is given by

$$u_\Psi = -\frac{\psi_0 R^{-1}\psi_1'}{\psi_1 R^{-1}\psi_1'} \tag{31}$$

where

$$\psi_0 = L_f\Psi[x(t)] + \sigma[x(t)] + \alpha[x(t)], \qquad \psi_1 = L_g\Psi[x(t)] \tag{32}$$

This is precisely the inverse optimal controller proposed in Ref. 25

Finally, before closing this section we consider a modified control law that takes into account the sample and hold nature of receding-horizon implementations.

Algorithm 2:

0) Data: a CLF $\Psi(x)$, an invariant region S_Ψ such that $0 \in \text{int}\{S_\Psi\}$, a horizon T, a sampling interval δT, a positive definite function $\sigma(\cdot)$.

1) If $x(t) \in S_\Psi$, $u_s(x) = \text{argmin}_u\{\|u\|: L_f\Psi + L_g\Psi u \leq -\sigma(x) < 0\}$.

2) If $x(t) \notin S_\Psi$, then $u_s(\tau) = u(\tau)$, $\tau \in [t, t + \delta T]$, where $u(\cdot)$ is given by

$$u = \underset{u}{\text{argmin}}\left\{\int_t^{T+t} (x'Qx + u'Ru)\,\text{d}t + \Psi[x(T+t)]\right\} \tag{33}$$

Subject to:

$$\begin{aligned}-\sigma[x(\lambda + t + T)] \geq\, & x(\lambda + t + T)'Qx(\lambda + t + T) \\ & + L_f\Psi|_{x(\lambda+t+T)} + u'(\lambda + t + T)Ru(\lambda + t + T) \\ & + L_g\Psi|_{x(\lambda+t+T)}u(\lambda + t + T) - x(t)'Qx(t) - u'(t)Ru(t) \\ & \text{for all} \quad 0 \leq \lambda \leq \delta T\end{aligned} \tag{34}$$

□

Lemma 2: The control law u_s renders the origin a globally stable equilibrium point of Eq. (5). Moreover, it is nearly optimal and coincides with the globally optimal control law when $\Psi(x) = V(x)$.

Proof: The proof, omitted for space reasons, follows along the same lines of the proof of Theorem 3.

V. Selecting Suitable CLFs

In principle, any of the methods available in the literature for finding CLFs such as feedback linearization and backstepping (see, for instance, Refs. 7 and 26) can be used to find the function $\Psi(\cdot)$. However, as shown in Lemma 1 the suboptimality level incurred by the proposed algorithm is roughly similar to the difference between the CLF $\Psi(\cdot)$ and the actual value function V. Thus, it is of interest to find CLFs that are close to V. In this section we briefly discuss an alternative, motivated by the empirically observed success of the SDRE method, briefly covered in the Appendix.

From Lemma 3 in the Appendix, it follows that $\Psi(x) = x'P(x)x$, where $P(x)$ denotes the solution to the SDRE, is a CLF in a neighborhood of the origin. Moreover, because the control law u_{sdre} is locally stabilizing, it can be easily shown that there exists T_0 (possibly depending on the initial condition) such for all $T > T_0$ the constraints (25) are feasible. It follows that $\Psi(x) = x'P(x)x$ is a suitable choice for the terminal penalty. Moreover from the properties of the SDRE method (see the Appendix), it follows that with this choice the control law satisfies all the necessary conditions for optimality as $\mathcal{O}[\|x(t+T)\|^2]$. Thus, we will expect that Algorithm 2 using $\Psi(x) = x'P(x)x$ will generate a nearly optimal control law, even when T is relatively small. In the next section we show that this is indeed the case with two examples.

VI. Illustrative Examples

In this section we illustrate our results with two examples. The first one is a simple academic example that can be solved analytically. Thus it can be used to analyze explicitly the source of performance degradation when using feedback linearization or the SDRE method and to show the advantages of the proposed approach. The second example, a realistic problem arising in the context of control of thrust-vectored aircraft, was used for benchmarking several nonlinear design methods in Ref. 10.

Example 1: Consider the following regulation problem:

$$\min_u\left\{J = \int_0^\infty x_2^2 + u^2\,\text{d}t\right\} \tag{35}$$

subject to:

$$\begin{aligned}\dot{x}_1 &= x_2 \\ \dot{x}_2 &= -x_1e^{x_1} + \tfrac{1}{2}x_2^2 + e^{x_1}u\end{aligned} \tag{36}$$

The optimal control law is given by

$$u_{\text{opt}} = -x_2 \tag{37}$$

with the corresponding optimal storage function

$$V(x) = x_1^2 + x_2^2e^{-x_1} \tag{38}$$

A feedback linearization design selected so that the closed-loop system has the same storage function as $\|x\| \to 0$ yields the following controller and Lyapunov function:

$$u_{\text{FL}} = (1 - e^{-x_1})x_1 - x_2e^{-x_1}(1 + 0.5 * x_2), \qquad V_{\text{FL}} = x_1^2 + x_2^2 \tag{39}$$

Note that $u_{\text{FL}} \cong u_{\text{opt}}$ only for small values of x_1 and x_2. Consider now the following SDC parameterization:

$$A = \begin{bmatrix} 0 & 1 \\ -e^{x_1} & \frac{1}{2}x_2\end{bmatrix}, \qquad B = \begin{bmatrix} 0 \\ e^{x_1}\end{bmatrix} \tag{40}$$

The solution to the corresponding SDRE is given by

$$P(x) = \begin{bmatrix} e^{x_1} & 0 \\ 0 & 1\end{bmatrix} p(x) \tag{41}$$

where

$$p(x) = \left[\left(x_2/2e^{x_1}\right) + \sqrt{1 + \left(x_2/2e^{x_1}\right)^2}\right]e^{-x_1} \tag{42}$$

with associated control action

$$u_{\text{sdre}} = -x_2\left[\left(x_2/2e^{x_1}\right) + \sqrt{1 + \left(x_2/2e^{x_1}\right)^2}\right] \tag{43}$$

Finally, it can also be shown that

$$x'P(x)x = pe^{x_1}V(x) \tag{44}$$

Thus $x'P(x)x$ gives a good estimate of $V(x)$, and $u_{\text{sdre}} \cong u_{\text{opt}}$ only when $(x_2/2e^{x_1}) \ll 1$.

Table 2 shows the different costs starting from the initial condition $x(0) = [-2 \;\; 2]$ for several controllers, with the corresponding trajectories shown in Figs. 2 and 3. The last two entries of the table correspond to the proposed controller using an horizon $T = 1$ s and as estimates of the value function $\Psi = V_{\text{FL}}$ and $\Psi = x'P(x)x$, respectively. In this case performance of both controllers is virtually identical to optimal.

Example 2: Consider again the simplified model of the thrust-vectored aircraft used in the Introduction. Table 3 shows the cost corresponding to the initial condition $\xi(0) = [0 \;\; 0 \;\; 0 \;\; 12.5 \;\; 0 \;\; 0]$, obtained using different controllers. The two lowest entries correspond to the proposed method using $T = 1$ s and $T_s = 0.5$ s and terminal penalties derived from JL and the SDRE methods, respectively. The latter virtually achieves optimal performance, whereas the former is only 2% suboptimal. This behavior can be explained by looking at Fig. 4, which shows the different portions of the cost as a function of the horizon, starting from the initial condition $\xi(0)$. These plots show that whereas $\Psi(x) = x'P(x)x$ gives initially a very

Table 2 Comparison of different methods for example 1

Method	Cost
Optimal	33.56
FL	95.11
SDRE	143.0
RH + FL	34.24
RH + SDRE	33.7

Table 3 Comparison of different methods for example 2

Method	Cost	
Exact	1115	
LQR[10]	1.1×10^5	
CLF[10]	2.53×10^4	
LPV[10 a]	1833	
SDRE	1640	
RH + JL	1142 (T = 1)	1321 (T = 0.4)
RH + SDRE	1117 (T = 1)	1310 (T = 0.4)

[a]Linear parameter varying.

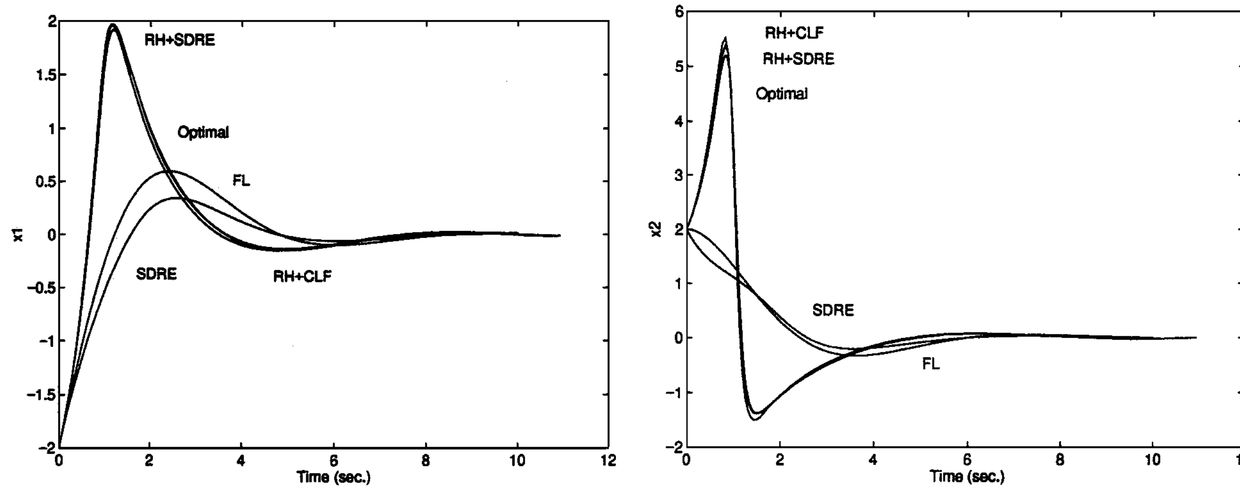

Fig. 2 State trajectories for Example 1.

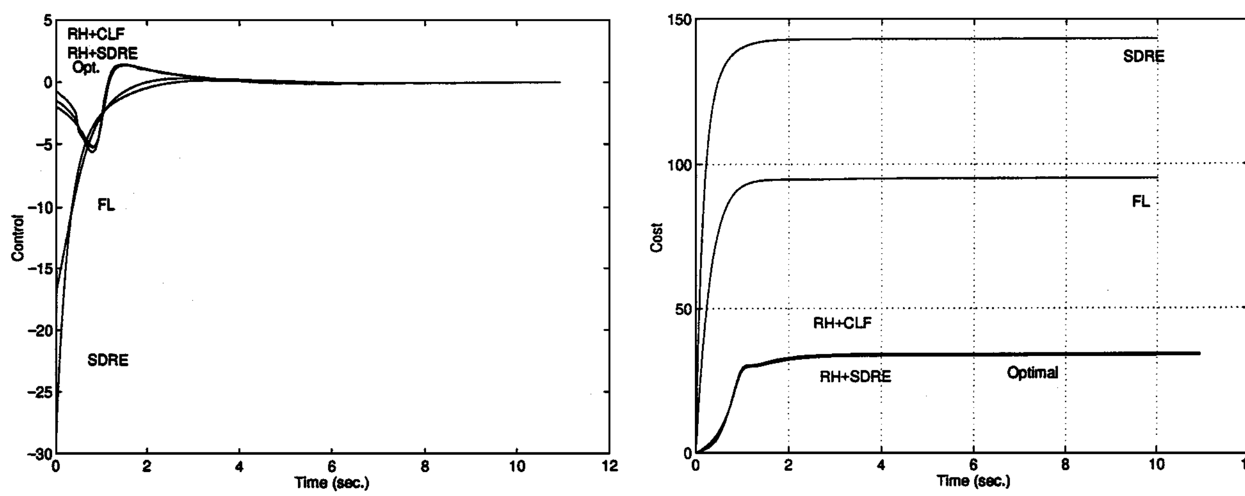

Fig. 3 Control effort and cost for Example 1

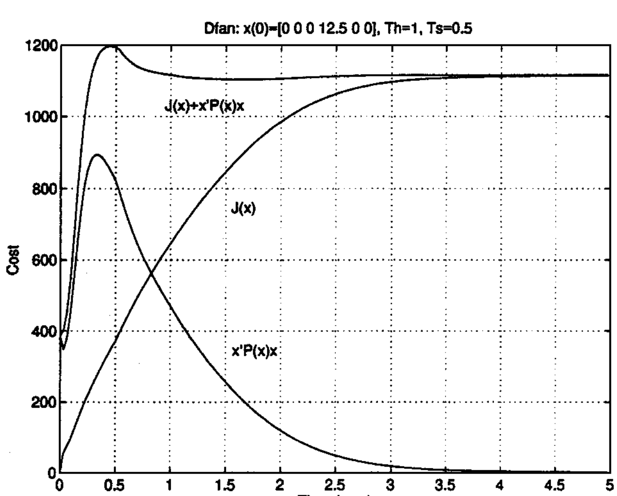

Fig. 4 Terms of the cost as function of the horizon.

poor estimate of the cost-to-go, the combination of $\Psi(x)$ and the explicit integral in Eq. (24) give a very good estimate if T is chosen ≥ 1 s. It is worth mentioning that a conventional receding-horizon controller [i.e., one obtained by setting $\Psi \equiv 0$ in Eq. (24)] with the same choice of horizon and sampling time fails to stabilize the system.

VII. Conclusions

In contrast with the case of linear plants, tools for simultaneously addressing performance and stability of nonlinear systems have emerged relatively recently. Recent counterexamples[10,27] illustrated the fact that although several commonly used techniques can successfully stabilize nonlinear systems the resulting closed-loop performance varies widely. Moreover, these performance differences are problem dependent, with performance of a given method ranging from (near) optimal to very poor.

In this paper we propose a new suboptimal regulator for affine nonlinear systems, based upon the combination of receding-horizon and control Lyapunov functions techniques. The main result of the paper shows that under certain relatively mild conditions this regulator renders the origin a globally asymptotically stable equilibrium point. In the limit as the horizon $T \longrightarrow 0$, the proposed control law reduces to the inverse optimal controller proposed by Freeman and Kokotovic.[25] Thus, these conditions are essentially equivalent to the existence of a CLF. Additional results in the paper show that the regulator is near optimal, provided that a good approximation to the optimal storage function is known in a neighborhood of the origin. These results were illustrated with two examples. In both cases the proposed controller outperformed several other commonly used techniques. Finally, we note that the finite approximation (15) is also valuable as a tool to speed-up off-line numerical computation of near optimal solutions by providing computationally inexpensive, yet suitable initial conditions for numerical optimization algorithms such as conjugate gradients.[11]

An issue that was not addressed in this paper is that of the computational complexity associated with solving the nonlinear optimization problem (24). Following Ref. 10, this complexity could be reduced by exploiting differential flatness to perform the optimization in flat space. Additional research being pursued includes the explicit incorporation of state and control constraints into the formalism and its extension to handle model uncertainty.

Appendix: SDRE Approach to Nonlinear Regulation

In this section we briefly cover the details of the SDRE approach developed by Cloutier and coworkers.[18] The main idea of the method is to recast the nonlinear system (5) into a state-dependent coefficient linear-like form

$$\dot{x} = A(x)x + B(x)u \qquad \text{(A1)}$$

and to solve pointwise along the trajectory the corresponding algebraic Riccati equation:

$$A'(x)P(x) + P(x)A(x) - P(x)B(x)R^{-1}(x)B'(x)P(x) + Q(x) = 0 \qquad \text{(A2)}$$

The suboptimal control law is given by $u_{\text{sdre}} = -\frac{1}{2}R^{-1}(x)B'P(x)x$, where $P(x)$ is the positive definite (pointwise stabilizing) solution of Eq. (A2). In the sequel we briefly review the properties of this control law. The corresponding proofs can be found in the appropriate references.

Lemma 3 (Ref. 18): Assume that $Q(x) = C'(x)C(x)$ and that there exists a neighborhood Ω of the origin where the pairs $\{A(x), B(x)\}$ and $\{A(x), C(x)\}$ are pointwise stabilizable and detectable respectively and all of the matrix functions involved are $\mathcal{C}^1$. Then the control law u_{sdre} renders the origin a locally asymptotically stable equilibrium point of the closed-loop system.

Lemma 4 (Ref. 18): The SDRE control law and its associated state and costate trajectories satisfy the following necessary condition for optimality:

$$\frac{\partial H}{\partial u} = 0$$

where $H = x'Q(x)x + u'R(x)u + \lambda'[f(x) + B(x)u]$ is the Hamiltonian of the system and where λ denotes the costates.

Lemma 5 (Ref. 18): Assume that the parametrization [Eq. (A1)] is stabilizable and all of the matrices involved along with their gradients are bounded in a neighborhood Ω of the origin. Then the SDRE control law and its associated state and costate trajectories asymptotically satisfy at a quadratic rate [i.e., $\|\lambda + \partial H/\partial x\| \longrightarrow 0$ as $\mathcal{O}(\|x\|^2)$ as $x \longrightarrow 0$] following necessary condition for optimality:

$$\dot{\lambda} = -\frac{\partial H}{\partial x}$$

in the sense that

$$\left\| \lambda + \frac{\partial H}{\partial x} \right\| \leq x'Ux$$

for some constant matrix $U > 0$ and all $x \in \Omega$.

Lemma 6 (Ref. 28): Let $P(x)$ denote a solution to the SDRE [Eq. (A2)]. If there exists a positive definite function $V(x)$ such that $\partial V(x)/\partial x = P(x)x$, then u_{sdre} is the globally optimal control law. {A necessary and sufficient condition for this to hold is that the Jacobian matrix $(\partial/\partial x)[P(x)x]$ is symmetric (see, for instance, Ref. 29. Sec. 3.1). In this case $V(x)$ can be computed as

$$V(x) = \int_0^x y'P(y) \cdot \mathrm{d}y$$

where this line integral is independent of the path.}

Acknowledgments

The authors are indebted to Yun Huang, Control and Dynamical Systems, California Institute of Technology, for providing the data for the ducted fan example. This work was supported in part by AFOSR under a summer faculty program and Grant F49620-00-1-0020, and by NSF under Grant ECS-9625920.

References

[1]Vidyasagar, M., "Optimal Rejection of Persistent Bounded Disturbances," *IEEE Transactions on Automatic Control,* Vol. AC-31, No. 6, 1986, pp. 527–535.

[2]Dorato, P., "A Survey of Robust Multiobjective Design Techniques," *Control of Uncertain Dynamic Systems,* edited by S. P. Bhattacharyya and L. H. Keel, CRC Press, Boca Raton, FL, 1991, pp. 249–259.

[3]Sznaier, M., and Dorato, P., *Proceedings of the 33rd IEEE CDC*, Organizers, special session on "Multiobjective Robust Control," Session FA02, 1994.

[4]Krstic, M., Kanellakopoulos, I., and Kokotovic, P., *Nonlinear and Adaptive Control Design*, Wiley, New York, 1995, Chap. 2.

[5]Artstein, Z., "Stabilization with Relaxed Controls," *Nonlinear Analysis*, Vol. 7, No. 11, 1983, pp. 1163–1173.

[6]Sontag, E. D., "A 'Universal' Construction of Artstein's Theorem on Nonlinear Stabilization," *Systems and Control Letters*, Vol. 13, No. 2, 1989, pp. 117–123.

[7]Freeman, R. A., and Primbs, J. A., "Control Lyapunov Functions, New Ideas from an Old Source," *Proceedings of the 35th IEEE CDC*, 1996, pp. 3926–3931.

[8]Qu, Z., "Robust Control of Nonlinear Uncertain Systems Without Generalized Matching Conditions," *IEEE Transactions on Automatic Control*, Vol. 40, No. 8, 1995, pp. 1453–1460.

[9]van Nieuwstadt, M. J., and Murray, R. M., "Real Time Trajectory Generation for Differentially Flat Systems," *International Journal of Robust and Nonlinear Control*, Vol. 8, No. 11, 1998 pp. 995–1020.

[10]Doyle, J. C., Freeman, R., and Krstic, M., "Nonlinear Control: Comparisons and Case Studies," Workshop #7, 1997 American Control Conf., June 1997.

[11]Cloutier, J. R., and Hull, R. A., "Periodically Preconditioned Conjugate Gradient Restoration Algorithm for Optimal Control: The Hybrid Approach," *Proceedings of the 1997 AIAA Guidance, Navigation, and Control Conference*, AIAA, Reston, VA, 1997, pp. 1–11.

[12]Doyle, J. C., "Nonlinear Benchmarks," *Proceedings of the 35th IEEE CDC*, Session FM02, 1996, pp. 3915–3947.

[13]van der Schaft, A. J., "$\mathcal{L}_2$-Gain Analysis of Nonlinear Systems and Nonlinear State Feedback $\mathcal{H}_\infty$ Control," *IEEE Transactions on Automatic Control*, Vol. 37, No. 6, 1992, pp. 780–784.

[14]Isidori, A., and Wang, K., "$\mathcal{H}_\infty$ Control via Measurement Feedback for General Nonlinear Systems," *IEEE Transactions on Automatic Control*, Vol. 37, No. 9, 1992, pp. 1283–1293.

[15]Ball, J. A., Helton, J. W., and Walker, M. L., "$\mathcal{H}_\infty$ Control for Nonlinear Systems with Output Feedback," *IEEE Transactions on Automatic Control*, Vol. 38, No. 4, 1993, pp. 546–559.

[16]Lu, W. M., and Doyle, J. C., "$\mathcal{H}_\infty$ Control of Nonlinear Systems: A Convex Characterization," *IEEE Transactions on Automatic Control*, Vol. 40, No. 9, 1995, pp. 1668–1775.

[17]Lu, W. M., and Packard, A., "Asymptotic Rejection of Persistent $\mathcal{L}_\infty$ Bounded Disturbances for Nonlinear Systems," *Proceedings of the 35th IEEE CDC*, 1996, pp. 2401–2406.

[18]Mracek, C. P., and Cloutier, J. R., "Control Design for the Nonlinear Benchmark Problem via the State-Dependent Riccati Equation Method," *International Journal of Robust and Nonlinear Control*, Vol. 8, No. 4–5, 1998, pp. 401–433.

[19]Sznaier, M., and Damborg, M., "Heuristically Enhanced Feedback Control of Constrained Discrete Time Linear Systems," *Automatica*, Vol. 26, No. 3, 1990, pp. 521–532.

[20]Rawlings, J. B., and Muske, K. R., "Stability of Constrained Receding Horizon Control," *IEEE Transactions on Automatic Control*, Vol. 38, No. 10, 1993, pp. 1512–1516.

[21]Mayne, D. Q., and Michalska, H., "Receding Horizon Control of Nonlinear Systems," *IEEE Transactions on Automatic Control*, Vol. 35, No. 7, 1990, pp. 814–824.

[22]Michalska, H., and Mayne, D. Q., "Robust Receding Horizon Control of Constrained Nonlinear Systems," *IEEE Transactions on Automatic Control,* Vol. 38, No. 11, 1993, pp. 1623–1633.

[23]Sznaier, M., and Damborg, M. J., "Suboptimal Control of Linear Systems with State and Control Inequality Constraints," *Proceedings of the 26th IEEE CDC,* 1987, pp. 761, 762.

[24]Sontag, E., "Mathematical Control Theory, Deterministic Finite Dimensional Systems," *Texts in Applied Mathematics*, 2nd ed., Vol. 6, Springer–Verlag, New York, 1998.

[25]Freeman, R. A., and Kokotovic, P. V., "Inverse Optimality in Robust Stabilization," *SIAM Journal on Control and Optimization*, Vol. 34, July 1996, pp. 1365–1391.

[26]Freeman, R. A., and Konotovic, P. V., *Robust Nonlinear Control Design: State Space and Lyapunov Techniques*, Birkhauser Boston, Cambridge, MA, 1996.

[27]Doyle, J. C., Primbs, J. A., Shapiro, B., and Nevistic, V., "Nonlinear Games: Examples and Counterexamples," *Proceedings 35th IEEE CDC*, 1996, pp. 3915–3920.

[28]Jacobson, D. H., *Extensions of Linear Quadratic Control, Optimization and Matrix Theory*, Academic, London, 1977.

[29]Khalil, H. K., *Nonlinear Systems*, 2nd ed., Prentice–Hall, Upper Saddle River, NJ, 1996.