IEEE0343

Stabilizing Receding Horizon Control of Nonlinear Systems: A Control Lyapunov Function Approach

Ali Jadbabaie, Jie Yu, and John Hauser [1]
Control and Dynamical Systems
Mail Code 107-81
California Institute of Technology
Pasadena, CA 91125, USA
{alij,jyu,hauser}@cds.caltech.edu

Abstract

In this paper, a modified version of the receding horizon control of nonlinear systems is proposed. The new approach is based on a finite horizon optimal control problem with an end point penalty. This method can be treated as an extension of recent results of De Nicolao *et al.* to the case where a globally stabilizing control law is available. Similar to the afore mentioned result, the end point penalty is the cost-to-go if a globally stabilizing control law is applied. The control law is computed *a priori* using a Control Lyapunov Function. It is further shown that stability is maintained if the vailable CLF is used as an end point penalty instead of the actual cost. The results are illustrated by applying this approach to the planar model of a ducted fan with a CLF obtained using Quasi LPV methods.

Keywords: Receding Horizon Control, nonlinear control design, optimal control, LPV methods, Control Lyapunov Functions.

1 Introduction

Recently there has been a rapidly growing interest in using receding horizon control schemes for control of nonlinear systems. This interest is mainly due to availability of faster and cheaper computers. This approach is based on solving a finite horizon optimal control problem over an interval $[t, t+T]$, and then applying a portion of the computed control action from $[t, t+\delta]$ where $\delta < T$. Successive application of this strategy will result in a feedback control law since the control action depends on the current state. Due to the fact that an optimization over finite horizon does not guarantee stability, application of this strategy requires some caution. Several researchers have come up with different methods to guarantee the closed-loop stability of the Receding horizon scheme, some of which will be summarized in what follows:

[1] On leave from University of Colorado at Boulder

Keerthi and Gilbert [5] imposed a terminal state equality constraint $x(t+T) = 0$. This would result in a finite-horizon optimization problem which turns out to be computationally demanding. Michalska and Mayne [8] ensured closed-loop stability by requiring that $x(t+T)$ enters a suitable neighborhood of the origin and then the control law is switched to a locally stabilizing linear controller. Another approach proposed by Parsini and Zopolli [12] and later by Chen and Allgöwer [3], is based on using a quadratic endpoint penalty of the form $ax(t+T)^T Px(t+T)$ for some $a > 0$ and some positive definite matrix P. In a more recent paper by De Nicolao *et al.* [10] stability of the receding horizon controller is guaranteed by using a possibly non quadratic end point penalty which is the cost incurred if a locally stabilizing linear control law is applied at the end of the time horizon T. The linear control law ensures local exponential stability of the equilibrium at $x = 0$, and it is assumed that the region of attraction of the linear controller is large enough that can be reached from the initial condition within the time interval $[0, T]$. However, this assumption may not be necessarily true. A different approach to receding horizon control, is due to Primbs *et al.* [13, 14]. In this approach, first a globally stabilizing control law is achieved by finding a global Control Lyapunov Function (CLF). Once the global CLF is obtained, stability of the receding horizon controller is guaranteed by including additional state constraints that require the derivative of the CLF along the receding horizon trajectory be negative and also that the decrease in the value of the CLF be less than the decrease if a controller obtained from the CLF is used instead. This approach is very attractive in a sense that guarantees global stability and it does not require some of the assumptions that the previous methods do. Also, depending on the structure of the nonlinear system, there are variety of methods for obtaining a global CLF. For example, if the system is feedback linearizable, feedback linearization techniques can be used to obtain the CLF, or if the system is in block feedback form [6], back stepping techniques can be used to obtain a CLF. Moreover, one can use Quasi LPV methods for a large class of affine in control nonlinear systems. However, the online computations for a finite horizon optimal control problem with state inequality constraints seem to be computationally expensive.

The approach used in this paper is a hybrid of the methods of Primbs *et al.* [13, 14] and De Nicolao *et al.*, i.e., similar to [10], an endpoint penalty is used which is equal to the cost incurred if the globally stabilizing control law obtained from the CLF is used after the end of horizon. An important note is that this controller is actually never implemented. It is merely used to calculate the cost that one would get if this controller was to be applied. To avoid calculating the cost incurred by applying the nonlinear control law, it is shown that the available CLF can be used as a terminal penalty instead. Removing the state inequality constraints and putting a penalty would lower the computation burden quite significantly. The effectiveness of this method is finally illustrated on the planar model of the Caltech ducted fan, which is an experimental testbed for flight control experiments.

This paper is organized as follows: In section 2 we review the definition of a control

Lyapunov function. In section 3, we present the receding horizon scheme with the terminal penalty when a globally stabilizing controller is available. In section 4, the terminal penalty is replaced by the Lyapunov function and it is shown that the receding horizon scheme will indeed guarantee stability. Furthermore, if the Lyapunov function is globally stabilizing, the receding horizon control law will be globally stabilizing as well. Section 5 deals with the special case that the nonlibeare system is affine in control. We use a Quasi LPV scheme to generate a Control Lyapunov function. To illustrate the effectiveness of this approach, the receding horizon strategy is applied to the planar model of Caltech ducted fan. Simulation results are presented in Section 6. Conclusions and some future work are discussed in Section 7.

2 Control Lyapunov Functions

Consider the nonlinear system

$$\dot{x} = f(x, u) \tag{1}$$

where $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, and $\dot{V}(x, u) = V_x(x).f(x, u)$. A *control Lyapunov function* (CLF) is a $\mathbf{C}^1$, proper, positive definite function $V : \mathbb{R}^n \to \mathbb{R}_+$ such that:

$$\inf_u \left[\dot{V}(x, u)\right] < 0 \tag{2}$$

If it is possible to make the derivative negative at every point by an appropriate choice of u, then we have achieved our goal and can stabilize the system with V a Lyapunov function for the closed loop that we choose. This is exactly the condition given in (2). It can be shown that the existence of a CLF for the system (1) is equivalent to the existence of a globally asymptotically stabilizing control law $u = k(x)$ which is smooth everywhere except possibly at $x = 0$ [1]. Moreover, one can calculate such a control law k when the system is affine in control, i.e. $f(x, u) = \tilde{f}(x) + g(x)u$, explicitly from $\tilde{f}$, g and V [16]. In the next section, we assume that a CLF and a stabilizing feedback have already been obtained. We also define a q-CLF for future use. consider a positive-definite function $q(x, u)$. A $\mathbf{C}^1$, proper, positive definite function $V : \mathbb{R}^n \to \mathbb{R}_+$ such that:

$$\inf_u \left[\dot{V}(x, u) + q(x, u)\right] \leq 0 \tag{3}$$

is called a q CLF. We proceed by presenting the receding horizon control with end point penalty in the next section.

3 Receding Horizon Control with End Point Penalty

Consider the following Finite Horizon Optimal Control Problem

$$J^*(t, x, T) = \inf_{u_{[t,t+T]}} \int_t^{t+T} q(x(\tau), u(\tau))d\tau + M(x(t + T))$$

$$s.t. \qquad \dot{x} = f(x,u)$$

Where $q(\cdot,\cdot)$ is a positive definite function of $x(\cdot), u(\cdot)$. The resulting open loop control trajectory $u^*(\cdot)$ is implemented until a new state update occurs, usually at pre-specified sampling intervals. Repeating these calculations yields a feedback control law. In other words, we are trying to find a feedback control law that would result in a cost close to solution of the infite horizon optimal control problem with a reasonable amount of computation. The premise of our paper is the same as that of Primbs *et al.* [13, 14] in the sense that assuming a Control Lyapunov Function (CLF) is already available, one can use the CLF in conjunction with a receding horizon strategy to further improve the performance without losing guarantees for stability.

To ensure closed-loop stability, we generalize the method of De Nicolao *et al.* by picking the cost incurred if the Quasi-LPV control law were to be applied at the end of the time horizon as endpoint penalty, i.e.

$$\begin{aligned}
M(x(t+T)) &= \int_{t+T}^{\infty} q(\phi_2(\tau; t+T, x^*(t+T), k(\cdot)), k(\phi_2(\tau; t+T, x^*(t+T), k(\cdot))))d\tau \\
x^*(t+T) &= \phi_1(t+T, t, x(t), u^*(\cdot; t, x, T) \\
u^*(\cdot; t, x, T) &= arg \inf_{u_{[t,t+T]}(\cdot)} J^*(t, x, T, u(\cdot), k(\cdot))
\end{aligned} \tag{4}$$

where $J(t,x,T,u,k)$ is the cost incurred if the controller u is applied for time $[t, t+T]$ and the cost associated with applying controller $u^k(x)$ from $[t+T, \infty)$ is used as a penalty. Also ϕ_1 is the flow of the vector field along the open-loop receding horizon trajectory u^*, and ϕ_2 is the flow along the feedback control law $k(\cdot)$ obtained in advance. One can also use Sontag's formula to obtain the feedback law directly from the CLF [16, 18] instead of using the stabilizing control law. As was mentioned earlier, this end point penalty is the cost incurred if the stabilizing controller is applied for the rest of the time after the end of the horizon. Using the argument presented in Lemma 3 and theorem 1 of [10] as well as proposition 1 in [7], we have the following theorem:

Theorem 1 *The Receding Horizon Optimal Control scheme (4) is asymptotically stabilizing if a stabilizing feedback control law $u^k(x)$ is available.*

Proof : Let $J^*(t,x,T)$ be the cost associated with the solution of the FHOCP introduced earlier. Our goal is to prove that $J^*(t,x,T)$ will qualify as a Lyapunov function for the closed loop system. Construct the following sub-optimal control strategy for the time interval $[t, t+T+\delta]$.

$$\tilde{u} = u^*_{[t,t+T]} + u^k_{[t+T,t+T+\delta]} \tag{5}$$

We have the following relation between J^* and J:

$$J(t+\delta, x^*(t+\delta), T, \tilde{u}(\cdot), k(\cdot)) = J^*(t,x,T) - \int_t^{t+\delta} q(\phi_1(\tau; t, x, u^*(\cdot)), u^*(\tau))d\tau$$

However, since $\tilde{u}$ is a suboptimal control strategy, we have the following:

$$J^*(t+\delta, x^*(t+\delta), T) \leq J(t+\delta, x^*(t+\delta), T, \tilde{u}(\cdot), k(\cdot)) \tag{6}$$

Therefore, we can conclude the following:

$$J^*(t+\delta, x^*(t+\delta), T) - J^*(t, x, T) \leq -\int_t^{t+\delta} q(x^*(\tau), u^*(\tau))d\tau \tag{7}$$

deviding both sides by δ and taking the $\limsup$ as δ goes to zero, we get

$$\limsup_{\delta\to 0} \frac{J^*(t+\delta, x^*(t+\delta), T) - J^*(t, x, T)}{\delta} \leq -q(x, u^*(t)) \leq -C_1||(x, u^*(t))||^2 \leq -C_1||x||^2 \tag{8}$$

where the last two inequalities follow from the assumption that the preformance index $q(\cdot,\cdot)$ is positive definite. Also, by definition of J^*, we have $J^*(t, 0, T) = 0\ \forall t,\ T > 0$. Therefore, we have shown that J^* qualifies as a Lyapunov function and asymptotic stability is guaranteed. Moreover, if the feedback control lawu^k is globally stabilizing, i.e., $V(x)$ is a global CLF, the above results are global. ■

4 CLF as Endpoint Penalty

Despite the fact that the argument given in the previous section ensures closed loop stability, calculation of the cost-to-go requires integration of the closed-loop system with the stabilizing controller $u^{k(\cdot)}(x)$ over a possibly long period. This would further increase the burden of online computation. One possible remedy would be to replace the actual cost-to-go with the CLF obtained from one of the several available methods. However, the question that one could raise is does there exist a stabilizing controller whose cost is equal to the value of CLF? One can argue that if there exist a controller that would incur a cost equal to the value of CLF, the same argument presented in the proof of Theorem 1 holds. Note that mere existence of such a controller is sufficient for the proof of Theorem 1 to remain valid. This is plausible specially when the CLF is an upper bound to the cost-to-go using the controller, since it is always possible to generate a controller that would make the cost worse, i.e., a controller which would perform poorer. Nevertheless even if such a controller does not exist, the stability of the receding horizon scheme can be guaranteed with the following theorem.

Theorem 2 *Consider the FHOCP (4) with the following terminal penalty:*

$$M(x(t+T)) = \mu V(x(t+T)) \tag{9}$$

where $V(\cdot)$ is a CLF obtained a priori and μ is a design parameter. Denote the optimal cost associated with this problem by J_V^*. Then there exists a ρ_0 such that for all $\mu \geq \mu_0$, J_V^* is a Lyapunov function for the closed loop system with the receding horizon feedback. Furthermore, if $V(\cdot)$ is a global CLF, so is J_V^*.

Proof: The proof argument is along the proof of Theorem 1. Again, we generate a sub-optimal control strategy $\tilde{u}$. The cost associated with applying $\tilde{u}$ given in (5) satisfies the following:

$$\begin{aligned} J_V^*(t,x,T) &= J_V(t+\delta, x^*(t+\delta), T, \tilde{u}(\cdot)) + \int_t^{t+\delta} q(x^*(t+\tau), u^*(\cdot))d\tau + \\ &\quad \mu\left[V(x^*(t+T)) - V(\phi_2(t+T+\delta; t+T, x^*(t+T), k(\cdot)))\right] \\ &- \int_{t+T}^{t+T+\delta} q(\phi_2(\tau - t - T; t+T, x^*(t+T), k(x^*(t+T)), k(\cdot))d\tau \quad (10) \end{aligned}$$

Since J_V is the sub- optimal strategy, $J_V^* \leq J_V$. Replacing J_V by the optimal value, and deviding both sides by δ and letting $\delta \to 0$, we obtain the following:

$$\begin{aligned} \limsup_{\delta\to 0} \frac{J_V^*(t+\delta, x^*(t+\delta), T) - J_V^*(t,x,T)}{\delta} &\leq -q(x, u^*(t)) \quad (11) \\ &+ \mu\dot{V}(x^*(t+T)) + q(x^*(t+T), k(x^*(t+T))) \end{aligned}$$

If $J_V^*(t,x,T)$ is $\mathbf{C^1}$ in t and x, then we can write the above inequality as:

$$\dot{J}_V^*(t,x,T) \leq -q(x,u^*(t)) + \mu\dot{V}(x^*(t+T)) + q(x^*(t+T), k(x^*(t+T))) \quad (12)$$

Since $V(x)$ is a Lyapunov function for the closed-loop system, we know that $\dot{V} \leq -C_3||x(\cdot)||^2$ for some $C_3 > 0$. Furthermore, $q(x,k(\cdot)) < C_4||x||^2$ for some $C_4 > 0$. Hence, we can always pick $\mu_0 > 0$ such that $\mu_0\dot{V}(x^*(t+T)) + q(x^*(t+T), k(\cdot))\forall x$. Note that $-q(x,u^*(t)) < -C_1||x||^2$. If V is a global CLF, then it also follows that J_V^* will be a global CLF as well. ■

Remark 1 *By picking the design parameter μ to be a large value ($\mu >> 1$), the penalty is essentially dominating the cost. For small values of μ, we have the receding horizon scheme without the penalty.*

Remark 2 *Note that the above argument is independent of the method which gerated the CLF.*

Remark 3 *If $V(\cdot)$ is a q-CLF, then an obvious choice of $\mu_0 = 1$ would guarantee the stability..*

In the next section, we use Quasi LPV methods to generate a q-CLF. Once the q-CLF is obtained the method discussed above can be used to improve on the cost.

4.1 Linear Parameter Varying (LPV) Methods

It can be shown that using LPV methods [17] the problem of minimizing an upper bound on the quadratic performance index

$$\int_0^\infty (x(\tau)^T Qx(\tau) + u(\tau)^T u(\tau))d\tau \quad (13)$$

subject to

$$\dot{x} = A(\rho(x))x + B(\rho(x))u \tag{14}$$

can be converted into an infinite dimensional convex optimization problem. Assume the underlying parameter ρ varies in the allowable set

$$\mathcal{F}_{\mathcal{P}}^{\nu} := \{\rho \in \mathcal{C}^1(\mathrm{I\!R}^+, \mathrm{I\!R}^{\mathrm{m}}) : \rho \in \mathcal{P},\ \underline{\nu}_{\mathrm{i}}(\rho) \leq \dot{\rho}_{\mathrm{i}} \leq \overline{\nu}_{\mathrm{i}}(\rho),\ \mathrm{i} = 1, \cdots, \mathrm{m}\} \tag{15}$$

where $\mathcal{P} \subset \mathrm{I\!R}^{\mathrm{m}}$ is a compact set. If there exists a positive definite $X(\rho)$ such that the following convex optimization is solvable:

$$\min_{X(\rho)>0, Z} tr(Z) \tag{16}$$

Subject to

$$\begin{bmatrix} -\sum_{i=1}^{m} \overline{\underline{\nu}}_i(\rho)\frac{\partial X}{\partial \rho_i} + A(\rho)X(\rho) + X(\rho)A^T(\rho) - B(\rho)B^T(\rho) & C^T(\rho)X(\rho) \\ X(\rho)C(\rho) & -I \end{bmatrix} < 0 \tag{17}$$

and

$$\begin{bmatrix} Z & I \\ I & X(\rho) \end{bmatrix} > 0 \tag{18}$$

for all $\rho \in \mathcal{P}$ where $C(\rho) = Q^{\frac{1}{2}}(\rho(x))$, then the closed-loop system is stable with the state feedback

$$u(x) = -B^T(\rho(x))X^{-1}(\rho(x))x \tag{19}$$

Moreover, we have an upper bound on the actual cost incurred by utilizing the above state feednback law.

$$V(x) = x^T X^{-1}(\rho(x))x \geq \int_0^{\infty} (x(\tau)^T Q x(\tau) + u(\tau)^T u(\tau))d\tau.$$

The notation $\sum_{i=1}^{m} \overline{\underline{\nu}}_i(\rho)$ in (17) means that every combination of $\overline{\nu}_i(\rho)$ and $\underline{\nu}_i(\rho)$ should be included in the inequality. For instance, when $m = 2$, $\overline{\nu}_1(\rho) + \overline{\nu}_2(\rho)$, $\overline{\nu}_1(\rho) + \underline{\nu}_2(\rho)$, $\underline{\nu}_1(\rho) + \overline{\nu}_2(\rho)$ and $\underline{\nu}_1(\rho) + \underline{\nu}_2(\rho)$ should be checked individually. In other words, (17) actually represents 2^m inequalities. Solving (17) involves riding the parameter space $\mathcal{P}$ and choosing a finite set of basis for $X(\rho)$ (see [17] for details).

5 Caltech Ducted Fan Model

The Caltech Ducted Fan is a small flight control experiment whose dynamics are representative of either a Harrier in hover mode or a thrust vectored aircraft such as F18-HARV or X-31 in forward flight [9]. This system has been used for a number of studies and papers. In particular, a comparison of several linear and nonlinear controllers was performed in [4, 2, 11]. In this section we describe the simple planar model

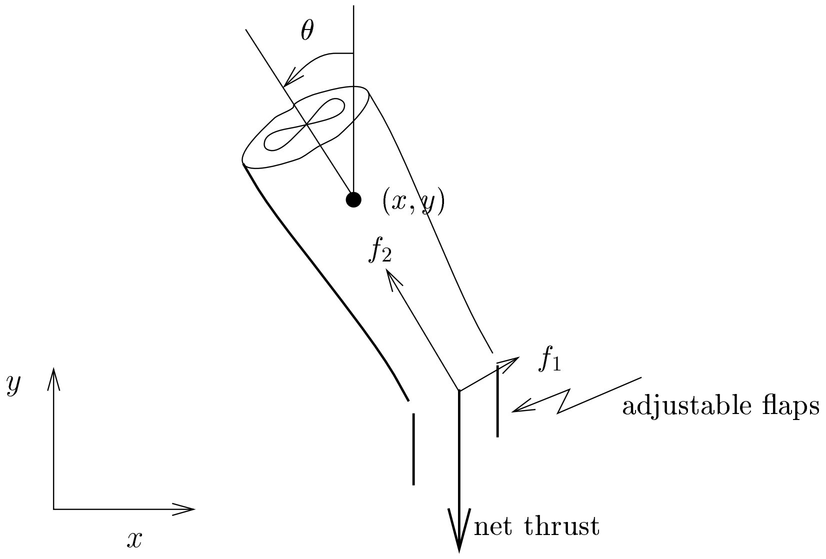

Figure 1: Planar ducted fan model

of the fan shown in Figure 1 which ignores the stand dynamics. This model is useful for initial controller design and would serve as a good testbed for purposes of this paper. Let (x, y, θ) denote the position and orientation of a point on the main axis of the fan that is distance l from the center of mass. We assume that the forces acting on the fan consist of a force f_1 perpendicular to the axis of the fan acting at a distance r and a force f_2 parallel to the axis of the fan. Assuming m, J, and g to be the mass of the fan, the moment of inertia, and the gravitational constant respectively, the equations of motion can be written as

$$\begin{aligned} m\ddot{x} &= -d\dot{x} + f_1\cos\theta - f_2\sin\theta \\ m\ddot{y} &= -d\dot{y} + f_1\sin\theta + f_2\cos\theta - mg \\ J\ddot{\theta} &= rf_1 \end{aligned} \tag{20}$$

where the drag terms are modeled as viscous friction with d being the viscous friction coefficient. The objective is to stabilize the system to a hover configuration. It is convenient to redefine the inputs so that the origin is an equilibrium point of the system with zero input. If we let $u_1 = f_1$ and $u_2 = f_2 - mg$.

$$\begin{aligned} m\ddot{x} &= -mg\sin\theta - d\dot{x} + u_1\cos\theta - u_2\sin\theta \\ m\ddot{y} &= mg(\cos\theta - 1) - d\dot{y} + u_1\sin\theta + u_2\cos\theta \\ J\ddot{\theta} &= ru_1 \end{aligned} \tag{21}$$

These equations are referred to as the *planar ducted fan equations.* The following quadratic cost function was used for comparison of different design techniques:

$$J = \int_0^\infty (x^T(t)Qx(t) + u^T u(t))dt$$

where Q is chosen to be

$$Q = diag(\begin{bmatrix} 10 & 1 & 10 & 1 & 1 & 1 \end{bmatrix}) \tag{22}$$

and

$$x(t) = \begin{bmatrix} x & \dot{x} & y & \dot{y} & \theta & \dot{\theta} \end{bmatrix} \tag{23}$$

are the states.We illustrate the performance of the closed loop system from the challenging initial condition

$$\begin{bmatrix} x & \dot{x} & y & \dot{y} & \theta & \dot{\theta} \end{bmatrix} = \begin{bmatrix} 5 & 5 & 5 & 0 & -0.9\pi/2 & 0 \end{bmatrix} \tag{24}$$

To obtain a CLF for the ducted fan using quasi LPV methods, we have to pick the varying parameter and a bound on parameter variation. For the ducted fan, $\rho = \theta$ is chosen as the varying parameter, and the operation range $\mathcal{P} = [-\frac{\pi}{2}, \frac{\pi}{2}]$. The bound on the rate variation on θ is set to be 10, i.e., $|\dot{\theta}| \leq 10$. Obvious parametrizations of f and g are given by setting $A(\rho)$ and $B(\rho)$ as follows:

$$A(\theta) = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & -\frac{d}{m} & 0 & 0 & -\frac{g \sin\theta}{\theta} & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -\frac{d}{m} & \frac{g(\cos\theta - 1)}{\theta} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad B(\theta) = \begin{bmatrix} 0 & 0 \\ \frac{\cos\theta}{m} & -\frac{\sin\theta}{m} \\ 0 & 0 \\ \frac{\sin\theta}{m} & \frac{\cos\theta}{m} \\ 0 & 0 \\ \frac{r}{j} & 0 \end{bmatrix} \tag{25}$$

A set of basis functions are chosen to compute $X(\rho)$, i.e., $X(\rho) = \sum_{i=1}^{5} c_i(\rho)X_i$ where X_i's are symmetric coefficient matrices (LMI variables) and $\{c_i(\rho)\}$ are fifth order Legendre polynomials on $\mathcal{P}$:

$$\{c_i(\rho)\} = \{1,\ \frac{2}{\pi}\theta,\ (3(\frac{2}{\pi}\theta)^2 - 1)/2,\ (5(\frac{2}{\pi}\theta)^3 - 3(\frac{2}{\pi}\theta))/2,\ (35(\frac{2}{\pi}\theta)^4 - 30(\frac{2}{\pi}\theta)^2 + 3)/2\}$$

Using the above CLF and a horizon length of 0.8 with a sampling time of 0.1, a total cost of 1382.5 was obtained, which is close to the cost of the infinite horizon problem which was 1363. The simulation results for the closed-loop system is shown in Figures 2 and 3.

6 Conclusions

The purpose of this paper was to extend some recent results for the stability of the receding horizon control scheme. It was shown that by using a specific end point penalty, global stability of receding horizon scheme is guaranteed. The results are based on the assumption that a (globally) stabilizing controller has already been obtained using the CLF obtained by Quasi-LPV methods.The end point penalty is the cost that would be incurred if the afore mentioned controller is applied at the end of horizon. Furthermore, it was shown that the end point penalty can be replaced by its upper bound. This algorithm was applied to the planar model of a ducted fan, and the results were shown to be satisfactory. One interesting direction for further research would be extending these results to the trajectory tracking problem as well as including input constraints.

References

[1] Z. Artstein. Stabilization with relaxed controls. *Nonlinear Anal.*, 7(11):1163–1173, 1983.

[2] B. Bodenheimer, P. Bendotti, and M. Kantner. Linear parameter-varying control of a ducted fan engine. *Int.J.Robust and Nonlinear Control*, 6:1023–1044, Nov 1996.

[3] H. Chen and F. Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. In *Procedding ECC,97*, 1997.

[4] M. Kantner, B. Bodenheimer, P.Bendotti, and R. M. Murray. An experimental comparison of controllers for a vectored thrust,ducted fan engine. In *Proceedings of the 1995 American Control Conference*, pages 1956–1961, Seattle, 1995.

[5] 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.

[6] M. Kristić, I. Kanellakopoulos, and P. Kokotović. *Nonlinear and Adaptive Control Design*. John Wiley & Sons, New York, 1995.

[7] L. Magni and R. Sepulchre. Stability margins of nonlinear receding-horizon control via inverse optimality. *Syst. and Cont. Lett.*, 32:241–245, 1997.

[8] H. Michalska and D.Q. Mayne. Robust receding horizon control of constrained nonlinear systems. *IEEE Trans. Aut. Control*, 38(11):1623–1633, November 1993.

[9] R.M. Murray. Modeling of the caltech ducted fan, class notes for cds 111. Technical report, California Institute of Technology, Control and Dynamical Systems 107–81, Pasadena, CA 91125, 1998.

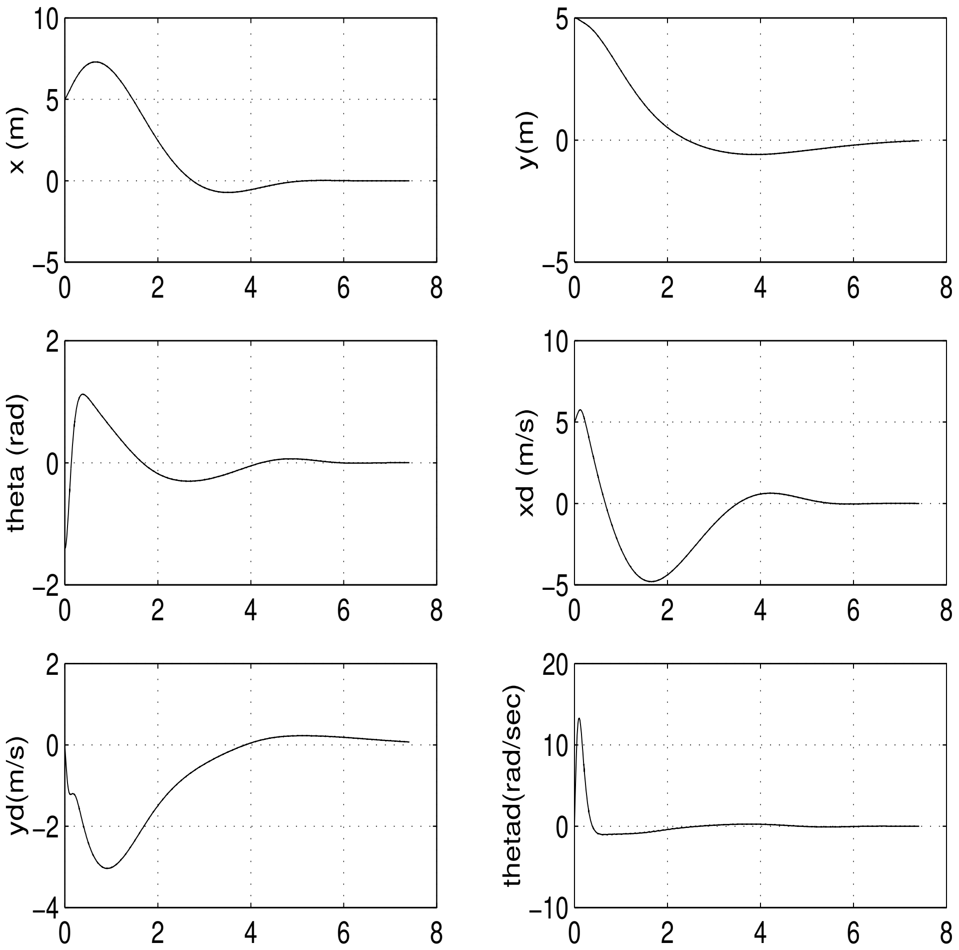

Figure 2: Initial condition response of the fan.

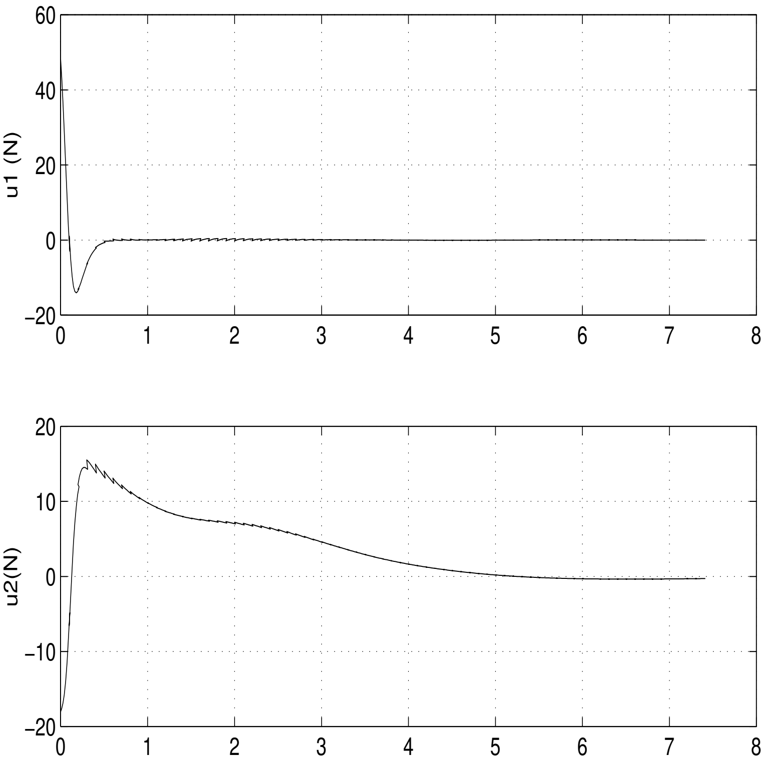

Figure 3: Control action.

[10] G. De Nicolao, L. Magni, and R. Scattolini. Stabilizing receding-horizon control of nonlinear time-varying systems. *IEEE Tran. Aut. Cont.*, 43(7):1030–1036, July 1998.

[11] M.van Nieuwstadt and R.M.Murray. Real time trjectory generation for differentially flat systems. In *Proceedings of the 1996 IFAC World Congress*, 1996.

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

[13] J. A. Primbs, V. Nevistić, and John C. Doyle. On Control Lyapunov Functions and Receding Horizon Control: Part I Background. *Submitted to IEEE TAC*, 1998.

[14] J. A. Primbs, V. Nevistić, and John C. Doyle. On Control Lyapunov Functions and Receding Horizon Control: Part II a receding horizon extension of pointwise min-norm controllers. *Submitted to IEEE TAC*, 1998.

[15] E.D. Sontag. A lyapunov-like characterization of asymptotic controllability. *SIAM J. Contr. Optimiz.*, 21(3):462–471, 1983.

[16] E.D. Sontag. A 'universal' construction of Artstein's theorem on nonlinear stabilization. *Syst. Contr. Lett.*, 13(2):117–123, 1989.

[17] Fen Wu, Hua Xin Yang, Andy Packard, and Greg Becker. Induced L_2-Norm control for lpv systems with bounded parameter variation. *Int. Journal of Nolinear and Robust Control*, vol.6:983–998, 1996.

[18] J. Yu, A. Jadbabaie, J. Primb, and Y. Huang. Comparison of nonlinear control designs for a ducted fan model. 1998. Submitted to IFAC World Congress.