Receding Horizon Control of The Caltech Ducted Fan: A Control Lyapunov Function Approach.*

Ali Jadbabaie†, Jie Yu†, and John Hauser‡
{alij,jyu,hauser}@cds.caltech.edu

Abstract

This paper deals with the application of receding horizon methods to the simplified model of a flight control experiment developed at Cal Tech. The dynamics of the system are representative of a Vertical Landing and Take off (VTOL) aircraft, such as a Harrier around hover. The adopted control methodology is a hybrid of receding horizon techniques and Control Lyapunov Function (CLF) based ideas. First a CLF is found, and then, by using the CLF as the terminal cost in the receding horizon optimization, stability is guaranteed. It is shown that if the horizon length is long enough, a simple choice of CLF, i.e., the one obtained from Jacobian linearization of the dynamics at hover, will achieve a good performance. However, if due to computational costs, longer time horizons are not possible, stability can be guaranteed by applying the CLF obtained using Quasi-LPV methods, as a terminal costs. Several numerical simulations for different time horizons are presented to illustrate the effectiveness of the discussed 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 methods. This interest is due to the fact that online computation has become cheaper and more efficient. In addition, numerical methods for solving optimal control problems have become more reliable. In view of the fact that despite these advances solving the exact optimal control problem, i.e., solving the Hamilton Jacobi equation is still prohibitively difficult, these sub optimal control strategies have become more and more attractive. The main idea behind receding horizon techniques is to turn the infinite horizon optimal control problem into a series of finite horizon optimizations. After the finite horizon optimization is solved by solving a two point boundary value problem, the first portion of the resulting control trajectory is applied. Repeating this process iteratively will result in a feedback law. However, since its introduction to the literature, receding horizon techniques have been suffering from lack of reasonable stability guarantees. Several researchers have addressed the stability problem and reported some success.

Keerthi and Gilbert [5] imposed a terminal state equality constraint $x(t+T) = 0$ where T is the horizon length. This would result in a finite-horizon optimization problem which turns out to be computationally demanding. Michalska and Mayne [7] 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 [11] and later by Chen and Allgöwer [2], 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.* [9] 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. One recent approach to the receding horizon problem is due to Primbs *et al.* [12, 13]. In this approach, stability is guaranteed by imposing a CLF based constraint on the optimization. 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 nega-

*Research supported in part by AFOSR and DARPA.

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

‡Electrical and Computer Engineering, University of Colorado, Boulder, CO 80309-0425; currently on leave at Caltech 104-44, Pasadena, CA 91125.

tive 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. One possible remedy to this problem was introduced by Jadbabaie *et al.* [3]. In this approach stability is guaranteed by using a CLF which is an upper bound on the cost-to-go. One possible choice of a CLF that has this property is the CLF obtained using LPV methods. This paper is focused on using LPV methods to generate a CLF and then Using that CLF as a terminal cost in receding horizon optimization. The results are applied to the simplified model of the Cal Tech Ducted Fan experiment. To make the problem more challenging, an initial condition is picked that would make the stabilization task harder.

This paper is organized as follows: In section 2, an overview of receding horizon method is resented and two choices for the terminal cost are discussed. In Section 3, a brief overview of Quasi LPV methods is presented, and the procedure of obtaining a CLF is briefly discussed. Section 4 deals with the description of the simplified model of the Caltech Ducted experiment around hover point. Simulations are presented in section 5. Finally the conclusions are presented in section 6 and some future work is discussed.

2. Receding Horizon Control with Terminal Cost

Consider the following Finite Horizon Optimal Control Problem

$$\begin{aligned} J^*(t,x,T) &= \inf_{u_{[t,t+T]}} \int_t^{t+T} q(x(\tau),u(\tau))d\tau + M(x(t+T)) \\ s.t.\ \dot{x} &= f(x,u) \end{aligned}$$

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 infinite horizon optimal control problem with a reasonable amount of computation. Following the results presented in [3], the premise of this paper is that if 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. The following theorem is a generalization of the results of De Nicolao *et al.* [9], and was proved in [3].

Theorem 1 *The Receding Horizon Optimal Control scheme (1) is asymptotically stabilizing if the cost incurred by applying a stabilizing control law obtained from the CLF is used as terminal cost.*

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. The following theorem [3] illustrates this point.

Theorem 2 *Consider the finite horizon optimal control problem (1) with the following terminal penalty:*

$$M(x(t+T)) = \mu V(x(t+T)) \tag{1}$$

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 μ_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^*.

Remark 1 *If $V(\cdot)$ is a CLF that is also an upper bound on the cost, i.e., $\dot{V} + q < 0$, then $\mu = 1$ is an obvious choice.*

A very interesting method of obtaining a CLF for a nonlinear systems that are affine in input and are point-wise linearly controllable, is Linear Parameter Varying methods. One major advantage of the CLF obtained using LPV methods is that it has the property described in Remark 1, i.e., the Value of the CLF is an upper bound on the cost-to-go. The following section gives a brief introduction to LPV methods.

2.1. Linear Parameter Varying (LPV) Methods

It can be shown that using LPV methods [14] the problem of minimizing an upper bound on the quadratic performance index

$$\int_0^\infty (x(\tau)^T Q x(\tau) + u(\tau)^T u(\tau)) d\tau \tag{2}$$

subject to

$$\dot{x} = A(\rho(x))x + B(\rho(x))u \tag{3}$$

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

$$\mathcal{F}_{\mathcal{P}}^{\nu} := \{\rho \in \mathcal{C}^1(\mathrm{I\!R}^+, \mathrm{I\!R}^{\mathrm{m}}) : \rho \in \mathcal{P},\ \underline{\nu}_{\mathrm{i}}(\rho) \leq \dot{\rho}_{\mathrm{i}} \leq \overline{\nu}_{\mathrm{i}}(\rho)\} \tag{4}$$

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:

$$\min_{X(\rho)>0, Z} tr(Z) \tag{5}$$

Subject to

$$\begin{bmatrix} \mathcal{L}(X(\rho)) & C^T(\rho)X(\rho) \\ X(\rho)C(\rho) & -I \end{bmatrix} < 0 \tag{6}$$

and

$$\begin{bmatrix} Z & I \\ I & X(\rho) \end{bmatrix} > 0 \tag{7}$$

Where

$$\mathcal{L}(X(\rho)) = -\sum_{i=1}^{m} \underline{\overline{\nu}}_i(\rho) \frac{\partial X}{\partial \rho_i} + A(\rho)X(\rho) + X(\rho)A^T(\rho) - B(\rho)B^T(\rho) \tag{8}$$

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

$$u(x) = -B^T(\rho(x))X^{-1}(\rho(x))x \tag{9}$$

Moreover, we have an upper bound on the actual cost incurred by utilizing the above state feedback 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 $\sum_{i=1}^m \underline{\overline{\nu}}_i(\rho)$ in (6) 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, (6) actually represents 2^m inequalities. Solving (6) involves riding the parameter space $\mathcal{P}$ and choosing a finite set of basis for $X(\rho)$ (see [14] for details).

3. 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 [8]. 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, 1, 10]. In this section we describe the simple planar 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. 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 follows:

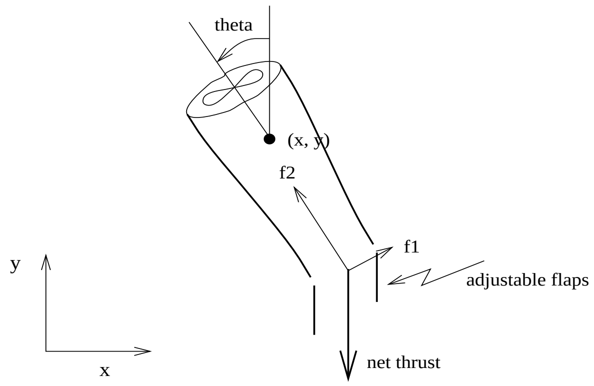

Figure 1: Planar ducted fan model

$$\begin{aligned} m\ddot{x} &= -d\dot{x} + f_1 \cos\theta - f_2 \sin\theta \\ m\ddot{y} &= -d\dot{y} + f_1 \sin\theta + f_2 \cos\theta - mg \\ J\ddot{\theta} &= r f_1 \end{aligned} \tag{10}$$

where the drag terms are modeled as viscous friction with d being the viscous friction coefficient. The numerical data for the experiment are as follows: $m = 11.2kg$, $g = 0.28m/sec^2$, $J = 0.0462kgm^2$, $r = 0.156m$, $d = 0.1Nsec$. Note that to the use of a counter weight for balancing the fan, the gravitational constant is not $9.8m/s^2$. 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 $u_2 = f_2 - mg$.

$$
\begin{aligned}
m\ddot{x} &= -mg\sin\theta - d\dot{x} + u_1\cos\theta - u_2\sin\theta \\
m\ddot{y} &= mg(\cos\theta - 1) - d\dot{y} + u_1\sin\theta + u_2\cos\theta \\
J\ddot{\theta} &= ru_1 \qquad (11)
\end{aligned}
$$

These equations are referred to as the *planar ducted fan equations.* The following quadratic cost function was used:

$$
\mathcal{J} = \int_0^\infty (\bar{x}^T(t)Q\bar{x}(t) + u^T(t)Ru(t))dt
$$

where $\bar{x} = [x\ \dot{x}\ y\ \dot{y}\ \theta\ \dot{\theta}]^T$, R and Q is chosen to be a diagonal matrix with the following diagonal terms:

$$
\begin{aligned}
Q &= diag([\ 0.5\quad 1\quad 5\quad 1\quad 5\quad 1\ ]) \qquad (12) \\
R &= diag([\ 2\quad 1\ ]) \qquad (13)
\end{aligned}
$$

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 parameterizations of f and g are given by setting $A(\rho)$ and $B(\rho)$ as follows:

$$
A(\theta) = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & -\frac{d}{m} & 0 & 0 & -\frac{g\sin\theta}{\theta} & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -\frac{d}{m} & \frac{g(\cos\theta-1)}{\theta} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}
$$

$$
B(\theta) = \begin{bmatrix} 0 & 0 \\ \frac{\cos\theta}{m} & -\frac{\sin\theta}{m} \\ 0 & 0 \\ \frac{\sin\theta}{m} & \frac{\cos\theta}{m} \\ 0 & 0 \\ \frac{r}{j} & 0 \end{bmatrix} \qquad (14)
$$

A set of polynomial 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 $\mathcal{P}$:

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

Once the CLF is obtained, it can be used as the terminal cost in a receding horizon optimization. Also As it was mentioned earlier, if the horizon length is long enough, the local CLF obtained by solving a Riccati equation for the linearized model of the system. The simulation results will be discussed in the next section.

4. Simulation Results

The following initial condition was chosen:

$$
[\ x\quad \dot{x}\quad y\quad \dot{y}\quad \theta\quad \dot{\theta}\ ] = [\ -1\quad 2.5\quad 0\quad 0\quad -45^\circ\quad 0\ ]
$$

Simulations were performed for several different horizon lengths. Also, to have a measure of how good the performance was, the optimal control trajectory was calculated by solving the optimization for the initial condition and a long horizon length. It can be seen from the simulation plots that a simple LQR controller performs very poorly on this system, where as the LPV controller does a better job. As we include on line optimization in the process, it can be seen that when a simple quadratic is chosen as the terminal cost, for long enough horizon, the phase curves converge to that of the optimal trajectory (Figure 1). This manifests the earlier results of De Nicolao *et al.*, Michalska and Mayne, and Chen and Allgöwer that if the horizon length is long enough, the CLF obtained by solving an Algebraic Riccati equation for the linearized system is going to be a good choice for terminal cost. However, if the horizon length is short, the receding horizon with LQR CLF does not give a good performance and in fact it is unstable. The x-y trajectories for several receding horizon iterations in this case are plotted in Figure 2. As it can be seen from Figure 3, the LPV CLF as the terminal cost does a better job if the horizon is not long.

An important point to keep in mind is that one might think that since stability is guaranteed by having a CLF as terminal cost, anything greater than that would be a legitimate choice. Simulation results indicate that this is not the case, and merely putting a quadratic penalty which is greater than the value of the CLF does not work. Figure 4 is the simulation of the receding horizon scheme with having a penalty of the form $\gamma||x(t+T)||^2$ where $gamma > 0$ is chosen such that $\gamma||x||^2 > V(x)\forall x$. Phase curves for this case are plotted in Figure 4. Finally, the ratio of the cost to the optimal case for different horizon lengths are plotted in Figure 5.

5. Conclusion

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, stability of receding horizon scheme is guaranteed. The results are based on the assumption that a 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

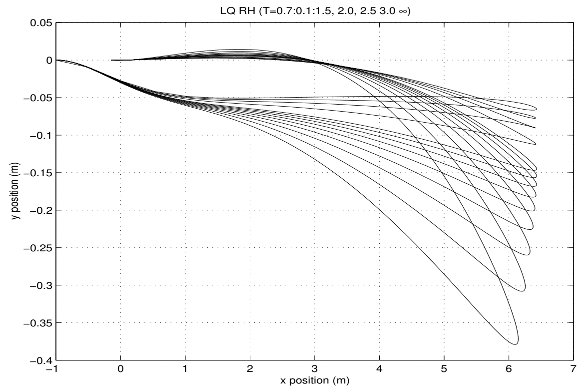

Figure 2: $x-y$ position of the fan for RH with quadratic CLF.

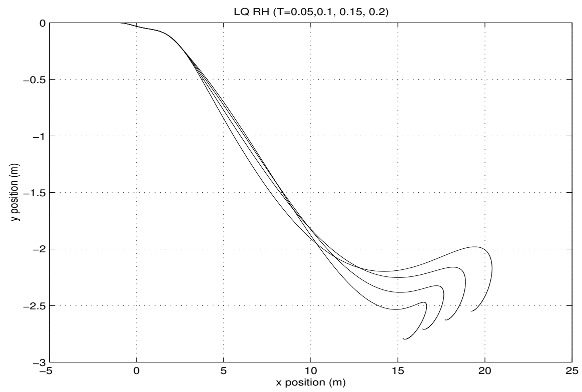

Figure 3: $x-y$ position of the fan for four receding horizon iterations (small horizon length).

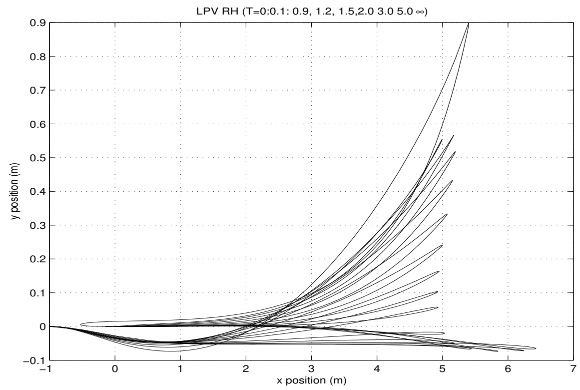

Figure 4: $x-y$ position of the fan for Receding horizon with the CLF from LPV.

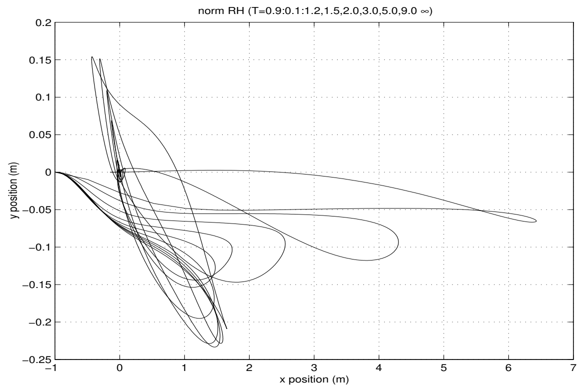

Figure 5: $x-y$ position of the fan for Receding horizon with $\gamma||x||^2$ as penalty.

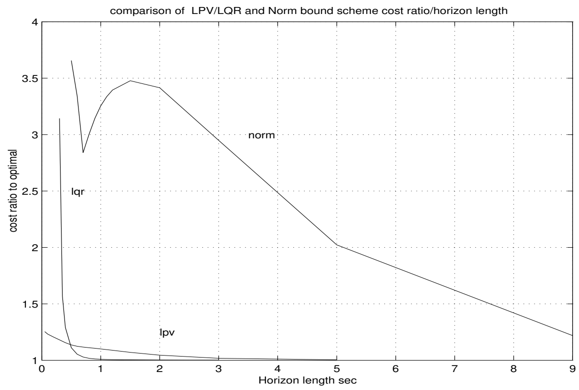

Figure 6: Ratio of the cost to the optimal vs. the horizon length.

can be replaced by its upper bound. This algorithm was applied to the planar model of a ducted fan. It was further shown that increasing the horizon length will improve the performance, so if a long time horizon is feasible in terms of computation, having a quadratic terminal cost obtained from the linearization is sufficient. However One should be careful not to make too many conclusions based on simulations of simple nonlinear models. The next step in this direction is to repeat these simulations for the full model including the aerodynamics as well as doing trajectory tracking rather than mere point stabilization.

References

[1] 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.

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

[3] A. Jadbabaie, J. Yu, and J. Hauser. Stabilizing receding horizon control of nonlinear systems: A control Lyapunov function approach. In *Submitted to ACC 99*, 1999.

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

[8] 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.

[9] 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.

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

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

[12] 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.

[13] 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.

[14] 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.