Applied Receding Horizon Control of the Caltech Ducted Fan¹

Ryan Franz²

Mark Milam³

John Hauser²

Abstract

This paper details the application of a constrained receding horizon control strategy to an indoor vectored-thrust flight experiment known as the Caltech Ducted Fan. The strategy is used to stabilize the experiment about one operating point, and step response and disturbance rejection are examined with different configurations and in comparison to a gain-scheduled LQR controller. Issues related to non-zero computation times, choice of horizon length and terminal cost are discussed. **Keywords:** real-time optimization, model predictive control, optimal control, nonlinear control, guidance.

1. Introduction

In receding horizon control, an open-loop trajectory is found by solving a finite-horizon optimal control problem starting from the current state. The controls of this trajectory are then applied for a certain fraction of the horizon length, after which the process is repeated. See [6] for a good review of recent work in this field. This approach has been used in the process control industry successfully for some time, where dynamics are relatively slow, but the computing power required and the tendency for naive implementations to destabilize a system have prevented its use on fast, stability critical nonlinear systems. Recently, however, theoretical results regarding stability properties of receding horizon controllers and ever increasing computing power have revived interest in the scheme.

The application of Receding Horizon Control to aerial vehicles has been proposed and analyzed by several researchers; [10], for example, provide simulation results for stabilization of a helicopter UAV about an open loop trajectory using receding horizon control. The strategy offers many benefits in this environment, such as the inherent ability to deal with constraints in the state and control. Examples of such constraints commonly encountered include static terrain obstacles, dynamic or pop-up threats and saturations on the actuators. Much of the existing work, however, has not addressed the fact that computation times cannot be ignored with the fast dynamics of most aerial platforms.

The goal of the work in this paper was twofold: first, to address issues of implementation with substantial computation times, and second, to provide a validation of theoretical results through implementation on an ac-

tual nonlinear experiment. The optimal control problem is solved using NTG (Nonlinear Trajectory Generation), a software package developed at Caltech [7]; a full aerodynamic model of the flying wing is used in order to demonstrate feasibility on a true UAV platform. A gain-scheduled LQR controller is used for comparison because of its good overall performance. This paper deals exclusively with stability, region of attraction and disturbance rejection properties; use of receding horizon in a maneuvering sense will be examined in a future paper.

The progression of the paper is as follows: section 2 provides theoretical background as well as some motivation for the choices made in terms of timing; section 3 describes in detail the two different timing methods used in the experiment; section 4 describes the actual experiment, and finally sections 6 and 7 show simulation and actual results before concluding.

2 Theoretical Background

This section presents a summary of relevant theory and motivates our choices for timing made in the sequel.

We begin by defining a system with an ideal feedback law, looking at it in a sampled data sense and then proposing four different methods of applying the feedback. Next we consider the effects of weakening the idealness of the system, where one method becomes the standard theoretical receding horizon formulation, one method becomes a naive implementation with non-zero runtimes, and the remaining two become candidates for actual implementation. We conjecture that the implementable methods will be stable for short enough runtimes and long enough horizons.

Our system is described by

$$\dot{y} = f(y, u) + g(y, u, w) \quad (1)$$

where $f(\cdot, \cdot)$ is the nominal (i.e., model) system vector field and $g(\cdot, \cdot, \cdot)$ describes the effect of the external disturbance w together with that portion of the system dynamics that is not explicitly modeled. Thus, for the purpose of control design, etc., we will use

$$\dot{x} = f(x, u) \quad (2)$$

as the *model* system. Now, suppose that

$$u = k(x) \quad (3)$$

is a state feedback that exponentially stabilizes the origin for the nominal system (2) and that $V(x)$ is a quadratic Lyapunov function proving such. For example, $k(\cdot)$ might arise as the solution to an infinite

¹Research supported in part by DARPA grant F33615-98-C-3613

²Electrical and Computer Engineering, University of Colorado, Boulder, CO 80309-0425

³Control and Dynamical Systems, Mail Code 107-81, California Institute of Technology, Pasadena, CA 91125

horizon optimization problem with $V(\cdot)$ as the corresponding minimum cost (to go). In the case that the perturbation is nonzero but can be bounded by a constant, one may use Lyapunov arguments to show that the state of the true closed loop system (1), (3) will converge to a neighborhood of the origin.

Next we construct a sampled data feedback structure such that at every time $t_k := k\delta$ we obtain a measurement $y_k := y(t_k)$. At every time step, we calculate a trajectory $x(\cdot; y_k)$, $u(\cdot; y_k)$ by simulating the closed loop model system (2), (3) for a length of time (either δ or 2δ seconds).

We propose the following four methods for applying the resulting *open loop* input trajectory to the actual system (1):

1. apply $u_{[0,\delta]}(y_k)$ (the control trajectory over the interval $t_{sim} \in [0, \delta]$ resulting from a simulation starting at y_k) over the interval $t \in [t_k, t_{k+1}]$. Note that implementing this option requires that the simulation be run in zero time.
2. apply $u_{[0,\delta]}(y_k)$ over the interval $t \in [t_{k+1}, t_{k+2}]$. Note that this option will always involve a delay.
3. apply $u_{[\delta,2\delta]}(y_k)$ over the interval $t \in [t_{k+1}, t_{k+2}]$.
4. apply $u_{[0,\delta]}(x(\delta; y_k))$ over the interval $t \in [t_{k+1}, t_{k+2}]$. Here $x(\delta; y_k)$ represents the state of the system x starting at y_k simulated ahead δ s.

When the system perturbation is identically zero $g(x, u, w) \equiv 0$, we see that options 1, 3, and 4 will be identical. Options 2, 3, and 4 are all implementable if the simulation computation can be completed in less than δ seconds (i.e., faster than real time). Because option 2 involves a delay (even in the no perturbation case), we propose that 3 and 4 will be the best methods with non-zero runtimes. Clearly the performance of the sampled data system schemes with nonzero perturbation will depend on the sample time δ .

As a next step, suppose that we compute the input trajectory $u(\cdot; y_k)$ by solving the finite horizon optimal control problem

$$J_T^*(y(t_k)) = \min_{u(\cdot)} \int_0^T q(x(\tau), u(\tau)) d\tau + V(x(T)), \quad (4)$$

$$\dot{x}(t) = f(x(t), u(t)), \quad x(0) = y(t_k)$$

where the incremental cost satisfies $q(x, u) \geq c_q(\|x\|^2 + \|u\|^2)$ with $c_q > 0$. If the terminal cost $V(\cdot)$ is chosen to be a control Lyapunov function (CLF) satisfying $\min_u ((\dot{V}) + q)(x, u) \leq 0$ on a neighborhood of the origin, option 1 (with $g(x, u, w) \equiv 0$) is the receding horizon control scheme $\mathcal{RH}(T, \delta)$, analyzed in [4]. Now allowing $g(x, u, w)$ to be nonzero, we discuss some stability properties of this structure.

As in the stability analysis of unperturbed receding horizon control [4], we will use $J_T^*(\cdot)$ as a Lyapunov function. Roughly speaking, we require that $J_T^*(y_k)$

be a strictly decreasing sequence, ensuring the convergence of the state to a (hopefully small) neighborhood of the origin.

Note that $J_T^*(\cdot)$ is Lipschitz continuous with constant K over the compact region of interest. The properties of $q(\cdot, \cdot)$ and $V(\cdot)$ ensure that

$$J_T^*(x(t_k + \delta)) \leq J_T^*(x(t_k)) - Q_\delta(x(t_k)) \quad (5)$$

where the decrement $Q_\delta(\cdot)$ is a positive definite function (given by integrating the optimal incremental cost over a δ second interval).

Suppose, now, that we apply the same open loop control $u(\cdot)$ (e.g., the just computed optimal $u(\cdot)$) to the real and model systems, (1) and (2), with potentially different initial conditions. By a standard argument (using the Bellman-Gronwall lemma, see [5]), we have

$$\|y(t_k + \delta) - x(t_k + \delta)\| \leq e^{L\delta} \|y(t_k) - x(t_k)\| + \frac{b}{L} (e^{L\delta} - 1) \quad (6)$$

where b is a bound on $\|g(y(t), u(t), w(t))\|$, $t \in [t_k, t_k + \delta]$, and L is a Lipschitz constant for $f(\cdot, \cdot)$.

Combining (5) and (6) and noting that $y(t_k) = x(t_k)$, we obtain

$$\begin{aligned} J_T^*(y(t_k + \delta)) &\leq J_T^*(x(t_k + \delta)) + K\|y(t_k + \delta) - x(t_k + \delta)\| \\ &\leq J_T^*(y(t_k)) - Q_\delta(y(t_k)) + K\frac{b}{L}(e^{L\delta} - 1). \end{aligned}$$

For small $\delta > 0$, we can bound the terms on the right according to $Q_\delta(x) \geq \frac{\delta}{2} c_q \|x\|^2$ and $K\frac{b}{L}(e^{L\delta} - 1) \leq 2Kb\delta$. We conclude that, for small δ , $J_T^*(y(t))$ will decrease provided that

$$\|y(t)\|^2 \geq \frac{4Kb}{c_q}. \quad (7)$$

This determines the *radius* r_b for an invariant sublevel set of $J_T^*(\cdot)$ to which the state of the true system will converge to under the scheme of option 1.

Finally, we extend this discussion to include situations in which $y(t_k) \neq x(t_k)$, which is the case in options 2, 3, and 4. In this case (6) necessarily contains an exponential (in δ) term multiplied by the error in the initial conditions. Performing analysis similar to that detailed above, we obtain the relation

$$\begin{aligned} J_T^*(y(t_k + \delta)) &\leq J_T^*(y(t_k)) - Q_\delta(y(t_k)) + K\frac{b}{L}(e^{L\delta} - 1) \\ &\quad + (K(1 + e^{L\delta}) + K_Q)\|y(t_k) - x(t_k)\| \end{aligned} \quad (8)$$

where K_Q is a Lipschitz constant for $Q_\delta(\cdot)$. Clearly, mismatches in the initial conditions lead to performance degradations, including an enlargement of the terminal set (r_b increased) as well as potential destabilization. It is therefore of prime importance to minimize the initial condition mismatch to the extent possible. We conjecture that option 2 does a poor job of this; indeed, even in the no perturbation case such an error is induced by delay. Accordingly, we study options 3 and 4 both in simulation and experimentally on the physical system in an attempt to determine which will provide the best performance.

Figure 1: Illustration of timing scheme without prediction

3 Timing and Optimization Formulation

In some applications of receding horizon, runtimes are insignificant compared to the dynamics of the system. This is not the case on most aerial platforms with current computing power. On our hardware we were able to achieve runtimes between 0.1s and 0.2s in most cases; we ordinarily run linear controllers at a minimum of 50Hz. Because of this, the preceding discussion is crucial. Before detailing our timing implementations, we first state explicitly the cost functions used in this paper, as defined in equation 4:

$$\begin{aligned} q(x(t), u(t)) &= \frac{1}{2} x_{err}^T(t) Q x_{err}(t) + \frac{1}{2} u_{err}^T(t) R u_{err}(t) \\ V(x(T)) &= \frac{1}{2} x_{err}^T(T) P x_{err}(T) \end{aligned} \quad (9)$$

We choose Q and R to be the same as weights used in the past to generate LQR gains with good performance, and P to be the corresponding solution to the algebraic Riccati equation resulting in a CLF terminal cost around a certain operating point. The choice of $x(0)$ for the optimization is dictated by the choice of timing scheme. We use two different strategies, corresponding to options 3 and 4 in the previous section, for choosing these initial constraints.

3.1 Option 3: No Prediction

The first scheme for choosing the initial constraints in the state is the simplest, as it involves no model prediction. Whenever a computation is triggered, the current state of the system is given as the initial constraint on the state trajectory for the optimization problem. By the time the computation is finished t_{sample} seconds later, however, the idea is that the system has changed significantly. To attempt to use a valid control, we simply discard the first t_{sample} seconds of the trajectory, hoping that the resulting start point will coincide roughly with where we were in the previous trajectory. Fig. 1 shows graphically how this process works on one of the states. In this case, the controls corresponding to the line labeled “Receding Horizon Reference Trajectory” are applied to the system. Some of our results apply these controls open-loop, while at other times we apply V_m open-loop while applying a PD feedback rule with θ and $\dot{\theta}$ to compute the δp used. Note that the figure exaggerates certain things for illustration—for example, the horizon length $t_{horizon}$ is in reality much longer than t_{sample} .

In our implementation, t_{sample} can either be set to

Figure 2: Illustration of timing scheme with prediction

some constant, or the computations can be run as “fast as possible”, meaning a new computation is triggered immediately after the last one has finished. In this case, t_{sample} varies with the runtime.

3.2 Option 4: With Prediction

The second scheme we examine attempts to minimize discontinuities by using prediction. When a computation is triggered, the current state of the fan is first used as the initial condition for a simulation in which the control trajectory of the previous computation is used as input. This simulation is run for some amount of time t_{sim} ; if a fixed period is being used, t_{sim} is simply equal to the t_{sample} , but if a “fast as possible” rule is used, t_{sim} is taken as an average of the past n runtimes. After the simulation is completed, the final values are passed as the initial constraints to the optimization. The resulting trajectory is output from the beginning. Fig. 2 shows this process graphically. Again, the controls corresponding to the line labeled “Receding Horizon Reference Trajectory” are applied to the system, and can be used either open-loop or with feedback around θ and $\dot{\theta}$.

3.3 Further Considerations

As with any timing scheme, there are necessarily discontinuities in the resulting control due to model mismatch and a non-zero sampling period. Early experience showed that some effort in minimizing these jumps was worthwhile. Accordingly, we use an inequality constraint on each optimization to achieve smoother control signals, $|u_{k+1}(0) - u_k(t_{sample})| < a$ for some a . If a fixed period is used, t_{sample} is simply equal to the period, but if a “fast as possible” rule is used, t_{sample} is taken as an average of the past n runtimes. This approach is compatible with both timing schemes discussed above; graphically, control trajectories always start near the previous trajectory.

Another consideration involves non-convergent trajectory computations. Unfortunately, not all trajectory computations are guaranteed to converge. Each computation is given the last computed trajectory as an initial guess, which is sometimes not good enough; also, some combinations of initial constraints and cost function are simply degenerate. If a computation returns certain signs of failure, the last good trajectory is simply continued and another computation is triggered. This will certainly fail if non-convergence happens frequently or repeatedly, as it has the effect of greatly increasing the sample time; in practice, this has not

Figure 3: Ducted Fan Inertial Schematic Diagram

Figure 4: Ducted Fan Inertial Coordinate Frame

been a large problem.

4 Experimental Setup

Fig. 3 shows an overview of the Caltech ducted fan. The experiment consists of a vertical stand and a horizontal boom which holds the actual wing. This setup enables flight on a cylinder of height 2.5 m and radius 2.35 m. Because of a mass of 12.5 kg and a maximum thrust of only 15 N, a counterweight is attached to the boom via a cable and pulleys which reduces the effective gravity to achieve $mg_{eff} = 7$ N. This allows the system to attain sizable vertical accelerations, while minimizing the force of potential crashes. Mechanical brakes in the vertical direction are used as well to aid in crash landings. Actuation of the ducted fan is accomplished in two ways: by controlling the speed of the propeller mounted inside the cowl, and by vectoring the resulting thrust via a servo controlled bucket.

Fig. 4 shows more explicitly the inertial coordinate frame used in this paper. In this frame, the axes are fixed to the ground, and the x and z directions represent horizontal and vertical inertial translations. θ represents the rotation of the ducted fan about the boom axis. All three of these variables are measured via rotary encoders, and the resulting signals are routed to the computing platform via sliprings.

The equations of motion are given by

$$\begin{aligned} m\ddot{x} + F_{X_a} - F_{X_b} \cos \theta - F_{Z_b} \sin \theta &= 0 \\ m\ddot{z} - mg_{eff} + F_{Z_a} + F_{X_b} \sin \theta - F_{Z_b} \cos \theta &= 0 \\ J\ddot{\theta} - M_a + \frac{1}{r_s} I_p \Omega \dot{x} \cos \theta - F_{Z_b} r_f &= 0 \end{aligned}$$

where

$$F_{X_a} = D \cos \gamma + L \sin \gamma, \quad F_{Z_a} = -D \sin \gamma + L \cos \gamma$$

are the aerodynamic forces. See [8] for a complete derivation of these equations. We chose a spatial representation of the equations of motion so that we can consider both hover and forward flight modes. F_{X_b} and F_{Z_b} are thrust vectoring body forces; $I_p = 2e^{-5} \text{ kg mm}^2$ and $\Omega = 1300 \text{ rad/s}$ are the moment of inertia and angular velocity of the ducted fan propeller, respectively. $J = .25 \text{ kg m}^2$ is the moment of inertia of the ducted fan about the boom, and $r_f = .35 \text{ m}$ is the distance from center of mass along the X_b axis to the effective application point of the thrust vectoring force. The angle of attack α is related to the pitch angle θ and the flight path angle γ by $\alpha = \theta - \gamma$, where the flight path angle can be derived from the spatial velocities by

$$\gamma = \arctan \frac{-\dot{z}}{\dot{x}}.$$

The lift (L), drag (D), and moment (M) are given by

$$L = qSC_L(\alpha), D = qSC_D(\alpha), \text{ and } M = \bar{c}SC_M(\alpha),$$

respectively. The dynamic pressure is given by $q = \frac{1}{2} \rho V^2$. The norm of the spatial velocity is denoted by V and ρ is the atmospheric density. The coefficients of lift ($C_L(\alpha)$), drag ($C_D(\alpha)$) and moment ($C_M(\alpha)$) were determined from a combination of wind tunnel and flight testing.

5 Trajectory Generation Methodology

There are three components to the trajectory generation methodology we propose. The first is to determine a parameterization (output) such that the equations of motion can be mapped to a lower dimensional space (output space). [2] gives information on finding this mapping if the system is flat. The idea is to map dynamic constraints to algebraic ones. Once this is done the cost and constraints can also be mapped to the output space. The second is to parameterize each component of the output in terms of an appropriate B-spline polynomial. Finally, sequential quadratic programming is used to solve for the coefficients of the B-splines that minimize the cost subject to the constraints in output space. See [7, 9] for more details on this approach. The NTG software package is an implementation of this concept. The user provides the cost and the constraints in terms of the outputs and their derivatives as well as the Jacobian of the cost and constraints with respect to each output and the maximum derivative that occurs in each output.

By using this methodology, we can sufficiently reduce the dimension of the nonlinear programming problem to make real-time computation possible. For our system we will choose as outputs $z_1 = x(t)$, $z_2 = z(t)$, $z_3 = \theta(t)$ in solving the problem posed in eq. 9. Given these outputs, their derivatives and the control trajectory can be computed easily.

6 Simulation Study

Because of the complex implementation involved in this experiment, a simulation was built in tandem with the real implementation which was functionally equivalent

horizon	predict		no predict	
	w/o θ	w/ θ	w/o θ	w/ θ
1.0s	0.4s	0.1s	0.15s	0.2s
1.5s	0.5s	0.1s	0.2s	0.4s
2.0s	0.65s	0.2s	0.3s	0.7s
2.5s	0.6s	0.45s	0.4s	0.6s
3.0s	0.5s	0.3s	0.4s	0.45s

Table 1: Maximum acceptable periods as determined in simulation

except for the use of a model instead of the real system. The model used is the same model used in the NTG code, but we quantize the outputs to emulate the effect of the encoders and use a Kalman filter estimator to estimate the derivatives. The simulation allows us to explore many different configurations without fear of damaging the hardware. Table 1 shows results of identifying the highest acceptable periods for different combinations of timing mode, horizon length and theta controller. The test used for these results was a 20 m step in x , a fairly demanding request which puts the fan into a forward flight state to test out the full features of the model. We were unable to design a gain-scheduled LQR controller which could perform this maneuver in an acceptable fashion. Acceptable results were chosen as stable and with few qualitative difference from the best results. The theta controller is simply a PD loop around computed theta trajectories which proved necessary on the experiment.

An interesting feature of these results is that the theta controller significantly improves the performance of the non-prediction timing scheme, while it has a negative effect on the prediction timing system. We attribute this to the fact that the prediction mode relies on knowing the exact inputs to the model as given by the last computation, while the theta controller changes those inputs.

In this range of horizon values, runtimes are consistently between 0.05 s and 0.35 s. Accordingly, we anticipated a much greater chance of success with greater horizon lengths.

7 Results

The same controllers which were used in the simulation study were implemented on the real ducted fan. We present here the results of two investigations: a study of maximum acceptable periods in different configurations for the step response, and an investigation of behavior in the face of severe disturbances.

7.1 Step Response

In order to examine the limits of the different configurations, we performed experiments similar to the simulation study in which we sought the maximum periods providing acceptable response to a step input. On the real experiment we used a smaller step of 5m, in order to prevent damage to the apparatus. Another difference from the simulation is that the fixed period chosen is only a lower bound on the actual period; the majority of calculations remain below the fixed period in all the runs, but there are still some which exceed the value due to limited computing power. Table 2 shows the values we obtained; a value of 0.0s means that only a "fast as possible" scheme was acceptable. We were unable

horizon	predict		no predict	
	w/o θ	w/ θ	w/o θ	w/ θ
1.0s	—	0.0s	—	0.2s
2.0s	—	0.3s	—	0.6s
3.0s	—	0.4s	—	0.3s

Table 2: Maximum acceptable periods as determined on the real experiment

to achieve stability without the θ controller. To compare performance with the scheduled LQR controller, figures 5 and 6 show the LQR and receding horizon step responses in x . The lines branching off of the receding horizon plot show the trajectories computed by the optimization routine. The LQR controller, which was highly tuned, shows a much longer rise time.

Figure 5: x response to a 5 m step using scheduled LQR controller

Figure 6: x response to a 6 m step input; horizon= 1 s, period= 0.1s, no prediction, w/ θ controller

7.2 Disturbance Rejection

To examine the disturbance rejection properties of the receding horizon controllers, we gave the fan consistent shoves involving a large torque on θ and a large force on x . We used the "fast as possible" scheme with all combinations of prediction mode and horizon length. All configurations were able to recover from the disturbances in a very reasonable manner. Figures 7 and 8 show the response in x to a series of these disturbances using both the scheduled LQR controller and

the receding horizon controller. Although it was very difficult to guarantee perfectly consistent disturbances, the plots show definite qualitative differences; the LQR controller takes longer to recover and is more tolerant of offsets.

Figure 7: x response to a series of disturbances, using scheduled LQR controller

Figure 8: x response to a series of disturbances; horizon= 3 s, "fast as possible", prediction, w/ θ controller

For comprehensive experimental results using the "fast as possible" scheme, see [1]. [1] also correlates different horizon lengths with runtimes.

8 Conclusion

The theoretical discussion in section 2 provides a framework for evaluating in a qualitative sense different choices for a timing scheme to deal with non-zero runtimes. In particular, two timing paradigms are chosen which seem reasonable. The simulation study showed the feasibility of both methods; the prediction mode shows longer maximum periods than the mode without prediction when the θ controller is not used, while the no prediction mode had longer maximum periods with the θ controller.

The real system concurred with these results to a certain degree; of particular interest is the fact that the best results (i.e., longest maximum periods) were obtained at the same point in the table, with 2s horizon,

using no prediction with the θ controller. We conjecture that the reason for the reduced performance with the 1s horizon was due to the fact that the runtimes were longer, often longer than the fixed period; this especially affects the prediction mode controller, since the model is predicted ahead by an average of past runtimes. The performance with the 3s horizon was most likely degraded by the reduced accuracy of the B-spline representation used by NTG with the increased trajectory length. We plan to continue to investigate the instability of the system without the θ controller; we attribute this to fast, unmodeled dynamics in the θ direction which require shorter periods than our computing power currently allows.

To summarize, we believe the results show a clear advantage over a finely tuned LQR controller in both step response and disturbance rejection. The receding horizon methodology developed has other advantages, such as in dealing with constraints and changes in model parameters. Current work involves expanding this methodology to provide robust, nonlinear trajectory tracking which could be used for entire mission management. We also are seeking stronger mathematical results which could indicate theoretically which timing mode to implement. Another intriguing possibility is to extend the prediction timing mode by feeding the optimization routine constant updates of what the best predicted state is. The predicted state would increase in accuracy while not changing the resulting optimization solution greatly.

References

- [1] W.B. Dunbar, M.B. Milam, R. Franz, and R. M. Murray. Model predictive control of a thrust-vector flight control experiment. In *15th IFAC World Congress*.
- [2] M. Fliess, J. Lévine, P. Martin, and P. Rouchon. A Lie-Bäcklund approach to equivalence and flatness of nonlinear systems. *IEEE Trans. Auto. Cont.*, 44(5):928-937, 1999.
- [3] Ali Jadbabaie and John Hauser. On the stability of unconstrained receding horizon control with a general terminal cost. *Systems and Control Letters*, submitted.
- [4] Ali Jadbabaie, Jie Yu, and John Hauser. Unconstrained receding-horizon control of nonlinear systems. *IEEE Trans. on Automatic Control*, 46:776-783, 2001.
- [5] Hassan Khalil. *Nonlinear Systems*. Prentice Hall, Upper Saddle River, NJ, 1996.
- [6] D.Q. Mayne, J. B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789-814, 2000.
- [7] M. B. Milam, K. Mushambi, and R. M. Murray. A new computational approach to real-time trajectory generation for constrained mechanical systems. In *IEEE Conference on Decision and Control*, 2000.
- [8] Mark Milam and R.M. Murray. A testbed for nonlinear flight control techniques: The caltech ducted fan. In *Conference on Control Applications*, 1999.
- [9] N. Petit, M. B. Milam, and R. M. Murray. Inversion based constrained trajectory optimization. In *5th IFAC symposium on nonlinear control systems*, 2001.
- [10] Leena Singh and James Fuller. Trajectory generation for a UAV in urban terrain, using nonlinear MPC. In *Proceedings of the American Control Conference*, 2001.