

Controller Synthesis for Constrained Flight Systems via Receding Horizon Optimization

Richard M. Murray William B. Dunbar
Reza Olfati Saber Lars B. Cremean

Division of Engineering and Applied Science
California Institute of Technology

Receding horizon control allows a blending of navigation and control functions at the inner and outer loop levels and significantly enhances the ability of the control system to react to complex dynamic and environmental constraints. In this paper, we explore some of the limits of receding horizon control, including the extent to which traditional control specifications can be cast as RHC problem specifications. Simulation results for a planar flight vehicle with representative flight dynamics illustrate the main features of the proposed approach.

Motivation and Philosophy

Over the past five years, the amount of computing power available in flight control systems has increased dramatically, and it has recently become possible to perform online optimization quickly enough that one can stabilize flight systems using receding horizon, optimal control.^{1,8} This paper advocates the use of such an approach as the primary mechanism for both inner and outer loop control of aircraft. Furthermore, we propose a design paradigm that utilizes the insights that control designers have gained for high performance, robust control design, while exploiting the ability of optimization-based control approaches to handle constraints and reconfigurable operation.

The basic philosophy that we propose is illustrated in Figure 1. We begin with a nonlinear system, including a description of the constraint set. We linearize this system about a representative equilibrium point and perform a linear control design using standard (modern) tools. Such a design gives provably robust performance around the equilibrium point and, more importantly, allows the designer to meet a wide variety of formal and informal performance specifications through experience and the use of sophisticated linear design tools.

This linear control law then serves as a *specification* of the desired control performance for the entire nonlinear system. We convert the control law specification into a receding horizon control formulation, chosen such that for the linearized system, the receding horizon controller gives comparable perfor-

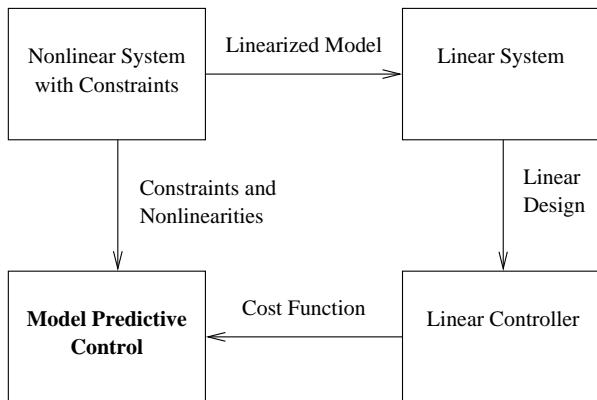


Fig. 1 Optimization-based control approach.

mance. However, because of its use of optimization tools that can handle nonlinearities and constraints, the receding horizon controller is able to provide the desired performance over a much larger operating envelope than the controller design based just on the linearization. Furthermore, by choosing cost formulations that have certain properties, we can provide proofs of stability for the full nonlinear system and, in some cases, the constrained system.

The advantage of the proposed approach is that it exploits the power of humans in designing sophisticated control laws in the absence of constraints with the power of computers to rapidly compute trajectories that optimize a given cost function in the presence of constraints. New advances in online trajectory generation serve as an enabler for this approach and their demonstration on representative flight control experiments shows their viability. This approach can be extended to existing nonlinear paradigms as well, as we describe in more detail

below.

To our knowledge, the basic philosophy that we propose has not been explored in detail in other flight control applications and is certainly not a mainstream approach to flight control design. It represents a shift in the basic methods for control law design of complex, nonlinear systems by focusing the designer on the fundamental (linearized) dynamics of the system and using computation to provide full-envelope, highly aggressive flight control laws.

This paper is organized as follows. We first provide a review of some of the prior work related to the approach described above. Following this review, we give some preliminary results on some of the possibilities and limits for our framework. These results analyze some simple cases that serve as both examples and bounds on what is achievable. We next present simulation results for a simple flight-like system that demonstrates the efficacy of our approach. Finally, we give our conclusions and a discussion of next steps.

Review of Relevant Previous Work

In this section we present a brief review of relevant work and indicate some of the prior studies in the same direction as what is proposed here. This review is not intended to be exhaustive, but rather to put into context some of the results that guide and bound our approach. Some of the material in this section is drawn from a recent book chapter by the authors.¹

Optimization-based control

Optimization-based control refers to the use of online, optimal trajectory generation as a part of the feedback stabilization of a (typically nonlinear) system. The basic idea is to use a *receding horizon* control technique: a (optimal) feasible trajectory is computed from the current position to the desired position over a finite time T horizon, used for a short period of time $\delta < T$, and then recomputed based on the new position.

A key advantage of optimization-based approaches is that they allow the potential for customization of the controller based on changes in *mission*, *condition*, and *environment*. Because the controller is solving the optimization problem online, updates can be made to the cost function, to change the desired operation of the system; to the model, to reflect changes in parameter values or damage to sensors and actuators; and to the constraints, to reflect new regions of the state space that must be avoided due to external influences. Thus, many of the challenges of designing controllers that are ro-

bust to a large set of possible uncertainties become embedded in the online optimization.

Development and application of receding horizon control (also called model predictive control, or MPC) originated in process control industries where plants being controlled are sufficiently slow to permit its implementation. An overview of the evolution of commercially available MPC technology is given by Qin and Badgwell¹⁴ and a survey of the current state of stability theory of MPC is given by Mayne et al.⁶ Closely related to the work in this paper, Singh and Fuller¹⁶ have used MPC to stabilize a linearized simplified UAV helicopter model around an open-loop trajectory, while respecting state and input constraints.

A number of approaches in receding horizon control employ the use of terminal state equality or inequality constraints, often together with a terminal cost, to ensure closed loop stability. In Primbs et al.,¹³ aspects of a stability-guaranteeing, global control Lyapunov function were used, via state and control constraints, to develop a stabilizing receding horizon scheme. Many of the nice characteristics of the CLF controller together with better cost performance were realized. Unfortunately, a global control Lyapunov function is rarely available and often not possible.

Motivated by the difficulties in solving constrained optimal control problems, we have developed an alternative receding horizon control strategy for the stabilization of nonlinear systems.³ In this approach, closed loop stability is ensured through the use of a terminal cost consisting of a control Lyapunov function that is an incremental upper bound on the optimal cost to go. This terminal cost eliminates the need for terminal constraints in the optimization and gives a dramatic speed-up in computation. Also, questions of existence and regularity of optimal solutions (very important for online optimization) can be dealt with in a rather straightforward manner.

Linear quadratic regulators

The philosophy presented here relies on the synthesis of an optimal control problem from specifications that are embedded in an externally generated controller design. This controller is typically designed by standard classical control techniques for a nominal plant, absent constraints. In this framework, the controller's performance, stability and robustness specifications are translated into an equivalent optimal control problem and implemented in a receding horizon fashion.

One central question that must be addressed when

considering the usefulness of this philosophy is: *Given a control law, how does one find an equivalent optimal control formulation?* The seminal paper by R. E. Kalman⁵ lays a solid foundation for this class of problems, known as *inverse optimality*. In this paper, Kalman considers the class of linear time-invariant (LTI) plants with full-state feedback and a single input variable, with an associated cost function that is quadratic in the input and state variables. These assumptions set up the well-known linear quadratic regulator (LQR) problem, by now a staple of optimal control theory.

In Kalman’s paper, the mathematical framework behind the LQR problem is laid out, and necessary and sufficient algebraic criteria for optimality are presented in terms of the algebraic Riccati equation, as well as in terms of a condition on the return difference of the feedback loop. In terms of the LQR problem, the task of synthesizing the optimal control problem comes down to finding the integrated cost weights Q and R given only the dynamical description of the plant represented by matrices A and B and of the feedback controller represented by K . Kalman delivers a particularly elegant frequency characterization of this map.⁵

There are two natural extensions of these results: extension to more general dynamical systems and extension to more general optimal control formulations. The contribution of this paper is the simultaneous extension of this approach to systems with constraints along with the extension to the more general receding horizon control framework. A first step in this approach is extension of inverse optimal results to the finite horizon case.

It is important to note that Kalman’s results are restricted to the infinite horizon case ($T \rightarrow \infty$) in addition to the assumptions of linearity, time-invariance and quadratic costs. This additional assumption is necessary to derive the results associated with the algebraic Riccati equation ($\dot{P} = 0$). However, we will show that through proper application of terminal costs, the same inverse optimality problem can be soundly addressed in the case of finite horizon length. This problem is addressed by the authors in this paper in the context of Kalman’s work; the review of these results will be made mathematically explicit in the next section.

Inverse optimality for nonlinear systems

The above results can be generalized to nonlinear systems, in which one takes a nonlinear control system and attempts to find a cost function such that the given controller is the optimal control with respect to that cost.

The history of inverse optimal control for nonlinear systems goes back to the early work of Moylan and Anderson.⁹ More recently, Sepulchre et al.¹⁵ showed that a nonlinear state feedback obtained by Sontag’s formula from a control Lyapunov function (CLF) is inverse optimal. The connections of this inverse optimality result to passivity and robustness properties of the optimal state feedback are discussed in Jankovic *et al.*⁴ The past research on inverse optimality does not consider the constraints on control or state. However, the results on the unconstrained inverse optimality justify the use of a more general nonlinear loss function in the integrated cost of a finite horizon performance index combined with a real-time optimization-based control approach that takes the constraints into account.

Preliminary results

In this section we describe some preliminary results that determine the applicability of the control design methodology described in the introduction.

State feedback and LQR

We begin by considering the question of when a given state feedback control law can be realized as an optimal controller for a given cost function. This is the problem that Kalman considered⁵ and our results are a variation on that work, focused on a receding horizon formulation.

We consider a linear system

$$\dot{x} = Ax + Bu \quad x \in \mathbb{R}^n, u \in \mathbb{R}^m \quad (1)$$

with state x and input u . We consider only the single input, single output case for now ($m = 1$). Given a control law

$$u = Kx$$

we wish to find a cost functional of the form

$$J = \int_0^T x^T Q x + u^T R u dt + x^T(T) P_T x(T) \quad (2)$$

where $Q \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{m \times m}$ define the integrated cost, $P_T \in \mathbb{R}^{n \times n}$ is the terminal cost, and T is the time horizon. Our goal is to find $P_T > 0$, $Q > 0$, $R > 0$, and $T > 0$ such that the resulting optimal control law is equivalent to $u = Kx$.

The optimal control law for the quadratic cost function (2) is given by

$$u = -R^{-1} B^T P(t),$$

where $P(t)$ is the solution to the Riccati ordinary differential equation

$$-\dot{P} = A^T P + PA - PBR^{-1}B^T P + Q \quad (3)$$

with terminal condition $P(T) = P_T$. In order for this to give a control law of the form $u = Kx$ for a constant matrix K , we must find P_T , Q , and R that give a constant solution to the Riccati equation (3) and satisfy $-R^{-1}B^T P = K$. It follows that P_T , Q and R should satisfy

$$\begin{aligned} A^T P_T + P_T A - P_T B R^{-1} B^T P_T + Q &= 0 \\ -R^{-1} B^T P_T &= K. \end{aligned} \quad (4)$$

We note that the first equation is simply the normal algebraic Riccati equation of optimal control, but with P_T , Q , and R yet to be chosen. The second equation places additional constraints on R and P_T .

Equation (4) is exactly the same equation that one would obtain if we had considered an infinite time horizon problem, since the given control was constant and hence $P(t)$ was forced to be constant. This infinite horizon problem is precisely the one that Kalman considered in 1964, and hence his results apply directly. Namely, in the single-input single-output case, we can always find a solution to the coupled equations (4) under standard conditions on controllability and observability.⁵ The equations can be simplified by substituting the second relation into the first to obtain

$$A^T P_T + P_T A - K^T R K + Q = 0.$$

This equation is linear in the unknowns and can be solved directly (remembering that P_T , Q and R are required to be positive definite).

The implication of these results is that any state feedback control law satisfying these assumptions can be realized as the solution to an appropriately defined receding horizon control law. Thus, we can implement the design framework summarized in Figure 1 for the case where our (linear) control design results in a state feedback controller.

The static state feedback that corresponds to the finite horizon optimal control problem in equation (2) is implemented in a receding horizon (RH) fashion by applying the control for any δ seconds, where $0 < \delta < T$. Since we have started with a feedback gain K that is known to be stabilizing, the receding horizon application of the feedback is trivially stabilizing. On the other hand, if one is solving the standard finite horizon LQ problem, i.e., given Q, R and P_T , find $K(t)$, the RH application of the time-varying feedback is not necessarily stabilizing, a result known for some time.¹⁰ In particular, in the limit as $\delta \rightarrow 0$, and one considers the receding horizon feedback $u(t) = -R^{-1}B^T P(0)x(t)$, a sufficient condition for stability is in terms of the terminal

cost weighting P_T , given an adjustment on the state weighting Q . The adjustment results in the ‘‘Fake Algebraic Riccati Equation’’.¹⁰

Finite horizon and CLF

We now consider the more general case of the nonlinear plants. In an inverse optimal control problem for a nonlinear system

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

the objective is to demonstrate that a given state feedback $u = k(x)$ is optimal for a cost functional

$$J = \int_0^T L(x, u) dt = \int_0^T W(x) + u^T R(x)u dt \quad (6)$$

where $W(x)$ is a positive definite loss function and $R(x) > 0$ is a positive definite weight. The terms of the integrated cost $L(x, u)$ need to be determined from the feedback $k(x)$ and a Lyapunov function $V(x)$ associated with the closed-loop system in (5).

Let us denote the directional derivative of smooth function $V(x)$ along $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ by the column vector

$$L_f V(x) = [\nabla V(x)f(x)]^T.$$

A smooth positive definite and radially unbounded function $V(x)$ is called a *control Lyapunov function* (CLF) for the nonlinear system in (5) if it satisfies the following property

$$L_g V(x) = 0 \implies L_f V(x) < 0, \quad \forall x \neq 0. \quad (7)$$

Sontag¹⁷ provided an explicit formula for a stabilizing state feedback $u = -s(x)L_g V(x)$ with $s(x) > 0$ given that $V(x)$ is a CLF for (5).

In the following, we demonstrate that the controller from Sontag’s formula is optimal for a finite horizon optimal control problem with integrated cost $L(x, u)$ in the form given in (6) and a terminal cost that is a CLF function.

Theorem 1. *Let $s(x)$ be a function defined as*

$$s(x) = \begin{cases} 1 + \frac{a(x) + \sqrt{a(x)^2 + (b(x)^T b(x))^2}}{b(x)^T b(x)} & b(x) \neq 0 \\ 1 & b(x) = 0, \end{cases} \quad (8)$$

with $a(x) = L_f V(x)$ and $b(x) = L_g V(x)$.¹⁷ Then, for all $\gamma \geq 1/2$, $V(x)$ is a Lyapunov function for the closed loop system with the state feedback

$$u = k_\gamma(x) := -\gamma s(x)L_g V(x). \quad (9)$$

Moreover, for all $\gamma \geq 1$, $u = k_\gamma(x)$ is inverse optimal for a finite horizon cost functional with a positive definite integrated cost that is quadratic in u and a terminal cost $V(x(T))$.

Proof. Given the controller $u = k_\gamma(x)$, after calculating \dot{V} along the solutions of (5) with $x \in \mathbb{R}^n$ and $u \in \mathbb{R}^m$, we get

$$\begin{aligned}\dot{V}(x) &= L_f V(x) + u^T L_g V(x) \\ &= a(x) - \gamma s(x) b(x)^T b(x).\end{aligned}$$

By definition of the CLF $V(x)$ and the function $s(x)$, if $b(x) = 0$ then $\dot{V} = a(x) = L_f V(x) < 0$, for all $x \neq 0$. Suppose $b(x) \neq 0$ and $a(x) \geq 0$ (otherwise, trivially $\dot{V} < 0$ for all $x \neq 0$). Then for all $\gamma \geq \frac{1}{2}$, we have

$$\begin{aligned}\dot{V} &= (1 - \gamma)a(x) - \gamma \sqrt{a(x)^2 + (b(x)^T b(x))^2} \\ &\quad - \gamma(b(x)^T b(x)) \\ \dot{V} &= \frac{(1 - 2\gamma)a(x)^2 - \gamma^2(b(x)^T b(x))^2}{(1 - \gamma)a(x) + \gamma \sqrt{a(x)^2 + (b(x)^T b(x))^2}} \\ &\quad - \gamma(b(x)^T b(x)) < 0, \quad \forall x \neq 0.\end{aligned}$$

In other words, setting $W(x, \gamma) := -\dot{V}(x)$, one gets $W(x, \gamma) > 0$, for all $x \neq 0$. Furthermore, $V(x)$ is a Lyapunov function for the closed loop system with $u = k_\gamma(x)$ for $\gamma \geq \frac{1}{2}$. To show that $u = k_\gamma(x)$ is an inverse optimal control, let us consider an intermediate cost functional given by

$$J_0 = \int_0^T (u + \gamma s(x) L_g V(x))^T R(x) \cdot (u + \gamma s(x) L_g V(x)) dt, \quad \gamma \geq 1 \quad (10)$$

with an obvious optimal control $u^* = -\gamma s(x) L_g V(x)$ that minimizes J_0 . The optimal value for u^* is $J_0^* = 0$. Following Jankovic,⁴ assume $2s(x)R(x) = I_n$ and let $\Phi(x, u)$ denote the integrand of J_0 . Due to $u^T L_g V(x) = \dot{V}(x) - L_f V(x)$, we get

$$\begin{aligned}\Phi(x, u) &= u^T R(x)u + \gamma^2 s(x)^2 (L_g V(x))^T R(x) \cdot \\ &\quad (L_g V(x)) + 2\gamma s(x) u^T R(x) L_g V(x), \\ &= u^T R(x)u + \frac{\gamma^2}{4} (L_g V(x))^T R(x)^{-1} \cdot \\ &\quad (L_g V(x)) + \gamma(\dot{V}(x) - L_f V(x)) \\ &= u^T R(x)u - \gamma[L_f V(x) + \\ &\quad (\frac{\gamma}{2} s(x) L_g V(x))^T L_g V(x)] + \gamma \dot{V}(x) \\ &= \gamma W(x, \frac{\gamma}{2}) + u^T R(x)u + \gamma \dot{V}(x).\end{aligned}$$

Since $\int_0^T \dot{V}(x) dt = V(x(T)) - V(x(0))$ and the optimal value function J_0^* is equal to zero, we get

$$V(x(0)) = \frac{1}{\gamma} \int_0^T L(x, u) dt + V(x(T)),$$

where $L(x, u) = \gamma W(x, \frac{\gamma}{2}) + u^T R(x)u$. Thus, $u^* = k_\gamma(x)$ is an optimal control for the following cost functional

$$J = \int_0^T W(x, \frac{\gamma}{2}) + \frac{1}{\gamma} u^T R(x)u dt + V(x(T)), \quad \gamma \geq 1 \quad (11)$$

with an optimal value function $J^* = V(x(0))$. \square

Remark 1. Relatively large values of $\gamma \gg 1$ correspond to relatively small weights $1/\gamma$ on the control. This is consistent with the intuition for the so-called “cheap control” case in a standard LQR problem.

Remark 2. The aforementioned analysis of inverse optimality of Sontag’s CLF-based controller is somewhat similar to the proof given in Jankovic⁴ for an infinite horizon inverse optimality of $k_\gamma(x)$ with $\gamma = 1$. Here, we have generalized the result to a finite horizon inverse optimality of $u = k_\gamma(x)$ for all $\gamma \geq 1$.

The existence of a CLF for the nonlinear system in (5) allows us to formulate an unconstrained finite horizon optimal control problem. This motivates us to attempt solving a constrained finite horizon optimal control problem with an integrated cost that is obtained from analysis of the unconstrained problem with a loss function $W(x) = -\dot{V}(x) > 0$ for all $x \neq 0$ and a quadratic penalty on control u .

Attitude alignment for multiple vehicles

We illustrate how these results can be used by considering the problem of aligning a group of robots in a common direction using a decentralized controller. Consider a group of n robots where θ_i denotes the heading angle of the i th robot. Let $\dot{\theta}_i = \omega_i$ for $i = 1, \dots, n$ be the attitude kinematics of each robot with the angular velocity as the control input. Assume that each robot can only use its own attitude and the attitude of its neighboring robots on a graph G . Let $A = [a_{ij}]$ be a nonnegative adjacency matrix of the graph with the property $A^T = A$. If robots i and j are neighbors (or adjacent) then $a_{ij} > 0$. Otherwise, $a_{ij} = 0$. The set of neighbors of i are denoted by $N_i = \{j : a_{ij} > 0\}$. The objective in the alignment problem is to design a controller $\omega = k(\theta)$ for the system $\dot{\theta} = \omega$ (where $\theta = (\theta_1, \dots, \theta_n)^T$) such that the closed loop system $\dot{\theta} = k(\theta)$ has a globally asymptotically stable equilibrium point θ^* with identical elements, i.e. $\theta_i^* = \theta_j^*$ for all i, j . Based on previous work by the authors,^{11,12} the following linear and distributed feedback

$$\omega_i = \sum_{j \in N_i} a_{ij} (\theta_j - \theta_i) \quad (12)$$

solves the alignment problem. This linear feedback can be rewritten as

$$\omega = -L\theta \quad (13)$$

where the matrix $L = [l_{ij}]$ is called *graph Laplacian* and its elements are defined by

$$l_{ij} = \begin{cases} \sum_{j=1}^n a_{ij}, & i = j; \\ -a_{ij}, & i \neq j. \end{cases} \quad (14)$$

The Laplacian L has a simple eigenvalue at zero associated with $e = (1, \dots, 1)^T$. All other eigenvalues of L are positive. It turns out that $V(\theta) = \theta^T L \theta$ satisfies the following property

$$V(\theta) = \sum_{i < j} a_{ij} (\theta_j - \theta_i)^2$$

and is a (weak) CLF for $\dot{\theta} = \omega$. Defining the loss function $W(\theta) = -\dot{V}(\theta) = \theta^T L^2 \theta$, one can show that the alignment rule in equation (12) is inverse optimal for the following linear-quadratic finite horizon cost functional

$$J = \int_0^T \theta^T L^2 \theta + \omega^T \omega dt + \theta_f^T L \theta_f, \quad \theta_f = \theta(T). \quad (15)$$

Now, consider the same attitude alignment problem under the constraint of bounded control inputs. It can be shown that the following positive definite and radially unbounded function

$$V_b(\theta) = \sum_{i < j} a_{ij} \psi(\theta_j - \theta_i) \quad (16)$$

with $\psi(z) = \sqrt{1 + z^2} - 1 \geq 0$ is a CLF for $\dot{\theta} = \omega$ and $\omega = -(\nabla V_b(\theta))^T$ is a uniformly bounded inverse optimal distributed control for an integrated cost in the form $L(\theta, \omega) = \|\nabla V_b(\theta)\|^2 + \sum_i \omega_i^2$. This is because $W(\theta) = -\dot{V}_b = -\nabla V_b \cdot (-\nabla V_b)^T = \|\nabla V_b\|^2$.

Simulation Study for Planar Flight Vehicle

In this section, we shall consider the dynamics of a simple planar flight vehicle that is propelled by two independently controlled ducted fan engines, as shown in Figure 2. The vehicle is nonlinear and underactuated, with control constraints. Denoting the configuration $(x, y, \theta) \in SE(2)$ and assuming viscous friction, the equations of motion of the vehicle are

$$\begin{aligned} m\ddot{x} &= -\eta\dot{x} + (F_s + F_p) \cos \theta \\ m\ddot{y} &= -\eta\dot{y} + (F_s + F_p) \sin \theta \\ J\ddot{\theta} &= -\psi\dot{\theta} + (F_s - F_p)r_J. \end{aligned}$$



Fig. 2 The Caltech Multi-Vehicle Wireless Testbed (MVWT).

The starboard and port fan forces are denoted F_s and F_p , respectively, with $F_s, F_p \in [0, F_{\max}]$, and r_J denotes the (common) moment arm of the forces. The physical parameter values are $m = 5.0$ kg, $J = 0.05$ kg-m², $\eta = 4.5$ kg/s, $\psi = 0.084$ kg-m²/s, $r_J = 0.123$ m, and the maximum force is $F_{\max} = 5.0$ N. For the example we consider, the polar coordinate form of the dynamics will be easier to handle, given as

$$\begin{aligned} m(r\ddot{\beta} + 2\dot{r}\dot{\beta}) &= -\eta r\dot{\beta} + (F_s + F_p) \sin(\theta - \beta) \\ m(\ddot{r} - r\dot{\beta}^2) &= -\eta\dot{r} + (F_s + F_p) \cos(\theta - \beta) \\ J\ddot{\theta} &= -\psi\dot{\theta} + (F_s - F_p)r_J. \end{aligned} \quad (17)$$

The polar coordinate state is denoted $Z = (r, \beta, \theta, \dot{r}, \dot{\beta}, \dot{\theta})$. The testbed floor on which the vehicle operates is bounded and has an obstacle in the center, requiring that the radius satisfy the constraint

$$\begin{aligned} R_{\min} &\leq r \leq R_{\max} \\ R_{\min} &= 0.8 \text{ m} \quad R_{\max} = 3.0 \text{ m}. \end{aligned} \quad (18)$$

An equilibrium point for (17) is any constant position with zero velocity; however, these dynamics are not linearly controllable around any such equilibrium point. To recover linear controllability, we consider the dynamics linearized around a reference trajectory $Z_r = (r_r, \beta_r, \theta_r, \dot{r}_r, \dot{\beta}_r, \dot{\theta}_r)$ that is a circular path with constant radius and angular velocity, parameterized as

$$\begin{aligned} (r_r, \beta_r, \theta_r, \dot{r}_r, \dot{\beta}_r, \dot{\theta}_r) &= \\ (\rho, \xi t + \beta_{r0}, \xi t + \theta_{r0}, 0, \xi, \xi) \quad \rho > 0, \xi \neq 0. \end{aligned} \quad (19)$$

The values for β_{r0} and θ_{r0} must satisfy the equation $\theta_{r0} = \beta_{r0} + \alpha_0 + \pi \text{sign}(\xi)/2$, $\alpha_0 = \arctan(m\xi/\eta)$,




Fig. 3 LQR: Cartesian state and control input responses to tracking a circular path. The dashed line is the reference trajectory and the solid line is the closed loop response.

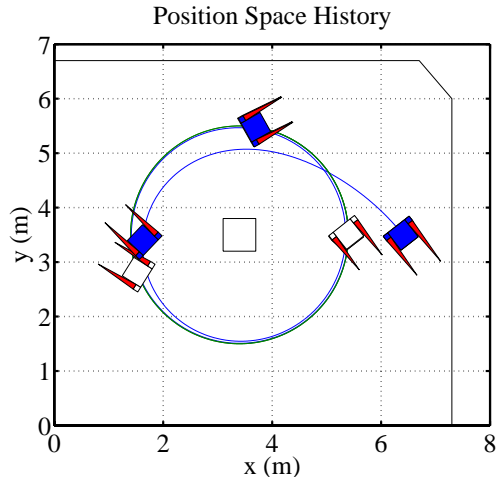


Fig. 4 LQR tracking and reference vehicle in position space on the multi-vehicle wireless testbed floor. The dark vehicle represents the tracking vehicle and the white vehicle the reference.




Fig. 5 RHC: Cartesian state and control input responses to tracking a circular path.

ence. The cone lengths are proportional to the fan forces. The vehicles are shown at time snapshots of 0, 5 and 20 seconds. At 20 seconds, the vehicles are on top of one another, meaning the tracking vehicle has met its objective.

We now compare the LQR controller with a RHC law, from the same initial condition. The horizon and update times are chosen to be $T = 3.0$ and $\delta = 1.0$ seconds. The RHC tracking performance is shown by the Cartesian state and control input responses in Figure 5. The response is quite similar to the LQR performance. An advantage of the RHC approach is that the initial values of the forces can be set as an initial constraint in the optimal control problem. For comparison, we set the initial force values in the first optimization to be approximately equal to the initial values of the LQR forces. Subsequent updates enforce the initial forces to be equal to the previous forces applied at the update time. The forces are thus continuous in time, although not differentiable at the receding horizon updates, as indicated by the figure.

The tracking vehicle and reference vehicle are shown in x, y space in Figure 6.




Fig. 6 RHC tracking and reference vehicle in position space on the multi-vehicle wireless testbed floor.

LQR vs. RHC: Example 2 (Nonlinear Regime)

Consider a second simulation example using the same reference trajectory. The initial condition of the vehicle is now given by $Z(0) = (\rho - 1, \beta_{r0} - 3\pi/4, \theta_{r0}, 0, 0, 0)$. The LQR tracking performance is shown by the Cartesian state and control input responses in Figure 3. The linear controller violates the state constraints in (18) and saturates the controls for at least 1 second. The violation of the state constraint is observable by the position trajectory in Figure 8. Although the simulation permits the vehicle to pass through the obstacle, the actual vehicle in the testbed environment would likely get stuck. Since the obstacle has flat edges and the fans are unidirectional, on several occasions we have observed that once a vehicle hits the obstacle, it remains there, only able to thrust into the fixed boundary. The LQR feedback substantially violates the input constraints, but the simulation enforces the saturation, as would the actual vehicle fans.

For the same horizon and update time as in the previous example, the RHC tracking performance is shown by the Cartesian state and control input responses in Figure 5 and the position trajectory in Figure 10. The latter figure shows that the state constraints are satisfied throughout the entire transient response, where the vehicle meets the tracking objective after roughly 15 seconds according to the state responses. Moreover, the RHC force inputs satisfy the thrust constraints. In Figure 10, the time snapshots of 0, 5 and 20 seconds are printed on the top of each vehicle, to clarify the location of the actual vehicle with respect to the reference vehicle. For the first 10 seconds, the actual vehicle is

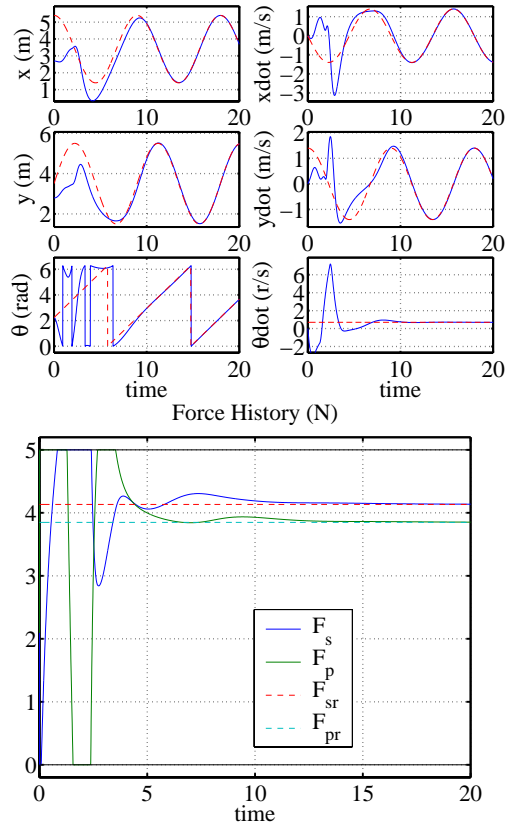


Fig. 7 LQR: Cartesian state and control input responses to tracking a circular path.




Fig. 8 LQR tracking and reference vehicle in position space on the multi-vehicle wireless testbed floor.




Fig. 9 RHC: Cartesian state and control input responses to tracking a circular path.

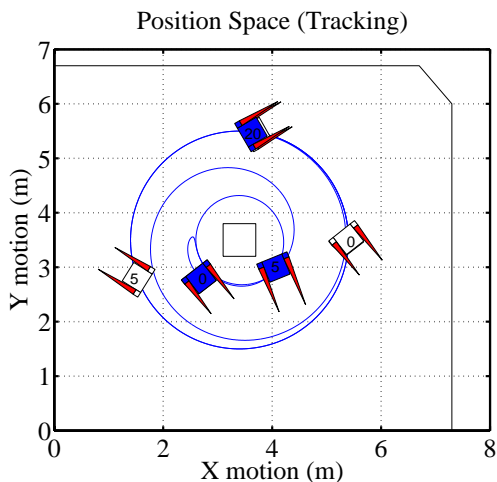


Fig. 10 RHC tracking and reference vehicle in position space on the multi-vehicle wireless testbed floor.

operating on the lower boundary of the constraint in equation (18).

By the simulation examples we have shown that the chosen RHC policy performs similarly to the LQR controller when the dynamics are in the linear regime. Further, the RHC policy is able to handle the state and control constraints directly, while such nonlinearities disrupt the LQR performance, and on the real testbed, generally lead to instability.

Summary and Future Work

In this paper we have presented a framework for control design that combines the advantages of modern linear and nonlinear control techniques with the computational power now available for use in optimization-based control. By using linear and nonlinear tools to generate a controller around a representative operating point, we can generate a specification for a receding horizon optimal control problem. This specification can then be used to generate optimization-based controllers capable of taking into account additional nonlinearities and input/state constraints.

We have focused in this paper on state space controller formulations, building on early work by Kalman on inverse optimal control. To be truly useful, the techniques must be extended to cover dynamic compensators, of the sort that modern control methods generate. One approach to doing this is to realize the controller as a generalized observer followed by state feedback. We can then extend the observer to the constrained nonlinear system and use the state space control gains to generate the cost function for optimization-based control. The viability of this approach is the subject of future work.

Acknowledgments This work has been funded in part by the DARPA Software Enabled Control (SEC) program, under contract F33615-98-C-3613. The authors would like to thank Petar Kokotovic, Andy Teel, and Steve Waydo for early discussions on this work.

References

- ¹R. M. Murray *et al.* Online control customization via optimization-based control. In T. Samad and G. Balas, editors, *Software-Enabled Control: Information Technology for Dynamical Systems*. IEEE Press, 2003.
- ²P. Gill, W. Murray, M. Saunders, and M. Wright. *User's guide for NPSOL 5.0: A FORTRAN package for nonlinear programming*. Systems Optimization Laboratory, Stanford University, Stanford, CA 94305, 1998.
- ³A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 46, May 2001.

⁴M. Jankovic, R. Sepulchre, and P. V. Kokotović. CLF based designs with robustness to dynamic input uncertainties. *Systems Control Letters*, 37:45–54, 1999.

⁵R. E. Kalman. When is a linear control system optimal? *Trans. ASME Journal of Basic Engineering*, 86:51–60, March 1964.

⁶D.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.

⁷M. Milam, K. Mushambi, and R. M. Murray. A new computational approach to real-time trajectory generation for constrained mechanical systems. *Proc. of the 39th IEEE Conf. on Decision and Control*, 1:845–551, 2000.

⁸Mark B. Milam, Ryan Franz, John E. Hauser, and Richard M. Murray. Receding horizon control of a vectored thrust flight experiment. *IEE Proceedings on Control Theory and Applications*, 2003. Submitted.

⁹P. J. Moylan and B. D. O. Anderson. Nonlinear regulator theory and an inverse optimal control problem. *IEEE Trans. on Automatic Control*, 18(5):460–454, 1973.

¹⁰G. De Nicolao and R.R. Bitmead. Fake Riccati equations for stable receding-horizon control. *Proc. of the European Control Conference*, 1997.

¹¹R. Olfati Saber and R. M. Murray. Agreement problems in networks with directed graphs and switching topology. *Conference on Decision and Control*, 2003. To appear.

¹²R. Olfati Saber and R. M. Murray. Consensus protocols for networks of dynamic agents. *Proc. of the American Control Conference*, 2003.

¹³J. A. Primbs, V. Nevistić, and J. C. Doyle. A receding horizon generalization of pointwise min-norm controllers. *IEEE Transactions on Automatic Control*, 45:898–909, 2000.

¹⁴S.J. Qin and T.A. Badgwell. An overview of industrial model predictive control technology. In J.C. Kantor, C.E. Garcia, and B. Carnahan, editors, *Fifth International Conference on Chemical Process Control*, pages 232–256, 1997.

¹⁵R. Sepulchre, M. Jankovic, and P. V. Kokotović. *Constructive Nonlinear Control*. Springer, London, 1997.

¹⁶L. Singh and J. Fuller. Trajectory generation for a UAV in urban terrain, using nonlinear MPC. In *Proceedings of the American Control Conference*, 2001.

¹⁷E.D. Sontag. A universal construction of Arstein’s theorem on nonlinear stabilization. *Systems Control Letters*, 13:117–123, 1989.