



May 2001

Unconstrained Receding-Horizon Control of Nonlinear Systems

Ali Jadbabaie

University of Pennsylvania, jadbabai@seas.upenn.edu

Jie Yu

California Institute of Technology

John Hauser

University of Colorado

Follow this and additional works at: https://repository.upenn.edu/ease_papers

Recommended Citation

Ali Jadbabaie, Jie Yu, and John Hauser, "Unconstrained Receding-Horizon Control of Nonlinear Systems", . May 2001.

Copyright 2001 IEEE. Reprinted from *IEEE Transactions on Automatic Control*, Volume 46, Number 5, May 2001, pages 776-783.

Publisher URL: <http://dx.doi.org/10.1109/9.920800>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

NOTE: At the time of publication, author Ali Jadbabaie was affiliated with the California Institute of Technology. Currently (March 2005), he is a faculty member in the Department of Electrical and Systems Engineering at the University of Pennsylvania.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ease_papers/89
For more information, please contact repository@pobox.upenn.edu.

Unconstrained Receding-Horizon Control of Nonlinear Systems

Abstract

It is well known that unconstrained infinite-horizon optimal control may be used to construct a stabilizing controller for a nonlinear system. In this note, we show that similar stabilization results may be achieved using *unconstrained* finite horizon optimal control. The key idea is to approximate the tail of the infinite horizon cost-to-go using, as terminal cost, an appropriate control Lyapunov function. Roughly speaking, the terminal control Lyapunov function (CLF) should provide an (incremental) upper bound on the cost. In this fashion, important stability characteristics may be retained *without* the use of terminal constraints such as those employed by a number of other researchers. The absence of constraints allows a significant speedup in computation. Furthermore, it is shown that in order to guarantee stability, it suffices to satisfy an improvement property, thereby relaxing the requirement that truly *optimal* trajectories be found. We provide a complete analysis of the stability and region of attraction/operation properties of receding horizon control strategies that utilize finite horizon approximations in the proposed class. It is shown that the guaranteed region of operation contains that of the CLF controller and may be made as large as desired by increasing the optimization horizon (restricted, of course, to the infinite horizon domain). Moreover, it is easily seen that both CLF and infinite-horizon optimal control approaches are limiting cases of our receding horizon strategy. The key results are illustrated using a familiar example, the inverted pendulum, where significant improvements in guaranteed region of operation and cost are noted.

Keywords

Lyapunov methods, control system synthesis, nonlinear control systems, optimal control, stability, control Lyapunov function, guaranteed region of operation, infinite horizon cost-to-go, inverted pendulum, region of attraction, stability characteristics, terminal cost, unconstrained finite horizon optimal control, unconstrained receding-horizon control, control Lyapunov functions (CLFs), model predictive control, nonlinear control design, receding horizon control

Comments

Copyright 2001 IEEE. Reprinted from *IEEE Transactions on Automatic Control*, Volume 46, Number 5, May 2001, pages 776-783.

Publisher URL: <http://dx.doi.org/10.1109/9.920800>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

NOTE: At the time of publication, author Ali Jadbabaie was affiliated with the California Institute of Technology. Currently (March 2005), he is a faculty member in the Department of Electrical and Systems Engineering at the University of Pennsylvania.

Unconstrained Receding-Horizon Control of Nonlinear Systems

Ali Jadbabaie, Jie Yu, and John Hauser

Abstract—It is well known that unconstrained infinite-horizon optimal control may be used to construct a stabilizing controller for a nonlinear system. In this note, we show that similar stabilization results may be achieved using unconstrained finite horizon optimal control. The key idea is to approximate the tail of the infinite horizon cost-to-go using, as terminal cost, an appropriate control Lyapunov function. Roughly speaking, the terminal control Lyapunov function (CLF) should provide an (incremental) upper bound on the cost. In this fashion, important stability characteristics may be retained without the use of terminal constraints such as those employed by a number of other researchers. The absence of constraints allows a significant speedup in computation. Furthermore, it is shown that in order to guarantee stability, it suffices to satisfy an improvement property, thereby relaxing the requirement that truly optimal trajectories be found. We provide a complete analysis of the stability and region of attraction/operation properties of receding horizon control strategies that utilize finite horizon approximations in the proposed class. It is shown that the guaranteed region of operation contains that of the CLF controller and may be made as large as desired by increasing the optimization horizon (restricted, of course, to the infinite horizon domain). Moreover, it is easily seen that both CLF and infinite-horizon optimal control approaches are limiting cases of our receding horizon strategy.

The key results are illustrated using a familiar example, the inverted pendulum, where significant improvements in guaranteed region of operation and cost are noted.

Index Terms—Control Lyapunov functions (CLFs), model predictive control, nonlinear control design, optimal control, receding horizon control.

I. INTRODUCTION

Receding horizon control strategies, also known as model predictive control (MPC), have become quite popular recently. This interest is partly due to the availability of faster and cheaper computers as well as efficient numerical algorithms for solving optimization problems. Another key advantage of these strategies is the potential ability to handle control saturations. Many of the successful applications of receding horizon control methods have been in the area of chemical process control. This is due, in part, to the fact that a number of important industrial chemical processes are open-loop stable so that stability is not a primary concern for many of these methods. However, application of these methods to unstable plants has appeared to be more difficult. Several researchers have attempted to address the problem of stability for receding horizon control to allow its application in stability critical areas.

Keerthi and Gilbert [10] imposed a terminal state equality constraint $x(t+T) = 0$. This results in a finite-horizon optimization problem which turns out to be computationally demanding. Michalska and Mayne [13] 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 Parisini and Zoppoli [15] and later by Chen

and Allgöwer [5], 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 Magni and Sepulchre [12] and later by De Nicolao *et al.* [14], stability of the receding horizon scheme is guaranteed (for continuous-time and discrete-time systems, respectively) by using a (possibly nonquadratic) 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 $[0, T]$. Moreover, it is assumed that the optimization is performed over *admissible* control sequences, i.e., control sequences which guarantee that at the end of the horizon the state has reached a suitable neighborhood of the origin which is an exponential stability region for the linear controller. In other words, a state inequality constraint is implicitly imposed.

An approach for the receding horizon control of globally stabilizable nonlinear systems was developed by Primbs *et al.* [16]. 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 to be negative and also that the decrease in the value of the CLF be greater than that obtained using the controller derived from the CLF. There are a variety of methods that can exploit system structure, e.g., differential flatness and backstepping, to construct suitable CLFs. Although the current results rely on a global CLF, it is clear that local versions of this approach may be developed.

An alternative approach was developed by the authors in [9], [8]. This approach obtains stability guarantees through the use of an *a priori* CLF as a terminal cost rather than by imposing state inequality (or equality) constraints. The attendant speedup in calculations can be dramatic. Moreover, stability continues to be guaranteed as long as the CLF is an (appropriate) upper bound on the cost-to-go.

The terminal *cost* should be thought of as an *approximation* to the (infinite horizon) value function rather than as a terminal *penalty*. Indeed, simulation results in [8] indicate that, contrary to what one might think, a mere upper bound on the cost-to-go does not generally provide an appropriate terminal cost.

Since it is rarely possible to obtain a global CLF (as most systems are not even globally stabilizable), it is desirable to be able to estimate the region of attraction of a receding horizon controller. In particular, one would like to know whether region of attraction (or operation) estimates for the receding horizon system contain those of the CLF controlled system and to what extent these regions may be expanded, e.g., to the regions for the infinite horizon controller, by increasing the horizon length.

These issues are completely addressed in this paper.

This paper is organized as follows. The problem setting is described in Section II. In Section III, we explore the important relationships between an infinite horizon optimal control problem and its finite horizon approximations and present the main results. The optimality condition is replaced by an improvement property in Section IV. The key results are illustrated in Section V using an inverted pendulum example. Finally, our conclusions are presented in Section VI.

II. PROBLEM SETTING

The nonlinear system under consideration is

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

Manuscript received May 5, 1999; revised February 11, 2000 and October 4, 2000. Recommended by Associate Editor G. De Nicolao. This work was supported in part by AFOSR and DARPA.

A. Jadbabaie and J. Yu are with Control and Dynamical Systems, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: alij@cds.caltech.edu; jyu@cds.caltech.edu).

J. Hauser is with the Electrical and Computer Engineering and Aerospace Engineering Departments, University of Colorado, Boulder, CO 80309-0425 USA (e-mail: hauser@schof.colorado.edu).

Publisher Item Identifier S 0018-9286(01)03642-X.

where the vector field $f: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is C^2 and possesses a linearly controllable critical point at the origin, e.g., $f(0, 0) = 0$ and $(A, B) := (D_1 f(0, 0), D_2 f(0, 0))$ is controllable. We require the set $f(x, \mathbb{R}^m) \subset \mathbb{R}^n$ to be convex for each $x \in \mathbb{R}^n$. Given an initial state x and a control trajectory $u(\cdot)$, the state trajectory $x^u(\cdot; x)$ is the (absolutely continuous) curve in \mathbb{R}^n satisfying

$$x^u(t; x) = x + \int_0^t f(x^u(\tau; x), u(\tau)) d\tau$$

for $t \geq 0$. We require that the trajectories of the system satisfy an *a priori* bound

$$\|x(t)\| \leq \beta(x, T, \|u(\cdot)\|_1) < \infty, \quad t \in [0, T]$$

where β is continuous in all variables and monotone increasing in T and $\|u(\cdot)\|_1 = \|u(\cdot)\|_{L_1(0, T)}$. Most models of physical systems will satisfy a bound of this type.

The performance of the system will be measured by a given incremental cost $q: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ that is C^2 and fully penalizes both state and control according to

$$q(x, u) \geq c_q(\|x\|^2 + \|u\|^2), \quad x \in \mathbb{R}^n, u \in \mathbb{R}^m$$

for some $c_q > 0$ and $q(0, 0) = 0$. We further require that the function $u \mapsto q(x, u)$ be convex for each $x \in \mathbb{R}^n$. These conditions imply that the quadratic approximation of q at the origin is positive-definite, $D^2 q(0, 0) \geq c_q I > 0$.

We will also suppose that f and q are sufficiently compatible to uniquely define a C^2 Hamiltonian for the (optimized) system. In particular, we will require that there is a C^2 function $\bar{u}^*: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^m$: $(x, p) \mapsto \bar{u}^*(x, p)$ providing a global minimum of the pre-Hamiltonian $K(x, p, u) := p^T f(x, u) + q(x, u)$ so that the Hamiltonian $H(x, p) := K(x, p, \bar{u}^*(x, p))$ is C^2 . Such a $\bar{u}^*(\cdot, \cdot)$ is locally guaranteed by the implicit function theorem (though we would require $f, q \in C^3$). Note that this condition is trivially satisfied for control affine f and quadratic q for then $u \mapsto K(x, p, u)$ is strictly convex and *real analytic*.

The cost of applying a control $u(\cdot)$ from an initial state x over the infinite-time interval $[0, \infty)$ is given by

$$J_\infty(x, u(\cdot)) = \int_0^\infty q(x^u(\tau; x), u(\tau)) d\tau.$$

The optimal cost (from x) is given by

$$J_\infty^*(x) = \inf_{u(\cdot)} J_\infty(x, u(\cdot))$$

where the control functions $u(\cdot)$ belong to some reasonable class of admissible controls (e.g., piecewise continuous). The function $x \mapsto J_\infty^*(x)$ is often called the *optimal value function* for the infinite horizon optimal control problem. For the class of f and q considered, we know that J_∞^* is a positive definite C^2 function on a neighborhood of the origin. This follows from the *geometry* of the corresponding Hamiltonian system [19], [20]. In particular, since $(x, p) = (0, 0)$ is a hyperbolic critical point of the Hamiltonian vector field $X_H(x, p) := (D_2 H(x, p), -D_1 H(x, p))^T$, the local properties of J_∞^* are determined by the linear-quadratic approximation to the problem and, moreover, $D^2 J_\infty^*(0) = P > 0$ where P is the stabilizing solution of the appropriate algebraic Riccati equation.

For practical purposes, we are interested in approximating the infinite horizon optimization problem with one over a finite horizon. In

particular, let V be a nonnegative C^2 function and define the finite horizon cost [from x using $u(\cdot)$] to be

$$J_T(x, u(\cdot)) = \int_0^T q(x^u(\tau; x), u(\tau)) d\tau + V(x^u(T; x))$$

and denote the optimal cost (from x) as

$$J_T^*(x) = \inf_{u(\cdot)} J_T(x, u(\cdot)).$$

As in the infinite horizon case, one can show, by geometric means, that J_T^* is locally smooth (C^2). Other properties, e.g., local positive definiteness, will depend on the choice of V and T .

Let Γ^∞ denote the domain of J_∞^* (the subset of \mathbb{R}^n on which J_∞^* is finite). It is not too difficult to show that the cost functions J_∞^* and J_T^* , $T \geq 0$ are continuous functions on Γ^∞ using the same arguments as in [1, Prop. 3.1]. We make the following assumption.

Standing Assumption (SA): The minimum value of cost functions J_∞^* , J_T^* , $T \geq 0$, is attained.

The assumption (SA) guarantees the existence of an optimal trajectory

$$(x_T^*(t; x), u_T^*(t; x)), \quad t \in [0, T]$$

such that

$$J_T(x, u_T^*(\cdot; x)) = J_T^*(x).$$

Continuity of $u_T^*(\cdot; x)$ follows directly from Pontryagin's Maximum Principle. This trajectory is not necessarily unique. In fact, in examples one finds two trajectories of equal (minimal) cost originating at points where J_T^* is only continuous (and not differentiable). Under assumptions of the sort given (convexity, boundedness, etc.), one can turn (SA) into a proposition. This involves the use of techniques from regularity theory and the direct methods of the calculus of variations, see [4] and [3].

It is easy to see that J_∞^* is proper on its domain so that the sublevel sets

$$\Gamma_r^\infty := \{x \in \Gamma^\infty : J_\infty^*(x) \leq r^2\}$$

are compact and path connected and moreover $\Gamma^\infty = \bigcup_{r \geq 0} \Gamma_r^\infty$. Note also that Γ^∞ may be a proper subset of \mathbb{R}^n since there may be states that cannot be driven to the origin. We use r^2 (rather than r) here to reflect the fact that our incremental cost is quadratically bounded from below. We refer to sub-level sets of J_T^* and V using

$$\Gamma_r^T := \text{path connected component of } \{x \in \Gamma^\infty : J_T^*(x) \leq r^2\} \\ \text{containing } 0$$

and

$$\Omega_r := \text{path connected component of } \{x \in \mathbb{R}^n : V(x) \leq r^2\} \\ \text{containing } 0.$$

III. INFINITE AND FINITE HORIZON OPTIMIZATION

In this section, we explore some of the relationships between an infinite-horizon optimal control problem and its finite-horizon approximations. We will show that the use of an appropriate terminal cost allows us to retain desirable features of the infinite-horizon problem.

A. Infinite Horizon Properties

What infinite horizon problem properties are interesting for finite horizon approximations and, in particular, are useful for receding

horizon strategies? This is a question that we intend to answer in this section.

Let $(x_\infty^*, u_\infty^*)(\cdot; x)$ be any optimal trajectory originating at x . Then, for any $\delta > 0$, we have

$$J_\infty^*(x_\infty^*(\delta; x)) = J_\infty^*(x) - \int_0^\delta q(x_\infty^*(\tau; x), u_\infty^*(\tau; x)) d\tau. \quad (1)$$

Since, by (SA), $\tau \mapsto u_\infty^*(\tau; x)$ is continuous for $\tau \geq 0$, we see that

$$\begin{aligned} \lim_{\delta \searrow 0} \frac{J_\infty^*(x_\infty^*(\delta; x)) - J_\infty^*(x)}{\delta} \\ = -q(x, u_\infty^*(0; x)) \leq -c_q \|x\|^2 \end{aligned}$$

so that J_∞^* possesses a well defined (negative-definite) directional derivative in each optimal direction $f(x, u_\infty^*(0; x))$. In fact, we may write

$$\dot{J}_\infty^*(x, u_\infty^*(0; x)) + q(x, u_\infty^*(0; x)) = 0 \quad (2)$$

where $\dot{J}_\infty^*(x, u)$ is the directional derivative of J_∞^* in the direction $f(x, u)$ (when it exists). [At points of differentiability, $\dot{J}_\infty^*(x, u) = D J_\infty^*(x) \cdot f(x, u)$].

We conclude that each sub-level set Γ_r^∞ , $r > 0$, is positively invariant under optimal actions, both incremental ($\delta > 0$) and infinitesimal. Also, in all cases, these sets are attracted to the origin exponentially fast. In particular, the (not necessarily unique) feedback $u = k_\infty(x) := u_\infty^*(0; x)$ exponentially stabilizes the origin.

B. Finite Horizon Properties

As noted above, one may use optimal (infinite-horizon) actions to provide a stabilizing feedback for a nonlinear system. It is natural to expect that a similar result would be possible using a finite-horizon optimization. For instance, one could implement a *receding horizon* scheme as follows. From the current state $x(t)$, obtain an optimal trajectory $(x_T^*, u_T^*)(\tau; x(t))$, $\tau \in [0, T]$, and use as feedback $u(t) = u_T^*(0; x(t))$. (This feedback is not uniquely defined at points where more than one optimal trajectory is available.) This approach requires one to continuously re-solve the finite horizon optimization. An alternative scheme is to solve the finite horizon optimization every $\delta > 0$ seconds and use the control trajectory $u_T^*(\tau; x(t))$, $\tau \in [0, \delta]$, to drive the system from $x(t)$ at time t to $x_T^*(\delta, x(t))$ at time $t + \delta$. [Practically speaking, a better idea is to use a local tracking controller to regulate the system about the desired trajectory $(x_T^*, u_T^*)(\tau; x(t))$, $\tau \in [0, \delta]$.] We will denote this receding horizon scheme as $\mathcal{RH}(T, \delta)$. One might also consider using a variable $\delta_k > 0$, which will be denoted as $\mathcal{RH}(T, \{\delta_k\})$. Note that the receding horizon strategy defined a (sampled data) *feedback* law in contrast with the one shot use of an open loop optimal trajectory.

In defining (unconstrained) finite-horizon approximations to the infinite-horizon problem, the key design parameters are the terminal cost function V and the horizon length T (and, perhaps also, the increment δ). What choices will result in success?

It is well known (and easily demonstrated with linear examples), that simple truncation of the integral (i.e., $V \equiv 0$) may have disastrous effects if $T > 0$ is too small. Indeed, although the resulting value function

may be nicely behaved, the ‘‘optimal’’ receding-horizon closed-loop system can be unstable!

A more considered approach is to make good use of a suitable terminal cost V . Evidently, the best choice for the terminal cost is $V(x) = J_\infty^*(x)$ since then the optimal finite and infinite horizon costs are the same. Of course, if *the* optimal value function were available there would be no need to solve a trajectory optimization problem. What properties of the optimal value function should be retained in the terminal cost? To be effective, the terminal cost must account for the discarded tail by ensuring that the origin can be reached from the terminal state $x^u(T; x)$ in an efficient manner (as measured by q). One way to do this is to use an appropriate control Lyapunov function (CLF).

To this end, suppose that V is a proper C^2 function satisfying $V(0) = 0$

$$V(x) \geq c_v \|x\|^2, \quad x \in \mathbb{R}^n$$

and that is compatible with the incremental cost in the sense that

$$\min_u (\dot{V} + q)(x, u) \leq 0 \quad (3)$$

on a neighborhood of $x = 0$. Here, $\dot{V}(x, u) := DV(x) \cdot f(x, u)$. Condition (3) (together with the properties of f and q) guarantees the existence of a C^1 feedback law stabilizing the origin. Indeed, the feedback

$$u = k_V(x) := \bar{u}^*(x, DV(x)^T) \quad (4)$$

does the job. Note that V can be thought of as a Control Lyapunov Function which is also an upper bound on the cost-to-go. [The definition of the CLF requires that only $\min_u \dot{V}(x, u) \leq 0$]. The maximum principle ensures that $V = J_\infty^*$ also satisfies (3) according to (2).

Continuity and properness of V guarantee the existence of a continuous nondecreasing function $r \mapsto \bar{c}_v(r)$ such that $V(x) \leq \bar{c}_v(r) \|x\|^2$ for all $x \in \Omega_r$ so that $x \notin \Omega_{r_0}$ implies that $\|x\|^2 \geq r_0^2 / \bar{c}_v(r_0)$. Also, let $r_v > 0$ be the largest r such that (3) is satisfied for all $x \in \Omega_r$. The following result provides a basis for the use of finite horizon optimization in a receding horizon control strategy (cf. [9]).

Theorem 1: Suppose that $x \in \mathbb{R}^n$ and $T > 0$ are such that

$$x_T^*(T; x) \in \Omega_{r_v}. \quad (5)$$

Then, for each $\delta \in [0, T]$, the optimal cost from $x_T^*(\delta; x)$ satisfies

$$J_T^*(x_T^*(\delta; x)) \leq J_T^*(x) - \int_0^\delta q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau. \quad (6)$$

Note that $(x_T^*, u_T^*)(\cdot; x)$ can be any optimal trajectory for the problem with horizon T .

Proof: Let $(\tilde{x}(t), \tilde{u}(t))$, $t \in [0, 2T]$, be the trajectory obtained by concatenating $(x_T^*, u_T^*)(t; x)$, $t \in [0, T]$, and $(x^k, u^k)(t - T; x_T^*(T; x))$, $t \in [T, 2T]$. Here, $(x^k, u^k)(s; x_0)$ is the closed-loop trajectory starting from x_0 at time $s = 0$

$$x^k(s; x_0) = x_0 + \int_0^s f(x^k(\tau; x_0), k(x^k(\tau; x_0))) d\tau$$

where $u = k(x)$ is any feedback law such that $(\dot{V} + q)(x, k(x)) \leq 0$ for $x \in \Omega_{r_v}$, e.g., that defined by (4).

Consider now the cost of using $\tilde{u}(\cdot)$ for T seconds beginning at an initial state $x_T^*(\delta; x)$, $\delta \in [0, T]$. We have

$$\begin{aligned} J_T(x_T^*(\delta; x), \tilde{u}(\cdot)) &= \int_{\delta}^{T+\delta} q(\tilde{x}(\tau), \tilde{u}(\tau)) d\tau + V(\tilde{x}(T+\delta)) \\ &= J_T^*(x) - \int_0^{\delta} q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau - V(x_T^*(T; x)) \\ &\quad + \int_T^{T+\delta} q(\tilde{x}(\tau), \tilde{u}(\tau)) d\tau + V(\tilde{x}(T+\delta)) \\ &\leq J_T^*(x) - \int_0^{\delta} q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau \end{aligned}$$

where we have used the fact that $q(\tilde{x}(\tau), \tilde{u}(\tau)) \leq -\dot{V}(\tilde{x}(\tau), \tilde{u}(\tau))$ for all $\tau \in [T, 2T]$. The result follows since the optimal cost satisfies $J_T^*(x_T^*(\delta; x)) \leq J_T(x_T^*(\delta; x), \tilde{u}(\cdot))$. \square

At this point, one is tempted to conclude that our approach to approximating the infinite horizon problem using a CLF terminal cost has been successful. After all, (6) is an appropriate approximation to (1) for invariance purposes. In fact, Theorem 1 is sufficient to conclude the desired invariance and attractiveness properties in the case that V is a *global* CLF, for then that pesky “if” condition (5) will be trivially satisfied.

The situation when V is but a local CLF is much more delicate. Indeed, we must determine conditions under which (6) will hold under iteration of the receding horizon map, i.e., whether $x_T^*(T; x_T^*(\delta; x)) \in \Omega_{r_v}$ holds. One way to ensure success is to solve a *constrained* optimization that imposes such a condition, see, e.g., [14], [13]. We will show that such an approach is unnecessary.

We begin with a surprising lemma that helps us control the behavior of the terminal state of optimal trajectories.

Lemma 2: Suppose that $x \in \Omega_r$, $r \leq r_v$. Then, $x_T^*(T; x) \in \Omega_r$ for every $T \geq 0$.

Proof: As before, let $(x^k, u^k)(t; x)$, $t \geq 0$, be the trajectory (starting at x) obtained using a feedback control $u = k(x)$ satisfying $(\dot{V} + q)(x, k(x)) \leq 0$ on Ω_{r_v} . The optimal cost with horizon $T \geq 0$ satisfies

$$\begin{aligned} J_T^*(x) &\leq \int_0^T q(x^k(\tau; x), u^k(\tau; x)) d\tau + V(x^k(T; x)) \\ &\leq \int_0^T -\dot{V}(x^k(\tau; x), u^k(\tau; x)) d\tau + V(x^k(T; x)) \\ &= V(x) \leq r^2. \end{aligned}$$

Thus,

$$\begin{aligned} V(x_T^*(T; x)) &= J_T^*(x) - \int_0^T q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau \\ &\leq J_T^*(x) \leq V(x) \leq r^2. \end{aligned} \quad \square$$

Note that Lemma 2 *does not* say that $x_T^*(t; x) \in \Omega_{r_v}$ for all $t \in [0, T]$ when $x \in \Omega_{r_v}$. This is false in general as simple examples show. Indeed, one might say that methods that attempt to maintain the invariance of Ω_r , $r \leq r_v$, are inefficient. (Moreover, adding constraints of that sort also drive up the computation cost.)

A key motivation for using online optimization is to *enlarge* the operating region for a controller. We are now in a position to show that

the receding horizon controller does at least as good a job as the CLF controller, from the point of view of theoretical operating region predictions.

Proposition 3: For all $T \geq 0$, $x \in \Gamma_{r_v}^T$ implies that $x_T^*(T; x) \in \Omega_{r_v}$. Moreover, $\Omega_{r_v} \subset \Gamma_{r_v}^T$ for all $T \geq 0$.

Proof: Let $T \geq 0$ and $x \in \Gamma_{r_v}^T$ and note that

$$V(x_T^*(T; x)) \leq r_v^2 - \int_0^T q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau \leq r_v^2.$$

The second statement was proved in the proof of Lemma 2. \square

We now show that application of the receding horizon strategy results in the exponential convergence of the trajectory to the origin:

Theorem 4: Let $T > 0$ and consider the use of a receding horizon scheme $\mathcal{RH}(T, \{\delta_k\})$ with each $\delta_k \in (0, T]$ and $\sum_{j=0}^k \delta_j \rightarrow \infty$ as $k \rightarrow \infty$. Then, for each $x_0 \in \Gamma_{\bar{r}(T)}^T$, the resulting trajectory converges to the origin exponentially fast, where $\bar{r}(T)$ is the largest radius such that for each $x_0 \in \Gamma_{\bar{r}(T)}^T$, $x_T^*(T; x_0) \in \Omega_{r_v}$.

Proof: Given $T > 0$, set c_T and m_∞ such that $J_T^*(x) \leq c_T \|x\|^2 \forall x \in \Gamma_{\bar{r}(T)}^T$ and $J_\infty^*(x) \geq m_\infty \|x\|^2 \forall x \in \Gamma_{\bar{r}(T)}^T$. Let $(x_{RH}(t), u_{RH}(t))$, $t \geq 0$, be the receding horizon trajectory originating from an arbitrary $x_0 \in \Gamma_{\bar{r}(T)}^T$ and define

$$W(t; x_0, u_{RH}(\cdot)) := \int_t^\infty q(x_{RH}(\tau), u_{RH}(\tau)) d\tau.$$

[The control trajectory $u_{RH}(\cdot)$ is piecewise continuous since each optimal control trajectory $u_T^*(t; x)$, $t \geq 0$, is continuous as a function of time]. As shown in Proposition 9 below, $W(t; x_0, u_{RH}(\cdot)) \leq J_T^*(x_{RH}(t))$. Also, since $\mathcal{RH}(T, \{\delta_k\})$ is a suboptimal strategy, $J_\infty^*(x_{RH}(t)) \leq W(t; x_0, u_{RH}(\cdot))$. Now, since $q(x_{RH}(t), u_{RH}(t)) \rightarrow 0$ as $t \rightarrow \infty$, we have

$$\begin{aligned} \frac{\partial}{\partial t} W(t; x_0, u_{RH}(\cdot)) &= -q(x_{RH}(t), u_{RH}(t)) \\ &\leq -c_q \|x_{RH}(t)\|^2 \\ &\leq -\frac{c_q}{c_T} J_T^*(x_{RH}(t)) \\ &\leq -\frac{c_q}{c_T} W(t; x_0, u_{RH}(\cdot)) \end{aligned}$$

so that $W(t; x_0, u_{RH}(\cdot)) \leq e^{-(c_q/c_T)t} W(0; x_0, u_{RH}(\cdot))$. The result follows since

$$\begin{aligned} m_\infty \|x_{RH}(t)\|^2 &\leq J_\infty^*(x_{RH}(t)) \\ &\leq W(t; x_0, u_{RH}(\cdot)) \\ &\leq e^{-(c_q/c_T)t} W(0; x_0, u_{RH}(\cdot)) \\ &\leq e^{-(c_q/c_T)t} J_T^*(x_0) \\ &\leq c_T e^{-(c_q/c_T)t} \|x_0\|^2. \end{aligned} \quad \square$$

Note that the optimal control $u_T^*(\cdot; x)$ is uniquely defined in a neighborhood of the origin since $J_T^*(\cdot)$ is locally C^2 so that the locally defined *instantaneous* receding horizon control $u = k_T(x) := u_T^*(0; x)$ (i.e., $\delta = 0$) defines a feedback providing local exponential stability of the origin. Indeed, the resulting feedback law is identical to that

obtained by solving the associated Hamilton–Jacobi–Bellman PDE. When there are states x possessing multiple optimal trajectories (as occurs in the example below), it is no longer clear that an instantaneous receding horizon control can be successfully employed. A complete examination of the properties of the resulting differential inclusion is beyond the scope of this paper. From a practical point of view, the restriction to $\delta > 0$ is quite sufficient as some computation time is always required.

Theorem 4 says that for every fixed $T > 0$, the receding horizon scheme using a T -horizon optimization is effective. What it does not say, in particular, is that we may vary T and expect a stable process, i.e., stability is not guaranteed (by our results) when different horizon lengths are allowed at each receding horizon iteration. In contrast, we note that one does not need to use a fixed δ when implementing a receding horizon scheme since (6) implies that $x_T^*(\delta; x) \in \Gamma_{r_v}^T$ for all $\delta \in (0, T]$.

One expects that the region of effectiveness should grow as the optimization horizon T is increased, eventually covering all of Γ^∞ . This cannot be done without increasing r beyond r_v as the following result on inclusions shows.

Proposition 5: Let $r > 0$ be given and suppose that $T > 0$ is such that

$$x_T^*(T; x) \in \Omega_{r_v}$$

for all $x \in \Gamma_r^T$. Then

$$\Gamma_r^T \subset \Gamma_r^{T_1}$$

for all $T_1 \geq T$ so that, in particular, $\Gamma_r^T \subset \Gamma_r^\infty$.

Proof: Using (an extended version of) $\tilde{u}(\cdot)$ from the proof of Theorem 1, we see that

$$\begin{aligned} J_{T_1}(x, \tilde{u}(\cdot)) &= \int_0^T q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau \\ &\quad + \int_T^{T_1} q(\tilde{x}(\tau), \tilde{u}(\tau)) d\tau + V(\tilde{x}(T_1)) \\ &\leq \int_0^T q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau \\ &\quad + V(x_T^*(T; x)) = J_T^*(x). \end{aligned}$$

It follows that $J_{T_1}^*(x) \leq J_T^*(x)$ for all $x \in \Gamma_r^T$ (cf. [12]). \square

An important question is whether there exists a suitable horizon length for any desired radius r . The following result guarantees the existence of a suitable optimization horizon for a given (desired) radius r .

Proposition 6: For any $r > 0$ there is a $T_v = T_v(r)$ such that

$$x_T^*(T; x) \subset \Omega_{r_v}$$

for all $x \in \Gamma_r^\infty$ and all $T \geq T_v(r)$. In particular, $x_T^*(T; x) \subset \Omega_{r_v}$ for all $x \in \Gamma_r^T$.

Proof: First, note that $J_T^*(x)$ is bounded (hence, well defined) on Γ_r^∞ for all $T \geq 0$ since

$$\begin{aligned} J_T^*(x) &\leq \int_0^T q(x_\infty^*(\tau; x), u_\infty^*(\tau; x)) d\tau + V(x_\infty^*(T; x)) \\ &\leq J_\infty^*(x) + b_v(r) \end{aligned}$$

where $b_v(r) := \max_{x \in \Gamma_r^\infty} V(x)$. Next, we note that, regardless of the horizon length T , the trajectory $x_T^*(\cdot; x)$ must enter the set Ω_{r_v} within a bounded interval of time. Indeed, let $x \in \Gamma_r^\infty$ and $T > 0$ be arbitrary and suppose that $x_T^*(t; x) \notin \Omega_{r_v}$ on an interval $t \in [0, t_1]$. In this case, the optimal cost satisfies

$$\begin{aligned} J_T^*(x) &= \int_0^T q(x_T^*(\tau; x), u_T^*(\tau; x)) d\tau + V(x_T^*(T; x)) \\ &\geq \int_0^{t_1} c_q \|x_T^*(\tau; x)\|^2 d\tau \\ &\geq \frac{c_q}{c_v(r_v)} r_v^2 t_1. \end{aligned}$$

Combining the two inequalities, we see that, for $T > 0$ sufficiently large, $x_T^*(\cdot; x)$ must enter Ω_{r_v} with the first arrival time $t_1(x, T)$ satisfying

$$t_1(x, T) \leq \bar{t}_1(r) := \frac{\bar{c}_v(r_v)}{c_q} \frac{r^2 + b_v(r)}{r_v^2}.$$

In particular, we see that using $T_v = \bar{t}_1(r) + \epsilon$, $\epsilon > 0$, guarantees the existence of times $t_1(x) < T_v$, $x \in \Gamma_r^\infty$, such that $V(x_{T_v}^*(t_1(x); x)) \leq r_v^2$. The result $x_{T_v}^*(T_v; x) \in \Omega_{r_v}$ follows by Lemma 2 completing the proof. \square

The following corollary follows immediately from the above Proposition.

Corollary 7: Let $x_0 \in \Gamma^\infty$ be arbitrary. There exist $r, T < \infty$ such that

- 1) $x_0 \in \text{int } \Gamma_r^T$;
- 2) $x_T^*(T; x) \in \Omega_{r_v}$ for all $x \in \Gamma_r^T$.

This also shows that Γ^∞ is an open set.

We are now prepared to present the following theorem.

Theorem 8: Let Λ be a compact subset of Γ^∞ . There is a $T < \infty$ such that Λ is contained in the exponential region of attraction for the receding horizon strategy $\mathcal{RH}(T, \delta)$ for every $\delta \in (0, T]$.

Proof: For each $x \in \Lambda$, let $U(x) = \text{int } \Gamma_{r(x)}^T$ where $T(x)$ and $r(x)$ are given by Corollary 7. The collection $\{U(x)\}_{x \in \Lambda}$ is an open cover of Λ . By compactness, there is a finite subcover $\{U(x_i)\}_{i \leq N}$. Setting $T_i = T(x_i)$ and $r_i = r(x_i)$ we see that

$$\Lambda \subset \bigcup_{i \leq N} \Gamma_{r_i}^{T_i} \subset \bigcup_{i \leq N} \Gamma_{r_i}^{T_m} \subset \Gamma_{r_m}^{T_m}$$

where $T_m = \max_i T_i$, $r_m = \max_i r_i$ and the last two inclusions follow from Proposition 5.

Setting $T = T_m$ (and $r = r_m$) we see that $x_T^*(T; x) \in \Omega_{r_v}$ for all $x \in \Gamma_r^T \supset \Lambda$. The result follows since (6) ensures that $x_T^*(\delta; x) \in \Gamma_r^T$ for all $\delta \in (0, T]$. \square

Theorem 8 tells us that we may make the effective operating region of a receding horizon control strategy as large as we like (relative to the infinite-horizon operating region). Of great importance is the fact that this result is obtained using finite-horizon optimization without imposing any constraints on the terminal cost.

The following result provides a performance guarantee for our receding horizon control strategies.

Proposition 9: Suppose that $T, r > 0$ are such that $x_T^*(T; x) \in \Omega_{r_v}$ for all $x \in \Gamma_r^T$. Let $x_0 \in \Gamma_r^T$ and consider a

trajectory $(x_{\text{rh}}(t), u_{\text{rh}}(t))$, $t \geq 0$, resulting from the use of a receding horizon strategy $\mathcal{RH}(T, \{\delta_k\})$ (with $\delta_k > 0$, $\sum_{k=0}^l \delta_k \rightarrow \infty$ as $l \rightarrow \infty$). Then, the cost of this strategy satisfies

$$J_{\infty}(x_0, u_{\text{rh}}(\cdot)) \leq J_T^*(x_0).$$

Proof: The receding-horizon strategy defines a sequence of points $\{x_k\}_{k=0}^{\infty}$ according to $x_{k+1} = x_T^*(\delta_k, x_k)$ where $\sum_{k=0}^l \delta_k \rightarrow \infty$ as $l \rightarrow \infty$. Using a telescoping argument, the result follows easily from the fact (Proposition 5) that $J_T^*(x_k) \leq J_{T-\delta_k}^*(x_k)$. \square

The above proposition generalizes the fact that

$$\begin{aligned} (\dot{V} + q)(x(t), u(t)) &\leq 0, \quad t \geq 0 \\ \implies \int_0^{\infty} q(x(\tau), u(\tau)) d\tau &\leq V(x(0)) \end{aligned}$$

when V is positive definite [implying $x(t) \rightarrow 0$]. In both cases, we obtain an upper bound on the cost for a *family* of trajectories.

We also point out that the cost of using a receding horizon control strategy approaches the infinite horizon cost as the horizon T is increased since $J_{\infty}^*(x_0) \leq J_{\infty}(x_0, u_{\text{rh}}(\cdot)) \leq J_T^*(x_0)$ and $J_T^*(x_0) \rightarrow J_{\infty}^*(x_0)$ as $T \rightarrow \infty$.

IV. RELAXING THE REQUIREMENT FOR OPTIMALITY

In the previous section, we have detailed the theoretical properties of *ideal* receding horizon strategies wherein a *global* minimum is computed at each step. Only in very special cases (e.g., linear dynamics, strictly convex cost, etc.) can one expect reliable (approximate) computation of a global minimum. It is the purpose of this section to illustrate one of the many ways in which this requirement may be relaxed. See Sckaert, Mayne, and Rawlings [18] for results of this nature for discrete-time systems.

Receding-horizon techniques produce a sequence of (state and control) trajectories with ever decreasing cost. Stabilization or, more precisely, convergence of the cost may be obtained by ensuring that there is sufficient *improvement* at each step. Thus, we may replace the optimality test at *each* step by a test for improvement *between* steps. The following result provides a sufficient condition to ensure convergence of the state to the origin.

Proposition 10: Fix T , $\delta > 0$ and let x_i , $u_i(\cdot)$, $i \geq 0$, be such that $x_{i+1} = x^{u_i}(\delta; x_i)$ and

$$J_T(x_{i+1}, u_{i+1}(\cdot)) \leq J_{T-\delta}(x_{i+1}, u_i(\cdot + \delta)). \quad (7)$$

Then, $x_i \rightarrow 0$ as $i \rightarrow \infty$.

Proof: Note that the sequence of costs $c_i := J_T(x_i, u_i(\cdot))$ is monotone decreasing and bounded from below. It follows that the incremental cost

$$\delta c_i = c_i - c_{i+1} \geq \int_0^{\delta} q(x^{u_i}(\tau; x_i), u_i(\tau)) d\tau$$

must go to zero as $i \rightarrow \infty$. This implies that $x_i \rightarrow 0$ since there is a $\kappa > 0$ such that $\int_0^{\delta} q(x^u(\tau; x), u(\tau)) d\tau \geq \kappa \cdot \min\{1, \|x\|^2\}$ for every $u(\cdot)$. \square

How may we ensure, at each step, the existence of an improving control $u_{i+1}(\cdot)$?

Proposition 11: Suppose that x_0 and $u_0(\cdot)$ are such that $x^{u_0}(T; x_0) \in \Omega_{r_v}$. Then, there exists a sequence of controls $\{u_i(\cdot)\}_{i=1}^{\infty}$ such that $x^{u_i}(\delta; x_i) = x_{i+1} \rightarrow 0$ as $i \rightarrow \infty$.

Proof: Given x_i , $u_i(\cdot)$, choose $u_{i+1}(\cdot)$ such that $x^{u_{i+1}}(T; x_{i+1}) \in \Omega_{r_v}$ and the improvement property (7) is satisfied. One choice is the control obtained by using the remainder of $u_i(\cdot)$ followed by a CLF feedback control (as in the proof of Theorem 1). \square

One may (and many have) use constrained optimization to solve, at each step, a feasibility problem of the sort indicated. In that regard, the above result shows that the problem will remain feasible if it is initially thus. Also, since feasible controls may be obtained *for free*, we may use any means whatsoever (including unconstrained optimization) in our search for better controls, accepting only those that satisfy both terminal and improvement conditions.

V. EXAMPLE

For the purpose of illustration, we consider the problem of balancing an inverted pendulum on a cart. We discard the states associated with the cart to allow two dimensional visualization. (Please note that this is a highly unrealistic system as it allows equilibria where the cart is experiencing continuous acceleration—the system is *for visualization only*.) The pendulum is modeled as a thin rod of mass m and length $2l$ (the center of mass is at distance l from pivot) riding on a cart of mass M with applied (horizontal) force u . The dynamics of the pendulum are then given by (with θ measured from the vertical up position)

$$\ddot{\theta} = \frac{g/l \sin \theta - m_r \dot{\theta}^2/2 \sin 2\theta - m_r/ml \cos \theta u}{4/3 - m_r \cos^2 \theta}$$

where $m_r = m/(m + M)$ is the mass ratio and g is the acceleration of gravity. Specific values used are $m = 2$ kg, $M = 8$ kg, $l = 1/2$ m, and $g = 9.8$ m/s².

System performance is measured using the quadratic incremental cost $q(x, u) = 0.1x_1^2 + 0.05x_2^2 + 0.01u^2$ where as usual the state is $(x_1, x_2) = (\theta, \dot{\theta})$. To obtain an appropriate control Lyapunov function, we modeled the system locally as a polytopic linear differential inclusion (PLDI) [2]. This approach is quite satisfactory for this simple (planar) system over a large range of angles. Working over a range of plus or minus 60°, we obtained the quadratic CLF $V(x) = x^T P x$ with

$$P = \begin{bmatrix} 151.57 & 42.36 \\ 42.36 & 12.96 \end{bmatrix}.$$

Simple numerical calculations (in low dimensions!) show that $r_v \approx 6.34$, that is, $\min_u (\dot{V} + q)(x, u)$ is negative on solid P -ellipses Ω_r with a radius $r < 6.34$. An optimization technique that can be adapted to the problem of computing r_v in higher dimensions can be found in [11].

By Theorem 4, we know that, for $T \geq 0$, $\Gamma_{r_v}^T$ is an invariant subset of the region of attraction for the receding-horizon controller $\mathcal{RH}(T, \delta)$ with $\delta \in (0, T]$. Fig. 1(a) depicts the set $\Gamma_{r_v}^T$ for $T = 0.3$, $r_v = 6.34$ together with trajectories $x_T^*(\cdot; x)$ for x on the boundary. Also shown is the set Ω_{r_v} . The inclusion $\Omega_{r_v} \subset \Gamma_{r_v}^T$ (Proposition 5) is evident as is the fact that $x_T^*(T, x) \in \Omega_{r_v}$ for $x \in \Gamma_{r_v}^T$. Fig. 1(b) provides a comparison of receding horizon trajectories [for $\mathcal{RH}(0.3, 0.05)$] with those obtained using the CLF controller $u = k(x) = \arg \min_w (\dot{V} + q)(x, w)$. Note that $\Gamma_{r_v}^T$ is not invariant under the CLF flow. As expected, the receding horizon trajectories do remain inside $\Gamma_{r_v}^T$.




Fig. 1. (a) Sublevel set Γ_r^T for $T = 0.3$ and $r = r_v = 6.34$ together with Ω_{r_v} . Also depicted are the trajectories $x_T^*(\cdot; x)$ for x on the boundary of Γ_r^T . (b) Receding horizon $\mathcal{RH}(0.3, 0.05)$ and CLF controller (dashed) trajectories.




Fig. 2. State and control trajectories (RH—solid and CLF—dashed) from $x_0 = (-3.5, 5.9)$.

We also note that the CLF controller often requires significantly more control authority. For example, as shown in Fig. 2, the CLF controller demands almost 15 times as much authority when stabilizing from $x_0 = (-3.5, 5.9)$. (The pair originating at $x_0 = (-4.0, 7.5)$ has a ratio greater than 35) This is not too surprising since the CLF controller was designed for angular deviations of perhaps 60 degrees and qualified on the set Ω_{r_v} . The chosen x_0 is well outside of the guaranteed CLF performance region. In contrast, a small optimization horizon ($T = 0.3$ compared with a convergence time of >1.5) allows the receding horizon controller to exploit its knowledge of the nonlinear system dynamics in this region.

In this case, we see that significant performance improvements are obtained through the use of a relatively inexpensive receding horizon strategy.

The appropriate finite horizon optimization problems were solved numerically using RIOTS [17] as well as some local codes that are under development.

VI. CONCLUSION

In this note, we have developed a family of receding horizon control strategies that obtain excellent stability and performance properties through the use of a CLF as terminal cost. This approach is quite natural, providing a happy medium between the use of a CLF controller

and an ideal infinite horizon controller. Of practical significance, we have shown that this approach does not require the introduction of terminal constraints (for stability), thereby eliminating a key source of computational burden. In fact, it appears that these computations may be made fast enough to allow their use even in challenging areas such as flight control. Furthermore, it was shown that in order to maintain stability, it suffices to improve on the cost, thereby eliminating the need for precise calculation of (globally) optimal trajectories. An interesting further research direction is the extension of these techniques to the case of the trajectory tracking for nonlinear systems. Of course, the situation is much more complicated since the problem of finding useful trajectories of a nonlinear system is itself a rather difficult problem. A first step in that direction is the use of *trajectory morphing* techniques [7], [6].

REFERENCES

- [1] M. Bardi and I. Capuzzo-Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton–Jacobi–Bellman Equations*. Norwell, MA: Birkhauser, 1997.
- [2] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. Philadelphia, PA: SIAM, June 1994, vol. 15, Studies in Applied Mathematics.
- [3] G. Buttazzo, G. Mariano, and S. Hildebrandt, *Direct Methods in the Calculus of Variations*. Oxford, U.K.: Oxford Univ. Press, 1998.
- [4] L. Cesari, *Optimization—Theory and Applications: Problems with Ordinary Differential Equations*. New York: Springer-Verlag, 1983.
- [5] H. Chen and F. Allgöwer, “A quasiinfinite horizon nonlinear model predictive control scheme with guaranteed stability,” *Automatica*, vol. 34, pp. 1205–1217, 1998.
- [6] J. Hauser and D. G. Meyer, “The trajectory manifold of a nonlinear control system,” in *IEEE Conf. Decision Control*, 1998.
- [7] —, “Trajectory morphing for nonlinear systems,” in *Amer. Control Conf.*, 1998.
- [8] A. Jadbabaie, J. Yu, and J. Hauser, “Receding horizon control of the Caltech ducted fan: A control Lyapunov function approach,” in *IEEE Conf. Control Appl.*, 1999.
- [9] —, “Stabilizing receding horizon control of nonlinear systems: A control Lyapunov function approach,” in *Amer. Control Conf.*, 1999.
- [10] S. Keerthi and E. Gilbert, “Optimal infinite-horizon feedback laws for a general class of constrained discrete-time systems: Stability and moving-horizon approximations,” *J. Optim. Theory Appl.*, pp. 265–293, 1988.
- [11] M. C. Lai and J. Hauser, “Computing maximal stability region using a given Lyapunov function,” in *Amer. Control Conf.*, San Francisco, CA, 1993, pp. 1500–1502.
- [12] L. Magni and R. Sepulchre, “Stability margins of nonlinear receding horizon control via inverse optimality,” *Syst. Control Lett.*, vol. 32, pp. 241–245, 1997.
- [13] H. Michalska and D. Q. Mayne, “Robust receding horizon control of constrained nonlinear systems,” *IEEE Trans. Automat. Contr.*, vol. 38, pp. 1623–1633, Nov. 1993.

- [14] G. De Nicolao, L. Magni, and R. Scattolini, "Stabilizing receding-horizon control of nonlinear time-varying systems," *IEEE Trans. Automat. Contr.*, vol. 43, pp. 1030–1036, July 1998.
- [15] T. Parisini and R. Zoppoli, "A receding horizon regulator for nonlinear systems and a neural approximation," *Automatica*, vol. 31, pp. 1443–1451, 1995.
- [16] J. A. Primbs, V. Nevistić, and J. C. Doyle, "A receding horizon generalization of pointwise min-norm controllers," *IEEE Trans. Automat. Contr.*, vol. 45, pp. 898–909, June 2000.
- [17] A. Schwartz, "Theory and implementation of numerical methods based on Runge–Kutta integration for optimal control problems," Ph.D. dissertation, University of California, Berkeley, 1996.
- [18] P. Sokaert, D. Mayne, and J. Rawlings, "Suboptimal model predictive control (feasibility implies stability)," *IEEE Trans. Automat. Contr.*, vol. 44, pp. 648–654, Mar. 1999.
- [19] A. J. van der Schaft, "On a state space approach to nonlinear H_∞ control," *Syst. Control Lett.*, vol. 116, pp. 1–8, 1991.
- [20] —, *L_2 -Gain and Passivity Techniques in Nonlinear Control*. London, U.K.: Springer-Verlag, 1994, vol. 218, Lecture Notes in Control and Information Sciences.

Anticipatory Iterative Learning Control for Nonlinear Systems with Arbitrary Relative Degree

Mingxuan Sun and Danwei Wang

Abstract—In this note, the anticipatory iterative learning control is extended to a class of nonlinear continuous-time systems without restriction on relative degree. The learning algorithm calculates the required input action for the next operation cycle based on the pair of input action taken and its resultant variables. The tracking error convergence performance is examined under input saturation being taken into account. The learning algorithm is shown effective even if differentiation of any order from the tracking error is not used.

Index Terms—Convergence, learning control, nonlinear systems, relative degree.

I. INTRODUCTION

Recently, rigorous analyses of continuous-time iterative learning control (ILC) have been developed, see, for example, [2]–[10]. In particular, a fundamental characteristic of a class of learning control design methodologies is examined in [5], which clarifies the necessity of the use of error derivative for systems without direct transmission term. In [6], this characteristic is further clarified for nonlinear continuous-time systems where error derivatives, the highest order is equal to the relative degree of the systems, are used to update the control input. ILC using the highest-order error derivatives only is termed D-type ILC. Numerical calculations might be required to obtain error derivatives for the implementation. However, the signals obtained by numerical differentiation will be very noisy if the measurement is contaminated with noise. ILC without using differentiation is referred to as P-type ILC. Several technical analyzes of P-type ILC are presented for nonlinear continuous-time systems

with relative degree one, by imposing somewhat strict restriction on system dynamics, for example, the passivity property [11] and the boundedness of derivative of the input-output coupling matrix [12],[13]. Most recently, in [1], a fundamental concept is introduced in parallel to the two basic schemes: D-type and P-type ILCs. This design approach has the anticipatory characteristic of the D-type ILC and the simplicity like P-type ILC. Results have been developed again for nonlinear continuous-time systems with relative degree one and experimental results are obtained in robotic systems. This approach is also studied in the form of noncausal filtering [9]. In this note, the anticipatory learning algorithm [1] is applied to systems with arbitrary relative degree. A definition of extended relative degree is presented to explore a causal property of the systems under consideration. The tracking error convergence results are established.

II. PROBLEM FORMULATION

Consider the class of nonlinear continuous-time systems described by the state-space equations

$$\dot{x}(t) = f(x(t)) + B(x(t))u(t) \quad (1)$$

$$y(t) = g(x(t)) \quad (2)$$

where $x \in R^n$, $u \in R^r$ and $y \in R^m$ denote the state, control input and output of the system, respectively. The functions $f(\cdot) \in R^n$, $B(\cdot) = [b_1(\cdot), \dots, b_r(\cdot)] \in R^{n \times r}$ and $g(\cdot) = [g_1(\cdot), \dots, g_m(\cdot)]^T \in R^m$ are smooth in their domain of definition and are known of certain properties only. This system performs repetitive operations within a finite time interval $[0, T]$. For each fixed $x(0)$, S denotes a mapping from $(x(0), u(t), t \in [0, T])$ to $(x(t), t \in [0, T])$ and O a mapping from $(x(0), u(t), t \in [0, T])$ to $(y(t), t \in [0, T])$. In these notations, $x(\cdot) = S(x(0), u(\cdot))$ and $y(\cdot) = O(x(0), u(\cdot))$. The control problem to be solved is formulated as follows. Given a realizable trajectory $y_d(t), t \in [0, T]$ and a tolerance error bound $\varepsilon > 0$, find a control input $u(t), t \in [0, T]$, by applying an ILC technique, so that the error between the output trajectory $y(t)$ and the desired one $y_d(t)$ is within the tolerance error bound, i.e., $\|y_d(t) - y(t)\| < \varepsilon, t \in [0, T]$, where $\|\cdot\|$ is the vector norm defined as $\|a\| = \max_{1 \leq i \leq n} |a_i|$ for an n -dimensional vector $a = [a_1, \dots, a_n]^T$. Throughout the paper, for a matrix $A = \{a_{ij}\} \in R^{m \times n}$, the induced norm $\|A\| = \max_{1 \leq i \leq m} \sum_{j=1}^n |a_{ij}|$. To solve this problem, we use the ILC in the form of the following anticipatory updating law [1]:

$$v_{k+1}(t) = \begin{cases} u_k(t) + \Gamma_k(t)e_k(t + \sigma), & \text{if } t \in [0, T - \sigma] \\ v_k(T - \sigma), & \text{if } t \in (T - \sigma, T] \end{cases} \quad (3)$$

$$u_k(t) = \text{sat}(v_k(t)) \quad (4)$$

where

$\sigma > 0$	small number;
k	number of operation cycle;
$e_k(t) = y_d(t) - y_k(t)$	output or tracking error;
$\Gamma_k(t) \in R^{r \times m}$	learning gain matrix piecewise continuous and bounded.

This updating law is based on the causal relationship between the control input and the system output to be specified in the next section. The time shift ahead in the tracking error installs the anticipatory characteristic in the updating law, where actuator saturation is taken into

Manuscript received March 14, 2000; revised September 9, 2000. Recommended by Associate Editor C. Wen.

The authors are with the School of Electrical and Electronic Engineering, Nanyang Technological University, Nanyang, 639798, Singapore.

Publisher Item Identifier S 0018-9286(01)04121-6.