A Globally Stabilizing Receding Horizon Controller for Neutrally Stable Linear Systems with Input Constraints[1]

Ali Jadbabaie[†], Claudio De Persis[†], and Tae-Woong Yoon[†][2]
[†]Department of Electrical Engineering
Yale University, New Haven, CT 06520
ali.jadbabaie@yale.edu

Abstract

It is well known that exponentially unstable linear systems can not be globally stabilized in the presence of input constraints. In the case where the linear system is neutrally stable, one can achieve global asymptotic stability using a particular Control Lyapunov Function (CLF)-based controller. Using this particular CLF as terminal cost in a receding horizon scheme, we obtain a receding horizon controller which globally stabilizes such systems. Contrary to previous results, the horizon length is fixed, and can be chosen arbitrarily. The resulting controller also outperforms the CLF controller, since it provides a lower cost as measured by a quadratic performance index.

1 Introduction

Receding horizon control, also known as model predictive control, is perhaps the most popular method for controlling constrained systems. Over the past decade, the problem of stability has completely been characterized and solved, even for nonlinear systems. For an excellent review of this literature, we refer the reader to [6].

In the case of linear systems, even more is known. For example, we now know that the optimal MPC controller is piece-wise affine [1, 7], and the resulting value function is piece-wise quadratic. Even more surprisingly, we know now that under some mild non-degeneracy assumptions, the resulting value function is continuously differentiable [5].

A major characteristic of receding horizon controllers is that they are *local* in nature, except for very special cases. However, one can grow the region of attraction of the controller by increasing the horizon length, up to the infinite horizon region of attraction. More importantly, even in the nonlinear case, one can often do so without the use of artificial terminal stability constraints, therefore making the optimizations easier to solve [3].

It is well known and can be easily demonstrated via examples, that exponentially unstable linear systems subject to input constraints can not be stabilized globally, therefore no receding horizon scheme can achieve global stability of constrained linear systems when the system is exponentially unstable. However, when the system is marginally stable, i.e, it has simple eigenvalues on the imaginary axis, the situation is different. Recent results [2] have indicated that it is possible to construct a globally asymptotically stabilizing receding horizon scheme for such systems when the horizon is

allowed to be infinite. This is shown [2] to be equivalent to a fixed, finite-horizon receding horizon scheme over any compact sets of initial conditions. However, in order to make the result global, a variable-horizon receding horizon scheme should be implemented.

Since there are CLF-based results in the literature which prove global asymptotic stability of neutrally stable systems subject to input constraints [4], it is natural to think that one can develop a receding horizon extension of such controllers for *any fixed and finite* horizon length, thus improving on the cost of existing schemes on the one hand, and providing a practically feasible receding horizon strategy on the other.

The purpose of this paper is exactly the above point; to provide a CLF-based receding horizon scheme with an arbitrary finite horizon length [3], such that it globally stabilizes input constrained linear systems with simple eigenvalues on the imaginary axis. We use the CLF obtained in [4] as terminal cost in the receding horizon scheme providing global asymptotic stability. Note that since the CLF is not quadratically bounded from above, even when the incremental cost is quadratically bounded from below, *global exponential* stability is not possible;it is in fact a well-known result that not even neutrally stable linear systems with input constraints can be globally exponentially stabilized.

It will be shown that the above scheme works for any positive horizon length, and its performance (measured by a quadratic cost) is always better than that of the CLF controller.

This paper is organized as follows: In section 2 we setup our notation, and set up the problem. Section 3 provides a review of a CLF-based receding horizon strategy. We provide a particular CLF for marginally stable linear systems subject to input constraints in section 4 and provide a receding horizon extension by using the CLF as terminal cost in a receding

horizon framework. We present our concluding remarks in section 5.

2 Notation and problem setting

We consider a nonlinear system

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

where the vector field $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ is C^2 and possesses a linearly stabilizable 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 stabilizable. We also require the set $f(x, \mathbb{R}^m)$ 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

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

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 $q : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ that is C^2 and fully penalizes both state and control according to

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

for some $c_q > 0$ and $q(0,0) = 0$. We further require the function $u \mapsto q(x, u)$ to be convex for each $x \in \mathbb{R}^n$. These conditions imply that the quadratic approximation of q at the origin is positive definite, $D^2 q(0,0) \geq c_q I > 0$. The input $u(\cdot)$ is constrained to be inside the compact convex set $\mathcal{U}$ for all times. All of the above conditions are trivially satisfied for a linear system with a quadratic performance index.

The performance of the system is measured by the cost function

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

where V is a proper C^1 function satisfying $V(0) = 0$, which is also a control Lyapunov function that is compatible with the incremental cost in the sense that

$$\min_{u \in \mathcal{U}} (\dot{V} + q)(x, u) \leq 0 \qquad (1)$$

on a neighborhood of $x = 0$. Here $\dot{V}(x, u) := DV(x) \cdot f(x, u)$. Let $\Omega_r :=$ path connected component of $\{x \in \mathbb{R}^n : V(x) \leq r^2\}$. Also let r_v be the largest r such that (1) is satisfied for all $x \in \Omega_r$. Note that if (1) is satisfied globally, $\Omega_{r_v} = \mathbb{R}^n$.

3 CLF-based receding horizon control

In receding horizon control, a finite horizon open-loop optimization problem is solved. The first portion of the resulting optimal control trajectory is applied to the system, for a small fraction of the horizon length $\delta << T$. This procedure is then repeated, resulting in a sample data feedback law. In order for the resulting controller to be stabilizing certain ingredients of the problem have to be chosen properly. In this paper, we use the approach in [3], the terminal cost V is chosen in accordance to (1). Usually, the domain over which (1) is satisfied is some compact neighborhood of the origin. It is precisely the lack of a function for which (1) is true globally, that makes the receding horizon scheme appealing. However, as we will see in section 4, when the system is linear and neutrally stable and the cost quadratic, i.e., $f(x, u) := Ax + Bu$ with A having a simple spectrum on the imaginary axis and $q(x, u) = x^T Q x + u^T R u$ with $Q \geq 0$, $R > 0$, such global CLF does exist. Before discussing how such CLF can be found for linear systems, we mention the following theorem, which acts as a basis for the stability of the

CLF-based receding horizons scheme when a global CLF is available, and is a special case of the results in [6, 3].

Theorem 1 *Consider the nonlinear system $\dot{x} = f(x, u)$. Suppose $V(x)$ is a global CLF satisfying (1) for some incremental cost $q(x, u)$. Then the input constrained receding horizon scheme is globally asymptotically stabilizing for any horizon length $T \geq 0$, and any sampling time $\delta \in (0, T]$. Furthermore, the closed-loop system is semi-globally exponentially stable.*

In the next section, we use the results of [4] to come up with a global CLF satisfying (1) when $f(x, u) := Ax + Bu$ with A marginally stable.

4 A non-quadratic CLF for constrained, neutrally stable linear systems

Let $\bar{k} > 0$ be a real number such that the cube $C_{\bar{k}}$ is contained in the convex set $\mathcal{U}$ and consider the multi-valued saturation function defined as (cf. e.g. [8]) a function $\text{sat}(\cdot) : \mathbb{R}^s \to \mathbb{R}^m$ with $\text{sat}(v_1, \ldots, v_s) = (\text{sat}_1(v_1), \ldots, \text{sat}_s(v_s))$ and $\text{sat}_i(\cdot)$ a saturation function for all $i = 1, \ldots, s$, that is a locally Lipschitz function $\text{sat}(\cdot) : \mathbb{R} \to \mathbb{R}$ such that: (i) $\text{sat}(0) = 0$ and $r\text{sat}(r) > 0$ for all $r \neq 0$; (ii) $|\text{sat}(r)| \leq \bar{k}/m$ for all r and there exists $\underline{k}$ such that $\liminf_{|r| \to \infty} |\text{sat}(r)| \geq \underline{k}$; (iii) $\text{sat}(\cdot)$ is differentiable in a neighborhood of the origin and $\text{sat}'(0) = 1$. The stabilizability of the pair (A, B) and the neutral stability of A yields the existence of matrices $P = P^T > 0$ and $M = M^T > 0$ such that $A^T P + PA \leq 0$ and $(A - BB^T P)^T M + M(A - BB^T P) = -2I$. Then the results of [4] show that for any positive number $\varepsilon > 0$ there exist positive numbers $\lambda_1, \lambda_2 > 0$ such that the function

$$V(x) = \lambda_1 x^T M x + \lambda_2 (x^T P x)^{3/2} \qquad (2)$$

satisfies the inequality

$$\dot{V}|_{u=\text{sat}(-B^T Px)} = \frac{\partial V}{\partial x}(Ax + B\text{sat}(-B^T Px)) \leq -\varepsilon |x|^2 .$$

Now note that there exists a positive number $\lambda > 0$ such that

$$(x^T Q x + u^T R u)|_{u=\text{sat}(-B^T P x)} = \lambda |x|^2 \; ,$$

since from the definition of saturation function $|\text{sat}(s)| \leq H|s|$, for some $H > 0$. Correspondingly to such λ choose λ_1, λ_2 in (2) in such a way that

$$\dot{V}|_{u=\text{sat}(-B^T P x)} \leq -\lambda |x|^2 \; .$$

Noting that by the definition of $\text{sat}(\cdot)$, $u = \text{sat}(-B^T P x) \in \mathcal{U}$ for all $x \in \mathbb{R}^n$, we immediately have

$$\min_{u \in \mathcal{U}} (\dot{V} + q)(x, u) \leq (\dot{V} + q)(x, u)|_{u=\text{sat}(-B^T P x)} \leq 0 \; ,$$

that is $V(\cdot)$ is a CLF satisfying (1). The following result is then an immediate consequence of the arguments above and Theorem 1:

Theorem 2 *Main result: Consider the controllable linear system $\dot{x} = Ax + Bu$, with A neutrally stable. Then the receding horizon scheme with the non-quadratic CLF (2) chosen as terminal cost is globally asymptotically (and semi-globally exponentially) stabilizing.*

Note that the above result is for any fixed horizon, and therefore stronger than those of [2] since the horizon length can be chosen arbitrarily. Also, when the horizon length is chosen to be zero, we recover the CLF controller of [4]. Furthermore, the performance of the receding horizon controller as measured by the quadratic performance index is always better than the mere CLF controller [3]. A possible disadvantage of this scheme is that the problem is no longer solvable by a quadratic program, since the terminal cost is non-quadratic. However, it should be noted that the terminal cost is still convex. There are a variety of sequential quadratic programing algorithms which one can use to solve the optimization efficiently.

5 Concluding remarks

We proposed a CLF-based receding horizon scheme which globally asymptotically stabilizes input-constrained linear systems that are open-loop neutrally stable. A non-quadratic CLF which is an upper bound on the cost-to-go is utilized as terminal cost in the receding horizon scheme. This results is an improvement over the existing results which require a variable horizon length to achieve global asymptotic stability.

References

[1] A. Bemporad, M. Morari, V. Dua, and N. Pistikopolos. The explicit linear quadratic regulator for constrained systems. Technical Report 99-16, ETH Zurich, 1999.

[2] D. Chmielewski and V. Manousiouthakis. On constrained infinite-time linear quadratic optimal control. *Systems and Control Letters*, 29:121–129, 1996.

[3] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 46, May 2001.

[4] W. Liu, Y. Chitour, and E. Sontag. On finite gain stabilizability of linear systems subject to input saturation. *SIAM Journal on Control and Optimization*, 34:1190–1219, 1994.

[5] D. Q. Mayne. Control of constrained dynamic systems. *European Jounral of Control, to appear*, 2001.

[6] D. Q. Mayne, J. B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[7] M. Seron, J. A. De Dona, and G. C. Goodwin. Global analytic model predictive control with input constraints. In *Proceedings of IEEE Conference on Decision and Control, Sydney, Australia*, 2000.

[8] A. R. Teel. A nonlinear small gain theorem for the analysis of control systems with

saturation. *IEEE Transactions on Automatic Control*, 41:1256–1270, 1996.