Relaxing the optimality condition in receding horizon control [1]

Ali Jadbabaie[2] and John Hauser[3]
{alij,hauser}@cds.caltech.edu

Abstract

Receding horizon control is based on iteratively solving an open-loop finite horizon optimization problem. Despite its success in a variety of industrial applications, theoretical issues such as stability were not completely addressed until recently. It was shown in [5] that by utilizing a suitable Control Lyapunov Function (CLF) as terminal cost, the stability of the receding horizon scheme can be guaranteed and the region of attraction of the receding horizon controller can be estimated. The key point in this approach, which made it different from others, was removal of additional stability constraints, hence making the optimizations much easier to solve. A requirement implied in the previous results was being able to solve the optimizations globally . In this paper, that assumption is removed and it is shown that the optimality can be replaced by an *improvement* property. Specifically, instead of requiring the trajectories to be optimal, it is required that a certain amount of decrease in the cost is obtained at each receding horizon iteration. It is further shown that there always exist a sequence of controls which guarantee the necessary decrease in the cost. A numerical example using the inverted pendulum is presented to illustrate this point.

Keywords: receding horizon control, nonlinear control design, model predictive control, optimal control, Control Lyapunov Functions, improvement property.

[1] Research supported in part by AFOSR and DARPA.

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

[3] Electrical and Computer Engineering, University of Colorado, Boulder, CO 80309-0425.

1 Introduction

The advent of faster and cheaper computing power as well as efficient numerical algorithms for solving large scale optimization problems, has made receding horizon control quite popular. Despite practical success in applying receding horizon control methodology to open loop stable plants, its application in stability critical situations has not been as successful. Stability of the receding horizon scheme has been tackled by a variety of researchers in the field.

These approaches used additional endpoint equality constraints [6] or endpoint inequality constraints [7, 8, 9, 10, 11], in order to guarantee closed-loop stability. Others, such as [3, 14, 12], showed (in the context of constrained linear systems) that for long enough horizon lengths, the receding horizon scheme can be stabilizing without any terminal constraint.

An alternative approach was developed by the authors in [4, 5]. This approach obtained stability guarantees through the use of an *a priori* obtained CLF as a terminal cost rather than by imposing state inequality (or equality) constraints. By utilizing a suitable control Lyapunov function which can be obtained off-line, one can guarantee the stability of the receding horizon scheme in a more efficient manner. In this setting, it was shown that the stability constraints are automatically satisfied, hence can be eliminated from the optimization. This resulted in a dramatic speedup in the computations.

Furthermore, it was shown in [5] that the region of attraction of the unconstrained receding horizon control law is always larger than that of the CLF and it can be grown further to contain any compact subset of the infinite horizon region of attraction by a suitable choice of the horizon length.

The analysis presented in [5] made explicit use of *optimality*, that is, the global minimum was required at each step. In this paper, similar to the results of Chen and Allgöwer [2], and Scokaert *et al.* [13], we propose an algorithm that requires only improvement of the cost in each iteration, rather than obtaining the true minimum. Hence, stability is guaranteed by application of a *sub optimal* receding horizon strategy. Specifically, we provide a test that can be applied to show that the computed solutions are sufficiently good to ensure stability. This paper is organized as follows:
Section 2 describes the problem setting, in which some infinite horizon properties are reviewed. In section 3 some of the results on stabilizing receding horizon control strategy with CLF as terminal cost are discussed. Our main result, i.e., relaxing the requirement for optimality is described in section 4. An example is presented in section 5 to illustrate the fact that locally optimal trajectories can be stabilizing. Finally, our conclusions are presented in section 6.

2 Problem Setting

The nonlinear system under consideration is

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

where the vector field $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ is C^2 and possesses an exponentially 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. For simplicity, we will strengthen this and require (A, B) to be controllable. 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 $t \geq 0$. Smoothness of f guarantees existence and uniqueness for x^u on small time intervals. (Note that global existence in time can only be guaranteed by appropriate choice of x and $u(\cdot)$.)

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 $c_q > 0$. This implies that the quadratic approximation of q at the origin is positive definite, $D^2 q(0,0) \geq c_q I > 0$. This could be weakened, e.g., to some observability/detectability condition.

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.

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 proposition 3.1 of [1].

It is easy to see that J_∞^* is proper on its domain so that the sub-level 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

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

containing 0, and

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

containing 0.

3 Unconstrained receding horizon control

In receding horizon control, an open-loop finite horizon optimization problem is solved to generate a control trajectory. Once the trajectory is obtained, a small portion of it is applied to the system. Successive repetition of this procedure results in a feedback law which can be thought of as an approximation to the solution of the infinite horizon problem. The main challenge in the receding horizon scheme is to somehow take the ignored portion of the cost function into account. As it was shown in [5], this can be done by first obtaining a suitable CLF which is compatible with the incremental cost and then use that as the terminal cost.

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

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

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

$$\min_u \ (\dot{V} + q)(x,u) \leq 0 \tag{1}$$

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

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 (1).

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 (1) 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. [5]).

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

$$x_T^*(T;x) \in \Omega_{r_v} \ . \tag{2}$$

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

$$J_T^*(x_T^*(\delta;x)) \leq J_T^*(x) - \int_0^\delta q(x_T^*(\tau;x), u_T^*(\tau;x)) \, d\tau \ . \tag{3}$$

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

Proof: See [5] for a complete proof. ■

The above theorem suggests that once the terminal condition (2) is satisfied, stability can be ensured by using the minimum cost J_T^* as a CLF candidate. Equation (3) ensures proper decrease, therefore theorem 1 is sufficient to conclude the desired invariance and attractiveness properties in the case that V is a *global* CLF.

It turns out that in the case where V is a local CLF [5], stability can be still guaranteed *without* imposing constraints. The resulting controller has region of attraction at least as big as that of the CLF, and can be grown to contain any subset of the region of attraction of the infinite horizon control law by increasing the horizon length, yet keeping the problem unconstrained.

The following theorem from [5] summarizes the main results.

Theorem 2 *Let $T \geq 0$. The feedback law*

$$u = k_T(x) := u_T^*(0;x)$$

exponentially stabilizes the origin for $\dot{x} = f(x,u)$. Moreover, the region of attraction contains $\Gamma^T_{r_v}$. More generally, $\Gamma^T_{r_v}$ is contained in the exponential region of attraction for the receding horizon scheme for every $\delta \in [0,T]$.

The above results made explicit use of optimality of the trajectories. From a practical point of view, even when the optimization problem is convex, one merely computes an approximation to the optimum. In the general non-convex case, which is certainly the case in receding horizon optimizations, there is no general way to insure that the computed solution is close to the global optimum. In the next section, we completely address this issue.

4 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 Scokaert, Mayne, and Rawlings [13] 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. Note that it is not required that the resulting trajectories be optimal, as long as there is enough decrease in the cost at each step.

Proposition 3 *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*

$$J_T(x_{i+1}, u_{i+1}(\cdot)) \leq J_{T-\delta}(x_{i+1}, u_i(\cdot + \delta)) . \quad (4)$$

Then $x_i \to 0$ as $i \to \infty$.

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 $i \to \infty$.

This implies that $x_i \to 0$ 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 $u(\cdot)$. ■

The above proposition gives conditions on the controls, under which stability of the closed-loop will be guaranteed. How may we ensure, at each step, the existence of an improving control $u_{i+1}(\cdot)$? The following proposition guarantees the existence of such controls.

Proposition 4 *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 $\{u_i(\cdot)\}_1^\infty$ such that $x^{u_i}(\delta; x_i) = x_{i+1} \to 0$ as $i \to \infty$.*

Proof: Given x_i, $u_i(\cdot)$, choose $u_{i+1}(\cdot)$ such that

$$x^{u_{i+1}}(T; x_{i+1}) \in \Omega_{r_v} \quad (5)$$

and the improvement property (4) is satisfied. Note that if x_i is inside Ω_{r_v}, condition (5) is already satisfied. One obvious choice is the control obtained by using the remainder of $u_i(\cdot)$ for $T - \delta$ seconds followed by a feedback control obtained from the CLF for δ seconds. Such a control will indeed guarantee the improvement property. (See proof of Theorem 1 in [5]). ■

One may use (and many have used) 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. Furthermore, note that if the starting point is inside Ω_{r_v} the trajectories will end up there at the end of the horizon length T, therefore they will remain feasible automatically. This point will be illustrated further in the following section.

5 Example

In this section, we use the familiar inverted pendulum on a cart example. The cart dynamics has been ignored to allow two dimensional visualization. Application of the receding horizon strategies to this example was demonstrated in [5]. 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 $g = 9.8\ \text{m/s}^2$.

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). 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 degrees, we obtained a 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 $\mathbb{R}^2$) 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 $r < 6.34$.

For small values of T, the set $\Gamma_{r_v}^T$ looks very much like the ellipse Ω_{r_v} which is reasonably well lined up with the stable manifold of the pendulum. As the value of T is increased, the ends of $\Gamma_{r_v}^T$ begin to open up, eventually wrapping back around toward the inverted equilibrium, indicating that it can be efficient (from a cost standpoint) to allow the pendulum to swing down before bringing it back up to the vertical position. Figure 1 depicts the nature of this wrap-around for $T = 2.0$ and $r = r_v = 6.34$. The set $\Gamma_{r_v}^T$ is shown without an overlap by plotting half of the set boundary which, together with trajectories starting on the boundary, provide an unwrapped view of the set. Figure 2 provides a close up view of the overlapping set $\Gamma_{r_v}^T$ together with the set Ω_{r_v}. At each point in the overlap region, there are (at least) two local minima. Strict use of the *global* optimum in a receding horizon strategy would indicate a preference for letting the pendulum fall in many situations where the pendulum can be brought back to the vertical quickly and for a reasonable, though suboptimal, cost.

Consider, for example, the use of a receding horizon strategy with $\delta = 0.1$ (and $T = 2.0$) starting at the initial condition $x(0) = x_0 = (-1.803, 8.413)$ with optimal cost $J_T^*(x_0) = r_v^2 = 40.1956$. The situation is depicted in Figure 3. After $\delta = 0.1$ seconds, we arrive at $x(\delta) = (-1.038, 6.887)$ and find two local minima with values 36.973 and 37.323, offering two potential strategies. As both costs are less than r_v^2, it is clear that either course will result in convergence to the inverted equilibrium. The resulting trajectories are shown in Figure 3. The evolution of the costs is shown in Figure 4 verifying its decreasing nature as well as the possibility of discrete jumps, indicating strict inequality in (4).

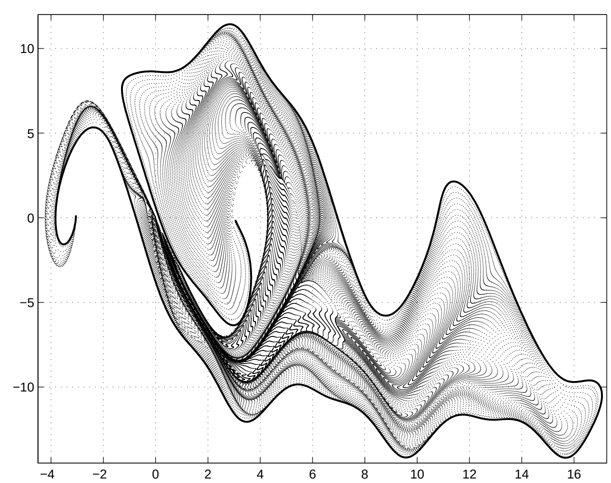

Figure 1: The sublevel set Γ_r^T for $T = 2.0$ and $r = r_v = 6.34$. Half of the boundary (together with trajectories) is shown in an *unwrapped* fashion to aid in understanding the overlapping nature of the set.

6 Conclusion

The purpose of this paper was to extend the results of [5] to the case were the optimizations are not solved exactly. Instead of requiring the receding horizon trajectories to be optimal in each iteration, a certain decrease in the value of the cost was required. Furthermore, it was shown that there always exists a controller which provides the proper decrease. The control trajectory consists of two parts, the first part is the tail of the trajectory obtained from the previous iteration and the second is a feedback obtained from the CLF. A numerical example using the inverted pendulum compared the use of locally optimal and globally optimal trajectories. Simulations indicate that there are regions in which more than one locally optimal trajectory exist and both of them are stabilizing. These results can have interesting implications in practical cases where there might not be enough time to even compute the locally optimal trajectories.

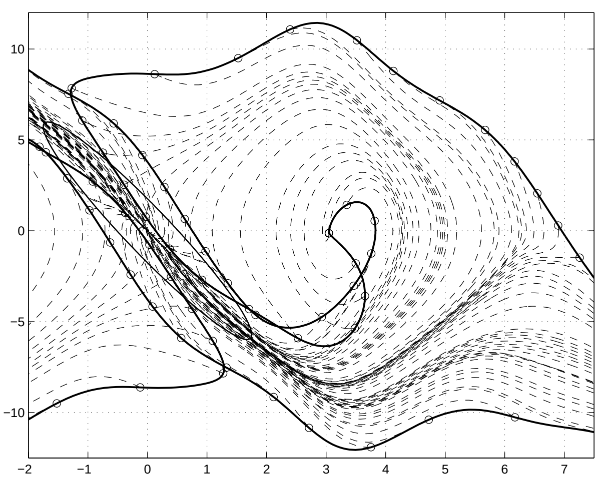

Figure 2: A closeup of the sublevel set Γ_r^T for $T = 2.0$ and $r = r_v = 6.34$ together with Ω_{r_v}. Also depicted are several locally optimal trajectories beginning on the boundary.

References

[1] M. Bardi and I. Capuzzo-Dolcetta. *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Birkhauser, Boston, 1997.

[2] H. Chen and F. Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Automatica*, 34:1205–1217, 1998.

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

[4] A. Jadbabaie, J. Yu, and J. Hauser. Receding horizon control of the Caltech ducted fan: A control Lyapunov function approach. In *IEEE Conference on Control Applications*, 1999.

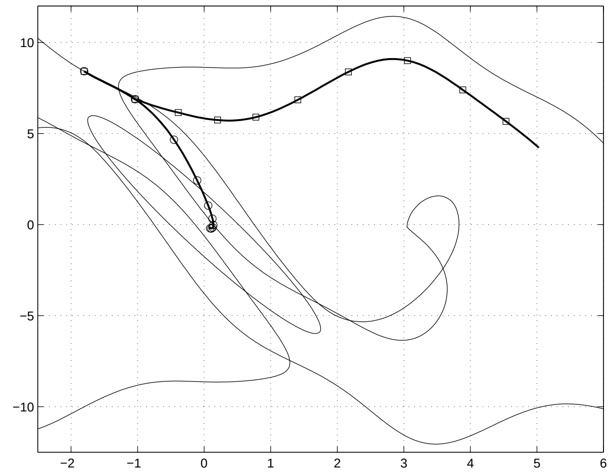

Figure 3: Receding horizon trajectories using $\delta = 0.1$ beginning at $x(0) = (-1.803, 8.413)$. At $x(\delta) = (-1.038, 6.887)$, local minima with costs of 36.973 (square) and 37.323 (circle) are found providing two different strategies.

[5] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control of nonlinear systems. In *IEEE Conference on Decision and Control*, 1999.

[6] S. Keerthi and E. Gilbert. Optimal infinite-horizon feedback laws for a general class of constrained discrete-time systems: Stability and moving-horizon approximations. *Journal of Optimization Theory and Applications*, pages 265–293, 1988.

[7] L. Magni and R. Sepulchre. Stability margins of nonlinear receding horizon control via inverse optimality. *Systems and Control Letters*, 32:241–245, 1997.

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

[9] G. De Nicolao, L. Magni, and R. Scattolini. Stabilizing receding-horizon control of nonlinear time-varying systems. *IEEE Transactions on Automatic Control*, 43(7):1030–1036, 1998.

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

[11] J. 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, June 2000.

[12] J.A. Primbs and V. Nevistic. Feasibility and stability of constrained finite receding horizon control. *Automatica*, 36(7):965–971, 2000.

[13] P. Scokaert, D. Mayne, and J. Rawlings. Suboptimal model predictive control (feasibility implies stability). *IEEE Transactions on Automatic Control*, 44:648–654, March 1999.

[14] P. Scokaert and J. Rawlings. Constrained linear quadratic regulation. *IEEE Transactions on Automatic Control*, 43:1163–1169, August 1999.

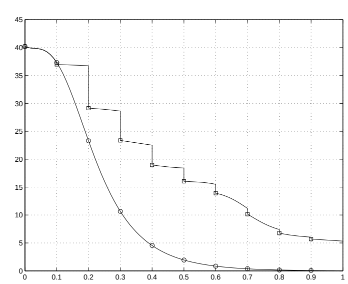

Figure 4: Evolution of the cost for the two strategies shown in Figure 3.