A Recentred Barrier for Constrained Receding Horizon Control

Adrian Wills*†and Will Heath‡

September 19, 2001

Abstract

Constrained receding horizon controllers are often designed to regulate the system state about some desired set-point subject to input constraints. This paper presents a class of receding horizon controllers which force the inputs to lie inside the constraint set by including a so-called 'recentred barrier function'. The significance of such a controller is that hard constraints are replaced with penalty type soft constraints. This results in a control law that 'backs off' near constraint boundaries. The degree to which this backing-off occurs, is directly related to the parameter value that characterises the class of controllers. For each parameter value there is a corresponding unconstrained minimisation problem. The associated control law is obtained in the standard manner by solving this problem at each time interval and applying the first control move to the system. This method is applicable to general convex objective functions with convex inequality constraints. We illustrate this idea by way of an example application to linear discrete-time plant models with linear and convex quadratic static input constraints.

1 Introduction

The cost function associated with receding horizon control is a means for determining that one control path is preferrable to another [XX]. The cost function is inherently linked to closed-loop dynamic behaviour and in certain cases this connection can be made explicit [XX]. In the case where constraints are deemed important to the closed-loop behaviour, then it is customary to represent them as a closed convex subset of a finite dimensional real vector space (Mayne *et al.*, 2000). The receding horizon control problem may then be expressed as: at each time interval with given system state, optimise the cost function over future control moves, subject to constraints, then apply the first control move to the system (Mayne *et al.*, 2000).

In this paper we propose instead to include the constraints as part of the cost function. In particular, use of a barrier function (which is quite standard for interior-point methods (Fiacco and McCormick, 1968)) ensures that the solution lies on the interior of the constraint set. This approach penalises control action near the constraint boundary more than control action away from the constraint boundary. Therefore, the resulting control law becomes 'cautious' when approaching a boundary point. Furthermore, the degree of cautiousness is directly related to the parameter value which governs the intensity of the barrier term. The class of controllers presented in this paper is characterised by the same parameter value, henceforth referred to by the positive scalar μ. Moreover, the more traditional receding horizon control problem is captured in the limit as $\mu \to 0$.

Given an arbitrary steady-state set-point on the interior of the constraint set, we develop a so-called *recentred barrier function*. If the closed-loop system is stable for the chosen parameter value, say μ_p, then this barrier function ensures convergence to the set-point in steady-state. This property is not guaranteed with a more general barrier - *even with integral action*. This recentring essentially shapes the

*Department of Electrical and Computer Engineering, University of Newcastle.
†Email: onyx@ecemail.newcastle.edu.au
‡Centre for Integrated Dynamics And Control (CIDAC), University of Newcastle.

barrier function such that the minimum occurs at the desired steady-state set-point.

To illustrate these ideas, we investigate the application of the proposed controller for the case of a linear plant model with linear and convex quadratic input constraints. We construct a *terminal constraint set* which is broadly based on (Chen and Allgower, 1998), and demonstrate global stability for this case.

The paper is organised as follows: section 2 introduces two key ideas, namely the use of barrier functions for constrained receding horizon control and the recentred barrier function. Section 3 provides a class of controllers which are based on the recentred barrier function. Section 4 provides a detailed example application of this approach to the case of linear discrete-time plant models with linear and convex quadratic constraints. Section 5 presents a simple example simulation to illustrate the key ideas of this method. Section 6 concludes the paper.

2 Barrier Function Preliminaries

2.1 Notation

This section introduces some preliminary notation and terminology relevant to barrier function methods. Let G be a convex constraint set for a finite dimensional real vector space $Z = \mathbb{R}^n$ defined as,

$$G := \{z \in Z : \; f_i(z) \le 0 \text{ for } i = 1, \ldots, M\}. \tag{1}$$

where each $f_i(z)$ is a continuous convex function of z. Let G^0 denote the interior of G defined as $G^0 := \{z \in Z : \; f_i(z) < 0 \text{ for } i = 1, \ldots, M\}$. It is assumed throughout this paper that $G^0 \neq \emptyset$ and G is bounded.

Let $F(z)$ denote a barrier function for G^0 (see (Nesterov and Nemirovskii, 1994)). Usually, $F(z)$ is chosen to be the logarithmic barrier:

$$F(z) = \begin{cases} -\sum_{i=1}^{M} \ln(-f_i(z)) & \text{if } z \in G^0 \\ \infty & \text{otherwise.} \end{cases} \tag{2}$$

See figure 2(a) for a simple illustration of this barrier function. We define the class of barrier generated minimisation problems as follows,

$$(\mathcal{CP}_\mu): \quad \underset{z \in Z}{\operatorname{argmin}} \left\{ f_0(z) + \mu F(z) \right\}, \tag{3}$$

where $f_0(z)$ is also a continuous convex function. For ease of exposition we consider the case where $f_0(z)$ is strictly convex. Let $z(\mu)$ be the solution to $(\mathcal{CP}_\mu)$; then the set of points $\{z(\mu) \in Z : \; \mu \in (0, \infty)\}$ is called the central path for the following optimisation problem,

$$(\mathcal{C}): \quad \underset{z \in Z}{\operatorname{argmin}} \; f_0(z) \;\; \text{s.t. } z \in G \tag{4}$$

In the sequel, we may refer to $(\mathcal{C})$ as the 'limiting case' optimisation problem. Indeed, in the limit as $\mu \to 0$, the solution of $(\mathcal{CP}_\mu)$ approaches the solution of $(\mathcal{C})$ (see e.g. (Ye, 1997), (Wright, 1997), (Nesterov and Nemirovskii, 1994), (Fiacco and McCormick, 1968)).

Since G is closed, convex and compact with non-empty interior, then a necessary and sufficient condition for optimality of (3) is $\nabla f_0(z(\mu)) = -\mu \nabla F(z(\mu))$. Let $z(\infty)$ be defined as $z(\infty) := \lim_{\mu \to \infty} z(\mu)$, then $z(\infty)$ is the unique point that minimises $F(z)$ (often termed the analytic centre of the constraint set G).

2.2 Recentred Barrier Function

In this section we motivate two crucial concepts. The first is to use $(\mathcal{CP}_\mu)$ with a fixed weighting parameter, say μ_p, to define an unconstrained receding horizon optimisation problem. The second is the recentred barrier function which ensures that if the system converges, then it converges to the set-point.

Conventionally when using interior-point methods for MPC, an ϵ-solution to $(\mathcal{C})$ is found at each time step (see e.g. (Wright, 1997) and (Rao *et al.*, 1998)). This can be *interpreted* as solving $(\mathcal{CP}_\mu)$ for a very small value of μ. Recall from section 2.1 that the central path is a uniquely defined set of points which are parameterised by a positive scalar $\mu \in (0, \infty)$. We propose to fix this parameter, to say μ_p, and solve the corresponding unconstrained minimisation problem $(\mathcal{CP}_{\mu_p})$. This approach differs from standard interior-point solutions in that the choice of μ_p may be quite large by comparison to ϵ. The receding horizon control law is then defined in the usual manner by selecting the first control move from this solution. In this sense, $(\mathcal{CP}_\mu)$ defines a class of receding horizon optimisation problems for $\mu \in (0, \infty)$, see figure 1. In terms of receding horizon control, this has the effect of weighting control moves more heavily near the constraint boundaries. Furthermore, we now have an unconstrained convex optimisation problem, which for μ_p large enough, can be solved via simple Newton iterations. A more versatile algorithm for solving $(\mathcal{CP}_{\mu_p})$ is presented in (Wills and Heath, 2001). Moreover, the limiting case $(\mathcal{C})$, is captured in a continuous manner as $\mu_p \to 0$.

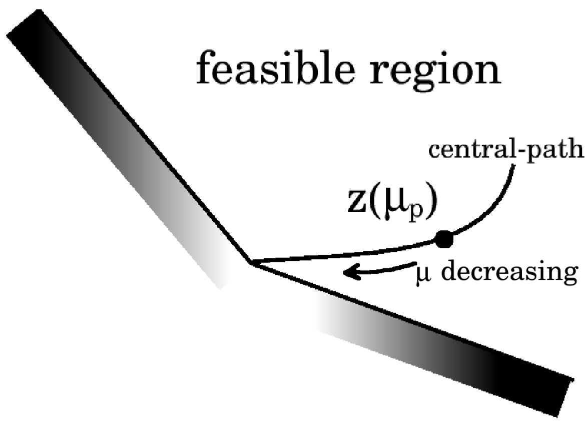

Figure 1: Central Path: We propose to stop at the point corresponding to the parameter value μ_p

A consequence of receding horizon control strategies is that the objective function $f_0(z)$ may change at each time step - yet the overall objective remains the same, i.e. that the state should converge to the set-point (providing the set-point is feasible). Suppose that the desired steady-state set-point is $z_d \in G^0$ (a strictly feasible point), and that z_d is the minimum of the cost function $f_0(z)$ in steady-state. Suppose also that we have a fixed weighting parameter $\mu_p \gg 0$. Ideally, if the system does converge, then it should converge to z_d, i.e. the minimum of $(\mathcal{CP}_{\mu_p})$ should be $z(\mu_p) = z_d$. However, since $f_0(z)$ and $F(z)$ are strictly convex functions, then $z(\mu_p)$ will lie somewhere between $z(\infty)$ and z_d (recall that $z(\infty)$ is the minimum of $F(z)$). Therefore, the points $z(\mu_p)$ and z_d will not coincide unless $z(\infty) = z_d$, i.e. unless the minimum of $F(z)$ occurs at z_d. This means that if the system does converge, then using the standard logarithmic barrier function with a fixed weighting parameter μ_p, does not guarantee convergence to the set-point z_d - *even with integral action.* One way to overcome this issue is to construct a new barrier function which satisfies this requirement. One possibility is as follows:

Definition 1 *Let $F_{z_d}(z)$ denote the gradient recentred barrier function for G^0, at the point $z_d \in \text{int}G$. We define $F_{z_d}(z)$ as,*

$$F_{z_d}(z) = F(z) - \langle \nabla F(z_d), z \rangle. \tag{5}$$

For the logarithmic barrier this may be explicitly given by

$$F_{z_d}(z) = F(z) + b_{z_d}^T z \tag{6}$$

where b_{z_d} is given by

$$b_{z_d} = \sum_{i=1}^{M} \frac{1}{f_i(z_d)} \nabla_z f_i(z_d). \tag{7}$$

Note that the gradient recentred barrier $F_{z_d}(z)$ is a continuous convex function on the interior of G. The minimum of $F_{z_d}(z)$ occurs at $\nabla F_{z_d}(z) = \mathbf{0}$, i.e. when $\nabla F(z) + b_{z_d} = \mathbf{0}$ which occurs at $z = z_d$. Moreover, it can easily be shown that $F_{z_d}(z_d) = 0$. Figure 2(b) provides a simple illustration of the gradient recentred barrier for recentring about the origin.

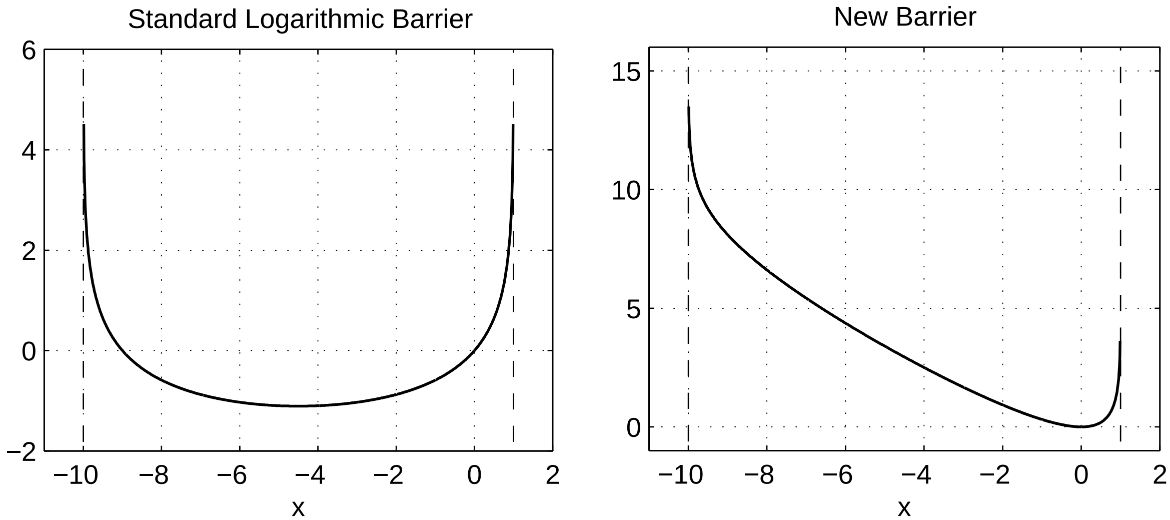

(a) Standard Logarithmic Barrier Function: This is the standard logarithmic barrier function for $-10 \leq x \leq 1$. Note that the minimum does not occur at the origin.

(b) Gradient Recentred Barrier Function: This is the new barrier function for $-10 \leq x \leq 1$. Note that the minimum occurs at the origin.

Figure 2: Comparison of logarithmic and gradient recentred barrier functions.

3 A Class of Receding Horizon Controllers

In this section we present a novel class of receding horizon controllers which are based on the gradient recentred barrier function (see Definition 1). As is the nature of receding horizon control, at each time interval we solve an optimisation problem to determine a sequence of controls, then apply the first control move to the system.

Let the system dynamics be described by the following equation,

$$x(k+1) = g(x(k), u(k)), \tag{8}$$

where $x(k) \in \mathbb{R}^n$ and $u(k) \in \mathbb{R}^m$ are the system state and inputs at time k respectively. For prediction horizon N, we define the input sequence space as $U = \mathbb{R}^{N \times m}$, and an input sequence $\mathcal{U}_k \in U$ at time k as,

$$\mathcal{U}_k := \{u(k), u(k+1), \ldots, u(k+N-1)\}\,. \tag{9}$$

Let $u_{ss} \in \mathbb{R}^m$ denote the desired steady-state input and let $\mathcal{U}_{ss} \in U$ denote the sequence of steady-state inputs formed by using u_{ss} for each element.

We consider the case where at time k, with given state $x(k)$, we may express the receding horizon control objective $f_0(\mathcal{U}_k)$ as a convex function of future inputs.

In this paper we only consider static input constraints. Let $\mathbb{U}$ denote a feasible domain for the individual control vectors $u(k)$. We define $\mathbb{U}$ as,

$$\mathbb{U} := \{u \in \mathbb{R}^m : f_i(u) \leq 0,\ \text{for } i = 1, \ldots, M\}\,, \tag{10}$$

where each $f_i(.)$ is a convex function of u. Furthermore, let $\mathbb{U}_s$ denote the set of admissible input sequences defined as,

$$\mathbb{U}_s := \{\mathcal{U} : u(j) \in \mathbb{U}, \text{ for } j = 0, \ldots, N-1\}, \tag{11}$$

where $u(j)$ denotes the j'th element of $\mathcal{U}$. Let $\mathbb{U}^0$ and $\mathbb{U}_s^0$ denote the interior of $\mathbb{U}$ and $\mathbb{U}_s$ respectively. Then for $u_{ss} \in \mathbb{U}^0$, let $F_{u_{ss}}(u)$ denote the recentred logarithmic barrier for $\mathbb{U}^0$ (see (6)). Furthermore, let $F^s_{u_{ss}}(\mathcal{U})$ denote the recentred logarithmic barrier for $\mathbb{U}_s^0$ defined as,

$$F^s_{u_{ss}}(\mathcal{U}) := \sum_{j=0}^{N-1} F_{u_{ss}}(u(j)), \tag{12}$$

where $u(j)$ denotes the j'th element of $\mathcal{U}$. In the sequel, we adopt a slight abuse of notation, and may refer to $\mathcal{U}$ as a sequence, or, as a vector given by $\mathcal{U} = [u^T(k), u^T(k+1), \ldots, u^T(k+N-1)]^T$.

Let the class of receding horizon optimisation problems be defined for the parameter $\mu \in (0, \infty)$,

$$(\mathcal{RC}_\mu): \quad \underset{\mathcal{U}\in U}{\text{argmin}}\ f_0(\mathcal{U}) + \mu F^s_{u_{ss}}(\mathcal{U}). \tag{13}$$

The associated receding horizon control law is constructed in the standard manner by selecting the first control move from the solution to $(\mathcal{RC}_\mu)$. An instance of this class of controllers will be termed a μ_p-controller, where μ_p refers to the fixed weighting parameter for the corresponding unconstrained minimisation problem $(\mathcal{RC}_{\mu_p})$.

Remark 3.1 *Our controller requires $u_{ss} \in \mathbb{U}^0$. This is a common assumption for many MPC controllers in the literature (Mayne* et al.*, 2000). Nevertheless, a common requirement in practice is that the desired steady-state should lie on or near the boundary of the constraint set (Qin and Badgwell, 1997). Indeed the desired steady-state is often derived as the solution of a separate linear, quadratic or semi-definite program (Qin and Badgwell, 1997) (Kassmann* et al.*, 2000) (Muske and Rawlings, 1993). In the context of this paper, a natural resolution is to find u_{ss} as a point on the corresponding central path (with associated parameter μ_{ss} being small but not infinitesimally so). This ensures u_{ss} lies on the interior; furthermore the consequent degradation in performance is bounded by the duality gap (Nesterov and Nemirovskii, 1994), and can be pre-specified to be arbitrarily small. This is the approach we adopt for simulations in section 5.*

Remark 3.2 *One benefit of the proposed method is that steady-state and dynamic performance have separate 'tuning knobs', namely μ_{ss} and μ_p respectively. This idea will be demonstrated in section 5 by way of simulation.*

4 Recentered Barrier MPC Example

In this section we give a detailed illustration of a μ_p-controller for a linear plant model and static input constraints - we consider both linear and convex quadratic constraints. In section 4.3 we provide a global stability proof for this case.

4.1 Model Predictive Control Preliminaries

This section describes some preliminary notation and definitions used to describe the limiting case receding horizon cost function that is considered throughout the remainder of this paper. We are concerned with linear time-invariant discrete-time state-space systems of the form,

$$\tilde{x}(k+1) = A\tilde{x}(k) + B\tilde{u}(k) \tag{14}$$

where $\tilde{x} \in \mathbb{R}^n$ and $\tilde{u} \in \mathbb{R}^m$. As mentioned in section 2, it is often desirable to regulate the system about some point $x_{ss} \in \mathbb{R}^n$. Corresponding to x_{ss} is a (not necessarily unique) steady-state input denoted by $u_{ss} \in \mathbb{R}^m$. We will often make reference to system (14) in its original co-ordinates *and* in the shifted co-ordinate form defined by letting $x(.) = \tilde{x}(.) - x_{ss}$ and $u(.) = \tilde{u}(.) - u_{ss}$.

For prediction horizon N, let $\mathcal{U}$ denote an input sequence, in shifted co-ordinate form, defined in the same manner as (9). Let $\mathcal{U}_{ss}$ denote the steady-state input sequence for given u_{ss}. Therefore, $\mathcal{U} + \mathcal{U}_{ss}$ refers to the corresponding input sequence in original co-ordinate form. Let $\mathbb{U}$ denote the feasible domain for individual input vectors $u + u_{ss}$ (see (10)), where in this case each $f_i(u + u_{ss})$ is a linear or convex quadratic function for $i = 1, \ldots, M$. Let $\mathbb{U}_s$ denote the corresponding set of admissible control sequences defined in the same manner as (11). Let $F_{u_{ss}}(u + u_{ss})$ denote a recentred logarithmic barrier for $\mathbb{U}^0$ (see (6)) and let $F^s_{u_{ss}}(\mathcal{U} + \mathcal{U}_{ss})$ denote a recentred logarithmic barrier for $\mathbb{U}^0_s$, defined in the same manner as (12). It is assumed that $\mathbb{U}$ is bounded and $\mathbb{U}^0 \neq \emptyset$ with $\mathbf{0} \in \mathbb{U}^0$. Moreover, we assume that $u_{ss} \in \mathbb{U}^0$.

In the sequel we take advantage of the shift invariant nature of the problem and let the current state $\mathbf{x}(k)$ be expressed as $\mathbf{x}(0)$. The limiting case receding horizon control cost with initial state $\mathbf{x}(0) = \mathbf{x}(k)$ and prediction horizon N is then,

$$J_N(x(0), \mathcal{U}) = \|x(N)\|^2_P + \sum_{i=0}^{N-1} \left(\|x(i)\|^2_Q + \|u(i)\|^2_R\right) \tag{15}$$

where Q, R, and P are positive definite and symmetric matrices.

4.2 Terminal Constraint Set X_f

A pivotal concept in modern MPC stability, is the terminal constraint set X_f, which is invariant under a local controller $\kappa_f(\cdot)$ (see Mayne *et. al.* (Mayne *et al.*, 2000)). In what follows, we define X_f in a similar way to Chen and Allgöwer (Chen and Allgower, 1998), but for the linear discrete case. This choice of terminal constraint set fits quite naturally into the current MPC example. However, we are not restricted to this particular choice.

Let Q and R be positive definite symmetric matrices. Furthermore, let $Q^* = \eta I + Q$ for some scalar $\eta \geq 0$. We consider the following discrete algebraic Ricatti equation,

$$P = (A - BK)^T P (A - BK) + Q^* + K^T R K \tag{16}$$

where K is given by,

$$K = (B^T P B + R)^{-1} B^T P A. \tag{17}$$

If system (14) is stabilisable, then there exists a unique positive definite symmetric matrix P, which solves (16), and a linear stabilising controller $\kappa_f(x) := -Kx$. Then, in the spirit of Chen and Allgöwer (Chen and Allgower, 1998), we have the following Lemma,

Lemma 4.2.1 *Let $u_{ss} \in \mathbb{U}^0$ and let $\mu_p > 0$.*

(a) For the stabilising controller $\kappa_f(x) = -Kx$ determined above, we can define a series of regions X_α, sets X_u and X_n and a scalar $\alpha_{max} > 0$ as,

$$X_\alpha := \{x \in \mathbb{R}^n : x^T P x < \alpha\} \tag{18}$$

$$X_u := \{x \in \mathbb{R}^n : -Kx + u_{ss} \in \mathbb{U}\} \tag{19}$$

$$X_n := \{x \in \mathbb{R}^n : \eta\|x\|^2 \geq \mu_p F_{u_{ss}}(-Kx + u_{ss})\} \tag{20}$$

$$\alpha_{max} := \max\{\alpha \in (0, \infty) : X_\alpha \subseteq \{X_u \cap X_n\}\} \tag{21}$$

(b) We can define an invariant terminal constraint set X_f as,

$$X_f := X_{\alpha_{max}} \tag{22}$$

Proof. Since $\eta \geq 0$ may be chosen arbitrarily large, and $\mu_p > 0$ may be chosen arbitrary small, then we can always obtain sets in the form of (18), (19) and (20) (since $u_{ss} \in \mathbb{U}^0$). Hence, $\alpha_{max} > 0$ and therefore (a) is true. Since (a) is true, then the Lyapunov stability theorem says that the set $X_{\alpha_{max}}$ is invariant and (b) is therefore true. □

Since the terminal state $x(N)$ can be expressed as a linear function of the current state $x(0)$, and future input sequence $\mathcal{U}$, we may express the terminal constraint as $h(\mathcal{U}) \leq 0$, where $h(\mathcal{U}) := x^T(N)Px(N) - \alpha_{max}$. Note that $h(\mathcal{U})$ is a convex quadratic function of $\mathcal{U}$ for given $x(0)$. Therefore, the implicit terminal constraint defined above can be trivially included by extending $\mathbb{U}_s$ in the following way. Let $\mathbb{U}_f$ denote the extended admissible constraint set defined as,

$$\mathbb{U}_f := \{\mathcal{U} \in \mathbb{U}_s : h(\mathcal{U}) \leq 0, \text{ for given } x(0)\}. \tag{23}$$

Let $F^f_{u_{ss}}(\mathcal{U} + \mathcal{U}_{ss})$ denote a recentred barrier function for $\mathbb{U}^0_f$ defined as,

$$F^f_{u_{ss}}(\mathcal{U} + \mathcal{U}_{ss}) := F^s_{u_{ss}}(\mathcal{U} + \mathcal{U}_{ss}) - \ln(-h(\mathcal{U})). \tag{24}$$

Since $h(.)$ is defined about u_{ss}, then it is not necessary to recentre this term.

Remark 4.1 *The above terminal constraint set is highly conservative on at least two accounts. Firstly, it is possible to use a more encompassing invariant set - such as the Output Admissible Set (Gilbert and Tan, 1991). Moreover, in the definition of X_n we have neglected other terms that depend on the initial state and previous control move. Inclusion of these terms may, in many cases, result in the best choice of η being zero.*

Remark 4.2 *In the case where $\mu_p > 0$, then our proposed design transforms the limiting case MPC law (with inequality constraints) to an unconstrained non-linear MPC law. This non-linear MPC law seems to fall under the general class considered by Jadbabaie* et. al. *(Jadbabaie* et al.*, 2001) (albeit that our controller is defined in the discrete domain whereas their discussion concerns continuous MPC control laws). Their results may be interpreted as saying that: for prediction horizon long enough, feasibility is guaranteed. Furthermore, for a prediction horizon long enough the final state is guaranteed to lie in the invariant terminal constraint set X_f, and therefore including a terminal constraint is unnecessary.*

4.3 Stability of Recentered Barrier MPC

This section provides a proof of global stability for the class of μ_p-controllers defined in section 3 as applied to the current example. The proof is broadly based on Mayne *et. al.* (Mayne *et al.*, 2000).

For given $\mu_p > 0$, and initial state $x(0)$, we define the optimal control sequence $\mathcal{U}^o_{\mu_p}$ as,

$$\mathcal{U}^o_{\mu_p} := \underset{\mathcal{U} \in U}{\operatorname{argmin}} \; J_N(x(0), \mathcal{U}) + \mu_p F^f_{u_{ss}}(\mathcal{U} + \mathcal{U}_{ss}) \tag{25}$$

Corresponding to $\mathcal{U}^o_{\mu_p}$ is an optimal state evolution denoted by $\mathcal{X}^o_{\mu_p} := \left\{x, x^o_{\mu_p}(1), \ldots, x^o_{\mu_p}(N)\right\}$. Let $J^o_{\mu_p}(x(0))$ denote the optimal cost given by,

$$J^o_{\mu_p}(x(0)) := J_N(x(0), \mathcal{U}^o_{\mu_p}) + \mu_p F^f_{u_{ss}}(\mathcal{U}^o_{\mu_p} + \mathcal{U}_{ss}). \tag{26}$$

Let X_N denote the set of states x such that $\forall x \in X_N$, there exists a feasible solution to (25). Assuming that the current state lies inside X_N, we wish to show that $J^o_{\mu_p}(x^o_{\mu_p}(1)) - J^o_{\mu_p}(x(0)) \leq 0$. Since $J^o_{\mu_p}(x^o_{\mu_p}(1)) \leq J_\phi(x^o_{\mu_p}(1), \tilde{\mathcal{U}})$ for any control sequence $\tilde{\mathcal{U}}$, it suffices to show that for some sequence $\tilde{\mathcal{U}}$, the following holds,

$$J_{\mu_p}(x^o_{\mu_p}(1), \tilde{\mathcal{U}}) - J^o_{\mu_p}(x(0)) \leq 0. \tag{27}$$

Result 4.1 *Let $\kappa_f(x)$ be the linear stabilising controller determined in Lemma 4.2.1 with associated $\mu_p > 0$ and $\eta \geq 0$. For given initial state $x(0) \in X_N$, we can define an admissible control sequence $\tilde{\mathcal{U}}$ as,*

$$\tilde{\mathcal{U}} := \{u^o_{\mu_p}(1), \ldots, u^o_{\mu_p}(N-1), \kappa_f(x^o_{\mu_p}(N))\} \tag{28}$$

Furthermore, the following holds,

$$J_{\mu_p}(x^o_{\mu_p}(1), \tilde{\mathcal{U}}) - J^o_{\mu_p}(x(0)) \leq 0. \tag{29}$$

Proof. Since $x(0) \in X_N$, there exists an admissible control sequence $\mathcal{U}^o_{\mu_p}$ which steers $x(0)$ to $x^o_{\mu_p}(N) \in X_f$. Given the initial state $x^0_{\mu_p}(1)$, then the truncated input sequence $\{u^o_{\mu_p}(1), \ldots, u^o_{\mu_p}(N-1)\}$, steers $x^0_{\mu_p}(1)$ to the same point $x^o_{\mu_p}(N) \in X_f$. By construction, the set X_f is invariant under the local controller $\kappa_f(x^o_{\mu_p}(N))$ and therefore, $(A - BK)x^o_{\mu_p}(N) \in X_f$. Furthermore, since $\kappa_f(x^o_{\mu_p}(N)) + u_{ss} \in \mathbb{U}$, then $\tilde{\mathcal{U}}$ is an admissible control sequence. It remains to show that (29) is true.

Using (15), (16), (17), (26) and Lemma 4.2.1, it is sufficient for (29) that,

$$-\eta\|x^o_{\mu_p}(N)\|^2 - \|x(0)\|^2_Q - \|u^o_{\mu_p}(0)\|^2_R + \mu_p[F^f_{u_{ss}}(\tilde{\mathcal{U}}) - F^f_{u_{ss}}(\mathcal{U}^o_{\mu_p})] \le 0$$

Since $\|x\|^2_P \ge \|x\|^2_{P-Q^*}$ for any $x \in X_f$ (see (16)), then

$$-\ln(\alpha_{max} - \|x^o_{\mu_p}(N)\|^2_{P-Q^*}) + \ln(\alpha_{max} - \|x^o_{\mu_p}(N)\|^2_P) \le 0$$

Moreover, since $F_{u_{ss}}(u^o_{\mu_p}(0))$, $\|x\|^2_Q$ and $\|u^o_{x,\mu_p}(0)\|^2_R$ are all non-negative, then it suffices to show that,

$$\eta\|x^o_{\mu_p}(N)\|^2 \ge \mu_p F_{u_{ss}}(\kappa_f(x^o_{\mu_p}(N)) + u_{ss}) \tag{30}$$

which from the definition of X_f is satisfied. □

Remark 4.3 *In the limiting case as $\mu_p \to 0$, then the above converges to the usual case which is well studied and known to be stable for $\eta = 0$, see e.g. (Mayne* et al.*, 2000).*

4.4 Discussion

It is worth noting that even when the state is inside the terminal constraint set X_f, the control action is *not* determined by the associated linear control law $\kappa_f(x)$. If we expand the recentred barrier function as a Taylor series around the steady-state point u_{ss} we find

$$F^f_{u_{ss}}(\mathcal{U} + \mathcal{U}_{ss}) = F^f_{u_{ss}}(\mathcal{U}_{ss}) + \mathcal{U}^T W \mathcal{U} + O(|\mathcal{U}|^3) \tag{31}$$

where $W = \nabla^2 F^f_{u_{ss}}(\mathcal{U}_{ss})$. Thus the *local* behaviour approximates that given by the standard closed-form solution to

$$J_N(x(0), \mathcal{U}) = \|x(N)\|^2_P + \sum_{i=0}^{N-1} \left(\|x(i)\|^2_{Q+W} + \|u(i)\|^2_R\right) \tag{32}$$

In the case of the recentred logarithmic barrier we may say (loosely) that as u_{ss} approaches a constraint, then $Q + W$ will grow in the direction perpendicular to the constraint. If such a feature is undesirable then we should choose a barrier such that $\nabla^2 F^f_{u_{ss}}(\mathcal{U}_{ss}) = 0$. This may be achieved by further modifying the recentred barrier, although care must be taken to preserve convexity. An obvious choice would be to construct a new barrier $F^*_{u_{ss}}(\mathcal{U})$ as

$$F^*_{u_{ss}}(\mathcal{U}) := \beta(\mathcal{U}) \times F^f_{u_{ss}}(\mathcal{U} + \mathcal{U}_{ss}) \tag{33}$$

where $\beta(\mathcal{U}) := \mathcal{U}^T D \mathcal{U}$ for some positive definite symmetric D. The Taylor series expansion of this barrier begins with terms of order $|\mathcal{U}|^4$.

5 Simulation Example

An important feature of the proposed class of controllers is that the control action 'backs-off' in a continuous manner as the manipulated variables approach their constraint boundaries. The degree to which this occurs is determined by the tuning parameter μ_p. This can be seen from the following trivial simulation example. The plant model is given by,

$$x_{k+1} = \begin{bmatrix} -0.3 & -0.8 \\ 0.5 & 0 \end{bmatrix} x_k + \begin{bmatrix} 1 \\ 0 \end{bmatrix} u_k, \quad y_k = [0.5\ 0]x_k \tag{34}$$

where the input signal is constrained to lie within simple bounds given by $-1 \leq u_k \leq 0.4$. For a prediction horizon of $N = 10$, we applied a step-disturbance to the system and the results are illustrated in figures 3(a) and 3(b). Figures 3(c) and 3(d) show another case where the step disturbance magnitude is large and the system is unable to return to the origin. This example illustrates the concept of separate 'tuning knobs' for the steady-state and dynamic cases.

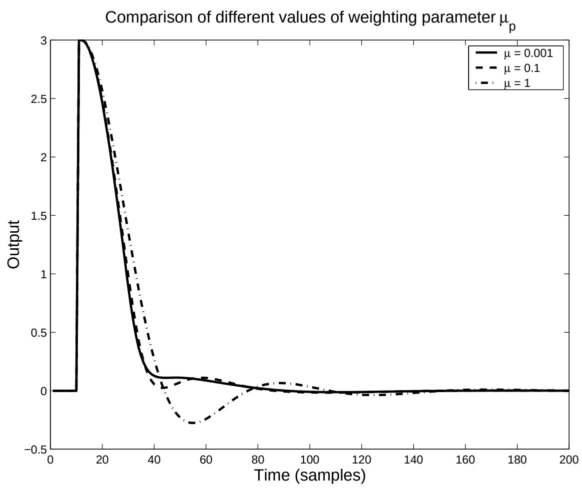

(a) Output signal: This figure shows that increasing the weighting parameter results in a controller that is more cautious. Note the marginal difference between $\mu_p = 0.001$ and $\mu_p = 0.1$.

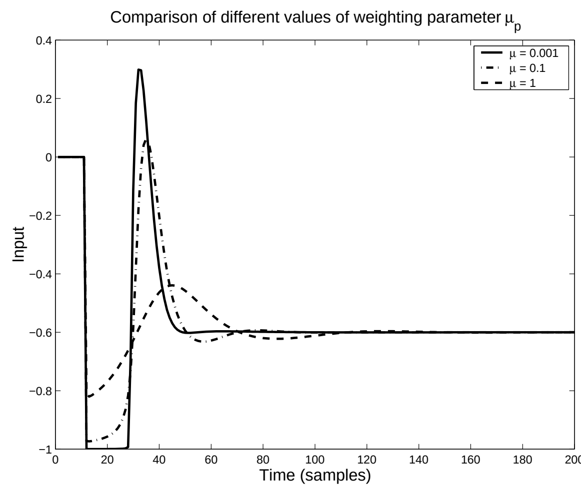

(b) Inputs signals: This figure illustrates the cautiousness of the control law as μ_p is increased. Note that away from the constraint boundary, the signals behave quite similarly.

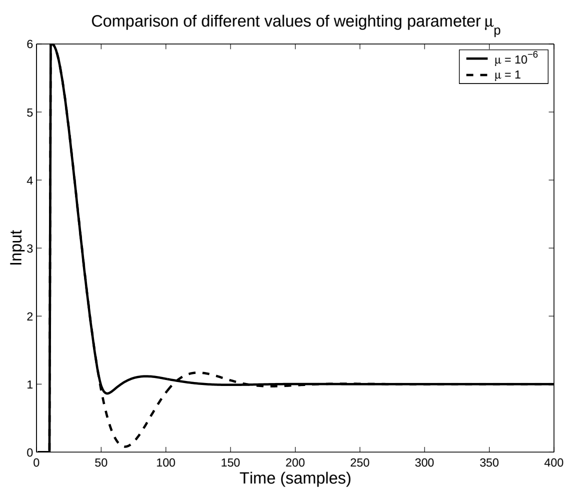

(c) Output signals: We see that both controllers converge to the same point, even though the parameter values differ by six orders of magnitude. This illustrates the benefit of recentring

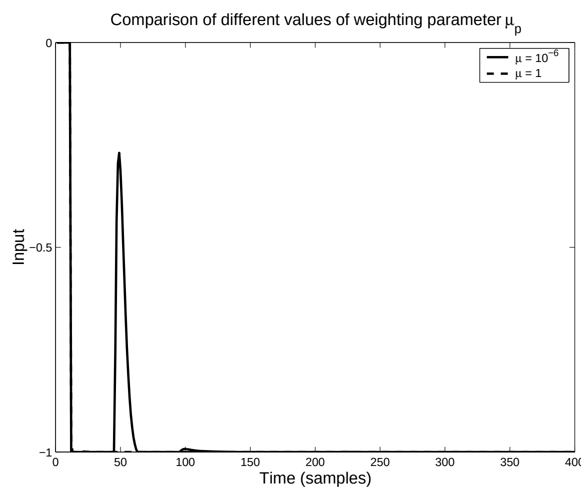

(d) Inputs signals: This figure shows that the input trajectories for the two controllers coincide in steady-state operation, even with a large difference between parameter values.

Figure 3: Response of different controllers to different output step disturbances.

6 Conclusion

In this paper we have presented a new class of receding horizon controllers based on the recentred barrier function. As illustrated in simulations, the significance of this class of controllers is that the control action becomes cautious near constraint boundaries. We conjecture that such behaviour is highly desirable in many physical systems where uncertainty exists near constraint boundaries. The degree of cautiousness is directly related to the positive weighting parameter which characterises the controller class. Moreover, it is possible to specify separate parameter values for the dynamic and steady-state cases. An interesting outcome from this paper is that a *constrained* receding horizon control problem is replaced with an *unconstrained* receding horizon control problem. Therefore, in the case of general non-linear objective functions with convex inequality constraints, using the barrier approach creates a problem that is essentially *no worse* than minimising the objective function by itself.

References

Chen, H. and F. Allgower (1998). A quasi-infinite nonlinear model predictive control scheme with guaranteed stability. *Automatica.*

Fiacco, A. V. and G. P. McCormick (1968). *Nonlinear Programming: Sequential Unconstrained Minimization Techniques.* John Wiley & Sons.

Gilbert, E. G. and K. T. Tan (1991). Linear systems with state and control constraints: The theory and application of maximal output admissible sets. *IEEE Transactions on Automatic Control.*

Jadbabaie, Ali, Jie Yu and John Hauser (2001). Unconstrained receding horizon control on nonlinear systems. *IEEE Transactions on Automatic Control* **46**(5), 776–783.

Kassmann, Dean E., Thomas A. Badgwell and Robert B. Hawkins (2000). Robust teady-state target calculation for model predictive control. *AIChE* **46**(5), 1007–1024.

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

Muske, K. R. and J. B. Rawlings (1993). Model predictive control with linear models. *AIChE Journal.*

Nesterov, Y. and A. Nemirovskii (1994). *Interior-point Polynomial Algorithms in Convex Programming.* SIAM Philidelphia.

Qin, S. Joe and Thomas A. Badgwell (1997). An overview of industrial model predictive control technology. *AIChE Symposium Series, 5th International Symposium on Chemical Process Control* **93**, 232–256.

Rao, C. V., S. J. Wright and J. B. Rawlings (1998). Application of interior point methods to model predictive control. *Journal of Optimization Theory and Applications* **99**(3), 723–757.

Wills, Adrian and William P. Heath (2001). Using a modified predictor-corrector algorithm for model predictive control. *Submitted for publication.*

Wright, S. J. (1997). Applying new optimization algorithms to model predictive control. *Chemical Process Control-V, CACHE, AIChE Symposium Series* **93**(316), 147–155.

Ye, Y. (1997). *Interior Point Algorithms - Theory and Analysis.* John Wiley & Sons Inc.