

An Exterior/Interior-point Approach to Infeasibility in Model Predictive Control

A. G. Wills and W. P. Heath

Centre for Integrated Dynamics and Control,

University of Newcastle, NSW, Australia.

Email: onyx@ecemail.newcastle.edu.au.

Abstract

In this paper we address the issue of infeasibility in Model Predictive Control. One popular way of dealing with infeasibility is to categorise constraints as being either hard or soft and include a penalty for soft constraint violations in the cost function. This penalty term can be made exact in the sense that all constraints will be satisfied if indeed possible. The exact penalty property is typically ensured by further including a linear term in the cost. If this term is not large enough, then the exact penalty property may be lost. We illustrate, by way of example, that choosing this term to be large may be detrimental to system dynamics. We propose a two stage process where feasibility is detected and then a Model Predictive Control problem is solved using this new information. Our development is based on the classical exterior/interior-point framework which allows for an intuitive tuning procedure.

1 Introduction

Model Predictive Control (MPC) is a receding horizon control strategy where an estimate of the current state is used as the initial state of a finite horizon optimal control problem, typically with constraints on the state and control action. Solving this problem at each time step gives a sequence of control moves. The current control action is chosen as the first control move from this sequence. Feedback enters this control strategy via an updating of the current state estimate. A number of surveys on MPC exist, including those by (Garcia *et al.* 1989, Muske and Rawlings 1993, Qin and Badgwell 1997, Chen and Allgöwer 1998, Mayne *et al.* 2000, Maciejowski 2002). In particular, Mayne *et al.* (2000) offer a comprehensive survey of MPC from a theoretical perspective.

It is often convenient to categorise the types of constraints encountered in MPC as either hard or soft. Hard constraints must be strictly satisfied while soft constraints should be satisfied if possible. If all constraints can be strictly satisfied, the optimisation problem is said to be feasible. By contrast, an infeasible problem occurs when not all constraints can be strictly satisfied. We will assume that the optimisation problem is always feasible if only hard constraints are included. Since soft constraints do not determine feasibility, we are always able to solve the problem under this assumption. Inclusion of constraints and their respective categorisation is a control engineering decision – see e.g. Qin and Badgwell (1997). For example, actuator limits are often included as hard constraints on the inputs. Meanwhile, state constraints are often deemed to be soft since strict satisfaction may not always be possible.

Scokaert and Rawlings (1999) summarise two different strategies found in the literature for handling the infeasible case (see also Maciejowski (2002)). The first strategy – originally proposed by Rawlings and Muske (1993) – determines a minimum time in the prediction horizon after which all constraints can be strictly satisfied. Prior to this time, only hard constraints are required to be strictly satisfied. Scokaert and Rawlings (1999) call this the “minimum-time” strategy. The second strategy includes a penalty term in the cost function which penalises soft constraint violations. This is called the “soft-constraint” strategy and is often formulated using an exact penalty. In this paper we are concerned with the soft-constraint strategy.

The desirable property of an exact penalty method is that all constraints are strictly satisfied if possible. If not, then violations of the soft constraints are penalised in the cost. Exact penalty functions are summarised in e.g. Fletcher (1987). Mayne and Polak (1987) use an exact penalty function algorithm to handle state and input constraints in a general optimal control setting. Their algorithm adaptively chooses a finite valued scaling term which results in an exact penalty. They do not address the issue of infeasibility.

Scokaert and Rawlings (1996) propose the use of exact penalty functions to handle the infeasible case (see also Scokaert and Rawlings (1999)). Rao *et al.* (1998) also consider the exact penalty approach in their discussion of interior-point methods for linear MPC. The exact penalty property, however is usually guaranteed by including a suitably chosen linear term in the cost function. If the linear term is not sufficiently large, then the exact penalty property may be lost. Unfortunately it is difficult to determine a suitable linear term a priori. Kerrigan and Maciejowski (2000) provide a lower bound for the linear term in the setting of multi-parametric quadratic programming and MPC.

If the MPC optimisation problem is feasible then – provided the linear term is sufficiently large – the exact penalty approach does not interact with the controller dynamics. If infeasible, however,

a tradeoff exists between satisfying constraints and returning to the origin of the system. By choosing the linear term to be large, emphasis is inherently placed on satisfying constraints. System dynamics may be consequentially awry.

In this paper, we propose a two step strategy which subsumes the exact penalty approach as a special case. The first step determines feasibility of the MPC optimisation problem. The second step solves a suitably chosen MPC problem depending on the result of the first step. Separating the problem into two stages allows for a more intuitive tuning procedure. Our discussion is based on the interior/exterior point approach presented in Fiacco and McCormick (1968).

Section 2 introduces the MPC formulation considered in this paper. Soft and hard constraints are separated and a related MPC problem is stated using logarithmic barrier functions for hard constraints and quadratic loss penalty functions for soft constraints. Section 3 presents a simple MPC algorithm which incorporates an intermediate step to detect feasibility. Section 4 illustrates the potential benefits of using this algorithm via a simulation example. Section 5 concludes this paper.

2 Barriers and Penalties and MPC

The MPC optimisation problem we consider in the paper is expressed as follows. Let $x(t) \in \mathbb{R}^{n_x}$ denote the current state. The system dynamics are expressed as,

$$x_{k+1} = f(x_k, u_k), \quad (1)$$

where $u_k \in \mathbb{R}^{n_u}$ is the system input vector. It is assumed that $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ is an affine function of the state and input, i.e. either the system is linear or is a linear approximation to a non-linear system (the linear approximation may include a constant offset term due to non-zero initial conditions).

Constraints on the states and inputs are represented mathematically as inclusion in the sets \mathbb{X} and \mathbb{U} respectively. A terminal constraint set is often included, this will be denoted by \mathbb{X}_F . The sets \mathbb{X} , \mathbb{X}_F and \mathbb{U} are expressed as follows,

$$\mathbb{X} = \{x \in \mathbb{R}^{n_x} : c_i^x(x) \leq 0, i = 1, \dots, M_x\}, \quad (2)$$

$$\mathbb{X}_F = \{x \in \mathbb{R}^{n_x} : c_i^F(x) \leq 0, i = 1, \dots, M_F\}, \quad (3)$$

$$\mathbb{U} = \{u \in \mathbb{R}^{n_u} : c_i^u(u) \leq 0, i = 1, \dots, M_u\}. \quad (4)$$

$$(5)$$

Each of the functions $c_i^x : \mathbb{R}^{n_x} \rightarrow \mathbb{R}$, $c_i^F : \mathbb{R}^{n_x} \rightarrow \mathbb{R}$ and $c_i^u : \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ is assumed at least to

be continuous and convex. It is assumed that $\mathbb{X}_F \subseteq \mathbb{X}$. Note that \mathbb{X} , \mathbb{X}_F and \mathbb{U} are convex sets. It is assumed that \mathbb{X} , \mathbb{X}_F and \mathbb{U} are closed with the further restriction that \mathbb{U} be compact. The interiors of \mathbb{X} , \mathbb{X}_F and \mathbb{U} are assumed to be non-empty.

It is convenient to define a family of sequence spaces denoted by \mathbb{S}_q^p where p and q are non-negative integers,

$$\mathbb{S}_q^p = \prod_{i=0}^p \mathbb{R}^q. \quad (6)$$

Let the positive integer N denote the prediction horizon. Define a sequence \mathcal{X} (called a state sequence) with $N + 1$ elements where each element is a vector of dimension n_x , i.e.

$$\mathcal{X} \in \mathbb{S}_{n_x}^N, \quad \mathcal{X} = \{x_0, \dots, x_N\}.$$

Similarly let \mathcal{U} denote a sequence (called an *input sequence*) with N elements where each element is a vector of dimension n_u ,

$$\mathcal{U} \in \mathbb{S}_{n_u}^{N-1}, \quad \mathcal{U} = \{u_0, \dots, u_{N-1}\}.$$

Define the MPC cost function as,

$$J(\mathcal{X}, \mathcal{U}) = F(x_N) + \sum_{i=0}^{N-1} \ell(x_i, u_i), \quad (7)$$

where the terminal cost $F : \mathbb{R}^{n_x} \rightarrow \mathbb{R}$ and stage cost $\ell : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ are assumed to be continuous and convex. It is further assumed that F is a positive semi-definite function and $\ell(\cdot, u)$ is a positive semi-definite function with $\ell(x, \cdot)$ a positive definite function. Assuming that the origin lies in the interior of \mathbb{X}_F and \mathbb{U} the MPC optimisation problem is expressed as,

$$\begin{aligned} (\mathcal{MPC}) : \quad (\mathcal{X}^*, \mathcal{U}^*) &= \arg \min_{\mathcal{X}, \mathcal{U}} J(\mathcal{X}, \mathcal{U}) \\ \text{s.t.} \quad x_0 &= x(t) \\ x_{k+1} &= f(x_k, u_k), \quad k = 0, \dots, N-1, \\ u_k &\in \mathbb{U}, \quad k = 0, \dots, N-1, \\ x_k &\in \mathbb{X}, \quad k = 1, \dots, N-1, \\ x_N &\in \mathbb{X}_F. \end{aligned}$$

It is straightforward to show that (\mathcal{MPC}) is a convex optimisation problem. Indeed J is convex, \mathbb{X} , \mathbb{X}_F and \mathbb{U} are convex and the equality constraints are affine.

Problem (\mathcal{MPC}) may not be feasible in general. In what follows we separate hard and soft constraints and develop a related problem to (\mathcal{MPC}) which is always well defined. The proviso

is that (\mathcal{MPC}) is assumed to be always feasible when only hard constraints are included. This is always achievable if only input constraints are considered to be hard (and the input constraint set is non-empty).

Let \mathcal{I}_H^u denote a set of indices representing the hard input constraints and let \mathcal{I}_S^u denote a set of indices representing soft input constraints. It is assumed that $\mathcal{I}_H^u \cap \mathcal{I}_S^u = \emptyset$ and $\mathcal{I}_H^u \cup \mathcal{I}_S^u = \{1, \dots, M_u\}$. Let \mathbb{U}_H and \mathbb{U}_S be the corresponding hard and soft input constraint sets. Let \mathcal{I}_H^x , \mathcal{I}_S^x , \mathbb{X}_H and \mathbb{X}_S be defined in a similar manner for state constraints. Let \mathcal{I}_H^F , \mathcal{I}_S^F , \mathbb{X}_{FH} and \mathbb{X}_{FS} be defined in a similar manner for terminal state constraints. It is assumed that the following problem is always feasible,

$$\begin{aligned}
(\mathcal{MPC}_H): \quad (\mathcal{X}^*, \mathcal{U}^*) &= \arg \min_{\mathcal{X}, \mathcal{U}} J(\mathcal{X}, \mathcal{U}) \\
\text{s.t.} \quad x_0 &= x(t) \\
x_{k+1} &= f(x_k, u_k), \quad k = 0, \dots, N-1, \\
u_k &\in \mathbb{U}_H, \quad k = 0, \dots, N-1, \\
x_k &\in \mathbb{X}_H, \quad k = 1, \dots, N-1, \\
x_N &\in \mathbb{X}_{FH}.
\end{aligned}$$

Define a new MPC cost function as follows,

$$J_{HS}(\mathcal{X}, \mathcal{U}, \mu, \nu) = F_{HS}(x_N, \nu) + \sum_{i=0}^{N-1} \ell_{HS}(x_i, u_i, \mu, \nu) \quad (8)$$

where

$$F_{HS}(x, \mu, \nu) = F(x) + \mu \sum_{i \in \mathcal{I}_H^F} B(c_i^F(x)) + \nu \sum_{i \in \mathcal{I}_S^F} P(c_i^F(x)). \quad (9)$$

In the above B is a barrier function and P is a penalty function with the scalar terms μ and ν being positive. We restrict our attention to the logarithmic barrier and quadratic-loss penalty functions. Hence B and P are given by,

$$B: \mathbb{R}_- \rightarrow \mathbb{R} : B(z) = -\ln(-z), \quad (10)$$

$$P: \mathbb{R} \rightarrow \mathbb{R}_+ \cap 0 : P(z) = (z + |z|)^2. \quad (11)$$

In the above \mathbb{R}_- refers to the negative half axis and \mathbb{R}_+ to the positive half axis. The function ℓ_{HS} is defined similarly as,

$$\begin{aligned}
\ell_{HS}(x, u, \mu, \nu) &= \ell(x, u) + \mu \sum_{i \in \mathcal{I}_H^x} B(c_i^x(x)) + \mu \sum_{i \in \mathcal{I}_H^u} B(c_i^u(u)) \\
&\quad + \nu \sum_{i \in \mathcal{I}_S^x} P(c_i^x(x)) + \nu \sum_{i \in \mathcal{I}_S^u} P(c_i^u(u)).
\end{aligned} \quad (12)$$

Recall that the interiors of \mathbb{X}_H , \mathbb{X}_{FH} and \mathbb{U}_H are assumed to be non-empty. Hence there exists a point for which the logarithmic barrier terms are well defined. A family of MPC optimisation problems can be defined as,

$$\begin{aligned}
 (\mathcal{MPC}_{\mu,\nu}) : \quad (\mathcal{X}(\mu,\nu), \mathcal{U}(\mu,\nu)) &= \arg \min_{\mathcal{X}, \mathcal{U}} J_{HS}(\mathcal{X}, \mathcal{U}, \mu, \nu) \\
 \text{s.t.} \quad x_0 &= x(t) \\
 x_{k+1} &= f(x_k, u_k), \quad k = 0, \dots, N-1.
 \end{aligned}$$

Note that J_{HS} is a convex function and $(\mathcal{MPC}_{\mu,\nu})$ is a convex optimisation problem for all $\mu \in (0, \infty)$ and all $\nu \in [0, \infty)$.

If (\mathcal{MPC}) is feasible, then Fiacco and McCormick (1968) give conditions for a sequence of μ 's and ν 's such that the solution to $(\mathcal{MPC}_{\mu,\nu})$ converges to the solution of (\mathcal{MPC}) (see Section 4.3 in Fiacco and McCormick (1968)). In essence, $\mu \rightarrow 0$ and $\nu \rightarrow \infty$. This gives the same result as using an exact penalty function in this case.

If (\mathcal{MPC}) is not feasible then the exact penalty approach can potentially associate severe penalties on constraint violation – which may be undesirable (see Section 4). The same effect is seen for large values of ν in this case.

In the next section we propose an intermediate step where feasibility of (\mathcal{MPC}) is determined and the value of ν along with hard and soft constraints are then chosen according to some design criterion – which reflects knowledge of the system. The resulting algorithm considers the case where $\mu \rightarrow 0$ but ν is chosen to be finite at each time step. We have considered the case where all constraints are hard and μ is fixed in Heath and Wills (2002) and Wills and Heath (2002).

3 Detecting Infeasibility

Detecting infeasibility of (\mathcal{MPC}) can be achieved by embedding the problem into a slightly bigger optimisation problem. Indeed, this type of embedding is often used in two-stage algorithms; the first stage finds an initial feasible point and the second stage solves the optimisation problem starting from this initial point. Fiacco and McCormick (1968) propose the use of such two-stage algorithms to solve problems like $(\mathcal{MPC}_{\mu,\nu})$.

One popular embedding is described as follows. Define the sets $\bar{\mathbb{X}}_S$, $\bar{\mathbb{X}}_{FS}$ and $\bar{\mathbb{U}}_S$ as,

$$\bar{\mathbb{X}}_S = \{(x, z) \in \mathbb{R}^{n_x} \times \mathbb{R} : c_i^x(x) \leq z, i \in \mathcal{I}_S^x\}, \quad (13)$$

$$\bar{\mathbb{X}}_{FS} = \{(x, z) \in \mathbb{R}^{n_x} \times \mathbb{R} : c_i^F(x) \leq z, i \in \mathcal{I}_S^F\}, \quad (14)$$

$$\bar{\mathbb{U}}_S = \{(u, z) \in \mathbb{R}^{n_u} \times \mathbb{R} : c_i^u(u) \leq z, i \in \mathcal{I}_S^u\}. \quad (15)$$

$$(16)$$

Define a feasibility optimisation problem as,

$$\begin{aligned} (\mathcal{F}) : \quad (\mathcal{X}^*, \mathcal{U}^*, z^*) &= \arg \min_{\mathcal{X}, \mathcal{U}, z} z \\ \text{s.t.} \quad x_0 &= x(t) \\ x_{k+1} &= f(x_k, u_k), \quad k = 0, \dots, N-1, \\ u_k &\in \mathbb{U}_H, \quad k = 0, \dots, N-1, \\ x_k &\in \mathbb{X}_H, \quad k = 1, \dots, N-1, \\ x_N &\in \mathbb{X}_{FH}, \\ (u_k, z) &\in \bar{\mathbb{U}}_S, \quad k = 0, \dots, N-1, \\ (x_k, z) &\in \bar{\mathbb{X}}_S, \quad k = 1, \dots, N-1, \\ (x_N, z) &\in \bar{\mathbb{X}}_{FS}, \\ z &\geq \gamma. \end{aligned}$$

In the above γ is some negative constant. If the solution to (\mathcal{F}) has $z^* \leq 0$ then (\mathcal{MPC}) is feasible. If $z^* > 0$ then (\mathcal{MPC}) is not feasible. It may be desirable to include a separate slack variable for each soft constraint to provide more information.

The value of z^* may be used to determine an appropriate value of ν . If for example $z^* > 0$, then the size of z^* shows the maximum constraint violation. If it is important that constraints are not violated then ν may be chosen large.

A simple algorithm for MPC with hard and soft constraints based on $(\mathcal{MPC}_{\mu, \nu})$ is given as follows,

Algorithm 3.1 *At each time step t , given the current state $x(t)$, perform the following steps,*

1. *Solve (\mathcal{F}) .*
2. *Determine a value of ν according to knowledge of the system.*
3. *Solve $(\mathcal{MPC}_{\mu, \nu})$ for ν from the previous step and $\mu \rightarrow 0$.*
4. *Let the system input be the first element of $\mathcal{U}^*(\mu, \nu)$.*

One possibility would be to choose $\nu \rightarrow \infty$ when $z^* \leq 0$ and ν constant when $z^* > 0$. We use an heuristic rule for ν in the simulation below.

4 Simulation

The simulation presented in this section illustrates the benefits of using an intermediate step to determine feasibility. For comparison we have used an exact penalty method with linear term large. Similar results would be obtained by fixing ν to be large.

The system is modelled as,

$$x_{k+1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 \end{bmatrix} x_k + \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} u_k \quad (17)$$

We are using the standard linear MPC formulation where the cost J is given by,

$$J(\mathcal{X}, \mathcal{U}) = \|x_N\|_{Q_N}^2 + \sum_{i=0}^{N-1} (\|x_i\|_Q^2 + \|u_i\|_R^2). \quad (18)$$

The prediction horizon is chosen at $N = 5$. The matrices Q and R are given by,

$$Q = \begin{bmatrix} 0.7 & & & \\ & 0.7 & & \\ & & 0.7 & \\ & & & 0.7 \end{bmatrix}, \quad R = \begin{bmatrix} 33 & \\ & 33 \end{bmatrix}. \quad (19)$$

The terminal cost matrix Q_N is given as the solution to the Discrete Algebraic Riccati Equation for state and input weighting matrices Q and R respectively. Constraints on the system are as follows. The input must satisfy $\|u\|_\infty \leq 0.5$. Let $\{x\}_i$ denote the i 'th element of the vector x . The state should satisfy $\{x\}_2 \geq 0$ and $\{x\}_4 \geq 0$. There are no terminal constraints.

State disturbance was included in the simulation in the form of additive zero mean white noise is zero with diagonal covariance matrix, where all diagonal elements are equal to 0.2.

Figure 1 shows a plot of the state elements. The linear term for slack variables was set to $1000e$, where e is a vector of all ones (of appropriate dimension). A further disturbance is introduced at $t = 20$ in $\{x\}_4$. This large disturbance violates the state constraint. Because the linear term was chosen large to ensure the exact penalty property, the cost associates large penalties for this constraint violation. Hence the emphasis of the controller is to regain feasibility. Consequently, the system overshoots the origin.

Figure 2 shows a plot of the state elements using Algorithm 3.1. Step 2 is given by the following heuristic. When the constraint violation is small, we choose to make ν large, but as the violation increases we allow ν to become correspondingly smaller. As such, we employ the following rule for ν at time t (denoted as $\nu(t)$),

$$\nu(t) = \min(\exp(-z^*), 10^6), \quad (20)$$

where $z^*(t)$ comes from the solution to (\mathcal{F}) at time t .

5 Conclusion

In this paper we discuss a concern with treating the feasible and infeasible cases in one optimisation problem. This is typical of exact penalty methods. Adopting a two stage process, where feasibility is firstly determined, allows for more general tuning rules that may be useful in practice. The exterior/interior-point framework of Fiacco and McCormick (1968) combined with the two stage process subsumes the exact penalty approach as a special case. The benefits of using a two stage process are illustrated via a simulation example. In this example we use an heuristic choice for ν (the soft constraint penalty scaling factor) which gives improved performance.

References

- Chen, H. and F. Allgöwer (1998). A quasi-infinite nonlinear model predictive control scheme with guaranteed stability. *Automatica* **34**(10), 1205–1217.
- Fiacco, A. V. and G. P. McCormick (1968). *Nonlinear Programming: Sequential Unconstrained Minimization Techniques*. John Wiley & Sons.
- Fletcher, R. (1987). *Practical Methods of Optimization*. John Wiley & Sons. New York.
- Garcia, C. E., D. M. Prett and M. Morari (1989). Model predictive control: theory and practice - a survey. *Automatica* **25**(3), 335–348.
- Heath, W. P. and A. G. Wills (2002). Recentred barriers for model predictive control with state constraints. *Non-linear Predictive Control Workshop at the University of Oxford*.
- Kerrigan, E. C. and J. M. Maciejowski (2000). Soft constraints and exact penalty functions in model predictive control. In: *Proceedings of the UKACC International Control Conference*. Cambridge, UK.

- Maciejowski, J. M. (2002). *Predictive Control with Constraints*. Pearson Education Limited.
- Mayne, D. Q. and E. Polak (1987). An exact penalty function algorithm for control problems with state and control constraints. *IEEE Trans. Automat. Control* **AC-32**(5), 380–387.
- 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* **39**(2), 262–287.
- Qin, S. Joe and T. 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.
- Rawlings, J. B. and K. Muske (1993). The stability of constrained receding horizon control. *IEEE Trans. Automat. Control* **38**(10), 1512–1516.
- Scokaert, P. O. M. and J. B. Rawlings (1996). Infinite horizon linear quadratic control with constraints. In: *Proceedings of the 13th IFAC World Congress*. San Fransisco, California.
- Scokaert, P. O. M. and J. B. Rawlings (1999). Feasibility issues in linear model predictive control. *AiChE J.* **45**(8), 1649–1659.
- Wills, A. G. and W. P. Heath (2002). A recentered barrier for constrained receding horizon control. In: *Proceedings of the 2002 American Control Conference*. Vol. 5. pp. 4177–4182.




Figure 1: This figure shows the second and last elements of the state at each time interval. An exact penalty formulation is used. After the disturbance at $t = 20$, the constraint remains in violation for 5 time steps. Note the large overshoot.




Figure 2: This figure shows the second and last elements of the state at each time interval. Algorithm 3.1 is used with the heuristic (20) for the value of ν at each time step. After the disturbance at $t = 20$, the constraint remains in violation for 7 time steps. The overshoot has been considerably reduced.