

SOFT CONSTRAINTS AND EXACT PENALTY FUNCTIONS IN MODEL PREDICTIVE CONTROL

Eric C. Kerrigan and Jan M. Maciejowski

Department of Engineering, University of Cambridge

Trumpington Street, Cambridge CB2 1PZ, United Kingdom, <http://www-control.eng.cam.ac.uk>

Fax : +44 1223 332662. Email: eck21@eng.cam.ac.uk, jmm@eng.cam.ac.uk

ABSTRACT

One of the strengths of Model Predictive Control (MPC) is its ability to incorporate constraints in the control formulation. Often a disturbance drives the system into a region where the MPC problem is infeasible and hence no control action can be computed. Feasibility can be recovered by softening the constraints using slack variables. This approach does not necessarily guarantee that the constraints will be satisfied, if possible. Results from the theory of exact penalty functions can be used to guarantee constraint satisfaction. This paper describes a method for computing a lower bound for the constraint violation penalty weight of the exact penalty function. One can then guarantee that the soft-constrained MPC solution will be equal to the hard-constrained MPC solution for a bounded subset of initial states, control inputs and reference trajectories.

Keywords : feasibility, Lagrange multipliers, multi-parametric quadratic programming

INTRODUCTION

The success of Model Predictive Control (MPC) in industry is primarily due to the ease with which constraints on the inputs and states can be included in the control problem formulation. However, sometimes a disturbance drives the plant into a state for which the control problem is infeasible and hence a new control input cannot be computed. Heuristic methods such as removing constraints or repeating the previously computed input, are not optimal and could lead to unpredictable closed-loop behaviour.

A more systematic method for dealing with infeasibility is to “soften” the constraints by adding slack variables to the problem [1, 2] where the size of the slack variables correspond to the size of the associated constraint violations. The slack variables are added to the MPC cost function and the optimiser searches for a solution which minimises the original cost function, while keeping the constraint violations as small as possible.

It is desirable that the solution to the soft-constrained MPC problem be the same as the solution to the original hard-constrained MPC problem, if the latter is feasible. The theory of exact penalty functions can be used to derive a lower bound for the violation weight [3, Sect. 14.3]. The

problem, however, is that in MPC this weight is dependent on the current state of the system and it is therefore necessary to calculate the lower bound for the weight for all possible states that the plant could be in.

To the authors’ knowledge, a systematic method for computing a lower bound has not yet been published. A naive and impractical solution would be to grid the state space region of interest and compute the optimal Lagrange multipliers at each point. This method is computationally demanding and due to the finite nature of the grid one cannot guarantee that the true lower bound on the weight has been found. As mentioned in [4], a conservative state-dependent upper bound might be obtainable by exploiting the Lipschitz continuity of the quadratic program [5]. However, it is unclear as to how exactly one would proceed to implement this for the entire feasible state space.

This paper shows how the Karush-Kuhn-Tucker (KKT) conditions can be used to compute a lower bound by solving a finite number of linear programs (LPs). This method is therefore computationally less demanding than gridding and provides a guarantee that the lower bound has been found.

Once a lower bound has been computed, the soft-constrained MPC problem can be set up. This new MPC problem will produce a result where the original hard-constrained MPC problem would have been infeasible. The important result is that one can guarantee that the soft- and hard-constrained MPC problems will produce the same result for the region in which the latter is feasible.

The paper starts by defining a standard reference tracking formulation of MPC. It is shown that the cost function and constraints of the resulting quadratic program (QP) are dependent on the current plant state, previous control input and current reference trajectory. More precisely, the MPC problem can be treated as a multi-parametric quadratic program (mp-QP) [6]. This allows one to gain additional insight into the structure of the problem.

Following this, exact penalty functions are introduced in order to find a condition on the lower bound for the violation weight. By introducing slack variables the non-smooth, exact penalty function can be converted into a smooth, soft-constrained QP problem.

A procedure for setting up an optimisation routine for computing a non-conservative lower bound for the violation weight is described. This weight guarantees the exactness of the penalty function over an *a priori* chosen subset of feasible states.

A simple example is presented to show how a soft-constrained mp-QP could be set up to have the same solution as the original hard-constrained mp-QP. The paper concludes with a summary of the results.

MODEL PREDICTIVE CONTROL

A standard formulation for MPC is described below. The cost function and constraints are shown to be dependent on an augmented system state vector, which includes the current state, previous control input and reference trajectory. The feasible region for the MPC problem is defined.

Standard Formulation

Consider the following discrete-time LTI state-space model:

$$\hat{x}(k+1|k) = A\hat{x}(k|k) + Bu(k) \quad (1a)$$

$$\hat{z}(k) = C\hat{x}(k|k) \quad (1b)$$

where $\hat{x}(k+i|k) \in \mathbb{R}^n$ denotes an estimate of the plant state at time $k+i$ made at time k ; $u(k) \in \mathbb{R}^m$ is the real input to the plant; the controlled variables are $\hat{z}(k) \in \mathbb{R}^p$. Note that $\hat{x}(k|k) = x(k)$ is the current plant state.

The cost function to be minimised is:

$$\begin{aligned} V(\xi(k), \Delta\mathcal{U}(k)) = & \sum_{i=1}^{H_p} \|\hat{z}(k+i|k) - r(k+i)\|_{Q(i)}^2 \\ & + \sum_{i=0}^{H_u-1} \|\Delta\hat{u}(k+i|k)\|_{R(i)}^2. \end{aligned} \quad (2)$$

The first term in (2) penalises deviations of the controlled variables from the reference trajectory $r(k+i)$ and the second term penalises changes in the control input $\Delta\hat{u}(k) \triangleq \hat{u}(k) - \hat{u}(k-1)$. H_p and H_u are the output and control horizons; $Q(i) \succeq 0$ and $R(i) \succ 0$ are the weights on tracking error and control action where $\|\alpha\|_Q^2 \triangleq \alpha^T Q \alpha$. It is assumed that $H_u \leq H_p$ and $\Delta\hat{u}(k+i|k) = 0$ for $i \geq H_u$.

The cost function can be rewritten as:

$$V(\xi(k), \Delta\mathcal{U}(k)) = \|\mathcal{Z}(k) - \mathcal{T}(k)\|_Q^2 + \|\Delta\mathcal{U}(k)\|_{\mathcal{R}}^2 \quad (3)$$

where

$$\mathcal{Z}(k) \triangleq \begin{bmatrix} \hat{z}(k+1|k) \\ \vdots \\ \hat{z}(k+H_p|k) \end{bmatrix}$$

$$\mathcal{T}(k) \triangleq \begin{bmatrix} r(k+1|k) \\ \vdots \\ r(k+H_p|k) \end{bmatrix}$$

$$\Delta\mathcal{U}(k) \triangleq \begin{bmatrix} \Delta\hat{u}(k|k) \\ \vdots \\ \Delta\hat{u}(k+H_u-1|k) \end{bmatrix}.$$

$\mathcal{Z}(k)$ has the form $\mathcal{Z}(k) = \Psi x(k) + \Upsilon u(k-1) + \Theta \Delta\mathcal{U}(k)$. The matrices $\Psi, \Upsilon, \Theta, \mathcal{Q}$ and \mathcal{R} are obtained by substituting (1) into (2) and collecting terms, before defining $\Gamma \triangleq [-\Psi \quad -\Upsilon \quad I]$. It is necessary to define the augmented state vector:

$$\xi(k) \triangleq \begin{bmatrix} x(k) \\ u(k-1) \\ \mathcal{T}(k) \end{bmatrix}.$$

Cost function (3) is usually minimised subject to linear inequality constraints on the inputs, states and outputs of the plant, possibly also including the reference trajectory $\mathcal{T}(k)$:

$$\Omega \begin{bmatrix} \Delta\mathcal{U}(k) \\ \mathcal{U}(k) \\ \mathcal{Z}(k) \\ \mathcal{X}(k) \\ \mathcal{T}(k) \end{bmatrix} \preceq \omega \quad (4)$$

where $\mathcal{U}(k)$, the vector of future control inputs, and $\mathcal{X}(k)$, the vector of future plant states, are defined in a similar fashion as above; Ω and ω are problem-dependent. The MPC problem then reduces to the following strictly convex QP problem:

$$\begin{aligned} \min_{\Delta\mathcal{U}(k)} \frac{1}{2} \Delta\mathcal{U}(k)^T \mathcal{H} \Delta\mathcal{U}(k) + \Delta\mathcal{U}(k)^T \mathcal{G} \xi(k) \\ + \xi(k)^T \mathcal{F} \xi(k) \end{aligned} \quad (5a)$$

subject to

$$E \Delta\mathcal{U}(k) \preceq f + G \xi(k) \quad (5b)$$

where $\mathcal{F} = \Gamma^T \mathcal{Q} \Gamma$, $\mathcal{G} = -2\Theta^T \mathcal{Q} \Gamma$, and $\mathcal{H} = 2(\Theta^T \mathcal{Q} \Theta + \mathcal{R})$; E , f and G are obtained by substituting (1) into (4) and collecting terms. The term involving \mathcal{F} in (5a) is usually dropped, since it does not affect the optimal solution $\Delta\mathcal{U}^*(k)$.

Only the first part of the solution is used in accordance with the receding horizon strategy and the implemented control input is therefore

$$u^*(k) = u(k-1) + [I_m \quad 0_{m \times (H_u-1)}] \Delta\mathcal{U}^*(k). \quad (6)$$

The state at the next time instant $x(k+1)$ is measured and the process of setting up the QP and calculating the new control action is repeated.

Note that both the cost function (5a) and constraints (5b) are dependent on $\xi(k)$, which includes the current state $x(k)$, past input $u(k-1)$ and the reference trajectory $\mathcal{T}(k)$. The MPC problem can therefore be treated as an mp-QP for which an explicit solution can be computed offline [6].

Feasibility and Invariance

The QP constraints (5b) define the set of feasible control sequence and augmented state pairs:

$$\mathbb{F} \triangleq \{(\xi, \Delta\mathcal{U}) : E\Delta\mathcal{U} \preceq f + G\xi\} \quad (7)$$

and the assumption that $\mathbb{F} \neq \emptyset$ is made. Often some of the constraints on ξ and $\Delta\mathcal{U}$ in \mathbb{F} are redundant and removing these will improve computation time both on-line and when computing the constraint violation weight, as described later. The values of ξ for which the QP problem is feasible, i.e. for which a feasible control sequence exists, is therefore defined as:

$$\Xi_{\mathbb{F}} \triangleq \{\xi : \exists \Delta\mathcal{U} \text{ such that } (\xi, \Delta\mathcal{U}) \in \mathbb{F}\} . \quad (8)$$

The MPC problem is not defined for any other combination of past control input, current state and reference trajectory, i.e. if $\xi \notin \Xi_{\mathbb{F}}$ then the QP problem is infeasible. The set $\Xi_{\mathbb{F}}$ can be seen to be the orthogonal projection of \mathbb{F} onto the ξ subspace and can therefore be computed using standard techniques, such as the Fourier-Motzkin elimination method [7].

In general, even for the case with no disturbances and model uncertainty, the set $\Xi_{\mathbb{F}}$ is not necessarily *positively invariant* for the closed-loop system. Since constraints can be satisfied if and only if the initial condition $\xi(k)$ is in a set which is *positively invariant* [8] for the closed-loop system, it is important to design the controller such that $\Xi_{\mathbb{F}}$ is invariant. If $H_p = H_u$, one can guarantee nominal feasibility for all time by requiring that the predicted terminal state $\hat{\xi}(k + H_p|k)$ lie in a control invariant set, as discussed in [9]. For simplicity, it is assumed that $\Xi_{\mathbb{F}}$ is positively invariant for the nominal closed-loop system and that a feasible sequence of reference trajectories is always chosen.

Assuming the above, it is still possible that a disturbance or modelling error could result in the system being driven to a state where the problem is infeasible and hence no solution exists. One possible way of dealing with this situation is to soften some or all of the constraints.

SOFT CONSTRAINTS

A straightforward way for softening constraints is to introduce slack variables which are defined such that they are non-zero only if the corresponding constraints are violated. If the original, hard-constrained solution is feasible, one would like the soft-constrained problem to produce the same control action. In order to guarantee this the weights in the cost function have to be chosen large enough such that the optimiser tries to keep the slack variables at zero, if possible. Exact penalty functions can be used to guarantee this behaviour [3, Sect. 14.3].

Exact Penalty Functions

The general non-linear, constrained minimisation problem can be stated as:

$$\min_{\theta} V(\theta) \quad (9a)$$

subject to

$$c(\theta) \preceq 0 . \quad (9b)$$

This optimisation problem can be recast into the following equivalent unconstrained, non-smooth penalty function minimisation:

$$\min_{\theta} V(\theta) + \rho \|c(\theta)^+\| \quad (10)$$

where ρ is the constraint violation penalty weight, the vector $c(\theta)^+$ contains the magnitude of the constraint violations for a given θ and $c_i^+ \triangleq \max(c_i, 0)$.

The concept of a *dual norm* is used in the condition on ρ for which the solution θ^* to (10) is equal to the solution to (9). The dual of a norm $\|\cdot\|$ is defined as

$$\|u\|_D \triangleq \max_{\|v\| \leq 1} u^T v . \quad (11)$$

It can be shown that the dual of $\|\cdot\|_1$ is $\|\cdot\|_{\infty}$ and vice versa, and that $\|\cdot\|_2$ is the dual of itself.

If θ^* denotes the optimal solution to (9) and λ^* is the corresponding optimal Lagrange multiplier vector, then the following well-known result gives a condition under which the solutions to (9) and (10) are equal:

Theorem 1. *If the penalty weight $\rho > \|\lambda^*\|_D$ and $c(\theta^*) \preceq 0$, then the solution θ^* to (9) is equal to the solution to (10).*

Proof. See [3, Thm. 14.3.1]. □

If $\rho > \|\lambda^*\|_D$, then (10) is called an *exact penalty function*. The cost function (10) is non-smooth and therefore not as easy to solve for as, say, a QP. One way to overcome this difficulty is to introduce slack variables into the problem.

Slack Variables as Soft Constraints

The non-smooth, unconstrained minimisation (10) can be cast into the following equivalent smooth, constrained problem which is a lot easier to compute:

$$\min_{\theta, \epsilon} V(\theta) + \rho \|\epsilon\| \quad (12a)$$

subject to

$$c(\theta) \preceq \epsilon \quad (12b)$$

$$0 \preceq \epsilon \quad (12c)$$

where ϵ are the slack variables representing the constraint violations, i.e. $\epsilon = 0$ if the constraints are satisfied. If $V(\theta) = V(\xi(k), \Delta\mathcal{U}(k))$ as in (3), $c(\theta) = E\Delta\mathcal{U} - f - G\xi$, $\theta = \Delta\mathcal{U}$ and $\|\epsilon\|_1$ or $\|\epsilon\|_\infty$ is used in (12), then the problem can be formulated as a QP and solved using standard techniques [1, 2].

Note that even though the l_2 -norm $\|\epsilon\|_2 \triangleq \sqrt{\epsilon^T \epsilon}$ will result in a non-smooth penalty function, one cannot formulate the soft-constrained MPC problem as a QP because $V(\xi(k), \Delta\mathcal{U}(k))$ is quadratic and $\|\epsilon\|_2$ has a square root. Using the l_2^2 quadratic norm $\|\epsilon\|_2^2 \triangleq \epsilon^T \epsilon$ one can express the problem as a QP, but this does not result in an exact penalty function since (10) will be smooth; it is the non-smoothness of the penalty function which allows it to be exact¹.

COMPUTING A LOWER BOUND FOR THE PENALTY WEIGHT

In MPC, the optimal solution $\Delta\mathcal{U}^*$ is dependent on the current augmented state ξ as can be seen in (5) and hence the corresponding Lagrange multiplier λ^* is also dependent on ξ . The lower bound for ρ is therefore dependent on ξ .

One would therefore have to calculate a lower bound for ρ which guarantees that the soft-constrained MPC will produce the same solution as the original hard-constrained MPC for all $\xi \in \Xi_{\mathbb{F}}$. Duality in optimisation theory provides some insight into the relation of the Lagrange multipliers to ξ .

KKT conditions for mp-QP Problems

The Lagrangian of optimisation problem (5) is

$$\begin{aligned} \mathcal{L}(\Delta\mathcal{U}, \lambda) = & \frac{1}{2} \Delta\mathcal{U}^T \mathcal{H} \Delta\mathcal{U} + \Delta\mathcal{U}^T \mathcal{G} \xi + \xi^T \mathcal{F} \xi \\ & + \lambda^T (E\Delta\mathcal{U} - f - G\xi) . \end{aligned} \quad (13)$$

A stationary point for the Lagrangian occurs when $\nabla_{\Delta\mathcal{U}} \mathcal{L}(\Delta\mathcal{U}, \lambda) = 0$, hence the corresponding KKT optimality conditions are:

$$\mathcal{H} \Delta\mathcal{U} + \mathcal{G} \xi + E^T \lambda = 0 \quad (14a)$$

$$\lambda \succeq 0, \lambda \in \mathbb{R}^q \quad (14b)$$

$$E\Delta\mathcal{U} - f - G\xi \preceq 0 \quad (14c)$$

$$\text{diag}(\lambda)(E\Delta\mathcal{U} - f - G\xi) = 0 \quad (14d)$$

where q is the minimal number of linear inequalities describing \mathbb{F} . Provided $H \succ 0$ (as is the case when $\mathcal{R} \succ 0$), one can solve for $\Delta\mathcal{U} = -\mathcal{H}^{-1}(\mathcal{G}\xi + E^T\lambda)$ and substitute it back into (14).

¹In [2], $\|\epsilon\|_2^2$ is added to the cost function, together with a weighted l_1 -norm; the l_1 -norm guarantees an exact penalty function and S is an extra tuning weight used to penalise the constraint violations.

A Non-conservative Lower Bound

The condition on the lower bound on ρ over all feasible ξ can now be stated as:

$$\rho > \max_{\xi, \lambda} \|\lambda\|_D \quad (15)$$

with the maximisation subject to the KKT optimality conditions (14) with $\Delta\mathcal{U}$ as above. This is the lowest bound on ρ that guarantees that the soft- and hard-constrained QP problems produce the same solution for all feasible ξ , since all points $(\Delta\mathcal{U}, \lambda)$ which satisfy the KKT conditions for a given ξ solve the corresponding strictly convex primal QP and dual problem. It can be shown that the optimal $\Delta\mathcal{U}^*$ and λ^* are uniquely defined continuous, piecewise affine functions of ξ [6]. The optimisation in (15) is difficult, since it is the maximisation of the norm of a piecewise affine function, which is not necessarily convex or concave.

It is also possible that the maximisation is unbounded. If the region $\Xi_{\mathbb{F}}$ is bounded, then the maximisation is bounded. However, $\Xi_{\mathbb{F}}$ is not necessarily bounded. From this point on, the optimisation is subject to the additional constraint $\xi \in \Xi_0$, where Ξ_0 is a polyhedron of initial conditions which is chosen such that $\Xi_0 \cap \Xi_{\mathbb{F}}$ is bounded. If Ξ_0 is a polytope², then $\Xi_0 \cap \Xi_{\mathbb{F}}$ is also a polytope, hence the maximisation is bounded.

The last constraint in (14) is the complementary slackness condition. Let $\check{\lambda}^j$ and $\tilde{\lambda}^j$ denote the Lagrange multiplier vectors for the j 'th set of inactive and active constraints as in [6]. Let $\check{E}^j, \check{f}^j, \check{G}^j$ and $\tilde{E}^j, \tilde{f}^j, \tilde{G}^j$ be the corresponding matrices extracted from E, f and G . Adopting the above the optimisation in (15) becomes³:

$$\max_{j, \xi, \check{\lambda}^j} \|\check{\lambda}^j\|_D \quad (16a)$$

subject to

$$\tilde{\lambda}^j \succeq 0, \xi \in \Xi_0, j \in \{1, 2, \dots, N\} \quad (16b)$$

$$(\check{E}^j \mathcal{H}^{-1} \mathcal{G} + \check{G}^j) \xi + \check{f}^j \succeq 0 \quad (16c)$$

$$\tilde{\lambda}^j = -(\tilde{E}^j \mathcal{H}^{-1} (\tilde{E}^j)^T)^{-1} (\tilde{f}^j + (\tilde{G}^j + \tilde{E}^j \mathcal{H}^{-1} \mathcal{G}) \xi) \quad (16d)$$

where N is the number of possible active and inactive constraint combinations. The norm $\|\check{\lambda}\|_D = \|\tilde{\lambda}\|_\infty$ is used if $\|\epsilon\|_1$ is used to penalise the constraint violations and $\|\check{\lambda}\|_D = \|\tilde{\lambda}\|_1$ if $\|\epsilon\|_\infty$ is used.

Remark 1. Note that $\tilde{\lambda}^j \succeq 0$ for each combination of active constraints. If $\|\check{\lambda}^j\|_\infty \triangleq \max_i |\check{\lambda}_i^j| =$

²A polytope is a bounded polyhedron.

³As in [6], it is assumed that the rows of \tilde{E}^j are linearly independent in order to guarantee that $(\tilde{E}^j \mathcal{H}^{-1} (\tilde{E}^j)^T)^{-1}$ exists. The fact that $\tilde{\lambda}^j = 0$ allows one to eliminate it from the equations. Note that one can also eliminate $\tilde{\lambda}^j$ from the cost function and constraints.

$\max_i \tilde{\lambda}_i^j$ is used in the maximisation (16a), a sequence of LPs solves $\max_{i,j,\xi,\tilde{\lambda}^j} \tilde{\lambda}_i^j$. Similarly, if $\|\tilde{\lambda}^j\|_1 \triangleq \sum_i |\tilde{\lambda}_i^j| = \sum_i \tilde{\lambda}_i^j$ is used, a sequence of LPs will solve $\max_{j,\xi,\tilde{\lambda}^j} \sum_i \tilde{\lambda}_i^j$.

For large systems with many constraints, this approach might seem computationally impractical, because of the large number of possible combinations of active constraints ($2^q - 1$). In practice, however, far fewer combinations of active constraints can actually occur over the feasible set, e.g. it is not possible for an input or state to be at both its upper and lower bound.

A method for computing the possible active constraint combinations that can occur over $\Xi_0 \cap \Xi_{\mathbb{F}}$ is given in [6]. The authors outline a method where the feasible space is divided into polytopes in which the same constraints on $\Delta\mathcal{U}$ become active at the solution. By solving the above-mentioned LPs over the corresponding polytopes, one can compute a lower bound for the penalty weight.

The authors of [6] also make some comments regarding the computational complexity and maximum number of possible active constraint combinations. However, for off-line design and analysis of the system computation speed is less of an issue. The method outlined here is more efficient than the ‘‘brute force’’ method of gridding and provides a guarantee that the lower bound has been found.

EXAMPLE

This section demonstrates how a soft-constrained mp-QP problem can be designed given a hard-constrained mp-QP. A simple example was chosen, in order that the reader can work out the solutions analytically and visualise the results easily. Consider the following hard-constrained mp-QP:

$$\min_{\theta \in \mathbb{R}} \theta^2 + \theta\xi + \xi^2 \quad (17a)$$

subject to

$$\theta \leq 1 + \xi \quad (17b)$$

$$\theta \geq -1 \quad (17c)$$

where the inequalities describe the feasible set \mathbb{F} . The feasible set for ξ is therefore given by $\Xi_{\mathbb{F}} = \{\xi : \xi \geq -2\}$, i.e. the hard-constrained mp-QP problem is infeasible for $\xi < -2$.

For the soft-constrained problem one can take $\|\lambda\|_D = \|\lambda\|_{\infty}$. If $\rho > \max_{\xi \geq -2} \|\lambda\|_{\infty}$ then the soft-constrained mp-QP⁴

$$\min_{\theta, \epsilon} \theta^2 + \theta\xi + \xi^2 + \rho\|\epsilon\|_1 \quad (18a)$$

⁴Since $\rho\|\epsilon\|_1 = \rho\mathbf{1}^T\epsilon$ if $\epsilon \succeq 0$, problem (18) can be written as a QP.

subject to

$$\theta \leq 1 + \xi + \epsilon_1 \quad (18b)$$

$$\theta \geq -1 - \epsilon_2 \quad (18c)$$

$$\epsilon \succeq 0 \quad (18d)$$

has the same solution as the hard-constrained mp-QP (17) for $\xi \geq -2$.

The first step is to define the regions in which the different combinations of constraints become active, using the KKT conditions (14). Considering all four possible combinations of active and inactive constraints, the analytic expressions for the Lagrange multipliers for all feasible $\xi \geq -2$ are:

$$\lambda = \begin{cases} \begin{bmatrix} -3\xi-2 \\ 0 \end{bmatrix} & \text{if } -2 \leq \xi \leq -\frac{2}{3} \\ \begin{bmatrix} 0 \\ 0 \end{bmatrix} & \text{if } -\frac{2}{3} \leq \xi \leq 2 \\ \begin{bmatrix} 0 \\ \xi-2 \end{bmatrix} & \text{if } \xi \geq 2 \end{cases}$$

The fourth combination, when both constraints are active, occurs only at $\xi = -2$ and this combination is therefore redundant.

The next step is to calculate $\max \|\lambda\|_{\infty}$ for the areas in which constraints are active. For $\xi \geq 2$, $\lambda_2 = \xi - 2$, hence $\max_{\xi \geq 2} \|\lambda\|_{\infty}$ is unbounded. One therefore has to bound ξ from above if ρ is to be finite. Restricting our search to $-2 \leq \xi \leq 4$ gives:

$$\max_{\xi} \|\lambda\|_{\infty} = \begin{cases} 4 & \text{if } -2 \leq \xi \leq -\frac{2}{3} \\ 0 & \text{if } -\frac{2}{3} \leq \xi \leq 2 \\ 2 & \text{if } -2 \leq \xi \leq 4 \end{cases}$$

The lower bound of the violation weight for the soft-constrained mp-QP is therefore

$$\rho > \max_{-2 \leq \xi \leq 4} \|\lambda\|_{\infty} = 4.$$

Choosing $\rho > 4$ guarantees that the soft-constrained mp-QP (18) solution θ_{soft}^* is equal to the solution θ_{hard}^* of the hard-constrained mp-QP (17) for all $-2 \leq \xi \leq 4$.

Figure 1 is a plot of the actual Lagrange multipliers of the hard-constrained mp-QP at the optimal solution as ξ is varied from -2 to 4, confirming that λ is a piecewise affine function of ξ . Figure 2 shows that the difference between the soft-constrained optimal solution and the hard-constrained optimal solution is zero for $\rho > 4$ over the range $-2 \leq \xi \leq 4$. The soft-constrained and hard-constrained solutions will differ for $\xi > 4$, depending on the actual value used for ρ . For $\xi < -2$ the hard-constrained mp-QP does not have a solution, while the soft-constrained mp-QP solution minimises the constraint violations.

CONCLUSIONS

A standard reference tracking formulation of MPC was given. The set of states for which the MPC




Figure 1: Lagrange multipliers of the hard-constrained optimal solution for $-2 \leq \xi \leq 4$




Figure 2: Plot showing that the difference between θ_{soft}^* and θ_{hard}^* for $-2 \leq \xi \leq 4$ is zero if $\rho > 4$

problem is feasible was defined. The importance of ensuring that the feasible set is positively invariant for the nominal closed-loop system was briefly discussed.

It was shown that both the cost function and the constraints of the resulting QP is dependent on the current state, previous control input and reference trajectory. This implies that the Lagrange multipliers are also dependent on these variables.

Exact penalty functions require that the constraint violation weight be larger than the norm of the Lagrange multiplier of the original optimisation problem. It is therefore necessary to compute the upper bound on this norm for all feasible combinations of current states, previous control inputs and reference trajectories.

A method for computing the upper bound on the norm of the Lagrange multipliers over a bounded subset of the feasible states was presented. The region of interest can be divided into polytopes in which different combinations of constraints become active at the solution. The problem of finding the maximum norm of the Lagrange multipliers reduces to solving a finite number of LPs. The maximum norm therefore lies at one of the vertices of these polytopes.

If the constraint violation weight that is used in the soft-constrained cost function is larger than the maximum norm, the solution is guaranteed to be equal to the hard-constrained solution for all feasible conditions that were considered.

FURTHER REMARKS

Thus far, in all examples, the authors have found the norm to be convex over $\Xi_{\mathbb{F}}$ and hence the maximum is obtained at one of the vertices of $\Xi_0 \cap \Xi_{\mathbb{F}}$. This might be related to the fact that the optimal value of the MPC cost function is convex over $\Xi_{\mathbb{F}}$. It might be that the Lagrange multipliers are related to the partial derivatives with respect to ξ of the optimal cost function, since the mp-QP is similar in structure to a perturbed QP when performing a local sensitivity analysis. The authors will appreciate any correspondence which confirms this or suggests otherwise.

ACKNOWLEDGEMENTS

The authors would like to thank Dr Alberto Bemporad, ETH Zürich, for proof-reading a draft of the paper and giving helpful comments and suggestions.

REFERENCES

- [1] N.M.C. de Oliveira and L.T. Biegler. Constraint handling and stability properties of model-predictive control. *AIChE Journal*, 40(7):1138–1155, July 1994.
- [2] P.O.M. Scokaert and J.B. Rawlings. Feasibility issues in model predictive control. *AIChE Journal*, 45(8):1649–1659, August 1999.
- [3] R. Fletcher. *Practical Methods of Optimization*. John Wiley & Sons, 2nd edition, 1987.
- [4] P.O.M. Scokaert and J.B. Rawlings. Infinite horizon linear quadratic control with constraints. In *Proceedings of the 13th Triennial IFAC World Congress*, pages 109–114, San Francisco, USA, 1996. IFAC.
- [5] W.W. Hager. Lipschitz continuity for constrained processes. *SIAM Journal of Control and Optimization*, 17(3):321–338, May 1979.
- [6] A. Bemporad, M. Morari, V. Dua, and E.N. Pistikopoulos. The explicit solution of model predictive control via multiparametric quadratic programming. In *Proceedings of the American Control Conference*, Chicago, USA, June 2000.
- [7] S.S. Keerthi and E.G. Gilbert. Computation of minimum-time feedback control laws for discrete-time systems with state-control constraints. *IEEE Transactions on Automatic Control*, AC-32(5):432–435, 1987.
- [8] F. Blanchini. Set invariance in control. *Automatica*, 35:1747–1767, 1999.
- [9] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789–814, 2000.