

Universitat Autònoma de Barcelona

Robustness aspects of Model Predictive Control

A DISSERTATION SUBMITTED IN
PARTIAL FULFILMENT FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY AT THE
UNIVERSITAT AUTÒNOMA DE BARCELONA

David Megías Jiménez

March 2000

Chapter 4

Robust constrained RHPC using min-max optimisation

4.1 Introduction

As deeply analysed in the previous chapter, the first few robustness in the MPC framework results were obtained for unconstrained, linear predictive controllers in the SISO case. In that situation, an LTI form of these controllers exists, what makes it possible to apply classical robustness analysis and design tools, such as the small gain theorem. However, these early results have serious limitations. To begin with, the assumption that the true plant is an LTI system does not allow non-linear and/or time-varying uncertainty. Hence the results presented above are valid only if non-linearities are mild and time variations are “slow” compared to the true system dynamics. In addition, as highlighted above, one of the most celebrated properties of MPC is constraint handling, and thus it is required that the robustness results be extended to the constrained case. The main aim of this chapter is to develop RHPC controllers to handle these two very significant problems.

To cope with non-linearities and/or time variability, stabilising approaches have been obtained in the non-linear MPC family (Mayne and Michalska, 1990; Michal-

ska and Mayne, 1993; Chen and Allgöwer, 1998b), and an extensive review of these methods can be found in (Chen and Allgöwer, 1998a). However, non-linear model predictive schemes involve an enormous computational burden to compute a single control move. This makes them impractical in many real control situations, especially if fast dynamics occur. As technology provides faster and faster computers, non-linear MPC applications will become more common, and it is likely that they completely replace linear model-based schemes wherever non-linearities cannot be neglected. However, the current technology and numerical algorithms do not allow for a widespread use of these new methods yet, and linear MPC are still the most suitable tool to handle many typical situations. In short, although very detailed non-linear models of the system are often available, it is usually required to resort to linear MPC controllers for computational reasons. This alternative requires the incorporation of some uncertainty description within the system model to be used by the control system.

In addition, it is worth pointing out that the RS conditions established in the previous chapter are restricted to the unconstrained case. This is indeed a serious limitation since, as already remarked, constraint handling is one of the most appealing properties of predictive controllers. Newer results are available taking into account the problem of robustness in the constrained case and a few relevant ones are highlighted below.

In (Allwright and Papavasiliou, 1991; Papavasiliou and Allwright, 1991; Allwright and Papavasiliou, 1992; Allwright, 1994) robust min-max methods based on *Finite Impulse Response* (FIR) models are presented. Such a min-max problem consists of obtaining the control schedule which minimises the maximum of a cost function as the impulse response of the system ranges over a *polytope* of possible impulse responses. A very efficient solution to this problem based on linear programming is provided in (Allwright, 1994). This solution can handle both time invariant and time varying uncertainty, and guarantees the accomplishment of constraints as far as the true impulse

response coefficients occur within the polytope. This method, though computationally efficient, have some inconveniences, the most relevant of which is that FIR models can only be used to represent (approximately) stable plants, a limitation which also affects the DMC controller of Cutler and Ramaker (1980). In addition, many impulse response coefficients (30–50 or even more) are often required to describe the dynamics of stable plants. The alternative min-max MPC solutions introduced below overcome this limitation.

Polytopic linear model descriptions of plant uncertainty, combined with *Linear Matrix Inequalities* (LMI) optimisation, are used in (Kothare *et al.*, 1996) to design robust min-max controllers which satisfy input, output and state constraint specifications. This control scheme computes each control move minimising the worst case cost function over the convex hull of a polytope of linear models, $\mathcal{G} = \text{Co}\{G_1, G_2, \dots, G_L\}$, subject to a set constraints. The true (unknown) model is assumed to be a linear combination of the polytope vertices:

$$G_0 = \sum_{i=1}^L \lambda_i G_i,$$

where $\lambda_i \geq 0$ for $1 \leq i \leq L$ and $\sum_{i=1}^L \lambda_i = 1$. Henceforth, this solution is referred to as the multi-model approach. If the G_i are taken to be linear state-space models, and if G_0 (possibly time-varying) is assumed to lie in \mathcal{G} , this approach can even handle (possibly in an overly conservative way) non-linear systems for which the Jacobian is known to lie in the polytope. The control optimisation problem can be formulated as

$$\Delta \mathbf{u}^{\text{opt}}(t) = \arg \min_{\Delta \mathbf{u}} \max_{G \in \mathcal{G}} J(t),$$

where $\Delta \mathbf{u}$ is a vector of control moves¹ and $J(t)$ is the cost (or objective) function. To account for non-linearities and/or time variability, the worst case must be evaluated letting the plant vary within the polytope at each sampling instant in the future.

¹In the formulation of (Kothare *et al.*, 1996), the control actions (\mathbf{u}) replace the control moves ($\Delta \mathbf{u}$), but the formulae involved are similar.

Addressed in that fashion, the min-max problem usually becomes intractable, but an efficient upper bound solution for state-space models using LMI optimisation is depicted in (Kothare *et al.*, 1996).

The methods of Kothare *et al.* (1996) have a few limitations. To begin with, if a constant setpoint tracking problem is considered, the results apply only to uncertain linear time-invariant systems (and thus are not valid for non-linear systems). In addition, these methods are formulated for state-space descriptions of the model and state feedback. It is a difficult task to apply this approach in controllers which use input/output models, since “pole” uncertainty cannot be expressed. On the other hand, “delay”, “gain” and “zero” uncertainties can be handled without much difficulty. When the (*integer*) dead-time can vary, it is enough to include all the possible delay values in the models (vertices) of the polytope. In Section A.3 of Appendix A the use of this representation when gain and zero uncertainties occur is presented. In addition, as shown in Section A.3, the plant family must often be extended to allow this description, leading to overly conservative results since systems which are not included in the initial delay/gain/zero uncertainty specifications must be taken into account. Moreover, this parametrisation cannot be used when uncertainty affects the system’s poles, which are, indeed, the most significant parameters of transfer function models as the open-loop dynamics are concerned.

Furthermore, several kinds of strong non-linearities, such as saturation, hysteresis, relay or dead-zone, can occur in a plant. In such a case, a polytope of linear models, as the ones used in (Kothare *et al.*, 1996), does not provide with an appropriate description, because a linearisation can be hard (or impossible) to obtain since the non-linear functions involved in the “true” system equations are not differentiable. Another remarkable issue is the influence of disturbances. Unmeasurable perturbations often occur in such a way that they are not included within the polytopic description, leading to inaccurate predictions and, possibly, to constraint violations. Obviously, this

potential hazard must be handled somehow. As a solution, constraints which are more stringent than necessary might be used, but this possibility can give rise to negative economical implications since the optimal operating point is nearly always close to (or on) the constraint boundary. For example, a chemical may need be produced with at least a 97% purity. Then, from an economical point of view, the most convenient setpoint would be to produce exactly at 97% purity. However, if the actual purity achieved by the process is of (say) only 96.5%, the chemical may have to be rejected, with the subsequent serious economical loss.

A *global uncertainty* description can be the solution of many practical control problems. Min-max algorithms, either for state-space (Scokaert and Mayne, 1998) or transfer function (Camacho and Bordóns, 1995) models, can be developed using this formulation, and it is even possible to write them as an efficiently solvable LP problem (Camacho and Bordóns, 1995) if 1-norm cost functions are used. A global uncertainty is an unknown (bounded) signal $\theta(t)$ which, added to the model, produces the true plant. This very simple concept is general enough to range over linear and non-linear, time varying and time invariant, stable and unstable uncertainty, which can affect the poles, the zeroes, the gain, the delay, or whatever parameter of the model. In addition, it perfectly describes (unmeasurable) disturbances. However, a global uncertainty parameter can involve some degree of conservativeness, since situations which are worse than those described by the specifications are often taken into account. A band updating procedure is presented in this chapter to “adapt” the min-max algorithm to the current uncertainty values. This procedure is intended to reduce this cautiousness, since the initial conservative settings of the uncertainty bands can be replaced, according to the uncertainty measurements, by tighter values.

The scope of this chapter is to formulate and test predictive control schemes based on a global uncertainty description, and to compare them to other robustness enhancing tools. The methods introduced below are requested to provide nominal stability,

and thus the basis of the controllers developed below are the controllers presented in Chapter 2 which guarantee the stability of the nominal closed-loop system, or are “very unlikely” to provide an unstable closed-loop system, such as the $QGPC_1^\infty$. Another requirement is to preserve stability in the presence of global uncertainty (as far as the uncertainty signal does not become unbounded). Last, but not least, the min-max controllers formulated below are requested to satisfy the constraints in spite of uncertainty. In the forthcoming sections, unless otherwise explicitly specified, the “min-max” prefix is used for methods based on the global uncertainty approach and not on the polytopic multi-model linear plant description.

This chapter is organised as follows. In Section 4.2, the global uncertainty description is introduced, and input/output models which incorporate this kind of uncertainty are provided. The solution of the min-max optimisation problem is tackled from both analytical and numerical points of view. In Section 4.3, the $QGPC_1^\infty$ presented in Chapter 2 is converted to a min-max formulation. This min-max controller can be implemented as a LP problem with *exactly the same number* of constraints as its precursor ($QGPC_1^\infty$), despite the min-max formulation. Section 4.4 presents a set of simulation results performed on linear and non-linear plants, and compares the effectiveness and efficiency of the newly formulated min-max controllers against various control strategies. Section 4.5 provides an analysis of the robustness of min-max MPC and the classical T -based controllers based on the statistical learning theory and Monte Carlo simulation. Finally, Section 4.6 finishes the chapter summarising the most significant concluding remarks.

The formulae provided in the sequel are written for the SISO case only for simplicity of notation, but all of them can be readily extended to the MIMO case in a straightforward manner.

4.2 Global uncertainty and min-max optimisation

As defined in (Camacho and Bordóns, 1995), a global uncertainty is an unknown (but bounded) signal $\theta(t)$ which adds up to the model output to produce the true system output. Thus, all sources of uncertainty are collected into the single parameter $\theta(t)$. In this very simple fashion, all kinds of uncertainties, namely linear or non-linear, time invariant or time varying, modelling errors and non-measurable disturbances, can be represented.

In the input/output domain, the global uncertainty $\theta(t)$ can be incorporated into a *Controlled Auto-Regressive Moving Average* (CARMA) model as follows:

$$A(q^{-1})y(t) = B(q^{-1})u(t-1) + \theta(t), \quad (4.1)$$

where A and B are known polynomials in the backward shift operator (defined in eqn.2.2) and $\theta(t)$ is the global uncertainty signal, which is assumed to be bounded by θ^- and θ^+ : $\theta^- \leq \theta(t) \leq \theta^+$, henceforth referred to uncertainty bounds/bands/limits.

The description of eqn.4.1 can be used if (bounded) step-like disturbances affect the output, but the occasional drift produced by disturbances cannot be handled by this formulation. To account for drift disturbances, the global uncertainty signal can be integrated, resulting in the CARIMA model

$$A(q^{-1})y(t) = B(q^{-1})u(t-1) + \frac{\theta(t)}{\Delta}. \quad (4.2)$$

In the sequel the CARMA model of eqn.4.1 is used, since the methods to be introduced consider (asymptotically) constant disturbances. Similar algorithms can be easily developed for drift disturbances using the model provided in eqn.4.2.

Unfortunately, the polynomial T , which can be used to improve the tune the cautiousness/performance trade-off of GPC-like predictive controllers, as shown in Chapter 3, is dropped in this formulation. The robustness features of min-max controllers must

then rely on the optimisation procedure, which should provide with some tuning knobs to recover some of the degrees of freedom which are associated to T .

4.2.1 Output predictions

For the system of eqn.4.1, output predictions $y(t+j|t)$ can be computed as

$$\begin{aligned} y(t+j|t) = & -a_1\hat{y}(t+j-1|t) - \dots - a_{n_a}\hat{y}(t+j-n_a|t) \\ & + b_1u(t+j-1|t) + \dots + b_{n_b}u(t+j-n_b|t) \\ & + \theta(t+j|t) + h_1\theta(t+j-1|t) + \dots + h_{j-1}\theta(t+1|t), \end{aligned} \quad (4.3)$$

where $\hat{y}(t+j|t)$ are predictions performed taking all the future global uncertainties to be zero, *i.e.* $\theta(t+j|t) = 0$ for all $j > 0$, and h_j denotes the j^{th} coefficient of the impulse response of the system $1/A$. Note that, with these definitions, $\hat{y}(t+k|t) = y(t+k|t)$ for all $k \leq 0$, since these are past output values which are not affected by the future uncertainties.

Now, the output predictions can be written as

$$\begin{aligned} y(t+j|t) = & f(t+j|t) + g_1\Delta u(t+j-1|t) + g_2\Delta u(t+j-2|t) + g_j\Delta u(t|t) \\ & + \theta(t+j|t) + h_1\theta(t+j-1|t) + \dots + h_{j-1}\theta(t+1|t), \end{aligned} \quad (4.4)$$

where $f(t+j|t)$ are free response predictions which are computed taking *both* the future control moves *and the future global uncertainty predictions* to be zero, *i.e.* $\Delta u(t+k|t) = 0$ and $\theta(t+k|t) = 0$ for all $k \geq 0$. Since the output predictions are an affine function of the global uncertainty values, the extreme (maximum and minimum) predictions of $y(t+j|t)$ occur for extreme values of $\theta(t+k|t)$, *i.e.* $\theta(t+k|t) = \theta^+$ or $\theta(t+k|t) = \theta^-$ for all $1 \leq k \leq j$.

4.2.1.1 Output predictions for GPC^∞ and QGPC_1^∞

The definitions given in Section 2.2.2 for the vectors $\Delta u(t)$, $y(t)$ and $f(t)$ and the dynamic matrix G are assumed to hold in this section. Now, the *predicted uncertainty vector* is defined as

$$\theta^{N_\theta}(t) = [\theta(t+1|t) \quad \theta(t+2|t) \quad \dots \quad \theta(t+N_\theta|t)]^T, \quad (4.5)$$

where N_θ is the *uncertainty prediction horizon*.

The prediction equation (eqn.4.3) allows to arrange the vector of predictions $y(t)$ as

$$y = f + G\Delta u + H_\theta \theta^N, \quad (4.6)$$

where the dynamic matrix H_θ is given by

$$H_\theta = \begin{bmatrix} 1 & 0 & \dots & 0 \\ h_1 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ h_{N-1} & h_{N-2} & \dots & 1 \end{bmatrix},$$

such that h_j is the j^{th} coefficient of the impulse response of the system $1/A$. Notice that the output predictions are affine in both Δu and θ^N .

It is also possible to define predictions on the unstable part of the output. Let \bar{A} and \tilde{A} be, respectively, the strictly stable and the unstable *monic* factors of A . The unstable part of the output of eqn.4.1 can be obtained as

$$\begin{aligned} \tilde{y}(t) &= \frac{B(q^{-1})}{\tilde{A}(q^{-1})} u(t-1) + \frac{1}{\tilde{A}(q^{-1})} \theta(t) \\ &= \bar{A} y(t). \end{aligned} \quad (4.7)$$

Now, if the definitions introduced in Section 2.2.2 for the vectors $\tilde{y}(t)$ and $\tilde{w}(t)$ and the dynamic matrix \tilde{G} are used, the predictions \tilde{y} can be arranged as

$$\tilde{y} = \tilde{f} + \tilde{G}\Delta u + \tilde{H}_\theta \theta^{N+n_{\tilde{a}}}, \quad (4.8)$$

for the dynamic matrix $\widetilde{\mathbf{H}}_\theta$ is defined as

$$\widetilde{\mathbf{H}}_\theta = \begin{bmatrix} \tilde{h}_{N-1} & \tilde{h}_{N-2} & \dots & 0 \\ \tilde{h}_N & \tilde{h}_{N-1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{h}_{N+n_a-1} & \tilde{h}_{N+n_a-2} & \dots & 1 \end{bmatrix},$$

where \tilde{h}_k is the k^{th} impulse response coefficient of the system $1/\tilde{A}(q^{-1})$, and the free response predictions on of the unstable part of the output $\tilde{f}(t+j|t)$ are computed taking both the future control moves and the future global uncertainties to be zero. Notice that the predictions $\tilde{\mathbf{y}}$ are affine in $\Delta \mathbf{u}$ and $\boldsymbol{\theta}^{N+n_a}$.

4.2.1.2 Output predictions for finite horizon controllers

The definitions for the vectors $\Delta \mathbf{u}(t)$, $\mathbf{y}_1(t)$, $\mathbf{y}_2(t)$, $\mathbf{w}_1(t)$, $\mathbf{w}_2(t)$ and the dynamic matrices \mathbf{G}_1 and \mathbf{G}_2 introduced in Section 2.2.1 are used here. Now let the future predicted uncertainty vectors $\boldsymbol{\theta}_1(t)$ and $\boldsymbol{\theta}_2(t)$ be specified as

$$\begin{aligned} \boldsymbol{\theta}_1(t) &= \boldsymbol{\theta}^{N_y-1}(t), \\ \boldsymbol{\theta}_2(t) &= \boldsymbol{\theta}^{N_2}(t), \end{aligned}$$

where $\boldsymbol{\theta}^{N_y-1}(t)$ and $\boldsymbol{\theta}^{N_2}(t)$ can be obtained using eqn.4.5 with $N_\theta = N_y-1$ and $N_\theta = N_2$ respectively. This notation allows to arrange the output prediction vectors $\mathbf{y}_1(t)$ and $\mathbf{y}_2(t)$ as

$$\begin{aligned} \mathbf{y}_1 &= \mathbf{f}_1 + \mathbf{G}_1 \Delta \mathbf{u} + \mathbf{H}_{\theta_1} \boldsymbol{\theta}_1, \\ \mathbf{y}_2 &= \mathbf{f}_2 + \mathbf{G}_2 \Delta \mathbf{u} + \mathbf{H}_{\theta_2} \boldsymbol{\theta}_2, \end{aligned} \tag{4.9}$$

where the dynamic matrices \mathbf{H}_{θ_1} and \mathbf{H}_{θ_2} are given by

$$\begin{aligned} \mathbf{H}_{\theta_1} &= \begin{bmatrix} h_{N_1-1} & h_{N_1-2} & \dots & 0 \\ h_{N_1} & h_{N_1-1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ h_{N_y-2} & h_{N_y-1} & \dots & 1 \end{bmatrix}, \\ \mathbf{H}_{\theta_2} &= \begin{bmatrix} h_{N_y-1} & h_{N_y-2} & \dots & 0 \\ h_{N_y} & h_{N_y-1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ h_{N_2-1} & h_{N_2-2} & \dots & 1 \end{bmatrix}, \end{aligned}$$

such that h_j is the j^{th} coefficient of the impulse response of the system $1/A$, and $h_0 = 1$ since A is assumed monic ($a_0 = 1$).

4.2.2 Min-max optimisation

An optimal control move vector can be obtained by solving the problem

$$\overline{\Delta \mathbf{u}^{\text{opt}}} = \arg \min_{\Delta \mathbf{u}} \max_{\boldsymbol{\theta}^{N_\theta} \in \Theta^{N_\theta}} J(t), \quad (4.10)$$

subject to the equality/inequality constraints associated to the controller. $J(t)$ is the cost function and Θ^{N_θ} is the convex hull of an *uncertainty polytope*, formed by 2^{N_θ} vertices for all possible combinations of extreme values of $\boldsymbol{\theta}^{N_\theta}(t)$ along the horizon N_θ , *i.e.*

$$\Theta^{N_\theta} = \text{Co} \left\{ \boldsymbol{\theta}_1^{N_\theta}, \boldsymbol{\theta}_2^{N_\theta}, \dots, \boldsymbol{\theta}_{2^{N_\theta}}^{N_\theta} \right\}$$

such that each component of $\boldsymbol{\theta}_i^{N_\theta}$ is either θ^+ or θ^- and $\boldsymbol{\theta}_i^{N_\theta} \neq \boldsymbol{\theta}_j^{N_\theta}$ if $i \neq j$. Notice that there are exactly 2^{N_θ} vertices $\boldsymbol{\theta}_i^{N_\theta}$. All possible global uncertainty predictions are a linear combination of the polytope vertices:

$$\boldsymbol{\theta}^{N_\theta} = \sum_{i=1}^{2^{N_\theta}} \lambda_i \boldsymbol{\theta}_i^{N_\theta}, \quad (4.11)$$

where $\lambda_i \geq 0$ for $1 \leq i \leq 2^{N_\theta}$ and $\sum_{i=1}^{2^{N_\theta}} \lambda_i = 1$. The convex hull Θ^{N_θ} is formed by infinitely many vectors $\boldsymbol{\theta}^{N_\theta}$ satisfying this definition, or

$$\Theta^{N_\theta} = \left\{ \boldsymbol{\theta}^{N_\theta} = \sum_{i=1}^{2^{N_\theta}} \lambda_i \boldsymbol{\theta}_i^{N_\theta} : \lambda_i \geq 0 \text{ for } 1 \leq i \leq 2^{N_\theta}, \sum_{i=1}^{2^{N_\theta}} \lambda_i = 1 \right\}.$$

As an example, a three-dimensional polytope is shown in Fig.4.1. In such a case the polytope vertices form a cube, and any point within the cube or on its surface is a potential future uncertainty vector.




Figure 4.1: Uncertainty polytope for $N_\theta = 3$

Any cost function can be used for $J(t)$ in this min-max formulation. The cost functions defined in Chapter 2 lead to the min-max GPC for $J_2(t)$ defined in eqn.2.4, the min-max CRHPC for $J_2(t)$ defined in eqn.2.5, the min-max GPC^∞ for $J_2(t)$ defined in eqn.2.6, the min-max 1-norm controllers for $J_1(t)$ as defined in eqn.2.34, and the min-max QGPC_1^∞ for $J_1(t)$ as defined in eqn.2.48. The formulae associated to the min-max (finite horizon) GPC and the GPC_1 can be found in (Camacho and Bordóns, 1995) and are omitted in the sequel.

In general, equality and inequality constraints must be taken into account in the min-max problem posed above. As discussed in Chapter 2, the former are required by methods with stability guarantees (*e.g.* the CRHPC and the GPC^∞) and must be handled by min-max controllers, whereas the latter are designed by the user to define closed-loop specifications or performance. The following section tackles the problem of constraint incorporation into the min-max optimisation framework.

4.2.2.1 Inequality constraints

As remarked in Chapter 2, the general inequality constraints can be written as

$$P\Delta u \leq r,$$

for some matrix P and vector r . This set of constraints can be divided into two separate parts:

$$P_u\Delta u \leq r_u,$$

which are independent of uncertainty (constraints on the control signal, control rate or control acceleration) and

$$P_\theta\Delta u \leq r_\theta,$$

which depend on the global uncertainty predictions (constraints on the output signal, output rate, output acceleration or internal states). As illustrated below, the matrix P_θ is independent of the global uncertainty predictions θ^{N_θ} , and the vectors r_θ are affine functions in θ^{N_θ} . Hence the extreme (most restrictive) constraints occur at the vertices of the polytope Θ^{N_θ} . Therefore these constraints need *only* be considered *at the polytope vertices*:

$$P_\theta\Delta u \leq r_i, \forall \theta_i^{N_\theta} \in \Theta^{N_\theta}. \quad (4.12)$$

For instance, consider the GPC[∞] with the output constraints $y(t+j|t) \leq y^+$ for $1 \leq j \leq N$. Using the prediction equations obtained above, these constraints can be arranged in vector form as

$$f + G\Delta u + H_\theta\theta_i^N \leq y^+\mathbf{1}, \forall \theta_i^N \in \Theta^N,$$

where $\mathbf{1}$ denotes the ones vector of appropriate dimension. Notice that the uncertainty prediction vectors θ^N are N -dimensional, and thus the polytope Θ^N consists of 2^N

vertices. Now, the inequality provided above can be written as

$$\mathbf{G}\Delta\mathbf{u} \leq \mathbf{y}^+ \mathbf{1} - \mathbf{f} - \mathbf{H}_\theta \boldsymbol{\theta}_i^N, \forall \boldsymbol{\theta}_i^N \in \Theta^N, \quad (4.13)$$

which specify $N2^N$ constraints (rows) to be included in the min-max optimisation problem, since the polytope Θ^N consists of 2^N vertices $\boldsymbol{\theta}_i^N$. However, note that all these constraints differ only in the term $\mathbf{H}_\theta \boldsymbol{\theta}_i^N$, hence it is possible to examine these vectors component by component and pick up the value which produces the more restrictive constraint. All the other constraints can be discarded since they are redundant. As an example, the set of constraints

$$\begin{aligned} x &\leq z + 3, \\ x &\leq z - 7, \\ x &\leq z + 1, \\ x &\leq z - 3, \end{aligned}$$

can be reduced simply to $x \leq z - 7$, and the rest are redundant. If this procedure is applied to all the rows of eqn.4.13, then it is possible to reduce the $N2^N$ constraints of to just N :

$$\mathbf{G}\Delta\mathbf{u} \leq \mathbf{y}^+ \mathbf{1} - \mathbf{f} - \widehat{\boldsymbol{\theta}}^N,$$

where $\widehat{\boldsymbol{\theta}}^N$ is formed taking the smallest components of the vectors $\mathbf{H}_\theta \boldsymbol{\theta}_i^N$. Notice that the right-hand side of the finally chosen N constraints include elements associated to different vertices.

In general, the $\dim(\mathbf{r}_i)2^N$ constraints of eqn.4.12 can be reduced to just $\dim(\mathbf{r}_i)$ using this procedure, where $\dim(\mathbf{v})$ denotes the dimension (or size) of the vector \mathbf{v} . Finally, these constraints can be written in the (standard) form

$$\mathbf{P}_\theta \Delta\mathbf{u} \leq \widehat{\mathbf{r}}_\theta,$$

with $\mathbf{P}_\theta = \mathbf{G}$ and $\widehat{\mathbf{r}}_\theta = \mathbf{y}^+ \mathbf{1} - \mathbf{f} - \widehat{\boldsymbol{\theta}}^N$,

Remark 4.1 It is worth pointing out that the constraint reduction procedure is a *Non-deterministic Polynomial* (NP) problem, since 2^{N_θ} vectors $\mathbf{H}_\theta \boldsymbol{\theta}_i^{N_\theta}$ must be examined so as to choose the most restrictive constraints. $\square\square\square$

4.2.2.2 Equality constraints

End-point equality constraints are common to the stabilising MPC controllers described in Chapter 2. In fact, these constraints are essential for the nominal stability proofs, and must be considered by the min-max controllers to preserve stability at least in the undisturbed case.

Prior to undertake the incorporation of these constraints into the min-max optimisation framework, the uncertainty prediction horizon N_θ must be defined as $N_2 = N + m$ in the CRHPC and by $N + n_{\tilde{a}}$ in the GPC $^\infty$. These definitions of N_θ include the predictions in the constraint window as well as the costing horizon.

For the GPC $^\infty$ (GPC $_1^\infty$ and QGPC $_1^\infty$) the equality constraints are as defined in eqn.2.20:

$$\tilde{\mathbf{y}}(t) = \tilde{\mathbf{w}}(t),$$

which, using the predictions in eqn.4.8, become

$$\tilde{\mathbf{G}}\Delta\mathbf{u} = \tilde{\mathbf{w}} - \tilde{\mathbf{f}} - \tilde{\mathbf{H}}_\theta \boldsymbol{\theta}^{N+n_{\tilde{a}}}.$$

For the CRHPC (CRHPC $_1$) the equality constraints must be enforced on the *whole* model output (not only the unstable part), as described in Section 2.2, and are provided by

$$\mathbf{y}_2(t) = \mathbf{w}_2(t),$$

which can be rearranged, using eqn.4.9, as

$$\mathbf{G}_2\Delta\mathbf{u} = \mathbf{w}_2 - \mathbf{f}_2 - \mathbf{H}_{\theta 2}\boldsymbol{\theta}^{N_2}.$$

Be that as it may, the end-point equality constraints should, in principle, be considered for all the infinitely many possible global uncertainties $\theta^{N_\theta} \in \Theta^{N_\theta}$ (not only the vertices $\theta_i^{N_\theta}$), *i.e.*, all the possible global uncertainty values described in eqn.4.11. With only a finite number (N_u) of decision variables, it is obviously impossible to enforce infinitely many equality constraints. To overcome this difficulty, here it is suggested to impose the equality constraints assuming that *all* the future global uncertainties occur at the middle point value, *i.e.*

$$\theta(t+j|t) = \bar{\theta} = \frac{\theta^- + \theta^+}{2},$$

which can be thought of as a sort of minimum variance estimate. Then, the vector $\bar{\theta}^{N_\theta}$ is defined as

$$\bar{\theta}^{N_\theta} = \left[\bar{\theta} \quad (N_\theta) \quad \bar{\theta} \right]^T,$$

the geometrical interpretation of which is nothing but the centre of the uncertainty polytope. Enforced for this so-called “average” vector, the end-point equality constraints become

$$\tilde{G}\Delta u = \tilde{w} - \tilde{f} - \tilde{H}_\theta \bar{\theta}^{N_\theta},$$

for the GPC[∞] and

$$G_2\Delta u = w_2 - f_2 - H_{\theta_2} \bar{\theta}^{N_2},$$

for the CRHPC.

Notice that if $\theta^+ = \theta^- = \theta = 0$ (undisturbed case) the constraints defined above are exactly those which ensure (nominal) stability. In fact, if the global uncertainty settles down to some steady-state value θ^* (as usually occurs for step-like disturbances) and the upper and lower bounds θ^- and θ^+ are both updated to equal θ^* , then $\bar{\theta} = \theta^*$ and the predictions become exact. In such a case, these end-point equality constraints provide both stability and offset-free setpoint tracking (Sjokaert, 1997).

4.2.2.3 2-norm cost functions

This section uses the GPC^∞ formulation provided in Chapter 2. Analogous results would be obtained for the CRHPC (or the more general formulation given in eqn.2.3), which are omitted here for brevity.

If the 2-norm is used, the GPC^∞ cost function to be optimised can be written in the form:

$$J_2(t) = (\mathbf{w} - \mathbf{f} - \mathbf{G}\Delta\mathbf{u} - \mathbf{H}_\theta\boldsymbol{\theta}^{N_\theta})^\top \boldsymbol{\Lambda} (\mathbf{w} - \mathbf{f} - \mathbf{G}\Delta\mathbf{u} - \mathbf{H}_\theta\boldsymbol{\theta}^{N_\theta}) + \Delta\mathbf{u}^\top \mathbf{R}\Delta\mathbf{u},$$

subject to equality ($\tilde{\mathbf{y}} = \tilde{\mathbf{w}}$) and possibly to inequality constraints, with the weighting matrices $\boldsymbol{\Lambda}$ and \mathbf{R} defined in Chapter 2.

Remark 4.2 The prediction horizon for the global uncertainty can be reduced to $N_\theta = N$, since the values beyond this point do not affect the predictions $y(t+j|t)$ for $1 \leq j \leq N$. Longer uncertainty horizons can be considered for the equality ($N_\theta = N + n_{\bar{a}}$) or the inequality constraints (if these are enforced beyond N), but the shape of the cost function is only affected by $\theta(t+j|t)$ for $1 \leq j \leq N$, and thus $N_\theta = N$ is assumed in the rest of this section. □□□

This cost function is quadratic in $\Delta\mathbf{u}$, and can be posed in the standard form:

$$J_2(t) = \Delta\mathbf{u}^\top \mathbf{A}_{\Delta\mathbf{u}} \Delta\mathbf{u} + \mathbf{b}_{\Delta\mathbf{u}}^\top \Delta\mathbf{u} + c_{\Delta\mathbf{u}},$$

for

$$\mathbf{A}_{\Delta\mathbf{u}} = \mathbf{G}^\top \boldsymbol{\Lambda} \mathbf{G} + \mathbf{R},$$

$$\mathbf{b}_{\Delta\mathbf{u}}^\top = 2(\mathbf{w} - \mathbf{f} - \mathbf{H}_\theta\boldsymbol{\theta}^{N_\theta})^\top \boldsymbol{\Lambda} \mathbf{G},$$

$$c_{\Delta\mathbf{u}} = (\mathbf{w} - \mathbf{f} - \mathbf{H}_\theta\boldsymbol{\theta}^{N_\theta})^\top \boldsymbol{\Lambda} (\mathbf{w} - \mathbf{f} - \mathbf{H}_\theta\boldsymbol{\theta}^{N_\theta}),$$

but it is also quadratic in $\boldsymbol{\theta}^{N_\theta}$:

$$J_2(t) = \boldsymbol{\theta}^{N_\theta \top} \mathbf{A}_\theta \boldsymbol{\theta}^{N_\theta} + \mathbf{b}_\theta^\top \boldsymbol{\theta}^{N_\theta} + c_\theta,$$

for

$$\begin{aligned} \mathbf{A}_\theta &= \mathbf{H}_\theta^\top \Lambda \mathbf{H}_\theta, \\ \mathbf{b}_\theta^\top &= 2(\mathbf{w} - \mathbf{f} - \mathbf{G}\Delta\mathbf{u})^\top \Lambda \mathbf{H}_\theta, \\ c_\theta &= (\mathbf{w} - \mathbf{f} - \mathbf{G}\Delta\mathbf{u})^\top \Lambda (\mathbf{w} - \mathbf{f} - \mathbf{G}\Delta\mathbf{u}) + \Delta\mathbf{u}^\top \mathbf{R}\Delta\mathbf{u}. \end{aligned}$$

Hence, due to the convexity of $J_2(t)$, the maximum of $J_2(t)$ as θ^{N_θ} ranges over the polytope Θ^{N_θ} necessarily occurs at one of the vertices of Θ^{N_θ} (Bazaraa and Shetty, 1979; Camacho and Bordóns, 1995). Then the optimisation problem reduces to consider 2^{N_θ} quadratic cost functions, one for each vertex $\theta_i^{N_\theta}$. This is true for any *convex* cost function (e.g. a 1-norm cost function).




Figure 4.2: Min-max solution for a pair of quadratic functions

An analytical solution of the min-max problem for quadratic cost functions is described next. Let $J_2^i(t)$ denote the quadratic cost function obtained for the polytope vertex $\theta_i^{N_\theta} \in \Theta^{N_\theta}$. Now consider a pair of cost functions $J_2^i(t)$ or $J_2^j(t)$ (for $N_u = 1$) as shown in Fig.4.2, obtained at the vertices $\theta_i^{N_\theta}$ and $\theta_j^{N_\theta}$ respectively. The solution of the subproblem

$$\Delta\mathbf{u}^*(t) = \arg \min_{\Delta\mathbf{u}} \max_{\theta^{N_\theta} \in \{\theta_i^{N_\theta}, \theta_j^{N_\theta}\}} J_2(t)$$

lies either on the intersection of the cost functions (a) or on one of the minima of $J_2^i(t)$ and $J_2^j(t)$ (b). Since the matrix $\mathbf{A}_{\Delta\mathbf{u}}$ is independent of θ^{N_θ} , the intersection

problem can be solved as the minimisation of either $J_2^i(t)$ or $J_2^j(t)$ subject to the equality constraint $J_2^i(t) = J_2^j(t)$, which can be written as

$$\mathbf{b}_{\Delta \mathbf{u}}^T(\boldsymbol{\theta}_i^{N_\theta}) \Delta \mathbf{u} + c_{\Delta \mathbf{u}}(\boldsymbol{\theta}_i^{N_\theta}) = \mathbf{b}_{\Delta \mathbf{u}}^T(\boldsymbol{\theta}_j^{N_\theta}) \Delta \mathbf{u} + c_{\Delta \mathbf{u}}(\boldsymbol{\theta}_j^{N_\theta})$$

or

$$\left[\mathbf{b}_{\Delta \mathbf{u}}(\boldsymbol{\theta}_i^{N_\theta}) - \mathbf{b}_{\Delta \mathbf{u}}(\boldsymbol{\theta}_j^{N_\theta}) \right]^T \Delta \mathbf{u} = c_{\Delta \mathbf{u}}(\boldsymbol{\theta}_j^{N_\theta}) - c_{\Delta \mathbf{u}}(\boldsymbol{\theta}_i^{N_\theta}), \quad (4.14)$$

apart from the (equality and inequality) constraints associated to the controller.

In the general case of dimension $N_u > 1$, the global solution lies either on any of the minima of $J_2^i(t)$ or in the intersection of up to $N_u + 1$ cost functions $J_2^i(t)$.

Remark 4.3 Each intersection defines an equality constraint in the minimisation, and each equality constraint reduces in one the degrees of freedom available for optimisation *i.e.* the dimension of the decision variables space. In other words, 2 cost functions can intersect producing a $(N_u - 1)$ -dimensional decision space, and $N_u + 1$ cost functions can intersect on a single point (0-dimensional space), but $N_u + 2$ cost functions do not (in general) intersect on a N_u -dimensional space. □□□

Remark 4.4 Actually, the equality constraints associated to the controller must be taken into account, since they reduce the dimension of the decision variables space. Thus, if the controller involves c equality constraints, the number of cost functions which intersect reduces to $N_u - c + 1$. □□□

A set of subproblems can be obtained forming all the combinations of up to $N_u - c + 1$ cost functions, and each of these can be solved by means of a QP problem subject to up to $N_u - c + 1$ equality constraints of the form of eqn.4.14 plus l equality constraints associated to the controller. Note that there are exactly

$$\binom{2^{N_\theta}}{0} + \binom{2^{N_\theta}}{1} + \cdots + \binom{2^{N_\theta}}{N_u - c + 1},$$

subproblems. Let $(J_2^*, \Delta \mathbf{u}^*)$ denote the solution of one of these subproblems. If $(J_2^*, \Delta \mathbf{u}^*)$ is such that any cost function value $J_2^k(\Delta \mathbf{u}^*)$ is lower than or equal to J_2^* for all $1 \leq k \leq N_\theta$, then $\Delta \mathbf{u}^*$ turns out to be a candidate global solution of the min-max problem. The global optimum is finally found as the one which provides the minimal J_2^* over all the candidate solutions. The following algorithm implements this analytical min-max optimisation method:

1. for all $1 \leq i \leq 2^{N_\theta}$ do
 - (a) Form all the combinations $\{J_2^{j_1}, \dots, J_2^{j_p}\}$ with up to $N_u - c$ ($p \leq N_u - c$) cost functions, including the empty one ($p = 0$), and with $j_k > i$ for all $1 \leq k \leq p$
 - (b) for all such combinations $\{J_2^{j_1}, \dots, J_2^{j_p}\}$ do
 - (i) $(\Delta \mathbf{u}^*, J_2^*) := (\arg \min_{\Delta \mathbf{u}} J_2^i(t), \min_{\Delta \mathbf{u}} J_2^i(t))$ subject to $J_2^i(t) = J_2^{j_k}(t)$ for all $1 \leq k \leq p$, and subject to the controller equality/inequality constraints
 - (ii) if $J_2^* \geq J_2^k(\Delta \mathbf{u}^*)$ for all $1 \leq k \leq 2^{N_\theta}$ then add $(\Delta \mathbf{u}^*, J_2^*)$ to the candidate list
- endfor
- endfor
2. Pick up the minimum from the candidate list:

$$\left(\overline{\Delta \mathbf{u}^{\text{opt}}}, \overline{J_2^{\text{opt}}}\right) := \min_{J_2^*} \{(\Delta \mathbf{u}^*, J_2^*)\};$$

return

Obviously this algorithm is NP (the number of iterations depends on 2^{N_θ}), and hence should only be used for small N_θ . In addition, it must be taken into account that this method is to be applied *on-line* at each sampling instant and thus the computational burden must be small to allow for the obtention of the optimal control move vector

within a sampling period. As an alternative, a numerical solution can be used to reduce the computational requirements. An equivalent formulation of eqn.4.10 is given by the non-linear programming problem:

$$v^{opt}, \overline{\Delta u^{opt}} = \arg \min_{v, \Delta u} v \text{ subject to } v \geq J_2^i(t) \text{ for } 1 \leq i \leq 2^{N_\theta},$$

and subject to the equality/inequality constraints associated to the controller. This problem can be solved using *Sequential Quadratic Programming* (SQP) methods, as described in (Schittowski, 1985; The Mathworks, 1997).

To compare the analytical and the numerical solutions in terms of computational burden, the algorithm described above and the non-linear programming solution have been tested. A 5-dimensional search space ($N_u = 5$) has been considered. For this experiment, the matrix $A_{\Delta u}$ (definite-positive), the vectors $b_{\Delta u}$ and the scalars $c_{\Delta u}$ have been randomly generated. In addition no constraints (either equality or inequality) have been taken into account.

		Solution		Ratio (t_A/t_N)
		Analytical (t_A)	Numerical (t_N)	
Polytope vertices (2^{N_θ})	2^1	1.0	4.0	0.25
	2^2	3.0	6.8	0.44
	2^3	37.8	6.8	5.60
	2^4	2903.2	8.3	350.84

Table 4.1: Normalised CPU time

The computational burden, measured in normalised CPU time (see Section 2.4.3), is shown in Table 4.1. The outcome of this experiment evidences that an analytical solution is only advisable for $N_\theta < 3$, since the CPU time it requires increases dramatically with N_θ . In many practical situations, the SQP solution should thus be preferred. Notice that the rate of increase of the computational burden with N_θ associated to the numerical solution is much lower than the result obtained with the analytical method.

However, it is worth pointing out that the example presented above is quite an ideal case, since no constraints have been taken into account. For the stabilising RHPC methods introduced in Section 2.2, at least the end-point equality constraints on the output or on the unstable part of the output, depending on whether CRHPC or GPC^∞ are used, should be included. In fact, these constraints (equality and inequality) must be incorporated into the sequential QP problems which are solved in the SQP formulation, what involves an increase in the computational burden. Thus, even the numerical solution can become unpractical for relatively fast systems.

More efficient min-max controllers can be obtained if 1-norm cost functions are used. An example of such a formulation is provided in (Camacho and Bordóns, 1995), where the min-max 1-norm finite horizon GPC is presented. In the next section, the 1-norm quasi-infinite horizon controller defined in Section 2.3.3 is formulated in the min-max framework.

4.3 Min-max QGPC_1^∞

Several examples are provided in the literature to show that 1-norm formulations of min-max MPC lead to methods which involve a very low computational burden (Allwright, 1994; Camacho and Bordóns, 1995). The objective of this section is to formulate a min-max controller which can be implemented as a single LP problem, for which very efficient standard solutions are available.

As shown in Section 2.5, the QGPC_1^∞ converges to the GPC_1^∞ when the control horizon is increased, and hence is “very likely” to provide with nominal stability, at least for large enough N_u . The GPC_1^∞ itself does not appear to be a convenient choice to obtain a min-max controller because of the requirement that the iterative algorithm presented in Section 2.3.2 is applied at each sampling instant. In the min-max framework, this iterative procedure involves the solution of several min-max problems

(twice the number of iterations needed for convergence) at each sampling instant, what would involve an excessively large computational burden. On the other hand, if the QGPC₁[∞] is used, just a single min-max problem needs be solved at each sampling instant. In addition, the QGPC₁[∞] leads to an appropriate nominal behaviour for the vast majority of systems, as discussed in Chapter 2.

Furthermore, in min-max controllers, it is of particular relevance to keep the prediction horizon short since the number of polytope vertices increases exponentially with the prediction (or uncertainty) horizon. As remarked in Chapter 2, the CRHPC₁ with short control/prediction horizons often leads to deadbeat-like behaviour, and thus does not appear to be a suitable candidate to yield a robust min-max controller. The (quasi) infinite horizon approach to stability is preferred to the solution provided by the CRHPC₁, since deadbeat-like behaviour is unlikely with the GPC₁[∞] and the QGPC₁[∞].

Taking into account all these considerations, the QGPC₁[∞] seems a suitable candidate to provide with an effective solution of most control problems of uncertain systems within the global uncertainty approach. Whenever the QGPC₁[∞] does not provide a convenient nominal closed-loop behaviour (instability or deadbeat-like responses), two possibilities are suggested:

1. Increase the control horizon N_u until the QGPC₁[∞] converges to the GPC[∞] (see Section 2.5). If the control horizon required for convergence is “not too long” then the min-max QGPC₁[∞] solves the optimisation problem in very short time.
2. Implement the min-max 2-norm GPC[∞] using a numerical solution to reduce the computational burden.

4.3.1 Solving the min-max problem

The QGPC₁[∞] cost function is given by eqn.2.48:

$$J_1(t) = \sum_{j=1}^{N-1} |e(t+j|t)| + \alpha \sum_{j=N-n_a+1}^N |e(t+j|t)| + \sum_{j=1}^{N_u} \rho(j) |\Delta u(t+j-1|t)|,$$

which, as discussed in Chapter 2, provides an upper bound of the infinite horizon cost function if the system-dependent weighting α is chosen as suggested in eqn.2.38.

Now the min-max QGPC₁[∞] is defined as the solution to the optimisation problem:

$$\overline{\Delta \mathbf{u}^{\text{opt}}}(t) = \arg \min_{\Delta \mathbf{u}} \max_{\theta^{N_\theta} \in \Theta^{N_\theta}} J_1(t) \text{ subject to } \begin{cases} \mathbf{P}_u \Delta \mathbf{u} \leq \mathbf{r}_u, \\ \mathbf{P}_\theta \Delta \mathbf{u} \leq \mathbf{r}_\theta, \\ \tilde{\mathbf{G}} \Delta \mathbf{u} = \tilde{\mathbf{w}} - \tilde{\mathbf{f}} - \tilde{\mathbf{H}}_\theta \bar{\boldsymbol{\theta}}^{N+n_a}. \end{cases} \quad (4.15)$$

As discussed in Section 4.2.2.2, the uncertainty horizon N_θ must extend to $N + n_a$ for the equality constraints.

As the predictions are affine functions of $\boldsymbol{\theta}^{N_\theta}$, the cost function and the inequality constraints $\mathbf{P}_\theta \Delta \mathbf{u} \leq \mathbf{r}_\theta$ need only be evaluated at the polytope vertices $\boldsymbol{\theta}_i^{N_\theta} \in \Theta^{N_\theta}$. Thus there are 2^{N_θ} sets of constraints of this form, one set for each vertex. The constraint reduction procedure depicted in Section 4.2.2.1 can be used to cut down the number of constraints.

Remark 4.5 Henceforth, it is assumed, unless otherwise explicitly specified, that the constraints $\mathbf{P}_u \Delta \mathbf{u} \leq \mathbf{r}_u$ and $\mathbf{P}_\theta \Delta \mathbf{u} \leq \mathbf{r}_\theta$ are enforced in the prediction horizon $t+1, t+2, \dots, t+N$ (not $t+N+n_a$). Hence the dimension of the uncertainty vectors $\boldsymbol{\theta}^N$ relevant to these constraints is N and not $N + n_a$. Uncertainty vectors of dimension $N + n_a$ are used only for the equality constraints. Notice that the constraint horizon might need to be increased even beyond $N + n_a$ to obtain infinite horizon-like constraints, which are needed to preserve the nominal stability property, as remarked in (Rawlings and Muske, 1993). However, a constraint horizon of N is enough for the vast majority of systems and, without loss of generality, is used hereafter. If the upper constraint

horizon is greater than N , the formulae provided below can be used introducing slight straightforward modifications. □□□

This optimisation problem can be solved, using LP tools, similarly as done in Section 2.3.3 for the T -based QGPC₁[∞]. However, the LP problem must consider the maximum, among all the polytope vertices, of the cost function $J_1(t)$. Let the variables $\sigma(j) \geq 0$ and $\beta(j) \geq 0$ be defined as in Section 2.3.3, then the problem of minimising Ψ subject to:

$$\begin{aligned}
 -\sigma(j) &\leq e(t+j|t) \leq \sigma(j), \quad 1 \leq j \leq N, \forall \theta^N \in \Theta^N \\
 -\beta(j) &\leq \Delta u(t+j-1|t) \leq \beta(j), \quad 1 \leq j \leq N_u, \\
 0 &\leq \sum_{j=1}^N \mu(j)\sigma(j) + \sum_{j=1}^{N_u} \rho(j)\beta(j) \leq \Psi,
 \end{aligned} \tag{4.16}$$

with the weighting sequence $\mu(j)$ defined in Section 2.3.3, and subject to the equality/inequality constraints associated to the controller, can be thought of as a LP version of the problem described in eqn.4.15. Notice that the variables $\sigma(j)$ stand for the maximum predicted tracking errors for all the possible uncertainty vectors within the polytope Θ^N . Now, since the cost function is (strictly) convex, this problem must be considered only at the vertices, and the constraints associated with the variables $\sigma(j)$ can be replaced by

$$-\sigma(j) \leq e(t+j|t) \leq \sigma(j), \quad 1 \leq j \leq N, \forall \theta_i^N \in \Theta^N.$$

Remark 4.6 This LP problem is, however, more conservative than the formulation of eqn.4.15, since the variables $\sigma(j)$ collect the worst case prediction errors for different vertices θ_i^N . Thus $\sigma(j_1)$ can occur for the vertex $\theta_{i_1}^N$, whereas $\sigma(j_2)$ occurs at a different vertex $\theta_{i_2}^N$ (with $i_1 \neq i_2$). In fact, due to the convexity of $J_1(t)$, the maximum of eqn.4.15 is to be found at a single vertex, and not at a combination of different vertices. To solve *exactly* the min-max problem of eqn.4.15, different sets of variables $\sigma_i(j)$ should

be used for different vertices, *i.e.*

$$\begin{aligned} -\sigma_i(j) &\leq e(t+j|t) \leq \sigma_i(j), \quad 1 \leq j \leq N, \forall \theta_i^N \in \Theta^N \\ 0 &\leq \sum_{j=1}^N \mu(j)\sigma_i(j) + \sum_{j=1}^{N_u} \rho(j)\beta(j) \leq \Psi, \quad \forall \theta_i^N \in \Theta^N. \end{aligned}$$

This modification yields an exact solution to the min-max problem, but the computational price is enormous, since the number of variables σ_i (and the associated constraints) is multiplied by 2^N . For computational reasons, it is thus advisable to solve the problem as formulated in eqn.4.15, although an unnecessarily conservative solution might be attained. □□□

The problem of minimising Ψ (eqn.4.16) can be posed as the linear programming problem:

$$\min_{\Psi, \sigma, \beta, \Delta u} \Psi, \text{ subject to } \left\{ \begin{array}{l} \sigma \geq G\Delta u + H_\theta \theta_i^N + f - w, \\ \sigma \geq -G\Delta u - H_\theta \theta_i^N - f + w, \\ P_\theta \Delta u \leq r_\theta, \\ \beta \geq \Delta u, \\ \beta \geq -\Delta u, \\ \Psi \geq \mu^T \sigma + \rho^T \beta, \\ P_u \Delta u \leq r_u, \\ \tilde{G}\Delta u = \tilde{w} - \tilde{f} - \tilde{H}_\theta \bar{\theta}^{N+n_a}, \\ \sigma \geq 0, \beta \geq 0, \Psi \geq 0, \end{array} \right\} \forall \theta_i^N \in \Theta^N$$

where the coefficients $\mu(j)$ are defined in Section 2.3.3 and the vectors μ , σ , ρ and β are defined in eqn.2.35. Notice that only the equality constraints on the unstable part of the output and the first few inequality constraints are uncertainty-dependent, whereas the rest of them are independent from the global uncertainty signal.

As usual, this problem can be written in the standard form as

$$\min_x c^T x \text{ subject to } Ax \leq b, \sigma \geq 0, \beta \geq 0, \Psi \geq 0,$$

with

$$x = \begin{bmatrix} \Delta u \\ \sigma \\ \beta \\ \Psi \end{bmatrix}, \quad c = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix},$$

and

$$\begin{aligned} \mathbf{A} &= [\mathbf{A}_1^T \quad \dots \quad \mathbf{A}_{2^N}^T \quad \mathbf{A}_u^T]^T, \\ \mathbf{b} &= [\mathbf{b}_1^T \quad \dots \quad \mathbf{b}_{2^N}^T \quad \mathbf{b}_u^T]^T, \end{aligned}$$

where the block matrices \mathbf{A}_i and vectors \mathbf{b}_i are defined as

$$\begin{aligned} \mathbf{A}_i &= \left[\begin{array}{c|c|c|c} \mathbf{G} & -\mathbf{I} & \mathbf{0} & \mathbf{0} \\ \hline -\mathbf{G} & -\mathbf{I} & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{P}_\theta & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{array} \right], & \mathbf{b}_i &= \left[\begin{array}{c} -\mathbf{H}_\theta \boldsymbol{\theta}_i^N - \mathbf{f} + \mathbf{w} \\ \mathbf{H}_\theta \boldsymbol{\theta}_i^N + \mathbf{f} - \mathbf{w} \\ \hline \mathbf{r}_i \end{array} \right], \\ \\ \mathbf{A}_u &= \left[\begin{array}{c|c|c|c} \mathbf{I} & \mathbf{0} & -\mathbf{I} & \mathbf{0} \\ \hline -\mathbf{I} & \mathbf{0} & -\mathbf{I} & \mathbf{0} \\ \hline \mathbf{0} & \boldsymbol{\mu}^T & \boldsymbol{\rho}^T & -1 \\ \hline \mathbf{P}_u & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \hline \tilde{\mathbf{G}} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \hline -\tilde{\mathbf{G}} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{array} \right], & \mathbf{b}_u &= \left[\begin{array}{c} \mathbf{0} \\ \mathbf{0} \\ \hline \mathbf{0} \\ \hline \mathbf{r}_u \\ \hline -\tilde{\mathbf{H}}_\theta \bar{\boldsymbol{\theta}}^{N+n_{\bar{a}}} - \tilde{\mathbf{f}} + \tilde{\mathbf{w}} \\ \hline \tilde{\mathbf{H}}_\theta \bar{\boldsymbol{\theta}}^{N+n_{\bar{a}}} + \tilde{\mathbf{f}} - \tilde{\mathbf{w}} \end{array} \right]. \end{aligned}$$

Now, since there are exactly 2^N pairs $(\mathbf{A}_i, \mathbf{b}_i)$, the total number of constraints can be obtained as $[2N + \dim(\mathbf{r}_1)]2^N + 2N_u + 1 + \dim(\mathbf{r}_u) + 2(n_{\bar{a}} + 1) + N + N_u + 1$ (the last three terms come from the constraints $\boldsymbol{\sigma} \geq \mathbf{0}$, $\boldsymbol{\beta} \geq \mathbf{0}$, and $\Psi \geq 0$). Notice that the block matrices \mathbf{A}_i are identical, and thus the constraint reduction procedure depicted in Section 4.2.2.1 can be used, as discussed in (Camacho and Bordóns, 1995). If the j^{th} row of the constraint blocks $\mathbf{A}_i \mathbf{x} \leq \mathbf{b}_i$ is considered, the only constraint which determines the feasible region is the one for which the j^{th} element of the vector \mathbf{b}_i is the lowest, and the other $2^N - 1$ constraints are redundant and can be discarded. Hence, this LP problem can be implemented reducing the number of constraints to just $2N + \dim(\mathbf{r}_1) + 2N_u + 1 + \dim(\mathbf{r}_u) + 2(n_{\bar{a}} + 1) + N + N_u + 1$, *i.e. exactly the same number as for the T -based QGPC₁[∞] of Section 2.3.3.*

It is worth pointing out that the constraint reduction is an NP problem, as already remarked in Section 4.2.2.1. If the uncertainty bounds θ^- and θ^+ are taken to be time invariant, this fact is not relevant since the constraint reduction procedure can

be performed off-line obtaining the vectors $\mathbf{H}_\theta \theta_i^N$ for all the polytope vertices $\theta_i^N \in \Theta^N$. This is the case of the controllers introduced in (Camacho and Bordóns, 1995). However, the closed-loop dynamics of the uncertainty signal $\theta(t)$ can be used to update the uncertainty bounds, as shown in the following section. If the uncertainty bands are modified on-line, the constraint reduction procedure must be carried out on-line, and thus this NP problem should be solved *at each sampling instant*. In that case, it is advisable to keep N (or N_θ), and consequently N_u , small for computational reasons. This implies a trade-off choice of N_u , since nominal stability with the QGPC₁[∞] requires a “large enough but relatively small” N_u , whereas computational issues suggest to choose a short control horizon.

4.3.2 A band updating algorithm

In this section a band updating algorithm to modify the lower and upper global uncertainty bounds, θ^- and θ^+ , on-line is suggested. To develop this algorithm the closed-loop uncertainty dynamics need be investigated.

To begin with, notice that the global uncertainty signal $\theta(t)$, as defined in eqn.4.1, can be measured on-line as the residue from the true to the predicted output, *i.e.*

$$\theta(t) = y(t) - \hat{y}(t|t-1),$$

where $\hat{y}(t|t-1)$ is the output prediction at time t using information available at time $t-1$, *performed assuming that the future global uncertainty is zero*. At time t , this one-step-ahead prediction can be obtained as

$$\hat{y}(t+1|t) = f(t+1|t) + g_1 \Delta u(t|t).$$

The following experiment has been carried out to illustrate the typical closed-loop behaviour of the residues $\theta(t)$. The lightly damped second-order system of eqn.A.10:

$$G(q^{-1}) = \frac{q^{-1}B}{A} = \frac{0.2358q^{-1} + 0.2319q^{-2}}{1 - 1.4835q^{-1} + 0.9512q^{-2}}$$

has been chosen as the nominal model. The poles of this system are located at $0.7418 \pm 0.6333j$ and the zero at -0.9832 . On the other hand, the true plant has been defined as

$$G_0(q^{-1}) = \frac{q^{-1}B_0}{A_0} = \frac{0.1091q^{-1} + 0.0751q^{-2}}{1 - 1.4835q^{-1} + 0.7467q^{-2}},$$

i.e. the true zero (-0.6882) and steady-state gain (0.7) are 0.7 times the nominal ones, and the true poles ($0.7418 \pm 0.4433j$) lead to a much more damped open-loop response compared to the nominal plant. The true output has been computed as

$$A_0(q^{-1})y(t) = B_0(q^{-1})u(t-1) + x(t), \quad (4.17)$$

where the (unmeasurable) additive disturbance $x(t)$ is 0 for $t < 51$ and 0.05 for $t \geq 51$. Notice that such a disturbance affects the internal states and not only the output of the system.

To obtain the closed-loop uncertainty signal, the QGPC₁[∞] has been chosen with the tuning knobs $N_u = 5$ and $\rho = 1$. The model of eqn.4.1 has been used for the prediction for the nominal plant introduced above and with the assumption $\theta^- = \theta^+ = 0$ for all t . In the nominal case and for these tuning parameters, the QGPC₁[∞] is indistinguishable from the truly infinite horizon GPC₁[∞] (see Section 2.4.3). Notice that the min-max optimisation procedure can be avoided since all the assumed polytope vertices reduce to a single point, *i.e.* $\Theta = \mathbf{0}$. With these settings, a step setpoint change of unit amplitude has been simulated at the fifth sample.

Fig.4.3 displays the closed-loop behaviour of the global uncertainty $\theta(t)$. This result provides with an example of the typical uncertainty dynamics, and the following overall features can be observed:

1. After the setpoint change ($t = 5$), the controller operation leads to relatively large control moves and, consequently, to a sharp change in the global uncertainty signal $\theta(t)$ due to the modelling errors.

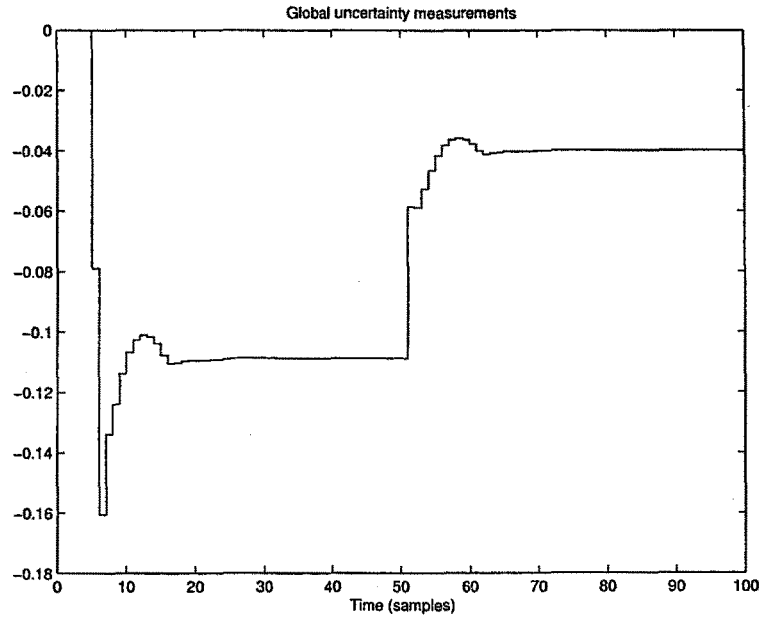


Figure 4.3: Global uncertainty signal

2. The global uncertainty signal oscillates for several samples and, after a while, settles down to some steady-state value.
3. When the disturbance enters the system ($t = 51$), the situation is similar to a set-point change, since the control activity introduced to compensate the disturbance implies an increase and some oscillation in the global uncertainty signal.
4. The disturbance leads to oscillations of the global uncertainty signal for a few samples, and finally the $\theta(t)$ settles down to a different steady-state value.

The lower and upper uncertainty bounds should be modified according to this closed-loop behaviour. However, it must be taken into account that whenever a band is violated by the true uncertainty signal, constraint violation or even instability may occur, since the true uncertainty can be worse than that assumed by the min-max method. Thus, any band updating algorithm should guarantee that the uncertainty bands are respected as much as possible. The following set of rules is intended to provide with an appropriate method to update the uncertainty bounds on-line:

1. Every time a setpoint change is introduced, the uncertainty bands must be widened to prevent that the uncertainty oscillations surpass the limits. This can be done by assigning a large, possibly very conservative, value to θ^- and θ^+ .
2. The bands must decrease and converge to the uncertainty values. When the global uncertainty is at steady-state, it is advisable that the lower and the upper bounds converge to the steady-state value, because exact predictions would be available, leading to offset-free setpoint tracking.
3. Unmeasurable disturbances are, by nature, unpredictable and they can lead to band (and possibly constraint) violations, which could even result in instability. Any time a band violation occurs, the lower and upper bounds must be modified (increased) accordingly so as to avoid future violations.

Remark 4.7 If the uncertainty limits are not respected by the global uncertainty signal, only the constraints which depend on $\theta(t)$, basically output and state constraints, can be affected (violated), whereas the input-like (input amplitude, input rate, input acceleration and so on) constraints are not to be influenced by these potential uncertainty band violations. This latter kind of constraints is totally independent of the uncertainty. □□□

The heuristic rules suggested above are illustrated in Fig.4.4. Notice that the upper uncertainty bound is violated when the disturbance enters the system, and thus the uncertainty bands are widened in order to cope with the future uncertainty. The band updating procedure suggested here modifies both the upper and the lower bounds whenever a band violation occurs. In the figure, the disturbance leads to an upper band violation and, when this situation is detected, the lower bound is also modified. The aim of this solution is twofold. To begin with, a perturbation usually leads to high-frequency uncertainty oscillations of large amplitude which can be enclosed by the updated uncertainty bounds if both limits are modified. On the other hand, the equality




Figure 4.4: Band updating procedure: global uncertainty $\theta(t)$ (solid) and uncertainty bounds (dotted)

constraints on the unstable part of the output are enforced for constant uncertainty at the average of the lower and upper bounds. Thus, if both θ^- and θ^+ are widened, a sudden change in the average $\bar{\theta}$ and, consequently, in the end-point equality constraints, is prevented.

Remark 4.8 When a disturbance involves uncertainty band violation, the constraints specified by the designer might be temporarily violated. If the band updating procedure manages to enclose the uncertainty signal, this possibility vanishes in the future. Disturbances may even lead to infeasibility problems, which can be handled using the methods described in (Scokaert, 1994; Álvarez and de Prada, 1997). Notice, however, that constraint violation and infeasibility situations cannot be avoided when unmeasured disturbances enter the system, since these are only detected by the feedback loop when the output is already disturbed. This drawback is common to all the constrained

MPC schemes. □□□

Finally, the suggested band updating algorithm (to be executed at each sampling instant t) can be implemented as:

1. $l(t) := \min\{\theta(t), \theta(t-1), \dots, \theta(t-M_\theta+1)\};$
 $h(t) := \max\{\theta(t), \theta(t-1), \dots, \theta(t-M_\theta+1)\};$
 $inc(t) := \max\{0, \theta^-(t) - \theta(t), \theta(t) - \theta^+(t)\};$
2. **if** $\theta^-(t) \leq \theta(t)$ **then**
 $\theta^-(t) := (1 - \mu)l(t) + \mu\theta^-(t-1) - inc(t)$
else
 $\theta^-(t) := \theta^-(t) - 2inc(t)$
endif
if $\theta^+(t) \geq \theta(t)$ **then**
 $\theta^+(t) := (1 - \mu)h(t) + \mu\theta^+(t-1) + inc(t)$
else
 $\theta^+(t) := \theta^+(t) + 2inc(t)$
endif

where $M_\theta \geq 1$ is the number of (past) uncertainty samples which are “remembered” by the updating algorithm, and $0 \leq \mu \leq 1$ is the pole of a first order filter with unit steady-state gain:

$$\frac{1 - \mu}{1 - \mu q^{-1}}. \quad (4.18)$$

According to the first step of the algorithm, $l(t)$ and $h(t)$ are assigned, respectively, the lowest and the highest last M_θ values of $\theta(t)$, and $inc(t)$ is the current amount of band violation (0 if the bands are respected). The procedure works as follows, if both bands are respected ($inc = 0$), they are updated in order to approach the extreme (maximum

and minimum) last M_θ values of $\theta(t)$ following a first order trajectory the dynamics of which depend on μ ($\mu \approx 1$ for slow rate and $\mu \approx 0$ for fast rate):

$$\theta^-(t) = \frac{1 - \mu}{1 - \mu q^{-1}} l(t), \quad \theta^+(t) = \frac{1 - \mu}{1 - \mu q^{-1}} h(t).$$

Apparently, higher-order filters might be used in this procedure, in a way that would recover some of the polynomial T filtering properties. However, a careful analysis of the effect of the filter of eqn.4.18 in the uncertainty bands dynamics discourages this possibility. To begin with, the filter of eqn.4.18 is not applied to the uncertainty signal itself, but the highest (h) and lowest (l) last M_θ values of $\theta(t)$. Hence the suggested band updating algorithm does not filter the modelling errors directly, and eqn.4.18 cannot be thought of as the counterpart of² S_p/T (Yoon and Clarke, 1995a; Megías, 1996; Megías *et al.*, 1997) in the T -based unconstrained MPC case analysed in Chapter 3. On the other hand, note that the variable *inc* is an input signal to the uncertainty bands (see the second step of the algorithm). Whenever *inc*(t) changes from 0 to some other value, this can be viewed as an impulse input to the uncertainty bands. It must be pointed out that the impulse response of (stable) first order filters is monotonically decreasing, but this is not true for greater order filters. The monotonicity of the impulse response coefficients of the filter which determines the band dynamics is quite a valuable property, since it speeds up the band convergence from the current values $\theta^-(t)$ and $\theta^+(t)$ to the (extreme) measurements $l(t)$ and $h(t)$, what avoids unnecessarily large band values for a few samples after a modification. As a consequence of these observations, only first order filters are used for the uncertainty band dynamics in this thesis.

Aside 4.1 *Compare the impulse response of the filters*

$$\frac{0.1}{1 - 0.9q^{-1}} \text{ and } \frac{0.01}{(1 - 0.9q^{-1})^2}.$$

End of Aside

²For stable systems, the choice $T = A(1 - \mu q^{-1})$ leads to $S_p/T = (1 - \mu)/(1 - \mu q^{-1})$.

If one of the two bands is violated at current time t ($inc > 0$), it is automatically updated by adding (or subtracting) *twice* the amount ($2inc$) it has been violated. In such a case, the other band is also added (or subtracted) the quantity inc .

Remark 4.9 In case of persistent (low-amplitude) noise, as always occurs in practical applications, a *looseness factor* can be introduced to keep the upper and lower bounds separated some distance at steady-state. The steady-state lower to upper bound offset should be greater than or equal to the maximum noise amplitude, what would guarantee constraint satisfaction. If the looseness factor is tuned to bound the noise amplitude tightly, then the setpoint can be brought closer to the constraint boundary, which usually determines the optimal operating condition, at least from an economical point of view. □□□

4.3.3 Tuning guidelines

The band updating procedure outlined above uses a few tuning knobs. The aim of these parameters is to adjust the robustness/performance trade-off. The two main tuning knobs to be chosen are the “past” uncertainty horizon M_θ and the band dynamics parameter μ . In addition, there are two secondary parameters to be set, namely, the band enlargement due to a setpoint change (band initialisation), and the number of samples the bands are “frozen” whenever a setpoint change occurs (band freezing), which are illustrated in Fig.4.4. As the band initialisation is concerned, some knowledge about the increase of the residues from the true output to the model predictions when a setpoint change occurs is needed to adjust this parameter. The most obvious possibility is to widen the uncertainty bounds a given percentage of the setpoint change. This alternative might be too simplistic for some non-linear systems, since the uncertainty behaviour often varies at different operating points, and consequently different setpoint changes can produce different uncertainty behaviour. If little knowledge about the residues is available, it is advisable to assign conservative (large) bound values which

keep the input activity low for a few samples. The number of samples the initial band values are frozen after a setpoint change is useful to allow the band updating algorithm to collect enough data to include, at least, a maximum and a minimum of the global uncertainty signal. The longer this period is, the more conservative the controller becomes. If a slow μ is chosen, this parameter does not need to be as large as the peak distance (p in Fig.4.4), since the initial values of θ^- and θ^+ are slowly modified (decreased), and the maximum and minimum of $\theta(t)$ are very likely to occur within the bands. In addition, note that the band updating procedure self-adjust the uncertainty bounds whenever a band violation occurs. Hence, these two parameters (band initialisation and band freezing) are not critical, but they help to avoid uncertainty band surpassing and, consequently, the constraint violations which could result as a consequence.

Now, tuning guidelines for the two main parameters, M_θ and μ , are proposed. In order to enclose the uncertainty oscillations, the past uncertainty horizon parameter should be chosen so as to “remember” at least a local maximum and a local minimum of the uncertainty signal. Thus M_θ is suggested to be greater than or equal to the peak to peak distance p (see Fig.4.4), *i.e.* $M_\theta \geq p$. As μ is concerned, this parameter determines the dynamics of the lower and upper uncertainty bounds. As a rule of thumb, it is advised that μ is chosen according to the decreasing rate of the enveloping curve which encloses the uncertainty signal. Hence μ should be tuned such that the uncertainty bounds dynamics are similar (or slower) than those of the uncertainty envelope.

Notice that the tuning guidelines for M_θ and μ , though intuitive, are somewhat difficult to apply since they require some *a priori* knowledge about the behaviour of the residues. Thus, the following issue must be tackled:

How can M_θ and μ be chosen for the first time when no knowledge about

$\theta(t)$ is available?

The most obvious way to overcome this difficulty is to perform a few experiments in order to obtain uncertainty data which can be used to tune both M_θ and μ . If a detailed (possibly non-linear) model of the true system is available, then the experiments can be simulated, and then these tuning parameters can be chosen as suggested in the guidelines reported above. If it is not possible to perform simulated experiments, large conservative values for M_θ and μ should be chosen at the beginning. These can be replaced by tighter (less conservative) choices in later adjustments to improve performance.

The closer the uncertainty bounds are to the envelope of $\theta(t)$, the less conservative the controller becomes. A convenient choice of M_θ and μ can help to improve performance, making it possible to operate in the proximity of the constraint boundary.

4.3.3.1 A self-tuning procedure

The tuning guidelines for M_θ and μ reported above can be implemented within the controller to provide with a self-tuning band updating algorithm. Using the past values of $\theta(t)$ it is quite an easy task to locate the last two local maxima and the last two local minima by examining at which samples the derivative (the difference) of $\theta(t)$, namely $\Delta\theta(t)$, changes from positive to negative (maximum) or from negative to positive (minimum). The distance (in samples) from maximum to maximum and from minimum to minimum can then be readily evaluated, and this information used to update M_θ .

As shown in Fig.4.5, let M_1 and M_2 denote, respectively, the last two local maxima occurring at the time samples t_{M1} and t_{M2} . Analogously, let m_1 and m_2 denote, respectively, the last two local minima and t_{m1} and t_{m2} their occurrence times. According to




Figure 4.5: Local minima and maxima for the self-tuning procedure

these definitions, M_θ can be chosen as

$$M_\theta = \max \{ M_\theta^{\min}, p_M, p_m \},$$

where M_θ^{\min} is a minimum value used for safety (*e.g.* 10 samples), $p_M = t_{M_1} - t_{M_2} + 1$ is the peak distance between the last two maxima, and $p_m = t_{m_1} - t_{m_2} + 1$ is the peak distance between the last two minima. If the modelling errors are LTI, then p_m and p_M must be identical (differing at most by 1 or 2), but in the more general case of non-linear and/or time-varying uncertainty distances p_m and p_M might be quite different.

The same information can be used to update the first order pole μ by measuring the damping factor ψ of the uncertainty curve. Once the peak distance p has been evaluated (as the maximum of p_M and p_m), ψ can be obtained from the last two pairs of local minima and maxima, since the following relation should (approximately) hold:

$$(M_1 - m_1) = \psi^p (M_2 - m_2),$$

and ψ can be estimated as

$$\psi \approx \left(\frac{M_1 - m_1}{M_2 - m_2} \right)^{\frac{1}{p}}.$$

Finally, the pole of the first order filter can be set as

$$\mu = \min \{ \max \{ \mu^{\min}, \psi \}, \mu^{\max} \},$$

where the limits μ^{\min} and μ^{\max} are introduced for safety, to avoid too fast or too slow dynamics.

Remark 4.10 The true situation is not usually as well-behaved as that shown in Fig.4.5. The contribution of noise and of secondary high-frequency oscillations often require some filtering or signal processing on $\theta(t)$ prior to proceed with the computation of M_θ and μ . This difficulty can become so severe in practice that the use of this self-tuning method would sometimes be discouraged unless signal processing tools are applied to filter out the rippling of $\theta(t)$. □□□

Note that it is always possible to adjust this self-tuning procedure so that the adaptive band updating method is at least as conservative as some given fixed setting. For instance, if $[M_\theta, \mu] = [10, 0.8]$ have been found to be mainly correct but some band violations occur, the self-tuning method can be used with $M_\theta^{\min} = 10$, $\mu^{\min} = 0.8$ and (e.g.) $\mu^{\max} = 0.95$ to obtain a behaviour which would certainly be at least as conservative as the fixed design $[M_\theta, \mu] = [10, 0.8]$. The self-tuning method is expected to head M_θ and μ in the appropriate direction, leading to a finer band updating algorithm and an improved performance.

4.4 Simulation results and comparative analyses

This section is devoted to perform simulated experiments with the min-max controllers formulated in this chapter in order to test this approach against other control strategies. In the sequel, simulations for both linear and non-linear systems are provided.

4.4.1 Comparative study of 1-norm and 2-norm min-max controllers

This section compares the performance and computational burden of two min-max controllers based on the global uncertainty approach, namely, the min-max 2-norm GPC^∞ and the min-max (1-norm) QGPC_1^∞ . The band updating procedure and the constraint reduction method reported in the last few sections have been used for both controllers. For the min-max GPC^∞ a numerical algorithm based on SQP methods has been used for solving the optimisation problem, since this approach does not involve the enormous computational burden associated to the analytical solution suggested in Section 4.2.2.3.

Let the true and nominal systems be those used in Section 3.4.3, *i.e.*

$$G(q^{-1}) = \frac{q^{-1}B}{A} = \frac{0.2358q^{-1} + 0.2319q^{-2}}{1 - 1.4835q^{-1} + 0.9512q^{-2}},$$

and

$$G_0(q^{-1}) = \frac{q^{-1}B_0}{A_0} = \frac{0.2973q^{-1} + 0.2923q^{-2}}{1 - 1.7802q^{-1} + 1.3698q^{-2}}.$$

Notice that the true system has two unstable poles, whereas the nominal plant is stable (though very lightly damped). Hence this is quite a challenging experiment, since the true and the nominal dynamics are quite different. Apart from these modelling errors, a disturbance $x(t)$ of magnitude 0.05 has been added the true system at time $t = 51$ samples, *i.e.* the true output is simulated as

$$A_0(q^{-1})y(t) = B_0(q^{-1})u(t-1) + x(t),$$

with

$$x(t) = \begin{cases} 0 & \text{if } t < 51, \\ 0.05 & \text{if } t \geq 51. \end{cases}$$

As already shown in Section 2.5, the QGPC_1^∞ is stable for all N_u for this nominal system, and cannot be distinguished from the truly infinite horizon GPC_1^∞ for $N_u > 2$.

Hence, if the tuning settings are chosen as $N_u = 5$ and $\rho = 1$, the QGPC_1^∞ provides nominal stability and becomes, *de facto*, identical to the GPC_1^∞ . The same tuning knobs $[N_u, \rho] = [5, 1]$ have been used for the 2-norm GPC^∞ .

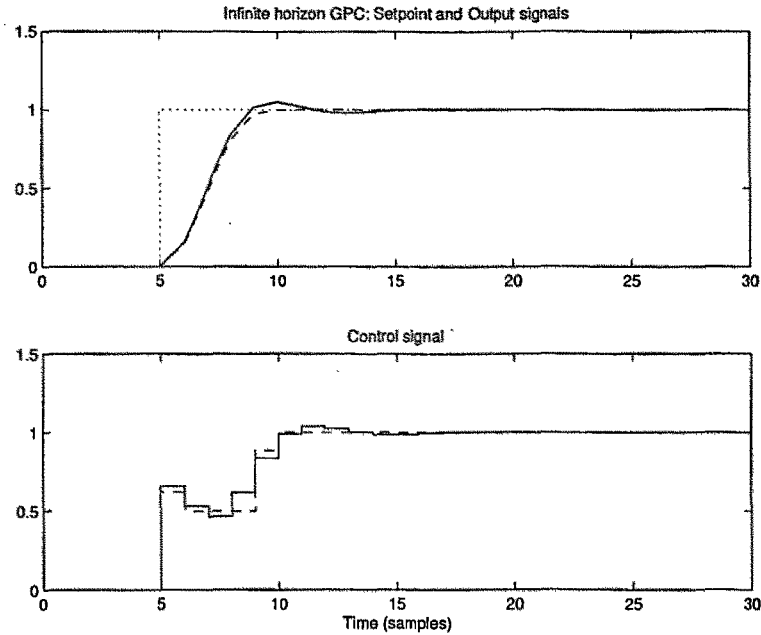


Figure 4.6: Nominal input/output responses for the GPC^∞ and the QGPC_1^∞ : $[N_u, \rho] = [5, 1]$

Fig.4.6, which is the same as Fig.3.12(a) of Chapter 3, compares the *nominal* closed-loop responses obtained with the 1-norm QGPC_1^∞ and the 2-norm GPC^∞ for these tuning settings. It can be observed that the nominal closed-loop behaviour obtained with both controllers is almost identical.

For the first few experiments, a setpoint change from 0 to 1 at time $t = 5$ samples is introduced. The initial uncertainty bound values are chosen as $\theta^-(5) = -0.2$ and $\theta^+(5) = 0.2$, *i.e.* a 20% of the setpoint change. In addition, these initial values are frozen for 4 samples after the setpoint change. Finally, the tuning knobs for the band updating procedure are set to $M_\theta = 10$ and $\mu = 0.9$.

The results obtained with the min-max GPC^∞ and the min-max QGPC_1^∞ are shown in Fig.4.7 and 4.8 respectively. It is quite remarkable that the closed-loop behaviour is,

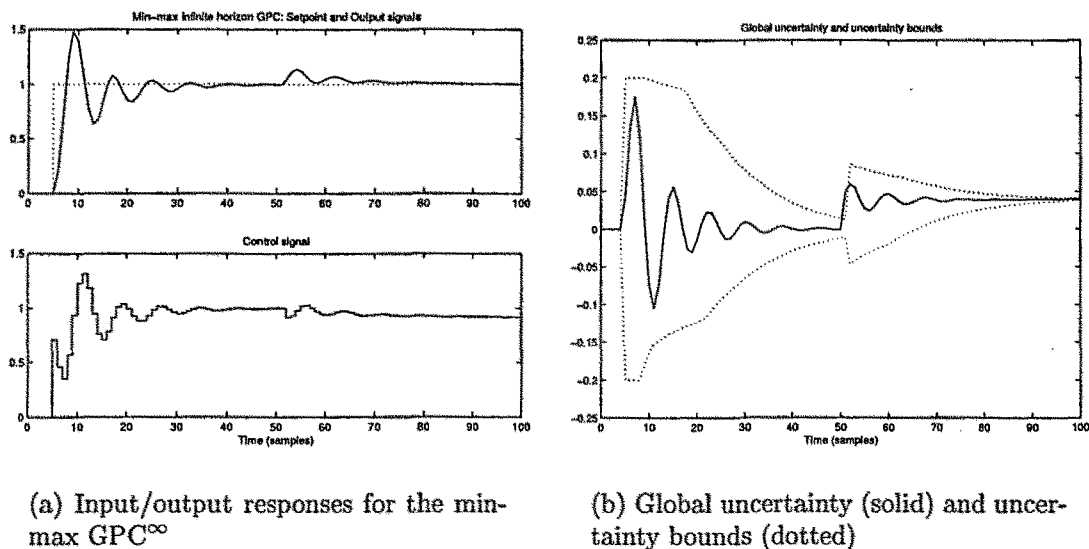



Figure 4.7: Closed-loop behaviour of the min-max (2-norm) GPC^∞

very approximately, the same for both controllers, and the following comments apply to both cases. First of all, notice that the global uncertainty increase associated to the setpoint change is always kept between the upper and the lower bounds (until the disturbance time $t = 51$). A few samples after the setpoint change, the band updating procedure modifies the assumed uncertainty limits, which converge to the true uncertainty values. The dynamics which shape this convergence are determined by the tuning parameters M_θ and μ . When the disturbance enters the system, there is a temporary (one-sample) band violation, which could result on constraint fulfilment problems in case that user-designed constraints had been used. Once the uncertainty values due to the disturbance are available ($t = 51$), the band updating procedure determines that the upper band has been violated and modifies the lower and upper bounds in order to keep $\theta(t)$ within the uncertainty bounds. No other violations occur afterwards. Finally, the uncertainty bounds converge to the steady-state value of $\theta(t)$, and a stable closed-loop system with offset-free setpoint tracking is achieved.

The convergence rate of the uncertainty bands dynamics is determined mainly by μ , and thus this parameter fixes the disturbance rejection properties of the controller.



(a) Input/output responses for the min-max QGPC_1^∞

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.8: Closed-loop behaviour of the min-max QGPC_1^∞

Only when the lower and upper bounds converge to the steady-state uncertainty value is the disturbance completely rejected. Therefore, as disturbance rejection is concerned, the role of μ is similar to that of the slowest root of the T polynomial (see Chapter 3). Finally, the peak distance of the signal $\theta(t)$ can be easily computed from the simulated experiments, and it has been found that p equals 8 or 9 samples. This justifies the choice $M_\theta = 10 \geq p$, which is consistent with the tuning guidelines suggested above.

Although the results are nearly the same with both controllers, the 2-norm min-max GPC^∞ has taken more than twice (about 2.2) the computation time of the 1-norm counterpart for this experiment. This difference becomes more serious for longer control horizons N_u (and thus longer prediction horizons N). Table 4.2 shows the normalised CPU times required by the 1-norm and 2-norm min-max controllers as a function of N_u (or N) for a simulation time of 50 samples. The computation times of Table 4.2 are normalised dividing by the CPU time required by the fastest experiment (the min-max QGPC_1^∞ with $N_u = 1$) so as to make the results as computer-independent as possible.

Remark 4.11 In these experiments, the optimisation routines used for the min-max controllers are the LP (function “lp”) and the SQP (function “constr”) algorithms

Horizon	N_u	1	2	3	4	5	6	7	8	9
	N	2	3	4	5	6	7	8	9	10
Normalised CPU time	1-norm (t_1)	1.0	1.6	2.5	3.9	8.0	11.3	20.8	22.5	38.6
	2-norm (t_2)	1.4	2.4	4.6	5.9	14.0	33.7	84.3	162.2	715.7
Ratio	t_2/t_1	1.4	1.5	1.8	1.5	1.7	3.0	4.0	7.2	19.0

Table 4.2: Normalised CPU time as a function of N_u (or N)

provided by the MATLAB *Optimisation Toolbox*. Both methods are interpreted and not compiled, and thus these data can be considered “fair”. In addition, in the SQP problem associated to the min-max GPC^∞ , the gradients of the cost function and the constraints have been used, making the SQP solution much faster and more reliable. If a compiled LP method is used for the min-max QGPC_1^∞ , e.g. the function “e04mbf” as implemented in the *NAG Foundation Toolbox*, the min-max QGPC_1^∞ becomes up to 15 (or more) times faster. □□□

Table 4.2 makes it clear that the 2-norm algorithm takes more computation time than the 1-norm controller always, from 1.4 up to about 20 times³. In addition this difference increases (dramatically) with the control horizon, and thus the 1-norm solution becomes a more convenient choice in most cases, since it allows more degrees of freedom for tuning the controller (N_u can be chosen greater preserving low computational burden).

Apart from computation time, there is another great advantage related to 1-norm min-max formulations, namely the use of LP tools instead of non-linear programming tools. The former can be exactly solved with a *finite* number of iterations which are known, *a priori*, from the number of design variables (N_u control moves plus the additional variables σ , β , and Ψ) and the number of constraints. Whenever the LP problem is feasible, the LP tools find a solution in very short time. On the other hand, there is no information about the number of iterations the SQP methods need to find

³The solution using compiled methods is several times faster than t_1 provided in the table.

a solution to the optimisation problem. It is even possible that SQP does not converge after a full sampling time, which is the all the time available for computing the next control move. Thus, as reliability is concerned, the on-line implementation of the LP solution is to be preferred to the SQP counterpart.

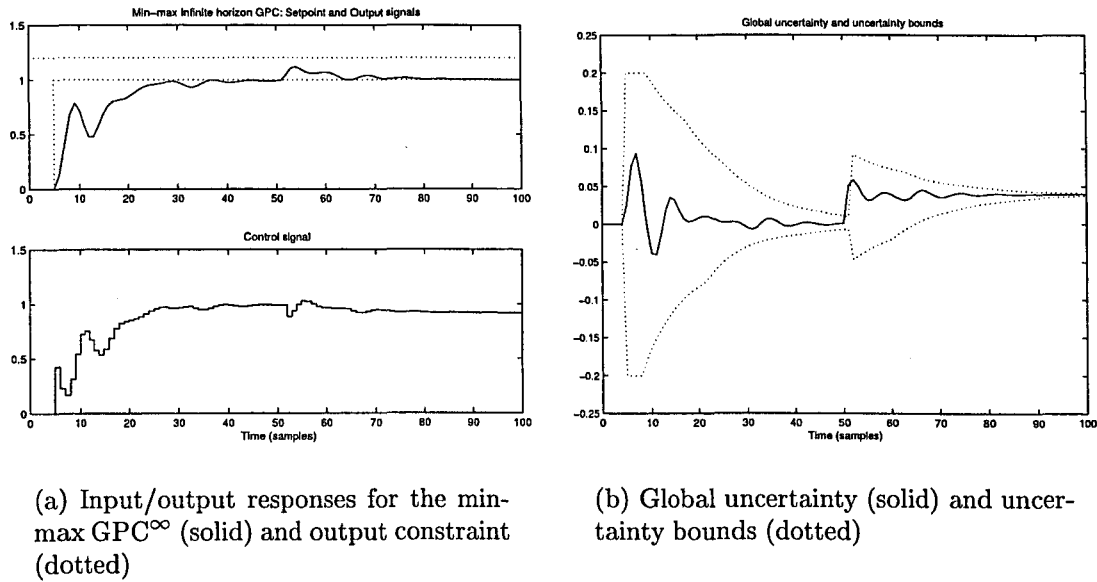


Figure 4.9: Closed-loop behaviour of the min-max (2-norm) GPC[∞] (constrained case)

In short, LP methods are faster and more reliable than the 2-norm min-max controllers. This is so clear an advantage that, for real applications, the nominal stability guarantees (the nominal QGPC₁[∞] is not always stabilising as remarked in Chapter 2) can often be overlooked. Moreover, as discussed above, the likelihood of nominal instability with the QGPC₁[∞] is remarkably small, at least for a wide class of systems. Nominal stability with the QGPC₁[∞] can be easily determined in simulation and thus, in the quite unlikely cases for which an unstable closed-loop systems arises, the control horizon can be increased until a stability is achieved (assuming that the convergence conjecture presented in Section 2.5 is true). Only when the control horizon required for stability is too large (*e.g.* $N_u > 12$) does the min-max two norm GPC[∞] appears as a better alternative, as far as it provides convenient performance with shorter control horizons.

Now the same experiments have been performed constraining the output such that $y(t) \leq 1.2$ for all t . The tuning settings used for these constrained examples are the same as the ones chosen for the unconstrained case. Fig.4.9 shows the closed-loop behaviour obtained with the min-max GPC^∞ . Despite the modelling errors and the disturbance, the output never violates the constraint. When the disturbance enters the system, the uncertainty bands are widened according to the band updating algorithm described above and the controller manages to satisfy the output constraint. Notice that the initial band values are too wide, since the uncertainty signal is quite far from the assumed bounds, especially as compared with the unconstrained case (Fig.4.7). This is a consequence of the output constraint, which results in lower control moves with respect to the unconstrained case and, consequently, to lower global uncertainty. If the initial band values were lower, better performance (less cautious control) would have arisen.

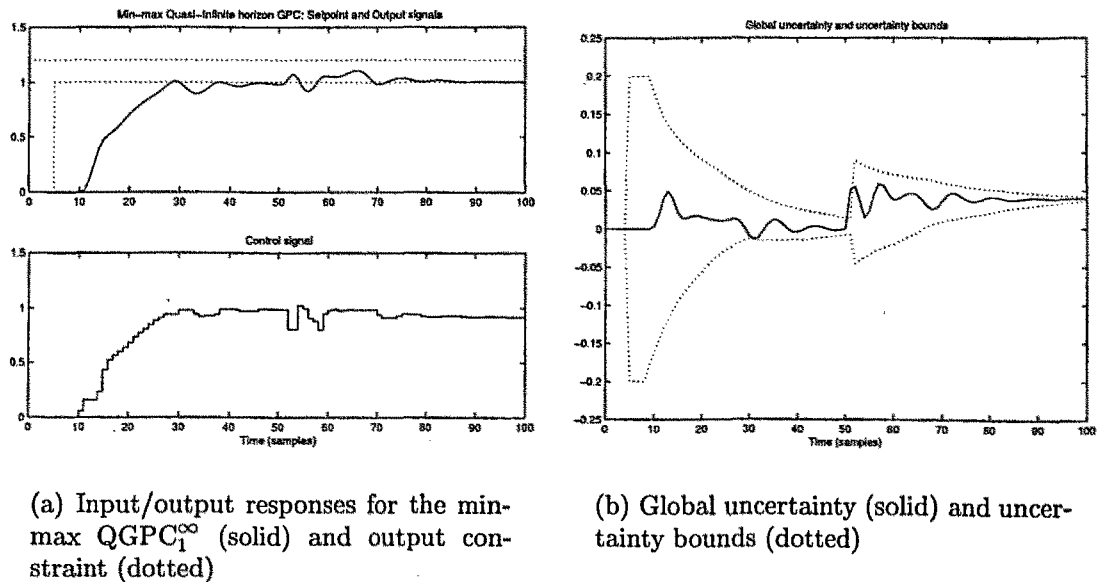



Figure 4.10: Closed-loop behaviour of the min-max QGPC_1^∞ (constrained case)

Fig.4.10 shows the closed-loop behaviour provided by the min-max QGPC_1^∞ , which is somewhat more cautious than that obtained with the 2-norm counterpart. A five-sample delay, which is not observed in the unconstrained case (Fig.4.8), shows up at

the output. This dead-time is due to the combination of the control move penalty ($\rho = 1$), which leads to tiny control moves, and the output constraint, which is quite close to the setpoint. Such a behaviour can be avoided by choosing a lower control move weight (e.g. $\rho = 0.1$) and/or narrower initial uncertainty bands (e.g. $\theta^- = -0.15$ and $\theta^+ = 0.15$).



(a) Input/output responses for the min-max QGPC_1^∞ (solid) and output constraint (dotted)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.11: Closed-loop behaviour of the min-max QGPC_1^∞ (constrained case)

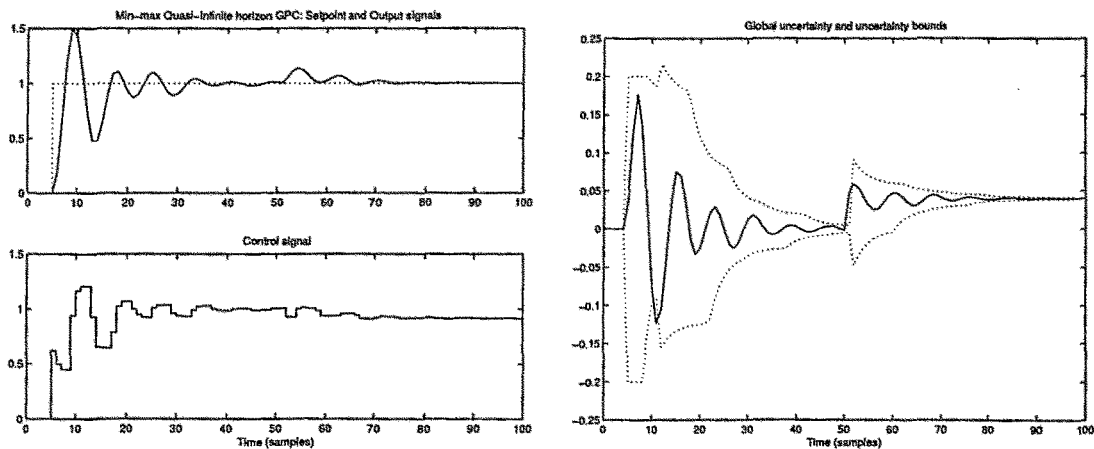
The closed-loop behaviour obtained with the min-max QGPC_1^∞ using the tuning knobs $\theta^-(5) = -0.15$, $\theta^+(5) = 0.15$ and $\rho = 0.1$, shown in Fig.4.11, evidences that the five-sample delay of Fig.4.10 is a consequence of the tuning settings. The narrower uncertainty bands and the lower penalty in the control moves results in an improved performance (less cautious control) compared to the previous result. As usual, a trade-off between robustness (cautiousness) and performance must be achieved.

4.4.2 The influence of the tuning parameters μ and M_θ

In this section the effect of the tuning parameters μ and M_θ is analysed. Although tuning guidelines are suggested above, there are some degrees of freedom to choose these

parameters which allow to adjust the balance between robustness and performance. The objective of this section is to show how these parameters affect the closed-loop behaviour.

The true and nominal systems presented in the previous section have been used for the experiments reported below. Since 1-norm and 2-norm controllers behave in a similar way as far as the tuning settings μ and M_θ are concerned, only the min-max QGPC $^\infty$ is analysed here, but the results extend to the min-max GPC $^\infty$ too.




(a) Input/output responses for the min-max QGPC $^\infty$ (solid) and output constraint (dotted)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.12: Closed-loop behaviour of the min-max QGPC $^\infty$ (unconstrained case)

First of all, the influence of μ is analysed. The experiment presented in Fig.4.8 is repeated but the first order trajectory determined by eqn.4.18 is generated with $\mu = 0.7$ instead of $\mu = 0.9$. The result, shown in Fig.4.12, is that the amplitude of the oscillations of the input/output responses and the uncertainty signal $\theta(t)$ is greater than those of Fig.4.8. Notice, also, that the uncertainty bands θ^- and θ^+ converge to the measurements of $\theta(t)$ quite faster than for the example presented in Section 4.4.1. The band dynamics are so fast that there is a lower band violation at the first few samples, but that situation is readily compensated by means of the band updating

procedure. Although a large difference is not evidenced, it is worth pointing out that the disturbance rejection dynamics are also faster with $\mu = 0.7$, since the lower and upper uncertainty bounds convergence rate is increased with respect to $\mu = 0.9$. In short, the closer the root μ is to the origin, the less conservative the control strategy becomes. This can even give rise to some uncertainty band violations, as occurs with this example, and caution should be taken to avoid such possibility, especially in the constrained case.



(a) Input/output responses for the min-max QGPC₁[∞] (solid) and output constraint (dotted)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.13: Closed-loop behaviour of the min-max QGPC₁[∞] (unconstrained case)

Now, the effect of the tuning parameter M_θ is analysed. The same experiment has been carried out using $\mu = 0.7$ (as in Fig.4.12) and $M_\theta = 20$, *i.e.* 20 past values (instead of 10) of the uncertainty signal $\theta(t)$ are examined at each sampling instant. The outcome of this experiment is shown in Fig.4.13, where it can be observed that the uncertainty bounds are farther from the uncertainty signal compared to Fig.4.12(b). The reason for such a behaviour is that this choice of M_θ implies that two maxima and two minima of $\theta(t)$ are “remembered” by the band updating algorithm. This is unnecessarily conservative and produces poorer performance compared to the case $M_\theta = 10$, shown in Fig.4.12(a), since wider oscillations and more sluggish input/output

responses result.

In short, this section shows that a convenient choice of μ and M_θ can help to adjust the compromise between cautiousness and performance. The closer the uncertainty bounds are to the uncertainty measurements, the less conservative the controller becomes, making it possible to obtain an improved performance.

4.4.3 Min-max controllers versus the classical T approach

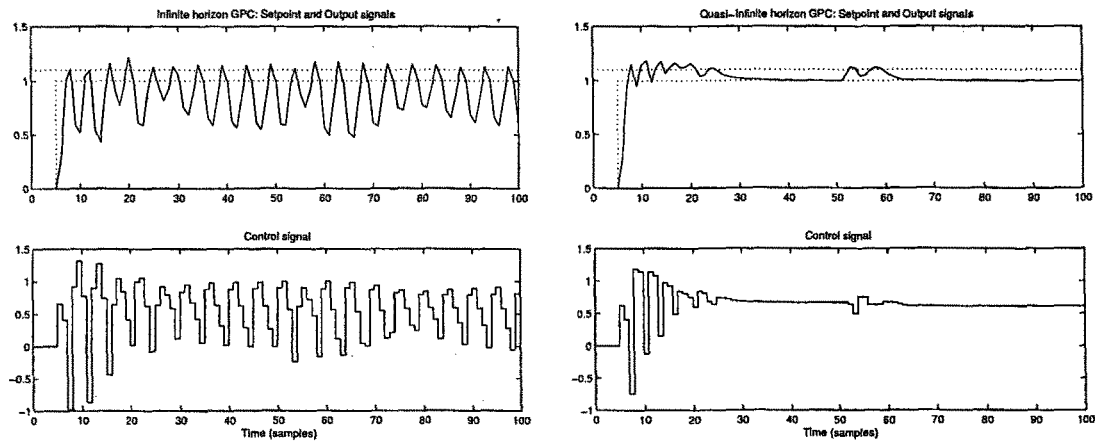
This section is focused on showing that the robustness-enhancing methods described for unconstrained MPC schemes in Chapter 3 fail to accomplish the expected closed-loop behaviour when hard (output) constraints are introduced.

4.4.3.1 Linear plant

In this section, the experiments have been performed using the same true and nominal systems as in Section 4.4.1, and the output constraint $y(t) \leq 1.1$, quite near the setpoint (1), has been enforced. The true steady-state gain has been increased a 50%, *i.e.* $G_0 = 1.5G_0$ in order to make it even more difficult to meet the constraint. The results obtained for both 1-norm and 2-norm algorithms are presented below.

The 2-norm T -based GPC^∞ and min-max GPC^∞ have been tuned with $N_u = 5$ and $\rho = 1$, whereas the 1-norm QGPC_1^∞ and min-max QGPC_1^∞ have been used with $N_u = 5$ and $\rho = 5$ to provide similar performance compared to the 2-norm case, especially as rise time is taken into account. For the min-max controllers, the band updating procedure has been implemented with $\theta^-(5) = -0.2$, $\theta^+(5) = 0.2$, $M_\theta = 10$ and $\mu = 0.9$. On the other hand, the polynomial $T = 1 - 0.9q^{-1}$ has been chosen for the standard (T -based) GPC^∞ and QGPC_1^∞ . With the methods depicted in Chapter 3, this choice of T can be shown to provide robust stability in the unconstrained case for the 2-norm

controller, as far as the inverse multiplicative uncertainty description is applied. Finally, it is worth pointing out that the (nominal) QGPC_1^∞ is indistinguishable from the GPC_1^∞ with these tuning settings, and hence the (min-max) QGPC_1^∞ can be considered as an efficient implementation of the (min-max) GPC_1^∞ and, of course, nominal stability is easily achieved.



(a) Input/output responses for the (2-norm) GPC^∞ (solid) and output constraint (dotted)


(b) Input/output responses for the QGPC_1^∞ (solid) and output constraint (dotted)

Figure 4.14: Closed-loop behaviour of the T -based GPC^∞ and QGPC_1^∞

Fig.4.14(a) shows that the closed-loop with the 2-norm T -based controller is on the verge of instability (if not unstable). In addition the output constraint is violated systematically by the maximum output peaks. On the other hand, the behaviour with the T -based QGPC_1^∞ , displayed in Fig.4.14(b), is a bit better as stability is concerned, but fails to satisfy the output constraint, which is violated several times after the setpoint change ($t = 5$) and the disturbance time ($t = 51$). Although the predictions made by both controllers are such that the output constraint is always respected, *i.e.* the predicted output is always lower than 1.1, the maximum peak reaches 1.2150 for the 2-norm controller and 1.1813 for the QGPC_1^∞ .

This experiment illustrates, indeed, the typical situation of the process industry. The system is expected to work near the constraint boundary (the setpoint is 1 and

the constraint is 1.1) for economical reasons, but modelling errors and/or disturbances can push the system beyond the scheduled limits. Hence, according to the results presented in Fig.4.14, the setpoint would have to be changed to say 0.9, farther from the constraint boundary to satisfy the limits. This problem can become even more serious if the constraints specify physical or security limits which must by no means be surpassed. In such a case, a severe malfunctioning or physical danger might be associated to constraint violations.



(a) Input/output responses for the min-max (2-norm) GPC^∞ (solid) and output constraint (dotted)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.15: Closed-loop behaviour of the min-max GPC^∞

Fig.4.15 shows the closed-loop behaviour obtained with a min-max 2-norm GPC^∞ . Notice that the output is always lower than the upper limit, in fact the maximum peak occurs at $1.0963 < 1.1$. Even when the additive disturbance enters the system, the output is readily compensated to avoid a constraint violation. Notice, also, that a constraint violation might only have occurred at time $t = 51$, the only instant at which the uncertainty bounds are not respected. However, the band updating algorithm redresses that situation, and the controller leads the output to a “safer” value for a few samples, to avoid constraint violation.

Finally, Fig.4.16 shows that the min-max $QGPC_1^\infty$ performs as effectively as the

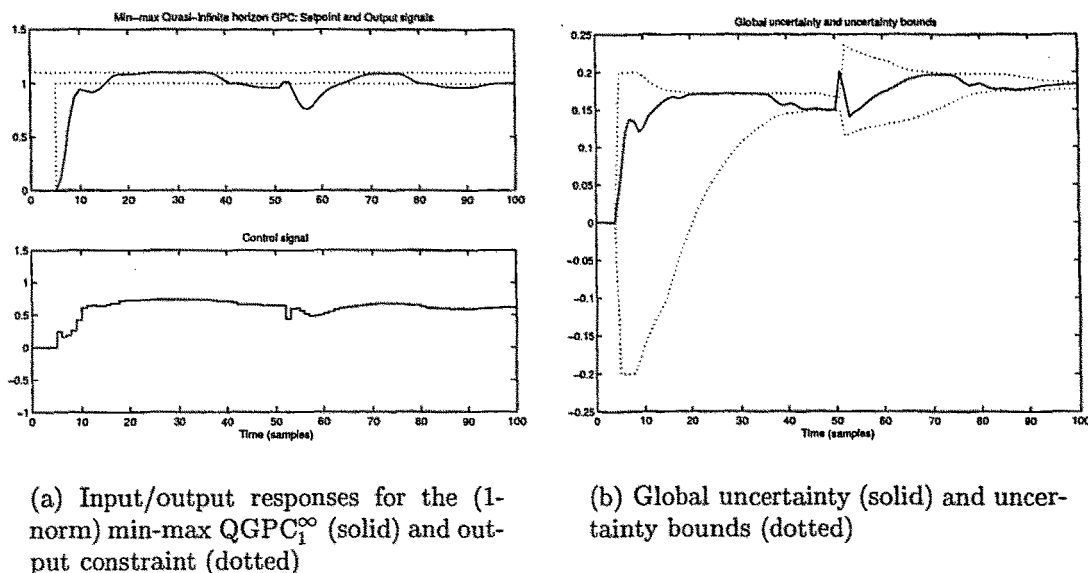


Figure 4.16: Closed-loop behaviour of the min-max QGPC_1^∞

2-norm counterpart (but with a lower computational burden as illustrated in Section 4.4.1), keeping the output below the upper constraint. In fact the maximum peak has been found at 1.0999, *i.e.* almost on the constraint boundary. Again, the output satisfies the constraint in spite of the additive disturbance and, of course, the modelling errors.

The examples reported in this section show the limitations of the classical robustness-enhancing tools of MPC when constraints are enforced. The results, shown for the T -based methods, extend to the Q -parametrisation and the T -optimisation schemes. In addition, it must be taken into account that the robust stability guarantees in the presence of input/output/state constraints are no longer valid, since the controller becomes non-linear (QP or LP must be used), and the robustness analysis performed through Chapter 3 relies on a linearity assumption. In fact, these simulation experiments show that the GPC^∞ does not manage to stabilise the true system, even though the RS conditions are satisfied. In short, the robustness methods for unconstrained MPC evidence two major flaws if they are applied in the constrained framework:

1. Firstly, although the internal predictions satisfy the constraints, these are often violated by the true output, and
2. secondly, the RS conditions and robustness analysis tools used in Chapter 3 cannot be applied.

On the other hand, the min-max methods based on a global uncertainty formulation, together with the band updating procedure suggested in Section 4.3.2, have proved successful in the few cases presented throughout this section, not only as stability is concerned, but especially as constraint satisfaction is taken into account. The band updating procedure is shown to drive the initial conservatism (band distance) closer to the actual measurements carried out on-line, what makes it possible to improve the performance and reach the setpoint with offset-free setpoint tracking. Moreover, the tuning guidelines given in Section 4.3.3, which are used in these first few examples, accomplish the expectations making the tuning task easier. However, these few examples may be biased and a deeper analysis is needed. The newly developed methods are difficult to analyse with the classical approach, as done for the unconstrained controllers in Chapter 3, since the min-max optimisation performed on-line is, intrinsically, a non-linear process. Obviously there is no closed-form solution for these controllers and the robustness analysis must be tackled from quite a different point of view. Section 4.5 is devoted to analyse the robustness of these min-max methods, whereas a few significant simulated experiments are provided in the sequel to contrast these min-max controllers with other control schemes suggested in the literature.

In the sequel, for computational reasons, 1-norm min-max controllers are preferred. As already shown in the examples so far described, the closed-loop behaviour of 1-norm and 2-norm min-max controllers is almost identical if appropriate tuning knobs are chosen, and thus the results provided below directly extend to the min-max GPC^∞ .

4.4.3.2 Non-linear plant with a saturation




Figure 4.17: Non-linear benchmark plant

The experiments presented in this section have been performed on the non-linear system described in Section A.5 of Appendix A, which is shown in Fig.4.17 (Fig.A.5). The nominal model, obtained ignoring the saturation and using a sampling time of $T_s = 0.05$ s with a ZOH on the input, is given by (eqn.A.10):

$$G(q^{-1}) = \frac{q^{-1}B}{A} = \frac{0.2358q^{-1} + 0.2319q^{-2}}{1 - 1.4835q^{-1} + 0.9512q^{-2}},$$


i.e. exactly the same nominal system as for the experiments presented in the previous few sections. As remarked in Appendix A, the main difficulty to control this system is the saturation embedded within an inner feedback loop. This kind of non-linearity is quite difficult to handle with linear control approaches (including linear MPC). Even the quite sophisticated methods described in (Kothare *et al.*, 1996) seem difficult to apply to this system, since this non-linearity cannot be overcome by obtaining a set of locally linearised models about different operating points. Such a drawback is common to all the strong non-linearities for which the associated functions are non-differentiable.

The most obvious way to surmount this difficulty is the use of anti-windup solutions, or to incorporate the constraint associated to the saturating amplifier ($-2 \leq x \leq 2$) explicitly into the control design problem. However, it is quite reasonable to assume that many of the very detailed models which are usually obtained as a prior step to design a control system are based on simplified physical laws. Thus, several strong non-linearities are often overlooked at the control design stage. These non-linearities are thus neglected in the model, and must be incorporated as “system uncertainty”

and handled with robustness tools. The saturation block embedded in the system of Fig.4.17 can be viewed as an illustrative example of this typical situation.

The next few experiments compare the behaviour obtained with the min-max QGPC_1^∞ and the T -based QGPC_1^∞ . For both controllers, the tuning parameters have been chosen as $N_u = 5$ and $\rho = 5$, which provide a smooth first-order nominal response. In addition, the polynomial $T = 1 - 0.9q^{-1}$ has been used for the latter. The setpoint changes from 0 to 55 at time $t = 5$ samples, and a constant additive disturbance of amplitude 3 affects the output for $t \geq 51$. The setpoint has been chosen such that the saturation is outside the linear region for quite long, leading to non-zero global uncertainty. Notice that when the input x to the saturating block is between -2 and 2 the modelling errors reduce to 0.

To complete the experiment setup, the band updating procedure of the min-max controller has been tuned using $\mu = 0.9$, $M_\theta = 10$ and initial band values $\theta^-(5) = -15$ and $\theta^+(5) = 15$, *i.e.* the amplitude of the uncertainty signal is expected to be less than a 30% of the setpoint change during the first few samples.



(a) Input/output responses for the min-max QGPC_1^∞ (solid) and T -based QGPC_1^∞ (dashed)

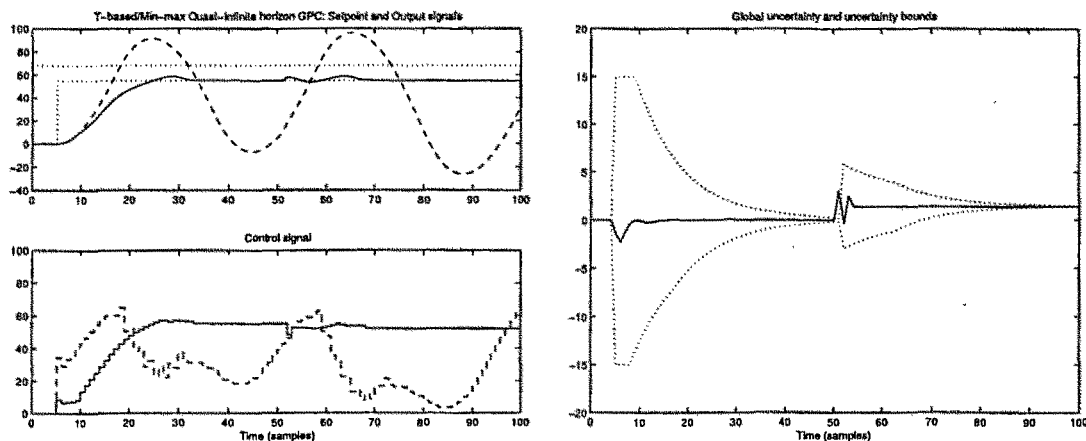
(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.18: Closed-loop behaviour of the min-max QGPC_1^∞ and the T -based QGPC_1^∞

Fig.4.18 shows the closed-loop behaviour obtained with both controllers. Notice, in Fig.4.18(b), that, apart from the disturbance time, the uncertainty signal is always within the lower and upper bounds, which finally converge to some steady-state value. The peak to peak distance of $\theta(t)$ cannot be observed in this experiment, since the controller leads the plant to the linear region (in which the global uncertainty equals 0) before a second maximum or minimum of $\theta(t)$ occurs. Fig.4.18(a) shows that the input/output responses obtained with the min-max method are smoother, less oscillatory, compared to the “classical” T -design. It may be argued that another choice of T such as $T = A(1 - 0.9q^{-1})$ (as suggested in Chapter 3) may lead to better results but, as a consequence of the proximity of the nominal system’s poles to the unit circle, such a choice produces an unstable closed-loop system the responses of which are not shown here for brevity. The disturbance rejection characteristics are similar for both the T -based and the min-max controllers, since the root of T and the parameter $\mu = 0.9$, which determines how fast the uncertainty bands converge to the last M_θ extreme uncertainty values, are identical. Needless to say, the results differ using other observer polynomials, and better closed-loop behaviour arises by using *e.g.* $T = (1 - 0.9q^{-1})^2$. However, this latter choice of T produces a more sluggish response, and the disturbance rejection response is slower compared to the min-max controller. Therefore the comparison provided in Fig.4.18 can be regarded as “fair”.

For the next example, the conditions of the last experiment (controllers, tuning settings, setpoint change, disturbances and so on) are re-created, but an output constraint $y(t) \leq 68$ for all t has been enforced. Fig.4.19(a) shows that the min-max controller produces a smooth response which never surpasses the constraint, whereas the T -based QGPC₁[∞] fails not only to satisfy the constraint, but even to provide with a stable closed-loop system. On the other hand, the choice $T = (1 - 0.9q^{-1})^2$ stabilises the system, but fails to satisfy the output constraint specification.

In Fig.4.19(b) it can be observed that, as a consequence of the output constraint,



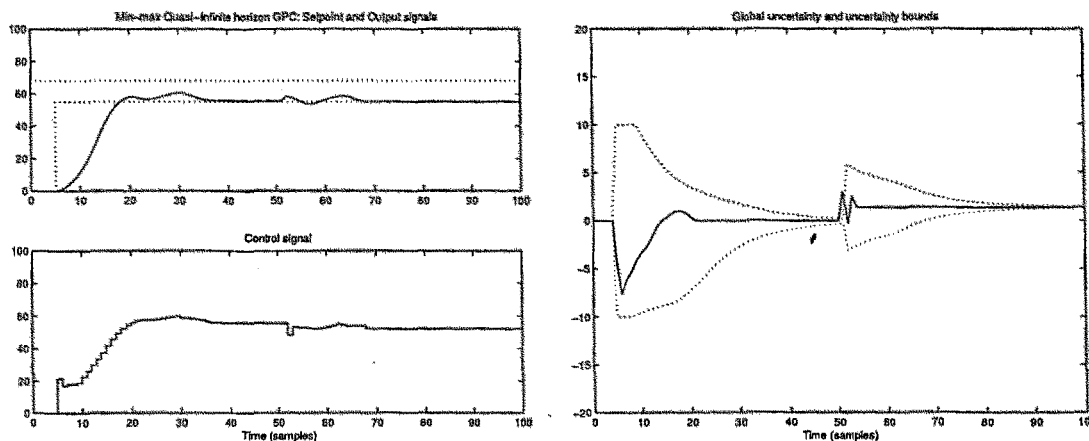
(a) Input/output responses for the min-max $QGPC_1^\infty$ (solid) and T -based $QGPC_1^\infty$ (dashed), and output constraint (dotted)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.19: Closed-loop behaviour of the min-max $QGPC_1^\infty$ and the T -based $QGPC_1^\infty$

the min-max controller is quite cautious and the small control efforts revert on very low values of the global uncertainty signal, which are quite far from the initially assumed bounds $[-15, 15]$. Hence the performance can be improved (a faster response can be obtained) if the initial band distance is narrowed. However, a conservative tuning can be justified in practical applications, especially for the first few experiments as little data about $\theta(t)$ are available. This experiment provides with a (simplified) instance of a real case, in which conservativeness leads to very cautious responses.

In order to show the effect of using lower initial uncertainty bands, the simulation results obtained with $\theta^-(5) = -10$ and $\theta^+(5) = 10$ are displayed in Fig.4.20. Notice that the input/output responses are faster, and the output reaches the setpoint in less than 15 samples (0.75 seconds). Once again, the output constraint is always respected. In addition, with these initial settings, the controller is less cautious compared to the previous experiment and, as a consequence of this, the amplitude of the uncertainty signal is greater, especially for the first few samples. The minimum of $\theta(t)$ occurs at -7.6213 , quite close to the initial lower bound (-10). As usual, a trade-off between



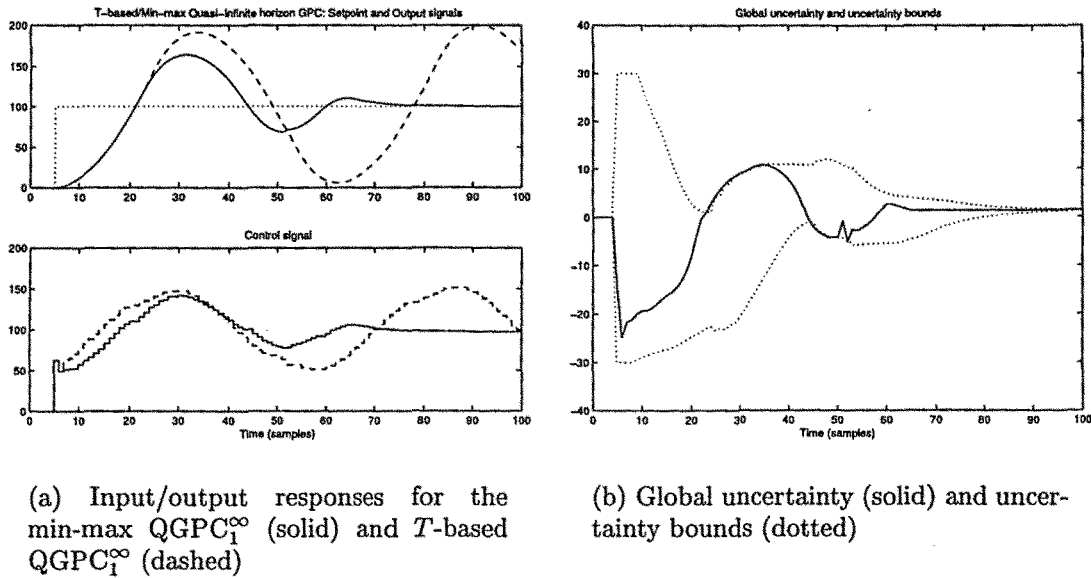
(a) Input/output responses for the min-max QGPC₁[∞] (solid) and output constraint (dotted)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.20: Closed-loop behaviour of the min-max QGPC₁[∞]

robustness (cautiousness) and performance must be reached. In fact, if lower initial bands are chosen, *e.g.* $\theta^- = -8$ and $\theta^+ = 8$ the bounds are violated shortly after the setpoint change, but the band updating procedure responds by widening the bands and avoiding future violations.

The next experiment has been carried out to show that the advantages of the min-max approach do not limit to the constrained case. The setpoint has been set to $w(t) = 100$ for $t \geq 5$, and no output constraints have been scheduled. The same controllers have been tested, setting the initial band values as $\theta^-(5) = -30$ and $\theta^+(5) = 30$ (a 30% of the setpoint change). The results are shown in Fig.4.21, where it can be observed that the closed-loop system with the T -based controller is unstable, whereas the min-max QGPC₁[∞] leads to a stable closed-loop system. Note that the min-max controller is not properly tuned, since the peak distance in $\theta(t)$ is greater than $M_\theta = 10$. In this example, this setting of M_θ does not cause major problems because no output constraints have been enforced, and the band updating algorithm manages to enclose the uncertainty signal despite the violations. Better results would arise with both the



(a) Input/output responses for the min-max QGPC_1^∞ (solid) and T -based QGPC_1^∞ (dashed)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.21: Closed-loop behaviour of the min-max QGPC_1^∞ and the T -based QGPC_1^∞

T -based schemes and the min-max methods if the tuning settings were chosen in a more conservative way, *e.g.* $T = (1 - 0.9q^{-1})^2$ and $M_\theta = 20$.

The experiments provided through this section illustrate that the tuning guidelines reported in Section 4.3.3 have proved useful and meaningful. The newly developed min-max controllers are shown to be a powerful control strategy not only for uncertain LTI systems, but also when the uncertainty comes from (difficult) non-linearities. In addition, this technique is able to meet the user designed constraints if the band updating procedure is well-tuned, even when disturbances and serious modelling errors occur. Notice that this feature is not possessed by the classical T -based controllers, which often involve constraint violations when system uncertainty takes place. When the control aim is to *stabilise* an uncertain system, examples can be provided for which the classical robust control synthesis methods (including the T -based MPC) overcome the min-max MPC based on the global uncertainty approach. However, when state/output constraint handling is a priority and uncertainty cannot be neglected, the min-max MPC strategy stands out as a more convenient choice since the classical MPC methods often


fail to accomplish the constraints. Therefore, the min-max methods depicted in this chapter yield a systematic methodology to synthesise robust MPC controllers when constraint handling is an essential control requirement.

4.4.4 Limitations of min-max MPC

This section presents a few examples for which the behaviour obtained with min-max controllers shows some unexpected characteristics. For example, as illustrated in Fig.4.10, the input/output responses obtained with min-max MPC controllers in the constrained case can exhibit large dead-times in some situations. These delays appear as a consequence of the tuning settings, and can be suppressed by choosing more appropriate tuning knobs. However, tuning settings which provide very similar nominal responses with the QGPC₁[∞] can lead to quite a different behaviour when the controller is implemented within the min-max framework.

As an example, consider the non-linear plant introduced in the previous section. Now the experiment which provides the outcome of Fig.4.19 is repeated changing just a single tuning knob: the control effort weighting ρ is reduced from 5 to 1. The rest of the tuning parameters remain unchanged, *i.e.* $N_u = 5$, $\theta^-(5) = -15$, $\theta^+(5) = 15$, $M_\theta = 10$ and $\mu = 0.9$. The setpoint changes from 0 to 55 at time $t = 5$ samples, and the output constraint $y(t) \leq 68$ is enforced. With these setup, it is expected that the lower penalty in the control moves leads to greater control moves and, consequently, the output response is expected to reach the setpoint faster than for the example shown in Fig.4.19.

The outcome of this experiment is shown in Fig.4.22. Surprisingly enough, the output exhibits an inverse response for the first few samples, what is absolutely unexpected since the nominal system is minimum-phase. This example is analogous to that shown in Fig.4.10, in which the dead-time is caused by the min-max controller and not



(a) Input/output responses for the min-max $QGPC_1^\infty$ (solid) and output constraint (dotted)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.22: Closed-loop behaviour of the min-max $QGPC_1^\infty$

by the system. Now, why a lower control move penalty produces an inverse response? The reason for such a behaviour is the fact that, at each sampling instant, a control move vector of dimension N_u is computed, and thus all the elements of $\Delta \mathbf{u}$ (not only the first one) contribute to the minimisation of the cost function. In this example the greater control activity allowed by $\rho = 1$ extends to all five ($N_u = 5$) elements of the control move vector and not only to the first one. At time $t = 5$, the *predicted* errors from $t + 1$ to $t + N$ are, in general, lower for $\rho = 1$ as expected, but the 1-step-ahead predicted tracking error ($t + 1$) is lower for $\rho = 5$. The inverse response is caused by the controller to be able to meet the output constraint, since the greater control moves produced by $\rho = 1$ lead the output predictions closer to the setpoint, and thus closer to the constraint. These considerations explain the apparent contradiction between the results shown in Fig.4.20 ($\rho = 5$) and 4.22 ($\rho = 1$). This behaviour is by no means specific to 1-norm formulations, but general to min-max MPC methods based on the global uncertainty approach. An almost identical result (inverse response) arises with the min-max 2-norm GPC^∞ if the tuning knob $\rho = 0.1$ is chosen whereas the other settings are identical to those used for the min-max $QGPC_1^\infty$. This behaviour might

be redressed by re-tuning the controller with a time varying weighting sequence $\rho(j)$ to pose a greater penalty on further control moves, which preserves nominal stability as remarked *e.g.* in Theorem 2.6. An alternative is the use of a time varying tracking error weighting $\mu(j)$ to pose a greater penalty on the first few predicted errors, but this choice affects the stability proof for the nominal case and, in addition, $\mu(j)$ are not used in (quasi) infinite horizon formulations (see Chapter 2).

This example, together with that shown in Fig.4.10, points out some drawbacks of the min-max methods described in this chapter. Unexpected closed-loop behaviour such as dead-times or inverse responses may appear with min-max controllers when (output) constraints are considered. As discussed in (Scokaert and Mayne, 1998), these peculiarities might be caused by the difficulty to handle an uncertainty polytope of 2^{N_θ} vertices with a single control move sequence Δu . For the example provided in this section, a 5-dimensional control move vector is expected to cope with $2^6 = 64$ different polytope vertices, what can be quite a challenging task for the controller. The next chapter exploits the feedback formulations of min-max controllers (Scokaert and Mayne, 1998) to keep the ratio between the degrees of freedom (control moves) and the polytope dimension 2^{N_θ} closer to unity. This alternative helps to avoid such an undesirable behaviour.

4.4.5 Self-tuning min-max QGPC₁[∞]

In this section the self-tuning algorithm suggested in Section 4.3.3.1 is tested against the non-linear benchmark system used in the previous few examples. The adaptive scheme is compared to a fixed design of the parameters M_θ and μ .

For this experiment the setpoint changes from 0 to 120 (which implies that the input to the linear block is saturated for quite long) at time $t = 5$ samples, and no constraints have been taken into account. The fixed min-max QGPC₁[∞] has been tuned




Figure 4.23: Closed-loop behaviour of the min-max QGPC₁[∞]

using $[N_u, \rho] = [5, 1]$, $\theta^-(5) = -30$, $\theta^+(5) = 30$, $M_\theta = 10$ and $\mu = 0.9$. With these settings the closed-loop behaviour is as shown in Fig.4.23, where it can be observed that the input/output responses are quite oscillatory, and the uncertainty signal surpasses the lower and upper bounds several times. The reason for such a behaviour is that the peak distance in the uncertainty signal is greater than $M_\theta = 10$, and the pole $\mu = 0.9$ of the first-order filter is not “slow enough” to avoid uncertainty band violations.

The self-tuning min-max QGPC₁[∞] has been designed using $M_\theta^{\min} = 10$, $\mu^{\min} = 0.7$ and $\mu^{\max} = 0.975$. With these settings the past uncertainty horizon is always greater than or equal to 10, *i.e.* in the worst case it would be at least as large as the fixed design presented above. On the other hand, the pole of the first-order filter (eqn.4.18) is constrained to be between 0.7 and 0.975. Notice that the fixed design $\mu = 0.9$ is included within this interval. Initially (at time $t = 5$ samples), M_θ and μ are chosen identical to the fixed design considered above, namely $M_\theta = 10$ and $\mu = 0.9$.

The result of this experiment is shown in Fig.4.24. The self-tuning algorithm adapts the values of μ and M_θ in such a way that the bands are violated fewer times compared

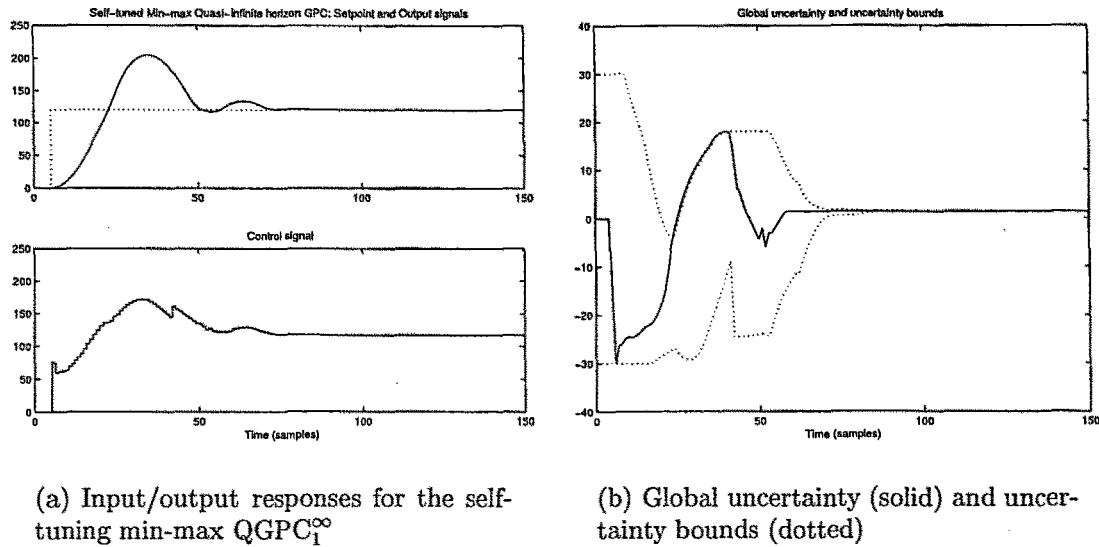


Figure 4.24: Closed-loop behaviour of the self-tuning min-max QGPC_1^∞

to the fixed design. As a consequence of the use of the self-tuning method, M_θ is modified on-line and the controller is more cautious. Therefore the input/output responses are much softer, with fewer oscillations, and the output settles down earlier than for the fixed design. However, as discussed in Section 4.3.3.1, the application of this method in practical control problems would require a pre-processing of the uncertainty signal $\theta(t)$ in order to get rid of noise and high-frequency oscillations. These two phenomena have been neglected in the example provided in this section.

4.4.6 The min-max approach versus non-linear MPC

In this section, an example is provided to show that the min-max linear MPC solution involves much less computational burden than non-linear MPC controllers. The CPU time of the latter methods can be several orders of magnitude larger than that of the min-max approach. Needless to say, non-linear MPC usually leads to better performance (if the nominal and true systems are identical), but the difference in CPU time is so large that the application of non-linear model-based techniques is often confined to very slow processes.

The experiments presented below have been performed using the non-linear benchmark system of Fig.4.17. This system can be described using the equations:

$$\begin{aligned} x &= \min \{ \max \{ -2, u - y \}, 2 \}, \\ \ddot{y} + \dot{y} &= 200x. \end{aligned} \tag{4.19}$$

Now, assume that this model and the true system to be controlled are exactly the same. Throughout this section no disturbances of any kind are considered. A non-linear (2-norm) finite horizon predictive controller (NLPC) can be formulated as the minimisation of the cost function

$$J_2(t) = \sum_{j=1}^{N_y} \mu(j) [w(t+j|t) - y(t+j|t)]^2 + \sum_{j=1}^{N_u} \rho(j) \Delta u^2(t+j-1|t),$$

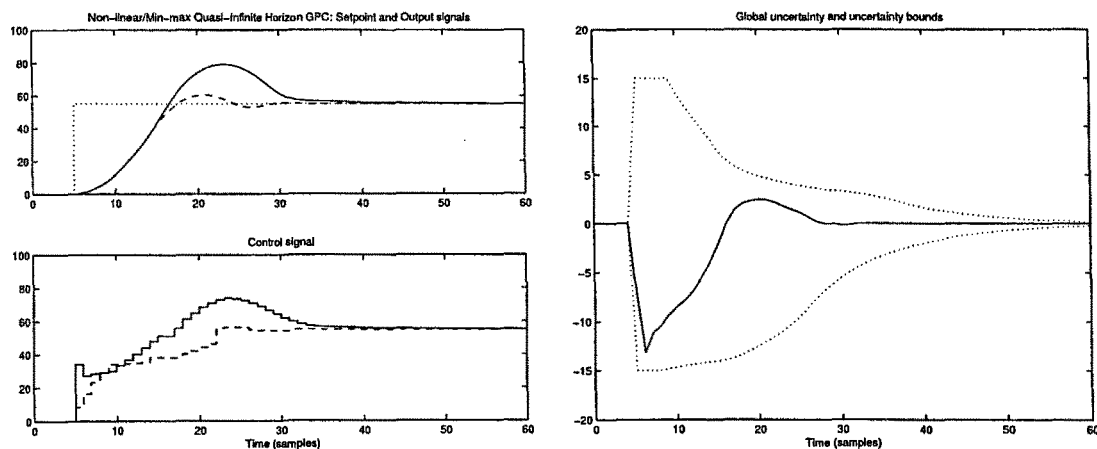
where the predictions $y(t+j|t)$ are computed, at each sampling instant, integrating the non-linear *Ordinary Differential Equations* (ODE) system provided by eqn.4.19. A Runge-Kutta 4-5 formulation (the Dormand-Price method) (The Mathworks, 1998) has been used to compute the output predictions. A finite prediction horizon ($N_y < \infty$) has been considered in this section. This choice does not guarantee stability, but the non-linear system can be easily stabilised with the NLPC since exact predictions are available, and thus no complications arise with the finite horizon strategy.

An SQP method has been used to minimise the cost function. Notice that, in this case, a single cost function evaluation is very time-consuming, since the differential equations need be integrated to obtain the output predictions. In addition, the gradient of $J_2(t)$ with respect to the control move vector $\Delta \mathbf{u}(t)$ cannot be obtained analytically, and thus the SQP method cannot be speeded up using (exact) gradient data.

The results obtained with both the NLPC and the QGPC₁[∞] are provided below. The min-max controller uses the internal model of eqn.A.10, which has been obtained neglecting the saturation. The tuning settings $[N_u, \rho] = [5, 5]$ and a sampling time of $T_s = 0.05$ seconds (the same used in the previous sections) have been chosen for both controllers. In addition, the NLPC has been tuned with $N_y = 15$ and $\mu = 1$, whereas

the min-max QGPC₁[∞] uses $M_\theta = 10$, $\mu = 0.9$, $\theta^-(5) = -15$ and $\theta^+(5) = 15$. For this experiment, the setpoint changes from 0 to 55 at time $t = 5$ samples.

With these conditions, a 100-sample simulation takes 2.0361 *hours* for the NLPC in a 400 MHz computer, whereas the min-max QGPC₁[∞] takes 23.6640 *seconds* using an interpreted LP algorithm, and 4.2760 seconds with a compiled routine. Thus the simulation with the non-linear MPC scheme takes more than 300 (1700 if a compiled LP routine is used) times the CPU time required by the min-max linear model approach. Notice, also, that this non-linear system is quite simple: a second-order ODE. Therefore, the computational burden would become enormous for a real MIMO process described by tens or hundreds of differential equations which must be integrated on-line.



(a) Input/output responses for the min-max QGPC₁[∞] (solid) and the NLPC (dashed)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 4.25: Closed-loop behaviour of the min-max QGPC₁[∞] and the NLPC

Fig.4.25 compares the closed-loop behaviour obtained with the NLPC and the QGPC₁[∞]. It is worth pointing out that, although the performance obtained with NLPC is somewhat finer than the one provided by the min-max controller (the overshoot is lower and the settling time is shorter), the closed-loop behaviour accomplished with

the min-max QGPC₁[∞] is quite acceptable, especially taking into account that this solution involves a much lower computational burden. In addition, it must be taken into account that the overshoot can be reduced by enforcing an output constraint with no (or negligible) increase in the CPU time.

This example highlights that the performance obtained with the min-max approach can be comparable to that provided by a non-linear MPC controller although the former requires much less computations. Thus the min-max controllers presented in this PhD thesis appear as a convenient candidate to handle non-linearities if fast dynamics are involved. In such a case, the non-linear MPC cannot be considered as a suitable alternative because of the enormous computational load. Obviously, if the non-linearities are of such a kind that they cannot be “represented” as a linear model plus some uncertainty, min-max MPC would not provide with an appropriate solution, and thus non-linear MPC or any other technique (*e.g.* feedback linearisation) should be used.

4.4.7 Global uncertainty versus multi-model descriptions

In this section the global uncertainty approach is contrasted with the polytopic multi-model linear plant description suggested in (Kothare *et al.*, 1996). The multi-model approach describes the plant family as the convex hull of several linear models $\mathcal{G} = \text{Co}\{G_1, G_2, \dots, G_L\}$, and the true (unknown) system is assumed to be a linear combination of the polytope vertices:

$$G_0 = \sum_{i=1}^L \lambda_i G_i,$$

where $\lambda_i \geq 0$ for $1 \leq i \leq L$ and $\sum_{i=1}^L \lambda_i = 1$. As already remarked in Section 4.1, in the input/output model framework, this kind of description can be used when uncertainty affects the numerator only, since pole uncertainty cannot be represented using this formulation. The optimal control move vector is obtained as the solution to the min-

max problem

$$\Delta \mathbf{u}^{\text{opt}}(t) = \arg \min_{\Delta \mathbf{u}} \max_{G \in \mathcal{G}} J(t),$$

where $J(t)$ is the cost function. As discussed in (Kothare *et al.*, 1996), it is enough to consider the worst case only at the polytope vertices $G_i \in \mathcal{G}$.

For GPC-like controllers based on input/output models, each vertex G_i is a discrete-time transfer function:

$$G_i = \frac{q^{-1}B_i(q^{-1})}{A_i(q^{-1})}.$$

In this section, the finite horizon (2-norm) GPC cost function (eqn.2.4):

$$J_2(t) = \sum_{j=N_1}^{N_2} \mu(j) [w(t+j|t) - y(t+j|t)]^2 + \sum_{j=1}^{N_u} \rho(j) \Delta u^2(t+j-1|t).$$

is considered. Note that an infinite horizon approach ($N_2 = \infty$) is difficult in this case, since the equality constraints on the unstable part of the output cannot be enforced *for all the vertices* of \mathcal{G} with a single control move vector. This could be a serious inconvenient, because nominal stability for finite horizon controllers is not guaranteed (see Chapter 2). The state-space methods depicted in (Kothare *et al.*, 1996) solve this difficulty and provide with an upper-bound solution of the infinite horizon problem with stability guarantees.

Now, for each vertex G_i , the cost function $J_2(t)$ can be written in the standard form:

$$J_2^i(t) = \Delta \mathbf{u}^T \mathbf{A}_i \Delta \mathbf{u} + \mathbf{b}_i^T \Delta \mathbf{u} + c_i,$$

with

$$\mathbf{A}_i = \mathbf{G}^{*T} \mathbf{M}^* \mathbf{G}^* + \mathbf{R},$$

$$\mathbf{b}_i^T = 2(\mathbf{w}^* - \mathbf{f}^*)^T \mathbf{M}^* \mathbf{G}^*,$$

$$c_i = (\mathbf{w}^* - \mathbf{f}^*)^T \mathbf{M}^* (\mathbf{w}^* - \mathbf{f}^*),$$

where the dynamic matrix \mathbf{G}^* , the free response vector \mathbf{f}^* , the setpoint vector \mathbf{w}^* and the weighting matrix \mathbf{M}^* are defined as

$$\begin{aligned}\mathbf{G}^* &= [\mathbf{G}_1^T \quad \mathbf{G}_2^T]^T, \\ \mathbf{f}^* &= [\mathbf{f}_1^T \quad \mathbf{f}_2^T]^T, \\ \mathbf{w}^* &= [\mathbf{w}_1^T \quad \mathbf{w}_2^T]^T, \\ \mathbf{M} &= \text{diag}[\mu(N_1), \mu(N_1 + 1), \dots, \mu(N_2)],\end{aligned}$$

to match the formulation of (Clarke *et al.*, 1987), and where \mathbf{G}_1 , \mathbf{G}_2 , \mathbf{f}_1 , \mathbf{f}_2 , \mathbf{w}_1 , \mathbf{w}_2 and \mathbf{R} are defined in Section 2.2.1 (Chapter 2). Finally, notice that the matrix \mathbf{G}^* and the vector \mathbf{f}^* depend on the vertex G_i . For each polytope vertex, the step response must be computed to obtain \mathbf{G}^* , and free response predictions must be performed on-line to form the vector \mathbf{f}^* . In order to obtain these free response predictions, the model

$$A_i(q^{-1})y(t) = B_i(q^{-1})u(t-1) + \frac{1}{\Delta}\xi(t),$$

is used, where $\xi(t)$ is a zero-mean stochastic noise. This CARIMA model is identical to that used by the controllers presented in Chapter 2 if $T = 1$ is chosen, and the noise term models random step-like disturbances. This kind of model provides with offset-free setpoint tracking for (asymptotically) constant disturbances.

The min-max optimisation problem can now be solved using the non-linear programming formulation:

$$v^{\text{opt}}, \Delta \mathbf{u}^{\text{opt}} = \arg \min_{v, \Delta \mathbf{u}} v \text{ subject to } v \geq J_2^i(t) \text{ for } 1 \leq i \leq L,$$

and subject to the constraints associated to the controller. Again, the solution can be found by means of convex optimisation tools, such as SQP (see Section 4.2.2.3). The resulting controller is a min-max multi-model GPC, referred to as MGPC hereafter.

Remark 4.12 Using this solution, it is assumed that the system is time invariant within the coincidence horizon $t + N_1, \dots, t + N_2$. In other words, this non-linear

programming problem finds the min-max solution for any true plant within \mathcal{G} as far as it does not vary within the prediction horizon. To account for non-linearities and/or time variability, the worst case should be evaluated letting the plant vary among all the polytope vertices at each sampling instant in the future. Addressed in that fashion, the min-max problem usually becomes intractable, since the number of cases to take into account grows exponentially with the prediction horizon, opening up as an L -ary tree, where L is the number of polytope vertices (see above). An upper bound solution for state-space models using LMI optimisation is presented in (Kothare *et al.*, 1996), but those results are difficult to extend to transfer function formulations. The examples provided below, consider only uncertain time-invariant systems, and thus the SQP solution to the min-max optimisation problem formulated above suffices. $\square\square\square$

The min-max QGPC $_{1}^{\infty}$ and the MGPC have been tested on the uncertain system provided in Section A.3:

$$G_0(q^{-1}) = K \frac{0.25}{(1-\phi)} \frac{q^{-1}(1-\phi q^{-1})}{1-1.4q^{-1}+0.65q^{-2}},$$

where the gain $0.5 \leq K \leq 1.5$ and the zero $0.4 \leq \phi \leq 0.6$ are uncertain parameters. As discussed in Appendix A, a family $\hat{\mathcal{G}}$ which includes all the plants which satisfy this definition (and a few more) can be described by the convex hull of a four-vertex polytope:

$$\begin{aligned} G_1 &= 0.2083 \frac{q^{-1}(1-0.4q^{-1})}{1-1.4q^{-1}+0.65q^{-2}}, & G_2 &= 0.2083 \frac{q^{-1}(1-0.6q^{-1})}{1-1.4q^{-1}+0.65q^{-2}}, \\ G_3 &= 0.9375 \frac{q^{-1}(1-0.4q^{-1})}{1-1.4q^{-1}+0.65q^{-2}}, & G_4 &= 0.9375 \frac{q^{-1}(1-0.6q^{-1})}{1-1.4q^{-1}+0.65q^{-2}}, \end{aligned}$$

The “nominal” system has been chosen as the one obtained with $K = 1$ and $\phi = 0.5$ (eqn.A.6), *i.e.*

$$G_m(q^{-1}) = \frac{0.5000q^{-1} - 0.2500q^{-2}}{1 - 1.4q^{-1} + 0.65q^{-2}}.$$

First of all, the *nominal* QGPC $_{1}^{\infty}$ and finite horizon GPC and have been tuned to provide “similar” closed-loop behaviour. The tuning knobs $[N_1, N_2, N_u, \mu, \rho] =$

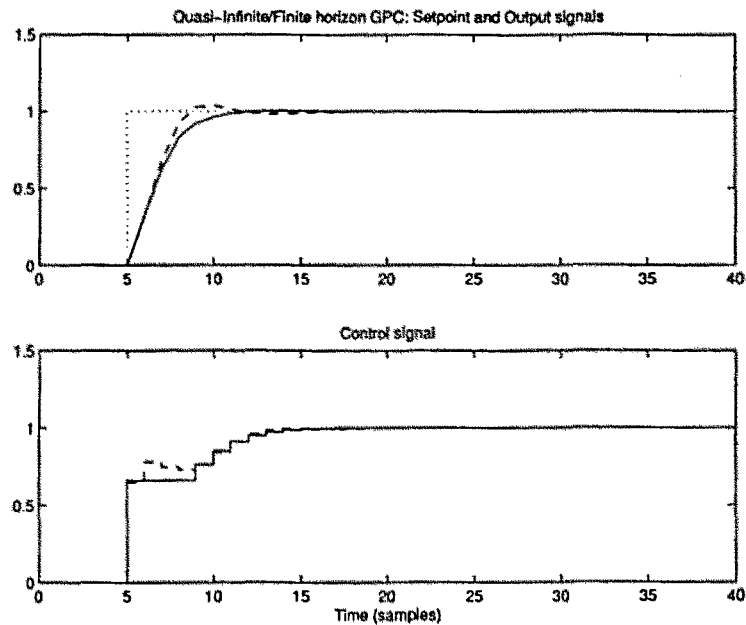


Figure 4.26: Nominal input/output responses for the GPC (dashed) and the QGPC_1^∞ (solid)

$[1, 40, 3, 1, 1]$ have been chosen for the GPC, whereas $[N_u, \rho] = [3, 1]$ have been used for the QGPC_1^∞ . With the tuning setting $N_2 = 40$, the 2-norm controller is (almost) indistinguishable from a truly infinite horizon GPC. As shown in Fig.4.26, the response obtained with the GPC is somewhat faster than the one provided by the 1-norm controller, and presents a low overshoot.

Now the min-max QGPC_1^∞ and the MGPC have been tested for a set of plants chosen among the family \mathcal{G} described in Section A.3. In these experiments, the setpoint changes from 0 to 1 at time $t = 5$ samples, and an additive disturbance $x(t)$ of magnitude 0.05 adds up to the system for $t \geq 51$, as described in eqn.4.17. The min-max QGPC_1^∞ has been tuned with $\theta^-(5) = -0.3$, $\theta^+(5) = 0.3$, $M_\theta = 10$ and $\mu = 0.9$. Fig.4.27 shows the closed-loop input/output responses obtained with both controllers. It is remarkable that the ones provided by the multi-model approach are quite similar for all the plants, whereas the dispersion of the responses obtained with the min-max QGPC_1^∞ is much larger. This feature is also observed in the experiments presented in (Camacho and Bordóns, 1995), where the advantages of using parametric

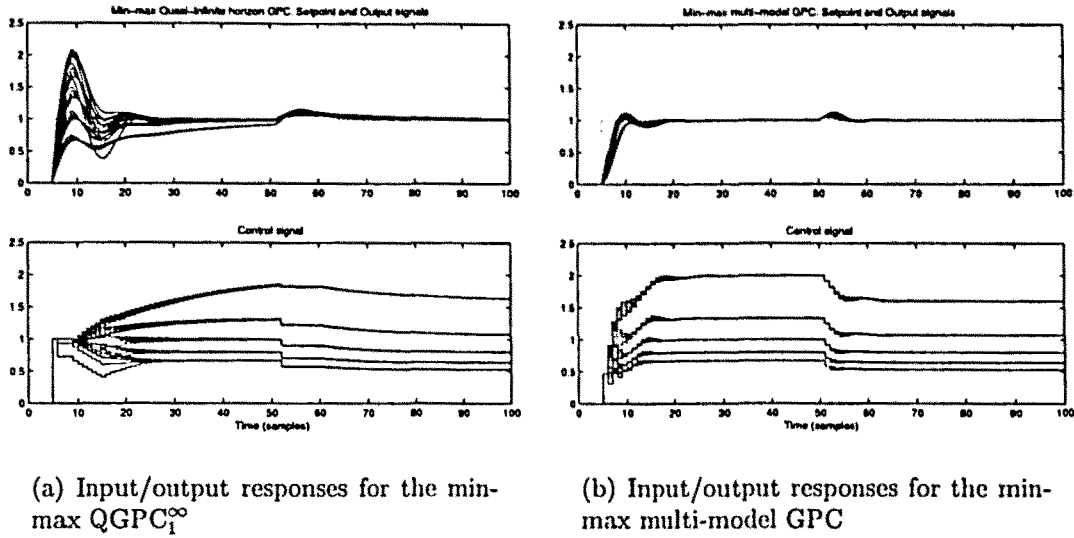


Figure 4.27: Closed-loop behaviour of the min-max QGPC₁[∞] and the min-max multi-model GPC

uncertainty are highlighted. In fact, it is not surprising that the the multi-model approach overcome the global uncertainty formulation in this case, since the multi-model description is much closer to the type of uncertainty presented in this example, whereas the global uncertainty description is much more general (*e.g.* it covers non-linearities). In addition, the MGPC is provided with *a priori* information about uncertainty at the design stage, whereas the min-max QGPC₁[∞] obtains *on-line* the uncertainty signal $\theta(t)$ and then responds according to the measurements. On the other hand, as computation time is concerned, the min-max QGPC₁[∞] takes, in average, 3.2379 seconds for a 100-sample simulation in a 400 MHz computer, compared to the 11.3800 seconds of the MGPC. That is, the latter involves 3.5 times the CPU-time required by the former.

Furthermore, it must be taken into account, that this example is quite favourable to the multi-model approach, since there are only four vertices, and the uncertainty affects the system's numerator only. In more typical situations, there could be (many) more than these four vertices, and thus the MGPC solution might become intractable, since the computation time would increase to a great extent. In addition, if uncertainty

affects the poles of the system, the MGPC solution is not possible, since a convex hull of linear plants never provides poles which are not present in the vertices. Moreover, if uncertainty is due to non-linearities and/or time variability, the MGPC solution cannot be implemented simply as the non-linear programming problem used here, since the predictions must be performed letting the plant vary among all the polytope vertices in the future. Besides, equality constraints (either on the unstable part or the whole output) are difficult to handle within the multi-model approach, and thus the way the stabilising CRHPC or GPC^∞ can be implemented in the multi-model context is an open question. Last, but not least, the MGPC solution is based on non-linear programming, compared to the simple LP approach of the min-max QGPC_1^∞ .

The solution presented in (Kothare *et al.*, 1996) solves some of these problems by obtaining an upper-bound solution based on LMI optimisation. However, those methods are based on state-space models and solve the regulation problem of driving the state to zero. If the more common setpoint tracking objective is pursued the methods described by Kothare *et al.* (1996) are restricted to LTI uncertainty.

In conclusion, the global uncertainty approach remains as the most appealing alternative to handle all kinds of uncertain systems within the GPC-like family. However, whenever a multi-model approach is possible (LTI uncertainty which affects the numerator only) with not too many vertices, the MGPC can become a convenient solution since it would lead, in general, to improved performance compared to the min-max global uncertainty methods.

4.4.8 Comparative study of \mathcal{H}_∞ and min-max MPC

In this section, the min-max predictive controllers described throughout this chapter are tested against a classical robust control design, namely an \mathcal{H}_∞ controller. The benchmark system chosen for this comparative analysis is the linear plant with gain,

zero and pole uncertainty provided in eqn.A.7 of Section A.4:

$$G_0(q^{-1}) = K \frac{(1 - \eta)(1 - \bar{\eta})}{(1 - \phi)} \frac{q^{-1}(1 - \phi q^{-1})}{(1 - \eta q^{-1})(1 - \bar{\eta} q^{-1})},$$

where the steady-state gain K lies in an interval centred at the nominal value $K_m = 1$ with a radius of a 20% (0.2), the zero ϕ is located in an interval centred at the nominal value $\phi_m = 0.6$ with a radius of a 10% (0.06), and the complex-conjugate poles occur within the circles centred at the nominal value $\eta_m = 0.6261 + 0.3130j$ with a radius of a 10% ($0.1 |\eta_m| = 0.1 \cdot 0.7 = 0.07$). *i.e.* the plant family is defined as

$$\mathcal{G} = \{G_0 : K = (1 + \Delta_K)K_m, \phi = (1 + \Delta_\phi)\phi_m, \eta = (1 + \Delta_\eta)\eta_m, \},$$

with $|\Delta_K| \leq 0.2$, $|\Delta_\phi| \leq 0.1$, and $|\Delta_\eta| \leq 0.1$. The nominal system is then obtained for $\Delta_K = \Delta_\phi = \Delta_\eta = 0$ leading to (eqn.A.8):

$$G(q^{-1}) = \frac{q^{-1}B(q^{-1})}{A(q^{-1})} = \frac{0.5944q^{-1} - 0.3567q^{-2}}{1 - 1.2522q^{-1} + 0.4900q^{-2}}.$$

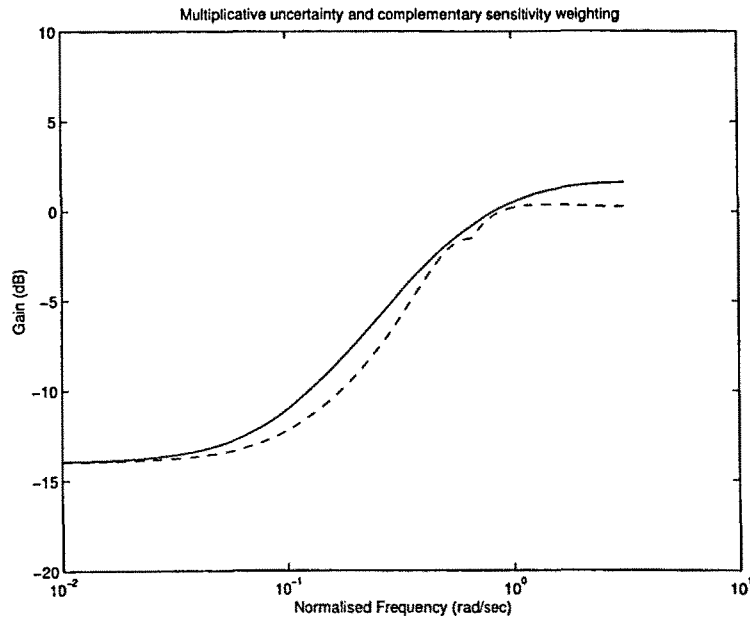


Figure 4.28: Multiplicative uncertainty bound (dashed) and complementary sensitivity weighting W_m (solid)

The maximum magnitude of multiplicative uncertainty (see Chapter 3) for this nominal system and the true plant family described above is shown in Fig.4.28 (dashed

line). A simple first order upper bound, displayed with a solid line, is provided by the transfer function

$$W_m(q^{-1}) = \frac{0.9692 - 0.8769q^{-1}}{1 - 0.5385q^{-1}},$$

which is used for the \mathcal{H}_∞ design as the complementary sensitivity weighting. Note, in Fig.4.28, that this choice of the weighting W_m is very close to the maximum magnitude of uncertainty. This is intended to avoid an excessively conservative design.

Now the robust stability condition for multiplicative uncertainty can be applied to design a controller which robustly stabilises the plant family described above. This condition, provided in Table 3.3, can be written in terms of the complementary sensitivity function:

$$\|T_u W_m\|_\infty < 1, \forall \omega_n \in [0, \pi].$$

In the SISO case, the complementary sensitivity at the output T_y can be used instead of T_u , since both are identical. The transfer function $T_y = GK(1 + GK)^{-1}$ can now be factorised as $T_y = GQ_T$ for $Q_T = K(1 + GK)^{-1}$, and thus the RS condition becomes

$$\|GQ_T W_m\|_\infty < 1 \Leftrightarrow \|Q_T\|_\infty < \left\| \frac{1}{GW_m} \right\|_\infty, \quad (4.20)$$

for all $0 \leq \omega_n \leq \pi$. Now, the filter Q_T can be designed to satisfy this inequality, since both G and W_m are known. Notice that Q_T is a factor of T_y , and thus the zeros and the poles of Q_T are viewed in the nominal closed-loop transfer function from the setpoint to the output ($T_y = GQ_T$) and in the nominal closed-loop transfer function from the setpoint to the input ($U_y = Q_T$). Therefore, slow (and of course unstable) poles and unstable zeroes must be avoided in Q_T . In addition, Q_T can be designed in order to cancel out the undesired open-loop dynamics of the nominal plant.

Once the transfer function Q_T is chosen, the controller K is found by solving the equation $Q_T = K(I + GK)^{-1}$ for K , which yields

$$K = \frac{Q_T}{1 - Q_T G}.$$

In order to achieve offset-free setpoint tracking for constant setpoints $w(t)$, the transfer function T_y must be 1 at low frequencies (condition ①) or, equivalently, the controller K must have a pole at 1 (condition ②). Needless to say, both conditions lead to the same result:

$$\textcircled{1} \quad T_y(1) = 1 \Leftrightarrow G(1)Q_T(1) = 1 \Leftrightarrow Q_T(1) = \frac{1}{G(1)}, \text{ or}$$

$$\textcircled{2} \quad 1 - Q_T(1)G(1) = 0 \Leftrightarrow Q_T(1) = \frac{1}{G(1)}.$$

Thus, the offset-free setpoint tracking requirement becomes a condition on the steady-state gain of Q_T . With all these guidelines, the filter Q_T can be chosen, on a pole-assignment basis, as

$$Q_T(q^{-1}) = \frac{B_Q(q^{-1})}{A_Q(q^{-1})} = \frac{1}{G(1)} \frac{B^*(1)A'_Q(1)}{A^*(1)B'_Q(1)} \frac{A^*(q^{-1})B'_Q(q^{-1})}{B^*(q^{-1})A'_Q(q^{-1})}, \quad (4.21)$$

subject to eqn.4.20, and where $B^*(q^{-1})/A^*(q^{-1})$ is a factor of $G(q^{-1})$ which consists of the dynamics (poles and zeroes) to be cancelled in the closed-loop transfer function, and $B'_Q(q^{-1})/A'_Q(q^{-1})$ are designed closed-loop dynamics. For open-loop unstable systems, this procedure cannot be directly applied since it is not possible to cancel out unstable poles using this approach.

The tuning settings for the min-max QGPC $_1^\infty$ and the \mathcal{H}_∞ controller have been chosen as follows. To begin with, the min-max QGPC $_1^\infty$ has been tuned using $[N_u, \rho] = [5, 1]$, $\theta^-(5) = -0.2$, $\theta^+(5) = 0.2$, $\mu = 0.7$ and $M_\theta = 10$. The setpoint changes from 0 to 1 at time $t = 5$ samples, and no disturbances have been taken into account (the comparison is made in terms of setpoint response only). The min-max QGPC $_1^\infty$ provides the nominal closed-loop input/output responses displayed in Fig.4.29(a). On the other hand, Q_T has been chosen as per eqn.4.21 with $A^* = A$, $B^* = 1$, $B'_Q = 1$ and $A'_Q = 1 - 1.1500q^{-1} + 0.3931q^{-2}$. Thus the open-loop poles are cancelled out and the closed-loop poles have been fixed at $0.5750 \pm 0.2500j$ in order to provide a

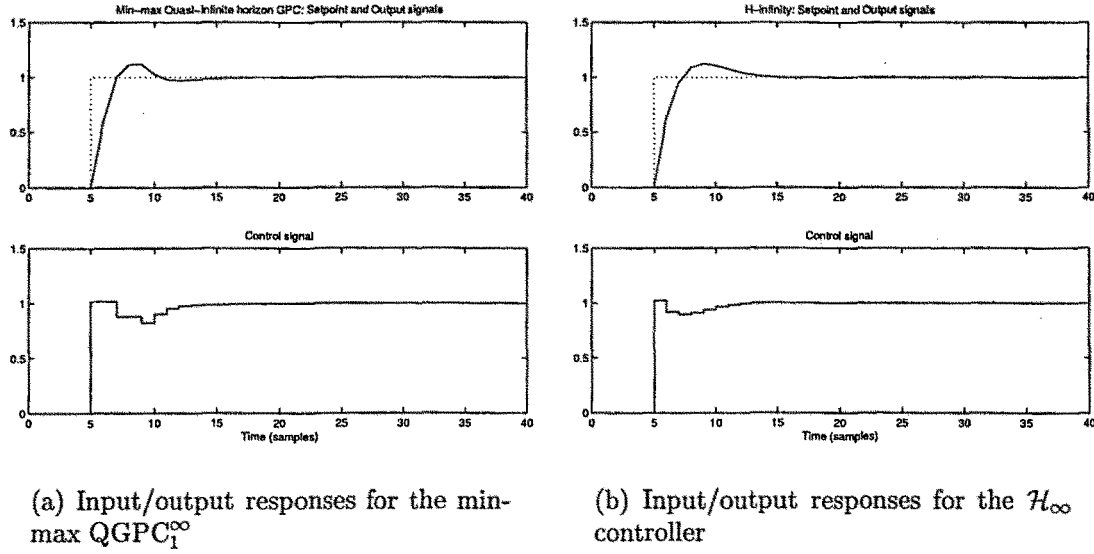


Figure 4.29: Nominal closed-loop behaviour of the min-max QGPC₁[∞] and the \mathcal{H}_∞ controller

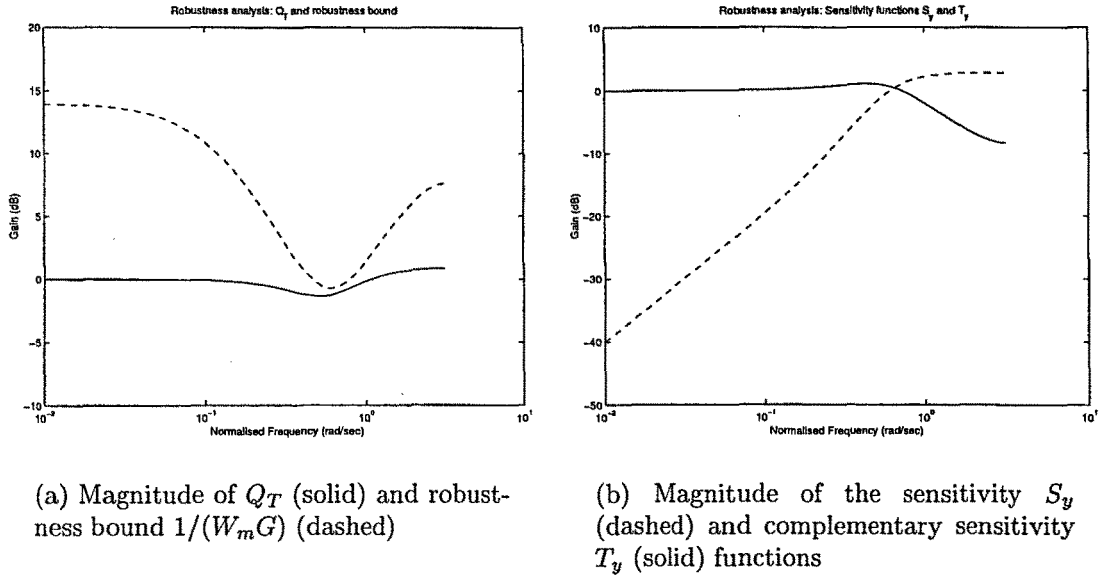
nominal closed-loop response which is similar to the min-max QGPC₁[∞], as displayed in Fig.4.29(b). Finally, the transfer function Q_T becomes

$$Q_T(q^{-1}) = \frac{1.0225 - 1.2804q^{-1} + 0.5010q^{-2}}{1 - 1.1500q^{-1} + 0.3931q^{-2}}$$

which yields the controller

$$K(q^{-1}) = \frac{1.0225 - 1.2804q^{-1} + 0.5010q^{-2}}{1 - 1.7578q^{-1} + 0.7578q^{-1}}.$$

As shown in Fig.4.30(a), this choice of Q_T satisfies the condition of eqn.4.20, and thus the resulting controller robustly stabilises the whole plant family \mathcal{G} . Notice that Q_T is quite close to the robustness bound, hence it is quite difficult to improve (nominal) performance preserving robust stability. The sensitivity S_y and complementary sensitivity T_y functions are shown in Fig.4.30(b). It is worth pointing out that S_y is slightly above the 0 dB line at high frequencies, producing an amplifying effect at that range. This could be inconvenient in case of high-frequency output disturbances $d_y(t)$, since these would be amplified at the output $y(t)$. However, since this example considers only the setpoint responses, this issue can be overlooked, though it could become relevant in real applications.


 Figure 4.30: Robustness analysis of the \mathcal{H}_∞ controller

Now, the true input/output responses obtained with the min-max QGPC $_1^\infty$ and the \mathcal{H}_∞ controller are compared in Fig.4.31. 250 true plants $G_j \in \mathcal{G}$ have been chosen, within the family described above, in such a way that the extreme cases (lowest and highest value of each parameter) are included in this experiment. Note that the envelope of the output responses is greater with the min-max QGPC $_1^\infty$ than with the \mathcal{H}_∞ controller. This may lead to think that the average performance provided by the min-max controller is lower than that obtained with the \mathcal{H}_∞ for this particular example. However, this is not the case. To show that the average performance is very approximately *the same* for both approaches, the following average performance index is defined:

$$J_{\text{perf}} = \frac{1}{N_p} \sum_{j=1}^{N_p} \sum_{t=1}^{n_t} [w(t) - y_j(t)]^2,$$

where N_p (250 in the example) is the number of true plants chosen within the family \mathcal{G} , n_t (50 in the example) is the simulation time, and $y_j(t)$ is the output response obtained with the plant G_j . This index takes into account the tracking errors only, but the control efforts might have been included as well. In this experiment, the

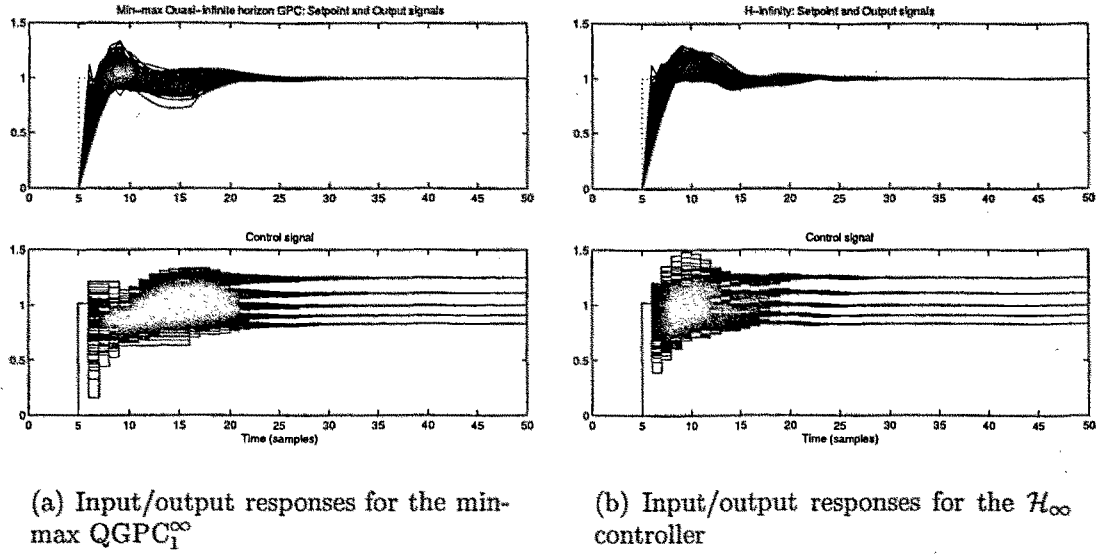


Figure 4.31: True closed-loop behaviour of the min-max QGPC₁[∞] and the \mathcal{H}_∞ controller performance index yields 1.2517 for the min-max QGPC₁[∞], and 1.2459 for the \mathcal{H}_∞ controller. Thus, the average performance is *almost identical* for both approaches. This is quite a remarkable result, since it must be taken into account that the \mathcal{H}_∞ controller is tightly tuned since:

1. The complementary sensitivity weighting W_m is very close to the upper bound of multiplicative uncertainty,
2. the transfer function Q_T has been designed such that it is quite close to the robustness bound $1/(W_m G)$ (see Fig.4.30), and
3. the \mathcal{H}_∞ controller has been chosen to provide nearly the same nominal input/output responses as the min-max QGPC₁[∞].

This implies that any attempt of improving robust performance with the \mathcal{H}_∞ controller can easily involve the violation of the robustness bound, and thus RS would no longer be guaranteed. In addition, if a better \mathcal{H}_∞ design were found, the min-max QGPC₁[∞] could be re-tuned to achieve similar (or better) performance.

In short, for this particular example, the min-max predictive controller does not only provide robust stability but, as performance is concerned, it closely parallels the results obtained with a tightly tuned \mathcal{H}_∞ controller. Taking into account that the key property of min-max MPC is constraint-handling, the result of the experiment presented in this section is aimed to illustrate that the behaviour of min-max MPC can be comparable to a classical robust control design, although min-max controllers are *not explicitly designed* to provide robust stability or robust performance *in the classical sense* (see Chapter 3).

4.5 Robustness analysis: the randomised approach

As pointed out above, no closed form exists for the min-max controllers introduced in this chapter, and thus it is quite a difficult issue to apply classical robustness analysis methods. In addition, the results obtained in Chapter 3 cannot be extended to the constrained case, since no LTI form of constrained MPC controllers exists, even though they are not implemented in the min-max framework. This section tackles the robustness analysis problem from the point of view of the statistical learning theory (Vidyasagar, 1997; Vidyasagar, 1998).

4.5.1 Fundamentals of the statistical learning theory

As discussed in (Vidyasagar, 1997; Vidyasagar, 1998), the statistical learning theory can be used to tackle the robust control analysis and design problems. The robustness analysis problem in the framework of the statistical learning theory can be formulated as follows:

Given a family of plants $\mathcal{G} = G(x)$ (parametrised in x) and a controller K , compute the probability p that the closed-loop behaviour is not “convenient”

where “convenient” can mean “stable” or can include other specifications, such as constraint satisfaction. Then K_1 is considered more robust than K_2 if the probability of inconvenient closed-loop behaviour for K_1 is lower than the one for K_2 .

It can become quite a difficult task to compute such a probability *analytically* for the whole plant family $G(x)$. An alternative is to *estimate* p using randomised methods. Consider a Bernoulli process, *i.e.* an experiment with two possible outcomes, namely “success” and “failure”. A coin-tossing is an appropriate example of such a process, for which “head” and “tail” can be taken as success and failure respectively (or conversely). It is widely known that the true probability of success in such a process is $p = 0.5$. If the coin-tossing experiment is repeated m times, the empirical probability \hat{p}_m can be computed as the ratio of the number of “heads” l among the m experiments. The greater m is chosen, the more probable is that $p \approx \hat{p}_m$, what can be loosely formulated as

$$\lim_{m \rightarrow \infty} \hat{p}_m = p.$$

This approach is known as Monte Carlo simulation.

The result “inconvenient closed-loop behaviour” can be thought of as the outcome of a Bernoulli process, and thus the following randomised procedure, detailed in (Vidyasagar, 1998), can be used to obtain an estimate \hat{p}_m of p :

1. Choose m plants $G(x_j)$ within the family $G(x)$, randomly generated according to some probability measure P_G .
2. Simulate the closed-loop system obtained with K and $G(x_j)$. If inconvenient closed-loop behaviour arises increase l .
3. Let the estimate or empirical probability \hat{p}_m be the fraction of plants among m which produce an inconvenient closed-loop behaviour, *i.e.* $\hat{p}_m = l/m$.

Now, given m , it is convenient to know how accurate is \hat{p}_m an estimate of p . In other words: what is the probability that $|\hat{p}_m - p| \leq \varepsilon$? The best answer currently available for this question is the Chernoff bound (Chernoff, 1952; Vidyasagar, 1998) according to which:

$$\text{prob}(|\hat{p}_m - p| \leq \varepsilon) \geq 1 - 2e^{-2m\varepsilon^2}, \quad (4.22)$$

where $\text{prob}(z)$ stands for the probability of the event z . Now, if one wishes to measure p with an accuracy ε and a confidence of $1 - \delta$, how many experiments need be performed? This is solved by using eqn.4.22:

$$2e^{-2m\varepsilon^2} \leq \delta,$$

or

$$m \geq \frac{1}{2\varepsilon^2} \log \left(\frac{2}{\delta} \right). \quad (4.23)$$

If m is chosen using this formula, it is possible to assert with confidence $1 - \delta$ that $|\hat{p}_m - p| \leq \varepsilon$. As an example (Vidyasagar, 1998), in order to be 99% ($\delta = 0.01$) sure that the empirical probability is within 0.05 ($\varepsilon = 0.05$) of the true value, it is enough to generate 1060 plants.

Notice that the bound provided in eqn.4.22 is only a sufficient condition, meaning that it may be possible to assert $|\hat{p}_m - p| \leq \varepsilon$ with probability higher than $1 - 2e^{-2m\varepsilon^2}$, or, in other words, than some m lower than the one provided in eqn.4.23 is usually enough. This can be easily checked with a coin-tossing process. After 1060 experiments the empirical probability of “head” has been found to be 0.4962, and the difference with the true probability is $3.7736 \cdot 10^{-3}$, quite lower than the 0.05 accuracy guaranteed with a 99% probability. This points out that the Chernoff bound can be quite a conservative approximation, but this is the tightest result currently available. In fact, the greater m becomes, the more conservative Chernoff bound seems to be.

As a matter of fact, the same theory can be used to show the conservativeness of the Chernoff bound. Consider the experiment: “compute the empirical probability \hat{p}_{m_1}

for the coin-tossing problem with $m_1 = 1060$ ". Now if $|\hat{p}_{m_1} - p| > 0.05$ call the result "success" and otherwise count it as "failure". Thus a new Bernoulli process has been defined, and let q be the probability of success for this new process. The Chernoff bound ensures that $q < 0.01$, and an estimate \hat{q}_{m_2} can be computed. If $\varepsilon = 0.005$ and a confidence level of 99% are chosen, the Chernoff bound yields $m_2 = 105967$. With these settings, the estimate $\hat{q}_{m_2} = 1.1985 \cdot 10^{-3}$ has been obtained. This seems to point out that the Chernoff bound is an order of magnitude (10^{-2} versus 10^{-3}) overly conservative for this particular example. Note that this result is only approximate to 0.005 of the true value (with a 99% confidence).

4.5.2 Robustness assessment using Monte Carlo simulation

The aim of this section is to how robust the min-max QGPC₁[∞] controller is compared to the T -based schemes through a set of randomised tests. This analysis has been performed only for predictive controllers because constraint satisfaction is one of the requirements. The unconstrained simulations are provided for completeness only. The parameters used to analyse the robustness of these controllers are, on the one hand, the empirical probability of obtaining an *unstable closed-loop system*, denoted as \hat{p}_m^u and, on the other hand, the empirical probability of obtaining *stable closed responses which satisfy the user-designed constraints*, referred to as \hat{p}_m^c . It is worth pointing out that a set of tuning knobs which provided robust stability for the *whole* plant family would thus be senseless for this comparison, since the empirical probability of instability would be 0 in either case. Hence the experiments have been carried out in such a way that the closed-loop performance obtained with the T -based and min-max controllers is similar.

The nominal model chosen in Section 4.4.1 is used here, and a 30% variation margin

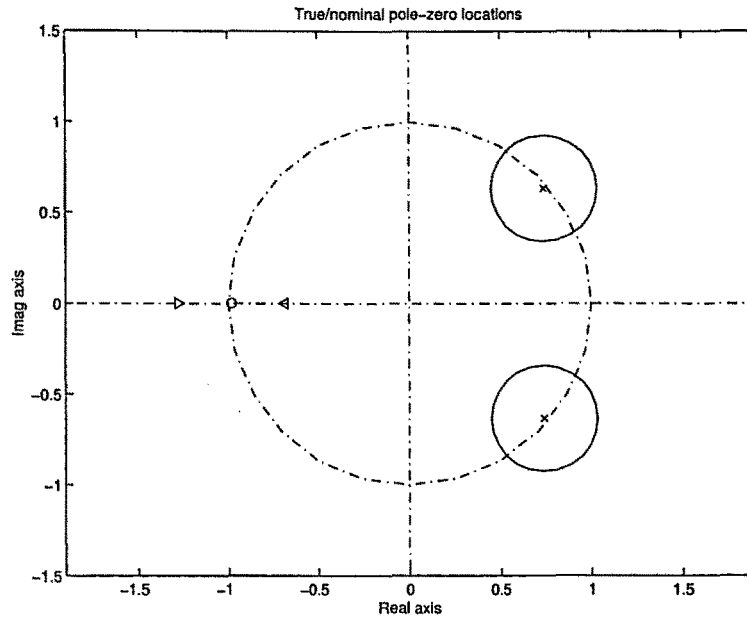


Figure 4.32: True and nominal pole-zero locations

has been considered to generate the true parameters. That is, given

$$G_0(q^{-1}) = K \frac{(1 - \eta)(1 - \bar{\eta})}{(1 - \phi)} \frac{q^{-1}(1 - \phi q^{-1})}{(1 - \eta q^{-1})(1 - \bar{\eta} q^{-1})},$$

the true steady-state gain K is assumed to vary in the interval $[0.7, 1.3]$ (the nominal steady-state gain is 1), the poles η and $\bar{\eta}$ are within a circle of radius $0.3 \cdot |0.7418 \pm 0.6333j|$ centred at $0.7418 \pm 0.6333j$ (the nominal poles), and the true zero lies in the interval $[-1.3 \cdot 0.9832, -0.7 \cdot 0.9832]$ (where -0.9832 is the nominal zero). In other words, $K = (1 + \Delta_K)K_m$ with $|\Delta_K| \leq 0.3$, $\eta = (1 + \Delta_\eta)\eta_m$ with $|\Delta_\eta| \leq 0.3$, and $\phi = (1 + \Delta_\phi)\phi_m$ with $|\Delta_\phi| \leq 0.3$, where the nominal values are given by $K_m = 1$, $\eta_m = 0.7418 + 0.6333j$ and $\phi_m = -0.9832$:

$$G(q^{-1}) = \frac{q^{-1}B}{A} = \frac{0.2358q^{-1} + 0.2319q^{-2}}{1 - 1.4835q^{-1} + 0.9512q^{-2}}.$$

The pole-zero loci of the nominal and true systems are shown in Fig.4.32, where “x” is used for the nominal poles, “o” is used for the nominal zero, the circles represent the true poles loci and the true zero lies between the “v” and “<” signs. Notice that the true plant family includes minimum phase, non-minimum phase, stable and unstable systems, and thus this example can be considered as quite a challenging benchmark

to measure the robustness of the min-max controllers introduced in this chapter. In addition, a disturbance of magnitude 0.05 has been added to the true system for $t \geq 51$ samples, *i.e.* the true output has been generated using eqn.4.17 with $x(t) = 0.05$ for $t \geq 51$.

It is worth pointing out that the robustness analysis methods depicted in Chapter 3 cannot be applied directly to this example, since no stable uncertainty description can be found to represent the whole plant family. More sophisticated techniques, such as the coprime factorisation, could be used (Skogestad and Postlethwaite, 1996) but they would only apply to the unconstrained case.

For this comparison, the tuning knobs $T_1 = 1 - 0.9q^{-1}$ and $\mu = 0.9$ have been chosen since they lead to similar responses, as disturbance rejection is concerned, for the T -based and the min-max controllers respectively (see above). In addition, the polynomial $T_2 = (1 - 0.9q^{-1})^2$ has been used to show that increasing the low-pass behaviour of the filter $1/T$ is not always convenient. In this case, including A as a factor of T leads to low robustness margins because the roots of A are very near the unit circle. For instance, the results obtained with $T = A(1 - 0.9q^{-1})$ are much poorer than the ones provided below.

Five predictive controllers (the GPC^∞ , the QGPC_1^∞ , each with T_1 and T_2 , and the min-max QGPC_1^∞), using the same nominal system, have been tested on 1060 plants randomly generated on the family described above. Hence the empirical probabilities computed are within 0.05 of the true value with a confidence of 99%. The true plants have been chosen within the intervals of a 30% parameter variation (poles, zero and gain, delay excluded), with uniform distribution. The tuning knobs of the controllers are $N_u = 5$, $\rho = 1$ and, for the min-max controller, $\theta^-(5) = -0.2$, $\theta^+(5) = 0.2$ (the setpoint change occurs at $t = 5$ samples) and $M_\theta = 10$.

The first experiment has been performed to assess the probability of obtaining an

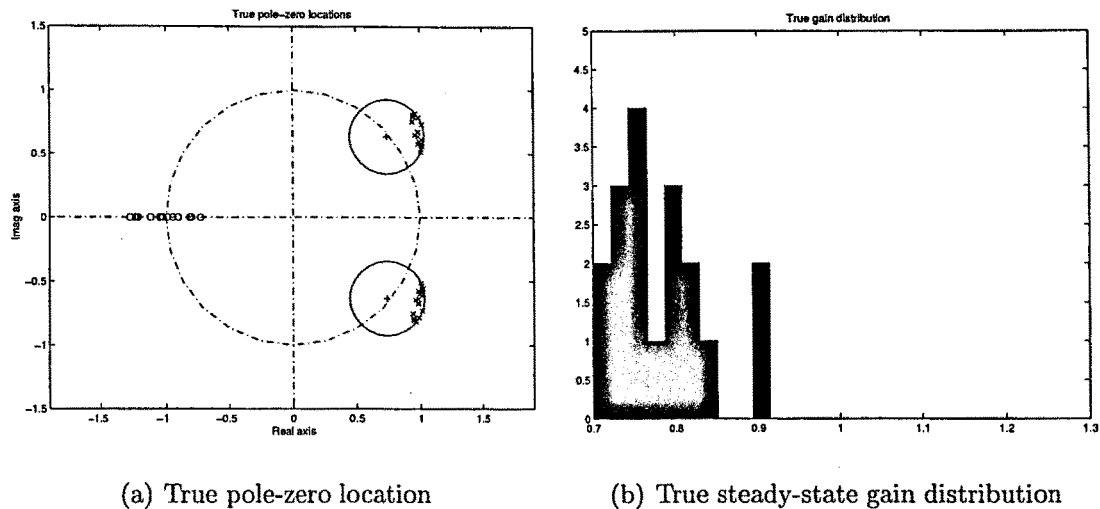


Figure 4.33: Plants leading to an unstable closed-loop system

that the true poles are almost on the boundary of the circle used to generate the true plant and that all these plants are open-loop *unstable* (whereas the nominal system is not). It is also remarkable that the true steady-state gain is always lower than the nominal value, which produces control moves which are lower than those required to stabilise the system. In addition, it is worth pointing out that if the nominal gain is closer to the true value the closed-loop systems turn to be stable with these plants too.

	GPC [∞]		QGPC ₁ [∞]		min-max QGPC ₁ [∞]
	T_1	T_2	T_1	T_2	
l^u	73	20	28	21	45
$\hat{p}_m^u = l^u/m$	0.0689	0.0189	0.0264	0.0198	0.0425
l^c	565	233	522	229	1007
$\hat{p}_m^c = l^c/m$	0.5330	0.2189	0.4925	0.2160	0.9500

Table 4.4: Randomised tests results (ii)

The randomised tests have been repeated enforcing an output constraint $y(t) < 1.2$ for all t . Table 4.4 shows the results of the experiments, where l^c is the number of *stable* closed-loop system for which the output constraint is respected, and \hat{p}_m^c is the empirical probability of stability *and* constraint satisfaction. It must be noticed that the incorporation of the constraint leads to a greater instability probability in the min-

	GPC [∞]		QGPC ₁ [∞]		min-max QGPC ₁ [∞]
	T ₁	T ₂	T ₁	T ₂	
l ^u	36	170	33	171	18
$\hat{p}_m^u = l^u/m$	0.0340	0.1604	0.0311	0.1613	0.0170

Table 4.3: Randomised tests results (i)

(which is identical to the GPC₁[∞] in the nominal case with these tuning settings) and the 2-norm GPC[∞] behave quite similarly. The QGPC₁[∞] seems to be a bit more robust than the 2-norm GPC[∞] with $T = T_1$ (other results of this experiment also point out that situation), but with $T = T_2$ the results are almost indistinguishable. This slight difference can be caused by the tuning knobs (N_u and ρ), since no indication exists to think that 1-norm controllers are intrinsically more robust than 2-norm counterparts. Needless to say, T_1 is a much suitable a choice than T_2 in this example, since it leads to much fewer unstable cases. Finally, it is worth point out that whenever a given true plant produces an unstable closed-loop system with the min-max controller, the closed-loop systems obtained with the T -based controllers are also unstable, *i.e.* the unstable cases obtained with the min-max QGPC₁[∞] are a subset of those of the T -based counterparts.

As an example, the min-max QGPC₁[∞] is unstable for the plant:

$$G_0(q^{-1}) = \frac{0.1177q^{-1} + 0.1440q^{-2}}{1 - 1.9733q^{-1} + 1.3054q^{-2}},$$

the poles of which are located at $0.9867 \pm 0.5761j$ (magnitude 1.1426), the zero at -1.2236 , and steady-state gain is 0.7881, *i.e.* the true system is unstable, non-minimum phase and the true steady-state gain is greater than the nominal value (1). In other words, this particular true plant provides an open-loop behaviour which is quite different from the nominal system.

Fig.4.33 shows the pole-zero and gain distributions for the 18 plants which lead to an unstable closed-loop system with all the controllers used in this experiment. Notice

unstable closed-loop system in the unconstrained case. Let p^u denote the probability of instability, l^u the number of unstable cases and \hat{p}_m^u the empirical probability computed as the ratio of unstable cases among the m experiments. Now a definition of “unstable closed-loop system” is needed. Here, a system is said to be stable in the *Bounded-Input/Bounded-Output* (BIBO) sense, *i.e.* if the system response to a bounded input (setpoint) is a bounded output. However, in simulation, the output of a system is always “bounded” in some sense (since no computer can handle infinitely large numbers). Here, the closed-loop system has been considered unstable if the maximum or minimum output values exceed some limits during the simulation time. Now, as the setpoint changes from 0 to 1, the system is taken to be stable if the output is bounded between -0.75 and 2.5 . Perhaps some stable closed-loop systems are counted as unstable with this approach, but performance would be so poor if these limits were not satisfied that they can be thought of as unstable anyway.

The results are shown in Table 4.3. Notice that most of the empirical probabilities are, in fact, lower than the accuracy parameter ($\varepsilon = 0.05$). If accuracy were increased up to (say) 0.005 , $m = 105967$ plants would be required, and performing that many simulations would take *a few days* with current computing capabilities. It is thus assumed that the values of Table 4.3 are, indeed, quite accurate, *i.e.* that the error between the true probabilities and the empirical ones is much lower $\varepsilon = 0.05$ (as occurs in the coin-tossing example presented above). In addition, there is no reason to think that the empirical probabilities computed are more “exact” for one controller than for another. Other simulations performed in the same conditions provide quite similar results.

The outcome of this experiment evidences that the min-max controller is, at least, as robust as the more robust T -based controller which has been found. In fact, if the results of Table 4.3 were exact, the min-max QGPC_1^∞ is about twice as robust as the T -based QGPC_1^∞ with $T = T_1$. It is also remarkable that the 1-norm QGPC_1^∞

max controller compared to the result of the unconstrained experiment (Table 4.3). The reason for such a behaviour is that, in order to satisfy the constraint for wide θ^- and θ^+ , the controller must lead the output to negative values, this makes θ^- and θ^+ increase and, thus, the output must be pushed even farther from the constraint. This process ends up causing instability. In fact, the probability of an unstable closed-loop system is higher than that obtained with most of the T -based approaches.

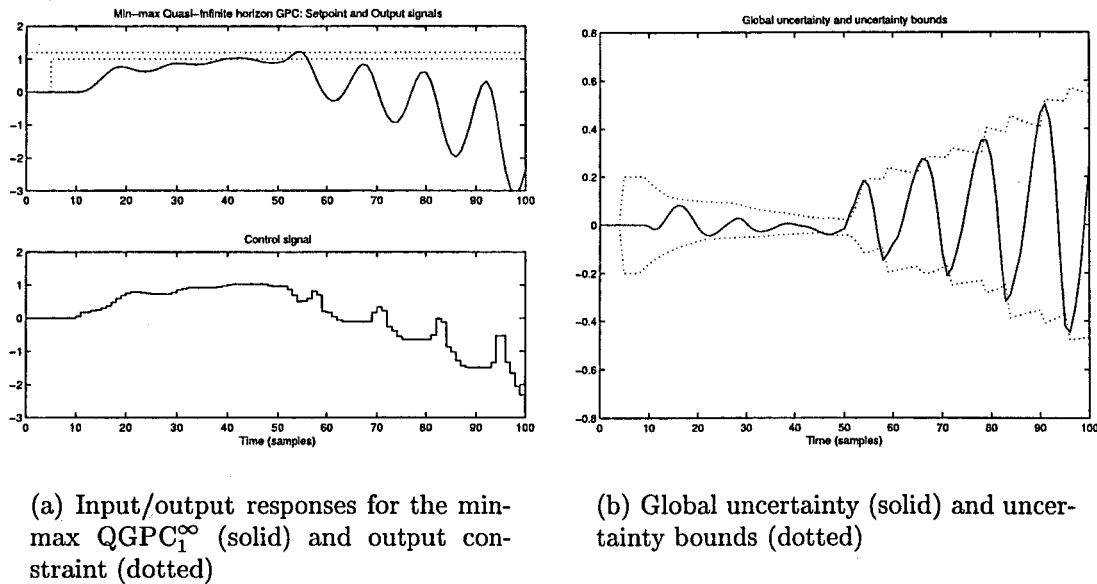


Figure 4.34: Closed-loop behaviour of the min-max QGPC₁[∞]

As an example of this peculiarity, the closed-loop behaviour of the min-max QGPC₁[∞] is shown in Fig.4.34 for the true plant

$$G_0(q^{-1}) = \frac{0.1352q^{-1} + 0.1461q^{-2}}{1 - 1.9734q^{-1} + 1.2657q^{-2}},$$

the poles of which are located at $0.9867 \pm 0.5405j$ (magnitude 1.1251), the zero at -1.0805 , and the steady-state gain is 0.9624. Apart from the initial delay which is due to the tuning settings, the closed-loop system is stable until the disturbance enters the system ($t = 51$), but the combination of the perturbation, the band updating algorithm, the output constraint, and the open-loop instability of the true plant, produces closed-loop instability. Notice that the oscillations of the global uncertainty signal increase

when the disturbance enters the system. Then the bands are widened and the output must be led further from the constraint. As a consequence, the uncertainty grows even more, the bands need be widened again, and so on. This situation might have been avoided with a more conservative band updating procedure. However, the sources of uncertainty are so extreme in this example (unstable, non-minimum phase plant versus stable minimum-phase nominal system) that such a behaviour is not expected in real control problems.

Now, as constraint satisfaction is considered, the 95% of the plants tested with the min-max controller have produced a stable closed-loop system for which the constraint is always respected, compared to the just a bit more than a 53% attained with the best T -based controller. The polynomial T_2 appears as quite an unsuitable choice in this respect, since the output breaks the constraint systematically whenever it is used. In addition, the maximum value of the output reached after the 1060 tests (taking into account the stable cases only) with the min-max controller is 1.2367 (quite close to the constraint) whereas with the T -based scheme the maximum value of the output has been found at 1.8731 and 2.3623 for the QGPC $_1^\infty$ with T_1 and T_2 respectively, and at 1.9085 and 2.3268 for the (2-norm) GPC $^\infty$ with T_1 and T_2 respectively. Notice that all of these values are unacceptably large, since the constraint specification is $y(t) < 1.2$ for all t . The T -based controllers can behave even worse. For instance, if the magnitude of the disturbance is 0.1 instead of 0.05, \hat{p}_m^c is still about 90% for the min-max QGPC $_1^\infty$, but it falls down to *less than 5%* (!) for the T -based schemes.

Moreover, the number of constraint violations has never exceeded 2 samples (out of 100) with the min-max scheme, whereas with the T -based controller there have been up to 31 constraint violations for a single experiment. The large probability of constraint satisfaction provided by the min-max QGPC $_1^\infty$ has been obtained at the price of a slightly lower stability results, since the T -based controllers are stable for a few more plants than the min-max scheme.

This situation might be improved by using feedback formulations of the min-max controllers (Scokaert and Mayne, 1998), as additional degrees of freedom are introduced to cope with different uncertainty sequences. This is particularly important when constraints are enforced. However, this solution does not lead to such a simple LP problem, as discussed in Chapter 5.

Finally, the outcome of these randomised experiments leads to the conclusion that the min-max QGPC₁[∞] suggested in this chapter, tested in quite a difficult environment, provides with a convenient solution in both the constrained and the unconstrained cases. The min-max approach overcomes the classical (well-tuned) T -based controllers, especially as constraint satisfaction, one of the most important issues in predictive control, is concerned.

4.6 Concluding remarks

In this chapter, MPC controllers based on min-max optimisation and a global uncertainty description are formulated and tested. To begin with, a CARMA model formulated in terms of a global uncertainty parameter (or signal) is provided in eqn.4.1, and the subsequent output predictions are derived using such a model. The min-max optimisation problem is undertaken for both 2-norm and 1-norm formulations. Taking into account the results of the robustness analysis for the unconstrained case provided in Chapter 3, this chapter is focused on infinite (or quasi-infinite) horizon controllers, since they provide nominal stability guarantees and their intrinsic robustness features are superior to other approaches such as the CRHPC.

The min-max optimisation problem for 2-norm controllers can be solved using either analytical or numerical methods, but the former tend to be untractable for usual settings ($N_u = 3$ or larger), and thus the latter is regarded as a more convenient alternative for practical applications. However, the numerical solution relies on non-linear

programming (SQP is used in this PhD thesis), and there are no *a priori* bounds of computation time. On the other hand, 1-norm min-max controllers are shown to lead to a simple LP solution which can be solved very efficiently with current optimisation tools. What is more, the CPU time required for a LP problem solution can be predicted (or bounded) prior to solving the problem. A comparative analysis of the computational burden has been performed, and the results show that the 1-norm approach is always more efficient.

The min-max optimisation using global uncertainties is solved assuming lower and upper bounds on the uncertainty signal, and the cost function is then minimised for worst case prediction. Since convex cost functions are used and the uncertainty regions (polytopes) are always convex, the maximum lies on a vertex of the polytope, and this property makes it much simpler the solution of the problem. In the existing methods which take this approach (Camacho and Bordóns, 1995; Scokaert and Mayne, 1998), it is assumed that the lower and upper uncertainty bounds are constant. In this thesis, the closed-loop behaviour of the uncertainty signal when modelling errors and/or disturbances occur is investigated. It is shown that the uncertainty signal oscillates whenever the setpoint changes or a disturbance enters the system and, later on, it settles down to some steady-state value. This property has been used to suggest a band updating algorithm which allows for the uncertainty bands to approach the measurements. Thus, the control algorithm is driven from cautious initial settings to more stringent counterparts. In addition, this approach guarantees that the (uncertainty-dependent) constraints are respected as far as the uncertainty bounds are not violated.

The novel methods, especially the min-max GPC^∞ and the min-max QGPC_1^∞ , have been extensively tested in quite tough situations, showing that they overcome the polynomial T approach when constraints are enforced. The classical T settings often fail not only to satisfy constraints, but even to provide a stable closed-loop system. In the unconstrained case, the min-max approach produces results which are comparable

(sometimes better) to the T -based controllers.

The min-max controllers are also tested on a strongly non-linear system. This approach provides with a remarkable trade-off between performance and computational burden, and again overcomes the solution obtained with the T -based controllers. As shown by means of an example, the min-max QGPC₁[∞] is several orders of magnitude faster than a non-linear MPC method, at the price of a lower performance.

The min-max methods have also been tested against robust multi-model controllers, such as those suggested by (Kothare *et al.*, 1996). It must be noticed that in transfer function formulations the min-max multi-model approach is difficult to apply, since pole-uncertainty cannot be easily handled. Nevertheless, a comparison with the global uncertainty approach is provided in this chapter, and it is shown that the multi-model methods yield better performance for a larger computational burden. However, the application of min-max multi-model methods for transfer function formulations is quite limited since uncertainty comes from many sources (strong non-linearities, changing poles) which cannot be handled by this approach.

A comparative analysis of the min-max predictive controllers formulated in this chapter with an \mathcal{H}_∞ controller is also provided. Although the min-max QGPC₁[∞] does not guarantee robust stability, it is shown to provide with almost identical results to those obtained with a tightly tuned \mathcal{H}_∞ controller. A performance index shows that both solutions are almost indistinguishable, which is quite a noteworthy achievement since the main feature of min-max methods is the ability to handle output and state constraints in spite of uncertainty and not to provide robust stability or performance in the classical sense.

A robustness analysis of the newly proposed controllers based on the statistical learning theory shows that these methods perform quite well on the most difficult situations, and appear as the most convenient choice when constraint handling in spite

of uncertainty is an indispensable control requirement. Again, the min-max controllers are shown to be at least as robust as the classical T -based formulations even when constraints are not enforced.

Finally, a few limitations of min-max MPC using global uncertainties are pointed out in this chapter. Some simulated experiments reveal that unexpected dead-times or inverse responses can occur. Although such a behaviour can be often redressed by re-tuning the controller, the reason for such a behaviour may stem from the difference between the degrees of freedom of the controllers and the amount of polytope vertices. This difference increases exponentially with the uncertainty horizon N_θ what makes it difficult to handle all the possible situations with a single control move vector. Among other considerations, the next chapter suggests some min-max controllers which overcome this difficulty.

Chapter 5

Trends and prospects of min-max RHPC

5.1 Introduction

As discussed in Chapter 4 (Section 4.4.4), there are some limitations or unexpected closed-loop characteristics related to the classical min-max MPC controllers based on the global uncertainty description. Scokaert and Mayne (1998) pointed out that some of the drawbacks related to min-max controllers are often associated to the use of a single control profile to handle all the possible disturbance (uncertainty) sequences.

As remarked in Scokaert and Mayne (1998), the standard min-max approach does not include the notion that feedback is present in the receding horizon implementation of the controller. It is also noted that the undesired closed-loop characteristics of the standard min-max MPC can be worse than the behaviour shown in Section 4.4.4, since infeasibility or even instability can occur due to the inability of the controller to handle, with a single control profile, all the potential uncertainty realisations. These considerations led to the formulation of a *feedback min-max* state-space MPC law which, for a simple first-order system subject to unmeasurable bounded disturbances, is shown to succeed where the classical min-max approach fails.

Consider, for example, the min-max QGPC₁[∞] as defined in the previous chapter. If $N_u = 4$ and $N = 5$, the single control move sequence

$$\Delta \mathbf{u}(t) = [\Delta u(t|t) \quad \Delta u(t+1|t) \quad \Delta u(t+2|t) \quad \Delta u(t+3|t)]^T,$$

is expected to cope with $2^5 = 32$ different uncertainty vertices (realisations):

$$\boldsymbol{\theta}_i^5 = [\theta(t+1|t) \quad \theta(t+2|t) \quad \theta(t+3|t) \quad \theta(t+4|t) \quad \theta(t+5|t)]^T,$$

with $\theta(t+k|t) = \theta^-$ or $\theta(t+k|t) = \theta^+$, and $\theta_i^5 \neq \theta_j^5$ if $i \neq j$ for all $1 \leq i, j \leq 32$. In addition, it must be taken into account that some of the degrees of freedom available in the control move vector $\Delta \mathbf{u}(t)$ are used to enforce the equality constraints on the unstable part of the output (see Chapter 4). In this example, if the nominal model is open-loop stable, a single equality constraint ($n_{\bar{a}} + 1 = 1$) must be enforced, and a component of $\Delta \mathbf{u}(t)$ is used to satisfy such an end-point constraint. In short, there are only *three* (four decision variables minus one equality constraint) available degrees of freedom in $\Delta \mathbf{u}(t)$, namely $\Delta u(t|t)$, $\Delta u(t+1|t)$ and $\Delta u(t+2|t)$, to cope with *thirty-two* different uncertainty realisations. Note that the classical min-max approach does not exploit the fact that new measurements of $\theta(t)$ will be available at time $t+1$, $t+2$ and so on. There is no real need of handling so many different possibilities with a single control profile.

In (Scockaert and Mayne, 1998) it is remarked that, apart from using different control profiles for different uncertainty realisations, a so-called “causality constraint” must be incorporated. Consider again the example provided above. The ideal approach would be to define 32 different hypothesised control move sequences, say $\Delta \mathbf{u}_i(t)$ with $1 \leq i \leq 32$, to cope with the 32 different uncertainty realisations. This would give $32 \times 5 = 160$ minus 32 (used to enforce the end-point equality constraints) degrees of freedom for optimisation, possibly much more than required. Now, if the problem were solved that way, there would be 32 different candidate first control moves to update the control signal. With no knowledge about the future uncertainty values, it would

be impossible to decide among that many candidates. The causality constraint is thus introduced to allow for just a single “next control move” for each possible internal process state, which depends on past input/output (or state) data only. In particular, this constraint ensures that all the hypothesised control move sequences start with the same first control move, since the current internal process state is obviously the same for all the different *future* uncertainty realisations. Needless, the feedback min-max approach involves a larger computational burden compared to the standard min-max methods, but the benefits as performance is concerned might justify the use of the feedback formulation.

In this chapter, the notion of feedback min-max MPC is exploited, and a feedback min-max (2-norm) GPC^∞ control law is suggested and tested. A feedback min-max (1-norm) QGPC_1^∞ using LP is also outlined. As shown below, if the 2-norm is used, the feedback min-max problem can be solved using non-linear programming tools, such as the SQP method. 2-norm cost functions are differentiable, and thus gradients (of both the cost function and the constraints) can be obtained. The use of gradients speeds up the convergence to the optimal point and provides the non-linear programming problem with better numerical properties. These gradients cannot be computed for 1-norm cost functions since these are not differentiable, and thus the SQP solution for the 1-norm case is not encouraged.

If a 1-norm cost function is used, such as that of the QGPC_1^∞ (eqn.2.48), it is possible to solve the optimisation problem by means of LP methods. However, in the feedback implementation of such a min-max controller, the number of variables and constraints would increase exponentially with the uncertainty horizon. A simpler solution, outlined in Remark 5.2, can be used to reduce the number of variables, but the number of constraints still depends on N_θ exponentially. The advantages and drawbacks of 1-norm implementations are similar to those of the standard min-max algorithms. That is, the CPU time can be known (or bounded) *a priori*, since LP tools

are applied, but the QGPC_1^∞ does not guarantee the stability of the nominal closed-loop system, though it is very unlikely that nominal stability problems occur. However, due to the number of constraints required by the LP problem, the computational advantages of feedback 1-norm implementations compared to 2-norm counterparts are not as significant as for the standard min-max approach.

Taking into account all these considerations, the next sections are focused on feedback min-max 2-norm implementations. The 1-norm case, which would lead to similar results, is not deeply analysed.

5.2 Feedback min-max GPC^∞

In this section the min-max GPC^∞ depicted in Chapter 4 is reformulated so as to apply different control profiles for different uncertainty realisations, subject to the causality constraint defined in Section 5.1.

As discussed above, if the settings $N_u = 4$ and $N = 5$ are assumed and the nominal model is open-loop stable, there are only three degrees of freedom in $\Delta \mathbf{u}(t)$ available to minimise the cost function, whereas the last control move, $\Delta u(t + 3|t)$ is needed to satisfy the end-point constraint. In such a case (Sokaert and Mayne, 1998), four different control move sequences, defined in terms of eleven decision variables, can be specified to face eight different potential future uncertainty realisations:

$$\textcircled{1} \quad \left. \begin{array}{l} \theta_1^3 = [\theta^- \ \theta^- \ \theta^-]^T \\ \theta_2^3 = [\theta^- \ \theta^- \ \theta^+]^T \end{array} \right\} \iff \Delta \mathbf{u}_1(t), \text{ with}$$

$$\Delta \mathbf{u}_1(t) = [\Delta u_1(t|t) \ \Delta u_1(t+1|t) \ \Delta u_1(t+2|t) \ | \ \Delta u_1(t+3|t)]^T,$$

$$\textcircled{2} \quad \left. \begin{array}{l} \theta_3^3 = [\theta^- \ \theta^+ \ \theta^-]^T \\ \theta_4^3 = [\theta^- \ \theta^+ \ \theta^+]^T \end{array} \right\} \iff \Delta \mathbf{u}_2(t), \text{ with}$$

$$\Delta \mathbf{u}_2(t) = [\Delta u_1(t|t) \ \Delta u_1(t+1|t) \ \Delta u_2(t+2|t) \ | \ \Delta u_2(t+3|t)]^T,$$

$$\textcircled{3} \left. \begin{array}{l} \theta_5^3 = [\theta^+ \ \theta^- \ \theta^-]^T \\ \theta_6^3 = [\theta^+ \ \theta^- \ \theta^+]^T \end{array} \right\} \iff \Delta \mathbf{u}_3(t), \text{ with}$$

$$\Delta \mathbf{u}_3(t) = [\Delta u_1(t|t) \ \Delta u_2(t+1|t) \ \Delta u_3(t+2|t) \ | \ \Delta u_3(t+3|t)]^T,$$

$$\textcircled{4} \left. \begin{array}{l} \theta_7^3 = [\theta^+ \ \theta^+ \ \theta^-]^T \\ \theta_8^3 = [\theta^+ \ \theta^+ \ \theta^+]^T \end{array} \right\} \iff \Delta \mathbf{u}_4(t), \text{ with}$$

$$\Delta \mathbf{u}_4(t) = [\Delta u_1(t|t) \ \Delta u_2(t+1|t) \ \Delta u_4(t+2|t) \ | \ \Delta u_4(t+3|t)]^T.$$

Note that these four control profiles start with the same first control move $\Delta u_1(t|t)$. This is a consequence of the causality constraint, since the past input/output data are obviously the same at this stage. However, at time $t+1$, $\theta(t+1|t)$ can be either (closer to) θ^- or (closer to) θ^+ , thus two possible control moves are considered at time $t+1$, either $\Delta u_1(t+1|t)$ for $\theta(t+1|t) = \theta^-$, or $\Delta u_2(t+1|t)$ for $\theta(t+1|t) = \theta^+$. This procedure is then repeated for the sampling instant $t+2$.

Now these four different hypothesised control move sequences, $\Delta \mathbf{u}_1(t)$, $\Delta \mathbf{u}_2(t)$, $\Delta \mathbf{u}_3(t)$ and $\Delta \mathbf{u}_4(t)$, can be defined in terms of an *extended control move vector*:

$$\begin{aligned} \Delta \mathbf{U}(t) = & [\Delta u_1(t|t) \ | \ \Delta u_1(t+1|t) \ \Delta u_2(t+1|t) \ | \\ & \Delta u_1(t+2|t) \ \Delta u_2(t+2|t) \ \Delta u_3(t+2|t) \ \Delta u_4(t+2|t) \ | \\ & \Delta u_1(t+3|t) \ \Delta u_2(t+3|t) \ \Delta u_3(t+3|t) \ \Delta u_4(t+3|t)]^T. \end{aligned}$$

Notice that the last four elements of the vector $\Delta \mathbf{U}(t)$ are used to enforce the equality constraint, and thus these do not contribute with additional degrees of freedom to handle different uncertainty realisations. Because of this, uncertainty vertices θ_i^3 consisting of three elements instead of four are considered above. Hence, if the last four decision variables are not taken into account (since these are needed to ensure the equality constraints), there is almost a one-to-one correspondence between the uncertainty realisations (8) and the degrees of freedom available for the controller (7). In general, there are 2^n uncertainty vertices and $2^n - 1$ control moves available for optimisation.

Therefore the ratio between different uncertainty realisations and degrees of freedom is always close to unity. This is a clear advantage compared to the classical min-max approach, for which there are $N_u - n_{\bar{a}} - 1$ degrees of freedom (control moves minus end-point constraints) to handle 2^{N_θ} uncertainty realisations, where $N_\theta \geq N \geq N_u + n_b - 1$ (see Chapter 2, eqn.2.14). That is, for standard min-max controllers the degrees of freedom increase linearly with N_u , whereas the polytope vertices increase exponentially with N_u . Hence, for standard min-max MPC methods, the greater the control horizon is chosen, the more difficult the optimisation task becomes.

Each different control profile can be selected from $\Delta U(t)$ by using a convenient *selection matrix* formed by zeroes and ones. For example, for $\Delta u_1(t)$ and $\Delta U(t)$ as defined above, it holds that

$$\Delta u_1(t) = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \Delta U(t).$$

In general, the control profile $\Delta u_i(t)$ can be obtained by means of a selection matrix S_i as $\Delta u_i(t) = S_i \Delta U(t)$.

For the general case, given N_u , $n_{\bar{a}} + 1$ (the number of end-point equality constraints) and N , the uncertainty horizon is defined as $N_\theta = N_u - n_{\bar{a}} - 1$, and the extended control move vector is formed by $2^{N_\theta} - 1$ control moves available for optimisation plus $2^{N_\theta - 1}$ control moves required to enforce the end-point equality constraints. Then $2^{N_\theta - 1}$ different hypothesised control move sequences are defined in terms of the extended control move vector. For the example introduced above, $N_\theta = 3$, $n_{\bar{a}} + 1 = 1$ and $N = 5$, there are $2^3 - 1 = 7$ control moves available for optimisation, whereas $2^{3-1} = 4$ control moves are used to enforce the end-point constraints, and thus the extended control move vector is formed by $7 + 4 = 11$ decision variables.

Note that, with this definition of N_θ , the uncertainty horizon is shorter than the prediction horizon N , and thus some hypothesis must be made about the future un-

certainty signal for $t + N_\theta + 1, t + N_\theta + 2, \dots, t + N$. In the example, since (output) predictions are performed up to $t + 5$ and these depend on the uncertainty signal from $t + 1$ to $t + 5$, uncertainty predictions for $t + 4$ and $t + 5$ are required, but the uncertainty vertices defined above consider only up to 3-step-ahead predictions. Here, it is suggested that the future global uncertainties from $t + N_\theta + 1$ to $t + N$ be assumed equal to the average of the lower and the upper uncertainty bounds, *i.e.*

$$\theta(t + j|t) = \bar{\theta} = \frac{\theta^- + \theta^+}{2}, \quad N_\theta + 1 \leq j \leq N,$$

which can be thought of as the “expected uncertainty value”.

Now the min-max GPC[∞] introduced in Section 4.2.2.3 can be reformulated so as to use different control sequences for different uncertainty realisations. As discussed in Chapter 4, the standard min-max GPC[∞] can be posed as

$$\overline{\Delta \mathbf{u}^{\text{opt}}} = \arg \min_{\Delta \mathbf{u}} \max_{\theta_i^{N_\theta} \in \Theta^{N_\theta}} J_2(t),$$

where $J_2(t)$ is the quadratic cost function:

$$J_2(t) = \left(\mathbf{w} - \mathbf{f} - \mathbf{G}\Delta \mathbf{u} - \mathbf{H}_\theta \theta_i^{N_\theta} \right)^T \Lambda \left(\mathbf{w} - \mathbf{f} - \mathbf{G}\Delta \mathbf{u} - \mathbf{H}_\theta \theta_i^{N_\theta} \right) + \Delta \mathbf{u}^T \mathbf{R} \Delta \mathbf{u},$$

subject to equality ($\tilde{\mathbf{y}} = \tilde{\mathbf{w}}$) and possibly to inequality constraints, with the weighting matrices Λ and \mathbf{R} defined in Chapter 2. Since the worst case occurs at a polytope vertex (see Chapter 4), the cost function $J_2(t)$ is evaluated only at the polytope vertices $\theta_i^{N_\theta} \in \Theta^{N_\theta}$.

This min-max problem can be rewritten in order to incorporate different control profiles and the causality constraint, leading to the problem

$$\overline{\Delta \mathbf{U}^{\text{opt}}} = \arg \min_{\Delta \mathbf{U}} \max_{\theta_i^{N_\theta} \in \Theta^{N_\theta}} J_2^i(t),$$

subject to

$$J_2^i(t) = \left(\mathbf{w} - \mathbf{f} - \mathbf{G}\Delta \mathbf{u}_j - \mathbf{H}_\theta \theta_i^{N_\theta} \right)^T \Lambda \left(\mathbf{w} - \mathbf{f} - \mathbf{G}\Delta \mathbf{u}_j - \mathbf{H}_\theta \theta_i^{N_\theta} \right) + \Delta \mathbf{u}_j^T \mathbf{R} \Delta \mathbf{u}_j,$$

with $\Delta u_j = S_j \Delta U$ and

$$j = \left\lceil \frac{i}{2} \right\rceil, \quad (5.1)$$

where $\lceil z \rceil$ denotes the lowest integer which is greater than or equal to z , subject to the equality constraints

$$\tilde{G} \Delta u = \tilde{w} - \tilde{f} - \tilde{H}_\theta \bar{\theta}^{N+n_{\bar{a}}},$$

possibly subject to the general inequality constraints

$$P \Delta u_j \leq r, \theta_i^{N_\theta} \in \left\{ \theta_{2j-1}^{N_\theta}, \theta_{2j}^{N_\theta} \right\},$$

and subject to the uncertainty constraints

$$\theta(t+j|t) = \bar{\theta}, j > N_\theta. \quad (5.2)$$

Remark 5.1 As discussed in Chapter 4, it is not possible to enforce the end-point equality constraints for infinitely many polytope vertices. Therefore, the vector $\bar{\theta}^{N+n_{\bar{a}}}$ is formed with all the components equal to the band average $\bar{\theta}$. $\square\square\square$

Taking into account the ‘‘uncertainty constraints’’ of eqn.5.2, the polytope vertices are formed as

$$\theta_i^{N_\theta} = \left[\theta(t+1|t) \quad \dots \quad \theta(t+N_\theta|t) \quad \bar{\theta} \quad \dots \quad \bar{\theta} \right]^T,$$

where the elements $\theta(t+j|t)$ equal θ^- or θ^+ for $1 \leq j \leq N_\theta = N_u - n_{\bar{a}} - 1$. The last few components of $\theta_i^{N_\theta}$, for $t + N_\theta + 1$ through $t + N$ (or further if required by the constraints), are assumed identical to the average $\bar{\theta}$ of the lower and upper uncertainty bounds. Of course, $\theta_i^{N_\theta} \neq \theta_j^{N_\theta}$ if $i \neq j$, what yields exactly 2^{N_θ} different vertices.

Although the notation is admittedly cumbersome, the interpretation of the formulae provided above is rather intuitive. Each control move profile $\Delta u_j(t)$ is required to handle two uncertainty vertices, namely $\theta_{2j-1}^{N_\theta}$ and $\theta_{2j}^{N_\theta}$. For the example presented

above, $\Delta \mathbf{u}_1(t)$ handles θ_1^3 and θ_2^3 . Conversely, the vertex $\theta_i^{N_\theta}$ is handled by the control move sequence $\Delta \mathbf{u}_{\lceil i/2 \rceil}$, what clarifies the meaning of eqn.5.1. In addition, the end-point equality constraints must be satisfied with all the control sequences $\Delta \mathbf{u}_j$ for $1 \leq j \leq 2^{N_\theta-1}$. The inequality constraints are defined such that each control profile copes with the associated uncertainty realisations. Finally, an equality constraint on the uncertainty beyond N_θ , as discussed above, is assumed. Note that this latter constraint may need to be enforced even further from $t + N$ in case that, for infinite horizon-like constraints, the constraint horizon need be placed beyond the upper prediction horizon (N).

It is worth pointing out that the causality constraint is *implicit* in the above formulation. The way the control sequences $\Delta \mathbf{u}_j(t)$ are selected from $\Delta \mathbf{U}(t)$ by means of the matrices \mathbf{S}_j enforces the causality constraint.

Now, although an analytical solution of this problem, analogous to that described for the standard min-max GPC $^\infty$ in Section 4.2.2.3, is not hard to obtain, such an alternative would be excessively CPU intensive. Fortunately, a numerical solution based on non-linear programming (*e.g.* SQP) can be used to find the optimal extended control move vector as

$$v^{\text{opt}}, \overline{\Delta \mathbf{U}^{\text{opt}}} = \arg \min_{v, \Delta \mathbf{U}} v \text{ subject to } v \geq J_2^i(t) \text{ for } 1 \leq i \leq 2^{N_\theta},$$

where $J_2^i(t)$ is the cost function associated to the uncertainty vertex $\theta_i^{N_\theta}$, subject to all the equality/inequality constraints reported above. The cost functions $J_2^i(t)$ specify the (non-linear) constraints $v \geq J_2^i(t)$. These cost functions (or constraints) can be easily rearranged to be defined in terms of $\Delta \mathbf{U}(t)$ rather than $\Delta \mathbf{u}_j(t)$ by making use of the relation $\Delta \mathbf{u}_j(t) = \mathbf{S}_j \Delta \mathbf{U}(t)$, which yields

$$J_2^i(t) = \left(\mathbf{w} - \mathbf{f} - \mathbf{G}\mathbf{S}_j\Delta\mathbf{U} - \mathbf{H}_\theta\theta_i^{N_\theta} \right)^T \Lambda \left(\mathbf{w} - \mathbf{f} - \mathbf{G}\mathbf{S}_j\Delta\mathbf{U} - \mathbf{H}_\theta\theta_i^{N_\theta} \right) + \Delta\mathbf{U}^T \mathbf{S}_j^T \mathbf{R} \mathbf{S}_j \Delta\mathbf{U}.$$

Note that all the constraints and the optimisation criterion (v) are differentiable with respect to the decision variables $\Delta U(t)$ and the additional degree of freedom v , a property which would not be satisfied if 1-norm cost functions (see Chapter 2) were used. The solution to this optimisation problem with non-linear programming methods, such as SQP, is straightforward. As already remarked, the use of gradient data is beneficial since, firstly, it speeds up the convergence of the optimisation methods and, secondly, the numerical properties of non-linear programming (such as SQP) is improved if the gradients are incorporated.

Remark 5.2 If a 1-norm cost function is used, similarly as done for the standard QGPC $_1^\infty$ in Section 4.3, an LP solution to the optimisation problem can be found. In principle, for each pair of uncertainty vertices $\theta_{2j-1}^{N_\theta}$ and $\theta_{2j}^{N_\theta}$, different σ variables should be defined. Hence the number of σ variables would be $N2^{N_\theta}$, *i.e.* the number of variables and constraints would increase exponentially with N_θ :

$$-\sigma_j \leq G\Delta u_j + H_\theta \theta_i^{N_\theta} + f - w \leq \sigma_j, \theta_i^{N_\theta} \in \left\{ \theta_{2j-1}^{N_\theta}, \theta_{2j}^{N_\theta} \right\},$$

for $j = 1, 2, \dots, 2^{N_\theta-1}$. A simpler (more conservative) solution could be obtained by using the same variables σ_j for all the control profiles Δu_j . Then, if the relation $\Delta u_j = S_j \Delta U$ is introduced, these constraints can be specified as

$$-\sigma \leq GS_j \Delta U + H_\theta \theta_i^{N_\theta} + f - w \leq \sigma, \theta_i^{N_\theta} \in \left\{ \theta_{2j-1}^{N_\theta}, \theta_{2j}^{N_\theta} \right\},$$

for $j = 1, 2, \dots, 2^{N_\theta-1}$. Each of these constraints might be examined row by row so as to apply the constraint reduction procedure depicted in Section 4.2.2.1, leading to the incorporation of $2N \cdot 2^{N_\theta-1}$ linear constraints to the LP problem. Such a solution might be less computationally intensive than the 2-norm counterpart based on SQP presented in this section, at least for a relatively small N_θ , but it is noted that the number of constraints would still increase exponentially with N_θ . □□□

5.3 Simulation results

In this section simulation results for linear and non-linear systems are presented to show how the feedback min-max implementation can overcome some of the drawbacks of the standard min-max MPC approach.

5.3.1 Linear example

First of all, the standard and feedback implementations of the min-max GPC^∞ are compared for a linear plant. This example is analogous to that presented in Fig.4.10 of Section 4.4.1.

Let the true and nominal systems be those used in Section 4.4.1, *i.e.*

$$G(q^{-1}) = \frac{q^{-1}B}{A} = \frac{0.2358q^{-1} + 0.2319q^{-2}}{1 - 1.4835q^{-1} + 0.9512q^{-2}},$$

and

$$G_0(q^{-1}) = \frac{q^{-1}B_0}{A_0} = \frac{0.2973q^{-1} + 0.2923q^{-2}}{1 - 1.7802q^{-1} + 1.3698q^{-2}}.$$

Thus the true system has two unstable poles, whereas the nominal plant is an open-loop stable lightly damped system. In addition, an additive disturbance of magnitude 0.05 (eqn.4.17) is assumed to enter the system at time $t \geq 51$.

In Fig.4.10 it is observed that a dead-time appears when the min-max (1-norm) QGPC_1^∞ is used with a particular choice of the tuning knobs. This undesirable behaviour can also occur with the min-max 2-norm. In the following experiment, a setpoint change from 0 to 1 at time $t = 5$ samples has been performed. The constraint $y(t) \leq 1.2$ for all t has been enforced, and the tuning knobs of the standard min-max GPC^∞ have been chosen as $[N_u, \rho] = [5, 0.1]$, $\theta^-(5) = -0.2$, $\theta^+(5) = 0.2$, $M_\theta = 10$, $\mu = 0.9$. Finally, the uncertainty bands are frozen for four samples after the setpoint change.




Figure 5.1: Closed-loop behaviour of the standard min-max GPC^∞

The outcome of this experiment is shown in Fig.5.1. It is noted that a dead-time appears in the output, similarly as occurs with the min-max QGPC_1^∞ in Fig.4.10. Although the output in the first few samples is slightly greater than zero, it remains far from the setpoint and only when the uncertainty bands are updated is the output taken closer to the setpoint. This example illustrates that this kind of dead-time responses are not a peculiarity of 1-norm formulations, but general to the standard min-max MPC approach irrespective of which norm is used in the cost function definition.

The same experiment, performed for the feedback implementation of the min-max GPC^∞ , yields the input/output responses displayed in Fig.5.2(a), and the uncertainty signal of Fig.5.2(b). Note that the improvement, as performance is concerned, is dramatic. The additional degrees of freedom available to handle different uncertainty realisations are exploited and the five-sample dead-time in the output response of the standard min-max MPC method is no longer found. These supplementary degrees of freedom revert on quite greater control efforts compared to the standard min-max case. The controller does not need to be excessively cautious and the output approaches the setpoint much faster than in Fig.5.1(b), with no risk as constraint handling is concerned,




Figure 5.2: Closed-loop behaviour of the feedback min-max GPC[∞]

since the uncertainty bands are always respected. In addition, it is worth pointing out that the uncertainty signal, as a consequence of the greater control efforts, is larger for the feedback implementation, but it is always kept within the assumed lower and upper bounds.

As the computational burden is concerned, the whole 100-sample simulation performed with the standard min-max GPC[∞] takes 58.5150 seconds on a 400 MHz computer, whereas the feedback min-max GPC[∞] takes 204.9150 seconds. Roughly speaking, the latter implementation involves four times the computational burden of the former. Of course, the same interpreted SQP algorithm has been used for both controllers. Note that the benefits as performance is regarded are quite appreciable, but the increase in CPU time is not too large.

5.3.2 Non-linear plant with a saturation

The next few experiments have been carried out for the non-linear system described in Section A.5 of Appendix A. These experiments are closely related to the results

presented in Section 4.4.4.

The nominal model has been chosen, neglecting the saturating amplifier, assuming a ZOH on the input, and using a sampling time of 0.05 seconds, as (eqn.A.10)

$$G(q^{-1}) = \frac{q^{-1}B}{A} = \frac{0.2358q^{-1} + 0.2319q^{-2}}{1 - 1.4835q^{-1} + 0.9512q^{-2}}$$

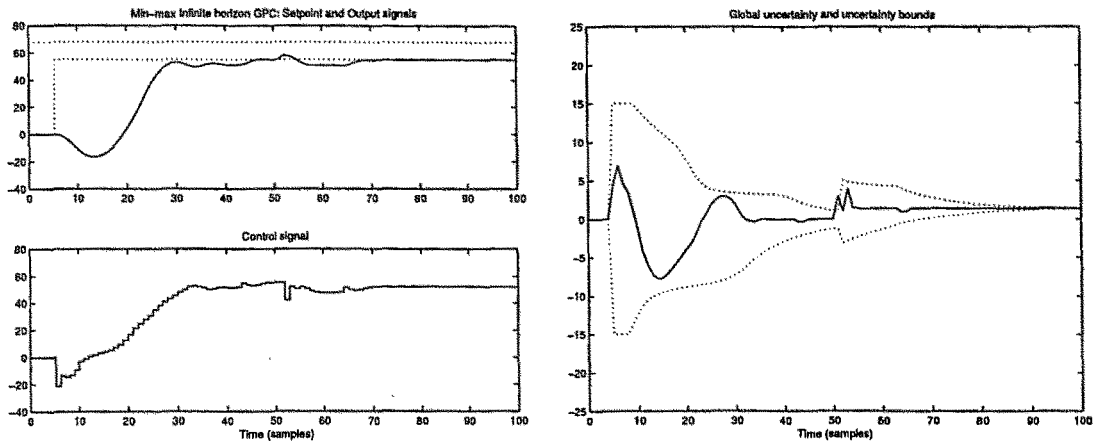
In addition, a constant disturbance of magnitude 3 is added to the true output for $t \geq 51$ samples.

As shown in Fig.4.22, the min-max QGPC₁[∞] can lead to an inverse response for a particular choice of the tuning settings. This drawback is analogous to the dead-time response obtained in the linear case shown in Fig.4.10. Again, as illustrated below, this peculiarity is not confined to 1-norm controllers.

In the next experiment, performed with the standard min-max GPC[∞], a setpoint change from 0 to 55 is scheduled at time $t = 5$ samples, and the output constraint $y(t) \leq 68$ for all t is enforced. The tuning knobs are chosen as $[N_u, \rho] = [5, 0.1]$, $\theta^-(5) = -15$, $\theta^+(5) = 15$, $M_\theta = 10$, $\mu = 0.9$, and the bands are frozen for four samples after the setpoint change.

The outcome of this experiment is displayed in Fig.5.3, where an inverse response, similar to that shown in the experiment of Fig.4.22, can be observed. Once again, it is remarked that such an undesirable closed-loop behaviour is not linked specifically to the 1-norm cost function definition. This kind of problems are common to all the standard min-max implementations.

Now the experiment is performed again with identical settings for the feedback implementation of the min-max controller. The results, shown in Fig.5.4, evidence that the feedback implementation overcomes the standard approach since the inverse response behaviour is completely suppressed.




(a) Input/output responses for the standard min-max GPC^∞ (solid) and output constraint (dotted)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 5.3: Closed-loop behaviour of the standard min-max GPC^∞

Similarly as for the linear example presented in the previous section, the additional degrees of freedom bestowed to the controller allow for greater control efforts and the output is led to the setpoint much faster than for the standard approach. Note, also, that the larger control moves also involve a larger uncertainty signal, which breaks the lower uncertainty bound shortly after the setpoint change. However, since the output is still far from the setpoint (and the constraint), this violation, which reverts on the enlargement of the uncertainty bands, does not involve any difficulty with the accomplishment of the output constraint.

The CPU time required by the 100-sample simulation for the feedback version of the min-max GPC^∞ is 388.8890 seconds on a 400 MHz computer, compared to the 173.8700 seconds required by the standard implementation. Surprisingly enough, the feedback implementation involves less than 2.25 times the computational burden of the classical min-max MPC controller. In such a case, the feedback min-max controller could be taken as a convenient alternative since the relatively small difference as CPU time is regarded leads to quite an improved performance.



(a) Input/output responses for the feedback min-max GPC^∞ (solid) and output constraint (dotted)

(b) Global uncertainty (solid) and uncertainty bounds (dotted)

Figure 5.4: Closed-loop behaviour of the feedback min-max GPC^∞

The advantages of the feedback implementation of min-max MPC are clearly highlighted with both the linear and the non-linear examples.

5.4 Conclusions

In this chapter, one of the most recent developments related to min-max MPC, namely the inclusion of the notion that feedback is present in the receding-horizon implementation of the controller, is presented. The results of Sokaert and Mayne (1998) have been extended to the min-max (2-norm) controllers presented in Chapter 4 in order to lessen some of the drawbacks related to the standard min-max implementations.

The feedback min-max controllers are formulated in such a way that different control profiles are considered for different uncertainty realisations (vertices). Noting that, in the future sampling instants, new measurements of the uncertainty signal will be available, it is not necessary to handle all the possibilities with a single hypothesised control move vector. This idea is combined with a causality constraint to ensure that

the next control move used to update the control signal is completely determined by the current internal process state, and thus does not depend on future (unavailable) data. Making use of these two concepts, it is possible to establish a near one-to-one correspondence between the degrees of freedom available for optimisation and the number of uncertainty realisations, a property which is not shared by the standard min-max approach, for which the difference between the degrees of freedom of the controller and the uncertainty realisations increases dramatically with the control and prediction horizons.

The optimisation problem can be solved using standard non-linear programming tools, such as SQP. It is also highlighted that an LP solution for the analogous 1-norm problem is possible, but the number of constraints depends on 2^{N_θ} , and thus the computational advantages of the feedback implementation are not as remarkable as for the standard min-max QGPC₁[∞]. In addition, the 2-norm GPC[∞] guarantees the stability of the nominal closed-loop system, and hence the feedback min-max GPC[∞] seems favoured with respect to the QGPC₁[∞] counterpart.

Simulation results are presented for both linear and non-linear systems, and it is shown that the feedback implementation overcomes the standard counterpart as performance is concerned. The additional degrees of freedom are used by the feedback min-max methods to avoid some of the problems related to the standard min-max controllers, such as unexpected inverse responses or dead-times which can occur in the constrained case.

Chapter 6

Conclusions and future research

6.1 Conclusions

MPC has become a mature control strategy from both the industrial and the academical points of view. The simplicity of the ideas behind MPC, together with the ability of these methods to tackle difficult control problems, such as complex dynamics, MIMO systems, constrained systems, etc., are the two main reasons for such a remarkable success.

Since the first few predictive controllers were suggested in the late 1970's, a great deal of different MPC control laws have emerged. Among these, the GPC controller of (Clarke *et al.*, 1987) gained the early recognition of the scientific and the industrial communities. Promptly after the suggestion of the GPC, the problems of stability and robustness were addressed.

Throughout this PhD research, the problems of stability, robustness and constraint handling for GPC-like controllers have been investigated in a combined manner.

In this thesis, the main stability results for MPC controllers based on input/output models are reviewed. Stability proofs for both the CRHPC and the infinite horizon GPC (GPC[∞]) are given. These results, based on the monotonicity of the optimal

cost function sequence, provide with an intuitive insight to the stability issue. Other stabilising approaches, such as the SGPC or the SIORHC, are not explicitly reviewed since these are proved to be theoretically equivalent to the CRHPC (under ideal conditions), although the SGPC enjoys better numerical properties. In addition, 1-norm counterparts of the classical 2-norm formulations have been obtained. Therefore, the 1-norm GPC (GPC_1), 1-norm CRHPC (CRHPC_1) and 1-norm GPC^∞ (GPC_1^∞) are formulated. These 1-norm MPC schemes can be implemented using very efficient standard LP routines. The stability proofs provided for the 2-norm case are extended for the 1-norm counterparts using the same monotonicity argument. Simulation results are provided to illustrate that stability is achieved and that a non-increasing sequence of optimal cost function values results.

Moreover, it is noted that the GPC_1^∞ control law requires the implementation of an iterative procedure until some convergence condition is satisfied. This iterative algorithm involves the solution of two LP problems at each iteration, and the associated computational burden can become excessively large. A less computationally intensive controller, the Quasi-infinite horizon 1-norm GPC (QGPC_1^∞), is thus suggested. This method, which minimises an upper bound of the truly infinite horizon costing problem, can be implemented using a single LP problem, and thus the iterative algorithm required by the GPC_1^∞ can be avoided. In addition, a convergence property from the QGPC_1^∞ to the GPC_1^∞ is conjectured and illustrated by means of examples.

The usefulness of 1-norm GPC-like controllers is the formulation of efficient robust constrained MPC methods based on min-max optimisation. Prior to undertake such a task, a robustness analysis of unconstrained GPC-like controllers is provided. This analysis is based on the LTI formulation of unconstrained 2-norm methods, which can be obtained following the procedure presented in (Bitmead *et al.*, 1990). These formulae make it possible to apply classical robust control results, such as the small gain theorem, which yields stability conditions for different uncertainty representations. In

this PhD thesis, the extension of this approach to the infinite horizon case is provided.

Several stabilising approaches, namely the CRHPC, the GPC^∞ and a softened version of the CRHPC, are compared in terms of nominal performance and robustness. It is noted that the infinite horizon approach provides with smooth input/output responses and convenient robustness margins for typical choices of the tuning knobs, whereas the CRHPC often leads to deadbeat-like closed-loop behaviour with poor robustness bounds. The classical approach to enhance robustness, via the heuristic T -design and the systematic Q -parametrisation methods, for LTI uncertainty are reviewed. In addition, a new robustness-enhancing method based on choosing the observer polynomial T by means of optimisation instead of heuristic rules, termed T -optimisation, is suggested. This new method is shown to overcome the classical T -design and Q -parametrisation approaches for a particular example.

The robustness of constrained MPC control laws is then addressed. The global uncertainty approach has been taken, since this kind of description can be used to describe all kinds of uncertainties, namely linear, non-linear, time invariant, time varying, stable, unstable, parametric, non-parametric, modelling errors, disturbances and so on. A global uncertainty is an unknown bounded signal which, added to the model output, yields the true system output. The min-max optimisation problem is defined as the computation of the control profile which minimises the maximum of a cost function as the future uncertainty ranges within the assumed lower and upper bounds in the future. Solutions for both the 1-norm and the 2-norm cases are provided. According to the results of the robustness analysis in the unconstrained case, the (quasi) infinite horizon approach is preferred, and thus the min-max (2-norm) GPC^∞ and the min-max min-max (1-norm) QGPC_1^∞ are formulated. The computational advantages of the latter with respect to the former, since the optimisation problem can be efficiently solved with LP tools in the 1-norm case, are remarked.

An algorithm to update the uncertainty bands on-line is suggested. This procedure starts with conservative settings of the lower and upper bounds, but these are replaced by more appropriate values when new uncertainty measurements become available. This band updating algorithm can be adjusted using a few parameters for which tuning guidelines are provided and tested. Simulation results, obtained with these controllers for different sources of uncertainty (linear, non-linear, time invariant, disturbances, etc.), illustrate the behaviour of these new methods which often overcome the classical MPC controllers, especially in the constrained case. A robustness analysis performed using Monte Carlo simulation is provided, and it is shown that the min-max QGPC₁[∞] successfully handles quite difficult control problems. The min-max approach manages to satisfy constraint specifications where the classical T -based methods fail.

Nevertheless, a few limitations of the standard min-max controllers are illustrated: unexpected dead-times or inverse responses can occur in the constrained case. The problems with the standard min-max approach can even include instability and infeasibility. Some of these drawbacks stem from the use of a single control profile to handle a large number of uncertainty sequences (or polytope vertices). The notion that feedback is present in the receding-horizon implementation of the controller can be exploited to allow for different control profiles to handle different uncertainty sequences. This idea, combined with a so-called “causality constraint” to avoid a multiplicity of choices for the next control move, gives rise to feedback implementations of the min-max controllers. Examples of the feedback min-max GPC[∞] are provided to show that this approach overcomes some of the drawbacks of the standard min-max MPC controllers at the price of a somewhat larger computational burden.

6.2 Possible directions for future research

The results obtained throughout this PhD research provide with some tools to tackle control problems of constrained uncertain systems. It has been shown that min-max MPC methods based on the global uncertainty approach can successfully handle difficult control problems and satisfy constraints in the presence of uncertainty with very low computational requirements. There are, however, several directions to further this research:

1. To begin with, robust stability conditions for the min-max controllers formulated in this thesis might be investigated. Although no stability problems should be expected whenever the uncertainty signal is kept within the assumed lower and upper bounds, such a signal might become unbounded and thus lead to instability despite the use of the band updating algorithm.
2. The band updating algorithm can be further refined, and more sophisticated self-tuning methods might be obtained.
3. In addition, the possibility of combining the global uncertainty approach with other descriptions, such as multi-model formulations, can provide with a methodology to tackle a difficult control problem using a two-step design: firstly a set of locally linearised models can be obtained at different operating points and, secondly, the global uncertainty approach could be applied about each of those models to allow for some (small) discrepancies. The combined advantages of these two descriptions might provide with an improved robust performance, and no difficulty should arise to satisfy constraint specifications.
4. New optimisation techniques might also be used to cut down the number of computations, especially for the min-max 2-norm controllers.

5. The MIMO formulation of the min-max methods depicted throughout this thesis requires some small adjustments, which can be easily addressed.
6. Finally, industrial trials performed with the suggested min-max methods can be carried out to assess the benefits of these controllers and to focus on practical implementation issues.

These are only a handful of the ideas to further this research but, undoubtedly, many other features concerning the robustness of constrained MPC will capture the attention of the control community in a near future.

Appendix A

Benchmark systems

A.1 Linear unstable nearly undetectable plant

One of the benchmark models used to analyse the properties of different controllers is the “Unstable GPC Example” provided by Bitmead *et al.* (1990). The system is described by the transfer function

$$G(q^{-1}) = \frac{q^{-1}B(q^{-1})}{A(q^{-1})} = \frac{q^{-1} - 1.999q^{-2}}{1 - 4q^{-1} + 4q^{-2}}, \quad (\text{A.1})$$

which is a non-minimum phase open-loop unstable system with both poles at 2. In addition, there is a near pole-zero cancellation, what implies that the system is almost undetectable. This benchmark system has been shown to raise objection to the classical GPC, since it is quite difficult to tune the GPC in order to obtain closed-loop stability for this process.

A.2 Linear stable plant with gain uncertainty

Some of the experiments which are carried out in this thesis consider a benchmark model consisting of a transfer function with the following structure:

$$G_0(s) = \frac{N_0(s)}{D_0(s)} = K(1 + \Delta_K) \frac{\omega_0^2 \omega_1}{\phi \omega_2} \frac{(\phi - s)(s + \omega_2)}{(s^2 + 2\zeta \omega_0 s + \omega_0^2)(s + \omega_1)}.$$

Where $\Delta_K \in [-0.5, 0.5]$ is a multiplicative uncertainty in the gain which may vary due to disturbances, changes in the process, etc. In the undisturbed case $\Delta_K = 0$. The other parameters are given in Table A.1.

K	ω_0	ζ	ω_1	ϕ	ω_2
-0.5	10	0.3	5	5	20

Table A.1: Parameters of the linear benchmark model of eqn.A.2

With the parameters of Table A.1, $G_0(s)$ becomes

$$G_0(s) = \frac{N_0(s)}{D_0(s)} = (1 + \Delta_K) \frac{2.5s^2 + 37.5s - 250}{s^3 + 11s^2 + 130s + 500}, \quad (\text{A.2})$$

which is a non-minimum phase, open-loop stable underdamped system with (negative) gain $-0.5(1 + \Delta_K)$, three poles and two zeroes. The poles are located at $-3 \pm 9.5394j$ and -5 , and the zeroes at -20 and 5 .

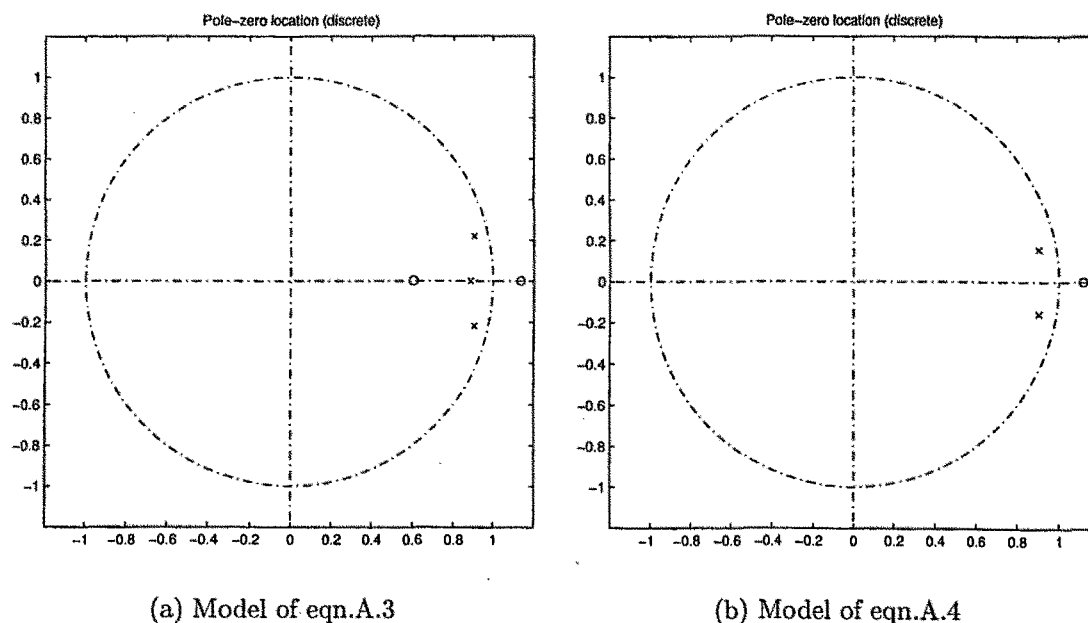


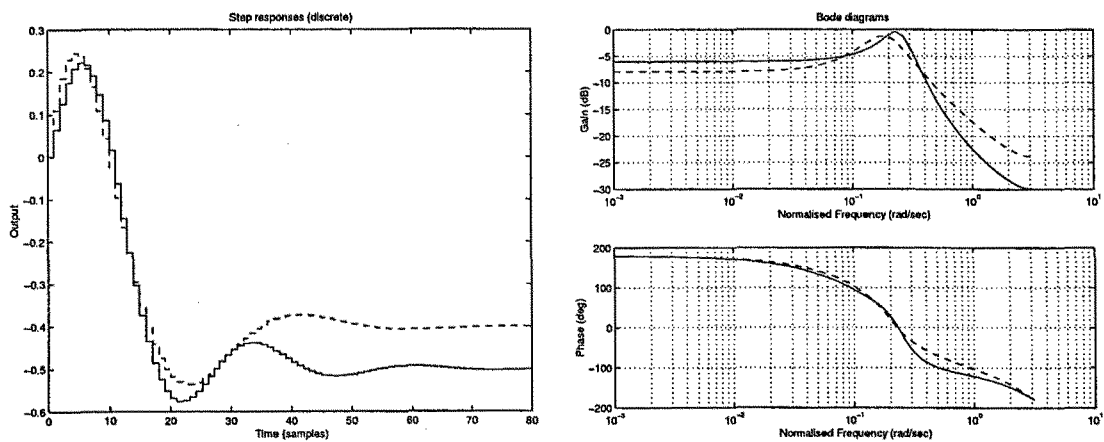
Figure A.1: Pole-zero location of the discrete benchmark system and the identified model

A sampling time of $T_s = 0.025$ s (in such a way that the Nyquist frequency is greater than 100 rad/s) and a *Zero-Order Hold* (ZOH) on the input have been chosen

to obtain the discrete-time model

$$G_0(q^{-1}) = \frac{q^{-1}B_0(q^{-1})}{A_0(q^{-1})} = (1 + \Delta_\kappa) \frac{q^{-1}(0.0639 - 0.1110q^{-1} + 0.0437q^{-2})}{1 - 2.6855q^{-1} + 2.4518q^{-2} - 0.7596q^{-3}}. \quad (\text{A.3})$$

The pole and zero locations of this discrete-time system are shown in Fig.A.1(a). A cross is used for the poles, at $0.9015 \pm 0.2192j$ and 0.6028 , whereas the zeroes, at 0.8825 and 1.1336 , are displayed with circle. Needless to say, the gain is the same as for the continuous system, namely $-0.5(1 + \Delta_\kappa)$.



(a) Step responses of the true system (solid) and the identified model (dashed)

(b) Frequency responses of the true system (solid) and the identified model (dashed)

Figure A.2: Comparison of the true system and model responses

In order to introduce modelling errors in the experiments, a second-order model has been identified for the system of eqn.A.2 with $\Delta_\kappa = 0$. This model has been found using the system identification procedure described by Whitfield (1986), which works with data taken from the frequency domain, that is, the true frequency response is measured at several points and then curve-fit optimisation is made using these data. For this example, ten frequencies have been chosen within the interval $[1, 100]$ rad/s, which roughly includes one decade before and one decade after the system's bandwidth frequency, which is at about 15 rad/s. The transfer function attained in this fashion is

$$G(s) = \frac{N(s)}{D(s)} = \frac{5.0182s - 23.4458}{s^2 + 6.8198s + 58.4235},$$

with a zero at 4.5899, two poles at $-3.4099 \pm 6.8408j$ and a steady-state gain of -0.4013 . With a sampling time of $T_s = 0.025$ s and a ZOH on the input, a discrete model can be obtained as

$$G(q^{-1}) = \frac{q^{-1}B(q^{-1})}{A(q^{-1})} = \frac{q^{-1}(0.1098 - 0.1232q^{-1})}{1 - 1.8098q^{-1} + 0.8432q^{-2}}, \quad (\text{A.4})$$

the poles of which are located at $0.9049 \pm 0.1563j$ and the zero at 1.1223, as shown in Fig.A.1(b). As analysed in (Megías, 1996) this is, among the models obtained trying several time and frequency domain identification methods, the best second-order model of the “true” system of eqn.A.3 in terms of frequency response fit.

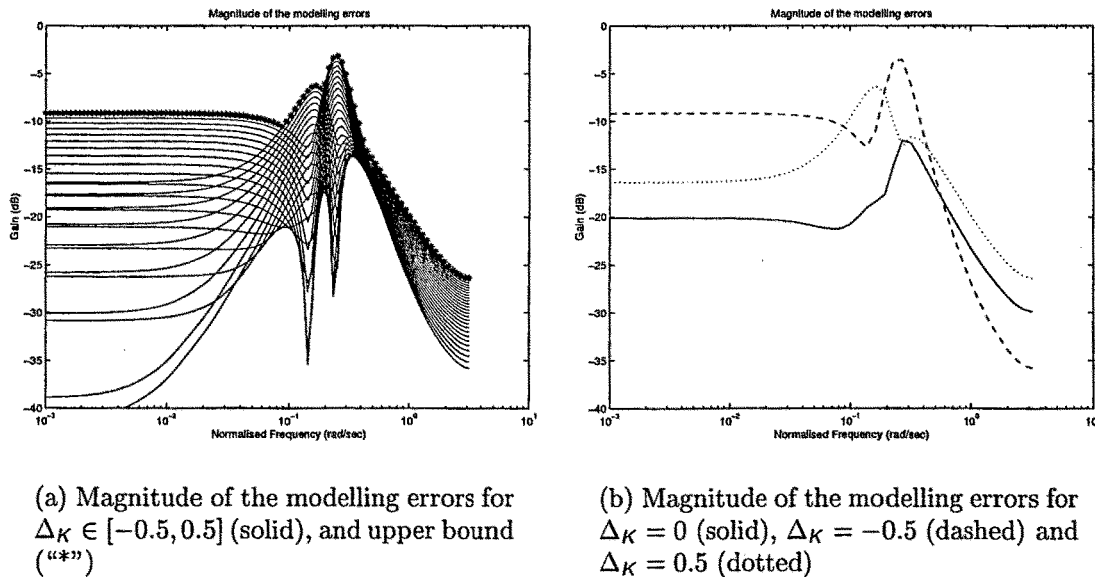


Figure A.3: Additive system uncertainty

Fig.A.2 compares the step (a) and frequency (b) responses of the true system (eqn.A.3 with $\Delta_K = 0$) and the identified model (eqn.A.4). It is observed that there is a small error in the gain, and that the magnitude of the frequency responses diverge at high frequencies. It must be taken into account that the frequency responses are influenced by the ZOH which is assumed in the inputs. In addition, notice that the frequency response for both gain and phase are plotted against normalised and not true frequencies. The normalised frequency ω_n is defined as $\omega_n = T_s\omega$, where ω is the

true frequency and T_s is the sampling time. Hence, the normalised Nyquist frequency is always π rad/s.

Finally, Fig.A.3 displays the magnitude of the modelling errors (additive system uncertainty) from $G_0(q^{-1})$ to $G(q^{-1})$:

$$|G_0(q^{-1}) - G(q^{-1})|_{q=e^{j\omega_n}} = \left| \frac{B_0(q^{-1})A(q^{-1}) - B(q^{-1})A_0(q^{-1})}{A_0(q^{-1})A(q^{-1})} \right|_{q=e^{j\omega_n}}$$

In Fig.A.3(a) the modelling errors are shown for several values of $\Delta_\kappa \in [-0.5, 0.5]$, and the maximum (upper bound) is emphasised by means of “*” signs. In Fig.A.3(b) the modelling errors are shown for $\Delta_\kappa = 0$, which is the best case of system uncertainty, for $\Delta_\kappa = 0.5$, which is the worst case at low frequencies and shows the highest peak, and for $\Delta_\kappa = -0.5$, which is the worst case at high frequencies. The modelling errors for $\Delta_\kappa = 0$ (used to identify the model) are remarkably small, since the magnitude curve is always below the -10 dB line. It is worth pointing out that, in any case, the modelling errors are maximal near the resonance frequency (about $2 \cdot 10^{-1}$ rad/s of normalised frequency) of the true and identified systems.

A.3 Gain and zero uncertainty

In this section, a second order system with gain and zero uncertainty is presented. Let the true plant be described as the second-order system

$$G_0(q^{-1}) = K \frac{0.25}{(1 - \phi)} \frac{q^{-1}(1 - \phi q^{-1})}{1 - 1.4q^{-1} + 0.65q^{-2}}, \quad (\text{A.5})$$

where the gain $0.5 \leq K \leq 1.5$ and the zero $0.4 \leq \phi \leq 0.6$ are uncertain parameters. The plant family which results of this definition is thus given by $\mathcal{G} = \{G_0 : K \in [0.5, 1.5], \phi \in [0.4, 0.6]\}$.

Some robust MPC formulations require that uncertainty is described as a polytope of linear plants such that any possible true plant can be expressed as a linear com-

ination of the polytope vertices. In the transfer function context, this is possible if *all* the uncertain parameters affect the numerator only. Notice that the plant definition provided above does not satisfy this condition, since the denominator possesses the uncertain factor $(1 - \phi)$. This difficulty can be overcome by using the following alternative plant definition:

$$\widehat{G}_0(q^{-1}) = \widehat{K} \frac{q^{-1}(1 - \phi q^{-1})}{1 - 1.4q^{-1} + 0.65q^{-2}},$$

where \widehat{K} and ϕ are uncertain parameters. The relation between the steady-state gain K and the parameter \widehat{K} is provided by

$$\widehat{K} = 0.25 \frac{K}{1 - \phi}.$$

Now the ranges $0.5 \leq K \leq 1.5$ and $0.4 \leq \phi \leq 0.6$ can be used to determine a range for \widehat{K} . The minimum value of \widehat{K} occurs for $\widehat{K} = 0.25 \cdot 0.5 / (1 - 0.4) = 0.2083$ whereas the maximum is found at $\widehat{K} = 0.25 \cdot 1.5 / (1 - 0.6) = 0.9375$. Now an extended plant family $\widehat{\mathcal{G}}$ can be defined as $\widehat{\mathcal{G}} = \{\widehat{G}_0 : \widehat{K} \in [0.2083, 0.9375], \phi \in [0.4, 0.6]\}$. Note that $\mathcal{G} \subset \widehat{\mathcal{G}}$, *i.e.* $\widehat{\mathcal{G}}$ includes some plants which are not in \mathcal{G} , but all the plants in \mathcal{G} are also in $\widehat{\mathcal{G}}$.

Any true plant $G_0 \in \widehat{\mathcal{G}}$, and thus any true plant $G_0 \in \mathcal{G} \subset \widehat{\mathcal{G}}$, can be obtained as a linear combination of the following four systems (vertices):

$$\begin{aligned} G_1 &= 0.2083 \frac{q^{-1}(1 - 0.4q^{-1})}{1 - 1.4q^{-1} + 0.65q^{-2}}, & G_2 &= 0.2083 \frac{q^{-1}(1 - 0.6q^{-1})}{1 - 1.4q^{-1} + 0.65q^{-2}}, \\ G_3 &= 0.9375 \frac{q^{-1}(1 - 0.4q^{-1})}{1 - 1.4q^{-1} + 0.65q^{-2}}, & G_4 &= 0.9375 \frac{q^{-1}(1 - 0.6q^{-1})}{1 - 1.4q^{-1} + 0.65q^{-2}}, \end{aligned}$$

i.e. given $G_x \in \mathcal{G}$ then

$$G_x = \lambda_1 G_1 + \lambda_2 G_2 + \lambda_3 G_3 + \lambda_4 G_4,$$

with $\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4 = 1$ and $\lambda_i \geq 0$ for $i = 1, 2, 3, 4$. It can be noted that the steady-state gain of G_2 and G_3 are 0.3333 and 2.25 respectively, both of them outside the interval $[0.5, 1.5]$ defined for K .

Now the coefficients λ_i can be easily obtained from \widehat{K} and ϕ as

$$\begin{aligned}\lambda_1 &= (1.2857 - 1.3714\widehat{K})(3 - 5\phi), & \lambda_2 &= (1.2857 - 1.3714\widehat{K})(-2 + 5\phi), \\ \lambda_3 &= (-0.2857 + 1.3714\widehat{K})(3 - 5\phi), & \lambda_4 &= (-0.2857 + 1.3714\widehat{K})(-2 + 5\phi),\end{aligned}$$

and it can be checked that $\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4 = 1$ and $\lambda_i \geq 0$ for $\widehat{K} \in [0.2083, 0.9375]$ and $\phi \in [0.4, 0.6]$. The plant family $\widehat{\mathcal{G}}$ can be described as the convex hull of the systems G_1, G_2, G_3 and G_4 , denoted by $\widehat{\mathcal{G}} = \text{Co}\{G_1, G_2, G_3, G_4\}$.

As an example, the system

$$G_m(q^{-1}) = \frac{0.5000q^{-1} - 0.2500q^{-2}}{1 - 1.4q^{-1} + 0.65q^{-2}}, \quad (\text{A.6})$$

obtained for $K = 1$ and $\phi = 0.5$, can be written as

$$G_m(q^{-1}) = 0.5 \frac{q^{-1}(1 - 0.5000q^{-1})}{1 - 1.4q^{-1} + 0.65q^{-2}},$$

i.e. $\widehat{K} = 0.5$ and $\phi = 0.5$, which yields the coefficients $\lambda_1 = 0.3$, $\lambda_2 = 0.3$, $\lambda_3 = 0.2$, $\lambda_4 = 0.2$.

A.4 Gain, zero and pole uncertainty

Similarly as done in the previous section, let a true plant be described as the second-order system

$$G_0(q^{-1}) = K \frac{(1 - \eta)(1 - \bar{\eta})}{(1 - \phi)} \frac{q^{-1}(1 - \phi q^{-1})}{(1 - \eta q^{-1})(1 - \bar{\eta} q^{-1})}, \quad (\text{A.7})$$

where the gain $K = (1 + \Delta_K)K_m$, the zero $\phi = (1 + \Delta_\phi)\phi_m$ and the pole $\eta = (1 + \Delta_\eta)\eta_m$ are uncertain parameters such that $K_m = 1$, $\phi_m = 0.6$, $\eta_m = 0.6261 + 0.3130j$ and

$$|\Delta_K| \leq 0.2,$$

$$|\Delta_\phi| \leq 0.1,$$

$$|\Delta_\eta| \leq 0.1,$$

i.e. there is a 20% maximum variation of the steady-state gain with respect to the central value ($K_m = 1$), and a 10% variation of the pole and the zero with respect to $\eta_m = 0.6261 + 0.3130j$ and $\phi_m = 0.6$ respectively.

Now the nominal system can be defined as the one for which each parameter (gain, zero and pole) occurs at the central value, *i.e.*

$$G(q^{-1}) = \frac{0.5944q^{-1} - 0.3567q^{-2}}{1 - 1.2522q^{-1} + 0.4900q^{-2}}. \quad (\text{A.8})$$




Figure A.4: Multiplicative uncertainty (solid) and upper uncertainty bound (“*”)

Fig.A.4 displays the magnitude of the modelling errors (multiplicative system uncertainty) from $G_0(q^{-1})$ to $G(q^{-1})$:

$$\left| \frac{G_0(q^{-1}) - G(q^{-1})}{G(q^{-1})} \right|_{q=e^{j\omega n}} = \left| \frac{B_0(q^{-1})A(q^{-1}) - B(q^{-1})A_0(q^{-1})}{A_0(q^{-1})B(q^{-1})} \right|_{q=e^{j\omega n}},$$

and the maximum (upper bound) is highlighted using “*” signs.




Figure A.5: Non-linear benchmark plant

A.5 Non-linear system with a saturation

Some of the control algorithms depicted in this thesis have been tested on the non-linear system of Fig.A.5 (Naslin, 1962; Thaler and Pastel, 1962), which can be described by the equations:

$$x = \min \{ \max \{ -2, u - y \}, 2 \},$$


$$\ddot{y} + \dot{y} = 200x.$$

The difficulty to control this system is caused by the saturating amplifier which appears inside the inner feedback loop. Saturators are a kind of hard non-linearity which commonly appear in real systems. A linear model of the system can be obtained, neglecting the saturation, as

$$G(s) = \frac{200}{s + s^2 + 200}, \quad (\text{A.9})$$

with poles at $-0.5 \pm 14.1333j$ and unit steady-state gain. As a consequence of the pole location, the open-loop response of this linear system is very lightly damped.

The different behaviour patterns provided by the true system of Fig.A.5 and the linear model of eqn.A.9 depend on the amplitude of the input signal. In Fig.A.6 the step responses of the non-linear system of Fig.A.5 and the linear model of eqn.A.9 for steps of amplitude 5 (a) and 15 (b) on the input are compared. The step change in the input occurs at $t = 0.1$ s. It can be observed that the higher the input amplitude is, the more different the step responses are. In fact, for input amplitudes lower than 3, the step responses are almost identical. Of course, this is due to the saturating amplifier, which is the only difference between the linear and the non-linear models.



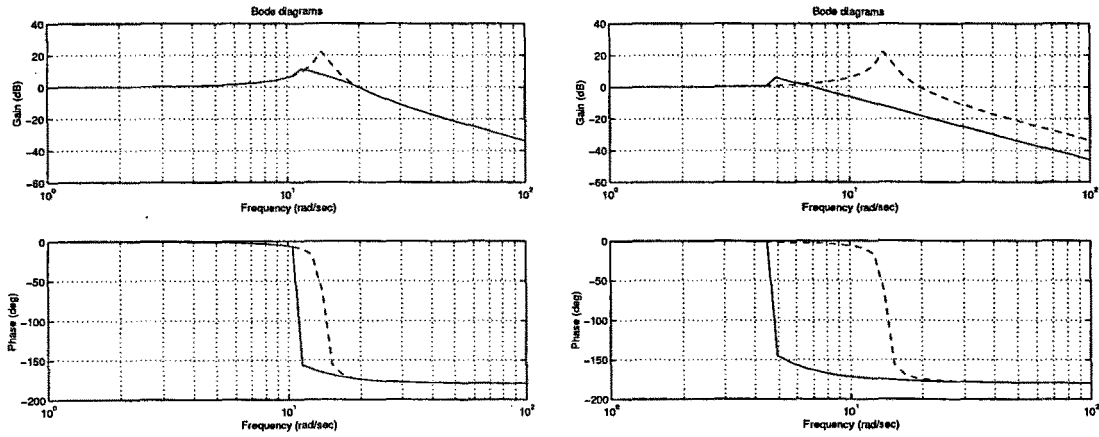
(a) Step responses of the non-linear system (solid) and the linear model (dashed). Amplitude of the input signal = 5

(b) Step responses of the non-linear system (solid) and the linear model (dashed). Amplitude of the input signal = 15

Figure A.6: Step responses of the non-linear system and the linear model for different input amplitudes

In Fig.A.7 the frequency responses of the non-linear and the linear systems are compared. The frequency response in the non-linear case has been measured using the algorithm described by Wellstead and Zarrop (1991), feeding the system with sine waves of different amplitudes at the input. Of course, the frequency response of the linear system is independent of the input amplitude, but the behaviour of the non-linear system is quite different in this aspect. Similarly as happens with the step responses, the higher the input amplitude is, the more distant the frequency responses become. In fact, for low amplitudes (0.25 or less) the frequency responses of both systems are identical, since the saturation has no effect in that case. It is worth remarking that the resonance frequency of the non-linear system decreases as the input amplitude increases, what can be easily observed by comparing Fig.A.7(a) and (b).

With a ZOH on the input and a sampling time of $T_s = 0.05$ s, which has been chosen such that each period of the open-loop step response is sampled about 20 times, the



(a) Frequency responses of the non-linear system (solid) and the linear model (dashed). Amplitude of the input signal = 1

(b) Frequency responses of the non-linear system (solid) and the linear model (dashed). Amplitude of the input signal = 10

Figure A.7: Frequency responses of the non-linear system and the linear model for different input amplitudes

linear model of eqn.A.9 turns out to be the discrete-time transfer function

$$G(q^{-1}) = \frac{q^{-1}B(q^{-1})}{A(q^{-1})} = \frac{0.2358q^{-1} + 0.2319q^{-2}}{1 - 1.4835q^{-1} + 0.9512q^{-2}}, \quad (\text{A.10})$$

the poles of which are located at $0.7418 \pm 0.6333j$ and the zero at -0.9832 . The modulus of the poles is 0.9753, very near the unit circle.

References

- Allwright, J. C. (1994). On min-max mode-based predictive control. In: Clarke (1994).
- Allwright, J. C. and G. C. Papavasiliou (1991). A reduced-size lp formulation for robust mpc using impulse responses. In: 30th *IEEE Conference on Decision and Control*. Brighton, UK.
- Allwright, J. C. and G. C. Papavasiliou (1992). On linear programming and robust model-predictive control using impulse-responses. *Systems and Control Letters* **18**, 159–164.
- Álvarez, T. and C. de Prada (1997). Handling infeasibilities in predictive control. *Computer and Chemical Eng.* **21**, S577–S582.
- Anderson, B. D. O. and J. B. Moore (1971). *Linear Optimal Control*. Prentice Hall. Englewood Cliffs, NJ, USA.
- Ansay, P. and V. Wertz (1997). Model uncertainties in GPC: A systematic two-step design. In: *Proc. of the European Control Conference*. Brussels, Belgium.
- Åström, K. J. (1970). *Introduction to stochastic control theory*. Academic Press. New York, NY, USA.
- Åström, K. J. and B. Wittenmark (1973). On self-tuning regulators. *Automatica* **9**, 185–199.

- Bazaraa, M. S. and C. M. Shetty (1979). *Nonlinear Programming*. John Wiley and Sons. Chichester, UK.
- Berber, R. and Kravaris, C., (Eds.) (1998). *Nonlinear Model Based Process Control*. Kluwer academic publishers. Dordrecht, Germany.
- Bitmead, R., M. Gevers and V. Wertz (1990). *Adaptive Optimal Control. The Thinking Man's GPC*. Prentice Hall. Sydney, Australia.
- Camacho, E. F. and C. Bordóns (1995). *Model Predictive Control in the Process Industry*. Springer-Verlag. London, UK.
- Camacho, E. F. and M. Berenguel (1994). Application of generalised predictive control to a solar power plant. In: Clarke (1994).
- Chang, T. S. and D. E. Seborg (1983). A linear programming approach to multivariable feedback control with inequality constraints. *Int. J. Control* **37**, 583–597.
- Chen, H. and F. Allgöwer (1998a). Nonlinear model predictive control schemes with guaranteed stability. pp. 465–494. In: Berber and Kravaris (1998).
- Chen, H. and F. Allgöwer (1998b). A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Automatica* **34**(10), 1205–1217.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals Math. Stat.* **29**, 493–507.
- Chischi, L. and E. Mosca (1994). Stabilizing predictive control: the singular transition. In: Clarke (1994).
- Clarke, D. W. and C. Mohtadi (1989). Properties of Generalized Predictive Control. *Automatica* **25**(6), 859–875.
- Clarke, D. W. and R. Scattolini (1991). Constrained receding-horizon predictive control. *IEE Proc. Part D* **138**(4), 347–354.

- Clarke, D. W., C. Mohtadi and P. S. Tuffs (1987). Generalized Predictive Control - Part I: The basic algorithm and Part II: Extensions and interpretations. *Automatica* **23**(2), 137–160.
- Clarke, D. W., (Ed.) (1994). *Advances in model-based predictive control*. Oxford University Press. Oxford, UK.
- Cristea, S. (1998). Control predictivo en el dominio delta. PhD thesis. Departamento de Ingeniería de Sistemas y Automática. Universidad de Valladolid. (In Spanish).
- Cutler, C. R. and B. C. Ramaker (1980). Dynamic matrix control—a computer control algorithm. In: *Proc. Joint Automatic Control Conference*. San Francisco, CA, USA.
- Cutler, C. R. and R. B. Hawkins (1987). Constrained multivariable control of a hydrocracker reactor. In: *Proc. of the American Control Conference*. Minneapolis, MN, USA. pp. 1014–1020.
- Dahleh, M. A. and I. J. Díaz-Bovillo (1995). *Control of Uncertain Systems*. Prentice Hall. Englewood Cliffs, NJ, USA.
- De Keyser, R. M. C. and A. R. Van Cauwenberghe (1985). Extended prediction self-adaptive control. In: *IFAC Symposium on Identification and System Parameter Estimation*. York, UK.
- De Nicolao, G., L. Magni and R. Scattolini (1996). Robust predictive control of systems with uncertain impulse response. *Automatica* **32**(10), 1475–1479.
- de Prada, C., J. Serrano, P. Vega and M. A. Piera (1994). A comparative study of DMC and GPC controllers. In: Clarke (1994).
- El Ghoumari, M. Y. (1998). Control predictivo no lineal de un robot de tres grados de libertad. Master's thesis. Departament d'Informàtica. Universitat Autònoma de Barcelona. (In Spanish).

- Fikar, M. and S. Engell (1997). Receding horizon predictive control based upon Youla-Kučera parametrization. *European J. Control* 3(4), 304–316.
- Fikar, M., S. Engell and P. Dostál (1999). Youla-Kučera parametrisation approach to predictive control. In: *Proc. of the European Control Conference*. Karlsruhe, Germany.
- Francis, B. A. (1987). *A Course in H_∞ Control Theory*. Springer-Verlag, Berlin, Germany.
- Garcia, C. E. and A. M. Morshedi (1984). Quadratic programmatic solution of dynamic matrix control. *Eng. Chem. Commun.* 46, 73–87.
- Garcia, C. E. and M. Morari (1982). Internal model control: Part 1. a unifying review and some new results. *Industrial Engineering and Chemical Process Design and Development* 21, 308–323.
- Hrissagis, K., O. Crisalle and M. Sznaier (1996). Robust unconstrained predictive control design with guaranteed nominal performance. *AIChE Journal* 42(5), 1293–1303.
- Kleinman, D. L. (1974). Stabilizing a discrete, constant, linear system with application to iterative methods for solving the Riccati equation. *IEEE Trans. on Automatic Control* 19(3), 252–254.
- Kothare, M. V., V. Balakrishnan and M. Morari (1996). Robust constrained model predictive control using linear matrix inequalities. *Automatica* 32(10), 1361–1379.
- Kouvaritakis, B., J. A. Rossiter and A. O. T. Chang (1992). Stable generalised predictive control: An algorithm with guaranteed stability. *IEE Proc. Part D* 139(4), 349–362.
- Kuznetsov, A. G. and D. W. Clarke (1996). The performance of generalised predictive control with interval constraints. *European J. Control* 2, 260–277.

- Kwakernnak, H. and R. Sivan (1972). *Linear Optimal Control Systems*. John Wiley and Sons. New York, NY, USA.
- Kwon, W. H. and A. E. Pearson (1978). On feedback stabilization of time-varying discrete linear systems. *IEEE Trans. on Automatic Control* **23**(3), 479–481.
- Maciejowski, J. M. (1989). *Multivariable Feedback Design*. Addison-Wesley. Reading, MA, USA.
- Martin, G. D., J. M. Caldwell and T. E. Ayril (1986). Predictive control applications for the petroleum refining industry. In: *Petroleum Refining Conf.*. Tokyo, Japan.
- Martín Sánchez, J. M. and J. Rodellar (1996). *Adaptive Predictive Control. From the concepts to plant optimization*. Prentice Hall. Hertfordshire, UK.
- Matsko, T. N. (1985). Internal model control for chemical recovery. *Chem. Eng. Progress* **81**(12), 46–51.
- Mayne, D. Q. and H. Michalska (1990). Receding horizon control of nonlinear systems. *IEEE Trans. on Automatic Control* **35**(7), 814–824.
- Megías, D. (1994). Control DMC de un reactor químico: aspectos de implementación. Bachelor's thesis. Departament d'Informàtica. Universitat Autònoma de Barcelona. (In Spanish).
- Megías, D. (1996). System uncertainty treatment in Generalized Predictive Control. Master's thesis. Departament d'Informàtica. Universitat Autònoma de Barcelona.
- Megías, D., J. Serrano and A. G. Kuznetsov (1999a). A systematic method to enhance the robustness of stabilising receding-horizon predictive controllers. In: *Proc. of the European Control Conference*. Karlsruhe, Germany.


- Megías, D., J. Serrano and C. de Prada (1996). Tratamiento de la incertidumbre en el GPC: diseño del polinomio T . In: *XVII Jornadas de Automática*. Santander, Spain. (In Spanish).
- Megías, D., J. Serrano and C. de Prada (1997). Uncertainty treatment in Generalized Predictive Control. In: *Proc. of the European Control Conference*. Brussels, Belgium.
- Megías, D., M. Y. El Ghoumari and J. Serrano (1999b). Extended linearised predictive control: Practical control algorithms for non-linear systems. In: *Proc. of the European Control Conference*. Karlsruhe, Germany.
- Mehra, R. K., R. Rouahni, J. Eterno, J. Richalet and A. Rault (1982). Model algorithmic control: review and recent development. In: *Eng. Foundation Conf. on Chemical Process Control II*. Sea Island, GA, USA. pp. 287–310.
- Michalska, H. and D. Q. Mayne (1993). Robust receding horizon control of nonlinear systems without differentiability of the optimal value function. *IEEE Trans. on Automatic Control* **38**(11), 1623–1633.
- Morari, M. and E. Zafiriou (1989). *Robust process control*. Prentice Hall. Englewood Cliffs, NJ, USA.
- Mosca, E. and J. Zhang (1992). Stable redesign of predictive control. *Automatica* **28**(6), 1229–1233.
- Mosca, E., G. Zappa and C. Manfredi (1984). Multistep horizon self-tuning controllers: the MUSMAR approach. In: *Proc. 9th IFAC World Congress*. Budapest, Hungary.
- Naslin, P. (1962). *Les régimes variables dans les systèmes linéaires et non linéaires*. Dunod. Paris, France. (In French).
- Nelder, J. and R. Mead (1965). A simplex method for function minimization. *The Computer Journal* **8**, 42–52.

- Oliveira, S. L. and M. Morari (1998). Contractive model predictive control with local linearization for nonlinear systems. pp. 403–431. In: Berber and Kravaris (1998).
- Oliveira, S. L., V. Nevistić and M. Morari (1995). Model predictive control for nonlinear systems subject to input constraints. Technical report. Automatic Control Laboratory, Swiss Federal Institute of Technology (ETH).
- Papavasiliou, G. C. and J. C. Allwright (1991). A descent algorithm for a min-max problem in model predictive control. In: 30th *IEEE Conference on Decision and Control*. Brighton, UK.
- Penrose, R. (1989). *The Emperor's New Mind*. Oxford University Press. Oxford, UK.
- Peterka, V. (1984). Predictor-based self-tuning control. *Automatica* **20**, 39–50.
- Propoi, A. I. (1963). Use of LP methods for synthesizing sampled-data automatic systems. *Autom. Remote Control*. (24).
- Qin, S. J. and T. A. Badgwell (1996). An overview of industrial model predictive control technology. In: *5th International Conference on Chemical Process Control - CPC V*.
- Rawlings, J. B. and K. R. Muske (1993). The stability of constrained receding horizon control. *IEEE Trans. on Automatic Control* **38**(10), 1512–1516.
- Richalet, J. (1993a). Industrial applications of model based predictive control. *Automatica* **29**(5), 1251–1274.
- Richalet, J. (1993b). *Pratique de la commande prédictive*. Hermes. Paris, France. (In French).
- Richalet, J., A. Rault, J. L. Testud and J. Papon (1978). Model predictive heuristic control: applications to industrial processes. *Automatica* **14**(5), 413–428.

- Robinson, B. D. and D. W. Clarke (1991). Robustness effects of a prefilter in generalised predictive control. *IEE Proc. Part D* **138**(1), 2–8.
- Rossiter, J. A. and B. Kouvaritakis (1993). Constrained stable generalised predictive control. *IEE Proc. Part D* **140**(4), 243–254.
- Rossiter, J. A. and B. Kouvaritakis (1994). Numerical robustness and efficiency of generalised predictive control algorithms with guaranteed stability. *IEE Proc. Part D* **141**(3), 154–162.
- Rossiter, J. A., B. Kouvaritakis and M. J. Rice (1998). A numerically robust state-space approach to stable-predictive control strategies. *Automatica* **34**(1), 65–73.
- Rossiter, J. A., J. R. Gossner and B. Kouvaritakis (1996). Infinite horizon stable predictive control. *IEEE Trans. on Automatic Control* **41**(10), 1522–1527.
- Rouhani, R. and R. K. Mehra (1982). Model algorithmic control (mac): basic theoretical properties. *Automatica* **18**, 401–414.
- Schittowski, K. (1985). NLQPL: A FORTRAN-subroutine solving constrained nonlinear programming problems. *Annals of Operations Research* **5**, 485–500.
- Scokaert, P. O. M. (1994). Constrained predictive control. PhD thesis. Department of Engineering Science, University of Oxford.
- Scokaert, P. O. M. (1997). Infinite horizon generalized predictive control. *Int. J. Control* **66**(1), 161–175.
- Scokaert, P. O. M. and D. Q. Mayne (1998). Min-max feedback model predictive control for constrained linear systems. *IEEE Trans. on Automatic Control* **43**(8), 1136–1142.

- Serrano, J. (1994). Aspectos de implementación y técnicas de robustez en CPBM. PhD thesis. Departament d'Informàtica. Universitat Autònoma de Barcelona. (In Spanish).
- Serrano, J., D. Megías and C. de Prada (1994). Control DMC de un reactor químico: Aspectos de implementación. In: *Seminario anual de Automática y Electrónica Industrial*. Tarragona, Spain. (In Spanish).
- Singh, S. (1997). *Fermat's last theorem*. Fourth Estate. London, UK.
- Skogestad, S. and I. Postlethwaite (1996). *Multivariable Feedback Control. Analysis and Design*. John Wiley and Sons. Chichester, UK.
- Soeterboek, A. R. M. (1992). *Predictive control. A unified approach*. Prentice Hall. Englewood Cliffs, New Jersey, USA.
- Staffans, O. J. (1993). The four-block model matching problem in l^1 and infinite-dimensional linear programming. *SIAM J. Control and Optimization* 31(3), 747–779.
- Thaler, G. J. and M. P. Pastel (1962). *Analysis and design of nonlinear feedback control systems*. McGraw-Hill. New York, NY, USA.
- The Mathworks (1997). *MATLAB Optimization Toolbox User's Guide*. Natick, MA, USA.
- The Mathworks (1998). *Simulink - Dynamic system simulation with MATLAB*. Natick, MA, USA.
- Vidyasagar, M. (1997). Statistical learning theory and its applications to randomized algorithms for robust controller synthesis. In: *Plenary Lectures and Mini-Courses of the European Control Conference* (G. Bastin and M. Gevers, Eds.). Chap. 6, pp. 161–189. Brussels, Belgium.

- Vidyasagar, M. (1998). Statistical learning theory and randomized algorithms for control. *December 1998 Special Issue of the IEEE Control Systems Magazine on Emerging Technologies* pp. 69–89.
- Wellstead, P. E. and M. B. Zarrop (1991). *Self-tuning systems control and signal processing*. John Wiley and Sons. Chichester, UK.
- Whitfield, A. H. (1986). Transfer function synthesis using frequency response data. *Int. J. Control* **43**(5), 1413–1425.
- Ydstie, B. E. (1984). Extended horizon adaptive control. In: *Proc. 9th IFAC World Congress*. Budapest, Hungary.
- Yoon, T.-W. (1994). Robust adaptive predictive control. PhD thesis. Department of Engineering Science, University of Oxford.
- Yoon, T.-W. and D. W. Clarke (1995a). Observer design in receding-horizon predictive control. *Int. J. Control* **61**(1), 171–191.
- Yoon, T.-W. and D. W. Clarke (1995b). A reformulation of receding-horizon predictive control. *Int. J. Systems Sci.* **26**(7), 1383–1400.
- Zadeh, L. A. and B. H. Whalen (1962). On optimal control and linear programming. *IRE Trans. on Automatic Control.* **7**(4).



Universitat Autònoma de Barcelona

Robustness aspects of Model Predictive Control

A DISSERTATION SUBMITTED IN
PARTIAL FULFILMENT FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY AT THE
UNIVERSITAT AUTÒNOMA DE BARCELONA

David Megías Jiménez

March 2000