CHAPTER 4: MODEL PREDICTIVE CONTROL.

4.1. INTRODUCTION.

Model Predictive Control (MPC) possesses many attributes which makes it a successful approach to industrial control design:

- Simplicity: The basic ideas of MPC do not require complex mathematics and are 'intuitive'.
- Richness: All of the basic MPC components can be tailored to the details of the problem in hand.
- Practicality: It is often the resolution of problems such as satisfying control- or output constraints, which determines the utility of a controller.
- Demonstrability: It works, as shown by many real applications in industry where MPC is routinely and profitably employed [20].

At present MPC is the most widely used multivariable control algorithm in the chemical process industry and in other areas [21]. While MPC is suitable for almost any kind of problem, it displays its main strength when applied to problems with:
- A large number of manipulated and controlled variables.
- Constraints imposed on both the manipulated and controlled variables.
- Changing control objectives and/or equipment (sensor/actuator) failure.
- Time delays.

The furnace model used for simulation purposes [13] consists of 17 states and 7 inputs (manipulated variables (MVs) and disturbances). Constraints exist on the controlled variables (CVs) of the EAF due to physical constraints and control objectives, and the MVs have limited ranges. MPC will therefore be used as the automatic control strategy to be compared to manual control as is currently used.

The remainder of this chapter will provide the theoretical background on MPC controller design and implementation.

4.2. BACKGROUND.

A conceptual diagram illustrating the principles of an MPC controller is shown in Figure 4.1 [20]. The heart of the controller is a model $M(\theta)$, parameterised by a set θ, which is used to predict the future behaviour of the plant. The prediction has two main components: The free response (f_r), being the expected behaviour of the output assuming zero future control actions, and the forced response (f_o), being the additional component of the output response due to the 'candidate' set of future controls (u). For a linear system, the total prediction can be calculated as $f_o + f_r$.

The reference sequence (r) is the target values the output should attain. The future system errors can then be calculated as $e = r - (f_o + f_r)$, where f_o, f_r and r are vectors of the appropriate dimensions.

An optimiser, having a user defined objective function J(e,u), is used to calculate the best set of future control actions by minimising the objective function, J(e,u). The optimisation is subject to constraints on the MVs and CVs.

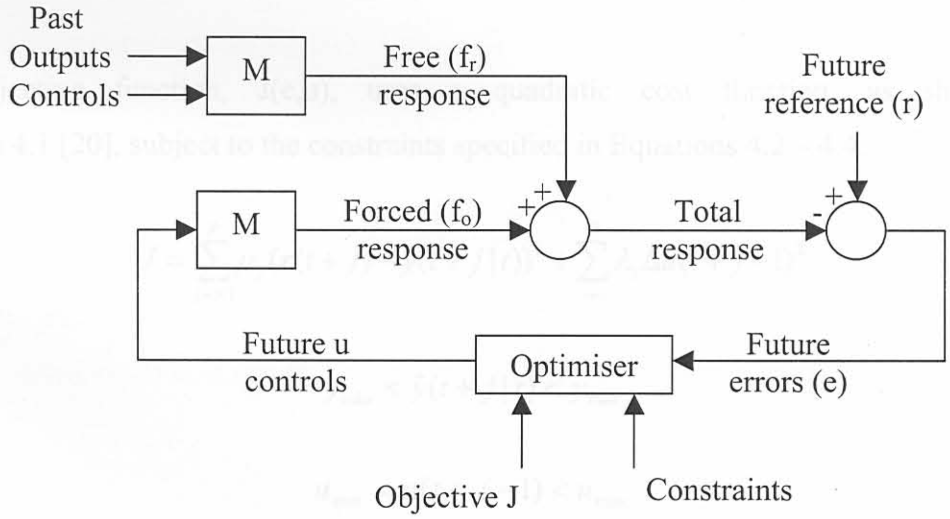

Figure 4.1. Basic Structure of MPC.

What makes MPC a closed loop control law is the use of the receding horizon approach. This implies that only the first of the set of control actions, u, is transmitted to the plant, after which the complete optimisation and prediction procedure is repeated, using the current plant output.

Another principle employed by MPC, is the use of horizons. The prediction horizon, P, specifies the number of future plant outputs to be calculated, using the model, M, the past control actions and the computed future control actions. The control horizon, C, specifies the number of future control actions to be calculated, in order to minimise the objective function, J(e,u), subject to the plant constraints. The future controls, u, will thus be a vector of dimension n × C, for n the number of manipulated variables. Only the first control actions (n × 1) will however be implemented, after which a new control sequence will be calculated.

The vectors f_r, f_o, and e will be vectors of dimension m × P, for m the number of plant outputs. The reference trajectory, r, has the dimension m × t, for t the total time for which the controller is implemented. For the calculation of e, a portion of r with dimension m × P is used to allow matrix manipulation.

The objective function, J(e,u), uses a quadratic cost function, as shown in Equation 4.1 [20], subject to the constraints specified in Equations 4.2 – 4.4.

$$J = \sum_{j=N1}^{P} \mu_j (r(t+j) - \hat{y}(t+j\,|\,t))^2 + \sum_{j=1}^{C} \lambda_j \Delta u(t+j-1)^2 \tag{4.1}$$

$$y_{\min} < \hat{y}(t+j\,|\,t) < y_{\max} \tag{4.2}$$

$$u_{\min} < u(t+j-1) < u_{\max} \tag{4.3}$$

$$\Delta u_{\min} < \Delta u(t+j-1) < \Delta u_{\max} \tag{4.4}$$

In Equation 4.1, $\hat{y}$ generally represents the future predictions of the system outputs, and $r - \hat{y}$ thus represents the predicted future errors, $\hat{e}$. Δu represents a differential control action. A differential value for u is preferred to an absolute value, as high frequency changes in u (Δu) tend to wear out actuators and might potentially cause instability. Constant high actuator values (u) however have no disadvantages to the actuators or stability, although it might influence plant operational cost (e.g. high feed rates) that need to be accounted for elsewhere.

μ_j and λ_j represents weights applied to the MVs and CVs. Weights applied to the outputs, μ_j, are used mainly to assign different priorities to different CVs. This is useful in ensuring that a CV that is much more critical than another enjoys the appropriate priority. For CVs with large differences in ranges, appropriate weights will ensure that relatively large deviations from variables with small nominal values enjoy a larger priority than relatively small deviations from variables with large nominal values.

Weights applied to variations in the MVs, λ_j, are used mainly for move suppression to prevent oscillatory behaviour. Increasing λ_j will prevent oscillation of the MVs, but large values of λ_j tend to slow down response times. Increasing λ_j thus trades system error minimisation against control signal variance [22].

In Equation 4.1, the differential manipulated variables are summed from 1 to C, the control horizon. The controlled variables are summed from N1 to P, the prediction horizon. For systems with dead time or inverse responses, the value of N1 is usually chosen large enough to prevent the inverse response, or unaffected response due to dead time, from being included in the cost function [22]. In the absence of dead time or inverse responses N1 = 1.

4.3. DESIGN STRATEGY.

The main tuning parameters are the control and prediction horizons (C and P) and the weights applied to the manipulated and controlled variables (μ and λ). Their functions will be discussed in turn.

The prediction horizon determines the number of predictions that are used in the optimisation calculations. Increasing the prediction horizon results in more conservative control action that has a stabilising effect, but it also increases the computational effort [23]. The predictions are furthermore just as good as the model used. A very large prediction horizon would thus be recommended only for a very good model and if feedback is limited.

The control horizon determines the number of future control actions that are calculated in the optimisation step to minimise the predicted errors. A large value for the control horizon, C, relative to the prediction horizon, P, tends to yield excessive control actions. A smaller value for C leads to a robust controller that is relatively insensitive to model errors [23]. Computational effort is also reduced by decreasing C.

A number of choices for the horizons have been suggested for the particular EAF model, all using a controller with a sampling interval of 1 s. Bekker [13] suggested using C = 2 and P = 6. This choice was based on defined criteria that had to be met and a trade-off between minimising computational effort and system error. Viljoen [10] continued on Bekker's research and controlled a different set of variables using additional manipulated variables. The control horizon was selected as C = 2 and the integral square error (ISE) between the setpoints and simulated plant outputs were determined for P between 5 and 8. It was found that P = 6 yielded the lowest ISE, as was also suggested by Bekker [13]. Oosthuizen [24] used a normalised ISE as criteria to determine the most suitable choices of C and P. A choice of C = 3 and P = 8 minimised the normalised ISE and was used effectively in simulations.

The choices of the three authors mentioned, all suggest that C should be relatively small compared to P (approximately 3 times smaller). Oosthuizen [24] also showed that no improvement is obtained by increasing P beyond 8, and that performance actually degrades due to modelling inaccuracies. Soeterboek suggested the following choices for N_1, C and P, for a system with dead time, d, a system order of n_A, a 5% settling time of t_s, a bandwidth of ω_b and a sampling time, $T_s = 2\pi/\omega_s$ [25]:

- $N_1 = d + 1$
- $C = n_A$
- P = integer (t_s/T_s) for a well-damped system.
- P = integer ($2\omega_s/\omega_b$) for a badly-damped system.

The computational effort is however not considered by Soeterboek, and might be an important consideration in the determination of C and P. The suggestions of Soeterboek [25] as well as the techniques described in [10, 13, 24] will be used to determine the most suitable horizons.

The other tuning parameters are the weights applied to the controlled and manipulated variables (μ and λ). A typical initial choice is $\mu = I$, the identity matrix of appropriate dimension and $\lambda = fI$, for f a tuning parameter, typically chosen as small as possible [25]. Variations in all the manipulated variables are penalised proportional to f. Increasing f thus causes less vigorous control [23]. The disadvantage of this approach is that all the manipulated variables and deviations of the controlled variables from the setpoints are penalised in equal proportions. This selection of μ and λ as discussed above would thus only be useful if the priorities of the controlled variables and the ranges of the manipulated variables are equal. Bekker [13] performed some trial and error tuning on μ and λ until the desired response was obtained. Oosthuizen [24] selected initial weights based on the ranges of the manipulated variables and the maximum errors of the controlled variables. Viljoen [10] introduced dynamic weighting that changes as the variable changes. In all three cases a lot of trial and error tuning had to be performed to get the required system response. None of these simulation studies [10, 13, 24] however took economic considerations (which complicates the weighting process further) into account.

Becerra *et al.* [26] presented three structures typically used to optimise plant performance economically, as illustrated in Figure 4.2.

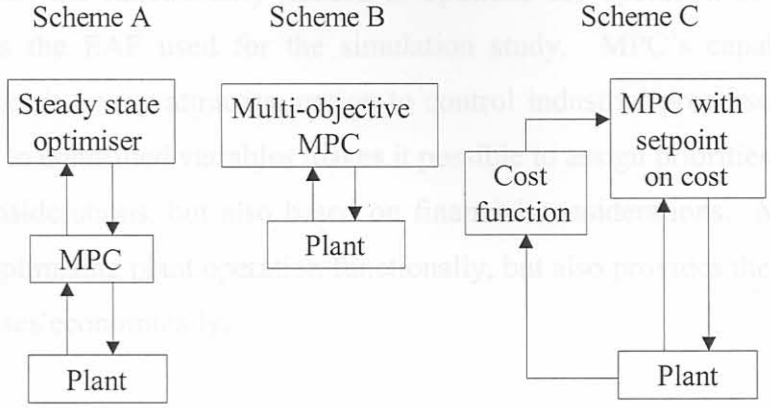

Figure 4.2. Implementation of MPC with economic objectives.

The scheme most commonly employed is Scheme A, where an upper level steady state optimiser provides setpoints for controlled variables and/or targets for manipulated variables to be used in the MPC algorithm. Scheme A has the disadvantage that plant operation is optimised only at steady state values, and no dynamic optimisation is possible. Schemes B and C addresses this problem to a certain extent, but have the disadvantage that some trade-offs need to be made between the functional and economic objectives.

This disadvantage can be avoided by implementing an MPC controller with only economic objectives. Instead of combining functional and economic objectives as suggested by Scheme B, the cost contributions of the MVs and CVs as discussed in Chapter 3 can be used to translate all functional objectives into economic objectives, prior to controller design. Minimisation of the MPC objective function would thus minimise plant operating cost, subject to the accuracy of the economic model.

The design and implementation of an economically and functionally efficient controller will be discussed in Chapter 5, using the structure of Scheme B and utilising the suggested tuning parameters as discussed in this chapter.

4.4. CONCLUSION.

MPC provides all the functionality needed to optimise the operation of a multivariable process such as the EAF used for the simulation study. MPC's capability to handle constraints makes it a very attractive option to control industrial processes. The weights that are applied to controlled variables makes it possible to assign priorities, not only based on physical considerations, but also based on financial considerations. MPC is thus not only useful in optimising plant operation functionally, but also provides the functionality to optimise processes economically.