

The Implicit Daisy-Chaining Property of Constrained Predictive Control

J.M.Maciejowski
Cambridge University Engineering Department
Cambridge CB2 1PZ, England
Tel: +44 1223 332732, Fax: +44 1223 332662
Email: jmm@eng.cam.ac.uk

26 January 1997
Revised 5 October 1997

Accepted for publication in Applied Mathematics and Computer Science, special issue on Data Processing and Process Control.

Abstract: Systems with redundant control actuators are sometimes arranged so that a new actuator comes into play if the one normally used becomes saturated. We call this ‘daisy-chaining’. Such an arrangement also provides a degree of fault-tolerance against actuator failures. This paper points out that a similar property is implicit in constrained predictive control, as it is usually formulated. Predictive control therefore has a degree of implicit fault-tolerance. In order to obtain this property the predictive control must have explicit constraints on the input levels (actuator positions) and the usual disturbance model, which results in integral action arising in the controller. The general considerations dealt with in the paper are illustrated by a simplified example, based on the liquefaction of natural gas.

Keywords: Fault-tolerance, Predictive control, Reconfiguration.

1 Introduction

In systems with redundant actuators, ‘daisy-chaining’ refers to an arrangement in which one actuator (manipulated variable) is used in normal operation, but another (or others) actuator is brought into operation if the first one saturates or fails. Figures 1 and 2 show, in block-diagram form, how this can be achieved. Figure 1 shows the use of a model of the saturation characteristic of an actuator, while figure 2 shows the use of feedback of the actuator position — for example, feedback of measured valve stem position, or of measured flow rate in a pipe. (It will be seen from these examples that the term ‘actuator’ is used in this paper to mean either a device, such as a valve or hydraulic ram, or the variable affected by such a device, such as a flow rate or control surface angle; the appropriate meaning is problem-dependent.)




Figure 1: Daisy-chained actuators

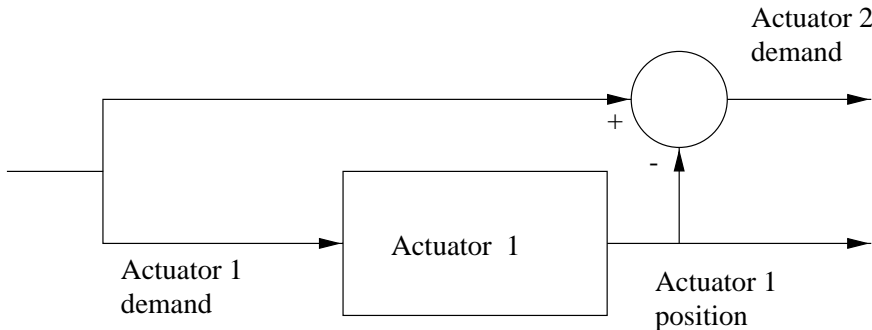


Figure 2: Daisy-chained actuators with actuator position feedback

It is apparent from these figures how daisy-chaining works in case an actuator saturates. It is probably less evident that these schemes are also effective in case the actuator used for normal operation fails, for example by getting stuck at a constant value. Figure 1 requires integral action in the controller to be effective in this case. The back-up actuator is not brought into operation immediately, but a persistent error in the controlled output leads to an increasing control signal from the integral action, until the daisy-chaining system ‘thinks’ that the normal actuator has become saturated, whereupon the back-up actuator is brought into play. The scheme shown in figure 2 does not necessarily need integral action to be effective in the event of failure, because the discrepancy between the required and actual actuator settings is immediately apparent to the daisy-chaining scheme. It should be pointed out that actuator position feedback can be ineffective for dealing with certain kinds of actuator failures. For example, the effect of valve stem travel being reduced from its normal range can be the same as that of a pipe becoming clogged downstream of the valve, but only the first failure will be detected by valve stem position feedback.

Daisy-chaining can be implemented by software, or by clever electromechanical or even purely mechanical arrangements. Conventional daisy-chaining must be deliberately added to a control system, may require additional hardware, and can provide only foreseen reconfiguration possibilities.

Constrained predictive control is a technology which is increasingly used in the process industries [1, 7]. The purpose of this paper is to demonstrate that constrained predictive control implicitly has some daisy-chaining capability, whenever there is a degree of redundancy in the plant actuators. Actually ‘daisy-chaining’ should be understood here in a slightly more general sense than the one used above, to mean a transfer of control action from faulty actuators to healthy ones. It may be, for instance, that in normal operations all the actuators are used to some extent. In this case, ‘daisy-chaining’ refers to the possibility of using the healthy actuators to a greater degree than normal when some actuators fail. In contrast with conventional daisy-chaining, the capability arises inherently from the usual constrained predictive control problem formulation, and can exploit any available redundancy, even in situations which had not been foreseen by the designer. This property enhances considerably the robustness of constrained predictive control schemes in the face of actuator saturation, and their tolerance to certain kinds of failure.

In this paper we first give a brief account of constrained predictive control, and discuss how integral action arises in the usual versions of this. We then discuss systems with redundant actuators, and concentrate on their steady-state characteristics. In section 4 we prove that a redundantly-actuated system with constrained predictive control (and with certain assumptions) cannot converge to an incorrect set-point, even if there are actuator failures, so long as those failures are compatible with achieving the required set-point. Finally, the ideas of the paper are illustrated by an example taken from gas liquefaction.

2 Constrained Predictive Control

In predictive control, an explicit ‘internal model’ is used to obtain predictions of plant behaviour over some future time interval, assuming some trajectory of control variables. The control variable trajectory is chosen by optimizing some aspect of system behaviour over this interval. Only an initial segment of the optimized control trajectory is implemented, after which the whole cycle of prediction and optimization is repeated, typically over an interval of the same length. The necessary computations are performed on-line. The optimization problem solved can include constraints, most often constraints on input levels and input rates of change, and constraints on levels of outputs and possibly internal (unmeasured) variables. Predictive control has hitherto been applied mostly in the process industries, where the explicit specification of constraints allows operation closer to constraints than standard controllers would permit, and hence operation at more profitable conditions.

Many formulations of predictive control assume that a linear time-invariant model is available in the form of a (multivariable) step or impulse response, and that predictions are generated by convolution: suppose that the multivariable step response sequence is given by $\{g_i : i = 0, 1, \dots\}$, that the (control) input vector at time k is $u(k)$ and that the (to be controlled) output vector at time k is $y(k)$. Also let $\Delta u(k) = u(k) - u(k-1)$ be the change in the input at time k . Then

the output is given by

$$y(k) = \sum_{i=-\infty}^k g_{k-i} \Delta u(i) + d(k) \quad (1)$$

where it has to be assumed that the open-loop system is asymptotically stable for this to be valid, and $d(k)$ is assumed to be a disturbance acting on the output.

In this case predictions of the output are computed by

$$\hat{y}(k+j|k) = \sum_{i=k+j-N}^{k+j} g_{k+j-i} \Delta u(i) + \hat{d}(k+j|k) \quad (2)$$

where N is a relatively large integer, and $\hat{d}(k+j|k)$ is some estimate of $d(k+j)$. Usually the disturbance is estimated as

$$\hat{d}(k|k) = y(k) - \hat{y}(k) \quad (3)$$

and it is assumed that future disturbances are the same as the current one:

$$\hat{d}(k+j|k) = \hat{d}(k|k). \quad (4)$$

The convolution model is an inefficient one, since the same model can be represented much more compactly in either transfer function or state-space form. Furthermore, representing the system by a model of this kind removes the restriction to stable models. However, the representation chosen for the model is not important here. Later we shall assume a state-space model.

Predictive control works by choosing control actions to minimise some cost function, such as

$$J(k) = \sum_{i=N_1}^{N_2} \|\hat{y}(k+i|k) - r(k+i)\|_Q^2 + \sum_{i=1}^{N_u} \|\Delta u(k+i)\|_R^2 \quad (5)$$

subject to constraints such as

$$|\Delta u_j(k+i)| \leq V_j \quad (6)$$

$$|u_j(k+i)| \leq U_j \quad (7)$$

$$|\hat{y}_j(k+i|k)| \leq Y_j \quad (8)$$

where $r(k)$ is some set-point trajectory for $y(k)$. It is assumed that the control signals are constant after the end of the optimisation horizon, namely that $\Delta u(k+i) = 0$ for $i > N_u$.

The cost function penalises non-zero changes $\Delta u(k)$ in the control signals, rather than the control signals $u(k)$ themselves, since the required steady-state values of $u(k)$ are not known in advance. Penalising non-zero $u(k)$ would ‘drag’ the control signals away from the required steady-state values, thus preventing integral action, for instance.

When a linear model and quadratic cost is used, the resulting controller is linear time-invariant providing that either no constraints are active, or that a fixed set of constraints is active. (For each such set, a different linear control law results.) Thus the control law can be a linear law for long periods of time. However, when hard constraints are approached the controller can behave

in a very nonlinear way. In particular, it may react mildly to a disturbance which drives the system away from constraints, but very sharply to a disturbance of similar magnitude but in the opposite direction, which drives the system towards constraints.

Predictive control, as described above, exhibits ‘integral action’, namely it has the capability of tracking constant set-point vectors without error in the steady-state, and of asymptotically rejecting constant disturbances [2]. In order to understand the ‘daisy-chaining’ property later, it is important to understand how this ‘integral action’ property of predictive control arises. It is often stated, erroneously, that the property is due to the inherent ‘integration’ in the predictive control law, since this computes the increments $\Delta u(t)$, whereas the control applied to the plant is $u(t)$. This is manifestly false, since this integration is exactly offset by the additional ‘differentiation’ inherent in computing Δu instead of u .

It is, in fact, the combination of penalising $\|\Delta u\|$ (rather than $\|u\|$) in the objective function, *and the disturbance model* which produces the integral action. The penalisation of $\|\Delta u\|$, rather than $\|u\|$, allows the optimal solution to be unprejudiced by any prior judgement about the required steady-state actuator settings. It is certainly necessary to the emergence of the integral action property, but it is far from being sufficient. Suppose that the controlled output vector y settles to a value different from the set-point vector. The predictive control applies a set of actuator settings which, according to its internal model, should cause the output vector to equal the set-point vector. An output disturbance is therefore estimated (as the difference between the two), *and is assumed to persist unchanged into the future*. The actuator settings are therefore adjusted, to correct for this perceived error. If an error still persists, this process is repeated, with the actuator settings increasing, until the output vector reaches the set-point vector. The fact that only differences in the input vector appear in the cost function implies that the optimization is indifferent to the steady-state settings of the actuators, and hence this process is not inhibited by the optimization. There is, therefore, no explicit integrating element in the predictive controller, but the same effect is achieved by an ‘integration’ of the estimated output disturbance, and computation of the control vector to counteract this increasing estimate.

3 Redundantly-actuated systems

Some control systems have redundant actuators, in the sense that there are more degrees of control freedom than control objectives. Examples of this arise in some modern aircraft, for example, with conventional control surfaces, such as rudders, sometimes consisting of several parts, and others, such as elevators and ailerons, being independently actuated, so that elevators can be used to produce roll torques and ailerons can be used to provide pitching moments (the opposite of their usual functions). Also, wing spoilers (air brakes) are commonly used asymmetrically to enhance the rolling moment at low speeds [3].

Examples in the process industries are less apparent, because such redundancy is less commonly designed in deliberately, except perhaps in some safety-critical sectors such as nuclear power plant. But inevitably there is considerable redundancy in a complex plant: malfunctioning of one valve may be compensated by another in a different part of the plant, or by changing the speed of a compressor, for example.

Of course it is rare for a control actuator to be completely redundant, in the sense that it has no function in normal operation. (Deliberate provision of back-up capability in safety-critical

applications is an exception.) But there may be redundancy with respect to certain control objectives. In particular, redundancy with respect to steady-state set-point tracking objectives is quite common, and it is redundancy in this sense that we will concentrate on in this paper.

Let u_s denote a steady-state (constant) value of the plant input vector, namely of the plant actuators. Similarly let d_s denote the value of a plant disturbance vector, again assumed constant, and let y_s denote the resulting steady-state value of the vector of controlled outputs — which we assume here to be constant. In general u_s , d_s and y_s are related by some nonlinear function P_s :

$$y_s = P_s(u_s, d_s). \quad (9)$$

We assume here that any factors affecting the steady-state, other than the actuators, are included in the disturbance vector d_s , and for simplicity we also assume that the steady-state map P_s does not change with time.

Consider the set $U(P_s, y_s, d_s)$ of all constant input vectors which result in the steady-state output vector y_s in the presence of the disturbance vector d_s :

$$U(P_s, y_s, d_s) = \{u_s : y_s = P_s(u_s, d_s)\}. \quad (10)$$

This set may consist of a single point, or a set of isolated points, or a continuum of points. For the purposes of this paper, we shall say that the plant is *redundantly actuated* if $U(P_s, y_s, d_s)$ is a manifold of dimension at least 1. Actually we will not need most of the properties of a manifold, such as smoothness at each point. In practical cases the set will be differentiable at least once in the neighbourhood of feasible plant operating points. If the plant is redundantly actuated, we shall call $U(P_s, y_s, d_s)$ the *steady-state manifold*. The dimension of the manifold measures the degree of redundancy: it is a 1-dimensional curve if there is 1 redundant actuator, a 2-dimensional surface if there are 2 redundant actuators, etc.

If the plant behaviour is linear, has m independently manipulated actuators, ℓ disturbances, and p controlled outputs to be held at set-points, then P_s is a $p \times (\ell + m)$ matrix. Let it be partitioned as $P_s = [P_{su}, P_{sd}]$, so that $y_s = P_{su}u_s + P_{sd}d_s$. If $m \geq p$ and $\text{rank}(P_{su}) = p$, then the steady-state manifold $U(P_s, y_s, d_s)$ is an affine variety of dimension $p - m$, which is a translation of the null space of P_{su} .

Suppose that the plant behaviour is linear. Let r_s be a constant set-point vector, and let u_{s0} be the point in $U(P_s, r_s, d_s)$ at which the plant is supposed to operate under normal conditions. If u is the actual value of the input vector (not necessarily constant with time), it will be helpful to consider the difference $u - u_{s0}$. Let the orthogonal projection of $u - u_{s0}$ onto $U(P_s, r_s, d_s)$ be v , and $w = (u - u_{s0}) - v$. That is, w is the ‘component’ of $u - u_{s0}$ in the steady-state manifold, and v is the ‘component’ orthogonal to this manifold. If u settles to a constant value u_{s1} , with corresponding values of v and w being v_s and w_s , respectively, then $\|v_s\| = 0$ indicates that the set-point vector is being tracked without error, and $\|w_s\|$ then measures the extent of any actuator reconfiguration that has taken place in order to achieve this. If the behaviour is nonlinear then one needs to define u and v with respect to the local tangent space of the steady-state manifold, and a suitable replacement for $\|w_s\|$ remains to be investigated — for some purposes the chordal distance in the Euclidean ‘input space’ will be appropriate, whereas for others a geodesic distance in the manifold may be more suitable.

4 The Daisy-Chaining Property

We focus here on daisy-chaining in response to actuator failures. Let us call the particular steady-state manifold obtained when $y_s = r_s$ the *set-point manifold*. The constraints on input levels restrict the feasible region of this set-point manifold. Call this feasible region U_1 . In the case of independent constraints on each component of the input vector, U_1 is a hyper-rectangle. Actuator failures also restrict the attainable region of the set-point manifold. Call this attainable region U_2 . In the case of the i 'th actuator being stuck at a particular value, U_2 is a section of the set-point manifold, orthogonal to the i 'th basis vector. If the intersection of U_1 and U_2 is not empty we say that the actuator failure is *compatible with the set-point specification*.

We assume that the controlled plant is linear and described by:

$$x(k+1) = A_p x(k) + B_p v(k) \quad (11)$$

$$y(k) = C_p x(k) \quad (12)$$

The internal model used by the predictive controller assumes a constant output disturbance:

$$z(k+1) = A_m z(k) + B_m u(k) \quad (13)$$

$$d(k+1) = d(k) \quad (14)$$

$$y(k) = C_m z(k) + d(k) \quad (15)$$

As in [6], the constant disturbance is modelled as a constant state, in order to discuss its estimation in a standard way. Assuming a constant input disturbance would be more natural for dealing with actuator failures, but we assume output disturbances for consistency with popular versions of predictive control, such as DMC and GPC.

Notice that the controller produces a control signal $u(k)$, whereas the plant receives the control signal $v(k)$. We model actuator failures as:

$$v(k) = S u(k) + \beta \quad (16)$$

where S is a square matrix and β is a constant vector. This allows failures to zero, to non-zero values, and actuator gain changes (but not changes in actuator dynamics). Jammed actuators are represented by $S = \text{diag}\{1 - s_i\}$, $\beta_i = s_i \gamma_i$, and $s_i \in \{0, 1\}$, where γ_i is the value at which the i 'th actuator is jammed. Notice also that for the time being we assume possibly quite different dynamics in the internal model and the plant.

We take the cost function to be as defined in (5), and we assume for simplicity that the only constraints are on the values of the control signal: $u(k) \in U$, where U is a compact, convex set. We further assume that the constraints on the real plant are the same as those known to the predictive controller: $v(k) \in U$. Now we assume that there is an actuator failure which is compatible with the set-point specification, namely that there exists $u_s \in U$, such that

$$x_s = A_p x_s + B_p (S u_s + \beta) \quad (17)$$

$$r_s = C_p x_s \quad (18)$$

Finally we assume that $Q > 0$, so that any deviation from the set-point is penalised.

With these assumptions, we will show that, with large enough N_u and N_2 , the only possible equilibrium of the closed-loop system is at $y_s = r_s$. This is analogous to the basic property

of integral control. To do this we will first show that at an equilibrium the predictor correctly predicts the steady-state plant output, and then that the optimiser can ‘see’ a better solution, unless it is already at the set-point. Finally we shall show that the presence of input constraints implies that the plant moves away from the equilibrium, unless it is already at the set-point.

Lemma 1 *Suppose that the controller output is constant, $u(k) = \bar{u}$, that the plant is at an equilibrium $x(k) = \bar{x}$, $y(k) = \bar{y}$, and that a standard, asymptotically stable Luenberger observer is used to estimate the model state $[z(k)^T, d(k)^T]^T$. Then the output predicted by the internal model converges to \bar{y} .*

Remark 1 *Note that this would be a completely standard result if we assumed that the plant and model were the same, namely $A_p = A_m$, etc. But we do not assume this here.*

Proof: The standard Luenberger observer is

$$\begin{bmatrix} \hat{z}(k+1|k) \\ \hat{d}(k+1|k) \end{bmatrix} = (\mathcal{A} - LC) \begin{bmatrix} \hat{z}(k|k-1) \\ \hat{d}(k|k-1) \end{bmatrix} + \mathcal{B}u(k) + Ly(k) \quad (19)$$

where

$$\mathcal{A} = \begin{bmatrix} A_m & 0 \\ 0 & I \end{bmatrix}, \quad \mathcal{B} = \begin{bmatrix} B_m \\ 0 \end{bmatrix}, \quad C = [C_m, I].$$

Since we assume that the observer is asymptotically stable, the estimated state converges (with k) to $[\bar{z}^T, \bar{d}^T]^T$, for some pair \bar{z}, \bar{d} . Hence the estimated output converges to $\hat{y} = C[\bar{z}^T, \bar{d}^T]^T$. Now, if L is partitioned conformally with $[z^T, d^T]^T$, so that $L = [L_z^T, L_d^T]^T$, then from (19) it follows that $L_d \hat{y} = L_d \bar{y}$. Now, assuming that L_d is nonsingular, it follows that $\hat{y} = \bar{y}$. Note that this is a harmless assumption, since if L_d were singular, then an arbitrarily small perturbation would make it nonsingular without losing asymptotic stability of the observer. (Also note, as pointed out by [6], that if the plant is open-loop stable and the observer used is a Kalman filter, then $L_d = I$.) Hence the predicted value of the output for all future times converges to \bar{y} , since this is obtained by iterating (19), which has a fixed point at $(\bar{z}, \bar{d}, \bar{u}, \bar{y})$. \square

Now we show that, if the actuator failure is compatible with the set-point specification, then there is always an output disturbance estimate which makes the control signal required (by the plant) from the controller, u_s , consistent with the set-point, r_s . This establishes that an equilibrium solution exists which gives the correct set-point, and is consistent with the model and constraints known to the controller. For simplicity we assume that the model does not contain any integrators.

Lemma 2 *There exists a pair (z_s, d_s) , such that*

$$z_s = A_m z_s + B_m u_s \quad (20)$$

$$r_s = C_m z_s + d_s \quad (21)$$

providing that $(I - A_m)^{-1}$ exists, where u_s is defined in (17).

Proof: Rewrite the equations as

$$\begin{bmatrix} I - A_m & 0 \\ C_m & I \end{bmatrix} \begin{bmatrix} z_s \\ d_s \end{bmatrix} = \begin{bmatrix} B_m u_s \\ r_s \end{bmatrix} \quad (22)$$

from which the result is obvious. \square

Remark 2 *If the model does contain integrators, then a solution will exist providing that*

$$\begin{bmatrix} B_m u_s \\ r_s \end{bmatrix} \in \text{span} \begin{bmatrix} I - A_m & 0 \\ C_m & I \end{bmatrix} \quad (23)$$

This is a compatibility condition relating the (failed) plant and the model, which essentially says that the integrators in the plant must be modelled correctly, or at least ‘compatibly’.

We now need a slightly different Lemma, which says that, if the plant-controller combination is sitting at an equilibrium, away from the set-point, then a pair (z_f, u_f) exists which is predicted to drive the output to the set-point, in the presence of the currently estimated disturbance \bar{d} . It is easy to see that such a pair must solve the equation:

$$\begin{bmatrix} I - A_m & -B_m \\ C_m & 0 \end{bmatrix} \begin{bmatrix} z_f \\ u_f \end{bmatrix} = \begin{bmatrix} 0 \\ r_s - \bar{d} \end{bmatrix} \quad (24)$$

and so a solution will certainly exist if

$$\begin{bmatrix} 0 \\ r_s - \bar{d} \end{bmatrix} \in \text{span} \begin{bmatrix} I - A_m & -B_m \\ C_m & 0 \end{bmatrix} \quad (25)$$

However, we need the solution to be *feasible*, that is, we also need $u_f \in U$. It seems necessary to introduce this as an independent assumption, unless stronger assumptions are made about the closeness of the model to the plant. In the following Lemmas we proceed to make a rather strong assumption in this respect: we assume that the plant and model matrices are identical.

Lemma 3 *Suppose that the plant and model are described by (11–12) and (13–15), respectively, and that $A_m = A_p$, $B_m = B_p$, and $C_m = C_p$. Suppose also that $(I - S)\beta = \beta$. Then under the conditions described in Lemma 1 the estimated output disturbance is*

$$\bar{d} = C_p(I - A_p)^{-1}B_p(I - S)(\beta - \bar{u}). \quad (26)$$

Remark 3 *The assumption on S and β is valid if one or more actuators are jammed. In this case S is an orthogonal projector, and β is orthogonal to the range space of S , so that $S\beta = 0$ and $(I - S)\beta = \beta$.*

Remark 4 *Note that $C_p(I - A_p)^{-1}B_p = P_{su}$, as defined in the previous section. So we can also write $\bar{d} = P_{su}(I - S)(\beta - \bar{u})$.*

Proof: From (11–12) and (16) we have

$$\bar{y} = C_p(I - A_p)^{-1}B_p(S\bar{u} + \beta) \quad (27)$$

and from (13–15) we have

$$\bar{y} = C_p(I - A_p)^{-1}B_p\bar{u} + \bar{d}. \quad (28)$$

Hence

$$\bar{d} = C_p(I - A_p)^{-1}B_p[\beta - (I - S)\bar{u}] \quad (29)$$

$$= C_p(I - A_p)^{-1}B_p(I - S)(\beta - \bar{u}). \quad (30)$$

□

Lemma 4 *Suppose that the assumptions of Lemma 3 hold. Suppose further that the plant actuator failure is compatible with the set-point specification and that $Su_s + (I - S)\bar{u} \in U$ (where u_s is defined in (17) and U is the set of admissible inputs). Then there exists an admissible u_f , as defined above.*

Remark 5 *The assumption that $Su_s + (I - S)\bar{u} \in U$ is rather restrictive. It holds, for example, if U is defined by component-wise restrictions on the magnitudes of actuator positions: $U_i^- \leq u_i(k) \leq U_i^+$. It should be possible to relax this assumption, and to exploit the assumption that U is a convex set which contains the points u_s , \bar{u} , $Su_s + \beta$, and $S\bar{u} + \beta$.*

Proof: From (24), Lemma 3, and (17–18) we have

$$P_{su}u_f = r_s - \bar{d} \quad (31)$$

$$= P_{su}(Su_s + \beta) - P_{su}(I - S)(\beta - \bar{u}) \quad (32)$$

$$= P_{su}[Su_s + (I - S)\bar{u}] \quad (33)$$

Hence $u_f = Su_s + (I - S)\bar{u}$ is a solution, and $u_f \in U$. □

Remark 6 *It may seem that this proof requires the controller to ‘know’ about the actuator failure, since we have used $r_s = P_{su}(Su_s + \beta)$. But this is not so. All that is required is that an admissible control input exists which the controller predicts would hold the plant at r_s in the absence of disturbances. $Su_s + \beta$ is such an input, and the controller can find it without knowing S or β .*

Lemma 5 *Suppose that the predictive controller and plant together are at an equilibrium as defined in Lemma 1. Let the corresponding value of the cost (5) be \bar{J} . Then, for large enough N_u and N_2 , there exists a sequence of admissible control moves $\{\Delta u(k+i) : i = 0, \dots, N_u\}$, with associated cost $J_u(k)$, such that $J_u(k) < \bar{J}$. (In this case an admissible control move $\Delta u(k)$ is one which results in $u(k) \in U$.)*

Proof: In Lemma 1 we established that, if the plant-controller combination is at an equilibrium, then the predicted output converges to the actual plant output. Therefore $J(k)$ converges (as k increases) to $\bar{J} = (N_2 - N_1)\|\bar{y} - r_s\|_Q^2 > 0$ if $\bar{y} \neq r_s$.

Suppose that the plant (and hence model) is open-loop stable and that $u(k+i) = u_f$ for all $i > 0$ (which is admissible), so that

$$J_u(k) = \sum_{i=N_1}^{N_2} \|\hat{y}(k+i|k) - r_s\|_Q^2 + \|u_f - \bar{u}\|_R^2 \quad (34)$$

and hence

$$\bar{J} - J_u(k) = (N_2 - N_1) \|\bar{y} - r_s\|_Q^2 - \sum_{i=N_1}^{N_2} \|\hat{y}(k+i|k) - r_s\|_Q^2 - \|u_f - \bar{u}\|_R^2. \quad (35)$$

The first term in this expression increases linearly with $N_2 - N_1$, while the second term converges to some finite limit as N_2 increases, since $\hat{y}(k+i|k)$ converges exponentially (with i) to r_s . Therefore $\bar{J} - J_u(k) > 0$ for N_2 large enough (assuming N_1 fixed).

Now suppose that the plant (and model) is unstable. Then the unstable modes must be driven to their new equilibrium values within N_u steps, and the control must remain at u_f thereafter. Note that without constraints this could certainly be done if N_u were at least as large as the number of unstable modes. With constraints on $u(k)$ it may be necessary to have a value of N_u larger than this, but there always exists some value of N_u for which it is possible. So in this case we have

$$\begin{aligned} \bar{J} - J_u(k) = & (N_2 - N_1) \|\bar{y} - r_s\|_Q^2 - \sum_{i=N_u+1}^{N_2} \|\hat{y}(k+i|k) - r_s\|_Q^2 - \\ & \sum_{i=N_1}^{N_u} \|\hat{y}(k+i|k) - r_s\|_Q^2 - \sum_{i=1}^{N_u} \|\Delta u(k+i)\|_R^2. \end{aligned} \quad (36)$$

Now the first two terms correspond to the first two terms in (35). The last two terms are new, but again are fixed for a fixed control sequence, so that again $\bar{J} - J_u(k) > 0$ if N_2 is large enough. \square

Remark 7 For more general constraint sets, involving state constraints, it would also be necessary to postulate large enough N_1 to ensure feasibility, and to consider questions of constrained stabilizability — see [8].

Remark 8 In general a large enough value of N_2 is also necessary to ensure closed-loop stability. Following the work of [8] and later developments, there seems little reason for using $N_2 < \infty$ nowadays [4, 5].

Remark 9 If the controller applied a control sequence that converged on $Su_s + (I - S)\bar{u}$, the value of u_f found in the proof of Lemma 4, then the plant would be driven to the correct set-point, since then we would have $Su_f = Su_s$. However, the controller is very likely to find what it thinks is a better steady-state control input. Any input which differs from this value of u_f by an element of the kernel of P_{su} will be predicted to give the same steady-state output. In most cases there will be a compact set of such inputs which will be admissible.

So far we have shown that if the plant-controller combination is at an incorrect equilibrium ($\bar{y} \neq r_s$) then the controller will eventually change the control sequence which it computes, and will therefore apply a new control signal $u(k+1)$ to the plant. Since the plant has an actuator failure, however, it could happen that the control signal received by the plant does not change (namely $Su(k+1) + \beta = S\bar{u} + \beta$), so that the plant could remain at the incorrect equilibrium.

In particular, this will occur if the controller attempts to drive the plant to the set-point by manipulating the failed actuator, which would be quite likely if the controller did not have a model of actuator constraints. As can be seen from Lemma 3, the estimated disturbance \bar{d} can be attributed to an input disturbance in a direction orthogonal to the range of S — not surprisingly, since that is exactly what an actuator failure is. It is therefore natural that the controller should try to correct for this by altering the input in the same direction, which unfortunately has no effect on the faulty plant. Now we show that this situation cannot persist indefinitely if the set of admissible controls, U , is bounded. Again, we make two assumptions to simplify the proof: that only one actuator has failed, and that the open-loop controller has no periodic modes. Both can probably be relaxed. Recall that, providing the solution remains unconstrained, the predictive controller is a linear, time-invariant system — see [9], for example.

Lemma 6 *Suppose that the predictive controller and plant together are at an equilibrium as defined in Lemma 1. Assume that the open-loop predictive controller, when all constraints are inactive, has no periodic modes. Assume also that $\dim(I - S) = 1$. Then, if U is bounded, the plant cannot remain at the equilibrium indefinitely (unless $\bar{y} = r_s$).*

Proof: Lemma 5 shows that a feasible input trajectory exists which would move the plant away from the equilibrium, if the internal model were correct. We assume that the predictive controller finds this solution and applies its initial move to the plant. Either the plant moves away from the equilibrium, or not. If not, the controller is effectively running open-loop, with the constant plant output appearing as a constant disturbance. Under the assumed conditions the open-loop controller is linear, time-invariant, and all inputs to it are constant. Its output therefore either converges to a constant, diverges, or converges to a periodic trajectory. Lemma 5 has already shown that it cannot converge to a constant. By assumption we exclude the possibility of a periodic trajectory. Its output must therefore diverge.

Since the set of admissible inputs, U , is bounded, the controller output $u(k)$ will eventually reach the boundary of U . As stated earlier, up to this point the changes in the controller output will be in $\text{span}(I - S)$, and by assumption this is a 1-dimensional subspace. But Lemma 4 showed that $u_f = Su_s + (I - S)\bar{u}$ is a feasible steady-state input which would hold the plant at the correct steady-state. Either the controller immediately applies an input signal to the plant for which the change is not in $\text{span}(I - S)$, in which case the plant output changes, or it maintains a constant input (which is at the boundary of U) for some time, in which case it will eventually compute a different input signal, by Lemma 5; this change in this signal will necessarily be not in $\text{span}(I - S)$, so the plant output will change. \square

Remark 10 *With more than one actuator failure one would have to deal with the possibility of the input ‘wandering around’ on the boundary of U while remaining in $\text{span}(I - S)$.*

Remark 11 *The assumption on periodic modes of the controller can probably be removed. But with more elaborate disturbance models, one might deliberately introduce periodic modes into the controller.*

Putting everything together we have proved the following:

Theorem 1 *Assume a plant defined by (11–12) with an actuator failure defined by (16), and a predictive controller which minimises (5) subject to the constraints $u(k) \in U$, using the internal model (13–15) and a Luenberger observer. U is a bounded, compact, convex set. A constant set-point is demanded, which is consistent with U , and the actuator failure is compatible with this set-point. Then, under the assumptions made in Lemmas 2, 3, 4, 5, and 6, the only possible equilibrium point of the closed-loop system is at the correct set-point.*

Remark 12 *We have not shown that the closed loop will necessarily converge to the correct set-point, but in a sense that is a different issue. We have been guided by the analogy with integral action, where the essential property is again that convergence to the wrong set-point is impossible.*

The essential ingredients of the property we have proved are the disturbance model, which gives, in essence, ‘integral action’, and a model of actuator constraints, which gives the ‘reconfiguration’, or ‘daisy-chaining’. Note that this reconfiguration occurs without any re-estimation of the internal model. Of course any additional information, such as direct measurement of the i ’th actuator’s setting, or model re-estimation showing that the i ’th actuator has become ineffective, can be used to make the reconfiguration occur more quickly.

5 Example

Figure 3 shows a simplified version of an ‘auto-cascade’ refrigeration process used for liquefying natural gas [10]. A mixture of two refrigerants is compressed by the compressor A and partially condensed in the condenser B . We shall call the refrigerants #1 and #2. (One can think of #1 as mostly propane and #2 as mostly methane, but the simple two-stage process we describe would not work with any practical refrigerants) The remaining vapour (mostly #2) is cooled in the heat exchanger C and condensed in heat exchanger D . The condensate from B (mostly #1) is cooled in heat exchanger C and expanded to low pressure through valve E , causing its temperature to fall (to about $-40\text{ }^\circ\text{C}$). The condensed #2 is expanded to low pressure through valve F , its temperature falling (to about $-160\text{ }^\circ\text{C}$) in consequence. After expansion, the resulting two-phase (vapour/liquid) #2 passes through heat exchanger D again, where the liquid phase evaporates, obtaining the required enthalpy of evaporation from the natural gas and from the #2 liquid coming from C , thus cooling them. The #2 vapour is then mixed with the two-phase #1 coming from valve E and passes through heat exchanger C , where the liquid #1 evaporates, cooling the incoming natural gas, the #2 vapour flowing from B , and the #1 liquid flowing from B . Finally the vapour mixture is returned to the compressor A . Note that the streams running from left to right in figure 3 become cooler through the process, whereas the stream running from right to left becomes warmer. (This simplified process is not, in fact, thermodynamically feasible, because the range of temperatures is too great to be attained by only two stages of heat exchange with practical refrigerants. In practice more heat exchangers are used, with more complex mixtures of refrigerants.)

The available control inputs are the settings of the two valves E and F , the power supplied to the compressor A , and the setting of the valve G regulating the flow of liquid natural gas (LNG).




Figure 3: Process for liquefying natural gas

For successful operation, it is important to control the composition and flow rate of the liquid in each of the evaporators, and to supply power to the compressor at a rate which maintains the required pressure ratio. If there is a shortage of liquid in the evaporators then the suction pressure falls and the compressor must raise the vapour through a larger pressure ratio, which is inefficient. On the other hand, if there is too much liquid in an evaporator then unevaporated liquid leaves the evaporator, which then gains enthalpy of evaporation at an unintended, and hence unproductive, point in the cycle. The composition on entry to each evaporator must also be correct, so that the profile of its evaporation curve will yield a positive temperature difference between the cooling and warming streams along the length of each heat exchanger [11]. The temperature of the LNG being delivered to storage also needs to be controlled within tight limits,

In practice, control is effected by regulating the level of liquid in the condenser B , the pressure ratio across the compressor A , and the temperature of the LNG being sent to storage. Since there are 4 control actuators and 3 controlled variables, the set-point manifold in this case is a 1-dimensional curve. The normal operating point on this curve is established by economic considerations, which usually require the production rate of LNG to be maximised, so the plant is run as near as possible to full capacity.

Suppose that valve E sticks in a ‘too-closed’ position. The mass flow rate through evaporator C reduces, causing the liquid level in B to rise. The plant can be returned to its set-point vector by reducing the compressor power (which essentially determines the product of refrigerant mass flow rate and pressure ratio) and closing valve F , producing the correct value of the pressure ratio, but at a reduced refrigerant flow rate. Now the rate of heat transfer from the natural gas to the refrigerant is reduced, so the LNG temperature rises, but this can be corrected by closing valve G so that the LNG delivery rate is reduced. Thus the correct set-points are held and the plant continues to operate, but at a less profitable level.

In [12] it is shown, for a nonlinear model of a more elaborate, 3-stage, version of this process, that such a valve failure is indeed corrected in the anticipated manner. In this case the prediction horizon was $N_2 = 40$ (with $N_1 = 1$) and the control horizon was $N_u = 8$, the control update

time being 10 seconds.

In this process, failure of valve E is less likely to occur than blockage of the pipe downstream of the valve, for example by ingress of water vapour and subsequent freezing. The effect on the process would be similar to the scenario examined above. However, whereas a valve failure may be detectable by valve stem position feedback, and hence corrected more quickly than by ‘integrating’ an estimated flow disturbance, a pipe blockage could not be detected in such a direct manner. Hence a pipe blockage could only be corrected by the predictive control analogue of the arrangement of figure 1.

The changes in the control actuator settings, and the reduced delivery rate, are signals which indicate that there is a problem, although it has otherwise been masked by successful reconfiguration of the controls. If the stuck position of valve E is too far away from its usual setting, so that the reconfiguration takes the process away considerably from its normal operating condition, then controlling the level in B and the pressure ratio may not hold the compositions and flow rates at satisfactory levels, in which case some ‘higher-level’ action may need to be taken to compute new set-points for the controlled variables.

In [13] it is shown that similar reconfiguration occurs in a realistic model of rudder failure in an aircraft, when predictive control is used.

6 Conclusion

The fact that constrained predictive controllers exhibit daisy-chaining when actuators saturate is no surprise. Indeed it has long been claimed as one of the benefits of the technology, that it can accommodate constraints on control signals. This paper has concentrated on showing that predictive controllers also exhibit ‘daisy-chaining’, or actuator reconfiguration, in response to some actuator failures. This has been previously reported [7], but the phenomenon has apparently not previously been analysed.

The paper has introduced formal definitions of the notions of redundant actuation and the steady-state manifold, in order to make precise treatment of the associated phenomena possible. A quantitative measure of actuator reconfiguration has been proposed, for the linear case at least, essentially by introducing coordinates for actuator settings in and orthogonal to the steady-state manifold.

It has been shown that, if an actuator failure occurs which is compatible with the set-point specification, then a constrained predictive controller with a model of the actuator saturation characteristics, and a disturbance model which gives rise to integral action, will reconfigure the actuators so as to track the set-point without error — although we have not examined convergence to the correct set-point in these circumstances.

Constrained predictive controllers therefore have a considerable degree of fault-tolerance built into them. Unlike the conventional schemes for daisy-chaining, such as those shown in figures 1 and 2, this can be achieved without special-purpose devices or algorithms, and can offer protection even against unanticipated failures.

References

- [1] Clarke,D.W. (ed), *Model-Based Predictive Control*, (Oxford: Oxford University Press), 1994.
- [2] Mosca,E, *Optimal, Predictive and Adaptive Control*, (Englewood Cliffs,NJ: Prentice-Hall), 1995.
- [3] Raymond,E.T. and Chenoweth,C.C, *Aircraft Flight Control Actuation System Design*, (Warrendale,PA: Society of Automotive Engineers Inc.), 1993.
- [4] Morari,M, Model predictive control: multivariable control technique of choice in the 1990s?, in: Clarke,D.W. (ed), *Model-Based Predictive Control*, (Oxford: Oxford University Press), 1994.
- [5] Morari,M, and Lee,J.H, Model predictive control: past, present and future, *Proc. Joint 6th Int'l Symp. on Process Systems Engineering (PSE97)*, 1997.
- [6] Muske,K.R, and Rawlings,J.B, Model predictive control with linear models, *AIChE Journal*, **39**, 262–287, 1993.
- [7] S.J.Qin and T.A.Badgwell, An overview of industrial model predictive control technology, *Proc. Conf. Chemical Process Control, CPC V*, Lake Tahoe, 1996.
- [8] Rawlings,J.B, and Muske,K.R, The stability of constrained receding horizon control, *IEEE Trans. Automatic Control*, **38**, 1512–1516, 1993.
- [9] Zafiriou,E, Robust model predictive control of processes with hard constraints, *Computers in Chemical Engineering*, **14**, 359–371, 1990.
- [10] Gosney,W.B, *Principles of Refrigeration*, (Cambridge: Cambridge University Press), 1982.
- [11] Linnett,D.T, and Smith,K.C, The process design and optimisation of a mixed refrigerant cascade plant, *Proc. Int'l. Conf. on Liquefied Natural Gas, London, March 1969*, (London: Inst. Mech. Eng.)
- [12] Rowe,C.A, The fault-tolerant capabilities of constrained model predictive control, *M.Phil. Thesis, Cambridge University*, 1997.
- [13] Maciejowski,J.M, Reconfigurable control using constrained optimization, *Proc. European Control Conf.*, Brussels, 1997.