INTERNATIONAL JOURNAL OF ADAPTIVE CONTROL AND SIGNAL PROCESSING
Int. J. Adapt. Control Signal Process. 2001; **15**:445–470 (DOI: 10.1002/acs.684)

Multiple model adaptive control with safe switching

Brian D. O. Anderson,[1] Thomas Brinsmead,[1,*,†] Daniel Liberzon,[2] and A. Stephen Morse[3]

[1] *Department of Systems Engineering, Research School of Information Sciences and Engineering, Australian National University, Canberra, ACT, 0200, Australia*
[2] *Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801, U.S.A.*
[3] *Department of Electrical Engineering, Yale University, New Haven, CT 06520-208267, U.S.A.*

SUMMARY

The purpose of this paper is to marry the two concepts of multiple model adaptive control and safe adaptive control. In its simplest form, multiple model adaptive control involves a supervisor switching among one of a finite number of controllers as more is learnt about the plant, until one of the controllers is finally selected and remains unchanged. Safe adaptive control is concerned with ensuring that when the controller is changed in an adaptive control algorithm, the frozen plant–controller combination is never (closed-loop) unstable. This is a non-trivial task since by definition of an adaptive control problem, the plant is not fully known. The proposed solution method involves a frequency-dependent performance measure and employs the Vinnicombe metric. The resulting safe switching guarantees depend on the extent to which a closed-loop transfer function can be accurately identified. Copyright © 2001 John Wiley & Sons, Ltd.

KEY WORDS: adaptive control; switching algorithms; stability properties; model reference control

1. INTRODUCTION

This paper combines two concepts in adaptive control, namely, multiple model adaptive control (MMAC) and safe adaptive control. We shall describe each of these concepts in turn. For convenience, we shall restrict attention throughout to scalar plants, although the ideas have validity for the multiple-input multiple-output (MIMO) case.

Multiple model adaptive control (see for example, References [1–9]) postulates that the unknown true plant either belongs to a given finite set of nominal plants, or is at least in some way close to one (or more) members of that set. (The set, for example, might comprise a finite set of linear time-invariant plants, and closeness might be reflected by a bounded, additive or multiplicative uncertainty model to account for possible non-linearity.)

* Correspondence to: Thomas Brinsmead, Department of Systems Engineering, Research School of Information Sciences and Engineering, Australian National University, Canberra, ACT, 0200, Australia.
† E-mail: Thomas.Brinsmead@anu.edu.au

Contract/grant sponsor: US Army Research Office, Far East
Contract/grant sponsor: Office of Naval Research; Contract/grant number: N00014-97-1-0946

Each nominal plant corresponds to a controller that is presumed to give satisfactory performance in conjunction with both the nominal plant, and the associated uncertainty ball (if any). What is meant by the term 'satisfactory performance' depends on the application. However, in general it connotes, but at the same time demands more than, closed-loop stability.

As well as the finite set of 'low-level' controllers, the adaptive controller also includes a 'high level' element, a supervisor, which switches between the implementation of controllers from the finite controller set in accord with some rule. The overall objective is to converge to the best controller for the true unknown plant after some finite time. If the true plant coincides with one of the nominal plants, there is an obvious candidate for a good controller, although depending on the performance evaluation criteria and the design method chosen for designing the controller set, this may not be the 'best' performing controller. If, furthermore, the true plant lies in an uncertainty ball around one (or more) nominal plants, the notion of the best controller may be even more ambiguous.

Like most (but not all) adaptive control algorithms, the method suggested here is initially formulated under the assumption that the plant is time-invariant. However, there is an underlying requirement that the adaptive controller has the capability to track time variations in the plant, which are generally relatively slow compared to the input–output dynamics.

The above qualitative description deserves a number of qualifying remarks:

(a) It is likely that the original uncertainty set for the unknown plant is quite large, if not infinite. The determination of the finite set of nominal plants is itself a non-trivial problem. How many should there be, and where should they be located? Such questions are addressed in Reference [1].
(b) One can regard the supervisor's first task as one of plant identification, or more accurately, hypothesis testing. Associated with each nominal plant P_i, there is an hypothesis, H_i, that the true plant lies in the uncertainty ball around P_i. The task of the supervisor is to determine the most likely hypothesis, and switch in the corresponding controller. Issues then arise regarding the effects of noise, errors in the hypothesis testing scheme, the time required to make a decision (with a low probability of error) and so on. The fact that the plant is in a closed loop is a complication. While it may be possible to estimate the closed-loop transfer function fairly accurately, this may not convey much information concerning the (open-loop) plant transfer function itself. The fact that the controller may be changed in the future presents an additional complication, since this will change the experimental conditions. The parameters of the hypothesis testing algorithm will also then need adjustment.
(c) After convergence to a 'best' nominal controller has been attained, it is possible to further fine tune the controller parameters with a view to improving performance [10].

We now indicate something about possible structures of the supervisor, recalling the supervisory control architecture of Reference [5]. Consider Figure 1. Noise signals are omitted for simplicity, as is the element constructing the difference signals $(y - y_i)$ from y and y_i. The unknown plant to be controlled is $\mathbf{P}$, and there are m nominal prescribed plants, $P_1, \ldots, P_m$, each associated with nominal prescribed controllers $C_1, \ldots, C_m$. At time t, the controller is $C_{\sigma(t)}$. The multiestimator is a linear system driven by the unknown plant input u and output y, with m outputs $y_1, \ldots, y_m$. These have the special property that if $\mathbf{P} = P_i$, then $y = y_i$ (after decay of initial condition effects, in the absence of noise, and given that all signals are bounded). Even with these specifications, there is still much freedom in the design of the multiestimator [5].

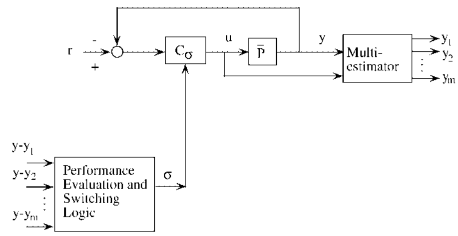

Figure 1. Outline of multimodel adaptive controller, with $\sigma \in \{1, \dots, m\}$.

The task of the performance evaluation and switching logic is to select the P_i closest to **P** and then to implement the corresponding C_i. This is achieved by monitoring the signals $(y - y_i)$, typically looking at an (exponentially weighted) $\mathscr{L}_2$ norm over the semi-infinite interval of past time. More precisely, in Reference [5] the switching logic relied on a particular choice of *monitoring signals* which were defined as follows:

$$\mu_i = \int_0^t e^{-\lambda(t-\tau)}(y_i - y)^2 \, d\tau \tag{1}$$

Then, at any instant of time t, the controller C_σ is implemented with $\sigma(t) \in \{1, 2, \dots, m\}$ taken to be $\arg\min_{i \in I} \mu_i(t)$, and controller $C_{\sigma(t)}$ is used. A dwell-time or hysteresis may be imposed to slow down the switching. One of the goals of this work is to investigate other possibilities for defining the monitoring signals, which would be more compatible with the objective of safe switching.

As argued by Morse *et al.* [5], if **P** coincides exactly with one of the nominal plants P_I, then, even if the switching process produces unbounded signals, the exponentially weighted $\mathscr{L}_2$ norm of $(y - y_I)$ should tend to zero. This is intuitively obvious if the switching process is guaranteed to produce *bounded* signals. Note that such a choice of monitoring signal μ_i, which, in effect gives a measurement of the closeness between the true plant and various possible models, is somewhat arbitrary. However, it does possess the following desirable property. Provided that the reference signal is persistently exciting so that the monitoring signals μ_i remain non-zero for $i \neq I$, C_I will eventually be selected.

The motivation for switching is now clear. Each C_I has been specifically designed to achieve satisfactory performance for the model P_I. When switching finally stops and a particular nominal model is settled upon, then assuming that the finally chosen nominal model is a good representation of the true plant, the actual closed-loop performance will also be close to the designed satisfactory performance.

Note that the switching algorithm does not rule out switching before the monitoring signal has converged. That is, the supervisor does not simply wait until one of the μ_i, say μ_I, has finally converged to zero with the rest non-zero, with the result that C_σ is switched only once, to C_I. This could deny the potential of improving performance much earlier. It would also only be really acceptable in the case that $\mathbf{P} = P_I$, and there is no possibility of the plant drifting with time, a situation that is perhaps unlikely. Of course, if the plant **P** to be controlled is (slowly) time varying, the controller may never settle down. One would hope, however, that the algorithm would insert, for the majority of the time at least, that controller which best controls the current plant **P**.

For such a supervisory structure, the theoretical issues revolve around the design of the multiestimator, the demonstration that all signals remain bounded irrespective of initial conditions, and the demonstration that $\sigma(t)$ converges in finite time. One must also consider the case where **P** does not coincide exactly with any of the nominal plants.

There is an important observation to make about the latter case. Suppose that controller C_J is inserted, and the best controller for the unknown **P** happens to be $C_I, I \neq J$. It may not be necessarily true that the (weighted $\mathscr{L}_2$) norm of $(y - y_I)$ will be the smallest among the norms of $(y - y_i)$ for $i = 1, \dots, m$. This is because the presence of C_J can change the experimental conditions, in the sense that it can make P_I look very far from **P**, whereas with C_I in the loop, P_I may look close to **P**. The concepts of plant identification and accuracy of plant approximation only make sense with respect to the particular controller attached to the plant, a point emphasized strongly in References [11, 12]. Another way of looking at this observation is as follows. The performance evaluation step really evaluates that plant in the set $\{P_1, P_2, \dots, P_m\}$ which would perform the best with the existing controller C_J in the set $\{C_1, C_2, \dots, C_m\}$, rather than the controller which would perform the best with the existing plant **P**. The index associated with the best plant and the best controller need not be the same.

This issue indicates the importance of a careful choice of metric with which to measure the 'closeness' of P_I and **P**. In this paper we make extensive use of the properties of the δ_v metric introduced by Vinnicombe [13]. This metric has the property (also shared by the gap metric [14]) that it induces the weakest topology in operator space such that feedback stability is robust property. The δ_v metric possesses an advantage over the gap metric in that it is less conservative in a well-defined sense (see Chapter 4 of Reference [15]). Robust stability and performance guarantees using the δ_v metric rely on a small-gain argument and impose conditions on particular winding numbers. It has recently been shown that these conditions are equivalent to various unstable pole count conditions which are necessary for more traditional small-gain robustness analysis [16]. We will exploit various robust stability and performance properties that may be guaranteed on consideration of the δ_v metric, in the following development.

We now make some remarks on safe adaptive control. In many adaptive control algorithms, although the plant is initially unknown, an explicit or implicit identified model of the plant is used to design a controller which is then connected in the course of executing the adaptive algorithm. This means both that (a) the true plant normally differs from the model used for controller design purposes and (b) the controller undergoes changes. Such controller change is potentially dangerous, since even if the original closed loop appears to be stable (that is, no signals are expanding in an alarming way), because the plant is not fully known, the new (frozen) closed loop may be unstable. Of course if this were to happen, the adaptive algorithm should be able to discover the inappropriateness of the controller and change it further after improved plant identification. This is the mechanism by which many adaptive control algorithms are able to guarantee that all signals remain bounded. This is however, a beguiling conclusion, if at times the 'frozen' controller–plant combination is unstable. It is a conclusion that, for example, allows 1 MA of current in a 1 kW motor. There even exist algorithms for which, although all signals are bounded, the finite bound itself may be arbitrarily large!

Safe adaptive control refers to adaptive control algorithms in which one guarantees *a priori* that any controller introduced will always yield a stable frozen closed loop when combined with the only partially known plant, that is, if the controller introduced at any time were to remain unchanged from that time on, the resulting time-invariant system would be stable [12]. In this paper, we explore how to achieve the safe adaptive control property for a multiple model adaptive control algorithm.

In safe adaptive control, changes of controller will generally be small (slow) changes. (We deal with the quantification of the size of controller change later in this paper.) It may be the case that, given a particular set of nominal plants P_i and controllers C_i, few or none of the theoretically possible changes from C_i to $C_j \neq C_i$ are small. Such a situation would indicate the desirability of expanding the set of nominal plants, and thus the controller set, which may or may not be practical.

The outline of this paper is as follows. In Section 2, we review the multiple model adaptive control framework, including the functioning of an estimator supervisor, introducing a variant to previous work [5]. One technique for securing safe switching is described in Section 3, and an alternative technique is given in Section 4. This is followed by an example in Section 5. Further issues are discussed in Section 6, with concluding remarks in Section 7. The appendix recalls some notions of robust stability and performance due in the main to Vinnicombe [13, 15].

2. PROBLEM FRAMEWORK AND ESTIMATION SUPERVISOR PROPERTIES

2.1. Assumptions

In this section, we derive some properties of the arrangement of Figure 1, investigating the situation where the switching supervisor is disconnected and the controller remains fixed. In order to enable better understanding of Figure 1 we make some assumptions, not all of which will be carried over to the safe adaptive control algorithm. These assumptions are as follows:

- (A1) The reference signal r is stationary with a wide band spectrum.
- (A2) No noise is present (this assumption will be relaxed later).
- (A3) The true plant $\mathbf{P}$ is linear and time-invariant.
- (A4) The controller C_σ is linear and time-invariant (and is not switched): the lack of switching assumption will be removed later.
- (A5) There are m nominal plants P_i which we identify with their transfer functions $P_i = n_i/d_i$ with n_i and d_i coprime polynomials. For some stable polynomial D, the part of the multiestimator linking $[y, u]$ to y_i is depicted in transfer function terms in Figure 2. Note that if the transfer function of $\mathbf{P}$ were equal to that of the nominal process model P_i then, in the absence of noises and disturbances, y_i would converge to y asymptotically.

2.2. Interpretation of performance evaluation

We now analyse the performance evaluation block, partially displayed in Figure 1. For notational convenience, we drop the subscript σ. Observe that the transfer function from r to $(y_i - y)$ is

$$W_i = \left(\frac{n_i}{D} - \frac{d_i}{D}\mathbf{P}\right)\frac{C}{1 + \mathbf{P}C} \tag{2}$$

(a) In the case that $(\mathbf{P}, C)$ is stable and r has power spectrum $\Phi_{rr}(\omega)$, the spectrum of $(y_i - y)$ is

$$\Phi_i(\omega) = \left|\frac{n_i}{D} - \frac{di}{D}\mathbf{P}\right|^2 \left|\frac{C}{1 + \mathbf{P}C}\right|^2 \Phi_{rr}(\omega) \tag{3}$$

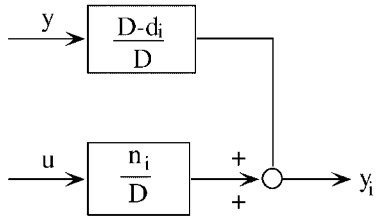

Figure 2. Constituent part of multiestimator.

Further, $\int_0^t (y_i - y)^2 \, d\tau$ will behave like $t\int_0^\infty \Phi_i(\omega)\, d\omega$, so that the error measure of Reference [5], namely μ_i of Equation (1), with λ a suitably small positive constant will behave like a multiple of the (low-pass filtered) integrated spectrum for large t. It is clear then that μ_i is a measure of the error between P_i and $\mathbf{P}$, involving looking at all frequencies, and introducing a frequency weighting which is dependent on Φ_{rr} and C. The measure also involves integrating over ω, rather than alternative functionals such as, for example, taking a supremum. To ensure stability, it may be advantageous to work with pointwise, rather than integral, measures of frequency domain quantities (see below).

(b) In the case that the closed loop is unstable, every signal $(y_i - y)$ will grow at a rate $\exp(at)$ for some $a > 0$ (perhaps with oscillation, and disregarding the special case of polynomial growth). It is then obvious that the value of μ_i is dominated by the effect of transfer function errors between P_i and $\mathbf{P}$ at only a single point (or complex conjugate pair of points) in the (right half) complex plane. In such a case, the μ_i measures are of limited utility as a basis for deciding which of a finite set of transfer functions P_i, $i = 1, \ldots, m$ is closest to $\mathbf{P}$.

(c) In the stable closed-loop case, the μ_i are bounded (or grow in proportion to t if $\lambda = 0$). In the unstable case they grow exponentially fast. This property should be the basis for deciding in the supervisor whether the closed loop is stable. Detecting instability would imply that the currently connected C is unacceptable, but, by the previous argument, is unhelpful for suggesting which replacement is appropriate.

(d) In the framework of Reference [5], μ_i is a *scalar* measure associated with each plant (and thus controller) possibility. In principle, however, investigating the frequency content of $(y_i - y)$ would allow a more sophisticated measure of the distance between each P_i and the true plant to be derived. With the restrictive assumptions that we have made, plus closed-loop stability, we could, in principle, identify the transfer function from r to $(y_i - y)$.

2.3. Relaxation of certain assumptions

We have first argued that under a set of restrictive assumptions, including an additional assumption of closed-loop stability on the $(\mathbf{P}, C)$ loop, we can identify the transfer function from r to $(y_i - y)$. This concept can be carried over with relaxation of the assumptions. Assumption (A1) can be relaxed to require a wide-band property (more comments can be found in Section 6), and (A2) can be relaxed to permit noise. Assumption (A3) can be relaxed to allow time-variation, on a scale much slower than the time-scale for identification, and (A4) can be relaxed to permit switching of the controllers, noting that transient effects that might be associated with the new initial conditions just after switching are allowed to decay. (A state-shared controller parameterization ensuring bumpless transfer has been proposed [7].)

The term 'identification' now means estimating a transfer function where the estimate is understood to have (in general) an error associated with it. The error may be defined by statistics, or by hard bounds, depending on the noise models and the particular identification method used. The term should not be taken to mean error free, or even extremely accurate identification of a transfer function (which, in a noisy situation may nevertheless be possible, given a long enough time interval).

Our interest is rather in spending no longer than necessary on identification; we need to simply pin down the transfer function in question sufficiently accurately that, with the aim of theories dealing with robust stability or performance, we can draw appropriate conclusions regarding stability or performance.

2.4. Performance evaluation using transfer functions

We first make the following key observation.

Lemma 1. Adopt assumptions $(A1)$–$(A6)$. Let $W_i(\mathrm{j}\omega)$ be the transfer function from r to $(y - y_i)$ in Equation (2), for the scheme in Figure 1. Let $\kappa(P_i(\mathrm{j}\omega), \mathbf{P}(\mathrm{j}\omega))$ be the chordal distance (see Appendix A) between P_i and $\mathbf{P}$ at $s = \mathrm{j}\omega$, and let $T(\mathbf{P}(\mathrm{j}\omega), C(\mathrm{j}\omega))$ be the generalized sensitivity matrix of $(\mathbf{P}, C)$, (Appendix A). Then

$$|W_i(\mathrm{j}\omega)| = \kappa(P_i(\mathrm{j}\omega), \mathbf{P}(\mathrm{j}\omega))\,\boldsymbol{\sigma}\left[T(\mathbf{P}(\mathrm{j}\omega), C(\mathrm{j}\omega))\right]\frac{|C|}{\sqrt{1+|C|^2}}\left|\frac{D_i}{D}\right| \tag{4}$$

where D_i is a stable polynomial satisfying $D_i^* D_i = n_i^* n_i + d_i^* d_i$.

Proof. Suppose that $\mathbf{P} = \mathbf{n}/\mathbf{d}$, with $\mathbf{n}$ and $\mathbf{d}$ coprime. Let $\mathbf{D}$ be stable so that $\mathbf{D}^*\mathbf{D} = \mathbf{n}^*\mathbf{n} + \mathbf{d}^*\mathbf{d}$. Then

$$\bar{\mathbf{G}} = \begin{bmatrix} \mathbf{n}/\mathbf{D} \\ \mathbf{d}/\mathbf{D} \end{bmatrix}$$

is a right normalized coprime factorization of $\mathbf{P}$. Let

$$\tilde{G}_i = \begin{bmatrix} -\dfrac{d_i}{D_i} & \dfrac{n_i}{D_i} \end{bmatrix}$$

denote the left normalized coprime factorization of $P_i = -n_i d_i^{-1}$. We shall use the easily verified fact that $|\mathbf{D}/\mathbf{d}| = \sqrt{1+|P|^2}$. Now

$$W_i = \left[\frac{n_i}{D_i} - \frac{d_i}{D_i}\mathbf{P}\right]\frac{D_i}{\mathbf{D}}\cdot\frac{C}{1+\mathbf{P}C}$$

$$= \begin{bmatrix} -\dfrac{d_i}{D_i} & \dfrac{n_i}{D_i} \end{bmatrix} \begin{bmatrix} \dfrac{\mathbf{n}}{D} \\ \dfrac{\mathbf{d}}{\mathbf{D}} \end{bmatrix} \frac{D_i}{D}\cdot\frac{C(\mathbf{D}/\mathbf{d})}{1+\mathbf{P}C}$$

and so at each frequency, by appealing to formulae in the appendix,

$$|W_i| = |\tilde{G}_i G| \quad \frac{|C|\sqrt{1+|\mathbf{P}|^2}}{|1+\mathbf{P}C|}\left|\frac{D_i}{D}\right|,$$

$$= \kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C)]\,\frac{|C|}{\sqrt{1+|C|^2}}\left|\frac{D_i}{D}\right|. \qquad \square$$

The interpretation of Equation (4) is crucial for the case when $(\mathbf{P}, C)$ is stabilizing. Evidently, by inverting the multiplicative scaling effect of $|D_i||C|/(|D|\sqrt{1+|C|^2})$, we observe that the error signal contains the frequency spectrum of the input shaped by $\kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C)]$. This quantity is critical for possibly allowing us to guarantee that C stabilizes P_i, given that it stabilizes $\mathbf{P}$. In a sense, the best P_i, *given the condition that C will be retained* (and a fulfilment of a winding number condition involving P_i and $\mathbf{P}$, see the appendix), is the one that keeps $|W_i||D|\sqrt{1+|C|^2}/(|D||C|)$ small across the whole spectrum.

In the event that $\mathbf{P} \in \{P_i, \cdots, P_m\}$, then of course $W_i = 0$ for some i. Once initial conditions in the multiestimator have decayed, it becomes evident how to model $\mathbf{P}$, irrespective of whether the controller C_i or any other controller completes the loop. On the other hand, if $\mathbf{P} \notin \{P_1, \ldots, P_m\}$, then the best P_i is yet to be identified. Although quantities μ_i provide a measure of an integrated version of $|W_i|^2$ weighted by $\Phi_{rr}(\omega)$, the above analysis indicates that if frequency by frequency (approximate) identification is possible, then a Vinnicombe distance criterion may aid choice of the best P_i. Robust performance, as well as robust stability may be considered in the same framework.

An alternative construction for the multiestimator exists which is more costly in terms of realization complexity, but which modestly simplifies the above calculations. This alternative involves replacing D in Figure 2 by D_i so that, in place of (4), one has the slightly simpler

$$W_i = \kappa(P_i, P)\boldsymbol{\sigma}[T(\mathbf{P}, C)]\frac{|C|}{\sqrt{}}1+|C|^2 \tag{5}$$

This may be achieved at the cost of no longer allowing all the multiestimator output signals to be constructed from a common state (with dimension twice the order of D). It will then be slightly easier to evaluate the m possible $\kappa(P_i, P)\boldsymbol{\sigma}[T(\mathbf{P}, C)]$, with the aim of choosing one which is the smallest (taking into account the whole spectrum). In a rough sense, the smallest (over i) of the functions on the left of (5) may identify the best controller in terms of performance; however, we will focus mainly on the safety issue.

Of course, there may not be a single index i which minimizes $|W_i|$ given by (4) or (5), and in a conventional multiple model adaptive control structure we cannot expect to identify the W_i exactly, but only to within some error bound. Further, they may be changing slowly, due to changing $\mathbf{P}$. Last, even though (as we have noted) quantities like $|W_i|$ measure the approximation error between P_i and $\mathbf{P}$, given a particular controller C, rather than measuring the best C_i to use for $\mathbf{P}$, we need to be able to use the $|W_i|$ for the purposes of selecting the controller to switch in. We must focus on the safety issue in orchestrating this switching.

2.5. Changing the controller

The goal of the estimator-based supervisor is not to find which $P_i \in \{P_1, \ldots, P_m\}$ is closest to the true plant $\mathbf{P}$ with the current controller, but what would make a better controller for $\mathbf{P}$ than the current one. We could select P_i corresponding to the smallest $\kappa(P_i, \mathbf{P})\boldsymbol{\sigma}\{T[(\mathbf{P}, C)]\}$ or an integrated version thereof, and then hope that C_i makes a good, and certainly stabilizing, controller for $\mathbf{P}$. In order to *guarantee* that C_i stabilizes $\mathbf{P}$ (that is, to ensure safe adaptive control, as reviewed in the introduction, see also the appendix) more must be established. One sufficient condition, given the hypothesis that $(\mathbf{P}, C)$ is stable, is that

$$\kappa(C, C_i)\boldsymbol{\sigma}[T(\mathbf{P}, C)] < 1 \quad \forall \omega$$

It is not trivial to verify this inequality on the basis of available data. This is the issue that will be addressed in the next section.

3. SAFE SWITCHING

Suppose that the controller $C = C_j$, the nominal controller associated with the plant P_j, is currently connected to the unknown plant $\mathbf{P}$. Suppose further that we believe, perhaps on the basis of some performance estimation information, that it would be appropriate to use C_I. To be assured of safety, we would like to check that

$$\kappa(C_I, C_j)\boldsymbol{\sigma}\{[T(\mathbf{P}, C_j)]\} < 1 \quad \text{for all } \omega \tag{6}$$

Provided a winding number condition involving C_I, C_j holds in addition to (6), C_I is certainly stabilizing (see Lemma A2 of Appendix A.1).

Since the C_i are all known, $\kappa(C_I, C_j)$ is also known. Thus, in order to check (6) we need only to evaluate $\boldsymbol{\sigma}\{[T(\mathbf{P}, C_j)]\}$. Although the performance estimator is not configured (or indeed, able to be configured) to yield $\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]$ directly, even with the embellishment of operating as a transfer function identifier, the performance robustness results of Appendix A.1 allow us to calculate an overbound for this quantity. The following lemma and its corollaries are variants of the small gain theorem or theorems of optimal robustness in the δ_v metric [13], specialized so that the various terms in the expressions are quantities that we can estimate.

Lemma 2. Let $\mathbf{P}$ be the true plant and let $\{P_1, \ldots, P_m\}$ and $\{C_1, \ldots, C_m\}$ be the collection of nominal plants and controllers. Suppose that C_j stabilizes $\mathbf{P}$ and P_i, and that for all ω

$$\kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C_j)] < \tfrac{1}{2} \tag{7}$$

Then

$$\boldsymbol{\sigma}[T(\mathbf{P}, C_j)] \leqslant \boldsymbol{\sigma}[T(P_i, C_j)]\{1 + \kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]\}. \tag{8}$$

Proof. Let us first note that Equation (7) implies

$$\kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(P_i, C_j)] < 1 \tag{9}$$

From Equation (23) in Lemma 6 and the fact that C_j stabilizes both P_i *and* $\mathbf{P}$, we have that (for all ω):

$$\boldsymbol{\sigma}[T(P_i, C_j)] \leqslant \frac{\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]}{1 - \kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]}$$

and so

$$\kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(P_i, C_j)] \leqslant \frac{\kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]}{1 - \kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]}$$

The above inequality and Equation (7) immediately yield Equation (9). Now using Equation (9) and again the fact that C_j stabilizes P_i *and* $\mathbf{P}$, there holds (for all ω)

$$\boldsymbol{\sigma}[T(\mathbf{P}, C_j)] \leqslant \frac{\boldsymbol{\sigma}[T(P_i, C_j)]}{1 - \kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(P_i, C_j)]} \tag{10}$$

Multiplication of the above by $1 - \kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(P_i, C_j)]$ and rearrangement yields (8). □

The lemma may be used in the following way. With C_j connected to $\mathbf{P}$, and giving a stable closed loop, using the multiestimator signals, one (approximately) identifies $\kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]$ as a function of frequency for each P_i which is stabilized by C_j (see Section 2.4). (Note that the information regarding whether (P_i, C_j) is stable is available *a priori*.) We then use Equation (8) to overbound $\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]$ giving the following corollary of Lemma 2.

Corollary 1. At each frequency ω,

$$\boldsymbol{\sigma}[T(\mathbf{P}, C_j)] \leqslant \min_{i \in \Im_j} \bar{\sigma}[T(P_i, C_j)]\{1 + \kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]\} \tag{11}$$

where $\Im_j \subset \{I, \dots, m\}$ satisfies $i \in \Im_j$ if P_i is stabilised by C_j and Equation (7) holds.

Denote by $\mathbf{T}_j$ the right-hand side of Equation (11). Of course, $\mathbf{T}_j$ is something that we can estimate, albeit with some error. Furthermore, by Lemma 5 a sufficient condition that C_I will stabilize $\mathbf{P}$ is that both $\text{wno}(1 + C_I^* C_j) + \eta(C_I) - \tilde{\eta}(C_j) = 0$ and that Equation (6) holds. In the light of Equation (10) this leads to the following corollary for a more conservative sufficient condition for stability, but one involving quantities that we can estimate.

Corollary 2. Given that the closed loop $[\mathbf{P}, C_j]$ is stable, then a sufficient condition for $[\mathbf{P}, C_I]$ to be stable is that $\text{wno}(1 + C_I^* C_j) + \eta(C_I) - \tilde{\eta}(C_j) = 0$ and

$$\kappa(C_I, C_j)\mathbf{T}_j < 1 \tag{12}$$

Notice that if the set $\{P_1, P_2, \dots, P_m\}$ is dense enough in the set of possible unknown plants, then Equation (7) will be straightforwardly satisfied for some i. Likewise, if the set $\{P_1, P_2, \dots, P_m\}$ is a fairly dense one, then with fixed j, the quantity $\kappa(C_i, C_j)$ will be small for some i.

Notice also that if for some particular $i = I$, there holds

$$\kappa(P_I, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C_j)] < \kappa(P_i, \mathbf{P})\boldsymbol{\sigma}[T(\mathbf{P}, C_j)]$$

for all $i \neq I$, then it would be logical to hypothesise that $\mathbf{P}$ is best modelled by P_I, when C_j is attached. It would then be natural to check both whether $\kappa(C_I, C_j,)\mathbf{T}_j < 1$, together with the winding number condition for C_I and C_j, in order to determine whether C_I can be safely implemented.

Finally, we remark that we cannot, in practice, expect measurements on the closed-loop system to yield exact values of quantities such as $|W_i|$ at each frequency. It is also well known [17] that, in the presence of noise in a stochastic framework, any method of identifying the frequency-domain quantity W_i is subject to variance which increases with the number of parameters to be identified and decreases with the time available for identification. Reduced identification variance and hence faster identification times can be achieved by reduction of the number of parameters in the model set, but, in general, only at the expense of increased identification bias.

In particular, for a least-squares transfer function identification, the mean-square identification error J_{MSE} is composed of a sum of a bias term J_{B} and a variance term J_{V} [17]. The variance term is $J_{\mathrm{V}}(\omega) \propto (n/N)\,\Phi_N(\omega)/\Phi_R(\omega)$ (Equation (12.35) in Reference [17]), asymptotically proportional to the number of parameters n, inversely proportional to the number of sampled data points (identification time interval) N and inversely proportional to the signal to noise ratio $\Phi_R(\omega)/\Phi_N(\omega)$. The bias term, on the other hand, generally decreases with n and is asymptotically independent of N.

Were exact values of W_i available, however, we could, in principle, identify $\mathbf{P}$ and test each C_i with $\mathbf{P}$. However, exact values are not necessary, since robust stability is ensured merely by the satisfaction of particular inequalities. Hence, in the presence of norm bounded noise, it is in principle possible to give hard error bounds on the identification error and hence give a hard guarantee of safe switching.

In the example of Section 5, we achieve perfectly satisfactory operation of the part of the algorithm that ensures safe switching. This is achieved even though the identification of $\mathbf{T}_j$ is subject to error due to noise, and with the checking of (12) at only a finite number of discrete frequency points.

4. ALTERNATIVE METHOD TO SAFETY

Alternative methods to guarantee safe switching, based on different *a priori* assumptions, exist. As explained in Section 2, and using the notation of that section, we can conceive of an identification (albeit with error) of

$$W_i = \left(\frac{n_i}{D} - \frac{d_i}{D}\mathbf{P}\right)\frac{C_j}{1 + \mathbf{P}C_j}$$

Depending on the location relative to the stability boundary of the zeros of D and d_i, an identification algorithm may allow one to determine with reasonable accuracy the frequency response of

$$\mathbf{V}_i :\stackrel{\text{def}}{=} (P_i - \mathbf{P}) \frac{C_j}{1 + \mathbf{P}C_j}$$

We will make use of the following inequality:

$$\left|\frac{\mathbf{P}C_j}{1 + \mathbf{P}C_j}\right| \leqslant \left|\frac{P_i C_j}{1 + P_i C_j}\right| + \left|\frac{(P_i - \mathbf{P})C_j}{1 + \mathbf{P}C_j}\right| \left(1 + \left|\frac{P_i C_j}{1 + P_i C_j}\right|\right) \tag{13}$$

This inequality is easily derived from the observation that

$$\frac{\mathbf{P}C_j}{1 + \mathbf{P}C_j} = \frac{P_i C_j}{1 + \mathbf{P}C_j} + \frac{(\mathbf{P} - P_i)C_j}{1 + \mathbf{P}C_j}$$

and

$$\frac{P_i C_j}{1 + \mathbf{P}C_j} = \frac{P_i C_j}{1 + P_i C_j} \left[1 + \frac{(P_i - \mathbf{P})C_j}{1 + \mathbf{P}C_j}\right]$$

Now we can state the following lemma (also a variant of the small gain theorem).

Lemma 3. Suppose that $\mathbf{P}$ and P_i are each stabilized by C_j. Then $\mathbf{P}$ is stabilized by

$$C_k = C_j(I + \Delta)$$

if the transfer functions C_k and C_j have the same number of right half-plane poles and

$$|\Delta| \left|\frac{\mathbf{P}C_j}{1 + \mathbf{P}C_j}\right| < 1 \quad \forall \omega \tag{14}$$

A sufficient condition for Equation (14) is that (for all ω)

$$|\Delta| \left[\left|\frac{P_i C_j}{1 + P_i C_j}\right| + \left|\frac{(P_i - \mathbf{P})C_j}{1 + \mathbf{P}C_j}\right| \left(1 + \left|\frac{P_i C_j}{1 + P_i C_j}\right|\right)\right] < 1 \tag{15}$$

Proof. The first part of the lemma is standard [15, 18]. In addition, if Equation (15) holds, then Equation (14) it can be seen to be also true by (13). □

We can take advantage of the above lemma in the following way. Using some identification algorithm, we estimate the frequency responses $\mathbf{V}_i$, which appear in Equation (15). All other quantities are known. Therefore, for each k with the property that C_k and C_j have the same number of right half-plane poles, we can check condition (15), albeit with estimated quantities replacing true quantities. If (15) does hold, then it is safe to insert C_k. Once again, one can consider all the P_i for which Equation (15) holds and take the right-hand side (at each frequency) to be the minimum (at each frequency) over i subject to (P_i, C_j) being stable. In this way, a less conservative sufficiency condition is obtained.

Whether this scheme or that of Section 3 will give better results will depend on the particular problem. It may be that only one of the two schemes can be used, depending on the pole distribution of the controllers.

5. EXAMPLE

In order to demonstrate the method, the controller switching scheme with safety, as described in Section 3, was implemented in Matlab. The original specifications are given here in continuous time, although the simulation was implemented in discrete time using the zero-order hold equivalent representation of continuous transfer functions with a sampling interval of $\Delta = 0.05$ s. The simulation period was 30 s.

The plant to be controlled was chosen as

$$\mathbf{P} = \frac{1.2\,(\frac{1}{2}s+1)(-\frac{1}{4}s+1)}{(\frac{2}{3}s+1)(\frac{1}{3}s+1)(\frac{1}{10}s+1)}$$

with a DC gain of $\mathbf{K} = 1.2$ and a non-minimum-phase zero at $\mathbf{z} = 4$. The control objective is to extend the bandwidth of the open-loop plant, that is for the complementary sensitivity transfer function to have a bandwidth that exceeds that of the open-loop plant, with regard to practical limitations imposed by the existence of a non-minimum-phase zero. There were 441 plant models used for the multiple model set. They were specified as

$$P_i = \frac{K_i(\frac{2}{3}s+1)(-(1/z_i)s+1)}{(\frac{1}{2}s+1)\ \ (\frac{1}{3}s+1)(\frac{1}{3}s+1)}$$

where the modeled DC gain varied from $K_i = 0.2$ to 2 in 20 logarithmically equally spaced intervals, and the modelled non-minimum-phase zero varied from $z_i = 1$ to 10, also in 20 logarithmically equally spaced intervals.

The controllers for each model were designed using discrete-time Q-synthesis (where, for stable nominal models G_o, the controller is calculated as $C = Q(1 - QG_0)^{-1}$, where the target complementary sensitivity function T_d defines Q by $T_d = QG_0$ [19]). The target closed-loop (complementary sensitivity) transfer function had a DC gain of $K_t = 0.95$ (this slight detuning avoided issues involved with calculating the quantity $\kappa(C_i, C_j)$ between marginally stable controllers, that is, controllers with integrators). The target complementary sensitivity had a target bandwidth of $p_t = \eta z_i$ where $\eta = \frac{3}{4}$ is a tuning factor indicating the aggressiveness of the design, which is limited by the (modelled) non-minimum-phase zero z_i. The target closed-loop denominator polynomials were specified as $(s+p_t)^{n_1}(s^2 + 2\zeta_t p_t s + p_t^2)^{n_2}$, with a damping factor $\zeta_t = 0.5$ and with $n_1 \in \{0, 1\}$ and n_2 chosen to give the appropriate polynomial order $n_1 + 2n_2$. Naturally, the target complementary sensitivity retained the modelled non-minimum-phase zero z_i, so that the final (continuous-time) target complementary sensitivity function becomes

$$T_{\mathrm{d}} = K_{\mathrm{t}} \frac{-s + z_i}{s + z_i} \frac{p_{\mathrm{t}}^{n_1 + 2n_2}}{(s + p_{\mathrm{t}})^{n_1} (s^2 + 2\zeta_{\mathrm{t}} p_{\mathrm{t}} s + p_{\mathrm{t}}^2)^{n_2}} \tag{16}$$

The controllers were implemented in discrete-time state space, in controller canonical form [19], with the state variable *not* reset between controller switchings.

In order to maintain persistent excitation, the reference signal was specified as a filtered version of a white noise signal with variance $\sigma_r^2 = 10^2$ that was created with the Matlab **randn** function. A triangular filter with unit weighting and window length of $\tau_{\mathrm{f}} = 2$ s was used to smooth the high-frequency components of the original white signal. White output noise n of variance $\sigma_n^2 = 0.5^2$ and white input noise d_u of variance $\sigma_u^2 = 1^2$ were also included in the simulation. The observed output y_m is thus

$$y_m = \mathbf{P}(u + d_u) + n$$

where u is the demanded control.

For the estimation of the transfer functions W_i (Equation (4)), a standard recursive least-squares [17] algorithm was employed in order to directly identify a parameter vector θ, consisting of the coefficients a_{ji}, b_{ji} of the discrete time representation of the fourth-order transfer function

$$W_i \approx \frac{b_{1i} z^{-1} + b_{2i} z^{-2} + b_{3i} z^{-3} + b_{4i} z^{-4}}{1 + a_{1i} z^{-1} + a_{2i} z^{-2} + a_{3i} z^{-3} + a_{4i} z^{-4}} \tag{17}$$

The regressor vectors ϕ were composed of low-pass filtered versions of past observed measurements y_m and demanded control u. The low-pass filter used to create the regressors had continuous poles at $\{-5, -6, -8, -8\}$, and no zeroes. The system identification covariance matrix was initialized at

$$P_0 = \begin{bmatrix} 10^2 I_4 & 0 \\ 0 & 10^5 I_4 \end{bmatrix}.$$

In addition, the recursive least-squares scheme that was implemented included the feature that the covariance matrix P was scaled by a reset factor of $10\times$, each time a controller switching took place. All other initial conditions were taken as zero, including the initial plant and controller states, multiestimator states, initial regressor vectors and estimates of parameters of the transfer functions W_i. No attempt was made to estimate either the parameter covariance or the typical frequency response uncertainty for the estimates of the transfer functions W_i using this identification method.

The multiestimator (Equation (2)) was designed with denominator D having (continuous) poles $p_{\mathrm{me}} = \{-2, -3, -4\}$. The multiestimator performance signals μ_i (Equation (1)), were designed with λ corresponding to a time constant of $\tau_\mu = 10$ s. The simulation was run several times with the supervisor initialized with the controller corresponding to one of two nominal models: one with $K_i = 0.2 \times 10^{0.80} \approx 1.26$ and $z_i = 10^{0.70} \approx 5.01$ and the other with $K_i = 0.2 \times 10^{0.75} \approx 1.12$ and $z_i = 10^{0.05} \approx 1.12$. We denote the first of these controllers K^{α}_{init} and the second K^{β}_{init}. Both of these initial controllers were stabilizing for the true system **P**.

The minimum time between controller switchings was set at $\tau_S = 1$. At any time after the minimum controller switching time, the controller suggested by the supervisor corresponded to the minimum of the μ_i. In order to check safety, condition (12) was investigated for a finite number of frequency points, with of course the *estimated* transfer functions $\mathbf{T}_j$ used in place of the actual values. This set of frequencies consisted of $n_\omega = 100$ data points spaced equally from 0 to the (normalized) Nyquist frequency of π (that is an unnormalized frequency of π/Δ). In order to reduce computational load, if after a time of τ_S, an alternative to the currently implemented controller that subsequently failed the safety checking routine was suggested, another interval of τ_S would be required before allowing the suggestion of an another controller. For comparison, simulations of the algorithm without the safety checking property were also conducted.

A summary of the (somewhat arbitrary) choices of parameters and design methods used for the example appears in Table I.

5.1. Results

The process was simulated and data recorded a total of 16 times with the particular choice of simulation parameters as described above. The results varied for alternative realization of the noise and reference signals. With the safety property enforced, as was to be expected the allowed controller switchings were much less frequent. For simulations both with and without the safety checking property, it was noticed that starting with K^{α}_{init} typically resulted in better performance and less frequent switching (as well as less chance of instability for the unsafe adaptive case) than using the initial controller K^{β}_{init}. A typical controller switching trajectory is shown in Figure 3 with an initial controller of K^{α}_{init}. For the purposes of interpreting the controller switching graph, the dotted line corresponds to the modelled DC gain, and the dashed line corresponds to the modeled non-minimum phase zero. The corresponding trajectory showing the output lagging the reference also appears in Figure 3.

In all the simulation runs with the safety checking property, the routine did not once allow the switching in of any controller which was destabilizing, although it did allow controllers with quite poor performance. See Figure 4 for such a simulation run, which was started with the initial controller K^{β}_{init}. (Note that both the time and output axes differ between the right and left figures.)

For simulations run without the safety checking property, conversely, as expected, the controller switchings allowed by the supervisor were much more frequent. A typical controller switching trajectory and output trajectory are shown in Figure 5. For comparison with Figure 3, these correspond to an initial controller K^{α}_{init}.

Most of the simulations that were run without safety checking exhibited some periods of poor performance. See, for example, Figure 6, which shows the trajectory for an initial controller K^{β}_{init}. In addition to resulting in poor performance, on a number of occasions, the controller corresponding to the minimum μ_i resulted in a destabilized closed loop (this was checked numerically, using the model of the true plant **P**), although the situation did not persist for more than a few periods τ_S. See Figure 7 for one such a simulation (initial controller K^{β}_{init}) where an unstable controller was introduced between 2 and 4 s (note the scale on the *y*-axis). However, even without the safety checking property, an unstable closed loop did not occur on every simulation run.

Table I.

Description	Symbol	Value
Sampling interval	Δ	0.05 s
Simulation Period	—	30 s
Continuous to discrete method	—	Zero-order hold equivalent
Real plant	**P**	$\dfrac{1.2(\frac{1}{2}s+1)(-\frac{1}{4}s+1)}{(\frac{2}{3}s+1)(\frac{1}{3}s+1)(\frac{1}{10}s+1)}$
Plant models	P_i	$\dfrac{K_i(\frac{2}{3}s+1)(-\frac{1}{z_i}s+1)}{(\frac{1}{2}s+1)(\frac{1}{3}s+1)(\frac{1}{3}s+1)}$
Model range DC gain	K_i	[0.2, 2]
Number of DC gain intervals	—	20 Logarithmically spaced
Model range NMP zero	z_i	[1, 10]
Number of NMPZ intervals	—	20 Logarithmically spaced
Number of models	—	$441 = 21 \times 21$
Controller design method	—	Q-synthesis
Closed-loop target DC gain	K_t	0.95
Closed-loop target bandwidth	p_t	$-\eta z_i$
Closed-loop target damping	ζ_t	0.5
CLTBW: OLNMPZ ratio	η	$\frac{3}{4}$
Controller implementation	—	Controller canonical form
Reference signal	r	Triangular filtered white noise
Reference covariance	σ_r^2	10^2
Ref. filter window time length	τ_f	2 s
Output noise covariance	σ_n^2	0.5^2
Input disturbance covariance	σ_u^2	1^2
Multiestimator poles	p_{me}	$\{-2, -3, -4\}$
Multiestimator time constant	τ_μ	10 s
Transfer function estimation scheme	—	Recursive least squares
Estimated parameters	$\hat{\theta}$	Discrete time TF coefficients
Regressor filter poles	—	$\{-5, -6, -8, -8\}$
Model order	—	4
Initial covariance matrix	P_0	diag $\{10^2 I, 10^5 I\}$
Covariance reset factor	—	10
Initial conditions	$x_{\mathbf{P}}(0)$, $x_K(0)$, $\hat{\theta}(0)$, $\phi(0)$, $\mu_i(0)$	zero I.C.
Controller minimum switching time	τ_S	1 s
Initial controller index	$\sigma(0)$ for K^{α}_{init}	corresponding to $K_i \approx 1.26$, $z_i \approx 5.01$
	for K^{β}_{init}	corresponding to $K_i \approx 1.12$, $z_i \approx 1.12$
Test frequency range		$[0, \pi]$
Number of test frequency intervals	n_ω	100, Linearly spaced

5.2. Discussion

We now investigate why, for this example, choosing the minimum over μ_i as a basis on which to select a controller, risks leading to instability. Keeping in mind that μ_i is approximately the (low-pass filtered) spectrum of Φ_i in Equation (3), we see from expression (2) for W_i that there is a weighting bias in favour of low frequencies introduced by the term $|C|/\sqrt{1+|C|^2}$ due to

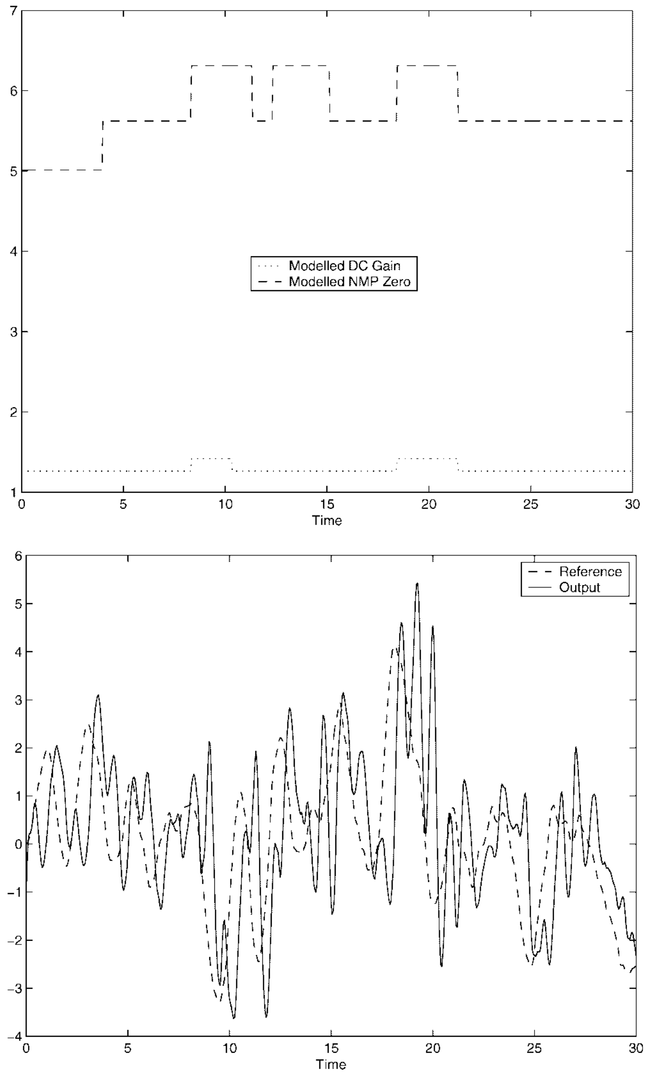

Figure 3. (Safe) Controller switching.

high-gain control near DC. This is exacerbated by the predominantly low-pass reference spectrum Φ_{rr}.

Note also that the pole at -10 in the true plant is modelled in all the P_i at -3. Thus, the true plant does not roll-off until much later in frequency than the models suggest. In addition, the true plant pole at -1.5 and zero at -2 is modelled by a pole at -2 and zero at -1.5. Therefore, if the gain of the plant and model match at low frequencies (as is tended to be promoted by the weighting bias in μ_i toward low frequencies) then the magnitude gain of the model at mid-range frequencies is higher than that of the true plant. This fact, in addition to the high-frequency pole

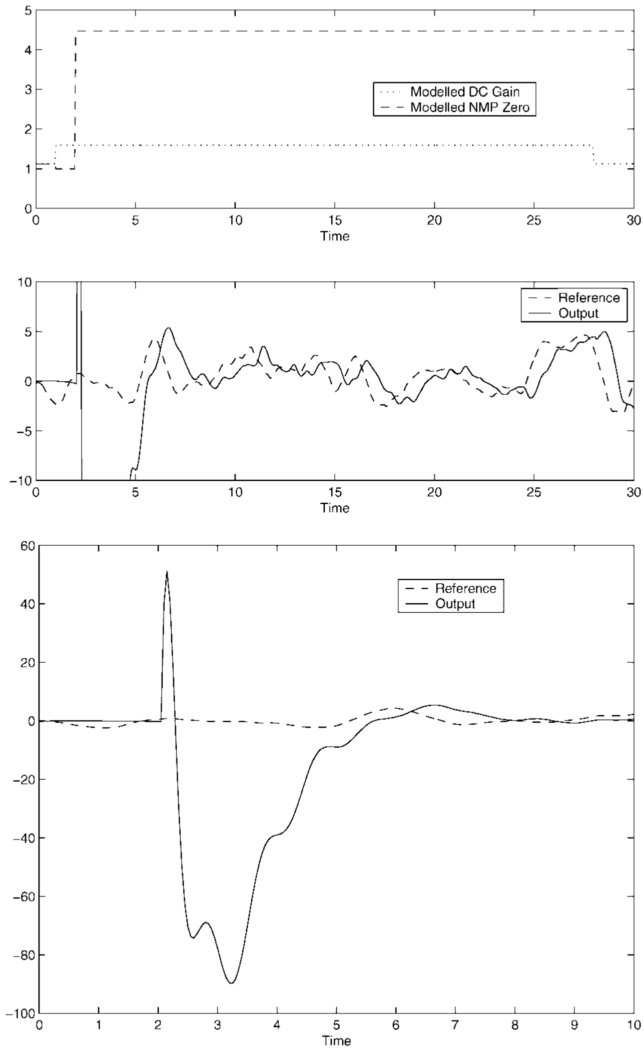

Figure 4. Example of poor performance: safe switching.

being modelled at a lower than actual frequency, promotes a tendency to estimate the non-minimum phase zero also at a higher frequency than actual. It can be seen that the controllers that were actually implemented in the simulations correspond to models with DC gains close to the correct value of 1.2 and non-minimum-phase zeroes greater than the correct value of $\mathbf{z} = 4$. The mismodelled (overmodelled) non-minimum-phase zero leads to attempts to implement a more aggressive (high bandwidth) closed loop than is justified, which, in combination with the mismodelled (undermodelled) pole results in a high risk of instability.

In contrast, a heuristic argument would suggest that for the particular purposes of stability (but also for performance) a more appropriate model choice would be one which is close to the plant

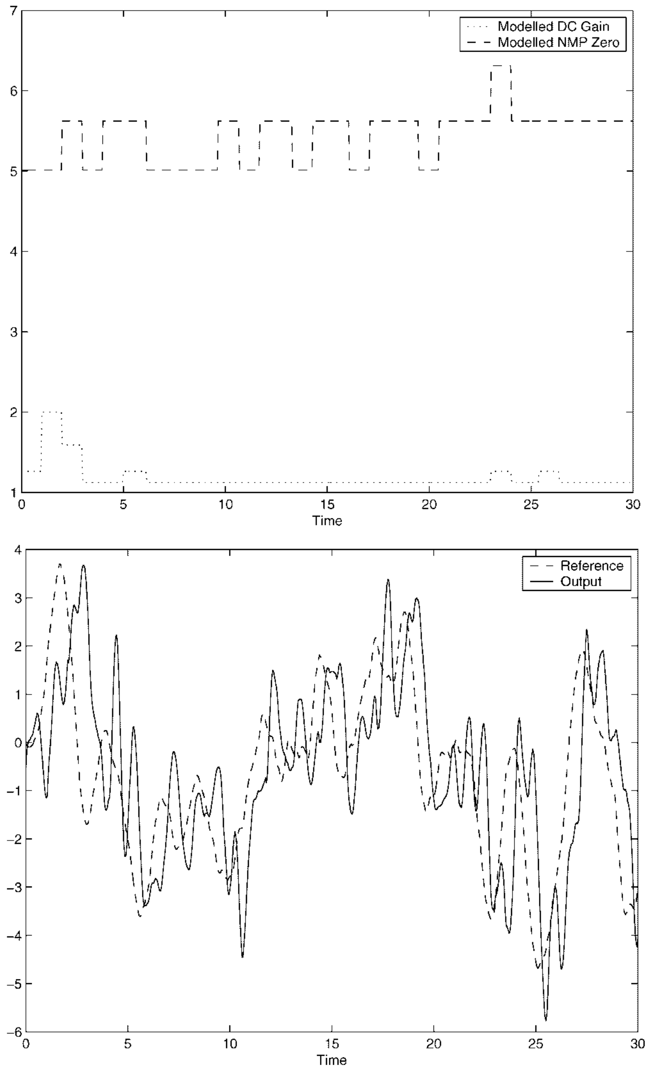

Figure 5. Controller switching (without safety).

near the closed-loop crossover frequency (see Figure 8). Because of the mismodelling of the low-frequency stable poles and zeroes, such an argument would suggest that underestimating the DC gain in order to more closely represent the mid-frequency ranges would be appropriate. Furthermore, it is safer to underestimate rather than overestimate the location of a non-minimum-phase zero, although the switching algorithm allowed such a situation to occur, even when embellished with safety checking.

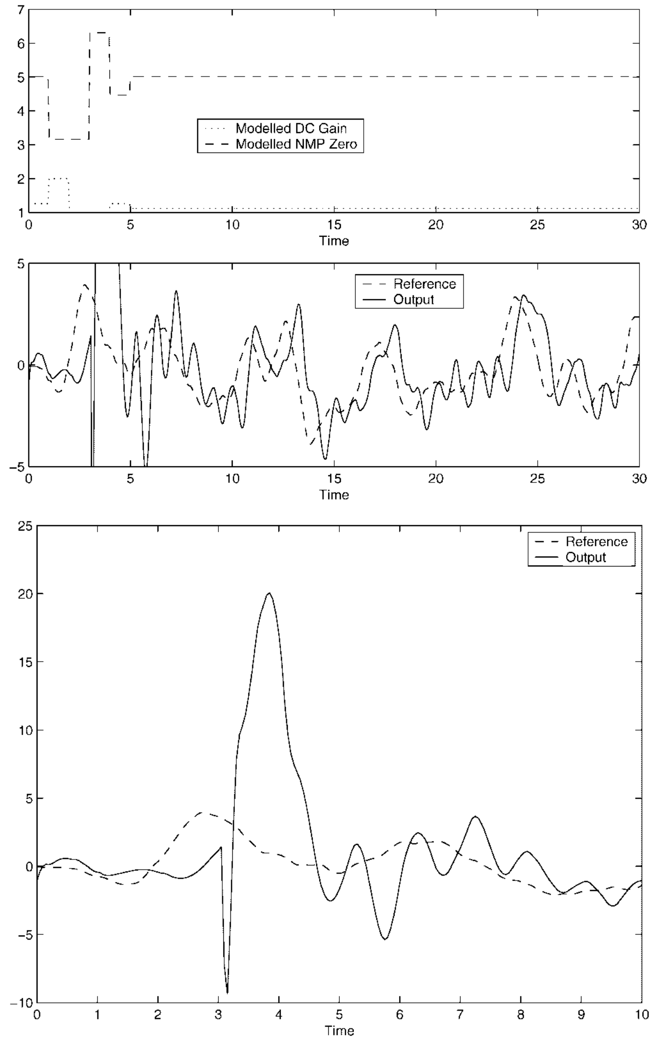

Figure 6. Example of poor performance (without safety).

6. FURTHER ISSUES

Requirements on r: In this paper we have assumed that r is a wideband signal. In the absence of more information about $\mathbf{P}$, this is a necessary requirement. Safe adaptive control is concerned with ensuring stability, and, for linear time-invariant plants, this is assessed by a condition which has to be checked at each frequency. In our case, an experimental form of checking is proposed, implying that excitation over the whole range of frequencies is required.

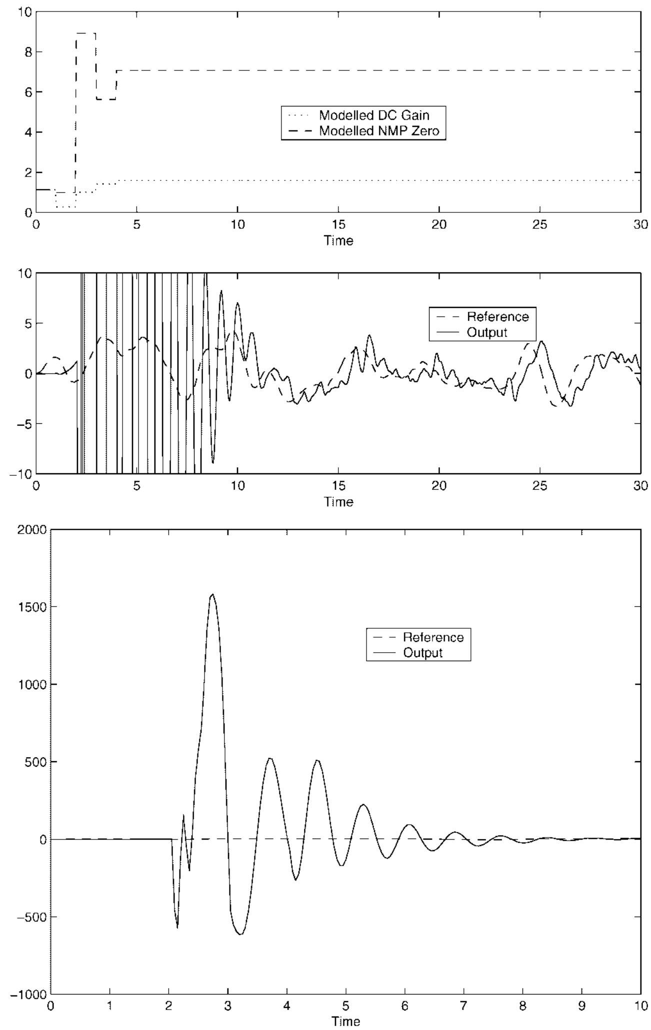

Figure 7. Example of temporary instability (without safety).

There are at least three exceptions to this requirement. If $\mathbf{P}$ is known to be a rational transfer function with degree no more than some integer k, then the knowledge of $\mathbf{P}$ at k complex frequencies is sufficient to determine $\mathbf{P}$ at all frequencies. This, however, is an unlikely situation. An alternative exception occurs if $\mathbf{P}$ is unknown up to some frequency Ω_l say, but for $|\omega| > \Omega_l$, one knows that $|\mathbf{P}| \leqslant (\omega^2 + \Omega_l^2)^{-1/2}$. Finally, a third exception is effectively a combination of the first two. If we know that $\mathbf{P} = \mathbf{P}_0(I + \Delta)$ or $\mathbf{P} = \mathbf{P}_0 + \Delta$ for some unknown rational $\mathbf{P}_0$ of known degree, and for some Δ lying within certain frequency-dependent bounds, then, once again, the requirement that r be a wideband signal may be relaxed.

The key in each case is to have r such that a combination of experimental data and *a priori* information is sufficient to assure the satisfaction of the stability conditions.

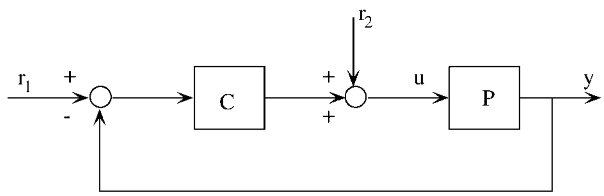

Figure 8. Closed loop used for defining the generalized sensitivity matrix.

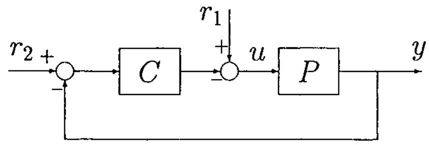

Figure 9. Closed-loop system (P, C).

Effect of noise: Although identification can never be exact when noise is present, exact identification is not necessarily required. We merely need assurance that certain inequalities involving transfer functions are fulfilled. The presence of noise will not invalidate the algorithm, but it will set limits on the switching speed, since noise will increase the time required to identify to any given degree of confidence.

Speed of switching: For the proposed estimation-supervision component of the algorithm, the switching speed will be limited for two reasons: with a given controller in place, a steady-state identification must be achieved, implying that transients must settle down. Second, at higher noise levels, longer periods may be needed to secure sufficiently accurate identification.

7. CONCLUSION

This paper has presented some possible methods for ensuring that controller adaptation in a multiple model adaptive control context proceeds cautiously and safely, in the spirit of earlier work [12]. We use Vinnicombe metric results to ensure that any change in controller is small enough so as not to result in a (frozen) unstable closed loop. A sufficient condition for guaranteeing safety can be checked, in principle, by the calculation of quantities based on available measurements. Note, however, that this requires the identification of an *estimate* of a particular transfer function, so that the safety guarantee is not absolute. However, given confidence estimates on the identification method, it would be possible to give a quantitative probabilistic guarantee. This would be a function of both the magnitude of noise and any bias introduced by choice of identification method. In principle, it would be possible to give absolute guarantees on switching safety if hard error bounds on the transfer function identification step were known.

In addition, practical considerations suggest that resorting to checking a finite number of frequencies is more appropriate, although the stability guarantee, in principle, requires checking a condition over an infinite number of frequencies. This issue could be tackled rigorously if an *a priori* bound on *complexity* [15] (roughly, the rate of change of the transfer function with frequency) of the true plant were known.

A demonstration of the safe switching multiple model algorithm was presented. This showed that the safe switching algorithm resulted in a conservative switching regime which indeed maintained (frozen) closed-loop stability, whereas the same supervisor without the safety check occasionally implemented a destabilizing controller.

The suggested method requires that the initial controller chosen is a stabilizing controller. Further research is required to determine methods to quickly find a stabilizing controller in the event that a destabilizing initial controller is implemented or in the unlikely event that a destabilizing controller is switched in during operation. Methods to detect an unstable closed loop are also required. Further research is also required on determining an appropriate switching time. This would be based on both the time interval required for securing sufficient confidence in identification, and the time for switching transients to have decayed appropriately.

APPENDIX A: ROBUST STABILITY AND PERFORMANCE RESULTS FOR SCALAR PLANT

A.1. Robust stability results

We summarize some ideas from Vinnicombe [13], for scalar plants. With $P_1(s)$ and $P_2(s)$ two plants with rational transfer functions, and with $P^*(s) = P(-s)$, define the chordal distance κ at $s = \mathrm{j}\omega$ by

$$\kappa[P_1(\mathrm{j}\omega), P_2(\mathrm{j}\omega)] \stackrel{\mathrm{def}}{=} \frac{|P_1(\mathrm{j}\omega) - P_2(\mathrm{j}\omega)|}{\sqrt{1 + |P_1(\mathrm{j}\omega)|^2}\sqrt{1 + |P_2(\mathrm{j}\omega)|^2}} \tag{A1}$$

For a rational function $X(\mathrm{j}\omega)$, let wno $[X]$ denote the number of encirclements of the origin made by $X(s)$ as s follows the standard Nyquist D-contour in a counterclockwise direction, indented into $\mathrm{Re}[s] > 0$ around any $\mathrm{j}\omega$-axis poles and closing on a semicircle in the right half plane. All open right half-plane poles and zeros of $X(s)$ must be enclosed by the contour. Associated with Equation (A1) is a winding number condition, namely

$$\mathrm{wno}(1 + P_2^* P_1) + \eta(P_1) - \tilde{\eta}(P_2) = 0 \tag{A2}$$

where $\eta(X)$ and $\tilde{\eta}(X)$ denote respectively, the number of poles of X in the open and closed right half-plane. If Equation (A2) holds, then

$$\delta_v(P_1, P_2) = \sup_\omega \kappa[P_1(\mathrm{j}\omega), P_2(\mathrm{j}\omega)]$$

is the v-gap matrix distance (Vinnicombe distance) between P_1 and P_2.

Alternative expressions involving normalized coprime factorizations are available. Let $P_i = n_i/d_i$ where n_i, d_i are coprime polynomials and let D_i be a stable polynomial with

$D_i^* D_i = n_i^* n_i + d_i^* d_i$. Setting

$$G_i = \begin{bmatrix} n_i/D_i \\ d_i/D_i \end{bmatrix}$$

and

$$\tilde{G}_i = [\,-d_i/D_i \quad n_i/D_i\,]$$

then

$$\kappa(P_1(j\omega), P_2(j\omega)) = |\tilde{G}_1(j\omega)\, G_2(j\omega)|$$

and Equation (A2) is equivalent (with $G_2^* :\overset{\text{def}}{=} G_2^{\mathrm{T}}(-s)$) to

$$\mathrm{wno}(G_2^* G_1) = 0$$

Let $C(s)$ denote the transfer function of a controller, and consider the 2×2 transfer function $T(P, C)$ from the pair $[r_1\, r_2]^{\mathrm{T}}$ to $[y\ u]^{\mathrm{T}}$ for the scheme of Figure 9, which is sometimes known as the generalized sensitivity matrix.

$$T(P, C) = \begin{bmatrix} P \\ 1 \end{bmatrix} (1 + CP)^{-1} [C \quad 1]$$

It is easily seen that

$$\boldsymbol{\sigma}[T(j\omega)] = \frac{\sqrt{1 + |P(j\omega)|^2}\sqrt{1 + |C(j\omega)|^2}}{|1 + P(j\omega) C(j\omega)|}.$$

The generalized stability margin is defined to be

$$b_{P,C} = \frac{1}{\sup_\omega \boldsymbol{\sigma}[T(j\omega)]}$$

if (P, C) is stable, and is zero otherwise.

The main robust stability results are then as follows [13, 15].

Lemma A1. Suppose (P_1, C_1) is a stable closed loop. Then (P_2, C_1) is stable for all plants P_2 satisfying $\delta_v(P_1, P_2) \leqslant \beta$ if and only if $b_{P_1, C_1} > \beta$. Also, (P_1, C_2) is stable for all controllers C_2 satisfying $\delta_v(C_1, C_2) \leqslant \beta$ if and only if $b_{P_1, C_1} > \beta$.

Lemma A2. Suppose (P_1, C_1) is a stable closed loop and that

$$\kappa[P_1(j\omega), P_2(j\omega)]\, \boldsymbol{\sigma}\{T(P_1(j\omega), C_1(j\omega))\} < 1 \quad \forall \omega \tag{A3}$$

Then (P_2, C_1) is stable if and only if Equation (A2) holds. If

$$\kappa[C_1(j\omega), C_2(j\omega)]\,\boldsymbol{\sigma}\{T[P_1(j\omega), C_1(j\omega)]\} < 1 \quad \forall\omega \tag{A4}$$

then (P_1, C_2) is stable if and only if

$$\text{wno}(I + C_2^* C_1) + \eta(C_1) - \tilde{\eta}(C_2) = 0 \tag{A5}$$

A.2. Robust performance results

Robust performance results are expressed in terms of $\boldsymbol{\sigma}[T(P_i, C_j)]$ for various choices of i, j [13, 15].

Lemma 6. Suppose (P_1, C_1) is a stable closed loop. If Equations (A2) and (A3) hold, then at each ω,

$$\boldsymbol{\sigma}[T(P_2, C_1)] \leqslant \frac{\boldsymbol{\sigma}[T(P_1, C_1)]}{1 - \kappa(P_1, P_2)\boldsymbol{\sigma}[T(P_1, C_1)]} \tag{A6}$$

If Equations (A1) and (A5) hold, then

$$\boldsymbol{\sigma}[T(P_1, C_2)] \leqslant \frac{\boldsymbol{\sigma}[T(P_1, C_2)]}{1 - \kappa(C_1, C_2)\boldsymbol{\sigma}[T(P_1, C_2)]}$$

ACKNOWLEDGEMENT

The authors acknowledge funding of this research by the US Army Research Office, Far East, and the Office of Naval Research, Washington, grant number N00014-97-1-0946.

REFERENCES

1. Anderson BDO, Brinsmead TS, De Bruyne F, Hespanha J, Liberzon D, Morse AS. Multiple model adaptive control, Part 1: finite coverings. *International Journal of Robust and Nonlinear Control* 2000; **10**(11–12):909–927.
2. Hespanha JP, Liberzon D, Morse AS, Anderson BDO, Brinsmead TS, De Bruyne F. Multiple model adaptive control, Part 2: switching. *International Journal on Robust and Nonlinear Control*, 2001, **11**(5): 479–496.
3. Hespanha JP, Morse AS. Stability of switched systems with average dwell time. *Proceedings of the 38th IEEE Conference on Decision and Control*, 1999; pp. 2655–2660.
4. Hespanha JP, Morse AS. Towards the high performance control of uncertain processes via supervision. *Proceedings of the 30th Conference on Information Sciences and Systems*, 1996; 405–410.
5. Morse AS. Supervisory control of families of linear set-point controllers — Part 1: exact matching. *IEEE Transactions on Automatic Control* 1996; **41**(10):1413–1431.
6. Morse AS. A bound for the disturbance-to-tracking-error gain of a supervised set-point control system. In *Perspectives in Control, Theory and Applications*, Normand-Cyrot D (ed.). Springer: Berlin, 1998; 23–41.
7. Morse AS. Control using logic-based switching. In *Trends in Control: A European Perspective*, Isidori A (ed.). Springer: Berlin, 1998; 69–113.
8. Pait FM, Kassab F Jr. Parallel algorithms for adaptive control: robust stability. In *Control Using Logic-Based Switching*, Morse AS (ed.). Springer: London, 1997.
9. Pait FM, Kassab F Jr. On a class of switched, robustly stable, adaptive systems. *International Journal on Adaptive Control and Signal Processing*, 2001, **15**(3):213–238.
10. Narendra KS, Balakrishnan J. Adaptive control using multiple models. *IEEE Transactions on Automatic Control* 1997; **42**(2):171–187.
11. Anderson BDO, Gevers M. Fundamental problems in adaptive control. In *Perspectives in Control, Theory and Applications*, Normand-Cyrot D (ed.). Springer: Berlin, 1998.

12. Lee WS, Anderson BDO, Mareels IMY, Kosut RL. On some key issues in the windsurfer approach to adaptive robust control. *Automatica* 1995; **31**(11):1619–1636.
13. Vinnicombe G. Frequency domain uncertainty and the graph topology. *IEEE Transactions on Automatic Control* 1993; **38**(9):1371–1383.
14. Georgiou TT, Smith MC. Optimal robustness in the gap metric. *IEEE Transactions on Automatic Control* 1990; **35**(6):673–687.
15. Vinnicombe G. *Uncertainty and Feedback*: $\mathscr{H}_\infty$ *Uncertainty and the* ν*-Gap Metric*, 1999, Imperial College Press.
16. Anderson BDO, Brinsmead TS. Connections between ν-gap, additive, multiplicative and inverse multiplicative robust stability conditions. *IEEE Conference on Decision and Control*, 2001, submitted.
17. Ljung L. *System Identification*, *A Theory for the User*. Prentice-Hall International: Enlgewood Cliffs, NJ, 1987.
18. Zhou K, Doyle JC, Glover K. *Robust and Optimal Control*. Prentice-Hall: Englewood Cliffs, NJ, 1996.
19. Goodwin GC, Graebe SE, Salgado ME. *Principles of Control System Design*. Prentice-Hall: Englewood Cliffs, NJ, 1998.

