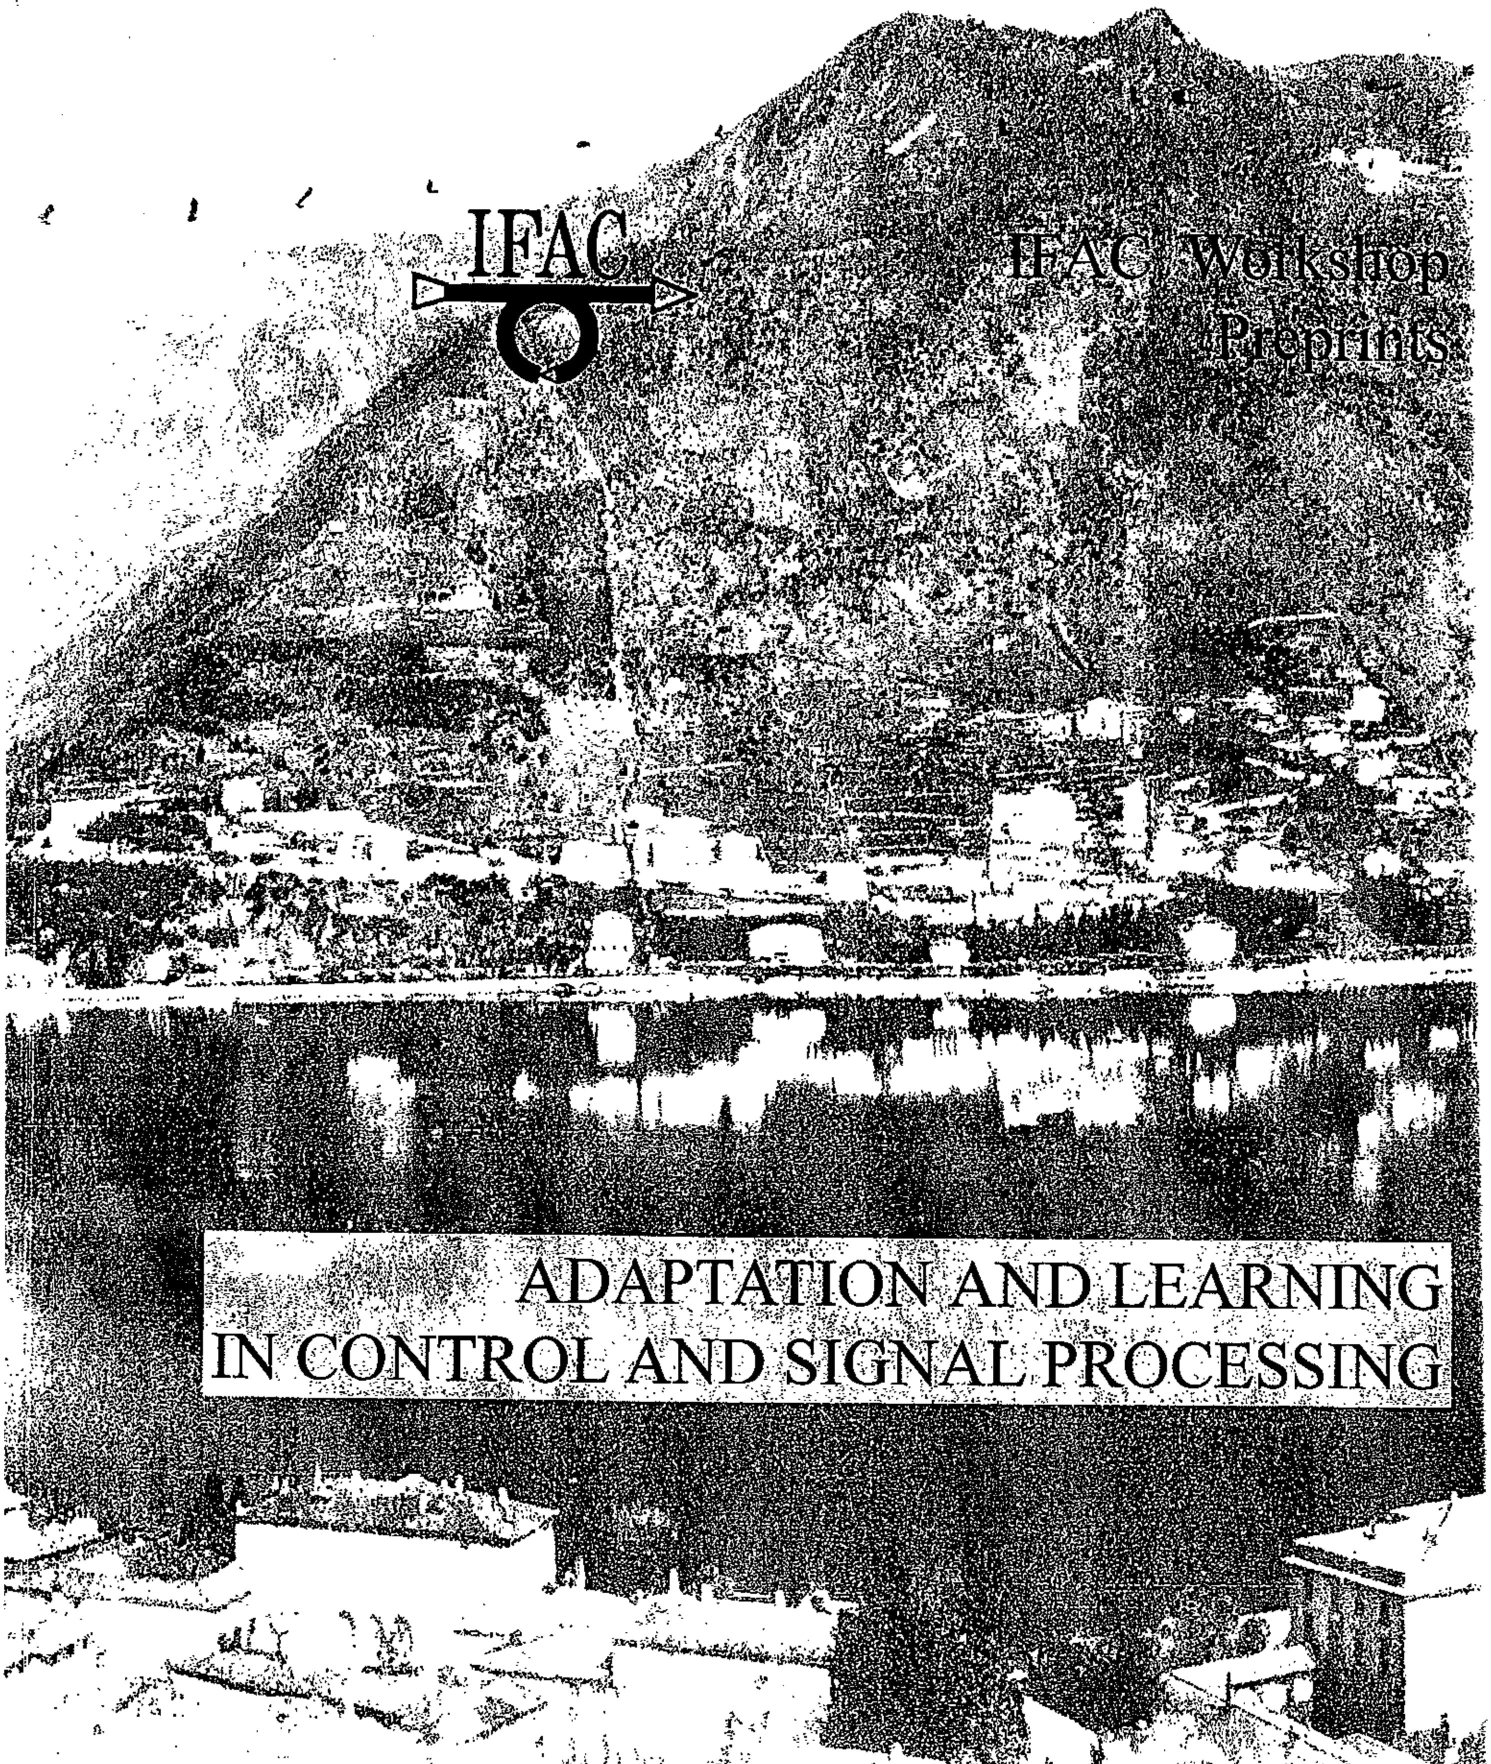

IFAC

IFAC Workshop Preprints

ADAPTATION AND LEARNING IN CONTROL AND SIGNAL PROCESSING

SAFE SWITCHING in MULTI-CONTROLLER IMPLEMENTATION [1]

Brian D. O. Anderson[†],
Thomas S Brinsmead[† 2],
Daniel Liberzon[‡], **A. Stephen Morse**[§]

† *Research School of Info. Sciences and Eng., Australian Nat. Uni., ACT 0200, Australia*

‡ *Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA*

§ *Department of Electrical Engineering, Yale Uni., New Haven, CT 06520-208267, USA*

Abstract- Multiple Model Adaptive Control involves a supervisor switching among one of a finite number of controllers as more is learnt about the plant. 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. A controller scheme is proposed that combines these principles and involves a frequency-dependent performance measure based on the Vinnicombe metric. Safe switching is guaranteed to the extent which closed loop transfer function identification is accurate.

Keywords– 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. **Multiple Model Adaptive Control** (Anderson et al., 2000; Hespanha et al., 2001; Hespanha and Morse, 1999; Morse, 1996, 1998; Pait and Kassab, 2001) postulates that an unknown true plant is within a small uncertainty ball of one (or more) members of a given finite set of nominal plant models.

A controller that gives satisfactory performance is designed for each nominal model and the associated uncertainty ball, giving a finite set of alternative nominal controllers. A "high level" supervisor then switches among the controllers from the finite controller set based on the available observations. 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. The notion of the "best" controller, however, may be ambiguous. Note the following remarks.

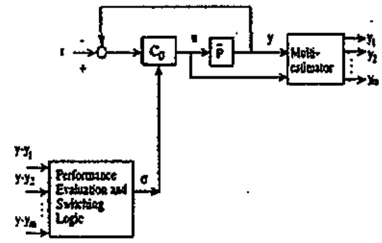

Fig. 1. Outline of multimodel adaptive controller, with $\sigma \in \{1, \cdots, m\}$

(a) The determination of the finite set of nominal plants is addressed in Anderson et al. (2000).
(b) One can regard the supervisor's first task as plant identification, or alternatively, testing a set of hypotheses that the true plant lies in the uncertainty ball around each nominal plant model. That the plant is in a closed loop and that the controller may be changed in the future complicate this identification.
(c) After convergence to one of the nominal controllers further controller tuning for performance improvement is possible (Narendra and Balakrishnan, 1997).

One possible supervisory control architecture from Morse (1996) is depicted in Figure 1. There are m nominal models, $P_1, \cdots, P_m$, for an unknown plant $\bar{P}$, each associated with nominal controllers $C_1, \cdots, C_m$. A multiestimator is driven by the unknown plant input u and output y, with m outputs $y_1, \cdots, y_m$ with the property that if $\bar{P} = P_i$, then $y = y_i$ (after decay of initial condition effects, in the absence of noise, provided all signals are bounded). At time t, the controller $C_{\sigma(t)}$ is implemented. In Morse (1996) the switching logic relies on ***monitoring signals***:

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

At each time t, the controller C_σ is implemented with $\sigma(t) \in \{1, 2, \ldots, m\}$ taken as $\arg\min_{i \in I} \mu_i(t)$. A dwell-time or hysteresis may be imposed to slow down the switching.

If $\bar{P}$ coincides exactly with one of the nominal plants P_I, then, even if the switching process produces unstable signals, provided the exogenous signals are persistently exciting, then the monitoring signals μ_i will remain non-zero for $i \neq I$ and C_I will eventually be selected (Morse, 1996). This work investigates other possibilities for the monitoring signals, which are more compatible a safe switching objective.

The controller may switch before the monitoring signal has converged since this enables the earlier improvement of performance. Of course, if $\bar{P}$ is (slowly) time-varying, the controller may never converge.

---

[1] The authors acknowledge funding of this research by the US Army Research Office, Far East, the Office of Naval Research, Washington, grant number N00014-97-1-0946.
[2] email: Thomas.Brinsmead@anu.edu.au, Phone +61-2-6125 8674, Fax: +61-2-6125 8660

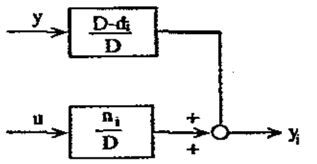

Fig. 2. Constituent part of multiestimator

For such a supervisory structure, the theoretical issues are: the multiestimator design; the boundedness of various signals; and the convergence of $\sigma(t)$ in finite time. One must also consider where $\bar{P}$ does not coincide exactly with any of the nominal plants.

Note that the concept of accuracy of plant approximation only makes sense with respect to the particular controller attached to the plant (Anderson and Gevers, 1998; Lee et al., 1995). The index associated with the "best" nominal model therefore depends upon the currently implemented controller and may not coincide with the index for the "best" controller.

It is hence important to make a careful choice of metric to measure the "closeness" of P_I and $\bar{P}$. We make extensive use of the δ_ν metric (Vinnicombe, 1993) which is related to, but less conservative in a well-defined sense (see Chapter 4 of Vinnicombe (2001)) than, the gap metric (Georgiou and Smith, 1990). Robust stability analysis based on the δ_ν metric relies on a small-gain argument for which more traditional small-gain results are a special case (Anderson and Brinsmead, 2001).

Many adaptive control algorithms, use an explicit or implicit identified model of the plant to design a controller. Hence (a) the true plant normally differs from the model used for controller design and (b) the controller undergoes changes. Such change is potentially dangerous, since even if the original closed loop appears stable the new (frozen) closed loop may be unstable. Of course this results in input signal excitation leading to more accurate identification and a different controller— the mechanism by which many adaptive control schemes guarantee that all signals remain bounded. This is of questionable utility if the "frozen" controller-plant combination is sometimes unstable. There are even algorithms for which all signals are bounded, but the bound is arbitrarily large!

Safe Adaptive Control refers to adaptive control algorithms in which any new controller is guaranteed *a priori* to yield a stable frozen closed loop when combined with the only partially known plant (Lee et al., 1995). In Safe Adaptive Control, controller changes will generally be small (slow). This requires that the set of controllers C_i and thus the set P_i of nominal plant models, is reasonably dense.

2. THE ESTIMATION SUPERVISOR

2.1 *Assumptions*

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, we make some assumptions, some of which will be relaxed later.

- (A1) The reference signal r is stationary with a wide band spectrum;
- (A2) No noise is present;
- (A3) The true plant $\bar{P}$ is linear and time-invariant;
- (A4) The controller C_σ is linear and time-invariant (and is not switched);
- (A5) The nominal plants P_i have transfer functions $P_i = n_i/d_i$ with n_i and d_i coprime polynomials. The part of the multiestimator linking $[y, u]$ to y_i is depicted in Figure 2 where D is a stable polynomial.

2.2 *Interpretation of performance evaluation*

The transfer function: r to $(y_i - y)$ is (Figure 1)

$$W_i = \left(\tfrac{n_i}{D} - \tfrac{d_i}{D}\bar{P}\right)\tfrac{C}{1+\bar{P}C} \qquad (2)$$

If $(\bar{P}, C)$ is stable and r has power spectrum $\Phi_{rr}(\omega)$, the spectrum of $(y_i - y)$ is

$$\Phi_i(\omega) = \left|\tfrac{n_i}{D} - \tfrac{d_i}{D}\bar{P}\right|^2 \left|\tfrac{C}{1+\bar{P}C}\right|^2 \Phi_{rr}(\omega), \qquad (3)$$

and $\int_0^t (y_i - y)^2 d\tau \approx t\int_0^\infty \Phi_i(\omega)d\omega$ so that μ_i in equation (1) will behave like an integrated spectrum, that is, a measure of the error between P_i and $\bar{P}$, frequency weighted by Φ_{rr} and a C dependent term, and integrated over ω. Point-wise, rather than integral, measures of frequency domain quantities, however, may be more useful (see below).

Since for a stable closed loop the μ_i are bounded and for the unstable case they grow exponentially fast, this property could be used to distinguish whether the closed loop is stable. Detecting instability, while indicating that the currently connected C is unacceptable, is unhelpful for suggesting an appropriate replacement. If the closed loop is unstable, then the signals $(y_i - y)$ will grow exponentially and μ_i will be dominated by the transfer function error between P_i and $\bar{P}$ at only a single point in the (right half) complex plane. In such a case, the μ_i measures are likely to be of limited utility.

In the framework of Morse (1996), μ_i is a *scalar* measure associated with each plant (and thus controller) possibility. In principle, however, investigating the frequency content of $(y_i - y)$ allows for a more sophisticated metric.

2.3 *Relaxation of Certain Assumptions*

Under the above restrictive assumptions, we can identify the transfer function from r to $(y_i - y)$ if the $[\bar{P}, C]$ loop is stable. Assumption (A1) can be relaxed (see 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 contoller switching, provided that transient effects are allowed to decay. (A state-shared controller parameterisation ensuring bumpless transfer has been proposed by Morse (1998).)

The term "identification" means finding a transfer function estimate which is understood to have associated error. The error may be described probabilistically, or by hard bounds, depending on the noise models and identification method. "Identification" is not necessarily error free, nor even extremely accurate (which may be possible even in noise, given sufficient identification time). We merely need an estimate that is sufficiently accurate that robust stability and performance analysis can be performed.

2.4 *Performance Evaluation using Transfer Functions*

We make the following key observation.

Lemma 1. Adopt assumptions A1-A6. Let $W_i(j\omega)$ be the transfer function from r to $(y-y_i)$ in equation (2), for the scheme in Figure 1. Let $\kappa(P_i(j\omega), \bar{P}(j\omega))$ be the chordal distance (Vinnicombe, 1993) between P_i and $\bar{P}$ at $s = j\omega$, and let $T(\bar{P}(j\omega), C(j\omega))$ be the generalised sensitivity matrix of $(\bar{P}, C)$. Then $|W_i(j\omega)| =$

$$\kappa(P_i(j\omega), \bar{P}(j\omega))\bar{\sigma}\left[T(\bar{P}, C)\right]\frac{|C|}{\sqrt{1+|C|^2}}\left|\frac{D_i}{D}\right|, \quad (4)$$

where $D_i^* D_i = n_i^* n_i + d_i^* d_i$ for stable polynomial D_i. The proof is omitted for brevity. See Anderson et al. (2001).

The error signal $|W_i(j\omega)|$ in (4) is proportional to the spectrum of the input shaped by $\kappa(P_i, \bar{P})\bar{\sigma}[T(\bar{P}, C)]$. This quantity is critical for possibly allowing us to guarantee that C stabilises P_i, given that it stabilises $\bar{P}$. In a sense, the best P_i, *given the condition that C will be retained* (and a winding number condition involving P_i and $\bar{P}$), is the one that keeps $|W_i||D|\sqrt{1+|C|^2}/|D||C|$ small across the whole spectrum.

Although the quantities μ_i provide a measure of an integrated version of $|W_i|^2$ weighted by $\Phi_{rr}(\omega)$, the above argument indicates that if (approximate) frequency-wise identification is possible, then a Vinnicombe distance criterion may aid choice of the best P_i. Robust performance, as well as robust stability may also be considered.

An alternative multiestimator construction modestly simplifies the above at the expense of realisation complexity. By replacing D in Figure 2 with D_i one has the slightly simpler $W_i = \kappa(P_i, P)\bar{\sigma}[T(\bar{P}, C)]\frac{|C|}{\sqrt{1+|C|^2}}$ instead of (4). Of course, we cannot expect to identify the W_i exactly, but only to within some error bound. Further, the W_i may be changing slowly, due to changing $\bar{P}$. Also, even though quantities like $|W_i|$ measure the approximation error between P_i and $\bar{P}$, given a particular controller C, rather than measuring the best C_i to use for $\bar{P}$, we still use the $|W_i|$ for controller selection. We now focus on the safe switching issue.

2.5 *Changing the Controller*

The goal of the estimator-based supervisor is not to find which P_i is closest to the true plant $\bar{P}$ with the current controller, but what would make a better controller for $\bar{P}$ than the current one. We could select P_i corresponding to the smallest $\kappa(P_i, \bar{P})\bar{\sigma}\{T[(\bar{P}, C)]\}$ or an integrated version thereof, and then hope that C_i makes a good, and certainly stabilising, controller for $\bar{P}$. In order to *guarantee* that C_i stabilises $\bar{P}$, that is, to ensure safe adaptive control a sufficient condition (Vinnicombe, 1993), given the hypothesis that $(\bar{P}, C)$ is stable, is that $\kappa(C, C_i)\bar{\sigma}[T(\bar{P}, C)] < 1 \quad \forall\omega$. In the following section we show how to verify this inequality on the basis of available data.

3. SAFE SWITCHING

We note the following lemma.

Lemma 2. (Vinnicombe, 1993) Suppose (P_1, C_1) is stable and

$$\kappa[P_1(j\omega), P_2(j\omega)]\bar{\sigma}\{T(P_1(j\omega), C_1(j\omega))\} < 1, \forall\omega. \quad (5)$$

Then (P_2, C_1) is stable if and only if

$$\text{wno}(1 + P_2^* P_1) + \eta(P_1) - \tilde{\eta}(P_2) = 0. \quad (6)$$

If $\kappa[C_1(j\omega), C_2(j\omega)]\bar{\sigma}\{T[P_1(j\omega), C_1(j\omega)]\} < 1, \quad (7)$

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

$$\text{wno}(I + C_2^* C_1) + \eta(C_1) - \tilde{\eta}(C_2) = 0. \quad (8)$$

If the controller C_j is currently connected to $\bar{P}$ and we believe that it would be preferable to use C_I, then to be assured of safety, we would like to check

$$\kappa(C_I, C_j)\bar{\sigma}\left\{[T(\bar{P}, C_j)]\right\} < 1 \text{ for all } \omega. \quad (9)$$

Provided a winding number condition involving C_I, C_j holds in addition to (9), then by Lemma 2 C_I is certainly stabilising.

Since the C_i are all known, $\kappa(C_I, C_j)$ is also known. Thus, in order to check (9) we need only to evaluate $\bar{\sigma}\{[T(\bar{P}, C_j)]\}$. Although the performance estimator is not configured to yield $\bar{\sigma}[T(\bar{P}, C_j)]$ directly, even if it were to operate as a transfer function identifier, the following lemma and its corollaries, variants of the small gain theorem, specialised so that the terms can be estimated, allow us to calculate an overbound for this quantity.

Lemma 3. Suppose (P_1, C_1) is stable. If equations (5) and (6) hold, then at each ω,

$$\bar{\sigma}[T(P_2, C_1)] \leq \frac{\bar{\sigma}[T(P_1,C_1)]}{1-\kappa(P_1,P_2)\bar{\sigma}[T(P_1,C_1)]}. \quad (10)$$

If equations (7) and (8) hold then $\bar{\sigma}[T(P_1, C_2)] \leq \frac{\bar{\sigma}[T(P_1,C_2)]}{1-\kappa(C_1,C_2)\bar{\sigma}[T(P_1,C_2)]}$.

Lemma 4. Let $\bar{P}$ be the true plant and let $\{P_1, \cdots, P_m\}$ and $\{C_1, \cdots, C_m\}$ be the collection of nominal plants and controllers. Suppose that C_j stabilises $\bar{P}$ and P_i, and that for all ω

$$\frac{1}{2} > \kappa(P_i, \bar{P})\bar{\sigma}[T(\bar{P}, C_j)]. \text{ Then} \quad (11)$$

$$\bar{\sigma}[T(\bar{P},C_j)] \leq \bar{\sigma}[T(P_i,C_j)]\{1+\kappa(P_i,\bar{P})\bar{\sigma}[T(\bar{P},C_j)]\}. \, (12)$$

The proofs are omitted for brevity.

The Lemma may be used in the following way. With C_j connected to $\bar{P}$ stable, one uses the multiestimator signals to (approximately) identify $\kappa(P_i, \bar{P})\bar{\sigma}[T(\bar{P}, C_j)]$ as a function of frequency for each P_i which is stabilised by C_j (see Section 2.4). Equation (12) then overbounds $\bar{\sigma}[T(\bar{P}, C_j)]$.

A better bounds is given by defining, at each ω: $\bar{\sigma}[T(\bar{P}, C_j)] \leq \bar{T}_j :\stackrel{\text{def}}{=}$

$$\min_{i \in \Im_j} \bar{\sigma}[T(P_i, C_j)] \, \{1 + \kappa(P_i, \bar{P})\bar{\sigma}[T(\bar{P}, C_j)]\} \quad (13)$$

where $\Im_j \subset \{I, \cdots, m\}$ satisfies $i \in \Im_j$ if P_i is stabilised by C_j and equation (11) holds.

We can estimate the right hand side of equation (13) albeit with some error. By Lemma 2 a sufficient condition that C_I will stabilise $\bar{P}$ is that both $\text{wno}(1 + C_I^* C_j) + \eta(C_I) - \bar{\eta}(C_j) = 0$ and (9) hold. In the light of (13) this leads to the following more conservative sufficient condition for stability, involving quantities that we can estimate.

Corollary 1. Given that $[\bar{P}, C_j]$ is stable, then a sufficient condition for $[\bar{P}, C_I]$ to be stable is that $\text{wno}(1 + C_I^* C_j) + \eta(C_I) - \bar{\eta}(C_j) = 0$ and $\kappa(C_I, C_j)\bar{T}_j < 1$.

Notice that if the finite set $\{P_1, P_2, \ldots, P_m\}$ is dense enough in the full uncertainty set of possible unknown plants, then (11) will be straightforwardly satisfied for some i and with fixed j, the quantity $\kappa(C_i, C_j)$ will be small for some i.

If for some particular $i = I$, there holds

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

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

Of course, we cannot, in practice, expect measurements on the closed loop system to yield exact values of $|W_i|$ at each frequency. It is also well-known (Ljung, 1987), that in the presence of stochastic noise, any parameter estimate is subject to variance which increases with the number of unknown parameters and decreases with the quantity of available identification data. Reduced identification variance and hence faster identification times can be achieved by reduction of the number of unknown parameters, but in general, only at the expense of increased identification bias.

If exact values of W_i were available, we could identify $\bar{P}$ and test each C_i with $\bar{P}$. However, exact values are not necessary, since robust stability is ensured merely by the satisfaction of particular inequalities. In the presence of norm bounded noise, however 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 satisfactory operation of the safe switching algorithm even though the identification of $\bar{T}_j$ is subject to noise, and with the checking of the inequality condition of Corollary 1 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, we (imperfectly) identify the frequency response $\bar{V}_i :\stackrel{\text{def}}{=} (P_i - \bar{P}) \frac{C_j}{1+\bar{P}C_j}$. We will make use of the following inequality:

$$\left|\frac{\bar{P}C_j}{1+\bar{P}C_j}\right| \leq \left|\frac{P_iC_j}{1+P_iC_j}\right| + \left|\frac{(P_i-\bar{P})C_j}{1+\bar{P}C_j}\right| \left(1 + \left|\frac{P_iC_j}{1+P_iC_j}\right|\right) \, (14)$$

which is derived from the observations $\frac{\bar{P}C_j}{1+\bar{P}C_j} = \frac{P_iC_j}{1+\bar{P}C_j} + \frac{(\bar{P}-P_i)C_j}{1+\bar{P}C_j}$ and $\frac{P_iC_j}{1+\bar{P}C_j} = \frac{P_iC_j}{1+P_iC_j}\left[1 + \frac{(P_i-\bar{P})C_j}{1+\bar{P}C_j}\right]$. The following lemma is a variant of the small gain theorem.

Lemma 5. Suppose that $\bar{P}$ and P_i are each stabilised by C_j. Then $\bar{P}$ is stabilised 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 $|\Delta|\left|\frac{\bar{P}C_j}{1+\bar{P}C_j}\right| < 1, \quad \forall\omega$. A sufficient condition for the above is that

$$|\Delta|\left[\left|\frac{P_iC_j}{1+P_iC_j}\right|+\left|\frac{(P_i-\bar{P})C_j}{1+\bar{P}C_j}\right|\left(1+\left|\frac{P_iC_j}{1+P_iC_j}\right|\right)\right]<1. \quad (15)$$

This lemma can be applied as follows. Using some identification algorithm, we estimate the frequency responses $\bar{V}_i$. All other quantities in (15) are known. We can check condition (15), with estimated quantities replacing true quantities. If C_k and C_j have the same number of right half plane poles and (15) holds, then C_k is safe. One can consider all the P_i for which (15) holds subject to (P_i, C_j) being stable, and perform a minimisation over i frequency by frequency, thus obtaining a less conservative sufficiency condition.

Whether this scheme or that of Section 3 will give better results will depend on the particular problem. Depending on the controller pole distribution one or other of the two schemes can be used.

5. EXAMPLE

In order to demonstrate the method, the controller switching scheme with safety, as described in Section 3, was implemented in Matlab. The plant to be controlled was chosen as $\bar{P} = \frac{1.2(\frac{1}{2}s+1)(-\frac{1}{4}s+1)}{(\frac{2}{3}s+1)(\frac{1}{3}s+1)(\frac{1}{10}s+1)}$ with 441 plant models used for the multiple model set, each with varying DC gains and non-minimum-phase zero location. The controllers for each model were designed using discrete time Q-synthesis (equivalent to pole placement), with the designed bandwidth dependent on the nominal non-minimum-phase zero location.

In order to maintain persistent excitation, the reference signal was filtered white noise. For the estimation of W_i in equation (4), a standard recursive least squares (Ljung, 1987) algorithm was employed in order to directly identify a vector of discrete time ARMA (auto-regressive, moving average) parameters, and the multiestimator equation (2) was used as the monitoring signal to suggest controller switchings.

The minimum controller switching time was $\tau_S = 1$. At any time after the minimum controller switching time, the controller suggested by the supervisor corresponded to the minimum μ_i. In order to check safety, the inequality condition of Corollary 1 was investigated for a finite number of frequency points, with of course the *estimated* transfer functions $\bar{T}_j$ used in place of the actual values. For comparison, simulations of the algorithm without the safety checking property were also conducted. The full specifications of the (somewhat arbitrary) choices

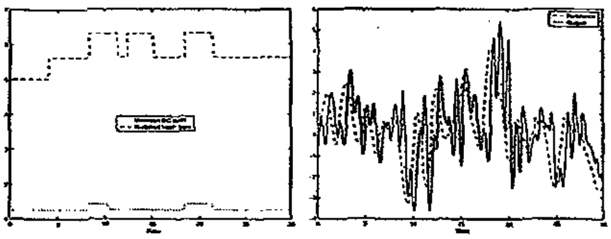

Fig. 3. (Safe) Controller Switching

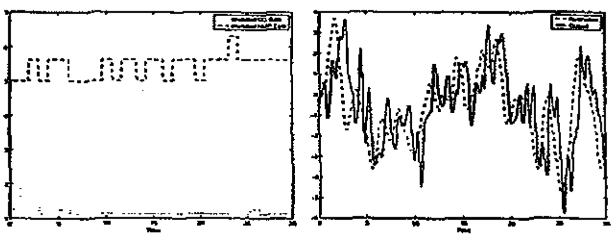

Fig. 4. Controller Switching (without safety)

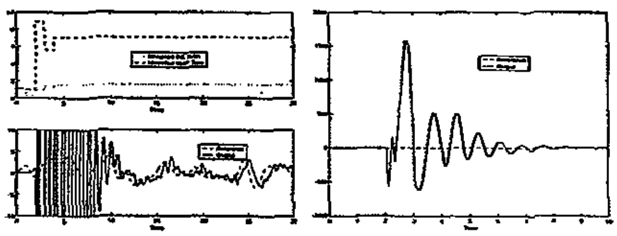

Fig. 5. Temporary Instability (without safety)

of parameters and design methods used for the example appear in Anderson et al. (2001).

5.1 *Results*

The process was simulated and data recorded a total of sixteen (16) times. With the safety property enforced, the allowed controller switchings were much less frequent. A typical controller switching trajectory is shown in Figure 3. For the purposes of interpreting the switching graph, the lines correspond to two modeled parameters: DC gain and non-minimum phase zero location. 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 destabilising, although it did allow controllers with quite poor performance.

For simulations run without safety checking, more frequent controller switching was allowed. A typical controller switching trajectory and output trajectory are shown in Figure 4. Most of the simulations that were run without safety checking exhibited some periods of poor performance. In addition to resulting in poor performance, on a number of occasions, the implemented controller resulted in a destabilised closed loop, although only for a few periods τ_S. See Figure 5 for a simulation with a destabilising controller between 2 and 4 seconds (note the y-axis scale).

6. FURTHER ISSUES

Requirements on r: In the absence of more *a priori* information about $\bar{P}$, it is necessary to assume that r is a wideband signal. Excitation over a wide range of frequencies is required to enable checking the safe switching condition. In general, r needs to be such that a combination of experimental data and *a priori* information is sufficient to assure the satisfaction of the stability conditions.

If $\bar{P}$ is known to be a rational transfer function with known degree k, excitation at k complex frequencies is sufficient to determine $\bar{P}$ at all frequencies. Alternatively, if $\bar{P}$ is unknown up to some frequency Ω_l say, but for $|\omega| > \Omega_l$, one knows that $|\bar{P}| \leq (\omega^2+\Omega_l^2)^{-\frac{1}{2}}$ or if we know that $\bar{P} = \bar{P}_o(I+\Delta)$ for some unknown rational $\bar{P}_o$ of known degree, and some Δ lying within certain frequency-dependent bounds, then the requirement that r be a wideband signal may be somewhat relaxed.

Speed of switching: The switching speed will be limited in order to secure accurate identification for two reasons: switching transients must settle down and noise may be present.

Effect of noise: Although identification in noise can never be exact, we merely need assurance that certain inequalities involving transfer functions are fulfilled. For any given degree of confidence, noise will increase the identification time and thus limits the switching speed.

7. CONCLUSION

This paper has presented some methods to ensure that controller adaptation in Multiple Model Adaptive Control proceeds cautiously and safely, in the spirit of Lee et al. (1995). Vinnicombe metric properties ensure that no controller switchings result in a (frozen) unstable closed loop. A sufficient condition for guaranteeing safety can be checked by available measurements requiring an *estimate* of a particular transfer function. Given absolute or probabilistic confidence estimates on the identification method, it would be possible to give a quantitative absolute or probabilistic guarantees on switching safety. In addition, although the stability guarantee in principle requires checking a condition over an infinite number of frequencies only a finite number can be checked in practice. *A priori* bounds on *complexity* (Vinnicombe, 2001) of the true plant would enable this issue to be addressed rigourously.

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 destabilising controller.

Open research issues include methods: to detect an unstable closed loop; to quickly find a stabilising controller the event that a destabilising controller is switched either initially or during operation; and to find an appropriate switching hysteresis time. This last issue involves a trade-off of fast supervisor response against identification confidence.

References

B. D. O. Anderson and M. Gevers. Fundamental problems in adaptive control. In Dorothée Normand-Cyrot, editor, *Perspectives in Control, Theory and Applications.* Springer, 1998.

B. D. O. Anderson and T. S. Brinsmead. Connections between ν-gap, additive, multiplicative and inverse multiplicative robust stability conditions. *IEEE Conf. Decision & Control,* 2001. submitted.

B. D. O. Anderson, T. S. Brinsmead, F. De Bruyne, J .P Hespanha, D. Liberzon, and A. S. Morse. Multiple model adaptive control, part 1: Finite coverings. *Int. JRNC,* 19:909–927, 2000.

B. D. O. Anderson, T. S. Brinsmead, J .P Hespanha, D. Liberzon, and A. S. Morse. Multiple model adaptive control with safe switching. *Int. J. Adapt. Cont. & Sig. Proc.,* 2001. to appear.

T. T. Georgiou and M. C. Smith. Optimal robustness in the gap metric. *IEEE Trans. Auto. Control,* 35(6):673–687, June 1990.

J. P. Hespanha, D. Liberzon, A. S. Morse, B. D. O. Anderson, T. S. Brinsmead, and F. De Bruyne. Multiple model adaptive control, part 2: Switching. *Int. J. Rob. & Nonl. C.,* 11:479–496, 2001.

J. P. Hespanha and A. Stephen Morse. Stability of switched systems with average dwell time. *Proc.* 38^{th} *IEEE CDC,* 2655–2660, 1999.

W. S. Lee, B. D. O. Anderson, I. M. Y. Mareels, and R. L. Kosut. On some key issues in the windsurfer approach to adaptive robust control. *Automatica,* 31(11):1619–1636, 1995.

L. Ljung. *System Identification, A Theory for the User.* Prentice-Hall Int., NJ, 1987.

A. S. Morse. Control using logic-based switching. In Alberto Isidori, editor, *Trends in Control: A European Perspective,* 69–113. Springer, 1998.

A. S. Morse. Supervisory control of families of linear set-point controllers— part 1: Exact matching. *IEEE TAC,* 41(10):1413–1431, Oct. 1996.

K. S. Narendra and J. Balakrishnan. Adaptive control using multiple models. *IEEE Trans. Auto. Control,* 42(2):171–187, Feb. 1997.

F. M. Pait and F. Kassab. On a class of switched, robustly stable, adaptive systems. *Int. JACSP,* 2001. to appear.

G. Vinnicombe. Frequency domain uncertainty and the graph topology. *IEEE Trans. Auto. Control,* 38(9):1371–1383, Sep. 1993.

G. Vinnicombe. *Uncertainty and Feedback:* $\mathcal{H}_\infty$ *Uncertainty and the* ν*-gap metric.*to appear, 2001.