in "Perspectives in Control, Theory and Applications", D. Normand-Cyrot Ed.; Springer, 1998. C.98

Fundamental Problems in Adaptive Control

B.D.O. Anderson[1] *and M. Gevers*[2]

[1] Research School of Information Sciences and Engineering, and Cooperative Research Centre for Robust and Adaptive Systems, Australian National University, Canberra, ACT 0200, Australia, e-mail: brian.anderson@anu.edu.au

[2] CESAME, Louvain University, Bâtiment Euler, B-1348 Louvain-la-Neuve, Belgium, e-mail: gevers@csam.ucl.ac.be

Summary. [1] The paper identifies three fundamental problems in adaptive control: the need to work with models of plants which may be very accurate but are virtually never exact; the inability to know, given an unknown plant, whether a desired control objective is practical or impractical, and the possibility of transient instability, or extremely large signals occurring before convergence. A technique is advanced for addressing these problems based on controller adjustment limited by Vinnicombe metric considerations.

1. Introduction

The aim of this paper is to highlight some fundamental problems, almost at the conceptual level, with adaptive control, and indicate, at least partially, how they can be resolved.

[1] The authors acknowledge funding of activities of the Cooperative Research Centre for Robust and Adaptive Systems by the Australian Commonwealth Government under the Cooperative Research Centres Program, funding of this research by the US Army Research Office, Far East, the Office of Naval Research, Washington, and the Belgian Programme on Inter-university Poles of Attraction, initiated by the Belgian State, Prime Minister's Office for Science, Technology and Culture.

Modern adaptive control is a creation of approximately the last 20 years. The earlier results of the period, as exemplified in text books such as [5], can be encapsulated in theorems starting with words like "Suppose P is an unknown linear system that is of known degree and minimum phase ... " and ending with words like " ... all signals in the closed-loop remain bounded for all time, the adaptive controller converges, and the performance index is minimized".

There are difficulties with this sort of result.

1.1 The problem of changing experimental conditions given accurate but inexact models

It is now well known that one can find two plants whose Nyquist diagrams or impulse responses are practically indistinguishable, and a controller for the two plants for which the closed-loop behaviours are enormously different - even unstable in one instance and stable in the other [7, 1, 8]. Quoting from [8], "Modelling Principle 1: arbitarily small modelling errors can lead to arbitarily bad closed-loop performance". The higher the performance sought of the controller, the more readily this phenomenon can occur. Conversely one can attach a stabilizing controller to two plants and observe what appear to be identical closed-loop behaviours, when the open-loop behaviours of the plants are quite different [7, 8]. To quote from [8] again, "Modelling Principle 2: larger open-loop modelling errors don't necessarily lead to larger closed-loop prediction errors". (As an example of the latter phenomenon, consider a plant P_0 which is a good low frequency approximation of a plant P_1, which has a high frequency resonance not possessed by P_0. Suppose that C is a controller stabilizing P_0 and yielding a bandwidth outside of which the high frequency resonance falls. Then very similar closed-loop behaviours of the loops (C, P_0) and (C, P_1) will be observed.)

Suppose now the two plants contemplated in the above observations are a "true plant" and a model. Then one can have a marvellous open-loop model of the true plant, but the model quality collapses on introducing a controller. Alternatively, one can have a model which in conjunction with a controller very accurately captures the closed-loop behaviour of a true plant controlled the same way. However, as an open-loop model, it could be very poor.

The issue here is change of experimental conditions. Models, and the task of finding them [identification] can only have their quality evaluated for a particular set of experimental conditions. Changing from open-loop operation to closed-loop operation with a specified controller is of course one change of experimental conditions. *So is any change of a controller a change of experimental conditions*

Unless a plant model is exact, high accuracy under one set of experimental conditions does not guarantee its efficacy under changed experimental conditions. It is intuitively clear that small changes of controller should probably avoid, or ameliorate, the problem of possible loss of efficacy of a model under changed experimental conditions. But how small is small? Adaptive control theory offers no clues as to how to answer this question.

To the extent that in adaptive control (or its batch processing first cousin, iterative control and identification), one changes controller from time to time, based on an assumed model of the plant, one must reckon with the possibility that under a changed controller, the assumed model will be ineffective (if it was derived to be effective for the controller used before the change).

1.2 The problem of impractical control objective

Some years ago, the question was examined of whether the attractive robustness properties (gain and phase margin) of a linear quadratic state feedback control system would extend to the output feedback case. A celebrated paper [2] established with a simple example that there are no guaranteed margins, and presented an example of an LQG optimal system with almost zero margins.

Now contemplate an adaptive control problem, where the unknown plant (though unknown to the controller) was given exactly by that used in [2], and the desired performance index was that of [2] also. What would happen? One would expect very large transient signal values (transient instability), and great difficulty in learning the correct controller, the problems being worse the higher the noise intensity. If the unknown plant (still unknown to the controller of course) was actually given by a small perturbation of the plant of [2], the problem could be far worse.

The issue here is that the control problem posed is a very (impossibly?) demanding one in the known plant case; put another way, the closed-loop system in the known plant case is a marvellous example of a nonrobust design. Attempts to overlay adaption must cause problems.

What makes the problem especially challenging is that the very unknownness of the plant at the outset logically prevents assessment of whether the proposed control objective is or is not feasible. Of course, practitioners of adaptive control have long accepted that one should not ask for the impossible; but this begs the question of how one might know what is impossible in the face of the unknownness.

The challenge is evidently to identify in the course of operation of an adaptive control algorithm that too much is being sought, and to back off on the control objective.

As a tool for addressing the changing experimental condition problem and the impractical control objective problem, we review in the next section results of G. Vinnicombe [9] on robust stability. In Section 3., we turn to ideas for resolving the problems, in a context of iterative identification and control [6]

1.3 The transient instability problem

Recall the pseudo-quotation from a prototypical adaptive control theorem " ... all signals in the closed-loop remain bounded for all time ... ". The apparently comforting words leave a gap so large that not just a truck but an ocean liner could fit through. After all, they allow values of 10^9 volts or amps or watts (even if we would prefer more modest levels). In fact, the situation can be worse. There exist adaptive control algorithms with the following property. Let ξ_0 denote the vector of initial conditions for an adaptive control problem and $z(t)$ the vector at time t of all signals of interest. Then for arbitrary $N > 0$, there exists ξ_0 with $\| \xi_0 \| \leq 1$ and $t_1 > 0$ such that $\| z(t_1) \| > N$, while $z(.)$ is bounded; the bound is not uniform in ξ_0.

The issue is that in the course of the algorithm, controllers are connected which produce unstable behaviour - but they are removed or changed before infinite time, and by the time $t \to \infty$, the controllers have become stabilizing. We term this problem transient instability.

It is not as fundamental as problems 1.1 and 1.2, and indeed certain algorithms have been developed specifically to exclude transient instability, see eg [6].

The problem is also linked with but is not the same as the problem of changing experimental conditions.

2. Vinnicombe Metric

Consider the feedback system of Figure 1.

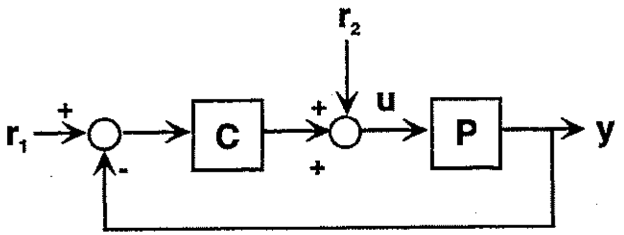

Fig. 2.1. Closed-loop system

The transfer function from $\binom{r_1}{r_2}$ to $\binom{y}{u}$ is

$$T(P,C) = \begin{bmatrix} P(I+CP)^{-1}C & P(I+CP)^{-1} \\ (I+CP)^{-1}C & (I+CP)^{-1} \end{bmatrix} \quad (2.1)$$

$$= \begin{bmatrix} P \\ I \end{bmatrix} (I+CP)^{-1} \begin{bmatrix} C & I \end{bmatrix}. \quad (2.2)$$

One can define a generalized stability margin by

$$b_{P,C} = \begin{cases} \|T(P,C)\|_\infty^{-1} & \text{if } (P,C) \text{ is stable,} \\ 0 & \text{otherwise.} \end{cases} \quad (2.3)$$

Here, $\|G\|_\infty = \sup_\omega \bar{\sigma}[G(j\omega)]$ and for scalar G, $\|G\|_\infty = \sup_\omega |G(j\omega)|$. Good designs correspond to $b_{P,C}$ which are well away from zero. For a scalar P and C, if $|PC(1+PC)^{-1}|$ is large at some frequency, then $|1 - PC(1+PC)^{-1}| = |(1+PC)^{-1}|$ is large at that frequency, and the closed-loop must be close to instability. If $|C(1+PC)^{-1}|$ is large at some frequency, then the plant is likely to be saturated through over-driving: perhaps a design has been attempted which seeks to achieve a wider closed-loop bandwidth than the plant open-loop bandwidth. It is also the case that if certain types of plant uncertainty are present, it is desirable for each of the entries of (1) to be small in order that the uncertainty not do too much damage. For example, it is well known that if $|PC(I+PC)^{-1}|$ is large at some frequency, one must be concerned about neglected high frequency dynamics in the plant, changing number of right half plane zeros, and output (sensor) errors. And if $|P(I+PC)^{-1}|$ is large, low frequency parameter errors and uncertainty about right half plane poles can cause problems, including output errors to input commands and disturbances, [10, 3].

A scalar measure of the difficulty of controlling a particular plant is given by

$$b_{\text{opt}}(P) = \sup_C b_{P,C} \quad (2.4)$$

and $b_{\text{opt}}(P)$ can be related by an elegant formula to the Hankel norm of a normalized coprime realization of P, [4]. This formula exposes clearly the fact that plants with right half plane poles and zeros are harder to control. The larger $b_{\text{opt}}(P)$ is, the smaller can $\| T(P,C)_\infty \|$ be made, i.e. the better one can achieve conflicting objectives of loop following, low sensitivity to noise and disturbances, avoidance of input saturation, and low sensitivity to plant modelling inaccuracies of certain types.

To understand better why $b_{P,C}$ deserves the label of stability margin, we need another concept. Vinnicombe [9] introduced a ν-gap metric defining a distance between any two linear time-invariant plants, stable or unstable, with the same input and output dimensions. For scalar plants, which will be our focus in this paper, the simplest definition is

$$\delta_\nu(P_1,P_2) = \| (I+P_2^*P_2)^{-\frac{1}{2}}(P_2-P_1)(I+P_1^*P_1)^{-\frac{1}{2}} \|_\infty < 1, \quad (2.5)$$

if a certain winding number condition is satisfied, and $\delta_\nu(P_1,P_2)=1$ otherwise [that $\delta_\nu(P_1,P_2)$ is in fact a metric is established in [9]; it is not *a priori* clear from (2.5).]

Now we can state some results of [9].

Proposition 2.1. *1. Given a nominal plant P_1 and a stabilizing compensator C, then (P_2,C) is stable for all plants P_2 satisfying $\delta_\nu(P_1,P_2)\le\beta$ if and only if $b_{P_1,C}>\beta$.*

2. Given a nominal plant P and a stabilizing compensator C_1, then (P,C_2) is stable for all compensators C_2 satisfying $\delta_\nu(C_1,C_2)\le\beta$ if and only if $b_{P,C_1}>\beta$.

Aside from the obvious point that 1 and 2 deal respectively with plant and controller adjustment, let us note several other points,

- If $b_{\text{opt}}(P_1)$ is very small (because P_1 is hard to control), then since $\beta < b_{P_1,C_1} < \sup_C b_{P_1,C} = b_{\text{opt}}(P_1)$, β must be small. So the extent of variation of P_1 to P_2 with guaranteed retention of stability with C_1 is small, as is the extent of variation of C_1 to C_2, with guaranteed retention of stability with P_1
- A more sophisticated version of the Proposition, not needed by us, can consider simultaneous variation of P_1 to P_2 and C_1 to C_2, with guarantees on closed-loop stability
- Proposition 2.1, part 1 implies that among the plants P_2 for which $\delta_\nu(P_1,P_2)=b_{P_1,C}$ one at least will not give closed-loop stability. *It does not mean* that all P_2 with $\delta_\nu(P_1,P_2)=b_{P_1,C}$ will yield an unstable closed-loop system. [Indeed, V Blondel has pointed out to us the example $P_1=s(s+1)^{-1}, P_2=(s+1)s^{-1}$, both of which are stabilized by $C=1$, while also $\delta_\nu(P_1,P_2)=1$].
- Proposition 1 does *not* connect tidily with the Youla-Kucera parametrization of all stabilizing controllers, or all plants stabilized by one controller. Exposing the details would take us too far afield here; suffice it to say that introducing a Youla-Kucera parameterization amounts to perturbing P_1 with a certain directionality and arbitrary large perturbations in a certain direction will not destroy stability, even though the perturbed P_2 may satisfy $\delta_\nu(P_1,P_2)>>b_{P_1,C}$
- Suppose C_1 stabilizes P and that C_2 replaces C_1, with the inequality $\delta_\nu(C_1,C_2)<b_{P,C_1}$ holding, so that C_2 necessarily stabilizes P. Then C_2 may result in a less robust design than C_1: In fact, it can be shown that, see [9]
$$\sin^{-1} b_{P,C_2}\ge \sin^{-1} b_{P,C_1}-\delta_\nu(C_1,C_2) \tag{2.6}$$
(and the equality sign may be obtained)
- instead of choosing C_2 to replace C_1 and restricting $\delta_\nu(C_1,C_2)$ in order to retain stability, one can focus on limiting the change of the closed-loop transfer function matrix. In fact, it can be shown that

$$\delta_\nu(C_1,C_2)\le\| T(P,C_1)-T(P,C_2)\|_\infty\le\frac{\delta_\nu(C_1,C_2)}{b_{P,C_1}b_{P,C_2}} \tag{2.7}$$

and

$$\| T(P,C_1)-T(P,C_2\|_\infty\le\frac{\| T(P,C_1)\|_\infty^2\,\delta_\nu(C_1,C_2)}{1-\| T(P,C_1)\|_\infty\,\delta_\nu(C_1,C_2)}. \tag{2.8}$$

3. Iterative Control and Identification

Iterative control and identification is an approach to adaptive control which decouples the identification and controller design steps. Typically:

- with controller C_i connected to the real system P (and yielding a stable closed-loop), one identifies a model, call it $\hat{P}_j$ of P
- using $\hat{P}_j$, one redesigns the controller, to obtain C_{i+1} and a certain designed or predicted closed-loop performance. The design may in some sense be cautious (i.e. such as to ensure C_{i+1} is not greatly different from C_i). If the combined (P,C_{i+1}) true closed-loop system offers performance like that predicted using $\hat{P}_j$ and C_{i+1}, then one can redesign, each time cautiously and each time obtaining improved performance, to find C_{i+2}, C_{i+3}, until a discrepancy occurs between what is observed in the real (P,C_{i+l}) loop and the design $(\hat{P}_j,C_{i+l})$ loop. (Here, l is the first integer at which the discrepancy occurs.)
- since the loops (P,C_{i+l}) and $(\hat{P}_j,C_{i+l})$ behave differently, one re-identifies to replace $\hat{P}_j$ by $\hat{P}_{j+1}$. Of course, the identification is done in closed-loop.

Then one repeats the cycle, until the desired performance is achieved, or it is evidently unattainable.

One particular version of the above advanced for plants whose worst instability is a pole (possibly repeated) at the origin is set out in [6]. There, the successive designs $C_{i+1},C_{i+2},\ldots$ were obtained by the IMC (internal model control) design concept, with successive broadening of the bandwidth of the desired closed-loop system. It is evident from [6] and related works that only modest jumps in the designed closed-loop bandwidth were made at each new controller design, this being a form of caution to address the changing experimental condition problem raised in Section 1.. One can think of the approach as one which through caution limits $\| T(\hat{P},C_1)-T(\hat{P},C_2)\|_\infty$. By (2.7), this limits $\delta_\nu(C_1,C_2)$.

3.1 Cautious controller adjustment: what is acceptable?

Recall that (P,C_i) is the real loop, and it behaves like $(\hat{P}_j,C_i)$. We wish to adjust C_i to C_{i+1} without encountering the changed experimental condition problem of Section 1.. Can we do this in a planned way, while not knowing

P? Yes, as we shall now explain, we can, on the basis that the (P, C_i) and $(\hat{P}_j, C_i)$ loops behave similarly.

What does it mean to say the (P, C_i) and $(\hat{P}_j, C_i)$ loops behave similarly? We shall take it to mean that

Assumption 3.1. *$T(P, C_i)$ is well approximated by $T(\hat{P}_j, C_i)$.*

An immediate consequence is that

$$b_{P,C_i} \simeq b_{\hat{P}_j,C_i}. \tag{2.9}$$

By the Vinnicombe theory, we can then conclude that if C_{i+1} satisfies

$$\delta_\nu(C_i, C_{i+1}) < b_{\hat{P}_j,C_i}, \tag{2.10}$$

then C_{i+1} will be stabilizing for P. A safety play is to take C_{i+1} such that

$$\delta_\nu(C_i, C_{i+1}) < k b_{\hat{P}_j,C_i}, \tag{2.11}$$

where k is a constant in $(0, 1)$, say 0.5.

3.2 Cautious controller adjustment: what is desirable?

In the previous subsection, we have discussed what, from the point of view of stability, is an acceptable controller C_{i+1} to replace C_i. However, if our goal is to achieve a certain closed-loop performance, by minimizing a performance index say, then we need to understand which C_{i+1} we should choose in a set defined like (2.11).

To fix ideas, we shall postulate

Assumption 3.2. *For the unknown plant P, the design goal is to obtain a stabilizing compensator C to minimize a performance index $J(P, C)$.*

Suppose also that we have a model $\hat{P}_j$ of P such that $T(\hat{P}_j, C_i) \simeq T(P, C_i)$. Let $\hat{P}_j$ have a right coprime realization ND^{-1} and C_i a right coprime realization UV^{-1}. The set of all stablizing compensators of $\hat{P}_j$ is given by [10].

$$\mathcal{C} = \left\{C(Q) : C(Q) = (U - DQ)(V + NQ)^{-1}\right\}, \tag{2.12}$$

where Q (the Youla-Kucera parameter) is an arbitrary stable proper transfer function. Let us make a further postulate, that is certainly fulfilled in the H_2 and H_∞ problems.

Assumption 3.3. *The performance index $J(\hat{P}_j, C(Q))$ for $C \epsilon \mathcal{C}$ depends on Q in a convex manner.*

[Convexity is actually linearity in H_2 and H_∞ problems.] Now we can find C_{i+1} in the following way. Suppose

$$C^*_{i+1} = \arg\min_{C(Q)} J(\hat{P}_j, C(Q)). \tag{2.13}$$

(When the minimum has to be replaced by an infimum, there is a minor adjustment to these calculations.) Let Q^* be such that

$$C^*_{i+1} = (U - DQ^*)(V + NQ^*)^{-1}. \tag{2.14}$$

To avoid trivialities, suppose that $Q^* \neq 0$, i.e. C_i does not minimize $J(\hat{P}_j, C)$.

If

$$\delta_\nu(C_i, C^*_{i+1}) \leq k b_{\hat{P}_j,C_i}, \tag{2.15}$$

with k the constant introduced at the end of the last section, choose

$$C_{i+1} = C^*_{i+1}. \tag{2.16}$$

Otherwise, consider the set

$$C(\alpha Q^*) = (U - \alpha DQ^*)(V + \alpha NQ^*)^{-1}, \qquad \alpha \epsilon [0, 1]. \tag{2.17}$$

Observe that $\alpha = 0$ corresponds to C_i, $\alpha = 1$ corresponds to C^*_{i+1}, and for all $\alpha \epsilon [0, 1]$, $C(\alpha Q^*)$ is stabilizing. Choose $\alpha \epsilon (0, 1)$ so that

$$\delta_\nu(C_i, C(\alpha Q^*)) = k b_{\hat{P}_j,C_i}. \tag{2.18}$$

Such an α exists, since δ_ν is a smooth function of α, taking values at $\alpha = 0$ of 0 and at $\alpha = 1$ of something in excess of $k b_{\hat{P}_j,C_i}$. Also, take

$$C_{i+1} = C(\alpha Q^*). \tag{2.19}$$

Evidently, this choice moves the controller *in the direction of* C^*_{i+1}, but not necessarily all the way; in fact, the movement is such as to retain the bound on δ_ν. Does this assist as far as the performance index is concerned? Yes, it does:

Proposition 3.1. *Suppose that the stable transfer function Q^* minimizes the performance index $J(\hat{P}_j, C(Q))$ which satisfies Assumption 3.3. Let $\alpha \epsilon (0, 1)$. Then*

$$J(\hat{P}_j, C(Q^*)) \leq J(\hat{P}_j, C(\alpha Q^*)) < J(\hat{P}_j, C_i),$$

(where C_i corresponds to $\alpha = 0$).

Proof. The left hand inequality follows by optimality of Q^*. For the right hand inequality observe that by the convexity property of J,

$$J(\hat{P}_j, C(\alpha Q^*)) \leq (1 - \alpha) J(\hat{P}_j, C_i) + \alpha J(\hat{P}_j, C(Q^*)) \tag{2.20}$$

$$< (1 - \alpha) J(\hat{P}_j, C_i) + \alpha J(\hat{P}_j, C_i) \tag{2.21}$$

$$= J(\hat{P}_j, C_i). \tag{2.22}$$

In [6] a performance index of the type described above was not used to determine the controller. Rather, the so called IMC design method was used, where one seeks a controller to achieve a standard closed-loop transfer function in which a single parameter, the bandwidth, appears.

This means that the controller which, in conjunction with a model $\hat{P}_j$, achieves a particular bandwidth is parameterisable by that bandwidth. It is again straightforward to compute a Vinnicombe distance between two such controllers and to set a limit on the change of bandwidth, in terms of $b_{\hat{P}_j,C_i}$.

Let us note that if C_i has been chosen to secure a closed-loop bandwidth exceeding that of the open loop plant P, the entry $C(1+PC)^{-1}$ of $T(P,C)$ will become large, in fact $O[|P^{-1}|]$ outside the plant bandwidth, and accordingly $b_{\hat{P}_j,C_i}$ will be small. This will limit the scope for further bandwidth expansion.

4. Addressing the fundamental problems

In this section, we shall explain how the design idea of Section 3. serves to address the fundamental problems raised in Section 1..

4.1 The problem of changing experimental conditions given accurate but inexact models

The key to addressing this problem is to limit the change in experimental conditions [through placing a bound on $\delta_\nu(C_i, C_{i+1})$] in such a way that the effects of the change of experimental conditions are guaranteed limited.

It is crucial that the bound, although relevant in predicting something about the interconnection of C_{i+1} with the true plant P, is computable in terms of the accurate (but inevitably inexact) model $\hat{P}_j$ - in terms of $\| T(\hat{P}_j, C_i) \|_\infty$ in fact. This is telling us about P precisely because $T(\hat{P}_j, C_i) \simeq T(P, C_i)$. The transfer function matrix $T(P, C_i)$ is not computable, and the closeness of $T(\hat{P}_j, C_i)$ to $T(P, C_i)$ can only be established on the basis of measured signals.

4.2 The problem of impractical control objective

Impractical objectives, technically, are those for which, were the plant known and the objective attained with a controller C, the quantity $\| T(P,C) \|_\infty$ would be very large.

Suppose a problem with impractical objectives is set. How will the ideas of the previous section handle this? One of several things may happen; e.g.

- the identification process, which should yield $T(\hat{P}_j, C_i) \cong T(P, C_i)$ and therefore $b_{\hat{P}_j,C_i} \cong b_{P,C_i}$, actually gives rise to errors in the transfer function estimate (the standard deviation of which may be a consequence of the identification algorithm) of such a magnitude that the equality $b_{\hat{P}_j,C_i} \simeq b_{P_j,C_i}$ cannot be relied upon. (Examples in [6] are of this type). Equivalently, measurement times required to obtaining needed estimates of acceptable quality may simply become excessive.
To understand this better, let us explain one situation where the issue may arise. Consider the scheme of Figure 2.1, and suppose the input r_2 is zero (as is quite frequent). With measurements of r, u and y it is clear that one can obtain straightforwardly an estimate of $PC(1+PC)^{-1}$. If there is additive measurement noise, then the quality of the estimate will be poor when the output signal to noise ratio is small. Now suppose that the closed-loop bandwidth is ω_0, and we consider estimating $PC(1+PC)^{-1}$ at a frequency $10\omega_0$ with r_1 comprising white noise. Then it is conceivable that the SNR at this frequency could be 0dB. So any estimate of $P(1+PC)^{-1}$ based on computing this as C^{-1} times the estimate of $PC(1+PC)^{-1}$ will have a sizeable percentage error (as the SNR is again 0dB). If P has a resonance at $10\omega_0$, $|P(1+PC)^{-1}|$ may be larger than its estimate and in fact the dominant term in $T(P,C)$ evaluated at $10\omega_0$. This means that if $\bar{\sigma}[T(10j\omega_0)]$ approximates $\| T(P,C) \|_\infty$, our estimate of this quantity, and indeed of $b_{P,C}$ may be poor.
- $\| T(\hat{P}_j, C_i) \|_\infty$ may be so large and thus $b_{\hat{P}_j,C_i}$ so small that the scope for adjusting C_{i+1} satisfying $\delta_\nu(C_i, C_{i+1}) \leq k b_{\hat{P}_j,C_i}$ is negligible.

4.3 The problem of transient instability

Elimination of transient instability problems in the scheme of Section 3. requires firstly an additional assumption.

Assumption 4.1. *A stabilizing compensator is known for the unknown plant.*

Given this assumption, the ideas of Section 3. guarantee avoidance of transient instability. This is because all the different controllers individually are stabilizing, and they are not switched so fast as to cause instability through loss of a quasi-time-invariance assumption.

5. Conclusions

In these conclusions, we summarize the key results of the paper, and indicate directions in which further development should be both possible and beneficial.

The main messages have been these:

- any adaptive control scheme needs to embrace modest steps and/or rate of change in the controller. This statement at the qualitative level is folklore, if not a truism. By exploiting ideas of the Vinnicombe metric, one can quantify such steps or rate of change bounds.

- a desired control objective may be impractical, and in an adaptive control conext, this may be unknown; accordingly, any adaptive control scheme needs the capability of recognizing this fact, and the capability of stopping short of the objective. For a wide variety of objectives, one can rely on successive iterates achieving designs which approach the optimum.
- any adaptive control scheme needs protection against the introduction of a controller, which for a finite interval of time, is destabilizing.

What now of shortcomings and future possibilities? Let us note three:

- The Vinnicombe metric is crude, in that it associates a single number with the whole frequency axis. It is virtually certain one could refine the ideas by working with quantities such as $\bar{\sigma}\,[T(P(j\omega), C(j\omega)]$ and $\bar{\sigma}\left[(I+P_2^*(j\omega)P_2(j\omega))^{-\frac{1}{2}}(P_2(j\omega)-P_1(j\omega))(I+P_1^*(j\omega)P_1(j\omega))^{-\frac{1}{2}}\right]$ There are, after all, refinements of the theories of [9], stated in [9], which use frequency dependent quantities. One would hope that such refinements would capture the rough idea [7] that a model of a true plant can provide a basis for a high performance control design if in the vicinity of the cross-over frequency, it is highly accurate
- One area in which adaptive controller application is desired arises when plants can undergo step changes. Immediately following the step change, the previously satisfactory controller may be so unsatisfactory as to be destabilizing. None of the ideas of this paper are helpful in that regard, and quite different approaches are needed. As a general rule, in the presence of noise, hypotheses can be learnt exponentially fast, and parameters learnt at a $\frac{1}{t}$ type of rate; it follows that approaches based on hypothesis testing rather than parameter estimation probably are needed to recover stability in the quickest possible time.
- In iterative control and identification, one may well be identifying the same plant with different controllers. There is an absence of precise theory on how to characterize the performance of an identification algorithm, and even how to design one in this situation. Broadly repeating, when the second controller in considered, and one runs an identification experiment, one does not want to throw away all the information obtained when the first controller was connected. More precisely, how best should one use it? Continuously-varying identification algorithms appear to handle this problem better, focussing as they do on recursive parameter identification; this however has clear disadvantages, for example if there is a risk of initially unexcited high frequency resonances in the plant, and even undermodelling of the plant order.

References

1. R. L. Dailey and M. S. Lukich, "Recent results in identification and control of a flexible truss structure", *Proc Amer Control Conf*, Georgia, 1988, pp 1468-1473.
2. J. C. Doyle, "Guaranteed margins in LQG Regulators", *IEEE Trans Auto Control*, Vol AC-23, 1978, pp 664-665.
3. J. C. Doyle, J. E. Wall and G. Stein, "Performance and robustness analysis for structured uncertainty", *Proc IEEE Conf on Decision and Control*, Orlando, 1982, pp 629-634.
4. K. Glover and D. McFarlane, "Robust stabilization of normalized coprime factor plant descriptions with H_∞ - bounded uncertainty", *IEEE Trans Auto Control*, Vol 34, 1989, pp 821-830.
5. G. C. Goodwin and K. S. Sin, *Adaptive Filtering, Prediction and Control*, Prentice Hall, New Jersey, 1984
6. 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*, Vol 31, 1995, pp 1619-1636.
7. R. Schrama, *Approximate Identification and Control Design*, PhD Thesis, Delft University of Technology, 1992.
8. R. E. Skelton, "Model Error Concepts in Control Design", *Int Journal Control*, Vol 49, 1989, pp 1725-1753.
9. G. Vinnicombe, "Frequency domain uncertainty and the graph topology", *IEEE Trans Auto Control*, Vol 38, 1993, pp 1571-1383.
10. K. Zhou, J. C. Doyle and K. Glover, *Robust and Optimal Control*, Prentice Hall, New Jersey, 1996.