2004 8th International Conference on
Control, Automation, Robotics and Vision
Kunming, China, 6-9th December 2004

Two Decades of Adaptive Control Pitfalls

Brian D.O. Anderson

National ICT Australia and Research School of Information Sciences and Engineering,
Australian National University
Canberra, ACT 0200, Australia
Email: brian.anderson@nicta.com.au

Abstract

Adaptive control is a very appealing technology, at least in principle. Yet its use has been conditioned by an attitude of distrustfulness on the part of some practitioners. In this paper, we review some concepts, the isolation of which was necessary to engender confidence in the technology. These include the unpredictable failure of the MIT rule; the bursting phenomenon, and how to prevent it; the notion that identification of a plant is only valid conceptually for a restricted range of controllers (with the implication that in adaptive control, certain controller changes may be hazardous); and the concept of multiple model adaptive control.

1 Introduction

Adaptive controllers are a fact of life, and have been for some decades. However, theory and practice have not always tracked one another. In this paper, we examine several instances of such a mismatch. These are:

- The MIT rule, an intuitively based gradient descent algorithm that gave unpredictable performance; satisfactory explanation of performance started to become possible in the 1980s.
- Bursting, a phenomenon of temporary instability in adaptive control algorithm implementation of a type observed in the 1970s; explanation and our understanding of avoidance mechanisms only became possible in the 1980s.
- Iterative controller re-design and identification, an intuitively appealing approach to updating controllers that came to prominence in the 1980s and 1990s, and which can lead to unstable performance. Explanation and an understanding of an avoidance mechanism came around 2000
- Multiple model adaptive control, another intuitively appealing approach to adaptive control with the potential to include non-linear systems. It too can lead to unstable performance; early theoretical development left untouched important issues of the number of controllers to be used, and their location in parameter space. These issues have only been dealt with in the last five years

These are not the only examples we might have picked on. Let us point out just two more practical difficulties, which most theory has left unaddressed, and which will not be explored further in this paper.

- There are many adaptive control theorems which run along the lines that if conditions A, B and C hold, then all signals in the loop are bounded and some form of convergence occurs. These theorems leaves open the theoretical probability of 1MA current in a 10W motor; it is not hard to provide simulation examples demonstrating this. The real point is that transient instability, usually from the temporary insertion of a destabilising controller, is a practically unacceptable phenomenon, but often not excluded by the theory.
- Much adaptive control theory postulates unknownness, to some degree at least, of the plant, together with a performance index, which should be minimized. Given that the plant is unknown, it may be the case that the performance index, at least in practical terms, can never be minimized. What if for the real unknown plant, the minimisation yields a controller giving rise to a closed-loop phase margin of .01 radians? Adaptive control theories for the most part sidestep any treatment of this issue.

The structure of the paper is as follows. In Section 2, we present a very high level statement of what constitutes adaptive control. The next four sections treat each of the four problems we first identified above. The final section contains concluding remarks.

2 Adaptive Control

Our aim in this section is to summarize material that is now to be found in scores of textbooks and hundreds of papers. We cite a few textbooks [1–6]. A conventional control loop appears as in Figure 1. The controller is determined using knowledge of the plant together with a list of specifications on the closed-loop performance: the specifications may include the requirement to minimise a certain performance index.

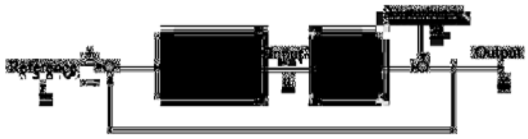

Figure 1: Conventional Control Loop

The first thing that is different in an adaptive control context is that the plant is initially unknown, or only partially known, or it may be slowly varying. Since in many cases, a single controller will not be able to deliver a satisfactory result for any possible plant, it is necessary to incorporate some learning capability in the controller. One can of course conceive of doing this whether or not the plant and controller are linear; for the most part here, we shall confine consideration to linear plants and controllers. A typical non-adaptive controller maps the error signal $r-y$ of Figure 1 into the plant input u in a causal timeinvariant manner:

$$\begin{aligned} \dot{x}_c &= A_c x_c + b_c(r-y) \\ u &= c_c x_c \end{aligned} \tag{1}$$

where A_c, b_c, c_c are constant matrices, and x_c is the state vector of the controller. In an adaptive controller some of the entries of A_c, b_c or c_c are adjusted. One high level scheme is shown in Figure 2

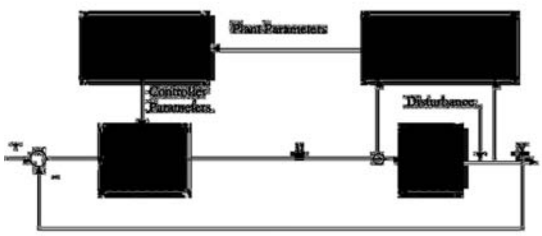

Figure 2: A high level architecture for adaptive control

The job of the identifier is to determine what the plant is, or at least determine an estimate of the plant. The control law calculator does on-line what would be done by the designer before controller introduction in a conventional control problem where the plant is known.

Most of the theory of adaptive control is concerned with writing down algorithms to do these various tasks, and to establishing that the algorithms do what is desired with no untoward side effects.

For future reference, we point out the existence of three time scales associated with the scheme of Figure 2:

- The time scale of the underlying closed-loop dynamics, with fixed plant parameters and controller parameters
- The time scale for identifying the plant
- The time scale associated with plant parameter variations

It is clear that the identification time scale needs to be faster than the plant variation time scale, else identification cannot keep up. It also turns out that it is harder to develop good adaptive controllers, which identify (and thus adjust the controllers) at a time scale comparable with that of the closed-loop dynamics. Interaction of the two processes can occur and generate instability. Nevertheless, operational requirements may force comparability of the time scales, for example in the case of a sudden component failure of a plant, where controller reduction must occur very fast to avoid catastrophe.

3 The MIT Rule

The MIT rule is a scalar parameter adjustment law which was proposed around 1960 for the adaptive control of a linear system modelled as a cascade of a linear stable plant and a single unknown gain [7, 8]. The adjustment law involves approximating a gradient-descent procedure seeking the minimum of an integral-squared performance criterion. The initial intended application was to the control of aircraft dynamics where the single unknown parameter was related to dynamic pressure.

In the history of adaptive control, the MIT Rule represented a watershed; it offered the possibility of adaptation for a useful application, the method was simply formulated, and apparently straightforward to apprehend in an intellectual sense. Performance however turned out to be unpredictable; explanations (as opposed to mere reporting of the performance) took some time to be achieved, see eg [9].

The basic set up is shown in Figure 3. The plant is $k_p Z_p(s)$, where k_p is unknown apart from its sign, and $Z_p(s)$ is a known stable transfer function. The basis of identification is that a cascade adjustable positive known gain $k_c(t)$ is introduced as shown, and the output of the upper arm is compared to the output of $k_m Z_p(s)$ when the same driving signal is applied; here, k_m is a known gain, with the same sign as k_p. Of course, if $k_p k_c(t) = k_m$ for all t, zero error will result and k_p will be given as k_m/k_c. If the error $e(t)$ is nonzero the idea is to adjust k_c to cause it to go to zero. The MIT rule is the rule of adjustment for k_c.

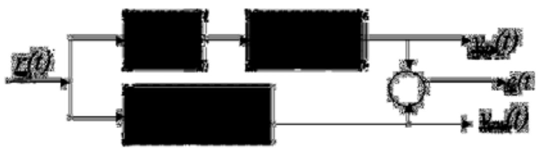

Figure 3: Basic Set-Up for MIT Rule

The idea is to use gradient descent, i.e.,

$$\dot{k}_c = -g \frac{\partial}{\partial k_c} \big[\frac{1}{2} e^2(t)\big] \tag{2}$$

where g is a positive gain constant. Equivalently,

$$\dot{k}_c = -g \big[y_p - y_m\big] y_m. \tag{3}$$

Sometimes this works, and sometimes it does not work. Figure 4 shows the result for a plant $Z_p(s) = (s+1)^{-1}$, with unit amplitude sinusoidal input of frequency ω, and variable adaptive gain g. How can one explain the instability?

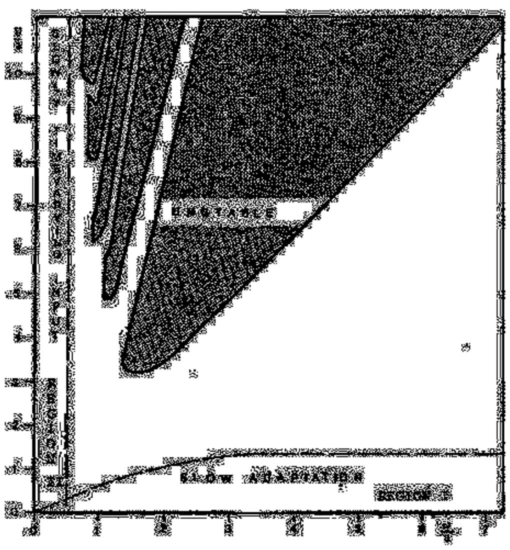

Figure 4: MIT Rule Performance for first order plant with sinusoidal input

Broadly speaking, the instability being displayed is a result of interaction of the adaptive loop dynamics with the plant dynamics, a phenomenon that is occurring when the time scales are comparable, but does not occur otherwise. We shall explain this in more detail.

Suppose first that $r(t) = R$, a constant. Then the MIT Rule (3) leads to a characteristic equation

$$s + g k_m k_p R^2 Z_p(s) = 0 \tag{4}$$

Figure 4 indicates no instability will arise for $Z_p(s) = (s+1)^{-1}$, but the seeds of an instability can be seen in this equation. If $Z_p(s)$ has a right half plane zero, then for high enough gR^2, an instability will occur. If $Z_p(s)$ is of the form $(s^2+as+b)^{-1}$ then again instability will occur if $g k_m k_p R^2 > ab$. Indeed, to ensure stability for all gR^2, one would need

$$\arg Z_p(j\omega) \in (-\pi/2, 3\pi/2) \;\; \forall \omega \in \mathbb{R} \tag{5}$$

If instead of a constant input one has a sinusoidal input, it is not possible to perform as simple an analysis. The differential equation for k_c is actually a Mathieu equation. Instability corresponds to a sort of resonance effect which can only arise when g exceeds a certain threshold. This is why in Figure 4, there are no instabilities for any ω provided g is small enough.

Below, we will develop further some of the above ideas. But, we first need to explain a second instability mechanism which arises when one adjusts the problem statement in a more practical direction. With plant $k_p Z_p(s)$, we have assumed that all the unknownness is in k_p. However, we can also suppose there is some unknownness (eg unmodelled high frequency dynamics) in $Z_p(s)$, and that our knowledge of $Z_p(s)$ is captured by a model transfer function $Z_m(s)$, which is like $Z_p(s)$, but not identical. The error signal we construct is shown in Figure 5. The same rule (2) is still adopted. The expanded form of (3) is now

$$\begin{aligned}\dot{k}_c = -g\big[Z_m(s) k_m r(t)\big]\big[Z_p(s) k_p k_c r(t)\big] + \\ g\big[Z_m(s) k_m r(t)\big]^2\end{aligned} \tag{6}$$

[The notation is of course suggestive, though not strictly proper]

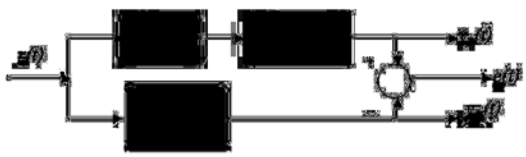

Figure 5: MIT Rule Set-up with Model distinct from Plant

One might argue that the MIT rule, being an adaptive rule, is meant to cope with uncertainties or inaccuracies and that there should therefore be some capability to deal with $Z_p(s)$ unequal to $Z_m(s)$. Figure 6 shows what happens; the regions of gain-frequency pairs giving instability have expanded, and while there is still protection at low frequencies, at high frequencies instability is guaranteed for any gain, no matter how small.

These phenomena can be understood with a tool called averaging theory; references [2] and [5] contain

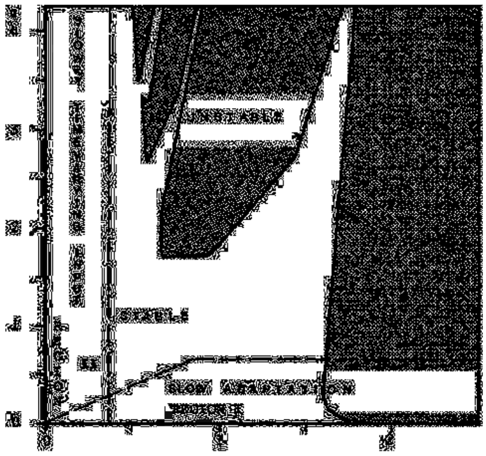

Figure 6: MIT Rule Performance with Model distinct from Plant

much material applying averaging theory to adaptive control. Averaging theory is a tool which is usable given separation of time scales of the closed-loop plant dynamics and the learning/adaptation rate. It indicates that if in (6) the gain g is small enough so that $\dot{k}$ is small, then the behaviour of (6) can be approximated by the behaviour of

$$\dot{k}_c^* = -g\big[Z_m(s)k_m r(t)\big]\big[Z_p(s)k_p r(t)\big]k_c^*(t) + g\big[Z_m(s)k_m r(t)\big]^2 \tag{7}$$

In (6) $k_c(t)$ is processed by the dynamics of $Z_p(s)$, before contributing on the right hand side to making up $\dot{k}_c$. In (7), this in not the case. As further intuition, note that k_p, which is constant, can be pulled out of the action on it by $Z_p(s)$; it is reasonable then that if k_c is nearly constant, the same conclusion remains true for it.

Now stability in (7) is assured if the average value of

$$a(t) := \big[Z_m(s)k_m r(t)\big]\big[Z_p(s)k_p r(t)\big] \tag{8}$$

is positive Since k_m and k_p are constant with the same sign, it is clear that we need $r(t)$ to have the bulk of its energy confined to those frequencies where $Z_m(s)$ and $Z_p(s)$ have similar frequency responses. Obviously, if $Z_p(s) = exp(-s)Z_m(s)$, and $r(t) = sin\omega t$, then for ω suitably large, $a(t)$ will have a negative average value; this accounts for the high frequency behaviour in Figure 6.

4 Bursting

In the early 1980s, scattered reports appeared of adaptive control systems which worked well for a long period say a week and then unexpectedly burst into an oscillation which then died away. In this section, we describe the reasons for the phenomenon, and indicate how such undesirable behaviour can be avoided.

Figure 7 shows the phenomenon; an adaptive controller is connected to a first order plant, and set-point control is sought.

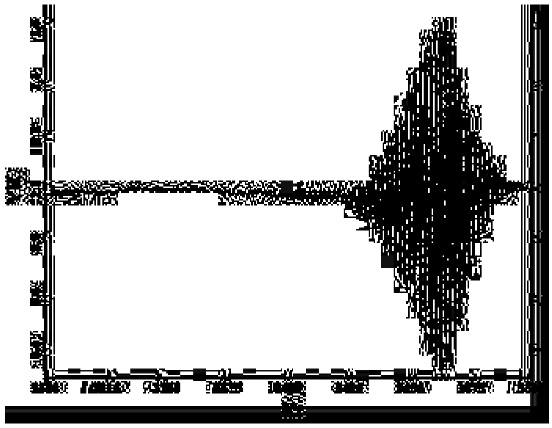

Figure 7: Plant output with adaptive controller connection during a "burst"

The plant in question is described by a differential equation

$$\dot{y} + cy = bu \tag{9}$$

for some constant, unknown b and c. The identifier block in Figure 2 has the task of using the measurements of u and y to determine b and c. Now with most measurements this is possible; but if u is a nonzero constant (as expected with set-point control), the identifier can be expected to identify the plant DC gain, ie the ratio b/c, but not to identify b and c separately. However the identifier does not know that the input is such that it cannot identify b and c separately; it simply runs an algorithm driven by u and y and producing quantities $\hat{b}$ and $\hat{c}$, estimates of b and c, which are supposed to have the property that $\hat{b} \to b$ and $\hat{c} \to c$ as time evolves.

A typical identifier by the way has the form

$$\begin{bmatrix} \dot{\hat{b}} \\ \dot{\hat{c}} \end{bmatrix} = \begin{bmatrix} \text{time function derived from u,y} \\ \text{second time function derived from u,y} \end{bmatrix} \tag{10}$$

and the error obeys

$$\frac{d}{dt} = -\begin{bmatrix} u \\ y \end{bmatrix} [u\ y] \begin{bmatrix} b - \hat{b} \\ c - \hat{c} \end{bmatrix} \tag{11}$$

In case u and y are both constant, the equation then implies

$$u(\hat{b} - b) + y(\hat{c} - c) \to 0 \tag{12}$$

while also $cy = bu$. It follows that $\hat{b}/\hat{c} \to b/c$.

With constant u and y the actual time trajectories followed by $\hat{b}$ and $\hat{c}$ will depend on initial conditions, drift, noise etc. In any case, the adaptive controller uses $\hat{b}$ and $\hat{c}$ to implement a control law. Suppose that the controller simply inserts a gain K with a view to having a high design bandwidth $d = c + bK$. So it will choose $\hat{K} = (d - \hat{c})/\hat{b}$ and the actual closed loop pole will be at $-c - b\hat{K}$. If $\hat{b}/\hat{c} = b/c$, this pole is then at $-bd/\hat{b}$. Since $\hat{b}$ and $\hat{c}$ are not separately constrained but can move round, randomly in many cases, the situation can arise that $-bd/\hat{b} > 0$; instability then occurs. But with instability comes richer signals, and much improved identification. Much improved identification then produces a stabilizing controller, and the whole set-up recovers, until the next time the drifting signals induce an unstable closed loop.

Obviously, one wants $\hat{b} - b$ and $\hat{c} - c$ to go to zero, preferably exponentially fast,to give protection against noise etc. It is nontrivial that a sufficient and virtually necessary condition for this is that for some α_1, α_2 and T all positive and for all s, there holds

$$\alpha_1 I < \int_s^{s+T} \begin{bmatrix} u(\sigma) \\ y(\sigma) \end{bmatrix} [u(\sigma)\ y(\sigma)]\, d\sigma < \alpha_2 I \quad (13)$$

This is termed a persistency of excitation condition. Such a condition was advanced simultaneously by several workers; the simplest proof is probably to be found in [11].

Condition [13] is far from straightforward to apprehend. It uses signals internal to the loop. Alternative conditions have been found, starting with [12], which are conditions on external signals and as such are much easier to verify or arrange to have fulfilled.

If we require the external (reference) signal to be a sum of a sinusoids (regarding a constant signal as a zero frequency sinusoid), then for this problem [13] is guaranteed by requiring that there be at least one complex sinusoid; more generally, if the plant has $m-1$ finite zeros and n poles (thus there are $m+n$ coefficients in the numerator and denominator polynomials of the transfer function to be identified), then there must be an external input signal that excites $m+n$ distinct frequencies (with 0 counting as one frequency and $\pm j\omega$ counting as two frequencies). Such a condition ensures that the generalization of [13] is fulfilled, which in turn means that the adaptive identifier can learn each unknown parameter. Of course, broadband noise will also qualify as a rich enough input signal.

As a practical issue then, set-point control and adaptation are incompatible. Remedies include turning off the adaptation where the input is constant, or superimposing a (presumably small) rich excitation on the external set-point signal to ensure the adaptation does not lead to bursting.

5 Iterative Control and Identification

A frequently advanced approach to adaptive control design is interative identification and controller redesign, see eg [13–15].

This is a form of adaptive control in which the tasks of identification and control are strictly separated. One iteration comprises: (a) identifying the plant with the current controller (b) redefining the controller on the basis of the identified model of the plant, and then implementing it on the real plant. At this point, the identification task is recommenced, with the old information being (largely at least) thrown away. If the plant is unknown initially, but constant, then one expects convergence, ie the controller settles down. However, if the plant is slowly varying, this will not happen, naturally.

While appealing conceptually, the above approach can lead to instability of the closed loop. In the remainder of this section, we shall explain why this is so.

At each iteration, there is delivered a model of the plant (ie the output of the identification process). We would term it a good model if a simulation of the model with a copy of the current controller behaves like the plant connected with the current controller; normally, we would expect the identification step to yield a good model. At this point, the algorithm changes the controller to better reflect the control objective. The controller change is determined by working with the current plant model, with the new controller being attached to the actual plant. One knows that the current controller connected to the plant will behave like the current controller connected to the (current) model. One wants the new controller connected to the plant to behave like the new controller connected to the current model.

This will not necessarily happen. Consider two transfer functions, a model transfer function $P_1 = (s+1)^{-1}$ and a plant transfer function $P_2 = (s+1)(0.1s+1)^{-1}$. Consider Figure 8. The left hand figure shows the open-loop step responses (ie the responses with a zero controller). The right hand figure shows the responses with two different constant gain controllers. It is evident that with gain 100, the closed-loop responses are very different.

The conclusion is that a model may be a good model of the plant with one controller, but it is not guaranteed to be a good model of the plant for all controllers. Therefore, if it is used as a basis for controller redesign, one even has the risk that a new controller, while fine with the model, could destabilise the plant.

There are at least two approaches to deal with this problem. One, due to Safonov and colleagues, eg [16, 17] is able to certify, for a wide class of con-

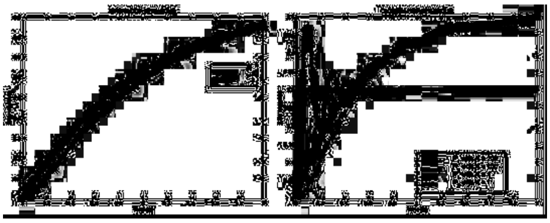

Figure 8: Comparison of Plant and Model responses in open loop and closed loop

trollers and despite the fact that the model is not identical with the plant, that insertion of the replacement controller will not produce closed-loop instability. The second uses gap metric ideas to identify controller changes which are small enough to not cause instability. If the iterative design calls for a big controller change, one moves from the current controller "in the direction of" but not all the way to the newly designed controller, ie one makes a "safe change", or one which will not induce instability, though it does improve performance, see eg [14, 15, 18, 19].

The windsurfing approach to adaptive control [9] (which is an interative identification and controller redesign approach) is a very good example; at each controller re-design step, one expands the closed-loop bandwidth by a small amount, which is consistent with making a safe and small controller change. The closed-loop bandwidth is expanded out to a design objective, or the algorithm indicates that, because of the identification uncertainty, there can be no guarantee that a further bandwidth expansion is safe, and pursuit of a wider closed-loop bandwidth should be abandoned.

This notion of having a flag in an adaptive control algorithm to indicate the inappropriateness of an originally posed objective is practically important, and missing from older adaptive control literature. Logic really demands it. If a plant is initially unknown or only partially unknown, a designer may not know a priori that a proposed design objective is or is not practically obtainable for the plant. Having the algorithm discover this is helpful.

6 Multiple Model Adaptive Control

Imagine a bus on a city street. The equations of motion have parameters which depend on the load, and the friction between the tyres and the road. The friction coefficient will depend on the road surface, including the amount of oil and/or water on the surface. This is an example of a plant which may have a complicated transfer function, or even a nonlinear description, while also containing a (frequently small) number of physical parameters which are unknown and/or changeable. Call such a plant $P(\lambda)$ where λ is the vector of plant parameters.

It would be nice if one could learn λ from measurements by some equations such as

$$\dot{\hat{\lambda}} = f(\hat{\lambda}, measurements) \tag{14}$$

with $\hat{\lambda} \to \lambda_{true}$. This may however be too hard, especially for nonlinear plants, other than in some specific cases.

A completely different approach, multiple model adaptive control, has been suggested to cope with this situation see eg [20–22]. Suppose the unknown parameter vector λ lies in a bounded closed simply-connected region Λ with $\bar{P}$ the associated true plant. Choose a representative set of values $\lambda_1, \lambda_2, ..., \lambda_N$ in the region, with associated plants $P_1, P_2, ...P_N$. Design N controllers C_i such that C_i gives good performance with P_i and plants "near" P_i.

The adaptive control algorithm works as follows. With one controller connected, run an algorithm which estimates at any time instant the particular representative model from $P_1, P_2, ...P_N$, call it P_I, which is the best model to explain the measurements of the inputs and outputs of $\bar{P}$. Then connect up C_I to replace the current controller. Hopefully after at most a finite number of switchings, the best controller for the plant is obtained, and indeed the thrust of early theoretical results on multiple model control was to establish conclusions such as this.

There are problems with this intuitively appealing framework at two levels:

- How many plants P_i should be chosen, how does one choose a representative set of plants $P_1, P_2, ...P_N$, and how can it be assured that the controllers C_i will give good performance for plants "near" P_i? (Indeed, what does "near" mean?)

- If controller C_J is connected, and it turns that P_I is the best explainer of $\bar{P}$ (a good model), there is no guarantee that after switching in of C_I to replace C_J, P_I will continue to be a good model of $\bar{P}$. Instead, P_K for some other index K might be a good model. (This point was explained in the previous section).

We postpone for several paragraphs dealing with the first issue. How can the second be resolved? The situation is analogous to that for iterative identification and control. First, instability may be encountered if the basic algorithm is used, for the same reasons as in iterative identification and control. Second, for broad classes of controllers, before a particular C_i is switched in, the methods of Safonov et al, see eg [16], allow prospective evaluation of its suitability, including closed-loop stability. Third, one can conceive of safe switching, ie depending on the quality of the identification of the loop comprising the current controller and true plant, one can determine which controllers are safe to switch in, and one can elect to switch in only one of them, when it is foreshadowed that were the true plant replaced by the best model of it, the controller being switched in would offer superior performance to the current controller.

The algorithm can indeed be modified along these lines [23], and safe switching results, with the penalty that switching occurs less frequently that with the unmodified algorithm (due to certain proposed switching being ruled out on safety grounds). This is a reasonable penalty to pay, given that the unmodified algorithm on occasions gives rise to connection of a destabilising controller.

How now can we deal with the first issue? This problem is addressed in [24]. In outline, one sequentially picks $P_1, P_2, ...P_N$ by a systematic procedure. Choose λ_1 and thus P_1 arbitrarily. Design C_1. Now determine an open ball B_1 around P_1 for which C_1 constitutes a satisfactory design. [One approach is to choose P to be in the ball if and only if $\delta_\nu(P, P_1) < 0.3b_{P_1,C_1}$. Here δ_ν denotes the ν-gap metric distance [25] between P and P_1, and b_{P_1,C_1} is the generalised stability margin, viz.

$$b_{P_1,C_1} = \| T(P_1, C_1) \|_\infty^{-1} \tag{15}$$

with

$$T_{P_1,C_1} = \begin{bmatrix} P_1 \\ I \end{bmatrix} \left[I - C_1 P_1 \right]^{-1} \left[-C_1 \; I \right] \tag{16}$$

Now choose a λ_2 near the outer limit of the ball B_1, and design C_2 using $P_2 = P(\lambda_2)$. Determine an open ball B_2 around P_2 for which C_2 is a good controller. Choose λ_3 near the outer limit of $B_1 \cup B_2$, and so on. Of course, the λ_i must be drawn from Λ. The procedure terminates at some finite N (by the Heine-Borel Theorem).

In [24], a plant collection is considered where there are two scalar parameter that can vary, a gain and a right-half-plane zero position. The scheme just proposed rapidly determines a collection of plants which turns out to be very similar to these obtained in [20], where a trial and error approach was used that must have been tedious.

7 Conclusions

Let us now summarise some of the key lessons from this survey of decades of adaptive control difficulties.

- The MIT rule confirms that keeping adaptation and plant dynamics time scales separate reduces the likelihood of problems. It also emphasises that one should model as well as possible, even if there is adaptive capability.

- The bursting study suggests not to use more parameters than one needs for modelling purposes; it also confirms that to learn, one needs satisfactory experimental conditions. If you want to keep learning, you need more excitation than a constant reference signal can provide.

- The iterative control re-design and identification study reminds us that a good model of a plant is only a good model for some controllers. This is also a message for multiple model adaptive control. Abrupt controller changes can introduce instability, even if the new controller is defined with what has been a good model. **Safe adaptive control** is one remedy.

- For multiple adaptive control, it is possible to pick representative models systematically, at least in the linear cases. Even then however, safe adaptive control procedures should be used, for the same reasons as in iterative control re-design and identification.

A non-linear adaptive control challenge is to push out the MMAC ideas, especially to provide a sound basis for picking representative models.

8 Acknowledgments

This work was supported by an Australian Research Council Discovery Projects Grant and by National ICT Australia, which is funded by the Australian Government's Department of Communications, Information Technology and the Arts and the Australian Research Council through the *Backing Australia's Ability* initiative and the ICT Centre of Excellence Program.

.

References

[1] GC Goodwin and KS Sin, *Adaptive Filtering, Prediction and Control*, Prentice Hall, Englewood Cliffs, NJ 1984

[2] BDO Anderson, RR Bitmead, CR Johnson Jr, PV Kokotovic, RL Kosut, IMY Mareels, L Praly, BD Riedle, *Stability of Adaptive Systems*, MIT Press, 1986

[3] KJ Astrom and B Wittenmark,*Adaptive Control*, Addison Wesley, 1989

[4] KS Narendra and AM Annaswamy, *Stable Adaptive Systems*, Prentice Hall, 1989

[5] S Sastry and M Bodson, *Adaptive Control*, Prentice Hall, 1989

[6] M Krstic, I Kanellakopoulos, P Kokotovic, *Nonlinear and Adaptive Control Design*, John Wiley, 1995

[7] HP Whitaker, "An Adaptive System for Control of the Dynamics Performances of Aircraft and Spacecraft", Inst Aeronautical Services, Paper 59-100, 1959

[8] PV Osbourne, HP Whitaker and A Kezer, "New Developments in the Design of Model Reference Adaptive Control Systems", Inst Aeronautical Services, Paper 61-39, 1961

[9] IMY Mareels, BDO Anderson, RR Bitmead, M Bodson and S Sastry, "Revisiting the MIT Rule for Adaptive Control", Proc 2nd IFAC Workshop on Adaptive Systems in Control and Signal Processing, 1986, 161-166

[10] BDO Anderson, "Adaptive systems, Lack of Persistency of Excitation and Bursting Phenomena", *Automatica* **21**, 1985, 247-258

[11] BDO Anderson, "Exponential Stability of Linear Equations Arising in Adaptive Identification", *IEEE Trans Auto Control*, **AC-22**, 1977, 83-88

[12] BDO Anderson and CR Johnson Jr, "Exponential Convergence of Adaptive Identification and Control Algorithms", *Automatica*, **18**, 1982, 1-13

[13] P Albertos, "Model-based Iterative Control Design", in *Iterative Identification and Control*, P Albertos and A Salas (Eds), Springer 2002, 121 142

[14] RR Bitmead, "Iterative Optimal Control Design", in *Iterative Identification and Control*, P Albertos and A Salas (Eds), Springer 2002, 167 184

[15] M Gevers, "Identification and Validation for Robust Control", in *Iterative Identification and Control*, P Albertos and A Salas (Eds), Springer 2002, 185 208

[16] A Pail and MG Safonov, "Model reference adaptive control using multiple controllers and switching", Proc 42nd IEEE CDC, Hawaii, 2003, 3256 3261

[17] FB Cabral and MG Safonov, Unfalsified Model Reference Adaptive Control using the ellispoid algorithm", Proc 42nd IEEE CDC, Hawaii, 2003, 3250-3255

[18] BDO Anderson and M Gevers, "Fundamental Problems in Adaptive Control", in *Perspectives in Control*, D Normand-Cyrot, Springer, Berlin, 1998, 9-21

[19] BDO Anderson, "Windsurfing Approach to Iterative Control Design", in *Iterative Identification and Control*, P Albertos and A Salas (Eds), Springer 2002, 143-166

[20] AS Morse, "Supervisory Control of Families of Linear Set-point Controllers part 1: Exact Matching", *IEEE Trans Auto Control*, **41**, 1996, 1413 1431

[21] AS Morse, "Control Using Logic-based Switching", in *Trends in Control: A European Perspective*, A Isidori (ed), Springer, 1998, 69-113

[22] KS Narendra and J Balakrishnan, "Adaptive Control Using Multiple Models", *IEEE Trans Auto Control*, **42**, 1997, 171-187

[23] JF Hespanha, D Liberzon, AS Morse, BDO Anderson, TS Brinsmead, and F De Bruyne, "Multiple Model Adaptive Control. Part 2: Switching", *International Journal of Robust & Nonlinear Control*, **11**, 2001, 479-496

[24] BDO Anderson, TS Brimsmead, F De Bruyne, J Hespanha, D Liberzon, AS Morse, "Multiple Model Adaptive Control. Part 1, Finite Controller Coverings", *International Journal of Robust and Nonlinear Control*, **10**, 2000, 909-929

[25] G Vinnicombe, *Uncertainty and Feedback: H_∞ loop shaping and the ν-gap metric*, Imperial College Press, 2000