

HISTORICAL, GENERIC AND CURRENT CHALLENGES OF ADAPTIVE CONTROL

Brian D. O. Anderson^{*,**,1} Arvin Dehghani^{*,**}

** Department of Information Engineering,
Research School of Information Sciences and Engineering,
The Australian National University, Canberra, ACT 0200,
Australia*

*** National ICT Australia Ltd., Locked Bag 8001,
Canberra, ACT 2601, Australia*

Abstract: This paper reviews three different types of challenges to adaptive control. The first group comprises challenges met in the subject's development. They include difficulties associated with the MIT rule, bursting, the Rohr's counterexample and unplanned instability in iterative identification and control. An understanding of these phenomena and mitigating strategies are now available. The second group comprises difficulties that are intrinsic to virtually any adaptive control algorithm, and that have frequently been overlooked. For example, if a plant is unknown, and a control objective is set, the objective may in practical terms be unachievable, and any adaptive control algorithm needs to deal with that possibility. The third group comprises some issues to which researchers are currently devoting significant attention, including multiple model adaptive control and model free design.

Keywords: Adaptive control, adaptive systems, multiple model adaptive control, model free design

1. INTRODUCTION

Adaptive control has been with us for some decades. Many books, for example (Anderson et al., 1986, Astrom and Wittenmark, 1989, Goodwin and Sin, 1984, Narendra and Annaswamy, 1988, Sastry and Bodson, 1994), and hundreds if not thousands papers have been written on the subject, the number of applications appears to be steadily increasing, and it remains an active research area. All this indicates a certain state of health. Nevertheless, the road to this point has not been a smooth one. Much theory was over-sold; experimentation or practical application did not always run according to plan; and practitioners are understandably cautious about the implementation of the technology.

In this paper, we describe a number of problems of adaptive control. The first collection of problems are those where specific difficulties became apparent only sometime after generation of the initial ideas. In this group, we include the MIT rule, bursting, the Rohr's counterexample, and iterative control and identification. The second set of problems we describe are generic problems; these constitute a list of obstacles that exist in nearly all adaptive control problems, and which any algorithm needs to overcome. They include the problem of impractical control objectives, the problem of transient instability, the problem of dealing with a suddenly unstable closed loop, and the problem of changing experimental conditions. Finally, we turn to current problems, associated

with particular approaches to adaptive control. We shall comment on multiple model adaptive control, and model-free adaptive control.

2. OLDER PROBLEMS OF ADAPTIVE CONTROL

2.1 The MIT rule

The MIT rule was proposed around 1960 (Osbourne et al., 1961, Whitaker, 1959) as a way of adjusting a scalar parameter in a control law associated with a plant modelled by a known linear system multiplied by a unknown and variable gain.

In more detail, suppose the plant is $k_p Z_p(s)$ where k_p is unknown apart from its sign, and $Z_p(s)$ is a known and stable transfer function. Set up a model with transfer function $k_m Z_p(s)$ where k_m is a known gain, with the same sign as k_p , and suppose that the model is driven by a signal r and produces an output y_m . Further assume that prior to the plant is located an *adjustable* gain $k_c(t)$, and that the cascade of this gain and the plant are also driven by r but produce an output y_p . Clearly, if we were to have $y_m = y_p$ for all time with a rich input r , there would have to hold $k_c k_p = k_m$, and then the value of the unknown gain k_p is effectively known. In the event that $y_m \neq y_p$, a gradient law aimed at minimizing $(y_p - y_m)^2$ as a function of k_c is used to adjust $k_c(t)$, as follows:

$$\begin{aligned} \dot{k}_c &= -\mu \frac{\partial}{\partial k_c} [(1/2)(y_p - y_m)^2] \\ &= -\mu [y_p - y_m] y_m \end{aligned} \quad (1)$$

¹ Corresponding author. *Firstname.Surname@anu.edu.au*

Here, μ is a positive gain. For some plants $Z_p(s)$, gains μ and input signals r , this works in the sense that $y_p - y_m \rightarrow 0$, (and so k_p is identified via $k_p k_c = k_m$). For other plants, there is instability. Why?

Thinking about this problem causes us to reflect on the underlying time scales in an adaptive control problem. There are in fact at least three: the time scale associated with the signals in the system in the absence of an adaptation process, and thus associated with the bandwidths of input and output and the transfer functions—in our case then this is the inverse of the bandwidth of r and of $Z_p(s)$; the time scale of the adaption process, which is governed by the size of the gain μ , and the time scale of the rate of change of the underlying dynamics, in this case the rate of change of k_p , which is environmentally induced.

The examples of instability arise even when k_p is constant. They also arise when the gain μ gets large; in fact, what is happening is that the two time scales (defined by r and/or $Z_p(s)$ on the one hand and μ on the other hand) become close. On occasions, this will be cause no damage. But as a general rule, in any adaptive control problem, seeking adaption with a time scale approaching that associated with the non-adapting system is asking for trouble. As a detailed study of the MIT rule (Mareels et al., 1986) reveals the problem becomes worse again when there additional uncertainties beyond k_p in the plant. For example, if the real plant has transfer function $k_p e^{-sT} Z_p(s)$ and the model neglects the time delay term, failure of the algorithm can occur at extremely low gains.

Analytical tools are largely lacking for coping with this unseparated time scale situation. On the other hand, if the time scales are separated, analytical tools revolving around *averaging theory*, see e.g. (Anderson et al., 1986, Kosut et al., 1987, Sastry and Bodson, 1994), allow good forecasting of what will happen, and indeed can be used to explain the instability occurring with time delays at low gains. The general conclusion is that *adaptive designs are more likely to be safe when there is a separation of time scales of the adaptation mechanism and the unadapted process*. The supplementary conclusion is that their stability properties are likely to be easier to analyze.

The third time scale mentioned above is associated with changing of the underlying plant parameters. Clearly, adaptation needs to be *faster* than this rate of change, if the rate of change is to be followed. Thus the three time scales are ordered.




Fig. 1. MIT Rule Set-up

2.2 Bursting

Approximately 25 years ago, scattered generally informal reports appeared of adaptive control systems that would work well for an extended period, and then burst into oscillation. Shortly thereafter the oscillation died away. What a worry to a plant manager! The phenomenon, known as bursting, was first explained in (Anderson, 1985).

A brief explanation follows. Suppose that an adaptive controller for a plant contains two adjustable parameters. In effect, the controller must learn appropriate values for those parameters using plant measurements, whether the aim of the controller is to position poles, minimize an LQ index, or do something else. Now if the plant is excited by a constant input, the only thing that can be learnt about the plant is its DC gain; consequently, there can be no basis for tuning *two* parameters in the control, though there may be a basis for tuning one, or tuning a single function of the two parameters. Nevertheless, the adaptive mechanism will tune the two parameters since that is what it is designed to do, and it will use both relevant information (corresponding to the plant DC gain) and irrelevant information (corresponding to the noise in the measurements). In due course, i.e. after some time, the cumulative effect of using the noise can lead the controller to parameter settings that destabilise the plant. At this point, the signals surrounding the plant cease to be constant, but are richer in frequency content, and as such, the adaptive mechanism now receives more relevant information about the plant than previously, and can then correctly adjust the controller parameters.

The take-home message from this brief discussion is that adaptation should only be attempted when there is adequate relevant information to allow tuning of parameters. Roughly speaking, if there are N parameters in a controller (or N parameters in an adjustable model of the plant which is used to determine a controller), there needs to be the equivalent of at least N frequencies (counting a sinusoid as two frequencies) in the signals whose values are used to drive the process determining the parameters. This is termed having *persistence of excitation* or *persistently exciting inputs*; See (Anderson, 1977, Anderson and Johnson Jr, 1982). The common situation of operating a plant in steady state (with constant input) in most situations will *not be enough* to ensure adaptation. Adaptation mechanisms should in general be switched off during steady state operation.

2.3 Rohr's counterexample

Rohr's counterexample, see (Astrom and Bohlin, 1966, Rohrs et al., 1985), reported some simulations of an adaptive controller which demonstrated unstable behavior. The argument advanced to explain this phenomenon was roughly as follows: nearly all theorems establishing convergence of an adaptive controller algorithm de-

pend (or at that time depended) on a hypothesis that a certain transfer function associated with a plant is positive real (Anderson et al., 1986, Astrom and Wittenmark, 1989, Goodwin and Sin, 1984, Narendra and Annaswamy, 1988, Sastry and Bodson, 1994); however, in real life, there are unmodelled dynamics in a plant which mean that the transfer function actually cannot be positive real. Therefore (the argument went, and illogically at this point) the adaptive control system will be unstable. Of course, even if a theorem hypothesis is not fulfilled, the conclusion of the theorem is *not necessarily* invalidated.

The real reasons for the instability of the Rohr's counterexample appear to be two, and they have already been referred to. First, there was lack of persistency of excitation. Second, there was an interaction of time scales (plant bandwidth and adaption gain), made worse by the necessarily inaccurate modelling of the plant at high frequencies. (By 'necessarily inaccurate modelling', we mean that the real plant had higher order than the model, and so the model could never, even by adjustment of its parameters, accurately capture the plant behavior, particularly as it turned out at high frequencies).

Rohr's counterexample was a false alarm, in two senses. First, adaptive control theory was not fatally undermined. Second, the arguments and explanations were themselves false.

2.4 Iterative control and identification

Many adaptive control algorithms proceed on the basis of adjusting the controller parameters continuously (in a continuous time framework) or at each time instant (in a discrete time framework). Such adjustment is based on an implicit or explicit identification of the plant. Iterative control and identification (Albertos and Sala, 2002, Gevers, 2005, 2000, Lee et al., 1993, Schrama, 1992, Van den Hof and Schrama, 1995, Zang et al., 1991) was an attempt to play with the time scales, and in some way separate the processes identifying and adjusting the control. In particular, over a block of time and with the controller fixed, data is collected and an identification of the plant is made using that data. Using that identified model of the plant, the controller is then designed or redesigned or its parameters are retuned, in a single instant so to speak. Any design criterion may be used, e.g. LQ. Then over the next block of time, a further set of data is collected, again with the controller fixed; an identification process is run, and at the end of the timeblock, a new model of the plant is available. This is used to adjust the controller once more. The process repeats. Success in such an algorithm is marked by convergence of the controller parameters to steady values (always assuming the underlying plant parameters are fixed; if they are varying, one would want the controller parameters to be able to track the plant parameters).

Like the MIT rule, this algorithm sometimes works satisfactorily and at other times it does

not. The worst form of unsatisfactory operation of the algorithm here is exemplified by the closed loop becoming unstable after the connection of a new controller. Like the MIT rule, an explanation for this phenomenon is needed, in order then that measures can be taken to avoid the problem.

In iterative identification and control, the average rate of controller variation is almost always much slower than the underlying closed-loop dynamics. Each identification interval is normally long enough that transients associated with any controller switching at the start of the interval die out early in the interval. Thus, despite the fact that the controller undergoes step changes, one cannot explain unsatisfactory behavior as a consequence of violating the separation of time scales dictum. It is however true that unsatisfactory behavior is more likely to be observed when the controller changes are large. The discussion in the next section dealing with the problem of changing experimental conditions sheds further light on the phenomenon, and offers a mitigating strategy.

2.5 Summary of the key points of the section

In this section, we have drawn attention to the desirability of separating the time scales associated with adaptation and the dynamics of the underlying process, and we have emphasised the need for persistency of excitation. However, there are further lessons to be learnt. It is always better to use an accurate model than not to do so: if there is time delay present and it is neglected, or if there are high frequency dynamics present and they are neglected, an algorithm is more likely to fall over. Our discussion has been couched in a framework of linear systems; obviously though, the conclusions are applicable also to nonlinear adaptive control problems.

3. GENERIC PROBLEMS OF ADAPTIVE CONTROL

In this section, we aim to set out certain general challenges that it would seem reasonable to have most adaptive algorithms address—although many do not. The term 'most adaptive algorithms' is meant to cover nonlinear and linear systems, continuous and discrete-time systems, systems with or without unmodelled or unmodellable dynamics, etc.

3.1 Impractical control objectives

Control engineers grounded in classical control know it is possible to formulate control design problems which in practical terms are not possible to solve. An inverted pendulum with more than two rods is a well known example; again, a plant with nonminimum phase zeros well inside the passband and unstable poles may be near

impossible to control, unless additional inputs or outputs are used; another famous example was provided in (Doyle, 1978) and so on. When the plant is initially known, as well as the control objective, it will generally become clear at some point in the design process, if not *ab initio*, that the control objective is impractical.

Now what happens in adaptive control? The catch is that a full description of the plant is lacking. There may be *no way* to decide on the basis of the *a priori* information that the projected design task is or is not practical. So what will happen if an adaptive control algorithm is run in such a case? At the least, the algorithm will not converge. At worst, an unstable closed loop will be established.

Evidently, there is a generic risk that an adaptive control problem with an associated control objective may involve an impractical control objective, and there is nothing that one can do before attempting the design to recognize yet alone deal with the situation. What then can be done? One should seek algorithms that will recognize that a practically impossible task has been set, and bring the adaptation process to a graceful conclusion with a flag marking inability to achieve the objective. Of course, it is desirable that the objective can be to some degree achieved. The majority of adaptive control algorithms do not offer this facility. Among those that do, we note algorithms associated with the ‘windsurfer’ approach to adaptive control (Anderson, 2002, Dehghani et al., 2004, Lee et al., 1993, 1995). The control objective here is often to push the closed-loop bandwidth out a great distance; if the plant high frequency dynamics are not known, one may not know whether or not a desired closed-loop bandwidth provided in the specifications is or is not achievable. The windsurfer approach proceeds by pushing out the closed-loop bandwidth from a small value in a number of modest steps. The algorithm will terminate early if and when it recognizes on the basis of collected data that it is not safe to push the closed-loop bandwidth further out than the currently obtained value.

3.2 Transient instability

A certain style of adaptive control paper and textbook typically includes one or more theorems which say that under assumptions A, B, C, D and E, and with use of algorithm F, all signals remain bounded and as time goes to infinity, such and such happens. All this may be true; the information may be valuable. But the information is *not enough* to allow a user to sleep peacefully knowing that the theorem will protect his or her multi-million dollar plant investment with an adaptive controller connected. This is because the stated conclusion does not rule out the possibility that at some time before time goes to infinity, a particular controller is connected to the plant with the frozen parameters of the controller such that at that instant of time, the plant-controller closed loop is unstable. What happens locally in time may be

rather like bursting: an instability is observed, and then, through the adaptive mechanism learning something, the instability is killed.

Another manifestation of this sort of difficulty arises in those algorithms, like an early algorithm on discrete-time adaptive control, that if a certain divisor, the value of which depends on measured data, happens to be zero, one replaces it by some nonzero quantity determined in accord with some rule. What happens if the divisor is instead 10^{-12} ? The answer is that the signals will be enormous.

Without protection against the introduction of transient instability, loosely defined as connecting a controller to a plant so that the frozen parameter values define an unstable closed loop, any adaptive control algorithm is dangerous from a practical point of view. The possibility of one million amps of armature current is not precluded. The windsurfer approach to adaptive control certainly attempts to rule out this possibility, and we discuss further aspects later in this paper. The term ‘safe adaptive control’ is a useful descriptor of those approaches or algorithms which rule out this behavior from the beginning.

It might be argued that a second cause of transient instability could be contemplated, even in the case when any controller connected to the plant would, with frozen parameter values, be stabilizing; for it is known that a time-varying closed-loop can be unstable, even with the property that it would be stable with frozen parameter values. Thus one could contemplate instability arising even when any connected controller is (with frozen values) stabilizing. However, for such instability to occur, the rate of time variation cannot be too slow. The explanation of instability could therefore be cast in terms of the time constants associated with adaptation starting to conflict with those associated with the normal plant dynamics—and this has already been dealt with above.

Securing safe adaptive control is far from straightforward. We return to the matter later, but let us note here an example of why it can be difficult. Imagine a plant for which a root locus diagram reveals that there are nonintersecting intervals of a scalar gain guaranteeing that an associated closed loop is stable. Suppose that a value of gain in one of these intervals optimizes a performance index, and suppose that the initial closed loop has the gain in a different interval. Between the two intervals lies an interval of destabilising gains. What form of controller can secure a safe adaptive controller? At the beginning, we obviously have to rule out one which smoothly changes the gain. But then we are forced to violate another principle we have advanced, of discouraging the use of fast variation of the adapted parameters, or maintaining time scale separation. The violation may only be local in time however, and thus not fatal.

3.3 Suddenly unstable closed loops

One of the potential uses of adaptive control is to rescue the situation when an unexpected change in the plant occurs—possibly due to a component failure or some kind of environmental step change. It may be that following such a change the closed loop is unstable. Generally speaking, an adaptive controller will have to correct the instability fast, before much damage is done. On the other hand, as we have argued above, fast changes in adaptive controllers may be dangerous, although if the change is one-off and thus localized in time, it may accordingly be sustainable. There remains however the problem of recognizing that instability is occurring and identifying what particular controller change should be implemented. Any *a priori* knowledge of possible fault modes may help; but without that, the adaptive control logic may have a hard task. If the plant is other than very low order, and undergoes one of a number of changes that produces a growing oscillation, masking just about everything else, identifying the new plant might be very difficult, and that may be a prerequisite for obtaining the controller.

3.4 Changing experimental conditions


In adaptive control, it is very frequently the case that a model is explicitly or implicitly constructed, and the performance of the model is in some way compared with the actual system. When the actual system comprises a plant in combination with a controller, the model may for example comprise an estimate of the plant obtained by identification in combination with the same controller. Or the model may be a desired closed-loop transfer function, and the actual system comprises the closed-loop formed by the actual plant and the controller.

It is common to talk about accurate modelling, or a good model. Though a truism, we must recall that such phraseology requires an *a priori* agreement on a measure of quality. Assuming plant and controller are linear, such a measure may or may not involve the relevant transfer functions alone; for the designer may only want a good model up to 10 Hz, or between 10 Hz and 20 Hz, perhaps because input signals are restricted in this way. In this case, the measure of approximation may involve a frequency weighted error between the relevant transfer functions. But this observation illustrates a second point of critical importance to us here: a model may be a good model with one set of experimental conditions, but not with another set; this would occur if for example, the two transfer functions formed from plant and controller and from model and controller were close from 0 to 10 Hz, but not from 10 Hz to 20 Hz, or vice versa.


All this is reasonably obvious. But there is a particular sense in which it becomes extremely important. Consider two transfer functions $P_1(j\omega)$ and $P_2(j\omega)$ and suppose that $|P_1(j\omega) - P_2(j\omega)| <$

.01 for all ω . Then it might be quite reasonable to term $P_2(j\omega)$ a good model of $P_1(j\omega)$. Now consider the closed-loop transfer functions obtained by attaching a controller with transfer function $C(j\omega)$ to each of $P_1(j\omega)$ and $P_2(j\omega)$. The associated transfer functions are $T_i(j\omega) = P_i(j\omega)C(j\omega)[1 + P_i(j\omega)C(j\omega)]^{-1}$. It is frequently forgotten that one closed-loop transfer function may then be a *very bad model* of the other! Figure 2(a) shows two open-loop responses for the plants $P_1(s) = 1/(s + 1)$ and $P_2(s) = 1/(s + 1)(.1s + 1)$. Figure 2(b) shows the responses with two constant gain controllers. Evidently, $P_1(s)$ is a good approximation of $P_2(s)$ in open loop, and with a small feedback gain, but a poor approximation with a large feedback gain. Changing from having no controller in the loop to having a controller in the loop is *another way of changing the experimental conditions*, and there is then no guarantee that what was a good model will remain so. Similarly, if two closed-loop transfer functions are close, it does *not* imply that the two open-loop transfer functions are close. Even more crucially for adaptive control, **if the closed-loop transfer functions are close with one controller, and the controller is changed, the two new closed-loop transfer functions may no longer be close.**

It is this observation which explains the potential difficulty with iterative identification and control. During an identification step, what is determined



(a) Plant and Model responses in Open Loop



(b) Plant and Model responses in Closed Loop

Fig. 2. Comparison of step responses

is not the true plant, but a good model of the true plant with the *current* controller, and associated external excitations. Even if the latter remain unaltered when the controller is changed, what was a good model of the plant may become a poor model after the change. Consequently, one could conceive of the model and the new controller defining a stable closed loop, while the plant and the new controller yield an unstable closed loop. With small controller changes, it is much less likely that this will happen. It is however a real possibility with large controller changes.

How can one determine the smallness or largeness of a controller change? In this paper, we will answer this question from more than one point of view. The first answer follows, and is drawn from (Anderson et al., 2001), which is a development of (Anderson and Gevers, 1998). We restrict attention to scalar plants here.

Let the notation $H(P, C)$ refer to the four transfer functions obtained by interconnecting P and C , viz., $(1 + PC)^{-1}$, $C(1 + PC)^{-1}$, $(1 + PC)^{-1}P$ and $(1 + PC)^{-1}PC$. Usually, we write

$$H(P, C) = \begin{bmatrix} \frac{PC}{1 + PC} & \frac{P}{1 + PC} \\ \frac{C}{1 + PC} & \frac{1}{1 + PC} \end{bmatrix} \quad (2)$$

At any frequency ω , let $\sigma[H(P, C)]$ denote the nonzero singular value of $H(P, C)$ (the other being zero, since H is singular). The complementary sensitivity is $H_{11}(P, C)$, and it is not hard to verify that

$$\sigma^2[H(P, C)] = \left[1 + \frac{1}{|C|^2} \right] |H_{11}(P, C)|^2 + [1 + |C|^2] [1 - |H_{11}(P, C)|^2] \quad (3)$$

For each frequency ω , denote the *chordal distance* between C_1 and C_2 by

$$\kappa(C_1, C_2) = \frac{|C_1 - C_2|}{[1 + |C_1|^2]^{1/2} [1 + |C_2|^2]^{1/2}} \quad (4)$$

The key result states that if C_1 is stabilizing, a sufficient condition for C_2 to be stabilizing is that

$$\kappa(C_1, C_2)\sigma[H(P, C_1)] < 1 \quad (5)$$

It is generally accepted that from the point of view of stability, a design is to be preferred with smaller $\sigma[H(P, C_1)]$. Plants with simultaneous right half plane poles and zeros typically have larger values of this quantity. The inequality makes clear that any change in moving from C_1 to C_2 which guarantees stability is necessarily smaller when the singular value is large, so that plants which are harder to control will only be tolerant of smaller controller changes.

Now to use this criterion, apparently we need to know $\sigma[H(P, C_1)]$, and this quantity involves the plant, which is unknown. But observe that if we have a good model of the plant with the controller C_1 in the loop, then we have approximate knowledge of the complementary sensitivity $H_{11}(P, C_1)$,

and then Eq. (3) gives us the singular value as a function of frequency. To the extent that even with a good model of the plant, the complementary sensitivity $H_{11}(P, C_1)$ is not known exactly, one should not expect that a controller C_2 that almost ensures equality in Eq. (5) would be safe to use. (Of course, it can be that a controller not satisfying the inequality would be safe to use; the inequality is only a sufficient condition for safety.)

In the next section, we shall examine a combined experimental and graphical way of checking that a controller C_2 can safely be inserted, when noisy frequency domain data for $H_{11}(P, C_1)$ are available.

Suppose now that a proposed controller change is ruled out as being too great, according to the above criterion. What can one do? This question is addressed in (Anderson and Gevers, 1998). It is common, though admittedly not universal, that the closed-loop design criterion involves a norm, typically \mathcal{H}_2 or \mathcal{H}_∞ . It is well known that the set of all stabilizing controllers of a particular transfer function, given one stabilizing controller, can be expressed as a set of transfer functions parametrised by a Youla parameter, and that the associated performance index is convex in that parameter. Thus suppose that with a current model of the plant, call it $\hat{P}(j\omega)$, and a current controller $C_1(j\omega)$, the set of stabilizing controllers is defined by

$$C(Q) = (U - DQ)(V + NQ)^{-1} \quad (6)$$

where U, V, N and D are all stable proper transfer functions determined by C_1 and \hat{P} , and Q ranges over the set of stable proper transfer functions.

Suppose that $C^* = C(Q^*)$ is a controller transfer function obtained by minimizing the performance index assuming the plant has transfer function equal to the model, i.e. \hat{P} . Suppose that $\kappa(C_1, C^*)$ fails the sufficiency test for stability of Eq (5). By continuity, the test will be satisfied if $C(Q^*)$ is replaced by $C(\alpha Q^*)$ for suitably small α . The closer α approaches 1 from below, the smaller will be the performance index computed with \hat{P} ; therefore, we should choose α as large as possible but so that $C_2 = C(\alpha Q^*)$ satisfies the stability requirement of Eq. (5), with some margin for error as noted above.

3.5 Summary of the key points of the section

A performance objective for an unknown plant may be practically unachievable, but the fact that this is so may be unknown. Any adaptive control algorithm needs to protect against the possibility. Adaptive algorithms also should never connect a destabilizing controller even temporarily-and risk doing so when the controller change is large, due to the fact that any explicit or implicit model of the plant may change from being a good (accurate) model to a poor model when the controller changes. Coping with sudden instabilities associ-

ated with a plant component failure is a major theoretical problem.

4. CURRENT APPROACHES TO ADAPTIVE CONTROL

4.1 Multiple Model Adaptive Control

Much adaptive control theory has been established for linear plants, and a great deal of the theory postulates that all coefficients in the numerator and denominator of a plant transfer function or a controller transfer function are unknown and are to be found. On the other hand, adaptive control problems arise for nonlinear as well as linear systems, and commonly, unknownness may stem from unknownness of one or more physical parameters, where the dimension of the unknown parameter vector, call it λ , is much less than the dimension of a vector formed by the coefficients of the numerator and denominator polynomials of a transfer function (always assuming the underlying plant is linear). Multiple model adaptive control (MMAC) (Anderson et al., 2000, Hespanha et al., 2001, Morse, 1996, 1995, Narendra and Balakrishnan, 1997) provides an approach to adaptive control which at least in principle can apply to nonlinear as well as linear systems, and which is particularly suited to situations where the dimension of the parameter vector containing the unknownness is small. (There are of course other approaches to nonlinear adaptive control, see e.g. (Krstic et al., 1995). The fundamental premise on which MMAC is built is that a finite set of controllers is obtained, with the property that any plant that could be encountered can be satisfactorily controlled by at least one controller of the set. Here, 'satisfactorily' does not mean optimally, but it may mean that a controller giving performance within $x\%$ of optimal can be obtained.

Thus suppose that the plants are indexed by λ , with $\lambda \in \Lambda$, the set of possible parameter values, which is typically but not always a bounded connected set. A representative set of values $\{\lambda_1, \lambda_2, \dots, \lambda_N\} \in \Lambda$ is obtained, with associated plants P_1, P_2, \dots, P_N ; then N controllers C_1, C_2, \dots, C_N are designed such that C_I gives good performance with P_i and for all plants near P_i . The adaptive control system works by including a supervisor, the task of which is to monitor performance with the current controller, and to implement a switch from the current controller if the performance is not regarded as satisfactory. The monitoring of performance is essentially achieved by looking at the plant input and output and determining which of P_1, P_2, \dots, P_N would be the best explainer, which is a form of identification via hypothesis testing. The premise is that after at most a finite number of switchings, the best controller will be found.

There are some fundamental problems in implementing this approach, as argued in, for example (Anderson, 2005). First, there is no guidance on how many plants P_i should be chosen, and where

they should be chosen. Second, if a particular controller C_J is connected and it turns out that P_I is the best explainer of the measurements of the input and output of the plant, there is no guarantee that after switching of C_J to C_I , P_I will necessarily remain the best explainer of the true plant, i.e. the best model—for the reasons advanced previously. In this situation, the supervisor will have made a switch to an incorrect controller. Third, there is a need to ensure that a destabilising controller is not connected to the plant; as argued in (Anderson et al., 2001), without taking some precautions in the supervisor to prevent this, it is perfectly possible to temporarily connect in a destabilising controller. Indeed, this reference provides for linear plants and controllers a scheme to ensure safe switching. The idea is to determine, based on the quality of identification of the closed loop comprising current controller and the plant, which other controllers are safe to switch in, and then only to allow the supervisor to make a switching of the controller when the safety test is met. This approach more or less takes care of the second problem mentioned above; for safe switching normally involves only a modest change of the controller, and then the likelihood of switching to an incorrect controller is reduced.

How does one deal with the first problem, of choosing the controller set in the first place? The issue is addressed in (Anderson et al., 2000) for linear plants. Roughly speaking, one value of the plant parameter is arbitrarily chosen, yielding P_1 . The controller C_1 is determined and then it is used to determine a ball of plants around P_1 in which control is acceptable. Then a plant P_2 at the boundary of this ball is determined and a controller C_2 determined for P_2 . A ball around P_2 is then found, and so on. Under suitable conditions, a finite number of such balls will cover the set of plants defined by Λ . In (Anderson et al., 2000), balls are characterised using ν -gap metrics. An approach better adapted to mixing adaptive control and optimal, robust control design can be found in (Athans et al., 2005, Fekri et al., 2006, 2007). Actually, these latter references construct the control as a weighted combination of the outputs of several controllers, while adjusting the weights on-line using measurements. When all weights but one go to zero and one weight goes to 1, then a single controller is identified as the correct one. There are some parallels with a rather old paper exploring this idea—see (Kosut and Anderson, 1988).

When it comes to using the MMAC approach with nonlinear systems, it remains difficult to see how one can select 'representative' values from the parameter set Λ , and this is the biggest roadblock. The general architecture should be little different from that applying the linear case, but deciding whether 7, 70 or 700 controllers is appropriate, and if so which ones, is a question with an elusive solution.

4.2 Model-free Adaptive Control

The preceding discussions on earlier adaptive control techniques and methodologies all included the notion of a ‘model’ of the true plant and the existence of an identification mechanism to provide an implicit or explicit plant model to the adaptive algorithm. A different approach to adaptive control is *model-free adaptive control*, which can potentially address some of the problems highlighted in the previous sections as it does not require any form of identification of a plant model. This approach exploits information in the real-time measurement data to forecast performance levels of all candidate controllers simultaneously, *before* insertion of any one of them into the actual closed-loop to replace the existing controller. Put another way, the algorithm purportedly evaluates what the performance levels of all candidate controllers would have been if they had been placed in the closed-loop.

As a result of an interplay between multiple model controllers and switching ideas and model-free adaptive control, the concept of Unfalsified Adaptive Control has emerged and been extensively investigated over the past decade (Cabral and Safonov, 1996, Safonov and Cabral, 2001, Safonov and Tsao, 1997, Stefanovic and Safonov, 2006, Wang et al., 2005). In this context, a controller is said to be falsified by measurement data when it has the following property: if it were in the feedback loop, a nominated performance requirement (which includes stability) on the closed-loop would fail. Otherwise the controller is said to be unfalsified (Safonov and Cabral, 2001, Safonov and Tsao, 1997). The ideas of unfalsified adaptive control promise appealing features and intuitive advantages. A celebrated property is the fact that the adaptive algorithm does not formally require “any prior assumption on the plant” and the plant can be of any order, stable or unstable, non-minimum phase, linear or non-linear; see e.g. (Stefanovic and Safonov, 2006, Wang et al., 2005). Indeed it is stated in such references that all one would need to use the algorithm are (i) a given cost function which satisfies certain prescribed properties, (ii) a set of bi-proper minimum-phase linear controllers (even if the plant is nonlinear), which includes at least one stabilizing controller, and (iii) a switching mechanism. Let us outline this approach, and then note some practical difficulties with it.

Consider the configuration of Fig. 3 and let P be an unknown plant with (u, y) the measurable plant input/output signals and r the reference signal. The signals r, u, y are all assumed to be square integrable over any bounded interval $[0, \tau] \in \mathbb{R}_+$. The time-invariant controllers C_1, C_2, \dots, C_N belong to the given finite candidate controller set \mathcal{C} . At any instant of time, the inserted controller is one in the set \mathcal{C} , while a supervisor can switch among the controllers of the set. The switching occurs in response to the measured data (u, y) and at discrete instants of time with a minimum interval (dwell-time) between successive switches.

More precisely, with \mathbf{Z}_τ denoting the truncated space of measured signals (u, y) from $t = 0$ to the current time τ , and with some abuse of notation, we denote the adaptive control law at time t as $\hat{C}(t, \mathbf{Z}_t)$; the law $\hat{C}(\cdot, \mathbf{Z})$ maps the function $\begin{bmatrix} r(\cdot) \\ y(\cdot) \end{bmatrix}$ (as in Fig. 3, $r - y = e$) to the function $u(\cdot)$. It is assumed that the switching stops after some final switching time t^f and the final controller is denoted by C^f .

The so-called fictitious reference signals $\tilde{r}_1, \dots, \tilde{r}_i$, are constructed using (u, y) data. Given a current time τ , a set of past measured data (u, y) obtained with a fixed plant over $[0, \tau]$, and a new candidate controller $C_i \in \mathcal{C}$ for future use, a fictitious reference signal \tilde{r}_i for this candidate controller is a hypothetical reference signal defined over $[0, \tau]$ that would have produced exactly the measured data (u, y) over $[0, \tau]$ had C_i been in the feedback loop (noise free) with the unknown plant during the entire time period $[0, \tau]$. This fictitious reference signal can be constructed by using $\tilde{r}_i = C_i^{-1}u + y$ for the configuration in Fig. 3, in which it is assumed for simplicity that noise, disturbance and initial conditions x_0 are all zero.

At any instance of time, a hysteresis switching mechanism (Stefanovic et al., 2005, Wang et al., 2005) *decides* which controller should be placed in the closed-loop, through choosing that controller corresponding to the least performance index.

For a candidate controller $C_i \in \mathcal{C}$, the performance index $V(C_i, u, y, \tau)$ is defined to be a mapping $V : \mathcal{C} \times \mathbf{Z}_\tau \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$, where \mathbf{Z}_τ is the truncated space of the measured signals (u, y) from initial to the current time τ . This is a measure of performance of the closed loop system over the interval $[0, \tau]$, had the controller C_i been in the closed-loop with the corresponding \tilde{r}_i as the reference signal and (u, y) as the measured data. High values correspond to poor performance. A typical example of such a performance index (Wang et al., 2005) is

$$V(C_i, u, y, t) = \max_{\tau \in [0, t]} \frac{\|u\|_\tau^2 + \|\tilde{r}_i - y\|_\tau^2}{\|\tilde{r}_i\|_\tau^2 + \alpha}, \quad \alpha \in \mathbb{R}_+ \quad (7)$$

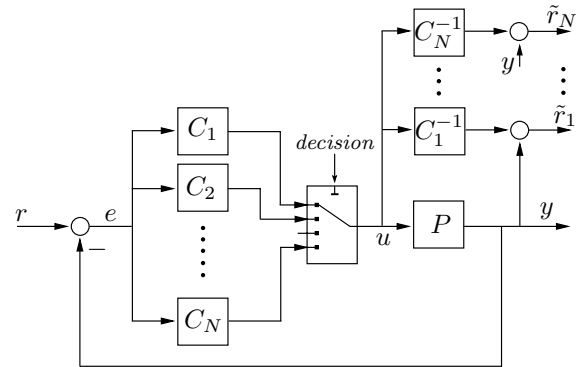


Fig. 3. An Adaptive Control Setup—Unfalsified Approach

which is similar to the \mathcal{H}_∞ mixed-sensitivity performance criterion for linear time-invariant systems (Maciejowski, 1989, Zhou and Doyle, 1998).

The stability and convergence of an unfalsified adaptive control system using multiple controllers and switching are studied in (Stefanovic and Safonov, 2006, Stefanovic et al., 2004, 2005, Wang and Safonov, 2005, Wang et al., 2005), in which it is stated that if the candidate controller set contains linear, finite dimensional, bi-proper and minimum-phase controllers (with a stabilizing one among them), and if the cost function is continuous and monotone increasing in time and has a certain cost-detectability property, then the unfalsified adaptive control algorithm can “reliably” identify controllers that can achieve closed-loop stability and performance specifications based on cost-minimization.

Unfalsified adaptive control is not yet mature and a number of questions are currently unanswered. A recent article (Dehghani et al., 2007a) airs caution in some circumstances and pinpoints some shortcomings of the unfalsified adaptive control approach.

First, the notions of ‘model-free’ and no ‘prior assumption on the plant’ are misleading and reflect a more purist view than most would see as appropriate. Practically speaking, one needs to incorporate as much knowledge of the plant as possible in order to be able to sensibly decide on a set of controllers, and the structure for the controllers in the set. An important but unanswered question is how much information about the plant is required to guarantee a practical and safe model-free adaptive control. Especially is this relevant if one needs to be able to guarantee in advance that the set of controllers contains a stabilizing controller.

Second, the unfalsified adaptive control approach unfortunately gives no guarantee of protection against inserting a destabilizing controller in the closed-loop; moreover, such a destabilizing controller can remain in the closed-loop (before being replaced by a stabilizing controller) for a long period of time resulting in very large control signals; for the example of (Dehghani et al., 2007a) a maximum value of 1.228×10^6 was recorded for the plant input signal $u(t)$ when the reference signal $r(t)$ was a sinusoid of magnitude 1. Indeed, one cannot even put a global upper bound on the time during which the destabilizing controller is attached.

Third, the closed-loop system can actually turn unstable but the instability need not be revealed by a value of the cost function tending to infinity with time, e.g. in (7). Note that a large value for the cost function in Equation (7) would normally be expected if the control signal or the tracking error were large. There, however, can be situations where the control signal $u(t)$ is growing large but the fictitious input \tilde{r}_i is growing large too, resulting in a cost function with no reflection of instability. A closer look will reveal that here we

have a system driven by (u, y) in which \tilde{r}_i is the response rather than having a system driven by the reference and u and y being the response. Ideally, what one would like to be able to do with different C_i is to compare performance indices which have a common reference signal but different u and y for each C_i . Unfortunately, here the comparison is between performance indices which have the same signals u and y but different input signals, and although the mapping $(u, y) \mapsto \tilde{r}_i$ is stable, the mapping $r \mapsto (u, y)$ could be unstable. Thus we unfortunately cannot expect to have instability reflected in the cost function if we never actually connect the destabilizing controller. It seems possible that this is not just a problem with the performance index suggested above, but indeed the whole class of performance indices for which the key theorems concerning model free adaptive control are proved.

In the sequel, we shall briefly review a recent result which offer guarantees of protection against insertion of a destabilizing controller in the closed-loop, in the absence of any form of identified model of the plant.

4.3 Validating controllers via closed-loop data

The term ‘safe adaptive control’ usually refers to adaptive control algorithms which guarantee *a priori* that the introduction of a new controller will always yield a stable closed-loop. The previous sections should have convinced the reader that in spite of copious research attempts, many state-of-the-art adaptive control design methodologies do not explicitly rule out the possibility of inserting a destabilizing controller in the closed-loop, and this deficiency can result in a catastrophe.

The problem of interest is the following: given a stable feedback control interconnection $[P, C_1]$ comprised of the unknown plant P and the stabilizing controller C_1 , and a limited amount of experimental data (e.g. noisy frequency response) obtained with C_1 , how can one verify—without actual insertion in the closed-loop—if the introduction of the new controller C_2 will stabilize the plant P ? The remedy is detailed in (Dehghani, 2007, Dehghani et al., 2007b, Lanzon et al., 2006) and the following theorem defines the experimental framework for the stability tests.

Theorem 1. Let $[P, C_1]$ be internally stable. Let $C_1 = \tilde{V}_1^{-1}\tilde{U}_1$ and $C_2 = \tilde{V}_2^{-1}\tilde{U}_2$ be left coprime factorizations over \mathcal{RH}_∞ (Vinnicombe, 2000, Zhou et al., 1996). Consider the configuration in Fig. 4 and define the mapping $T : r \mapsto z$ to be

$$T = \begin{bmatrix} -\tilde{U}_2 & \tilde{V}_2 \end{bmatrix} \begin{bmatrix} P(I - C_1P)^{-1} \\ (I - C_1P)^{-1} \end{bmatrix} \tilde{V}_1^{-1} \quad (8)$$

and let \arg denote the unwrapped phase (McGowan and Kuc, 1982). Then the following are equivalent:

- a) $[P, C_2]$ is internally stable;
- b) $T^{-1} \in \mathcal{RH}_\infty$;

- c) $\det T(j\omega) \neq 0 \forall \omega$ and $\text{wno} \det T = 0$, where $\text{wno} X$ is short for winding number, i.e. the number of integral multiples of 2π by which the argument of $X(j\omega)$ changes as ω moves from $-\infty$ to ∞ , under the assumption that X is nonzero for all ω ;
- d) $\det T(j\omega) \neq 0 \forall \omega$ and $\arg \det T(j\infty) = \arg \det T(j0)$.

Note that the plant is unknown and hence one cannot explicitly construct T in a closed form. However, the stable mapping $T : r \rightarrow z$ can be studied in a safe—in which no instability can occur—experimental framework, as shown in Fig. 4. The required properties of T can be experimentally inferred by utilizing the reference signal r and the constructed output signal z (computed as a filtered version of the measured signals $\begin{bmatrix} y \\ u \end{bmatrix}$). Observe that the transfer function T can be written as

$$T = \tilde{V}_2(I - C_2P)(I - C_1P)^{-1}\tilde{V}_1^{-1} \quad (9)$$

for which under the assumptions that $\tilde{V}_1(j\infty) = \tilde{V}_2(j\infty) = I$ and the transfer functions PC_1 and PC_2 are strictly proper, we have $\det T(j\infty) = 1$. Thus, $\det T(j\infty)$ is strictly positive and known and will be used as a datum for the verification of condition (d) in Theorem 1.

Moreover, T can be written as

$$T = I + T' \quad (10)$$

$$T' = \begin{bmatrix} -(\tilde{U}_2 - \tilde{U}_1) & (\tilde{V}_2 - \tilde{V}_1) \end{bmatrix} \begin{bmatrix} P(I - C_1P)^{-1} \\ (I - C_1P)^{-1} \end{bmatrix} \tilde{V}_1^{-1} \quad (11)$$

which shows that T is the sum of a known term (i.e. I) and a term which is strictly proper. Hence, it can be expected that measuring the frequency response of T up to a frequency, ω_0 say, where the response of T' has nearly vanished is enough to characterize the full frequency response of T in the direction of the following results. The following test (limited for convenience in its presentation here to the SISO case) measures the frequency response up to a certain finite frequency ω_0 . Importantly, the measurement can tolerate significant error as its purpose is simply to facilitate computation of a certain integer multiple of 2π (as both $T(j0)$ and $T(j\infty)$ are real) corresponding to the phase change.

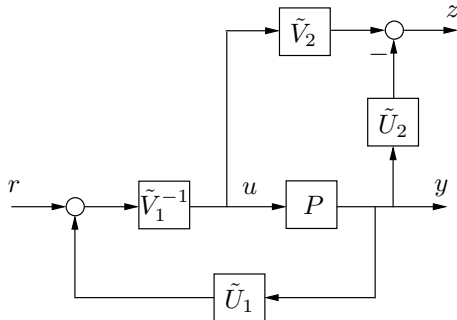


Fig. 4. Experimental setting

Theorem 2. Adopt the hypothesis of Theorem 1 and additionally, suppose that T is scalar. Let $T' = T - 1$, with $T'(j\omega) \rightarrow 0$ as $\omega \rightarrow \infty$.

Let $\omega_0 \in [0, \infty)$ be a frequency such that $|T'(j\omega)| < 1 \forall \omega \geq \omega_0$, i.e.

$$\left| \tilde{V}_1^{-1}(j\omega) \left| \frac{-(\tilde{U}_2 - \tilde{U}_1)P + (\tilde{V}_2 - \tilde{V}_1)}{1 - C_1P} (j\omega) \right| \right| < 1 \quad \forall \omega \geq \omega_0 \quad (12)$$

Then the following condition is equivalent to condition (d) in Theorem 1

$$\begin{cases} T(j\omega) \neq 0 \quad \forall \omega & \text{and} \\ \pi \left\lfloor \frac{\arg T(j\omega_0)}{\pi} \right\rfloor = \arg T(j0) \end{cases} \quad (13)$$

where $\lfloor \cdot \rfloor$ denotes the closest integer.

The aforementioned validation test aims to verify internal stability with the introduction of a new controller C_2 by utilizing a limited amount of experimental data obtained from the stable closed-loop interconnection $[P, C_1]$. Theorem 2 proposes a type of phase test analogous to the Nyquist criterion and utilizes the noisy frequency response information of the closed-loop mapping T up to a finite frequency ω_0 to check if C_2 will stabilize the unknown plant P . This is all achieved despite the restrictive assumption that no *a priori* information about the plant is available. This promises to address the so-called transient instability problem in the context of multiple model adaptive control (MMAC) and iterative identification and control ideas discussed earlier.

Note that the controller implementation in Fig. 4 is a general form of implementing a controller in a feedback loop and is referred to as ‘observer-form’ controller implementation (Dehghani, 2007, Dehghani et al., 2006, Vinnicombe, 2000). However, the results of Theorem 1 and Theorem 2 can be extended to include experimental settings involving alternative controller implementations. When C_1 is in the forward path, and C_2 is potentially to replace it, provided C_1 and C_2 are bi-proper and minimum-phase, one can study the stable mapping $T_1 : r \rightarrow \tilde{r}$ with $\tilde{r} = C_2^{-1}u + y$, as a replacement for T , in a safe experiment. Similarly, for the feedback-path implementation of C_1 , if C_1 and C_2 have no right half-plane poles, the stable mapping $T_2 : r \rightarrow \tilde{w}$ with $\tilde{w} = u + Cy$ can be examined by utilizing the results of this subsection and satisfying the assumptions of Theorem 2.

The following example illustrates the effectiveness of the stability test proposed in Theorem 2.

4.3.1. An Example Let the unknown plant be given by

$$P = \frac{-186.66(s-5)(s+4.5)}{(s+10)^2(s+7)(s+6)}$$

and let C_1 be a stabilizing controller, $H(P, C_1) \in \mathcal{RH}_\infty$, given by

$$C_1 = \frac{0.021(s+10.92)(s+8.87)(s+7.31)(s+5.93)}{(s^2 + 8.6s + 19.84)(s^2 - 0.603s + 5.34)}$$

with a left coprime factorization, $C_1 = \tilde{V}_1^{-1}\tilde{U}_1$,

$$\tilde{V}_1 = \frac{(s^2 + 8.603s + 19.84)(s^2 - 0.602s + 5.34)}{(s^2 + 8.64s + 19.97)(s^2 + 1.83s + 6.96)}$$

satisfying $\tilde{V}_1(j\infty) = 1$, and

$$\tilde{U}_1 = \frac{0.021(s+10.92)(s+8.87)(s+7.31)(s+5.93)}{(s^2 + 8.64s + 19.97)(s^2 + 1.83s + 6.96)}$$

Further, suppose that the data collected from the closed-loop suggests the use of a new controller C_2 given by

$$C_2 = \frac{0.33(s + 0.586)(s + 2.99)(s + 3.416)}{(s + 2)(s^2 + 2.26s + 3.52)}$$

with a left coprime factorization, $C_2 = \tilde{V}_2^{-1}\tilde{U}_2$,

$$\tilde{V}_2 = \frac{(s + 2)(s^2 + 2.26s + 3.52)}{(s + 1.87)(s^2 + 2.81s + 3.712)}$$

satisfying $\tilde{V}_2(j\infty) = 1$, and

$$\tilde{U}_2 = \frac{0.33(s + 0.586)(s + 2.99)(s + 3.416)}{(s + 1.87)(s^2 + 2.81s + 3.712)}$$

Setting up the experimental configuration of Fig. 4 for simulation and utilizing Theorem 2, the simulation results in Fig. 5a show that


$$|T - 1| \leq 1 \quad \forall \omega \geq 1.27 \text{ rad/s}$$

Since $\arg T(j0) = 0$ and $\arg T(j\omega_0) = -0.285\pi$, as shown in Fig. 5b, the condition in Theorem 2 holds and hence C_2 is stabilizing. The Nyquist plot of $T(j\omega)$ is shown in Figure 5c. Indeed, computing $H(P, C_2)$ shows that C_2 is stabilizing. Obviously, but importantly, there could be substantial uncertainty in the frequency response data, and yet the same conclusion could be drawn with confidence.


5. CONCLUSIONS

In the first part of this paper, we have presented some adaptive control problems to which solutions have been found. Nevertheless, we should remember that these problems were not foreseen, and they delayed and indeed imperilled the implementation of adaptive control. It was seen as a theory being oversold by theorists, who should draw a lesson. The second part of the paper advanced a number of generic problems arising in almost any adaptive control system, and thus problems which ought to be taken into account for almost any algorithm. If one particular problem stands out, it is the task of securing safe adaptive control, or never even temporarily establishing an instability. Last, we focused on two more recent technologies, multiple model adaptive control, and model-free adaptive control. In relation to the former, we noted the importance of building in a safety constraint, and we noted the considerable difficulties in a practical extension to nonlinear problems.


Concerning model-free adaptive control, there remain a number of difficulties, which need to be theoretically argued away or established as being inconsequential through convincing experimental examples involving plants beyond laboratory scale.



(a) Magnitude Response of $T' = T - 1$



(b) Phase Response of T



(c) Nyquist plot of $T(j\omega)$

Fig. 5. Magnitude and Phase responses

ACKNOWLEDGMENT

This work was supported by ARC Discovery-Projects Grant (DP0664427) and National ICT Australia. National ICT Australia is funded by the Australian Government's Department of Communications, Information Technology, and the Arts and the Australian Research Council through Backing Australia's Ability and the ICT Research Centre of Excellence programs.

REFERENCES

- P. Albertos and A. Sala. *Iterative Identification and Control*. Springer-Verlag, London, UK, 2002.
- B. D. O. Anderson. Adaptive systems, lack of persistency of excitation and bursting phenomena. *Automatica*, 21(3):247–258, May 1985.
- B. D. O. Anderson. Windsurfing approach to iterative control design. In P. Albertos and A. Sala, editors, *Iterative Identification and Control*, pages 143–166. Springer-Verlag, London, 2002.
- B. D. O. Anderson. Failures of adaptive control theory and their resolution. *Communications in Information and Systems*, 5(1):1–20, 2005.
- B. D. O. Anderson. Exponential stability of linear equations arising in adaptive identification. *IEEE Transactions on Automatic Control*, 22(1):83–88, February 1977.
- B. D. O. Anderson and M. Gevers. Fundamental problems in adaptive control. In D. Normand-Cyrot, editor, *Perspectives in Control Theory and Application*, pages 9–21. Springer-Verlag, Berlin, 1998.
- B. D. O. Anderson and C. R. Johnson Jr. Exponential convergence of adaptive identification and control algorithms. *Automatica*, 18(1):1–13, January 1982.
- B. D. O. Anderson, R. R. Bitmead, C. R. Johnson, Jr., P. V. Kokotovic, R. L. Kosut, I. Mareels, L. Praly, and B. D. Riedle. *Stability of Adaptive Systems*. MIT Press, Cambridge, MA, 1986.
- B. D. O. Anderson, T. Brinsmead, F. De Bruyne, J. Hespanha, D. Liberzon, and A. S. Morse. Multiple model adaptive control. part 1: Finite controller coverings. *International Journal of Adaptive Control and Signal Processing*, 10(11):909–929, August 2000.
- B. D. O. Anderson, T. Brinsmead, D. Liberzon, and A. S. Morse. Multiple model adaptive control with safe switching. *International Journal of Adaptive Control and Signal Processing*, 15(5):445–470, July 2001.
- K. J. Astrom and T. Bohlin. Numerical identification of linear dynamic systems from normal operating records. In P. H. Hammond, editor, *Theory of Self-Adaptive Control Systems*, pages 96–111. Plenum Press, NY, USA, September 1966.
- K. J. Astrom and B. Wittenmark. *Adaptive Control*. Addison-Wesley Publishing Company, 1989.
- M. Athans, S. Fekri, and A. Pascoal. Issues on robust adaptive feedback control. In *Plenary paper presented by M. Athans at 16th IFAC World Congress*, Prague, Czech Republic, July 2005.
- F. B. Cabral and M. G. Safonov. A falsification perspective on model reference adaptive control. In *Proceedings of the 35th IEEE Conference on Decision and Control*, volume 3, pages 2982–2985, Kobe, Japan, 1996.
- A. Dehghani. *Robust Adaptive Control Schemes*. PhD thesis, The Australian National University, Australia, 2007.
- A. Dehghani, A. Lanzon, and B. D. O. Anderson. An \mathcal{H}_∞ algorithm for the windsurfer approach to adaptive robust control. *International Journal of Adaptive Control and Signal Processing*, 18(8):607–628, October 2004.
- A. Dehghani, A. Lanzon, and B. D. O. Anderson. A two degree of freedom \mathcal{H}_∞ control design method for robust model matching. *International Journal of Robust and Nonlinear Control*, 16(10):467–483, July 2006.
- A. Dehghani, B. D. O. Anderson, and A. Lanzon. Unfalsified adaptive control: A new controller implementation and some remarks. In *Proceedings of the 2007 European Control Conference*, Kos, Greece, July 2007a.
- A. Dehghani, B. D. O. Anderson, A. Lanzon, and A. Lecchini. Verifying stabilizing controllers via closed-loop noisy data: MIMO case. In *Submitted to the 46th IEEE Conference on Decision and Control*, New Orleans, USA, December 2007b.
- J. C. Doyle. Guaranteed margins for LQG regulators. *IEEE Transactions on Automatic Control*, 23:756–757, 1978.
- S. Fekri, M. Athans, and A. Pascoal. Issues, progress and new results in robust adaptive control. *International Journal of Adaptive Control and Signal Processing*, 20(10):519–579, December 2006.
- S. Fekri, M. Athans, and A. Pascoal. Robust multiple-model adaptive control (rmmac): A case study. *International Journal of Adaptive Control and Signal Processing*, 21(1):1–30, February 2007.
- M. Gevers. Identification for control: from the early achievements to the revival of experiment design. *European Journal of Control*, 11(4-5):335–352, 2005.
- M. Gevers. A decade of progress in iterative control design: from theory to practice. In *Symp. on Advanced Control of Chemical Processes*, volume 2, pages 677–694, Pisa, Italy, 2000.
- G. C. Goodwin and K. S. Sin. *Adaptive Filtering: Prediction and Control*. Prentice Hall, first edition, 1984.
- J. Hespanha, D. Liberzon, A. S. Morse, B. D. O. Anderson, T. Brinsmead, and F. De Bruyne. Multiple model adaptive control. part 2: switching. *International Journal of Adaptive Control and Signal Processing*, 11(5):479–496, April 2001.
- R. Kosut, B. Anderson, and I. Mareels. Stability theory for adaptive systems: Method of averaging and persistency of excitation. *IEEE Transactions on Automatic Control*, 32(1):26–34, January 1987.
- R. L. Kosut and B. D. O. Anderson. Adaptive control via finite modelling and robust control. In *Proceedings of the IFAC Workshop on Robust Adaptive Control*, pages 91–95, August 1988.
- M. Krstic, I. Kanellakopoulos, and P. Kokotovic. *Nonlinear and Adaptive Control Design*. John Wiley and Sons, 1995.
- A. Lanzon, A. Lecchini, A. Dehghani, and B. D. O. Anderson. Checking if controllers are stabilizing using closed-loop data. In *Proceedings of the 45th IEEE Conference on Decision and Control*, San Diego, USA, December 2006.
- W. S. Lee, B. D. O. Anderson, R. L. Kosut, and I. M. Y. Mareels. A new approach to adaptive robust control. *International Journal of Adaptive Control and Signal Processing*, 7:183–211, 1993.
- 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*, 31:1619–1636, 1995.
- J. M. Maciejowski. *Multivariable feedback design*. Addison-Wesley, UK, 1989.
- I. M. Y. Mareels, B. D. O. Anderson, R. R. Bitmead, M. Bodson, and S. Sastry. Revisiting the mit rule for adaptive control. In *2nd IFAC Workshop on Adaptive Systems in Control and Signal Processing*, pages 161–166, Lund, Sweden, 1986.
- R. McGowan and R. Kuc. A direct relation between a signal time series and its unwrapped phase. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 30(5):719–726, October 1982.
- A. S. Morse. Supervisory control of families of linear set-point controllers part i. exact matching. *IEEE Transactions on Automatic Control*, 41(10):1413–1431, October 1996.
- A. S. Morse. Control using logic-based switching. In A. Isidori, editor, *Trends in Control: A European Perspective*, pages 69–113. Springer-Verlag, London, July 1995.
- K. S. Narendra and A. M. Annaswamy. *Stable Adaptive Systems*. Prentice Hall, New Jersey, USA, 1988.
- K. S. Narendra and J. Balakrishnan. Adaptive control using multiple models. *IEEE Transactions on Automatic Control*, 42(2):171–187, February 1997.
- P. V. Osbourne, H. P. Whitaker, and A. Kezer. New developments in the design of model reference adaptive control systems. *Inst. Aeronautical Services*, 1961. Paper 61–39.
- C. Rohrs, L. Valavani, M. Athans, and G. Stein. Robustness of continuous-time adaptive control algorithms in the presence of unmodeled dynamics. *IEEE Transactions on Automatic Control*, 30(9):881–889, September 1985.
- M. G. Safonov and F. B. Cabral. Fitting controllers to data. *Systems and Control Letters*, 43(4):299–308, July 2001.
- M. G. Safonov and T. C. Tsao. The unfalsified control concept and learning. *IEEE Transactions on Automatic Control*, 42(6):843–847, June 1997.
- S. Sastry and M. Bodson. *Adaptive Control: Stability, Convergence, and Robustness*. Prentice Hall, 1994.
- R. J. P. Schrama. Accurate identification for control: The necessity of an iterative scheme. *IEEE Transactions on Automatic Control*, 37:991–994, 1992.
- M. Stefanovic and M. G. Safonov. Guaranteeing safety of switching adaptive control systems. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 2813–2818, San Diego, CA, USA, December 2006.
- M. Stefanovic, R. Wang, and M. G. Safonov. Stability and convergence in adaptive systems. In *Proceedings of American Control Conference*, Boston, MA, July 2004.
- M. Stefanovic, A. Paul, and M. G. Safonov. Safe adaptive switching control through infinite controller set: Stability and convergence. In *Proceedings of the 16th IFAC World Congress*, Prague, Czech Republic, July 2005.
- P. M. J. Van den Hof and R. J. P. Schrama. Identification and control: Closed-loop issues. *Automatica*, 31(12):1751–1770, December 1995.
- G. Vinnicombe. *Uncertainty and Feedback: \mathcal{H}_∞ loop-shaping and the ν -gap metric*. Imperial College Press, 2000.
- R. Wang and M. G. Safonov. Stability of unfalsified adaptive control using multiple controllers. In *Proceedings of American Control Conference*, Portland, Oregon, June 2005.
- R. Wang, A. Paul, M. Stefanovic, and M. G. Safonov. Cost-detectability and stability of adaptive control systems. In *Proceedings of the 44th IEEE Conference on Decision and Control and the 2005 European Control Conference*, Seville, Spain, December 2005.
- H. P. Whitaker. An adaptive system for control of the dynamics performances of aircraft and spacecraft. *Inst. Aeronautical Services*, 1959. Paper 59–100.
- Z. Zang, R. R. Bitmead, and M. Gevers. Iterative model refinement and control robustness enhancement. In *Proceedings of the 30th IEEE Conference on Decision and Control*, pages 279–284, Brighton, UK, December 1991.
- K. Zhou and J. C. Doyle. *Essentials of Robust Control*. Prentice-Hall, 1998.
- K. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice-Hall, 1996.