ADAPTIVE SYSTEMS IN CONTROL AND SIGNAL PROCESSING 1986

Proceedings of the 2nd IFAC Workshop
Lund, Sweden, 1–3 July 1986

Edited by

K. J. ÅSTRÖM

and

B. WITTENMARK

Lund Institute of Technology, Sweden

Published for the

INTERNATIONAL FEDERATION OF AUTOMATIC CONTROL

by

PERGAMON PRESS

OXFORD · NEW YORK · BEIJING · FRANKFURT
SÃO PAULO · SYDNEY · TOKYO · TORONTO

REVISITING THE MIT RULE FOR ADAPTIVE CONTROL

Iven M. Y. Mareels, Brian D. O. Anderson, Robert R. Bitmead, Marc Bodson* and Shankar S. Sastry*

Department of Systems Engineering, Research School of Physical Sciences, Australian National University, Canberra, A.C.T., 2600, Australia
**Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720, USA*

ABSTRACT

The MIT rule is a scalar parameter adjustment law which was proposed in 1961 for the model reference adaptive control of linear systems modeled as the cascade of a known stable plant and a single unknown gain. This adjustment law was derived by approximating a gradient descent procedure for an integral error squared performance criterion. For the early part of the 1960s this rule was the basis of many adaptive control schemes and a considerable wealth of practical experience and engineering folklore was amassed.

The MIT rule is in general not globally convergent nor stable but has a performance determined by several factors such as algorithm gain, reference input magnitude and frequency, and the particular transfer function appearing in the cascade. These restrictions on the MIT rule slowly came to be discerned through experimentation and simulation but effectively were without theoretical support until some novel algorithm modifications and stability analysis, so-called Lyapunov redesign, due to Parks. Our aim in this paper is to pursue a theoretical analysis of the original MIT rule to support the existing simulation evidence and to indicate mechanisms for treating questions of robustness of MIT-rule-based adaptive controllers with undermodelling effects.

The techniques that we apply to this problem centre on root locus methods, Nyquist methods and the application of the theory of averaging. Stability and instability results are presented and, using pertinent theories for different regimes of the gain-frequency plane, we approximate the experimentally derived stability margins, but for a broader signal class than simply periodic inputs. The mechanisms of instability and stability for these adaptive systems are highlighted and allow us to enunciate guidelines for the MIT rule to work. It is a pleasing by-product of this theoretical analysis that these guidelines coincide to a large degree with those advanced in earlier times on experimental and heuristic grounds.

1. INTRODUCTION

The MIT rule of adaptive control is a scalar parameter adjustment law which was formulated in the late fifties and early sixties as a model reference adaptive control law for linear systems modelled as a cascade of known stable plant and a single unknown gain. The names generally associated with this formulation are Whitaker, Osburn and Kezer [1,2] and the initial intended application was to the control of aircraft dynamics where the single unknown system parameter was related to dynamic pressure. The basis for this adaptive control law was an explicit performance criterion minimization carried out by an on-line gradient search.

In the history of adaptive control - or at least in its folklore - the MIT rule represents a watershed. The method was simply formulated in the model-reference framework, was easily appreciated, and was applicable. Consequently, this approach to self-optimizing systems was taken up by theorists and practitioners alike as a potential route to enhanced perfOormance. In application trials with aircraft dynamics, however, the MIT rule adaptive controller led to unpredicted instability with a considerable associated loss of face of and confidence in ad hoc adaptive control rules.

Simulation studies provided some guidelines [3] to the rule's stability properties and indicated the likely complexity of any analysis. Also, engineering guidelines or rules of thumb were developed to indicate factors affecting performance [4]. Indeed the broad principles enunciated by Donalson and Leondes [4] bear a disturbingly close resemblance to currently emerging "modern" notions of suitable operating conditions for the applicability of adaptive controllers in general, particularly with regard to averaging techniques [5-8] and time-scale separation. Theoretical tools dealing with the MIT rule have been lacking, however, and it is in this area where we are now able to reappraise and affirm the earlier results on the MIT rule by utilising essentially these latter techniques. In particular, we address the stability issues of the rule to describe the underlying stability and instability mechanisms.

The broad field of adaptive control has moved on from the criterion minimization approach primarily to a stability based rationale. This "Lyapunov redesign" of adaptive control schemes was originally proposed by Butchart and Shackcloth [9] and beautifully espoused, extended and promulgated by Parks [10]. (The Lyapunov redesign also heralded the appearance of strictly positive real conditions in adaptive control.) Our aim here is not to develop new adaptive control schemes but rather, by revisiting the well-documented MIT rule, to demonstrate the efficacy of some recently developed tools and to assess their agreement with experimental and simulation evidence.

2. THE MIT RULE

The setup under consideration is depicted in Figure 1. The plant to be controlled is modelled by a known, time-invariant, linear system with transfer function $Z_p(s)$ in cascade with an unknown scalar gain k_p of known sign, here assumed (without loss of generality) to be positive. The control objective is to adjust a feedforward control gain k_c so that the plant output $y_p(t)$ tracks the reference model output $y_m(t)$ determined by the parallel model system with transfer function $k_m Z_p(s)$. Here, k_m is a known gain, assumed positive for convenience. The bounded reference input signal $r(t)$ is the same to both the reference model and the controller plant systems. (In the original aircraft dynamics problem k_p was related to the dynamic pressure which approximately alters the aircraft dynamics in this fashion and changes with altitude and Mach number.)

The MIT rule is derived by attempting to minimize the integral squared error

$$V = \lim_{T\to\infty} \frac{1}{2T} \int_0^T e^2(t,k_c)dt \quad (2.1)$$

where $e(t,k_c)$ is the output error

$$e(t,k_c) = y_p(t) - y_m(t). \quad (2.2)$$

Whitaker proposed to adjust k_c using a gradient formula to attempt to minimize V, viz.

$$\dot{k}_c = -g \frac{\partial[\frac{1}{2}e^2(t,k_c)]}{\partial k_c} \quad (2.3)$$

$$= -g\ e[Z_p(s)\{k_p r\}] \quad (2.4)$$

This may be implemented as

$$\dot{k}_c = -g\ [y_p(t) - y_m(t)]y_m(t) \quad (2.5)$$

where the sign of g is the same as that of k_p - here assumed positive - and whose magnitude scales the adaptation speed. This is the MIT rule for adaptive feedforward control.

[We have used above a somewhat transparent mixture of notations with which we should be careful. When Z(s) is a strictly stable transfer function, Z(s){r(t)} represents the time signal generated by passing r(t) into Z(s). We ignore initial condition effects since they decay exponentially and our analysis is linear. These seeming abuses may be formalized easily but provide a convenient notation.]

Other variants on the rule (2.5) are equally possible by considering windowed criteria in place of (2.1) such as

$$V(k_c,t) = \begin{cases} \frac{1}{2T}\int_{t-T}^{t} e(k_c,\tau)^2 d\tau, \text{ rectangular window} & (2.6) \\ \frac{1}{2}\int_0^t e^{-\beta(t-\tau)}e(k_c,\tau)^2 d\tau, \text{ exponential window} & (2.7) \end{cases}$$

or by considering different controller structures. These former variants yield adaptive laws:

$$\dot{k}_c = \begin{cases} -g \int_{t-T}^{t} y_m(\tau)[y_p(\tau) - y_m(\tau)]d\tau & (2.8) \\ -g \int_0^t e^{-\beta(t-\tau)}y_m(\tau)[y_p(\tau)-y_m(\tau)]d\tau & (2.9) \end{cases}$$

The latter variants generally add complexity without new insight. We shall now move on to consider instability mechanisms for these schemes.

3. INSTABILITY MECHANISMS

Because the most remarkable feature of the MIT rule was its unpredicted instability, we begin our analysis by first investigating mechanisms to cause this behaviour. We study three main mechanisms: high adaptation gain, resonance effects and model mismatch.

3.1 Large Adaptation Gains

The gain of the MIT rule (2.5) (or (2.8), (2.9)) scales with g and the magnitude squared of the reference signal r(t). Thus this instability can arise due to either large algorithm gain or large reference inputs. To demonstrate the possibility of high gain instability we consider a simple constant input r(t) = R. This corresponds to set-point regulation.

When r(t) = R and k_p is constant the MIT rule reduces to

$$\dot{k}_c(t) = -g[Z_p(s)k_p k_c(t)R - Z_p(s)k_m R]Z_p(s)k_m R$$
$$= -gk_m k_p R^2\ Z_p(s)k_c(t) + g\ k_m{}^2 R^2 \quad (3.1)$$

where we have taken $Z_p(0) = 1$, without loss of generality. This in turn may be written as

$$k_c(s) = \frac{gk_m^2R^2}{s+gk_m k_p R^2 Z_p(s)} \quad (3.2)$$

and root locus arguments may be applied directly to establish the boundedness of $k_c(t)$. The gain parameter is $gk_m k_p R^2$ and our first global result is:

Lemma 1: The MIT rule with r(t) ≡ R has infinite gain margin (i.e. for all positive g and R, the adaptive law is stable independent of k_p) if and only if

$$\frac{-\pi}{2} < \arg Z_p(j\omega) < \frac{3\pi}{2} \quad \forall\omega\epsilon\mathbb{R} \quad (3.3)$$

Remark 3.1: As our simple structure produces a linear differential equation for $k_c(t)$, the converse of this lemma is that, unless (3.3) is satisfied, there exist adaptation gains, g, and constant reference signal levels, R, which drive the adaptive control scheme unstable. Condition (3.3) will not be satisfied by many $Z_p(s)$ containing nonminimum-phase zeros and/or having relative degree greater than one. For example, with $Z_p(s) = (s+1)^{-2}$, (3.1) is unstable for $gR^2 > 2(k_m k_p)^{-1}$. It is worthwhile to note, however, that (3.3) is satisfied by all strictly positive real $Z_p(s)$.

Remark 3.2: Whenever (3.1) is exponentially stable for r(t) = R the adaptive feedforward gain k_c becomes asymptotically optimal as desired, i.e. $k_c \to k_m k_p{}^{-1}$ as $t\to\infty$.

Remark 3.3: Notice the nonlinear manner in which the input enters the algorithm gain - proportionally to R^2. The nonlinear dependence of the algorithm's behaviour on the input is at the core of the problem in understanding the MIT rule. This becomes more pronounced in later developments.

Remark 3.4: For the alternative algorithms (2.8), (2.9) equivalent conditions can be derived although only that from (2.9) communicates much: the exponentially weighted MIT rule has infinite gain margin for constant r(t) = R if and only if $Z_p(s-\beta)$ satisfies (3.3). This modification [and also the modification appropriate for (2.8)] only serve to make stability more difficult, so that the algorithm may not perform adequately even with strictly positive real $Z_p(s)$.

3.2 Resonance Effects

By considering periodic input signals r(t), as opposed to constant inputs in the previous subsection, another broader class of instabilities is displayed. We shall proceed by using an example.

Let $k_p = k_m = 1$ and $Z_p(s) = 1/(s+1)$, which incidentally is SPR. This avoids the high gain instability as $Z_p(s)$ satisfies (3.3). Now we take r(t) = cosωt and investigate the effect of altering ω. The MIT rule (2.5) may then be written in terms of x_1, the state of the plant $Z_p(s)$, and $x_2(t) = k_c(t)-1$, the parameter error, as

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} -1 & \cos\omega t \\ -g\ y_m(t) & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \quad (3.4)$$

where

$$y_m(t) = \mathrm{Re}[Z_p(j\omega)e^{j\omega t}] = (\omega^2+1)^{-\frac{1}{2}} \cos(\omega t - \tan^{-1}(\omega)).$$

Equation (3.4) is a linear ordinary differential equation with periodic coefficients - note the similarity with the classical Mathieu equation. Its stabilty properties may be studied using Floquet theory as was done for this adaptive control problem by James using numerical integration methods [3]. The results displayed in Figure 2 depict the stability domain in the frequency (ω)-gain (g) parameter plane and exhibit the extreme complexity of the stability/instability boundary characteristic of this class of equations - again recall the Mathieu equation and the extraordinary difficulty of describing analytically its stability properties. The gain margin for this $Z_p(s)$, which is infinite when ω=0, is drastically reduced around the cut-off frequency (ω=1) of the plant.

More complicated $Z_p(s)$ demonstrate this and more complicated behaviour. Similarly, introduction of forgetting factors and integration produces more complex analysis, and produces more complex behaviour. Replacing $Z_p(s)$ by $p/(s+p)$ is equivalent to scaling ω or g respectively and so does not alter affairs. Again, further examples of resonance phenomena or "pumping up" can be derived with different periodic $r(t)$.

3.3 Modelling Errors

Under normal circumstances, the plant's transfer function $Z_p(s)$ is not exactly known to the designer and only an approximate representation $Z_m(s)$ is available. This approximation must be used to generate the desired reference trajectory $y_m(t)$. These modelling errors alter the previous analysis and, in particular, can drastically affect control performance. In the previous example, if $y_m(t) = [e^{-s}/(s+1)]r(t)$ and $Z_p(s)$ remains as before, then the resulting stability/instability boundary is as in Figure 3. Notice in particular that the algorithm becomes unstable for $\omega \approx \pi/2$ and that the stability margin is further reduced.

These three instability mechanisms demonstrate seemingly disparate phenomena which do not augur well for the adequate performance of the MIT rule in almost any situation. However, a rudimentary analysis of these examples indicates that operation with small adaptation gain g may be necesseary for good performance. Small g corresponds to slow adaptation which will be specifically pursued in greater detail to rescue the MIT rule.

4. STABILITY ANALYSIS VIA AVERAGING: RESCUING THE MIT RULE

Abiding by the warnings of the instability mechanisms of the previous seciton, we shall seek now to consider the case of time-scale separation between the plant and the adaptation. This restriction allows us to use averaging and/or singular perturbation techniques to obtain some intuitively appealing sufficient requirements for good performance.

Throughout this section we assume that the parallel model transfer function $Z_m(s)$ is only an approximation of the plant transfer function $Z_p(s)$, and we shall derive both stability and instability results where possible. We consider two different types of timescale separation where alternatively (i) the adaptation is slow relative to plant and reference signals and (ii) the reference input is slow relative to the plant and adaptation.

4.1 Slow Adaptation

Writing

$$y_m(t) = k_m Z_m(s)\{r(t)\} \tag{4.1}$$

$$y_p(t) = Z_p(s)\{k_p k_c(t) r(t)\} \tag{4.2}$$

the MIT rule (2.5) is

$$\dot{k}_c = -g[Z_m(s)\{k_m r(t)\}][Z_p(s)\{k_p k_c r(t)\}] + g[Z_m(s)\{k_m r(t)\}]^2 \tag{4.3}$$

Assuming that g is small, i.e. k_c is slowly time-varying, it is reasonable to approximate (4.3) by formally treating k_c as a constant in the right hand side, i.e., with k_c^* approximating k_c,

$$\dot{k}_c^* = -g[Z_m(s)\{k_m r(t)\}][Z_p(s)\{k_p r(t)\}]k_c^*(t) + g[Z_m(s)\{k_m r(t)\}]^2 \tag{4.4}$$

For sufficiently small g (4.4) and (4.3) will have similar stability properties. In particular, exponential stability or instability of (4.4) for sufficiently small g will imply the same for (4.3), see [6,8].

Notice now that (4.4) is a linear time-varying first-order differential equation whose stability properties are assessed quite readily. We have

Lemma 2: The homogeneous part of (4.4) is exponentially stable for bounded $r(t)$ and stable $Z_m(s)$ and $Z_p(s)$ if and only if:

$$\liminf_{T\to\infty} \frac{1}{T}\int_0^T [Z_m(s)\{k_m r(t)\}][\,Z_p(s)\{k_p r(t)\}]dt > 0 \tag{4.5}$$

Remark 4.1: If the homogeneous part of (4.4) is exponentially stable, $r(t)$ is bounded, and $Z_m(s)$ is stable, then $k_c^*(t)$ converges to a bounded limiting function $k_c^\infty(t)$ independent of initial conditions but not necessarily constant.

Remark 4.2: The condition (4.5) is a variant on the usual persistence of excitation conditions - both $Z_m(s)$ and $Z_p(s)$ are now involved. For stable Z_m, Z_p we still require $r(t)$ to be persistently exciting but (4.5) embodies an additional requirment that the energy in $r(t)$ be localised where Z_m and Z_p have similar frequency responses, or at least phase responses.

We may use averaging theory to transfer this result directly to (4.3).

Theorem 1: Under the condition that $Z_m(s)$ and $Z_p(s)$ are strictly stable, that $r(t)$ is bounded and that (4.5) is satisfied, there exists a positive constant g^* such that for all $g\varepsilon(0,g^*)$ the gain k_c adjusted by the MIT rule (2.5) is bounded and converges exponentially fast to $k_c^*(t) + O(g)$ as $t\to\infty$.

Remark 4.3: The constant g^* above may be quantified in terms of α in (4.5) and further characteristics of $r(s)$. Averaging theory (or, alternatively, singular perturbation theory) permits us only to look with surety up to g^* and gives us no further information about the properties of (4.3) in terms of (4.4). Thus this boundary may be quite conservative.

Remark 4.4: Subject to small g we have an equivalent instability result that, under conditions of strictly stable $Z_m(s)$, $Z_p(s)$ and bounded $r(t)$ with the integral on the left of (4.5) having strictly negative limit superior, there exists g^* such that $g\varepsilon(0,g^*)$ yields k_c unbounded.

Remark 4.5: The time invariance of k_p has not been invoked up to this point and, given the usual rationale of adaptive systems of adjustment to slowly-varying parameter values, one can envisage seeking to admit k_p variation, provided at least that k_p does not change sign. Using the same averaging principles in allowing k_p to varying more slowly than the adaptation speed, it is possible to split the timescales into three distinct components and (4.5) becomes:

$$\liminf_{T\to\infty} \frac{1}{T}\int_0^T Z_m(s)\{r(t)\}Z_p(s)\{r(t)\}dt \geq \alpha' > 0 \tag{4.6}$$

We now ask under what conditions (4.6) is satisfied and for clarity consider this condition for inputs of the kind

$$r(t) = \sum_{i=-N}^{N} \alpha_i e^{j\omega_i t}, \quad \alpha_i = \alpha^*_{-i}, \quad \omega_i = -\omega_{-i} \tag{4.7}$$

i.e. almost periodic inputs. With such r(t) (4.6) becomes

$$|\alpha_0|^2 + 2\sum_{i=1}^{N} |\alpha_i|^2 \mathrm{Re}[Z_m^*(j\omega_i)Z_p(j\omega_i)] \tag{4.8}$$

This condition is the direct analogue of that of Riedle and Kokotovic [7] for more familiar adaptive schemes and indicates that the MIT rule will be stable for small g for sufficiently close Z_m and Z_p provided r(t) has frequency content concentrated about values of similar response of Z_m and Z_p. This is summarized in:

Corollary 1: The MIT rule for feedforward control (2.5) for almost periodic input r(t) (given by (4.7)) and strictly stable $Z_m(s)$, $Z_p(s)$ is stable for sufficiently small g if (4.8) is satisfied.

For arbitrary almost periodic r(t) the MIT rule is stable for sufficiently small g if

$$\sup_{\omega \in R} |\arg Z_m(j\omega) - \arg Z_p(\omega)| < \frac{\pi}{2} \tag{4.9}$$

Clearly (4.8) is a weaker condition that (4.9) in that (4.9) implies (4.8). This relaxation of the phase condition to require restrictions on the input at the allowance of violations of (4.9) is equivalent to the relaxation of SPR conditions in adaptive systems [7,8].

Instability results for small g may be simply derived when the quantity on the left of (4.8) is negative. Thus for small adaptation gain and almost periodic inputs, these results delineate a sharp stability-instability boundary. The instability here is due to model errors since the small adaptation gain precludes the high gain and resonance phenomena.

As far as control performance is concerned, it is relatively straightfoward to demonstrate that the integral squared error criterion is no longer optimized with the control law (4.3), even in the limit, unless the model errors are zero on the spectral lines of r(t).

The analysis so far has concentrated on timescale separation involving slow adaptation compared with plant and input signals. These averaging theory results predict the simulation results in Figure 2 denoted by region I indicating stable behaviour for all ω and small g, and predict instability for $\pi/2 \le \omega \le 3\pi/2$ and small gain as depicted in Figure 3. We now move on to a different timescale separation admitting results in a different region of ω,g-space.

4.2 Slowly Time-varying Inputs

Starting from the MIT rule description (2.5) and now assuming that r(t) and k_p are slowly time varying we may appeal once again to averaging and attempt to obtain the extension of the root locus arguments used in Section 3 with constant r(t). Just as (4.4) approximates (4.3) with g small, we can find the following equation to approximate (4.3) with r(t) and k_p slowly varying.

$$\dot{\bar{k}}_c = -gk_mk_pr^2(t)[Z_p(s)\{\bar{k}_c(t)\}] + gk_m{}^2r^2(t) \tag{4.10}$$

where we have again taken $Z_m(0) = Z_p(0) = 1$. In justifying the approximation, in addition to $d/dt(k_pr)$ being small, we require $Z_m(s)$ and $Z_p(s)$ strictly stable. We do not require that g be small. Then $k_c(t)$ and $\bar{k}_c(t)$ will be close as $t\to\infty$ in a way to be made precise shortly.

The stability properties of (4.10) may be derived simply as an extension of the root-locus analysis of subsection 3.1. We have

Lemma 3: Provided the zeros of

$$s + gk_mR^2Z_p(s) \tag{4.11}$$

have real parts less than a negative constant, $-\sigma$, for some positive g and k_m, and for all $R_0^2 \le R^2 \le R_1^2$ where R_0 and R_1 are fixed positive quantities, there exists a positive constant $\rho(\sigma)$ such that the homogeneous part of (4.10) is exponentially stable for all inputs r(t) and gains k_p satisfying

$$R_0^2 \le k_pr^2 \le R_1^2 \tag{4.12}$$

$$\int_s^{s+1} \left|\frac{d}{dt}[k_pr(t)]\right| dt \le \rho, \quad \forall\, s \ge 0. \tag{4.13}$$

Remark 4.5: Notice that (4.12) requires r(t) to be non zero and that the range of r(t) depends on the values taken by k_p. This result represents a slight modification of [11,pp125-127] in that (4.13) is an integral bound on the derivative magnitude which can allow jumps as opposed to a strict overbound. This result follows directly from the proof in [11] using integration by parts before appealing to Gronwall's Lemma.

We may now apply the results of Lemma 3 to ascertain the behaviour of the MIT rule with slowly-varying inputs.

Theorem 2: Under the conditions of Lemma 3, and the condition that $Z_m(s)$ and $Z_p(s)$ are strictly stable, there exists a positive ρ^* such that for $k_pr(t)$ satisfying (4.12) and (4.13) for $\rho \in (0,\rho^*)$, the MIT rule (2.5) is stable. Moreover,

$$k_c(t) \to \bar{k}_c(t) + O(\rho) \text{ as } t\to\infty$$

exponentially fast.

Remark 4.6: Theorem 3 explains the stability properties of the MIT rule for slowly time-varying inputs. In particular, it predicts the stability in regions II of Figures 2 and 3. It does not specifically address the problem of model errors since we have good low frequency matching. (Recall that $Z_p(0) = Z_s(0) = 1$.) This result may be seen as providing sufficient conditions to avoid high gain instability for very low frequency inputs. Lemma 1 arises as a special case.

Remark 4.7: Under the conditions of Theorem 2, the MIT rule is not only stable but close to optimal in the integral squared error sense, since $\bar{k}_c(t)\to k_m/k_p$ as $t\to\infty$. The results of this subsection complement those of the previous subsections in that they deal with a different design circumstance and cover a different part of the stability regimes depicted in Figures 2 and 3.

5. CONSEQUENCES AND CONCLUSIONS

The MIT rule has been presented in its simplest manifestation of scalar feedforward gain adjustment and we have pieced together a collection of stability and instability results and examples from which it behoves us to draw some more generic conclusions.

The guiding principles arising from our analysis of instability mechanisms and stability theorems are that (a) the timescales of our adaptive system need to be sufficiently distinct and (b) the model and plant transfer functions need to be well matched over the frequency range of dominant power in the input. These conditions have also been recently espoused as good engineering sense in [8] with regard to a broad class of adaptive algorithms, and were used to establish local robustness of possibly nonlinear schemes. There is a good historical perspective to be gained by comparing these ideas based on recent application of sound theoretical tools to engineering principles voiced in [4], of 1963, where timescale separation is routinely invoked to allow a heuristic analysis to proceed as well as to accommodate engineering design intuitions. To a very large extent our results are a reaffirmation of these fundamental principles and concepts of adaptation in general. The most realistic scenario is that of slow adaptation where the stability objectives for the MIT rule are achieved, per Theorem 1, by taking small adaptation gain g, and insisting that r(t), $Z_m(s)$, $Z_p(s)$ jointly satisfy (4.5). This latter condition, in turn, requires that r(t) have a persistence property and that its dominant energies should be located at frequency bands where Z_m is close to Z_p. Indeed, these may very well be envisaged as conditions necessary for the well-posedness of an adaptive solution using the MIT rule.

The case of slowly time-varying inputs is also of interest in that it highlights how these theoretical tools may be applied to gradually validate more of the stability region generated using numerical methods and Floquet thoery. It helps us to isolate the generic aspects of the problem which may then be incorporated into engineering design guidelines.

Although our aim here has primarily been to use the MIT rule as a showcase for these methods, we should reiterate that more recently developed algorithms are amenable to these analytical tools as well. A desire to avoid the instability mechanisms of the MIT rule was behind the development of these new algorithms as is discussed by Parks [10,12]. The primary modification there is to replace criterion minimization by Lyapunov function specification and the effect on the adaptive law is, typically, to replace (2.5) by

$$k_c = -g\, r(t)(y_p(t) - y_m(t))$$

and to utilize SPR $Z_p(s)$ to allow generation of a quadratic Lyapunov function.

We have revisited the MIT rule for adaptive feedforward control and displayed some of its instability mechamisms and produced hard stability results relying on time-scale separation and good model selection. In some ways this has demonstrated the reasons for the loss of confidence in this rule and also suggested potential remedies necessary to resuscitate it.

ACKNOWLEDGEMENTS

The authors are grateful to Professor Petar Kokotovic for suggesting this endeavour of applying averaging theory to revisit the MIT rule.

REFERENCES

[1] H.P. Whitaker, "An adaptive system for control of the dynamics performance of aircraft anhd spacecraft", Inst. Aeronautical Sciences, Paper 59-100, 1959.

[2] P.V. Osburn, H.P. Whitaker and A. Kezer, "New developments in the design of model reference adaptive control systems", Inst. Aeronautical Sciences, Paper 61-39, 1961.

[3] D.J.G. James, "Stability of a model reference control system", AIAA J., Vol.9, pp950-952, 1971.

[4] D.D. Donalson and C.T. Leondes, "A model referenced parameter tracking technique for adaptive control systems", IEEE Trans. Appl. Ind., pp241-262, 1963.

[5] R.L. Kosut, B.D.O. Anderson and I.M.Y. Mareels, "Stability theory for adaptive systems: methods of averaging and persistence of excitation", IEEE Trans. Auto. Control, to appear.

[6] L. Fu, M. Bodson and S. Sastry, "New stability theorems for averaging and their application to the convergence analysis of adaptive identification and control schemes", Proc. 24th IEEE Conf. on Decision and Control, Fort Lauderdale, pp473-477, 1985.

[7] B.D. Riedle and P.V. Kokotovic, "A stablity/ instability boundary for disturbance-free slow adaptation and unmodelled dynamics", Proc. 23rd IEEE Conf. on Decision and Control, Las Vegas, NV, pp998-1001, 1984.

[8] B.D.O. Anderson, R.R. Bitmead, C.R. Johnson Jr, P.V. Kokotovic, R.L. Kosut, I.M.Y. Mareels, L. Praly and B.D. Riedle, Stability of Adaptive Systems: Passivity and Averaging Analysis, MIT Press, Boston, MA, 1986.

[9] R.L. Butchart and B. Shackloth, "Synthesis of model reference adaptive control systems by Lyapunov's second method", Proc. IFAC Symp. on Adaptive Control, Teddington, UK, 1965.

[10] P.C. Parks, "Lyapunov redesign of model reference adaptive control systems", IEEE Trans. Auto. Control, Vol.AC-11, pp362-367, 1966.

[11] C.A. Desoer and M. Vidyasagar, Feedback Systems: Input/Output Properties, Academic Press, New York, NY, 1975.

[12] C.C. Hang and P.C. Parks, "Comparative studies of model reference adaptive control systems", IEEE Trans. Auto. Control, Vol.AC-18, pp419-428, 1973.

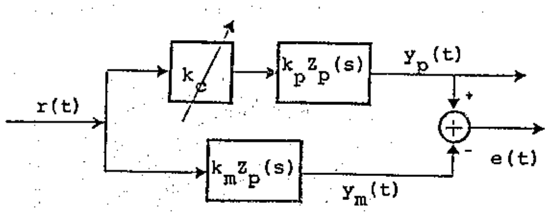

Figure 1: The MIT rule's problem

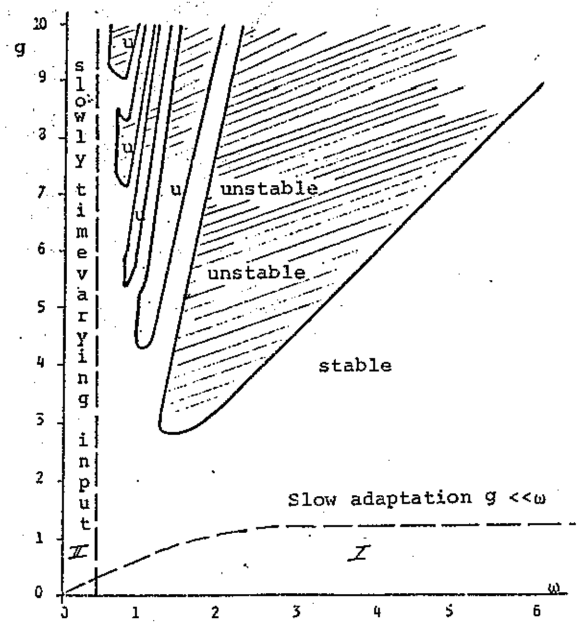

Figure 2. Stability/Instability boundary

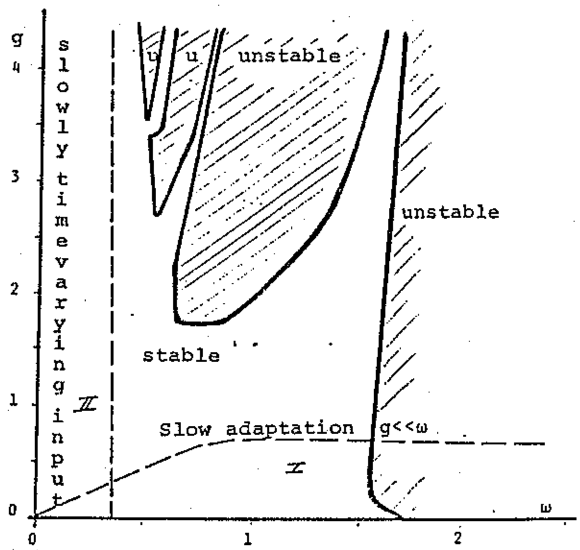

Figure 3. Stability/Instability boundary