WORKSHOP ON APPLICATIONS OF ADAPTIVE SYSTEMS THEORY

June 15-17, 1983
Yale University
New Haven, Connecticut

ROBUSTNESS OF REDUCED-ORDER ADAPTIVE MODEL FOLLOWING

Robert L. Kosut
Integrated Systems, Inc.
151 University Avenue
Palo Alto, CA 94301

C. Richard Johnson, Jr.
Cornell University
School of Electrical Engineering
Ithaca, NY 14853

Brian D. O. Anderson
Australian National University
Camberra, Australia

ABSTRACT

This paper examines some issues in available stability theory regarding reduced-order adaptive model following. Given that no parametrization exists such that the tracking error is identically zero for all admissible reference trajectories, then the objective of convergence to a point in the parameter space is replaced by a set convergence, i.e., local stability. It is shown that local stability is dependent on the stability of a linear time-varying system obtained by linearization of the nonlinear adaptive system about a tuned trajectory in the convergent set. Conditions are presented that insure stability of the time-varying system.

INTRODUCTION

Adaptive control is termed reduced-order when the controller structure being adaptively parameterized is of inadequate dimension to allow exact satisfaction of the desired control objective. This issue is of undeniable practical importance because no actual plant is truly linear and finite-dimensional. The order of the plant model, used for controller order specification, will always be insufficient to exactly describe the plant behavior with any parameterization; however, present adaptive control theory requires that the model order allow complete plant description or equivalently that the controller order is adequate to achieve the desired dynamic control objective [1],[2]. In a deterministic setting for adaptive model-following, this implies the ability to zero the model-following error for all possible reference signals. Simulations have shown that when this sufficient order condition is violated, adaptive controllers with arbitrary initializations can become unstable, even when a controller parameterization exists that would closely approximate the desired behavior [3]. These simulate circumstances of undesirable behavior are in sharp contrast with successful applications of adaptive control where reduced-order modeling is unavoidable. A principal aim of reduced-order adaptive control research is to close this gap between theory and practice by specifying conditions under which adaptive controllers retain stability despite reduced-order use.

The purpose of this paper is to discuss some recent theoretical results [4]-[6] concerning the stability of reduced-order adaptive control. The next section presents a heuristic discussion of the source and meaning of these results with respect to the objective of adaptive model-following. The latter sections present the formal statement of these results. The conclusion cites open theoretical issues for continued study in reduced-order adaptive control.

DISCUSSION OF THEORETICAL RESULTS

In adaptive model-following, the model-following (or tracking) error is used as a measure of the success of the adaptive controller parameterization. If this tracking error is nonzero, it (or a filtered version of it) is commonly used as a multiplicative component of the correction term adjusting the controller parameter. The other principal factors of this correction term are a step-size weighting and the information in the controller multiplying the controller parameters to achieve the current control action. The latter, the information vector, is presumed to take on almost any value. The former should be bounded away from zero to allow reaction to tracking errors due to plant parameter variations. With reduced-order application, the premise is that no controller parameterization exists that will cause this tracking error to be identically zero for all possible reference trajectories forcing the model and thereby generating model outputs to be tracked; therefore, if the step-size is bounded away from zero and the information vector is sufficiently rich, the correction term is always nonzero. Thus one should not demand convergence of the reduced-order adaptive controller parameterization to a point in the parameter space; however, such convergence to a point has been the principal theoretical thrust of adaptive control studies. If the step-size is small, still allowing some tracking of plant parameter time-variation, and the reduced-order controller can be tuned to provide small tracking error, this 'jitter'' in the controller parameter adaption about the ''tuned' parameterization can be small. A reasonable question to then ask of adaptive control is whether or not this type of behavior can be theoretically verified. Applications indicate that the answer is yes. But the theoretician's concern is whether or not this local stability can only be obtained with the addition of protective mechanisms added in practice but presently absent from the adaptive controller being investigated analytically.

To answer this local stability question analytically, we shall invoke two concepts. The first is exponential stability in ideal, full-order usage [4]. The second is the idea of initialization near a tuned parameterization [5],[6]. These two concepts can be combined in the following manner. The system equations describing the propagation of the combined controller parameter error and plant output error form a complex, homogeneous, nonlinear, time-varying system that can be proven to be exponentially stable in ideal use given persistent excitation by the external reference signal. Form nonlinear stability theory [7], such a homogeneous, i.e., unforced, error system will retain stability, i.e., a bounded state, given a satisfactorily small input. In this case the state of the error system is the parameter and output tracking errors with respect to a tuned controller. In other words, the model for comparison is not the actual desired model but the best achievable approximation to this model. Thus the adaptive controller could generate zero output tracking and parameter errors with respect to this tuned model. The assertion is that the difference between this tuned model output and the desired model output is small. It is this tuned error that now enters the formerly homogeneous error system. Exponential convergence in the homogeneous case implies a bounded-input, bounded-state property for this forced error system, i.e., local stability.

In addition to persistent excitation, a strictly positive real (SPR) condition arises. Common to full-order adaptive model-following theory is the requirement that the desired model times the filtering of the tracking error used in the adaptive algorithm must be SPR. This is needed to use the error system structure shown in Fig. 1.

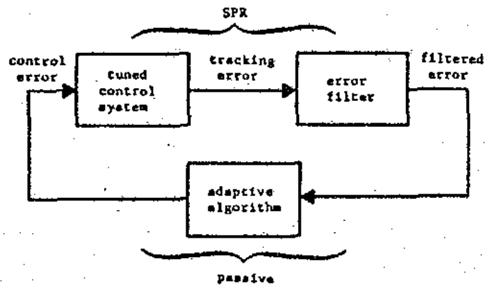

Figure 1. Unforced Error System

If the forward path is SPR and the feedback path is passive--which is true for the typical parameter estimator correction form—then this system is globally asymptotically stable. Thus its output--the filtered prediction error--converges to zero, the state of the forward path--the tracking error--converges to zero, the output of the feedback path--the control input error--converges to zero, and the state vector of the feedback path--the controller parameter error--remains bounded. If, in addition, the information vector (which is the output vector of the state description of the passive feedback path) is persistently exciting to avoid orthogonality with the adaptive algorithm parameter estimate error state vector, then the parameter estimate error converges to zero for the control input error to converge to zero. Furthermore, all of these convergence occur exponentially fast.

Given reduced-order application, the control input error now also includes a term equal to the difference between the tuned error and the control input that would cause the plant to behave like the desired model. This 'input' to Fig. 1 can be translated to the output of the intermediate point in the forward path of Fig. 1 by considering the tracking error between the desired model output and the tuned control system output. This signal was hypothesized to be small by assuming that an accurate, though not exact, reduced-order tuning exists. Filtering this difference signal by the error filtering block in Fig. 1 generates Fig. 2.

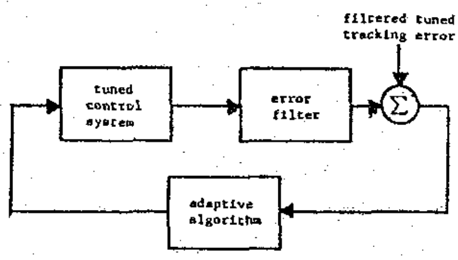

Figure 2. Forced Error System

Fig. 2 is in the form investigated in [4]. In [4] it was shown that the smaller is the perturbing filtered tuned tracking error the larger is the region in which the adaptive algorithm can be initialized and stability retained. Conversely, if the ''best' reduced-order controller tuning results in a larger tracking error with respect to the desired model, then the adaptive controller must be initialized much closer to this best tuned parameterization. Admittedly, such a result does not have a hugely practical applicability, since, if the tuned parameterization were sufficiently well-known, adaptive control might not be used; however, the practical strength of adaptive control in real applications can be seen as, based on an initially adequate controller tuning providing the fine-tuning accommodating modest plant parameter drift. In fact, this is the role of adaptive control in many existing applications. A weakness of the outlined theory is that the transfer function required to be SPR drifts as the plant drifts. Thus the fixed error filtering guaranteeing SPR satisfaction for the initial tuning may become inadequate with time. This suggests, as hinted earlier, that additional algorithm refinements are needed to expand the practical applicability of present, theoretically studied adaptive controllers. The practicality of these theoretical results are further diluted by the present conservatism in our ability to establish how small is small enough with regard to the tolerable tracking error of the tuned control. On a more positive note, however, such results substantiate that existing adaptive control schemes do possess some degree of robustness in reduced-order use.

The preceding formulation does suggest asking different questions that have some practical significance. For example, in many process control applications, the reference signal forcing the model is not persistently exciting. A rough definition of persistent excitation is that the reference signal needs as many distinct sinusoidal components as the reduced order assumed for the plant-controller combination. Step reference inputs are common in process control. A step can be viewed as a single, zero-frequency sinusoid. Clearly any controller that stabilized the plant and provides the correct DC gain for the control system will asymptotically track the desired model step response. In this case, the tuned tracking error converges to zero. In [5] it was proven that if this tuned tracking error, and thus the error system perturbation in Fig. 2, converges to zero, then any reduced-order adaptive controller initialization will asymptotically achieve this zero tracking error with bounds on the adaptive control system behavior in the interim. Unfortunately, these bounds can be so large as to be impractical. Also, the problem of SPR satisfaction remains. Another result from[5] expanding the adaptive controller initialization range is to switch to an algorithm with a contracting integration once the tracking error is excessive. Though global stability is retained, practicality is severly limited. Such a scheme presupposes an ability to prespecify a range for the controller parameterization that will lead to a stable fixed control system; furthermore, this must be related to the allowable tracking error range. Again the adaptive control system behavior, though bounded, may exceed a reasonable operating range.

By focusing more on local stability properties rather than a expansion of the initialization range, the more restrictive SPR condition can be eliminated. Practically, it is more reasonable to expect that the controller may be adequately initialized (or re-initialized) that it is to expect the seemingly arbitrary SPR condition to be satisfied. Heuristically, if the adaptive algorithm correction term remains small then the adaptive algorithm should remain near this proper initialization. Formalizing this trivial idea is difficult, thought a first attempt has been provided in [6]. Again, the needed theoretical assumptions are lacking in their practicality.

Clearly more work is needed. The next section formalizes some of the heuristic claims of this section. The last section considers directions for extending these results to enhance their practical reasonableness and narrow the theory-practice gap.

3. Adaptive Model Error

In this section we present an adaptive error model that will be used in the subsequent analysis. To facilitate the development of the error model, consider the simple model reference adaptive controller (MRAC) depicted in Fig. 3 with:

Uncertain Plant

$$y = d + Pu \qquad\qquad (3.1a)$$
$$d := \text{external disturbance}$$

Reference Model

$$y_r = Rr \qquad\qquad (3.1b)$$
$$r := \text{reference command}$$

Adaptive Control

$$u = -(\hat{\theta}_1,\ \hat{\theta}_2)\begin{pmatrix} y \\ -r \end{pmatrix} := -\hat{\theta}' z \qquad\qquad (3.1c)$$
$$\hat{\theta} := \text{adaptive gains}$$

Adaptation Law

$$D\hat{\theta} = Bze,\ B = B' > o,$$
$$(D\hat{\theta})(t) := \begin{cases} \dot{\hat{\theta}}(t), & \text{continuous} \\ \hat{\theta}_t - \hat{\theta}_{t-1}, & \text{discrete} \end{cases} \qquad\qquad (3.1d)$$
$$e := y - y_r$$

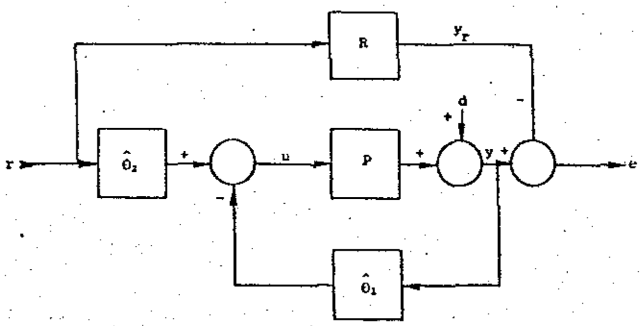

Figure 3. MRAC System

Define the <u>adaptive gain error</u> by

$$\tilde{\theta} := \hat{\theta} - \theta_* \tag{3.2}$$

where $\theta_* = (\theta_{*1}, \theta_{*2})'$ is a constant vector of <u>tuned</u> gains; i.e., the values that would be selected if the plant P were known. Using (3.2) we can rewrite (3.1c) as

$$u = -\theta_*' z - v \tag{3.3}$$

where v is the <u>adaptive control error</u> signal. An equivalent representation of the model following adaptive system (Fig. 3) is given by the adaptive error model Fig. 4 described by:

$$e = e_* - H_{ev} v \tag{3.4a}$$

$$z = z_* - H_{zv} v \tag{3.4b}$$

$$v = z'\tilde{\theta} \tag{3.4c}$$

$$\tilde{\theta} = \tilde{\theta}_0 + Lze \tag{3.4d}$$

where (e_*, z_*) are the outputs of the <u>tuned system</u>, as shown in Figure 5; $\tilde{\theta}_0$ is the initial value of the adaptive gain error, and H_{ev}, H_{zv}, L are subsystem <u>interconnection</u> operators. This error model (3.4) is very general and applies in structure to practically all adaptive controllers and filters [5]. In general (H_{ev}, H_{zv}, L) can be nonlinear, time-varying, and infinite-dimensional. In the simple case here (Fig. 3.1), they are linear-time-invariant (LTI) with rational proper trhnsfer functions. In particular, the tuned signals are:

$$e_* = (1+P\theta_{*1})^{-1} d + [(1+P\theta_{*1})^{-1} P\theta_{*2} - R] r \tag{3.5a}$$

$$z_* = (e_* + r, -r)' \tag{3.5b}$$

and the interconnections are:

$$H_{ev} = (1 + P\theta_{*1})^{-1} P \tag{3.6a}$$

$$H_{zv} = (\, (1 + P\theta_{*1})^{-1} P,\ 0\,)' \tag{3.6b}$$

$$L = \begin{cases} (1/s)B, & \text{continuous} \\ (z/(z-1))B, & \text{discrete} \end{cases} \tag{3.6c}$$

Proofs of global stability utilize the fact that the map $e \mapsto v$ is passive (see e.g., [1], [2]). Thus, if H_{ev} is strictly positive real (SPR), then the Passivity Theorem (see e.g. [7], p. 182) guarantees L_2-stability of the map $e_* \mapsto (v, e)$, even though $\tilde{\theta} \in L_{\infty e}$. Boundedness of e_* and z_*, together with stability of H_{zv} ensure that $\tilde{\theta}$ is bounded. Since it is virtually impossible to maintain $H_{ev} \in$ SPR despite unmodeled dynamics [3]-[5], it is more meaningful to develop conditions for local stability that are not dependent on the SPR condition.

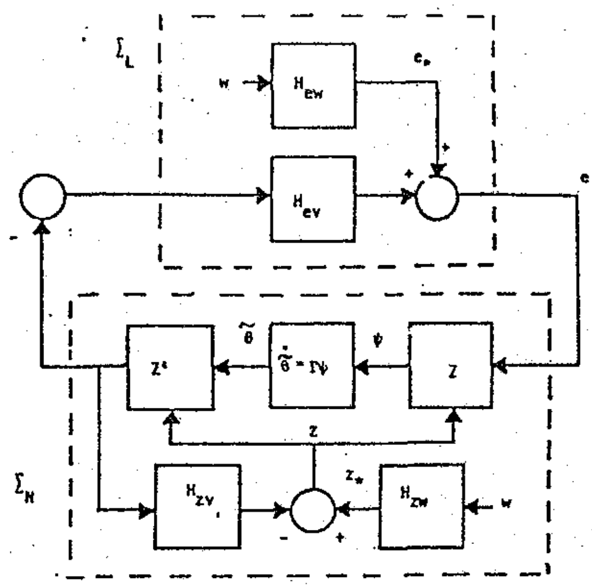

Figure 4. Adaptive Error System

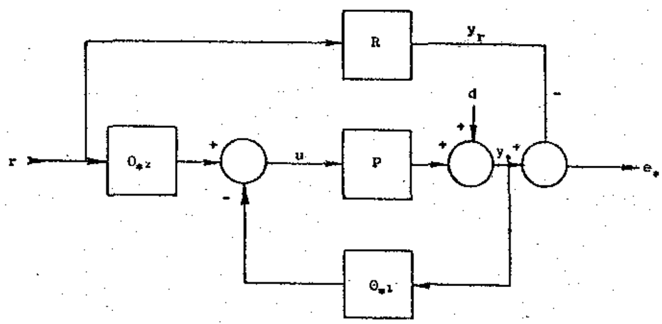

Figure 5. Tuned System

Local Error Model

The error model (3.4) can be transformed to a more useful form for local stability analysis; i.e.,

$$x = \bar{x} - Gf(x) \tag{3.7}$$

where the quantities above are defined below by

$$x := \begin{bmatrix} \tilde{e} \\ \tilde{z} \\ \tilde{\theta} \end{bmatrix} := \begin{bmatrix} e - e_* \\ z - z_* \\ \theta - \theta_* \end{bmatrix}, \quad f(x) := \begin{bmatrix} \tilde{z}'\,\tilde{\theta} \\ \tilde{z}\ \tilde{e} \end{bmatrix} \tag{3.7a}$$

$$\bar{x} := \begin{bmatrix} \bar{e} \\ \bar{z} \\ \bar{\theta} \end{bmatrix} := \begin{bmatrix} -H_{ev} z_*' \tilde{\theta} \\ -H_{zv} z_*' \tilde{\theta} \\ (I + LM)^{-1} \tilde{\theta}_0 + K z_* e_* \end{bmatrix} \tag{3.7b}$$

$$G := \begin{bmatrix} H_{ev}[1 - z_*'K(z_*H_{ev} + e_*H_{zv})] & H_{ev}z_*'K \\ H_{zv}[1 - z_*'K(z_*H_{ev} + e_*H_{zv})] & H_{zv}z_*'K \\ K(z_*H_{ev} + e_*H_{zv}) & -K \end{bmatrix} \quad (3.7c)$$

$$K := (I + LM)^{-1}L,\ M := z_*H_{ev}z_*' + e_*H_{zv}z_*' \quad (3.7d)$$

The model (3.7) is arrived at by separating the nonlinear cross product terms in $f(x)$ from the linear terms in x . We shall refer to x as the response of the linearized system (3.7b). This is almost identical to the linearized system studied by Rohrs et al. [3], which was arrived at by a 'final approach analysis.' Note that in this case the linearized system is the input to the nonlinear system (3.7). Also, the operators K and G are time-varying due to their dependence on the tuned signals (e_*, z_*).

4. Conditions for Local Stability

If $\bar{x}$ in (3.7b) is small, and the nonlinear function $Gf(x)$ is suitably restricted, then intuitively, x would remain in the neighborhood of $\bar{x}$. However, $\bar{x}$ is small if $\bar{\theta}$ is small, i.e., if the linearized response of the adaptive parameters remain in the neighborhood of the tuned parameters θ_*. Thus, it is desirable that the linearized system is stable. It turns out]6] that conditions which ensure stabilty of the linearized system are the same as those required to ensure the stability of the map G. Specifically, if the feedback system shown in Fig. 6 is internally stable, i.e., the system $S: (u_1, u_2) \mid\!\!\longrightarrow (e_1, e_2, y_1, y_2)$ is stable, then the linearized system and the map G are also stable, provided that e_* and z_* are bounded, and that H_{ev} and H_{zv} are stable. Those latter provisions are guaranteed if the tuned system is stable, e.g., see equations (3.5) - (3.6).

The fundamental question is whether S is stable without H_{ev} being SPR? The answer is in the affirmative, particularly if z_* is persistently exciting—which is not the only mechanism to guarantee stability of S [6]. An interesting case is when H_{ev} is 'almost' SPR, i.e., H_{ev} is well approximated over a restricted frequency range by an SPR function $\bar{H}_{ev}$. Thus,

$$\mathrm{Re}[\bar{H}_{ev}(j\omega)] > 0\ ,\ \text{for all } \omega \quad (4.1a)$$

and

$$|\tilde{H}_{ev}(j\omega)| \leq \varepsilon,\quad \omega\ \varepsilon[\omega_1, \omega_2] \quad (4.1b)$$

where

$$\tilde{H}_{ev} := H_{ev} - \bar{H}_{ev} \quad (4.1c)$$

Outside of $[\omega_1, \omega_2]$ the transfer function $\tilde{H}_{ev}$ is assumed to be stable, although not necessarily small. We can now state the following result, which is similar to Lemma 4.3 in [6], and the proof can be found there.

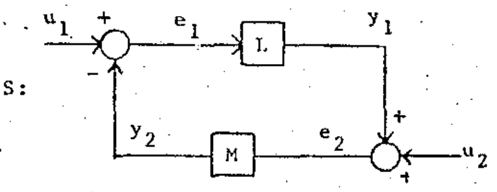

Figure 6. Feedback System S.

Theorem: Local Stability by Persistent Excitation

Let z_* be persistently exciting (see, e.g., [8]) in the sense that the system

$$\dot{x} = -z_*\bar{H}_{ev}z_*'x \quad (4.2)$$

is exponentially stable, i.e., for all bounded initial conditions $x(0)$, there exists positive constants m and λ such that

$$||x(t)|| \leq me^{-\lambda t}\ ||x(0)|| \quad (4.3)$$

Define the operator Q by

$$Q := B(z_*\tilde{H}_{ev}z_*' - e_*H_{zv}z_*') \quad (4.4)$$

with B from (3.1d). Then, the system S is stable if

$$(m/\lambda)\ \gamma_\infty(Q) < 1 \quad (4.5)$$

where $\gamma_\infty(Q)$ is the L_∞-gain of Q.

Discussion

The significance of this theorem is that H_{ev} need not be SPR, but only approximately SPR in the sense defined by (4.1). The remaining issue is the 'smallness' of Q. Since e_* is the output of the tuned system, it is expected that this is small, and so the term $e_*H_{zv}z_*$ in Q would also be small. To ensure that the remaining term $z_*\tilde{H}_{ev}z_*'$ is small requires a deeper analysis involving the selection of z_* (and $\bar{H}_{ev}$). In other words, the behavior and structure of the tuned system directly impacts the adaptive system, i.e., performance of the adaptive system can be determined by examining the tuned system.

4. Concluding Remarks

In this paper we have addressed some of the technical issues in the analysis and synthesis of adaptive reduced-order control. In particular, we focused on two issues: (1) a non-zero, but bounded tuned error signal, denoted by e_*; and (2) the SPR condition--actually the lack of SPR condition on a particular operator, denoted by H_{ev}. In practical situations, it is a certainty that e_* is only bounded and H_{ev} is not SPR. Proofs of stability—particularly global stability—rely on $H_{ev}\ \varepsilon$ SPR and $e_* \to 0$ asymptotically. This results in a discrepancy between theory and practice, which can be eliminated—as argued here—by developing conditions for local stability. To support this argument we have presented a stability result which utilizes persistent excitation together with the notion of an approximate SPR condition on H_{ev}. This result in incomplete in that further conditions on signal size, etc., are required to guarantee stability of the complete adaptive system [6]. However, we have presented a system, S, whose stability guarantees the stability of the 'linearized' adaptive system.

An intriguing conjecture to make at this point is that stability of S is both necessary and sufficient for stability of the adaptive system. To prove this conjecture requires a deeper understanding of the tuned system. Suppose, for example, that the tuned system is unique, i.e., only one set of parameters exists which stabilizes the system. In this case, the conjecture is strongly reinforced. Unfortunately, it is easily argued that there is no unique tuned system, but rather a collection of tuned systems whose performance is similar. The question then is which one to use in the analysis. In other words, where is the adaptation going, i.e., can we define the terminal set of adaptive parameters.

Along these lines, it is natural to assume that the tuned system involves constant tuned parameters. However, when e_* is only bounded, the tuned parameters are not constants. In fact, it would be surprising if these were constant. For example, in optiam control theory, the optimal controller (e.g., quadratic cost) is time-varying if the system is time-varying. Thus, the tuned parameters should also be time-varying.

Although these latter issues have not been pursued here, the input/output view as presented allows for these considerations under one framework. Thus, questions of reduced-order control and robustness to unmodeled dynamics involve a similar treatment, independent of whether the plant (or controller) is linear, nonlinear, continuum, and/or discrete. To do this, however, requires a quantification of uncertainty, that is, a more specific knowledge about the plant—at least a set description of the plant. The simpler the set, the more 'intelligent' the adaptive controller, i.e., one that is cautious in the face of uncertainty. However, this logically results in a low-performance system. To gain high performance would therefore require more imbedded intelligence about the specific location of plant uncertainty.

References

[1] K.S. Narendra, Y.H. Lin and L.S. Valavani, 'Stable Adaptive Controller Design, Part II: Proof of Stability,' IEEE Trans. Aut. Contr., Vol. AC-25, pp. 440-448, June 1980.

[2] G.C. Goodwin, P.J. Ramadge and P.E. Caines, 'Discrete-Time Multivariable Adaptive Control,' IEEE Trans. Aut. Contr., Vol. AC-25, pp. 449-456, June 1980.

[3] C. Rohrs, L. Valavani, M. Athans and G. Stein, 'Analytical Verification of Undesirable Properties of Direct Model Reference Adaptive Control Algorithms,' Proc. 20th IEEE CDC, San Diego, CA, December 1981., pp. 1272-1284.

[4] B.D.O. Anderson and C.R. Johnson, Jr., 'On Reduced-Order Adaptive Output Error Identification and Adaptive IIR Filtering,' IEEE Trans. Aut. Contr., Vol. AC-27, No. 4, Aug. 1981, pp. 927-933.

[5a] R.L. Kosut and B. Friedlander, 'Performance Robustness Properties of Adaptive Control Systems,' Proc. 21st IEEE CDC, Orlando, FL, December 1982, pp. 18-23.

[5b] R.L. Kosut and B. Friedlander, 'Robust Adaptive Control: Conditions for Global Stability,' submitted to IEEE Trans. on Aut. Cont.

[6] R.L. Kosut, 'Robust Adaptive Control: Conditions for Local Stability,' Proc. 1982 ACC, San Francisco, CA, June 1982.

[7] C.A. Desoer and M. Vidysagar, Feedback Systems: Input-Output Properties, Academic, 1975.

[8] B.D.O. Anderson, 'Exponential Stability of LInear Equations Arising in Adaptive Identification,' IEEE Trans. Aut. Contr., Vol. AC-22, No. 1, pp. 83-88, February 1977.