CONVERGENCE RESULTS FOR WIDROW'S ADAPTIVE CONTROLLER

Brian D.O. Anderson and Richard M. Johnstone

Department of Systems Engineering, Research School of Physical Sciences, Institute of Advanced Studies, Australian National University, P.O. Box 4, Canberra, A.C.T. 2600, Australia.

Abstract. Widrow and colleagues have proposed an adaptive control scheme, for the task of following a prescribed reference trajectory, which relies on adaptively identifying a finite impulse response model of the inverse of the unknown plant. The structure of the proposed controller is unlike any other proposed, and no theoretical convergence results have so far been demonstrated, despite excellent simulation experience. In this paper, a convergence result is established, which basically demands a persistency of excitation condition on the reference trajectory. This condition in fact guarantees exponentially fast decay of errors to zero. Without it however, it is possible for convergence not to occur (in contrast to some other adaptive control algorithms). Key tools in establishing the conclusion are provided by new results on the propagation of the persistent excitation property through time-invariant and slowly time-varying systems.

Keywords. Adaptive control; control system analysis; nonlinear control systems; stability; time-varying systems.

1. INTRODUCTION

In this paper, we analyze in some detail the convergence properties of a little known algorithm for adaptive control due to Widrow, McCool and Medoff (1978). As simulations in (Widrow, McCool and Medoff, 1978) indicate, this algorithm appears robust, and capable of handling the reference trajectory following problem for nonminimum phase plants. A key restriction, not indicated in the paper, is that the algorithm applies only to stable plants; also, when applied to plants other than all-pole plants, it inevitably involves some approximations. Convergence conditions are also not described. The main contribution of this paper is to establish such conditions. Not surprisingly, especially in view of recent results (Anderson and Johnson, 1981) on the Goodwin-Ramadge-Caines (Goodwin, Ramadge and Caines, 1980) algorithm, persistency of excitation of the reference trajectory plays a key role. We can in fact exhibit failure of convergence, at least for one example, when this property is lacking (an observation which is not true for the Goodwin-Ramadge-Caines algorithm), while if the property holds convergence of parameter and tracking errors occurs exponentially fast (an observation which is true for the Goodwin-Ramadge-Caines algorithm).

A key tool in establishing the result is provided in Section 4, where we establish that if a time-invariant linear finite-dimensional plant has a persistently exciting input, its output has the same property; further, if the plant is slowly varying, the conclusion remains true.

A number of the ideas of this paper also address issues related to approximation in adaptive control, see Section 6.

2. REVIEW OF THE WIDROW ADAPTIVE CONTROL ALGORITHM

As an exact algorithm, the Widrow scheme only applies to a stable autoregressive (all-pole, no zero) plant. (We shall comment later on extensions.) For such a plant, there exists a moving average (MA) system which will act as a right or left inverse-with-delay. Thus if the plant is

$$y_{k+1} = a_1 y_k + \dots + a_n y_{k-n+1} + b_1 u_{k-d+1} \tag{2.1}$$

one has

$$u_{k-d} = b_1^{-1} y_k - b_1^{-1} a_1 y_{k-1} + \dots + b_1^{-1} a_n y_{k-n} \tag{2.2}$$

Widrow's idea can be summed up using Fig. 1. A (normalized or unnormalized) LMS or other algorithm is used to adaptively identify the inverse (2.2), positioned after the plant. A replica of the inverse is positioned in front of the plant, and driven by the reference trajectory $\{y_k^*\}$, which is assumed to be known d time units into the future (because of the delay d in the plant, see (2.1), such advance knowledge of the reference trajectory is essential for any adaptive scheme).

The Normalized LMS Algorithm

This algorithm, which is identical with an equation error algorithm for a finite-impulse-response plant, works as follows. Denote the correct inverse by

$$u_{k-d} = \bar{\alpha}_1 y_k + \bar{\alpha}_2 y_{k-1} + \ldots + \bar{\alpha}_{n+1} y_{k-n} \tag{2.3}$$

(so that $\bar{\alpha}_1 = b_1^{-1}$, etc.), and an approximation by

$$\hat{u}_{k-d} = \alpha_1^k y_k + \alpha_2^k y_{k-1} + \ldots + \alpha_{n+1}^k y_{k-n} \tag{2.4}$$

Then the α_i^k can be updated by the normalized LMS algorithm

$$\alpha_i^{k+1} = \alpha_i^k - \rho^k \frac{y_{k+1-i}}{\varepsilon + \sum_{j=0}^{n} y_{k-j}^2} (\hat{u}_{k-d} - u_{k-d}) \tag{2.5}$$

where ε and ρ^k are positive constants; also, ρ^k must lie in the interval (0, 2), and is often chosen around 1. With the definitions

$$\underline{\phi}_k = [y_k \; y_{k-1} \; \ldots \; y_{k-n}]^T \quad \underline{\tilde{\theta}}_k = [\alpha_1^k - \bar{\alpha}_1 \ldots \alpha_{n+1}^k - \bar{\alpha}_{n+1}]^T \tag{2.6}$$

we have the fundamental parameter error equation

$$\underline{\tilde{\theta}}_{k+1} = [\underline{I} - \frac{\rho_k \underline{\phi}_k \underline{\phi}_k^T}{\varepsilon + \|\underline{\phi}_k\|^2}] \underline{\tilde{\theta}}_k \tag{2.7}$$

Of course, as already indicated, the signals u_k are generated by driving a replica of the adjustable model with $\{y_k^*\}$, thus

$$u_k = \alpha_1^k y_{k+d}^* + \alpha_2^k y_{k+d-1}^* + \ldots + \alpha_{n+1}^k y_{k+d-n}^* \tag{2.8}$$

3. PRELIMINARY CONSIDERATION OF CONVERGENCE

Let us note two potential difficulties with the algorithm:

(a) Regard k=0 as the initial time, and suppose that initial conditions are zero. In particular, suppose that $u_k=0$ for all $k<0$, so that $y_k=0$ for all $k<d$ and $\hat{u}_k=0$ for all $k<0$. Also, suppose that we initialize the α_i^0 by setting $\alpha_i^0=0$ for all i. Now irrespective of what values are assumed by $\{y_k^*\}$, we will have $u_0=0$ [see (2.8)], i.e. $u_k=0$ for $k \leq 0$, and then $y_k=0$ for $k \leq d$, and $\hat{u}_k=0$ for $k \leq 0$. Consequently, from (2.5), we get $\alpha_i^1=0$ for all i. Repeating the argument shows that $u_k=0$, $y_k=0$, $\hat{u}_k=0$ and $\alpha_i^k=0$ for all k and i. Thus there is no reference trajectory tracking. This behaviour is caused by zero initial conditions (or for that matter a long enough sequence of zero $\{y_k^*\}$ values), together with the adapting coefficients α_i^k becoming all zero.

(b) Consider now a plant with n=2, and suppose that again we have the zero initial conditions on u_k, y_{k_0} and $\hat{u}_k$ as described above, while $\alpha_1^0=1$, $\alpha_2^0=-\frac{1}{2}$, $\alpha_3^0=-\frac{1}{2}$. Suppose further that $y_k^*=1$ for all k. Then we obtain $u_0=0$, and then arguing as in (a), we conclude that the α_i^k will not change, and that $y_k=0$ for all k. Again, there is no reference trajectory tracking. The difficulty here is really associated with the failure of the signal $\{y_k^*\}$ to be persistently exciting.

Our main result is that if we prevent all the α_i^k from becoming zero, and if the signal $\{y_k^*\}$ is persistently exciting in an appropriate sense, then the algorithm works satisfactorily. In particular, the α_i^k approach the steady state values $\bar{\alpha}_i$ exponentially fast, and $y_k - y_k^*$ approaches zero exponentially fast. In broad outline, the argument is as follows. (Reference to Fig. 1 may be helpful)

(a) if $\{y_k\}$ is persistently exciting, the α_i^k will converge exponentially fast as will $u_k - \hat{u}_k$ (a well-known result, see e.g. (Weiss and Mitra, 1979; Bitmead and Anderson, 1980) viewed in a nonstandard way);

(b) if $\{u_k\}$ is persistently exciting, then $\{y_k\}$ will also have this property (a new result, but tied to recent results on persistently exciting signals, (Anderson and Johnson, 1981);

(c) if $\{y_{k+d}^*\}$ is persistently exciting, and adaption had occurred, i.e. the α_i^k were all constant $\{u_k\}$, would be persistently exciting [same argument as (b)]. The argument of (b) can be extended to the varying α_i^k case: because the variation of the α_i^k inevitably slows down during the adaption process. One can still argue (new result) that persistency of excitation of y_{k+d}^* implies the same property for u_k, provided that the α_i^k cannot all go to zero;

(d) one can combine the ideas of (b) and (c) into one idea; a simple linear equation links $\{y_k^*\}$ to $\{y_k\}$, with time-varying adaptive coefficients. But if the time-variation is slow enough, (and this is a consequence of the adaption equations) persistency of excitation of $\{y_k^*\}$ implies the same property for $\{y_k\}$.

4. PERSISTENCY OF EXCITATION

We shall now consider some results in a more general form than we need. The first result connects persistency of excitation of inputs and outputs of a time-invariant plant.

Theorem 4.1: Suppose that a plant is described by

$$y_k+a_1y_{k-1}+\dots+a_ny_{k-n} = b_du_{k-d}+\dots+b_mu_{k-m} \quad (4.1)$$

that the plant is stable, and that for all j, and some S>0, q≥0 for which S ≥ m+n-d+q, and some $\mu_1,\mu_2>0$

$$\mu_2 I \geq \sum_{k=j+n}^{j+S} \begin{bmatrix} u_{k-d} \\ \vdots \\ u_{k-m-q} \end{bmatrix} [u_{k-d}\dots u_{k-m-q}] \geq \mu_1 I>0 \quad (4.2)$$

Let $a_{max}=\max\{|a_i|, 1\}$, and $\|b\|^2 = \sum_{i=d}^{m} b_i^2$. Then there exists $\lambda_2 > 0$ and η depending only on m-d+1 and q such that for all j

$$\lambda_2 I \geq \sum_{k=j}^{j+S} \begin{bmatrix} y_k \\ \vdots \\ y_{k-q} \end{bmatrix} [y_k\dots y_{k-q}] \geq \frac{\eta\mu_1\|b\|^2}{(S-n+1)(n+1)^2 a_{max}^2} \quad (4.3)$$

The proof is omitted, due to restrictions on length. The ideas in the proof are similar to those in (Anderson and Johnson, 1981).

Next, we consider a modification of the above result to cope with slowly time-varying plants; the idea is that if (4.1) is slowly enough time-varying, persistency of excitation is retained.

Theorem 4.2: Suppose that

a) $y_k+a_1^k y_{k-1}+\dots+a_n^k y_{k-n}=b_d^k u_{k-d}+\dots+b_m^k u_{k-m}$ (4.4)

b) The a_i^k are bounded and vary so as to retain stability, i.e. $\|u_k\| < c_1 \ \forall\ k$ implies $\|y_k\| < c_2c_1 \ \forall\ k$ and for some fixed c_2.

c) $\max_k \sum_{i=d}^{m} (b_i^k)^2 < \infty$ and $\min_k \sum_{i=d}^{m} (b_i^k)^2 = \|b\|^2_{min} > 0$.

d) For all j and some S > 0, q ≥ 0 with S≥ m+n-d+q, and some μ_1, μ_2 > 0, (4.2) holds

e) let $(\Delta a)_{max} = \max_{\substack{i\in[1,n] \\ k\in[j+n-q,j+S] \\ j}} |a_i^{j+n} - a_i^k|$

$$(\Delta b)_{max} = \max_{\substack{i\in[d,m] \\ k\in[j+n-q,j+S] \\ j}} |b_i^{j+n} - b_i^k|$$

$$\Delta = \max[(\Delta a)_{max}, (\Delta b)_{max}]$$

(Thus Δ is a measure of time variation of the plant).

If Δ is such that

$$\Delta < \eta' \frac{\|b\|_{min}\sqrt{\mu_1}}{N\sqrt{S-n+1}} \quad (4.5)$$

where η' depends only on m,n,d and q and N is a global bound on $|u_k|$, $|y_k|$, existing by (b) and (d), then there exist $\lambda_1,\lambda_2>0$ such that for all j,

$$\lambda_2 I \geq \sum_{j}^{j+S} \begin{bmatrix} y_k \\ \vdots \\ y_{k-q} \end{bmatrix} [y_k\dots y_{k-q}] \geq \lambda_1 I \quad (4.6)$$

where

$$\sqrt{\lambda_1} < \left[\frac{\eta\|b\|_{min}\sqrt{\mu_1}}{\sqrt{S-n+1}} - \frac{\eta N}{\eta'}\Delta\right]\frac{1}{(n+1)\max_{i,k}[1,|a_k^i|]} \quad (4.7)$$

Again, the proof is omitted, due to restrictions on length.

5. CONVERGENCE OF THE WIDROW SCHEME

Consider again Fig. 1, and observe that y_k^* and y_k are related by the following equation:

$$b_1^{-1}y_k - b_1^{-1}a_1y_{k-1} - \dots - b_1^{-1}a_ny_{k-n}=u_{k-d} = \alpha_1^{k-d}y_k^*+\alpha_2^{k-d}y_{k-1}^*+\dots+\alpha_{n+1}^{k-d}y_{k-n}^* \quad (5.1)$$

With (5.1) paralleling (4.4), we see that conditions (a) and (b) of Theorem 4.2 are clearly fulfilled. The adaptive algorithm ensures the α_i^k are bounded, as is well known. We shall show later how one secures

$$\sum_{i=1}^{n+1}(\alpha_i^{k-d})^2 \geq \|b\|^2_{min} > 0 \quad (5.2)$$

for all k. Thus condition (c) will be satisfied. Condition (d) translates to a requirement that y_k^* be persistently exciting; note that by virtue of (5.1), the quantities m,n, and d in Theorem 4.2 become n,n and 0. We shall take q=n. Accordingly, S≥3n. If we can show that the a_i^k vary slowly enough, (4.5) will hold, and persistent excitation of $\{y_k\}$ will follow as per (4.6).

It remains therefore to show how (5.2) may be secured, and exhibit the asserted slow variation of the α_i^k.

We can secure (5.2) by adjusting the quantities ρ_i^k in the adaptive update equations (2.5). We require one a priori assumption.

Assumption: A quantity σ>0 is known such that the unknown $\bar{\alpha}_i$ satisfy,

$$\sum_i \bar{\alpha}_i^{2} > \sigma > 0$$

We set $\rho_i^k = \rho$ a fixed constant in (0,2) unless this would result in $\sum_i |\alpha_i^{k+1}|^2 < 0.1\sigma$, in which case we will then adjust ρ^k, ensuring that $\rho^k \in (0,2)$, to cause $\sum_i |\alpha_i^{k+1}|^2 > 0.1\sigma$. We omit further details.

Now we consider the question of time-variation of the α_i^k. Bearing in mind (2.6), it is equally acceptable to look at the variation of $\tilde{\theta}_k$. Now it is easily established from (2.7) that $\| \tilde{\underline{\theta}}_k \|^2$ is monotone decreasing and that

$$\frac{\underline{\phi}_k^T \tilde{\underline{\theta}}_k}{\varepsilon + \|\underline{\phi}_k\|^2} \rightarrow 0$$

Consequently, for any fixed j, $\|\tilde{\underline{\theta}}_{k+j} - \tilde{\underline{\theta}}_k\| \rightarrow 0$ as $k \rightarrow \infty$. Thus, at least for "large" times, $\{y_k\}$ will be persistently exciting. In fact, persistency of excitation can be expected beginning at comparatively early times. The only way we will loose persistency of excitation is basically when $\|\tilde{\underline{\theta}}_{k+1} - \tilde{\underline{\theta}}_k\|$ is too big. Now for constant ρ^k

$$\| \tilde{\underline{\theta}}_{k+1} - \tilde{\underline{\theta}}_k \|^2 < \rho^2 \frac{(\underline{\phi}_k^T \tilde{\underline{\theta}}_k)^2}{\varepsilon + \|\underline{\phi}_k\|^2}$$

while also

$$\| \tilde{\theta}_k \|^2 - \| \tilde{\theta}_{k+1} \|^2 \geq (2\rho - \rho^2) \frac{(\phi_k^T \tilde{\underline{\theta}}_k)^2}{\varepsilon + \|\underline{\phi}_k\|^2}$$

It follows that when $\| \tilde{\theta}_{k+1} - \tilde{\theta}_k \|$ is large, $\| \tilde{\theta}_k \|$ is most rapidly decreasing. Thus $\| \tilde{\theta}_{k+1} - \tilde{\theta}_k \|$ can only exceed a prescribed bound a finite number of times, i.e. there will be only a finite number of time instants at which persistency of excitation is lost. At and around these time instants divergence does not occur; rather, (exponential) convergence just temporarily ceases.

6. BEHAVIOUR OF THE ALGORITHM FOR PLANTS WITH ZEROS

Simulations in (Widrow, McCool and Medoff, 1978) indicate that the proposed algorithm is in fact not restricted in its practical application to all-pole plants, or even minimum phase plants (though it is restricted to stable plants). The broad theoretical thinking behind this conclusion is as follows.

(a) any nonminimum phase plant with no zeros on the unit circle can be approximated by an all pole plant, perhaps with much greater delay.[1]

(b) with persistently exciting $\{y_k^*\}$, the whole adaptive system in an ideal case is exponentially stable, and therefore tolerant of departures from the idealizing assumptions. Approximation of the plant can be viewed as one such departure.

In more detail, let us assume simply that the unknown plant is time-invariant and stable, and that there exists an approximate finite-impulse-response inverse. Thus

$$u_{k-d} = \bar{\alpha}_1 y_k + \ldots + \bar{\alpha}_{n+1} y_{k-n} + \eta_k$$

where η_k is the error involved in making the approximation. Note that any finite-dimensional plant with no zeros on the unit circle has an approximate inverse of this form where, by taking d and n large enough, we can assume $\eta_k \leq K \max_{\ell \leq k} \{|y_\ell|\}$ where K is arbitrarily small.

We update the estimates of the $\tilde{\alpha}_i$ as before. Thus (2.4), (2.5) still apply. [1]With the definitions of (2.6), we obtain for the parameter error vector equation

$$\tilde{\underline{\theta}}_{k+1} = \left[\underline{I} - \frac{\rho_k \underline{\phi}_k \underline{\phi}_k^T}{\varepsilon + \|\underline{\phi}_k\|^2} \right] \tilde{\underline{\theta}}_k + \frac{\rho_k \underline{\phi}_k}{\varepsilon + \|\underline{\phi}_k\|^2} \eta_k \qquad (6.1)$$

The previously homogeneous equation (2.7) is replaced by the forced equation (6.1); in general, (6.1) cannot have a solution which approaches zero. Nevertheless, if $|\eta_k|$ is small for all k, and if the homogeneous part of (6.1) is exponentially stable, $|\tilde{\theta}_k|$ will be small, with a bound depending on the $|\eta_k|$ bound.

Proceeding as before, we define

$$u_k = \alpha_1^k y_{k+d}^* + \alpha_2^k y_{k+d-1}^* + \ldots + \alpha_{n+1}^k y_{k+d-n}^* \qquad (6.2)$$

while the plant maps the sequence $\{u_k\}$ into $\{y_k\}$ in a time-invariant fashion. We denote this by:

$$y_k = P(u_k) \qquad (6.3)$$

Equations (6.1) through (6.3) are coupled equations describing the adaptive algorithm. Notice that the presence of $\{\eta_k\}$ affects $\{\phi_k\}$, for from (6.1), it is clear that $\{\eta_k\}$ affects $\{\tilde{\theta}_k\}$, from (6.2) that $\{\eta_k\}$ affects $\{u_k\}$, from (6.3) that $\{\eta_k\}$ affects $\{y_k\}$ and thus $\{\phi_k\}$, since ϕ_k is just a collection of y_j stacked together. To derive a boundedness result is therefore not completely straightforward since (6.1) is not simply a linear equation

---

[1]Compute the 2-sided inverse z-transform of the plant and truncate it in both time directions. Do a time translation to get a finite impulse response for the inverse-with-delay approximating plant.

with external input, but we shall outline such a procedure.

Consider the coupled set of equations (6.1) through (6.3) with $\eta_k \equiv 0$. Note that (6.3) is unchanged, so that an equation of the form

$$u_{k-d} = \tilde{\alpha}_1 y_k + \dots + \tilde{\alpha}_{n+1} y_{k-n}$$

is not being assumed to hold. With $\eta_k \equiv 0$, the equations do not describe any adaptive system, but they will allow us to draw some conclusions about the adaptive system we are studying. The same arguments as in the last section allow us to conclude that $\tilde{\theta}_k \to 0$ exponentially fast if y_k^* is persistently exciting (but not that $u_k - \hat{u}_k \to 0$ exponentially fast); for (6.1) with $\eta_k \equiv 0$ implies that the variations in $\tilde{\theta}_k$ and thus the α_i^k for $i=1, \dots, n+1$ become slower and slower. Then (6.2) and (6.3) imply y_k is persistently exciting. Then (6.1) with $\eta_k \equiv 0$ implies exponential convergence of $\tilde{\theta}_k$ to zero.

When η_k is not zero in (6.1), we can appeal to ideas of (Desoer and Vidyasagar, 1975) to conclude that there exists constant M such that if $|\eta_k| < M$ for all k, $\limsup_{k\to\infty} \|\tilde{\theta}_k\| < N \limsup_{k\to\infty} |\eta_k|$ for some further constant N, i.e. the error in $\tilde{\theta}_k$ introduced by a nonzero η_k is bounded by the size of the $\{\eta_k\}$ sequence, so long as some absolute bound on this sequence is not exceeded.[2] In effect, we can conclude a bounded-input, bounded-output stability result given an exponential stability result, as is possible in a purely linear situation. What is additional to the purely linear situation is the introduction of the absolute bound M. The implication for the adaptive control problem is that if the approximation of the plant inverse by a finite impulse response involves too crude an approximation, then the whole algorithm could collapse.

7. CONCLUSIONS

There are perhaps three points in the paper deserving emphasis. First, persistency of excitation of $\{y_k^*\}$ plays an even more critical role than in (Anderson and Johnson, 1981; Goodwin, Ramadge and Caines, 1980). Without it, not only is exponential convergence lost, but maybe any convergence at all. Second, we have discussed how persistency of excitation can propagate through a system, even a time-varying system; this is a potentially valuable idea. Third, we have argued how certain adaptive control results can be made approximate, in the vein of (Desoer and Vidyasagar, 1975; Johnson and Anderson, 1981; Anderson and Johnstone, 1981).

Several interesting ideas for future work present themselves. One could attempt to replace the finite-impulse response model of the inverse by a pole-zero model, identified via an equation-error type of algorithm. This should be straightforward. More difficult would be the elimination of the constraint that the plant be stable while retaining the basic Widrow idea.

[2]The full argument is omitted, due to its length.

REFERENCES

Anderson, B. D. O., and C. R. Johnson (Jr.) (1981). Exponential Convergence of Adaptive Identification and Control Algorithms, Automatica, to appear. Also, Proc. 8th IFAC Triennial World Congress, Kyoto.

Anderson, B. D. O., and R. M. Johnstone (1981). Robust Lyapunov Stability Results and Adaptive Systems, Proc. IEEE Conference on Decision and Control, to appear.

Bitmead, R. R., and B. D. O. Anderson (1980). Performance of Adaptive Estimation Algorithms in Dependent Random Environments, IEEE Trans. Autom. Control, AC-25, 788 - 794.

Desoer, C. A., and M. Vidyasagar (1975). Feedback Systems : Input-Output Properties, Academic Press, New York.

Goodwin, G. C., P.J. Ramadge, and P. E. Caines (1980). Discrete-time Multivariable Adaptive Control, IEEE Trans. Autom. Control, AC-25, 449 - 456.

Johnson, C. R. (Jr.), and B. D. O. Anderson (1981). On Reduced Order Adaptive Output Error Identification and Adaptive IIR Filtering, IEEE Trans. Autom. Control, to appear.

Weiss, A., and D. Mitra (1979). Digital Adaptive Filters : Conditions for Convergence, Rates of Convergence, Effects of Noise and Errors Arising from the Implementation, IEEE Trans. Info. Theory, IT-25, 637 - 652.

Widrow, B., J. M. McCool, and B. P. Medoff (1978). Adaptive Control by Inverse Modelling, Proc. 1978 Asilomar Conference.

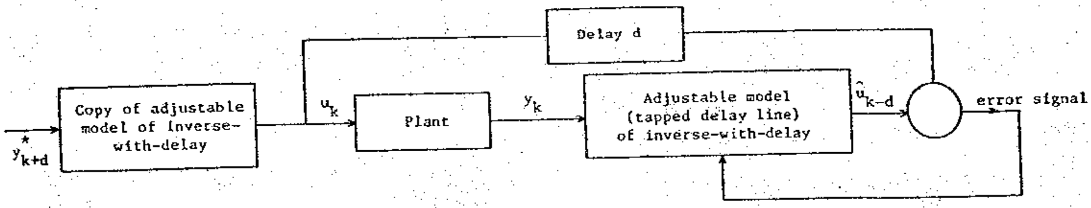

Figure 1: Idea of the Widrow adaptive control scheme for following a reference trajectory; arguments on signals indicate what is observed at time k.