ADAPTIVE INVERSE CONTROL

Bernard Widrow

Department of Electical Engineering, Information Systems Laboratory, Stanford University, Stanford, CA 94305, USA

Abstract. Adaptive control is seen as a two part problem, control of plant dynamics and control of plant noise. The parts are treated separately.

An unknown plant will track an input command signal if the plant is driven by a controller whose transfer function approximates the inverse of the plant transfer function. An adaptive inverse identification process can be used to obtain a stable controller, even if the plant is nonminimum phase. A model-reference version of this idea allows system dynamics to closely approximate desired reference model dynamics. No direct feedback is used, except that the plant output is monitored and utilized in order to adjust the parameters of the controller.

Control of internal plant noise is accomplished with an optimal adaptive noise canceller. The canceller does not affect plant dynamics, but feeds back plant noise in a way that minimizes plant output noise power.

Keywords. Adaptive control, modeling, identification, inverse modeling, noise cancelling, deconvolution, adaptive inverse control.

INTRODUCTION

This paper presents a very brief description of a means of using adaptive filtering techniques for solution of adaptive control problems. The emphasis of this paper is on systems concepts rather than mathematic analysis. A full paper by B. Widrow and E. Walach is in preparation which will address the issues more completely.

Many problems in adaptive control can be divided into two parts: (a) control of plant dynamics, and (b) control of plant noise. Very often, a single system is utilized to achieve both of these control objectives. The approach of this paper treats each problem separately. Control of plant dynamics can be achieved by preceding the plant with an adaptive controller whose transfer function is the inverse of that of the plant. Control of plant noise can be achieved by an adaptive feedback process that minimizes plant output noise without altering plant dynamics.

ADAPTIVE FILTERS

An adaptive digital filter, shown in Fig. 1, has an input, an output, and another special input called the "desired response." The desired response input is sometimes called the "training signal."

The adaptive filter contains adjustable parameters that control its impulse response. These parameters could, for example, be variable weights connected to the taps of a tapped delay line or to internal points of a lattice filter. There are many ways to configure such a filter.

The adaptive filter also incorporates an "adaptive algorithm" whose purpose is to automatically adjust the parameters to minimize some function of the error (usually mean square error). The error is defined as the difference between the desired response and the actual filter response. Many such algorithms exist, a number of which are described in the textbook by Widrow and Stearns [1].

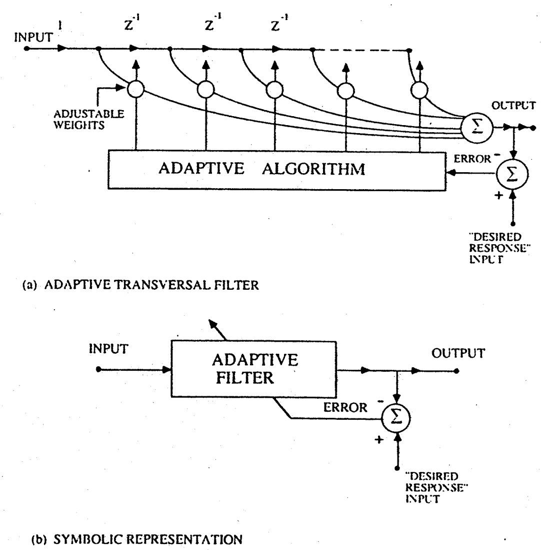

FIG. 1 AN ADAPTIVE FILTER

For the purposes of this paper, the adaptive filter may be considered to be like the one shown in Fig. 1a, an adaptive tapped delay line or transversal filter. With fixed weights, this is a linear finite-impulse-response (FIR) digital filter having a transfer function with only zeros, no poles in the finite z-plane. This will be the basic building block for the adaptive systems to be described below.

DIRECT PLANT IDENTIFICATION

Adaptive plant modeling or identification is an important function. Figure 2 illustrates how this can be done with an adaptive FIR filter. The plant input signal is the input to the adaptive filter. The plant output signal is the desired response for the adaptive filter. The adaptive algorithm minimizes mean square error, causing the model $\hat{P}$ to be a best least squares match to the plant P for the given input signal and for the given set of weights allocated to $\hat{P}$.

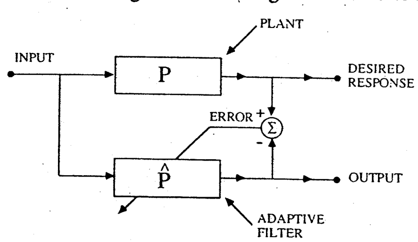

FIG 2. DIRECT IDENTIFICATION

INVERSE PLANT IDENTIFICATION

Another important function is inverse plant identification. This technique is illustrated in Fig. 3. The plant input is as before. The plant output is the input to the adaptive filter. The desired response for the adaptive filter is the plant input in this case. Minimizing mean square error causes the adaptive filter $\hat{P}^{-1}$ to be a best least squares inverse to the plant P for the given input spectrum and for the given set of weights of the adaptive filter. The adaptive algorithm attempts to make the cascade of plant and adaptive inverse behave like a unit gain. This process is often called deconvolution.

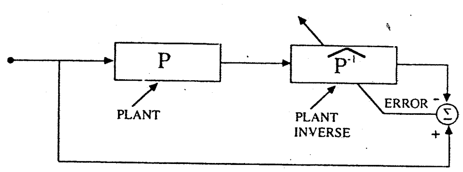

FIG. 3 INVERSE IDENTIFICATION

For sake of argument, the plant can be assumed to have poles and zeros. An inverse, if it also had poles and zeros, would need to have zeros where the plant had poles and poles where the plant had zeros. Making an inverse would be no problem except for the case of a nonminimum phase plant. It would seem that such an inverse would need to have unstable poles, and this would be true if the inverse were causal. If the inverse could be noncausal as well as causal however, then a two-sided stable inverse would exist for all linear time-unvariant plants in accord with the theory of two-sided Laplace transforms.

A causal FIR filter can approximate a delayed version of the two-sided plant inverse, and an adaptive FIR filter can self adjust to this function. The method is illustrated in Fig. 4. The time span of the adaptive filter (the number of weights multiplied by the sampling period) can be made adequately long so that the mean square error of the optimized inverse would be a small fraction of the plant input power. To achieve this objective with a nonminimum phase plant, the delay Δ needs to be chosen appropriately.

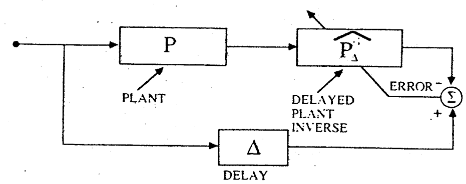

FIG. 4 DELAYED INVERSE IDENTIFICATION

The choice is generally not critical.

The inverse filter is used as a controller in the present scheme, so that Δ becomes the response delay of the controlled plant. Making Δ small is generally desirable, but the quality of control depends upon the accuracy of the inversion process which sometimes requires Δ to be of the order of half the length of the adaptive filter.

A computer simulation experiment has been done to illustrate the effectiveness of the inversion process. Figure 5a shows the impulse response of a nonminimum phase plant having a small transport delay. Figure 5b shows the impulse response of the best least squares inverse with a delay of $\Delta = 26$ sample periods. The error power was less than 5% of the plant input power. Figure 5c is a convolution of the plant and its inverse impulse response. The result is essentially a unit impulse at a delay of 26, with small "sidelobes" elsewhere.

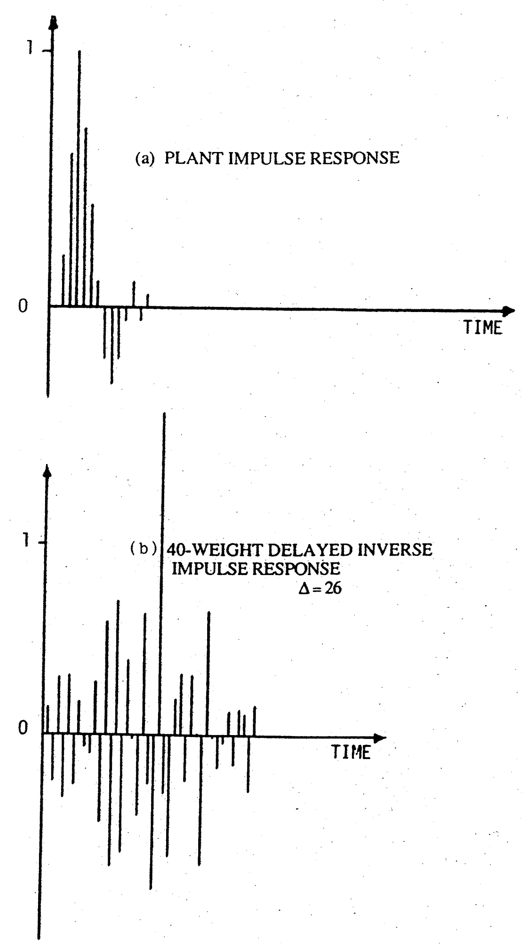

FIG. 5 A PLANT AND ITS DELAYED INVERSE

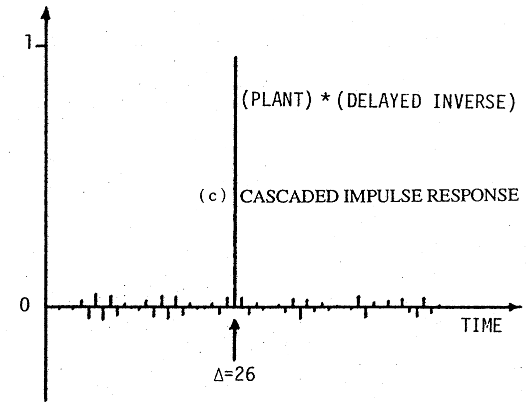

FIG. 5 (CONT'D) A PLANT AND ITS DELAYED INVERSE

A model-reference inversion process is shown in Fig. 6. A reference model is used in place of the delay of Fig. 4. Minimizing mean square error with the system of Fig. 6 causes the cascade of the plant and its "model-reference inverse" to closely approximate the response of a model M. Much is known about the design of model reference systems [2]. The model is chosen to give a desirable response to the overall system. Some delay may need to be incorporated into the model in order to achieve low error.

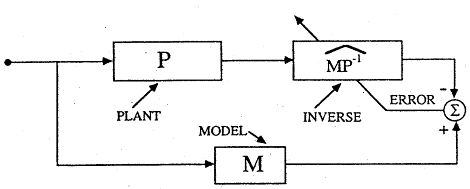

FIG. 6 MODEL-REFERENCE PLANT INVERSE

Thus far, the plant has been noise free. Plant noise creates difficulty for the inverse modeling schemes of Figs. 3, 4, and 6. The noise will bias the inverse solutions. Wiener filter theory tells why: Plant output noise goes directly into the inputs of the adaptive filters, biasing the input covariance matrices of these filters.

To avoid this problem, the scheme of Fig. 7 can be used. A direct modeling process yields $\hat{P}$. Wiener filter theory shows why $\hat{P}$ is unbiased: The plant noise does not affect the input to $\hat{P}$ and therefore does not influence its covariance matrix. The noise is added to the desired response of $\hat{P}$, i.e. to the plant output. But the plant output noise is not correlated with the plant input. The result is that for the adaptive filter, the plant noise does not affect the crosscorrelation between the desired response and the adaptive filter input. Therefore, the Wiener solution for $\hat{P}$ is unbiased.

Now using an exact copy of $\hat{P}$ in place of P, an off line process is shown in Fig. 7 which calculates the model reference plant inverse. The off line process can run faster than real time, so that as $\hat{P}$ is calculated, the model reference inverse is immediately obtained.

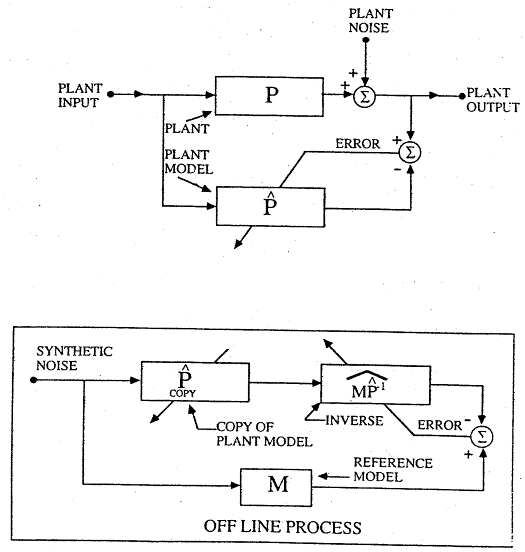

FIG. 7 MODEL-REFERENCE INVERSE OF NOISY PLANT

ADAPTIVE CONTROL OF PLANT DYNAMICS

Now having the plant inverse, it can be used as a controller to provide a driving function for the plant. This simple idea is illustrated in Fig. 8. Error analysis for this structure has been done and will be presented in the paper by Widrow and Walach. Many simulation examples have been done, with consistent good results. The idea works. The plant must be stable, and the plant zeros should not be very close to the $j\omega$-axis in the s-plane (analog) or to the unit circle in the z-plane (digital), otherwise a very long inverse filter would be required.

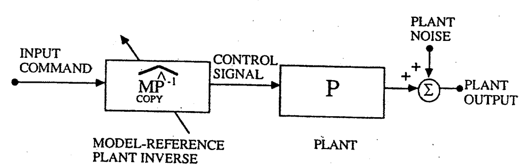

FIG. 8 DYNAMIC CONTROL OF NOISY PLANT

ADAPTIVE PLANT-NOISE CANCELLING

The system of Fig. 8 only controls and compensates for plant dynamics. The noise appears at the plant output unabated. The only way that the plant output noise can be reduced is to obtain this noise from the plant output and process it, then feed it back into the plant input. The system shown in Fig. 9 does this.

In Fig. 9, an exact copy of $\hat{P}$ is fed the same input signal as the plant P. The output of this $\hat{P}$ copy is subtracted from the plant output. Assuming that $\hat{P}$ has a dynamic response essentially identical to that of the plant P, the difference in the outputs is a close estimate of the plant noise. This noise is filtered by Q and then subtracted from the plant input. The filter Q is generated by an off line process that delivers new values of Q almost instantaneously with new values of $\hat{P}$.

The filter Q is the best inverse (without delay) of $\hat{P}$, essentially the best inverse of P (without delay). The "synthetic noise" should have a spectral character like that of the plant noise. It will be shown in the Widrow and Walach paper that the noise cancelling system of Fig. 9 adapts and converges to minimize the plant noise at the plant output. As such, it is an optimal linear least squares system. There is no way to further reduce the plant noise.

The system of Fig. 9 appears to be a feedback system. However, if $\hat{P}$ is dynamically the same as P, the transfer function around the loop is zero. The transfer function from the "control signal" input point to the "plant output" point is that of the plant alone. Thus, the noise canceller does not affect the plant dynamics.

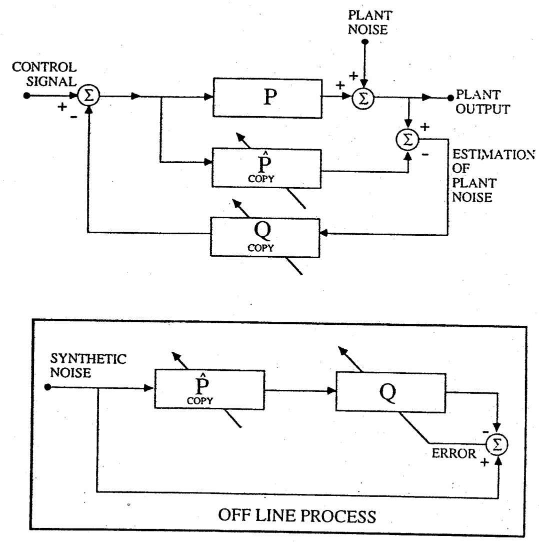

FIG. 9 ADAPTIVE NOISE CANCELLER FOR NOISY PLANT

Figure 10 shows results of a computer plant noise cancellation experiment. The plant in this case was minimum phase. Almost perfect noise cancellation is possible with a minimum phase plant, and this is evident from the experiment.

ADAPTIVE INVERSE CONTROL

The system of Fig. 11 combines all of the parts, allowing control of plant dynamics and control of plant noise. The entire system is called "adaptive inverse control."

In Fig. 11, dither noise is used in the plant identification process to obtain $\hat{P}$. This should be done in cases where the plant input signal is not persistent or is otherwise unsuitable for plant identification. Also in Fig. 11, one can see a "panic button" for breaking the noise cancelling feedback. Emergency conditions could develop if the plant P suddenly underwent massive changes in dynamics. Its model $\hat{P}$ would require time to catch up, and in the meanwhile the whole noise canceller could go unstable. Pushing the panic button saves the situation, and releasing it as soon as $\hat{P}$ converges to P causes plant noise cancelling to be resumed.

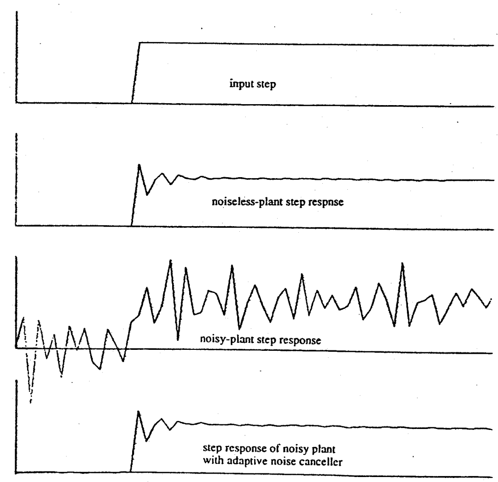

FIG. 10 OPTIMAL ADAPTIVE NOISE CANCELLING WITH A MINIMUM PHASE PLANT

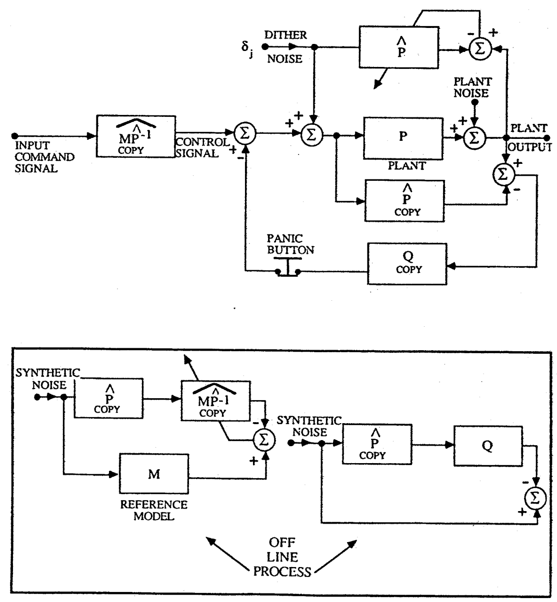

FIG. 11 ADAPTIVE INVERSE CONTROL

CONCLUSION

Methods for adaptive control of plant dynamics and for control of plant noise have been described. For their proper application, the plant must be stable, and the plant zeros should not be extremely close to the $j\omega$-axis of the s-plane or the unit circle of the z-plane. An unstable plant could first be stabilized with feedback, then adaptively controlled. The feedback approach could also be used to move plant zeros if required. Proper design of such feedback is a subject of current research.

REFERENCES

Widrow, B. and Stearns, S. D. (1985). Adaptive Signal Processing. New Jersey: Prentice-Hall.

Landau, I. D. (1979). Adaptive Control. The Model Reference Approach, Control and Systems Theory Series, Vol. VIII. Marcel Dekker, New York.

Widrow, B., McCool, J., and Medoff, B. (1978). Adaptive control by inverse modeling. Twelfth Asilomar Conference on Circuits, Systems, and Computers.

Widrow, B., Shur, D., and Shaffer, S. (1981). On adaptive inverse control. Fifteenth Asilomar Conference on Circuits, Systems, and Computers.

Anderson, B. D. O. and Johnstone, R. M. (1981). Convergence results for Widrow's adaptive controller. IFAC Conf. on System Identification.

Widrow, B. (1986). Adaptive inverse control. Second IFAC Workshop on Adaptive Systems in Control and Signal Processing, Sweden.

Astrom, K. J. and Wittenmark, B. (1973). On self-tuning regulators. Automatica Vol. 9, No. 2.

Astrom, K. J. and Wittenmark, B. (1974). Analysis of self-tuning regulator for nonminimum phase systems. IFAC Symposium on Stochastic Control, Budapest.

Astrom, K. J. and Wittenmark, B. (1980). Self-tuning controllers based on pole-zero placement. Proc. IEE, Vol. 127, Pt.D., No. 3, pp. 120-130.

Landau, I. D. (1974). A survey of model reference adaptive techniques-theory and applications. Automatica, Vol. 10, pp. 353-379.

Landau, I. D. (1976). Unbiased recursive identification using model reference adaptive techniques. IEEE Transactions on Automatic Control, Vol. 21.

Martin-Sanchez, J. M. (1976). A new solution to adaptive control, Proc. IEEE, Vol. 64, No. 8.

Nakamura, K. and Yoshida, Y. (1973). Learning dual control under complete state information. A paper presented at NSF Workshop on Learning System Theory and its Applications, Florida.

Tse, E. and Athans, M. (1972). Adaptive stochastic control for a class of linear systems. IEEE Tansactions on Automatic Control. pp. 38-51.

Widrow, B. and others. (1975). Adaptive noise cancelling: principles and applications. Proc. IEEE, Vol. 63, pp. 1692-1716.

Astrom, K. J. (1970). Introduction to Stochastic Control, New York, Academic Press.

Kailath, T. (1980). Linear Systems, Prentice-Hall.

Lee, D. T. L., Morf, M., and Friedlander, B. (1981). Recursive least squares ladder estimation algorithms. IEEE Trans. Circuits Syst., Vol. CAS-28, No. 6, pp. 467-481.

Griffiths, L. J. (1977). A continuously adaptive filter implemented as a lattice structure. Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing, Hartford, CT, pp. 683-686.

Friedlander, B. (1982). Lattice filters for adaptive processing. Proc. IEEE, Vol. 70, No. 8, pp. 829-867.

Bitmead, R. R. (1983). Convergence in distribution of LMS-type adaptive parameter estimates. IEEE Trans. Automat. Contr., Vol. AC-28, No. 1, pp. 54-60.

Macchi, O. and Eweda, E. (1983). Second-order convergence analysis of stochastic adaptive linear filtering. IEEE Trans. Automat. Contr., Vol. AC-28, No. 1, pp. 76-85.

Cioffi, J. M. and Kailath, T. (1984). Fast, recursive-least-squares, transversal filters for adaptive filtering. IEEE Acoust. Speech, Signal Processing, Vol. ASSP-32.