ADAPTIVE INVERSE CONTROL BASED ON NONLINEAR ADAPTIVE FILTERING

Bernard Widrow[1], **Gregory Plett**[2], **Edson Ferreira**[3] **and Marcelo Lamego**[4]

Information Systems Lab., EE Dep., Stanford University

Abstract: Many problems in adaptive control can be divided into two parts; the first part is the control of plant dynamics, and the second is the control of plant disturbance. 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 plant disturbance can be achieved by an adaptive feedback process that minimizes plant output disturbance without altering plant dynamics. The adaptive controller is implemented using adaptive filters.

Keywords: Adaptive Control, Inverse Control, Adaptive Filters, Neural Networks, Nonlinear Systems.

1. INTRODUCTION

At present, the control of a dynamic system (the "plant") is generally done by means of feedback. This paper proposes an alternative approach that uses adaptive filtering to achieve feedforward control. Precision is attained because of the feedback incorporated in the adaptive filtering. The control of plant dynamic response is treated separately, without compromise, from the optimal control of the disturbance. All of the required operations are based on adaptive filtering techniques (Widrow and Walach, 1996). Following the proposed methodology, knowledge of adaptive signal processing allows one to go deeply into the field of adaptive control.

In order for adaptive inverse control to work, the plant must be stable. If the plant is not stable, then conventional feedback methods should be used to stabilize it. Generally, the form of this feedback is not critical and would not need to be optimized. If the plant is stable to begin with, no feedback would be required. If the plant is linear, a linear control system would generally be used. The transfer function of the controller converges to the reciprocal of that of the plant. If the plant is minimum phase, an inverse is easily obtained. If the plant is non-minimum phase, a delayed inverse can be obtained. The delay in the inverse results in a delay in overall system response, but this is inevitable with a non-minimum-phase plant. The basic idea can be used to implement "model-reference control" by adapting the cascaded filter to cause the overall system response to match a pre-selected model response.

Disturbance in a linear plant, whether minimum phase or non-minimum phase, can be optimally controlled by a special circuit that obtains the disturbance at the plant output, filters it, and feeds it back into the plant input. The circuit works in such a way that the feedback does not alter the plant dynamic response. So disturbance control and control of dynamic response can be accomplished separately. The same ideas work for MIMO systems as well as SISO systems.

Control of nonlinear plants is an important subject that raises significant issues. Since a nonlinear plant does not have a transfer function, how could it have an inverse? By using a cascade of the nonlinear adaptive filter with the nonlinear plant, the filter can learn to drive the plant as if it were the plant's inverse. This works surprisingly well for a range of training and operation signals. Control of dynamic response and

---

[1] Professor.
[2] Ph.D. student.
[3] Visiting Professor sponsored by CAPES/UFES, Brazil.
[4] Ph.D. student sponsored by CNPq/UFES, Brazil.

plant disturbance can be done.

This paper introduces adaptive inverse control by first discussing adaptive filters. Then, inverse plant modeling for linear plants is described. The ideas are extended to nonlinear control, examples are presented and conclusions made.

2. 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".

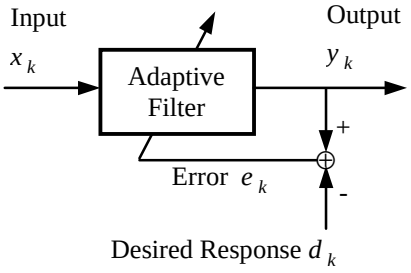

Fig. 1. *Symbolic representation of an adaptive transversal filter adapted by the LMS algorithm*

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. The filter would thus be FIR, finite impulse response.

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 text-books by Widrow and Stearns (1985) and by Haykin (1996).

3. INVERSE PLANT MODELING

The plant's controller will be an inverse of the plant. Inverse plant modeling of a linear SISO plant is illustrated in Fig 2. The plant input is its control signal. The plant output, shown in the figure, is the input of an adaptive filter. The desired response for the adaptive filter is the plant input (sometimes delayed by a modeling delay, Δ). Minimizing mean square error causes the adaptive filter $\hat{P}^{-1}$ to be the best least squares inverse to the plant P for the given input spectrum. The adaptive algorithm attempts to make the cascade of plant and adaptive inverse behave like a unit gain. This process is often called deconvolution. With the delay Δ incorporated as shown, the inverse will be a delayed inverse.

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 non-minimum 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 non-causal as well as causal, however, then a two-sided stable inverse would exist for all linear time-invariant plants in accord with the theory of two-sided z-transforms. For useful realization, the two-sided inverse response would need to be delayed by Δ. A causal FIR filter can approximate the delayed version of the two-sided plant inverse. The time span of the adaptive filter (the number of weights multiplied by the sampling period) should be made adequately long, and the delay Δ needs to be chosen appropriately. The choice is generally not critical.

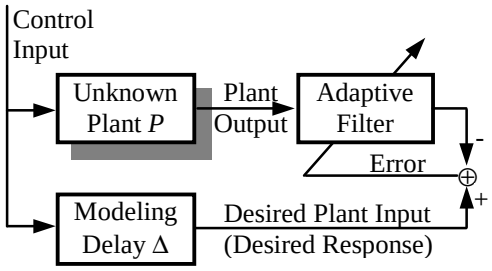

Fig.2. *Delayed inverse modeling of an unknown plant*

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 on the accuracy of the inversion process, which sometimes requires Δ to be of the order of half the length of the adaptive filter.

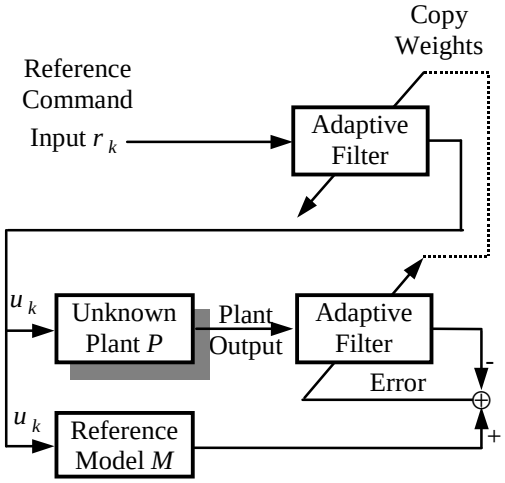

Fig. 3. *Adaptive inverse model control system*

A model-reference inversion process is incorporated in the feedforward control system of Fig. 3. A reference model is used in place of the delay of Fig. 2. Minimizing mean square error with the system of Fig. 3 causes the cascade of the plant and its "model-

reference inverse" to closely approximate the response of the reference-model *M*. Much is known about the design of model-reference systems (Landau, 1979). The model is chosen to give a desirable response for the overall system.

Thus far, the plant has been treated as disturbance free. But, if there is disturbance, the scheme of Fig. 4 can be used. A direct plant modeling process, not shown, yields $\hat{P}$, a close fitting FIR model of the plant. The difference between the plant output and the output of $\hat{P}$ is essentially the plant disturbance.

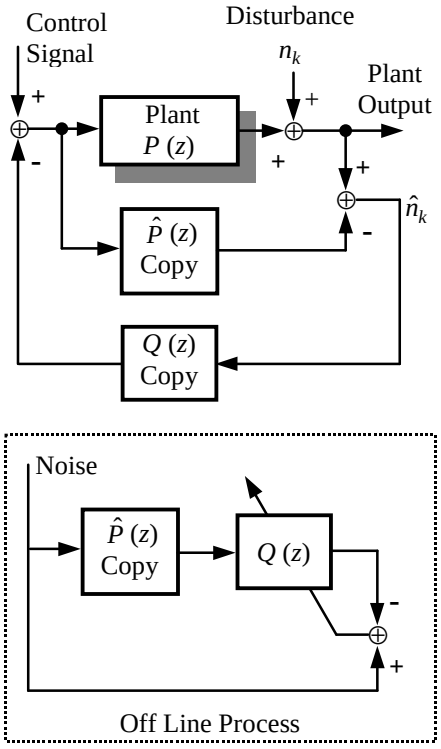

Fig. 4. *Optimal adaptive plant disturbance canceler*

Now, using a digital copy of $\hat{P}$ in place of P, an off line process, shown in Fig. 4, calculates the best least-squares plant inverse Q. The off line process can run much faster than real time, so that as $\hat{P}$ is calculated, the inverse Q is immediately obtained. The disturbance is filtered by digital copy of Q and subtracted from the plant input. For linear systems, the scheme of Fig. 4 has been shown to be optimal in the least-squares sense (Widrow and Walach, 1996).

To illustrate the effectiveness of adaptive inverse control, a non-minimum phase plant has been simulated, and its impulse response is shown in Fig. 5(a). the output of this plant and the output of its reference model are plotted in Fig. 5(b), showing dynamics tracking when the command input signal is a random first-order Markov process. The gray line is the desired output and the black line is the actual plant output. Tracking is quite good. With disturbance added to the plant output, Fig. 5(c) shows the effect of disturbance cancelation. Both the desired and actual plant outputs are plotted in the figure, and they become close when the canceler is turned on, at 300 samplings.

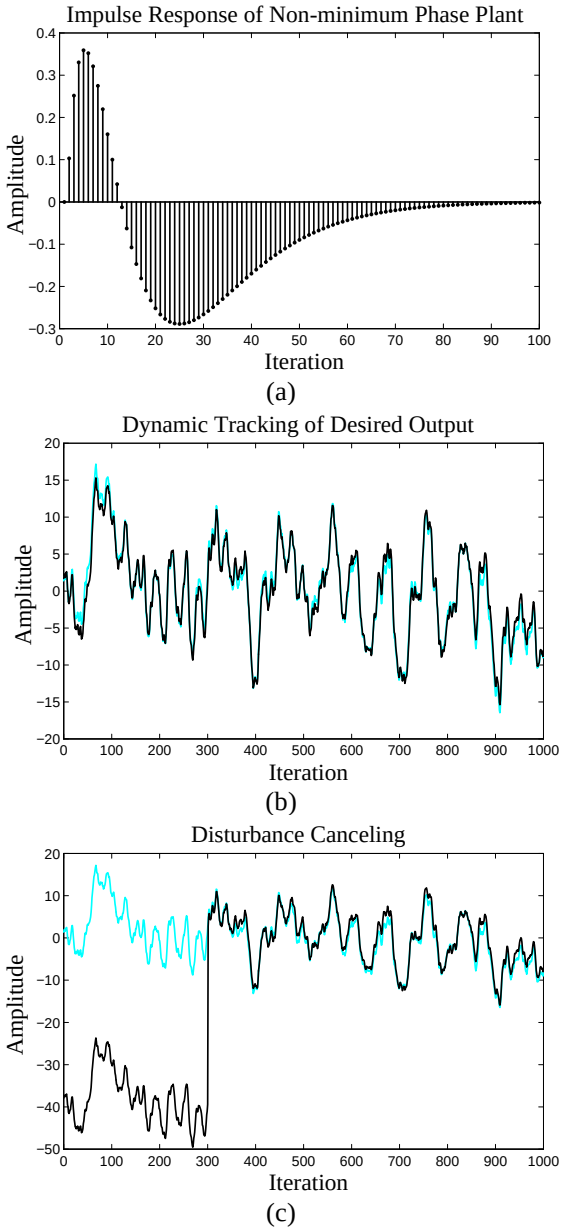

Fig. 5. *(a) Impulse response of the non-minimum phase plant used in simulation; (b) Dynamics tracking of desired output by actual plant output when the plant was not disturbed. (c) Cancelation of plant disturbance.*

4. NONLINEAR ADAPTIVE INVERSE CONTROL WITH NEURAL NETWORKS

Nonlinear inverse controllers can be used to control nonlinear plants. Although the theory is in its infancy, experiments can be done to demonstrate this. A nonlinear adaptive filter is shown in Fig. 6. It is composed of a neural network whose input is a tapped delay line connected to the exogenous input signal. In addition, the input to the network might include a tapped delay line connected to its own output signal. This type of nonlinear filter is called a Nonlinear Auto-Regressive with eXogeneous Input *(NARX)* filter, and has recently been shown to be a universal dynamical system (Siegelmann, et al., 1997). Algorithms such as

real-time-recurrent-learning *(RTRL)* (Williams and Zipser, 1989) and the backpropagation-through-time *(BPTT)* (Werbos, 1990) may be used to adapt the weights of the neural network to minimize the mean squared error. If the feedback connections are omitted, the familiar backpropagation may be used (Werbos, 1974), (Rumelhart, et al., 1986). In the nonlinear adaptive inverse control scheme of Fig.7, such filters are used as the plant emulator and controller.

Nonlinear systems do not commute. Therefore, the simple and intuitive block-diagram method of Figs. 2 and 3, for adapting a controller to be the inverse of the plant, will not work if the plant is nonlinear. Instead, a lower-level mathematical approach is taken. We use an extension of the RTRL learning algorithm to train the controller. This method can be briefly summarized using the notation of ordered derivatives, proposed by Werbos

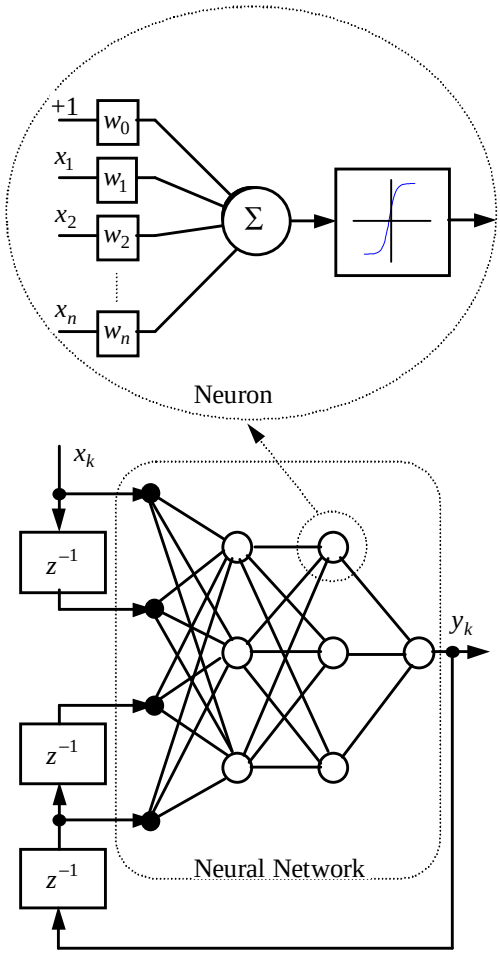

Fig. 6. *An adaptive nonlinear filter composed of a tapped delay line and a three-layer neural network.*

(1974). The goal is to adapt the weights of the controller to minimize the mean squared error of the output of the system. We use the fact that the controller computes a function of the form

$$u_k = g(u_{k-1}, u_{k-2}, ..., u_{k-m}, r_k, r_{k-1}, ..., r_{k-q}, W),$$

where W are the weights of the controller's neural network. We also use the fact that the plant model computes a function of the form

$$y_k = f(y_{k-1}, y_{k-2}, ..., y_{k-n}, u_k, u_{k-1}, ..., u_{k-p}).$$

The weights of the controller are adapted using steepest descent. The change in the weights at each time step is in the negative direction to the gradient of the system error with respect to the weights of the controller. To find the gradient, we use the chain-rule expansion for ordered derivatives

$$\frac{\partial^+ \|e_k\|^2}{\partial W} = -2e_k \frac{\partial^+ y_k}{\partial W}$$

$$\frac{\partial^+ u_k}{\partial W} = \frac{\partial u_k}{\partial W} + \sum_{j=1}^{m} \left(\frac{\partial u_k}{\partial u_{k-j}} \right) \left(\frac{\partial^+ u_{k-j}}{\partial W} \right) \tag{1}$$

$$\frac{\partial^+ y_k}{\partial W} = \sum_{j=0}^{p} \left(\frac{\partial y_k}{\partial u_{k-j}} \right) \left(\frac{\partial^+ u_{k-j}}{\partial W} \right) + \sum_{j=1}^{n} \left(\frac{\partial y_k}{\partial y_{k-1}} \right) \left(\frac{\partial^+ y_{k-j}}{\partial W} \right) \tag{2}$$

Each of the terms in Eqs. (1) and (2) is either a Jacobian matrix, which may be calculated using the *dual-subroutine* (Werbos, 1992) of the backpropagation algorithm, or is previously calculated value of $\partial^+ u_k / \partial W$ or $\partial^+ y_k / \partial W$.

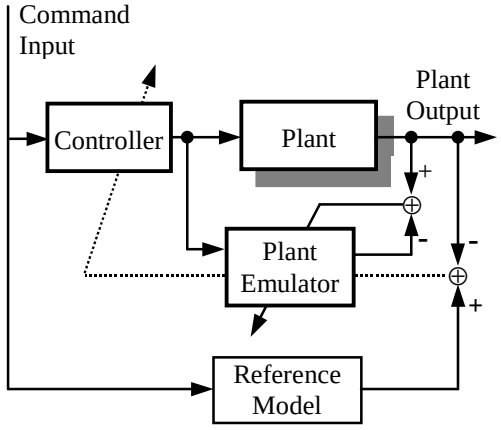

Fig. 7. *A method for adapting a nonlinear controller*

To be more specific, the first term in Eq. (1) is the partial derivative of the controller's output with respect to its weights. This term is one of the Jacobian matrices of the controller and may be calculated with the dual subroutine of the backpropagation algotithm.

The second part of Eq. (1) is a summation. The first term of the summation is the partial derivative of the controller's current output with respect to a previous output. However, since the controller is externally recurrent, this previous output is also a current input.

Therefore the first term of the summation is really just a partial derivative of the output of the controller with respect to one of its inputs. By definition, this is a submatrix of the Jacobian matrix for the network, and may be computed using the dual-subroutine of the backpropagation algorithm.
The second term of the summation in Eq. (1) is the ordered partial derivative of a previous output with respect to the weights of the controller. This term has already been computed in a previous evaluation of Eq. (1), and need not be re-computed.

A similar analysis may be performed to determine all of the terms required to evaluate Eq. (2). After calculating these terms, the weights of the controller may be adapted using the weight-update equation

$$\Delta W_k = 2\mu e_k \frac{\partial^+ y_k}{\partial W}$$

Continual adaptation will minimize the mean squared error at the system output.

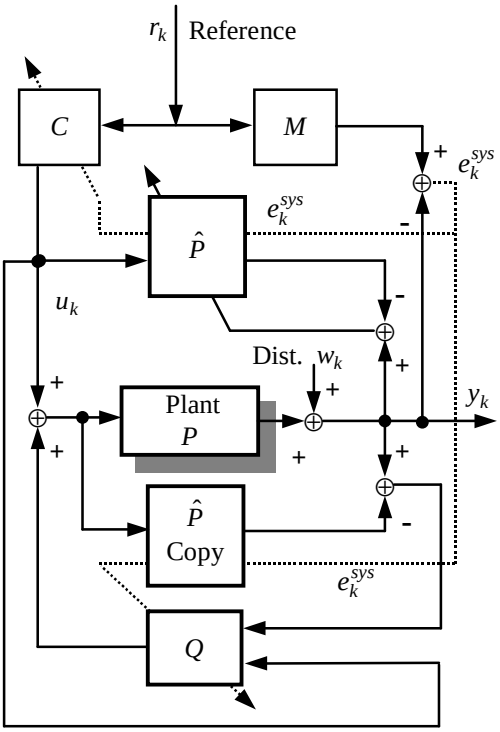

Fig. 8. *A fully integrated nonlinear adaptive inverse control scheme.*

Disturbance canceling for a nonlinear system is performed by filtering an estimate of the disturbance with the nonlinear filter *Q* and adding the filter's output to the control signal. An additional input to *Q* is the control signal to the plant u_k, to allow the disturbance canceler knowledge of the plant state. The same algorithm which was used to adapt the controller can be used to adapt the disturbance canceling filter. The entire control system is shown in Fig. 8.

An interesting discrete-time nonlinear plant has been studied by Narendra and Parthasarathy (1990)

$$y_k = \frac{y_{k-1}}{1+y_{k-1}^2} + u_{k-1}^3 .$$

The method just described for adapting a controller and disturbance canceler were simulated for this plant, and the results are presented here.

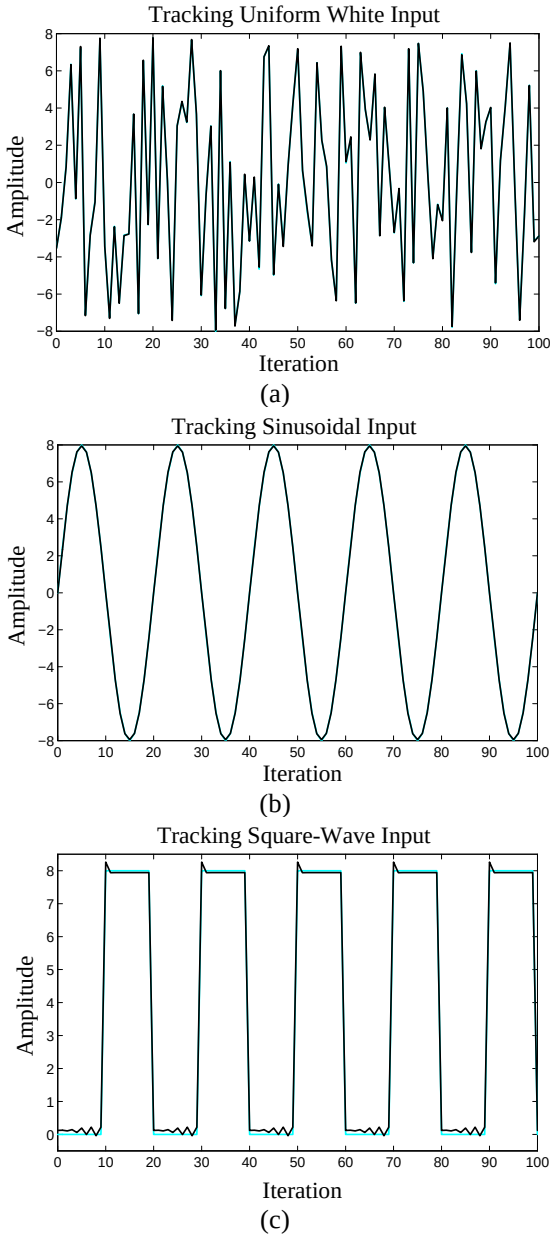

Fig. 9. *Feedforward control of a nonlinear system. The controller was trained with uniform distributed (white) random input (a). Plots (b) and (c) show the plant tracking sinusoidal and square waves, having been previously trained with the random imput.*

With the reference model being a simple delay, and the command input being an i.i.d (independent and identically distributed) uniform process, the system adapted and learned to track the model output. The result is shown in Fig. 9(a). The desired plant output (gray line) and the true plant output (solid line) are shown, at the end of training, when the training signal was used to drive the controller. The gray line is

completely covered by the black line, indicating near-perfect control. With the weights fixed at their trained values, the next two plots show the generalization ability of the controller. After training with the random input, the adaptive process was halted. With no further training, the system was tested with inputs of different character in order to demonstrate the generalization ability of the controller. The first test was a sine-wave command input. Tracking was surprisingly good, as shown in Fig. 9(b). Again, without further training, the system was tested with a square-wave command input, and the results, shown in Fig. 9(c), are excellent.

A disturbance canceler was also trained for this plant, were the disturbance was a first-order Markov signal added to the plant output. Fig.10 shows the results of disturbance cancelation. The power of the system error is plotted versus time. The disturbance canceler was turned on at iteration 500. Dramatic improvement may be seen.

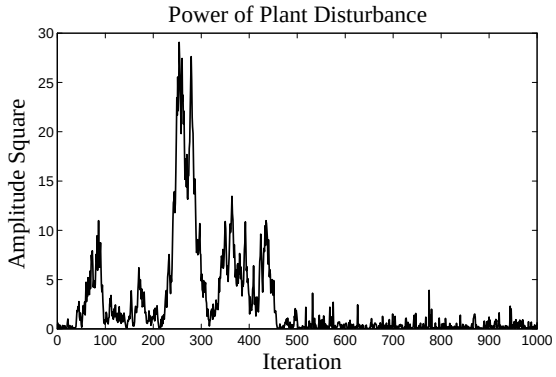

Fig. 10. *Cancelation of plant disturbance for a nonlinear plant. The disturbance canceler was turned on at iteration 500.*

CONCLUSIONS

Adaptive control is seem as a two part problem, (a) a control of plant dynamics, and (b) control of plant disturbance. Conventionally, one uses feedback control to treat both problems simultaneously. Tradeoffs and compromises are necessary to achieve good solutions, however.

The method proposed here, based on inverse control, treats the two problems separately without compromise. The method applies to SISO and MIMO linear plants, and to nonlinear plants.

An unknown linear 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 non-minimum 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 by an adaptive algorithm to adjust the parameters of the controller. Although nonlinear plants do not have transfer functions, the same idea works well for nonlinear plants.

Control of internal plant disturbance is accomplished with an adaptive disturbance canceler. The canceler does not affect plant dynamics, but feeds back plant disturbance in a way that minimizes plant output disturbance power. This approach is optimal for linear plants and works surprisingly well with nonlinear plants.

A great deal of work will be needed to gain greater understanding of this kind of behavior, but the prospects for useful and unusual performance and for development of this new approach seem very promising.

REFERENCES

Haykin, S. (1996). *Adaptive Filter Theory.* Prentice Hall, third edition, Upper Saddle River, NJ.

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

Narendra, K. S. and K. Parthasarathy (1990). Identification and control of dynamical systems using neural networks. *IEEE Transactions on Neural Networks*, **Vol. 1(1)**, March.

Rumelhart, D. E., G. E. Hinton and R. J. Williams (1986). Learning internal representations by error propagation. *Parallel Distributed Processing.* (D. E. Rumelhart and J. L McClelland editors), volume 1, chapter 8. MIT Press, Cambridge, MA.

Siegelmann, H. T., B. B. Horne and C. L. Giles (1997). Computational capabilities of recurrent NARX neural networks. *IEEE Transactions on Systems, Man and Cybernetics — Part B: Cybernetics,* **Vol. 27(2)**, April, pp. 208–215.

Werbos, P. (1974). *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences.* PhD thesis, Harvard University, Cambridge, MA, August.

Werbos, P. (1990). Backpropagation through time: What it does and how to do it. *Proceedings of the IEEE*, **Vol. 78(10)**, October, pp. 1545–1680.

Werbos, P. (1992). Neurocontrol and supervised learning: An overview and evaluation. *Handbook of Intelligent Control: Neural, Fuzzy and Adaptive Approaches.* (D. White and D. Sofge editors), chapter 3. Van Nostrand Reinhold, New York.

Widrow, B. and S. D. Stearns (1985). *Adaptive Signal Processing.* Prentice Hall, Englewood Cliffs, NJ.

Widrow, B. and E. Walach (1996). *Adaptive Inverse Control.* Prentice Hall PTR, Upper Saddle River, NJ.

Williams, R. J. and D. Zipser (1989). Experimental analysis of the real-time recurrent learning algorithm. *Connection science,* **Vol. 1(1),** pp.87–111.