Chapter 1

The Adaptive Inverse Control Concept

1.0 INTRODUCTION

Adaptive filtering techniques have been successfully applied to adaptive antenna systems [1–20]; to communications problems such as channel equalization [21–30] and echo cancelation in long-distance telephony [31–39]; to interference canceling [40–46]; to spectral estimation [47–57]; to speech analysis and synthesis [58–60]; and to many other signal processing problems. It is the purpose of this book to show how adaptive filtering algorithms can be used to achieve adaptive control of unknown and possibly time varying systems.

The system to be controlled, usually called the "plant," may be noisy, that is, subject to disturbances, and for the most part it may be unknown in character.[1] The plant and its internal disturbances may be time variable in an unknown way. In some cases, the plant might even be unstable. Adaptive control systems for such plants would be advantageous over fixed systems since the parameters of adaptive systems can be adjusted or tailored to the unknown and varying requirements of the plant to be controlled. Adaptivity finds a natural area of application in the control field [88].

In the past two decades or so, many hundreds of papers have been published on adaptive control systems in the *Transactions* of the IEEE Control Systems Society, in *Automatica*, in the IFAC (International Federation for Automatic Control) journals and conference proceedings, and elsewhere. At the same time, a very large number of papers on adaptive signal processing and adaptive array processing have appeared in the *Transactions* of the IEEE Signal Processing Society, Antennas and Propagation Society, Communications Society, Circuits and Systems Society, Aerospace and Electronics Society, the *Proceedings of the IEEE*, and elsewhere. Many books have been published on these subjects. The two schools of thought, adaptive controls and adaptive signal processing, have developed almost independently. The control theorists have by and large studied adaptive control using

[1] Some prior knowledge of the character of the plant and its internal disturbances will be needed in order to establish proper control. For example, at least a rough idea of the transient response time of the plant would be required in order to model it adaptively. Some idea of how rapidly the plant characteristics change for plants that vary over time would be needed. Some knowledge of the plant disturbance would be useful, such as disturbance power level at the plant output. Detailed knowledge of the plant and its disturbances would not be required however.

state variable feedback coupled through variable parameter networks to regulate unknown plants and to control their disturbances. The signal processing people have been working on problems that for the most part involve adapting weights of transversal filters by gradient methods and employing the resultant adaptive filters to systems without feedback (except for feedback in the adaptive process itself). The signal processing people have found a great number of practical applications for their work, and so have the adaptive control people.

The goal of this book is not to bridge the gap between these two schools of thought but to attack certain problems in adaptive control from an alternative point of view using the methodology of adaptive signal processing. The result is what we call "adaptive inverse control."

We begin with a discussion of direct modeling (or identifying) the characteristics of the unknown plant using simple adaptive filtering methods. Then we show how similar methods, with some modification but in a different configuration, can be used for inverse modeling (or equalization or deconvolution). Inverse plant models can be used to control plant dynamics. Next we show how both direct and inverse models can be used in the same adaptive process to minimize the effects of plant disturbance. In this development, we assume that the plant is completely controllable and observable, that it can (in a quasistatic sense) be represented in terms of an input-output transfer function (albeit an unknown one), and that the plant is stable (if unstable, someone has previously applied stabilization feedback). The plant may be either minimum-phase or nonminimum-phase.

The basic ideas of adaptive inverse control have been under development at Stanford University over the course of many years. The earliest related work is described in a paper by Widrow on blood pressure regulation [61]. Subsequent work is reported in several papers that were presented at Asilomar conferences [62, 63]. A Ph.D. dissertation by Shmuel Schaffer was concerned with model-reference adaptive inverse control [64]. A tutorial on the work is given by Widrow and Stearns [65]. The first paper on adaptive inverse control including adaptive plant disturbance canceling was presented by Widrow and Walach in 1983 at the First IFAC Workshop in Control and Signal Processing in San Francisco [66]. The second presentation was by Widrow in 1986 in a keynote talk at the Second IFAC Workshop on Adaptive Systems in Control and Signal Processing, University of Lund, Sweden [67]. There have been almost no other publications on inverse control and disturbance canceling until recently. Several recent publications in the neural network literature have appeared concerning nonlinear adaptive inverse control [95, 96, 97].

1.1 INVERSE CONTROL

A conventional control system like the one illustrated in Fig. 1.1 uses feedback, sensing the response of the plant to be controlled, comparing this response to a desired response, and using the difference to excite an actuator or controller whose output drives the plant input to cause the plant output to follow the desired response more closely.

The system of Fig. 1.1 has unity feedback and is often called a *follow-up* system since the objective is that the plant output follow the input signal or the *command input*. Any difference between the plant output and the command input signal is an *error signal* sensed by the controller which amplifies and filters it to drive the plant to reduce the error.

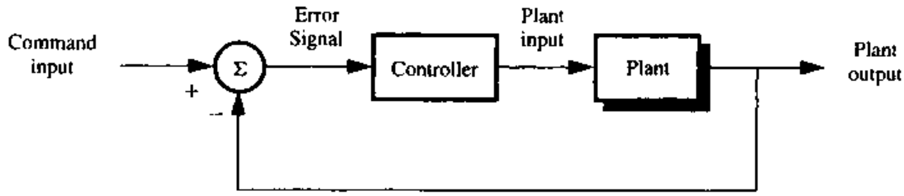

Figure 1.1 A conventional feedback control system.

The use of feedback must be done in a careful way to prevent instability and to achieve satisfactory dynamic response. When the plant characteristics are time variable or nonstationary, it is sometimes necessary to design the controller to vary with the plant. A common objective in doing this would be to minimize the mean square of the error. But achieving this objective is generally difficult. If one knew the plant characteristics versus time, one might be able to determine the best controller versus time. Not knowing the plant, an identification process could be used to estimate plant characteristics over time, and these characteristics could be used to determine the controller over time. Another idea would be to parametize the controller and vary the parameters to directly minimize mean square error. The difficulty with this approach is that, regardless of how the controller is parametized, the mean square error versus the parameter values would be a function not having a unique extremum and one that could easily become infinite if the controller parameters were pushed beyond the brink of stability.

The objective of the present work is to take an alternative look at the subject of adaptive control. The approach to be developed, adaptive inverse control, in some sense involves open-loop control and it is quite different from the feedback-control approach in Fig. 1.1. We attempt to develop a form of adaptive control that is simple, robust, and precise. With some knowledge of the subject of adaptive filtering, adaptive inverse control is easy to understand and use in practice.

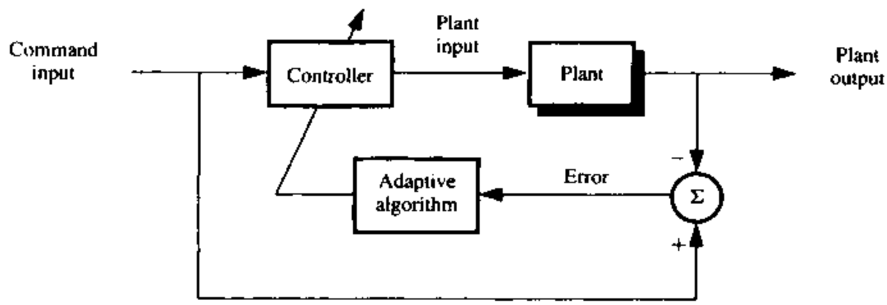

Figure 1.2 Basic concept of adaptive inverse control.

The basic idea of adaptive inverse control is to drive the plant with a signal from a controller whose transfer function is the inverse of that of the plant itself. The idea is illustrated with the system of Fig. 1.2. The objective of this system is to cause the plant output to follow the command input. Since the plant is generally unknown, it is necessary to adapt or to adjust the parameters of the controller in order to create a true plant inverse. An error signal, the difference between the plant output and the command input, is used by an adaptive algorithm to adjust the controller's parameters to minimize the mean square of this error.

Referring to Fig. 1.2, the controller in this diagram can be thought of as a filter having an input and an output. This controller has adjustable parameters, and adjustability is indicated by the arrow going through the box. Control of the adjustable parameters is done by means of an "adaptive algorithm," which is driven by the error signal. A usual objective for the adaptive algorithm would be minimization of mean square error, the error being the difference between the plant output and the command input.

Comparing the system of Fig. 1.1 with that of Fig. 1.2, *minimizing mean square error is done in the first case by using the error signal directly in a feedback process for forming the plant input signal, whereas in the second case the error signal is used in a feedback process to control the parameters of the controller and is not fed back directly to the plant input. The first case is feedback control and the second case is feedforward control. In both cases, feedback is used to ensure precise system responses.*

If adapting the controller of Fig. 1.2 were to make the error small, the controller would have become an inverse of the plant. The cascade of the controller and plant would thereby have a combined transfer function matching a gain of unity.

We assume that the plant is linear and that it varies slowly so that it is quasistatically stationary. We assume that the controller has converged and that it too is linear and quasistatically stationary. The dynamic characteristics of both the plant and controller may be represented by transfer functions, and the transfer function of the controller would be the reciprocal of that of the plant.

If the plant has internal delay, the inverse controller may have difficulty in overcoming it. The controller would need to be a predictor. Furthermore, if the plant is nonminimum-phase (transfer function zeros in the right half of the s-plane or outside the unit circle in the z-plane), then the inverse controller would want to have poles in the right half of the s-plane or outside the unit circle in the z-plane. Such an inverse would normally be unstable. Means for overcoming these difficulties are developed in Chapter 5.

The system of Fig. 1.2 illustrates the basic concept of adaptive inverse control. How the adaptive algorithm works to control the parameters of the controller is a matter to be developed below.

Sometimes it is desired that the plant output track not the command input itself but a delayed or smoothed version of the command input. The system designer would generally know the smoothing characteristic to be used. A smoothing model can be readily incorporated into the adaptive inverse control concept. How this may be done is illustrated in Fig. 1.3. The smoothing model is generally designated as the *reference model* in the control theory field. Thus, the system of Fig. 1.3 may be called a *model-reference* adaptive inverse control system. The model-reference idea is due to Whitaker, and an early reference is [89]. A recent reference is [88].

The reference model is chosen to have the same dynamic response that the designer would like for the entire system. Referring to Fig. 1.3, it is evident that this result would be obtained by once again adapting the controller to cause the mean square error to be low. In this case, the cascade of the controller and the plant after convergence would have a dynamic response like that of the reference model. The product of the controller and plant transfer functions would closely approximate the transfer function of the reference model.

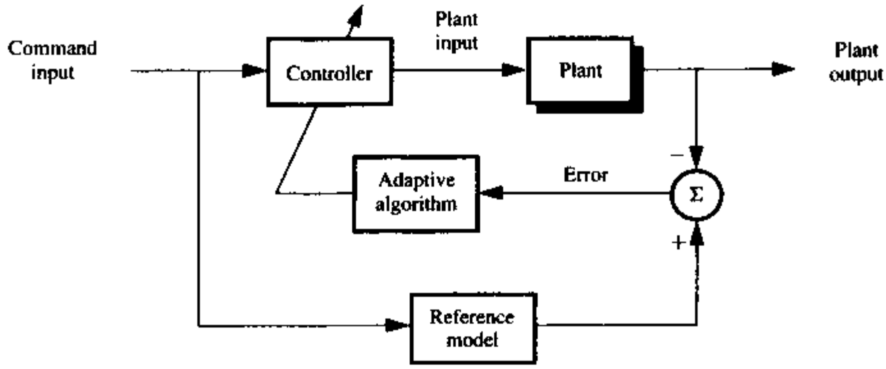

Figure 1.3 Model-reference adaptive inverse control.

Plant noise and disturbance present a problem for the adaptive inverse control approach. Lack of feedback from the plant output back to the plant input permits internal plant noise and disturbance to exist unchecked at the plant output. Various signal processing methods for noise canceling have been developed [40–46], and with some modification they have been applied to the cancelation of plant noise and disturbance. The basic scheme is shown in Fig. 1.4. Use is made of both a plant model and a plant inverse model. The plant model has the same transfer function as the plant, while the plant inverse model has a transfer function which is the reciprocal of that of the plant. How these models can be obtained will be described below. For the present argument, we assume that these models exist.

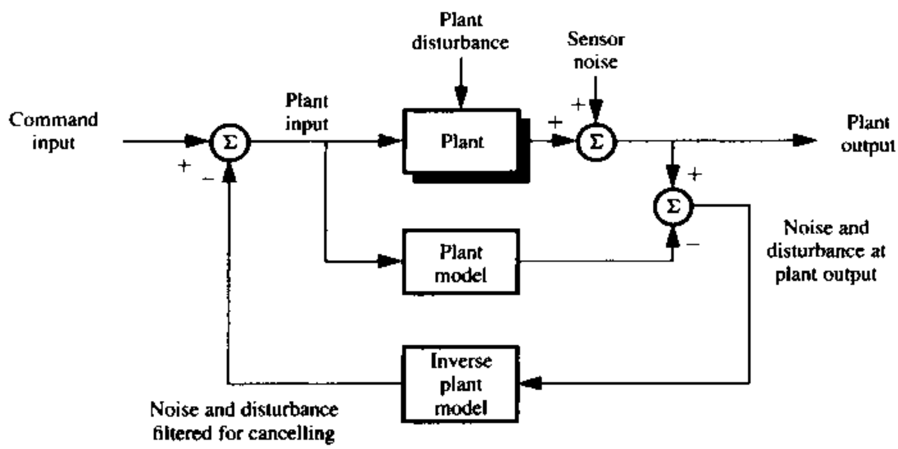

Figure 1.4 Canceling plant noise and disturbance.

In the control literature, plant disturbance is often represented as an additive noise at the plant input. Sensing the plant output is done with a detector or sensor that may be noisy. Sensor noise is often represented as an additive noise at the plant output. In the system of Fig. 1.4, the plant noise and disturbance are separated from the plant's dynamic output response. The plant input drives both the plant and its model (which is free of noise and disturbance). The difference between the plant output and the plant model output is the plant noise and disturbance as they appear at the plant output. The sum of the plant output noise and disturbance is used to drive the inverse plant model to generate filtered noise and disturbance for subtraction from the plant input. The ultimate effect is to cancel noise and disturbance at the plant output [67].

With almost perfect direct and inverse models, one can show that the transfer function from the point of injection of the plant sensor noise to the plant output point is close to zero. This implies that the sum of the plant noise and plant disturbance will be highly attenuated at the plant output. One can show furthermore that the dynamic response of the plant is essentially unchanged even while the plant's noise and disturbance are canceled by the disturbance canceling feedback. It is interesting to note that with perfect direct and inverse models, the system of Fig. 1.4 will be an unusual feedback system. The feedforward link of the major loop will have zero gain when the dynamics of the plant model perfectly balances that of the plant itself.

An adaptive inverse control system including the plant noise and disturbance canceling features of Fig. 1.4 and the model-reference control features of Fig. 1.3 are shown in Fig. 1.5. In a practical system of this type, separate adaptive processes would be needed to obtain the plant model, the inverse plant model, and the controller.

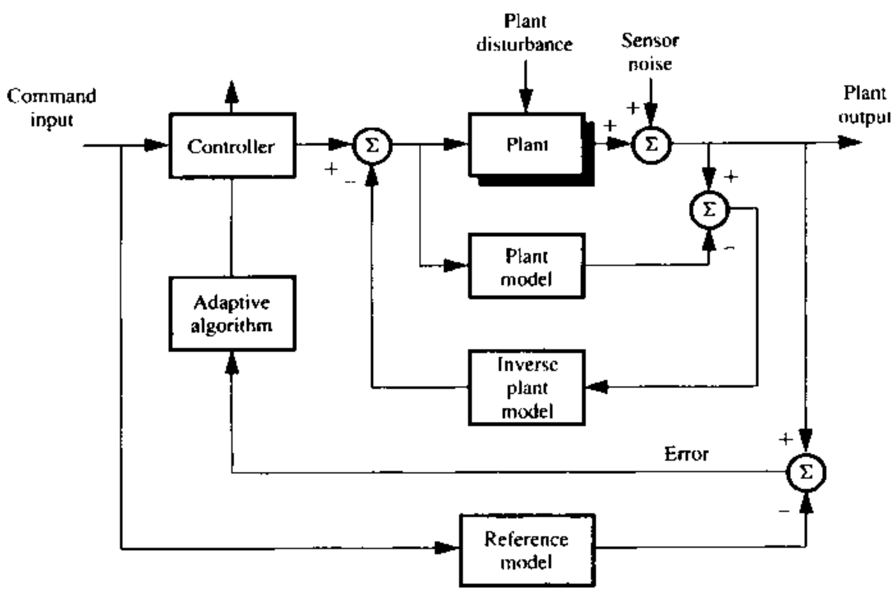

Figure 1.5 Model-reference adaptive inverse control system with plant noise and disturbance canceling.

With regard to ordinary control systems, like the one shown in Fig. 1.1, control of plant noise and disturbance is done with feedback. But by incorporating feedback for noise and disturbance control, the plant dynamics are inevitably altered. A compromise is generally required in the design process to obtain good dynamic response and good noise and disturbance control all at the same time.

Our approach is a different one in that it involves separate adaptive processes for (a) the control of plant dynamics and (b) for the control of plant noise and disturbance. By handling the problems in this way, the corresponding adaptive subsystems are relatively simple, easy to analyze, and easy to optimize.

The system of Fig. 1.5 represents our approach to adaptive control in an overall way. In the remainder of this text, we develop a variety of practical adaptive algorithms for implementation of the various subsystems of Fig. 1.5, and we analyze their performance and study how they interact in the overall control context. Simulation results are presented to verify theory and to demonstrate workability.

1.2 SAMPLE APPLICATIONS OF ADAPTIVE INVERSE CONTROL

The objective of this section is to illustrate the application of adaptive inverse control to a variety of control problems. These problems are in some sense classic in nature, and they will be discussed from many points of view in the chapters to follow. Here we see without much explanation results of control by adaptive inverse control techniques applied to a set of exemplary problems.

1.2.1 Dynamic Control of a Minimum-Phase Plant

The plant to be controlled has the transfer function

$$\frac{s+0.5}{(s+1)(s-1)}. \tag{1.1}$$

This plant is minimum-phase and is unstable. The first step is to stabilize it with feedback. A root-locus diagram is shown in Fig. 1.6. It is clear from this diagram that the plant can be stabilized by making use of the simple unity feedback system of Fig. 1.7, by setting the loop gain within the stable range $\infty > k > 2$. The loop gain was set to $k = 4$ for this control experiment. The closed loop transfer function is minimum-phase and has two poles in the left half of the s-plane.

The plant and its stabilization are continuous (analog) systems. The adaptive inverse control part, as it would be in the real world, is discrete (digital). A diagram of the complete system is shown in Fig. 1.8, including the necessary analog-to-digital conversion (ADC) and digital-to-analog conversion (DAC) components. The command input is sampled and is fed to both the inverse controller and the reference model. The controller output is converted to analog form, using a zero-order hold, to drive the plant and its stabilization loop. The error signal used to adapt the inverse controller is discrete. This is the difference between the reference model output and the sampled plant output. The system of Fig. 1.8 is modeled after that of Fig. 1.3.

Referring to Fig. 1.8, we define the *discretized equivalent plant* to be the discrete transfer function from the digital input to the DAC, to the sampled output of the plant. This includes the plant and its stabilization loop, shown in Fig. 1.7, and the DAC, and the sampler at the plant output. The impulse response of the discretized equivalent plant is shown in Fig. 1.9. The chosen sampling rate was 10 samples per second.

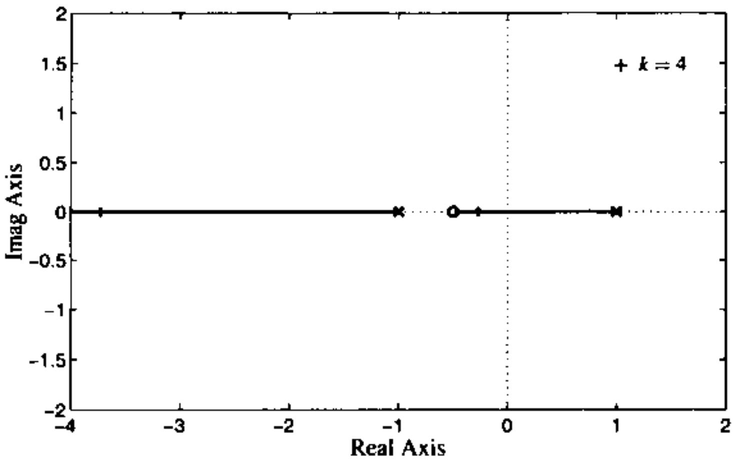

Figure 1.6 Root-locus of minimum-phase plant with proportional feedback.

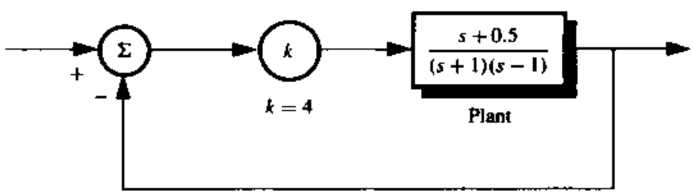

Figure 1.7 Minimum-phase plant stabilized with proportional feedback, $k = 4$.

The reference model chosen for this experiment was a one-pole digital filter with a one-second time constant. Its impulse response is shown in Fig. 1.10. The objective was to cause the overall system response, from the samples of the command input to the samples of the plant output, to be as close to the response of the reference model as possible in the least squares sense. The sampled command input was a first-order Markov process, generated by filtering white noise with a one-pole one-second time constant digital filter. The controller was allowed to have 100 weights, and the theoretically optimal impulse response for that many weights is shown in Fig. 1.11. The learned impulse response of the inverse controller is shown in Fig. 1.12. Notice the similarity that it has to the optimal impulse response.

It is useful to compare the plant output to the reference model output. At the beginning, the inverse controller is learning and not yet performing well. The plant output and the reference model output are shown over the first 200 samples in Fig. 1.13. They do not track. It is too early in the process.

The effects of learning can be observed in Fig. 1.14. The plant output begins to track properly after about 3,000 samples. This corresponds to about 300 seconds in real time. The entire sequence has 100,000 samples, corresponding to 10,000 seconds in real time. A comparison of the two outputs over the last 200 samples of the sequence is shown in Fig. 1.15.

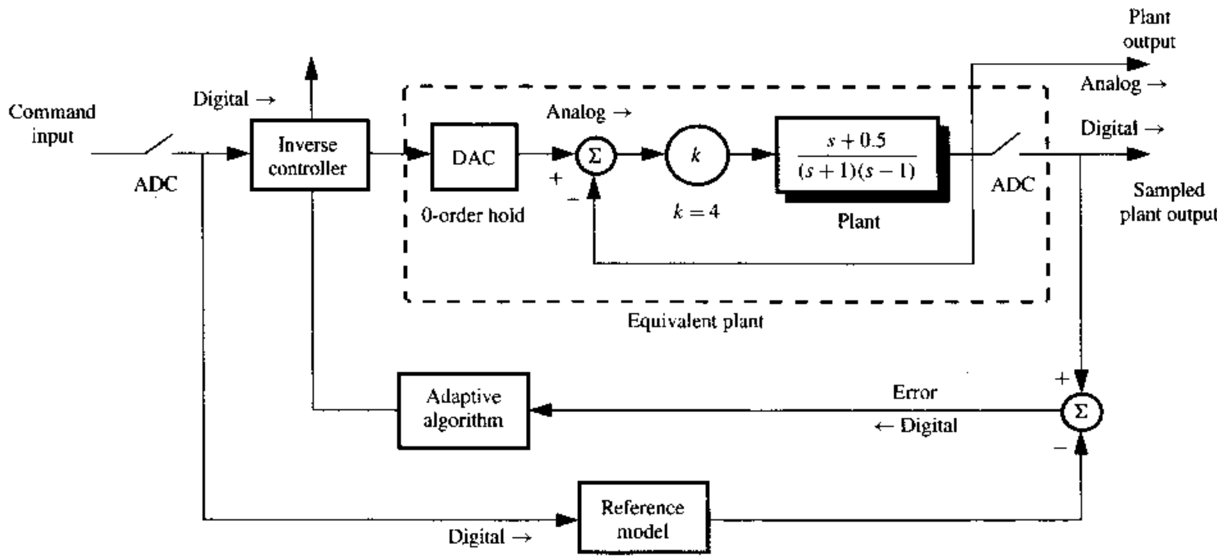

Figure 1.8 Model-reference adaptive inverse control of a stabilized minimum-phase plant.

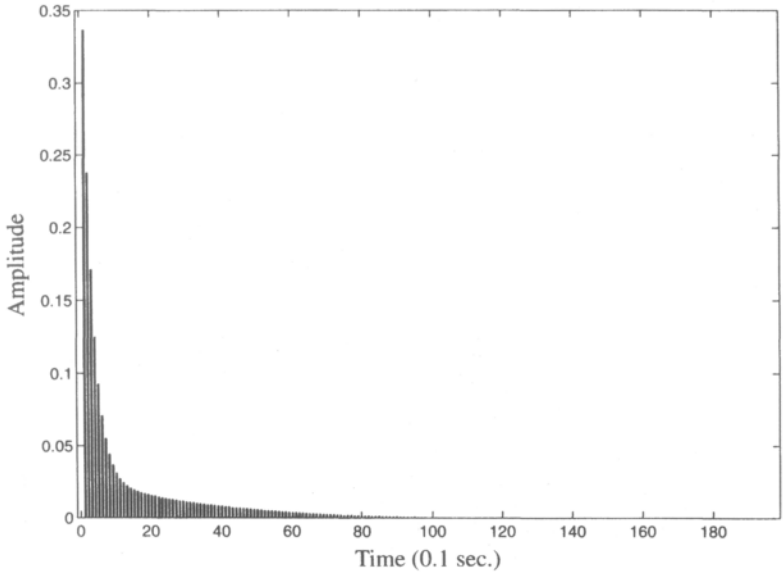

Figure 1.9 Impulse response of discretized equivalent minimum-phase plant.

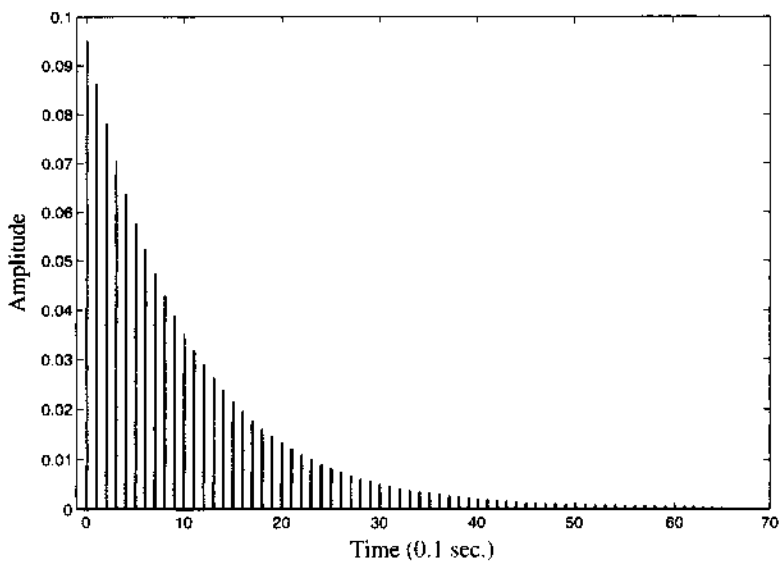

Figure 1.10 Impulse response of reference model.

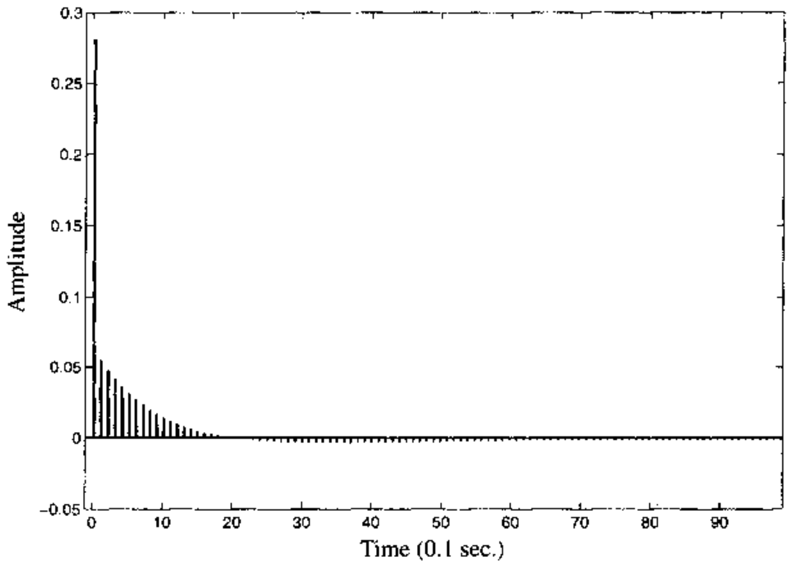

Figure 1.11 Impulse response of optimal 100-weight inverse controller.

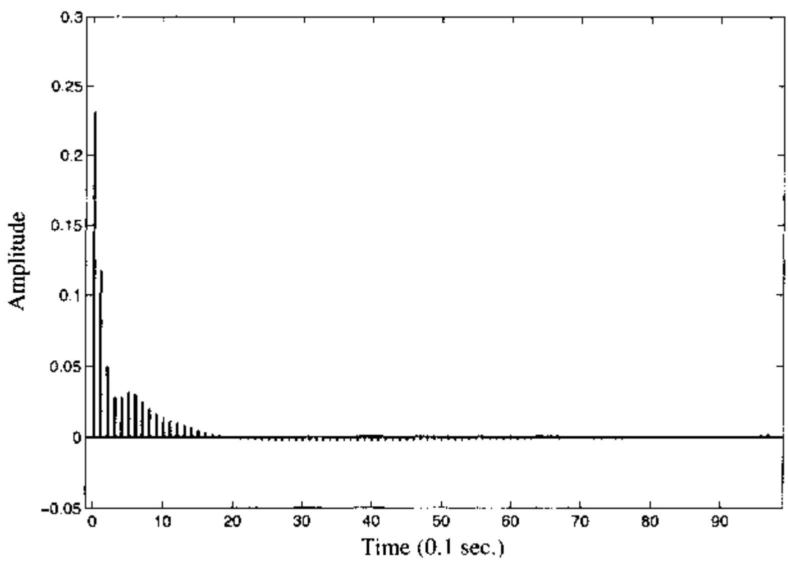

Figure 1.12 Impulse response of the adaptive controller after 100,000 learning iterations.

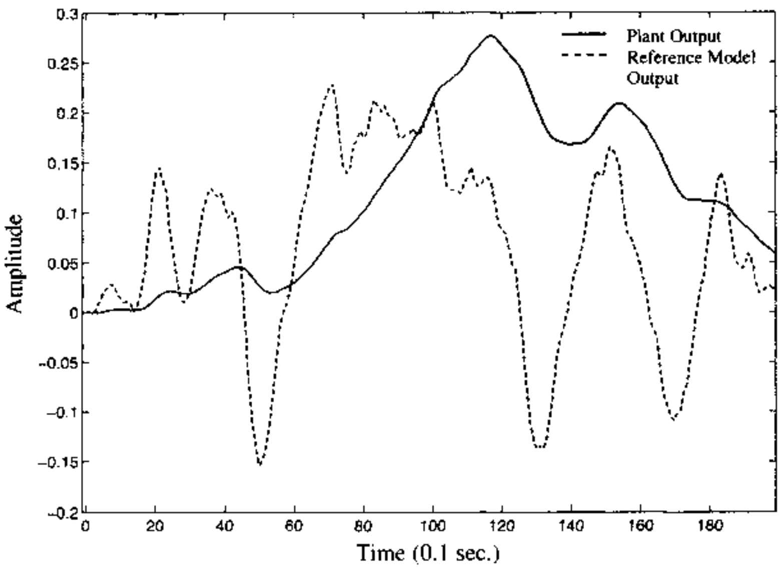

Figure 1.13 Plant output and reference model output during first 200 samples with minimum-phase plant.

Tracking can be seen to be essentially perfect. Model-reference inverse control does indeed work in this case.

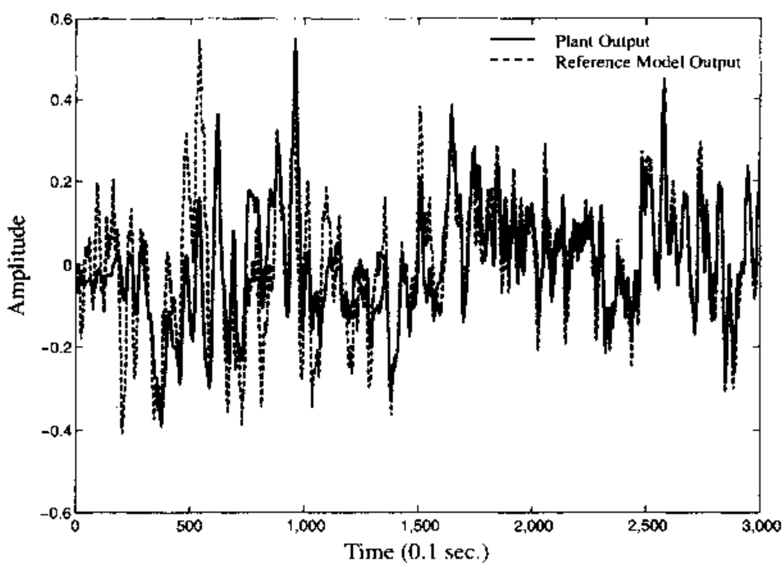

Figure 1.14 Plant output and reference model output over sequence of 3,000 samples, demonstrating controller learning.

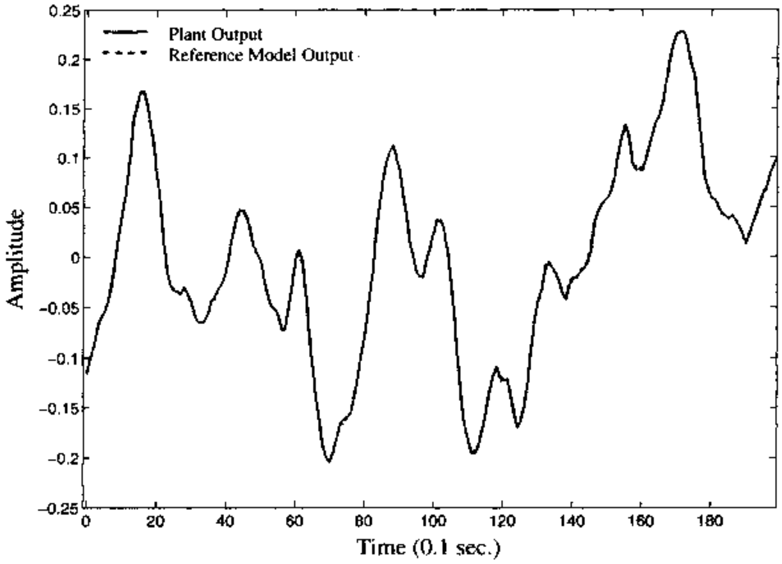

Figure 1.15 Plant output and reference model output during last 200 samples of sequence of 100,000 samples.

1.2.2 Dynamic Control of a Nonminimum-Phase Plant

For this experiment, a more challenging one, the plant to be controlled has the transfer function

$$\frac{(s-0.5)}{(s+1)(s-1)}. \tag{1.2}$$

Since this plant is unstable, the first step is once again to stabilize it with feedback. This plant cannot be stabilized with simple proportional feedback, like that of Fig. 1.7, but it can be stabilized by using feedback with a compensating network. The compensating network used in this experiment had the transfer function

$$\frac{(s+1)}{(s+7)(s-2)}. \tag{1.3}$$

A root-locus plot is shown in Fig. 1.16, and the stabilization feedback diagram is shown in Fig. 1.17. The compensating network has a zero at $s = -1$, and poles at $s = 2$ and -7. The stable range for the loop gain is $\infty > k > 20$. For this experiment, k was chosen to be $k = 24$. The result is a stabilized plant that remains nonminimum-phase.

1.2.3 Canceling Disturbance in the Minimum-Phase Plant

The stabilized plant was incorporated into an adaptive inverse control system like the one shown in Fig. 1.8. The sampling rate was chosen to be 10 samples per second. The reference model was a one-pole digital filter with a one-second time constant, but in this case the reference model included an eight-second delay. The impulse response of the reference

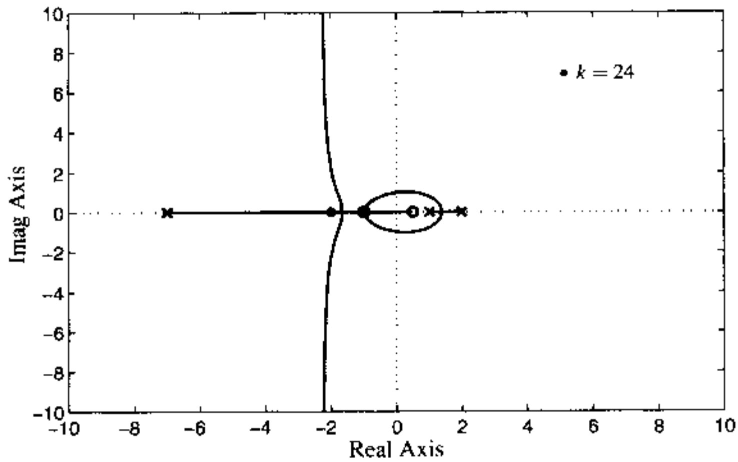

Figure 1.16 Root-locus plot for nonminimum-phase plant with compensating network.

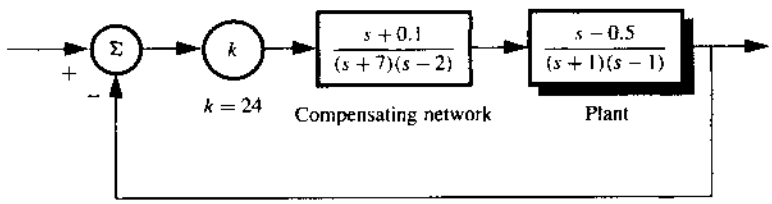

Figure 1.17 Nonminimum-phase plant with compensating network stabilization. $k = 24$.

model including the delay is shown in Fig. 1.18. The delay is necessary for inverse control of a nonminimum-phase plant. This will be explained in Chapter 5.

The impulse response of the discretized stabilized plant, which includes the plant and its stabilizer (shown in Fig. 1.17), and the DAC and the sampler at the plant output, are shown in Fig. 1.19. This is the impulse response from the input to the DAC to the sampled output of the plant.

The impulse response of the converged inverse controller is shown in Fig. 1.20. This controller has 150 weights. It was obtained from a learning algorithm. The form and shape of this impulse response could not be inferred intuitively.

To observe and demonstrate workability, we compare the plant output with the reference model output when the entire system is driven by a command input. For this experiment, the command input was a first-order Markov process generated by applying white noise to a one-pole digital filter with a one-second time constant. Figure 1.21 shows plots of plant output and reference model output for the first 200 samples (the first 20 seconds) of a training sequence. Figure 1.22 shows 1,000 samples of this sequence. Learning can be observed with good tracking after about 500 samples, corresponding to 50 seconds of operation. The last 200 samples of the 20,000-sample sequence are shown in Fig. 1.23. The

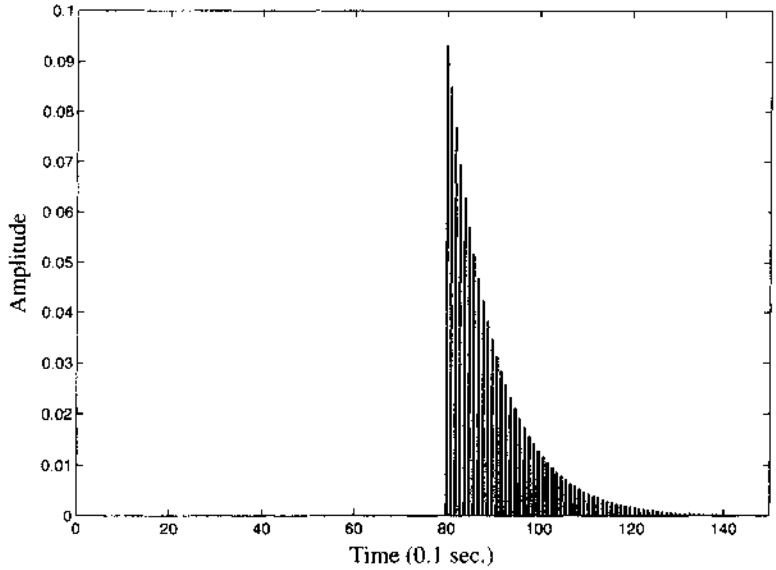

Figure 1.18 Impulse response of reference model, including a delay of 8 seconds.

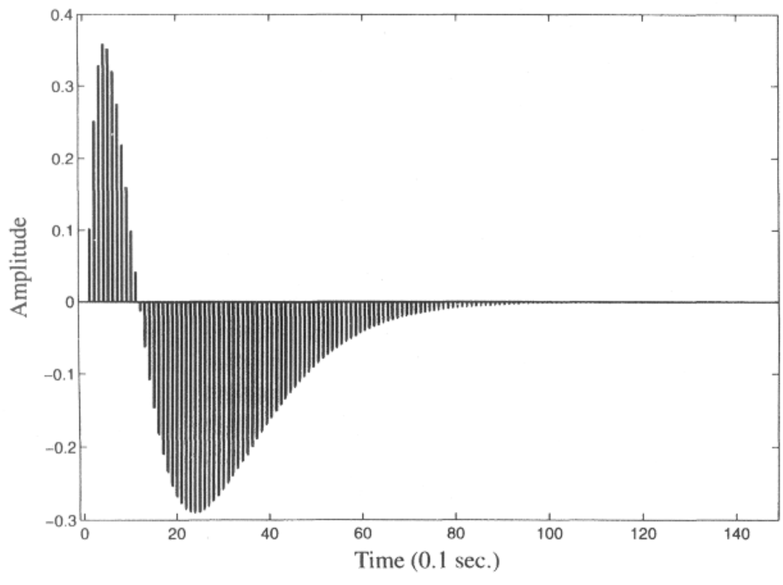

Figure 1.19 Impulse response of discretized stabilized nonminimum-phase plant with $k = 24$.

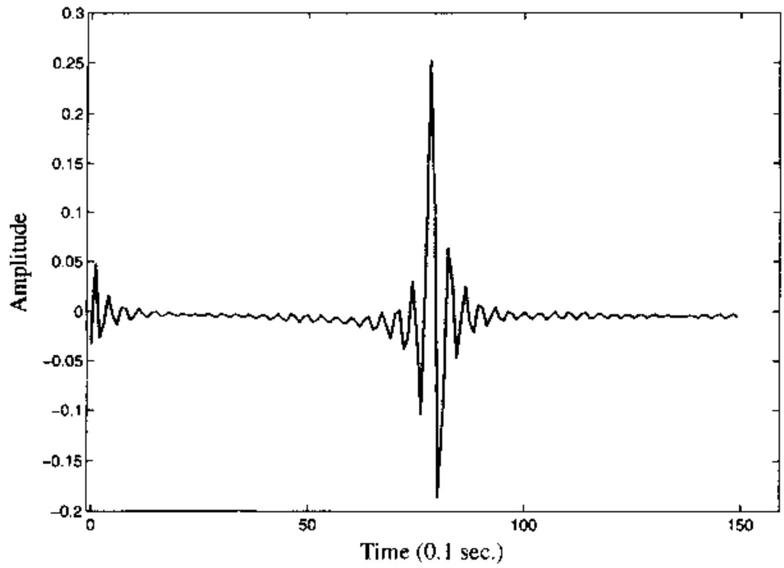

Figure 1.20 Converged impulse response of adaptive inverse controller for nonminimum-phase plant.

tracking is excellent, demonstrating precise control of a nonminimum-phase plant by means of adaptive inverse control, once learning has taken place.

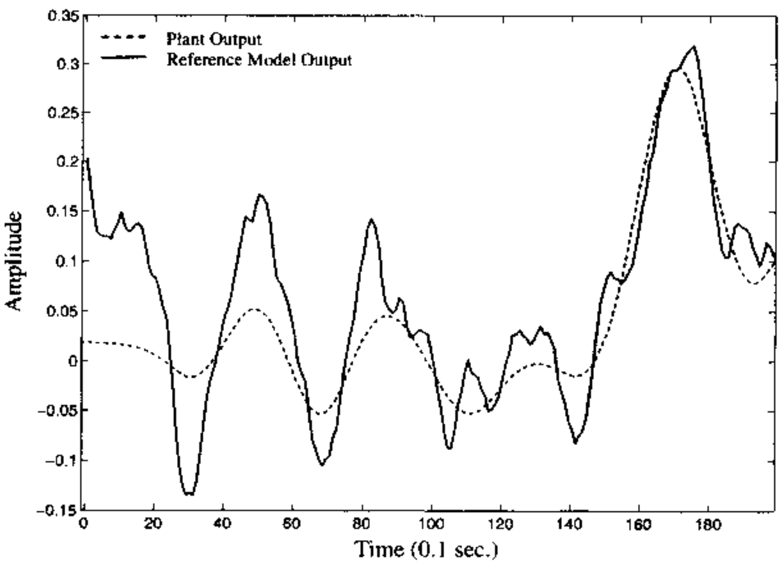

Figure 1.21 Plant output and reference model output during first 200 samples with nonminimum-phase plant.

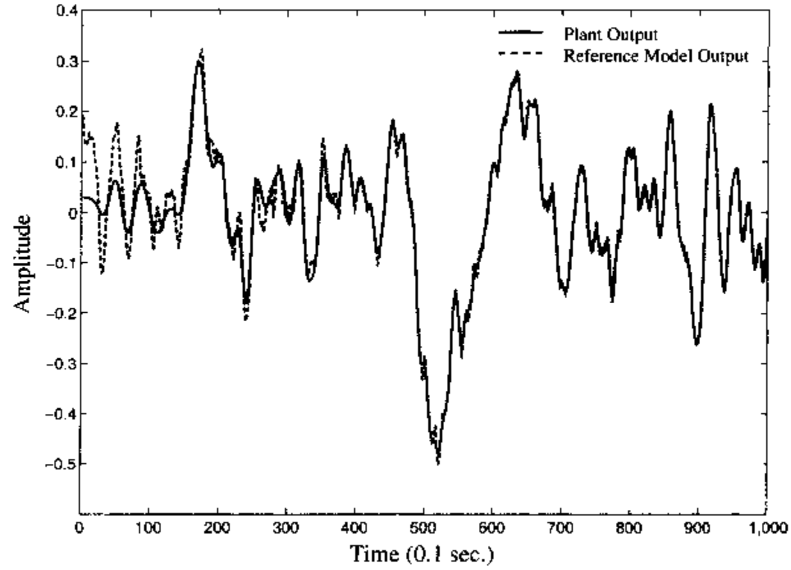

Figure 1.22 Plant output and reference model output during sequence of 1,000 samples, demonstrating learning with a nonminimum-phase plant.

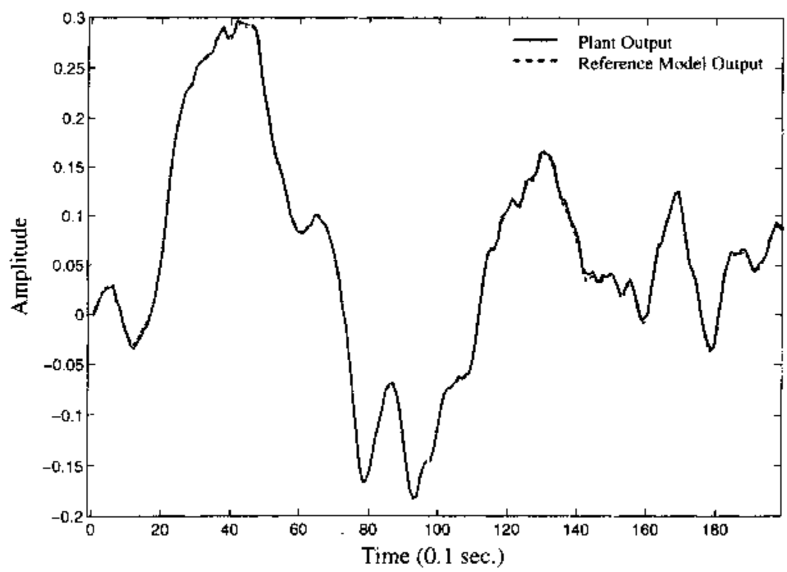

Figure 1.23 Plant output and reference model output during last 200 samples of training sequence with nonminimum-phase plant.

1.2.4 Canceling Disturbance in the Minimum-Phase Plant

Since both plant disturbance and sensor noise show up at the plant output, it is convenient for our purposes to lump them together and to simply call the combined effect plant disturbance, represented as an additive disturbance at the plant output. We shall do this throughout the book.

The minimum-phase plant described by (1.1) was stabilized in accord with the block diagram of Fig. 1.7. Noise is now injected into the plant output to simulate plant disturbance. An adaptive noise canceler fashioned after the one of Fig. 1.4 is used in this experiment to minimize the disturbance appearing at the plant output. The stabilized plant and its noise canceler are integrated into a model-reference adaptive inverse control system like the one diagrammed in Fig. 1.5. The resulting system is shown in Fig. 1.24.

Referring to Fig. 1.24, an adaptive model was made of the equivalent plant, and the converged impulse response is shown in Fig. 1.25. An adaptive inverse model of the equivalent plant was made for use in the disturbance canceler, and its converged impulse response is shown in Fig. 1.26.

The inverse controller was adapted, and its converged impulse response is shown in Fig. 1.27. It is useful to note that although this controller was adapted to control a noisy equivalent plant, and the learning process took place in the presence of plant disturbance, the converged dynamics of the controller were essentially unaffected by the plant disturbance. The controller impulse response of Fig. 1.12 was adapted for the same equivalent plant, but without plant disturbance and without an adaptive disturbance canceler. Comparison of Figs. 1.12 and 1.27 shows that both impulse responses are almost the same. The separation of the functions of *dynamic control* and *plant disturbance control* is a subject to be discussed in detail in Chapters 8 and 9. This separation is a characteristic of the adaptive inverse control approach.

The effectiveness of the adaptive disturbance canceler can be assessed from an inspection of Fig. 1.28. The plant output disturbance (which is computed as the difference between the plant output and what the plant output would be if there were no disturbance) is squared at each sample time and plotted in Fig. 1.28. There is no averaging in this plot, as the squared values are plotted over time. The disturbance canceling feedback loop was open until the five-thousandth sample in the time sequence. During this epoch, there was more than enough time for adaptive modeling of the equivalent plant and adaptive inverse modeling of this equivalent plant. Both models were required by the disturbance canceler. Then the loop was closed and the immediate noise reduction became apparent. It will be shown in Chapter 8 that this form of disturbance canceler is optimal in the least squares sense even when the plant is under dynamic control.

The response of the overall system is exemplified by the plots of Figs. 1.29–1.32. Figure 1.29 shows a comparison of the outputs of the reference model and the plant at the beginning of the learning sequence. The controller has not yet converged, and this causes overall error. Also, the disturbance canceler has not been turned on yet. Plant output disturbance contributes further to the overall error. Figure 1.30 shows the entire sequence of 10,000 error samples. The controller has learned its function after the first several hundred samples. One can see the results of convergence of the controller at the beginning of the plot. Error persists, however, because of the plant disturbance. The disturbance canceler is turned on at 5,000 samples, and then the error is reduced to its lowest possible level. In Fig. 1.31, the

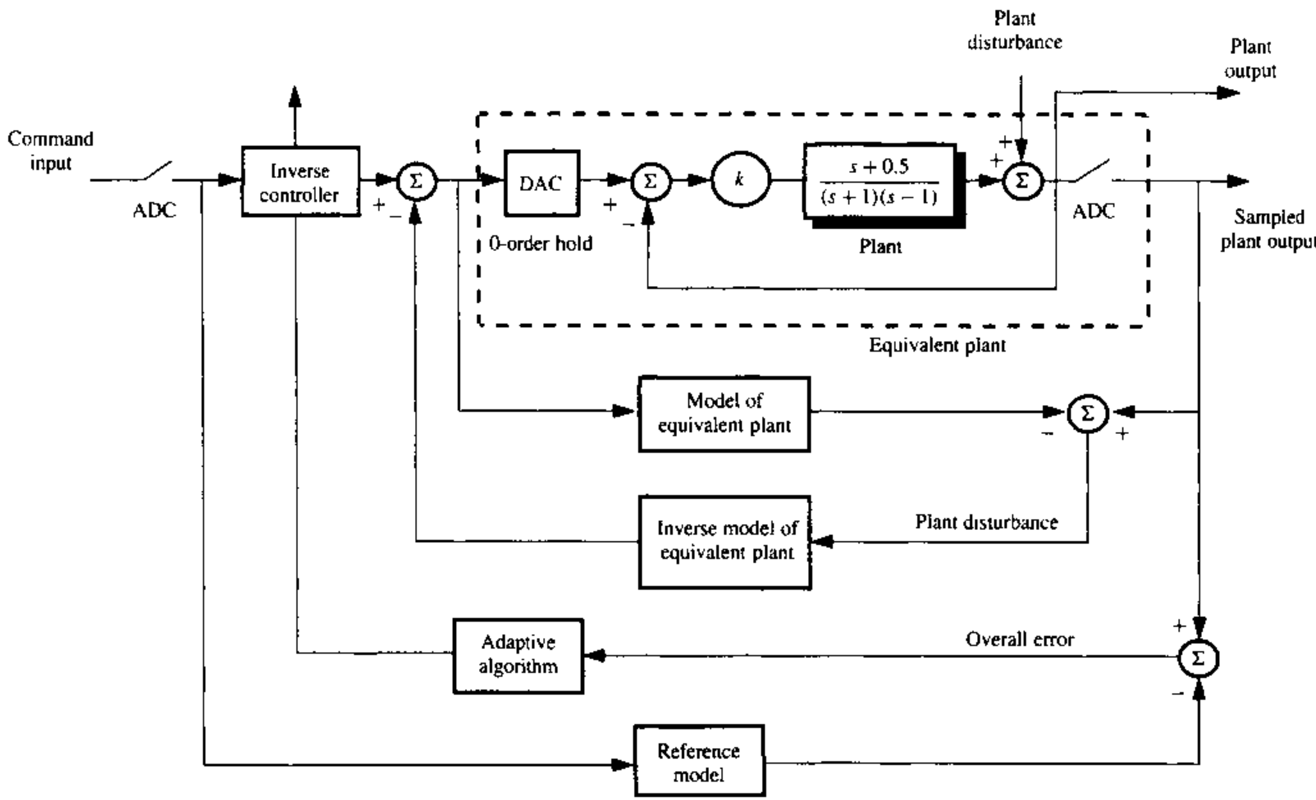

Figure 1.24 Model-reference inverse control system for minimum-phase plant with adaptive plant disturbance canceler.

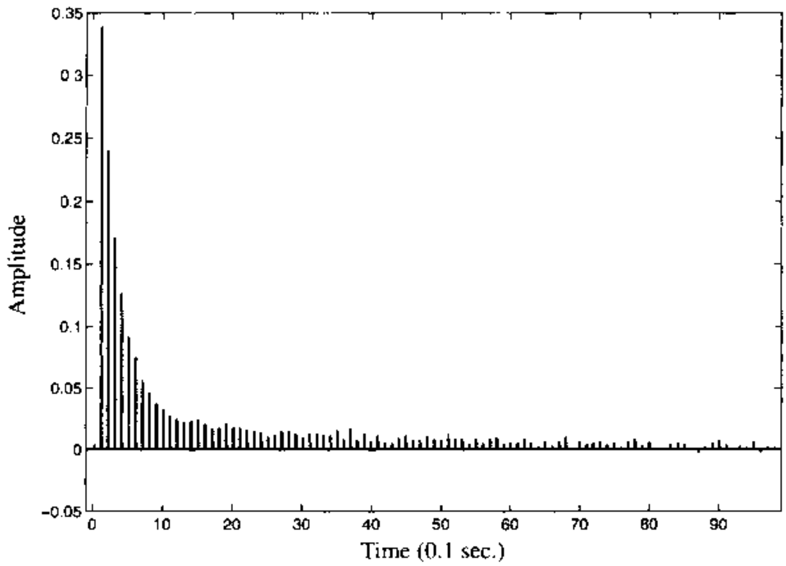

Figure 1.25 Impulse response of converged adaptive model of equivalent plant (minimum-phase).

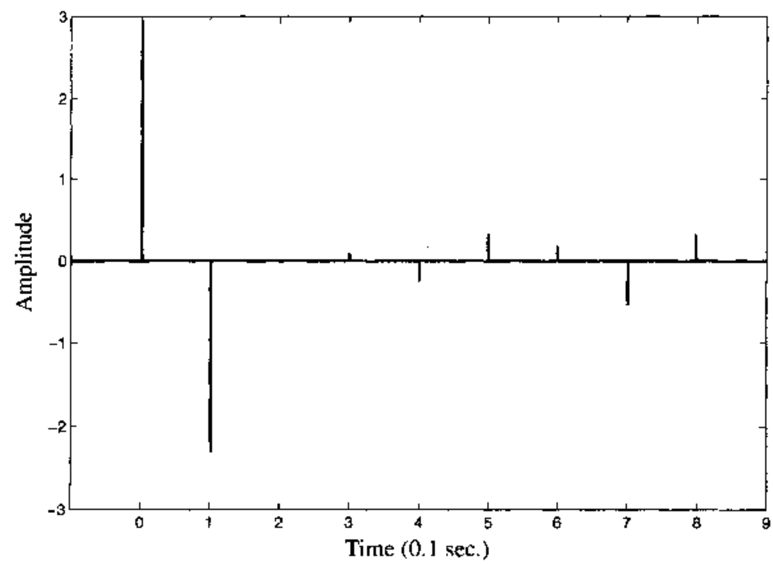

Figure 1.26 Converged impulse response of adaptive inverse model of equivalent plant (minimum-phase), for use in plant disturbance canceler.

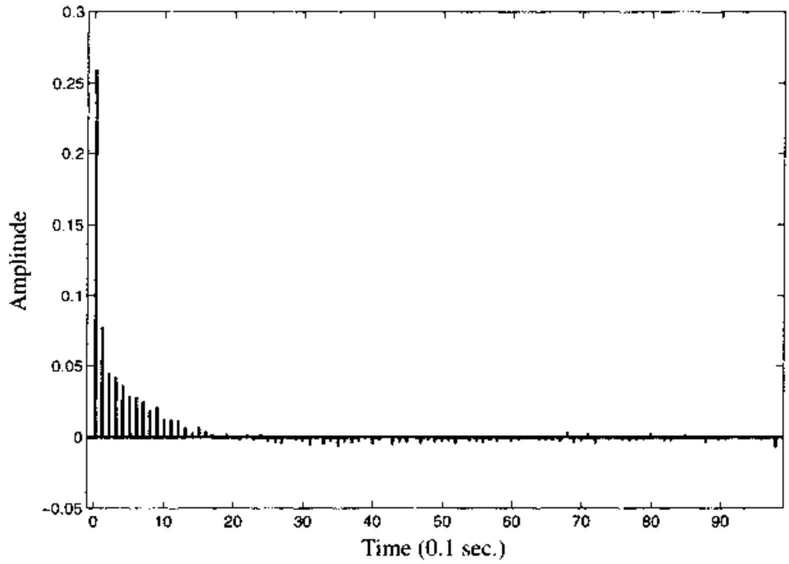

Figure 1.27 Convergence of model-reference inverse controller for equivalent plant (minimum-phase). Plant disturbance was present during learning process.

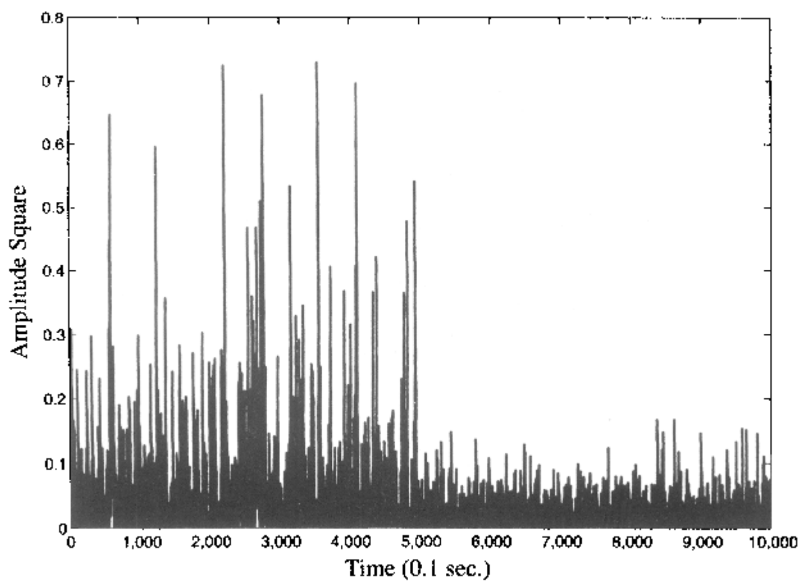

Figure 1.28 Square of plant output disturbance versus time. The disturbance canceler was turned on at the five-thousandth sample.

square of this error is graphed over time, without averaging. Figure 1.32 shows the model-reference output and the plant output over the last 200 samples in the sequence. Tracking between them is almost perfect, except for the small residual of uncanceled plant disturbance. The controller has long since converged and provides the proper dynamic control over the plant, and the disturbance canceler is turned on and doing an optimal job of canceling plant disturbance. The plant is under control!

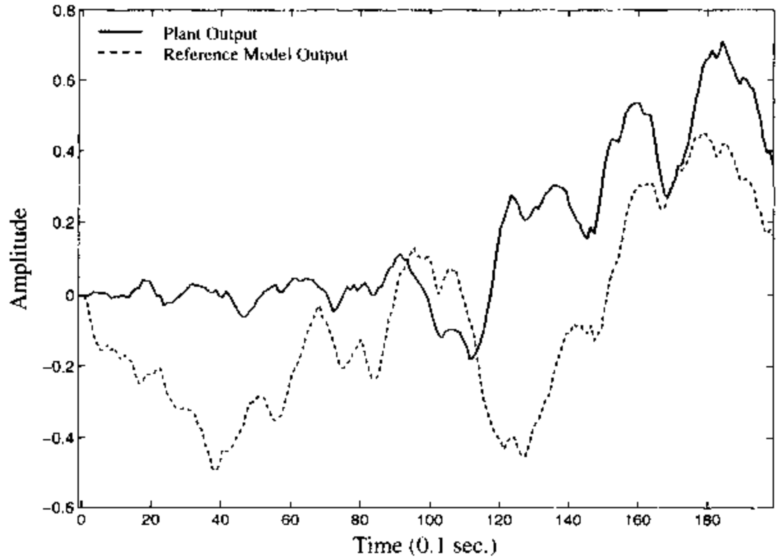

Figure 1.29 Comparison of reference model output with plant output at the beginning of the learning sequence. Plant disturbance is present.

.3 AN OUTLINE OR ROAD MAP FOR THIS BOOK

Section 1.1 has provided an overview of the idea of adaptive inverse control. This approach treats dynamic control of the plant as a separate problem from that of control of plant disturbance:

- Dynamic control is effected to cause the overall system response from command input signal to plant output signal to match that of a selected reference model.
- Plant disturbance is controlled to minimize the mean square of the plant output noise and disturbance.

Treating these as two separate problems is a very effective approach since the processes of dynamic control and noise and disturbance control can be optimized without one compromising the other.

The basic block diagrams of Section 1.2 are highly simplified. There are many details that need to be considered before one can successfully apply adaptive inverse control to

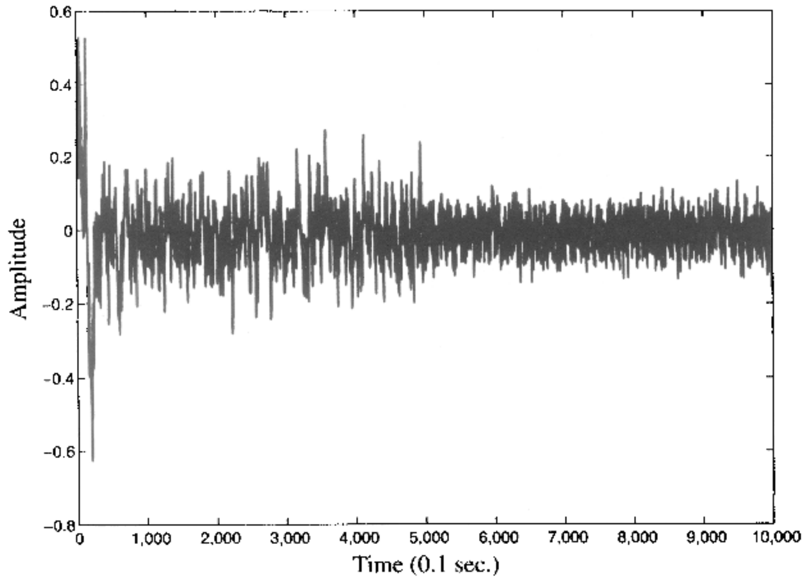

Figure 1.30 Entire learning sequence of 10,000 samples. Plot of error, equal to difference between reference model output and plant output. Disturbance canceler turned on at 5,000 samples.

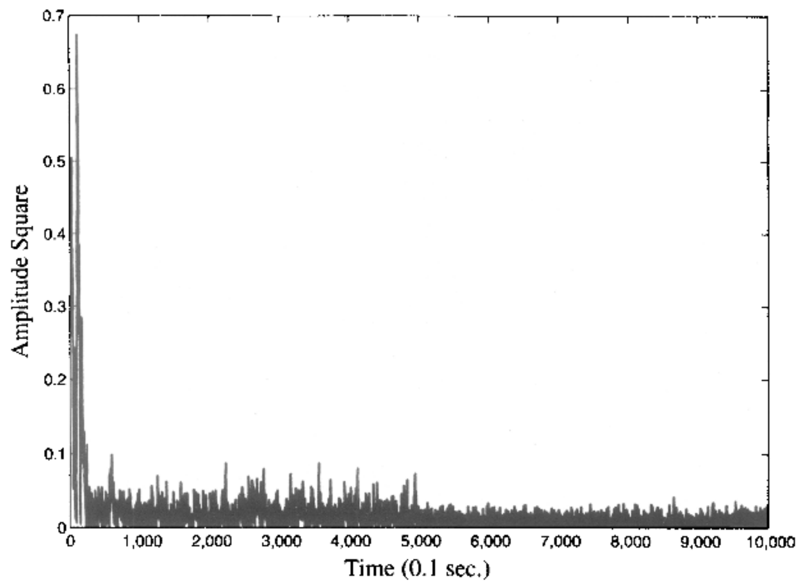

Figure 1.31 Square of overall system error plotted without averaging, over entire 10,000-sample learning sequence. Disturbance canceler turned on at 5,000 samples.

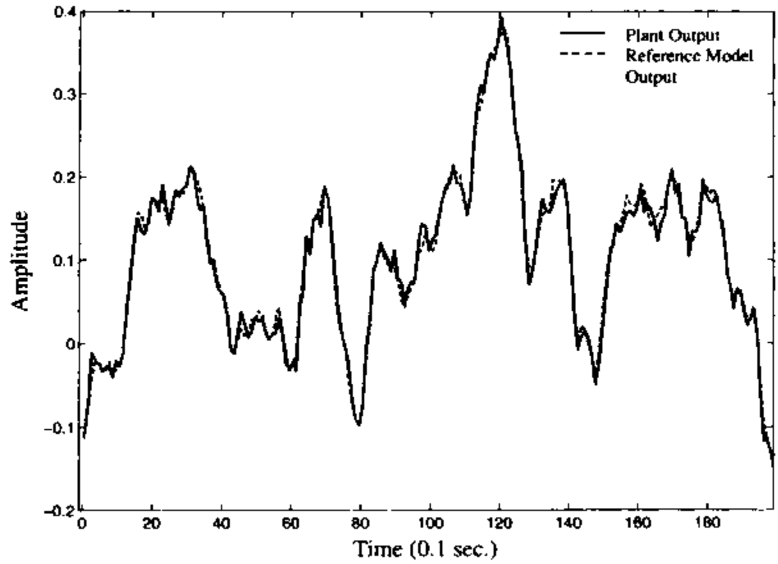

Figure 1.32 Comparison of reference model output with plant output over last 200 samples of training sequence. Controller has long ago converged, and plant disturbance is being canceled.

practical problems. Our goal is to explain some of the many different techniques that are available for the practical application of adaptive inverse control. The theory behind the various techniques is also derived and explained.

This book contains 12 chapters and 8 appendices. Since we expect that the subject will be new to most readers, we thought it appropriate to include in this introductory chapter a "road map" of the book. We are addressing both signal processing engineers and control engineers. The methodology taught here, the mathematical techniques and the adaptive signal processing techniques, will be more familiar to signal processing people. On the other hand, the problems that are being discussed are more familiar to control people.

We next describe the book and explain how it is organized.

Chapter 1: The Adaptive Inverse Control Concept

Chapter 1 introduces the idea of adaptive control and explains the need for it. It describes the more usual approaches to adaptive control and outlines the adaptive inverse control idea and its history. This chapter explains in an overall way how adaptive inverse control manages to control plant dynamics, and separately, plant noise or plant disturbances. It briefly reviews the subject matter of each of the chapters of this book. The objective is to aid the control engineer and the signal processing engineer in gaining an understanding and a perspective on adaptive inverse control and its applications.

Chapter 2: Wiener Filters

Wiener filters are best linear least squares filters which are very useful for prediction, estimation, interpolation, signal and noise filtering, and so on. To design them, prior knowledge of the appropriate statistical properties of the input signal(s) is required. The theory of Wiener filtering is also very useful in determining the asymptotic behavior of adaptive filters. For this reason, we present a brief chapter on the subject.

This chapter is a review of Wiener theory from the discrete point of view. In discrete form, Wiener theory is appropriate for analyzing digital adaptive filters. Causal and noncausal Wiener filters are studied. Both forms are used to analyze the behavior of adaptive filters as inverse controllers and plant disturbance cancelers. A very simple approach to causal Wiener filtering was devised by Bode and Shannon [68]. Shannon-Bode theory is explained in detail. This theory is extremely useful in the development of adaptive inverse control.

Chapter 3: Adaptive LMS Filters

This chapter reviews the theory of adaptive digital filtering, which is essential to the development of adaptive inverse control. This subject is discussed in many papers [1–60], among them several by B. Widrow and co-authors. Adaptive filters are discussed at length in several textbooks such as Widrow and Stearns [65], Haykin [69, 70] Cowan and Grant [71], and Treichler and colleagues [72].

The chapter begins with the idea of an adaptive filter, a tapped delay line with variable coefficients or tap weights driven by the LMS algorithm of Widrow and Hoff [73–77]. This is a gradient algorithm based on the method of steepest descent. The use of this algorithm to adjust the weights to minimize mean square error is described in the context of several practical applications. An important application is that of plant identification, necessary for adaptive control. Wiener theory is used to describe asymptotic adaptive behavior. The speed of convergence and the effects of gradient noise (caused by obtaining gradients with finite amounts of input-signal data) are analyzed. Fast adaptation causes noisy weights which, in excess, can result in poor performance.

The efficiency of adaptive algorithms is discussed. Two algorithms operating with the same level of noise in their weights can be compared. The one that converges faster is the more efficient. It is shown that when the input of an adaptive filter is stationary and the eigenvalues are equal or close in value, the LMS algorithm performs at a level of efficiency close to a theoretical maximum. When the eigenvalues are highly disparate, an algorithm similar to LMS but based on Newton's method would be needed to approach the theoretical maximum in efficiency. For nonstationary inputs, the conventional steepest-descent LMS algorithm performs with close to maximum efficiency [77].

This chapter concludes with a description of adaptive filtering as applied to adaptive noise canceling. The principles of adaptive noise canceling are derived. Experimental results are presented illustrating the effectiveness of noise canceling techniques to problems in adult electrocardiography (removing 60-Hz interference) and to fetal electrocardiography (removing interference from the maternal heart to reveal tiny signals from the fetal heart).

Chapter 4: Adaptive Modeling

The idea of LMS adaptive filtering for plant identification was discussed in Chapter 3. This chapter analyzes the sources of error in adaptive plant modeling, such as the effects of (a) weight noise, (b) inadequate input signal or dither signal, (c) plant noise, and so forth. Expressions are obtained for the covariance of the error or the difference between the plant's impulse response and the adaptive model's impulse response. These expressions are used in subsequent chapters to analyze the performance of adaptive inverse control systems.

Chapter 5: Inverse Plant Modeling

Direct adaptive modeling or plant identification was discussed in detail in Chapter 4. Chapter 5 discusses inverse adaptive modeling. The direct model has a transfer function similar to that of the plant being modeled. The inverse model has a transfer function like the reciprocal of the plant transfer function.

Forming a stable inverse transfer function is easy when the plant is minimum-phase. But doing this for a nonminimum-phase plant requires a two-sided Wiener impulse response which, to be causal, needs to be delayed and truncated. A theoretically optimal solution is obtained with the help of Shannon-Bode theory. The theory shows how approximate inverses can be formed for nonminimum-phase plants. Although such inverses are delayed, they can be very effective.

Model-reference inverses can be formed in a similar way. The result is that the cascade of the plant and its inverse develops an impulse response equal to that of a selected *reference model*. The inverse filter becomes the controller for the plant. The entire control system then has a dynamic response like that of the reference model. This is model-reference adaptive inverse control.

If the plant being controlled is subject to disturbance, the plant inverse may be obtained by adapting the inverse filter against a direct model of the plant instead of against the plant itself. Gradient noise in the weights of both the direct model and the inverse model affect the accuracy of the inverse model. Expressions are obtained for the variance of the error at the plant output due to noise in the inverse filter's weights. This knowledge is very important to the design of the adaptive controller.

Chapter 6: Adaptive Inverse Control

Schemes for adaptive inverse control were first presented in Chapter 1, but as shown there, they lack sufficient detail to be used in practice.

Practical model-reference control systems are shown in Chapter 6 for both disturbance-free plants and for plants with enough disturbance to interfere with the inverse modeling process. To alleviate the effects of plant disturbance, direct modeling is done first, and the inverse modeling process is done offline using a disturbance-free plant model to form the inverse. The idea of a dither signal, used in the direct modeling process, is introduced here.

If the dither is strong, good direct and inverse modeling is the result. But the dither adds to the plant output disturbance. A theory is developed to optimize the power of the dither to minimize the summed effects of the control error and plant disturbance at the plant output.

Adaptive control of blood pressure in experimental animals is reported. Successful experiments were performed regulating a dog's blood pressure in real time. The animal was in an induced state of shock. Under computer control, life-sustaining administration of therapeutic drugs took place. Dynamic control of blood pressure was accomplished by adaptive inverse control.

Chapter 7: Other Configurations for Adaptive Inverse Control

Two other algorithms for finding the inverse controller when plant disturbance is present are described in Chapter 7. These algorithms and variations on them are called *filtered-X* and *filtered-ϵ* LMS algorithms. Their advantages and disadvantages relative to each other and to the algorithms of Chapter 6 are discussed in detail. Conditions for stability are derived. Learning rates are obtained, as are expressions for weight noise variance. All of this can be used to predict the performance of adaptive inverse control systems.

A practical problem, canceling noise that leaks through earphones in a high-noise environment, illustrates application of both the filtered-X algorithm and the filtered-ϵ algorithm. This signal processing problem is very much like a control problem.

Chapter 8: Plant Disturbance Canceling

In the previous chapters a variety of techniques were developed for achieving precise dynamic control, even for plants subject to disturbance. These methods control plant dynamics but do nothing to control, reduce, or eliminate plant disturbance, however. An optimal adaptive scheme for reducing or eliminating plant disturbance is described and analyzed in this chapter. Its learning rate is determined. Shannon-Bode theory [68] is used to prove optimality and to derive expressions for the power of the residual plant output disturbance after adaptive cancelation. It is proven for linear systems that no other method can reduce the mean square of the plant output disturbance to a lower level. Expressions are derived for overall system output mean square error for simultaneous application of adaptive inverse control and adaptive plant disturbance canceling.

A simulation of an adaptive disturbance canceling system is reported. Application of this system to real-time aircraft ride control is illustrated. Fast control of the airplane's ailerons could reduce the vertical component of disturbance due to random updrafts and downdrafts.

Chapter 9: System Integration

An entire control system consisting of an inverse controller, a plant to be controlled, and an adaptive disturbance canceler for the plant is described and analyzed in Chapter 9. A surprising result develops from the analysis. If the plant model $\widehat{P}(z)$, which is used to obtain both the controller $\widehat{C}(z)$ and the noise-canceling feedback filter $Q(z)$, does not perfectly represent the plant $P(z)$, the errors in $\widehat{P}(z)$ and the resulting errors in $\widehat{C}(z)$ and $Q(z)$ combine in such a way that the overall system transfer function from command input to plant output is unaffected by the errors in $\widehat{P}(z)$. In other words, small errors in $\widehat{P}(z)$ cause small errors in the design of the controller $\widehat{C}(z)$. In addition, small errors in $\widehat{P}(z)$ cause small changes in the transfer function of the plant $P(z)$ in feedback connection with its disturbance canceler. The effects of these two transfer function errors cancel, leaving the overall system transfer function unaffected by small errors in the plant model $\widehat{P}(z)$.

The effect of small errors in $\widehat{P}(z)$ upon the disturbance canceler's ability to cancel plant disturbance is small. This result was obtained in Chapter 8. The effectiveness of the adaptive disturbance canceler is therefore not significantly impaired by small errors in $\widehat{P}(z)$.

Once the plant model $\widehat{P}(z)$ has been obtained, the rest of the system falls into place. Offline processes can be used to compute $Q(z)$ for the plant disturbance canceler, and to compute the controller $\widehat{C}(z)$. $\widehat{P}(z)$ is also used in the plant disturbance canceler.

The determination of $\widehat{P}(z)$ is not critical, however. Small errors in $\widehat{P}(z)$ only slightly degrade the performance of the disturbance canceler and have no effect on the precision of the dynamic control of the plant. This is an important result. Although the plant is being controlled by an open-loop controller, feedback in the adaptive algorithms and interactions among them result in precise dynamic control and close to optimal disturbance canceling.

Chapter 10: Multiple-Input Multiple-Output (MIMO) Adaptive Inverse Control Systems

MIMO systems are used when a physical plant to be controlled has multiple actuators, all of which have interacting effects on the process, and multiple sensors. Modeling and controlling MIMO systems are much more complicated than doing the same with SISO (single-input single-output systems).

This chapter begins with a review of digital signal processing for MIMO systems. The review section introduces notation used throughout the chapter. MIMO systems may be represented with block diagrams and flow graphs. Each signal line carries a bundle of signals, that is, a signal vector. Transfer functions are matrices. The rules of matrix algebra apply, and transfer functions are not commutable. This has a considerable effect on the design of adaptive algorithms to do plant modeling, inverse modeling, and plant disturbance canceling, all of which are described for MIMO systems.

Two methods for devising MIMO inverse controllers are presented in Chapter 10. One is an algebraic technique, and the other is based on the filtered-ϵ algorithm. Methods for adaptive plant disturbance canceling are also described.

Integrated MIMO systems incorporating inverse dynamic control and plant disturbance canceling in a single control system are shown. There are many ways to do this, several of which are illustrated in Chapter 10.

An application is described concerning the problem of reducing noise in a volume of space by using several loudspeakers and several sensing microphones in a MIMO controller configuration. Control of noise in an airplane cabin is illustrated here as an example. These are signal processing problems which are mathematically identical to control problems.

Chapter 11: Nonlinear Adaptive Inverse Control

The purpose of this chapter is to show how to do adaptive "inverse control" with nonlinear plants of both the SISO and MIMO types. Although nonlinear dynamic plants generally do not have inverses, techniques like inverse control can still be used.

Many of the rules of MIMO systems apply to nonlinear systems, such as noncommutability of filtering operations. An additional rule for nonlinear systems is that plant behavior should be modeled only with input signal character and power level corresponding to those of the actual plant input. Scaling and linearity do not work.

Inverse control of nonlinear plants and plant noise and disturbance canceling require the use of nonlinear adaptive filters. This chapter shows how to make nonlinear filters as tapped delay lines with Volterra networks [93], [94] connecting the tap signals to the variable weights, or with neural networks connected to the tap signals. Both of these methods process input signals nonlinearly to make output signals. A Volterra filter may be adapted by the LMS algorithm. A neural network filter would generally be adapted by the *backpropagation* algorithm of Werbos [78], and Rumelhart and colleagues [79], [80]. Backpropagation is the most widely used training algorithm for neural networks worldwide. It is an outgrowth of and a substantial generalization of the LMS algorithm.

How to do nonlinear plant modeling, with and without dither, is described in this chapter. When using dither, the superposed natural plant command signals, which could be nonstationary and which sometimes could be larger or sometimes smaller than the dither, mix nonlinearly in the plant. Simple direct modeling is not so simple. Means of dealing with this interaction are explained.

Once the plant model has been obtained, an inverse controller can be devised. Doing this with the filtered-ϵ LMS algorithm is illustrated. Care is exercised to ensure that all adaptive filters during training have input signals that have the right amplitude levels and the right signal characteristics.

Nonlinear plant noise canceling is demonstrated next. Online and offline processes for obtaining $\widehat{P}$, $\widehat{P}^{-1}$, Q, and $\widehat{C}$ are shown. The chapter ends with a block diagram for an integrated nonlinear control system which could be SISO or MIMO. The system incorporates both plant disturbance canceling and nonlinear inverse dynamic control.

Chapter 12: Pleasant Surprises

This chapter summarizes the principal theoretical results of adaptive inverse control. So many of these results were unexpected and worked out so nicely that we called the chapter "Pleasant Surprises."

Appendix A: Stability and Misadjustment of the LMS Adaptive Filter

This appendix summarizes learning theory for adaptive transversal filters based on the LMS adaptive algorithm. Key issues are learning rate, stability, and effects of noise in the weights (misadjustment). Most of this work has been published elsewhere by Widrow and colleagues [73–77]. More precise stability conditions and formulas for misadjustment has been reported by Horowitz and Senne [81]. Simplified derivations are given here that provide substantially the same results as those of Horowitz and Senne.

The work presented in this appendix is used throughout the book and is brought together here for the convenience of the reader.

Appendix B: Comparative Analyses of Dither Modeling Schemes A, B, and C

Basically, there are four ways of doing plant modeling which are illustrated in this book. The first method uses the natural plant input signals encountered during normal system operation to do the modeling. The second method, scheme A, adds a dither signal to the natural plant input to augment this signal and cause the combination to be "persistently exciting" [82]. The objective is to ensure that all frequency components are present in the modeling signal. The third method, scheme B, addresses an issue that could arise from nonstationarity of the natural plant input signal or if this signal has a highly nonuniform spectral character. Modeling might in some cases be done better with the dither alone, without the natural signal. With scheme B, the input signal to the adaptive modeling filter is pure dither. Carrying the idea one step further, the fourth method, scheme C, removes the natural plant output component (due to the natural plant input) from the "desired response" or training signal of the adaptive modeling filter. Scheme C uses only dither for the input and for the desired response of the adaptive plant model.

All four of these methods are used throughout the book. They have advantages and disadvantages relative to each other. This appendix develops ranges of stable operating conditions, learning rates, and misadjustment for all of the methods.

Appendix C: A Comparison of the Self-Tuning Regulator of Åström and Wittenmark with the Techniques of Adaptive Inverse Control

Some of the finest work in adaptive control has been done by Åström and Wittenmark. Their self-tuning regulator [82–88] is known worldwide for its simplicity, elegance, and utility. It can be used to control plant disturbance and plant dynamics, but it does not control them independently like adaptive inverse control. The purpose of Appendix C is to compare these two approaches to adaptive control in order to appraise their relative strengths and weaknesses.

The self-tuning regulator can be represented as a system consisting of a plant, a feedback controller, and a feedforward controller. The feedback controller can be designed to minimize plant output noise and disturbance power. But the feedback changes the dynamics. Compensation can be made for this, and at the same time, the overall response can be made to match a model response $M(z)$ by properly choosing the feedforward input controller.

It is shown in this appendix that an optimal feedback controller for a self-tuning regulator can always be chosen, and the resulting feedback portion of the system will always be stable. The requirements on the feedforward controller may turn out to make it unstable, however. (This is unstable in the sense of having poles outside the unit circle in the z-plane.) When this happens, one can design a stable feedforward controller that can approximately realize the required transfer function with a delayed response. The self-tuning regulator is indeed a general methodology.

Adaptive inverse control can always realize the optimal plant noise canceler without experiencing instability due to feedback. Furthermore, the ideal inverse controller will be stable, with its poles inside the unit circle if the plant is minimum-phase. This ideal controller can be realized approximately by a transversal filter without delay. A delayed response is only required when the plant is nonminimum-phase. Adaptive inverse control is a general methodology.

Adaptive inverse control and the self-tuning regulator are two completely different approaches to the problem of adaptive control of plant dynamics and plant disturbance. Their methods of adaptation differ, their adaptive filters are of different structure, their errors in dynamic response are different, but their abilities to cancel plant disturbance should be equal. Which approach is better, more robust, faster to converge, easier to implement, easier to design and debug, and easier to understand? The answers are probably problem dependent.

Appendix D: Adaptive Inverse Control for Unstable Linear SISO Plants

If the plant to be controlled is unstable, it is impossible to control it with adaptive inverse control. The reason for this is that a feedforward controller, even though it is adaptive, will leave the plant unstable.

The first step in the utilization of adaptive inverse control for an unstable plant is stabilization with feedback. This appendix shows that the choice of stabilization feedback is not critical as long as the unstable plant is stabilized. The plant together with its feedback stabilizer should be treated as a unit, as a stable *equivalent plant*. The equivalent plant can be outfitted with an adaptive disturbance canceler and with an inverse controller, just like an ordinary stable plant.

If two different feedback filters can stabilize the plant, it is shown in this appendix that the minimal plant output disturbance (after adaptive disturbance cancelation) is the same for both stabilizing filters. Since no one stabilizer does better than any other, the choice of stabilizer is not critical. All one needs is some form of feedback filter that will stabilize the plant, and the adaptive disturbance canceler will deliver optimal performance.

The same is true for the design of the inverse controller. The choice of feedback stabilization is immaterial. One can always design an inverse controller for the stabilized plant. If the inverse controller needs delay for its proper realization, the required delay will not depend on the choice of the stabilization filter. Also, the length of the inverse controller will not depend on this choice.

The conclusion of Appendix D is that plant instability is no impediment to the use of adaptive inverse control. One needs to know only enough about the plant to design a feedback stabilizer for it. The design could be accomplished by experimentation. The design is not at all critical as long as the plant is stabilized. Once stabilized, the plant and its stabilizer should be treated as a stable equivalent plant, and adaptive inverse control should be applied in the usual way. Although a wide variety of stabilizer designs would generally be acceptable, forming any one design would not always be a straightforward process and indeed may require some effort, particularly if the unstable plant is nonlinear and/or MIMO.

Appendix E: Orthogonalizing Adaptive Algorithms: RLS, DFT/LMS, and DCT/LMS

This appendix was prepared by Dr. Francçise Beaufays, based on her Ph.D. dissertation research in the Department of Electrical Engineering at Stanford University. Extreme eigenvalue spread of the autocorrelation matrix of the input to an adaptive filter causes convergence problems for the LMS algorithm. To speed up convergence and maintain stability, she proposes a new adaptive algorithm, the DCT/LMS.

She begins with a brief discussion of the RLS (recursive least squares) algorithm. This algorithm is most often very effective, but it is complicated and has its own stability problems.

The next algorithm discussed is the DFT/LMS. This algorithm first Fourier transforms the signals at the taps of a filter's tapped delay line. The purpose is to achieve orthogonalization, which the DFT does well, except for the phenomenon of *leakage*. Once the tap signals are orthogonalized (approximately) and power normalized, they are weighted and summed to produce the filter output. The weights are adapted by conventional LMS. Although adaptation is done simply with steepest descent, the behavior is similar to adaptation with Newton's method.

Finally, the DCT/LMS algorithm is discussed. Structurally, this algorithm is very similar to DFT/LMS, with the DCT (digital cosine transform) substituted for the DFT. Ms. Beaufays shows that the DCT/LMS algorithm performs significantly better than the DFT/LMS algorithm when the adaptive filter input is first-order Markov, a type of signal that is common in signal processing and controls. This type of signal comes from applying white noise to a one-pole filter.

The DFT/LMS and DCT/LMS algorithms are easy to implement and are computationally robust. All of their parts, DFT or DCT, power normalization (like AGC in a radio or TV), and LMS algorithm are "bulletproof." These algorithms should enjoy greater acceptance and application in the future.

Appendix F: A MIMO Application: An Adaptive Noise Canceling System Used for Beam Control at the Stanford Linear Accelerator Center

This appendix was prepared by Dr. Tom Himel, a research physicist at the Stanford Linear Accelerator Center (SLAC). It is based on his experience with control of the beam of a two-mile long linear accelerator on the Stanford campus whose output drives positrons and electrons in opposite directions along the arcs of a circular collider. To achieve collisions suitable for physics research, the electron and positron beams must be controlled in position to within several microns of each other after each travels distances of about three miles.

Control of the beam is critical. The U.S. Department of Energy, the sponsor, spends millions of dollars a year operating the accelerator and has spent many hundreds of millions building it over the past 25 years. Physicists from all over the world depend on this machine for their research.

Until recently, the beam was controlled by many servo loops positioned along it, with manual control of set points. Intercoupling between stages has always been a problem for overall control. The beam output of one stage is the beam input for the next. Adaptive techniques are now used to obtain precise local beam positioning and at the same time, to obtain decoupling between stages.

The corrective methodology is based on adaptive noise canceling. The adaptive disturbance canceler at each stage is an eight-input, eight-output MIMO system.

The adaptive system (implemented in software) has been installed and working for more than a year, 24 hours a day, 7 days a week. The result is much better accelerator operation. It is automatic, without the need of human operator intervention, and the frequency of collision events has increased. This is an operational system, no longer a laboratory demonstration.

Appendix G: Thirty Years of Adaptive Neural Networks: Perceptron, Madaline, and Backpropagation

This appendix is a reprint of a paper by Widrow and Lehr that was published in the September 1990 issue of the *Proceedings of the IEEE*. For the reader who is familiar with the LMS algorithm and adaptive filters, this paper introduces the subject of neural networks. It shows how the LMS algorithm can be extended to form the *backpropagation* algorithm, the most widely used learning procedure for neural networks.

Appendix H: Neural Control Systems

This appendix shows how to construct nonlinear adaptive inverse control systems by making use of neural networks and the backpropagation algorithm. It also describes applications of neural control systems.

One such application is the truck backer-upper. A simulated trailer truck is steered by a neural controller while backing to a loading platform. After many backing runs from many different initial conditions, the controller learns to steer the truck by making many mistakes and learning what not to do. Once learning is complete, the truck can be placed in almost any initial state, states not previously encountered, and the controller steers the truck while driving backward to the loading dock. The truck backer-upper is an example of how a nonlinear controller can learn and, in a real sense, design itself.

Principles of inverse neural control similar to those incorporated in the truck backing system have been used to control the functioning of electric furnaces and for chemical process control. These applications are described in some detail and are not laboratory exercises. They are industrial applications of high commercial significance and are examples of what is currently being called "intelligent control systems."

This completes the road map of the book. We now invite you to travel down the road with us. We wish you a productive journey.

Bibliography for Chapter 1

[1] B. WIDROW, K.M. DUVALL, R.P. GOOCH, and W.C. NEWMAN, "Signal cancellation phenomena in adaptive antennas: Causes and cures," *IEEE Trans. Antennas Propag.*, Vol. AP-30 (May 1982), pp. 469–478.

[2] C.W. JIM, "A comparison of two LMS constrained optimal array structures," *Proc. IEEE*, Vol. 65 (December 1977), pp. 1730–1731.

[3] L.J. GRIFFITHS and C.W. JIM, "An alternative approach to linearly constrained adaptive beamforming," *IEEE Trans. Antenna Propag.*, Vol. AP-30 (January 1982), pp. 27–34.

[4] R.P. GOOCH, "Adaptive pole-zero array processing," in *Proc. 16th Asilomar Conf. Circuits Syst. Comput.*, Santa Clara, CA, November 1982.

[5] P. HOWELLS, "Intermediate frequency side-lobe canceller," U.S. Patent 3202 990, August 24, 1965.

[6] Special issue on adaptive antennas, *IEEE Trans. Antennas Propag.*, Vol. AP-24, No. 5 (September 1976).

[7] W.F. GABRIEL, "Adaptive arrays—An introduction," *Proc. IEEE*, Vol. 64 (February 1976), pp. 239–272.

[8] B. WIDROW, "Adaptive antenna systems," *Proc. IEEE*, Vol. 55, No. 12 (December 1967), pp. 2143–2159.

[9] L.J. GRIFFITHS, "A simple adaptive algorithm for real-time processing in antenna arrays," *Proc. IEEE*, Vol. 57, No. 10 (October 1969), pp. 1696–1704.

[10] O.L. FROST, III, "An algorithm for linearly constrained adaptive array processing," *Proc. IEEE*, Vol. 60, No. 8 (August 1972), pp. 926–935.

[11] C.L. ZAHM, "Applications of adaptive arrays to suppress strong jammers in the presence of weak signals," *IEEE Trans. Aerosp. Electron. Systems*, Vol. AES-9 (1973), pp. 260–271.

[12] R.T. COMPTON, JR., R.J. HUFF, W.G. SWARNER, and A.A. KSIENSKI, "Adaptive arrays for communciation systems: An overview of research at the Ohio State University," *IEEE Trans. Antennas Propag.*, Vol. AP-24, No. 5 (September 1976), pp. 599–606.

[13] R.T. COMPTON, JR., "An experimental four-element adaptive array," *IEEE Trans. Antennas Propag.*, Vol. AP-24, No. 5 (September 1976), pp. 697–706.

[14] L.E. BRENNAN, AND I.S. REED, "Convergence rate in adaptive arrays," Technology Service Corp., Rep. TSC-PD-177-4, January 13, 1978.

[15] W.D. WHITE, "Adaptive cascade networks for deep nulling," *IEEE Trans. Antennas Propag.*, Vol. AP-26, No. 3 (May 1978), pp. 396–402.

[16] S.P. APPLEBAUM, "Adaptive arrays," *IEEE Trans. Antennas Propag.*, Vol. AP-24 (September 1976), pp. 585–598.

[17] E. WALACH, "On superresolution effects in maximum-likelihood antenna arrays," *IEEE Trans. Antennas Propag.*, Vol. AP-32, No. 3 (March 1984), pp. 259–263.

[18] L.J. GRIFFITHS, "A comparison of multidimensional Wiener and maximum-likelihood filters for antenna arrays," *Proc. IEEE (Letters)*, Vol. 55 (November 1967), pp. 2045–2047.

[19] N.K. JABLON, "Effect of element errors on half-power beamwidth of the Capon adaptive beamformer," *IEEE Trans. on Circuits and Systems*, Vol. CAS-34, No. 7 (July 1987), pp. 743–751.

[20] N.K. JABLON, "Steady state analysis of the generalized sidelobe canceller by adaptive noise cancelling techniques," *IEEE Trans. Antenna Propag.*, Vol. AP-34 (March 1986), pp. 330–337.

[21] R.W. LUCKY, "Automatic equalization for digital communication," *Bell Syst. Tech.*, Vol. 44 (April 1965), pp. 547-588.

[22] R.W. LUCKY, "Techniques for adaptive equalization of digital communication systems," *Bell Syst. Tech.*, Vol. 45 (February 1966), pp. 255–286.

[23] A. GERSHO, "Adaptive equalization of highly dispersive channels for data transmission," *Bell Syst. Tech.*, Vol. 48 (January 1969), pp. 55–70.

[24] M.M. SONDHI, "An adaptive echo canceller," *Bell Syst. Tech.*, Vol. 46 (March 1967), pp. 497–511.

[25] J.G. PROAKIS, *Digital communications* (New York: McGraw-Hill, 1983).

[26] R.W. LUCKY, J. SALZ, and E.J. WELDON, JR., *Principles of data communication* (New York: McGraw-Hill, 1968).

[27] R.D. GITLIN, J.F. HAYES, and S.B. WEINSTEIN, *Data communications principles* (New York: Plenum Press, 1992).

[28] S. QURESHI, "Adaptive equalization—a comprehensive review," *IEEE Communications Magazine* (March 1982), pp. 9–16.

[29] J.G. PROAKIS, "Adaptive digital filters for equalization of telephone channels," *IEEE Trans. Audio Electroacoust.*, Vol. AU-18, No. 2 (June 1970), pp. 484–497.

[30] J.G. PROAKIS and J.H. MILLER, "An adaptive receiver for digital signaling through channels with intersymbol interference," *IEEE Trans. Info. Theory*, Vol. IT-15, No. 4 (July 1969), pp. 484–497.

[31] S.J. CAMPANELLA, H.G. SUYDERHOUD, and M. ONUFRY, "Analysis of an adaptive impulse response echo canceller," *Comsat Tech. Rev.*, Vol. 2, No. 1 (Spring 1972), pp. 1–37.

[32] D.L. DUTTWEILER, "A twelve-channel digital echo canceller," *IEEE Trans. Communications*, Vol. COM-26 (May 1978), pp. 647–653.

[33] D.L. DUTTWEILER, and Y.S. CHEN, "Performance and features of a single chip FLSI echo canceller," in *Proc. NTC 79*, Washington, DC, November 17–19, 1979.

[34] M.M. SONDHI and D.A. BERKLEY, "Silencing echoes on the telephone network," *Proc. IEEE*, Vol. 68 (August 1980), p. 948.

[35] F.K. BECKER and H.R. RUDIN, "Application of automatic transversal filters to the problem of echo suppression," *Bell Syst. Tech.*, Vol. 45 (December 1966), pp. 1847–1850.

[36] M.M. SONDHI, and A.J. PRESTI, "A self-adaptive echo canceller," *Bell Syst. Tech.*, Vol. 45 (December 1966), pp. 1851–1854.

[37] D.D. FALCONER, K.H. MUELLER, and S.B. WEINSTEIN, "Simultaneous two-way data transmission over a two-wire circuit," in *Proc. NTC*, Dallas, December 1976.

[38] T.R. ROSENBURGER and E.J. THOMAS, "Performance of an adaptive echo canceller operating in a noisy, linear time-invariant environment," *Bell Syst. Tech.*, Vol. 50 (March 1971), p. 785.

[39] N. DEMYTKO, and L.K. MACKECHNIE, "A high speed digital adaptive echo canceller," *Aust. Telecommun. Rev.*, Vol. 7 (1973), pp. 20–27.

[40] B.WIDROW, *et al.*, "Adaptive noise cancelling: principles and applications," *Proc. IEEE*, Vol. 63, No. 12 (December 1975), pp. 1692–1716.

[41] J. GLOVER, "Adaptive noise cancelling of sinusoidal interference," Ph.D. diss., Stanford Electronics Lab., Stanford University, Stanford, CA (December 1975).

[42] J. GLOVER, "Adaptive noise cancelling applied to sinusoidal interferences," *IEEE Trans. Acoust., Speech, Signal Processing*, Vol. ASSP-25, No. 6 (December 1977), pp. 484–491.

[43] M.R. SAMBUR, "Adaptive noise cancelling for speech signals," *IEEE Trans. Acoust., Speech, Signal Processing*, Vol. ASSP-26 (October 1978), p. 419.

[44] M.J. SHENSHA, "Non-Wiener solutions of the adaptive noise canceller with a noisy reference," *IEEE Trans. Acoust., Speech, Signal Processing*, Vol. ASSP-28 (August 1980), p. 468.

[45] L.J. GRIFFITHS, "An adaptive lattice structure for noise-cancelling applications," in *Proc. ICASSP-78*, p. 87.

[46] L.J. GRIFFITHS, and R.S. MEDAUGH, "Convergence properties of an adaptive noise cancelling lattice structure," *Proc. 1978 IEEE Conf. Decision and Control*, San Diego, CA (January 1979), pp. 1357–1361.

[47] L.J. GRIFFITHS, "Rapid measurement of instantaneous frequency," *IEEE Trans. Acoust., Speech, Signal Processing*, Vol. ASSP-23 (April 1975), pp. 209–222.

[48] J.P. BURG, "Maximum entropy spectral analysis," presented at the 37th Annual Meeting, Soc. Exploration Geophysicists, Oklahoma City, OK, 1967.

[49] J.R. ZEIDLER, *et al.*, "Adaptive enhancement of multiple sinusoids in uncorrelated noise," *IEEE Trans. Acoust., Speech, Signal Processing*, Vol. ASSP-26 (June 1978), p. 240.

[50] D.R. MORGAN, "Response of a delayed-constrained adaptive linear predictor filter to a sinusoid in white noise," in *Proc. 1981 ICASSP*, p. 271.

[51] J.A. EDWARDS, and M.M. FITELSON, "Notes on maximum-entropy processing," *IEEE Trans. Info. Theory*, Vol. IT-19 (March 1973), p. 232.

[52] R.T. LACOSS, "Data adaptive spectral analysis methods," *Geophysics* (August 1971), p. 661.

[53] A. VAN DEN BOS, "Alternative interpretation of maximum entropy spectral analysis," *IEEE Trans. on Info. Theory*, Vol. IT-17 (July 1971), p. 493.

[54] A. POPOULIS, "Maximum entropy and spectral estimation: A review," *IEEE Trans. Acoust., Speech, Signal Processing*, Vol. ASSP-29 (December 1981), p. 1176.

[55] W.F. GABRIEL, "Spectral analysis and adaptive array superresolution techniques," *Proc. IEEE*, Vol. 68 (June 1980).

[56] B. FRIEDLANDER, "Lattice methods for spectral estimation," *Proc. IEEE*, Vol. 70, No. 9 (September 1982), pp. 990–1017.

[57] B. WIDROW, and S.D. STEARNS, *Adaptive signal processing* (Englewood Cliffs, NJ: Prentice Hall, 1985), chap. 12.

[58] M. MORF, A. VIEIRA, and D.T. LEE, "Ladder forms for identification and speech processing," *Proc. 1977 IEEE Conf. Decision and Control*, New Orleans, LA (December 1977), pp. 1074–1078.

[59] J.I. MAKHOUL and L.K. COSELL, "Adaptive lattice analysis of speech," *IEEE Trans. Acoust., Speech, Signal Processing*, Vol. ASSP-29 (June 1981), p. 654.

[60] F. ITAKURA and S. SAITO, "Digital filtering techniques for speech analysis and synthesis," in *Proc. 7th Int. Conf. Acoust.*, Vol. 3, Paper 25C-1, 1971, p. 261.

[61] B. WIDROW, "Adaptive model control applied to real-time blood-pressure regulation," in *Pattern Recognition and Machine Learning; proceedings*, ed. K.S. Fu (New York: Plenum Press, 1971), pp. 310–324.

[62] B. WIDROW, J. MCCOOL, and B. MEDOFF, "Adaptive control by inverse modeling," in *Conf. Rec. of 12th Asilomar Conference on Circuits, Systems and Computers*, Santa Clara, CA, November 1978, pp. 90–94.

[63] B. WIDROW, D. SHUR, and S. SHAFFER, "On adaptive inverse control," in *Conf. Rec. of 15th Asilomar Conference on Circuits, Systems and Computers*, Santa Clara, CA, November 1981, pp. 185–189.

[64] S. SHAFFER, "Adaptive inverse-model control," Ph.D. diss., Stanford University, August 1982.

[65] B. WIDROW, and S.D. STEARNS, *Adaptive signal processing* (Englewood Cliffs, NJ: Prentice Hall, 1985).

[66] E. WALACH and B. WIDROW, "Adaptive signal processing for adaptive control," in *IFAC Workshop on Adaptive Systems in Control and Signal Processing*, San Francisco, CA, 1983.

[67] B. WIDROW, "Adaptive inverse control," in *Second IFAC Workshop on Adaptive Systems in Control and Signal Processing*, Lund, Sweden, July 1986.

[68] H.W. BODE, and C.E. SHANNON, "A simplified derivation of linear least square smoothing and prediction theory," *Proc. IRE*, Vol. 38 (April 1950), pp. 417–425.

[69] S. HAYKIN, *Introduction to adaptive filters* (New York: Macmillan, 1984).

[70] S. HAYKIN, *Adaptive filter theory* (Englewood Cliffs, NJ: Prentice Hall, 1986).

[71] C.F.N. COWAN, and P.M. GRANT, *Adaptive filters* (Englewood Cliffs, NJ: Prentice Hall, 1985).

[72] J.R. TREICHLER, C.R. JOHNSON, JR., and M.G. LARIMORE, *Theory and design of adaptive filters* (New York: John Wiley, 1987).

[73] B. WIDROW, and M.E. HOFF, "Adaptive switching circuits," in *IRE WESCON Conv. Rec.*, Pt. 4, 1960, pp. 96–104.

[74] B. WIDROW, and J.M. MCCOOL, "A comparison of adaptive algorithms based on the methods of steepest descent and random search," *IEEE Trans. Antennas Propag.*, Vol. AP-24, No. 5 (September 1967), pp. 615–637.

[75] B. WIDROW, J.M. MCCOOL, M.G. LARIMORE, and C.R. JOHNSON, Jr., "Stationary and nonstationary learning characteristics of the LMS adaptive filter," *Proc. IEEE*, Vol. 64, No. 8 (August 1976), pp. 1151–1162.

[76] B. WIDROW, "Adaptive filters," in *Aspects of Network and Systems Theory*, ed. R.E. Kalman and N. De Claris (New York: Holt, Rinehart and Winston, 1970), pp. 563–587.

[77] B. WIDROW, and E. WALACH, "On the statistical efficiency of the LMS algorithm with nonstationary inputs," *IEEE Trans. Info. Theory—Special Issue on Adaptive Filtering*, Vol. 30, No. 2, Part 1 (March 1984), pp. 211–221.

[78] P. WERBOS, *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*, Ph.D. diss., Harvard University, August 1974.

[79] D.E. RUMELHART, and J.L. MCCLELLAND, *Parallel Distributed Processing*, Vols. I and II (Cambridge, MA: M.I.T. Press, 1986).

[80] B. WIDROW and M.A. LEHR, "30 years of adaptive neural networks: perceptron, madaline, and backpropagation," *Proceedings of the IEEE*, Vol. 78, No. 9 (September 1990), pp. 1415–1442.

[81] L.L. HOROWITZ, and K.D. SENNE, "Performance advantage of complex LMS for controlling narrow-band adaptive arrays," *IEEE Trans. on Circuits and Systems*, Vol. CAS-28, No. 6 (June 1981), pp. 562–576.

[82] L. LJUNG, *System identification: Theory for the user* (Englewood Cliffs, NJ: Prentice Hall, 1987).

[83] K. ÅSTRÖM, *Introduction to stochastic control* (New York: Academic Press, 1970).

[84] K. ÅSTRÖM, and B. WITTENMARK, "On self-tuning regulators," *Automatica*, Vol. 9, No. 2 (1973).

[85] K. ÅSTRÖM, and B. WITTENMARK, "Analysis of a self-tuning regulator for nonminimum-phase systems," in *IFAC Symp. on Stochastic Control*, Budapest, Hungary, 1974.

[86] K. ÅSTRÖM, U. BORISSON, L. LJUNG, and B. WITTENMARK, "Theory and application of self tuning regulators," *Automatica*, Vol. 13 (1977), pp. 457–476.

[87] K. ÅSTRÖM, and B. WITTENMARK, "Self-tuning controllers based on pole-zero placement," *Proc. IEEE*, Vol. 127, Pt. D, No. 3 (May 1980), pp. 120–130.

[88] K. ÅSTRÖM, and B. WITTENMARK, *Adaptive control* (Reading, Mass: Addison-Wesley, 1989).

[89] P.V. OSBURN, H.P. WHITAKER, and A. KEZER, "New developments in the design of adaptive control systems," Paper No. 61-39, Institute of Aeronautical Sciences, February 1961.

[90] G.F. FRANKLIN, J.D. POWELL, and A. EMAMI-NAEINI, *Feedback control of dynamic systems* (Reading, Mass.: Addison-Wesley, 1986).

[91] G.F. FRANKLIN, J.D. POWELL, and M.L. WORKMAN, *Digital control of dynamic systems*, 2nd ed. (Reading, Mass.: Addison-Wesley, 1990).

[92] K. OGATA, *Modern control engineering*, 2nd ed. (Englewood Cliffs: Prentice Hall, 1990).

[93] D.F. SPECHT, "Generation of polynomial discriminant functions for pattern recognition," *IEEE Trans. on Electronic Computers*, Vol. EC-16, No. 3 (June 1967), pp. 308–319.

[94] V.J. MATHEWS, "Adaptive polynomial filters," *IEEE Signal Processing Magazine*, Vol. 8, No. 3 (July 1991), pp. 10–26.

[95] D. PSALTIS, A. SIDERIS, and A.A. YAMAMURA, "A multilayered neural network controller," *IEEE Control Systems Magazine*, Vol. 8 (April 1988), pp. 17–21.

[96] K.J. HUNT, and D. SBARBARO, "Neural networks for nonlinear internal model control," *IEEE Proceedings-D*, Vol. 138, No. 5 (September 1991), pp. 431–438.

[97] K.J. HUNT, D. SBARBARO, R. ZBIKOWSKI, and P.J. GANTHROP, "Neural networks for control systems—a survey," *Automatica*, Vol. 28, No. 6 (November 1992), pp. 1083–1112.