

Lehrstuhl für Elektrische Antriebssysteme
Technische Universität München

Adaptive Control in the Presence of Disturbances

Matthias J. Feiler

Vollständiger Abdruck der von der Fakultät für Elektrotechnik und Informationstechnik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktor-Ingenieurs

genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Dr.-Ing. Georg Färber

Prüfer der Dissertation:

1. Univ.-Prof. Dr.-Ing. Dr.-Ing. h.c. Dierk Schröder
2. Univ.-Prof. D.Sc. h.c. Kumpati S. Narendra, Ph.D.
YALE University, U. S. A.

Die Dissertation wurde am 28.06.2004 bei der Technischen Universität München eingereicht und durch die Fakultät für Elektrotechnik und Informationstechnik am 29.07.2004 angenommen.

Preface

The essence of adaptive control is to observe an unknown process, derive a hypothesis of how to control the process in order to obtain best performance and continuously update this hypothesis based on new information. In this thesis, I investigate ways to improve performance even as the information is corrupted by disturbances affecting the process.

I would like to thank Professor D. Schröder for having introduced me to this exciting field of research and for having provided his continuous support and interest during the years of my doctoral work. His open-mindedness and vision was the incentive for me to pursue new ways of intelligent control for mechatronic systems.

To my teacher and mentor, Professor K. S. Narendra, I express my profound gratitude for having accepted me as his graduate student even though I never officially enrolled at Yale. I thank him for having opened my eyes to the complex nature of research, for being such a brilliant thinker and such a dedicated teacher.

The research reported in this doctoral thesis was made possible by SIEMENS AG, which generously provided me with an Ernst von Siemens-stipend over a period of three years. In particular, I would like to thank Dr. Höller who decided that I should get the stipend, Mrs. Kiehnlein and my technical advisors Dr. Schierling, Prof. Dr. Wesselak and Dr. Bauer. Our intense discussions are responsible for my “coming down to earth” in the last chapter of this thesis. Furthermore, I gratefully acknowledge financial support by the German Academic Exchange Service (DAAD) and by Professor Narendra’s funding agencies who sponsored my regular visits to the Center for Systems Science at Yale University, U.S.A.

It is a great pleasure to thank my colleagues and friends, in particular, Frank Dietrich, Christian Westermaier, Hans Schuster and Martin Rau for the numerous enriching and delightful discussions we had. I am equally indebted to A. Angermann, G. Hoffstätter, W. Ebert and – most importantly – our secretary Leslie Patri, for helping me with the many administrative questions that one encounters in our institution.

I can think of no other person more directly affected by the preparation and completion of this thesis than my wife Elisabeth. She has gone through the many late working-hours and weekends as much as I have. Her continuing support and understanding was my key source of inspiration. To Elisabeth, I dedicate this work.

Augsburg, June 20, 2004

Matthias Feiler

Kurzzusammenfassung

Schwerpunkt dieser Arbeit ist die Entwicklung einer theoretischen Grundlage für den praktischen Entwurf adaptiver Systeme. In der Industrie bestehen adaptive Regler meist aus der Erweiterung eines Standard-Regelkreises durch ein Parameterschätzverfahren. Störgrößen machen eine korrekte Bestimmung der Parameter jedoch unmöglich. Es wird aufgezeigt, wie Systeme mit unbekanntem Parametern geregelt werden können, ohne dabei auf die Ergebnisse einer Parameterschätzung zurückzugreifen. Die zusätzlich wirksamen, unbekanntem Störungen werden in drei Klassen unterteilt, für die je eine Methode zur Unterdrückung erarbeitet wird: externe Störungen, nichtmodellerte Dynamik und zeitabhängige Parameter. Im letzten Fall wird ein Ansatz basierend auf multiplen adaptiven Modellen verwendet, der während regelmäßiger Forschungsaufenthalte in Yale entwickelt wurde. Experimentelle Ergebnisse an einem Zweimassen-Versuchsstand schließen die Arbeit ab.

Abstract

The emphasis of this thesis is on developing a theoretical framework for the practical design of adaptive systems. The state of the art in industrial adaptive control is to combine a parameter estimator with an (existing) linear control-loop. The success of the approach depends upon the knowledge of the physical parameters after completion of the estimation process. In the presence of disturbances, though, a correct estimation of the parameters is impossible. It is first shown how a system with unknown parameters can be controlled adaptively, without relying on parameter convergence. Following this, three classes of disturbances and the corresponding methods to reject them are developed: external disturbances, unmodelled dynamics and time-variations. In the latter case, a new algorithm based on multiple adaptive models is presented which was developed at Yale University during regular visits of the author at the Center for Systems Science. An experimental study of an adaptively controlled two-mass system concludes the thesis.

Contents

1	Introduction	1
2	Parameter Estimation	9
2.1	Projection Algorithm	11
2.1.1	Elementary Properties	12
2.1.2	Parameter Convergence	13
2.2	Orthogonalized Projection Algorithm	17
2.3	Recursive Least-Squares Algorithm	18
2.4	Arbitrary Relative Degree	22
2.5	An Example from Industry	23
3	Adaptive Control	27
3.1	Linear Adaptive Control	28
3.1.1	Exact Tracking	28
3.1.2	Auto-Regressive Moving-Average Representation	31
3.1.3	Adaptive Control	33
3.1.4	Proof of Stability	34
3.2	Nonlinear Adaptive Control Using Neural Networks	36
3.2.1	Nonlinear Normal Form	38
3.2.2	Control Using Input-Output Data	39
3.2.3	Practical Neurocontrol	43
4	Disturbance Rejection	49

4.1	Linear Disturbance Rejection	50
4.1.1	Deterministic Disturbances	50
4.1.2	Stochastic Disturbances	54
4.2	Nonlinear Disturbance Rejection	66
5	Design Considerations	71
5.1	Choice of the Reference Model	73
5.2	Input Constraints	80
5.2.1	Optimization of the Adaptive Gain	83
5.2.2	Discussion	88
5.3	Unmodelled Dynamics	91
5.3.1	The Full-Order Case	93
5.3.2	The Reduced-Order Case	94
6	Rapidly Time-Varying Systems	101
6.1	Multi-Model Adaptive Control	102
6.1.1	General Methodology	103
6.1.2	Models	105
6.1.3	Switching and Tuning	106
6.1.4	Control	107
6.1.5	Benefits	108
6.2	Proof of Stability	109
6.2.1	Case (i): All adaptive models	110
6.2.2	Case (ii): One adaptive model and one fixed model	111
6.2.3	Case (iii): (N-2) fixed models and 2 adaptive models	112
7	Self-Organization	115
7.1	Introduction	116
7.1.1	Qualitative Description	118
7.1.2	The method	118
7.1.3	Statement of the Problem	119

7.2	Parameter Convergence in Static Systems	121
7.2.1	Time-invariant Environment	123
7.2.2	Elementary Properties	125
7.2.3	Cycle vs. Instantaneous Dynamics	130
7.2.4	Qualitative Behavior of the Algorithm	135
7.2.5	Restriction to a 1-dimensional subspace	146
7.3	Dynamic Systems	152
7.3.1	Identification	153
7.3.2	Certainty Equivalence Control	159
7.4	Future Research	163
8	Application: Two-Mass System	167
8.1	Model of the Plant	168
8.2	Design of the Adaptive Controller	171
8.3	Experimental Results	176
8.3.1	General performance	177
8.3.2	Modification of the Reference System	180
8.3.3	Optimization of the adaptive gain	183
8.3.4	Disturbance Rejection	184
8.3.5	Adaptive Control using Multiple Models	189
	Nomenclature	193

Chapter 1

Introduction

To control a process means to influence its behavior in order to achieve a desired goal. A control system is one in which some quantities of interest are directed towards a prescribed value despite of disturbances affecting the system. Examples can be found in virtually any field of engineering and science. They range from ordinary applications such as the temperature control in buildings to sophisticated flight control of an aircraft or motion control of a robot manipulator. A short history serves to elucidate the train of thought that led to the scientific exploration of systems and control and illustrates how the qualifier “adaptive” entered the picture.

The central tool of control is the use of feedback which is introduced to correct for deviations of a system from some desired behavior. If the system as well as the environment in which it operates is known completely, it is possible to control the system using only feedforward signals. The rationale for using feedback is that in the presence of perturbations, the output of the system may be altered in an unexpected way. It is then necessary to detect such a deviation and adjust the control input appropriately. The concept of feedback relies upon the fundamental elements: measurement, comparison and adjustment. It formalizes the common-sense procedure of monitoring the output of a system, comparing it with its desired value and respond to errors by adjusting the input. Emerging from mechanical engineering, control theory became an independent discipline in the late 1940’s, when Wiener introduced the term “cybernetics” as the science of communication and control in the animal and the machine. According to Wiener, the homeostasis (equilibrium) of the body temperature is maintained through the use of feedback control. In the early 1960’s, Bellman felt that control theory had already become a mathematical discipline that could exist independent of its applications (Bellman, 1961 [8]). The mathematical foundations of control theory can be traced back to the work of J.C. Maxwell who analyzed the stability of Watt’s flyball governor

in terms of differential equations describing the dynamics of the system (1868). His technique was to linearize the differential equations of motion to find the characteristic equation of the system. He studied the effect of the system parameters on stability and showed that the system is stable if the roots of the characteristic equation have negative real parts. In 1892, A.M. Lyapunov published a seminal paper on the “general problem of stability of motion” which can be regarded as the foundation of modern stability theory (Lyapunov, 1892 [42]). He not only introduced the basic definitions of stability that are in use today, but also proved many of its fundamental theorems. The basic philosophy of his so-called second or direct method comes from a physical observation: if the total energy of a system is continuously dissipated, then the system, whether it is linear or nonlinear, must eventually settle down to an equilibrium point. Thus, the stability of a system can be determined without actually solving the associated differential equations. Control theory made significant strides, with the use of frequency domain methods and Laplace transforms in the 1930s and 1940s and the development of optimal control methods and state space analysis in the 1950s and 1960s. The design techniques developed emphasize the importance of a mathematical model describing the process to be controlled. The model of a system must be “simple enough” so that it can be analyzed with available mathematical techniques, and “accurate enough” to describe the important aspects of the relevant dynamical behavior. As dynamical systems became more complex and controllers had to be faster and, at the same time, more accurate, simple models were replaced by more complex ones. Uncertainty, regarding inputs, parameter values and structure of the system increasingly entered the picture. When the uncertainties in the plant and environment are large, fixed feedback controllers were found to be inadequate. As a consequence, researchers became interested in new areas of automatic control, such as adaptive, self-optimizing and self-organizing systems.

The term adaptation is defined in biology as the “advantageous conformation of an organism to changes in the environment”. The term was introduced in control theory by Drenick and Shahbender [16] in 1957, who defined an adaptive system as one that monitors its own performance and adjusts its parameters in the direction of better performance. It was soon realized that the definition lacks uniqueness since feedback itself has the ability to adjust some quantities in order to improve the performance of a system. In fact, designing a system that performs well even in the presence of uncertainties was the principal driving force for inventing feedback. From a practical point of view, an adaptive system can be regarded as one that copes with large changes in process dynamics and disturbance characteristics. A second aspect is that adaptive techniques can be used to provide automatic tuning of controllers. In this thesis, adaptive systems are defined as the special class of nonlinear systems that arises when controlling a linear system with unknown parameters. The system can be thought of

as having two feedback loops. One corresponding to the controller structure that is used when the system is known and an additional one by which the parameters are adjusted. In some cases, the adjustment loop is regarded as a (vanishing) perturbation of the linear time-invariant closed-loop system. In fact, as the parameters converge, the adaptive system becomes linear asymptotically. Another important observation is that the adjustment loop is usually slower than the state feedback loop.

The part of adaptive control theory that is well established deals with parametric uncertainties in the representation of the system. These are due to the lack of knowledge of system parameters. Another class of uncertainty arises when structural information, e.g. regarding the order of the plant, is missing. By its very nature, structural uncertainty cannot be modelled and is hard to analyze. The questions that have been addressed in this context regard the extent to which parameter adaptive control can be applied even when structural uncertainty is present. The principal difficulty is to establish robust stability of the adaptive system. Adaptive solutions have been derived under fairly idealistic conditions. Researchers in the early 1980's have shown that very simple adaptive system may become unstable if the strong assumptions made in the stability proofs are violated. Numerous examples were provided in which adaptive strategies performed poorly, and it was generally felt that adaptive algorithms cannot be used for practical adaptive control. Following a paper by (Rohrs et al., 1985 [69]) an effort was made to analyze the mechanism that leads to instability and to establish conditions for robust stability. A viewpoint that is widely agreed upon is that disturbances may cause the estimates to drift along a solution manifold in parameter space until stability is lost (Åström, 1995 [4]). This manifold is given by the set of estimates that yield a zero identification error. To stop the migration process, conditions have to be imposed on the nature of the reference input. If the input signal is sufficiently rich the solution manifold reduces to an equilibrium point—which would prevent the drift process.

In (Anderson et al., 1986 [5]) an attempt was made to develop a uniform theory for the robust stability of adaptive systems. Most of the theoretical results obtained rely on approximations of the nonlinear equations describing the adaptive process. The separation of fast system and slow adaptive variables is used in averaging theory to derive conditions for stability when the parameters are already close to their tuned values (Kokotovic et al., 1985 [38]). All (fast) system variables are collected in the regression vector ϕ . Parameter adaptation depends upon the product of the regression vector and the estimation error ϵ . The key idea in averaging is to approximate the product $\phi\epsilon$ by its average value $\text{avg}\{\phi\epsilon\}$ calculated under the assumption that the parameters are constant. The local behavior of the solution is thus obtained by linearization of the averaged equations. The analysis is valid only in a neighborhood of the tuned solution, i.e. assuming that the estimates are already close to

those parameters which result in a stable closed-loop system. Furthermore, the averaged equations may possess equilibria (obtained when $\text{avg}\{\phi\epsilon\} = 0$) even if the exact equations do not have an equilibrium. The consequence of the analysis presented in (Anderson et al., 1986 [5]) for the design of an adaptive control system is two-fold: first, keep the adaptive gain small in order to decouple the time-scales of regressor and estimator. Second, enhance the degree of persistence of excitation by adding small magnitude sinusoids to the reference input.

In contrast to the work discussed above, here we are not interested in developing a theory of robust stability for adaptive systems. Due to the nonlinear nature of the problem such a theory has proven to be very elaborate while having few consequences for the adaptive control practitioner. In fact, most of the implications, such as keeping the adaptive gain small, are rather straightforward and do not require a universal theory. Moreover, the motivation for using adaptive control in industrial applications comes from the fact that adaptation (hopefully) improves the performance of the system when large uncertainty is present. While stability is a major concern for the control theorist, it does not in itself provide the reason why an adaptive controller should be used instead of a linear one, which, in many practical cases, achieves the same goal of stabilizing the system. A number of modifications of the adaptive law have been proposed in the literature, including dead-zone (Peterson and Narendra 1982, [64]), σ -modification (Ioannou and Kokotovic, 1983 [32]) and parameter projection (Praly, 1984 [68]). The idea is to eliminate the pure integral action of the adaptive law and to keep the parameter error bounded. In this thesis we will not delve into any of these approaches but merely *assume* that the system can be adaptively controlled in a stable fashion (using any of the above methods). Our focus is to investigate conditions under which the *performance* of the system can be improved when the latter is affected by three specific classes of disturbances:

- external disturbances,
- unmodelled dynamics, and
- rapid time-variations.

In chapters **two** and **three** of the thesis the foundations of adaptive control are presented consisting of the problem of parameter estimation and control based on the certainty equivalence principle. The chapters are tutorial in nature and aim to provide the reader with all aspects of adaptive control assuming that the plant operates under ideal conditions. Emphasis is placed on the theoretical requirements needed to establish stability. It is demonstrated that the methods proposed for linear systems with unknown coefficients have nonlinear counterparts provided that the state of the system evolves within a neighborhood of an equilibrium

point. Once a nonlinear control law which meets the control objectives has been shown to exist, a neural network is used to implement it. The theory presented is applied to control problems that arise in the context of electrical drive systems.

Chapter **four** deals with the first class of disturbances in the above list. The objective is to reject unknown external disturbances completely, provided that the latter can be expressed as the output of a homogeneous difference equation. A new development is presented in the case where the disturbance generating system is driven by a white noise sequence. It is shown that a stochastic disturbance $v(k)$ affecting the input of the system can be rejected using a controller of augmented order, if $v(k)$ is highly correlated with its past values. In addition, extensions to the nonlinear domain are discussed.

In chapter **five**, a new method for minimizing the control effort in adaptive systems is introduced. The success of adaptive control in a practical environment depends critically on the nature of the control input. Design considerations, which assist in obtaining a smooth input are therefore at the center of our discussion. Our first step is to demonstrate that the dead-beat approach which is the preferred design technique in most theoretical papers on adaptive control can be replaced by more practicable designs. In particular, it is shown that by exploiting some very basic prior information about the plant, the sensitivity of the control loop to noisy measurement can be reduced. Our second, and major step is to develop a new procedure for optimizing the adaptive gain used in the parameter update algorithm. The principal idea is that, in certainty equivalence control, the input $u(k)$ is computed using the most recent parameter estimate, i.e. its *a posteriori* value $\hat{\theta}(k)$. But this means that, when updating $\hat{\theta}(k-1)$, one may determine an *optimal* $\hat{\theta}(k)$ in the sense that control based on this parameter minimizes the control effort. The set of permissible $\hat{\theta}(k)$ s is defined by the regression vector and the adaptive gain $0 < \eta < 2$. We formulate a constrained optimization problem which minimizes the cost due to the rate $|u(k) - u(k-1)|$ and magnitude $|u(k)|$ of the input by determining the optimal gain η^* at every instant of time. If the open-loop plant is stable, the procedure provides an effective way of limiting actuator stress. Since the optimized control input is also less oscillatory, the method enhances robustness to unmodelled high-frequency dynamics.

Chapter **six** is devoted to the problem of rapid time-variations in systems with unknown parameters which may be due to subsystem failures or large and sudden load variations. The objective is to outline the principal ideas of the multiple model architecture which has been established in recent literature (e.g. [54], [57]) and upon which our analysis in the following chapter is based. An example serves to highlight the importance of appropriately locating the models in a time-varying environment. The idea is that at least on model should be

close to the parameters that the plant is likely to assume. It is clear that the procedure is based on considerable prior information regarding the time-variation of the plant.

In chapter **seven** we consider the case where such prior information is not available. In other words, all models are located far from the plant. To keep the problem analytically tractable, it is assumed that the plant parameter θ assumes a finite number N of constant values contained in a compact set $S \subset \mathbb{R}^p$ but switches between the elements of S at random instants of time. Each parameter corresponds to an environment in which the plant operates, e.g. corresponding to a specific disturbance affecting the system. Even in this special case we have to assume that the time-variation is *regular* in some sense. In particular, it is assumed that the failures and load variations that occurred in the past are likely to occur in the future as well. A learning algorithm is proposed by which N models are updated simultaneously and in a competing fashion. The algorithm causes the configuration of N models to evolve gradually towards the configuration of the plant parameters. It is not known a priori which model converges to which plant parameter but no two models converge to same parameter. Although inspired by earlier work on motor learning and control [29], the algorithm as well as the questions we raise regarding the stability of the adaptive system are new and have not been addressed in the literature. The convergence is found to be a complex nonlinear process which cannot be analyzed using standard arguments from stability theory. In fact, the behavior of the algorithm is seen to be different from all adaptive processes known. In our analysis, numerous special cases and problems are considered which highlight the many aspects of the algorithm and set the stage for a succinct treatment of the problem. Our contribution is to provide qualitative insights into the nature of convergence and to develop a mathematical framework within which the problem can be addressed. The proof of convergence remains an open question.

In chapter **eight**, the effectiveness of the adaptive schemes developed in this thesis can be verified under real operating conditions. The objective is to control the speed of a rotational mass which is coupled to an electrical drive through a flexible shaft. The inertias of the drive and the load as well as the spring and damping constants are assumed to be unknown. Deterministic disturbances, in the form of an (unknown) piecewise constant or sinusoidal torque, are applied on the load side. The model used for control is of third-order and reflects only the (dominant) mechanical part of the two-mass system. Hence, robustness to unmodelled dynamics is a prerequisite for obtaining a stable closed-loop system. The performance of the system is improved using an optimized, time-varying adaptive gain as well as the multi-model approach of chapter six (in the case of a piece-wise constant disturbance). In a addition to the methods developed in the theoretical chapters of the thesis, a procedure for constraining the parameter estimates to a convex region in parameter space is presented,

which further enhances the robustness of the adaptive system. In all these cases, a self-tuning controller with fast and accurate response is obtained. The results serve to highlight the effectiveness of adaptive schemes in an industrial control problem.

Chapter 2

Parameter Estimation

When the parameters of a system are unknown, algorithms to estimate them naturally assume fundamental importance. Their discussion therefore forms the starting point of our investigation. The objective of this chapter is to review some of the standard estimation algorithms available in the literature (see e.g. Isermann 1992 [35]). However, since our ultimate goal is not system identification but adaptive control, the presentation focuses on those properties of the estimation algorithms that will become useful when the estimator is used as part of the adaptive controller.

Given a physical process, our objective is to derive a mathematical description for that process. Since we are mainly interested in dynamical systems, such a description will assume the form of a set of differential equations. If sufficient prior information is available, the coefficients may be determined using the laws of physics. If, on the other hand, the process is unknown, it may be possible to determine the coefficients by performing a series of experiments and observing the nature of the system's response. The problem then consists in adjusting the parameters of an appropriately chosen model, such as to match its response to that of the actual system. The success of this procedure critically depends upon the choice of the model structure as well as on the algorithm used to adjust the parameters.

Choice of a model structure: Clearly, a complete description of any physical system would have to be of infinite order. In practice however, one is only interested in the response of the system to specific input signals within a limited frequency range. This enables one to separate dominant and parasitic dynamics and describe the behavior of the system using reduced-order models. Moreover, for practical purposes, the system may be assumed to be linear. However, any physical system becomes nonlinear beyond a certain range of operation. Prior information is needed in order to decide whether the model should be linear or nonlinear.

The selection of the model is greatly influenced by the experience of the designer and depends upon the specific application. In the following, we assume that the model structure is given. The input–output characteristics of a large class of linear and nonlinear systems can be written in the following simple form:

$$y(k) = \phi(k-1)^T \theta_0 \quad (2.1)$$

where $\phi(k)$ is a linear or nonlinear function of a finite sequence of inputs and/ or outputs:

$$\begin{aligned} U(k) &= \{u(k), u(k-1), \dots, u(k-m)\} \\ Y(k) &= \{y(k), y(k-1), \dots, y(k-n)\} \end{aligned} \quad (2.2)$$

$\phi(k) \in \mathbb{R}^p$ (here: $p = m + n + 2$) is referred to as the regression vector. $\theta_0 \in \mathbb{R}^p$ corresponds to the unknown parameter vector which has to be estimated, and $y(k)$ is a scalar output, measured at instant of time k .

Example 2.1 The logistic equation, a model of population growth,

$$x(k+1) = a x(k)[1 - x(k)] = [x(k) \quad x^2(k)] \begin{bmatrix} a \\ -a \end{bmatrix} \quad (2.3)$$

is of the form (2.1). It displays chaotic behavior for e.g. $a = 3.8$.

Estimation Algorithm: The objective here is to use experimental data in order to determine the unknown parameters. In the *off-line* case, data is available prior to analysis and treated as a block of information. There is no constraint on the duration of the estimation process. In the *on-line* case, sequential data is used to update the parameter estimates recursively, using a relatively simple algorithm. A class of online–algorithms which has proven successful in practical applications has the following general form:

$$\hat{\theta}(k) = \hat{\theta}(k-1) + M_P(k-1)\phi(k-1)\epsilon(k) \quad (2.4)$$

where $\hat{\theta}(k)$ is the estimate of θ_0 at instant of time k . $M_P(k-1)$ denotes the algorithm gain and $\epsilon(k)$ is the modelling error

$$\epsilon(k) = y(k) - \hat{y}(k) \quad (2.5)$$

where $\hat{y}(k)$ is the output of the identification model

$$\hat{y}(k) = \phi(k-1)^T \hat{\theta}(k-1) \quad (2.6)$$

Depending upon the precise definitions of $\hat{\theta}(\cdot)$, $M_P(\cdot)$ and $\epsilon(k)$, the algorithm may take different forms. Some of them are discussed below.

2.1 Projection Algorithm

In this case,

$$M_P(k) = \frac{1}{\phi(k)^T \phi(k)} \quad (2.7)$$

The algorithm can be motivated geometrically as follows. Given $y(k)$ and $\phi(k-1)$, all possible parameter values θ^* satisfying equation (2.1) lie on the hyperplane

$$H = \{\theta^* \mid y(k) = \phi(k-1)^T \theta^*\} \quad (2.8)$$

Among all these candidate values for θ_0 , we choose $\hat{\theta}(k)$ as the element on H which is closest to $\hat{\theta}(k-1)$. In other words, $\hat{\theta}(k)$ is the orthogonal projection of $\hat{\theta}(k-1)$ onto H . Notice that the signed distance of $\hat{\theta}(k-1)$ to H is given by

$$d_{\hat{\theta}}(k-1) = -\frac{\phi(k-1)^T \tilde{\theta}(k-1)}{\|\phi(k-1)\|} \quad (2.9)$$

where $\tilde{\theta}(k-1) = \hat{\theta}(k-1) - \theta_0$. Since $\phi(k-1)$ is orthogonal on H , the projection operator used to update $\hat{\theta}(k-1)$ can be defined as follows

$$\text{Proj}\{\tilde{\theta}(k-1)\} = \frac{\phi(k-1)}{\|\phi(k-1)\|} d_{\hat{\theta}}(k-1) = -\frac{\phi(k-1) \phi(k-1)^T}{\phi(k-1)^T \phi(k-1)} \tilde{\theta}(k-1)$$

This is illustrated in figure (2.1).




Figure 2.1: Projection Algorithm

We obtain

$$\hat{\theta}(k) = \hat{\theta}(k-1) - \frac{\phi(k-1) \phi(k-1)^T}{\phi(k-1)^T \phi(k-1)} \tilde{\theta}(k-1) \quad (2.10)$$

Note that from (2.1) and (2.6) we have

$$\epsilon(k) = \phi(k-1)^T[\theta_0 - \hat{\theta}(k-1)] = -\phi(k-1)^T\tilde{\theta}(k-1) \quad (2.11)$$

A problem with the projection algorithm is that there is a (remote) possibility of division by zero when $\phi(k-1) = 0$. This can be easily overcome by adding a small positive constant in the denominator of (2.10)

$$\hat{\theta}(k) = \hat{\theta}(k-1) + \eta \frac{\phi(k-1)\epsilon(k)}{c + \phi(k-1)^T\phi(k-1)} \quad (2.12)$$

where $0 < \eta < 2$ is the adaptive gain and $c > 0$, usually $c = 1$.

2.1.1 Elementary Properties

When the parameters are updated according to equation (2.12), we obtain:

$$\begin{aligned} \text{(i)} \quad & \|\tilde{\theta}(k)\| \leq \|\tilde{\theta}(k-1)\| \leq \|\tilde{\theta}(0)\| \\ \text{(ii)} \quad & \lim_{N \rightarrow \infty} \sum_{k=1}^N \frac{\epsilon(k)^2}{c + \phi(k-1)^T\phi(k-1)} < \infty \end{aligned} \quad (2.13)$$

Property (i) corresponds to the fact that the parameter error is a nonincreasing function of time. In other words, $\hat{\theta}(k)$ is never farther from θ_0 than $\hat{\theta}(0)$. Property (ii) implies that the modelling error, $\epsilon(k)$, when appropriately normalized, is square summable. This can be used to show that the error does not grow faster than the signals contained in the regression vector ϕ . The following arguments are used in the proof of (i) and (ii): Subtracting θ_0 from both sides of equation (2.12) we obtain a nonlinear difference equation describing the evolution of the parameter error vector $\tilde{\theta}(k)$:

$$\tilde{\theta}(k) = \tilde{\theta}(k-1) - \eta \frac{\phi(k-1)\phi(k-1)^T}{c + \phi(k-1)^T\phi(k-1)} \tilde{\theta}(k-1) \quad (2.14)$$

Now, introduce the Lyapunov-function $V(k) = \|\tilde{\theta}(k)\|^2$. Using (2.11) we obtain

$$\|\tilde{\theta}(k)\|^2 - \|\tilde{\theta}(k-1)\|^2 = \eta \left[-2 + \eta \frac{\phi(k-1)^T\phi(k-1)}{c + \phi(k-1)^T\phi(k-1)} \right] \frac{\epsilon(k)^2}{c + \phi(k-1)^T\phi(k-1)} \quad (2.15)$$

Since, $0 < \eta < 2$, and $c > 0$, the bracketed expression is strictly less than zero. This already proves property (i). Hence, $\|\tilde{\theta}(k)\|^2$ is a bounded nonincreasing function which can be written as the following sum:

$$\|\tilde{\theta}(k)\|^2 = \|\tilde{\theta}(0)\|^2 + \eta \sum_{i=1}^k \left[-2 + \eta \frac{\phi(i-1)^T\phi(i-1)}{c + \phi(i-1)^T\phi(i-1)} \right] \frac{\epsilon(i)^2}{c + \phi(i-1)^T\phi(i-1)} \quad (2.16)$$

Since $\|\tilde{\theta}(k)\|^2$ is nonnegative whereas the sum is over negative semi-definite terms, property (ii) follows. It can be shown that a third property holds, which states that two successive parameter estimates approach each other as $k \rightarrow \infty$:

$$(iii) \quad \lim_{k \rightarrow \infty} \|\hat{\theta}(k) - \hat{\theta}(k-1)\| = 0 \quad (2.17)$$

The properties have great significance for the proof of stability in adaptive control since they have been derived under extremely weak conditions (in particular, $\phi(k)$ does not even have to be bounded). Note, however, that we have not said anything about $\hat{\theta}(k)$ necessarily converging to θ_0 . Further conditions on the nature of the signals contained in the regression vector $\phi(k)$ have to be imposed to ensure parameter convergence.

Example 2.2 In order to illustrate the above properties consider the problem of identifying a linear time-invariant system given by:

$$y(k+1) = -a_0y(k) - a_1y(k-1) + b_0u(k) \quad (2.18)$$

where $\theta_0 = [-a_0 \quad -a_1 \quad b_0]^T$ corresponds to the vector of unknown parameters and $\phi(k) = [y(k) \quad y(k-1) \quad u(k)]^T$. The plant is tested open-loop with a constant input signal $u(k) \equiv 5$. The simulation in figure (2.2) illustrates that $\tilde{\theta}(k) = \theta_0 - \hat{\theta}(k)$ is a nonincreasing sequence and that $\epsilon^2(k)/1 + \phi(k-1)^T \phi(k-1) \rightarrow 0$ irrespective of whether the plant (2.18) is stable or not. Note however that in both cases the parameter error does not converge to zero. Each discrete time-instant k corresponds to an interval of length $T_S = 10 \text{ ms}$, i.e. $t = kT_S$.

2.1.2 Parameter Convergence

The idea in establishing parameter convergence is to show that the parameter error $\tilde{\theta}(\cdot)$ is appropriately reflected in the identification error, such that if the latter is zero, the parameter error is zero as well. Intuitively, this means that $\tilde{\theta}(k-1)$ is observable through the identification error $\epsilon(k) = y(k) - \phi(k-1)^T \hat{\theta}(k-1)$.

In order to see this, we set $c = 0$ in equation (2.14) and write:

$$\tilde{\theta}(k) = H(k-1) \tilde{\theta}(k-1) \quad (2.19)$$

where

$$H(k-1) = I - \frac{\phi(k-1)\phi(k-1)^T}{\phi(k-1)^T \phi(k-1)} \quad (2.20)$$

The solution at instant of time $k > 1$ can be given in the following form:

$$\tilde{\theta}(k) = \Phi(k, k_0) \tilde{\theta}(k_0) \quad (2.21)$$

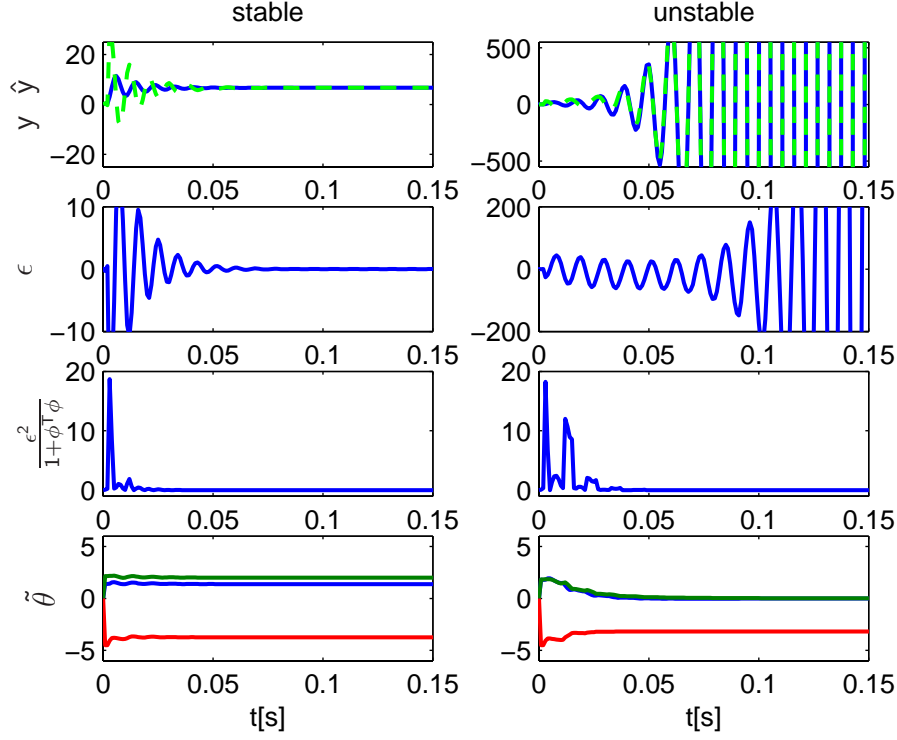


Figure 2.2: Properties of the Projection Algorithm

with the state transition matrix

$$\Phi(k, k_0) = H(k_0 + k - 1) \dots H(k_0 + 1) H(k_0) = \prod_{i=0}^{k-1} H(k_0 + i) \quad (2.22)$$

In the error equation (2.11), which is repeated here for convenience, $\phi(k-1)$ is a time-varying vector which maps the state $\tilde{\theta}(k-1)$ onto the output $\epsilon(k)$.

$$\epsilon(k) = -\phi(k-1)^T \tilde{\theta}(k-1) \quad (2.23)$$

Let us introduce the normalized identification error $\bar{\epsilon}(\cdot)$, [27]:

$$\bar{\epsilon}(k) = -\frac{\phi(k-1)^T \tilde{\theta}(k-1)}{\|\phi(k-1)\|} \quad (2.24)$$

The following matrix can be shown to be the “observability Gramian” for the state space system (2.19) with output equation (2.24):

$$G_O(k_0, k_0 + T) = \sum_{i=0}^{T-1} \frac{\Phi(k_0 + i, k_0)^T \phi(k_0 + i) \phi(k_0 + i)^T \Phi(k_0 + i, k_0)}{\phi(k_0 + i)^T \phi(k_0 + i)} \quad (2.25)$$

The projection algorithm is globally exponentially convergent to θ_0 provided that the following observability condition is satisfied

$$G_O(k_0, k_0 + T) \geq cI \quad c > 0 \quad \text{for all } k_0 \text{ and some fixed } T > 0 \quad (2.26)$$

To prove the statement, consider the Lyapunov function $V(k) = \tilde{\theta}(k)^T \tilde{\theta}(k)$. We know that

$$\begin{aligned} V(k+1) &= \tilde{\theta}(k)^T H(k)^T H(k) \tilde{\theta}(k) \\ &= V(k) - \frac{\tilde{\theta}(k)^T \phi(k) \phi(k)^T \tilde{\theta}(k)}{\phi(k)^T \phi(k)} \end{aligned}$$

Summing up T terms we obtain

$$\begin{aligned} V(k+T) &= V(k) - \tilde{\theta}(k)^T G_O(k, k+T) \tilde{\theta}(k) \\ &\leq (1-c)V(k) \quad c > 0 \end{aligned}$$

Since $V(k+T) > 0$, we have $c < 1$. Notice that,

$$1 - c < e^{-c}$$

Hence, $V(k+T) \leq e^{-c} V(k)$, or

$$V(k_n T) \leq e^{-ck_n} V(0) \quad (2.27)$$

where the index k_n refers to steps of length T . Thus $V(k)$ and hence $\tilde{\theta}(k)$ converges to zero exponentially fast. ■

It can be shown that the observability condition (2.26) is equivalent to the following simpler expression:

$$\sum_{i=0}^{T-1} \frac{\phi(k+i) \phi(k+i)^T}{\phi(k+i)^T \phi(k+i)} \geq cI \quad c > 0 \quad (2.28)$$

for all k and some fixed $T > 0$.

Intuitively, the condition means that over a fixed time interval T , the regression vector must have a nonzero projection along any direction of the parameter space. In other words, if $\tilde{\theta} \in \mathbb{R}^p$ then $\phi(k), \dots, \phi(k+T-1)$ must span the whole space \mathbb{R}^p . The condition is also known as *persistence of excitation* which reflects the fact that ϕ must be rich enough to excite all modes of the system that is to be identified. Notice that the condition refers to the regression vector not to the input of the system. The excitation can be provided by the input but may be lost in closed-loop since feedback introduces dependencies among the variables appearing in the regression vector (for further details, see e.g. Åström, Wittenmark 1995 [4]).

Example 2.3 In the above example the input to the system is modified to contain three distinct frequencies $u(k) = \sum_{i=1}^3 \sin(kT_S \omega_i)$, where T_S is the sampling time. It is seen that due to the richness of the probing signal, the parameter error $\tilde{\theta}$ converges to (a neighborhood of) zero. It should be kept in mind that even though the persistence of excitation condition is satisfied the convergence rate may be very slow. $\tilde{\theta}(k)$ vanishes for $k \rightarrow \infty$.




Figure 2.3: Parameter convergence due to richness of the input signal

Example 2.4 The fact that persistence of excitation need not always be due to the control input, can be appreciated if we return to the system (2.3) in example 2.1. It is well known that for $0 \leq a < 3$ the solution tends to a fixed point, while for values $a \geq 3$ it undergoes a sequence of bifurcations. Points of periods ≥ 2 appear, and at $a = 4$ we find points of all periods in the solution of the system, i.e. given any integer $i > 0$ we can find an initial condition $x(k_0)$ such that $x(k_0 + i) = x(k_0)$. The bifurcations are referred to as the *period-doubling route to chaos* (see e.g. Guckenheimer 1983, [28]). A model of the form $\hat{x}(k+1) = \phi(k)^T \hat{\theta}(k)$, with $\phi(k) = [x(k) \ x(k)^2]^T$ was used to identify the system. It is clear that at $a = 4$ the regression vector is sufficiently rich, causing the parameters in figure (2.4) to converge to the correct values while for $a < 3$, the identification error is zero without the parameters converging to the correct values.

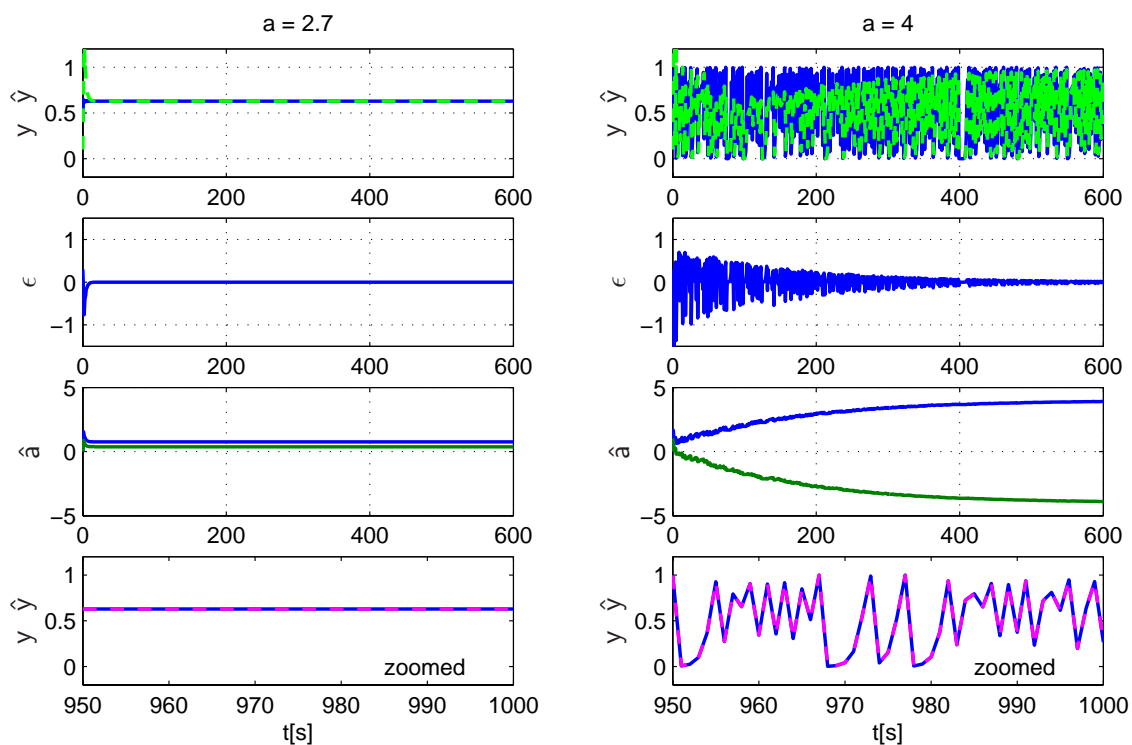


Figure 2.4: Identification of the logistic map

2.2 Orthogonalized Projection Algorithm

If the persistence of excitation condition (2.28) holds and the vector $\phi(k)$ is orthogonal to all its previous values $\phi(1), \dots, \phi(k-1)$ then the parameters converge in a finite number of steps. This can be verified in figure (2.1) where the parameters converge after only two projections, provided that $\phi(k)$ is orthogonal to $\phi(k-1)$. This suggests an improved version of the projection algorithm in which the directions along which the parameters are updated form a sequence of orthogonal vectors. It can be shown that there exists a time-varying matrix $P(k-1)$ such that $P(k-1)\phi(k)$ is the component of $\phi(k)$ which is orthogonal to $\phi(1), \dots, \phi(k-1)$. The algorithm is given by

$$\hat{\theta}(k) = \hat{\theta}(k-1) + \frac{P(k-2)\phi(k-1)\epsilon(k)}{\phi(k-1)^T P(k-2)\phi(k-1)} \quad (2.29)$$

where

$$P(k-1) = P(k-2) - \frac{P(k-2)\phi(k-1)\phi(k-1)^T P(k-2)}{\phi(k-1)^T P(k-2)\phi(k-1)} \quad (2.30)$$

with $P(0) = I$. In the event that $\phi(k-1)^T P(k-2)\phi(k-1) = 0$, we may set $\hat{\theta}(k) = \hat{\theta}(k-1)$ and $P(k-1) = P(k-2)$ in order to exclude unbounded solutions. The vector $P(k-2)\phi(k-1)$

in equation (2.29) is the component of $\phi(k-1)$ which is orthogonal to all previous $\phi(\cdot)$ vectors. A proof of this statement can be found in (Goodwin 1984, [27]). Its principal ideas are as follows. First it is shown that $P(k)P(k) = P(k)$, i.e. the matrix $P(k)$ is idempotent. Using equation (2.30) it is proven that

$$P(0) \dots P(k) = P(k) \quad (2.31)$$

The second step consists in showing that the null-space of $P(k)$ is equivalent to the span of $\phi(1) \phi(2) \dots \phi(k)$, i.e.

$$\ker[P(k)] = \text{span}[\phi(1) \phi(2) \dots \phi(k)] \quad (2.32)$$

Hence, $P(k)\phi(k) = 0$. But this means that $P(k)\phi(k-i) = P(k) \dots P(k-i)\phi(k-i) = 0$. Hence,

$$\phi(k-i)^T P(k-1)\phi(k) = \phi(k-i)^T P(k-i) \dots P(k-1)\phi(k) = 0 \quad (2.33)$$

This proves that $P(k-1)\phi(k)$ is orthogonal to all previous vectors $\phi(k-i)$, $i = 1, \dots, k-1$. In addition, it can be shown that $\tilde{\theta}(k)$ is orthogonal to $[\phi(1) \phi(2) \dots \phi(k-1)]$. This implies that if the regression vectors span the whole parameter space, $\tilde{\theta}(\cdot)$ must be zero. In other words, the parameter estimates converge to θ_0 in m steps provided that

$$\text{rank}[\phi(1) \phi(2) \dots \phi(m)] = p \quad (2.34)$$

where p is the dimension of the parameter space, i.e. $\theta_0 \in \mathbb{R}^p$.

2.3 Recursive Least-Squares Algorithm

The algorithm above has the usual disadvantage of a possible division by zero. By adding a small positive constant in the denominator of equation (2.29) we obtain a modified version which corresponds to the well known *Recursive Least-Squares Algorithm*:

$$\hat{\theta}(k) = \hat{\theta}(k-1) + \frac{P(k-2)\phi(k-1)\epsilon(k)}{1 + \phi(k-1)^T P(k-2)\phi(k-1)} \quad (2.35)$$

where

$$P(k-1) = P(k-2) - \frac{P(k-2)\phi(k-1)\phi(k-1)^T P(k-2)}{1 + \phi(k-1)^T P(k-2)\phi(k-1)} \quad (2.36)$$

with the initial estimate $\hat{\theta}(0)$, and any positive definite matrix $P(-1)$. If we compare (2.35) to (2.12) we see that the basic structure of the algorithms are equivalent and that $P(k-1)$ may be regarded as a time-varying matrix adaptive gain. Hence, it comes as no surprise that

the algorithm has properties very similar to those of the projection scheme. For convenience, these properties are listed below.

Given the system (2.1), if the parameters in (2.6) are updated according to the recursive least-squares algorithm (2.35) and (2.36) we have

$$\begin{aligned}
& \text{(i)} \quad \|\tilde{\theta}(k)\|^2 \leq \kappa \|\tilde{\theta}(0)\|^2 \quad k \geq 1 \\
& \text{(ii)} \quad \lim_{N \rightarrow \infty} \sum_{k=1}^N \frac{\epsilon(k+1)^2}{1 + \phi(k)^T P(k-1) \phi(k)} < \infty \\
& \text{(iii)} \quad \lim_{k \rightarrow \infty} \|\hat{\theta}(k) - \hat{\theta}(k-l)\| = 0 \quad \text{for any finite } l
\end{aligned} \tag{2.37}$$

where κ is the condition number of $P(-1)^{-1}$

$$\kappa = \frac{\lambda_{\max}[P(-1)^{-1}]}{\lambda_{\min}[P(-1)^{-1}]}$$

The proof is analogous to the one used in section (2.1.1) and is sketched here for completeness. Using $V(k) = \tilde{\theta}(k)^T P(k-1)^{-1} \tilde{\theta}(k)$ as a Lyapunov-function, we have

$$V(k) - V(k-1) = -\frac{\epsilon(k)^2}{1 + \phi(k-1)^T P(k-2) \phi(k-1)} \tag{2.38}$$

Hence, for all $k \geq 1$,

$$\tilde{\theta}(k)^T P(k-1)^{-1} \tilde{\theta}(k) \leq \tilde{\theta}(0)^T P(-1)^{-1} \tilde{\theta}(0) \tag{2.39}$$

Using the matrix inversion lemma (see e.g. Kailath 1980, [36]), (2.36) may be rewritten as

$$P(k)^{-1} = P(k-1)^{-1} + \phi(k) \phi(k)^T \tag{2.40}$$

It follows that the smallest eigenvalue $\lambda_{\min}[P(k)^{-1}]$ is nondecreasing. This implies

$$\begin{aligned}
\lambda_{\min}[P(-1)^{-1}] \|\tilde{\theta}(k)\|^2 & \leq \lambda_{\min}[P(k-1)^{-1}] \|\tilde{\theta}(k)\|^2 \\
& \leq \tilde{\theta}(k)^T P(k-1)^{-1} \tilde{\theta}(k) \\
& \leq \tilde{\theta}(0)^T P(-1)^{-1} \tilde{\theta}(0) \\
& \leq \lambda_{\max}[P(-1)^{-1}] \|\tilde{\theta}(0)\|^2
\end{aligned} \tag{2.41}$$

This establishes property (i). Next, let us sum the right hand side of (2.38) from 1 to N

$$V(N) = V(0) - \sum_{i=1}^N \frac{\epsilon(i)^2}{1 + \phi(i-1)^T P(i-2) \phi(i-1)} \tag{2.42}$$

Property (ii) follows from the fact that $V(N)$ is nonnegative. Finally, using straightforward arguments which can be found in (Goodwin, 1984 [27]), we have that the change of the parameter estimates tends to zero (Property (iii)).

In most practical applications the recursive least squares algorithm is the preferred estimation scheme since it is both fast and robust to measurement noise. The algorithm has fast initial convergence rate but slows down dramatically when the covariance matrix P gets small after a few iterations. This ensures parameter convergence even in the presence of a zero-mean, white noise signal. When the parameters are time-varying, it may be necessary to reset the matrix P at various times (preferably when one expects that a parameter change has occurred). This is accomplished by the least-squares algorithm with covariance resetting. Let $R = \{k_1, k_2, k_3, \dots\}$ be the instants of time at which the matrix is reset. Then for all $k \notin R$ equation (2.36) is used to update P . Otherwise, for $k = k_i \in R$

$$P(k) = cI \quad c \gg 1 \quad (2.43)$$

Since in between the reset instants the algorithm behaves as in the original case it is easy to show that it retains the above properties (i) – (iii). However, prior knowledge as to when parameter changes are likely to happen is required to determine the set of reset instants R . An alternative approach which is used frequently in time-varying contexts is to apply greater weight to more recent data by means of a “forgetting factor”. An algorithm which exponentially discards old data is given by:

$$\hat{\theta}(k) = \hat{\theta}(k-1) + \frac{P(k-2) \phi(k-1) \epsilon(k)}{\lambda(k-1) + \phi(k-1)^T P(k-2) \phi(k-1)} \quad (2.44)$$

with

$$P(k-1) = \frac{1}{\lambda(k-1)} \left[P(k-2) - \frac{P(k-2) \phi(k-1) \phi(k-1)^T P(k-2)}{\lambda(k-1) + \phi(k-1)^T P(k-2) \phi(k-1)} \right] \quad (2.45)$$

where $0 < \lambda(\cdot) < 1$. It has been shown in (Åström, 1980 [1]) that the forgetting factor λ must be time-varying in order to avoid “burst” phenomena if the data are not persistently exciting. In such a case the parameter estimator can become unstable since the P matrix begins to grow whenever the excitation is insufficient. Suggestions as to how to adjust the forgetting factor depending on the information content of the data have been made in (Åström, 1980 [1]; Cordero, 1981 [13]).

Finally, a condition similar to (2.28) can be given under which the parameter estimates, which are updated according to equations (2.35) and (2.36) converge to the true value θ_0 . Again, the regression vector has to be nonvanishing:

$$\lim_{k \rightarrow \infty} \sum_{i=0}^k \phi(i) \phi(i)^T = \infty \quad (2.46)$$

Example 2.5 A comparison of the different estimation schemes presented so far is most illuminating if the number of unknown parameters is large. The reason is that in a low-dimensional parameter space few steps are needed to attain a region in which the estimation error is small. In figure (2.5) the comparison is carried out using an unknown linear fourth-order system (i.e. eight unknown parameters). The probing input was again chosen to be a combination of sine-waves of different frequencies. It is seen that with the projection algorithm the identification error tends to zero but convergence is very slow. In the case of the orthogonalized projection algorithm we observe a zero crossing of $\phi(k-1)P(k-2)\phi(k-1)$ resulting in an infinite adaptation gain and instability. This is avoided when using the recursive least squares algorithm by which a fast and accurate response is obtained.

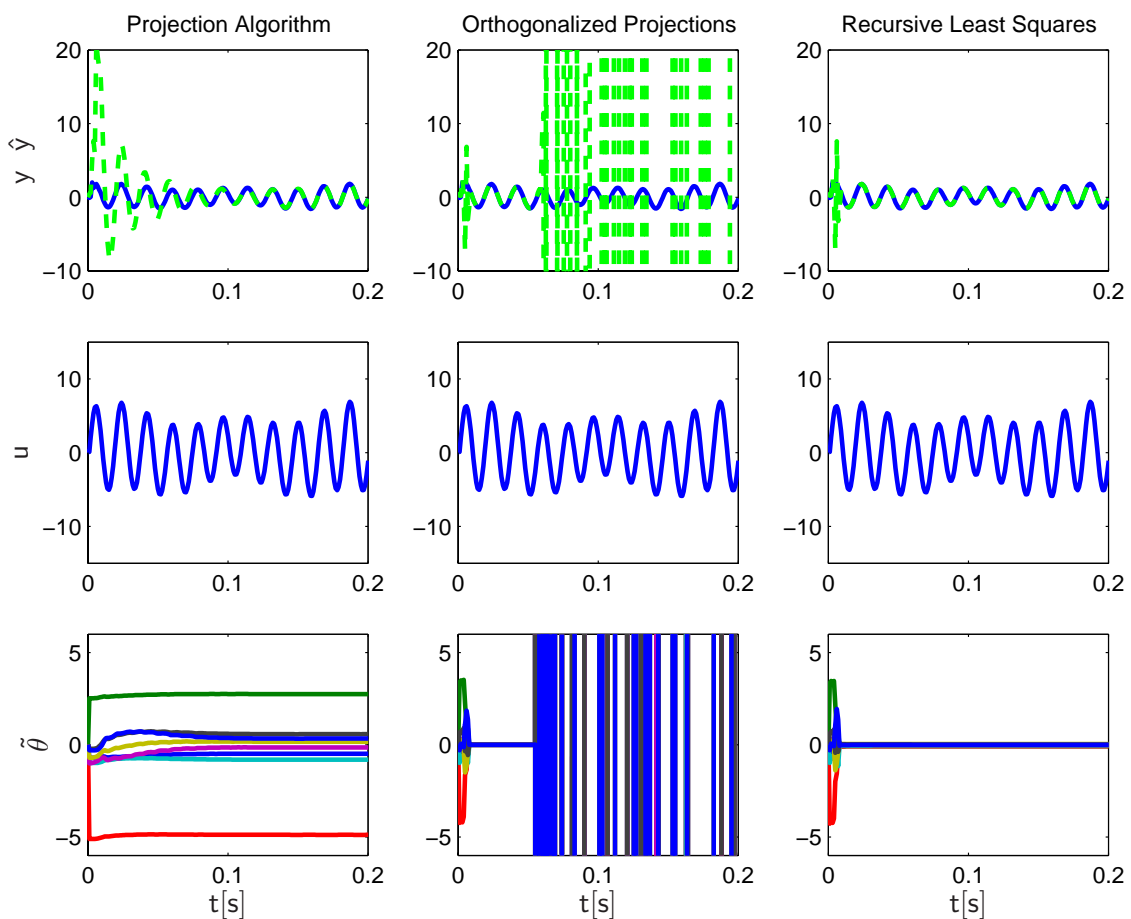


Figure 2.5: Comparison of different parameter estimation schemes

2.4 Arbitrary Relative Degree

In discrete time, the relative degree of a system has an intuitive meaning. It corresponds to the number d of sampling units after which an input signal $u(\cdot)$ affects the output $y(\cdot)$ of the system. In the above choice of the model structure (2.1) it has been assumed that $d = 1$. When d is arbitrary, the equation may be modified as follows:

$$y(k) = \phi(k-d)^T \theta_0 \quad (2.47)$$

where $\phi(k-d)$ is a function of the vectors $Y(k-d)$ and $U(k-d)$ defined in equation (2.2). Equivalently,

$$y(k) = H(q^{-1})\phi(k)^T \theta_0 \quad (2.48)$$

where $H(q^{-1}) = q^{-d}$ is a known filter transfer function. The estimation model is obtained by replacing θ_0 by the most recent parameter estimate, i.e. $\hat{\theta}(k-1)$. We end up with an error equation of the following form:

$$\epsilon(k) = -H(q^{-1})\phi(k)^T \tilde{\theta}(k-1) \quad (2.49)$$

where $\tilde{\theta}(k-1) = \hat{\theta}(k-1) - \theta_0$. This means that the estimation error $-\phi(k)^T \tilde{\theta}(k-1)$ is available as the output of some transfer function $H(q^{-1})$. Depending on whether this transfer function is strictly positive real or not equation (2.49) corresponds to error model three and four respectively, according to the classification introduced in (Narendra, 1989 [51]). The central question in both cases is how the parameters are to be updated such that the estimation algorithm is stable. If the transfer function $H(q^{-1})$ in equation (2.49) is not strictly positive real (i.e. error model four) a method involving the concept of the *augmented error* which is due to (Monopoli, 1974 [46]) can be used to derive a stable estimation algorithm. Define

$$\bar{\epsilon}(k) = \epsilon(k) + \tilde{\epsilon}(k) \quad (2.50)$$

where $\bar{\epsilon}(k)$ corresponds to the augmented error and $\tilde{\epsilon}(k)$ to an auxiliary signal of the form

$$\tilde{\epsilon}(k) = H(q^{-1})[\phi(k)^T \hat{\theta}(k)] - [H(q^{-1})\phi(k)]^T \hat{\theta}(k-1) \quad (2.51)$$

The reasoning behind this definition is that by means of the augmented error the standard model structure is recovered. This can be seen as follows: Since θ_0 is time-invariant we have

$$H(q^{-1})[\phi(k)^T \theta_0] = [H(q^{-1})\phi(k)]^T \theta_0 \quad (2.52)$$

This can be used to show that

$$\bar{\epsilon}(k) = -[H(q^{-1})\phi(k)]^T \tilde{\theta}(k-1) \quad (2.53)$$

which corresponds to the standard error equation (2.11). Using $H(q^{-1}) = q^{-d}$ we obtain:

$$\bar{\epsilon}(k) = -\phi(k-d)^T \tilde{\theta}(k-1) \quad (2.54)$$

Hence, by replacing $\epsilon(\cdot)$ by $\bar{\epsilon}(\cdot)$ and $\phi(k-1)$ by $[H(q^{-1})\phi(k)]$ in the above algorithms, the convergence analysis proceeds as in the standard case. It follows that $\bar{\epsilon}(k) \rightarrow 0$, i.e. the *augmented error* tends to zero. The question is what happens to the *actual estimation error* $\epsilon(k)$? From the properties of the estimation algorithms it is known that $\|\hat{\theta}(k) - \hat{\theta}(k-1)\| \rightarrow 0$. This implies that the estimation model is asymptotically time-invariant so that equality (2.52) holds with θ_0 replaced by $\hat{\theta}(k) = \text{const}$. It follows that the auxiliary error $\tilde{\epsilon}(k)$ defined in (2.51) tends to zero. In other words, $\epsilon(k) \rightarrow \bar{\epsilon}(k) \rightarrow 0$ as $k \rightarrow \infty$.

2.5 An Example from Industry

So far, parameter estimation has been introduced as a way of identifying unknown dynamical systems of the form (2.1). However, the ultimate objective is to use the parameter estimates in order to control the system. A straightforward idea would be to wait until the parameters converge and then design a (linear) standard controller assuming that the parameters have converged to the true values. The approach is highly attractive to industry since it seems to be a natural extension of existing (and well tested, hence trustworthy) design techniques to systems with unknown parameters.

Consider the problem of controlling the speed of a rotational mass with inertia J . A typical requirement is to keep the speed n of the mass constant even in the presence of an unknown load which is represented by the signal v . A discrete-time description of the system can be given as follows:

$$n(k+1) = n(k) + \frac{T_S}{J}[u(k) - v(k)] \quad (2.55)$$

where T_S is the sampling time and $u(k)$ is the torque generated by an electrical drive. Assuming that the torque of the load $v = v_0$ is constant, we build an identification model of the form

$$\hat{n}(k+1) = n(k) + \hat{\theta}_1(k) u(k) + \hat{\theta}_2(k) \quad (2.56)$$

The objective is that the parameter estimates $\hat{\theta}_1$ and $\hat{\theta}_2$ converge to the actual parameters of the system, i.e.

$$\begin{aligned} \hat{\theta}_1(k) &\rightarrow \frac{T_S}{J} \\ \hat{\theta}_2(k) &\rightarrow -\frac{T_S}{J} v_0 \end{aligned} \quad (2.57)$$

A stationary value $n(k) = n_0 = \text{const.}$ in equation (2.55) is obtained when $u(k) = v_0$. The torque u_0 required at the input of the system is equal to the torque of the load, which is regarded as a parameter in this context. Hence, $v_0 = u_0$ is known. On the other hand, any variation of the inertia J is not identifiable since the net torque driving the system is zero. In order to see the effect of the inertia we have to cause a change in speed n . If, e.g. $u(k) = u_0 > v_0$ we have constant acceleration and y increases linearly. Yet, it is still impossible to separate the effects of J and v_0 by observing the output of the system. The reason can be given as follows:

An input signal $u(k)$ is persistently exciting of order p if its two-sided spectrum is nonzero at p points (or more). The order of the excitation has to be at least equal to the number of unknown parameters. In order to separate the effects of the inertia and the load, it is necessary to increase the frequency content of the input signal e.g. by setting $u(k) = \sin(\omega_0 T_S k)$. The two-sided spectrum of the sine is nonzero at $\pm\omega_0$ and suffices to identify $p = 2$ unknown parameters. This is illustrated in figure (7.31) which displays the evolution of the estimates in the (two-dimensional) parameter space. Once the parameters have converged,




Figure 2.6: Parameter convergence for input signals of different orders of excitation

the physical parameters can be recovered. By a slight abuse of notation (“ $\hat{\theta}(\infty)$ ” corresponds to the value of $\hat{\theta}$ for which the estimation error is small in a sense specified by the designer) we may write:

$$J = \frac{T_S}{\hat{\theta}_1(\infty)} \quad v = \frac{\hat{\theta}_2(\infty)}{\hat{\theta}_1(\infty)} \quad (2.58)$$

Once J and v are known, the design of the controller proceeds using the usual frequency domain criteria known from classical linear control theory.

Two comments are in order.

- The approach is valid if identification and control are performed separately, i.e. the loop is closed only after completion of the estimation procedure. In principle, this is possible since speed control systems are open-loop stable. In practice, the primary motivation for using adaptive algorithms is to improve the performance of an existing (but poorly tuned) standard controller. The idea is to fine-tune the controller in an automatic fashion based on the information obtained while the plant is in operation. But this means that the parameters have to be adjusted even while controlling the plant. A popular viewpoint is to disregard the time-variation of the plant and to treat the system as a linear one. The argument provided is that adaptation is a slow process when compared to the evolution of the state variables of the system. The system is considered to be stable if the associated linear time-invariant system, obtained by “freezing” the current parameter estimates, is stable. The following example demonstrates that this is a hazardous approach which may lead to erroneous conclusions about the stability of a system.

Example 2.6 Given a (continuous-time) system with time-varying parameters:

$$\dot{x} = A(t)x \quad (2.59)$$

with

$$A(t) = \begin{bmatrix} -1 - 9 \cos^2(6t) + 6 \sin(12t) & 12 \cos^2(6t) + 4.5 \sin(12t) \\ -12 \sin^2(6t) + 4.5 \sin(12t) & -1 - 9 \sin^2(6t) - 6 \sin(12t) \end{bmatrix}$$

The characteristic polynomial is time-invariant

$$\chi(\lambda) = \lambda^2 + 11\lambda + 10 \quad (2.60)$$

and has the eigenvalues $\lambda_1 = -1, \lambda_2 = -10$. Hence, the “frozen system” is stable at every instant of time. One may be tempted to conclude that also the original system is stable. This is in contrast to the solution curve depicted in figure (2.7) which starts at the origin and tends to infinity.

- If the identification is indeed carried out in closed-loop the excitation necessary for parameter convergence is generated through the reference input. This may lead to a conflict since the signals required to identify the system are not always compatible with the process requirements. This is particularly true if the number of unknown parameters is large. If, in the example, the rotational mass tracks a piecewise constant reference speed, identification is only possible when the speed changes from one




Figure 2.7: Unstable linear time-varying system

constant value to another. In [39], an interesting solution has been proposed which involves monitoring the reference signal. The parameters of the controller are initialized such that the closed-loop system is stable and are only updated when a change in the reference signal takes place.

Parameter estimation as discussed in this chapter is fundamental for designing stable adaptive controllers. However, the example illustrates that the approach most commonly encountered in industrial applications is to employ parameter estimators as a mere supplement to an existing control setup. The concept is based on linear design principles and uses the knowledge of the physical parameters of the system (acquired through the estimation process). As a consequence, the identification procedure must take place off-line and no control action can be taken before the parameters converge. If on-line identification is used instead, the resulting problems of stability are often ignored. Furthermore, the controller cannot be applied with confidence unless one ensures that sufficiently rich probing signals are applied to the system. Prior testing may be required to see which probing signals result in fast convergence. The persistence of excitation condition has been introduced uniquely as a consequence of the design procedure which is linear in nature. It is seen that an unknown system can be controlled in a stable fashion using much less information. This is at the center of our discussion in the next chapter.

Chapter 3

Adaptive Control

The control of a dynamical system whose parameters are unknown is certainly among the most fascinating achievements of modern control theory. Concepts like “learning”, “self-tuning” and “artificial intelligence” have their roots in the solution of the adaptive control problem. Many different approaches to adaptive control have been proposed in the literature which can be broadly classified under two headings: Self-Tuning Regulators (STR) and Model Reference Adaptive Control (MRAC). In the latter case, the basic idea is to cause the system to behave like a given reference model. It is seen that the two approaches are closely related. In fact, a self-tuning regulator in which process zeros are cancelled can be interpreted as a model-reference adaptive controller. Stability of the resulting adaptive closed-loop system remained an open question for almost a decade and was finally resolved by (Narendra et al. 1980 [59]), (Goodwin et al. 1980 [26]) and (Morse 1980, [48]).

The design of an adaptive control system is conceptually simple. It consists of a parameter estimator in combination with a controller structure that would result in a stable closed-loop system if the parameters were known. In the design process, the estimates are used as if they were the true parameters. This approach is commonly referred to as the *Certainty Equivalence Principle*. By means of this principle it is possible to conceive a wide range of combinations of estimators and controllers. We will be focused on some of them which have proven stability and convergence properties. The design methodology adopted throughout the thesis may be classified an indirect one, since the evaluation of the control law is indirectly achieved via an identification model. The *indirect* method offers the possibility of extending the design to the case where multiple identification models are used. This will be of interest in Chapter 6. We shall see that even though the design relies on the parameters of an underlying identification model, it is not necessary that the estimates converge to the true parameters of the system. This is in contrast to the method presented at the end of the last

chapter where knowledge of the physical parameters of the system was mandatory. In some cases, it is possible to parameterize the system directly in terms of the controller parameters. In such a case the controller is directly estimated and hence obtained using a *direct* approach.

3.1 Linear Adaptive Control

Our discussion starts with a system given by the linear state equations

$$\begin{aligned}x(k+1) &= Ax(k) + bu(k) \\ y(k) &= c^T x(k)\end{aligned}\tag{3.1}$$

where $x(k) \in \mathbb{R}^n$ is the state of the system and $u(k), y(k) \in \mathbb{R}$ is the scalar input and output variable respectively. A is a $n \times n$ matrix, b and c are $n \times 1$ vectors. The system is assumed to be both controllable and observable. The objective is to design a controller such that $y(k)$ tracks an arbitrary bounded reference output $y^*(k)$ even as the parameters of the system contained in A , b and c are unknown. The following presentation is guided by the paper (Cabrera and Narendra, 1999 [9]). We first assume that the parameters of the system are known and consider the *exact tracking problem*.

3.1.1 Exact Tracking

Given the linear system (3.1), determine an integer N and a controller

$$u(k) = f_b^T x(k) + f_f y^*(k + N)\tag{3.2}$$

where $f_b \in \mathbb{R}^{n \times 1}$ and $f_f \in \mathbb{R}$ such that the closed-loop system satisfies:

1. $x(k+1) = [A + b f_b^T]x(k)$ is asymptotically stable.
2. For every initial condition $x(0)$ and arbitrary bounded reference sequence $\{y^*(k)\}$, $|y(k) - y^*(k)| = 0$ for $k \geq N$, i.e $y(k)$ is identical to $y^*(k)$ after N steps.

It turns out that N is equal to the relative degree d of the system, where d is defined as the integer satisfying the conditions $c^T b = c^T A b = \dots = c^T A^{d-2} b = 0$, $c^T A^{d-1} b \neq 0$. If no such integer exists, the transfer function of the system is identically zero, and the relative degree is defined to be ∞ .

A convenient way to analyze the system is to transform it into normal form. It is well known that for a linear system (3.1) with relative degree d there exists a global change of coordinates

$z = T^{-1}x$, where T^{-1} is a nonsingular matrix defined by:

$$T^{-1} = \begin{bmatrix} U \\ V \end{bmatrix}_{n \times n} \quad \text{with } U = \begin{bmatrix} c^T \\ c^T A \\ \vdots \\ c^T A^{d-1} \end{bmatrix} \quad V = \begin{bmatrix} v_1^T \\ v_2^T \\ \vdots \\ v_{n-d}^T \end{bmatrix} \quad (3.3)$$

where $\mathcal{U} := \text{span } U^T$ is a linear subspace of \mathbb{R}^n and V is a basis matrix for \mathcal{V} , which is a completion of \mathcal{U} to \mathbb{R}^n , i.e. $\mathcal{U} \oplus \mathcal{V} = \mathbb{R}^n$. In addition, V is chosen such that $v_i^T b = 0$ for $i = 1 \dots n - d$. The *normal form* of the system (3.1) is obtained as:

$$\begin{aligned} z_1(k+1) &= z_2(k) \\ z_2(k+1) &= z_3(k) \\ &\vdots \\ z_d(k+1) &= c^T A^d T z(k) + c^T A^{d-1} b u(k) \\ z_{d+1}(k+1) &= v_1^T A T z(k) \\ &\vdots \\ z_n(k+1) &= v_{n-d}^T A T z(k) \\ y(k) &= z_1(k) \end{aligned} \quad (3.4)$$

where $c^T A^{d-1} b \neq 0$ is the high-frequency gain of the system. The following notation is used: The first d components of the state vector $z(k)$ are denoted by $\xi(k)$ and correspond to

$$\xi(k) = [z_1(k) \ z_2(k) \ \dots \ z_d(k)]^T = [y(k) \ y(k+1) \ \dots \ y(k+d-1)]^T \quad (3.5)$$

whereas $\zeta(k) = [z_{d+1}(k) \ \dots \ z_n(k)]^T$ refers to the $n_d = n - d$ remaining state variables. Clearly, $z(k) = [\xi(k)^T \ \zeta(k)^T]^T$. From the state equation (3.4) we obtain the subsystem,

$$\zeta(k+1) = Q\zeta(k) + P\xi(k) \quad (3.6)$$

where $Q \in \mathbb{R}^{n_d \times n_d}$ and $P \in \mathbb{R}^{n_d \times d}$. If we constrain the output of the system to zero, $\xi(k) \equiv 0$, equation (3.6) is homogenous and referred to as the zero-dynamics of the system. The zero-dynamics describe the evolution of the internal state variables and play a fundamental role in the solution of the tracking problem based on inverse control. This is seen as follows:

Using the normal form, it is clear that $z_d(k+1) = y(k+d)$ is the variable directly affected by the control input. Assuming that the desired output $y^*(k+d)$ is known at instant of time k , the control input has the following form

$$u(k) = \frac{1}{c^T A^{d-1} b} [y^*(k+d) - c^T A^d x(k)] \quad (3.7)$$

The control law results in exact tracking after d steps. However, this does not yet imply that the closed-loop system is stable. Since $\xi(k)$ is bounded, the boundedness of $z(k)$ is assured if the internal states $\zeta(k)$ are bounded. From (3.6) it is clear that $\zeta(k)$ is bounded if and only if Q is an asymptotically stable matrix. In other words, the zero dynamics are asymptotically stable. Since the system is linear, this corresponds to the fact that it is also minimum-phase. In fact, if we assume that (3.1) is a minimal realization of the transfer function $H(z) = c^T(zI - A)^{-1}b$, then the eigenvalues of Q correspond to the zeros of $H(z)$.

Example 3.1 Given a third-order system of the form,

$$\begin{aligned} x(k+1) &= \begin{bmatrix} a_{11} & a_{12} & 0 \\ 0 & a_{22} & a_{23} \\ a_{31} & 0 & a_{33} \end{bmatrix} x(k) + \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} u(k) \\ y(k) &= [1 \ 0 \ 0] x(k) \end{aligned} \quad (3.8)$$

it is readily verified that $c^T b = 0$ and $c^T A b = a_{12}$, i.e. the relative degree is $d = 2$. The matrix which transforms the system into normal form is obtained as

$$T^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ a_{11} & a_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (3.9)$$

where $v_1^T = [0 \ 0 \ 1]$ in the last row of the matrix completes the linear subspace spanned by the first two rows. In addition, $v_1^T b = 0$. Carrying out the transformation yields

$$\begin{aligned} z(k+1) &= \begin{bmatrix} 0 & 1 & 0 \\ -a_{11}a_{22} & a_{11} + a_{22} & a_{12}a_{23} \\ a_{31} & 0 & a_{33} \end{bmatrix} z(k) + \begin{bmatrix} 0 \\ a_{12} \\ 0 \end{bmatrix} u(k) \\ y(k) &= [1 \ 0 \ 0] z(k) \end{aligned} \quad (3.10)$$

The control law, defined in equation (3.7), which results in exact tracking after two steps reads

$$u(k) = \frac{1}{a_{12}} \left(y^*(k+2) - \begin{bmatrix} a_{11}^2 & a_{11}a_{12} + a_{12}a_{22} & a_{12}a_{23} \end{bmatrix} x(k) \right) \quad (3.11)$$

The stability of the closed loop system depends upon the stability of the zero-dynamics which are given by the subsystem

$$z_3(k+1) = a_{33}z_3(k) + a_{31}z_1(k) \quad (3.12)$$

If $|a_{33}| < 1$ the system (3.8) is minimum-phase. This corresponds to the fact that the zero of the transfer function

$$H(z) = \frac{a_{12}(z - a_{33})}{z^3 - (a_{11} + a_{33} + a_{22})z^2 + (a_{11}a_{33} + a_{11}a_{22} + a_{22}a_{33})z - a_{31}a_{12}a_{23} - a_{11}a_{22}a_{33}} \quad (3.13)$$

is strictly stable. Hence, the control law (3.11) solves the exact tracking problem.

3.1.2 Auto-Regressive Moving-Average Representation

In most practical control problems, the state variables are not accessible and control has to be based only on inputs and outputs. In the following we derive a representation of the plant which expresses the output at any instant as a linear combination of the past n values of the input and n values of the output. From (3.1) it follows that

$$\begin{bmatrix} y(k) \\ y(k+1) \\ y(k+2) \\ \vdots \\ y(k+n-1) \end{bmatrix} = \begin{bmatrix} c^T \\ c^T A \\ c^T A^2 \\ \vdots \\ c^T A^{n-1} \end{bmatrix} x(k) + \begin{bmatrix} 0 \\ c^T b u(k) \\ c^T A b u(k) + c^T b u(k+1) \\ \vdots \\ c^T A^{n-2} b u(k) + \dots + c^T b u(k+n-2) \end{bmatrix} \quad (3.14)$$

Since the system is observable, the matrix

$$W_O = \begin{bmatrix} c^T \\ c^T A \\ c^T A^2 \\ \vdots \\ c^T A^{n-1} \end{bmatrix} \quad (3.15)$$

is nonsingular. Using the notation $Y_n(k) := [y(k), y(k+1), \dots, y(k+n-1)]^T$ and $U_n(k) := [u(k), u(k+1), \dots, u(k+n-1)]^T$ to denote a sequence of input/output values we obtain:

$$x(k) = W_O^{-1} Y_n(k) - W_O^{-1} L_U [U_{n-1}(k)] := L[Y_n(k), U_{n-1}(k)] \quad (3.16)$$

where $L_U [U_{n-1}(k)]$ corresponds to the last term in (3.14). Since $y(k+n) = c^T x(k+n)$,

$$y(k+n) = c^T A^n L[Y_n(k), U_{n-1}(k)] + \sum_{i=0}^{n-1} c^T A^{n-1-i} b u(k+i) \quad (3.17)$$

Finally, shifting the time axis by $n-1$ instants we obtain the autoregressive moving average model (ARMA) of the plant:

$$y(k+1) = \sum_{i=0}^{n-1} \bar{a}_i y(k-i) + \sum_{i=0}^{n-1} \bar{b}_i u(k-i) \quad (3.18)$$

where \bar{a}_i and \bar{b}_i are constants depending on the coefficients of the linear equation (3.17). It is seen that the system is completely determined by $2n$ parameters. The input at time k affects the output at the next instant of time, i.e. $d = 1$. In the case of a higher relative degree

$d > 1$, the expressions $c^T b, c^T A b, \dots, c^T A^{d-2} b$ in equation (3.14) are all zero. Equivalently, the first $d - 1$ elements of the last sum in (3.18) are zero. If the time axis in (3.18) is shifted by $d - 1$ instants we obtain

$$y(k+d) = \sum_{i=1-d}^{n-d} \bar{a}_i y(k-i) + \sum_{i=0}^{n-d} \bar{b}_i u(k-i) \quad (3.19)$$

Equation (3.18) can again be used to express $y(k+1), y(k+2), \dots, y(k+d-1)$ successively in terms of the past values of the inputs and outputs. This results in the general ARMA representation of the system (3.1) having arbitrary relative degree $d \geq 1$:

$$y(k+d) = \sum_{i=0}^{n-1} a_i y(k-i) + \sum_{i=0}^{n-1} b_i u(k-i) \quad (3.20)$$

It is seen that the ARMA model can be written using the following shorthand notation introduced in chapter 2:

$$y(k+d) = \phi(k)^T \theta_0 \quad (3.21)$$

where $\phi(k)$ is the regression vector of the system defined as

$$\phi(k) = \begin{bmatrix} Y_{n'}(k) \\ U_{n'}(k) \end{bmatrix} \quad (3.22)$$

with an obvious definition of $Y_{n'}(k)$ and $U_{n'}(k)$ similar to the one used in (3.16). θ_0 contains the $2n$ parameters of the system:

$$\theta_0 = [a_0, a_1, \dots, a_{n-1}, b_0, b_1, \dots, b_{n-1}]^T \quad (3.23)$$

Our control objective is to track an arbitrary bounded output sequence $\{y^*(k+d)\}$. If the parameters are known, the equation

$$y^*(k+d) = \phi(k)^T \theta_0 \quad (3.24)$$

implicitly defines the control law. We obtain,

$$\begin{aligned} u(k) &= g[y^*(k+d), Y_{n'}(k), \bar{U}_{n'}(k), \theta_0] \\ &= 1/b_0 [y^*(k+d) - a_0 y_k - \dots - a_{n-1} y_{k-n+1} - b_1 u_{k-1} - \dots - b_{n-1} u_{k-n+1}] \end{aligned} \quad (3.25)$$

where $\bar{U}_{n'}(k) = [u(k-1), u(k-2) \dots u(k-n+1)]$. The application of this control law results in exact tracking after d steps and bounded signals provided that the system is minimum-phase. If the parameters are unknown, how should (3.25) be modified? In particular, how should the parameters be updated such that the closed-loop system is stable? These are the questions addressed in the *adaptive control problem*.

3.1.3 Adaptive Control

The problem may be stated as follows: Determine a control law $g[y^*, Y_{n'}, \bar{U}_{n'}, \hat{\theta}]$, dependent on the estimated parameter values $\hat{\theta}$, as well as an adaptive procedure to update $\hat{\theta}$, such that

$$\lim_{k \rightarrow \infty} |y(k) - y^*(k)| = 0 \quad (3.26)$$

and all the signals in the system remain bounded. Certain assumptions have to be made concerning the plant in order to be able to determine a solution. These are summarized as follows:

- an upper bound on the order n of the system is known,
- the relative degree d is known exactly,
- the zeros of the system lie strictly inside the unit circle.

The last assumption is equivalent to the fact that the system be minimum-phase. It can be dropped when the problem is one of regulating the output at a constant value (set-point control) or if the desired output is a specified periodic signal. Here, we are interested in asymptotic tracking of an *arbitrary* bounded reference signal.

The solution to the adaptive control problem is based on a simple and intuitive concept, known as the *certainty equivalence principle*. It states that, at every instant of time, the parameter estimates are used as if they were the true parameters of the system. This enables us to combine any parameter estimation algorithm with an arbitrary stabilizing control law. In particular, $u(k)$ can be obtained using equation (3.24) with θ_0 replaced by $\hat{\theta}(k)$:

$$y^*(k+d) = \phi(k)^T \hat{\theta}(k) \quad (3.27)$$

Any of the parameter estimation algorithms introduced in chapter 1 can be used to generate $\hat{\theta}(k)$. Recall, that since the system is not explicitly parameterized in the control parameters, our approach is indirect, i.e. the control law is obtained indirectly from the system model. In other words, $\hat{\theta}(k)$ is determined from an underlying estimation procedure using the following estimation model:

$$\hat{y}(k) = \phi(k-d)^T \hat{\theta}(k-1) \quad (3.28)$$

Some care must be taken regarding the time index of the parameter estimate $\hat{\theta}(\cdot)$. If the current desired output is $y^*(k)$, an appropriate control input must have been generated d instants of time ahead, i.e. at $k-d$. At this stage, parameter estimates up to $\hat{\theta}(k-d)$ are available, see equation (3.27). This is in contrast to the identification problem where $u(\cdot)$

need not be determined but is regarded as a given input signal. Hence, at instant of time k , the control input $u(\cdot)$ up to $k - 1$ has already been determined. Consequently, a more recent parameter estimate is available, i.e. $\hat{\theta}(k - 1)$ in (3.28). Since equation (3.27) is also the equation for the closed loop system, we obtain the following relationship between the control error $e(k) = y(k) - y^*(k)$ and the identification error $\epsilon(k) = y(k) - \hat{y}(k)$:

$$e(k) = \epsilon(k) + \phi(k - d)^T [\hat{\theta}(k - 1) - \hat{\theta}(k - d)] \quad (3.29)$$

The last term measures the rate of change of $\hat{\theta}(k)$ and decays as the parameters converge. Notice that it is zero whenever $d = 1$. In such a case the control and identification errors are equivalent. The design procedure may be summarized as follows:

1. Determine the control law (3.25) which results in exact tracking if the parameters are known and use $\hat{\theta}(k)$ instead of θ_0 .

$$u(k) = g[y^*(k + d), Y_{n'}(k), \bar{U}_{n'}(k), \hat{\theta}(k)] \quad (3.30)$$

2. Update the parameter estimates using one of the algorithms described in chapter 1 which are of the form

$$\hat{\theta}(k) = \hat{\theta}(k - 1) + M_P(k - 1) \phi(k - 1) \epsilon(k) \quad (3.31)$$

where M_P is a (matrix) adaptation gain.

It turns out that, depending on the reference input, there exists a large number of parameters values θ^* that result in a stable closed-loop system and $\lim_{k \rightarrow \infty} |y(k) - y^*(k)| = 0$. In other words, it is not necessary that the parameters converge to θ_0 to solve the asymptotic tracking problem. This will become clear in the following proof of stability.

3.1.4 Proof of Stability

The tracking error is defined as $e(k) = y(k) - y^*(k)$. Using (3.29) we see that

$$\frac{|e(k)|}{[1 + \phi(k - d)^T \phi(k - d)]^{\frac{1}{2}}} \leq \frac{|\epsilon(k)|}{[1 + \phi(k - d)^T \phi(k - d)]^{\frac{1}{2}}} + \frac{\|\phi(k - d)\| \|\hat{\theta}(k - 1) - \hat{\theta}(k - d)\|}{[1 + \phi(k - d)^T \phi(k - d)]^{\frac{1}{2}}}$$

From the discussion of the properties of the parameter estimation algorithms in chapter 2 it follows that

$$\frac{|\epsilon(k)|}{[1 + \phi(k - d)^T \phi(k - d)]^{\frac{1}{2}}} \rightarrow 0 \quad (3.32)$$

and

$$\|\hat{\theta}(k-1) - \hat{\theta}(k-d)\| \rightarrow 0 \quad (3.33)$$

Hence,

$$\lim_{k \rightarrow \infty} \frac{|e(k)|}{[1 + \phi(k-d)^T \phi(k-d)]^{\frac{1}{2}}} = 0 \quad (3.34)$$

The fact that the system is minimum-phase means that its inverse is strictly stable. If h is the impulse response of the inverse system, we have that $h \in l_1$, i.e. there exist some fixed $m_1 > 0$ and $m_2 > 0$ such that for all $k > 0$,

$$|u(k-d)| \leq m_1 + m_2 \max_{1 \leq \kappa \leq k} |y(\kappa)| \quad (3.35)$$

In other words, the input is linearly bounded by the output of the system¹. It then follows from the definition of $\phi(k)$ in (3.22) that

$$\|\phi(k-d)\| \leq m_3 + m_4 \max_{1 \leq \kappa \leq k} |y(\kappa)| \quad \text{for some } m_3 > 0, m_4 > 0 \quad (3.36)$$

Since $|y(k)| = |e(k) + y^*(k)| \leq |e(k)| + |y^*(k)|$, and $|y^*(k)|$ is bounded, we have

$$\|\phi(k-d)\| \leq c_1 + c_2 \max_{1 \leq \kappa \leq k} |e(\kappa)| \quad \text{for some } c_1 > 0, c_2 > 0 \quad (3.37)$$

In equation (3.34), the norm of the regression vector $\phi(k-d)$ can be either bounded for all k , or grow in an unbounded fashion. In the former case, it directly follows that $|e(k)| \rightarrow 0$ as $k \rightarrow \infty$. If $\|\phi(k-d)\|$ grows in an unbounded fashion, then, by equation (3.37), it cannot grow faster than $|e(k)|$. However, this leads to a contradiction of equation (3.34). Hence, $\|\phi(k-d)\|$ is bounded and

$$\lim_{k \rightarrow \infty} |e(k)| = 0 \quad (3.38)$$

Comment: Inequality (3.35) can also be obtained under slightly more general conditions.

- All modes of the inverse system lie inside *or on* the unit circle. Any modes on the unit circle are of multiplicity one, *and*
- all controllable modes of the inverse system lie strictly inside the unit circle.

This is true because the uncontrollable state variables are bounded if the roots on the unit circle are simple. Thus, by superposition, only the controllable part of the system needs to be considered [27]. The relaxation will permit us to extend the adaptive procedure to a system of augmented order in chapter 4.1.

¹Let S be the linear space of all sequences $\{x_k\}_{k \geq 0}$. l_p is defined as the space of sequences for which the series $\sum_{k=0}^{\infty} |x_k|^p$ converges. It is a subspace of S for all $p \geq 1$: $l_p = \{x \in S \mid \sum_{k=0}^{\infty} |x_k|^p < \infty\}$. A system satisfying (3.35) is said to be finite gain stable (Desoer and Vidyasagar, 1975 [15]).

Example 3.2 A two-mass system consists of two rotational masses coupled by a flexible shaft. The objective is to control the speed of one mass by the torque applied to the other mass. The dynamics of the system may be described by a third-order differential equation. Furthermore, it is easy to see that the system is minimum-phase (for more details see chapter 8). All parameters of the system are unknown. It turns out that if a zero-order hold sampling method is employed to obtain a discrete-time state space description of the system, all entries of the corresponding vectors c and b are nonzero. This means that the relative degree of the sampled system is $d = 1$, since $cb \neq 0$. Furthermore, if the speed of one of the masses is used as the output variable, the system is observable. This is all the information needed to control the system using an underlying estimation model of the following structure:

$$\hat{y}(k+1) = \phi(k)^T \hat{\theta}(k) \quad (3.39)$$

where $\phi(k) = [y(k) \ y(k-1) \ y(k-2) \ u(k) \ u(k-1) \ u(k-2)]^T$ and $\hat{\theta}(k) \in \mathbb{R}^6$. Here, $u(\cdot)$ corresponds to the torque applied to one mass, and $\hat{y}(\cdot)$ is an estimate of the speed of the other mass. If $\hat{y}(k+1)$ is replaced by the desired value $y^*(k+1)$, equation (3.39) implicitly defines the control law. The simulation results contained in figure (3.1) illustrate that – starting from arbitrary initial estimates $\hat{\theta}(0)$ – the adaptive control law results in stable, asymptotic tracking of a bounded reference signal. The result is independent of the fact whether the parameters converge to the true values (as on the right side of figure 3.1) or not as in the case of a sine reference.

3.2 Nonlinear Adaptive Control Using Neural Networks

The results derived above apply to linear systems. In many cases linear approximations may prove to be adequate for purposes of analysis and synthesis (of the controller). However, all physical systems become nonlinear beyond a certain range of operation. The question naturally arises as to whether the concepts and tools developed for linear systems may be extended to the nonlinear domain. In such a case, the class of nonlinear systems, the domain in which the representation of the system is valid, and assumptions made for identification and control need to be clarified.

It is assumed that the nonlinear system to be controlled can be described by the state equations

$$\begin{aligned} x(k+1) &= f[x(k), u(k)] \\ y(k) &= h[x(k)] \end{aligned} \quad (3.40)$$

where $x(k) \in \mathbb{R}^n$ is the state of the system, $u(k) \in \mathbb{R}$ denotes the input, and $y(k)$ the output of the system. $f : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ and $h : \mathbb{R}^n \rightarrow \mathbb{R}$ are assumed to be unknown smooth functions

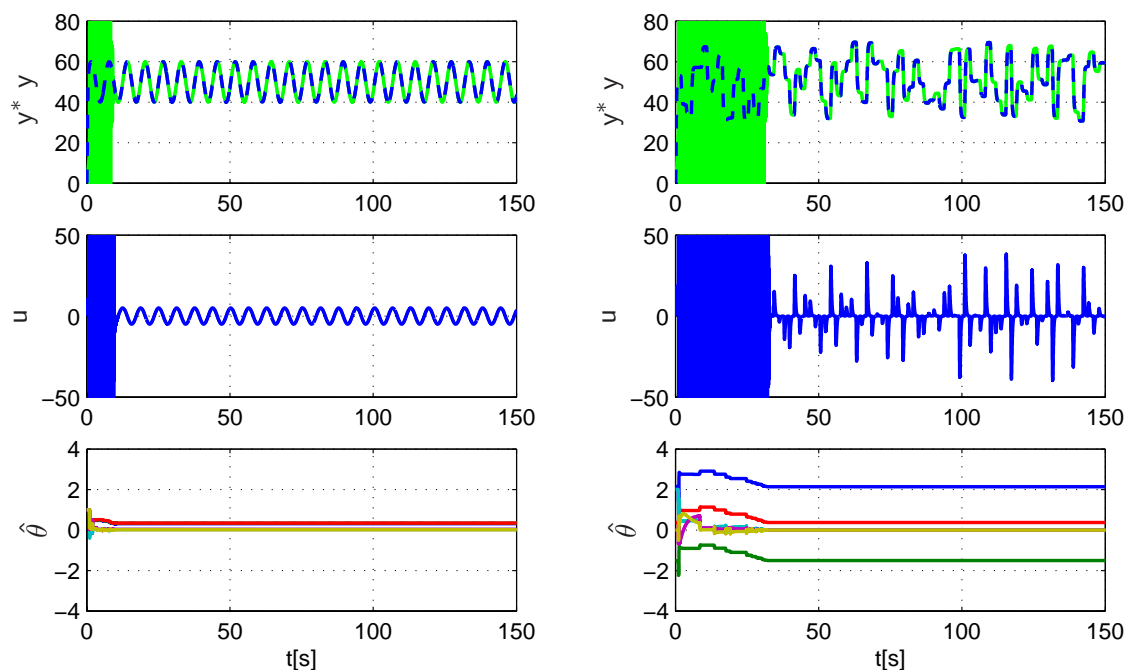


Figure 3.1: Adaptive control of a two-mass system

$f, h \in \mathcal{C}^\infty$. (\mathcal{C}^n is the class of n -times continuously differentiable functions). Further, it is assumed that $f[0, 0] = 0, h[0] = 0$, i.e. the origin is an equilibrium state of (3.40).

As in the linear case, the objective is to track an arbitrary output sequence $\{y^*(k)\}$ using a nonlinear controller. Only the input $u(k)$ and output $y(k)$ are assumed to be accessible. The questions that arise in this context are whether a suitable mathematical input/output-representation exists, what structure should be used for the controller and the manner in which identification and control are to be carried out. Considerable work has been reported on all the above questions and we include in this section those that serve to highlight the principal ideas of neurocontrol. A precise statement of the problem may be given as follows. Given the system (3.40), determine an integer N , a scalar $\Delta > 0$, a neighborhood \mathcal{X} of the origin in \mathbb{R}^n and an analytic function $g : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$, with $g(0, 0) = 0$, defining the control law

$$u(k) = g[x(k), y^*(k + N)] \quad (3.41)$$

such that the closed-loop system satisfies the following conditions:

1. $x(k + 1) = f[x(k), g[x(k), 0]]$ is asymptotically stable.
2. For every initial condition $x(0) \in \mathcal{X}$, and every reference sequence $\{y^*(k)\}$ bounded by Δ for all k , $|y(k) - y^*(k)| = 0$ for $k \geq N$.

Paralleling our discussion in the linear case, we have stated the *exact tracking problem*. In the asymptotic tracking problem, the second item is replaced by the condition that the error tends to zero asymptotically, i.e. $\lim_{k \rightarrow \infty} |y(k) - y^*(k)| = 0$. It is seen that the exact tracker is unique while the asymptotic tracker is not. However, the problems are equivalent in that the conditions for obtaining a stable tracking control law are the same in both cases. Determining the conditions that (3.40) has to satisfy to ensure the existence of a control law $g[x(\cdot), y^*(\cdot)]$ is the key to the solution of the nonlinear adaptive control problem [9].

All of the following statements are local in nature, i.e. they depend upon the linearization of the nonlinear system around the equilibrium point:

$$\begin{aligned} x(k+1) &= Ax(k) + bu(k) \\ y(k) &= c^T x(k) \end{aligned} \quad (3.42)$$

where $\left. \frac{\partial f(x,u)}{\partial x} \right|_{0,0} = A \in \mathbb{R}^{n \times n}$, $\left. \frac{\partial f(x,u)}{\partial u} \right|_{0,0} = b \in \mathbb{R}^n$ and $\left. \frac{\partial h(x)}{\partial x} \right|_0 = c \in \mathbb{R}^n$. It is assumed that the linearized system is both controllable and observable, so that the nonlinear system is also controllable and observable in a neighborhood of the origin in \mathbb{R}^{n+1} .

3.2.1 Nonlinear Normal Form

In order to derive a normal form for linear systems, the concept of the relative degree was found to have central importance. It corresponds to the number of sampling instants, after which the effect of a control input $u(k)$ can be felt at the output. For nonlinear systems, a local definition of the relative degree d may be derived.

The system (3.40) is said to have relative degree d at $(x, u) = (0, 0)$ if there exists a neighborhood $\mathcal{X} \times \mathcal{U}$ of the equilibrium state such that for all $(x, u) \in \mathcal{X} \times \mathcal{U}$:

$$\begin{aligned} \frac{\partial(h \circ f^i \circ f[x(k), u(k)])}{\partial u(k)} &= 0 \quad \text{for } 0 \leq i \leq d-2 \\ &\neq 0 \quad \text{for } i = d-1 \end{aligned} \quad (3.43)$$

where f^i is the i -times composition of f . If there exists a $i^* \geq 0$ and a neighborhood $\mathcal{X}_* \times \mathcal{U}_*$ of $(0, 0)$ such that

$$\begin{aligned} \frac{\partial(h \circ f^{i^*} \circ f[x(k), u(k)])}{\partial u(k)} &= 0 \quad \text{at } (0, 0) \\ &\neq 0 \quad \text{for } (x, u) \in \mathcal{X}_* \times \mathcal{U}_* \setminus \{(0, 0)\} \end{aligned} \quad (3.44)$$

the relative degree is not well defined. Since our ultimate aim is to develop general methods for controlling nonlinear systems, we shall assume in what follows that the relative degree

of the system is well defined and known. In such a case it is possible to obtain a system of local coordinates in which the system has normal form:

$$\begin{aligned}
z_1(k+1) &= z_2(k) \\
&\vdots \\
z_d(k+1) &= F[\zeta(k), \zeta_0(k), u(k)] \\
\zeta_0(k+1) &= D[\zeta(k), \zeta_0(k), u(k)] \\
y(k) &= z_1(k)
\end{aligned} \tag{3.45}$$

where $\zeta = [z_1, \dots, z_d]^T$, $\zeta_0 = [z_{d+1}, \dots, z_n]^T$ and $z = [\zeta^T \ \zeta_0^T]^T$. The equation is the equivalent of the normal form defined in (3.4) for linear systems. $F : \mathbb{R}^d \times \mathbb{R}^{(n-d)} \times \mathbb{R} \rightarrow \mathbb{R}$ and $D : \mathbb{R}^d \times \mathbb{R}^{(n-d)} \times \mathbb{R} \rightarrow \mathbb{R}^{(n-d)}$ are analytic functions of their arguments. Also, $\left. \frac{\partial F(\zeta, \zeta_0, u)}{\partial u} \right|_{0,0} \neq 0$.

The intuition gained in the linear case suggests that $\zeta_0(k)$ in equation (3.45) corresponds to the internal state variables of the system. This can be verified as follows. Given an initial state $x(0) \in \mathcal{X} \subset \mathbb{R}^n$, assume that a feedback controller has been designed such that $y(k) \equiv 0$. The transformed system reduces to

$$z_d(k+1) = F[0, \zeta_0(k), u(k)] = 0 \tag{3.46}$$

$$\zeta_0(k+1) = D[0, \zeta_0(k), u(k)] \tag{3.47}$$

The control input $u(k)$ corresponding to the problem of zeroing the output is only a function of $\zeta_0(k)$. From (3.47), we obtain $u(k) = g_0[\zeta_0(k)]$ which leads to

$$\zeta_0(k+1) = D[0, \zeta_0(k), g_0[\zeta_0(k)]] \quad \zeta_0(0) \in \mathcal{X}_0 \tag{3.48}$$

The autonomous dynamical system (3.48) is called the zero dynamics of the system. We *assume* that the partial state vector $\zeta_0(k)$ evolves in some neighborhood \mathcal{X}_0 of the origin in order to ensure the existence of the corresponding maps. It describes the dynamics of the system when the input and the initial condition are jointly chosen such as to make the output identically zero. Asymptotic stability of the zero-dynamics implies that the internal state of the system is bounded. As seen in the linear case, this is a prerequisite for a solution of the tracking problem to exist.

3.2.2 Control Using Input-Output Data

In section 3.1.2 it was shown that if a linear system is observable, it can be described using only input-output data. The corresponding statement relating to the observability of the nonlinear system (3.40) are substantially more complex. However, it has been shown

in (Levin and Narendra 1996, [44]) that if the linearization (3.42) is observable, then the nonlinear system is observable in a neighborhood \mathcal{X} of the equilibrium state. This implies that any state $x(k) \in \mathcal{X}$ can be uniquely determined by probing the system with an input sequence of sufficient length and observing the corresponding output sequence. From equation (3.40) we have

$$\begin{aligned} y(k) &= h[x(k)] = \Psi_1[x(k)] \\ y(k+1) &= h[f(x(k), u(k))] = \Psi_2[x(k), u(k)] \\ &\vdots \\ y(k+n-1) &= h \circ f^{n-1}[\cdot] = \Psi_n[x(k), u(k), u(k+1), \dots, u(k+n-2)] \end{aligned} \quad (3.49)$$

where f^{n-1} denotes the $(n-1)$ -times iterated composition of f . Using the notation $Y_n(k) := [y(k), y(k+1), \dots, y(k+n-1)]^T$ and $U_n(k) := [u(k), u(k+1), \dots, u(k+n-1)]^T$, we may express (3.49) as follows

$$Y_n(k) = \Psi[x(k), U_{n-1}(k)] \quad (3.50)$$

where $\Psi : \mathcal{X} \times \mathcal{U}_{n-1} \rightarrow \mathcal{Y}_n$ and $\mathcal{U}_{n-1}, \mathcal{Y}_n$ are the sets of all input/output sequences of length $n-1$ and n respectively. The key question is whether an inverse of the function Ψ exists so that, as in the linear case, the state $x(k)$ can be explicitly expressed in terms of $Y_n(k)$ and $U_{n-1}(k)$. It is known that a continuously differentiable mapping f is invertible in a neighborhood of any point x_0 at which its linearization is invertible. Furthermore, $f(x, y) = 0$ can be solved for x in terms of y in a neighborhood of any point (x_0, y_0) at which $f(x_0, y_0) = 0$ and $\frac{\partial f}{\partial x} \neq 0$. The following theorems make these statements precise (see e.g. Rudin 1964, [70])

The inverse function theorem

Let U be an open set in \mathbb{R}^n and let $f : U \rightarrow \mathbb{R}^n$ be a \mathcal{C}^k function with $k \geq 1$. If the Jacobian $\frac{\partial f}{\partial x}|_{x_0}$ is invertible at some point $x_0 \in U$, then there exists an open neighborhood \mathcal{V} of x_0 in U such that $f : \mathcal{V} \rightarrow f(\mathcal{V})$ is invertible with a \mathcal{C}^k inverse. \square

The implicit function theorem

Let U be an open set in $\mathbb{R}^n \times \mathbb{R}^m$ and let $f : U \rightarrow \mathbb{R}^n$ be a \mathcal{C}^k function with $k \geq 1$. Consider a point $(x_0, y_0) \in U$ where $x_0 \in \mathbb{R}^n$ and $y_0 \in \mathbb{R}^m$ with $f(x_0, y_0) = 0$. If the Jacobian $\frac{\partial f}{\partial x}|_{x_0, y_0}$ is invertible at (x_0, y_0) , then there exist open sets $\mathcal{V}_n \subset \mathbb{R}^n$ and $\mathcal{V}_m \subset \mathbb{R}^m$ with $(x_0, y_0) \in \mathcal{V}_n \times \mathcal{V}_m \subset U$ such that to every $y \in \mathcal{V}_m$ there corresponds a unique x such that

$$(x, y) \in U \quad \text{and} \quad f(x, y) = 0$$

If this x is defined to be $\psi(y)$, then ψ is a \mathcal{C}^k mapping $\psi : \mathcal{V}_m \rightarrow \mathcal{V}_n$ with $\psi(y_0) = x_0$ and

$$f(\psi(y), y) = 0 \quad \text{for } y \in \mathcal{V}_m$$

□

Hence, equation (3.50) can be solved for $x(k)$ provided that the Jacobian $\frac{\partial \Psi}{\partial x(k)}|_{0,0}$ is invertible. But the Jacobian is merely the observability matrix of the linearized system and hence (3.40) is locally observable if (3.42) is observable. Since the latter has been assumed to be true it follows that

$$x(k) = \psi[Y_n(k), U_{n-1}(k)] \quad (3.51)$$

in a neighborhood of the equilibrium state and $\psi \in \mathcal{C}^1 : \mathbb{R}^n \times \mathbb{R}^{n-1} \rightarrow \mathbb{R}^n$. From (3.49) we have,

$$\begin{aligned} y(k+n) &= \Psi_{n+1}[x(k), U_n(k)] \\ &= \Psi_{n+1}[\psi[Y_n(k), U_{n-1}(k)], U_n(k)] \\ &:= \bar{F}[y(k+n-1), \dots, y(k), u(k+n-1), \dots, u(k)] \end{aligned} \quad (3.52)$$

This is equivalent to

$$y(k+1) = \bar{F}[y(k), \dots, y(k-n+1), u(k), \dots, u(k-n+1)] \quad (3.53)$$

which is an exact representation of the nonlinear plant valid in some neighborhood of the equilibrium point $(0, 0)$. The equation is referred to as the nonlinear auto-regressive moving-average (NARMA) model of the system, see (Narendra, 1996 [50]). It shall be emphasized that the representation exists only if the inputs and outputs of the system are restricted to a neighborhood in which the implicit function theorem can be applied.

In the preceding section a normal state space representation of the nonlinear system was derived which served to highlight the concepts of relative degree and zero dynamics. The following discussion illustrates how the relative degree and the zero dynamics can be retrieved in the NARMA representation of the system. Following the same steps as above, but starting with the normal form (3.45) rather than (3.40) it can be shown that the following input-output representation exists locally in the neighborhood of an equilibrium point:

$$y(k+1) = \bar{\mathcal{F}}[y(k), y(k-1), \dots, y(k-n+1), u(k-d+1), \dots, u(k-n+1)] \quad (3.54)$$

As a consequence of the relative degree $d > 1$, the equation does not depend on the most recent $d-1$ values of the input. The representation (3.54) exists provided that the system has a well defined relative degree. In order to obtain (3.54), the internal states $\zeta_0(k)$ have to be expressed in terms of the inputs and outputs of the system. If the linear model (3.42) is observable, it can be shown (using the implicit function theorem) that there exists an

analytic map $\phi : \mathbb{R}^n \times \mathbb{R}^{(n-d)} \rightarrow \mathbb{R}^{(n-d)}$, valid for some $(Y_n, U_{n-d}) \in \mathcal{Y}_n \times \mathcal{U}_{n-d}$ and $\zeta_0 \in \mathcal{X}_0$, such that

$$\zeta_0(k) = \phi[Y_n(k), U_{n-d}(k)] \quad (3.55)$$

Suppose a sequence of inputs $U_{n-d}^0 = [u^0(k-d+1), \dots, u^0(k-n+1)]$ has been chosen, such that $y(k) \equiv 0$, for all $k > 0$, then the internal states $\zeta_0(k) = \phi[0, U_{n-d}^0(k)]$ evolve as the zero dynamics of the system and the zeroing sequence satisfies

$$\bar{\mathcal{F}}[0, 0, \dots, 0, u^0(k-d+1), \dots, u^0(k-n+1)] = 0 \quad (3.56)$$

This difference equation is referred to as the input-output zero dynamics of the system. As in the linear case, we construct a d -step ahead predictor by expressing $y(k+1), \dots, y(k+d-1)$ successively in terms of Y_n and U_n using the NARMA representation (3.54) for each $y(k+i)$, $i \in \{1, \dots, d-1\}$:

$$y(k+d) = \mathcal{F}[y(k), \dots, y(k-n+1), u(k), \dots, u(k-n+1)] \quad (3.57)$$

This is the nonlinear equivalent of equation (3.20). In the above statement of the exact tracking problem we were looking for conditions under which the control law defined in (3.41) exists. We are now in a position to state these conditions.

If the system (3.40) has a well defined relative degree and asymptotically stable zero dynamics then there exists an analytic function \mathcal{G} of the form

$$u(k) = \mathcal{G}[y(k), \dots, y(k-n+1), u(k-1), \dots, u(k-n+1), y^*(k+d)] \quad (3.58)$$

which solves the exact tracking problem in a neighborhood of the origin. The conditions are seen to be equivalent to those in the linear exact tracking problem. In order to obtain the control law, $y(k+d)$ in equation (3.57) is replaced by its desired value $y^*(k+d)$. The implicit function theorem then assures the existence of the control law \mathcal{G} in equation (3.58). It solves the exact tracking problem provided that the system has a stable inverse. The role of the zero dynamics is to assure that the system state remains in a neighborhood of the equilibrium state in which the implicit function theorem is valid. The closed-loop system can be written as the sum of an unforced system and a perturbation term depending on $y^*(k+d)$ and the internal states of the system. Using the normal form (3.45) it is easy to see that asymptotic stability of the zero dynamics implies asymptotic stability of the autonomous part of the closed-loop system. The fact that the system state remains in some neighborhood then follows from Malkin's theorem on stability under perturbations [9].

The emphasis of this section has been in deriving conditions for the existence of nonlinear control laws which solve the exact tracking problem. This may be regarded as the algebraic part of the adaptive control problem. The analytic part consists in using artificial neural networks to approximate the corresponding maps. This will be discussed in following section.

3.2.3 Practical Neurocontrol

Any discrete-time dynamical system can be represented using linear functions, such as summations, multiplication by constants and delays, as well as nonlinear functions. A massively parallel interconnection of many of these basic functions is commonly referred to as a neural network. Architectures which have found wide application are the radial basis function network (RBFN) and the multi-layer perceptron network with at least one hidden layer (MNN). It has been shown that both types are capable of approximating any continuous function on a compact set to any degree of accuracy (see Cybenko 1989, [14]; Hornik 1989, [30]; Chen 1991, [10]). While other approximators, such as polynomials, splines, trigonometric series, and orthogonal functions share the same property, neural networks represent a very conveniently parameterized class of nonlinear maps. The fact that a relatively small number of parameters is needed to represent even complex functions, makes them attractive for identification and control of nonlinear systems.

The key idea of neurocontrol is to use neural networks in order to approximate the nonlinear maps involved in the solution of the exact tracking problem. This is part of a three step design procedure which is briefly described below.

Step 1: At this stage, we are interested in identification only. A neural network is trained to approximate the input–output behavior of the plant. It has been demonstrated that the dynamical behavior of the plant can be described by means of the static map $\mathcal{F} : \mathcal{Y}_n \times \mathcal{U}_n \rightarrow \mathbb{R}$ defined in equation (3.57). The map depends upon the vector

$$\phi(k) = \begin{bmatrix} Y_{n'}(k) \\ U_{n'}(k) \end{bmatrix} \quad (3.59)$$

introduced in (3.22). In much the same way as in the linear case (cf. equation 3.28) the identification model uses the vector $\phi(k)$ to generate an estimate $\hat{y}(k)$ of the output of the plant except that in this case, a neural network is used to parameterize the solution:

$$\hat{y}(k) = \mathcal{N}_1[Y_{n'}(k-d), U_{n'}(k-d), \hat{\theta}_1(k-1)] \quad (3.60)$$

The network parameters $\hat{\theta}_1$ are trained using the identification error $\epsilon(k) = \mathcal{N}_1 - y(k)$. The adjustment is carried out along the negative gradient of the cost function $J(\hat{\theta}_1) = \epsilon(\hat{\theta}_1)^2$,

$$\hat{\theta}_1(k) = \hat{\theta}_1(k-1) - 2\eta \frac{\partial \epsilon(k)}{\partial \hat{\theta}_1(k-1)} \epsilon(k) \quad (3.61)$$

with a sufficiently small step size η , i.e. $0 < \eta \ll 1$. Details regarding training methods based on gradients can be found in (Widrow 1990, [75]; Narendra 1991, [60]; Werbos 1990, [76]).




Figure 3.2: Neurocontrol

Step 2: This step consists in determining the map \mathcal{G} defined in equation (3.58). It can be initiated once the identification error $\epsilon(\cdot)$ in step 1 is small. A second network \mathcal{N}_2 is used providing an input $u(k)$ to the first network \mathcal{N}_1 . Figure (3.2) displays the resulting closed-loop system. \mathcal{N}_2 is of the form

$$u(k) = \mathcal{N}_2[y(k), y(k-1), \dots, y(k-n+1), u(k-1), \dots, u(k-n+1), y^*(k+d)] \quad (3.62)$$

Its parameters $\hat{\theta}_2$ are trained such as to minimize the control error $e(k) = y(k) - y^*(k)$. It is clear that once the error is zero, $\hat{\theta}_2$ is constant and \mathcal{N}_2 represents the control law for exact tracking.

Step 3: The reason for working with \mathcal{N}_1 instead of the plant during the training phase in step 2 is that it is not clear whether \mathcal{N}_2 stabilizes the system even as its parameters are being adjusted. A proof of stability paralleling the one presented for linear systems in section 3.1.4 is missing to date. Once the control error $e(\cdot)$ is small, one has sufficient confidence to apply the control law represented by \mathcal{N}_2 to the actual system. Step 3 then consists in replacing \mathcal{N}_1 by the plant. Step 2 is omitted whenever the control input can be computed algebraically from (3.60). In (Narendra 1997 [61]) it has been suggested that a general NARMA representation could be approximated by a model in which the control input $u(k)$ appears linearly. The use of such models simplifies the design of neurocontrollers substantially.

Example 3.3 The flux linkage ψ_E generated in the stator of a separately excited DC machine is a nonlinear function of the exciting current i_E . The nonlinearity is due to the saturation of the magnets at high currents, which is often approximated by a sigmoid static magnetizing curve $h(i_E)$, while hysteresis effects are neglected. In discrete-time, the follow-

ing state space model is obtained

$$\begin{aligned}\psi_E(k+1) &= \psi_E(k) + T_S u_E(k) - T_S R_E i_E(k) \\ i_E(k) &= h^{-1}[\psi_E(k)]\end{aligned}\tag{3.63}$$

R_E is the resistance of the field circuit while $T_S = 1\text{ ms}$ denotes the sampling time. R_E is assumed to be known but the magnetizing curve $h[\cdot]$ is unknown. By reducing the flux, a DC motor can be operated above base speed. Since ψ_E is not directly measured, the objective is to control the exciting current i_E which is the output variable in (3.63). The voltage $u_E(k)$ is generated by a converter which –given the above sampling time– can be approximated by a unit delay: $u_E(k+1) = u(k)$. The standard way for controlling i_E is to linearize (3.63) at a given operating point $(\bar{u}_E, \bar{\psi}_E)$ and to design a PI–control loop which is usually tuned to the “absolute value optimum” (see Schröder 2001, [71]). This may lead to the following problem. At low currents, the differential inductance $\frac{d\Psi_E}{di_E}$ is large, so a large controller gain k_p is needed to cause the output i_E to respond fast and accurately to e.g. a step input i_E^* . As i_E increases, the flux saturates, i.e. the inductance $\frac{d\Psi_E}{di_E}$ decreases. If the gain of the controller is kept at a high level the closed loop system may become unstable. Clearly, k_p has to be adjusted. There are many technical realizations of this adjustment which can be found in (Schröder, 2001 [71]). From a theoretical point of view, we are looking for a nonlinear map $u = \mathcal{G}[i_E^*, i_E, \cdot]$ with small gain at high currents and large gain at small currents. As seen in the following simulation, a neural network can be used to implement such a map, even if the exact shape of the saturation function $h[i_E]$ is unknown. For the purpose of illustration, we assume that i_E^* is a square wave and that the onset of the saturation occurs at low currents ($< 1/10$ of the nominal current i_{E0}). From a practical point of view, much weaker assumptions could be made (in particular, a DC drive behaves almost linearly up to $1/2 i_{E0}$) but the settings allow us to study the effect of the nonlinearity in a general field circuit.

An input–output representation of the system is obtained directly from (3.63)

$$i_E(k) = h^{-1}[b_0 u(k-2) + h[i_E(k-1)] - a_1 i_E(k-1)]\tag{3.64}$$

which is of the form (3.54). In order to derive the $d = 2$ step ahead predictor, (3.64) is successively used to express $i_E(k+1)$ and $i_E(k+2)$ in terms of past values, following the procedure described in section 3.2.2. This leads to the map

$$i_E(k+2) = \mathcal{F}[i_E(k), u(k-1), u(k)]\tag{3.65}$$

Note that no Jacobians are involved in checking observability and, hence, existence of the map \mathcal{F} . Furthermore, the system has no zero dynamics. This assures the existence of a

(global) control law

$$u(k) = \mathcal{G}[i_E^*(k+2), i_E(k), u(k-1)] \quad (3.66)$$

which is approximated by the neural net. In figure (3.3), the PI-controller is tuned moderately in order to prevent instability. This results in poor tracking performance. If a neurocontroller is used instead, the output of the system i_E follows the reference current i_E^* much more closely, even while keeping the amplitude of the input u_E within the bounds 500...600V prescribed by the power converter. \square

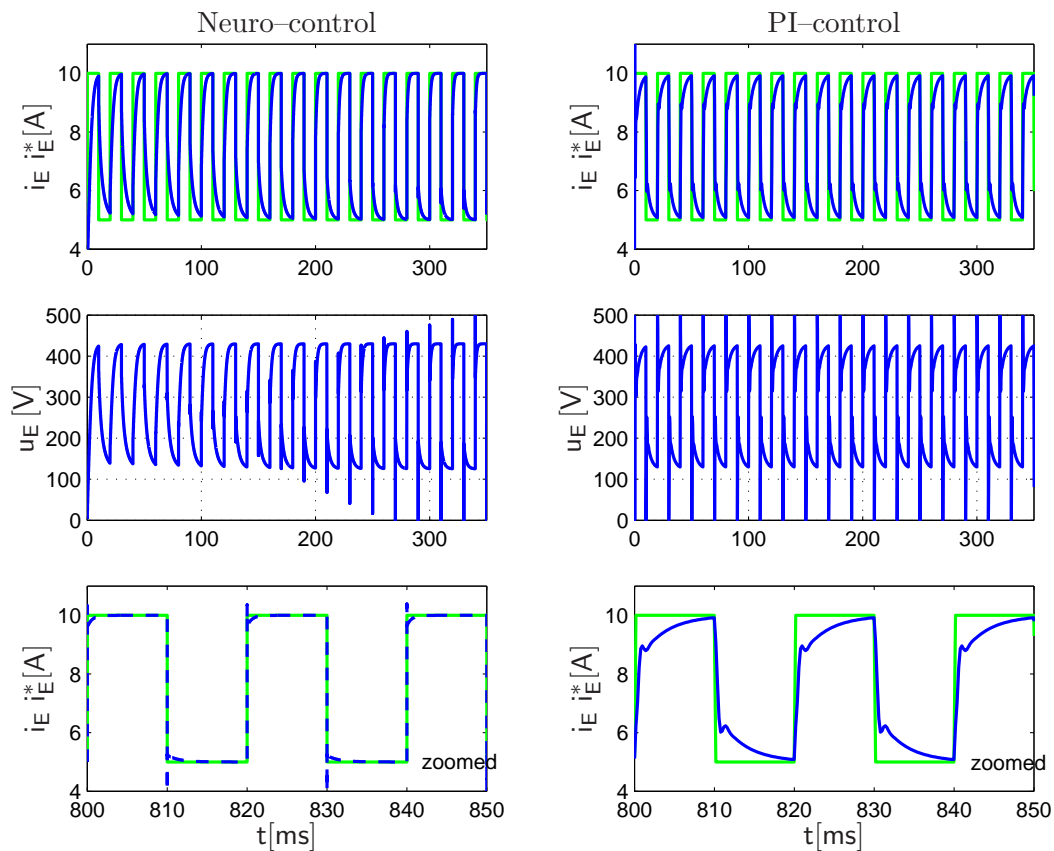


Figure 3.3: Field current control of a separately excited DC drive; left: neurocontrol, right: PI-control

The adaptive control methodology involving neural networks offers a unified solution for a very general class of unknown, nonlinear systems. However, due to the local nature of the implicit function theorem, the results are restricted to a neighborhood of the equilibrium state. It has been shown (Aeyels, 1981 [7]), that by increasing the number of input-output measurements up to $(2n + 1)$, the system becomes globally observable. The implications of this result for both identification and control need further investigation. Even when a

solution is known to exist (in some neighborhood of the equilibrium state), it is generally impossible to demonstrate stability of the adaptive law. In most cases, gradient methods are used which do not have proven stability properties. A class of systems for which stability of the adaptive closed loop system can be established is given by

$$y(k+1) = \sum_{i=0}^l a_i f_i[y(k), \dots, y(k-n+1)] + \sum_{i=0}^{n-1} b_i u(k-i) \quad (3.67)$$

where the functions $f_i(\cdot)$, $i = 0 \dots l$ are assumed to be globally Lipschitz and known. In this case, the regression vector assumes the form

$$\phi(k) = [f_1[\cdot], \dots, f_l[\cdot], u(k), \dots, u(k-n+1)] \quad (3.68)$$

and the procedure described in (3.1) for linear systems can be applied.

Chapter 4

Disturbance Rejection

The previous chapter dealt with the problem of controlling a system with unknown, constant parameters. In this chapter we consider the case where, in addition, the system is subject to unknown, external disturbances. It is seen that under certain conditions, the problem can be translated to the previous one if a controller of extended order is used. The approach is based on the well-known fact from linear control theory that deterministic disturbances containing a finite number of frequencies can be completely rejected by placing appropriate poles in the feed-forward path of the control loop. According to the internal model principle (Francis and Wonham 1976 [24]), this can be regarded as a procedure of expanding the system by a “disturbance model” which generates an additional input that compensates the effect of the disturbance.

The main purpose of this chapter is to discuss how this linear design rule extends to the case where the parameters of the plant and the disturbance model are unknown. It is seen that both deterministic and stochastic disturbances can be rejected using the same method of augmenting the state space of the closed-loop system. In the stochastic case, the performance improvement is seen to depend upon the impulse response of the disturbance model. If the latter decays fast, the disturbance affecting the system is almost pure white noise which cannot be eliminated. Pseudo-linear regression algorithms can be employed in this case to minimize the variance of the resulting output error. If, on the other hand, the impulse response is slowly decaying, the disturbance is highly correlated with its past values and can be rejected almost completely, as in the deterministic case. In the nonlinear domain, a similar procedure can be adopted provided that certain properties of the nonlinear system hold. These properties, defined in terms of the linearized system, guarantee the existence of an input-output representation of the augmented system. The main theoretical questions that arise in this context are similar to the ones discussed in chapter 3. It is

seen that nonlinear disturbance rejection is possible if the state of the system remains in a neighborhood of the origin. In summary, whenever an external disturbance can be described as the output of an unforced system of known order, it can be eliminated completely. The importance of this result lies in the fact that the proposed method not only guarantees stability under perturbations but also compensates for the perturbation, even as the latter is not known completely.

4.1 Linear Disturbance Rejection

4.1.1 Deterministic Disturbances

Let the system (3.1) be affected by an external, bounded, deterministic disturbance $v(k)$,

$$\begin{aligned}\Sigma : \quad x(k+1) &= Ax(k) + bu(k) + b_v v(k) \\ y(k) &= c^T x(k)\end{aligned}\tag{4.1}$$

where $x(k) \in \mathbb{R}^n$ is the state of the system, $A \in \mathbb{R}^{n \times n}$ the system matrix, $b \in \mathbb{R}^{n \times 1}$ and $b_v \in \mathbb{R}^{n \times 1}$ are input vectors and $y(k)$ a scalar output, i.e. $c \in \mathbb{R}^{n \times 1}$. The signal $v(k)$ is assumed to be the output of the following unforced system

$$\begin{aligned}\Sigma_v : \quad x_v(k+1) &= A_v x_v(k) \\ v(k) &= c_v^T x_v(k)\end{aligned}\tag{4.2}$$

where $x_v(k) \in \mathbb{R}^{n_v}$, $A_v \in \mathbb{R}^{n_v \times n_v}$ and $c_v \in \mathbb{R}^{n_v \times 1}$. The equation may be regarded as a model which generates the external disturbance. If A_v is a Schur matrix (with eigenvalues inside the unit circle), $x_v(k)$ and $v(k)$ tend to zero asymptotically. If, on the other hand, A_v is unstable, $x_v(k)$ will grow in an unbounded fashion and the control needed to reject the disturbance will also be unbounded. Since our interest is in bounded disturbances, we assume that A_v is a stable matrix with simple eigenvalues on the unit circle. Hence $v(k)$ can be expressed as a finite sum of sinusoidal signals (including a constant). Our objective in this case is to determine $u(k)$ in the composite system $\Sigma \circ \Sigma_v$ such that $\lim_{k \rightarrow \infty} |y(k) - y^*(k)| = 0$ where $y^*(k)$ is the desired output and $y^*(k+d) = r(k)$ is known at instant of time k .

The solution involves determining the input-output representation of both Σ and Σ_v . The key step then is to use the homogeneous difference equation obtained from (4.2) in order to eliminate $v(k)$ from (4.1). This is best illustrated by the following example where $n = n_v = 1$.

Example 4.1

$$y(k+1) = ay(k) + v(k) + u(k)\tag{4.3}$$

$$v(k) = a_v v(k-1) \quad a_v = \pm 1\tag{4.4}$$




Figure 4.1: Disturbance as output of homogeneous difference equation

If $v(k)$ is known, the input $u(k)$ can be computed from equation (4.3) to achieve exact tracking. From equations (4.3) and (4.4) we have $v(k) = a_v v(k-1) = a_v [y(k) - ay(k-1) - u(k-1)]$, i.e. the value of $v(k)$ depends only on past values of $y(\cdot)$ and $u(\cdot)$. By choosing

$$u(k) = y^*(k+1) - (a + a_v)y(k) + a_v a y(k-1) + a_v u(k-1) \quad (4.5)$$

exact tracking is achieved in two steps. Note that $v(k)$ is not known explicitly but is completely defined by its value at time $(k-1)$, which, in turn can be expressed in terms of past values of both the input and the output. \square

The idea is readily generalized to the case where the system as well as the disturbance model are of higher order. The input–output representations obtained in this case are:

$$\Sigma : y(k+d) = \sum_{i=0}^{n-1} a_i y(k-i) + \sum_{i=0}^{n-1} b_i u(k-i) + \sum_{i=0}^{n-1} c_i v(k-i) \quad (4.6)$$

$$\Sigma_v : v(k+1) = \sum_{i=0}^{n_v-1} \alpha_i v(k-i) \quad (4.7)$$

The steps to eliminate the disturbance from Σ by means of the autonomous system Σ_v are as follows: If $n > n_v$, (4.7) is used to express $v(k-n+1), \dots, v(k-n_v)$ in (4.6) in terms of $v(k) \dots v(k-n_v+1)$, otherwise this step is skipped. Then, (4.6) is solved for $v(k-n_v+1)$ and the result used to express the r.h.s. of (4.7) in the more recent values of $v(\cdot)$ and the inputs and outputs of the system. By shifting the time axis (4.7) backwards by one step we again obtain an expression for $v(k-n_v+1)$ which is used to eliminate $v(k-n_v+1)$ from (4.6). The procedure is repeated n_v times until the r.h.s of the second equation is expressed completely in terms of the inputs and outputs of the plant. In the final stage, the first equation (4.6) contains only the current value $v(k)$ which, in turn is eliminated using (4.7). The resulting equation represents the disturbance–free composite system. With n_v being the order of the disturbance generating system, n_v steps were needed to eliminate $v(\cdot)$. Consequently, the

dimension of the system has increased by n_v :

$$\bar{\Sigma} : \quad y(k+d) = \sum_{i=0}^{n+n_v-1} a_i y(k-i) + \sum_{i=0}^{n+n_v-1} b_i u(k-i) \quad (4.8)$$

Clearly, the equation is of the form (3.20) and the same control law as in case without external disturbances can be applied, except that it now depends on the past $(n+n_v-1)$ values of the signals of the system:

$$u(k) = \frac{1}{b_0} \left[y^*(k+d) - \sum_{i=0}^{n+n_v-1} a_i y(k-i) - \sum_{i=1}^{n+n_v-1} b_i u(k-i) \right] \quad (4.9)$$

The procedure is more transparent if we use an equivalent representation of the system (4.1) given by

$$A(q^{-1})y(k) = q^{-d}B(q^{-1})u(k) + q^{-d}G(q^{-1})v(k) \quad (4.10)$$

where

$$\begin{aligned} A(q^{-1}) &= 1 - a_1 q^{-1} - \dots - a_{n_A} q^{-n_A} \\ B(q^{-1}) &= b_0 + \dots + b_{n_B} q^{-n_B} \\ G(q^{-1}) &= g_0 + \dots + g_{n_G} q^{-n_G} \end{aligned}$$

are polynomials in the delay operator q^{-1} . The relative degree d is assumed to be known and $B(q^{-1})$ is a Hurwitz polynomial. The disturbance $v(k)$ is generated by homogenous system of the form

$$\begin{aligned} v(k) &= [d_1 q^{-1} + \dots + d_{n_D} q^{-n_D}]v(k) \\ &:= [1 - D(q^{-1})]v(k) \end{aligned} \quad (4.11)$$

If we solve equation (4.10) for $v(k)$ and use the result to substitute $v(k)$ in (4.11) we obtain

$$DA y(k) = q^{-d}DB u(k) \quad (4.12)$$

which corresponds to a system model the order of which has been augmented in order to account for the presence of the disturbance. For ease of notation, the explicit dependence of the polynomials on q^{-1} has been omitted. Replacing $y(k)$ in equation (4.12) by $y^*(k)$ and solving for $u(k)$ yields the equivalent of control law (4.9).

$$u(k-d) = \frac{1}{b_0} \left[y^*(k) - [1 - DA]y(k) - [DB - b_0]u(k-d) \right] \quad (4.13)$$

The application of the control law results in the closed-loop system:

$$y(k) = y^*(k) + q^{-d}GD v(k) \quad (4.14)$$

The effect of using a controller of augmented order is immediately evident, since, according to (4.11), $Dv(k) = 0$. If the parameters of the system are unknown, the control law is

determined on the basis of an underlying identification model. This model is of augmented order:

$$\hat{y}(k) = \phi(k-d)^T \hat{\theta}(k-1) \quad (4.15)$$

where $\phi(\cdot) \in \mathbb{R}^{n+n_v}$ and $\hat{\theta}(\cdot) \in \mathbb{R}^{n+n_v}$. Since (4.8) has been shown to be a valid representation of the composite system $\Sigma \circ \Sigma_v$, we conclude that there exists a constant parameter vector θ^* for which $y(k) = y^*(k)$. It is seen that the zeros of the augmented system are given by the original zeros plus the poles of the disturbance model. Since the latter is stable, the zeros of the polynomial DB in equation (4.12) lie inside or on the closed unit circle and the zeros of the transfer function $q^{-d}B/A$ obtained from (4.12) after pole-zero cancellation, i.e. its controllable modes, lie strictly inside the unit circle. In view of the comment made in chapter 3.1.4 the same arguments as in the disturbance-free case can be used to proof stability of the adaptive controller based on the augmented system.

As seen in equation (4.10), the relative degree d is the same as in the original system. The orders n and n_v must be known. In most practical situations, prior information about the plant as well as the disturbance is available. As an example, if the disturbances are harmonic, n_v is equal to the (expected) number of nonzero points in the two-sided spectrum of the disturbance signal. Using the parameter estimates $\hat{\theta}(\cdot) = [\hat{a}_0, \dots, \hat{a}_{n+n_v-1}, \hat{b}_0, \dots, \hat{b}_{n+n_v-1}]^T$ obtained from (4.15) the control law reads:

$$u(k) = \frac{1}{\hat{b}_0(k)} \left[y^*(k+d) - \sum_{i=0}^{n+n_v-1} \hat{a}_i(k) y(k-i) - \sum_{i=1}^{n+n_v-1} \hat{b}_i(k) u(k-i) \right] \quad (4.16)$$

Since the controller is based on the augmented system $\bar{\Sigma}$, exact cancellation of the disturbance is achieved. It is clear that the method asymptotically introduces poles in the feed-forward path of the control-loop which correspond to the poles of the disturbance model. As an example, the effect of a sinusoidal disturbance with frequency ω can be nulled by introducing a pair of complex conjugate poles at $z = e^{\pm i\omega T_s}$. The resulting disturbance transfer function has a notch at ω . These facts are well-known from linear systems theory. It is also known that the linear method fails if the frequency of the disturbance is not known exactly. The benefit of the adaptive version is that the frequency need not be known but the parameters are tuned automatically to generate a notch at the appropriate frequency such that, in any case, $\lim_{k \rightarrow \infty} |y(k) - y^*(k)| = 0$. In addition, slow drifts of the disturbance frequency can be tracked.

Example 4.2 Assume that the two-mass system of example (3.2) is subject to a sinusoidal disturbance of unknown frequency ω . The unforced system generating $v(\cdot)$ is given by

$$\Sigma_v : v(k+1) = 2 \cos(\omega h) v(k) - v(k-1) \quad (4.17)$$

This is already more than we need to know. In fact, all that is needed to design an adaptive controller eliminating $v(\cdot)$ is the order $n_v = 2$ of the disturbance generating system. Figure (4.2) displays the performance of the adaptive system, as the output is required to track a piecewise constant signal. The frequency of the disturbance is constant in the first 2 plots and (slowly) time-varying in the last row of figure 4.2.

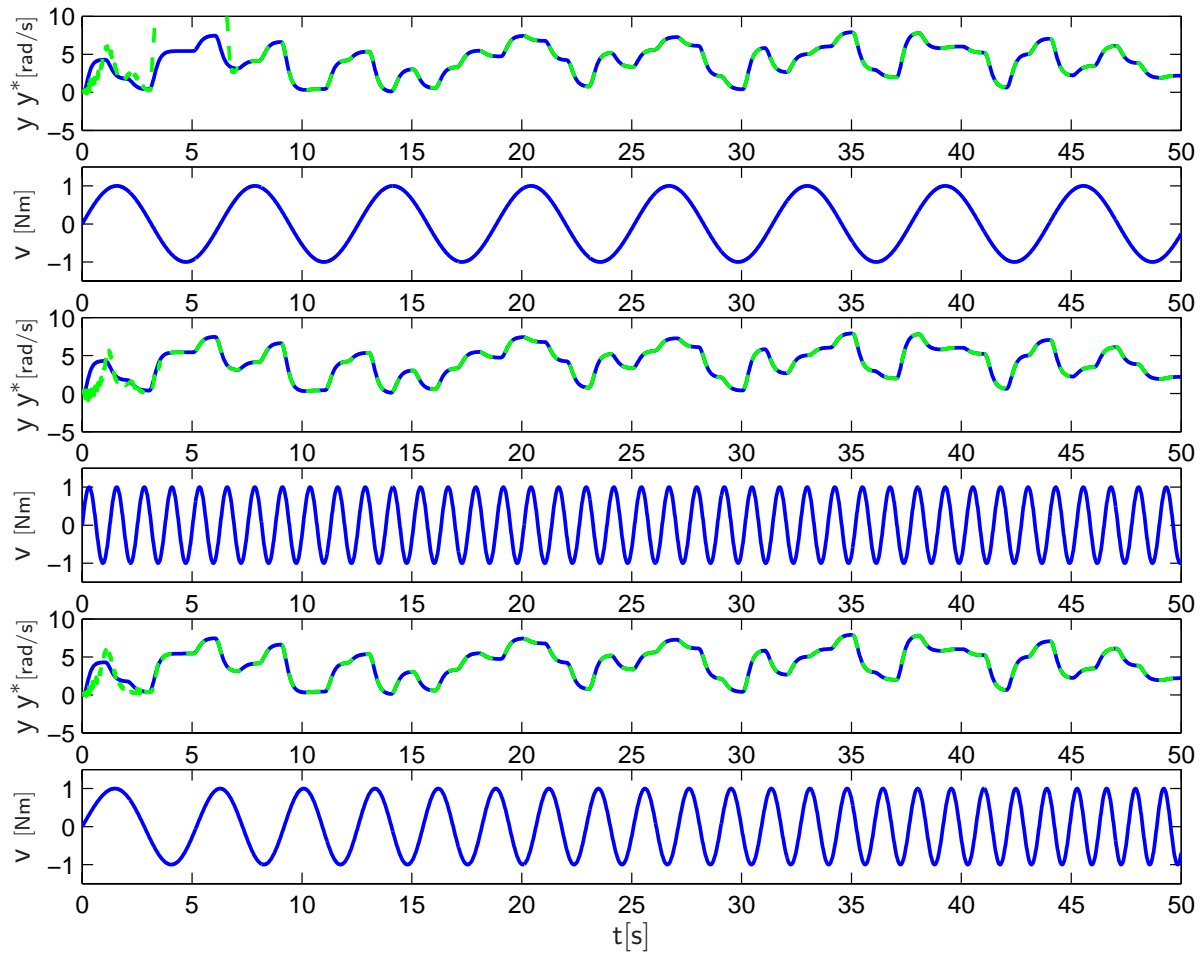


Figure 4.2: Adaptive rejection of disturbances with unknown frequency $\omega = 1 \dots 5$ rad/s

4.1.2 Stochastic Disturbances

In the following, we will assume that the linear plant described in equation (4.1) is affected at the input by correlated noise $v(k)$. As in the previous section, $v(k)$ is the output of a disturbance generating system. In contrast to equation (4.2), Σ_v is driven by a white noise

sequence $\{w(k)\}$.

$$\begin{aligned}\Sigma_v : \quad x_v(k+1) &= A_v x_v(k) + b_w w(k) \\ v(k) &= c_v^T x_v(k)\end{aligned}\tag{4.18}$$

where $b_w \in \mathbb{R}^{n_v \times 1}$ and $w(k)$ has the following properties:

$$E\{w(k)|k-1\} = 0 \quad \text{a.s.}\tag{4.19}$$

$$E\{w(k)^2|k-1\} = \sigma^2 \quad \sigma^2 < \infty \quad \text{a.s.}\tag{4.20}$$

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N w^2(k) < \infty \quad \text{a.s.}\tag{4.21}$$

In other words, the white noise has zero conditional mean, finite variance σ^2 and is mean square bounded. As in the previous section, we assume that Σ and Σ_v have an input-output representation. The problem is to design a control law $u(k)$ such that $y(k)$ follows a desired reference signal $y^*(k)$ as closely as possible. Due to the presence of the noise the control error $e(k) = [y(k) - y^*(k)]$ cannot be made zero, but its expected value can be minimized.

The effect of $w(k)$ on $v(k)$ obviously depends upon the parameters of Σ_v , i.e. the matrix A_v and the vectors b_w and c_v . We shall successively consider the control strategies when all the parameters are known, and when the parameters are unknown and have to be estimated on-line. In the former case we have a linear stochastic control problem, and in the latter case we have a linear stochastic adaptive control problem. Further, before proceeding to solve the two problems analytically, we shall discuss qualitatively the nature of the disturbance, and the conditions under which significant improvement in performance can be expected.

A Qualitative Analysis

Proceeding as in section 4.1.1, it is possible to obtain an augmented input-output model $\bar{\Sigma}$ of order $n + n_v$:

$$y(k+1) = \sum_{i=0}^{n+n_v-1} \bar{a}_i y(k-i) + \sum_{i=d-1}^{n+n_v-1} \bar{b}_i u(k-i) + \sum_{i=d-1}^{n+n_v-1} \bar{c}_i w(k-i)\tag{4.22}$$

In order to obtain the d -step ahead predictor form, the time axis is shifted and the outputs $y(k+d-1) \dots y(k+1)$ are expressed in terms of past values, see chapter 3.1.2.

$$\bar{\Sigma} : \quad y(k+d) = \sum_{i=0}^{n+n_v-1} a_i y(k-i) + \sum_{i=0}^{n+n_v-1} b_i u(k-i) + \sum_{i=0}^{n+n_v-1} c_i w(k-i)\tag{4.23}$$

From equation (4.23), it is clear that all past values of the white noise (i.e. $w(k-i)$, $i = 0, \dots, n + n_v - 1$) affect the output at time $k+d$. The question that has to be addressed




Figure 4.3: Correlated noise affecting a linear system

is the extent to which the control input at instant k can compensate for these values of the noise. This is discussed by means of the following simple examples:

Order of the disturbance model $n_v = 0$

First, let the order of the disturbance generating system be $n_v = 0$:

$$\bar{\Sigma} : \quad y(k+1) = ay(k) + u(k) + w(k) \quad (4.24)$$

Note that $w(k)$ is an unknown random variable. Since $w(k)$ is not known at instant of time k , no control input $u(k)$ can be defined that will compensate for $w(k)$. However, solving the difference equation we have

$$y(k) = a^k y_0 + \sum_{i=0}^{k-1} a^{k-1-i} [u(i) + w(i)] \quad (4.25)$$

with the initial value $y(0) = y_0$. We see that $y(k)$ is affected by all previous values $w(i)$, $i = 0 \dots (k-1)$ of the white noise, and therefore contains information about the disturbance. By choosing the control input $u(k) = r(k) - f_b y(k)$ a tracking controller is realized that cancels all past disturbance values but not the present one. f_b is the feedback gain. The control error becomes:

$$e(k+1) = y(k+1) - r(k) = (a - f_b) y(k) + w(k) \quad (4.26)$$

The variance of the control error, $E\{e^2(k+1)\} = (a - f_b)^2 E\{y^2(k)\} + \sigma^2$, has a minimal value for $f_b = a$. Hence, the best choice for the control law in the case $n_v = 0$ is obtained by simply ignoring the presence of the white noise $w(k)$ in $\bar{\Sigma}$.

Order of the disturbance model $n_v > 0$

Proceeding to a more interesting case, we now assume that the order of the disturbance generating system Σ_v is greater than zero. The signal $v(k)$ is correlated with $v(k-1)$, i.e.

$E\{v(k)v(k-1)\} \neq 0$. Once again, considering simple models for both Σ and Σ_v , we have for example,

$$\begin{aligned} \Sigma : \quad y(k+1) &= ay(k) + u(k) + v(k) \\ \Sigma_v : \quad v(k+1) &= a_v v(k) + b_w w(k) \end{aligned} \quad (4.27)$$

The transfer function of the disturbance model Σ_v is $F_v(z) = \frac{v(z)}{w(z)} = \frac{b_w}{z-a_v}$, with $|a_v| < 1$. It is clear that Σ_v may be thought of as an IIR-filter with white-noise input. For different values of the eigenvalue of Σ_v the nature of the output and the maximum values it assumes are different. For purposes of comparison, we would require the range of values assumed by the disturbance to be the same in all cases considered. We assume that the high frequency gain $b_w = b_w(a_v)$ can be determined experimentally for this purpose. Three values of the pair $[b_w, a_v]$ and the corresponding evolution of $v(k)$ for a specified white noise sequence $\{w(k)\}$ are displayed in figure (4.4).




Figure 4.4: Correlated noise $v(k)$ for different values of the pair $[b_w, a_v]$ in equation (4.27): (a) $[1, 0]$, (b) $[0.35, -0.8]$, (c) $[0.09, 0.99]$

The impulse response of the filter Σ_v in equation (4.27) is given by

$$h(k) = b_w a_v^{k-1} \quad k \geq 1 \quad (4.28)$$

Assuming that $v(0) = 0$, the output $v(k)$ can be determined as

$$v(k) = \sum_{i=0}^{k-1} h(k-i) w(i) = b_w \sum_{i=0}^{k-1} a_v^{k-1-i} w(i) \quad (4.29)$$

Hence, $|v(k)| = |b_w| \left| \sum_{i=0}^{k-1} a_v^{k-1-i} w(i) \right|$. If the eigenvalue a_v is close to ± 1 , the impulse response $h(k)$ decays slowly and the contribution of the last sum to the magnitude of $v(k)$ is large. Consequently, $|b_w|$ has to be reduced in order to keep $v(k)$ within the same bounds for all eigenvalues $a_v \in (-1 \dots 1)$. If $|a_v|$ increases the correlation of $v(k)$ with its past values also increases:

$$\begin{aligned} E\{v(k)v(k-1)\} &= a_v E\{v^2(k-1)\} + b_w E\{w(k-1)\} \cdot E\{v(k-1)\} \\ &= a_v E\{v^2(k-1)\} \quad \text{because of (4.19)} \end{aligned} \quad (4.30)$$

At the same time, $|b_w|$ is reduced. Hence, the effect of the noise in the determination of disturbance signal $v(k)$ becomes negligible. In summary we have,

- $|a_v| \approx 0$: $v(k)$ is (delayed) pure white noise
- $|a_v| < 1$: $v(k)$ is colored noise
- $|a_v| \approx 1$: $v(k)$ depends mainly on its past values

If $|a_v| \approx 0$ and $|b_w| = 1$, the disturbance affecting the system is delayed pure white noise. The response cannot be improved by a controller of augmented order. On the other hand, if $|a_v| \approx 1$, the gain $|b_w|$ is small and the resulting output highly correlated with its past values. From equation (4.29) we know that $v(k)$ is affected by all previous values of the white noise, i.e. $E\{v(k)w(k-i)\} \neq 0, i = 1 \dots k$. Thus, the effect of past values of $w(\cdot)$ can be indirectly observed through the autoregressive part of the disturbance model Σ_v . Since $|b_w| \approx 0$, this part dominates the effect of the current white noise input. In this case, a controller based on the augmented system rejects disturbances almost completely. In the example, the elimination of $v(k)$ yields:

$$\bar{\Sigma} : \quad y(k+1) = (a_v + a)y(k) - a a_v y(k-1) - a_v u(k-1) + u(k) + b_w w(k-1) \quad (4.31)$$

Since $|b_w| \approx 0$, the term due to the white noise $w(k-1)$ is neglected in the control law:

$$\Gamma : \quad u(k) = y^*(k+1) - (a_v + a)y(k) + a a_v y(k-1) + a_v u(k-1) \quad (4.32)$$

If $v(k)$ is highly correlated with its past values the impact of the sum of the past values of the noise is large with respect to present ones. In such a case a substantial improvement of the performance is obtained, if the order of the controller is augmented. The same qualitative behavior is observed if Σ_v is a general n_v - order system as defined in (4.18). Its impulse response is given by

$$h(k) = c_v^T A_v^{k-1} b_w \quad k \geq 1$$

If the eigenvalues of A_v are close to the unit circle, the impulse response decays slowly and the contribution of past values of the noise is large. At the same time, the elements of the vectors c_v and b_w have to be made small such that the maximum amplitude of the disturbance is the same for all stable eigenvalues of A_v . This is best illustrated by considering a disturbance generating system of order $n_v = 2$. Supposing that Σ_v has a pair of complex conjugate eigenvalues which are close to the unit circle, the output $v(k)$ is a damped sinusoid, which is the natural response of the homogeneous part of the filter, while the contribution of $w(k)$ to the output $v(k)$ is small. Therefore, the disturbance can be treated as the output of a homogeneous system and expansion of the state space can be used to reject them.

Example 4.3 Let the system Σ and the disturbance model Σ_v be two second-order systems:

$$\begin{aligned} \Sigma : \quad x(k+1) &= \begin{bmatrix} 0 & 0.2 \\ 1 & 0.1 \end{bmatrix} x(k) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(k) + \begin{bmatrix} 0 \\ 0.5 \end{bmatrix} v(k) \\ y(k) &= [1 \quad 2] x(k) \end{aligned} \quad (4.33)$$

and

$$\begin{aligned} \Sigma_v : \quad x_v(k+1) &= \begin{bmatrix} \alpha & -\beta \\ \beta & \alpha \end{bmatrix} x_v(k) + \begin{bmatrix} b_{v1} \\ b_{v2} \end{bmatrix} w(k) \\ v(k) &= [c_{v1} \quad c_{v2}] x_v(k) \end{aligned} \quad (4.34)$$

The eigenvalues of Σ and Σ_v are $z_1 = 0.5$, $z_2 = -0.4$ and $z_{1v,2v} = \alpha \pm i\beta$ respectively. In the simulation, z_{v1}, z_{v2} are placed arbitrarily within the unit circle corresponding to different degrees of correlation of $v(k)$ with its past values. The objective is to investigate if a control law based only on the autoregressive part of Σ_v would yield acceptable results. Since the gain $|v|/|w|$ also depends on the location of the eigenvalues, the correction factors c_{v1} and c_{v2} were chosen such that the peak-to-peak variation of $v(k)$ remained the same throughout the experiment in order to obtain comparable results. Since both systems Σ and Σ_v are observable the corresponding ARMA representations exist. After the elimination of $v(k)$ we obtain an extended 4th-order system $\bar{\Sigma}$ of the form:

$$\begin{aligned} \bar{\Sigma} : \quad y(k+1) &= a_0 y(k) + a_1 y(k-1) + a_2 y(k-2) + a_3 y(k-3) + b_0 u(k) + \\ &+ b_1 u(k-1) + b_2 u(k-2) + b_3 u(k-3) + c_2 w(k-2) + c_3 w(k-3) \end{aligned} \quad (4.35)$$

where a_i, b_i, c_i are constants depending only on the location of the eigenvalues $z_{1v,2v}$ of the disturbance model. With $y^*(k+1) = 0.5 \sin\left(\frac{2\pi k}{10}\right) + 0.5 \sin\left(\frac{2\pi k}{20}\right)$ being the reference signal, the control law for asymptotic tracking of the reference output is given by:

$$\begin{aligned} u(k) &= \frac{1}{b_0} [y^*(k+1) - a_0 y(k) - a_1 y(k-1) - a_2 y(k-2) - a_3 y(k-3) - \\ &- b_1 u(k-1) - b_2 u(k-2) - b_3 u(k-3)] \end{aligned} \quad (4.36)$$

Again, the contribution of the white noise to the current output is neglected. The simulation results are displayed in figure (4.5). In the first case, the eigenvalues of Σ_v are $z_{1v} = z_{2v} = 0$ and hence the disturbance $v(k)$ affecting Σ is pure white noise $v(k) = z^{-2}w(k)$. Since there is no correlation of the noise $v(k)$ with its past values, the performance of the controller is poor. A similar result is obtained in the second case, where z_{1v}, z_{2v} are close to zero. In the subsequent cases, the eigenvalues are placed closer to the unit circle corresponding to a large time constant of the impulse response of the filter Σ_v . Since previous values of the output $v(k)$ are now the dominant part of the disturbance, almost complete rejection is achieved using a controller (4.36) based on the augmented system. This becomes particularly evident in the last case, where a pair of complex conjugate eigenvalues with $|z_{v1,2}| \approx 1$ is considered and $v(k)$ is almost harmonic. \square

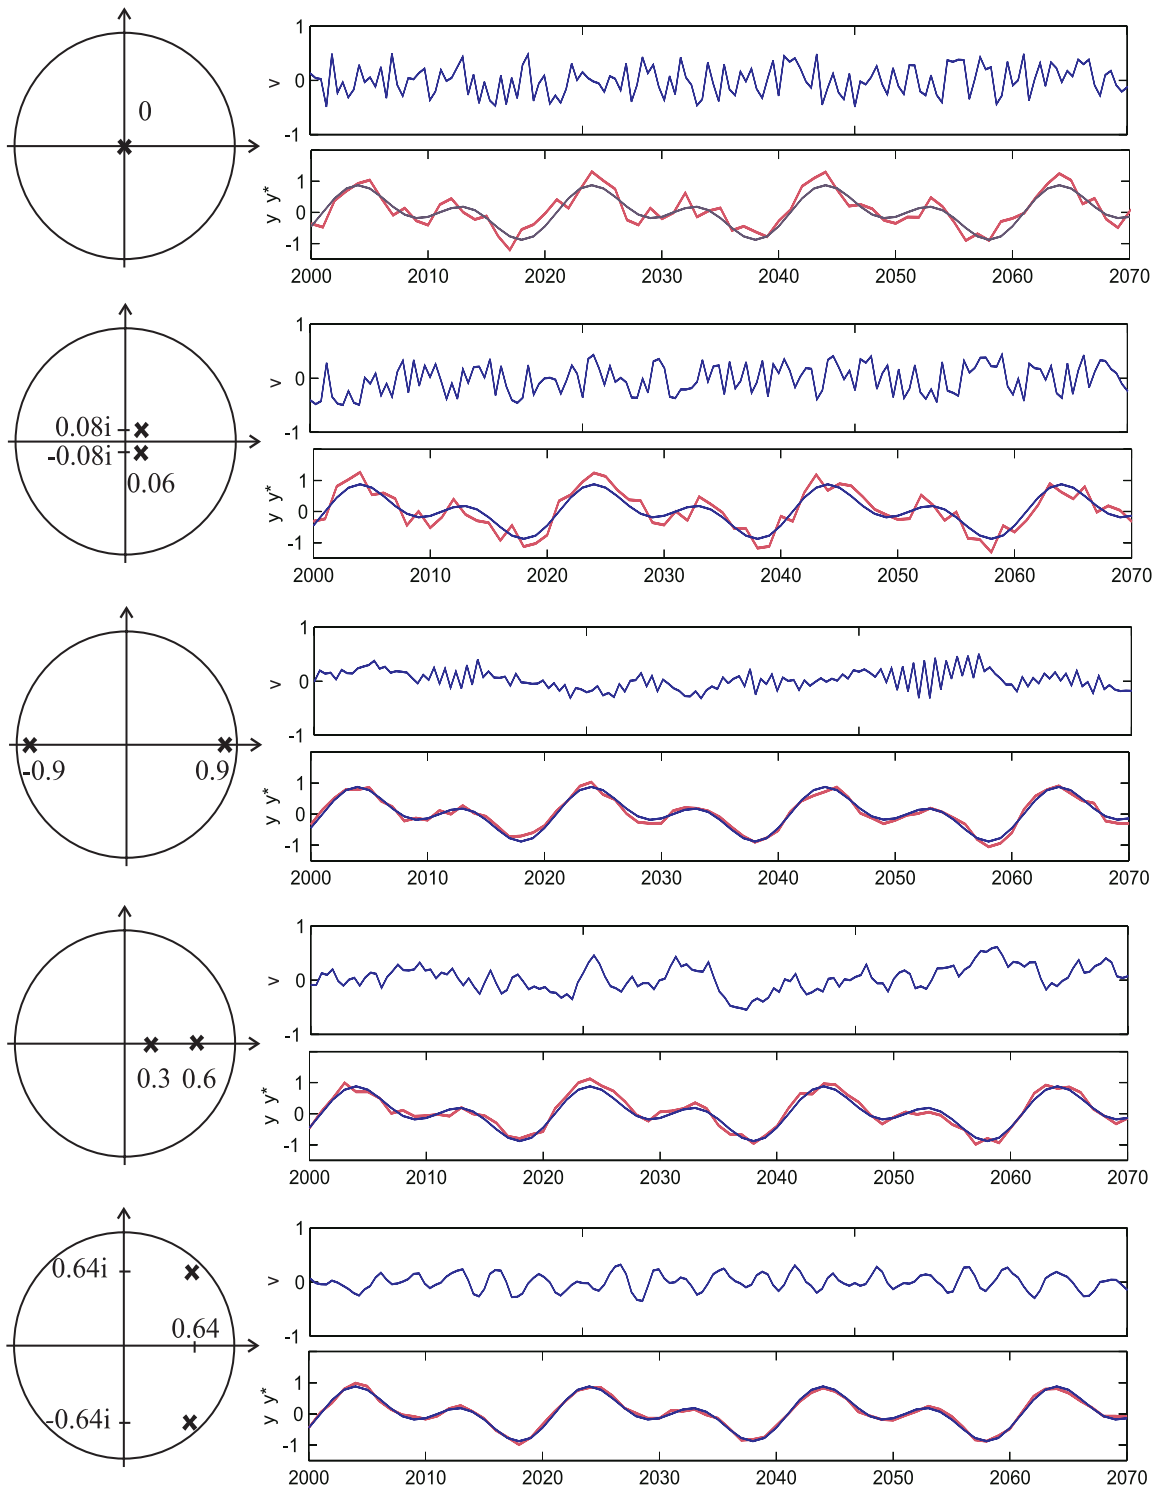


Figure 4.5: Plant output y versus reference output y^* for various locations of the poles of the disturbance model

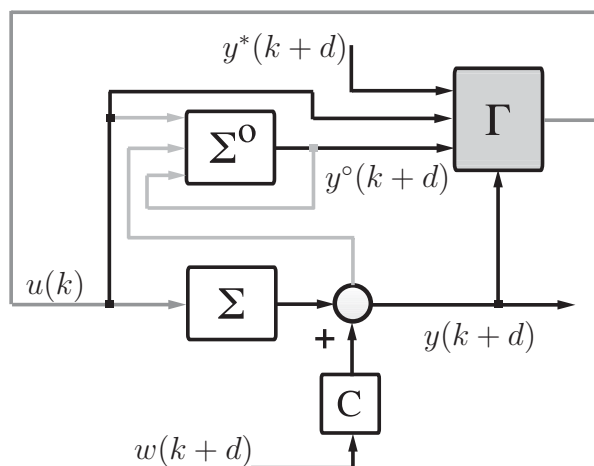


Figure 4.6: Minimum variance control

Stochastic Adaptive Control

In principle, an adaptive version of the above control laws can be obtained by simply replacing the coefficients in (4.32) or (4.36) by their estimated values. A question arises, as to how the estimation model should be chosen. If the noise terms are neglected in the definition of the regression vector, the resulting output error will have zero mean since $E\{w(k)\} = 0$ for all $k > 0$, but its variance may be large. In this paragraph we briefly discuss the stochastic adaptive control problem and present one of its solutions.

In studying this problem it is standard to assume that the white noise sequence $\{w(k)\}$ not only drives the model generating the stochastic disturbance input $v(k)$ but also affects the output of the system $y(k)$ directly. Let us start with a simple example:

$$\Sigma : \quad y(k+1) = ay(k) + cw(k) + w(k+1) + u(k) \quad (4.37)$$

with $|a|, |c| < 1$. $w(k)$ cannot be used in the definition of the estimation model $\hat{\Sigma}$, since it is not measurable. The idea is to use the model error $\epsilon(k) = y(k) - \hat{y}(k)$ instead of $w(k)$:

$$\hat{\Sigma} : \quad \hat{y}(k+1) = \hat{a}y(k) + \hat{c}\epsilon(k) + u(k) \quad (4.38)$$

The motivation for using $\epsilon(k)$ to replace $w(k)$ is that, if the parameters converge to their true values, the residual identification error $\epsilon(k) = y(k) - \hat{y}(k)$ is white noise $w(k)$. By virtue of the certainty equivalence principle, the control law is obtained by replacing $\hat{y}(k+1)$ in (4.38) by $y^*(k+1)$:

$$\Gamma : \quad u(k) = y^*(k+1) - \hat{a}y(k) - \hat{c}\epsilon(k) \quad (4.39)$$

It is obvious that no $u(k)$ at time k can compensate for the noise $w(k+1)$ at time $k+1$. However, Γ asymptotically cancels out the effect of the noise $w(k)$ at instant of time k . The

expected value of the squared error satisfies $E\{[y(k+1) - y^*(k+1)]^2 | k\} = E\{w^2(k+1) | k\} = \sigma^2$ as $k \rightarrow \infty$, and equation (4.39) is referred to as a minimum variance controller. \square

The concept can be generalized as follows. The plant is usually given as an ARMAX (ARMA with eXogenous input) model:

$$A(q^{-1})y(k) = q^{-d}B(q^{-1})u(k) + C(q^{-1})w(k) \quad (4.40)$$

where A, B and C are polynomials in the delay operator q^{-d} , of degrees n_A, n_B and n_C respectively:

$$\begin{aligned} A(q^{-1}) &= 1 + a_1q^{-1} + a_2q^{-2} + \dots + a_{n_A}q^{-n_A} \\ B(q^{-1}) &= b_0 + b_1q^{-1} + b_2q^{-2} + \dots + b_{n_B}q^{-n_B} \\ C(q^{-1}) &= 1 + c_1q^{-1} + c_2q^{-2} + \dots + c_{n_C}q^{-n_C} \end{aligned} \quad (4.41)$$

It is assumed that the roots of $C(q^{-1})$ lie strictly inside the unit circle. $\{w(k)\}$ is a white noise sequence with the properties (4.19) to (4.21). Since $C(q^{-1})$ is monic, $w(k)$ directly affects the output $y(k)$ at a given instant k . The control input is obtained by minimizing the mean-square tracking error $E\{[y(k+d) - y^*(k+d)]^2\}$. It can be shown (see Goodwin, 1984 [27]) that this is equivalent to

$$\min_{u(k)} E\{[y^o(k+d|k) - y^*(k+d)]^2\} \quad (4.42)$$

where $y^o(k+d|k)$ is the optimal d -step ahead prediction of $y(k)$:

$$y^o(k+d|k) = E\{y(k+d)|k\} \quad (4.43)$$

In certainty equivalence control, the estimated parameters are used to determine $u(k)$. From the discussion in chapter 3.1 it is clear that the control input is based upon an underlying estimation process. More precisely, $u(k)$ is defined implicitly by an equation of the form $y^*(k+d) = \phi(\cdot)^T \hat{\theta}(\cdot)$ where the regression vector $\phi(\cdot)$ is yet to be defined. If the parameter estimates are determined such that $\phi(\cdot)^T \hat{\theta}(\cdot)$ is equal to $y^o(k+d|k)$, then the corresponding $u(k)$ clearly minimizes the expression in (4.42). In other words, the optimum is attained whenever $\phi(\cdot)^T \hat{\theta}(\cdot) \rightarrow y^o(k+d|k)$.

Numerous methods have been proposed in the identification literature to deal with the problem of parameter estimation in a stochastic environment (see e.g. Landau, 1998 [43]). As an example, we present the extended least-squares (ELS) algorithm which belongs to the class of pseudo linear regression algorithms. The distinctive feature of this class of algorithms is that the components of the regression vector depend upon previous values of the estimated parameters. As described above, the idea is to design an estimator producing an error $\epsilon(\cdot)$ that becomes white noise asymptotically.

Equation (4.40) can be rewritten as

$$y(k) = [1 - A(q^{-1})]y(k) + q^{-d}B(q^{-1})u(k) + C(q^{-1})w(k) \quad (4.44)$$

Adding and subtracting $[C(q^{-1}) - 1]\epsilon(k)$ we obtain

$$\begin{aligned} y(k) &= [1 - A]y(k) + q^{-d}B u(k) + C w(k) + [C - 1]\epsilon(k) - [C - 1]\epsilon(k) \\ &= \phi(k-1)^T \theta_0 + C w(k) - [C - 1]\epsilon(k) \end{aligned} \quad (4.45)$$

using

$$\begin{aligned} \phi(k-1) &= [y(k-1), \dots, y(k-n_A), u(k-d), \dots, u(k-d-n_B), \epsilon(k-1), \dots, \epsilon(k-n_C)]^T \\ \theta_0 &= [-a_1, \dots, -a_{n_A}, b_0, \dots, b_{n_B}, c_1, \dots, c_{n_C}]^T \end{aligned}$$

The elements of θ_0 are the coefficients of the polynomials in equation (4.41). The identification model reads:

$$\hat{y}(k) = \phi(k-1)^T \hat{\theta}(k-1) \quad (4.46)$$

Subtracting (4.46) from (4.45) we obtain:

$$\epsilon(k) = \frac{1}{C(q^{-1})} \phi(k-1)^T [\theta_0 - \hat{\theta}(k-1)] + w(k) \quad (4.47)$$

Using $\epsilon(k)$, the parameters are updated according to

$$\hat{\theta}(k) = \hat{\theta}(k-1) + P(k-1) \phi(k-1) \epsilon(k) \quad (4.48)$$

where $P(k-1)$ is a matrix satisfying $P(k)^{-1} = P(k-1)^{-1} + \phi(k)\phi(k)^T$, $P(-1)^{-1} > 0$. Note that $\hat{y}(k)$ in equation (4.46) is called the *a priori* prediction, obtained *before* the parameters have been updated. The algorithm can be improved by making use of the most recent parameter estimate. We obtain the *a posteriori* prediction:

$$\bar{y}(k) = \phi(k-1)^T \hat{\theta}(k) \quad (4.49)$$

A posteriori prediction errors are sometimes used in the definition of the regression vector to simplify the convergence analysis of the algorithm. However, the distinction is of secondary importance. In any case, the entries of the regression vector depend upon previous values of the estimated parameters. It is seen that $\epsilon(k)$ will become white noise asymptotically if the parameter error $\tilde{\theta}(k-1) = [\theta_0 - \hat{\theta}(k-1)] \rightarrow 0$ or, alternatively, the first term in equation (4.47) tends to zero.

The a posteriori estimate $\bar{y}(k)$ obtained by the ELS algorithm is the optimal one-step-ahead prediction of $y(k)$, i.e. $\bar{y}(k) = E\{y(k)|k-1\}$. Equation (4.49) allows us to determine the effect of the past values of the noise on the future response of the plant. With the justification

provided above, the adaptive controller Γ is designed on the basis of the optimal prediction of the system output. The control input is obtained by replacing the predicted output $\bar{y}(k)$ by the desired output $y^*(k)$.

$$y^*(k) = \phi(k-1)^T \hat{\theta}(k) \quad (4.50)$$

When solving this equation for $u(k)$ it follows that all past noise terms are cancelled out. Hence, the control law is the stochastic equivalent of the deadbeat-control law obtained from (3.27) in the deterministic case. The residual error has minimum variance. To ensure convergence of Γ to the minimum variance controller the following conditions have to be satisfied:

1. The orders n_A , n_B and n_C of the polynomials in equation (4.41) are known.
2. The relative degree d is known and equal to 1 (in this case).
3. $B(z^{-1})$ has all roots inside the unit circle.
4. $[1/C(z^{-1}) - \frac{1}{2}]$ is positive real.

While conditions 1 to 3 are the same as in the deterministic case, condition 4 is characteristic for estimation techniques based on pseudo linear regressions. It ensures that the extended least-squares algorithm has convergence properties similar to those of its deterministic counterpart (2.37). For further insight, see e.g. (Ljung 1999 [45]). Condition 2 can be relaxed to allow for $d > 1$. Subject to the above assumptions regarding the system and assumptions (4.19) to (4.20) regarding the noise we obtain that both $u(k)$ and $y(k)$ are mean square bounded and that the output $y(k)$ of the plant tracks $y^*(k)$ with minimum variance:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N E[y(k) - y^*(k) | k-1]^2 = \sigma^2 \quad (4.51)$$

where σ^2 is the variance of the white noise $w(k)$ affecting the system.

Comment: The augmented system (4.23) can be written in the form (4.40), except that $C(q^{-1})$ is not monic, since the $v(\cdot)$ in equation (4.1) does not directly affect the output $y(\cdot)$. The derivation of the ELS algorithm was based on the assumption that the a posteriori estimation error becomes white noise asymptotically. Hence, the fact that $C(q^{-1})$ is monic is critical and we have to assume that $\{w(k)\}$ directly affects the output.

Example 4.4 Given a second-order plant subject to a stochastic input disturbance v and affected at the output by white noise w ,

$$[1 + a_1 q^{-1} + a_2 q^{-2}] y(k) = q^{-1} [b_0 + b_1 q^{-1}] u(k) + q^{-1} g_0 v(k) + w(k) \quad (4.52)$$

where $v(\cdot)$ is the output of a first-order system driven by white noise:

$$[1 + d_1 q^{-1}] v(k) = q^{-1} w(k) \quad (4.53)$$




Figure 4.7: Adaptive minimum-variance control

The augmented system obtained by eliminating $v(k)$ in equation (4.52) reads

$$\begin{aligned} [1 + (d_1 + a_1)q^{-1} + (a_2 + d_1a_1)q^{-2} + d_1a_2q^{-3}]y(k) = \\ = q^{-1}[b_0 + (b_1 + b_0d_1)q^{-1} + d_1b_1q^{-2}]u(k) + [1 + d_1q^{-1} + g_0q^{-2}]w(k) \end{aligned} \quad (4.54)$$

and is of the form (4.40). All coefficients are assumed to be unknown, but $C(q^{-1}) = 1 + d_1q^{-1} + g_0q^{-2}$ satisfies the conditions stated in point 3 of the above list. The $d = 1$ step ahead identification model is of the form $\hat{y}(k+1) = \phi(k)^T \hat{\theta}(k)$ where

$$\phi(k) = [y(k) \ y(k-1) \ y(k-2) \ u(k) \ u(k-1) \ u(k-2) \ \epsilon(k) \ \epsilon(k-1) \ \epsilon(k-2)]^T \quad (4.55)$$

$\hat{\theta}(k) \in \mathbb{R}^9$ is the vector of parameter estimates and $\epsilon(k) = \hat{y}(k) - y(k)$. $\hat{\theta}(k)$ contains the parameters upon which the adaptive minimum variance control law is based:

$$\begin{aligned} u(k) = 1/\hat{\theta}_4(k) [y^*(k+1) - \hat{\theta}_1(k)y(k) - \hat{\theta}_2(k)y(k-1) - \hat{\theta}_3(k)y(k-2) \\ - \hat{\theta}_5(k)u(k-1) - \hat{\theta}_6(k)u(k-2) - \hat{\theta}_7(k)\epsilon(k) - \hat{\theta}_8(k)\epsilon(k-1) - \hat{\theta}_9(k)\epsilon(k-2)] \end{aligned} \quad (4.56)$$

The simulation displayed in figure (4.7) reveals, that the closed loop system tracks an arbitrary reference input and that the control error becomes white asymptotically. Due to the presence of noise, the degree of excitation is superior to the one obtained when no disturbances are present. As a consequence, the parameter error $\tilde{\theta}$ converges to zero. In the last row, the SPR condition has been verified, i.e. $[C(z^{-1}) - 1/2]$ has been evaluated for all $z^{-1} = e^{i\omega T_S}$, $\omega T_S = -\pi \dots \pi$.




Figure 4.8: Nonlinear oscillation generated by a first-order discrete-time system

4.2 Nonlinear Disturbance Rejection

In this section, we consider the case where the plant Σ as well as the disturbance generating system Σ_v is nonlinear:

$$\begin{aligned} \Sigma : \quad x(k+1) &= f[x(k), u(k), v(k)] \\ y(k) &= h[x(k)] \\ \Sigma_v : \quad x_v(k+1) &= f_v[x_v(k)] \\ v(k) &= h_v[x_v(k)] \end{aligned} \quad (4.57)$$

where $f : \mathbb{R}^n \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^n$, $h : \mathbb{R}^n \rightarrow \mathbb{R}$ and $f_v : \mathbb{R}^{n_v} \rightarrow \mathbb{R}^{n_v}$, $h_v : \mathbb{R}^{n_v} \rightarrow \mathbb{R}$ are unknown \mathcal{C}^1 -functions of their arguments. As in the linear case, the objective is to determine a control input $u(k)$ such as to make $x_v(k)$ unobservable through the output $y(k)$. Our interest is in bounded nonvanishing disturbances $v(k)$. In the nonlinear domain, the class of unforced systems Σ_v which generate a disturbance signal having the required properties is not immediately evident. We assume that Σ_v is of the following form:

$$\Sigma_v : \quad v(k+1) = \alpha v(k) - \beta f_v[v(k)] \quad (4.58)$$

where $\alpha > 1$ and β are constant parameters. Since the linear part of the system (4.58) is unstable, β has to be chosen such that the nonlinearity $f_v[v(k)]$ prevents $v(k)$ from growing in an unbounded fashion. If so, $v(k)$ lies in a compact set $S \subset \mathbb{R}$. The initial value $v(0)$ has to be chosen from this set.

Example 4.5

$$\Sigma_v : \quad v(k+1) = 3v(k) - 2v(k)^3 \quad (4.59)$$

For this particular choice of the parameters and for the initial condition $v(0) \in (-\sqrt{2}; \sqrt{2}) \setminus \{-1, 0, 1\}$, Σ_v generates a bounded nonvanishing output, shown in figure (4.8). The figure

also displays the linear and nonlinear parts of equation (4.59) in a one-dimensional map which illustrates the mechanism: Starting from an initial point $v(0)$ near the origin, $v(k)$ at first grows linearly until the nonlinear function becomes dominant for larger values of $v(k)$. The parameter β is chosen such that $v(k)$ is mapped back to a point $v^*(0)$ within the linear, unstable region. Since $v^*(0)$ does not necessarily coincide with the original point $v(0)$, the oscillation generated by Σ_v may not be periodic. In fact, systems of the form (4.58) with $\alpha > 1$ may give rise to quasiperiodic and chaotic motion. A classical example is the logistic map

$$v(k+1) = 4v(k) - 4v^2(k)$$

for values $v(0) \in [0, 1]$. Hence, in order to represent e.g. a periodic disturbance, the function $f_v(\cdot)$ in equation (4.58) has to be appropriately chosen. \square

The principle of nonlinear disturbance rejection is the same as in the linear case and relies upon a system representation of augmented order. This representation is obtained by eliminating $v(k)$ from the input–output model of the plant. In section (3.2), the existence of an input–output model of the nonlinear state–vector representation was seen to depend on the observability of the linearized system. A similar condition is needed in order to derive an input–output model of the composite system $\Sigma \circ \Sigma_v$. The linearized equations are given by:

$$\begin{aligned} x(k+1) &= Ax(k) + bu(k) + b_v v(k) \\ x_v(k+1) &= A_v x_v(k) \\ y(k) &= c^T x(k) \\ v(k) &= c_v^T x_v(k) \end{aligned} \tag{4.60}$$

where $A = \frac{\partial f(x,u,v)}{\partial x} \Big|_0 \in \mathbb{R}^{n \times n}$, $b = \frac{\partial f(x,u,v)}{\partial u} \Big|_0 \in \mathbb{R}^n$, $b_v = \frac{\partial f(x,u,v)}{\partial v} \Big|_0 \in \mathbb{R}^n$, $c = \frac{\partial h(x)}{\partial x} \Big|_0 \in \mathbb{R}^n$ and $A_v = \frac{\partial f_v(x_v)}{\partial x_v} \Big|_0 \in \mathbb{R}^{n_v \times n_v}$, $c_v = \frac{\partial h_v(x_v)}{\partial x_v} \Big|_0 \in \mathbb{R}^{n_v}$. The Jacobians are evaluated at the origin of the composite system. In the linear case, it was seen that if the pairs $[c, A]$ and $[c_v, A_v]$ are observable, the system has a linear input–output representation of dimension $(n + n_v)$, given by equation (4.8). Under the same conditions, a nonlinear input–output map of the form

$$\bar{\Sigma} : \quad y(k+d) = \mathcal{F}[Y_{n+n_v}(k), U_{n+n_v}(k)] \tag{4.61}$$

exists locally in the neighborhood of the origin in extended state–space \mathbb{R}^{n+n_v} . The map describes a system of extended order where $Y_{n+n_v}(k) = [y(k), \dots, y(k - \langle n + n_v - 1 \rangle)]$, $U_{n+n_v}(k) = [u(k), \dots, u(k - \langle n + n_v - 1 \rangle)]$ and $d \geq 1$. Given $\bar{\Sigma}$, the problem of disturbance rejection consists in determining a tracking control law such that the closed–loop system is asymptotically stable and tracks any bounded reference trajectory $\{y^*(k)\}$. Since $\bar{\Sigma}$ represents a system of augmented order, from which the explicit dependence on $v(k)$ has

been removed, the design of the controller proceeds as in the disturbance-free case. As seen in chapter (3.2) a solution to the problem exists provided that the relative degree d is well defined and the zero-dynamics are stable. Since the disturbance rejection problem is solved

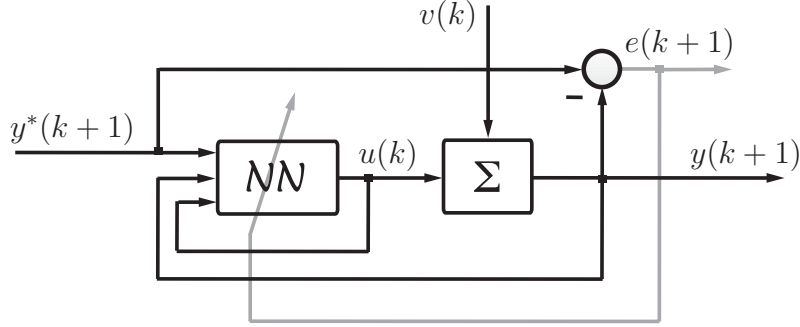


Figure 4.9: Neurocontrol in the presence of disturbances

by designing a controller based on the NARMA-model of $\bar{\Sigma}$ we have to assume that the values of inputs $\{u(k)\}$ and outputs $\{y(k)\}$ do not leave the neighborhood $\mathcal{U} \times \mathcal{Y}$ in which the NARMA-model is valid. This implies that the reference trajectory $\{y^*(k+d)\}$ lies in that neighborhood.

Example 4.6 The following example reveals some practical aspects while designing a neurocontroller for disturbance rejection. We consider the second-order system,

$$\Sigma : \quad y(k+1) = \frac{3y^2(k)[1-v^2(k)] \tanh y(k-1) + u(k)}{1+y^2(k)+y^2(k-1)} \quad (4.62)$$

In the simulation, the bounded disturbance $v(k)$ was chosen to be the filtered output of equation (4.59), described by the following homogenous input-output map:

$$\Sigma_v : \quad v(k+1) = 0.2v(k) + 0.2v(k-1) + [v^2(k-1) - 1][v(k) - v(k-1)] \quad (4.63)$$

From the above discussion it is clear that there exists a local NARMA representation of the composite system which is of augmented order $n + n_v = 4$. However, control based this model was found to result in poor performance. Hence the order of the NARMA-model was further increased. This can be thought of as a way of increasing the domain in which the augmented input-output representation is valid. If the composite system is of 5th-order, almost complete disturbance rejection was obtained:

$$\bar{\Sigma} : \quad y(k+1) = \mathcal{F}(Y_5(k), U_5(k)) \quad (4.64)$$

with $Y_5(k) = [y(k), \dots, y(k-4)]^T$ and $U_5(k) = [u(k), \dots, u(k-4)]^T$. A multilayer neural network (MNN) of the form $\mathcal{N}_{10,25,35,1}^3$ (i.e. a three-layered NN with a single output, 35




Figure 4.10: Tracking performance using a neurocontroller; (i): no disturbance rejection, (ii): disturbance rejection with insufficient number of input-output measurements, (iii): application of the control law defined in equation (4.65).

neurons in the first hidden layer, 25 neurons in the second hidden layer, and a 10-dimensional input vector) is used to identify the augmented system (4.64). After 350000 iterations with a step size of $\eta = 0.1$, perfect matching of the model and plant is achieved. Once the identification process is complete, a second neural net is trained to approximate the nonlinear map representing the controller:

$$\bar{\Gamma} : \quad u(k) = \mathcal{N}_c(y^*(k+1), y(k), \dots, y(k-4), u(k-1), \dots, u(k-4)) \quad (4.65)$$

The network is of the form $\mathcal{N}_{10,15,25,1}^3$ and is trained for 100000 steps with a step size of $\eta = 0.01$. Note that, since the controller $\bar{\Gamma}$ is in the feedback-loop of a dynamical system $\bar{\Sigma}$, the parameter adjustments in the second neural network have to be carried out using dynamic backpropagation. However, dynamic gradient methods are computationally intensive, so a static method is used to generate approximate gradients. The resulting tracking performance is shown in figure (4.10) where the reference trajectory was chosen to be $y^*(k+1) = \sin(2\pi k/10) + \sin(2\pi k/20)$. \square

When comparing the results of the last two chapters one realizes, that most arguments from linear theory have nonlinear counterparts. Although the results are valid only in a certain

neighborhood of the equilibrium state, they provide the justification for transferring linear methods to certain nonlinear problems. In fact, the problem of rejecting disturbances which are the output of an unforced system was solved for both linear and nonlinear systems by using the same concept of extending the order of the controller.

While dealing with stochastic disturbances, it was seen that the above method results in good performance provided that the disturbance is highly correlated with its past values. In addition, it was shown that the estimation error could be used to cancel out the effect of past values of the noise. This resulted in a minimum variance control law. It is only natural to assume, that a similar method exists in the nonlinear domain, if the latter is restricted to a neighborhood of an equilibrium state. Further work is needed to investigate whether the variance of the control error can be minimized even as the plant is nonlinear

$$y(k) = f[y(k-1), \dots, y(k-n_A), u(k-1), \dots, u(k-n_B), w(k), \dots, w(k-n_C)] \quad (4.66)$$

and $\{w(k)\}$ is a white noise sequence.

Chapter 5

Design Considerations

Adaptive controllers are quite complicated devices. Although many successful applications to various practical problems exist (see e.g. [51], Chap. 11), there are no general rules which cover all aspects of a given design problem. The available theory has to be complemented by simulations and experiments before using adaptive schemes in real-world processes. In the following chapter, we set out to address some of the practical issues of adaptive control of linear systems with unknown parameters. Among the many questions that arise we focus on the ones which directly build upon the theory presented thus far.

As a first step, we discuss an alternative to the deadbeat concept which was used extensively in the preceding chapters. Due to its simplicity, deadbeat-control is theoretically appealing but, on the other hand, impractical since it requires the system to respond accurately over the entire frequency range $[0 \dots 1/2T_S]$ Hz. This inevitably leads to an excessively large control effort. It is seen that within the model reference framework it is possible to impose an (almost) arbitrary closed-loop behavior onto the system. Thus, as far as the choice of the controller structure is concerned, the design process becomes amenable to linear design rules. It shall be noted that obviously the linear viewpoint is admissible only if we assume that the parameters have converged to some constant values. During the adaptation process the system is nonlinear and offers little insight into the nature of the signals generated in closed loop. Even if the system is stable the control input may be transiently large and oscillatory due to the time-variation of the parameters. In some cases, it may prove advantageous to reduce the speed of adaptation such that the system behaves (almost) like a linear one. In others, the control effort is large due to an unfavorable initialization of the parameter estimates. In such a case the adaptation speed has to be increased such that the parameter attain stable regions in parameter space. A systematic way of selecting the “right” adaptive gain is presented in the second part of this chapter. In particular, we develop a novel approach

to the problem of limiting the actuator stress in adaptive control. The requirements on the rate and magnitude of the input are typically large during the initial phase of parameter convergence. This motivates the use of an optimization algorithm which minimizes the rate and magnitude of the input through a variation of the adaptive gain η . In certainty equivalence control, the input signal $u(k)$ is defined in terms of the most recent estimate $\hat{\theta}(k)$. This means that, given $\hat{\theta}(k-1)$, the update $\hat{\theta}(k)$ can be determined such as to result in a minimal control effort. How this is to be carried out even while assuring the stability of the adaptive closed-loop system is demonstrated in section 5.2.1.

A substantial reduction of the control effort can also be achieved if — based on prior information about the plant — the estimates are initialized close to the actual parameters. The idea is that, with small transient errors, the system behaves almost like a linear one. However, a phenomenon which has been observed frequently in practice is that even if the estimates $\hat{\theta}$ are initialized close to the actual parameters θ_0 , they typically “jump away” before converging, i.e. the parameter error $\tilde{\theta} = \hat{\theta} - \theta_0$ is large during the initial stages of parameter convergence no matter how small $\tilde{\theta}(0)$ is. The effect is caused by unmodelled dynamics. Any physical system can be modelled only partially so that, in general, the order of the system is larger than that of the model. This mismatch generates a disturbance signal whose magnitude depends upon the input of the system. If unmodelled dynamics are excited, the identification error is large, causing the parameter estimates to diverge from the actual values. It was soon realized that the presence of unmodelled dynamics may provoke instability of the adaptive system. The origin of this instability lies in the fact that the error between the plant and the model is not exclusively due to unknown parameters but due to a *structural* uncertainty regarding the order of the plant. An approach to keep the parameter estimates bounded is to introduce a relative dead-zone (Kreisselmeier 1986, [41]) the size of which depends upon a suitably defined modelling error. The size of the dead-zone (which also gives a bound on the tracking error) can be reduced if the modelling error is small.

An interesting application of the above optimization procedure is to smooth the control input by imposing an appropriate rate constraint on $u(k)$. The idea is that high-frequency components contained in the input signal are attenuated, such that no unmodelled dynamics (which are typically due to fast parasitics) are excited. The optimized adaptive scheme enhances robustness to unmodelled dynamics, since it affects the nature of the input generated in the adaptive loop. In the example included in section 5.3, the modelling error was so small that the dead-zone could be completely dispensed with.

5.1 Choice of the Reference Model

The design steps presented in chapter 3 follow a very simple procedure: In order to determine the control law, the output $\hat{y}(k+d)$ of the identification model is replaced by the reference output $y^*(k+d)$. The equation thus obtained implicitly defines $u(k)$. When all parameters converge to their true values, the procedure clearly corresponds to implementing a deadbeat-controller. The resulting closed-loop system has a pole of multiplicity d at $z=0$. It is well known that deadbeat-control is impractical since the cancellation of the process poles entails an excessive control effort and high sensitivity to measurement noise. An important question is whether the design procedure can be modified such as to obtain a more general class of controllers. In many applications, PI-controllers are employed. If the parameters of the system are known, the controller is tuned to optimal values using standard procedures, e.g. the symmetrical optimum [71]. When the system is unknown, it is *not even clear* whether the PI-controller is capable of stabilizing the system.

In model reference adaptive control (MRAC) the question of stability is resolved by requiring the closed-loop system to match a given, stable reference system. In fact, the deadbeat approach corresponds to the most primitive case where the reference model consists of a simple d -step delay. A large variety of control laws can be conceived through the choice of different reference models. The objective of the following paragraph is to outline the principal design steps involved in adaptive control based on an arbitrary reference model. The problem may be subdivided into two parts:

1. Determine a controller, parameterized in the vector of unknown parameters θ , such that there exists a constant $\theta \equiv \theta_0$ for which the closed-loop system matches a given reference system.
2. Determine an adaptive procedure by which the vector of parameter estimates $\hat{\theta}(\cdot)$ is to be updated such that $\lim_{k \rightarrow \infty} |y(k) - y^*(k)| = 0$

An equivalent representation of the system (3.20) is given by

$$A(q^{-1})y(k) = q^{-d}B(q^{-1})u(k) \quad (5.1)$$

where $A(q^{-1})$, $B(q^{-1})$ are polynomials in the delay operator q^{-1} of degrees n_A and n_B respectively as in equation (4.10) or (4.41). Item one of the above list corresponds to the algebraic part of the problem which consists in finding an appropriate controller structure. The reference system is given by

$$C(q^{-1})y^*(k) = q^{-d^*}D(q^{-1})r(k) \quad (5.2)$$

where $r(k)$ is a bounded piecewise-continuous function of time. For the control law to be causal, the relative degree d^* of the reference system must be greater than or equal to that of the plant, i.e. $d \leq d^*$. Since the plant is linear the controller will also be linear,

$$M(q^{-1})B(q^{-1})u(k) = z^{-[d^*-d]}D(q^{-1})r(k) - N(q^{-1})y(k) \quad (5.3)$$

with unknown polynomials $M(q^{-1})$, $N(q^{-1})$ of degrees n_M and n_N that have yet to be determined.

The analytic part, which is stated in item 2, deals with the problem of determining adaptive laws by which the parameters of the controller are adjusted such that they evolve towards the desired values and the output of the system asymptotically tracks the reference output while all signals of the system remain bounded.

Algebraic Part

The first step in the adaptive control design is to demonstrate that the controller structure has enough freedom to ensure the existence of a solution assuming that the parameters are known. The point here is to show that there is a constant parameter vector θ_0 for which the closed-loop system matches the reference system. Inserting the control law (5.3) into (5.1) we obtain

$$[M(q^{-1})A(q^{-1}) + q^{-d}N(q^{-1})]y(k) = q^{-d^*}D(q^{-1})r(k) \quad (5.4)$$

where $M(q^{-1})$ and $N(q^{-1})$ are unknown polynomials in the unit delay operator q^{-1} . Comparing this to the reference system (5.2), $M(q^{-1})$ and $N(q^{-1})$ have to be determined such that the following equality holds:

$$M(q^{-1})A(q^{-1}) + q^{-d}N(q^{-1}) = C(q^{-1}) \quad (5.5)$$

The existence of two polynomials $M(q^{-1})$ and $N(q^{-1})$ that satisfy the equation is guaranteed by the following fact known as Bézout's lemma:

Let $P(q^{-1})$ and $Q(q^{-1})$ be polynomials of degrees n_P and n_Q respectively which are relatively prime. Then polynomials $M(q^{-1})$ and $N(q^{-1})$ exist such that for any arbitrary polynomial $C(q^{-1})$ the following equality holds:

$$M(q^{-1})Q(q^{-1}) + P(q^{-1})N(q^{-1}) = C(q^{-1}) \quad (5.6)$$

In our case, $P(q^{-1}) = q^{-d}$ and the monic polynomial $A(q^{-1})$ are clearly coprime. In the adaptive context we are not interested in actually determining the polynomials. All we need to show is their existence. For the purpose of analysis, though, it may be of interest to

compute $M(q^{-1})$ and $N(q^{-1})$. Let us first consider the case where $C(q^{-1})$ is monic and the relative degree of the system is $d = 1$. A straightforward calculation yields $M(q^{-1}) = 1$ and $N(q^{-1}) = q[C(q^{-1}) - A(q^{-1})]$.

For an arbitrary $C(q^{-1})$ and $d > 1$ the computation is more involved. Choosing polynomials of degrees $n_M \leq d - 1$ and $n_N \leq n - 1$, we write

$$\begin{aligned} M(q^{-1}) &= m_0 + m_1 q^{-1} + \cdots + m_{d-1} q^{-d+1} \\ N(q^{-1}) &= n_0 + n_1 q^{-1} + \cdots + n_{n-1} q^{-n+1} \end{aligned}$$

The closed-loop denominator polynomial $C(q^{-1})$ has the general form

$$C(q^{-1}) = c_0 + c_1 q^{-1} + \cdots + c_{n^*} q^{-n^*}$$

with degree $n^* \leq d + n - 1$. Recalling that

$$A(q^{-1}) = 1 - a_1 q^{-1} - \cdots - a_{n_A} q^{-n_A}$$

we equate the coefficients of both sides of equation (5.5). In the following vector notation the i -th row contains the coefficients pertaining to the i -th power of the delay operator q^{-i} .

$$\begin{bmatrix} -a_n & 0 & \cdots & \cdots & 0 & | & 1 & 0 & \cdots & \cdots & 0 \\ -a_{n-1} & -a_n & 0 & \cdots & 0 & | & 0 & 1 & \cdots & \cdots & 0 \\ \vdots & & & & \vdots & | & \vdots & & & & \vdots \\ -a_{n-d} & \cdots & \cdots & \cdots & -a_{n-1} & | & 0 & 0 & \cdots & 1 & 0 & \cdots & 0 \\ \vdots & & & & \vdots & | & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & \cdots & \cdots & \cdots & -a_{d-1} & | & 0 & \cdots & & & & & 0 \\ 0 & 1 & \cdots & \cdots & -a_{d-2} & | & & & & & & & \vdots \\ & & \ddots & & \vdots & | & \vdots & & & & & & \vdots \\ 0 & \cdots & 0 & 1 & & | & 0 & \cdots & & & & & 0 \end{bmatrix} \begin{bmatrix} m_{d-1} \\ m_{d-2} \\ \vdots \\ \vdots \\ m_0 \\ n_{n-1} \\ \vdots \\ \vdots \\ n_0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ c_{n^*} \\ c_{n^*-1} \\ \vdots \\ \vdots \\ \vdots \\ c_0 \end{bmatrix}$$

Since q^{-d} and $A(q^{-1})$ are relatively prime, the matrix is nonsingular and hence unique coefficients m_i , $i = 0, \dots, d-1$ and n_j , $j = 0, \dots, n-1$ can be determined such that equation (5.5) is satisfied.

Example 5.1 Consider the second-order system

$$[1 - a_1 q^{-1} - a_2 q^{-2}] y(k) = q^{-2} b_0 u(k)$$

which has relative degree $d = 2$. A controller which shifts the closed-loop poles of the system to arbitrary locations may be given in the following form

$$[m_0 + m_1 q^{-1}] b_0 u(k) = d_0 r(k) - [n_0 + n_1 q^{-1}] y(k) \quad (5.7)$$




Figure 5.1: Closed-loop pole assignment: (i) plant, (ii) closed-loop system, (iii) closed-loop system after pole-zero cancellations

Its parameters are obtained by solving

$$\begin{bmatrix} -a_2 & 0 & 1 & 0 \\ -a_1 & -a_2 & 0 & 1 \\ 1 & -a_1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} m_1 \\ m_0 \\ n_1 \\ n_0 \end{bmatrix} = \begin{bmatrix} c_3 \\ c_2 \\ c_1 \\ c_0 \end{bmatrix} \quad (5.8)$$

where c_i , $i = 0 \dots 3$ are the coefficients of the reference system:

$$[c_0 + c_1q^{-1} + c_2q^{-2} + c_3q^{-3}]y(k) = q^{-2}d_0r(k) \quad (5.9)$$

The resulting controller coefficients are

$$\begin{aligned} m_0 &= c_0 \\ m_1 &= c_1 + a_1c_0 \\ n_0 &= c_2 + a_1c_1 + (a_1^2 + a_2)c_0 \\ n_1 &= c_3 + a_2c_1 + a_2a_1c_0 \end{aligned} \quad (5.10)$$

and may be collected in the vector $\theta_0 = [m_1 \ m_0 \ n_1 \ n_0]$. How the poles of an unstable plant are shifted to match the pole-zero map of the reference system can be verified in figure (5.1). Unstable zeros are used to cancel the poles of the plant which lie outside of the unit circle. In the adaptive case, the values of the controller coefficients cannot be determined since the system is unknown. However, by equation (5.8), θ_0 is known to exist. \square

Analytic Part

The solution to the analytic part is based on a simple and intuitive concept, the *certainty equivalence principle* which was already introduced in chapter 3.1. At every instant of time,

the parameters values generated by an estimator are used as if they corresponded to the true parameters of the system. From equation (5.5) we know that

$$C(q^{-1})y(k) = M(q^{-1})A(q^{-1})y(k) + q^{-d}N(q^{-1})y(k) \quad (5.11)$$

which is equivalent to

$$C(q^{-1})y(k) = q^{-d} [M(q^{-1})B(q^{-1})u(k) + N(q^{-1})y(k)] \quad (5.12)$$

because of (5.1). Collecting past values of inputs and outputs of the system, as in (2.1), we obtain

$$C(q^{-1})y(k) = \phi(k-d)^T \theta_0 \quad (5.13)$$

where

$$\begin{aligned} \phi(k-d) &= q^{-d} [u(k) \quad \dots \quad u(k-n_B-n_M) \quad y(k) \quad \dots \quad y(k-n_N)]^T \\ \theta_0 &= [m_0 b_0 \quad \dots \quad m_{n_M} b_{n_B} \quad n_0 \quad \dots \quad n_{n_N}]^T \end{aligned}$$

Replacing θ_0 by the vector of parameter estimates $\hat{\theta}(k)$ yields

$$C(q^{-1})\hat{y}(k) = \phi(k-d)^T \hat{\theta}(k-1) \quad (5.14)$$

If the identification error is defined as $\epsilon(k) = C(q^{-1})[y(k) - \hat{y}(k)]$, the error equation has the same form as in (2.11):

$$\epsilon(k) = C(q^{-1})y(k) - \phi(k-d)^T \hat{\theta}(k) = \phi(k-d)^T [\theta_0 - \hat{\theta}(k-1)] \quad (5.15)$$

Hence, any of the standard algorithms introduced in chapter 1 can be used to adjust the parameters. The central step in the controller design is to replace $\hat{y}(k)$ in equation (5.14) by the desired output $y(k)^*$. We obtain

$$C(q^{-1})y^*(k) = \phi(k-d)^T \hat{\theta}(k-d) \quad (5.16)$$

which implicitly defines the controller $u(k-d)$. By virtue of the certainty equivalence principle we regard the parameter estimates in equation (5.16) to be the true parameters and solve for $u(k-d)$. The principle has the effect of relating the control error $\epsilon(k) = C(q^{-1})[y(k) - y^*(k)]$ to the identification error. More precisely,

$$e(k) = \epsilon(k) + \phi(k-d)^T [\hat{\theta}(k-1) - \hat{\theta}(k-d)] \quad (5.17)$$

which is obtained by adding and subtracting $C(q^{-1})y^*(k)$ in equation (5.15). Since (5.17) is of the form (3.29) the proof of stability proceeds as in chapter 3.1.4. The simulation in figure (5.2) illustrates that the success of adaptive pole placement depends upon the nature of the

input signal. After the parameters have converged to some constant values, the closed-loop system matches the reference model only in the case of a random input. However, it is also demonstrated, using Bode plots (column 3 of figure 5.2), that the frequency response matches that of the reference system at the frequencies at which the system has been excited (see e.g. row 2 of the figure: $r(k) = \sin(\omega_0 T_S k)$ where ω_0 corresponds to the eigenfrequency of the reference model).

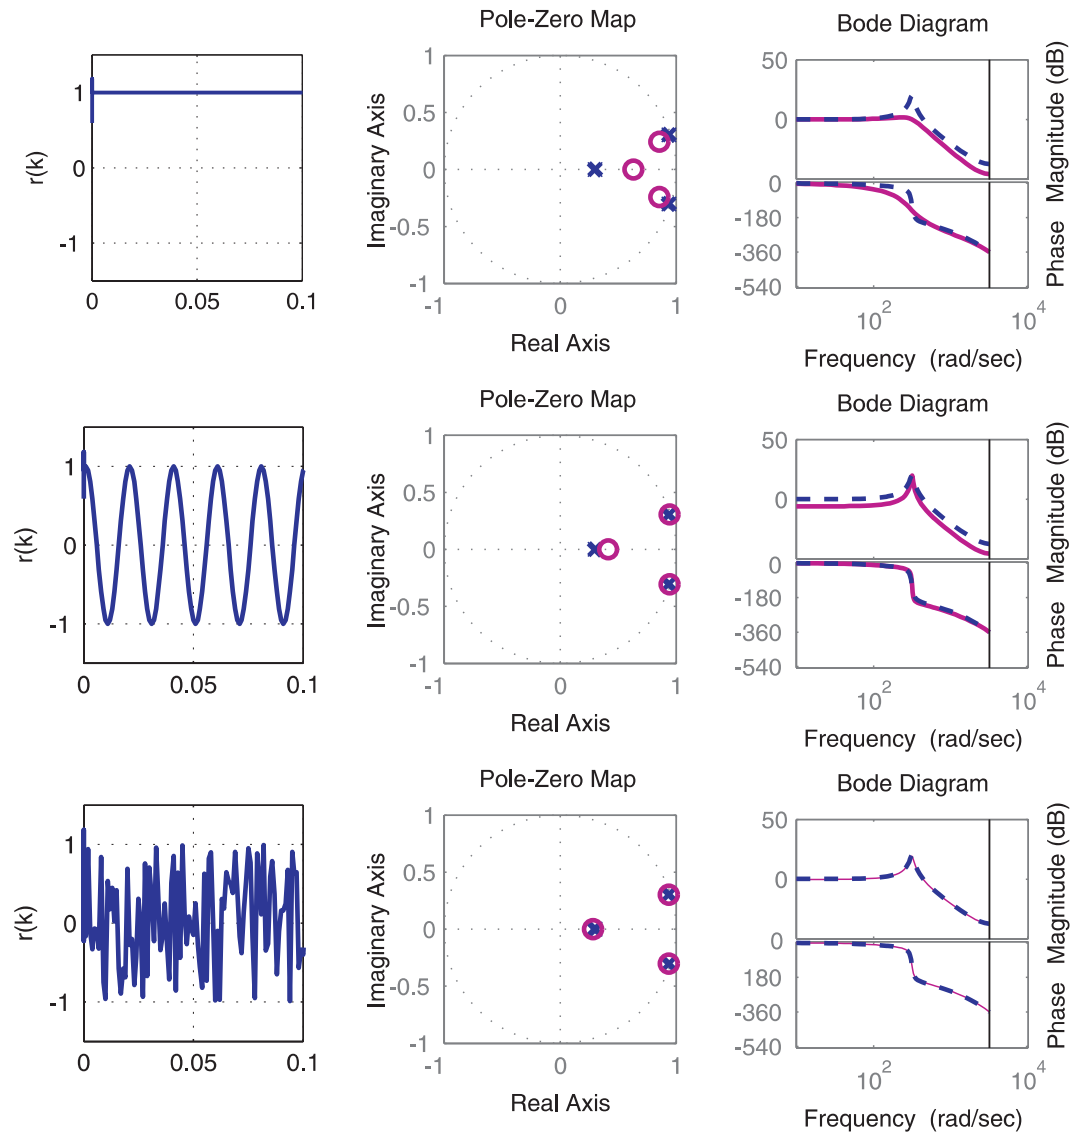



Figure 5.2: Adaptive Pole Assignment

If the system is observable, i.e. a representation of the form (5.1) exists, the poles of the closed-loop system can be placed at arbitrary locations through the choice of a reference model. This means that the whole spectrum of linear design rules can be applied to optimize

Figure 5.3: Example 2: Root locus for $\mu = 0 \dots 1$

the behavior of the closed-loop system. As an example, we include here a way of making the closed-loop system insensitive to measurement noise. It should be kept in mind that, obviously, the property holds only once the parameter have converged.

Example 5.2 Let the reference system be given by

$$[A(q^{-1}) + \mu [1 - A(q^{-1})]] y^*(k) = q^{-1} \mu r(k) \quad (5.18)$$

where $A(q^{-1})$ corresponds to the eigen-dynamics of the plant and $\mu \in [0 \dots 1]$. For $\mu = 1$, the reference model imposes deadbeat behavior, whereas for $\mu = 0$ the poles of the plant remain unchanged. Using the Diophantine equation (5.5), an immediate choice of the controller polynomials is $M(q^{-1}) = 1$ and

$$N(q^{-1}) = \mu q [1 - A(q^{-1})] \quad (5.19)$$

The control law reads

$$u(k) = \frac{\mu}{B(q^{-1})} r(k) - \mu \frac{q [1 - A(q^{-1})]}{B(q^{-1})} [y(k) + w(k)] \quad (5.20)$$

where $w(k)$ is the output measurement (white) noise. If the output is highly corrupted by noise and the plant itself is stable it may be advisable to have a low gain μ in the feedback path of the system. In the limit, $C(q^{-1}) = A(q^{-1})$ and $\mu = 0$. But this means that the eigen-dynamics of the reference model are the same as those of the plant. Even though $A(q^{-1})$ is assumed to be unknown, in practice there is always some prior knowledge about the plant. In the case of the two mass system, example 3.2, it is known e.g. that $A(q^{-1})$ has a pole at 1 and a complex conjugate pair of poles somewhere inside the unit circle. $C(q^{-1})$ can be chosen accordingly. If $C(q^{-1})$ is parameterized as above, the root locus diagram provides an immediate relationship between the smallness of μ and the fact that the poles of $C(q^{-1})$

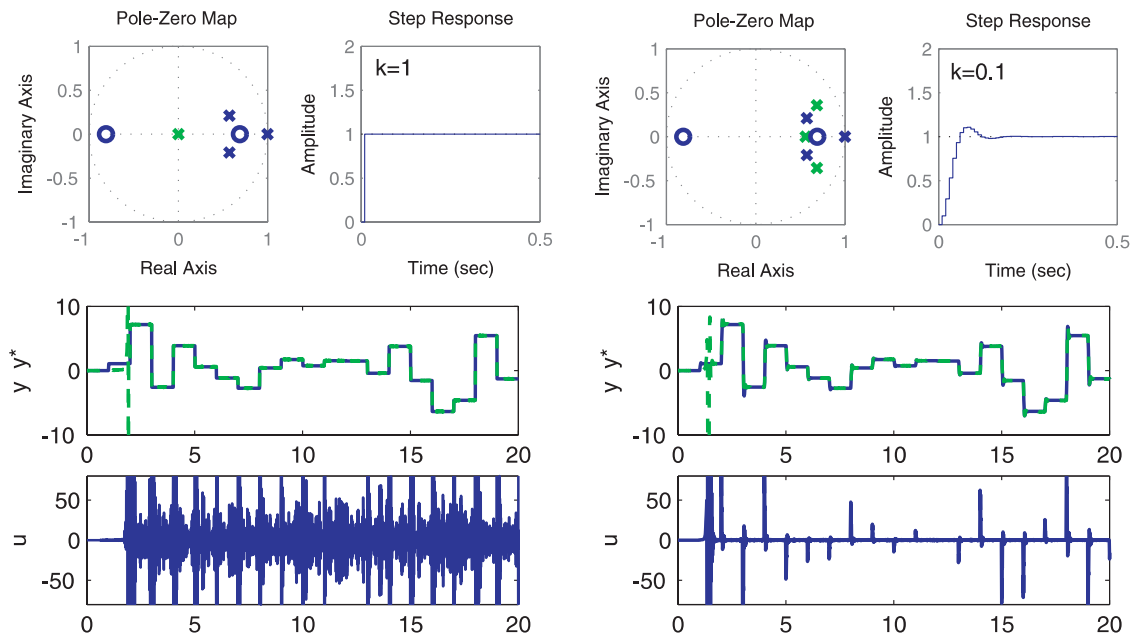


Figure 5.4: Adaptive rejection of measurement noise

lie close to the poles of the original system (\rightarrow see figure 5.3). Now suppose that, exploiting our prior knowledge about the two-mass system, we select $C(q^{-1})$ as described above, and adapt the parameters using an RLS algorithm, we observe that the noise contained in the input signal $u(k)$ is reduced substantially. The response to a piecewise constant reference signal is only slightly altered and corresponds to the step response of a second-order system with small overshoot, as seen in the upper right part of figure 5.4. Clearly, a noise-free control input $u(k)$ can only be expected if the parameters converge to the true values. The necessary excitation is provided by $u(k)$ itself, which contains the measurement noise. As the parameters converge, the noise is damped and the degree of excitation decreases.

5.2 Input Constraints

A common aspect to all real-world control systems is that the control variables are constrained due to the limited power of the actuator. Such limits have the form of amplitude, rate or acceleration constraints and constitute a most commonly encountered class of nonlinearities. Ignoring their presence may have severe consequences on the performance of the system or cause instability. If the actuator limitations are formulated in terms of a saturation nonlinearity in the feedback path of the linear system, the well known circle and

Popov criteria provide frequency–domain sufficient conditions for the stability of the origin of the overall system (see Narendra and Taylor, 1973 [62]). These classical criteria may be applied graphically in the case of a single–input single–output system. More recently, results have been obtained regarding the fundamental issue of determining conditions under which a system is stabilizable with bounded input (Sontag, 1984 [72]).

In the context of optimal control, the problem can be formulated as one of minimizing a cost–functional, e.g. the squared tracking error, subject to a set of inequality constraints. The standard controller used in most adaptive schemes is determined from

$$y^*(k) = [1 - A(q^{-1})] y(k) + q^{-d} B(q^{-1}) u(k) \quad (5.21)$$

and minimizes the cost function

$$J(k) = \frac{1}{2} [y(k) - y^*(k)]^2 \quad (5.22)$$

at every instant of time $k > 0$. The resulting input corresponds to a one–step ahead control law, since $J(k)$ is minimized on a one–step basis rather than on the infinite time horizon. If hard constraints are present, the optimization is subject to inequalities representing, e.g. an amplitude limitation u_0 of the actuator.

$$|u(k)| \leq u_0 \quad (5.23)$$

If the minimization of (5.22) is carried out subject to (5.23) it is not clear whether the resulting control input stabilizes the system. A resort is to convert (5.23) into a soft constraint and solve an unconstrained optimization problem. In (Goodwin, 1984 [27]), such a procedure has been adopted using the quadratic cost–function

$$J'(k) = \frac{1}{2} \{ [y(k) - y^*(k)]^2 + \gamma u(k-d)^2 \} \quad (5.24)$$

The penalty γ on the control effort is increased until the amplitude constraint $u(k-d)^2 < u_0^2$ is satisfied. Conditions on the stability of the *optimized* closed–loop system can be stated in terms of the system polynomials $A(q^{-1})$ and $B(q^{-1})$ and the weight on the control effort γ . The implicit assumption made in this case is that the optimal control law ensures that the constraint is always satisfied. Hence, stability is analyzed in the absence of the saturation nonlinearity. However, no guarantee can be given that u remains inside the prescribed bounds for all $k > 0$ as $k \rightarrow \infty$. A crucial assumption is therefore that the system remains stable even if the constraint is not satisfied. This is especially true if the controller is adaptive. The control input in this case is determined on the basis of the estimated parameters. Hence, the control effort varies as the estimates evolve. In course of the estimation process, the

constraints may be violated since the correct value of the weight γ on the control effort cannot be determined ahead of time. Hence, establishing stability of the adaptive system including actuator saturation is essential before any optimization techniques can be applied.

An early effort to deal with the effect of saturation in adaptive systems is reported in (Monopoli, 1975 [47]) where the concept of the augmented error is used to represent the effect of an excessive control signal (i.e. with magnitude greater than a prescribed value). A more recent result is presented in (Annaswamy and Kárason, 1995 [6]). The authors show that the control input generated by a standard adaptive one-step ahead controller in conjunction with a magnitude-limiting saturation nonlinearity results in global stability if the open-loop plant is stable. Moreover, if the open-loop plant (in discrete time) has at least one pole on or outside the unit circle, the adaptive system has bounded solutions provided that the parameters as well as the system state is initialized within a compact set. It is shown that if the initial conditions exceed certain values, the adaptive system may become unstable with saturated inputs if the open-loop system is unstable.

In the following, we assume that the open-loop system is stable. Our aim is to improve the performance of the adaptive controller by optimizing the system with respect to the control effort using a similar cost function as in (5.24). However, the drawback of the procedure in [27] is that optimality holds only asymptotically, i.e. once the parameters converge. This is because the optimal controller structure is derived assuming that the parameters are known. If the parameters are unknown, the most one can expect is that the control input u approaches the optimal control input u^* asymptotically (which indeed has been shown in [27]). The procedure –which is intrinsically linear– does not take into account the cost that is due to the adaptation process. Hence, an excessive control effort due to an unfavorable evolution of the parameter estimates $\hat{\theta}(k)$ cannot be prevented. The procedure is limited to deriving an optimal controller structure. As seen in the preceding section, such a structure can also be found through a proper choice of the reference model. In the following, we assume that the optimal structure has been determined, i.e. at θ_0 , all constraints are satisfied. We address the question whether the adaptive process can be optimized such as to cause minimal control effort. Given $\hat{\theta}(k-1)$, we determine the most favorable $\hat{\theta}(k)$ in the sense that it minimizes the cost function:

$$J[\eta, r(k)] = \gamma_1 u(k)^2 + \gamma_2 \Delta u(k)^2 \quad (5.25)$$

where $\Delta u(k) = |u(k) - u(k-1)|$ measures the rate of change of the input signal and $\gamma_1, \gamma_2 \in \mathbb{R}$ are weighting factors. The novelty here is that the optimization is carried out at every instant of time using the adaptive gain η (not the control input u) as an optimization variable. Given a parameter estimate $\hat{\theta}(k-1)$, this results in an optimal choice of the update $\hat{\theta}(k)$,

i.e. the control law determined from $\hat{\theta}(k)$ minimizes the above performance criterion.

5.2.1 Optimization of the Adaptive Gain

The idea of optimizing the adaptive gain η is motivated by the observation that the control input generated in the adaptive loop is large and oscillating primarily during the initial phase of parameter convergence when the amount by which the parameters are adjusted is large. The system is strongly nonlinear and none of the approximate methods (e.g. averaging), which assume quasi-linearity, apply. In view of this practical and analytical difficulty, an obvious choice would be to reduce the speed of convergence, i.e. to set the adaptive gain η close to zero. However, the improvement was soon found to be inconsistent, since in some cases, the parameters appearing in the denominator of the control law were close to zero, and a large adaptive gain would have proven advantageous to decrease the magnitude of the control input.

It is clear that without knowing the system it is impossible to predict the kind of control input required to meet a given control objective. Moreover, if the parameters of the controller are time-varying, and no assumptions on the nature of the time-variation are made, straightforward tools for the analysis of the system do not exist. Undoubtedly, adaptation is a nonlinear process and offers little insight into the signals generated in closed loop. On the other hand, notice that the overall adaptive procedure consists of two components which both result from simple optimization problems. First, the controller itself which minimizes a criterion defined in terms of the tracking error, second, the parameter estimator which –in the case of the projection algorithm presented in chapter 2.1– minimizes the cost function

$$J_P = \frac{1}{2} \|\hat{\theta}(k) - \hat{\theta}(k-1)\|^2 \quad (5.26)$$

subject to

$$y(k+d-1) = \phi(k-1)^T \hat{\theta}(k) \quad (5.27)$$

where d denotes the relative degree of the system. The minimization of J_P in equation (5.26) results in

$$\hat{\theta}(k) = \hat{\theta}(k-1) + \eta(k) \frac{\phi(k-1) \epsilon(k+d-1)}{1 + \phi(k-1)^T \phi(k-1)} \quad (5.28)$$

where usually $\eta(k) \equiv 1$ and $\epsilon(k)$ is the identification error defined in (5.15). The optimization is performed on a one-step ahead basis, i.e. based on the current estimate the new estimate is determined such that it is closest to the current one while satisfying the system equation (5.27). This geometric viewpoint does offer some insight into the nature of the sequence $\{\hat{\theta}_k\}$ generated by the estimator. However, it does not yet relate to the performance of the

adaptive system, in particular to the control input which results if the estimates are used in the feedback loop as parameters of the controller. Such a relationship will be established in the following paragraph.

The control input at instant of time k is defined implicitly by equation (5.16) or, equivalently,

$$C y^*(k+d) = \phi(k)^T \hat{\theta}(k) \quad (5.29)$$

and is based on the most recent estimate $\hat{\theta}(k)$. Using the definition of the reference system (5.2) we obtain

$$q^{d-d^*} D r(k) = \hat{\theta}_1(k) u(k) + \hat{\theta}_2(k)^T \bar{\phi}(k) \quad (5.30)$$

where the vectors ϕ and $\hat{\theta}$ have been partitioned as follows: $\phi(k) = [u(k) \ \bar{\phi}(k)^T]^T$ and $\hat{\theta}(k) = [\hat{\theta}_1(k) \ \hat{\theta}_2(k)^T]^T$. Substituting $\hat{\theta}_1(k)$ and $\hat{\theta}_2(k)$ in (5.30) by the right hand side of equation (5.28) and solving for $u(k)$ yields

$$u(k) = \frac{\alpha_0(k) - \eta(k) \alpha_1(k)}{\beta_0(k) + \eta(k) \beta_1(k)} \quad (5.31)$$


where the scalars $\alpha_0, \alpha_1, \beta_0, \beta_1$ are given by:

$$\begin{aligned} \alpha_0(k) &= q^{d-d^*} D r(k) - \hat{\theta}_2(k-1)^T \bar{\phi}(k) \\ \alpha_1(k) &= \frac{\epsilon(k+d-1) \bar{\phi}(k-1)^T \bar{\phi}(k)}{1 + \phi(k-1)^T \phi(k-1)} \\ \beta_0(k) &= \hat{\theta}_1(k-1) \\ \beta_1(k) &= \frac{\epsilon(k+d-1) u(k-1)}{1 + \phi(k-1)^T \phi(k-1)} \end{aligned} \quad (5.32)$$

This means that $u(k)$ is a function of the adaptive gain $\eta(k)$. If $\eta(k) \neq 1$ the resulting parameter update $\hat{\theta}(k)$ does not correspond to a minimum of J_P defined in (5.26). However, from the arguments provided in chapter 2.1.1 it is clear that η may be chosen arbitrarily within the interval $0 < \eta < 2$ without impairing the convergence properties of the algorithm. In figure (5.5), the allowed range of η is illustrated for a system with two unknown parameters θ_1 and θ_2 . It is immediately clear that the parameter error $\|\tilde{\theta}(k-1)\|$ does not increase as long as $0 \leq \eta \leq 2$ thereby satisfying elementary property (i) in chapter 2.1.1. Since this is the property upon which the proof of stability in chapter 3.1.4 is based, the adaptive gain may be any number $\eta \in (0; 2)$. This additional degree of freedom can be utilized to improve the signal properties of the control input. We formulate the following static optimization problem:

Given $\hat{\theta}(k-1)$ and the regression vector $\phi(k-1)$ as well as $r(k)$ and $\bar{\phi}(k)$ defined in (5.30), determine $\eta^*(k) \in (0; 2)$ such that

$$J[\eta^*(k)] = \gamma_1 u(k)^2 + \gamma_2 \Delta u(k)^2 \rightarrow \min \quad (5.33)$$

Figure 5.5: Allowed values of η using a projection algorithm

Example 5.3 Given a system of the form

$$y(k) = a_0 y(k-1) + b_0 u(k-1) \quad (5.34)$$

with unknown parameters a_0, b_0 . For simplicity, assume that the reference model reads:

$$y^*(k) = r(k-1) \quad (5.35)$$

Let $k = 1$. The estimates of a_0 and b_0 are initialized at $\hat{\theta}_1(0) = \hat{\theta}_2(0) = 0$. Hence, the estimation error satisfies $\epsilon(1) = y(1)$. We assume that $\phi(0) = [u(0) \ y(0)]^T$ is an arbitrary vector satisfying $\|\phi(0)\| = 1$. Further, let $r(0) = 1$. Then, $\alpha_0(1) = 1$, $\alpha_1(1) = \frac{1}{2} y(1)^2 y(0)$, $\beta_0(1) = 0$ and $\beta_1(1) = \frac{1}{2} y(1) y(0)$ in the above definition (5.32). The control input is given by

$$u(1) = \frac{1 - \frac{\eta}{2} y(1)^2 y(0)}{\eta y(1) y(0)} \quad (5.36)$$

The cost $J[\eta, r(0) = 1]$ generated by $u(1)$ is visualized by rotating the vector $\phi(0)$ (where $\|\phi(0)\| = 1$) around the origin in the $\hat{\theta}_1$ - $\hat{\theta}_2$ plane. For every given $\phi(0)$, the adaptive gain η assumes 100 evenly spaced values within the interval $[0; 5]$. The orientation of $\phi(0)$ determines the direction, and η the amount by which the parameters are updated, according to the equation

$$\hat{\theta}(1) = \frac{\eta}{2} \phi(0) y(1) \quad (5.37)$$

The cost associated with $\hat{\theta}(1)$ is displayed in figure (5.6) as a function of $\hat{\theta}_1(1)$ and $\hat{\theta}_2(1)$. The light shaded area corresponds to those $\hat{\theta}(1)$ which are obtained using η within the (allowed) range $(0; 2)$. All other estimates are marked as the dark area. From the figure it is clear that for any regression vector $\phi(0)$ there exists a favorable choice of parameter estimates

which minimizes $J[\eta, r(0) = 1]$. However, if $\phi(0)$ is parallel to the $\hat{\theta}_1$ -axis, then $\hat{\theta}_2(1) = 0$ or, equivalently, the denominator of (5.36) is zero. As expected, $J[\eta, r(0) = 1] \rightarrow \infty$ in this case, irrespective of the value of η . \square




Figure 5.6: Cost due to input magnitude and rate as a function of $\hat{\theta}(1)$

As seen in the example, J has a vertical asymptote at $\eta_A = -\beta_0(k)/\beta_1(k)$

$$\begin{aligned} \eta \rightarrow \eta_A^- : J &\rightarrow +\infty \\ \eta \rightarrow \eta_A^+ : J &\rightarrow +\infty \end{aligned} \quad (5.38)$$


Furthermore, a straightforward calculation in which we use (5.31) to substitute $u(k)$ in equation (5.33) shows that

$$\left. \frac{dJ}{d\eta} \right|_{\eta_{opt}} = 0$$

at

$$\eta_{opt}(k) = \frac{(\gamma_1 + \gamma_2)\alpha_0(k) - \gamma_2 \beta_0(k) u(k-1)}{(\gamma_1 + \gamma_2)\alpha_1(k) + \gamma_2 \beta_1(k) u(k-1)} \quad (5.39)$$

which corresponds to a global minimum since $J[\eta_{opt}] = 0$ and $J[\eta] > 0$ for all $\eta \in \mathbb{R} \setminus \{\eta_{opt}\}$. However, η_{opt} may not satisfy the inequality constraint. In order to determine the constrained

Figure 5.7: Constrained (*) and unconstrained (o) minimum of J

optimal value $\eta^* \in (0; 2)$, we distinguish four cases which are illustrated in figure (5.7). The graphs displayed correspond to a cross section through the above "cost landscape" along a given direction $\phi(0)$. The case number can be found in the top left part of each subplot. Cases 1 and 2 correspond to the situation where the vertical asymptote at η_A and the unconstrained minimum η_{opt} lie outside the interval $I = (0; 2)$. Define the compact interval $I \supset \bar{I} = [0.1; 1.9]$. The best choice for η^* is to be equal to the upper (lower) bound of \bar{I} , respectively. In case 3, we still have $\eta_A \notin I$ but η_{opt} lies on the same side of the vertical asymptote as I . Since J is convex on (η_A, ∞) , $\eta^* = \max\{0.1, \min(\eta_{opt}, 1.9)\}$. Case 4 corresponds to the (rare) case that $\eta_A \in I$. If $\eta_{opt} \notin I$ we set $\eta^* = \operatorname{argmin}\{J[0.1], J[1.9]\}$ otherwise $\eta^* = \eta_{opt}$. The optimal value η^* is indicated in each case at the position of the constrained minimum (*).

The optimization is carried out at instant of time k , using the most recent output measurement $y(k)$ which is contained in $\bar{\phi}(k)$ of equation (5.32). At the same time, $u(k)$ is available:

$$u(k) = \frac{\alpha_0(k) - \eta^*(k) \alpha_1(k)}{\beta_0(k) + \eta^*(k) \beta_1(k)} \quad (5.40)$$

Proposition: The adaptive controller defined in (5.40) is stable and optimal with respect to the instantaneous cost function (5.33).

Proof: Clear from the properties of the parameter estimation algorithm and the fact that

$\eta^*(k) \in \bar{I} \subset I = (0; 2)$ for all $k > 0$. ■

5.2.2 Discussion

At every instant of time k , $k > 0$, a static optimization problem is solved in order to determine $u(k)$. Once the parameters converge, the control effort is independent of η . Notice that at every instant of time an optimization problem is solved without awareness of the cost-to-go. In particular, the way in which the current control input $u(k)$ affects the future evolution of the cost is not accounted for in the optimization procedure. To obtain an exact solution, the objective should be to minimize the control effort over the infinite time-horizon rather than on a one-step basis. This is a dynamic optimization problem which consists in determining the optimal trajectory $\bar{\eta}^*(k)$, $k \in [0, \infty)$ which causes $\hat{\theta}(k)$ to converge to θ_0 while minimizing the effort of the associated adaptive tracking controller. However, it is seen that the cost-to-go depends upon θ_0 and the future regression vectors $\phi(k)$ which are both unavailable at instant of time k , since the system is unknown. We have to ask the question whether the sequence $\{\eta_k^*\}$ obtained by static optimization is a valid approximation of the optimal trajectory $\bar{\eta}^*(k)$. While this is a hard question the following reflection may be helpful: The critical point is that the cost may become large at some future instant $k + N$, $N > 0$ due to an unfavorable choice of $\eta(k)$ at instant k . Suppose the system is unstable, then a large input may be called for to prevent the output of the system from growing. In such a case, it is clearly futile to minimize the control effort at instant of time k . For an even larger control input will be needed at a later stage, when the state variables of the system have already grown in magnitude. The solution in this case would be not to optimize the adaptive gain until the parameters have moved into a stable region in parameter space. The problem is that there is no way to decide this at the current instant k since the large cost arises only at a later stage $k + N$. In addition, the parameters can only be updated using the information contained in the regression vector and it is obviously not clear where the stable regions in parameter space are. To avoid such difficulties, we assume that the plant is open-loop stable, i.e. there exist positive constants n_1, n_2 such that,

$$|y(k)| \leq n_1 + n_2 \max_{1 \leq \kappa \leq k-d} |u(\kappa)| \quad (5.41)$$

Further, it is reasonable to assume that there exists a class \mathcal{K} -function¹ $\alpha(\cdot)$ such that $\alpha(|y(k)|) \leq |u(k)|$. This general inequality refers to *any* stabilizing controller and states that the magnitude of $u(k)$ required to keep the solutions bounded is at least as large as a quantity

¹A function $\alpha : [0, \infty) \rightarrow [0, \infty)$ is said to be of class \mathcal{K} if it is continuous, strictly increasing and $\alpha(0) = 0$, see e.g. (Sontag, 1998 [73]).

that increases with the magnitude of the output $y(k)$. As an example, the linear stabilizer $u = -ay$ (for relative-degree-one systems and a sufficiently large) satisfies the inequality. The lower bound on $|u|$ is critical since we are discussing the possibility of a future increase of the cost due to an unfavorable choice of the current adaptive gain $\eta(k)$. But the lower bound $\alpha(|y(k)|)$ on the magnitude of u is non-growing if the open-loop plant is stable. Actually, it is bounded by the input u itself as seen in (5.41). Hence, the minimization of the current $u(k)$ does not entail a future increase of the magnitude of the control input required to stabilize the system. But this means that the one-step optimization results in a trajectory which is arbitrarily close to the optimal trajectory (determined only in hindsight). The condition was that the plant is open-loop stable. Note that most practical systems are stable to begin with, and if they are not, an underlying non-adaptive controller can be used to stabilize them. Extensive computer simulations have consistently shown that the optimization procedure results in a substantial reduction of the control effort.

Example 5.4 We assume that the system is given by equation (5.34) with unknown parameters a_0 , $|a_0| \leq 1$ and b_0 and the reference model is defined by

$$[1 + c_1 q^{-1}] y^*(k) = q^{-1} r(k) \quad (5.42)$$

with input $r(k) = 10 \sin(22\pi k T_S)$. The control law obtained from (5.31) is given by:

$$u(k) = \frac{[r(k) - \hat{\theta}_2(k-1)][1 + y(k-1)^2 + u(k-1)^2] - \eta(k) [y(k-1)\epsilon(k) y(k)]}{\hat{\theta}_1(k-1) [1 + y(k-1)^2 + u(k-1)^2] + \eta(k) [u(k-1)\epsilon(k)]} \quad (5.43)$$

The objective is to determine $\eta(k)$ such that the control effort $J(k) = 0.1u(k)^2 + 0.9\Delta u(k)^2$ is minimized at every instant of time. This would result in an optimal evolution of the parameter estimates, while at the same time, the control error tends to zero. On the right-hand side of figure (5.8), a constant gain was used $\eta(k) \equiv 1$. The output y approaches the output y^* of the reference system and the parameters are seen to converge to the correct values $\hat{\theta}_1 \rightarrow b_0 = 2$ and $\hat{\theta}_2 \rightarrow (c_1 - a_0) = -0.5$. However, the control effort spent during the initial transient phase (i.e. before the parameters converge), is significant (row 4 in the figure) and the corresponding input $u(k)$ large and oscillatory. If the optimal gain $\eta^*(k)$ is determined at every instant of time and used in equation (5.43) to compute the control input, the same output performance is achieved with almost zero effort, as displayed in the right column of the figure. It is seen, that parameter convergence is much slower in the optimized case. On one hand, this further emphasizes the fact that high-frequency components of the control input have been removed effectively. For, if the control effort is minimized, the degree of persistency of excitation also decreases which implies that the time constant $1/c$ in equation (2.27) increases. On the other hand, if the designer is interested in determining

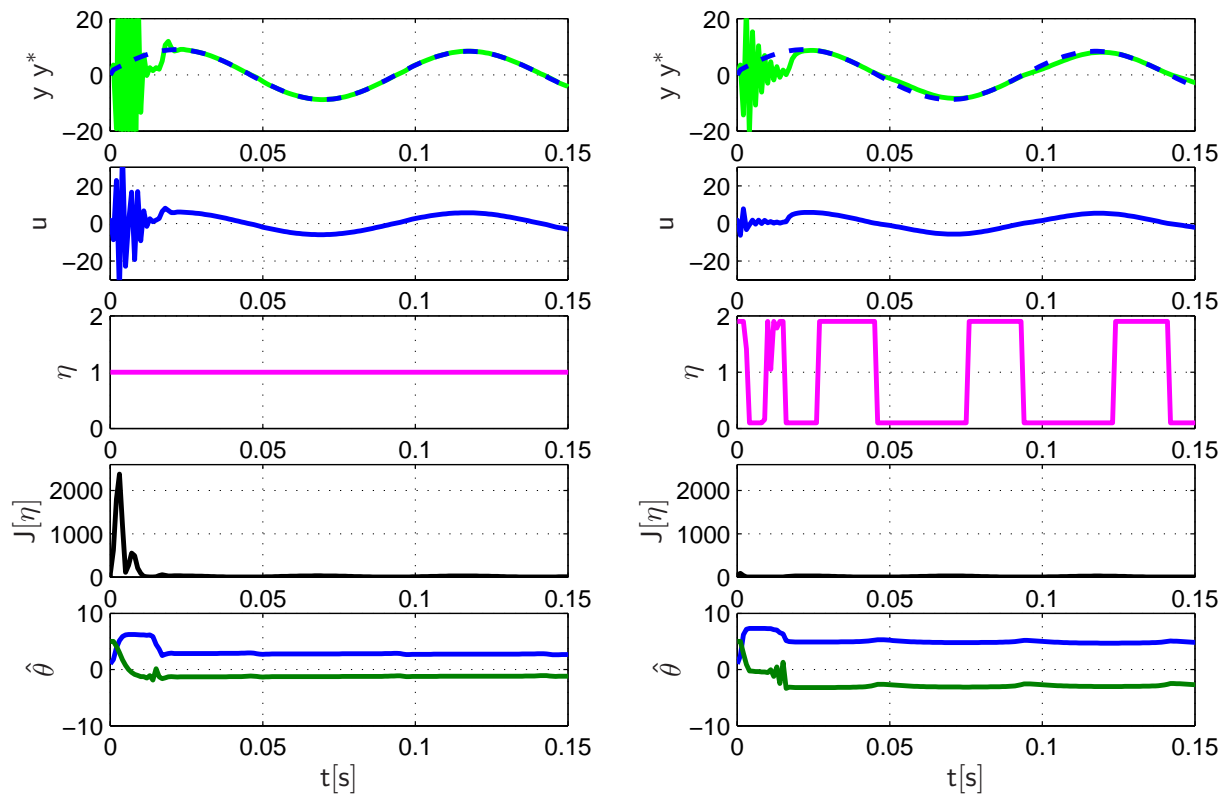


Figure 5.8: Minimizing the control effort of a first-order adaptive system

the actual physical parameters of the system even while controlling the plant adaptively, a large and excitatory control input $u(k)$ will be needed and the weights of the cost function in the optimization procedure have to be chosen appropriately. Hence, there is a trade-off between fast convergence using a rich input $u(k)$ or slow convergence with minimal control effort. This is particularly true in a higher-dimensional parameter space as seen in the next experiment.

Example 5.5 The objective is to control the speed of a two-mass system with unknown parameters $\theta_0 \in \mathbb{R}^6$, as in example (3.2). The control input needed for adaptive control was found to be large and oscillatory. The result is repeated here for convenience, on the left part of figure (5.9). The control effort –measured by means of a cost function of the form (5.33) with $\gamma_1 = 0$ and $\gamma_2 = 1$ – is excessively large ($> 10^{11}$) in the case with constant adaptive gain $\eta(k) \equiv 1$ and vanishes if $\eta(k) = \eta^*(k)$. Since the degree of excitation is high in the former case, the parameters are seen to converge fast and the input $u(k)$ corresponds to that of a linear controller afterwards. It should be kept in mind, though, that in a practical environment we cannot afford to apply inputs of the form displayed the left part of the figure over an extended period of time. Hence, the optimization of the adaptive gain is mandatory

in this case. The same experiment is again repeated in chapter 8 (see figure 8.12) where the optimization procedure is tested under real–world conditions.

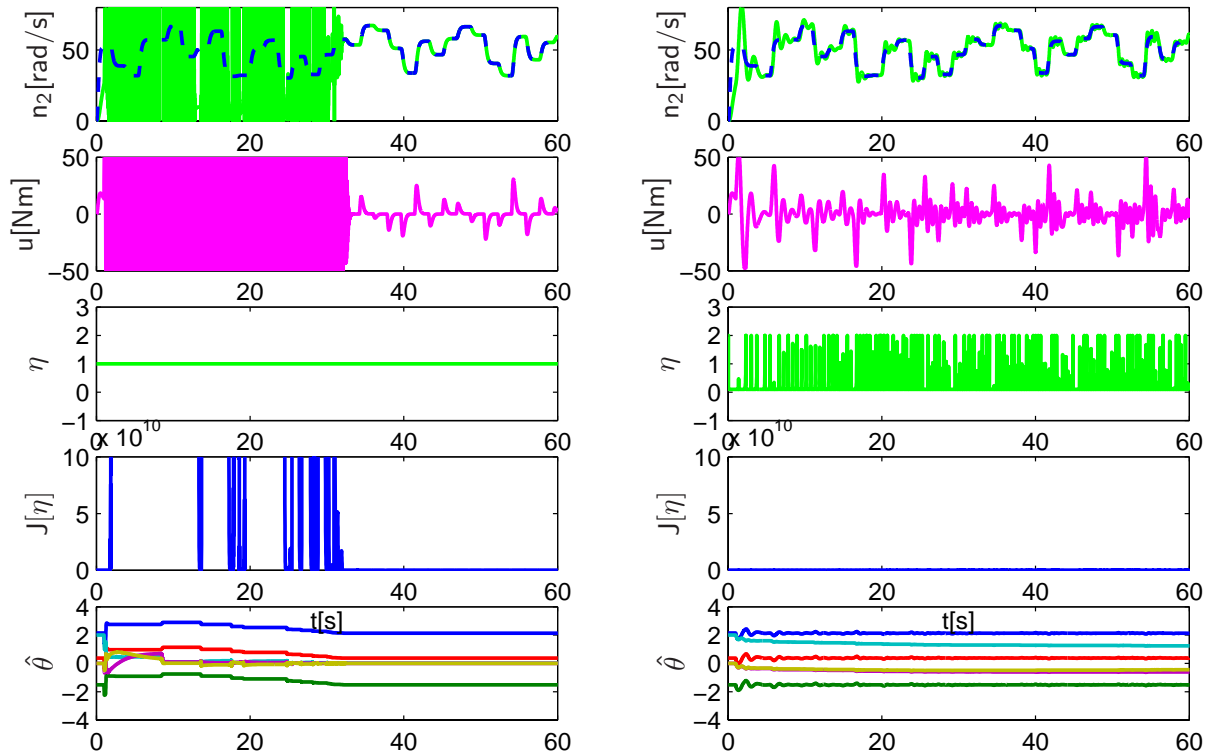


Figure 5.9: Optimal adaptive control of an unknown two–mass system (simulation)

5.3 Unmodelled Dynamics

The control effort spent to achieve a desired objective was at the center of the discussion in the previous section. In this section, we highlight another aspect which is also related to the nature of the control input generated by an adaptive system. It has been recognized very early that adaptive systems are highly sensitive to unmodelled dynamics (see e.g. Ioannou, 1983 [32]; Rohrs, 1985 [69]). If the control input has significant power in a high–frequency range, the presence of parasitic modes may cause the adaptive procedure to become unstable. This non–robustness has its roots in the restrictive assumptions made in the proof of stability, i.e. that the plant order and relative degree be known. Instability phenomena have been investigated using linearization (Rohrs, 1985 [69]), averaging (Åström, 1983 [2]) and singular perturbation techniques (Kokotovic, 1986 [37]) and numerous methods have been developed to prevent signals from growing in an unbounded fashion.

A large body of work is available on general robustness of adaptive systems assuming that the system is subject to bounded disturbances. The basic idea is to eliminate the pure integral action of the adaptive law. In (Egardt 1979, [17]; Peterson 1982, [64]) a dead-zone is introduced which consists in stopping the adaptation when the control error becomes smaller than some prescribed value. The drawback of this approach is that knowledge of an upper bound on the disturbance is needed to implement the dead-zone. A conservative bound leads to large tracking errors, even when the disturbances are small or zero. No prior information is needed in the case of the σ -modification scheme introduced by (Ioannou, 1983 [32]). An additional term $-\sigma\hat{\theta}(k)$ is used in the parameter update law to introduce “leakage”, i.e. convert the pure integral into a first-order system. However, as in the case of the dead-zone, zero residual tracking errors cannot be guaranteed when the disturbances are removed. A much celebrated approach proposed by (Narendra, 1987 [52]) is to replace σ by a term proportional to the output error $|e(k)|$. The rationale for using such a term is that it tends to zero with the output error. Hence, when there is no disturbance, the correction term tends to zero as well. Yet another approach requires the reference input to be persistently exciting of sufficient order. If no disturbances are present, this results in exponential stability of the overall adaptive system (Morgan and Narendra, 1977 [49]). Hence, arguing from linear systems theory, a bounded disturbance would lead to a bounded error. However, care must be taken since the behavior of the nonlinear perturbed system cannot be concluded directly from that of the unperturbed system. As shown in (Narendra, 1983 [53]), the statement is valid if the degree of persistent excitation is greater in some sense than the amplitude of the disturbance.

When unmodelled dynamics are present, global stability cannot be guaranteed by any of the above methods. The unmodelled dynamics acts as a disturbance which can no longer be assumed to be bounded since it results from a structural uncertainty. Even stating the problem is not straightforward since the nature of the disturbance varies significantly according to the kind of parasitics contained in a system. The most commonly encountered practical examples are fast actuator dynamics or high-frequency oscillations of mechanical components.

Any model reflects only the dominant part of a physical system. In other words, the system to which the control is applied is invariably of larger order than what one cares to model. In (Ioannou 1987, [34]), a class of unmodelled dynamics is considered which consists of both additive and multiplicative plant perturbations. The plant is given in the following form

$$G(s) = G_0(s)[1 + \mu\Delta_2] + \mu\Delta_1 \quad (5.44)$$

where $G_0(s)$ is the transfer function of the modelled part of the plant. Δ_1, Δ_2 is an additive

and multiplicative perturbation, respectively and $\mu > 0$ a small number. Robustness is achieved by using a modified adaptive law which employs a normalizing signal and a $-\sigma\theta$ term. Prior information is needed regarding an upper bound on the norm of the controller parameters and a bound on the stability margin of the poles of the unmodelled dynamics. The idea of a normalizing signal also appears in (Kreisselmeier 1986, [41]; Praly 1983 [67]). In addition, the authors propose a projection technique which keeps the parameter estimates bounded within a specified sphere in parameter space. The approach requires bounds on the unknown parameters as well as on the plant zeros. Alternatively, a relative dead-zone can be introduced in order to keep the parameter error bounded.

In the following section, a different viewpoint is taken. Our objective is to demonstrate — using simulations — that an optimal control input, such as the one obtained in the previous section, may be the key to a simple and intuitive solution to the problem. We assume that the effect of unmodelled dynamics can be described by an external disturbance $\xi(k)$ with $|\xi(k)| \leq \xi_0$. The upper bound ξ_0 depends upon the frequency content of the input to the system. The idea is based on the well-known fact that if the latter is low in the high-frequency range, no unmodelled dynamics are excited and, hence, ξ_0 can be made small. Parts of this section have been presented at a conference (Feiler, 2003 [19]).

5.3.1 The Full-Order Case

In the following, the dependency of the polynomials A, B, C, \dots on q^{-1} is omitted for notational convenience. In the ideal case, it is assumed that the order of the system $\deg[A] = n$ is known.

$$A y(k) = q^{-d} B u(k) \quad (5.45)$$

As explained in chapter 5.1, the structure of the controller is determined using Bézout's identity $C = M A + q^{-d} N$. Multiplying both sides of the equation by $y(k)$ and using (5.45) we obtained

$$C y(k) = q^{-d} [M B u(k) + N y(k)] \quad (5.46)$$

When $y(\cdot)$ is replaced by $y^*(\cdot)$ the equation implicitly defines the controller. We introduced the shorthand notation

$$C y(k) = \phi(k-d)^T \theta_0 \quad (5.47)$$

where $\phi(k-d)$ contains past values of $y(\cdot)$ and $u(\cdot)$ and θ_0 the corresponding (unknown) coefficients. θ_0 is replaced by a vector of parameter estimates $\hat{\theta}(\cdot)$ which are updated using one of the estimation algorithms presented in chapter 2. The key quantity for identification

is the error signal $\epsilon(k)$ which contains all the information about the unknown parameter vector θ_0 :

$$\begin{aligned}\epsilon(k) &= C y(k) - \phi(k-d)^T \hat{\theta}(k-1) \\ &= -\phi(k-d)^T \tilde{\theta}(k-1)\end{aligned}\tag{5.48}$$

where $\tilde{\theta} = \hat{\theta} - \theta_0$ denotes the parameter error. The proof of stability, which was presented in chapter 3.1.4, resides on the fundamental property of the parameter estimator, that the identification error ϵ cannot grow faster than the signals of the system contained in the regression vector ϕ . More precisely, recall that if ϵ is defined as in (5.48), we have

$$\lim_{k \rightarrow \infty} \frac{\epsilon(k)^2}{1 + \phi(k-d)^T \phi(k-d)} = 0\tag{5.49}$$

Using (5.49) and the fact that the system is minimum-phase it was shown that the assumption that ϕ grows without bound leads to a contradiction and that the control error tends to zero.

5.3.2 The Reduced-Order Case

In the reduced-order case, only the order $\bar{n} < n$ of the dominant part of the system (and its relative degree d) is assumed to be known.

$$\bar{A} \bar{y}(k) = q^{-d} \bar{B} u(k)\tag{5.50}$$

where $\deg(\bar{A}) = \bar{n} < n$. Notice, that the relative degree d is the same as in (5.45). The output $\bar{y}(\cdot)$ of the reduced-order system is merely of theoretical interest and cannot be measured. Proceeding as in the ideal case, we determine polynomials \bar{M} and \bar{N} such that

$$C = \bar{M} \bar{A} + q^{-d} \bar{N}\tag{5.51}$$

If we had $y(k) = \bar{y}(k)$ (no unmodelled dynamics) we could use $\bar{A} y(k) = q^{-d} \bar{B} u(k)$ from equation (5.50) and write

$$\begin{aligned}C y(k) &= q^{-d} [\bar{M} \bar{B} u(k) + \bar{N} y(k)] \\ &= \phi(k-d)^T \bar{\theta}_0\end{aligned}\tag{5.52}$$

Notice that $\bar{\theta}_0 \in \mathbb{R}^{2\bar{n}}$ is of reduced dimensionality. In our case, $y(k) = q^{-d} \frac{B}{A} u(k) \neq \bar{y}(k)$ and we obtain:

$$\begin{aligned}C y(k) &= q^{-d} [\bar{M} B \bar{A} / A u(k) + \bar{N} y(k)] \\ &= \phi(k-d)^T \bar{\theta}_0 + \xi(k-d)\end{aligned}\tag{5.53}$$

This is the equivalent of equation (5.46) where

$$\xi(k) = \bar{M} \left[\frac{\bar{A} B - \bar{B} A}{A} \right] u(k) \quad (5.54)$$

which quantifies the effect of unmodelled dynamics. As above, $\bar{\theta}_0$ is replaced by an estimation vector $\hat{\theta}$. We obtain the identification error:

$$\begin{aligned} \epsilon(k) &= C y(k) - \phi(k-d)^T \hat{\theta}(k-1) \\ &= \phi(k-d)^T [\bar{\theta}_0 - \hat{\theta}(k-1)] + \xi(k-d) \\ &= -\phi(k-d)^T \tilde{\theta}(k-1) + \xi(k-d) \end{aligned} \quad (5.55)$$

As expected, the identification error is not only due to the estimation error $\phi(k-d)^T \tilde{\theta}(k-1)$ but also due to unmodelled dynamics $\xi(k-d)$. Now, suppose that the same procedure as in the ideal case is used to derive the adaptive controller, i.e. using certainty equivalence, $u(k)$ is determined from

$$C y^*(k) = \phi(k-d)^T \hat{\theta}(k-d) \quad (5.56)$$

where $\hat{\theta}(\cdot) \in \mathbb{R}^{2\bar{n}}$ is the estimation vector which is updated on the basis of the identification error $\epsilon(k)$. The question that needs to be addressed is whether there exists a constant parameter vector θ_0 , for which $\epsilon(k) = 0$. It is clear that no such vector exists in general. However, within a limited frequency range, $|\epsilon(k)|$ can be made small. A second question is whether an equation of the form (5.49) can be established in the reduced-order case. Let us investigate this in the context of the projection algorithm, introduced in chapter 2.1. Property (ii), which was restated in (5.49), was established by showing that the norm of the parameter error $\|\tilde{\theta}(k)\|$ forms a nonincreasing sequence. In the reduced-order case, we have $\tilde{\theta}(k) = \hat{\theta}(k) - \bar{\theta}_0$ and obtain

$$\|\tilde{\theta}(k)\|^2 - \|\tilde{\theta}(k-1)\|^2 = 2\eta \frac{\phi(k-d)^T \tilde{\theta}(k-1) \epsilon(k)}{1 + \phi(k-d)^T \phi(k-d)} + \eta^2 \frac{\phi(k-d)^T \phi(k-d) \epsilon^2(k)}{[1 + \phi(k-d)^T \phi(k-d)]^2} \quad (5.57)$$

The r.h.s is equivalent to

$$\eta \left[\frac{\xi^2(k-d)}{\epsilon^2(k)} - 1 - \frac{[\phi(k-d)^T \tilde{\theta}(k-1)]^2}{\epsilon^2(k)} + \eta \frac{\phi(k-d)^T \phi(k-d)}{1 + \phi(k-d)^T \phi(k-d)} \right] \frac{\epsilon^2(k)}{1 + \phi(k-d)^T \phi(k-d)} \quad (5.58)$$

From the expression it is clear that there exists a class \mathcal{K} function $\sigma(\|\cdot\|)$ such that if

$$\xi(k-d)^2 \leq \sigma(\|\phi(k-d)^T \tilde{\theta}(k-1)\|^2) \quad (5.59)$$

the bracketed expression in (5.58) is always negative. Whenever $\xi(k-d) = 0$, we have that $\epsilon(k) = -\phi(k-d)^T \tilde{\theta}(k-1)$ and we recover equation (2.15). If the inequality holds we have that

$$\|\tilde{\theta}(k)\|^2 = \|\tilde{\theta}(0)\|^2 + \eta \sum_{\kappa=0}^k \rho(k) \frac{\epsilon^2(k)}{1 + \phi(k-d)^T \phi(k-d)}$$

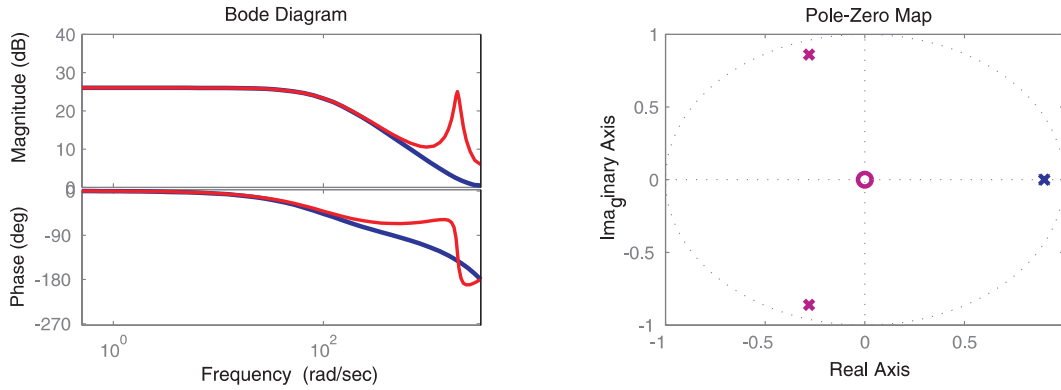


Figure 5.10: Full- and reduced order system

where $\rho(k) < 0$ for all $k > 0$. This means that the normalized identification error tends to zero as in (5.49) using the same argumentation as in chapter 2.1.1, i.e. the fact that $\|\tilde{\theta}(k)\|^2$ is nonnegative. Inequality (5.59) implies that $\xi \rightarrow 0$ when the parameters converge. Going back to the definition of ξ in equation (5.54) we see that $\xi(\cdot) = 0$ whenever $[\bar{A} B - \bar{B} A]$ vanishes. In frequency domain, this corresponds to the fact that the transfer functions $H(z^{-1})$, and $\bar{H}(z^{-1})$ of the full- and the reduced-order system are equal:

$$H(z^{-1}) = \frac{z^{-d}B(z^{-1})}{A(z^{-1})} = \frac{z^{-d}\bar{B}(z^{-1})}{\bar{A}(z^{-1})} = \bar{H}(z^{-1}) \quad (5.60)$$

where $z^{-1} = e^{-j\omega T_s}$ and T_s is the sampling time. The equality holds when no unmodelled dynamics are excited by the input signal. By assumption, ξ is due to high-frequency parasitics. Hence, the frequency responses of $H(e^{-j\omega T_s})$ and $\bar{H}(e^{-j\omega T_s})$ can be matched up to a maximum frequency ω^* . (5.60) is satisfied if the power spectrum of the input signal is zero for all $\omega > \omega^*$.

The following simulation demonstrates that the method developed in section (5.2) provides an effective way of removing high-frequency components of the control input. The idea is to minimize the rate $\Delta u(k)$ of the input signal.

Example 5.6 Consider the third order system

$$[1 - a_1q^{-1} - a_2q^{-2} - a_3q^{-3}]y(k) = q^{-1}b_0 u(k) \quad (5.61)$$

which is to be approximated by the first-order system

$$[1 - \bar{a}q^{-1}]y(k) = q^{-1}\bar{b}u(k) \quad (5.62)$$

The full-order system has a (dominant) pole at $z_1 = 0.9$ and a pair of complex conjugate poles at $z_{2,3} = 0.9 e^{\pm 1.885i}$ corresponding to an eigenfrequency of $f_0 = 300Hz$ of the parasitic

system. The Bode diagram and pole zero map of the corresponding transfer-functions are displayed in figure (5.10). The reference system is selected to be:

$$[1 - c_1 q^{-1}] y^*(k) = q^{-1} d_0 r(k) \quad (5.63)$$

where $c_1 = 0.3$ and $r(k) = 10 \sin(14\pi k T_S) + 5 \sin(18\pi k T_S)$ with sampling time $T_S = 1ms$.

The controller will be determined from the first-order system and used for the third-order system. The control law reads

$$u(k) = \frac{d_0 r(k) - \hat{\theta}_2(k) y(k)}{\hat{\theta}_1(k)} \quad (5.64)$$

where $\hat{\theta}_{1,2}$ are the unknown controller parameters which are to be adjusted such that $\hat{\theta}_1 \rightarrow \bar{b}$ and $\hat{\theta}_2 \rightarrow [\bar{a} - c_1]$.




Figure 5.11: Linear and adaptive controller in the presence of unmodelled dynamics

The simulations contained in figure (5.11) reveal an interesting phenomenon: In both cases, the parameters have been initialized at the correct values $\hat{\theta}(0) = [\bar{b} \ (\bar{a} - c_1)]^T$. On the left hand side, no adaptation is allowed, i.e. the parameters are fixed at their initial values. It is seen that the presence of the parasitic poles causes a deviation $|y - y^*| \ll 1$ but the overall system is stable. This corresponds to our expectations since, with fixed parameters, the system is linear. If we allow the adjustment of the parameters (right part of the figure), the system behaves well initially but becomes unstable due to a drift of the parameters. The drift is caused by an error due to unmodelled dynamics which becomes large whenever the control input $u(\cdot)$ is oscillatory and contains frequencies which excite the parasitic system.

The effect can be suppressed by optimizing the adaptive gain with respect to the rate of the input signal $\Delta u(k)$. The cost function J in equation (5.33) is minimized at every instant of time using $\gamma_1 = 0.1$ and $\gamma_2 = 0.9$. The resulting time-varying adaptive gain $0.1 \leq \eta^* \leq 1.9$ is used to update the parameter estimates $\hat{\theta}(k) \in \mathbb{R}^2$ upon which the control law (5.64) is based.

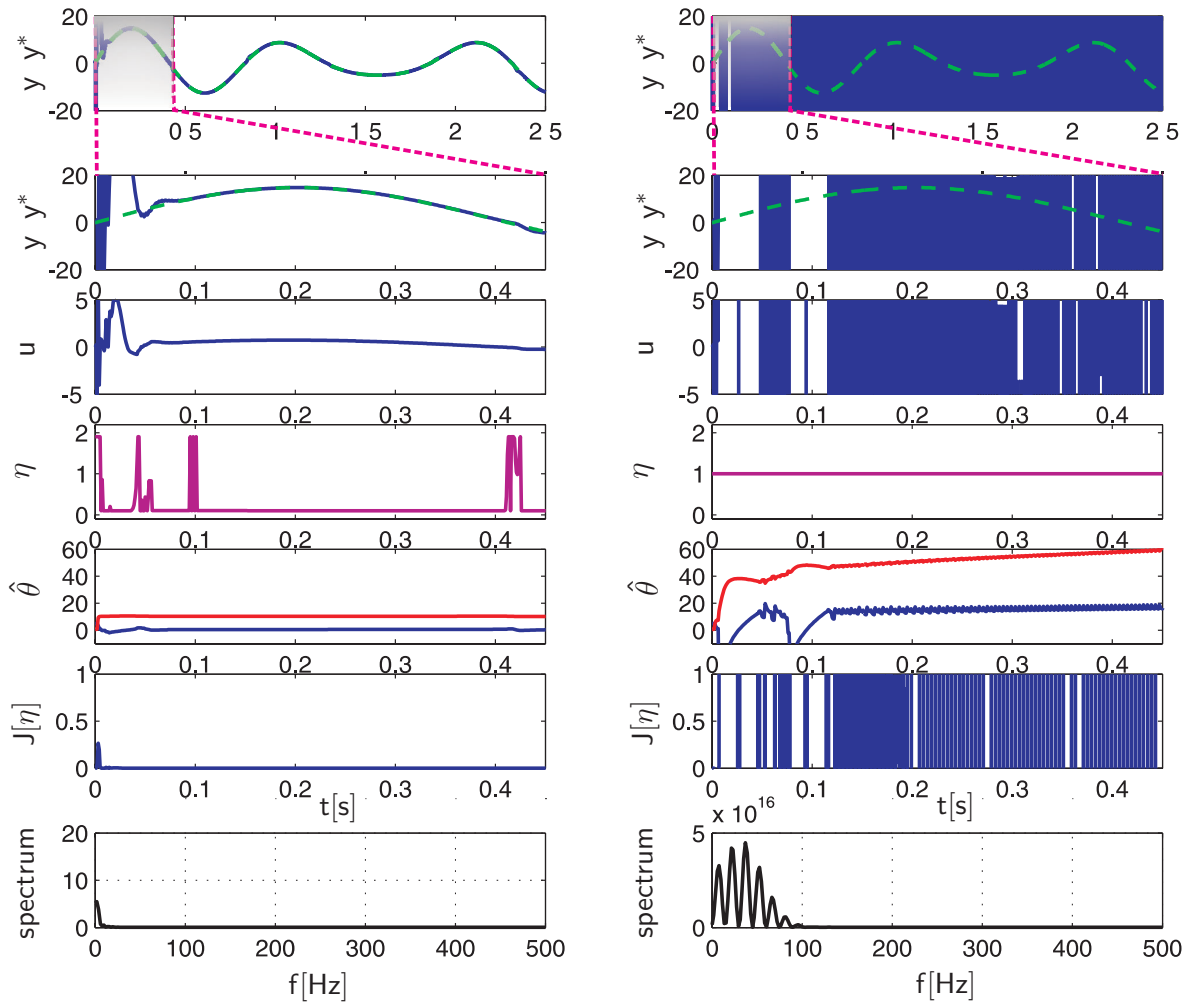


Figure 5.12: Adaptive control in the presence of unmodelled dynamics using time-varying (left) and constant (right) adaptive gain η

In figure (5.12), the first 450 milliseconds of the simulation are displayed in a zoomed section (rows 2-7). The tracking performance is excellent in the case with time-varying adaptive gain whereas signals grow in an unbounded fashion if a constant adaptive gain is used. In both cases, the parameter estimate $\hat{\theta}(0) = [5 \ 0.5]$ was initialized far from the true values. While in the former, the optimization prevents $u(k)$ from becoming large, in the latter case $u(k)$ is found to be large and oscillatory. As a consequence, the contribution of the unmodelled

dynamics to the identification error is large, causing the parameter estimates to drift (without bound). Compare row 5 on the left and right side of the figure to appreciate the effect of the optimization on the evolution of the parameters. Rows 6 and 7 of the figure display the cost and the power spectrum associated with the control input. If a time-varying gain is used the resulting input hardly produces any cost and has nonzero power only at frequencies corresponding to the reference input $r(k)$. With constant gain $\eta = 1$, the input has significant power at all frequencies up to Nyquist frequency. \square

In this chapter it has been demonstrated that the performance of adaptive schemes can be improved through an appropriate choice of the controller structure and the adaptive gain. The emphasis was placed on obtaining a smooth control input which satisfies magnitude and rate constraints imposed by the actuators of the system. An interesting connection to robustness in the presence of unmodelled dynamics was pointed out in the last section. The stability properties of the adaptive system are the same in the regular and the optimized case. However, even without unmodelled dynamics, the optimized version results in better performance in the sense that the same tracking response can be achieved with less control effort. The point to note is that it is the *improved performance* which helps us to obtain *stability* when the system is subject to structural disturbances. Future work is required to derive a quantitative measure of the robustness obtained through the optimization of the control input. In particular,

- How an upper bound ω^* on the frequency content of the control input translates into a rate constraint in time-domain (assuming that the input is a continuous function of the system states and the parameters, i.e. no switching is allowed) has to be investigated.
- The optimization has to be combined with one of the well-known robustification techniques, such as signal normalization and relative dead-zone. Even though the introduction of a dead-zone was not found to be necessary in the above example and the experiments carried out in chapter 8, it becomes indispensable when deriving stability guarantees in the presence of unmodelled dynamics. This is because the cost function (5.33) formulates only a soft constraint, i.e. it is not known a priori *to what extent* high-frequency components of the control input are attenuated.

Chapter 6

Rapidly Time–Varying Systems

The third class of disturbances considered in this thesis is due to large and sudden changes in the parameters of the system. In fact, one of the primary reasons for considering adaptive control in practical applications is to compensate for large variations of the plant parameters. Despite this fact, a coherent theory of adaptive control exists only when the unknown system is time–invariant. The accepted philosophy is that if an adaptive system is fast and accurate when the plant parameters are constant but unknown, it would also prove satisfactory when the parameters vary with time, provided the latter occurred on a relatively slower time–scale.

The analytical difficulties encountered while dealing with unknown, time–varying systems are due to the following factors. If $\theta : \mathbb{N}_0 \rightarrow \mathbb{R}^{2n}$ is the (time–dependent) control parameter vector which is used to compensate for variations in plant parameters, it must first be shown that a vector function $\theta^* : \mathbb{N}_0 \rightarrow \mathbb{R}^{2n}$ exists such that the control error is zero when $\theta(k) \equiv \theta^*(k)$. In practice, this is not easy to accomplish since the analysis of the resulting system involves time–varying operators. Even when $\theta^*(k)$ has been shown to exist, the derivation of a stable adaptive law is difficult since the error equations involved are nonlinear and non–autonomous. Using the parameter error $\tilde{\theta}(k) = [\hat{\theta}(k) - \theta^*(k)]$ one may attempt to apply a standard adaptive algorithm to adjust $\hat{\theta}(k)$. Due to the time–variation of θ^* , an additional term $\Delta\theta^*(k) = [\theta^*(k) - \theta^*(k-1)]$ appears on the r.h.s of the update equation which makes the system non–autonomous. Moreover, convergence has to be proven (in a pointwise fashion) in function space.

Results have been obtained in the continuous–time case under the assumption that the unknown plant parameter vector is the output of an asymptotically stable linear time–invariant system with constant input. In this case, stability was established without modifications of the adaptive law (see Narendra, 1989 [51]). Simulation studies have been carried out for

the case where the dynamics of the plant parameters is governed by a linear second-order system. When the frequency of the parameter variation is low, a standard adaptive law was seen to result in a small output error. For general time-variations, a robust adaptive control approach has been proposed by (Tsakalis, 1987 [74]). Discrete-time results have been obtained by (Kreisselmeier, 1982 [40]) and (Narendra, 1987 [58]). In the latter case, a bound on the parameter $\theta^*(k)$ and its variation $\Delta\theta^*(k)$ was established for which the standard adaptive law ensures stability.

All the above results are based on the assumption that the parameter variation is slow in some sense. If this is not satisfied, but the parameter perturbations are small in magnitude, robust control is generally preferred. However, as systems become more complex, situations where the parameters vary both rapidly and by large amounts arise with increasing frequency. Typical examples are mechanical processes with large variations in load, actuator failures or transition control tasks in chemical systems. It was to cope with such situations that a methodology based multiple models has been introduced by (Narendra and Balakrishnan, 1992 [54]). The general methodology as well as the principal results are summarized below (see also Narendra et al., 2003 [23]).

6.1 Multi-Model Adaptive Control

From the very beginning, the interest of adaptive control theorists was centered around adaptation in changing environments. However, due to mathematical tractability, they confined their attention to time-invariant systems with unknown parameters. The control algorithms that have proven stability properties are characterized by an adaptive law which is given in the form of a differential equation. Once the evolution of the parameters towards their true values (or –more generally– towards a solution manifold) is governed by a differential equation, the long-term behavior of the adaptive system can be studied in terms of a dynamical system and becomes amenable to powerful tools from the qualitative theory of nonlinear differential equations (e.g. convergence of the parameters is studied in terms of the stability of an associated dynamical system). It is clear, though, that this may not be the fastest way to determine the parameters since adaptation thus defined is inherently incremental. Note that when the unknown parameters appear nonlinearly, it may even be impossible to find a differential equation that describes the evolution of the parameter estimates towards their true values. In these cases, the parameters are determined by optimizing an error function, e.g. $\hat{\theta} = \operatorname{argmin}\{\epsilon(\hat{\theta})^2\}$, and the resulting trajectory in parameter space is generally a non-smooth curve. It is in the same spirit that one may wish to allow for switching while

adapting the parameters of an unknown system.

The advantage of switching would be to provide fast response to time-variations occurring in the system. Even, if no such variations take place, the tracking error is quite often unacceptably large and oscillatory during the initial phase of parameter convergence, when there are large errors in the estimates. This observation provided the motivation for developing an optimized version of the adaptive scheme in chapter 5.2. In this chapter we present a methodology, which –in addition to providing improved response in the case of time-invariant parameters– opens up new ways of controlling systems with large and sudden variations in the parameters.

Since the method was originally proposed in 1992, many other approaches have been reported in the adaptive control literature. However, only a few have been demonstrated to be stable. Since it is well known that efficient design methods for different classes of control systems can be developed only when their stability properties are well understood, we focus our attention on stability issues in multi-model based adaptive control. In particular, we will be interested in the assumptions that have to be made to assure stability. In a time-varying environment, the effectiveness of the approach depends critically on the location of the models. How adaptation and learning are to be combined to accomplish the latter efficiently, will be at the center of our discussion in chapter 7.

6.1.1 General Methodology

In this section the motivation of the approach as well as its basic concepts are presented. The style is entirely qualitative and sets the stage for the quantitative analysis undertaken in section 6.2. The basic idea for using multiple models comes from biology: every biological system is faced with a multiplicity of choices at any instant of time. As the environment changes, it demonstrates an ability to rapidly modify its strategy such as to maintain optimal performance. Such an ability involves recognizing the specific situation that has arisen and taking an appropriate control action which is selected from a set of available strategies. In order to acquire a repertoire of strategies the biological system learns from past experience. Hence, it is the ability of the system to learn and store information, and combine it with adaptation that is responsible for its performing satisfactorily in rapidly varying situations. It is this feature of biological system behavior that underlies the idea of multi-model adaptive control.

A mathematical description that captures the essential aspects of a physical system is generally referred to as a *model* of that system. It is clear that the kind of models generated for a

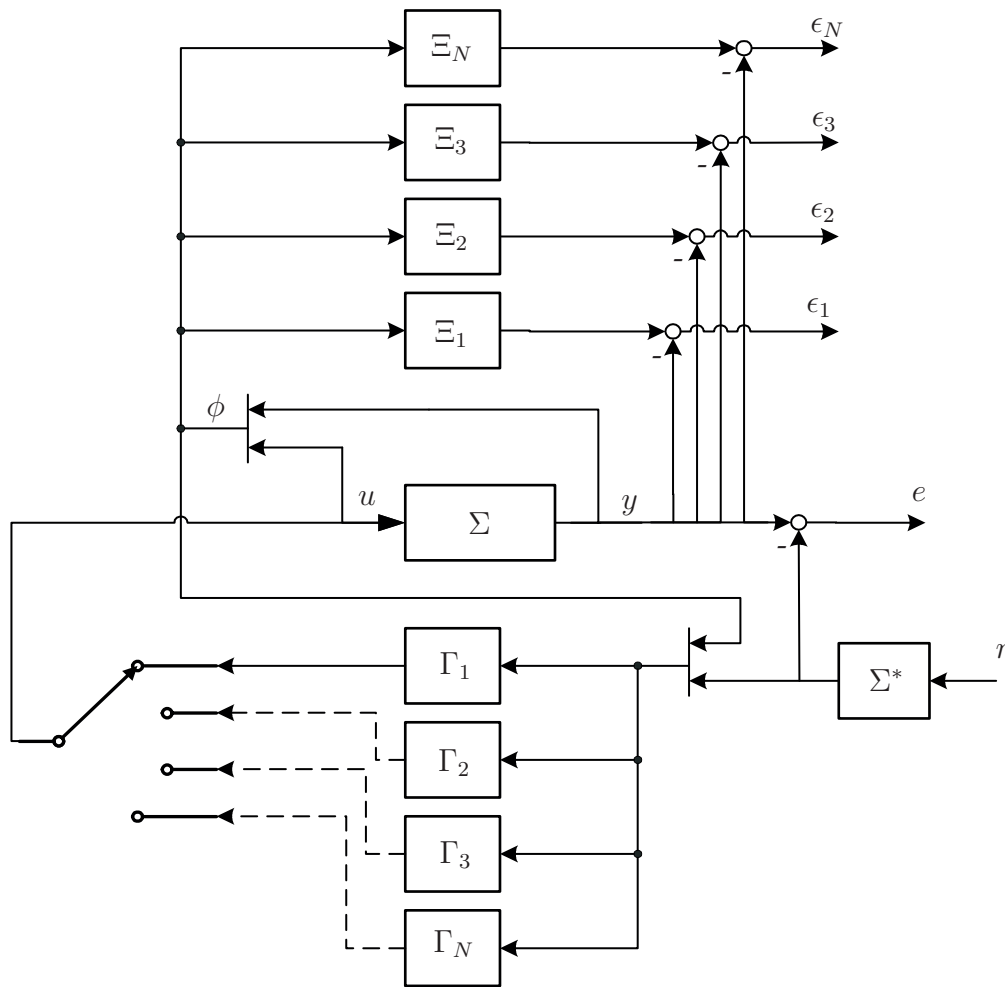


Figure 6.1: Multi-Model Adaptive Control (MMST)

given system may vary significantly according to what the designer *regards to be* its essential characteristics. Since it is reasonable to assume that diverse models may be appropriate for describing different environments, the use of multiple models arises naturally. The basic architecture for adaptive control based on multiple models (MMST) may be described as follows [57]:

The structure of the control system is shown in figure (6.1). The plant Σ to be controlled has an input u and output y . A reference model with piecewise continuous input r provides the desired output y^* , and the objective is to cause the control error $e = y - y^*$ to tend to zero, (or lie within specified bounds) for large values of time. N identification models $\Xi_1, \Xi_2, \dots, \Xi_N$ are used in parallel to estimate the parameters of the plant each of which generates an estimated output \hat{y}_i , $i = 1, 2, \dots, N$. The estimation error of the j th model Ξ_j is defined as $\epsilon_j = y - \hat{y}_j$. Corresponding to each model Ξ_j there exists a controller Γ_j

such that Ξ_j together with Γ_j in the feedback-path behaves like the reference model Σ^* . At every instant, based on a switching criterion, one of the model/controller pairs, e.g. $[\Xi_j, \Gamma_j]$, is chosen. Consequently, the output u_j of the j th controller is used to control the plant. Given prior information about the plant (e.g. that the latter is linear, nonlinear, stochastic, slowly time-varying etc.), the design problem is to choose the models Ξ_j and the controllers Γ_j together with the rules for switching to obtain best performance, and demonstrate that the overall system is stable.

For mathematical convenience, as well as for a precise definition of the control problem, it is assumed that the plant and all the identification models (unless otherwise stated) can be parameterized in the same fashion. If the unknown plant parameter is a vector θ_0 and the estimates of θ_0 given by the models are $\hat{\theta}_i$, we assume that θ_0 and $\hat{\theta}_i (i = 1, 2, \dots, N) \in S \subset \mathbb{R}^p$ where S is a compact set. For ease of exposition, we shall refer to θ_0 as the plant and $\hat{\theta}_i$ as a model. The general problem of MMST can also be considered as one of choosing $\hat{\theta}_i (i = 1, 2, \dots, N)$ so that for any $\theta_0 \in S$ the control objectives can be achieved and the overall system is stable. In the paragraphs that follow, further details concerning the models, the switching criterion, and the manner in which the control input is to be computed are discussed.

6.1.2 Models

As stated in the introduction, a large variety of models can be chosen including continuous-time or discrete-time, linear and nonlinear models. If the plant is subjected to disturbances, a model of extended order may be called for. Note that if models with feedback are used in the estimation procedure, there is the possibility that the models become unstable for some values of their parameters. In such cases, even when the controlled plant is stable and has a bounded output, the outputs of the models may grow in an unbounded fashion. To avoid this, the models are chosen to be of the *series-parallel type* [51], i.e. the regression vector contains measured input-output values of the plant and is the same for all the models.

Models with constant parameters $\hat{\theta}_i$ are referred to as fixed models while those which are continuously updated based on input-output data are referred to as adaptive models. Fixed models require very little computational overhead and are mainly used to provide better initial conditions for parameter estimation. Using computer simulations, the use of $N - 2$ fixed models and two adaptive models was found to be a reasonable compromise between computational complexity and performance in many adaptive problems [55]. If a fixed model $\hat{\theta}_i$ is selected (according to a performance criterion) at any instant k_0 , an adaptive model with

parameter vector $\hat{\theta}_a$ is initiated using $\hat{\theta}_i$ as an initial condition. If at a later time k_1 a different fixed model $\hat{\theta}_{i+1}$ is chosen, the adaptive model $\hat{\theta}_a$ is discarded and a new one is initiated at $\hat{\theta}_{i+1}$. Details regarding this procedure can be found in (Narendra and Balakrishnan 1997, [57]). If fixed models are to play an important role in the adaptive process their parameters must lie close to those of the plant, even as the latter varies with time. Since the plant parameter vector θ_0 is unknown and can vary arbitrarily, this is possible only if a very large number of fixed models is used within the compact set S . Alternatively, the location of the models is determined based on past experience assuming that the latter is available. Yet another approach, which does not require any prior information about the location of θ_0 , will be presented in chapter 7.

6.1.3 Switching and Tuning

Tuning is the process of incrementally adjusting the parameters of the control or estimation model. This is the method used in classical adaptive control where the evolution of the parameters is described by a differential equation. When multiple models are used, parameters can change discontinuously. To motivate the need for switching in an adaptive control system, we consider the case where it is known a priori that the plant can assume only one of two values (i.e. the plant parameter vector satisfies: $\theta_0 \in \{\theta_1, \theta_2\}$). Corresponding to each of the above, it is also known that controllers Γ_1 and Γ_2 exist such that θ_i together with Γ_i in the feedback path matches a stable reference model. It is further assumed that the plant θ_i together with controller $\Gamma_j (j \neq i)$ results in instability. If the plant were to switch rapidly between θ_1 and θ_2 the adaptive method used must be able to detect the change in the plant and switch to the appropriate controller to avoid instability. Since only the inputs and outputs of the plant are assumed to be known, detection of the change in the plant has to be concluded from the output estimation errors ϵ_i . Switching is desirable to react to rapid changes in the plant characteristics and avoid instability. Since the number of models is finite, while the number of values that the plant can assume within S is infinite, tuning is necessary if the control error is to tend to zero asymptotically. The essence of MMST is to combine switching and tuning efficiently such as to improve the performance of the closed-loop system while keeping all signals bounded. The control input u resulting from this procedure is in general piecewise-continuous.

A crucial role in the design of MMST is played by the switching criterion which determines when to switch from one model to another and which of the models should be the new one. A number of different performance indices can be defined based on the identification error $\epsilon(k)$ to determine which of the models best fits the plant at any instant. These may assume

the following forms:

$$\begin{aligned}
 \text{(i)} \quad J_i(k) &= \epsilon_i^2(k) \\
 \text{(ii)} \quad J_i(k) &= \sum_{\nu=0}^k \epsilon_i^2(\nu) \\
 \text{(iii)} \quad J_i(k) &= \alpha \epsilon_i^2(k) + \beta \sum_{\nu=0}^k \epsilon_i^2(\nu) \\
 \text{(iv)} \quad J_i(k) &= \alpha \epsilon_i^2(k) + \beta \sum_{\nu=0}^k \rho^{k-\nu} \epsilon_i^2(\nu)
 \end{aligned} \tag{6.1}$$

The criterion (i) represents instantaneous values of the square of the output error, (ii) the integral of the former, (iii) a linear combination of (ii) and (iii), and (iv) a modification of (iii) which includes a forgetting factor determined by the parameter $\rho \in [0; 1]$. Based on one of the criteria (i)–(iv) the model/controller pair $[\Xi_j, \Gamma_j]$ is used at every instant of time using $J_j(k) = \min_i J_i(k)$. Criterion (i) invariably results in rapid switching between controllers while relatively slow switching is achieved using (ii). As a consequence, the control errors are larger in the latter case. For satisfactory operation in rapidly time-varying environments, criterion (iv) is generally preferred and the parameter ρ is selected to achieve a compromise between speed and performance.

6.1.4 Control

The control input to the system is computed based on the parameters of the model that performs best according to one of the above criteria. The approach of using the parameters of an estimation model to determine the control input is referred to as an indirect one, in contrast to the direct approach where the controller parameters are identified directly based on the control error e . From the above discussion it is clear that the indirect approach is inherent to the MMST methodology. While a multiplicity of models can be used concurrently to estimate the plant, only one controller can be used at any instant to determine the control input.

Notice that the overall performance of the system will be judged on the basis of the *control error* e , whereas the choice of the controller at any instant is based on a performance index that depends upon the *estimation error* ϵ . As seen in chapter 3.1.3, equation (3.29), the control error deviates from the identification error by an amount depending upon the time-variation of the parameter. If a fixed controller is used, the two are equivalent. However, in general, the controller has to be adaptive in order for e to vanish. But this means that

the control based on the model with the smallest estimation error need not result in the smallest control error. This is the principal difficulty encountered when attempting to derive a quantitative measure of the performance improvement obtained through MMST.

6.1.5 Benefits

While being a natural extension of conventional adaptive control (as considered in chapter 3) to time-varying situations, the MMST methodology offers a number of intuitively appealing and practically relevant advantages:

- Its principal use is to detect changes in the environment (e.g. due to time-varying disturbances) and initiate appropriate control action.
- When the information needed to design an adaptive controller (e.g. its relative degree, order of the plant) is missing, multiple models can be used to obtain it.
- An application, which is pointed out in [11], is to combine the advantages of different models. One of the models, designed analytically, may assure stability, while another, designed heuristically, may have better performance. A proper combination of the two results in a stable system with improved performance.
- By placing fixed models in the neighborhood of the points in parameter space which the plant is likely to assume, the time for the parameters to converge can be reduced substantially and, hence, performance improved. Recall that, in the absence of a persistently exciting input, the models do not have to be close to the plant, since the identification error is small anywhere near the hypersurface defined in equation (2.8). But if the models are close, the error will be small for any input.

Example 6.1 Given a discrete-time plant described by,

$$y(k+1) = a(k)y(k) + b(k)u(k) \quad (6.2)$$

where the plant parameter vector $\theta(k) = [a(k) \ b(k)]$ is piecewise constant and switches between the elements of a finite set $S = \{\theta_1, \theta_2, \theta_3, \theta_4, \theta_5\}$ of unknown vectors. As indicated in figure (6.2), θ_2 corresponds to an unstable plant, while θ_4 is stable with eigenvalue $z = -1$, and all other parameter vectors are strictly stable. The objective is to control the plant such that the output tracks the desired output $y^*(k+1) = \sin(2\pi k/100) + \sin(2\pi k/90)$ with small errors even as the plant parameters change discontinuously from one element in S to another at random instants of time. The switching is assumed to be governed by an

ergodic Markov chain, where $p = 0.9$ is the probability of remaining at a given parameter θ_i , $i = 1 \dots 5$ and $q = 0.025$ is the probability of switching to one of the other parameters.

The first row of the figure displays the nature of the time-variation of the plant. Part (i) shows the response of the system using classical adaptive control with a single model. The performance is clearly unacceptable and the parameters do not converge anywhere. In part (ii), the location of the plant parameters is assumed to be known so that five fixed models having parameters identical to the plant can be used. At any instant, the current plant parameter is detected and the system is controlled using the parameters associated with the best model. In order to ensure a fast response, the instantaneous switching criterion is used to decide which of the models performs best. It is seen that an error occurs at the instants following a switching of the plant parameters, which persists over an interval of length d , where d is the relative degree of the system. To avoid the error, the “new” control input u would have to be computed d instants before a switching occurs. It is clear that this is impossible since the switching sequence is unknown. Hence, the error is inherent. In part (iii), the fixed models are replaced by adaptive ones which are initialized within a small neighborhood $N = \{\hat{\theta} \mid \|\hat{\theta} - \theta_i\| < \varepsilon, i = 1 \dots 5\}$ of the plant parameters, i.e. $\varepsilon = 0.5$. If the identification error of a given model $\hat{\theta}_i$, $i = 1 \dots 5$, is small, this model is updated while all other models retain their current position in parameter space. The combined performance criterion (iv) in equation (6.1) with $\alpha = \beta = 1/2$ and $\rho = 0.3$ was found to result in acceptable performance. In addition, observe that the parameters approach the true parameters of the system. Finally, in part (iv) of the figure the experiment is repeated, but in this case, the models are initialized in a larger neighborhood, $\varepsilon = 1$. The performance is seen to be similar to case (i). The example demonstrates, that the performance of the MMST approach in a time-varying environment critically depends on the location of the models and, hence, the prior information about the plant. \square

6.2 Proof of Stability

Stability proofs have been derived for linear deterministic and stochastic systems as well as certain classes of nonlinear systems [23], [63]. In all these cases, the unknown parameters are assumed to be constant. The principal stability question encountered in MMST can be qualitatively described by considering two models Ξ_1 and Ξ_2 and two controllers Γ_1 and Γ_2 . The estimated output given by the two models is \hat{y}_1 and \hat{y}_2 respectively and the corresponding estimation errors are ϵ_1 and ϵ_2 . We assume that the plant Σ is stable with controller Γ_1 in the feedback path and unstable with Γ_2 . If, based on the switching criterion, Ξ_2 is chosen

at a specific instant, the controller Γ_2 would be used in the feedback path resulting in an unstable closed loop. The question that has to be addressed is how the switching criterion should be chosen such as to result in the controller switching to Γ_1 in a finite time. In what follows, we present the proof of linear deterministic adaptive control using multiple models which builds upon the arguments presented in chapter 3.1.4 in the single model case. Three different configurations of models are considered:

- (i) N adaptive models
- (ii) One fixed model and one adaptive model
- (iii) $N - 2$ models and 2 adaptive models

The plant is described by the deterministic equation

$$y(k+d) = \sum_{i=0}^{n-1} a_i y(k-i) + \sum_{i=0}^{n-1} b_i u(k-i) \quad (6.3)$$

where the **constant** parameters a_i and b_i are unknown. As in 3.1.3, the plant is assumed to be minimum-phase and its relative degree d (delay) and order is known. The objective of the control is to determine a bounded input $u(k)$ such that the output of the plant $y(k+d)$ asymptotically tracks a given bounded reference output $y^*(k+d)$. It is assumed that $y^*(k+d)$ is known at time k (or alternately if $y^*(k+d) = r(k)$, $r(k)$ is specified). The reference model in this case is a pure delay of d units, i.e. it has the transfer function z^{-d} .

6.2.1 Case (i): All adaptive models

From our discussion in chapter 3.1.4 it is clear that a single model is sufficient to control an unknown plant in a stable fashion. The objective of this paragraph is to show that this is also the case when multiple models are used. More importantly, the closed-loop system is stable even as the switching between the models (and, hence, the corresponding controllers) is carried out in a random fashion.

If N models are used to estimate the plant parameters, and at time k the i th model $\hat{\theta}_i$ is chosen, the control input $u(k)$ is computed from the equation

$$y^*(k+d) = \phi(k)^T \hat{\theta}_i(k) \quad (6.4)$$

The control error at time k is given by

$$e(k) = \epsilon_i(k) + \phi(k-d)^T [\hat{\theta}_i(k-1) - \hat{\theta}_i(k-d)] \quad (6.5)$$

where $\epsilon_i(k)$ is the identification error of the i th model, see equation (3.29). If the model at the next instant is chosen randomly as $\hat{\theta}_j$, the control error e satisfies an equation similar to (6.5) with i replaced by j . This process can be repeated at every instant and a model chosen randomly from $\hat{\theta}_i$, $i = 1 \dots N$. The important fact to note is that the regression vector is the same for all of them. The parameters of all the models are estimated using the same algorithm, only the initial conditions are different. From the properties of the estimation algorithm (see e.g. chapter 2.1.1) we know that $\|\hat{\theta}_i(k-1) - \hat{\theta}_i(k-d)\| \rightarrow 0$ for all i . Hence both terms on the right hand side of equation (6.5), when normalized, i.e. $\frac{\epsilon_i(k)}{[1+\phi(k-d)^T\phi(k-d)]^{\frac{1}{2}}}$ and $\frac{\phi(k-d)^T[\hat{\theta}_i(k-1)-\hat{\theta}_i(k-d)]}{[1+\phi(k-d)^T\phi(k-d)]^{\frac{1}{2}}}$ tend to zero. Hence, using the same arguments as in the single model case, it follows that $e(k)$ grows at a slower rate than $\|\phi(k)\|$, $\|\phi(k)\|$ is bounded, and $\lim_{k \rightarrow \infty} e(k) = 0$.

In the proof, any one of the adaptive controllers is chosen randomly to control the system. This implies that with multiple models, the question of stability can be decoupled from that of performance and the switching procedure can be based entirely on the latter. Notice, that the proof can be directly extended to the case where adaptive models are either introduced or removed, provided that at least one model (referred to as a free-running model) is not disturbed during the adaptive process.

6.2.2 Case (ii): One adaptive model and one fixed model

If the plant is time-invariant, the benefit of having multiple models is that if the latter are initialized at different locations in parameter space, one model $\hat{\theta}_i$, $i = 1 \dots N$, may be close to the plant and will result in a smaller error ϵ_i . However, if no adaptive model is close to the plant, there may be no improvement in performance. It is clear that keeping all models adaptive requires considerable computational overhead and is not efficient since a single adaptive model was found to be adequate for stability in section 3.1.4. The idea is to replace $N - 1$ adaptive models by fixed ones. These can be thought of as providing convenient initial conditions for adaptation. In this section we consider the case of $N = 2$ models, one adaptive and one fixed.

The adaptive and the fixed model result in error equations $\epsilon(k) = \tilde{\theta}(k)^T \phi(k)$ and $\epsilon_f(k) = \tilde{\theta}_f^T \phi(k)$ respectively where, in general, only the error of the adaptive model $\epsilon(k)$ tends to zero. If $|e_f(0)| < |e(0)|$ and switching criterion (iii) in equation (6.1) is used, the system will initially start with the fixed model, and the control corresponding to it will be applied. However, since $J_i(k)$ is bounded while $J_f(k)$ (of the fixed model) grows monotonically with k , the system will switch to the adaptive model in a finite time. In the non-generic case that

the fixed model is identical to the plant, $\epsilon_f(k)$ is zero and no switching occurs. The above arguments are still valid if an arbitrary number N of fixed models are used. Notice that $\epsilon_f(k)$ can be zero on a subsequence $\{k_t\}$, $t = 1 \dots \infty$ because the regression vector $\phi(k_t)$ is orthogonal to $\tilde{\theta}_f$ at those instants of time. If criterion (i) is used, switching will not stop in finite time since $\epsilon(k_t) > \epsilon_f(k_t)$. This emphasizes the importance of using an integral criterion in a time-invariant situation.

6.2.3 Case (iii): (N-2) fixed models and 2 adaptive models

If the fixed model in the previous paragraph is closer to the plant than the adaptive model (as determined by the performance criterion J), the system switches to it and switches back to the adaptive model once $\epsilon(k)$ has become small. However, faster tracking can be obtained if a new adaptive model is initialized at the same value as the fixed model. In view of the above discussion, the fixed models have to be located based on the past performance of the plant. One adaptive model is free-running and is included to assure stability. A second model is initiated whenever a fixed model is found to be superior. A large number of fixed models may be needed to obtain one model with distinguished performance at which the adaptation can be initiated. The reinitialization of the second adaptive model is central for improving the performance of the system. If at any instant k_0 , $J_f(k_0)$ is found to be a minimum, the adaptive model is initialized using the parameter vector θ_f and the performance index $J_f(k_0)$. The adaptive process is continued. If, at a later instant k_1 , a different fixed model (with parameter $\theta_g \neq \theta_f$) is superior, the adaptive model is discarded and a new one is initialized at θ_g and $J_g(k_1)$.

The introduction of the additional (re-initialized) model does not adversely affect the stability of the overall system. Let the system switch between fixed and adaptive models at every instant of time. From the arguments provided in section 6.2.2, this can only last for a finite interval of time, since the performance index of the free-running adaptive model will eventually become smaller than the indices of all the fixed models. It is seen that the existence of the free-running model is the key argument for convergence. If the plant is rapidly time-varying, the error of a single adaptive model cannot tend to zero and the arguments are substantially more difficult. In this case, all models can be made adaptive and the switching criterion modified by using a finite window (i.e. $\sum_{k-T}^k e^{2(\nu)}$) or a forgetting factor $0 < \rho \leq 1$. In the ideal case, $J_i(k) \rightarrow 0$ for all adaptive models, while $J_f(k) > 0$ for the fixed models. The fundamental assumption made in this context is that at least one adaptive model is close to every parameter that the plant is likely to assume. This has already been observed in cases (iii) and (iv) of example 6.1.

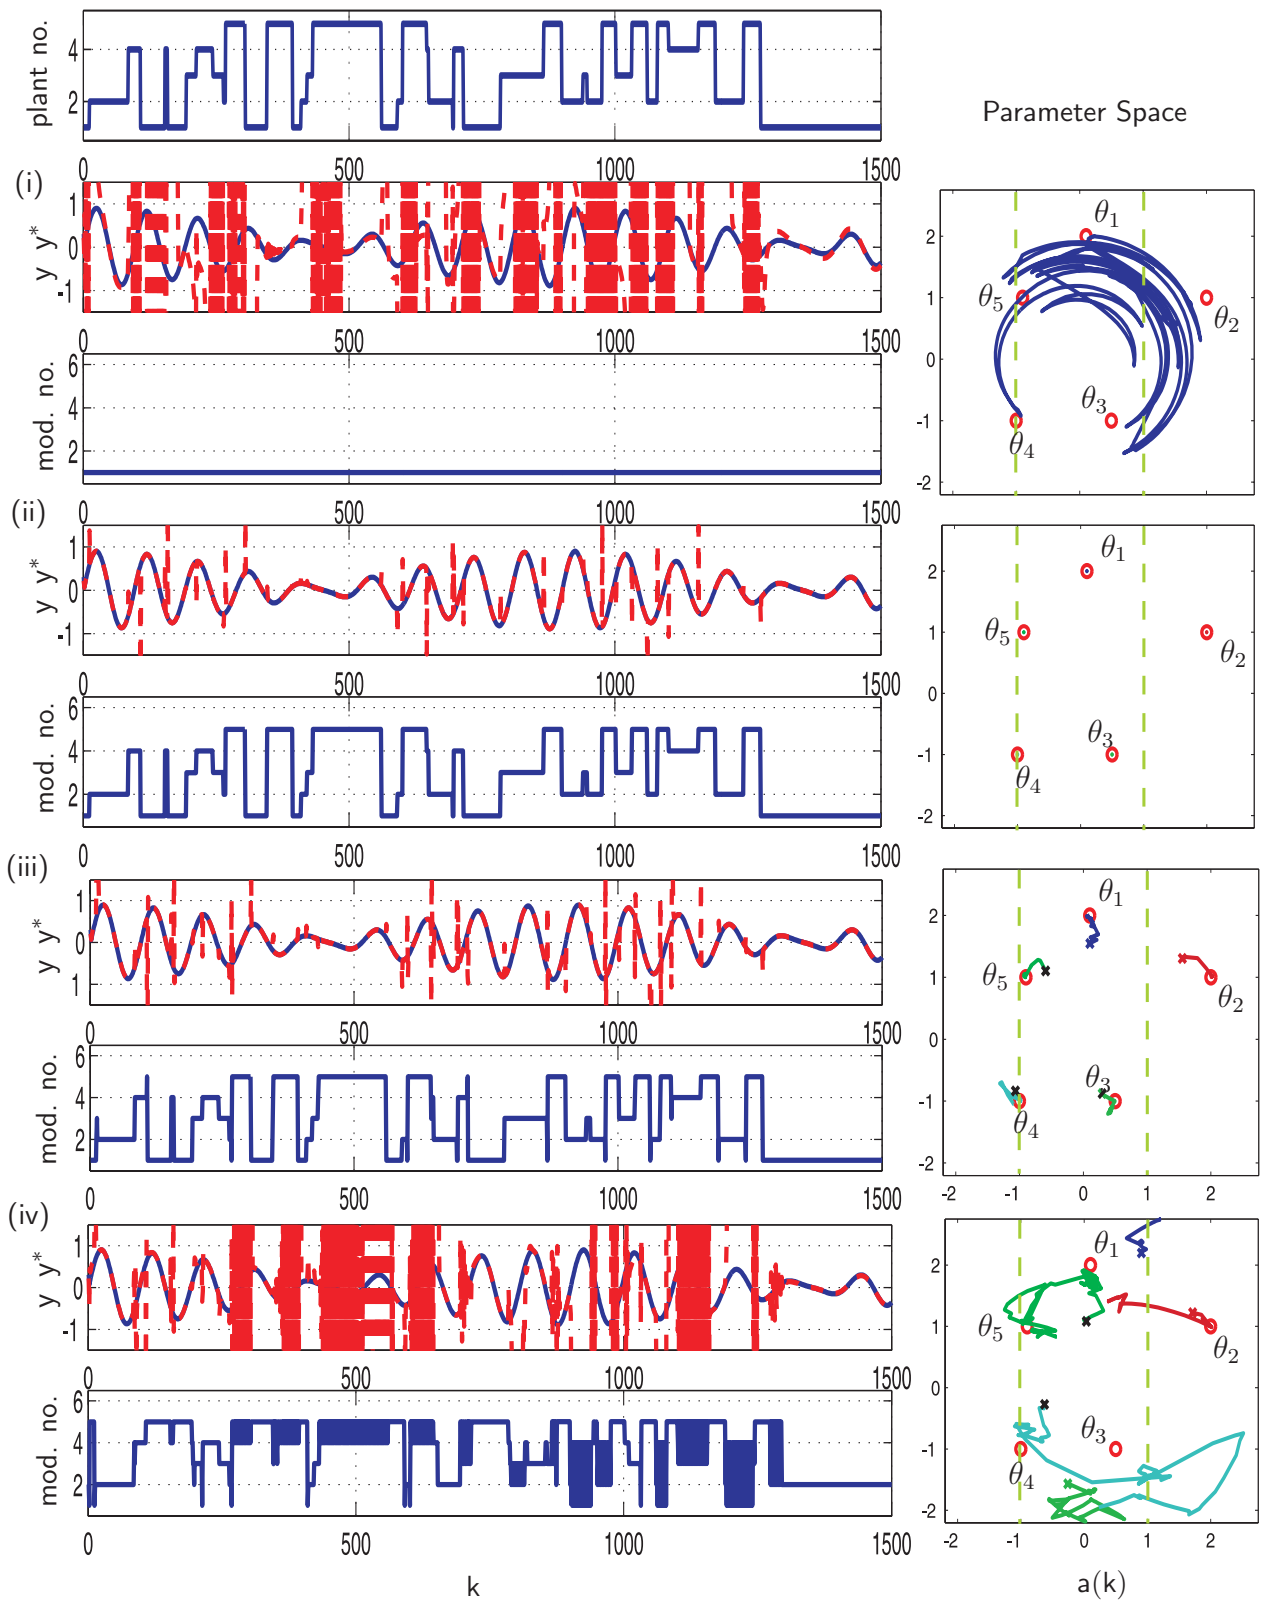


Figure 6.2: Multi-Model Adaptive Control (MMST)

Chapter 7

Self–Organization

In this section we pursue our discussion on adaptive control of rapidly time–varying systems. The parameter vector θ of the system is assumed to switch randomly between the elements of a finite set $S = \{\theta_1, \theta_2, \dots, \theta_N\}$ where θ_j , $j \in \Omega = \{1, 2, \dots, N\}$ are unknown constant vectors and the instants at which the switching takes place are also unknown. The integer N (or an upper bound on N) is assumed to be known. As seen in the previous chapter, a single adaptive model is –in general– not capable of assuring stability of the time–varying system. If multiple ($\geq N$) adaptive models are used one has to show that the errors of at least N such models tend to zero in order to obtain a bounded output error. In this context, the assumption that the models are initialized close to the plant parameters was seen to be of fundamental importance.

In this chapter, no such constraint is imposed upon the initialization of the models. A method is proposed by which N models are simultaneously updated in order to identify the elements of the set S . The approach relies upon a time–varying adaptive gain, the magnitude of which depends upon the errors of all the models with respect to the plant parameter that is in existence at a given instant. Extensive simulation studies reveal that for any initial condition, the N identification models converge to N different plants contained in the set S . The convergence is found to be a complex, nonlinear process. Our principal objective is to prove theoretically that the proposed learning algorithm results in the convergence. This is a formidable problem, and the investigations of the past two years have demonstrated that numerous special cases and problems have to be addressed before the appropriate mathematical tools are in place for dealing with the problem theoretically in a succinct fashion.

In the following, detailed analyses of the various issues are presented, often with a qualitative flavor. Such qualitative explanations are essential to an understanding of the theoretical diffi-

culties involved. The emphasis is on algebraic systems and on providing simple explanations for the manner in which the N models converge to the N plants. The insights provided by the convergence in such static systems motivates the use of similar procedures in dynamical systems. Parts of the following discussion were presented at the twelfth Yale Workshop on Adaptive and Learning Systems (Feiler and Narendra, 2003 [21]).

7.1 Introduction

Even in the non-adaptive case, theoretical results on the stability of time-varying systems are sparse. They are restricted to special cases, such as vanishing or periodic time-variations. In adaptive control, our goal is to use feedback in order to match the output of an unknown time-varying system to a known time-invariant (reference) system. The indirect approach consists in adjusting an identification model to obtain a small error and use the parameters of the model for control. The following example illustrates that in a time-varying context, its success in assuring bounded solutions critically depends upon the location of the models.

$$x(k+1) = a(k)x(k) + u(k) \quad a(k) \in S = \{a, -a\}, \quad a > 1 \quad (7.1)$$

The plant switches between the values $\pm a$ periodically with period $2T$. The stabilizing controller is defined as $u(k) = -\hat{a}(k)x(k)$ where $\hat{a}(k) \in \{\hat{a}, -\hat{a}\}$, $\hat{a} > 0$ are the parameters of two (fixed) models. We determine the size of the neighborhood $U = \{\hat{a} : |a - \hat{a}| \leq \varepsilon\}$ within which the models have to be located such that the closed-loop system has bounded solutions. A straightforward calculation shows that,

$$|x(k+T)| = |(a + \hat{a}) \cdot (a - \hat{a})^{[T-1]}| |x(k)| \stackrel{!}{\leq} |x(k)| \quad \forall k > 0 \quad (7.2)$$

At every instant, the controller is determined based on the model that performs best. Due to the relative degree $d = 1$ of the plant, an erroneous control law will be applied at the instant following a switching. This results in an unstable gain $(a + \hat{a}) > 1$ in equation (7.2). With an appropriately chosen $\hat{a} > 0$, a stable gain $(a - \hat{a}) < 1$ is obtained over the remaining $T - 1$ instants of the switching interval. As the switching frequency increases the contribution of the unstable gain becomes large. By solving the inequality in (7.2) for the worst case $T = 2$, we obtain,

$$\sqrt{a^2 - 1} \leq \hat{a} \leq \sqrt{a^2 + 1} \quad (7.3)$$

which determines the size of the neighborhood U and, hence, the location of the fixed models. If adaptive models are used they must enter the neighborhood U in a single step if the control law is to result in bounded solutions and the switching interval is $T = 2$. In general, *at least*

p steps are needed for the adaptive model to converge to an unknown vector $\theta \in \mathbb{R}^p$ where p is the number of unknown parameters. This motivates the use multiple models, switching and tuning. If there is a fixed model $\hat{\theta}_i$ in the neighborhood U_j of every plant θ_j , $j = 1 \dots N$, we may switch to $\hat{\theta}_i$ whenever the corresponding plant is in existence. Adaptation is thereby achieved in a single step regardless of the dimensionality p of the parameter space. In view of the above discussion, the closed-loop system will have bounded solution even if the switching is fast ($T = 2$).

If the parameters of an unknown plant vary slowly, or if the parameters vary discontinuously but remain approximately constant for a long period of time, classical adaptive control can be used with a single model. If, however, the plant variations are both rapid and large, the MMST approach has been found to ensure bounded solutions provided that at least one of the N fixed models is sufficiently close to the plant so that adaptation from that model results in satisfactory performance. It is clear that the choice of the models is an important consideration while using the MMST approach. In papers dealing with the latter it is claimed that the choice of the models would be based on the past performance of the system—in other words, it involves a learning process [56].

It is obvious that the ways in which plant parameters can vary are infinite and would correspond to the entire class of bounded piecewise continuous functions. To make the problem analytically tractable we consider only the case where θ assumes a finite number N of constant values, i.e. $\theta \in S = \{\theta_1, \theta_2, \dots, \theta_N\}$. Even in this special case, θ can switch randomly between the elements of S . The switching is assumed to be *regular* in the sense that every element of the set is assumed at least once within an interval of finite length $T^* > 0$, and the plant remains constant over an interval of minimum length $T > 0$. Again, to keep the arguments simple, we assume that $\theta(k)$ is periodic with period NT , where $\theta(k) \equiv \theta_j$, $j \in \Omega$, for T units of time. It is assumed that N is specified but that the elements of S and the interval T are unknown. The first objective then is to determine the values θ_j asymptotically. A second objective is to define a control law, based on the best estimate, such that the overall system is stable and all signals remain bounded. In focusing on this problem, we are influenced by earlier work by (Kawato et al., 2001 [29]) on motor learning and control. However, unlike our predecessors, our primary interest is in the stability of systems which learn and adapt at the same time.

Even though the description that we have given thus far corresponds roughly to the problems encountered in practice, we would like to abstract from this situation a problem that can be posed analytically whose solution would contribute greatly to our understanding of the control of unknown and rapidly time-varying dynamical systems. The following is a succinct

statement of such a problem.

7.1.1 Qualitative Description

The elements of the vector θ represent the unknown parameters of a linear dynamical system. θ can assume one of a finite number N of constant values θ_j , $j \in \Omega$, and switches between them at random instants of time but remains constant over an interval $T = [k_t, k_{t+1}]$. If $\theta(k) = \theta_j$ at a given instant, we shall refer to it as the j^{th} environment. If k_t and the j^{th} environment, which exists over the interval T are known (note that j and not θ_j is assumed to be known), the identification problem becomes a relatively straightforward one. In such a case, N models can be used to identify the N values θ_j , $j \in \Omega$, with only one model being updated at every instant of time (knowing that the j^{th} environment exists at that instant). The problem becomes truly difficult when neither k_t nor the environment j which is in existence are known at any instant. It is then not evident which of the N models has to be updated. Our objective is to adapt N models simultaneously so that every model converges to a *different* element of the set S , and the controlled system is stable.

7.1.2 The method

Assuming that $\theta(k) \in S = \{\theta_1, \theta_2, \dots, \theta_N\}$, where $\theta_j, j \in \Omega = \{1, 2, \dots, N\}$ are unknown constant vectors, N estimation models are used at every instant to generate estimates of the elements of S . Under the assumption that $\theta(k)$ assumes every value $\theta_j \in S$ at least once in a finite interval of length T , the objective is to determine the class of algorithms that will result in convergence of each $\hat{\theta}_i(k)$ to one of the unknown parameter vectors θ_j , such that for each element of S there is one model and no two models converge to the same element.

The idea is that to each model $\hat{\theta}_i$ an adaptive gain is assigned which depends on its relative performance (i.e. when compared to all the other models). More precisely, the step-size, by which parameter vector $\hat{\theta}_i$ is updated depends both on its estimation error $\epsilon_i(k)$ as well as the errors $\epsilon_l(k)$, $l \neq i$ of all the other models. If there is an $l \neq i$ such that $\epsilon_l(k)$ is less than $\epsilon_i(k)$, the step-size $\eta_i(k)$ is decreased. If, in turn, $\epsilon_i(k)$ is the smallest among all the estimation errors then $\eta_i(k)$ is large. From the above discussion it is clear that in order to control the system in a stable fashion, the models must be distributed in parameter space, such that for every plant parameter θ_j , there is a model contained in its neighborhood U_j . The motivation for introducing performance-dependent step-sizes is precisely to prevent all models from converging to the same point. But this implies, that only the model that is close to a given neighborhood U_j should be updated using a large step-size. This is in contrast to the fact

that ϵ_i may become small as a result of the scalar product $\epsilon_i(k) = -\phi(k-d)^T[\hat{\theta}_i(k-1) - \theta]$ even while the parameter error $\tilde{\theta}_i(k-1) = [\hat{\theta}_i(k-1) - \theta]$ is large at that instant. To prevent this, the performance has to be assessed by means of a criterion $J_i(k)$ defined in (6.1) as a function of ϵ_i and its past values.

Finally, notice that the algorithm is such that the distribution of the models takes place without help by a supervisor/ coordinator. This is possible since every model has full information about the performance of all its competitors. Hence, the title of this chapter, self-organization (of multiple adaptive models).

7.1.3 Statement of the Problem

Consider a linear time-varying dynamical system,

$$y(k+d) = \sum_{\nu=0}^{n-1} a_\nu(k) y(k-\nu) + \sum_{\nu=0}^{n-1} b_\nu(k) u(k-\nu) \quad (7.4)$$

In other words, $y(k+d) = \phi(k)^T \theta(k)$, using the time-varying plant parameter vector

$$\theta(k) = [a_0(k), \dots, a_{n-1}(k), b_0(k), \dots, b_{n-1}(k)]^T$$

and $\phi(k) = [y(k), \dots, y(k-n+1), u(k), \dots, u(k-n+1)]^T$. It is known that $\theta(k) \in S = \{\theta_1, \theta_2, \dots, \theta_N\}$ where $S \subset \mathbb{R}^{2n}$. M adaptive models are defined,

$$\hat{y}_i(k+d) = \phi(k)^T \hat{\theta}_i(k) \quad i = 1 \dots M, \quad M \geq N \quad (7.5)$$

where $\hat{\theta}_i(k)$ and $\hat{y}_i(k+d)$ are estimates of $\theta \in S$ and $y(k+d)$ respectively. To keep our arguments simple, we set $M = N$ (\rightarrow see section 7.2.2 for a discussion regarding the number of models required for convergence). In order to update the models $\hat{\theta}_i(k)$, any standard estimation algorithm of the form (2.4) can be used

$$\hat{\theta}_i(k) = \hat{\theta}_i(k-1) + M_P(k) \phi(k-d) \epsilon_i(k) \quad (7.6)$$

where $\epsilon_i(k) = y(k) - \hat{y}_i(k)$ is the estimation error which satisfies

$$\epsilon_i(k) = -\phi(k-d)^T [\hat{\theta}_i(k-1) - \theta] \quad (7.7)$$

In view of its simplicity and geometric interpretation, we choose the parameter projection algorithm, i.e.

$$M_P(k) = \frac{\eta_i(k)}{\phi(k-d)^T \phi(k-d)} \quad (7.8)$$

While the projection vector $\phi(k-d)$ is the same for all models, its length varies according to the error $\epsilon_i(k)$ and step-size $\eta_i(k)$ associated to every model $\hat{\theta}_i(k)$, $i \in \Omega$. The step-size is a function of the performance of $\hat{\theta}_i(k)$ relative to all the other models:

$$\eta_i(k) = \eta_0 \frac{\frac{1}{J_i(k)}}{\frac{1}{J_1(k)} + \frac{1}{J_2(k)} + \cdots + \frac{1}{J_N(k)}} \quad 0 < \eta_0 \leq 1 \quad (7.9)$$

As stated earlier, the performance index $J_i(k)$ is used to determine which model is closest to a given plant θ_j . In general, this cannot be concluded from the instantaneous error $\epsilon_i(k)$ but requires observing a sequence of estimation errors $\{\epsilon_i(k), \dots, \epsilon_i(k-T_0)\}$ over an interval of sufficient length T_0 . This will be made more precise at a later stage (in section 7.3.1). For the moment we assume that the performance index provides an unambiguous measure of the parameter error, i.e. the model with the smallest performance index $J_i(k) = \min_{s \in \Omega} J_s(k)$ corresponds to the model that is closest to θ_j in parameter space, $\tilde{\theta}_i(k) = \min_{s \in \Omega} \|\hat{\theta}_s(k) - \theta_j\|$. As a consequence of (7.9), only the closest model is updated by a large amount and, hence, its advantage over the other models consolidated. Notice that the nonlinear function (7.9) does not define a “winner takes it all”-policy since the *sum* of all step-sizes equals η_0 , i.e. $\eta_l(k)$, $l \neq i$, is zero only in the limit when $J_i(k) \rightarrow 0$ as $k \rightarrow \infty$ (if indeed $\hat{\theta}_i$ converges to θ_j). In the case of convergence, there must be a stage after which $\hat{\theta}_i$ enters the critical neighborhood U_j of θ_j . In order for the closed-loop system to be stable, there must be one of the N models in *every* such neighborhood U_j , $j = 1, \dots, N$. If the corresponding parameters are used to define the control law, we obtain bounded solutions even as the system switches randomly between the elements of the set S . In addition, the tracking error tends to zero, as the models converge to their respective parameters θ_j , $j = 1, \dots, N$. As seen previously, the control action will be incorrect over an interval of length d after a switching in the environment has occurred. Suppose that $\theta(k-1) \equiv \theta_1$ and $\theta(k) \equiv \theta_2$. Then the control input at instant $k-d$ should have been computed based on the parameter estimate at instant k , i.e. θ_2 . However, at $k-d$ it is not known that the environment changes at k nor which $\theta_j \in S$ exists at that instant. A critical assumption is therefore that the interval T over which the plant is constant is sufficiently large in order to allow for the controller to compensate for its erroneous first d steps during the remaining $T-d$ instants of time.

The following assumptions [A] are made:

1. The system is of known, constant order n and relative degree d .
2. It is minimum-phase for every constant value $\theta \in S$.
3. $\theta(k)$ assumes every value $\theta \in S$ at least once in a finite interval of length T^* and remains constant over an interval $T \geq \max\{T_0, d+1\}$.

Hence, while the system is rapidly time-varying in the sense that the parameter variations are rapid and large, it is slowly time-varying when compared to the impulse responses of its time-invariant subsystems. Our objective is twofold:

1. Prove that the algorithm defined in equations (7.6), (7.8) in conjunction with the time-varying adaptation gain (7.9) results in asymptotic convergence of each $\hat{\theta}_i(k)$ to one of the unknown parameter vectors θ_j , such that for each element of S there is a model and no two models converge to the same element.
2. Determine a bounded control input based on the model that performs best (according to a criterion that has to be specified) such that the tracking error tends to zero except for finitely many instants at which the environment switches.

7.2 Parameter Convergence in Static Systems

Our first objective is to identify a set of plants. While this is a standard problem if the elements of the set are identified one by one, the specific problem here is that the identification is carried out for all plants simultaneously. In the problem statement, an algorithm to update the parameters has been proposed consisting of a standard projection scheme (7.6) and a time-varying adaptive gain (7.9). The justification for using this algorithm comes from extensive simulation studies in which the algorithm was seen to be effective in solving the simultaneous identification problem. The simulations will be presented at a later stage. Here, we are primarily interested in proving that the algorithm indeed results in the convergence of a set of models to the set of plants.

While attempting to prove convergence, it is useful to consider static systems which are substantially simpler analytically and at the same time retain the principal features of the original problem. When no dynamics are present, the parameters are adjusted at every instant based on their distances to the prevailing plant parameter θ which is an element of the set $S = \{\theta_1, \theta_2, \dots, \theta_N\}$. If $\hat{S} = \{\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_N\}$ is the set of model parameter vectors the error vector between elements of S and \hat{S} is given by $\tilde{\theta}_{ij}$ where $\tilde{\theta}_{ij} = \hat{\theta}_i - \theta_j$. The first index i refers to the i th model, whereas the second index refers to the j th environment, which corresponds to a point in \mathbb{R}^p , e.g. $p = 2n$. Let the distance of the i th model to the j th point be denoted by $d_{ij} = \|\tilde{\theta}_{ij}\|$. For convenience of notation the plant index j is dropped whenever no particular reference is made to any specific element in S . We assume that an algorithm can be determined by which the norms of all the error vectors are reduced at every

instant according to the equation

$$\tilde{\theta}_i(k+1) = \tilde{\theta}_i(k) - \eta_i(k)\tilde{\theta}_i(k) \quad (7.10)$$

where the step-size $\eta_i(k)$ depends upon the distances $d_i(k)$, $i \in \Omega$ of all models. We define:

$$\eta_i(k) = \eta_0 \frac{\frac{1}{d_i^q(k)}}{\frac{1}{d_1^q(k)} + \frac{1}{d_2^q(k)} + \cdots + \frac{1}{d_N^q(k)}} \quad q > 1 \quad (7.11)$$

In contrast to (7.9), the step-size is directly defined in terms of the parametric distance $d_i(k)$. Since d_i is known for all $i \in \Omega$ the location of an element $\theta \in S$ in p -dimensional parameter space could also be determined using $p+1$ distance measurements (provided that $N \geq p+1$). However, as stated above, the static algorithm is merely an abstraction of the situation encountered in the dynamical case. Hence, all elements of the set S are identified simultaneously by sequentially updating the position of N adaptive models, the self-organization of which we are interested in studying. In intuitive terms, all models move concentrically towards the point that is currently in existence. The amount by which they update their positions is governed by equation (7.11). The objective is to show that the models separate and converge to different elements of the set S .

The plant parameters $\theta(k)$ assume a finite number N of constant values, i.e. $\theta \in S = \{\theta_1, \theta_2, \dots, \theta_N\}$. We assume that $\{\theta(k)\}$ be an ergodic Markov process with state space S . Because of ergodicity, there is a $T^* > 0$ so that the plant assumes every element of S at least once over an interval of finite length T^* , i.e. for every θ_j , $j \in \Omega$, there is a $t \in [k, k+T^*]$ for which $\theta(t) \equiv \theta_j$. T^* is referred to as a *cycle*. As an example, we assume that $\theta(k)$ is periodic with period $T^* = TN$, where $\theta(k) \equiv \theta_j \in S$ for T units of time. In such a case, the algorithm defined in equations (7.10) and (7.11) yields:

Proposition 7.1 For every $\theta_j \in S$ there is a unique $\hat{\theta}_i \in \hat{S}$ which satisfies:

$$\lim_{k \rightarrow \infty} \|\hat{\theta}_i(k) - \theta_j\| = 0$$

□

Each model converges to a different element of S . Since no prior information is available about the possible location of the parameters $\theta_j \in S$, $j \in \Omega$ we assume that the model parameters are initialized at random, i.e. no model is distinguished with respect to any element of S . Simulation studies have revealed that the parameters converge in a non-monotonic fashion. Hence, no straightforward arguments exist, by which parameter convergence can be established. An additional difficulty arises from the fact, that it is not clear a priori which model converges to which parameter. Indeed, depending on initial conditions and the environment *any* model may converge to *any* of the parameter vectors θ_j .

In the following, we present a detailed analysis of those aspects of the algorithm which serve to set up a mathematical framework within which the convergence problem can be addressed. The proof itself remains an open question.

7.2.1 Time-invariant Environment

In the definition of the time-varying adaptive gain the inverse of the distance of any model to the point that is currently in existence is raised to the power q . It is seen that q critically influences the convergence properties of the algorithm. This can be appreciated even in the case where no switching occurs, i.e. the environment is time-invariant.

Let $\{\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_M\}$ be a collection of models. Simulations studies have revealed that for $q \in [0, 1]$ all models converge to the same point whereas for $q \in (1, \infty)$ only the one that is closest converges. In the second case, the algorithm was tested for $q \leq 12$, since any larger exponent leads to small numbers which cannot be handled numerically by the computer.

In the following, we provide a proof that the value $q = 1$ indeed separates two qualitatively different ways of convergence. Without loss of generality, we assume that the number of models is $M = 3$ and $d_1 < d_2 < d_3$. We obtain the nonlinear coupled difference equations:

$$d_1(k+1) = d_1(k) - [d_2(k)d_3(k)]^q d_1(k)/d_S(k) \quad (7.12)$$

$$d_2(k+1) = d_2(k) - [d_1(k)d_3(k)]^q d_2(k)/d_S(k) \quad (7.13)$$

$$d_3(k+1) = d_3(k) - [d_1(k)d_2(k)]^q d_3(k)/d_S(k) \quad (7.14)$$

where $d_S(k) = [d_2(k)d_3(k)]^q + [d_1(k)d_3(k)]^q + [d_1(k)d_2(k)]^q$ is the normalizing term. The time index k will be omitted wherever no confusion arises.

First notice, that since $\hat{\theta}_1$ is closest it must converge for any q . This can be seen as follows. Dividing equation (7.12) by $[d_2d_3]^q$ we obtain:

$$d_1(k+1) = d_1(k) - \frac{d_1(k)}{1 + \left[\frac{d_1}{d_2}\right]^q + \left[\frac{d_1}{d_3}\right]^q} \quad (7.15)$$

In view of the above assumption, $\left[\frac{d_1}{d_2}\right] < 1$ and $\left[\frac{d_1}{d_3}\right] < 1$. Hence,

$$d_1(k+1) < d_1(k) - 1/3 d_1(k) \quad (7.16)$$

But this means that $0 \leq d_1(n) < (2/3)^n d_1(0)$, and hence

$$\lim_{n \rightarrow \infty} d_1(n) = 0 \quad (7.17)$$

In proving the second property, i.e. the fact that $\lim_{k \rightarrow \infty} d_2(k) \neq 0$, $\lim_{k \rightarrow \infty} d_3(k) \neq 0$, for all $q > 1$, we rely on the following lemma.

Lemma 7.1 Given the scalar difference equation

$$x(k+1) = [1 - m(k)]x(k) \quad x(0) > 0 \quad (7.18)$$

where $0 < m(k) < 1$ for all $k \in [0, \infty)$. If $m(k) < 1/k^q$ and $q = 2$ there exists a constant $\bar{c} > 0$ such that $\lim_{k \rightarrow \infty} x(k) \rightarrow \bar{x} > \bar{c}x(0)$. \square

Proof. Let $m(k) = 1/k^q$ and $k > 1$. The solution of (7.18) can be computed explicitly as

$$x(n) = \prod_{k=2}^n \left[1 - \frac{1}{k^q} \right] x(0) \quad (7.19)$$

For $q = 1$ we obtain,

$$x(n) = \left[\frac{n-2}{n-1} \right] \left[\frac{n-3}{n-2} \right] \cdots \frac{3}{4} \frac{2}{3} \frac{1}{2} x(0) = \frac{1}{n-1} x(0) \quad (7.20)$$

Hence, $x(n) \rightarrow 0$ as $n \rightarrow \infty$. In contrast to that, for $q = 2$ the infinite product becomes,

$$\begin{aligned} x(2n-1) &= \left[\frac{(2n-1)^2 - 1}{(2n-1)^2} \right] \cdot \left[\frac{(2n-2)^2 - 1}{(2n-2)^2} \right] \left[\frac{n-1}{2n-3} \right] \cdots \cdot \frac{15}{16} \frac{8}{9} \cdot \frac{3}{4} \cdot x(0) \\ &= \left[\frac{n}{(2n-1)} \right] \cdot 1 \cdots \cdots \cdots \cdot 1 \cdot 1 \cdot x(0) \\ &= [n/(2n-1)] x(0) \end{aligned} \quad (7.21)$$

As $n \rightarrow \infty$ we obtain $x(n) \rightarrow 1/2 x(0)$. In other words, the above inequality holds with $\bar{c} = 1/2$. \blacksquare

In the following, we show that $\hat{\theta}_2$ and $\hat{\theta}_3$ satisfy the conditions of the lemma. The update equations for the models (7.13) and (7.14) are of the form (7.18) where m corresponds to $[d_1 d_3]^q / d_S$ and $[d_1 d_2]^q / d_S$ respectively. In the case of $\hat{\theta}_2$, we divide (7.13) by $[d_3 d_2]^q$ and obtain

$$m(k) = \frac{\left[\frac{d_1}{d_2} \right]^q}{1 + \left[\frac{d_1}{d_2} \right]^q + \left[\frac{d_1}{d_3} \right]^q} < \left[\frac{d_1(k)}{d_2(k)} \right]^q \quad (7.22)$$

Hence, if we can show, that $d_1^q(k)/d_2^q(k) < 1/k^q$ and $q = 2$, then $\hat{\theta}_2$ satisfies the conditions of the above lemma. The evolution of d_1/d_2 is governed by the following equation:

$$\frac{d_1(k+1)}{d_2(k+1)} = \frac{(d_2^q + d_3^q) d_1^q d_1(k)}{(d_1^q + d_3^q) d_2^q d_2(k)} \quad (7.23)$$

A straightforward calculation shows that the first term on the right hand side is strictly less than one, for any fixed $q > 0$. Hence there exists a constant $a > 1$ such that

$$\frac{d_1(k+1)}{d_2(k+1)} < \frac{1}{a} \frac{d_1(k)}{d_2(k)} \quad (7.24)$$

for all $k > 0$. This is equivalent to

$$\frac{d_1(n)}{d_2(n)} < \left[\frac{1}{a}\right]^n \cdot \frac{d_1(0)}{d_2(0)} < \left[\frac{1}{a}\right]^n \quad (7.25)$$

since $d_1(0)/d_2(0) < 1$. But this means, using (7.22), that there is a $\bar{k} > 0$ such that,

$$m(k) < \left[\frac{1}{a^q}\right]^k < \frac{1}{k^q} \quad \text{for } k > \bar{k} \text{ and any fixed } q > 0 \quad (7.26)$$

The last inequality holds because any geometric sequence decays faster than a harmonic sequence. Hence, the conditions of the lemma are fulfilled, so

$$\lim_{k \rightarrow \infty} d_2(k) = \begin{cases} 0, & q = 1 \\ \bar{d}_2 > 0, & q = 2 \end{cases} \quad (7.27)$$

By the same arguments, it can be established that also d_3 does not tend to zero if $q = 2$. Hence, only $\hat{\theta}_1$, i.e. the closest model, converges.

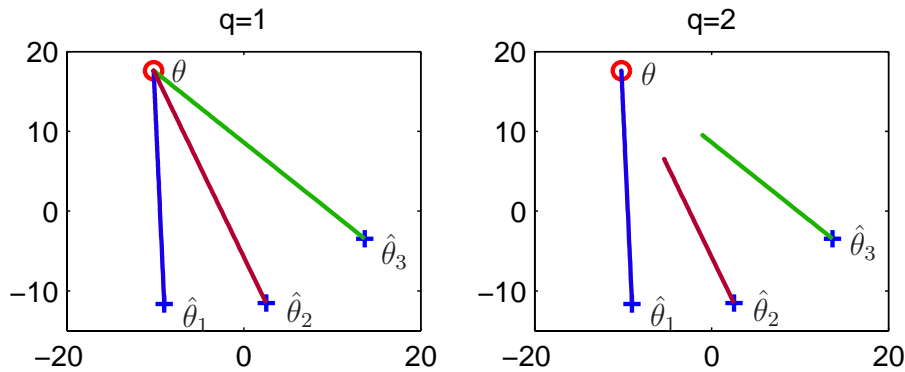


Figure 7.1: Converge of the models for different values of the parameter q

7.2.2 Elementary Properties

Even when the environment is constant the algorithm displays an important property: The models compete for the largest adaptive gain and only the closest one converges if $q = 2$. It is clear that this is a prerequisite for the convergence in a time-varying environment which involves a sorting process, i.e. no two models must converge to the same point.

Proving that the sorting process indeed takes place even when the environment switches, is substantially more difficult than the analysis performed in the time-invariant case. In the following section, some elementary properties of the algorithm are collected which are easily accessible.

Convergence to a Convex Hull

In order to motivate the algorithm, we assumed that all the models are initialized far from the plants in parameter space. Hence, none of the models is distinguished with respect to any plant and we expect all models to behave similarly. Given the set $S = \{\theta_1, \theta_2, \dots, \theta_N\}$ of plants, we define a convex hull \mathcal{H} of $S \subset \mathbb{R}^p$ as follows:

$$\mathcal{H} = \{\lambda_1\theta_1 + \dots + \lambda_N\theta_N \mid \theta_i \in S, \lambda_i \in [0, 1] \text{ and } \sum_{i=1}^N \lambda_i = 1, i = 1 \dots N\}$$

The following statement refers to models lying outside the convex hull \mathcal{H} :

Lemma 7.2 Given $i \in \Omega$ with $\hat{\theta}_i \notin \mathcal{H}$, then $\hat{\theta}_i$ converges monotonically to \mathcal{H} , i.e. $\text{dist}(\hat{\theta}_i(k), \mathcal{H}) := d_{i\mathcal{H}}(k) \rightarrow 0$ monotonically as $k \rightarrow \infty$. \square

Proof. From the algorithm we see that every model reduces its distance to an element contained in \mathcal{H} at every instant $k > 0$. The lemma follows from the definition of the distance of a model $\hat{\theta}_i$ to the set \mathcal{H}

$$d_{i\mathcal{H}}(k) := \min_{\xi \in \mathcal{H}} \|\hat{\theta}_i(k) - \xi\|$$

■

Hence, without loss of generality, all models can be initialized inside the convex hull (\rightarrow see figure 7.2).




Figure 7.2: Convex hull of the plants

Ordering of the steps

The previous statement holds for the set of all models. In this paragraph we examine the dynamics of an individual model. The quantities of interest are the distances $d_{ij} = \|\tilde{\theta}_{ij}\| =$

$\|\hat{\theta}_i - \theta_j\|$ between all the models and all the plants. If, at a given instant, the plant parameters are $\theta(k) \equiv \theta_j$ we obtain from equation (7.10),

$$d_{ij}(k+1) = d_{ij}(k) - \eta_{ij}(k)d_{ij}(k) \quad (7.28)$$

for all $i \in \Omega$ and any fixed $j \in \Omega$. As before, the plant index j is omitted whenever no particular reference to any element $\theta \in S$ is made. The exponent q in the definition of the adaptation step-size $\eta(k)$ in equation (7.11), was found to be a critical parameter. For $q > 1$, we observed competitive behavior among the models in the case where the environment is constant. This relates to a fundamental property, which can be stated as follows: The amount by which the models reduce their distances to the prevailing plant parameter is ordered according to their relative distances from that parameter. We will refer to such an amount as a “step” which is defined by:

$$s_i(k) = d_i(k+1) - d_i(k) = \eta_i(k)d_i(k) \quad (7.29)$$

Using the definition of η_i from equation (7.11) we obtain

$$s_i(k) = \frac{d_1^q \dots d_{[i-1]}^q d_i d_{[i+1]}^q \dots d_N^q}{\sum_{k=1}^N \prod_{l \neq k} d_l^q} \quad (7.30)$$

Let the indices $i \in \Omega$ be assigned such that

$$d_1 < d_2 < \dots < d_N \quad (7.31)$$

The following lemma regarding the ordering of the steps holds:

Lemma 7.3

$$d_i < d_{i+1} \Rightarrow s_i > s_{i+1} \quad \forall i \in \Omega \quad \text{iff } q > 1 \quad (7.32)$$

Proof. From equation (7.30) we have:

$$\begin{aligned} s_i &= \frac{1}{d_S} d_1^q \dots d_{i-1}^q d_i d_{i+1}^q \dots d_N^q \\ s_{i+1} &= \frac{1}{d_S} d_1^q \dots d_i^q d_{i+1} d_{i+2}^q \dots d_N^q \end{aligned} \quad (7.33)$$

where $d_S = \sum_{k=1}^N \prod_{l \neq k} d_l^q$. We obtain

$$\frac{s_i}{s_{i+1}} = \frac{d_{i+1}^{q-1}}{d_i^{q-1}} > 1 \quad \text{since } d_{i+1} > d_i \text{ and } q > 1$$

■




Figure 7.3: Evolution of the parameters for different values of q

The simulation in figure (7.3) illustrates the importance of the condition $q > 1$ in the case of three models and three plants. At every step, the formation of the models is ordered with respect to the prevailing plant parameter θ , in the sense that the ratio of the distances $\frac{d_{i+1}}{d_i} > 1$ increases:

$$\left[\frac{d_{i+1}}{d_i} \right] \Big|_{k+1} = \frac{1 - \eta_{i+1}}{1 - \eta_i} \frac{d_{i+1}}{d_i} \Big|_k = \underbrace{\frac{d_{i+1}^q \sum_{l \neq i+1} d_l^q}{d_i^q \sum_{l \neq i} d_l^q}}_{>1} \frac{d_{i+1}}{d_i} \Big|_k > \frac{d_{i+1}}{d_i} \Big|_k \quad (7.34)$$

The inequality follows from the fact that $d_{i+1} > d_i$. At any step, the algorithm performs an elementary sorting operation with respect to the current plant parameter. Our aim is to prove that the subsequent application of an elementary sorting operation to the set of models results in overall organization of the models.




Figure 7.4: The effect of a step to θ_1 in two different configurations

It is clear that a step towards the current plant parameter also affects the distances to all the other parameters. While the former results in an advantage of the closest model it is

not clear whether it contributes to the overall organization of the models. As illustrated in figure (7.4) this entirely depends on the geometry of the model–plant configuration. If θ_1 and $[\theta_2, \theta_3]$ are on opposite sides of $\hat{\theta}_i$, the fact that d_{i+1}/d_i increases at a step to θ_1 results in a reinforcement of the “pole position” of $\hat{\theta}_i$ with respect to θ_1 while $\hat{\theta}_{i+1}$ remains closer to θ_2 . In the right part of the figure the same step seems to be detrimental for the overall organization since $\hat{\theta}_i$ gains with respect to θ_1 and θ_2 . The step results in “mixing”, i.e. it is not clear, whether $\hat{\theta}_i$ eventually converges to θ_1 or θ_2 . Since no assumptions on the location of the plant are made there exists an infinite number of possible configurations. This constitutes the hard part of the proof.

Equilibrium Set

An equilibrium corresponds to the situation where the amount by which all models are updated in parameter space is zero. In other words,

$$s_{ij}(k) = 0 \quad (7.35)$$

for all $i, j \in \Omega$ and all $k > 0$. It is clear that, for (7.35) to hold, at least one of the factors in the numerator expression of (7.30) must be zero. Since S consists of N distinct values, a model cannot have zero distance to more than one plant.

Without loss of generality assume that $s_{11} = 0$ because $d_{11} = 0$, i.e. $\|\hat{\theta}_1 - \theta_1\| = 0$. This implies $d_{1j} \neq 0$ for all $j \setminus \{1\}$. Once the plant switches to θ_2 , the distance d_{i2} , $i \neq 1$ of a *different* model must be zero in order to obtain $s_{12} = 0$. The same is true for every $\theta \in S$. Since the numerator in equation (7.30) is the product of N such distances we conclude:

If *one* model $\hat{\theta}_i$ is in an equilibrium state, i.e. $d_i(k+1) = d_i(k) \forall k > 0$, then for *every* $\theta_j \in S$ there must be a model $\hat{\theta}_l$ for which $\|\hat{\theta}_l - \theta_j\| = 0$, for all $j, l \in \Omega$, i.e. for every plant there must be a model with zero distance. It follows that all models attain the equilibrium set simultaneously, i.e. they *converge simultaneously*. Note that the number of models M may exceed the number of plants N .

Lemma 7.4 If $M < N$, no model converges. If $M = N$, all models converge simultaneously and have zero distance to the plants asymptotically. Finally, if $M > N$, all models converge simultaneously but only N of them have zero distance asymptotically (\rightarrow figure 7.5). \square

The invariant set is defined as

$$E = \{d_{ij} \mid d_{f(i)j} = 0 \text{ whenever } f(i) = j, \text{ for all } i, j \in \Omega\}$$

where $f(\cdot) : \Omega \rightarrow \Omega$ is a rearrangement of the model indices such that every model has the same index as the plant to which it converges. Note that since *any* model may converge to

any plant (depending on the initial conditions) there are $N!$ possible rearrangements of the model indices.

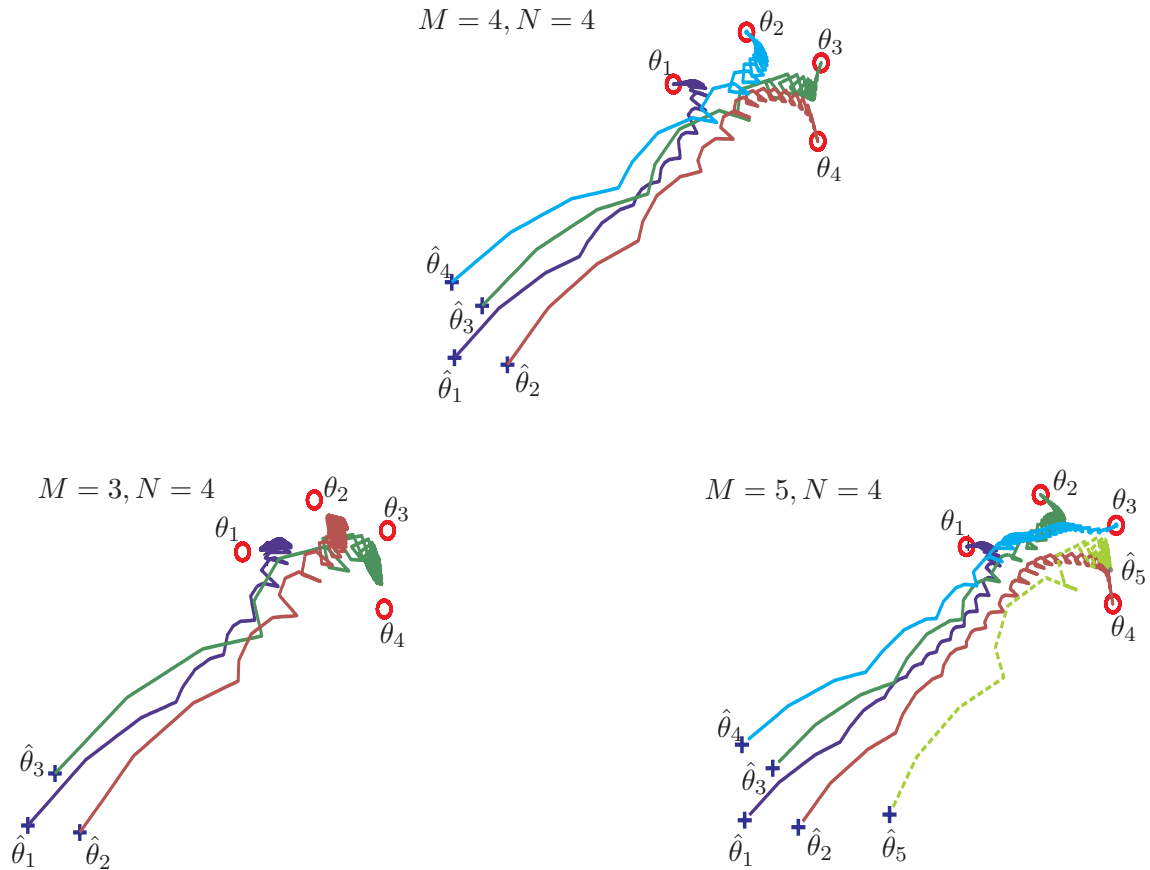


Figure 7.5: For convergence, $M \geq N$

7.2.3 Cycle vs. Instantaneous Dynamics

The dynamics of the models in parameter space is given by equations (7.10) and (7.11). Since the environment switches the right-hand side of the resulting state equations is discontinuous. Moreover, since the switching is time-dependent, the system is nonautonomous. Under the assumption that the switching is periodic, we may determine the interval T over which the environment is constant and define a new time scale where one instant corresponds to an interval of length NT which will be referred to as a “cycle”. The resulting mapping (corresponding to a Poincaré-section of the original orbit \rightarrow see e.g. [12]) is autonomous.

In view of the definition of the step-size η , we propose a mathematical description of the system using only distances. It is clear that the state of the system at a given instant can




Figure 7.6: Geometry of a model-plant configuration

be described by the distances d_{ij} between the elements of the set of plants S and the set of models \hat{S} . If $M = N$, we obtain N^2 state variables. However, at a given instant, only N of the N^2 distances are “effective states”, in the sense that their evolution is described by a difference equation of the form (7.28). If $\theta(k) \equiv \theta_j$ at a given instant, the state of the system consists of N distances $d_{1j}(k), \dots, d_{Nj}(k)$. All other distances $d_{i\Lambda}(k)$, $\Lambda = \Omega \setminus \{j\}$, are output variables which correspond to the distance of model $\hat{\theta}_i$ to all points which are not in existence at instant k . However, their role changes once $\theta(k) = \theta_l$, $l \in \Lambda$.

In summary, only N distances are effective states at a given instant but all N^2 distances assume the role of an “effective state” within a cycle. The output is a function not only of the state but depends also on the geometry of the model-plant configuration. This is illustrated by the following example.

Example 7.1 Choose $N = 2$ points $\theta_1, \theta_2 \in \mathbb{R}^2$ and locate $M = 2$ models $\hat{\theta}_1(0), \hat{\theta}_2(0)$ arbitrarily in \mathbb{R}^2 . Define $L = \|\theta_2 - \theta_1\|$. We consider the dynamics of $\hat{\theta}_1$ over one cycle $NT = [0, 2]$ where $\theta(0) \equiv \theta_1$ and $\theta(1) \equiv \theta_2$, i.e. $T = 1$. Equivalent equations hold for $\hat{\theta}_2$. When θ_1 is in existence, the dynamics are described by:

$$\begin{aligned} d_{11}(1) &= [1 - \eta_{11}(0)] d_{11}(0) \\ d_{12}(1) &= \sqrt{d_{11}^2(1) + L^2 - 2 d_{11}(1) L \cos \alpha_1(1)} \end{aligned} \quad (7.36)$$

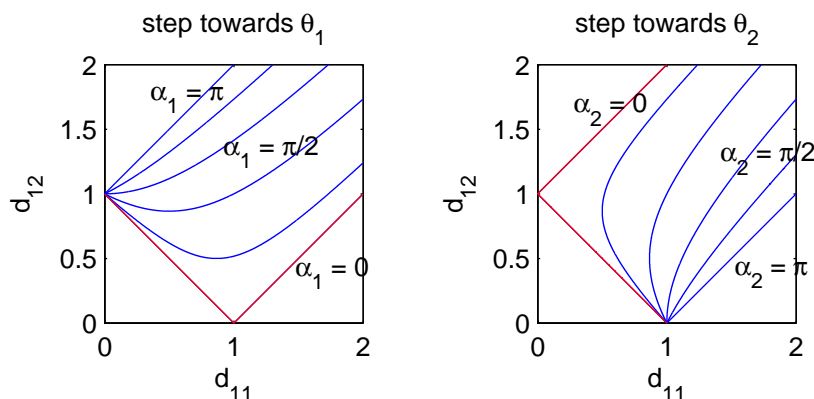


Figure 7.7: Output function for different values of the angle α_i at θ_i

Here, d_{11} is a state variable and d_{12} can be regarded as the output of the system since it is given by an algebraic equation which depends on d_{11} . In addition, the output is a function of the angle α_1 at θ_1 which, in turn, depends on the initial condition $\hat{\theta}_1(0)$, see figure (7.6). As mentioned above, geometry plays a key role in the convergence problem. A step towards θ_1 reduces the distances to both points if $\alpha > \pi/2$, while for $\alpha < \pi/2$ and $d_{11}^2 + d_{12}^2 < 1$ (i.e. inside the *Thales*-circle of diameter L) a step towards θ_1 results in an increase of distance to θ_2 . This is illustrated in figure (7.7), using $L = 1$ for simplicity. At the next instant, θ_2 is in existence and we obtain:

$$\begin{aligned} d_{12}(2) &= [1 - \eta_{12}(1)] d_{12}(1) \\ d_{11}(2) &= \sqrt{d_{12}^2(2) + L^2 - 2 d_{12}(2) L \cos \alpha_2(2)} \end{aligned} \quad (7.37)$$

In this case, d_{12} is the state and d_{11} is the output of the system. Equivalent statements about the dependency of the output value on the angle α_2 at θ_2 as above hold and are illustrated in the right part of figure (7.7). The angle α_2 , in turn, is determined by $\hat{\theta}_1(1)$, i.e. it depends upon the previous step in which θ_1 was in existence. Hence, it is possible to determine $d_{11}(2)$ from $d_{11}(0)$ and $d_{12}(1)$. However, the map is not unique since $\eta_{11}(0)$ and $\eta_{12}(1)$ have not been specified yet. From definition (7.11) it is clear that $\eta_{11}(0)$ depends not only on $d_{11}(0)$ but also on the distance of the second model, $d_{21}(0)$. The same is true for $\eta_{12}(1)$ which depends on $d_{12}(1)$ and $d_{22}(1)$. But this means that, when determining $d_{11}(2)$, the evolutions of both models $\hat{\theta}_1$ and $\hat{\theta}_2$ enter the equation. Using (7.36) and (7.37) and the corresponding equations for $\hat{\theta}_2$ it is possible to compute a mapping $f_{cyc} : \mathbb{R}^4 \rightarrow \mathbb{R}^2$:

$$[d_{11}, d_{12}]^+ = f_{cyc}([d_{11}, d_{12}, d_{21}, d_{22}]) \quad (7.38)$$

where the superscript $+$ refers to the distances after a cycle of length 2 (since $T = 1$). The simulation in figure (7.8) demonstrates that the phase trajectory of the system has the shape

of the output function (\rightarrow figure 7.7) since at a given instant only one of the distances is a state variable. Note that for different switching periods and identical initial conditions the models converge to two different equilibrium points. While $[d_{11}, d_{12}] \rightarrow [0, 1]$ in the left simulation, $[d_{11}, d_{12}] \rightarrow [1, 0]$ in the right simulation when $T = 2$. The objective is to show that the ω -limit set of the dynamical system defined by (7.38) is equal to one of the limit points.




Figure 7.8: Convergence of 2 models to the vertices of a convex set

The algebraic expression for f_{cyc} is very complex and does not offer any insight into the behavior of the corresponding dynamical system. Because of lemma 7.2 the models monotonically converge to the line, which represents the convex hull of θ_1, θ_2 . Hence, we may assume that $\hat{\theta}_1, \hat{\theta}_2$ have already been initialized on the line connecting θ_1 and θ_2 . The dependency on geometry disappears, since $\alpha_1 \equiv 0$ and $\alpha_2 \equiv 0$ for all $k > 0$. In this case it is reasonable to present an algebraic expression for f_{cyc} . Again, to keep the expression simple, we assume that $T = 1$ and $\eta_0 = 1$, where η_0 is the nominal step-size used in (7.11). In contrast to the general case, only two distances $d_{11}(0)$ and $d_{21}(0)$ are necessary to uniquely determine the solution at the end of a cycle, since $d_{12} = L - d_{11}$ and $d_{22} = L - d_{21}$. We obtain $[d_{11}, d_{21}]^+ = f_{cyc}([d_{11}, d_{21}])$:

$$d_{11}^+ = \frac{d_{11}^3}{d_{21}^2 + d_{11}^2} + \left[L - \frac{d_{11}^3}{d_{21}^2 + d_{11}^2} \right] \left[1 + \frac{[Ld_{21}^2 + Ld_{11}^2 - d_{11}^3]^2}{[Ld_{21}^2 + Ld_{11}^2 - d_{21}^3]^2} \right]^{-1} \quad (7.39)$$

$$d_{21}^+ = \frac{d_{21}^3}{d_{21}^2 + d_{11}^2} + \left[L - \frac{d_{21}^3}{d_{21}^2 + d_{11}^2} \right] \left[1 + \frac{[Ld_{21}^2 + Ld_{11}^2 - d_{21}^3]^2}{[Ld_{21}^2 + Ld_{11}^2 - d_{11}^3]^2} \right]^{-1} \quad (7.40)$$

The evolution of the models in parameter space is described by nonlinear coupled difference equations (7.39) and (7.40) in which the dependency on the environment has been removed. Note that the equation is only valid for $T = 1$. It is clear that this representation is equivalent to the “instantaneous” equations (7.36) and (7.37). However, it may offer analytical advantages, since the system is autonomous. On the other hand, the algebra involved is substantially more complex than in the case where individual steps are considered.

A central question regards the stability and region of attraction of an equilibrium point.

Example 7.2 Suppose that the models are initialized close to θ_1 , and θ_2 , i.e. $d_{11} < \varepsilon$ and $d_{21} > L - \varepsilon$. The question is whether the equilibrium point $[d_{11}, d_{21}] = [0, 1]$ is asymptotically stable. If so, we are interested in its region of attraction, i.e. we ask for the largest ε such that

$$\begin{aligned} d_{11}^+ &< d_{11} \quad \forall d_{11} < \varepsilon \\ d_{21}^+ &> d_{21} \quad \forall d_{21} > L - \varepsilon \end{aligned} \quad (7.41)$$

A straightforward calculation, in which we replace d_{11} in equation (7.39) by ε and d_{21} by $L - \varepsilon$ yields $\varepsilon < \varepsilon_0 = 0.4$. Figure (7.9) illustrates that $d_{11}^+ - d_{11} < 0$ for all $\varepsilon < \varepsilon_0$ and $d_{21}^+ - d_{21} > 0$ for all $\varepsilon < 3/2 \varepsilon_0$. Hence, f_{cyc} is a contraction mapping whenever $\varepsilon < \min\{\varepsilon_0, 3/2 \varepsilon_0\} = \varepsilon_0$ and $[0, 1]$ is asymptotically stable with region of attraction $A = \{[d_{11}, d_{21}] \mid d_{11} < \varepsilon_0, d_{21} > 1 - \varepsilon_0\}$. It is clear that the computations are substantially more complex when the period over which the environment is constant is $T > 1$ or when the number of plants and models is $N > 2$.




Figure 7.9: Relative decrease (increase) of d_{11} (d_{21}) as a function of ε

The model update equations (7.28) define a dynamical system. Its positive orbit is described by the sequence $\{d_k\}_{k=0}^{\infty}$ where $d_k \in \mathbb{R}^{N^2}$. When considering cycles, we are looking at a subsequence of the form $\{d_{k_t}\}_{t=0}^{\infty}$ with $k_t \in \mathbb{N}$, $k_t \rightarrow \infty$ as $t \rightarrow \infty$ and $(k_t - k_{t+1}) = T^* = NT$.

The objective is to show that $\{d_{k_t}\}_{t=0}^{\infty}$ converges to an element $\bar{d} \in E$ in the ω -limit set E , i.e. one of the $N!$ equilibria defined above. For if this is the case, the “full” sequence $\{d_k\}_{k=0}^{\infty}$ also converges to $\bar{d} \in E$ since \bar{d} is invariant, i.e. the system initialized at \bar{d} remains there for all k even if $k \neq k_t$. This is clear from our discussion in section 7.2.2. In particular, there is no closed orbit contained in the finite set E .

If some property \mathcal{P} can be proven for $\{d(k_t)\}$ it is only because another property \mathcal{P}' which is related to \mathcal{P} holds for the original sequence $\{d(k)\}$. In order to understand (and exploit) the relationship between \mathcal{P} and \mathcal{P}' , considerable insight into the qualitative behavior of the algorithm is needed. This is provided in the next section.

7.2.4 Qualitative Behavior of the Algorithm

Extensive simulation studies have been carried out which resulted in asymptotic convergence of N models to N plants independent of the value of N (\rightarrow figure 7.11). In this section, some of the characteristic features of the algorithm are collected. The central observation is that the configuration of the models gradually becomes more ordered. This property was already shown to depend on lemma 7.3, i.e. the fact that the steps of the models are ordered according to their relative distances to the prevailing plant parameter. Since the equation describing the dynamics of the models become very complex when N is large, an effort was made to establish convergence using only lemma 7.3 and pure logic. Although, the approach turned out to be inconclusive it provides a great deal of insight into the different stages of the ordering process.

Convergence Logic

Knowing that the plant parameter assumes the value of every element of S at least once in an interval of finite length T^* the objective is to determine the logic behind the ordering process which ultimately leads to convergence. Intuitively, “order” refers to the fact that all the models attain a distinguished position with respect to different elements in S .

Definition A model $\hat{\theta}_i$ has a distinguished position with respect to a point $\theta_j \in S$ if there is a $\bar{k} > 0$ such that

$$\|\hat{\theta}_i(k) - \theta_j\| = \min_{s \in \Omega} \|\hat{\theta}_s(k) - \theta_j\| \quad \forall k > \bar{k}$$

$\hat{\theta}_i$ is the closest among all the models with respect to θ_j . The key point to note is that it remains closest for all $k \in (\bar{k}, \infty)$. The definition is prompted by an observation that has

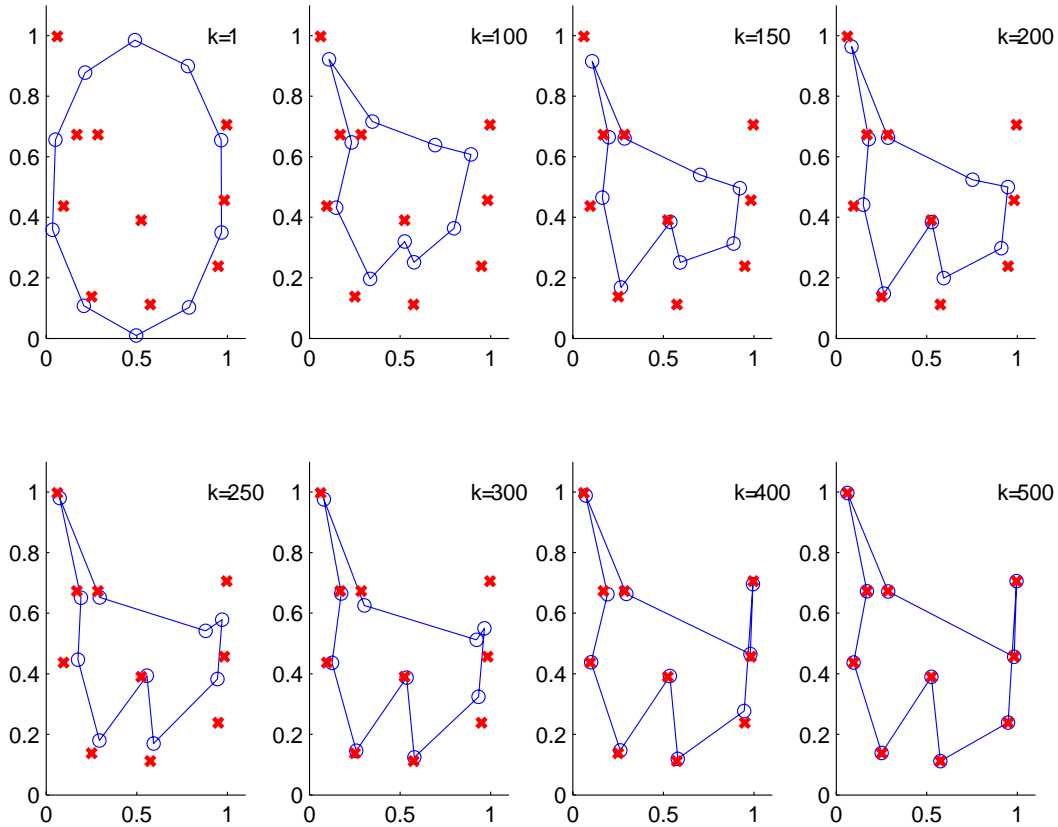


Figure 7.10: Convergence of 10 models to 10 plants (periodic environment with $T = 2$)

been made through simulation studies that if a model remains closest to a particular point it also converges to that point. The question is, whether a model $\hat{\theta}_i$ which is distinguished with respect to θ_j at instant k will be replaced by another model $\hat{\theta}_{i+1}$ at any future instant of time. The fact that the role of being distinguished passes from one model to another will be referred to as an “overtake”. It is intuitively clear that there can be no model $\hat{\theta}_i$ which remains in a distinguished position with respect to more than one point. For, as figure (7.7) illustrates, the closer $\hat{\theta}_i$ approaches one point, the more it increases its distance from the other points. Hence, if $\hat{\theta}_i$ has multiple distinguished positions initially, there must be an overtake. The critical step in proving self-organization is to show that there will be only a finite number of overtakes. In other words, we have to show that there will be a stage when every model remains in a distinguished position for all future instants of time. We summarize:

Fact 1: If $\hat{\theta}_i$ is distinguished with respect to θ_j , it converges to θ_j . □

This can be seen as follows: The fact that $\hat{\theta}_i$ is distinguished implies (by lemma 7.3) that its step towards θ_j is the largest among the models: $s_i > s_l, \forall l \in \Omega/\{i\}$, where s_i was defined

in equation (7.29). It does not yet imply that, at the end of a cycle, $\hat{\theta}_i$ ends up closer to θ_j than it was before. This will be true only eventually, when the number of possibilities for $\hat{\theta}_i$ to depart from θ_j has decreased. Suppose (in contradiction) that the distance $d_{ij} = \|\hat{\theta}_i - \theta_j\|$ increases.

Fact 2: d_{ij} does not increase for all $k > 0$, unless $\hat{\theta}_i$ is overtaken with respect to θ_j . \square

An overtake with respect to θ_j means that another model $\hat{\theta}_f$ appears with the property that $d_{fj} < d_{ij}$. Suppose that d_{ij} increases but no such model $\hat{\theta}_f$ appears. The increase of d_{ij} is due to the steps s_{il} towards all the other points θ_l , $l \in \Omega/\{j\}$. The exact amount of the increase is given by the projection of s_{il} onto the radius of the circle around θ_j which we denote by σ_{il} , see figure (7.11). This means that the distance of $\hat{\theta}_i$ to the set $S/\{\theta_j\}$ decreases in some sense. However, $\hat{\theta}_i$ cannot be closest to all the points θ_l , $l \in \Omega/\{j\}$. It follows, that σ_{il} is bounded by the step of some other model which is closer to θ_l than $\hat{\theta}_i$ is. As $k \rightarrow \infty$, the step $s_{ij} = \sigma_{ij}$ to θ_j is the only quantity which is not bounded by the distance of any other model. Hence, there must be a stage k_0 after which d_{ij} decreases.




Figure 7.11: Evolution of $\|\hat{\theta}_i - \theta_j\|$ over a cycle

If, on the other hand, $\hat{\theta}_f$ takes over with respect to θ_j and *remains* close to it, d_{ij} cannot decrease but the same argument as before holds for $\hat{\theta}_f$, i.e. there must be a stage after which d_{fj} decreases. The most adverse situation is when $\hat{\theta}_i$ and $\hat{\theta}_f$ interchange their role of being the closest model with respect to θ_j over a number of cycles. In such a case, we have to show that the oscillation cannot persist. In other words, the conditions for the existence of

overtakes have to be investigated.

Fact 3: A sufficient condition for an overtake is that θ_i is closest to more than one point. \square

Actually, a necessary condition would be more useful to prove the non-existence of overtakes. However, such a condition could not be established since it requires a more detailed description of the model configuration (such as the state-space description discussed in section 7.2.3). Much effort was spent in the attempt to establish that an overtake can lead to at most $N - 2$ future overtakes. The idea was that if $\hat{\theta}_f$ and $\hat{\theta}_i$ interchange their roles, they must end up close to each other. Additional information is provided by the fact, that an overtake can only happen indirectly, i.e. at a step to a point contained in $S/\{\theta_j\}$. Supposing that $\hat{\theta}_f$ took over $\hat{\theta}_i$ at a step to θ_k , the conjecture was that –after the overtake– the pair $\hat{\theta}_f - \hat{\theta}_i$ is close to the line connecting θ_j and θ_k .




Figure 7.12: Overtakes

Figure (7.12) displays a simulation of the case $N = 3$. The models are initialized ($k = 0$) such that $\hat{\theta}_1$ (marked by an \times) is closest to θ_1 and θ_3 . Hence, by fact 3, there must be an overtake, which indeed takes place at instant of time $k = 150$. $\hat{\theta}_2$ (marked as \bullet) is now closest to θ_1 . At $k = 330$, $\hat{\theta}_3$ (marked as \circ) takes over $\hat{\theta}_2$ which completes the total number of $N - 1 = 2$ overtakes. On the left side of the figure, the index of the model which is closest to θ_1 is displayed as a function of time. A cycle consists of 3×10 steps and is marked by a dotted line (—).

The arguments provided by pure logic are weak and qualitative in nature. In particular, they fail to establish that the degree of organization in the model-plant configuration increases. It is clear that a quantitative “measure of organization” is needed which adequately describes the fact that the models gradually attain distinguished positions. Such a measure could be the distance between the models, or the ratio of the distances of different models to a

point. In the case of 3 models, an intuitive quantity is given by the area of the triangle formed by the 3 models which attains a maximum when the models converge. However, a typical phenomenon observed in simulation studies is that it is always possible to find initial conditions such that the chosen quantity behaves in a non-monotonic fashion initially (\rightarrow figure 7.13). Coming back to figure (7.12), it can be verified that the instants at which the area decreases (roughly) coincide with the instants at which an overtake takes place.




Figure 7.13: Self Organization in the 3/3 case: The area of the inner triangle (as well as its perimeter) decreases initially and then increases monotonically up to the point where the inner and outer triangle match.

Nonmonotonic Convergence

Most quantities which provide an intuitive measure of the overall organization of the models turned out to evolve in a nonmonotonic fashion. However, it can be shown that in the neighborhood of an equilibrium point, the solutions converge monotonically. In order to illustrate this we define a (reduced) state-space consisting of the diagonal elements d_{ii} of the $N \times N$ matrix $[d_{ij}]_{i,j \in \Omega}$ of all possible distances between the models and the plants (\rightarrow figure 7.15). Every model $\hat{\theta}_i$ is associated with a point θ_i in parameter space and is initialized in an ε -neighborhood of θ_i . If ε is sufficiently small, the distance of the models to their respective points decreases, i.e. at the end of a cycle we have $d_{ii}^+ < d_{ii}$ for every $i \in \Omega$.

As seen in section 7.2.3, the computation of such an ε requires knowledge of the map $d_{ii}^+ = f_{cyc}(d_{ij})$ $i, j \in \Omega$ which depends on all N^2 distances. In order to avoid the algebraic

complexity of determining f_{cyc} , the following simplification of the algorithm is proposed. Instead of updating the models at every instant of time, their location is fixed during a cycle. The adjustment vectors which result at every instant are summed up and used to update the location of the model only after the cycle is complete. The modification can be thought of as a worst-case scenario. For, if indeed every step contributes to convergence, the improvement is less pronounced if only the average over the steps within a cycle is used. The algebraic complexity, in turn, is reduced substantially.




Figure 7.14: Monotonic convergence

Example 7.3 The models are initialized at a distance ε from the corners of the triangle in figure (7.14) with side lengths L_1, L_2, L_3 respectively. The condition for $d_{ii}^+ < d_{ii}$ is that the step towards θ_i is larger than to any other point $\theta_j, j \neq i$. If the models are updated at the end of a cycle, the condition can be restated in terms of a single step, since the configuration is symmetric, i.e. $d_{ii} = \varepsilon$ for all $i \in \Omega$. In figure (7.14), $s_2^* \geq s_2$ and $s_3^* \geq s_3$ denote the largest possible steps of the models $\hat{\theta}_2$ and $\hat{\theta}_3$ towards θ_1 . Because of the symmetry, the relative steps s_2^*/s_1 and s_3^*/s_1 are equal to the maximum relative amounts by which the first model $\hat{\theta}_1$ departs from θ_1 (at a step towards θ_2 and θ_3),

$$\frac{s_i^*}{s_1} = \frac{s_{1i}}{s_{ii}} = \frac{\varepsilon}{L_i - \varepsilon} \quad i = 2, 3 \quad (7.42)$$

The last equality follows from the elementary properties of the algorithm and $q = 2$. Using

an upper bound s^* on this amount we require that the following inequality holds:

$$s_1 > 2s^* \quad s^* = \max\{s_2^*, s_3^*\} \tag{7.43}$$

Since $L_3 = \min\{L_1, L_2, L_3\}$, the closest competitor of $\hat{\theta}_1$ is $\hat{\theta}_3$, i.e. $s^* = s_3$. We know that,

$$\frac{s_1}{s_3} = \frac{d_3}{d_1} = \frac{L_3 - \varepsilon}{\varepsilon} \tag{7.44}$$

Evaluating (7.43) gives $\varepsilon < L_3/3$. □

The example illustrates that there is an ε^* such that convergence is monotonic for all $\varepsilon < \varepsilon^*$. If we increase ε beyond ε^* , the distances d_{ii} , $i \in \Omega$ evolve in a non-monotonic fashion. For certain initial conditions, the solutions do not tend to the origin in $[d_{ii}]$ -space. From the discussion in section (7.2.2) it is clear that there are 6 equilibria, corresponding to $N!$ possible permutations of the models and plants. As an example, θ_4^* in figure (7.15) corresponds to the situation where $\hat{\theta}_1$ converges to θ_3 , $\hat{\theta}_2$ to θ_2 and $\hat{\theta}_3$ to θ_1 , i.e. $d_{11} = 1, d_{22} = 0$ and $d_{33} = 1$.




Figure 7.15: $[d_{11}, d_{22}, d_{33}]$ - space

If the models in the above example are initialized far from the vertices of the triangle, it is not clear where they are going to converge to. The regions of attractions of the 6 equilibria overlap, i.e. the neighborhood of a given $\hat{\theta}_1(0)$ with $\hat{\theta}_1 \rightarrow \theta_1^*$ contains initial conditions which give rise to trajectories converging to θ_i^* , $i \neq 1$. The trajectories in $[d_{11}, d_{22}, d_{33}]$ -space are not unique, since $N^2 = 9$ distances are needed to locate the models in parameter space. Hence it is impossible to display the different regions in figure (7.15). However, we assume that they have complex (maybe fractal) borders.

The following experiment was performed to illustrate the behavior of the algorithm if the models are not initialized close to any equilibrium point. In this experiment, the models

are arranged at the vertices of a small equilateral triangle (\rightarrow see figure 7.16) which is shifted along an even grid in parameter space, covering a total of 120 000 positions each corresponding to a different initial condition. The algorithm is started but only the evolution of one model $\hat{\theta}_1$ is traced. If $\hat{\theta}_1$ ends up at θ_1 the corresponding initial condition is marked by \bullet , if it ends up at θ_2 , \bullet is used and if it is θ_3 the initial condition is marked by \bullet in figure (7.17).




Figure 7.16: Experiment: models form a small equilateral triangle

The highly erratic nature of convergence appears to be similar to the one discovered in an experiment performed by (Peitgen, 1986 [65]), the so-called “magnetic pendulum”: The movement of an oscillating iron ball (corresponding to the first model $\hat{\theta}_1$) over three magnets is studied and the regions of attraction of each of the magnets (corresponding to one of the above equilibria) is determined. It is seen that if the iron ball is released in the neighborhood of one magnet, it comes to rest over the same magnet. If, in turn, it is not close to a particular magnet it may stop over any of them. The magnetic pendulum and the convergence problem studied here have the common property that complex dynamics arise from a competitive situation. In the first case, an iron ball is subject to multiple competing forces. In the second, the plants compete for one of the models contained in the configuration $[\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3]$.

The regions of attractions for the magnetic pendulum are displayed in figure (7.18). They possess a beautiful symmetric structure which has been studied in the literature [66]. When comparing this to figure (7.17) it should be kept in mind that the latter is the result of a projection of a 9-dimensional space onto the plane. Moreover, the vector-field is time-varying and depends on the location of three models $\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3$. Hence, from a dynamical systems point of view the convergence problem investigated here is substantially more complex.




Figure 7.17: Regions of attraction, $\hat{\theta}_1 \xrightarrow{\bullet} \theta_1$, $\hat{\theta}_1 \xrightarrow{\bullet} \theta_2$, $\hat{\theta}_1 \xrightarrow{\bullet} \theta_3$

The fact that no simple borders exist between the regions of attraction of the $N!$ isolated equilibria suggests that the models are attracted non-uniformly, i.e. the speed by which a model approaches an equilibrium point is not uniquely determined by its initial condition but varies with time. Intuitively, if $\hat{\theta}_1$ lies in the region of attraction of θ_1 it initially seems to come close to the other vertices of the triangle before decreasing its distance to θ_1 . This is in line with our previous observation, that convergence is non-monotonic.

Existence of a special point M^*

When the models $\hat{\theta}_i$ are initialized inside an ε^* -neighborhood of θ_i , they are already in a “distinguished” position in the sense of the definition in chapter 7.2.4. In other words, the models have departed from each other and the sorting process is almost over. Let us consider the reverse situation, i.e. the models are initialized at the same point with zero distance to each other: $\hat{\theta}_1(0) = \hat{\theta}_2(0) = \dots = \hat{\theta}_N(0) := \hat{\theta}(0)$. An interesting phenomenon can




Figure 7.18: Magnetic Pendulum: Regions of Attraction

be observed in this case. As expected, the models do not converge to any of the equilibria. This is clear from the fact that due to their initialization, the “stacked” models behave like a single model $\hat{\theta}$, and hence $M = 1$ and $M < N$ (\rightarrow lemma 7.4). It is seen, however, that there exists a unique limit cycle to which the models $\hat{\theta}(k)$ converge as $k \rightarrow \infty$. This cycle has the same period NT as the time-variation of the plant. Hence, if we look at the dynamics after a cycle (i.e. its Poincaré section), a characteristic point M^* is obtained to which *all* models $\hat{\theta}(k) \in \mathbb{R}^p$ converge. The importance of M^* is that it is the same for all initial conditions $\hat{\theta}(0) \in \mathbb{R}^p$. In figure (7.19), the convergence of $\hat{\theta}$ to M^* is illustrated displaying both the instantaneous solution (marked by $-$) and the cycle (marked by $-$). Further, even if the models are initialized with a nonzero but small distance from each other they invariably seem to approach M^* before converging to the vertices of the triangle. On the right side of figure (7.19), the trajectories of $\hat{\theta}_1$ are shown which come close to M^* before converging. The corresponding trajectories of $\hat{\theta}_2$ and $\hat{\theta}_3$ are hidden. Two stages can be distinguished in the convergence process. In the initial stage, the models merely “position” themselves (in

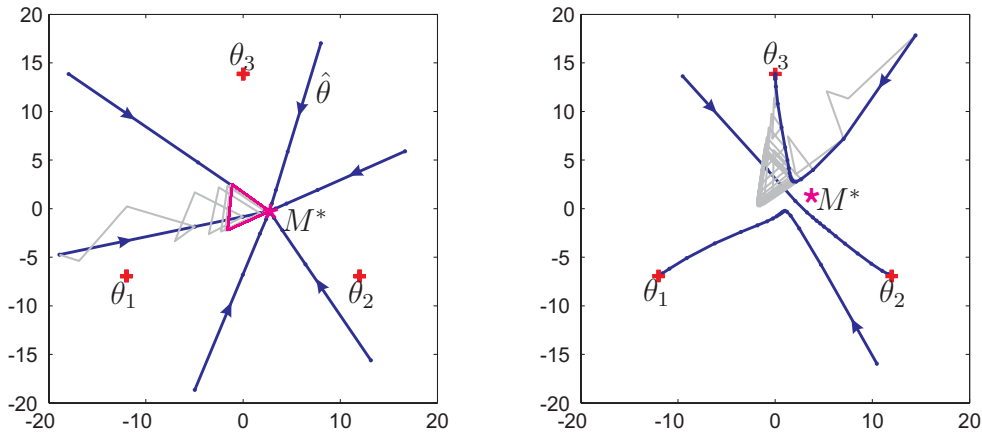


Figure 7.19: Existence of a characteristic point M^*

the neighborhood of M^*) without noticeably increasing the order of the configuration. In the second stage, the models diverge from M^* thereby approaching the plants located at the vertices of the triangle. A similar observation can be made for systems involving $N > 3$ plants and models. It suggests that M^* is a saddle-type equilibrium point. The models first follow a stable direction leading to M^* and then are repelled from M^* along the unstable direction.




Figure 7.20: Convergence in two stages

Figure (7.20) illustrates that, with fixed $\hat{\theta}_2$ and $\hat{\theta}_3$, a small variation of the initialization $\hat{\theta}_1(0)$ causes the first model $\hat{\theta}_1$ to converge to three different plants. The evolution of the $\hat{\theta}_1$ -trajectory depends sensitively on initial conditions but invariably approaches M^* before tending to one of the equilibria. The right-hand side of the figure shows that the models

$[\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3]$ also collectively approach M^* and expand only after attaining a neighborhood of M^* .

7.2.5 Restriction to a 1-dimensional subspace

So far, no restrictions have been imposed on the location of the models and the plants. In this section, we assume that the plants lie on a straight line in \mathbb{R}^p on which all the models have been initialized. This simplified situation does have relevance for the dynamic case, provided that J_i in (7.9) is defined in terms of the instantaneous error. This will be elaborated on in section 7.3.1. The motivation comes from the fact that in a standard parameter estimation scheme, the parameters are updated along the regression vector formed by measured inputs and outputs of the system. If there are multiple adaptive models they all use the same vector. Only the length of the vector differs according to the respective errors of the models. With an appropriately chosen J_i , the plants lie on a straight line in p -dimensional space and the dynamics of the models can be studied in a 1-dimensional subspace defined by that line.




Figure 7.21: 1-dimensional subspace

It is clear that the restriction to the line reduces the variety of model-plant configurations. In particular, the output equations which reflect the dependence on geometry (see example 7.1) are simplified. Still, the algebraic complexity of the equations describing the dynamics over a cycle remains, as seen in the 2/2-case (two models and two plants) in section 7.2.3. A key point to note is stated in the following lemma.

Lemma 7.5 In a 1-dimensional subspace, the system has only one equilibrium point. \square

Proof. From lemma 7.3, we know that the steps of the models are ordered. This implies that two models on a line cannot cross each other. Hence, the models converge in the order in which they have been initialized. With reference to figure (7.21), the equilibrium is attained whenever $\|\hat{\theta}_i - \theta_i\| = 0$, $i = 1, \dots, 4$, i.e. in d_{ii} -space, the only equilibrium is the origin.

Note that –despite the geometric simplicity– overtakes, as defined in section 7.2.4 do take place. In figure (7.21), $\hat{\theta}_4$ is close to both θ_3 and θ_4 . Hence, according to fact 3 in the above

section an overtake (effectuated by $\hat{\theta}_3$) with respect to θ_3 must occur. However, $\hat{\theta}_3$ can never be closer to θ_4 than $\hat{\theta}_4$, i.e. the order is fixed. Even though the complexity of the problem has been reduced drastically, the key problem –consisting of the nonmonotonic behavior of the algorithm– remains.

2/2 case \rightarrow 3/3 case

A fact frequently encountered by researchers in systems theory is that the hardest theoretical questions arise at the transition from 2nd-order to 3rd-order systems. In some sense, this is found to be the case even in this problem. Strictly speaking, the state of the system depends upon N^2 distances, where N is the number of models and plants. However, as discussed earlier, only N distances are “effective states” at a given instant. It is therefore reasonable to refer to N as the order of the system.

2/2 case: Initialize 2 models $\hat{\theta}_1, \hat{\theta}_2$ at arbitrary locations on the line connecting θ_1 and θ_2 . The state of the system is given by four distances, $d_{ij}(k) = \|\hat{\theta}_i(k) - \theta_j\|$, $i, j \in \{1, 2\}$. Hence, in view of the above comment, we refer to this as the second-order case.




Figure 7.22: 2/2–case

We define $x(k) = d_{21}(k) - d_{11}(k) = d_{12}(k) - d_{22}(k)$, where the last equal sign is due to the fact that the models are on the line. Without loss of generality, assume $d_{11}(0) < d_{21}(0)$, i.e. $x(0) > 0$. If the environment is periodic with period $T = 2$ and $\theta(0) \equiv \theta_1$ and $\theta(1) \equiv \theta_2$ we obtain:

$$\begin{aligned} x(1) &= x(0) + [\eta_{11}(0)d_{11}(0) - \eta_{21}(0)d_{21}(0)] \\ x(2) &= x(1) + [\eta_{22}(1)d_{22}(1) - \eta_{12}(1)d_{12}(1)] \end{aligned} \quad (7.45)$$

From (7.11) we have (using $q = 2$),

$$\Delta x = \eta_{11}d_{11} - \eta_{21}d_{21} = \frac{d_{11}d_{21}(d_{21} - d_{11})}{d_{11}^2 + d_{21}^2} \geq 0 \quad \text{if } \theta(k) = \theta_1 \quad (7.46)$$

$$\Delta x = \eta_{22}d_{22} - \eta_{12}d_{12} = \frac{d_{12}d_{22}(d_{12} - d_{22})}{d_{12}^2 + d_{22}^2} \geq 0 \quad \text{if } \theta(k) = \theta_2 \quad (7.47)$$

It is seen that $\Delta x = 0$ whenever $d_{ij} \in E = \{d_{11} = 0 \wedge d_{22} = 0\}$ which is equivalent to the fact that $x = \|\theta_1 - \theta_2\| = L$. Hence, $\{x(k)\}$ is a monotonically increasing sequence bounded

above by L . Suppose that $\lim_{k \rightarrow \infty} x(k) = l < L$. Then $x(k+1) = l + \Delta x(k)$ with $\Delta x(k) > 0$ for all $d_{ij} \in \mathbb{R}_0^+ \setminus E$. Hence, $\lim_{k \rightarrow \infty} x(k) = L$. We see that the distance x between the models increases at every step by a finite amount independent of which plant is in existence. Thus, $L - x(k)$ serves as a Lyapunov function by which the convergence of the models to E can be established.

3/3-case: The quantities of interest are the distances between the plants $L_1 = \|\theta_1 - \theta_2\|$ and $L_2 = \|\theta_2 - \theta_3\|$ as well as the distances between the models $x_1 = d_{21} - d_{11} > 0$ and $x_2 = d_{31} - d_{21} > 0$. By the same arguments as in the 2/2-case, x_1 and x_2 are monotonically increasing as long as there is no plant in between two models. However, by the very fact that the distance between the model increases, there must be a stage where $\hat{\theta}_2$ and $\hat{\theta}_3$ are on different sides of θ_2 . This is illustrated in figure (7.23).




Figure 7.23: 3/3-case

When θ_2 is in existence, only x_1 increases whereas x_2 decreases by an unknown amount. If $\hat{\theta}_2$ remains to the left of θ_2 , x_1 is monotonically increasing and the arguments of the 2/2-case may be applied to prove convergence of $\hat{\theta}_1$ to θ_1 and $\hat{\theta}_2$ to θ_2 . From the discussion in chapter 7.2.2, it is clear that this implies that also $\hat{\theta}_3$ must converge, since all models converge simultaneously. However, whether $\hat{\theta}_2$ crosses θ_2 or not depends on the distances of all three models to θ_3 since a crossing can only occur when θ_3 is in existence. In the adverse situation, $\hat{\theta}_2$ crosses and remains to the right of θ_2 even at a step towards θ_1 . Again, if it were possible to establish that $\hat{\theta}_2$ remains there *forever*, the same arguments as above could be used, to show that x_2 is the quantity which monotonically increases. The critical point is that while such an argument may exist, it cannot be established using the distance between the models. Hence, the ideas used for the second-order system cannot be transferred to the third order case and $V = L_i - x_i(k)$, $i = 1, 2$ fails to be a Lyapunov-function.

Monotonic quantity

The above discussion reinforces the fact that a monotonic quantity is needed in order to establish convergence. The simulations of the previous chapter provided evidence of the existence of a critical point M^* . In view of its location near the center of gravity of the plants

it may be regarded as a “center of convergence” which temporarily attracts all models. It furthermore marks an intermediate stage which separates the contraction phase (where all models accumulate near M^*) from the expansion phase (where the models approach their respective plants). The question addressed in this paragraph is whether such a critical point M^* exists even on the line connecting the plants and if so, whether it helps us to overcome the difficulties caused by the apparent non-monotonicity of all quantities relevant for the proof of convergence.

In order to determine M^* , all models have to be initialized at the same point $\hat{\theta}(0)$, i.e. on top of each other. The resulting trajectory is referred to as the $\hat{\theta}_{M^*}$ -trajectory (knowing that $\hat{\theta}_{M^*} \rightarrow M^*$). The evolution of $\hat{\theta}_{M^*}$ may be uniquely described by the distance $d_{M^*}(k) = \|\theta_1 - \hat{\theta}_{M^*}(k)\|$. Since the models have zero relative distance, they behave like a single model and the update equation is linear:

$$d_{M^*}(k+1) = \left[1 - \frac{\eta_0}{N}\right] d_{M^*}(k) \quad (7.48)$$

where N is the number of models. If $N = 3$ and the environment switches at every instant




Figure 7.24: Evolution of $\hat{\theta}_{M^*}$

of time, we may determine $\hat{\theta}_{M^*}^+$ at the end of a cycle as

$$d_{M^*}^+ = \left[1 + \frac{1}{3}\eta_0^2 - \eta_0 - \frac{1}{27}\eta_0^3\right] d_{M^*} + \frac{1}{3}\eta_0 L_2 - \frac{1}{9}\eta_0^2 L_1 + \frac{2}{3}\eta_0 L_1 \quad (7.49)$$

L_1, L_2 are the distances among the plants. The eigenvalue of the linear first-order equation (7.49) lies strictly inside the unit circle for all $\eta_0 > 0$. Hence $d_{M^*} \rightarrow \bar{d}_{M^*} = \|\theta_1 - M^*\|$ without overshoot, where

$$\bar{d}_{M^*} = \frac{3L_2 - \eta_0 L_1 + 6L_1}{\frac{1}{3}\eta_0^2 - 3\eta_0 + 9}$$

marks the critical point M^* , to which all trajectories converge if the models are initialized on top of each other.

Let us now consider an arbitrary initialization of the models, where $\hat{\theta}_1(0) \neq \hat{\theta}_2(0) \neq \hat{\theta}_3(0)$. In the previous section an attempt was made to define a Lyapunov-function in terms of the distances x_1 and x_2 between the models. However, the approach was found to be inconsistent since neither of the quantities evolved in a monotonic fashion *for all* initial conditions. This is re-examined in the following experiment. The quantity of interest is defined to be $x = d_{31} - d_{11}$, i.e. the distance between the two “boundary” models $\hat{\theta}_3$ and $\hat{\theta}_1$ (\rightarrow see figure 7.25). The objective is to see whether x or a quantity related to x increases in a monotonic fashion such that the same argument as in the above $2/2$ -case can be used to establish convergence. The step towards the intermediate plant θ_2 is again found to be critical. If the associated model $\hat{\theta}_2$ is sufficiently close to θ_2 it provides an upper bound on the update steps of the boundary models towards θ_2 . It is seen, though, that there is a large number of initial conditions $\hat{\theta}_2(0)$ for which x decreases when evaluated after a cycle. This is not surprising

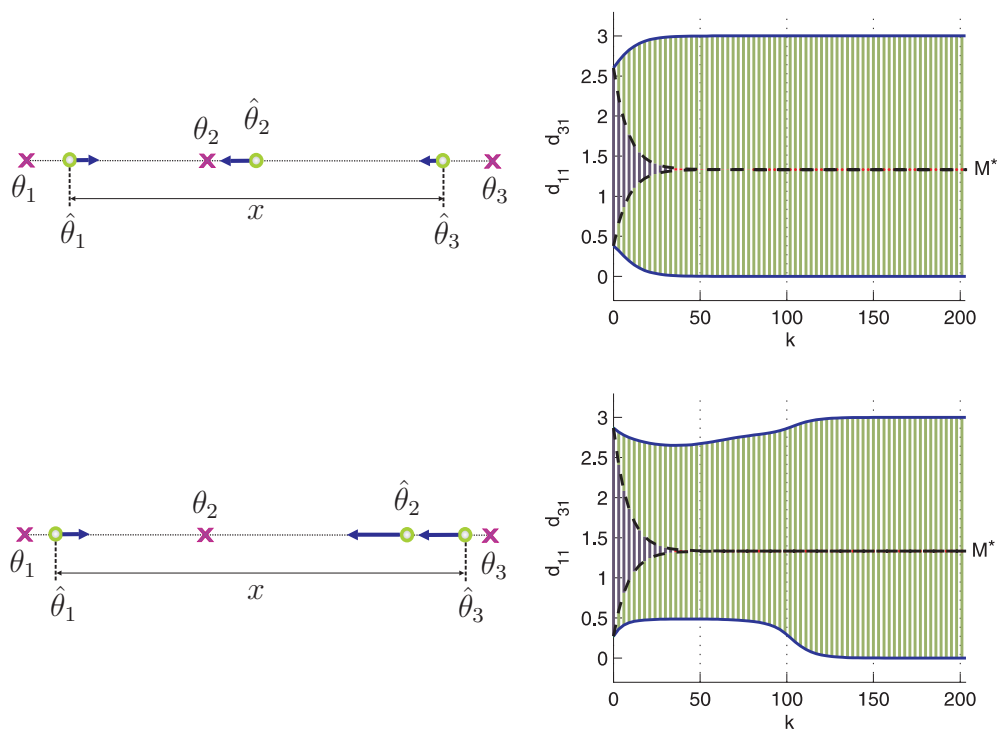


Figure 7.25: Contraction and separation part (of the convergence process)

since from the previous discussion we know that convergence takes place in two stages. If the models are initialized in an ϵ -neighborhood of their respective plants, the first stage is less pronounced, since the “degree of organization” among the models is already high. This

can be observed in the upper part of figure (7.25) – the models directly enter the expansion phase and x increases monotonically. In the lower part, the two models $\hat{\theta}_2$ and $\hat{\theta}_3$ are close to each other while $x_1 = d_{21} - d_{11}$ is large. This “unbalanced” initialization clearly corresponds to a low degree of organization. In such a case, the models are attracted by M^* before expanding, which accounts for the decrease of x in figure (7.25). We ask the question how this situation compares to the case where the models start with zero organization, i.e. zero distance among each other. To this end, two θ_{M^*} -trajectories (defined above) are initialized at $\hat{\theta}_1(0)$ and $\hat{\theta}_3(0)$. In other words, we obtain $\hat{\theta}_{1M^*}$ and $\hat{\theta}_{3M^*}$ which start at the same points as the trajectories of the boundary models. The idea is that if the actual model configuration is only slightly better than the stacked configuration, the corresponding trajectories must diverge from $\hat{\theta}_{iM^*}$, $i \in \{1, 3\}$, in a monotonic fashion. If the initial organization among the models is poor (but nonzero) it is plausible that the models first follow $\hat{\theta}_{iM^*}$ and then gradually separate. While the separation of the models (measured by means of the distance $x = d_{31} - d_{11}$) occurs only once the models have reached a neighborhood of M^* , the divergence from $\hat{\theta}_{iM^*}$ takes place independent of the model location. $\hat{\theta}_{iM^*}$, $i \in \{1, 3\}$, serve as worst-case “reference” trajectories for the boundary models $\hat{\theta}_1$ and $\hat{\theta}_3$. No other trajectory converges to M^* at a faster rate. But this implies that the distance between $\hat{\theta}_{iM^*}$ and $\hat{\theta}_i$ increases at every step. Hence if we subtract $x_r = \|\hat{\theta}_{3M^*} - \hat{\theta}_{1M^*}\|$ from x , we obtain a **monotonically increasing quantity**. This is illustrated in the right column of figure (7.25), where the dark-shaded part corresponds to the distance x_r between the reference trajectories. Intuitively, we are subtracting the part that is due to the contraction phase (of the convergence process) from the evolution of x .

The qualitative description that we have given so far is the quintessence of a large number of numerical experiments. It provides the basis for a quantitative treatment of the problem. Such an analysis has been initiated but was incomplete at the time of the preparation of the thesis and, hence, is not included here. The main difficulty is due to the fact that the dynamics of the models is governed by a composition of N maps each one consisting of N nonlinear coupled difference equations. Hence, even in a low-order case ($N = 3$) the algebraic expression for the evolution of the models cannot be handled easily and does not offer any structure unless it is appropriately simplified. It is this simplification that needs further investigation.

The key result in our study of static systems is that the interplay between the “contraction” and the “separation”-process is responsible for the non-monotonic convergence observed in the simulations. As mentioned previously, it suggests the existence of a saddle point at M^* . This could be verified in the 2/2-case where the map $[d_{11}, d_{21}]^+ = f_{cyc}([d_{11}, d_{21}])$ has been determined, see equations (7.39) and (7.40).

Evaluating the Jacobian at $d_{M^*} = 2/3 L$ we obtain

$$\left. \frac{\partial f_{cyc}}{\partial d_{i1}} \right|_{i=1,2} = \begin{bmatrix} 5/4 & -1 \\ -1 & 5/4 \end{bmatrix} \quad (7.50)$$

which indeed has a stable eigenvalue $\lambda_1 = 1/4$ with eigenvector $[1 \ 1]$ and an unstable one $\lambda_2 = 9/4$ with eigenvector $[1 \ -1]$. In figure (7.26), the solutions in $[d_{11}, d_{21}]$ -space are displayed. The line $d_{11} = d_{21}$ separates two invariant regions. Υ corresponds to the situation where $\hat{\theta}_1$ is initialized closer to θ_1 than $\hat{\theta}_2$, in $\bar{\Upsilon}$ the situation is reversed. Invariance of the regions follows from lemma 7.5.




Figure 7.26: Existence of a saddle point

7.3 Dynamic Systems

In static systems the distances in parameter space were directly used to update the location of the models. In real-world problems, we have to use the output of a dynamic (i.e. containing at least one integrator) system in order to adjust the parameters. In general, a sequence of output measurements is necessary to gain information about the parameter error and generate an appropriate update vector. The length of this vector is determined by the performance-dependent adaptation gain defined in (7.9). The question addressed in this chapter is whether a self-organization process –similar to the one observed in static systems– takes place, such that each model converges to a *different* element of the set

$S = \{\theta_1, \theta_2, \dots, \theta_N\}$. In particular, conditions have to be determined under which the dynamic system has the same convergence properties as its static counterpart.

Depending upon the choice of the performance criterion in (7.9), different static analogies are obtained. If J_i is an integral-type criterion and the plant is constant over a sufficiently long interval T , the update direction of the models is concentrical with respect to the current plant parameter vector $\theta \in S$. If, in turn, an instantaneous criterion is used, the models evolve along a single direction and the static analogy would be that all plants and models are confined to a line (\rightarrow section 7.2.5). It is seen that the choice of the performance criterion varies significantly according to the nature of the time-variation and the size of the set S . In principle, it is possible to obtain convergence of all the models, even as the plant switches randomly within the elements of the set S .

The ultimate objective is to determine a control law that stabilizes the system and matches the output of the unknown time-varying system to the output of a known time-invariant reference system. In contrast to the identification problem, it is impossible that the tracking error tends to zero. As pointed out previously, the error will be nonzero at the instants following a switching (as a consequence of the relative degree of the system). The principal question is whether a certainty equivalence controller based on the parameter estimates guarantees bounded solutions even if the models are far from the plants.

7.3.1 Identification

In the following, we examine the equivalence between the static and the dynamic identification problem. Let us start with an example.

Example 7.4 Consider the first-order, time-varying system

$$y(k+1) = \theta(k)y(k) + u(k) \quad (7.51)$$

where $\theta(k)$ assumes a finite number of constant unknown parameters $\theta_j \in S \subset \mathbb{R}$, $j \in \Omega$. Using estimation models of the form,

$$\hat{y}_i(k+1) = \hat{\theta}_i(k)y(k) + u(k), \quad i \in \Omega \quad (7.52)$$

we obtain the identification error:

$$\epsilon_i(k+1) = y(k) [\hat{\theta}_i(k) - \theta(k)] \quad (7.53)$$

If $y(k+1)$ is the measured output, and $y(k) \neq 0$, equation (7.53) can be solved for the parameter error $\tilde{\theta}_i(k) = [\hat{\theta}_i(k) - \theta(k)]$ and has a unique solution. We obtain the update rule

for the parameter estimates

$$\tilde{\theta}_i(k+1) = \tilde{\theta}_i(k) - \eta_i(k) \tilde{\theta}_i(k) \quad (7.54)$$

which is the same as equation (7.10). Hence, in this one-dimensional case, the static and dynamic cases are equivalent. \square

In general, when $S \subset \mathbb{R}^p$, $p > 1$, it is impossible to uniquely determine $\tilde{\theta}_i(k)$ from the error equation

$$\epsilon_i(k) = -\phi(k-d)^T \tilde{\theta}_i(k-1) \quad (7.55)$$

However, under the additional assumption that the regression vector $\phi(\cdot)$ is persistently exciting, the error signal $\epsilon_i(\cdot)$ contains sufficient information about $\tilde{\theta}_i(\cdot)$.

Integral performance criterion

For the moment, let us assume that ϕ is persistently exciting. Notice, that while the excitation is usually provided by the (reference) input to the system, an additional source of excitation is opened up by the switching of the plant parameters. Hence the assumption is not a strong one. It implies that over a finite interval of length T_0 , the regression vector ϕ has a nonzero projection along *any* direction in parameter space. Hence, after T_0 steps we have enough equations to determine the unique solution of equation (7.55). Persistence of excitation corresponds to the fact that the matrix $G_O(k_0, k_0 + T)$ defined in (2.25) satisfies:

$$G_O(0, T_0) \geq cI \quad c > 0 \quad \text{for some fixed } T_0 > 0 \quad (7.56)$$

where $k_0 = 0$. The condition is equivalent to the fact that the system given by (7.6) and (7.7) is *observable on the interval* $[0, T_0]$. This means that it is possible to uniquely determine $\tilde{\theta}_i(0)$ from the knowledge of the output $\epsilon_i(k)$ over the interval $k \in [0, T_0]$. In other words, if ϵ_1 and ϵ_2 are the error signals of two models then

$$\epsilon_1[k, 0, \tilde{\theta}_1] = \epsilon_2[k, 0, \tilde{\theta}_2] \quad \forall k \in [0, T_0] \quad \Rightarrow \quad \tilde{\theta}_1 = \tilde{\theta}_2$$

Hence, in order to determine the parameter error of a model at instant k it is necessary to observe its output error over an interval of length T_0 . If (7.56) holds for some fixed $T_0 > 0$ then the parameter projection algorithm is globally exponentially convergent. In such a case, $\|\tilde{\theta}_i\|$ forms a strictly monotonically decreasing sequence and $\|\tilde{\theta}_i(k)\| \rightarrow 0$. Equation (7.56) suggests the definition of a performance index, defined in terms of a finite sequence of output errors:

$$J_i(k) = \sum_{\nu=k-T_0}^k \epsilon_i^2(\nu) \quad (7.57)$$




Figure 7.27: Dynamic parameter adjustment

If the interval over which the plant is constant is sufficiently long, $T \gg T_0$, the criterion provides a well-defined measure of the parametric distances. These are the quantities needed in the definition of the adaptive gain (7.9). However, no attention has yet been paid to the fact that the system is time-varying. The fact that signals are summed over a window of length T_0 implies that the result contains previous errors that do not always correspond to the plant that is currently in existence. In a time-varying environment it is impossible to obtain $\tilde{\theta}_i(k)$ unless the instants at which the plant switches from one element in S to another are known a priori which, of course, corresponds to the trivial situation. Similarly, with the above performance criterion, it is not possible to correctly determine the closest model in parameter space at every instant of time. To some extent, the problem can be taken care of by using a performance index with exponential forgetting factor:

$$J_i(k) = \alpha \epsilon_i^2(k) + \beta \sum_{\nu=k-T_0}^k \rho^{k-\nu} \epsilon_i^2(\nu) \quad \rho \in [0, 1] \quad (7.58)$$

While (7.58) enables us to select any blend between an instantaneous and the integral performance criterion, prior information about the nature of the time-variation is needed to select an appropriate one.

Figure (7.27) illustrates the case where $\hat{\theta}_3$ turns out to be the best model when comparing its performance J_3 (defined in 7.57) to the performance of the other models. The plant

satisfies $\theta(k) \equiv \theta_3$ over an interval of length T . The gain η_3 assigned to $\hat{\theta}_3$ is large at every update step during which θ_3 is in existence. Once the plant switches, the largest gain will not immediately pass to another model since, again, T_0 observations are necessary to assess the performance of all the models when a new plant is in existence. As a consequence, the vector obtained by summing up T_0 steps does not exactly coincide with $\tilde{\theta}_{33} = [\hat{\theta}_3 - \theta_3]$, since it will inherently contain a number of erroneous steps. If $\theta(k) \equiv \theta_3$ over a long interval $T \gg T_0$, the share of these steps will be small and the sum vector be a good approximation of the static case, where $\tilde{\theta}_{33}$ is used to update the location of the models.

Instantaneous performance criterion

The introduction of a forgetting factor in equation (7.58) is motivated by the fact that the resulting performance criterion reflects a time-variation more promptly than a criterion where the errors of past and present instants are weighted equally. As ρ becomes small, the impact of past errors decreases. In this section, we investigate the properties of the algorithm when $\rho = 0$, i.e. only instantaneous errors are used to assess the performance of the models. According to equation (7.5), the regression vector ϕ is the same for all models. This is a consequence of the series-parallel architecture used in the definition of multiple models. The measured output is used instead of the estimated one in the definition of the regression vector. The models are merely predictors of the plant output and do not possess eigen-dynamics. The parameters of the models are updated according to the equation:

$$\hat{\theta}_i(k) = \hat{\theta}_i(k-1) + \eta_i(k) \frac{\phi(k-d)\epsilon_i(k)}{\phi(k-d)^T \phi(k-d)} \quad (7.59)$$

where

$$\eta_i(k) = \eta_0 \frac{\frac{1}{\epsilon_i^q(k)}}{\frac{1}{\epsilon_1^q(k)} + \frac{1}{\epsilon_2^q(k)} + \dots + \frac{1}{\epsilon_N^q(k)}} \quad q > 1 \quad (7.60)$$

using

$$J_i(k) = \epsilon_i^q(k) \quad (7.61)$$

It follows that the parameters of all the models are adjusted using the same projection vector. In other words, at instant of time k , all models move in same direction $\phi(k-d)$ — or opposite to $\phi(k-d)$ depending upon the sign of $\epsilon_i(k) = -\phi(k-d)^T \tilde{\theta}_i(k-1)$. Let us first assume that the regression vector is constant.

Example 7.5 The time-varying moving-average system

$$y(k+1) = \theta(k)^T \phi(k) \quad (7.62)$$

where $\theta(k) \in S = \{\theta_1, \theta_2, \theta_3\}$, $S \subset \mathbb{R}^2$ and $\phi(k) = [u(k) \ u(k-1)]^T$ is to be identified using $N = 3$ models. We set $u(k) \equiv 1 \ \forall k > 0$. Since each element of S consists of two unknown parameters, $\phi(k)$ is *not* sufficiently exciting. The models converge to a point on the hyperplane

$$H_j = \{\hat{\theta} \mid \phi(k-d)^T[\hat{\theta} - \theta_j] = 0\} \quad \forall k \in [k, k+T] \quad (7.63)$$

H_j represents the set of solution points associated to an element $\theta_j \in S$, $j = 1 \dots 3$, since the identification error vanishes at any point on H_j , for all $k \in [k, k+T]$ where θ_j is in existence. In the present case, H_j corresponds to the line through θ_j and orthogonal to $\phi(k) \equiv \text{const.}$ (see figure 7.28). But this means that the convergence problem can be reduced to a one-dimensional subspace corresponding to the line parallel to ϕ . \square




Figure 7.28: Line analogy

The parameters take a single step to converge to a point on H_j if they are updated according to (7.59) and $\eta_i \equiv 1$. All models use the same direction $\phi(k-d)$ to update their parameters. For the analysis of the case where $\eta_i \neq 1$ is defined as (7.60), it is possible to focus on the component of the parameter error $\tilde{\theta}_i(\cdot)$ which is orthogonal to H_j . This component is given by

$$\tilde{\vartheta}_i(k-1) = -\frac{\phi(k-d) \epsilon_i(k)}{\|\phi(k-d)\|^2}$$

Using (7.55), we may rewrite $\tilde{\vartheta}_i(k-1)$ as

$$\tilde{\vartheta}_i(k-1) = \frac{\phi(k-d)}{\|\phi(k-d)\|} \|\tilde{\theta}_i(k-1)\| \cos[\alpha(k-1)] \quad (7.64)$$

where $\alpha(k-1)$ is the angle between $\phi(k-d)$ and $\tilde{\theta}_i(k-1)$. The models differ only by their distance $\delta_i(k) = \|\tilde{\theta}_i(k)\| \cos \alpha(k)$ from the hyperplane H_j where $\delta_i \in \mathbb{R}$. The evolution of $\tilde{\vartheta}_i(k)$ may be described as (\rightarrow see equation 7.59):

$$\tilde{\vartheta}_i(k) = \tilde{\vartheta}_i(k-1) - \eta_i(k) \tilde{\vartheta}_i(k-1) \quad (7.65)$$

Since $\tilde{\vartheta}_i(k)$ is known, the equation is equivalent to the static update equation (7.10), where all models and plants evolve on a straight line.

Now, suppose that $\phi(k)$ is *not* constant. Since all models use the *same* regression vector, a similar argumentation as above could be used to reduce the problem to the 1-dimensional subspace defined by the line parallel to $\phi(k)$. However the situation is more difficult. This is because, the set of solution points defined in (7.63) is smaller if $\phi(k)$ is time-varying. It collapses to a single point when $\phi(k)$ is persistently exciting. If the instantaneous error is used in the definition of the step-size $\eta_i(k)$ in (7.60), a model which is far from the solution point may misleadingly be assigned a large step-size. It is hard to determine how this affects the overall convergence process. In the special case where $\phi(k)$ forms a sequence of orthogonal vectors (as in the case of the orthogonal projection algorithm, chapter 2.2), we obtain a set of independent convergence processes associated with each direction of the regression vector $\phi(k)$, $k > 0$. In figure (7.28) we consider two directions, marked as ϕ and $P\phi$, where $\phi \perp P\phi$. Any adjustment along $P\phi$ does not alter the location of the models on ϕ and vice versa. In other words, the adjustments are completely decoupled and can be studied separately, as if the associated regression vectors were constant. This has two implications: First, the analysis of the overall convergence is simplified, since it can be carried out using the “on-the-line” static analogy for each direction separately. Secondly, new equilibrium points appear at the intersections of the hyperplanes corresponding to each θ_j , which are indicated by a “+” in figure (7.28). They result from the fact that for a given model $\hat{\theta}_i$ to converge it suffices that its identification error ϵ_i is zero for *one* of the regression vectors, i.e. either $\phi^T[\hat{\theta}_i - \theta_j] = 0$ or $P\phi^T[\hat{\theta}_i - \theta_j] = 0$, while any of the other models can be used to satisfy the other error equation. In other words, every model converges to a vector $\bar{\theta}_i$, $i \in \{1, 2, 3\}$, whose elements coincide with the elements of the *collection* $[\theta_1, \theta_2, \theta_3]$ of the true plants but not necessarily to the elements of a *single* vector θ_j , $j \in \{1, 2, 3\}$. Additional processing may be required (after convergence) before the identification result can be applied for control.

7.3.2 Certainty Equivalence Control

As stated in the introduction, the objective is to control a linear time-varying system where $\theta(k)$ assumes a finite number N of constant, unknown values, i.e. $\theta \in S = \{\theta_1, \theta_2, \dots, \theta_N\} \subset \mathbb{R}^p$. From equation (7.4) we have,

$$y_j(k+d) = \phi(k)^T \theta_j \quad (7.66)$$

where $\theta_j \in S$ is the plant parameter that is in existence over an interval $T = [k_t, k_{t+1}]$, $t > 0$. As in the static case, multiple adaptive models generate the estimates of the plant parameters. The models have the following form,

$$\hat{y}_i(k+d) = \phi(k)^T \hat{\theta}_i(k) \quad (7.67)$$

where $\hat{\theta}_i \in \mathbb{R}^p$, $i \in \Omega$ which are updated simultaneously using the performance-dependent adaptive gain $\eta_i(k)$ defined above. The control law is determined at every instant based on the model that performs best according to the switching criterion (6.1). Note that the criterion for selecting the controller may be different from the criterion used in the definition of $\eta_i(k)$. The objective is to cause $y(k)$ to track an arbitrary bounded signal $y^*(k)$ that is the output of a time-invariant reference system. If, at a given instant, $\hat{\theta}_i^*$ are the parameters corresponding to the best model, the control input is determined assuming that $\hat{\theta}_i^*$ are the true values of the plant parameters. In other words, $u(k)$ is computed from the equation

$$y^*(k+d) = \phi(k)^T \hat{\theta}_i^*(k) \quad (7.68)$$

The design procedure is seen to be identical to the one used in conventional adaptive control. By virtue of the certainty equivalence principle, the model parameters are regarded as the true parameters of the system. As seen in chapter 3.1.4, the proof of stability relies upon the fact that the estimation error cannot grow faster than the signals of the system. This, in turn, was seen to depend upon the property of the estimation algorithm, that the parameter error $\tilde{\theta}(k)^T \tilde{\theta}(k)$ is a non-increasing function of time (\rightarrow equation 2.16). If the plant is time-varying it is generally impossible to define a non-increasing sequence of parameter errors $\tilde{\theta}$, unless there is a *set of models* that converges to the N distinct parameters that the plant assumes. If the parameter error is then defined to be the error of the closest model, there must be a stage after which the error does not increase. But this is precisely, what is accomplished by the above control law defined in terms the best estimate $\hat{\theta}_i^*$. By virtue of the simultaneous identification scheme introduced earlier, all models converge simultaneously to the set S . Unlike regular MMST (chapter 6), the models converge even if there are not initially distinguished with respect to any plant parameter. This enables us to define a

parameter error $\tilde{\theta}^*(k) = [\hat{\theta}_i^*(k) - \theta_j]$, $i, j \in \Omega$, the norm of which is non-increasing after some stage $k > \bar{k}$, independent of which plant is in existence at that instant. The only exception are the instants following a switching of the plant. Since this error is only due to the relative degree of the system (not the fact that $\|\tilde{\theta}^*\|$ increases), similar arguments as in the time-invariant case can be used to proof stability of the closed-loop system. The assumptions made at the beginning of this chapter are found to be relevant. In particular, it must be assumed that the system is minimum-phase for every $\theta_j \in S$ and that the plant remains constant over at least $T = \max\{T_0, d + 1\}$ instants of time. The last condition reduces to $T > d$ if an instantaneous switching criterion is used for the definition of the step-size and the selection of the controller.

Example 7.6 Given the second-order system

$$y(k+1) = a_0(k)y(k) + a_1(k)y(k-1) + u(k-1) \quad (7.69)$$

where $a(k) \in \mathbb{R}^2$ assumes one out of five unknown values $S = \{\theta_1, \theta_2, \theta_3, \theta_4, \theta_5\}$ randomly as in example 6.1. The parameter region for which the open-loop system is stable is indicated by the triangle in figure (7.29). Only θ_2 and θ_3 correspond to stable plants. We compare the performance of a controller based on the regular multi-model switching and tuning approach (MMST) of chapter 6 to the simultaneous identification approach (SIC) introduced in this chapter. The objective is to track a desired output of the form $y^*(k+1) = \arctan(5 \sin \frac{2\pi k}{150}) + \frac{1}{2} \sin \frac{2\pi k}{60}$. In both cases, a certainty equivalence controller is determined using the parameters of the model that performs best. To avoid frequent switching, the sum over $T_0 = 10$ subsequent errors is used as a switching criterion together with a forgetting factor of $\rho = 0.9$, i.e.

$$J_i(k) = \sum_{\nu=k-T_0}^k (0.9)^{k-\nu} \epsilon_i^2(\nu) \quad (7.70)$$

In MMST, only the best model (according to (7.70)) is updated while the location of the other models is retained. In the SIC-approach all models are updated at every instant using a performance-dependent step-size $\eta_i(k)$ as defined in (7.9). The closest model is determined by means of the criterion (7.70), with $T_0 = 3$.

The top part of figure (7.29) displays the evolution of the models in parameter space. Only two models are found to be adaptive in the MMST case. The approach clearly fails since the plant assumes five distinct values which cannot be represented by two models. In SIC, all models – which are initialized far from the plants – are adapted and gradually approach the five points. Even though the trajectories are found to be highly erratic and convoluted,

there is a stage \bar{k} after which all models have entered a neighborhood of their respective plants. The parameter error associated with each model–plant pair is non–increasing, for all $k > \bar{k}$, and the above argumentation can be used to obtain a stabilizing control law. As a consequence, all signals are bounded in the SIC–case and the tracking error tends to zero except for the instants following a switching. In the lower part of figure (7.29) the outputs y^* and y of the reference and the actual system are displayed as well as the input u generated by the adaptive controller. We focus on a rather advanced time–interval $k \in [7000, 8000]$ (models have reached the above neighborhoods). Observe that on the right part of the figure (SIC) the sequence of model indices clearly reflects the time–variation of the plant. \square .

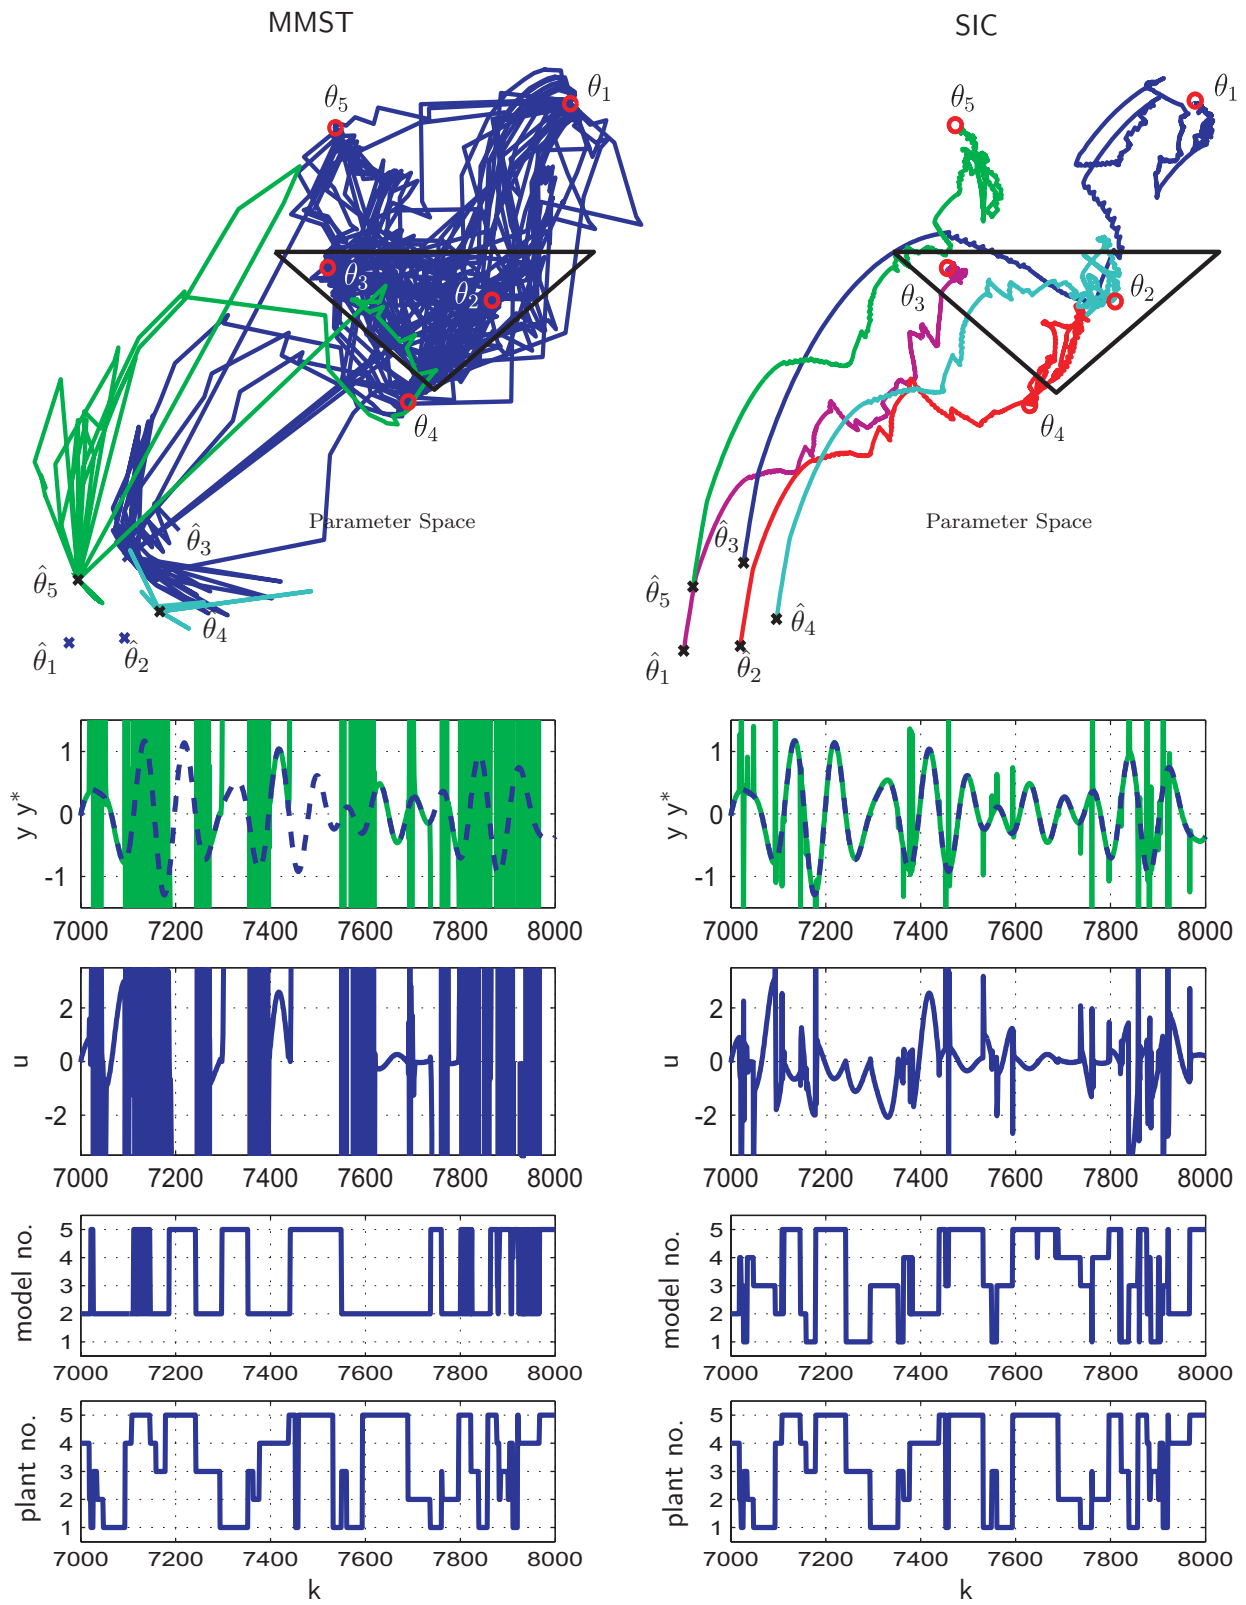


Figure 7.29: Tracking in a time-varying environment

7.4 Future Research

In this chapter, a new algorithm for locating multiple adaptive models in a time-varying environment has been introduced. The emphasis of our discussion has been on developing a mathematical framework within the subsequent problems of **simultaneous identification** and **control** can be addressed. Furthermore, by providing numerous qualitative insights into the behavior of the algorithm we have prepared the ground for a quantitative analysis. The methodology opens up a large number of avenues for future research. Some of these are stated below.

Static Systems

- Convergence was seen to take place in two stages: contraction and separation (of N adaptive models). However, the equations describing the evolution of the models are too complex as to allow us to exploit this observation analytically. Can the equations be made simpler (e.g. in a transformed coordinate system)?
- The simulations suggest the existence of a saddle point M^* . Can this be verified analytically? Do local properties (in the neighborhood of M^*) provide any insight?
- Convergence has been investigated assuming that the time-variation is periodic. What arguments are needed to show that N models converge if the plant switches randomly between the elements of the set S and the switching is governed by an ergodic Markov




Figure 7.30: Convergence of 3 estimates to 3 fixed random points in \mathbb{R}^2 with the following property: A point stays where it is with probability 0.8 and switches to one of the two other points with probability 0.1. In all three simulations the same initial conditions for the estimates were used.

chain (\rightarrow figure 7.30)?

- How must the algorithm be modified if the points $\theta_j \in S$, $j = 1 \dots N$ are replaced by compact sets?

Dynamic Systems

- The assumption that the system assumes each value contained in the set S at least once within a finite interval of length T was critical. An environment which satisfies this condition may be referred to as “persistently exciting”. Is it possible to define persistence of excitation for time-varying systems the way it is defined in classical adaptive control (for time-invariant systems)? The fact that the system is switching does not necessarily imply that the regression vector is persistently exciting as the following example indicates.

Example 7.7 Consider the system

$$y(k+1) = [y(k) \quad u(k) \quad u(k-1)] \theta(k)$$

where $\theta(k) \in S = \{\theta_1, \theta_2, \theta_3\}$, $S \subset \mathbb{R}^3$. Suppose $u(k) \equiv 1$, $\forall k > 0$. Then ϕ is time-varying because of $y(k)$ –but not persistently exciting. \square .

- To what extent can the results obtained for static systems be applied to the dynamic case? Why does the proposed scheme converge at all, if the performance is defined in terms of the instantaneous error?
- Supposing that the models incidentally lie close to the plant how does the performance of the proposed scheme compare to adaptive control using multiple models, switching and tuning (MMST)?
- How must the algorithm be modified if the set S is infinite and the time-variation of the plant is continuous, i.e. $\theta(k)$ is an (arbitrary) continuous function of time? Is it nevertheless possible to control the system in a stable fashion if only a finite number of models is used?

Nonlinear Systems

The situation where the time-varying system is nonlinear is discussed in some detail. Suppose that the system is given by

$$y(k+1) = f_i[Y(k)U(k)] \tag{7.71}$$

where $U(k) = [u(k), \dots, u(k-n+1)]$ and $Y(k) = [y(k), \dots, y(k-n+1)]$ is a sequence of inputs and outputs respectively. Equation (7.73) is known as the nonlinear autoregressive moving-average (NARMA) representation of a nonlinear system having delay $d = 1$. The mapping $f : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ represents a one-step-ahead predictor for the output of the system. Let us assume that the system assumes one of a finite set of right hand sides $S = \{f_1, f_2, \dots, f_N\}$ at least once over an interval of finite length T . As in the previous section, the objective is to estimate the elements of the set S and to build a controller based on the best estimate. For simplicity, we assume that the NARMA model may be approximated using a simplified model in which the control input occurs linearly [61]. This permits an easy algebraic computation of the control input.

$$y_j(k+1) = f_j[Y(k), U(k)] + g_j[Y(k), \bar{U}(k)] u(k) \quad (7.72)$$

where $\bar{U}(k) = [u(k-1), \dots, u(k-n+1)]$. As in the previous sections, N adaptive models are used to generate estimates of the maps f_j, g_j respectively (where $j \in \Omega$). Since the system is nonlinear a radial basis function network (RBFN) is used to approximate the mappings using input-output data. Each model consists of 2 RBFN's:

$$\hat{y}_i(k+1) = \mathcal{F}_i(x)^T \hat{\theta}_i(k) + \mathcal{G}_i(x)^T \hat{\psi}_i(k) u(k) \quad (7.73)$$

where $\mathcal{F}(\cdot), \mathcal{G}(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$ is a vector of radial basis functions defined on a n -dimensional input space. The network parameters $\hat{\theta}_i$ and $\hat{\psi}_i$ are updated using the algorithm (7.59), (7.60) with ϕ replaced by \mathcal{F} and $\mathcal{G}u(k)$ respectively. A certainty equivalence controller is then defined by

$$u(k) = \frac{y^*(k+1) - \mathcal{F}_{i^*}(x)^T \hat{\theta}_{i^*}(k)}{\mathcal{G}_{i^*}(x)^T \hat{\psi}_{i^*}(k)} \quad (7.74)$$

where the index i^* refers to the model that best fits the plant according to one of the performance criteria defined above.

Example 7.8 Consider an unknown nonlinear, time-varying plant which assumes one of the following right hand sides periodically with period $T = 1500$ s

$$y(k+1) = \begin{cases} \frac{y(k) u(k-1)}{1 + y(k)^2} + u(k) \\ \frac{\cos y(k)}{1 + u(k-1)^2} + u(k) \\ \frac{\sin^3 y(k)}{1 + \frac{1}{2} \cos^2 y(k-1)} + u(k) \end{cases} \quad (7.75)$$

Since the system is linear in $u(k)$, the control input may be directly computed according to equation (7.74) where $\mathcal{G}_{i^*}(x)^T \hat{\psi}_{i^*}(k) \equiv 1 \forall k > 0$. The RBFN used to approximate the first

term in (7.75) consists of a set of radial basis functions of the form

$$\mathcal{F}_l = \exp[-0.5(x - c_l)^T \text{diag}\{25\}(x - c_l)], \quad l = 1 \dots 16$$

where $x \in \mathbb{R}^2$ and c_l are equally spaced centers over a square $[-1 \ 1] \times [-1 \ 1]$. Figure

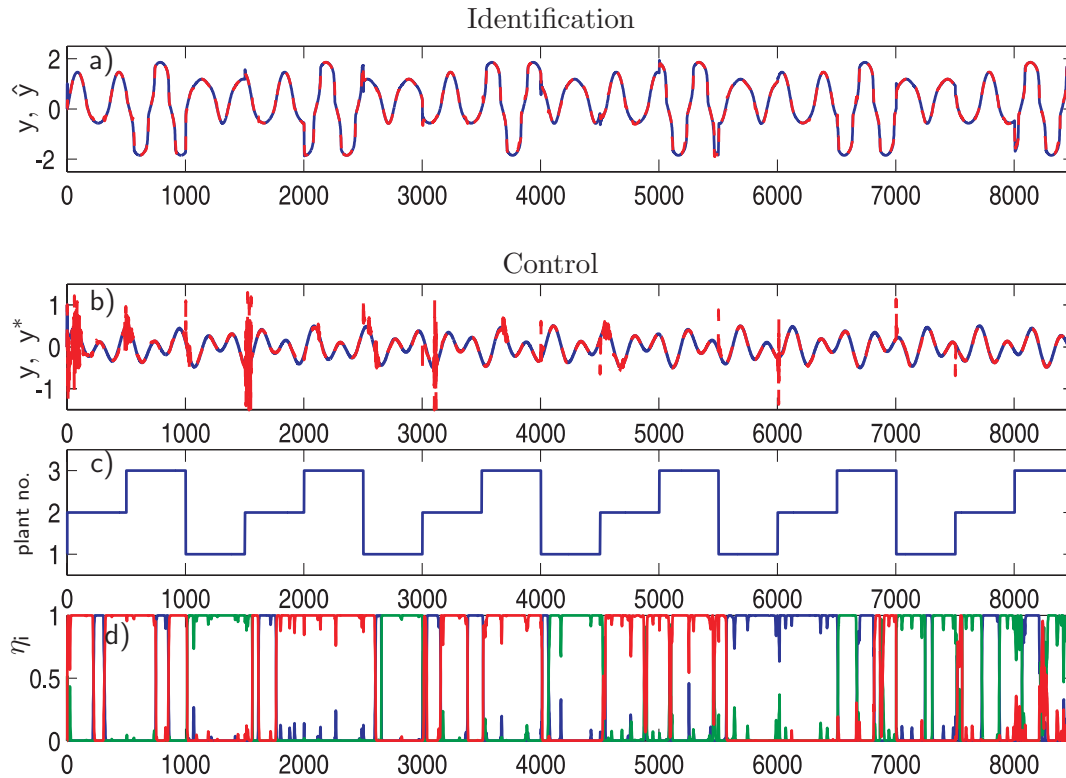


Figure 7.31: Identification and Control of a nonlinear time-varying system

(7.31) displays the results of an identification procedure (a) with $u(k) = \sin \frac{2\pi k}{350}$ as well as the performance of a tracking controller (b) for the desired output $y^*(k+1) = 0.2 \sin \frac{2\pi k}{400} + 0.3 \sin \frac{2\pi k}{225}$. In (c) and (d) it is seen that different RBFNs specialize in the different right-hand sides that the plant assumes periodically. \square

Chapter 8

Application: Two–Mass System

In order to demonstrate that the theory contained in this thesis can be applied successfully to an industrial control problem we present an experimental study performed on a servo drive system. Industrial processes often require electrical drives with high dynamic performance. Servo drive applications unavoidably include mechanical transfer elements such as shafts, gears or belts. A perfect coupling would consist of an absolutely rigid shaft, no friction in the bearings and gears with no backlash. However, such elements do not exist in practice. The imperfections introduced by the transfer elements have to be considered in high dynamic speed drives.




Figure 8.1: A Two-Mass System

In particular, an elastic coupling may cause resonant frequencies to lie within the bandwidth of the speed controller resulting in degraded response and the onset of weakly-damped oscillations of the mechanical load. Drives with considerable transmission elasticity and small friction and backlash effects may be efficiently modelled by a linear two-mass system. This is the kind of system considered in this chapter. It consists of two rotational masses of inertias J_1 and J_2 , the first one being the mass of a motor and the second one corresponding to the mass of the load. Typically, $J_1 \ll J_2$. The masses are coupled through a flexible transmission shaft which is modelled as a spring-damper element with stiffness c_X and viscous damping coefficient d_X . A schematic representation of the system is provided in

figure (8.1). A torque v may be applied on the load side acting as an (unknown) disturbance signal. The objective is to control the speed of the load n_2 through the motor torque u such that n_2 tracks an arbitrary bounded reference speed n_2^* while the parameters of the system are assumed to be unknown.

8.1 Model of the Plant

Before deriving the mathematical model of the plant, let us take a look at the hardware components involved. The experimental plant consists of two synchronous drives with integrated encoders, two voltage source power converters, a brake unit, A/D (analog to digital) and D/A converters as well as a computer on which the adaptive control is implemented. The first motor is used to generate the drive torque while the second motor generates disturbance signals.

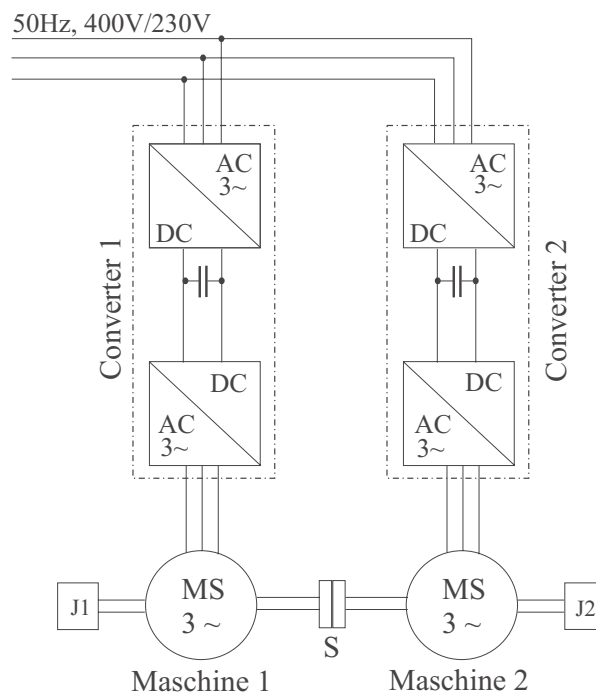


Figure 8.2: Electrical configuration of the plant

Synchronous Drives: The synchronous motors used in this system are excited by “rare earth” permanent magnets which result in a power density superior to DC motors because the limiting effects of the mechanical commutator are absent. The nominal power of each motor is equal to 4.8 kW, the nominal torque generated 23 Nm and the nominal speed 2000

r.p.m. An integrated encoder measures the position of the shaft with a resolution of 2048 bars/rev. The speed signal is obtained by differentiation.

Power Converters: The actuator is given by a voltage source PWM thyristor converter with DC-link voltage of 510 . . . 620 V. If the link voltage exceeds 620 V it is connected to the brake resistor contained in the brake unit. On the machine-side, the converter outputs pulse-width modulated rectangular voltages the fundamental components of which are prescribed by the voltage command signal. Despite the superimposed harmonics (caused by the modulator) the resulting currents in the electrical drive can be considered close to sinusoidal since the switching frequency of 5 . . . 7.5 kHz is relatively high.

The principle of controlling the currents in the synchronous motors is similar to the field orientated control of an induction motor except that no flux model is needed to estimate the rotor flux and slip frequency. Since the latter is obviously zero the rotor position is directly used as the angle of reference. Field orientation has the effect of rendering the dynamic behavior of the synchronous machine similar to that of the DC machine which is substantially simpler since the interaction of flux and armature current is absent. Hence, these quantities can be controlled directly and independently. In a typical configuration, the torque generating current i_q is controlled by an underlying control loop while the direct component i_d is regulated to zero. The response of the current control loops may be approximated by lag elements with a time constant of 2 ms or less. Since the minimal sampling interval used in this setup is 10 ms the time constant is negligible. Hence, the power converter together with the synchronous drive may be viewed as a torque-generating unit responding to command signal with practically no delay. In other words, the electrical components of the plant are part of the fast, parasitic dynamics of the system.

The torque is thus assumed to be directly available at the input u of the two-mass system which, in turn, corresponds to a (damped) linear oscillator with eigenfrequency around 10 Hz. Hence, the mechanical components represent the dominant dynamics of the plant. These are the ones for which a mathematical model is built. In (Schröder 2001, [71]) we find a continuous-time state space model for the elastic coupling of two rotational masses:

$$\begin{bmatrix} \dot{n}_1 \\ \Delta\dot{\varphi} \\ \dot{n}_2 \end{bmatrix} = \begin{bmatrix} -\frac{d_X}{J_1} & -\frac{c_X}{J_1} & \frac{d_X}{J_1} \\ 1 & 0 & -1 \\ \frac{d_X}{J_2} & \frac{c_X}{J_2} & -\frac{d_X}{J_2} \end{bmatrix} \begin{bmatrix} n_1(t) \\ \Delta\varphi(t) \\ n_2(t) \end{bmatrix} + \begin{bmatrix} \frac{u(t)}{J_1} \\ 0 \\ -\frac{v(t)}{J_2} \end{bmatrix} \quad (8.1)$$

With reference to figure (8), the state is given by the angular velocity n_1 of the motor and the velocity n_2 of the load. $\Delta\varphi$ denotes the angle between the masses. The parameters c_X , d_X , J_1 and J_2 are assumed to be unknown. However, it is clear from the laws of physics that all parameters must be positive. The transfer-function from the input u to the output

n_2 of the system has a zero at $s = -c_X/d_X$ which is negative real since $c_X > 0$ and $d_X > 0$, i.e. the system is minimum-phase. Since the adaptive controller will be implemented on a computer, it is useful to work with a discrete-time model. The state and input variables are sampled using a zero-order hold method, i.e. the variables are assumed to be constant over the sampling interval $T_S = 10 \text{ ms}$. For shorter intervals, $T_S < 10 \text{ ms}$ we may get unstable zeros even though the continuous-time plant is minimum-phase. This corresponds to the well-known fact that rapid sampling may give rise to non-stably invertible systems (\rightarrow see Åström et al. 1984 [3] and Goodwin, et al. 1986 [25]). If $x(t) = [n_1(t) \Delta\phi(t) n_2(t)]^T$ and \check{A} , \check{b} denote the system matrix and input vector in (8.1) respectively we obtain

$$x(t_{k+1}) = e^{\check{A}(t_{k+1}-t_k)} x(t_k) + \int_{t_k}^{t_{k+1}} e^{\check{A}(t_{k+1}-\tau)} d\tau \check{b} u(t_k) \quad (8.2)$$

The discrete-time state space model of the plant is therefore given by

$$\begin{aligned} x(k+1) &= \bar{A} x(k) + \bar{b} u(k) \\ n_2(k) &= [0 \ 0 \ 1] x(k) \end{aligned} \quad (8.3)$$

which is of the form (3.1) where $\bar{A} = e^{\check{A}T_S}$ and $\bar{b} = \int_0^{T_S} e^{\check{A}\bar{\tau}} d\bar{\tau} \check{b}$ using $\bar{\tau} = t_{k+1} - \tau$. The system is completely observable through the output n_2 which is the only measured signal in our plant. We may therefore express the 3-dimensional state vector of the discrete system in terms of 3 subsequent output measurements. After some standard calculation (see chapter 3.1.1) we obtain an auto-regressive moving-average (ARMA) model of the two-mass system:

$$n_2(k+1) = \sum_{i=0}^2 a_i n_2(k-i) + \sum_{i=0}^2 b_i u(k-i) + \sum_{i=0}^2 c_i v(k-i) + \xi(k+1) \quad (8.4)$$

Clearly, a_i, b_i, c_i are nonlinear functions of the physical parameters J_1, J_2, c_X, d_X . An additional disturbance signal $\xi(k+1)$ is present which reflects the nonlinear effects such as dry friction, nonlinear spring characteristics as well as the parasitic electrical part of the plant, in short, all effects that have not been included into the model. We assume that ξ is small compared to the state of the overall system, such that if the latter is bounded, ξ is bounded as well. For notational convenience we define

$$\begin{aligned} \theta_0 &= [a_0 \quad a_1 \quad a_2 \quad b_0 \quad b_1 \quad b_2]^T \\ \phi(k) &= [y(k) \quad y(k-1) \quad y(k-2) \quad u(k) \quad u(k-1) \quad u(k-2)]^T \end{aligned} \quad (8.5)$$

and obtain

$$n_2(k+1) = \phi(k)^T \theta_0 + \sum_{i=0}^2 c_i v(k-i) + \xi(k+1) \quad (8.6)$$

8.2 Design of the Adaptive Controller

Summing up the prior information about the plant, we know that the system is minimum-phase, of third order and delay one. Given a reference model with delay $d^* \geq 1$ and an arbitrary bounded input, e.g. $r(k) = n_{20} + \sin(\omega T_S k)$, the objective is to design a controller which tracks the output of the reference model $n_2^*(k) = q^{-d^*} \frac{D(q^{-1})}{C(q^{-1})} r(k)$ while keeping all signals in the system bounded:

$$\lim_{k \rightarrow \infty} |n_2(k) - n_2^*(k)| \leq \chi \quad (8.7)$$

$$\|\phi(k)\| \leq \bar{\phi} < \infty \quad \text{for all } k > 0 \quad (8.8)$$

where $\chi > 0$ is some small value which is zero if $\xi = 0$. Let the closed-loop transfer function be denoted by $H(q^{-1})$. If $H^*(q^{-1})$ is the transfer function of the reference system note that our objective is to match $H(q^{-1})r(k)$ and $H^*(q^{-1})r(k)$ as closely as possible which is different from matching the transfer functions themselves, i.e. $H(e^{-j\omega T_S}) \equiv H^*(e^{-j\omega T_S})$, for all $\omega \in [0, \pi/T_S]$. The last equality may only be obtained if the regression vector $\phi(k)$ is persistently exciting of sufficient order. If no disturbances are present, the only source of excitation is the reference input $r(k)$. Note, however, that the nature of the reference signal is dictated by the process requirements not the mathematical ones (related to the richness of $\phi(k)$). In some cases, a pseudo-binary noise signal is added to the reference signal to enhance excitation. It should be kept in mind, though, that the presence of unmodelled dynamics generally interdicts such an action since it invariably leads to the problem of matching a low-order model to a high-order plant. The design proceeds in three steps:

Step 1: Controller structure

In chapter 5.1, it was shown that the reference model may be chosen to be an arbitrary stable system with relative degree greater than or equal to d . In the experiment we proceed as in example (5.2) where the denominator of the reference model was defined to be

$$C(q^{-1}) = A(q^{-1}) + \mu [1 - A(q^{-1})] \quad (8.9)$$

and $\mu \in [0 \dots 1]$. $A(q^{-1})$ is chosen to be *some* polynomial having a complex conjugate pair of poles and an additional pole at $z = 1$. Here, we have exploited the prior knowledge that the two-mass system consists of a linear oscillator and a pure integrator. For $\mu = 1$, the reference model corresponds to a deadbeat control system which causes the load speed $n_2(k)$ to be equal to its desired value after only $d = 1$ instant of time. For the moment, let us

assume that $v(k) \equiv 0$, for all $k > 0$. If the parameters were known we would set

$$C(q^{-1})n_2^*(k) = \phi(k-1)^T\theta_0 \quad (8.10)$$

and solve for $u(k)$. Since the parameters are unknown, an identification model has to be set up which generates estimates of θ_0 . In view of the previous equation, an obvious choice for such a model is

$$C(q^{-1})\hat{n}_2(k) = \phi(k-1)^T\hat{\theta}(k-1) \quad (8.11)$$

where

$$\hat{\theta}(k) = [\hat{a}_0(k) \ \hat{a}_1(k) \ \hat{a}_2(k) \ \hat{b}_0(k) \ \hat{b}_1(k) \ \hat{b}_2(k)]^T \quad (8.12)$$

The identification error is equal to

$$\begin{aligned} \epsilon(k) &= C(q^{-1})n_2(k) - \phi(k-1)^T\hat{\theta}(k-1) \\ &= \phi(k-1)^T[\theta_0 - \hat{\theta}(k-1)] + \xi(k) \end{aligned} \quad (8.13)$$

It contains the effects of both the parametric error

$$\tilde{\theta}(k) = \hat{\theta}(k) - \theta_0 \quad (8.14)$$

and the residual error $\xi(k)$ due to unmodeled dynamics. As above, we set

$$C(q^{-1})n_2^*(k) = \phi(k-1)^T\hat{\theta}(k-1) \quad (8.15)$$

and solve for $u(k)$. By virtue of the *certainty equivalence principle* the estimates $\hat{\theta}(k)$ are used as if they were the actual parameters θ_0 of the system. This has the effect of relating the control error to the identification error by an equation of the form (3.29). In this case, the control error $e(k)$ becomes

$$e(k) = C(q^{-1})n_2(k) - C(q^{-1})n_2^*(k) = -\phi(k-1)^T\tilde{\theta}(k-1) + \xi(k) \quad (8.16)$$

which is actually equivalent to the identification error (8.13) because of $d = 1$. Since the system is minimum-phase the adaptive controller defined by equation (8.15) is stable provided that ξ is sufficiently small in the sense of our discussion in section 5.3.

If the torque on the load side is nonzero but constant, i.e. $v(k) = v_0$, it may be regarded as an (unknown) parameter. The system equation reads

$$\begin{aligned} n_2(k+1) &= \phi(k)^T\theta_0 + c_V v_0 + \xi(k+1) \\ &= \phi_1(k)^T\theta_1 + \xi(k+1) \end{aligned} \quad (8.17)$$

where $c_V = c_0 + c_1 + c_2$ from equation (8.4),

$$\begin{aligned} \theta_1 &= [a_0 \quad a_1 \quad a_2 \quad b_0 \quad b_1 \quad b_2 \quad c_V v_0]^T \\ \phi_1(k) &= [y(k) \ y(k-1) \ y(k-2) \ u(k) \ u(k-1) \ u(k-2) \ 1]^T \end{aligned} \quad (8.18)$$

In this case, the adaptive controller is defined by

$$C(q^{-1})n_2^*(k) = \phi_1(k-1)^T \hat{\theta}(k-1) \quad (8.19)$$

where $\hat{\theta}(k) \in \mathbb{R}^7$.

A third case has been studied, in which $v(k) = \sin(\omega hk + \varphi)$ where the frequency ω and phase φ of the disturbance are unknown. As discussed in chapter 4.1.1, complete elimination of the disturbance can be achieved if the order of the controller is augmented. It has been shown that a disturbance-free plant representation exists which is obtained by eliminating $v(k)$ in equation (8.6) using a second-order homogenous system. For the convenience of the reader we restate this fact. Setting $\xi(k) = 0$, a representation of the plant (8.6) is given by

$$A(q^{-1})n_2(k) = q^{-1}B(q^{-1})u(k) + q^{-1}G(q^{-1})v(k) \quad (8.20)$$

where

$$\begin{aligned} A(q^{-1}) &= 1 - a_0q^{-1} - a_1q^{-2} - a_2q^{-3} \\ B(q^{-1}) &= b_0 + b_1q^{-1} + b_2q^{-2} \\ G(q^{-1}) &= c_0 + c_1q^{-1} + c_2q^{-2} \end{aligned}$$

The disturbance signal $v(k)$ is generated by a homogenous system of the form

$$\begin{aligned} v(k) &= [d_1q^{-1} + d_2q^{-2}]v(k) \\ &:= [1 - D(q^{-1})]v(k) \end{aligned} \quad (8.21)$$

If we solve equation (8.20) for $v(k)$ and use the result to substitute $v(k)$ in (8.21) we obtain

$$n_2(k+1) = [1 - D(q^{-1})A(q^{-1})]n_2(k+1) + D(q^{-1})B(q^{-1})u(k) \quad (8.22)$$

which corresponds to a system model the order of which has been augmented in order to account for the presence of an unknown sinusoidal disturbance. The system equation now reads

$$n_2(k+1) = \phi_2(k)^T \theta_2 + \xi(k+1) \quad (8.23)$$

where

$$\theta_2 = \begin{bmatrix} d_1 + a_0 & d_2 - d_1a_0 + a_1 & -d_2a_0 - d_1a_1 + a_2 & -d_2a_1 - d_1a_2 & -d_2a_2 & \dots \\ \dots & b_0 & b_1 - d_1b_0 & b_2 - d_1b_1 - d_2b_0 & -d_1b_2 - d_2b_1 & -d_2b_2 \end{bmatrix}^T$$

$$\phi_2(k) = \begin{bmatrix} y(k) & y(k-1) & y(k-2) & y(k-3) & y(k-4) & \dots \\ \dots & u(k) & u(k-1) & u(k-2) & u(k-3) & u(k-4) \end{bmatrix}^T$$

It is seen from the composition of θ_2 , that even if the parameters of the plant were known, the augmented system is unknown if the coefficients d_1 , d_2 of the disturbance model are not available. In such a case, we design an adaptive controller of the form

$$C(q^{-1})n_2^*(k) = \phi_2(k-1)^T \hat{\theta}(k-1) \quad (8.24)$$

where $\hat{\theta}(k) \in \mathbb{R}^{10}$. It follows that the principal design objective in step 1 is the choice of an appropriate order of the identification model which forms the basis for the adaptive controller. This choice depends strongly on the prior information regarding the order of the undisturbed plant as well as the nature of the disturbance (constant or sinusoidal).

Step 2: Parameter estimation

In order to adjust the parameter vector $\hat{\theta}$ on the basis of the identification error ϵ , any one of the estimation algorithms introduced in chapter 2 can be used. As expected, best results are obtained with the recursive least square algorithm which is both fast and robust to measurement noise. However, the algorithm cannot be used without taking precautions. It is precisely the speed of the algorithm that causes problems. Even if we chose the initial estimate $\hat{\theta}(0)$ close to the actual parameters θ_0 , the algorithm may cause the estimates to “jump away” from θ_0 because of the perturbation term ξ in the identification error. Experimental data indicate that the jump phenomenon is particularly pronounced if the control input generated in the adaptive loop is large and oscillatory. This is in line with our expectations since the perturbation ξ is large whenever unmodelled dynamics are excited by a high-frequency input. The fact that the parameters jump or drift away from the actual parameters may have fatal consequences on the stability of the system which have been observed by many adaptive control practitioners (see e.g. IEE Colloquium 1997 [31]). In order to prevent instability, two alternative methods have been employed in this experiment:

- Select the reference model according to equation (8.9) with $0 < \mu \leq 0.1$.
- Constrain the parameter estimates to lie in a convex region in parameter space.

The rationale behind the first method is that – assuming that the system is known – the control law resulting from (8.9) is given by

$$u(k) = \frac{\mu}{B(q^{-1})} r(k) - \frac{q[1 - A(q^{-1})]}{1/\mu B(q^{-1})} n_2(k) \quad (8.25)$$

At sampling intervals of $T_s = 10 \dots 100 \text{ ms}$ the coefficients of the denominator polynomial $B(q^{-1})$ are close to zero. In the adaptive context this is critical since it causes the corresponding estimates to evolve within a small neighborhood of zero. This inevitably leads

to a (transiently) large control input u . If μ in the above equation is a small number, the effect can be compensated for. Even though no guarantee can be given that the parameter estimates remain away from zero, the performance is improved substantially, since the actual coefficients of the denominator polynomial are far from zero.

An interesting result has been obtained when applying the second method (Feiler et al., 2003 [20]). In this case, the reference model is of deadbeat-type, i.e. $\mu = 1$. Instead of shifting the actual coefficients of the polynomial $B(q^{-1})$, their estimated values are constrained to lie in a convex region of parameter space, such that if the system were frozen at a given instant, $\hat{B}(q^{-1})$ would be stable. Even though the modification cannot be justified theoretically, the nature of the control input was seen to improve drastically if the parameters are constrained. From an intuitive viewpoint, we may argue that, as the matrix adaptation gain P tends to zero, the adaptive controller is almost time-invariant. Suppose then that $\hat{B}(q^{-1})$ in the denominator of the controller is unstable. This is not meaningful, since our system is known to be minimum-phase so all its zeros lie inside the unit circle. Hence it is reasonable to constrain the parameters as described. The closed convex region Υ within which the estimates $\hat{\theta}$ are required to lie is defined as follows:

$$\Upsilon = \{ \hat{\theta} \mid \hat{\theta}_{[4]}z^2 + \hat{\theta}_{[5]}z + \hat{\theta}_{[6]} \neq 0, \forall |z| \geq 1 \} \quad (8.26)$$

where $\hat{\theta}_{[i]}$ denotes the i -th component of $\hat{\theta}$ and $[\hat{\theta}_{[4]} \hat{\theta}_{[5]} \hat{\theta}_{[6]}] = [\hat{b}_0 \hat{b}_1 \hat{b}_2]$, see equation (8.12). In a two-dimensional subspace of \mathbb{R}^7 , Υ has the shape indicated in figure 8.3. It is




Figure 8.3: Convex region Υ

essential to guarantee that the parameter error does not increase even though the parameter estimates evolve in a closed region in \mathbb{R}^7 . The procedure described in (Goodwin 1984, [27]) is adopted here to ensure convergence of the RLS-algorithm. The parameters are updated

as follows:

$$\hat{\theta}(k) = \begin{cases} \hat{\theta}(k-1) + \frac{P(k-2)\phi(k-1)\epsilon(k)}{1 + \phi^T(k-1)P(k-2)\phi(k-1)} & \text{if } \hat{\theta} \in \Upsilon \\ P(k-1)^{1/2}\hat{\rho}(k)' & \text{if } \hat{\theta} \notin \Upsilon \end{cases} \quad (8.27)$$

where $\hat{\rho}(k)'$ is the projection of $\hat{\rho}(k) = P(k-1)^{-1/2}\hat{\theta}(k)$ onto the boundary of $\bar{\Upsilon}$, and $\Upsilon \xrightarrow{P^{-1/2}} \bar{\Upsilon}$. The reason for working in a transformed coordinate basis is that we have to assure that the quantity $V(k) = \tilde{\theta}(k)^T P(k-1)^{-1}\tilde{\theta}(k) = \tilde{\rho}(k)^T \tilde{\rho}(k)$, defined in chapter 2.3, is non-increasing. We transform the coordinate basis of the parameter space by the following linear time-varying transformation

$$\hat{\rho}(k) = P(k-1)^{-1/2}\hat{\theta}(k) \quad (8.28)$$

This involves computing the square root of the inverse matrix $P(k-1)^{-1}$ at every instant of time k . The square root is well-defined since $P(k-1)^{-1}$ is positive definite for all $k > 0$. With the orthogonal matrix U which diagonalizes P^{-1} , i.e. $P^{-1}U = U \text{diag}\{\lambda_1(k-1) \dots \lambda_7(k-1)\}$, we obtain at every instant of time:

$$P^{-1/2} = U \text{diag}\{\sqrt{\lambda_1}, \dots, \sqrt{\lambda_7}\} U^T \quad (8.29)$$

The image $\bar{\Upsilon}(k)$ of the convex region Υ under the linear transformation $P^{-1/2}$ is again convex. We now project $\hat{\rho}(k)$ orthogonally onto the boundary of $\bar{\Upsilon}(k)$ to obtain $\hat{\rho}(k)'$ for which

$$\|\hat{\rho}(k)' - \rho_0\|^2 \leq \|\hat{\rho}(k) - \rho_0\|^2 \quad (8.30)$$

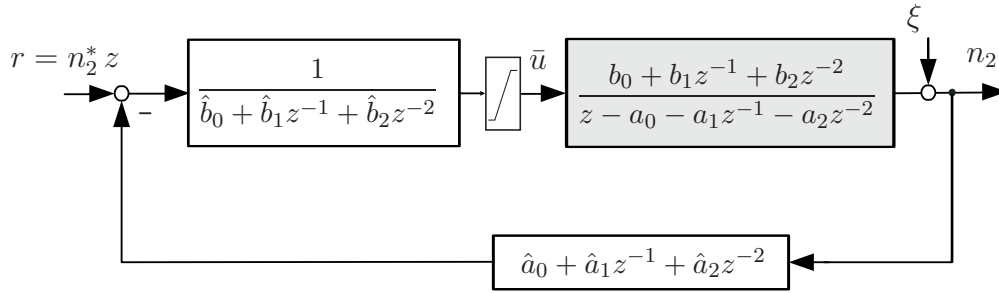
since any point on the boundary of $\bar{\Upsilon}(k)$ is closer to the true parameter $\rho_0 = P(k-1)^{-1/2}\theta_0$ in the transformed space than a point outside $\bar{\Upsilon}(k)$. Moreover, by projecting orthogonally, we pick the $\hat{\rho}(k)'$ on $\bar{\Upsilon}(k)$ which is closest to $\hat{\rho}(k)$. Finally, we put $\hat{\theta}(k)' = P(k-1)^{1/2}\hat{\rho}(k)'$. The use of transformed coordinates has the effect of assuring that

$$\tilde{\theta}(k)'^T P(k-1)\tilde{\theta}(k)' \leq \tilde{\theta}(k)^T P(k-1)\tilde{\theta}(k) \quad (8.31)$$

which is equivalent to the fact that $V(k)$ does not increase as required. Once the estimates have converged to some constant values $\hat{\theta}$, the closed-loop system may be represented using transfer-functions as displayed in figure 8.4.

8.3 Experimental Results

The tests were performed on the system displayed in figure 8.5, the principal technical features of which have been presented in chapter 8.1. The adaptive algorithms have been

Figure 8.4: Closed loop system for constant control parameters $\hat{\theta}$

implemented in a real-time environment using MATLAB[®] and XPCTARGET[®]. The speed of the second mass n_2 is the quantity to be controlled by a torque u generated in the machine connected to the first mass; n_2 is required to follow the speed n_2^* of the reference system as closely as possible.




Figure 8.5: Two-Mass System

8.3.1 General performance

Figure (8.6) displays the response of the adaptively controlled system to a set of reference signals. The results are collected column-wise. The sampling time used in this experiment was $T_S = 0.1$ s. Column one displays the tracking performance using a sinusoidal reference input (with offset): $r(k) = 50 + 5 \sin(0.2\pi kT_S) + 5 \sin(0.4\pi kT_S)$. In column two, the frequency of the sinusoid is increased. In column three, the system is seen to follow a random reference input $r(k) = 100 + w(k)$, w being a random number within the interval $[-15, 15]$. The plant is initially controlled using (arbitrary) fixed parameters. At $t = 10$ s adaptation is turned on, i.e. the parameters are adjusted. It is seen that, in all three cases, the output

n_2 almost perfectly tracks the reference output (see row 4 of the figure) while keeping the input torque u bounded, $|u| \leq 23Nm$ (row 2).

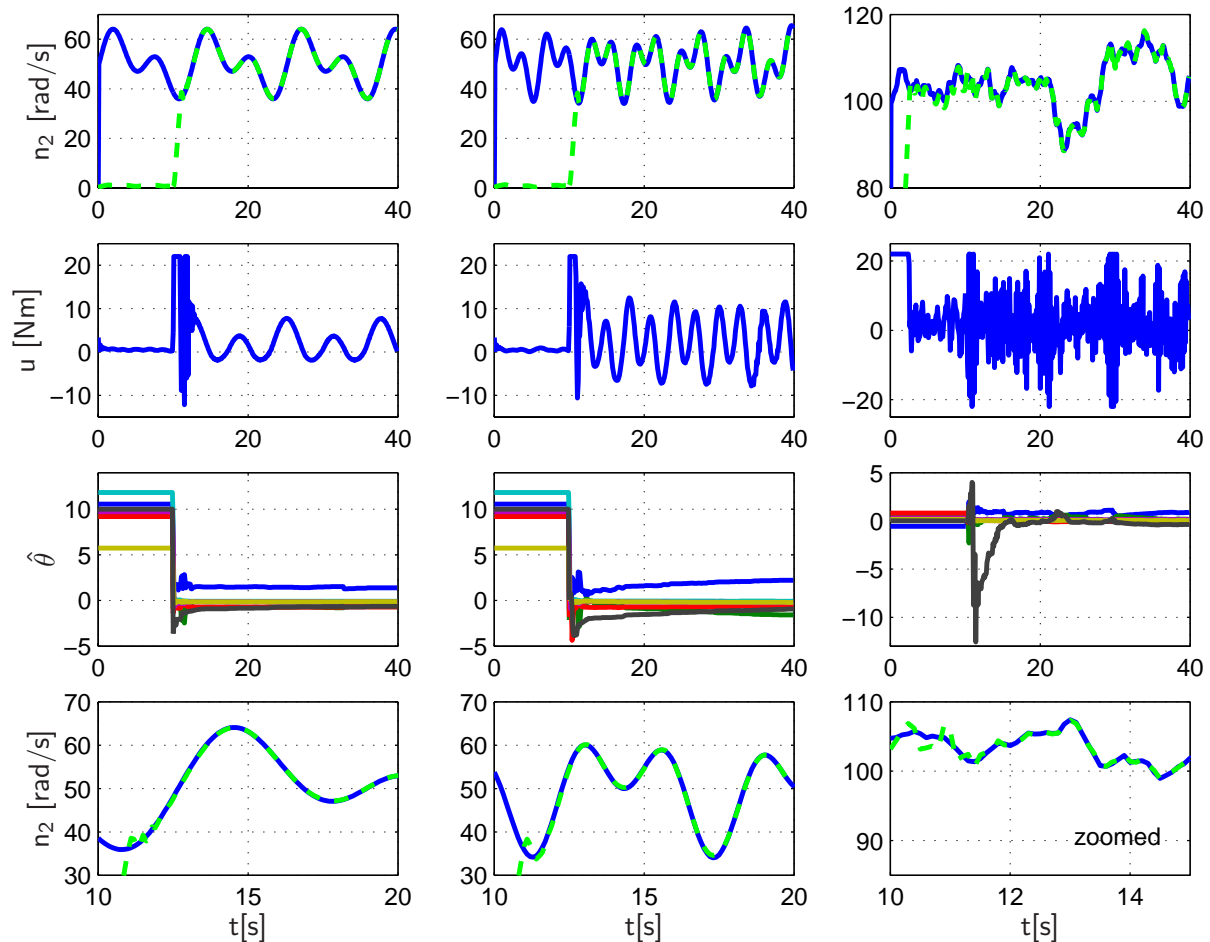


Figure 8.6: Closed loop system for constant control parameters $\hat{\theta}$

It shall be noted that the results (particularly regarding the evolution of the input u) could not have been obtained without constraining the parameters to the convex region defined in (8.26). This is illustrated in the next experiment, figure 8.7. In the left column, the controller works without constraining the parameter estimates to the convex region Υ . Assuming that the parameters evolve on a much slower time-scale than the state of the system, the adaptation is frozen at every instant of time and the roots of the polynomial $\hat{B}(q^{-1})$ in (8.25) – corresponding to the poles z_1 and z_2 of the controller transfer function – displayed in row 3 of the figure. There is one pole which lies (close but) outside the unit circle. The control input u is large and oscillatory thereby exciting unmodelled dynamics. Still, the signals remain bounded due to the presence of the actuator saturation $|u| \leq u_0$. Also, notice that the system is open-loop stable. Hence n_2 cannot grow without bound if the average input

u is zero. On the right hand side, the same experiment is performed but the parameter estimates are constrained to lie in Υ at every instant of time. The effect can be observed in row 3: as soon as the parameters hit the boundary of Υ , i.e. the controller poles become unstable, the parameter estimates are projected back in the interior of Υ in such a way that the poles lie at $|z| = 0.7$, i.e. inside the unit circle. The evolution of the control input u changes dramatically.

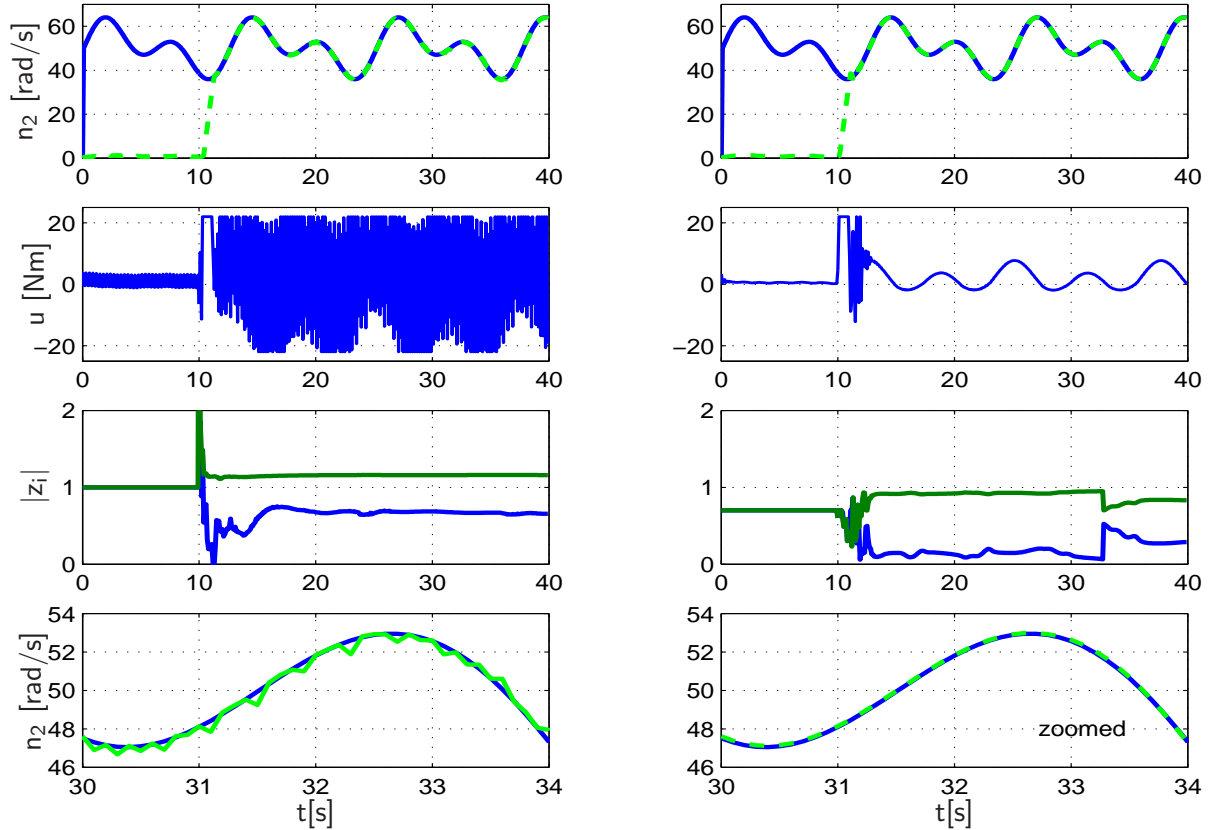


Figure 8.7: Effect of constraining the parameters to the convex set Υ (I)

To further underscore the fact that by constraining the parameters we do not merely keep the parameter error bounded but significantly alter the nature of the control input, the experiment is repeated and the results reported in figure (8.8). In contrast to above, the parameters are constrained in both cases. In the left column, the parameter estimates have been constrained to lie within Υ_1 , which was chosen in such a way as to keep the poles of the (frozen) controller transfer function within a circle of radius 1.1, i.e. $|z_i| \leq 1.1$, $i = 1, 2$ and in the right column Υ_2 was chosen to ensure that $|z_i| \leq 1$. This seemingly small difference results in a substantial reduction of the control effort spent to achieve identical tracking performance. The result obviously needs to be analyzed further analytically to be of any

theoretical significance. Here we merely document the improvement which was found to be consistent over all (>50) the experiments performed.

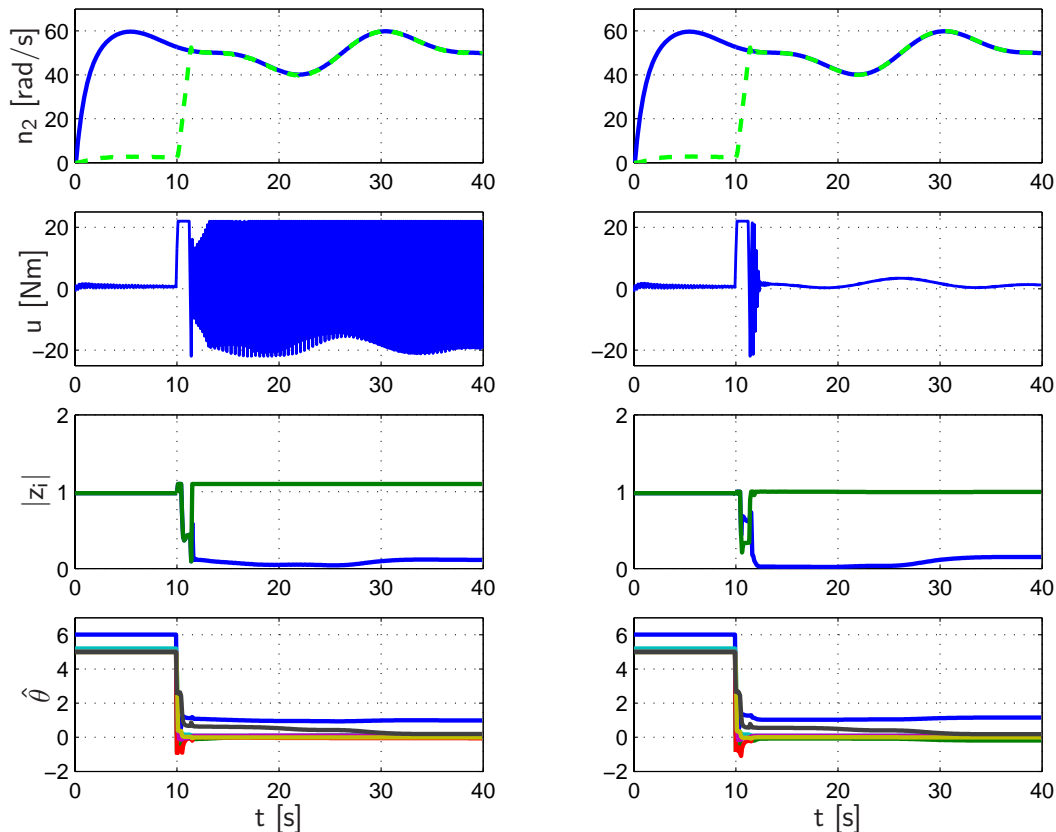


Figure 8.8: Effect of constraining the parameters to the convex set Υ (II)

Remark: Turning on the adaptation at $t = 10$ s is merely for the purpose of better illustrating the transient response of the adaptive system. Further, the effect of static friction can be suppressed by starting the adaptation only once the system is in motion.

8.3.2 Modification of the Reference System

For the purpose of analysis, the true physical parameters of the two-mass system are of interest. The inertias of the two aluminum flywheels can be calculated from the geometry of the construction to be $J_1 = 0.166 \text{ kg m}^2$ and $J_2 = 0.333 \text{ kg m}^2$. The spring constant of the flexible shaft is determined experimentally and is given by $c = 410 \text{ Nm/rad}$. Even with considerable insight into the mechanical properties of the plant it is impossible to correctly determine the damping which acts as a speed-dependent braking torque and hence cannot be separated from the effect of Coulomb friction. Its value is simply set to be $d = 11 \text{ Nms/rad}$

by fitting the model to measured input–output data. This results in an eigenfrequency of the oscillator at $f_0 = 10 \text{ Hz}$. The (approximate) knowledge of the parameters of the system has the following implications on our design.

The sampling time $T_S = 0.1 \text{ s}$ chosen in the previous experiments was inappropriate since it violates the sampling theorem which requires $1/T_S \geq 2f_0$. Hence, provided the natural oscillation of the system is excited (e.g. by a step input), the true output signal cannot be reconstructed from the sampled output (aliasing). As a consequence, the sampling interval has to be decreased. However, if the sampling time is small, another difficulty arises. Using the approximated parameters obtained above, we generate a discrete–time model of the plant in the form of an ARMA model, as in equation (8.20). The coefficients $a_i, b_i, i \in \{0, 1, 2\}$ are displayed in figure (8.9). It is seen that the coefficients b_i tend to zero as $T_S \rightarrow 0$.




Figure 8.9: Dependence of the coefficients of the discrete–time model on T_S

This is critical since $\hat{B}(q^{-1}) = \hat{b}_0 + \hat{b}_1 q^{-1} + \hat{b}_2 q^{-2}$ is in the denominator of the control law (8.25). The expected difficulty here is that if the true coefficients b_i assume values close to zero, the estimated coefficients \hat{b}_i evolve in a neighborhood of zero. While this neighborhood may be contained in the region Υ for which the polynomial $\hat{B}(q^{-1})$ is stable, the gain of the resulting controller becomes excessively large. Hence, no improvement can be expected by constraining the parameter estimates. It is also impossible to define a new region Υ_n by excluding an open neighborhood of zero from Υ (i.e. to bound the parameters $\hat{b}_i, i \in \{0, 1, 2\}$ from below), since the true parameters may lie outside that region. It is the latter aspect that can be corrected through an appropriate choice of the reference system.

As shown above, if the reference model is chosen according to (8.9) a factor $1/\mu$ appears in the denominator of the control law. As $\mu \rightarrow 0$, the gain of the control law decreases, since the closed–loop behavior approaches that of the open–loop plant. In the experiment, displayed in figure (8.11), such a reference model has been chosen using the physical parameters obtained




Figure 8.10: Control of the two-mass system using $T_S = 0.01$ s: (i)–fixed controller, (ii)–adaptive controller with $\mu = 1$ (deadbeat), (iii)– adaptive controller with $\mu = 0.1$ (generalized reference model)

above in order to determine $A(q^{-1})$ in equation (8.9). The response of the system is shown using a sinusoidal and a piecewise constant reference input. The sampling time is set to $T_S = 0.01$ s. In order to demonstrate that the above physical parameters are merely approximate ones, a non-adaptive controller of the same structure (8.25) was initialized at the “true” parameters and $\mu = 0.1$ in part (i) of the figure. The resulting control input is unacceptably large and oscillating between positive and negative saturation of the actuator and had to be turned off after a few seconds in order to avoid rupture of the shaft. A similar performance is observed in the adaptive case if the reference model is chosen to be a pure delay (part ii). This confirms our expectations, since for $T_S = 0.01$ s the coefficients b_i are already close to zero. Furthermore, it is clear that, by decreasing the sampling time, we have increased the frequency range over which the closed-loop system has to respond as a one-step delay. This can be corrected by setting $\mu = 0.1$, as in (iii). It is the only case where the performance of the system is found to be satisfactory. As we take a closer look at the evolution of the

signals, we observe that the control input is still oscillatory. So far, a recursive least-squares algorithm has been used to update the parameters. If it is substituted by a projection algorithm, we observe the same qualitative behavior, but the use of this simple algorithm opens up the possibility of optimizing the adaptive system with respect to the control effort. This is at the center of our attention in the next experiment.




Figure 8.11: Case (iii), input signals

8.3.3 Optimization of the adaptive gain

As discussed in chapter 5.2, the control effort can be measured by means of a cost function

$$J(k) = \gamma_1 u(k)^2 + \gamma_2 \Delta u(k)^2 \quad (8.32)$$

Using the adaptive gain η , introduced in equation (2.12), as an optimization variable, a procedure was outlined to determine the sequence $\{\eta^*(k)\}$, such that $J(k)$ assumes a minimal value at every instant of time $k > 0$. In order to ensure stability, η^* was constrained to lie within the interval $(0, 2)$. If the procedure is applied to the real system we obtain the results displayed in figure (8.12). In the left column, the adaptive gain satisfies $\eta \equiv 1 \forall k > 0$, while in the right column the gain has been optimized with respect to the cost function J where $\gamma_1 = 0.05$ and $\gamma_2 = 0.95$.

The objective is to track a piecewise constant, random signal of the form $r(k) = 50 + w(k)$ where $w \in [-20, 20]$ with minimal control effort. In row 5 of the figure it is seen that the

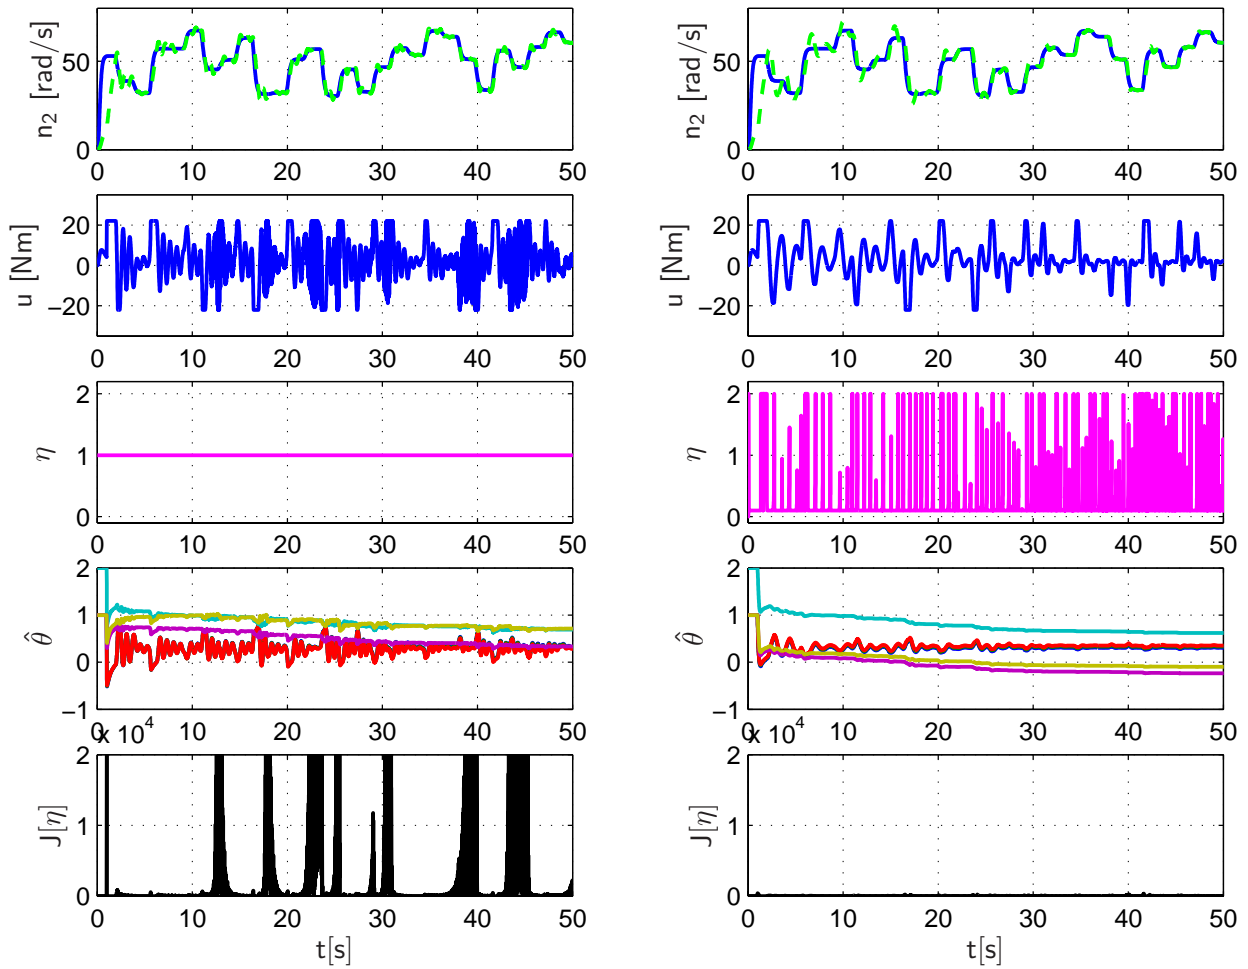


Figure 8.12: Tracking of a **random** reference input using a standard (left) and optimized (right) adaptive scheme

optimized version of the adaptive controller hardly produces any cost. As a consequence, the input signal u is less oscillatory than in the case where $\eta \equiv 1$ while the performance of the output n_2 is equivalent in both cases. The same experiment is repeated for a sinusoidal reference input, in figure (8.13). It is seen that by optimizing the gain and, hence, the evolution of the estimates in parameter space, the same output performance as in the standard case can be obtained with much less control effort.

8.3.4 Disturbance Rejection

The very fact that adaptive control is applied to a real-world process implies that it has to cope with disturbances such as noise and unmodelled dynamics. The modifications of the

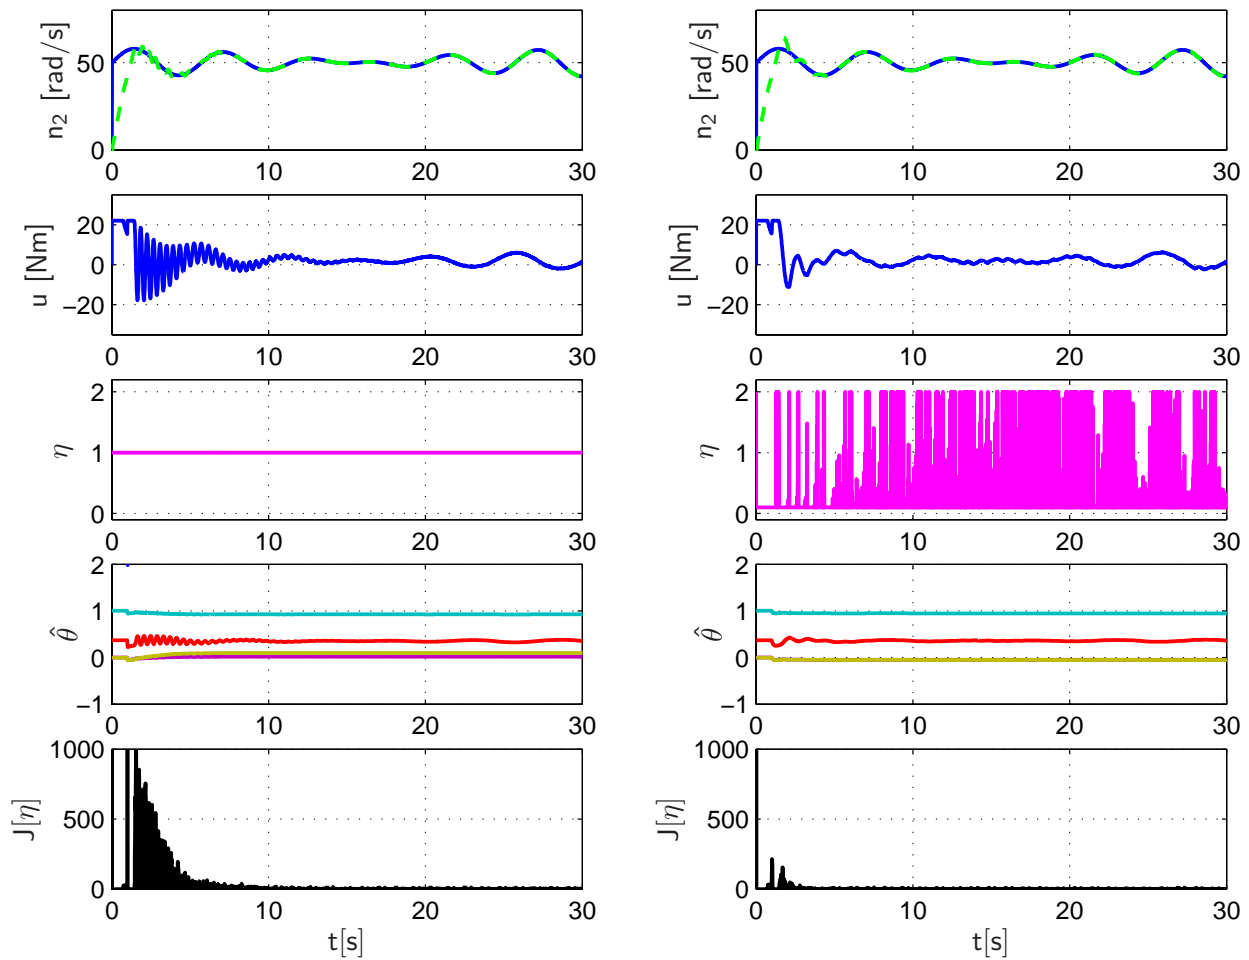


Figure 8.13: Tracking of a **sinusoidal** reference input using a standard (left) and optimized (right) adaptive scheme

adaptive scheme regarding the region in which the parameter estimates evolve as well as the step-size by which they are updated may be viewed as an effort to reject disturbances which are intrinsic to real-world applications.

In the presence of unmodelled dynamics, the nature of the control input is not only a question of performance but also of stability. As discussed in chapter 5.3, if the control input excites the unmodelled part of the plant, a disturbance signal results which may cause the estimates to drift to an unstable region in parameter space. Such a disturbance may be regarded a “hidden” one, since it does not appear when the parameters are known and a linear control concept is used (unless the reference input itself has significant high-frequency content).

In the following experiments we are interested in “explicit” deterministic disturbances, which enter the system in the form of a disturbance torque v on the load-side of the two-mass

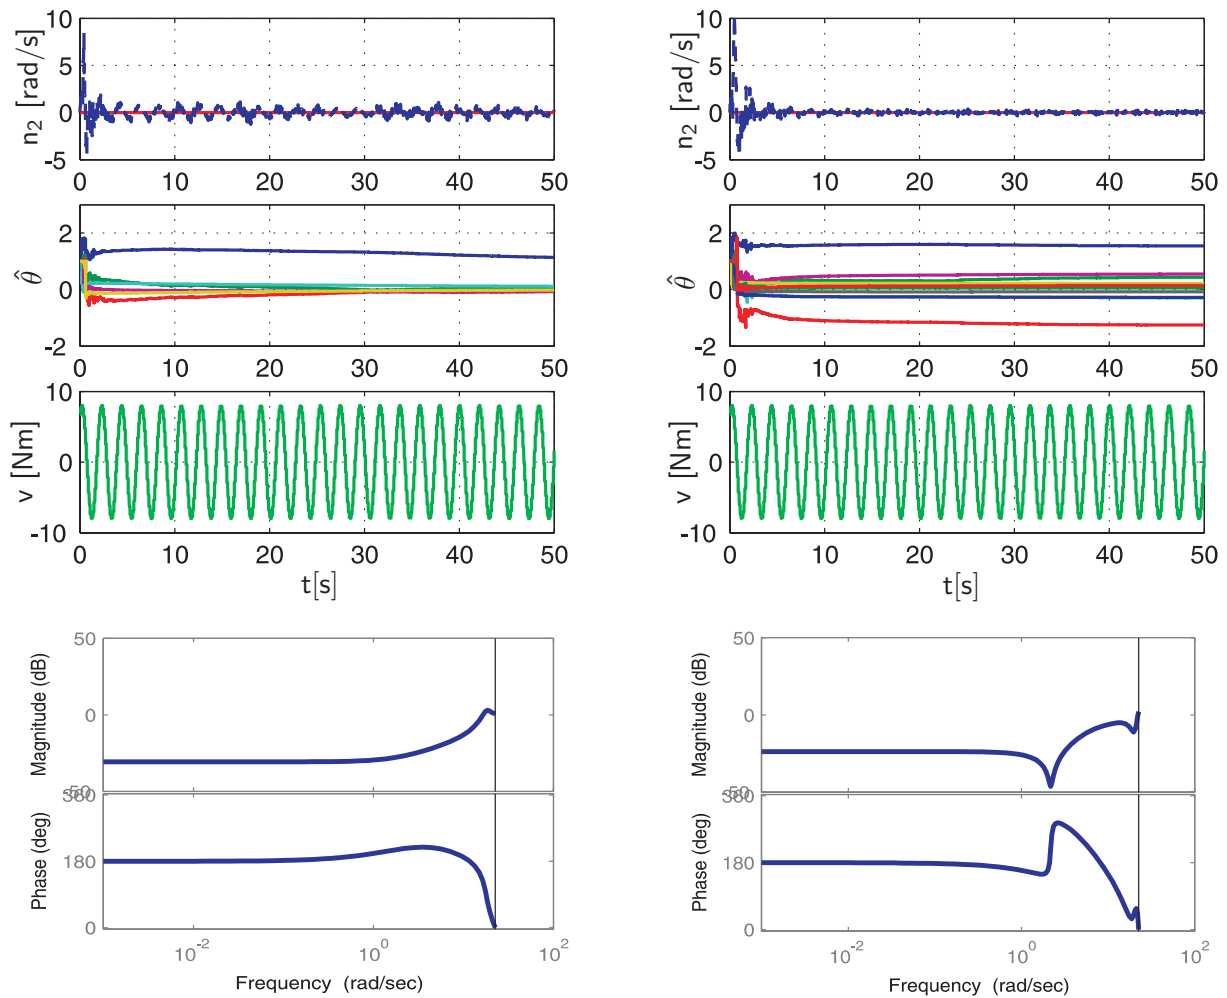


Figure 8.14: Rejection of a sinusoidal disturbance torque v with unknown frequency

system. Such a disturbance typically arises due to variations of the load or due to an out-of-balance rotational mass. It may consequently be assumed to be either piecewise-constant or sinusoidal. In the latter case, an approach has been presented (\rightarrow chapter 4.1.1) which consisted in augmenting the state-space of the system by the dimension of a disturbance generating system. This can be appreciated in the next experiment displayed in figure (8.14). The objective here is to regulate the speed n_2 of the output to zero in spite of the presence of a sinusoidal disturbance of unknown frequency. The amplitude of the disturbance torque is equal to almost half the maximum torque generated by the electrical drives $|v| = 10 \text{ Nm}$. To the left of the figure the same controller as in the disturbance-free case is used, i.e. the order of the system has not been augmented. It is seen that the output error remains nonzero. In the right part of the figure the order has been increased by two. Notice that we have a

total of ten unknown parameters in this case. It is seen that the disturbance can be rejected completely.

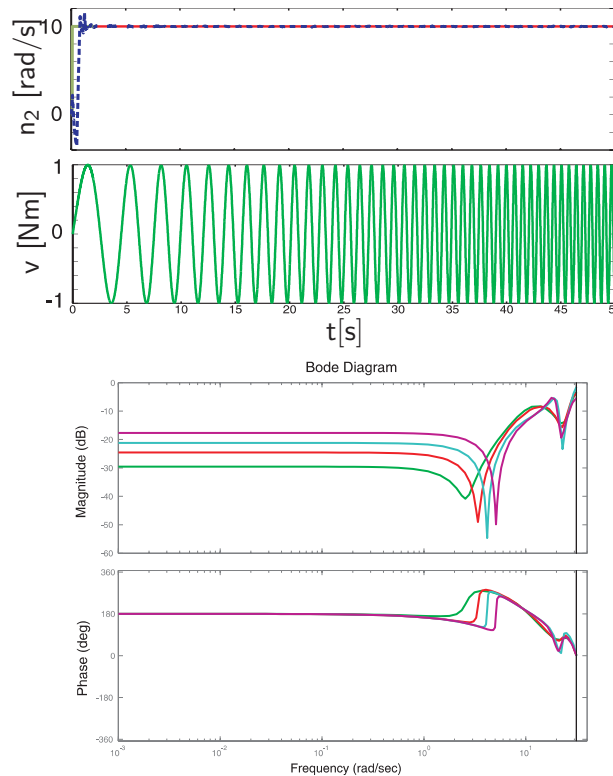


Figure 8.15: Rejection of a time-varying sinusoidal disturbance torque v

The last two rows of the figure contain a Bode diagram of the disturbance transfer function obtained after the parameters converge to some constant values. In the augmented system we find a notch at the disturbance frequency $\omega_v = 3 \text{ rad/s}$. Notice that since the controller is adaptive the location of the notch is adjusted such that it always corresponds to ω_v . In the non-augmented case, the parameters cannot converge anywhere if the disturbance persists. Hence it does not make any sense to determine the frequency response of the system. However, for comparison purposes, the Bode diagram has been computed using an average of the parameter estimates obtained in the interval $[40 \text{ s} \dots 50 \text{ s}]$. The same experiment is conducted using a sinusoidal disturbance with (slowly) time-varying frequency. It is seen that the notch in the disturbance transfer function tracks the disturbance frequency ω_v , figure (8.15). Again, the Bode diagram is determined from a (time-invariant) average of the parameter estimates $\hat{\theta}$ over a finite window. Naturally, the above procedure for minimizing the control effort can also be adopted in combination with a controller of augmented order. In order to illustrate this, the adaptive gain η of a 5th-order adaptive controller is optimized.

The two-mass system is required to track a piecewise-constant random signal *and* reject a sinusoidal disturbance of amplitude $|v| = 10 \text{ Nm}$. In the left column of figure (8.16), with constant gain $\eta \equiv 1$, significant control effort is spent to achieve the goal. In the right column, with time-varying gain $\eta^*(k)$, the effort (displayed in row 5) is minimized and the resulting control input less oscillatory.

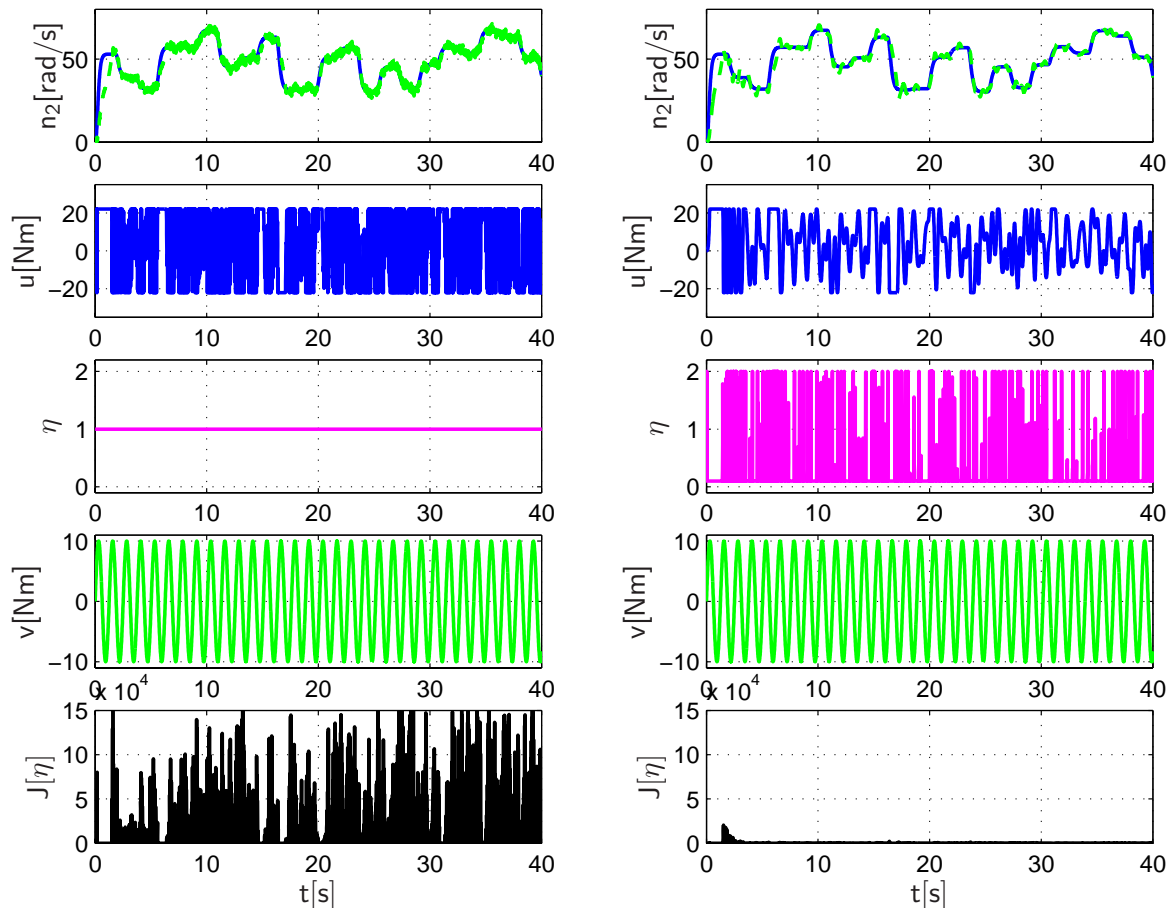


Figure 8.16: Rejection of a sinusoidal disturbance with minimal control effort

If the disturbance is piecewise constant, the order of the controller does not have to be augmented. As indicated in equation (8.18), the constant disturbance can be treated as an unknown parameter. In other words, the disturbance is converted into a parametric uncertainty. The parameter is assumed to be constant over an interval of some minimum length but then changes discontinuously from one unknown level to another. If the time-variation occurs infrequently, regular adaptive control is able to cope with the problem as illustrated by the experiment contained in figure (8.17).

If the recursive least-squares algorithm is used to adjust the parameters, a forgetting factor




Figure 8.17: Rejection of a piecewise constant disturbance

has to be introduced, or –alternatively– the matrix adaptation gain “revitalized”. Such revitalization is carried out by resetting the gain preferably at those instants at which a time–variation occurs. This requires a mechanism for detecting the variation. Secondly, it is seen that if the parameters do not remain constant over a sufficiently long interval, the plant becomes unstable. It is in such rapidly time–varying situations that the multi–model approach (MMST) is found to be effective.

8.3.5 Adaptive Control using Multiple Models

A central question in MMST regards the location of the fixed models. In the following experiment, $N = 25$ models were located in parameter space on an evenly spaced scale along the $\hat{\theta}_{[7]}$ –axis, where $\hat{\theta}_{[7]}$ denotes the element of the vector $\hat{\theta}$ corresponding to the constant disturbance v_0 . The models represent 25 different levels of v_0 within a range of $-10 Nm$ to $+10 Nm$ including a model for $v_0 = 0 Nm$. All other parameters $\hat{\theta}_{[1]}, \dots, \hat{\theta}_{[6]}$ of the 25 models were initialized at the same values obtained from previous experiments where no disturbance was present. The switching criterion was chosen to be of the form (iii) defined in equation (6.1), except that the summation is carried out over a window of finite

length, i.e. $J_i(k) = 0.5 \epsilon_i^2(k) + 0.5 \sum_{\nu=k-T_0}^k \epsilon_i^2(\nu)$, where $T_0 = 5$. Let us first observe the performance of the system if a single adaptive model is used (\rightarrow figure 8.18). The control




Figure 8.18: Adaptive control in the presence of disturbances using a single model

input is oscillatory and the signals remain bounded only because the amplitude of the input was hard-limited by a saturation nonlinearity. With the above choice of multiple models the performance can be improved substantially. At every instant $k > 0$, an adaptive model is initialized using the parameters of the fixed model that performs best, i.e. $\hat{\theta}(k) = \theta_f$ where $J_f(k) = \min_{i \in \{1 \dots 25\}} J_i(k)$ and discarded once another fixed model θ_g surpasses the performance of the adaptive model. At this point, a new adaptive model is initialized at θ_g . Since the fixed models lie close to the plant in parameter space, the choice of the models almost directly reflects the time-variation of the load. This can be verified in row 4 of figure (8.19). The MMST approach enables us to detect changes in the environment and provide a quick response by initiating an appropriate control action based on the best model. Whenever the adaptive model is reinitialized at some fixed point in parameter space, the corresponding gain matrix used in the recursive least squares algorithm is reset, so that adaptation from that point is fast. The performance of the system is seen to be far superior

to the single model case both from the point of view of the output error as well as the nature of the control input.




Figure 8.19: Adaptive control in the presence of disturbances using multiple models

This concludes our discussion of the experimental results as well as the thesis itself. Future work needs to be done regarding the prior information required to improve performance in a time-varying environment. The procedure introduced in chapter seven, which employs self-organization and learning principles to locate multiple models effectively in parameter space represents a most promising avenue of future research. Specific future problem areas have been formulated at the end of chapters five and seven and are not repeated here.

Having investigated the behavior of adaptive systems in both theory and practice it is the conviction of the author that the *nonlinear* nature of adaptive schemes is wrongly regarded as an obstacle and should be exploited instead. In chapter five, the step-size η is considered an additional design variable by which the performance of the system could be optimized. This is impossible if the parameters are trained off-line, i.e. using a linear control approach combined with a parameter estimator as in the industrial example discussed in chapter two. Even though the proof of stability (presented in chapter three) offers less insights than the stability

criteria obtained for linear systems, a lot of intuition was developed throughout the thesis regarding the behavior of adaptive systems. Chapters four, five and six develop new strategies for rejecting disturbances and improving the performance even when large uncertainties are present. The solutions proposed are not derived from linear design methodologies but use distinctly *nonlinear* features to generate a new kind of system behavior. Still, our choice of algorithms is very limited and mainly dictated by mathematical tractability. Adaptation as it occurs in biology has a much more far-reaching character regarding the way information is processed, stored and used for control. Emulating this ability will be the goal of future investigations. This requires new adaptive configurations, such as the one investigated in chapter seven. Again, it is a *nonlinear* function that governs the evolution of multiple adaptive models in such a way that they converge to different points in parameter space.

The increasing appearance of nonlinear strategies in many theoretical developments is in sharp contrast to what is currently being implemented in actual real-world applications. Linear strategies represent the preferred design technique despite the fact that, in view of growing demands on system performance, new features such as self-tuning and adaptation are called for. The acceptance of nonlinear control schemes is traditionally low since no intuition has been acquired regarding the behavior of the controller in a real-world environment. In this thesis, the emphasis was on investigating the behavior of adaptive controllers in the presence of disturbances. The choice of the problem classes was itself motivated by practical applications and includes external disturbances, unmodelled dynamics and time-variations. The experimental results presented in chapter eight strongly support the effectiveness of the proposed solutions. No outcome seems more desirable to the author than having increased the esteemed reader's motivation to further explore the potential and intuition behind the *nonlinear* process called adaptation.

Nomenclature

η	adaptive gain (step-size of the parameter estimator)
η^*	optimized adaptive gain
μ	positive real variable
ϕ	regression vector
θ_0	true system parameters
$\hat{\theta}$	parameter estimates (general)
$\hat{\theta}_i$	parameter estimates (fixed model)
$\hat{\theta}_a$	parameter estimates (adaptive model)
$\hat{\theta}_{M^*}$	critical trajectory of parameter estimates tending to M^* (chapter 7)
θ^*	constant parameters, e.g. after convergence
ε	small positive number defining a neighborhood around a critical point
$\epsilon(k)$	identification error
$\bar{\epsilon}(k)$	augmented error
$\xi(k)$	residual error due to unmodelled dynamics
χ	upper bound on the control error
ζ_0	state variables corresponding to the zero dynamics
ω	frequency [rad/s]
Σ	plant
Ξ	model
Σ_v	disturbance model
$\bar{\Sigma}$	composite system $\Sigma \circ \Sigma_v$
Γ	controller
Υ	region in parameter space
Ω	index set: $\Omega = \{i \in \mathbb{N} \mid i \leq N, N > 0\}$
A	system matrix $A \in \mathbb{R}^{n \times n}$ (state-space representation)
b	input vector $b \in \mathbb{R}^{n \times 1}$
c	output vector $c \in \mathbb{R}^{1 \times n}$
$A(q^{-1})$	polynomials in q^{-1} , similarly: $B(q^{-1}), C(q^{-1}), \dots$
c_X	spring constant
d_X	damping constant
d	relative degree
d_{ij}	$\ \hat{\theta}_i - \theta_j\ $, i.e. distance of model i to plant j
$e(k)$	control error

f, g, h	nonlinear maps
i	integer, index
k	discrete time instant
k_0	initial time
p	dimension of parameter space
n	order of the system
\bar{n}	reduced order of the system
n_v	order of the disturbance model
n_X	$\deg[X]$, where X is a polynomial
s_i	amount by which model $\hat{\theta}_i$ is updated (“step”)
t	time (in seconds)
q^{-1}	delay operator (discrete-time)
$v(k)$	deterministic disturbance
v_0	constant disturbance
\mathcal{F}	nonlinear map (NARMA-model)
\mathcal{G}	nonlinear control law
\mathcal{H}	convex hull
\mathcal{N}	neural network approximating \mathcal{F} or \mathcal{G}
$\mathcal{U}, \mathcal{X}, \mathcal{Y}$	neighborhoods of the origin
$H(\cdot)$	transfer function
I	finite interval of time
J	performance criterion
K	interval over which the plant parameters remain constant
M_P, P	matrix adaptation gains
M^*	critical model configuration (chapter 7)
N	number of estimation models (MMAC)
S	set of parameters assumed by a time-varying plant
T	interval of finite length
T_S	sampling time, i.e. $t = kT_S$
$V(k)$	Lyapunov function
x	state vector
z	state vector (in transformed coordinates)
MMST	Multiple Models, Switching & Tuning
SIC	Simultaneous Identification and Control

Bibliography

- [1] ÅSTRÖM, K.J.: *Self-Tuning Regulators – design principles and applications*, in "Applications of Adaptive Control". ed. K.S. Narendra and R.V. Monopoli, Academic Press, New York, 1980.
- [2] ÅSTRÖM, K.J.: *Analysis of Rohrs counterexamples to adaptive control*. Proc. 22nd IEEE Conference on Decision and Control, San Antonio, TX, 1983.
- [3] ÅSTRÖM, K.J.; HAGANDER, P.; STERNBY J.: *Zeros of Sampled Systems*, Automatica, Vol.20, No. 1, pp. 31–38, 1984.
- [4] ÅSTRÖM, K.J.; WITTENMARK, B.: *Adaptive Control*, Addison-Wesley Publishing Company, 1995.
- [5] ANDERSON, B.D.O.; BITMEAD, R.R.; JOHNSON, C.R.; KOKOTOVIC, P.V.; KOSUT, R.L.; MAREELS, I.M.Y; PRALY, L.; RIEDLE, B.D.: *Stability of adaptive systems: Passivity and averaging analysis*. MIT Press, 1986.
- [6] ANNASWAMY, A.M.; KÁRASON, S.P: *Discrete-Time Adaptive Control in the Presence of Input Constraints*, Automatica, 31, 1421–1431, 1995.
- [7] AEYELS, D.: *Generic observability of differentiable systems*, SIAM J. Control and Optimization, vol. 19, pp. 595–603, 1981.
- [8] BELLMAN, R.E.: *Adaptive Control Processes – A Guided Tour*. Princeton University Press, Princeton 1961.
- [9] CABRERA, J.B.D; NARENDRA, K.S.: *Issues in the application of neural networks for tracking based on inverse control*. IEEE Transactions on Automatic Control, vol. 44, no. 11, pp. 2007-2027, 1999.
- [10] CHEN, S.; COWAN, C.F.N; GRANT, P.M.: *Orthogonal least squares algorithm for radial basis function networks*. IEEE Transactions on Neural Networks, vol. 2, pp. 302-309, 1991.

- [11] CHEN, L.; NARENDRA, K.S.: *Nonlinear adaptive control using neural networks and multiple models*. Automatica, 37(8), pp. 1245-1255, 2001.
- [12] CLARKE, R.S.: *Dynamical Systems: stability, symbolic dynamics and chaos*. CRC Press, Boca Raton, 1995.
- [13] CORDERO, A.O.; MAYNE, D.Q.: *Deterministic convergence of a self-tuning regulator with variable forgetting factor*. Proc. IEE, Vol. 128, No. 1, pp. 19-23, 1981.
- [14] CYBENKO, G.: *Approximation by superposition of a sigmoidal function*. Mathematics of Control, Signals and Systems, 2, pp. 303-314, 1989.
- [15] DESOER, C.A.; VIDYASAGAR, M.: *Feedback systems: Input-output properties*. Academic Press, New York, 1975.
- [16] DRENICK, R.F.; SHAHBENDER, R.A.: *Adaptive Servomechanisms*. AIEE Transactions 76: pp. 286-292, 1957.
- [17] EGARDT, B.: *Stability of Adaptive Controllers*. Springer Verlag, Berlin 1979.
- [18] FEILER, M.J.; NARENDRA, K.S.: *Adaptive Control in the Presence of Disturbances*. Technical Report No. 0101, Center for Systems Science, Yale University, New Haven, CT, USA, January 2001.
- [19] FEILER, M.; SCHRÖDER, D.: *Adaptive Regelung mit Modellen reduzierter Ordnung*. In: GMA-Workshop 2003 Neuere Verfahren der Regelungstechnik, Interlaken, Schweiz, 28.09.-01.10.2003. Hrsg.: VDI/VDE-Gesellschaft. Düsseldorf: VDI-Verlag GmbH, 2003.
- [20] FEILER, M.; WESTERMAIER, C.; SCHRÖDER, D.: *Adaptive Speed Control of a Two-Mass System*. In: Proceedings of the 2003 IEEE International Conference on Control Applications p. 1112-1117, Istanbul, Turkey, 2003.
- [21] FEILER, M.J.; NARENDRA, K.S.: *Adaptive Control of Rapidly Time-Varying Systems*. In: Proceedings of the 12th Yale Workshop on Adaptive and Learning Systems, p. 33-42, Yale University, New Haven, CT, USA, 2003.
- [22] FEILER, M.J.; NARENDRA, K.S.: *Simultaneous Identification and Control of Linear Time-Varying Plants using Multiple Models*. Technical Report No. 0401, Center for Systems Science, Yale University, New Haven, CT, USA, March 2004.
- [23] NARENDRA, K.S.; DRIOLLET, O.A.; FEILER, M.J.; KOSHY, G.: *Adaptive Control Using Multiple Models, Switching and Tuning*. International Journal of Adaptive Control and Signal Processing, vol. 17, pp.87-102, 2003.

- [24] FRANCIS, B.A; WONHAM, W.M.: *The internal model principle of control theory*, Automatica, Vol. 12, pp. 457–465.
- [25] GOODWIN, G.C.; LOZANO LEAL, R.; MAYNE, D.Q.; MIDDLETON, R.H.: *Rapprochement between Continuous and Discrete Model Reference Adaptive Control*. Automatica, Vol. 22, No. 2, pp. 199–207, 1986.
- [26] GOODWIN, G.C.; RAMADGE, P.J.; CAINES, P.E.: *Discrete-Time Multivariable Adaptive Control*. IEEE Transactions on Automatic Control, vol. 25, no. 3, 449–456, 1980.
- [27] G. C. GOODWIN, K.S. SIN: *Adaptive Filtering, Prediction and Control*. Prentice-Hall, 1984.
- [28] GUCKENHEIMER, J.; HOLMES, P.: *Nonlinear Oscillations, Dynamical Systems and Bifurcations of Vector Fields*. Springer Verlag, New York, 1983.
- [29] HARUNO, M.; WOLPERT, D.M. AND KAWATO M.: *MOSAIC Model for Sensorimotor Learning and Control*. Neural Computation, 13:2201–2220, 2001.
- [30] HORNIK, K.; STINCHCOMBE, M.; WHITE, H.: *Multi-layer feedforward networks are universal approximators*, Neural Networks, vol. 2, pp. 359–366, 1989.
- [31] IEE COLLOQUIUM ON: *Adaptive Controllers in Practice*, Digest No.1997/176, Coventry, UK, 1997.
- [32] IOANNOU, P.A; KOKOTOVIC, P.V.: *Adaptive Systems with Reduced Models*. Springer Verlag, New York 1983.
- [33] IOANNOU, P.A.; SUN, J.: *Robust Adaptive Control*. Prentice-Hall, Upper Saddle River, NJ, 1996.
- [34] IOANNOU, P.A; TSAKALIS,K.S.: *A robust direct adaptive controller*. IEEE Transactions on Automatic Control, vol. 31, pp.1033-1043, 1987.
- [35] ISERMANN, R.: *Identifikation dynamischer Systeme 1,2*. Springer Verlag. 1992.
- [36] KAILATH, TH.: *Linear Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1980.
- [37] KOKOTOVIC, P.V.; KHALIL, H.K.; O'REILLY, J.: *Singular Perturbation Methods in Control: Analysis and Design*. Academic Press, New York, 1986.

- [38] KOKOTOVIC, P.V.; RIEDLE, B.; PRALY, L.: *On a stability criterion for continuous slow adaptation*. Systems & Control Letters, vol. 6, pp. 7–14, 1985.
- [39] KÖHLER, S.: *Entwurf eines robusten Drehzahlreglers*. Diplomarbeit, Friedrich–Alexander Universität Erlangen–Nürnberg, 2001.
- [40] KREISSELMEIER, G.; NARENDRA, K.S.: *Stable model reference adaptive control in the presence of bounded disturbances*. IEEE Transactions on Automatic Control, vol. 27, pp. 1169–1175, 1982.
- [41] KREISSELMEIER, G.; ANDERSON, B.D.O.: *Robust Model Reference Adaptive Control*. IEEE Transactions on Automatic Control, vol. 31, pp. 127–134, 1986.
- [42] LYAPUNOV, A.M.: *Problème général de la stabilité du mouvement*. 1892, in Russian, 1907: French Translation, photo–reproduced as Annals of Mathematics Study No. 17, Princeton University Press, Princeton, New Jersey, 1949.
- [43] LANDAU, I.D.; LOZANO, R. AND M’SAAD, M.: *Adaptive Control*. Springer Verlag, London, 1998.
- [44] LEVIN, A.U; NARENDRA, K.S.: *Control of nonlinear dynamical systems using neural networks –Part II: Observability, identification, and control*. IEEE Transactions on Neural Networks, vol. 7, pp.30–42, 1996.
- [45] LJUNG, L.: *System Identification – Theory for the User*. Prentice–Hall, Englewood Cliffs, NJ, 1999.
- [46] MONOPOLI, R.V.: *Model Reference Adaptive Control with an augmented error signal*. IEEE Transactions on Automatic Control, vol. 19, pp. 474–482, 1974.
- [47] MONOPOLI, R.V.: *Adaptive Control for systems with hard saturation*. Proc. IEEE Conf. on Decision and Control, Houston, TX, pp. 841–843, 1975.
- [48] MORSE, A.S.: *Global stability of parameter adaptive control systems*. IEEE Transactions on Automatic Control, vol. 25, no. 3, pp. 433–439, 1980.
- [49] MORGAN, A.P.; NARENDRA, K.S.: *On the stability of nonautonomous differential equations $\dot{x} = [A + B(t)]x$ with skew–symmetric matrix $B(t)$* . SIAM Journal of Control and Optimization, vol. 15, pp.163–176, 1977.
- [50] NARENDRA, K.S.: *Neural networks for control– theory and practice*. Proceedings of the IEEE, vol. 84, pp.1385–1406, 1996.

- [51] NARENDRA, K.S.; ANNASWAMY, A.M.: *Stable Adaptive Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1989.
- [52] NARENDRA, K.S.; ANNASWAMY, A.M.: *A new adaptive law for robust adaptive control without persistent excitation*. IEEE Transactions on Automatic Control, vol. 32, pp.134-145, 1987.
- [53] NARENDRA, K.S.; ANNASWAMY, A.M.: *Persistent Excitation and Robust Adaptive Algorithms*. Proc. of the 3rd Yale Workshop on Appl. of Adaptive Systems Theory, Yale University, pp. 11-18, 1983.
- [54] NARENDRA, K.S.; BALAKRISHNAN, J.: *Improving transient response of adaptive control systems using multiple models and switching*. Proc. of the 7th Yale Workshop on Adaptive and Learning Systems, Yale University, 1992.
- [55] NARENDRA, K.S.; BALAKRISHNAN, J.: *Improving transient response of adaptive control systems using multiple models and switching*. IEEE Transactions on Automatic Control, 39(9), pp. 1861-1866, 1994.
- [56] NARENDRA, K.S.; BALAKRISHNAN, J. AND CILIZ M.K.: *Adaptation and Learning Using Multiple Models, Switching and Tuning*. IEEE Control Systems Magazine, 15(3):37-51, June 1995.
- [57] NARENDRA, K.S.; BALAKRISHNAN, J.: *Adaptive Control Using Multiple Models*. IEEE Transactions on Automatic Control, Vol. 42, No. 2, Feb. 1997.
- [58] NARENDRA, K.S.; LEE, A.M.: *Stable direct adaptive control of time-varying discrete-time systems*. Technical Report No. 8720, Center for Systems Science, Yale University, New Haven, 1987.
- [59] NARENDRA, K.S.; LIN, Y.H.; VALAVANI, L.S.: *Stable Adaptive Controller Design, Part II: Proof of Stability*. IEEE Transactions on Automatic Control, Vol. 25, No. 3, pp. 440-449, 1980.
- [60] NARENDRA, K.S.; PARTHASARATHY, K.: *Gradient methods for the optimization of dynamical systems containing neural networks*. IEEE Transactions on Neural Networks, vol. 2, pp. 252-262, 1991.
- [61] NARENDRA, K.S.; MUKHOPADHYAY, S.: *Adaptive Control Using Neural Networks and Approximate Models*. IEEE Transactions on Neural Networks, vol. 8, pp. 475-485, 1997.

- [62] NARENDRA, K.S.; TAYLOR, J.H.: *Frequency Domain Criteria for Absolute Stability*. Academic Press, New York, September 1973.
- [63] NARENDRA, K.S.; XIANG, C.: *Adaptive Control of Discrete-Time Systems Using Multiple Models*. IEEE Transactions on Automatic Control, Vol. 45, No. 9, Sept. 2000.
- [64] PETERSON, S.; NARENDRA, K.S.: *Bounded error adaptive control*. IEEE Transactions on Automatic Control, vol. 27, pp.1161-1168, 1982.
- [65] PEITGEN, H.O.; RICHTER, P.H.: *The Beauty of Fractals. Images of Complex Dynamical Systems*, Springer-Verlag, 1986.
- [66] PEITGEN, H.O.; SAUPE, D. (EDS.): *The Science of Fractal Images*, Springer-Verlag, 1988.
- [67] PRALY, L.: *Robustness of model reference adaptive control*. Proc. of the 3rd Yale Workshop on Appl. of Adaptive Systems Theory, Yale University, 1983.
- [68] PRALY, L.: *Robust model reference adaptive controllers-Part I: Stability Analysis*. Proc. of the 23rd IEEE Decision and Control Conference, pp. 1009-1014, Dec. 1984.
- [69] ROHRS, C.; VALAVANI, L.; ATHANS, M. AND STEIN, G.: *Robustness of continuous-time adaptive control algorithms in the presence of unmodeled dynamics*. IEEE Transactions on Automatic Control, vol. 30, pp.881-889, 1985.
- [70] RUDIN, W.: *Principles of Mathematical Analysis*. McGraw-Hill Mathematical Series, 1964.
- [71] SCHRÖDER, D.: *Elektrische Antriebe – Regelung von Antriebssystemen*. Springer-Verlag, Berlin Heidelberg, 2. Auflage, 2001.
- [72] SONTAG, E.D.: *An algebraic approach to bounded controllability of linear systems*. Int. J. Control, vol. 39, pp. 181-188, 1984.
- [73] SONTAG, E.D.: *Mathematical Control Theory, Deterministic Finite Dimensional Systems*. 2nd edition, Springer Verlag, New-York 1998.
- [74] TSAKALIS, K.; IOANNOU, P.: *Adaptive Control of linear time-varying plants: A new model reference control structure*. Technical Report No. 86-10-1, University of Southern California, 1987.
- [75] WIDROW, B.; LOHR, M.A.: *30 Years of Adaptive Neural Networks: Perception, Madeline, and Backpropagation* Proc. IEEE, 1990, vol. 78, pp. 1415-1442.

- [76] WERBOS, P.J.: *Backpropagation through time: What it does and how to do it*. Proc. IEEE, 1990, vol. 78, pp. 1550-1560.