

Hidden Markov Model

Signal Processing and Control

Iain B. Collings

January 1995

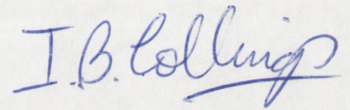
*A thesis submitted for the degree of Doctor of Philosophy
of the Australian National University*

Department of Systems Engineering
Research School of Information Sciences and Engineering
The Australian National University

Statement of Originality

These doctoral studies were conducted with Professor John B. Moore as supervisor, and Dr Iven M.Y. Mareels and Dr. Matthew R. James as advisors.

The work presented in this thesis is the result of original research carried out by myself, in collaboration with others, while enrolled in the Department of Systems Engineering as a Doctor of Philosophy student. It has not been submitted for any other degree or award in any other university or educational institution.

A handwritten signature in blue ink that reads "I.B. Collings". The signature is written in a cursive style with a prominent underline for the name.

Iain B. Collings

January 1995

Acknowledgements

I would like to thank my supervisor John Moore for his enthusiasm and encouragement, his never ending desire to generate new ideas, his insights and guidance, and for teaching me the value of co-operative research. I thank Vikram Krishnamurthy for introducing me to the area of hidden Markov models and for stimulating discussions on a great many of life's questions. I have also enjoyed the exciting working environment provided by all the staff and students of the department, especially those residing in the lakeview and rank4 student rooms.

A number of people have made useful comments and contributions to the results obtained in this thesis. In addition to John and Vikram, I would particularly like to thank Robert Elliott, Lang White, Matt James and Iven Mareels for their insight.

During my studies I have been fortunate to visit, for extended periods, the University of Alberta, Canada (Department of Statistics and Applied Probability), the Australian Defence Science Technology Organisation (Communications Division), and the University of Bremen, Germany (Institute of Dynamical Systems). I would like to thank everyone concerned with these visits for their enthusiasm and support.

Last, but by no means least, I would like to thank Carolyn Bull for her support and friendship throughout my doctorate, as well as her cute smile.

Abstract

The theory of Markov processes has many applications in areas such as communications systems, speech processing, biological signal processing, and pattern recognition. In such systems the underlying information is *hidden* in a noisy environment, resulting in system models being termed *hidden Markov models*, (HMMs). This thesis develops new and novel information-state techniques for HMM identification, adaptive estimation, and control. A number of application areas are also explored.

New HMM identification schemes are developed, based on recursive prediction error techniques. Reduced computational complexity schemes are also generated. The resulting algorithms are shown to have better asymptotic performance than previously known schemes. An important feature of these algorithms, is that they are on-line. This is true for most of the techniques presented in this thesis. It is also of fundamental importance for real time systems.

Applications of the new HMM signal processing techniques to communications systems are investigated for the cases of quadrature amplitude modulated (QAM) and differentially encoded phase shift keyed (DPSK) signals. Such systems are shown to require extensions to the basic hidden Markov model representation, as they result in hybrid models with both discrete-range and continuous-valued states. The approach of this thesis is to use conditional coupling of schemes which would otherwise be optimal, to generate state and parameter estimation algorithms for the HMM. The optimal schemes which are used include the Kalman filter and modifications to standard HMM filtering schemes. Extensive simulation studies have been carried out to confirm convergence, robustness, and the superior performance of these adaptive HMM signal processing techniques.

This thesis also considers control problems. Risk-sensitive, or exponential performance criteria, regulation and tracking results are derived for both linear systems and HMMs. Simulation studies are presented which demonstrate the effect on control, of variations in risk sensitivity. Connections to risk-sensitive filtering are also discussed.

Preface

This thesis is divided into 7 chapters.

- Chapter 1 introduces the topics and problems which are considered in this thesis. It presents an introduction to the area of hidden Markov models (HMMs), and gives a review of standard filtering techniques. It also introduces communications systems and risk-sensitive filtering and control problems, thereby providing motivation for the study of HMMs.
- Chapter 2 considers the problem of parameter identification for an HMM. The model is formulated in such a way as to allow standard nonlinear recursive prediction error identification techniques to be applied, thus generating on-line algorithms. Simulation studies are used to demonstrate the superior performance of this technique, compared with current schemes.
- Chapter 3 presents a modification to the parameter identification algorithms of Chapter 2, which results in reduced complexity schemes in certain cases. Simulation studies are again used to demonstrate the performance of the algorithms.
- Chapters 4 and 5 consider applications of HMM signal processing to communication systems. Such systems require adaptive estimation of time-varying parameters. The problem of adaptive demodulation for quadrature amplitude modulated (QAM) and M-ary differential phase shift keyed (MDPSK) signals in noisy fading channels is addressed. Models for such systems are nonlinear, however, with the application of HMMs, they can be represented in a new way which is bi-linear in terms of states and parameters, and thus allows the application of new coupled conditional linear information state filters.
- Chapter 6 presents solutions to risk-sensitive control problems. Both linear systems and HMM systems are considered. Simulation results are presented demonstrating the effect of variations to the desired amount of risk. In addition, connections between risk-sensitive control and filtering problems are discussed.
- Conclusions are presented in Chapter 7.

The following is a list of publications in refereed journals and conference proceedings, based on the work presented in this thesis. In some cases the conference papers contain material overlapping with the journal publications.

Journal Papers

- I.B. COLLINGS, V. KRISHNAMURTHY AND J.B. MOORE, *On-line Identification of Hidden Markov Models via Recursive Prediction Error Techniques*, IEEE Trans. on Signal Processing, Vol. 42, No. 12, pp. 3535-3539, December 1994.
- I.B. COLLINGS AND J.B. MOORE, *An HMM Approach to Adaptive Demodulation of QAM Signals in Fading Channels*, Int. Jour. of Adaptive Control and Signal Processing, Vol. 8, No. 5, pp. 457-474, October 1994.
- I.B. COLLINGS AND J.B. MOORE, *An Adaptive Hidden Markov Model Approach to FM and M-ary DPSK Demodulation in Noisy Fading Channels*, Signal Processing, accepted for publication, December 1994.
- I.B. COLLINGS, M.R. JAMES AND J.B. MOORE, *An Information State Approach to Risk-Sensitive Tracking Problems*, submitted to the Jour. of Mathematical Systems, Estimation, and Control.

Conference Papers

- I.B. COLLINGS, V. KRISHNAMURTHY AND J.B. MOORE, *Recursive Prediction Error Techniques for Adaptive Estimation of Hidden Markov Models*, in Proc. of the 12th World Congress IFAC, Sydney, Australia, Vol. V, pp. 423-426, July 1993.
- I.B. COLLINGS AND J.B. MOORE, *Adaptive Demodulation of QAM Signals in Noisy Fading Channels*, in Proc. of the 2nd Int. Conf. on Intelligent Signal Processing and Communication Systems : ISPACS, Sendai, Japan, pp. 99-104, October 1993.
- I.B. COLLINGS AND J.B. MOORE, *Adaptive HMM filters for Signals in Noisy Fading Channels*, in Proc. of the Int. Conf. on Acoustics, Speech & Signal Processing : ICASSP '94, Adelaide, Australia, Vol. 3, pp. 305-308, April 1994.
- I.B. COLLINGS, M.R. JAMES AND J.B. MOORE, *An Information State Approach to Linear/Risk-Sensitive/Quadratic/Gaussian Control*, in Proc. of the 33rd IEEE Conf. on Decision and Control, Orlando, USA, Vol. 4, pp. 3802-3807, December 1994.
- I.B. COLLINGS AND J.B. MOORE, *Identification of Hidden Markov Models with Grouped State Values*, submitted to the 1995 IFAC conference on Youth Automation.

Contents

Statement of Originality	i
Acknowledgements	iii
Abstract	v
Preface	vii
Glossary of Symbols	xiii
1 Introduction	1
1.1 Hidden Markov Models	3
1.2 Model Identification	6
1.3 Parameter Estimation	9
1.4 Communication Systems	10
1.5 Risk-Sensitive Filtering and Control	14
1.6 Algorithms Developed	17
1.7 Thesis Structure	18
2 Identification of Hidden Markov Models	21
2.1 Introduction	21
2.2 Problem Formulation	23
2.2.1 State Space Signal Model	23
2.2.2 Model Parameterisation	26
2.2.3 Parameterised Information-State Signal Model	27
2.3 Identification Algorithms	29
2.3.1 General Prediction Error Algorithms	29
2.3.2 Off-line Gauss-Newton Prediction Error Method	30
2.3.3 The Recursive Prediction Error Algorithm	30
2.4 Gradient Vector and Projection Calculations	33

2.4.1	Parameterisation on a Simplex	34
2.4.2	Parameterisation on a Sphere	36
2.4.3	RKL Scheme with Parameterisation on a Sphere	37
2.5	Implementation Considerations and Simulations	38
2.5.1	Implementation Considerations	38
2.5.2	Simulation Studies	41
2.6	Conclusions	44
2.7	Figures and Tables	44
3	Identification of HMMs with Grouped States	55
3.1	Introduction	55
3.2	Problem Formulation	57
3.2.1	State Space Signal Model	57
3.2.2	Reduced Order State Space Signal Model	59
3.2.3	Model Parameterisation	60
3.2.4	Parameterised Information State Signal Model	61
3.3	Gradient Vector and Projection Calculations	63
3.4	Implementation Considerations and Simulations	65
3.4.1	Implementation Considerations	65
3.4.2	Simulation Studies	66
3.5	Conclusions	70
3.6	Figures and Tables	70
4	HMM Processing for QAM Digital Communication Systems	75
4.1	Introduction	75
4.2	Quadrature Amplitude Modulation (QAM)	79
4.2.1	Signal Model	79
4.2.2	Channel Model	80
4.2.3	Observation Model	82
4.2.4	State Space Signal Model	82
4.2.5	Conditional Information-State Signal Model	86
4.3	Adaptive HMM Algorithms	89
4.3.1	Adaptive HMM/EKF Scheme	89
4.3.2	Adaptive HMM/KF Schemes	91
4.4	Coloured Noise Case	93

4.5	Robustness Issues	94
4.6	Implementation Considerations and Simulations	95
4.7	Conclusions	98
4.8	Figures	99
5	HMM Processing for Differentially Phase Modulated Communication Systems	103
5.1	Introduction	103
5.2	MDPSK and FM Signal Models	104
5.2.1	Initial Signal Model Formulation	104
5.2.2	Channel Model	105
5.2.3	Observation Model	106
5.2.4	State Space Signal Model	106
5.2.5	Conditional Information-State Signal Models	112
5.3	Higher Order Message Models	116
5.4	Adaptive HMM Algorithms	117
5.4.1	Adaptive HMM/EKF Scheme	118
5.4.2	Adaptive HMM/KF Schemes	119
5.5	Simulation Studies	121
5.6	Conclusions	123
5.7	Figures	124
6	Risk Sensitive Control Problems	129
6.1	Introduction	129
6.2	Linear Systems	131
6.2.1	State Space Model	132
6.2.2	Cost	132
6.2.3	Information State	133
6.2.4	Alternate Cost Representation	135
6.2.5	Dynamic Programming	136
6.2.6	Dynamic Programming Solution	137
6.3	Hidden Markov Models	139
6.3.1	State Space Model	139
6.3.2	Cost	140
6.3.3	Information State	140
6.3.4	Alternate Cost and Dynamic Programming	141

6.4	Constant Reference Input Case	142
6.4.1	Control Integrator Approach	142
6.4.2	Reference Model Integrator Approach	144
6.5	Simulation Studies	145
6.6	Risk-Sensitive Filtering Interpretations	148
6.6.1	Linear Systems	148
6.6.2	Hidden Markov Models	150
6.6.3	Dual Control	151
6.7	Conclusion	152
6.8	Figures	152
7	Conclusion	157
7.1	Overview of Thesis	157
7.1.1	State Space Models for HMMs	158
7.1.2	On-Line Algorithms	159
7.1.3	Information-State Signal Models	159
7.1.4	Conditional Coupled Filters	160
7.1.5	Robust Processing	161
7.2	Algorithms	161
7.3	Summation	162
7.4	Future Research	163
	Appendices	165
A	Gradients in a Manifold	165
B	Theorem Proofs	167
B.1	Proof of Theorem 6.1	167
B.2	Proof of Theorem 6.4	168
	Bibliography	171

Glossary

Sets and Spaces

\mathbb{C}	The complex numbers.
\mathcal{F}_k	The complete σ -field generated by a Markov process X_k .
\mathbb{R}	The real numbers.
\mathcal{S}	The set of N -dimensional unit vectors $\{e_1, e_2, \dots, e_N\}$, where $e_i = (0, \dots, 0, 1, 0, \dots, 0)' \in \mathbb{R}^N$ with 1 in the i^{th} position.
\mathcal{Y}_k	The σ -field generated by observations y_k .
Y_k	The sequence of outputs (y_0, \dots, y_k) .
\mathcal{Z}^+	The set of non-negative integers.

Notation and Operators

$\langle x, y \rangle$	The Euclidean inner product of x and y in \mathbb{R}^n .
$\mathbf{1}_N$	The N -dimensional column vector containing all 1's.
A	The transition probability matrix.
A'	The transpose of a matrix A .
\mathbf{b}	The vector of symbol probabilities.
B	The symbol probability matrix.
$\text{diag}(\mu_1, \dots, \mu_n)$	The matrix with diagonal elements (μ_1, \dots, μ_n) and all other elements zero.
$\text{diag}(x)$	The matrix with the vector x on its diagonal and all other elements zero.
$E[\cdot]$	The expectation operator.
I_N	The $N \times N$ identity matrix.
$N[m, \sigma^2]$	The normal distribution, mean m , and standard deviation σ .
σ_w	Standard deviation of white Gaussian noise.
X_k	A finite-discrete Markov state indicator vector.
x_k	A continuous-range time-varying parameter vector.
$\hat{x}_{y z}$	The expected value of x at time y , given z . i.e. $E[x_y z]$.
θ	Used to represent both a continuous state constant parameter vector, and the risk-sensitive parameter.

Abbreviations

AGC	Automatic gain control.
EKF	Extended Kalman filter.
EM	Expectation maximisation.
FM	Frequency modulation.
HMM	Hidden Markov model.
i.i.d.	Independent and identically distributed.
KF	Kalman filter.
LQG	Linear-quadratic-Gaussian.
M-ary DPSK	Differentially phase shift keyed modulation with M levels.
MAP estimate	Maximum <i>a posteriori</i> estimate.
ML	Maximum likelihood.
PEE	Parameter estimate error = $\sqrt{(1/I) \sum_{i=1}^I (x - \hat{x})^2}$.
PLL	Phase locked loop.
QAM	Quadrature amplitude modulation.
RKL	Recursive Kullback-Leibler.
RLQG	Risk-sensitive linear-quadratic-Gaussian.
RPE	Recursive prediction error.
WGN	White Gaussian noise.

Chapter 1

Introduction

Most systems, from natural processes to high-technology communication networks, can be considered to consist of an underlying, or hidden, structure coupled with a mechanism by which this structure is observed. Many such systems can be viewed in an *input/output* framework, where control signals (or inputs) are applied to the system, and system responses (or outputs) are collected. Often the first aim of systems analysis is to model accurately a system's unknown internal structure, using known input/output sequences. This form of system modelling is commonly termed *system identification*. Closely related to system identification are the areas of *signal processing* and *control*. The task of signal processing is to take a system model and use it to determine the input sequence corresponding to a given output sequence, and at the same time improve the model. For control problems the aim is to use the system model and output sequence, to generate a controlling input sequence which will cause the actual system to perform desired tasks. System identification, signal processing and control problems are, to a large extent, interdependent. By considering the interconnections a large number of tasks can be undertaken on a wide range of system types.

In seeking to classify adequately the underlying structure of a system, it is often useful to model system behaviour in terms of internal variables, termed *states*, which define the structure and are associated with the dynamics of the system. Depending on the class of system, the dynamics (or relationship governing the evolution of the state) can be nonlinear or linear, time-varying or time-invariant, deterministic or stochastic, and continuous-time or discrete-time. The states themselves can be real valued or complex valued, and belong to either a finite set or a continuous range. The

observation process, which generates an output sequence, is generally a noisy nonlinear function of the state process. The task of system identification is to choose the appropriate combination of internal variable set, state process, and observation process, which best models the actual given system.

The class of systems to be considered in this thesis has internal dynamics which are governed by Markov processes. Models for such Markov systems are nonlinear and stochastic. In the case that the internal states belong to a finite set of discrete values, and the observations are nonlinear functions of these finite-discrete states plus stochastic noise, the Markov models are termed *hidden Markov models* (HMMs). Such models are of great importance as the finite-discrete nature of the HMM has applications to many systems, for example digital communication systems and discrete event systems, where the states belong to a finite set of discrete values. This thesis deals mainly with *mixed state hidden Markov models* which contain both finite-discrete and continuous-range states. It considers problems involving model identification, signal processing, and system control, along with practical applications for these mixed state models.

The study of HMMs has grown out of the theory of Markov processes, which was introduced by A.A. Markov in the early 1900's. The first presentation of the explicit form of Markov dependence came in his 1906 work [Markov 1906], and was followed a year later by a paper introducing the basic concepts [Markov 1907]. Markov's pioneering work on dependent random variables was later published as a collected series of papers [Markov 1951]. Following its introduction, applications of Markov models were immediately abundant. Markov himself used the techniques to model the succession of consonants and vowels in the literary work "Eugene Onegin" [Markov 1913, Markov 1924]. Since then the Markov model has been used to represent an ever increasing range of systems and has provided a rich area for research.

Early investigations into the model tended to focus on the statistical properties and uniqueness questions associated with the Markov chain. It was not until much later that the HMM was developed in order to address the task of filtering in noisy environments. However, following the realisation that computation requirements for discrete-state Markov processing were, at the time, restrictively high, the excitement tended to wane. With the advent of digital computers, interest was renewed in the HMM, and applications abounded. As digital systems have replaced their analog predecessors, the finite-discrete nature of the HMM has led to its increased relevance. Recently, HMMs have been used in a wide variety of applications, including speech processing [Rabiner 1989, Levinson *et al.* 1983], frequency tracking [Streit and Barrett 1990, Xianya and

Evans 1991], biological signal processing [Chung *et al.* 1991], and image recognition [Bellegarda *et al.* 1994].

In this thesis, signal processing applications to communication systems are considered for mixed state hidden Markov models. The HMM techniques provide a new and structured way of approaching many of the challenges in modern digital communication systems. Also, the mixture of finite-discrete and continuous-range states in mixed state HMMs is found to be directly applicable, and of great use. In addition to signal processing applications, control problems are also considered in this work. An area of control theory termed *risk-sensitive* control is investigated for both linear system and HMM tracking problems. Risk-sensitive policies are useful for adding robustness to optimal controllers in the presence of uncertain models. The resulting control schemes apply similar techniques to those employed in the signal processing application to communication systems. Both these application areas (communications and tracking) are important not only to demonstrate the usefulness of HMMs for a wide range of problems not previously considered, but also to present significant new results where gains are achieved over standard approaches to the various problems.

This chapter discusses the hidden Markov model, and outlines standard techniques for parameter identification and state estimation. This is followed by an introduction to recursive prediction error (RPE) and extended Kalman filter (EKF) techniques for parameter identification and estimation. A discussion on communication systems is then presented with emphasis on why hidden Markov models are both applicable and of great benefit. Next, problems of control are discussed and the concept of risk-sensitive control introduced. Finally an outline of the thesis structure is given.

1.1 Hidden Markov Models

An n^{th} -order Markov process is a stochastic process for which the probability distribution of the present state in a sequence is a function of the model parameters and the n previous states, and is independent of all history prior to that. For a hidden Markov model, the Markov process can only be observed via another stochastic process which produces a sequence of *observations*, or outputs resulting from the underlying Markov process. Calculating the expected value of these outputs requires knowledge of only the present Markov state, and is again independent of all previous history.

The Markov processes considered in this thesis are first order, finite-discrete state, discrete-time, homogeneous processes, denoted X_k (where k is the discrete time variable). The phrase *finite-discrete state* implies that the state is restricted to be one of a finite set of discrete values. The term homogeneous means that the parameters are invariant of time (for a thorough treatment of the properties of such processes, see Kemeny and Snell [1960]). For a first order hidden Markov model, the parameters which define the model are: the state values, the transition probabilities, the initial state probability distribution, and the variance of the observation noise.

The main motivation behind considering this particular class of HMMs, is that they can be used to represent many systems arising in a wide range of modern digital technologies, such as computer networks and telecommunication systems. There are, of course, a number of other reasons. One is that the discrete nature of the formulation allows current high speed digital signal processing techniques to be applied, in order to generate practical algorithms. Another is that the computational requirements of higher order Markov models are much greater than for the first order models considered here. And a third reason is that, as is evident from simulation studies, even in the cases where systems are not strictly first order, the first order assumption often still provides a reasonable approximation.

An important aspect of the approach to hidden Markov modelling taken in this thesis, is the method by which internal states are represented. The internal state of the system at any time k , X_k , is represented by an indicator vector from the set of possible orthogonal unit indicator vectors $S = \{e_i\}$, $i = 1, \dots, N$, where N is the total number of states, and $e_i = (0, \dots, 0, 1, 0, \dots, 0)' \in \mathbb{R}^N$ where 1 appears in the i^{th} position. Associated with each indicator vector is a state value. By way of an example, consider a binary digital signal which has voltage levels of ± 1 volt. For this case $S = \{(0\ 1)', (1\ 0)'\}$. When $X_k = (0\ 1)'$ it is indicating that the state value is +1 volt, and when $X_k = (1\ 0)'$ it is indicating that the state value is -1 volt.

The motivation for introducing such notation (termed the indicator vector formulation) is that the model dynamics can be shown to be characterised by the following linear relationship [Segall 1976]:

$$X_{k+1} = \mathbf{A}'X_k + M_{k+1} ,$$

where \mathbf{A} is the transition probability matrix defined by $\mathbf{A} \triangleq (a_{ij})$, for $a_{ij} = P(X_{k+1} = e_j \mid X_k = e_i)$, and M_k is a Martingale increment. The observation process associated with the model

is given by

$$y_k = h(X_k) + w_k ,$$

where w_k is an independent and identically distributed (i.i.d.) noise process and $h(\cdot)$ is the output function. The observation symbol probabilities are defined by $b_k(i) \triangleq P(y_k | X_k = e_i)$ and, in the case where w_k is Gaussian white noise, are distributed by $b_k(i) \sim N[h(e_i), \sigma_w^2]$ where σ_w^2 is the variance of the Gaussian distribution.

The main benefit to this state space representation of the Markov process, is that it allows any nonlinear function of the Markov state, X_k , (in particular, the observations) to be rewritten as a linear function of that state, as follows:

$$y_k = [h(e_1), \dots, h(e_N)]X_k + w_k .$$

While the above indicator vector formulation has been used widely in the statistical literature for analysis of Markov chains, it does not seem to have been applied previously to engineering problems involving *hidden* Markov chains. In this thesis, the indicator vector representation is used in solving problems of signal processing and control for systems modelled by HMMs. By considering such a representation, many important insights can be gained into a wide variety of HMM processing problems. This thesis presents new algorithms which in many cases have distinct advantages over current approaches. The benefits arise to a large extent because the HMM is now in a standard linear systems type representation. This gives new clues relating to the application of finite dimensional filters.

Filtering, or state estimation, of HMMs can be expressed in terms of a conditional expectation of the state, X_k , given the observations, y_k . This leads to the concept of an *information-state* which, as the name suggests, provides information about the state of the system [Kumar and Varaiya 1986] (p. 81). In the case of HMMs, the information-state is a probability distribution vector representing the conditional probability of each state. The most likely state is determined by the largest element of this probability distribution vector. When considering the limiting case, for linear systems in white Gaussian noise, when the set of real numbers are quantised into an infinite number of states, then the information-state is an infinite dimensional vector with a Gaussian distribution. In this case the largest element of the information-state is in fact the minimum variance estimate. In the

more realistic case of finite dimensional systems, the largest element of the information-state is simply termed the maximum *a posteriori* (MAP) estimate.

An important aspect of the work in this thesis is the extensive use of information-states. The application based problems which are considered, are each formulated with the aid of the indicator vector representation in such a way as to enable the full information-state to be used. This is in contrast to more traditional approaches to the problems which just use MAP estimates, and are therefore throwing away information about the quality of that estimate.

This thesis makes use of the new state space HMM representation in conjunction with information state filters, to investigate on-line model identification, adaptive parameter estimation and risk-sensitive control problems. The term *on-line* refers to the case when a new state, or system, estimate is made each time a new measurement is made. The on-line feature is important for applications such as mobile communication systems, and follows in many ways from the state space approach. The techniques presented in this thesis are finite dimensional. They are also new, novel, and in many cases provide distinct advantages over current techniques. The schemes generated are, of course, not specific to communication systems and have a wide range of other applications, such as on-line speech processing and on-line frequency tracking.

1.2 Model Identification

Model identification and estimation are important aspects of signal processing. These two terms apply to the technique of measuring a signal output from a system, and using it to generate a model for that system. Estimation is the term used when the model parameters are time-varying, while identification is applied to the time-invariant parameter case. Both are performed in two stages, firstly the model set must be chosen, and secondly model parameters need to be determined in order to locate a system within the model set which best approximates the true system behaviour. This thesis deals with model identification and adaptive estimation within the class of hidden Markov models. Such tasks are of great interest since they are essential in the solutions to a wide variety of engineering problems, including the communication systems considered later in this thesis.

A popular approach to HMM identification, is the maximum likelihood approach of the expectation maximisation (EM) algorithm. This technique was first developed by Baum and his colleagues

[Baum *et al.* 1970], however, the definitive reference is the paper by Dempster *et al.* [1977]. The EM algorithm is an off-line approach and consists of two steps for each iteration, or *pass* through a batch of data. The first step is an *expectation* step, in which a conditional expectation of some log-likelihood function of the state sequence is generated. The second step is a *maximisation* step with the aim of maximising the conditional expectation, of step one, over all possible models in the model set. The expectation step involves what is commonly known as the *forward-backward algorithm* and the maximisation step makes use of what are commonly called the *Baum-Welch re-estimation equations*. Such an approach provides estimates of the model which have the attractive property that they are guaranteed not to decrease in likelihood with each pass through the data. Unfortunately there are limitations to the EM technique. The algorithm only converges linearly with the number of iterations, and the scheme requires off-line processing of data. The linear characteristic is not a major limitation, but the off-line processing aspect implies that the EM algorithm can not be used for problems where the signal must be analysed in real-time, or on-line, such as for communication systems.

A modification to the EM algorithm for HMMs is presented in Krishnamurthy and Moore [1993] in which on-line processing is achieved. This technique applies a two step procedure, similar to the EM algorithm, as each new piece of data arrives, instead of at the completion of a pass through a batch of data. The techniques are derived using stochastic approximations to maximise the Kullback-Leibler information measure and are therefore termed recursive Kullback-Leibler (RKL) optimisation schemes. This approach is a generalisation, for the Markov case, of algorithms proposed by Titterton [1984] and Weinstein *et al.* [1990], for estimating finite state chains in white Gaussian noise (WGN).

With the application of smoothing and forgetting, the on-line EM techniques can be seen to converge to the correct model in a single *forward* pass through a batch of data. However, a problem arises with these algorithms. The equations presented for updating the estimates of the transition probabilities associated with the Markov process (given by the matrix A from Section 1.1), are not constrained to ensure that these estimates do not become negative, a fundamental requirement for probability measures. This, coupled with the linear convergence properties of the EM technique, suggests the need for further investigation into the on-line HMM identification problem.

Another method for on-line identification of nonlinear systems is the recursive prediction error (RPE) technique presented in Ljung and Söderström [1983]. While it is designed for continuous-

range Gauss-Markov processes, in this thesis it is used in a new approach to the identification of HMMs. The RPE technique is discrete-time, nonlinear and finite-dimensional, which makes it ideal for use with HMMs. As with the EM algorithm, this approach is based on off-line techniques. It is formulated by considering the minimum variance of the prediction error, (that is, the error between the actual measurement and the predicted measurement, based on the best model estimate at the time). The minimisation of this error function provides an updated estimate of the system model with each new measurement. The resulting scheme is expressed as a gradient type algorithm, where the gradient is given by the derivative of the prediction error function. By means of appropriate parameterisation, the scheme is applied in this thesis to HMMs to achieve on-line estimates for the model. An important feature of this new RPE approach to HMMs is the fact that the parameters are considered as states, and take values in a continuous range. This leads to the notion of *mixed-state* HMMs where some states are finite-discrete and some are continuous-range. It is this formulation which allows RPE techniques to be applied.

Convergence analysis of RPE techniques is most fully understood for identification of linear stochastic systems. However, results are also available for stochastic nonlinear parameter estimation problems [Ljung and Söderström 1983, Ljung 1977, Moore and Weiss 1979]. In general, the RPE technique, as with the algorithm in Krishnamurthy and Moore [1993], is not globally converging but rather, converges to the nearest local minimum of the prediction error variance cost function. The RPE scheme is based on a quadratically convergent off-line scheme, therefore improved asymptotic convergence rates are expected for RPE/HMM schemes, over RKL/HMM schemes, but at a computational cost. The computational complexity of the RPE scheme can however be reduced by appropriate de-coupling assumptions, while maintaining the same asymptotic estimates; however initial transients may be degraded.

One must be careful, however, when applying RPE techniques to HMM identification, that successive model estimates preserve the properties of an HMM. In particular, the estimated state transition probabilities (elements of \mathbf{A}) must remain positive and the rows of \mathbf{A} must sum to 1. In this thesis, these requirements are rewritten as smooth equality constraints which define a smooth compact manifold of parameters. Results from differential geometry are used to restrict the gradients, derived for the RPE schemes, to this manifold. In order to ensure the positivity constraint is satisfied, the HMM is parameterised by the square root of the transition probabilities. Thus, the parameter estimates generated can be positive or negative, while the transition probability estimates will always be positive. This implies that the new parameterisation now lives on an N -

dimensional sphere leading to a straightforward application of differential geometry techniques. The approach taken in this thesis is new to recursive identification of HMMs, and has proven to be of great benefit. Simulation studies have shown that in all cases, the application of differential geometry to the RPE schemes, improves convergence and allows for a large increase in initial conditioning errors.

1.3 Parameter Estimation

When considering time-varying systems, many of the assumptions made in generating the identification schemes discussed above, do not apply. The main implication is that convergence results are more difficult to obtain, and sub-optimal algorithms are often required, especially when considering hidden Markov model estimation.

For stochastic linear systems, the optimal state estimator is the Kalman filter (KF), which was presented most completely by Anderson and Moore [1979]. Like the nonlinear RPE scheme, the KF is finite-dimensional. Unlike the RPE scheme, it applies to both time invariant and time varying systems. The KF generates conditional mean estimates and associated covariance estimates for states with linear dynamics driven by white Gaussian noise. For nonlinear systems, the KF can be modified to generate the extended Kalman filter (EKF). At each iteration, the EKF makes linearisations around the conditional mean estimate for the state, and then uses standard KF methods to update the estimate. As with the KF, this technique is finite dimensional.

The KF and EKF approaches are well known and understood. In fact the KF has been used almost exclusively for a wide variety of linear stochastic systems, since its inception in 1960 [Kalman 1960] (see also [Kalman 1963]). When used for parameter estimation, it is known to converge to the true model, even for time-varying systems, as long as certain observability and controllability conditions hold [Anderson and Moore 1979]. For nonlinear systems, convergence of estimates is somewhat more of an issue. In order to ensure that the EKF scheme has bounded errors, more stringent conditions apply [La Scala *et al.* 1993].

For the mixed-state HMM systems considered in this thesis, it is possible to write the observations in a bi-linear form with respect to the state and the parameters. This is due to the indicator vector representation of the Markov process, discussed earlier. The term bi-linear relates to equations which would be linear in one variable if the other was constant, and *vice versa*. Such

a representation allows coupled conditional linear filters to be applied, in conjunction with the nonlinear HMM filter. The term *coupled conditional filter* is given to a filtering scheme where separate filters for each parameter, or state, are run in parallel, and where each filter generates outputs which are conditioned on the estimates given by the other filters. Through this coupled conditional approach, the difficulties of increased computations encountered with the nonlinear EKF scheme are avoided. Unfortunately, while the algorithms generated are more suited to implementation, it has proved difficult to obtain meaningful analytical results describing their asymptotic properties. Therefore, simulation studies are used to demonstrate performance.

1.4 Communication Systems

Communication systems provide a rich area for application of HMM signal processing techniques. In particular, digital systems (for which the discrete-time signal takes on only a finite number of discrete values) fall directly into the framework of the HMM. In the environment of modern computing systems, where on-line HMM processing has become computationally feasible, such applications provide exciting areas for research. This thesis demonstrates some of the advantages that can be gained through the use of HMM filtering in communication systems.

In this thesis hidden Markov models are used to tackle the important problem of tracking fading channels associated with the transmission of digital signals. The term *fading channel* refers to a transmission medium which has continually varying properties. Such channels arise, for example, in mobile communication systems [Vucetic and Du 1992, Loo and Secord 1991] and indoor radio communication systems [Hashemi 1993], where the signal path from transmitter to receiver varies as the receiver changes location or when objects move into the path of the signal. The resulting task is to estimate accurately these channel variations, in order to ensure the correct message is decoded. Much work has been done on signals in Rayleigh fading channels. Some current approaches to the demodulation problem include the use of maximum likelihood (ML) Matched filters [Proakis 1983, Pahlavan and Matthews 1990], and maximum likelihood sequence estimation Viterbi algorithms [Viterbi 1967, Forney Jr. 1973]. Recently, variations to these algorithms have been developed (for example, by Lodge and Moher [1990]) incorporating the ideas of sequence estimation. A summary of results for phase modulated signals can be found in Haeb and Meyr [1989], where optimal solutions are shown to be infinite dimensional, and where sub-optimal schemes are discussed involving decision directed maximum *a posteriori* (MAP) estimates. The

HMM techniques of this thesis can be seen, in many cases, to have advantages over these ML and MAP techniques, particularly with regards to robustness.

At its most abstract level, a communication system, depicted in Figure 1.1, consists of a transmitter, a channel, and a receiver. The sub-components of these three elements vary, depending on the type of messages to be sent. In this thesis digital signals are considered, and the transmitter consists of an encoder and a modulator, while the receiver consists of a demodulator and a decoder. It is the task of the demodulator to estimate the fading channel properties, and therefore, it is in its design that HMM signal processing is applied.

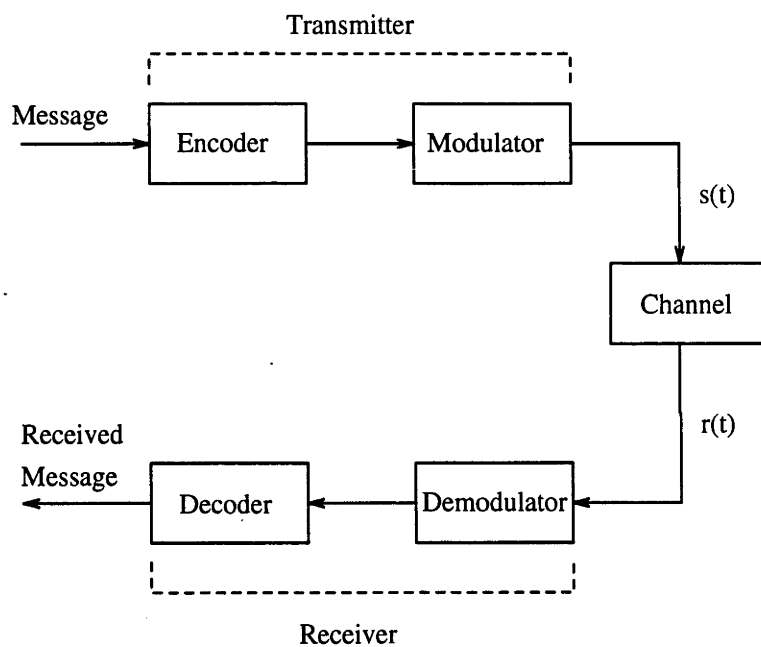


Figure 1.1: Signal transmission system

In order to design the demodulator, it is of course necessary to have an understanding of the transmitter's operation. The transmitter obtains a message signal, then encodes it, modulates it and transmits it. In general, the message could be any signal in a digital form, for example FAX signals or quantised speech signals. The encoder, which processes these messages, has two main purposes: to ensure that transmission through the channel is efficient, and to add redundancy so that errors can be corrected at the receiver. These two tasks are carried out by *source* encoding and *channel* encoding respectively. The source encoder acts to randomise the digital message to ensure that the maximum possible amount of information can be transmitted per bit. If the digital signal has equally probable, or independent and identically distributed (i.i.d.), symbols (for example, in the binary case, the symbols 1 and 0 are equally likely) then the signal cannot

be predicted and therefore all of the signal contains information. If the symbols are not equally probable, then part of the signal is redundant (because it can be predicted). The channel encoder, on the other hand, adds redundancy to the signal, but does so without compromising the efficiency greatly. Redundancy is added so that errors can be more easily detected and corrected at the receiver. Common channel encoders include parity check encoders [Lin 1970], convolutional encoders [Proakis 1983], and more recently, trellis encoders [Biglieri *et al.* 1991].

Once the message is encoded it is then modulated for transmission through the channel. Modulation turns the discrete-time digital signal into an appropriate continuous-time signal. This thesis is primarily interested in applying HMM techniques to two digital transmission schemes, namely quadrature amplitude modulation (QAM) and M-ary differential phase shift keyed (MDPSK) modulation. Both signal classes are commonly used in current communication systems. Other modulation schemes can also be considered in the HMM framework, and although they are not presented in this thesis, they do serve to demonstrate the wide applicability of the mixed-state HMM approach presented here.

Quadrature amplitude modulation is a transmission scheme characterised by a signal which is *quadrature* in nature, and for which the real and imaginary components of the message are transmitted as two *amplitudes* which *modulate* the quadrature and in-phase sinusoidal carriers. The resulting transmitted signal, $s(t)$, has the following form:

$$s(t) = A_c[m^R(t) \cos(2\pi ft + \theta) + m^I(t) \sin(2\pi ft + \theta)] ,$$

where the carrier amplitude A_c , frequency f , and phase θ are constant, and $m(t) = m_k$ for $t = [t_k, t_{k+1})$, where t_k arises from regular sampling, and m_k is the complex value of the message symbol. Therefore, each digital symbol is represented by a different complex number.

In contrast, phase shift keyed modulation schemes carry the message solely in the phase of the signal. They are therefore termed constant amplitude, or constant modulus, transmission schemes. For differential PSK schemes, the message is carried in the difference between successive phases, rather than by the magnitude of the absolute phase value. M-ary DPSK schemes allow M different symbols and thus M different discrete phase jumps. In this case, the transmitted signal has the following form:

$$s(t) = A_c \cos [2\pi f_c t + \theta(t)] ,$$

where $\theta(t) = \theta_k$ for $t = [t_k, t_{k+1})$ and $\theta_k = \sum_{\ell=0}^k f_{\ell}$, where f_k is the frequency shift representing the message symbol, and where the carrier amplitude A_c , and frequency f_c are constant.

The mapping between the output of the source encoder and the sequence of symbols the modulator can transmit, depends of course on the channel encoder. Consider the case of convolutional codes. For such schemes, the code words, corresponding to the modulator symbols, are generated through linear operations on source bits. An example is shown in Figure 1.2 where a binary input signal, from the source encoder, is passed through a four element shift register. In this case, the adders perform modulo two addition. The result is that each possible combination of bits in the register corresponds to one of four output values. Therefore, if we consider a QAM system, a four symbol constellation is required in the modulator, where the complex symbols take values $m_k \in \{\pm c \pm jc\}$, for some constant c .




Figure 1.2: Example of convolutional coding scheme

It is at this point that the usefulness of the HMM formulation becomes apparent. If the combination of bits in the shift register is defined as the state of the message, then each possible combination can be assigned an indicator vector. It is clear that the next state of the shift register depends only on the current state, and the next input bit from the source encoder, and therefore falls directly into the HMM framework. The transition matrix, A , of the resulting Markov process is automatically defined by the probability distribution of the bit sequence, and the possible transitions of state as a new bit enters the register and one drops off the end. In the case of parity check channel encoders, the state of the message can be defined by the modulator symbol (eg. m_k in the case of QAM), and again the HMM formulation applies, as the next message symbol will depend on the current one, and the next input from the source encoder.

With either of these definitions, the HMM demodulator can be designed to adaptively track fading channel characteristics, by applying the mixed-state HMM estimation approach of coupled

conditional filters, described in Section 1.3. Here, an on-line HMM filter is used to estimate the signal state (a discrete quantity) while a Kalman filter is used to track the channel variations (a continuous quantity). The two filters are coupled, as each depends on output from the other. A feature of this technique is that the full information-state (in this case the probability distribution vector which is the output from the HMM filter) is used to condition the Kalman filter. This is in contrast to current techniques which use only the most likely message value for conditioning, and thus do not make use of estimate uncertainty.

Many of the current approaches to demodulation, discussed previously, make use of the above definition of message state, whereby the state represents the contents of the shift register, however they all perform maximum-likelihood based estimation as opposed to, in the HMM case, evaluating the full probability distribution for the expected value of the state. Such ML and MAP techniques are only optimal if the channel is non-fading, and if the input bit sequence from the source encoder is i.i.d. (actually, under such ideal conditions, the HMM filter with MAP estimate (discussed in Section 1.1) is identical to the Matched filter). When the system involves fading channels, and non-i.i.d. input signals, the HMM techniques presented in this thesis can be shown to have distinct advantages, when it comes to signal decoding, channel parameter tracking, and model identification.

1.5 Risk-Sensitive Filtering and Control

An exciting area of current research in stochastic systems, is risk-sensitive filtering [Speyer *et al.* 1992] and control [Jacobson 1973]. Such approaches include an exponential operator in the performance criteria. They lead to an optimal solution, similar to the case for Kalman filtering and linear-quadratic-Gaussian (LQG) control; however in addition, the sensitivity to risk can be varied. One application area for risk-sensitive control has been economics, where risk-sensitivity is termed *hedging* or *risk-aversion*, for example in Karp [1988] and Caravani [1986]. These papers illustrate that advantages can be gained from the risk-sensitive approach, for problems such as dynamic trading and futures market prediction.

For a risk-sensitive policy, decisions (be they control or filtering) are made, based on a desired amount of risk. For example, if a model is known to be accurate, then only a small amount of risk will be involved in decision making. However, if the model is inaccurate then a higher amount of risk would be involved, and so steps should be taken to account for this. Traditional

optimal control techniques employ averaging procedures which assume disturbances are acting in some average manner. More recent robust (for example H_∞) controllers account for maximum disturbances resulting in so called “worst case” designs. The risk-sensitive controller allows for a variation in the robustness of the design, by considering changes in the sensitivity to risk.

As in the HMM application to communication systems discussed in Section 1.4, information-states again play an important role in risk-sensitive techniques. For standard LQG control [Anderson and Moore 1989], a separation principle holds allowing the *state* feedback problem to be solved, and the state estimate, which is calculated separately, is substituted in place of the state. For risk-sensitive problems, no such principle holds and so the full information-state must be fed into the controller. In addition, a key to the technique is that the information-state actually contains information about the cost incurred, as well as the state estimate [Bensoussan and van Schuppen 1985]. This is the vital concept which allows the risk-sensitive performance criteria to be expressed in terms of the information-state, as opposed to the actual state. Therefore, in the control context, output feedback control can be implemented *optimally*, since the information-state, and hence the performance criteria, is a function of the observations, and not the state directly. As is seen here, the concept of an *information-state* covers more than just conditional expectations of the state.

Another important aspect of this work on risk-sensitive systems, is the way in which reference probability measures are defined. For many problems in stochastic systems, advantages can be gained by translating the measure by which information is evaluated, from “real-world” measures, to ones which are more attractive mathematically. The changes of probability measure used in this thesis are discrete-time versions of Girsanov’s theorem [Segall 1976]. The technique involves defining a new reference probability measure for which the observations are i.i.d. random variables. It can be shown that under such a measure change, nonlinear information-states can be re-expressed in terms of *new* un-normalised information-states, which can in turn be evaluated by linear equations [Elliott 1993].

By way of example, consider that $E[\cdot]$ is the expectation operator under the usual reference probability measure, and that $\bar{E}[\cdot]$ is the expectation operator under the new measure where the observations are independent. Consider the following:

$$E[x] = \int x dP = \int x \frac{dP}{d\bar{P}} d\bar{P} = \bar{E}[\Lambda x],$$

where $\Lambda \triangleq dP/d\bar{P}$ is the change of probability measure. Therefore, expectations under one

measure can be expressed in terms of expectations under another measure, with the inclusion of the change of probability terms. For filtering problems involving conditional probabilities, one has the following version of Bayes' theorem:

$$E[x|y] = \frac{\overline{E}[\Lambda x|y]}{\overline{E}[\Lambda]}.$$

In this thesis similar techniques to those in Elliott [1993] are used to compute $\overline{E}[\Lambda x|y]$ using linear equations and the quantity of real interest, namely $E[x|y]$, is generated by normalising this *new* information-state.

Such measure changes can be used for both linear and nonlinear systems analysis. This thesis uses techniques presented in James *et al.* [1994], to derive finite-dimensional output-feedback risk-sensitive linear-quadratic-Gaussian (RLQG) control results. Solutions for this linear case are obtained in a non-separated form, and thus provide insight to the infinite-dimensional nonlinear solution of James *et al.* [1994]. The resulting equations are compared to those in Whittle [1981] where linear techniques were used to solve the regulator problem. They are also shown to specialise to standard LQG results under limiting conditions.

In addition to regulation, this thesis considers applying a risk-sensitive policy to tracking problems. Augmentations to the plant model are presented, in order to achieve zero steady state error in the constant reference signal case. For such problems, integrators are included in the controller design. One advantage to this approach is that it leads to an increase in the allowable range over which the sensitivity to risk can be varied. Unfortunately, however, adding an integrator greatly reduces the benefit gained by such variations.

In addition to the RLQG solution, the techniques are applied to finite-discrete HMM tracking problems, where a finite dimensional information-state is derived. The control solution in this HMM case, is however infinite dimensional, and so sub-optimal solutions are discussed. In each of the tracking cases, simulation examples are presented which demonstrate the effect of varying the controller's sensitivity to risk.

Finally, the risk-sensitive filtering problem is shown to be solved by considering the tracking problem, and reformulating the performance criteria. For tracking problems, the task is to choose a control which causes the state to be close to a desired value. It is possible to reformulate this into a filtering problem, for which the task is to pick a value close to the true state. Using such a

technique, this thesis derives the risk-sensitive filtering equations directly from the information-state equations obtained for the tracking problem. Such a connection leads to the possibility of considering more difficult, and as yet unsolved, problems of coupled risk-sensitive filtering and risk-sensitive control.

1.6 Algorithms Developed

This section presents a list of the algorithms developed in this thesis:

- **On-Line Identification of HMMs :** New recursive prediction error (RPE) based algorithms are developed to perform on-line identification of HMMs. The schemes perform better than previous algorithms in terms of asymptotic convergence, as well as having a number of other attractive properties, for example they automatically ensure that transition probability estimates are positive.
- **Reduced Complexity On-Line Identification of HMMs :** A reduction in computational complexity can be achieved in cases where states are clustered in groups. An algorithm is presented, which is again based on the RPE algorithm, for on-line identification in such cases.
- **Adaptive Demodulation of QAM signals :** New HMM based algorithms are developed for demodulation of digital QAM signals transmitted through flat fading channels. The schemes are especially applicable for mobile communication systems where fast fading environments exist. The schemes are on-line, adaptive, and make use of new information-state ideas. Simulation studies are carried out which demonstrate the advantages of such an approach when compared current schemes.
- **Adaptive Demodulation of MDPSK signals :** Information-state HMM techniques are used to demodulate digital M-ary DPSK signals in fading channels. The resulting algorithms are shown also to be applicable to analogue FM signals with appropriate quantisation. As in the QAM case, the schemes are on-line and adaptive.
- **Risk-Sensitive Tracking for Linear Systems :** An algorithm is presented for tracking with output feedback and a risk-sensitive control policy. State space augmentations are also discussed in order to incorporate integrators for the purpose of obtaining zero steady

state error. While a similar tracking result has been derived previously, the technique by which the solution is generated in this thesis provides a much simplified framework for implementation, as well as giving more insight into nonlinear systems.

- **Risk-Sensitive Tracking for HMMs** : The risk-sensitive approach is applied to tracking for hidden Markov model systems, to allow for variations in the robustness of HMM controllers. Numerical solutions to the dynamic programming control problem are also discussed.

1.7 Thesis Structure

A brief outline of the progression of ideas in this thesis is as follows :

Hidden Markov Model Identification

In Chapter 2 an on-line state and parameter identification scheme for HMMs, with states in a finite-discrete set, is developed using RPE techniques. The parameters of interest are the transition probabilities and discrete state values of a Markov chain, and the noise density associated with the observations. In general any independent and identically distributed (i.i.d.) noise density could be considered, however, in this work, zero mean white Gaussian noise (WGN) is used. In contrast to the more familiar off-line EM algorithm which is a fixed-interval “forward-backward” multi-pass approach, the RPE on-line schemes presented here, have significantly reduced memory requirements. In addition, a reduced computational complexity algorithm is developed for the case where state values are clustered into groups.

Implementation aspects of the proposed algorithms are discussed, and simulation studies on three and four state Markov chains are presented to show that the algorithms provide competitive, and in fact asymptotically faster converging, estimates to those of earlier proposed on-line HMM schemes based on recursive Kullback-Leibler (RKL) measure optimisation. Also, an improved version of the RKL scheme is proposed with a parameterisation that ensures positivity of transition probability estimates.

Adaptive Parameter Estimation with HMMs

In Chapter 4 the techniques of extended Kalman filtering and HMM signal processing are combined to adaptively demodulate QAM signals in noisy fading channels. The approach is to formulate the QAM signal model and channel parameter model, in a mixed finite-discrete state and continuous

state framework. The *mixed state model* is then re-formulated in terms of conditional information-states, using HMM theory. This leads to models which are amenable to standard EKF or related techniques. Adaptive state and parameter estimation schemes are devised based on the assumption that the transmission channel introduces time-varying gain and phase changes, modelled by a stochastic linear system, and has additive Gaussian noise. The case of white observation noise is considered, as well as generalisations to cope with coloured noise. The adaptive HMM approach results in a practical finite-dimensional algorithm for state and parameter estimation, consisting of a coupled continuous state KF and finite state HMM filter. A more sophisticated EKF scheme with an HMM sub filter is also discussed. Simulation studies demonstrate the ability to estimate the signal, and track effectively the rapidly time-varying channel parameters, in noisy conditions. Comparisons with traditional techniques are also presented. These serve to demonstrate the advantages of the mixed state HMM approach.

In Chapter 5 EKF and HMM signal processing, techniques are again blended, this time in order to demodulate MDPSK signals in noisy fading channels. Two adaptive HMM approaches are formulated, both consist of a continuous state KF coupled with finite-discrete state HMM filters. The first proves computationally intensive, the second incorporates decoupling which achieves practical finite-dimensional algorithms. The technique used is to represent MDPSK signals with state space signal models where some of the state components belong to a finite-discrete set and others are in a continuous range. As with the QAM case, HMM signal processing is then applied to re-formulate these models as conditional information-state models from which the KF/HMM coupled filters are derived. Simulation studies demonstrate the ability to estimate the MDPSK signal, and track time-varying channel parameters.

Risk-Sensitive Control

In Chapter 6 the information-state approach is used to obtain solutions to risk-sensitive quadratic control problems. Specifically, the case of tracking a desired trajectory, is considered. Results are presented for linear discrete-time models with Gaussian noise, and also for finite-discrete state, discrete-time hidden Markov models with continuous-range observations. Using the information-state approach, the tracking solution is obtained without appealing to a certainty equivalence principle. Limit results are discussed which demonstrate the link to standard linear quadratic Gaussian control. Also included is a section on achieving zero steady state error with risk-sensitive control policies. Simulation studies are presented to show some advantages gained via the use of a

risk-sensitive approach. Finally, a discussion on risk-sensitive filtering is presented. The filtering problem is shown to be solved by considering the control problem with a slightly different cost function. This result leads to the notion of dual risk-sensitive filtering and risk-sensitive control.

Chapter 2

Identification of Hidden Markov Models

2.1 Introduction

This chapter considers the problem of on-line model identification for finite-discrete state, discrete-time hidden Markov models (HMMs) with continuous-range observations. As mentioned in Chapter 1, HMMs with states in a finite-discrete set have been widely applied in areas such as speech processing [Rabiner 1989], biological signal processing [Chung *et al.* 1991] and pattern recognition [Bellegarda *et al.* 1994]. In such cases, it is of primary importance to identify the parameters of the Markov model, so as to gain a better understanding of the system concerned. The Expectation Maximisation (EM) algorithm [Dempster *et al.* 1977, Baum *et al.* 1970] is a popular off-line technique for obtaining maximum likelihood estimates of the HMM parameters. However, a limitation of the off-line EM methods is the ‘curse of dimensionality’ which arises because the computational effort, speed, and memory requirements are in proportion to the square of the number of possible states of the Markov chain. Memory requirements are also proportional to the length of data being processed. As a consequence of these limitations, there is incentive to explore on-line (sequential) algorithms to seek improvements in terms of memory requirements and computational speed, and also to cope with the possibility of time-varying HMM parameters.

The key contribution of this chapter is to reformulate HMMs, with states in a finite-discrete set and with unknown parameters, in such a way that Recursive Prediction Error (RPE) techniques can be applied for model identification. The quadratically convergent RPE scheme is shown to be a superior alternative to the linearly convergent recursive Kullback-Leibler (RKL) scheme presented in Krishnamurthy and Moore [1993].

This chapter is mainly concerned with estimating a finite state Markov chain in white Gaussian noise (WGN). The parameters to be estimated are the transition probabilities and state values of the Markov chain, and the noise variance associated with the measurements. To derive the RPE based algorithms, it is necessary to first formulate the HMM in an N -dimensional state space form (where N is the number of Markov states). In terms of traditional linear system terminology, the $N \times N$ transition probability matrix corresponds to the system matrix, while the vector of Markov state values forms the observation matrix. The next step is to formulate an alternative representation of the finite-discrete state HMM, based on the conditional filtered state estimate of the Markov chain. In this case, the filtered estimate, termed information state, lies in a continuous range, and is given by the “forward” variable, α_k , of the EM algorithm [Rabiner 1989]. This nonlinear estimator based model is characterised in terms of the HMM parameters.

A requirement of RPE formulation presented here, is that, in addition to computing α_k , it is necessary to evaluate the first derivatives of α_k with respect to the HMM parameters. Recursive techniques, similar to the forward recursion for α_k , are proposed for computing these derivatives. The computational complexity required for these recursions, is the same as that for α_k (that is, $O(N^2)$ per time instant).

In this chapter, two model parameterisations are considered for handling transition probabilities, constrained by the Markov property that rows of the system matrix must add to one. The first is with transition probabilities constrained to a simplex, the second is with the square root of transition probabilities constrained to the surface of a sphere in \mathbb{R}^N . In both cases the derivatives are restricted to the tangent space of the simplex or manifold. It is shown that an advantage is gained by working on the sphere, instead of the simplex, because the surface of the sphere is smooth, and the estimates of transition probabilities are assured to be non-negative. Other parameterisations are also possible (for example a polar co-ordinate representation on the sphere) however they are less attractive since most result in much greater complexity in the calculation of gradients for the RPE scheme. This chapter also presents an RKL scheme on the sphere which

eliminates the inherent possibility, in Krishnamurthy and Moore [1993], of negative transition probability estimates.

For an additional improvement to the estimates, it is also possible to modify the RPE formulation to incorporate fixed-lag smoothing. Simulations have shown that using fixed lag variable estimates, such as γ_k (the smoothed state estimate of a Markov chain [Rabiner 1989]), and first derivatives, can improve the transient properties of the estimates. Of course fixed-lag smoothing adds computational requirements proportional to the lag.

In the actual implementation of the HMM/RPE algorithm it is necessary to scale or normalise the filtered variables and their derivatives to keep their values in the numerical range of a computer. This chapter employs a scaling procedure which scales without affecting the parameter estimates.

Simulation examples are presented to illustrate the algorithms. These examples show that the proposed schemes can satisfactorily identify HMM parameters and that their performance is better, asymptotically, than that of RKL schemes.

This chapter is organised as follows: Section 2.2 formulates the HMM and details the estimation objectives. It also provides an outline of RPE techniques and shows their relevance to HMM estimation. In Section 2.4, RPE based recursive algorithms for HMM estimation are derived, working on both the simplex and the sphere for the constrained transition probability estimates. Also, an RKL scheme on the spherical parameter constraint manifold is proposed. In Section 2.5 implementation considerations are discussed, and simulation examples are given. Finally, conclusions are presented in Section 2.6.

2.2 Problem Formulation

This section presents the HMM signal model, details estimation objectives, reviews the RPE algorithm, and points out its relevance in adaptive HMM estimation.

2.2.1 State Space Signal Model

Let X_k be a discrete-time homogeneous, first order Markov process belonging to a finite-discrete set. The state space of X , without loss of generality, can be identified with the set of unit vectors

$\mathbf{S} = \{e_1, e_2, \dots, e_N\}$, $e_i = (0, \dots, 0, 1, 0, \dots, 0)' \in \mathbb{R}^N$ with 1 in the i^{th} position. The transition probability matrix is

$$\mathbf{A} = (a_{ij}) \quad , \quad 1 \leq i, j \leq N \quad , \quad \text{where } a_{ij} = P(X_{k+1} = e_j \mid X_k = e_i) \quad ,$$

so that

$$E[X_{k+1} \mid X_k] = \mathbf{A}' X_k \quad ,$$

where $E[\cdot]$ denotes the expectation operator. Of course $a_{ij} \geq 0$, $\sum_{j=1}^N a_{ij} = 1$, for each i . It is also of use to denote $\{\mathcal{F}_l, l \in \mathbb{Z}^+\}$ to be the complete filtration generated by X . That is, for any $k \in \mathbb{Z}^+$, \mathcal{F}_k is the complete σ -field generated by $X_k, l \leq k$.

Definition 2.1: [Adams and Guillemin 1986] Consider events A_i , in a sample space Ω . Now, a σ -field, \mathcal{F} , is defined to be a collection of subsets of Ω which contains the empty set, and is closed under complements and countable unions. That is, given $\{A_i\}_{i=1}^{\infty}$ in \mathcal{F} , then $\bigcup_{i=1}^{\infty} A_i$ is also in \mathcal{F} .

□

Lemma 2.1 *The dynamics of X_k are given by the state equation*

$$X_{k+1} = \mathbf{A}' X_k + M_{k+1} \quad , \quad (2.1)$$

where M_{k+1} is a $(\mathbf{A}, \mathcal{F}_k)$ martingale increment, in that $E[M_{k+1} \mid \mathcal{F}_k] = 0$.

Proof : [Segall 1976]

$$E[M_{k+1} \mid \mathcal{F}_k] = E[X_{k+1} - \mathbf{A}' X_k \mid X_k, \mathbf{A}'] = E[X_{k+1} \mid X_k, \mathbf{A}'] - \mathbf{A}' X_k = 0 \quad .$$

■

Assume that X_k is hidden, that is, indirectly observed by measurements y_k . The *observation process* y_k has the form

$$y_k = g(X_k) + w_k \quad , \quad (2.2)$$

where without loss of generality, since X_k is in a discrete set, $g(X_k) = \langle g, X_k \rangle$, with $\langle \cdot, \cdot \rangle$ denoting the inner product in \mathbb{R}^N , and $g \in \mathbb{R}^N$ is the vector of state values of the Markov chain. It is also of use to define $Y_k \triangleq (y_0 \dots y_k)$. Assume that w_k is independent and identically distributed (i.i.d.), with a zero mean Gaussian distribution, so that $w_k \sim N[0, \sigma_w^2]$, and $E[w_{k+1} \mid \mathcal{F}_k \vee \mathcal{Y}_k] = 0$ where

\mathcal{Y}_l is the σ -field generated by y_k , $k \leq l$. It is now readily seen that $E[M_{k+1} | \mathcal{F}_k \vee \mathcal{Y}_k] = 0$. Because w_k is white, the independence property

$$E[y_k | X_{k-1} = e_i, \mathcal{F}_{k-2}, \mathcal{Y}_{k-1}] = E[y_k | X_{k-1} = e_i], \quad (2.3)$$

holds and is essential for formulating the problem as an HMM.

Now, let $\mathbf{b}_k = (b_k(1), \dots, b_k(N))'$ denote the vector of parameterised probability densities (which will loosely be called symbol probabilities), where $b_k(i) \triangleq b(y_k, g_i) = P[y_k | X_k = e_i, \theta]$, and

$$b(y_k, g_i) = \frac{1}{\sqrt{2\pi\sigma_w^2}} \exp\left(-\frac{(y_k - g_i)^2}{2\sigma_w^2}\right). \quad (2.4)$$

Also, the assumption is made that the initial state probability vector for the Markov chain $\underline{\pi} = (\pi_i)$ is defined from $\pi_i = P(X_1 = e_i)$. The HMM is denoted $\lambda = (\mathbf{A}, g, \underline{\pi}, \sigma_w^2)$.

Remark 2.1: Figure 2.1 gives an example of a realization of the observation process, y_k , for a 3 state HMM. The state values are $g = [0, 1, 2]'$. □




Figure 2.1: Example of observation process (3 state HMM)

2.2.2 Model Parameterisation

Consider that λ is parameterised by an unknown vector θ so that $\lambda(\theta) = (\mathbf{A}(\theta), g(\theta), \underline{\pi}, \sigma_w^2(\theta))$. Two parameterisations are proposed, the dimensions of which are $N_\theta = N + N^2 + 1$, representing N state values, N^2 transition probabilities, and the noise variance associated with the observations.

1. Parameterisation on a Simplex :

$$\theta = (g_1, \dots, g_N, a_{11}, \dots, a_{1N}, a_{21}, \dots, a_{NN}, \sigma_w^2)' , \quad (2.5)$$

subject to the equality and inequality constraints defined by the simplex

$$\Delta^N := \left\{ a_{ij} : \sum_{j=1}^N a_{ij} = 1 \text{ and } a_{ij} \geq 0 \right\} . \quad (2.6)$$

2. Parameterisation on a Sphere :

$$\theta = (g_1, \dots, g_N, s_{11}, \dots, s_{1N}, s_{21}, \dots, s_{NN}, \sigma_w^2)' , \quad (2.7)$$

where $a_{ij} = s_{ij}^2$. The benefit of the second parameterisation (2.7), over the first (2.5), is that the constraint manifold is differentiable at all points. In addition, (2.7) has only the equality constraint of the sphere surface, \mathbf{S}^{N-1} , where

$$\mathbf{S}^{N-1} := \left\{ s_{ij} : \sum_{j=1}^N s_{ij}^2 = 1 \right\} , \quad (2.8)$$

and therefore ensures each transition probability estimate is positive.

For either parameterisation the following state space signal model applies:

$$\boxed{\begin{aligned} X_{k+1} &= \mathbf{A}'(\theta) X_k + M_{k+1} \\ y_k &= \langle g(\theta), X_k \rangle + w_k \end{aligned}} \quad (2.9)$$

This signal model is not, however, in a form suitable for application of RPE techniques to achieve estimates of θ and X_k from the measurements y_k (due to the finite-discrete nature of the states, X_k). Such a model is now developed.

2.2.3 Parameterised Information-State Signal Model

Let $\hat{X}_{k|\theta}$ denote the conditional filtered state estimate of X_k at time k , that is,

$$\hat{X}_{k|\theta} \triangleq E[X_k | \mathcal{Y}_k, \theta]. \quad (2.10)$$

Also, define $\underline{1}$ to be the column vector containing all ones, and the *information-state* (termed “forward” variable in Rabiner [1989]) $\alpha_{k|\theta} \triangleq (\alpha_{k|\theta}(1), \dots, \alpha_{k|\theta}(N))'$, is defined by

$$\alpha_{k|\theta}(i) \triangleq P(Y_k, X_k = e_i | \theta). \quad (2.11)$$

Lemma 2.2 $\hat{X}_{k|\theta}$, as defined in (2.10), can be expressed in terms of the column vector $\alpha_{k|\theta}$, defined in (2.11), as follows:

$$\hat{X}_{k|\theta} = \langle \alpha_{k|\theta}, \underline{1} \rangle^{-1} \alpha_{k|\theta}. \quad (2.12)$$

Proof : From (2.10),

$$\begin{aligned} \hat{X}_{k|\theta} = E[X_k | \mathcal{Y}_k, \theta] &= \sum_{i=1}^N e_i P(X_k = e_i | Y_k, \theta) \\ &= \begin{pmatrix} P(X_k = e_1 | Y_k, \theta) \\ \vdots \\ P(X_k = e_N | Y_k, \theta) \end{pmatrix} = \frac{1}{P(Y_k | \theta)} \begin{pmatrix} P(Y_k, X_k = e_1 | \theta) \\ \vdots \\ P(Y_k, X_k = e_N | \theta) \end{pmatrix} \\ &= \langle \alpha_{k|\theta}, \underline{1} \rangle^{-1} \alpha_{k|\theta}. \end{aligned}$$

■

Lemma 2.3 The information-state $\alpha_{k|\theta}$ can be computed using the following “forward” recursion [Rabiner 1989]:

$$\alpha_{k+1|\theta} = \mathbf{B}(y_{k+1}, \theta) \mathbf{A}'(\theta) \alpha_{k|\theta}, \quad (2.13)$$

where $\mathbf{B}(y_k, \theta) = \text{diag}(b(y_k, g_1), \dots, b(y_k, g_N))$.

Proof :

$$\begin{aligned}
\alpha_{k+1|\theta}(j) &= P(Y_{k+1}, X_{k+1} = e_j | \theta) \\
&= P(y_{k+1}, Y_k, X_{k+1} = e_j | \theta) \\
&= P(y_{k+1} | Y_k, X_{k+1} = e_j, \theta) P(Y_k, X_{k+1} = e_j | \theta) \\
&= P(y_{k+1} | X_{k+1} = e_j, \theta) \sum_{i=1}^N P(Y_k, X_{k+1} = e_j, X_k = e_i | \theta) \\
&= b(y_{k+1}, g_j) \sum_{i=1}^N P(X_{k+1} = e_j | Y_k, X_k = e_i, \theta) P(Y_k, X_k = e_i | \theta) \\
&= b(y_{k+1}, g_j) \sum_{i=1}^N P(X_{k+1} = e_j | X_k = e_i, \theta) P(Y_k, X_k = e_i | \theta) \\
&= b(y_{k+1}, g_j) \sum_{i=1}^N a_{ij} \alpha_{k|\theta}(i)
\end{aligned}$$

Writing this in vector notation gives the result. ■

It is now necessary to express the observations, y_k , in terms of the un-normalised conditional estimates, $\alpha_{k|\theta}$.

Lemma 2.4 *The conditional measurements $y_{k|\theta}$ are given by*

$$y_{k|\theta} = \langle g(\theta), \langle \alpha_{k-1|\theta}, \mathbf{1} \rangle^{-1} \mathbf{A}' \alpha_{k-1|\theta} \rangle + n_k ,$$

where α_k is defined in (2.11) and n_k is a $(\theta, \mathcal{Y}_{k-1})$ martingale increment.

Proof : Following standard arguments since $\alpha_{k|\theta}$ is measurable with respect to $\{\theta, \mathcal{Y}_k\}$, $E[w_{k+1} | \mathcal{Y}_k] = 0$ and $E[M_{k+1} | \mathcal{Y}_k] = 0$, then

$$\begin{aligned}
E[n_{k+1} | \theta, \mathcal{Y}_k] &= E[\langle g(\theta), X_{k+1} \rangle + w_{k+1} - \langle g(\theta), \langle \alpha_{k|\theta}, \mathbf{1} \rangle^{-1} \mathbf{A}' \alpha_{k|\theta} \rangle | \theta, \mathcal{Y}_k] \\
&= \langle g(\theta), \mathbf{A}' \hat{X}_{k|\theta} \rangle - \langle g(\theta), \langle \alpha_{k|\theta}, \mathbf{1} \rangle^{-1} \mathbf{A}' \alpha_{k|\theta} \rangle = 0 .
\end{aligned}$$

■

In summary, the parameterised estimator based signal model for an HMM parameter θ , and with states $\alpha_{k|\theta}$, is given by

$$\begin{aligned}
\alpha_{k+1|\theta} &= \mathbf{B}(y_{k+1}, \theta) \mathbf{A}'(\theta) \alpha_{k|\theta} \\
y_{k|\theta} &= \langle g(\theta), \langle \alpha_{k-1|\theta}, \mathbf{1} \rangle^{-1} \mathbf{A}' \alpha_{k-1|\theta} \rangle + n_k
\end{aligned}$$

(2.14)

where n_k is a $(\theta, \mathcal{Y}_{k-1})$ martingale increment. This signal model is now in a suitable form to apply RPE algorithms to achieve, simultaneously, state and parameter estimates, on-line. Of course, an alternative signal model in terms of normalised estimates could be formulated as:

$$\hat{X}_{k+1|\theta} = \frac{\mathbf{B}(y_{k+1}, \theta) \mathbf{A}'(\theta) \hat{X}_{k|\theta}}{\langle \mathbf{1}, \mathbf{B}(y_{k+1}, \theta) \mathbf{A}'(\theta) \hat{X}_{k|\theta} \rangle}, \quad (2.15a)$$

$$y_{k|\theta} = \langle g(\theta), \mathbf{A}'(\theta) \hat{X}_{k-1|\theta} \rangle + n_k. \quad (2.15b)$$

However, this model turns out to be less convenient for deriving an RPE scheme.

2.3 Identification Algorithms

2.3.1 General Prediction Error Algorithms

The off-line minimum variance prediction error task is to evaluate the following [Ljung and Söderström 1983]:

$$\min_{\theta} V(\theta), \text{ where } V(\theta) = \frac{1}{2} E[y_k - \hat{y}_{k|\theta, k-1}]^2, \quad (2.16)$$

and $\hat{y}_{k|\theta, k-1} \triangleq E[y_k | \theta, \mathcal{Y}_{k-1}]$. This process yields a set of global minima $\{\theta^*\}$, where $\theta^* = \operatorname{argmin}_{\theta} V(\theta)$. It is known that if the signal generating system is in the model set, then under appropriate observability, (identifiability), conditions the true system parameter $\theta^o \in \{\theta^*\}$. Moreover, if in addition, the signal model is uniquely parameterised, then $\{\theta^*\}$ contains one element, namely θ^o . For a finite state sequence of length T , and under ergodicity assumptions, where by time averages are equivalent to expectations, the task (2.16) is approximated as

$$\min_{\theta} V_T(\theta), \text{ where } V_T(\theta) = \frac{1}{2T} \sum_{k=1}^T (y_k - \hat{y}_{k|\theta, k-1})^2, \quad (2.17)$$

to yield estimates $\{\hat{\theta}_T\}$. As discussed in Ljung and Söderström [1983], under reasonable regularity conditions, $\{\hat{\theta}_T\}$ converges with probability 1, to a minimum of $\bar{V}(\theta) = \lim_{T \rightarrow \infty} E[V_T(\theta)]$, as T tends to infinity. In addition, for appropriate initial conditions, the estimates are consistent, in that $\lim_{T \rightarrow \infty} \{\hat{\theta}_T\} \equiv \{\theta^*\}$. Moreover, if $d^2 \bar{V}(\theta)/d\theta^2$ is invertible then the asymptotic rate of convergence is proportional to \sqrt{T} .

2.3.2 Off-line Gauss-Newton Prediction Error Method

The Gauss-Newton Prediction Error method [Ljung and Söderström 1983] is a gradient search approach to finding a local minima of (2.16). At the end of each successive pass, p , through the data, used to evaluate $V_T(\theta)$ and its derivatives, the parameter estimate, $\hat{\theta}$, is updated by

$$\hat{\theta}_{p+1} = \hat{\theta}_p + \gamma \left(\frac{d^2V(\theta)}{d\theta^2} \Big|_{\theta=\hat{\theta}_p} \right)^{-1} \frac{dV(\theta)}{d\theta} \Big|_{\theta=\hat{\theta}_p}. \quad (2.18)$$

For suitable γ ($\gamma \rightarrow 1$ as $p \rightarrow \infty$), $V(\hat{\theta}_{p+1}) \leq V(\hat{\theta}_p)$ and a local minimum is found. Under ergodicity, the algorithm (2.18) is approximated by the practical algorithm

$$\hat{\theta}_{p+1} = \hat{\theta}_p + \gamma \left(\frac{1}{T} \sum_{t=1}^T \psi_{\theta_p} \psi'_{\theta_p} \right)^{-1} \left(\frac{1}{T} \sum_{t=1}^T \psi_{\theta_p} \epsilon_{\theta_p} \right), \quad (2.19)$$

where $\psi_{\theta} = dV_T(\theta)/d\theta$ and $\epsilon_{\theta_p} = y_k - \hat{y}_{k|\theta_p, k-1}$. In the case of θ parameterised on a sphere, as in (2.7) and (2.8), the gradient vector, ψ , is evaluated in the tangent space of the sphere.

The computational cost of the off-line scheme (2.19) is of order TN_{θ}^2 for each pass through a T point observation sequence. Also it is necessary to invert the $N_{\theta} \times N_{\theta}$ matrix, $(d^2V(\theta)/d\theta^2)$, at the end of each pass.

2.3.3 The Recursive Prediction Error Algorithm

In this section a recursive version of the prediction scheme is applied to the signal model (2.14). In the RPE case, at each time update k , there is no attempt to calculate $V(\hat{\theta})$ or $V_T(\hat{\theta})$ exactly, but rather by approximations. Thus the estimate for α_k is recursively computed at each iteration, using obvious notation, as follows:

$$\hat{\alpha}_{k+1|\hat{\Theta}_k} = \mathbf{B}(y_{k+1}, \hat{\theta}_k) \mathbf{A}'(\hat{\theta}_k) \hat{\alpha}_{k|\hat{\Theta}_{k-1}}, \quad (2.20)$$

where $\hat{\theta}_k$ is the recursive estimate of the parameter vector of unknown coefficients based on \mathcal{Y}_k and $\hat{\Theta}_k \triangleq \{\hat{\theta}_1, \dots, \hat{\theta}_k\}$. Let $\hat{y}_{k|\hat{\Theta}_{k-2}, \hat{\theta}_{k-1}}$ denote a predicted output at time k based on measurements up to time $k-1$.

$$\hat{y}_{k|\hat{\Theta}_{k-2}, \hat{\theta}_{k-1}} = \langle g(\hat{\theta}_{k-1}), \langle \hat{\alpha}_{k-1|\hat{\Theta}_{k-2}}, \mathbf{1} \rangle^{-1} \mathbf{A}'(\hat{\theta}_{k-1}) \hat{\alpha}_{k-1|\hat{\Theta}_{k-2}} \rangle. \quad (2.21)$$

Now, the RPE parameter update equation is [Ljung and Söderström 1983] (p. 94):

$$\hat{\theta}_{k+1} = \Gamma_{proj} \{ \hat{\theta}_k + \gamma_{k+1} R_{k+1}^{-1} \psi_{k+1} \hat{n}_{k+1|\hat{\theta}_{k-1}, \hat{\theta}_k} \} \quad (2.22)$$

where

$$\hat{n}_{k|\hat{\theta}_{k-2}, \hat{\theta}_{k-1}} = y_k - \hat{y}_{k|\hat{\theta}_{k-2}, \hat{\theta}_{k-1}}, \quad (2.23)$$

$$R_{k+1} = R_k + \gamma_{k+1} (\psi_{k+1} \psi_{k+1}' - R_k), \quad (2.24)$$

or by the matrix inversion lemma (see Equation (B.1)),

$$R_{k+1}^{-1} = \frac{1}{1 - \gamma_{k+1}} \left(R_k^{-1} - \frac{\gamma_{k+1} R_k^{-1} \psi_{k+1} \psi_{k+1}' R_k^{-1}}{(1 - \gamma_{k+1}) + \gamma_{k+1} \psi_{k+1}' R_k^{-1} \psi_{k+1}} \right). \quad (2.25)$$

Here γ_k is a gain sequence (often referred to as step size) satisfying,

$$\gamma_k \geq 0, \quad \sum_{k=1}^{\infty} \gamma_k = \infty, \quad \sum_{k=1}^{\infty} \gamma_k^2 < \infty. \quad (2.26)$$

Also ψ_k is the gradient

$$\psi_k \triangleq \left(-\frac{d\hat{n}_{k|\hat{\theta}_{k-2}, \theta}}{d\theta} \right) \Big|_{\theta=\hat{\theta}_{k-1}}, \quad (2.27)$$

where $\hat{n}_{k|\hat{\theta}_{k-2}, \theta} = y_k - \hat{y}_{k|\hat{\theta}_{k-2}, \theta}$ and

$$\hat{y}_{k|\hat{\theta}_{k-2}, \theta} = \langle g, \langle \hat{\alpha}_{k-1|\hat{\theta}_{k-2}, \theta}, \mathbf{1} \rangle^{-1} \mathbf{A}' \hat{\alpha}_{k-1|\hat{\theta}_{k-2}, \theta} \rangle. \quad (2.28)$$

R_k is the Hessian, or covariance matrix, approximation. The notation $\Gamma_{proj}\{\cdot\}$ represents a projection into the so called stability (in this case constraint) domain. This is discussed later in Section 2.4.

Implementing (2.24) is computationally expensive, since a matrix inversion is required. It is therefore more attractive to implement (2.25), thereby reducing computation requirements to the order N_θ^2 . To further reduce the computational cost of the RPE scheme it is possible to consider that, as k increases, R_k becomes diagonally dominated. A scheme of complexity order N_θ can be derived by assuming that the covariance matrix is always diagonal. The resulting update equation for each of the diagonal elements of R_k is given by

$$r_{k+1}(i) = r_k(i) + \gamma_{k+1} (\psi_{k+1}^2(i) - r_k(i)) \quad 1 \leq i \leq N_\theta, \quad (2.29)$$

where $R_k = \text{diag}(r_k(1), \dots, r_k(N_\theta))$. Simulations have shown that this technique does not appear to reduce the asymptotic estimate of the RPE scheme, however initial transients may be degraded. Also, in comparison to the full RPE scheme, it is generally necessary to choose initial conditions which are closer to the true values.

Convergence Properties : The RPE algorithm has the following convergence properties under reasonable ergodicity assumptions: The estimate $\hat{\theta}_k$ converges with probability one, to a local minimum of $V(\theta)$ as k approaches infinity, where $V(\theta)$ is defined in (2.16), (see convergence analysis in Ljung and Söderström [1983], Dupuis and Kushner [1989], Kushner and Shwartz [1984]). Unfortunately however, a rigorous verification of the ergodicity type conditions is not yet complete for the HMM case.

While it is not possible to make claims of global convergence, it should be noted that in the simulation studies of Section 2.5.2, the only times convergence was to a point other than the global minimum of $V(\theta)$, was when all state level estimates of the HMM converged to the same value, or when the observation noise variance was so small that the covariance matrix, R_k , was out of the range of the computer, for initial estimates far from the true value. These are of course easily detected situations. However, in cases where initial estimates were far from the true parameters, the evolution of the estimates was sometimes too erratic or too slow to be useful.

The RPE technique is based on a Newton off-line re-estimation approach which is known to be quadratically convergent. In contrast, the RKL technique [Krishnamurthy and Moore 1993] is based on the off-line EM algorithm which is known to be linearly convergent. While the theory has not yet been fully developed for convergence rates of either recursive scheme, at least it is possible to say that from simulations, after initial transients, the RPE scheme has an error function, $1/k \sum_{t=1}^k |\hat{\theta} - \theta|$, which appears proportional to $1/\sqrt{k}$, compared to the RKL scheme which appears proportional to $1/\sqrt[4]{k}$, (for noise variance 1). On the other hand, the RKL scheme has a computational effort, for parameter updates, of order N_θ (where N_θ is defined in Section 2.2.2), whereas the RPE scheme has effort of order N_θ^2 . The RPE scheme can, however, often be implemented using decoupling assumptions to gain reduced computational complexity of order N_θ without adversely effecting the asymptotic estimates.

It is also worth noting that under the usual ergodicity assumptions, the convergence result holds, for the RPE scheme, whether or not the true system belongs to the model set, $\{\hat{\theta}\}$. This means that even if the system is of higher dimension than parameterised, (2.22) will choose an appropriate

best approximation in the sense of minimising $V(\theta)$. This is a valuable robustness property, and implies the RPE technique is not dependent on a precise set of model assumptions.

A final point to mention is that prediction error algorithms, such as RPE, are based on the assumption that observations consist of two parts: a predictable part and an unpredictable Gaussian distributed random part. The implication of this for Markov chains (which jump randomly between states in a finite-discrete set as opposed to the random-walk nature of linear systems in Gaussian noise) is that such algorithms may not be applicable in the low noise case. Under such conditions the random jumps from state to state will be much larger than the deviations due to the Gaussian observation noise, and hence can not be approximated by a Gaussian assumption. Fortunately in low noise conditions, with the aid of smoothing, the random jumps can be predicted and this problem overcome, although the techniques are somewhat ad-hoc (see Example 3.5).

Increased Step Size and Averaging : Equations (2.22) and (2.24) show how the gain sequence, γ_k , scales the update of both R_k and $\hat{\theta}_k$. Apart from satisfying (2.26) it can be any function. Generally, it has the form $\gamma_k = \gamma_0/k^n$, $n \in \mathbb{R}$. This form is used under the assumption that as k increases, the estimates improve and need to be updated less. In the derivation of (2.24), $\gamma_k = 1/k$ is assumed. In practise, for this case, γ_k tends to become too small too quickly, and does not allow fast convergence for initial estimates chosen far from the minimum error point. To overcome this problem, a method suggested by Polyak and Juditsky [1992] is used whereby a larger step size is applied (that is $0 \leq n \leq 1$), and then the estimates are averaged. Averaging is used to get a smoother estimate, as the larger step will mean higher sensitivity to noise, and also to ensure that the third requirement in (2.26) remains satisfied. In the simulations $n = 0.5$ was chosen.

A further heuristic technique is to start averaging only after the estimates are in a close range of the true value, as indicated by a small error value, $V(\theta)$. However, while this may improve the transient response, it will have no effect on the asymptotic properties.

2.4 Gradient Vector and Projection Calculations

In this section gradient and projection calculations for the RPE based algorithms (2.22)-(2.27) are presented, to estimate the HMM parameters, θ . These algorithms update $\hat{\theta}_k$ and R_k recursively, and require $\hat{y}_k|\hat{\theta}_{k-1}$ and ψ_k to be evaluated at each update. In Section 2.4.1 ψ_k is evaluated for the parameterisation on a simplex, given by (2.5). In Section 2.4.2 the parameterisation on a sphere

(2.7), is employed. The required projections, Γ_{proj} , are also given in each case. In Section 2.4.3 the RKL scheme is presented with the parameterisation on a sphere, as opposed to the simplex as is given in Krishnamurthy and Moore [1993], again to ensure the transition probability estimates remain non-negative.

2.4.1 Parameterisation on a Simplex

The derivative vector, ψ_k , defined in (2.27), is given, for $m, n \in [1, \dots, N]$, in the case of the first parameterisation (2.5), by

$$\psi_k = \left. \frac{\partial \hat{y}_k | \hat{\Theta}_{k-2, \theta}}{\partial \theta} \right|_{\theta = \hat{\theta}_{k-1}} = \left(\left. \frac{\partial \hat{y}_k | \hat{\Theta}_{k-2, \theta}}{\partial g_m}, \frac{\partial \hat{y}_k | \hat{\Theta}_{k-2, \theta}}{\partial a_{mn}}, \frac{\partial \hat{y}_k | \hat{\Theta}_{k-2, \theta}}{\partial \sigma_w^2} \right)' \right) \Big|_{\theta = \hat{\theta}_{k-1}}, \quad (2.30)$$

where $\hat{y}_k | \hat{\Theta}_{k-2, \theta}$ is given in (2.28). In the remainder of this section the dependence of $\hat{\alpha}_k$ on $\hat{\Theta}_{k-1}$ is omitted and $N_k \triangleq \langle \hat{\alpha}_k, \mathbf{1} \rangle^{-1}$.

The derivatives with respect to the discrete-state values, g_i , are obtained by differentiating (2.28) to yield

$$\boxed{\frac{\partial \hat{y}_{k+1} | \hat{\Theta}_{k-1, \theta}}{\partial g_m} = N_k \langle a_{(\cdot)m}, \hat{\alpha}_k \rangle + N_k \langle g, \mathbf{A}' \eta_k(m) \rangle - N_k^2 \langle \mathbf{1}, \eta_k(m) \rangle \langle g, \mathbf{A}' \hat{\alpha}_k \rangle} \quad (2.31)$$

where $a_{(\cdot)m} = (a_{1m}, \dots, a_{Nm})'$, and $\eta_k(m) = \partial \hat{\alpha}_k / \partial g_m$ is the N -dimensional vector given by the following recursive equation:

$$\eta_{k+1}(m) = \mathbf{B}(y_{k+1}, \theta) \mathbf{A}' \eta_k(m) + \text{diag}(e_m) \left(\frac{y_{k+1} - g_m}{\sigma_w^2} \right) \mathbf{B}(y_{k+1}, \theta) \mathbf{A}' \hat{\alpha}_k. \quad (2.32)$$

The derivatives with respect to transition probabilities, $0 < a_{ij} < 1$, are given by

$$\boxed{\frac{\partial \hat{y}_{k+1} | \hat{\Theta}_{k-1, \theta}}{\partial a_{mn}} = N_k g_n \hat{\alpha}_k(m) + N_k \langle g, \mathbf{A}' \xi_k(m, n) \rangle - N_k^2 \langle \mathbf{1}, \xi_k(m, n) \rangle \langle g, \mathbf{A}' \hat{\alpha}_k \rangle} \quad (2.33)$$

where $\xi_k(m, n) = \partial \hat{\alpha}_k / \partial a_{mn}$ is the N -dimensional vector given by the following recursive equation:

$$\xi_{k+1}(m, n) = \mathbf{B}(y_{k+1}, \theta) \mathbf{A}' \xi_k(m, n) + \hat{\alpha}_k(m) b(y_{k+1}, g_n) e_n. \quad (2.34)$$

The derivatives (2.33) and (2.34) are not defined along the “edges” of the simplex, that is for $a_{ij} = 0$ or 1, so that in practise, a projection from the simplex boundary to a neighbourhood inside the simplex, is required.

The derivative with respect to the measurement noise variance, σ_w^2 , is given by

$$\boxed{\frac{\partial \hat{y}_{k+1} | \hat{\theta}_{k-1, \theta}}{\partial \sigma_w^2} = N_k \langle g, \mathbf{A}' \rho_k \rangle - N_k^2 \langle \mathbf{1}, \rho_k \rangle \langle g, \mathbf{A}' \hat{\alpha}_k \rangle} \quad (2.35)$$

where $\rho_k = \partial \hat{\alpha}_k / \partial \sigma_w^2$ is the N -dimensional vector given by the following recursive equation:

$$\rho_{k+1} = \mathbf{B}(y_{k+1}, \theta) \mathbf{A}' \rho_k + \delta(y_{k+1}) \mathbf{B}(y_{k+1}, \theta) \mathbf{A}' \hat{\alpha}_k, \quad (2.36)$$

where

$$\delta(y_{k+1}) = \text{diag} \left(\frac{(y_{k+1} - g_1)^2}{2 \sigma_w^4}, \dots, \frac{(y_{k+1} - g_N)^2}{2 \sigma_w^4} \right) - \text{diag} \left(\frac{1}{2 \sigma_w^2} \mathbf{1} \right).$$

In the case of this first parameterisation, projection onto $\sum_{j=1}^N \hat{a}_{ij} = 1$ is straightforward:

$$\hat{a}_{ij}(k+1) = \left(\sum_{j=1}^N \hat{a}_{ij}^U(k+1) \right)^{-1} \hat{a}_{ij}^U(k+1), \quad (2.37)$$

where $\hat{a}_{ij}^U(k)$ is the un-normalised RPE estimate.

There is, however, no guarantee that $\hat{a}_{ij} \geq 0$, either before, or after, projection. One approach to achieve this end, is to reduce the step size at any iteration which violates the condition. Notice too that the projection in this case, using the first parameterisation (2.5), is not an orthogonal projection. It does however have the property that projection of $\hat{a}_{ij}^U \geq 0$ onto $\sum_{j=1}^N \hat{a}_{ij} = 1$ achieves $\hat{a}_{ij} \geq 0$, in contrast to the orthogonal projections which may violate $\hat{a}_{ij} \geq 0$ after projection.

The computational requirements can be further reduced by estimating only $N - 1$ transition probabilities for each state value, and using the constraint $\sum_{j=1}^N a_{ij} = 1$ to determine the remaining value. Actually it is a little disconcerting that, for such a scheme, there is a free choice of which a_{ij} 's to leave out of the parameter vector, θ . However, experience has shown that it is better to leave out the smallest element in each row of \mathbf{A} .

2.4.2 Parameterisation on a Sphere

In the case of the second parameterisation (2.7) with s_{ij} rather than its square a_{ij} , then

$$\psi_k = \left. \frac{\partial \hat{y}_k |_{\hat{\theta}_{k-2}, \theta}}{\partial \theta} \right|_{\theta = \hat{\theta}_{k-1}} = \left(\left. \frac{\partial \hat{y}_k |_{\hat{\theta}_{k-2}, \theta}}{\partial g_m}, \frac{\partial \hat{y}_k |_{\hat{\theta}_{k-2}, \theta}}{\partial s_{mn}}, \frac{\partial \hat{y}_k |_{\hat{\theta}_{k-2}, \theta}}{\partial \sigma_w^2} \right)' \right|_{\theta = \hat{\theta}_{k-1}}, \quad (2.38)$$

and again (2.31)(2.32)(2.35) and (2.36) apply, however (2.33) and (2.34) are formulated with a modification to the derivative calculation. Given that the parameterisation implies s_{ij} lives on the surface of a unit sphere in \mathbb{R}^N , more appropriate expressions for (2.33) and (2.34) can be gained by projecting gradients onto the tangent space of the surface of this sphere. The projected gradient can be written as follows, (see Appendix A):

$$\boxed{\frac{\partial \hat{y}_{k+1} |_{\hat{\theta}_{k-1}, \theta}}{\partial s_{mn}} = 2N_k \hat{\alpha}_k(m) s_{mn} g' \left(e_n - \text{diag}(s_{m(\cdot)}) s'_{m(\cdot)} \right) + N_k \langle g, \mathbf{A}' \xi_k(m, n) \rangle - N_k^2 \langle \mathbf{1}, \xi_k(m, n) \rangle \langle g, \mathbf{A}' \hat{\alpha}_k \rangle} \quad (2.39)$$

where $s_{m(\cdot)} = (s_{m1}, \dots, s_{mN})$ and $\xi_k(m, n) = \partial \hat{\alpha}_k / \partial s_{mn}$ is the N -dimensional vector given by the following recursive equation

$$\xi_{k+1}(m, n) = \mathbf{B}(y_{k+1}, \theta) \mathbf{A}' \xi_k(m, n) + 2\hat{\alpha}_k(m) \mathbf{B}(y_{k+1}, \theta) s_{mn} \left(e_n - \text{diag}(s_{m(\cdot)}) s'_{m(\cdot)} \right) \quad (2.40)$$


In achieving an update estimate of s_{ij} at time $k+1$ via (2.38), denoted here $\hat{s}_{ij}(k+1)$, there is a required projection $\Gamma_{proj}\{\cdot\}$ into the constraint domain (2.8), the surface of a unit sphere in \mathbb{R}^N . Observe that this domain is also a stability domain. Thus in updating \hat{s}_{ij} , first an unconstrained update, denoted \hat{s}_{ij}^U , is derived then projected onto the sphere as follows

$$\hat{s}_{ij}^2(k+1) = \frac{\left(\hat{s}_{ij}^U(k+1) \right)^2}{\sum_{j=1}^N \left(\hat{s}_{ij}^U(k+1) \right)^2}, \quad (2.41)$$

to achieve $\sum_{j=1}^N \hat{s}_{ij}^2(k+1) = 1$ as required.

Recall, as described in Section 2.2.1, with the parameterisation on a sphere, the constraint manifold is differentiable at all points and the positivity of transition probability estimates is inherently guaranteed. Further, the projection in (2.41), is orthogonal. Simulation studies show that the parameterisation on a sphere yields much improved estimates compared to the parameterisation on a simplex.

Remark 2.2: Figure 2.2 shows the HMM/RPE scheme in block diagram form. It indicates the way in which the standard HMM filter and RPE identification algorithm have been coupled. \square



Where $f_1(\cdot)$ is given implicitly in (2.32)(2.34) (2.36) and (2.40) and $f_2(\cdot)$ is given implicitly in (2.31)(2.33) (2.35) and (2.39)

Figure 2.2: HMM/RPE scheme for HMM identification

2.4.3 RKL Scheme with Parameterisation on a Sphere

The RKL scheme presented by Krishnamurthy and Moore [1993] with parameterisation on a simplex, can now be modified for the parameterisation (2.7). It is of benefit to do this because even though the RKL update equations for a_{ij} on a simplex, have the property that the incremental terms, when summed over j , add to zero, there is no constraint to ensure $a_{ij} \geq 0$. Employing (2.7), and the approach given in Appendix A for gradients on a manifold, leads to the following modifications to the score vector, S_A , and the Fisher information matrix (FIM), I^A , as defined in

Krishnamurthy and Moore [1993].

$$s_A(i, j) = g_j^{(i)} = \frac{2\zeta_{k+1|k+1}(i, j)}{s_{ij}} - 2\gamma_{k+1|k+1}(i)s_{ij}, \quad (2.42a)$$

$$p_A(i, j) = \mu_j^{(i)} = \sum_{t=1}^{k+1} \left(\frac{2\zeta_{t|k+1}(i, j)}{s_{ij}^2} + 2\gamma_{t|k+1}(i) \right), \quad (2.42b)$$

where the score vector for the transition probabilities is now given by the vector $S_A = (S_A(1), \dots, S_A(N))'$ for $S_A(i) = (s_A(i, 1), \dots, s_A(i, N))$, and the FIM is given by $I^A = \text{blockdiag}(P_1, \dots, P_N)$ where now, $P_i = \text{diag}(p_A(i, 1), \dots, p_A(i, N))$. Also, $\gamma_{k|k}(i) = P(X_k = e_i | Y_k, \theta)$ is the conditional probability of the state at time k , and $\zeta_{k|k}(i, j) = P(X_k = e_i, X_{k+1} = e_j | Y_k, \theta)$ is the conditional probability of transition from time k to $k+1$. In contrast to the RKL scheme on a simplex, where I^A is block diagonal, here it is easily seen that I^A is strictly diagonal. Consequently, the computational complexity of this RKL scheme is still of order N_θ , being that for RKL on a simplex.

The update for the transition probabilities on a sphere, is now given by

$$s_{ij}(k+1) = s_{ij}(k) + \frac{g_j^{(i)}}{\mu_j^{(i)}}, \quad (2.43)$$

where $g_j^{(i)}$ and $\mu_j^{(i)}$ are given in (2.42a) and (2.42b) respectively. As with the RPE parameterisation on a sphere, the projection (2.41) is required.

2.5 Implementation Considerations and Simulations

Simulations of the RPE scheme, as applied to the HMM formulation presented in this chapter, have been carried out on three-state and four-state Markov chains. This section presents some practical implementation aspects which should be considered in simulations, and also the results of the simulations.

2.5.1 Implementation Considerations

This section considers important implementation aspects such as scaling, multi-pass processing, introduction of fixed-lag smoothing, stability, and re-initialisation to cope with slowly time varying models.

Scaling : From (2.13) (2.32)(2.34) and (2.36) it is noted that as k increases, $\alpha_k(i)$ and derivatives of $\alpha_k(i)$ decrease exponentially. It is necessary to scale estimates of these variables to ensure they remain within the numerical range of the computer. In order to develop a consistent scaling strategy the scaling factor must cancel out in the final update equations, (2.22) and (2.24), to ensure it has no effect on the estimates.

The following strategy for scaling is proposed, based on techniques in Rabiner [1989]. Let α_k be the actual un-scaled forward variable defined in (2.13), let $\tilde{\alpha}_k$ be the un-scaled updated version based on the previous scaled version, and let $\bar{\alpha}_k$ be the scaled updated version based on the previous scaled version, that is:

$$\bar{\alpha}_k(i) = C_k \tilde{\alpha}_k(i), \quad (2.44)$$

where

$$C_k \triangleq \frac{1}{\sum_{i=1}^N \tilde{\alpha}_k(i)}. \quad (2.45)$$

It follows from (2.44) that

$$\bar{\alpha}_k(i) = C_T \alpha_k(i) \text{ with } C_T = (C_k C_{k-1} \dots C_0), \quad (2.46)$$

and for the derivative terms (η, ξ, ρ) similar expressions can be derived. For example (2.32)(in scalar notation) gives

$$\tilde{\eta}_{k+1}(j) = \sum_{i=1}^N \tilde{\eta}_k(i) a_{ij} b(y_{k+1}, g_j) + \sum_{i=1}^N \tilde{\alpha}_k(i) a_{ij} \frac{y_{k+1} - g_j}{\sigma_w^2} b(y_{k+1}, g_j), \quad (2.47)$$

which leads to

$$\bar{\eta}_k(i) = C_T \eta_k(i). \quad (2.48)$$

Similarly, $\bar{\xi}_k(i) = C_T \xi_k(i)$ and $\bar{\rho}_k(i) = C_T \rho_k(i)$. From (2.46) and (2.48) it can be shown, by direct substitution into (2.32), that evaluating derivatives of ψ_{k+1} with respect to g_i , using $\bar{\alpha}_k$ and $\bar{\eta}_k$, yields identical results to the case where no scaling has taken place. Similarly it can be shown that scaling does not affect the estimates for the transition probabilities and noise variance.

Multi-pass Processing : Multiple passes through the data can be used to improve the estimates. One method is to divide the data up into sections. The algorithm is applied to each section a number of times before the next section of data is used. The initial estimate at each successive pass is given

by the output from the previous pass and is therefore an improved initial estimate. Asymptotically this will have no effect, however initially it may be a useful technique for improving transients.

Fixed-Lag Smoothing Algorithms : A fixed-lag smoothed estimate of the state can be computed by replacing α_k with α_k^Δ , for which the i^{th} element is given by $\alpha_k^\Delta(i) \triangleq P(Y_{k+\Delta}, X_k | \theta)$. It is shown in Krishnamurthy and Moore [1993] that α_k^Δ can be computed as

$$\alpha_k^\Delta = \alpha_k(i) \beta_{k|k+\Delta}(i), \quad (2.49)$$

where Δ is the fixed lag and $\beta_{k|k+\Delta}$ is given by

$$\beta_{k+\Delta|k+\Delta}(i) = 1,$$

$$\beta_{t|k+\Delta}(i) = \sum_{j=1}^N \beta_{t+1|k+\Delta}(j) a_{ij} b(y_{t+1}, g_j), \text{ for } t = k + \Delta - 1, k + \Delta - 2, \dots, k, \quad (2.50)$$

for $\beta_{k|k+\Delta} \triangleq P(y_k, y_{k+1}, \dots, y_{k+\Delta} | X_k, \theta)$.

Of course a small memory requirement is imposed because $\Delta + 1$ data points must be stored, namely $(y_k, y_{k+1}, \dots, y_{k+\Delta})$.

The justification for this substitution is based on the observation [Krishnamurthy and Moore 1993] that smoothing improves estimates by giving a better approximation of the Markov chain's current state. In simulation studies it was found that, in many cases, parameter estimates were better and convergence was improved, when (2.49) was implemented.

Initialisation and Projection : Initialisation needs to be addressed to insure all assumptions are satisfied. The main requirements are that: initial parameter estimates are close to the actual values and satisfy the model constraints, initial gradients are normal to the surface of the sphere in \mathbb{R}^N (for (2.39) and (2.40)), and the covariance matrix is initialised in such a way that the parameter error remains bounded.

The second two requirements are satisfied in the scheme presented here, by not updating the parameter estimates in the first instance, until reasonable estimates of gradients and covariance matrix are obtained. It was found that in the order of N^2 data points were required to obtain such a state. In place of this method, a projection can be applied in the first instance to ensure estimates

do not change substantially from the initialised values. These schemes can be implemented to ensure $\hat{\theta}_{k+1}$ is in a small neighbourhood of $\hat{\theta}_k$, as is required for RPE.

From simulation studies it was found that two other factors effected the performance of this HMM/RPE scheme. They are the choice of the initial step size and the covariance matrix. Fortunately, however, the 'best' choice of γ_0 and R_0 was observed to depend only on the observation noise variance, which is often assumed to be known. If other initialisations were used, the estimates still converged to the true values, although it generally took many more iterations. Therefore, in cases where the observation noise variance is known, the transient performance of the scheme can be improved by appropriate selection of γ_0 and R_0 .

Re-initialisation Of Variables : As noted previously, the choice of γ_k is based on the assumption that as k increases, the estimates are better. In the case of time-varying HMMs, γ_k may become too small to allow for large jumps in signal parameters. One way of tracking time-varying HMMs is to periodically reinitialise γ_k , especially when large step changes are involved. For the same reason, it is also useful to reinitialise R_k and ψ_k . The frequency of re-initialisation is chosen as a trade off between good steady state estimation and time-varying tracking ability. Techniques for tracking quickly time-varying parameters are discussed in Chapter 4.

2.5.2 Simulation Studies

Presented here are results of simulation studies using computer generated finite state Markov Chains in WGN.

Example 2.1: A three state Markov Chain embedded in WGN was generated with parameter values $a_{ii} = 0.9$, $a_{ij} = (1 - a_{ii})/(N - 1)$ for $i \neq j$, $g = [0, 1, 2]'$, $\sigma_w = 1$. Figure 2.3 and Figure 2.4 show parameter estimates for this data. The RPE algorithm was used with parameterisation (2.7), the projection scheme (2.41), and the covariance matrix updated by (2.25). The algorithm was initialised with parameter estimates, $\hat{a}_{ii} = 0.1$, $\hat{g} = [0.5, 0.6, 0.7]'$. This scheme gives a complete update by accounting for all the cross coupling between estimates, rather than implementing decoupling of the covariance matrix as discussed in Section 2.3.3, equation (2.29), however there is little difference in the estimates generated for the two cases, as shown later in Example 2.5. In addition, for the covariance matrix recursion, $\gamma_k = 1/k$, as is strictly the case from its derivation. However, for the parameter update, $\gamma_k = 1/\sqrt{k}$, and averaging is implemented to achieve faster convergence, as discussed in Section 2.3.3.

Notice that after only 10000 points, the estimates for both the state values and the transition probabilities are close to the true values, even though they were initialised with a large error. By way of comparison, the RKL scheme presented in Krishnamurthy and Moore [1993], was applied to the same data. The estimates shown, using the RPE scheme (Figure 2.3 and Figure 2.4), converge to the neighbourhood of the correct values, in a comparable time to the RKL scheme, however, asymptotic convergence is faster.

Example 2.2: A four state Markov Chain embedded in WGN was generated with parameter values $a_{ii} = 0.9$, $a_{ij} = (1 - a_{ii})/(N - 1)$ for $i \neq j$, $g = [0, 1, 2, 3]'$, $\sigma_w = 1$. Figure 2.5 and Figure 2.6 show the parameter estimates for this data, where the computational effort is of course greater than that for Example 2.1. Initial estimates are, again, able to be taken far from the true values, however more iterations than for the three state case, are required before the estimates are in the region of the correct values (10^5 points compared to 10^4). These estimates were generated with the same scheme as was used in Example 1, with initial estimates $\hat{a}_{ii} = 0.1$, $\hat{g} = [0.5, 1, 1.5, 2]'$. HMMs with up to six states have been tested, and no limit to the number of states is anticipated.

Example 2.3: For a comparison of performance under the high noise case, a three state Markov Chain embedded in WGN was generated with parameter values $a_{ii} = 0.9$, $a_{ij} = (1 - a_{ii})/(N - 1)$ for $i \neq j$, $g = [0, 1, 2]'$, $\sigma_w = 4$. Figure 2.7 and Figure 2.8 give parameter estimates for this data using the same scheme as in Example 2.1, and initialised with $\hat{a}_{ii} = 0.1$, $\hat{g} = [0.5, 1.4, 1.6]'$. In this case the parameter estimates take longer to reach the neighbourhood of the correct values (2×10^5), and there is greater sensitivity to initial estimates. A longer averaging length would be useful in order to gain smoother estimates for this high noise case.

Example 2.4: Figure 2.9 and Figure 2.10 show the estimates for a three state Markov Chain under the assumption that six states are present. The data used is the same as for Example 2.1. It can be seen that even if the algorithm is computed for a higher order model, the estimates will still move to the correct values. It should be noted that the transition probability estimates have estimated the values for a six state system with $g = [0, 0, 1, 1, 2, 2]'$. If model reduction is applied to the estimated values, then the true three state transition probability matrix is generated.

Example 2.5: This example demonstrates the feasibility of using the reduced complexity RPE algorithm derived by assuming the covariance matrix is diagonally dominated, as discussed in Section 2.3.3. The parameters of the HMM are $a_{ii} = 0.7$, $a_{ij} = (1 - a_{ii})/(N - 1)$ for $i \neq j$, $g = [0, 1, 2]'$, $\sigma_w = .5$. The initial conditions for the algorithms are $\hat{a}_{ii} = 0.1$, $\hat{g} = [0.5, 1.4, 1.6]'$.

Figures 2.11 and 2.12 show parameter estimates for this data, generated using the full RPE algorithm as used in the previous examples. Figures 2.13 and 2.14 show parameter estimates generated using a strictly diagonal covariance matrix. The same data sequence was used in both cases.

Comparing the results from the two schemes shows that there is very little difference in the estimates. It can be seen, however, from other simulation studies, that the range of initial conditions which result in convergence, is smaller for the reduced complexity scheme. Also, the transient response of the estimates is generally worse. These are the trade-offs which need to be made when considering such reduced complexity algorithms.

Empirical comparison of the convergence rates of RPE and RKL : When focussing on the asymptotic convergence rates, as discussed in Section 2.4, simulations show that the RKL scheme asymptotically converges at a rate which is the square root of that for the RPE scheme. Figure 2.15 and Figure 2.16 show the error function defined in Section 2.3.3, plotted against time, k , for the RPE and RKL schemes respectively. The data used is the same as that used in Example 2.1. Although it is not shown, the RPE scheme with covariance matrix decoupling, displays the same asymptotic convergence as the full RPE scheme.

Parameter Estimation Error Studies : Presented in Tables 2.1 to 2.6 are results of simulations carried out using two-state Markov chains in WGN. Each table is generated from 50 simulations and the error function used (The parameter estimate error (PEE)) is given by $PEE(\hat{x}) = \sqrt{(1/50) \sum_{i=1}^{50} (x - \hat{x})^2}$. A two-state process was used in order to give meaning to the noise variance, when presented in the form of a signal to noise ratio (SNR). The parameters of the Markov chain are $g = [0, 1]'$ and $a_{ii} = 0.9$. The SNR is therefore given by $10 \log(1/\sigma_w^2)$. Initial parameter estimates used in generating Tables 2.1 to 2.4 were $g = [0.4, 0.6]'$ and $a_{ii} = 0.5$. The tables show that the HMM/RPE algorithms converge to the correct values, even for high noise, and also for a wide range of initial conditions. In some cases the state value estimates collapsed to a single state, when this occurred the algorithm was re-initialised.

2.6 Conclusions

In this chapter an RPE based on-line parameter and state identification scheme for HMMs has been derived. The algorithms presented are memory efficient and of computational complexity $O(N_\theta^2)$ for the RPE estimate and $O(N^2)$ for the HMM information-state evaluation. Two parameterisations of the HMM have been considered, one on a simplex, the other on a sphere. Simulation studies have shown that advantages are gained by working on the sphere and projecting gradients into the tangent space. An existing RKL scheme has also been reformulated with the parameterisation on a sphere. In addition, implementation aspects for the RPE scheme have been discussed, and the ability to estimate parameters of both three and four state Markov chains in WGN has been demonstrated. Simulation studies have also shown that, in many cases, decoupling assumptions can be made which reduce the complexity of the RPE scheme to $O(N_\theta)$.

2.7 Figures and Tables

The figures and tables for this chapter are now presented.




Figure 2.3: Level estimates of 3 state Markov chain

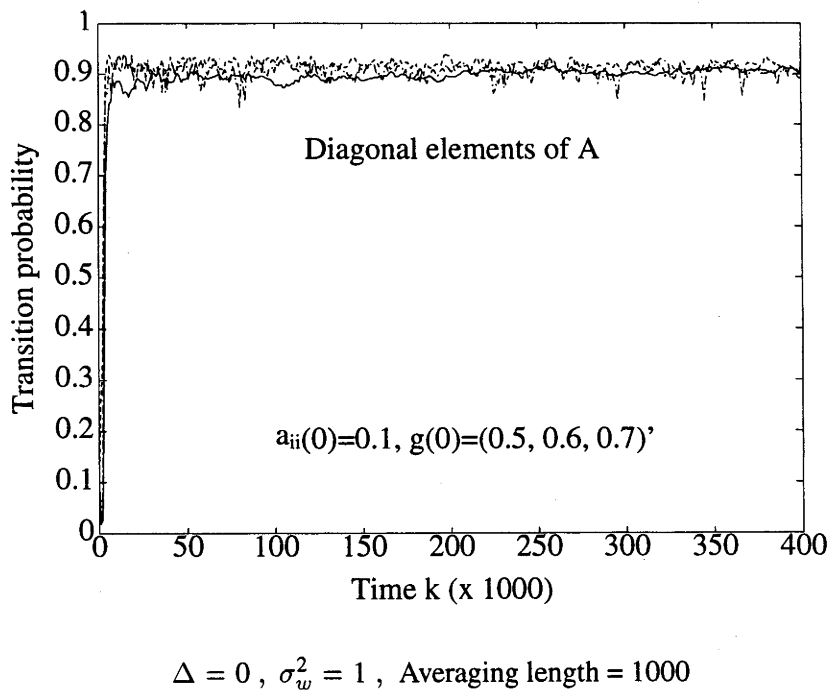

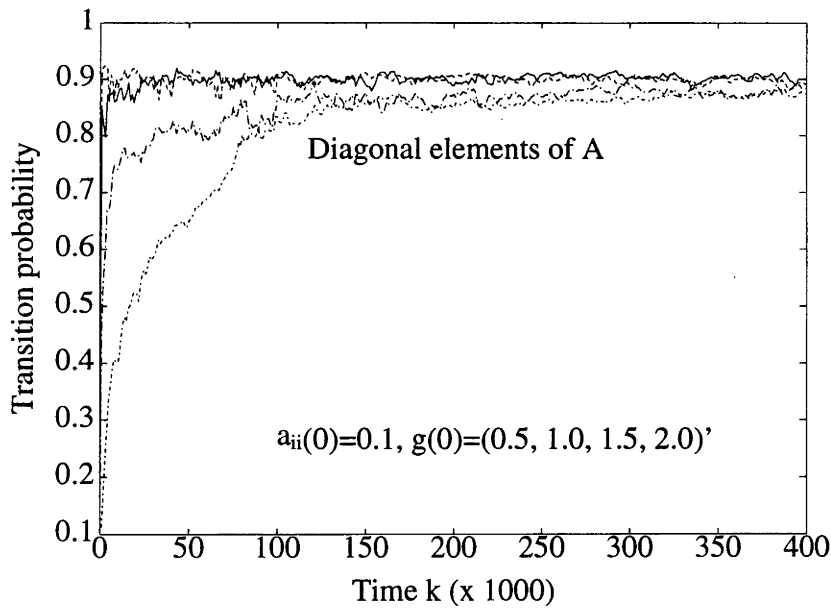


Figure 2.4: Transition probability estimates of 3 state Markov chain



$\Delta = 0, \sigma_w^2 = 1, \text{ Averaging length} = 1000$

Figure 2.5: Level estimates of 4 state Markov chain



$\Delta = 0, \sigma_w^2 = 1, \text{ Averaging length} = 1000$

Figure 2.6: Transition probability estimates of 4 state Markov chain




Figure 2.7: Level estimates of 3 state Markov chain in high noise case

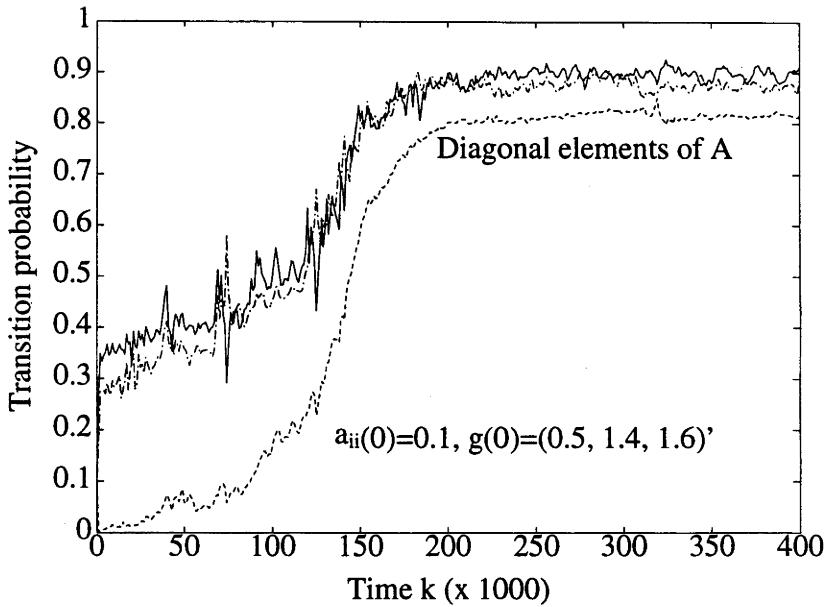


Figure 2.8: Transition probability estimates of 3 state Markov chain in high noise case




Figure 2.9: Level estimates of 3 state Markov chain under 6 state assumption

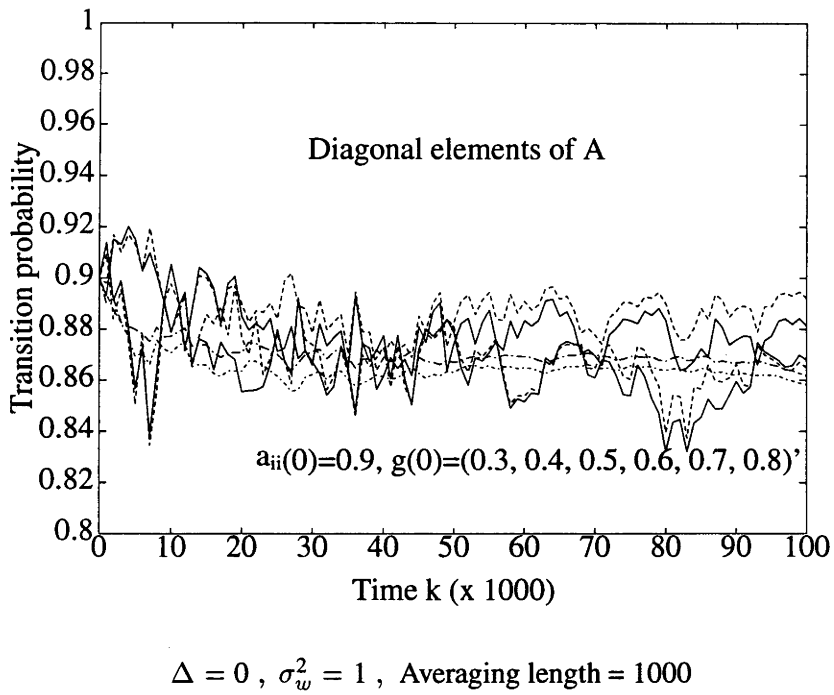




Figure 2.10: Transition probability estimates of 3 state chain under 6 state assumption



$\Delta = 0, \sigma_w = .5, \text{Averaging length} = 1000$

Figure 2.11: Level estimates 3 state HMM



$\Delta = 0, \sigma_w = .5, \text{Averaging length} = 1000$

Figure 2.12: Transition probability estimates of 3 state HMM




Figure 2.13: Level estimates for reduced complexity RPE scheme






Figure 2.14: Transition probability estimates for reduced complexity RPE scheme



Derived from Figure 2.3

Figure 2.15: Error function for RPE scheme



Derived from RKL scheme with same data as for Figure 2.3

Figure 2.16: Error function for RKL scheme

Iterations	PEE(\hat{g}_1)	PEE(\hat{g}_2)	PEE(\hat{a}_{11})	PEE(\hat{a}_{22})
25000	0.085	0.088	0.042	0.047
50000	0.058	0.041	0.017	0.013
75000	0.045	0.041	0.015	0.010
100000	0.036	0.033	0.011	0.012

Table 2.1: Parameter estimation error (PEE) for SNR = 0dB

Iterations	PEE(\hat{g}_1)	PEE(\hat{g}_2)	PEE(\hat{a}_{11})	PEE(\hat{a}_{22})
25000	0.231	0.165	0.076	0.103
50000	0.026	0.128	0.056	0.068
75000	0.185	0.115	0.043	0.029
100000	0.162	0.127	0.043	0.027

Table 2.2: PEE for SNR = -6.0dB

Iterations	PEE(\hat{g}_1)	PEE(\hat{g}_2)	PEE(\hat{a}_{11})	PEE(\hat{a}_{22})
25000	0.238	0.224	0.063	0.121
50000	0.239	0.187	0.046	0.101
75000	0.219	0.186	0.040	0.084
100000	0.194	0.185	0.037	0.065

Table 2.3: PEE for SNR = -9.5dB

Iterations	PEE(\hat{g}_1)	PEE(\hat{g}_2)	PEE(\hat{a}_{11})	PEE(\hat{a}_{22})
25000	0.271	0.238	0.131	0.146
50000	0.245	0.232	0.159	0.137
75000	0.226	0.190	0.127	0.110
100000	0.210	0.210	0.099	0.085

Table 2.4: PEE for SNR = -12.0dB

$\hat{a}_{ii}(0)$	PEE(\hat{g}_1)	PEE(\hat{g}_2)	PEE(\hat{a}_{11})	PEE(\hat{a}_{22})
0.7	0.107	0.115	0.038	0.023
0.5	0.160	0.095	0.063	0.031
0.3	0.182	0.113	0.067	0.045
0.1	0.143	0.077	0.057	0.033

Results after 25000 iterations, $\hat{g}(0) = [0, 1]'$, SNR = 0dB

Table 2.5: PEE for variations in initial trans. prob. estimates

$\hat{g}_1(0)$	$\hat{g}_2(0)$	PEE(\hat{g}_1)	PEE(\hat{g}_2)	PEE(\hat{a}_{11})	PEE(\hat{a}_{22})
0.1	0.9	0.081	0.094	0.078	0.099
0.3	0.7	0.078	0.070	0.022	0.025
0.5	0.5	0.120	0.135	0.073	0.083

Results after 25000 iterations, $\hat{a}_{ii}(0) = 0.9$, SNR = 0dB

Table 2.6: PEE for variations in initial level estimates

Chapter 3

Identification of HMMs with Grouped States

3.1 Introduction

This chapter presents an algorithm which addresses the issue of computational complexity for on-line identification of hidden Markov models (HMMs). In Chapter 2, a recursive prediction error based scheme was presented for HMMs, for which the computational effort was in the order of the square of the number of parameters. This chapter considers HMMs for which the state values, g , of the Markov chain, are *clustered* into groups. This allows a reformulation of the Markov model and results in a sub-optimal reduced order identification scheme. The computational complexity of this new scheme is much lower than that for the full scheme of Chapter 2. Actually, an exact definition of clustering is not discussed, rather, a general identification technique is presented for which the computational requirements are greatly reduced when the state values are divided in some way into groups. The applicability to certain types of cluster patterns is tested via simulation studies.

An important motivation for this work comes from recent developments in multi-resolution communication systems. For such systems, essential information (for example, outlines in an image) is sent with a certain signalling scheme, while more detailed, but not essential, information (for example, shading in an image) is sent with a higher resolution signalling scheme, therefore re-

quiring higher resolution in the receiver in order to be decoded. In high noise conditions it may be that the high resolution signal can not be decoded, yet the low resolution signal, carrying the essential information, is successfully decoded. One example is with quadrature amplitude modulated (QAM) signals (as discussed later in Chapter 4) where the symbol constellation can be modified by dividing it into groups with non-equidistant symbols (and hence varying degrees of resolution). The general techniques of this chapter can be applied to such QAM systems with clustered state values in order to reduce the computational complexity of demodulation.

Much work has been carried out into grouping states associated with Markov chains (see for example, Simon and Ando [1961]). Techniques such as stochastic complementation [Meyer 1989] are sophisticated methods of producing reduced complexity representations of Markov chains which have large numbers of states. They have been used mainly to evaluate steady-state probability distributions [Cao and Stewart 1985] and reduced order controllers for Markov systems [Aldaheri and Khalil 1991, Delebecque and Quadrat 1981]. When there exist only weak interactions between groups of states, these techniques provide very accurate reduced order approximations [Courtois 1975]. The procedures, however, require knowledge of the transition probabilities, which are of course unknown in the case of model identification. Therefore, more straight forward state lumping techniques are of interest for the on-line identification problem considered in this thesis.

In this chapter a lumping procedure is proposed which produces a reduced order representation of the Markov chain, suitable for use with RPE parameter estimation techniques. It is an exact technique, according to the definition of lumpability in Kemeny and Snell [1960](p. 124), for Markov chains which have the property that for each state within a group, there is an equal probability of making a transition to any other given group. The technique is also relevant for Markov chains for which the transition probability matrix is doubly stochastic. However, in this case the result is only approximate, being based on the average, or steady-state, behaviour of the process. Other more general Markov chains can also be considered although knowledge of the steady state probability distribution, and hence the transition probability matrix, is required for such cases.

The general approach to identification in this chapter is the same as in Chapter 2, where the HMM is modelled in such a way as to allow an RPE algorithm to be applied. The parameters of interest are the state values and transition probabilities of the Markov chain, and the noise variance

associated with the measurements. In this chapter the transition probabilities are parameterised on a sphere so as to ensure that the derivatives are smooth, and that the estimates remain positive.

A key to the ideas presented in this chapter is the subdivision of the Markov chain, and allocation of associated *flag* states, representing a lumping of all states not in the particular group. This extra flag state is added to each group of state values to indicate whether the actual state is in the group or not, at each time instant. This is an essential idea which allows a reduced complexity HMM/RPE identification scheme to be developed.

Simulation examples are presented to illustrate the algorithms. These examples show that the proposed schemes can satisfactorily identify HMM parameters. They also illustrate that the distance between clusters, the distance between state values within clusters, and the noise level, all contribute to the achievable performance.

This chapter is organised as follows: Section 3.2 formulates the HMM and details the reduced order model. In Section 3.3 RPE based recursive algorithms for HMM estimation are derived, working on the sphere for the constrained transition probability estimates. In Section 3.4 simulation examples are given. Finally, conclusions are presented in Section 3.5.

3.2 Problem Formulation

This section presents modifications to the HMM signal model of Chapter 2, for the case where state values are clustered into distinct regions.

3.2.1 State Space Signal Model

Let X_k be a discrete-time homogeneous, first order Markov process belonging to a finite-discrete set, as in Section 2.2.1. Let N be the number of states in the Markov process, and let L be the number of groups into which the states are partitioned. Let the i^{th} group have $K^{(i)}$ elements and define

$$H^{(i)} = \sum_{\ell=1}^{i} K^{(\ell)}. \quad (3.1)$$

Now, the state space of X , without loss of generality, can be identified with the set of unit vectors $\mathbf{S} = \{e_1, e_2, \dots, e_N\}$, where $e_i = (0, \dots, 0, 1, 0, \dots, 0)' \in \mathbb{R}^N$ with 1 in the i^{th} position. The

dynamics of X_k are given, from Lemma 2.1, as follows:

$$X_{k+1} = \mathbf{A}' X_k + M_{k+1}, \quad (3.2)$$

where \mathbf{A} is the $N \times N$ transition probability matrix with elements $a_{ij} = P(X_{k+1} = e_j | X_k = e_i)$.

It is partitioned according to the grouping of states, in the following way:

$$\mathbf{A} = \begin{pmatrix} A_{11} & A_{12} & \dots & A_{1L} \\ A_{21} & A_{22} & \dots & A_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ A_{L1} & A_{L2} & \dots & A_{LL} \end{pmatrix}, \quad (3.3)$$

where A_{ii} is a square matrix of dimension $K^{(i)} \times K^{(i)}$. Of course $a_{ij} \geq 0$ and $\sum_{j=1}^N a_{ij} = 1$, for each i . It is also of use to denote $\{\mathcal{F}_l, l \in \mathcal{Z}^+\}$ the complete filtration generated by X , that is, for any $k \in \mathcal{Z}^+$, \mathcal{F}_k is the complete σ -field generated by $X_k, l \leq k$.

Also, consider the observation process defined in (2.2):

$$y_k = g(X_k) + w_k, \quad (3.4)$$

where without loss of generality, since X_k is in a discrete set, $g(X_k) = \langle g, X_k \rangle$, and $g \in \mathbb{R}^N$ is the vector of state values of the Markov chain. In this chapter it is assumed that the state values can be grouped as follows:

$$g = ((g^{(1)})', (g^{(2)})', \dots, (g^{(L)})')'. \quad (3.5)$$

Assume that w_k is independent and identically distributed (i.i.d.), with a zero mean Gaussian distribution, so that $w_k \sim N[0, \sigma_w^2]$. Also, let $\mathbf{b}_k = (b_k(1), \dots, b_k(N))'$ denote the vector of parameterised probability densities (which will loosely be called symbol probabilities), where $b_k(i) \triangleq b(y_k, g_i) = P[y_k | X_k = e_i, \theta]$. The explicit expression for $b(y_k, g_i)$ is given in (2.4).

In addition, the assumption is made that the initial state probability vector for the Markov chain $\underline{\pi} = (\pi_i)$ is defined from $\pi_i = P(X_1 = e_i)$. As in Chapter 2, the HMM is denoted $\lambda = (\mathbf{A}, g, \underline{\pi}, \sigma_w^2)$.

3.2.2 Reduced Order State Space Signal Model

In order to decouple the Markov chain, and hence obtain a reduced order signal model, it is necessary to introduce the concept of flags. For each group of state values, an additional *flag* state is defined which indicates, at each time, k , whether or not the state value of the Markov chain belongs to that particular group. In addition, a new state vector, and a corresponding set of indicator vectors, are defined for each group.

To illustrate the concept, consider a four state chain divided into two groups each with two elements. If $X_k = (0, 1, 0, 0)'$ for some k , then the corresponding grouped indicator vectors would be:

$$\begin{pmatrix} 0 \\ 1 \\ \text{---} \\ 0 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 0 \\ 0 \\ \text{---} \\ 1 \end{pmatrix},$$

where the bottom element of each vector is the flag. If the flag is zero then the Markov state value is in that group, if the flag is 1, then the Markov state value is in one of the other groups. In effect, the flag state represents a lumping together of all the states not in the particular group.

More precisely, this can be expressed in the following way. Let $X_k^{(i)}$ be the state vector associated with the i^{th} group. In this work, $X_k^{(i)}$ is called a *grouped state vector*. The dimension of $X_k^{(i)}$ is $K^{(i)} + 1$, as it represents the state values in the group, plus the flag. For the remainder of this chapter, let m and n be integers in the range $[1, N]$, and let x , z , t and r , be given by the following expressions:

$$x \triangleq m - H^{(t-1)} \quad \text{for } t \text{ such that } H^{(t-1)} + 1 \leq m \leq H^{(t)} \quad (3.6)$$

$$z \triangleq n - H^{(r-1)} \quad \text{for } r \text{ such that } H^{(r-1)} + 1 \leq n \leq H^{(r)} \quad (3.7)$$

where $H^{(i)}$ is defined in (3.1).

Now, the state space of $X_k^{(i)}$ can be identified with the set of unit vectors $\mathbf{S}^{(i)} = \{e_1^{(i)}, \dots, e_{K^{(i)}+1}^{(i)}\}$, where $e_j^{(i)} = (\dots, 0, 1, 0, \dots)' \in \mathbb{R}^{K^{(i)}+1}$, with 1 in the j^{th} position. Now, as seen in the illustration previously, only one grouped state vector will take a value other than $(\dots, 0, 0, 0, \dots, 1)'$, at each time, k . Specifically, it will be the grouped vector $X_k^{(t)} = e_x^{(t)}$, where t and x are given in (3.6), and where $X_k = e_m$.

The dynamics of each of the i grouped state vectors are given by the following equation, for $1 \leq i \leq L$:

$$X_{k+1}^{(i)} = (\mathbf{A}^{(i)})' X_k^{(i)} + M_{k+1}^{(i)}, \quad (3.8)$$

where $M_k^{(i)}$ depends on M_k (see Remark 3.1), and

$$\mathbf{A}^{(i)} = \begin{pmatrix} A_{ii} & A_{12}^{(i)} \\ A_{21}^{(i)} & A_{22}^{(i)} \end{pmatrix}, \quad (3.9)$$

where A_{ii} is defined in (3.3), and

$$A_{12}^{(i)} = \sum_{\substack{\ell=1 \\ \ell \neq i}}^L \sum_{j=1}^{K^{(\ell)}} A_{i\ell}(\cdot, j), \quad (3.10)$$

$$A_{21}^{(i)} = \frac{1}{J^{(i)}} \sum_{\substack{\ell=1 \\ \ell \neq i}}^L \sum_{j=1}^{K^{(\ell)}} A_{\ell i}(j, \cdot), \quad (3.11)$$

$$A_{22}^{(i)} = \frac{1}{J^{(i)}} \sum_{\substack{\ell=1 \\ \ell \neq i}}^L \sum_{n=1}^L \sum_{j=1}^{K^{(\ell)}} \sum_{x=1}^{K^{(n)}} A_{\ell n}(j, x), \quad (3.12)$$

where $A_{i\ell}(\cdot, j)$ is the j^{th} column of the matrix $A_{i\ell}$, $A_{\ell i}(j, \cdot)$ is the j^{th} row of the matrix $A_{\ell i}$, $A_{\ell n}(j, x)$ is the element in the j^{th} row and x^{th} column of the matrix $A_{\ell n}$, and

$$J^{(i)} = \sum_{\substack{\ell=1 \\ \ell \neq i}}^L K^{(\ell)}. \quad (3.13)$$

Remark 3.1: Equation (3.8) strictly holds only for lumpable or weakly lumpable Markov chains [Kemeny and Snell 1960]. In other cases, (3.8) is used as an approximation. The validity of such an approximate approach is discussed in Lindqvist [1978], and can be tested via simulation studies. \square

3.2.3 Model Parameterisation

As in Chapter 2, consider that λ is parameterised by an unknown vector θ so that $\lambda(\theta) = (\mathbf{A}(\theta), g(\theta), \pi, \sigma_w^2(\theta))$. In this chapter, only the parameterisation on the sphere is proposed, due to the superior performance gained over the parameterisation on the simplex as observed in Chapter 2. The dimension of θ is $N_\theta = N + N^2 + 1$, representing N state values, N^2 transition

probabilities, and the noise variance associated with the observations. The important aspect of the work in this chapter is that the identification scheme involves decoupling which greatly reduces the computational complexity.

1. Parameterisation on a Sphere :

$$\theta = (g_1, \dots, g_N, s_{11}, \dots, s_{1N}, s_{21}, \dots, s_{NN}, \sigma_w^2)' , \quad (3.14)$$

where $a_{ij} = s_{ij}^2$. As with the parameterisation (2.7), (3.14) has only the equality constraint of the sphere surface, S^{N-1} , where

$$S^{N-1} := \left\{ s_{ij} : \sum_{j=1}^N s_{ij}^2 = 1 \right\} , \quad (3.15)$$

and therefore ensures each transition probability estimate is positive.

The following state space signal model applies:

$$\begin{array}{l} X_{k+1}^{(1)} = (\mathbf{A}^{(1)}(\theta))' X_k^{(1)} + M_{k+1}^{(1)} \\ \vdots \\ X_{k+1}^{(L)} = (\mathbf{A}^{(L)}(\theta))' X_k^{(L)} + M_{k+1}^{(L)} \end{array} \quad (3.16)$$

$$y_k = \sum_{i=1}^L \left\langle \begin{pmatrix} g^{(i)}(\theta) \\ 0 \end{pmatrix}, X_k^{(i)} \right\rangle + w_k$$

This signal model is not, however, in a form suitable for application of RPE techniques to achieve estimates of θ and X_k from the measurements y_k (due to the finite-discrete nature of the states, $X_k^{(i)}$). Such a model is now developed.

3.2.4 Parameterised Information State Signal Model

Let $\hat{X}_{k|\theta}^{(i)}$ denote the conditional filtered state estimate of $X_k^{(i)}$ at time k , that is,

$$\hat{X}_{k|\theta}^{(i)} \triangleq E[X_k^{(i)} | \mathcal{Y}_k, \theta] . \quad (3.17)$$

Let $\mathbf{1}$ be defined to be the column vector containing all ones, and the information state $\alpha_{k|\theta}^{(i)}$ be a $K^{(i)} + 1$ dimension column vector with it's j^{th} element given by

$$\alpha_{k|\theta}^{(i)}(j) \triangleq P(Y_k, X_k^{(i)} = e_j^{(i)} | \theta). \quad (3.18)$$

Observe from Lemma 2.2, that $\hat{X}_{k|\theta}^{(i)}$ can be expressed in terms of the column vector $\alpha_{k|\theta}^{(i)}$, as follows:

$$\hat{X}_{k|\theta}^{(i)} = \langle \alpha_{k|\theta}^{(i)}, \mathbf{1} \rangle^{-1} \alpha_{k|\theta}^{(i)}. \quad (3.19)$$

Here, $\alpha_{k|\theta}^{(i)}$ is computed using the following ‘‘forward’’ recursion:

$$\alpha_{k+1|\theta}^{(i)} = \mathbf{B}^{(i)}(y_{k+1}, \theta) (\mathbf{A}^{(i)}(\theta))' \alpha_{k|\theta}^{(i)}, \quad (3.20)$$

where $\mathbf{B}^{(i)}(y_k, \theta) = \text{diag}(b(y_k, g_1^{(i)}), \dots, b(y_k, g_{K^{(i)}}^{(i)}), \sum_{\forall \ell \notin \{H^{(i-1)}+1, \dots, H^{(i)}\}} b(y_k, g_\ell))$.

Remark 3.2: It can be seen from these equations that the computational complexity associated with evaluating the information-states is in the order of $\sum_{i=1}^L (K^{(i)} + 1)^2$. This is compared to N^2 for the scheme of Chapter 2. In most cases, $\sum_{i=1}^L (K^{(i)} + 1)^2 \ll N^2$. For example, if $K^{(i)} = C$, $\forall i$, then the computational requirements of the grouped algorithm presented in this chapter are less than those for the full scheme when the number of groups, L , satisfies the following inequality: $L > 1 + \frac{2}{C} + \frac{1}{C^2} > 2$. \square

It is now necessary to express the observations, y_k , in terms of the un-normalised conditional estimates, $\alpha_{k|\theta}^{(i)}$.

Lemma 3.1 *The conditional measurements $y_{k|\theta}$ are given by*

$$y_{k|\theta} = \sum_{i=1}^L \left\langle \begin{pmatrix} g^{(i)}(\theta) \\ 0 \end{pmatrix}, \frac{(\mathbf{A}^{(i)})' \alpha_{k-1|\theta}^{(i)}}{\langle \alpha_{k-1|\theta}^{(i)}, \mathbf{1} \rangle} \right\rangle + n_k,$$

where $\alpha_k^{(i)}$ is defined in (3.18) and n_k is a $(\theta, \mathcal{Y}_{k-1})$ martingale increment.

Proof : The proof follows the same arguments as that for Lemma 2.4. \blacksquare

In summary, the parameterised estimator based signal model for an HMM parameter θ , and with

states $\alpha_{k|\theta}^{(i)}$, is given by

$$\begin{aligned}
 \alpha_{k+1|\theta}^{(1)} &= \mathbf{B}^{(1)}(y_{k+1}, \theta) (\mathbf{A}^{(1)}(\theta))' \alpha_{k|\theta}^{(1)} \\
 &\vdots \\
 \alpha_{k+1|\theta}^{(L)} &= \mathbf{B}^{(L)}(y_{k+1}, \theta) (\mathbf{A}^{(L)}(\theta))' \alpha_{k|\theta}^{(L)} \\
 \\
 y_{k|\theta} &= \sum_{i=1}^L \left\langle \begin{pmatrix} g^{(i)}(\theta) \\ 0 \end{pmatrix}, \frac{(\mathbf{A}^{(i)})' \alpha_{k-1|\theta}^{(i)}}{\langle \alpha_{k-1|\theta}^{(i)}, \mathbf{1} \rangle} \right\rangle + n_k
 \end{aligned} \tag{3.21}$$

where n_k is a $(\theta, \mathcal{Y}_{k-1})$ martingale increment. This signal model is now in a suitable form to apply RPE algorithms to achieve, simultaneously, state and parameter estimates, on-line.

3.3 Gradient Vector and Projection Calculations

In this section gradient and projection calculations for the RPE based algorithms (2.22)-(2.27) are presented, to estimate the HMM parameters, θ . As in Chapter 2, these algorithms update $\hat{\theta}_k$ and R_k recursively, and require $\hat{y}_{k|\hat{\theta}_{k-1}}$ and ψ_k to be evaluated at each update, where $\hat{\Theta}_k \triangleq \{\hat{\theta}_1, \dots, \hat{\theta}_k\}$, and

$$\hat{y}_{k|\hat{\theta}_{k-2}, \hat{\theta}_{k-1}} = \sum_{i=1}^L \left\langle \begin{pmatrix} g^{(i)}(\hat{\theta}_{k-1}) \\ 0 \end{pmatrix}, \frac{(\mathbf{A}^{(i)}(\hat{\theta}_{k-1}))' \hat{\alpha}_{k-1|\hat{\theta}_{k-2}}^{(i)}}{\langle \hat{\alpha}_{k-1|\hat{\theta}_{k-2}}^{(i)}, \mathbf{1} \rangle} \right\rangle. \tag{3.22}$$

The derivative vector, ψ_k , defined in (2.27), is given, for $m, n \in [1, \dots, N]$, in the case of parameterisation (3.14), by

$$\psi_k = \left. \frac{\partial \hat{y}_{k|\hat{\theta}_{k-2}, \theta}}{\partial \theta} \right|_{\theta=\hat{\theta}_{k-1}} = \left(\frac{\partial \hat{y}_{k|\hat{\theta}_{k-2}, \theta}}{\partial g_m}, \frac{\partial \hat{y}_{k|\hat{\theta}_{k-2}, \theta}}{\partial s_{mn}}, \frac{\partial \hat{y}_{k|\hat{\theta}_{k-2}, \theta}}{\partial \sigma_w^2} \right)' \Big|_{\theta=\hat{\theta}_{k-1}}, \tag{3.23}$$

where $\hat{y}_{k|\hat{\theta}_{k-2}, \theta}$ is given in (3.22) by replacing $\hat{\theta}_{k-1}$ with θ . In the remainder of this section the dependence of $\hat{\alpha}_k^{(i)}$ on $\hat{\Theta}_{k-1}$ is omitted, and $N_k^{(i)} \triangleq \langle \hat{\alpha}_k^{(i)}, \mathbf{1} \rangle^{-1}$, where

$$\hat{\alpha}_{k+1|\hat{\Theta}_k}^{(i)} = \mathbf{B}^{(i)}(y_{k+1}, \hat{\theta}_k) (\mathbf{A}^{(i)}(\hat{\theta}_k))' \hat{\alpha}_{k|\hat{\Theta}_{k-1}}^{(i)}. \tag{3.24}$$

The derivatives with respect to the discrete-state values, g_i , are obtained by differentiating (3.22) to yield the following:

$$\frac{\partial \hat{y}_{k+1} \hat{\theta}_{k-1, \theta}}{\partial g_m} = N_k^{(t)} \left(\begin{array}{c} A_{tt}(\cdot, x) \\ A_{21}^{(t)}(x) \end{array} \right)' \hat{\alpha}_k^{(t)} + \sum_{i=1}^L \left(N_k^{(i)} \left\langle g^{(i)}, \left(\begin{array}{c} A_{ii} \\ A_{21}^{(i)} \end{array} \right)' \eta_k^{(i)}(m) \right\rangle - (N_k^{(i)})^2 \langle \mathbf{1}, \eta_k^{(i)}(m) \rangle \left\langle g^{(i)}, \left(\begin{array}{c} A_{ii} \\ A_{21}^{(i)} \end{array} \right)' \hat{\alpha}_k^{(i)} \right\rangle \right) \quad (3.25)$$

where $A_{tt}(\cdot, x) = (A_{tt}(1, x), \dots, A_{tt}(K^{(t)}, x))'$, for x and t defined in (3.6), and where $\eta_k^{(t)}(m) = \partial \hat{\alpha}_k^{(t)} / \partial g_m$ is the $(K^{(t)} + 1)$ -dimensional vector given by the following recursive equation:

$$\eta_{k+1}^{(i)}(m) = \mathbf{B}^{(i)}(y_{k+1}, \theta) (\mathbf{A}^{(i)})' \eta_k^{(i)}(m) + \text{diag}(e_\ell^{(i)}) \left(\frac{y_{k+1} - g_m}{\sigma_w^2} \right) b(y_{k+1}, g_m) (\mathbf{A}^{(i)})' \hat{\alpha}_k^{(i)}, \quad (3.26)$$

where

$$e_\ell^{(i)} = \begin{cases} e_x^{(i)} & \text{if } H^{(i-1)} + 1 \leq m \leq H^{(i)} \\ e_{K^{(i)}+1}^{(i)} & \text{otherwise,} \end{cases} \quad (3.27)$$

The projected derivatives with respect to the square root of the transition probabilities, s_{ij} , are given by

$$\frac{\partial \hat{y}_{k+1} \hat{\theta}_{k-1, \theta}}{\partial s_{mn}} = 2F s_{mn} \text{diag}(g) \left(e_n - \text{diag}(s_{m(\cdot)}) s'_{m(\cdot)} \right) + \sum_{i=1}^L \left(N_k^{(i)} \left\langle g^{(i)}, \left(\begin{array}{c} A_{ii} \\ A_{21}^{(i)} \end{array} \right)' \xi_k^{(i)}(m, n) \right\rangle - (N_k^{(i)})^2 \langle \mathbf{1}, \xi_k^{(i)}(m, n) \rangle \left\langle g^{(i)}, \left(\begin{array}{c} A_{ii} \\ A_{21}^{(i)} \end{array} \right)' \hat{\alpha}_k^{(i)} \right\rangle \right) \quad (3.28)$$

where

$$F = [f_1(K^{(1)} + 1), \dots, f_{t-1}(K^{(t-1)} + 1), f_t(x), f_{t+1}(K^{(t+1)} + 1), \dots, f_L(K^{(L)} + 1)] , \quad (3.29)$$

for $f_i(j)$ defined by

$$f_i(j) = G_i N_k^{(i)} \hat{\alpha}_k^{(i)}(j) \mathbf{1}'_{K^{(i)}}, \quad (3.30)$$

with

$$G_i = \begin{cases} 1 & \text{if } H^{(t-1)} + 1 \leq i \leq H^{(t)} \\ 1/J^{(i)} & \text{otherwise} \end{cases} \quad (3.31)$$

where t and x are given in (3.6).

Also, $\xi_k^{(i)}(m, n) = \partial \hat{\alpha}_k^{(i)} / \partial s_{mn}$ is the $(K^{(i)} + 1)$ -dimensional vector given by the following recursive equation:

$$\begin{aligned} \xi_{k+1}^{(i)}(m, n) = & \mathbf{B}^{(i)}(y_{k+1}, \theta) (\mathbf{A}^{(i)})' \xi_k^{(i)}(m, n) \\ & + 2C_\ell \hat{\alpha}_k^{(i)}(\ell) s_{mn} \mathbf{B}^{(i)}(y_{k+1}, \theta) E^{(i)} \left(e_n - \text{diag}(s_{m(\cdot)}) s'_{m(\cdot)} \right) \end{aligned} \quad (3.32)$$

where

$$E^{(i)} = \begin{bmatrix} 0 & \dots & 0 & I_{K^{(i)}} & 0 & \dots & 0 \\ \mathbf{1}'_{K^{(1)}} & \dots & \mathbf{1}'_{K^{(i-1)}} & 0 \mathbf{1}'_{K^{(i+1)}} & \dots & \mathbf{1}'_{K^{(L)}} \end{bmatrix},$$

$$C_\ell = \begin{cases} 1 & \text{if } r = t \\ 1/J^{(r)} & \text{if } r \neq t \end{cases} \quad \text{and} \quad \hat{\alpha}_k^{(r)}(\ell) = \begin{cases} \hat{\alpha}_k^{(r)}(x) & \text{if } r = t \\ \hat{\alpha}_k^{(r)}(K^{(r)} + 1) & \text{if } r \neq t \end{cases}.$$

As in Section 2.4.2, in achieving an update estimate of s_{ij} at time $k + 1$ via (2.22), denoted here $\hat{s}_{ij}(k + 1)$, there is a required projection $\Gamma_{proj}\{\cdot\}$ into the constraint domain (3.15), the surface of a unit sphere in $\mathbb{R}^{H^{(L)}}$. This is performed using (2.41) to achieve $\sum_{j=1}^{H^{(L)}} \hat{s}_{ij}^2(k + 1) = 1$ as required.

3.4 Implementation Considerations and Simulations

Simulations of the RPE scheme, as applied to the grouped HMM formulation presented in this chapter, have been carried out on six-state and nine-state Markov chains. This section presents some practical implementation aspects which should be considered in simulations, and also the results of the simulations.

3.4.1 Implementation Considerations

As with the scheme presented in Chapter 2, important implementation aspects include scaling, multi-pass processing, fixed-lag smoothing, stability, and re-initialisation. Such topics have previously been discussed in Section 2.5.1. This section discusses issues specific to the grouped HMM problem.

Initialisation and Projection : Initialisation needs to be addressed to insure all assumptions are satisfied. The main requirements are that: initial parameter estimates are close to the actual values and satisfy the model constraints, initial gradients are normal to the surface of the sphere in $\mathbb{R}^{H^{(L)}}$

(for (3.28) and (3.32)), and the covariance matrix is initialised in such a way that the parameter error remains bounded.

In Chapter 2 simulation studies showed that initial estimates could be chosen far from the actual values. In this chapter, where the grouping assumption has been made, more stringent conditions apply. However, it has been observed that as a general rule, as long as the initial estimates are grouped in a similar way to the actual states, then convergence will occur.

De-coupling of the Covariance Matrix : An important aspect of the grouping procedure presented in this chapter, is that the covariance matrix of the RPE estimate, R_k , can be assumed to be approximately block diagonal. In fact, it can be seen that there is little degradation to the performance, by considering the matrix to be strictly block diagonal. This is because in this grouped case, interactions within groups are much stronger than those between groups. In fact it is often possible to consider the matrix to be strictly diagonal, as in (2.29).

3.4.2 Simulation Studies

Presented here are results of simulation studies using computer generated finite state Markov chains in WGN.

Example 3.1: A six state Markov chain embedded in WGN was generated with parameter values

$$\mathbf{A} = \begin{pmatrix} .957 & .02 & .02 & .001 & .001 & .001 \\ .02 & .957 & .02 & .001 & .001 & .001 \\ .02 & .02 & .957 & .001 & .001 & .001 \\ .001 & .001 & .001 & .957 & .02 & .02 \\ .001 & .001 & .001 & .02 & .957 & .02 \\ .001 & .001 & .001 & .02 & .02 & .957 \end{pmatrix}, \quad \mathbf{g} = \begin{pmatrix} -4 \\ -3 \\ -2 \\ 2 \\ 3 \\ 4 \end{pmatrix},$$

and $\sigma_w = 1$. This Markov chain satisfies the conditions of lumpability [Kemeny and Snell 1960], and as such, there is no approximation in the definition of the grouped transition probability matrices. In this example, the HMM was divided into two groups and the RPE algorithm was used to estimate the state values, with the transition probabilities being assumed to be known. The algorithm was initialised with parameter estimates, $\hat{g} = [-2.6, -2.5, -2.4, 2.4, 2.5, 2.6]'$. The reduced complexity RPE scheme is used, where the covariance matrix is assumed to be strictly diagonal, as discussed in Section 3.4.1. In addition, for the covariance matrix recursion, $\gamma_k = 1/k$,

as is strictly the case from its derivation. However, for the parameter update, $\gamma_k = 1/\sqrt{k}$, and averaging is implemented to achieve faster convergence, as discussed in Section 2.3.3.

Figure 3.1 shows parameter estimates for this data. It also shows, with a dotted line, the actual noise-free data, $\langle g, X_k \rangle$, to indicate that the grouping of states results in the RPE algorithm only making significant updates to state-value estimates in the group which is active at the time.

Example 3.2: A nine state Markov Chain embedded in WGN was generated with parameter values

$$A_{ii} = \begin{pmatrix} .97 & .012 & .012 \\ .012 & .97 & .012 \\ .012 & .012 & .97 \end{pmatrix}, \quad A_{ij} = \begin{pmatrix} .002 & .001 & 0 \\ 0 & .002 & .001 \\ .001 & 0 & .002 \end{pmatrix} \quad \text{for } i \neq j,$$

$g = (1, 2, 3, 7, 8, 9, 13, 14, 15)'$, and $\sigma_w = 0.5$. The HMM was divided into three groups and the RPE scheme was implemented with the same conditions as in Example 3.1. The initial state value estimates were $\hat{g} = (1.8, 2, 2.2, 7.8, 8, 8.2, 13.8, 14, 14.2)'$.

Figure 3.2 shows parameter estimates for this data. This figure demonstrates that the scheme works well for Markov chains which do not satisfy the strict conditions of lumpability [Kemeny and Snell 1960], and are thus more widely applicable in practise. It also shows that when more states are considered, more data is required, as is expected from persistence of excitation requirements associated with the RPE scheme [Ljung and Söderström 1983].

Example 3.3: A four state Markov chain embedded in WGN was generated with parameter values

$$\mathbf{A} = \begin{pmatrix} .9 & .08 & .01 & .01 \\ .08 & .9 & .01 & .01 \\ .01 & .01 & .9 & .08 \\ .01 & .01 & .08 & .9 \end{pmatrix}, \quad g = \begin{pmatrix} -2 \\ -1 \\ 1 \\ 2 \end{pmatrix},$$

and $\sigma_w = 0.5$. This Markov chain satisfies the conditions of lumpability [Kemeny and Snell 1960], and as such, there is no approximation in the definition of the grouped transition probability matrices. In this example, the HMM was divided into two groups and the RPE algorithm was used to estimate the transition probabilities as well as the state values. A diagonal covariance matrix

was used, as discussed in Section 3.4.1. The algorithm was initialised with parameter estimates,

$$\hat{\mathbf{A}} = \begin{pmatrix} .5 & .48 & .01 & .01 \\ .48 & .5 & .01 & .01 \\ .01 & .01 & .5 & .48 \\ .01 & .01 & .48 & .5 \end{pmatrix}, \quad g = \begin{pmatrix} -1.7 \\ -1.3 \\ 1.3 \\ 1.7 \end{pmatrix}.$$

Figures 3.3 and 3.4 show parameter estimates for this data. The figures demonstrate that the estimates converge to the true values. It was found, as in Example 3.1, that as long as the initial state value estimates are grouped similarly to the true values, then convergence will occur for a wide range of initial transition probability estimates. The rate of convergence for the estimates of each group, of course depends on the characteristics of the excitation of that group.

Example 3.4: A four state Markov chain embedded in WGN was generated with parameter values the same as those in Example 3.3, except that in this case higher noise variance is considered, where $\sigma_w = 1$. The results are presented in Figures 3.5 and 3.6. This example shows that even under these higher noise conditions, and with a reduced complexity RPE scheme, the grouped HMM/RPE algorithm satisfactorily estimates both the state values and the transition probabilities.

Example 3.5: A six state Markov chain embedded in WGN was generated with parameter values

$$A_{ii} = \begin{pmatrix} .8 & .16 \\ .16 & .8 \end{pmatrix}, \quad A_{ij} = \begin{pmatrix} .01 & .01 \\ .01 & .01 \end{pmatrix} \quad \text{for } i \neq j,$$

$g = (1, 2, 6, 7, 11, 12)'$, and $\sigma_w = 0.5$. The HMM was divided into three groups and the RPE scheme was used to estimate the transition probabilities, with the state values being assumed to be known. The RPE scheme was implemented with the same conditions as in Example 3.1. The initial transition probability estimates were

$$\hat{A}_{ii} = \begin{pmatrix} .1 & .86 \\ .86 & .1 \end{pmatrix}, \quad \hat{A}_{ij} = \begin{pmatrix} .01 & .01 \\ .01 & .01 \end{pmatrix} \quad \text{for } i \neq j,$$

In this example, it was found to be advantageous to apply a threshold to the error $\hat{n}_{k|\hat{\Theta}_{k-1}}$. This is due to the fact that the groups of state values are far apart in relation to the Gaussian observation

noise variance (as discussed at the end of Section 2.3.3). The threshold was chosen to be twice the standard deviation, σ_w .

Figure 3.7 shows parameter estimates for this data. The figure demonstrates that the estimates converge to the true values. There is no limit to the number of states which can be estimated, however the length of data required for accurate estimates increases as the number of states increases.

Example 3.6: A six state Markov chain embedded in WGN was generated with parameter values

$$A_{ii} = \begin{pmatrix} .8 & .097 & .097 \\ .097 & .8 & .097 \\ .097 & .097 & .8 \end{pmatrix}, \quad A_{ij} = \begin{pmatrix} .002 & .002 & .002 \\ .002 & .002 & .002 \\ .002 & .002 & .002 \end{pmatrix} \quad \text{for } i \neq j,$$

$g = (-6, -5, -4, 4, 5, 6)'$, and $\sigma_w = 0.3$. The HMM was divided into two groups and the RPE scheme was used to estimate the transition probabilities with the state values being assumed known. The RPE scheme was implemented with the same conditions as in Example 3.1. The initial transition probability estimates were

$$\hat{A}_{ii} = \begin{pmatrix} .5 & .247 & .247 \\ .247 & .5 & .247 \\ .247 & .247 & .5 \end{pmatrix}, \quad \hat{A}_{ij} = \begin{pmatrix} .002 & .002 & .002 \\ .002 & .002 & .002 \\ .002 & .002 & .002 \end{pmatrix} \quad \text{for } i \neq j,$$

In this example it was found that transition probabilities relating to state values which were surrounded by other state values in the same group (that is, they were not the boundary state values), required larger values of γ_k in the RPE algorithm. In this example, the γ_k was multiplied by 10 for estimates of a_{ij} for $i = 2$ and $i = 5$.

Figure 3.8 shows parameter estimates for this data.

3.5 Conclusions

In this chapter a reduced order RPE based on-line parameter identification scheme for grouped HMMs has been derived. The algorithms presented are memory efficient and of lower computational complexity than the scheme of Chapter 2. Implementation aspects for the RPE scheme have been discussed, and the ability to estimate parameters of both six and nine state Markov chains in WGN has been demonstrated.

3.6 Figures and Tables

The figures and tables for this chapter are now presented.




Figure 3.1: Level estimates of 6 state Markov chain divided into 2 groups




Figure 3.2: Level estimates of 9 state Markov chain divided into 3 groups




Figure 3.3: Level estimates of 4 state Markov chain divided into 2 groups

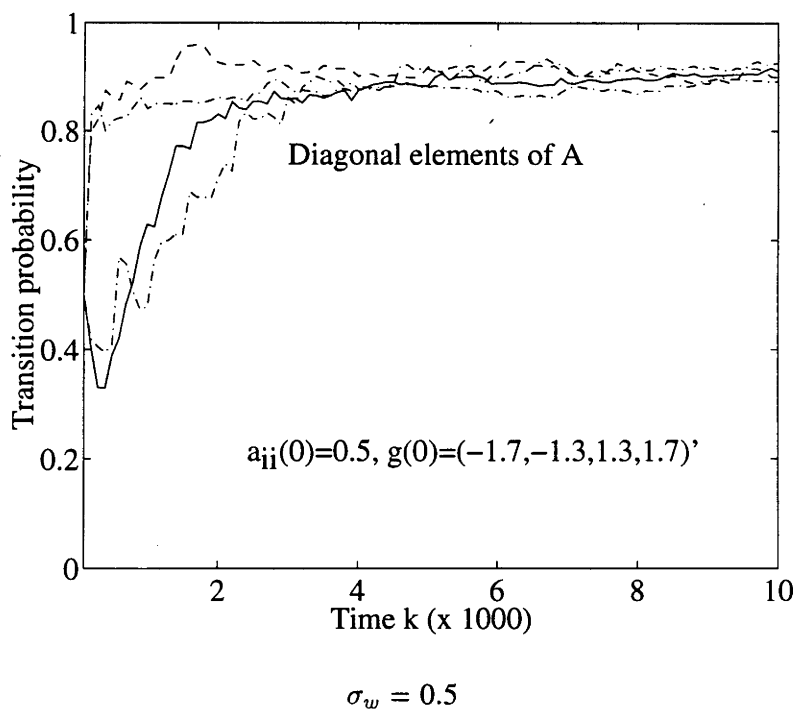


Figure 3.4: Transition probability estimates of 4 state Markov chain divided into 2 groups




Figure 3.5: Level estimates in higher noise conditions

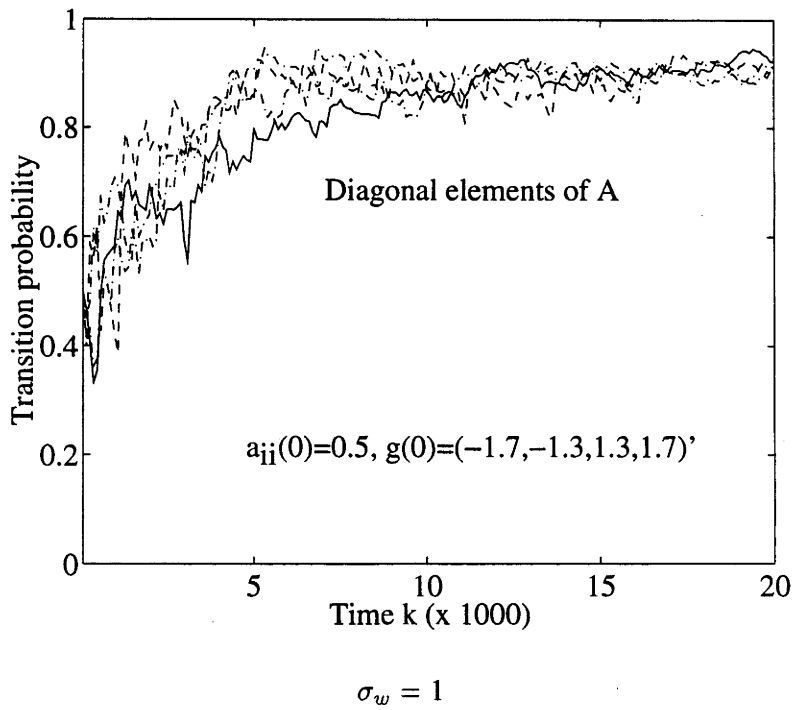


Figure 3.6: Transition probability estimates in higher noise conditions




Figure 3.7: Transition probability estimates with 6 states divided into 3 groups




Figure 3.8: Transition probability estimates with 6 states divided into 2 groups

Chapter 4

HMM Processing for QAM Digital Communication Systems

4.1 Introduction

This chapter considers the problem of fading channels in digital communication systems. In contrast to the model identification work of Chapters 2 and 3, this problem requires consideration of adaptive estimation schemes. Some quite general, new and novel, hidden Markov model signal processing techniques are developed to adaptively track transmission channel parameters and estimate the digital signal. The specific modulation scheme considered, is that of quadrature amplitude modulation (QAM), which is used extensively in communication systems. Such a study illustrates the application potential for hidden Markov model signal processing techniques.

The problem of fading channels can be the limiting factor in many communication systems. For example, a multi-path Rayleigh fading mobile telephone channel can introduce amplitude gain, and phase shifts, to the transmitted signal, making it unrecognisable at the receiver. Demodulation of signals under fading conditions requires adaptive estimation of the transmission channel characteristics. The traditional signal model formulation for modulated digital signals leads to a nonlinear task for estimating both the signal and the channel distortions. This is commonly performed using a matched filter (MF), or Viterbi algorithm, for state estimation, and an analog PLL operating in tandem with an AGC, for channel estimation, as discussed in Bingham

[1988](Ch. 5,6). However, these ML and MAP techniques are known to be far from optimal under fading conditions. Other more recent techniques involve the use of training sequences (as used for example in the European Global System for Mobile Communication (GSM) digital standard). Unfortunately, however, these schemes are also sub-optimal. The main drawback being that up to 20% of the transmission time can be required for sending the training data. Optimal schemes, on the other hand, are inherently infinite dimensional and are thus impractical, (as discussed for example in Haeb and Meyr [1989]). Also, they may not be robust to modelling errors. The challenge, therefore, is to devise sub-optimal robust and efficient demodulation schemes to cope with fading signals, particularly in the case where the message symbols are not equally-probable.

Recent approaches to the fading problem for QAM signals have involved the use of pilot symbol-aided schemes [Sampei and Sunaga 1993] and alterations to the QAM signal constellation [Webb *et al.* 1991]. These deal mainly with the transmitter in an effort to improve the bit-error-rates (BERs). The approach in this chapter is to apply hidden Markov modelling at the receiver, and as such, these HMM schemes can be implemented in tandem with the above techniques. It should be pointed out that this HMM approach does not require any modification to the transmitter.

In tackling demodulation using recent techniques in stochastic and adaptive systems, it is worth recalling the role of the Kalman filter (KF) and extended Kalman filter (EKF) (The EKF turns out to be a PLL in disguise). Examples of the use of the EKF are given in Anderson and Moore [1979] (target tracking p. 53, frequency modulation p. 200). More recent schemes have been developed coupling Kalman filtering techniques with maximum likelihood sequence estimation, for continuous phase modulated (CPM) signals in fading channels [Lodge and Moher 1990]. The approach of this thesis, for QAM signals, is to use EKF techniques coupled with hidden Markov model filtering.

To date, HMM filters have been widely applied in areas such as speech processing [Rabiner 1989] and biological signal processing [Chung *et al.* 1991], however these applications have involved off-line analysis. An important aspect of the work in this chapter is the application of new on-line HMM processing techniques to problems involving time-varying parameters.

In this chapter, KF/EKF techniques and HMM signal processing techniques are coupled in an adaptive HMM approach, to estimate both the signal, and the time-varying transmission channel parameters, on-line. The HMM filter is ideal for signals which do not have equally-probable (or i.i.d.) message symbols, as is the case with coded signals for example. Coding techniques

such as convolutional coding [Proakis 1983] (p.441), produce signals which are not i.i.d. and as such display Markov properties. Also, recent investigations in trellis coding (reported in Du and Vucetic [1991]) suggest that a similar situation arises for trellis coded signals. Actually, in the un-coded case, with equally-probable message symbols, the HMM filter with a maximum *a posteriori* (MAP) estimate, is identical to the matched filter, which is known to be optimal for non-fading i.i.d. digital signals. However, when fading is present, even in the i.i.d. case, there is an advantage to the non-MAP estimate HMM approach (which makes use of the full information state rather than MAP estimates), because more information about the message is being fed back to the channel tracking algorithm.

In this chapter the technical approach is to work with the signal in a discrete set, and associate with this signal a finite-discrete state vector X_k . X_k is an indicator vector for the signal. In the case of parity check encoders, each of the allowable values of X_k would represent one of the QAM signal constellation points. For convolutional codes, as mentioned in Chapter 1, each allowable value of X_k would represent a different combination of bits in the shift register. The states X_k are assumed to be first order Markov with known transition probability matrix A and state values Z . This is a reasonable assumption given that the coding scheme is known. In less friendly communication environments, where codes may not be known, the techniques of Chapter 2 could be used in conjunction with these tracking schemes, to identify parameters associated with the coding algorithm.

Associated with the channel are time-varying parameters (gain, phase shift, and noise colour), which are modelled as states x_k , in a continuous range, $x_k \in \mathbb{R}^n$. The channel parameters arise from a known linear time-invariant stochastic system. In this chapter, state space models are formulated involving a mixture of the states X_k and x_k , and are termed *mixed state models*. These are reformulated using HMM filtering theory to achieve a nonlinear representation with a state vector consisting of α_k and x_k , where α_k is an un-normalised information state, representing a discrete state conditional probability density for X_k . These reformulated models are termed *conditional information state models*. Next, the EKF algorithm, or some derivative scheme, can be applied for state estimation of this innovations representation, thereby achieving both signal and channel estimation. The resulting adaptive HMM algorithms appear either as coupled KF and HMM filters, or as a more sophisticated EKF with an HMM filter as a sub filter.

This new idea of using information states as opposed to MAP estimates, or matched filters, is an important aspect of this work. It allows information about the statistics of the state estimate to

be used, in contrast to the more traditional MAP estimate approaches discussed for example in Haeb and Meyr [1989]. It is also important to note that while the algorithms presented in this thesis involve simultaneous coupled state estimation and channel tracking (that is, with each new observation a state estimate is made and the channel estimate adjusted accordingly) the techniques can easily be extended to sequence estimation, in order to generate new information-state versions of the traditional MAP based Viterbi algorithms. These extensions can be considered to be similar to more recent soft-output Viterbi algorithms (SOVA), where accuracy information is supplied along with the MAP sequence estimate. The advantage however, in an information-state sequence estimator, is that through the techniques of this chapter, a systematic method is in place for the use of such accuracy information in the channel tracker.

In addition to reformulating the QAM signal representation, a non-intuitive channel representation is employed. Rather than work directly with a linear stochastic model for channel gain and phase shift, it is proposed to formulate the channel in terms of a linear stochastic model with the state being the real and imaginary components of the channel. Working in rectangular co-ordinates instead of polar co-ordinates allows the models to be written in a familiar bi-linear state space form driven by Gaussian noise. This bi-linear representation facilitates the application of a Kalman filter, as opposed to a nonlinear PLL. Unfortunately the rectangular co-ordinate representation introduces coupling between the two noise sources in the model. This coupling is, however, well understood.

When the channels are time-invariant (non-fading), the EKF and derivative schemes are equivalent to the recursive prediction error (RPE) approach for HMM identification and estimation, the subject of Chapter 2. There are quite solid theoretical foundations in the RPE case, giving confidence of asymptotic optimality, with quadratic convergence rates. When the channels are fading, however, the problem falls, in general, within the context of EKF theory which is less developed. Therefore, strong theoretical convergence results are not sought here, save that it is expected from known theory that in the low noise case, the EKF is near optimal after initial transients.

This chapter is organised as follows : In Section 4.2 the QAM signal model is formulated in the HMM framework. In Section 4.3 the HMM/EKF and HMM/KF adaptive algorithms are presented. Coloured noise is considered in Section 4.4. Section 4.6 gives simulation examples which demonstrate good tracking ability for fast changing channels. Finally, conclusions are presented in Section 4.7.

4.2 Quadrature Amplitude Modulation (QAM)

Digital information grouped into fixed length bit strings, or out-putted from a convolutional encoder, is frequently represented by suitably spaced points in the complex plane. Quadrature amplitude modulation (QAM) transmission schemes are based on such a representation. In this section, the usual QAM signal model is presented, followed by a reformulation so as to apply hidden Markov model (HMM) and extended Kalman filtering (EKF) methods.

This section considers the set of states to be the set of complex message symbols constituting the QAM constellation. This is the situation arising with parity check encoders, and is presented here in order to simplify notation. The state could easily have been considered to be the contents of a shift register, as in the convolutional coded case, however, whilst the approach is the same, the resulting equations are slightly more complicated.

4.2.1 Signal Model

Let m_k be a complex discrete-time signal, ($k \in \mathcal{Z}^+$), where for each k ,

$$m_k \in \mathbf{Z} = \{z^{(1)}, \dots, z^{(2^N)}\}, \text{ where } z^{(i)} \in \mathbf{C}, N \in \mathcal{Z}^+. \quad (4.1)$$

In addition, the vector z is defined by

$$z = z^R + \mathbf{j}z^I = (z^{(1)}, \dots, z^{(2^N)})' \in \mathbf{C}^{2^N}. \quad (4.2)$$

For digital transmission, each element of \mathbf{Z} is used to represent a string of N bits. In the case of QAM, each of the complex elements, $z^{(i)}$, is chosen so as to generate a rectangular grid of equally spaced points in the complex space \mathbf{C} . A 16 state ($N = 4$) QAM signal constellation is illustrated in Figure 4.1.

Note that at any time, k , the message, $m_k \in \mathbf{Z}$, is complex valued and can be represented in either polar or rectangular form, as follows:

$$m_k = \rho_k \exp[\mathbf{j} \Upsilon_k] = m_k^R + \mathbf{j}m_k^I. \quad (4.3)$$

The real and imaginary components of m_k can be used to generate piece-wise constant time signals, $m(t) = m_k$ for $t = [t_k, t_{k+1})$, where t_k arises from regular sampling. The messages are




Figure 4.1: 16 state QAM Signal Constellation

then modulated and transmitted in quadrature as a QAM bandpass signal

$$s(t) = A_c[m^R(t) \cos(2\pi ft + \theta) + m^I(t) \sin(2\pi ft + \theta)] , \quad (4.4)$$

where the carrier amplitude A_c , frequency f , and phase θ are constant. This transmission scheme is termed QAM because the signal is *quadrature* in nature, where the real and imaginary components of the message are transmitted as two *amplitudes* which *modulate* quadrature and in-phase carriers.

4.2.2 Channel Model

The QAM signal is passed through a channel which can cause amplitude and phase shifts, as for example, in fading channels due to multiple transmission paths. The channel model considered in this chapter is appropriate for narrow band digital transmission systems, such as time-division multiple access (TDMA) mobile communication systems. In these cases the bandwidth is narrow enough that fading occurs evenly across the band, or in other words, it is non-frequency selective, or *flat fading*. Therefore, the channel has a delta impulse response, but with a time varying amplitude and phase shift to the modulating carrier, and hence no inter-symbol interference (ISI) takes place. In cases where ISI is present, the task of tracking fading channels is more difficult. HMM techniques can be applied to the ISI problem, however such techniques are not presented in this thesis, as they are the topic of current research.

In the flat fading case, the channel can be modelled by a multiplicative disturbance, $g(t)$, resulting in a discrete time baseband disturbance g_k ,

$$g_k = \kappa_k \exp[\mathbf{j} \phi_k] = g_k^R + \mathbf{j} g_k^I, \quad (4.5)$$

which introduces time-varying gain and phase changes to the signal. The time variations in g_k are realistically assumed to be slow in comparison to the message rate.

Channel State - Cartesian Co-ordinate Representation : In this co-ordinate system, the vector x_k is associated with the real and imaginary parts of g_k .

$$x_k = \begin{pmatrix} \kappa_k \cos \phi_k \\ \kappa_k \sin \phi_k \end{pmatrix} = \begin{pmatrix} g_k^R \\ g_k^I \end{pmatrix}. \quad (4.6)$$

Channel State - Polar Co-ordinate Representation : An alternative to Cartesian co-ordinates, in the complex plane, is the more traditional polar co-ordinate representation.

$$x_k^P = \begin{pmatrix} \kappa_k \\ \phi_k \end{pmatrix}. \quad (4.7)$$

As mentioned previously, the Cartesian co-ordinates allow the observations to be written in a form which enables linear Kalman filtering to be applied, while the polar co-ordinates require the nonlinear suboptimal PLL for phase estimation. The practical benefits of each approach are discussed later in Section 4.6.

Assumption on Channel Fading Characteristics : Consider that the dynamics of x_k , from (4.6), are given by

$$x_{k+1} = F x_k + v_{k+1}, \quad v_k = N[0, Q_k], \quad (4.8)$$

for some known F , (usually with $\lambda(F) < 1$, where λ indicates eigen-values, to avoid unbounded x_k , and typically with $F = fI$ for some scalar $0 \ll f < 1$). In polar co-ordinates, (4.7), a corresponding model is

$$\begin{aligned} \kappa_{k+1} &= f_\kappa \kappa_k + v_{k+1}^\kappa, \quad \text{where } v_{k+1}^\kappa \text{ is Rayleigh distributed } [\mu_\kappa, \sigma_\kappa^2], \\ \phi_{k+1} &= f_\phi \phi_k + v_{k+1}^\phi, \quad \text{where } v_{k+1}^\phi \text{ is Uniformly distributed over } [0, 2\pi), \end{aligned}$$

and typically, $0 \ll f_\kappa < 1$ and $0 \ll f_\phi < 1$.

For both channel models it is assumed that the variation associated with the magnitude of the channel gain, κ , and the phase shift, ϕ , are independent, with variances given by σ_κ^2 and σ_ϕ^2 respectively. It follows from Anderson and Moore [1979](p.53), that the covariance matrix of the Cartesian channel model noise vector v_k , is given by

$$Q_k = E[v_k v_k'] \simeq \begin{bmatrix} \sigma_\kappa^2 \cos^2 \phi_k + \kappa_k^2 \sigma_\phi^2 \sin^2 \phi_k & (\sigma_\kappa^2 - \kappa_k^2 \sigma_\phi^2) \sin \phi_k \cos \phi_k \\ (\sigma_\kappa^2 - \kappa_k^2 \sigma_\phi^2) \sin \phi_k \cos \phi_k & \sigma_\kappa^2 \sin^2 \phi_k + \kappa_k^2 \sigma_\phi^2 \cos^2 \phi_k \end{bmatrix}. \quad (4.9)$$

For the remainder of this chapter the Cartesian channel model will be used, as it allows the system to be written in the familiar state space form.

4.2.3 Observation Model

The baseband output of the channel, corrupted by additive noise w_k , is therefore given by

$$y_k = g_k m_k + w_k. \quad (4.10)$$

Assume that $w_k \in \mathbf{C}$ has i.i.d. real and imaginary parts, w_k^R and w_k^I respectively, each with zero mean Gaussian distributions, so that $w_k^R, w_k^I \sim N[0, \sigma_w^2]$.

In vector notation the observations have the form:

$$\begin{pmatrix} y_k^R \\ y_k^I \end{pmatrix} = \begin{pmatrix} m_k^R & -m_k^I \\ m_k^I & m_k^R \end{pmatrix} \begin{pmatrix} g_k^R \\ g_k^I \end{pmatrix} + \begin{pmatrix} w_k^R \\ w_k^I \end{pmatrix}. \quad (4.11)$$

4.2.4 State Space Signal Model

Consider the following assumption on the message sequence.

Assumption on Message Signal

$$m_k \text{ is a first order homogeneous Markov process.} \quad (4.12)$$

Remark 4.1: This assumption enables the signal to be considered in a Markov model framework, and thus allows Markov filtering techniques to be applied. As discussed earlier, it is a reasonable assumption on the signal, given that error correcting coding has been employed in transmission. Of course un-coded i.i.d. signals can be considered in this framework too, since a Markov chain with a transition probability matrix which has all elements the same, gives rise to an i.i.d. process. \square

Let the vector X_k be an indicator function associated with m_k . The state space of X_k , without loss of generality, can be identified with the set of unit vectors $\mathbf{S} = \{e_1, e_2, \dots, e_{2^N}\}$, where $e_i = (0, \dots, 0, 1, 0, \dots, 0)' \in \mathbb{R}^{2^N}$ with 1 in the i^{th} position, so that

$$m_k = z' X_k, \quad (4.13)$$

where z is defined in (4.2). Under Assumption (4.12) the transition probability matrix associated with the message, m_k , in terms of X_k , is

$$\mathbf{A} = (a_{ij}), \quad 1 \leq i, j \leq 2^N, \quad \text{where } a_{ij} = P(X_{k+1} = e_j \mid X_k = e_i),$$

so that

$$E[X_{k+1} \mid X_k] = \mathbf{A}' X_k.$$

Of course $a_{ij} \geq 0$, $\sum_{j=1}^{2^N} a_{ij} = 1$, for each i . In addition, denote $\{\mathcal{F}_l, l \in \mathcal{Z}^+\}$ to be the complete filtration generated by X , that is, for any $k \in \mathcal{Z}^+$, \mathcal{F}_k is the complete σ -field generated by $X_k, l \leq k$.

From Lemma 2.1 it can be seen that the dynamics of X_k are given by the state equation

$$X_{k+1} = \mathbf{A}' X_k + M_{k+1}, \quad (4.14)$$

where M_{k+1} is a $(\mathbf{A}, \mathcal{F}_k)$ martingale increment.

As noted previously, in the case of quadrature modulated signals, the states represented by X_k are each characterised by a complex value, $z^{(i)}$, corresponding to the unit vector $e_i \in \mathbf{S}$. These are termed the state values of the Markov chain.

The *observation process* from (4.11), for the Cartesian channel model (4.6), can be expressed in terms of the state X_k as

$$\begin{pmatrix} y_k^R \\ y_k^I \end{pmatrix} = \begin{pmatrix} (z^R)'X_k & -(z^I)'X_k \\ (z^I)'X_k & (z^R)'X_k \end{pmatrix} \begin{pmatrix} g_k^R \\ g_k^I \end{pmatrix} + \begin{pmatrix} w_k^R \\ w_k^I \end{pmatrix}, \quad (4.15)$$

or equivalently, with the appropriate definition of $h(\cdot)$,

$$\mathbf{y}_k = h(X_k) x_k + \mathbf{w}_k, \quad \mathbf{w}_k \sim N[0, R_k]. \quad (4.16)$$

Note that, $E[w_{k+1}^R | \mathcal{F}_k \vee \mathcal{Y}_k] = 0$ and $E[w_{k+1}^I | \mathcal{F}_k \vee \mathcal{Y}_k] = 0$, where \mathcal{Y}_l is the σ -field generated by y_k , $k \leq l$. In addition, let $Y_k \triangleq (y_0 \dots y_k)$. It is usual to assume that w_k^R and w_k^I are independent so that the covariance matrix associated with the measurement noise vector \mathbf{w}_k has the form

$$R_k = \begin{bmatrix} \sigma_{w^R}^2 & 0 \\ 0 & \sigma_{w^I}^2 \end{bmatrix}. \quad (4.17)$$

Remark 4.2: In the previous chapters, the symbol R_k was used to denote the covariance matrix of the parameter estimate, as is traditional when considering RPE techniques. This chapter, however, considers Kalman filtering techniques. For such schemes it is usual to use the symbol R_k to denote the covariance matrix associated with the measurement noise, while using Σ_k for the covariance matrix of the state estimate. \square

It is now readily seen that

$$E[M_{k+1} | \mathcal{F}_k \vee \mathcal{Y}_k] = 0. \quad (4.18)$$

In order to demonstrate the attractiveness of the Cartesian channel model, the properties of the indicator function, X_k , are now used to express the observations (4.16) in a bi-linear form with respect to X_k and x_k .

$$\begin{aligned} \mathbf{y}_k &= h(X_k) x_k + \mathbf{w}_k \\ &= [h(e_1)x_k, h(e_2)x_k, \dots, h(e_{2N})x_k]X_k + \mathbf{w}_k \\ &= H' [I_{2N} \otimes x_k]X_k + \mathbf{w}_k, \end{aligned} \quad (4.19)$$

where $H' = [h(e_1), \dots, h(e_{2N})]$ and “ \otimes ” denotes a Kronecker product. The observations (4.19) are now in a form which is bi-linear in X_k and x_k .

It is possible to define the vector of parameterised probability densities (which will loosely be called symbol probabilities), as $\mathbf{b}_k = (b_k(1), \dots, b_k(2^N))'$, for $b_k(i) \triangleq P[y_k | X_k = e_i, x_k]$, where

$$b_k(i) = \frac{1}{2\pi\sigma_w^2} \exp \left(-\frac{\{y_k^R - [(z^R)'e_i g_k^R - (z^I)'e_i g_k^I]\}^2}{2\sigma_w^2} - \frac{\{y_k^I - [(z^I)'e_i g_k^R + (z^R)'e_i g_k^I]\}^2}{2\sigma_w^2} \right). \quad (4.20)$$

Here it is assumed that $\sigma_{w^R} = \sigma_{w^I} = \sigma_w$. Because w_k^R and w_k^I are independently distributed,

$$E(y_k | X_{k-1} = e_i, \mathcal{F}_{k-2}, \mathcal{Y}_{k-1}) = E(y_k | X_{k-1} = e_i), \quad (4.21)$$

which is essential for formulating the problem as an HMM, parameterised by the fading channel model parameter x_k .

To summarise, the following lemma is presented,

Lemma 4.1 *Under assumptions (4.12) and (4.8), the QAM signal model (4.1) to (4.10) has the following state space representation, in terms of the 2^N dimension finite-discrete state message indicator function, X_k , and the continuous state associated with the fading channel characteristics, x_k :*

$$\begin{array}{l} X_{k+1} = \mathbf{A}' X_k + M_{k+1} \\ x_{k+1} = F x_k + v_{k+1} \\ \mathbf{y}_k = H' [I_{2^N} \otimes x_k] X_k + \mathbf{w}_k \end{array} \quad (4.22)$$

■

Remarks 4.3: 1. If x_k is known, then the model specialises to an HMM denoted $\lambda = (\mathbf{A}, \mathbf{Z}, \underline{\pi}, \sigma_w^2, x_k)$, where $\underline{\pi} = (\pi_i)$, defined from $\pi_i = P(X_1 = e_i)$, is the initial state probability vector for the Markov chain.

2. If X_k is known then the model specialises to a linear state space model.

3. By way of comparison, for the polar co-ordinate channel representation (4.7), the observation process can only be expressed in terms of a linear operator on the channel gain, with a nonlinear

operator on the phase. Thus if X_k and ϕ_k , or κ_k and ϕ_k , are known then the model specialises to a linear state space model, but not if X_k and κ_k are known and ϕ_k is unknown.

4. In Figure 4.2 the output constellation is presented, with signal to noise ratio SNR = 6dB, from a channel with sinusoidal characteristics given by

$$\kappa_k = 1 + 0.5 \sin(3\pi k/1000)$$

$$\phi_k = 0.75\pi \cos(10\pi k/1000)$$

The plots show 1000 data points at each of the constellation points for times $k = 200$ and $k = 450$, and give an indication of how the channel affects the QAM signal constellation. \square




Figure 4.2: 16 state QAM Signal Constellation output from channel

4.2.5 Conditional Information-State Signal Model

Let $\hat{X}_{k|\mathcal{X}}$ denote the conditional filtered state estimate of X_k at time k , given the channel parameters $\mathcal{X}_k = \{x_0, \dots, x_k\}$, that is,

$$\hat{X}_{k|\mathcal{X}} \triangleq E[X_k | \mathcal{Y}_k, \mathcal{X}_k]. \quad (4.23)$$

Let $\mathbf{1}$ be the column vector containing all ones, and the information state, $\alpha_{k|\mathcal{X}} = (\alpha_{k|\mathcal{X}}(1), \dots, \alpha_{k|\mathcal{X}}(2^N))'$ be such that the i^{th} element

$$\alpha_{k|\mathcal{X}}(i) \triangleq P(Y_k, X_k = e_i | \mathcal{X}_k). \quad (4.24)$$

Observe that $\hat{X}_{k|\mathcal{X}}$ can be expressed in terms of $\alpha_{k|\mathcal{X}}$ by

$$\hat{X}_{k|\mathcal{X}} = \langle \alpha_{k|\mathcal{X}}, \mathbf{1} \rangle^{-1} \alpha_{k|\mathcal{X}}. \quad (4.25)$$

Here $\alpha_{k|\mathcal{X}}$ is computed using the following “forward” recursion:

$$\alpha_{k+1|\mathcal{X}} = \mathbf{B}(y_{k+1}, x_{k+1}) \mathbf{A}' \alpha_{k|\mathcal{X}}, \quad (4.26)$$

where $\mathbf{B}(y_{k+1}, x_{k+1}) = \text{diag}(b_{k+1}(1), \dots, b_{k+1}(2^N))$ and $b_k(i)$ is defined in (4.20).

The observations, y_k , are now expressed in terms of the un-normalised conditional information state, $\alpha_{k|\mathcal{X}}$.

Lemma 4.2 *The conditional measurements $\mathbf{y}_{k|\mathcal{X}}$ are defined by*

$$\mathbf{y}_{k|\mathcal{X}} = H' [I_{2^N} \otimes x_k] \langle \alpha_{k-1|\mathcal{X}}, \mathbf{1} \rangle^{-1} \mathbf{A}' \alpha_{k-1|\mathcal{X}} + n_{k|\mathcal{X}}, \quad (4.27)$$

where $\alpha_{k|\mathcal{X}}$ is defined in (4.24) and $n_{k|\mathcal{X}}$ is a $(\mathcal{X}_k, \mathcal{Y}_{k-1})$ martingale increment. In addition, the co-variance matrix of the conditional noise term $n_{k|\mathcal{X}}$ is given by

$$R_n = \sigma_w^2 I + H' [I_{2^N} \otimes x_k] \{ \hat{X}_{k|\mathcal{X}}^D - \hat{X}_{k|\mathcal{X}} \hat{X}_{k|\mathcal{X}}' \} [I_{2^N} \otimes x_k]' H, \quad (4.28)$$

where $\hat{X}_{k|\mathcal{X}}^D$ is the matrix which has \hat{X}_k on its diagonal, and all other elements zero.

Proof : Following standard arguments, since $\alpha_{k|\mathcal{X}}$ is measurable with respect to $\{\mathcal{X}_k, \mathcal{Y}_k\}$, $E[w_{k+1}^R | \mathcal{Y}_k] = 0$, $E[w_{k+1}^I | \mathcal{Y}_k] = 0$ and $E[M_{k+1} | \mathcal{Y}_k] = 0$, then

$$\begin{aligned} E[n_{k|\mathcal{X}} | \mathcal{X}_k, \mathcal{Y}_{k-1}] &= E[H' [I_{2^N} \otimes x_k] X_k + \mathbf{w}_k \\ &\quad - H' [I_{2^N} \otimes x_k] \langle \alpha_{k-1|\mathcal{X}}, \mathbf{1} \rangle^{-1} \mathbf{A}' \alpha_{k-1|\mathcal{X}} | \mathcal{X}_k, \mathcal{Y}_{k-1}] \\ &= H' [I_{2^N} \otimes x_k] \left(\mathbf{A}' \hat{X}_{k-1|\mathcal{X}} - \langle \alpha_{k-1|\mathcal{X}}, \mathbf{1} \rangle^{-1} \mathbf{A}' \alpha_{k-1|\mathcal{X}} \right) = 0. \end{aligned}$$

Also,

$$\begin{aligned} R_n &= E[n_k^2 | \mathcal{X}_k, \mathcal{Y}_{k-1}] \\ &= E[(w_k + H' [I_{2^N} \otimes x_k] (X_k - \frac{\mathbf{A}' \alpha_{k-1|\mathcal{X}}}{\langle \alpha_{k-1|\mathcal{X}}, \mathbf{1} \rangle}))^2 | \mathcal{X}_k, \mathcal{Y}_{k-1}] \\ &= E[w_k^2 | \mathcal{X}_k, \mathcal{Y}_{k-1}] + \\ &\quad E[H' [I_{2^N} \otimes x_k] (X_k - \hat{X}_{k|\mathcal{X}}) (X_k - \hat{X}_{k|\mathcal{X}})' [I_{2^N} \otimes x_k]' H | \mathcal{X}_k, \mathcal{Y}_{k-1}] \\ &= \sigma_w^2 I + H' [I_{2^N} \otimes x_k] E[(X_k - \hat{X}_{k|\mathcal{X}}) (X_k - \hat{X}_{k|\mathcal{X}})' | \mathcal{X}_k, \mathcal{Y}_{k-1}] [I_{2^N} \otimes x_k]' H \\ &= \sigma_w^2 I + H' [I_{2^N} \otimes x_k] \{ \hat{X}_{k|\mathcal{X}}^D - \hat{X}_{k|\mathcal{X}} \hat{X}_{k|\mathcal{X}}' \} [I_{2^N} \otimes x_k]' H. \end{aligned}$$

■

In summary, the following lemma is presented,

Lemma 4.3 *The state space representation (4.22) can be reformulated to give the following conditional information state signal model, with states $\alpha_{k|\mathcal{X}}$,*

$$\begin{aligned} \alpha_{k+1|\mathcal{X}} &= \mathbf{B}(y_{k+1}, x_{k+1})\mathbf{A}'\alpha_{k|\mathcal{X}} \\ x_{k+1} &= Fx_k + v_k \\ \mathbf{y}_{k|\mathcal{X}} &= H'[I_{2N} \otimes x_k] \langle \alpha_{k-1|\mathcal{X}}, \mathbf{1} \rangle^{-1} \mathbf{A}'\alpha_{k-1|\mathcal{X}} + n_{k|\mathcal{X}} \end{aligned} \quad (4.29)$$

■

Remarks 4.4: 1. When $F \equiv I$ and $v \equiv 0$, then x_k is constant. Under these conditions, the problem of channel state estimation reduces to one of parameter identification, and recursive prediction error techniques can be used, as in Chapter 2. However, when x_k is not constant, an EKF or some derivative scheme is required for parameter tracking, as in the following section.

2. By way of comparison, for the polar co-ordinate channel representation (for which the observations are nonlinear in terms of the channel phase parameter), it is possible to consider the case when ϕ_k is quantized into a discrete set of values, and is assumed to be Markov with indicator function $X_k^\phi \in \{e_1, e_2, \dots\}$. A conditional filtered estimate of X_k^ϕ can then be generated by the same means as used for the conditional filtered estimate of X_k , (4.25). The reformulated information state signal model, for the polar channel model case, is now given, in obvious notation by

$$\begin{aligned} \alpha_{k+1} &= \mathbf{B}^P(y_{k+1}, \kappa_{k+1}, \alpha_{k+1}^\phi)\mathbf{A}'\alpha_k, \\ \alpha_{k+1}^\phi &= \mathbf{B}^\phi(y_{k+1}, \kappa_{k+1}, \alpha_{k+1})(\mathbf{A}^\phi)'\alpha_k^\phi, \\ \kappa_{k+1} &= f_\kappa\kappa_k + v_\kappa^k, \\ \mathbf{y}_k &= H'_P(\langle \alpha_{k-1}, \mathbf{1} \rangle^{-1} \mathbf{A}'\alpha_{k-1})(\kappa_k) H'_\phi(\langle \alpha_{k-1}^\phi, \mathbf{1} \rangle^{-1} (\mathbf{A}^\phi)'\alpha_{k-1}^\phi) + n_k, \end{aligned} \quad (4.30)$$

where $H'_P = [h_P(e_1) \dots h_P(e_{2N})]$, $H'_\phi = [h_\phi(e_1) \dots h_\phi(e_{L_\phi})]$, and

$$\begin{aligned} h_P(\cdot) &= \langle z_\rho, \cdot \rangle \exp[\mathbf{j}\langle z_\Upsilon, \cdot \rangle], \\ h_\phi(\cdot) &= \exp[\mathbf{j}\langle z_\phi, \cdot \rangle], \end{aligned} \quad (4.31)$$

where z_ϕ is the vector containing the discrete values of ϕ , z_ρ and z_γ are vectors containing the magnitudes and phases respectively of the QAM signal constellation, L_ϕ is the number of discrete values of ϕ , and $X_k^\phi \in \mathcal{S} = \{e_1, \dots, e_{L_\phi}\}$ is the indicator function associated with ϕ_k so that when $\phi_k = z_\phi^{(i)}$, $X_k^\phi = e_i$. Also, $\mathbf{B}^P(y_{k+1}, \kappa_{k+1}, \alpha_{k+1}^\phi) = \text{diag}(b_{k+1}^P(1), \dots, b_{k+1}^P(2^N))$, where $b_{k+1}^P(i) \triangleq b^P(y_{k+1}, e_i, \kappa_{k+1}, \alpha_{k+1}^\phi)$, and

$$b_k^P(i) = \frac{1}{2\pi\sigma_w^2} \exp\left(\frac{-\{y_k - h_P(e_i)(\kappa_k) H'_\phi \alpha_k^\phi\}^2}{2\sigma_w^2}\right). \quad (4.32)$$

In addition, the matrix $\mathbf{B}^\phi(y_{k+1}, \alpha_{k+1}, \kappa_{k+1}) = \text{diag}(b_{k+1}^\phi(1), \dots, b_{k+1}^\phi(L_\phi))$, is defined where $b_{k+1}^\phi(i) \triangleq b^\phi(y_{k+1}, e_i, \alpha_{k+1}, \kappa_{k+1})$, and

$$b_k^\phi(i) = \frac{1}{2\pi\sigma_w^2} \exp\left(\frac{-\{y_k - H'_P \alpha_k(\kappa_k) h_\phi(e_i)\}^2}{2\sigma_w^2}\right). \quad (4.33)$$

The key property which facilitates estimator construction is that now, with the quantised discrete-state assumption on ϕ , in this polar co-ordinate representation for the channel, the measurements are tri-linear in X_k , κ and X_k^ϕ . Problems arise, however, in the choice of \mathbf{A}^ϕ associated with the phase quantisation, with the quantisation error itself, and with the increased computational requirement resulting from the need to evaluate the recursion for α_k^ϕ . \square

4.3 Adaptive HMM Algorithms

Two adaptive HMM schemes are presented here, the first is referred to as an HMM/EKF scheme, and is a full nonlinear scheme for the information state signal model (4.29) with the augmented vector $(\alpha_k, x_k)'$. The second scheme is referred to as the HMM/KF scheme, and is with a simplification assumption which results in a KF for channel estimation, coupled with an HMM filter for signal state estimation.

4.3.1 Adaptive HMM/EKF Scheme

Let $\underline{x}_k = (\alpha_k, x_k)'$, then (4.29) can be written as

$$\underline{x}_{k+1} = f_k(\underline{x}_k) + g_k(\underline{x}_k)v_k, \quad (4.34a)$$

$$y_k = h_k(\underline{x}_k) + n_k, \quad (4.34b)$$

where the nonlinear functions are given by

$$f_k(\underline{x}_k) = \begin{pmatrix} \mathbf{B}(x_k)\mathbf{A}'\alpha_k \\ Fx_k \end{pmatrix}, \quad g_k(\underline{x}_k) = \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

$$h_k(\underline{x}_k) = h(\alpha_{k|k-1})x_k.$$

Assuming that the nonlinearities are smooth functions, they can be expanded in Taylor series about the conditional means $\hat{\underline{x}}_{k|k}$ and $\hat{\underline{x}}_{k|k-1}$. With an assumption that higher order terms can be ignored, as would be the case when $\hat{\underline{x}}_{k|k}$ is close to \underline{x}_k , (4.34a) and (4.34b) can be written as follows:

$$\underline{x}_{k+1} = F_k \underline{x}_k + G_k v_k + u_k, \quad (4.35a)$$

$$y_k = H_k' \underline{x}_k + n_k + z_k, \quad (4.35b)$$

where $F_k = \partial f_k / \partial \underline{x}_k$, $G_k = \partial g_k / \partial \underline{x}_k$, $H_k = \partial h_k / \partial \underline{x}_k$, $u_k = f_k(\hat{\underline{x}}_{k|k}) - F_k \hat{\underline{x}}_{k|k}$ and $z_k = h_k(\hat{\underline{x}}_{k|k-1}) - H_k' \hat{\underline{x}}_{k|k-1}$.

The EKF equations for (4.34a) and (4.34b) are the KF equations for (4.35a) and (4.35b), now summarised:

$$\hat{\underline{x}}_{k|k} = \hat{\underline{x}}_{k|k-1} + K_k [y_k - h_k(\hat{\underline{x}}_{k|k-1})], \quad (4.36a)$$

$$\hat{\underline{x}}_{k+1|k} = f_k(\hat{\underline{x}}_{k|k}), \quad (4.36b)$$

$$K_k = \Sigma_{k|k-1} H_k [H_k' \Sigma_{k|k-1} H_k + R_k]^{-1}, \quad (4.36c)$$

$$\Sigma_{k|k} = \Sigma_{k|k-1} - \Sigma_{k|k-1} H_k [H_k' \Sigma_{k|k-1} H_k + R_k]^{-1} H_k' \Sigma_{k|k-1}, \quad (4.36d)$$

$$\Sigma_{k+1|k} = F_k \Sigma_{k|k} F_k' + G_k Q_k G_k', \quad (4.36e)$$

where (4.36c) gives the Kalman gain and (4.36d) and (4.36e) are the Riccati equations. Figure 4.3 gives a block diagram for this adaptive HMM scheme, when switch 1 is closed and switch 2 is in the top position. If switch 1 was in the open position then the HMM/KF scheme given below would result. Further assumptions can be made for simplification if the maximum *a posteriori* estimate of α_k were used, indicated by having switch 2 in the lower position. This approach would be similar to using the matched filter, where only the most likely message symbol is used, and not the full information state.

Remarks 4.5: 1. This HMM/EKF scheme suffers from the fact that through (4.36a), the update for α_k requires a further projection to ensure positivity of each element. This adds undesired nonlinearities to the model and provides further incentive to consider the HMM/KF scheme presented below, where this problem does not arise.

2. The filter here is in fact a smoothed filter in the sense that $f_k(\underline{x}_k)$ is actually $f_{k+1}(\underline{x}_k)$ due to the dependence of $\mathbf{B}(\underline{x}_k)$ on y_{k+1} . This again provides incentive to consider the HMM/KF scheme presented below, where this problem does not arise. □

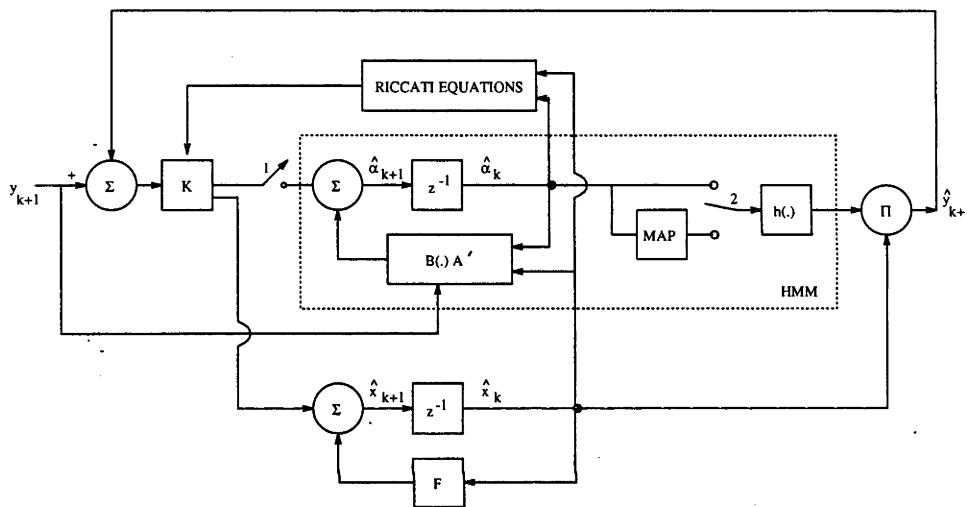


Figure 4.3: Adaptive HMM/EKF scheme

4.3.2 Adaptive HMM/KF Schemes

This scheme can be viewed as a derivative of the above HMM/EKF scheme by setting the Kalman gain term, associated with the $\hat{\alpha}_k$ update, to zero. The rationale for this is that in the case where the channel parameters are constant, this term in fact does go to zero asymptotically. Indeed setting it to zero under constant parameter conditions leads to the RPE scheme as used in Chapter 2, for which there are strong theoretical foundations. If the channel is only slowly varying, then it is expected that the components of the Kalman gain associated with the $\hat{\alpha}_k$ update, will be asymptotically small. There is then a temptation and some rational, to neglect these terms for the simplicity of the resulting scheme, which is now described in more detail.

The HMM estimator for the signal information state, α_k , conditioned on the channel estimate sequence $\{\hat{x}_k\}$, is given by

$$\hat{\alpha}_{k+1|\hat{x}_k} = \mathbf{B}(y_{k+1}, \hat{x}_k) \mathbf{A}' \hat{\alpha}_{k|\hat{x}_{k-1}}, \quad (4.37a)$$

$$\hat{X}_{k|\hat{x}_{k-1}} = \langle \alpha_{k|\hat{x}_{k-1}}, \mathbf{1} \rangle^{-1} \alpha_{k|\hat{x}_{k-1}}. \quad (4.37b)$$

The Kalman filter equations for the channel parameter, x_k , conditioned on the indicator state estimates \hat{X}_k , are

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k [y_k - H'_k \hat{x}_{k|k-1}], \quad (4.38a)$$

$$\hat{x}_{k+1|k} = F \hat{x}_{k|k}, \quad (4.38b)$$

$$K_k = \Sigma_{k|k-1} H_k [H'_k \Sigma_{k|k-1} H_k + R_k]^{-1}, \quad (4.38c)$$

$$\Sigma_{k|k} = \Sigma_{k|k-1} - \Sigma_{k|k-1} H_k [H'_k \Sigma_{k|k-1} H_k + R_k]^{-1} H'_k \Sigma_{k|k-1}, \quad (4.38d)$$

$$\Sigma_{k+1|k} = F \Sigma_{k|k} F' + Q_k, \quad (4.38e)$$

where

$$H'_k = \partial(H' [I_{2N} \otimes x_k] \hat{X}_k) / \partial x_k, \quad (4.39)$$

and R is the covariance matrix of the noise on the observations w_k given in (4.17), Q is the covariance matrix of v_k , given in (4.9), and Σ is the covariance matrix of the channel parameter estimate \hat{x}_k , (x_k is defined in (4.6)). Figure 4.4 shows the scheme in block form.

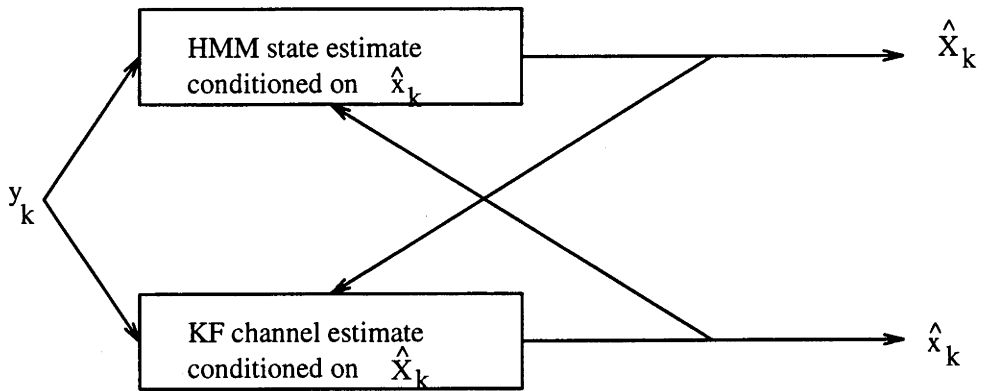


Figure 4.4: Adaptive HMM/KF scheme

Remarks 4.6: 1. A further sub-optimal HMM/KF scheme can be generated by conditioning the KF on a maximum *a posteriori* probability estimate of the message state, \hat{X}_k^{MAP} . Here

$$H'_k = \partial(H'[I_{2N} \otimes x_k] \hat{X}_k^{MAP}) / \partial x_k. \quad (4.40)$$

Figure 4.5 shows this scheme in block form. In fact hybrid versions can be derived by setting the

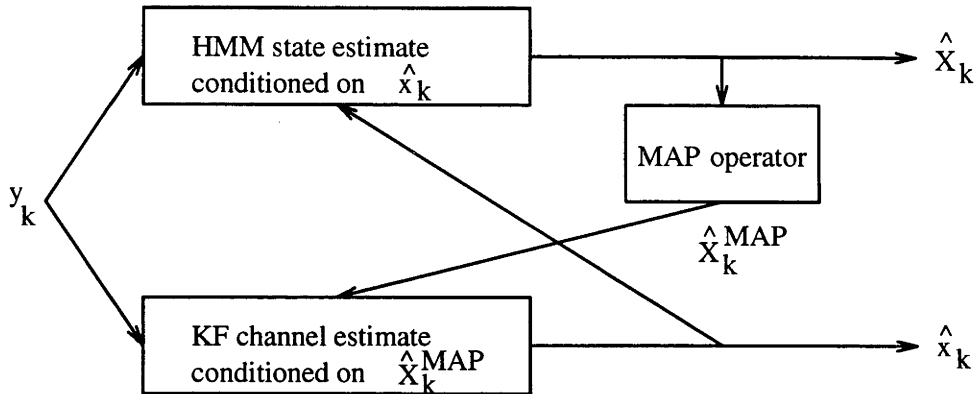


Figure 4.5: Adaptive HMM/KF scheme with MAP approximation

small valued, that is low probability, elements of \hat{X}_k to zero and re-normalising.

2. Extensions to this simultaneous coupled state estimation and channel tracking technique, to sequence estimation techniques are straight forward. In a batch processing situation the off-line HMM schemes of Rabiner [1989] could be used to gain estimates of the sequence from forward and backward passes through the data. In an on-line delayed sequence estimation technique, the smoothing approach of (2.49) and (2.50) would be used to estimate a sequence of length Δ , and then the new information state α_k^Δ could be used in place of α_k for conditioning of the channel estimate. \square

4.4 Coloured Noise Case

In the coloured noise case, it is reasonable to work with the following signal model involving a moving average of white noise as follows:

$$y_k = h(\cdot)x_k + \mathbf{w}_k + c_1\mathbf{w}_{k-1} + \dots + c_n\mathbf{w}_{k-n}. \quad (4.41)$$

The task of estimating the noise coefficients, c_i , is carried out by augmenting the state vector x_k by a vector $x_k^w = (w_{k-1}, w_{k-2}, \dots)'$. An example is given for the case $n = 3$. The state vector x_k^w is the vector of noise values, $x_k^w = (w_{k-1}, w_{k-2}, w_{k-3})'$, and the vector of noise coefficients is $\theta = (c_1, c_2, c_3)'$.

$$\begin{pmatrix} x_{k+1} \\ x_{k+1}^w \end{pmatrix} = \begin{pmatrix} F & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_k \\ x_k^w \end{pmatrix} + \begin{pmatrix} v_k \\ w_k \\ 0 \\ 0 \end{pmatrix}, \quad (4.42a)$$

$$y_k = [h(\cdot) \ \theta'] \begin{pmatrix} x_k \\ x_k^w \end{pmatrix} + w_k. \quad (4.42b)$$

The ascribed estimation task can now be solved with an EKF, or derivative KF, where the state vector is now the augmented vector, (α_k, x_k, x_k^w) . If θ is unknown, it can be adaptively estimated using standard RPE/EKF ideas.

4.5 Robustness Issues

Due to the inherently sub-optimal nature of these adaptive HMM algorithms, it is necessary to consider robustness issues. In the case of the HMM/EKF scheme above, some robustness is gained by the fact that the Kalman gain term acts on the whole augmented vector \hat{x}_k . A problem however is that the dimension of \hat{x}_k is too large for practical implementation. The HMM/KF schemes presented in Section 4.3.2, which are practical derivatives of the HMM/EKF scheme, effectively set the Kalman gain terms associated with the respective α 's, to zero. This results in coupled conditional estimators, which are used to condition the other estimates. Unfortunately, however, there is no theory for convergence for these coupled schemes, when dealing with time varying parameters. The validity of such an approach can only be tested via simulation studies.

It should be noted that, through the use of information-state techniques, some degree of robustness is inherent in these schemes, when compared to standard MAP estimate approaches. This is due to the fact that, instead of only feeding the most likely signal state into the channel tracker, with information-state techniques, the full state probability distribution is fed in. This effectively means that a measure of the error involved in the state estimate is taken into account, thus making the scheme more robust.

In an effort to address further the inevitable robustness question, it is possible to look to the standard procedures from Kalman filtering. A widely practised method for adding robustness, is to model the estimate errors, due to incorrect conditioning, as noise in the observations. This procedure can also be used with the adaptive HMM techniques presented in this chapter. By adding extra noise to the observation model, the vector of parameterised probability densities (symbol probabilities) will be more uniform. That is, the diagonal “observation” update matrix, $B(\cdot)$, in the “forward” procedure, (4.26), for the information state α_k , will place less emphasis on the observations. An additional method for adding robustness to the adaptive HMM scheme, is to assume the probability of remaining in the same state is higher than it actually is. That is, by using a more diagonally dominant transition probability matrix A . This will also have the effect of placing less importance on the observations, through the “forward” procedure for the discrete state estimate α_k .

These robustness techniques are of course an attempt to counter estimation errors in high noise. They therefore restrict the ability of the estimates to track quickly varying parameters, as the rapid changes will effectively be modelled as noise. There is here, as in all cases, a trade off to be made between robustness and adaptive tracking ability.

4.6 Implementation Considerations and Simulations

In this section results are presented which demonstrate the ability of the adaptive HMM/KF scheme to demodulate QAM signals in noisy fading channels. For comparisons, the standard MF/AGC/PLL scheme, which is diagrammatically represented in Figure 4.6, is used (similar to the LMS algorithm presented in Pahlavan and Matthews [1990]). Viterbi schemes could also be used, however, they require some degree of off-line, or delayed, sequence estimation. An advantage of the approach in this chapter is that no delay is required.

The signal considered is a 16 state QAM signal with a strong dependence from one message symbol to the next, (as is the case with some convolutional codes, or if oversampling were to be used). The channel characteristics for Example 4.1 are deterministic while those for Examples 4.2 to 4.5 are given by a low pass filtered (LPF) white Gaussian noise stochastic process. The variance of the Gaussian process is 1 for amplitude variations, and 5 for phase variations. The bandwidth of the LPF is W_c times the bit rate (W_c is different in each example). An example of a real valued stochastic channel is shown in Figure 4.9 for $W_c = 0.1$. These variations are very fast in the

case of FAX and modem applications, but are more reasonable in applications involving mobile communications and indoor communication channels [Hashemi 1993, Loo and Secord 1991].

Two main points can be gained from the following examples, the first is that under these non-equally-probable message symbol conditions, the HMM filter is a major improvement over the MF, the second point is that the Cartesian and polar co-ordinate systems can each have their advantages, depending on the channel conditions. Computationally, the MF is of course less taxing, however for mobile communications under the conditions (16-QAM, 19.2 kB/sec, $f_c=1800$ MHz, car travelling at 100 km/h and with one channel update every 120 samples), the processing power required for the HMM/KF approach is only 10 MFlops, which is reasonable with current DSP technology. Therefore the approach presented in this chapter is computationally feasible, and is seen to out-perform the traditional scheme for the case of non-equally-probable messages. In the equally-probable message case, when a MAP operator is used in the HMM filter, as in Figure 4.5, the HMM approach is identical to the traditional MF scheme. There is still an advantage to the HMM approach using the full information state, because more information is being fed back to the Kalman filter or AGC/PLL combination.

Example 4.1: A 16 state QAM signal was generated under assumption (4.12) with parameter values $a_{ii} = 0.95$, $a_{ij} = (1 - a_{ii})/(N - 1)$ for $i \neq j$, $(z^{(i)})^R = \pm 0.01976 \pm 0.03952$, $(z^{(i)})^I = \pm 0.01976 \pm 0.03952$. The channel characteristics used were deterministic, as opposed to the stochastic low pass filtered white noise channels used in the following examples. The deterministic channel gives a more easily repeatable test, and allows results to be displayed in a manner which more clearly shows tracking ability of these schemes. The channel characteristics used in this example were given by

$$\begin{aligned}\kappa_k &= 1 + 0.5 \sin(3\pi k/1000), \\ \phi_k &= 0.75\pi \cos(10\pi k/1000),\end{aligned}$$

and the signal to noise ratio (SNR) associated with the observations, in the absence of fading, is $\text{SNR} = (E_b/\sigma_w^2) = 16\text{dB}$, where E_b is the energy per bit associated with the transmitted signal. Of course much lower SNRs can be accommodated in the presence of more slowly varying channels, and it should be noted that the SNR effectively varies as the channel fades. The lowest effective SNR in this example occurs at $k = 500$ where $\text{SNR} = 10\text{dB}$. This example is used to demonstrate the HMM/KF scheme of Figure 4.4. The results are presented in Figure 4.7 and Figure 4.8, and show that even though the channel changes quite quickly, good estimates are generated. Figure 4.7 shows the true channel values and the estimated values in real and imaginary format, that is,

exactly as estimated from (4.38a) to (4.38e). Figure 4.8(a) shows the actual channel amplitude κ_k , and the estimate of this, generated from the estimates in Figure 4.7. Likewise, Figure 4.8(b) shows the actual channel phase shift ϕ_k and the estimate of this, generated from the estimates in Figure 4.7. Small glitches can be seen in the amplitude and phase estimates at points where tracking was slow and the received channel amplitude was low, but the recovery after this burst of errors seems to be quite good. It is natural that the estimates during these periods be worse, since the noise on the observations is effectively greater when $\kappa_k < 1$, as seen from the signal model (4.22).

Example 4.2: This example demonstrates the ability of the HMM/KF adaptive algorithm to demodulate a 16-QAM signal, in the presence of a real valued stochastic channel. The signal parameter values are given in Example 4.1. The results for this example are displayed in Figure 4.10, where signals of length 50000 data points have been used to generate bit-error-rate (BER) values. The simulations assume that 90 degree phase invariant coding is used. A comparison is given to the conventional MF/AGC/PLL system (of course the PLL is not required since the channel is real valued). It can be seen that the HMM/KF scheme provides distinct advantages over the traditional scheme. As noted before, the case of $W_c = 0.1$ is one of severe fading, and it is seen that even under such conditions, the HMM/KF scheme performs well.

Example 4.3: In this example, it is demonstrated that the HMM approach with MAP operator, is identical to the MF approach in the case of equally-probable message signals. The discrepancies which can be seen between the two schemes, in the results of Figure 4.11, are due to the Cartesian approach compared to the polar approach of the traditional scheme. It seems that under these channel conditions, in the high SNR case, the polar approach is better than the Cartesian approach. Such a comparison is the subject of the next two examples.

Example 4.4: This example demonstrates the ability of the HMM/KF adaptive algorithm to demodulate a 16-QAM signal, in the presence of a complex valued stochastic channel. The signal characteristics are the same as for Example 4.1. Results for this example are displayed in Figure 4.12. It can be seen again that the HMM approach has significantly better performance than the traditional scheme involving the MF. Here the results for the adaptive HMM approach, when formulated in the polar representation, are also presented. For this case an AGC/PLL scheme for the channel parameter was implemented (note that this is not the quantised approach presented in Remark 4.4). It can be seen that under these conditions, the nonlinearities in the PLL approach are not detrimental, and in fact the HMM/AGC/PLL approach performs better than the HMM/KF scheme.

Example 4.5: This example investigates the relative benefits of the Cartesian channel parameterisation versus the polar representation. The signal characteristics are the same as for Example 4.1 and the results are displayed in Figure 4.13. In the previous example, the channel phase shift varied more slowly than in this example. It can be seen that under the more stringent conditions presented here, the nonlinearities in the PLL approach are detrimental, and the HMM/KF approach performs better than the HMM/AGC/PLL scheme.

From these examples it can easily be seen that the HMM approach is more suited, than the MF, to signals with non-equally-probable message symbols. Also, depending on the channel characteristics, the Cartesian co-ordinate representation can provide improvements over the traditional polar representation. Such improvements are most apparent in conditions of rapidly varying phase where the nonlinearities associated with the PLL are detrimental to performance.

4.7 Conclusions

In this chapter derivations have been given for adaptive HMM on-line state and parameter estimation schemes for QAM signals in fading communications channels. A key element of the HMM approach, which appears to be quite powerful, is to work with mixed finite-discrete and continuous range state models. These are reformulated via HMM filtering theory as conditional information state models. The resulting adaptive algorithms blend EKF and HMM techniques. They are based on optimal techniques, but are inevitably sub-optimal. Simulation studies are presented which show the ability to effectively track time-varying channel parameters for QAM signals, and demonstrate advantages over traditional approaches to the fading channel problem.

4.8 Figures

The figures for this chapter are now presented.

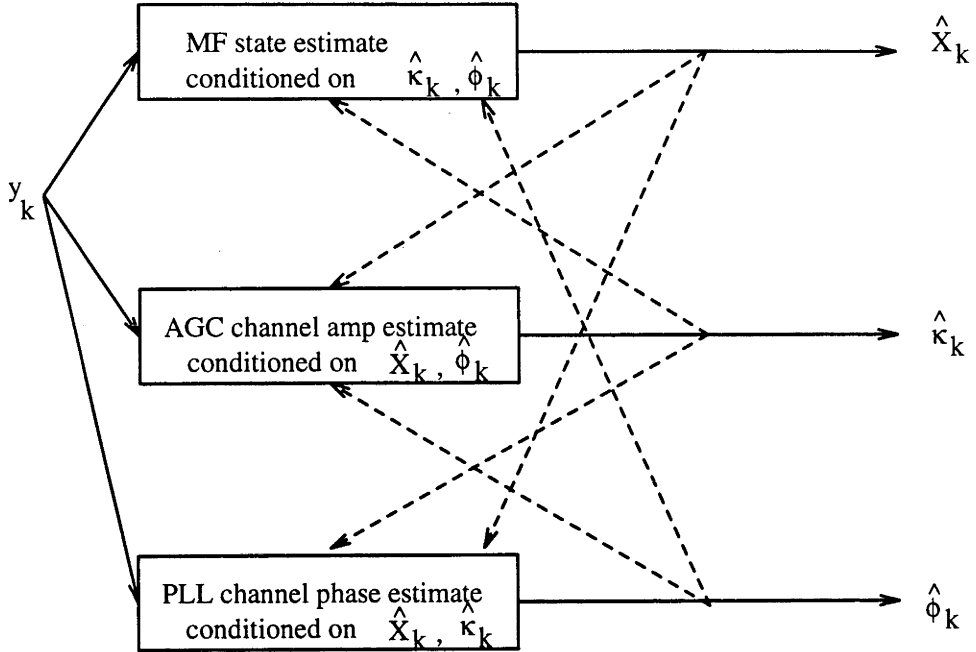


Figure 4.6: Standard MF/AGC/PLL scheme




Figure 4.7: Real and imaginary channel gain for 16-QAM




Figure 4.8: Channel amplitude gain (κ) and phase rotation (ϕ_k) estimates




Figure 4.9: Stochastic channel gain, $W_c = 0.1$




Figure 4.10: BER v SNR for real valued channels




Figure 4.11: BER v SNR for complex valued channels with equally-probable message symbols




Figure 4.12: BER v SNR for complex valued channels




Figure 4.13: BER v SNR for complex valued channels

Chapter 5

HMM Processing for Differentially Phase Modulated Communication Systems

5.1 Introduction

This chapter addresses the problem of demodulating differentially phase modulated signals under conditions of fading transmission channels. The HMM approach presented, is similar to that of the previous chapter, however some significant extensions to the signal model are required in this case.

When addressing the problem of frequency, or differentially phase, modulated signal demodulation in noisy fading channels, it is of interest to note the historical role of the EKF. EKF techniques have previously been applied to problems of continuous time frequency modulated (FM) signal demodulation (for example, in Anderson and Moore [1979]). In these cases, however, the channel is considered to be known and constant. Under such conditions the EKF turns out to be a PLL in disguise, where the Riccati equations are decoupled from the state estimate and can therefore be solved off-line. The case of fading channels is however more difficult.

The specific signal models considered in this chapter are digital M-ary differential phase shift keyed (MDPSK) signals, and analog frequency modulated (FM) signals. As in the previous chapter, the approach taken is to couple EKF and HMM filtering techniques. The finite-discrete nature of the Markov model is ideal for modelling the digital, or quantised analog, signal while the EKF is used to estimate the continuously varying complex channel. As in the previous chapter, an advantage is gained by using the full probability distributions for the states, obtained from the HMM filter, instead of more traditional MAP estimate techniques.

Unfortunately, in the analog FM case, the discrete-state nature of the formulation implies quantisation errors. However, if the quantisation is fine, and sampling rate high, then these errors can be minimised such that the processing gains, due to optimal HMM filtering, will overcome the quantisation disadvantages. In addition to DPSK and FM signals, the HMM schemes presented in this chapter have direct application to digital frequency shift keyed (FSK) and phase shift keyed (PSK) signals. In fact the PSK model has the same form as the quadrature amplitude modulation (QAM) model, which was the subject of Chapter 4.

This chapter is organised as follows : In Section 5.2 the state space signal model for the MDPSK and FM signals is formulated. Section 5.4 reviews Kalman filtering and presents estimation objectives. In Section 5.5 simulation examples are given which demonstrate good tracking ability for fast changing channels. Finally, some conclusions are presented in Section 5.6.

5.2 MDPSK and FM Signal Models

Digital M-ary differential phase shift keyed (MDPSK) and analog frequency modulated (FM) signals are common methods for information transmission. Such signals carry the information message in the frequency component of the signal. This section first presents the usual MDPSK and FM signal models and then proposes a reformulation so as to apply hidden Markov model (HMM) and extended Kalman filtering (EKF) techniques.

5.2.1 Initial Signal Model Formulation

Let f_k be a real valued discrete-time signal, where for each k ,

$$f_k \in \mathbf{Z}_f = \{z_f^{(1)}, \dots, z_f^{(L_f)}\}, \quad z_f^{(i)} = (i/L_f)\pi \in \mathbf{R}, \quad L_f \in \mathcal{Z}^+. \quad (5.1)$$

Also, let

$$z_f = (z_f^{(1)}, \dots, z_f^{(L_f)})' \in \mathbb{R}^{L_f}. \quad (5.2)$$

Therefore each of the $L_f \in \mathcal{Z}^+$ elements of \mathbf{Z}_f is an *equally spaced* frequency in the range $[0, \pi)$. Now note that at any time, k , the message, $f_k \in \mathbf{Z}_f$, is real valued, and can be used to generate piece-wise constant time signals, $f(t) = f_k$ for $t = [t_k, t_{k+1})$. For transmission, the instantaneous frequency is varied linearly with the baseband signal $f(t)$, giving the following transmitted signal.

$$s(t) = A_c \cos [2\pi f_c t + \theta(t)] \quad , \quad \theta(t) = \int_0^t f(\tau) d\tau, \quad (5.3)$$

where the carrier amplitude A_c , and frequency f_c are constant, and $\theta(t)$ is the phase of the signal. For the formulation which follows, it is convenient to represent the signal in customary complex baseband notation, relevant when the signal is sampled in a quadrature and in-phase manner.

$$s(t) = A_c \exp[\mathbf{j}\theta(t)] \quad , \quad s_k = A_c \exp[\mathbf{j}\theta_k]. \quad (5.4)$$

Here the amplitude, A_c , is a known constant, and

$$\theta_k = (\theta_{k-1} + f_k)_{2\pi}, \quad (5.5)$$

where $(\cdot)_{2\pi}$ denotes modulo 2π addition.

5.2.2 Channel Model

The channel model used in this chapter is the same as that presented in the previous chapter (see Section 4.2.2).

Channel State - Cartesian Co-ordinate Representation : In this co-ordinate system, it is possible to work with the vector x_k associated with the real and imaginary parts of g_k .

$$x_k = \begin{pmatrix} \kappa_k \cos \phi_k \\ \kappa_k \sin \phi_k \end{pmatrix} = \begin{pmatrix} g_k^R \\ g_k^I \end{pmatrix}. \quad (5.6)$$

The dynamics of x_k are given in (4.8).

5.2.3 Observation Model

The baseband output of the channel, corrupted by additive noise w_k , is given in discrete-time, by

$$y_k = g_k s_k + w_k, \quad (5.7)$$

where g_k is defined in (4.5). Assume that $w_k \in \mathbf{C}$ has i.i.d. real and imaginary parts, w_k^R and w_k^I respectively, each having a zero mean Gaussian distribution, so that $w_k^R, w_k^I \sim N[0, \sigma_w^2]$.

Reformulating (5.7) in matrix notation gives

$$\begin{pmatrix} y_k^R \\ y_k^I \end{pmatrix} = \begin{pmatrix} A_c \cos \theta_k & -A_c \sin \theta_k \\ A_c \sin \theta_k & A_c \cos \theta_k \end{pmatrix} \begin{pmatrix} g_k^R \\ g_k^I \end{pmatrix} + \begin{pmatrix} w_k^R \\ w_k^I \end{pmatrix}. \quad (5.8)$$

5.2.4 State Space Signal Model

For the MDPSK and FM signals, a discrete-time state space signal model is now generated. Consider the following assumption on the message signal.

Assumption on Message Signal

$$f_k \text{ is a first order homogeneous Markov process.} \quad (5.9)$$

Remarks 5.1: 1. This assumption enables the signal to be considered in a Markov framework, and thus allows Markov filtering techniques to be applied. It is a reasonable assumption on the FM signal if the transition probability matrix is chosen to be diagonally dominated, Toeplitz, and circulant. In the case of digital MDPSK signals, the assumption is valid, given that error correcting coding has been employed in transmission.

2. Higher order message signal models are discussed later in Section 5.3. □

Let the vector X_k^f be an indicator function associated with f_k . The state space of X^f , without loss of generality, can be identified with the set of unit vectors $\mathbf{S} = \{e_1^f, e_2^f, \dots, e_{L_f}^f\}$, where $e_i^f = (0, \dots, 0, 1, 0, \dots, 0)' \in \mathbf{R}^{L_f}$ with 1 in the i^{th} position, so that

$$f_k = z_f' X_k^f. \quad (5.10)$$

Under Assumption (5.9) the transition probability matrix associated with f_k , in terms of X_k^f , is

$$\mathbf{A}^f = (a_{ij}^f), \quad 1 \leq i, j \leq L_f, \quad \text{where } a_{ij}^f = P(X_{k+1}^f = e_j^f \mid X_k^f = e_i^f),$$

so that

$$E[X_{k+1}^f \mid X_k^f] = \mathbf{A}' X_k^f,$$

Of course $a_{ij}^f \geq 0$, $\sum_{j=1}^{L_f} a_{ij}^f = 1$, for each i . In addition, denote $\{\mathcal{F}_l, l \in \mathcal{Z}^+\}$ to be the complete filtration generated by X^f , that is, for any $k \in \mathcal{Z}^+$, \mathcal{F}_k is the complete σ -field generated by $X_k^f, l \leq k$.

From Lemma 2.1 it can be seen that the dynamics of X_k^f are given by the state equation

$$X_{k+1}^f = (\mathbf{A}^f)' X_k^f + M_{k+1}^f. \quad (5.11)$$

As noted previously, the states represented by X^f are each characterised by a real value, $z_f^{(i)}$, corresponding to the unit vector $e_i^f \in \mathbf{S}$. These values are termed state values of the Markov chain.

Now, let the set of discrete phase values be given by:

$$\mathbf{Z}_\theta = \{z_\theta^{(1)}, \dots, z_\theta^{(L_\theta)}\}, \quad \text{where } z_\theta^{(i)} = 2\pi i/L_\theta \in \mathbf{R}, \quad (5.12)$$

and define the corresponding vector

$$z_\theta = (z_\theta^{(1)}, \dots, z_\theta^{(L_\theta)})' \in \mathbf{R}^{L_\theta}. \quad (5.13)$$

Lemma 5.1 Consider the discrete-state message $f_k \in \mathbf{Z}_f$ from (5.1), the phase θ_k from (5.5) and the set \mathbf{Z}_θ of (5.12). Now, $\theta_k \in \mathbf{Z}_\theta \forall k$, iff $\theta_0 \in \mathbf{Z}_\theta$ and $L_\theta = 2nL_f$, for some $n \in \mathcal{Z}$

Proof: For any $a \in \{1, \dots, L_\theta\}$ and $b \in \{1, \dots, L_f\}$

$$\begin{aligned} \theta_{k+1} &= (\theta_k + f_k)2\pi \\ &= \left(a \frac{2\pi}{L_\theta} + b \frac{2\pi n}{L_\theta}\right)_{2\pi} \in \mathbf{Z}_\theta. \end{aligned}$$

■

It is now possible to define a vector $X_k^\theta \in \mathcal{S} = \{e_1^\theta, \dots, e_{L_\theta}^\theta\}$ to be an indicator function associated with θ_k , so that when $\theta_k = z_\theta^{(i)}$, $X_k^\theta = e_i^\theta$. Now given (5.5), it can be seen that X_{k+1}^θ is a “rotation” on X_k^θ by an amount determined from X_{k+1}^f . In particular,

$$X_{k+1}^\theta = [\mathbf{A}^\theta(X_{k+1}^f)]' X_k^\theta, \quad (5.14)$$

where $\mathbf{A}^\theta(\cdot)$ is a transition probability matrix given by

$$\mathbf{A}^\theta(X_{k+1}^f)' = S^{r_{k+1}}, \text{ where } r_k = [1, 2, \dots, L_f] X_k^f, \quad (5.15)$$

and S is the rotation operator

$$S = \begin{pmatrix} 0 & 0 & \dots & 0 & 1 \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{pmatrix}. \quad (5.16)$$

Now define $x_k^{f\theta}$ to be the state associated with the augmentation of f_k and θ_k ,

$$x_k^{f\theta} = \begin{pmatrix} f_k \\ \theta_k \end{pmatrix}, \quad (5.17)$$

and $X_k^{f\theta}$ to be the $L_f \times L_\theta$ dimension indicator function associated with $x_k^{f\theta}$,

$$X_k^{f\theta} = X_k^f \otimes X_k^\theta. \quad (5.18)$$

As before, the state space of $X_k^{f\theta}$ can be identified with the set of unit vectors $\mathcal{S} = \{e_1^{f\theta}, \dots, e_{L_f \times L_\theta}^{f\theta}\}$, so that $X_k^{f\theta}$ uniquely determines X_k^f and X_k^θ :

$$X_k^f = X_k^f(X_k^{f\theta}), \quad (5.19a)$$

$$X_k^\theta = X_k^\theta(X_k^{f\theta}). \quad (5.19b)$$

Lemma 5.2 Given X_k^f from (5.11), and the relationship (5.5), then $X_k^{f\theta}$, defined in (5.18), is a first order Markov process given by the equation

$$X_{k+1}^{f\theta} = (\mathbf{A}^{f\theta})' X_k^{f\theta} + M_{k+1}^{f\theta}. \quad (5.20)$$

Here, $M_{k+1}^{f\theta}$ is a $(\mathbf{A}^{f\theta}, \mathcal{F}_k^{f\theta})$ martingale increment, in that $E[M_{k+1}^{f\theta} | \mathcal{F}_k^{f\theta}] = 0$ where $\mathcal{F}_k^{f\theta}$ is the complete σ -field generated by $X_k^{f\theta}$, $l \leq k$, $\{\mathcal{F}_l^{f\theta}, l \in \mathcal{Z}^+\}$. Also, $\mathbf{A}^{f\theta} = (a_{ij}^{f\theta})$ and $\mathbf{A}^f = (a_{ij}^f)$. In addition, from (5.19a) and (5.19b), let $e_{n(i)}^f$ and $e_{m(i)}^\theta$ be the respective state space indicator vectors corresponding to $e_i^{f\theta}$. Now, $\mathbf{A}^{f\theta}$ is given from \mathbf{A}^f as follows

$$a_{ij}^{f\theta} = \delta(X_k^\theta - (X_k^f + X_{k-1}^\theta)) \times a_{n(i),n(j)}^f, \quad 1 \leq i, j \leq L_\theta \times L_f. \quad (5.21)$$

Proof : That $X_k^{f\theta} = X_k^f \otimes X_k^\theta$ is a first order Markov process follows, since both X_{k+1}^f and X_{k+1}^θ depend only on X_k^f and X_k^θ , and not on $X_{k-1}^f, X_{k-1}^\theta, \dots, X_0^f, X_0^\theta$ (as can be seen from (5.11) and (5.14)). Application of Lemma 2.1 leads directly to equation (5.20). For a proof of (5.21) the following steps are presented (here the obvious superscripts on the respective unit vectors, e , are omitted):

$$\begin{aligned} & P(X_k^{f\theta} = e_j | X_{k-1}^{f\theta} = e_i, X_{k-2}^{f\theta} = e_h, \dots) \\ &= P(X_k^f = e_{n(j)}, X_k^\theta = e_{m(j)} | X_{k-1}^f = e_{n(i)}, X_{k-2}^f = e_{n(h)}, \dots, X_{k-1}^\theta = e_{m(i)}, X_{k-2}^\theta = e_{m(h)}, \dots) \\ &= P(X_k^\theta = e_{m(j)} | X_k^f = e_{n(j)}, X_{k-1}^f = e_{n(i)}, X_{k-1}^\theta = e_{m(i)}) \times \\ & \quad P(X_k^f = e_{n(j)} | X_{k-1}^f = e_{n(i)}, X_{k-2}^f = e_{n(h)}, \dots, X_{k-1}^\theta = e_{m(i)}, X_{k-2}^\theta = e_{m(h)}, \dots) \\ & \text{from (5.5),} \\ &= \delta(X_k^\theta - (X_k^f + X_{k-1}^\theta)) \times P(X_k^f = e_{n(j)} | X_{k-1}^f = e_{n(i)}) \\ &= \delta(X_k^\theta - (X_k^f + X_{k-1}^\theta)) \times a_{n(i),n(j)}^f. \end{aligned}$$

where all arguments are interpreted modulo 2. ■

The observation process (5.8) can now be expressed in terms of the state X_k^θ .

$$\begin{pmatrix} y_k^R \\ y_k^I \end{pmatrix} = \begin{pmatrix} A_c \cos[z_\theta' X_k^\theta] & -A_c \sin[z_\theta' X_k^\theta] \\ A_c \sin[z_\theta' X_k^\theta] & A_c \cos[z_\theta' X_k^\theta] \end{pmatrix} \begin{pmatrix} g_k^R \\ g_k^I \end{pmatrix} + \begin{pmatrix} w_k^R \\ w_k^I \end{pmatrix}, \quad (5.22)$$

or equivalently with the appropriate definition of $h_\theta(\cdot)$,

$$\begin{aligned} \mathbf{y}_k &= h_\theta(X_k^\theta) x_k + \mathbf{w}_k \quad , \quad \mathbf{w}_k = N[0, R_k] \\ &= [h_\theta(e_1^\theta) x_k \quad h_\theta(e_2^\theta) x_k \dots h_\theta(e_{L_\theta}^\theta) x_k] X_k^\theta + \mathbf{w}_k \\ &= H_\theta' [I_{L_\theta} \otimes x_k] X_k^\theta + \mathbf{w}_k \quad , \end{aligned} \quad (5.23)$$

where the augmented matrix $H_\theta' = [h_\theta(e_1^\theta) \dots h_\theta(e_{L_\theta}^\theta)]$. Note that the Cartesian co-ordinates for the channel model, allow the observations to be written in a form which is linear in both X_k^θ and x_k . Also, $E[w_{k+1}^R | \mathcal{F}_k \vee \mathcal{Y}_k] = 0$ and $E[w_{k+1}^I | \mathcal{F}_k \vee \mathcal{Y}_k] = 0$, where \mathcal{Y}_l is the σ -field generated by y_k , $k \leq l$. It is usual to assume that w^R and w^I are independent so that the covariance matrix associated with the measurement noise vector \mathbf{w}_k has the form

$$R_k = \begin{bmatrix} \sigma_{w^R}^2 & 0 \\ 0 & \sigma_{w^I}^2 \end{bmatrix} . \quad (5.24)$$

It is now readily seen that

$$E[M_{k+1}^f | \mathcal{F}_k \vee \mathcal{Y}_k] = 0 , \quad (5.25a)$$

$$E[M_{k+1}^{f\theta} | \mathcal{F}_k \vee \mathcal{Y}_k] = 0 . \quad (5.25b)$$

The vector of parameterised probability densities can now be defined by $\mathbf{b}_k^\theta = (b_k^\theta(1), \dots, b_k^\theta(L_\theta))'$, for $b_k^\theta(i) \triangleq P[y_k | X_k^\theta = e_i^\theta, x_k]$, where

$$b_k^\theta(i) = \frac{1}{2\pi\sigma_w^2} \exp \left(- \frac{[y_k^R - (A_c \cos[z'_\theta e_i^\theta] g_k^R - A_c \sin[z'_\theta e_i^\theta] g_k^I)]^2}{2\sigma_w^2} - \frac{[y_k^I - (A_c \sin[z'_\theta e_i^\theta] g_k^R + A_c \cos[z'_\theta e_i^\theta] g_k^I)]^2}{2\sigma_w^2} \right) . \quad (5.26)$$

In summary, the following lemma is presented

Lemma 5.3 Under assumption (5.9) and (4.8), the signal model (5.1) to (5.7) has the following state space representation, in terms of the $L_f \times L_\theta$ dimension indicator function $X_k^{f\theta}$,

$$\begin{aligned}
 X_{k+1}^{f\theta} &= (\mathbf{A}^{f\theta})' X_k^{f\theta} + M_{k+1} \\
 x_{k+1} &= Fx_k + v_{k+1} \\
 \mathbf{y}_k &= H'_\theta [I_{L_\theta} \otimes x_k] X_k^\theta(X_k^{f\theta}) + \mathbf{w}_k
 \end{aligned} \tag{5.27}$$

or equivalently, in terms of the L_f and L_θ indicator functions X_k^θ and X_k^f respectively,


$$\begin{aligned}
 X_{k+1}^f &= (\mathbf{A}^f)' X_k^f + M_{k+1}^f \\
 X_{k+1}^\theta &= [\mathbf{A}^\theta(X_{k+1}^f)]' X_k^\theta \\
 x_{k+1} &= Fx_k + v_{k+1} \\
 \mathbf{y}_k &= H'_\theta [I_{L_\theta} \otimes x_k] X_k^\theta + \mathbf{w}_k
 \end{aligned} \tag{5.28}$$

■

Remarks 5.2: 1. Observe that both models are in terms of the channel parameters (states) x_k in a continuous range, and in terms of indicator functions (states) which belong to a finite-discrete set, being the vertices of a simplex.

2. The first model (5.27) has linear state dynamics with measurements bi-linear in x_k and $X_k^{f\theta}$. The second model (5.28) has linear dynamics for the states x_k and X_k^f , but X_{k+1}^θ is bi-linear in X_k^θ and X_{k+1}^f . The measurements are bi-linear in x_k and X_k^θ .

3. In Figure 5.1 the channel output for 5000 data points is presented, with channel noise variance $\sigma_w^2 = 0.00001$, and channel dynamics $\kappa(t) = 1 + 0.5 \sin(2\pi t/5000)$, $\phi(t) = 0.04\pi \cos(2\pi t/5000)$. In the simulations of Section 5.5, much more rapidly changing channel dynamics are used. Figure 5.1 is used here merely to illustrate the time-varying nature of the channel's effect on the signal constellation. □



$$\kappa_k = 1 + 0.5 \sin(2\pi k/5000), \phi_k = 0.04\pi \cos(2\pi k/5000), \text{SNR} = 2.4 \text{ dB}$$

Figure 5.1: FM signal constellation output from channel

5.2.5 Conditional Information-State Signal Models

First Model : Let $\hat{X}_{k|\mathcal{X}}^{f\theta}$ denote the conditional filtered state estimate of $X_k^{f\theta}$ at time k , given the channel parameter sequence $\mathcal{X}_k = \{x_0, \dots, x_k\}$, that is,

$$\hat{X}_{k|\mathcal{X}}^{f\theta} \triangleq E[X_k^{f\theta} | \mathcal{Y}_k, \mathcal{X}_k]. \quad (5.29)$$

It is possible to define $\mathbf{1}$ to be the column vector containing all ones, and the information state $\alpha_{k|\mathcal{X}}^{f\theta}$ is such that the i^{th} element

$$\alpha_{k|\mathcal{X}}^{f\theta}(i) \triangleq P(Y_k, X_k^{f\theta} = e_i | \mathcal{X}_k). \quad (5.30)$$

Observe that $\hat{X}_{k|\mathcal{X}}^{f\theta}$ can be expressed in terms of $\alpha_{k|\mathcal{X}}^{f\theta}$ by

$$\hat{X}_{k|\mathcal{X}}^{f\theta} = \langle \alpha_{k|\mathcal{X}}^{f\theta}, \mathbf{1} \rangle^{-1} \alpha_{k|\mathcal{X}}^{f\theta}. \quad (5.31)$$

Here $\alpha_{k|\mathcal{X}}^{f\theta}$ is computed using the following “forward” recursion:

$$\alpha_{k+1|\mathcal{X}}^{f\theta} = \mathbf{B}^{f\theta}(y_{k+1}, x_{k+1}) (\mathbf{A}^{f\theta})' \alpha_{k|\mathcal{X}}^{f\theta}, \quad (5.32)$$

where $\mathbf{B}^{f\theta}(y_{k+1}, x_{k+1}) = \text{diag}(b_{k+1}^{f\theta}(1), \dots, b_{k+1}^{f\theta}(L_f \times L_\theta))$, and where $b_{k+1}^{f\theta}(i)$ is defined by $b_{k+1}^{f\theta}(i) \triangleq b^\theta(y_{k+1}, e_{m(i)}^\theta(e_i^{f\theta}), x_{k+1})$, for $e_{m(i)}^\theta(e_i^{f\theta})$ defined in Lemma 5.2.

It is now possible to express the observations, y_k , in terms of the un-normalised conditional estimates, $\alpha_{k-1|\mathcal{X}}^{f\theta}$.

Lemma 5.4 *The conditional measurements $\mathbf{y}_{k|\mathcal{X}}$ are defined by*

$$\mathbf{y}_{k|\mathcal{X}} = H'_{f\theta} [I_{L_f \times L_\theta} \otimes x_k] \langle \alpha_{k-1|\mathcal{X}}^{f\theta}, \mathbf{1} \rangle^{-1} (\mathbf{A}^{f\theta})' \alpha_{k-1|\mathcal{X}}^{f\theta} + n_{k|\mathcal{X}}, \quad (5.33)$$

where $\alpha_{k|\mathcal{X}}^{f\theta}$ is defined in (5.32), $H'_{f\theta} = [h_\theta(e_{m(1)}^\theta(e_1^{f\theta})) \dots h_\theta(e_{m(L_f \times L_\theta)}^\theta(e_{L_f \times L_\theta}^{f\theta}))]$, and $n_{k|\mathcal{X}}$ is a $(\mathcal{X}_k, \mathcal{Y}_{k-1})$ martingale increment.

Proof : Following standard arguments since $\alpha_{k|\mathcal{X}}^{f\theta}$ is measurable with respect to $\{\mathcal{X}_k, \mathcal{Y}_k\}$, $E[w_{k+1}^R | \mathcal{Y}_k] = 0$, $E[w_{k+1}^I | \mathcal{Y}_k] = 0$ and $E[M_{k+1}^{f\theta} | \mathcal{Y}_k] = 0$, then

$$\begin{aligned} E[n_{k|\mathcal{X}} | \mathcal{X}_k, \mathcal{Y}_{k-1}] &= E[H'_\theta [I_{L_\theta} \otimes x_k] X_k^\theta (X_k^{f\theta}) + \mathbf{w}_k \\ &\quad - H'_{f\theta} [I_{L_f \times L_\theta} \otimes x_k] \langle \alpha_{k-1|\mathcal{X}}^{f\theta}, \mathbf{1} \rangle^{-1} (\mathbf{A}^{f\theta})' \alpha_{k-1|\mathcal{X}}^{f\theta} | \mathcal{X}_{k-1}, \mathcal{Y}_{k-1}] \\ &= H'_{f\theta} [I_{L_f \times L_\theta} \otimes x_k] \left((\mathbf{A}^{f\theta})' \hat{X}_{k-1|\mathcal{X}}^{f\theta} - \langle \alpha_{k-1|\mathcal{X}}^{f\theta}, \mathbf{1} \rangle^{-1} (\mathbf{A}^{f\theta})' \alpha_{k-1|\mathcal{X}}^{f\theta} \right) \\ &= 0. \end{aligned}$$

■

In summary, the following lemma is presented,

Lemma 5.5 *The state space representation (5.27) can be reformulated to give the following conditional information state signal model, with states $\alpha_{k|\mathcal{X}}^{f\theta}$,*

$\begin{aligned} \alpha_{k+1 \mathcal{X}}^{f\theta} &= \mathbf{B}^{f\theta}(y_{k+1}, x_{k+1}) (\mathbf{A}^{f\theta})' \alpha_{k \mathcal{X}}^{f\theta} \\ x_{k+1} &= F x_k + v_k \\ \mathbf{y}_{k \mathcal{X}} &= H'_{f\theta} [I_{L_f \times L_\theta} \otimes x_k] \langle \alpha_{k-1 \mathcal{X}}^{f\theta}, \mathbf{1} \rangle^{-1} (\mathbf{A}^{f\theta})' \alpha_{k-1 \mathcal{X}}^{f\theta} + n_{k \mathcal{X}} \end{aligned}$	(5.34)
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------

■

This fading channel signal model is now in standard state space form to allow the application of extended Kalman filtering.

Second Model : Let $\hat{X}_{k|\mathcal{X},\mathcal{X}^{\theta}}^f$ and $\hat{X}_{k|\mathcal{X},\mathcal{X}^f}^{\theta}$ denote the conditional filtered state estimates of X_k^f and X_k^{θ} respectively, where $\mathcal{X}_k = \{x_0, \dots, x_k\}$, $\mathcal{X}_k^{\theta} = \{X_0^{\theta}, \dots, X_k^{\theta}\}$, and $\mathcal{X}_k^f = \{X_0^f, \dots, X_k^f\}$.

In particular,

$$\hat{X}_{k|\mathcal{X},\mathcal{X}^{\theta}}^f \triangleq E[X_k^f | \mathcal{Y}_k, \mathcal{X}_k, \mathcal{X}_{k-1}^{\theta}], \quad (5.35a)$$

$$\hat{X}_{k|\mathcal{X},\mathcal{X}^f}^{\theta} \triangleq E[X_k^{\theta} | \mathcal{Y}_k, \mathcal{X}_k, \mathcal{X}_k^f]. \quad (5.35b)$$

It is possible to define $\mathbf{1}$ to be the column vector containing all ones, and the information states, $\alpha_{k|\mathcal{X},\mathcal{X}^{\theta}}^f$ and $\alpha_{k|\mathcal{X},\mathcal{X}^f}^{\theta}$, be such that their i^{th} elements are given respectively by

$$\alpha_{k|\mathcal{X},\mathcal{X}^{\theta}}^f(i) \triangleq P(Y_k, X_k^f = e_i^f | \mathcal{X}_k, \mathcal{X}_{k-1}^{\theta}), \quad (5.36a)$$

$$\alpha_{k|\mathcal{X},\mathcal{X}^f}^{\theta}(i) \triangleq P(Y_k, X_k^{\theta} = e_i^{\theta} | \mathcal{X}_k, \mathcal{X}_k^f). \quad (5.36b)$$

Observe that $\hat{X}_{k|\mathcal{X},\mathcal{X}^{\theta}}^f$ and $\hat{X}_{k|\mathcal{X},\mathcal{X}^f}^{\theta}$ can be expressed in terms of $\alpha_{k|\mathcal{X},\mathcal{X}^{\theta}}^f$ and $\alpha_{k|\mathcal{X},\mathcal{X}^f}^{\theta}$ respectively, by

$$\hat{X}_{k|\mathcal{X},\mathcal{X}^{\theta}}^f = \langle \alpha_{k|\mathcal{X},\mathcal{X}^{\theta}}^f, \mathbf{1} \rangle^{-1} \alpha_{k|\mathcal{X},\mathcal{X}^{\theta}}^f, \quad (5.37a)$$

$$\hat{X}_{k|\mathcal{X},\mathcal{X}^f}^{\theta} = \langle \alpha_{k|\mathcal{X},\mathcal{X}^f}^{\theta}, \mathbf{1} \rangle^{-1} \alpha_{k|\mathcal{X},\mathcal{X}^f}^{\theta}. \quad (5.37b)$$

Here $\alpha_{k|\mathcal{X},\mathcal{X}^{\theta}}^f$ is computed using the following “forward” recursion [Rabiner 1989]:

$$\alpha_{k+1|\mathcal{X},\mathcal{X}^{\theta}}^f = \mathbf{B}^f(y_{k+1}, x_{k+1}, X_k^{\theta}) (\mathbf{A}^f)' \alpha_{k|\mathcal{X},\mathcal{X}^{\theta}}^f, \quad (5.38)$$

where $\mathbf{B}^f(y_{k+1}, x_{k+1}, X_k^{\theta}) = \text{diag}(b_{k+1}^f(1), \dots, b_{k+1}^f(L_f))$, and where $b_{k+1}^f(i) \triangleq P[y_{k+1} | X_{k+1}^f = e_i^f, x_{k+1}, X_k^{\theta}]$, where

$$b_k^f(i) = \frac{1}{2\pi\sigma_w^2} \exp\left(-\frac{[y_k^R - (A_c \cos[z_{\theta}'([\mathbf{A}^{\theta}(e_i^f)]' X_{k-1}^{\theta})]g_k^R - A_c \sin[z_{\theta}'([\mathbf{A}^{\theta}(e_i^f)]' X_{k-1}^{\theta})]g_k^I)]^2}{2\sigma_w^2} - \frac{[y_k^I - (A_c \sin[z_{\theta}'([\mathbf{A}^{\theta}(e_i^f)]' X_{k-1}^{\theta})]g_k^R + A_c \cos[z_{\theta}'([\mathbf{A}^{\theta}(e_i^f)]' X_{k-1}^{\theta})]g_k^I)]^2}{2\sigma_w^2}\right). \quad (5.39)$$

Also, $\alpha_{k|\mathcal{X},\mathcal{X}^f}^\theta$ is conveniently computed using (5.14) to give the following “forward” recursion:

$$\begin{aligned}\alpha_{k+1|\mathcal{X},\mathcal{X}^f}^\theta &= \mathbf{B}^\theta(y_{k+1}, x_{k+1}) \left(\sum_{i=1}^{L_f} [\mathbf{A}^\theta(e_i^f)]' X_{k+1}^f(i) \right) \alpha_{k|\mathcal{X},\mathcal{X}^f}^\theta \\ &= \mathbf{B}^\theta(y_{k+1}, x_{k+1}) \mathcal{A}^\theta[X_{k+1}^f \otimes I_{L_\theta}] \alpha_{k|\mathcal{X},\mathcal{X}^f}^\theta,\end{aligned}\quad (5.40)$$

where $\mathbf{B}^\theta(y_{k+1}, x_{k+1}) = \text{diag}(b_{k+1}^\theta(1), \dots, b_{k+1}^\theta(L_\theta))$, and where $b_k^\theta(i)$ is given in (5.26). Also, $\mathcal{A}^\theta = [[\mathbf{A}^\theta(e_1^f)]' \dots [\mathbf{A}^\theta(e_{L_f}^f)]']$.

It is now possible to express the observations, y_k , in terms of the un-normalised conditional estimates, $\alpha_{k-1|\mathcal{X},\mathcal{X}^\theta}^f$.

Lemma 5.6 *The conditional measurements $\mathbf{y}_{k|\mathcal{X},\mathcal{X}^\theta}$ are defined by*

$$\mathbf{y}_{k|\mathcal{X},\mathcal{X}^\theta} = H_\theta' [I_{L_\theta} \otimes x_k] \mathcal{A}^\theta[\langle \alpha_{k-1|\mathcal{X},\mathcal{X}^\theta}^f, \mathbf{1} \rangle^{-1} (\mathbf{A}^f)' \alpha_{k-1|\mathcal{X},\mathcal{X}^\theta}^f \otimes I_{L_\theta}] X_{k-1}^\theta + n_{k|\mathcal{X},\mathcal{X}^\theta}, \quad (5.41)$$

where $n_{k|\mathcal{X},\mathcal{X}^\theta}$ is a $(\mathcal{X}_k, \mathcal{X}_{k-1}^\theta, \mathcal{Y}_{k-1})$ martingale increment.

Proof: Following standard arguments since $\alpha_{k|\mathcal{X},\mathcal{X}^\theta}^f$ is measurable with respect to $\{\mathcal{X}_k, \mathcal{X}_{k-1}^\theta, \mathcal{Y}_k\}$, $E[w_{k+1}^R | \mathcal{Y}_k] = 0$, $E[w_{k+1}^I | \mathcal{Y}_k] = 0$ and $E[m_{k+1} | \mathcal{Y}_k] = 0$, then

$$\begin{aligned}E[n_{k|\mathcal{X},\mathcal{X}^\theta} | \mathcal{X}_k, \mathcal{X}_{k-1}^\theta, \mathcal{Y}_{k-1}] &= E[H_\theta' [I_{L_\theta} \otimes x_k] X_k^\theta + \mathbf{w}_k \\ &\quad - H_\theta' [I_{L_\theta} \otimes x_k] \mathcal{A}^\theta[\langle \alpha_{k-1|\mathcal{X},\mathcal{X}^\theta}^f, \mathbf{1} \rangle^{-1} (\mathbf{A}^f)' \alpha_{k-1|\mathcal{X},\mathcal{X}^\theta}^f \otimes I_{L_\theta}] X_{k-1}^\theta | \mathcal{X}_k, \mathcal{X}_{k-1}^\theta, \mathcal{Y}_{k-1}] \\ &= H_\theta' [I_{L_\theta} \otimes x_k] \left(\mathcal{A}^\theta[\hat{X}_{k|\mathcal{X},\mathcal{X}^\theta}^f \otimes I_{L_\theta}] X_{k-1}^\theta \right. \\ &\quad \left. - \mathcal{A}^\theta[\langle \alpha_{k-1|\mathcal{X},\mathcal{X}^\theta}^f, \mathbf{1} \rangle^{-1} (\mathbf{A}^f)' \alpha_{k-1|\mathcal{X},\mathcal{X}^\theta}^f \otimes I_{L_\theta}] X_{k-1}^\theta \right) \\ &= 0.\end{aligned}$$

In practise, however, as noted above, there is no access to X_{k-1}^θ , but its conditional expectation, $\alpha_{k-1|\mathcal{X},\mathcal{X}^f}^\theta$. Therefore the conditional measurement for this second model, $\mathbf{y}_{k|\mathcal{X},\mathcal{X}^\theta}$, which can be used in practise, does not have a martingale increment noise term, $n_{k|\mathcal{X},\mathcal{X}^\theta}$. In addition, the covariance matrix, R_n , of $n_{k|\mathcal{X},\mathcal{X}^\theta}$, is of higher magnitude than that of w_k . The exact form of R_n

is however too complicated for presentation in this chapter, and application based estimates of R_n can be used when implementing these algorithms.

In summary, the following lemma is presented,

Lemma 5.7 *The state space representation (5.28) can be reformulated to give the following conditional information state signal model, with states $\alpha_{k|\mathcal{X},\mathcal{X}^\theta}^f$ and $\alpha_{k|\mathcal{X},\mathcal{X}^f}^\theta$,*

$$\begin{aligned}
 \alpha_{k+1|\mathcal{X},\mathcal{X}^\theta}^f &= \mathbf{B}^f(y_{k+1}, x_k, X_k^\theta) (\mathbf{A}^f)' \alpha_{k|\mathcal{X},\mathcal{X}^\theta}^f \\
 \alpha_{k+1|\mathcal{X},\mathcal{X}^f}^\theta &= \mathbf{B}^\theta(y_{k+1}, x_{k+1}) \mathcal{A}^\theta[\langle \alpha_{k|\mathcal{X},\mathcal{X}^\theta}^f, \mathbf{1} \rangle^{-1} (\mathbf{A}^f)' \alpha_{k|\mathcal{X},\mathcal{X}^\theta}^f \otimes I_{L_\theta}] \alpha_{k|\mathcal{X},\mathcal{X}^f}^\theta \\
 x_{k+1} &= Fx_k + v_k \\
 \mathbf{y}_{k|\mathcal{X}} &= H_\theta' [I_{L_\theta} \otimes x_k] \mathcal{A}^\theta\left[\frac{(\mathbf{A}^f)' \alpha_{k-1|\mathcal{X},\mathcal{X}^\theta}^f}{\langle \alpha_{k-1|\mathcal{X},\mathcal{X}^\theta}^f, \mathbf{1} \rangle} \otimes I_{L_\theta}\right] \frac{\alpha_{k-1|\mathcal{X},\mathcal{X}^f}^\theta}{\langle \alpha_{k-1|\mathcal{X},\mathcal{X}^f}^\theta, \mathbf{1} \rangle} + n_{k|\mathcal{X}}
 \end{aligned} \tag{5.42}$$

■

This fading channel signal model is now in standard state space form to allow the application of extended Kalman filtering.

5.3 Higher Order Message Models

Lemma 5.7 provides insight into methods for coping with higher order analog signal models, and thus allows Assumption (5.9) to be relaxed. To do this it is necessary to continue quantising the range of phases, while allowing the model of the frequencies to be in a continuous range, and in vector form. Therefore the state space representation of the frequency message is no longer a first order system. Also, the quantisation errors are now involved with the phase estimate. This approach allows the message frequency model to be in a continuous range, while still employing the attractive optimal filtering of the HMM filter for the phase. The following state space model

applies,

$$\begin{aligned}
 x_{k+1}^f &= F_k^f x_k^f + v_{k+1}^f \\
 X_{k+1}^\theta &= [\mathbf{A}^\theta(X_k^f(h'x_k^f))] X_k^\theta \\
 x_{k+1} &= Fx_k + v_{k+1} \\
 \mathbf{y}_k &= H_\theta' [I_{L_\theta} \otimes x_k] X_k^\theta + \mathbf{w}_k
 \end{aligned} \tag{5.43}$$

where F_k^f is the function associated with the dynamics of the frequency message given by the state x_k^f , the scalar message frequency is given by $h'x_k^f$, and X_k^f is the quantised frequency in state space form.

Following the steps presented in Section 5.2, for the signal under Assumption (5.9), an information-state signal model can be generated for this higher order state space signal model. As with the previous information-state signal model, this higher order model also results in an HMM/EKF scheme similar to that presented in the following section.

5.4 Adaptive HMM Algorithms

Two adaptive HMM schemes are presented here for the MDPSK and FM signal model. The technique used here for linking HMM signal state estimation with EKF channel parameter tracking is the same as was presented in the previous chapter for the QAM signal model. The equations are, however, slightly more complex. The first scheme is referred to as the HMM/EKF scheme (and is a full nonlinear EKF scheme for the information state signal models (5.34) and (5.42), with the augmented vectors given below), and the second scheme is referred to as the HMM/KF scheme, (and is with a simplification assumption which results in a KF for channel estimation, coupled with HMM filters for signal state estimation). In both cases, it is important to remember that unlike standard approaches, such as matched filters or Viterbi decoders which use MAP estimates of the digital state, this technique utilises the full information state probability distribution vector so as to feed back information about the reliability of the estimate as well as the most likely state. Also in this section implementation aspects are discussed when considering MDPSK and FM signals.

5.4.1 Adaptive HMM/EKF Scheme

Let

$$\underline{x}_k = (\alpha_k^{f\theta}, x_k)', \quad \text{for model (5.34)}, \quad (5.44a)$$

$$= (\alpha_k^f, \alpha_k^\theta, x_k)', \quad \text{for model (5.42)}. \quad (5.44b)$$

Then (5.34) and (5.42) can be written as

$$\underline{x}_{k+1} = f_k(\underline{x}_k) + g_k(\underline{x}_k)v_k, \quad (5.45a)$$

$$y_k = h_k(\underline{x}_k) + n_k, \quad (5.45b)$$

where the nonlinear functions are given by

$$\left\{ \begin{array}{l} f_k(\underline{x}_k) = \begin{pmatrix} \mathbf{B}^{f\theta}(x_k)(\mathbf{A}^{f\theta})'\alpha_k^{f\theta} \\ Fx_k \end{pmatrix}, \quad g_k(\underline{x}_k) = \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix} \\ h_k(\underline{x}_k) = H'_{f\theta} [I_{L_f \times L_\theta} \otimes x_k] \langle \alpha_{k-1}^{f\theta}, \mathbf{1} \rangle^{-1} (\mathbf{A}^{f\theta})'\alpha_{k-1}^{f\theta} \end{array} \right\}, \quad \text{for model (5.34)}, \quad (5.46a)$$

and

$$\left\{ \begin{array}{l} f_k(\underline{x}_k) = \begin{pmatrix} \mathbf{B}^f(x_k)(\mathbf{A}^f)'\alpha_k^f \\ \mathbf{B}^\theta(x_k)\mathcal{A}^\theta[\langle \alpha_k^f, \mathbf{1} \rangle^{-1} (\mathbf{A}^f)'\alpha_k^f \otimes I_{L_\theta}] \alpha_k^\theta \\ Fx_k \end{pmatrix}, \quad g_k(\underline{x}_k) = \begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix} \\ h_k(\underline{x}_k) = H'_\theta [I_{L_\theta} \otimes x_k] \mathcal{A}^\theta \left[\frac{(\mathbf{A}^f)'\alpha_{k-1}^f}{\langle \alpha_{k-1}^f, \mathbf{1} \rangle} \otimes I_{L_\theta} \right] \frac{\alpha_{k-1}^\theta}{\langle \alpha_{k-1}^\theta, \mathbf{1} \rangle} \end{array} \right\}, \quad \text{for model (5.42)}. \quad (5.46b)$$

As for the adaptive HMM algorithm with an EKF presented in the previous chapter, the EKF equations for (5.45a) and (5.45b) are given in (4.36a) to (4.36e).

Figure 4.3 applies here, and gives a block diagram for this adaptive HMM scheme, for the first model. Figure 5.2 gives a block diagram for this adaptive HMM scheme, for the second model, when switch 1 and 2 are closed. As with Figure 4.3, if switch 1 and 2 were in the open position then the HMM/KF scheme given below would result.

Remark 5.3: This HMM/EKF scheme suffers from the same problems discussed in the previous chapter in relation to the QAM signal model (see Remarks 4.5). These implementation difficulties again provide incentive to explore sub-optimal schemes such as the HMM/KF algorithm which involves coupling separated state and parameter estimators. \square

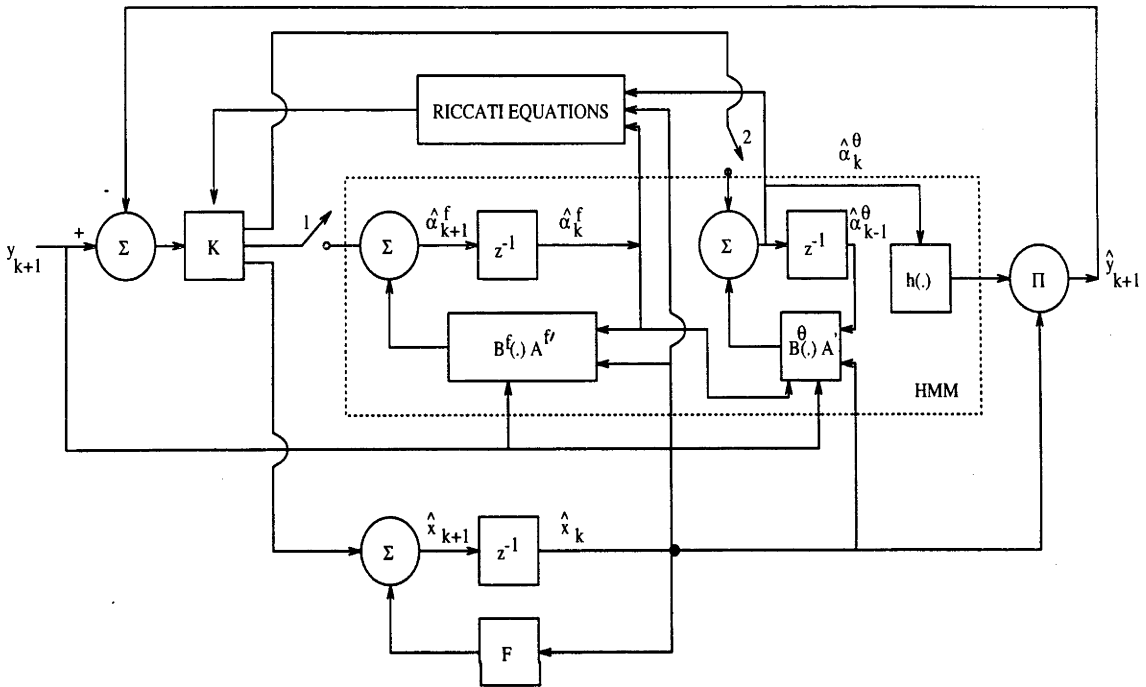


Figure 5.2: EKF/HMM scheme for adaptive HMM filter for second model

5.4.2 Adaptive HMM/KF Schemes

This scheme can be viewed as a derivative of the HMM/EKF scheme above by setting the Kalman gain term, associated with the $\hat{\alpha}_k^{f\theta}$, $\hat{\alpha}_k^f$, $\hat{\alpha}_k^\theta$ updates, to zero. The rationale for this is the same as that in Chapter 4, and is based on the fact that if the parameters were constant then the HMM/RPE scheme of Chapter 2 applies, for which there are strong theoretical foundations. If the channel is only slowly varying, then it is expected that the components of the Kalman gain associated with the $\hat{\alpha}_k^{f\theta}$, $\hat{\alpha}_k^f$, $\hat{\alpha}_k^\theta$ updates, will be asymptotically small. The resulting scheme can be viewed as coupled conditional HMM filters together with a conditional Kalman filter as follows:

In the case of (5.34), the HMM estimator for the signal state, $\hat{\alpha}_k^{f\theta}$, conditioned on the channel estimate sequence $\{\hat{x}_k\}$, is given by

$$\hat{\alpha}_{k+1|k}^{f\theta} = \mathbf{B}^{f\theta}(y_{k+1}, \hat{x}_k) (\mathbf{A}^{f\theta})' \hat{\alpha}_{k|\hat{x}_{k-1}}^{f\theta}, \quad (5.47a)$$

$$\hat{X}_{k|\hat{x}_{k-1}}^{f\theta} = \langle \alpha_{k|\hat{x}_{k-1}}^{f\theta}, \underline{1} \rangle^{-1} \alpha_{k|\hat{x}_{k-1}}^{f\theta}. \quad (5.47b)$$

In the case of (5.42), the conditional HMM estimators for the signal states, $\hat{\alpha}_k^f, \hat{\alpha}_k^\theta$, are given by,

$$\hat{\alpha}_{k+1|\hat{x}_k, \hat{\alpha}_k^\theta}^f = \mathbf{B}^f(y_{k+1}, \hat{x}_k, \hat{\alpha}_k^\theta) (\mathbf{A}^f)' \hat{\alpha}_{k|\hat{x}_{k-1}, \hat{\alpha}_{k-1}^\theta}^f, \quad (5.48a)$$

$$\hat{X}_{k|\hat{x}_{k-1}, \hat{\alpha}_{k-1}^\theta}^f = \langle \alpha_{k|\hat{x}_{k-1}, \hat{\alpha}_{k-1}^\theta}^f, \underline{1} \rangle^{-1} \alpha_{k|\hat{x}_{k-1}, \hat{\alpha}_{k-1}^\theta}^f, \quad (5.48b)$$

and

$$\hat{\alpha}_{k+1|\hat{x}_k, \hat{\alpha}_k^f}^\theta = \mathbf{B}^\theta(y_{k+1}, \hat{x}_k, \hat{\alpha}_k^f) (\mathbf{A}^\theta)' \hat{\alpha}_{k|\hat{x}_{k-1}, \hat{\alpha}_{k-1}^f}^\theta, \quad (5.49a)$$

$$\hat{X}_{k|\hat{x}_{k-1}, \hat{\alpha}_{k-1}^f}^\theta = \langle \alpha_{k|\hat{x}_{k-1}, \hat{\alpha}_{k-1}^f}^\theta, \underline{1} \rangle^{-1} \alpha_{k|\hat{x}_{k-1}, \hat{\alpha}_{k-1}^f}^\theta. \quad (5.49b)$$

The Kalman filter equations for the channel parameter, x_k , conditioned on the indicator state estimates, $\hat{X}_k^{f\theta}$ in the first case, and \hat{X}_k^f and \hat{X}_k^θ in the second case, are given in (4.38a) to (4.38e), and are reproduced here:

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k [y_k - H_k' \hat{x}_{k|k-1}], \quad (5.50a)$$

$$\hat{x}_{k+1|k} = F \hat{x}_{k|k}, \quad (5.50b)$$

$$K_k = \Sigma_{k|k-1} H_k [H_k' \Sigma_{k|k-1} H_k + R_k]^{-1}, \quad (5.50c)$$

$$\Sigma_{k|k} = \Sigma_{k|k-1} - \Sigma_{k|k-1} H_k [H_k' \Sigma_{k|k-1} H_k + R_k]^{-1} H_k' \Sigma_{k|k-1}, \quad (5.50d)$$

$$\Sigma_{k+1|k} = F \Sigma_{k|k} F' + Q_k, \quad (5.50e)$$

where

$$H'_k = \begin{cases} \partial(H'_{f\theta} [I_{L_f \times L_\theta} \otimes Fx_{k-1}] \langle \alpha_{k-1}^{f\theta}, \underline{1} \rangle^{-1} (\mathbf{A}^{f\theta})' \alpha_{k-1}^{f\theta}) / \partial x_k & \text{for model (5.34) ,} \\ \partial(H'_\theta [I_{L_\theta} \otimes Fx_{k-1}] \mathcal{A}^\theta [\frac{(\mathbf{A}^f)' \alpha_{k-1}^f}{\langle \alpha_{k-1}^f, \underline{1} \rangle} \otimes I_{L_\theta}] \frac{\alpha_{k-1}^\theta}{\langle \alpha_{k-1}^\theta, \underline{1} \rangle}) / \partial x_k & \text{for model (5.42) ,} \end{cases} \quad (5.51)$$

and R is the covariance matrix of the noise on the observations w_k given in (5.24), Q is the covariance matrix of v_k , given in (4.9), and Σ is the covariance matrix of the channel parameter estimate \hat{x}_k , (x_k is defined in (5.6)).

A further sub-optimal HMM/KF scheme can be generated by using the state space signal models (5.34) and (5.42), and estimating the KF conditioned on a maximum *a priori* probability estimates $(\hat{\alpha}_k^f)^{MAP}$, $(\hat{\alpha}_k^\theta)^{MAP}$ and $(\hat{\alpha}_k^{f\theta})^{MAP}$. In fact hybrid versions can be derived by setting the small valued, that is low probability, elements of $(\hat{\alpha}_k^f)$, $(\hat{\alpha}_k^\theta)$ and $(\hat{\alpha}_k^{f\theta})$ to zero and re-normalising.

5.5 Simulation Studies

Example 5.1: In this example an 8-DPSK signal was demodulated using the techniques presented in this chapter. It is assumed that a coding scheme was employed in transmission, leading to the following signal properties. The signal is of amplitude $A_c = 0.2$, and assumption (5.9) holds with $a_{ii}^f = 0.95$, $a_{ij}^f = (1 - 0.95)/(L_f - 1)$ for $i \neq j$, where $L_f = 8$ in this example. The channel characteristics used were stochastic channels, generated from low pass filtered white Gaussian noise. The variance of the Gaussian process is 1 for amplitude variations, and 5 for phase variations. The bandwidth of the LPF is W_c times the bit rate, where values of W_c are shown in Figure 5.3. The estimation scheme used here is the de-coupled HMM/KF scheme implemented on the 8-DPSK signal model given in (5.42). For comparisons to the HMM/KF technique, the standard MF/AGC/PLL scheme is used, (diagrammatically represented in Figure 4.6), where the MF estimates the phase, and then successive phases are subtracted in order to gain estimates for the frequency. Computationally, the MF is of course less taxing, however for mobile communications under the conditions (8-DPSK, 19.2 kB/sec, $f_c=1800$ MHz, car travelling at 100 km/h), the processing power required for the HMM/KF approach is only 3 MFlops, which is reasonable with current DSP technology. Therefore the approach presented in this chapter is computationally feasible, and is seen in this example to out-perform the traditional scheme for

the case of non-equally-probable messages, while being identical to the traditional scheme in the equally-probable message case.

The results for this example are displayed in Figure 5.3, where signals of length 50000 data points have been used to generate bit-error-rate (BER) values. It should be noted that the case of $W_c = 0.04$ is one of severe fading, and it is seen that even under such conditions, the HMM/KF scheme performs well.

Example 5.2: In order to demonstrate the application of the HMM techniques to analog signals, a frequency modulation scheme, under assumption (5.9) with $a_{ii}^f = 0.95$, was generated with $L_f = 16$. θ was quantised into $L_\theta = 32$ values, under Lemma 5.1. The signal is of amplitude $A_c = 0.2$. The channel characteristics used were deterministic as opposed to the stochastic, low pass filtered white noise channels used in Example 5.1. The deterministic channel allows the results to be displayed in a manner which more clearly shows tracking ability of these schemes. The channel characteristics were given by

$$\begin{aligned}\kappa_k &= 1 + 0.5 \sin(3\pi k/1000), \\ \phi_k &= 0.75\pi \cos(10\pi k/1000),\end{aligned}$$

and the signal to noise ratio associated with the observations in the absence of fading is $\text{SNR} = (E_b/\sigma_w^2) = 2.4\text{dB}$, where E_b is the energy per bit associated with the transmitted signal, if the signal were a 16-FSK digital signal. Of course much lower SNRs can be accommodated in the presence of more slowly varying channels, and it should be noted that the SNR effectively varies as the channel fades. The lowest effective SNR in this example occurs at $k = 500$ where $\text{SNR} = 1.8\text{dB}$. The channels presented here are very fast in the case of FAX and modem applications, but are more reasonable in applications involving mobile communications and indoor communication channels [Hashemi 1993, Loo and Secord 1991].

The estimation scheme used here is the de-coupled HMM/KF scheme implemented on the FM signal model given in (5.42). The simulations which were carried out using the signal representation in terms of the $L_f \times L_\theta$ dimension vector $\alpha_k^{f\theta}$ were found to be too computationally taxing. The results for the de-coupled scheme are presented in Figure 5.4 to Figure 5.6 and show that even though the channel changes quite quickly, good estimates are generated. Figure 5.4 shows the true channel values and the estimated values in real and imaginary format, that is, exactly as estimated from (5.50a) to (5.50e). Figure 5.5 shows the actual channel amplitude κ_k , and the

estimate of this, generated from the estimates in Figure 5.4. Likewise, Figure 5.6 shows the actual channel phase shift ϕ_k and the estimate of this, generated from the estimates in Figure 5.4. These results show sudden phase shifts, seen as glitches in the phase estimate in Figure 5.6. These can be any multiple of π/L_f due to the symmetry of the phase quantisation. In this case, there is tracking degradation over the period where channel amplitude is less than one. It is natural that the estimates during these periods be worse, since the noise on the observations is effectively greater when $\kappa_k < 1$.

Example 5.3: In this example the FM signal parameters and channel characteristics were the same as those used in Example 5.2. The SNR, without fading, was however SNR = 1.4dB. Note again that at the maximum fading point, $k = 500$, the SNR is effectively SNR = 0.8dB. The phase estimate results are presented Figure 5.7 and show that the phase shift estimates are not able to track the actual phase shifts as accurately as in the lower noise case of Example 5.2. This example demonstrates the effect on robustness when the observation noise is assumed to give SNR = 0.4dB without fading, an order of magnitude greater than the actual value. For this robust case, the channel gain estimate shows no appreciable difference from the non-robust estimate, however the channel phase shift estimate shown in Figure 5.8 shows improved tracking ability.

5.6 Conclusions

In this chapter adaptive HMM on-line state and parameter estimation schemes have been derived for MDPSK and FM signals in fading communication channels. The schemes which blend EKF and HMM techniques, are based on optimal techniques, but are inevitably sub-optimal. Simulation studies are presented which show the ability to track time varying channel parameters, for both MDPSK and FM signals. Comparisons to traditional schemes demonstrate performance advantages for the HMM/KF techniques presented.

5.7 Figures

The figures for this chapter are now presented.




Figure 5.3: BER v SNR for 8-DPSK




Figure 5.4: Real and imaginary channel gain for FM




Figure 5.5: Channel amplitude gain (κ) estimates for FM




Figure 5.6: Channel phase rotation (ϕ_k) estimates for FM





Figure 5.7: Channel phase rotation (ϕ_k) estimates for FM



SNR = 1.4 dB but with SNR = 0.4 dB used in algorithm

Figure 5.8: Channel phase rotation (ϕ_k) estimates for FM

Chapter 6

Risk Sensitive Control Problems

6.1 Introduction

This chapter addresses optimal regulation and tracking problems for both linear systems and hidden Markov models. Specifically, it considers cost criteria involving *risk-sensitive* policies. Recently there has been much interest in risk-sensitive control techniques. Such policies lead to an optimal solution, similar to the case for linear quadratic Gaussian (LQG) control. However, with a risk-sensitive policy, the controller's sensitivity to risk can be varied. One application area for risk-sensitive control has been economics, where risk-sensitivity is termed *hedging* or *risk-aversion*, for example see Karp [1988] and Caravani [1986]. These papers illustrate that advantages can be gained from the risk-sensitive approach, for problems such as dynamic trading and futures market prediction.

The particular risk-sensitive control policy considered in this chapter involves an exponential in the cost function. This approach was first presented in Jacobson [1973], when considering the risk-sensitive LQG problem with state feedback. Jacobson demonstrated a link between exponential performance criteria and deterministic differential games. He showed that the risk-sensitive approach provides a method for varying the the robustness of the controller, and noted that in the case of no risk, or *risk-neutral*, the well known LQG solution [Anderson and Moore 1989] would result.

The discrete-time risk-sensitive linear-quadratic-Gaussian (RLQG) output feedback control solution was first presented in Whittle [1981], where use was made of a risk-sensitive version of the certainty equivalence principle. This allowed the state estimation and control optimisation to be decoupled, solved separately, and then re-coupled. The continuous-time case was solved in Bensoussan and van Schuppen [1985], where a technique which generalises to the nonlinear case, was used. More recent developments in risk-sensitive control have included a solution to the output feedback control problem for nonlinear systems using information-state techniques [James *et al.* 1994]. The solution is of course infinite dimensional, but does not require the use of a certainty equivalence principle.

This chapter presents the output feedback RLQG solution, derived via the methods in Bensoussan and van Schuppen [1985], James *et al.* [1994]. It considers specifically, the case of tracking a desired trajectory. The resulting equations are shown to be consistent with those presented in Whittle [1981], and in the “risk-neutral” case, consistent with the standard LQG solution. In addition, methods are discussed for achieving zero steady state error for tracking with risk-sensitive control policies.

The solution to the discrete-time hidden Markov model (HMM) risk-sensitive tracking problem is also presented. In this case a finite-dimensional information-state is derived. However, the control solution requires an infinite-dimensional dynamic programming problem to be solved. Fortunately though, it is possible to discretise the information-state space and thus obtain approximate solutions.

The key to the technique used in this chapter is that an information-state is chosen in such a way that it represents both a state estimate and the cost incurred to the time of the estimate, as in Bensoussan and van Schuppen [1985]. A change of reference probability measure is used to arrive at a linear recursive update equation for this information-state. Then, dynamic programming methods are employed to obtain the solution to the control problem, with the cost having been re-formulated in terms of the information-state. This derivation is fundamentally different to Whittle’s approach [Whittle 1981], being more closely linked to the work in Bensoussan and van Schuppen [1985].

An important feature of this chapter is that it presents a finite dimensional solution to the risk-sensitive output feedback control problem in the LQG case. It therefore provides a finite-dimensional example of the quite general infinite-dimensional controllers derived in James *et al.* [1994], and gives insight to the nonlinear control solution. The presentation of results for tracking

with hidden Markov model systems demonstrates a nonlinear situation where a finite dimensional information-state can be derived. The dynamic programming solution is infinite dimensional, but can be solved approximately, as discussed later in Section 6.3.4.

This work is, in part, an extension of the work presented in [Aggoun *et al.* 1994], where bi-linear systems were considered. In the work of this chapter, however, the tracking solution is discussed in addition to regulation, and the control solution is solved explicitly. Simulation studies are presented in an effort to demonstrate the effect of variations in the controller's sensitivity to risk. Various tracking problems are considered to show the advantages of the risk-sensitive approach.

Finally, the problem of risk-sensitive filtering is discussed. The filtering solution is shown to be derived directly from the tracking equations, with a slightly different interpretation of the cost function. Having demonstrated the close link between the information-state (used for control purposes) and the risk-sensitive filter, the idea of dual control can then be considered, as discussed in Section 6.6.

This chapter is organised as follows: Sections 6.2 and 6.3 present risk-sensitive tracking results for linear systems and HMMs respectively. Approaches for tackling constant reference inputs are considered in Section 6.4. In Section 6.5, simulation studies are presented, followed by a discussion on risk-sensitive filtering in Section 6.6. Finally, conclusions are presented in Section 6.7.

6.2 Linear Systems

This section considers the risk-sensitive tracking problem for discrete-time linear systems in Gaussian noise. Such problems are termed here *risk-sensitive-linear-quadratic-Gaussian* (RLQG) problems. The case of time-invariant systems is presented, however in this finite-time framework the result is equally applicable to time-varying systems.

6.2.1 State Space Model

Consider the following discrete-time system on the probability space (Ω, \mathcal{F}, P) with complete filtration $\{\mathcal{F}_k\}$:

$$\begin{aligned}x_{k+1} &= Ax_k + Bu_k + v_k, \\y_{k+1} &= Cx_k + w_k, \\z_{k+1} &= Dx_k,\end{aligned}\tag{6.1}$$

over the finite time interval $k = 0, 1, \dots, T$. The state of the system is represented by the process x . The observable part of the system is represented by the process y . This chapter will consider the problem of output tracking, where the desired trajectory will be denoted by \tilde{z} . The process which is to follow \tilde{z} is defined by z . The random variables v_k and w_k have normal densities $\psi \sim N(0, \Sigma)$ and $\phi \sim N(0, \Gamma)$ respectively, where Σ and Γ are $n \times n$ and $p \times p$ positive definite matrices. The control, u , takes values in \mathbb{R}^m . The complete filtration generated by (y_0, \dots, y_k) is denoted by \mathcal{Y}_k , and the admissible controls u are the set of \mathbb{R}^m -valued $\{\mathcal{Y}_k\}$ adapted processes. Also, write $U_{k,l}$ for the set of such control processes defined on the interval k, \dots, l .

In order to reformulate the system model (6.1), a new probability measure, \bar{P} , can be defined by setting

$$\Lambda_{0,k} = \frac{dP}{d\bar{P}} \Big|_{\mathcal{F}_k} = \prod_{\ell=1}^k \lambda_\ell, \quad \text{where } \lambda_k = \frac{\psi(x_k - Ax_{k-1} - Bu_{k-1})\phi(y_k - Cx_{k-1})}{\psi(x_k)\phi(y_k)}.\tag{6.2}$$

Here, $\Lambda_{0,k}$ is an \mathcal{F}_k martingale, and $E[\Lambda_{0,k}] = 1$. Now, under \bar{P} , x_k and y_k are two sequences of independent, normally distributed random variables with densities ψ and ϕ respectively. This reformulated model results in a linear recursion for the un-normalised information-state, as seen later in Section 6.2.3.

6.2.2 Cost

The cost function for the risk-sensitive control problem is given, for any admissible control $u \in U_{0,T-1}$, by

$$J(u) = E \left[\exp \theta \left\{ \Psi_{0,T-1} + \frac{1}{2} x_T' M_T x_T \right\} \right]\tag{6.3}$$

$$= \bar{E} \left[\Lambda_{0,T} \exp \theta \left\{ \Psi_{0,T-1} + \frac{1}{2} x_T' M_T x_T \right\} \right],\tag{6.4}$$

where

$$\Psi_{j,k} \triangleq \sum_{\ell=j}^k \frac{1}{2} [x'_\ell M x_\ell + u'_\ell N u_\ell + (\tilde{z}_{\ell+1} - D x_\ell)' Q (\tilde{z}_{\ell+1} - D x_\ell)] . \quad (6.5)$$

Here, $\theta > 0$ is a real number and represents the amount of risk in the control policy. For small values of θ , approaching zero, the effect is to make control decisions assuming the stochastic disturbances are acting in an average manner. For larger values of θ , the control is effectively more conservative, or in other words, has a higher sensitivity to risk.

It can be shown [James *et al.* 1994] that in the limit when θ approaches zero, the cost function (6.4) is identical to the more familiar cost function considered for LQG control [Anderson and Moore 1989]. This implies that the LQG solution will be generated by simply setting θ equal to zero in the final RLQG equations (as demonstrated in the following sections).

6.2.3 Information State

In this section finite dimensional recursions are presented for the *information-state*. In the case of risk-sensitive control, it is convenient to also include a component of the cost in the information-state. This is an important concept and is a more general type of information-state than those used in previous chapters. For the formulation presented here, the information-state is again a probability distribution (it can be compared to the ‘past stress’ in Whittle [1981]). For small values of θ , approaching zero, the mean and variance of the information-state become the state and covariance estimates for the linear Kalman filter. This can be seen in the following definition by setting θ equal to zero.

For any admissible control u , consider the measure

$$\alpha_k(x) dx \triangleq \overline{E}[\Lambda_{0,k} \exp(\theta \Psi_{0,k-1}) I(x_k \in dx) | \mathcal{Y}_k] , \quad (6.6)$$

where $I(\cdot)$ is the indicator function.

Lemma 6.1 *The information-state $\alpha_k(x)$, as defined in (6.6), obeys the following recursion:*

$$\alpha_{k+1}(x) = \phi^{-1}(y_{k+1}) \int_{\mathbf{R}^n} \phi(y_{k+1} - C\xi) \exp(\theta \Psi_{k,k}) \psi(x - A\xi - B u_k) \alpha_k(\xi) d\xi . \quad (6.7)$$

Proof :

$$\begin{aligned}
\alpha_{k+1}(x)dx &= \overline{E}[\Lambda_{0,k+1} \exp(\theta\Psi_{0,k}) I(x_{k+1} \in dx) | \mathcal{Y}_{k+1}] \\
&= \overline{E}[\lambda_{k+1} \Lambda_{0,k} \exp(\theta\Psi_{k,k}) \exp(\theta\Psi_{0,k-1}) I(x_{k+1} \in dx) | \mathcal{Y}_{k+1}], \\
\alpha_{k+1}(x) &= \int_{\mathbf{R}^n} \dots \int_{\mathbf{R}^n} \frac{\phi(y_{k+1} - Cx_k)}{\phi(y_{k+1})} \exp(\theta\Psi_{k,k}) \psi(x - Ax_k - Bu_k) \\
&\quad \Lambda_{0,k} \exp(\theta\Psi_{0,k-1}) d\overline{P}(x_0, \dots, x_k) \\
&= \phi^{-1}(y_{k+1}) \int_{\mathbf{R}^n} \phi(y_{k+1} - C\xi) \exp(\theta\Psi_{k,k}) \psi(x - A\xi - Bu_k) \alpha_k(\xi) d\xi.
\end{aligned}$$

■

Theorem 6.1 *The information-state $\alpha_k(x)$ is an un-normalised Gaussian density given by*

$$\alpha_k(x) = \alpha_k(x, \chi_k) = Z_k \exp(-1/2)[(x - \mu_k)' R_k^{-1} (x - \mu_k)], \quad (6.8)$$

where the new information-state $\chi_k \triangleq (\mu_k, R_k, Z_k)$, and μ_k , R_k^{-1} and Z_k are given by the following algebraic recursions:

$$\begin{aligned}
\mu_{k+1} &= R_{k+1} [\Sigma^{-1} Bu_k \\
&\quad + \Sigma^{-1} A a_k^{-1} (R_k^{-1} \mu_k - A' \Sigma^{-1} Bu_k + C' \Gamma^{-1} y_{k+1} - \theta D' Q \tilde{z}_{k+1})] \\
R_{k+1}^{-1} &= \Sigma^{-1} - \Sigma^{-1} A a_k^{-1} A' \Sigma^{-1} \\
Z_{k+1} &= Z_k |\Sigma|^{-\frac{1}{2}} |a_k|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\right) [\gamma_k - \mu'_{k+1} R_{k+1}^{-1} \mu_{k+1}]
\end{aligned} \quad (6.9)$$

where

$$a_k = C' \Gamma^{-1} C - \theta(M + D' Q D) + A' \Sigma^{-1} A + R_k^{-1}, \quad (6.10)$$

$$\begin{aligned}
\gamma_k &= u'_k (-\theta N + B' \Sigma^{-1} B) u_k + \mu'_k R_k^{-1} \mu_k - \theta \tilde{z}'_{k+1} Q \tilde{z}_{k+1} \\
&\quad - (\mu'_k R_k^{-1} - u'_k B' \Sigma^{-1} A + y'_{k+1} \Gamma^{-1} C - \theta \tilde{z}'_{k+1} Q D) a_k^{-1} \\
&\quad (R_k^{-1} \mu_k - A' \Sigma^{-1} Bu_k + C' \Gamma^{-1} y_{k+1} - \theta D' Q \tilde{z}_{k+1}), \quad (6.11)
\end{aligned}$$

under the condition that a_k and R_k be positive definite for all k .

Proof : Due to the linearity of the dynamics, and the fact that v_k and w_k are independent and normally distributed, it can be seen that $\alpha_k(x)$ is an un-normalised Gaussian density. The recursions for μ_k , R_k^{-1} and Z_k are obtained by evaluating the integral in (6.7). An outline of the derivation is given in Appendix B. ■

Further matrix manipulations yield the following, more familiar, expressions:

$$\begin{aligned}
 \mu_{k+1} &= A\mu_k + Bu_k \\
 &\quad + A\tilde{K}_k [C'\Gamma^{-1}(y_{k+1} - C\mu_k - \theta\Gamma(C')^{-1}D'Q\tilde{z}_{k+1}) + \theta(M + D'QD)\mu_k] \\
 \tilde{K}_k &\triangleq (R_k^{-1} + C'\Gamma^{-1}C - \theta(M + D'QD))^{-1} \\
 R_{k+1} &= \Sigma + A\tilde{K}_k A'
 \end{aligned}
 \tag{6.12}$$

which can be compared to the result presented in Whittle [1981] for the case where $Q = 0$.

Limit Result

Equations (6.12) can be re-expressed in the following form:

$$\begin{aligned}
 \mu_{k+1} &= A\mu_{k|k} + Bu_k \\
 \mu_{k|k} &\triangleq \mu_k + K_k [y_{k+1} - C\mu_k - \theta(\Gamma(C')^{-1}D'Q\tilde{z}_{k+1} - \Gamma(C')^{-1}(M + D'QD)\mu_k)] \\
 K_k &\triangleq (R_k^{-1} - \theta(M + D'QD))^{-1}C'[C(R_k^{-1} - \theta(M + D'QD))^{-1}C' + \Gamma]^{-1} \\
 R_{k+1} &= \Sigma + AR_{k|k}A' \\
 R_{k|k} &\triangleq R_k - K_k C R_k
 \end{aligned}
 \tag{6.13}$$

In the case when θ approaches zero, it can easily be seen that the equations in (6.13) reduce to the standard Kalman filter equations [Anderson and Moore 1979] (p. 40).

6.2.4 Alternate Cost Representation

This section shows that the cost function can be expressed in terms of the information-state. This allows the optimisation problem to be solved by dynamic programming, without any appeal to a certainty equivalence principle. It can however be viewed as a slightly different form of the separation principle ([Anderson and Moore 1989] p. 218), since due to the fact that the cost is

actually in terms of the information-state, there is no need to evaluate the state feedback solution and then substitute the state estimate. The system is not however fully separated because the information-state, as defined in (6.6), contains terms from the control cost function.

Theorem 6.2 For any admissible control u , the risk sensitive cost can be expressed in the form

$$J(u) = \bar{E} [\langle \alpha_T(\cdot, \chi_T), \beta_T \rangle] , \quad (6.14)$$

where $\langle f(\cdot), q(\cdot) \rangle = \int_{\mathbf{R}^n} f(z)q(z)dz$ and $\beta_T(x) \triangleq \exp(\frac{\theta}{2}x'M_Tx)$.

Proof : Now, from (6.4) it is possible to see that

$$\begin{aligned} J(u) &= \bar{E} \left[\Lambda_{0,T} \exp(\theta\Psi_{0,T-1}) \exp(\frac{\theta}{2}x'_T M_T x_T) \right] \\ &= \bar{E} \left[\bar{E} [\Lambda_{0,T} \exp(\theta\Psi_{0,T-1}) \beta_T(x_T) | \mathcal{Y}_T] \right] \\ &= \bar{E} \left[\int_{\mathbf{R}^n} \beta_T(x) \alpha_T(x) dx \right] \\ &= \bar{E} [\langle \alpha_T(\cdot, \chi_T), \beta_T \rangle] . \end{aligned}$$

■

6.2.5 Dynamic Programming

Following James *et al.* [1994] it is possible to see that the alternative control problem can be solved using dynamic programming. Suppose that at some time k , $0 < k < T$, the information-state χ_k is $\chi = (\mu, R, Z)$.

The value function for this control problem is [Aggoun *et al.* 1994, James *et al.* 1994]:

$$V(\chi, k) = \inf_{u \in \mathcal{U}_{k,T-1}} \bar{E} [\langle \alpha_k, \beta_k \rangle | \alpha_k = \alpha(\chi)] , \quad (6.15)$$

where β_k is an adjoint process defined, for $k \leq T - 1$, by

$$\beta_k(x) = \bar{E} [\Lambda_{k+1,T} \exp(\theta\Psi_{k,T-1}) \exp(\frac{\theta}{2}x'_T M_T x_T) | x_k = x, \mathcal{Y}_T] . \quad (6.16)$$

The adjoint process is different to the 'future stress' in [Whittle 1981], as it relates to output feedback, not state feedback.

Theorem 6.3 [Aggoun et al. 1994, James et al. 1994] *The value function satisfies the recursion*

$$V(\chi, k) = \inf_{u \in U_{k,k}} \bar{E}[V(\chi_{k+1}(\chi_k, u, y_{k+1}), k+1) | \chi_k = \chi] \quad (6.17)$$

and $V(\chi, T) = \langle \alpha_T(\cdot, \chi), \beta_T \rangle$. ■

6.2.6 Dynamic Programming Solution

Theorem 6.4 *The value function is the exponential of a quadratic in μ ,*

$$V(\chi, k) = Z_k \exp(\theta/2) [\mu'_k S_k^a \mu_k + 2S_k^b \mu_k + S_k^c], \quad (6.18)$$

and the optimal control is linear in μ ,

$$\boxed{u_k^{min} = -(N + B' \tilde{S}_{k+1} B)^{-1} B' [\tilde{S}_{k+1} \tilde{A} \mu_k + S_{k+1}^b + \theta K_k^b]}, \quad (6.19)$$

where

$$\begin{aligned} \tilde{S}_{k+1} &= ((S_{k+1}^a)^{-1} - \theta \tilde{\Gamma}_k \tilde{\Gamma}'_k)^{-1}, & \tilde{A} &= A \rho^{-1}, \\ K_k^a &= (N + B' \tilde{S}_{k+1} B)^{-1} B' \tilde{S}_{k+1} \tilde{A}, & \tilde{M} &= (M + D' Q D) \rho^{-1}, \\ K_k^b &= S_{k+1}^a \tilde{\Gamma}_k \delta^{-1} (\tilde{\Gamma}'_k S_{k+1}^b - \Theta D' Q \tilde{z}_{k+1}), & \sigma &= N + B' \tilde{S}_{k+1} B, \\ \tilde{\Gamma}_k &= A \tilde{K}_k C' \Gamma^{-1} C \Theta, & \delta &= I - \theta \tilde{\Gamma}'_k S_{k+1}^a \tilde{\Gamma}_k, \\ \Theta &= [(C' \Gamma C)^{-1} + \rho^{-1} R_k]^{1/2}, & \rho &= I - \theta R_k (M + D' Q D). \end{aligned}$$

Also, S_k^a and S_k^b are given by the following backwards recursions:

$$\boxed{\begin{aligned} S_k^a &= \tilde{M} + \tilde{A}' S_{k+1}^a (I + B N^{-1} B' S_{k+1}^a - \theta \tilde{\Gamma}_k \tilde{\Gamma}'_k S_{k+1}^a)^{-1} \tilde{A} \\ S_k^b &= \tilde{A}' S_{k+1}^b - (I + \theta \tilde{M} R_k) D' Q \tilde{z}_{k+1} + \tilde{A}' \tilde{S}_{k+1} B \sigma^{-1} B' (S_{k+1}^b + \theta K_k^b) + \theta \tilde{A}' K_k^b \end{aligned}} \quad (6.20)$$

under the condition that $(I - \theta \tilde{\Gamma}'_k S_{k+1}^a \tilde{\Gamma}_k)$ is positive definite for all k , and C is positive definite except in the cases where $C = D$ or $D = 0$.

Proof : By evaluating the dynamic programming equation (6.17) for $V(\chi, T - 1)$ it can be seen that the value function is the exponential of a quadratic in μ . The remainder of the proof is outlined in Appendix B, and is essentially an evaluation of the dynamic programming equation (6.17), with appropriate variable transformations. ■

Remark 6.1: The condition that C be positive definite, is a manifestation of the variable transformation used in order to present the results in a form which more readily demonstrates the link to standard LQG results. As can be seen from the exclusion when $D = 0$, the condition only applies to the tracking part of the solution, (that is the S_k^b and K_k^b recursions). It is possible to solve the dynamic programming problem without such a variable transformation and thus remove the condition on C . □

In order to demonstrate consistency with the results presented in Whittle [1981], where an appeal was made to a certainty equivalence principle, and $Q \equiv 0$, we now set

$$\Pi_k = S_k^a [I + \theta R_k S_k^a]^{-1}, \quad (6.21)$$

and $Q = 0$, which results in the following recursion for Π_k

$$\Pi_k = M + A' [\Pi_{k+1}^{-1} + B N^{-1} B' - \theta \Sigma]^{-1} A \quad (6.22)$$

under the condition that $(I - \theta R_k \Pi_k)$ is positive definite for all k .

Substitution of (6.22) into (6.19), yields

$$u_k^{min} = -N^{-1} B' (\Pi_{k+1}^{-1} + B N^{-1} B' - \theta \Sigma)^{-1} A [I - \theta R_k \Pi_k]^{-1} \mu_k \quad (6.23)$$

where the term $[I - \theta R_k \Pi_k]^{-1} \mu_k$ is sometimes referred to as the minimum stress estimate.

Limit Result

In the case where θ approaches zero, it can easily be seen that \tilde{S}_{k+1} , \tilde{A} and \tilde{M} approach S_{k+1}^a , A and $M + D'QD$ respectively, and the following equations result from manipulations to (6.19) and (6.20):

$$\begin{aligned}
u_k^{min} &= -(N + B'S_{k+1}^a B)^{-1} B' [S_{k+1}^a A \mu_k + S_{k+1}^b] \\
S_k^a &= M + D' Q D + A' [S_{k+1}^a - S_{k+1}^a B (N + B'S_{k+1}^a B)^{-1} B' S_{k+1}^a] A \\
S_k^b &= (A - B(N + B'S_{k+1}^a B)^{-1} B' S_{k+1}^a A)' S_{k+1}^b - D' Q \tilde{z}_{k+1}
\end{aligned} \tag{6.24}$$

These are the standard LQG equations, as presented for example in Anderson and Moore [1989] (p. 32 and p. 81).

6.3 Hidden Markov Models

In this section the risk-sensitive tracking result is presented for hidden Markov models. Such systems are discrete time and have finite-discrete states. Consider the case of continuous valued observations.

6.3.1 State Space Model

Let X_k be a discrete-time homogeneous, first order Markov process belonging to a finite-discrete set. As in previous chapters, the state space of X_k , without loss of generality, can be identified with the set of unit vectors $\mathbf{S} = \{e_1, e_2, \dots, e_n\}$, $e_i = (0, \dots, 0, 1, 0, \dots, 0)' \in \mathbb{R}^n$ with 1 in the i^{th} position [Segall 1976]. Now consider that the process is defined on the probability space (Ω, \mathcal{F}, P) with complete filtration $\{\mathcal{F}_k\}$. Suppose, in this control setting, there is a family of generators $\mathbf{A}(u) = (a_{ij}(u))$, $1 \leq i, j \leq n$ where $a_{ij}(u) = P(X_{k+1} = e_j \mid X_k = e_i, u)$ so that $E[X_{k+1} \mid X_k, u] = \mathbf{A}'(u)X_k$. These generators depend on the admissible controls, u . Of course $a_{ij}(u) \geq 0$, $\sum_{j=1}^n a_{ij}(u) = 1$, for each i . In this chapter consider the case of continuous valued observations y_k , and desired trajectories \tilde{z}_k . The state space model for the HMM is given by

$$\begin{aligned}
X_{k+1} &= \mathbf{A}'(u) X_k + M_{k+1}, \\
y_k &= c(X_k) + w_k, \\
z_k &= d(X_k),
\end{aligned} \tag{6.25}$$

where M_{k+1} is a $(\mathbf{A}(u), \mathcal{F}_k)$ martingale increment, in that $E[M_{k+1} \mid \mathcal{F}_k] = 0$. Also, the random variable w_k has normal density $\phi \sim N(0, \Gamma)$, where Γ is a $p \times p$ positive definite matrix.

In order to reformulate the system model (6.25), a new probability measure, \bar{P} , can be defined by setting

$$\Lambda_{0,k} = \frac{dP}{d\bar{P}} \Big|_{\mathcal{F}_k} = \prod_{\ell=1}^k \lambda_{\ell}, \quad \text{where } \lambda_k = \frac{\phi(y_k - c(X_k))}{\phi(y_k)}. \quad (6.26)$$

Here, $\Lambda_{0,k}$ is an \mathcal{F}_k martingale and $E[\Lambda_{0,k}] = 1$. Now under \bar{P} , y_k is a sequence of independent, normally distributed random variables with density ϕ . This reformulated model results in a linear recursion for the un-normalised information-state, as in Section 6.3.3.

6.3.2 Cost

The cost function for the risk-sensitive control problem for HMMs is given, for any admissible control $u \in U_{0,T-1}$, by

$$J(u) = E \left[\exp \theta \left\{ \Psi_{0,T-1} + \frac{1}{2} X_T' M_T X_T \right\} \right] \quad (6.27)$$

$$= \bar{E} \left[\Lambda_{0,T} \exp \theta \left\{ \Psi_{0,T-1} + \frac{1}{2} X_T' M_T X_T \right\} \right], \quad (6.28)$$

where

$$\Psi_{j,k} \triangleq \sum_{\ell=j}^k \frac{1}{2} [X_{\ell}' M X_{\ell} + u_{\ell}' N u_{\ell} + (\bar{z}_{\ell} - d(X_{\ell}))' Q (\bar{z}_{\ell} - d(X_{\ell}))]. \quad (6.29)$$

In this case, where X_k is an indicator vector, the cost has a connotation slightly different to the minimum variance controller of Section 6.2. However, consider a limiting example where the set of real numbers is quantised into an infinite set of Markov states. In this case, X_k is infinite dimensional. If $M = C'C$ in (6.29) then the result is a minimum variance controller. This infinite dimensional example gives some insight into the motivation behind considering (6.29) as the cost function in this HMM problem.

6.3.3 Information State

As in Section 6.2.3, an information-state is presented which includes a component of the cost. Unlike the linear case, however, for HMMs the information-state is a probability distribution *vector* of dimension n .

For any admissible control u , consider the measure

$$\alpha_k(e_i) \triangleq \bar{E}[\Lambda_{0,k} \exp(\theta \Psi_{0,k-1}) \langle X_k, e_i \rangle | \mathcal{Y}_k]. \quad (6.30)$$

Theorem 6.5 The information-state $\alpha_k = (\alpha_k(e_1), \dots, \alpha_k(e_n))'$, as defined in (6.30), obeys the following recursion

$$\boxed{\alpha_{k+1} = \mathcal{B}_k \mathbf{A}'(u) \mathcal{D}_k \alpha_k} \quad (6.31)$$

where

$$\mathcal{B}_k = \text{diag} \left(\frac{\psi(y_{k+1} - c(e_1))}{\psi(y_{k+1})}, \dots, \frac{\psi(y_{k+1} - c(e_n))}{\psi(y_{k+1})} \right), \quad (6.32)$$

$$\mathcal{D}_k = \text{diag} \left(\exp \frac{\theta}{2} [e_1' M e_1 + u_k' N u_k + (\tilde{z}_k - d(e_1))' Q (\tilde{z}_k - d(e_1))], \dots, \exp \frac{\theta}{2} [e_n' M e_n + u_k' N u_k + (\tilde{z}_k - d(e_n))' Q (\tilde{z}_k - d(e_n))] \right) \quad (6.33)$$

Proof :

$$\begin{aligned} \alpha_{k+1}(e_i) &= \bar{E}[\Lambda_{0,k+1} \exp(\theta \Psi_{0,k}) \langle X_{k+1}, e_i \rangle | \mathcal{Y}_{k+1}] \\ &= \bar{E}[\lambda_{k+1} \Lambda_{0,k} \exp(\theta \Psi_{k,k}) \exp(\theta \Psi_{0,k-1}) X_k' \mathbf{A}(u) e_i | \mathcal{Y}_{k+1}] \\ &= \bar{E} \left[\frac{\psi(y_{k+1} - c(e_i))}{\psi(y_{k+1})} \right. \\ &\quad \left. \left(\sum_{j=1}^n \langle X_k, e_j \rangle \exp \frac{\theta}{2} [e_j' M e_j + u_k' N u_k + (\tilde{z}_k - d(e_j))' Q (\tilde{z}_k - d(e_j))] \right) \right. \\ &\quad \left. \left(\sum_{j=1}^n a_{ji}(u) \langle X_k, e_j \rangle \right) \Lambda_{0,k} \exp(\theta \Psi_{0,k-1}) | \mathcal{Y}_{k+1} \right] \\ &= \frac{\psi(y_{k+1} - c(e_i))}{\psi(y_{k+1})} \sum_{j=1}^n \exp \frac{\theta}{2} [e_j' M e_j + u_k' N u_k + (\tilde{z}_k - d(e_j))' Q (\tilde{z}_k - d(e_j))] \\ &\quad a_{ji}(u) \alpha_k(e_j). \end{aligned}$$

Writing this in matrix notation gives the result. ■

6.3.4 Alternate Cost and Dynamic Programming

For the HMM case, the cost (6.28) can be expressed in a separated form, as for the linear case in Theorem 6.2, with appropriate notational changes. The dynamic programming solution is likewise obtained from Theorem 6.3. Unfortunately for the HMM, the solution to the dynamic programming equation cannot be evaluated in terms of finite-dimensional Riccati equations (as is the case for linear systems in Section 6.2.6). The solution for the HMM system requires a search over all possible control values for each backwards step, and for each possible value of the information-state. Therefore while the HMM case results in a finite dimensional information-state, it unfortunately has an infinite dimensional solution to the dynamic programming problem. It is

possible however to solve an approximate dynamic programming problem by making practical numerical approximations such as quantising the information-state space. This can be computationally feasible in some cases since the information-state is known to have positive elements. Also, a normalised version of the information-state can be used in the Dynamic programming problem since the following property is known to hold in the HMM case [James and Elliott n.d.] :

$$V(c\alpha, k) = cV(\alpha, k) . \quad (6.34)$$

In Section 6.5 an example of such an approximate dynamic programming solution is presented for this risk-sensitive HMM case.

6.4 Constant Reference Input Case

This section investigates the case where \tilde{z}_k is a constant value. Under such conditions it is possible to design an optimal controller with zero steady state error. This section considers the discrete-time linear system of Section 6.2.

Consider the cost function given in (6.4). Note that for this general function there exist some trade-offs which do not allow zero steady state error to be achieved. For example, (6.5) penalises deviations of x_k from zero, while at the same time penalising deviations of Dx_k from \tilde{z}_k , these are conflicting objectives. Also, the control u_k is penalised for deviations from zero when it is known that, in steady state, it must be a constant non-zero value for this constant reference input case. These considerations indicate that the tracking problem must be reformulated.

6.4.1 Control Integrator Approach

A standard method for obtaining zero steady state error, is to introduce an integrator in the forward path of the control loop. This technique can be used in this risk-sensitive case with a few minor adjustments to the control policy. Figure 6.1 shows the block diagram for the control system presented in the preceding sections of this chapter. By introducing an integrator and augmenting the state, as in Figure 6.2, it is possible to obtain a more appropriate cost function.

In this section the control u_k is chosen to be an extra state, and a new control \tilde{u}_k is defined. The state of the augmented system is then given by

$$\tilde{x}_k = \begin{pmatrix} x_k \\ x_k^c \end{pmatrix}, \quad (6.35)$$

where $x_k^c = u_k = \sum_{i=1}^k \tilde{u}_i$. This augmented state is an un-normalised Gaussian density, and is given from equations (6.9) to (6.11), with appropriate re-definitions for the augmented system, by

$$\begin{aligned} \tilde{x}_{k+1} &= \begin{pmatrix} A & B \\ 0 & I \end{pmatrix} \tilde{x}_k + \begin{pmatrix} 0 \\ I \end{pmatrix} \tilde{u}_k + \begin{pmatrix} I \\ 0 \end{pmatrix} v_k, \\ y_{k+1} &= (C \ 0) \tilde{x}_k + w_k, \\ z_{k+1} &= (D \ 0) \tilde{x}_k. \end{aligned} \quad (6.36)$$

The cost function to be considered is now given, for $\tilde{z}_k = \tilde{z}$, by (6.4) where $\Psi_{j,k}$ is re-defined as

$$\Psi_{j,k} \triangleq \sum_{\ell=j}^k \frac{1}{2} [\tilde{u}'_{\ell} N \tilde{u}_{\ell} + (\tilde{z} - (D \ 0) \tilde{x}_{\ell})' Q (\tilde{z} - (D \ 0) \tilde{x}_{\ell})]. \quad (6.37)$$

It can easily be seen that for this cost function there are no conflicting objectives, and as such zero steady state error can be achieved.

Unfortunately, however, there exist some hidden problems. The first is that the new state \tilde{x}_k has zero state noise and as such results in a singular filtering problem. This can be overcome by assuming there exists some noise of variance ϵ and then taking the limit of the information-state, as ϵ approaches zero. As can be seen from (6.12), the limit exists with \tilde{K}_k re-defined as

$$\tilde{K}_k \triangleq R_k (I + C' \Gamma^{-1} C R_k - \theta (M + D' Q D) R_k)^{-1}. \quad (6.38)$$

The second problem is that in the case of modelling errors, even with the augmented system, zero steady state error is not necessarily achieved. This is due to the fact that there exists a term in the optimal control law (6.19), which is not proportional to the state estimate μ_k . If this term is not calculated correctly, as would be the case with modelling errors, then the control \tilde{u}_k would drive the output z_k to an incorrect steady state value. Since the observations, y_k , are being fed back, not the tracking value, z_k , there is no non-zero error term to drive the controller to zero steady state offset.

Although zero steady state error may not be achieved in certain cases, there is still an advantage to applying the integrator approach. In the risk-sensitive case, when modelling errors are present, it is possible to achieve a lower minimum variance cost than for the case of LQG control (as can be seen in Section 6.5). One problem however is that the step response can be undesirable for values of θ which are too large. By augmenting the system with an integrator, the step response will be smoothed out, resulting in a risk-sensitive policy which has both a lower minimum variance cost and an acceptable step response.

6.4.2 Reference Model Integrator Approach

An extra point to note is that in the scheme presented so far, it is necessary to have prior knowledge of the constant reference input signal, \bar{z} . An approach for removing this assumption, commonly used in LQG tracking systems, is to model the reference \bar{z} by a first order integrator,

$$\begin{aligned} x_{k+1}^r &= x_k^r + v_k^r, \\ \tilde{z}_{k+1} &= \tilde{D}x_k^r. \end{aligned} \quad (6.39)$$

This would of course slow the response of the system, but would have the advantage of zero-steady state error in conditions of uncertain models. The augmented state vector is given by

$$\tilde{x}_k = \begin{pmatrix} x_k \\ x_k^c \\ x_k^r \end{pmatrix}. \quad (6.40)$$

This new augmented state is again an un-normalised Gaussian density, and is given from equations (6.9) to (6.11), with appropriate re-definitions for the augmented system, by

$$\begin{aligned} \tilde{x}_{k+1} &= \begin{pmatrix} A & B & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix} \tilde{x}_k + \begin{pmatrix} 0 \\ I \\ 0 \end{pmatrix} \tilde{u}_k + \begin{pmatrix} v_k \\ 0 \\ v_k^r \end{pmatrix}, \\ y_{k+1} &= (C \ 0 \ -\tilde{D}) \tilde{x}_k + w_k, \\ z_{k+1} &= (D \ 0 \ -\tilde{D}) \tilde{x}_k. \end{aligned} \quad (6.41)$$

The cost function to be considered is given by (6.4) where $\Psi_{j,k}$ is re-defined as

$$\Psi_{j,k} \triangleq \sum_{\ell=j}^k \frac{1}{2} \left[\tilde{u}'_{\ell} N \tilde{u}_{\ell} + \tilde{x}'_{\ell} (D \ 0 \ -\tilde{D})' Q (D \ 0 \ -\tilde{D}) \tilde{x}_{\ell} \right]. \quad (6.42)$$

In this case there are no terms in (6.42) which are linear in \tilde{x}_k , and as such the optimal control will be proportional to the state estimate μ_k and have no extra terms (that is, S_k^b and K_k^b will not appear). In fact the solution to the dynamic programming problem for this augmented system is given in Theorem 6.4 with the following substitutions (in addition to those for the augmented system representation (6.41)):

$$\tilde{M} = (D \ 0 \ -\tilde{D})' Q (D \ 0 \ -\tilde{D}) , \quad \tilde{Q} = 0 . \quad (6.43)$$

Due to the purely proportional feedback nature of this solution, it can now be seen that it is possible to obtain zero steady state error even in the case of modelling errors, as there is no longer a constant offset term contributed by K_k^b . Unfortunately, however, the initial transient will suffer due to the fact that the controller is no longer able to anticipate the step in the reference input, as it is now assumed to be unknown.

One final point to note is that this second augmentation can be used without the first augmentation, and zero steady state error will result for the case where $N = 0$, (this is termed *cheap control*). However, undesirable oscillations in the transient response will increase, compared to the situation where an integrator is present in the forward path.

6.5 Simulation Studies

Simulation studies are now presented to demonstrate the effect of variations to the risk-sensitive parameter θ .

Example 6.1: This example demonstrates a case where modelling errors are present. The true system is given by the following parameters:

$$\begin{aligned} A &= \begin{bmatrix} -0.2 & 1 \\ -0.2 & 0 \end{bmatrix}, & B &= \begin{bmatrix} 0.9 \\ -0.6 \end{bmatrix}, & C &= [1 \ 0], & M &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \\ T &= 100, & \Sigma &= \begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix}, & N &= 0.1, & \Gamma &= 0.01, \\ & & & & Q &= 100, & & \end{aligned}$$

and the trajectory to be followed, \tilde{z}_k , is a unit step at $k = 20$. The modelling error is introduced

by assuming in the design that A is given by

$$A = \begin{bmatrix} -0.8 & 1 \\ -0.8 & 0 \end{bmatrix}.$$

Table 6.1 gives values of the LQG, minimum variance, cost function (that is, $\Psi_{0,T-1} + 0.5x_T' M_T x_T$) averaged over 100 simulation runs. It can be seen that in the case where no modelling error is present, of course $\theta = 0$ gives lowest cost. However, when the error is introduced, a higher value of θ gives a lower minimum variance cost. This example displays an advantage of the risk-sensitive approach in the presence of modelling errors.

Unfortunately, the sample path properties may not improve with a lower minimum variance cost, as one would wish, especially if θ is too large. Here, *too large* will depend on the type of modelling error, and will of course be unknown to the designer. Figure 6.3 shows a typical sample run for the case of no modelling errors. It shows that the cost function chosen, for the tracking task considered, results in little difference in tracking errors between the LQG and risk-sensitive policies. Figure 6.4 shows a typical sample run for the case where modelling errors are present. Even though the minimum variance cost is lower for the risk-sensitive policy, the tracking performance might not be as desirable, having much greater oscillations in the transient response. Therefore the desirability of a risk-sensitive approach cannot be measured purely by the minimum variance cost.

Example 6.2: This example demonstrates the case of a constant reference input, where an integrator is added in the forward path of the control design. This is done in an effort to combat the undesired oscillations in the transient response, seen in the previous example. The system is the same as in Example 6.1, but with $M = \underline{0}_{3 \times 3}$. When no modelling errors are present, zero steady state is achieved. When errors are introduced to the model it is not possible to have zero steady state error, however Figures 6.5 and 6.6 demonstrate that there are still advantages to the integrator approach. The modelling error in these figures is the same as that in Example 6.1. As can be seen from the figures, the addition of an integrator effectively increases the usable range of risk-sensitive parameter values, θ , by smoothing the step response. However, there is less cost benefit from varying θ .

Example 6.3: This example demonstrates the case of a constant reference input, where an integrator is added in the forward path of the control design as well as using a model for the reference. As can be seen from Figure 6.7, zero steady state error is achieved for both modelling

errors and no modelling errors, however it is at the expense of the speed of transient response. The other point to note from this example is that the benefit from a risk-sensitive control policy is reduced when integrators are added. This is due to the fact that the LQG cost function is much smaller when zero steady state error is achieved, compared to when it is not achieved. The result is that the effect of varying θ is less, as the variation is over a less steep region of the exponential curve.

It should also be noted that it is only for certain types of modelling error, that a cost benefit is derived from a risk-sensitive approach. In many cases the LQG cost is either larger for a risk sensitive controller, or it is only smaller for a certain range of parameter values. Therefore, the desirability of a risk sensitive control policy must be determined via simulation studies, from case to case, due to its dependency on the type of modelling error present.

One final point to note is that the LQG design is not optimal for the doubly augmented system, due to the fact that the reference is modelled by an integrator (6.39) when in fact it is a deterministic signal. This is effectively an unavoidable modelling error due to the augmented design, and as such the LQG solution is not optimal with this augmentation. In fact, for the example considered, the cost is actually less for the risk-sensitive solution than for the LQG solution, even in the case where A is known precisely (that is, the line labelled 'no modelling error' in Figure 6.7).

Example 6.4: This simulation study presents an example of an approximate solution to the risk-sensitive HMM control problem of Section 6.3. The system is given by the following parameters

$$A = \begin{bmatrix} 0.1 & 0.9 \\ 0.9 & 0.1 \end{bmatrix} + \begin{bmatrix} u & -u \\ -u & u \end{bmatrix}, \quad c(X_k) = \left\langle \begin{bmatrix} 0 \\ 1 \end{bmatrix}, X_k \right\rangle, \quad d(X_k) = \left\langle \begin{bmatrix} 0 \\ 1 \end{bmatrix}, X_k \right\rangle,$$

$$M = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad N = 0.1, \quad \Gamma = 2,$$

$$Q = 0, \quad T = 100,$$

and in this example there is no trajectory to be followed, (that is, $\tilde{z}_k = 0$). Note from this system that in an uncontrolled situation, the output trajectory, z_k , will tend to oscillate between the values 0 and 1 at each discrete time instant. From the definition of M it can be seen that the control objective is to force z_k to the value zero.

In this example the dynamic programming problem is solved by quantising the normalised information-state, α_k^q , into six discrete values,

$$\alpha_k^q \in \{(\ell \times 0.2, m \times 0.2)'\}, \quad 0 \leq \ell, m \leq 5, \quad \ell + m = 5, \quad (6.44)$$

and allowing only three possible control values

$$u_k \in \{0.1, 0.3, 0.5\}. \quad (6.45)$$

The dynamic programming problem is then solved by evaluating the cost which minimises the value function for each possible information-state, α_k^q , at each step backwards in time, k .

In Table 6.2 the steady state control values which result from the approximate solution to the dynamic programming problem are presented. These values demonstrate the effect of the risk-sensitive parameter on the control policy. It can be seen that as the risk-sensitive parameter increases, the information-state must be increasingly more confident of the true state, before the controller is willing to apply a large control value. This example therefore demonstrates the robustness property gained from increasing the sensitivity to risk.

6.6 Risk-Sensitive Filtering Interpretations

In this section the risk-sensitive filtering problem is presented in order to demonstrate its connection with the control problem considered in the preceding sections.

6.6.1 Linear Systems

In this section the risk-sensitive filtering problem is shown to be solved by the same equations derived previously for the tracking problem, but with a slight reinterpretation of the cost function. These results for the linear filtering case have already been solved, without the control interpretation, in Speyer *et al.* [1992], and for the nonlinear case in Dey and Moore [1994].

To see the connection to the control problem, the risk-sensitive filtering cost function is now presented.

$$J_k(\hat{x}_k) = \overline{E} \left[\Lambda_{0,k} \exp \theta \Psi_{0,k}(\hat{x}_k) | \mathcal{Y}_k, \hat{\mathcal{X}}_{k-1} \right], \quad (6.46)$$

where

$$\Psi_{j,k}(\hat{x}_k) \triangleq \sum_{\ell=j}^k \frac{1}{2} [(x_\ell - \hat{x}_\ell)' Q (x_\ell - \hat{x}_\ell)] , \quad (6.47)$$

and \hat{x}_ℓ is the risk-sensitive state estimate of x_ℓ . Note here that the cost is an expectation conditioned on the set of observations. This is due to the fact that the filtering problem is one of optimisation in the forward direction (as opposed to the control problem which is an optimisation in the backwards direction), and as such, the previous observations will be available when the optimisation procedure is carried out at each iteration.

Comparing (6.46) and (6.47) to (6.4) and (6.5) it can easily be seen that (6.47) can be obtained from (6.5) by replacing $\tilde{z}_{\ell+1}$ by \hat{x}_ℓ , and setting $D \equiv I$ and $B \equiv M \equiv N \equiv 0$. Therefore, with the same definition of information-state given in (6.6), the result presented in Theorem 6.1 holds here as well, with the appropriate replacement of symbols. Equation (6.9) can be compared to the equations in Speyer *et al.* [1992] which differ only slightly since the prediction problem was considered, rather than the filtering solution presented here. From the information-state, α_k , the risk-sensitive state estimate, \hat{x}_k , is obtained, as in Dey and Moore [1994], by solving the forward optimisation task of the following theorem.

Theorem 6.6 *The risk-sensitive state estimate, \hat{x}_k , defined by*

$$\hat{x}_k = \underset{\zeta}{\operatorname{argmin}} J_k(\zeta) , \quad (6.48)$$

is given by the mean, μ_k , of the information-state, α_k , defined in (6.6), where $\Psi_{j,k}$ is given in (6.47), $\tilde{z}_{\ell+1}$ is replaced by \hat{x}_ℓ , $D \equiv I$, and $B \equiv M \equiv N \equiv 0$.

Proof :

$$\begin{aligned} \hat{x}_k &= \underset{\zeta}{\operatorname{argmin}} \bar{E} \left[\Lambda_{0,k} \exp \theta \Psi_{0,k}(\zeta) | \mathcal{Y}_k, \hat{\mathcal{X}}_{k-1} \right] \\ &= \underset{\zeta}{\operatorname{argmin}} \int_{\mathbf{R}^n} \dots \int_{\mathbf{R}^n} \exp \theta \Psi_{k,k}(\zeta) \Lambda_{0,k} \exp \theta \Psi_{0,k-1} d\bar{P}(x_0, \dots, x_k) \\ &= \underset{\zeta}{\operatorname{argmin}} \int_{\mathbf{R}^n} \exp \theta \Psi_{k,k}(\zeta) \alpha_k(z) dz \\ &= \underset{\zeta}{\operatorname{argmin}} \int_{\mathbf{R}^n} \exp \frac{\theta}{2} (z - \zeta)' Q (z - \zeta) Z_k \exp(-1/2) (z - \mu_k)' R_k^{-1} (z - \mu_k) dz \\ &= \underset{\zeta}{\operatorname{argmin}} C_1 \exp(1/2) \left\{ \theta \zeta' Q \zeta - \mu_k' R_k^{-1} \mu_k - (\mu_k' R_k^{-1} - \theta \zeta' Q) a^{-1} (R_k^{-1} \mu_k - \theta Q \zeta) \right\} \\ &= -(\theta Q - \theta^2 Q a^{-1} Q)^{-1} \theta Q a^{-1} R_k^{-1} \mu_k \\ &= -(a - \theta Q)^{-1} R_k^{-1} \mu_k \\ &= \mu_k , \end{aligned}$$

where $a = (\theta Q - R_k^{-1})$, C_1 is constant with respect to ζ , and a is positive definite. It can now be seen that the risk-sensitive state estimate is given by the mean, μ_k , of the information-state, α_k . ■

The recursion for this estimate is easily obtained by rearranging the expression for μ_k from (6.12), with the appropriate substitutions for the filtering problem. The following equation results:

$$\mu_{k+1} = A\mu_k + A(R_k^{-1} + C'\Gamma^{-1}C - \theta Q)^{-1}[C'\Gamma^{-1}(y_{k+1} - C\mu_k) - \theta(\mu_k - \hat{x}_k)]. \quad (6.49)$$

The final recursive equation is obtained by setting \hat{x}_k equal to μ_k , from Theorem 6.6 as follows:

$$\hat{x}_{k+1} = A\hat{x}_k + A(R_k^{-1} + C'\Gamma^{-1}C - \theta Q)^{-1}[C'\Gamma^{-1}(y_{k+1} - C\hat{x}_k)]. \quad (6.50)$$

This filtering result corresponds, as noted before, to the prediction result presented in Speyer *et al.* [1992].

6.6.2 Hidden Markov Models

In this section the control interpretation is used to derive the risk sensitive filtering results for HMMs. For such systems the risk-sensitive state estimate \hat{X}_k is obtained by solving the following optimisation task:

$$\hat{X}_k = \underset{\hat{X}}{\operatorname{argmin}} \bar{E}[\Lambda_{0,k} \exp \theta \Psi_{0,k} | \mathcal{Y}_k, \hat{X}_{k-1}], \quad (6.51)$$

where $\Psi_{0,k}$ is defined in (6.29), and α_k , \mathcal{B}_k , \mathbf{A} and \mathcal{D}_k are defined, as for the control problem, in (6.30) and Theorem 6.5, and all of these terms require the following re-definitions: replace \tilde{z}_k by \hat{X} , and set $d(\cdot) = (\cdot)$, and $B \equiv M \equiv N \equiv 0$, as in the previous section. The following theorem gives the solution for the risk-sensitive estimate \hat{X}_k .

Theorem 6.7 *The risk-sensitive estimate \hat{X}_k , defined in (6.51), is given by the following recursive equation:*

$$\hat{X}_k = e_{\hat{X}^*}, \text{ where } \hat{X}^* = \underset{\hat{X}}{\operatorname{argmin}} \langle \mathbf{1}, \mathcal{D}_k \alpha_k \rangle. \quad (6.52)$$

with the above definitions of α_k and \mathcal{D}_k .

Proof : The term to be minimised in (6.52) is given by:

$$\begin{aligned} & \overline{E}[\Lambda_{0,k} \exp \theta \Psi_{0,k} | \mathcal{Y}_k, \hat{\mathcal{X}}_{k-1}] \\ &= \overline{E} \left[\exp \frac{\theta}{2} (\hat{X} - X_k)' Q (\hat{X} - X_k) \Lambda_{0,k} \exp \theta \Psi_{0,k-1} | \mathcal{Y}_k, \hat{\mathcal{X}}_{k-1} \right] \\ &= \sum_{j=1}^N \exp \frac{\theta}{2} (\hat{X} - e_j)' Q (\hat{X} - e_j) \alpha_k(j) . \end{aligned}$$

Writing this in matrix notation gives the result. ■

6.6.3 Dual Control

Now that the risk sensitive filtering result has been shown to be obtained from the control result, via an alteration to the cost function, the idea of dual control can be considered. Dual control issues result when a risk-sensitive filter is coupled to a risk-sensitive controller. These ideas are not fully developed but are noted here to provide an insight to future research opportunities.

Consider the following cost function:

$$J(u) = \overline{E} [\Lambda_{0,T} \exp \Psi_{0,T} | \mathcal{Y}_T] , \quad (6.53)$$

where

$$\Psi_{j,k} \triangleq \sum_{\ell=j}^k \frac{1}{2} [\theta_1 x'_\ell M x_\ell + \theta_1 u'_\ell N u_\ell + \theta_2 (x_\ell - \hat{x}_\ell)' Q_\ell (x_\ell - \hat{x}_\ell)] , \quad (6.54)$$

and where

$$\begin{aligned} Q_i &= Q \quad \text{if } i \leq k , \\ &= 0 \quad \text{if } i > k . \end{aligned}$$

In this case, the cost function combines the task of filtering and control. With Q_i defined as it is, the control equations of Theorem 6.4 continue to hold while the filtering equations of Theorem 6.1 likewise hold with replacement of \tilde{z}_{k+1} by \hat{x}_k , and setting $D \equiv I$ as before. By varying θ_1 and θ_2 it is possible to trade off control and filtering objectives. For dual control, where current control decisions are made with the aim of improving filtering so that future control will be better, the situation is more complicated as Q_i is constant for all i .

6.7 Conclusion

In this chapter the solution to the linear risk-sensitive quadratic Gaussian control problem has been presented. Results have been derived for the case of tracking a desired trajectory. The solution to the dynamic programming problem has been achieved without the need to appeal to a certainty equivalence principle, and hence it gives insight to the solution for nonlinear systems. Limit results have also been presented which demonstrate the link to standard linear quadratic Gaussian control. Also, the solution to the problem of risk-sensitive tracking for hidden Markov models has been presented, as well as a discussion on achieving zero steady state error with risk-sensitive control policies. Simulation studies were presented in order to show some advantages of the risk-sensitive approach. Finally, the risk-sensitive filtering solution was derived from the tracking equations, to demonstrate the links between filtering and control.

6.8 Figures

The figures for this chapter are now presented.

$\times 10^2$	$\theta = 0$ (LQG)	$\theta = 0.1$	$\theta = 0.15$
No model error	4.714	4.715	4.716
With model error	9.363	6.076	6.593

Table 6.1: Error analysis for risk-sensitive control

		α^q					
		$\begin{pmatrix} 0 \\ 1 \end{pmatrix}$	$\begin{pmatrix} 0.2 \\ 0.8 \end{pmatrix}$	$\begin{pmatrix} 0.4 \\ 0.6 \end{pmatrix}$	$\begin{pmatrix} 0.6 \\ 0.4 \end{pmatrix}$	$\begin{pmatrix} 0.8 \\ 0.2 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$
θ	0.1	0.1	0.1	0.1	0.3/0.5*	0.5	0.5
	1	0.1	0.1	0.1	0.1	0.5	0.5
	10	0.1	0.1	0.1	0.1	0.1	0.5

* control oscillates between the two values

Table 6.2: Risk-sensitive HMM control values

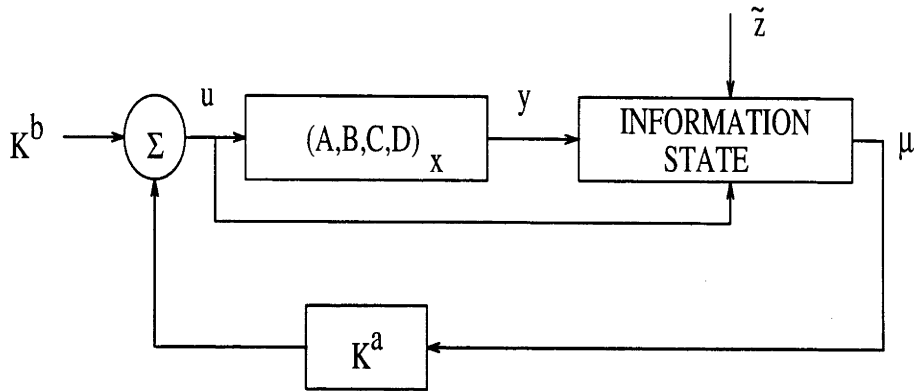


Figure 6.1: Block diagram for standard control policy

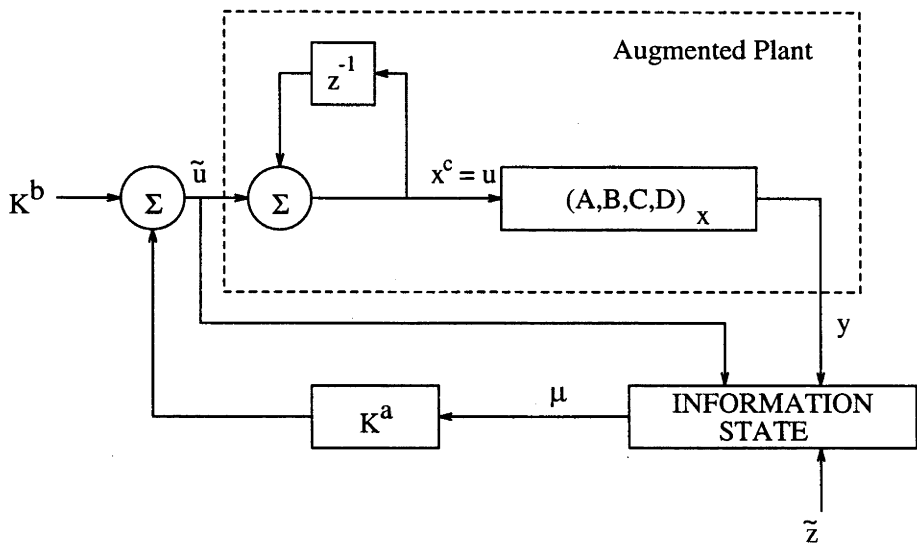


Figure 6.2: Block diagram for constant reference input case




Figure 6.3: No modelling errors




Figure 6.4: With modelling errors




Figure 6.5: Augmented system with modelling errors, $\theta = 0.15$




Figure 6.6: Augmented system with modelling errors, $\theta = 0.17$

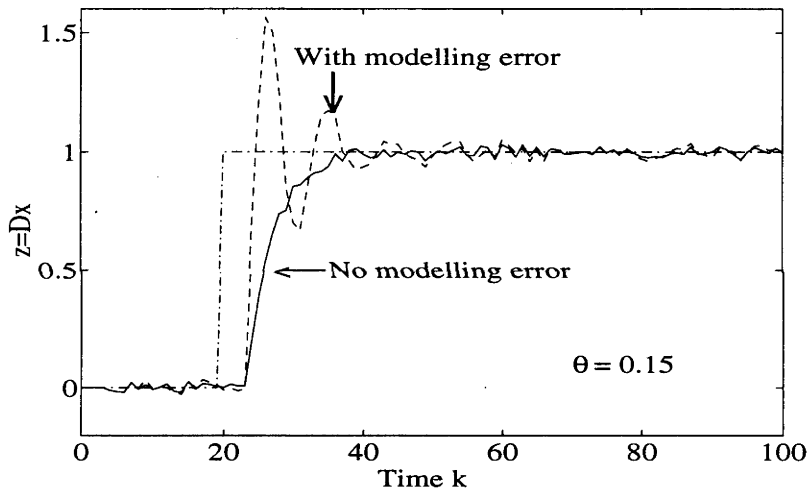


Figure 6.7: Doubly augmented system, $\theta = 0.15$

Chapter 7

Conclusion

7.1 Overview of Thesis

While a great deal of study has been carried out on Markov chains, the field of hidden Markov model (HMM) research is relatively new, and has many open problems. This thesis has considered issues arising in HMM signal processing and control. Specific problems tackled include on-line identification of HMMs, adaptive estimation of mixed state HMMs, risk-sensitive control of both linear systems and hidden Markov models, and the development of reduced computational complexity algorithms.

In addition to the more theoretical issues, there is also enormous potential for the application of HMM based algorithms in a wide range of areas. This thesis considered an application of HMMs to digital communication systems. Such systems have a finite number of discrete states and therefore fall directly into the HMM framework. Specifically, solutions have been presented for the problem of fading transmission channels. Another application of HMM techniques was considered in the area of control and trajectory tracking. Risk-sensitive techniques were developed for HMMs, thus allowing the sensitivity to errors, be they modelling or system noise, to be varied. This thesis provides new and novel solutions to a variety of problems, and serves in many ways to highlight the application potential for HMM signal processing and control.

The main ideas and contributions of this thesis are summarised in the following sections.

7.1.1 State Space Models for HMMs

One of the main contributions of this thesis has been to reformulate the standard HMM representation, into a form more akin to the state space formulations of systems theory. The reformulated state space model for HMMs has been used throughout this thesis. It is an important feature which facilitates the development of new and novel signal processing and control techniques for systems which are Markov in nature.

To achieve the reformulation, the Markov process was represented by a dynamical state equation with a Martingale increment stochastic noise term. This allowed the observations to be written in standard systems notation. In carrying out this reformulation, it was necessary to represent the Markov states by a set of unit indicator vectors. A resulting feature is that any nonlinear function of the state can be represented in a linear form. This is crucial for the signal processing and control schemes developed.

Once this step has been made, all HMM based systems can be viewed as a nonlinear state process, observed via a linear function of the state. In the cases considered in this thesis, the observations were also able to be expressed linearly in terms of the possibly time-varying continuous-range model parameters. The resulting models had observations which were bi-linear in the states and the parameters. Such an approach has two main advantages. The first is that the bi-linear representation allows optimal linear schemes to be conditionally coupled, for state and parameter estimation. The second is that the state space representation of the HMM provides a more systematic approach for the application of HMM techniques to many and varied problems.

In addition to the state space representation, and motivated by the problem of model identification, this thesis presented a new parameterisation for the HMM. This involved slightly reformulating the Markov model, in terms of the square root of transition probabilities. RPE based algorithms were then developed for model identification, which ensured that the estimates for transition probabilities always remain positive. The reformulated model was also applied to a previous EM based scheme, to again ensure positivity.

In conjunction with the new parameterisation, a projection was used, drawing on techniques from differential geometry, to ensure the RPE gradients were evaluated in the tangent space of a sphere.

The sphere was defined by the requirement that the sum of squares of the reformulated transition probabilities must equal one, across each row of the matrix A . By making use of the new parameterisation and related gradient projection, on-line algorithms have been generated which provide competitive, and in fact asymptotically faster converging, estimates to those of earlier proposed on-line HMM identification schemes.

7.1.2 On-Line Algorithms

To date, with few exceptions, hidden Markov model signal processing algorithms have been based on off-line expectation-maximisation algorithms. This thesis, however, has presented a new and systematic approach to on-line HMM signal processing. Such methods can be applied to many systems requiring real-time analysis. Specifically, results have been generated for problems of model identification, and adaptive parameter estimation. This thesis has considered an application to digital communication systems, for which the on-line capabilities are essential. Comparisons with current schemes have shown that, while the HMM approach is generally more computationally intensive, advantages can often be gained, especially in the case of coded signals. As a result of the work in this thesis, real-time signal processing for HMM based systems, be they communication systems, frequency tracking systems, or a host of other possibilities, have become a realistic prospect.

7.1.3 Information-State Signal Models

A key to the results presented throughout this thesis is the use of information-states. In the first instance, by defining the information-state to be a probability distribution of the Markov state, a new signal model has been developed for the HMM. The main benefit from this reformulation is that the resulting signal model no longer consists of discrete-valued states. The concept, therefore, led to algorithms which mixed discrete-valued states and continuous-range states. Standard nonlinear filters, such as the EKF, were then be applied directly, to estimate the states and parameters simultaneously.

By applying these ideas to signal decoding in the case of fading transmission channels, algorithms have been developed which are, amongst other things, generally more robust than current schemes. Traditional approaches to the problem employ MAP estimate decision directed techniques, such as matched filters, or maximum likelihood sequence estimation schemes. By considering an

information-state reformulation, this thesis has presented algorithms for which the complete probability distribution is used. This implies that not only is the estimate of the state fed to the channel tracker, but so is a measure of the accuracy of that estimate. The result is that a degree of robustness is inherent in the new algorithms.

In the latter work on risk-sensitive control, use was made of information-states to obtain optimal output feedback results without needing to appeal to a certainty equivalence principle. The information-state used in this case consisted not only of a probability distribution for the state, but it also contained information about the control cost incurred. Again, a key to the technique is that the full information-state, as opposed to a state estimate or MAP estimate, is fed to the controller.

7.1.4 Conditional Coupled Filters

The application of conditionally coupled filters to HMM based systems is an important aspect of the signal processing schemes developed in this thesis. Since the systems considered are in general nonlinear, and consist of a mixture of discrete-valued and continuous-range states, optimal processing is invariably infinite dimensional. Of course, sub-optimal schemes, such as the extended Kalman filter, can be used, however problems arise when they are applied to information-state representations, as discussed in Remark 4.5. Instead, this thesis has presented conditionally coupled Kalman and HMM filters in a practical approach to adaptive state-and-parameter estimation.

The first step taken in generating such algorithms was to apply separate filters for estimation of the state and parameters. Each filter was optimal, based on the assumption that all other information was known. Of course this is not the case, so each filter was conditioned on outputs from the others. An advantage of this approach is that practical schemes have been produced for implementation in real systems, as opposed to the infinite dimensional optimal schemes which would otherwise have arisen.

The HMM/KF conditionally coupled filtering approach has been applied, in this thesis, to the problem of fading channels in digital communication systems. The techniques developed are, however, quite general and can be applied to any Markov based system to generate practical signal processing algorithms.

7.1.5 Robust Processing

When developing sub-optimal algorithms, it is always necessary to consider the issue of robustness. In the signal processing work of this thesis, robustness gains have been achieved, over current schemes, through the use of information-state techniques, as discussed in Section 7.1.3. In the latter work on control, robustness was used as a motivation for considering a risk-sensitive approach. By the inclusion of an exponential operator in the cost function, a control strategy is developed for which the sensitivity to deviations from desired trajectories can be varied. Limiting arguments have shown that at one extreme, an average or minimum variance policy results, while in another limiting case, robust or H_∞ control results. This thesis has considered risk-sensitive control of HMMs in an effort to robustify known HMM control techniques.

7.2 Algorithms

This section presents a summary of the algorithms developed in this thesis.

- **On-Line Identification of HMMs :** In Chapter 2, RPE techniques were used, along with a reformulation of the standard hidden Markov model, to generate a new on-line identification algorithm for HMMs. The resulting scheme was used to estimate the state values and transition probabilities of the Markov process, as well as the variance of the observation noise. The estimates were shown, via simulation studies, to be quadratically convergent, and to converge to the true values for a wide variety of initial conditions.
- **Reduced Complexity On-Line Identification of HMMs :** A grouping assumption was used in Chapter 3, in order to reduce the computational complexity of the RPE based algorithm of Chapter 2. The resulting algorithm was shown to converge to the true values for a wide variety of initial transition probability estimates, as long as the initial state value estimates were grouped in the same way as the true values.
- **Adaptive Demodulation of QAM signals :** In Chapter 4, new adaptive algorithms were developed for digital QAM signals transmitted through flat fading channels. Two schemes were presented, one linking HMM and EKF techniques, the other linking HMM and KF techniques. The HMM/KF coupled algorithm proved more practical, and simulation studies were carried out which demonstrated the advantages of such an approach over current schemes.

- **Adaptive Demodulation of MDPSK signals** : The techniques of Chapter 4, were applied in Chapter 5, to the case of digital M-ary DPSK signals in fading channels. The resulting algorithms were also shown to be applicable to analogue FM signals with appropriate quantisation.
- **Risk-Sensitive Tracking for Linear Systems** : In Chapter 6, an algorithm was presented for tracking with output feedback and a risk-sensitive control policy. State space augmentations were presented in order to incorporate integrators for the purpose of obtaining zero steady state error. Simulation studies explored many of the advantages to the risk-sensitive approach.
- **Risk-Sensitive Tracking for HMMs** : The risk-sensitive approach was applied to tracking for hidden Markov model systems, to allow for variations in the robustness of HMM controllers. Numerical solutions to the dynamic programming control problem were discussed and a simulation study was used to demonstrate the effect of varying the sensitivity to risk.

7.3 Summation

This thesis has provided a new approach to HMM signal processing and control. The problems considered include model identification, adaptive parameter estimation, and risk-sensitive control. Applications for these HMM techniques, range from the digital communication systems and tracking problems which were discussed in the thesis, to other areas of current interest, such as finite/infinite impulse response filters, pulse train de-interleaving, and image processing. More generally, any system which has states belonging to a finite-discrete set can be approximated by an HMM, and can benefit from the techniques presented in this thesis.

7.4 Future Research

This section presents some open problems and a number of application areas of current interest, related to the work in this thesis.

- **Analytical Convergence Analysis** : To date, while the identification and estimation schemes presented demonstrate good performance in simulation studies, they do not have strong analytical convergence results. It would be interesting to investigate such results further. In the case of RPE, there are good indications to suggest convergence properties can be obtained, based on parallel arguments to those in [Ljung and Söderström 1983]. The problem is of course that in the case of HMMs, a Martingale increment term is associated with the state dynamics, as opposed to white noise in more traditional RPE applications. For the time-varying adaptive estimation case of Chapters 4 and 5, unfortunately there are no such analytical indications, except to say that simulation studies appear very promising.
- **HMM System Theory** : An even more challenging aim in the study of HMM signal processing and control, is the development of a complete systems theory for the models. Unfortunately, much work has been carried out into this problem with few results to date (for example, see [Picci 1978]). However, given the state space representation presented in this thesis, it seems possible that some progress could be made in the direction of establishing parallel results to those of linear systems theory.
- **Filtered HMMs** : The study of communication systems presented in this thesis, has lead to the emergence of a wide variety of problems which are ideally suited to HMM processing techniques. One particular area of interest is the processing of signals transmitted through IIR and FIR filters. By reformulating these systems into the, now familiar, HMM framework, it becomes possible once again to use coupled conditional filters for state and parameter estimation. In addition, and of greater interest for future research, the possibility arises here, of applying the reference probability measure techniques of [Elliott 1993] in a filtering problem, as opposed to the control problem to which they have previously been applied in Chapter 6. In the case of filtered HMMs, estimates, other than the standard state estimate, are required (such as the number of times a state has been active). It seems extremely likely that the Elliott approach may prove useful, and result in new and novel approaches to practical filtering problems.

- **Frequency Tracking** : Recently, algorithms have been developed for frequency tracking using HMM signal processing techniques [Streit and Barrett 1990, Xianya and Evans 1991]. These schemes employed Fourier transform and MAP estimate techniques on blocks of data, in order to track the time varying frequency of a signal. After considering the results presented in this thesis, it would be of interest to investigate two aspects of this frequency tracking problem. The first would be the effect of an information-state approach, as opposed to the MAP estimates used previously. The second would be a possible application of the MDPSK work of Chapter 5 to the problem, thus removing the need for batch processing and Fourier transform processing.
- **Non-Linear Filtering via HMM Techniques** : As mentioned in Section 6.3.2, HMMs can be used as quantised representations of continuous-range variables. In the case of linear systems in Gaussian noise, if the HMM had infinitely small quantisations over the whole space, then the probability distribution out-putted from the HMM filter would be a Gaussian with mean and variance the same as those given by the corresponding Kalman filter. Of course, the computational requirements of such an approach would not be feasible, however, it would be very interesting to investigate the way in which such information-state techniques could be applied to the quantisation of more general nonlinear systems. Techniques such as adaptive quantisation and state aggregation could prove useful in generating quantised filters which would challenge traditional schemes such as the EKF, in certain nonlinear environments.

Appendix A

Gradients in a Manifold

In the case where parameters, x , are constrained to a manifold, M , it is necessary to constrain any derivatives, evaluated with respect to those parameters, to the tangent space of M . This is achieved by the appropriate projection of the derivative in Euclidean space, to the tangent space of the manifold. In the case where $M = \mathbf{S}^{N-1}$, the tangent space at a point x , $x \in \mathbf{S}^{N-1}$, is $\{q \in \mathbf{R}^N \mid q'x = 0\}$, the set of points perpendicular to x . Thus, the projection of $Df = (\partial f / \partial x_i)_{i=1, \dots, N}$, the Euclidean derivative, onto the tangent space of M , is given by

$$\nabla f' = Df' - (Df'x)x' \quad (\text{A.1})$$

The technique for constraining derivatives, as applied in (2.39) (2.40) and (2.42b), will now be outlined using (2.39) as an example. In this case, x_i consists of the elements of s_{ij} from (2.8). It is important, in these recursive functions, to note that if the value of the derivative at the first iteration is zero (therefore automatically satisfying the tangent space requirements), then for all future iterations using the projected gradient, terms which rely on past derivative values, will also be tangent to the manifold. This allows the following result to be generated by taking the Euclidean derivative of (2.28) with respect to t and neglecting the terms involving derivatives of $\hat{\alpha}$.

$$\frac{d\hat{y}_{k+1}}{dt} = 2 \left(\frac{\hat{\alpha}_k(i)}{\langle \hat{\alpha}_k, \mathbf{1} \rangle} g' \text{diag}(x_i) \right) \frac{dx_i}{dt} \quad (\text{A.2})$$

Using (A.1) the projected gradient vector $\nabla f'$ can be evaluated.

$$\nabla f' = 2 \frac{\hat{\alpha}_k(i)}{\langle \hat{\alpha}_k, \mathbf{1} \rangle} g' \text{diag}(x_i) (I_N - x_i x_i') \quad (\text{A.3})$$

The individual elements of this vector notation are written out explicitly, and are given in (2.39).

Appendix B

Theorem Proofs

B.1 Proof of Theorem 6.1

Substitution of (6.8) into (6.7) yields the following:

$$\begin{aligned}
 \alpha_{k+1}(x) &= \phi^{-1}(y_{k+1}) \int_{\mathbf{R}^n} \phi(y_{k+1} - C\xi) \exp(\theta\Psi_{k,k}) \psi(x - A\xi - Bu_k) \alpha_k(\xi) d\xi \\
 &= \phi^{-1}(y_{k+1}) (2\pi)^{-m/2} |\Gamma|^{-1/2} (2\pi)^{-n/2} |\Sigma|^{-1/2} Z_k \\
 &\quad \int_{\mathbf{R}^n} \exp\left(-\frac{1}{2}\right) \left\{ (y_{k+1} - C\xi)' \Gamma^{-1} (y_{k+1} - C\xi) \right. \\
 &\quad \left. - \theta(\xi' M \xi + u'_k N u_k + (D\xi - \tilde{z}_{k+1})' Q (D\xi - \tilde{z}_{k+1})) + \right. \\
 &\quad \left. (x - A\xi - Bu_k)' \Sigma^{-1} (x - A\xi - Bu_k) + (\xi - \mu_k)' R_k^{-1} (\xi - \mu_k) \right\} d\xi \\
 &= \mathcal{A} (2\pi)^{-n/2} \int_{\mathbf{R}^n} \exp\left(-\frac{1}{2}\right) \left\{ \xi' a_k \xi - 2b'_k \xi + c_k \right\} d\xi,
 \end{aligned}$$

where

$$\begin{aligned}
 \mathcal{A} &= \phi^{-1}(y_{k+1}) (2\pi)^{-m/2} |\Gamma|^{-1/2} |\Sigma|^{-1/2} Z_k, \\
 a_k &= C' \Gamma^{-1} C - \theta(M + D' Q D) + A' \Sigma^{-1} A + R_k^{-1}, \\
 b'_k &= \mu'_k R_k^{-1} + (x - Bu_k)' \Sigma^{-1} A + y'_{k+1} \Gamma^{-1} C - \theta \tilde{z}'_{k+1} Q D, \\
 c_k &= (x - Bu_k)' \Sigma^{-1} (x - Bu_k) - \theta u'_k N u_k - \theta \tilde{z}'_{k+1} Q \tilde{z}_{k+1} + y'_{k+1} \Gamma^{-1} y_{k+1} + \mu'_k R_k^{-1} \mu_k.
 \end{aligned}$$

Now,

$$\begin{aligned}
\alpha_{k+1}(x) &= \mathcal{A}|a_k|^{-1/2} \exp\left(-\frac{1}{2}\right) \{c_k - b'_k a_k^{-1} b_k\} \\
&= Z_k |\Sigma|^{-1/2} |a_k|^{-1/2} \exp\left(-\frac{1}{2}\right) \{c_k - b'_k a_k^{-1} b_k - y'_{k+1} \Gamma^{-1} y_{k+1}\} \\
&= Z_k |\Sigma|^{-1/2} |a_k|^{-1/2} \exp\left(-\frac{1}{2}\right) \{x' \beta_1 x - 2\beta'_2 x + \gamma_k\} \\
&= Z_{k+1} \exp\left(-\frac{1}{2}\right) \{(x - \beta_1^{-1} \beta_2)' \beta_1 (x - \beta_1^{-1} \beta_2)\},
\end{aligned}$$

where

$$\begin{aligned}
Z_{k+1} &= Z_k |\Sigma|^{-1/2} |a_k|^{-1/2} \exp\left(-\frac{1}{2}\right) [\gamma_k - \mu'_{k+1} R_{k+1}^{-1} \mu_{k+1}], \\
R_{k+1}^{-1} &= \beta_1 = \Sigma^{-1} - \Sigma^{-1} A a_k^{-1} A' \Sigma^{-1}, \\
\mu_{k+1} &= \beta_1^{-1} \beta_2 = R_{k+1} [\Sigma^{-1} B u_k \\
&\quad + \Sigma^{-1} A a_k^{-1} (R_k^{-1} \mu_k - A' \Sigma^{-1} B u_k + C' \Gamma^{-1} y_{k+1} - \theta D' Q \tilde{z}_{k+1})], \\
\gamma_k &= u'_k (-\theta N + B' \Sigma^{-1} B) u_k + \mu'_k R_k^{-1} \mu_k - \theta \tilde{z}'_{k+1} Q \tilde{z}_{k+1} \\
&\quad - (\mu'_k R_k^{-1} - u'_k B' \Sigma^{-1} A + y'_{k+1} \Gamma^{-1} C - \theta \tilde{z}'_{k+1} Q D) a_k^{-1} \\
&\quad (R_k^{-1} \mu_k - A' \Sigma^{-1} B u_k + C' \Gamma^{-1} y_{k+1} - \theta D' Q \tilde{z}_{k+1}).
\end{aligned}$$

Here ends the proof of Theorem 6.1. Equations (6.12) and (6.13), which follow the theorem, are obtained from the above equations with the aid of the matrix inversion lemma [Kailath 1980] (p. 656):

$$[A + BCD]^{-1} = A^{-1} - A^{-1} B [DA^{-1} B + C^{-1}]^{-1} DA^{-1}, \quad (\text{B.1})$$

for any nonsingular A ($m \times m$) and C ($n \times n$) matrices.

B.2 Proof of Theorem 6.4

This section presents an outline of the proof for Theorem 6.4. The proof is obtained by solving the dynamic programming recursion of Theorem 6.3. This could be done most directly by substituting (6.9) for the information state, and evaluating the integral in the dynamic programming recursion. However, in this section a variable transformation is used to enable the result to be displayed in a form similar to the standard LQG result (where a separation principle allows the *state* feedback problem to be solved, and the state estimate to be substituted in place of the state). From the final equations generated by this proof, the LQG result can be obtained directly, by setting θ equal to zero.

The proof proceeds as follows:

Consider the expression for Z_{k+1} in (6.9). The term in the exponential can be written in the following way:

$$\begin{aligned}\Xi &\triangleq \gamma_k - \mu'_{k+1} R_{k+1}^{-1} \mu_{k+1} \\ &= \mu_k R_k^{-1} \mu_k - \theta u'_k N u_k - \theta \tilde{z}'_{k+1} Q \tilde{z}_{k+1} \\ &\quad - (\mu'_k R_k^{-1} + y'_{k+1} \Gamma^{-1} C - \theta \tilde{z}'_{k+1} Q D) \tilde{K}_k (R_k^{-1} \mu_k + C' \Gamma^{-1} y_{k+1} - \theta D' C' Q \tilde{z}_{k+1}),\end{aligned}$$

where \tilde{K}_k is defined in (6.12). After expanding and collecting terms, the following expression can be reached:

$$\begin{aligned}\Xi &= -\theta u'_k N u_k - \theta \tilde{z}'_{k+1} Q \tilde{z}_{k+1} - \mu'_k \Phi_{xx} \mu_k + 2\mu'_k \Phi_{xv} \hat{v}_k + \hat{v}'_k \Phi_{vv} \hat{v}_k \\ &\quad - (y'_{k+1} - \theta \tilde{z}'_{k+1} Q D C^{-1} \Gamma) \Gamma^{-1} (y_{k+1} - \theta \Gamma (C')^{-1} D' Q \tilde{z}_{k+1}),\end{aligned}$$

where

$$\begin{aligned}\hat{v}_k &= y_{k+1} - C \mu_k - \theta \Gamma (C')^{-1} D' Q \tilde{z}_{k+1}, \\ \Phi_{xx} &= \theta (M + D' Q D) [I + \theta \tilde{K}_k (M + D' Q D)], \\ \Phi_{xv} &= \theta (M + D' Q D) \tilde{K}_k C' \Gamma^{-1}, \\ \Phi_{vv} &= \Gamma^{-1} - \Gamma^{-1} C \tilde{K}_k C' \Gamma^{-1}.\end{aligned}$$

Now Ξ can be written in the following form:

$$\begin{aligned}\Xi &= -\theta u'_k N u_k - \theta \mu'_k \tilde{M} \mu_k + \tilde{v}'_k \tilde{v}_k - \theta \tilde{z}'_{k+1} Q \tilde{z}_{k+1} \\ &\quad - (y'_{k+1} - \theta \tilde{z}'_{k+1} Q D C^{-1} \Gamma) \Gamma^{-1} (y_{k+1} - \theta \Gamma (C')^{-1} D' Q \tilde{z}_{k+1}),\end{aligned}$$

where

$$\begin{aligned}\tilde{v}_k &\triangleq \Phi_{vv}^{1/2} [\hat{v}_k - \Phi_{vv}^{-1} \Phi'_{xv} \mu_k], \\ \tilde{M} &= (1/\theta) (\Phi_{xx} + \Phi_{xv} \Phi_{vv}^{-1} \Phi'_{xv}) \\ &= (M + D' Q D) [I - \theta R_k (M + D' Q D)]^{-1}.\end{aligned}$$

Rewriting the recursion for μ_k in terms of the new variable \tilde{v}_k , yields the following expression:

$$\mu_{k+1} = \tilde{A} \mu_k + B u_k + \tilde{\Gamma}_k \tilde{v}_k,$$

where

$$\begin{aligned}\tilde{A} &= A [I - \theta R_k (M + D' Q D)], \\ \tilde{\Gamma}_k &= A \tilde{K}_k C' \Gamma^{-1} C \Theta, \\ \Theta &= [(C' \Gamma C)^{-1} + [I - \theta R_k (M + D' Q D)]^{-1} R_k]^{1/2}.\end{aligned}$$

This concludes the variable transformation for the information state. The dynamic programming problem is now solved by substituting (6.18) and Ξ into the recursion in Theorem 6.3 and evaluating the integral:

$$V(\chi, k) = \inf_{u \in U_{k,k}} \int_{\mathbf{R}^m} V(\chi_{k+1}(\chi_k, u, y_{k+1}), k+1) \psi^{-1}(y_{k+1}) dy_{k+1}.$$

By applying the transformation

$$\begin{aligned} y_{k+1} &= \Phi_{vv}^{-1/2} \tilde{v}_k + C\mu_k + \theta\Gamma(C')^{-1}D'Q\tilde{z}_{k+1} + \Phi_{vv}^{-1}\Phi'_{xv}\mu_k, \\ d\tilde{v}_k/dy_{k+1} &= \Phi_{vv}^{1/2}, \end{aligned}$$

and completing the square in the integral, an expression is gained which is minimised by the selection of u_k according to (6.19). The recursions in (6.20) are obtained by substituting this value for u_k back into the expression for the value function, rearranging the expression into the general form, and matching up the terms. The details are omitted here as the procedure for completing the square is similar to that used in Section B.1.

Bibliography

- ADAMS M. AND V. GUILLEMIN (1986). *Measure theory and probability*. Wadsworth & Brooks/Cole Advanced Books and Software. Monterey, California.
- AGGOUN L., A. BENSOUSSAN, R. J. ELLIOTT AND J. B. MOORE (1994). Finite-dimensional quasi-linear risk-sensitive control. *Systems and Control Letters*, to appear.
- ALDHAHERI R. W. AND H. K. KHALIL (1991). Aggregation of the policy iteration method for nearly completely decomposable Markov chains. *IEEE Trans. on Automatic Control*. Vol. 36, No. 2, pp. 178–187.
- ANDERSON B. D. O. AND J. B. MOORE (1979). *Optimal filtering*. Prentice-Hall. New Jersey.
- ANDERSON B. D. O. AND J. B. MOORE (1989). *Optimal control : linear quadratic methods*. Prentice-Hall. New Jersey.
- BAUM L. E., T. PETRIE, G. SOULES AND N. WEISS (1970). A maximisation technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Annals of Mathematical Statistics*. Vol. 41, No. 1, pp. 164–171.
- BELLEGGARDA J. R., D. NAHAMOO, K. S. NATHAN AND E. J. BELLEGGARDA (1994). Supervised hidden Markov modeling for on-line handwriting recognition. in *Proc. of the Int. Conf. on Acoustics, Speech and Signal Processing : ICASSP 94, Adelaide, Australia*. Vol. 5. pp. 149–152.
- BENSOUSSAN A. AND J. H. VAN SCHUPPEN (1985). Optimal control of partially observable stochastic systems with an exponential-of-integral performance index. *SIAM Jour. on Control and Optimization*. Vol. 23, pp. 599–613.
- BIGLIERI E., D. DIVSALAR, P. J. MCLANE AND M. K. SIMON (1991). *Introduction to trellis-coded modulation with applications*. Macmillan Publishing Company. New York.

- BINGHAM J. A. C. (1988). *The theory and practice of modern design*. John Wiley & Sons. New York.
- CAO W. AND W. J. STEWART (1985). Iterative aggregation/disaggregation techniques for nearly uncoupled Markov chains. *Jour. of the Association for Computing Machinery*. Vol. 32, No. 3, pp. 702–719.
- CARAVANI P. (1986). On extending linear quadratic control theory to non-symmetric risky objectives. *Jour. Economic Dynamics & Control*. Vol. 10, pp. 83–88.
- CHEN H. F. AND L. GUO (1986). Convergence rate of least-squares identification and adaptive control for stochastic systems. *Int. Jour. Control*. Vol. 44, No. 5, pp. 1459–1476.
- CHUNG S. H., V. KRISHNAMURTHY AND J. B. MOORE (1991). Adaptive processing techniques based on hidden Markov models for characterising very small channel currents buried in noise and deterministic interferences. *Philosophical Transactions of the Royal Society, Lond., Series B*. Vol. 334, pp. 357–384.
- COURTOIS P. J. (1975). Error analysis in nearly-completely decomposable stochastic systems. *Econometrica*. Vol. 43, No. 4, pp. 691–709.
- DELEBECQUE F. AND J. QUADRAT (1981). Optimal control of Markov chains admitting strong and weak interactions. *Automatica*. Vol. 17, No. 2, pp. 281–296.
- DEMPSTER A. P., N. M. LAIRD AND D. B. RUBIN (1977). Maximum likelihood estimation from incomplete data via the EM algorithm. *Jour. Royal Statistical Society, Series B*. Vol. 39, pp. 1–38.
- DEY S. AND J. B. MOORE (1994). Risk-sensitive filtering and smoothing via reference probability method. submitted.
- DUEL-HALLEN A. AND C. HEEGARD (1989). Delayed decision-feedback sequence estimation. *IEEE Trans. on Communications*. Vol. 37, No. 5, pp. 428–436.
- DU J. AND B. VUCETIC (1991). New 16-QAM trellis codes for fading channels. *Electronics Letters*. Vol. 27, No. 12, pp. 1009–1010.
- DUPUIS P. AND H. J. KUSHNER (1989). Stochastic approximation and large deviations : upper bounds and w.p.1 convergence. *SIAM Jour. on Control and Optimization*. Vol. 27, No. 5, pp. 1108–1135.

- ELLIOTT R. J. (1993). A general recursive discrete time filter. *Jour. of Applied Probability*. Vol. 30, pp. 575–588.
- FORNEY JR. G. D. (1973). The Viterbi algorithm. *Proc. of the IEEE*. Vol. 61, No. 3, pp. 268–278.
- GOODWIN G. C. AND K. S. SIN (1984). *Adaptive filtering prediction and control*. Prentice-Hall. New Jersey.
- HAEB R. AND H. MEYR (1989). A systematic approach to carrier recovery and detection of digitally phase modulated signals on fading channels. *IEEE Trans. on Communications*. Vol. 37, No. 7, pp. 748–754.
- HASHEMI H. (1993). The indoor radio propagation channel. *Proc. of the IEEE*. Vol. 81, No. 7, pp. 943–967.
- HAYKIN S. (1983). *Communication systems*. second ed.. John Wiley & Sons. New York.
- IOSIFESCU M. (1980). *Finite Markov processes and their applications*. John Wiley & Sons. Chichester.
- JACOBSON D. H. (1973). Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Trans. on Automatic Control*. Vol. AC-18, No. 2, pp. 124–131.
- JAMES M. R. AND R. J. ELLIOTT (n.d.). Risk-sensitive and risk-neutral control for continuous-time hidden Markov models. preprint.
- JAMES M. R., J. S. BARAS AND R. J. ELLIOTT (1994). Risk-sensitive control and dynamic games for partially observed discrete-time systems. *IEEE Trans. on Automatic Control*. Vol. AC-39, No. 4, pp. 780–792.
- JUANG B. H. AND L. R. RABINER (1985). A probabilistic distance measure for hidden Markov models. *AT&T Technical Journal*. Vol. 64, No. 2, pp. 391–408.
- KAILATH T. (1980). *Linear systems*. Prentice-Hall. New Jersey.
- KALMAN R. E. (1960). A new approach to linear filtering and prediction problems. *Jour. Basic Engineering, Trans. ASME, Series D*. Vol. 82, No. 1, pp. 35–45.
- KALMAN R. E. (1963). New methods in Wiener filtering theory. in *Proc. Symp. Engineering Appl. Random Functions Theory and Probability*. John Wiley and Sons, Inc., New York.

- KALMAN R. E. AND R. S. BUCY (1961). New results in linear filtering and prediction theory. *Jour. Basic Engineering, Trans. ASME, Series D*. Vol. 83, No. 3, pp. 95–108.
- KARP L. S. (1988). Dynamic hedging with uncertain production. *Int. Economic Review*. Vol. 29, pp. 621–637.
- KEMENY J. G. AND J. L. SNELL (1960). *Finite Markov chains*. D. Van Nostrand Company, Inc., Princeton, New Jersey.
- KRICHAGINA N. V., R. S. LIPTSER AND E. Y. RUBINOVICH (1984). Kalman filter for Markov processes. in *Steklov Seminar 1984 : Statistics & Control of Stochastic Processes*. pp. 197–213.
- KRISHNAMURTHY V. AND J. B. MOORE (1993). On-line estimation of hidden Markov model parameters based on the Kullback-Leibler information measure. *IEEE Trans. on Signal Processing*. Vol. 41, No. 8, pp. 2557–2573.
- KUMAR P. R. AND P. VARAIYA (1986). *Stochastic systems*. Prentice-Hall. New Jersey.
- KUSHNER H. J. AND A. SHWARTZ (1984). An invariant measure approach to the convergence of stochastic approximations with state dependent noise. *SIAM Jour. on Control and Optimization*. Vol. 22, No. 1, pp. 13–27.
- LA SCALA B. F., R. R. BITMEAD AND M. R. JAMES (1993). Conditions for stability of the extended Kalman filter and their application to the frequency tracking problem. Submitted.
- LEE E. A. AND D. G. MESSERSCHMITT (1988). *Digital communication*. Kluwer Academic Publishers. Boston.
- LEVINSON S. E., L. R. RABINER AND M. M. SONDDHI (1983). An introduction to the application of the theory of probabilistic functions of a Markov process to automatic speech recognition. *The Bell System Technical Journal*. Vol. 62, No. 4, pp. 1035–1074.
- LINDQVIST B. (1978). On the loss of information incurred by lumping states of a Markov chain. *Scand. Jour. Statist.* Vol. 5, pp. 92–98.
- LIN S. (1970). *An introduction to error correcting codes*. Prentice-Hall. Englewood Cliffs, New Jersey.
- LJUNG L. (1977). Analysis of recursive stochastic algorithms. *IEEE Trans. on Automatic Control*. Vol. AC-22, No. 4, pp. 551–575.

- LJUNG L. (1987). *System identification : Theory for the user*. Prentice-Hall. New Jersey.
- LJUNG L. AND T. SÖDERSTRÖM (1983). *Theory and practice of recursive identification*. MIT Press. Cambridge, Massachusetts.
- LODGE J. H. AND M. L. MOHER (1990). Maximum likelihood sequence estimation of CPM signals transmitted over Rayleigh flat-fading channels. *IEEE Trans. on Communications*. Vol. 38, No. 6, pp. 787–794.
- LOO C. AND N. SECORD (1991). Computer models for fading channels with applications to digital transmission. *IEEE Trans. on Vehicular Technology*. Vol. 40, No. 4, pp. 700–707.
- MAHER B. (1992). Phase-locked loop FM demodulators. *Electronics Australia*. pp. 54–55.
- MARKOV A. A. (1906). Extension of the law of large numbers to dependent variables. *Izv. Fiz-Mat. Obšč. pri Kazansk. Univ. (2 ser.)(Russian)*. Vol. 15, No. 4, pp. 135–156.
- MARKOV A. A. (1907). Investigation of an important case of dependent trials. *Izvestia Acad. Nauk. SPB (ser. I)(Russian)*. Vol. 6, pp. 61–80.
- MARKOV A. A. (1913). An example of statistical analysis of the text of “Eugene Onegin”, illustrating the associations of trials into a chain. *IAN (ser. 6)(Russian)*. Vol. 7, pp. 153–162.
- MARKOV A. A. (1924). *Probability calculus*. 4th ed.. (Russian). Moscow, Gosizdat.
- MARKOV A. A. (1951). *Selected works*. Izd. Akad. Nauk. SSSR, (Russian). Moscow.
- MEYER C. D. (1989). Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems. *SIAM Review*. Vol. 31, No. 2, pp. 240–272.
- MOORE J. B. AND H. WEISS (1979). Recursive prediction error methods for adaptive estimation. *IEEE Trans. on Systems, Man, and Cybernetics*. Vol. SMC-9, No. 4, pp. 197–205.
- PAHLAVAN K. AND J. W. MATTHEWS (1990). Performance of adaptive matched filter receivers over fading multipath channels. *IEEE Trans. on Communications*. Vol. 38, No. 12, pp. 2106–2113.
- PICCI G. (1978). On the internal structure of finite-state stochastic processes. in *Recent Developments in Variable Structure Systems*. Vol. 162. New York : Springer Verlag.
- POLYAK B. T. AND A. B. JUDITSKY (1992). Acceleration of stochastic approximation by averaging. *SIAM Jour. on Control and Optimization*. Vol. 30, No. 4, pp. 838–855.
- PROAKIS J. G. (1983). *Digital communications*. second ed.. McGraw-Hill. New York.

- RABINER L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. of the IEEE*. Vol. 77, No. 2, pp. 257–285.
- ROMANOVSKY V. I. (1970). *Discrete Markov chains*. Wolters-Noordhoff Publishing. The Netherlands.
- SAMPEI S. AND T. SUNAGA (1993). Rayleigh fading compensation for QAM in land mobile radio communications. *IEEE Trans. on Vehicular Technology*. Vol. 42, No. 2, pp. 137–147.
- SEGALL A. (1976). Recursive estimation from discrete-time point processes. *IEEE Trans. on Information Theory*. Vol. IT-22, No. 4, pp. 422–431.
- SIMON H. A. AND A. ANDO (1961). Aggregation of variables in dynamic systems. *Econometrica*. Vol. 29, No. 2, pp. 111–138.
- SPEYER J. L., C. FAN AND R. N. BANAVAR (1992). Optimal stochastic estimation with exponential cost criteria. in *Proc. of the 31st Conf. on Decision and Control*. pp. 2293–2298.
- STREIT R. L. AND R. BARRETT (1990). Frequency line tracking using hidden Markov models. *IEEE Trans. on Acoustics, Speech and Signal Processing*. Vol. 38, No. 4, pp. 586–598.
- TITTERINGTON D. M. (1984). Recursive parameter estimation using incomplete data. *Jour. of the Royal Statistical Society, Series B*. Vol. 46, No. 2, pp. 257–267.
- TREES H. L. VAN (1968-1971). *Detection, estimation and modulation theory*. John Wiley & Sons. New York.
- VITERBI A. J. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans. on Information Theory*. Vol. IT-13, pp. 260–269.
- VUCETIC B. AND J. DU (1992). Channel modeling and simulation in satellite mobile communication systems. *IEEE Jour. on Selected Areas in Communications*. Vol. 10, No. 8, pp. 1209–1218.
- WEBB W. T., L. HANZO AND R. STEELE (1991). Bandwidth efficient QAM schemes for Rayleigh fading channels. *IEE Proc. I, Communications, Speech and Vision*. Vol. 138, No. 3, pp. 169–175.
- WEINSTEIN E., M. FEDER AND A. V. OPPENHEIM (1990). Sequential algorithms for parameter estimation based on the Kullback-Leibler information measure. *IEEE Trans. on Acoustics Speech and Signal Processing*. Vol. 38, No. 9, pp. 1652–1654.

- WHITTLE P. (1981). Risk-sensitive linear/quadratic/gaussian control. *Advances in Applied Probability*. Vol. 13, pp. 764–777.
- WHITTLE P. (1990). *Risk-sensitive optimal control*. John Wiley & Sons. Chichester, New York.
- WIDROW B. AND S. D. STEARNS (1985). *Adaptive signal processing*. Prentice-hall. Englewood Cliffs, New Jersey.
- XIANYA X. AND R. J. EVANS (1991). Multiple target tracking and multiple frequency line tracking using hidden Markov models. *IEEE Trans. on Signal Processing*. Vol. 39, No. 12, pp. 2659–2676.
- ZEITOUNI O. AND A. DEMBO (1988). Exact filter for the estimation of the number of transitions of finite-state continuous-time Markov processes. *IEEE Trans. on Info. Theory*. Vol. 34, No. 4, pp. 890–893.