

UC San Diego

UC San Diego Previously Published Works

Title

Adaptive Optimal Control The Thinking Man's GPC

Permalink

<https://escholarship.org/uc/item/01z3w6d7>

Authors

Bitmead, Robert R

Gevers, Michel

Wertz, Vincent

Publication Date

1990

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Adaptive Optimal Control

The Thinking Man's GPC

*To Jan, Naomi and Bronwen
-R.R.B.*

*To Marie-Antoinette
-M.G.*

*To Colette, Nicolas, Sylvain and Isaline
-V.W.*

Prentice Hall International Series
in Systems and Control Engineering

M. J. Grimble, Series Editor

BANKS, S. P., *Control Systems Engineering: Modelling and Simulation, Control Theory and Microprocessor Implementation*
BANKS, S. P., *Mathematical Theories of Nonlinear Systems*
BENNETT, S., *Real-time Computer Control: An Introduction*
BITMEAD, GEVERS, WERTZ, *Adaptive Optimal Control*
CEGRELL, T., *Power Systems Control*
COOK, P. A., *Nonlinear Dynamical Systems*
LUNZE, J., *Robust Multivariable Feedback Control*
SODERSTROM, T., and STOICA, P., *System Identification*
WARWICK, K., *Control Systems: An Introduction*
PATTON, R., CLARK, R. N., FRANK, P. M. (editors), *Fault Diagnosis in Dynamic Systems*

Adaptive Optimal Control

The Thinking Man's GPC

Robert R. Bitmead
*Australian National University
Canberra, Australia*

Michel Gevers
*Université Catholique de Louvain
Louvain-la-Neuve, Belgium*

Vincent Wertz
*Université Catholique de Louvain
Louvain-la-Neuve, Belgium*



Prentice Hall

New York London Toronto Sydney Tokyo Singapore

© 1990 by Prentice Hall of Australia Pty Ltd

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission of the publisher.

2 3 4 5 94 93 92

Printed and bound in Australia by Impact Printing Pty Ltd, Brunswick, Victoria.

ISBN 0 13 013277 2.

**National Library of Australia
Cataloguing-in-Publication Data**

Bitmead, Robert R.
Adaptive optimal control: the thinking man's GPC

Bibliography.
Includes index.
ISBN 0 13 013277 2.

I. Adaptive control systems. I. Gevers, Michel.
II. Wertz, Vincent. III. Title.

629.836

**Library of Congress
Cataloguing-in-Publication Data**

Bitmead, Robert R.
Adaptive optimal control: the thinking man's GPC/Robert R.
Bitmead, Michel Gevers, Vincent Wertz.
p. cm.
Includes bibliographical reference (p.) and index.
ISBN 0 13 013277 2.

I. Adaptive control systems. I. Gevers, Michel. II. Wertz,
Vincent. III. Title
I. Title.

TJ217.848 1990 90-7531
629.8'36--dc20 CIP

Prentice Hall, Inc., *Englewood Cliffs, New Jersey*
Prentice Hall Canada, Inc., *Toronto*
Prentice Hall Hispanoamericana, SA, *Mexico*
Prentice Hall of India Private Ltd, *New Delhi*
Prentice Hall International, Inc., *London*
Prentice Hall of Japan, Inc., *Tokyo*
Prentice Hall of Southeast Asia Pty Ltd, *Singapore*
Editora Prentice Hall do Brasil Ltda, *Rio de Janeiro*



PRENTICE HALL

A division of Simon & Schuster

Contents

Preface	ix
1 The Scene, the Props, the Players	1
1.1 Introduction and Purpose	1
1.2 A Jaundiced View of Adaptive Control History	3
1.3 Further Perspectives	5
1.3.1 Robustness	5
1.3.2 Recent Trends	6
1.4 A Gedankenexample	8
1.4.1 Open Loop Identification	10
1.4.2 Control Law Selection	13
1.4.3 Closed Loop Identification	14
1.4.4 Summary	15
1.5 The Audience	16
1.6 A Brief Tour	17
2 Generalized Predictive Control	21
2.1 Introduction	21
2.1.1 Motivation	21
2.1.2 What is Predictive Control?	23
2.1.3 A Brief Historical Perspective	24
2.2 The GPC Method of Clarke et al.	27
2.3 Optimal Prediction and GPC Solution	29
2.4 A Simple GPC Example	32
2.5 Closed Loop Formulae	34
2.6 GPC Based on a ‘Performance Model’	36
2.7 Conclusions	42

3	Linear Quadratic Gaussian Optimal Control	43
3.1	Introduction	43
3.2	The Linear Quadratic Regulator	46
3.2.1	The Finite Horizon Regulator	46
3.2.2	The Infinite Horizon Regulator	47
3.2.3	The Receding Horizon Regulator	47
3.3	The Linear Quadratic Tracking Problem	48
3.4	The Linear Optimal State Estimator	52
3.4.1	The Kalman Predictor (KP)	52
3.4.2	The Kalman Filter (KF)	53
3.5	Optimal Filter Design with Disturbance Models	54
3.6	LQG Controllers	57
3.6.1	The Composite System and the LQ Objective	57
3.6.2	Observers for the Composite System	59
3.7	Examples	60
3.8	Closed Loop Control Formulae	70
3.9	GPC as LQG	75
3.9.1	Control Criterion Equivalence	75
3.9.2	An Example	77
3.10	Conclusions	80
4	Stability and Performance Properties of Receding Horizon LQ Control	83
4.1	Introduction	83
4.2	Monotonicity and Stability of Receding Horizon LQ Control	85
4.2.1	Stability via the ARE	85
4.2.2	Monotonicity Properties of the RDE	89
4.2.3	Stability via Monotonicity of the RDE	90
4.3	Stabilizing Feedback Strategies	93
4.3.1	Alternative Forms of the RDE	94
4.3.2	The Stabilizing Controllers of Kwon, Pearson and Kleinman	96
4.4	Mid Chapter Conclusion	98
4.5	Comparative Performance of LQ Schemes	98
4.6	Stability Properties of GPC	101
4.6.1	An Unstable GPC Example	102
4.6.2	Time-varying Strategies in Receding Horizon Control	105
4.6.3	The Use of N_u in GPC Stability	106
4.6.4	Stability Theorems of Clarke and Mohtadi	107
4.7	Conclusions	110

5	Robust LQG Design — Features for Adaptive Control	113
5.1	Introduction	113
5.1.1	A First Hint at the Adaptation/Robustness Interplay	114
5.1.2	A Guided Tour of LQG Robustness Theory	115
5.2	Robustness of Unity Feedback Systems	117
5.3	LQ and KF Robustness — Return Difference Equalities	124
5.3.1	The LQ Return Difference Equality	124
5.3.2	The KP Return Difference Equality	127
5.3.3	The EDRs and LQ, KP Robustness	127
5.4	Robustness of LQG Control — Loop Transfer Recovery	130
5.4.1	Loop Transfer Recovery Rationale	131
5.4.2	LQG Controller Transfer Functions	132
5.4.3	The Discrete-time LTR Theory of Maciejowski	135
5.5	An LQG/LTR Example	141
5.6	Conclusion	145
6	Recursive Least Squares Identification in Adaptive Control	147
6.1	Introduction	147
6.2	Prediction Error Identification	149
6.2.1	Off-line Prediction Error Identification: a Refresher	149
6.2.2	Recursive Least Squares Identification	151
6.3	Frequency Domain Properties of the Identified Model	153
6.3.1	Open Loop Identification	154
6.3.2	Closed Loop Identification	155
6.4	Recursive Identification in Closed Loop Control — Local Theory	158
6.4.1	Heuristic Motivation	158
6.4.2	Potential Convergence Point	160
6.4.3	Integral Manifolds and Slow Adaptation	164
6.5	Recursive Identification in Closed Loop — Global Methods	168
6.5.1	Normalization and Deadzones	169
6.5.2	Projection and Leakage	170
6.5.3	Covariance Resetting	170
6.5.4	A Global Comment	171
6.6	Conclusion	172
7	A Candidate Robust Adaptive Predictive Controller	175
7.1	Introduction	175
7.2	The Certainty Equivalence Control Law	177
7.2.1	The Plant, Noise and Reference Models	177

7.2.2	Kalman Filter Design	179
7.2.3	LQ State-variable Feedback Design	181
7.2.4	The LQG Controller	182
7.3	The System Parameter Identifier	183
7.4	The Candidate — A Summary	186
7.5	The Platform	187
7.5.1	LQG Controller Properties	187
7.5.2	RLS Identifier Properties	188
7.5.3	Manipulation of the Candidate Controller	193
7.6	Computer Studies and Examples	194
7.6.1	The Gedankenexample Revisited	194
7.6.2	The Working Example Revisited	198
7.7	Conclusion	204
8	Le Jugement Dernier	205
8.1	Introduction — The Final Analysis	205
8.2	Adaptation and Stability Robustness	206
8.2.1	Linear Stability Robustness	208
8.2.2	Closed Loop Identification and Stability Robustness	212
8.3	Adaptation and Performance	215
8.4	Forethoughts on a Postscript	217
8.5	Extensions and Generalizations	220
8.5.1	Candidate Alternative Optimal Control Laws	220
8.5.2	Alternative Identification Methods	224
8.6	Refinements and Theoretical Support	226
8.6.1	Theoretical Refinements	227
8.6.2	The Rohrs Examples	228
8.6.3	Adaptive Control versus Robust Control	230
8.7	Coda	231
	References	233
	Index	241

Preface

Adaptive Control is a progressive and fascinating field of control systems design and analysis that continues to be particularly ebullient because of the intellectual appeal of its notions of *self-tuning* or *self-optimizing*. Added to this technical attraction is the enthralling feature that, absence of complete theoretical support notwithstanding, many practical implementations of adaptive control have been working on industrial systems. Our goal in this book is to explore new connections between adaptive control theory and practice which permit fuller contact of adaptive control design with modern robust control design, thereby improving the scope of the former and the motivation of the latter.

About the Title

The vehicle that we use for our theoretical study is Adaptive Linear Quadratic Gaussian Optimal Control, hence *Adaptive Optimal Control* in our title. The practical connection is via the recently popular Generalized Predictive Control adaptive control method, which has proven remarkably successful in applications, therefore *GPC*. The final unexplained portion of the title, *The Thinking Man*¹, reflects our aim to introduce new theoretical tools to bring this style of adaptive control design into more complete harmony with existing alternative robust design methods, allowing a more balanced and explicable design phase. By developing our presentation from this specific Adaptive Optimal Control/GPC stance, we are able to move to the more general understanding more easily in the later parts of the book.

Our subtitle, while sounding a little self-inflated, is not intended to denigrate GPC in any way. Indeed, the motivation for the writing of this book is that GPC has proven so successful in practice. We have taken up the challenge to build a theoretical support for this and similar algorithms and, in doing so, it is clear that we must build upon and extend the advances

¹The authors are not attempting to be sexist in selecting such a title, merely to reflect their experience that no Thinking Woman uses GPC.

of those before. So, if we do sound a little negative towards GPC at times, please regard this as straight talking permissible through familiarity rather than as veiled calumniations.

The Audience and Prerequisite Knowledge

The perceived audience for this work falls into several categories reflecting their classification as: practitioners interested in acquiring some theory associated with the design of adaptive control algorithms, theoreticians concerned with studying the connection between methods which have arisen from practical applications and their more mathematical analyses, graduate students of control systems interested in both disciplines. The audience is identified more fully in the first chapter after some motivational examples, when some guide to reading is given.

The book has been written to provide a study of adaptive control but the explicit techniques treated along the way include: Linear Quadratic optimal control and estimation (Kalman filtering), recursive identification, some Linear Systems theory, and robustness arguments. On the whole, we have endeavored to be complete in our derivations where this is pertinent to understanding the material, but some previous exposure to Linear Quadratic control and estimation and to recursive identification would be advantageous. Naturally, a goodly number of suitable references of such source material are provided.

Acknowledgments

First in the list of people for whom acknowledgment, thanks and tribute are due for their suffering during the writing of this tome, are our families and close friends who have endured absences overseas, late late nights at home and the office, and days of neglect. Allied with this group of mistreated people deserving of an explanation must also come our graduate students and colleagues. For their careful technical reading, which was unremittingly iconoclastic just when we needed it most, we delight in thanking Brian Anderson, Georges (Jojo) Bastin and Geoff Williamson. Their efforts have aided us very considerably in honing, tweeking and, where necessary, completely rewriting our explanations and rationalizing our philosophical diatribes. Many others have given of their time and efforts to help us formulate better our approach and technique. Notable amongst these are Greg Allen, Laurent Praly, Claude Samson, Jan Willem Polderman and two of the lead players in GPC, David Clarke and Coorous Mohtadi. Finally, we thank the Groupe de Recherche Coordonnée (GRECO) SARTA and the boys of Grenoble (Ioan Landau, Luc

Dugard, Jean-Michel Dion and Mohamed M'Saad), Rick Johnson, Robert Kosut, Petar Kokotovic and Guy Champion for their support — technical, emotional and moral (as in definition 8 of the Macquarie Dictionary Revised Edition 1985, and in neither the religious nor ethical sense).

Bonne Lecture, Mates.

*Bob, Mig and Vinnie,
Canberra, February 1990.*

Chapter 1

The Scene, the Props, the Players

1.1 Introduction and Purpose

Adaptive Control describes a body of theoretical and practical engineering methodology in which one combines the design of feedback controllers based on plant system models with the on-line estimation of the model parameters or the controller parameters using input and output data measurements. Adaptive Optimal Control then is meant to refer to just such an adaptive control scheme but with the controller design being carried out using the methods of Optimal Control Theory. There is nothing new in this idea and, indeed, in this book we take the most direct and mundane Recursive Least Squares parameter identifiers (RLS) and Linear Quadratic Gaussian control laws (LQG) to concoct an adaptive controller. This would appear to be more anachronistic than avant-garde, more conservative than risqué. Our goal, however, is to explore the intriguing interconnection between the control law selection and the identification rule operating simultaneously in closed loop. These specific selections of RLS and LQG provide both an explicit workplace and a pleasing and direct connection to the currently much favored practical adaptive methods of Generalized Predictive Control (GPC). The validating features of this work are the tying of modern theories of Linear System Identification and Robust Linear Controller Design onto methods developed for their practical and intuitive appeal in adaptive control. The final extensions carry this further to permit the rapid acquisition of modern control design principles into adaptive control in a logical fashion.

Because Adaptive Control consists of the coupling of a parameter estimator and a control law design schema, it is necessary in analysing the behavior

of an adaptive control scheme that we consider two complementary features: the effect of the identified parameter value upon the controller and hence on the closed loop performance, and the effect of the closed loop control on the parameter estimation. It is this explicit interplay between control law selection and parameter identification that is the focal point of this book and which distinguishes it from previous works. These works have tended to concentrate on either the identifier or the control law and have relied to some extent on the insensitivity of one component from the effects of the other. The joint analysis of these separate effects presented here leads us to more specific design guidelines for adaptive control.

The message of our work is that by appealing to robust control design theory this interplay above between on-line control design and closed loop identification can be harnessed in such a way that the combined adaptive control/identification scheme possesses greater robustness margins than the separate designs of control law and identification algorithm would have allowed. In order not to confuse the issues with vague generalities, we carry through our analysis and study of consequences for the specific class of Adaptive (Linear Quadratic) Optimal Control algorithms rather than attempting to encompass all potential schemes at once. Our starting point for development and motivation is a particular member of this family of algorithms, the currently popular Generalized Predictive Control (GPC) method, which is a practically engendered methodology based on intuitively appealing design rules that has met with considerable industrial acceptance. We use this procedure as the vehicle for our investigations and to make concrete our specific commentary concerning adaptive control algorithms.

The theoretical machinery of our approach combines the following tools:

- stability of optimal control laws in their various guises: finite horizon, receding horizon, infinite horizon;
- stability margins of optimal control laws — design for robustness;
- frequency domain characterization of parameter estimation in closed loop;
- integral manifold methods for the analysis of slow adaptation, linking adaptation dynamics to closed loop performance.

Our ambition is to arrive at the specification of an adaptive control procedure with guaranteed closed loop stability margins with respect to both plant undermodeling and to initial condition effects, and possessing a very small number of design parameters linked to closed loop performance in an

intuitively comprehensible way. The results arising from this analysis are an explicit rationalization of the adaptive optimal control approach. This incorporates specific guidelines pertaining to the above techniques with the goal of assuring closed loop stability margins for the adaptively controlled system. These results thus make full contact between adaptive control design and recent robust nonadaptive control design. The perversity of our methods is that the resultant adaptive control scheme preserves many of the hallmarks and simplicities of GPC, albeit at the specific exclusion of other features.

The development proceeds in the following order. The GPC control law schema is presented in Chapter 2 together with some of its embellishments. Linear Quadratic Gaussian (LQG) control design methods are the subject of Chapter 3, with the incorporation of GPC as a special case. Closed loop stability and performance issues of LQG (and thus of GPC) are treated in Chapter 4. Chapter 5 completes our analysis of nonadaptive control properties with the presentation of robustness results for linear systems, made specific for LQG control via Loop Transfer Recovery (LTR) techniques. Chapter 6 has a dual motive: to examine closed loop parameter estimation with undermodeling, and to study how on-line adaptation affects the closed loop behavior. On the basis of these results the interplay between parameter estimation and control law selection is utilized to derive a coherent adaptive optimal control design in Chapter 7. This is our Candidate Robust Adaptive Predictive Controller. Chapter 8 explores the extension of these results into the fuller arena of Adaptive Control with choices of control law possibly being different from LQG, choices of the identification criterion being different from one-step-ahead prediction, treatment of tracking of slow plant variations, and a critique of Adaptive Control versus Robust Nonadaptive Control. This last chapter draws out from our specific study of Adaptive LQG Control to present in much fuller generality (but of course less detail) the extension of the theory into a broader context.

1.2 A Jaundiced View of Adaptive Control History

The motivation for writing this book stemmed from several impolite but (we think) astute observations of the authors. These were concerned with the nature of adaptive control theory and practice and centered around the realization that, with few exceptions, adaptive control theoreticians and practitioners were looking at adaptive control from angles so wide apart that they were seemingly not addressing the same issues.

Historically, the applied adaptive control utilizers have focused upon the specification of a control law suitable for adaptive applications. That is, the emphasis has been upon the statement of a control law design procedure, starting from the prescription of a linear model, which would yield a computationally manageable method capable of delivering specified closed loop performance. Thus these researchers operated within the rules of finding a super control law able to overcome the vagaries of any identification method. The discussion then centered upon issues such as the relative merits of pole placement versus predictive control strategies, and, within this latter family of controllers, the development of ever more sophisticated control criteria designed with a view to coping with wider and wider classes of plants and perturbations. The rôle of the identifier in this approach is merely to provide the plant parameter estimates and its specific design is not treated openly as a major issue.

Contrariwise, the theoretical advances in adaptive control have recently centered upon the use of normalization, relative deadzones, leakage and other devices in the identifier to achieve the best possible parameter estimate properties, essentially independently from the particular control law chosen for the nominal closed loop. By cajoling and forcing the parameter estimates to conform to certain growth conditions whatever the control law, the global boundedness of system signals is asserted. Thus, these individuals concentrated upon the development of a super identifier able to leap over any controller in a single bound.

Our view is heretical in both camps. We roughly ignore computational issues to a large extent, since these problems tend to disappear with technological advances, and we do not concern ourselves too much with global parameter convergence properties and vague statements of performance such as ultimate boundedness. Rather, we home in upon the development of a joint criterion for the performance of the identifier stage and the control design stage of adaptive control. Our results are local in nature but are synthetic and synergistic because they suggest how available design choices ought to be made to reflect either *a priori* knowledge or ultimate closed loop objectives. We believe that our new contribution lies precisely in the development of a combined adaptive identification and control strategy in which the design of the identifier and its filters fully utilizes the controller properties to allow slackness in the model in the frequency ranges where this does not deteriorate the closed loop stability margins too much, and conversely to design the controller in a way that takes into account the frequency ranges in which the identified model has greatest validity. In doing so we hope to provide the missing link in adaptive control.

1.3 Further Perspectives

1.3.1 Robustness

Much of the terminology of modern control systems theory uses the notion of robustness and, as we too shall be appealing to this concept, it is as well to make some remarks about robustness at this early stage so that later statements not be deemed too vague.

A property is *robust* if, should the property hold at one point, it holds in an open neighborhood of that point [Lau79a]. A similar notion for dynamical systems is *structural stability*, where small changes in the system do not produce large changes in the solutions [GH83]. Thus ideas of robustness require the specification of a topology (or, better still, a metric) in order to define properly what is meant by a neighborhood of a system or of a signal. The property of greatest significance in the study of adaptive control systems for which robustness is desired is the explicit boundedness of closed loop signals. The small changes in the dynamical systems should refer to variations in the plant system to be controlled, so that robustness of the entire adaptive control system implies the preservation of adequate closed loop performance in spite of the target plant being only roughly approximated by the parametrized model set.

For stable linear systems the metric above may be given by the L_∞ norm or maximum value of the frequency response. (For unstable systems a derived L_∞ metric, the graph metric, is used.) Thus robustness is typically measured by the degree of tolerance to the distortion of the frequency response of the systems involved. By dint of linearity, the initial conditions do not play a part in robustness here.

Adaptive Control systems are nonlinear by nature (since the feedback control gain is a function of system output and yet multiplies this output) and so linear notions of robustness of closed loop stability will be inappropriate or, at best, incomplete. Initial conditions need to be taken into account and robustness defined in terms of the function norm, typically an L_p norm, of the deviations between different solutions over the entire time interval. Robustness here refers to the notion that small changes in initial conditions and/or structural properties of the system elements preserve solution properties such as signal bounds and yield only small variation in closed loop performance including adaptive transients. If the adaptive controller is left fixed then these small structural variations may still be measured in terms of frequency response deviations of the linear plant system.

1.3.2 Recent Trends

As keen students of the history and development of adaptive control theory and practice, we have observed several recent trends which indicate that the theoretical development of a general adaptive control method based upon LQG and RLS might not be so stultifying an idea as we seem to suggest earlier.

- Adaptive control practitioners have proposed a series of control design procedures, well exemplified by the Generalized Predictive Control (GPC) method of Clarke *et al.* [CMT87], which have demonstrated that a wide range of difficult process control applications problems are amenable to regulation by an adaptive controller possessing a specifically structured optimal control law with only a handful of adjustable design variables, coupled with an RLS identifier.
- Linear (nonadaptive) control design methods have adopted the goal of robustness, i.e. the ability to perform adequately in spite of the plant being only nominally modeled by its parametrized description. The robustness measures are usually given in frequency domain terms, intersecting with classical notions of gain margins, phase margins and other quantifiers.
- These linear robustness studies have also addressed issues of the robustness of LQG control and of the re-posing of gnostic H_∞ control problems as more standard LQG problems.
- This robustness crusade has also been taken up by the adaptive control community, albeit with differing ground rules, to develop techniques of parameter adaptation which behave well in the face of a variety of system operating condition changes, including on-line feedback control law modification.
- Recent developments in parameter identification theories have also yielded frequency domain modeling formulations which indicate a potential interplay with the linear control robustness ideas.

Indeed, it is our major aim in this work to explore fully the mechanism for the joint application of RLS and LQG techniques to compose adaptive controllers which demonstrate a synergism in their robustness. That is, the robustness of the combination is greater than that achievable by their separate application, because the RLS procedure without cleverly chosen LQG feedback control identifies a model which is inappropriate for feedback

control purposes and, equally, the LQG controller without a well fitting model is unable to achieve adequate performance.

The tools for carrying through this synthesis of adaptive controllers will be to exploit the frequency domain theory of robustness of linear controllers, specifically the theory of LQG Loop Transfer Recovery, simultaneously with the frequency domain theory of linear transfer function identification. But these are the tools of Linear Systems theory, and adaptive control systems are *per se* nonlinear. Thus new methods of describing the near linear behavior of slowly adapting schemes need also to be generated in order that the liberty be granted to appeal contemporaneously to these schools of thought pertaining to robust control and system identification.

In order not to overload the benevolent reader and to place the controller synthesis in its proper technical context, we embark upon our presentation by developing those aspects of the major disciplines which play a central part in either the realization or the support of our controller design method. Thus this book contains chapters dealing with LQG Control, Stability of Receding Horizon Controllers, Robustness of Linear Systems and Parameter Identification, in which we endeavor to convey both a working knowledge of the techniques involved and an understanding of some of the basic principles shoring up the theory. The confluence of ideas takes place in Chapter 7.

Historically, there have been a number of researchers who have proposed the implementation of adaptive LQG controllers in various guises, including Minimum Variance controllers [ÅW73], Generalized Minimum Variance controllers [CG79], and full blown LQG. Most notable amongst these latter are Samson [Sam82], Peterka [Pet84] and Grimble [Gri84]. These people have recognized the inherent design strengths of LQG in its ability to cope with non-minimum phase systems, delays, and relatively high order models, while still preserving comprehensible design variables. Indeed, Peterka adopted his opening stance as one of implementing LQ control via predictors. These methods, however, were often considered as numerically unwieldy in their day. (Our belief is that this is no longer true.) In any event, they were focused more strongly on the properties of the control law, including its implied robustness, without examining the effect of the identification stage. For example, in the early work of Grimble [Gri84] the separate strategies of implicit LQ adaptive control and explicit adaptive control were delineated. Explicit LQ adaptive control, which is the approach advocated in this book, equates with indirect adaptive control where an explicit parametrized plant model is generated and then an LQG problem solved for the control law. Implicit adaptive control, which was implicitly held to be preferable, was a direct approach in which one attempted to generate the LQ adaptive control

law via a signal based minimization in place of the identifier. As the material of the book unfolds, one should be able to interpret the developments as an endorsement and completion of the work of these researchers and as a demonstration of how more modern control design techniques might also be taken logically on board into Adaptive Control, through the joint analysis of their control and identifier interactions.

Throughout the book, however, we derive indefatigable support for our machinations from the recent practical development and acceptance of Predictive Adaptive Control techniques in many process control applications. As mentioned above, we treat GPC as our datum point in the realm of these empirical methods and devote considerable efforts to explaining the connection of such procedures to the broader field of LQG/RLS adaptive control. Our rationale is that these sometimes heuristic techniques have been vindicated by experience and, moreover, have indicated to us just what are the salient aspects of functional adaptive control.

Hence, we draw inspiration and unassailable support from these experimentally generated techniques, while at the same time attempting to provide the theory both to explain their successful behavior and to formalize their design methodology within the framework of adaptive LQG/RLS control. Our dichotomy of stance then becomes apparent — GPC is simultaneously the genesis and the butt of our derivations, the departure point and the way station, the focal point and the aberration. In fact, it is our excuse.

Our reliance upon the GPC methods for support and enlightenment, before dispensing with them altogether in favor of more cerebral pursuits, is the reason for our subtitle.

1.4 A Gedankenexample

To help motivate proceedings and to aid in *a priori* comprehension of our goals in this book, we shall now present a simple example which assists the demonstration of the issues which are the major subjects of our treatise. We shall consider the identification and control of a simple second order system using a first order model. While this situation is much simpler than adaptive control, because we deliberately decouple the identification stage and the control stage, the potentially deleterious effects of incompatible objectives of these stages will be demonstrated. Further, in line with *gedanken-* things, this Gedankenexample is rather extreme in its conception in order that a *reductio ad absurdum* eventuate to illustrate the issues. Later in the book more realistic examples will also be presented.

Suppose that we are provided with a plant system, $P(z)$, to be controlled,




Figure 1.1: Frequency response magnitudes for $P(z)$ (—) and $\hat{P}(z)$ (- - -)

which is described by the following transfer function,

$$P(z) = \frac{z^{-1} + 1.2z^{-2}}{1 - 1.6z^{-1} + 0.68z^{-2}}. \quad (1.1)$$

This is a low-pass, stable, but non-minimum phase system of second order. Further suppose that, as we do not know the plant description, we attempt to model it by a first order model,

$$\hat{P}(z) = \frac{bz^{-1}}{1 + az^{-1}}, \quad (1.2)$$

where a and b are the adjustable parameters, the selection of which is the job of the identifier. The control objective is to place the closed loop pole at the point $z = d$ inside the unit circle. Here we take $d = 0.5$, which is a rather severe objective but, for a Gedankenexample at least, not too outrageous for an already stable low-pass plant such as P .

The discrete frequency response magnitudes of $P(z)$ and a candidate $\hat{P}(z)$ are shown in Figure 1.1. The phases are shown in Figure 1.2. The effect of altering the parameter a through positive values is to affect the spread of the main lobe of the frequency response. The effect of altering b is to change the height of the main lobe. The two systems are roughly




Figure 1.2: Frequency response phases for $P(z)$ (—) and $\hat{P}(z)$ (- - -)

comparable at low frequencies — a phase problem is introduced at high frequencies. It is apparent from the shape of the frequency responses at low frequencies that these two parameters do not permit perfect matching at all points. A compromise is struck between fits at respective frequency values by an identifier.

1.4.1 Open Loop Identification

Here we consider the consequences of identification of a first order model, which is parametrized as $\hat{P}(z)$ above, to the actual second order plant $P(z)$ in several different experimental circumstances. Our aim is to demonstrate the effect of the input signal on this undermodeled fit.

Write the difference equation equivalent to (1.1) with plant input u_t and output y_t ,

$$y_{t+1} = 1.6y_t - 0.68y_{t-1} + u_t + 1.2u_{t-1}. \quad (1.3)$$

If the input sequence $\{u_t\}$ is chosen to be a constant, \bar{u} , then, by stability, the measured plant output sequence $\{y_t\}$ tends exponentially to a constant,

$$\begin{aligned} y_t &= \bar{y} \\ &= P(1)\bar{u} = 27.5\bar{u}. \end{aligned}$$

At the same time we are attempting to fit a model to $P(z)$ and it is clear that a perfect fit is possible, provided we satisfy

$$\hat{P}(1) = \frac{b}{1+a} = P(1) = 27.5. \quad (1.4)$$

That is, any parameter values a and b on a linear variety

$$b = 27.5 + 27.5a$$

would be acceptable as an identified model with this input. This choice of input sequence is inadequate to resolve both parameters, i.e. it leads to nonidentifiability of the model. The difficulty is tied to the fact that a constant input does not excite enough dynamical modes to provide sufficient information about the system to the identifier. Such an input signal is referred to as nonpersistently exciting for the model. We shall throughout this work take measures to ensure that our plant input signals are always persistently exciting since, otherwise, it is difficult to have confidence in the plant information content of the measured data.

For a model complexity of n parameters, one typically requires that the input spectrum contain sufficiently many, at least n in this case, distinct complex frequencies [SB89] in order that it be persistently exciting. Let us consider what happens here when $\{u_t\}$ is chosen to consist of a single frequency sinusoid (i.e. two complex conjugate frequencies), which is a persistently exciting signal for this model,

$$u_t = \cos(\omega t).$$

Now we have that, exponentially,

$$y_t = \operatorname{Re} \left(P(e^{j\omega}) \right) \cos(\omega t) - \operatorname{Im} \left(P(e^{j\omega}) \right) \sin(\omega t). \quad (1.5)$$

Similarly, for the model we would have,

$$y_t^m = \operatorname{Re} \left(\hat{P}(e^{j\omega}) \right) \cos(\omega t) - \operatorname{Im} \left(\hat{P}(e^{j\omega}) \right) \sin(\omega t), \quad (1.6)$$

and a simple calculation shows that a unique solution for a and b exists for which $y_t = y_t^m$,

$$a = -\frac{\operatorname{Re}(P(e^{j\omega}))}{\operatorname{Im}(P(e^{j\omega}))} \sin(\omega) - \cos(\omega) \quad (1.7)$$

$$b = -\frac{\operatorname{Im}(P(e^{j\omega}))}{\sin(\omega)} (1 + a^2 + 2a \cos(\omega)). \quad (1.8)$$

Frequency (radians)	a	b	k
0.1	-0.8005	6.1268	-0.0490
0.2	-0.8657	6.3712	-0.0574
0.4	-1.1979	7.6170	-0.0916
0.8	-14.3764	57.0363	-0.2433
1.0	4.1190	-12.3213	-0.3749

Table 1.1: Identified plant parameters and computed gain

Some sample values are given in Table 1.1 where it is clear that, although a unique identified parameter value exists for any particular input frequency, the specific value is determined by the value of the frequency. The model actually fits the true plant system transfer function exactly at that input frequency. Also shown are the computed controller gains, k , for a closed loop pole at $d = 0.5$.

To determine in detail what happens with model fits based on recursive identification methods with more realistic and complex inputs is really the province of Chapter 6, but here we consider the effect of identification with a system input

$$u_t = \cos(0.1 t) + 0.1 \cos(t). \quad (1.9)$$

(This input has considerable high frequency content — but this is only a Gedankenexample.) Now we find that an exact fit of the model to the plant input–output data is impossible and a compromise must be struck between a fit at frequency $\omega = 0.1$ and at $\omega = 1$. This compromise is found by weighting the respective fits at each frequency by the signal energy at that frequency. The resulting parameters from a Least Squares fit with the input (1.9) is

$$a = -0.8330 \quad \text{and} \quad b = 5.2663, \quad (1.10)$$

yielding a control gain $k = -0.0632$. It should be noted that the model error is not zero at either frequency nor are the parameters a simple modification of their values at the individual separate frequencies.

The point to be made at this juncture is that, in situations of under-modeling, the input signal spectrum determines the identified model based essentially on fits of the model to the true plant at points of support of the spectrum of the input. With inputs containing more complex frequencies than there are free parameters the fit involves the compromised weighted fit at all available frequencies. This important notion of the identified model depending upon the input spectrum will be the linchpin of our derived adaptive control approach.

1.4.2 Control Law Selection

Based upon the identified model, it is natural to construct the feedback controller to cause the identified model to behave in a desirable fashion. Here we suppose that this control objective is to place the closed loop pole (the model is only first order) at the point inside the unit circle, $z = d$. To effect this control one selects the feedback law

$$u_t = ky_t + r_t \quad (1.11)$$

$$= \frac{(a+d)}{b}y_t + r_t, \quad (1.12)$$

where $\{r_t\}$ is an external reference signal. The value chosen for k reflects both the identified model parameters and the overall control objective. Recall that, although designed for the model $\hat{P}(z)$, this controller will be applied to the plant $P(z)$.

The achieved closed loop characteristic polynomial for this controller applied to the real plant $P(z)$ (1.1) is therefore

$$\phi_{cl}(z) = z^2 - (k + 1.6)z + (0.68 - 1.2k). \quad (1.13)$$

The stability of the closed loop is assured, provided

$$0.0364 > k > -0.2667, \quad (1.14)$$

that is, if b is positive,

$$-a + 0.0364b > d > -a - 0.2667b. \quad (1.15)$$

If we take as a design objective that d should equal, say, 0.5 then we see that closed loop stability for $P(z)$ requires

$$-0.5 - 0.2667b < a < -0.5 + 0.0364b.$$

For our identifications with a single sinusoid at frequency less than approximately 0.85 and with the two sinusoids weighted as in (1.9), this condition is satisfied and so the resultant controller would be stabilizing with the real plant. We see, however, that the mixing of distinct frequencies in the input signal can yield significantly different controller values and too high a frequency of the input can even yield unstable control laws for the actual plant. This will be even more apparent when closed loop identification is broached shortly.

The points to note from this inequality above are that the identified model parameters, a and b , and the control objective, encapsulated by d ,

jointly determine the stability and performance of the resultant closed loop, together with the actual plant of course. Clearly, the stability margins in (1.15) reflect the importance of the identification experiment and the specific control design upon the closed loop robustness. This, too, will be a central theme of our study which shall be made much more precise and formal from Chapter 5 onwards. But now we examine the effect of the control law upon the identification in closed loop.

1.4.3 Closed Loop Identification

If such a fixed control law as (1.12) is applied to the plant system then the spectrum of the input signal to the plant, $\{u_t\}$, is determined not just by the external reference but also by the closed loop plant including the control law. Indeed,

$$u_t = [1 - kP(z)]^{-1} r_t, \quad (1.16)$$

so that the relative components of the spectrum of $\{r_t\}$ are altered in their proportions in the construction of $\{u_t\}$.

We remark on several features:

- if $\{r_t\}$ is identically zero, then subject to closed loop stability in (1.16) the control signal $\{u_t\}$ will be exponentially zero and the earlier non-identifiability issues arise. Thus, to achieve persistency of excitation it is necessary to insist upon an exciting reference signal.
- if $\{r_t\}$ is a constant, \bar{r} , then again, with stability, u_t tends exponentially to a constant and persistence of excitation difficulties occur.
- if $\{r_t\}$ is a single frequency sinusoid, then with stability the plant input is asymptotically a sinusoid of the same frequency and so unique identifiability is possible at this frequency value.
- the interesting (and generic) case arises when $\{r_t\}$ has a spectrum with more than one frequency present. In this case, as described above, the identification fits a model with a criterion representing a compromise between those frequencies comprising $\{u_t\}$. Further, the precise nature of this compromise is dictated by the relative energies of these frequencies in $\{u_t\}$.
- the effect of the controller on identification in closed loop now becomes more apparent. Equation (1.16) shows how the actual plant and the control law combine to distort the relative energy content of the different frequencies in $\{r_t\}$. In this way the controller exerts a direct

influence upon the identified parameters, which in turn determine the control law and therefore the energy distribution of $\{u_t\}$. We already divine the interplay which is the focus of this book.

We return now to the example involving two frequency components in the input to an open loop Least Squares identification but now, instead of u_t being given by (1.9), we have

$$r_t = \cos(0.1 t) + 0.1 \cos(t). \quad (1.17)$$

We consider that the plant now operates in closed loop with the feedback control gain $k = -0.0632$, which is that value resulting from the open loop identification with u_t given by (1.9) and which yields a stable closed loop. The weightings on the respective frequency components in the achieved plant input then causes

$$u_t = 10.3761 \cos(0.1 t + \phi_1) + 3.1034 \cos(t + \phi_2),$$

where ϕ_1 and ϕ_2 are corresponding phase shifts which are unimportant here. We note that the ratio of the low input frequency weighting to the high input frequency weighting has been reduced from 10 in the open loop case to approximately 3 in this closed loop case. The identified parameters in the plant with this input are now

$$a = -0.9722 \quad \text{and} \quad b = 1.7528,$$

which demonstrates a shifting emphasis to the model fit at the higher frequency. More importantly, computation of a new feedback controller based on these identified parameter values yields a controller gain $k = -0.2694$ just outside the stability region for the plant. That is, in this quasi-adaptive situation where we identify with a fixed controller, compute a new controller, re-identify under this new feedback and so on, we see that the closed loop stability is quickly lost due to the frequency shift of the model induced by the control law.

1.4.4 Summary

The Gedankenexample of this section is not yet fully an adaptive controller. Rather, it is a separate analysis of a closed loop identifier and of a control design for a system incorporating undermodeling. The purpose of this example is to illustrate the interplay between these separate elements and to foreshadow the likely sources of problems in an adaptive controller, where closed loop identification is conducted simultaneously with the control design based on the identified parameter. The major aspects of the example are summarized as follows:

- Parameter identification of a plant requires persistently exciting inputs in order to admit the resolution of best fitting parameter values.
- The specific parameter value selected by an identifier in the case of undermodeling is determined by the spectral properties of the plant input signal and the relative energies in different frequency bands.
- The existence of a closed loop controller has the effect of altering the relative balance between closed loop plant input spectral energies. Thus, the control law affects the potential convergence points or neighborhoods of the identifier.
- The identified parameters determine the feedback control law through the control design schema. These parameters plus the design objective jointly influence the achieved closed loop performance and stability.
- In adaptive control, there is an interplay amongst the identifier and the control law schema which possesses the possibility of either supporting or frustrating the robust stability of the adaptive closed loop.
- The example, itself, demonstrated a quasi-adaptive experiment in which closed loop identification plus iterated control design led eventually to instability.

It is the consideration of these larger issues in adaptive control which is our brief here.

1.5 The Audience

In keeping with our practical/theoretical dichotomy outlined earlier, we identify two distinct classes of readers for this book and, naturally, propose two different ways to read this work. These broad readership classifications are:

- Practitioners of adaptive GPC style predictive control who would like to know of the body of theoretical material linking their algorithms to mainstream (adaptive and nonadaptive) control design and analysis. For these readers we hope that the book provides a point of access to, or bridge between, adaptive control practice and nonadaptive robust control design.
- Theoreticians interested in appreciating the practically significant advances in adaptive control design, but from their own standpoint. Our aim is to provide these readers with an appreciation of the practical directions of adaptive control.

For the first class of readers we believe the major interest lies in the interpretation of GPC methods within the framework of LQG/RLS, and of the points of contact with linear robustness theory. Their goal could be to comprehend the features underlying the Candidate Robust Adaptive Predictive controller proposed in Chapter 7 and to appreciate the design philosophy underpinning the suggestions made. Their interest need not extend in the first instance to the derivational material nor the technicalities of, for example, ensuring the adequate behavior of parameter identifiers operating in indirect adaptive control. Therefore we have sprinkled the chapters with convincing examples verifying our claims, and with conclusions to our technically developmental chapters taking the form of executive summaries providing encapsulation of the major points of these sections. This should provide access to the design principles relatively quickly and painlessly.

Other readers should, on the other hand, delight in the derivations and theory presented in each section, even though this aspect of our material will be restricted to only those features having direct bearing upon the global thesis of this work — the synthesis of a theoretically supported adaptive control law employing the major features of existing practical methods and current theoretical viewpoints. For these readers, we have included a moderate level of mathematical strictness, especially inasmuch as this is required for the rigorous formalization of the transmogrification of the two linear theories, robustness and identification, to the realm of adaptive control. For them we have included both proofs and technical arguments. Our recommendation is, however, that they skip the examples, since these might otherwise confuse them.

The existence of these two levels of consciousness amongst our potential readership was confirmed in our meetings with colleagues during conferences and colloquia. The former class of readers were often concerned with ‘lifting their game’ theoretically to attempt to comprehend what they had found to be critical in practice, while the latter class of readers’ worries were typically with justifying their funding from corporate sponsors and the lending of weight to their pursuit of excellence. We hope to assist both.

1.6 A Brief Tour

Here we give a short rundown of that material contained in each of the chapters of the book. The purpose is to identify the contents and also to aid the different readers to assess just what sections are pertinent to their pursuits and desires.

Chapter 2 — Generalized Predictive Control

The GPC method of Clarke *et al.* is introduced and studied more or less in its natural habitat of *ad hoc* adaptive control, where its historical connections are made evident, as are the underlying design features of tracking controllers, output predictors and disturbance models. An example is provided and a variant due to Irving is advanced. These formulations are made to familiarize the gentle reader with the context of such GPC procedures before launching into more elevated treatments of these ideas.

Chapter 3 — Linear Quadratic Gaussian Optimal Control

Linear Quadratic Gaussian Optimal Control is advanced in a standard design setting for state-space systems. Firstly the LQ regulator problem is tackled and notions of LQ horizons are introduced before extension to tracking problems is considered within the same framework. The dual problem of Kalman prediction and filtering is then broached and similarly extended to include disturbance models. LQ control and Kalman filter design are jointly incorporated to yield descriptions of LQG feedback controllers. Following this, the GPC control law is re-examined to reveal its nature as a subset of finite horizon LQG. This then sets the stage for the full exploitation of this interpretation of GPC as adaptive LQG optimal control.

Chapter 4 — Stability and Performance Properties of Receding Horizon LQ Control

With the demonstration of GPC as a receding horizon LQG strategy under our belts, we progress to the first formal application of LQ stability theory to GPC in an attempt to derive analytical and design methods to ensure closed loop stability and performance of the GPC strategy. The recent Fake Algebraic Riccati Techniques of Poubelle are applied to effect the presentation of guaranteed stable well performing controllers based on state-variable methods. Some comparative points are made between the standard GPC solution and what might be achievable by more erudite procedures, which need only be slightly more numerically demanding. We begin to discern the thematic focus of the book.

Chapter 5 — Robust LQG Design – Features for Adaptive Control

Since robustness will be a central property desired of our closed loop system operating under adaptive LQG and RLS, we present in this chapter the rudiments of linear control system robustness in general and then the specifics of

LQG robustness enhancement techniques such as Loop Transfer Recovery. The fundamental robustness theorem indicates how modeling properties and controller features are traded to secure closed loop robust stability, and the prospect for an adaptive controller to take advantage of this is hinted at.

Chapter 6 — Recursive Least Squares Identification in Adaptive Control

To complement the linear system robustness results developed in Chapter 5, the closed loop behavior of Recursive Least Squares identification is studied here from a frequency domain modeling viewpoint. There are two main features: the inclusion of a feedback control law dependent upon the parameters causes a change in the effective identification criterion minimized, while the nonlinearity introduced in adaptive control leads to a fundamental change in the nature of the dynamics. These issues are resolved via quantifications on adaptation speed and on the ultimate modeling accuracy and sensitivity.

Chapter 7 — A Candidate Robust Adaptive Predictive Controller

The synthesis of a Candidate Robust Adaptive Predictive controller is made based upon the preceding robustness and estimation theories, in such a fashion that the combined system possesses a mutually supporting control law and identification criterion. The rationale is presented for the design advanced and, perhaps not too surprisingly, it bears many of the hallmarks of the familiar GPC procedures, except that within this framework the design choices are more apparent.

Chapter 8 — Le Jugement Dernier

Following the specifics of the candidate controller design via LQG and RLS, the questions are now raised about how these methods might be more generally extended. Linear system stability robustness is presented in a standardized fashion for several different classes of model uncertainty and for general feedback control laws. The earlier LQG-specific results of Chapter 5 are extended and then conclusions drawn about identifier behavior and feedback system stability. Following this we develop the connection between closed loop performance and adaptation, since after ensuring stability this is the ultimate objective for adaptive control. We conclude with a short exploration of further issues. Particulars dealt with include the outlook for adaptive H_∞ control and for the use of more esoteric robust control design strategies to be adopted in the adaptive context. The future consequences of our design philosophy are briefly speculated on without too much brouhaha.

Chapter 2

Generalized Predictive Control

2.1 Introduction

2.1.1 Motivation

Before anybody throws the objection at us, let us make a frank admission: there is no solid justification for commencing this book with a chapter on Generalized Predictive Control, or GPC as it is now commonly called, given that we will later diverge to a considerable degree from this control design method to wander into the more powerful methods of Adaptive Optimal Control. If we decide to start with an introduction into predictive control methods, it is not just to assert the authors' rights to discuss any subject they fancy. The facts of industrial applied adaptive control are that GPC has found a market niche in providing a justifiable, sophisticated yet flexible and comprehensible adaptive controller design platform, which has met with acceptance from practitioners and theoreticians alike. We shall use this feature of GPC to motivate our foray into Adaptive Optimal Control from a more general viewpoint and with the provision of a theoretical basis for the selection of design variables. In view of the global messages we would like to convey in this book, there are three good reasons to start with a chapter on predictive control.

Firstly, even though GPC will be shown to be a simple (some would say simplistic) formulation of adaptive optimal control, it exhibits many of the salient features of the more sophisticated adaptive LQG controller: they are both derived from quadratic optimal control criteria based on a designed trade-off between control performance and control energy consumption, lead-

ing to dynamic output feedback controllers with built-in observer designs. Secondly, the rather parsimonious design criterion of GPC has quickly made practitioners realize that to achieve the goals it was set to achieve required ever more sophisticated modifications and add-ons. At the same time it was discovered that the manipulation of the design parameters, which at first were thought to have intuitively clear effects on overall performance, proved to be much harder to understand than was initially expected. This raised more and more questions about closed loop stability, tracking performance, robustness to unmodeled dynamics and other dynamical properties. This makes GPC an excellent starting point and motivational example for our subsequent analysis of the broader class of LQG controllers. Finally, the ultimate result of our robustness analysis of adaptive LQG methods turns out to vindicate many of the simple design principles behind GPC, while also suggesting that others should be dropped or modified. Thus, the predictive control methods, which were initially derived on the basis of intuitive rather than theoretical ideas with the overall simplicity being of paramount importance, will be shown in this book to have some remarkable features which can be theoretically justified on the basis of robust control design theory coupled with adaptive recursive least squares identification. This probably explains the practical successes obtained by these adaptive control methods and the tenacity of their proponents in their pursuit of ever more sophisticated modifications to the initially simple design criterion in order to rescue the basic schema.

But there is a fourth motivation, that is more connected to the authors' self-imposed rôle within the adaptive control community than to the requirements of a tutorial presentation of the material of this book. As theoreticians in adaptive control who keep close contacts with our more practically oriented colleagues, we were surprised to discover that predictive control practitioners were claiming remarkable industrial successes and developing ever more complicated variations of these predictive control methods without much theoretical support, while most theoreticians were staying on the fence and were looking at these bizarre developments with a certain degree of condescension. Our initial desire therefore was to understand and analyse GPC, without any preconceptions or prejudices, and since we believe that we have now understood, we think we can provide a useful service by exposing our understanding and by following in this book the path that led us from GPC to LQG, from closed loop stability to closed loop robustness, and from the nonadaptive to the adaptive versions of these methodologies.

Our aim is certainly not to write a comprehensive survey on predictive control methods and their many variants. It is therefore unlikely that the

predictive control practitioners will find much in this chapter that they don't already know. For the others, we believe that it is a chapter well worth reading, since it will both enhance their general culture in an arena that has proved its importance, and it will provide motivation for many of the questions that will be addressed in the later chapters.

2.1.2 What is Predictive Control?

Predictive methods in adaptive control refer to a collection of control design formulations that pose control criteria at a given time explicitly in terms of predictions of future plant outputs and sometimes of future plant inputs. Because such predictions become more difficult as they become more distant in time, these criteria are typically finite horizon optimal control criteria.

The history of predictive methods in adaptive control is mixed with many suggestions for their use in various guises coming independently from several quarters. Early versions bearing such esoteric names as IDentification and COMmand (IDCOM), Dynamic Matrix Control (DMC), Internal Model Control (IMC), Predictor-based Self Tuning Control, Extended Horizon Adaptive Control (EHAC), Model Algorithmic Control (MAC), MUSMAR, EPSAC, etc. were suggested by Richalet *et al.* [RRTP78], Cutler and Ramaker [CR80], Garcia and Morari [GM82], Peterka [Pet84], Ydstie [Yds84], Rouhani and Mehra [RM72], Mosca *et al.* [MZL89], Greco *et al.* [GMMZ84], De Keyser and Van Cauwenberghe [dKvC85], Maurath *et al.* [MSM85] and others. All these methods have certain features in common which distinguish them from previous design philosophies — the solution of a finite horizon optimization problem at each time instant implemented in a receding horizon way, the incorporation of plant output predictions, the provision of a small number of design parameters connected to various degrees with the closed loop dynamics. We shall not fall into the trap of trying to decide who did what first, but we refer to Garcia *et al.* [GPM89] for a comprehensive survey of theoretical derivations and practical implementations and successes of these predictive methods.

The version which appears to have had the most acceptability is that derived by Clarke, Mohtadi and Tuffs [CMT87] and called GPC for Generalized Predictive Control. The moniker GPC has since been adopted as the popular collective denomination of the whole class of long range predictive methods in adaptive control. We shall adopt the GPC method of Clarke *et al.* both as our representative of this set of control designs and as our particular reference datum in developing the explanations and comparisons of this book. Thus we regard GPC both as a specific algorithm and as a sobriquet for the class of related procedures. Most of the results which we

develop within this work do not hinge upon specific features of one particular algorithm of that class.

The practical success of adaptive predictive control methods is well reported in the literature: see for example Clarke [Cla88], Garcia and Prett [GP86], Martin *et al.* [MCA86], Richalet *et al.* [RRTP78], Seborg *et al.* [SSE86] and others. As with many practical adaptive control design procedures, predictive methods focus upon the control law design component of the adaptive controller almost to the exclusion of mentioning the identification component. To quote from one of the producers of the original GPC algorithm,

The ‘Achilles’ heel’ of its self-tuning version, however, is the identification algorithm which was originally introduced more as an *afterthought* than as an integral part of the design [SMS90].

We shall endeavour to make the connection between controller design and adaptive identifier design in the development of our theory. We shall thereby indicate the potential synergism between these components, which might explain their performance in applications, and we shall further indicate how this might be improved. First, we spend some more time on the history of the GPC algorithm as presented by Clarke, Mohtadi and Tuffs since this will give credence to some of the conclusions that we shall obtain later in the book.

2.1.3 A Brief Historical Perspective

The antecedent of GPC is the well-known Minimum Variance Controller, described in a nonadaptive version by Åström [Åst70] and which formed the basis of the famous Self-Tuning Regulator of Åström and Wittenmark [ÅW73], a variant of which was also analysed by Goodwin, Ramadge and Caines [GRC80]. This controller is obtained by minimizing, for a given linear input–output model, the following criterion,

$$J(u, t) = \text{E} \left\{ (y_{t+1} - r_{t+1})^2 \right\}, \quad (2.1)$$

where y_t is the output of the system, r_t is a reference signal and $\text{E}(\cdot)$ denotes expectation. This criterion is minimized at time t by the selection of the control signal, u_t . At time $t + 1$ a new problem is solved for u_{t+1} . The criterion (2.1) is based on the implicit assumption that the plant model has a unit delay; it would indeed be senseless to attempt to minimize the tracking error at time $t + 1$ if the control u_t affects the plant output only at time $t + d$, say, with $d > 1$. If the plant has a delay d , the error in

(2.1) is replaced by $y_{t+d} - r_{t+d}$. Since handling plants with delays causes no additional conceptual complication (provided this delay is known), and does not alter the qualitative consequences of the analysis, we shall in future consider only the case of plants with a unit delay.

As is well known, this control strategy works only for minimum phase systems (i.e. systems with stable plant zeros). For non-minimum phase plants this control law suffers from demanding excessive control input in order to effect the optimal output variance, i.e. the controller achieves its performance through the cancellation of the plant zeros (including the unstable zeros) which leads to a loss of internal stability of the feedback system. One way which has been suggested to make the same strategy also work for non-minimum phase systems is to modify slightly the criterion (2.1) through the inclusion of a penalty on the control signal as well as on the output. This yields the new criterion,

$$J(u, t) = \text{E} \left\{ (y_{t+1} - r_{t+1})^2 + \lambda u_t^2 \right\}, \quad (2.2)$$

which has been called Generalized Minimum Variance control (GMV) [CG79] and which implements a one-step-ahead optimal control law. While this strategy has the potential to produce an internally stabilizing control law, this still is not guaranteed for specific choices of λ . Indeed, even if the plant is known perfectly well, the stability analysis of this controller requires resort to root locus techniques.

Another frequent modification made to this GMV control law is the statement of the optimization problem, not in terms of u_t , but in terms of Δu_t , where Δu_t is the incremental control input of the system,

$$\begin{aligned} \Delta u_t &= u_t - u_{t-1} \\ J(u, t) &= \text{E} \left\{ (y_{t+1} - r_{t+1})^2 + \lambda (\Delta u_t)^2 \right\}. \end{aligned} \quad (2.3)$$

At time t , this criterion is then minimized with respect to Δu_t . The reason for using incremental inputs in the criterion is that (2.2) does not admit zero static error in the case of a non-zero constant reference unless the open loop plant contains an integrator, which would allow y_t to remain at a non-zero constant value with the control input being zero. For the standard GMV problem one always penalizes non-zero u_t even in the case of tracking a non-zero reference. One feature of using Δu_t in the control law specification is that explicit bounds on u_t may be difficult to obtain. We note finally that this new control strategy appears really as a variant of the MV controller. Thus the idea is generally to keep λ as small as possible in order to remain as close as possible to the objective of maintaining the output variance minimal,

while still maintaining closed loop stability. The positive weighting λ is included simply to prevent control signal explosion.

Since this controller again fails for some unstable and non-minimum phase systems, and particularly for systems with poorly known delays, an additional extension was made by Clarke *et al.*, and led to the GPC which minimizes the following criterion:

$$J(u, t) = E\left\{ \sum_{j=N_1}^{N_2} [y_{t+j} - r_{t+j}]^2 + \lambda \sum_{j=1}^{N_u} [\Delta u_{t+j-1}]^2 \right\} \quad (2.4)$$

$$\text{subject to } \Delta u_{t+i} = 0, \quad i = N_u, \dots, N_2.$$

This minimization produces $\Delta u_t, \Delta u_{t+1}, \dots, \Delta u_{t+N_u-1}$, but only Δu_t is actually applied. At time $t+1$ a new minimization problem is solved. This implementation is called Receding Horizon Control. For a known time-invariant system, this yields a time-invariant controller.

The rationale for this modification of the control objective should be interpreted as an embellishment of the MV and GMV rules to attempt to persuade the solution to provide both adequate performance and asymptotic stability for a wider class of potential plants. The optimization is performed on several $(N_2 - N_1 + 1)$ successive future output values taking into account several (N_u) future incremental control actions. The presumption is that this algorithm will work provided the ‘true’ delay is included in the interval between N_1 and N_2 . Usually N_1 is chosen as the delay or a lower bound of its estimate; it is thus not really used as a design parameter. The inclusion of several future control values is made in order that longer observation of the signal will mitigate against its becoming unbounded. The addition of the constraint on the control increment after a certain time is rationalized on the grounds of encouraging the controller to achieve its performance quickly, but justified on the basis of the computational features of the solution. We note, however, that whatever the reasoning behind this scheme and its relations, it still represents the statement of a finite horizon (N_2) criterion, the solution of which is implemented and assessed over the infinite horizon, via stability and tracking variance evaluation.

What is remarkable about the control criterion (2.4) is that it possesses a significant level of complexity sufficient to make it capable of producing effective controllers for enormous ranges of candidate plants, while the criterion itself depends upon the specification of only three main design parameters, N_2 , N_u and λ . This is a hallmark of many sophisticated (some would say ‘sophistical’) control design procedures, such as LQG to be discussed later, where the constructivist phase of the design is transposed from the solution for the direct parameters of the controller to the issue of specifying those

free variables of the problem formulation. The particular solution for the controller from this point then becomes simply computational. The requirement for these design variables is that their influence upon the dynamical properties of the solution should be easily assessed and, further, that guidelines exist for their selection to effect certain closed loop properties. This is only partly true for GPC.

This chronology clearly shows the link between GPC and Minimum Variance control and explains why in all subsequent uses of this algorithm small values of λ have generally been advocated. Notice that, so far, little has been said about the adaptive nature of these algorithms. This is not surprising when considering the Minimum Variance Regulator which has in fact been introduced in a nonadaptive context [Åst70], but is a little more peculiar for the subsequent two algorithms which have originally been presented specifically as adaptive control algorithms. However, as already noted, in the presentations of GMV [CG79] and GPC [CMT87] the adaptive nature actually does not play any major rôle, with the certainty equivalence principle being invoked after a nonadaptive presentation of the control algorithm. For ease of presentation, we shall also describe the GPC algorithm as a nonadaptive control method at first, but we stress the fact that it is the adaptive nature of the complete algorithm, i.e. the interplay between a recursive prediction error identification method and the control design based on the minimization of the predictive criterion, which actually gives credit to the GPC method and for which the method has rightly received its bouquets for practical control applications.

2.2 The GPC Method of Clarke *et al.*

The GPC method is a control procedure which is applicable to both single-input/single-output and multi-input/multi-output processes. Indeed, essentially all of the theory treated within this book will be applicable to multi-input/multi-output processes. It is merely for the sake of presentational simplicity that we deal with single-input/single-output systems. We begin by supposing that a model of the linear (or linearized) plant is given in the following ARIMAX (Auto-Regressive Integrated Moving-Average eXogenous input) form:

$$A(q^{-1})y_t = B(q^{-1})u_{t-1} + \frac{C(q^{-1})}{\Delta}\xi_t, \quad (2.5)$$

where, as usual, u_t , y_t and ξ_t are the plant input signal, output signal and disturbance process, respectively, at time t , and A , B and C are polynomials in the unit delay operator, q^{-1} , with A and C monic.

The rôle of the Δ operator ($\Delta = 1 - q^{-1}$) is to ensure integral action in the controller in order to cancel the effect of step output disturbances. The disturbance signal ξ_t may be either a deterministic or a stochastic signal but, because of the Δ operator, its mean value is assumed to be zero. In the development of the GPC approach, ξ_t is assumed stochastic; the polynomial $C(q^{-1})$ can then always be taken as a stable polynomial, since only the spectral properties of the signal $(C(q^{-1})/\Delta)\xi_t$ influence the predictions of future values of y_t . Thus the nature of the plant description (2.5) is to model the output as being corrupted by the effects of an additive random walk process. While this may not be a realistic model, its effect upon the controllers derived from it will be to force the ability to reject step output disturbances. This is entirely a reflection of the known nature of load perturbations arising in the process control industry where these control laws have found such successful application. Thus the GPC incorporates this noise model directly into its formulation in order to tailor its response for particular circumstances. This is a manifestation of the Internal Model Principle.

The system described by equation (2.5) can be equivalently represented by the following equation:

$$A(q^{-1})\Delta y_t = B(q^{-1})\Delta u_{t-1} + C(q^{-1})\xi_t. \quad (2.6)$$

The differencer Δ and the C polynomial itself then play a rôle in the derivation of plant output predictors defined below.

With this linear model specified, a partially constrained quadratic optimal control criterion is posed in terms of the incremental inputs and outputs (and not explicitly the state) of the plant. The cost function to be minimized is:

$$J(u, t) = E\left\{ \sum_{j=N_1}^{N_2} [y_{t+j} - r_{t+j}]^2 + \lambda \sum_{j=1}^{N_u} [\Delta u_{t+j-1}]^2 \right\} \quad (2.7)$$

$$\text{subject to } \Delta u_{t+j} = 0, \quad j = N_u, \dots, N_2$$

where N_1 is the minimum costing horizon, N_2 is the maximum costing horizon, and N_u is the control costing horizon. The signal r_t is the reference signal which we desire the system output to track. The positive constant λ weights the relative importance of control and tracking error energies. The use of the expectation in (2.7) is made to indicate that the control values chosen are calculated conditioned on data available up to and including time t and presuming the stochastic disturbance model. Thus the control design component of GPC involves the solution of a standard finite horizon optimal control problem. We shall return in Chapter 3 to discuss further the connections between the GPC problem formulation and this

more familiar optimal control problem. As such, the GPC control strategy is an open loop control policy, since N_u future control increments are computed explicitly through the minimization of (2.7). However, at time t , one solves this optimization problem with criterion $J(u, t)$ for the control strategy $\{\Delta u_{t+j}, j = 0, \dots, N_u - 1\}$ but one applies only the first element, $u_t = u_{t-1} + \Delta u_t$. A new problem is solved at time $t + 1$, with criterion $J(u, t + 1)$ and solution $\{\Delta u_{t+1+j}, j = 0, \dots, N_u - 1\}$, and only u_{t+1} is applied. In a nonadaptive implementation with a time-invariant model, this leads to a time-invariant controller, as already noted. This re-solution of the optimization problem at each instant indicates how this seemingly open loop strategy is in fact implemented in closed loop. The incorporation of all preceding plant input and output information into the generation of the predictions necessary for the solution of these optimizations is the mechanism for closing the loop. This will be amply demonstrated shortly by calculation and by example.

2.3 Optimal Prediction and GPC Solution

To solve the problem posed by the minimization of (2.7), we have to compute a set of j -step ahead predictions of the output y_{t+j} , for $j = N_1, \dots, N_2$, based on information known at time t and on the future values of the control increments, which we will later choose so that the GPC criterion J is optimized. These predictions involve the use of Diophantine¹ equations arising from the plant ARIMAX model as is standard in the theory of prediction of such stochastic processes, see [ÅW84] and [GS84]. Specifically, to compute the j -step ahead output prediction one solves the following Diophantine equation,

$$C(q^{-1}) = E_j(q^{-1})A(q^{-1})\Delta + q^{-j}F_j(q^{-1}), \quad (2.8)$$

where each of the variables is polynomial in q^{-1} and $\deg(E_j) = j - 1$. Using (2.6) and dropping the explicit arguments in q^{-1} , we then obtain

$$y_{t+j} = \frac{B}{A}u_{t+j-1} + E_j\xi_{t+j} + \frac{F_j}{A\Delta}\xi_t.$$

¹Recent nationalistic historical investigations have led to discussions of the priority of certain mathematicians in the study of these equations over integral domains. Vidyasagar [Vid85] has attributed them to Aryabhata rather than to Diophantus. We have applied different investigative methods, more attuned to our view of history, and conducted a poll. We would like to reveal here that, by popular acclaim, these equations were first written down by a Mr Eddy Van Compennolle of Linkebeek, Belgium, in 1973.

If we replace ξ_t using (2.6), this yields

$$y_{t+j} = \frac{F_j}{C}y_t + \frac{E_j B}{C}\Delta u_{t+j-1} + E_j \xi_{t+j},$$

where the last term contains information which is independent from signals measurable at time t . It is then obvious that the minimum variance prediction of y_{t+j} given data known at time t is obtained by replacing the last term by zero, yielding

$$\hat{y}_{t+j} = \frac{F_j}{C}y_t + \frac{E_j B}{C}\Delta u_{t+j-1}. \quad (2.9)$$

In this expression \hat{y}_{t+j} is a function of known signal values at time t and also of future control inputs which have yet to be computed. We then use a second Diophantine equation to distinguish between past and future control values,

$$E_j(q^{-1})B(q^{-1}) = G_j(q^{-1})C(q^{-1}) + q^{-j}\Gamma_j(q^{-1}), \quad (2.10)$$

which yields the following expression for the prediction,

$$\hat{y}_{t+j} = G_j\Delta u_{t+j-1} + \Gamma_j u_{t-1}^f + F_j y_t^f, \quad (2.11)$$

with u_t^f and y_t^f being filtered versions of Δu_t and y_t ,

$$u_t^f = C^{-1}(q^{-1})\Delta u_t, \quad (2.12)$$

$$y_t^f = C^{-1}(q^{-1})y_t. \quad (2.13)$$

Equivalently,

$$\hat{y}_{t+j} = G_j(q^{-1})\Delta u_{t+j-1} + \hat{y}_{t+j|t} \quad (2.14)$$

where $\hat{y}_{t+j|t}$ is the free response prediction of y_{t+j} assuming that future control increments after time $t-1$ will be zero,

$$\hat{y}_{t+j|t} = \Gamma_j(q^{-1})u_{t-1}^f + F_j(q^{-1})y_t^f. \quad (2.15)$$

The first set of Diophantine equations (2.8) are classical to define optimal predictions [Åst70]. It can be shown that the C polynomial arising in (2.9) determines the dynamics of the observer equivalent to that obtained by a steady-state Kalman filter. The idea behind this remark is that instead of taking the C polynomial as the one delivered by some identification procedure of a real process, it can also be chosen by the user as specifying the observer dynamics.

One may interpret (2.8) and (2.10) as implementing polynomial division, with j determining the degree of the quotient required. Thus one has for (2.8),

$$C(A\Delta)^{-1} = E_j + q^{-j}F_j(A\Delta)^{-1},$$

and hence, substituting for E_j in (2.10),

$$\begin{aligned} CB(A\Delta)^{-1} &= BE_j + q^{-j}BF_j(A\Delta)^{-1} \\ &= G_jC + q^{-j}\Gamma_j + q^{-j}BF_j(A\Delta)^{-1} \\ B(A\Delta)^{-1} &= G_j + q^{-j}\Gamma_jC^{-1} + q^{-j}BF_j(A\Delta C)^{-1}. \end{aligned} \quad (2.16)$$

Note that these multiple Diophantine equations may all be solved recursively, i.e. using simple order update expressions, which results in considerable computational savings over their individual and several calculation [FDR88]. It can also be seen from (2.16) that the polynomial G_j contains the first j Markov parameters g_i of the plant model transfer function $B/A\Delta$.

Define the vector \mathbf{f} , composed of the ‘free response’ predictions,

$$\mathbf{f} = [\hat{y}_{t+1|t}, \hat{y}_{t+2|t}, \dots, \hat{y}_{t+N_2|t}]^T, \quad (2.17)$$

i.e. the predictions of \hat{y}_{t+k} , $k = 1, \dots, N_2$, given $\{u_{s-1}, y_s; s \leq t\}$ assuming that $\{u_{t+k} = 0, k = 0, \dots, N_2 - 1\}$. Next define the vector of future control increments $\tilde{\mathbf{u}}$ (recall that we have set $\Delta u_{t+j} = 0$ for $j \geq N_u$),

$$\tilde{\mathbf{u}} = [\Delta u_t, \Delta u_{t+1}, \dots, \Delta u_{t+N_u-1}]^T, \quad (2.18)$$

and define the vector of predicted controlled plant outputs,

$$\hat{\mathbf{y}} = [\hat{y}_{t+1}, \hat{y}_{t+2}, \dots, \hat{y}_{t+N_2}]^T. \quad (2.19)$$

From the prediction equations (2.14) the predicted input–output relationship of the plant can be written as the vector equation,

$$\hat{\mathbf{y}} = \mathbf{G}\tilde{\mathbf{u}} + \mathbf{f}, \quad (2.20)$$

where the matrix \mathbf{G} is composed of the impulse response parameters, g_i , of the plant model $B/A\Delta$,

$$\mathbf{G} = \begin{pmatrix} g_0 & 0 & \dots & 0 \\ g_1 & g_0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ g_{N_u-1} & g_{N_u-2} & \dots & g_0 \\ \vdots & \vdots & & \vdots \\ g_{N_2-1} & g_{N_2-2} & \dots & g_{N_2-N_u} \end{pmatrix}. \quad (2.21)$$

(The dimensions of \mathbf{G} are $N_2 \times N_u$, since we have taken into account the constraints on Δu_{t+j} for $j \geq N_u$ and have taken N_1 equal to one for simplicity in (2.21). The effect of altering N_1 is to delete rows from the top of G . Also, we have denoted the first delayed impulse response parameter g_0 since, by assumption, the plant has unit delay.) The quadratic minimization of (2.7) now becomes a direct problem of linear algebra, with

$$J = (\hat{\mathbf{y}} - \mathbf{r})^T (\hat{\mathbf{y}} - \mathbf{r}) + \lambda \tilde{\mathbf{u}}^T \tilde{\mathbf{u}} \quad (2.22)$$

and the solution of the future incremental control vector $\tilde{\mathbf{u}}$ is

$$\tilde{\mathbf{u}} = (\mathbf{G}^T \mathbf{G} + \lambda \mathbf{I})^{-1} \mathbf{G}^T (\mathbf{r} - \mathbf{f}), \quad (2.23)$$

where \mathbf{r} has been obviously defined from the reference signal as

$$\mathbf{r} = [r_{t+1}, r_{t+2}, \dots, r_{t+N_2}]^T. \quad (2.24)$$

Equation (2.23) yields the future control increments for times t to $t + N_u - 1$ as an open loop strategy based upon information available at time t . The mechanism utilized for closing the loop and forcing a feedback control in GPC is to implement only the first element of $\tilde{\mathbf{u}}$, i.e. Δu_t , and then to recompute the solution to the optimal control problem again for the next step using data available at time $t + 1$ in the specification of \mathbf{f} . This procedure is known as Receding Horizon Control and is, by now, a classical method [Tho75]. With regard to (2.23), the effect of implementing a receding horizon control law is that the control gain calculated in (2.23) remains fixed and only the vectors \mathbf{f} and \mathbf{r} are updated from time instant to time instant.

The combined closed loop control strategy of GPC is embodied in the design stage, incorporating the calculation of the gain matrices and filters, together with the implementation involving signal construction of \mathbf{f} from (2.15)–(2.17) and control computation, $\tilde{\mathbf{u}}$, via (2.23), with only Δu_t actually implemented. We next devote a short section to the presentation of a simple example which should help to clarify the picture.

2.4 A Simple GPC Example

We consider here the GPC control of a plant with ARIMAX description,

$$y_t = 1.7y_{t-1} - 0.7y_{t-2} + 0.9\Delta u_{t-1} - 0.6\Delta u_{t-2} + \xi_t, \quad (2.25)$$

with ξ_t a zero mean white noise signal. With the notation of (2.5), this corresponds with

$$\begin{aligned} A(q^{-1}) &= 1 - 0.7q^{-1} \\ B(q^{-1}) &= 0.9 - 0.6q^{-1} \end{aligned}$$

Note that we have taken the C polynomial unity in this plant model, and that the differentiator Δ has been incorporated into the A polynomial. Further, we consider prediction and control horizon parameters to be $N_1 = 1$, $N_2 = 3$ and $N_u = 3$, and take the control weighting constant $\lambda = 0.1$.

Predicted plant outputs up to a range of three are therefore needed. These may be derived by back substitution into the plant model (2.25), with $\hat{\xi}_{t+i} = 0$, $i \geq 1$, as follows:

$$\begin{aligned}\hat{y}_{t+1} &= 1.7y_t - 0.7y_{t-1} + 0.9\Delta u_t - 0.6\Delta u_{t-1} \\ \hat{y}_{t+2} &= 1.7\hat{y}_{t+1} - 0.7y_t + 0.9\Delta u_{t+1} - 0.6\Delta u_t \\ &= 2.19y_t - 1.19y_{t-1} + 0.9\Delta u_{t+1} + 0.93\Delta u_t - 1.02\Delta u_{t-1} \\ \hat{y}_{t+3} &= 1.7\hat{y}_{t+2} - 0.7\hat{y}_{t+1} + 0.9\Delta u_{t+2} - 0.6\Delta u_{t+1} \\ &= 2.533y_t - 1.533y_{t-1} + 0.9\Delta u_{t+2} + 0.93\Delta u_{t+1} + 0.951\Delta u_t \\ &\quad - 1.314\Delta u_{t-1},\end{aligned}$$

or by the computation of the predictor polynomials à la [CMT87],

$$\begin{aligned}1 &= E_1(q^{-1})A(q^{-1})\Delta + q^{-1}F_1(q^{-1}) \\ 1 &= 1(1 - 1.7q^{-1} + 0.7q^{-2}) + q^{-1}(1.7 - 0.7q^{-1})\end{aligned}$$

and similarly, following the schema (2.8) and (2.10),

$$\begin{aligned}1 &= (1 + 1.7q^{-1})(1 - 1.7q^{-1} + 0.7q^{-2}) + q^{-2}(2.19 - 1.19q^{-1}) \\ 1 &= (1 + 1.7q^{-1} + 2.19q^{-2})(1 - 1.7q^{-1} + 0.7q^{-2}) \\ &\quad + q^{-3}(2.533 - 1.533q^{-1}),\end{aligned}$$

$$\begin{aligned}E_1B &= 0.9 - 0.6q^{-1} \\ E_2B &= 0.9 + 0.93q^{-1} - 1.02q^{-2} \\ E_3B &= 0.9 + 0.93q^{-1} + 0.951q^{-2} - 1.314q^{-3}.\end{aligned}$$

This yields finally the form of (2.20),

$$\begin{aligned}\begin{pmatrix} \hat{y}_{t+1} \\ \hat{y}_{t+2} \\ \hat{y}_{t+3} \end{pmatrix} &= \begin{pmatrix} 0.9 & 0 & 0 \\ 0.93 & 0.9 & 0 \\ 0.951 & 0.93 & 0.9 \end{pmatrix} \begin{pmatrix} \Delta u_t \\ \Delta u_{t+1} \\ \Delta u_{t+2} \end{pmatrix} \\ &\quad + \begin{pmatrix} 1.7y_t - 0.7y_{t-1} - 0.6\Delta u_{t-1} \\ 2.19y_t - 1.19y_{t-1} - 1.02\Delta u_{t-1} \\ 2.533y_t - 1.533y_{t-1} - 1.314\Delta u_{t-1} \end{pmatrix},\end{aligned}$$

where the matrices G , $\tilde{\mathbf{u}}$ and \mathbf{f} are explicitly identified. The gain matrix may now be calculated,

$$(G^T G + \lambda I_3)^{-1} G^T = \begin{pmatrix} 0.8947 & 0.0929 & 0.0095 \\ -0.8316 & 0.8091 & 0.0929 \\ -0.0766 & -0.8316 & 0.8947 \end{pmatrix}.$$

Further carrying through the calculation of the first row above with the coefficients of y_t , y_{t-1} and Δu_{t-1} in the vector \mathbf{f} given earlier yields the receding horizon closed loop control law,

$$\begin{aligned} \Delta u_t = & -0.644\Delta u_{t-1} + 1.7483y_t - 0.7513y_{t-1} \\ & + 0.8947r_{t+1} + 0.0929r_{t+2} + 0.0095r_{t+3}. \end{aligned} \quad (2.26)$$

This is a time-invariant controller, since the model is time-invariant. With the specification of the closed loop control via (2.26), we divine the connection between the GPC formulation and more normal linear controller designs. In particular we see, with this specific plant and controller specification, that the controller achieved is a second order dynamical system acting on the plant output and on the reference r_t . We shall return to study this example from an optimal control viewpoint in the following chapter and we now further motivate the reason for this book by developing some closed loop formulae for this Generalized Predictive Control strategy.

2.5 Closed Loop Formulae

We will now derive some closed loop formulae for the Generalized Predictive Controller in the nonadaptive case in order to show how the design parameters N_1 , N_2 , N_u and λ might affect the stability of the controlled plant. We first define α_i , $i = 1, \dots, N_2$, as the coefficients of the first row of the matrix in (2.23),

$$(\alpha_1 \quad \dots \quad \alpha_{N_2}) = (1 \quad 0 \quad \dots \quad 0) (\mathbf{G}^T \mathbf{G} + \lambda \mathbf{I})^{-1} \mathbf{G}^T. \quad (2.27)$$

It is then easy to see that the first component of the vector equation (2.23) can be rewritten as

$$\Delta u_t = - \sum_{i=1}^{N_2} \alpha_i \frac{\Gamma_i}{C} \Delta u_{t-1} - \sum_{i=1}^{N_2} \alpha_i \frac{F_i}{C} y_t + \sum_{i=1}^{N_2} \alpha_i r_{t+i}. \quad (2.28)$$

This then leads to the controller equation

$$(C + \sum_{i=1}^{N_2} \alpha_i \Gamma_i q^{-1}) \Delta u_t = (C \sum_{i=1}^{N_2} \alpha_i q^{-N_2+i}) r_{t+N_2} - (\sum_{i=1}^{N_2} \alpha_i F_i) y_t, \quad (2.29)$$

or, with obvious definitions for the polynomials R , T_1 and S ,

$$R\Delta u_t = CT_1 r_{t+N_2} - Sy_t. \quad (2.30)$$

This shows that GPC is a way of synthesizing linear feedback controllers by means of an optimization criterion instead of, say, a pole placement design. Of course, every linear controller, whatever the synthesis technique, will have a structure like (2.30).

Next consider the closed loop which is obtained by combining the plant model (2.6) with the controller equation (2.30). This immediately gives the following closed loop,

$$(A\Delta R + BSq^{-1})y_t = BCT_1 r_{t+N_2-1} + CR\xi_t. \quad (2.31)$$

Using the definitions of R , S , and T_1 and after some calculations, one can show that:

$$\begin{aligned} A\Delta R + BSq^{-1} &= A\Delta C + \sum_{i=1}^{N_2} \alpha_i (A\Delta \Gamma_i + BF_i) q^{-1} \\ &= C(A\Delta + \sum_{i=1}^{N_2} \alpha_i (B - A\Delta G_i) q^{i-1}) \\ &= CA_c. \end{aligned} \quad (2.32)$$

Hence, equation (2.31) becomes:

$$y_t = \frac{BT_1}{A_c} r_{t+N_2-1} + \frac{R}{A_c} \xi_t \quad (2.33)$$

where we see that the C polynomial cancels in the closed loop transfer functions (as is usually the case with observer dynamics) and that the stability and performance of the closed loop system are governed by the roots of the A_c polynomial. How the roots of this polynomial are affected by different choices of the design parameters N_1 , N_2 , N_u and λ is hard to tell from the definition of A_c . The case of $N_u = 1$ has been studied in [WZ87], giving rise to a root-locus analysis. In practice, approaches have been made to optimize on-line the value of the design parameter λ in order to get reasonable values for the roots of A_c but this is not always possible. Notice also that $A_c(1) = \sum_{i=1}^{N_2} \alpha_i B(1)$ and $(BT_1)(1) = \sum_{i=1}^{N_2} \alpha_i B(1)$ so that the static gain of the closed loop transfer function from reference to output is always 1.

Coming back to the rôle of the C polynomial as an observer polynomial, we point out that if a simpler model is chosen for the plant, with $C = 1$, and if an observer polynomial T_0 , chosen by the user, is used instead of

C in the Diophantine (Aryabhata/Van Compernelle) equations (2.8) and (2.10), these observer dynamics will also disappear in the closed loop transfer function from r_t to y_t but not in the noise to output transfer function. A much more thorough discussion of the rôle of the observer polynomial T_0 and its use as a design choice can be found in [Moh88]. In the last section of this chapter, we shall show how some modifications of this GPC control strategy can result in prescribed closed loop dynamics between reference and output. The comfort to be drawn from this section and the preceding example is that the GPC control law produces (in the nonadaptive case) linear time-invariant controllers indistinguishable from those perhaps designed by other methods.

2.6 GPC Based on a ‘Performance Model’

A variant of the GPC algorithm has been proposed by Irving, Falinower and Fonte and later studied by several authors [IFF86], [GWZ87], [MSD87]. It has been coined by its authors ‘GPC based on a performance model’, but we claim no responsibility for the relevance of that name. We shall see, in fact, that it implements a pole placement design objective rather than a performance optimization objective. The reasons for presenting it here are manifold.

- It is representative, or symptomatic, of the variations made to control design formulations within the GPC framework under the banner of performance enhancement and/or stability promises. It is included for the sake both of contemporaneity and of completeness, but absolutely no imposition is made upon the reader necessarily to absorb fully this material, since it can be subsumed by our later adaptive LQG analysis.
- It implements a pole placement technique, namely a procedure with precise and separate specification of the closed loop tracking and regulation dynamics.
- It possesses the (potential at best, apocryphal at worst) robustness associated with predictive control or linear quadratic control strategies.
- It appears to have had some success in practice, and to have admitted assertedly stable designs via a deadbeat strategy, but without the apparent inherent problems. These claims are supported by empirical and/or anecdotal evidence, and so deserve our attention.
- It is not supported by an existent robustness theory, although that presented in this book may help to justify it.

- It provides a motivational example where the theory needs to come forward to embrace practically engendered methodology.
- Its formulation in terms of reference models and noise models, applied on top of (effectively) singular control techniques, demonstrates close similarities to the available methods for the pursuit of LQG control strategies that will be further pursued in this book.

To make the variation of GPC more precise, the on-line solution of the Van Compernelle equation

$$AR + q^{-1}BS = A_m,$$

which appears in pole placement methods, often leads to numerical problems due to near common factors in the A and B polynomials. These may happen in the transient phase of the identification algorithm or may be produced by overparametrization. Hence, this equation should be avoided in an adaptive context, and is replaced in the following method by a ‘performance model’.

The stability of GPC controllers is hard to ensure, unless deadbeat control settings are chosen for the design parameters of the criterion (2.7). However, deadbeat control often leads to wild control inputs unless, as is the case here, this deadbeat control strategy is applied on a filtered performance model instead of the process model itself.

We start from the process model (2.6) and define directly a tracking reference model for r_t by the following equation,

$$r_t = \frac{BT}{A_m}n_t, \quad (2.34)$$

where n_t is the input of the reference model, and where A_m is a monic polynomial chosen by the user in order to have desired tracking dynamics: A_m contains the desired closed loop poles. T is a scalar constant which serves to ensure unit static gain. We again drop the dependency on q^{-1} for notational convenience. Note that we have chosen to keep all the process zeros in the closed loop transfer function. Well damped zeros could of course be canceled if desired (see [ÅW84]).

We then denote by u_t^r the control input that would produce an output y_t identical to the output r_t of the reference model if there were no noise. Since the open loop plant is $q^{-1}B/A$, u_t^r is defined as

$$u_t^r = \frac{AT}{A_m}n_{t+1}. \quad (2.35)$$

Note that (2.34) is the desired closed loop reference (n_t) to output (y_t) transfer function, which is classical in conventional pole placement techniques, but

that (2.35) is the corresponding closed loop reference to input transfer function which is usually not explicitly used but which plays a central rôle in the subsequent development. We now define deviations from the reference values by

$$e_t^y = y_t - r_t \quad (2.36)$$

$$e_t^u = \Delta(u_t - u_t^r), \quad (2.37)$$

and filtered versions of these deviations,

$$e_t^{yf} = A_f e_t^y \quad (2.38)$$

$$e_t^{uf} = A_f e_t^u, \quad (2.39)$$

where A_f is a stable monic polynomial which will influence the regulation closed loop transfer function, as will be shown within a few minutes to the fast reader.

The classical model reference method, which in our case simplifies to a pole placement method since the reference model and the plant model have the same zeros, then amounts to finding the linear controller which sets e_t^y to zero (and simultaneously also sets e_t^u to zero), but this involves the solution of the previously mentioned Van Compernelle equation. Instead of forcing e_t^y and e_t^u identically to zero to achieve exact model matching, the way which will be followed here is to minimize the following criterion,

$$J(u, t) = E\left\{ \sum_{j=N_1}^{N_2} [e_{t+j}^{yf}]^2 + \lambda \sum_{j=1}^{N_u} [e_{t+j-1}^{uf}]^2 \right\} \quad (2.40)$$

$$\text{subject to } e_{t+j}^{uf} = 0, \quad j = N_u, \dots, N_2.$$

Thus we have transposed a pole positioning aim into a GPC framework. When using (2.40) one will still apply a receding horizon strategy: N_u future control actions are defined by the optimization of the criterion but only the first one is actually applied at time t . In order to minimize (2.40) we need a model which relates e_t^{uf} to e_t^{yf} and which is called the ‘performance model’. Multiplying both sides of (2.6) by A_f , subtracting $A_f A \Delta r_t$ from both sides and taking into account (2.34), (2.35), (2.38), (2.39) yields

$$A \Delta e_t^{yf} = B e_{t-1}^{uf} + A_f C \xi_t. \quad (2.41)$$

Optimal predictions for e_{t+j}^{yf} can then be obtained using the following Diophantine equations

$$A_f(q^{-1})C(q^{-1}) = E_j(q^{-1})A(q^{-1})\Delta + q^{-1}F_j(q^{-1}), \quad (2.42)$$

$$E_j(q^{-1})B(q^{-1}) = G_j(q^{-1})A_f(q^{-1})C(q^{-1}) + q^{-j}\Gamma_j(q^{-1}). \quad (2.43)$$

In a way entirely similar to previous derivations, we obtain the following equation for the prediction

$$A_f C \hat{e}_{t+j}^{yf} = A_f C G_j e_{t+j-1}^{uf} + \Gamma_j e_{t-1}^{uf} + F_j e_t^{yf}. \quad (2.44)$$

We then define the following quantities, in complete analogy with (2.15), (2.17), (2.18) and (2.19)

$$\hat{e}_{t+j|t}^{yf} = \Gamma_j (A_f C)^{-1} e_{t-1}^{uf} + F_j (A_f C)^{-1} e_t^{yf} \quad (2.45)$$

$$\mathbf{f} = [\hat{e}_{t+1|t}^{yf}, \hat{e}_{t+2|t}^{yf}, \dots, \hat{e}_{t+N_2|t}^{yf}]^T \quad (2.46)$$

$$\mathbf{e}^u = [e_t^{uf}, e_{t+1}^{uf}, \dots, e_{t+N_u-1}^{uf}]^T, \quad (2.47)$$

$$\hat{\mathbf{e}}^y = [\hat{e}_{t+1}^{yf}, \hat{e}_{t+2}^{yf}, \dots, \hat{e}_{t+N_2}^{yf}]^T. \quad (2.48)$$

The predicted input–output relationship of the performance model can then be written as the vector equation,

$$\hat{\mathbf{e}}^y = \mathbf{G} \mathbf{e}^u + \mathbf{f},$$

where the matrix \mathbf{G} is the same as in (2.21). The control criterion becomes

$$J = \hat{\mathbf{e}}^{y^T} \hat{\mathbf{e}}^y + \lambda \mathbf{e}^{u^T} \mathbf{e}^u, \quad (2.49)$$

and its solution is given by

$$\mathbf{e}^u = -(\mathbf{G}^T \mathbf{G} + \lambda \mathbf{I})^{-1} \mathbf{G}^T \mathbf{f}. \quad (2.50)$$

Since again a receding horizon strategy is sought, only the first value of \mathbf{e}^u is applied and, using (2.27), we get

$$e_t^{uf} = - \sum_{j=1}^{N_2} \alpha_j \hat{e}_{t+j|t}^{yf}.$$

Taking into account (??), this leads to

$$\begin{aligned} (A_f C + \sum_{j=1}^{N_2} \alpha_j \Gamma_j q^{-1}) e_t^{uf} &= - \sum_{j=1}^{N_2} \alpha_j F_j e_t^{yf} \\ (A_f C + \sum_{j=1}^{N_2} \alpha_j \Gamma_j q^{-1}) \Delta(u_t - u_t^r) &= - \sum_{j=1}^{N_2} \alpha_j F_j (y_t - r_t). \end{aligned}$$




Figure 2.1: Controller structure for the ‘performance model’ GPC

If we define

$$R = (A_f C + \sum_{j=1}^{N_2} \alpha_j \Gamma_j q^{-1})$$

$$S = \sum_{j=1}^{N_2} \alpha_j F_j,$$

we then have a classical linear expression for the controller,

$$R\Delta u_t = -S(y_t - r_t) + R\Delta u_t^r. \quad (2.51)$$

This reveals the structure of the control system as shown in Figure 2.1. The control input consists of a feedback term, which is based on the error between the output of the process and that of the reference model, and a feedforward term which is computed by this reference model.

It is worth computing explicitly the closed loop transfer functions arising from this control strategy. We combine (2.34), (2.35), (2.6) and (2.51) to yield

$$(AR\Delta + q^{-1}BS)y_t = (AR\Delta + q^{-1}BS)r_t + CR\xi_t. \quad (2.52)$$

Then, if and only if $(AR\Delta + q^{-1}BS)$ has all its roots inside the unit circle, this is equivalent to

$$y_t = r_t + \frac{CR}{AR\Delta + q^{-1}BS}\xi_t$$

$$= \frac{BT}{A_m}n_t + \frac{CR}{AR\Delta + q^{-1}BS}\xi_t, \quad (2.53)$$

which means that the closed loop tracking transfer function approaches asymptotically the desired reference model, independently of the process parameters and the particular design variables of the criterion, while the regulation dynamics do depend on these parameters and design variables. These conclusions of course hold only if the model is capable of describing the true plant exactly. In such case, the optimal value of the criterion is just the variance of the noise contribution and is independent of the design variables λ , N_u and N_2 . In fact, the idea behind the criterion (2.40) is not to achieve a compromise between minimization of the tracking error and minimization of the control energy, but it is to minimize the error of the pole placement Diophantine equations. Hence, the GPC with a ‘performance model’ is not an honest predictive control method, but a devious (and clever) way of doing pole placement in disguise. It reformulates the solution of the sometimes poorly conditioned Diophantine equations as a better conditioned quadratic minimization problem, thereby counting on the superior numerical properties of the latter over the former. In the ideal case of perfect modeling, the two solutions are identical.

We further note that the GPC minimization (2.40) provides the feedback component of the design. The tracking property is achieved by the reference model which provides a feedforward contribution [Wit85]. Further, straightforward calculations show that

$$\begin{aligned} AR\Delta + q^{-1}BS &= A_f C [A\Delta + \sum_{j=1}^{N_2} \alpha_j (B - A\Delta G_j) q^{j-1}] \\ &= A_f C A_c. \end{aligned}$$

Hence one attempts to choose the design variables of the criterion in order that the dynamics associated with the roots of A_c become negligible compared to those of A_m and A_f , i.e. they are relatively fast. Certain choices of GPC design parameters permit the nominal placement of A_c poles close to zero, and these could be used here. We shall investigate such settings further in Section 4.6.4.

We have now shown an example of how the use of reference and noise models in GPC style formulations is promulgated, and further, how these methods are manipulated to allow a single design philosophy (GPC) to perform different tasks. Once again, we remark that the genesis of these methods in adaptive control practice is reason enough for their interest, but also strong motivation for their analysis in a broader framework.

2.7 Conclusions

Our aim in this chapter has been to present the Generalized Predictive Control algorithm and one of its recent variants which motivates our investigations into the robustness properties of predictive control algorithms in general, especially in an adaptive context. The main ideas which we have illustrated in this chapter are shared by most of the long range input-output predictive control methods and can be summarized as follows:

- These algorithms have been developed as extensions of and in the line of Minimum Variance control strategies. Hence, the weighting on control signals in the criterion, when applicable, is usually chosen small provided closed loop stability is accomplished.
- A main feature of these control methods is that they can handle many different control problems, on a wide range of plants, with a manageable number of design variables in the criterion. These variables have to be specified by the user depending upon their prior knowledge of the plant and control objectives and constraints.
- These control strategies have been used in several industrial applications and have shown good performance and a certain degree of robustness with respect to overparametrization or poorly known delays.
- These algorithms are nothing other than finite horizon quadratic optimization strategies which can be (and will be) recast in an LQG control framework in the next chapter.
- A theoretical analysis in order to assess the influence of the design variables (prediction and control horizons, weighting factors) on the stability of the closed loop control system can only be done for very specific choices of these design variables. The absence of stability, performance and robustness guarantees is the most serious failing of the strategy.

Having exposed the reader to these ideas of the establishment adaptive control practitioner, we shall next indulge in a little sedition to cobble these applications methods onto mainstream LQG control theory. As a matter of fact, we shall show in passing that GPC is a particular example of LQG controller for special choices of weighting matrices.

Chapter 3

Linear Quadratic Gaussian Optimal Control

3.1 Introduction

Linear Quadratic Gaussian¹ (LQG) Optimal Control laws are firmly established bastions of state-space-based linear control systems design. They are well understood, well tried and, some would say, old-fashioned and/or superseded by newer optimal control methods such as H_∞ . Nevertheless, it is just these features in their linear application which make LQG controllers of great practical significance, and their rôle is to underpin our development of new adaptive predictive controllers, extensible later (Chapter 8) to adaptive H_∞ optimal control. We present their requisite properties here.

In Chapter 2 we have provided some motivational support for the use of criterion-based controller design methods. GPC is one particularly simple optimal control design criterion, formulated in an input–output framework, which has met with popular acclaim. One of the often quoted advantages claimed for GPC is its particularly simple computational solution which, unlike state-space formulated optimal control problems, does not require the solution of a Riccati equation.

Whereas they are undoubtedly simple computationally, the GPC control formulae do not lend themselves easily to an analysis of the closed loop

¹Throughout this book, we shall use the terms Linear Quadratic Gaussian (LQG) to denote the interconnection of linear state estimation with linear state-estimate feedback. These feedback/observer combinations need not necessarily have arisen from LQ optimal control and Kalman filtering. This means, in particular, that the acronym LQG will be used without necessarily the requirement of a Gaussian noise assumption, or possibly even without any noise assumption.

stability and performance properties of GPC such as: the computation of closed loop poles, the amount of unmodeled dynamics that can be tolerated while still preserving closed loop stability, or the comparison between the designed performance and the achieved performance when receding horizon controllers are used.

One way to analyse the stability and performance properties of GPC is to embed it into the more powerful framework of Linear Quadratic Gaussian (LQG) optimal control, since GPC can be, and will be, recast as a special case of LQG. Therefore this chapter will review the central features of feedback controller design via Linear Quadratic (LQ) optimal control design plus Kalman Filter (KF) or Kalman Predictor (KP) state-estimate design. Chapter 4 will then be devoted to a thorough analysis of the stability and performance of LQ receding horizon controllers, with a particular, almost maternal, attention paid to GPC.

However, exposing the reader, you, to a wide range of basic features about LQ optimal control, Kalman predictors and filters, and LQG design for the sole purpose of studying the properties of GPC controllers would be like using a ferry boat to cross the Nile.² Despite our meritorious rescue efforts, GPC, at least in its present-day formulation, will have to be jettisoned on the grounds of the stability analysis of Chapter 4. The central theme of our book will thereafter turn to the application of the fuller design methods of LQG controllers and Least Squares (LS) parameter estimation within the context of adaptive control systems. The present chapter will therefore introduce the notations, lay down the mathematical foundations, and present the formulae necessary for the analysis to follow in Chapters 4, 5 and 6. Even though some readers might like to indulge in the derivation of the various expressions, if only to check whether the authors are foolproof formulae producers, some others may prefer only to have a cursory glance at this chapter and to use it mainly as a reference guide for the later chapters.

There is a wealth of literature available on the subject of Linear Quadratic Optimal Control and Estimation, see for example [AM71], [AM79], [AM90], [KS72], [Lew86a], [Lew86b], which explores the derivation of results and methods, as well as their application to time-varying systems, and their varied behaviors in the case of special (sometimes pathological) choices of designer specifications. Our needs here are less ambitious since, in adaptive control, one normally deals with controller designs based upon time-invariant models, frozen at each iteration at the present value of the parameter estimate, generally without pathological choices in cost criterion specification.

²As is well known, the Nile (called Nil in French) is a 1.5m wide creek in Belgium that flows close to the University of Louvain la Neuve, where much of this book was written.

LQG control is formulated in terms of state-space models of the plant system rather than simpler input–output models. Consequently, we suppose that we are provided with a state-space model of the linear plant, which we further assume to be strictly proper,

$$x_{t+1} = Fx_t + Gu_t + w_t \quad (3.1)$$

$$y_t = Hx_t + v_t. \quad (3.2)$$

Here x_t is the n -dimensional state vector of the plant, u_t the m -dimensional control input vector, y_t the p -dimensional measured plant output vector, and w_t and v_t are zero mean white noise processes of appropriate dimensions representing state disturbance and measurement noise respectively.

We shall begin by considering the LQ optimal control problem, which focuses upon the behavior of the state x_t in response to the control signal u_t . This problem is thus associated with the equation (3.1) alone and does not involve the output measurement equation (3.2) directly. Although we shall see later how specific choices of design matrices can be made to produce an LQ solution which reflects input–output properties and so is invariant to the particular state coordinate basis chosen. Firstly the LQ regulator problem will be treated, where the aim is to select the control sequence so as to maintain x_t suitably close to zero. This will then be extended to the tracking problem more closely related to GPC, where x_t is controlled in a fashion to force the tracking by y_t of a reference trajectory, r_t . Whether the objective is to regulate the state or the output of a system or to track a reference trajectory, LQ optimal control is a versatile control philosophy that admits a variety of cost criteria leading to different control laws, time-varying or time-invariant, with differing computational loads. GPC is a receding horizon strategy, as we have seen in Chapter 2. The formal derivation of a receding horizon LQ control law necessitates the solution of a finite horizon LQ control problem, while the asymptotic stability and performance of the resulting closed loop are most naturally studied in the framework of infinite horizon LQ theory. Therefore, we shall present the finite horizon, infinite horizon and receding horizon problem formulations. Chapter 4 will focus on the stability and performance characteristics of the associated closed loops.

Since the stability properties are dependent only upon the regulator problem features, we begin with the LQ regulator problem, and subsequently show that the tracking problem can be recast into a regulator framework by state augmentation.

3.2 The Linear Quadratic Regulator

3.2.1 The Finite Horizon Regulator

The quadratic cost criterion which we seek to minimize with respect to the sequence u_t, \dots, u_{t+N-1} in the LQ regulator problem is given by

$$J(N, x_t) = \mathbb{E} \left\{ x_{t+N}^T P_0 x_{t+N} + \sum_{j=0}^{N-1} \left\{ x_{t+j}^T Q_{c, N-j-1} x_{t+j} + u_{t+j}^T R_{c, N-j-1} u_{t+j} \right\} \right\}, \quad (3.3)$$

where \mathbb{E} denotes expectation and P_0 , $Q_{c,j}$ and $R_{c,j}$ are non-negative definite symmetric matrices, and where perfect state (x_t) measurement has been assumed. This is a finite horizon (N) optimal control problem which we specify for the state-space model (3.1) above. We note that we have used expectations with respect to the stochastic uncertainties here to generate our cost function. We could as easily have treated the deterministic LQ problem ($w_t = 0$) with deterministic cost. The solution is invariant to this feature.

The solution of the LQ optimal regulator problem may be given directly in closed loop form as follows. One iterates the Riccati Difference Equation (RDE),

$$P_{j+1} = F^T P_j F - F^T P_j G (G^T P_j G + R_{c,j})^{-1} G^T P_j F + Q_{c,j}, \quad (3.4)$$

from the initial condition P_0 and implements the feedback control sequence given by

$$\begin{aligned} u_{t+N-j} &= -(G^T P_{j-1} G + R_{c,j-1})^{-1} G^T P_{j-1} F x_{t+N-j}, \\ &= K_{j-1} x_{t+N-j}, \quad j = 1, \dots, N, \end{aligned} \quad (3.5)$$

where F , G and $R_{c,j}$ arise from the control problem, P_j is the matrix solution of the RDE, and x_t is the state of the plant. The evaluated optimal cost, $J(N, x_t)^*$, is given as follows by P_N , which itself does not explicitly appear in the control law,

$$J(N, x_t)^* = x_t^T P_N x_t. \quad (3.6)$$

We note that the direction of iteration of the RDE (3.4) is reverse-time compared with the direction of evolution of the plant (3.1). This can be seen directly by observing the index of the gain in (3.5). In the next chapter we shall focus more closely on the dynamics of the RDE and so it suits us better to consider its evolution also as being forward in time. Hence our need for complicated indexing in (3.3) and (3.5).

3.2.2 The Infinite Horizon Regulator

In the general formulation of LQ optimal control problems, there is no need to presume the underlying plant time-invariant — the plant matrices F , G , H in (3.1) may change as functions of time just as the weighting matrices $Q_{c,j}$ and $R_{c,j}$ may in the criterion (3.3). If, however, we consider a problem with both time-invariant plant and with constant weighting matrices Q_c and R_c , and if we then allow N to approach infinity, then (3.3) (divided by N) yields an infinite horizon optimal control problem which is also well posed:

$$J(x_t) = \lim_{N \rightarrow \infty} \frac{1}{N} J(N, x_t). \quad (3.7)$$

More will be said later concerning interconnections between the LQ problems over these finite and infinite horizons.

If a stationary infinite horizon LQ problem is formulated as in (3.7) then the solution above alters in that, subject to P_0 being non-negative definite, R_c being positive definite, $[F, G]$ being stabilizable, and $[F, Q_c^{1/2}]$ being detectable, P_j converges to a constant matrix P_∞ , as j goes to infinity, which is the maximal solution of the Algebraic Riccati Equation (ARE),

$$P_\infty = F^T P_\infty F - F^T P_\infty G (G^T P_\infty G + R_c)^{-1} G^T P_\infty F + Q_c. \quad (3.8)$$

A stationary control law is therefore defined,

$$\begin{aligned} u_{t+j} &= -(G^T P_\infty G + R_c)^{-1} G^T P_\infty F x_{t+j} \\ &= K x_{t+j}. \end{aligned} \quad (3.9)$$

Note that the differing directions of evolution of the state equation and the RDE effectively move the RDE initial condition P_0 into the infinitely remote future for this problem. Thus the control gain in (3.9) is time-invariant.

3.2.3 The Receding Horizon Regulator

A direct examination of the finite horizon discrete-time LQ regulator problem discloses several simple computational features. Firstly, since the cost function $J(N, x_t)$ revolves solely around the selection of N control values, the optimal control sequence may, in principle, be found by finite-dimensional optimization, as in GPC. Secondly, the RDE (3.4) may be explicitly iterated from P_0 to P_N using simple linear algebra. By contrast, the infinite horizon problem at first sight involves an infinite-dimensional optimization or the solution of an algebraic matrix equation, the ARE (3.8), whose solution typically requires eigenvalue methods for symplectic matrices. Balancing these

computational niceties of the finite horizon regulator are the facts that its solution for the control signal is not obviously extensible to the infinite interval, it is not time-invariant state feedback over the interval, and if derived via finite-dimensional optimization will typically be specified as open loop control values.

One method proposed to admit the calculation simplicity of finite horizon LQ methods while addressing an infinite horizon implementation and preserving the time-invariance of the infinite horizon feedback is the receding horizon LQ regulator. In this formulation one chooses u_t as the first element of the finite horizon solution minimizing $J(N, x_t)$, u_{t+1} as the first element of the solution minimizing $J(N, x_{t+1})$, and so on. Comparison with the description of the finite horizon regulator above demonstrates immediately that, for this receding horizon strategy, one has

$$\begin{aligned} u_t &= u_t \left(\arg \min_{u_t, \dots, u_{t+N-1}} J(N, x_t) \right) \\ &= -(G^T P_{N-1} G + R_{c, N-1})^{-1} G^T P_{N-1} F x_t \\ &= K_{N-1} x_t. \end{aligned} \quad (3.10)$$

We shall re-address receding horizon methods much more fully in the next chapter, since they underlie much of the thinking behind GPC. Our next immediate task is, however, to extend this LQ regulator analysis to address the Linear Quadratic Tracking Problem similar to that of the GPC criterion. To the surprise of at least one of the authors, one may achieve this simply by extending the plant model (3.1) to include an elementary description of the evolution of the reference signal and posing an LQ regulator problem for the combined system.

3.3 The Linear Quadratic Tracking Problem

We alter the finite horizon criterion (3.3) to include deviations of the output from a reference signal r_t as follows:

$$\begin{aligned} J(N, x_t) &= E \{ (y_{t+N} - r_{t+N})^T P_0 (y_{t+N} - r_{t+N}) \\ &\quad + \sum_{j=0}^{N-1} \{ (y_{t+j} - r_{t+j})^T Q_{c, N-j-1} (y_{t+j} - r_{t+j}) + u_{t+j}^T R_{c, N-j-1} u_{t+j} \} \}. \end{aligned} \quad (3.11)$$

Note that P_0 and $Q_{c, j}$ are not the same as those of (3.3): in particular, they have different dimensions. Recognizing that r_t is an externally prescribed

signal, we may incorporate an artificial, dimension N , state model for it,

$$x_{t+1}^r = F^r x_t^r + G^r n_t \quad (3.12)$$

$$r_t = H^r x_t^r \quad (3.13)$$

with n_t zero mean white noise of dimension p , and

$$F^r = \begin{pmatrix} 0 & I & 0 & \dots & 0 \\ 0 & 0 & I & \dots & 0 \\ \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & 0 & \dots & I \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix} \quad (3.14)$$

$$G^r = (0 \ 0 \ 0 \ \dots \ I)^T \quad (3.15)$$

$$H^r = (I \ 0 \ 0 \ \dots \ 0). \quad (3.16)$$

Here the 0 and I in the matrices F^r , G^r , H^r indicate matrices of dimension $p \times p$. Thus the sequence of elements $\{r_{t+i}, i = 0, \dots, N-1\}$ is ‘stacked’ into the initial condition, x_t^r , and the further distant future elements enter the stack via n_t , which is unpredictable at time t . Now we may recast the LQ tracking criterion (3.11) as a regulation criterion (3.3) with state

$$x_t^{track} = \begin{pmatrix} x_t \\ x_t^r \end{pmatrix}.$$

The state equation for x_t^{track} is then the direct sum of the state equations for x_t and x_t^r ,

$$x_{t+1}^{track} = \begin{pmatrix} F & 0 \\ 0 & F^r \end{pmatrix} x_t^{track} + \begin{pmatrix} G \\ 0 \end{pmatrix} u_t + \begin{pmatrix} I & 0 \\ 0 & G^r \end{pmatrix} \begin{pmatrix} w_t \\ n_t \end{pmatrix}, \quad (3.17)$$

while $y_t - r_t$ is given by

$$y_t - r_t = (H \quad -H^r) x_t^{track}.$$

Note that this state model is not controllable from u_t but that it is stabilizable since the x_t^r subsystem is deadbeat and noninteracting by construction. With the P_0 , $Q_{c,j}$ and $R_{c,j}$ matrices replaced by

$$P_0^{track} = \begin{pmatrix} H^T \\ -H^{rT} \end{pmatrix} P_0 (H \quad -H^r) \quad (3.18)$$

$$Q_{c,j}^{track} = \begin{pmatrix} H^T \\ -H^{rT} \end{pmatrix} Q_{c,j} (H \quad -H^r) \quad (3.19)$$

$$R_{c,j}^{track} = R_{c,j}, \quad (3.20)$$

in (3.3), the tracking criterion (3.11) can now be rewritten as

$$\begin{aligned}
 J(N, x_t^{track}) &= \mathbb{E} \left\{ x_{t+N}^{track T} P_0^{track} x_{t+N}^{track} \right. & (3.21) \\
 &+ \left. \sum_{j=0}^{N-1} \left[x_{t+j}^{track T} Q_{c, N-j-1}^{track} x_{t+j}^{track} + u_{t+j}^T R_{c, N-j-1}^{track} u_{t+j} \right] \right\},
 \end{aligned}$$

plus additional terms not influenced by the controls. The tracking criterion has thus been recast as an equivalent regulation criterion. The tracking control law then is derived by solving this regulation problem to provide a feedback strategy involving linear feedback of the full state x_t^{track} .

The LQ tracking problem has a solution which incorporates into the feedback law the future knowledge of the reference signal for N steps. In this fashion it is manifest how one may transform, via such a formulation, the finite horizon LQ tracking problem into the finite horizon LQ regulation of the composite system. The introduction of a tracking objective into an N -step LQ problem necessitates the augmentation of the state by N elements, thereby increasing the computational burden of solution. An infinite horizon tracking problem is not well posed in these terms because of the presumption that r_t be prescribed into the infinitely remote future. In these circumstances, it is still possible, however, to develop meaningful tracking problems by adding finite-dimensional x_t^r to the augmented state model and then solving an infinite horizon regulation problem. The dimension of the x^r component of the state vector then corresponds to the number of advanced data available for the reference signal. Future reference values outside this scope are modeled by the unpredictable white noise process n_t and so are automatically replaced by zeros. Naturally, receding horizon LQ tracking problems are easily posed and are, in fact, closest in spirit to GPC. This feature will be reinforced by a computational example shortly.

For tracking of constant reference signals, one may replace the tapped-delay-line model structure (3.12), (3.13) by the following one-dimensional state model

$$x_{t+1}^r = x_t^r = r.$$

This corresponds to choosing $F^r = I$, $H^r = I$ and $G^r = 0$.

It is worth making several further remarks concerning the differences between the two LQ problems considered in these last two subsections. The additional states x_t^r augmenting the LQ regulator problem above are not reachable from the input process u_t , but they are stable by construction. If one partitions the solution of the augmented RDE conformably with the

state equation (3.17),

$$P = \begin{pmatrix} P^{11} & P^{12} \\ P^{12T} & P^{22} \end{pmatrix}, \quad (3.22)$$

then one observes that the control gain matrix $K = (K^x \ K^r)$ depends only upon P^{11} and P^{12} . Specifically,

$$\begin{aligned} K_j &= -(G^T P_j^{11} G + R_{c,j})^{-1} (G^T P_j^{11} F \quad G^T P_j^{12} F^r) \\ &= (K_j^x \ K_j^r), \end{aligned} \quad (3.23)$$

and

$$\begin{aligned} u_{t+N-j} &= K_{j-1} x_{t+N-j}^{track} \\ &= K_{j-1}^x x_{t+N-j} + K_{j-1}^r x_{t+N-j}^r. \end{aligned}$$

Similarly, the RDE decomposes, with P^{11} satisfying the regulation RDE,

$$P_{j+1}^{11} = F^T P_j^{11} F - F^T P_j^{11} G (G^T P_j^{11} G + R_{c,j})^{-1} G^T P_j^{11} F + H^T Q_{c,j} H, \quad (3.24)$$

with initial condition $P_0^{11} = H^T P_0 H$, and P^{12} satisfying a coupled Lyapunov equation,

$$\begin{aligned} P_{j+1}^{12} &= (F - G(G^T P_j^{11} G + R_{c,j})^{-1} G^T P_j^{11} F)^T P_j^{12} F^r - H^T Q_{c,j} H^r \\ &= (F - G K_j^x)^T P_j^{12} F^r - H^T Q_{c,j} H^r, \end{aligned} \quad (3.25)$$

with initial condition $P_0^{12} = -H^T P_0 H^r$.

In the case of constant weighting matrices, Q_c and R_c , the sequences P_j^{11} and P_j^{12} will converge to steady-state values P_∞^{11} and P_∞^{12} provided P_0 is non-negative definite, R_c is positive definite, $[F, G]$ is stabilizable and $[F, Q_c^{1/2} H]$ is detectable. (Indeed, it follows from the stability of F^r that these last two conditions imply the stabilizability and detectability respectively of the corresponding matrix pairs for the augmented system (3.17).) The time-varying finite horizon tracking controller (3.23)–(3.25) can then be replaced by the corresponding time-invariant controller, through the substitution of P_∞^{11} and P_∞^{12} for P_j^{11} and P_j^{12} in (3.23). The dimension of x_t^r in this steady-state tracking problem is then chosen as discussed above. It is apparent from the relations (3.24) and (3.25) and the nonreachability of the state x_t^r that the stability properties of this system are dictated solely by the properties of the solution of the regulator problem. This will be illustrated by an example later in the chapter.

We next move on to study the derivation of optimal state estimators or observers.

3.4 The Linear Optimal State Estimator

To implement the feedback control laws of LQ optimal control requires that the state x_t be available at time t for construction of the control signal. (The reference signal r_t , and hence x_t^r , is always available for measurement by assumption, as is the control signal u_t .) As the state is not always perfectly measurable, an alternative control law is used,

$$\begin{aligned} u_{t+N-j} &= -(G^T P_{j-1} G + R_c)^{-1} G^T P_{j-1} F \hat{x}_{t+N-j} \\ &= K_{j-1} \hat{x}_{t+N-j}, \end{aligned} \quad (3.26)$$

or

$$\begin{aligned} u_{t+N-j} &= -(G^T P_\infty G + R_c)^{-1} G^T P_\infty F \hat{x}_{t+N-j} \\ &= K \hat{x}_{t+N-j}, \end{aligned} \quad (3.27)$$

(or their tracking equivalents) where \hat{x}_t is a state estimate produced by an observer. We shall now consider the features of Least Squares optimal state estimation for (3.1) where the available signals, $\{u_t\}$ and $\{y_t\}$, are used to generate \hat{x}_t .

3.4.1 The Kalman Predictor (KP)

The theory of linear least squares optimal observer design is dual to that of LQ controller design [KS72]. Thus the methodology of LQ control may be immediately transposed to observer construction. The starting point for the design of a state estimator is, naturally, a model for the evolution of the state (3.1),

$$\begin{aligned} x_{t+1} &= Fx_t + Gu_t + w_t \\ y_t &= Hx_t + v_t, \end{aligned}$$

where w_t represents the state disturbance and v_t the output noise process. The usual assumptions concerning the noise processes in this model are that w_t and v_t are mutually independent, zero mean white noises with covariance,

$$\mathbb{E} \left\{ \begin{pmatrix} w_t \\ v_t \end{pmatrix} \begin{pmatrix} w_t^T & v_t^T \end{pmatrix} \right\} = \begin{pmatrix} Q_o & 0 \\ 0 & R_o \end{pmatrix}, \quad (3.28)$$

and that the initial state, x_0 , is independent of $\{w_t, t \geq 0\}$ and $\{v_t, t \geq 0\}$ and is also zero mean with covariance Σ_0 . Also the control signal, u_t , is perfectly known. If we denote by $\hat{x}_{t|t-1}$ the linear estimate of x_t given $\{u_{t-1}, u_{t-2}, \dots, u_0\}$ and $\{y_{t-1}, y_{t-2}, \dots, y_0\}$ and an optimal estimator

is sought to minimize the criterion $E\{(x_t - \hat{x}_{t|t-1})^T(x_t - \hat{x}_{t|t-1})\}$, then $\hat{x}_{t|t-1}$ is given by the Kalman predictor,

$$\hat{x}_{t+1|t} = (F - M_t^P H)\hat{x}_{t|t-1} + Gu_t + M_t^P y_t, \quad (3.29)$$

where

$$M_t^P = F\Sigma_t H^T (H\Sigma_t H^T + R_o)^{-1}, \quad (3.30)$$

and Σ_t satisfies the optimal filtering Riccati Difference Equation,

$$\Sigma_{t+1} = F\Sigma_t F^T - F\Sigma_t H^T (H\Sigma_t H^T + R_o)^{-1} H\Sigma_t F^T + Q_o. \quad (3.31)$$

Further, as $t \rightarrow \infty$, and provided Σ_0 is non-negative definite, R_o is positive definite, $[H, F]$ is detectable and $[F, Q^{1/2}]$ is stabilizable, then Σ_t tends to the maximal solution, Σ_∞ , of the filtering Algebraic Riccati Equation,

$$\Sigma_\infty = F\Sigma_\infty F^T - F\Sigma_\infty H^T (H\Sigma_\infty H^T + R_o)^{-1} H\Sigma_\infty F^T + Q_o, \quad (3.32)$$

and M_t^P tends towards the constant M^P ,

$$M^P = F\Sigma_\infty H^T (H\Sigma_\infty H^T + R_o)^{-1}. \quad (3.33)$$

3.4.2 The Kalman Filter (KF)

The Kalman (one-step-ahead) predictor produces the best linear estimate, in a least squares sense, of x_t given input signal data $\{u_{t-1}, u_{t-2}, \dots, u_0\}$ and output data $\{y_{t-1}, y_{t-2}, \dots, y_0\}$. To pass from the Kalman predictor to the Kalman filter, one seeks the best linear estimate, $\hat{x}_{t|t}$, of x_t given data $\{u_{t-1}, u_{t-2}, \dots, u_0\}$ and $\{y_t, y_{t-1}, \dots, y_0\}$. This filtered estimate is directly computed from the predicted estimate plus the new data, via

$$\begin{aligned} \hat{x}_{t|t} &= \hat{x}_{t|t-1} + M_t^F (y_t - H\hat{x}_{t|t-1}) \\ &= (I - M_t^F H)\hat{x}_{t|t-1} + M_t^F y_t, \end{aligned} \quad (3.34)$$

with

$$M_t^F = \Sigma_t H^T (H\Sigma_t H^T + R_o)^{-1}. \quad (3.35)$$

This, coupled with the time update equation, equivalent to (3.29),

$$\hat{x}_{t+1|t} = F\hat{x}_{t|t} + Gu_t, \quad (3.36)$$

then yields the recursive formula for the Kalman filter

$$\hat{x}_{t+1|t+1} = (I - M_{t+1}^F H)F\hat{x}_{t|t} + (I - M_{t+1}^F H)Gu_t + M_{t+1}^F y_{t+1}. \quad (3.37)$$

We note that the Kalman Filter and Kalman Predictor gains are related as follows:

$$M_t^P = FM_t^F, \quad M_\infty^P = FM_\infty^F.$$

The equations (3.29) and (3.37) are the recursive formulae for the generation of least squares optimal linear state estimates with single observer delay and with zero observer delay, i.e. with direct feedthrough. The gain matrices, M_t^P and M_t^F , of this design method arise from the solution of a Riccati difference equation (RDE) (3.31) running forwards in time from initial condition Σ_0 . Completely analogously to the LQ control problem, if we allow the initial condition to become remote then the observer gain tends to a constant value. One may equally well have begun with a fixed observer gain, M^P or M^F , in mind and specified the state estimators

$$\hat{x}_{t+1|t} = (F - M^P H)\hat{x}_{t|t-1} + Gu_t + M^P y_t, \quad (3.38)$$

or

$$\hat{x}_{t+1|t+1} = (I - M^F H)F\hat{x}_{t|t} + (I - M^F H)Gu_t + M^F y_{t+1}. \quad (3.39)$$

The stability of these observers depends upon the selection of M and the detectability of $[F, H]$.

3.5 Optimal Filter Design with Disturbance Models

The use of a white noise assumption in (3.1) for the modeled state noise w_t is a useful design choice, since its rôle is to capture the variability of the state in spite of the control action. However, it is frequently the case in industrial control problems that the other noise assumptions are very much too idealized, with output disturbances consisting of steps, ramps and sinusoids as well as white or colored noise. In such circumstances, it is possible to incorporate the knowledge of the unwanted output disturbance into the explicit state model above and thereby to have the Kalman filter design automatically attempt to eliminate their influence. In this fashion the Internal Model Principle [FW76] is asserted in the controller design, in a manner entirely similar to the GPC technique of forcing an integrator into the controller by modeling the measurement noise as passing through an integrating linear system.

The control system designer could possess specific knowledge of the nominal spectrum, $\Phi_{vv}(z)$, of the output disturbance process, v_t . If $\Phi_{vv}(z)$ can be

approximated reasonably well by a finite degree rationale spectrum $\hat{\Phi}_{vv}(z)$ then a stable state-space model could be derived from this by factorizing $\hat{\Phi}_{vv}(z)$ as

$$\hat{\Phi}_{vv}(z) = [I + H^d(zI - F^d)^{-1}G^d]R_o^m[I + H^d(z^{-1}I - F^d)^{-1}G^d]^T.$$

Such a factorization is always possible and can, in fact, be performed using the ARE [AM90]. This procedure yields a noise model similar to the state model (3.1) above except for the appearance of the same process, q_t , as state and as output noise. Where the spectrum of v_t is strictly non-zero on the unit circle, it is also possible to generate a noise model identical in form to (3.1) and (3.2) with uncorrelated driving noises, p_t and q_t ,

$$x_{t+1}^d = F^d x_t^d + G^d p_t \quad (3.40)$$

$$v_t = H^d x_t^d + q_t, \quad (3.41)$$

with x_t^d a k -dimensional state vector. Such a noise model pertains provided the measurement noise, v_t , has a sufficiently large unpredictable part modeled by q_t above. Equally, the noise model could be derived directly from physical data as, for example, could be done by modeling a step disturbance as being due to the unknown initial state of an integrator or by modeling sinusoidal disturbances as being due to the presence of a marginally stable oscillatory system generating part of v_t . These latter processes formally do not possess spectra but do admit treatment via an explicit noise model as above.

We may now use the dual of the technique of Section 3.3 used to generate a tracking LQ controller from a regulator and write a composite state equation by combining (3.1)–(3.2) and (3.40)–(3.41) with new state

$$x_t^m = \begin{pmatrix} x_t \\ x_t^d \end{pmatrix}. \quad (3.42)$$

The total state evolution and output generation equations are now

$$x_{t+1}^m = \begin{pmatrix} F & 0 \\ 0 & F^d \end{pmatrix} x_t^m + \begin{pmatrix} G \\ 0 \end{pmatrix} u_t + \begin{pmatrix} I & 0 \\ 0 & G^d \end{pmatrix} \begin{pmatrix} w_t \\ p_t \end{pmatrix}, \quad (3.43)$$

$$y_t = (H \quad H^d) x_t^m + q_t, \quad (3.44)$$

which again have the form of the standard state estimation problem

$$x_{t+1}^m = F^m x_t^m + G^m u_t + L^m w_t^m \quad (3.45)$$

$$y_t = H^m x_t^m + v_t^m, \quad (3.46)$$

with the obvious assignments of symbols,

$$F^m = \begin{pmatrix} F & 0 \\ 0 & F^d \end{pmatrix} \quad (3.47)$$

$$G^m = \begin{pmatrix} G \\ 0 \end{pmatrix} \quad (3.48)$$

$$H^m = (H \quad H^d) \quad (3.49)$$

$$L^m = \begin{pmatrix} I & 0 \\ 0 & G^d \end{pmatrix} \quad (3.50)$$

$$w_t^m = \begin{pmatrix} w_t \\ p_t \end{pmatrix} \quad (3.51)$$

$$v_t^m = q_t \quad (3.52)$$

$$\begin{aligned} Q_o^m &= L^m \mathbf{E}(w_t^m w_t^{mT}) L^{mT} \\ &= \begin{pmatrix} \mathbf{E}(w_t w_t^T) & 0 \\ 0 & G^d \mathbf{E}(p_t p_t^T) G^{dT} \end{pmatrix} \end{aligned} \quad (3.53)$$

$$R_o^m = \mathbf{E}(q_t q_t^T). \quad (3.54)$$

The stationary optimal state estimator with direct feedthrough from output to control signal, i.e. Kalman filter, for x_t^m is then obtained [AM79] as

$$\begin{aligned} \hat{x}_{t+1}^m &= (I - M^F H^m) F^m \hat{x}_t^m + (I - M^F H^m) G^m u_t \\ &\quad + M^F y_{t+1}, \end{aligned} \quad (3.55)$$

with M^F given by the solution, Σ^m , of a filtering ARE like (3.32),

$$M^F = \Sigma^m H^{mT} (H^m \Sigma^m H^{mT} + R_o^m)^{-1}. \quad (3.56)$$

Otherwise, a stationary optimal state estimator without direct feedthrough, i.e. Kalman predictor, for x_t^m is obtained by replacing the matrix triple (H, F, G) by (H^m, F^m, G^m) and M^F by $F M^F$ in (3.38).

The presentation of the observer construction above including the noise model is quite standard in the Kalman filtering literature [AM79] and, closely similar to the LQ control theory, much of it is directly applicable in an adaptive context. However, there still exist many variables requiring setting or selection before a design could be commenced and implemented. It will be our objective later to specify how generic choices of some of the design parameters may be made in order that the observer depend upon a manageable number of parameters and that it be relatively easily calculable. Also, because our ambit is the study of adaptive LQG control, and adaptive systems are *ipso facto* input-output oriented, we shall also present methods later which are state coordinate system independent.

3.6 LQG Controllers

Linear Quadratic Gaussian controllers are generated through the interconnection of a linear state-variable feedback control law and a linear state estimator. Indeed, in the case where the plant is linear and exactly given by the state model (3.1) and where the noises w_t and v_t are zero mean gaussian processes as is the initial state x_0 , this interconnection provides the optimal dynamic output feedback control of the plant, where optimality is measured according to the LQ criterion. This is the so-called ‘Separation Principle’ of linear optimal control [KS82]. In this section we shall investigate further the underlying structure of this combination, especially when reference and disturbance models are included.

We begin by considering the inclusion of the reference and measurement noise models to yield a composite system for LQ control. We then move on to treat the associated state estimation issues. In both these cases, the reachability and observability properties of the augmented models will come into play to admit simplified solutions, where the ARE solutions may be partitioned. Following this analysis, we shall then investigate expressions for the overall controller transfer functions in Section 3.8.

3.6.1 The Composite System and the LQ Objective

If we start with the system state description (3.1)–(3.2) and take advantage of the measurement noise model (3.40)–(3.41) and the reference trajectory model (3.12)–(3.13) associated with the specification of the tracking control objective (3.11), then we may compose a super state, \bar{x}_t , description of the combined process as follows:

$$\bar{x}_t = \begin{pmatrix} x_t \\ x_t^r \\ x_t^d \end{pmatrix}. \quad (3.57)$$

This super state satisfies the equation

$$\begin{aligned} \bar{x}_{t+1} &= \begin{pmatrix} x_{t+1} \\ x_{t+1}^r \\ x_{t+1}^d \end{pmatrix} \\ &= \begin{pmatrix} F & 0 & 0 \\ 0 & F^r & 0 \\ 0 & 0 & F^d \end{pmatrix} \begin{pmatrix} x_t \\ x_t^r \\ x_t^d \end{pmatrix} + \begin{pmatrix} G \\ 0 \\ 0 \end{pmatrix} u_t + \begin{pmatrix} I & 0 & 0 \\ 0 & G^r & 0 \\ 0 & 0 & G^d \end{pmatrix} \begin{pmatrix} w_t \\ n_t \\ p_t \end{pmatrix} \\ &= \bar{F}\bar{x}_t + \bar{G}u_t + \bar{L}\bar{w}_t. \end{aligned} \quad (3.58)$$

The LQ tracking control objective (3.11) then yields a super state regulator problem (3.3) with weighting matrices \bar{P}_0 , $\bar{Q}_{c,j}$ and $\bar{R}_{c,j}$ given by

$$\bar{P}_0 = \begin{pmatrix} H^T \\ -H^{r^T} \\ H^{d^T} \end{pmatrix} P_0 \begin{pmatrix} H & -H^r & H^d \end{pmatrix} \quad (3.59)$$

$$\bar{Q}_{c,j} = \begin{pmatrix} H^T \\ -H^{r^T} \\ H^{d^T} \end{pmatrix} Q_{c,j} \begin{pmatrix} H & -H^r & H^d \end{pmatrix} \quad (3.60)$$

$$\bar{R}_{c,j} = R_{c,j}. \quad (3.61)$$

The linear state-variable feedback solution to this LQ regulation problem,

$$\begin{aligned} u_{t+N-j} &= -(\bar{G}^T \bar{P}_{j-1} \bar{G} + \bar{R}_{c,j-1})^{-1} \bar{G}^T \bar{P}_{j-1} \bar{F} \hat{x}_{t+N-j} \\ &= \begin{pmatrix} K_{j-1}^x & K_{j-1}^r & K_{j-1}^d \end{pmatrix} \begin{pmatrix} \hat{x}_{t+N-j} \\ \hat{x}_{t+N-j}^r \\ \hat{x}_{t+N-j}^d \end{pmatrix}, \end{aligned} \quad (3.62)$$

produces a feedback control signal, u_t , which accounts for the present state position, the given future trajectory of the reference to be tracked, and the predictable component of the corrupting measurement noise.

The solution of the above composite LQ regulator problem would, at first sight, appear to involve the solution of a correspondingly large Riccati equation. However, as with the earlier LQ regulator problem involving only the plant model and reference model, the solution decomposes. In particular, if we divide \bar{P} conformably with the state \bar{x}_t ,

$$\bar{P} = \begin{pmatrix} P^{11} & P^{12} & P^{13} \\ P^{12^T} & P^{22} & P^{23} \\ P^{13^T} & P^{23^T} & P^{33} \end{pmatrix}, \quad (3.63)$$

then the block structure of \bar{G} forces the control gain matrix K only to depend upon P^{11} , P^{12} and P^{13} . Specifically,

$$\begin{aligned} K_j &= -(G^T P_j^{11} G + R_{c,j})^{-1} (G^T P_j^{11} F \quad G^T P_j^{12} F^r \quad G^T P_j^{13} F^d) \\ &= (K_j^x \quad K_j^r \quad K_j^d). \end{aligned} \quad (3.64)$$

Similarly, the RDE decomposes with P^{11} satisfying the regulation RDE (3.24), P^{12} satisfying the Lyapunov equation (3.25) and P^{13} satisfying the Lyapunov equation,

$$\begin{aligned} P_{j+1}^{13} &= (F - G(G^T P_j^{11} G + R_{c,j})^{-1} G^T P_j^{11} F)^T P_j^{13} F^d + H^T Q_{c,j} H^d \\ &= (F + G K_j^x)^T P_j^{13} F^d + H^T Q_{c,j} H^d, \end{aligned} \quad (3.65)$$

with initial condition $P_0^{13} = H^T P_0 H^d$.

A moment's reflection upon the reachability subspace of (3.58) allows the reader to realize that the control signal is able only to exert its influence upon the plant state, x_t , and that the other state components, x_t^r and x_t^d , really serve a rôle of providing predictive information to this control via the state-estimate feedback law (3.62). These ancillary states are themselves not reachable from the input u_t .

In the case of constant weighting matrices Q_c and R_c the sequences P_j^{11} , P_j^{12} and P_j^{13} will converge to steady-state values under the same stabilizability and detectability conditions (as earlier) that P_j^{11} converges. The time-varying tracking control gain (3.64) can then be replaced by the corresponding steady-state gain

$$\begin{aligned} K &= - \left(G^T P_\infty^{11} G + R_c \right)^{-1} \left(G^T P_\infty^{11} F \quad G^T P_\infty^{12} F^r \quad G^T P_\infty^{13} F^d \right) \\ &= \left(K^x \quad K^r \quad K^d \right). \end{aligned}$$

3.6.2 Observers for the Composite System

The composite super state model (3.58) contains the plant's state, the reference state and the noise process state. One may associate with this state model a combined measurement equation reflecting the available data,

$$\bar{y}_t = \begin{pmatrix} y_t \\ x_t^r \end{pmatrix} \tag{3.66}$$

$$= \begin{pmatrix} H & 0 & H^d \\ 0 & I & 0 \end{pmatrix} \begin{pmatrix} x_t \\ x_t^r \\ x_t^d \end{pmatrix} + \begin{pmatrix} I \\ 0 \end{pmatrix} q_t \tag{3.67}$$

$$= \bar{H} \bar{x}_t + \bar{v}_t. \tag{3.68}$$

A Kalman filter may now be constructed for this (super) state model and measurement equation, which is precisely in the form of (3.1). An immediate simplification is, however, possible because clearly one portion of the state, viz x_t^r , is available noise-free since it represents the future values of the reference trajectory. One proceeds exactly as in Section 3.5 to construct a Kalman filter observer with direct feedthrough to produce state estimates via (3.55),

$$\begin{pmatrix} \hat{x}_t \\ \hat{x}_t^r \\ \hat{x}_t^d \end{pmatrix} = \begin{pmatrix} (\hat{x}_t^m)_{1,\dots,n} \\ x_t^r \\ (\hat{x}_t^m)_{n+1,\dots,n+k} \end{pmatrix}. \tag{3.69}$$

It is clear from the preceding material that, through the invocation of the composite system, it is possible to formulate the complete LQ tracking

problem with colored measurement noise disturbances as a combined LQ plus KF, i.e. LQG, regulation task. For the control law we have

$$\begin{aligned}
 u_t &= (K^x \quad K^r \quad K^d) \begin{pmatrix} \hat{x}_t \\ \hat{x}_t^r \\ \hat{x}_t^d \end{pmatrix} \\
 &= (K^x \quad K^d) \begin{pmatrix} \hat{x}_t \\ \hat{x}_t^d \end{pmatrix} + K^r x_t^r \\
 &= (K^x \quad K^d) \begin{pmatrix} \hat{x}_t \\ \hat{x}_t^d \end{pmatrix} + K^r (zI - F^r)^{-1} G^r n_t. \quad (3.70)
 \end{aligned}$$

For the N -step-lookahead tracking problem one may take $n_t = r_{t+N}$. While, for the partial observer with direct feedthrough (3.55),

$$\begin{pmatrix} \hat{x}_{t+1} \\ \hat{x}_{t+1}^d \end{pmatrix} = (I - M^F H^m) F^m \begin{pmatrix} \hat{x}_t \\ \hat{x}_t^d \end{pmatrix} + (I - M^F H^m) G^m u_t + M^F y_{t+1}, \quad (3.71)$$

with M^F , F^m , G^m , H^m coming from (3.56) and the composite disturbance and plant model (3.47)–(3.49). Thus, we arrive at a unified description of the LQG controller as being composed of a dynamic state estimator with its gain matrix, M^P or M^F , and of a nondynamic linear state-variable feedback law with gain K . If either infinite horizon or receding horizon designs are performed then these gains are constant. This total feedback control system is then a finite-dimensional time-invariant linear dynamical system which has two inputs, the measured plant output y_t and the N -step-ahead reference value r_{t+N} , and a single output, the control signal u_t . We shall broach the issue of the nature of the transfer function of this dynamical system in Section 3.8; but first, for those readers whose patience we have already abused with our deluge of formulae, we now illustrate some typical behaviors of LQG controllers, and the rôle of their design parameters, via a numerical example.

3.7 Examples

We consider two examples. The first is a simple, minimum phase second order plant and the second is a more difficult third order non-minimum phase plant derived by discretizing a continuous-time plant. We shall portray the separate design effects of choice of control weighting, λ , choice of the noise model, and the effect of the prediction horizon, N_r , in the construction of a reference model.

A Simple System

The first system we consider is described in state-space form by the following matrices:

$$\begin{aligned} F &= \begin{pmatrix} -0.5 & 1 \\ -0.5 & 0 \end{pmatrix} \\ G &= \begin{pmatrix} 0.9 \\ -0.6 \end{pmatrix} \\ H &= (1 \ 0). \end{aligned}$$

This represents a simple minimum phase stable second order process with transfer function

$$\begin{aligned} &\frac{0.9 z^{-1} - 0.6 z^{-2}}{1 + 0.5 z^{-1} + 0.5 z^{-2}} \\ &= 0.9 \times \frac{1 - 0.6667 z^{-1}}{(1 + (0.25 + j0.6614)z^{-1})(1 + (0.25 - j0.6614)z^{-1})}. \end{aligned}$$

For the moment, suppose that the reference model is given by the third order process

$$\begin{aligned} F^r &= \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \\ G^r &= \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \\ H^r &= (1 \ 0 \ 0). \end{aligned}$$

This means that, in the LQ criterion, the deviation with respect to the next three reference values is penalized or, equivalently, that the control at time t is computed using knowledge of the next three references. The disturbance model is described by

$$F^d = 1, G^d = 1, H^d = 1.$$

That is, we postulate an integrator to take into account step-like load disturbances.

In the spirit of the earlier sections, we choose to solve an infinite time LQG problem with $Q_c = (H \ -H^r \ H^d)^T (H \ -H^r \ H^d)$ and $R_c = \lambda$ as




Figure 3.1: Step response for $\lambda = 0.1$

weighting matrices for the control criterion and with the following weighting matrices for the Kalman Filter

$$Q_o^m = \begin{pmatrix} 1 & 0 \\ 0 & \rho \end{pmatrix}$$

and

$$R_o^m = \rho.$$

We take $\rho = 1$ in all subsequent simulations. Because we consider cases where the true plant is used in both the design and the closed loop simulation, the observer (KF/KP) dynamics cancel and so do not affect the performance. Later we shall see that they do play a part in the closed loop robustness, where the plant model used for controller design is different from the true system.

We describe now a few cases which illustrate some features of the LQ problem.

1. We take $\lambda = 0.1$ and look at the response of the process to relatively rapid setpoint changes. As can be seen from Figure 3.1 the desired setpoint is not reached. This might be due to excessive weighting on the control signal in the LQ criterion.




Figure 3.2: Step response for $\lambda = 0.0001$

2. We perform the same simulation with $\lambda = 0.0001$. Now, we have almost zero static error in step response, as can be seen from Figure 3.2. This illustrates something of the rôle of λ in determining the output performance. Consideration of the control signals in this case would also show that the control effort increases as the weighting diminishes.
3. For the same value of lambda, Figure 3.3 now illustrates the response of the process to step changes of the disturbance acting on the output of the process. The dotted line represents the disturbance v_t acting on the output of the plant. Thanks to the disturbance model which contains an integrator, the disturbance is rejected after an initial transient perturbation.
4. Figure 3.4 and Figure 3.5 illustrate the response of the system to set point changes and to output disturbances respectively, under similar conditions to those in 2 and 3, when there is no disturbance model. The response to setpoint changes remains unchanged, but now the controlled system is unable to reject the disturbances.

These few simulations serve only the purpose of illustrating the influence of weighting matrices of the LQ criterion and the rôle of the disturbance



Figure 3.3: Disturbance rejection for $\lambda = 0.0001$ 

Figure 3.4: Step response without disturbance model




Figure 3.5: Disturbance rejection without disturbance model

model in the observer. It should be noted that zero static error in the step response could of course be obtained if the incremental control action Δu_t was used in the criterion, instead of u_t (compare with the GPC algorithm). We now change the plant to observe the effects of design variables in more difficult circumstances.

A More Difficult Plant

Consider the following plant, $P(z)$, which will arise in the subsequent material relatively frequently, and so is denoted ‘The Working Example’:

$$P(z) = \frac{-0.05359 z^{-1} + 0.5775 z^{-2} + 0.5188 z^{-3}}{1 - 0.6543 z^{-1} + 0.5013 z^{-2} - 0.2865 z^{-3}}. \quad (3.72)$$

This is a sampled data version of the continuous-time plant

$$\mathcal{P}(s) = \frac{-10s + 10}{(s + 1)(s^2 + 1.5s + 10)},$$

with sampling interval $T_s = 0.5$. Note that the continuous-time plant $\mathcal{P}(s)$ is non-minimum phase, as is its discretized version $P(z)$. The value of ρ is 1 and N_r , the dimension of the tapped-delay-line reference model, is taken initially to be 10.




Figure 3.6: The Working Example output response for $\lambda = 0.1$

We again perform two simulations indicating the effect on the performance and control signals of the weighting λ . Figures 3.6 and 3.7 depict the plant output and control signals with $\lambda = 0.1$ while Figures 3.8 and 3.9 show the corresponding signals for $\lambda = 0.0001$. The obvious point here is that the reduced penalty upon the control signal is reflected in a lessening of the constraints on its deviations. Thus wild control signals can often be associated with very small λ values, particularly with non-minimum phase plants. This problem is known as the bounded peaking problem of singular optimal control and has been extensively studied [FG78].

The preceding examples were performed with the reference model being finite impulse response and of dimension $N_r = 10$. This was helpful in demonstrating the effects of λ for this example. Now we shall briefly investigate the rôle of N_r in this problem. Specifically, we take $N_r = 3$ and compare the output response keeping $\lambda = 0.0001$. Figure 3.10 demonstrates that this smaller value of N_r dramatically affects the achievable performance.

The reasoning for this behavior rests in the precise infinite horizon LQ problem being solved here. At time t the controller is being selected to cause the plant output, y_t , to track the infinite reference signal based upon only N_r available future values, i.e. $\{r_{t+1}, r_{t+2}, r_{t+3}, \dots, r_{t+N_r}, 0, 0, \dots\}$. As the plant is relatively difficult to control, for small values of N_r the LQ tracking solution is found by focusing more on the achievable zero part of the tra-




Figure 3.7: The Working Example control signal for $\lambda = 0.1$




Figure 3.8: The Working Example output response for $\lambda = 0.0001$




Figure 3.9: The Working Example control signal for $\lambda = 0.0001$




Figure 3.10: The Working Example output response for $N_r = 3$ and for $\lambda = 0.0001$




Figure 3.11: The Working Example output response with an integrator reference model

jectory than on expending efforts to follow the short non-zero component. Intermediate values of N_r between 3 and 10 display intermediate performance. As this parameter, N_r , is central to the predictive control idea, it is evident that careful selection needs to be made and that thought needs to be given to a delicate interpretation of implicitly defined tracking criteria.

One method which springs to mind to overcome dimensionality problems associated with large N_r values is to replace the finite impulse response model for the reference with, say, a first order integrator model

$$r_{t+1} = r_t + w_t.$$

This captures the roughly constant nature of r_t . The control system then is driven with a difference signal $w_t = r_{t+1} - r_t$. Figure 3.11 shows the achieved performance with just such a control law. The steady-state step response error now goes to zero but at the price of slower transient response.

These examples illustrate some of the available design variables of LQ design based on predictive models, and indicate certain areas for care in the selection of criteria.

3.8 Closed Loop Control Formulae

Our goal in this section will be to develop simple formulae for the transfer functions linking the control signal, u_t , in the LQ or LQG feedback loop with the external signals coming into the loop, namely the reference signal, r_t , and the plant output measurement noise signal, v_t . These formulae are not new, like much of the work presented in this chapter, but they are difficult to find in textbooks, and we shall need them in our later analysis of the robust interplay between LQG control and recursive least squares identification in closed loop. Indeed, these formulae will enable us to examine how the spectral distribution of the noise and the reference signal affect the spectral distribution of the input signal of the plant, which in turn influences the convergence domain of the identified model. We shall therefore derive these formulae here for several possible configurations of the feedback observer/controller connection; we examine successively full-state feedback, state-estimate feedback with an observer without direct feedthrough (hereafter called Kalman predictor), and state-estimate feedback with a direct feedthrough observer (hereafter called Kalman filter). In order not to assault the sensibilities of the reader, we start off gently by first considering the easier case where there is no plant measurement noise. We establish the dynamical connections between the reference signal, r_t , and the control input, u_t , for this case.

Ideal Plant with Full State Feedback

Consider the state model,

$$\begin{aligned}x_{t+1} &= Fx_t + Gu_t, \\y_t &= Hx_t,\end{aligned}\tag{3.73}$$

with linear state-variable feedback involving direct measurements of the state vector, x_t ,

$$u_t = Kx_t + r_t,\tag{3.74}$$

where r_t is the external reference signal. Then, substituting (3.74) into (3.73), we have

$$x_{t+1} = (F + GK)x_t + Gr_t,$$

or, in transfer function terms,

$$x_t = (zI - F - GK)^{-1}Gr_t.$$

Resubstituting into (3.74) then yields

$$u_t = [I + K(zI - F - GK)^{-1}G]r_t \quad (3.75)$$

$$= [I - K(zI - F)^{-1}G]^{-1}r_t. \quad (3.76)$$

Thus we see that the transfer function linking the LQ control signal with the reference is given, in (3.76), by the inverse of a transfer function known as the return difference, which shall recur frequently in later developments. We now turn to the case where an observer is inserted into the loop.

Ideal Plant with State-Estimate Feedback and Kalman Predictor

When an observer is used in place of full state measurement, the control equation (3.74) is replaced by

$$\hat{x}_{t+1} = (F - MH)\hat{x}_t + Gu_t + My_t, \quad (3.77)$$

$$u_t = K\hat{x}_t + r_t. \quad (3.78)$$

A transfer function description of (3.77) is simply derived as

$$\hat{x}_t = (zI - F + MH)^{-1}Gu_t + (zI - F + MH)^{-1}My_t. \quad (3.79)$$

Whence, using $y_t = H(zI - F)^{-1}Gu_t$, the system transfer function, we have

$$\begin{aligned} \hat{x}_t &= (zI - F + MH)^{-1}Gu_t \\ &\quad + (zI - F + MH)^{-1}MH(zI - F)^{-1}Gu_t, \\ &= (zI - F + MH)^{-1}\{(zI - F) + MH\}(zI - F)^{-1}Gu_t \\ &= (zI - F + MH)^{-1}(zI - F + MH)(zI - F)^{-1}Gu_t \\ &= (zI - F)^{-1}Gu_t. \end{aligned} \quad (3.80)$$

Note that the observer poles, which are the zeros of $\det(\lambda I - F + MH)$ and are obviously presumed to be stable by design, are canceled to reach (3.80) above.

Continuing now, one substitutes from (3.80) into (3.78) to produce

$$u_t = K\hat{x}_t + r_t \quad (3.81)$$

$$= K(zI - F)^{-1}Gu_t + r_t \quad (3.82)$$

$$= [I - K(zI - F)^{-1}G]^{-1}r_t. \quad (3.83)$$

Several comments are in order at this point.

- The transfer function linking the LQG control signal, u_t , to the reference signal, r_t , is exactly the same as that, (3.76), derived in the case of application of full state feedback, except that the veracity of the result depends upon the cancellation of the observer dynamics in (3.80). That is, again the return difference inverse arises as the transfer function from reference input to control signal.
- Even though the observer poles cancel in the transfer function from r_t to u_t , the effect of these canceled dynamics will be felt in the decay of initial condition transients.
- The exact cancellation occurs only under the condition that the plant model, on which basis the observer is designed, is an exact description of the actual plant. In Chapter 5 we shall study the effect of unmodeled dynamics on the closed loop behavior of LQG controlled plants. It will be the object of Section 5.4.2 to compute new LQG controller formulae for the more realistic situation of a plant/model mismatch, in which case the observer poles do not cancel.

We next turn to the case where the observer contains a direct feedthrough term.

Ideal Plant with State-Estimate Feedback and Kalman Filter

Here the observer with delay, (3.77), is replaced by a direct feedthrough observer

$$\hat{x}_{t+1} = (F - MHF)\hat{x}_t + (G - MHG)u_t + My_{t+1}. \quad (3.84)$$

Rewriting this equation using (3.73) yields

$$\begin{aligned} \hat{x}_{t+1} &= (F - MHF)\hat{x}_t + (G - MHG)u_t + MHF x_t + MHG u_t \\ &= (F - MHF)\hat{x}_t + MHF x_t + G u_t. \end{aligned}$$

We may now substitute $x_t = (zI - F)^{-1}G u_t$ for the state, producing

$$\hat{x}_{t+1} = (F - MHF)\hat{x}_t + MHF(zI - F)^{-1}G u_t + G u_t$$

or

$$\begin{aligned} \hat{x}_t &= (zI - F + MHF)^{-1}[MHF(zI - F)^{-1} + I]G u_t \\ &= (zI - F + MHF)^{-1}[MHF + zI - F](zI - F)^{-1}G u_t \\ &= (zI - F)^{-1}G u_t. \end{aligned} \quad (3.85)$$

This is now identical to (3.80) and we may proceed as before to obtain the identical relation (3.76). Note that, as in the previous case with no direct feedthrough, the derivation of this transfer function involves the cancellation of the observer dynamics to reach (3.85). Once again, the design assumes that these poles are chosen to be stable.

We now consider the case where the plant output contains measurement noise v_t and construct the transfer function from v_t to u_t .

Plant with Measurement Noise: State-Estimate Feedback and Kalman Predictor

To effect this analysis we must replace the plant output equation by

$$y_t = Hx_t + v_t, \quad (3.86)$$

where v_t is the additive plant output measurement noise. Thus in the preceding calculations y_t must be replaced by $H(zI - F)^{-1}Gu_t + v_t$. One may proceed from this point as before or may use the linearity of the situation. Replacing y_t by (3.86) in (3.79) and using linearity, it follows from (3.80) that

$$\hat{x}_t = (zI - F)^{-1}Gu_t + (zI - F + MH)^{-1}Mv_t. \quad (3.87)$$

Substituting \hat{x}_t in (3.78) yields

$$u_t = [I - K(zI - F)^{-1}G]^{-1}\{r_t + K(zI - F + MH)^{-1}Mv_t\}. \quad (3.88)$$

The transfer function between the output noise, v_t , and the input signal, u_t , in this case of state-estimate feedback with a Kalman predictor, is thus given by

$$u_t = [I - K(zI - F)^{-1}G]^{-1}K(zI - F + MH)^{-1}Mv_t. \quad (3.89)$$

Plant with Measurement Noise: State-Estimate Feedback and Kalman Filter

When the state estimate is computed through a direct feedthrough observer, we can use linearity again to obtain, using (3.84) and (3.86),

$$\hat{x}_t = (zI - F)^{-1}Gu_t + (zI - F + MHF)^{-1}Mv_{t+1}. \quad (3.90)$$

Therefore, using (3.78) again, we get

$$\begin{aligned} u_t &= [I - K(zI - F)^{-1}G]^{-1}\{r_t + K(zI - F + MHF)^{-1}Mv_{t+1}\} \\ &= [I - K(zI - F)^{-1}G]^{-1} \\ &\quad \times \{r_t + [KM + K(F - MHF)(zI - F + MHF)^{-1}M]v_t\}. \end{aligned} \quad (3.91)$$

The conclusions of this section are as follows.

- The transfer function linking the reference input, r_t , and the closed loop control signal, u_t , in a linear state-variable feedback system is given by the inverse of the control return difference matrix. This result holds true whether the full state is measured or whether a state estimate is generated by an observer, with or without direct feedthrough. This is because the dynamics of the observer cancel in the closed loop for both observers. Since these observer dynamics are stable by design, their cancellation does not introduce direct stability problems. However, as remarked above, this exact cancellation evanesces when there is a plant/model mismatch.
- As for the transfer functions linking the measurement noise, v_t , to the closed loop control signal, u_t , direct comparison of (3.88) and (3.91) shows that they differ in the cases of observers possessing a direct feedthrough and of those that do not, leading to different closed loop noise gain properties. The bandwidth of the disturbance component of u_t is greater in the case of direct feedthrough. Thus the abilities of these two observers to influence the design are different.

A fuller analysis of filtering and prediction theory, which is beyond our brief here, indicates that the use of observers with direct feedthrough (and hence without propagation delay) admits superior control performance with noise [AM79]. In the analysis of robust LQ control, with which we shall deal in Chapter 5, this feature also has been established [Mac85], [IT86].

We note here that the results derived above refer to two specific features of linear state and state-estimate feedback. That is, we have shown that, at least in the case where the plant model is an exact description of the actual system, the control signal is always the same function of the reference signal, independently of the particular LQ or LQG feedback structure, and we have exhibited the effects of the observer dynamics on the transmission of the output measurement noise to the control input in the closed loop. The introduction of a (possibly fictitious) output measurement noise is one method of studying the robustness properties of the control system. Thus, for example, the requirement for careful observer design only becomes apparent after such a signal is introduced. More shall be said about robust LQG design and the rôle of the observer in Chapter 5.

3.9 GPC as LQG

It is immediately apparent from the comparison between the optimization criteria (2.7) and (3.11) that the GPC problem formulation should fit within the framework of the LQ problem. Indeed, the underlying system is linear and the cost is a quadratic tracking criterion. We conduct this section of our treatise in two parts: firstly we develop the formal theoretical connections which place the control law specification of GPC as a receding horizon LQG control law. Then, secondly, we re-present the example of Section 2.4 of control law computation by GPC methods, but this time we carry through the calculations using the equivalent LQG techniques. Naturally, this exercise merely verifies our theoretical knowledge but is included to convey something of the concrete nature of these esoteric design rules and also to persuade the intransigent reader who does not trust mathematics.³

3.9.1 Control Criterion Equivalence

Let the input–output model of Chapter 2, (2.6), be described by the equivalent state-space model,

$$\begin{aligned}\bar{x}_{t+1} &= F\bar{x}_t + G\Delta u_t + \bar{w}_t \\ y_t &= H\bar{x}_t + \bar{v}_t.\end{aligned}$$

The finite horizon control signal $\tilde{\mathbf{u}}$ (see (2.18)), resulting from the minimization of the GPC criterion (2.7) with zero tracking $\{r_t\} = 0$ regulation objective and with constant λ , is the same as that obtained by the minimization of the LQ regulation criterion

$$\begin{aligned}J(N, x_t) &= \mathbf{E}\left\{x_{t+N}^T P_0 x_{t+N} + \sum_{j=0}^{N-1} \left\{x_{t+j}^T Q_{c, N-j} x_{t+j} \right. \right. \\ &\quad \left. \left. + \Delta u_{t+j}^T R_{c, N-j} \Delta u_{t+j} \right\}\right\},\end{aligned}\tag{3.92}$$

provided the following substitutions are made:

$$N_2 = N\tag{3.93}$$

$$Q_{c, t} = \begin{cases} H^T H & \text{if } t = 0, \dots, N_2 - N_1 \\ 0 & \text{if } t = N_2 - N_1 + 1, \dots, N_2 - 1 \end{cases}\tag{3.94}$$

³Readers who are offended by the notion of an example confirming that which has already been established via mathematics should take heart in the fact that the authors resisted the temptation to include also a multitude of simulation studies still further evidencing the same result.

$$R_{c,t} = \begin{cases} \infty I & \text{if } t = 0, \dots, N_2 - N_u - 1 \\ \lambda I & \text{if } t = N_2 - N_u, \dots, N_2 - 1 \end{cases} \quad (3.95)$$

$$P_0 = H^T H. \quad (3.96)$$

For the tracking problem, we have equivalent specifications of the weighting matrices via (3.18)–(3.20). We note that the formal mechanism employed to enforce the control constraints of (2.7) is to select the control weighting as infinite for controls after time N_u . We also note that this formal requirement of infinite weightings is easily realized and accommodated within the LQ gain computations (3.4) and (3.5), because for those steps where the control weighting is infinite the RDE becomes simply a Lyapunov equation. Thus, the criteria (2.7) and (3.92) with the above substitutions are well posed and identical. Therefore the control sequences $\{\Delta u_{t+j}, j = 0, \dots, N_u - 1\}$ obtained through the minimization of these two criteria are identical, provided they are computed on the basis of the same information. The GPC solution $\tilde{\mathbf{u}}$ of (2.23) is based on the information available at time t , i.e. the vector \mathbf{f} of (2.17) is constructed using $\{u_s, y_s, s \leq t\}$. Hence the solution sequence $\{\Delta u_{t+j}, j = 0, \dots, N_u - 1\}$ of GPC is computed in open loop on the basis of information available at time t . The solution sequence (3.26) of LQG, rewritten here in incremental form,

$$\Delta u_{t+N-j} = K_{j-1} \hat{x}_{t+N-j|t+N-j}, \quad j = 1, \dots, N$$

is computed under the presumption that the control at time $t+k$ uses all information $\{u_s, y_s, s \leq t+k\}$ available at time $t+k$. Therefore the two solution sequences $\{\Delta u_{t+j}, j = 0, \dots, N_u - 1\}$ obtained from GPC and LQG need not coincide exactly except for $j = 0$, where their computation is based on the same information. We note that another instance where these solutions coincide is in the case of purely deterministic systems, $\bar{w}_t = 0, \bar{v}_t = 0$, since in such cases $\hat{x}_{t+k|t+k}$ can be computed exactly from $\hat{x}_{t|t}$ and the intervening control signal. The distinction between GPC and LQG formulations has been explored with considerable subtlety by Peterka [Pet89].

Now, as identified earlier, the GPC control law does not consist of the complete application of the finite horizon control signal but rather it implements a receding horizon strategy. That is to say, only the control signal value Δu_t from the above solution is applied at time t and the solution of a new finite horizon problem is computed for the next instant. The import of our discussion above is that the controls Δu_t resulting from the minimization of the GPC and LQG criteria are the same, since they are constructed based on the same information. In terms of the LQ problem it is then clear

that the GPC control signal is given by the receding horizon LQ law,

$$\Delta u_t^{GPC} = -(G^T P_{N-1} G + \lambda I)^{-1} G^T P_{N-1} F \hat{x}_t, \quad (3.97)$$

where P_{N-1} is derived as the solution of the RDE (3.4) with variables as specified above in (3.94)–(3.96) and where \hat{x}_t is the filtered state estimate, $\hat{x}_{t|t}$, based on $\{u_s, y_s, s \leq t\}$.

Hence, we see that GPC implements a stationary, i.e. time-invariant, feedback law derived from a particular LQ criterion. That is, the control gain in (3.97) is fixed by the RDE solution to the optimization problem. To complete the establishment of GPC as a subset of LQ control it behoves us now to consider the observer generated by GPC via the Diophantine recursions and predictor equations (2.8)–(2.13).

We make the remark here that the rôle of the predictions \mathbf{f} in GPC (see (2.17)) is entirely analogous to the rôle of the state estimate $\hat{x}_{t|t}$ in the LQ formulation, just as the role of \mathbf{r} is identical to that of x_t^r . Indeed, \mathbf{f} contains the ‘free response’ prediction of the output given $\{u_s, y_s, s \leq t\}$ assuming that future inputs are zero, while $\hat{x}_{t|t}$, being based on the same information from the past, contains sufficient information also to compute these free response predictions. This general equivalence between observers and predictors has been explored in [ÅW84], where it is shown that the eigenvalues of $F - M^P H$ in the observer equation (3.38) are the same as the zeros of the polynomial C used to construct the predictors via (2.8). Since $F - M^P H = F(I - M^F H)$, it follows that the eigenvalues of $F - M^F H F$ in the observer equation (3.39) are also the same as the zeros of C . In particular, if $C(q^{-1})$ is chosen to be unity (i.e. $C^*(q) = q^{N_2} C(q^{-1})$ is q^{N_2}), then the equivalent observer will be deadbeat. The introduction of a non-trivial C polynomial will allow the control signal to be written as a linear combination of u_t^f and y_t^f from (2.12) and (2.13). These issues of the use of the polynomial in (2.8) to determine observer dynamics are advanced in [CMT87] where the choice of C is involved as part of the control design and not necessarily as a part of the explicit system modeling. This viewpoint will be reinforced in Chapter 7, where we shall discuss a potentially robust LQG design attuned to the adaptive context, which includes the formal selection of observer properties.

3.9.2 An Example

We return to the system and control problem used in Section 2.4 now to amplify and consolidate our claims of equivalence between GPC and LQG methods. Firstly we begin by solving the LQ regulator problem connected

with GPC before augmenting this to the full tracking problem. This allows us to demonstrate some of our LQ claims as well.

The plant of (2.25) is

$$y_t = 1.7y_{t-1} - 0.7y_{t-2} + 0.9\Delta u_{t-1} - 0.6\Delta u_{t-2} + \xi_t, \quad (3.98)$$

which has a minimal state-space description in observable canonical form,

$$\begin{aligned} \begin{pmatrix} x_{t+1}^1 \\ x_{t+1}^2 \end{pmatrix} &= \begin{pmatrix} 1.7 & 1 \\ -0.7 & 0 \end{pmatrix} \begin{pmatrix} x_t^1 \\ x_t^2 \end{pmatrix} + \begin{pmatrix} 0.9 \\ -0.6 \end{pmatrix} \Delta u_t \\ y_t &= (1 \ 0)x_t + \xi_t. \end{aligned}$$

That is, taking our information from Section 2.4, we have the following assignments:

$$\begin{aligned} F &= \begin{pmatrix} 1.7 & 1 \\ -0.7 & 0 \end{pmatrix} & G &= \begin{pmatrix} 0.9 \\ -0.6 \end{pmatrix} \\ H &= (1 \ 0) & Q_c &= \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \\ R_c &= 0.1 & P_0 &= \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \\ N_1 &= 1 & N_2 &= N_u = 3. \end{aligned} \quad (3.99)$$

We compute the solution of the RDE (3.4) from initial condition P_0 until P_2 ,

$$P_2 = \begin{pmatrix} 1.3566 & 0.2354 \\ 0.2354 & 0.1705 \end{pmatrix}.$$

Next the control gain is computed as in (3.97),

$$\begin{aligned} K_2 &= -(G^T P_2 G + R_c)^{-1} G^T P_2 F \\ &= (-1.7483 \quad -1.0733). \end{aligned}$$

This is then followed by the computation of a direct feedthrough observer in the form of (3.56). Write

$$M^F = \begin{pmatrix} m_1 \\ m_2 \end{pmatrix},$$

and note that the characteristic equation of the observer satisfies,

$$\det(zI - F + M^F H F) = z^2 + (m_2 - 1.7(1 - m_1))z + 0.7(1 - m_1).$$

Thus to create a deadbeat observer we take

$$M^F = \begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

yielding a direct feedthrough observer as in (3.55),

$$\hat{x}_{t+1} = \begin{pmatrix} 0 & 0 \\ -0.7 & 0 \end{pmatrix} \hat{x}_t + \begin{pmatrix} 0 \\ -0.6 \end{pmatrix} \Delta u_t + \begin{pmatrix} 1 \\ 0 \end{pmatrix} y_{t+1}. \quad (3.100)$$

When this observer is coupled to the control law $\Delta u_t = K_2 \hat{x}_t$ with K_2 above, we achieve the controller transfer function

$$\begin{aligned} \Delta u_t &= K_2 \left(zI - F + M^F H F - (G - M^F H G) K_2 \right)^{-1} M^F y_{t+1} \\ &= \frac{1.7483q - 0.7513}{q(q + 0.644)} y_{t+1} \\ &= -0.644 \Delta u_{t-1} + 1.7483 y_t - 0.7513 y_{t-1}, \end{aligned}$$

which is identical to the feedback compensator part derived from the GPC formulation in (2.26).

To develop this into a full tracking problem, and this solution could have been addressed from the start, it suffices to consider the RDE with

$$\begin{aligned} F &= \begin{pmatrix} 1.7 & 1 & 0 & 0 & 0 & 0 \\ -0.7 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \\ G &= (0.9 \quad -0.6 \quad 0 \quad 0 \quad 0 \quad 0)^T \\ P_0 = Q_c &= \begin{pmatrix} 1 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \end{aligned}$$

The RDE is iterated twice to produce P_2 , and the feedback gain K_2 is computed via (3.5),

$$K_2 = (-1.7483 \quad -1.0733 \quad 0 \quad 0.8947 \quad 0.0929 \quad 0.0095).$$

Thus the closed loop control strategy by the LQ tracking method is

$$\begin{aligned} \Delta u_t &= -0.644 \Delta u_{t-1} + 1.7483 y_t - 0.7513 y_{t-1} \\ &\quad + 0.8947 r_{t+1} + 0.0929 r_{t+2} + 0.0095 r_{t+3}, \end{aligned}$$

which is identical to (2.26).

3.10 Conclusions

Our aim has been to present a brief treatment of the principal features of LQG control design which will have immediate application in our treatment of the adaptive control of plants modeled by linear systems. In particular, we have chosen to display, in the previous section, how the general LQG setting subsumes the GPC control design.

The main points raised by way of this development may be summarized as follows.

LQ Control

- An LQ regulator problem commences from a linear state-space plant model and a quadratic cost function. This cost function may be finite horizon or infinite horizon. The infinite horizon regulator with stationary weighting matrices produces a time-invariant control law.
- Receding horizon LQ control laws may also be derived where the control value at time t , u_t , is the first element in the solution of a fixed finite horizon problem solved at each time, t . This procedure also produces a time-invariant control law.
- The control sequence solving an LQ regulator problem is given by linear state-variable feedback. These feedback gains are derived from the solution of the Riccati Difference Equation or the Algebraic Riccati Equation. The stability properties of these solutions will be the topic of the next chapter.
- An LQ tracking objective may be recast, via the use of state augmentation, as an LQ regulation problem. The computational burden of the solution of this larger-dimensional problem splits conveniently into two smaller subproblems.

State Estimation and LQG

- The implementation of an LQ control law requires the use of a state estimator/observer. One may derive this observer as the dual to the construction of the LQ control law, the Kalman predictor.
- With a strictly proper plant, i.e. possessing a zero at infinity, one may use either an observer with delay, the Kalman predictor, or an observer with direct feedthrough, the Kalman filter, without needing to address issues of algebraic loops in the feedback system.

- It is possible to design state estimators for processes subject to colored measurement noise. This requires the augmentation of the state-space model for the plant to include noise coloring states. This is dual to the inclusion of a tracking signal in LQ and introduces similar computational issues. Both the plant state estimate and the predictable component of the noise are produced by the observer.
- By combining LQ linear state-variable feedback and KF observer construction, one produces an LQG controller. If the gains of these separate elements are constant, then the controller is a linear time-invariant system.
- The constructs of reference and noise models from LQ and KF separately carry over directly to the LQG design. In particular, the block decomposition of the solution may be preserved and the control input signal separated into components related to each subsystem.

GPC and LQG

- The GPC control law specification may be subsumed within the LQG framework as a receding horizon LQ tracking controller with specific choices for the cost function weighting matrices.
- The mechanism for forcing the GPC controller to admit zero tracking error, optimizing over Δu_t rather than over u_t , may be simply included in the LQ systematization.
- Since the GPC control law is independent of a state coordinate system, we see that the choice of $Q_c = H^T H$ in the LQ criterion (which applies to the state equation involving only F and G) has the effect of forcing a state-space coordinate-free controller. That is, this choice of Q_c provides the mechanism for incorporating the measurement equation information into the control law selection. As adaptive control is primarily concerned with input-output features of the linear system, where particular state-variable realizations may not have any special significance since the plant model itself is derived from input-output data, it is an important feature of LQG control that the controllers produced be capable of being coordinate-free.
- The prediction vector \mathbf{f} of the GPC formulation contains the same information as the state $\hat{x}_{t|t}$ of the direct feedthrough observer of LQG. Their design, therefore, and the selection of the GPC polynomial $C(z)$

can be reposed as a Kalman filtering issue, where this choice is connected to presumed noise models.

- While GPC is presented via the receding horizon design philosophy, its concerns are definitely associated with the infinite horizon properties of stability and detuned minimum variance performance. The artefacts of choosing design variables, such as $Q_{c,j}$, $R_{c,j}$, N_2 , N_u and λ , for the finite horizon problem should really be assessed on the infinite horizon.

Following on from this presentation of the principal features of LQG control in our adaptive setting and its ability to encompass the empirically supported GPC rule, we are next beholden to investigate the asymptotic stability properties of these schemes. This is the subject of Chapter 4, where we shall study intimately both the stability and performance properties of receding horizon LQ and infinite horizon LQ strategies. In this way we shall provide a theoretical basis for these studies and make more complete the rôle of LQG theory in guaranteeing the adequate behavior of GPC. Through this connection we open the floodgates for decades of accumulated LQG and classical control design to inundate the practical adaptive control arena. But more of this in the fullness of time ...

Chapter 4

Stability and Performance Properties of Receding Horizon LQ Control

4.1 Introduction

In Chapter 3, we have identified three different manifestations of linear quadratic regulation design criteria, an infinite horizon regulation criterion, a finite horizon criterion and a receding horizon one. We have then shown that GPC can be recast into a receding horizon LQ control problem, with a particular, coordinate-free, choice of the criterion. While the asymptotic stability of the infinite horizon controller is both well publicized and a rather obvious consequence of its formulation as an asymptotic problem, and while that of the finite horizon controller is irrelevant given its open loop and time-varying nature, that of the receding horizon controller would appear *a priori* to have no built-in justification.

Indeed, as receding horizon LQ control is a feedback strategy effected through the application of a stationary feedback law applied prospectively over an infinite interval, questions of stability naturally arise while solutions are sometimes slow to emerge. One of the major drawbacks of GPC control is that closed loop stability is not guaranteed, and that attempts at producing stability results for GPC on the basis of its explicit input–output description have been remarkably unsuccessful, usually necessitating the abandonment of a specific control performance.

Having established the exact state-space LQ equivalent of GPC, we are now in a position to bring to bear the powerful tools and properties of Riccati equations on the resolution of these stability questions for receding horizon

LQ controllers in general, and for GPC in particular. This is the starting point of the analysis developed in this chapter, where we will use recently developed properties of the solutions of the Riccati equations to enforce closed loop stability of receding horizon controllers. The reason why GPC in its generic formulation may produce unstable closed loops whatever the games played with the finite horizons N_u and N_2 will then become transparent, and so will the necessity, for guaranteed closed loop stability, to resort to the more classical state-space LQ formulation with appropriate choices of the weighting matrices. Of course, in a nonadaptive implementation, the closed loop stability of any designed controller can always be verified before actually closing the loop on the plant, and nominal stability margins can be computed as well. But in an adaptive implementation the advantage of using controller designs with guaranteed nominal stability margins becomes self-evident.

Somewhat unsurprisingly perhaps, our results will show that receding horizon controllers can be designed with guaranteed asymptotic closed loop stability provided their cost function weighting matrices can be chosen in such a way that the problem becomes equivalent to an infinite horizon LQ problem with different (but related) weightings. Further, we provide a simple test for stability with the receding horizon controller.

Our journey of discovery through the stability properties of receding horizon control will then eventually lead us to the conclusion that, if closed loop stability is a key requirement (and we don't know of many practical problems where it is not), one might as well resort to the tested and proven infinite horizon LQ controller design formulation.

An LQ formulation of a controller design procedure incorporates an obvious design and performance evaluation objective. One side-effect of our analysis, and an added bonus to the reader, is that we shall confront the issue of how the achieved performance of a receding horizon strategy compares to the theoretical optimal performance over the infinite horizon.

So much for introduction, but what are the goods? We begin in this chapter by studying the closed loop stability of receding horizon LQ controllers by the use of monotonicity of solutions of Riccati Difference Equations (RDE). One particular way of enforcing this monotonicity is a scheme due to Kwon and Pearson [KP78], which we present thereafter. The stability behavior (and possible misbehavior) of GPC is then illustrated. Finally, the comparative performances of receding horizon and infinite horizon LQ schemes are presented.

4.2 Monotonicity and Stability of Receding Horizon LQ Control

As remarked above, finite horizon LQ feedback control strategies do not guarantee closed loop asymptotic stability. Indeed, the mere question of asymptotic stability being associated with a finite horizon controller is clearly vacuous since a finite horizon criterion is not designed to deliver such a property, but this issue has direct interest with receding horizon strategies. Therefore questions need to be asked concerning mechanisms for guaranteeing asymptotic stability properties of receding horizon controllers.

4.2.1 Stability via the ARE

The fundamental LQ asymptotic stability result derives from the stationary infinite horizon regulator problem and the properties of the solution of the Algebraic Riccati Equation (3.8) (ARE). We have:

Theorem 4.1 [dSGG86]

Consider the ARE associated with an infinite horizon LQ control problem,

$$P = F^T P F - F^T P G (G^T P G + R)^{-1} G^T P F + Q \quad (4.1)$$

where

- $[F, G]$ is stabilizable,
- $[F, Q^{1/2}]$ has no unobservable modes on the unit circle,
- $Q \geq 0$ and $R > 0$.

Then

- there exists a unique, maximal, non-negative definite symmetric solution \bar{P} .
- \bar{P} is a unique stabilizing solution, i.e.

$$F - G(G^T \bar{P} G + R)^{-1} G^T \bar{P} F \quad (4.2)$$

has all its eigenvalues strictly within the unit circle.

We shall call the solution \bar{P} above the *stabilizing* solution of the ARE. We note that (4.2) is the state transition matrix of the closed loop system when the stationary control law (3.9) is used. Hence Theorem 4.1 is the fundamental closed loop stability result for infinite horizon LQ control.

There exist other asymptotic stability results concerning the stability of the infinite horizon LQ optimal feedback system (3.5) as a time-varying linear control system. But the hopes of receding horizon strategists rest firmly on the following convergence result

Theorem 4.2

Consider the ARE (4.1) above and its stabilizing solution \bar{P} , and consider the RDE

$$P_{t+1} = F^T P_t F - F^T P_t G (G^T P_t G + R)^{-1} G^T P_t F + Q. \quad (4.3)$$

Then, provided $[F, G]$ is stabilizable, $[F, Q^{1/2}]$ is detectable and $P_0 \geq 0$, $P_t \rightarrow \bar{P}$ as $t \rightarrow \infty$.

The conventional wisdom (see for example [CMT87], [Wou77], [WKR79]) of receding horizon has been to invoke the above results to argue that, provided N is taken sufficiently large, $[F, G]$ is stabilizable and $[F, Q^{1/2}]$ is detectable, one will ensure that $F - G(G^T P_t G + R)^{-1} G^T P_t F$ has its eigenvalues all within the unit circle, for any $t \geq N$. Indeed, this was the central motivation for the replacement of the one-step-ahead GMV controller by the family of long range predictive controllers. The issue however has always been: how big a value of N needs to be taken and further how can this be affected by choice of P_0 ? We shall now briefly present some results, due primarily to Poubelle, which yield connections between the RDE solutions and the ARE stability results.

We begin by rewriting the RDE as an ARE,

$$P_t = F^T P_t F - F^T P_t G (G^T P_t G + R)^{-1} G^T P_t F + \bar{Q}_t. \quad (4.4)$$

This equation plays such a pivotal rôle in this theory that Mademoiselle Poubelle christened it the Fake Algebraic Riccati Equation (FARE) — a convention we preserve. Notice that this is not so much a rewriting of the RDE as a definition for \bar{Q}_t :

$$\bar{Q}_t = P_t - F^T P_t F + F^T P_t G (G^T P_t G + R)^{-1} G^T P_t F \quad (4.5)$$

or, in terms of differences of successive solution matrices,

$$\bar{Q}_t = Q - (P_{t+1} - P_t). \quad (4.6)$$

Clearly, while we have not altered the RDE in viewing it as a masquerading ARE, we do have the immediate result following from Theorem 4.1 and (4.4).

Theorem 4.3

Consider the FARE (4.4), or (4.5) defining the matrix \bar{Q}_t . If $\bar{Q}_t \geq 0$, $R > 0$, $[F, G]$ is stabilizable, $[F, \bar{Q}_t^{1/2}]$ is detectable, then P_t is stabilizing, i.e.

$$\bar{F}_t = F - G(G^T P_t G + R)^{-1} G^T P_t F \quad (4.7)$$

has its eigenvalues all strictly within the unit circle.

This theorem forms the centerpiece of our developments to translate stability properties from the ARE arena to the RDE. In particular, we shall explore the implications of the defining relation (4.6) for \bar{Q}_t in terms of the difference between successive solution matrices. We note that, unfortunately, Theorem 4.3 does not admit an *if and only if* statement (see [BGPK85]).

Referring to Section 3.2.3, we note that the asymptotic stability of the closed loop obtained from applying the N -step receding horizon LQ controller (3.10) to the system (3.1) rests upon the eigenvalue locations of the closed loop state transition matrix

$$\bar{F}_{N-1} = F - G(G^T P_{N-1} G + R)^{-1} G^T P_{N-1} F. \quad (4.8)$$

This matrix will have all its eigenvalues strictly inside the unit circle if the associated FARE (4.4), with $t = N - 1$, satisfies the prescriptions of Theorem 4.3.

We observe, therefore, that just as the finite horizon LQ controller is naturally connected to the associated RDE, the infinite horizon LQ controller is naturally connected to the associated ARE, while the receding horizon LQ controller is naturally connected to the associated FARE. The closed loop asymptotic stability properties of the infinite and receding horizon controllers are derived from the properties of the associated ARE and FARE. More precisely, assuming that $[F, G]$ is stabilizable and that $R > 0$, we have the following Table 4.1.

We note that, given an N -step receding horizon control criterion with weighting matrices Q and R , the asymptotic stability properties of the controller are identical to those of an associated infinite horizon LQ controller with the same R but a different Q , i.e. Q is replaced by \bar{Q}_{N-1} . We shall see shortly that one way of guaranteeing closed loop stability of the receding horizon controller is to enforce monotonic decrease of the solution sequence P_t of the RDE, thereby forcing $\bar{Q}_{N-1} \geq Q$, (see (4.6)). Whence we notice that, having set up a receding horizon LQ problem with a designed Q , the stability properties of the corresponding closed loop will be those of an infinite horizon LQ problem with a larger Q . We shall return in Section 4.5 to the consequences of this observation on the performance of the corresponding closed loop system.

Optimal control problem	Riccati equation	Stability condition (under $[F, G]$ stabilizable, $R > 0$)
Finite horizon N	RDE	Not relevant
Infinite horizon	ARE	$Q \geq 0$ and $[F, Q^{1/2}]$ detectable
Receding horizon N	FARE	$\bar{Q}_{N-1} \geq 0$ and $[F, \bar{Q}_{N-1}^{1/2}]$ detectable

Table 4.1: Control problems and stability conditions

We now return to Theorem 4.3 and examine separately how the two conditions on \bar{Q}_t ($\bar{Q}_t \geq 0$ and $[F, \bar{Q}_t^{1/2}]$ detectable) can be made to hold from conditions on Q and the RDE. As detectability of matrix pairs is central to our arguments, we first state the following simple lemma.

Lemma 4.1

Given two non-negative definite symmetric matrices Q_1 and Q_2 satisfying

$$Q_1 \leq Q_2$$

then $[F, Q_1^{1/2}]$ detectable implies $[F, Q_2^{1/2}]$ detectable.

Proof The detectability assumption on $[F, Q_1^{1/2}]$ requires that for every eigenvalue, λ , of F which is greater in magnitude than unity and for its associated left eigenvector w , $wQ_1^{1/2} \neq 0$. By the inequality of the lemma statement, we also see that $wQ_1^{1/2} \neq 0$ implies that $wQ_2^{1/2} \neq 0$. \mathcal{CQFD}

We now draw the connection between decreasing properties of the solution of the RDE and stability of \bar{F}_t from (4.7).

Corollary 4.1 [BGPK85]

If the RDE, with $[F, G]$ stabilizable, has $[F, Q^{1/2}]$ detectable and if P_t is nonincreasing at t , i.e.

$$P_{t+1} \leq P_t,$$

then \bar{F}_t of (4.7) is stable.

This result follows directly from (4.6) and Lemma 4.1. We shall see shortly that the monotonicity properties are strongly determined by the initial conditions of the RDE, and so stability can sometimes be assured for an entire sequence of $\{\bar{F}_t\}$.

4.2.2 Monotonicity Properties of the RDE

We begin by stating a recent pleasing result of de Souza [dSo89] on the comparative properties between solutions of similar RDEs.

Lemma 4.2 [dSo89]

Let P_t^1 and P_t^2 be the solutions of two RDEs (4.3) with the same F , G and R matrices but possibly different Q s: Q^1 and Q^2 respectively. Then, the matrix

$$\tilde{P}_t = P_t^2 - P_t^1$$

satisfies the equation

$$\tilde{P}_{t+1} = \bar{F}_t^{1T} \tilde{P}_t \bar{F}_t^1 - \bar{F}_t^{1T} \tilde{P}_t G (G^T P_t^2 G + R)^{-1} G^T \tilde{P}_t \bar{F}_t^1 + \tilde{Q} \quad (4.9)$$

or

$$\tilde{P}_{t+1} = \bar{F}_t^{1T} \tilde{P}_t \bar{F}_t^1 - \bar{F}_t^{1T} \tilde{P}_t G (G^T \tilde{P}_t G + \tilde{R}_t)^{-1} G^T \tilde{P}_t \bar{F}_t^1 + \tilde{Q},$$

where

$$\begin{aligned} \bar{F}_t^1 &= F - G(G^T P_t^1 G + R)^{-1} G^T P_t^1 F \\ \tilde{Q} &= Q_2 - Q_1 \\ \tilde{R}_t &= G^T P_t^1 G + R. \end{aligned}$$

This lemma represents a simplification of the results of Nishimura [Nis67] and [BGPK85] and is a discrete-time counterpart to those of [PBG88]. Sources close to the President have revealed to us that a similar result was published in [Sam80], a copy of which we do not possess.

Directly from Lemma 4.2 we have the following results, the first of which extends the continuous-time observations of Kailath [Kai75].

Theorem 4.4 [BGPK85]

If the non-negative definite solution P_t of the RDE (4.3) is monotonically nonincreasing at one time, i.e.

$$P_{t+1} \leq P_t, \quad \text{for some } t,$$

then P_t is monotonically nonincreasing for all subsequent times,

$$P_{t+k+1} \leq P_{t+k}, \quad \text{for all } k \geq 0.$$

Proof Take P_t^1 as P_t , P_t^2 as P_{t+1} , $Q^1 = Q^2$ in Lemma 4.2. Then, from (4.9), we see that \tilde{P}_t being nonpositive definite at time t implies that it remains nonpositive definite and therefore we have the monotonicity of the theorem statement. CQFD

Theorem 4.5 [BGPK85]

If the solution P_t of the RDE (4.3) is monotonically nondecreasing at one time, i.e.

$$P_{t+1} \geq P_t, \quad \text{for some } t,$$

then P_t is monotonically nondecreasing for all subsequent times,

$$P_{t+k+1} \geq P_{t+k}, \quad \text{for all } k \geq 0.$$

Proof Take P_t^1 as P_{t+1} , P_t^2 as P_t , $Q^1 = Q^2$ in Lemma 4.2. CQFD

These monotonicity conditions on the solution may be extended to include monotonicity of the differences between successive solutions, using the same device.

Theorem 4.6 [dSo89]

If the solution P_t of the RDE (4.3) has a nonpositive definite second difference at time t , i.e.

$$P_{t+2} - 2P_{t+1} + P_t \leq 0,$$

then, for all $k \geq 0$,

$$P_{t+k+2} - 2P_{t+k+1} + P_{t+k} \leq 0.$$

With these monotonicity results under our wings we now turn back to their implications for the stability of closed loop systems derived from the solutions of the RDE.

4.2.3 Stability via Monotonicity of the RDE

By appealing to the monotonicity/stability connection in Corollary 4.1 we may directly develop the following theorems on stability of \bar{F}_t of (4.7) thereby rederiving the combined results of [BGPK85], [PPGB86], [PBG88], [dSo89].

Theorem 4.7

Consider the RDE (4.3). If

- $[F, G]$ is stabilizable;
- $[F, Q^{1/2}]$ is detectable;
- $P_{t+1} \leq P_t$ for some t ;

then \bar{F}_k , given by (4.7) with P_k , is stable for all $k \geq t$.

Proof The monotonicity of P_t at one step implies, by Theorem 4.4, the monotonicity at all subsequent steps. Corollary 4.1 then completes the proof.

CQFD

We have the following immediate consequence:

Corollary 4.2 [PBG88]

Consider the RDE (4.2). If

- $[F, G]$ is stabilizable;
- $[F, Q^{1/2}]$ is detectable;
- $\bar{Q}_0 \geq Q$;

then \bar{F}_t given by (4.7) is stable for all $t \geq 0$.

Proof The third assumption implies $P_1 \leq P_0$, and the result then follows from the previous theorem.

CQFD

The above results link the monotonic decreasing nature of the solution of the RDE with the solution's stability properties. This will form the core of our stability arguments to be advanced later in this chapter where initial conditions will be sought which cause the P_t sequence to be monotonically nonincreasing *ab initio*, i.e. which force $\bar{Q}_0 \geq Q$, thereby yielding stability of \bar{F}_t for all t . This will then imply that receding horizon controllers will produce closed loop stability whatever the horizon N used in the criterion. These stability aspects of monotonically decreasing P_t have been used in [BTP86] to derive stable filtering designs based effectively only on the choice of P_0 . Here we also discuss possible stability issues when P_t is not monotonically nonincreasing but where the difference $P_{t+1} - P_t$ is monotonically nonincreasing. We quote the following result from de Souza [dSo89]:

Theorem 4.8

Let P_t be the solution of the RDE (4.3) and \bar{Q}_t be defined by (4.5). If for some t :

- $[F, G]$ is stabilizable;
- $[F, \bar{Q}_t^{1/2}]$ is detectable;
- $P_{t+2} - 2P_{t+1} + P_t \leq 0$;

then \bar{F}_k is stable for all $k \geq t$.

Proof The key idea used is that the subsequent nonpositive definiteness of the second difference of P_k is implied by the theorem conditions and Theorem 4.6. Thus after time t , $P_{k+1} - P_k$ is a nonincreasing sequence of symmetric matrices. Therefore \bar{Q}_k is nondecreasing and so always greater than \bar{Q}_t . *CQFD*

We have now established the theoretical basis for the derivation of stability conditions for solutions of the RDE. As was demonstrated in the preceding chapter, the GPC receding horizon control law implements a fixed linear state-variable feedback gain computed as the solution of a finite horizon LQ problem, that is as the end solution of a finite horizon RDE. The stability properties of this and similar controllers will be addressed in the next section by using these monotonicity ideas. But we conclude this section with some fallacious conjectures in the spirit of [BGPK85] to help clarify matters.

Recall that $P_t \geq P_\infty$ is a necessary condition for the RDE solution to be monotonically nonincreasing at time t . This follows trivially from Theorem 4.4 and Theorem 4.2. This, however, is not sufficient.

Fallacious Conjecture 4.1

If P_∞ is stabilizing and $P_0 > P_\infty$, then P_0 is stabilizing.

Counterexample¹ Take

$$F = \begin{pmatrix} 1 & 0 \\ c & 10 \end{pmatrix}, G = \begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

$$P_0 = \begin{pmatrix} a & 0 \\ 0 & a \end{pmatrix}$$

and Q any positive definite matrix. Then, by Theorem 4.1, P_∞ is stabilizing. But

$$F - G(G^T P_0 G + rI)^{-1} G^T P_0 F = \begin{pmatrix} \frac{r}{a+r} & 0 \\ c & 10 \end{pmatrix},$$

which is always unstable whatever the values of a or r or $c \neq 0$. Further, a can always be chosen such that $P_0 > P_\infty$. *CQFD*

The above fallacious conjecture serves to provide a warning against simply demanding large P_0 to achieve stability. We next provide a fallacious conjecture which gives more positive support to finding stabilizing feedbacks.

Fallacious Conjecture 4.2

$F - G(G^T P_0 G + rI)^{-1} G^T P_0 F$ is stabilizing only if $\bar{Q}_0 > 0$.

¹Literally *exemple de comptoir*.

Counterexample Take, as in Fallacious Conjecture 3 of [BGPK85],

$$F = \begin{pmatrix} 0 & 2 \\ 2 & 0 \end{pmatrix}, G = \begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

$$P_0 = \begin{pmatrix} 16 & 0 \\ 0 & 16 \end{pmatrix},$$

for which

$$\bar{F}_0 = \begin{pmatrix} 0 & 2/17 \\ 2 & 0 \end{pmatrix},$$

which is stable, and

$$\bar{Q}_0 = \begin{pmatrix} -48 & 0 \\ 0 & 12.2 \end{pmatrix}.$$

CQFD

Hence, while $\bar{Q}_0 \geq Q \geq 0$ guarantees the closed loop stability for all $t \geq 0$ by Corollary 4.2, it is not a necessary condition for closed loop stability. We shall now move on to consider design methods which are capable of producing guaranteed stable closed loop systems in receding horizon control via the monotonicity of an associated RDE solution.

4.3 Stabilizing Feedback Strategies

The questions naturally arising from the monotonicity results of the previous section focus on how one could (or should) choose initial conditions, P_0 , for the RDE that satisfy $\bar{Q}_0 \geq Q$, in order to achieve a monotonic nonincreasing sequence of P_t and, thereby, closed loop stability whatever the choice of finite horizon for the receding horizon LQ strategy. We shall now address these issues and include in our discussion two earlier methods proposed for the design of constant stabilizing feedback gains, due to Kleinman [Kle74] and to Kwon and Pearson [KP78].

Clearly, from Theorem 4.1, if closed loop asymptotic stability is desired by choice of constant state feedback gain, then this may be achieved by constructing an infinite horizon LQ problem satisfying the stabilizability and detectability conditions and then solving an ARE for \bar{P} . In earlier days this route was considered computationally prohibitive compared with iterating the RDE and applying a receding horizon controller. This is still in part true, especially for adaptive applications, and motivates the consideration of these issues now. Before being able to proceed easily, we firstly develop some necessary technical machinery.

4.3.1 Alternative Forms of the RDE

It is apparent from the convergence theory of Theorem 4.2, and the monotonicity results of Theorem 4.4, that any attempt to achieve stability properties of the solution P_t of the RDE via monotonic nonincreasing behavior needs to commence with P_0 greater than \bar{P} . Indeed, one might be tempted to suggest that arbitrarily large P_0 could always yield the desired stability. This need not be true, as has been studied in [BGPK85] and treated in Fallacious Conjecture 1. However, there does exist an intelligent choice for large P_0 which we explain shortly. This involves the choice of effectively infinite P_0 as implemented by specifying

$$P_0^{-1} = W_0 = 0, \quad (4.10)$$

and then iterating an equation for $P_t^{-1} = W_t$, which is also an RDE. This might appear to contradict Fallacious Conjecture 1, but it will be seen that it can be roughly interpreted as insisting upon a finer structure of an infinite initial condition. We begin by simply deriving this alternative RDE for W_t .

Rewrite the RDE iteration (4.3) as the following coupled equations:

$$P_{t+1}^* = P_t - P_t G (G^T P_t G + R)^{-1} G^T P_t \quad (4.11)$$

$$P_{t+1} = F^T P_{t+1}^* F + Q \quad (4.12)$$

for P_t and P_t^* . Then, presuming for the moment that F , P_t , P_t^* and R are invertible, denote

$$W_t = P_t^{-1}, \quad W_t^* = P_t^{*-1}. \quad (4.13)$$

This yields, via the matrix inversion lemma, the recursions

$$P_{t+1}^{*-1} = P_t^{-1} + G R^{-1} G^T \quad (4.14)$$

$$P_{t+1}^{-1} = \left[F^T (P_{t+1}^* + F^{-T} Q F^{-1}) F \right]^{-1} \quad (4.15)$$

$$= F^{-1} \left[P_{t+1}^* + F^{-T} Q F^{-1} \right]^{-1} F^{-T} \quad (4.16)$$

$$= F^{-1} P_{t+1}^{*-1} F^{-T} - F^{-1} P_{t+1}^{*-1} F^{-T} Q^{1/2}$$

$$\times [I + Q^{1/2} F^{-1} P_{t+1}^{*-1} F^{-T} Q^{1/2}]^{-1} Q^{1/2} F^{-1} P_{t+1}^{*-1} F^{-T} \quad (4.17)$$

or, in terms of W_t ,

$$\begin{aligned} W_t &= F^{-1} W_t^* F^{-T} - F^{-1} W_t^* F^{-T} Q^{1/2} \\ &\quad \times [I + Q^{1/2} F^{-1} W_t^* F^{-T} Q^{1/2}]^{-1} Q^{1/2} F^{-1} W_t^* F^{-T}, \end{aligned} \quad (4.18)$$

or

$$W_{t+1}^* = W_t + GR^{-1}G^T \quad (4.19)$$

$$\begin{aligned} &= F^{-1}W_t^*F^{-T} - F^{-1}W_t^*F^{-T}Q^{1/2} \\ &\times [I + Q^{1/2}F^{-1}W_t^*F^{-T}Q^{1/2}]^{-1}Q^{1/2}F^{-1}W_t^*F^{-T} + GR^{-1}G^T. \end{aligned} \quad (4.20)$$

That is, W_t^* satisfies an RDE similar to (4.3) with the assignments

$$F \sim F^{-T}, \quad G \sim F^{-T}Q^{1/2}, \quad R \sim I, \quad Q \sim GR^{-1}G^T. \quad (4.21)$$

The closed loop matrix \bar{F}_t of (4.7) may be written as

$$\begin{aligned} \bar{F}_t &= F - G(G^T P_t G + R)^{-1}G^T P_t F \\ &= F - G(G^T W_t^{-1}G + R)^{-1}G^T W_t^{-1}F \\ &= F - GR^{-1}W_{t+1}^{*-1}F. \end{aligned} \quad (4.22)$$

These alternative forms of writing the RDE and the solution to the LQ optimal control problem are probably more familiar as duals of optimal filtering equations, where the equivalent of W_t^* arises as the inverse of the Kalman filter covariance in the information filter formulation [AM79]. The relationship between P_t and P_t^* is equivalent to that of the covariances of the Kalman predictor and of the Kalman filter respectively. In our LQ problem the important feature of these issues here is that they permit us to work alternatively with either scheme. It is worth remarking that it is possible to pursue this approach even in the case of singular F and/or R , but this might take us a little too far afield.

We shall see that taking zero initial conditions for the W_t^* equations yields a method of correctly using infinite initial conditions for P_0 , which entails monotonic nonincreasing P_t , and hence stability. We complete this subsection by including the following result on the rank properties of W_t when the RDE (4.18) is iterated from a zero initial condition.

Theorem 4.9

Consider the solution W_t of the RDE (4.18)–(4.20), under the assumption that F and R are full rank and with initial condition $W_0 = 0$. Then, provided $[F, G]$ is controllable (not just stabilizable), W_t will be full rank for all $t \geq n$, where n is the dimension of F .

Proof Suppose that W_n is not full rank. Then there exists a vector, x , such that $W_n x = 0$. Now consider the LQ optimal control problem associated with the RDE of the theorem. The criterion involves ‘ Q ’ and ‘ R ’ matrices

$GR^{-1}G^T$ and I respectively. Therefore, since $x^T W_n x$ represents the optimal cost and this is zero, the optimal control signal is zero and the state trajectory is $\{x, Fx, \dots, F^{n-1}x\}$, which in turn produces zero cost with the state weighting matrix $GR^{-1}G^T$. This clearly contradicts the controllability assumption on the pair $[F, G]$. *CQFD*

We are now in a position to investigate stabilizing feedback gain strategies using these results and the earlier monotonicity ideas.

4.3.2 The Stabilizing Controllers of Kwon, Pearson and Kleinman

Kleinman presented the following novel method for stabilizing a discrete-time linear, time-invariant system in state-variable form in [Kle74] in 1974. This was a discrete-time counterpart to his continuous-time method of 1970 [Kle70]. The method is based on Gramian construction.

Theorem 4.10

Consider the state-variable system of dimension n

$$x_{t+1} = Fx_t + Gu_t$$

and assume that $[F, G]$ is controllable. Then, for any $N \geq n$,

$$L = -R^{-1}G^T[(F^T)^N S_N^{-1}F^N]F \quad (4.23)$$

$$= -(R + G^T V_N G)^{-1}G^T V_N F \quad (4.24)$$

with

$$S_N = \sum_{i=0}^N F^i G R^{-1} G^T (F^T)^i \quad (4.25)$$

$$V_N = (F^T)^N S_{N-1}^{-1} F^N \quad (4.26)$$

yields the matrix $F + GL$ with all its eigenvalues strictly inside the unit circle.

We note that the structure of L above leads us to suspect the existence of an underlying interpretation of this controller as an LQ solution.

Indeed, if F is invertible, we see directly from (4.23) and (4.25) that the gain L may be rewritten as

$$L = -R^{-1}G^T W_N^*{}^{-1}F, \quad (4.27)$$

where

$$W_N^* = \sum_{i=0}^N F^{-i} G R^{-1} G^T (F^T)^{-i}. \quad (4.28)$$

From (4.28) it is apparent that W_N^* is obtained by iterating N times the following (Lyapunov) equation with zero initial conditions,

$$W_{t+1}^* = F^{-1}W_t^*(F^T)^{-1} + GR^{-1}G^T. \quad (4.29)$$

Further, comparing (4.20) to (4.29), we see that this is identical to the RDE for W_t^* with $Q = 0$. That is to say, the Kleinman feedback gain is identical to that derived from an LQ optimal control problem with $Q = 0$ and $W_0^* = 0$. We now note that this W_t^* from (4.28) is guaranteed to be monotonically nondecreasing, which in turn, appealing to Theorem 4.9 and Corollary 7.7.4 of [HJ85], implies that the equivalent P_t^* and P_t sequences are well defined and monotonically nonincreasing for all $t \geq n$. Thus the stability of this controller follows from Theorem 4.7. If F is not invertible then [KP78] shows that the stability still holds.

While these results of Kleinman are clearly interesting in that they provide a mechanism for achieving stability, they do not yield in any sense a solution which approximates that of a legitimate LQ problem involving specific Q and R . Further, if the Kleinman solution is to be used as the initial condition for iteration of the RDE with Q , stability of this latter solution is not guaranteed. Indeed, because the Kleinman solution is Q -independent, there is no surety that $P_{Kleinman} \geq P_\infty^Q$, where this latter quantity is the steady-state stabilizing solution with $Q \neq 0$. For an extension which does achieve monotonicity we turn to Kwon and Pearson of 1978 [KP78].

Theorem 4.11

Consider the state-variable system above of dimension n and assume that $R > 0$, $Q \geq 0$, F is invertible and $[F, G]$ is controllable. Further consider W_t^ the solution of the RDE (4.20) with initial condition $W_0^* = 0$. Then*

$$K_t = -R^{-1}G^T W_{t+1}^{*-1} F \quad (4.30)$$

$$= -(R + G^T W_t^{-1} G)^{-1} G^T W_t^{-1} F \quad (4.31)$$

yields $F + GK_t$ with all its eigenvalues strictly inside the unit circle for all $t \geq n$.

Proof Starting from $W_0^* = 0$ and iterating the RDE for one step we have $W_1^* = GR^{-1}G^T$ and hence $W_0^* \leq W_1^*$. Thus the W_t^* sequence is monotonically nondecreasing and W_t will be full rank for all $t \geq n$ by Theorem 4.9. We may then appeal directly to Theorem 4.7 for the stability result as in the previous proof. CQFD

We see that this feedback strategy both subsumes the Kleinman result and, further, identifies a mechanism for initializing a specific finite horizon

LQ optimal control problem in such a fashion as to ensure asymptotic stability if implemented as a receding horizon law. The choice of initial condition $W_0^* = 0$ may be interpreted as the effective method for choosing P_0 as infinity in order to achieve guaranteed monotonic decrease.

4.4 Mid Chapter Conclusion

So far, so good.

4.5 Comparative Performance of LQ Schemes

One of the major attractions of LQ methods, and of GPC in particular, is that the optimality criterion presents a direct connection between problem specification and achieved controller performance. The weighting of state deviations and control energies may be relatively easily conceptualized in terms of achieving, for example, more sluggish control dynamics by choosing higher control cost or more vigorous closed loop behavior by increasing Q . It is the relative magnitudes of Q and R which dictate asymptotic performance, and guidelines for their prospective selection may be easily developed.

Our aim in this section is to present a single result on the achieved performance of LQ control schemes (as opposed to the designed performance) and then apply this result to receding horizon designs. That is, we shall examine by how much the performance, measured by the achieved quadratic cost with the designer's choice of Q and R , deteriorates when the corresponding optimal control law is replaced by a suboptimal one obtained from the optimization of a quadratic criterion with a different Q and R . This is exactly the situation that arises in receding horizon control. Since receding horizon control, as opposed to finite horizon control, is a design based on long term or stationary application, the obvious way to evaluate the performance of a receding horizon controller designed with a Q and R weighting, is to compute the actual (i.e. achieved) infinite horizon cost criterion with the same Q and R . Now, we have observed in Section 4.3 that the strategy to enforce asymptotic stability of the closed loop system, when receding horizon LQ controllers with designed Q and R are used, is to force the controller to be the optimal solution of an associated infinite horizon optimal control problem with the same R and a larger Q . Therefore, it is important to evaluate by how much the replacement of the designer's choice of Q by a larger \bar{Q} deteriorates the performance of the receding horizon controller.

We begin by defining three successive quadratic performance measures

associated with the single state equation (3.1)

$$x_{t+1} = Fx_t + Gu_t,$$

and with a single infinite horizon quadratic criterion having constant Q and R

$$J(Q, R) = \lim_{N \rightarrow \infty} \frac{1}{N} \left(\sum_{t=0}^{N-1} \{x_t^T Q x_t + u_t^T R u_t\} \right). \quad (4.32)$$

Definition 1

The optimal control performance J_{opt} is the minimal value of $J(Q, R)$,

$$J_{opt} = \min_{\mathbf{u}} J(Q, R).$$

From our earlier theory, this $J(Q, R)$ is achieved by applying the LQ optimal feedback control derived by solving the ARE with matrices F, G, Q, R .

Definition 2

The designed control performance J_{des} is the optimal value of the LQ criterion (4.32) associated with a design problem using a possibly different weighting matrix pair \bar{Q}, \bar{R} ,

$$J_{des} = \min_{\mathbf{u}} J(\bar{Q}, \bar{R}).$$

We denote the corresponding control sequence by $u_{des,t}$,

$$\mathbf{u}_{des} = \arg \min_{\mathbf{u}} J(\bar{Q}, \bar{R}).$$

Again, our earlier theory says that this J_{des} is computable by solving an ARE with matrices F, G, \bar{Q} and \bar{R} .

Definition 3

The achieved control performance J_{ach} is the value of $J(Q, R)$ computed when the control law designed with matrices \bar{Q} and \bar{R} is applied to the above state-variable system,

$$J_{ach} = \lim_{N \rightarrow \infty} \frac{1}{N} \left(\sum_{t=0}^{N-1} \{x_t^T Q x_t + u_{des,t}^T R u_{des,t}\} \right).$$

Thus J_{ach} is the achieved performance, as measured by the $J(Q, R)$ criterion, of an optimal controller designed with differing weighting matrices \bar{Q} and \bar{R} .

We have the following result from [PPGB86]:

Lemma 4.3

Consider two AREs (4.1) with the same F , G and R matrices but with differing Q matrices, Q_1 and Q_2 respectively. Denote their respective stabilizing solutions by P_1 and P_2 . Then

$$Q_1 \geq Q_2 \quad \text{implies} \quad P_1 \geq P_2.$$

This is easily derived from the Riccati equation comparison formulae used earlier but, more surprisingly, it has the following (earlier!) extension due to Nishimura [Nis67] in the case of optimal filtering, the dual of which we now present.

Theorem 4.12

With the above definitions of performance measures, we have

$$\bar{Q} \geq Q, \bar{R} \geq R \quad \text{implies} \quad J_{des} \geq J_{ach} \geq J_{opt}.$$

The first inequality, $J_{des} \geq J_{ach}$, is apparent from the optimality criterion (4.32) itself, since u_t and x_t are identical in the calculation of J_{des} and J_{ach} but the weighting matrices are different. The second inequality stems trivially from the optimality. Some further discussion of this result for optimal filtering is included in [AM79].

We remark here that a complementary result involving $\bar{Q} \leq Q$ is not possible because of a lack of ordering between J_{des} and J_{ach} . Further, if \bar{Q} is allowed to become negative (presaging a difficulty later on), even stability need not be achieved, in spite of the earlier Fallacious Conjecture 2.

An important consequence of Theorem 4.12 for receding horizon LQ problems can be derived if we consider the following facts. Even if a finite time receding horizon LQ problem is posed, with Q and R the matrices in the criterion, the choice of these weighting matrices is usually made on long term (infinite time or asymptotic) considerations, so that the user effectively has an infinite time problem in mind. Hence the optimal performance he thinks of is given by J_{opt} as defined previously. The FARE (4.4) tells us that solving the finite time receding horizon problem with Q and R matrices amounts to solving an infinite time LQ problem with matrices \bar{Q}_{N-1} and R , for which the performance could be characterized by

$$J_{des} = x_t^T P_{N-1}^{Q,R} x_t = x_t^T P_{\infty}^{\bar{Q}_{N-1},R} x_t.$$

Still, the achieved performance in the long term, as measured by the initial choice of weighting matrices, will be J_{ach} . Now, if one uses such strategies that guarantee that $\bar{Q}_{N-1} \geq Q$, one obtains

$$J_{opt} \leq J_{ach} \leq J_{des},$$

and so one achieves guaranteed stability (by the FARE) and an upper bound for the performance which is given by J_{des} . Further, as N tends to ∞ we squeeze the upper bound on performance onto the lower bound. If a receding horizon strategy is used which does not imply that $\bar{Q}_{N-1} \geq Q$, then neither stability nor performance can be certified. We reiterate here that the predictability of the closed loop behavior is one of the advantages of infinite horizon LQ methods and thus it is important to preserve it in any approximating scheme such as receding horizon.

4.6 Stability Properties of GPC

We have now run much of our course in setting the scene for a full LQ re-assessment of the nonadaptive GPC control law based on receding horizon control. Our path has led from Chapters 2 and 3 with the interpretation of GPC as an explicit receding horizon LQ law implemented with specific direct feedthrough observers, through to this chapter where stability and performance properties of general receding horizon strategies were established together with methods for assuring these properties via monotonicity arguments. It next behoves us to examine the GPC control law in the light of its satisfaction or dissatisfaction of these criteria. Recall that Theorem 4.7 states that a receding horizon strategy constructed with monotonically non-increasing P_t will yield an asymptotically stable closed loop. We have the following contrary result for GPC.

Theorem 4.13

Consider the solution, P_t , of the RDE (3.4) associated with the GPC control law having $N_u = N_2$, $N_1 = 1$, i.e. $Q_t = H^T H$, $R_t = \lambda I$ and $P_0 = H^T H$. Then, for all $t \geq 0$,

$$P_{t+1} \geq P_t,$$

i.e. P_t is always monotonically nondecreasing.

Proof The RDE associated with this choice of N_u , N_2 and N_1 is a constant coefficient RDE,

$$P_{t+1} = F^T P_t F - F^T P_t G (G^T P_t G + \lambda I)^{-1} G^T P_t F + H^T H. \quad (4.33)$$

The solution of this RDE from initial condition $P_0 = H^T H$ is identical to that produced from initial condition $P_{-1} = 0$, by inspection. Since this entails $P_0 \geq P_{-1}$, we have from Theorem 4.5 that P_t is monotonically non-decreasing for all time. CQFD

The import of this theorem is that it dictates that the GPC formulation with full control horizon, $N_u = N_2$, cannot produce monotonically decreasing P_t , whatever the value of N (N_2) chosen. This should then make us somewhat suspicious about GPC's ability always to produce stabilizing control laws. We do remark here that while GPC does not produce an explicit calculation for P_t , the efforts of Chapter 3 have shown that it is implicitly defined by P_{N-1} (see (3.97)). The properties of the solution of the RDE (4.33) at $t = N - 1$ therefore directly determine the stability and performance of the GPC control law. In addition to potential closed loop instability, it also follows from Theorem 4.13 that GPC fails to inherit the nice performance properties discussed in the previous section. In fact, with monotonicity going the wrong way, it is hard to say anything about the achieved performance of the receding horizon GPC controller as measured against its optimal infinite horizon cost criterion except as N goes to infinity. We now provide an example which illustrates controller misbehavior in GPC.

4.6.1 An Unstable GPC Example

Consider the second order system with transfer function polynomials given by

$$\begin{aligned} A(q^{-1}) &= 1 - 4q^{-1} + 4q^{-2} = (1 - 2q^{-1})^2 \\ B(q^{-1}) &= q^{-1} - 1.999q^{-2}. \end{aligned}$$

It has a state-variable description as follows:

$$\begin{aligned} F &= \begin{pmatrix} 2 & 0 \\ 1 & 2 \end{pmatrix} \\ G &= (1 \ 0)^T \\ H &= (1 \ 0.001). \end{aligned}$$

Note that this system is unstable with open loop poles both at -2. Further, it possesses a near pole-zero cancellation, which translates into this state-variable model being nearly undetectable.

We select as control cost weighting factor $\lambda = 0.1$ and compute the solution of the GPC RDE (4.33) from $P_0 = H^T H$. We also compute the closed loop pole positions, λ_{cl} , resulting from the corresponding GPC control

law. These are listed below for several representative values.

$$P_0 = \begin{pmatrix} 1 & 0.001 \\ 0.001 & 0.000001 \end{pmatrix}, \quad \lambda_{cl} = 0.1819, \quad 1.9990$$

$$P_1 = \begin{pmatrix} 1.3640 & 0.00136 \\ 0.00136 & 0.00000136 \end{pmatrix}, \quad \lambda_{cl} = 0.1367, \quad 1.9990$$

$$P_2 = \begin{pmatrix} 1.3731 & 0.00137 \\ 0.00137 & 0.00000137 \end{pmatrix}, \quad \lambda_{cl} = 0.1358, \quad 1.9990$$

$$P_{10} = \begin{pmatrix} 1.3732 & 0.00137 \\ 0.00137 & 0.00000140 \end{pmatrix}, \quad \lambda_{cl} = 0.1358, \quad 1.9990$$

$$P_{20} = \begin{pmatrix} 1.3813 & 0.01644 \\ 0.01644 & 0.02809 \end{pmatrix}, \quad \lambda_{cl} = 0.1358, \quad 1.9881$$

$$P_{25} = \begin{pmatrix} 4.2486 & 5.3616 \\ 5.3616 & 9.9924 \end{pmatrix}, \quad \lambda_{cl} = 0.1358, \quad 0.6772$$

$$P_\infty = \begin{pmatrix} 5.7870 & 8.2294 \\ 8.2294 & 15.3385 \end{pmatrix}, \quad \lambda_{cl} = 0.1358, \quad 0.5003.$$

It is clear from this example that the convergence speed of the solution of the RDE is very slow from an initial condition $P_0 = H^T H$ when the pair $[F, H]$ is almost nondetectable. The first stabilizing controller is achieved only after 25 iterations, i.e. $N_2 = 25$. Convergence effectively to the steady-state value P_∞ has occurred by the thirtieth iteration. This is illustrated in Figure 4.1, which shows the achieved closed loop poles with a receding horizon controller as the horizon is altered. The rather remarkable amount of time before stability is achieved is due to the almost nondetectability of the plant model.

This example illustrates several things. Firstly, it clearly shows the monotonically nondecreasing properties of the GPC RDE solution, at least when $N_u = N_2$. Secondly, it shows that monotonically nondecreasing solutions of the RDE do not necessarily provide closed loop stability. And thirdly, it shows the potential difficulty arising from the choice of initial condition matrix P_0 as $H^T H$ when detectability problems exist. For this (albeit contrived) example it is clear that the application of a receding horizon strategy based upon GPC would be unwise.

There are several ways to solve this dilemma. Probably the simplest way is to just utilize an infinite horizon LQG controller by solving an ARE. This provides guaranteed closed loop stability; the closed loop eigenvalues for




Figure 4.1: Closed loop poles as a function of horizon

this example are 0.1358 and 0.5003, as shown above. Alternatively, one can utilize a receding horizon LQ controller by, for example, selecting a different Q and/or P_0 matrix. Both solutions of course require more computations than GPC, but with the speed of modern day computers this is rarely a decisive handicap. For comparison, here we carry through the computation of the Kwon-Pearson controller, which yields a stable closed loop after two steps, i.e. for any $N_2 \geq 2$ as opposed to $N_2 \geq 25$ for GPC.

For the calculation of the Kwon-Pearson controller for the above system, we have, from (4.20) with zero initial condition and Theorem 4.11,

$$\begin{aligned}
 F^{-1} &= \begin{pmatrix} 0.5 & 0 \\ -0.25 & 0.5 \end{pmatrix} \\
 W_3^* &= \begin{pmatrix} 10.73 & -0.389 \\ -0.389 & 0.249 \end{pmatrix} \\
 K_2 &= (-0.1085 \quad -0.7710) \\
 F + GK_2 &= \begin{pmatrix} 0.392 & -0.771 \\ -0.25 & 0.5 \end{pmatrix},
 \end{aligned}$$

yielding closed loop eigenvalues $\lambda_{cl} = 0.0034, 0.8881$. The point to remark here is that closed loop stability is guaranteed by the latter controller without further selection of design parameters or iteration on, say, N_2 .

4.6.2 Time-varying Strategies in Receding Horizon Control

The above results indicate some negative aspects of the GPC algorithm in that the P_t sequence will always be monotonically nondecreasing, which does not augur well for guaranteeing stability. The ‘ace up the sleeve’ of GPC, however, is that its formulation admits the use of time-varying weighting matrices in the derivation of its control law. In particular, the ploy of using several steps with $R = \infty$ is frequently asserted in GPC by selecting N_u appropriately. This is what we mean by ‘time-varying strategies’ in receding horizon control. Although originally introduced to achieve computational efficiency, this tool might have the potential of alleviating some problems by accelerating the initial increase of P_t , as we shall now explore. We shall investigate what stability properties might be achieved by this devious act.

We study the possible effects of combining a number of initial steps of the RDE with infinite R value, followed by steps of the finite R RDE. Specifically, we are looking for indications of how this can effect stability properties of the final solution. When $R = \infty$ we have the following version of the RDE (4.3), which now has the form of a Lyapunov equation,

$$P_{t+1} = F^T P_t F + Q, \quad (4.34)$$

possessing the obvious solution,

$$P_t = (F^T)^t P_0 F^t + \sum_{i=0}^{t-1} (F^T)^i Q F^i. \quad (4.35)$$

For stable F the solution P_t above converges as $t \rightarrow \infty$. For unstable F , P_t does not converge. With infinite R (and Q finite), the control signal will always be zero and so questions of stabilizability are not sensible. Nevertheless, this may yield a useful procedure for developing an initial condition from which to iterate the RDE with finite R .

We have some elementary results on the solutions of such infinite R RDEs.

Theorem 4.14

Consider the infinite R RDE (4.34). If $0 \leq P_0 \leq Q$, then P_t is monotonically nondecreasing for all t .

Proof From the theorem conditions and using (4.35) we see immediately that $P_1 = F^T P_0 F + Q \geq Q \geq P_0$, and so P_t is therefore monotonically nondecreasing for all t by Theorem 4.5. *CQFD*

Note that this result is not necessarily true with strict increase, i.e. $P_{t+1} >$

P_t , as may be seen by taking $F = 0$ for example. Similarly, if $P_0 \geq Q$ it need not be true.

Theorem 4.15

Denote the solution of the RDE with infinite R and initial condition P_0^∞ by P_t^∞ , and denote the solution with finite R by P_t^R . Then if $P_0^\infty \geq P_0^R \geq 0$ we have

$$P_t^\infty \geq P_t^R.$$

Proof Subtracting the two RDEs (4.3) and (4.34) we have

$$P_{t+1}^\infty - P_{t+1}^R = F^T(P_t^\infty - P_t^R)F + F^T P_t^R G(G^T P_t^R G + R)^{-1} G^T P_t^R F. \quad (4.36)$$

The theorem statement follows from the non-negativity of the right hand side above. *CQFD*

A simple corollary of the above theorem is that, if P_∞^∞ exists (which it will if and only if F is stable), then $P_0^\infty \geq P_0^R$ implies $P_\infty^\infty \geq P_\infty^R$. Thus one might envisage that the choice of several steps of $R = \infty$ might yield a P_t value exceeding P_∞^R and thereby the possibility of producing a monotonically decreasing sequence of P_{t+k} from that point onwards. We have the following remarkable new result ('Funny Result 4A'), whose proof we omit in order not to appear too funny.

Theorem 4.16

With the same notation as above, provided $P_\infty^R \geq 0$ exists, $[F, G^T P_\infty^R]$ is observable, and either

- $[F, Q^{1/2}]$ is stabilizable, or
- $P_0^\infty > 0$,

then there exists a K such that for all $k > K$

$$P_k^\infty \geq P_\infty^R.$$

We shall next investigate how this observation plays a part in justifying the GPC use of N_u .

4.6.3 The Use of N_u in GPC Stability

Recall that the use of a N_u smaller than N_2 in GPC corresponds to initial RDE steps being implemented with $R = \infty$, followed by iterations with a finite R . It was shown in the previous subsection that these procedures yield

a solution matrix, P_t^∞ , which always exceeds the equivalent finite R solution, P_t^R . Additionally, subject to mild assumptions, P_t^∞ is greater than or equal to the steady-state value P_∞^R for all t exceeding a finite constant K . The strategy advocated in the use of these initial infinite weightings therefore is to allow a fixed number of steps to be taken with infinite R and then to switch to finite R iterations of the RDE, operating under the presumption that this latter phase of the RDE will be initialized from above the steady-state solution and therefore be likely to yield monotonically decreasing P_t . While this strategy certainly points in the right direction, it does not really solve the problem of closed loop stability. The difficulty, as with GPC in general, is that it is essentially impossible to ascertain how many steps of each kind are necessary to achieve the desired result. To be fair, we admit here that the application of such steps in GPC, which is a consequence of the choice $N_u < N_2$, is advanced mostly on computational grounds.

In addition to the ‘playing of games’ with the N_u parameter, GPC aficionados have proposed other schemes to attempt to guarantee closed loop stability via time-varying strategies within the horizon. This includes the manipulations of Gorez, Wertz and Zhou [GWZ87], who consider altering the value of N_1 to this effect.

4.6.4 Stability Theorems of Clarke and Mohtadi

Mohtadi and Clarke [MC86], [Moh87], have advanced three theorems for GPC stability by judicious choice of the parameter values involved. We shall present these here and attempt to interpret them in our terms and in terms of their efficacy of application. As we shall see, each relies on limiting properties of specific LQ solutions and these limiting cases need not really reflect desirable control goals.

Their first theorem is concerned with the case where $N_2 \rightarrow \infty$ which, in RDE terms, is subsumed under the auspices of Theorem 4.2 and Theorem 4.1.

Theorem 4.17 [MC86]

If the system (3.1) with $\dim F = n \times n$ is stabilizable and detectable, then the closed loop under GPC control is stable if:

- $N_u, N_2 \rightarrow \infty$ with $N_u = N_2$ and $\lambda > 0$, or
- $N_u, N_2 \rightarrow \infty$ with $N_u = N_2 - n + 1$ and $\lambda > 0$, provided there is no plant zero on the stability boundary.

As remarked above, the first part of this theorem is easily absorbed into our earlier theorems on the RDE convergence and ARE stability. The second

part, after some cursory contemplation, is subsumed similarly because both N_u and N_2 go to infinity and the boundary effects of $N_u \neq N_2$ play no part. The purpose of the requirement of no plant zeros on the stability boundary would appear to be specious. The reason for discussing this controller selection in [MC86] is because of its connection to earlier proposed predictive controllers of [Pet84].

The second theorem of [MC86] forces the GPC LQ controller to yield a deadbeat feedback strategy by selecting $N_1 = n$ (the state dimension), $N_2 \geq 2n - 1$, $N_u = n$ and $\lambda = 0$.

Theorem 4.18 [MC86]

The GPC controller results in a stable deadbeat system if

- $N_1 = n$, $N_2 \geq 2n - 1$, $N_u = n$ and $\lambda = 0$.

The strategy here is best seen by considering the time interval $(0, N_2)$. During the first part of the interval, $(0, n - 1)$, the LQ weights are $Q = 0$ and $R = 0$ with initial condition $P_0 = H^T H$. Thus any control action is cost-free during this period, and further state deviations are not penalized. During the second half of the interval, (n, N_2) , the control cost is infinite but the state cost is non-zero. The optimal (zero cost!) control is clearly a deadbeat strategy which zeroes the state over the first n steps. While this stable controller is derivable from within the GPC formalism, its use does beg the question of whether such a circuitous route to this feedback strategy has any advantage over an explicit deadbeat computation.

The third theorem deals with finite timescales for the GPC and so is the most interesting from the perspective of receding horizon LQ control.

Theorem 4.19 [MC86]

The closed loop under GPC is stable if

- *The state-variable model (3.1) is completely controllable and observable with state dimension n ,*
- $N_1 \geq n$, $N_u = N_2 - N_1 + 1 \geq n$, and
- $\lambda = \epsilon$ where ϵ is a sufficiently small positive constant.

We remark that, while the theorem does not rely upon limiting arguments with time indices, it still requires the selection of a ‘vanishingly small’ constant ϵ . Nevertheless, some analysis does indicate how stability is achieved and consideration of the above example will illustrate some remaining difficulties with its application.

Referral to the GPC/LQ specifications (3.94)–(3.96) for the problem outlined in the theorem shows that, in RDE terms, the solution is generated by iterating the RDE in two stages. From $P_0 = H^T H$ the RDE is iterated for N_u steps with $R = \infty I$ and $Q = H^T H$. Then $N_2 - N_u$ steps are performed with $R = \epsilon I$ and $Q = 0$. After the first stage we have

$$P_{N_u} = \sum_{i=0}^{N_u} (F^T)^i H^T H F^i, \quad (4.37)$$

as was illustrated in (4.35). By observability, this solution will have full rank provided $N_u \geq n - 1$. Then to implement the $Q = 0$ stages, we revert to the inverse formulation of the RDE in terms of W_t^* .

Specifically, we define $W_{N_u} = P_{N_u}^{-1}$ from (4.13) and then $W_{N_u+1}^* = W_{N_u} + GR^{-1}G^T$ from (4.19). The RDE (4.20) is subsequently iterated for $N_2 - N_u$ steps, but since $Q = 0$ (4.20) becomes a Lyapunov equation (4.29)

$$W_{t+1}^* = F^{-1}W_t^*(F^T)^{-1} + GR^{-1}G^T. \quad (4.38)$$

The solution to this equation is therefore directly calculable as

$$\begin{aligned} W_{N_2+1}^* &= F^{-m} \left\{ \sum_{i=0}^{N_u} (F^T)^i H^T H F^i \right\}^{-1} (F^T)^{-m} \\ &\quad + \epsilon^{-1} \sum_{i=0}^{m-1} F^{-i} G G^T (F^T)^{-i}, \end{aligned} \quad (4.39)$$

where $m = N_2 - N_u$.

It is relatively easy to calculate examples where this W_t^* is not monotonically nondecreasing at $t = N_2 + 1$, and so stability is not verifiable by this means here. Nevertheless, it is apparent from (4.39) that, for small enough ϵ , $W_{N_2+1}^*$ is dominated by the term (4.28) of the Kleinman controller. As $\epsilon \rightarrow 0$, while $W_{N_2+1}^* \not\rightarrow W_{Kleinman}$, because of the ϵ^{-1} factor, from (4.30) one sees that $K_{N_2+1} \rightarrow K_{Kleinman}$. The stability properties of this controller are achieved by the stability of the Kleinman controller and the continuity of the closed loop poles with respect to the state feedback gains. We may now return to the simple second order GPC example above to illustrate this mechanism for stability.

With the earlier example of GPC, which has system order $n = 2$, we take $N_1 = 2$, $N_2 = 4$, $N_u = 2$ and controller LQ weightings (Q, R) successively as $(H^T H, \infty I)$, $(H^T H, \infty I)$, $(0, \epsilon I)$, $(0, \epsilon I)$. A computer analysis using Pro-Matlab shows that, in order for closed loop stability to be achieved, it is necessary to take $\epsilon \leq 5.2 \times 10^{-13}$! Thus we see that, even with this example

which is not too badly detectable, one must take an infinitesimal value for ϵ to gain closed loop stability by this procedure. Further, the controller that will be implemented will be close to that of the Kleinman scheme, which does not bear a close resemblance in performance terms to a general LQ specification since it does not involve Q .

In summary it is apparent that the systematic techniques provided for the GPC method so far to deliver stability and/or closed loop performance are either asymptotic in nature or involve the careful selection of various constants in a way which would be hard to formalize in a general context connected with a realistic performance measure. Our recommendation later will be that, as stability for receding horizon LQ strategies can only be easily guaranteed with monotonicity arguments, and these in turn state that the receding horizon Riccati solution is a solution of a (Fake) algebraic Riccati equation, one might just as well solve the ARE to generate the LQ control law. This ensures both stability and connection to desirable performance measures for the designed system. In the next chapter, we will also see that stability robustness properties additionally are available.

4.7 Conclusions

The attractive feature of GPC controller design is that it yields access to the sophisticated properties of Linear Quadratic control without necessitating the complete expertise of the applications engineer in this subject. This is achieved by precluding some of the design choices or recouching them in terms of process parameters, i.e. by choosing $Q = H^T H$. The implication of this, as has been evidenced in this chapter, is that it can be difficult to guarantee both closed loop stability and performance. Indeed, while the assumption is that the GPC receding horizon strategy should approximate an infinite horizon LQ law with its guaranteed stability property, it is quite apparent that this need not be achieved. Further, attempts to coax GPC to behave more adequately appear to suffer from either being asymptotic in nature or perhaps dependent upon iterative selection of various parameters in a trial and error type mode.

Our thesis here is that LQ state-variable feedback design, incorporating either receding horizon methods and the RDE or, preferably, infinite horizon methods and solution of the ARE together with observer design, may be formulated in a similar fashion to GPC requiring the same level of expertise of the applications personnel. As a matter of fact, after spending the next two chapters analysing the robustness properties of infinite horizon LQG controllers in an adaptive control implementation, we shall make a case in

Chapter 7 for the use of a specific adaptive LQG controller whose design choices, while as simple as those of GPC, are theoretically justified by the combined theories of LQG robustness and of RLS identification.

In addition to guarantees of closed loop stability and performance, together with computable stability margins as we shall see in the next chapter, the LQ formalism admits considerably more flexibility through the manipulation of observer dynamics, quadratic cost weightings without artificial constraint, and frequency weighting if such modifications are feasible and desired. Of course, the price to pay for this added flexibility is in the cost of computation of the controller, which must be performed at every iteration of the adaptive control algorithm, but it is our belief that there are many instances where the computational constraint is not an active one. This is becoming increasingly more so with improved, cheaper computing facilities and given the noncritical dimensions of typical applications processes, where GPC has found a niche. Additionally, efficient algorithms are at present available for the computations required in the solution of LQ problems, such as the millstone Schur algorithm of Laub [Lau79b] for the solution of an ARE equation.

Having concluded our stability and performance analysis of receding horizon controllers, and established the merits, from a stability point of view, of full infinite horizon LQ designs, we shall from now on abandon receding horizon strategies and GPC, and devote our attention to the robustness properties of infinite horizon LQ and LQG controllers.

Chapter 5

Robust LQG Design — Features for Adaptive Control

5.1 Introduction

With this chapter we mark our first excursions into the field of Robust Linear Control Systems. We shall return to this subject in Chapter 7 to refine the issues associated specifically with LQG robustness in an adaptive context. Later, in Chapter 8, we develop both a more general linear robustness theory and its implications for less specific adaptive control laws. Thus, for the moment, we shall concentrate upon the detailed presentation of LQG robustness theory and leave in abeyance (rather briefly) the fuller treatment of robust linear control. Our reasoning behind this is to explore reasonably completely the connections between the current GPC adaptive control laws (viewed now as a subset of LQG) and the existing linear robustness theories for LQG. Our analysis will focus upon LQG controllers made robust with Loop Transfer Recovery. Thus our tack is to develop firstly a thorough and detailed treatment of controller robustness for LQG adaptive optimal control, and only then in Chapter 8 to wax lyrical about more esoteric and fanciful, broader principles. By this later stage, the reader should have accreted the requisite intellectual baggage and intricate specific experience to deal comfortably with the technical discussion which treats, at one fell swoop, a plethora of control methods.

5.1.1 A First Hint at the Adaptation/Robustness Interplay

Having established a solid theoretical justification for the use of infinite horizon LQG control laws, we now proceed to establish the properties that these controllers should possess in order to be utilized in an adaptive loop. The adaptive version of LQG control consists of the interconnection of a parameter identifier operating on the plant input and output signals and producing plant model parameter estimates, and a control design stage where this estimated model is used as the basis for the generation of the control law. In this framework, the LQG design needs to be performed frequently as the estimates are refined or updated.

One of the key features of adaptive control systems is that the incorporation of an on-line system identifier coupled to the control law design module of necessity involves an admission that the system modeling will, at times, be poor. This poor quality has a double origin: most often, the parametrized model can only at best approximate the actual plant and, in addition, the estimated parameters do not coincide with the best possible parameters within the chosen model structure, either because they have not converged or because the actual plant dynamics are varying slowly in time. As a consequence of this, it is critical that the control law chosen as the basis for the adaptive computation should be robust to certain modeling errors. The study of the ability of a control law to preserve closed loop stability and performance in the face of inexact plant knowledge is the province of *robustness*. Because of the inherent inaccuracy of the identified plant model in an adaptive context, robustness of the underlying linear control law is a critical ingredient of any adaptive controller.

In line with our definition of robustness explained in Section 1.3, this means that, should the adaptation be stayed at some nominal parameter value, the resulting linear controller should be capable of stabilizing a neighborhood of plants about this nominal value. Of central importance in the consideration of linear system robustness is the definition of an appropriate topology within which this neighborhood is described and hence to delineate the class of allowable perturbations to the plant for which the controller can still provide stability. In this fashion the controller robustness properties can be tied in to the demands placed upon the adaptive identifier, which determines the accuracy of the nominal plant model. Thus we see the emergence of the first real insight into the interplay necessary between the control law selection and the plant identifier design which is the central focus of this book, and which makes up the (so-called) ‘Adaptive Control Problem’.

It is with this objective in mind that we now move on to consider some existing theories which develop robustness results for LQ and LQG control.

These notions will be formulated in the frequency domain and will produce joint modeling and control requirements indicating the respective conditions necessary to achieve robust stability. To complement this in the adaptive context, recent frequency domain formulations of the Least Squares (LS) system identification criterion, due to Ljung [Lju87] and others, will be studied in Chapter 6 to determine how the closed loop signals generated by a control strategy will affect the properties of a model fitted by a concurrent LS identifier, and how the identifier filters will affect the inaccuracy in the identified model. We shall then be able in Chapter 7 to specify unified goals for the control and identification components of an adaptive LQG controller. In Chapter 8 we then extend and complement the specific LQG analysis to cope with more general control laws, but at the price of somewhat less focused conclusions.

5.1.2 A Guided Tour of LQG Robustness Theory

As with much of our preceding theory, our aim in this chapter will be to navigate a path through a body of existing literature on the subject of the robustness of feedback control and then to interpret this specifically in the light of how it impinges upon adaptive LQG methods. Many of the results to be stated here are not new but, equally, many of their derivations in discrete-time are not freely available. Throughout, our goal will be to identify what are the crucial properties affecting control robustness and, later, to investigate how these might be affected and manipulated in the design of adaptive LQG or adaptive predictive controllers. Our presentation of the robustness theory of LQG controllers will proceed along the following path.

- We shall begin by recalling some classical theory on the robustness of unity feedback control systems. The question posed in such theory is as follows: if a nominal plant model, in a unity feedback loop, produces a stable closed loop, how large is the neighborhood of plants around this nominal model that can also be stabilized in the same unity feedback loop? The main result is a criterion that exactly connects the amount of relative plant model error that is allowed at any frequency with the return difference margin, at that frequency, of the nominal plant. The return difference of a plant in a feedback loop will be defined in due time; the result roughly tells us that the return difference margin of the nominal plant model must be large where the relative plant model error is large in the passband of the plant. This is a specific robustness criterion of many possible criteria (some of which will be presented in Chapter 8) which is appropriate when the nature of the plant model

error centers on gain variations in the passband, rather than, say, a robustness criterion which revolves around roll-off errors or stopband phase for example. It is robustness to this class of uncertainty which is demonstrable for LQ and LQG systems.

- Turning then to full state feedback LQ controllers and to Kalman predictors, the designs of which are dual of one another, we shall observe that a key Linear Quadratic design equality, called the Return Difference Equality, provides us with a constant (i.e. frequency independent) and computable lower bound on the stability margin for an LQ controlled system or for a Kalman predictor. In view of its important rôle, we shall call this Return Difference Equality the EDR (for *Égalité de la Différence du Retour*) in order not to confuse it with the Riccati Difference Equation, RDE.
- We shall then observe that when an observer or Kalman filter is included in the loop, i.e. when full state feedback is replaced by state-estimate feedback, the stability margin guaranteed to the designer by the EDR in the case of LQ control, may evanesce in this situation of LQG control. This is not to say that LQG feedback laws have no robustness to unmodeled dynamics; it just points to the fact that the tool for providing us with robustness bounds in the case of LQ controllers, the EDR, is no longer there to produce guaranteed margins in the case of LQG controllers.
- Finally, we shall show that, in the case of LQG design, Loop Transfer Recovery (LTR) techniques can be called on to rescue the guaranteed margins offered by full state feedback LQ design, with the guarantees only holding for minimum phase, minimum delay systems. In continuous time LQG design, these techniques consist of choosing either the noise variances of the Kalman filter or the weighting matrices of the LQ control criterion in such a way that the open loop transfer function of the combined ‘plant model/LQG controller’ converges to the corresponding open loop transfer function of the Kalman predictor or LQ controller, thereby inheriting their EDR-derived stability margins. As we shall show, in the discrete-time case that is of interest to us in this book, this situation is slightly more complicated: the guaranteed stability margin of the Kalman predictor only (and not that of the LQ controller) can be recovered by a specific choice of control weighting function matrices.

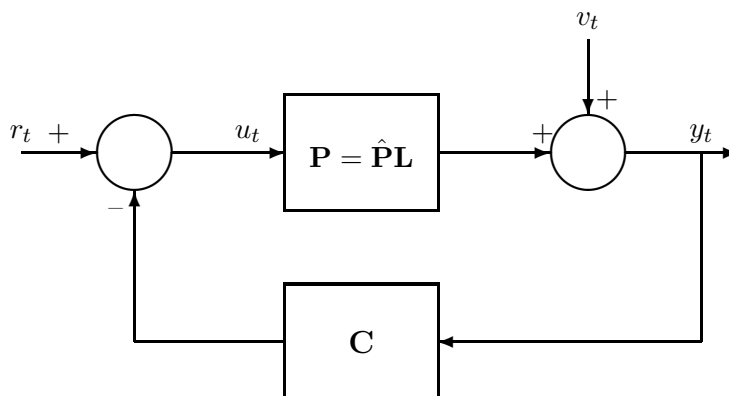


Figure 5.1: Feedback system

5.2 Robustness of Unity Feedback Systems

We begin by considering the underlying formulation of the unity feedback robustness problem à la Lehtomaki *et al.* [LSA81]. We reiterate that this represents just one of many possible robustness formulations, but it is arguably the most appropriate for LQG systems.

We suppose that our linear plant system is described by

$$y_t = P(z)u_t + v_t, \quad (5.1)$$

where y_t , u_t and v_t are the plant output, input and measurement noise signals, respectively, and $P(z)$ is the *true* or *actual* linear plant transfer function. Based on input–output measurements, however, we presume that we have an identified plant model or *nominal* plant transfer function $\hat{P}(z)$. For the moment we treat the case where $\hat{P}(z)$ is fixed and consider characterizations of circumstances under which stabilizing feedback controllers designed for $\hat{P}(z)$ maintain stability also for the actual plant $P(z)$.

Suppose that the linear controller is described by

$$u_t = -C(z)y_t + r_t, \quad (5.2)$$

where r_t is an external reference signal and $C(z)$ is the controller transfer function. The feedback system combining the plant (5.1) and the controller (5.2) can then be represented as in Figure 5.1.

Note that in an adaptive context such a controller would be designed based upon the nominal plant $\hat{P}(z)$ and applied to the actual plant $P(z)$.




Figure 5.2: Plant/controller cascade unity feedback system

We denote the cascaded plant/controller pairs as

$$\mathcal{G}(z) = C(z)P(z) \quad (5.3)$$

$$\hat{\mathcal{G}}(z) = C(z)\hat{P}(z). \quad (5.4)$$

Further, we presume that the nominal model, $\hat{P}(z)$, differs from the actual plant, $P(z)$, by a multiplicative perturbation, $L(z)$, i.e. we have

$$P(z) = \hat{P}(z)L(z) \quad (5.5)$$

and, hence,

$$\mathcal{G}(z) = \hat{\mathcal{G}}(z)L(z). \quad (5.6)$$

With the definitions made above, the closed loop system (5.1)–(5.2) can be depicted as a unity feedback system as shown in Figure 5.2, where $L(z)$ represents the multiplicative modeling error if the actual feedback system is considered, and $L(z) = I$ if the nominal closed loop system is considered. Thus $[1 + \mathcal{G}(z)]^{-1}\mathcal{G}(z)$ is the achieved closed loop transfer function, while $[1 + \hat{\mathcal{G}}(z)]^{-1}\hat{\mathcal{G}}(z)$ is the designed closed loop transfer function. The robust stability question is: under what conditions does stability of the designed closed loop imply stability of the achieved closed loop? The transfer function $[1 + \hat{\mathcal{G}}(z)]$ is called the *return difference* of the nominal model, and it will play an important rôle in how much plant model error is tolerated while still preserving closed loop stability.

It will prove to be convenient also to consider unity feedback configurations where the output signal is fed back into the loop, in which case $\mathcal{G}(z)$ represents the controller/plant cascade,

$$\mathcal{G}(z) = P(z)C(z) \quad (5.7)$$

$$\hat{\mathcal{G}}(z) = \hat{P}(z)C(z). \quad (5.8)$$




Figure 5.3: Controller/plant cascade unity feedback system

In such case, the multiplicative perturbation between $P(z)$ and $\hat{P}(z)$ will be assumed to take the following form

$$P(z) = L(z)\hat{P}(z) \quad (5.9)$$

and, similarly,

$$\mathcal{G}(z) = L(z)\hat{\mathcal{G}}(z). \quad (5.10)$$

With these definitions, the closed loop system can now be depicted as shown in Figure 5.3. Note that the closed loop stability properties do not depend in any way upon the ordering of the plant and the controller in these pictures.

Whatever the chosen configuration, we suppose that $\hat{\mathcal{G}}(z)$ has a state-variable description (recall that $P(z)$ is assumed strictly proper throughout this book):

$$\begin{aligned} \zeta_{t+1} &= A\zeta_t + B\nu_t, \\ \xi_t &= C\zeta_t. \end{aligned}$$

With $\nu_t = 0$ and the configuration of Figure 5.3, ν_t is the error signal, while ξ_t is the output, y_t , of the nominal model; with $\nu_t = 0$ and the configuration of Figure 5.2, ν_t is the input applied to the plant (or plant model), while ξ_t is the computed control signal that is fed back. The nominal open loop poles are then defined to be the zeros of $\hat{\phi}_{ol}(z)$, where

$$\hat{\phi}_{ol}(z) = \det(zI - A). \quad (5.11)$$

Note that, so far, no assumption of controllability or observability has been made, so the open loop poles therefore will include all poles canceled by open loop zeros as well.

When unity output feedback is applied to this system, we have

$$\nu_t = -\xi_t + r_t = -C\zeta_t + r_t,$$

so that the nominal closed loop has a state-variable description,

$$\begin{aligned}\zeta_{t+1} &= (A - BC)\zeta_t + Br_t \\ \xi_t &= C\zeta_t.\end{aligned}$$

Thus we define the nominal model closed loop poles to be the zeros of $\hat{\phi}_{cl}(z)$, where

$$\hat{\phi}_{cl}(z) = \det(zI - A + BC). \quad (5.12)$$

Again, this definition of closed loop poles will include those which might be canceled by the plant zeros.

Since the nominal model closed loop transfer function is given by $[I + \hat{\mathcal{G}}(z)]^{-1}\hat{\mathcal{G}}(z)$, we have the following characterization of the nominal feedback system's return difference, obtained by noting that the poles of $[I + \hat{\mathcal{G}}(z)]$ are precisely the poles of $\hat{\mathcal{G}}(z)$, and the zeros of $[I + \hat{\mathcal{G}}(z)]$ are the closed loop poles:

$$\det[I + \hat{\mathcal{G}}(z)] = \frac{\hat{\phi}_{cl}(z)}{\hat{\phi}_{ol}(z)} = \frac{\det(zI - A + BC)}{\det(zI - A)}. \quad (5.13)$$

Poles which are uncontrollable or unobservable in the state-variable description will not be moved by the output feedback and so will cancel between the polynomials $\hat{\phi}_{ol}(z)$ and $\hat{\phi}_{cl}(z)$.

The Nyquist stability criterion is familiar to most people only in the continuous-time case, where the Principle of the Argument is used with a contour encompassing all possible finite plant poles in the right half of the complex plane (see for example [FH77]). Thus, a semicircular 'D'-shaped contour is chosen with sufficiently large radius. In discrete time the criterion is actually easier to apply because the contour chosen is simply required to be the unit circle, suitably indented inwards around open loop poles on the circle. Denote this contour by Ω . Then we have:

Theorem 5.1

The nominal closed loop system (i.e. with $L(z) = I$) will be internally asymptotically stable, i.e. $\hat{\phi}_{cl}(z)$ will have no zeros outside or on the unit circle, if and only if the number of counter-clockwise encirclements of zero by $\det[I + \hat{\mathcal{G}}(z)]$, as z traverses the contour Ω in an anticlockwise sense, equals the number of zeros of $\hat{\phi}_{ol}(z)$ outside or on the unit circle.

Proof The Principle of the Argument.

CQFD

As is clear in the theorem statement, the above result applies only to the stability of the nominal system. However, following the line of argument of Lehtomaki *et al.* [LSA81] or Rosenbrock [Ros74], it is possible to extend this Nyquist stability result to describe the robust stability, i.e. $L(z) \neq I$, of the closed loop by insisting that the nominal and perturbed Nyquist diagrams not deviate too greatly in order that their encirclements of the origin remain fixed. We present the following result for the case where $\mathcal{G}(z) = L(z)\hat{\mathcal{G}}(z)$; an identical result holds of course when $\mathcal{G}(z) = \hat{\mathcal{G}}(z)L(z)$.

Theorem 5.2

Let $\mathcal{G}(z) = L(z)\hat{\mathcal{G}}(z)$ denote the perturbed system and let $\phi_{ol}(z)$ and $\phi_{cl}(z)$ denote its open loop and closed loop characteristic equations with unity feedback as above. Then $\phi_{cl}(z)$ will have no zeros outside or on the unit circle if:

1. $\phi_{ol}(z)$ and $\hat{\phi}_{ol}(z)$ have the same number of zeros outside the unit circle;
2. $\phi_{ol}(z)$ and $\hat{\phi}_{ol}(z)$ have the same unit circle zeros;
3. $\hat{\phi}_{cl}(z)$ has no zeros outside or on the unit circle;
4. $\det[I + (1 - \epsilon)\hat{\mathcal{G}}(z) + \epsilon\mathcal{G}(z)] \neq 0$ for all $z \in \Omega$ and for all $\epsilon \in [0, 1]$.

Proof Because $\mathcal{G}(z)$ and $\hat{\mathcal{G}}(z)$ are both finite for $z \in \Omega$, the Nyquist diagram of $\det[I + (1 - \epsilon)\hat{\mathcal{G}}(z) + \epsilon\mathcal{G}(z)]$ begins, for $\epsilon = 0$, with that of $\hat{\mathcal{G}}(z)$ and is deformed continuously into that of $\mathcal{G}(z)$, for $\epsilon = 1$. To increment or decrement the number of encirclements of the origin by this deformed diagram as ϵ changes requires that it pass either through zero or through infinity. The second condition of the theorem statement prevents the passage through infinity, while the fourth condition prevents zeros. The other two conditions force the stability of $\phi_{cl}(z)$ to be implied by the Nyquist diagram of $\phi_{ol}(z)$ having the same number of encirclements as that of $\hat{\phi}_{ol}(z)$. *CQFD*

In order to use this result concerning the deformation of the Nyquist diagram of the nominal model into that of the plant system, we need to introduce the following two lemmata of [LSA81], whose proofs we include for completeness.

Lemma 5.1

For a constant complex square matrix A , denote by $\bar{\sigma}(A)$ and $\sigma(A)$ the maximum and minimum singular values of A , i.e. $\lambda_{max}^{1/2}(A^H A)$ and $\lambda_{min}^{1/2}(A^H A)$ respectively.¹

¹ A^H denotes the complex conjugate transpose of the matrix A .

For constant square matrices G and L , with L invertible,

$$\det(I + GL) \neq 0$$

if

$$\bar{\sigma}(L^{-1} - I) < \underline{\sigma}(I + G). \quad (5.14)$$

Proof Write

$$I + GL = [(L^{-1} - I)(I + G)^{-1} + I](I + G)L,$$

and note that (5.14) implies that $\|(L^{-1} - I)(I + G)^{-1}\|_2 < 1$, which in turn implies that the term in brackets above is nonsingular. The nonsingularity of L and $(I + G)$ follows from the lemma conditions. *CQFD*

Lemma 5.2

If L is a square matrix and $P(\epsilon)$ is defined as

$$P(\epsilon) = (1 - \epsilon)I + \epsilon L,$$

for $\epsilon \in [0, 1]$, then

$$\bar{\sigma}(L^{-1} - I) < \alpha \leq 1 \quad \Rightarrow \quad \bar{\sigma}(P(\epsilon)^{-1} - I) < \alpha. \quad (5.15)$$

Proof Rewrite $\bar{\sigma}(P(\epsilon)^{-1} - I) < \alpha$ as

$$\begin{aligned} \alpha^2 P^H(\epsilon)P(\epsilon) - (P(\epsilon) - I)^H(P(\epsilon) - I) \\ = \epsilon^2[\alpha^2 L^H L - (L - I)^H(L - I)] \\ + \alpha^2(1 - \epsilon)[(1 - \epsilon)I + \epsilon(L + L^H)] > 0. \end{aligned}$$

By the assumption on $\bar{\sigma}(L^{-1} - I)$, the first term on the right hand side above is positive definite. Further, with $\alpha \leq 1$, this condition also yields

$$L + L^H > I + (1 - \alpha^2)L^H L > 0,$$

so that the second term on the right hand side is non-negative definite. *CQFD*

These two lemmata are worded for constant matrices but, by applying their conclusions at each point on the Nyquist contour Ω , we immediately have the following theorem:

Theorem 5.3

The closed loop characteristic polynomial $\phi_{cl}(z)$ of the actual system has no zeros outside or on the unit circle if:

- $\phi_{ol}(z)$ and $\hat{\phi}_{ol}(z)$ have the same number of zeros outside the unit circle;
- $\phi_{ol}(z)$ and $\hat{\phi}_{ol}(z)$ have the same unit circle zeros;
- $\hat{\phi}_{cl}(z)$ has no zeros outside or on the unit circle;
- $\bar{\sigma}(L^{-1}(z) - I) < \min(\alpha(z), 1)$ at each $z \in \Omega$, where

$$\alpha(z) \triangleq \underline{\sigma}(I + \hat{\mathcal{G}}(z)). \quad (5.16)$$

We make the technical observation here that, if one assumes that the unstable part of P and \hat{P} are identical, then L and L^{-1} will both be stable and the term $\min(\alpha(z), 1)$ can be replaced by just $\alpha(z)$. This will be revealed in Chapter 8.

The importance of this theorem is that it shows that a closed loop control system, as illustrated in Figure 5.1, will remain stable when the open loop plant/controller or controller/plant combination is perturbed multiplicatively, provided this perturbation is suitably small as measured by $\bar{\sigma}(L^{-1}(z) - I)$ and as compared to the nominal plant/controller or controller/plant return difference $(I + \hat{\mathcal{G}}(z))$. Further, since in the case of Figure 5.2, say,

$$\begin{aligned} L^{-1}(z) - I &= [\hat{\mathcal{G}}(z)L(z)]^{-1}[\hat{\mathcal{G}}(z) - \hat{\mathcal{G}}(z)L(z)] \\ &= \mathcal{G}^{-1}(z)[\hat{\mathcal{G}}(z) - \mathcal{G}(z)], \end{aligned} \quad (5.17)$$

we have the clear interpretation of Theorem 5.3 as demonstrating the robustness of feedback control systems to multiplicative perturbations provided the relative, or percentage, error of the frequency response of the plant/controller combination is kept small. As remarked earlier, this relative error is typically only easily bounded in the closed loop passband, making this criterion of primary use in ensuring stability when passband gain uncertainties exist. It further imposes conditions on the nominal controlled plant model that the quantity α remain sufficiently large, especially where the model is inaccurate. Since $\hat{\mathcal{G}}(z)$ is directly determined by the controller design and $L(z)$ can be shaped by the identification algorithm, as we shall see in Chapter 6, we see that robustness can be enhanced by choosing the controller $C(z)$ so that the nominal (i.e. designed) return difference has its minimum singular value large at the frequencies where the model is inaccurate, and by shaping the identification algorithm in such a way that the relative plant/model error is small where the designed $\alpha(z)$ is small. The designs of the controller and the identifier can therefore be made to reinforce one another so that the total robustness is larger than the sum of those delivered by two individual

designs. Finally, we note that for single input–single output systems the result of Theorem 5.3 states simply that the Nyquist diagram of the cascaded plant/controller should remain well clear of the minus one point.

We may now state a set of prioritized criteria for controller design in an adaptive control scheme.

- Stabilize the nominal closed loop system.
- Achieve adequate closed loop performance for the nominal plant with respect to tracking of references and disturbance rejection.
- Maximize robustness so that stability and performance are preserved in the face of modeling errors.

In operational terms these requirements are reflected in conditions on the chosen control law schema for the adaptive controller. The final point above indicates that, after satisfaction of the first two conditions, a control law which endeavors to keep the return difference frequency response large is preferred.

5.3 LQ and KF Robustness — Return Difference Equalities

5.3.1 The LQ Return Difference Equality

Our analysis of unity feedback systems has led us to conclude that, provided some conditions on zeros of the nominal and actual systems are satisfied, the controller based on the nominal plant will stabilize a class of neighboring plants if, at every frequency, the relative plant uncertainty is smaller than the smallest singular value, $\alpha(z)$, of the nominal return difference. It clearly follows that if one can choose a controller design that makes $\alpha(z)$ large, then one can allow for a large amount of plant model uncertainty, i.e. this controller will stabilize a large neighborhood of plants around the nominal one. We shall now show that full state feedback LQ controllers, as well as Kalman predictors, have the property of a guaranteed stability margin, namely the smallest singular value of their corresponding return difference is strictly positive everywhere around the unit circle, and it can actually be computed. This guaranteed stability margin of LQ controlled systems and of Kalman predictors can be shown to be endowed upon them by virtue of an important algebraic equation, called the Return Difference Equality (or EDR for short, for reasons explained earlier).

We commence here with a short development of the Return Difference Equality, which can be considered as the fundamental frequency domain equality of infinite horizon LQ optimal control, since many of the robustness results of LQ control stem from it. We subsequently present the dual EDR of the stationary Kalman predictor (KP). These equalities form the linchpin of the LQG robustness theory to follow. Specifically, we shall use these two equalities to derive global bounds on the particular LQ or KP return difference matrix, $\sigma(I + \mathcal{G}(z))$, appearing in the earlier general robustness analysis. We do, however, forewarn the good-natured reader that the next two subsections consist chiefly of unenjoyable algebraic manipulations.

Take the Algebraic Riccati Equation of LQ optimal control (3.8),

$$P = F^T P F - F^T P G (G^T P G + R_c)^{-1} G^T P F + Q_c, \quad (5.18)$$

and rewrite this in terms of the quantity

$$S = G^T P G + R_c, \quad (5.19)$$

and the feedback gain (see (3.9))

$$K = -(G^T P G + R_c)^{-1} G^T P F = -S^{-1} G^T P F, \quad (5.20)$$

yielding

$$P = F^T P F - K^T S K + Q_c.$$

From which description, we may establish the following identity directly by multiplying and equating like powers of z ,

$$Q_c - K^T S K = (z^{-1}I - F)^T P (zI - F) + (z^{-1}I - F)^T P F + F^T P (zI - F). \quad (5.21)$$

Thus,

$$\begin{aligned} (z^{-1}I - F)^{-T} Q_c (zI - F)^{-1} &= P + P F (zI - F)^{-1} \\ &\quad + (z^{-1}I - F)^{-T} F^T P + (z^{-1}I - F)^{-T} K^T S K (zI - F)^{-1}, \end{aligned}$$

and hence, adding R_c to both sides of the above multiplied on each side by G^T and G respectively,

$$\begin{aligned} &R_c + G^T (z^{-1}I - F)^{-T} Q_c (zI - F)^{-1} G \\ &= R_c + G^T \{P + P F (zI - F)^{-1} + (z^{-1}I - F)^{-T} F^T P \\ &\quad + (z^{-1}I - F)^{-T} K^T S K (zI - F)^{-1}\} G \\ &= (R_c + G^T P G) + G^T P F (zI - F)^{-1} G \\ &\quad + G^T (z^{-1}I - F)^{-T} F^T P G + G^T (z^{-1}I - F)^{-T} K^T S K (zI - F)^{-1} G \\ &= S - S K (zI - F)^{-1} G - G^T (z^{-1}I - F)^{-T} K^T S \\ &\quad + G^T (z^{-1}I - F)^{-T} K^T S K (zI - F)^{-1} G \\ &= [I - K (z^{-1}I - F)^{-1} G]^T S [I - K (zI - F)^{-1} G]. \end{aligned}$$

This latter equality is known as the *Return Difference Equality*,

$$\begin{aligned} R_c + G^T(z^{-1}I - F)^{-T}Q_c(zI - F)^{-1}G \\ = [I - K(z^{-1}I - F)^{-1}G]^T(G^T P G + R_c)[I - K(zI - F)^{-1}G]. \end{aligned} \quad (5.22)$$

Recall that the closed loop poles of our state-variable feedback system are given by the zeros of $\det(\lambda I - F - GK)$. In the case of full state LQ feedback, $I - K(zI - F)^{-1}G$ is the return difference $I + G(z)$. This follows directly from the state equation (3.73) and the control equation (3.74): see also (3.76). We shall call this return difference $I + \mathcal{G}_{LQ}(z)$. Thus,

$$\mathcal{G}_{LQ}(z) = -K(zI - F)^{-1}G. \quad (5.23)$$

The zeros of the return difference are related to the closed loop poles directly by the following formula:

Lemma 5.3

$$[I - K(zI - F)^{-1}G]^{-1} = I + K(zI - F - GK)^{-1}G. \quad (5.24)$$

Proof The Matrix Inversion Lemma [Kai80] states that

$$[A + BCD]^{-1} = A^{-1} - A^{-1}B[C^{-1} + DA^{-1}B]^{-1}DA^{-1}.$$

Here we take $A = I$, $B = -K$, $C = (zI - F)^{-1}$, and $D = G$ to yield the lemma statement. *CQFD*

It follows directly from (5.24) that the closed loop poles are the zeros of the return difference.

The importance of the Return Difference Equality has been recognized by many researchers, both in discrete-time and in continuous-time, using the corresponding version. Even though we shall be using it here to demonstrate that it leads to guaranteed and computable stability margins for LQ controllers, it is probably worth mentioning that the EDR can also be used for the computation of the optimal control gain; this offers an alternative to the solution of the ARE. Indeed, its expression demonstrates that the solution of the infinite horizon LQ control problem may be obtained by computing the left hand side of the EDR, since this depends only upon the open loop plant and the LQ weighting matrices, and then finding the minimum phase spectral factor of this term with the same poles as the open loop plant and with unity gain at $z = \infty$. Once this spectral factor has been computed, the feedback control gain matrix, K , may be explicitly retrieved from the return difference subject to controllability conditions. Lemma 5.3, together with

the stability properties of infinite horizon LQ control, Theorem 4.1, dictates that it is the stable zeros of the above factor that need to be included in the computation. By observation, the poles of the return difference are the same as those of the plant.

5.3.2 The KP Return Difference Equality

Before turning to the use of the EDR of LQ optimal control for the computation of stability margins, we now consider the dual stationary optimal estimation problem as in Chapter 3. Recall the filtering ARE (3.32),

$$\Sigma = F\Sigma F^T - F\Sigma H^T(H\Sigma H^T + R_o)^{-1}H\Sigma F^T + Q_o, \quad (5.25)$$

and the predictor gain M^P computation (3.33),

$$M^P = F\Sigma H^T(H\Sigma H^T + R_o)^{-1}. \quad (5.26)$$

By direct comparison with the EDR of LQ control, i.e. by duality, we have the dual EDR of optimal prediction,

$$\begin{aligned} R_o + H(zI - F)^{-1}Q_o(z^{-1}I - F)^{-T}H^T \\ = [I + H(zI - F)^{-1}M^P](H\Sigma H^T + R_o)[I + H(z^{-1}I - F)^{-1}M^P]^T. \end{aligned} \quad (5.27)$$

In the case of the Kalman predictor, $I + H(zI - F)^{-1}M^P$ is the return difference $I + \mathcal{G}(z)$. The corresponding open loop transfer function will therefore similarly be given a specific name,

$$\mathcal{G}_{KP}(z) = H(zI - F)^{-1}M^P. \quad (5.28)$$

The LQ and KP EDRs will play a central rôle in our subsequent theory from two points of view. The KP EDR will serve a pivotal task in the development of a robustness result for the Kalman predictor which we shall endeavor to carry over to the complete LQG controlled system through appeals to the recent theory of LQG Loop Transfer Recovery. Secondly, the LQ EDR will allow us to describe the spectral properties of the closed loop input signal to the plant operating in feedback with the LQG controller, a feature intimately coupled to the integrity of adaptive LQG control since it links the control law properties to those of the parallel identifier.

5.3.3 The EDRs and LQ, KP Robustness

Having defined the open loop transfer functions $\mathcal{G}_{LQ}(z)$ and $\mathcal{G}_{KP}(z)$ of the LQ controller and Kalman predictor problems via (5.23) and (5.28), respectively, we shall now appeal to their corresponding Return Difference Equalities to derive lower bounds on the quantity $\alpha(z)$ defined in (5.16). It will

then follow that unity feedback systems whose open loop transfer functions are equal to $\mathcal{G}_{LQ}(z)$ or $\mathcal{G}_{KP}(z)$ inherit the corresponding stability margins.

Recall that the LQ EDR is as follows:

$$\begin{aligned} R_c + G^T(z^{-1}I - F)^{-T}Q_c(zI - F)^{-1}G \\ = [I - K(z^{-1}I - F)^{-1}G]^T(G^T P G + R_c)[I - K(zI - F)^{-1}G]. \end{aligned}$$

Also, notice that with full state feedback LQ control and the identity (5.23), we have

$$[I + \mathcal{G}_{LQ}(z)] = [I - K(zI - F)^{-1}G]. \quad (5.29)$$

This leads us directly to the following:

Theorem 5.4

Consider the LQ control system with full state feedback above, for which we assume $[F, G]$ to be stabilizable, $[F, Q_c]$ to be detectable and R_c to be full rank, then one has the following bound on the minimum singular value of the return difference matrix. There exists a positive constant $\bar{\alpha} \leq 1$ such that

$$\sigma(I - K(zI - F)^{-1}G) \geq \bar{\alpha} > 0, \forall z \in \Omega. \quad (5.30)$$

Proof From EDR it is clear that, for $z \in \Omega$,

$$[I - K(z^{-1}I - F)^{-1}G]^T(G^T P G + R_c)[I - K(zI - F)^{-1}G] > R_c,$$

since the remaining term is a spectrum and hence non-negative definite on the unit circle. Further, from the theorem conditions and the work of de Souza, Gevers and Goodwin [dSGG86], a finite maximal positive definite solution, P , of the ARE (5.18) exists. Therefore there exists a positive scalar, β , such that $G^T P G + R_c < \beta I$, and so (5.30) holds with

$$\bar{\alpha} = \underline{\sigma}(R_c)/\beta.$$

We also note that necessarily $\bar{\alpha} \leq 1$.

CQFD

We note that, since the solution of the LQ problem is usually effected via the computation of the solution P of the ARE, the bound $\bar{\alpha}$ comes as an easy by-product. In the work of Shaked [Sha86] a more precise bound for $\bar{\alpha}$ is found which includes explicit bounds on P , G , F and Q_c . Our result here is quantitatively poorer, but qualitatively the same. Since our goal is to derive conceptual rules for adaptive predictive control design, we shall make do with this more conservative bound. In the continuous-time theory one does not have the term involving $G^T P G$ in the left hand side of

the EDR and so one immediately has that $\bar{\alpha} = 1$ in the scalar case, and an equivalent result in terms of the maximal and minimal eigenvalues of R in the multivariable case. Presumably, if one samples a continuous-time plant sufficiently fast, then $G^T P G + R \approx R$ as well.

The importance of Theorem 5.4 is that it shows that if a LQ controller is designed on the basis of a nominal (estimated) plant, $\hat{P}(z)$, then the corresponding nominal return difference, $I + \hat{G}(z)$, will have its smallest singular value lower bounded by $\bar{\alpha}$. Hence this LQ controller will be capable of stabilizing a neighborhood of plants with a robustness margin at least as large as $\bar{\alpha}$.

With this bound on the frequency response of the return difference, we may now appeal directly to Theorem 5.3 to produce the following result on LQ robustness.

Theorem 5.5

The LQ optimal controller designed on the basis of the nominal model will produce an actual closed loop system with denominator $\phi_{cl}(z)$ having no poles outside or on the unit circle provided:

- $\phi_{ol}(z)$ and $\hat{\phi}_{ol}(z)$ have the same number of zeros outside the unit circle;
- $\phi_{ol}(z)$ and $\hat{\phi}_{ol}(z)$ have the same unit circle zeros;
- $[F, G]$ is stabilizable, $[F, Q_c]$ is detectable and R_c has full rank;
- with $\gamma(z) \triangleq \bar{\sigma}(L^{-1}(z) - I)$, either of the following hold, with $\bar{\alpha}$ from Theorem 5.4,
 - $Q_c > 0$ and $\gamma(z) \leq \min(\bar{\alpha}, 1), \forall |z| = 1$;
 - $\hat{\phi}_{ol}(z) \neq 0$ and $\gamma(z) < \min(\bar{\alpha}, 1), \forall |z| = 1$.

We remark that the final complication in the last condition of the theorem statement is concerned simply with preventing unit circle poles of the closed loop. If Q_c is of full rank then the stability theorems of [dSGG86] will force the unit circle zeros of $\hat{\phi}_{ol}(z)$ to be moved inside the circle by an LQ control. The third condition of the theorem is included to guarantee the stability and robustness of the nominal LQ closed loop.

Taking the dual Kalman predictor return difference,

$$[I + \mathcal{G}_{KP}(z)] = [I + H(zI - F)^{-1}M^P], \tag{5.31}$$

one achieves the immediate dual result producing a stability margin for the Kalman predictor, using the dual Return Difference Equality (5.27).

Theorem 5.6

Consider the Kalman predictor above, for which we assume $[F, Q_o]$ stabilizable, $[F, H]$ detectable and R_o full rank. Then there exists a positive constant $\bar{\delta}$ such that

$$\sigma \left(I + H(zI - F)^{-1} M^P \right) \geq \bar{\delta} > 0, \forall z \in \Omega. \quad (5.32)$$

This particular theorem, rather than its LQ dual Theorem 5.4, will act as the hub of our robustness theory to be developed since, as we shall see, the KP robustness is more easily preserved for discrete-time systems than is the LQ counterpart. Naturally, a direct KP equivalent of Theorem 5.5 can immediately be stated.

Having now established that a full state feedback LQ controller or a Kalman predictor have return differences that have a guaranteed stability margin via the inequalities (5.30) and (5.32), respectively, we have now laid the foundations for the robust design of LQG systems. Indeed, the strategy will be to design either the controller or the Kalman filter to achieve closeness of the LQG return difference, $I + \hat{\mathcal{G}}(z)$, to that of either the LQ return difference $I + \hat{\mathcal{G}}_{LQ}(z)$ or the KP return difference $I + \hat{\mathcal{G}}_{KP}(z)$, which will *ipso facto* imply that they recover their closed loop stability robustness properties. The design procedure to achieve this recovery is the subject of Loop Transfer Recovery, which we examine now.

5.4 Robustness of LQG Control — Loop Transfer Recovery

The central general robustness conditions of Theorem 5.3, upon which we shall focus with our study of LQG, are

$$\sigma(I + \hat{\mathcal{G}}(z)) \stackrel{\Delta}{=} \alpha(z) > 0, \quad (5.33)$$

and

$$\bar{\sigma}(L^{-1}(z) - I) < \min(\alpha(z), 1), \quad (5.34)$$

where $\hat{\mathcal{G}}(z)$ is the cascaded plant/controller or controller/plant pair, so that $I + \hat{\mathcal{G}}(z)$ is the return difference, which we have identified earlier in (5.29) and (5.31), and $L(z)$ is the multiplicative plant perturbation, so that $L^{-1}(z) - I$ is the relative plant error between the nominal plant model and the actual plant as derived in (5.17). This latter condition (5.34) is a matter of achieving a system model with good percentage error, which problem is the subject of Chapters 6 and 7. The former condition is, however, a matter of controller design, which we now address.

5.4.1 Loop Transfer Recovery Rationale

The results of the preceding subsection deal primarily with the robust stability of an LQ controller or of the KP system, even though the robustness theory of Section 5.2 is demonstrably applicable to any feedback system capable of being written as a unity feedback structure, which includes LQG. It is just the feature of LQ control without observers or of the Kalman predictor without feedback control that the EDR leads immediately to the satisfaction of (5.33) for a particular α . The difficulties in extending the results to LQG arise in quantifying the properties of the return difference when an observer is included in the open loop transfer function, since then the $\hat{\mathcal{G}}(z)$, which we shall derive in the next section, differs from those above for which the EDR may be applied directly. Stated otherwise, no EDR is available to us when $\hat{\mathcal{G}}(z)$ is obtained from a LQG controller.

Doyle [Doy78] and Doyle and Stein [DS79] have analysed the effects of observers on the robustness of LQG designs in a continuous-time setting with the conclusion that, with the incorporation of an observer into the LQ feedback law, much or all robustness of the LQ design can be lost. In terms of the transfer functions above, that is to say that the robustness of the closed loop to perturbations $L(s)$, inherent in the LQ state-variable feedback design open loop matrix $\hat{\mathcal{G}}(s)$, need not be preserved when $\hat{\mathcal{G}}(s)$ is replaced by a form including an observer. This is not to say that LQG controllers have no robustness to unmodeled dynamics, but only that, unlike the full state feedback LQ design, there is no tool available, such as the EDR, to guarantee robustness margins.

The topic of Loop Transfer Recovery (LTR) is devoted to the formulation of Kalman observer design methodologies (by specific choices of R_o and Q_o) and/or state-variable feedback design methodologies (by specific choices of R_c and Q_c) which attempt to cause the open loop transfer LQG matrix $\hat{\mathcal{G}}(s)$ to approach either that of the LQ, $\hat{\mathcal{G}}_{LQ}(s)$, or that of the Kalman predictor design, $\hat{\mathcal{G}}_{KP}(s)$, and thereby to ‘recover’ the consequent degree of robustness associated with this latter system. A potential mechanism to achieve this effect was first identified by Kwakernaak [Kwa69] and later reformulated by Doyle [Doy78] as follows. The observer is designed by choosing the Kalman filter for the system with $Q_o = GG^T$ but with measurement noise covariance matrix, R_o , tending to zero. Thus, an asymptotically singular optimal filtering problem is posed. For minimum phase plant systems (i.e. those systems possessing neither finite nor excessive infinite zeros outside the unit circle), arbitrarily closed loop recovery is possible by designing a Kalman filter observer in this way [SA87], although the method achieves its robustness through the use of very large observer gains. This means that by

designing the observer this way and letting R_o tend to zero, the open loop LQG transfer matrix $\hat{\mathcal{G}}(s)$ will converge to the open loop LQ transfer matrix $\hat{\mathcal{G}}_{LQ}(s)$ obtained with full state feedback. Also, because in continuous time one does not draw a distinction between the Kalman predictor and the Kalman filter, the LQ recovery and the KF/KP recovery are completely dual and so LQ robustness recovery via observer design may be replaced by KP robustness recovery using LQ design. Naturally, the price of this robustness will be with the closed loop performance of the nominal system.

Having discussed something of the continuous-time LTR history and philosophy, we shall next turn to the development of discrete-time LTR results following the methodology of Maciejowski [Mac85]. To do this, we first need to derive the expressions of the open loop transfer functions $\hat{\mathcal{G}}(z)$ for the various situations of LQG design.

5.4.2 LQG Controller Transfer Functions

We shall combine the observer and control law equations (3.70) and (3.71) to derive complete expressions for the open loop transfer function $\hat{\mathcal{G}}(z)$ of the LQG controller in a unity feedback loop. In pursuing the derivation of this transfer function, we shall suppress the explicit appearance of the separate reference model and noise model for the sake of clarity. Indeed, the distinction between plant states and the states of the other models need only be drawn for explanatory reasons — the critical features of the super-state model being its stabilizability and detectability.

The dynamical equations of an observer/controller pair are encapsulated by

$$u_t = K\hat{x}_t, \quad (5.35)$$

for the control law and, for an observer with no direct feedthrough,

$$\hat{x}_{t+1} = (F - M^P H)\hat{x}_t + Gu_t + M^P y_t, \quad (5.36)$$

while, for an observer with direct feedthrough,

$$\begin{aligned} \hat{x}_{t+1} &= (F - M^F H F)\hat{x}_t + (G - M^F H G)u_t + M^F y_{t+1} \\ &= [F - M^F H F + (G - M^F H G)K]\hat{x}_t + M^F y_{t+1} \\ &= (I - M^F H)(F + GK)\hat{x}_t + M^F y_{t+1}. \end{aligned} \quad (5.37)$$

In order to identify the nature of the LQG controller transfer function vis-à-vis the LQ controller with full state feedback we consider the LQG controller in a unity feedback loop, and compute the forward path transfer functions for the plant/controller and controller/plant cascades. That

is, these cascade transfer functions appear as the forward path of a unity feedback description of the LQG controlled system. We first consider the plant/controller case, where the computed control signal $-K\hat{x}_t$ is fed back into the system, and we compute the transfer function from plant input u_t to the computed control $-K\hat{x}_t$. This is the situation depicted in Figure 5.2. We have the following lemma:

Lemma 5.4

Consider the transfer function, $\hat{G}(z)$, from system input, u_t , to computed control output, $z_t = -K\hat{x}_t$, arising in a unity feedback representation of the LQG controlled plant (see Figure 5.2). Then,

- with an observer possessing no direct feedthrough term

$$\hat{G}_{LQ}^{KP}(z) = -K(zI - F + M^P H - GK)^{-1} M^P H (zI - F)^{-1} G, \quad (5.38)$$

- with an observer possessing a direct feedthrough term

$$\begin{aligned} \hat{G}_{LQ}^{KF}(z) &= -K \\ &\times \left\{ I + (I - M^F H)(F + GK)[zI - (I - M^F H)(F + GK)]^{-1} \right\} \\ &\times M^F H (zI - F)^{-1} G. \end{aligned} \quad (5.39)$$

Proof The first part follows by substituting $u_t = K\hat{x}_t$ into the observer equation (5.36), since this is the value of u_t presumed by the observer. Then z_t is gained by equating $z_t = -K\hat{x}_t$.

The second part follows as the first, and is repeated here for completeness. From (5.37) we have

$$\begin{aligned} \hat{x}_{t+1} &= (F - M^F H F)\hat{x}_t + (G - M^F H G)u_t + M^F y_{t+1} \\ &= [F - M^F H F + (G - M^F H G)K]\hat{x}_t + M^F y_{t+1}. \end{aligned}$$

Therefore the transfer function from y_t to z_t (see Figure 5.2) is given by

$$\begin{aligned} K(z) &= -zK(zI - F + M^F H F - (G - M^F H G)K)^{-1} M^F \\ &= -zK(zI - (I - M^F H)(F + GK))^{-1} M^F \\ &= -K[I - z^{-1}(I - M^F H)(F + GK)]^{-1} M^F \\ &= -K[I + z^{-1}(I - M^F H)(F + GK) \\ &\quad + z^{-2}[(I - M^F H)(F + GK)]^2 + \dots] M^F \\ &= -K[I + (I - M^F H)(F + GK)[zI - (I - M^F H)(F + GK)]^{-1} M^F. \end{aligned}$$

CQFD

The point of this lemma is to compare the open loop transfer functions $\hat{\mathcal{G}}_{LQ}^{KP}(z)$ and $\hat{\mathcal{G}}_{LQ}^{KF}(z)$ of LQ control when an observer is used to reconstruct the state, i.e. LQG control is used, with the transfer function $\hat{\mathcal{G}}_{LQ}(z)$ of (5.23) when full state feedback is applied.

In a fashion completely dual to the above derivation, one may derive equivalent open loop transfer functions for a unity feedback LQG controller with the plant output y_t being preserved as the feedback signal. The resulting transfer functions are achieved simply by forcing the observer/controller block back around to the input of the loop. In this case it is the Kalman predictor (or observer) return difference which appears as an underlying object modified by the presence of the control law, rather than the control return difference of the previous case.

Lemma 5.5

Consider the transfer function, $\hat{\mathcal{G}}(z)$, from the error input signal, e_t , to the system output, y_t , arising in a unity feedback representation of the LQG controlled plant (see Figure 5.3). Then,

- with a Kalman predictor or observer possessing no direct feedthrough term

$$\hat{\mathcal{G}}_{KP}^{LQ}(z) = -H(zI - F)^{-1}GK(zI - F + M^P H - GK)^{-1}M^P, \quad (5.40)$$

- with a Kalman filter or observer possessing a direct feedthrough term

$$\begin{aligned} \hat{\mathcal{G}}_{KF}^{LQ}(z) &= -H(zI - F)^{-1}GK \\ &\times \left\{ I + (I - M^F H)(F + GK)[zI - (I - M^F H)(F + GK)]^{-1} \right\} \\ &\times M^F. \end{aligned} \quad (5.41)$$

This formulation follows directly as in the previous lemma or simply by recognizing that the explicit appearance of the output signal can be achieved by forcing the observer/controller part of the transfer function back around the loop from the previous case of input feedback. Again, the idea of this lemma is to illustrate the differences between the open loop transfer functions $\hat{\mathcal{G}}_{KP}^{LQ}(z)$ and $\hat{\mathcal{G}}_{KF}^{LQ}(z)$ when a LQ controller is inserted into the system and the open loop transfer function $\hat{\mathcal{G}}_{KP}(z)$ (see (5.28)) of the Kalman predictor without feedback control input, for which a guaranteed stability margin has been shown to exist.

At this stage we should ask what the distinction is between these transfer function calculations and those already presented in Section 3.8. The answer

is ‘not much really’, since they are the same except for the fact that the earlier results absorb the open loop plant transfer function into the control calculation to permit the algebraic cancellation of observer poles. Here in the robustness study, our goal is to examine the effects of incorrect open loop models upon the closed loop stability. Thus it is not permissible to regard y_t as precisely $H(zI - F)^{-1}Gu_t$ when H , F , G are the parameters of the nominal (or estimated) model on the basis of which all the stability margins are computed. The distinction is that the formulae of Section 3.8 are *closed loop* formulae incorporating cancellations, while those here are a computation of the forward path or open loop transfer functions of the controller/plant cascade which does not possess the cancellation unless the above substitution for y_t is made.

5.4.3 The Discrete-time LTR Theory of Maciejowski

We present the Loop Transfer Recovery design methodology as converted to the discrete-time situation by Maciejowski [Mac85]. We shall begin by considering the KP robustness recovery using LQ design, and comment upon the LQ version via KP or KF design later. The alert reader should be aware of the fact that, in a practical situation, including the adaptive situation to which the present robustness results will ultimately be applied, all the models $P(z)$ and open loop plant/controller transfer functions $\mathcal{G}(z)$ to be considered in this subsection are the nominal ones. However, since LTR theory is about recovering one transfer function from another by proper design, it does not matter whether the transfer functions are actual plant or model transfer functions, and we shall therefore not use *chapeaux* on top of these transfer functions.

We first recall that in discrete time, the ideal, guaranteed robust $\hat{\mathcal{G}}_{KP}(z)$ associated with the Kalman predictor is given by

$$\hat{\mathcal{G}}_{KP}(z) = H(zI - F)^{-1}M^P. \quad (5.42)$$

The achieved LQG closed loop $\hat{\mathcal{G}}_{KF}^{LQ}(z)$, however, with the Kalman filter and LQ control included, is

$$\begin{aligned} \hat{\mathcal{G}}_{KF}^{LQ}(z) &= -H(zI - F)^{-1}GK \\ &\times \left\{ I + (I - M^F H)(F + GK)[zI - (I - M^F H)(F + GK)]^{-1} \right\} \\ &\times M^F, \end{aligned} \quad (5.43)$$

where we recall that $M^P = FM^F$. The LTR paradigm is to design a feedback gain K such that $\hat{\mathcal{G}}_{KF}^{LQ}(z)$ approaches $\hat{\mathcal{G}}_{KP}(z)$ for z on the unit circle, and

thus the LQG system inherits the robustness of the Kalman predictor. This design proceeds by taking LQ control design with

$$Q_c = H^T H \quad (5.44)$$

$$R_c \rightarrow 0. \quad (5.45)$$

We start by considering the solution of the LQ problem with Q_c given by (5.44) and $R_c = 0$.

Lemma 5.6

Provided the open loop plant model, $P(z) = H(zI - F)^{-1}G$, satisfies

$$\det(HG) \neq 0, \quad (5.46)$$

the solution of the stationary LQ optimal control problem with weighting matrices,

$$Q_c = H^T H \quad (5.47)$$

$$R_c = 0, \quad (5.48)$$

has ARE solution

$$P = H^T H, \quad (5.49)$$

and control gain,

$$K = -(HG)^{-1}HF. \quad (5.50)$$

Proof With the parameters of the lemma statement one has the ARE,

$$P = F^T P F - F^T P G (G^T P G)^{-1} G^T P F + H^T H.$$

By direct substitution of $P = H^T H$ into the right hand side we have

$$\begin{aligned} & F^T P F - F^T P G (G^T P G)^{-1} G^T P F + H^T H \\ &= F^T H^T H F - F^T H^T H G (G^T H^T H G)^{-1} G^T H^T H F + H^T H \\ &= H^T H \\ &= P, \end{aligned}$$

where the invertibility of HG is used explicitly to reduce the expression. K may then be computed as

$$\begin{aligned} K &= -(G^T P G)^{-1} G^T P F \\ &= -(G^T H^T H G)^{-1} G^T H^T H F \\ &= -(HG)^{-1} HF. \quad \mathcal{CQFD} \end{aligned}$$

We have imposed no minimum phase assumption in Lemma 5.6, but it is of course well known that if $P(z)$ is non-minimum phase, the control gain (5.50) leads to internal instability. We now immediately apply this explicit solution of the singular LQ problem to demonstrate that full loop transfer recovery is achieved under a minimum phase condition.

Theorem 5.7

With the same conditions applying as in Lemma 5.6 we have the following equality

$$\hat{\mathcal{G}}_{KF}^{LQ}(z) = \hat{\mathcal{G}}_{KP}(z). \quad (5.51)$$

Thus provided the open loop plant has no zeros in $|z| \geq 1$ and is minimum delay, i.e. $\det(HG) \neq 0$, we have full loop transfer recovery with internal stability by selecting $Q_c = H^T H$ and $R_c = 0$.

Proof Denote by Π the matrix $G(HG)^{-1}H$. Then, clearly, one has $H(I - \Pi) = 0$. Further, $F + GK = (I - \Pi)F$, and so the closed loop system matrix is in the null space of H , i.e. $H(F + GK) = 0$. Now write the controller/observer transfer function of the KF/LQ pair as

$$zK[zI - (I - M^F H)(F + GK)]^{-1}M^F = zK(zI - F - GK)^{-1}M^F.$$

Next we use the relations $GK = -\Pi F$ and $M^P = FM^F$ and rewrite

$$\begin{aligned} & \hat{\mathcal{G}}_{KF}^{LQ}(z) - \hat{\mathcal{G}}_{KP}(z) \\ &= H(zI - F)^{-1}[-zGK(zI - F - GK)^{-1}M^F - M^P] \\ &= H(zI - F)^{-1}[z\Pi F(zI - F + \Pi F)^{-1} - F]M^F \\ &= H(zI - F)^{-1}\{z\Pi[zI - F(I - \Pi)]^{-1} - I\}FM^F \\ &= H(zI - F)^{-1}(F - zI)(I - \Pi)[zI - F(I - \Pi)]^{-1}FM^F \\ &= -H(I - \Pi)[zI - F(I - \Pi)]^{-1}FM^F \\ &= 0. \end{aligned}$$

This establishes (5.51). However, notice from the expressions (5.42) and (5.43) that the poles of $\hat{\mathcal{G}}_{KP}(z)$ occur only at the open loop poles of the plant $P(z)$. Therefore, by default, the poles of the controller/observer must have been canceled by the plant zeros through this particular selection of the LQ control law. Thus, although the identity (5.51) still pertains, for plants $P(z)$ possessing zeros outside $|z| < 1$ it does so by unstable pole/zero cancellation, which violates internal stability. *CQFD*

We now offer some remarks concerning this Loop Transfer Recovery result.

- The LTR property of Theorem 5.7 is effected by placing the controller poles at the open loop plant zeros — hence the need for stable plant zeros. If one performs a singular optimal control, not by setting $R_c = 0$ but by taking $R_c \rightarrow 0$, then the controller poles tend to the stable plant zeros plus the stable reflections in the unit circle of those zeros outside the circle. To see this, observe that the controller/observer poles, which are guaranteed stable for any LQG designed system, are the poles of $zK(zI - F - GK)^{-1}M^F$. Now, by Lemma 5.3, these poles are the zeros of the return difference $I - K(zI - F)^{-1}G$. Since R_c is close to zero, it then follows from the EDR (5.22) that the stable factor of the right hand side of this EDR will contain the stable zeros of the plant model and the reflections of the unstable ones. Thus exact or full LTR with $R_c = 0$ relies upon the minimum phase property of the plant. To preserve internal stability with non-minimum phase plants, one must take $R_c \rightarrow 0$. Only for minimum phase systems does this solution converge to that for $R_c = 0$.
- Maciejowski [Mac85] makes the point that for many discrete-time systems the non-minimum phase zeros due to sampling (as opposed to those due to a non-minimum phase continuous-time system) lie close to the negative real axis and, therefore, have an effect on the design which falls outside the control bandwidth. He makes the point that, while LTR is only guaranteed for minimum phase, minimum delay plants, it frequently works well for other systems — the limiting feature being how small one may take the value of R_c before striking too serious a performance deterioration. We should note that the effected control strategy is still LQG but with specific choices of Q_c and R_c . The LTR component is simply a robustness enhancement feature.
- One might suppose that, by some universal duality theorem, the LTR of an LQ controller via singular Kalman filtering should be possible. This is not the case because, by observation, the LQ control of a strictly proper linear system is dual to the Kalman predictor design for that system. The Kalman filter is not dual to LQ control, as may be seen by observing the noninterchangeable rôles of GK and $M^F H$ in the LQ/KF transfer function. As remarked earlier, this is a peculiarity of the discrete-time case because in continuous time the distinction between KP and KF evanesces. Thus LTR is only achievable with this class of plants via KF design followed by LQ recovery. Since the designer is frequently more concerned with output regulation than control input properties, this need not be such a restriction.

- The condition $\det(HG) \neq 0$ is a requirement that the open loop plant be minimum delay. The effect of violating this condition can be seen from (5.42) and (5.43). The ideal KP transfer function will in general possess a unit delay and, by construction, the LQ controller is causal, hence should the plant $P(z)$ possess more than a unit delay then it is impossible to recover fully the ideal transfer function. If sample-and-hold devices are used in developing a discrete plant from a continuous one then $\det(HG) \neq 0$ is generic, except for systems with a significant deadtime. In this latter case, again a limit to achievable recovery exists and one is limited by both the achievable level of recovery and by the undesirable signal values introduced into the control.
- Similarly to the above remark, the use of the Kalman predictor in place of the Kalman filter causes the LQG controller to become strictly causal and hence $\hat{G}(z)$ to be the product of two strictly proper systems, again making the complete recovery impossible. In other words, $\hat{G}_{KP}(z)$ cannot be recovered from $\hat{G}_{KP}^{LQ}(z)$, whatever the LQ control design.
- Ishihara and Takeda [IT86] have investigated what type of return difference is recoverable from LQG systems if one incorporates a delay into the control law. They use the results of Mita [Mit85] concerning the LQ optimal control law with controller delay and the Kalman predictor to show that if a singular KP is used with the LQ delay control law then one recovers the LQ delay return difference. Similarly, if singular LQ with delay is applied to the KP, then the KP with controller delay return difference is recovered. These results are technically interesting and indicate a duality not present in the theory above. The connection between these delay controllers and achievable robustness has also been investigated [Ish88].
- It is frequently the case that the plant and achievable control objectives are known to be bandlimited to a region $\omega < \mathcal{W}_s$ beyond which modeling and attempts to move the closed loop gain crossover are unlikely to succeed. This information may be taken into account in determining the control law and the model class to support robustness. We shall discuss this further in later chapters.
- The robustness ‘recovered’ by the LTR methodology is precisely that robustness of the Kalman predictor without control input. The theory does not state that this robustness is necessarily better or worse than that of any particular LQG design incorporating this KP or KF. It is

only the guaranteed minimum level of robustness which is assured by this procedure.

- The cost of robustness achieved by this (and probably most other methods yielding controllers of a similar complexity) is that the closed loop performance may be degraded in the pursuit of robust stability in the face of modeling errors. We know that for the nominal system the combination of LQ control plus KF yields the optimal feedback controller according to the LQ criterion. The LTR design, however, deliberately detunes the LQ part of the controller away from the ‘natural’ values of Q_c and R_c in order to achieve robustness. Thus the optimality of LQG is lost. Clearly a compromise between robustness and performance is needed.
- In adaptive control one usually proceeds from input–output data to a plant model and thence to controller design. One of the favorable features of LQG/LTR design is that, although it involves a state-variable description, the resulting solution for the LQ control law is state-variable coordinate basis independent, i.e. is an input–output description, because Q_c is chosen to be $H^T H$. Notice how this intersects with the GPC control law specification interpreted as LQG (3.93)–(3.96).
- With small values of R_c taken in the LQG/LTR formulation, we see several undesirable effects with non-minimum phase plants. Firstly, the LTR is not ensured completely. Secondly, the control signals in this case can become excessive as the weighting is reduced. Thirdly and more subtly, as the control weighting R_c is reduced, the controller gain increases and the closed loop bandwidth increases. This increases the likelihood of meeting stability problems because of the failure of the robustness conditions in this larger bandwidth. We shall see examples of these issues later in this chapter and in Chapter 7.

Having developed the LTR theory for LQG robustness in discrete time à la Maciejowski, we shall next perform a brief computational example to reinforce some of these issues. In an adaptive control context, where the underlying system model is not precisely known and, indeed, is expected to be changing with time, one must adopt tactics cautiously. We shall return to the design issues for LQG/LTR controllers in the adaptive context in Chapter 7, where the design variables specified will allow a simple supervision of the controller design stage.

5.5 An LQG/LTR Example

We reconsider the two systems of Section 3.7 — a simple second order minimum phase plant and a more difficult non-minimum phase third order plant, the working example.

A Simple Example

We treat again the system described in state-space form by the following matrices:

$$\begin{aligned} F &= \begin{pmatrix} -0.5 & 1 \\ -0.5 & 0 \end{pmatrix} \\ G &= \begin{pmatrix} 0.9 \\ -0.6 \end{pmatrix} \\ H &= (1 \ 0), \end{aligned}$$

and expose it to various LQG and LTR treatments to demonstrate that which has been established by theory above. Note that this plant is minimum phase, minimum delay and stable, and so is a prime candidate for guaranteed robustness recovery with LTR.

We performed several LQG designs with constant $R_o = \rho = 1$, $Q_o = GG^T$, $Q_c = H^T H$ and varying values of $R_c = \lambda$. Additionally, the transfer function $\hat{\mathcal{G}}_{KF}^{LQ} = \hat{P}C$ was computed for the Kalman predictor with this plant and this Q_o , R_o . The data are plotted as discrete Nyquist diagrams with the dotted curve representing the uncontrolled KP function $\hat{\mathcal{G}}_{KP}$.

Figure 5.4 illustrates several features. Firstly, the dotted curve of $\hat{\mathcal{G}}_{KP}$ displays the maintenance of a fixed distance from the point -1. This establishes the implicit KP robustness discussed in Section 5.3. The $\hat{\mathcal{G}}_{KF}^{LQ} = \hat{P}C$ of the LQG system, however, is considerably further from minus one essentially at all frequencies. The implication here is that the LQG system is rather more robust to multiplicative passband modeling uncertainties than its KP equivalent.

Figure 5.5 shows that, with decreasing value of λ (here 0.01), the two transfer functions approach. With $\lambda = 10^{-5}$ (see Figure 5.6) the functions effectively coincide. As we saw in Chapter 3, the choice of such small values of λ with this plant does not necessarily lead to excessive control gains but does tend to produce good tracking properties. Therefore we see that good tracking performance here is associated with diminished robustness, although a particular LQG performance criterion need not reflect a lesser robustness than its LTR version.




Figure 5.4: $\hat{\mathcal{G}}_{KF}^{LQ} = \hat{P}C$ Nyquist plot with $\lambda = 1$ versus $\hat{\mathcal{G}}_{KP}$




Figure 5.5: $\hat{\mathcal{G}}_{KF}^{LQ} = \hat{P}C$ Nyquist plot with $\lambda = 0.01$ versus $\hat{\mathcal{G}}_{KP}$




Figure 5.6: $\hat{\mathcal{G}}_{KF}^{LQ} = \hat{P}C$ Nyquist plot with $\lambda = 10^{-5}$ versus $\hat{\mathcal{G}}_{KP}$

A More Difficult Plant

Next we conduct a similar experiment for the more difficult plant, the ‘Working Example’:

$$P(z) = \frac{-0.05359 z^{-1} + 0.5775 z^{-2} + 0.5188 z^{-3}}{1 - 0.6543 z^{-1} + 0.5013 z^{-2} - 0.2865 z^{-3}}. \quad (5.52)$$

This is non-minimum phase and third order and was shown in Section 3.7 to demonstrate control difficulties with LQ control having light control weighting, R_c .

Figure 5.7 shows the Nyquist plots of the Kalman predictor loop transfer function, $\hat{\mathcal{G}}_{KP}$ (dotted curve), and that of the LQG system designed with $\lambda = 1$. Two features are immediately apparent. Neither system’s loop function tends particularly closely to minus one over the entire range, nor do the two curves remain close. For this value of λ one could argue that the KP system is less robust than the LQG system.

Figure 5.8 contains the Nyquist plot with $\lambda = 0.01$, and Figure 5.9 shows the same plots for $\lambda = 10^{-5}$. Further reduction of λ does not result in any further approach of the two curves. Nevertheless, it is apparent that the two curves do lie considerably closer for small λ values than for large values. It is also evident from Figure 5.9 that the loop function of the LQG system




Figure 5.7: $\hat{\mathcal{G}}_{KF}^{LQ} = \hat{P}C$ Nyquist plot with $\lambda = 1$ for the Working Example versus $\hat{\mathcal{G}}_{KP}$




Figure 5.8: $\hat{\mathcal{G}}_{KF}^{LQ} = \hat{P}C$ Nyquist plot with $\lambda = 0.01$ for the Working Example versus $\hat{\mathcal{G}}_{KP}$




Figure 5.9: $\hat{G}_{KF}^{LQ} = \hat{P}C$ Nyquist plot with $\lambda = 10^{-5}$ for the Working Example versus \hat{G}_{KP}

actually lies further from the -1 point than its KP counterpart and so is potentially the more robust system. What is masked by this figure is that its concomitant control signal in, say, its step response will be very large.

These two examples serve to display some of the features of linear robustness via LQG/LTR. Further, since robustness is concerned with vaguely specified uncertainties, we also see that this theory is not really very useful for predicting the detailed properties of individual systems. This reinforces the notion that robustness comes at a price in terms of performance and that robustness *per se* is not the major control objective. Methodologies such as LTR should really be regarded as ancillary features for the enhancement of basic designs.

5.6 Conclusion

The aim in LTR robust controller design is to produce an open loop transfer function $\hat{G}(z)$ which satisfies the requirements of maintaining $\sigma(I + \hat{G}(e^{i\omega}))$ sufficiently positive. That this may be achieved by LQ control with full state feedback is evidenced by Theorem 5.5. The potential loss of LQ robustness with the inclusion of state estimators is demonstrated by the examples of

Doyle and others [Doy78]. However, it still may be possible to generate a robust feedback law using observers since the loss of margins is not a generic property. Equally, control laws other than LQ can be contemplated in robust applications. The point to be drawn here is that the automatic design of robust feedback control laws, i.e. without operator intervention and iteration, is difficult but it is possible to specify some simple LQG controller formulations in which this design can take place without involving too many tunable parameters. This will be the subject of Chapter 7, where we incorporate the robust control law design with the parameter identification methods to be described in the next chapter. For the moment, though, we move on to consider the identification component of the Adaptive Controller.

Chapter 6

Recursive Least Squares Identification in Adaptive Control

6.1 Introduction

In most publications on adaptive control, the effort is spent either on the control design with only scant attention paid to the properties of the parameter estimator, or on the parameter estimator with little consideration given to the way the regulator affects the parameter estimator or the transfer function estimates. And yet, because the identification is being performed in closed loop (that is, there is explicit dependence of the plant input signal upon the plant output and the identification is performed on the plant *in situ* in a closed loop, not that the plant is identified from input and output measurement of the closed loop), and because the parameters of the regulator are being continuously adapted, there is an intimate interplay between the mechanisms of parameter adaptation, the properties of the estimated models and the stability of the closed loop system. (This was demonstrated in part by our Gedankenexample of Chapter 1.) It is our purpose in this chapter to unveil part of this intimacy and, without excessive voyeurism, to display the bare essentials about the interconnection between the stability properties of the closed loop system, the spectral properties of the reference signals and of the noises, and the prefilters.

Our discussion in this chapter is entirely focused on the nonexact modeling case, i.e. the case where the parametrized model can at best approximate the true plant. This situation, where the true system is not in the model set, has been given a variety of exotic names in the control literature, such as

mismodeling, nonexact modeling, restricted complexity model, plant-model mismatch, neglected dynamics or — with reference to the French ‘dynamique négligée’ — dynamical nightie. In Chapter 5 we have shown that a control design that stabilizes a plant model will also stabilize the true plant provided the relative error between the true plant and the model is bounded, at each frequency, by a quantity that depends on both the plant and the controller. In order to apply fully this robustness theory in an adaptive framework, it is therefore important to be able to characterize the plant-model mismatch over the frequency range when the model arises as the result of an identification algorithm. This is the object of the present chapter.

The first part of our analysis will be nonadaptive. To set the stage and reveal our good schooling in identification theory, we start with a very brief review of off-line prediction error identification with a Least Squares criterion. We then use Ljung’s [Lju87] frequency domain expression for the asymptotic bias distribution to examine the effect of the input signal, the noise and the filters on the asymptotic plant-model mismatch in the case of open loop identification. This expression is then extended to the somewhat more difficult case where the identification is performed in closed loop with a fixed (i.e. constant) regulator. The influence of the regulator transfer functions and of the reference signal spectrum will then come to light.

When the plant-model parameters are adjusted on-line, and when the regulator parameters are simultaneously adjusted as a function of these plant-model parameters according to a regulator design criterion, this changing feedback law will influence the signals in the loop and therefore the convergence point, if any, of our closed loop parameter adaptation law. Precautions must then be taken to guarantee that the parameter vector estimate converges to a point that is close enough to the minimizing argument of a correctly frequency weighted off-line prediction error criterion and for which closed loop stability can still be ensured. Recall the central robust stability condition (5.14) for linear systems,

$$\bar{\sigma}(L^{-1} - I) < \underline{\sigma}(I + \hat{\mathcal{G}}).$$

In adaptive control, the nominal plant model fit, and therefore the frequency domain characteristics of its error, is determined by an identifier, while the robustness of the control design for this nominal system is determined by the control law schema. Thus the identification plays the part of ensuring that the left hand side above is sufficiently small, while the control law design is concerned with maintaining the right hand side sufficiently large.

A proper way to analyse this situation involving both the identifier and controller effects is to study the global nonlinear system made up of all

the dynamical equations of the closed loop control system, together with the parameter adaptation equations, and to study the asymptotic behavior of this complete nonlinear system. The powerful tools of averaging theory and of time scale separation can be used to show that, provided certain precautions are taken, the solution of the parameter update equation will converge to a point that is close to the asymptotic solution of the off-line parameter estimation criterion. These precautions take the form of slow adaptation, persistency of excitation, restriction on the minimal achievable modeling error, low sensitivity of the regressors to the parameter vector, amongst other features.

In Section 2 we shall give an exceedingly brief introduction to prediction error identification, the aim being mainly to introduce notations and to set the stage for the presentation of the bias formulae of Section 3. These formulae give a precise description of how the plant-model mismatch is distributed over frequency when there is such a mismatch. They will also tell us how to influence the model obtained through the minimization of an off-line prediction error criterion by a proper choice of data filters. Section 4 then examines the behavior of a recursive identification algorithm operating in an adaptive feedback loop using integral manifold theory. This is a local analysis which presumes that the parameter is adapted slowly and that the algorithm is initialized with a parameter value that is not too far removed from an optimal one. In Section 5 we shall briefly examine the basic ingredients that need to be utilized when the algorithm is extended for global convergence.

6.2 Prediction Error Identification

6.2.1 Off-line Prediction Error Identification: a Refresher

We consider an input–output description of the systems and models described in Chapters 2 and 3. We assume that there is a true system given by

$$y_t = P(z)u_t + v_t, \quad (6.1)$$

where $P(z)$ is a strictly proper rational transfer function matrix, u_t is the control input and v_t is an unmeasurable disturbance acting on the output y_t . We shall assume that $\{u_t\}$ and $\{v_t\}$ are quasistationary: see Ljung [Lju87]. A signal $\{z_t\}$ is called quasistationary if the following conditions hold for all t, s ,

$$|\mathbb{E}z_t| \leq C, \quad (6.2)$$

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{\tau=t}^{t+N} z_\tau z_{\tau-s} \triangleq R_z(s), \quad |R_z(s)| \leq C, \quad (6.3)$$

where C is a finite constant and E denotes expectation with respect to all the stochastic components in the signals. We then define the spectrum of $\{z_t\}$ as

$$\Phi_z(\omega) = \sum_{s=-\infty}^{\infty} R_z(s) e^{-j\omega s}. \quad (6.4)$$

For the system (6.1) we assume that the signals $\{y_t\}$, $\{u_t\}$ and $\{v_t\}$ possess spectra $\Phi_y(\omega)$, $\Phi_u(\omega)$ and $\Phi_v(\omega)$, respectively. Thus, for the moment, we presume that stability of $P(z)$ holds.

We also consider a parametrized model set

$$\{\hat{P}(z, \theta), \hat{H}(z, \theta), \theta \in \mathcal{D}_\theta\}. \quad (6.5)$$

A particular model in that model set will be described by

$$y_t = \hat{P}(z, \theta)u_t + \hat{H}(z, \theta)q_t, \quad (6.6)$$

for a particular value of θ , where θ in R^d is a parameter vector, $\{q_t\}$ is white noise and $\hat{P}(z, \theta)$ and $\hat{H}(z, \theta)$ are, respectively, strictly proper and proper rational transfer function matrices, with $\hat{H}(z, \theta)$ and $\hat{H}^{-1}(z, \theta)$ having all their poles strictly inside the unit circle. We shall for simplicity assume that y_t is a scalar signal, although this is not essential to our future developments. For brevity of notation, we shall in future often refer to a model $\hat{P}(z, \theta)$, $\hat{H}(z, \theta)$, for a particular value of θ in \mathcal{D}_θ , as the model θ . Associated with the model (6.6) is the one-step-ahead predictor

$$\hat{y}_{t|t-1}(\theta) = \hat{H}^{-1}(z, \theta)\hat{P}(z, \theta)u_t + [I - \hat{H}^{-1}(z, \theta)]y_t. \quad (6.7)$$

It is worth noting that in very many cases the predictor $\hat{y}_{t|t-1}$ can also be written in pseudolinear regression form

$$\hat{y}_{t|t-1}(\theta) = \phi_t^T(\theta)\theta, \quad (6.8)$$

where $\phi_t(\theta)$ is a d -vector of pseudoregressors. This form is useful for the derivation of algorithms minimizing a positive function of the prediction errors. The prediction error is then

$$\begin{aligned} \epsilon_t(\theta) &= y_t - \hat{y}_{t|t-1}(\theta) \\ &= \hat{H}^{-1}(z, \theta)[(P(z) - \hat{P}(z, \theta))u_t + v_t]. \end{aligned} \quad (6.9)$$

In the special case of models that are linear in the parameters, such as autoregressive models with exogenous inputs (usually called ARX models), ϕ_t in (6.8) depends only on the data and not on the parameter vector: it then becomes a true regressor, rather than a pseudoregressor. The prediction error $\epsilon_t(\theta)$ is also linear in the parameters in this case.

Perhaps the most common way to identify a ‘best’ model in the model set (6.5) is to select the parameter vector θ that minimizes

$$V_N(\theta, \eta) = \frac{1}{N} \sum_{t=1}^N [\epsilon_t^f(\theta, \eta)]^2, \quad (6.10)$$

where $\epsilon_t^f(\theta, \eta)$ denotes the prediction errors filtered through a stable linear filter with transfer function $D(z, \eta)$:

$$\epsilon_t^f(\theta, \eta) = D(z, \eta)\epsilon_t(\theta). \quad (6.11)$$

The parameter vector η indicates that the filter may depend on another set of parameters that are tuned in order to achieve a desirable objective. The parameter estimate is then defined as

$$\hat{\theta}_N = \arg \min_{\theta \in \mathcal{D}_\theta} V_N(\theta, \eta), \quad (6.12)$$

where \mathcal{D}_θ in R^d is a subset of admissible values.

Except for models that are linear in the parameters (such as ARX models), the solution of (6.12) cannot be stated explicitly as a closed form expression. A battery of iterative algorithms exists that can be used to compute an off-line solution of this minimization problem: see [Lju87].

6.2.2 Recursive Least Squares Identification

Alternatively to iterative off-line computation, $\hat{\theta}$ can also be computed recursively or on-line. This means that, at time t and with $\hat{\theta}_{t-1}$ available, one new prediction error is computed from the observed output y_t and the predicted output $\hat{y}_{t|t-1}$:

$$\epsilon_t(\hat{\theta}_{t-1}) = [y_t - \hat{y}_{t|t-1}(\hat{\theta}_{t-1})]. \quad (6.13)$$

The filtered prediction error is obtained by filtering these ϵ_t by $D(z, \eta)$, as anybody could have guessed:

$$\epsilon_t^f(\hat{\theta}_{t-1}, \eta) = D(z, \eta)[y_t - \hat{y}_{t|t-1}(\hat{\theta}_{t-1})]. \quad (6.14)$$

We notice that the computation of $\epsilon_t^f(\hat{\theta}_{t-1}, \eta)$ through the filtering operation (6.14) involves past prediction errors $\epsilon_t(\hat{\theta}_{t-1})$, $\epsilon_{t-1}(\hat{\theta}_{t-2})$, $\epsilon_{t-2}(\hat{\theta}_{t-3})$, \dots that are computed with different $\hat{\theta}$. This is typical of any recursive identification scheme. The parameter vector is then updated as follows:

$$\hat{\theta}_t = \hat{\theta}_{t-1} + \gamma_t R_t^{-1} \psi_t^f(\hat{\theta}_{t-1}, \eta) \epsilon_t^f(\hat{\theta}_{t-1}, \eta), \quad (6.15)$$

where

$$\psi_t^f(\hat{\theta}_{t-1}, \eta) = D(z, \eta) \frac{\partial \hat{y}_{t|t-1}(\hat{\theta}_{t-1})}{\partial \theta}, \quad (6.16)$$

γ_t is a positive scalar that determines the step size of the adaptation, and R_t can take various forms. In a more general context one could take different filters in (6.14) and (6.16) for the error and the regressor respectively (see [Joh88]). The algorithm (6.15) becomes a gradient algorithm for $R_t = I$. It becomes a recursive Gauss-Newton algorithm when R_t is updated as

$$R_t = R_{t-1} + \gamma_t \left[\psi_t^f(\hat{\theta}_{t-1}, \eta) \psi_t^{fT}(\hat{\theta}_{t-1}, \eta) - R_{t-1} \right]. \quad (6.17)$$

Where ϕ_t in (6.8) is independent of θ , as in the case with ARX models, ψ_t^f is equal to the filtered regression vector, which is also independent of θ :

$$\psi_t^f(\hat{\theta}_{t-1}, \eta) = \phi_t^f(\eta) = D(z, \eta) \phi_t. \quad (6.18)$$

The algorithm (6.14)–(6.17), with ψ replaced by ϕ , then becomes the well-known recursive least squares algorithm for linear-in-the-parameter models. For the sake of simplicity, in our further developments we shall use the term Recursive Least Squares for any recursive algorithm that results from minimizing a Least Squares cost function recursively in the general framework just described, whether the model is linear in the parameters or not. We shall deal later with some of the many variations that have been proposed for the general algorithm (6.14)–(6.17) in order to improve either its robustness or its tracking capabilities, particularly when used in adaptive control. What will be important for our further exposé is the rôle played by the adaptation gain, γ_t , which determines the speed of adaptation. This gain can be explicitly related to an exponential forgetting factor for the LS criterion, and in an identification application (as opposed to an adaptive control application) its selection represents a trade-off between tracking ability and noise rejection. In an adaptive control context, other considerations in addition to tracking such as closed loop stability may prevail, as we shall see later in this chapter.

6.3 Frequency Domain Properties of the Identified Model

As we stated earlier in this chapter, we shall be primarily concerned with the analysis of the case where the chosen model structure is unable to represent the true system for any of its parameter values. This is almost always the case in practice, because the true system is of much higher complexity than the available model structures. It is then important to be able to characterize the misfit between the true system and the estimated model. The misfit between the true system and the best possible model within the model set (i.e. the model that would be obtained if an infinite number of data were available and the model were fitted off-line) is called the bias. As it turns out, it is possible to give a precise characterization of the distribution of this bias in the frequency domain, and to extract from this characterization the means of affecting this bias distribution by a suitable manipulation of the input signal spectrum, the filter $D(z, \eta)$ and, possibly, the choice of the noise model, $\hat{H}(z, \theta)$.

When the identification experiment is carried out in closed loop with a fixed regulator, the input signal spectrum is of course a function both of this regulator and of the reference input, and again it is important to examine how these quantities will affect the distribution of the bias of the estimated model. The situation becomes more complicated, but more crucial from a convergence and closed loop stability point of view, when the regulator parameters are themselves functions of the model parameters, as is the case in an indirect adaptive control situation, where the regulator parameters are continuously adjusted on-line as a function of the estimated model parameters. We shall later be appealing to the timescale separation of slowly adapting systems to apply the results for fixed closed loop controllers to adaptive control.

In this section, we shall present the formulae for the asymptotic (as the number of data tends to infinity) distribution over frequency of the bias of an estimated model fitted with a filtered LS criterion and quasistationary signals. We shall first do this for the case where identification takes place in open loop, then for the case of closed loop identification. The derivation of the asymptotic bias expressions which we shall present is due to the pioneering work of Wahlberg and Ljung [WL86]. In the next section, we shall then examine the consequences of these results in the case of simultaneous identification and control, i.e. in the case of adaptive control. All our formulae will be derived for the simple case of a single input–single output system. The extension to the multivariable case is reasonably straightforward, but

the multivariable formulae do not add any additional insight that cannot be gained from the simpler scalar formulae.

6.3.1 Open Loop Identification

Under reasonable conditions [Lju87], the RLS parameter estimate $\hat{\theta}$ will converge asymptotically to a value θ^* defined as follows:

$$\theta^* = \arg \min_{\theta \in \mathcal{D}_\theta} \bar{V}(\theta, \eta), \quad (6.19)$$

where

$$\begin{aligned} \bar{V}(\theta, \eta) &= \lim_{N \rightarrow \infty} EV_N(\theta, \eta) \\ &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N E[\epsilon_t^f(\theta, \eta)]^2. \end{aligned} \quad (6.20)$$

Denoting by $\Phi_\epsilon^f(\omega)$ the spectrum of the filtered prediction errors ϵ_t^f (see (6.11)), (6.3)–(6.4) and using Parseval's identity allows us to re-express the limiting value θ^* as follows:

$$\theta^* = \arg \min_{\theta \in \mathcal{D}_\theta} \int_{-\pi}^{\pi} \Phi_\epsilon^f(\omega) d\omega. \quad (6.21)$$

Using (6.9) and (6.11), the filtered prediction errors can be written as

$$\epsilon_t^f(\theta, \eta) = D(z, \eta) \hat{H}^{-1}(z, \theta) [(P(z) - \hat{P}(z, \theta))u_t + v_t]. \quad (6.22)$$

When the system operates in open loop, u_t and v_t are independent, and therefore (6.21)–(6.22) yields

$$\theta^* = \arg \min_{\theta \in \mathcal{D}_\theta} \int_{-\pi}^{\pi} \left[|P(e^{j\omega}) - \hat{P}(e^{j\omega}, \theta)|^2 \Phi_u(\omega) + \Phi_v(\omega) \right] \frac{|D(e^{j\omega}, \eta)|^2}{|\hat{H}(e^{j\omega}, \theta)|^2} d\omega. \quad (6.23)$$

The asymptotic model can therefore be seen as a compromise between fitting the input–output transfer function, $\hat{P}(z)$, to the true transfer function, $P(z)$, in a frequency weighted norm, and fitting the noise model spectrum $\hat{H}(z, \theta)$ to the error spectrum of $\{\epsilon_t^f\}$. A thorough discussion of the consequences of this result and of ways of manipulating this asymptotic bias can be found in either [WL86] or in [Lju87]. Let us just note that, in the case of a fixed noise model ($\hat{H}(e^{j\omega})$ independent of θ), the formula shows that the open loop transfer function is fitted to the true transfer function with a frequency weighting that is proportional to the input spectrum and the data filter

spectrum, and inversely proportional to the fixed noise model spectrum. (Note that under these assumptions $\Phi_v(\omega)$ does not play any part in the minimization.) This allows the user, by a proper choice of these design quantities, to decide in which frequency bands she wants a tight fit between the estimated model and the true plant transfer function, and in which frequency bands she tolerates more slackness.

6.3.2 Closed Loop Identification

We now consider the case where the identification is performed while the system is operating in closed loop. The formulae for the frequency distribution of the asymptotic model bias in the closed loop identification case have been derived in [Gun88]. We shall here give slightly different formulae that will be better suited to our subsequent discussion of their application to the case of LQG regulators. Suppose for the time being that the regulator is fixed, and call ρ the vector of regulator parameters. In order subsequently to specialize our analysis to the case of LQG regulators, we shall consider two degree of freedom regulators of the following form,

$$u_t = F_1(z, \rho)r_t - F_2(z, \rho)y_t, \quad (6.24)$$

where r_t and y_t are, respectively, the reference signal and the measured output of the plant and F_1 and F_2 are the transfer functions of the two components of the two degree of freedom regulator, parametrized by a parameter vector ρ . Replacing u_t in (6.22) by this regulator expression and using (6.1) yields

$$\begin{aligned} \epsilon_t^f(\theta, \eta, \rho) &= D(z, \eta)\hat{H}^{-1}(z, \theta) \left[\Delta P(z, \theta) \frac{F_1(z, \rho)}{1 + F_2(z, \rho)P(z)} r_t \right. \\ &\quad \left. + \frac{1 + F_2(z, \rho)\hat{P}(z, \theta)}{1 + F_2(z, \rho)P(z)} v_t \right] \end{aligned} \quad (6.25)$$

$$\begin{aligned} &= D(z, \eta)\hat{H}^{-1}(z, \theta) [\Delta P(z, \theta)W_1(z, \rho)r_t \\ &\quad + (\Delta P(z, \theta)W_2(z, \rho) + 1)v_t], \end{aligned} \quad (6.26)$$

where

$$\begin{aligned} \Delta P(z, \theta) &= P(z) - \hat{P}(z, \theta) \\ W_1(z, \rho) &= \frac{F_1(z, \rho)}{1 + F_2(z, \rho)P(z)} \end{aligned} \quad (6.27)$$

$$W_2(z, \rho) = -\frac{F_2(z, \rho)}{1 + F_2(z, \rho)P(z)}. \quad (6.28)$$

It is worth noting that $W_1(z, \rho)$ and $W_2(z, \rho)$ are the transfer functions from r_t to u_t and from v_t to u_t , respectively, in the closed loop system,

$$u_t = W_1(z, \rho)r_t + W_2(z, \rho)v_t. \quad (6.29)$$

For a system operating in closed loop with the regulator (6.24), the limiting value of θ^* is therefore given by the following expression:

$$\theta^* = \arg \min_{\theta \in \mathcal{D}_\theta} \bar{V}(\theta, \eta, \rho), \quad (6.30)$$

where

$$\bar{V}(\theta, \eta, \rho) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N \mathbb{E}[\epsilon_t^f(\theta, \eta, \rho)]^2, \quad (6.31)$$

with $\epsilon_t^f(\theta, \eta, \rho)$ defined by (6.25) or (6.26). This yields the following expressions for the frequency distribution of the limiting model in closed loop:

$$\begin{aligned} \theta^* &= \arg \min_{\theta \in \mathcal{D}_\theta} \int_{-\pi}^{\pi} \left[|\Delta P(e^{j\omega}, \theta)|^2 \frac{|F_1(e^{j\omega}, \rho)|^2}{|1 + F_2(e^{j\omega}, \rho)P(e^{j\omega})|^2} \Phi_r(\omega) \right. \\ &\quad \left. + (|\Delta P(e^{j\omega}, \theta)|^2 \frac{|F_2(e^{j\omega}, \rho)|^2}{|1 + F_2(e^{j\omega}, \rho)P(e^{j\omega})|^2} + 1) \Phi_v(\omega) \right] \frac{|D(e^{j\omega}, \eta)|^2}{|\hat{H}(e^{j\omega}, \theta)|^2} d\omega. \quad (6.32) \\ &= \arg \min_{\theta \in \mathcal{D}_\theta} \int_{-\pi}^{\pi} \left[|\Delta P(e^{j\omega}, \theta)|^2 |W_1(e^{j\omega}, \rho)|^2 \Phi_r(\omega) \right. \\ &\quad \left. + (|\Delta P(e^{j\omega}, \theta)|^2 |W_2(e^{j\omega}, \rho)|^2 + 1) \Phi_v(\omega) \right] \frac{|D(e^{j\omega}, \eta)|^2}{|\hat{H}(e^{j\omega}, \theta)|^2} d\omega. \\ &= \arg \min_{\theta \in \mathcal{D}_\theta} \int_{-\pi}^{\pi} \left[|\Delta P|^2 (|\hat{W}_1|^2 \Phi_r + |\hat{W}_2|^2 \Phi_v) |\hat{W}_A|^2 + \Phi_v \right] \frac{|D|^2}{|H|^2} d\omega, \quad (6.33) \end{aligned}$$

where

$$\hat{W}_1(z) = \frac{F_1(z, \rho)}{1 + F_2(z, \rho)\hat{P}(z)} \quad (6.34)$$

$$\hat{W}_2(z) = -\frac{F_2(z, \rho)}{1 + F_2(z, \rho)\hat{P}(z)} \quad (6.35)$$

$$\hat{W}_A(z) = \frac{1 + F_2(z, \rho)\hat{P}(z)}{1 + F_2(z, \rho)P(z)}. \quad (6.36)$$

Note, at this point, that the weighting functions \hat{W}_1 and \hat{W}_2 depend only upon the controller design and not upon the actual plant $P(z)$. The influence

of the true plant on this frequency domain optimization is via the factor \hat{W}_A of (6.36) which, as will be shown in Chapter 8, is the ratio of achieved to designed sensitivity functions. This factor should, for well modeled and robustly controlled systems, be close to one over the passband of the closed loop system. If, as is reasonable, we take $|\hat{W}_A| \approx 1$ over this band, then the effect of the control design on the identification becomes clear from (6.33).

This expression shows exactly how a parametrized model of the form (6.6) approximates the true system in the limit as the number of data tends to infinity when that true system is not in the model set and when the identification is performed in closed loop with the regulator (6.24) operating. Even though it only provides an implicit expression for the minimizing θ^* , it clearly exhibits the influence over the frequency distribution of the asymptotic bias of the noise and reference signal spectra Φ_v and Φ_r , of the prefilter $D(z, \eta)$ and of the regulator transfer functions $F_1(z, \rho)$ and $F_2(z, \rho)$.

Finally, we can describe the limiting model in the special case of LQG control (see Chapter 3). It suffices to replace $\hat{W}_1(z, \rho)$ and $\hat{W}_2(z, \rho)$ by their expressions from (3.88) and (3.91). In the case where observers with direct feedthrough are used, which is the simpler case to write down, we have from (3.91):

$$\hat{W}_1(z) = [1 - K^x(zI - \hat{F})^{-1}\hat{G}]^{-1}K^r(zI - F^r)^{-1}G^r z^{N+1} \quad (6.37)$$

$$\hat{W}_2(z) = [1 - K^x(zI - \hat{F})^{-1}\hat{G}]^{-1}K^m(zI - \hat{F}^m + M^F \hat{H}^m \hat{F}^m)^{-1}M^F z, \quad (6.38)$$

with H^m , F^m defined in Section 3.5 and $K^m \triangleq \begin{pmatrix} K^x & K^d \end{pmatrix}$. The formula (6.33) is the fuller analysis of closed loop identification promised with the Gedankenexample in Section 4 of Chapter 1. The reader is invited to review this computational example for moral support at this stage.

Here we draw the reader's attention to a crucial point in the foregoing presentation which will recur at several stages in the forthcoming development. That is, that the controller transfer functions $\hat{W}_1(z)$ and $\hat{W}_2(z)$ are computed based upon the identified model $\hat{P}(z, \theta)$ and not upon the actual plant $P(z)$. Hence we have made explicit the matrices with *chapeaux* in (6.37) and (6.38). The dependence upon the identified parameters of the weighting factors in the RLS equivalent integral represents a further complication of proceedings where the minimization in (6.33) becomes an opaque implicit optimization problem. What is required to alleviate matters and reassert the transparency of this frequency domain analysis is to insist that these weighting functions are, to first order, independent of θ , at least in the neighborhood of θ^* . That is, we require that the achieved closed loop be

roughly impervious to local changes in θ . That is not to say that the prediction error becomes independent of θ but that the nondifferenced signals do so. This is clearly a robustness requirement on the applied control law and a statement with implications for local validity of the heuristic interpretation of these formulae. We shall return again to this issue later but flag at this stage that, just as identifier performance affects the controller through its robustness, so does the controller and its robustness arise as critical for the adequate predictable behavior of the identifier.

6.4 Recursive Identification in Closed Loop Control — Local Theory

6.4.1 Heuristic Motivation

We have seen in the previous section that, when identification is performed in closed loop in an off-line manner, the regulator influences the frequency spectrum of the plant input, and therefore affects the bias distribution of the estimated model. This is apparent from the expression (6.30)–(6.33), which we rewrite in a concise way as a mapping from the regulator and filter parameters to the asymptotic value θ^* ,

$$\theta^* = \Pi(\rho, \eta). \quad (6.39)$$

This mapping brings out quite clearly the fact that different regulator parameters and different filters will asymptotically lead to different estimated models. We stress that this mapping is representative of an off-line identification applied to a closed loop system with a fixed regulator and a fixed filter.

In an adaptive control-type situation, the regulator parameters will be continuously adjusted as a function of the on-line estimates $\hat{\theta}_t$. Depending on the designer's choice, the filter parameters may either be fixed, or they may also be recursively updated as a function of the $\hat{\theta}_t$. This means that the parameter vectors ρ and η may be functions of the recursively estimated $\hat{\theta}_t$:

$$\rho_t = \rho(\hat{\theta}_t) \quad \eta_t = \eta(\hat{\theta}_t). \quad (6.40)$$

As a consequence there is no guarantee, in an adaptive closed loop, that the solution of the recursive parameter estimator will converge, and, even if it does so, it is not clear what the meaning of the convergence point might be. To get a handle on this question, it is perhaps easiest to try and first formulate the off-line problem of which the recursive scheme is an approximate

iterative minimization. It is natural to consider the minimization problem (6.30) in which η and ρ are replaced by $\eta(\theta)$ and $\rho(\theta)$, respectively,

$$\theta^* = \arg \min_{\theta \in \mathcal{D}_\theta} \bar{V}(\theta, \eta(\theta), \rho(\theta)). \quad (6.41)$$

This is a highly nonlinear minimization problem, and it is not clear that it has a unique (or any) solution in general, or that its solution θ^* , if it exists, yields a stabilizing closed loop when it is used to design an LQG regulator. We shall henceforth call a stabilizing solution, or stabilizing model, a model θ for which the corresponding LQG regulator yields a stabilizing closed loop when applied to the real plant. Our first task in this section will be to specify, via the LQG criterion, a domain \mathcal{D}_θ for which the problem (6.41) has a stabilizing solution.

Having defined a minimizing and stabilizing value θ^* of the off-line minimization problem (6.41), we now assert that the solution of the recursive parameter estimation problem, within an adaptive closed loop, need not converge near such a solution or may even diverge to infinity. The mechanism for such ‘bad craziness’¹ can be intuitively understood as follows. Assume that at time t the plant model parameter estimates are given by $\hat{\theta}_t$. Let the corresponding regulator parameters be ρ_t ; assume for simplicity that the filter parameters are constant. The new regulator parameters will affect the new input signal, which itself affects the new parameter estimate $\hat{\theta}_{t+1}$. As the recursions proceed, the regulator may progressively push the input signal spectrum into high frequency bands, say, that will produce a $\hat{\theta}$ model for which the closed loop system has a poor stability margin. This will then produce an input signal spectrum with even higher frequency contents, producing a model of even worse quality, which eventually can become unstable. The reader is referred back to the Gedankenexample of Chapter 1 for partial evidence of just this effect.

Our next task, therefore, will be to prove that if the adaptive algorithm is initialized with a value $\hat{\theta}_0$ that is close enough to the optimal value θ^* of the off-line minimization problem (6.41), and if the adaptation gain γ of the parameter estimator is small enough, then the trajectory of the parameter estimation algorithm will remain within a set of stabilizing models, with the additional assumption of persistency of excitation of the external reference signal. And further, the estimate will converge to the neighborhood (and not necessarily a single point in this neighborhood) of a point that is close to the optimal solution of the off-line problem. We will appeal here to the powerful integral manifold theory developed for similar problems by Riedle and Kokotovic [RK86a], [RK86b].

¹*op. cit.* *Fear and Loathing in Las Vegas*, Hunter S. Thompson.

6.4.2 Potential Convergence Point

As promised above (promises are one thing we do like to keep), we now set ourselves the task of defining a domain \mathcal{D}_θ for which the minimization problem (6.41) will yield a stabilizing solution. When the true system is contained in the model set and when the external reference signal is persistently exciting, the identified model $\hat{P}(z, \hat{\theta})$ will, under reasonable conditions, converge to the true system $P(z)$, and therefore, if the regulator has been designed by a lucid control engineer, the closed loop will be stable. However, in the case of model mismatch, and unless special precautions are taken, there is no reason why a model that minimizes a prediction error criterion should also minimize in any way a control criterion and therefore stabilize the plant. Since our aim is to promote the use of indirect LQG control, we shall define a desirable θ^* model, which is the solution of a minimization problem (6.41) via the solution of an LQG model selection problem.

Consider that the true plant is described by (6.1), and that a model set has been selected together with corresponding state-variable descriptions (see Chapter 3 and (6.5)), $\{\hat{P}(z, \theta), \hat{H}(z, \theta), \theta \in \mathcal{D}_\theta\}$. Consider now that we have decided on a particular LQG criterion for our control design problem that only involves measurable signals (i.e. inputs, outputs and reference signals) such as (3.11). For any model θ in our model set, the solution of the LQG problem leads to a particular regulator, parametrized by a vector, say $\rho(\theta)$, and the application of that regulator to the actual plant yields a corresponding LQ cost $J(\theta)$. Note that $J(\theta)$ is computed from the LQ criterion using the measured inputs and outputs. It is the achieved cost, but with a regulator computed on the basis of the model $\hat{P}(z, \theta), \hat{H}(z, \theta)$. Hence the notation $J(\theta)$.

We then define the best model within our model set as the model that produces a regulator that will minimize the achieved LQ cost, as measured on the controlled plant. We call θ^{**} the parameter vector that defines that model:

$$\theta^{**} = \arg \min_{\theta \in \mathcal{D}_\theta} J(\theta). \quad (6.42)$$

Note that by construction, and provided the standard stabilizability assumption is satisfied and the LQG control problem has been selected properly, θ^{**} must correspond to a stabilizing controller.

When the true plant is in the model set, the minimization of the criterion (6.42) and that of any reasonably formulated prediction error criterion will, under suitable conditions, yield the same (i.e. exact) model. In the case of unmodeled dynamics, θ^{**} need not be obtainable as the solution of a prediction error criterion minimization, for the simple reason that the criteria

are entirely different. We shall now show that, provided the model set is such that the plant-model mismatch is not too large, one can define a model θ^* , obtained via the minimization of the prediction error criterion (6.41), that is close enough to the θ^{**} model also to be stabilizing, provided the domain of admissible values \mathcal{D}_θ is suitably restricted. In the final step of our analysis, we will then use averaging theory to show that the recursive parameter estimate will converge to a neighborhood of the stabilizing value θ^* , provided the adaptation is slow enough and the initial estimate is sufficiently close. That is, if one starts from a reasonable parameter value then one's solution tends to still more advantageous values.

We first rewrite the system (6.1) as

$$y_t = \phi_t^T \theta^{**} + \zeta_t + v_t, \quad (6.43)$$

where ϕ_t is a regressor of past inputs and outputs, θ^{**} is defined above, v_t is as in (6.1) and ζ_t denotes the unmodeled dynamics. Note that (6.43) is as much a definition for ζ_t as a rewriting of the plant description. When the plant is operating under feedback control with a stabilizing LQG regulator computed via a model θ , ϕ_t and ζ_t in (6.43) are functions of θ , and will therefore be denoted by $\phi_t(\theta)$ and $\zeta_t(\theta)$. We shall then assume that v_t is uncorrelated with both $\phi_t(\theta)$ and $\zeta_t(\theta)$. As in all our analysis in this chapter, we assume that the reference signals and the external perturbations are such that the signals in the time-invariant closed loop system are quasistationary. We can then introduce the following notations,

$$m = E[|\phi_t(\theta^{**})|] \quad (6.44)$$

$$\lambda = \lambda_{min} E[\phi_t(\theta^{**}) \phi_t^T(\theta^{**})], \quad (6.45)$$

where λ_{min} denotes the smallest eigenvalue. We shall also from now on denote by $\bar{V}(\theta)$ the value of the cost (6.31), where it is understood that ρ and, possibly, η are functions of the model θ , with the regulator computed as the solution of an LQG criterion.

We now introduce an important assumption on the amount of unmodeled dynamics that is tolerated in the model θ^{**} .

Assumption 6.1

There exists a closed hypersphere

$$B_r(\theta^{**}) = \{\theta : |\theta - \theta^{**}| \leq r\}$$

*centered on θ^{**} with radius r such that*

1. *for all $\theta \in B_r(\theta^{**})$ the closed loop system is stable,*

2. there exist positive constants α, β, δ, k such that for all $\theta \in B_r(\theta^{**})$

$$\begin{aligned} \mathbb{E}|\zeta_t(\theta^{**})| &\leq \alpha(\mathbb{E}|\phi_t(\theta^{**})| + k), & \mathbb{E}\left|\frac{\partial\phi_t(\theta)}{\partial\theta}\right| &< \beta, \\ \mathbb{E}\left|\frac{\partial\zeta_t(\theta)}{\partial\theta}\right| &< \delta \end{aligned} \quad (6.46)$$

with α, β , and δ small enough so that

$$r\lambda > 2[\beta mr^2 + \delta mr + \beta\delta r^2 + \alpha(m+k)(m + \beta r + \delta)]. \quad (6.47)$$

Assumption 6.1 has the following implications:

- the first part assumes that around the ‘best’ system θ^{**} there is a neighborhood of stabilizing models. This is a very reasonable assumption if the closed loop regulator is computed using a robust design methodology as recommended in Chapter 5;
- the second part is a constraint on keeping the unmodeled dynamics small enough relative to the regressor (α, k) , as well as an assumption of smoothness of both the regressor (β) and the unmodeled dynamics of the closed loop system (δ) with respect to the model parameter θ . We notice that satisfaction of the constraint (6.47) also hinges on the amount of persistence of excitation of the regressor vector $\phi_t(\theta^{**})$ through the parameter λ ;
- the two parts of the assumption really address the existence of B_r and then define the radius of the hypersphere through the combined constraints of closed loop stability, small unmodeled dynamics and smoothness. Condition (6.47) is the limiting constraint upon the applicability of this analysis to establish robust stability of adaptive control for a particular system;
- finally, we notice that the satisfaction of the constraint (6.47) is the critical condition for the following lemma.

The model θ^{**} has been defined through the minimization of an LQ cost over all models in a set \mathcal{D}_θ . We now show that, under Assumption 6.1, close to θ^{**} (i.e. inside the hypersphere $B_r(\theta^{**})$) there exists a stabilizing model θ^* that is defined through the minimization of an off-line least squares prediction error criterion.

Lemma 6.1

Let θ^{**} be defined as the LQ optimal parameter value as in (6.42) and let $B_r(\theta^{**})$ satisfy Assumption 6.1. Then there exists a θ^* , an interior point of $B_r(\theta^{**})$, defined as the following solution of a prediction error problem:

$$\theta^* = \arg \min_{\theta \in B_r(\theta^{**})} \bar{V}(\theta, \rho(\theta), \eta(\theta)). \quad (6.48)$$

Proof Since θ^* is in the closed hypersphere $B_r(\theta^{**})$ by construction, the proof consists of showing that it cannot be on the boundary. Let θ_0 be an arbitrary point on the boundary of $B_r(\theta^{**})$; we shall show that $\bar{V}(\theta_0) > \bar{V}(\theta^{**})$. Recall that $\bar{V}(\theta)$ has to be interpreted, as indicated above Assumption 6.1, as the value of (6.31) for a particular θ . Consider first the model θ^{**} and its corresponding optimal prediction error cost $\bar{V}(\theta^{**})$. From (6.43), it follows that the predicted value $\hat{y}_t(\theta^{**})$ can be written as

$$\hat{y}_t(\theta^{**}) = \phi_t^T(\theta^{**})\theta^{**}. \quad (6.49)$$

The prediction error follows immediately:

$$\epsilon_t(\theta^{**}) = y_t - \hat{y}_t(\theta^{**}) = \zeta_t(\theta^{**}) + v_t. \quad (6.50)$$

Therefore, with $D = 1$, and using the independence assumption on $\{\zeta_t\}$ and $\{v_t\}$,

$$\bar{V}(\theta^{**}) = \text{E} [\zeta_t^2(\theta^{**}) + v_t^2]. \quad (6.51)$$

Consider now any model θ_0 on the boundary. Then

$$\begin{aligned} \bar{V}(\theta_0) &= \text{E}[y_t - \hat{y}_t(\theta_0)]^2 \\ &= \text{E} \left[\phi_t^T(\theta_0)\tilde{\theta} + \zeta_t(\theta_0) + v_t \right]^2, \end{aligned} \quad (6.52)$$

where $\tilde{\theta} = \theta_0 - \theta^{**}$. Using Taylor's formula with remainder, we can write

$$\begin{aligned} \phi_t(\theta_0) &= \phi_t(\theta^{**}) + \tilde{\theta}^T \frac{\partial \phi_t}{\partial \theta}(\theta_1) \\ \zeta_t(\theta_0) &= \zeta_t(\theta^{**}) + \tilde{\theta}^T \frac{\partial \zeta_t}{\partial \theta}(\theta_2), \end{aligned} \quad (6.53)$$

where θ_1 and θ_2 are two intermediate points in $B_r(\theta^{**})$ between θ_0 and θ^{**} . $\bar{V}(\theta_0)$ can then be written as

$$\begin{aligned} \bar{V}(\theta_0) &= \text{E} [\zeta_t(\theta^{**}) + v_t]^2 \\ &+ \text{E} \left[\phi_t^T(\theta^{**})\tilde{\theta} + \tilde{\theta}^T \frac{\partial \phi_t}{\partial \theta}(\theta_1)\tilde{\theta} + \tilde{\theta}^T \frac{\partial \zeta_t}{\partial \theta}(\theta_2) \right]^2 \\ &+ 2\text{E} [\zeta_t(\theta^{**}) + v_t] \left[\phi_t^T(\theta^{**})\tilde{\theta} + \tilde{\theta}^T \frac{\partial \phi_t}{\partial \theta}(\theta_1)\tilde{\theta} + \tilde{\theta}^T \frac{\partial \zeta_t}{\partial \theta}(\theta_2) \right]. \end{aligned} \quad (6.54)$$

We recognize the first term as being $\bar{V}(\theta^{**})$, and we denote the sum of the other two terms by S . Noting that $|\tilde{\theta}| = r$, recalling that v_t is uncorrelated with both $\phi_t(\theta)$ and $\zeta_t(\theta)$, and using the Cauchy-Schwartz inequality, we can write

$$\begin{aligned} S &\geq \text{E} \left[\phi_t^T(\theta^{**})\tilde{\theta} + \tilde{\theta}^T \frac{\partial \phi_t}{\partial \theta}(\theta_1)\tilde{\theta} + \tilde{\theta}^T \frac{\partial \zeta_t}{\partial \theta}(\theta_2) \right]^2 \\ &\quad - 2\alpha r(k+m)(m+\beta r+\delta) \\ &\geq \lambda r^2 - 2[\beta m r^3 + \delta m r^2 + \beta \delta r^3 + \alpha r(k+m)(m+\beta r+\delta)]. \end{aligned} \tag{6.55}$$

It follows directly that $S > 0$ if the inequality (6.47) is satisfied. *CQFD*

The point of Lemma 6.1 has been to show that, under conditions of smoothness and limited unmodeled dynamics, there exists a closed hypersphere $B_r(\theta^{**})$ surrounding θ^{**} , the interior of which contains a stabilizing model θ^* that can be obtained as the solution of an off-line prediction error identification problem, with the search domain suitably restricted to that hypersphere. The model θ^{**} can then be seen as a vehicle for suitably defining the model θ^* as the solution of an off-line identification problem. The trick that made this vehicle deliver its goods was to impose a constraint (Assumption 6.1) that restricted the amount of allowable unmodeled dynamics. This causes the minimum of the prediction error criterion and the optimal control performance to be related. We should note that conditions like Assumption 6.1 are very standard in the literature on robust indirect adaptive control: see for example [Sam82]. The churlish reader will undoubtedly expect us at this point not to launch into the ubiquitous integral manifold theory — but we will anyway. This provides the connection between the existence of the solutions to the off-line minimization and their attractivity for slowly adapting recursive estimators.

6.4.3 Integral Manifolds and Slow Adaptation

Having defined θ^* , our analysis will proceed by demonstrating that, with a suitably restricted search domain and a sufficiently small adaptation gain, the solution of our recursive prediction error algorithm, used in an adaptive closed loop, will converge to a neighborhood of stabilizing models around θ^* . The thrust of our argument will be as follows. We first describe the nonlinear equations of the complete adaptive control system, and argue that these can be split up into fast θ -dependent ‘state’ equations and slow parameter update equations, provided the adaptation gain γ is small enough. The average behavior of the parameter update equations can be described by an

ordinary differential equation (ODE), whose convergence point θ_0 is close to θ^* if the unmodeled dynamics are not too important and if there is sufficient excitation. Under suitable conditions similar to those of Assumption 6.1, the solution of the parameter update equation can then be shown by integral manifold arguments to converge, from a suitable region of initial conditions, to a limiting solution, which itself is close to θ_0 . This solution to which we converge need not be a point but is a function of time which remains close to θ_0 . The application of integral manifold theory to the analysis of adaptive control systems is due to Riedle and Kokotovic [RK86a], [RK86b], and we shall refer to their work for proofs and details.

The full set of dynamical equations of the adaptive closed loop system can be written in compact form as follows:

$$\Xi_{t+1} = A(\theta_t)\Xi_t + B(\theta_t)n_t, \quad \Xi \in R^s, \quad (6.56)$$

$$\theta_{t+1} = \theta_t + \gamma f_t(\theta_t, \Xi_t), \quad \theta \in R^d, \quad (6.57)$$

where Ξ includes the states of the true plant, of the plant model, of the observer, of the regulator and of the filters, n_t denotes a vector made up of all the external signals (i.e. reference signals and noises), while the parameter update equation (6.57) is just another expression for the recursive least squares equation (6.15) with a constant gain γ . Examples of adaptive control algorithms rewritten in this global form can be found in [RK86b], [ABJ+86] and others. In Adaptive Optimal Control θ_t in (6.57) is clearly the parameter update. The dependence of $A(\cdot)$ and $B(\cdot)$ on θ includes the solution of LQ and KF AREs, for example.

With θ^* defined as in the previous section, we now make the following assumption:

Assumption 6.2

1. *There exists a compact set Θ containing θ^* and constants $\lambda \in (0, 1)$ and $K_1 \geq 1$ such that $\forall \theta \in \Theta$ and $\forall t \geq 0$*

$$|A(\theta)^t| \leq K_1 \lambda^t. \quad (6.58)$$

2. *There exist constants c , c_1 and c_2 such that the frozen parameter response*

$$\nu_t(\theta) = \sum_{j=0}^{\infty} A^j(\theta)B(\theta)n_{t-j-1} \quad (6.59)$$

and its sensitivity $\frac{\partial \nu_t}{\partial \theta}(\theta)$ satisfy

$$\begin{aligned} |\nu_t(\theta)| &\leq c, & \left| \frac{\partial \nu_t}{\partial \theta}(\theta) \right| &\leq c_1 \\ \left| \frac{\partial \nu_t}{\partial \theta}(\theta) - \frac{\partial \nu_t}{\partial \theta}(\theta^*) \right| &\leq c_2 |\theta - \theta^*| \end{aligned} \quad (6.60)$$

for all $t \in Z$ and all $\theta, \theta^* \in \Theta$.

3. The function $f_t(\theta, \Xi)$ in (6.57) is bounded, Lipschitz in θ and Ξ uniformly with respect to t , $\theta \in \Theta$ and Ξ in compact sets.

The first part of Assumption 6.2 says that around the ‘optimal’ model θ^* there is a compact set of models that also produce exponentially stable closed loops. This is a reasonable assumption given that, if in the selection of our model θ^* the filter $D(z, \theta)$ has been chosen so as to maximize the robustness of the ensuing closed loop to unmodeled dynamics according to the prescriptions of Chapter 5 (see also Chapters 7 and 8), then this will have the effect of maximizing the size of the compact set around θ^* for which closed loop stability can be guaranteed. The second and third parts of Assumption 6.2 are essentially smoothness assumptions on the nonlinear model (6.56)–(6.57).

The integral manifold theory of Riedle and Kokotovic is based on a time-scale separation between the dynamical equations for Ξ and for θ . The idea is that, if the gain γ is small enough, the solution of the parameter update law can be approximated by the solution of an ‘averaged’ equation,

$$\bar{\theta}_{t+1} = \bar{\theta}_t + \gamma \bar{f}(\bar{\theta}_t), \quad (6.61)$$

where \bar{f} is obtained by averaging f over t with $\bar{\theta}$ fixed. Asymptotically stable solutions of (6.61) can in turn be related to solutions of the ODE,

$$\frac{d\hat{\theta}}{dt} = \bar{f}(\hat{\theta}). \quad (6.62)$$

We then have the following important stability result, which we paraphrase in words, leaving out the precise values of the bounds: see [RK86a], [RK86b] for details and exact values of the bounds.

Theorem 6.1

Suppose that θ^* , defined by (6.48), is a local minimum,

$$\theta^* = \arg \min_{\theta \in B_K(\theta^*)} \bar{V}(\theta, \eta(\theta), \rho(\theta)), \quad (6.63)$$

of $\bar{V}(\theta)^2$, defined by (6.31), with $B_K(\theta^*) \triangleq \{\theta \in R^d : |\theta - \theta^*| \leq K\}$ such that the vector ψ_t^f is persistently exciting $\forall \theta \in B_K(\theta^*)$. If $\bar{V}(\theta^*)$ is small enough, then

1. the ODE (6.62) has an asymptotically stable equilibrium point θ^0 such that

$$|\theta^0 - \theta^*| \leq b\bar{V}(\theta^*) \text{ for some finite } b; \quad (6.64)$$

2. given $\chi > 0$, there exists a sufficiently small $\gamma^*(\chi)$ such that, for $\gamma \in (0, \gamma^*)$, the equation (6.57) possesses a bounded uniformly asymptotically stable solution $\tilde{\theta}_t(\gamma)$ which is close to θ^0 ,

$$\lim_{\gamma \rightarrow 0} |\tilde{\theta}_t(\gamma) - \theta^0| = 0; \quad (6.65)$$

3. every solution $\theta_t(\gamma)$ of (6.57) with $\theta_0(\gamma) \in B_{K-\chi}(\theta^*)$ satisfies, for $\gamma \in (0, \gamma^*)$,

$$\theta_t(\gamma) \in B_K(\theta^*), \quad \lim_{t \rightarrow \infty} |\theta_t(\gamma) - \tilde{\theta}_t(\gamma)| = 0. \quad (6.66)$$

The constant b in the first part of the theorem is proportional to the maximum value, over all the models in the hypersphere $B_K(\theta^*)$, of the average value of the Euclidean norm of the regressors, and is inversely proportional to the average amount of excitation in the filtered regressors $\psi_t^f(\theta)$.

The main conclusion to be drawn from this result is that the solution of the recursive parameter adaptation algorithm, implemented in an adaptive loop, will converge close to the solution θ^* of the off-line problem (6.48), provided the following conditions hold.

1. The plant-model mismatch is small enough; this is embodied in the condition that $\bar{V}(\theta^*)$ must be small enough.
2. The filtered regressors are persistently exciting.
3. The initial condition of the parameter update algorithm is sufficiently close to the ‘optimal’ value θ^* .
4. The models are sufficiently smooth functions of θ around θ^* .
5. The gain of the parameter update algorithm is sufficiently small.

²For reasons of conciseness we shall again delete the explicit dependence of $\bar{V}(\theta)$ on $\rho(\theta)$ and on $\eta(\theta)$.

Further, since under similar assumptions θ^* is close to the performance criterion minimizing θ^{**} , we will have θ converging to the neighborhood of θ^{**} .

The fuller analysis of the quantification of both the limiting and transient response of both θ_t and Ξ_t stemming from the above treatment follows by appeal to the fixed point methods of [ABJ+86]. The achievable quantifications are characterized by the level of undermodeling and the level of persistence of excitation.

6.5 Recursive Identification in Closed Loop — Global Methods

The results of the preceding section convey the picture of the local behavior of a recursive identification scheme operating in an adaptive closed loop. The term ‘local’ refers to the fact that the initial parameter estimate value is required to reside within a fixed distance from the set of final attraction. The reward for analysis under this hypothesis is that rather precise descriptions of the performance of the adaptive system are achievable together with bounds and rates as in the previous section and in [ABJ+86]. The alternative class of possible results is concerned with *global* dynamics of the identifier, where arbitrarily large initial conditions are permitted and ultimate convergence to a set of suitable parameters established. Typically, however, this set is not necessarily small nor is it feasible to supply overbounds on transient performance during this convergence [KMABJ87]. We study some of the rudimentary features of algorithm modifications made to achieve global properties. We do this in order to provide a more complete picture of the measures which might need to be contemplated to force adaptive controllers to function adequately in a broader setting. Unfortunately, the interface between global and local theories is still somewhat hazy, since it is not at present feasible to propose an algorithm which has guaranteed properties from both camps. Certainly, because performance in spite of the control law for these global methods is *de rigueur*, they are seditious in our current development, but are added for completeness and comparison.

In order to establish the central features of the interplay between control design and identification in Adaptive Control, we have so far explicitly needed to appeal to the local theory so that we have available suitably accurate descriptions of the limiting parameter set in closed loop. However, it is naturally desirable to guarantee convergence of the estimated parameters from a rather large initial set to a region from which the local results could be invoked. The difficulty which arises in such global theory is that,

in adaptive feedback systems, the presence of the nonlinear feedback path can lead to wildly unstable systems and thereby extreme ill-conditioning of the signals, in that one (the most unstable) dynamical mode of the plant greatly dominates all input–output data. The response then is to modify the RLS algorithm to cope with this effect. Useful references for these classes of modifications are [Pra88], [Joh88] and [MGHM88].

6.5.1 Normalization and Deadzones

One mechanism proposed to deal with this underinformative data at high signal values is *Normalization*, an approach usually attributed to Egardt [Ega79], Samson [Sam83] and Praly [Pra82] (see also [Pra86]). The key idea is to replace the error signal and the regressor in RLS by derived quantities which are guaranteed bounded by dividing by a normalization signal constructed from past plant inputs and outputs.

Typically one selects positive constants $\bar{\rho}$ and μ and computes

$$\rho_t = \mu\rho_{t-1} + \max\left(|\psi_t^f|, \bar{\rho}\right), \quad (6.67)$$

where μ is connected with a stability margin known (or assumed) to be achievable for the actual plant. One then replaces the filtered prediction error, ϵ_t^f , and filtered regressor, ψ_t^f , in the RLS updates (6.15) and (6.17) by

$$\bar{\psi}_t^f = \frac{\psi_t^f}{\rho_t} \quad (6.68)$$

$$\bar{\epsilon}_t^f = \frac{\epsilon_t^f}{\rho_t}. \quad (6.69)$$

The growth properties of the solution ρ_t of (6.67) with respect to the closed loop signals dictates the boundedness of the normalized signals and, therefore, limits the potential overindulgence of the identifier during unstable adaptation transients. Naturally, this interference in the operation of RLS affects its ability to respond to these transients and so wild excursions in the state Ξ_t are possible during initial phases.

An ideologically similar algorithm modification to cope with similar problems is the *Relative Deadzone*, due (at least in discrete time) to Kreisselmeier and Anderson [KA86]. Here the normalized prediction error above is replaced by its normalized deadzone counterpart,

$$\tilde{\epsilon}_t^f = \mathcal{F}\left(\bar{\epsilon}_t^f\right), \quad (6.70)$$

with

$$\mathcal{F}(x) = \begin{cases} x - d_0, & \text{if } x > d_0 \\ 0, & \text{if } |x| \leq d_0 \\ x + d_0, & \text{if } x < -d_0 \end{cases} \quad (6.71)$$

where d_0 is a constant threshold below which the normalized error is ignored.

One intriguing aspect of systems involving deadzones is that they reject those errors which are too small. Such errors could arise due to normalization with very large signals or, equally, with underexcitation. Thus they frequently have the capacity to switch off during quiescent periods when persistence is lacking. Performance issues still need to be adequately resolved. Initial results are given in [Pra88]. Other embellishments such as directional dependent deadzones and normalization have also been proposed.

6.5.2 Projection and Leakage

As the bulk of the difficulty in ensuring globally adequate behavior of the identifier occurs when attempting to handle local instability, several methods have been advanced which apply *a priori* knowledge of a region of the parameter space in which closed loop stability is assured to restart adaptation from new conditions when instability is detected.

Projection is concerned with forcing a restart of the estimator at a suitable reinitialization after stability is lost for the previous parameter value. This stability could refer either to that of the actual closed loop as detected by signal magnitudes or to that of a computed model. Various suggestions have been brought forward for the precise means of restarting, including projection to a single fixed point, random restart in a given set, backtracking in the history of the adaptation. Each method has its pros and cons.

Leakage refers to the artificial addition of terms in the RLS update which encourage the parameter estimates to tend towards a known reasonable point. In Adaptive Control this is typically introduced as a method for overcoming the drift introduced by additive output disturbances acting upon underexcited systems (see [IK83]). An interesting error-dependent leakage has been suggested by Narendra and Annaswamy [NA87] based on the magnitude of the local error. Unfortunately, the local dynamics of such an adaptive controller demonstrate very complex nonlinear effects.

For both the above methods, information on the potential convergence points of the identification algorithm is required which permits incorporation of these data into the design at the cost of some globality. The performance effects of projection are very difficult to assess, while leakage introduces an offset to the identified parameter even with ideal model matching. A

fatalistic assessment is that these are just the requisite costs to be borne for having the desirable features.

6.5.3 Covariance Resetting

In adaptive RLS estimation problems operating in conditions with variable signal levels there are two effects which can disturb the correct performance of the identification algorithm, viz signal levels being too large and signal levels being too small. The difficulty in either of these conditions is that the step size matrix R_t in (6.17) is affected too greatly by recent signals. Signals being small causes R_t to become large, thereby making the algorithm step size too active. Signals being large causes the opposite effect through the step size turning off.

To overcome these effects, it is common practice to restrict the variability of the R_t matrix by artificial means, so that it neither becomes too large nor too small, yet still preserves a measure of responsiveness to data variations. Such methods might consist of the addition of positive definite terms to the recursion (6.17) and/or the inclusion of negative definite terms which switch in when signal levels become excessive. Many of these types of modifications are discussed in [GS84], [SGM88].

6.5.4 A Global Comment

The theme of our work up to this section has been the coupling between controller objective and identifier objectives in formulating an adaptive control law. This ‘interplay’ has been demonstrated in the local robust stability framework. The central property of the modifications above aimed at global convergence of the identifier is that the parameter estimation algorithm is ‘robustified’ to the effects of the control law. Thus these methods reek of the philosophy of which we complained earlier, i.e. the exclusive focus upon only one component of the adaptive controller, and so are not consistent with our message. The difficulties associated with the global approaches *vis-à-vis* our coupled local approach are that their ability to quantify the achieved closed loop performance is reduced because the effect of the closed loop controller on the parameter set to which the parameter identifier converges is not recognized. The dual issue is that the control law is chosen to be robust to a whole class of potential identified systems without specific regard to the finer modeling structure.

Nevertheless, these global techniques are included for several reasons:

1. They give an insight into the distinction between the two approaches.

2. They indicate what measures might need to be adopted to guarantee that the parameter value reaches a set in which one may appeal to the robustness arguments from an initial value reasonably distant.
3. They display some techniques for handling the difficulties associated with underexcitation.

One might interpret all of these adjustments of this section as a means of coping with persistence of excitation problems, due to either excessive or insufficient signal energy. Our presumption in the local theory is that persistence of excitation is assured whenever the adaptation is active. To us this seems a reasonable assumption, since otherwise the loop signals do not contain sufficient information about the plant to be a good basis on which to design a feedback control law.

One of the few works which also studies the interplay (*Conflict or Conflux*) between the parameter estimation algorithm and the control law is that of Polderman [Pol87]. Although his conclusions, regarding the LQ control laws, seem rather contradictory to the main claim of this treatise, we must stress that his analysis relies on the assumptions that perfect modeling is used (no unmodeled dynamics) and that there is no persistently exciting reference.

The real comparison between the local and global approaches may also be viewed as one of compromise between initial assumptions — initial condition restrictions are traded against structural assumptions on the class of plants and control laws.

6.6 Conclusion

The main message from Chapter 5 is that an off-line design can be guaranteed robust provided a measure of the relative plant-model mismatch can be upperbounded, at each frequency, by a measure of the feedback system's return difference. In order to utilize fully those results in an adaptive context, it is necessary in this chapter to answer two questions:

1. Given that there is a plant-model mismatch, how can the identified model be shaped in a way as to maximize our chances of satisfying this robustness inequality? In layman's language, one might rephrase this question as 'given that I know my model is incapable of representing the true system, how can I organize my identification design in such a way that the model will only be poor at frequencies where it doesn't hurt?'

2. Given that in an adaptive control loop all the computations are performed on-line, and that the closed loop is now described by a complex set of nonlinear equations, are the asymptotic expressions and the design guidelines that were derived in an off-line identification context still valid when the identification is performed within a constantly changing closed loop?

Our presentation of the frequency domain expressions for the asymptotic models was essentially aimed at formulating the equations to allow a response to the first question, while our examination of the consequences of the integral manifold theory of Riedle and Kokotovic was necessary to go from conclusions and design guidelines based on these off-line results to conclusions and design guidelines for the recursive situation. The detailed connections between the resolutions of these two separate issues have been left a little understated because their main purpose is more for theoretical support in order that they later become illuminating. The main lessons to be drawn from this chapter from a practical point of view can briefly be summarized as follows.

- When the plant is not in the model set, the asymptotic formulae for the limit model tell us how the plant-model mismatch is distributed over frequency.
- The implicit definition of an RLS optimal parameter as minimizing a frequency domain integral allows one to view the explicit effects of loop gain on the level of excitation at the plant input. The increase or decrease in this excitation has an obvious effect upon the identification of the plant model at various frequencies. This notion (informally stated here) will later be tied to the need to identify best the plant in certain frequency regions.
- These formulae show that, once the regulator structure has been chosen and given that the reference signal spectrum is usually imposed by the user, the plant-model mismatch can be shaped to fit the robustness needs by a proper choice of the data filter $D(z)$. This will be more fully exploited in Chapters 7 and 8.
- In order for the solution of the adaptive parameter estimator to converge to a solution that is both stabilizing and close to the one that would be expected from a robustness-engineered off-line design, several precautions must be taken. Essentially they are that the initial condition must not be too far away from the ‘optimal’ solution of the

off-line problem, that the absolute plant-model mismatch must not be too large, that the adaptation gain must be kept small enough, and that the reference signal must be persistently exciting.

- Techniques do exist which (theoretically at least) permit adaptation without restriction on either or both the initial conditions and the signal excitation, but these do not take advantage of the potential synergistic coupling between identifier and control law.

With this chapter we have reached the stage where analysis can give way to synthesis and we may concentrate the tools and techniques acquired in the preceding chapters to design a candidate robust adaptive control law. Of critical concern in this venture is the understanding gained here on the local (linearized) behavior of the nonlinear adaptive scheme which permits direct access to the linear robustness principles and off-line identification interpretations. The next chapter concentrates on this synthesis.

Chapter 7

A Candidate Robust Adaptive Predictive Controller

7.1 Introduction

At last the analytical portion of this *magnum opus* is essentially behind us and we next commence the synthetic component by proposing an adaptive controller coupling the robustness theory of LQG/LTR linear control law design (i.e. Chapter 5) with Least Squares identification modeling ideas (Chapter 6). Here it is necessary to combine both facets of the adaptive controller, the identifier and the control law, in order that the synergism of these disparate techniques be brought out.

Recall that the fundamental closed loop stability robustness requirement is given by Theorem 5.3 and

$$\bar{\sigma}[L^{-1}(z) - I] < \min\left(\sigma[I + \hat{\mathcal{G}}(z)], 1\right), \quad \forall |z| \in \Omega, \quad (7.1)$$

where Ω denotes the unit circle. Here $L(z)$ is the multiplicative deviation of the nominal or model transfer function matrix, $\hat{P}(z)$, from that of the actual or true plant system, $P(z)$, and $\hat{\mathcal{G}}(z)$ is the cascaded designed plant/controller or controller/plant transfer function. Hence $L(z)$ is intimately entwined with the system identification aspects of the plant modeling, while $\hat{\mathcal{G}}(z)$ is equally closely enmeshed with the control law. Our thesis now must be to consider how the specification of an adaptive controller can incorporate the manipulation of the identification properties and the tinkering with the controller to cause the behavior of each to enhance separately

the performance of the other. This must operate in closed loop, and should yield considerable leverage to preserve the stabilizing and performance properties of the overall feedback system. As is evident from the subject matter advanced to the reader so far, we intend to present a candidate adaptive controller based upon LQG feedback control and RLS system identification. Thus there will be certain similarities to the features of GPC. Indeed, many of the observed useful properties of GPC in applications can be well explained and analysed by the LQG methods presented so far. It is mainly for sentimental and historical reasons (which, by now, should be apparent) that we shall retain the vestigial term *predictive* for our candidate adaptive controller, although it should also by now be clear to even the most skeptical reader that LQG controllers are just as predictive as GPC, the mechanism for computing predictions being the Kalman predictor (or Kalman filter) in LQG, and equivalent Diophantine equations for output predictions in GPC.

The controller which we propose here is a ‘Certainty Equivalence’ or ‘Indirect’ controller, i.e. it consists of a parameter identifier and a control law design module with the control law using the estimated plant parameters as if they were the actual parameters, albeit with account being taken of their probable inaccuracy. (An alternative class of controller, which we do not treat, is ‘Direct Adaptive Control’, where the controller parameters themselves are estimated rather than generating an explicit plant model.) We commence by considering the control law on its own and then move on to state the identification method used, before explicitly analysing their interplay. It can be seen from the statement that our control law will consist of LQG plus RLS, that already the scope for design choice has been considerably narrowed. A central feature of our Candidate Robust Adaptive Predictive controller will be the simplification still further of the available design variables to be small enough in number to provide the engineer on the shop floor sufficient flexibility in tuning the behavior of the adaptive system in response to the observed properties or to *a priori* information about the plant system, while avoiding the requirement for the provision of too many *ad hoc* variable selections. This is a critical feature of any control design methodology which entertains pretensions of industrial applications.

The controller we are presenting here is a candidate general purpose robust adaptive controller based on ideas from predictive control. Following the presentation of the candidate, the platform of support will be detailed. The vote, however, must be left for the control populace.

7.2 The Certainty Equivalence Control Law

Here we discuss the aspects of the sequence of design steps to be taken in going from *a priori* knowledge and on-line estimated data provided by the identifier to the generation of a complete LQG controller. We assume that the control objective is embodied in the specification of an output tracking problem with a given reference trajectory, r_t , which is available to the control law N steps ahead of time. While the primary objective is reference tracking, control signal energy and peaking are also of concern as is the internal stability and robustness of the closed loop.

7.2.1 The Plant, Noise and Reference Models

Parameter identification schemes provide, naturally, a model of the plant system under study as evaluated by on-line data measurements. For us with LQG control in mind, this necessitates the passage from a parameter estimate vector, $\hat{\theta}$, for the plant model to an equivalent state-space description (3.1) and (3.2). That is, we define system matrices $F(\hat{\theta})$, $G(\hat{\theta})$, $H(\hat{\theta})$ from the estimate. We shall not be investigating specific methods for defining these matrices, i.e. realization procedures, because we shall be presenting a controller which, as much as possible, has closed loop properties independent of the state-variable realization chosen.

Also, these identification schemes may provide the user with a noise model describing the spectral properties of the measurement noise associated with the plant and its model. When such a noise model is available, one could take it as a suitable specification of the noise state matrices $F^d(\hat{\theta})$, $G^d(\hat{\theta})$, $H^d(\hat{\theta})$ of (3.40) and (3.41) used in the construction of the state estimator part of an LQG controller. Even when a noise modeling is explicit in the identification algorithm, however, there is no necessity that it be taken *holus bolus* into the control design. Instead, the users may substitute a noise model of their own.

In the process control industry there appear to be two central features militating against the use of these on-line estimated noise models. Firstly, the plant disturbances tend to be rather sporadic and so may not be well captured by such an identifier with a moderate forgetting factor. The inclusion of a noise model also increases the parameter dimension, which affects the convergence speed of the identifier. Secondly, it is often the case that essential features of the output disturbance are *a priori* known. This information could consist of knowledge of the spectral limits of the noise, or that the plant suffers step change or periodic load disturbances, for example. This information can then be used to construct (or concoct) a noise model

analogously to the procedure with a noise model provided by the identifier. Historical experience with modern adaptive control, such as it is, seems to indicate that *a priori* specification of the noise model is a preferred option amongst applications engineers. This notion has strong links to the inclusion of the Δ operator in the plant/noise model for GPC (see 2.5) with the intent of forcing the introduction of an integrator into the closed loop.

The reference signal, r_t , too may have a model associated with its evolution which allows us to take advantage of its availability for control signal construction a fixed number, N , of steps ahead of time. In Section 3.3 we have taken this model to be given by a delay line driven by a zero mean white noise input. Alternative models, taking into account expected sluggishness of references or non-zero offsets, may also be constructed by using, say, integrated random walk-type models or other constructions. Our point here is that these models should either be chosen by the designer to reflect likely specific features of the reference or else chosen to use the N -step-ahead property of the model (3.14).

From this point, we presume that the information passing from the identifier to the control design stage consists of the plant model matrices $F = F(\hat{\theta})$, $G = G(\hat{\theta})$, $H = H(\hat{\theta})$ and that the noise model matrices F^d , G^d , H^d are either specified by the identifier or *a priori* by the designer. The reference model matrices F^r , G^r , H^r are presumed stipulated by the designer. The composite state-space model available for LQG/LTR control design is then (see Chapter 3):

$$\begin{pmatrix} x_{t+1} \\ x_{t+1}^r \\ x_{t+1}^d \end{pmatrix} = \begin{pmatrix} F & 0 & 0 \\ 0 & F^r & 0 \\ 0 & 0 & F^d \end{pmatrix} \begin{pmatrix} x_t \\ x_t^r \\ x_t^d \end{pmatrix} + \begin{pmatrix} G \\ 0 \\ 0 \end{pmatrix} u_t + \begin{pmatrix} I & 0 & 0 \\ 0 & G^r & 0 \\ 0 & 0 & G^d \end{pmatrix} \begin{pmatrix} w_t \\ n_t \\ p_t \end{pmatrix} \quad (7.2)$$

$$\begin{pmatrix} y_t \\ r_t \end{pmatrix} = \begin{pmatrix} H & 0 & H^d \\ 0 & H^r & 0 \end{pmatrix} \begin{pmatrix} x_t \\ x_t^r \\ x_t^d \end{pmatrix} + \begin{pmatrix} I \\ 0 \end{pmatrix} q_t. \quad (7.3)$$

The definition of an LQG control law on the basis of this state model requires the additional specification of several design variables, namely the noise covariance matrices Q_o and R_o (see Section 3.4) and the cost criterion weighting matrices Q_c and R_c (see Section 3.2). The selection of these variables in a manner that is both consistent with our previous developments and that leads to a minimal (but sensible) number of free parameters for design is the object of the next two subsections.

To complete the adaptive control design, where the matrices F , G , H , and possibly F^d , G^d , and H^d , are estimated on-line, the step size gain γ , the filter $D(z)$ of the RLS identification algorithm and the reference input spectrum, Φ_r , need to be specified. This will be done in Section 7.3.

7.2.2 Kalman Filter Design

Normally, one would address the issue of LQ control law selection before broaching the subject of attempting to generate a state estimator. However, our viewpoint has been colored by the robustness material presented in Chapter 5, where robust LQG design in discrete time proceeds via the construction of the Kalman filter as the object capable of delivering closed loop stability robustness, and then to the control law whose specifications are determined by the requirements of Loop Transfer Recovery. Hence, we shall indulge our newfound prejudice by starting with the KF design.

Since the state, x_t^r , of the reference model need not be estimated, it suffices for Kalman filter design to consider the composite state description combining the noise and plant models, (3.43) and (3.44), which we rewrite here for the sake of the reader's comfort,

$$x_{t+1}^m = \begin{pmatrix} F & 0 \\ 0 & F^d \end{pmatrix} x_t^m + \begin{pmatrix} G \\ 0 \end{pmatrix} u_t + \begin{pmatrix} I & 0 \\ 0 & G^d \end{pmatrix} \begin{pmatrix} w_t \\ p_t \end{pmatrix}, \quad (7.4)$$

$$\triangleq F^m x_t^m + G^m u_t + L^m w_t^m \quad (7.5)$$

$$y_t = (H \quad H^d) x_t^m + q_t, \quad (7.6)$$

$$\triangleq H^m x_t^m + v_t^m. \quad (7.7)$$

Recall that x_t^m is the composite state defined in (3.42). From here one solves the filtering ARE,

$$\Sigma = F^m \Sigma F^{mT} - F^m \Sigma H^{mT} (H^m \Sigma H^{mT} + R_o)^{-1} H^m \Sigma F^{mT} + Q_o, \quad (7.8)$$

for Σ and then computes the KF gain M^F via

$$M^F = \Sigma H^{mT} (H^m \Sigma H^{mT} + R_o)^{-1} \triangleq \begin{pmatrix} M^x \\ M^d \end{pmatrix}. \quad (7.9)$$

The KF state estimator then may be constructed:

$$\begin{aligned} \hat{x}_{t+1}^m &= (I - M^F H^m) F^m \hat{x}_t^m + (I - M^F H^m) G^m u_t \\ &\quad + M^F y_{t+1}. \end{aligned} \quad (7.10)$$

The system matrices above are provided directly from the identifier so that those variables remaining to be specified before this KF design can be

carried through are the design covariance matrices Q_o and R_o . Referring to Section 3.5, we see that these covariance matrices are given by

$$\begin{aligned} Q_o &= L^m \mathbb{E}(w_t^m w_t^{mT}) L^{mT} = \begin{pmatrix} \mathbb{E}(w_t w_t^T) & 0 \\ 0 & G^d \mathbb{E}(p_t p_t^T) G^{dT} \end{pmatrix} \\ R_o &= \mathbb{E}(q_t q_t^T), \end{aligned}$$

where w_t is the process noise of the plant, and p_t and q_t are the process and measurement noises of the noise model, all assumed zero mean, white and mutually independent. Since p_t and q_t both drive the noise v_t corrupting the plant output, the noise coloring will change if we allow independent variation of $\mathbb{E}(p_t p_t^T)$ and $\mathbb{E}(q_t q_t^T)$. Therefore we take both covariances to be identical and absorb the potential difference into the definition of G^d . Further, we take

$$R_o = \rho I, \quad (7.11)$$

where ρ is a non-negative design constant which scales the measurement noise power versus the plant process noise power, $\mathbb{E}(w_t w_t^T)$, which we shall take to be constant.

In terms of plant state estimation, the comparison between $\mathbb{E}(w_t w_t^T)$ and ρI indicates that this represents a signal to noise ratio property for the plant output measurement, y_t . Separate scaling of both values achieves nothing that could not otherwise have arisen by selecting ρ and leaving $\mathbb{E}(w_t w_t^T)$ fixed. Thus, once $\mathbb{E}(w_t w_t^T)$ is specified, ρ becomes our sole design parameter for the KF component of LQG.

Harking back to Chapter 3 and the GPC controller interpreted as LQG, recall that the LQ control design problem depends centrally upon the system matrices F and G , themselves, which depend greatly upon the particular state-space coordinate basis chosen for the realization. One method to force coordinate system independence of the controller was to select $Q_c = H^T H$, which includes the output matrix into the LQ criterion and also reflects a tracking objective. For the Kalman filter here we choose the dual of this scheme,

$$\mathbb{E}(w_t w_t^T) = G G^T,$$

which causes the KF to exhibit input–output properties which will remove the state coordinate system dependence from the LQG controller.

Summarizing this Kalman filter design stage, we choose:

1. From the parameter vector $\hat{\theta}$ provided by the identifier, construct a realization $[F, G, H]$ for the plant.
2. Either from the identifier or from *a priori* information, construct a disturbance model realization $[F^d, G^d, H^d]$.

3. Select the design variables R_o and Q_o as follows:

$$R_o = \rho I \quad (7.12)$$

$$Q_o = \begin{pmatrix} GG^T & 0 \\ 0 & \rho G^d G^{dT} \end{pmatrix}. \quad (7.13)$$

4. Solve the filtering ARE (7.8) and compute the KF gain, M^F , via (7.9).

7.2.3 LQ State-variable Feedback Design

Given the plant, noise and reference model specifications from the identifier and the user, we may begin with the super-state model (7.2) embodying all these subsystems and proceed to design the LQ state-variable feedback control law analogously to the KF design above. We use the following tracking criterion in which the reference, r_t , is assumed to be known up to N steps ahead,

$$J(N, x_t) = \mathbb{E} \sum_{j=0}^{N-1} \left\{ (y_{t+j+1} - r_{t+j+1})^T (y_{t+j+1} - r_{t+j+1}) + \lambda u_{t+j}^T u_{t+j} \right\}. \quad (7.14)$$

This corresponds to the following design choices for the weighting matrices, \bar{Q}_c and \bar{R}_c , on the super-state and control of the model (7.2),

$$\bar{Q}_c = \begin{pmatrix} H^T \\ -H^r{}^T \\ H^d{}^T \end{pmatrix} (H \quad -H^r \quad H^d), \quad \bar{R}_c = \lambda I. \quad (7.15)$$

Compare this with Section 3.6. Further, in line with the stability arguments developed in Chapter 4 and the robustness properties in Chapter 5, we now use the corresponding infinite horizon regulator gains, which are associated with allowing the cost horizon in (7.14) to tend to infinity while preserving the property that the reference is only available N steps ahead. These are obtained by using the super-state model (7.2) and the design choices above in (7.15) and by specializing the formulae of Sections 3.6 and 3.3. The LQ regulator ARE yields

$$P^{11} = F^T P^{11} F - F^T P^{11} G (G^T P^{11} G + \lambda I)^{-1} G^T P^{11} F + H^T H, \quad (7.16)$$

which is connected with the solution of the output regulation problem for the plant system with neither noise model nor reference model. From this we define the regulation gain

$$K^x = -(G^T P^{11} G + \lambda I)^{-1} G^T P^{11} F, \quad (7.17)$$

and the ancillary gain matrices, K^r and K^d , via the Lyapunov equations

$$P^{12} = (F + GK^1)^T P^{12} F^r - H^T H^r \quad (7.18)$$

$$P^{13} = (F + GK^1)^T P^{13} F^d + H^T H^d, \quad (7.19)$$

and

$$K^r = -(G^T P^{11} G + \lambda I)^{-1} G^T P^{12} F^r, \quad (7.20)$$

$$K^d = -(G^T P^{11} G + \lambda I)^{-1} G^T P^{13} F^d. \quad (7.21)$$

The LQ design component of the adaptive system thus consists of the following steps:

1. Solve the LQ regulator ARE (7.16) for P^{11} with $Q_c^{11} = H^T H$ and $R_c = \lambda I$.
2. Compute the LQ regulator gain K^x via (7.17).
3. Solve the Lyapunov equations (7.18) and (7.19) for P^{12} and P^{13} .
4. Compute the gains K^r and K^d via (7.20) and (7.21).

7.2.4 The LQG Controller

The LQG controller is now specified by the plant matrices $[F, G, H]$, the reference and disturbance model matrices $[F^r, G^r, H^r]$ and $[F^d, G^d, H^d]$, and the various gain matrices M^F, K^x, K^r, K^d . Denote

$$\begin{aligned} K^m &\triangleq (K^x \quad K^d) \\ \tilde{F} &\triangleq (I - M^F H^m)(F^m + G^m K^m) \\ &= \begin{pmatrix} I - M^x H & -M^x H^d \\ -M^d H & I - M^d H^d \end{pmatrix} \begin{pmatrix} F + GK^x & GK^d \\ 0 & F^d \end{pmatrix}. \end{aligned}$$

We now take the estimator state equation (7.10) and substitute

$$u_t = K^m \hat{x}_t^m + K^r x_t^r, \quad (7.22)$$

to yield

$$\hat{x}_{t+1}^m = \tilde{F} \hat{x}_t^m + (I - M^F H^m) G^m K^r x_t^r + M^F y_{t+1}.$$

Therefore,

$$\hat{x}_t^m = (zI - \tilde{F})^{-1} (I - M^F H^m) G^m K^r x_t^r + (zI - \tilde{F})^{-1} M^F y_{t+1}.$$

Substituting this last expression into (7.22) produces the transfer function of the LQG feedback controller,

$$\begin{aligned}
 u_t = & \left[I + K^m(zI - \tilde{F})^{-1}(I - M^F H^m)G^m \right] K^r(zI - F^r)^{-1}G^r r_{t+N} \\
 & + K^m M^F y_t + K^m \tilde{F}(zI - \tilde{F})^{-1}M^F y_t.
 \end{aligned} \tag{7.23}$$

It is important to observe that the controller gain matrices K^x and K^d and the observer gain matrices M^x and M^d all have an effect on the reference r_t to input u_t transfer function.

This controller may equally well be written in state-variable form as follows:

$$\begin{aligned}
 \begin{pmatrix} x_{t+1}^r \\ x_{t+1}^c \end{pmatrix} &= \begin{pmatrix} F^r & 0 \\ (I - M^F H^m)G^m K^r & \tilde{F} \end{pmatrix} \begin{pmatrix} x_t^r \\ x_t^c \end{pmatrix} + \begin{pmatrix} G^r & 0 \\ 0 & \tilde{F} M^F \end{pmatrix} \begin{pmatrix} n_t \\ y_t \end{pmatrix} \\
 u_t &= (K^r \quad K^m) \begin{pmatrix} x_t^r \\ x_t^c \end{pmatrix} + K^m M^F y_t.
 \end{aligned} \tag{7.24}$$

In this latter form one sees how the adaptive control law may be carried along in a recursive manner. As the identifier produces new parameter values for the system matrices F , G , H , one computes associated gains M^F , K^x , K^r , K^d and the new controller matrix \tilde{F} . These matrices are updated at each step by re-solving the ARE but the state x_t^c is carried over from time instant to time instant, providing the storage between parameter changes.

7.3 The System Parameter Identifier

The aspect of our candidate adaptive controller complementary to the control law statement is the presentation of the identification module which operates in closed loop and provides, to the control design section, estimates of the system model, as well as perhaps a noise model if this is desired. It is now our task to make precise the design-variable choices for the RLS scheme.

Recall the recursive estimate update from (6.15),

$$\hat{\theta}_t = \hat{\theta}_{t-1} + \gamma R_t^{-1} \psi_t^f(\hat{\theta}_{t-1}, \eta) \epsilon_t^f(\hat{\theta}_{t-1}, \eta), \tag{7.25}$$

where $\psi_t^f(\hat{\theta}_{t-1}, \eta)$ is the filtered regressor for the algorithm, η are the parameters of the filter, γ is the step size of the adaptation (here chosen to be constant), and R_t can be chosen $R_t = I$ for a gradient algorithm or iterated as

$$R_t = R_{t-1} + \gamma [\psi_t^f(\hat{\theta}_{t-1}, \eta) \psi_t^{fT}(\hat{\theta}_{t-1}, \eta) - R_{t-1}], \tag{7.26}$$

for the recursive Gauss-Newton algorithm. This algorithm stems from the writing of a plant output predictor in pseudolinear regression form

$$\hat{y}_{t|t-1}(\theta) = \phi_t^T(\theta)\theta, \quad (7.27)$$

and computing the θ value, which minimizes the filtered prediction error squared criterion (6.10),

$$V_N(\theta, \eta) = \frac{1}{N} \sum_{t=1}^N \left(D(z, \eta) \left(y_t - \hat{y}_{t|t-1}(\theta) \right) \right)^2. \quad (7.28)$$

The (pseudo-)regressor for the algorithm is then specified by

$$\psi_t^f(\hat{\theta}_{t-1}, \eta) = \phi_t^f(\eta) = D(z, \eta)\phi_t. \quad (7.29)$$

What are the available design choices for such an identifier?

1. The class of plant models specifies the range of potential $\hat{P}(z, \theta)$ deriving from the estimator.
2. The class of noise models, $\hat{H}(z, \theta)$, affects the information produced by the identifier and, via the formula (6.33), the frequency weighting associated with the plant fit.
3. The filtering, $D(z, \eta)$, introduced into the selection criterion (7.28), alters the frequency weighting again as shown in (6.33). We remark that $D(z, \eta)$ and $\hat{H}(z, \theta)$ affect the plant-model mismatch ΔP only through the ratio D/\hat{H} .
4. The spectral properties of the reference signal, r_t , if they are available to the designer, allow him or her to distribute the identification weighting in (6.33).
5. The algorithm step size, γ , determines the convergence and tracking speed and affects the ability to appeal to averaging and convergence theories of Chapter 6.
6. The algorithm structure, e.g. gradient or Gauss-Newton, affects the convergence speed and numerical complexity of the identifier.
7. Other features specific to adaptive control identifiers discussed briefly in Chapter 6, such as normalization, relative deadzones, leakage and projection, operate in such a fashion as to constrain the algorithm behavior in times of manifestly incorrect estimates, closed loop instability or quiescence of the signal excitation in the loop.

For our adaptive controller we have the following selections:

1. We take the plant model structure as having been specified by the users in the light of their system knowledge. This is not a feature which will be universally available for all plants.
2. We presume that the noise model is fixed *a priori* by the designer to reflect known disturbance features or else is selected to be unity. If $H(z) = 1$, then we have an Output Error style of model, although we do not advocate the use of Output Error identifiers over Equation Error methods because of convergence issues to be mentioned in Chapter 8.
3. The filter, $D(z)$, which filters the signals before they are passed to the identifier, is chosen according to the following guidelines, which are treated in a more general context in Chapter 8:
 - the rolloff in the frequency response of D should roughly correspond to the rolloff of the reference trajectory, Φ_r , in order that the closed loop performance be enhanced by adaptation.
 - if *a priori* knowledge is available about the value $\mathcal{G}(z)$ of the cascade of the *actual* plant and the nominal controller, then $D(z)$ should be chosen to allow some measure of identification up to the gain crossover, if this is known approximately beforehand. This allows the adaptation to enhance the stability robustness.
 - Otherwise we choose $D(z) \approx H(z)$ for frequencies below these cutoffs. This reduces the noise perturbation (bias, essentially) of the equation error identifier.
4. The reference spectrum applied during periods of adaptation is chosen to consist of the desired tracking trajectory suitably modified to ensure persistence of excitation over a bandwidth extending to a reasonable estimate of the desired closed loop bandwidth and at a level which causes Φ_r to dominate Φ_v over this interval. $D(z)$ may then be chosen to predistort these low frequency signals before the identifier to remove excessive coloring.
5. The step size, γ , is chosen to be sufficiently small so that the adaptation after large initial transients proceeds at a time scale considerably slower than the closed loop dynamics. For tracking time variations, γ must be chosen sufficiently large that plant parameter variations occur at a time scale considerably slower than that of adaptation.

6. A Gauss-Newton algorithm is chosen; that is, RLS is chosen rather than a gradient algorithm.
7. Such other features as are deemed desirable and/or necessary to safety-jacket the algorithm in its transient or quiescent phases are specified by the user. We refer to Section 6.5 for a brief description of the most commonly used tools.

Here then is our candidate robust adaptive predictive controller. The astute reader will already have divined the reasoning behind many of the choices made above based both on the earlier theory and on our poorly disguised prejudices revealed along the way. We now attempt to justify this candidate by presenting less of its ideology and more of its pragmatic policies.

7.4 The Candidate — A Summary

The complete candidate adaptive controller consists of the interconnection of the following components:

- An RLS parameter identifier, with
 - slow adaptation, i.e. small but not infinitesimal step size γ ,
 - fixed noise model $H(z)$ based on prior knowledge, if any,
 - filters $D(z)$ rolling off outside the reference bandwidth and around the desired closed loop bandwidth, which predistort to account for the noise model and reference over this band,
 - a reference signal, r_t , consisting of the desired trajectory plus (if necessary) a perturbation which contains sufficient spectral support to dominate the output disturbance over the closed loop bandwidth,
 - additional features such as normalization, deadzones and leakage as necessary.
- An LQG controller, with
 - a Kalman filter design based on a coupled plant and noise model, with measurement noise power determined by ρ and plant process noise covariance GG^T , yielding KF gain M^F ,
 - an LQ tracking control law design based on coupled plant, noise and reference models with state weighting being $H^T H$ and control weighting λ , yielding feedback gain K ,

- a state-variable formulation (7.24) amenable to adaptive control generation of the control input signal even as the parameter estimates change.

7.5 The Platform

Now we cast back to the overall requirements of our adaptive control scheme of providing closed loop stability and good performance as measured by a minimum variance (LQ) objective function. The subsequent objectives were then to keep the control signal reasonably well bounded and to make the closed loop system robust to slight model imperfections so that slow variations and small nonlinearities do not upset the control performance too greatly. It is in the light of these objectives that we must assess the candidate controller.

7.5.1 LQG Controller Properties

Since the global aim of the controller is to achieve minimum variance output tracking, the choice of LQ control gain as λI , with λ small compared to the state weighting $H^T H$, is well consistent with this goal. Since LQG control is chosen, stability of the nominal system is assured. Also, through manipulation of the design variable λ , one may balance the achievement of the control objective against control signal magnitudes, thereby effecting the design compromise between these issues. But further, this particular choice of LQ criterion weightings is precisely that advocated by the Loop Transfer Recovery Theory of LQG in Chapter 5 in order that the closed loop stability robustness be assured of being as close as possible to that of the Kalman predictor without a control input. Thus this controller design attempts to produce a plant model/controller cascade $\hat{\mathcal{G}}(z)$ which, *ipso facto*, satisfies

$$\sigma[1 + \hat{\mathcal{G}}(z)] > \bar{\delta},$$

where $\bar{\delta}$ is related to the open loop plant and to ρ (see Theorem 5.6). The controller, operating in the neighborhood of the actual plant, exhibits some of the desirable features required of our adaptive solution. We reiterate not only our caveat of Chapter 5, that LTR is only guaranteed to work for minimum phase plants, but also the experience that often LTR can help in achieving reasonable robustness for many other plants. Certainly, it provides a logical manner in which to search a range of potentially robust controllers.

As a further benefit of a robust control design, as opposed to, say, an algebraic design such as pole-positioning, the closed loop system in the neighborhood of its $\hat{\theta}$ -parametrized operating point should have not only stability

but also a smoothness of variability of loop signals. That is, the sensitivity of the closed loop to small parameter changes in the locality of $\hat{\theta}$ will be small. This feature will be of importance in appealing to the estimator properties developed in the previous chapter, and to be reassessed next.

7.5.2 RLS Identifier Properties

Next we study whether the other half of the adaptive controller supports these features. The robustness study of Chapter 5 required that for closed loop stability we have

$$\bar{\sigma}[L^{-1}(z) - I] < \min(\sigma[I + \hat{\mathcal{G}}(z)], 1), \quad \forall |z| \in \Omega, \quad (7.30)$$

where $L(z)$ is the multiplicative deviation of the nominal or model system, $P(z, \hat{\theta})$, from the actual or true plant system, $P(z)$, and $\hat{\mathcal{G}}(z)$ is the cascade mentioned above. Thus $L(z)$ is a function of the θ choice and, in fact, may be written as (5.17),

$$\begin{aligned} L^{-1}(z) - I &= \mathcal{G}^{-1}(z)[\hat{\mathcal{G}}(z) - \mathcal{G}(z)] \\ &= P^{-1}(z)[\hat{P}(z, \hat{\theta}) - P(z)]. \end{aligned} \quad (7.31)$$

Therefore, in order to achieve closed loop robustness under this formulation, we need to ensure that the *relative* plant model error is kept small by the identifier.

To examine the issue of the size of the error in (7.31) we appeal to the examination of Chapter 6, which expressed the prediction error minimization criterion in frequency response terms (see (6.32), (6.33)). Specifically, under the assumption of a fixed noise model $H(z)$ as advocated above, we have the following closed loop prediction error criterion for the plant model parameter, θ^* , about which the identifier parameter $\hat{\theta}$ remains,

$$\begin{aligned} \theta^* &= \arg \min_{\theta \in \mathcal{D}_\theta} \int_{-\pi}^{\pi} \left[|\Delta P(e^{j\omega}, \theta)|^2 \frac{|F_1(e^{j\omega}, \rho)|^2}{|1 + F_2(e^{j\omega}, \rho)P(e^{j\omega})|^2} \Phi_r(\omega) \right. \\ &\quad \left. + (|\Delta P(e^{j\omega}, \theta)|^2 \frac{|F_2(e^{j\omega}, \rho)|^2}{|1 + F_2(e^{j\omega}, \rho)P(e^{j\omega})|^2} + 1) \Phi_v(\omega) \right] \frac{|D(e^{j\omega}, \eta)|^2}{|\hat{H}(e^{j\omega}, \theta)|^2} d\omega. \\ &= \arg \min_{\theta \in \mathcal{D}_\theta} \int_{-\pi}^{\pi} \left[|\Delta P|^2 (|\hat{W}_1|^2 \Phi_r + |\hat{W}_2|^2 \Phi_v) |\hat{W}_A|^2 + \Phi_v \right] \frac{|D|^2}{|H|^2} d\omega, \end{aligned} \quad (7.32)$$

where

$$\hat{W}_1(z) = \frac{F_1(z, \rho)}{1 + F_2(z, \rho)\hat{P}(z)} \quad (7.33)$$

$$\hat{W}_2(z) = -\frac{F_2(z, \rho)}{1 + F_2(z, \rho)\hat{P}(z)} \quad (7.34)$$

$$\hat{W}_A(z) = \frac{1 + F_2(z, \rho)\hat{P}(z)}{1 + F_2(z, \rho)P(z)}. \quad (7.35)$$

From (7.32) we see that the transfer function $\hat{W}_1(z)\hat{W}_A(z)$ frequency weights the reference signal spectrum in the model fitting criterion, while the noise spectrum is weighted through $\hat{W}_2(z)\hat{W}_A(z)$ but also appears on its own as a disturbing influence upon the model fitting criterion, since it does not always multiply the plant model error. This equation indicates how the noise signal, v_t , may be beneficial in helping to excite the plant via the feedback control, thereby improving the model fit, even though it compromises this fit through the presence of its unpredictable part in the system output. Analysis of the structure of $\hat{W}_2(z)\hat{W}_A(z)$ shows the rôle of both the Kalman filter design and the LQ control in determining the influence of v_t .

From our point of view, however, the use solely of the noise to power the model fit is both difficult to analyse (although this can be done) and somewhat perverse, because we believe that the adaptive control paradigm is best enunciated when sufficient information levels are present in the plant signals. We therefore assume from this point that, during periods of adaptation, the reference signal dominates the plant measurement noise within the bandwidth of the closed loop plant,

$$\Phi_r(\omega) \gg \Phi_v(\omega).$$

After all, this is just what the suggested selection of reference is for adaptation. Thus the model fitting criterion will be altered to be approximately

$$\theta^* \approx \arg \min_{\theta \in D_\theta} \int_{-\pi}^{\pi} |P(e^{j\omega}) - P(e^{j\omega}, \theta)|^2 \frac{|\hat{W}_1(e^{j\omega})\hat{W}_A(e^{j\omega})D(e^{j\omega})|^2}{|H(e^{j\omega})|^2} \Phi_r(\omega) d\omega,$$

at least over the reference/closed loop bandwidth. The selection of D rolling off sharply after this permits us to extend the integral above to the full $(-\pi, \pi]$ range.

For an LQG controller, the expression of $\hat{W}_1(z)$ has been computed in (6.37) and we see that the parameter identification criterion incorporates an overall frequency weighting determined by

$$\frac{\hat{W}_1(z)\hat{W}_A(z)D(z)}{H(z)} = [1 - K^x(zI - \hat{F})^{-1}\hat{G}]^{-1} \times K^r(zI - F^r)^{-1}G^r z^{N+1}$$

$$\times \hat{W}_A \times \frac{D}{H} \quad (7.36)$$

$$\triangleq F_1(z) \times F_2(z) \times F_3(z) \times F_4(z). \quad (7.37)$$

We shall examine these four terms in turn, studying their particular forms in the case of our particular candidate controller choices.

The Frequency Weighting $F_1(z)$

In our case this transfer function is directly the inverse of the LQ control law return difference,

$$F_1(z) = [1 - K^x(zI - F)^{-1}G]^{-1}. \quad (7.38)$$

To expose the nature of this transfer function in our application we return with deference to the return difference equality (5.22), which stems directly from the regulation ARE (7.16),

$$\begin{aligned} \lambda + G^T(z^{-1}I - F)^{-T}H^T H(zI - F)^{-1}G \\ = [I - K^x(z^{-1}I - F)^{-1}G]^T(G^T P^{11}G + \lambda)[I - K^x(zI - F)^{-1}G], \end{aligned} \quad (7.39)$$

where we have substituted the selections $R_c = \lambda I$ and $Q_c = H^T H$. Clearly, this may be rewritten in more familiar terms:

$$\lambda + P(z^{-1}, \hat{\theta})P(z, \hat{\theta}) = F_1^{-1}(z^{-1})\Lambda F_1^{-1}(z), \quad (7.40)$$

with Λ being the positive definite matrix $(G^T P^{11}G + \lambda)$.

Now recall that our operating conditions in this adaptive controller are such that we desire minimum variance regulation. That is, we attempt as much as is possible to take λ small in relation to the output weighting in the LQ control criterion. This causes the return difference, F_1^{-1} , to have poles identical to the plant poles and to have zeros at the stable plant zeros and at the inverse of the unstable plant zeros (see Section 5.2). Thus, we have, for small λ ,

$$|F_1(e^{j\omega})| \approx |P(e^{j\omega}, \hat{\theta})|^{-1} \times \beta, \quad \forall \omega \in [-\pi, \pi], \quad (7.41)$$

where β is a constant.

The Frequency Weighting $F_2(z)$

Recall that $F_2(z) = K^r(zI - F^r)^{-1}G^r z^{N+1}$, with K^r derived in the solution of the LQ tracking criterion with a reference model, $[F^r, G^r, H^r]$, via (3.23)

and (3.25). For the reference model matrices chosen in (3.14), one may explicitly evaluate $F_2(z)$, which we briefly demonstrate. Clearly, the term z^{N+1} plays no part in the effect of $F_2(z)$ as a frequency weighting on $|z| = 1$.

Equation (3.25) can be rewritten as

$$P^{12} = (F + GK^x)^T P^{12} F^r - H^T H^r,$$

which in turn possesses the obvious solution,

$$P^{12} = - \sum_{j=0}^{N-1} (F + GK^x)^{jT} H^T H^r F^{r^j}.$$

Here we have used the nilpotence of F^r to terminate the sum. Then,

$$K^r = (G^T P^{11} G + \lambda)^{-1} G^T \sum_{j=0}^{N-1} (F + GK^x)^{jT} H^T H^r F^{r^{j+1}},$$

and, hence, ignoring the z^{N+1} ,

$$\begin{aligned} F_2(z) &= K^r (zI - F^r)^{-1} G^r \\ &= \text{const} \times \sum_{j=0}^{N-1} G^T (F + GK^x)^{jT} H^T H^r F^{r^{j+1}} (zI - F^r)^{-1} G^r \\ &= \text{const} \times \sum_{j=0}^{N-1} G^T (F + GK^x)^{jT} H^T z^{-N+j}. \end{aligned} \quad (7.42)$$

One sees from (7.42) that $F_2(z)$ consists of the first $N-1$ impulse response terms of the closed loop plant written in reverse order. Thus for large values of N one has $F_2(z)$ approximately equal to the closed loop reference to output transfer function, while for $N = 1$, i.e. no lookahead, one has $F_2(z) = 1$. Our argument here is that, for LQ tracking problems requiring long lookahead, one would normally expect the controller to achieve a greater closed loop bandwidth than the open loop plant. Therefore the ability to identify the open loop gain crossover point is unimpaired by allowing $F_2(z) \neq 1$. We admit that these statements are a little loose and woolly but our assertion is that $F_2(z)$ does not upset the other identification weightings much.

The Frequency Weighting $F_3(z)$

The weighting $F_3(z)$ is $\hat{W}_A(z)$ which, in turn, is the ratio of the achieved sensitivity function to the designed sensitivity function. We shall reconsider this more fully in the next chapter but, for the moment, simply observe that,

within the passband of the achieved or designed plant/controller cascade (these should be close), this function has a value very close to 1. Therefore, its effect on the model fitting should be minimal. Outside this band, the rolloff from D should obviate the need to consider \hat{W}_A 's behavior further.

The Frequency Weighting $F_4(z)$

The transfer function $F_4(z)$ is given simply as $D(z)/H(z)$, and with our choices for $D(z)$ we have this function being close to one or $\Phi_r^{-\frac{1}{2}}$ over the closed loop bandwidth. Outside this bandwidth, F_4 rolls off rapidly.

The Total RLS Frequency Weighting

Combining these four frequency weightings for the RLS based plant identifier, we see that the effective closed loop identification criterion asserted by the candidate adaptation in league with the candidate controller is that the identified parameter θ will approach and stay near the value θ^* which satisfies, for representative choices of variables above,

$$\theta^* = \arg \min_{\theta \in D_\theta} \int_{-\pi}^{\pi} \frac{|P(e^{j\omega}) - P(e^{j\omega}, \theta)|^2}{|P(e^{j\omega}, \theta)|^2} \Phi_r(\omega) |D'(e^{j\omega})|^2 d\omega, \quad (7.43)$$

where D' is the remaining weighting in D after removal of H effects, which effectively limits the regime of modeling to focus on the reference bandwidth, \mathcal{W}_r .

Now making the connection to our robustness requirement (7.1) via the identity (7.31) that

$$L^{-1}(z, \hat{\theta}) - 1 = (P(z) - P(z, \hat{\theta})) P^{-1}(z),$$

we see that

$$\theta^* \approx \arg \min_{\theta \in D_\theta} \int_{-\mathcal{W}_r}^{\mathcal{W}_r} [L^{-1}(e^{j\omega}, \theta) - 1]^2 \frac{|P(e^{j\omega})|^2}{|P(e^{j\omega}, \theta)|^2} \Phi_r(\omega) d\omega. \quad (7.44)$$

That is to say, the closed loop identification of the plant operating under the LQG control law will automatically find that value of θ which provides a model close to the plant but which minimizes the ℓ_2 norm of $(L^{-1} - 1)$ on the unit circle over the closed loop bandwidth.

7.5.3 Manipulation of the Candidate Controller

Our thesis has been that the candidate adaptive controller yields a closed loop control law which attempts to meet the requirement that $|1 + \mathcal{G}| > \alpha$ on the unit circle for some $\alpha > 0$. Simultaneously, the identifier, via the implicit frequency weighting induced by the control law, produces a model that satisfies the requirement that $|L^{-1} - 1|$ is maintained small on the unit circle in a suitable frequency region. It is by these methods that the overall controller achieves its robust stability, while attempting to meet the minimum variance objective. This is the synergism of control law and identifier of which we spoke earlier. In this section we shall next move on to consider how this robustness goal may be further enhanced.

The robustness of the control law stage is achieved by two features. Firstly, the Kalman filter is designed to have a degree of robustness, and secondly, the LQG/LTR methodology is used to strive to achieve a closed loop having the equivalent KP robustness.

One procedure to influence the robustness at this point would be to attempt to select the design variable ρ to maximize the KF robustness. A cursory analysis of the KP EDR (5.27) shows that, at least for stable plants, a more robust KP is obtained by choosing a bigger R_o than necessary. This causes a more cautious predictor and therefore yields a KF having slower response and better robustness. Clearly, one is sacrificing the closed loop performance for the sake of ameliorating the robustness. Nevertheless, increasing ρ remains as the primary tool in the hands of the designer to manipulate the closed loop performance.

The design variable complementary to ρ is the LQ control weighting λ . The LQG/LTR theory dictates that the selection of λ should be made after the selection of ρ and should, indeed, be to take λ as small as is feasible before the control signal in the closed loop begins to fail to comply with its design limits. Thus ρ is chosen first and then a suitably small value for λ is sought.

The remaining method for the manipulation of the closed loop robustness is the incorporation of *a priori* knowledge about the plant. Here we see how the specification of the filter $D(z)$ can achieve this. Recall that the robustness criterion (7.1) does not require that maximum over all z of the left hand side be dominated by the minimum of the right hand side. (This is effectively what is achieved by the candidate controller.) Rather, the inequality need only be satisfied at each value of z . Thus the need is not for $L^{-1} - 1$ to be small everywhere on the unit circle, but to be small when $1 + \mathcal{G}$ is small. To encourage this preferential fitting, we suggest that, if an estimate of $1 + \mathcal{G}$ is known *a priori*, then filtering by D possessing a factor

$(1 + \mathcal{G})^{-1}$ is appropriate.

7.6 Computer Studies and Examples

We now present something which is inevitable in the course of a work such as this — computer studies (more or less) verifying the properties claimed for our adaptive controller. This section is divided into two subsections, similar to the presentation of examples in Chapter 3, where we consider a simple, relatively easily controlled non-minimum phase plant (the Gedankenexample of Chapter 1 actually) and then treat the more intransigent Working Example, which is a much more difficult proposition.

7.6.1 The Gedankenexample Revisited

Recall that the system of the Gedankenexample is

$$P(z) = \frac{z^{-1} + 1.2z^{-2}}{1 - 1.6z^{-1} + 0.68z^{-2}}, \quad (7.45)$$

which is a low-pass, stable, but non-minimum phase system of second order. The model set, $\{\hat{P}(z, \theta)\}$, consists of first order plants parametrized by two parameters,

$$\hat{P}(z) = \frac{bz^{-1}}{1 + az^{-1}}. \quad (7.46)$$

In Section 1.4 the reader was asked to trace through a Gedankenexperiment in which a sequence of closed loop identifications and control designs was made that demonstrated a tendency for the controller to excite unduly the high frequency modes of the undermodeled plant. This caused a drift in the identified parameters to fit models at higher and higher frequencies, eventually leading to instability because of the unsuitability of the high frequency fit for controller design. This Gedankenexperiment might realistically be called a preadaptive controller since the identifier never really runs in an adaptive closed loop. We shall now show that the truly adaptive controller performs in essentially the same fashion.

We took the above plant, $P(z)$, and model set of first order systems. Coupled to this we chose a control law design which was LQG with

$$\lambda = 0.1, \quad \rho = 0.01.$$

A tapped delay line reference model was taken with $N_r = 3$. There was no measurement noise, and a reference signal was selected consisting of one cycle of a square wave of width 30 followed by one cycle of a square wave of




Figure 7.1: System output, Gedankenexample with $D(z) = 1$

width 10, repeated. The identifier was an equation error scheme with step size $\gamma = 5 \times 10^{-2}$. This identifier was run with the plant in open loop for the first 50 time steps, after which the loop was closed.

The first experiment was performed without the inclusion of the filter D , i.e. $D = 1$. The resulting system output is shown in Figure 7.1. The parameters are displayed in Figure 7.2. The key feature of these plots is precisely that phenomenon presaged in Section 1.4. That is, the controller has the effect of increasing the system bandwidth in closed loop, thereby forcing the parameters to begin an inexorable drift towards high frequency fits which produce unstable controllers. This drift may be seen in Figure 7.2. Recall that the identifier operates in open loop up to time 50, so the drift is evident from this time until the explosion of the signals. What really is surprising about this experiment is that the adaptive controller misbehaves so rapidly.

If the root cause of plant misbehavior is the induced drift towards higher frequencies, then our analyses (especially of Chapter 6) suggest that the insertion of $D(z)$ nonunity should provide a means to alleviate the problem. Consequently, demonstrating great bravado and faith in our own methods we chose $D(z)$ to be a second order Butterworth filter with a cutoff frequency at 0.1 radians. The resultant signals are displayed in Figure 7.3 (system output plus reference), Figure 7.4 (identified parameters), and Figure 7.5 (system




Figure 7.2: Identified parameters for Gedankenexample with $D(z) = 1$




Figure 7.3: System output, Gedankenexample with $D(z)$ second order Butterworth filter with cutoff at 0.1 radians




Figure 7.4: Identified parameters, Gedankenexample with $D(z)$ second order Butterworth filter with cutoff at 0.1 radians

control input).

There are several conclusions to be drawn from these simulations. Firstly, as is clear, stability and reasonable performance are achieved in fairly short time. The system signals show the strong ringing typical of second order systems driven by rather high control gains — this is the source of the poor modeling with $D = 1$. Nevertheless, with a low-pass $D(z)$ it is possible to ensure that the modeling only takes place at lower frequencies, which ameliorates the parameter drift problems. Other experiments were performed, such as moving the cutoff of D from 0.1 radians to 0.3 radians. This had the predictable effect of destabilizing the adaptive loop, albeit somewhat more slowly than with $D = 1$. Similarly, a very small value (10^{-10}) was tried for λ and D with cutoff at 0.1 radians. The result was again instability.

From this simple example we see that much of our intuition has been supported and, indeed, one is able to predict what measures will assist stability of the adaptive loop and what actions will jeopardize stability. We shall next treat the recidivistic Working Example.




Figure 7.5: System input, Gedankenexample with $D(z)$ second order Butterworth filter with cutoff at 0.1 radians

7.6.2 The Working Example Revisited

Recall the plant for the Working Example of Sections 3.7 and 5.5

$$P(z) = \frac{-0.05359 z^{-1} + 0.5775 z^{-2} + 0.5188 z^{-3}}{1 - 0.6543 z^{-1} + 0.5013 z^{-2} - 0.2865 z^{-3}}. \quad (7.47)$$

This plant, being non-minimum phase, presents difficulties for LQG control and for the LTR path to robustness. These have been demonstrated earlier. Our task now is to endeavor to produce an adaptive controller which can be tuned sensibly to perform well with this plant.

We begin by showing the best achievable LQG performance for this plant under our experimental conditions. Thus, we designed an LQG controller for the Working Example plant using the exact knowledge of its transfer function. This controller was designed with $\lambda = 0.01$, $\rho = 0.01$, $N_r = 10$ (the reference model lookahead), $\gamma = 5 \times 10^{-2}$, neither process nor output noise, and the same reference as above.

The resulting adaptive controller's response is shown in Figure 7.6. Since P is modeled exactly by a \hat{P} , the parameters converge precisely to their correct values rather quickly. Therefore the system response after time 170 approximately represents the best performance of the LQG design with these design-variable values. Subsequent behavior should be compared to this.




Figure 7.6: System output, Working Example with exact model

Also note the anticipatory nature of the controller which is built into our reference model.

Our next experiment was to restrict the class of models $\{\hat{P}\}$ to the set of second order plants,

$$\hat{P}(z) = \frac{b_1 z^{-1} + b_2 z^{-2}}{1 + a_1 z^{-1} + a_2 z^{-2}}.$$

We then performed trials with the above experimental conditions and design variables.

With the choice of the identifier filter $D(z) = 1$, the system performance was, surprisingly, not unstable, but did demonstrate significantly worse tracking and overshoot than the ideal case above. The output is shown in Figure 7.7, where several high frequency artefacts appear in the system response. These are presumably due to two main factors: the attempt for the second order model to fit the high frequency behavior of a third order (nasty) plant, and the appearance of adaptation dynamics in the loop response which indicates the sensitivity of this achieved closed loop to the parameter values. These parameters are shown in Figure 7.8. While the variations in the parameters are not particularly great, their effect upon the system response is large. This variability of the system is further evidenced by Figure 7.9, which depicts the estimated plant step response as




Figure 7.7: System output, Working Example with $D(z) = 1$, second order model




Figure 7.8: Identified parameters, Working Example with $D(z) = 1$




Figure 7.9: Evolution of the identified plant step response for the Working Example with $D = 1$

it evolves under adaptation. The non-minimum phase feature is captured by the negative going step response but the variation with input signal is marked.

To round out our presentation our final example is the Working Example with a choice of filter $D(z)$, which is in tune with the suggested *modus operandi* of this chapter. The D selected is again a Butterworth low-pass filter of second order but with cutoff at 0.3 radians. With the insertion of such a filter into the identifier we produce the response illustrated in Figure 7.10. Figure 7.10 shows the achievement of a control performance very close to the ideal circumstance of exact modeling. This is essentially because the model fit of \hat{P} to P is best in the region of reference power. The parameters are displayed in Figure 7.11, indicating about as much variability as earlier with $D(z) = 1$, but now the closed loop system is more robust to these parameter changes, as is further demonstrated by plotting the identified plant step response in this case (see Figure 7.12).

As a final display we show a plot of the closed loop controller/plant transfer function versus the identifier filter D in Figure 7.13.

Further trials were carried out under various other conditions. The remarkable thing about these trials, from our point of view, was that the Working Example worked.




Figure 7.10: System output, Working Example with $D(z)$ second order Butterworth filter with cutoff at 0.3 radians




Figure 7.11: Identified parameters, Working Example with $D(z)$ second order Butterworth filter with cutoff at 0.3 radians




Figure 7.12: Evolution of the identified plant step response for the Working Example with D second order Butterworth low-pass filter




Figure 7.13: Frequency response magnitude plots of $\mathcal{G} = CP$ and D

7.7 Conclusion

The Candidate Robust Adaptive Predictive controller consists of control law design selections which coincide with those recommended by the theory of LQG/LTR robust linear control. Further, these choices are consistent with the objective of minimum variance control design. Coupled with this choice of certainty equivalence control law design is a choice of identifier filter and structure. With this choice one observes how the closed loop identifier operating with this control law yields a model which automatically fulfils its half of the robustness bargain by keeping the appropriate relative model error small. This is the specific model error whose smallness dovetails with the LQG/LTR controller properties to produce a robust closed loop.

The central features of the design are summarized in Section 7.4 and the manipulation of the design variables λ and ρ is discussed in Section 7.5. The major consequence of this study is that we have investigated how one may proceed from control law design, closed loop robustness studies, and closed loop parameter estimation to construct a unified adaptive control procedure which uses the mutually supportive features of the control design and the identifier. By coupling these distinct components of the design, we believe that we have demonstrated how sophisticated robust control procedures can be intelligently incorporated into the adaptive milieu. Further, with the explicit connections between our approach and GPC methods, developed in practice, we have strong circumstantial evidence for the veracity of many of our propositions.

At this stage in proceedings we have effectively explored in detail the intricacies of adaptive LQG optimal control design methodology. But in charting this course we have touched upon elements from broader areas of control. Our next avenue for progress is to investigate methods for placing this (we hope) appealing theory onto a less specific pedestal. In other words, 'Let's generalize'.

Chapter 8

Le Jugement Dernier¹

8.1 Introduction — The Final Analysis

In the final analysis, it is the closed loop performance of any control system design, be it adaptive or otherwise, which determines the success or failure of the method. Consequently, it is a serious aspect of an analysis or synthesis theory to be able to comment on the performance properties of the resultant controller. It is this subject which we address chiefly in this chapter with regard to adaptive controllers.

It is necessary to differentiate between three kinds of criteria of a closed loop control system:

- **Desired** properties are those associated with a specification of features which would be advantageous if achievable. They might include perfect tracking and disturbance rejection, for example. These reflect the desires of the designer and so are available, even if perhaps only being subjectively or qualitatively specified.
- **Designed** properties are those determined by the explicit controller design as performed using a model, \hat{P} , of the plant. Since this is a computed property it is nominally known to the designer through her choice of particular methods and quantitative specifications.
- **Achieved** properties of the closed loop deal with the resultant effects measured when the designed controller is connected into the feedback

¹We reiterate here for our francophone readers that our chapter title is not meant to indicate an arrogant self-assessment of the contents of this book but rather is a reference to the site of some of the formative philosophizing for (and escape from) this very material. The address: 165 Chaussée de Haecht, Schaerbeek, Bruxelles B-1030, Belgium. The telephone number: int + 32 2 217 9597.

loop with the actual plant, P . Assessment of these properties can only be performed by physical experiment but it is these quantities that are the object of the earlier desires.

In adaptive control of the indirect certainty equivalence style that we treat, it is the rôle of the control law schema to embody the desired objectives and to attempt to achieve the designed properties of the system, while accounting for model inaccuracies in the achieved closed loop by dint of robustness conditions encapsulated in the desired properties. The control law half of the adaptive controller is, in our view, focused primarily on the desired and designed properties and, further, one may view the incorporation of a robustness objective as a detuning or relaxation of these goals to admit ahead of time the likely deviation between designed and achieved features. The identification component, however, may be seen as operating more at the interface between designed and achieved properties, since the aim of this half of the adaptive controller is to adjust the parametric model in response to measured (i.e. experimental) actual system responses. The underlying motivation for adaptation should therefore be to avoid some of the performance degradation associated with robust design by allowing the on-line monitoring of achieved performance. It is this performance-based adaptation which we seek ultimately to characterize in this chapter.

A brief tour of the chapter discloses the following itinerary. We begin by reassessing the connections between robust stability and adaptation for control laws more general than those LQG/LTR laws of Section 5.4. This serves to clarify the rôle of closed loop identification design in achieving stability. Then we carry on to study the interaction of closed loop identification and achieved performance, again for general control laws. This exposes both the behavior of the adaptation in affecting the performance and the part played in influencing this behavior through judicious selection of the filters and reference signal. From this juncture our developmental task is almost complete and so we turn, albeit rather briefly, to propose some extensions to and refinements of the earlier work to indicate potential areas for application. We then gracefully bow out with a comment or two on the insights (hopefully) gained en route to the close. This represents the denouement of the treatise.

8.2 Adaptation and Stability Robustness

Recall that the advantage of LQG/LTR control laws in adaptive control rested with the following properties:

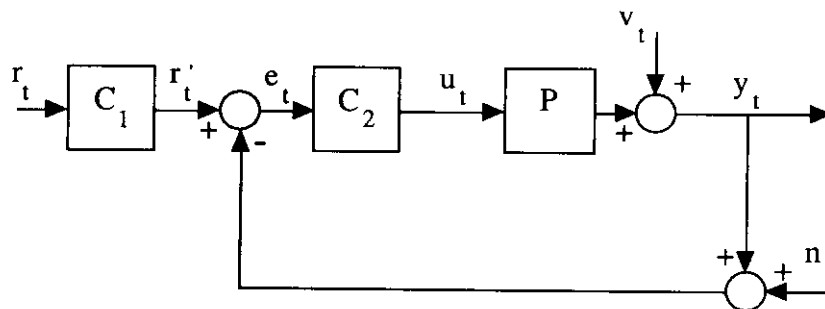


Figure 8.1: Two degree of freedom control system

- The control law, as a linear time-invariant controller, possesses a guaranteed degree of robustness to linear plant model errors.
- The class of plant model errors, to which robustness is achieved, is able to be characterized in a useful fashion. Specifically, LQG/LTR control laws provide robustness to multiplicative plant errors, at least when the plant is minimum phase.
- The control law's effect upon the closed loop plant input spectrum is also characterized simply, so that the interplay between the identifier and the controller is assessable.

In searching for alternative control law schema, each of the above aspects needs to be addressed. Specifically, we shall endeavor to study general linear controller designs from the standpoint of their desired and designed (i.e. determined with $\hat{P}(z, \theta)$) closed loop performance somewhat irrespective of their detailed computational approach. We shall treat separately the issues of robustness of closed loop stability and robustness of closed loop performance. Although, obviously, the latter is contingent upon the former, performance enhancement is the ultimate goal of adaptation and so we investigate mechanisms for achieving this. We shall roughly follow [BA90] in recasting these adaptive control issues in the broader context. The closed loop system under consideration includes the general two degree of freedom controller of Chapter 6, (6.24), depicted in Figure 8.1 with a minor notational alteration ($F_1 = C_2 C_1$, $F_2 = C_2$), and with the explicit addition of sensor noise, n_t , distinct from the output disturbance, v_t .




Figure 8.2: Unity feedback system for stability analysis

8.2.1 Linear Stability Robustness

Our aim here is to present as briefly and rapidly as possible a range of alternatives and extensions to the linear stability robustness theory of Section 5.2 based upon the early work of Lehtomaki *et al.* Recall that we still have the unity feedback structure of that section and that the closed loop stability only depends upon the controller element $C_2(z)$ in Figure 8.1. Therefore, for the moment we consider just the unity feedback loop of Figure 8.2 and concentrate on embellishments of the earlier theory, drawing more or less freely on the formulations of Lunze [Lun89] and Morari and Zafriou [MZ89].

Plant Perturbations

The treatment of Chapter 5 centered upon robustness to multiplicative plant perturbations, where the actual plant, $P(z)$, and the model of the plant, $\hat{P}(z)$, were related via

$$P(z) = \hat{P}(z)(I + L_M(z)) \quad (8.1)$$

$$L_M(z) = \hat{P}^{-1}(z)(P(z) - \hat{P}(z)). \quad (8.2)$$

(That is, $I + L_M(z)$ here equals the $L(z)$ of (5.5) of Chapter 5.) Alternative descriptions of plant-model mismatch are also possible based upon additive perturbations,

$$P(z) = \hat{P}(z) + L_A(z) \quad (8.3)$$

$$L_A(z) = P(z) - \hat{P}(z), \quad (8.4)$$




Figure 8.3: Actual plant as a perturbed model structure

or upon feedback perturbations,

$$P(z) = \hat{P}(z) (I + L_F(z) \hat{P}(z))^{-1} \tag{8.5}$$

$$L_F(z) = P^{-1}(z) - \hat{P}^{-1}(z). \tag{8.6}$$

Diagrammatic representations of the actual plant in terms of the model, $\hat{P}(z)$, and the perturbation, $L_i(z)$, $i = M, A, F$, are given in Figure 8.3. We shall focus upon the multiplicative and additive perturbations mostly but we do remark that still further descriptions are possible (see [Lun89]), of which the feedback description is taken as an example. The aim is to select that perturbation which best reflects plant knowledge and/or modeling uncertainty. We make the following assumption which was not invoked in the derivations of Chapter 5:

Assumption 8.1

The plant perturbation, $L_i(z)$, is asymptotically stable.

Note that if $P(z)$ and $\hat{P}(z)$ have the same unstable parts, then $L_M(z)$, $L_M^{-1}(z)$ and $L_A(z)$ are all stable.




Figure 8.4: Feedback interconnection of the perturbation

Feedback Reformulation

Taking the lead from the redrawings of Figure 8.3, it is possible to replace the plant, $P(z)$, in the unity feedback system Figure 8.2 by its equivalent model and perturbation. The closed loop then, itself, may be redrawn as the feedback interconnection of the perturbation, $L_i(z)$, in isolation with a feedback transfer function, $M_i(z)$, composed solely of a rearrangement of the plant model, $\hat{P}(z)$, and the controller, $C_2(z)$. For example, for an additive perturbation, $L_A(z)$, one has, referring to Figure 8.4,

$$\begin{aligned} u_t &= -C_2(z) \left(\hat{P}(z)u_t + z_t \right) \\ &= - \left(I + C_2(z)\hat{P}(z) \right)^{-1} C_2(z)z_t, \end{aligned}$$

yielding Figure 8.5 with

$$M_A(z) = \left(I + C_2(z)\hat{P}(z) \right)^{-1} C_2(z). \quad (8.7)$$

Similarly, one arrives at equivalent forms for multiplicative and feedback uncertainties,

$$M_M(z) = \left(I + C_2(z)\hat{P}(z) \right)^{-1} C_2(z)\hat{P}(z) \quad (8.8)$$

$$M_F(z) = \left(I + C_2(z)\hat{P}(z) \right)^{-1} \hat{P}(z). \quad (8.9)$$




Figure 8.5: Perturbation feedback system

Closed Loop Stability

We reiterate that for each of the above perturbation descriptions, the equivalent feedback, $M_i(z)$, to the $L_i(z)$ is determined completely by the designed closed loop, via $\hat{P}(z)$ and $C_2(z)$. We make the following assumption:

Assumption 8.2

The closed loop system designed on the basis of the model $\hat{P}(z)$, i.e. with zero plant perturbation, is internally stable so that no uncontrollable nor unobservable unstable modes are present in the designed closed loop. Thus the following transfer functions are asymptotically stable:

$$\begin{aligned} & (I + C_2(z)\hat{P}(z))^{-1}, & (I + C_2(z)\hat{P}(z))^{-1} C_2(z), \\ & (I + C_2(z)\hat{P}(z))^{-1} \hat{P}(z), & (I + C_2(z)\hat{P}(z))^{-1} C_2(z)\hat{P}(z). \end{aligned}$$

This is a standard reasonable assumption describing the well posedness of the ideal design [Fra87].

We have then, directly from the Nyquist stability theorem:

Theorem 8.1

Suppose that Assumptions 8.1 and 8.2 hold, then the closed loop unity feedback system of Figure 8.2 is asymptotically stable, provided

$$\bar{\sigma} (L_i(e^{j\omega})M_i(e^{j\omega})) < 1 \quad \text{for all } \omega \in (-\pi, \pi]. \quad (8.10)$$

The presentation in Theorem 8.1 invites comparison with the earlier Theorem 5.3. For multiplicative perturbations both theorems purport to

give sufficient conditions for robust stability. The correspondence between the two conditions may be seen by noting the following identities, with L defined as in Chapter 5 and $\mathcal{G} = C_2\hat{P}$,

$$\begin{aligned} I + \mathcal{G}L &= \left[(L^{-1} - I)(I + \mathcal{G})^{-1} + I \right] (I + \mathcal{G})L \\ L_M \times M_M &= (L - I) \times (I + \mathcal{G})^{-1}\mathcal{G}. \end{aligned}$$

Whence

$$L_M M_M + I = \left[(L^{-1} - I)(I + \mathcal{G})^{-1} + I \right] L.$$

The stability Assumption 8.1, that P and \hat{P} have the same unstable parts, implies that the multiplicative perturbation, L_M , and its inverse, L_M^{-1} , both be stable because $P = \hat{P}(I + L_M)$. In light of this it follows that the two robustness conditions are equivalent. Indeed, under the previous assumption, the additional 1 in the inequality of Theorem 5.3 may be removed.

We next turn to the assessment of the interplay between these general linear stability robustness conditions and closed loop RLS identification.

8.2.2 Closed Loop Identification and Stability Robustness

Recall our earlier analysis of Chapter 7, where the closed loop identification studies of Chapter 6 were applied specifically to LQG/LTR by making the connections:

- the controllers possess inherent robustness to multiplicative modeling errors made evident via the EDR,
- the closed loop input spectrum, $\Phi_u(\omega)$, is effectively given by the reference spectrum filtered through the inverse of the plant,
- the structure of the RLS identification criterion, $V(\theta)$, when interpreted in the frequency domain and with this class of input, was such as to minimize the multiplicative modeling error in line with the requirements of the robust stability theory.

Now we reinterpret the frequency domain formulation of $V(\theta)$ simply in terms of specific $\{L_i(z), M_i(z)\}$ pairs for the various plant perturbation classes.

We refer to the earlier relation (6.32), which we repeat here for clarity with the notational alteration ($F_1 = C_2 C_1$, $F_2 = C_2$):

$$V(\theta) = \frac{1}{4\pi} \int_{-\pi}^{\pi} \left\{ \left| \Delta P(e^{j\omega}, \theta) \frac{C_2(e^{j\omega})C_1(e^{j\omega})}{1 + C_2(e^{j\omega})P(e^{j\omega})} \right|^2 \Phi_r(\omega) \right\} \quad (8.11)$$

$$+ (|\Delta P(e^{j\omega}, \theta) \frac{C_2(e^{j\omega})}{1 + C_2(e^{j\omega})P(e^{j\omega})}|^2 + 1) \Phi_v(\omega) \left. \vphantom{\frac{C_2(e^{j\omega})}{1 + C_2(e^{j\omega})P(e^{j\omega})}} \right\} \frac{|D(e^{j\omega})|^2}{|\hat{H}(e^{j\omega}, \theta)|^2} d\omega.$$

Here $\Delta P(z, \theta) = P(z) - \hat{P}(z, \theta)$ and, while still using the scalar form for legibility, we have grouped the respective transfer functions slightly differently.

We now draw the reader's attention to the integrand of (8.11), which we rewrite in terms of the $L_i M_i$ as follows. We observe that

$$[P - \hat{P}] \frac{C_1 C_2}{1 + C_2 P} = \begin{cases} L_A M_A \frac{1 + C_2 \hat{P}}{1 + C_2 P} C_1 \\ L_M M_M \frac{1 + C_2 \hat{P}}{1 + C_2 P} C_1 \\ L_F M_F \frac{1 + C_2 \hat{P}}{1 + C_2 P} P C_1 C_2 \end{cases} . \quad (8.12)$$

Next define transfer functions,

$$W_A = W_M = \frac{1 + C_2 \hat{P}}{1 + C_2 P} \quad (8.13)$$

$$W_F = \frac{1 + C_2 \hat{P}}{1 + C_2 P} P C_2, \quad (8.14)$$

to yield the reconstrued closed loop identification criterion in terms of the linear robustness elements,

$$V(\theta) = \frac{1}{4\pi} \int_{-\pi}^{\pi} \left\{ |L_i(e^{j\omega}, \theta) M_i(e^{j\omega})|^2 |W_i(e^{j\omega})|^2 \left(|C_1(e^{j\omega})|^2 \Phi_r(\omega) + \Phi_v(\omega) \right) + \Phi_v(\omega) \right\} \frac{|D(e^{j\omega})|^2}{|\hat{H}(e^{j\omega}, \theta)|^2} d\omega. \quad (8.15)$$

The import of (8.15) is that it reinterprets the closed loop identification criterion directly in terms of quantities of interest for closed loop stability robustness. Thus, the selection of reference excitation signal, r_t , and identification filter, $D(z)$, are immediately addressable in terms of the strictions of Theorem 8.1 that $|L_i(e^{j\omega}) M_i(e^{j\omega})| < 1$. The novelties here are:

- The identification criterion in terms of additive uncertainty is identical to that for multiplicative uncertainty. Therefore the extension of the results of Section 5.2 is relatively easily made. (The robustness of LQG to additive perturbations, though, is not easily established.) With a feedback description of uncertainty, however, matters are less clear and this really explains why the case was introduced only for comparison.

- The weighting terms, $W_A(z)$ and $W_M(z)$, of (8.13) are composed of the ratio of modeled and actual return differences, which is nominally unity over the bandwidth of interest.
- The signal weighting for the plant identification, $|C_1|^2\Phi_r + \Phi_v$, is nominally a known spectrum. That is, $|C_1|^2\Phi_r$ is the spectrum of the known reference signal passing through the feedforward controller and Φ_v is an approximately modeled additive factor which frustrates the minimization of $|L_iM_i|$ when $\Phi_r(\omega) < \Phi_v(\omega)$. This occurs because of Φ_v 's appearance in two places in (8.15) — weighting $|L_iM_i|$ and as an additive positive term. Typically we have that $\Phi_r \gg \Phi_v$ over $\omega \leq \mathcal{W}_r$.
- The identification filter, $D(z)$, still remains available for the adjustment of the bandwidth of the model fit through alleviating improper distortions introduced by the reference, the controller and the noise model. (See Figure 8.6.)
 - In order to encourage the minimization of $|L_iM_i|_\infty$ in the adaptation, $D(z)$ should be chosen roughly to flatten the weighting on ΔP in (8.15).
 - Recall from the discussion of Chapter 5 that the requirement that $|L_iM_i| < 1$ is frequently only an active constraint over a frequency band $\omega < \mathcal{W}_s$, beyond which the smallness of this product is assured by model and plant rolloff. $D(z)$ should only weight the model fit over this band and should, itself, roll off after this to prevent noise effects due to Φ_v beyond this frequency.
 - The reference signal spectrum, Φ_r , is typically low-pass and the designed closed loop bandwidth would normally significantly exceed this value. Thus $D(z)$ should compensate for a lack of reference excitation between the reference bandwidth, \mathcal{W}_r , and \mathcal{W}_s , an interval which will include the closed loop gain crossover.
 - While the underlying desired reference signal is low-pass, it is necessary that r_t possess some significant frequency components throughout the band \mathcal{W}_s for stability robustness to be achieved. This is a persistence of excitation condition which is necessary to ensure that adaptation remains robust.

Summary

We pause now to reflect briefly upon the nature of those results just established in terms of their part in extending the adaptive LQG/LTR theory of earlier chapters.

We began by considering some different formulations of the robust stability problem for linear systems, firstly by writing down sample descriptions of plant uncertainty, *viz* additive, multiplicative and feedback. When the uncertainty is described by a stable transfer function, one may then appeal to the Nyquist stability theorem to provide a robust stability condition expressed solely in terms of the uncertainty, $L_i(z)$, and a designed transfer function, $M_i(z)$, evaluated on the unit circle. By rewriting the frequency domain RLS identification criterion in terms of a weighted integral of $L_i M_i$, it then becomes possible to display directly the influence of the closed loop identification and its design variables upon the robust stability of the linearized closed adaptive loop.

These conclusions are an extension of the earlier LQG results because the specific origin of the controller is irrelevant. As we shall see, they are very general and apply to H_∞ or other control strategies as much as to LQG. The point worth making here is that the LQG/LTR study motivated by GPC comparisons fits within this framework but lends itself mostly to the robustness approach using multiplicative descriptions of plant uncertainty. Mechanisms for extension to, say, H_∞ designs for additive uncertainty are immediately apparent. Tied to the robust stability problem is the problem of ensuring robust performance. This is the next issue to be broached.

8.3 Adaptation and Performance

The first requirement of system performance is that the closed loop be stable. But this is clearly not the full story, and the greater part of control system design is associated with efforts to achieve this with intransigent plants. Indeed, the *raison d'être* for Adaptive Control is its potential enormous performance enhancement features. Our task now is to study the robustness of closed loop performance for adaptively controlled systems very much in line with the immediately above treatment of stability robustness before combining the guidelines of both.

We refer again to Figure 8.1 and the signals described therein. There are many performance measures for the closed loop system, so that in practice a control design often ends up as an implicit multicriterion optimization problem. For example, if one's goal is to have the plant output, y_t , track the reference signal, r_t , and/or to reject the effects of the output disturbance, v_t , then the natural objective function to consider is related to the plant sensitivity function,

$$S(z) = [I + P(z)C_2(z)]^{-1}, \quad (8.16)$$

since this is the transfer function from the filtered reference, r'_t , to the error

signal, e_t , or the transfer function from the disturbance, v_t , to output, y_t . Alternatively, if the feedback of sensor noise, n_t , is a problem, then the complementary sensitivity function,

$$T(z) = I - S(z) = P(z)C_2(z)[I + P(z)C_2(z)]^{-1}, \quad (8.17)$$

is more appropriate, being the transfer function from sensor noise, n_t , to the output, y_t , and from r'_t to y_t . More usually, a combination of both behaviors is desired along with, say, limits upon control signal gains and other performance measures.

Our LQG designs of Chapter 3 are examples of design methods aimed at common objectives of reference tracking and disturbance rejection, using the knowledge of the signal spectra to effect the design. The figure of merit in such a case would typically be associated then with properties of the closed loop sensitivity. As we are posing questions of robust performance with adaptation, one is interested in issues of the ability to keep the sensitivity function close to a nominal design value (usually zero) over a frequency range of interest, typically the bandwidth of the reference signal.

We define the following two sensitivity functions. (We limit ourselves to the scalar transfer function case for brevity again.)

$$S(z) = (1 + C_2(z)P(z))^{-1}, \quad \hat{S}(z) = (1 + C_2(z)\hat{P}(z))^{-1}. \quad (8.18)$$

These are the *achieved* sensitivity and the *designed* sensitivity function. Our next observation is that the integrand of the RLS criterion (8.11) may be reorganized in terms of these sensitivities:

$$\begin{aligned} [P - \hat{P}] \frac{C_2}{1 + C_2P} &= \frac{PC_2}{1 + C_2P} - \frac{\hat{P}C_2}{1 + C_2\hat{P}} + \frac{\hat{P}C_2}{1 + C_2\hat{P}} - \frac{\hat{P}C_2}{1 + C_2P} \\ &= T - \hat{T} + \hat{P}C_2(\hat{S} - S) \\ &= (1 + \hat{P}C_2)(\hat{S} - S). \end{aligned} \quad (8.19)$$

From (8.19) we may substitute into the RLS criterion (8.11) yielding an alternate to (8.15),

$$V(\theta) = \frac{1}{4\pi} \int_{-\pi}^{\pi} \left\{ |(1 + \hat{P}C_2)(\hat{S} - S)|^2 (|C_1|^2 \Phi_r + \Phi_v) + \Phi_v \right\} \frac{|D|^2}{|\hat{H}|^2} d\omega. \quad (8.20)$$

As with the formulation of (8.15), the consequences of (8.20) are direct:

- The (desired) tracking and disturbance rejection performance objectives are achieved by keeping $|S(e^{j\omega})| \ll 1$ for $\omega \leq \mathcal{W}_r$, where \mathcal{W}_r is the bandwidth of the reference signal, r_t .

- The design objective is to keep $|\hat{S}(e^{j\omega})| \ll 1$ for $\omega \leq \mathcal{W}_r$. That is,

$$|\hat{P}(e^{j\omega})C_2(e^{j\omega})| \gg 1, \quad \text{for } \omega \leq \mathcal{W}_r, \quad (8.21)$$

and to keep $|1 + \hat{P}C_2|$ smaller at high frequencies. This objective is to be achieved by proper designed controller selection, C_2 .

- The high frequency aspects of the selection of the identification filter, $D(z)$, are basically determined by the requirements of the stability robustness conditions. At low frequencies, however, it should be chosen also to flatten distortions introduced by excessive weight in $\hat{P}C_2$.
- The stability robustness conditions would lean towards the selection of a reference signal, r_t or r'_t , which is rather broadband to account for the need to have $|L_i M_i|$ small over the entire modeling bandwidth, \mathcal{W}_s . This conflicts with the performance requirement of tracking r_t over a lesser bandwidth, \mathcal{W}_r . Indeed, the stability requirement, which must take precedence, dictates that the reference be artificially extended in bandwidth, thereby penalizing the achieved performance. It is our view that this is the price to pay for adaptation.

The totality of recommendations for the design of $D(z)$ and Φ_r from both a stability robustness and performance robustness point of view are encapsulated in Figure 8.6.

8.4 Forethoughts on a Postscript

With Chapter 7 we (finally) broached the subject of synthesis of adaptive control schemes on the basis of our joint local analysis of the identifier and the control law schema. Our message that this design has the potential for the mutual support and synergism of these two components has been established by the construction of just such an adaptive controller example using an LQG/LTR controller coupled with an RLS identifier. This example demonstrated the amusing feature of the subsumption of several features of the adaptive GPC methodology within its scope, thereby lending weight to the theoretical pursuit of more coherent adaptive control design methodologies and giving technical support and direction to the algorithmic development of practical adaptive controllers. It is our feeling that the strong family ties between the CRAP controller and the GPC vindicate our whole development. These LQG/LTR results have been carried over to more general control laws in the preceding sections and the notion of achieved performance-based adaptation has been introduced. We have now




Figure 8.6: Example frequency regimes for definition of $D(z)$

reached the climax of our work — the main course is finished — and we should delicately wind up proceedings with a tasty but light dessert. This we do by raising some perennial questions of research,

‘*So What?*’² and ‘*Where do we go from here?*’.

There are several directions for the pursuit of answers to these questions which will be explored, with the intention of evaluating briefly what impact the insights gained earlier might have in the broader discipline of Adaptive Control. These fall into two categories —

Extensions and Generalizations

- What other (more) modern robust control design techniques might be directly amenable to reassessment in the light of our new viewpoint? In particular, can recent advances in, say, H_∞ control design be of direct applicability in adaptive control?
- Conversely to the preceding question, can modifications to the identification procedure be of possible benefit with LQG or some other feedback law? Especially of interest here is an exploration of the use of multiple-step-ahead prediction error criteria in the identification stage, which would really make these Predictive Adaptive Controllers.
- Do there exist numerically advantageous paths to achieve these adaptive controllers? These methods could exhibit either reduced computational complexity or reduced computational sensitivity.

Refinements and Theoretical Supports

- Can one delineate more usefully the class of plants for which Adaptive Control is suitable from a robustness and performance perspective?
- What modifications are necessary to the theory to be able to accommodate truly time-varying plants, since this is the *raison d’être* of Adaptive Control? Further, how does one resolve the L_2 frequency domain reformulation of the RLS criterion with the L_∞ robust stability requirements?

²The answer is not 42.

- How do these methods avoid the pitfalls, if indeed they do, of other adaptive controllers as exemplified in, say, the work of Egardt [Ega79] and Rohrs *et al.* [RVAS85]?
- Do the methods proposed give any insight into the outcome of the World Heavyweight Title Fight, Robust Control versus Adaptive Control?

The extension questions will be addressed first to indicate some directions available for natural development from the earlier theory. Specifically here we treat the first two questions but, unfortunately, leave in abeyance the issue of numerically sophisticated algorithm description. (Because this is not our expertise, we leave the generation of fast and numerically stable adaptive control algorithms as an exercise for the reader.) Following this we shall turn inwards again to approach the questions of refinement and interpretation of the theoretical statements.

8.5 Extensions and Generalizations

8.5.1 Candidate Alternative Optimal Control Laws

We discuss two prospective Adaptive Optimal Control Techniques. We deliberately steer away from strictly algebraic control techniques such as pole placement, certain model reference strategies and geometric disturbance decoupling controllers because of their propensity for non-robustness and discontinuity. This is exemplified by issues of near non-coprimeness of factors in the solution of the Diophantine equations of pole placement via Sylvester matrix techniques. As a broad value judgment, functional analytic control design methods tend to admit easier connection to (nonlinear) robustness than purely algebraic methods. This, in part, explains our restriction of subject matter to Adaptive Optimal Control.

Frequency Weighted Adaptive LQ Optimal Control

The LQ control objective is stated in terms of the plant and control signal weightings, Q_c and R_c ,

$$J(u) = \sum_{t=1}^{\infty} x_t^T Q_c x_t + u_t^T R_c u_t. \quad (8.22)$$

Desired properties of the closed loop performance, such as disturbance rejection, reference tracking and decoupling must be massaged so as to fit into the

LQ criterion formalism. A recent advance in cajoling differing goals into the LQ framework has been to incorporate explicit frequency shaping properties into the LQ cost function. Thus (8.22) is replaced by

$$J(u, \omega) = \sum_{t=1}^{\infty} x_t^T Q_c(e^{j\omega}) x_t + u_t^T R_c(e^{j\omega}) u_t, \quad (8.23)$$

where now the weighting matrices Q_c and R_c have been replaced by frequency dependent weighting functions $Q_c(e^{j\omega})$ and $R_c(e^{j\omega})$ which reflect the designer's desires for frequency bands in which control action should be concentrated.

Being non-negative definite frequency dependent functions, $Q_c(e^{j\omega})$ and $R_c(e^{j\omega})$ are actually spectra. Thus, we take

$$Q_c(z) = W_1^T(z^{-1}) Q_c^0 W_1(z) \quad \text{and} \quad R_c(z) = W_2^T(z^{-1}) R_c^0 W_2(z),$$

for given finite dimensional frequency weighting transfer functions, i.e. filters, $W_1(z)$ and $W_2(z)$. The solution to the frequency weighted LQ problem then follows from the EDR. Write the left side of the EDR as

$$\begin{aligned} R_c(z) + G^T(z^{-1} - F^T)^{-1} Q_c(z) (zI - F)^{-1} G \\ = W_2^T(z^{-1}) \{ R_c^0 + W_2^{-T}(z^{-1}) G^T (z^{-1}I - F^T)^{-1} W_1^T(z^{-1}) Q_c^0 \\ \times W_1(z) (zI - F)^{-1} G W_2^{-1}(z) \} W_2(z), \end{aligned} \quad (8.24)$$

which admits the interpretation of the frequency weighted LQ problem as a modified higher order LQ problem, which may be solved via spectral factorization/ARE techniques to find the frequency dependent gain $K(z)$ that solves the problem. Similar treatments of frequency weighted Kalman filters also are directly possible. The reader is referred to [MM87], [AM90] for details.

What is of interest here in adaptive control is that these frequency weighted LQ controllers are derived via EDRs and, being but glorified higher order LQ problems plus filters, possess guaranteed robustness margins for (frequency weighted) relative model errors as well as an easily ascertained effect on the closed loop signal spectra. Thus, provided $W_1(z)$ and $W_2(z)$ are full rank on the unit circle and unity at infinity, this design process has the capacity to fall easily into our theory and engender a whole bevy of frequency weighted CRAP controllers. The *quid pro quo* for this enormous added flexibility is, of course, increased computational complexity. As remarked earlier, we ride roughshod over all issues of a numerical nature because these are frequently resolved by inexorably marching technology.




Figure 8.7: General 2-input–2-output H_∞ problem formulation

H_∞ Optimal Adaptive Control

Probably the most popular, or at least most exciting, recent robust linear control law design method is H_∞ Optimal Control [Fra87]. Here the design criterion is deliberately specified with a closed loop stability robustness issue in mind. We present a précis of the features of H_∞ control which best display its dovetailing with our theory so far. We apologize to ‘person in the street’ who has lasted this far for our elevated tone in preaching really to the cognoscenti here, but our desire is (like Polonius) to be brief, or at least witty.

Although a large number of linear system problem types may be recast or coerced into the H_∞ framework, the standard ones concentrate on achieving guaranteed robust stability in spite of plant perturbations (or *unstructured uncertainties* in the vernacular) which are classified only in terms of their H_∞ norm. Thus no further structural information is given about these perturbations and the H_∞ problems are regarded as minimax problems. General formulations of H_∞ problems are possible and revolve around a 2(vector) input–2(vector) output system, the connection between the second output vector and the second input vector of which encapsulates the feedback control. The first vector input represents disturbances and references, and the first vector output consists of errors and measured output signals. In this general situation, the H_∞ problem has a solution which is strikingly structurally similar to LQG — at least in continuous time. The controller is composed of state-variable feedback and state-variable estimation analogues derived from the solution of AREs but with sign-indefinite R_c and R_o . The

parallels are then drawn between H_∞ Optimal Control and Linear Quadratic Game Theory [DGKF89].

In addition to this remarkable similarity in form between the H_∞ solution and the LQG solution shown by Doyle *et al.*, Fujii and Khargonekar [FK88] have explored explicit (but not really dangerous) liaisons between H_∞ controllers and LQ controllers, including the occurrence of Return Difference Equality analogues.

The full development of H_∞ control in discrete time is, as yet, not complete but its analysis so far is tantalizingly suggestive of great possibilities [LGW89]. We make the following observations about H_∞ control from an adaptive viewpoint:

- H_∞ possesses explicit characterization of the class of plant knowledge required to ensure closed loop stability. Internal stability is guaranteed by construction.
- Through the generalized (four block) problem specification, the H_∞ formalism is amenable to application for a large variety of common control problems including, *inter alia*, disturbance rejection, reference tracking and control penalty. In this way, it has much in common with LQG.
- Since H_∞ problems are minimax or worst-case analyses, the emphasis is placed on maximum transfer function magnitudes, or ℓ_2 system gains, rather than ℓ_2 signal energies which arise in LQG control and RLS identification. This is the essence of the distinction between H_∞ and LQG.
- By incorporating weighting transfer functions into the optimization criterion, frequency weighted H_∞ problems are readily posed and solved. This overcomes the strict minimax property of H_∞ at the expense of system order.
- Algorithmic procedures for the solution of H_∞ problems bear a great structural similarity to the LQG solution methods. Both involve the solution of an ARE and a dual ARE which differ between the methods according to the sign-definiteness of R_c and R_o .
- The tools of discrete-time H_∞ theory are at present underdeveloped. The continuous-time return difference equality objects exist but it is not yet clear how they should be applied to characterize closed loop plant input spectra. It is also unclear to what Loop Transfer Recovery

in H_∞ controllers might correspond, or for what they might prove useful.

- Numerically, by grace of the ARE connection above, H_∞ problems should be only about as difficult to compute as LQG. Existence questions for positive definite stabilizing solutions of the ARE arise. At present these methods for controller design can yield a controller achieving a given H_∞ bound, if such a controller exists, but cannot necessarily provide the infimal control law.

It is apparent from our commentary above that Adaptive H_∞ Optimal Control is an enthralling prospect since its robustness is manifestly more easily tunable than LQG, although perhaps at the price of increased difficulty in guaranteeing performance measures. At the time of writing the analytic machinery for discrete H_∞ was only in its infancy but, of the criteria necessary for appeal to our interplay theory, H_∞ clearly possesses strict robustness margins linking the plant model accuracy to closed loop stability. The missing link for the complete inclusion of H_∞ is the need to be able to relate fully the closed loop plant input spectrum to the control design schema.

8.5.2 Alternative Identification Methods

While one would be tempted to draw a full analogy between the quadratic performance indices of LQG and RLS to support that a universal balance should exist between the control and identification components in Adaptive Control, there remains a fundamental distinction in that the control law stability robustness depends upon system ℓ_2 gains, i.e. L_∞ bounds on the frequency response, while the identification is concerned with data dependent error signal energies, leading to L_2 frequency response modeling as in Chapter 6. Thus the controller robustness condition is expressed in terms of a maximum frequency response magnitude value while the identification is expressed as an integral squared frequency response fit. Hence, in looking for potential identification method choices there need not necessarily be any counterparts to available control modifications. Nevertheless, we do consider two alternative identification criteria with which to replace the one-step-ahead prediction error RLS procedure. Many of the issues are raised and dealt with rather well in [Lju87], to which we refer the reader for supporting material.

Equation Error versus Output Error

The distinction between Equation Error and Output Error linear system modeling lies in the construction of the regressor vector. In Equation Error one uses a regression vector composed of past plant inputs, u_t , past plant outputs, y_t , and past noise estimates, $\hat{e}_{t|t-1}$, to compute the parameter estimate, θ . In Output Error the regressor is comprised only of past inputs, u_t , and past estimated plant outputs, \hat{y}_t , arising from the parametrized model. Thus Output Error corresponds to the Prediction Error methods earlier with a fixed noise model, $\hat{H}(z, \theta) = 1$. Equation Error has a more general model.

The respective advantages of these schemes are as follows. Equation Error permits the identification of both a plant input–output model and a noise model for the corrupting measurement noise. On the other hand, Output Error methods sacrifice the noise model for the sake of an estimated plant model which suffers considerably less from bias problems. Thus, one observes that the choice between the two methods might well be seen as a trade-off between performance and robustness — a good noise model is essential to achieve adequate disturbance rejection, while a close plant model irrespective of the noise is needed for stability robustness. Alternatively, the availability (or not) of prior information about the disturbance properties might be a decisive criterion for the choice between these two methods.

In terms of the frequency domain interpretation of these two schemes, the off-line behavior of Equation Error may be transformed to be identical to that of Output Error by correct selection of the filtering $D(z, \eta)$. For the recursive, on-line schemes however, there are some significant differences. Output Error is nonlinear in the parameters because it involves an explicit filtering by a parameter dependent transfer function and it has more difficult convergence properties. Equation Error is simpler to run but can need to rely upon the selection of its filters to achieve adequate models from which to perform control design. It is our view that the convergence difficulties of Output Error methods would militate against their use in all but the most reliably modeled (or modelable) processes. The absence of serious nonlinearities in Equation Error estimation, coupled with the need to include only linear filtering, indicates that it is probably best suited to circumstances where frequency based thinking figures strongly in the design.

Multiple-step-ahead Predictors

From its formulation, one would expect that Predictive Control techniques involving long range horizons would rely upon the quality of their plant output predictions to achieve adequate control performance. The longer the

horizon, the more critical this should become. Hence, people have proposed several times that system models be fitted not on the standard one-step-ahead LS criterion but on a LS criterion using many-steps-ahead prediction errors. Additionally Mosca *et al.* [MZL89] and others have suggested that predictors of all necessary scopes be fitted on-line to generate the \mathbf{f} vector of predictions in GPC.

From an off-line identification point of view, one can see that the effect of altering the modeling criterion from one step ahead to more is manifested in the frequency domain formulae of Chapter 6, through the alteration of the noise model [Lju87]. Thus, from an off-line perspective, the same effect could be achieved by filtering since $\hat{H}(z, \theta)$ and $D(z, \eta)$ play similar rôles. As a general rule, the larger the prediction horizon becomes the more emphasis is placed in the model fit at low frequencies. Further, in the limit that the horizon tends to infinity, one recovers the noise model free Output Error version of the criterion. This is discussed in detail by Ljung [Lju87].

These modifications to the identification law have some interpretation of their effects given by this style of off-line analysis but naturally, in a practical adaptive closed loop, the peculiarities of each approach would come to the fore. A meaningful comparison between transient behaviors and the dynamics of the recursive algorithms is beyond our capabilities here. However, the techniques of Chapter 6 show that sensible guidelines for performance in the neighborhood of stationary solutions can be gained.

8.6 Refinements and Theoretical Support

Apart from hypothesizing about where these Adaptive Optimal Control ideas might be headed, we also recognize the need to clean up where we have been to *préciser* the theory in order to improve its applicability. Central amongst the objections leveled at local adaptive control theories are its conservatism and focus upon a single fixed ‘nominal’ parameter value about which the linearization is performed. We shall consider the question of how these methods apply to the underlying problem of adaptation, the tracking of slowly time-varying plants. In addition to these introspections we also deliver a few ideas on some classical battles: the Rohrs examples, and Robust Control versus Adaptive Control.

8.6.1 Theoretical Refinements

Interpretation of Robustness Conditions

Critical to the analysis of our theory of robust interplay between identifier and controller are the conditions of Assumption 6.1 where sufficient requirements for local estimator convergence to a neighborhood of a nominal value are stated. These conditions consist of two components, sensitivity conditions for the linearization to be valid and persistence of excitation conditions connected with the uniform contractivity of the estimation algorithm.

The sensitivity conditions deal with the magnitudes of various partial derivatives in the region around the nominal point. Thus the closed loop control signal and the closed loop output performance need to be smooth in this region. This is a realistic assumption because with adaptation one expects the controller values to be changing smoothly. There is a natural connection between this requirement and the robustness properties of the time-invariant control law. The feature of the analysis which is distinct from a standard linearization treatment is that an explicit upper bound appears on the unmodeled dynamics and therefore on the achieved closed loop performance. Hence, this property delineates the class of feasible systems for the application of this type of adaptive control. That is, the existence of an adequately fitting model for the plant in the model set needs to hold — simple linearizability is not sufficient.

The other class of conditions arising in the theorem are devoted to persistence of excitation properties for the closed loop system. These conditions are undesirable from a practical viewpoint, since they imply the deliberate perturbation of the plant reference signal which will necessarily upset potentially optimal performance. It is our view that such conditions are fundamentally unavoidable during periods of adaptation in order to guarantee the integrity of the data in terms of its content of information about the underlying plant system. The existence of an approximate underlying positive real system, as is the basis for the development in [ABJ+86], is hidden in the description of the relatively bounded modeling error ζ_t . Thus the combined interpretation here of these conditions is very much in accordance with that advanced in [ABJ+86], where adequate modeling of the controlled nominal system is presumed along with persistence of the reference at the frequencies where the model is fitted best.

Time-varying Systems

As we have already alluded, the motive force behind adaptive control systems is frequently the desire to control time-varying plants, where wear on parts,

environmental effects, raw material changes and other operating conditions can cause a significant change in plant dynamics but where good regulation is still desired. The work of [ABJ+86] demonstrates that, with slowly adapting controllers as we have postulated here, the capacity to cope with time variations depends upon these variations being on a time scale which is still slower than that of adaptation. In such a circumstance, the earlier analysis goes through without appreciable modification — the time-varying problem is treated as a sequence of time-invariant ones with local (in time) behavior being governed by the linearized methods presented here.

L_∞ and L_2 Connections

It should be clear to even the most casual reader that our statements about closed loop stability robustness rely upon L_∞ frequency response bounds being satisfied, while the rewriting of the parameter estimation criterion in terms of frequency response indicates only the minimization of an L_2 error. There therefore appears to be a leap of faith required to interpret closed loop adaptation as providing a palliative to stability robustness and/or closed loop performance. Thankfully, however, in a recent paper [CBG89] Caines and Baykal-Gürsoy have studied conditions under which L_∞ norm convergence is implied by Least Squares (L_2) consistency. The main point is that, subject to smoothness or analyticity conditions, convergence takes place in either norm. While our circumstances are both different and less formal than [CBG89], the same general principle applies.

Our presentation so far has homed in upon the crucial effects of closed loop insensitivities in the neighborhood of θ^* . This may be seen in the various rôles played by the robustness requirements in the identification, performance and convergence sections. Robustness or insensitivity conditions may be directly linked to smoothness properties of the frequency responses involved (see for example [Hil59]). This provides yet another link in the chain that binds the control law selection to the identifier.

8.6.2 The Rohrs Examples

The examples of Rohrs *et al.* [RVAS85] were computer-based demonstrations of the fragility of some of the global stability theorems associated with adaptive control in the early 1980s. Although many of the instability mechanisms had been described earlier in the book by Egardt [Ega79], the concrete numerical examples which were easily repeatable helped to turn the adaptive control community collectively towards the issue of robustness — without this latter term ever being defined satisfactorily. Because of this historical

importance, these examples (called ‘counterexamples’ by some) have become a *de facto* benchmark for the consideration of adaptive controller robustness. While these examples are continuous time and neither LQG nor GPC, the paths to instability that they exemplify are all too real in our situation and it behoves us to comment upon our approach to dealing with them.

The Rohrs examples are centrally of two types. Both deal with a third order plant, with one low frequency dominant pole and two well damped high frequency modes, being modeled by a first order system and operating under adaptive model reference control.

In the first example, a single sinusoid is used as the input to the reference model. The frequency of this sinusoid, however, is exactly at that frequency where the phase of the third order system passes through $\pi/2$ radians and the consequent model fit can only be achieved by the selection of infinite parameter values. This is an extreme example of the effects demonstrated by the Gedankenexample of Chapter 1 where alterations in the reference frequency were shown to yield significantly different identified parameter values, and hence significantly different controller gains.

In terms of our theory here, we note that this reference is persistently exciting for the identification of a first order (two parameter) system but that, for the true plant under consideration, the conditions of Assumption 6.1 fail to hold. Specifically, the constraint upon the unmodeled dynamics ζ_t with respect to the modeled portion at a stabilizing value θ^{**} fails to hold. These notions of frequency dependent signal and modeling conditions have been further explored in [ABJ+86] and in the heuristic analysis of Åström [Åst83]. Specific measures that may be taken to reduce these effects, as suggested by these theories, are either to select a control law capable of stabilizing the third order plant when given the first order model fitting at this reference frequency or to inject significant levels of low frequency signals into the reference. A few moments’ thought will indicate that this particular adaptive control problem has been deliberately ill-posed to highlight just these failings.

The second Rohrs example deals with a similar system driven by a constant reference trajectory, which is not persistently exciting for a two parameter model, but with the output measurement corrupted by a sinusoidal disturbance of infinitesimal amplitude. The dynamics here are governed by the feature that a linear manifold in the parameter space (\mathbb{R}^2 here) describes the degenerate class of all parametrizations which permit exact model fit with just the constant reference — such a property was demonstrated with the Gedankenexample as well. The effect of the additive output disturbance, with the model reference criterion, is to force the identified parameter to

‘slide’ along this linear manifold in a vain attempt to zero the plant output error. In this way, the parameters escape to infinity linearly with time.

Our response here is that the message of adaptation only under persistent excitation has been ignored and this is a typical consequence. Indeed, additive output disturbances may be replaced by numerical roundoff to yield the same effect [And85] without violating the conditions of global convergence. Other measures that can deal with this circumstance are adequate leakage [IK83], [HC87], and deadzones.

At this point it is as well to reiterate the comment that not all systems are amenable to adaptive control and, of those that are, their amenability can very well be dependent upon experimental (i.e. signal) conditions. These signal based conditions for the viability of adaptive control in the local sphere have been explored in [ABJ+86].

8.6.3 Adaptive Control versus Robust Control

The Rohrs examples of the preceding subsection were, at the time, a source of considerable debate about the future of adaptive control and its possible contribution *vis-à-vis* the re-emergent field of robust linear controller design, such as that of LQG/LTR as in Chapter 5. Therefore, this is probably a sensible place to reopen the wound and consider a response in the light of the interplay theory presented so far.

Certainly, the robustness of linearizable (i.e. slow) adaptive control does not extend beyond that achievable via explicit robust linear design, since the latter is concocted specifically with the point of maximizing robustness. The penalty for robustness, however, is often felt in reduced performance. Thus an adaptive controller might well outperform a fixed robust controller through its ability to identify more suitable parameters for a less robust but higher performance control law. This was at the heart of the performance based interpretation of the adaptation. Similarly, the same effect can be envisaged with slowly time-varying plants. The cost of adaptation, however, is seen in the lengths to which one must go in order to ensure robust adaptation (in the nonlinear sense). Specifically, persistence of excitation needs to be asserted in frequency bands of central interest, and this could have an obvious detrimental consequence with, say, constant reference tracking because the reference signal is artificially disturbed to maintain the plant information in the measured signals.

The choice between Robust Linear Control and Adaptive Control therefore should be made on the basis of suitability of the circumstances for each, together with the feasibility of performing the significant design exercises. Incorporating explicitly the robust design methodology into the adap-

tive framework, as has been done with LQG/LTR in this oeuvre, opens the prospect of the design of control laws adopting the best of each subdiscipline.

8.7 Coda

We have now run our course with the examination of Adaptive Optimal Control. We took our lead from the practically derived GPC adaptive control and ran with this through a sequence of generalizations and expansions to include LQG, linear system robustness, LTR, frequency response interpretation of closed loop adaptive identification in terms of stability robustness and of achieved performance. This led us to the realization of the two *forces majeures* in adaptive control:

- the effect of the identification of parameters upon the achieved control performance and robustness,
- the effect of the closed loop controller upon the parametric model identified through the influence on the plant input spectrum.

In Chapter 7 we drew these revelations and themes together to propose an adaptive control law, based on LQG and RLS, which possessed a synergism due to the interplay between the controller and identifier acting in concert to support robustness. This was further extended in Chapter 8 to a more general setting, in which the suite of available design variables for the control law part and the identification rule part of an adaptive controller was made apparent and discussed from a synthesis stance.

This is ultimately the message of the book, that the theories of robust linear control system design and of closed loop parameter identification exist and are continually being advanced, and a sensible approach to the design of adaptive controllers must take into account an appreciation of both. We have endeavored to provide a vehicle for the translation of linear systems design tools and identification techniques into the adaptive control sphere. This, we hope, should create a conduit for the rapid acquisition of improved methodologies in adaptive control and for the encouragement of the development of control design and system identification methods better reflecting the needs of adaptive control.

We thank you for your attention.

References

- [**ABJ+86**] B.D.O. Anderson, R.R. Bitmead, C.R. Johnson Jr, P.V. Kokotovic, R.L. Kosut, I.M.Y. Mareels, L. Praly and B.D.Riedle, *Stability of Adaptive Systems: Passivity and Averaging Analysis*, MIT Press, Cambridge MA, USA, 1986.
- [**AM71**] B.D.O. Anderson and J.B. Moore, *Linear Optimal Control*, Prentice-Hall, Englewood Cliffs NJ, USA, 1971.
- [**AM79**] B.D.O. Anderson and J.B. Moore, *Optimal Filtering*, Prentice-Hall, Englewood Cliffs NJ, USA, 1979.
- [**AM90**] B.D.O. Anderson and J.B. Moore, *Optimal Control: Linear Quadratic Methods*, Prentice-Hall, Englewood Cliffs NJ, USA, 1990.
- [**And85**] B.D.O. Anderson, ‘Adaptive Systems, Lack of Persistency of Excitation and Bursting Phenomena,’ *Automatica*, vol. 21, pp. 247–258, 1985.
- [**Åst70**] ³ K.J. Åström, *Introduction to Stochastic Control Theory*, Academic Press, New York, USA, 1970.
- [**Åst83**] K.J. Åström, ‘Analysis of Rohrs’ Counterexample to Adaptive Control,’ *Proc. 22nd I.E.E.E. Conference on Decision and Control*, San Antonio TX, USA, pp. 982–987, 1983.
- [**ÅW73**] K.J. Åström and B. Wittenmark, ‘On Self-tuning Regulators,’ *Automatica*, vol. 9, pp. 185–199, 1973.
- [**ÅW84**] K.J. Åström and B. Wittenmark, *Computer Controlled Systems: Theory and Design*, Prentice-Hall, Englewood Cliffs NJ, USA, 1984.
- [**BA90**] R.R. Bitmead and G.H. Allen, ‘Adaptive Control and Robust Linear Design,’ Research Report, Dept of Systems Engineering, Australian National University, 1990.
- [**BGPK85**] R.R. Bitmead, M. Gevers, I.R. Petersen and R.J. Kaye, ‘Monotonicity and Stabilizability Properties of Solutions of the Riccati Difference Equation: Propositions, Lemmas, Theorems, Fallacious Conjectures and Counterexamples,’ *Systems and Control Letters*, vol. 5, pp. 309–315, 1985.

³A note to our Swedish readership. We apologize for not obeying the ordering of the Swedish alphabet in listing Å at the beginning of the references rather than at the end, where it more properly should be found. We know you are adaptable types.

- [BTP86] R.R. Bitmead, A.C. Tsoi and P.J. Parker, 'A Kalman Filtering Approach to Short Time Fourier Analysis,' *I.E.E.E. Transactions on Acoustics Speech and Signal Processing*, vol. ASSP-34, pp. 1493–1501, 1986.
- [CBG89] P.E. Caines and M. Baykal-Gürsoy, 'On the L_∞ Consistency of L_2 Estimators,' *Systems and Control Letters*, vol. 12, pp. 71–76, 1989.
- [CG79] D.W. Clarke and P.J. Gawthrop, 'Self-tuning Control,' *Proc. IEE*, vol. 123, pp. 633–640, 1979.
- [Cla88] D.W. Clarke, 'Application of Generalized Predictive Control to Industrial Processes,' *I.E.E.E. Control Systems Magazine*, vol. 8, pp. 49–55, 1988.
- [CMT87] D.W. Clarke, C. Mohtadi and P.S. Tuffs, 'Generalized Predictive Control: Parts i and ii,' *Automatica*, vol. 23, pp. 137–160, 1987.
- [CR80] C.R. Cutler and B.L. Ramaker, 'Dynamic Matrix Control: a Computer Control Algorithm,' *Joint American Control Conference*, San Francisco CA, USA, 1980.
- [DGKF89] J.C. Doyle, K. Glover, P.P. Khargonekar and B.A. Francis, 'State-space Solutions to Standard H_2 and H_∞ Control Problems,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-34, pp. 831–847, 1989.
- [dKvC85] R.M.C. de Keyser and A.R. van Cauwenberghe, 'Extended Prediction Self-adaptive Control,' *Proc. 7th IFAC Symposium on Identification and System Parameter Estimation*, York, UK, 1985.
- [Doy78] J.C. Doyle, 'Guaranteed Margins for LQG Regulators,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-23, pp. 756–757, 1978.
- [DS79] J.C. Doyle and G. Stein, 'Robustness with Observers,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-24, pp. 607–611, 1979.
- [dSGG86] C.E. de Souza, M. Gevers and G.C. Goodwin, 'Riccati Equations in Optimal Filtering of Nonstabilizable Systems having Singular State Transition Matrices,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-31, pp. 831–838, 1986.
- [dSo89] C.E. de Souza, 'Monotonicity and Stabilizability results for the Solution of the Riccati Difference Equation,' *Proc. Workshop on the Riccati Equation in Control, Systems and Signals*, S. Bittanti (ed.), Como, Italy, pp. 38–41, 1989.
- [Ega79] B. Egardt, *Stability of Adaptive Controllers*, Springer-Verlag, Berlin, FRG, 1979.
- [FDR88] G. Favier, D. Dubois and C. Rougerie, 'A Review of j -step ahead Predictors,' *8th IFAC Symposium on Identification and System Parameter Estimation*, Beijing, P.R. China, 1988.
- [FG78] B.A. Francis and K. Glover, 'Bounded Peaking in the Optimal Linear Regulator with Cheap Control,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-23, pp. 608–617, 1978.

- [FH77] T.E. Fortmann and K.L. Hitz, *An Introduction to Linear Control Systems*, Marcel-Dekker, New York, USA, 1977.
- [FK88] T. Fujii and P.P. Khargonekar, 'Inverse Problems in H_∞ Control Theory and Linear Quadratic Differential Games,' *27th I.E.E.E. Conference on Decision and Control*, Austin TX, USA, pp. 26–31, 1988.
- [Fra87] B.A. Francis, *A Course in H_∞ Control Theory*, Springer-Verlag, Berlin, FRG, 1987.
- [FW76] B.A. Francis and W.M. Wonham, 'The Internal Model Principle of Control Theory,' *Automatica*, vol. 12, pp. 457–465, 1976.
- [GH83] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems and Bifurcations of Vector Fields*, Springer-Verlag, New York, USA, 1983.
- [GM82] C.E. Garcia and M. Morari, 'Internal Model Control: Part 1. A Unifying Review and Some New Results,' *Industrial Engineering and Chemical Process Design and Development*, vol. 21, pp. 308–323, 1982.
- [GMMZ84] C. Greco, G. Menga, E. Mosca and G. Zappa, 'Performance Improvements of Self Tuning Controllers by Multistep Horizons: the MUSMAR Approach,' *Automatica*, vol. 20, pp. 681–700, 1984.
- [Gri84] M.J. Grimble, 'Implicit and Explicit LQG Self-tuning Controllers,' *Automatica*, vol. 20, pp. 661–669, 1984.
- [GP86] C.E. Garcia and D.M. Prett, 'Advances in Industrial Model Predictive Control,' in *Chemical Process Control – CPC3*, M. Morari and T.J. McAvoy (eds.), pp. 249–294, CAChe and Elsevier, Amsterdam, Netherlands, 1986.
- [GPM89] C.E. Garcia, D.M. Prett and M. Morari, 'Model Predictive Control: Theory and Practice—a Survey,' *Automatica*, vol. 25, pp. 335–348, 1989.
- [GRC80] G.C. Goodwin, P.J. Ramadge and P.E. Caines, 'Discrete-time Multivariable Adaptive Control,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-25, pp. 449–456, 1980.
- [GS84] G.C. Goodwin and K.S. Sin, *Adaptive Filtering, Prediction and Control*, Prentice-Hall, Englewood Cliffs NJ, USA, 1984.
- [Gun88] S. Gunnarsson, *Frequency Domain Aspects of Modeling and Control in Adaptive Systems*, PhD Thesis, Linköping University, Sweden, 1988.
- [GWZ87] R. Gorez, V. Wertz and K.Y. Zhu, 'On a Generalized Predictive Control Algorithm,' *Systems and Control Letters*, vol. 9, pp. 369–377, 1987.
- [HC87] L. Hsu and R. Costa, 'Bursting Phenomena in Continuous-time Adaptive Systems with σ -modification,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-32, pp. 84–86, 1987.
- [Hil59] E. Hille, *Analytic Function Theory, vols I and II*, Chelsea, New York, USA, 1959.

- [HJ85] R.A. Horn and C.A. Johnson, *Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1985.
- [IFF86] E. Irving, C.M. Falinower and C. Fonte, 'Adaptive Generalized Predictive Control with Multiple Reference Models,' *2nd IFAC Workshop on Adaptive Systems in Control and Signal Processing*, Lund, Sweden, 1986.
- [IK83] P.A. Ioannou and P.V. Kokotovic, *Adaptive Systems with Reduced Models*, Springer-Verlag, Berlin, FRG, 1983.
- [Ish88] T. Ishihara, 'Robust Stability Bounds for a Class of Discrete-time Regulators with Computation Delays,' *Automatica*, vol. 24, pp. 697-700, 1988.
- [IT86] T. Ishihara and H. Takeda, 'Loop Transfer Recovery Techniques for Discrete-time Optimal Regulators using Prediction Estimators,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-31, pp. 1149-1151, 1986.
- [Joh88] C.R. Johnson Jr, *Lectures on Adaptive Parameter Estimation*, Prentice-Hall, Englewood Cliffs NJ, USA, 1988.
- [KA86] G. Kreisselmeier and B.D.O. Anderson, 'Robust Model Reference Adaptive Control,' vol. AC-31, pp. 127-133, 1986.
- [Kai75] T. Kailath, 'Some New Algorithms for Recursive Estimation in Constant Linear Systems,' *I.E.E.E. Transactions on Information Theory*, vol. IT-21, pp. 750-760, 1975.
- [Kai80] T. Kailath, *Linear Systems*, Prentice-Hall, Englewood Cliffs NJ, USA, 1980.
- [Kle70] D.L. Kleinman, 'An Easy Way to Stabilize a Linear Constant System,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-15, p. 692, 1970.
- [Kle74] D.L. Kleinman, 'Stabilizing a Discrete Constant Linear System, with Application to Iterative Methods for Solving the Riccati Equation,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-19, pp. 252-254, 1974.
- [KMABJ87] R.L. Kosut, I.M.Y. Mareels, B.D.O. Anderson, R.R. Bitmead and C.R. Johnson Jr, 'Transient analysis of adaptive control,' *Proc. 10th World Congress IFAC*, Munich, July 1987.
- [KP78] W.H. Kwon and A.E. Pearson, 'On Feedback Stabilization of Time-varying Discrete Linear Systems,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-23, pp. 479-481, 1978.
- [KS72] H. Kwakernaak and R. Sivan, *Linear Optimal Control Systems*, Wiley Interscience, New York, USA, 1972.
- [Kwa69] H. Kwakernaak, 'Optimal Low Sensitivity Linear Feedback Systems,' *Automatica*, vol. 5, p. 276, 1969.
- [Lau79a] A.J. Laub, 'An Inequality and some Computations Related to the Robust Stability of Linear Dynamical Systems,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-24, pp. 318-320, 1979.

- [Lau79b] A.J. Laub, 'A Schur Method for Solving Algebraic Riccati Equations,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-24, pp. 913-921, 1979.
- [Lew86a] F.L. Lewis, *Optimal Control*, Wiley Interscience, New York, USA, 1986.
- [Lew86b] F.L. Lewis, *Optimal Estimation*, Wiley Interscience, New York, USA, 1986.
- [LGW89] D.J.N. Limebeer, M. Green and D. Walker, 'Discrete-time H_∞ Control,' *28th I.E.E.E. Conference on Decision and Control*, Tampa FL, USA, pp. 392-396, 1989.
- [Lju87] L. Ljung, *System Identification: Theory for the User*, Prentice-Hall, Englewood Cliffs NJ, USA, 1987.
- [LSA81] N.A. Lehtomaki, N.R. Sandell Jr and M. Athans, 'Robustness Results in Linear Quadratic Gaussian Based Multivariable Control Design,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-26, pp. 75-92, 1981.
- [Lun89] J. Lunze, *Robust Multivariable Feedback Control*, Prentice-Hall International, Hemel Hempstead, UK, 1989.
- [Mac85] J.M. Maciejowski, 'Asymptotic Recovery for Discrete-time Systems,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-30, pp. 602-605, 1985.
- [MC86] C. Mohtadi and D.W. Clarke, 'Generalized Predictive Control, LQ or Pole-placement: a Unified Approach,' *25th I.E.E.E. Conference on Decision and Control*, Athens, pp. 1536-1541, 1986.
- [MCA86] G.D. Martin, J.M. Caldwell and T.E. Ayril, 'Predictive Control Applications for the Petroleum Refining Industry,' *Japan Petroleum Institute — Petroleum Refining Conference*, Tokyo, Japan, 1986.
- [MGHM88] R.H. Middleton, G.C. Goodwin, D.J. Hill and D.Q. Mayne, 'Design Issues in Adaptive Control,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-33, pp. 50-58, 1988.
- [Mit85] T. Mita, 'Optimal Digital Feedback Control Systems Counting Computation Time of Control Laws,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-30, pp. 542-548, 1985.
- [MM87] J.B. Moore and D.L. Mingori, 'Robust Frequency-shaped LQ Control,' *Automatica*, vol. 23, pp. 641-646, 1987.
- [Moh87] C. Mohtadi, *Advanced Self-tuning Algorithms*, PhD Thesis, University of Oxford, 1987.
- [Moh88] C. Mohtadi, 'The Rôle of Prefiltering in Parameter Estimation and Control,' *Proc Workshop on Adaptive Control Analysis for Industrial Use*, Lodge Kananaskis, Alberta, Canada, 1988.
- [MSD87] M. M'Saad, G. Sanchez and M. Duque, 'A Performance Oriented Multivariable Adaptive Generalized Predictive Controller,' *Proc 10th IFAC World Congress*, Munich, FRG, 1987.

- [MSM85] P.R. Maurath, D.E. Seborg, D.A. Mellichamp, 'Predictive Controller Design by Principal Component Analysis,' *American Automatic Control Conference*, Boston MA, USA, 1985.
- [MZ89] M. Morari and E. Zafriou, *Robust Process Control*, Prentice-Hall, Englewood Cliffs NJ, USA, 1989.
- [MZL89] E. Mosca, G. Zappa and J.M. Lemos, 'Robustness of Multipredictor Adaptive Regulators: MUSMAR,' *Automatica*, vol. 25, pp. 521–529, 1989.
- [NA87] K.S. Narendra and A.M. Annaswamy, 'A New Adaptive Law for Robust Adaptation without Persistent Excitation,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-32, pp. 134–145, 1987.
- [Nis67] T. Nishimura, 'On the *a priori* Information in Sequential Estimation Problems,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-12, pp. 123–125, 1967.
- [PBG88] M.A. Poubelle, R.R. Bitmead and M. Gevers, 'Fake Algebraic Riccati Techniques and Stability,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-33, pp. 379–381, 1988.
- [Pet84] V. Peterka, 'Predictor-based Self-tuning Control,' *Automatica*, vol. 20, pp. 39–50, 1984.
- [Pet89] V. Peterka, 'Predictive and LQG Optimal Control: Equivalences, Differences and Improvements,' *Workshop on Control of Uncertain Systems*, Bremen, FRG, 1989.
- [Pol87] J.W. Polderman, *Adaptive Control and Identification: Conflict or Conflux?*, PhD Thesis, University of Groningen, Netherlands, 1987.
- [PPGB86] M.A. Poubelle, I.R. Petersen, M. Gevers and R.R. Bitmead, 'A Miscellany of Results on an Equation of Count J.F. Riccati,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-31, pp. 651–654, 1986.
- [PR82] L. Praly and M. Redjah, 'An Indirect Adaptive Control Scheme for Disturbed MIMO Systems', *Proc. 21st I.E.E.E. Conference on Decision and Control*, Orlando FL, USA, 1982.
- [Pra82] L. Praly, *MIMO Stochastic Adaptive Control*, CAI Report, École des Mines, Fontainebleau, France, 1982.
- [Pra86] L. Praly, 'Global Stability of a Direct Adaptive Control Scheme with Respect to a Graph Topology,' in *Adaptive and Learning Systems: Theory and Applications*, K.S. Narendra (ed.), Plenum Press, New York, USA, 1986.
- [Pra88] L. Praly, 'Oscillatory Behavior and Fixes in Adaptive Linear Control: a Worked Example,' *Proc. IFAC Workshop on Robust Adaptive Control*, Newcastle, Australia, 1988.
- [RK86a] B.D. Riedle and P.V. Kokotovic, 'Integral Manifolds of Slow Adaptation,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-31, pp. 316–323, 1986.

- [RK86b] B.D. Riedle and P.V. Kokotovic, 'Stability Bounds for Slow Adaptation: an Integral Manifold Approach,' *Proc. 2nd IFAC Workshop on Adaptive Systems in Control and Signal Processing*, Lund, Sweden, 1986.
- [RM72] R. Rouhani and R.K. Mehra, 'Model Algorithmic Control: Basic Theoretic Perspectives,' *Automatica*, vol. 18, pp. 401–414, 1972.
- [Ros74] H.H. Rosenbrock, *Computer-aided Control System Design*, Academic Press, New York, USA, 1974.
- [RRTP78] J. Richalet, A. Rault, J.L. Testud and J. Papon, 'Model Predictive Heuristic Control: Application to Industrial Processes,' *Automatica*, vol. 14, pp. 413–428, 1978.
- [RVAS85] C.E. Rohrs, L.S. Valavani, M. Athans and G. Stein, 'Robustness of Continuous-time Adaptive Control Algorithms,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-30, pp. 881–889, 1985.
- [SA87] G. Stein and M. Athans, 'The LQG/LTR Procedure for Multivariable Feedback Control Design,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-32, pp. 105–114, 1987.
- [Sam80] C. Samson, Thèse de Docteur Ingénieur, Université de Rennes I, France, 1980.
- [Sam82] C. Samson, 'Stability Analysis of Adaptively Controlled Systems Subject to Bounded Disturbances,' *Automatica*, vol. 19, pp. 81–86, 1982.
- [Sam83] C. Samson, *Problèmes en Identification et Commande de Systèmes Dynamiques*, PhD Thesis, Université de Rennes I, France, 1983.
- [SB89] S.S. Sastry and M. Bodson, *Adaptive Control: Stability, Convergence and Robustness*, Prentice-Hall, Englewood Cliffs NJ, USA, 1989.
- [SGM88] M.E. Salgado, G.C. Goodwin and R.H. Middleton, 'Modified Least Squares Algorithm Incorporating Exponential Resetting and Forgetting,' *International Journal of Control*, vol. 47, pp. 477–491, 1988.
- [Sha86] U. Shaked, 'Guaranteed Stability Margins for the Discrete-time Linear Quadratic Optimal Regulator,' *I.E.E.E. Transactions on Automatic Control*, vol. AC-31, pp. 162–165, 1986.
- [SMS90] D.S. Shook, C. Mohtadi and S.L. Shah, 'Identification for Long Range Predictive Control,' *submitted for publication*, 1990.
- [SSE86] D.E. Seborg, S.L. Shah and T.F. Edgar, 'Adaptive Control Strategies for Process Control: a Survey,' *AIChE Journal*, vol. 32, pp. 881–913, 1986.
- [Tho75] Y.A. Thomas, 'Linear Quadratic Optimal Estimation and Control with Receding Horizon,' *Electronics Letters*, vol. 11, pp. 19–21, 1975.
- [Vid85] M. Vidyasagar, *Control System Synthesis: a Factorization Approach*, MIT Press, Cambridge MA, USA, 1985.

-
- [**WGZ87**] V. Wertz, R. Gorez and K.Y. Zhu, ‘A New Generalized Predictive Controller: Application to the Control of Processes with Uncertain Deadtime,’ *26th I.E.E.E. Conference on Decision and Control*, Los Angeles CA, USA, 1987.
- [**Wit85**] B. Wittenmark, ‘Design of Digital Controllers, the Servo Problem,’ in *Applied Digital Control*, S.G. Tzafestas (ed.), Elsevier, Amsterdam, Netherlands, 1985.
- [**WKR79**] B. Wittenmark and P. Krishna Rao, ‘Comments on “Single-step versus Multistep Performance Criteria for Steady State SISO Systems”,’ *I.E.E.E. Transactions on Automatic Control*, vol. AC-24, pp. 140–141, 1979.
- [**WL86**] B. Wahlberg and L. Ljung, ‘Design Variables for Bias Distribution in Transfer Function Estimation,’ *I.E.E.E. Transactions on Automatic Control*, vol. AC-31, pp. 134–144, 1986.
- [**Wou77**] W.R.E. Wouters, ‘Single-step versus Multistep Performance Criteria for Steady State SISO Systems,’ *I.E.E.E. Transactions on Automatic Control*, vol. AC-22, pp. 272–279, 1977.
- [**Yds84**] B.E. Ydstie, ‘Extended Horizon Adaptive Control,’ *Proc. 9th IFAC World Congress*, Budapest, Hungary, 1984.

Index

- Achieved cost, 160
- Adaptation gain, 152, 159, 164, 167, 184, 185
- Adaptation, performance-based, 206
- Adaptive H_∞ optimal control, 224
- Adaptive LQG control, 2, 7
- Additive perturbation, 208, 213
- Algebraic Riccati equation (ARE), 47, 85
 - fake, 88
 - filtering, 53, 127, 179
 - LQ control, 125
 - LQ regulator, 181
 - stabilizing solution, 85
- ARIMAX model, 27
- ARX model, 151
- Aryabhatta, 29
- Asymptotic model, 154, 156
- Averaging, 164, 166

- Bad craziness, 159
- Bias distribution, 153, 157, 158

- Candidate robust adaptive predictive (CRAP) controller, 176, 186, 217
 - summary, 186
- Certainty equivalence controller, 176, 177
- Certainty equivalence principle, 27, 206
- Chaussée de Haecht, 205
- Closed loop bandwidth, 185
- Closed loop formulae, 34, 40, 70
 - full state feedback, 70
 - with Kalman filter, 72, 73
 - with Kalman predictor, 71, 73
- Closed loop identification, 14, 155

- Composite model, 1780, 179
- Composite system, 57
- Control horizon, 28
- Cost function, 24–26, 28, 38, 46
- Covariance resetting, 171

- Deadbeat control, 36, 37, 108
- Deadbeat observer, 77
- Deadzone, 169, 184
- Detectability, 88
- Diophantine equation, 29, 37, 38
- Disturbance model, 54, 180
- Disturbance spectra, 54
- DMC, 23
- Dynamical nightie, 148

- EDR, 116
- EHAC, 23
- EPSAC, 23
- Equation error, 185
- Equation error model, 225

- Fallacious conjecture, 92
- FARE, 86
- Feedback control, 46, 58
- Feedback perturbations, 209
- Feedforward, 41
- Finite horizon, regulator, 46
- Finite horizon, tracking problem, 48
- Finite-dimensional optimization, 47
- Free response prediction, 30, 77
- Frequency domain properties, 153, 212
- Frequency weighted norm, 154
- Frequency weighting, 184, 193
 - in LQ control, 220
 - of Kalman filter, 221
- Funny result, 106

- Gain crossover, 185
- Gaussian, 43, 57
- Gedankenexample, 8, 194
- Generalized Minimum Variance control, 25
- Global theory, 168
- GPC, 1, 2, 6, 8, 21, 23, 26, 27, 176, 217
 - instability, 102
 - performance model, 36
 - stability, 101, 107
- Gramian, 96
- H_∞ control, 215, 219, 222
- IDCOM, 23
- IMC, 23
- Impulse response, 31
- Incremental inputs, 25
- Indirect adaptive control, 176, 206
- Infinite horizon, regulator, 47, 181
- Infinite-horizon, tracking problem, 50
- Initial condition, 167
- Input-output model, 149
- Integral action, 28
- Integral manifold theory, 164
- Internal model principle, 28, 54
- Interplay, estimation/control, 6, 15, 147, 171, 176, 206, 212, 227
- Interplay, adaptation/robustness, 114
- Kalman filter, 53, 56
- Kalman filter, design, 179, 180
- Kalman predictor, 52, 56
- L_∞ norm, 5
- Leakage, 170, 184
- Least squares identification, 183
- Linear feedback controller, 35
- Load disturbances, 177
- Load perturbations, 28
- Local theory, 158, 168
- Loop transfer recovery (LTR), 116, 130, 131, 135, 137
 - example, 141
- LQ regulator, 46
- LQ tracking problem, 48
- LQG control, 1, 43, 57, 176, 187
 - design, 181
 - controller, 182
 - transfer function, 132
 - comparative performance, 98
 - loop transfer recovery, 187
 - robustness, 113, 115
- Lyapunov equation, 51, 58, 105, 182
- MAC, 23
- Markov parameters, 31
- Matrix inversion lemma, 126
- Measurement noise, 180
- Mid chapter conclusion, 98
- Minimum delay, 138, 139, 141
- Minimum phase system, 25, 61, 131, 137, 138, 141
- Minimum variance control, 24, 187
- Mismodeling, 148
- Model set, 150, 160, 161, 184
- Monotonicity conditions, 89
- Multi-step-ahead prediction, 219, 225
- Multicriterion optimization, 215
- Multiplicative perturbation, 118, 119, 123, 130, 213
- Multiplicative perturbations, 207, 208
- MUSMAR, 23
- Neglected dynamics, 148
- Nile, 44
- Noise model, 177
- Non-minimum phase system, 9, 25, 138, 140, 143, 194
- Nonidentifiability, 11
- Normalization, 169, 184
- Nyquist criterion, 120
- Observer, 59
 - direct feedthrough, 54, 56, 72–74, 132
 - dynamics, cancellation, 35, 72–74
 - no direct feedthrough, 53, 56, 71, 73, 132
- One-step-ahead predictor, 150
- Open Loop Identification, 10
- Open loop identification, 154

- Optimal control, 43
- Optimal cost, 46
- Optimal prediction, 29, 30
- Optimal state estimator, 51
- Output error, 185
- Output error model, 225
- Overparametrization, 37

- Parameter drift, 194, 195, 197
- Parameter update, 152
- Performance, 215
 - achieved, 99
 - designed, 99
 - of LQ schemes, 98
 - optimal, 99
- Performance model, 38
- Persistence of excitation, 11, 14, 16, 159, 162, 167, 172, 185, 189, 214, 217, 227
- Plant model, 177
- Plant-model mismatch, 148, 161, 167, 184, 188, 208
- Pole placement, 9, 36–38
- Poubelle, 86
- Practitioners, 4, 16, 21, 22
- Prediction error, 150
- Prediction error identification,
 - off-line, 149
 - recursive, 151
- Predictive control, 23, 176
- President, 89
- Principle of the argument, 120
- Process noise, 180
- Projection, 170, 184
- Pseudolinear regression, 150

- Quasistationary signal, 149

- Receding horizon control, 26, 29, 32, 83
 - abandon, 111
 - LQG, 75, 77
 - monotonicity, 85
 - performance, 98, 100
 - regulator, 47
 - stability, 85, 87, 88, 93
 - time-varying strategies, 105
- Recursive least squares, 1, 151, 176, 186, 188
 - frequency weighting, 192
- Reference model, 49
- Reference spectrum, 184
- Regulator, two degree of freedom, 155, 207
- Relative error, 123, 130
- Restricted complexity model, 148
- Return difference, 71, 74, 118, 120, 126, 127, 190, 214
 - equality, 116, 124, 126
 - frequency weighted LQ, 221
 - Kalman predictor, 127
 - LQ controller, 128
 - Kalman predictor, 129, 134
 - minimum singular value, 130
- Riccati difference equation (RDE), 46
 - augmented, 50
 - convergence, 86
 - filtering, 53
 - inverse formulation, 94
 - monotonicity, 89
 - stability, 90
- Robust linear control, 113
- Robustness, 5, 188, 193, 206
 - margin, 129, 131
 - Kalman predictor, 135, 139
 - unity feedback system, 117, 123

- Self-tuning regulator, 24
- Sensitivity condition, 227
- Sensitivity function, 191, 215
 - achieved, 216
 - complementary, 216
 - designed, 216
- Separation principle, 57
- Signal spectra, 154
- Signal to noise ratio, 180
- Singular optimal control, 138
- Singular optimal filtering, 131, 138
- Slow adaptation, 164
- Spectral factorization, 55
- Spectrum, 150
- Stability, 83, 188, 211

margin, 130, 134
Stabilizing feedback controller, 93, 96,
160
Stabilizing model, 162, 164
State-space model, 45
Step reference, 50, 69
Structural stability, 5
Suboptimal control, 98
Synergism, 6, 175, 193, 217

Theoreticians, 4, 16, 22
Time-invariant controller, 34
Time-varying system, 219, 226, 227
Time-varying weighting matrices, 105
Tracking reference model, 37

Undermodeling, 9
Unity feedback system, 118, 210, 211
Unmodeled dynamics, 161, 164
Unstructured uncertainty, 222

Van Compernelle, 29

Working example, 65, 198