## Dynamic Programming

 and Optimal Control Volume I THIRD EDITIONDimitri P. Bertsekas
Massachusetts Institute of Technology

WWW site for book information and http://www.athenasc.com

Athena Scientific, Belmont, Massachusetts

Athena Scientific
Post Office Box 805
Nashma, NH 03061-0805
U.S.A.

ErnaH: info@athenasc.com
WWW: http://www.athenasc.co:m

Cover Design: Ann Gallager, www.gallagerdesign.com
(C) 2005, 2000, 1995 Dimitri P. Bertsekas

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or mlormation storage and retrieval) without permission in writing from the publisher.

## ABOUT THE AUTHOR

Dimitri Bertsekas studied Mechanical and Electrical Engineering at the National Technical University of Athens, Greece, and obtained his Ph.D. in system science from the Massachusetts Institute of Technology. He has held faculty positions with the Engineering-Economic Systems Dept., Stanford University, and the Electrical Engineering Dept. of the University of Illinois, Urbana. Since 1979 he has been teaching at the Electrical Engineering and Computer Science Department of the Massachusetts Institute of Technology (M.LT.), where he is currently McAfee Professor of Engineering.

His research spans several fields, including optimization, control, la,rgescale computation, and data communication networks, and is closely tied to his teaching and book authoring activities. He has written llUInerous research papers, and thirteen books, several of which are used as textbooks in MIT classes. He consults regularly with private industry and has held editorial positions in several journals

Professor Bertsekas was awarded the INFORMS 1997 Prize for Research Excellence in the Interface Between Operations Research and Computer Science for his book "Neuro-Dynamic Programming" (co-authored with John Tsitsiklis), the 2000 Greek National Award for Operations Research, and the 2001 ACC John R. Ragazzini Education Award. In 2001, he was elected to the United States National Academy of Engineering.

## Publisher's Cataloging-in-Publication Data

Bertsekas, Dimitri P.
Dynamic Programming and Optimal Control
Includes Bibliography and Index

1. Mathematical Optimization. 2. Dynamic Programming. L Title.

QA402.5 .13465 2005 00-91281
ISBN 1-886529-26-4

## ATHENA SCIENTIFIC OPTIMIZATION AND COI\1PUTATION SERIES

1. Convex Analysis and Optimization, by Dimitri P. Bertsekas, with Angelia Nedić and Asuman E. Ozdaglar, 2003, ISBN 1-886529-$45-0,560$ pages
2. Introduction to Probability, by Dimitri P. Bertsekas and John N. Tsitsiklis, 2002, ISBN 1-886529-40-X, 430 pages
3. Dynamic Programming and Optimal Control, Two-Volume Set, by Dimitri P. Bertsekas, 2005, ISBN 1-886529-08-6, 840 pages
4. Nonlinear Programming, 2nd Edition, by Dimitri P. Bertsekas, 1999, ISBN 1-886529-00-0, 791 pages
5. Network Optimization: Continuous and Discrete Models, by Dimitri P. Bertsekas, 1998, ISBN 1-886529-02-7, 608 pages
6. Network Flows and Monotropic Optimization, by R. Tyrrell RockareUar, 1998, ISBN 1-886529-06-X, 634 pages
7. Introduction to Linear Optimization, by Dimitris Bertsimas and John N. Tsitsiklis, 1997, ISBN 1-886529-19-1, 608 pages
8. Parallel and Distributed Computation: Numerical Methods, by Dimitri P. Bertsekas and John N. Tsitsiklis, 1997, ISBN 1-886529-01-9, 718 pages
9. Neuro-Dynamic Programming, by Dimitri P. Bertsekas and John N. Tsitsiklis, 1996, ISBN 1-886529-10-8, 512 pages
10. Constra,ined Optimization and Lagrange Multiplier Methods, by Dimitri P. Bertsekas, 1996, ISBN 1-886529-04-3, 410 pages
11. Stochastic Optirnal Control: The Discrete-Time Case, by Dimitri P. Bertsekas and Steven E. Shreve, 1996, ISBN 1-886529-03-5, 330 pages

## Contents

1. The Dynamic Programming Algorithm
1.1. Introduction
p. 2
1.2. The Basic Problem. . . . . . . . . . p. 12
1.3. The Dynamic Programming Algorithm . p. 18
1.4. State Augmentation and Other Reformulations p. 35
1.5. Some Mathematical Issues
p. 42
1.6. Dynamic Prograrnming and Minimax Control p. 46
1.7. Notes, Sources, and Exercises . . . . . . .
2. Deterministic Systems and the Shortest Path Probleln
2.1. Finite-State Systems and Shortest Paths p. 6
2.2. Some Shortest Path Applications p. 68
2.2.1. Critical Path Analysis
p. 68
2.2.2. Hidden Markov Models and the Viterbi Algorithm p. 70
2.3. Shortest Path Algorithms . . . . . . . . . . . p. 77
2.3.1. Label Correcting Methods. . . . . . . . p. 78
2.3.2. Label Correcting Variations - $\boldsymbol{A}^{*}$ Algorithm p. 87
2.3.3. Branch-and-Bound . . . . . . . .
p. 88
2.3.4. Constrained and Multiobjective Problems p. 91
2.4. Notes, Sources, and Exercises .
p. 97
3. Deterministic Continuous-Time Optimal Control
3.1. Continuous-Time Optimal Control
p. 106
3.2. The Hamilton-Jacobi-Bellman Equation
3.3. The Pontryagin Minimum Principle
3.3.1. An Informal Derivation Using the HJB Equation
3.3.2. A Derivation Based on Variational Ideas
3.3.3. Minimum Principle for Discrete-Time Problems
p. 131
3.4.1. Fixed Terminal State p. 131
3.4.2. Free Initial State
3.4.3. Free Terminal Time ..... p. 135
3.4.4. Time-Varying System and Cost ..... p. 1383.4.5. Singular Problems . .p. 139
3.5. Notes, Sources, and Exercises ..... p. 142
4. Problellls with Perfect State Information
4.1. Linear Systems and Quadratic Cost ..... p. 148
4.2. Inventory Controlp. 162
4.3. Dynamic Portfolio Analysis ..... p. 170
4.4. Optimal Stopping Problems ..... p. 176
4.5. Scheduling and the Interchange Argument ..... p. 186
4.6. Set-Membership Description of Uncertainty ..... p. 190
4.6.1. Set-Membership Estimation ..... p. 191
p. 197
4.6.2. Control with Unknown-but-Bounded Disturbances ..... p. 201
5. Problems with Imperfect State Information
5.1. Reduction to the Perfect Information Case
p. 218
p. 218
5.2. Linear Systems and Quadratic Cost ..... p. 229
5.3. Minimum Variance Control of Linear Systems ..... p. 236 ..... p. 251
5.4. Sufficient Statistics and Finite-State Markov Chains
5.4. Sufficient Statistics and Finite-State Markov Chains
5.4.1. The Conditional State Distribution ..... p. 252
5.4.2. Finite-State Systems ..... p. 258
5.5. Notes, Sources, and Exercises ..... p. 270
6. 

## Control

6.1. Certainty Equivalent and Adaptive Control ..... p. 283
6.1.1. Caution, Probing, and Dual Control ..... p. 289
6.1.2. Two-Phase Control and Identifiability ..... p. 291
6.1.3. Certainty Equivalent Control and Identifiability ..... p. 293
6.1.4. Self-Tuning Regulators ..... p. 298
6.2. Open-Loop Feedback Control ..... p. 300
6.3. Limited Lookahead Policies ..... p. 304
6.3.1. Performance Bounds for Limited Lookahead Policies p. 305
6.3.2. Computational Issues in Limited Lookahead ..... p. 310
6.3.3. Problem Approximation - Enforced Decomposition ..... p. 312
6.3.4. Aggregation ..... p. 319
6.3.5. Parametric Cost-to-Go Approximation ..... p. 325
6.4. Rollout Algorithms. ..... p. 335
6.4.1. Discrete Deterministic Problems ..... p. 342
6.4.2. Q-Factors Evaluated by Simulation ..... p. 361
6.4.3. Q-Factor Approximation ..... p. 363

| 6.5. Model Predictive Control and Related Methods | p. 366 |
| :--- | :--- |
| 6.5.1. Rolling Horizon Approximations . . . | p. 367 |
| 6.5.2. Stability Issues in Model Predictive Control | p. 369 |
| 6.5.3. Restricted Structure Policies . . | p. 376 |
| 6.6. Additional Topics in Approximate DP | p. 382 |
| 6.6.1. Discretization . . . . . . . | p. 382 |
| 6.6.2. Other Approximation Approaches | p. 384 |
| 6.7. Notes, Sources, and Exercises . . . . . | p. 386 |

7. Introduction to Infinite Horizon Problems

| 7.1. An Overview . . . . . . . . | p. 402 |
| :--- | ---: |
| 7.2. Stochastic Shortest Path Problems | p. 405 |
| 7.3. Discounted Problems . . . . . . | p. 417 |
| 7.4. Average Cost per Stage Problems | p. 421 |
| 7.5. Semi-Markov Problems . . . | p. 435 |
| 7.6. Notes, Sources, and Exercises . . |  |

Appendix A: Mathematical Review
A.1. Sets . p. 459
A.2. Euclidean Space. p. 460
A.3. Matrices . . . .
p. 460
A.4. Analysis . . .
A.5. Convex Sets and Functions p. 467

Appendix B: On Optimization Theory

| B.1. Optimal Solutions . . . . . . . | p. 468 |
| :--- | :--- | :--- | :--- |
| B.2. Optimality Conditions . . . . . . | p. 470 |
| B.3. Minimization of Quadratic Forms | p. 471 |

Appendix C: On Probability Theory
$\begin{array}{ll}\text { C.1. Probability Spaces. . . } \\ \text { C. } 472 \\ \text { C.2. Random Variables } & \text { p. } 473\end{array}$
C.3. Conditional Probability p. 475

Appendix D: On Finite-State Markov Chains
D.1. Stationary Markov Chains
p. 477
D.2. Classification of States
D.3. Limiting Probabilities
p. 479
D.4. First Passage Times .
p. 479
p. 480

## nemun E: Kalman Filtering

E.1. Least-Squares Estimation . p. 481
E.2. Linear Least-Squares Estimation p. 483
E.3. State Estimation Kalman Filter p. 491
E.4. Stability Aspects . . . . . . . p. 496
E.5. Gauss-Markov Estimators p. 499
E.6. Deterministic Least-Squares Estimation
p. 501

Appendix F: Modeling of Stochastic Linear Systems

| F.1. Linear Systems with Stochastic Inputs | p. 503 |
| :--- | :--- |
| F.2. Processes with Rational Spectrum | p. 504 |
| F.3. The ARMAX Model . . . . . . | p. 506 |

Appendi G: Forrnulating Problems of Decision Under Uncertainty

| G.1. The Problem of Decision Under Uncertainty | p. 507 |
| :--- | :--- | :--- |
| G.2. Expected Utility Theory and Risk . . | p. 511 |
| G.3. Stoehastic Optimal Control Problems | p. 524 |

References p. 529
Index ... p. 541

## CONTENTS OF VOLUIVIE II

1. Infinite Horizon - Discounted Problems
1.1. Minimization of Total Cost Introduction
1.2. Discounted Problems with Bounded Cost per Stage
1.3. Finite-State Systems - Computational Methods
1.3.1. Value Iteration and Error Bounds
1.3.2. Policy Iteration
1.3.3. Adaptive Aggregation
1.3.4. Linear Programming
1.3.5. Limited Lookahead Policies
1.4. The Role of Contraction Mappings
1.5. Scheduling and Multiarmed Bandit Problems
1.6. Notes, Sources, and Exereises
2. Stochastic Shortest Path Problems
2.1. Main Results
2.2. Computational Methods
2.2.1. Value Iteration
2.2.2. Policy Iteration
2.3. Simulation-Based Methods
2.3.1. Policy Evaluation by Monte-Carlo Simulation
2.3.2. Q-Learning
2.3.3. Approximations
2.3.4. Extensions to Discounted Problems
2.3.5. The Role of Parallel Computation
2.4. Notes, Sources, and Exereises
3. Undiscounted Problems
3.1. Unbounded Costs per Stage
3.2. Linear Systems and Quadratic Cost
3.3. Inventory Control
3.4. Optimal Stopping
3.5. Optimal Gambling Strategies
3.6. Nonstationary and Periodic Problems
3.7. Notes, Sourees, and Exercises
4. Average Cost per Stage Problems
4.1. Preliminary Analysis
4.2. Optimality Conditions
4.3. Computational Methods
4.3.1. Value Iteration
4.3.2. Policy Iteration
4.3.3. Linear Programming
4.3.4. Simulation-Based Methods
4.4. Infinite State Space
4.5. Notes, Sources, and Exercises
5. Continuous-Time Problems
5.1. Uniformization
5.2. Queueing Applications
5.3. Semi-Markov Problems
5.4. Notes, Sources, and Exercises

## References

## Preface

This two-volume book is based on a first-year graduate course on dynamic programming and optimal control that I have taught for over twenty years at Stanford University, the University of Illinois, and the Massachusetts Institute of Technology. The course has been typically attended by students from engineering, operations research, economics, and applied mathematics. Accordingly, a principal objective of the book has been to provide a unified treatment of the subject, suitable for a broad audience. In particular, problems with a continuous character, such as stochastic control problems, popular in modern control theory, are simultaneously treated with problems with a discrete character, such as Markovian decision problems, popular in operations research. Furthermore, many applications and examples, drawn from a broad variety of fields, are discussed.

The book may be viewed as a greatly expanded and pedagogically improved version of my 1987 book "Dynamic Programming: Deterministic and Stochastic Models," published by Prentice-Hall. I have included much new material on deterministic and stochastic shortest path problems, as well as a new chapter on continuous-time optimal control problems and the Pontryagin Minimum Principle, developed from a dynamic programming viewpoint. I have also added a fairly extensive exposition of simulationbased approximation techniques for dynamic programming. These techniques, which are often referred to as "neuro-dynamic programming" or "reinforcement learning," represent a breakthrough in the practical application of dynamic programming to complex problems that involve the dual curse of large dimension and lack of an accurate mathematical model. Other material was also augmented, substantially modified, and updated.

With the new material, however, the book grew so much in size that it became necessary to divide it into two volumes: one on finite horizon, and the other on infinite horizon problems. This division was not only. natural in terms of size, but also in terms of style and orientation. The first volume is more oriented towards modeling, and the second is more oriented towards mathematical analysis and computation. I have included in the first volume a final chapter that provides an introductory treatment of infinite horizon problems. The purpose is to make the first volume self-
contained for instructors who wish to cover a modest amount of infinite horizon material in a course that is primarily oriented towards modeling, conceptualization, and finite horizon problems,

Many topics in the book are relatively independent of the others. For example Chapter 2 of Vol. I on shortest path problems can be skipped without loss of continuity, and the same is true for Chapter 3 of Vol. I, which deals with continuous-time optimal control. As a result, the book can be used to teach several different types of courses.
(a) A two-semester course that covers both volumes.
(b) A one-semester course primarily focused on finite horizon problems that covers most of the first volume.
(c) A one-semester course focused on stochastic optimal control that covers Chapters 1, 4, 5, and 6 of Vol. I, and Chapters 1, 2, and 4 of Vol. II.
(d) A one-semester course that covers Chapter 1, about $50 \%$ of Chapters 2 through 6 of Vol. I, and about $70 \%$ of Chapters 1, 2, and 4 of Vol. II. This is the course I usually teach at MIT.
(e) A one-quarter engineering course that covers the first three chapters and parts of Chapters 4 through 6 of Vol. I.
(f) A one-quarter mathematically oriented course focused on infinite horizon problems that covers Vol. II.
The mathematical prerequisite for the text is knowledge of advanced calculus, introductory probability theory, and matrix-vector algebra. A summary of this material is provided in the appendixes. Naturally, prior exposure to dynamic system theory, control, optimization, or operations research will be helpful to the reader, but based on my experience, the material given here is reasonably self-contained.

The book contains a large number of exercises, and the serious reader will benefit greatly by going through them. Solutions to all exercises are compiled in a manual that is available to instructors from the author. Many thanks are due to the several people who spent long hours contributing to this manual, particularly Steven Shreve, Eric Loiederman, Lakis Polyrnenakos, and Cynara Wu.

Dynamic programming is a conceptually simple technique that can be adequately explained using elementary analysis. Yet a mathematically rigorous treatment of general dynamic programming requires the complicated machinery of measure-theoretic probability. My choice has been to bypass the complicated mathematics by developing the subject in generality, while claiming rigor only when the underlying probability spaces are countable. A mathematically rigorous treatment of the subject is carried out in my monograph "Stochastic Optimal Control: The Discrete Time

Case," Academic Press, 1978,t coauthored by Steven Shreve. This monograph complements the present text and provides a solid foundation for the subjects developed somewhat informally here.

Finally, I am thankful to a number of individuals and institutions for their contributions to the book. My understanding of the subject was sharpened while I worked with Steven Shreve on our 1978 monograph. My interaction and collaboration with John Tsitsildis on stochastic shortest paths and approximate dynamic programming have been most valuable. Michael Caramanis, Emmanuel Fernandez-Gaucherand, Pierre Humblet, Lennart Ljung, and John Tsitsiklis taught from versions of the book, and contributed several substantive comments and homework problems. A number of colleagues offered valuable insights and information, particularly David Castanon, Eugene Feinberg, and Krishna Pattipati. NSF provided research support. Prentice-Hall graciously allowed the use of material from my 1987 book. Teaching and interacting with the students at MIT have kept up my interest and excitement for the subject.

Dimitri P. Bertsekas
Spring, 1995
t Note added in the 3rd edition: This monograph was republished by Athena Scientific in 1996, and can also be freely downloaded from the author's www site: http://web.mit.edu/dimitrib/www/home.html.

## Preface to the Second Edition

T'his second edition has expanded by nearly $30 \%$ the coverage of the original. Most of the new material is concentrated in four areas:
(a) In Chapter 4, a section was added on estimation and control of systems with a non-probabilistic (set membership) description of uncertainty. This subject, a personal favorite of the author since it was the subject of his $1971 \mathrm{Ph} . \mathrm{D}$. thesis, has become popular, as minimax and $H_{\text {oo }}$ control methods have gained increased prominence.
(b) Chapter 6 was doubled in size, to reflect the popularity of suboptimal control and neuro-dynamic programming methods. In particular, the coverage of certainty equivalent, and limited lookahead methods has been substantially expanded. Furthermore, a new section was added on neuro-dynamic programming and rollout algorithms, and their applications in combinatorial optimization and stochastic optimal control.
(c) In Chapter 7, an introduction to continuous-time, semi-Markov decisieHl problems was added in a separate last section.
(d) A new appendix was included, which deals with various formulations of problems of decision under uncertainty. The foundations of the minimax and expected utility approaches are framed within a broader context, and some of the aspects of utility theory are discussed.
There are also miscellaneous additions and improvements scattered throughout the text, and a more detailed coverage of deterministic problems is given in Chapter 1. Finally, a new internet-based feature was added to the book, which extends its scope and coverage. Many of the theoretical exercises have been solved in detail and their solutions have been posted in the book's www page
http://www.athenasc.com/dpbook.html
These exercises have been marked with the symbol (www)
I would like to express my thanks to the many colleagues who contributed suggestions for improvement of the second edition.

## Preface to the Third

The third edition contains a substantial amount of new material, particularly on approximate dynamic programming, which has now become one of the principal focal points of the book. In particular:
(a) The subject of minimax control was developed in greater detail, including a new section in Chapter 1, which connects with new material in Chapter 6.
(b) The section on auction algorithms for shortest paths in Chapter 2 was eliminated. These methods are not currently used in dynamic programming, and a detailed discussion has been provided in a chapter from the author's Network Optimization book. This chapter can be freely downloaded from
http://web.mit.edu/dimitrib/www/net.html
(c) A section was added in Chapter 2 on dynamic programming and shortest path algorithms for constrained and multiobjective problems.
(d) The material on sufficient statistics and partially observable Markov decision problems in Section 5.4 was restructured and expanded.
(e) Considerable new material was added in Chapter 6:
(1) An expanded discussion of one-step lookahead policies and associated performance bounds in Section 6.3.1.
(2) A discussion of aggregation methods and discretization of conti-nuous-state problems (see Subsection 6.3.4).
(3) A discussion of model predictive control and its relation to other suboptimal control methods (see Subsection 6.5.2).
(4) An expanded treatment of open-loop feedback control and related methods based on a restricted structure (see Subsection 6.5.3).

I have also added a few exercises, and revised a few sections while preserving their essential content. Thanks are due to Haixia Lin, who worked out several exercises, and to Janey Yu, who reviewed some of thenew sections and gave me valuable feedback.

Dimitri P. Bertsekas
http://web.mit.edu/dimitrib/www/home.html Summer 2005

## The Dynamic Programming

## Algorithm

## Contents

1.1. Introduction
p. 2
1.2. The Basic Problem
p. 12
1.3. The Dynamic Programming Algorithm
p. 18
1.4. State Augmentation and Other Reformulations
1.5. Some Mathematical Issues
1.6. Dynamic Programming and Minimax Control
1.7. Notes, Sources, and Exercises.

Life can only be understood going backwards,
but it lllust be lived going forwards.

## Kierkegaard

### 1.1 INTRODUCTION

This book deals with situations where decisions are made in stages. The outcome of each decision may not be fully predictable but can be anticipated to some extent before the next decision is made. The objective is to minimize a certain cost a mathematical expression of what is considered an undesirable outcome.

A key aspect of such situations is that decisions cannot be viewed in isolation since one must balance the desire for low present cost with the undesirability of high future costs. The dynamic programming technique captures this tradeoff. At each stage, it ranks decisions based on the sum of the present cost and the expected future cost, assuming optimal decision making for subsequent stages.

There is a very broad variety of practical problems that can be treated by dynamic programming. In this book, we try to keep the main ideas uncluttered by irrelevant assumptions on problem structure. To this end, we formulate in this section a broadly applicable model of optimal control of a dynamic system over a finite number of stages (a finite horizon). This model will occupy us for the first six chapters; its infinite horizon version will be the subject of the last chapter as well as Vol. II.

Our basic model has two principal features: (1) an underlying discretetime dynamic system, and (2) a cost function that is additive over time. The dynamic system expresses the evolution of some variables, the system's "state", under the influence of decisions made at discrete instances of time. T'he system has the form

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}, u_{k}\right), \quad k=0,1, \ldots, N-1,
$$

where
$k$ indexes discrete time,
$x_{k}$ is the state of the system and summarizes past information that is relevant for future optimization,
$u_{k}$ is the control or decision variable to be selected at time $k$,
$w_{k}$ is a random parameter (also called disturbance or noise depending on the context),
$N$ is the horizon or number of times control is applied,
and $f_{k}$ is a function that describes the system and in particular the mechanism by which the state is updated.

The cost function is additive in the sense that the cost incurred at time $k$, denoted by $g_{k}\left(x_{k}, u_{k}, w_{k}\right)$, accumulates over time. The total cost is

$$
g N(X N)+\sum_{k=O}^{N-1} g k\left(X k, U k, w_{k}\right)
$$

where $g N(X N)$ is a terminal cost incurred at the end of the process. However, because of the presence of ' $W k$, the cost is generally a random variable and cannot be meaningfully optimized. We therefore formulate the problem as an optimization of the expected cost

$$
E\left\{g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(x_{k}, u_{k}, w_{k}\right)\right\},
$$

where the expectation is with respect to the joint distribution of the random variables involved. The optimization is over the controls $u_{0}, u_{1}, \ldots, u_{N-1}$, but some qualification is needed here; each control $U k$ is selected with some knowledge of the current state $x k$ (either its exact value or some other related information).

A more precise definition of the terminology just used will be given shortly. We first provide some orientation by means of examples.

## Example 1.1.1 (Inventory Control)

Consider a problem of ordering a quantity of a certain item at each of $N$ periods so as to (roughly) meet a stochastic demand, while minimizing the incurred expected cost. Let us denote
$X k$ stock available at the beginning of the $k$ th period,
$U k$ stock ordered (and immediately delivered) at the beginning of the kth period,
$w_{k}$ demand during the kth period with given probability distribution.
We assume that $w_{0}, w_{1}, \ldots, w_{N-1}$ are independent random variables, and that excess demand is backlogged and filled as soon as additional inventory becomes available. Thus, stock evolves according to the discrete-time equation

$$
x_{k+1}=x_{k}+u_{k}-w_{k},
$$

where negative stock corresponds to backlogged demand (see Fig. 1.1.1).
The cost incurred in period $k$ consists of two components:
(a) A cost $r\left(x_{k}\right)$ representing a penalty for either positive stock $X k$ (holding cost for excess inventory) or negative stock $x k$ (shortage cost for unfilled demand).
so as to minimize the expected cost. The meaning of $\mu_{k}$ is that, for each $k$ and each possible value of $X k$,

$$
\mu_{k}\left(x_{k}\right)=\text { amount that should be ordered at time } k \text { if the stock is } X k .
$$

The sequence $\pi \quad\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}$ will be referred to as a policy or control law. For each 'if, the corresponding cost for a fixed initial stock $x_{0}$ is

$$
J_{\pi}\left(x_{0}\right)=E\left\{R\left(x_{N}\right)+\sum_{k=0}^{N-1}\left(r\left(x_{k}\right)+c \mu_{k}\left(x_{k}\right)\right)\right\}
$$

and we want to minimize $J_{\pi}\left(x_{0}\right)$ for a given $X o$ over all $\pi$ that satisfy the constraints of the problem. This is a typical dynamic programming problem. We will analyze this problem in various forms in subsequent sections. For example, we will show in Section 4.2 that for a reasonable choice of the cost function, the optimal ordering policy is of the form

$$
\mu_{k}\left(x_{k}\right)= \begin{cases}S_{k}-x_{k} & \text { if } X k<S_{k} \\ 0 & \text { otherwise }\end{cases}
$$

where $S k$ is a suitable threshold level determined by the data of the problem. In other words, when stock falls below the threshold $S k$, order just enough to bring stock up to $S k$.

The preceding example illustrates the main ingredients of the basic problem formulation:
(a) A discrete-time system of the form

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}, w_{k}\right)
$$

where $f_{k}$ is some function; for example in the inventory case, we have $f k\left(X k, U k, w_{k}\right)=x_{k}-1-u_{k}-w_{k}$.
(b) Independent random parameters $w_{k}$. This will be generalized by allowing the probability distribution of $w_{k}$ to depend on $X k$ and $U k$; in the context of the inventory example, we can think of a situation where the level of demand $w_{k}$ is influenced by the current stock level $X k$.
(c) A control constraint; in the example, we have $u_{k} \geq 0$. In general, the constraint set will depend on $X k$ and the time index $k$, that is, $u_{k} \mathrm{E} U k(X k)$. To see how constraints dependent on $X k$ can arise in the inventory context, think of a situation where there is an upper bound $B$ on the level of stock that can be accommodated, so $U k \leq B \quad x k .{ }^{\prime}$
(d) An additive cost of the form

$$
E\left\{g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(x_{k}, z_{k}, w_{k}\right)\right\}
$$

where $g k$ are some functions; in the inventory example, we have

$$
g_{N}\left(x_{N}\right) \quad g_{k}\left(x_{k}, u_{k}, w_{k}\right)=r\left(x_{k}\right)+c u_{k}
$$

(e) Optimization over (closed-loop) policies, that is, rules for choosing Uk for each $k$ and each possible value of $x k$.

## Discrete-State and Finite-State Problems

In the preceding example, the state $x k$ was a continuous real variable, and it is easy to think of multidimensional generalizations where the state is an n-dimensional vector of real variables. It is also possible, however, that the state takes values from a discrete set, such as the integers

A version of the inventory problem where a discrete viewpoint is more natural arises when stock is measured in whole units (such as cars), each of which is a significant fraction of $x k, U k$, or $W k$. It is more appropriate then to take as state space the set of all integers rather than the set of real numbers. The form of the system equation and the cost per period will, of course, stay the same.

Generally, there are many situations where the state is naturally discrete and there is no continuous counterpart of the problem. Such situations are often conveniently specified in terms of the probabilities of transition between the states. What we need to know is $\operatorname{Pij}(u, k)$, which is the probability at time $k$ that the next state will be $j$, given that the current state is $i$, and the control selected is $u$, Le.,

$$
p_{i j}(u, k) \quad P\left\{x_{k+1}=j \mid x_{k}=i, u_{k}=u\right\} .
$$

This type of state transition can alternatively be described in terms of the discrete-time system equation

$$
x_{k+1}=w_{k}
$$

where the probability distribution of the random parameter $W k$ is

$$
P\left\{w_{k}=j \mid x_{k}=i, u_{k}=u\right\}=p_{i j}(u, k)
$$

Conversely, given a discrete-state system in the form

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}, w_{k}\right)
$$

together with the probability distribution $P k\left(\left.W k\right|_{X k, U k)}\right.$ of $W k$, we can provide an equivalent transition probability description. The corresponding transition probabilities are given by

$$
P_{i j}(u, k)=P_{k}\left\{W_{k}(i, u, j) \mid x_{k}=i, u_{k}=u\right\}
$$

where $W(i, u, j)$ is the set

$$
W_{k}(i, u, j)=\left\{w \mid j=f_{k}(i, u, w)\right\}
$$

Thus a discrete-state system can equivalently be described in terms of a difference equation or in terms of transition probabilities. Depending on the given problem, it may be notationally or mathematically more convenient to use one description over the other.

The following examples illustrate discrete-state problems. The first example involves a deterministic problem, that is, a problem where there is no stochastic uncertainty. In such a problem, when a control is chosen at a given state, the next state is fully determined; that is, for any state $i$, control $u$, and time $k$, the transition probability $\operatorname{Pij}(u, k)$ is equal to 1 for a single state $j$, and it is 0 for all other candidate next states. The other three examples involve stochastic problems, where the next state resulting from a given choice of control at a given state cannot be determined a priori.

## Example 1.1.2 (A Deterministic Scheduling Problem)

Suppose that to produce a certain product, four operations must be performed on a certain machine. The operations are denoted by A, B, C, and D. We assume that operation B can be performed only after operation A has been performed, and operation $D$ can be performed only after operation $B$ has been performed. (Thus the sequence CDAB is allowable but the sequence CDBA is not.) The setup cost $C_{m n}$ for passing from any operation $\pi /$, to any other operation $n$ is given. There is also an initial startup cost $S_{A}$ or $S c$ for starting with operation A or C , respectively. The cost of a sequence is the sum of the setup costs associated with it; for example, the operation sequence ACDB has cost

$$
S_{A}+C_{A C}+C_{C D}+C_{D B} .
$$

We can view this problem as a sequence of three decisions, namely the choice of the first three operations to be performed (the last operation is determined from the preceding three). It is appropriate to consider as state the set of operations already performed, the initial state being an artificial state corresponding to the beginning of the decision process. The possible state transitions corresponding to the possible states and decisions for this problem is shown in Fig. 1.1.2. Here the problem is deterministic, Le., at a given state, each choice of control leads to a uniquely determined state. For example, at state AC the decision to perform operation D leads to state ACD with certainty, and has cost $C C D$. Deterministic problems with a finite number of states can be conveniently represented in terms of transition graphs' such as the one of Fig. 1.1.2. The optimal solution corresponds to the path that starts at the initial state and ends at some state at the terminal time and has minimum sum of arc costs plus the terminal cost. We will study systematically problems of this type in Chapter 2.


Figure 1.1.2 The transition graph of the deterministic scheduling problem of Exarnple 1.1.2. Each arc of the graph corresponds to a decision leading from some state (the start node of the arc) to some other state (the end node of the arc). The corresponding cost is shown next to the arc. The cost of the last operation is shown as a terminal cost next to the terminal nodes of the graph.

## Exarnple 1.1.3 (Machine Replacement)

Consider a problem of operating efficiently over $N$ time periods a machine that can be in anyone of $n$ states, denoted $1,2, \ldots, n$. We denote by $g(i)$ the operating cost per period when the machine is in state $i$, and we assume that

$$
g(l) \leq g(2) \leq \ldots \leq g(n)
$$

The implication here is that state $i$ is better than state $i+1$, and state 1 corresponds to a machine in best condition.

During a period of operation, the state of the machine can become worse or it may stay unchanged. We thus assume that the transition probabilities

$$
P_{i j}=P\{\text { next state will be } \mathrm{j} \mid \text { current state is } \mathrm{i}\}
$$

satisfy

$$
p_{i j}=0 \quad \text { if } j<\mathrm{i}
$$

We assume that at the start of each period we know the state of the machine and we must choose one of the following two options:
(a) Let the machine operate one more period in the state it currently is.
(b) Repair the machine and bring it to the best state 1 at a cost $R$.

We assume that the machine, once repaired, is guaranteed to stay in state 1 for one period. In subsequent periods, it may deteriorate to states $j>1$ according to the transition probabilities $P l j$.

Thus the objective here is to decide on the level of deterioration (state) at which it is worth paying the cost of machine repair, thereby obtaining the benefit of smaller future operating costs. Note that the decision should also be affected by the period we are in. For example, we would be less inclined to repair the machine when there are few periods left.

The system evolution for this problem can be described by the graphs of Fig. 1.1.3. These graphs depict the transition probabilities between various pairs of states for each value of the control and are known as transition probability graphs or simply transition graphs. Note that there is a different graph for each control; in the present case there are two controls (repair or not repair).


Figure 1.1.3 Machine replacement example. Transition probability graphs for each of the two possible controls (repair or not repair). At each stage and state i, the cost of repairing is $R+g(l)$, and the cost of not repairing is $g(i)$. The terminal cost is 0 .

## ExanJlple 1.1.4 (Control of a Queue)

Consider a queueing systern with room for $n$ customers operating over $N$ time periods. We assume that service of a customer can start (end) only at the beginning (end) of the period and that the system can serve only one customer at a time. The probability $P m$ of $m$ customer arrivals during a period is given, and the numbers of arrivals in two different periods are independent. Customers finding the system full depart without attempting to enter later. The system offers two kinds of service, fast and slow, with cost per period cf and cs, respectively. Service can be switched between fast and slow at the beginning of each period. With fast (slow) service, a customer in service at the beginning of a period will terminate service at the end of the period with probability $q f$ (respectively, $q s$ ) independently of the number of periods the customer has been in service and the number of customers in the system ( $q f>q s$ ). There is a cost rei) for each period for which there are i customers in the system. There is also a terminal cost $R(\mathrm{i})$ for i customers left in the system at the end of the last period.

The problem is to choose, at each period, the type of service as a function of the number of customers in the system so as to minimize the expected total cost over $N$ periods. One expects that when there is a large number of customers i in queue, it is better to use the fast service, and the question is to find the values of $i$ for which this is true.

Here it is appropriate to take as state the number $i$ of customers in the system at the start of a period and as control the type of service provided. Then, the cost per period is rei) plus $C_{f}$ or $c_{s}$ depending on whether fast or slow service is provided. We derive the transition probabilities of the system.

When the system is empty at the start of the period, the probability that the next state is $j$ is independent of the type of service provided. It equals the given probability of.$j$ customer arrivals when $j<n$,

$$
p_{0 j}\left(u_{f}\right)=p_{0 j}\left(u_{s}\right)=p_{j}, \quad \mathrm{j}=0,1, \ldots, n-1,
$$

and it equals the probability of $n$ or more customer arrivals when $j=n$,

$$
p_{0 n}\left(u_{f}\right)=p_{0 m}\left(u_{s}\right)=\sum_{m=n} p_{m} .
$$

When there is at least one customer in the system (i>0), we have

$$
\begin{gathered}
p_{i j}\left(u_{f}\right)=0, \quad \text { if } \mathrm{j}<\mathbf{i}-\mathbf{I} \\
p_{i j}\left(u_{f}\right)=q_{f} p_{0}, \quad \text { if } j=\mathrm{i} \quad 1,
\end{gathered}
$$

$p_{i j}\left(u_{f}\right)=P\{j \quad i+1$ arrivals, service completed $\}$
$+P\{j-\mathrm{i}$ arrivals, service not completed $\}$
$=q_{f} p_{j-i+1}+(1-q f) P . i-i, \quad$ if $\mathrm{i} \quad 1<j<n-1$,
$p_{i(n-1)}\left(u_{f}\right)=q_{f} \sum_{r n=n-i} p_{m}+\left(1-q_{f}\right) p_{n-1-i}$,

$$
p_{i n}\left(u_{f}\right)=\left(\begin{array}{ll}
1 & \text { qj}
\end{array}\right) \sum_{r n=n-i}^{o o} p m .
$$

The transition probabilities when slow service is provided are also given by these formulas with $u_{f}$ and $q f$ replaced by $u_{s}$ and qs, respectively.

## Example 1.1.5 (Optimizing a Chess Match Strategy)

A player is about to playa two-game chess match with an opponent, and wants to maximize his winning chances. Each game can have one of two outcomes:
(a) A win by one of the players (1 point for the winner and 0 for the loser).
(b) A draw ( $1 / 2$ point for each of the two players).

If the score is tied at $1-1$ at the end of the two games, the match goes into sudden-death mode, whereby the players continue to play until the first time one of them wins a game (and the match). The player has two playing styles and he can choose one of the two at will in each game, independently of the style he chose in previous games.
(1) Timid play with which he draws with probability $p_{d}>0$, and he loses with probability $\left(1-p_{d}\right)$.
(2) Bold play with which he wins with probability ' $\mathrm{P} w$, and he loses with probability $(1-P w)$.

Thus, in a given game, timid play never wins, while bold play never draws. The player wants to find a style selection strategy that maximizes his probability of winning the match. Note that once the match gets into sudden death, the player should play bold, since with timid play he can at best prolong the sudden death play, while running the risk of losing. Therefore, there are only two decisions for the player to make, the selection of the playing strategy in the first two games. Thus, we can model the problem as one with two stages, and with states the possible scores at the start of each of the first two stages (games), as shown in Fig. 1.1.4. The initial state is the initial score $0-0$. The transition probabilities for each of the two different controls (playing styles) are also shown in Fig. 1.1.4. There is a cost at the terminal states: a cost of -1 at the winning scores $2-0$ and $1.5-0.5$, a cost of 0 at the losing scores $0-2$ and $0.5-1.5$, and a cost of $-p_{w}$ at the tied score $1-1$ (since the probability of winning in sudden death is $P w$ ). Note that to maximize the probability $P$ of winning the match, we must minimize -Po

This problem has an interesting feature. One would think that if $p_{w}<$ $1 / 2$, the player would have a less than $50-50$ chance of winning the match, even with optimal play, since his probability of losing is greater than his probability of winning anyone game, regardless of his playing style. This is not so, however, because the player can adapt his playing style to the current score, but his opponent does not have that option. In other words, the player can use a closed-loop strategy, and it will be seen later that with optimal play, as determined by the dynamic programming algorithm, he has a better than

50-50 chance of winning the match provided $p w$ is higher than a threshold value $\bar{p}$, which, depending on the value of $P d$, may satisfy $\bar{p}<1 / 2$.


0-0


1st Game / Timid Play
(2-0)


0-2


1st Game / Bold Play


2nd Game / Bold Play

Figure 1.1.4 Chess match example. Transition probability graphs for each of the two possible controls (timid or bold play). Note here that the state space is not the same at each stage. The terminal cost is -1 at the winning final scores 2-0 and $1.5-0.5,0$ at the losing final scores $0-2$ and $0.5-1.5$, and $-P w$ at the tied score 1-1.

### 1.2 THE BASIC PROBLEM

We now formulate a general problem of decision under stochastic uncertainty over a finite number of stages. This problem, which we call basic, is central in this book. We will discuss solution methods for this problem
based on dynamic programming in the first six chapters, and we will ex-tenc-lour analysis to versions of this problem involving an infinite number of stages in the last chapter and in Vol. II of this work.

The basic problem is very general. In particular, we will not require that the state, control, or random parameter take a finite number of values or belong to a space of $n$-dimensional vectors. A surprising aspect of dynamic programming is that its applicability depends very little on the nature of the state, control, and random parameter spaces. For this reason $\mathrm{is}_{\mathrm{i}}$ is convenient to proceed without any assumptions on the structure of these spaces; indeed such assumptions would become a serious impediment later.

## Basic Problem

We are given a discrete-time dynamic system

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}, w_{k}\right), \quad k=0,1, \ldots, N-1
$$

where the state $x k$ is an element of a space $S k$, the control $u_{k}$ is an element of a space $O k$ :, and the random "disturbance" $w_{k}$ is an element of a space Dk.

The control $u_{k}$ is constrained to take values in a given nonempty subset $U(X k)$ C $O k$, which depends on the current state $x_{k}$; that is, $u_{k} \mathrm{E}$ $U_{k}\left(x_{k}\right)$ for all $X k E S k$ and $k$.

The random disturbance $w_{k}$ is characterized by a probability distribution $\left.\left.P k C^{\prime}\right|_{X k,} u_{k}\right)$ that may depend explicitly on $X k$ and $u_{k}$ but not on values of prior disturbances $w_{k-1}, \ldots, w_{0}$.

We consider the class of policies (also called control laws) that consist of a sequence of functions

$$
\pi=\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}
$$

where $\mu_{k}$ maps states $X k$ into controls $u_{k}=\mu_{k}\left(x_{k}\right)$ and is such that $\mu_{k}\left(x_{k}\right) \mathrm{E} U k(X k:)$ for all $x_{k} \mathrm{E}$ Sk. Such policies will be called admissible.

Given an initial state $x_{0}$ and an admissible policy $\pi \quad\left\{\mu_{0}, \ldots, \mu_{N--I}\right\}$ the states $x k$ and disturbances $w_{k}$ are random variables with distributions defined through the system equation

$$
\begin{equation*}
x_{k+1}=f_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right), \quad k=0,1, \ldots, N \quad 1 \tag{1.1}
\end{equation*}
$$

Thus, for given functions $g k, k=0,1, \ldots, N$, the expected cost of $\pi$ starting at $X o$ is

$$
J_{\pi}\left(x_{0}\right)=E\left\{g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right\}
$$

where the expectation is taken over the random variables $w k$ and $x k$. An optimal policy $\pi^{*}$ is one that minimizes this cost; that is,

$$
J_{\pi^{*}}\left(x_{0}\right)=\min _{\pi \in I I} J_{\pi}\left(x_{0}\right)
$$

where II is the set of all admissible policies.
Note that the optimal policy $\pi^{*}$ is associated with a fixed initial state $x_{0}$. However, an interesting aspect of the basic problem and of dynamic programming is that it is typically possible to find a policy $\pi^{*}$ that is simultaneously optimal for all initial states.

The optimal cost depends on $X o$ and is denoted by $J *(x o)$; that is,

$$
J^{*}\left(x_{0}\right)=\min _{\pi \in \Pi} J_{\pi}\left(x_{0}\right)
$$

It is useful to view $\mathrm{J}^{*}$ as a function that assigns to each initial state $X o$ the optimal cost $J^{*}\left(x_{0}\right)$ and call it the optimal cost function or optimal value function. $\dagger$

## The Role and Value of Information

We noted earlier the distinction between open-loop minimization, where we select all controls $I I, 0, \ldots, U N-l$ at once at time 0 , and closed-loop minimization, where we select a policy $\left\{\mu_{0}, \ldots ., \mu_{N-1}\right\}$ that applies the control $\mu_{k}\left(x_{k}\right)$ at time $k$ with knowledge of the current state $x_{k}$ (see Fig. 1.2.1). With closed-loop policies, it is possible to achieve lower cost, essentially by taking advantage of the extra information (the value of the current state). The reduction in cost may be called the value of the information and can be significant indeed. If the information is not available, the controller cannot adapt appropriately to unexpected values of the state, and as a result the cost can be adversely affected. For example, in the inventory control example of the preceding section, the information that becomes available at the beginning of each period $k$ is the inventory stock $x k$. Clearly, this information is very important to the inventory manager, who will want to adjust the amount $U k$ to be purchased depending on whether the current stock $x k$ is running high or low.

[^0]

Figure 1.2.1 Information gathering in the basic problem. At each time $k$ the controller observes the current state $x_{k}$ and applies a control $U_{k}=\mu_{k}\left(x_{k}\right)$ that depends on that state.

## Example 1.2.1

To illustrate the benefits of the proper use of information, let us consider the chess match example of the preceding section. There, a player can select timid play (probabilities $P d$ and $1-P d$ for a draw and a loss, respectively) or bold play (probabilities $p w$ and $1 \quad p w$ for a win and a loss, respectively) in each of the two games of the match. Suppose the player chooses a policy of playing timid if and only if he is ahead in the score, as illustrated in Fig. 1.2.2; we will see in the next section that this policy is optimal, assuming $P d>P w$. Then after the first game (in which he plays bold), the score is 1-0 with probability $p w$ and $0-1$ with probability $1 \quad P w$. In the second game, he plays timid in the former case and bold in the latter case. Thus after two games, the probability of a match win is $p w P d$, the probability of a match loss is $(1-P w) 2$, and the probability of a tied score is $P w(1-P d)+(1-P w) P w$, in which case he has a probability $p w$ of winning the subsequent sudden-death game. Thus the probability of winning the match with the given strategy is

$$
p_{w} p_{d}+p_{w}\left(p_{w}\left(1-p_{d}\right)+\left(1-p_{w}\right) p_{w}\right)
$$

which, with some rearrangement, gives

$$
\begin{equation*}
\text { Probability of a match win }=p_{w}^{2}(2-P w)+P w(l \quad P w) P d . \tag{1.2}
\end{equation*}
$$

Suppose now that $p w<1 / 2$. Then the player has a greater probability of losing than winning anyone game, regardless of the type of play he uses. From this we can infer that no open-loop strategy can give the player a greater than $50-50$ chance of winning the match. Yet from Eq. (1.2) it can be seen that with the closed-loop strategy of playing timid if and only if the player is ahead in the score, the chance of a match win can be greater than 50-50, provided that $p_{w}$ is close enough to $1 / 2$ and $P d$ is close enough to 1 . As an example, for $p_{w}=0.45$ and $P d=0.9$, Eq. (1.2) gives a match win probability of roughly 0.53 .


Figure 1.2.2 Illustration of the policy used in Example 1.2.1 to obtain a greater than 50-50 chance of winning the chess match and associated transition probabilities. The player chooses a policy of playing timid if and only if he is ahead in the score.

To calculate the value of information, let us consider the four open-loop policies, whereby we decide on the type of play to be used without waiting to see the result of the first game. These are:
(1) Play timid in both games; this has a probability $p_{d}^{2} p_{w}$ of winning the match.
(2) Play bold in both games; this has a probability $p_{w}^{2}+2 p_{w}^{2}(1-P w)=$ $p_{w}^{2}(3 \quad 2 p w)$ of winning the match.
(3) Play bold in the first game and timid in the second game; this has a probability $P w P d+p_{w}^{2}(1-P d)$ of winning the match.
(4) Play timid in the first game and bold in the second game; this also has a probability $P w P d+p_{w}^{2}\left(\begin{array}{ll}1 & P d\end{array}\right)$ of winning the match.
The first policy is always dominated by the others, and the optimal open-loop probability of winning the match is

Open-loop probability of win $=\max \left(p_{w}^{2}(3-2 p w), P w P d+p_{w}^{2}(1-P d)\right)$

$$
\begin{equation*}
=p_{w}^{2}+P w(l-P w) \max (2 p w, p d) \tag{1.3}
\end{equation*}
$$

Thus if $p_{d}>2 p w$, we see that the optimal open-loop policy is to play timid in one of the two games and play bold in the other, and otherwise it is optimal to play bold in both games. For $P w=0.45$ and $p_{d}=0.9$, Eq. (1.3) gives an optimal open-loop match win probability of roughly 0.425 . Thus, the value of the information (the outcome of the first game) is the difference of the optimal closed-loop and open-loop values, which is approximately $0.53-0.425=0.105$.

More generally, by subtracting Eqs. (1.2) and (1.3), we see that

$$
\begin{aligned}
\text { Value of information }= & p_{w}^{2}(2-P w)+P w\left(1 \quad p_{w}\right) p_{d} \\
& -p_{w}^{2}-P w(l-P w) \max (2 p w, p e t) \\
= & P w(l-P w) \min (p w, p d-P w) .
\end{aligned}
$$

It should be noted, however, that whereas availability of the state information cannot hurt, it may not result in an advantage either. For instance, in deterministic problems, where no random disturbances are present, one can predict the future states given the initial state and the sequence of controls. Thus, optimization over all sequences $\left\{u_{0}, u_{1}, \ldots, u_{N}-1\right\}$ of controls leads to the same optimal cost as optimization over all admissible policies. The same can be true even in some stochastic probleuls (see for example Exercise 1.13). This brings up a related issue. Assuming no information is forgotten, the controller actually knows the prior states and controls $x_{0}, u_{0}, \ldots, x_{k-1}, u_{k-1}$ as well as the current state $x_{k}$. Therefore, the question arises whether policies that use the entire system history can be superior to policies that use just the current state. The answer turns out to be negative although the proof is technically complicated (see [BeS78]). The intuitive reason is that, for a given time $k$ and state $x_{k}$, all future expected costs depend explicitly just on $x k$ and not on prior history.

## Encoding Risk in the Cost Function

As mentioned above, an important characteristic of stochastic problems is the possibility of using information with advantage. Another distinguishing characteristic is the need to take into account risk in the problem formulation. For example, in a typical investment problem one is not only interested in the expected profit of the investment decision, but also in its variance: given a choice between two investments with nearly equal expected profit and markedly different variance, most investors would prefer the investment with smaller variance. This indicates that expected value of cost or reward need not be the most appropriate yardstick for expressing a decision maker's preference between decisions.

As a more dramatic example of the need to take risk into account when formulating optimization problems under uncertainty, consider the so-called St. Petersburg paradox. Here, a person is offered the opportunity of paying $x$ dollars in exchange for participation in the following game: a fair coin is flipped sequentially and the person is paid $2^{k}$ dollars, where $k$ is the number of times heads have come up before tails come up for the first time. The decision that the person must make is whether to accept or reject participation in the game. Now if he accepts, his expected profit
from the game is

$$
\sum_{k=O}^{\infty 0} \frac{1}{2^{k}:+1} \cdot 2^{k}-x=00
$$

so if his acceptance eriterion is based on maximization of expected profit, he is willing to pay any amount $x$ to enter the game. This, however, is in strong disagreement with observed behavior, due to the risk element involved in entering the game, and shows that a different formulation of the problem is needed. The formulation of problems of deeision under uncertainty so that risk is properly taken into aeeount is a deep subject with an interesting theory. An introduction to this theory is given in Appendix G. It is shown in particular that minimization of expected cost is appropriate under reasonable assumptions, provided the cost function is suitably chosen so that it properly eneodes the risk preferences of the deeision maker.

### 1.3 THE DYNAMIC PROGRAMMING ALGORITHM

The dynamie programming (DP) technique rests on a very simple idea, the principle of optimality. The name is due to Bellman, who contributed a great deal to the popularization of DP and to its transformation into a systematic tool. Roughly, the principle of optimality states the following rather obvious fact.

## Pri $-:$ of Optimality

Let $\pi^{*} \quad\left\{\mu_{0}^{*}, \mu_{1}^{*}, \ldots, \mu_{N}^{*}-I\right\}$ be an optimal policy for the basic problem, and assume that when using $\pi^{*}$, a given state $x i$ occurs at time i with positive probability. Consider the subproblem whereby we are at $X_{i}$ at time i and wish to minimize the "cost-to-go" from time i to time $N$

$$
E\left\{g_{N}\left(x_{N}\right)+\sum_{k=i}^{N-1} g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right\} .
$$

Then the truncated poliey $\left\{\mu_{i}^{*}, \mu_{i+1}^{*}, \ldots, \mu_{N^{-}}^{*}\right\}$ is optimal for this subproblem.

The intuitive justification of the prineiple of optimality is very simple. If the truncated policy $\left\{\mu_{i}^{*}, \mu_{i+1}^{*}, \ldots, \mu_{N^{-}}^{*}\right.$ I $\}$ were not optimal as stated, we would be able to reduce the cost further by switching to an optimal policy for the subproblem once we reach $x i$. For an auto travel analogy, suppose that the fastest route from Los Angeles to Boston passes through Chicago. The principle of optimality translates to the obvious fact that the Chicago to Boston portion of the route is also the fastest route for a trip that starts from Chicago and ends in Boston.

The principle of optimality suggests that an optimal policy can be constructed in piecemeal fashion, first constructing an optimal policy for the "tail subproblem" involving the last stage, then extending the optimal policy to the "tail subproblem" involving the last two stages, and continuing in this manner until an optimal policy for the entire problem is constructed. The DP algorithm is based on this idea: it proceeds sequentially, by solving all the tail subproblems of a given time length, using the solution of the tail subproblems of shorter time length. We introduce the algorithm with two examples, one deterministic and one stochastic.

## The DP Algorithm for a Deterministic Scheduing sumarine

Let us consider the scheduling example of the preceding section, and let us apply the principle of optimality to calculate the optimal schedule. We have to schedule optimally the four operations A, B, C, and D. The transition and setup costs are shown in Fig. 1.3.1 next to the corresponding arcs.

According to the principle of optimality, the "tail" portion of an optimal schedule must be optimal. For example, suppose that the optimal schedule is CABD. Then, having scheduled first C and then A, it must be optimal to complete the schedule with BD rather than with DB. With this in mind, we solve all possible tail subproblems of length two, then all tail subproblems of length three, and finally the original problem that has length four (the subproblems of length one are of course trivial because there is only one operation that is as yet unscheduled). As we will see shortly, the tail subproblems of length $k+1$ are easily solved once we have solved the tail subproblems of lengt.h $k$, and this is the essence of the DP technique.

Tail Subproblems of Length 2: These subproblems are the ones that involve two unscheduled operations and correspond to the states AB, AC, CA, anel CD (see Fig. 1.3.1)

State $A B$ : Here it is only possible to schedule operation C as the next operation, so the optimal cost of this subproblem is 9 (the cost of scheduling C after B, which is 3, plus the cost of scheduling Dafter C, which is 6).
State $A C$ : Here the possibilities are to (a) schedule operation 13 and then D, which has cost 5 , or (b) schedule operation D anel then B, which has cost 9 . The first possibility is optimal, and the corresponding cost of the tail subproblem is 5, as shown next to node AC in Fig 1.3.1.

State CA: Here the possibilities are to (a) schedule operation 13 and then D , which has cost 3 , or (b) schedule operation D and then 13 , which has cost 7. The first possibility is optimal, and the correspond-


Figure 1.3.1 Transition graph of the deterministic scheduling problem, with the cost of each decision shown next to the corresponding arc. Next to each node/state we show the cost to optimally complete the schedule starting from that state. This is the optimal cost of the corresponding tail subproblem (ef. the principle of optimality). The optimal cost for the original problem is equal to 10 , as shown next to the initial state. The optimal schedule corresponds to the hick-line arcs.
ing cost of the tail subproblem is 3 , as shown next to node CA in Fig. 1.3.1.

State $C D$ : Here it is only possible to schedule operation A as the next operation, so the optimal cost of this subproblem is 5 .

Tail Subproblems of Length 3: These subproblems can now be solved using the optimal costs of the subproblems of length 2 .

State A: Here the possibilities are to (a) schedule next operation B (cost 2 ) and then solve optimally the corresponding subproblem of length 2 (cost 9 , as computed earlier), a total cost of 11 , or (b) schedule next operation $C$ (cost 3 ) and then solve optimally the corresponding subproblem of length 2 (cost 5 , as computed earlier), a total cost of 8 . The second possibility is optimal, and the corresponding cost of the tail subproblem is 8 , as shown next to node A in Fig. 1.3.1.
State $C$ : Here the possibilities are to (a) schedule next operation A (cost 4) and then solve optimally the corresponding subproblem of length 2 (cost 3, as computed earlier), a total cost of 7 , or (b) schedule next operation $D$ (cost 6 ) and then solve optimally the corresponding
subproblem of length 2 (cost 5, as computed earlier), a total cost of 11. The first possibility is optimal, and the corresponding cost of the tail subproblem is 7, as shown next to node A in Fig. 1.3.1.
Original Problem of Length 4: The possibilities here are (a) start with operation A (cost 5 ) and then solve optimally the corresponding subproblem of length 3 (cost 8 , as computed earlier), a total cost of 13 , or (b) start with operation C (cost 3 ) and then solve optimally the corresponding subproblem of length 3 (cost 7 , as computed earlier), a total cost of 10 . The second possibility is optimal, and the corresponding optimal cost is 10 , as shown next to the initial state node in Fig. 1.3.1.

Note that having computed the optimal cost of the original problem through the solution of all the tail subproblems, we can construct the optimal schedule by starting at the initial node and proceeding forward, each time choosing the operation that starts the optimal schedule for the corresponding tail subproblem. In this way, by inspection of the graph and the computational results of Fig. 1.3.1, we determine that CABD is the optimal schedule.

## The DP Algorithm for the Inventory Control lixample

Consider the inventory control example of the previous section. Similar to the solution of the preceding deterministic scheduling problem, we calculate sequentially the optimal costs of all the tail subproblems, going from shorter to longer problems. The only difference is that the optimal costs are computed as expected values, since the problem here is stochastic.
Ta'il Subproblems of Length 1: Assume that at the beginning of period $N$ _ 1 the stock is $X N-I$. Clearly, no matter what happened in the past the inventory manager should order the amount of inventory that minimizes over $U N-I \geq$ the sum of the ordering cost and the expected tenninal holding/shortage cost. Thus, he should minimize over UN-I the sum CUN-I $+E_{\{R(X N)\} \text {, which can be written as }}$

$$
C U N-I+\underset{w_{N-1}}{\mathrm{E}}\left\{R\left(X N-I+u_{N-1}-w_{N-1}\right)\right\} .
$$

Adding the holding/shortage cost of period $N \quad 1$, we see that the optimal cost for the last period (plus the terminal cost) is given by

$$
\begin{aligned}
J_{N-1}\left(x_{N-1}\right)= & r(x N-1) \\
& +\min _{u_{N-1} \geq 0}\left\{\text { CUN }_{-1}+\underset{w_{N-1}}{\mathrm{E}}\left\{R\left(X N-1+u_{N-1}-w_{N-1}\right)\right\}\right]
\end{aligned}
$$

Naturally, IN-/ is a function of the stock $X N-/$. It is calcula,ted either analytically or numerically (in which case a table is used for computer
storage ofthe function $I N-1$ ). In the process of caleulating $I N-I$, we obtain the optimal inventory policy $\mu_{N-1}^{*}\left(x_{N-1}\right)$ for the last period: $\mu_{N-1}^{*}\left(x_{N-1}\right)$ is the value of $u_{N-1}$ that minimizes the right-hand side of the preceding equation for a given value of $X N-1$.
Tail Subproblems of Length 2: Assume that at the beginning of period $N 2$ the stock is $x_{N-2}$. It is clear that the inventory manager should order the amount of inventory that minimizes not just the expected cost of period $N-2$ but rather the
(expected cost of period $N-2)+($ expected cost of period $N-1$, given that an optimal policy will be used at period $N_{-} 1$ ),
which is equal to

$$
r\left(x_{N-2}\right)+c u_{N-2}+E\left\{J_{N-1}\left(x_{N-1}\right)\right\}
$$

Using the system equation $x_{N-1}=X N-2+U N-2-W N-2$, the last term is also written as $I N-I(X N-2+U N-2 \quad W N-2)$.

Thus the optimal cost for the last two periods given that we are at state $x_{N-2}$, denoted IN-2(XN-2), is given by

$$
\begin{aligned}
= & T(X N-2) \\
& +\min _{u_{N-2} \geq 0}\left[c u_{N-2}+{\underset{W N-2}{ }}_{E}\left\{I N-l\left(X N-2+u_{N-2}-W N-2\right)\right\}\right]
\end{aligned}
$$

Again $J_{N-2}\left(x_{N-2}\right)$ is caleulated for every $X N-2$. At the same time, the optimal policy $\mu_{N-2}^{*}\left(x_{N-2}\right)$ is also computed.
Tail Subproblems of Length $N-k$ : Similarly, we have that at period $k$, when the stock is $x_{k}$, the inventory manager should order $U k$ to minimize
(expected cost of period $k)+\left(\right.$ expected cost of periods $k+1, \ldots, N_{-}$, given that an optimal policy will be used for these periods).

By denoting by $J k(X k)$ the optimal cost, we have

$$
\begin{equation*}
J_{k}\left(x_{k}\right)=r\left(x_{k}\right)+\min _{u_{k} \geq 0}\left[c u_{k}+\underset{w_{k}}{E}\left\{J_{k+1}\left(x_{k}+u_{k}-w_{k}\right)\right\}\right] \tag{1.4}
\end{equation*}
$$

which is actually the dynamic programming equation for this problem.
The functions $J_{k}\left(x_{k}\right)$ denote the optimal expected cost for the tail subproblem that starts at period $k$ with initial inventory $X k$. These functions are computed recursively backward in time, starting at period $N_{-} 1$ and ending at period 0 . The value $J o\left(x_{0}\right)$ is the optimal expected cost when the initial stock at time 0 is $x_{0}$. During the caleulations, the optiInal
policy is simultaneously computed from the minimization in the right-hand side of Eq. (1.4).

The example illustrates the main advantage offered by DP. While the original inventory problem requires an optimization over the set of policies, the DP algorithm of Eq. (1.4) decomposes this problem into a sequence of minimizations carried out over the set of controls. Each of these minimizations is much simpler than the original problem.

## The DP Algorithm

We now state the DP algorithm for the basic problem and show its optimality by translating into mathematical terms the heuristic argument given above for the inventory example.

Proposition 1.3.1: For every initial state $X o$, the optimal cost $J^{*}(x o)$ of the basic problem is equal to $J o(x o)$, given by the last step of the following algorithm, which proceeds backward in time from period $N-1$ to period 0 :

$$
\begin{gather*}
J_{N}\left(x_{N}\right)=g_{N}\left(x_{N}\right),  \tag{1.5}\\
J k(X k)=\min _{U K E U k(X k)} E_{W k}^{E}\left\{9 k\left(X k^{\prime \prime} U k, W k\right)+J_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\} \\
k=0,1, \ldots, N-1,
\end{gather*}
$$

(1.6)
where the expectation is taken with respect to the probability distribution of $w_{k}$, which depends on $X k$ and $U_{k}$. :Furthermore, if $u_{k}^{*}=\mu_{k}^{*}\left(x_{k}\right)$ minimizes the right side of Eq. '(1.6) for each $X k$ and $k$, the policy $\pi^{*}=\left\{\mu_{0}^{*}, \ldots, \mu_{N}^{*}-\mathrm{I}\right\}$ is optimal.

Proof: t For any admissible policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots, \mu_{N-1}\right\}$ and each $k=$ $0,1, \ldots, N-1$, denote $\pi^{k}=\left\{\mu_{k}, \mu_{k+1}, \ldots, \mu_{N-1}\right\}$. For $k \quad 0,1, \ldots, N-1$, let $J_{k}^{*}\left(x_{k}\right)$ be the optimal cost for the $(N-k)$-stage problem that starts at state $X k$ and time $k$, and ends at time $N$,
t Our proof is somewhat informal and assumes that the functions $J_{k}$ are well-defined and finite. For a strictly rigorous proof, some technical mathematical issues must be addressed; see Section 1.5. These issues do not arise if the disturbance $w_{k}$ takes a finite or countable number of values and the expected values of all terms in the expression of the cost function (1.1) are well-defined and finite for every admissible policy $\pi$.

For $k \quad I V$, we define $J_{N}^{*}\left(x_{N}\right)=g N(X N)$. We will show by induction that the functions $J_{k}^{*}$ are equal to the functions $J k$ generated by the DP algorithm, so that for $k=0$, we will obtain the desired result.

Indeed, we have by definition $J_{N}^{*}=J N=g N$. Assume that for some $k$ and all $X k+l$, we have $J_{k+1}^{*}\left(x_{k+1}\right)=J_{k+l}(X k+l)$. Then, since $\pi^{k} \quad\left(\mu_{k}, \pi^{k+1}\right)$, we have for all $x k$

$$
\begin{aligned}
& J_{k}^{*}\left(x_{k}\right)=\min _{\left(\mu_{k}, \pi^{k+1}\right) w_{k}, \ldots, w_{N-1}} E\left\{g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right. \\
& \left.+g_{N}\left(x_{N}\right)+\sum_{i=k+1}^{N-1} g_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), w_{i}\right)\right\}
\end{aligned}
$$

$$
\begin{aligned}
& \left.+\min _{\pi^{k+1}}\left[E_{w \mathrm{k}+\mathrm{I}, \ldots, \mathrm{IIJN-I}}\left\{9 N(X N)+\sum_{i=k+1}^{N-1} g i\left(X i, \mu_{i}\left(x_{i}\right),{ }^{\prime} W i\right)\right\}\right]\right\} \\
& \left.=\min _{\mu_{k}} E_{\mathrm{IUk}}\left\{9 k\left(X k^{\prime} \mu_{k}\left(x_{k}\right), w_{k}\right)+J_{k+1}^{*} U k\left(X k^{\prime} \mu_{k}\left(x_{k}\right), W k\right)\right)\right\} \\
& =\min _{\mu_{k}} \underset{w_{k}}{E}\left\{g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)+J_{k+1}\left(f_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right)\right\} \\
& =\min _{u_{k} \in U_{k}\left(x_{k}\right)} \underset{w_{k}}{E}\left\{9 k\left(X k,{ }^{\prime} U k,{ }^{\prime} W k\right)+J k+l\left(f k\left(X k, U k,{ }^{\prime} W_{k}\right)\right)\right\} \\
& =J k(X k),
\end{aligned}
$$

cOInpleting the induction. In the second equation above, we moved the minimum over $\pi^{k+1}$ inside the braced expression, using a principle of optimalityargument: "the tail portion of an optimal policy is optimal for the tail subproblem" (a more rigorous justification of this step is given in Section 1.5). In the third equation, we used the definition of $J_{k+1}^{*}$, and in the fourth equation we used the induction hypothesis. In the fifth equation, we converted the minimization over $\mu_{k}$ to a minimization over $u_{k}$, using the fact that for any function $F$ of $x$ and $u$, we have

$$
\min _{\mu \in M} F(x, \mu(x))=\min _{U E U(x)} F(x, u),
$$

where $M$ is the set of all functions $\mu(x)$ such that $\mu(x) \mathrm{E} U(x)$ for all $x$. Q.E.D.

The argument of the preceding proof provides an interpretation of $J_{k}\left(x_{k}\right)$ as the optimal cost for an ( $N-\mathrm{k}$ )-stage problem starting at state $x_{k}$ and time $k$, and ending at time $N$. We consequently call $J k(X k)$ the cost-to-go at state $X k$ and time $k$, and refer to $J k$ as the cost-to-go function at time $k$.

Ideally, we would like to use the DP algorithm to obtain closed-form expressions for $J_{k}$ or an optimal policy. In this book, we will discuss a large number of models that admit analytical solution by DP. Even if such models rely on oversimplified assumptions, they are often very useful. They may provide valuable insights about the structure of the optimal solution of more complex models, and they may form the basis for suboptimal control schemes. Furthermore, the broad collection of analytically solvable models provides helpful guidelines for modeling: when faced with a new problem it is worth trying to pattern its model after one of the principal analytically tractable models.

Unfortunately, in many practical cases an analytical solution is not possible, and one has to resort to numerical execution of the DP algorithm. This may be quite time-consuming since the minimization in the DP Eq. (1.6) must be carried out for each value of $X k$. The state space must be discretized in some way if it is not already a finite set. The computational requirements are proportional to the number of possible values of $X k$, so for complex problems the computational burden may be excessive. Nonetheless, DP is the only general approach for sequential optimization under uncertainty, and even when it is computationally prohibitive, it can serve as the basis for more practical suboptimal approaches, which will be discussed in Chapter 6.

The following examples illustrate some of the analytical and computational aspects of DP.

## Example 1.3.1

A certain material is passed through a sequence of two ovens (see Fig. 1.3.2). Denote

Xo: initial temperature of the material,
$X k, k=1,2$ : temperature of the material at the exit of oven $k$,
$u_{k-1}, k=1,2$ : prevailing temperature in oven $k$.
We assume a model of the form

$$
x_{k+1}=(1-a) x_{k}+a u_{k}, \quad k \quad 0,1,
$$

where $a$ is a known scalar from the interval $(0,1)$. The objective is to get the final temperature $x 2$ close to a given target $T$, while expending relatively little energy. This is expressed by a cost function of the form

$$
r\left(x_{2}-T\right)^{2}+u_{0}^{2}+u_{1}^{2}
$$

where $\boldsymbol{r}>{ }^{\mathrm{O}}$ is a given scalar. We assume no constraints on $u_{k}$. (In reality, there are constraints, but if we can solve the unconstrained problem and verify that the solution satisfies the constraints, everything will be fine.) The problem is deterministic; that is, there is no stochastic uncertainty. However,


Figure 1.3.2 Problem of Example 1.3.1. The temperature of the material evolves according to $X k+l=\left(\begin{array}{ll}1 & a) x k+a u_{k} \text {, where } a \text { is some scalar with }\end{array}\right.$ $O<a<1$.
such problems can be placed within the basic framework by introducing a fictitious disturbance taking a unique value with probability one.

We have $N=2$ and a terminal cost $92(X 2)=r(x 2-T) 2$, so the initial condition for the DP algorithm is [ef. Eq. (1.5)]

$$
J_{2}\left(x_{2}\right)=r\left(x_{2}-T\right)^{2}
$$

For the next-to-last stage, we have [cf. Eq. (1.6)]

$$
\begin{aligned}
J_{1}\left(x_{1}\right) & =\min _{u_{1}}\left[u_{1}^{2}+J 2(X 2)\right] \\
& =\min _{u_{1}}\left[u_{1}^{2}+J 2((1-a) x l+a U I)\right]
\end{aligned}
$$

Substituting the previous form of $J_{2}$, we obtain

$$
\begin{equation*}
J_{1}\left(x_{1}\right)=\min _{u_{1}}\left[u_{1}^{2}+r\left((1-a) x_{1}+a u_{1}-T\right)^{2}\right] \tag{1.7}
\end{equation*}
$$

This minimization will be done by setting to zero the derivative with respect to $u_{1}$. This yields

$$
0 \quad 2 n 1+2 r a((1-a) x l+a u l-T) \text {, }
$$

and by collecting terms and solving for $u_{1}$, we obtain the optimal temperature for the last oven:

$$
\mu_{1}^{*}(\mathrm{Xl})=\begin{gathered}
r a(T-(1-a) x l) \\
1+r a^{2}
\end{gathered}
$$

Note that this is not a single control but rather a control function, a rule that tells us the optimal oven temperature $u_{1}=\mu_{1}^{*}\left(x_{1}\right)$ for each possible state $X l$.

By substituting the optimal $u_{1}$ in the expression (1.7) for $J_{1}$, we obtain

$$
\begin{aligned}
J_{1}\left(x_{1}\right) & =\frac{r^{2} a^{2}\left((1-a) x_{1}-T\right)^{2}}{\left(1+r a^{2}\right)^{2}}+r\left((1-a) x_{1}+\frac{r a^{2}\left(T-(1-a) x_{1}\right)}{1+r a^{2}}-T\right)^{2} \\
& =\frac{r^{2} a^{2}\left((1-a) x_{1}-T\right)^{2}}{\left(1+r a^{2}\right)^{2}}+r\left(\frac{r a^{2}}{1+r a^{2}}-1\right)^{2}\left((1-a) x_{1}-T\right)^{2} \\
& =\frac{r\left((1-a) x_{1}-T\right)^{2}}{1+r a^{2}}
\end{aligned}
$$

We now go back one stage. We have [ef. Eq. (1.6)]

$$
J_{0}\left(x_{0}\right)=\min _{u_{0}}\left[u_{0}^{2}+J_{1}\left(x_{1}\right)\right]=\min _{u_{0}}\left[u_{0}^{2}+J_{1}\left((1-a) x_{0}+a u_{0}\right)\right]
$$

and by substituting the expression already obtained for $J_{l}$, we have

$$
J o(x o)=\operatorname{mim}_{u_{0}}\left[\begin{array}{c}
2 \\
\left.u_{0}+\begin{array}{c} 
\\
u_{0}\left((1-a)^{2} x_{0}+(1-a) a u o-T\right) 2 \\
\underline{1}+r a^{2}
\end{array}\right]
\end{array}\right.
$$

We minimize with respect to $u_{0}$ by setting the corresponding derivative to zero. We obtain

$$
O=2 u_{0}+\frac{2 r(1}{+r a^{2}}=\underline{a) a((1-a) 2 x o} \frac{\left. \pm(1=a) a u_{0}=T\right)}{1}=T
$$

This yields, after some calculation, the optimal temperature of the first oven:

$$
\mu_{0}^{*}\left(x_{0}\right)=\frac{r(l-a) a\left(T=(1-a) 2 x_{0}\right)}{1+r a^{2}\left(1+(1-a)^{2}\right)}
$$

The optimal cost is obtained by substituting this expression in the formula for Jo. This leads to a straightforward but lengthy calculation, which in the end yields the rather simple formula

$$
J_{0}\left(x_{0}\right)=\frac{r\left((1-a)^{2} x_{0}-T\right)^{2}}{1+r a^{2}\left(1+(1-a)^{2}\right)}
$$

This completes the solution of the problem.

One noteworthy feature in the preceding example is the facility with which we obtained an analytical solution. A little thought while tracing the steps of the algorithm will convince the reader that what simplifies the solution is the quadratic nature of the cost and the linearity of the system equation. In Section 4.1 we will see that, generally, when the system is linear and the cost is quadratic, the optimal policy and cost-to-go function are given by closed-form expressions, regardless of the number of sta.ges $N$.

Another noteworthy feature of the example is that the optimal policy remains unaffected when a zero-mean stochastic disturbance is added in the system equation. To see this, assume that the material's temperature evolves according to

$$
x_{k+1}=(1-a) x_{k}+a u_{k}+w_{k}, \quad k=0,1
$$

where wo, wi are independent random variables with given distribution, zero mean

$$
E\{w o\}=E\{W l\}=0,
$$

and finite variance. Then the equation for $J l$ ref. Eq. (1.6)] becomes

$$
\begin{aligned}
& J_{1}\left(x_{1}\right) \quad \min _{u_{1}} \underset{w_{1}}{E}\left\{u_{1}^{2}+r((I-a) x l+a U l+\mathrm{wI}-T) 2\right\} \\
& =\min _{u_{1}}\left[u_{1}^{2}+r((l \quad a) x l+a U l-T) 2\right. \\
& +2 r E\left\{w_{1}\right\}\left(\left(\begin{array}{ll}
1 & a) x l+a u_{1}-T
\end{array}\right)+r E\left\{w_{1}^{2}\right\}\right] .
\end{aligned}
$$

Since $E\{W l\}=0$, we obtain

$$
J_{1}\left(x_{1}\right)=\min _{u_{1}}\left[u_{1}^{2}+r\left((1-a) x_{1}+a u_{1}-T\right)^{2}\right]+r E\left\{w_{1}^{2}\right\}
$$

Comparing this equation with Eq. (1.7), we see that the presence of WI has resulted in an additional inconsequential term, $\operatorname{re}\left\{w_{1}^{2}\right\}$. Therefore, the optimal policy for the last stage remains unaffected by the presence of wi, while $J I(X I)$ is increased by the constant term $r E\left\{w_{1}^{2}\right\}$. It can be seen that a similar situation also holds for the first stage. In particular, the optimal cost is given by the same expression as before except for an additive constant that depends on $E\left\{w_{0}^{2}\right\}$ and $E\left\{w_{1}^{2}\right\}$.

If the optimal policy is unaffected when the disturbances are replaced' by their means, we say that certainty equivalence holds. We will derive certainty equivalence results for several types of problems involving a linear system and a quadratic cost (see Sections 4.1, 5.2, and 5.3).

## Hxample 1.3.2

To illustrate the computational aspects of DP, consider an inventory control problem that is slightly different from the one of Sections 1.1 and 1.2. In particular, we assume that inventory $U k$ and the demand $W k$ are nonnegative integers, and that the excess demand ( $W k-X k-U k$ ) is lost. As a result, the stock equation takes the form

$$
x_{k+1}=\max \left(0, x_{k}+u_{k}-w_{k}\right)
$$

We also assume that there is an upper bound of 2 units on the stock that can be stored, i.e. there is a constraint $x k+U k \leq 2$. The holding/storage cost for the kth period is given by

$$
\left(x_{k}+u_{k}-w_{k}\right)^{2}
$$

implying a penalty both for excess inventory and for unmet demand at the end of the kth period. The ordering cost is 1 per unit stock ordered. Thus the cost per period is

$$
g_{k}\left(x_{k}, u_{k}, w_{k}\right)=u_{k}+\left(x_{k}+u_{k}-w_{k}\right)^{2}
$$

The terminal cost is assumed to be 0 ,

$$
g_{N}\left(x_{N}\right)=0
$$

The planning horizon $N$ is 3 periods, and the initial stock xo is 0 . The demand $W k$ has the same probability distribution for all periods, given by

$$
p(W k=0)=0.1, \quad P(W k=1)=0.7, \quad p(W k=2)=0.2
$$

The system can also be represented in terms of the transition probabilities $P_{i j}(u)$ between the three possible states, for the different values of the control (see Fig. 1.3.3).

The starting equation for the DP algorithm is

$$
J_{3}\left(x_{3}\right)=0
$$

since the terminal state cost is 0 [ef. Eq. (1.5)]. The algorithm takes the form [cf. Eq. (1.6)]
$J k(X k)=\min _{\substack{0 \leq u_{k} \leq 2-x_{k} \\ u k=O, I, 2}}$ WEk$\left\{u_{k}+\left(X k+u_{k}-w_{k}\right)^{2}+J k+1\left(\max \left(\mathrm{O}, X k+u_{k}-W k\right)\right)\right\}$,
where $k=0,1,2$, and $X k, u_{k}, W k$ can take the values 0,1 , and 2 .
Period 2: We compute $J_{2(X 2)}$ for each of the three possible states. We have

$$
\begin{aligned}
J_{2}(0) & =\min _{u_{2}=0,1,2} E\left\{u_{w_{2}}+\left(u_{2}-w_{2}\right)^{2}\right\} \\
& =\min _{u_{2}=0,1,2}\left[u_{2}+0.1\left(u_{2}\right)^{2}+0.7\left(u_{2}-1\right)^{2}+0.2\left(u_{2}-2\right)^{2}\right] .
\end{aligned}
$$

We calculate the expectation of the right side for each of the three possible values of $U 2$ :

$$
\begin{aligned}
& U 2=0: E\{\cdot\}=0.7 \cdot 1+0.2 \cdot 4 \quad 1.5, \\
& U 2=1: E\{\cdot\}=1+0.1 \cdot 1+0.2 \cdot 1=1.3, \\
& u_{2}=2: E\{.\}=2+0.1 .4+0.7 .1
\end{aligned}
$$

Hence we have, by selecting the minimizing $u_{2}$,

$$
J_{2}(0)=1.3, \quad \mu_{2}^{*}(0)
$$

For $x 2=1$, we have

$$
\begin{align*}
\boldsymbol{J}_{2}(1) & \left.=\min _{\mathrm{u} 2=0,1}{\underset{w}{2}}_{E}^{\{ } \mathrm{U} 2+\left(1+u_{2}-W 2\right) 2\right\} \\
& =\min _{\mathrm{u} 2=\mathrm{O}, 1}\left[U 2+0.1\left(1+u_{2}\right)^{2}+0.7(\mathrm{U} 2) 2+-1-2\left(u_{2} \quad 1\right) 2\right]
\end{align*}
$$



Stock purchased $=0$

Stock = 20
Stock $=2$


Stock purchased $=1$
Stocke 0
Stock $=2$

Stock purchased $=2$

| Stock | Stage 0 <br> Cost-to-go | Stage 0 <br> Optimal <br> stock to <br> purchase | Stage 1 <br> Cost-to-go | Stage 1 <br> Optimal <br> stock to <br> purchase | Stage 2 <br> Cost-to-go | Stage 2 <br> Optimal <br> stock to <br> purchase |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 3.7 | 1 | 2.5 | 1 | 1.3 | 1 |
| 1 | 2.7 | 0 | 1.5 | 0 | 0.3 | 0 |
| 2 | 2.818 | 0 | 1.68 | 0 | 1.1 | 0 |

Figure 1.3.3 System and DP results for Example 1.3.2. The transition probability diagrams for the different values of stock purchased (control) are shown. The numbers next to the arcs are the transition probabilities. The control $u=1$ is not available at state 2 because of the limitation $X k-1-u_{k} \leq 2$. Similarly, the control $u=2$ is not available at states 1 and 2 . The results of the DP algorithm are given in the table.

The expected value in the right side is

$$
\begin{aligned}
& u_{2}=0: E\{\cdot\}=0.1 \cdot 1+0.2 \cdot 1=0.3 \\
& u_{2}=1: E\{\cdot\}=1+0.1 \cdot 4+0.7 \cdot 1=2.1
\end{aligned}
$$

Hence

$$
J_{2}(1)=0.3, \quad \mu_{2}^{*}(1)=0
$$

For x2 2 , the only admissible control is $u_{2}=0$, so we have

$$
\begin{gathered}
J_{2}(2)=\underset{\mathbf{1 U 2}}{E}\left\{\left(\begin{array}{ll}
(2 & w_{2}
\end{array}\right)^{2}\right\}=0.1 \cdot 4+0.7 \cdot 1=1.1 \\
J_{2}(2)=1.1, \quad \mu_{2}^{*}(2)=0
\end{gathered}
$$

Period 1: Again we compute $J_{l(X 1)}$ for each of the three possible states $\mathrm{x} 1=0,1,2$, using the values $J 2(0), J 2(1), J_{2}(2)$ obtained in the previous period. For xı 0, we have

$$
J_{1}(0)=\min _{u_{1}=0,1,2} \underset{w_{1}}{E}\left\{u_{1}+\left(u_{1}-w_{1}\right)^{2}+J_{2}\left(\max \left(0, u_{1}-w_{1}\right)\right)\right\}
$$

$u_{1}=0: E\{\cdot\}=0.1 . J_{2}(0)+0.7(1+J 2(0))+0.2(4-J 2(0)) \quad 2.8$,
$u_{1}=1: E\{\cdot\}=1+0.1\left(1+J_{2}(1)\right)+0.7 \cdot J_{2}(0)+0.2\left(1+J_{2}(0)\right)=2.5$,
$u_{1}=2: E\{\cdot\}=2+0.1(4+J 2(2))+0.7(1+J 2(1))+0.2 \cdot J_{2}(0)=3.68$,

$$
J_{l}(0)=2.5, \quad \mu_{1}^{*}(0)=1
$$

For $\mathrm{x}_{1}=1$, we have

$$
J_{1}(1)=\min _{u_{1}=0,1} \underset{w_{1}}{E}\left\{u_{1}+\left(1+u_{1}-w_{1}\right)^{2}+J_{2}\left(\max \left(0,1+u_{1}-w_{1}\right)\right)\right\}
$$

$$
u_{1}=0: E\{\cdot\}=0.1\left(1+J_{2}(1)\right)+0.7 \cdot J_{2}(0)+0.2\left(1+J_{2}(0)\right)=1.5
$$

$u_{1} \quad 1: E\{\cdot\}=1+0.1\left(4+J_{2}(2)\right)+0.7\left(1+J_{2}(1)\right)+0.2 \cdot J_{2}(0) \quad 2.68$,

$$
J_{l}(1) \quad 1.5, \quad \mu_{1}^{*}(1)=0
$$

For $\mathrm{x}_{1}=2$, the only admissible control is $u_{1}=0$, so we have

$$
\begin{aligned}
J_{l}(2) & =\underset{\boldsymbol{w}_{\mathbf{1}}}{E}\left\{\left(2-w_{1}\right) 2+J_{2}\left(\max \left(\mathrm{O}, 2 \quad w_{1}\right)\right)\right\} \\
& =0.1\left(4+J_{2}(2)\right)+0.7\left(1-J_{2}(1)\right)+0.2 \cdot J_{2}(0) \\
& =1.68
\end{aligned}
$$

$$
J_{l}(2) \quad 1.68, \quad \mu_{1}^{*}(2) \quad 0 .
$$

Period 0: Here we need to compute only $\mathrm{Jo}(\mathrm{O})$ since the initial state is known to be 0 . We have

$$
J o(O)=\min _{\mathrm{uO}=0,1,2} \underset{W_{0}}{E}\left\{11 \mathrm{O}+\left(u_{0} \quad \text { wo }\right) 2-J_{1}\left(\max \left(0, u_{0} \quad \text { WO }\right)\right)\right\}
$$

$u_{0} \quad 0: E\{\}=.0.1 . J_{1}(0)+0.7\left(1+J_{1}(0)\right)+0.2\left(4+J_{1}(0)\right)=4.0$,
$u_{0}=1: E\{\cdot\}=1+0.1\left(1+J_{1}(1)\right)+0.7 . J_{1}(0)+0.2\left(1+J_{1}(0)\right)=3.7$,
$u_{0}=2: E\{\cdot\}=2+0.1(4+J l(2))+0.7(1+J 1(1))+0.2 \cdot J_{1}(0)=4.818$,

$$
\mathrm{Jo}(\mathrm{O})=3.7, \quad \mu_{0}^{*}(0)=1
$$

If the initial state were not known a priori, we would have to compute in a similar manner $\mathrm{J}_{\mathrm{O}}(1)$ and $\mathrm{J}_{\mathrm{O}}(2)$, as well as the minimizing Uo. The reader may verify (Exercise 1.2 ) that these calculations yield

$$
\begin{array}{cc}
\mathrm{Jo}(1)=2.7, & \mu_{0}^{*}(1) \quad 0 \\
J o(2)=2.818, & \mu_{0}^{*}(2)=0
\end{array}
$$

Thus the optimal ordering policy for each period is to order one unit if the current stock is zero and order nothing otherwise. The results of the DP algorithm are given in tabular form in Fig. 1.3.3.

## Example 1.3.3 (Optimizing a Chess Match Strategy)

Consider the chess match example of Section 1.1. There, a player can select timid play (probabilities $P d$ and $1-P d$ for a draw or loss, respectively) or bold play (probabilities $p_{w}$ and $1-p_{w}$ for a win or loss, respectively) in each game of the match. We want to formulate a DP algorithm for finding the policy that maximizes the player's probability of winning the match. Note that here we are dealing with a maximization problem. We can convert the problem to a minimization problem by changing the sign of the cost function, but a simpler alternative, which we will generally adopt, is to replace the minimization in the DP algorithm with maximization.

Let us consider the general case of an N -game match, and let the state be the net score, that is, the difference between the points of the player minus the points of the opponent (so a state of 0 corresponds to an even score). The optimal cost-to-go function at the start of the kth game is given by the dynamic programming recursion

$$
\begin{align*}
J_{k}\left(x_{k}\right)=\max [ & P d J k+1(X k)+(1-P d) J k+1(X k-1) \\
& \left.P w J k+1(X k+1)+\left(1-p_{w}\right) J_{k+1}\left(x_{k}-1\right)\right] \tag{1.8}
\end{align*}
$$

The maximum above is taken over the two possible decisions:
(a) Timid play, which keeps the score at $X k$ with probability $P d$, and changes $x_{k}$ to $x_{k} \quad 1$ with probability $1 \quad P d$.
(b) Bold play, which changes $X k$ to $X k+1$ or to $X k-1$ with probabilities $p_{w}$ or $(1-P w)$, respectively.
It is optimal to play bold when
$p_{w} J_{k+1}\left(x_{k}+1\right)+\left(1-p_{w}\right) J_{k+1}\left(x_{k}-1\right) \geq p_{d} J_{k+1}\left(x_{k}\right)+\left(1-p_{d}\right) J_{k+1}\left(x_{k}-1\right)$
or equivalently, if

$$
\begin{equation*}
\frac{P w}{P d}>\frac{J k+1}{J k+1(X k+1)} \frac{(X k)}{J k+1} \frac{(X k}{J k+1(X k-1)} \tag{1.9}
\end{equation*}
$$

The dynamic programming recursion is started with

$$
J_{N}\left(x_{N}\right)= \begin{cases}1 & \text { if } X N>0  \tag{1.10}\\ p_{w} & \text { if } X N=0 \\ 0 & \text { if } X N<0\end{cases}
$$

In this equation, we have $I N(O)=p_{w}$ because when the score is even after $N$ games $(X N=0)$, it is optimal to play bold in the first game of sudden death.

By executing the DP algorithm (1.8) starting with the terminal condition (1.10), and using the criterion (1.9) for optimality of bold play, we find the following, assuming that $P d>p_{w}$ :

$$
\begin{aligned}
J_{N-1}\left(x_{N-1}\right)= & 1 \text { for } x_{N-1}>1 ; \quad \text { optirnal play: either } \\
I N-1(1)= & \max \left[p d+(1-P d) P w, p_{w}+\left(1 \quad p_{w}\right) p_{w}\right] \\
& P d+\left(1-p_{d}\right) p_{w} ; \quad \text { optimal play: timid } \\
I N-1(0)= & p_{w} ; \quad \text { optimal play: bold } \\
\mathbf{J}_{N-\mathbf{1}}(-1)= & p_{w}^{2} ; \quad \text { optimal play: bold }
\end{aligned}
$$

$$
I N-l(X N-1)=0 \text { for } X N-1<-1 ; \quad \text { optimal play: either. }
$$

Also, given $I N-l(X N-1)$, and Eqs. (1.8) and (1.9) we obtain

$$
\begin{aligned}
I N-2(O) & =\max \left[P d P W+\left(1-p_{d}\right) p_{w}^{2}, p w\left(p_{d}+(1-P d) P w\right)+\left(1-p_{w}\right) p_{w}^{2}\right] \\
& =p w\left(P w+\left(p_{w}+P d\right)(l \quad P w)\right)
\end{aligned}
$$

and that if the score is even with 2 games remaining, it is optirnal to play bold. Thus for a 2-game match, the optimal policy for both periods is to play timid if and only if the player is ahead in the score. The region of pairs $(P w, P d)$ for which the player has a better than $50-50$ chance to win a 2 -game match is

$$
R_{2}=\left\{\left(p_{w}, p_{d}\right) \mid J_{0}(0)=p_{w}\left(p_{w}+\left(p_{w}+p_{d}\right)\left(1-p_{w}\right)\right)>1 / 2\right\}
$$

and, as noted in the preceding section, it includes points where $p w<1 / 2$.

## Example 1. 3.4 (Finite-State Systems)

We mentioned earlier (d. the examples in Section 1.1) that systems with a finite number of states can be represented either in terms of a discretetime system equation or in terms of the probabilities of transition between the states. Let us work out the DP algorithm corresponding to the latter caSe. We assume for the sake of the following discussion that the problem is stationary (i.e., the transition probabilities, the cost per stage, and the control constraint sets do not change from one stage to the next). Then, if

$$
p_{i j}(u)=P\left\{x_{k+1}=j \mid x_{k}=i, u_{k}\right.
$$

are the transition probabilities, we can alternatively represent the system by the system equation (cf. the discussion of the previous section)

$$
x_{k+1}=w_{k}
$$

where the probability distribution of the disturbance $W k$ is

$$
P\left\{w_{k}=j \mid X k=\mathrm{i}, U k=U\right\}=\operatorname{Pij}(U)
$$

Using this system equation and denoting by $g(i, u)$ the expected cost per stage at state $i$ when control $U$ is applied, the DP algorithm can be rewritten as

$$
J_{k}(i)=\min _{u \in U(i)}\left[g(i, u)+E\left\{J_{k+1}\left(w_{k}\right)\right\}\right]
$$

or equivalently (in view of the distribution of $W k$ given previously)

$$
\left.J_{k}(i)=\min _{u \in U(i)}\left[9\left(i^{\prime} U\right)+\sum_{\mathrm{j}} \operatorname{Pij}(U) J k+l U\right)\right] .
$$

As an illustration, in the machine replacement example of Section 1.1, this algorithm takes the form

$$
\begin{gathered}
J_{N}(i)=0, \quad \mathrm{i}=1, \ldots, n, \\
J_{k}(i)=\min \left[R+g(1)+J_{k+1}(1), g(i)+\sum_{j=i}^{n} p_{i j} J_{k+1}(j)\right] .
\end{gathered}
$$

The two expressions in the above minimization correspond to the two available decisions (replace or not replace the machine).

In the queueing example of Section 1.1, the DP algorithm takes the form

$$
J_{N}(i)=R(i), \quad i=0,1, \ldots, n
$$

$\left.\left.. J k(i)=\min [r e i)+c f+\sum_{j=0}^{n} p_{i j}\left(u_{f}\right) . \mathrm{Jk}+\mathrm{l}(\mathrm{j}), r e i\right)+c_{s}+\sum_{j=0}^{n} p_{i j}\left(u_{s}\right) J_{k+1}(\mathrm{j})\right]$.
The two expressions in the above minimization correspond to the two possible decisions (fast and slow service).

Note that if there are $n$ states at each stage, and $U(i)$ contains as many as $m$ controls, the minimization in the right-hand side of the DP algorithm requires, for each (i, $k$ ), as many as a constant multiple of $m n$ operations. Since there are $n N$ state-time pairs, the total number of operations for the DP algorithm is as large as a constant multiple of $m n^{2} N$ operations. By contrast, the nmnber of all policies is exponential in $n N$ (it is as large as $m^{n N}$ ), so a brute force approach which enumerates all policies and compares their cost, requires an exponential number of operations in $n N$.

### 1.4 STATE AUGMENTATION AND OTHER REFORIVJ[ULATIONS

We now discuss how to deal with situations where some of the assumptions of the basic problem are violated. Generally, in such cases the problem can be reformulated into the basic problem format. This process is called state augmentation because it typically involves the enlargement of the state space. The general guideline in state augmentation is to include in the enlarged state at time $k$ all the information that is known to the controller at time $k$ and can be used with advantage in selecting $u_{k}$. Unfortunately, state augmentation often comes at a price: the reformulated problem may have very complex state and/or control spaces. We provide some examples.

Time Lags
In many applications the system state $x_{k+1}$ depends not only on the preceding state $x_{k}$ and control $u_{k}$ but also on earlier states and controls. In other words, states and controls influence future states with some time lag. Such situations can be handled by state augmentation; the state is expanded to include an appropriate number of earlier states and controls.

For simplicity, assume that there is at most a single period time lag in the state and control; that is, the system equation has the form

$$
\begin{equation*}
x_{k+1}=f_{k}\left(x_{k}, x_{k-1}, u_{k}, u_{k-1}, w_{k}\right), \quad k=1,2, \ldots, N-1 \tag{1.11}
\end{equation*}
$$

$$
x_{1}=f o\left(x o, u_{0}, w o\right)
$$

Time lags of more than one period can be handled similarly.
If we introduce additional state variables $Y k$ and $S k$, and we make the identifications $Y k=x_{k-1}, S k=u_{k-1}$, the system equation (1.11) yields

| $X k+l)$ | $f k(X k, Y k,: U k, S k, W k))$ |
| :---: | :---: |
| $Y k+l$ | $x_{k}$ |
| $S k+l$ | $\left(\begin{array}{l}k\end{array}\right.$ |

By defining $\tilde{x}_{k} \quad(X k, Y k, S k)$ as the new state, we have

$$
\tilde{x}_{k+1}=\tilde{f}_{k}\left(\tilde{x}_{k}, u_{k}, w_{k}\right),
$$

where the system function $\tilde{f}_{k}$ is defined from Eq. (1.12). By using the preceding equation as the system equation and by expressing the cost function in terms of the new state, the problem is reduced to the basic problem without time lags. Naturally, the control $u_{k}$ should now depend on the new state $\tilde{x}_{k}$, or equivalently a policy should consist of functions $\mu_{k}$ of the current state $X k$, as well as the preceding state $X k-l$ and the preceding control $u_{k-1}$.

When the DP algorithm for the reformulated problem is translated in terms of the variables of the original problem, it takes the form

$$
J_{N}\left(x_{N}\right)=g_{N}\left(x_{N}\right)
$$

$I N-I\left(x_{N-1}, X N-2, u_{N-2}\right)$

$$
\begin{aligned}
& \min _{u_{N-1}} E U N-l(X N-l) \\
& \quad \underset{w N-l}{E}\left\{9 \mathrm{~N}-\mathrm{I}\left(\mathrm{XN}-\mathrm{I}, u_{N-1}, w_{N-1}\right)\right. \\
& \left.\quad+I N\left(J N-l\left(X N-I, X N-2, u_{N-1}, U N-2,{ }^{\prime} W N-l\right)\right)\right\}
\end{aligned}
$$

$$
\begin{aligned}
& J_{k}\left(x_{k}, X k-l, u_{k-1}\right)=\min _{u_{k} \in U_{k}\left(x_{k}\right)} \underset{W k}{E}\left\{g_{k}\left(x_{k}, u_{k}, W k\right)\right. \\
& \left.\quad+\mathrm{Jk}+1\left(f k\left(X k, x_{k-1}, u_{k}, U k-I, W k\right), X k, u_{k}\right)\right\}, \quad k=1, \ldots, \mathrm{~N}-2,
\end{aligned}
$$

$$
J_{0}\left(x_{0}\right) \min _{u_{0} \in U_{0}\left(x_{0}\right)} \underset{w o}{E}\{90(X O, U o, w o)+J I(!O(X O, u o, w o), x o, u o)\} .
$$

Similar reformulations are possible when time lags appear in the cost; ${ }^{\text {' }}$ for example, in the case where the cost has the form

$$
E\left\{g_{N}\left(x_{N}, x_{N-1}\right)+g_{0}\left(x_{0}, u_{0}, w_{0}\right)+\sum_{k=1}^{N-1} g_{k}\left(x_{k}, x_{k-1}, u_{k}, w_{k}\right)\right\} .
$$

The extreme case of time lags in the cost arises in the nonadditive form

$$
E\left\{g_{N}\left(x_{N}, x_{N-1}, \ldots, x_{0}, u_{N-1}, \ldots, u_{0}, w_{N-1}, \ldots, w_{0}\right)\right\}
$$

Then, the problem can be reduced to the basic problem format, by taking as augmented state

$$
\tilde{x}_{k}=\left(x_{k}, x_{k-1}, \ldots, x_{0}, u_{k-1}, \ldots, u_{0}, w_{k-1}, \ldots, w_{0}\right)
$$

and $E\left\{g_{N}\left(\tilde{x}_{N}\right)\right\}$ as reformulated cost. Policies consist of functions $\mu_{k}$ of the present and past states $X k, \ldots, X o$, the past controls ' $U k-I, \ldots, u o$, and the past disturbances $w_{k-1}, \ldots$, Wo. Naturally, we must assume that the past disturbances are known to the controller. Otherwise, we are faced with a problem where the state is imprecisely known to the controller. Such problems are known as problems with imperfect state information and will be discussed in Chapter 5.

## Correlated Disturbances

We turn now to the case where the disturbances 'Wk are correlated over time. A common situation that can be handled efficiently by state augrnentation arises when the process ' $W 0, \ldots, W N-I$ can be represented as the output of a linear system driven by independent random variables. As an example, suppose that by using statistical methods, we determine that the evolution of $W k$ can be modeled by an equation of the form

$$
w_{k}=\lambda w_{k-1}+\xi_{k}
$$

where $\lambda$ is a given scalar and $\left\{\xi_{k}\right\}$ is a sequence of independent random vectors with given distribution. Then we can introduce an additional state variable

$$
Y k=' W k-r
$$

and obtain a new system equation

$$
\binom{x_{k+1}}{y_{k+1}}=\binom{f_{k}\left(x_{k}, u_{k}, \lambda y_{k}+\xi_{k}\right)}{\lambda y_{k}+\xi_{k}}
$$

where the new state is the pair $\tilde{x}_{k}=(X k, Y k)$ and the new disturbance is the vector $\xi_{k}$.

More generally, suppose that 'Wk can be modeled by

$$
w_{k}=C_{k} y_{k+1}
$$

where

$$
y_{k+1}=A_{k} y_{k}+\xi_{k}, \quad k \quad 0, \ldots, \mathbf{N - 1}
$$

$A_{k}, O k$ are known matrices of appropriate dimension, and $\xi_{k}$ are independent random vectors with given distribution (see Fig. 1.4.1). By viewing $Y k$ as an additional state variable, we obtain the new system equation

$$
\binom{x_{k+1}}{y_{k+1}}=\binom{f_{k}\left(x_{k}, u_{k}, C_{k}\left(A_{k} y_{k}+\xi_{k}\right)\right)}{A_{k} y_{k}+\xi_{k}}
$$



Figure 1.4.1 Representing correlated disturbances as the output of a linear system driven by independent random vectors.

Note that in order to have perfect state information, the controller must be able to observe $Y k$. Unfortunately, this is true only in the minority of practical cases; for example when $C_{k}$ is the identity matrix and $w k-l$ is observed before $u_{k}$ is applied. In the case of perfect state information, the DP algorithm takes the form

$$
\begin{gathered}
I N(X N, Y N)=g_{N}\left(x_{N}\right) \\
J_{k}\left(x_{k}, y_{k}\right)=\min _{u_{k} \in U_{k}\left(x_{k}\right)} \underset{\xi_{k}}{E}\left\{g_{k}\left(x_{k}, u_{k}, O_{k}\left(A_{k} y_{k}+\xi_{k}\right)\right)\right. \\
\left.+J k+l\left(f k\left(X k^{\prime} u_{k}, C k\left(A k Y k+\xi_{k}\right)\right), A k Y k+\xi_{k}\right)\right\}
\end{gathered}
$$

## Forecasts

Finally, consider the case where at time $k$ the controller has access to a forecast $Y k$ that results in a reassessment of the probability distribution of $W k$ and possibly of future disturbances. For example, $Y k$ may be an exact prediction of 'Wk or an exact prediction that the probability distribution of 'Wk is a specific one out of a finite collection of distributions. Forecasts of interest in practice are, for example, probabilistic predictions on the state of the weather, the interest rate for money, and the demand for inventory.

Generally, forecasts can be handled by state augmentation although the reformulation into the basic problem format may be quite complex. We will treat here only a simple special case.

Assume that at the beginning of each period $k$, the controller receives an accurate prediction that the next disturbance ' $W k$ will be selected according to a particular probability distribution out of a given collection of distributions $\left\{Q_{1}, \ldots, Q m\right\}$; that is, if the forecast is $i$, then $W k$ is selected according to $Q, \therefore$. The a priori probability that the forecast will be $i$ is denoted by $P i$ and is given.

For instance, suppose that in our earlier inventory example the demand $w_{k}$ is determined according to one of three distributions $Q l, Q 2$, and $Q_{3}$, corresponding to "srnall," "medium," and "large" demand. Each of the three types of demand occurs with a given probability at each time period, independently of the values of demand at previous time periods. However, the inventory manager, prior to ordering $u_{k}$, gets to know through a forecast the type of demand that will occur. (Note that it is the probability distribution of demand that becomes known through the forecast, not the demand itself.)

The forecasting process can be represented by means of the equation

$$
Y k+l=\xi_{k},
$$

where $y_{k+1}$ can take the values $1, \ldots, ' m$, corresponding to the $m$ possible forecasts, and $\xi_{k}$ is a random variable taking the value $i$ with probability
p'i. The interpretation here is that when $\xi_{k}$ takes the value $i$, then 'Wk-l-1 will occur according to the distribution Qi.

By combining the system equation with the forecast equation $Y k+l=$ $\xi_{k}$, we obtain an augmented system given by

$$
\binom{x_{k+1}}{y k+\mathbb{1}}=\binom{f_{k}\left(x_{k}, u_{k}, w_{k}\right)}{\xi_{k}}
$$

The new state is

$$
\tilde{x}_{k}=\left(X k^{\prime} Y k\right)^{\prime}
$$

and because the forecast $Y k$ is known at time $k$, perfect state information prevails. The new disturbance is

$$
\tilde{w}_{k}=\left(' W k, \xi_{k}\right)
$$

and its probability distribution is determined by the distributions $Q i$ and the probabilities $P i$, and depends explicitly on $\tilde{x}_{k}$ (via $Y k$ ) but not on the prior disturbances.

Thus, by suitable reformulation of the cost, the problem can be cast into the basic problem format. Note that the control applied depends on both the current state and the current forecast. The DP algorithm takes the form

$$
\begin{align*}
& I N(X N, Y N)=9 N(X N) \\
& J_{k}\left(x_{k}, y_{k}\right) \quad \min _{u_{k} \in U_{k}\left(x_{k}\right)} \underset{W k}{E}\left\{g_{k}\left(x_{k}, u_{k}, W k\right)\right. \\
& \\
& \left.\quad+\sum_{i=1}^{m} P i J k+1\left(f_{k}\left(x_{k}, u_{k}, W k\right), \text { i }\right) \mid Y k\right\} \tag{1.13}
\end{align*}
$$

where $Y k$ may take the values $1, \ldots$, 'm, and the expectation over $w_{k}$ is taken with respect to the distribution $Q Y k^{\prime}$

It should be clear that the preceding formulation admits several extensions. One example is the case where forecasts can be influenced by the control action and involve several future disturbances. However, the price for these extensions is increased complexity of the corresponding DP algorithm.

## Simplification for Uncontrollable State Uomponents

When augmenting the state of a given system one often ends up with composite states, consisting of several components. It turns out that if some of these components cannot be affected by the choice of control, the DP algorithm can be simplified considerably, as we will now describe.

Let the state of the system be a composite $(X k, J / k)$ of two components $X k$ and $Y k$. The evolution of the main component, $X k$, is affected by the control $u_{k}$ according to the equation
where the probability distribution $P k\left(W k \mid X k, Y k, u_{k}\right)$ is given. The evolution of the other component, $Y k$, is governed by a given conditional distribution $P k(Y k \mid X k)$ and cannot be affected by the control, except indirectly through $X k$. One is tempted to view $Y k$ as a disturbance, but there is a difference: $Y k$ is observed by the controller before applying $U k$, while ' $W k$ occurs after $u_{k}$ is applied, and indeed $W k$ may probabilistically depend on $u_{k}$.

We will formulate a DP algorithm that is executed over the controllable component of the state, with the dependence on the uncontrollable component being "averaged out." In particular, let $J k(X k, Y k)$ denote the optimal cost-to-go at stage $k$ and state $(X k, Y k)$, and define

$$
\hat{J}_{k}\left(x_{k}\right)=\underset{y_{k}}{E}\left\{J_{k}\left(x_{k}, y_{k}\right) \mid x_{k}\right\}
$$

We will derive a DP algorithm that generates $\mathbf{J}_{k(X k)}$.
Indeed, we have

$$
\begin{aligned}
& \hat{J}_{k}\left(x_{k}\right)=E_{Y k}\{J k(X k, Y k) \mid X k\} \\
& =E_{Y k}\left\{\min _{u_{k} \in U_{k}\left(x_{k}, y_{k}\right)} E W_{k}, X k+l, Y k+l\{9 k(X k, Y k, U k, W k)\right. \\
& +\left.\left.\mathcal{J}_{k+1}\left(x_{k+1}, y_{k+1}\right)\right|_{\left.X k^{\prime} Y k, U k\right\}}\right|_{X k\}} \\
& E_{Y k\{ } \min _{u_{k} \in U_{k}\left(x_{k}, y_{k}\right)} E_{w_{k}, x_{k+1}\{ }\left\{9 k\left(X k^{\prime} V k, U k, W k\right)\right. \\
& \left.+\left.E_{y_{k+1}}\left\{J_{k+1}\left(x_{k+1}, Y k+1\right) \mid x_{k+1}\right\}\right|_{X k, Y k,} u_{k}\right\}\left.\right|_{x k\}},
\end{aligned}
$$

and finally

$$
\begin{align*}
& \hat{J}_{k}\left(x_{k}\right)=\underset{y_{k}}{E}\left\{\operatorname { m i n } _ { u _ { k } \in U _ { k } ( x _ { k } , y _ { k } ) } { \underset { W k } { E } } _ { E } \left\{9 k\left(X k^{\prime} V k, ' U k, W k\right)\right.\right. \\
&\left.\left.+J_{k+1}(f k(X k, Y k, U k, W k))\right\} 1 x_{k}\right\} \tag{1.14}
\end{align*}
$$

The advantage of this equivalent DP algorithm is that it is executed over a significantly reduced state space. For example, if $X k$ takes $n$ possible values and $Y k$ takes $m$ possible values, then DP is executed over $n$ states instead of $n m$ states. Note, however, that the minimization in the righthand side of the preceding equation yields an optimal control law as a function of the full state ( $X k^{\prime} Y k$ ).

As an example, consider the augmented state resulting from the incorporation of forecasts, as described earlier. Then, the forecast $y_{k}$ represents
an uncontrolled state component, so that the DP algorithm can be simplified as in Eq. (1.14). In particular, by defining

$$
\mathbf{J}_{k}\left(x_{k}\right)=\sum_{\mathrm{i}=1}^{m} p_{i} J_{k}\left(x_{k}, ' i\right), \quad k=0,1, \ldots, N-1,
$$

and

$$
\hat{J}_{N}\left(x_{N}\right)=g_{N}\left(x_{N}\right)
$$

we have, using Eq. (1.13),

$$
\begin{aligned}
\left.\mathbf{J} k(x k)=\sum_{i=1}^{m} P i \min _{u_{k} \in U_{k}\left(x_{k}\right)}\right)_{w k}^{E}\{ & 9 k(x k, U k, W k) \\
& \left.+J_{k+1}(f k(X k, U k, W k)) \mid \cup k=\dot{\imath}\right\}
\end{aligned}
$$

which is executed over the space of $X k$ rather than $X k$ and ' $U k$.
Uncontrolled state components often occur in arrival systems, such as queueing, where action must be taken in response to a random event (such as a customer arrival) that cannot be influenced by the choice of control. Then the state of the arrival system must be augmented to include the random event, but the DP algorithm can be executed over a smaller space, as per Eq. (1.14). Here is another example of similar type.

## Exanlple 1.4.1: (Tetris)

Tetris is a popular video game played on a two-dimensional grid. Each square in the grid can be full or empty, making up a "wall of bricks" with "holes" and a "jagged top". The squares fill up as blocks of different shapes fall from the top of the grid and are added to the top of the wall. As a given block falls, the player can move horizontally and rotate the block in all possible ways, subject to the constraints imposed by the sides of the grid and the top of the wall. The falling blocks are generated independently according to some probability distribution, defined over a finite set of standard shapes. The game starts with an empty grid and ends when a square in the top row becomes full and the top of the wall reaches the top of the grid. When a row of full squares is created, this row is removed, the bricks lying above this row move one row downward, and the player scores a point. The player's objective is to maximize the score attained (total number of rows removed) within $N$ steps or up to termination of the game, whichever occurs first.

We can model the problem of finding an optimal tetris playing strategy as a stochastic DP problem. The control, denoted by $u$, is the horizontal positioning and rotation applied to the falling block. The state consists of two components:
(1) The board position, I.e., a binary description of the full/empty status of each square, denoted by $x$.
(2) The shape of the current falling block, denoted by $y$.

There is also an additional termination state which is cost-free. Once the state reaches the termination state, it stays there with no change in cost.

The shape $y$ is generated according to a probability distribution $p(y)$, independently of the control, so it can be viewed as an uncontrollable state component. The DP algorithm (1.14) is executed over the space of $x$ and has the intuitive form

$$
\mathrm{j}_{k}(x) \quad \sum_{v} p(y) \max _{u}\left[g(x, y, u)+\mathrm{J}_{\mathrm{k}+1}(f(x, y, u)) \mathrm{J}, \quad \text { for all } x,\right.
$$

where $g(x, y, u)$ and $f(x, y, u)$ are the number of points scored (rows removed), and the board position (or termination state) when the state is $(x, y)$ and control $u$ is applied, respectively. Note, however, that despite the simplification in the DP algorithm achieved by eliminating the uncontrollable portion of the state, the number of states $x$ is enormous, and the problem can only be addressed by suboptimal methods, which will be discussed in Chapter 6 and in Vol. II.

### 1.5 SOME $1 \backslash!1 A T H E I \backslash 1 A T I C A L ~ I S S U E S$

Let us now discuss some technical issues relating to the basic problem formulation and the validity of the DP algorithm. The reader who is not mathematically inclined need not be concerned about these issues and can skip this section without loss of continuity.

Once an admissible policy $\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}$ is adopted, the following sequence of events is envisioned at the typical stage $k$ :

1. The controller observes $X k$ and applies $u_{k}=j l^{\prime} k(X k)^{\prime}$
2. The disturbance $W k$ is generated according to the given distribution $P_{k}(\cdot \mid X k, 11 k(X k))^{\prime}$
3. The cost $g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)$ is incurred and added to previous costs.
4. The next state $x_{k+1}$ is generated according to the system equation

$$
x_{k+1}=f_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)
$$

If this is the last stage $(k=N-1)$, the terminal cost $9 \mathrm{~N}(x \mathrm{~N})$ is added to previous costs and the process terminates. Otherwise, $k$ is incremented, and the same sequence of events is repeated at the next stage.
For each stage, the above process is well-defined and is couched in precise probabilistic terms. Matters are, however, complicated by the need to
view the cost as a well-defined random variable with well-defined expected value. The framework of probability theory requires that for each policy we define an underlying probability space, that is, a set $\Omega$, a collection of events in $\Omega$, and a probability measure on these events. In addition, the cost must be a well-defined random variable on this space in the sense of Appendix C (a measurable function from the probability space into the real line in the terminology of measure-theoretic probability theory). For this to be true, additional (measurability) assumptions on the functions $f k, 9 k$, and $\mu_{k}$ may be required, and it may be necessary to introduce a.clditional structure on the spaces $S k, O k$, and $D_{k}$. Furthermore, these assumptions may restrict the class of admissible policies, since the functions $\mu_{k}$ may be constrained to satisfy additional (measurability) requiremen'ts.

Thus, unless these additional assumptions and structure are specified, the basic problem is formulated inadequately from a mathematical point of view. Unfortunately, a rigorous formulation for general state, control, and disturbance spaces is well beyond the mathematical framework of this introductory book and will not be undertaken here. Nonetheless, it turns out that these difficulties are mainly technical and do not substantially affect the basic results to be obtained. For this reason, we find it convenient to proceed with informal derivations and arguments; this is consistent with rnost of the literature on the subject.

We would like to stress, however, that under at least one frequently satisfied assumption, the mathematical difficulties mentioned above disappear. In particular, let us assume that the disturbance spaces $D_{k}$ are all countable and the expected values of all terms in the cost are finite for every admissible policy (this is true in particular if the spaces $D_{k}$ are finite sets). Then, for every admissible policy, the expected values of all the cost terms can be written as (possibly infinite)sums involving the probabilities of the elements of the spaces $D k$.

Alternatively, one may write the cost as

$$
\begin{equation*}
J_{\pi}\left(x_{0}\right)=\underset{x_{1}, \ldots, x_{N}}{E}\left\{g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} \tilde{g}_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right)\right)\right\} \tag{1.15}
\end{equation*}
$$

where

$$
\tilde{g}_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right)\right)=\underset{w_{k}}{E}\left\{g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right) \mid x_{k}, \mu_{k}\left(x_{k}\right)\right\},
$$

with the preceding expectation taken with respect to the distribution $\mathrm{P}_{\mathrm{k}}$ (. I $\left.X k, \mu_{k}\left(x_{k}\right)\right)$ defined on the countable set $D k$. Then one may take as the basic probability space the Cartesian product of the spaces $\tilde{S}_{k}, k \quad 1, \ldots, \mathrm{~N}$, given for all $k$ by

$$
\tilde{S}_{k+1}=\left\{x_{k+1} \in S_{k+1} \mid x_{k+1}=f_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right), x_{k} \in \quad, w_{k} \in D_{k}\right\}
$$

where $\tilde{S}_{0}=\{x o\}$. The set $\tilde{S}_{k}$ is the subset of all states that can be reached at tirne $k$ when the policy $\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}$ is used. Because the disturbance spaces $D k$ are countable, the sets $\tilde{S}_{k}$ are also countable (this is true since the union of any countable collection of countable sets is a countable set). The system equation $X k+l=f k\left(X k^{\prime} \mu_{k}\left(x_{k}\right), w_{k}\right)$, the probability distributions $P_{k}\left(\cdot \mid X k, \mu_{k}\left(x_{k}\right)\right)$, the initial state $X o$, and the policy $\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}$ define a probability distribution on the countable set $\tilde{S}_{1} \times \ldots, \times \tilde{S}_{N}$, and the expected value in the cost expression (1.15) is defined with respect to this latter distribution

Let us now give a more detailed proof of the validity of the DP algorithm (Prop. 1.3.1). We assume that the disturbance $W k$ takes a finite or countable number of values and the expected values of all terms in the expression of the cost function are finite for every admissible policy 1r. Furthermore, the functions $J k(X k)$ generated by the DP algorithm are finite for all states $X k$ and times $k$. We do not need to assume that the minimum over $u_{k}$ in the definition of $J k(X k)$ is attained by some $u_{k} \mathrm{E} U(X k)^{\prime}$

For any admissible policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots, \mu_{N-1}\right\}$ and each $k=$ $0,1, \ldots, N \quad 1$, denote $\pi^{k}=\left\{\mu_{k}, \mu_{k+1}, \ldots, \mu_{N-1}\right\}$. For $k=0,1, \ldots, N \quad 1$, let $J_{k}^{*}(x k)$ be the optimal cost for the $(N-k)$-stage problem that starts at state $1: \mathrm{k}$ and time $k$, and ends at time $N$; that is,

$$
J_{k}^{*}\left(x_{k}\right)=\min _{\pi^{k}} E\left\{g_{N}\left(x_{N}\right)+\sum_{i=k}^{N-1} g_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), w_{i}\right)\right\} .
$$

For $k=N$, we define $J_{N}^{*}\left(x_{N}\right)=9 N(X N)$. We will show by induction that the functions $J_{k}^{*}$ are equal to the functions $J k$ generated by the DP algorithm, so that for $k=0$, we will obtain the desired result.

For any $\epsilon>0$, and for all $k$ and $X k$, let $\mu_{k}^{\epsilon}\left(x_{k}\right)$ attain the minimum in the equation

$$
\begin{equation*}
J_{k}\left(x_{k}\right)=\min _{u_{k} \in U_{k}\left(x_{k}\right)} E\left\{g k\left(X k, u_{k}, W k\right)+J k+l\left(f k\left(x k, u_{k}, W k\right)\right)\right\} \tag{1.16}
\end{equation*}
$$

$$
k \quad O, I, \ldots, N-1
$$

within E; that is, for all $X k$ and $k$, we have $\mu_{k}^{\epsilon}\left(x_{k}\right) \mathrm{E} U k(X k)$ and

$$
\begin{equation*}
\underset{W k}{E}\left\{g_{k}\left(x_{k}, \mu_{k}^{\epsilon}\left(x_{k}\right), w_{k}\right)+f_{k+1}\left(f_{k}\left(x_{k}, \mu_{k}^{\epsilon}\left(x_{k}\right), w_{k}\right)\right)\right\} \leq J k(X k)+\mathrm{E} \tag{1.17}
\end{equation*}
$$

Let $J_{k}^{\epsilon}\left(x_{k}\right)$ be the expected cost starting at state $X k$ at time $k$, and using the policy $\left\{\mu_{k}^{\epsilon}, \mu_{k+1}^{\epsilon}, \ldots, \mu_{N-1}^{\epsilon}\right\}$. We will show that for all $X k$ and $k$, we have

$$
\begin{gather*}
J k(X k) \leq J_{k}^{\epsilon}\left(x_{k}\right) \leq J k(X k)+(N \quad \mathrm{k}) \mathrm{E}  \tag{1.18}\\
J_{k}^{*}(X k) \leq J_{k}^{\epsilon}\left(x_{k}\right) \leq J_{k}^{*}(X k)+(N-k) \mathrm{E}  \tag{1.19}\\
J k(X k)=J k(X k) \tag{1.20}
\end{gather*}
$$

It is seen using Eq. (1.17) that the inequalities (1.18) and (1.19) hold for $k=N-1$. By taking $\epsilon \rightarrow 0$ in Eqs. (1.18) and (1.19), it is also seen that $J_{N-1}=J_{N-1}^{*}$. Assume that Eqs. (1.18)-(1.20) hold for index $k+1$. We will show that they also hold for index $k$.

Indeed, we have

$$
\begin{aligned}
& J_{k}^{\epsilon}\left(x_{k}\right)=\underset{W k}{E}\left\{g k\left(X k^{\prime} \mu_{k}^{\epsilon}\left(x_{k}\right), W k\right)+J_{k+1}^{\epsilon}\left(f k\left(x_{k}, \mu_{k}^{\epsilon}\left(x_{k}\right), W k\right)\right)\right\} \\
& \left.\quad \leq{\underset{W W}{ }}_{E}^{\{g k}\left(X k^{\prime} \mu_{k}^{\epsilon}\left(x_{k}\right), w_{k}\right)+J k+1\left(f k\left(x_{k}, \mu_{k}^{\epsilon}\left(x_{k}\right), W k\right)\right)\right\}+(N \quad k-\mathrm{I}) \mathrm{E} \\
& \quad \leq J_{k}\left(x_{k}\right)+\epsilon+(N-k \quad 1) \mathrm{E} \\
& \quad=J_{k}\left(x_{k}\right)+(N-\mathrm{k}) \mathrm{E}
\end{aligned}
$$

where the first equation holds by the definition of $J_{k}^{\epsilon}$, the first inequality holds by the induction hypothesis, and the second inequality holds Eq. (1.17). We also have

$$
\begin{aligned}
J_{k}^{\epsilon}\left(x_{k}\right) & ={\underset{W k}{E}\left\{g k\left(X k^{\prime} \mu_{k}^{\epsilon}\left(x_{k}\right), W k\right)+J_{k+1}^{\epsilon}\left(f k\left(X k^{\prime} \mu_{k}^{\epsilon}\left(x_{k}\right), W k\right)\right)\right\}} \geq{\underset{W}{W k}}_{E}\left\{g_{k}\left(x_{k}, \mu_{k}^{\epsilon}\left(x_{k}\right), w_{k}\right)+J_{k+1}\left(f_{k}\left(x_{k}, \mu_{k}^{\epsilon}\left(x_{k}\right), w_{k}\right)\right)\right\} \\
& \geq \min _{u_{k} \in U\left(x_{k}\right)} \underset{w_{k}}{\mathbb{E}}\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+J_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\} \\
& =J_{k}\left(x_{k}\right)
\end{aligned}
$$

where the first inequality holds by the induction hypothesis. Combining the preceding two relations, we see that Eq. (1.18) holds for index $k$.

For every policy $\pi \quad\left\{\mu_{0}, \mu_{1}, \therefore, \mu_{N-1}\right\}$, we have

$$
\begin{aligned}
& J_{k}^{\epsilon}\left(x_{k}\right)=\underset{W k}{E}\left\{g k(X k, P k(X k), W k)+J_{k+1}^{\epsilon}\left(f_{k}\left(x_{k}, \mu_{k}^{\epsilon}\left(x_{k}\right), w_{k}\right)\right)\right\} \\
& \leq \underset{w k}{E}\left\{g k\left(X k^{\prime} \mu_{k}^{\epsilon}\left(x_{k}\right), W k\right)+-J k+{ }_{1}\left(f_{k}\left(X k^{\prime} \mu_{k}^{\epsilon}\left(x_{k}\right), w_{k}\right)\right)\right\}+(N-k-\mathrm{I}) \mathrm{E} \\
& \leq \min _{u_{k} \in U\left(x_{k}\right)} \underset{W k}{E}\left\{9 k\left(X k^{\prime} U k, W k\right)+-J k+l\left(f k\left(X k^{\prime} U k, W k\right)\right)\right\}+(N-k) \epsilon \\
& \leq \min _{u_{k} \in U\left(x_{k}\right) W k} \underset{W}{E}\left\{g k\left(X k^{\prime} u_{k}, w_{k}\right)+J_{k+1}\left(f_{k}\left(x_{k}, u_{k}, W k\right)\right)\right\}+(N-k) \epsilon \\
& \leq \underset{w k}{E}\left\{g k\left(X k^{\prime} l l k(X k), W k\right)+J 7 f k+l\left(f k\left(X k^{\prime} \mu_{k}\left(x_{k}\right), W k\right)\right)\right\}+(N-k) \epsilon \\
& =J_{\pi^{k}}\left(x_{k}\right)+(N-k) \epsilon,
\end{aligned}
$$

where the first inequality holds by the induction hypothesis, and the second inequality holds by Eq. (1.17). Taking the minimum over $\pi^{k}$ in the preceding relation, we obtain for all $X k$

$$
J_{k}^{\epsilon}\left(x_{k}\right) \leq J_{k}^{*}\left(x_{k}\right)+(N-k) \epsilon
$$

We also have by the definition of $J_{k}^{*}$, for all $X k$,

$$
J_{k}^{*}\left(x_{k}\right) \leq J k(X k)
$$

Combining the preceding two relations, we see that Eq. (1.19) holds for index $k$. Finally, Eq. (1.20) follows from Eqs. (1.18) and (1.19), by taking $\epsilon \rightarrow 0$, and the induction is complete.

Note that by using $\epsilon=0$ in the relation

$$
J_{0}^{\epsilon}\left(x_{k}\right) \leq J o(X k)+N \epsilon
$$

[ef. Eq. (1.19)], we see that a policy that attains the minimum for all $x_{k}$ and $k$ in Eq. (1.16) is optimal.

In conclusion, the basic problem has been formulated rigorously, and the DP algorithm has been proved rigorously only when the disturbance spaces Do, $\ldots, D_{N-1}$ are countable sets, and the expected values of all the cost expressions associated with the problem and the DP algorithm are finite. In the absence of these assumptions, the reader should interpret subsequent results and conclusions as essentially correct but mathematically imprecise statements. In fact, when discussing infinite horizon problems (where the need for precision is greater), we will make the countability assumption explicit.

We note, however, that the advanced reader will have little difficulty in establishing most of our subsequent results concerning specific finite horizon applications, even if the countability assumption is not satisfied. This can be done by using the DP algorithm as a verification theorem. In particular, if one can find within a subset of policies III (such as those satisfying certain measurability restrictions) a policy that attains the minimum in the DP algorithm, then this policy can be readily shown to be optImal within $\tilde{\Pi}$. This result is developed in Exercise 1.12, and can be used by the mathematically oriented reader to establish rigorously many of our subsequent results concerning specific applications. For example, in linearquadratic problems (Section 4.1) one determines from the DP algorithm a poliey in dosed form, which is linear in the current state. When $w_{k}$ can take uncountably many values, it is necessary that admissible policies consist of Borel measurable functions $\mu_{k}$. Since the linear policy obtained from the DP algorithm belongs to this class, the result of Exercise 1.12 guarantees that this policy is optimal. For a rigorous mathematical treatment of DP that resolves the associated measurability issues and supplements the present text, see the book by Bertsekas and Shreve [BeS78].

### 1.6 DYNAIVIIC PROGRAIVIMING AND MINIMAX CONTROL

The problem of optimal control of uncertain systems has traditionally been treated in a stochastic frarnework, whereby all uncertain quantities are described by probability distributions, and the expected value of the cost is
minimized. However, in many practical situations a stochastic description of the uncertainty may not be available, and one may have information with less detailed structure, such as bounds on the magnitude of the uncertain quantities. In other words, one may know a set within which the uncertain quantities are known to lie, but may not know the corresponding probability distribution. Under these circumstances one may use a minimax approach, whereby the worst possible values of the uncertain quantities within the given set are assumed to occur.

The minimax approach for decision making under uncertainty is described in Appendix G and is contrasted with the expected cost approach, which we have been following so far. In its simplest form, the corresponding decision problem is described by a triplet (II, $W, J$ ), where II is the set of policies under consideration, $W$ is the set in which the uncertain quantities are known to belong, and $J$ : II x VII $\mapsto[-00,+00]$ is a given cost function. The objective is to

$$
\operatorname{minimize} \max _{w E W} J(\pi, w)
$$

over all $\pi \mathrm{E}$ II.
It is possible to formulate a minimax counterpart to the basic problem with perfect state information. This problem is a special case of the abstract minimax problem above, as discussed more fully in Appendix G. Generally, it is unusual for even the simplest special cases of this problem to admit a closed-form solution. However, a computational solution using DP is possible, and our purpose in this section is to describe the corresponding algorithm.

In the framework of the basic problem, consider the case where the disturbances $w_{0}, w_{1}, \ldots, w_{N-1}$ do not have a probabilistic description but rather are known to belong to corresiponding given sets $W k\left(x_{k}, u_{k}\right) \subset D k$, $k=0,1, \ldots, \mathbf{N}-1$, which may depend on the current state $x_{k}$ and control $u_{k}$. Consider the problem of finding a policy $\pi \quad\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}$ with $\mu_{k}\left(x_{k}\right) \mathrm{E} U k(X k)$ for all $X k$ and $k$, which minimizes the cost function

$$
J_{\pi}\left(x_{0}\right)=\max _{\substack{w_{k} \in W_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right)\right) \\ k=0,1, \ldots, N-1}}\left[g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right] .
$$

The DP algorithm for this problem takes the following form, which resembles the one corresponding to the stochastic basic problem (rnaximization is used in place of expectation):

$$
\begin{equation*}
J_{N}\left(x_{N}\right)=g_{N}\left(x_{N}\right) \tag{1.21}
\end{equation*}
$$

$$
\begin{equation*}
J_{k}\left(x_{k}\right)=\min _{u_{k} \in U\left(x_{k}\right) w_{k} \in W_{k}\left(x_{k}, u_{k}\right)} \max \left[g_{k}\left(x_{k}, u_{k}, w_{k}\right)+J_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right] \tag{1.22}
\end{equation*}
$$

This algorithm can be explained by using a principle of optimality type of argurnent. In particular, we consider the tail subproblem whereby we are at state $x_{k}$ at time $k$, and we wish to minimize the "cost-to-go"

$$
\max _{\substack{w_{i} \in w_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right)\right) \\ i=k, k+1, \ldots, N-1}}\left[g_{N}\left(x_{N}\right)+\sum_{i=k}^{N-1} g_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), w_{i}\right)\right]
$$

and we argue that if $\pi^{*}=\left\{\mu_{0}^{*}, \mu_{1}^{*}, \ldots, \mu_{N}^{*}-\mathrm{I}\right\}$ is an optimal policy for the rninimax problem, then the truncated policy $\left\{\mu_{k}^{*}, \mu_{k+1}^{*}, \ldots, \mu_{N}^{*}-I\right\}$ is optimal for the tail subproblem. The optimal cost of this subproblem is $J_{k}\left(x_{k}\right)$, as given by the DP algorithm (1.21)-(1.22). The algorithm expresses the intuitively clear fact that when at state $x k$ at time $k$, then regardless of what happened in the past, we should choose $U k$ that minimizes the worst/maximum value over $w_{k}$ of the sum of the current stage cost plus the optimal cost of the tail subproblem that starts from the next state.

We will now give a mathematical proof that the DP algorithm (1.21)(1.22) is valid, and that the optimal cost is equal to $J o(x o)$. For this it is necessary to assume that $J k(X k)>-\infty$ for all $X k$ and $k$. This is analogous to the assumption we made in the preceding section for the validity of the DP algorithm under stochastic disturbances, i.e., that the values $J k(X k)$ generated by the DP algorithm are finite for all states $X k$ and stages $k$. The following lemma provides the key argument.

Lemma 1.6.1: Let $f: W \rightarrow X$ be a function, and $M$ be the set of all functions $\mu: \mathrm{X} \rightarrow U$, where $W, \mathrm{X}$, and $U$ are some sets. Then for any functions $G o: W \rightarrow(-00,00]$ and $G l: \mathrm{X} \times U \rightarrow(-00,00]$ such that

$$
\min _{u E U} \mathrm{G}_{\mathrm{l}}(f(w), u)>-\infty, \quad \text { for all } w \mathrm{E} W
$$

we have
$\min _{\mu \in M} \max _{w \in W}[\mathrm{GO}(\mathrm{W})+\mathrm{Gl}(f(w), \mu(f(w)))]=\max _{w \in W}\left[\mathrm{Go}(\mathrm{w})+\operatorname{rning}_{u \in U} \mathrm{Gl}(f(w), \mathrm{U})\right]$.

Proof: We have for all $\mu \mathrm{E}$ I!!

$$
\max _{w E W}[G o(w)+G 1(f(w), \mu(f(w)))] \geq \max _{w \in W}\left[G o(w)+\min _{u E U} G 1(f(w), u)\right]
$$

and by taking the minimum over $\mu \mathrm{E} M$, we obtain
$\min _{\mu \in M} \max _{w E W}\left[G_{0}(w)+G_{1}(f(w), \mu(f(w)))\right] \geq \max _{w \in W}\left[G o(w)+\min _{u E U} 01(f(w), u)\right]$.

To show the reverse inequality, for any $\epsilon>0$, let $\mu_{\epsilon} \mathrm{E}$ VIV be such that

$$
G_{1}\left(f(w), \mu_{\epsilon}(f(w))\right) \leq \min _{\mathfrak{u} U} G_{1}(f(w), u)+\mathrm{E}, \quad \text { for all } w \mathrm{E} W
$$

[Such a $\mu_{\epsilon}$ exists because of the assumption minuEu $O_{l}(f(w), u)>-00$.] Then

$$
\begin{aligned}
\min _{\mu \in M} \max _{w \in W}[O o(w) & +G 1(f(w), \mu(f(w)))] \\
& \leq \max _{w E W}\left[O o(w)+01\left(f(w), \mu_{\epsilon}(f(w))\right)\right] \\
& \leq \max _{w E W}\left[G o(w)+\min _{u E U} 01(f(w), u)\right]+\text { Е. }
\end{aligned}
$$

Since $\epsilon>0$ can be taken arbitrarily small, we obtain the reverse to Eq. (1.23), and the desired result follows. Q.E.D.

To see how the conclusion of the lemma can fail without the condition min'uEu Gl (.f $f(w), u$ ) >-00 for all $w$, let $u$ be a scalar, let $w \quad$ (WI, w2) be a two-dimensional vector, and let there be no constraints on $u$ and $w$ ( $U=\Re, W=\Re \times \Re$, where $\Re$ is the real line). Let also

$$
\left.\left.O o(w)=w_{1}, \quad f e w\right)=w_{2}, \quad G 1(f(w), u)=f e w\right)+u
$$

Then, for all $\mu \mathrm{E}$ M we have,

$$
\max _{w \in \mathbb{W}}\left[G o(w)+G_{1}(f(w), \mu(f(w)))\right]=\max _{w_{1} \in \Omega, w_{2} \in \Re}\left[w_{1}+w_{2}+\mu\left(w_{2}\right)\right]=00
$$

so that

$$
\min _{\mu \in M} \max _{w \in W}\left[G o(w)+\mathrm{G}_{1}(f(w), \mu(f(w)))\right] \quad 00 .
$$

On the other hand,

$$
\max _{w \in \mathbb{W}}\left[G o(w)+\min _{u \in \mathscr{Y}} \mathrm{G} 1(J(w), u)\right]=\max _{w_{1} \in \Re, w_{2} \in \Re}\left[\mathrm{WI}+\min _{u \in \Re}[\mathrm{w} 2+u]\right]=-\mathrm{oo}
$$

since $\min _{u \in \mathfrak{\Re}}\left[w_{2}+u\right]=-\infty$ for all $W 2$.
We now turn to proving the DP algorithm (1.21)-(1.22). The proof is similar to the one for the DP algorithm for stochastic problems. The optimal $\operatorname{cost} J^{*}(x o)$ of the problem is given by

$$
\begin{aligned}
J^{*}(x o)= & \min _{\boldsymbol{\mu}_{\mathbf{0}} \cdots} \min _{\mu_{N-1}} \max _{w_{0} \in W\left[x_{0}, \mu_{0}\left(x_{0}\right)\right]} \cdots_{w_{N-1} \in W\left[x_{N-1}, \mu_{N-1}\left(x_{N-1}\right)\right]} \\
& {\left[\sum_{k=0}^{N-1} g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)+9 N(X N)\right] }
\end{aligned}
$$

$$
\min _{\mu_{0}} \cdot . . \min _{\mu_{N-2}}\left[\min _{\mu_{N-1}} \max _{w_{0} \in W\left[x_{0}, \mu_{0}\left(x_{0}\right)\right]} \cdots_{w_{N-2} \in W\left[x_{N-2}, \mu_{N-2}\left(x_{N-2}\right)\right]}\right.
$$

$$
\sum_{[k=O}^{N-2} g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)+{\underset{w_{N-1} \in W\left[x_{N-1}, \mu_{N-1}\left(x_{N-1}\right)\right]}{ }}_{\max }
$$

$$
\left.\left.\left[g_{N-1}\left(x_{N-1}, \mu_{N-1}\left(x_{N-1}\right), w_{N-1}\right)+I N(X N)\right]\right]\right] .
$$

We can interchange the minimum over $\mu_{N-1}$ and the maximum over $W o, \ldots, w_{N-2}$ by applying Lemma 1.6 .1 with the identifications

$$
w=(w o, \text { WI, } \ldots, w N-2), \quad u=u_{N-1}, \quad j(W) \quad X N-I,
$$

$\mathrm{GO}(w) \begin{cases}\sum_{k=0}^{N-2} g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right) & \text { ifwk E } W_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right)\right) \text { for all } k, \\ \infty & \text { otherwise },\end{cases}$

$$
G_{1}(f(w), u)= \begin{cases}\hat{G}_{1}(f(w), u) & \text { if } u \mathrm{E} \mathrm{UN}_{-1}(J(w)) \\ \infty & \text { otherwise }\end{cases}
$$

where

$$
\begin{aligned}
\hat{G}_{1}(J(w), u)=\max _{w_{N-1} \in W_{N-1}(f(w), u)} & {[9 N-I(J(w), u, w N-I)} \\
& \left.+I N\left(I N-I\left(f(W), u, w_{N-1}\right)\right)\right]
\end{aligned}
$$

to obtain

$$
\begin{align*}
J^{*}\left(x_{0}\right)= & \min _{\mu_{0}} \cdots \min _{\mu_{N-2}} \\
& \operatorname{mon}_{0} \in W_{\left[x_{0}, \mu_{0}\left(x_{0}\right)\right]} \quad w_{N-2} \in W\left[x N_{-2}, \mu_{N-} 2(x N-2)\right]  \tag{1.24}\\
& {\left[\sum_{k=0}^{N-2} g_{k}\left(X k, \mu_{k}\left(x_{k}\right), w_{k}\right)+J_{N-1}\left(x_{N--1}\right)\right] . }
\end{align*}
$$ plication of Lemma 1.6.1) is implied by the assumption $/ N-I(X N-I)>-\infty$ for all $X N-I$. Now, by working with the expression for $J^{*}(x 0)$ in Eq. (1.24), and by similarly continuing backwards, with $N-1$ in place of $N$, etc., after $N$ steps we obtain $J^{*}\left(x_{0}\right)=J_{0}\left(x_{0}\right)$, which is the desired relation. The line of argument just given also shows that an optimal policy for the minimax problem can be constructed by minimizing in the right-hand side of the DP Eq. (1.22), similar to the case of the DP algorithm for the stochastic basic problem.

Unfortunately, as mentioned earlier, there are hardly any interesting examples of an analytical, closed-form solution of the DP algorithm (1.21)-(1.22). A cornputational solution, requires roughly comparable effort to the one of the stochastic DP algorithm. Instead of the expectation operation, one must carry out a maximization operation for each $X k$ and $k$.

Minimax control problems will be revisited in Chapter 4 in the context of reachability of target sets and target tubes (Section 4.6.2), and in Chapter 6 in the context of competitive games and computer chess (Section 6.3), and model predictive control (Section 6.5).

Dynamic programming is a simple mathematical technique that has been used for many years by engineers, mathematicians, and social scientists in a variety of contexts. It was Bellman, however, who realized in the early fifties that DP could be developed (in conjunction with the then appearing digital computer) into a systematic tool for optimization. In his influential books [Be157], [BeD62]' Bellman demonstrated the broad scope of DP and helped streamline its theory.

Following Bellman's works, there was much research on DP. In particular, the mathematical and algorithmic aspects of infinite horizon problems were extensively investigated, extensions to continuous-time problems were formulated and analyzed, and the mathematical issues discussed in Section 1.5, relating to the formulation of DP problems, were addressed. In addition, DP was used in a very broad variety of applications, ranging from many branches of engineering to statistics, economics, finance, and some of the social sciences. Samples of these applications will be given in subsequent chapters.

## EXERCISES

## 1.1

Consider the system

$$
x_{k+1}=x_{k}+u_{k}+w_{k}, \quad k=0,1,2,3,
$$

with initial state $x_{0} \quad 5$, and the cost function

$$
\sum_{k=0}^{3}\left(x_{k}^{2}+u_{k}^{2}\right)
$$

Apply the DP algorithm for the following three cases:
(a) The control constraint set $U k(X k)$ is $\{u \mid 0 \leq X k \quad u \leq 5$, $\boldsymbol{u}$ : integer $\}$ for all $X k$ and $k$, and the disturbance $W k$ is equal to zero for all $k$.
(b) The control constraint and the disturbance $w_{k}$ are as in part (a), but there is in addition a constraint $x_{4}=5$ on the final state. Hint: For this problem you need to define a state space for $x 4$ that consists of just the value $x_{4}=5$, and also to redefine $U_{3(X 3)^{\prime}}$ Alternatively, you may use a terminal cost $94(X 4)$ equal to a very large number for $X 4 \neq 5$.
(c) The control constraint is as in part (a) and the disturbance $W k$ takes the values -1 and 1 with equal probability $1 / 2$ for all $X k$ and $u_{k}$, except if $x_{k}+u_{k}$ is equal to 0 or 5 , in which case $U!k=0$ with probability 1 .

## 1.2

Carry out the calculations needed to verify that $J o(l)=2.67$ and $J o(2)=2.608$ in Example 1.3.2.

## 1.3

Suppose we have a machine that is either running or is broken down. If it runs throughout one week, it makes a gross profit of $\$ 100$. If it fails during the week, gross profit is zero. If it is running at the start of the week and we perform preventive maintenance, the probability that it will fail during the week is 0.4 . If we do not perform such maintenance, the probability of failure is 0.7 . However, maintenance will cost $\$ 20$. When the machine is broken down at the start of the week, it may either be repaired at a cost of $\$ 40$, in which case it will fail during the week with a probability of 0.4 , or it may be replaced at a cost of $\$ 150$ by a new machine that is guaranteed to run through its first week of operation. Find the optimal repair, replacement, and maintenance policy that maximizes total profit over four weeks, assuming a new machine at the start of the first week.

## 1.4

A game of the blackjack variety is played by two players as follows: Both players throw a die. The first player, knowing his opponent's result, may stop or may throw the die again and add the result to the result of his previous throw. He then may stop or throw again and add the result of the new throw to the sum of his previous throws. He may repeat this process as many times as he wishes. If his sum exceeds seven (Le., he busts), he loses the game. If he stops before exceeding seven, the second player takes over and throws the die successively until the sum of his throws is four or higher. If the sum of the second player is over seven, he loses the game. Otherwise the player with the larger sum wins, and in case of a tie the second player wins. The problem is to determine a stopping strategy for the first player that maximizes his probability of winning for each possible initial throw of the second player. Formulate the problem in terms of DP and find an optimal stopping strategy for the case where the second player's initial throw is three. Hint: Take $N=6$ and a state space consisting of the following 15 states:

## x1: busted

$$
x^{I}+^{i}: \text { already stopped at sum } \mathrm{i}(1 \leq i \leq 7)
$$

$x^{8+i}$ : current sum is i but the player has not yet stopped $(1 \leq \mathrm{i} \leq 7)$.
The optimal strategy is to throw until the sum is four or higher.

## 1.5 (Compu.ter Assignment)

In the classical game of blackjack the player draws cards knowing only one card of the dealer. The player loses upon reaching a sum of cards exceeding 21. If the player stops before exceeding 21, the dealer draws cards until reaching 17 or higher. The dealer loses upon reaching a sum exceeding 21 or stopping at a lower sum than the player's. If player and dealer end up with an equal sum no one wins. In all other cases the dealer wins. An ace for the player may be counted as a 1 or an 11 as the player chooses. An ace for the dealer is counted as an 11 if this results in a sum from 17 to 21 and as a 1 otherwise. Jacks, queens, and kings count as 10 for both dealer and player. We assume an infinite card deck so the probability of a particular card showing up is independent of earlier cards.
(a) For every possible initial dealer card, calculate the probability that the dealer will reach a sum of $17,18,19,20,21$, or over 21.
(b) Calculate the optimal choice of the player (draw or stop) for each of the possible combinations of dealer's card and player's sum of 12 to 20. Assume that the player's cards do not include an ace.
(c) Repeat part (b) for the case where the player's cards include an ace.

## 1.6 (Discounted Cost per Stage)

In the framework of the basic problem, consider the case where the cost is of the form

$$
\underset{\substack{w_{k} \\ k=0,1_{,}, \ldots, N-1}}{E}\left\{\alpha^{N} g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} \alpha^{k} g_{k}\left(x_{k}, u_{k}, w_{k}\right)\right\}
$$

where $\alpha$ is a discount factor with $0<\alpha<1$. Show that an alternate form of the DP algorithm is given by

$$
V N(X N)=9 N(X N),
$$

## 1.7 (Exponential Cost Function)

In the framework of the basic problem, consider the case where the cost is of the form

$$
\underset{\substack{w_{k} \\ k=0,1, \ldots, N-1}}{E}\left\{\exp \left(g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}
$$

(a) Show that the optimal cost and an optimal policy can be obtained from the DP-like algorithm

$$
J_{N}\left(x_{N}\right)=\exp \left(g_{N}\left(x_{N}\right)\right)
$$

$$
J_{k}\left(x_{k}\right)=\min _{u_{k} \in U_{k}\left(x_{k}\right)} \underset{w_{k}}{E}\left\{J_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right) \exp \left(g_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}
$$

(b) Define the functions $\operatorname{Vk}(X k)=\operatorname{InJk}(x k)$. Assume also that $9 k$ is a function of $X k$ and $u_{k}$ only (and not of $W k$ ). Show that the above algorithm can be rewritten as

$$
\begin{gathered}
V_{N}\left(x_{N}\right)=g_{N}\left(x_{N}\right) \\
V_{k}\left(x_{k}\right)=\min _{u_{k} \in U_{k}\left(x_{k}\right)}\left\{g_{k}\left(x_{k}, u_{k}\right)+\ln \underset{w_{k}}{E}\left\{\exp \left(V_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right)\right\}\right\}
\end{gathered}
$$

Note: The exponential cost function is an example of a risk-sensitive cost junction that can be used to encode a preference for policies with a small variance of the cost $9 N(X N)+\sum_{k=0}^{N-1} 9 k\left(X k, u_{k}, W k\right)$. The associated problems have a lot of interesting properties, which are discussed in several sources, e.g., Whittle [Whi90], Fernandez-Gaucherand and Markus [FeM94]' James, Baras, and Elliott [JBE94], Basar and Bernhard [BaB95].

## 1.8 (Terrninating Process)

In the framework of the basic problem, consider the case where the system evolution terminates at time $i$ when a given value $\bar{w}_{i}$ of the disturbance at time $i$ occurs, or when a termination decision $u_{i}$ is made by the controller. If termination occurs at time $i$, the resulting cost is

$$
\mathrm{T}+\sum_{k=0} 9 k\left(X k, u_{k}, W k\right)
$$

where T is a termination cost. If the process has not terminated up to the final time IV, the resulting cost is $9 N(X N)+\sum_{k=0}^{N-1} 9 k\left(X k,{ }^{\prime} U k, W k\right)$. Reformulate the problem into the framework of the basic problem. Hint: Augment the state space with a special termination state.

## 1.9 (Multiplicative Cost)

In the framework of the basic problem, consider the case where the cost has the multiplicative form

$$
\underset{\substack{w_{k} \\ \mathrm{k}=\mathrm{O}, \mathbf{I}, \ldots, \mathrm{~N}-\mathbf{I}}}{E}\left\{9 N(X N) .9 N-1(X N-1, U N-I, W N-1)^{\prime \prime} \quad 90(\text { Хo, uo, wo })\right\} .
$$

Develop a DP-like algorithm for this problem assuming that $9 k(X k, U k, W k) \geq$ for all $X k, u_{k}, W k$, and $k$.

### 1.10

Assume that we have a vessel whose maximum weight capacity is $z$ and whose cargo is to consist of different quantities of $1 V$ different items. Let $\mathrm{V}_{i}$ denote the value of the ith type of item, w'i the weight of ith type of itern, and $x_{i}$ the number of items of type i that are loaded in the vessel. The problem is to find the most valuable cargo, i.e., to maximize $\sum_{i=1}^{N} x_{i} v_{i}$ subject to the constraints $\sum_{i=1}^{N} x_{i} w_{i} \leq z$ and $x_{i}=0,1,2, \ldots$ Formulate this problem in terms of DP.

### 1.11

Consider a device consisting of IV stages connected in series, where each stage consists of a particular component. The components are subject to failure, and to increase the reliability of the device duplicate components are provided. For $j=1,2, \ldots, I V$, let $(1+m j)$ be the number of components for the $j$ th stage, let $p_{j}(m j)$ be the probability of successful operation of the $j$ th stage when $(1+m j)$ components are used, and let Cj denote the cost of a single component at the $j$ th stage. Formulate in terms of DP the problem of finding the number of cornponents at each stage that maximize the reliability of the device expressed by

$$
p_{1}\left(m_{1}\right) \cdot p_{2}\left(m_{2}\right) \cdots p_{N}\left(m_{N}\right)
$$

subject to the cost constraint $\sum_{j=1}^{N} \operatorname{cjmj} \leq A$, where $A>$ a is given.

### 1.12 (Minimization over a Subset of Policies)

This problem is primarily of theoretical interest (see the end of Section 1.5). Consider a variation of the basic problem whereby we seek

$$
\min _{\pi \in \bar{\Pi}} J K(X O),
$$

where $\tilde{\Pi}$ is some given subset of the set of sequences $\left\{\mu_{0}, \mu_{1}, \ldots, \mu_{N-1}\right\}$ of functions $\mu_{k}: S k \rightarrow C k$ with $\mu_{k}\left(x_{k}\right)$ E $U k(X k)$ for all $X k$ E $S k$. Assume that

$$
\pi^{*}=\left\{\mu_{0}^{*}, \mu_{1}^{*}, \ldots, \mu_{N-1}^{*}\right\}
$$

belongs to $\tilde{\Pi}$ and attains the minimum in the DP algorithm; that is, for all $k \quad 0,1, \ldots, \mathrm{IV}-1$ and $X k \mathrm{E} S k$,

$$
\begin{aligned}
J k(X k)= & \underset{w_{k}}{E}\left\{g_{k}\left(x_{k}, \mu_{k}^{*}\left(x_{k}\right), w_{k}\right)+J_{k+1}\left(f_{k}\left(x_{k}, \mu_{k}^{*}\left(x_{k}\right), w_{k}\right)\right)\right\} \\
& \min _{u_{k} \in U_{k}\left(x_{k}\right)} \underset{w_{k}}{E}\left\{9 k(X k, U k, W k)+J_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\},
\end{aligned}
$$

with $I N(X N)=9 N(X N)$. Assume further that the functions $J_{k}$ are real-valued and that the preceding expectations are well-defined and finite. Show that $1 T^{*}$ is optimal within $\tilde{\Pi}$ and that the corresponding optimal cost is equal to Jo(xo).

### 1.13 (Semilinear Systems)

Consider a problem involving the system

$$
x_{k+1}=A_{k} x_{k}+f_{k}\left(u_{k}\right)+w_{k},
$$

where $X k E \Re^{n}, f k$ are given functions, and $A_{k}$ and $\psi_{k}$ are random $n \times n$ matrices and n-vectors, respectively, with given probability distributions that do not depend on $X k, U k$ or prior values of $A_{k}$ and $w_{k}$. Assume that the cost is of the form

$$
\underset{\substack{A_{k}, w_{k} \\ k=0,1, \ldots, N-1}}{E}\left\{c_{N}^{\prime} x_{N}+\sum_{k=0}^{N-1}\left(c_{k}^{\prime} x_{k}+9^{\prime}(, 1 .(X k))\right)\right\}
$$

where $C k$ are given vectors and $g k$ are given functions. Show that if the optimal cost for this problem is finite and the control constraint sets $U k(X k)$ are independent of $X k$, then the cost-to-go functions of the DP algorithm are affine (linear plus constant). Assuming that there is at least one optimal policy, show that there exists an optimal policy that consists of constant functions $\mu_{k}^{*} i$ that is, $\mu_{k}^{*}\left(x_{k}\right)=$ constant for all $X k$ E $\Re^{n}$.

### 1.14

A farmer annually producing $X k$ units of a certain crop stores $(1-U k) X k$ units of his production, where $0 \leq U k \leq 1$, and invests the remaining $U k X k$ units, thus increasing the next year's production to a level $X k+l$ given by

$$
x_{k+1}=x_{k}+w_{k} u_{k} x_{k}, \quad k=0,1, \ldots, N-1 .
$$

The scalars $w_{k}$ are independent random variables with identical probability distributions that do not depend either on $X k$ or $U k$. Furthermore, $E\left\{w_{k}\right\}=\bar{w}>0$. The problem is to find the optimal investment policy that maximizes the total expected product stored over $N$ years

$$
\underset{\substack{w_{k} \\ k=0,1, \ldots, N-1}}{E}\left\{X N+\sum_{k=0}^{N-1}(1-U k) X k\right\}
$$

Show the optimality of the following policy that consists of constant functions:
(a) If $\bar{w}>1, \mu_{0}^{*}\left(x_{0}\right)=. .,=\mu_{N-1}^{*}(X N-l)=1$.
(b) If $0<\overline{\boldsymbol{w}}<\operatorname{liN}, \mu_{0}^{*}\left(x_{0}\right)=\ldots=\mu_{N-1}^{*}\left(x_{N-1}\right)=0$.
(c) If $\operatorname{li} N \leq \bar{w} \leq 1$,

$$
\begin{gathered}
\mu_{0}^{*}\left(x_{0}\right)=\cdots=\mu_{N-\bar{k}-1}\left(x_{N-\bar{k}-1}\right)=1 \\
\mu_{N-\bar{k}}^{*}\left(x_{N-\bar{k}}\right)=. .,=\mu_{N-1}^{*}\left(x_{N-1}\right)=0
\end{gathered}
$$

where $\vec{k}$ is such that $1 /(\bar{k}+1)<\bar{w} \leq 1 / \bar{k}$.

### 1.15

Let $X k$ denote the number of educators in a certain country at time $k$ and let $Y k$ denote the number of research scientists at time $k$. New scientists (potential educators or research scientists) are produced during the kth period by educators at a rate $\gamma_{k}$ per educator, while educators and research scientists leave the field due to death, retirement, and transfer at a rate $\delta_{k}$. The scalars $\gamma_{k}, k=0,1, \ldots, N-1$, are independent identically distributed random variables taking values within a closed and bounded interval of positive numbers. Similarly $\boldsymbol{\delta}_{\boldsymbol{k}}, k \quad O, l, \ldots, N-1$, are independent identically distributed and take values in an interval $\left[\delta, \delta^{\prime}\right]$ with $\mathrm{O}<\delta \leq \delta^{\prime}<1$. By means of incentives, a science policy maker can determine the proportion $U k$ of new scientists produced at time $k$ who become educators. Thus, the number of research scientists and educators evolves according to the equations

$$
\begin{gathered}
x_{k+1}=\left(1-\delta_{k}\right) x_{k}+u_{k} \gamma_{k} x_{k} \\
y_{k+1}=\left(1-\delta_{k}\right) y_{k}+\left(1-u_{k}\right) \gamma_{k} x_{k}
\end{gathered}
$$

The initial numbers $X o$, yo are known, and it is required to find a policy

$$
\left\{\mu_{0}^{*}\left(x_{0}, y_{0}\right), \ldots, \mu_{N-1}^{*}\left(x_{N-1}, y_{N-1}\right)\right\}
$$

with

$$
0<\alpha \leq \mu_{k}^{*}\left(x_{k}, y_{k}\right) \leq \beta<1, \quad \text { for all } x k, Y k, \text { and } k
$$

which maximizes $E_{\gamma_{k}, \delta_{k}}\{Y N\}$ (Le., the expected final number ofresearch scientists after $N$ periods). The scalars $\alpha$ and $f J$ are given.
(a) Show that the cost-to-go functions $J_{k}\left(x_{k}, Y k\right)$ are linear; that is, for some scalars $\xi_{k}, \zeta_{k}$,

$$
J_{k}\left(x_{k}, y_{k}\right)=\xi_{k} x_{k}+\zeta_{k} y_{k}
$$

(b) Derive an optimal policy $\left\{\mu_{0}^{*}, \ldots, \mu_{N-1}^{*}\right\}$ under the assumption

$$
E\left\{\gamma_{k}\right\}>E\left\{\delta_{k}\right\}
$$

and show that this optimal policy can consist of constant functions.
(c) Assume that the proportion of new scientists who become educators at time $k$ is $U k+\boldsymbol{\epsilon}_{\boldsymbol{k}}$ (rather than $U k$ ), where $\boldsymbol{\epsilon}_{\boldsymbol{k}}$ are identically distributed independent random variables that are also independent of $\gamma_{k}, \delta_{k}$ and take values in the interval $[-\alpha, 1-\beta]$. Derive the form of the cost-to-go functions and the optimal policy.

Given a sequence of matrix multiplications

$$
M_{1} M_{2} \cdots M_{k} M_{k+1} \cdots M_{N}
$$

where each $M_{k}$ is a matrix of dirnension $n k \times n k+l$, the order in which multiplications are carried out can make a difference. For example, if $n l=1, n 2=10$, $n_{3}=1$, and $n 4=10$, the calculation $\left(\left(M_{1} M_{2}\right) M_{3}\right)$ requires 20 scalar multiplications, but the calculation $\left(M_{1}\left(M_{2} M_{3}\right)\right)$ requires 200 scalar multiplications (multiplying an $\mathrm{m} \times n$ matrix with an $n \times k$ matrix requires mnk scalar multiplications).
(a) Derive a DP algorithm for finding the optimal multiplication order [any order is allowed, including orders that involve multiple partial products each consisting of two or more adjacent matrices, e.g., ((N!hM2)(M3M4))j. Solve the problem for $\mathrm{N}=3, n l=2, n 2=10, n 3=5$, and $n 4=\mathrm{I}$.
(b) Derive a DP algorithm for finding the optimal multiplication order within the class of orders where at each step, we maintain only one partial product that consists only of adjacent matrices, e.g., $\left(\left(M_{1}\left(M_{2} M_{3}\right)\right) M_{4}\right)$.

### 1.17

The paragraphing problem deals with breaking up a sequence of N words of given lengths into lines of length $A$. Let $w, \ldots, w \mathrm{~N}$ be the words and let $L_{1}, \ldots, L_{\mathrm{N}}$ be their lengths. In a simple version of the problem, words are separated by blanks whose ideal width is $b$, but blanks can stretch or shrink if necessary, so that a line $w_{i}, w_{i+1}, \ldots, W_{i+k}$ has length exactly $A$. The cost associated with the line is $(k+I) / b^{\prime}-b \mid$, where $/ / \quad\left(A-L_{i}-\ldots \quad L i+k\right) /(k+1)$ is the actual average width of the blanks, except if we have the last line ( $\mathrm{N}=\mathrm{i}+k$ ), in which case the cost is zero when $b^{\prime} \geq b$. Formulate a DP algorithm for finding the minimum cost separation. Hint: Consider the subproblems of optimally separating $w_{i}, \ldots, w_{N}$ for $i=1, \ldots, N$.

### 1.18 [Shr81]

A decision maker must continually choose between two activities over a time interval $[0, \mathrm{Tj}$. Choosing activity i at time $t$, where $\mathrm{i}=1,2$, earns reward at a rate $g_{i}(t)$, and every switch between the two activities costs $\mathrm{c}>0$. Thus, for example, the reward for starting with activity 1 , switching to 2 at time $t 1$, and switching back to 1 at time $t 2>t /$ earns total reward

$$
\int_{0}^{t_{1}} g_{1}(t) d t+\int_{t_{1}}^{t_{2}} g_{2}(t) d t+\int_{t_{2}}^{T} g_{1}(t) d t-2 c .
$$

We want to find a set of switching times that maximize the total reward. Assume that the function $91(t)-92(t)$ changes sign a finite number of times in the interval $[0, T j$. Formulate the problem as a finite horizon problem and write the corresponding DP algorithm.
1.19 (Garnes)
(a) Consider a smaller version of a popular puzzle game. Three square tiles numbered 1,2 , and 3 are placed in a $2 \times 2$ grid with one space left empty The two tiles adjacent to the empty space can be moved into that space, thereby creating new configurations. Use a DP argument to answer the question whether it is possible to generate a given configuration starting from any other configuration.
(b) From a pile of eleven matchsticks, two players take turns removing one or four sticks. The player who removes the last stick wins. Use a DP argument to show that there is a winning strategy for the player who plays first.

### 1.20 (The Counterfeit Coin Problem.)

We are given six coins, one of which is counterfeit and is known to have different weight than the rest. Construct a strategy to find the counterfeit coin using a two-pan scale in a minimum average number of tries. Hint: There are two initial decisions that make sense: (1) test two of the coins against two others, and (2) test one of the coins against one other.

### 1.21 (Regular Polygon Theorem)

According to a famous theorem (attributed to the ancient Greek geometer Zenodorus), of all $N$-side polygons inscribe'd in a given circle, those that are regular (all sides are equal) have maximal area.
(a) Prove the theorem by applying DP to a suitable problem involving sequential placement of N points in the circle.
(b) Use DP to solve the problem of placing a given number of points on a subarc of the circle, so as to maximize the area of the polygon whose vertices are these points, the endpoints of the subarc, and the center of the circle.

### 1.22 (Inscribed Polygon of IVIaximal Perhneter)

Consider the problem of inscribing an N -side polygon in a given circle, so that the polygon has maximal perimeter.
(a) Formulate the problem as a DP problem involving sequential placernent of N points in the circle.
(b) Use DP to show that the optimal polygon is regular (all sides are equal).

### 1.22 (Monotonicity Property of DP)

An evident, yet very important property of the DP algorithm is that if the terminal cost 9 N is changed to a uniformly larger cost $\bar{g} \mathrm{~N}$ [Le., $9 \mathrm{~N}\left(x_{\mathrm{N}}\right) \leq \bar{g}_{\mathrm{N}}\left(x_{\mathrm{N}}\right)$ for all $X_{N} J$, then the last stage cost-to-go $J_{N-1}\left(X_{N}-1\right)$ will be uniformly increased. More generally, given two functions $J_{k+1}$ and $\mathbf{J}_{\mathrm{k}+1}$ with $J_{k+1}\left(x_{k+1}\right) \leq J k+l(X k+l)$ for all $x_{k+1}$, we have, for all $X k$ and $u_{k} \mathrm{E} U k(X k)$,

$$
\begin{aligned}
\underset{w_{k}}{E}\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)\right. & \left.+J k+l\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\} \\
& \leq \underset{w_{k}}{E}\left\{g_{k}\left(x_{k}, u_{k}, W k\right)+J_{\mathrm{k}+1}(J k(X k, U k, W k))\right\}
\end{aligned}
$$

Suppose now that in the basic problem the system and cost are time invariant; that is, $S k \equiv S, O k \equiv O, D k \equiv \mathrm{D}, f k \equiv f, U k \equiv U, P k \equiv P$, and $g k \equiv 9$ for some $S, C, D, f, U, P$, and $g$. Show that if in the DP algorithm we have $I N-l(X) \leq I N(X)$ for all $x E S$, then

$$
J_{k}(x) \leq J_{k+1}(x), \quad \text { for all } x E S \text { and } k
$$

Similarly, if we have $I N-l(X) \geq I N(X)$ for all $X E S$, then

$$
J_{k}(x) \geq J_{k+1}(x), \quad \text { for all } x E S \text { and } k
$$

### 1.24 (Traveling Repairman Problem)

A repairman must service $n$ sites, which are located along a line and are sequentially numbered $1,2, \ldots, n$. The repairman starts at a given site $s$ with $1<s<n$, and is constrained to service only sites that are adjacent to the ones serviced so far, Le., if he has already serviced sites $i, i+1, \ldots, j$, then he may service next only site $i-1$ (assuming $1<\mathrm{i}$ ) or site $\mathrm{j}+1$ (assuming $\mathrm{j}<n$ ). There is a waiting $\operatorname{cost} c_{i}$ for each time period that site $i$ has remained unserviced and there is a travel cost $i i j$ for servicing site $j$ immediately after servicing site i. Formulate a DP algorithm for finding a minimum cost service schedule.

### 1.25

An unscrupulous innkeeper charges a different rate for a room as the day progresses, depending on whether he has many or few vacancies. His objective is to maximize his expected total income during the day. Let $x$ be the number of empty rooms at the start of the day, and let $y$ be the number of customers that will ask for a room in the course of the day. We assume (somewhat unrealistically) that the innkeeper knows $y$ with certainty, and upon arrival of a customer, quotes one of $r n$ prices $r_{i}, i=1, \ldots, r n$, where $0<T^{11} \leq T 2 \leq \ldots \leq r_{m}$. A quote of a rate $r_{i}$ is accepted with probability $P i$ and is rejected with probability $\mathbf{1}-p i$, in which case the customer departs, never to return during that day,
(a) Formulate this as a problem with $y$ stages and show that the maximal expected income, as a function of $x$ and $y$, satisfies the recursion

$$
J(x, y)=\max _{i=1, \ldots, m}\left[p_{i}\left(r_{i}+J(x-1, y-1)\right)+\left(\overline{1}-p_{i}\right) J(x, y-1)\right]
$$

for all $x \geq 1$ and $y \geq 1$, with initial conditions

$$
J(X, O)=.1(0, y)=0, \quad \text { for all } x \text { and } y
$$

Assuming that the product PiTi is monotonically nondecreasing with i, and that $P i$ is monotonically nonincreasing with $i$, show that the innkeeper should always charge the highest rate $T_{\cdot m}$
(b) Consider a variant of the problem where each arriving customer, with probability $p_{i}$, offers a price $T i$ for a room, which the innkeeper may accept or reject, in which case the customer departs, never to return during that day. Show that an appropriate DP algorithm is

$$
J(X, y)=\sum_{i=1}^{m} p_{i} \max [\mathrm{Ti}+J(x \quad 1, Y \quad 1), J(x, y-1) \mathrm{J},
$$

with initial conditions

$$
J(X, O)=J(O, y)=0, \quad \text { for all } x \text { and } y
$$

Show also that for given $x$ and $y$ it is optimal to accept a customer's offer if it is larger than some threshold $\bar{r}(x, v)$. Hint: This part is related to DP for uncontrollable state components (d. Section 1.4).

### 1.26 (Investing in a Stock)

An investor observes at the beginning of each period $k$ the price $X k$ of a stock and decides whether to buy 1 unit, sell 1 unit, or do nothing. There is a transaction cost c for buying or selling. The stock price can take one of $n$ different values $v^{1}, \ldots, v^{n}$ and the transition probabilities $p_{i j}^{k}=P\left\{X k+l=v^{j} \mid X k=v^{i}\right\}$ are known. The investor wants to maximize the total worth of his stock at a fixed final period $N$ minus his investment costs from period 0 to period $N-1$ (revenue from a sale is viewed as negative cost). We assume that the function

$$
P_{k}(x)=E\left\{x_{N} \mid x_{k}=x\right\}-x
$$

is monotonically nonincreasing as a function of $x$; that is, the expected profit from a purchase is a nonincreasing function of the purchase price.
(a) Assume that the investor starts with $N$ or more units of stock and an unlimited amount of cash, so that a purchase or sale decision is possible at each period regardless of the past decisions and the current price. For every period $k$, let $\underline{x}_{k}$ be the largest value of $x$ E $\left\{\mathrm{VI}, \ldots, v^{n}\right\}$ such that
$\operatorname{Pk}(X)>c$, and let $\bar{x}_{k}$ be the smallest value of $x \mathrm{E}\left\{\mathrm{VI}, \ldots, v^{n}\right\}$ such that $P_{k}(x)<-c$. Show that it is optimal to buy if $x_{k} \leq \underline{x}_{k}$, sell if $\bar{x}_{k} \leq x k$, and do nothing otherwise. Hint: Formulate the problem as one of maximizing

$$
E\left\{\sum_{k=0}^{N-1}\left(u_{k} P_{k}\left(x_{k}\right)-c\left|u_{k}\right|\right)\right\}
$$

where $u_{k} \mathrm{E}\{-1,0, \mathrm{I}\}$.
(b) Formulate an efficient DP algorithm for the case where the investor starts with less than $N$ units of stock and an unlimited amount of cash. Show that it is still optimal to buy if $x k \leq \underline{x}_{k}$ and it is still not optimal to sell if $x_{k}<\bar{x}_{k}$. Could it be optimal to buy at any prices $X k$ greater than $\underline{x}_{k}$ ?
(c) Consider the situation where the investor initially has $N$ or more units of stock and there is a constraint that for any time $k$ the number of purchase decisions up to $k$ should not exceed the number of sale decisions up to $k$ by more that a given fixed number $m$ (this models approximately the situation where the investor has a limited initial amount of cash). Formulate an efficient DP algorithm for this case. Show that it is still optimal to sell if $\bar{x}_{k} \leq x_{k}$ and it is still not optimal to buy if $\underline{x}_{k}<\chi k$.
(d) Consider the situation where there are restrictions on both the initial amount of stock as in part (b), and the number of purchase decisions as in part (c). Derive a DP algorithm for this problem.
(e) How would the analysis of (a)-(d) be affected if cash is invested at a given fixed interest rate?

## Deterministic Systems and the

## Shortest Path Problern

## Contents

2.1. Finite-State Systems and Shortest Paths
2.2. Some Shortest Path Applications . . . p. 68
2.2.1. Critical Path Analysis . . . . . p. 68
2.2.2. Hidden Markov Models and the Viterbi Algorithm p. 70
2.3. Shortest Path Algorithms
p. 77
2.3.1. Label Correcting Methods . . . . . . . p. 78
2.3.2. Label Correcting Variations - $\boldsymbol{A}^{*}$ Algorithm p. 87
2.3.3. Branch-and-Bound
2.3.4. Constrained and Multiobjective Problems
p. 88
4. Notes, Sources, and Exercises.
p. 97

In this chapter, we focus on deterministic problems, that is, problems where each disturbance ' $W k$ can take only one value. Such problems arise in many important contexts and they also arise in cases where the problem is really stochastic but, as an approximation, the disturbance is fixed at some typical value; see Chapter 6.

An important property of deterministic problems is that, in contrast with stochastic problems, using feedback results in no advantage in terms of cost reduction. In other words, minimizing the cost over admissible policies $\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}$ results in the same optimal cost as minimizing over sequences of control vectors $\left\{u_{0}, \ldots, u_{N-1}\right\}$. This is true because given a policy $\left\{\mu_{0}, \ldots \mu_{N-1}\right\}$ and the initial state $X o$, the future states are perfectly predictable through the equation

$$
x_{k+1}=f_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right)\right), \quad k=0,1, \ldots, \mathbf{N}-1,
$$

and the corresponding controls are perfectly predictable through the equation

$$
u_{k}=\mu_{k}\left(x_{k}\right), \quad k=0,1, \ldots, \mathbf{N}-\mathbf{I} .
$$

Thus, the cost achieved by an admissible policy $\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}$ for a deterministic problem is also achieved by the control sequence $\left\{u_{0}, \ldots, u_{N-1}\right\}$ defined above. As a result, we may restrict attention to sequences of controls without loss of optimality.

The difference just discussed between deterministic and stochastic problems often has important computational implications. In particular, in a deterministic problem with a "continuous space" character (states and controls are Euclidean vectors), optimal control sequences may be found by deterministic variational techniques to be discussed in Chapter 3, and by widely used iterative optimal control algorithms such as steepest descent, conjugate gradient, and Newton's method (see e.g., nonlinear programming texts such as Bertsekas [Ber99) or Luenberger [Lue84)). These algorithms, when applicable, are usually more efficient than DP. On the other hand, DP has a wider scope of appl-icability s-ince it can handle difficult constraint sets such as 'integer or discrete sets. Furthermore, DP leads to a globally optimal solution as opposed to variational techniques, for which this cannot be guaranteed in general.

In this chapter, we consider deterministic problems with a discrete character for which variational optimal control techniques are inapplicable, so that specialized forms of DP are the principal solution methods.

### 2.1 FINITE-STATE SYSTEMS AND SHORTEST PATHS

Consider a deterministic problem where the state space $S k$ is a finite set for each $k$. Then at any state $X k$, a control $U k$ can be associated with a
transition from the state $X k$ to the state $f k\left(x k, u_{k}\right)$, at a cost $g_{k}\left(x_{k}, u_{k}\right)$. Thus a finite-state deterministic problem can be equivalently represented by a graph such as the one of Fig. 2.1.1, where the arcs correspond to transitions between states at successive stages and each arc has an associated cost. To handle the final stage, an artificial terrninal node $t$ has been added. Each state $x N$ at stage $N$ is connected to the terminal node $t$ with an arc having cost $g N(X N)$. Control sequences correspond to paths originating at the initial state (node $s$ at stage 0 ) and terminating at one of the nodes corresponding to the final stage $N$. If we view the cost of an arc as its length, we see that a determinist'ic finite-state problem is equivalent to finding a minimum-length (or shortest) path from the initial node $s$ of the graph to the terminal node $t$. Here, by a path we mean a sequence of arcs of the form $\left(j_{1}, j_{2}\right),(j 2, j 3), \ldots,\left(j_{k-1}, j_{k}\right)$, and by the length of a path we mean the sum of the lengths of its arcs.


Figure 2.1.1 Transition graph for a deterministic finite-state system. Nodes correspond to states. An arc with start and end nodes $x k$ and $x k+l$, respectively, corresponds to a transition of the form $x k+1 \quad f_{k}(X k, U k)$. We view the cost $9 k(X k, U k)$ of the transition as the length of this arc. The problem is equivalent to finding a shortest path from the initial node $s$ to the terminal node $t$.

Let us denote
$a_{i j}^{k}=$ Cost of transition at stage $k$ from state i E to state $j \mathrm{E} S k+l$, $a_{i t}^{N}=$ Terminal cost of state i E SN [which is $9 N(i)$ ),
where we adopt the convention $a_{i j}^{k}=00$ if there is no control that moves the state from $i$ to $j$ at stage $k$. The DP algorithm takes the form

$$
\begin{gather*}
I N(i)=a_{i t}^{N}, \quad i \mathrm{E} \mathrm{SN},  \tag{2.1}\\
J k(i)=\min _{J E S k+l}\left[a_{i j}^{k}+J_{k+1}(j)\right], \quad i E S k, \quad \mathrm{k}=\mathrm{O}, 1, \ldots, \mathrm{~N}-1 .
\end{gather*}
$$

The optimal cost is $J o(s)$ and is equal to the length of the shortest path from $s$ to $t$.

## A Forward DP Algorithm

The preceding algorithm proceeds backward in time. It is possible to derive an equivalent algorithm that proceeds forward in time by means of the following simple observation. An optimal path from $s$ to $t$ is also an optimal path from t to $s$ in a "reverse" shortest path problem where the direction of each aTc is reversed and its length is left unchanged. The DP algorithm corresponding to this "reverse" problem starts from the states xı E 81 of stage 1, proceeds to states $x_{2} \mathrm{E} 82$ of stage 2, and continues all the way to states $x_{N}$ E $8 N$ of stage $\mathbf{N}$. It is given by

$$
\begin{equation*}
\tilde{J}_{N}(j)=a_{s j}^{0}, \tag{2.3}
\end{equation*}
$$

$$
\begin{equation*}
\left.\min _{i \in S_{N-k}} l a_{1 j-}^{N-}{ }^{k}+\tilde{J}_{k+1}\binom{( }{\imath}\right], \quad j \mathrm{E} 8_{N-k+l,} \quad k=1,2, \ldots, \mathrm{~N}-1 \tag{2.4}
\end{equation*}
$$

1'he optimal cost is

$$
\tilde{J}_{0}(t)=\min _{i \in S_{N}}\left[a_{i t}^{N}+\tilde{J}_{1}(i)\right] .
$$

The backward algorithm (2.1)-(2.2) and the forward algorithm (2.3)-(2.4) yield the same result in the sense that

$$
J o(s)=J o(t)
$$

and an optimal control sequence (or shortest path) obtained from anyone of the two is optimal for the original problem. We may view $\widetilde{J}_{k}(j)$ in Eq. (2.4) as an optimal cost-to-arrive to state $j$ from the initial state $s$. This should be contrasted with $J k(i)$ in Eq. (2.2), which represents the optimal cost-to-go from state $i$ to the terminal state $t$.

An important use of the forward DP algorithm arises in real-time applications where the stage $k$ problem data are unknown prior to stage $k$, and are revealed to the controller just before stage $k$ begins. An example will be given in connection with the state estimation of Hidden Markov Models in Section 2.2.2. Note that to derive the forward DP algorithm, we used the shortest path formulation, which is available only for deterministic problems. Indeed, for stochastic problems, there is no analog of the forward DP algorithm.

In conclusion, a deteTm'in'ist'ic finite-state pTOblem is equivalent to a special type of shoTtest path pTOblem and can be solved by either the ordinary (backward) IJP algoTithrn or by an alternat'ive forward DP algorithm. It is also interesting to note that any shortest path pTOblem can be posed as a determ:inistic finite-state DP pTOblem, as we now show.

Converting a Shortest Path Problem to a Deterministic $\mathbb{F i n i t e}-$ State Problem

Let $\{\mathrm{I}, 2, \ldots, N, \mathrm{t}\}$ be the set of nodes of a graph, and let aij be the cost of moving from node $\boldsymbol{i}$ to node $j$ (also referred to as the length of the arc joining $\boldsymbol{i}$ and $j$ ). Node $t$ is a special node, which we call the dest'inat'ion. We allow the possibility aij $=$ oo to account for the case where there is no arc joining nodes $i$ and $j$. We want to find a shortest path from each node $i$ to node t , i.e., a sequence of moves that minimizes total cost to get to $t$ from each of the nodes $1,2, \ldots, N$.

For the problem to have a solution, it is necessary to make an assumption relating to cycles, i.e., paths of the form $\left(i, j_{1}\right),\left(j_{1}, j 2\right), \ldots,(\mathrm{j} k, \mathrm{i})$ that start and end at the same node. We must exclude the possibility that a cycle has negative total length. Otherwise, it would be possible to decrease the length of some paths to arbitrarily small values simply by adding more and more negative-length cycles. We thus assume that all cycles have nonnegative length. With this assumption, it is clear that an optimal path need not take more than N moves, so we may limit the number of rnoves to $N$. We formulate the problem as one where we require exactly $N$ moves but allow degenerate moves from a node 'i to itself with cost aii $=0$, We denote for $\boldsymbol{i}=1, \ldots, N, k=0, \ldots, N-1$,

$$
J_{k}(i)=\text { optimal cost of getting from } \boldsymbol{i} \text { tot in } \mathrm{N}-k \text { moves. }
$$

Then the cost of the optimal path from $\boldsymbol{i}$ to t is $\mathrm{Jo}^{( }$( $i$ ).
It is possible to formulate this problem within the framework of the basic problem and subsequently apply the DP algorithm. For simplicity, however, we write directly the DP'equation, which takes the intuitively clear form
optimal cost from $\boldsymbol{i}$ to $t$ in $\mathrm{N} \quad k$ moves

$$
=\min _{j=1, \ldots, N}[a i j+(\text { optimal cost from } j \text { tot in } N \quad k-1 \text { moves })]
$$

or

$$
J_{k}(i)=\min _{j=1, \ldots, N}\left[a_{i j}+J_{k+1}(j)\right], \quad k \quad 0,1, \ldots, N--2,
$$

with

$$
J_{N-1}(i)=a_{i t}, \quad i=1,2, . ., N
$$

'1'he optimal policy when at node $\boldsymbol{i}$ after $k$ moves is to move to a node $j^{*}$ that minimizes $a i j+J_{k}+l(j)$ over all $\}=1, \ldots, N$. If the optimal path obtainecl from the algorithm contains degenerate moves from a node to itself, this simply means that the path involves in reality less than N moves.

Note that if for some $k>0$, we have $J k$ (i' $)=J_{k+1}(i)$ for all i, then subsequent DP iterations will not change the values of the cost-to-go
$\left[J_{k-m}(i)=J_{k}(i)\right.$ for all $\mathrm{m}>0$ and $\left.i\right]$, so the algorithm can be terminated with $J k(i)$ being the shortest distance from i to $t$, for all i .

To demonstrate the algorithm, consider the problem shown in Fig. 2.1.2(a) where the costs $a_{i j}$ with $i \neq j$ are shown along the connecting line segments (we assume a'ij $=a j^{\prime} i$ )' Figure 2.1.2(b) shows the cost-to-go $J k(i)$ at each $i$ and $k$ together with the optimal paths.

(a)

(b)

Figure 2.1.2 (a) Shortest path problem data. The destination is node 5. Arc lengths are equal in both directions and are shown along the line segments connecting nodes. (b) Costs-to-go generated by the DP algorithm. The number along stage $k$ and state i is $J k(i)$. Arrows indicate the optimal moves at each stage and node. The optimal paths are

$$
1 \rightarrow 5, \quad 2 \rightarrow 3 \rightarrow 4 \rightarrow 5, \quad 3 \rightarrow 4 \rightarrow 5, \quad 4 \rightarrow 5 .
$$

### 2.2 SOlVIE SHORTEST PATH APPLICATIONS

The shortest path problem appears in many diverse contexts. We provide some examples.

### 2.2.1 Critical Path Analysis

Consider the planning of a project involving several activities, some of which must be completed before others can begin. The duration of each activity is known in advance. We want to find the time required to complete the project, as well as the critical activities, those that even if slightly delayed will result in a corresponding delay of completion of the overall project.

The problem can be represented by a graph with nodes $1, \ldots, N$ such as the one shown in Fig. 2.2.1. Here nodes represent completion of some
phase of the project. An arc ( $i, j$ ) represents an activity that starts once phase i is completed and has known duration $t^{\prime} i j>0$. A phase (node) $j$ is completed when all activities or arcs $(i, j)$ that are incoming to $j$ are completed. The special nodes 1 and $N$ represent the start and end of the project. Node 1 has no incoming arcs, while node $N$ has no outgoing arcs. Furthermore, there is at least one path from node 1 to every other node.


Figure 2.2.1 Graph of an activity network. Arcs represent activities and are labeled by the corresponding duration. Nodes represent completion of some phase of the project. A phase is completed if all activities associated with incoming arcs at the corresponding node are completed. The project is completed when all phases are completed. The project duration time is the length of the longest path from node 1 to node 5 , which is shown with thick line.

An important characteristic of an activity network is that it is acyclic; that is, it has no cycles. This is inherent in the problem formulation and the interpretation of nodes as phase completions.

For any path $p=\left\{(1, j 1),\left(j_{1}, j_{2},\right), \ldots,(j k, i)\right\}$ from node 1 to a node i, let $D_{p}$ be the duration of the path defined as the sum of durations of its activities; that is,

$$
D_{p}=t_{1 j_{1}}+t_{j_{1} j_{2}}+\ldots+t j k^{\prime} i
$$

Then the time $T_{i}$ required to complete phase i is

$$
T_{i}=\max _{\substack{\text { paths } \mathrm{p} \\ \text { from } 1 \text { to } i}} D_{p}
$$

Thus to find $T_{i}$, we should find the longest path from 1 to ,i. This problem may also be viewed as a shortest path problem with the length of each $\operatorname{arc}(i, j)$ being $-t_{i j}$. In particular, finding the duration of the project is equivalent to finding the shortest path from 1 to $N$. This path is also calleel a critical path. It can be seen that a delay by a given amount in the completion of one of the activities on the critical path will delay the completion of the overall project by the same amount. Note that because the network is acyclic, there can be only a finite number of paths from 1 to any $i$, so that at least one of these paths corresponds to the maximal path duration Ti.

Let us denote by Sl the set of phases that do not depend on completion of a,ny other phase, and more generally, for $k=1,2, \ldots$, let $S k$ be the set

Sk $\quad\{i l$ all paths from 1 to $i$ have $k$ arcs or less $\}$,
with So $\{1\}$. The sets $S k$ can be viewed as the state spaces for the equivalent DP problem. Using maximization in place of minimization while changing the sign of the arc lengths, the DP algorithm can be written as

$$
T_{i}=\max _{\substack{(i, i) \text { such that } \\ j \in S_{k-1}}}\left[t_{j i}+T_{j}\right], \quad \text { for all i E } S_{k} \text { with i } \notin S_{k-1} .
$$

Note that this is a forward algorithm; that is, it starts at the origin 1 and proceeds towards the destination $N$. An alternative backward algorithm, which starts at $N$ and proceeds towards 1, is also possible, as discussed in the preceding section.

As an example, for the activity network of Fig. 2.2.1, we have

$$
\begin{gathered}
\text { So } \quad\{\mathrm{I}\}, \quad \mathrm{S} 1=\{\mathrm{I}, 2\}, \quad S 2 \quad\{\mathrm{I}, 2,3\}, \\
S 3=\{1,2,3,4\}, \quad S 4=\{1,2,3,4,5\} .
\end{gathered}
$$

A calculation using the preceding formula yields

$$
T_{1}=0, T_{2}=3, T_{3}=4, T_{4}=6, T_{5}=10
$$

'1'he critical path is $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5$.

### 2.2.2 Hidden Markov Models and the Viterbi Algorithm

Consider a Markov chain with a finite number of states and given state transition probabilities Pij. Suppose that when a transition occurs, the states corresponding to the transition are unknown (or "hidden") to us, but instead we obtain an observation that relates to that transition. Given a sequence of observations, we want to estimate in some optimal sense the sequence of corresponding transitions. We are given the probability $r(z ; i, j)$ of an observation taking value $z$ when the state transition is from $i$ to $j$. We assume independent observations; that is, an observation depends only on its corresponding transition and not on other transitions. We are also given the probability $\pi_{i}$ that the initial state takes value i. The probabilities $P^{\prime} i j$ a,nd $r(z ; \mathrm{i}, j)$ are assumed to be independent of time for notational convenience. The methodology to be described admits a straightforward extension to the case of time-varying system and observation probabilities.

Markov chains whose state transitions are imperfectly observed according to the probabilistic mechanism just described are called Hidden Markov lYlodels (HlVINIs for short) or partially observable Markov chains.

In Chapter 5 we will discuss the control of such Markov chains in the context of stochastic optimal control problems with imperfect state information. In the present section, we will focus on the problem of estimating the state sequence given a corresponding observation sequence. This is an important problem that arises in a broad variety of practical contexts.

We use a "most likely state" estimation criterion, whereby given the observation sequence $Z N=\{Z l, Z 2, \ldots, Z N\}$, we adopt as our estimate the state transition sequence $\hat{X} N=\left\{\hat{x}_{0}, \hat{x}_{1}, \ldots, \hat{x} N\right\}$ that maximizes over all $X N=\{x o, X l, \cdots, X N\}$ the conditional probability $p(X N \mid Z N)$. We will show that $\hat{X}_{N}$ can be found by solving a special type of shortest path problem that involves an acyclic graph.

We have

$$
p\left(_{N}^{X} Z_{N}\right)=\frac{p\left(X^{N}, Z N\right)}{p(Z N)}
$$

where $p(X N, Z N)$ and $p(Z N)$ are the unconditional probabilities of occurrence of ( $X N, Z N$ ) and $Z N$, respectively. Since $p(Z N)$ is a positive constant once $Z N$ is known, we can maximize $p\left(X_{N}, Z N\right)$ in place of $P(X N \mid Z N)$. The probability $P(X N, Z N)$ can be written as

$$
\begin{aligned}
p(X N, Z N) & =p\left(x_{0}, x_{1}, \quad, X N, Z l, Z 2, \ldots, Z N\right) \\
& =\pi_{x_{0}} p\left(x_{1}, \quad, X N, Z l, Z 2, \quad, Z N \mid x_{0}\right) \\
& =1 f_{x_{0}} p\left(X l, Z l l x_{0}\right) p\left(x_{2}, \quad, X N, Z 2, \ldots, Z N \mid x_{0}, x_{1}, z_{1}\right) \\
& =\pi_{x_{0}} p_{x_{0} x_{1}} r\left(z_{1} ; x_{0}, x_{1}\right) p\left(x_{2}, \ldots, x_{N}, z_{2}, \ldots, Z N \mid x_{0}, x_{1}, z_{1}\right) .
\end{aligned}
$$

This calculation can be continued by writing

$$
\begin{aligned}
p(X 2, & \left.\ldots, X N, Z 2, \ldots, Z N \mid \mathrm{Xo}, x_{1}, z_{1}\right) \\
& =p\left(X 2, Z 21 x_{0}, x_{1}, z_{1}\right) p\left(x_{3}, \quad, x_{N}, z_{3}, \ldots, Z N \mid x_{0}, x_{1}, z_{1}, x_{2}, z_{2}\right) \\
& =P \operatorname{PXX} 2 T\left(Z 2 ; x_{1}, X 2\right) p\left(X 3, \quad, X N, Z 3, \cdots, Z N \mid x_{0}, x_{1}, Z l, x_{2}, Z 2\right),
\end{aligned}
$$

where for the last equation, we used the independence of the observations, Le., $p\left(Z 2 \mid X o, x_{1}, X 2, z_{1}\right)=T\left(Z 2 ; X l, x_{2}\right)$. Combining the above two relations,

$$
\begin{aligned}
p(X N, Z N) \quad & \pi X O P X O X I r\left(z l ; x_{0}, X 1\right) P X I X 2 r\left(Z 2 ; x_{1}, x_{2}\right) \\
& \cdot p\left(x_{3}, \ldots, x_{N}, z_{3}, \ldots, Z N \mid x_{0}, x_{1}, z_{1}, x_{2}, z_{2}\right)
\end{aligned}
$$

and continuing in the same manner, we obtain

$$
\begin{equation*}
p(X N, Z N)=\pi_{x_{0}} \bigcap_{k=1}^{N} P X k_{-} l X k r\left(Z k ; x_{k-1}, x_{k}\right) . \tag{2.5}
\end{equation*}
$$

We now show how the maximization of the above expression can be viewed as a shortest path problem. In particular, we construct a graph of state-time pairs, called the trellis diagram, by concatenating $N+1$ copies
of the state space, and by preceding and following them with dummy nodes $s$ and $t$, respectively, as shown in Fig. 2.2.2. The nodes of the kth copy correspond to the states $X k-l$ at time $k-1$. An arc connects a node $X k-l$ of the $k$ th copy with a node $X k$ of the $(k+1)$ st copy if the corresponding transition probability $P X K-I X k$ is positive. Since maximizing a positive cost function is equivalent to maximizing its logarithm, we see from Eq. (2.5) tllat, given the observation sequence $Z N=\{Z 1, Z 2, \ldots, Z N\}$, the problem of maximizing $p(X N, Z N)$ is equivalent to the problem

$$
\operatorname{minimize}-\operatorname{In}\left(\pi_{x_{0}}\right) \quad \sum_{\mathrm{k}=1}^{N} \operatorname{In}\left(P X k_{-} I X k r\left(z_{k} ; x_{k-1}, X k\right)\right)
$$

over all possible sequences $\{X o, \mathrm{x} 1, \ldots, \mathrm{XN}\}$.

By assigning to an arc $(s, x o)$ the length $-\operatorname{In}(\pi x Q)$, to an arc $(x N, \mathrm{t})$ the length 0 , and to an arc $(X k-l, X k)$ the length $-\operatorname{In}\left(P X k_{-} I X k r\left(z_{k} ; X k-l, X k\right)\right)$, we see that the above minimization problem is equivalent to the problem of finding the shortest path from $s$ to $t$ in the trellis diagram. This shortest path defines the estimated state sequence $\left\{\hat{x}_{0}, \hat{x}_{1}, \ldots, \hat{x} N\right\}$.


Figure 2.2.2 State estimation of an HMM viewed as a problem of finding a shortest path from sto $t$. Length of arcs from s to states $X_{0}$ is $\operatorname{In}(\pi x o)$, and length of arcs from states $X N$ to $t$ is zero. Length of an arc from a state $X k-I$ to $X k$ is $-\operatorname{In}\left(P X k \_I X k T(Z k ; X k-I, X k)\right)$, where $Z k$ is the kth observation.

In practice, the shortest path is most conveniently constructed sequentially by forward DP, that is, by first calculating the shortest distance from $s$ to each node X 1 , then using these distances to calculate the shortest distance from $s$ to each node $x 2$, etc. In particular, suppose that we have computed the shortest distances $D k(X k)$ from $s$ to all states $X k$ on the basis of the observation sequence $Z k$, and suppose that the new observation $Z k+l$ is obtained. Then the shortest distances $D k+l(X k+l)$ from $s$ to any state
$x_{k+1}$ can be computed by the DP recursion

$$
D_{k+1}\left(x_{k+1}\right) \quad \min _{\substack{\text { all } x_{k} \text { such that } \\ p_{x_{k}} x_{k+1}>0}}\left[D_{k}\left(x_{k}\right)-\ln \left(p_{x_{k}: x_{k+1}} r\left(z_{k+1} ; x_{k}, x_{k+1}\right)\right)\right] .
$$

The initial condition is $D o(x o)=-\operatorname{In}\left(1 \mathrm{f}_{\mathrm{xo}}\right)^{\prime}$ The final estimated state sequence $\hat{X}_{N}$ corresponds to the shortest path from $s$ to the final state $\hat{x} N$ that minimizes $D N(x N)$ over the finite set of possible states $x N$. An advantage of this procedure is that it can be executed in real time, as soon as each new observation is obtained.

There are a number of practical schemes for estimating a portion of the state sequence without waiting to receive the entire observation sequence $Z N$, and this is useful if $Z N$ is a long sequence. For example, one can check fairly easily whether for some $k$, all shortest paths from $s$ to states $X k$ pass through a single node in the subgraph of states $x o, \ldots, x_{k-1}$. If so, it can be seen from Fig. 2.2.3 that the shortest path from $s$ to that node will not be affected by reception of additional observations, and therefore the subsequence of state estimates up to that node can be determined without waiting for the remaining observations.


Figure 2.2.3 Estimating a portion of the state sequence prior to receiving the entire observation sequence. Suppose that the shortest paths from $s$ to all states $x k$ pass through a single node m . If an additional observation is received, the shortest paths from s to all states $x k+l$ will continue to pass through m . Therefore, the portion of the state sequence up to node $m$ can be safely estimated because additional observations will not change the initial portion of the shortest paths from s up to $m$.

The shortest path-based estimation procedure just described is known as the Viterbi algorithm, and finds numerous applications in a variety of contexts. An example is speech recognition, where the basic goal is to transcribe a spoken word sequence in terms of elementary speech units called phonemes. One possibility is to associate the states of the HHM with
phonemes, and given a sequence of recorded phonemes $Z N=\left\{Z l^{\prime} \ldots, Z N\right\}$, to find a phonemic sequence $\hat{X}_{N}=\left\{\hat{x}_{1}, \ldots, \hat{x}_{N}\right\}$ that maximizes over all $\mathrm{X}_{\mathrm{N}} \quad\left\{x_{1}, \ldots, x_{N}\right\}$ the conditional probability $p(X N \mid Z N)$. The probabilities $T\left(Z k ; x_{k-1}, x_{k}\right)$ and $P X k-I X k$ can be experimentally obtained, if necessary by specialized "training" for each speaker that uses the speech recognition system. The Viterbi algorithm can then be used to find the most likely phonemic sequence. There are also other HMMs used for word and sentence recognition, where only phonemic sequences that constitute words from a given dictionary are considered. We refer the reader to Rabiner [Rab89] and Picone [Pic90] for a general review of HMMs applied to speech recognition and for further references to related work. It is also possible to use similar models for computerized recognition of handwriting.

The Viterbi algorithm was originally developed as a scheme for decoding data after transmission over a noisy communication channel. The following example describes this context in some detail.

## นленияュン 2.2.1 (Convolutional Coding and Decoding)

When binary data are transmitted over a noisy communication channel, it is often essential to use coding as a means of enhancing reliability of communication. A common type of coding method, called convolutional coding, converts a source-generated binary data sequence
$\left\{\mathrm{WI}, w_{2}, \ldots\right\}, \quad W k E\{0,1\}, \quad k=1,2, \ldots$,
into a coded sequence $\{Y I, Y 2, \ldots\}$, where each $Y k$ is an $n$-dimensional vector with binary coordinates, called codewoTd,

$$
y_{k}=\left(\begin{array}{c}
y_{k}^{\frac{1}{2}} \\
\vdots \\
y_{k}^{n}
\end{array}\right), \quad y_{k}^{i} \mathrm{E}\{0,1\}, \quad \mathrm{i}=1, \ldots, n, k=1,2, \ldots
$$

The sequence $\{Y I, Y 2, \ldots\}$ is then transmitted over a noisy channel and gets transformed into a sequence $\{\mathrm{ZI}, \mathrm{z2}, \ldots$,$\} , which is then decoded to yield the$ decoded data sequence $\left\{\hat{w}_{1}, \hat{w}_{2}, \ldots\right\}$; see Fig. 2.2.4. The objective is to design the encoder/decoder scheme so that the decoded sequence is as close to the original as possible.


The problem just discussed is central in information theory and can be approached in several different ways. In a particularly popular and eflective technique called convolutional coding, the vectors $Y k$ are related to $W k$ via equations of the form

$$
\begin{gather*}
y_{k}=C x_{k-1}+d w_{k}, \quad k=1,2, \ldots,  \tag{2.6}\\
x_{k}=A x_{k-1}+b w_{k}, \quad k=1,2, \ldots, \quad x_{0}: \text { given }, \tag{2.7}
\end{gather*}
$$

where $X k$ is an m-dimensional vector with binary coordinates, which we view as state, and $\mathrm{C}, \mathrm{d}, A$, and bare $n \times \mathrm{m}, n \times 1, \mathrm{~m} \times m$, and $m \times 1$ matrices, respectively, with binary coordinates. The products and the sums involved in the expressions $C X k-I+d W k$ and $A x_{k-1}+b W k$ are calculated using modulo 2 arithmetic.

As an example, let $\mathrm{m}=2, n=3$, and

$$
C=\left(\begin{array}{ll}
1 & 0 \\
0 & 1 \\
0 & 1
\end{array}\right), \quad A \quad\left(\begin{array}{ll}
0 & 1 \\
1 & 1
\end{array}\right), \quad d=\left(\begin{array}{l}
1 \\
1 \\
1
\end{array}\right), \quad b=\binom{0}{1}
$$

Then the evolution of the system (2.6)-(2.7) can be representecl by the diagram shown in Fig. 2.2.5. Given the initial Xo, this diagram can be used to generate the codeword sequence $\{\mathrm{YI}, Y 2, \ldots\}$ corresponding to a data sequence $\left\{\mathrm{WI}, w_{2}, \ldots.\right\}$. For example, when the initial state is $x_{0} 00$, the clata sequence

$$
\{W I, W 2, W 3, W 4\}=\{1,0,0,1\}
$$

generates the state sequence

$$
\{X O, X I, X 2, X 3, X 4\}=\{00,01,11,10,00\},
$$

and the codeword sequence

$$
\left\{y_{1}, y_{2}, y_{3}, y_{4}\right\}=\{111,011,111,011\}
$$



Figure 2.2.5 State transition diagram for convolutional coding. The binary number pair on each arc is the data/codeword pair $W k / V k$ for the corresponding transition. So for example, when $X k-l \quad 01$, a zero data bit ( $W k=0$ ) effects a transition to $X k=11$ and generates the codeword 001 .

Assume now that the characteristics of the noisy transmission channel are such that a codeword $Y$ is actually received as $z$ with known probability $p(z \mid V)$, where $z$ is any n-bit binary number. We assume independent errors so that

$$
\begin{equation*}
p(Z N \mid Y N)=\prod_{k=l}^{N} p(Z k \mid Y k) \tag{2.8}
\end{equation*}
$$

where $Z N=\{Z 1, \ldots, Z N\}$ is the received sequence and $Y N=\left\{Y I, \ldots, y_{N}\right\}$ is the transmitted sequence. By associating the codewords $Y$ with state transitions, we formulate a maximum likelihood estimation problem, whereby we want to find a sequence $\hat{Y} N=\left\{\hat{y}_{1}, \hat{y}_{2}, \ldots, \hat{y} N\right\}$ such that

$$
p\left(Z_{N} \mid \dot{Y}_{N}\right)=\max _{Y_{N}} p\left(Z_{N} \mid Y_{N}\right)
$$

The constraint on $Y N$ is that it must be a feasible codeword sequence (Le., it must correspond to some initial state and data sequence, or equivalently, to a sequence of arcs of the trellis diagram).

Let us construct a trellis diagram by concatenating N state transition diagrams and appending dummy nodes sand $t$ on its left and right sides, which are connected with zero-length arcs to the states $x Q$ and $X N$, respectively. By using Eq. (2.8), we see that, given the received sequence $Z N=\left\{Z l^{\prime} Z 2, \ldots, Z N\right\}$, the problem of maximizing $p(Z N \mid Y N)$ is equivalent to the problem

$$
\begin{aligned}
& \operatorname{minimize} \\
& \text { over all binary sequences }\{Y \mid, Y 2, \ldots, Y N\} \text {. }
\end{aligned}
$$

This is equivalent to a problem of finding a shortest path in the trellis diagram from $s$ to $t$, where the length of the arc associated with the codeword $Y k$ is $-\operatorname{In}(p(z k \mid Y k))$, and the lengths of each arc incident to a dummy node is zero. From the shortest path and the trellis diagram, we can then obtain the corresponding data sequence $\left\{\hat{w}_{1}, \ldots, \hat{\boldsymbol{w}} N\right\}$, which is accepted as the decoded data.

The maximum likelihood estimate $\hat{Y}_{N}$ can be found by solving the corresponding shortest path problem using the Viterbi algorithm. In particular, the shortest distances $D_{k+1}(X k+l)$ from $s$ to any state $X k+l$ are computed by the DP recursion

$$
\left.D_{k+1}\left(x_{k+1}\right)=\min _{\substack{\text { all } x k \text { such that } \\(x k, x k+l) \text { is an arc }}} f D k(X k)-\operatorname{In}\left(p\left(z k+l \backslash y_{k+1}\right)\right)\right]
$$

where $Y k+l$ is the codeword corresponding to the $\operatorname{arc}\left(X k^{\prime} x_{k+1}\right)$. The final state $\hat{x} N$ on the shortest path is the one that minimizes $\mathrm{D} N(x N)$ over $x N$.

### 2.3 SHORTEST PATH ALGORITHMS

We have seen that shortest path problems and deterministic finite-state optimal control problems are equivalent. The computational implications of this are twofold.
(a) One can use DP to solve general shortest path problems. Note that there are several other shortest path methods, some of which have superior theoretical worst-case performance to DP. However, DP is often preferred in practice, particularly for problems with an acyclic graph structure and also when a parallel computer is available.
(b) One can use general shortest path methods (other than DP) for deterministic finite-state optimal control problems. In most cases, DP is preferable to other shortest path methods, because it is tailored to the sequential nature of optimal control problems. However, there are important cases where other shortest path methods are preferable.
In this section we discuss several alternative shortest path methods. We motivate these methods by focusing on shortest path problems with a very large number of nodes. Suppose that there is only one origin and only one destination, as in shortest path problems arising from deterministic optimal control (cf. Fig. 2.1.1). Then it is often true that most of the nodes are not relevant to the shortest path problem in the sense that they are unlikely candidates for inclusion in a shortest path between the given origin and destination. Unfortunately, however, in the DP algorithm every node and arc will participate in the computation, so there arises the possibility of other more efficient methods.

A similar situation arises in some search problems that are common in artificial intelligence and combiriatorial optimization. Generally, these problems involve decisions that can be broken down into stages. With proper reformulation, the decision stages can be made to correspond to arc selections in a shortest path problem, or to the stages of a DP algorithm. We provide some examples.

## Example 2.3.1 (The Four Queens Problem)

Four queens must be placed on a $4 \times 4$ portion of a chessboard so that no queen can attack another. In other words, the placement must be such that every row, column, or diagonal of the $4 \times 4$ board contains at most one queen. Equivalently, we can view the problem as a sequence of problems; first, placing a queen in one of the first two squares in the top row, then placing another queen in the second row so that it is not attacked by the first, and similarly placing the third and fourth queens. (It is sufficient to consider only the first two squares of the top row, since the other two squares lead to symmetric positions.) We can associate positions with nodes of an acyclic graph where the root node $s$ corresponds to the position with no queens and the terminal nodes correspond to the positions where no additional queens can be placed
without some queen attacking another. Let us connect each terminal position with an artificial node $t$ by means of an arc. Let us also assign to all arcs length zero except for the artificial arcs connecting terminal positions with less than four queens with the artificial node $t$. These latter arcs are assigned the length oo (see Fig. 2.3.1) to express the fact that they correspond to deadend positions that cannot lead to a solution. Then, the four queens problem reduces to finding a shortest path from node $s$ to node $t$.

Note that once the nodes ofthe graph are enumerated the problem is essentially solved. In this $4 \times 4$ problem the number of nodes is small. However, we can think of similar problems with much larger memory requirements. For example, there is an eight queens problem where the board is $8 \times 8$ instead of 4 x 4 .

## Exarnple 2.3.2 (The Traveling Salesman Problem)

An important model for scheduling a sequence of operations is the classical traveling salesman problem. Here we are given $N$ cities and the mileage between each pair of cities. We wish to find a minimum-mileage trip that visits each of the cities exactly once and returns to the origin node. To convert this problem to a shortest path problem, we associate a node with every sequence of $n$ distinct cities, where $n \leq \mathrm{N}$. The construction and arc lengths of the corresponding graph are explained by means of an example in Fig. 2.3.2. The origin node $s$ consists of city A, taken as the start. A sequence of $n$ cities $(n<\mathrm{N})$ yields a sequence of $(n+1)$ cities by adding a new city. Two such sequences are connected by an arc with length equal to the mileage between the last two of the $n+1$ cities. Each sequence of $N$ cities is connected to an artificial terminal node $t$ with an arc having length equal to the distance from the last city of the sequence to the starting city A. Note that the number of nodes grows exponentially with the number of cities.

In the shortest path problem that we will consider in this section, there is a special node $s$, called the origin, and a special node $t$, called the destinat'lon. We will assume a single destination, but the methods to be discussed admit extensions to the case of multiple destinations (see Exercise 2.6). A node $j$ is called a child of node $i$ if there is an arc $(i, j)$ connecting i with $j$. The length of $\operatorname{arc}(i, j)$ is denoted by aij and we assume that all arcs have nonnegative length. Exercise 2.7 deals with the case where all cycle lengths (rather than arc lengths) are assumed nonnegative. We wish to find a shortest path from origin to destination.

### 2.3.1 Label Correcting Methods

We now discuss a general type of shortest path algorithm. The idea is to progressively discover shorter paths from the origin to every other node i, and to maintain the length of the shortest path found so far in a variable $d_{i}$ called the label of i. Each time $d_{i}$ is reduced following the discovery of a


Figure 2.3.1 Shortest path formulation of the four queens problern. Symmetric positions resulting from placing a queen in one of the rightmost squares in the top row have been ignored. Squares containing a queen have been darkened. All arcs have length zero except for those connecting dead-end positions to the artificial terminal node.
shorter path to i , the algorithm checks to see if the labels $d j$ of the children $j$ of $i$ can be "corrected," that is, they can be reduced by setting them to


Artificial Terminal Node t


Figure 2.3.2 Example of a shortest path formulation of the traveling salesman problem. The distance between the four cities A, B, C, and D are shown in the table. The arc lengths are shown next to the arcs.
$d_{i}+a_{i j}$ [the length of the shortest path to $i$ found thus far followed by arc $(i, j)]$. The label $d_{t}$ of the destination is maintained in a variable called UPPER, which plays a special role in the algorithm. The label $d_{s}$ of the origin is initialized at 0 and remains at 0 throughout the algorithm. The labels of all other nodes are initialized at oo, i.e., $d_{i}=$ oo for all $\mathrm{i} \neq s$.

The algorithm also makes use of a list of nodes called OPEN (another name frequently used is candidate list). The list OPEN contains nodes that are currently active in the sense that they are candidates for further examination by the algorithm and possible inclusion in the shortest path. Initially, OPEN contains just the origin node $s$. Each node that has entered OPEN at least once, except $s$, is assigned a "parent," which is some other node. The parent nodes are not necessary for the computation of the shortest distance; they are needed for tracing the shortest path to the origin after the algorithm terminates. The steps of the algorithm are as follows (see also Fig. 2.3.3):

## Label Correcting Algorithm

Step 1: Remove a node i from OPEN and for each child. 7 of i, execute step 2.
Step 2: If $d i+a i j<\min \{d j$, UPPER $\}$, set $d: i=d_{i}+$ aij and set ito be the parent of $j$. In addition, if $j \neq t$, place $j$ in OPEN if it is not already in OPEN, while if $j=t$, set UPPER to the new value $d_{i}+$ ait of $d_{t}$.

Step 3: If OPEN is empty, terminate; else go to step 1.


Figure 2.3.3 Diagrammatic illustration of the label correcting algorithm, with an interpretation of the tests for insertion of a node into the OPEN list.

It can be seen by induction that, throughout the algorithm, $d_{j}$ is either oo (if node $j$ has not yet entered the OPEN list), or else it is the length of some path from $s$ to $j$ consisting of nodes that have entered the OPEN list at least once. In the latter case, the path can be constructed by tracing backward the parent nodes starting with the parent of node $j$. Furthermore, UPPER is either oo, or else it is the length of some path from $s$ to $t$, and consequently it is an upper bound of the shortest distance from $s$ to $t$. The idea in the algorithm is that when a path from $s$ to $j$ is discovered, which is shorter than those considered earlier (eli + aij $<d j$ in step 2), the value of $d_{j}$ is accordingly reduced, and node $j$ enters the OPEN list so that paths passing through $j$ and reaching the children of $j$ can be taken into account. It makes sense to do so, however, only when the path considered has a chance of leading to a path from $s$ to $t$ with length smaller than the upper bound UPPER of the shortest distance from $s$ to $t$.

In view of the nonnegativity of the arc lengths, this is possible only if the path length $d_{i}+a_{i j}$ is smaller than UPPER. This provides the rationale for entering $j$ into OPEN in step 2 only if $d_{i}+^{-}$aij < UPPER (see Fig. 2.3.3).

Tracing the steps of the algorithm, we see that it will first remove node $s$ from OPEN and sequentially examine its children. If $t$ is not a child of $s$, the algorithm will place all children j of $s$ in OPEN after setting $d_{j} \quad$ asj. If $t$ is a child of $s$, then the algorithm will place all children j of $s$ examined before $t$ in OPEN and will set their labels to asj; then it will examine $t$ and set UPPER to ast; finally, it will place each of the remaining children j of $s$ in OPEN only if asj is less than the current value of UPPER, which is CLst. The algorithm will subsequently take a child $\mathrm{i} \neq t$ of $s$ from OPEN, and sequentially place in OPEN those of its children $\mathrm{j} \neq t$ that satisfy the criterion of step 2 , etc. Note that the origin 8 can never reenter OPEN because $d_{s}$ cannot be reduced from its initial value of zero. Also, by the rules of the algorithm, the destination can never enter OPEN. When the algorithrn terminates, we will show shortly that a shortest path can be obtained by tracing backwards the parent nodes starting from $t$ and going towards $s$. Figure 2.3.4 illustrates the use of the algorithm to solve the traveling salesman problem of Fig. 2.3.2.

The following proposition establishes the validity of the algorithm.

Proposition 2.3.1: If there exists at least one path from the origin to the destination, the label correcting algorithm terminates with UPPER equal to the shortest distance from the origin to the destination. Otherwise the algorithm terminates with UPPER $=00$.

Proof: We first show that the algorithm will terminate. Indeed, each time a node j enters the OPEN list, its label is decreased and becomes equal to the length of some path from 8 to $\mathbf{j}$. On the other hand, the number of distinct lengths of paths from 8 to $j$ that are smaller than any given number is finite. The reason is that each path can be decomposed into a path with no repeated nodes (there is a finite number of distinct such paths), plus a (possibly empty) set of cycles, each having a nonnegative length. 'I'herefore, there can be only a finite number of label reductions, implying that the algorithm will terminate.

Suppose that there is no path from $s$ to $t$. Then a node i such that $(i, t)$ is an arc cannot enter the OPEN list, because as argued earlier, this would establish that there is a path from 8 to i , and therefore also a path from 8 to $t$. Thus, based on the rules of the algorithm, UPPER can never be reduced from its initial value of oo.

Suppose now that there is a path from 8 to $t$. Then, since there is a finite number of distinct lengths of paths from 8 to $t$ that are smaller than any given number, there is also a shortest path. Let $(8, j 1, j 2, \ldots, j k, t)$


Artificial Terminal Node t

| Iter. No. | Node Exiting OPEN | OPEN at the End of Iteration | UPPER |
| :---: | :---: | :---: | :---: |
| 0 | - | 1 | Oo |
| 1 | 1 | $2,7,10$ | 00 |
| 2 | 2 | $3,5,7,10$ | 00 |
| 3 | 3 | $4,5,7,10$ | 00 |
| 4 | 4 | $5,7,10$ | 43 |
| 5 | 5 | $6,7,10$ | 43 |
| 6 | 6 | 7,10 | 13 |
| 7 | 7 | 8,10 | 13 |
| 8 | 8 | 9,10 | 13 |
| 10 | 9 | 10 | 13 |

Figure 2.3.4 The algorithm applied to the traveling salesman problem of Fig. 2.3.2. The optimal solution ABDC is found after examining nodes 1 through 10 in the figure in that order. The table shows the successive contents of the OPEN list.
be a shortest path and let $d^{*}$ be the corresponding shortest distance. We will show that the value of UPPER upon termination must be equal to $d^{*}$. Indeed, each subpath $(8, j 1, \ldots, j m), \mathrm{m}=1, \ldots, k$, of the shortest path $\left(s, j_{1}, j_{2}, \ldots, j k, t\right)$ must be a shortest path from 8 to $j_{m}$. If the value of

UPPER is larger than $d^{*}$ at termination, the same must be true throughout the algorithrn, and therefore UPPER will also be larger than the length of all the paths $\left(8, j 1, .{ }^{\prime}, j m\right), \mathrm{m}=1, \ldots, k$, throughout the algorithm, in view of the nonnegative arc length assumption. It follows that node $j k$ will never enter the OPEN list with djk equal to the shortest distance from 8 to $j k$, since in this case UPPER would be set to $d^{*}$ in step 2 immediately following the next time node $j k$ is examined by the algorithm in step 2. Similarly, and using also the nonnegative length assumption, this means that node $j k-1$ will never enter the OPEN list with djk_/ equal to the shortest distance from 8 to $j k-1$. Proceeding backward, we conclude that 11 never enters the OPEN list with $d_{j 1}$ equal to the shortest distance from 8 to $j_{1}$ [which is equal to the length of the arc $\left.(8, j 1)\right]$. This happens, however, at the first iteration of the algorithm, obtaining a contradiction. It follows that at termination, UPPER will be equal to the shortest distance from 8 to $t$ Q.E.D.

From the preceding proof, it can also be seen that, upon termination of the algorithm, the path constructed by tracing the parent nodes backward from $t$ to 8 has length equal to UPPER, so it is a shortest path from 8 to $t$. Thus the algorithm yields not just the shortest distance but also a shortest path, provided that we keep track of the parent of each node that enters OPEN.

An important property of the algorithm is that nodes $j$ for which $d i+$ $a_{i j} \geq$ UPPER in step 2 will not enter OPEN in the current iteration, and may possibly not enter in any subsequent iteration. As a result the number of nodes that enter OPEN may be much smaller than the total number of nodes. Furthermore, if a good lower bound to the shortest distance from 8 to $t$ (or the shortest distance itself) is known, the computation can be terminated once UPPER reaches that bound within an acceptable tolerance. This is useful, for example, in the four queens problem, where the shortest distance is known to be zero or infinity. Then the algorithm will terminate once UPPER $=0$, when a solution is found for the first time.

## Specific Label Correcting Methods

There is considerable freedom in selecting the node to be removed from OPEN at each iteration. This gives rise to several different methods. The following are some of the most important (the author's textbooks on network optimization [Ber91a], [Ber98a] contain a fuller account of label correcting methods and their analysis; [Ber91a] contains several computer code implementations).
(a) Breadth-first search, also known as the Bellman-Ford method, which adopts a first-in/first-out policy; that is, the node is always removed from the top of OPEN and each node entering OPEN is placed at the
bottom of OPEN. [Here and in the following methods except (c), we assume that OPEN is structured as a queue.]
(b) Depth-first search, which adopts a last-in/fi.rst-out policy; that is, the node is always removed from the top of OPEN and each node entering OPEN is placed at the top of OPEN. One motivation for this method is that it often requires relatively little memory. For example, suppose that the graph has a tree-like structure whereby there is a unique path from the origin node to every node other than the destination as shown in Fig. 2.3.5. Then the nodes will enter OPEN only once and in the order shown in Fig. 2.3.5. At anyone time, it is only necessary to store a small portion of the graph as shown in Fig. 2.3.6.

Origin Node s
1


Destination Node t


Figure 2.3.5 Searching a tree in depth-first fashion. The numbers next to the nodes indicate the order in which nodes exit the OPEN list.

Figure 2.3.6 Memory requirements of depthfirst search for the graph of Fig. 2.3.5. At the time the node marked by the checkmark exits the OPEN list, only the solicl-line portion of the tree is needed in memory. The dotted-line portion has been generated and purged from memory based on the rule that for a graph where there is only one path from the origin to every node other than the destination, it is unnecessary to store a node once all of its successors are out of the OPEN list. The brokenline portion of the tree has not yet been generated.
(c) Best-first search, which at each iteration removes from OPEN a node with minimum value of label, Le., a node $i$ with

$$
d_{i}=\min _{j \text { in OPEN }} d j .
$$

This method, also known as Dijkstm's method or label setting method, has a particularly interesting property. It can be shown that in this method, a node will enter the OPEN list at most once (see Exercise 2.4). The drawback of this method is the overhead required to find at each iteration the node of OPEN that has minimum label. Several sophisticated methods have been developed to carry out this operation efficiently (see e.g., Bertsekas [Ber98a]).
(d) D'Esopo-Pape method, which at each iteration removes the node that is at the top of OPEN, but inserts a node at the top of OPEN if it has already been in OPEN earlier, and inserts the node at the bottom of OPEN otherwise.
(e) Small-Label-First (SLF) method, which at each iteration removes the node that is at the top of OPEN, but inserts a node $i$ at the top of OPEN if its label $d_{i}$ is less than or equal to the label $d j$ of the top node $j$ of OPEN; otherwise it inserts $i$ at the bottom of OPEN. This is a low-overhead approximation to the best-first search strategy. As a complement to the SLF strategy, one may also try to avoid removing nodes with relatively large labels from OPEN using the following device, known as the Large-Label-Last (LLL) stmtegy: at the start of an iteration, the top node of OPEN is compared with the average of the labels of the nodes in OPEN and if it is larger, the top node is placed at the bottom of OPEN and the new top node of OPEN is similarly tested against the average of the labels. In this way the removal of nodes with relatively large labels is postponed in favor of nodes with relatively small labels. The extra overhead required in this method is small: maintain the sum of the labels and the number of nodes in OPEN. When starting a new iteration, the ratio of these numbers gives the desired average. There are also several other methods, which are based on the idea of examining nodes with small labels first (see Bertsekas [Ber93], [Ber98a], and Bertsekas, Guerriero, and Musmanno [BGM96] for detailed descriptions and computational studies).
Generally, it appears that for nonnegative arc lengths, the number of iterations is reduced as the method is more successful in removing from OPEN nodes with a relatively small label. For a supporting heuristic argument, note that for a node $j$ to reenter OPEN, some node $i$ such that $d_{i}+a_{i j}<e l j$ must first exit OPEN. Thus, the smaller $d j$ was at the previous exit of.i from OPEN, the less likely it is that $d_{i}+$ aij will subsequently become less than $d j$ for some node $i$ in OPEN and arc ( $i, j$ ). In particular,
if $d j \leq$ min'i in OPEN $d_{i}$, it is impossible that subsequent to the exit of $j$ from OPEN we will have $d_{i}+a^{\prime} i j<d j$ for some i in OPEN, since the arc lengths $a^{\prime} i j$ are nonnegative. The SLF and other related but more sophisticated methods, often require a number of iterations, which is close to the minimum (the one required by the best-first method). However, they can be much faster than the best-first method, because they require much less overhead for determining the node to be removed from OPEN.

### 2.3.2 Label Correcting Variations - $A^{*}$ Algorithm

The generic label correcting algorithm need not be started with the initial conditions $d_{s} \quad 0$ and $d i \quad$ oo for $\mathrm{i} \neq s$ in order to work correctly. It can be shown, similar to Prop. 2.3.1, that one can use any set of initial labels $d i$ such that for each node $\mathrm{i}, d i$ is either oo or else it is the length of some path from $s$ to $i$. The scalar UPPER may be taken to be equal to $d_{t}$ and the initial OPEN list may be taken to be the set $\{\mathrm{i} \operatorname{l} d i<00\}$.

This kind of initialization is very useful if, by using heuristics or a known solution of a similar shortest path problem, we can construct a "good" path

$$
P=\left(s, i_{l}, \ldots, \quad t\right)
$$

from $s$ to $t$. Then we can initialize the algorithm with

$$
d_{i}= \begin{cases}\text { Length of portion of path } P \text { from } s \text { to } i & \text { if i E } P, \\ \infty & \text { if } \mathrm{i} \notin P,\end{cases}
$$

with UPPER equal to $d t$, and with the OPEN list equal to $\{s, i 1, \ldots, i k\}$. If $P$ is a near-optimal path and consequently the initial value of UPPER is near its final value, the test for future admissibility into the candidate list will be relatively tight from the start of the algorithm and many unnecessary entrances of nodes into OPEN may be saved. In particular, it can be seen that all nodes whose shortest distances from the origin are greater or equal to the length of $P$ will never enter the candidate list.

Another possibility, known as the $A^{*}$ algorithm, is to strengthen the test di + aij < UPPER that node $j$ must pass before it can be placed in the OPEN list in step 2. This can be done if a positive underestimate $h j$ of the shortest distance of node $j$ to the destination is available. Such an estimate can be obtained from special knowledge about the problem at hand. We may then speed up the computation substantially by placing a node $j$ in OPEN in step 2 only when

$$
d i+a_{i j}+h j<\text { UPPER }
$$

(instead of $d_{i}+a_{i j}<$ UPPER). In this way, fewer nodes will potentially be placed in OPEN before termination. Using the fact that $h j$ is an underestimate of the true shortest distance from $j$ to the destination, it can be seen
that nodes $j$ such that $d_{i}+$ aij $+h_{j} \geq$ UPPER need not enter OPEN, and the argument given in the proof of Prop. 2.3.1 shows that the algorithm using the preceding test will terminate with a shortest path.

The $A^{*}$ algorithm is just one way to sharpen the test $d_{i}+$ aij $<$ UPPER for admission of node $j$ into the OPEN list. An alternative idea is to try to reduce the value of UPPER by obtaining for the node $j$ in step 2 an upper bound mj of the shortest distance from $j$ to the destination t (for example the length of some path from $j$ to $t$ ). Then if $d j+m j<U P P E R$ after step 2 , we can reduce UPPER to $d j+m j$, thereby making the test for future admissibility into OPEN more stringent. This idea is used in some versions of the branch-and-bound algorithm, one of which we now describe.

### 2.8.3 Branch-and-Bound

Consider a problem of minimizing a cost function $f(x)$ over a finite set of feasible solutions $X$. We have in mind problems where the number of feasible solutions is very large, so an enumeration and comparison of these solutions is impractical, e.g., the traveling salesman problem of Example 2.3.2. The idea of the branch-and-bound method is to avoid a complete enumeration by discarding solutions that, based on certain tests, have no chance of being optimal. This is similar to label correcting methods, where based on various tests that use the value of UPPER and other data, the insertion in the OPEN list of some nodes is avoided. In fact, we will see that the branch-and-bound method can be viewed as a form of label correcting method.

The key idea of the branch-and-bound method is to partition the feasible set into smaller subsets, and then use certain bounds on the attainable cost within some of the subsets to eliminate from further consideration other subsets. The rationale for this is captured in the following simple observation.

## Bounding Principle

Given the problem of minimizing $f(x)$ over $x \mathrm{E} X$, and two subsets Y1 C $X$ and $Y 2$ e $X$, suppose that we have bounds

$$
\underline{f}_{1} \leq \min _{\mathrm{xEY1}} f(x), \quad \bar{f}_{2} \geq \min _{x \in Y_{2}} f(x)
$$

Then, if $\bar{f}_{2} \leq \quad$ the solutions in $Y /$ may be disregarded since their cost cannot be smaller than the cost of the best solution in $Y_{2}$.

The branch-ancl-bouncl method calculates suitable upper and lower bounds, and uses the bounding principle to eliminate from consideration substantial portions of the feasible set. To describe the method, we use
an acyclic graph with nodes that correspond on a one-to-one basis with a collection $\mathcal{X}$ of subsets of the feasible set $X$. We require the following:

1. $X \in \mathcal{X}$ (i.e., the set of all solutions is a node).
2. For each solution $x$, we have $\{x\}$ E $\mathcal{X}$ (i.e., each solution viewed as a singleton set is anode).
3. Each set $Y \mathrm{E} \mathcal{X}$ that contains more than one solution $x \mathrm{E} X$ is partitioned into sets $\mathrm{Y}_{1}, \ldots \cdot, Y_{n} \mathrm{E} \mathcal{X}$ such that $Y_{i} \neq \mathrm{Y}$ for all i:

$$
\bigcup_{i=1}^{n} Y_{i}=Y
$$

The set $Y$ is called the parent of $Y I, \ldots, Y_{n}$, and the sets $Y_{1}, \ldots, Y_{n}$ are called the children of Y.
4. Each set in $\mathcal{X}$ other than $X$ has at least one parent.

The collection of sets $\mathcal{X}$ defines an acyclic graph with root node the set of all feasible solutions $X$ and terminal nodes the singleton solutions $\{x\}, x \mathrm{E} X$ (see Fig. 2.3.7). The arcs of the graph are those that connect parents $Y$ and their children $Y_{i}$. Suppose that for every nonterminal node $Y$ there is an algorithm that calculates upper and lower bounds $\underline{f}_{Y}$ and $\bar{f}_{\mathrm{y}}$ for the minimum cost over Y:

$$
\underline{f} y \leq \min _{x \in Y} f(x) \leq 1 y
$$

Assume further that the upper and lower bounds are exact for each singleton solution node $\{x\}$ :

$$
\underline{f}_{\{x\}}=f(x)=\bar{f}_{\{x\}}, \quad \text { for all } x \mathrm{E} X
$$



Figure 2.3.7 An acyclic graph corresponding to a branch-and-bound algorithm. Each node (subset) except those consisting of a single solution is partitioned into several other nodes (subsets).

Define now the length of an arc involving a parent $Y$ and a child $Y_{i}$, to be the lower bound difference

$$
\underline{f}_{Y_{i}} \quad \underline{f}_{Y}
$$

Then the length of any path from the origin node $X$ to any node Y is $\underline{f}_{Y}$. Since $\underline{f}_{\{x\}}=f(x)$ for all feasible solutions $x \mathrm{E} X$, it is clear that rninimizing $f(x)$ over $x \mathrm{E} X$ is equivalent to finding a shortest path from the origin node to one of the singleton nodes $\{x\}$.

Consider now a variation of the label correcting method, where in addition we use our knowledge of the upper bounds $\bar{f}_{Y}$ to reduce the value of UPPER. Initially, OPEN contains just $X$, and UPPER equals $\bar{f} x$.

## Branch-and-Bound AIgoritluTI

Step 1: Remove a node $Y$ from OPEN. For each child $Y_{j}$ of $Y$, do the following: If $f_{Y_{j}}<$ UPPER, then place $Y_{j}$ in OPEN. If in addition $\bar{f}_{Y_{j}}<$ UPPER, then set UPPER $=\bar{f}_{Y_{j}}$, and if $Y_{j}$ consists of a single solution, mark that solution as being the best solution found so far.
Step 2:
Test) If OPEN is nonempty, go to step I.
Otherwise, tenninate; the best solution found so far is optimal.

An alternative termination step 2 for the preceding algorithm is to set a tolerance $\epsilon>0$, and check whether UPPER and the minimum lower bound $f_{Y}$ over all sets $Y$ in the OPEN list differ by less than $\epsilon$. If so, the algoritlun is terminated, and SOHle set in OPEN must contain a solution that is within $\epsilon$ of being optimal. There are a number of other variations of the algorithm. For example, if the upper bound ! $y$ at a node is actually the cost $f(x)$ of some element $x \mathrm{E} \mathrm{Y}$, then this element can be taken as the best solution found so far whenever $\bar{f}_{Y}<$ UPPER in step 2. Other variations relate to the method of selecting a node from OPEN in step I. For example, two strategies of the best-first type are to select the node with minimal lower or upper bound. Note that it is neither practical no'r necessary to generate a priori the acyclic graph of the branch-and-bound method. Instead, one may adaptively decide on the order and the manner in which the parent sets are partitioned into children sets based on the progress of the algorithm.

We finally note that to apply branch-and-bound effectively, it is important to have good algorithms for generating upper and lower bounds at each node. These bounds should be as sharp as practically possible. Then, fewer nodes will be admitted into OPEN, with attendant computational savings. Typically, continuous optimization problems (usually linear or network optimization problems) are used to obtain lower bounds to the optimal costs of the restricted problems $\operatorname{minxEY} f(x)$, while various heuristics are used to construct corresponding feasible solutions whose costs can
be llsed as upper bounds. We refer to textbooks such as Nemhauser and Wolsey [NeW88], Bertsimas and Tsitsiklis [BeT97], Bertsekas [Ber98a], and Wolsey [Wol98] for fuller accounts.

### 2.3.4 Const.rained and Multiobjective Problems

In some shortest path contexts, there may be constraints on the resources required to traverse the optimal path, such as limits on time, fuel, etc. For example, there could be a restriction that the total time to travel through the optimal path $P$ should not exceed a given threshold $T$, Le.,

$$
\sum_{(i, j) E P} \tau_{i j} \leq T,
$$

where $\tau_{i j}$ is the time required to traverse arc $(i, j)$. Similarly, there could be a safety constraint, whereby the probability of being able to traverse the path $P$ safely should be no less than a given threshold. Here, we assume that traversal of an arc ( $i, j$ ) will be safe with a given probability Pij. Assuming probabilistic independence of the safety of arc traversa.ls, the probability that traversal of a path $P$ will be safe is the product $\Pi_{(i, j) \in P} p_{i j}$. The requirement that this probability is no less than a given threshold 13 translates to a path length constraint of the form

$$
\sum_{(i, j) E P} \ln \left(p^{\prime} i j\right) \geq \ln (j 3)
$$

We represent path constraints in the generic form

$$
\begin{equation*}
\sum_{(i, j) \in P} c_{i j}^{m} \leq b^{m}, \quad m=/, \ldots, M \tag{2.9}
\end{equation*}
$$

where $c_{i j}^{m}$ is the amount of $m$ th resource required to traverse arc $(i, j)$, and $b^{m}$ is the total amount of mth resource available. Thus the problem is to find a path $P$ that starts at the origin $s$, ends at the destination $t$, satisfies the constraints (2.9), and minimizes

$$
\sum_{(i, j) \in P} a^{i j} .
$$

We refer to this problem as the constrained shoTtest path problem, and we note that it is closely related to the constraint feasibility problem, where we simply want to find a path $P$ that satisfies the constraints (2.9). In particular, the constraint feasibility problem is the special case of the constrained shortest path problem where a'ij 0 for all arcs ( $i, j$ ).

Conversely, the constrained shortest path problem is equivalent to the constraint feasibility problem involving the constraints (2.9) and the additional constraint

$$
\sum_{(\mathrm{i}, \mathrm{j}) E P} a_{\imath j}^{\prime \prime}<L^{*}
$$

where $L^{*}$ is the optimal path length (which, however, is generally unknown).

Another, closely related problem is the multiobjective shortest path problem, where we want to find a path $P$ that simultaneously makes all the lengths

$$
\sum_{(\mathrm{i}, \mathrm{j}) \mathrm{EP}} c_{i j}^{m}, \quad \mathrm{~m}=1, \ldots, M
$$

"small," in a sense that we will now make precise. In particular, for any set $S \subset \Re^{M}$, let us call a vector $x=\left(x_{1}, \ldots, X M\right) \mathrm{E} S$ noninferior if $x$ is not dominated by any vector $Y=\left(y_{1}, \ldots, Y M\right) \mathrm{E} S$, in the sense that

$$
Y m \leq x_{m}, \quad m=I, \ldots, M
$$

and with strict inequality for at least one $m$ (see Fig. 2.3.8). More generally, given a problem with multiple cost functions $f_{1}(x), \ldots, f M(X)$ and a constraint set $X$, we say that $x$ is a noninferior solution if the vector of costs of $x$, i.e., $\left(f_{1}(x), \ldots, f M(X)\right)$, is a noninferior vector of the set of attainable costs

$$
\left\{\left(f_{1}(y), \ldots, f_{M}(y)\right) \mid y \in X\right\}
$$

Note that given a finite set of solutions, there is at least one noninferior solution. Furthermore, one can extract the set of noninferior solutions with a simple algorithm: sequentially test all solutions and discard those that are dominated by some not yet discarded solution, until no more solutions can be discarded.

The multiobjective shortest path problem is to find a noninferior path $P$, Le., one for which there is no other path pl satisfying

$$
\sum_{(\mathrm{i}, \mathrm{j}) E P^{\prime}} c_{c^{\prime} i j}^{m} \leq \sum_{(\mathrm{i}, \mathrm{j}) \mathrm{EP}} \dot{c}^{m} i j, \quad ' \mathrm{~m}=1, \ldots, M
$$

and with strict inequality for at least one $m$. Note that the constraint feasibility problem has a solution if and only if the subset of IYmltiobjective/noninferior solutions that satisfy the constraints (2.9) is nonempty. It follows that the constraint feasibility problem can be easily solved once the set of all noninferior paths is computed. Similarly, the constrained shortest path problem can be solved by casting it as a multiobjective shortest path problem, where the multiple objectives correspond to the cost and the constraints. Given the set of all noninferior solutions, one obtains an optimal solution of the constrained shortest path problem (provided a feasible solution exists), by selecting a path from this set that satisfies the constraints and minimizes the cost. Because of the connections outlined above, the three problems, constrained shortest path, constraint feasibility, and multiobjective, fundamentally share the same mathematical structure, and can be addressed with similar methodology.

## lVlultiobjective DP Problems

We have seen that shortest path problems and deterministic finite-state DP problems are equivalent, so it is not surprising that the rnethodology for multiobjective and/or constrained shortest path problems is shared by multiobjective and/or constrained deterministic finite-state DP problems. A multiobjective version of such a problem involves a single controlled deterministic finite-state system

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}\right)
$$

and multiple cost functions of the form

$$
g_{N}^{m}\left(x_{N}\right)+\sum_{k=O}^{N-1} g_{k}^{m}\left(x_{k}, U k\right), \quad \text { m } \quad 1, \ldots, l V I
$$

Let us provide an extension of the DP algorithm that finds the set of all noninferior solutions to the multiobjective deterministic DP problem involving the above system and cost functions. This algorithm proceeds backwards from the terminal time, and calculates for each stage $k$ and state $X k$, the set of noninferior control sequences for the tail (multiobjective) subproblem that starts at state $X k$. The algorithm is based on a fairly evident extension of the principle of optimality:

If $\left\{u_{k}, u_{k+1}, \ldots, u_{N-1}\right\}$ is a noninferior control sequence for the tail subproblem that starts at $x_{k}$, then $\left\{u_{k+1}, \ldots, u_{N-1}\right\}$ is a noninferior control sequence for the tail subproblem that starts at $f k\left(X k, u_{k}\right)$.

This allows the convenient calculation of the set of noninferior solutions of tail subproblems using the sets of noninferior solutions of shorter subproblems.

More specifically, let $F k(X k)$ be the set of all M-tuples of costs-to-go

$$
\begin{equation*}
\left(g_{N}^{1}\left(x_{N}\right)+\sum_{i=k}^{N-1} g_{i}^{1}\left(x_{i}, u_{i}\right), \ldots, g_{N}^{M}\left(x_{N}\right)+\sum_{i=k}^{N-1} g_{i}^{M}\left(x_{i}, u_{i}\right)\right) \tag{2.10}
\end{equation*}
$$

which correspond to feasible control sequences $\left\{u_{k}, \ldots, u_{N^{-}} \mathrm{I}\right\}$ that start at $x_{k}$ and are noninferior in the following sense: There is no other feasible control sequence $\left\{u_{k}{ }^{\prime}, \ldots, u_{N^{-}} I^{\prime}\right\}$ with corresponding state sequence $\left\{x_{k}^{\prime}, \ldots, x_{N}^{\prime}\right\}$ (where $X k^{\prime}=x_{k}$ ) such that

$$
g_{N}^{m}\left(x_{N^{\prime}}\right)+\sum_{i=k}^{N-1} g_{i}^{m}\left(x_{i}^{\prime}, u_{i}^{\prime}\right) \leq g_{N}^{m}\left(x_{N}\right)+{ }_{i=k}^{\mathrm{N}-\mathbf{I}} g_{i}^{m}\left(x_{i}, u_{i}\right), \quad \mathrm{m}=1, \ldots, M
$$

and with strict inequality for at least one m. Note that $F k(X k)$ is a finite set (since the control constraint set is finite, which implies that the set of control sequences is finite). The sets $F k(X k)$ are generated by an algorithm that starts at the terminal time $N$ with $\mathcal{F} N(X N)$ consisting of just the vector of terminal costs,

$$
F_{N}\left(x_{N}\right)=\left\{\left(g_{N}^{1}\left(x_{N}\right), \ldots, g_{N}^{M}\left(x_{N}\right)\right)\right\}
$$

and proceeds backwards according to the following process: Given the set $\mathcal{F}_{k+1}\left(x_{k+1}\right)$ for all states $X k+l$, the algorithm generates for each state $X k$ the set of vectors

$$
\left(g_{k}^{1}\left(x_{k}, u_{k}\right)+c^{1}, \ldots, g_{k}^{M}\left(x_{k}, u_{k}\right)+c^{M}\right)
$$

such that

$$
\left(c^{1}, \ldots, c^{M}\right) \in \mathcal{F}_{k+1}\left(f_{k}\left(x_{k}, u_{k}\right)\right), \quad u_{k} \in U_{k}\left(x_{k}\right)
$$

then it obtains $F k(X k)$ by extracting the noninferior subset, Le., by discarding from this set the vectors that are dominated by other vectors.

After $N$ steps, this algorithm yields $\mathcal{F}_{0}(x o)$, the set of all noninferior $M$-tuples of costs-to-go starting at initial state $x o$. Note that the algorithm is similar to the ordinary DP algorithm: it just maintains a set of noninferior $1 \backslash t$-tuples of costs-to-go at each state $X k$, rather than a single cost-to-go. Note also that by a similar argument to the one of Section 2.1, it is possible to eonstruct a forward version of this algorithm.

## constrained DP Problems

Let us now consider a related eonstrained DP problem with the same controlled system

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}\right)
$$

where we want to minimize the cost function

$$
\begin{equation*}
g_{N}^{1}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}^{1}\left(x_{k}, u_{k}\right) \tag{2.11}
\end{equation*}
$$

subject to the constraints

$$
\begin{equation*}
g_{N}^{m}\left(x_{N}\right)+\sum_{k=0}^{\mathbf{N}-\mathbf{I}} g_{k}^{m}\left(x_{k}, u_{k}\right) \leq b^{n}, \quad \mathrm{~m} \quad 2, \ldots, M \tag{2.12}
\end{equation*}
$$

We can solve this problem by finding the set of noninferior solutions/control sequences of the multiobjective DP problem involving the costs (2.10), by extracting from this set the subset of solutions/control sequences that satisfy the constraints (2.12), and by selecting from this subset a solution/control sequence that minimizes the cost (2.11).

However, we ean enhance this algorithm by discarding at tIle earliest opportunity control sequences that cannot be part of a feasible control sequence. The rationale for this is related to the ideas underlying label eorrecting methods and the $A^{*}$ algorithm in partieular (d. Section 2.3.2). For $\mathrm{m}=2, \ldots, M$, let $\widetilde{J}_{k}^{m}\left(x_{k}\right)$ be the optimal eost to arrive to $x_{k}$ from the given initial state $X 0$ with cost per stage equal to $g_{i}^{m}\left(x_{i}, u_{i}\right)$. This is the minimal value of

$$
\sum_{i=0}^{k-1} g_{k}^{m}\left(x_{i}, u_{i}\right), \quad m \quad 2, \ldots, 111[
$$

subject to the constraint that the state at the kth stage is $X k$, and can be calculated by the forward DP algorithm of Section 2.1. Consider now a DP-like algorithm that generates for each state and stage, a subset of $M$ tuples. It starts at the terminal time $N$ with the set $F N(X N)$ that consists of just the vector of terminal costs,

$$
\mathcal{F}_{N}\left(x_{N}\right)=\left\{\left(g_{N}^{1}\left(x_{N}\right), \ldots, g_{N}^{M}\left(x_{N}\right)\right)\right\} .
$$

It proceeds backwards as follows: given the set $\mathcal{F}_{k+1}\left(x_{k+1}\right)$ for each state $X k+l$, it generates for each state $X k$ the set of $1 \backslash I I$-tuples

$$
\begin{equation*}
\left(g_{k}^{1}\left(x_{k}, u_{k}\right)+c^{1}, \ldots, g_{k}^{M}\left(x_{k}, u_{k}\right)+c^{M}\right) \tag{2.13}
\end{equation*}
$$

such that

$$
\left(c^{1}, \ldots, c^{M}\right) \in \mathcal{F}_{k+1}\left(f_{k}\left(x_{k}, u_{k}\right)\right), \quad u_{k} \in U_{k}\left(x_{k}\right)
$$

and

$$
\begin{equation*}
\tilde{J}_{k}^{m}\left(x_{k}\right)+g_{k}^{m}\left(x_{k}, u_{k}\right)+c^{m} \leq b^{m}, \quad \mathrm{~m}=2, \ldots, M \tag{2.14}
\end{equation*}
$$

then it obtains $\mathcal{F}_{k}\left(x_{k}\right)$ by extracting the noninferior subset, Le., by discarding from this set the elements that are dominated by other elements.

Note that $1 \backslash 1$-tuples of the form (2.13) that violate the condition (2.14) correspond to paths that cannot be feasible, so they can be safely excluded from further consideration [in fact $\tilde{J}_{k}^{m}\left(x_{k}\right)$ may be replaced by any conveniently available underestimate in the condition (2.14); using $\tilde{J}_{k}^{m}\left(x_{k}\right)$ rnakes this condition as sharp as possible]. The set $F o(x o)$ obtained by the algorithm after $N$ steps consists of $M$-tuples of the cost and the constraint function values

$$
g_{N}^{m}\left(x_{N}\right)+g_{k=O}^{\mathrm{N}-\mathbf{I}} g_{k}^{m}\left(x_{k}, u_{k}\right), \quad m=\boldsymbol{I}, \ldots, M
$$

which correspond to all the feasible solutions that are noninferior. The first component of an M-tuple in $F O(x O)$ corresponds to the cost (2.11). An element of $\mathcal{F}_{0}(x o)$ whose first component is minimal is an optimal solution of the constrained shortest path problem. The advantage of using the criterion (2.14) is that it allows us to discard as early as possible infeasible solutions, and accordingly reduce the size of the sets $F k(x k)$ and the attendant computation.

Note also that if an upper bound, call it UPPER, is known for the optimal path length, it can be used to introduce the additional constraint

$$
g_{N}^{1}\left(x_{N}\right)+\sum_{k=O}^{\mathbf{N}-\mathbf{I}} g_{k}^{1}\left(x_{k}, u_{k}\right) \leq \text { UPPER }
$$

and to make the test (2.14) more effective by augmenting it with the additional inequality

$$
\tilde{J}_{k}^{1}\left(x_{k}\right)+g_{k}^{1}\left(x_{k}, u_{k}\right)+c^{1} \leq \mathrm{UPPLR}
$$

Any $M$-tuple (c$\left.{ }^{1}, \ldots, e M\right)$ that violates this condition corresponds to paths that cannot be optimal, so it can be safely excluded from further consideration. This idea may be further enhanced by introducing schemes to reduce UPPER as the algorithm progresses, similar to label correcting methods.

Clearly multiobjective and constrained DP algorithms require quite a bit more computation and storage than ordinary DP for the same system.

For this reason, there have been many efforts to develop approximate solution methods. These methods are outside our scope and we refer to the literature on the subject.

The preceding DP algorithms for multiobjective and constrained problems are easily adapted to shortest path problems. One possibility is to use the transformation described in Section 2.1 to reformulate the shortest path problem as a (multiobjective or constrained) deterministic finite-state problem. The latter problem can in turn be solved using the DP-like algorithms just described. It is also possible to use versions of label correcting algorithms, including the $A^{*}$ variant, for multiobjective and constrained shortest path problems. A label of a node now is not just a single number, but rather it is an M -dimensional vector with components that correspond to the $N / l$ cost functions; see the end-of-chapter references.

### 2.4 NOTES, SOURCES, AND EXERCISES

Work on the shortest path problem is very extensive. Literature surveys are given by Dreyfus [Dre69], Deo and Pang [DeP84]' and Gallo and Pallottino [GaP88]. For a detailed textbook treatment of shortest paths, see Bertsekas [Ber98a] (the chapter on shortest paths of this book is www-accessible), and also [Ber91a], which contains a variety of associated computer codes.

For a treatment of critical path analysis, see Elmaghraby [Elm78]. A tutorial survey of Hidden Markov Models is given by Rabiner [Rab89]. The Viterbi algorithm, first proposed in [Vit67], is also discussed by Forney [For73]. For applications in comrhunication systems, see Proakis and Salehi [PrS94]' and Sklar [8k188]. For applications in speech recognition, see Rabiner [Rab89] and Picone [Pic9D]. For applications in data network routing, see Bertsekas and Gallager [BeG92].

Label correcting methods draw their origin from the works of Bellman [Be157] and Ford [For56]. The D'Esopo-Pape algorithm appeared in [Pap74] and is based on an earlier suggestion of D'Esopo. For a discussion of various implementations of Dijkstra's algorithm, see Bertsekas [Ber98a]. The 8LF method and some variations were proposed by the author in [Ber93]; see also Bertsekas, Guerriero, and Musmanno [BGIVI:96], where the LLL strategy as well as implementations on a parallel computer of various label correcting methods are discussed. The $A^{*}$ method was proposed by Hart, Nilsson, and Raphael [HNR68] (with corrections in [HNR72]). See also the texts by Nilsson [Nil71], [Ni180], and Pearl [Pea84], which provide a broader discussion of the application of shortest path methods in artificial intelligence.

The Dijkstra algorithm has been extended to continuous space shortest path problems by Tsitsildis [Tsi75]. The SLF/LLL methods have also
been sirnilarly extended by Bertsekas, Guerriero, and Musmanno [BGM95], and by Polymenakos, Bertsekas, and Tsitsiklis [PBT98].

There is extensive literature on exact and approximate solution methods for constrained and multiobjective shortest path and DP problems. Analogs of label correcting and Dijkstra-like methods were proposed by Vincke [Vin74] and Hansen [Han80], respectively; see also Jaffe [Jaf84] and Martins [Mar84]. Recent work includes Guerriero and Musmanno [GuMO1], who investigate analogs of the SLF/LLL methods, and give many references and computational results. For a multiobjective version of the $A^{*}$ method, see Stewart and White [StW91], who also survey earlier work.

## EXERCISES

## 2.1

Find a shortest path from each node to node 6 for the graph of Fig. 2.4.1 by using the DP algorithm.


Figure 2.4.1 Graph for Exercise 2.1. The arc lengths are shown next to the arcs.

## 2.2

Find a shortest path fronl node 1 to node 5 for the graph of Fig. 2.4.2 by using the label correcting method of Section 2.3.1.
2.3

Air transportation is available between $n$ cities, in some cases directly and in others through intermediate stops and change of carrier. The airfare between cities


Figure 2.4.1 Graph for Exercise 2.2. The arc lengths are shown next to the arcs.
$i$ and $j$ is denoted by $a \cdot i j$. We assume that $a_{i j}=a j i$, and for notational convenience, we write $a i j=0$ if there is no direct flight between $i$ and $j$. The problem is to find the cheapest airfare for going between two cities perhaps through intermediate stops. Let $n=6$ and al2 $=30$, als $=60, a_{14}=25$, als $=\mathrm{a} 16=00$, ${ }_{a 23}={ }_{0,24}={ }_{0,25} \quad$ oo, $0,26=50,0,34=35,0,35 \quad 0,36=00,0,45=15, a_{46}=00$, $a_{56}=15$. Find the cheapest airfare from every city to every other city by using the DP algorithm.

## 2.4 (Dijkstra's Algorithm for Shortest Paths)

Consider the best-first version of the label correcting algorithm of Section 2.3.1. Here at each iteration we remove from OPEN a node that has minimum label over all nodes in OPEN.
(a) Show that each node $j$ will enter OPEN at most once, and show that at the time it exits OPEN, its label $d_{j}$ is equal to the shortest distance from s to $\mathbf{j}$. Hint: Use the nonnegative arc length assumption to argue that in the label correcting algorithm, in order for the node $i$ that exits OPEN to reenter, there must exist another node $k$ in OPEN with $d_{k}+a_{k i}<d_{i}$.
(b) Show that the number of arithmetic operations required for termination is bounded by $c N^{2}$ where $N$ is the number of nodes and c is some constant.

## 2.5 (Label Correcting for Acyclic Graphs)

Consider a shortest path problem involving an acyclic graph. Let Sk be the set of nodes $i$ such that all paths from the origin to $i$ have $k$ arcs or less and at least one such path has $k$ arcs. Consider a label correcting algorithm that removes from OPEN a node of $S k$ only if there are no nodes of $8_{1}, \ldots, S_{k-1}$ in OPEN. Show that each node will enter OPEN at most once. How does this result relate to the type of shortest path problem arising from deterministic DP (d. Fig. 2.1.1)?

## 2.6 (Label Correcting with Multiple Destinations)

Consider the problem of finding a shortest path from node $s$ to each node in a subset $T$, assuming that all arc lengths are nonnegative. Show that the following modified version of the label correcting algorithm of Section 2.3.1 solves the problem. Initially, UPPER $=00, d_{s}=0$, and $d_{i}=$ oo for all $\mathrm{i} \neq \mathrm{s}$.

Step 1: Remove a node i from OPEN and for each child $\boldsymbol{j}$ of $\mathbf{i}$, execute step 2. Step 2: If $d_{i}-\frac{-}{}$ aij $<\min \left\{\mathrm{el}_{j}\right.$, UPPER $\}$, set $e l j=d_{i}-\mathrm{l}$ aij, set $\boldsymbol{i}$ to be the parent of $\boldsymbol{j}$, and place $\boldsymbol{j}$ in OPEN if it is not already in OPEN. In addition, if $\boldsymbol{j} \in \boldsymbol{T}$, set UPPER $=$ maXtET $e l_{t}$
Step 3: If OPEN is empty, terminate; else go to step 1.
Prove a termination property such as the one of Prop. 2.3.1 for this algorithm.

## 2.7 (Label Correcting with Negative Arc Lengths)

Consider the problem of finding a shortest path from node $s$ to node $t$, and assume that all cycle lengths are nonnegative (instead of all arc lengths being nonnegative). Suppose that a scalar $u_{j}$ is known for each node $j$, which is an underestimate of the shortest distance from $\boldsymbol{j}$ to $t\left(u_{j}\right.$ can be taken $-\infty$ if no underestimate is known). Consider a modified version of the typical iteration of the label correcting algorithm of Section 2.3.1, where step 2 is replaced by the following:
IVfodified Step 2: If $d_{i}-\mathrm{I}-a i j<\min \{d j$, UPPER - ' $U j\}$, set $d j=d i-\mathrm{I}-a i j$ and set $i$ to be the parent of j . In addition, if $\mathrm{j} \neq t$, place j in OPEN if it is not already in OPEN, while if $\mathrm{j}=t$, set UPPER to the new value $d_{i}-1$ ait of $d t$.
(a) Show that the algorithm terminates with a shortest path, assuming there is at least one path from $s$ to $t$ (d. Prop. 2.3.1).
(b) Why is the algorithm of Section 2.3 .1 a special case of the one of this exercise?

## 2.8

We have a set of $N$ objects, denoted $1,2, \ldots, N$, which we want to group in clusters that consist of consecutive objects. For each cluster i, $i-1, \ldots, j$, there is an associated cost $a i j$. We want to find a grouping of the objects in clusters such that the total cost is minimum. Formulate the problem as a shortest path problem, and write a DP algorithm for its solution. (Note: An example of this problem arises in typesetting programs, such as TEXjLATEX, that break down a paragraph into lines in a way that optimizes the paragraph's appearance.)

## 2.9 (Shortest Path Tour Problem [BeC04])

Consider a problem of finding a shortest path from a given ongm node $s$ to a given destination node $t$ in a graph with nonnegative arc lengths. However, there is the constraint that the path should successively pass through at least one node from given node subsets $\mathrm{T}_{1}, T_{2}, \cdots, T N$ (Le., for all $k$, pass through some node from the subset $T k$ after passing through at least one node of the subsets $\left.T_{1}, \ldots, T k-I\right)$.
(a) Formulate this as a dynamic programming problem.
(b) Show that a solution can be obtained by solving a sequence of ordinary shortest path problems, each involving a single origin and multiple destinations.

### 2.10 (Two-Sided Dijkstra Algorithm [Nic66])

Consider a problem of finding a shortest path from a given origin node $s$ to a given destination node $t$ in a graph with nonnegative arc lengths. Consider an algorithm that maintains two subsets of nodes, Wand $V$, with the following properties:
(1) $s E W$ and $t \mathrm{E} V$.
(2) If' E Wand $\mathrm{j} \notin W$, then the shortest distance from $s$ to ${ }^{\prime} i$ is less than or equal to the shortest distance from $s$ to $\mathbf{j}$.
(3) If $\boldsymbol{i}$ E $V$ and $\mathrm{j} \notin V$, then the shortest distance from i to $t$ is less than or equal to the shortest distance from j to $t$.
At each iteration the algorithm adds a new node to $W$ and a new node to $V$ (the Dijkstra algorithm can be used for this purpose), and terminates when Wand $V$ have a node in common. Let $\boldsymbol{d}_{\boldsymbol{i}}^{s}$ be the shortest distance from $s$ to ' u using paths all the nodes of which, with the possible exception of i , lie in $W$ ( $d_{i}^{s}=$ oo if no such path exists), and let $\boldsymbol{d}_{i}^{t}$ be the shortest distance from i to $t$ using paths all the nodes of which, with the possible exception of i , lie in $V\left(d_{i}^{t} \quad\right.$ oo if no such path exists).
(a) Show that upon termination, the shortest distance $D_{s t}$ from $s$ to t is given by

$$
D_{s t}=\min _{i \in W}\left\{d_{i}^{s}-\mathrm{F} d_{i}^{t}\right\}=\min _{i \in}\left\{d_{i}^{s}+d_{i}^{t}\right\}_{i}=\min \left\{d_{i}^{s}-1-d_{i}^{t}\right\}
$$

(b) Show that the conclusion of part (a) holds if the algorithm is terminated once the condition

$$
\operatorname{mam}_{i \in W}\left\{d S+d_{i}^{t}\right\} \leq \max _{i \in W} d_{i-\mathrm{I}-\max _{\mathrm{iEV}}^{S}} d_{i}^{t}
$$

holds, even if the sets Wand $V$ have no node in common.

### 2.11 (DP on Two Parallel Processors [Las85])

Formulate a DP algorithm to solve the deterministic problern of Section 2.1 on a parallel computer with two processors. One processor should execute a forward algorithm and the other a backward algorithm.

### 2.12 (Doubling Algorithms)

Consider a deterministic finite-state problem that is time-invariant in the sense that the state and control spaces, the cost per stage, and the system equation
are the same for each stage. Let $J k(X, y)$ be the optimal cost to reach state $y$ at time $k$ starting from state $x$ at time 0 . Show that for all $k$

$$
J_{2 k}(x, y)=\min _{z}\left\{J_{k i}(x, z)+J_{k}(z, y)\right\}
$$

Discuss how this equation may be used with advantage to solve problems with a large number of stages.

### 2.13 (Distributed Shortest Path Computation [Ber82a]) (www)

Consider the problem of finding a shortest path from nodes $1,2, \ldots, N$ to nodet, and assume that all arc lengths are nonnegative and all cycle lengths are positive. Consider the iteration

$$
\begin{gathered}
t_{i}^{k}+{ }^{l}=\operatorname{mm}_{\boldsymbol{J}}\left[a_{i j}+d_{j}^{k}\right], \quad \mathrm{i}=1,2, \ldots, N \\
d_{t}^{k+1}=0
\end{gathered}
$$

(a) It was shown in Section 2.1 that if the initial condition is $d_{i}^{0}=$ oo for $i=1, \ldots, N$ and $d_{t}^{0}=0$, then the iteration (2.15) yields the shortest distances in $N$ steps. Show that if the initial condition is $d_{i}^{0}=0$, for all $i=1, \ldots, N, t$, then the iteration (2.15) yields the shortest distances in a finite mnnber of steps.
(b) Assume that the iteration

$$
\begin{equation*}
d_{i}:=\min _{j}\left[a_{i j}+d_{j}\right] \tag{2.16}
\end{equation*}
$$

is executed at node $i$ in parallel with the corresponding iteration for $d j$ at every other node $\mathbf{j}$. However, the times of execution of this iteration at the various nodes are not synchronized. Furthermore, each node i communicates the results of its latest computation of $d_{i}$ at arbitrary times with potentially large communication delays. Therefore, there is the possibility of a node executing iteration (2.16) several times before receiving a communication from every other neighboring node. Assume that each node never stops executing iteration (2.16) and transmitting the result to the other nodes. Show that the values $d_{i}^{T}$ available at time $T$ at the corresponding nodes $i$ are equal to the shortest distances for all $T$ greater than a finite time $\bar{T}$. Hint: Let $\bar{d}_{i}^{k}$ and $\underline{d}_{i}^{k}$ be generated by iteration (2.16) when starting from the first and the second initial conditions in part (a), respectively. Show that for every $k$ there exists a time $T_{k}$ such that for all $T \geq T_{k}$ and $k$, we have $\underline{d}_{i}^{k} \leq d_{i}^{T} \leq \bar{d}_{i}^{k}$. Note: For a detailed analysis of asynchronous iterations, including algorithms for shortest paths and DP, see Bertsekas and Tsitsiklis [BeTS9], Ch. 6. Distributed asynchronous shortest path algorithms find extensive application in the problem of packet routing in (lata communication networks. For a related discussion and analysis, see 13ertsekas and Gallager [13eG92], Ch. 5.

### 2.14 (Shortest Paths for an Infinite Number of Nodes)

Consider the shortest path problem of Section 2.3, except that the number of nodes in the graph may be countably infinite (although the nurnber of outgoing $\operatorname{arcS}$ from each node is finite). We assume that the length of each arc is a positive integer. Furthermore, there is at least one path from the origin node $s$ to the destination node $t$. Consider the label correcting algorithm as stated and initialized in Section 2.3.1, except that UPPER is initially set to some integer that is an upper bound to the shortest distance from $s$ to $t$. Show that the algorithm will terminate in a finite number of steps with UPPER equal to the shortest distance from s tot. Hint: Show that there is a fmite lllunber of nodes whose shortest distance from $s$ does not exceed the initial value of UPPER.

### 2.15 (Path Bottleneck Problem)

Consider the framework of the shortest path problem. For any path $P$, define the bottleneck arc of P as an arc that has maximum length over all arcs of $P$. Consider the problem of finding a path whose length of bottleneck arc is minimum, among the paths connecting an origin node and a destination node. Develop and justify an analog of the label correcting method of Section 2.3.1. Hint: Replace $d_{i}+$ aij with $\max \left\{d_{i}, a^{\prime} i j\right\}$ '

### 2.16

Air transportation is available between all pairs of $n$ cities, but because of a perverse fare structure, it may be more economical to go from one city to another through intermediate stops. A cost-minded traveler wants to find the minimum cost fare to go from an origin city $s$ to a destination city $t$. The airfare between cities $i$ and j is denoted by $a_{i j}$, and for the mth intermediate stop, there is a stopover cost $\mathrm{c}_{\mathrm{m}}$ (aij and $\mathrm{c}_{\mathrm{m}}$ are assumed positive). Thus, for example, to go from $s$ to $t$ directly it costs ast, while to go from $s$ to $t$ with intermediate stops at cities $i_{1}$ and $i_{2}$, it costs $a_{s i_{1}}+\mathrm{c} 1+a_{i_{1} i_{2}}+c_{2}+a^{\prime} i 2 t$.
(a) Formulate the problem as a shortest path problem, and identify the nodes, arcs, and arc costs.
(b) Write a corresponding DP algorithm that finds an optimal solution in $n-2$ stages.
(c) Assume that $c_{\boldsymbol{m}}$ is the same for all m . Devise a rule for detecting that an optimal solution has been found before iteration $n-2$ of the DP algorithm.

### 2.17

A businessman operates out of a van that he sets up in one of two locations on each day. If he operates in location $\boldsymbol{i}$ (where $i=1,2$ ) on day $k$, he makes a known and predictable profit, denoted $r_{k}^{i}$. However, each time he moves from one location to the other, he pays a setup cost c . The businessman wants to maximize his total profit over $N$ days.
(a) Show that the problem can be formulated as a shortest path problem, and write the corresponding DP algorithm.
(b) Suppose he is at location i on day $k$. Let

$$
R_{k}^{i}=r_{k}^{\bar{i}}-r_{k}^{i}
$$

where $\overline{\boldsymbol{i}}^{\text {denotes the location that is not equal to i. Show that if } R_{k}^{i} \leq 0 \text { it }}$ is optimal to stay at location i , while if $R_{k}^{i} \geq 2 c$, it is optimal to switch.
(c) Suppose that on each day there is a probability of rain Pi at location $\boldsymbol{i}_{\boldsymbol{i}}$ independently of rain in the other location, and independently of whether it rained on other days. If he is at location $i$ and it rains, his profit for the day is reduced by a factor $\beta_{i}$. Can the problem still be formulated as a shortest path problem? Write a DP algorithm.
(d) Suppose there is a possibility of rain as in part (c), but the businessman receives an accurate rain forecast just before making the decision to switch or not switch locations. Can the problem still be formulated as a shortest path problem? Write a DP algorithm.
eterministic Continuous-Time

## Optimal ntrol

## Contents

3.1. Continuous-Time Optimal Control ..... p. 106
3.2. The Hamilton-Jacobi-Bellman Equation ..... p. 1093.3. The Pontryagin Minimum Principlep. 115
3.3.1. An Informal Derivation Using the HJB Equation ..... p. 115
3.3.2. A Derivation Based on Variational Ideas ..... p. 125
3.3.3. Minimum Principle for Discrete-Time Problems ..... p. 129
3.4. Extensions of the Minimum Principle ..... p. 131 ..... p. 131
3.4. L Fixed Terminal State
3.4. L Fixed Terminal State
3.4.2. Free Initial State ..... p. 135 ..... p. 135
3.4.3. Free Terminal Time
3.4.3. Free Terminal Time
3.4.4. Time-Varying System and Cost ..... p. 138
3.4.5. Singular Problems ..... p. 139
3.5. Notes, Sources, and Exercises ..... p. 142

In this chapter, we provide an introduction to continuous-time deterministic optimal control. We derive the analog of the DP algorithm, which is the Hamilton-Jacobi-Bellman equation. Furthermore, we develop a celebrated theorem of optimal control, the Pontryagin Minimum Principle and its variations. We discuss two different derivations of this theorem, one of which is based on DP. We also illustrate the theorem by means of examples.

### 3.1 CONTINUOU8-TII\!IE OPTIMAL CONTROL

We consider a continuous-time dynamic system

$$
\begin{gathered}
\dot{x}(t)=f(x(t), u(t)), \quad \mathrm{O}_{\leq \mathrm{t}} \leq T \\
x(0): \text { given }
\end{gathered}
$$

where $x(t) \mathrm{E} \Re^{n}$ is the state vector at time $\mathrm{t}, j ;(t) \mathrm{E} \Re^{n}$ is the vector of first order time derivatives of the states at time $t, u(t) \mathrm{E} U \mathrm{C} \Re^{m}$ is the control vector at time $t, U$ is the control constraint set, and $T$ is the terminal time. The components of $f, x, \dot{x}$, and $u$ will be denoted by $f i, X i, \dot{x}_{i}$, and $U i$, respectively. Thus, the system (3.1) represents the $n$ first order differential equations

$$
\frac{d X i(t}{d t} \underline{\underline{2}} f_{i}(x(t), u(t)), \quad \mathrm{i}=1, \ldots, n o
$$

We view $j ;(t), x(t)$, and $u(t)$ as column vectors. We assume that the system function $I i$ is continuously differentiable with respect to $x$ and is continuous with respect to $u$. The admissible control functions, also calle of control trajectories, are the piecewise continuous functions $\left\{\left.u(t)\right|_{t E[0,} \prod_{n}\right.$ with $u(t) \mathrm{E} U$ for all $\mathrm{E}[0, T)$.

We should stress at the outset that the subject of this chapter is highly sophisticated, and it is beyond our scope to develop it according to high standards of mathematical rigor. In particulat, we assume that, for any admissible control trajectory $\{u(t) \mid t \mathrm{E}[0$, $\mid$ n, the system of differential equations (3.1) has a unique solution, which is denoted $\{x u(t) \mid$ $t ; \mathrm{E}[0, T)\}$ and is called the corresponding state tmjectory. In a more rigorous treatment, the issue of existence and uniqueness of this solution would have to be addressed more carefully.

We want to find an admissible control trajectory $\{u(t)$ It $\mathrm{E}[0$, Thl, which, together with its corresponding state trajectory $\left\{\left.x(t)\right|_{t \in[0,} \cap\right.$, minimizes a cost function of the form

$$
h(x(T))+\int_{0}^{T} g(x(t), 1 t(t)) d t
$$

where the functions 9 and $h$ are continuously differentiable with respect to $x$, and 9 is continuous with respect to $u$.

## Example 3.1.1 (Motion Control)

A unit mass moves on a line under the influence of a force $u$. Let $\mathrm{xl}(t)$ and $X 2(t)$ be the position and velocity of the mass at time $t$, respectively. From a given ( $\mathrm{X1}(0), X 2(0)$ ) we want to bring the mass "near" a given final position-velocity pair ( $\bar{x}_{1}, \bar{x}_{2}$ ) at time $T$. In particular, we want to

$$
\text { minimize }\left|x_{1}(T)-\bar{x}_{1}\right|^{2}+\left|x_{2}(T)-\bar{x}_{2}\right|^{2}
$$

subject to the control constraint

$$
|u(t)| \leq 1, \quad \text { for alIt } \mathrm{E}[0, T]
$$

The corresponding continuous-time system is

$$
\dot{x}_{1}(t)=X 2(t), \quad \dot{x}_{2}(t)=u(t)
$$

and the problem fits the general framework given earlier with cost functions given by

$$
\begin{aligned}
& h(x(T))=\left|x_{1}(T)-\bar{x}_{1}\right|^{2}+\left|X 2(T)-\bar{x}_{2}\right|^{2} \\
& g(x(t), u(t))=0, \quad \text { for all } t \mathrm{E}\left[0,1^{\prime}\right]
\end{aligned}
$$

There are many variations of the problem; for example, the final position and/or velocity may be fixed. These variations can be handled by various reformulations of the general continuous-time optimal control problem, which will be given later.

## Example 3.1.2 (Resource Allocation)

A producer with production rate $x(t)$ at time $t$ may allocate a portion $u(t)$ of his/her production rate to reinvestment and $1-u(t)$ to production of a storable good. Thus $x(t)$ evolves according to

$$
\dot{x}(t)=\gamma u(t) x(t)
$$

where $\gamma>0$ is a given constant. The producer wants to maximize the total amount of product stored

$$
\int_{0}^{x}(1-n(t)) x(t) d l
$$

subject to

$$
0 \leq u(t) \leq 1, \quad \text { for all } t \mathrm{E}[0, T]
$$

The initial production rate $x(O)$ is a given positivenurnber.

## 

Calculus of variations problems involve finding (possibly multidimensional) curves $x(t)$ with certain optimality properties. They are among the most celebrated problems of applied mathematics and have been worked on by many of the illustrious mathematicians of the past 300 years (Euler, Lagrange, Bernoulli, Gauss, etc.). We will see that calculus of variations problems can be reformulated as optimal control problems. We illustrate this reformulation by a simple example.

Suppose that we want to find a minimum length curve that starts at a given point and ends at a given line. The answer is of course evident, but we want to derive it by using a continuous-time optimal control formulation. Without loss of generality, we let $(0, \alpha)$ be the given point, and we let the given line be the vertical line that passes through ( $T, 0$ ), as shown in Fig. 3.1.1. Let also $(t, x(t))$ be the points of the curve $(0 \leq t \leq T)$. The portion of the curve joining the points ( $t, x(t »)$ and $(t+d t, x(t+d t »)$ can be approximated, for small $d t$, by the hypotenuse of a right triangle with sides $d t$ and $x(t) d t$. Thus the length of this portion is

$$
\sqrt{(d t)^{2}+(\dot{x}(t))^{2}(d t)^{2}}
$$

which is equal to

$$
\sqrt{1+(X(t)) 2} d t
$$

T'he length of the entire curve is the integral over $[0, T]$ of this expression, so the problem is to

$$
\begin{aligned}
& \operatorname{minimize} \int_{0}^{T} \sqrt{1+(\dot{x}(t))^{2}} d l \\
& \text { subject to } x(O)=\alpha
\end{aligned}
$$

To refonnulate the problem as a continuous-time optimal control problem, we introduce a control $u$ and the system equation

$$
\dot{x}(t)=u(t), \quad x(O)=\alpha
$$

Our problem then becomes

$$
\operatorname{minimize} \int_{0}^{T} \sqrt{1+("(i)) 2} d l \text {, }
$$

This is a problem that fits our continuous-time optimal control framework.


Figure 3.L1 Problem of finding a curve of minimum length from a given point to a given line, and its formulation as a calculus of variations problem.

We will now derive informally a partial differential equation, which is satisfied by the optimal cost-to-go function, under certain assumptions. This equation is the continuous-time analog of the DP algorithm, and will be motivated by applying DP to a discrete-time approximation of the continuoustime optimal control problem.

Let us divide the time horizon $[0, T]$ into $N$ pieces using the discretization interval

$$
\delta=\stackrel{T}{N^{\prime}}
$$

We denote

$$
\begin{array}{lrlrl}
x k & x(k \delta), & k & =0,1, & , N, \\
u_{k}=u(k \delta), & ' . k & =0,1, & , N,
\end{array}
$$

and we approximate the continuous-time system by

$$
x_{k+1}=x_{k}+f\left(x_{k}, u_{k}\right) \cdot \delta
$$

and the cost function by

$$
h\left(x_{N}\right)+\sum_{k=()}^{N-J} g\left(x_{k}, u_{k}\right) \cdot \delta
$$

We now apply DP to the discrete-time approximation. Let
$J^{*}(t, x)$ : Optima.l cost-ta-go a.t time $t$ and state $x$ for the continuous-time problem,
$\tilde{J}^{*}(t, x)$ : Optimal cost-to-go at time $t$ and state $x$
for the discrete-time approximation.

The DP equations are

$$
J^{*}(N o, x)=h(x)
$$

$\tilde{J}^{*}(k \delta, x)=\min _{u \in U}\left[g(x, u) \cdot \delta+\tilde{J}^{*}((k+1) \cdot \delta, x+j(x, u) \cdot \delta)\right], \quad k=0, \ldots, N-1$.
Assuming that $\tilde{J}^{*}$ has the required differentiability properties, we expand it into a first order Taylor series around $(k \delta, X)$, obtaining

$$
\begin{aligned}
\tilde{J} *((k+1) \cdot \delta, x+f(x, u) \cdot \delta)=\tilde{J}^{*}(k \delta, x) & +\nabla_{t} \tilde{J}^{*}(k \delta, x) \cdot \delta \\
& +\nabla_{x} \tilde{J}^{*}(k \delta, x)^{\prime} j(x, u) \cdot \delta+o(\delta),
\end{aligned}
$$

$\underset{\nabla}{\text { where }} o(\delta)$ represents second order terms satisfying $\lim _{\delta \rightarrow 0} \mathrm{O}(\delta) j_{\delta}=0$, $\nabla_{\mathrm{t}}$ denotes partial derivative with respect to $t$, and $\nabla x$ denotes the $n$ dilnensional (column) vector of partial derivatives with respect to $x$. Substituting in the DP equation, we obtain

$$
\begin{aligned}
\tilde{J}^{*}(k \delta, x)=\operatorname{IY}_{u \in U} \operatorname{lin}\left[\mathrm{~g}(\mathrm{x}, \mathrm{u}) \cdot \delta+\tilde{J}^{*}(k \delta, x)\right. & +\nabla_{t} \tilde{J}^{*}(k \delta, x) \cdot \delta \\
& \left.+\nabla_{x} \tilde{J}^{*}(k \delta, x)^{\prime} j(x, u) \cdot \delta+0(0)\right]
\end{aligned}
$$

Canceling $\tilde{J}^{*}(k \delta, x)$ from both sides, dividing by $\delta$, and taking the limit as $\delta \rightarrow 0$, while assuming that the discrete-time cost-to-go function yields in the limit its continuous-time counterpart,

$$
\lim _{k \rightarrow \infty, \delta \rightarrow 0, k \delta=t} \tilde{J}^{*}(k \delta, x)=J^{*}(t, x), \quad \text { for all } t, x
$$

we obtain the following equation for the cost-to-go function $J^{*}(t, x)$ :

$$
0 \quad \min _{u \in U}\left[\mathrm{~g}(\mathrm{x} ;, \mathrm{u})+\nabla_{t} J^{*}(t, x)+" V x J^{*}(t, x)^{\prime} j(x, u)\right], \quad \text { for all } \mathrm{t}, x \text {, }
$$

with the boundary condition $J^{*}(T, x)=h(x)$.
This is the Hamilton-Jacobi-Bellman (HJB) equat'ion. It is a partial clifferential equation, which should be satisfied for all time-state pairs $(t, x)$ by the cost-to-go function $J^{*}(t, x)$, based on the preceding informal derivation, which assumed among other things, differentiability of $J^{*}(t, x)$. In fact we do not Imow a priori that $J^{*}(t, x)$ is differentiable, so we do not know if $J^{*}(t, x)$ solves this equation. However, it turns out that if we can solve the HJB equation analytically or computationally, then we can obtain an optimal control policy by minimizing its right-hand-side. This is shown in the following proposition, whose statement is reminiscent of a corresponding statement for discrete-time DP: if we can execute the DP algorithm, which may not be possible due to excessive computational requirements, we can find an optimal policy by minimization of the right-hand side.

Proposition 3.2.1: (Sufficiency Theorem) Suppose $V(t, x)$ is a solution to the HJB equation; that is, $V$ is continuously differentiable in $t$ and $x$, and is such that

$$
\mathrm{O}=\min _{u \in \cup}\left[\mathrm{~g}(\mathrm{x}, \mathbf{u})+\nabla_{t} V(t, x)+\nabla_{x} V(t, x)^{\prime} f(x, u)\right], \quad \text { for all } t, x,
$$

$$
\begin{equation*}
V(T, x)=h(x), \quad \text { for all } x \tag{3.2}
\end{equation*}
$$

(3.3)

Suppose also that $\mu^{*}(t, x)$ attains the minimum in Eq. (3.2) for all $t$ and $x$. Let $\left\{x^{*}(t)\right.$ It $\mathrm{E}[O, T J\}$ be the state trajectory obtained from the given initial condition $x(O)$ when the control trajectory $u^{*}(t)=$ $p^{*}\left(t, x^{*}(t)\right), \mathrm{t} \mathrm{E}[0, T]$ is used [that is, $\mathrm{x}^{*}(\mathrm{O})=x(0)$ and for all $t \mathrm{E}$ $[O, T], \dot{x}^{*}(t) \quad j\left(x^{*}(t), p^{*}\left(t, x^{*}(t)\right)\right)$; we assume that this differential equation has a unique solution starting at any pair $(t, x)$ and that the control trajectory $\left\{p,{ }^{*}\left(t, x^{*}(t)\right)\right.$ It $\mathrm{E}[O, T J\}$ is piecewise continuous as a function of $t]$. Then $V$ is equal to the optimal cost-to-go function, I.e.,

$$
V(t, x)=J^{*}(t, x), \quad \text { for all } t, x
$$

Furthermore, the control trajectory $\left\{\left.u^{*}(t)\right|_{t \in[0, T]\}}\right.$ is optimal.

Proof: Let $\left\{\left.\hat{u}(t)\right|_{t \mathrm{E}[0, T n \text { be any admissible control trajectory and let }}\right.$ $\{\hat{x}(t) \mid t \mathrm{E}[0, T n$ be the corresponding state trajectory. From Eq. (3.2) we have for all $t \mathrm{E}[0, T]$

$$
\mathbf{0} \leq g(\hat{x}(t), \hat{u}(t))+\nabla_{t} V(t, \hat{x}(t))+\nabla_{x} V(t, \hat{x}(t))^{\prime} f(\hat{x}(t), \hat{u}(t))
$$

Using the system equation $\dot{\hat{x}}(t)=j(\hat{x}(t), \hat{u}(t))$, the right-hand side of the above inequality is equal to the expression

$$
g(\hat{x}(t), \hat{u}(t))+\frac{d}{d t}(V(t, \hat{x}(t)))
$$

where $d j d t(\cdot)$ denotes total derivative with respect to $t$. Integrating this expression over $t \mathrm{E}[0, T]$, and using the preceding inequality, we obtain

$$
0 \leq \int_{0}^{T} g(\hat{x}(t), i l(t)) d t+V\left(1^{\prime}, \hat{x}(T)\right) \quad V(O, x(O))
$$

Thus by using the terminal condition $V(T, x)=h(x)$ of Eq. (3.3) and the initial condition $\hat{x}(0)=x(0)$, we have

$$
V(O, x(O)) \leq h(\hat{x}(T))+\int_{0}^{T} g(\hat{x}(t), \hat{u}(t)) d t
$$

If we use $u^{*}(t)$ and $X^{*}(t)$ in place of $\hat{u}(t)$ and $\hat{x}(t)$, respectively, the preceding inequalities becomes equalities, and we obtain

$$
v(0, x(O))=h\left(x^{\prime}(T)\right)+\int_{0}^{T} g\left(x^{\prime \prime}\left(t j, u^{*}(t)\right) d t\right.
$$

Therefore the cost corresponding to $\left\{u^{*}(t)\right.$ It $\left.\mathrm{E}[O, T]\right\}$ is $V(O, x(O))$ and is no larger than the cost corresponding to any other admissible.control trajectory $\left\{\hat{u}(t)\right.$ It $\mathrm{E}\left[0, T n\right.$. It follows that $\left\{u^{*}(t)\right.$ It $\left.\mathrm{E}[0, T]\right\}$ is optimal and that

$$
V(O, x(O))=J^{*}(O, x(O))
$$

We now note that the preceding argument can be repeated with any initial time $t \mathrm{E}[0, T J$ and any initial state $x$. We thus obtain

$$
V(t, x)=J^{*}(t, x), \quad \text { for all } t, x
$$

Q.E.D.

## Example 3.2.1

To illustrate the HJB equation, let us consider a simple example involving the scalar system

$$
x(t)=u(t)
$$

with the constraint $|u(t)| \leq 1$ for all $t \mathrm{E}[0, \mathrm{TJ}$. The cost is

$$
\frac{1}{2}(x(T))^{2}
$$

The HJB equation here is

$$
\begin{equation*}
\left.\mathrm{O}=\min _{|u| \leq 1}\left[\nabla_{t} V e t, x\right)+\nabla_{x} V(t, x) u\right], \quad \text { for all } t, x \tag{3.4}
\end{equation*}
$$

with the terminal condition

$$
\begin{equation*}
V(T, x)=\frac{1}{2} x^{2} \tag{3.5}
\end{equation*}
$$

There is an evident candidate for optimality, namely moving the state towards 0 as quickly as possible, and keeping it at 0 once it is at 0 . The corresponding control policy is

$$
\mu^{*}(t, x)=-\operatorname{sgn}(\mathrm{x})= \begin{cases}\text { I } & \text { if } x<0  \tag{3.6}\\ 0 & \text { if } x=0 \\ -1 & \text { if } x>0\end{cases}
$$

For a given initial time $t$ and initial state $x$, the cost associated with this policy can be calculated to be

$$
\begin{equation*}
J^{*}(t, x)=\frac{1}{2}(\max \{0, \mathrm{Ix} 1-(\mathrm{T}-t)\}) 2 \tag{3.7}
\end{equation*}
$$

## $J^{*}(t, x)$



Figure 3.2.1 Optimal cost-to-go function $J^{*}(t, x)$ for Example 3.2.1.

This function, which is illustrated in Fig. 3.2.1, satisfies the terrninal condition (3.5), since $J^{*}(T, x)=(1 / 2) x^{2}$. Let us verify that this function also satisfies the HJB Eq. (3.4), and that $u=-\operatorname{sgn}(x)$ attains the minimum in the righthand side ofthe equation for all $t$ and $x$. Proposition 3.2.1 will then guarantee that the state and control trajectories corresponding to the policy $\boldsymbol{\mu}^{*}(t, x)$ are optimal.

Indeed, we have

$$
\begin{gathered}
\nabla_{t} J^{*}(t, x)=\max \{\mathrm{O}, 1 \mathrm{x} 1-(T-t)\} \\
\nabla_{x} J^{*}(t, x)=\operatorname{sgn}(\mathrm{x}) \cdot \max \{0, \mathrm{xx} 1-(T-i)\}
\end{gathered}
$$

Substituting these expressions, the HJB Eq. (3.4) becomes

$$
\mathrm{O}=\min _{|u| \leq 1}[1+\operatorname{sgn}(\mathrm{x}) \cdot u] \max \{\mathrm{O},|x|-(T-t)\}
$$

which can be seen to hold as an identity for all $(i, x)$. Purthermore, the minimum is attained for $u=-\operatorname{sgn}(x)$ : We therefore conclude based on Prop. 3.2.1 that $J^{*}(t, x)$ as given by Eq. (3.7) is indeed the optimal cost-to-go function, and that the policy defined by Eq. (3.6) is optimal. Note, however, that the optimal policy is not unique. Based on Prop. 3.2.1, any policy for which the minimum is attained in Eq. (3.8) is optimal. In particular, when $I X(i) 1 \leq T-i$, applying any control from the range $[-1,1)$ is optimal.

The preceding derivation generalizes to the case of the cost

$$
h(x(T »),
$$

where $h$ is a nonnegative differentiable convex function with $h(O)$
0 . The corresponding optimal cost-to-go function is

$$
J^{*}(t, x)= \begin{cases}f h(x-(T-t>) & \text { if } x>T-i \\ h(x+(T-t)) & \text { if } x<-(T-i), \\ 0 & \text { if } \mathrm{Ixl} \leq T-t\end{cases}
$$

and can be similarly verified to be a solution of the HJB equation.

## Hxample 3.2.2 (Linear-Quadratic Problelns)

Consider the n -dimensional linear system

$$
\dot{x}(t)=A x(t)+B u(t)
$$

where $A$ and $B$ are given matrices, and the quadratic cost

$$
x(T)^{\prime} Q l^{\prime} x(T)+\int_{0}^{T}\left(x(t)^{\prime} Q x(t)+u(t)^{\prime} R u(t)\right) d t
$$

where the matrices QT and $Q$ are symmetric positive semidefinite, and the rnatrix $R$ is symmetric positive definite (Appendix A defines positive definite and semidefinite matrices). The HJB equation is

$$
\begin{equation*}
\mathrm{O}=\min _{u \in \Re^{m}}\left[x^{\prime} Q x+u^{\prime} R u+\nabla_{t} V(t, x)+\nabla_{x} V(t, x)^{\prime}(A x+B u)\right] \tag{3.9}
\end{equation*}
$$

$$
V(T, x)=x^{\prime} Q_{T} x
$$

Let us try a solution of the form

$$
\text { Vet, } x)=x_{l} K(t) x, \quad K(t): n \times n \text { symmetric, }
$$

and see if we can solve the HJB equation. We have $\nabla_{x} V(t, x)=2 K(t) x$ and $\nabla_{t} V(t, x)={ }_{x l} \dot{K}(t) x$, where $\dot{K}(t)$ is the matrix with elements the first order derivatives of the elements of $K(t)$ with respect to time. By substituting these expressions in Eq. (3.9), we obtain

$$
\begin{equation*}
\mathrm{O}=\min _{u}\left[x \prime Q x+u^{\prime} R u+x ı \dot{K}(t) x+2 x^{\prime} K(t) A x+2 x^{\prime} K(t) B u\right] \tag{3.10}
\end{equation*}
$$

The minimum is attained at a $\boldsymbol{u}$ for which the gradient with respect to $\boldsymbol{u}$ is zero, that is,
or

$$
2 B^{I} K(t) x+2 R u=0
$$

$$
\begin{equation*}
u=R_{-}{ }^{l} B{ }^{I} K(t) x \tag{3.11}
\end{equation*}
$$

Substituting the minimizing value of $u$ in Eq. (3.10), we obtain

$$
0 \quad x^{\prime}\left(\dot{K}(t)+K(t) A+A I K(t)-K(t) B R_{-}{ }^{1} B{ }^{I} K(t)+Q\right) x, \quad \text { for all }(t, x)
$$

Therefore, in order for Vet, $x)=x_{l} K(t) x$ to solve the HJB equation, $K(t)$ must satisfy the following matrix differential equation (known as the continuous-time Riccati equation)

$$
\begin{equation*}
\dot{K}(t)=-K(t) A \quad A I K(t)+K(t) B R-^{l} B I^{I} K(t) \quad Q \tag{3.12}
\end{equation*}
$$

with the terminal condition

$$
\begin{equation*}
K(T)=Q T \tag{3.13}
\end{equation*}
$$

Reversing the argument, we see that if $K(t)$ is a solution of the Riccati equation (3.12) with the boundary condition (3.13), then Vet, $x$ ) $=x / K(t) x$ is a solution of the HJB equation. Thus, by using Prop. 3.2.1, we conclude that the optimal cost-to-go function is

$$
J^{*}(t, x)=x I K(t) x
$$

F'urthermore, in view of the expression derived for the control that minimizes in the right-hand side of the HJB equation [ef. Eq. (3.11)], an optimal policy is

$$
\mu^{*}(t, x)=-R^{-1} B^{\prime} K(t) x
$$

## THE PONTRYAGIN MINHMUM PRINCIPLE

In this section we discuss the continuous-time and the discrete-time versions of the Minimum Principle, starting with a DP-based informal argument.

### 3.3.1 An Informal Derivation Using the HJB Equation

Recall the HJB equation

$$
\begin{equation*}
\mathrm{O}=\min _{u E \bar{U}}\left[g(x, u)+\nabla_{t} J^{*}(t, x)+\nabla_{x} J^{*}(t, x)^{\prime} f(x, u)\right], \text { for all } t, x \tag{3.14}
\end{equation*}
$$

$$
\begin{equation*}
\left.J^{*}(T, x)=h e x\right), \quad \text { for all } x \tag{3.15}
\end{equation*}
$$

We argued that the optimal cost-to-go function $J^{\star}(t, x)$ satisfies this equation under some conditions. Furthermore, the sufficiency theorem of the preceding section suggests that if for a given initial state $x(0)$, the control trajectory $\left\{u^{*}(t)\right\} \mathrm{It} \mathrm{E}[0, \mathrm{~T}$ is optimal with corresponding state trajectory $\left\{x^{*}(t)\right.$ It $\left.\mathrm{E}[O, T]\right\}$, then for all $t \mathrm{E}[O, T]$,

$$
u^{*}(t) \quad \operatorname{argmin}\left[g E U\left(x^{*}(t), u\right)+\nabla_{x} J^{*}\left(t, x^{*}(t)\right)^{\prime} f\left(x^{*}(t), u\right)\right]
$$

Note that to obtain the optimal control trajectory via this equation, we do not need to know $\nabla x J^{*}$ at all values of $x$ and $t$; it is sufficient to know $\nabla x J^{*}$ at only one value of $x$ for each $t$, that is, to know only $\nabla x J^{*}\left(t, x^{*}(t)\right)$.

The Minimum Principle is basically the preceding Eq. (3.16). Its application is facilitated by streamlining the computation of $\nabla_{x} J^{*}\left(t, x^{*}(t)\right)$. It turns out that we can often calculate $\nabla_{x} J^{*}\left(t, x^{*}(t)\right)$ along the optimal state trajectory far more easily than we can solve the HJB equation. In particular, $\nabla x J^{*}\left(t, x^{*}(t)\right)$ satisfies a certain differential equation, called the adjo'int equation. We will derive this equation informally by differentiating the HJB equation (3.14). We first need the following lemma, which indicates how to differentiate functions involving minima.

Lernma 3.3.1: Let $F(t, x, u)$ be a continuously differentiable function of $t \mathrm{E} \Re, x \mathrm{E} \Re^{n}$, and $\boldsymbol{u} \mathrm{E} \Re^{m}$, and let $U$ be a convex subset of $\Re^{m}$. Assume that $J L^{*}(t, x)$ is a continuously differentiable function such that

$$
J 1^{*}(t, x)=\arg \min _{u E U} F(t, x, u), \quad \text { for all } t, x
$$

Then

$$
\begin{array}{ll}
\nabla_{\boldsymbol{t}}\left\{\min _{u \in U} F(t, x, u)\right\}=\nabla t F\left(t, x, \mu^{*}(t, x)\right), & \text { for all } \mathrm{t}, x \\
\nabla x\left\{\min _{u \in U} F(t, x, u)\right\}=\nabla x F\left(t, x, \mu^{*}(t, x)\right), & \text { for all } t, x
\end{array}
$$

[Note: On the left-hand side, $\nabla_{\boldsymbol{t}}\{\cdot\}$ and $\nabla_{\boldsymbol{x}}\{\cdot\}$ denote the gradients of the function $G(t, x)=\min _{u \in U} F(t, x, u)$ with respect to $t$ and $x$, respectively. On the right-hand side, $\nabla_{t}$ and $\nabla_{x}$ denote the vectors of partial derivatives of $F$ with respect to $t$ and $x$, respectively, evaluated at $\left(t, x, \mu^{*}(t, x)\right)$.]

Proof: For notational simplicity, denote $y=(t, x), F(y, u)=F(t, x, u)$, and $\mu^{*}(y)=\mu^{*}(t, x)$. Since $\min _{u \in U} F(y, u)=F\left(y, ; L^{*}(y)\right)$,

$$
\nabla\left\{\min _{u \in U} \mathrm{~F}(\mathrm{y}, \mathrm{u})\right\}=\nabla_{y} F\left(y, \mu^{*}(y)\right)+\nabla \mu^{*}(y) \nabla_{u} F\left(y, \mu^{*}(y)\right) .
$$

We will prove the result by showing that the second term in the righthand side above is zero. This is true when $U=\Re^{m}$, because then $\mu^{*}(y)$ is an unconstrained minimum of $F(y, u)$ and $\nabla_{u} F\left(y, \mu^{*}(y)\right)=0$. More generally, for every fixed $y$, we have

$$
\left(u-\mu^{*}(y)\right)^{\prime} \nabla_{u} F\left(y, \mu^{*}(y)\right) \geq 0, \quad \text { for all } u \mathrm{E} U,
$$

[see Eq. (B.2) in Appendix B]. Now by Taylor's Theorem, we have that when $y$ changes to $y+\Delta y$, the minimizing $\mu^{*}(y)$ changes from $\mu^{*}(y)$ to some vector $\mu^{*}(y+\Delta y)=\mu^{*}(y)+\nabla \mu^{*}(y)^{\prime} \Delta y+o(\|\Delta y\|)$ of $U$, so

$$
\left(\nabla \mu^{*}(y)^{\prime} \Delta y+o(\|\Delta y\|)\right)^{\prime} \nabla_{u} F\left(y, \mu^{*}(y)\right) \geq 0, \quad \text { for all } \Delta y
$$

implying that

$$
\nabla \mu^{*}(y) \nabla u F\left(y, \mu^{*}(y)\right)=0
$$

## Q.E.D.

Consider the HJB equation (3.14), and for any $(t, x)$, suppose that $\mu^{*}(t, x)$ is a control attaining the minimum in the right-hand side. We make the restrictive assumptions that $U$ is a convex set, and that $\mu^{*}(t, x)$ is continuously differentiable in $(t, x)$, so that we can use Lemma 3.3.1. (We note, however, that alternative derivations of the Minimum Principle do not require these assumptions; see Section 3.3.2.)

We differentiate both sides of the fIJB equation with respect to $x$ and with respect to $t$. In particular, we set to zero the gradient with respect to $x$ and t of the function

$$
g\left(x, \mu^{*}(t, x)\right)+\nabla_{t} J^{*}(t, x)+\nabla_{x} J^{*}(t, x)^{\prime} f\left(x, \mu^{*}(t, x)\right)
$$

and we rely on Lemma 3.3.1 to disregard the terms involving the derivatives of $\mu^{*}(t, x)$ with respect to $t$ and $x$. We obtain for all $(t, x)$,

$$
\begin{align*}
& O=V x g\left(X, f L^{*}(t, X)\right)++\nabla_{x t}^{2} J^{*}(t, x)+\nabla_{x x}^{2} J^{*}(t, x) f\left(x, \mu^{*}(t, x)\right)  \tag{3.17}\\
&+\nabla x f\left(x, \mu^{*}(t, \mathrm{x})\right) \mathrm{V}_{x} J^{*}(t, x), \\
& O=\nabla_{t t}^{2} J^{*}(t, x)+\nabla_{x t}^{2} J^{*}(t, x)^{\prime} f\left(x, \mu^{*}(t, x)\right) \tag{3.18}
\end{align*}
$$

where $\nabla_{x} f\left(x, \mu^{*}(t, x)\right)$ is the matrix

$$
\nabla_{x} f=\left(\begin{array}{cc}
\frac{\partial f_{1}}{\partial x_{1}} & \underline{B}=\left(\begin{array}{ll} 
\\
\vdots & \\
\frac{\partial f_{1}}{\partial x_{n}} & \frac{\partial f_{n}}{\partial x_{n}}
\end{array}\right) .
\end{array}\right.
$$

with the partial derivatives evaluated at the argument $\left(x, \mu^{*}(t, x)\right)$.
The above equations hold for all $(t, x)$. Let us specialize them along an optimal state and control trajectory $\left\{\left(x^{*}(t), u^{*}(t)\right) \backslash \mathrm{t}\right.$ E $\left.[0, T]\right\}$, where $u^{*}(t)=\mu^{*}\left(t, x^{*}(t)\right)$ for all E (O,Tl. We have for all t ,

$$
\dot{x}^{*}(t)=\mathrm{J}\left(\mathrm{x}^{*}(t), u^{*}(t)\right)
$$

so that the term

$$
\nabla_{x t}^{2} J^{*}\left(t, x^{*}(t)\right)+\nabla_{x x}^{2} J^{*}\left(t, x^{*}(t)\right) f\left(x^{*}(t), u^{*}(t)\right)
$$

in Eq. (3.17) is equal to the following total derivative with respect to $t$

$$
\frac{d}{d t}\left(\nabla_{x} J^{*}\left(t, x^{*}(t)\right)\right)
$$

Similarly, the term

$$
\nabla_{l t}^{2} J^{*}\left(t, x^{*}(t)\right)+\nabla_{x t}^{2} J^{*}\left(t, x^{*}(t)\right)^{\prime} f\left(x^{*}(t), u^{*}(t)\right)
$$

in Eq. (3.18) is equal to the total derivative

$$
\frac{d}{d t}\left(\nabla \mathrm{t}{ }^{\star}\left(t, x^{*}(t)\right)\right)
$$

Thus, by denoting

$$
\begin{equation*}
p(t)=V x J^{*}\left(t, x^{*}(t)\right) \tag{3.19}
\end{equation*}
$$

Eq. (3.17) becomes

$$
\begin{equation*}
\dot{p}(t)=-V x f\left(x^{*}(t), u^{*}(t)\right) p(t)-\nabla_{x} g\left(x^{*}(t), u^{*}(t)\right) \tag{3.21}
\end{equation*}
$$

and Eq. (3.18) becomes

$$
\dot{p}_{0}(t)=\mathrm{O}
$$

or equivalently,

$$
\begin{equation*}
p o(t)=\text { constant }, \quad \text { for all } t \mathrm{E}[0, T] . \tag{3.22}
\end{equation*}
$$

Equation (3.21) is a system of $n$ first order differential equations known as the adjoint equation. From the boundary condition

$$
J^{*}(T, x)=h(x), \quad \text { for all } x
$$

we have, by differentiation with respect to $x$, the relation $V x J^{*}(T, x)$ $\nabla h(x)$, and by using the definition $V x J^{*}\left(t, x^{*}(t)\right)=p(t)$, we obtain

$$
\begin{equation*}
p(T)=V h\left(x^{*}(T)\right) . \tag{3.23}
\end{equation*}
$$

Thus, we have a terminal boundary condition for the adjoint equation (3.21).

To summarize, along optimal state and control trajectories $X^{*}(t)$, $u^{*}(t), t \mathrm{E}[0, T]$, the adjoint equation (3.21) holds together with the boundary condition (3.23), while Eq. (3.16) and the definition of $p(t)$ imply that $u^{*}(t)$ satisfies

$$
\begin{equation*}
u^{*}(t)=\arg \min _{u \in U}\left[g\left(x^{*}(t), u\right)+p(t)^{\prime} f\left(x^{*}(t), u\right)\right] \tag{0,T}
\end{equation*}
$$

## Harniltonian Formulation

Motivated by the condition (3.24), we introduce the Hamiltonian function mapping triplets $(x, u, p) \mathrm{E} \Re^{n} \times \Re^{m} \times \Re^{n}$ to real numbers and given by

$$
H(x, u, p)=g(x, u)+p^{\prime} f(x, u)
$$

Note that both the system and the adjoint equations can be compactly written in terms of the Hamiltonian as

$$
\dot{x}^{*}(t) \quad \nabla_{p} H\left(x^{*}(t), u^{*}(t), p(t)\right), \quad \dot{p}(t)=-V x f I\left(x^{*}(t), u^{*}(t), p(t)\right)
$$

We state the Minimum Principle in terms of the Hamiltonian function.

Proposition 3.3.1: (Minimum Principle) Let $\left\{u^{*}(t)\right.$ It E [0, TJ $\}$ be an optimal control trajectory and let $\left\{x^{*}(t) \mid t \mathrm{E}[0, T]\right\}$ be the corresponding state trajectory, i.e.,

$$
\dot{x}^{*}(t)=f\left(x^{*}(t), u^{*}(t)\right), \quad x^{*}(0)=x(0): \text { given. }
$$

Let also $p(t)$ be the solution of the adjoint equation

$$
\dot{p}(t)=-\nabla_{x} H\left(x^{*}(t), u^{*}(t), p(t)\right)
$$

with the boundary condition

$$
p(T)=V h\left(x^{*}(T)\right)
$$

where $h(\cdot)$ is the terminal cost function. Then, for all $t \mathrm{E}[0, T]$,

$$
u^{*}(t)=\underset{u \in U}{\operatorname{argmin}} H\left(x^{*}(t), u, p(t)\right)
$$

Furthermore, there is a constant 0 such that

$$
H\left(x^{*}(t), u^{*}(t), p(t)\right)=0, \quad \text { for all } t \mathrm{E}[0, T]
$$

All the assertions of the Minimum Principle have been (informally) derived earlier except for the last assertion. To see why the Hamiltonian function is constant for $t \mathrm{E}[0,7]$ along the optimal state and control trajectories, note that by Eqs. (3.14), (3.19), and (3.20), we have for all $t \mathrm{E}[0, T]$

$$
\left.H\left(x^{*}(t), u^{*}(t), p(t)\right)=-V t J^{*}\left(t, x^{*}(t)\right)=-p o e t\right)
$$

and poet) is constant by Eq. (3.22). We should note here that the Hamiltonian function need not be constant along the optimal trajectory if the system and cost are not time-independent, contrary to our assumption thus far (see Section 3.4.4).

It is important to note that the Minimum Principle provides neces sary optimality conditions, so all optimal control trajectories satisfy these conditions, but if a control trajectory satisfies these conditions, it is not necessarily optimal. Further analysis is needed to guarantee optimality. One method that often works is to prove that an optirnal control trajectory exists, and to verify that there is only one control trajectory satisfying the conditions of the Minimum Principle (or that all control trajectories satisfying these conditions have equal cost). Another possibility to conclude optimality arises when the system function $f$ is linear in $(x, u)$, the
constraint set $U$ is convex, and the cost functions hand $g$ are convex. Then it can be shown that the conditions of the Minimum Principle are both necessary and sufficient for optimality.

The Minimum Principle can often be used as the basis of a numerical solution. One possibility is the two-point boundary problem method. In this method, we use the minimum condition

$$
u^{*}(t)=\arg \min _{u \in U} H\left(x^{*}(t), u, p(t)\right)
$$

to express $u^{*}(t)$ in terms of $x^{*}(t)$ and pet). We then substitute the result into the system and the adjoint equations, to obtain a set of $2 n$ first order diflerential equations in the components of $x^{*}(t)$ and $\left.p e t\right)$. These equations can be solved using the split boundary conditions

$$
\left.x^{*}(0)=x(O), \quad \text { pe } T\right)=V h\left(x^{*}(T)\right)
$$

The number of boundary conditions (which is $2 n$ ) is equal to the number of differential equations, so that we generally expect to be able to solve these differential equations numerically (although in practice this may not be simple).

Using the Minimum Principle to obtain an analytical solution is possible in many interesting problems, but typically requires considerable creativity. We give some simple examples.

## wxample 3.3.1 (Calculus of Variations Continued)

Consider the problem of finding the curve of minimum length from a point $(0, \boldsymbol{\alpha})$ to the line $\{(T, y) \operatorname{ly} \in \Re\}$. In Section 3.1 , we formulated this problem as the problem of finding an optimal control trajectory $\left\{\left.u(t)\right|_{t \mathrm{E}[0, T \eta}\right.$ that minimizes

$$
\int_{0}^{T} \sqrt{1+(u(t))^{2}} d t
$$

subject to

$$
\dot{x}(t)=u(t), \quad \mathrm{x}(\mathrm{O})=\alpha .
$$

Let us apply the preceding necessary conditions. The Hamiltonian is

$$
H(x, u, p)=\sqrt{ } 1 \pm \underline{u}^{2}+p u
$$

and the adjoint equation is

$$
\dot{p}(t)=0, \quad p e T)=0
$$

It follows that

$$
\text { pet })=0, \quad \text { for all } t \mathrm{E}[0, T]
$$

so minimization of the Hamiltonian gives

$$
u^{*}(t)=\operatorname{argmin} \sqrt[u \in \Re]{ } \sqrt{1+u^{2}}=0, \quad \text { for all } t \mathrm{E}[0, T] .
$$

Therefore we have $\dot{x}^{*}(t)=$ © for all $t$, which implies that $x^{*}(t)$ is constant. Using the initial condition $x^{*}(0)=\alpha$, it follows that

$$
x^{*}(t)=\alpha, \quad \text { for all } t \mathrm{E}[0, \mathrm{~T} J
$$

We thus obtain the (a priori obvious) optimal solution, which is the horizontal line passing through $(0, \alpha)$. Note that since the Minimum Principle is only a necessary condition for optimality, it does not guarantee that the horizontal line solution is optimal. For such a guarantee, we should invoke the linearity of the system function, and the convexity of the cost function. As rnentioned (but not proved) earlier, under these conditions, the Minimum Principle is both necessary and sufficient for optimality.

## Example 3.3.2 (Resource Allocation Continued)

Consider the optimal production problem (Example 3.1.2). We want to maximize

$$
\int_{0}^{T}(1-u(t)) x(i) d t
$$

subject to

$$
\begin{gathered}
\mathrm{O} \leq u(t) \leq 1, \quad \text { for all } t \mathrm{E}[0, T]^{\prime} \\
\dot{x}(t)=\gamma u(t) x(t), \quad x(O)>0: \text { given. }
\end{gathered}
$$

The Hamiltonian is

$$
H(x, u, p)=(1-u) x+p \gamma u x
$$

The adjoint equation is

$$
\begin{aligned}
\dot{p}(t) & =-\gamma u^{*}(t) p(t)-1+u^{*}(t) \\
p e T) & =0
\end{aligned}
$$

Maximization of the Hamiltonian over $u \mathrm{E}[0,1]$ yields

$$
\boldsymbol{u}^{*}(t)= \begin{cases}0 & \text { if } p(t)<\frac{1}{\gamma} \\ 1 & \text { if } p(t) \geq \frac{1}{\gamma}\end{cases}
$$

Since $p e T)=0$, for $t$ close to $T$ we will have $p(t)<1 / \gamma$ and $u^{*}(t)=0$. Therefore, for $t$ near $T$ the adjoint equation has the form $\dot{p}(t)=-]$ and $p(t)$ has the form shown in Fig. 3.3.1.


Figure 3.3.2 Form of the adjoint variable $p(t)$ and the optimal control in the resource allocation example.
Figure 3.3.1 Form of the adjoint variable $p(t)$ for $t$ near $T$ in the resource allocation example.

Thus, near $t=T, p(t)$ decreases with slope -1 . For $t=T-1 / \gamma, p(t)$ is equal to $1 / \gamma$, so $u^{*}(t)$ changes to $u^{*}(t)=1$. It follows that for $t<T-1 / \gamma$, the adjoint equation is

$$
\dot{p}(t)=-I P(t)
$$

or

$$
p(t) \quad e^{-\gamma t} . \text { constant. }
$$

Piecing together $p(t)$ for $t$ greater and less than $T-1 / \gamma$, we obtain the form shown in Fig. 3.3.2 for $p(t)$ and $u^{*}(t)$. Note that if $T<1 / \gamma$, the optimal control is $u^{*}(t)=0$ for all $t \mathrm{E}[0, \mathrm{TJ}$; that is, for a short enough horizon, it does not pay to reinvest at any time.

## Exarnple 3.3.3 (A Linear-Quadratic Problem)

Consider the one-dimensional linear system

$$
\dot{x}(t)=a x(t)+b u(t)
$$

where $a$ and $b$ are given scalars. We want to find an optirnal control over a given interval $[0, T]$ that minimizes the quadratic cost

$$
\frac{1}{2} q \cdot(x(T))^{2}+\frac{1}{2} \int_{0}^{T}(u(l))^{2} d t
$$

where $q$ is a given positive scalar. There are no constraints on the control, so we have a special case of the linear-quadratic problem of Example 3.2.2. We will solve this problem via the Minimum Principle.

The Hamiltonian here is

$$
H(x, u, p)=\frac{1}{2} u^{2}+p(a x+b u)
$$

and the adjoint equation is

$$
\dot{p}(t)=-a p(t),
$$

with the terminal condition

$$
p(T) \quad q x^{*}(T)
$$

The optimal control is obtained by minimizing the Hamiltonian with respect to $u$, yielding

$$
\begin{equation*}
u^{*}(t)=\arg \min _{u}\left[\frac{1}{2} u^{2}+p(t)\left(a x^{*}(t)+b u\right)\right] \quad-b p(t) \tag{3.25}
\end{equation*}
$$

We will extract the optimal solution from these conditions using two different approaches.

In the first approach, we solve the two-point boundary value problem discussed following Prop. 3.3.1. In particular, by eliminating the control from the system equation using Eq. (3.25), we obtain

$$
\dot{x}^{*}(t)=a x^{*}(t)-b^{2} p(t)
$$

Also, from the adjoint equation, we see that

$$
p(t)=e^{-a t} \xi, \quad \text { for all } t \mathrm{E}[0, \mathrm{~T}]^{\prime}
$$

where $\xi=p(O)$ is an unknown parameter. The last two equations yield

$$
\begin{equation*}
\dot{x}^{*}(t)=a x^{*}(t)-b^{2} e^{-a t} \xi . \tag{3.26}
\end{equation*}
$$

This differential equation, together with the given initial condition $x^{*}(0)$ $x(O)$ and the terminal condition

$$
x *(T)=\frac{e^{-a T} \xi}{q}
$$

(which is the terminal condition for the adjoint equation) can be solved for the unknown variable $\xi$. In particular, it can be verified that the solution of the differential equation (3.26) is given by

$$
x^{*}(t)=x^{(0)} \mathrm{e}^{a t-+} \frac{b^{2} \xi}{2 a}\left(-a t-e^{a t}\right)
$$

and $\xi$ can be obtained from the last two relations. Given $\xi$, we obtain $p(t)=$ $e^{-a t} \xi$ and from $p(t)$, we can then determine the optimal control trajectory as $\boldsymbol{u}^{*}(t)=-b p(t), t \mathrm{E}[0, T]$ [ef. Eq. (3.25)].

In the second approach, we basically derive the Riccati equation encountered in Example 3.2.2. In particular, we hypothesize a linear relation between $x^{*}(t)$ and $p(t)$, that is,

$$
K(t) x^{*}(t)=p(t), \quad \text { for all } t \mathrm{E}[0, T]^{\prime}
$$

and we show that $K(t)$ can be obtained by solving the Riccati equation. Indeed, from Eq. (3.25) we have

$$
\boldsymbol{u}^{*}(t)=-b K(t) x^{*}(t)
$$

which by substitution in the system equation, yields

$$
\dot{\boldsymbol{x}}^{*}(t)=\left(\begin{array}{ll}
a & b^{2} K(t)
\end{array}\right) x^{*}(t)
$$

By differentiating the equation $K(t) x^{*}(t)=p(t)$ and by also using the adjoint equation, we obtain

$$
\dot{K}(t) x^{*}(t)+K(t) \dot{x}^{*}(t)=\dot{p}(t)=-a p(t)=-a K(t) x *(t)
$$

By combining the last two relations, we have

$$
k(t) x *(t)-+K(t)\left(a-b^{2} K(t)\right) x *(t)=-a K(t) x *(t),
$$

from which we see that $K(t)$ should satisfy

$$
\dot{K}(t)=-2 a K(t)+b^{2}(K(t))^{2}
$$

This is the Riccati equation of Example 3.2.2, specialized to the problem of the present example. This equation can be solved using the terminal condition

$$
K(T)=q
$$

which is implied by the terminal condition $p(T)=q x^{*}(T)$ for the adjoint equation. Once $K(t)$ is known, the optimal control is obtained in the closedloop form $u^{*}(t)=-b K(t) x^{*}(t)$. By reversing the preceding arguments, this control can then be shown to satisfy all the conditions of the Minimum Principle.
3.3.2 A Derivation Based on Variational JIdeas

In this subsection we outline an alternative and more rigorous proof of the Minimum Principle. This proof is primarily directed towards the advanced reader, and is based on making small variations in the optimal trajectory and comparing it with neighboring trajectories.

For convenience, we restrict attention to the case where the cost is

$$
h(x(T))
$$

The more general cost

$$
\begin{equation*}
h(x(T))+\int_{0}^{T} g(x(t), u(/)) d t \tag{3.27}
\end{equation*}
$$

can be reformulated as a terminal cost by introducing a new state variable $y$ and the additional differential equation

$$
\begin{equation*}
\ddot{y}(t)=g(x(t), u(t)) \tag{3.28}
\end{equation*}
$$

The cost then becomes

$$
\begin{equation*}
h(x(T))+y(T) \tag{3.29}
\end{equation*}
$$

and the Minimum Principle corresponding to this terminal cost yields the Minimum Principle for the general cost (3.27).

We introduce some assumptions:

Convexity Assumption: For every state $x$ the set

$$
D=\{f(x, u) \mid u \in U\}
$$

is convex

The convexity assumption is satisfied if $U$ is a convex set and $f$ is linear in $u$ [and 9 is linear in $\boldsymbol{u}$ in the case where there is an integral cost of the form (3.27), which is reformulated as a terminal cost by using the additional state variable $y$ of Eq. (3.28)]. Thus the convexity assumption is quite restrictive. However, the Minimum Principle typically holds without the convexity assumption, because even when the set $D=\{f(x, u) \mid u \mathrm{E} U\}$ is nonconvex, any vector in the convex hull of $D$ can be generated by quick alternation between vectors from $D$ (for an example, see Exercise 3.10). This involves the complicated mathematieal concept of randomized or relaxed controls and will not be discussed further.

Regularity Assumption: Let $u(t)$ and $u^{*}(t), t$ E [ $\left.0 T T\right]$, be any two admissible control trajectories and let $\left\{x^{*}(t) \mid\right.$ t E $[0$, In be the state trajectory corfesponding to $u^{*}(t)$. For any $\epsilon \mathrm{E}[0,1]^{\prime}$ the solution $\left\{x_{\epsilon}(t)\right.$ It $\mathrm{E}[0, \mathrm{Il}$ of the system

$$
\begin{equation*}
\dot{x}_{\epsilon}(t)=(1-\epsilon) f\left(x_{\epsilon}(t), u^{*}(t)\right)+\epsilon f\left(x_{\epsilon}(t), u(t)\right), \tag{3.30}
\end{equation*}
$$

with $x_{\epsilon}(0)=x^{*}(0)$, satisfies

$$
\begin{equation*}
x E(t)=x^{*}(t)+\epsilon \xi(t)+o(\epsilon) \tag{3.31}
\end{equation*}
$$

where $\{\xi(t)$ It $\mathrm{E}[0, \mathrm{~T}]\}$ is the solution of the linear differential system

$$
\begin{equation*}
\dot{\xi}(t) \quad \nabla_{x} f\left(x^{*}(t), u^{*}(t)\right) \xi(t)+f\left(x^{*}(t), u(t)\right)-f\left(x^{*}(t), u^{*}(t)\right) \tag{3.32}
\end{equation*}
$$

with initial condition $\xi(0) \quad 0$.

The regularity assumption "typically" holds because from Eq. (3.30) we have

$$
\begin{aligned}
\dot{x}_{\epsilon}(t)-\dot{x}^{*}(t)= & f\left(x E(t), u^{*}(t)\right)-f\left(x^{*}(t), u^{*}(t)\right) \\
& +\epsilon\left(f\left(x_{\epsilon}(t), u(t)\right) \quad f\left(x_{\epsilon}(t), u^{*}(t)\right)\right),
\end{aligned}
$$

so from a first order Taylor series expansion we obtain

$$
\begin{aligned}
\delta \dot{x}(t)= & \nabla f\left(x^{*}(t), u^{*}(t)\right)^{\prime} \delta x(t)+\mathrm{O}(\|\mathrm{Jx}(\mathrm{t})\|) \\
& +\epsilon\left(f\left(x_{\epsilon}(t), u(t)\right)-f\left(x_{\epsilon}(t), u^{*}(t)\right)\right)
\end{aligned}
$$

where

$$
\delta x(t)=x €(t)-x^{*}(t) .
$$

Dividing by $\epsilon$ and taking the limit as $\epsilon \rightarrow 0$, we see that the function

$$
\xi(t)=\lim _{\epsilon \rightarrow 0} \delta x(t) / \epsilon, \quad t \mathrm{E}[0, \mathrm{~T}]
$$

should "typically" solve the linear system of differential equations (3.32), while satisfying Eq. (3.31).

In fact, if the system is linear of the form

$$
\dot{x}(t)=A x(t)+B u(t)
$$

where $A$ a,nd $B$ are given matrices, it can be shown that the regularity assumption is satisfied. To see this, note that Eqs. (3.30) and (3.32) take the forms

$$
\dot{x}_{\epsilon}(t)=A x €(t)+B u^{*}(t)+\epsilon B\left(u(t) \quad u^{*}(t)\right),
$$

and

$$
\dot{\xi}(t)=A \xi(t)+B\left(u(t)-\mathbf{'}^{*}(t)\right),
$$

respectively. Thus, taking into account the initial conditions $x_{\epsilon}(0) \quad x^{*}(0)$ and $\xi(0)=0$, we see that

$$
x_{\epsilon}(t)=x^{*}(t)+\epsilon \xi(t), \quad t \mathrm{E}[0, \mathrm{~T}]
$$

so the regularity condition (3.31) is satisfied.
We now prove the Minimum Principle assuming the convaxity and regularity assumptions above. Suppose that $\left\{q^{n}(t) \mid t \in[0, \Pi\right.$ is an optimal control trajectory, and let $\left\{X^{*}(t) \operatorname{It~}_{\mathrm{E}}[0, \mathrm{l} \|\right.$ be the corresponding state trajectory. Then for any other admissible control trajectory $\{u(t)\}$ $t \mathrm{E}[0, \mathrm{~T}]\}$ and any $\epsilon \mathrm{E}[0,1]$, the convexity assumption guarantees that for each $t$, there exists a control $\bar{u}(t) \mathrm{E} U$ such that

$$
f(x €(t), u(t))=(1-\epsilon) f\left(x_{\epsilon}(t), u^{*}(t)\right)+\epsilon f\left(x_{\epsilon}(t), u(t)\right)
$$

Thus, the state trajectory $\{x \in(t) \mid t \mathrm{E}[0, \mathrm{~T}]\}$ of $\mathrm{q} q$. (3.30) corresponcls to the admissible control trafectory $\{\bar{u}(\mathrm{t}) \mid t \in[0,1 \mathrm{l}$. Hence, using the optimality of $\left\{x^{*}(t)\right.$ It $E[0, \ \cap$ and the regularity assumption, we have

$$
\begin{aligned}
h\left(x^{*}(T)\right) \leq & h\left(x_{\epsilon}^{*}(T)\right) \\
& h\left(x^{*}(T)+\epsilon \xi(T)+. o(\epsilon)\right) \\
= & h\left(x^{*}(T)\right)+\epsilon \nabla h\left(x^{*}(T)\right)^{\prime} \xi(T)+o(\epsilon),
\end{aligned}
$$

which implies that

$$
\begin{equation*}
\nabla h\left(x^{*}(T)\right)^{\prime} \xi(T) \geq 0 \tag{3.33}
\end{equation*}
$$

Using a standard result in the theory of linear difFerential equations (see e.g. [CoL65]), the solution of the linear differential system (3.32) can be written in closed form as

$$
\begin{equation*}
\xi(t)=\Phi(t, \tau) \xi(\tau)+\int_{\tau}^{t} \Phi(t, \tau)\left(f\left(x^{*}(\tau), u(\tau)\right) \quad f\left(x^{*}(\tau), u^{*}(\tau)\right)\right) d \tau \tag{3.34}
\end{equation*}
$$

where the square matrix $\Phi$ satisfies for all $t$ and $T$,

$$
\begin{gather*}
\frac{\partial \Phi(t, \tau)}{\partial \tau}-\Phi(t, \tau) \nabla_{x} f\left(x^{*}(\tau), u^{*}(\tau)\right)^{\prime}  \tag{3.35}\\
\Phi(t, t) \quad I
\end{gather*}
$$

Since $\xi(0)=0$, we have from Eq. (3.34),

$$
\begin{equation*}
\xi(T)=\int_{0}^{T} \Phi(T, t)\left(f\left(x^{\prime}(t), u(t)\right)-f\left(x^{*}(t), u^{*}(t)\right)\right) d t \tag{3.36}
\end{equation*}
$$

Define

$$
\begin{equation*}
\left.p e T)=\nabla h\left(x^{*}(T)\right), \quad \text { pet }\right)=\Phi(T, t)^{\prime} p(T), \quad t \in(0, T] \tag{3.37}
\end{equation*}
$$

By differentiating with respect to $t$, we obtain

$$
\left.p^{\prime(t)}=\frac{\partial \Phi(T, t)}{\partial t} \frac{t}{2} p e T\right)
$$

Combining this equation with Eqs. (3.35) and (3.37), we see that $p(t)$ is generated by the dIfferentIal equatiOn

$$
\dot{p}(t)=-\nabla_{x} f\left(x^{*}(t), u^{*}(t)\right) p(t)
$$

with the terminal condition

$$
p e T)=\nabla h\left(x^{*}(T)\right)
$$

This is the adjoint. equation corresponding to $\left\{\left(x^{*}(t), u^{*}(t)\right) I t \mathrm{E}[0, T n\right.$.
Now, to obtam the Mímilnum Principle, we note that from Eqs. (3.33) $(3.36)$, and (3.37) we have

$$
\begin{align*}
& \leq p(T)^{\prime} \xi(T) \\
& \left.\left.=p(T)^{\prime} \int_{0}^{T} \Phi(T, t)\left(f\left(x^{*}(t), u(t)\right)-f(x \ll t), u \ll t\right)\right)\right) e l t  \tag{3.38}\\
& =\int_{0}^{T} p(t)^{\prime}\left(f\left(x^{*}(t), u(t)\right)-f\left(x^{\prime}(t), u^{*}(t)\right)\right) e l t
\end{align*}
$$

from which it can be shown that for all $t$ at which $\boldsymbol{u}^{*(.)}$ is contI' nuous, we
lave

$$
\begin{equation*}
p(t)^{\prime} f\left(x^{*}(t), u^{*}(t)\right) \leq p(t)^{\prime} f\left(x^{*}(t), u\right), \quad \text { for all } u \mathrm{E} U . \tag{3.39}
\end{equation*}
$$

Indeed, if for some $\hat{u} \mathrm{E} U$ and to $\mathrm{E}[0, T)$, we have

$$
p(t o)^{\prime} f\left(x^{*}\left(t_{0}\right), u^{*}\left(t_{0}\right)\right)>p(t o)^{\prime} f\left(x^{*}(t o), \hat{\boldsymbol{u}}\right)
$$

while $\left\{\left.u^{*}(t)\right|_{t \in[0, T n}\right.$ is continuous at $t o$, we would also have

$$
p(t)^{\prime} f\left(x^{*}(t), u^{*}(t)\right)>p(t)^{\prime} f\left(x^{*}(t), \hat{u}\right)
$$

for all $t$ in some nontrivial interval 1 containing to. By taking

$$
u(t)= \begin{cases}\hat{u} & \text { for } t \mathrm{E} 1 \\ u^{*}(t) & \text { for } t \notin 1\end{cases}
$$

we would then obtain a contradiction of Eq. (3.38)
We have thus proved the Minimum Principle (3.39) under the convexity and regularity assumptions, and the assumption that there is only a terminal cost $h(x(T))$. We have also seen that in the case where the constraint set $U$ is convex and the system is linear, the convexity and regularity assumptions are satisfied. To prove the Minimum Principle for the more general integral cost function (3.27), we can apply the preceding development to the system of differential equations $\dot{x}=f(x, u)$ augmented by the additional Eq. (3.28) and the equivalent terminal cost (3.29). The corresponding convexity and regularity assumptions are automatically satisfied if the constraint set $U$ is convex and the system function $f(x, u)$ as well as the cost function $g(x, u)$ are linear. This is necessary in order to maintain the linearity of the augmented system, thereby maintaining the validity of the regularity assumption.

### 3.3.3 Minimum Principle for Discrete-Tinne Problems

In this subsection we briefly derive a version of the Minimum Principle for discrete-time deterministic optimal control problems. Interestingly, it is essential to make some convexity assumptions in order for the Minimum Principle to hold. For continuous-time problems these convexity assumptions are typically not needed, because, as mentioned earlier, the differential system can generate any $\dot{x}(t)$ in the convex hull of the set of possible vectors $f(x(t), u(t))$ by quick alternation between different controls (see for example Exercise 3.10).

Suppose that we want to find a control sequence ( $\left.u_{0}, \prime / / \prime, \cdots, u_{N-1}\right)$ and a corresponding state sequence ( $X O, X l, \ldots, x_{N}$ ), which minimize

$$
J(u)=g N(; r: N)+\sum_{k=O}^{N-I} g_{k}\left(x_{k}, \cdot I l k\right)^{\prime}
$$

subject to the discrete-time system constraints

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}\right), \quad k=0, \ldots, N-1, \quad x o: \text { given }
$$

and the control constraints

$$
u_{k} \in U_{k} \subset \Re^{m}, \quad k=0, \ldots, N
$$

We first develop an expression for the gradient $\nabla J\left(u_{0}, \ldots, u_{N-1}\right)$. We have, using the chain rule,

$$
\begin{aligned}
\nabla \cup N-1 J\left(u_{0}, \ldots, u_{N-1}\right)= & \nabla_{u_{N-1}}\left(9 N\left(f N-1\left(I: N-I, u_{N-1}\right)\right)\right. \\
& \left.\left.+g N-] \quad, u_{N-1}\right)\right) \\
= & \nabla_{u_{N-1}} f_{N-1} . \nabla g_{N}+\nabla_{u_{N-1}} g_{N-1}
\end{aligned}
$$

where all gradients are evaluated along the control trajectory (uo, ..., UN-1) and the corresponding state trajectory. Similarly, for all $k$,

$$
\begin{align*}
\nabla_{u_{k}} J\left(u_{0}, \ldots, u_{N-1}\right)= & \nabla_{u_{k}} f_{k} \cdot \nabla_{x_{k+1}} f_{k+1} \cdots \nabla_{x_{N-1}} f_{N-1} \cdot \nabla g_{N} \\
& +\nabla_{u_{k}} f_{k} \cdot \nabla_{x_{k+1}} f_{k+1} \cdots \nabla_{x_{N-2}} f_{N-2} \cdot \nabla_{x_{N-1}} g_{N-1} \\
& +\nabla_{u_{k}} f_{k} \cdot \nabla_{x_{k+1}} g_{k+1} \\
& +\nabla_{u_{k}} g_{k} \tag{3.40}
\end{align*}
$$

which can be written in the form

$$
\nabla_{u_{k}} J\left(u_{\ell}, \ldots, u_{N-1}\right)=\nabla_{u_{k}} f_{k} \cdot p_{k+1}+\nabla_{u_{k}} g_{k}
$$

for an appropriate vector $P k+l$, or

$$
\begin{equation*}
\nabla_{u_{k}} J\left(u_{0}, \ldots, u_{N-1}\right)=\nabla_{u_{k}} H_{k}\left(x_{k}, u_{k}, p_{k+1}\right) \tag{3.41}
\end{equation*}
$$

where $11_{k}$ is the Hamiltonian function defined by

$$
H_{k}\left(x_{k}, u_{k}, p_{k+1}\right)=g_{k}\left(x_{k}, u_{k}\right)+p_{k+1}^{\prime} f_{k}\left(x_{k}, u_{k}\right)
$$

It can be seen from Eq. (3.40) that the vectors $P k+l$ are generated backwards by the discrete-time adjoint equation

$$
p_{k}=\nabla_{x_{k}} f_{k} \cdot p_{k+1}+\nabla_{x_{k}} g_{k}, \quad k=1, \ldots, \mathbf{N}-1
$$

with terminal condition

$$
p_{N}=\nabla g_{N}
$$

We will assume that the constraint sets $U k$ are convex, so that we can apply the optimality condition

$$
{ }_{\mathrm{k}=\mathrm{O}}^{\mathrm{N}-\mathbf{I}} \nabla_{u_{k}} J\left(u_{0}^{*}, \ldots u_{N-1}^{*}\right)^{\prime}\left(u_{k} \quad u_{k}^{*}\right) \geq 0
$$

for all feasible (uo, ..., $\boldsymbol{u}_{\boldsymbol{N}-1}$ ) (see Appendix B). This condition can be decomposed into the N conditions
$\nabla_{u_{k}} J\left(u_{0}^{*}, \ldots, u_{N-1}^{*}\right)^{\prime}\left(u_{k}-u_{k}^{*}\right) \geq 0, \quad$ for all $u_{k} \mathrm{E} U k, k=0, \ldots, \mathbf{N}-1$
We thus obtain:

## Proposition 3.3.2: (Discrete-Time Minimllm

Suppose that $\left(u_{0}^{*}, u_{1}^{*} \ldots, u_{N-1}^{*}\right)$ is an optimal control trajectory and that $\left(x_{0}^{*}, x_{1}^{*}, \ldots, x_{N}^{*}\right)$ is the corresponding state trajectory. Assume also that the constraint sets $U k$ are convex. Then for all $k=0, \ldots, \mathrm{~N} \_1$, we have

$$
\begin{equation*}
\nabla_{u_{k}} H_{k}\left(x_{k}^{*}, u_{k}^{*}, p_{k+1}\right)^{\prime}\left(u_{k}-u_{k}^{*}\right) \geq 0, \quad \text { for all } u_{k} \in U_{k} \tag{3.43}
\end{equation*}
$$

where the vectors $P I, \ldots, P N$ are obtained from the adjoint equation

$$
p_{k}=\nabla_{x_{k}} f_{k} \cdot p_{k+1}+\nabla_{x_{k}} g_{k}, \quad k=1, \ldots, \mathbf{N} \quad 1,
$$

with the terminal condition

$$
p_{N}=\nabla g_{N}\left(x_{N}^{*}\right)
$$

The partial derivatives above are evaluated along the optimal state and control trajectories. If, in addition, the Hamiltonian $\boldsymbol{H}_{\boldsymbol{k}}$ is a convex function of $u_{k}$ for any fixed $X k$ and $p_{k+1}$, we have

$$
u_{k}^{*}=\arg \min _{u_{k} \in U_{k}} H_{k}\left(x_{k}^{*}, u_{k}, p_{k+1}\right), \quad \text { for all } k \quad 0, \ldots, \mathrm{~N} \quad 1
$$

Proof: Equation (3.43) is a restatement of the necessary condition (3.42) using the expression (3.41) for the gradient of $J$. If $11_{k}$ is convex with respect to $U k$, Eq. (3.42) is a sufficient condition for the minimum condition (3.44) to hold (see Appendix B). Q.E.D.

We now consider some variations of the continuous-time optimal control problem and derive corresponding variations of the Minimum Principle.

### 3.4.1 Fixed Terminal State

Suppose that in addition to the initial state $x(O)$, the final state $x(T)$ is given. Then the preceding informal derivations still hold except that the terminal condition $J^{*}(T, x)=h(x)$ is not true anynlOre. In effect, here we have

$$
J^{*}(T, x)= \begin{cases}0 & \text { if } x=x(T) \\ \infty & \text { otherwise }\end{cases}
$$

Thus $J^{*}(T, x)$ cannot be differentiated with respect to $x$, and the terminal boundary condition $p(T)=\nabla h\left(x^{*}(T)\right)$ for the adjoint equation does not hold. However, as compensation, we have the extra condition

$$
x(T) \text { : given, }
$$

thus maintaining the balance between boundary conditions and unknowns. If only some of the terminal states are fixed, that is,

$$
x i(T) \text { : given, } \quad \text { for all } i \mathrm{E} I,
$$

where $I$ is some index set, we have the partial boundary condition

$$
P l^{( } T^{)}=\frac{8 h\left(x^{*}(T)\right),}{\partial x_{j}}, \quad \text { for all } \mathrm{j} \notin I,
$$

for the adjoint equation.
Example ${ }^{3}$. 4.1
Consider the problem of finding the curve of minimum length connecting two points $(0, \alpha)$ and ( $T,\{3$ ). This is a fixed endpoint variation of Example 3.3.1 in the preceding section. We have

$$
\begin{gathered}
\dot{x}(t)=u(t), \\
x(O)=\alpha, \quad x(T)=\{3,
\end{gathered}
$$

and the cost is

$$
\int_{0}^{T} \sqrt{1+(U(t)) 2} d t
$$

The adjoint equation is

$$
j J(t)=0
$$

implying that

$$
p(t)=\text { constant }, \quad \text { for all } t \mathrm{E}[0, T]
$$

Minimization of the Hamiltonian,

$$
\min _{u \in \oiint}\left[\sqrt{1} \pm u^{2}+p(t) u\right]
$$

yields

$$
u^{*}(t)=\text { constant }, \quad \text { for all } t \mathrm{E}[0, T]
$$

'rhus the optimal trajectory $\left\{X^{*}(t) \mid t \in[0, \mathrm{Tn}\right.$ is a straight line. Since this trajectory must pass through $(0, \alpha)$ and $(T,\{3)$, we obtain the (a priori obvious) optimal solution shown in Fig. 3.4.1.


Figure 3.4.1 Optimal solution of the problem of connecting the two points $(0, \alpha)$ and ( $T, \beta$ ) with a minimum length curve (cf. Example 3.4.1).

## Example 3.4.2 (The Brachistochrone Problem)

In 1696 Johann Bernoulli challenged the mathematical world of his time with a problem that played an instrumental role in the development of the calculus of variations: Given two points $A$ and $B$, find a curve connecting $A$ and $B$ such that a body moving along the curve under the force of gravity reaches $B$ in minimum time (see Fig. 3.4.2). Let $A$ he $(0,0)$ and $B$ be $(T,-b)$ with $b>0$. Then it can he seen that the problem is to find $\{x(t) \mid \mathrm{t} \mathrm{E}[0, T]\}$ with $x(O)=0$ and $x(T)=b$, which minimizes

$$
\int_{0}^{T} \frac{\sqrt{1+(\dot{x}(t))^{2}}}{\sqrt{2 \gamma x(t)}} d t
$$

where $\gamma$ is the acceleration due to gravity. Here $\{(t,-x(t))$ It $\mathrm{E}[O, T J\}$, is the desired curve, the term) $1+(\dot{x}(t))_{2} d t$ is the length of the curve from $x(t)$ to $x(t+d t)$, and the term $\sqrt{2 \gamma x(t)}$ is the velocity of the body upon reaching the level $x(t)$ [if m and $v$ denote the mass and the velocity of the body, the kinetic energy is $m v^{2} / 2$, which at level $x(t)$ must be equal to the change in potential energy, which is $m \gamma x(t)$; this yields $v=\sqrt{2} \gamma r(t)]$.

We introduce the system $\dot{x}=u$, and we obtain a fixed terminal state problem $[x(O)=0$ and $x(T)=b]$. Letting

$$
g(x, u)=\frac{\sqrt{1} \pm u^{2}}{\sqrt{2 \gamma x}}
$$

the Hamiltonian is

$$
H(x, u, p)=g(x, u)+p u
$$

We minimize the Hamiltonian by setting to zero its derivative with respect to $u$ :

$$
p e t)=-\nabla_{u} g\left(x^{*}(t), u^{*}(t)\right)
$$



Figure 3.4.2 Formulation and optimal solution of the brachistochrone problem.

We know from the Minimum Principle that the Hamiltonian is constant along an optimal trajectory, Le.,

$$
g\left(x^{*}(t), u^{*}(t)\right)-\nabla_{u} g\left(x^{*}(t), u^{*}(t)\right) u^{*}(t)=\text { constant, } \quad \text { for allt } \mathrm{E}[0, T] .
$$

Using the expression for $g$, this can be written as
$\sqrt{1 \pm\left(u^{*}(t)\right)^{2}}$

$$
\frac{\left(u^{*}(t)\right)^{2}}{\sqrt{1+\left(u^{*}(t)\right)^{2}} \sqrt{2 \gamma x^{*}(t)}}=\text { constant, } \quad \text { for all } t \mathrm{E}[0, T]
$$

or equivalently

$$
\frac{1}{\sqrt{1+\left(u^{*}(t)\right) 2} \sqrt{2 \gamma x^{*}(t)}}=\text { constant, } \quad \text { for all } t \mathrm{E}[0, T]
$$

Using the relation $\dot{x}^{*}(t)=u^{*}(t)$, this yields

$$
x^{*}(t)\left(1+x^{*}(t) 2\right)=C, \quad \text { for all } t E[O, T]
$$

for some constant $C$. Thus an optimal trajectory satisfies the differential equation

$$
\dot{x}^{*}(t)=\overline{\mathrm{C}}_{x^{*}(t)}^{\frac{x^{*}(t)}{}} \quad \text { for all } t \mathrm{E}[0, T]
$$

The solution of this differential equation was known at Bernoulli's time to be a cycloid; see Fig. 3.4.2. The unknown pargmeters of the cycloid are determined by the boundary conditions $X^{*}(0)=$ and $x^{*}(T)=b$.

### 3.4.2 Free Initial State

If the initial state $x(O)$ is not fixed but is subject to optimization, we have

$$
J^{*}\left(O, x^{*}(O)\right) \leq J^{*}(O, x), \quad \text { for all :1; E } \Re^{n}
$$

yielding

$$
\nabla_{x} J^{*}\left(0, x^{*}(0)\right)
$$

and the extra boundary condition for the adjoint equation

$$
p(O)=0
$$

Also if there is a cost $\ell(x(0))$ on the initial state, i.e., the cost is

$$
\ell(x(0))+\int_{0}^{T} g(x(t), u(t)) d t+h(x(T))
$$

the boundary condition becomes

$$
p(O) \quad-\nabla \ell\left(x^{*}(0)\right)
$$

This follows by setting to zero the gradient with respect to $x$ of $\ell(x)+$ $\mathrm{J}(\mathrm{O}, x)$, Le.,

$$
\left.\nabla_{x}\{\ell(x)+J(O, x)\} l x=x * c o\right)
$$

### 3.4.3 Free Terminal Time

Suppose the initial state and/or the terminal state are given, but the terminal time $T$ is subject to optimization.

Let $\left\{\left(x^{*}(t), u^{*}(t)\right) \mid \mathrm{t} \mathrm{E}(0, T J)\right.$ be an optimal state-control trajectory pair and let $T^{*}$ be the optimal terminal time. Then if the terminal time were fixed at $T^{*}$, the pair $\left\{\left(u^{*}(t), x^{*}(t)\right)\right.$ It $\mathrm{E}\left[\mathrm{O}, \mathrm{T}^{*} \mathrm{~J}\right\}$ would satisfy the conditions of the Minimum Principle. In particular,

$$
u^{*}(t)=\arg \min _{u \in U} H\left(x^{*}(t), u, p(t)\right), \quad \text { for alIt } \mathrm{E}\left[0, T^{*}\right]
$$

where $p(t)$ is the solution of the adjoint equation. What we lose with the terminal time being free, we gain with an extra condition derived as follows.

We argue that if the terminal time were fixed at $T^{*}$ and the initial state were fixed at the given $x(O)$, but instead the initial time were subject to optimization, it would be optimal to start at $t=0$. This mea, ns that the first order variation of the optimal cost with respect to the initial time must be zero;

$$
' \operatorname{VtJ} J^{*}(t, x *(t)) l t=o=0
$$

The I-IJB equation can be written along the optimal trajectory as

$$
\nabla_{t} J^{*}\left(t, x^{*}(t)\right)=-H\left(X^{*}(t), 1 t^{*}(t), p(t)\right), \quad \text { for all } t \mathrm{E}\left[O, T^{*}\right)
$$

[cf. Eqs. (3.14) and (3.19)), so the preceding two equations yield

$$
H\left(x^{*}(O), u^{*}(O), p(O)\right)=0
$$

Since the Hamiltonian was shown earlier to be constant along the optimal trajectory, we obtain for the case of a free terminal time

$$
H\left(x^{*}(t), u^{*}(t), p(t)\right)=0, \quad \text { for all } \mathrm{t} \mathrm{E}\left[0, T^{*}\right)
$$

## Example 3.4,3 (IVIinimum-Time Problem)

A unit mass object moves horizontally under the influence of a force $u(t)$, so that

$$
\ddot{y}(t)=u(t)
$$

where yet) is the position of the object at time $t$. Given the object's initial position $y(O)$ and initial velocity $\dot{y}(0)$, it is required to bring the object to rest (zero velocity) at a given position, say zero, while using at most unit magnitude force,

$$
-1 \leq u(t) \leq 1, \quad \text { for all } t
$$

We want to accomplish this transfer in minimum time. Thus, we want to

$$
\operatorname{minimize} T=\int_{0}^{T} 1 d t
$$

Note that the integral cost, $g(x(t), u(t)) \equiv 1$, is unusual here; it does not depend on the state or the control. However, the theory does not preclude this possibility, and the problem is still meaningful because the terminal time $T$ is free and subject to optimization.

Let the state variables be

$$
\mathrm{x} 1(t)=y e t), \quad x_{2}(t)=\dot{y}(t)
$$

so the system equation is

$$
\dot{x}_{1}(t)=x_{2}(t), \quad \dot{x}_{2}(t)=u(t)
$$

The initial state $(X I(0), X 2(0))$ is given and the terminal state is also given

$$
x I(T)=0, \quad x 2(T)=0
$$

If $\left\{u^{*}(t)\right.$ It E $\left[0, T \eta\right.$ is an optimal control trajectory, $u^{*}(t)$ must minimize the Hamiltonian for each $t$, i.e.,

$$
u^{*}(t)=\arg \min _{-1 \leq u \leq 1}\left[1+p_{1}(t) x_{2}^{*}(t)+p_{2}(t) u\right]
$$

Therefore

$$
u^{*}(t)= \begin{cases}\underline{1}_{-1} & \text { if } p 2(t)<0 \\ \text { if } P 2(t) \geq 0\end{cases}
$$

The adjoint equation is

$$
\dot{p}_{1}(t)=0, \quad \dot{p}_{2}(t)=-p_{1}(t)
$$

so

$$
p_{1}(t)=c_{1}, \quad p_{2}(t)=c_{2}-c_{1} t
$$

where CI and C2 are constants. It follows that $\{p 2(t) \mid t \in[0, T$ has one of the four forms shown in Fig. 3.4.3(a); that is, $\{p 2(t) \mid t E[0, \$ switches at most once in going from negative to positive or reversely. [Note that it is not possible for $P 2(t)$ to be equal to 0 for all $t$ because this implies that $\mathrm{Pl}(t)$ is also equal to 0 for all $t$, so that the Hamiltonian is equal to 1 for all $t$; the necessary conditions require that the Hamiltonian be 0 along the optimal trajectory.] The corresponding control trajectories are shown in Fig. 3.4.3(b). The ronclusion is that, for each $t, u^{*}(t)$ is either +1 or -1 , and $\left\{u^{*}(t)\right.$ It $\mathrm{E}[0, \mathrm{l}$ has at most one switching point in the interval $[0, T]$.


P2(1)
(a)


(b)

Figure 3.4 .3 (a) Possible forms of the adjoint variable $p_{2}(t)$. (b) Corresponding forms of the optimal control trajectory

To determine the precise form of the optimal control trajectory, we use the given initial and final states. For $u(t) \equiv \zeta$, where $\zeta= \pm 1$, the system evolves according to

$$
x_{1}(t)=x_{1}(0)+x_{2}(0) t+\frac{\zeta}{2} t^{2}, \quad x_{2}(t)=x_{2}(0)+
$$

By eliminating the time $t$ in these two equations, we see that for all $t$

$$
x_{1}(t)-\frac{1}{2 \zeta}\left(X 2(t)^{2}=x 1(0)-\frac{1}{2 \zeta}(X 2(0))^{2}\right.
$$

Thus for intervals where $u(t) \equiv 1$, the system moves along the curves where $X l(t)-\frac{1}{2}\left(x_{2}(t)\right)^{2}$ is constant, shown in Fig. 3.4.4(a). For intervals where $u(t) \equiv-1$, the system moves along the curves where $\mathrm{x} 1(t)+\frac{1}{2}(X 2(t))_{2}$ is constant, shown in Fig. 3.4.4(b).

(a)

(b)

Figure 3.4.4 State trajectories when the control is $u(t) \equiv 1$ [Fig. (a)] and when the control is ' $u(t)$ - 1 [Fig. (b)].

To bring the system from the initial state $(X I(0), X 2(0))$ to the origin with at most one switch in the value of control, we must apply control according to the following rules involving the switching c'urve shown in Fig. 3.4.5.
(a) If the initial state lies above the switching curve, use $u^{*}(t) \equiv-1$ until the state hits the switching curve; then use $u^{*}(t) \equiv 1$ until reaching the origin.
(b) If the initial state lies below the switching curve, use $u^{*}(t) \equiv 1$ until the state hits the switching curve; then use $u^{*}(t) \equiv-1$ until reaching the origin.
(c) If the initial state lies on the top (bottom) part of the switching curve, use $u^{*}(t) \equiv-1\left[u^{*}(t) \equiv 1\right.$, respectively until reaching the origin.

### 3.4.4 Thne-Varying System and Cost

If the system equation and the integral cost depend on the time $t$, Le.,

$$
\dot{x}(t)=f(x(t), 11,(t), t)
$$



Figure 3.4.5 Switching curve (shown with a thick line) and closed-loop optimal control for the minimmn time example.

$$
\operatorname{cost}=h(x(T))+\int_{0}^{T} g(x(t), u(t), t) d t
$$

we can convert the problem to one involving a time-independent system and cost by introducing an extra state variable $y(t)$ representing time:

$$
\begin{gathered}
\dot{y}(t)=1, \quad y(0)=0 \\
\dot{x}(t)=f(x(t), l l,(t), y(t)), \quad x(0): \text { given, } \\
\operatorname{cost}=h(x(T))+\int^{\prime} T g(x(t), u(t), y(t)) d t .
\end{gathered}
$$

After working out the corresponding optimality conditions, we see that they are the same as when the system and cost are time-independent. The only difference is that the Hamiltonian need not be constant along the optimal trajectory.

### 3.4.5 Singular Problelms

In some cases, the minimum condition

$$
\begin{equation*}
\left.u^{*}(t)=\arg \min _{u E U} H\left(x^{*}(t), u \text { pet }\right), t\right) \tag{3.45}
\end{equation*}
$$

is insufficient to determine $u^{*}(t)$ for all $t$, because the values of $x^{*}(t)$ and $p(t)$ are such that $H\left(x^{*}(t), u, p(t), t\right)$ is independent of $u$ over a nontrivial interval of time. Such problems are called singular. Their optimal trajectories consist of portions, called Teg'ular aTCS, where $u^{*}(t)$ can be determined from the minimum condition (3.45), and other portions, called singular arcs, which can be determined from the condition that the Hamiltonian is independent of $u$.

## Example 3.4.4 (Road Construction)

Suppose that we want to construct a road over a one-dimensional terrain whose ground elevation (altitude measured from some reference point) is known and is given by $z(t), t \mathrm{E}[0, T]$. The elevation of the road is denoted by $x(t)$, $t$ E [0,1'], and the difference $x(t)-z(t)$ must be made up by fill-in or excavation. It is desired to minimize

$$
\frac{1}{2} J_{0}^{[T}(x(t) \quad z(t))^{2} d t
$$

subject to the constraint that the gradient of the road $\dot{x}(t)$ lies between $-a$ and $a$, where $a$ is a specified maximum allowed slope. Thus we have the constraint

$$
|u(t)| \leq a, \quad t \mathrm{E}\left[0,1^{\prime}\right],
$$

where

$$
\dot{x}(t)=u(t), \quad t \mathrm{E}\left[0,1^{\prime}\right]
$$

The adjoint equation here is

$$
\dot{p}(t) \quad-x *(t)+z(t)
$$

with the terminal condition

$$
p(T)=0
$$

Minimization of the Hamiltonian

$$
H(x *(t), u, p(t), t)=\frac{1}{2}\left(x^{*}(t)-Z(t)\right) 2+p(t) u
$$

with respect to $u$ yields

$$
u^{*}(t)=\arg \min _{|u| \leq a} p(t) u
$$

for all $t$, and shows that optimal trajectories are obtained by concatenation of three types of arcs:
(a) Regular arcs where $p(t)>0$ and $u^{*}(t)=-a$ (maximum downhill slope arcs).
(b) H.egular arcs where $p(t)<0$ and $u^{*}(t)=a$ (maximum uphill slope arcs).
(c) Singular arcs where $p(t)=0$ and $u^{*}(t)$ can take any value in $[-a, a]$ that maintains the condition $p(t)=0$. From the adjoint equation we see that singular arcs are those along which $p(t)=0$ and $x^{*}(t)=z(t)$, Le., the road follows the ground elevation (no fill-in or excavation). Along such arcs we must have

$$
\dot{z}(t)=u^{*}(t) \mathrm{E}[-a, \mathrm{aJ} .
$$



Figure 3.4.6 Graphical method for solving the road construction exarnple. The sharply uphill (downhill) intervals $\bar{I}$ (respectively, $\underline{I}$ ) are first identified, and are then embedded within maximum uphill (respectively, downhill) slope regular arcs $\bar{V}$ (respectively, $V$ ) within which the total fill-in is equal to the total excavation. The regular arcs are joined by singular arcs where there is no fill-in or excavation. The graphical process is started at the endpoint $t=T$.

Optimal solutions can be obtained by a graphical method using the above observations. Consider the sharply uphill intervals $\overline{\boldsymbol{I}}$ such that $\dot{z}(t) \geq a$ for all $t \mathrm{E} \bar{I}$, and the sharply downhill intervals $I$ such that $\dot{z}(t) \leq-\boldsymbol{a}$ for alIt E $\underline{I}$. Clearly, within each sharply uphill interval $\bar{I}$ the optimal slope is $u^{*}(t)=a$, but the optimal slope is also equal to $a$ within a larger maximum uphill slope interval $\bar{V} \supset \bar{I}$, which is such that $p(t)<0$ within $\bar{V}$ and

$$
p\left(t_{1}\right)=p\left(t_{2}\right)=0
$$

at the endpoints $t l$ and $t 2$ of $\bar{V}$. In view of the fornl of the adjoint equation, we see that the endpoints $t l$ and $t 2$ of $\bar{V}$ should be such that

$$
\int_{t_{1}}^{t_{2}}\left(z(t)-x^{*}(t)\right) d t \quad 0
$$

that is, the total fill-in should be equal to the total excavation within (see Fig. 3.4.6). Similarly, each sharply downhill interval $I$ should be contained within a larger maximum downhill slope interval $\underline{V} \supset I$, which is such that $p(t)>0$ within $V$, while the total fill-in should be equal to the total excavation within $V$, (see Fig. 3.4.6). Thus the regular arcs consist of the intervals $\bar{V}$ and $\underline{V}$ described above. Between the regular arcs there can be one or IllOre singular arcs where $x^{*}(t) \quad z(t)$. The optimal solution can be pieced together starting at the endpoint $t=1^{\prime}$ [where we know that $p(T)=0$ ], and proceeding backwards.

## NOTES, SOURCES, AND EXERCISES

The calculus of variations is a classical subject that originated with the works of the great mathematicians of the 17 th and 18 th centuries. Its rigorous developrnent (by modern mathematical standards) took place in the 1930s and 19/10s, with the work of a group of mathematicians that originated mostly from the University of Chicago; Bliss, McShane, and Hestenes are some of the most prominent members of this group. Curiously, this development preceded the development of nonlinear programming by many years. $\dagger$ The modern theory of deterministic optimal control has its roots primarily in the work of Pontryagin, Boltyanski, Gamkrelidze, and Mishchenko in the 1950s [PBG65]. A highly personal but controversial historical account of this work is given by Boltyanski in [BMS96]. The theoretical and applications literature on the subject is very extensive. We give three representative references: the book by Athans and Falb [AtF66] (a classical extensive text that includes engineering applications), the book by Hestenes [Hes66] (a rigorous mathematical treatment, containing important work that predates the work of Pontryagin et a1.), and the book by Luenberger [LIIe69) (which deals with optimal control within a broader infinite dimensional context). The author's nonlinear programming book [13er99] gives a detailed treatment of optimality conditions and computational methods for discrete-time optimal control.

## $\mathbb{E X E R} \mathbb{C I S E S}$

## 3.1

Solve the problem of Example 3.2.1 for the case where the cost function is

$$
(X(T))_{2}+{\underset{.10}{T}}_{\boldsymbol{F}^{T}}(u(t))_{2} d t
$$

Also, calculate the cost-to-go function $J^{*}(t, x)$ and verify that it satisfies the HJB equation.
t In the 30 s and 40 s journal space was at a premium, and finite-dimensional optimization research was thought to be a simple special case of the calculus of variations, thus insufficiently challenging or novel for publication. Indeed the modern optimality conditions of finite-dimensional optimization subject to equality and inequality constraints were first developed in the 1939 Master's thesis by Karush, but first appeared in a journal quite a few years later under the names of other researchers.

## 3.2

A young investor has earned in the stock market a large amount of money Sand plans to spend it so as to maximize his enjoyment through the rest of his life pithout working. He estimates that he will live exactly T more years and that wis capital $x(t)$ should be reduced to zero at time T , i.e., $x(T)=0$. Also he models the evolution of his capital by the differential equation

$$
\frac{d x(t)}{d t}=\alpha x(t)-u(t)
$$

where $x(O) \quad S$ is his initial capital, $\alpha>0$ is a given interest rate, and $u(t) \geq 0$ is his rate of expenditure. The total enjoyment he will obtain is given by

$$
\int_{0}^{T} e^{-\beta t} \sqrt{u(t)} d t
$$

Here $\beta$ is some positive scalar, which serves to discount future enjoyment. Find the optimal $\left\{\left.u(t)\right|_{t \in[0, T)\} \text {. }} ^{\text {. }}\right.$

## 3.3

Consider the system of reservoirs shown in Fig. 3.5.1. The system equations are

$$
\begin{aligned}
& \dot{x}_{1}(t)=-X I(t)+u(t), \\
& \dot{x}_{2}(t)=\mathrm{XI}(t),
\end{aligned}
$$

and the control constraint is $0 \leq u(t) \leq 1$ for all $t$. Initially

$$
\mathrm{Xl}(0)=\mathrm{X} 2(0)=0 .
$$

We want to maximize $X 2(1)$ subject to the constraint $x_{1}(1) \quad 0.5$. Solve the problem.


Figure 3.5.1 Reservoir system for Exercise 3.3 .

Work out the minimum-time problem (Example 3.4.3) for the case where there is friction and the object's position moves according to

$$
\ddot{y}(t)=-a \dot{y}(t)+u(t)
$$

where $a>0$ is given. Hint: The solution of the system

$$
\begin{aligned}
& \mathrm{JJI}(t)=0 \\
& \dot{p}_{2}(t)=-P I(t)+a p 2(t)
\end{aligned}
$$

is

$$
P I(t)=\operatorname{PI}(0)
$$

$$
P 2(t)=\frac{1}{a}(1-e a t) P I(O)+e^{a t} p 2(0)
$$

The trajectories of the system for $u(t) \equiv-1$ and $u(t) \equiv 1$ are sketched in Fig.


Figure 3.5.2 State trajectories of the system of Exercise 3.4 for $u(t) \equiv-1$ and
$u(t) \equiv 1$.

## 3.5 (Isoperirnetric Problem)

Analyze the problem of finding a curve $\left\{\left.x(t)\right|_{t} \mathrm{E}[0, \mathrm{~T} \|\right.$ that maximizes the
area under $x$,

$$
\int_{0}^{T} x(t) d t
$$

Sec. 3.5 Notes, Sources, and Exercises
subject to the constraints

$$
\mathrm{x}(\mathrm{O})=a, \quad x(T) \quad b, \quad \int_{0}^{T} \sqrt{1+(\dot{x}(t))^{2}} d t=L
$$

where $a, b$, and $L$ are given positive scalars. The last constraint is known as an isoperimetric constraint; it requires that the length of the curve be L. Hint: Introduce the system $\dot{x}_{1}=u, \dot{x}_{2} \quad \sqrt{1} \pm \underline{u^{2}}$, and view the problem as a fixed terminal state problem. Show that the sine of the optimalu" ( $t$ ) depends linearly on $t$. Under some assumptions on $a, b$, and $L$, the optimal curve is a circular arc.

## 3.6 (L'Môpital's Problern)

Let $a, b$, and $T$ be positive scalars, and let $A=(0, a)$ and $13 \quad(T, b)$ be two points in a medium within which the velocity of propagation of light is proportional to the vertical coordinate. Thus the time it takes for light to propagate from $A$ to $B$ along a curve $\{x(t) \mid t \mathrm{E}[0, T J\}$ is

$$
\int_{0}^{T} \sqrt{1 \pm \frac{(X(t)) 2}{C x(t,)}} d t
$$

where $C$ is a given positive constant. Find the curve of minimum travel time of light from $A$ to 13 , and show that it is an arc of a circle of the form

$$
X(t) 2+(t \quad d) 2=\mathrm{D}
$$

where $d$ and D are some constants.

## 3.7

A boat moves with constant unit velocity in a stream moving at constant velocity $s$. The problem is to find the steering angle $u(t), 0 \leq t \leq T$, which minimizes the time $T$ required for the boat to move between the point $(0,0)$ to a given point ( $a, b$ ). The equations of motion are

$$
\dot{x}_{1}(t)=s+\cos u(t), \quad \dot{x}_{2}(t)=\sin u(t)
$$

where $\mathrm{X} 1(t)$ and $X 2(t)$ are the positions of the boat paraUel and perpendicular to the stream velocity, respectively. Show that the optimal solution is to steer at a constant angle.

A unit mass object moves on a straight line fronl a given initial position $\mathrm{xl}(0)$ and velocity $\mathrm{X} 2(0)$. Find the force $\{u(t) \mid t \mathrm{E}[\mathrm{O}, \mathrm{IJ}\}$ that brings the object at time 1 to rest $[X 2(1)=0]$ at position $\mathrm{x} 1(1)=0$, and minimizes

$$
\int_{0}^{1}(u(t))^{2} d t
$$

Use the Minimum Principle to solve the linear-quadratic problem of Example 3.2.2. Hint: Follow the lines of Example 3.3.3.

### 3.10 (On the Need for Convexity Assumptions)

Solve the continuous-time problem involving the system $\dot{x}(t)=u(t)$, the terminal cost $(x(T))_{2}$, and the control constraint $u(t)=-$ lor 1 for all $t$, and show that the solution satisfies the Minimum Principle. Show that, depending on the initial state Xo , this may not be true for the discrete-time version involving the system $x_{k+1} \quad x_{k}+u_{k}$, the terminal cost $x_{N}^{2}$, and the control constraint $\boldsymbol{u}_{\boldsymbol{k}}=-1$ or 1 for all $k$.

### 3.11

Use the discrete-time Minimum Principle to solve Exercise 1.14 of Chapter 1, assuming that each $w_{k}$ is fixed at a known deterministic value.

### 3.12

Use the discrete-time Minimum Principle to solve Exercise 1.15 of Chapter 1, assuming that $\gamma_{k}$ and $\delta_{k}$ are fixed at known deterministic values.

### 3.13 (Lagrange Multipliers and the Minimum Principle)

Consider the discrete-time optimal control problem of Section 3.3.3, where there are no control constraints $\left(U=\Re^{m}\right)$. Introduce a Lagrange multiplier vector $P k+l$ for each of the constraints

$$
f_{k}\left(x_{k}, u_{k}\right)-x_{k+1}=0, \quad k=0, \ldots, N-1
$$

and form the Lagrangian function

$$
g N(X N)+\sum_{\mathrm{k}=\mathrm{O}}^{N-l}\left(9 k\left(X k, u_{k}\right)+p_{k+1}^{\prime}\left(J k\left(X k^{\prime} u_{k}\right)-X k+l\right)\right)
$$

(cL Appendix B). View both the state and the control vectors as the optimization variables of the problem, and show that by differentiation of the Lagrangian function with respect to $X k$ and $u_{k}$, we obtain the discrete-time Minimum Principle.

## Problems with Perfect State

## Information

## Contents

| 4.1. Linear Systems and Quadratic Cost | p. 148 |
| :--- | :--- |
| 4.2. Inventory Control . . . . | p. 162 |
| 4.3. Dynamic Portfolio Analysis | p. 170 |
| 4.4. Optimal Stopping Problems | p. 176 |
| 4.5. Scheduling and the Interchange Argument | p. 186 |
| 4.6. Set-IVlembership Description of Uncertainty | p. 190 |
| 4.6.1. Set-Membership Estimation | p. 191 |
| 4.6.2. Oontrol with Unknown-but-Bounded Disturbances | p. 197 |
| 4.7. Notes, Sources, and Exercises. . . . . . . . . . . . | p. 201 |

In this chapter we consider a number of applications of discrete-time stochastic optimal control with perfect state information. These applications are special cases of the basic problem of Section 1.2 and can be addressed via the DP algorithm. In all these applications the stochastic nature of the disturbances is significant. For this reason, in contrast with the deterministic problems of the preceding two chapters, the use of closed-loop control is essential to achieve optimal performance.

### 4.1 LINEAR SYSTEMS AND QUADRATIC COST

In this section we consider the special case of a linear system

$$
x_{k+1}=A_{k} x_{k}+B_{k} u_{k}+w_{k}, \quad k=0,1, \ldots, N-1
$$

and the quadratic cost

$$
\underset{\substack{w_{k} \\ k=0,1, \ldots, N \cdots 1}}{E_{1}}\left\{x_{N}^{\prime} Q_{N} x_{N}+\sum_{k=0}^{N-1}\left(x_{k}^{\prime} Q_{k} x_{k}+u_{k}^{\prime} R_{k} u_{k}\right)\right\}
$$

In these expressions, $x_{k}$ and $U k$ are vectors of dimension nand m , respectively, and the matrices $A_{k}, 13_{k}, Q k, R k$ are given and have appropriate dimension. We assume that the matrices $Q k$ are positive semidefinite symmetric, and the matrices $R_{k}$ are positive definite symmetric. The controls $U k$ are unconstrained. The disturbances $\omega k$ are independent random vectors with given probability distributions that do not depend on $X k$ and $U k$. Furthermore, each $w k$ has zero mean and finite second moment.

The problem described above is a popular formulation of a regulation problem whereby we want to keep the state of the system close to the origin. Such problems are common in the theory of automatic control of a motion or a process. The quadratic cost function is often reasonable because it induces a high penalty for large deviations of the state from the origin but a relatively small penalty for small deviations. Also, the quadratic cost is frequently used, even when it is not entirely justified, because it leads to a nice analytical solution. A number of variations and generalizations have similar solutions. For example, the disturbances $w k$ could have nonzero means and the quadratic cost could have the form
$E\left\{\left(x_{N}-\bar{x}_{N}\right)^{\prime} Q_{N}\left(x_{N}-\bar{x}_{N}\right)+\sum_{k=0}^{N-1}\left(\left(x_{k}-\bar{x}_{k}\right)^{\prime} Q_{k}\left(x_{k}-\bar{x}_{k}\right)+u_{k}^{\prime} R_{k} u_{k}\right)\right\}$,
which expresses a desire to keep the state of the system close to a given trajectory $\left(\bar{x}_{0}, \bar{x}_{1}, \ldots, \bar{x} N\right)$ rather than close to the origin. Another generalized version of the problem arises when $A K^{\prime} 13 k$ are independent random
matrices, rather than being known. This case is considered at the end of this section.

Applying now the DP algorithm, we have

$$
\begin{gather*}
J_{N}\left(x_{N}\right)=x_{N}^{\prime} Q_{N} x_{N} \\
J_{k}\left(x_{k}\right)=\min _{u_{k}} E\left\{x_{k}^{\prime} Q_{k} x_{k}+u_{k}^{\prime} R_{k} u_{k}+J_{k+1}\left(A_{k} x_{k}+B_{k} u_{k}+w_{k}\right)\right\} \tag{4.1}
\end{gather*}
$$

It turns out that the cost-to-go functions $J k$ are quadratic and as a result the optimal control law is a linear function of the state. These facts can be verified by straightforward induction. We write Eq. (4.1) for $k=N-1$,

$$
\begin{aligned}
J_{N-1}\left(x_{N-1}\right)=\min _{\cup N-1} E\{ & x_{N-1}^{\prime} Q_{N-1} x_{N-1}+u_{N-1}^{\prime} R_{N-1} u_{N-1} \\
& +\left(A_{N-1} x_{N-1}+B_{N-1} u_{N-1}+w_{N-1}\right)^{\prime} Q_{N} \\
& \left.\cdot\left(A N-1 X N-1+B_{N-1} u_{N-1}+{ }^{\prime} W N-1\right)\right\}
\end{aligned}
$$

and we expand the last quadratic form in the right-hand side. We then use the fact $E\left\{w_{N-1}\right\}=0$ to eliminate the term $E\left\{w_{N-1}^{\prime} Q_{N}\left(A_{N-1} x_{N-1}+\right.\right.$ $\left.\left.B_{N-1} u_{N-1}\right)\right\}$, and we obtain

$$
\begin{aligned}
J_{N-1}\left(x_{N-1}\right)= & x_{N-1}^{\prime} Q_{N-1} x_{N-1}+\min _{U N-1}\left[u_{N-1}^{\prime} R_{N-1} u_{N-1}\right. \\
& \left.+u_{N-1}^{\prime} B_{N-1}^{\prime} Q_{N} B_{N-1} u_{N-1}+2 x_{N-1}^{\prime} A_{N-1}^{\prime} Q N 13 N-1 \cup N-1\right] \\
& +x_{N-1}^{\prime} A_{N-1}^{\prime} Q_{N} A_{N-1} x_{N-1}+E\left\{w_{N-1}^{\prime} Q_{N} w_{N-1}\right\}
\end{aligned}
$$

By differentiating with respect to $U N-1$ and by setting the derivative equal to zero, we obtain

$$
\left(R_{N-1}+B_{N-1}^{\prime} Q_{N} B_{N-1}\right) u_{N-1}=-B_{N-1}^{\prime} Q_{N} A_{N-1} x_{N-1}
$$

The matrix multiplying $U N_{-\mathbf{1}}$ on the left is positive definite (and hence invertible), since $R N-/$ is positive definite and $B_{N-1}^{\prime} Q_{N} B_{N-1}$ is positive semidefinite. As a result, the minimizing control vector is given by

$$
u_{N-1}^{*}=-\left(R_{N-1}+B_{N-1}^{\prime} Q_{N} B_{N-1}\right)^{-1} B_{N-1}^{\prime} Q_{N} A_{N-1} x_{N-1}
$$

By substitution into the expression for $I N-1$, we have

$$
J_{N-1}\left(x_{N-1}\right)=x_{N-1}^{\prime} K_{N-1} x_{N-1}+E\left\{w_{N-1}^{\prime} Q_{N} w_{N-1}\right\}
$$

where by straightforward calculation, the matrix $K_{N_{-1}}$ is verified to be

$$
\begin{aligned}
K_{N-1}= & A_{N-1}^{\prime}\left(Q N-Q_{N} B_{N-1}\left(B_{N-1}^{\prime} Q N B N-I+R_{N-1}\right)^{-1} B_{N-1}^{\prime} Q N\right) A_{N-1} \\
& +Q N-I .
\end{aligned}
$$

The matrix $K N-I$ is clearly symmetric. It is also positive semidefinite. To see this, note that from the preceding calculation we have for $x \mathrm{E} \Re^{n}$

$$
\begin{aligned}
& x^{\prime} K_{N-1} x=\min _{u}\left[\mathrm{x}^{\prime} \mathrm{QN}-\mathrm{IX}+u^{\prime} R_{N-1} u\right. \\
&\left.\quad+(A N-I X+B N-I U)^{\prime} Q N\left(A N-I X+B_{N-1} u\right)\right]
\end{aligned}
$$

Since $Q N-I, R N^{-}{ }_{1}$, and $Q N$ are positive semidefinite, the expression within brackets is nonnegative. Minimization over $\boldsymbol{u}$ preserves nonnegativity, so it follows that $x^{\prime} K \mathrm{~N}-1 x \geq 0$ for all $x \mathrm{E} \Re^{n}$. Hence](N-1 is positive semidefinite.

Since $/ \mathbb{N}-$ I is a positive semidefinite quadratic function (plus an inconsequential constant term), we may proceed similarly and obtain from the DP equation (4.1) the optimal control law for stage $N-2$. As earlier, we show that $\mathrm{JN}-2$ is a positive semidefinite quadratic function, and by proceeding sequentially, we obtain the optimal control law for every $k$. It has the form

$$
\begin{equation*}
\mu_{k}^{*}\left(x_{k}\right)=L_{k} x_{k} \tag{4.2}
\end{equation*}
$$

where the gain matrices $L k$ are given by the equation

$$
L_{k}=-\left(B_{k}^{\prime} K_{k+1} B_{k}+R_{k}\right)^{-1} B_{k}^{\prime} K_{k+1} A_{k}
$$

and where the symmetric positive semidefinite matrices $K k$ are given recursively by the algorithm

$$
\begin{equation*}
K_{N}=Q_{N} \tag{4.3}
\end{equation*}
$$

$$
\begin{equation*}
K_{k}=A_{k}^{\prime}\left(K_{k+1}-K_{k+1} B_{k}\left(B_{k}^{\prime} K_{k+1} B_{k}+R_{k}\right)^{-1} B_{k}^{\prime} K_{k+1}^{\prime}\right) A_{k}+Q_{k} \tag{4.4}
\end{equation*}
$$

Just like DP, this algorithm starts at the terminal time $N$ and proceeds backwards. The optimal cost is given by

$$
J o(x o)=x_{0}^{\prime} K_{0} x_{0}+\sum_{k=O}^{N-I} E\left\{w_{k}^{\prime} K_{k+1} W k\right\}
$$

The control law (4.2) is simple and attractive for implementation in engineering applications: the current state $x k$ is being fed back as input through the linear feedback gain matrix $L k$ as shown in Fig. 4.1.1. This accounts in part for the popularity of the linear-quadratic formulation. As we will see in Chapter 5, the linearity of the control law is still maintained even for problems where the state $x k$ is not completely observable (imperfect state information).


Figure 4.1.1 Linear feedback structure of the optimal controller for the linearquadratic problem.

## The Riccati Equation and Its Asymptotic Behavior

Equation (4.4) is called the d,tscrete-t'tme Riccati equation. It plays an important role in control theory. Its properties have been studied extensively and exhaustively. One interesting property of the Riccati equation is that if the matrices $A_{k}, B_{k}, Q k, R k$ are constant and equal to $A, B, Q$, $R$, respectively, then the solution $K k$ converges as $k \rightarrow-$ oo (under mild assumptions) to a steady-state solution $K$ satisfying the algebra'lc Riccati equation

$$
\begin{equation*}
K=A^{\prime}\left(K-K B\left(B^{\prime} K B+R\right)-I B^{\prime} K\right) A+ \tag{4.5}
\end{equation*}
$$

This property, to be proved shortly, indicates that for the system

$$
x_{k+1}=A x_{k}+B u_{k}+w_{k}, \quad k=0,1, \ldots, N \quad 1,
$$

and a large number of stages $N$, one can reasonably a, pproximate the control law (4.2) by the control law $\left\{\mu^{*}, \mu^{*}, \ldots, \mu^{*}\right\}$, where

$$
\begin{equation*}
\mu^{*}(x)=L x \tag{4.6}
\end{equation*}
$$

$$
L=-\left(B^{\prime} K B+R\right)-I B^{\prime} K A
$$

and ] ( solves the algebraic Riccati equation (4.5). This control law is stationary; that is, it does not change over time.

We now turn to proving convergence of the sequence of matrices $\{K k\}$ generated by the Riccati equation (4.4). We first introduce the notions of controllability and observability, which are very important in control theory.

Definition 4.1.1: A pair $(A, B)$, where $A$ is an $n \times n$ matrix and $B$ is an $n \times m$ matrix, is said to be controllable if the $n \times n m$ matrix

$$
[B, A B, A 2 B, \ldots, A n-I B]
$$

has full rank (Le., has linearly independent rows). A pair ( $A, O$ ), where $A$ is an $n \times n$ matrix and 0 an $\mathrm{m} \times n$ matrix, is said to be observable if the pair $\left(A I, O^{\prime}\right)$ is controllable, where $A^{\prime}$ and $G^{\prime}$ denote the transposes of $A$ and $G$, respectively.

One may show that if the pair $(A, B)$ is controllable, then for any initial state $x_{0}$, there exists a sequence of control vectors Un, UI,..., Un-I that force the state $x_{n}$ of the system

$$
x_{k+1}=A x_{k}+B u_{k}
$$

to be equal to zero at time $n$. Indeed, by successively applying the above equation for $k=n-1, n \quad 2, \ldots, 0$, we obtain

$$
x_{n}=A n x o+B U n-1+A B u_{n}-2+\ldots+A n-i B u o
$$

or equivalently

$$
x_{n}-A^{n} x_{0}=\left(B, A B, \ldots, A^{n-1} B\right)\left(\begin{array}{c}
\left.U_{n} n_{-2} I\right)  \tag{4.7}\\
\vdots \\
u_{0}
\end{array}\right)
$$

If $(A, B)$ is controllable, the matrix $(B, A B, \ldots, A n-I B)$ has full rank and as a result the right-hand side of Eq. (4.7) can be made equal to any vector in $\Re^{n}$ by appropriate selection of ( $u 0, \mathrm{U}, \ldots, u_{n-1}$ ). In particular, one can choose (uo, $\mathrm{U}, \ldots, u_{n-1}$ ) so that the right-hand side of Eq. (4.7) is equal to - Anxo, which implies $x_{n}=0$. This property explains the name "controllable pair" and in fact is often used to define controllability.

The notion of observability has an analogous interpretation in the context of estimation problems; that is, given measurements $z_{0}, \mathbf{Z I}, \cdots, Z n-1$ of the form $Z k=G x k$, it is possible to infer the initial state $x_{o}$ of the system $x_{k+1}=A X k^{\prime}$ in view of the relation

$$
\left(\begin{array}{c}
z_{n-1} \\
\vdots \\
z_{1} \\
z_{0}
\end{array}\right)=\left(\begin{array}{c}
C A^{n-1} \\
\vdots \\
C A \\
C
\end{array}\right) x_{0} .
$$

Alternatively, it can be seen that observability is equivalent to the property that, in the absence of control, if $O x k \rightarrow 0$ then $x k \rightarrow 0$.

The notion of stability is of paramount importance in control theory. In the context of our problem it is important.tha, the stationary control law (4.6) results in a stable closed-loop system; that is, in the absence of input disturbance, the state of the system

$$
x_{k+1}=(A+B L) x_{k}, \quad k=0,1, \ldots,
$$

tends to zero as $k \rightarrow 00$. Since $x k=(A+B L) k x o$, it follows that the closed-loop system is stable if and only if $(A+B L) k \rightarrow 0$, or equivalently (see Appendix A), if and only if the eigenvalues of the matrix $(A+B L)$ are strictly within the unit circle.

The following proposition shows that for a stationary controllable system and constant matrices $Q$ and $R$, the solution of the Riccati equation (4.4) converges to a positive definite symmetric matrix $K$ for an arbitrary positive semidefinite symmetric initial matrix. In addition, the proposition shows that the corresponding closed-loop system is stable. The proposition also requires an observability assumption, namely, that $Q$ can be written as $C^{\prime} G$, where the pair $(A, G)$ is observable. Note that if $r$ is the rank of $Q$, there exists an $r$ x $n$ matrix $G$ of rank $r$ such that $Q=G^{\prime} G$ (see Appendix A). The implication of the observability assumption is that in the absence of control, if the state cost per stage $x_{k}^{\prime} Q x_{k}$ tends to zero or equivalently $C X k \rightarrow 0$, then also $X k \rightarrow 0$.

To simplify notation, we reverse the time indexing of the Riccati equation. Thus, $P k$ in the following proposition corresponds to $K N-k$ in Eq. (4.4). A graphical proof of the proposition for the case of a scalar system is given in Fig. 4.1.2.

Proposition 4.4.1: Let $A$ be an $n \times n$ niatrix, $B$ be an $n \times{ }^{n} n$ matrix, $Q$ be an $n \times n$ positive semidefinite symmetric matrix, and $R$ be an $\mathrm{m} \times \mathrm{m}$ positive definite symmetric matrix. Consider the discrete-time Riccati equation

$$
\begin{equation*}
P_{k+1}=A^{\prime}\left(P_{k}-P_{k} B\left(B^{\prime} P_{k} B+R\right)^{-1} B^{\prime} P_{k}\right) A+Q, \quad k=0,1, \ldots \tag{4.8}
\end{equation*}
$$

where the initial matrix $P o$ is an arbitrary positive semidefinite symmetric matrix. Assume that the pair $(A, B)$ is controllable. Assume also that $Q$ may be written as $C^{\prime} G$, where the pair $(A, G)$ is observable. Then:
(a) There exists a positive definite symmetric matrix $P$ such that for every positive semidefinite symmetric initial matrix $P_{0}$ we have

$$
\lim _{k \rightarrow \infty} P_{k}=P
$$

Furthermore, P is the unique solution of the algebraic matrix equation

$$
\begin{equation*}
P=A^{\prime}\left(P-P B\left(B^{\prime} P B+R\right)-l B I P\right) A+Q \tag{4.9}
\end{equation*}
$$

within the class of positive semidefinite symmetric matrices.
(b) The corresponding closed-loop system is stable; that is, the eigenvalues of the matrix

$$
\begin{equation*}
D \quad A+B L, \tag{4.10}
\end{equation*}
$$

where

$$
\begin{equation*}
L=-\left(B^{\prime} P B+R\right)-l B^{\prime} P A, \tag{4.11}
\end{equation*}
$$

are strictly within the unit circle.

Proof: The proof proceeds in several steps. First we show convergence of the sequence generated by Eq. (4.8) when the initial matrix $P o$ is equal to zero. Next we show that the corresponding matrix $D$ of Eq. (4.10) satisfies $D k \rightarrow 0$. Then we show the convergence of the sequence generated by Eq. (4.8) when Po is any positive semidefinite symmetric matrix, and finally we show uniqueness of the solution of Eq. (4.9).

Initial MatTix Po $=0$. Consider the optimal control problem of finding $u_{0}, u_{1}, \ldots, u_{k-1}$ that minimize

$$
\sum_{i=0}^{k-1}\left(x_{i}^{\prime} Q x_{i}+u_{i}^{\prime} R u_{i}\right)
$$

subject to

$$
x_{i+1}=A x_{i}+B u_{i}, \quad \mathrm{i}=0,1, \ldots, k \quad l,
$$

where $x_{0}$ is given. The optimal value of this problem, according to the theory of this section, is $x_{0}^{\prime} P_{k}(0) x_{0}$,
where $\operatorname{Pk}(O)$ is given by the Riccati equation (4.8) with $P o=0$. For any control sequence ( $u_{0}, u_{1}, \ldots, u_{k}$ ) we have

$$
\sum_{i=0}^{k-1}\left(x_{i}^{\prime} Q x_{i}+u_{i}^{\prime} R u_{i}\right) \leq \sum_{i=0}^{k}\left(x_{i}^{\prime} Q x_{i}+u_{i}^{\prime} R u_{i}\right)
$$

and hence

$$
\begin{aligned}
x_{0}^{\prime} P_{k}(0) x_{0} & =\min _{\substack{u_{i} \\
\mathrm{i}=\mathrm{O}, \ldots, k-1}} \sum_{i=0}^{k-1}\left(x_{i}^{\prime} Q x_{i}+u_{i}^{\prime} R u_{i}\right) \\
& \leq \min _{\substack{u_{i} \\
\mathrm{i}=\mathrm{O}, \ldots, \mathrm{k}}} \sum_{i=0}^{k}\left(x_{i}^{\prime} Q x_{i}+u_{i}^{\prime} R u_{i}\right) \\
& =x_{0}^{\prime} P_{k+1}(0) x_{0}
\end{aligned}
$$



Figure 4.1.2 Graphical proof of Prop. 4.4.1 for the case of a scalar stationary system (one-dimensional state and control), assuming that $A \neq 0, B \neq 0, Q>0$, and $R>0$. The Riccati equation (4.8) is given by

$$
P_{k+i}=A^{2}\left(\begin{array}{ll}
P_{k} & \frac{B^{2} P_{k}^{2}}{B^{2} P_{k}+} R_{-}
\end{array}\right)+Q
$$

which can be equivalently written as

$$
P_{k+1}=F\left(P_{k}\right),
$$

where the function $F$ is given by

$$
P(P)=\stackrel{A^{2} R P}{--+-R+Q .}
$$

Because $F$ is concave and monotonically increasing in the interval ( $-R / B 2, o 0$ ), as shown in the figure, the equation $P=F(P)$ has one positive solutionP* and one negative solution $\tilde{P}$. The Riccati iteration $P_{k}+_{1} \quad F\left(P_{k}\right)$ converges to $P^{*}$ starting anywhere in the interval ( $P, \circ \circ$ ) as shown in the figure.
where both minimizations are subject to the system equation constraint $x_{i+1}=A x_{i}+B u_{i}$. Furthermore, for a fixed $x_{0}$ and for every $k, x_{0}^{\prime} P_{k}(0) x_{0}$ is bounded from above by the cost corresponding to a control sequence fhat forces $X o$ to the origin in $n$ steps and applies zero control after that. Such a sequence exists by the controllability assumption. Thus the sequence $\left\{x_{0}^{\prime} P_{k}(O) x o\right\}$ is nonclecreasing with respect to $k$ and bounded from above, and therefore converges to some real number for every $x_{0} \mathrm{E} \Re^{n}$. It follows that the sequence $\{P h(O)\}$ converges to some matrix $P$ in the sense that each of the sequences of the elements of $\operatorname{Pk}(O)$ converges to the corresponcl-
ing elements of $P$. To see this, take $X Q=(1,0, \ldots, 0)$. Then $x_{0}^{\prime} P_{k}(0) x_{0}$ is equal to the first diagonal element of $P k(O)$, so it follows that the sequence offlrst diagonal elements of $\operatorname{Pk}(O)$ converges; the limit of this sequence is the first diagonal element of $P$. Similarly, by taking $X Q=(0, \ldots, 0,1,0, \ldots, 0)$ with the 1 in the ith coordinate, for $\mathrm{i}=2, \ldots, n$, it follows that all the diagonal elements of $\operatorname{Pk}(0)$ converge to the corresponding diagonal elements of $P$. Next take $X Q=(1,1,0, \ldots, 0)$ to show that the second elements of the first row converge. Continuing similarly, we obtain

$$
\lim _{k \rightarrow \infty} P k(O)=P
$$

where $\operatorname{Pk}(O)$ are generated by Eq. (4.8) with $P Q=0$. Furthermore, since $\operatorname{Pk}(O)$ is positive semidefinite and symmetric, so is the limit matrix $P$. Now by taking the limit in Eq. (4.8) it follows that $P$ satisfies

$$
P=A^{\prime}\left(P-P B\left(B^{\prime} P B+R\right)-l B^{\prime} P\right) A+Q .
$$

In addition, by direct calculation we can verify the following useful equality

$$
\begin{equation*}
P \quad D^{\prime} P D+Q+L^{\prime} R L \tag{4.12}
\end{equation*}
$$

where $D$ and $L$ are given by Eqs. (4.10) and (4.11). An alternative way to derive this equality is to observe that from the DP algorithm corresponding to a finite horizon N we have for all states $x N-k$

$$
\begin{aligned}
x_{N-k}^{\prime} P_{k+1}(0) x_{N-k}=x_{N-k}^{\prime} Q x_{N-k} & +\mu_{N-k}^{*}\left(x_{N-k}\right)^{\prime} R \mu_{N-k}^{*}\left(x_{N-k}\right) \\
& +x_{N-k+1}^{\prime} \operatorname{Pk}(O) X N-k+l .
\end{aligned}
$$

By using the optimal controller expression $\mu_{N-k}^{*}\left(x_{N-k}\right)=L N-k X N-k$ and the closed-loop system equation $X N-k+l=(A+B L N-k) X N-k$, we thus obtain

$$
\begin{equation*}
P_{k+1}(0)=Q+L_{N-k}^{\prime} R L_{N-k}+\left(A+B L_{N-k}\right)^{\prime} P_{k}(0)\left(A+B L_{N-k}\right) \tag{4.13}
\end{equation*}
$$

Equation (4.12) then follows by taking the limit as $k \rightarrow$ oo in Eq. (4.13).
Stability of the Closed-Loop System. Consider the system

$$
\begin{equation*}
x_{k+1}=(A+B L) x_{k}=D x_{k} \tag{4.14}
\end{equation*}
$$

for an arbitrary initial state $X Q$. We will show that $X k \rightarrow$ as $k \rightarrow 0$. We have for all $k$, by using Eq. (4.12),

$$
x_{k+1}^{\prime} P x_{k+1}-x_{k}^{\prime} P x_{k}=x_{k}^{\prime}\left(D^{\prime} P D-P\right) X k=-x_{k}^{\prime}\left(Q+L^{\prime} R L\right) x_{k}
$$

Hence

$$
\begin{equation*}
x_{k+1}^{\prime} P x_{k+1}=x_{0}^{\prime} P x_{0}-{ }_{i=Q}^{k} x_{i}^{\prime}\left(Q+L^{\prime} R L\right) X i^{\prime} \tag{4.15}
\end{equation*}
$$

rrhe left-hand side of this equation is bounded below by zero, so it follows that

$$
\lim _{k \rightarrow \infty} x_{k}^{\prime}\left(Q+L^{\prime} R L\right) X k=0
$$

Since $R$ is positive definite and $Q$ may be written as $C^{\prime} C$, we obtain

$$
\begin{equation*}
\lim _{k \rightarrow \infty} O X k=0, \quad \lim _{k \rightarrow \infty} L X k=\lim _{k \rightarrow \infty} \mu^{*}\left(x_{k}\right)=0 \tag{4.16}
\end{equation*}
$$

The preceding relations imply that as the control asymptotically becomes negligible, we have $\lim _{k \rightarrow \infty} O X k=0$, and in view of the observability assumption, this implies that $X k \rightarrow 0$. To express this argument more precisely, let us use the relation $X k+l=(A+B L) x_{k}$ [d. Eq. (4.14)], to write

$$
\left|\begin{array}{cc}
C\left(x_{k+n-1}-\sum_{i=1}^{n-1} A^{i-1} B L X k+n-i-l\right)  \tag{4.17}\\
C(X k+n-2 & \left.\sum_{i=1}^{n-2} A \mathbf{A i - 1} B L X k+n-i-2\right) \\
C(X k+l-B L x k) \\
O X k
\end{array}\right|\left(\begin{array}{c}
C A^{n-1} \\
C A^{n-2} \\
\vdots \\
C A \\
C
\end{array}\right) x_{k}
$$

Since $L X k \rightarrow$ by Eq. (4.16), the left-hand side tends to zero and hence the right-hand side tends to zero also. By the observability assumption, however, the matrix multiplying $X k$ on the right side of (4.17) has full rank. It follows that $x k \rightarrow 0$.

Positio Definiteness of $P$. Assume the contrary, i.e., there exists some $X Q \neq$ such that $x_{0}^{\prime} P x_{0}=0$. Since $P$ is positive semidefinite, from Eq. (4.15) we obtain

$$
x_{k}^{\prime}\left(Q+L^{\prime} R L\right) x_{k}=0, \quad \mathbf{k}=\mathbf{O}, 1, \ldots
$$

Since $X k \rightarrow 0$, we obtain $x_{k}^{\prime} Q x_{k}=x_{k}^{\prime} C^{\prime} C x_{k} \quad$ and $x_{k}^{\prime} L^{\prime} R L x_{k}=0$, or

$$
C x_{k}=0, \quad L x_{k}=0, \quad k=0,1, \ldots
$$

Thus all the con(1)ols $\mu^{*}(X k)=L X k$ of the closed-loop system are zero while we have $C X k=$ for all $k$. Based on the observability assumption, we will show that this implies $X Q=0$, thereby reaching a contradiction. Indeed, consider Eq. (4.17) for $k=0$. By the preceding equalities, the left-ha,ncl side is zero and hence

$$
O=\left(\begin{array}{cc}
C A^{n-1} & \\
\vdots & \\
O \mathbf{A} & \\
\boldsymbol{C} &
\end{array}\right.
$$

Since the matrix multiplying $X Q$ above has full rank by the observability assumption, we obtain $X Q=0$, which contradicts the hypothesis $X Q \neq 0$ and proves that $P$ is positive definite.

Arbitrary Initial Matrix Po. Next we show that the sequence of matrices $\left\{P_{k}\left(P_{0}\right)\right\}$, defined by Eq. (4.8) when the starting matrix is an arbitrary positive sermdefinite symmetric matrix $P o$, converges to $P$ $\lim _{k \rightarrow \infty} P^{k}(0)$. Indeed, the optimal cost of the problem of minimizing

$$
\begin{equation*}
x_{k}^{\prime} P_{0} x_{k}+\sum_{i=0}^{k-I}\left(x_{i}^{\prime} Q x_{i}+u_{i}^{\prime} R u_{i}\right) \tag{4.18}
\end{equation*}
$$

subject to the system equation $X i+I=A X^{\prime} i+B U i$ is equal to $x_{0}^{\prime} P_{k}\left(P_{0}\right) x_{0}$. Hence we ha, ve for every $X o \mathrm{E} \Re^{n}$

$$
x_{0}^{\prime} P_{k}(0) x_{0} \leq x_{0}^{\prime} P_{k}\left(P_{0}\right) x_{0}
$$

Consider now the cost (4.18) corresponding to the controller $\mu\left(x_{k}\right)=U k=$ $L x_{k}$, where $L$ is defined by Eq. (4.11). This cost is

$$
x_{0}^{\prime}\left(D^{k^{\prime}} P_{0} D^{k}+\sum_{i=0}^{k-1} D^{i^{\prime}}\left(Q+L^{\prime} R L\right) D^{i}\right) x_{0}
$$

and is greater or equal to $x_{0}^{\prime} P_{k}\left(P_{0}\right) x_{0}$, which is the optimal value of the cost (4.18). Hence we have for all $k$ and $x \mathrm{E}_{R^{n}}$

$$
x^{\prime} P k(O) x \leq x^{\prime} P_{k}\left(P_{0}\right) x \leq x^{\prime}\left(D^{k^{\prime}} P_{0} D^{k}+\sum_{i=0}^{k-1} D i^{\prime}\left(Q+L^{\prime} R L\right) D i\right) x
$$

We have proved that

$$
\lim _{k \rightarrow \infty} P_{k}(0)=P
$$

and we also have, using the fact $\lim _{k \rightarrow \infty} D k^{f} P o D k=0$, and the relation $Q+L^{\prime} R L \quad P-D^{\prime} P D$ fcf. Eq. (4.12)),

$$
\begin{align*}
\lim _{k \rightarrow \infty}\left\{D^{k^{\prime}} P_{0} D_{k}\right. & \left.+\sum_{i=0}^{k-1} D^{i^{\prime}}\left(Q+L^{\prime} R L\right) D i\right\} \\
& =\lim _{k \rightarrow \infty}\left\{\sum_{i=0}^{k-1} D i^{\prime}\left(Q+L^{\prime} R L\right) D i\right\}  \tag{4.19}\\
& =\lim _{k \rightarrow \infty}\left\{\sum_{i=0}^{k-1} D^{i^{\prime}}\left(\mathrm{P}-\mathrm{D}^{\prime} \mathrm{P} D\right) D i\right\}
\end{align*}
$$

p.

Combining the preceding three equations, we obtain

$$
\lim _{k \rightarrow \infty} P k\left(P_{O}\right)=P
$$

for an arbitrary positive semidefinite symmetric initial matrix Po.
Uniqueness of Solution. If $\tilde{P}$ is another positive semidefinite symmetric solution of the algebraic Riccati equation (4.9), we have $P_{k}(\tilde{P})=\tilde{P}$ for all $k=0,1, \ldots$ From the convergence result.just proved, we then obtain

$$
\lim _{k \rightarrow \infty} P_{k}(\tilde{P})=P
$$

implying that $\tilde{P} \quad P$ Q.E.D.
The assumptions of the preceding proposition can be relaxed somewhat. Suppose that, instead of controllability of the pair $\quad B$ ), we assume that the system is stabilizable in the sense that there exists an $m \times n$ feedback gain matrix $G$ such that the closed-loop system $x_{k+1} \quad(A+B G) x_{k}$ is stable. Then the proof of convergence of $\operatorname{Pk}(O)$ to some positive semidefinite $P$ given previously carries through. [We use the stationary control law $\mu(x)=G x$ for which the closed-loop system is stable to ensure that $x_{0}^{\prime} P_{k}(0) x_{0}$ is bounded.] Suppose that, instead of observability of the pair $(A, O)$, the system is assumed detectable in the sense that $A$ is such that if $u_{k} \rightarrow 0$ and $O X k \rightarrow 0$ then it follows that $X k \rightarrow 0$. (This essentially means that instability of the system can be detected by looking at the measurement sequence $\left\{z_{k}\right\}$ with $Z k=C x_{k}$.) Then Eq. (4.16) implies that $X k \rightarrow 0$ and that the system $x_{k+1}=(A+B L) X k$ is stable. The other parts of the proof of the proposition follow similarly, with the exception of positive definiteness of $P$, which cannot be guaranteed anymore. (As an example, take $A=0, B \quad 0,0=0, R>0$. Then both the stabilizability and the detectability assumptions are satisfied, but $P=0$.)

To summarize, if the controllability and observability assumptions of the proposition are replaced by the preceding stabilizability and detectability assumptions, the conclusions of the proposition hold with the exception of positive definiteness of the limit matrix $P$, which can now only be guaranteed to be positive semidefinite.

## Random System Matrices

We consider now the case where $\left\{A_{O}, B_{O}\right\}, \ldots,\left\{A N-I, B_{N-1}\right\}$ are not known but rather are independent random matrices that are also independent of 'WO, 'WI, .... 'WN-1. Their probability distributions are given, and they are assumed to have finite second moments. This problem falls again within the framework of the basic problem by considering as disturbance at each time $k$ the triplet $\left(A_{k}, B_{k}, w_{k}\right)$. The DP algorithm is written as

$$
J_{N}\left(x_{N}\right)=x_{N}^{\prime} Q_{N} x_{N}
$$

$J_{k}\left(x_{k}\right)=\min _{u_{k} w_{k}, A_{k}, B_{k}}\left\{x_{k}^{\prime} Q_{k} x_{k}+u_{k}^{\prime} \Gamma_{k} u_{k}+J_{k+1}\left(A_{k} x_{k}+B_{k} u_{k}+w_{k}\right)\right\}$.

Calculations very similar to those for the case where $A k, B k$ are not random show that the optimal control law has the form

$$
\mu_{k}^{*}\left(x_{k}\right)=L_{k} x_{k}
$$

where the gain matrices $L k$ are given by

$$
L_{k}=-\left(R_{k}+E\left\{B_{k}^{\prime} K_{k+1} B_{k}\right\}\right)^{-1} E\left\{B_{k}^{\prime} K_{k+1} A_{k}\right\}
$$

and where the matrices $K k$ are given by the recursive equation

$$
K_{N}=Q_{N}
$$

$$
\begin{align*}
K k= & E\left\{A_{k}^{\prime} K_{k+1} A_{k}\right\} \\
& -E\left\{A_{k}^{\prime} K_{k+1} B_{k}\right\}\left(R_{k}+E\left\{B_{k}^{\prime} K_{k+1} B_{k}\right\}\right)^{-1} E\left\{B_{k}^{\prime} K_{k+1} A_{k}\right\}+\underset{(4.20)}{Q k} \tag{4.20}
\end{align*}
$$

In the case of a stationary system and constant matrices $Q k$ and $R k$ it is not necessarily true that the above equation converges to a steady-state solution. This is demonstrated in Fig. 4.1.3 for a scalar system, where it is shown that if the expression

$$
T=E\left\{A^{2}\right\} E\left\{B^{2}\right\}-(E\{A\})^{2}(E\{B\})^{2}
$$

exceeds a certain threshold, the matrices $K k$ diverge to oo starting from any nonnegative initial condition. A possible interpretation is that if there is a lot of uncertainty about the system, as quantified by $T$, optimization over a long horizon is meaningless. This phenomenon has been called the uncertainty threshold principle; see Athans, Ku , and Gershwin [AGK77], and Ku and Athans [KuA77].

## On Certainty Equivalence

We close this section by making an observation about the simplifications that arise when the cost is quadratic. Consider the minimization over $u$ of

$$
\underset{w}{E}\{(a x+b u+W) 2\}
$$

where $a$ and $b$ are given scalars, $x$ is known, and $w$ is a random variable. The optimum is attained for

$$
u^{*}=-\left(\frac{a}{b}\right) x-\left(\frac{1}{b}\right) E\{w\}
$$

Thus $u^{*}$ depends on the probability distribution of $w$ only through the mean $E\{w\}$. In particular, the result of the optimization is the same as


Figure 4.1.3 Graphical illustration of the asymptotic behavior of the generalized Riccati equation (4.20) in the case of a scalar stationary system (one-dimensional state and control). Using $P_{k}$ in place of $K N-K$, this equation is written as

$$
P_{k+1}=\bar{F}\left(P_{k}\right)
$$

where the function $\tilde{F}$ is given by

$$
\begin{aligned}
& T=E\left\{A^{2}\right\} E\left\{B^{2}\right\}-(E\{A\}) 2(E\{B\}) 2 .
\end{aligned}
$$

If $T=0$, as in the case where $A$ and $B$ are not random, the Riccati equation becomes identical with the one of Fig. 4.1.2 and converges to a steady-state. Convergence also occurs when $T$ has a small positive value. However, as illustrated in the figure, for $T$ large enough, the graph of the function $\tilde{F}$ and the 45 -degree line that passes through the origin do not intersect at a positive value of $P$, and the Riccati equation diverges to infinity.
for the corresponding deterministic problem where $w$ is replaced by $E\{w\}$. This property is called the certainty equivalence principle and appears in various forms in many (but not all) stochastic control problems involving linear systems and quadratic cost. For the first problem of this section, where $A_{k}, B_{k}$ are known, certainty equivalence holds because the optimal control law (4.2) is the same as the one that would be obtained from the corresponding deterministic problem where $w k$ is not random but rather is known and is equal to zero (its expected value). However, for the problem where $A_{k}, B_{k}$ are random, the certainty equivalence principle does not hold, since if one replaces $A_{k}, B k$ with their expected values in Eq. (4.20) the resulting control law need not be optimal.

### 4.2 INVENTORY CONTROL

We consider now the inventory control problem discussed in Sections 1.1 and 1.2. We assume that excess demand at each period is backlogged and is filled when additional inventory becomes available. This is represented by negative inventory in the system equation

$$
x_{k+1}=x_{k}+u_{k}-w_{k}, \quad k=0,1, \ldots, \mathrm{~N}-1 .
$$

We assume that the demands $W k$ take values within some bounded interval and are independent. We will analyze the problem for the case of a holding/shortage cost of the form

$$
r(x)=p \max (O,-x)+h \max (O, x),
$$

where $p$ and $h$ are given nonnegative scalars. Thus the total expected cost to be minimized is

$$
E\left\{\sum_{k=0}^{N-1}\left(c u_{k}+r\left(x_{k}+u_{k}-w_{k}\right)\right)\right\} .
$$

We assume that the purchase cost per unit stock c is positive and that $p>c$. The last assumption is necessary for the problem to be well posed; if c , the purchase cost per unit, were greater than $p$, the depletion cost per unit, it would never be optimal to buy new stock at the last period and possibly in earlier periods. Much of the subsequent analysis generalizes to the case where $r$ is a convex function that grows to infinity with asymptotic slopes $p$ and $h$ as its argument tends to -oo and oo, respectively.

By applying the DP algorithm, we have

$$
J_{N}\left(x_{N}\right)=0,
$$

$$
\begin{equation*}
J_{k}\left(x_{k}\right)=\min _{u_{k} \geq 0}\left[c u_{k}+H\left(x_{k}+u_{k}\right)+E\left\{J_{k+1}\left(x_{k}+u_{k}-w_{k}\right)\right\}\right], \tag{4.21}
\end{equation*}
$$

where the function $H$ is defined by

$$
H(y)=E\{r(y-W k)\}=p E\{\max (\mathrm{O}, W k-y)\}+h E\{\max (\mathrm{O}, y-W k)\}
$$

Actually, $H$ depends on $k$ whenever the probability distribution of $W k$ depends on $k$. To simplify notation, we do not show this dependence and assume that all demands are identically distributed, but the following analysis carries through even when the demand distribution is time-varying. The function $H$ can be seen to be convex, since $r(y-W k)$ is convex in $y$ for each fixed $W k$, and taking expectation over $W k$ preserves convexity.

By introducing the variable $Y k=X k+U k$, we can write the DP Eq. (4.21) as

$$
\begin{equation*}
J_{k}\left(x_{k}\right)=\min _{y_{k} \geq x_{k}} G_{k}\left(y_{k}\right)-c x_{k} \tag{4.22}
\end{equation*}
$$

where

$$
\begin{equation*}
G_{k}(y)=c y+H(y)+E\left\{J_{k+1}(y-w)\right\} . \tag{4.23}
\end{equation*}
$$

We will prove shortly that the function $G k$ is convex, but for the moment let us assume this convexity. Suppose that $G_{k}$ has an unconstrained minimum with respect to $y$, denoted by $S k$ :

$$
S_{k}=\arg \min _{y \in \Re} G_{k}(y)
$$

Then, in view of the constraint $Y k \geq X k$, it is seen that a minimizing $Y k$ in Eq. (4.22) equals $S k$ if $X k<S k$, and equals $X k$ otherwise [since by convexity, $G k(Y)$ cannot decrease as $Y$ increases beyond $S k]$. Using the reverse transformation $U k=Y k-X k$, we see that the minimum in the DB equation (4.21) is attained at $U k=S k-X k$ if $X k<S k$, and at $U k=$ otherwise. An optimal policy is determined by the sequence of scalars $\left\{S o, S l, \ldots, S_{N-1}\right\}$ and has the form

$$
\mu_{k}^{*}\left(x_{k}\right)=\left\{\begin{array}{lll}
S_{k} & X k & \text { if } X k<S k  \tag{4.24}\\
0 & & \text { if } X k \geq S k
\end{array}\right.
$$

Thus, the optimality of the policy (4.24) will be proved if we can show that the cost-to-go functions $J k$ [and hence also the functions $G k$ of Eq. (4.23)] are convex, and furthermore $\lim _{|y| \rightarrow \infty} G k(Y)=00$, so that the minimizing scalars $S k$ exist. We proceed to show these properties inductively.

We have that $\mathrm{J}_{\mathrm{N}}$ is the zero function, so it is convex. Since $\mathrm{c}<p$ and the derivative of $H(y)$ tends to $-p$ as $y \rightarrow-$ oo, we see that $G N-1(y)$ [which is $c y+H(y)$ ] has a derivative that becomes negative as $y \rightarrow-$ oo and becomes positive as $y \rightarrow 00$ (see Fig. 4.2.1). Therefore

$$
\lim _{|y| \rightarrow \infty} G N-1(y)=00
$$

Thus, as shown above, an optimal policy at time $N-1$ is given by

$$
\mu_{N-1}^{*}(X N-1)=\{9 N-1-X N-1 \quad \text { if } X N-1<S N-1
$$

Furthermore, from the DP equation (4.21) we have



Figure 4.2.1 Structure of the cost-to-go functions when the fixed cost is zero.
which is a convex function because $H$ is convex and $8 \mathrm{~N}-1$ minimizes $c y+$ $H(y)$ (see Fig. 4.2.1). Thus, given the convexity of $I N$, we were able to prove the convexity of IN-I. Furthermore,

$$
\lim _{|y| \rightarrow \infty} I N-l(y)=00
$$

This argument can be repeated to show that for all $k \quad \mathrm{~N}-2, \ldots, 0$, if $\mathrm{Jk}_{\mathrm{k}+1}$ is convex, $\lim _{|y| \rightarrow \infty} J k+l(y)=00$, and $\lim _{|y| \rightarrow \infty} G k(Y)=00$, then we have

$$
J_{k}\left(x_{k}\right)= \begin{cases}c\left(S_{k}-x_{k}\right)+H\left(S_{k}\right)+E\left\{J_{k+1}\left(S_{k}-w_{k}\right)\right\} & \text { if } X k<8 k \\ H\left(x_{k}\right)+E\left\{J_{k+1}\left(x_{k}-w_{k}\right)\right\} & \text { if } X k \geq 8 k\end{cases}
$$

where $8 k$ is an unconstrained minimum of $G k$. Furthermore, $J k$ is convex, $\lim _{|y| \rightarrow \infty} J k(Y)=00$, and $\lim _{|y| \rightarrow \infty} G k-l(Y)=00$. Thus, the optimality proof of the policy (4.24) is completed.

## Positive Fixed Cost and ( $s, 8$ ) Policies

We now turn to the more complicated case where there is a positive fixed cost $K$ associated with a positive inventory order. Thus the cost for order-
ing inventory $U \geq$ is

$$
\mathrm{C}(u)= \begin{cases}K+c u & \text { if } u>0 \\ O & \text { if } u=0\end{cases}
$$

The DP algorithm takes the form

$$
\begin{gathered}
I N(X N)=0, \\
J k(X k)=\min _{u_{k} \geq 0}[C(U k)+H(X k+U k)+E\{\mathrm{Jk}+1(X k+U k-W k)\}]^{\prime}
\end{gathered}
$$

with $H$ defined as earlier by

$$
H(y)=E\{r(y-w)\}=p E\{\max (\mathrm{O}, w-y)\}+h E\{\max (\mathrm{O}, y-w)\} .
$$

Consider again the functions

$$
G_{k}(y)=c y+H(y)+E\left\{J_{k+1}(y-w)\right\} .
$$

Then $J k$ is written as

$$
J k(X k)=\min \left[G k(X k), \min _{u_{k}>0}[K+G k(X k+U k)]\right]-c x_{k}
$$

or equivalently, through the change of variable $Y k=X k+u_{k}$,

$$
J_{k}\left(x_{k}\right)=\min \left[G_{k}\left(x_{k}\right), \min _{y_{k}>x_{k}}\left[K+G_{k}\left(y_{k}\right)\right]\right]-c x_{k}
$$

If, as in the case where $K=0$, we could prove that the functions $G k$ are convex, then it would not be-difficult to verify [see part (d) of the following Lemma 4.2.1] that the policy

$$
\mu_{k}^{*}(X k)=\left\{\delta \delta-X k \quad \begin{array}{ll}
\text { if } x k<s k  \tag{4.25}\\
& \text { if } x_{k} \geq s_{k}
\end{array}\right.
$$

is optimal, where $8 k$ is a value of $Y$ that minimizes $G k(Y)$ and $S k$ is the smallest value of $Y$ for which $G k(y)=K+G k(8 k)^{\prime}$ A policy of the form (4.25) is known as a multiperiod ( $s, 8$ ) policy.

Unfortunately, when $K>0$ it is not necessarily true that the functions $G_{k}$ are convex. This opens the possibility of $G_{k}$ having the form shown in Fig. 4.2.2, where the optimal policy is to order ( $8-x$ ) in interval I, zero in intervals II and IV, and $(\tilde{8}-x)$ in interval IlL However, we will show that even though the functions $G_{k}$ may not be convex, they have the property
$K+G_{k}(z+y) \geq G_{k}(y)+z\left(\frac{G_{k}(y)-G_{k}(y-b)}{b}\right), \quad$ for all $z \geq 0, b>0, y$.

This property is called $K$-convexity and was first used by Scarf [Sca60] to show the optimality of multiperiod $(8, B)$ policies. Now if K -convexity holds, the situation shown in Fig. 4.2.2 is impossible; if yo is the local maximum in the interval III, then we must have, for sufficiently small $b>0$,

$$
\left.\underline{G k(y o)}=\frac{G k(Y O}{b}=\underline{b}\right)>0,
$$

and from Eq. (4.26) it follows that

$$
K+G_{k}(\tilde{S}) \geq G_{k}\left(y_{0}\right)
$$

which contradicts the construction shown in Fig. 4.2.2. More generally, it will be shown by using part (d) of the following Lemma 4.2.1 that if the K-convexity Eq. (4.26) holds, then an optimal policy takes the ( $8, B$ ) form (4.25).


Figure 4.2.2 Potential form of the function $G k$ when the fixed cost is nonzero. If $\mathrm{G} k$ had the form shown in the figure, the optimal policy would be to order $(8-x)$ in interval I, zero in intervals II and IV, and $(\tilde{S}-x)$ in interval III. The use of $K$-convexity allows us to show that the form of $\mathrm{G} k$ shown in the figure is impossible.

Definition 4.2.1: We say that a real-valued function 9 is $K$-convex, where $K \geq 0$, if

$$
K+g(z+y) \geq g(y)+z\left(g(y) \quad \frac{9(y-b))}{} \quad \text { for all } z \geq 0, b>0, Y .\right.
$$

Some properties of $K$-convex functions are provided in the following lemma. Part (d) of the lemma essentially proves the optimality of the $(8,8)$ policy (4.25) when the functions $G k$ are K-convex.

## Lemma 4.2.1:

(a) A real-valued convex function 9 is also O-convex and hence also K-convex for all $K \geq 0$.
(b) If $g l(Y)$ and $g 2(y)$ are K-convex and L-convex ( $K \geq 0, L \geq 0)$, respectively, then $\alpha g_{1}(y)+p g 2(y)$ is $(\alpha K+\mathrm{pL})$-convex for all $\alpha>0$ and $p>0$.
(c) If $g(y)$ is K -convex and $w$ is a random variable, then $E w\{g(y$ $w)\}$ is also $K$-convex, provided $E_{w}\{I g(y-w) I\}<00$ for all $y$.
(d) If 9 is a continuous K-convex function and $g(y) \rightarrow 00$ as $\mid y!\rightarrow 00$, then there exist scalars 8 and 8 with $8 \leq 8$ such that
(i) $g(8) \leq g(y)$, for all scalars $y$;
(ii) $g(8)+K=g(8)<g(y)$, for all $y<8$;
(iii) $g(y)$ is a decreasing function on ( $-\mathbf{O O}, s$ );
(iv) $g(y) \leq g(z)+K$ for all $y, z$ with $8 \leq Y \leq z$.

Proof: Part (a) follows from elementary properties of convex functions, and parts (b) and (c) follow from the definition of a $K$-convex function. We will thus concentrate on proving.part (d).

Since 9 is continuous and $g(y) \rightarrow 00$ as lyl $\rightarrow$ oo, there exists a minimizing point of $g$. Let 8 be such a point. Also let $s$ be the smallest scalar $z$ for which $z \leq 8$ and $g(8)+K=g(z)$. For all $y$ with $y<s$, we have from the definition of $K$-convexity

$$
\left.K+g(8) \geq g(s)+\frac{\boldsymbol{B}}{-8}-\boldsymbol{s}(s)-g(y)\right) .
$$

Since $K+g(8)-g(8)=0$, we obtain $g(8) \quad g(y) \leq 0$. Since $y<8$ and 8 is the smallest scalar for which $g(8)+K=g(8)$, we must have $g(8)<g(y)$ and (ii) is proved. To prove (iii), note that for YI $<Y 2<8$, we have

$$
K+g(8) \geq g(Y 2)+\underset{Y 2-y_{1}}{\boldsymbol{B}-}\left(g(Y 2)-g\left(y_{1}\right)\right)
$$

Also from (ii),

$$
g(Y 2)>g(8)+
$$

and by adding these two inequalities we have

$$
\mathrm{O}>\frac{8}{Y 2}-\frac{Y 2}{\mathrm{YI}}\left(g(Y 2)-g\left(y_{1}\right)\right)
$$

from which we obtain $g\left(y_{1}\right)>g(Y 2)$, thus proving (iii). To prove (iv), we note that it holds for $y=z$ as well as for either $y=8$ or $y=s$. There remain two other possibilities: $8<Y<z$ and $s<y<8$. If $8<Y<z$, then by $K$-convexity

$$
K+g(z) \geq g(y)+\frac{z-\mathcal{Y}}{y-}(g(y)-g(8)) \geq g(y)
$$

and (iv) is proved. If $s<Y<8$, then by K-convexity

$$
g(s)=K+g(8) \geq g(y)+{ }_{y}^{S-}(g(y)-g(s))
$$

from which

$$
\left(1+\frac{S-y}{y-s}\right) g(s) \geq\left(1+\frac{S-\underline{y}}{\bar{y}-s}\right) g(y)
$$

and $g(s) \geq g(y)$. Noting that

$$
g(z)+K \geq g(8)+K=g(s)
$$

it follows that $g(z)+K \geq g(y)$. Thus (iv) is proved for this case as well. Q.E.D.

Consider now the function $\mathrm{Gn}_{\mathrm{N}}$ - :

$$
G N-1(y)=e y+H(y)
$$

Clearly, GN-1 is convex and hence by part (a) of Lemma 4.2.1, it is also K-convex. We have

$$
I N-1(X)=\min \left[G N-1(X), \min _{y>x}[\mathrm{~K}+G N-1(y) J]-e x\right.
$$

and it can be seen that

$$
J_{N-1}(x)= \begin{cases}K+G_{N-1}\left(S_{N-1}\right)-c x & \text { for } x<S N-1  \tag{4.27}\\ G_{N-1}(x)-c x & \text { for } x \geq S N-1\end{cases}
$$

where $8 N-1$ minimizes $G N-1(y)$ and $S N-1$ is the smallest value of $y$ for which $G N-1(y)=K+G N-1(8 N-1)$. Note that since $K>0$, we have $S N-1 \neq 8 N-1$ and furthermore the derivative of $G N-1$ at $S N-1$ is negative. As a result the left derivative of $\mathrm{JN}_{\mathrm{N}}-1$ at $S \mathrm{~N}-1$ is greater than the right derivative, as shown in Fig. 4.2.3, and $\mathrm{J}_{\mathrm{N}-1}$ is not convex. However, we will show that $I N-1$ is K -convex based on the fact that $G N-1$ is K-convex. To this end we must verify that for all $z \geq 0, b>0$, and $y$, we have

$$
\begin{equation*}
K+J_{N-1}(y+z) \geq J_{N-1}(y)+z\left(\frac{J_{N-1}(y)-J_{N-1}(y-b)}{b}\right) \tag{4.28}
\end{equation*}
$$



Figure 4.2.3 Structure of the cost-to-go function when fixed cost is nonzero.

We distinguish three cases:
Case 1: $y \geq S N-1$. If $Y-b \geq S N-1$, then in this region of values of $z, b$, and $y$, the function $I_{N-1}$, by Eq. (4.27), is the sum of a K-convex function and a linear function. Hence by part (b) of Lemma 4.2.1, it is K-convex and Eq. (4.28) holds. If $y-b<S N-1$, then in view of Eq. (4.27) we can write Eq. (4.28) as

$$
\begin{aligned}
K+G N-1(y+z) & -e(y+z) \geq G N-1(y)-e y \\
& +z \underline{(G-1} \underline{N-1} \underline{(y)}=\underline{c y}=\underline{G N b-l} \underline{(S N-l)} \pm \underline{c(y}=\underline{b}))
\end{aligned}
$$

or equivalently

$$
\begin{equation*}
\left.K+G N-1(y+z) \geq G N-1(y)+\ddot{z(G N-l(Y)}-\frac{G_{N-1}\left(s_{N-1}\right)}{b}\right) \tag{4.29}
\end{equation*}
$$

Now if $y$ is such that $G N-1(y) \geq G_{N-1}\left(s_{N-1}\right)$, then by K-convexity of $G N-1$ we have

$$
\begin{aligned}
K+G_{N^{-}} l(y+z) & \left.\geq G N-1(y)+z \frac{(G N-1(y)}{y}-\frac{G_{N-1}\left(s_{N-1}\right)}{S N-h_{1}}\right) \\
& >G_{N-1}(1)+z\left(G N-1(y)-G_{b} N-l(S N-d)\right.
\end{aligned}
$$

Thus Eq. (4.29) and hence also Eq. (4.28) hold. If $y$ is such that $G \mathrm{~N}-1(y)<$ $G_{N-1}\left(s_{N-1}\right)$, then we have

$$
\begin{aligned}
K+G N-1(y+z) & \geq K+G_{N-1}\left(S_{N-1}\right) \\
& =G N-1(S N-1) \\
& >G_{N-1}(y) \\
& \geq G_{N-1}\left(\frac{y}{2}+Z^{\left(G N-1(y)-\frac{G_{N-1}\left(s_{N-1}\right)}{b}\right)} .\right.
\end{aligned}
$$

So for this case, Eq. (4.29) holds, and hence also the desired K-convexity Eq. (4.28) holds

Case 2: $y \leq y+z \leq S N-l$. In this region, by Eq. (4.27), the function $I N-I$ is linear and hence the K-convexity Eq. (4.28) holds.

Case 3: $y<S N-I<y+z$. For this case, in view of Eq. (4.28), we can write the K-convexityEq. (4.28) as

$$
\begin{aligned}
K+G_{N-1}(y+z) & -c(y+z) \geq G_{N-1}\left(s_{N-1}\right)-c y \\
& +z\left(\frac{G_{N-1}\left(s_{N-1}\right)-c y-G_{N-1}\left(s_{N-1}\right)+c(y-b)}{b}\right)
\end{aligned}
$$

or equivalently

$$
K+G N-I(Y+z) \geq G N-I(S N-I)
$$

which holds by the definition of $S N-I$.
We have thus proved that K -convexity and continuity of $G N-I$, together with the fact that $G N-1(y) \rightarrow \infty$ as $|y| \rightarrow 00$, imply K-convexity of $I N-I$. In addition, $I N-l$ can be seen to be continuous. Now using Lemma 4.2.1, it follows from Eq. (4.23) that $G N-2$ is a $K$-convex function. Furthermore, by using the boundedness of $W N-2$, it follows that $G N-2$ is continuous and, in addition, $G N-2(Y) \rightarrow 00$ as lyl $\rightarrow 00$. Repeating the preceding argument, we obtain that $I N-2$ is K-convex, and proceeding similarly, we prove K-convexity and continuity of the functions $G k$ for all $k$, as well as that $G k(y) \rightarrow o o$ as $|y| \rightarrow o o$. At the same time [by using part (d) of Lemma 4.2.1] we prove optimality of the multiperiod ( $s, S$ ) policy of Eq. (4.25).

The optimality of policies of the $(S, S)$ type can be proved for several other inventory problems (see Exercises 4.3 to 4.10).

### 4.3 DYNAMIC PORTFOLIO ANALYSIS

Portfolio theory deals with the question of how to invest a certain amount of wealth among a collection of assets, perhaps over a long time interval. One approach, to be discussed in this section, is to assume that an investor makes a decision in each of several successive time periods with the objective of maximizing final wealth. We will start with an analysis of a single-period model and then extend the results to the multiperiod case.

Let $X o$ denote the initial wealth of the investor and assume that there are $n$ risky assets, with corresponding random rates of return el, ... en among which the investor can allocate his wealth. The investor can also invest in a riskless asset offering a sure rate of return $s$. If we denote by UI, $\ldots, U n$ the corresponding amounts invested in the $n$ risky assets and by
(xo UI - ... Un) the amount invested in the riskless asset, the wealth at the end of the first period is given by

$$
\mathrm{x} 1=s(x o-\mathrm{UI} \quad \ldots-U n)+\sum_{\mathrm{i}=1}^{n} e i U i
$$

or equivalently

$$
\mathrm{x} 1=S X O+\sum_{\mathrm{i}=1}^{n}\left(e_{i}-S\right) U i^{\prime}
$$

The objective is to maximize over UI, $\ldots, U n$,

$$
E\left\{U\left(x_{1}\right)\right\}
$$

where $U$ is a known function, referred to as the utility junction for the investor (Appendix G contains a discussion of utility functions and their significance in the formulation of optimization problems under uncertainty). We assume that $U$ is concave and twice continuously differentiable, and that the given expected value is well defined and finite for all $X o$ and $U i$. We will not impose constraints on UI, ... Un. This is necessary in order to obtain the results in convenient form. A few additional assumptions will be made later.

Let us denote by $u_{i}^{*}=\mu^{i^{*}}(x o), \mathrm{i}=1, \ldots, n$, the optimal amounts to be invested in the $n$ risky assets when the initial wealth is $X o$. We will show that when the utility function satisfies

$$
\begin{gather*}
U^{\prime}(x) \\
\underline{U I I}(X)=a+b x, \quad \text { for all } x, \tag{4.30}
\end{gather*}
$$

where $U^{\prime}$ and $U^{\prime \prime}$ denote the first and second derivatives of $U$, respectively, and $a$ and $b$ are some scalars, then the optimal portfolio is given by the linear policy

$$
\begin{equation*}
\mu^{i *}\left(x_{0}\right)=\alpha^{i}\left(a+b s x_{0}\right), \quad \mathrm{i}=1, \ldots, n \tag{4.31}
\end{equation*}
$$

where $\alpha^{i}$ are some constant scalars. Furthermore, if $J(x o)$ is the optimal value of the problem

$$
J(x o)=\max _{\forall i} E\left\{U\left(x_{1}\right)\right\}
$$

then we have

$$
\begin{equation*}
\underline{J^{\prime}(x o)} \underset{I I(X O)}{ }=\frac{a}{s}+b x o, \quad \text { for all } X o . \tag{4.32}
\end{equation*}
$$

It can be verified that the following utility functions $U(x)$ satisfy condition (4.30):
exponential: $e-x j a$, for $b=0, a>0$,
logarithmic: $\operatorname{In}(x+a)$, for $b=1$,
power: $(l /(b-l))(a+b x)^{1-(1 / b)}$, for $b \neq 0, b \neq 1$.

Only concave utility functions from this class are admissible for our problem. Furthermore, if a utility function that is not defined over the whole real line is used, the problem should be formulated in a way that ensures that all possible values of the resulting final wealth are within the domain of definition of the utility function.

To show the desired relations, let us hypothesize that an optimal portfolio exists and is of the form

$$
\mu^{i *}\left(x_{0}\right)=\alpha^{i}\left(x_{0}\right)\left(a+b s x_{0}\right)
$$

where $\alpha^{i}\left(x_{0}\right), \mathrm{i}=1, \ldots, n$, are some differentiable functions. We will prove that dai(xo)jdxo $=0$ for all Xo, implying that the functions $\boldsymbol{\alpha}^{i}$ must be constant, so the optimal portfolio has the linear form (4.31).

We have for every $X_{o}$ and $\mathrm{i}=1, \ldots, n$, by the optimality of $\mu^{i^{*}}(x o)$,

$$
\begin{align*}
\frac{d E\left\{U\left(x_{1}\right)\right\}}{d u_{i}} & =E\left\{U f\left(s x_{O}+{ }_{\mathbf{j}=1}^{\prime} t(e j-s) a j(X O)\left(a+b s X_{O}\right)\right)(e i-s)\right\} \\
& =0 \tag{4.33}
\end{align*}
$$

Differentiating each of the $n$ equations above with respect to Xo yields

$$
\begin{gather*}
E\left\{\left(\begin{array}{c}
\left(e_{1}-s\right)^{2} \cdots\left(e_{1}-s\right)\left(e_{n}-s\right) \\
\vdots \\
\left(e_{n}-s\right)\left(e_{1}-s\right) \cdots\left(e_{n}-s\right)^{2}
\end{array}\right) U^{\prime \prime}\left(x_{1}\right)\left(a+b s x_{0}\right)\right\}\left(\begin{array}{c}
\frac{d \alpha^{1}\left(x_{0}\right)}{d x_{0}} \\
\vdots \\
\frac{d \alpha^{n}\left(x_{0}\right)}{d x_{0}}
\end{array}\right) \\
=-\left(\begin{array}{c}
E\left\{U^{\prime \prime}\left(x_{1}\right)\left(e_{1}-s\right) s\left(1+\sum_{i=1}^{n}\left(e_{i}-s\right) \alpha^{i}\left(x_{0}\right) b\right)\right\} \\
\vdots \\
E\left\{U^{\prime \prime}\left(x_{1}\right)\left(e_{n}-s\right) s\left(1+\sum_{i=1}^{n}\left(e_{i}-s\right) \alpha^{i}\left(x_{0}\right) b\right)\right\}
\end{array}\right) \tag{4.34}
\end{gather*}
$$

Using relation (4.31), we have

$$
\begin{align*}
U^{\prime \prime}\left(x_{1}\right)=- & \left.a+b\left(s x o+\sum_{i=1}^{n} \frac{U^{\prime}\left(x_{1}\right)}{(e i} s\right) a i(x o)(a+b s x o)\right) \\
& \frac{U^{\prime}\left(x_{1}\right)}{\left(a+b s x_{0}\right)\left(1+\sum_{i=1}^{n}\left(e_{i}-s\right) \alpha^{i}\left(x_{0}\right) b\right)} \tag{4.35}
\end{align*}
$$

Substituting in Eq. (4.34) and using Eq. (4.33), we have that the right-hand side of Eq. (4.34) is the zero vector. The matrix on the left in Eq. (4.34), except for degenerate cases, can be shown to be nonsingular. Assuming that it is indeed nonsingular, we obtain

$$
\frac{d \alpha^{i}\left(x_{0}\right)}{d x_{0}}=0, \quad \mathrm{i}=1, \ldots, n
$$

and $\alpha^{i}(x o)=\alpha^{i}$, where $\alpha^{i}$ are some constants, thus proving the optimality of the linear policy (4.31).

We now prove Eq. (4.32). We have

$$
\begin{aligned}
J(x o) & =E\left\{U\left(x_{1}\right)\right\} \\
& =E\left\{U\left(s\left(1+\sum_{i=1}^{n}(e i-s) \alpha^{i} b\right) x_{o}+\sum_{i=1}^{n}(e i-s) \alpha^{i} a\right)\right\}
\end{aligned}
$$

and hence

$$
\begin{gather*}
J^{\prime}(x o)=E\left\{U^{\prime}(x,) s\left(1 \nmid \sum_{i=1}^{n}\left(e_{i}-s\right) \alpha^{i} b\right)\right\} \\
J^{\prime \prime}(x o)=E\left\{U^{\prime \prime}(x,) s^{\prime}\left(1+\sum_{i=1}^{n}\left(e_{i}-s\right) \alpha^{i} b\right) 2\right\} \tag{4.36}
\end{gather*}
$$

The last relation after some calculation using Eq. (4.35) yields

$$
\begin{equation*}
J^{\prime \prime}(x o)=-\frac{E\left\{U f(X l) S\left(1+\sum_{i=1}^{n}\left(e_{i}-s\right) \alpha^{i} b\right)\right\} s}{a+b s x o} \tag{4.37}
\end{equation*}
$$

By combining Eqs. (4.36) and (4.37), we obtain the desired result:

$$
\begin{aligned}
& J f(X O) \\
& -\underline{J^{\prime \prime}(x o)}=\frac{a}{s}+b x o .
\end{aligned}
$$

## The Multiperiod Problem

We now extend the preceding one-period analysis to the multiperiod case. We will assume that the current wealth can be reinvested at the beginning of each of $N$ consecutive time periods. We denote
$X k$ : the wealth of the investor at the start of the kth period,
$u_{i}^{k}$ : the amount invested at the start of the kth period in the ith risky asset,
$e_{i}^{k}$ : the rate of return of the ith risky asset during the kth period,
$S k$ : the rate of return of the riskless asset during the kth period.
We have the system equation

$$
X k+l=S k X k+\sum_{i=1}^{n}\left(e_{i}^{k}-s_{k}\right) u_{i}^{k}, \quad k=O, l, \ldots, N-1 .
$$

We assume that the vectors $e^{k}=\left(e_{1}^{k}, \ldots, e_{n}^{k}\right), k=0, \ldots, \mathrm{~N}-1$, are independent with given probability distributions that yield finite expected values throughout the following analysis. The objective is to maximize $E\{U(X N)\}$, the expected utility of the terminal wealth $X N$, where we assume that $U$ satisfies

$$
\begin{gathered}
U /(X) \\
--/ X X)
\end{gathered}=a+b x, \quad \text { for all } x
$$

Applying the DP algorithm to this problem, we have
$J k(X k)$

$$
\begin{gather*}
J_{N}\left(x_{N}\right)=U\left(x_{N}\right) \\
\max _{u_{1}^{k}, \ldots, u_{n}^{k}} E\left\{J_{k+1}\left(S k X k+\sum_{i=1}^{n}\left(e_{i}^{k}-s_{k}\right) u_{i}^{k}\right)\right\} \tag{4.38}
\end{gather*}
$$

From the solution of the one-period problem, we have that the optimal policy at period $N$ - 1 has the form

$$
\mu_{N-1}^{*}\left(x_{N-1}\right)=\alpha_{N-1}\left(a+b s_{N-1} x_{N-1}\right)
$$

where $\alpha_{N-1}$ is an appropriate n -dimensional vector. Furthermore, we have

$$
\frac{J_{N-1}^{\prime}(x)}{J_{N-1}^{\prime \prime} X}=\underset{S N-l}{a_{-}}+b_{X}
$$

Hence, applying the one-period result in the DP Eq. (4.38) for the next to the last period, we obtain the optimal policy

$$
\mu_{N-2}^{*}\left(x_{N-2}\right)=\alpha_{N-2}\left(\frac{a}{s_{N-1}}+b s_{N-2} x_{N-2}\right)
$$

where $\alpha_{N^{-2}}$ is again an appropriate n -dimensional vector.
Proceeding similarly, we have for the kth period

$$
\begin{equation*}
\mu_{k}^{*}\left(x_{k}\right)=\alpha_{k} \quad \underbrace{a}_{s_{N-1} \ldots S k+l}+b S k X k) \tag{4.39}
\end{equation*}
$$

where $\alpha_{k}$ is an n-dimensional vector that depends on the probability distributions of the rates of return $e_{i}^{k}$ of the risky assets and are determined by maximization in the DP Eq. (4.38). The corresponding cost-to-go functions satisfy

$$
\begin{equation*}
-\frac{J_{k}^{\prime}(x)}{J_{k}^{\prime \prime}(x)} \quad \bar{S} \bar{N}-\stackrel{a}{l}-\ldots S k \boldsymbol{b} \boldsymbol{x}, \quad k=0,1, \ldots, \mathrm{~N}-1 \tag{4.40}
\end{equation*}
$$

Thus it is seen that the investor, when faced with the opportunity to sequentially reinvest his wealth, uses a policy similar to that of the singleperiod case. Carrying the analysis one step further, it is seen that if the utility function $U$ is such that $a=0$, that is, $U$ has one of the forms

$$
\begin{aligned}
& \text { In } x, \quad \text { for } b \quad l, \\
& \underline{(b} \underline{=} \underline{1})(b x)^{1-(1 / b)}, \quad \text { for } b \neq 0, b \neq 1,
\end{aligned}
$$

then it follows from Eq. (4.39) that the investor acts at each stage $k$ as if he were faced with a single-period investment characterized by the rates of return $S k$ and $e_{i}^{k}$, and the objective function $E\{U(X k+l)\}$. This policy whereby the investor can ignore the fact that he will have the opportunity to reinvest his wealth is called a myopic policy.

Note that a myopic policy is also optimal when $S k=1$ for all $k$, which means that wealth is discounted at the rate of return of the riskless asset. Furthermore, it has been shown by Mossin [Mos68] that when $a=0$ a myopic policy is optimal even in the more general case where the rates of return $S k$ are independent random variables, and for the case where forecasts on the probability distributions of the rates of return $e_{i}^{k}$ of the risky assets become available during the investment process (see Exercise 4.14).

It turns out that even in the more general case where $a \neq 0$ only a small amount of foresight is required on the part of the decision maker. It can be seen [compare Eqs. (4.38)-(4.40)] that the optimal policy (4.39) at period $k$ is the one that the investor would use in a single-period problem to maximize over $u_{i}^{k}, \mathrm{i} \quad 1, \ldots, n$,

$$
E\left\{U\left(s_{N-1} \cdots s_{k+1} x_{k+1}\right)\right\}
$$

subject to

$$
X k+l=S k X k+\sum_{i=1}^{n}\left(e_{i}^{k}-s_{k}\right) u_{i}^{k}
$$

In other words, the investor at period $k$ should maximize the expected utility of wealth that results if amounts $u_{i}^{k}$ are invested in the risky assets in period $k$ and the resulting wealth $X k+l$ is subsequently invested exclusively in the riskless asset during the remaining periods $k+1, \ldots, \mathbf{N}-1$. This is known as a partially myopic policy. Such a policy can also be shown to be optimal when forecasts on the probability distributions of the rates of return of the risky assets become available during the investment process (see Exercise 4.14).

Another interesting aspect of the case where $a \neq 0$ is that if $S k>1$ for all $k$, then as the horizon becomes increasingly long ( $\mathrm{N} \rightarrow 00$ ), the policy in the initial stages approaches a myopic policy [compare Eqs. (4.39) and (4.40)]. Thus, for $S k>1$, a partially myopic policy becomes asymptotically myopic as the horizon tends to infinity.

### 4.4 OPTIMAL STOPPING PROBLEMS

Optimal stopping problems of the type that we will consider in this and subsequent sections are characterized by the availability, at each state, of a control that stops the evolution of the system. Thus at each stage the decision maker observes the current state of the system and decides whether to continue the process (perhaps at a certain cost) or stop the process and incur a certain loss. If the decision is to continue, a control must be selected from a given set of available choices. If there is only one choice other than stopping, then each policy is characterized at each period by the stopping set, that is, the set of states where the policy stops the system.

## Asset Selling

As a first example, consider a person having an asset (say a piece of land) for which he is offered an amount of money from period to period. We assume that the offers, denoted wo, $W I, \ldots, W N-I$, are random and independent, and take values within some bounded interval of nonnegative numbers ( $W k=0$ could correspond to no offer received during the period). If the person accepts an offer, he can invest the money at a fixed rate of interest $r>0$, and if he rejects the offer, he waits until the next period to consider the next offer. Offers rejected are not renewed, and we assume that the last offer $W N-1$ must be accepted if every prior offer has been rejected. The objective is to find a policy for accepting and rejecting offers that maximizes the revenue of the person at the Nth period.

The DP algorithm for this problem can be derived by elementary reasoning. As a modeling exercise, however, we will embed the problem in the framework of the basic problem by specifying the system and cost. We define the state space to be the real line, augmented with an additional state (call it $T$ ), which is a termination state. By writing that the system is at state $x k \quad T$ at some time $k \leq \mathrm{N}-1$, we mean that the asset has already been sold. By writing that the system is at a state $x k \neq T$ at some time $k \leq N-1$, we mean that the asset has not been sold as yet and the offer under consideration is equal to $x_{k}$ (and also equal to the kth offer $w k-l$ ). We take $x o=0$ (a fictitious "null" offer). The control space consists of two elements $u^{l}$ and $u^{2}$, corresponding to the decisions "sell" and "do not sell," respectively. We view $w k$ as the disturbance at time $k$.

With these conventions, we may write a system equation of the form

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}, w_{k}\right), \quad k=0,1, \ldots, \mathbf{N}-1
$$

where the function $f k$ is defined via the relation

$$
x_{k+1}= \begin{cases}T & \text { if } x k=T, \text { or if } x k \neq T \text { and } U k=u^{l} \text { (sell), } \\ w_{k} & \text { otherwise. }\end{cases}
$$

Note that a sell decision at time $k\left(U k \quad u^{1}\right.$ ) accepts the offer $W k-l$, and that no explicit sell decision is required to accept the last offer $w N-1$, as it must be accepted by assumption if the asset has not yet been sold. The corresponding reward function may be written as

$$
\underset{\substack{w_{k} \\ k=0,1, \ldots, N-1}}{E}\left\{g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(x_{k}, u_{k}, w_{k}\right)\right\}
$$

where

$$
\begin{gathered}
g_{N}\left(x_{N}\right)= \begin{cases}x_{N} & \text { if } x N \neq T \\
0 & \text { otherwise }\end{cases} \\
g_{k}\left(x_{k}, u_{k}, w_{k}\right)= \begin{cases}(1+r)^{N-k} x_{k} & \text { if } x k \neq T \text { and } U k=u^{1}(\text { sell }) \\
0 & \text { otherwise }\end{cases}
\end{gathered}
$$

Based on this formulation we can write the corresponding DP algorithm:

$$
\begin{align*}
& J_{N}\left(x_{N}\right)= \begin{cases}\text { oxN } & \text { if } x_{N} \neq T, \\
& \text { if } X N=T,\end{cases}  \tag{4.41}\\
& J_{k}\left(x_{k}\right)= \begin{cases}\max \left[(1+r)^{N}-k X k, E\left\{J_{k+1}(W k)\right\}\right] & \text { if } x k \neq T, \\
0 & \text { if } x k=T .\end{cases} \tag{4.42}
\end{align*}
$$

In the above equation, $(1+r) N-k X k$ is the revenue resulting from decision $u^{l}$ (sell) when the offer is $x k$, and $E\{J k+l(W k)\}$ represents the expected revenue corresponding to the decision $u^{2}$ (do not sell). Thus, the optimal policy is to accept an offer if it is greater than $E\left\{\mathrm{~J}_{k+1}(W k)\right\} /(1+\mathrm{r}) N-k$, which represents expected revenue discounted to the present time:

$$
\begin{array}{ll}
\text { accept the offer } x_{k} & \text { if } x_{k}>\alpha_{k} \\
\text { reject the offer } x k & \text { if } x_{k}<\alpha_{k}
\end{array}
$$

where

$$
\alpha k=\frac{E\{, J k+l(W k)\}}{(l+r) \bar{N}-\bar{k}} .
$$

When $x k=\alpha_{k}$, both acceptance and rejection are optimal. Thus the optimal policy is determined by the scalar sequence $\left\{\alpha_{k}\right\}$ (see Fig. 4.4.1).

## Properties of the Optimal Policy

We will now derive some properties of the optimal policy with some further analysis. Let us assume that the offers $w k$ are identically distributed, and to simplify notation, let us drop the time index $k$ and denote by $E_{w}\{$.$\} the$ expected value of the corresponding expression over $w k$, for all $k$. We will then show that
$\alpha_{k} \geq \alpha_{k+1}, \quad$ for all $k$,


Figure 4.4.1 Optimal policy for accepting offers in the asset selling problem.
which expresses the intuitive fact that if an offer is good enough to be acceptable at time $k$, it should also be acceptable at time $k+1$ when there will be one less chance for improvement. We will also obtain an equation for the limit of $\alpha_{k}$ as $k \rightarrow-\infty$.

Let us introduce the functions

$$
V k(X k)=\begin{gathered}
J k(X k) \\
\left(1 \pm \underline{1^{\prime}}\right) N-k^{\prime}
\end{gathered} \quad X k \neq T
$$

It can be seen from Eqs. (4.41) and (4.42) that

$$
\begin{equation*}
V N(X N)=X N \tag{4.43}
\end{equation*}
$$

$$
\begin{equation*}
V k(X k)=\max \left[x_{k},(1+r)^{-1} \underset{w}{E}\left\{V_{k+1}(w)\right\}\right], \quad k=0,1, \ldots, N-1 \tag{4.44}
\end{equation*}
$$

and that

$$
\alpha_{k}=\begin{gathered}
E w\{V k+l(w)\} \\
1+1^{\prime}
\end{gathered}
$$

To prove that $\alpha_{k} \geq \alpha_{k+1}$, note that from Eqs. (4.43) and (4.44), we have

$$
V_{N-1}(x) \geq V_{N}(x), \quad \text { for all } x \geq 0
$$

Applying Eq. (4.44) for $k=N^{\prime}-2$ and $k=N-1$, and using the preceding inequality, we obtain for all $x \geq 0$

$$
\begin{aligned}
V_{N-2}(x) & =\max \left[x,(1+r)^{-1} \underset{w}{E}\left\{V_{N-1}(w)\right\}\right] \\
& \geq \max \left[X,(1+\mathrm{r})^{-1} \underset{w}{E}\{V N(W)\}\right] \\
& =\operatorname{VN}-l(X) .
\end{aligned}
$$

Sec. 4.4

Continuing in the same manner, we see that

$$
V_{k}(x) \geq V_{k+1}(x), \quad \text { for all } x \geq 0 \text { and } k
$$

Since $\alpha_{k}=E_{w}\{V k+l(W)\} /\left(l+l^{\prime}\right)$, we obtain $\alpha_{k} \geq \alpha_{k+1}$, as desired.
Let us now see what happens when the horizon $N$ is very large. From the algorithm (4.43) and (4.44) we have

$$
\begin{equation*}
V k(X k)=\max \left(x k, \alpha_{k}\right) \tag{4.45}
\end{equation*}
$$

Hence we obtain

$$
\begin{aligned}
\alpha_{k} & =\begin{array}{c}
1 \\
1+1^{\prime}
\end{array}\left\{V_{k+1}(W)\right\} \\
& =-\underset{+l^{\prime}}{\mathbf{1}}-\int_{0}^{\alpha_{k+1}} \alpha_{k+1} d P(w)+-\underset{l^{\prime}}{\mathbf{1}}-\int_{\alpha_{k+1}}^{\infty} w d P(w)
\end{aligned}
$$

where the function $P$ is defined for all scalars $\lambda$ by

$$
P(\lambda)=\operatorname{Prob}\{w<\lambda\}
$$

The difference equation for $\alpha_{k}$ may also be written as

$$
\alpha_{k}=\begin{gather*}
P\left(\alpha_{k+1}\right)  \tag{4.46}\\
1+\iota
\end{gather*} \alpha_{k+1}+\frac{1}{1+\iota} \int_{\alpha_{k+1}}^{\infty} w d P(w), \quad \text { for all } k
$$

with $\alpha_{N}=0$
Now since we have

$$
\begin{gathered}
\mathrm{O} \frac{P(\alpha)}{1+} \frac{1}{1},-\frac{1}{-1}<\mathbf{1}, \quad \text { for all } \alpha \geq 0, \\
\boldsymbol{O}<\frac{1}{1}+r \int_{\alpha_{k+1}}^{\infty} w d P(w)<\frac{E}{\left.1+\frac{w}{\prime}\right\}} \quad \text { for all } k,
\end{gathered}
$$

it can be seen, using the property $\alpha_{k} \geq \alpha_{k+1}$, that the sequence $\left\{\alpha_{k}\right\}$ generated (backward) by the difference equation (4.46) converges (as $k \rightarrow$ $-\infty)$ to a constant $\bar{\alpha}$ satisfying

$$
(1+r) \bar{\alpha}=P(\bar{\alpha}) \bar{\alpha}+\int_{\bar{\alpha}}^{\infty} w d P(w)
$$

This equation is obtained from Eq. (4.46) by taking limits as $k \rightarrow-\infty$ and by using the fact that $P$ is continuous from the left.

Thus, when the horizon $N$ tends to become longer, the optimal policy for every fixed $k \geq 1$ can be approximated by the stationary policy

$$
\begin{array}{ll}
\text { accept the offer } x_{k} & \text { if } X k>\overline{\boldsymbol{\alpha}}, \\
\text { reject the offer } x_{k} & \text { if } X k<\overline{\boldsymbol{\alpha}} .
\end{array}
$$

The optimality of this policy for the corresponding infinite horizon problem will be shown in Section 7.3.

## ancuman with a Deadline

Let us consider another stopping problem that has similar nature. Assume that a certain quantity of raw material is required by a certain time. If the price of this material fluctuates, there arises the problem of deciding whether to purchase at the current price or wait a further period, during which the price may go up or down. We thus want to minimize the expected price of purchase. We assume that successive prices $W k$ are independent and identically distributed with distribution $P(W k)$, and that the purchase must be made within $N$ time periods.

This problem and the earlier asset selling problem have obvious similarities. Let us denote by

$$
x k+1=w k
$$

the price prevailing at the beginning of period $k+1$. We have similar to the earlier problem the DP algorithm

$$
\begin{gathered}
I N(X N)=X N, \\
J k(X k)=\min \left[X k^{\prime} E\{J k+1(W k)\}\right] .
\end{gathered}
$$

Note that $J_{k(X k)}$ is the optimal cost-to-go when the current price is $X k$ and the material has not been purchased yet. To be strictly formal, we should introduce a termination state $T$, to which the system moves following a purchasing decision and at which the system subsequently stays at no cost. A nonzero cost is incurred only when moving from $x k$ to $T$; this cost is equal to $x k$. Thus the cost-to-go from the termination state $T$ is 0 , and for this reason it was neglected in the preceding DP equation.

The optimal policy is given by

$$
\text { purchase if } x k<\alpha_{k},
$$

do not purchase if $x k>\alpha_{k}$,
where

$$
\alpha_{k} \quad E\{J k+1(W k)\}
$$

Similar to the asset selling problem, the thresholds $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{N-1}$ can be obtained from the discrete-time equation

$$
\alpha_{k}=\alpha_{k+1}\left(1-P\left(\alpha_{k+1}\right)\right)+\int_{0}^{\alpha_{k+1}} w d P(w)
$$

with the terminal condition

$$
\alpha_{N-1} \quad \int_{0}^{\infty} w d P(w)=E\{w\}
$$



Figure 4.4.2 Structure of the cost-to-go function $J_{N-1}\left(x_{N-1}\right)$ when prices are correlated.

## The Case of Correlated Prices

Consider now a variation of the purchasing problem where we do not assume that the successive prices $w o, \ldots, w_{N-1}$ are independent. Instead, we assume that they are correlated and can be represented as the state of a linear system driven by independent disturbances (cf. Section 1.4). In particular, we have

$$
w_{k}=x_{k+1}, \quad k=0,1, \ldots, N \quad 1,
$$

with

$$
x_{k+1}=\lambda x_{k}+\xi_{k}, \quad X o=0,
$$

where $\lambda$ is a scalar with $\leq \lambda<1$ and $\xi_{0}, \xi_{1}, \ldots, \xi_{N-1}$ are independent identically distributed random variables taking positive values with given probability distribution. As discussed in Section 1.4, the DP algorithm under these circumstances takes the form

$$
J_{N}\left(x_{N}\right)=x_{N}
$$

$$
J_{k}\left(x_{k}\right)=\min \left[x_{k}, E\left\{J_{k+1}\left(\lambda x_{k}+\xi_{k}\right)\right\}\right]
$$

where the cost associated with the purchasing decision is $X k$ and the cost associated with the waiting decision is $E\left\{{ }_{J k+1}\left(\lambda x_{k}+\xi_{k}\right)\right\}$.

We will show that in this case the optimal policy is of the same type as the one for independent prices. Indeed, we have

$$
J_{N-1}\left(x_{N-1}\right)=\min \left[x_{N-1}, \lambda x_{N-1}+\bar{\xi}\right]
$$

where $\bar{\xi}=E\left\{\xi_{N-1}\right\}$. As shown in Fig. 4.4.2, an optimal policy at time $N-1$ is given by

$$
\text { purchase if } X N-I<\alpha_{N-1}
$$

$$
\text { do not purchase if } X N-I>\alpha_{N-1}
$$

where $\alpha_{N-1}$ is defined from the equation $\alpha_{N-1} \quad \lambda \alpha_{N-1}+\bar{\xi}$ :

$$
\alpha_{N-1}=\frac{\bar{\xi}}{1-\lambda}
$$

Note that

$$
J_{N-1}(x) \leq J_{N}(x), \quad \text { for all } x
$$

and that $\mathbf{I N}-\mathbf{I}$ is concave and increasing in $x$. Using this fact in the DP algorithm, one may show (the monotonicity property of DP; Exercise 1.23 in Chapter 1) that

$$
J_{k}(x) \leq J_{k+1}(x), \quad \text { for all } x \text { and } k
$$

and that $J k$ is concave and increasing in $x$ for all $k$. Furthermore, since $\bar{\xi}=E\left\{\xi_{k}\right\}>0$ for all $k$, one can show that

$$
E\left\{J_{k+1}\left(\xi_{k}\right)\right\}>0, \quad \text { for all } k
$$

These facts imply (as illustrated in Fig. 4.4.3) that the optimal policy for every period $k$ is of the form

$$
\text { purchase if } X k<\alpha_{k}
$$

$$
\text { do not purchase if } X k>\alpha_{k}
$$

where the scalar $\alpha_{k}$ is the unique positive solution of the equation

$$
x=E\left\{J_{k+1}\left(\lambda x+\xi_{k}\right)\right\}
$$

Note that the relation $J k(X) \leq J k+I(X)$ for all $x$ and $k$ implies that

$$
\alpha_{k-1} \leq \alpha_{k} \leq \alpha_{k+1}, \quad \text { for all } k
$$

and hence (as one would expect) the threshold price to purchase increases as the deadline gets closer.


Figure 4.4.3 Determining the optimal policy when prices are correlated.

## General Stopping Problems and the One-Step-Lookahead Rule

We now formulate a general type of N -stage problem where stopping is mandatory at or before stage $N$. Consider the stationary version of the basic problem of Chapter 1 (state, control, and disturbance spaces, disturbance distribution, control constraint set, and cost per stage are the same for all times). Assume that at each state $X k$ and at time $k$ there is available, in addition to the controls $U k \mathrm{E} U(X k)$, a stopping action that forces the system to enter a termination state at a cost $t(X k)$ and subsequently remain there at no cost. The terminal cost, assuming stopping has not occurred by the last stage, is $t(X N)^{\prime}$ Thus, in effect, we assume that the termination cost will always be incurred either at the end of the horizon or earlier.

The DP algorithm is given by

$$
\begin{gather*}
J_{N}\left(x_{N}\right)=t\left(x_{N}\right)  \tag{4.47}\\
J_{k}\left(x_{k}\right)=\min \left[t\left(x_{k}\right), \min _{u_{k} \in U\left(x_{k}\right)} E\left\{g\left(x_{k}, u_{k}, w_{k}\right)+J_{k+1}\left(f\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}\right] \tag{4.48}
\end{gather*}
$$

and it is optimal to stop at time $k$ for states $x$ in the set

$$
T k=\left\{x, t(x) \leq \min _{u E U(x)} E\{g(x, u, w)+J k+l(!(X, U, w))\}\right\}
$$

We have from Eqs. (4.47) and (4.48)

$$
J_{N-1}(x) \leq J_{N}(x), \quad \text { for all } x
$$

and using this fact in the DP equation (4.48), we obtain inductively

$$
J_{k}(x) \leq J_{k+1}(x), \quad \text { for all } x \text { and } k .
$$

[We are making use here of the stationarity of the problem and the monotonicity property of DP (Exercise 1.23 in Chapter 1).] Using this fact and the definition of $T k$ we see that

$$
T o \mathrm{C} \ldots \text { С } T_{k} \mathrm{C} T k+l \text { С } \ldots \text { С } T N-I .
$$

We will now consider a condition guaranteeing that all the stopping sets $T k$ are equal. Suppose that the set $T N-l$ is absorbing in the sense that if a state belongs to $T N_{-1}$ and termination is not selected, the next state will also be in $T N-1$ :

$$
\begin{equation*}
f(x, u, w) \text { E } T N-I, \quad \text { for all } x \in T N-I, U \mathrm{E} U(x), w . \tag{4.50}
\end{equation*}
$$

We will show that equality holds in Eq. (4.49) and for all $k$ we have

$$
T k=T N-l \quad\left\{x \in S \mid t(x) \leq \min _{u E U(x)} E\{9(x, u, w)+t(f(X, U, w))\}\right\}
$$

To see this, note that by the definition of $T N-I$, we have

$$
I N-l(X)=t(x), \quad \text { for all } x \in T N-I,
$$

and using Eq. (4.50) we obtain for $x$ E $T N-l$

$$
\begin{aligned}
\min _{u E U(x)} E\{g(x, u, w) & +I N-l(f(x, u, w))\} \\
& =\min _{u E U(x)} E\{g(x, u, w)+t(f(x, u, w))\} \\
& \geq t(x)
\end{aligned}
$$

Therefore, stopping is optimal for all $X N-2$ E $T N$-lor equivalently $T N-l$ C $T N-2$. This together with Eq. (4.49) implies $T N-2=T N-I$. Proceeding similarly, we obtain $T k=T N-l$ for all $k$.

In conclusion, if condition (4.50) holds (the one-step stopping set $T N-l$ is absorbing), then the stopping sets $T k$ are all equal to the set of states for which it is better to stop rather than continue for one more stage and then stop. A policy of this type is known as a one-step lookahead pol'icy. Such a policy turns out to be optimal in several types of applications. We provide next some examples. Additional examples are given in the exercises, and in Vol. Chapter 3.

## Example 4.4.1 (Asset Selling with Past Offers Retained)

Consider the asset selling problem discussed earlier in this section with the difference that rejected offers can be accepted at a later time. Then if the asset is not sold at time $k$ the state evolves according to

$$
x_{k+1}=\max \left(x_{k}, w_{k}\right)
$$

instead of $X k+l=W k$. The DP equations (4.43) and (4.44) then become

$$
V_{N}\left(x_{N}\right)=x_{N}
$$

$$
V k(X k)=\max \left[X k^{\prime}(1+r)-l E\{V k+1(\max (x k, W k))\}\right] .
$$

The one-step stopping set is

$$
T_{N-1}=\{x \mid x \geq(1+r)-l E\{\max (\mathrm{x}, w)\}\}
$$

It is seen [compare with Eqs. (4.45) and (4.46)] that an alternative characterization is

$$
\begin{equation*}
\text { TN-I }=\{x \mid x \geq \bar{\alpha}\} \tag{4.51}
\end{equation*}
$$

where $\bar{\alpha}$ is obtained from the equation

$$
(1+r) \bar{\alpha}=P(\bar{\alpha}) \bar{\alpha}+\int_{\bar{\alpha}}^{\infty} w d P(w)
$$

Since past offers can be accepted at a later date, the effective offer available cannot decrease with time, and it follows that the one-step stopping set (4.51) is absorbing in the sense of Eq. (4.50). Therefore, the one-step lookahead stopping rule that accepts the first offer that equals or exceeds $\bar{\alpha}$ is optimal. Note that this policy is independent of the horizon length N .

## Example 4.4.2 Rational Burglar [Whi82])

A burglar may at any night $k$ choose to retire with his accumulated earnings $X k$ or enter a house and bring home an amount $W k$. However, in the latter case he gets caught with probability $p$, and then he is forced to terminate his activities and forfeit his earnings thus far. The amounts $W k$ are independent, identically distributed with mean $\bar{w}$. The problem is to find a policy that maximizes the burglar's expected earnings over N nights.

We can formulate this problem as a stopping problem with two actions (retire or continue) and a state space consisting of the real line, the retirement state, and a special state corresponding to the burglar getting caught. The DP algorithm is given by

$$
\begin{gathered}
J_{N}\left(x_{N}\right)=x_{N} \\
J_{k}\left(x_{k}\right)=\max \left[x_{k},(1-p) E\left\{J_{k+1}\left(x_{k}+w_{k}\right)\right\}\right] .
\end{gathered}
$$

The one-step stopping set is

$$
T_{N-1}=\{x \mid x \geq(1-p)(x+\bar{w})\}=\left\{x \left\lvert\, x \geq \frac{(1-p) \bar{w}}{p}\right.\right\}
$$

(more accurately this set together with the special state corresponding to the burglar's arrest). Since this set is absorbing in the sense of Eq. (4.50), we see that the one-step lookahead policy by which the burglar retires when his earnings reach or exceed $(1-p) \bar{w} / p$ is optimal. The optimality of this policy for the corresponding infinite horizon problem will be demonstrated in Vol. II, Chapter 3.

### 4.5 SCHEDULING AND THE INTERCHANGE ARGUMENT

Suppose one has a collection of tasks to perform but the ordering of the tasks is subject to optimal choice. As examples, consider the ordering of operations in a construction project so as to minimize construction time or the scheduling of jobs in a workshop so as to minimize machine idle time. In such problems a useful technique is to start with some schedule and then to interchange two adjacent tasks and see what happens. We first provide some examples, and we then formalize mathematically the technique.

## Example 4.5.1 (The Quiz Problem)

Consider a quiz contest where a person is given a list of $N$ questions and can answer these questions in any order he chooses. Question i will be answered correctly with probability $p i$, and the person will then receive a reward $R_{i}$. At the first incorrect answer, the quiz terminates and the person is allowed to keep his previous rewards. The problem is to choose the ordering of questions so as to maximize expected rewards.

Let i and j be the kth and $(k+1)$ st questions in an optimally ordered list

$$
L=\left(i_{0}, \ldots, i_{k-1}, i, j, i_{k+2}, \ldots, i_{N-1}\right)
$$

Consider the list

$$
L^{\prime}=(i a, \ldots, i k-1, j, i, i k+2, \ldots, i N-1)
$$

obtained from $L$ by interchanging the order of questions $i$ and $j$. We compare the expected rewards of Land $L^{\prime}$. We have
$\mathrm{E}\{$ reward of $L\}=\mathrm{E}\left\{\right.$ reward of $\left.\left\{i a, \ldots, i_{k-1}\right\}\right\}$

$$
+P i a^{\prime \prime \prime} P i_{k_{-} l}(P i R i+P i P j R j)
$$

+ Pia'" Pik_1PiPjE\{reward of $\left.\left\{i k+2, \ldots, i_{N-1}\right\}\right\}$
$\mathrm{E}\left\{\right.$ reward of $\left.L^{\prime}\right\}=\mathrm{E}\left\{\right.$ reward of $\left.\left\{i a, \ldots, i_{k-1}\right\}\right\}$

$$
\begin{aligned}
& \text { +Pia'" Pik-l }(p j R j+p j P i R i) \\
& + \text { Pia' }^{\prime \prime \prime} P^{\prime} i k_{-} l P j P i E\left\{\text { reward of }\left\{i k+2, \ldots, i_{N-1}\right\}\right\} .
\end{aligned}
$$

Since $L$ is optimally ordered, we have
$\mathrm{E}\{$ reward of $L\} \geq \mathrm{E}\left\{\right.$ reward of $\left.L^{\prime}\right\}$,
so it follows from these equations that

$$
p_{i} R_{i}+p_{i} p_{j} R_{j} \geq p_{j} R_{j}+p_{j} p_{i} R_{i}
$$

or equivalently

$$
\frac{P i R i}{1-P i}>p j R j
$$

Therefore, to maximize expected rewards, questions should be answered in decreasing order of $p_{i} R_{i} /\left(\begin{array}{ll}1 & P i\end{array}\right)$.

## Example 4.5.2 (Job Scheduling on a Single Processor)

Suppose we have $N$ jobs to process in sequential order with the ith job requiring a random time $T_{i}$ for its execution. The times $T_{1}, \ldots, T N$ are independent. If job i is completed at time $t$, the reward is $\boldsymbol{\alpha}^{t} R i$, where $\boldsymbol{\alpha}$ is a discount factor with $0<\alpha<1$. The problem is to find a schedule that maximizes the total expected reward.

It can be seen that the state for this problem is just the collection of jobs yet to be processed. Indeed, because the execution times $T_{i}$ are independent, and also because future costs are multiplicatively affected through discounting by the completion times of preceding jobs, the optimization of the scheduling of future jobs is unaffected by the completion times of preceding jobs. As a result these times need not be included in the state; this would not be so if either the times $T_{i}$ were correlated or if the reward for completing job i at time $t$ were not $\alpha^{t} R_{i}$ but instead had a general dependence on $t$. Now, given that the state is the collection of jobs yet to be processed, it is clear that an optimal policy can be mapped into an optimal job schedule (ia, ... , iN-1).

Suppose that $L=(i a, \ldots, i k-1, \mathrm{i}, \mathrm{j}, i k+2, \ldots, i N-1)$ is an optimal job schedule, and consider the schedule $L^{\prime}=(i a, \ldots, i k-1, j, i, i k+2, \ldots, i \mathrm{~N}-1)$ obtained by interchanging i and $\mathbf{j}$. Let $t k$ be the time of completion of job ik-1. We compare the rewards of the schedules Land $L^{\prime}$, similar to the preceding example. Since the reward for completing the remaining jobs $i k+2, \ldots, \mathrm{iN}-1$ is independent of the order in which jobs $i$ and $j$ are executed, we obtain

$$
E\left\{\alpha^{t_{k}+T_{i}} R_{i}+\alpha^{t_{k}+T_{i}+T_{j}} R_{j}\right\} \geq E\left\{\alpha^{t_{k}+T_{j}} R_{j}+\alpha^{t_{k}+T_{j}+T_{i}} R_{i}\right\}
$$

Since $t k, T_{i}$, and $T_{j}$ are independent, this relation can be written as

$$
\begin{aligned}
E\left\{\alpha^{t_{k}}\right\}\left(E\left\{\alpha^{T_{i}}\right\} R_{i}\right. & \left.+E\left\{\alpha^{T_{i}}\right\} E\left\{\alpha^{T_{j}}\right\} R_{j}\right) \\
& \geq E\left\{\alpha^{t_{k}}\right\}\left(E\left\{\alpha^{T_{j}}\right\} R_{j}+E\left\{\alpha^{T_{j}}\right\} E\left\{\alpha^{T_{i}}\right\} R_{i}\right)
\end{aligned}
$$

from which we finally obtain

$$
\frac{E\left\{\alpha^{T_{i}}\right\} R_{i}}{1-E\left\{\alpha^{T_{i}}\right\}} \geq \frac{E\left\{\alpha^{T_{j}}\right\} R_{j}}{1-E\left\{\alpha^{T_{j}}\right\}}
$$

It follows that scheduling jobs in order of decreasing $E\left\{\alpha^{T_{i}}\right\} R_{i} /\left(1-E\left\{\alpha^{T_{i}}\right\}\right)$ maximizes expected rewards. The structure of the optimal policy is identical with the one we derived for the preceding quiz contest example (identify $E\left\{\alpha^{T_{i}}\right\}$ with the probability pi of answering correctly question i).

## Example 4.5.3 (Job Scheduling on Two Processors in Series)

Consider the scheduling of $N$ jobs on two processors $A$ and $B$, such that $B$ accepts the output of $A$ as input. Job i requires known times ai and $b_{i}$ for processing in $A$ and $B$, respectively. The problem is to find a schedule that minimizes the total processing time.

To formulate the problem into the form of the basic problem, we increment discrete time at the moments when processing of a job is completed at machine $A$ and the next job is started. We take as state at time $k$ the collection of jobs $X_{k}$ that remain to be processed at $A$, together with the backlog of work $T k$ at machine $B$, that is, the amount of time that the jobs currently at $B$ need to clear $B$. Thus if ( $X_{k}, T k$ ) is the state at stage $k$ and job i is completed at machine $A$, the state changes to $\left(X_{k+1}, T k+l\right)$ given by

$$
X_{k+1}=X_{k}-\{i\}, \quad \tau_{k+1}=b_{i}+\max \left(0, \tau_{k}-a_{i}\right)
$$

The corresponding DP algorithm is

$$
J_{k}\left(X_{k}, \tau_{k}\right)=\min _{i \in X_{k}}\left[a_{i}+J_{k+1}\left(X_{k}-\{i\}, b_{i}+\max \left(0, \tau_{k}-a_{i}\right)\right)\right]
$$

with the terminal condition

$$
I N(0, T N)=T N,
$$

where $\varnothing$ is the empty set.
Since the problem is deterministic, there exists an optimal open-loop schedule

$$
\left\{i_{0}, \ldots, i_{k-1}, i, j, i_{k+2}, \ldots, i_{N-1}\right\}
$$

By arguing that the cost of this schedule is no worse than the cost of the schedule

$$
\left\{i_{0}, \ldots, i_{k-1}, j, i, i_{k+2}, \ldots, i_{N-1}\right\}
$$

obtained by interchanging $i$ and $j$, it can be verified that

$$
\begin{equation*}
J_{k+2}\left(X_{k}-\{i\}-\{j\}, \tau_{i j}\right) \leq J_{k+2}\left(X_{k}-\{i\}-\{j\}, \tau_{j i}\right), \tag{4.52}
\end{equation*}
$$

where $T_{i j}$ and $T_{j i}$ are the backlogs at machine $B$ at time $k+2$ when i is processed before $j$ and $j$ is processed before $i$, respectively, and the backlog at time $k$ was $T k$. A straightforward calculation shows that

$$
\begin{align*}
& \tau_{i j}=b_{i}+b_{j}-a_{i}-a_{j}+\max \left(\tau_{k}, a_{i}, a_{i}+a_{j}-b_{i}\right),  \tag{4.53}\\
& \tau_{j i}=b_{j}+b_{i}-a_{j}-a_{i}+\max \left(\tau_{k}, a_{j}, a_{j}+a_{i}-b_{j}\right) \tag{4.54}
\end{align*}
$$

Clearly, $J k+2$ is monotonically increasing in $T$, so from Eq. (4.52) we obtain

$$
\tau_{i j} \leq \tau_{j i}
$$

In view of Eqs. (4.53) and (4.54), this relation implies two possibilities. The first is

$$
\begin{aligned}
& \tau_{k} \geq \max \left(a_{i}, a_{i}+a_{j}-b_{i}\right), \\
& \tau_{k} \geq \max \left(a_{j}, a_{j}+a_{i}-b_{j}\right),
\end{aligned}
$$

in which case $T_{i j}=T_{j i}$ and the order of $i$ and $j$ makes no difference. (This is the case where the backlog at time $k$ is so large that both jobs $i$ and $j$ will find $B$ working on an earlier job.) The second possibility is that

$$
\max \left(a_{i}, a_{i}+a_{j}-b_{i}\right) \leq \max \left(a_{j}, a_{j}+a_{i}-b_{j}\right)
$$

which can be seen to be equivalent to

$$
\min \left(a_{i}, b_{j}\right) \leq \min \left(a_{j}, b_{i}\right)
$$

A schedule satisfying these necessary conditions for optimality can be constructed by the following procedure:

1. Find mini min $(a i, b i)^{\prime}$
2. If the minimizing value is an $a$ take the corresponding job first; if it is a $b$, take the corresponding job last.
3. Repeat the procedure with the remaining jobs until a complete schedule is constructed.
To show that this schedule is indeed optimal, we start with an optimal schedule. We consider the job io that minimizes min(ai' $b_{i}$ ) and by successive interchanges we move it to the same position as in the schedule constructed previously. It is seen from the preceding analysis that the resulting schedule is still optimal. Similarly, continuing through successive interchanges and maintaining optimality throughout, we can transform the optimal schedule into the schedule constructed earlier. We leave the details to the reader.

The Interchange Argument $^{\text {a }}$
Let us now consider the basic problem of Chapter 1 and formalize the interchange argument used in the preceding examples. The main requirement is that the problem has structure such that there exists an open-loop policy that is optimal, that is, a sequence of controls that performs as well or better than any sequence of control functions. This is certainly true in deterministic problems as discussed in Chapter 1, but it is also true in some stochastic problems such as those of Examples 4.5.1 and 4.5.2.

To apply the interchange argument, we start with an optimal sequence

$$
\left\{u o, \ldots, u_{k-1}, \bar{u}, \tilde{u}, u_{k+2}, \ldots, U N-I\right\}
$$

and focus attention on the controls $\bar{u}$ and $\tilde{u}$ applied at times $k$ and $k+1$, respectively. We then argue that if the order of $\bar{u}$ and $\tilde{u}$ is interchanged the expected cost cannot decrease. In particular, if $X_{k}$ is the set of states that can occur with positive probability starting from the given initial state $X o$ and using the control subsequence $\left\{u o, \ldots, u_{k-1}\right\}$, we must have for all Xk E $X k$

$$
\begin{align*}
& E\left\{g k(X k, \bar{u}, W k)+g k+l\left(\bar{x}_{k+1}, \tilde{u}, W k+l\right)+J_{k+2}^{*}\left(\bar{x}_{k+2}\right)\right\} \\
& \quad \leq E\left\{9 k(X k, \tilde{u}, W k)+g k+l\left(\tilde{x}_{k+1}, \bar{u}, W k+l\right)+J k+2(X k+2)\right\} \tag{4.55}
\end{align*}
$$

where $\bar{x}_{k+1}$ and $\bar{x}_{k+2}$ (or $\tilde{x}_{k+1}$ and $\tilde{x}_{k+2}$ ) are the states subsequent to $X k$ when $U k \quad \bar{u}$ and $U k+l=\tilde{u}$ (or $U k=\tilde{u}$ and $U k+l=\bar{u}$, respectively) are applied, and $J K^{+}{ }_{2}(\cdot)$ is the optimal cost-to-go function for time $k+2$.

Relation (4.55) is a necessary condition for optimality. It holds for every $k$ and every optimal policy that is open-loop. There is no guarantee that this necessary condition is powerful enough to lead to an optimal solution, but it is worth considering in some specially structured problems. Generally in scheduling problems, algorithms that aim to improve a suboptimal schedule through a sequence of interchanges, may not provide an optimal solution, but are often the basis for successful heuristics.

### 4.6 SET-MEMBERSHIP DESCRIPTION OF UNCERTAINTY

In this section, we focus on problems where the uncertain quantities are described by their membership in given sets rather than probability distributions. This type of description is appropriate in minimax control problems, as discussed in Section 1.6. Our purpose in this section is to analyze some basic problems of estimation and control involving uncertainty with a set-membership description. These problems are conceptually important, and arise in several contexts, including the model predictive control methodology discussed in Section 6.5. However, their general solution can be computationally difficult. We will discuss some easily implementable approximations, involving linear systems and ellipsoidal descriptions.

### 4.6.1 Set-Membership Estimation

Suppose that we are given a linear dynamic system of the type considered in Section 4.1 but without a control $(U k \equiv 0)$ :

$$
x_{k+1}=A_{k} x_{k}+w_{k}, \quad k=0,1, \ldots, \mathbf{N}-1,
$$

where $X k E \Re^{n}$ and $W k E \Re^{n}$ denote the state and disturbance vectors, respectively, and the matrices $A k$ are known. Suppose also that at each time $k$, we receive a measurement $Z k E \Re^{s}$ of the form

$$
z_{k}=C_{k} x_{k}+v_{k}
$$

where $V k E \Re^{s}$ is an (unknown) observation noise vector, and the matrix $C k$ is given.

An important and generic problem is to estimate the value of $X k$, given the observations $\mathrm{zl}, \ldots, Z k$, accumulated up to time $k$. The uncertain quantities here are the initial state $x o$, the system disturbances $W o, \cdot \cdot, W N-l$, and the observation noise vectors vi, $\ldots, V N$, When the joint probability distribution for these vectors is given, one may calculate the conditional distribution of $X k$ given $Z 1, \ldots, Z k$, and from this, obtain estimates such as for example the conditional expectation $E\{X k \mid \mathrm{zl}, \ldots, Z k\}$. This approach leads to a rich theory, centered around the Kalman filtering algorithm, which is described in detail in Appendix E.

Suppose now that, instead of a probability distribution, we have a set $\mathcal{R}$ within which the vector of unknown quantities

$$
r=\left(x o, W o, \ldots, W N-1, \mathrm{VI}^{\prime}, \ldots, v N\right)
$$

is known to belong. The state $X k$ can be expressed in terms of $r$ using the system equation as

$$
X k=A k-1 \ldots A \text { oxo }+\sum_{\mathrm{i}=0}^{\mathrm{k}-1} A k^{-} 2 \ldots A_{i+1} W i
$$

or more abstractly as

$$
x_{k}=L_{k} r
$$

where $L k$ is an appropriate matrix. Thus, knowing that $r \mathrm{E} \mathcal{R}$ and before any measurements are received, the state $X k$ is known to belong to the set

$$
\mathcal{X}_{k}=L_{k} \mathcal{R}=\left\{L_{k} r \mid r \in \mathcal{R}\right\}
$$

Each measurement $z_{i}$, when recerved, roctruts the set of possible values of $r$ to be such that $z_{i}=C_{i} L_{i} r+v_{i}$ or

$$
z_{i}=E_{i} r
$$

for an appropriate matrix Ei . Thus, with each new measurement, the set of possible vectors $r$ is further restricted, and so is the set of possible states. In particular, given the measurements $\mathrm{Zl}, \ldots, Z k$, the set of possible vectors $r$ is given by

$$
\begin{equation*}
\mathcal{R}_{k}\left(z_{1}, \ldots, z_{k}\right)=\mathcal{R} \cap\left\{r \mid z_{1}=E_{1} r\right\} \cap \cdots \cap\left\{r \mid z_{k}=E_{k} r\right\} \tag{4.56}
\end{equation*}
$$

and by a linear transformation, yields the set of possible states $X k$ as

$$
\begin{equation*}
\mathcal{X}_{k}\left(z_{1}, \ldots, z_{k}\right)=L_{k} \mathcal{R}_{k}\left(z_{1}, \ldots, z_{k}\right) \tag{4.57}
\end{equation*}
$$

The procedure just described is straightforward, and can easily be extended to systems and measurements that are nonlinear. The difficulty, however, is to specify conveniently the sets $\mathcal{R}_{k}\left(\mathbf{Z 1}{ }^{\prime} \ldots, Z k\right)$ and/or $X k(Z I, " " Z k)$. There are only few special cases where these sets admit a simple description, e.g., one that involves a finite set of numbers. The most interesting of these are:
(a) The polyhedral case, where the set $\mathcal{R}$ is a polyhedron (a set specified by a finite number of linear inequalities). Then, it can be seen that the sets $\operatorname{Rk}(Z I, ' ", Z k)$ and/or $\mathcal{X}_{k}\left(z_{1}, \ldots, Z k\right)$ are also polyhedra. The reason is that the intersection of a polyhedron with a linear manifold (a translated subspace) is a polyhedron, and a linear transformation of a polyhedron yields another polyhedron [cf. Eqs. (4.56) and (4.57)].
(b) The ellipsoidal case, where the set $\mathcal{R}$ is an ellipsoid (a linearly transformed sphere - a more specific description is given later). Then, it can be shown that the sets $\operatorname{Rk}(Z I, \ldots, Z k)$ and/or $\mathcal{X}_{k}\left(z_{1}, \ldots, Z k\right)$ are also ellipsoids. Similar to the polyhedral case, the reason is that the intersection of an ellipsoid with a linear manifold is an ellipsoid, and a linear transformation of this ellipsoid yields another ellipsoid.
The polyhedral case is interesting in some cases, but suffers from a quick explosion of the computational requirements to describe the associated polyhedra, as $k$ increases. We will focus instead on the ellipsoidal case, and we will use DP methods to derive easily implementable algorithms that resemble the Kalman filtering algorithm described in Appendix E.

## Energy Constraints

We first consider the most favorable case where the set of possible states $\mathcal{X}_{k}(\mathrm{Z}, \ldots, Z \mathrm{~K})$ turns out to be an ellipsoid. Suppose that the vector of unknown quantities $r$ is known to belong to a set of the form
$\mathcal{R}=\left\{r \mid\left(x_{0}-\hat{x}_{0}\right)^{\prime} S^{-1}\left(x_{0}-\hat{x}_{0}\right)+\sum_{i=0}^{N-1}\left(w_{i}^{\prime} M_{i}^{-1} w_{i}+v_{i+1}^{\prime} N_{i+1}^{-1} v_{i+1}\right) \leq 1\right\}$,
where $S, I V l$, and $N i$ are positive definite symmetric matrices, and $\hat{x}_{0}$ is a given vector. This set is a bounded ellipsoid. Much of the analysis that follows carries through in the case of more general ellipsoids, but for simplicity, we restrict attention to bounded ellipsoids.

Let us use DP to derive the set of possible states and show that it is an ellipsoid of the form

$$
\mathcal{X}_{k}\left(z_{1}, \ldots, z_{k}\right)=\left\{x_{k} \mid\left(x_{k}-\hat{x}_{k}\right)^{\prime} \Sigma_{k}^{-1}\left(x_{k}-\hat{x}_{k}\right) \leq 1-\delta_{k}\right\},
$$

where
$\Sigma_{k}$ is a positive definite symmetric matrix that is independent of the observations $\mathrm{Zl}, \ldots, Z k$,
$\hat{x}_{k}$ is a vector that depends on $\mathrm{Zl}, \ldots, Z k$,
$\delta_{k}$ is a positive scalar that depends on $\mathrm{Zl}, \ldots, Z k$.
We observe that a vector $\xi$ belongs to $\mathcal{X}_{k}\left(z_{1}, \ldots, Z k\right)$ if and only if there exist $X o$ and $W o, \ldots, W k-I$ such that

$$
\begin{aligned}
\left(x o-\hat{x}_{0}\right)^{\prime} S-l\left(x o-\hat{x}_{0}\right) & +w_{i=\mathrm{O}}^{k-1} w_{i}^{\prime} M_{i}^{-1} w_{i} \\
& +\sum_{\mathrm{i}=\mathrm{O}}^{k-1}\left(z_{i+1}-C i+l X_{H}\right) \cdot N_{i+1}^{-1}(\mathrm{ZH} 1-C H 1 X i+l) \leq 1
\end{aligned}
$$

while $\xi$ is equal to the vector $X k$ that is generated at the kth stage by the system

$$
\begin{equation*}
x_{i+1}=A_{i} x_{i}+w_{i}, \quad \mathrm{i}=0, \ldots, k-1 . \tag{4.58}
\end{equation*}
$$

Thus, we have $\xi \mathrm{E} \mathcal{X}_{k}\left(z_{1}, \ldots, Z k\right)$ if and only if $V_{k}(\xi) \leq 1$, where $V_{k}(\xi)$ is the optimal cost of the problem of minimizing the quadratic cost

$$
\begin{aligned}
(x o-X O)^{\prime} S-l\left(x O-\hat{x}_{0}\right) & +\sum_{i=0}^{k-1} w_{i}^{\prime} M_{i}^{-1} w_{i} \\
& +\sum_{i=0}^{k-1}\left(z_{i+1}-C i+1 X i+1\right)^{\prime} N_{i+1}^{-1}(\mathrm{ZH} 1-C H I X H 1)
\end{aligned}
$$

subject to the system equation constraint (4.58) and the terminal condition $x k=\xi$ (here the $W_{i}$ are viewed as the controls/minimization variables). Thus

$$
\begin{equation*}
\mathcal{X}_{k}\left(z_{1}, \ldots, z_{k}\right)=\left\{\xi \mid V_{k}(\xi) \leq 1\right\} \tag{4.59}
\end{equation*}
$$

As the analysis of Section 4.1 suggests, the function $V_{k}$ is quadratic in $\xi$, and can be calculated by a DP recursion. This is because in the above problem the system is linear and the cost is quadratic. Thus the set of
possible states $\mathcal{X}_{k}\left(z_{1}, \ldots, Z k\right)$ of Eq. (4.59) is an ellipsoid. To calculate the matrix and center of this ellipsoid, we can use DP. Since here the terminal state $X k$ is specified to be equal to the given $\xi$, we should use a forward DP algorithm and view $V_{k}(\xi)$ as an optimal cost to arrive at $\xi$ by optimal choice of $X o$ and wo, $\ldots, W k-l$ in the system (4.58). Using the reasoning employed in Section 2.1, for $\mathbf{i}=1, \ldots, k$, we have the forward recursion

$$
\begin{aligned}
V_{i}\left(x_{i}\right)= & \min _{\substack{w_{i-1}, x_{i-1} \\
X i=A i-1 X i-1+\mathrm{w}_{\mathrm{i}-1}}}\left\{V_{i-1}\left(x_{i-1}\right)+w_{i-1}^{\prime} M_{i-1}^{-1} w_{i-1}\right. \\
& \left.\quad+(Z i-C i X i)^{\prime} N_{i}^{l}(Z i-C i X i)\right\} \\
= & \min _{x_{i-1}}\left\{V_{i-1}\left(x_{i-1}\right)+\left(X i-A_{i-1} x_{i-1}\right)^{\prime} M_{i-1}^{-1}(X i-A i-l X i-I)\right. \\
& \left.\quad+(Z i-C i X i)^{\prime} N_{i}^{-1}(Z i-C i X i)\right\}
\end{aligned}
$$

starting with the initial condition

$$
V o(x o)=(x o-X O)^{\prime} S-I\left(x O-\hat{x}_{0}\right) .
$$

At the kth step of the recursion, we obtain the set of possible states $\mathcal{X}_{k}\left(z_{1}, \ldots, Z k\right)$ ofEq. (4.59).

Rather than provide the detailed derivation, we leave it for the reader to verify by induction the formula

$$
V_{k}\left(x_{k}\right)=\left(x_{k}-\hat{x}_{k}\right)^{\prime \Sigma_{k}^{-1}}\left(x_{k}-\hat{x}_{k}\right)+\delta_{k}
$$

where $\hat{x}_{k}$ and $\Sigma_{k}$ are generated by the recursions

$$
\begin{gather*}
\hat{x}_{k}=A_{k-1} \hat{x}_{k-1}+\Sigma_{k} C_{k}^{\prime} N_{k}^{-1}\left(z_{k}-C_{k} A_{k-1} \hat{x}_{k-1}\right),  \tag{4.60}\\
\Sigma_{k}=\left(\hat{\Sigma}_{k}^{-1}+C_{k}^{\prime} N_{k}^{-1} C_{k}\right)^{-1},  \tag{4.61}\\
\hat{\Sigma}_{k}=A_{k-1} \Sigma_{k-1} A_{k-1}^{\prime}+M_{k-1}, \tag{4.62}
\end{gather*}
$$

with the initial condition

$$
\Sigma_{0}=S
$$

and $\delta_{k}$ is given by

$$
\begin{equation*}
\delta_{k}=\sum_{\mathrm{i}=1}^{k}\left(z_{i}-C i A i-I X i-I\right)^{\prime}\left(C_{i} \hat{\Sigma}_{i} C_{i}^{\prime}+\mathrm{Ni}\right)-1(Z i-C i A i-I X i-I) . \tag{4.63}
\end{equation*}
$$

There are several variations of the estimation problem discussed above for which we refer to the sources given at the end of the chapter.

## Instantaneous Constraints

We now consider a different type of set description of the uncertainty In particular, we assume that the initial state, the system disturbances, and the observation noise vectors are independently constrained to lie in ellipsoids. In other words, we know that

$$
\begin{gather*}
x_{0}^{\prime} S^{-1} x_{0} \leq 1,  \tag{4.64}\\
w_{i}^{\prime} M_{i}^{-1} w_{i} \leq 1, \quad \mathrm{i}=0, \quad, N-1,  \tag{4.65}\\
v_{i+1}^{\prime} N_{i+1}^{-1} \mathrm{Vi}_{\mathrm{i}+1} \leq 1, \quad \mathrm{i}=0, \quad, N-1, \tag{4.66}
\end{gather*}
$$

where $S, M i$, and $N i$ are given symmetric positive definite matrices. Thus the vector

$$
r=(x o, w o, \ldots, W N-I, V I, \ldots, V N)
$$

is known to belong to the set

$$
\begin{aligned}
& \mathcal{R}=\left\{r \mathrm{I}\left(x o-\hat{x}_{0}\right)^{\prime} S-I\left(X O-\hat{x}_{0}\right) \leq 1, w_{i}^{\prime} M_{i}^{-1} w_{i} \leq 1,\right. \\
& \left.v_{i+1}^{\prime} N_{i+1}^{-1} v_{i+1} \leq 1, \mathrm{i}=0, \ldots, N-\mathrm{I}\right\}
\end{aligned}
$$

For this case, the set of possible states $\mathcal{X}_{k}\left(\mathrm{Zl}^{\prime} \ldots, Z k\right)$ is not an ellipsoid, but can be bounded by an ellipsoid, by bounding the set $\mathcal{R}$ with an el lipsoid $\underline{\mathcal{R}}$, and by bounding $\mathcal{X}_{k}\left(z_{1}, \ldots, Z k\right)$ with the ellipsoid $\underline{\mathcal{X}_{k}}\left(z_{1}, \ldots, Z k\right)$ that corresponds to $\underline{\mathcal{R}}$ as in the preceding case of energy constraints.

In particular, we observe that if $X o, W o, \ldots, W N-I, V I, \ldots, V N$ satisfy the instantaneous constraints of Eqs. (4.64), (4.65), (4.66), then they also satisfy the energy constraint

$$
\begin{equation*}
\sigma\left(x_{0}-x O\right)^{\prime} S--I\left(x O-\hat{x}_{0}\right)+\sum_{i=0}^{\mathrm{N}-\mathbf{1}}\left(\mu_{i} w_{i}^{\prime} M_{i}^{-1} w_{i}+\nu_{i+1} v_{i+1}^{\prime} N_{i+1}^{-1} v_{i+1}\right) \leq 1 \tag{4.67}
\end{equation*}
$$

where $\sigma, \mu_{i}, \nu_{i+1}$ are any positive scalars satisfying

$$
\sigma+\sum_{\mathrm{i}=\mathrm{O}}^{\mathrm{N}-\mathbf{1}}\left(\mu_{i}+V i+\mathrm{I}\right)=1
$$

We thus replace the instantaneous constraints of Eqs. (4.64), (4.65), (4.66) with the energy constraint (4.67), and we obtain a bounding ellipsoid of the form

$$
\overline{\mathcal{X}}_{k}\left(z_{1}, \ldots, z_{k}\right)=\left\{x_{k} \mid\left(x_{k}-\hat{x}_{k}\right)^{\prime} \Sigma_{k}^{-1}\left(x_{k}-\hat{x}_{k}\right) \leq 1-\delta_{k}\right\},
$$

where $\hat{x}_{k}$ and $\Sigma_{k}$ are generated by the recursions given earlier for the energy constraint case, after we replace $S$ with $S / \sigma, M i$ with $M_{i} / \mu_{i}$, and Ni
with $N_{i} / \nu_{i}$. The formulas obtained in this way are simplified if we write $\sigma, \mu_{i}, \nu_{i+1}$ in the following form:

$$
\begin{gathered}
\sigma=\left(1 \quad \beta_{0}\right)\left(1-\gamma_{1}\right)\left(1-\beta_{1}\right)\left(1-\gamma_{2}\right) \ldots\left(1-\beta_{k-1}\right)\left(1-\gamma_{k}\right) \\
\mu_{0}=\beta_{0}\left(1 \quad \gamma_{1}\right)\left(1-\beta_{1}\right)\left(1-\gamma_{2}\right) \cdots\left(1-\beta_{k-1}\right)\left(1-\gamma_{k}\right) \\
V I=\gamma_{1}\left(1-\beta_{1}\right)\left(1-\gamma_{2}\right) \cdots\left(1-\beta_{k-1}\right)\left(1-\gamma_{k}\right) \\
\mu_{k-1}=\beta_{k-1}\left(1-\gamma_{k}\right) \\
\nu_{k}=\gamma_{k}
\end{gathered}
$$

where $\beta_{i-1}, \gamma_{i}, i=1, \ldots, k$ are any scalars with

$$
0<\beta_{i-1}<1, \quad O<\gamma_{i}<1
$$

It is easy to see that for the scalars $\left(J, \mu_{i}, \nu_{i+1}\right.$ given by the above equations, we have $\sigma+\sum_{i=0}^{N-1}\left(\mu_{i}+V i+l\right)=1$.

Now, by writing the estimator equations (4.60)-(4.63), with $S, N I_{i}$, and $N i$ replaced by $S /\left(s, M_{i} / \mu_{i}\right.$, and $N_{i} / V i$, respectively, we can obtain after straightforward manipulation a bounding ellipsoid of the form

$$
\left.\underline{\mathcal{X}} k(Z I, \ldots, Z k)=\left\{X k \left\lvert\, \begin{array}{ll}
X k & \hat{x}_{k}
\end{array}\right.\right)^{\prime} \Sigma_{k}^{-1}\left(X k-\hat{x}_{k}\right) \leq 1-\delta_{k}\right\}
$$

where

$$
\begin{gathered}
\hat{x}_{k}=A k-1 X k-l+\gamma_{k} \Sigma_{k} C_{k}^{\prime} N_{k}^{-1}\left(z_{k}-C_{k} A_{k-1} \hat{x}_{k-1}\right) \\
\Sigma_{k}=\left(\left(1 \quad \gamma_{k}\right) \hat{\Sigma}_{k}^{-1}+\gamma_{k} C_{k}^{\prime} N_{k}^{-1} C_{k}\right)-1 \\
\hat{\Sigma}_{k}=\left(1-\beta_{k-1}\right)^{-1} A_{k-1} \Sigma_{k-1} A_{k-1}^{\prime}+\beta_{k-1}^{-1} M_{k-1}
\end{gathered}
$$

with the initial condition

$$
\Sigma_{0}=S
$$

and $\delta_{k}$ is generated by the equation

$$
\begin{aligned}
\delta_{k}= & \left(1-\beta_{k-1}\right)\left(1-\gamma_{k}\right) \delta_{k-1}+(Z k \quad C k A k-l X k-l)^{\prime} \\
& \left(\left(1-\gamma_{k}\right)^{-1} C_{k} \hat{\Sigma}_{i} C_{k}^{\prime}+\gamma_{k}^{-1} N_{k}\right)-1(Z k-C k A k-l X k-l)
\end{aligned}
$$

with the initial condition

$$
\delta_{0}=0
$$

We omit the verification of the above equations because it is tedious, and we refer to the cited references. Note that the estimators for both cases of energy and instantaneous constraints bear close resemblance to the Kalman filtering algorithm described in Appendix E. An interesting problem variant is when the system equation has the form $X k+l=X k$. In this case, the problem is to use linear measurements to estimate the initial state $X o$, which can be viewed as an unknown parameter vector. It can then be shown under mild assumptions that $\Sigma_{k} \rightarrow 0$ as $k \rightarrow 00$, so that the parameter vector is identified with arbitrary accuracy as the number of measurements increases.

### 4.6.2 Control with Unknown-but-Bounded $\mathbb{D}$ istarbances

We now consider a problem of control when the uncertain quantities are described by their membership in given sets. We consider the system

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}, w_{k}\right)
$$

where as usual $X k$ is the state, $U k$ is the control to be selected from a set $U k(X k)$, and $W k$ is a disturbance. However, instead of probability distributions, we only know that $W k$ belongs to a given set $W_{k}\left(x_{k}, U k\right)$, which may depend on the current state $X k$ and control $U k$.

Often in control problems one is interested in keeping the state of the system close to a desired trajectory, in spite of the effects of the disturbances. We can formulate such a problem as one of finding a policy $\pi=\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}$ with $\mu_{k}\left(x_{k}\right) \mathrm{E} U k(X k)$ for all $X k$ and $k$, such that for each $k=1,2, \ldots, N$, the state $X k$ of the closed-loop system

$$
x_{k+1}=f_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)
$$

belongs to a given set $X k$, called the target set at time $k$.
We may view the set sequence $\{X l, X 2, \ldots, X N\}$ as a "tube" within which the state must stay, even under the worst possible choice of the disturbances $W k$ from within the corresponding sets $W_{k}\left(X k^{\prime} \mu_{k}\left(x_{k}\right)\right)$. Accordingly we refer to this problem as the problem of reachability of a target tube.

One may formulate this problem as a minimax control problem (cf. Section 1.6), where the cost at stage $k$ is

$$
g k(X k)= \begin{cases}0 & \text { if } X k \in X k, \\ 1 & : \uparrow \\ 1 & X k \notin X \\ k\end{cases}
$$

With this choice, the optimal cost-to-go from a given initial state $x Q$ is the minimum number of violations of the target tube constraints $X k$ E $X k$ that can occur when the $W k$ are optimally chosen, subject to the constraint $W k$ E $W k(X k, U k)$, by an adversary wishing to maximize the number of violations. In particular, if $J k(X k)=0$ for some $X k E X k$, there exists a policy such that starting from $X k$, the subsequent system states $X i, i=k+1, \ldots$, are guaranteed to be within the corresponding sets $X i$.

It can be seen that the set

$$
\bar{X}_{k}=\left\{x_{k} \mid J_{k}\left(x_{k}\right)=0\right\}
$$

is the set that we must reach at time $k$ in order to be able to maintain the state within the subsequent target sets. Accordingly, we refer to $\underline{X} k$ as the effective target set at time $k$. We can generate the sets $\underline{X} k$ with a backwards recursion, which is derived from the DP algorithm for minimax problems
(see Section 1.6) but can also be easily justified from first principles. In particular, we start with

$$
\begin{equation*}
X N=X N \tag{4.68}
\end{equation*}
$$

and for $k=0,1, \ldots, N-1$, we have
$\underline{X_{k}}={ }_{\{X k} \mathrm{E} X_{k} \mid$ there exists $U k \mathrm{E} U k(X k)$ such that

$$
\begin{equation*}
f k(X k, U k, W k) \mathrm{E} \underline{X_{k}+1}, \text { for all } W k \mathrm{E} W k(X k, U k) \jmath . \tag{4.69}
\end{equation*}
$$

## Example 4.6.1

Consider the scalar linear system

$$
x_{k+1}=2 x_{k}+u_{k}+w_{k}
$$

and the target tube $\left\{X l, X_{2}, \ldots \cdot X_{\dot{N}}\right\}$, where for all $k$,

$$
X_{k}=[-1,1]
$$

We want to keep the state within this tube by using controls $U k$ that belong to the set $U k=[-1,1]$, and in spite of the effects of the disturbances $W k$ that can take any values in the set $[-1 / 2,1 / 2]$.

Let us construct the effective target sets $\underline{X} k$ by using the DP recursion of Eqs. (4.68) and (4.69). We have $\underline{X N}=[-1,1]$ ' and

$$
\begin{aligned}
& \underline{X_{N-1}} \quad\{x \text { I for some } U \mathrm{E} \quad \text { 1] we have } \\
& \quad-1 \leq 2 x+u+w \leq 1 \text { for all } w \mathrm{E}[-1 / 2,1 / 2 \mathrm{~J}\}
\end{aligned}
$$

We see that for $x$ and $u$ to satisfy $-1 \leq 2 x+u+w \leq 1$ for all $w \mathrm{E}[-1 / 2,1 / 2]^{\prime}$ it is necessary and sufficient that

$$
-\frac{1}{2} \leq 2 x+u \leq \frac{1}{2}
$$

so $u$ must be chosen (with knowledge of $x$ ) to satisfy

$$
-1 \leq u \leq 1, \quad-\frac{1}{2}-2 x \leq u \leq \frac{1}{2}-2 x
$$

For existence of such a $u$, the intervals $[-1,1]$ and $[-1 / 2-2 x, 1 / 2-2 x]$ must have nonempty intersection, which can be seen to be true if and only if

$$
3 \leq x \leq \frac{3}{4}
$$

Thus, to be able to reach the set $X_{N}$ at time $N$, the state $X_{N_{-1}}$ must belong to the (effective target) set

$$
\bar{X}_{N-1}=\left[-\frac{3}{4}, \frac{3}{4}\right]
$$

We similarly proceed to construct We have

$$
\begin{aligned}
& \underline{x N-2=\{x \text { I for some } u \mathrm{E}[-1,1] \text { we have }} \begin{aligned}
& -3 / 4 \leq 2 x+u+w \leq 3 / 4 \text { for all } w E[-1 / 2,1 / 2 \mathrm{~J}\}
\end{aligned}
\end{aligned}
$$

and we see that $u$ must be chosen so that

$$
-1 \leq u \leq 1, \quad-\frac{1}{4}-2 x \leq u \leq \frac{1}{4}-2 x
$$

For existence of such a $u$, the intervals $[-1,1]$ and $[-1 / 4-2 x, 1 / 4-2 x]$ must have nonempty intersection, which can be seen to be true if and only if

$$
-\frac{5}{8} \leq x \leq \frac{5}{8}
$$

Thus, to be able to reach the effective target set $\underline{X} N-1$ at time $N-1$, the state $X N-2$ must belong to the set

$$
\bar{X}_{N-2}=\left[\begin{array}{rr}
5 & 5 \\
-8 & 8 \\
8
\end{array}\right]
$$

The above calculations illustrate the form of the algorithm that yields the effective target set $\underline{X} k$ for every $k$. We have

$$
\bar{X}_{k}=\left[-\alpha_{k}, \alpha_{k}\right]
$$

where the scalars $\alpha_{k}$ satisfy the recursion

$$
\alpha_{k}=\frac{\alpha_{k+1}}{2}+\frac{1}{4}, \quad k=O, 1, \ldots, N-1
$$

with the starting condition

$$
\alpha_{N}=1 .
$$

In order to guarantee reachability of the given target tube, the initial state $X o$ should belong to the interval $\left[-\alpha_{0}, \alpha_{0}\right]$. Note that the scalars $\alpha_{k}$ are monotonically decreasing as $k \rightarrow-\infty$, and we have $\alpha_{k} \rightarrow 1 / 2$. Thus, if the initial state $X o$ is in the interval $[-1 / 2,1 / 2]^{\prime}$ then given any horizon length $N$, there is a policy that keeps the state of the system within the set $[-1,1]$. In fact, it can be seen that the linear stationary policy $\{\mu, \mu, \ldots$.$\} , where$

$$
\mu(x)=-2 x
$$

keeps the state of the system in the interval $[-1 / 2,1 / 2]^{\prime}$ provides the initial state belongs to that intervaL It can also be seen that if the initial state does not belong to the interval $[-1 / 2,1 / 2]^{\prime}$ then there is a large enough horizon length $N$, such that for every admissible policy, a sequence of feasible disturbances exists that will force the state to be outside the target set $[-1,1]$ at some time $k \leq N$. These observations can be generalized for the case oflinear
systems and ellipsoidal constraint sets (see the discussion and the references given below).

In general, it is not easy to characterize the effective target sets $\underline{X} k$. However, similar to the estimation problem of the preceding subsection, a few special cases involving the linear system

$$
x_{k+1}=A_{k} x_{k}+B_{k} u_{k}+w_{k}, \quad k=0,1, \ldots, N-1,
$$

where $A_{k}$ and $B_{k}$ are given matrices, are amenable to exact or approximate computational solution. One such case is when the sets $X_{k}$ are ellipsoids, and the sets $U k(X k)$ and $W k(X k, U k)$ are also ellipsoids that do not depend on $X k$ and ( $X k^{\prime} U k$ ), respectively. In this case, the effective target sets $\underline{X} k$ are not ellipsoids, but can be approximated by inner ellipsoids $X_{k}$ C $X k$ (this requires that the ellipsoids $U k$ have sufficiently large size, for otherwise the target tube may not be reachable and the problem may not have a solution). Furthermore, the state trajectory $\{X l, X 2, \ldots, X N\}$ can be maintained within the ellipsoidal tube

$$
\left\{\tilde{X}_{1}, \tilde{X}_{2}, \ldots, \tilde{X}_{N}\right\}
$$

by using a linear control law (compare with the preceding example). We outline the main algorithm in Exercise 4.31, and refer to the author's thesis work, [Ber71], [BeR71b], for a more detailed analysis.

Another case of interest is when the sets $X_{k}$ are polyhedral, and the sets $U k(X k)$ and $W_{k}\left(x_{k}, U k\right)$ are also polyhedral, and independent of $X k$ and ' $U k$ ' Then the effective target sets are polyhedral and can be computed by linear programming methods.

An important special case is when the problem is stationary and $f k^{\prime}$ $X k, U k$, and $W k$ do not depend on $k$. Then it can be shown that the effective target sets satisfy

$$
\bar{X}_{k} \subset \bar{X}_{k+1}, \quad \text { for all } k
$$

The intersection $\square_{t=0}^{N} \bar{X} k$ depends on the size of the horizon $N$ and "decreases" to the set

$$
X_{\infty}=\cap_{N=1}^{\infty} \cap_{k=0}^{N} \underline{X k}
$$

as $N$ increases to oo. We may view $X_{\infty}$ as the set within which the state can be kept for an arbitrarily large (but finite) number of time periods. Paradoxically, under some unusual circumstances, there may be states in $X_{\infty}$ starting from which it may be impossible to remain within $X_{\infty}$ for an indefinitely long horizon, i.e., an infinite number of time periods. Conditions that preclude this possibility have been investigated by the author in [Ber72a]. References [Ber71] and [Ber72a] contain a methodology for constructing ellipsoidal inner approximations to $X_{\infty}$ and an associated linear control law for the case where the system is linear, and the sets $X k, U k$, and $W_{k}$ are ellipsoids.

### 4.7 NOTES, SOURCES, AND EXERCISES

The certainty equivalence principle for dynamic linear-quadratic problems was first discussed by Simon [Sim56]. His work was preceded by Theil [The54], who considered a single-period case, and Holt, Modigliani, and Simon [HMS55], who considered a deterministic case. Similar problems were independently considered by Kalman and Koepcke [KaK58], Joseph and Tou [JoT61], and Gunckel and Franklin [GuF63]. The linear-quadratic problem is central in control theory; see the special issue [IEE71], which contains hundreds of references.

The literature on inventory control stimulated by the pioneering paper of Arrow et al. [AHM51] is also voluminous. An important work summarizing most of the research up to 1958 is Arrow, Karlin, and Scarf [AKS58]. Veinott [Vei66] also surveys the early work on the subject. The proof of optimality of $(8, S)$ policies in the case of nonzero fixed costs is due to Scarf [Sca60].

Most of the material in Section 4.3 is taken from Mossin [Mos68]; see also Hakansson [Hak70], [Hak71], and Samuelson [Sam69]. Many applications of DP in economics are described in Sargent [Sar87], and Stokey and Lucas [StL89].

The material of Section 4.4 is largely drawn from White [Whi69]. Example 4.5.1 is given by Ross [Ros70], Example 4.5 .2 is given by Ross [Ros83], and Example 4.5.3 is due to Weiss and Pinedo [WeP80]. An extensive reference on scheduling is Pinedo [Pin95].

The problem of state estimation with a set-membership description of the uncertainty was first formulated and addressed by vVitsenhausen in his Ph.D. work at MIT [Wit66], and also in the paper [Wit68]. The material given here follows closely the author's Ph.D. thesis [Ber71], where the problem of estimation for the case of an energy constraint was first formulated and solved. The estimator given in Section 4.6.1 for the case of instantaneous constraints was also first derived in the author's thesis using the method given here, and a steady-state analysis was also given. A state estimator using ellipsoidal approximations for the case of instantaneous constraints was first proposed by Schweppe [Sch68], [Sch74]. This estimator, however, has several drawbacks relative to the estimator given here. In particular, the associated matrix $\Sigma_{k}$ depends on the observations $z, \ldots, z_{k}$, and need not converge to a steady state as $k \rightarrow 00$.

Continuous-time versions of the estimators of Section 4.6, as well and other variants of the estimation problem (the prediction and smoothing problems) with a set-membership description of the uncertainty were first given by Bertsekas and Rhodes [BeR71a]. Kurzhanski and Valyi [KuV97] provide an account of set-membership estimation. Deller [DeI89] surveys applications in signal processing. Kosut, Lau, and Boyd [KLB92] discuss applications in system identification. The state estimation problem can also be addressed as a minimax problem that involves in part a probabilis-
tic description of the uncertainty. Basar [Bas91] describes the relations between this approach and the set-membership approach.

The target tube reachability problem was first formulated by the author in his Ph.D. thesis [Ber71]; see also the papers [BeR71b] and [Ber72a], and Exercises 3.23 and 3.24 of Vol. II. The associated recursion of Eqs. (4.68)-(4.69) for the effective target sets, and methods for approximating this recursion were also given in these references (see Exercise 4.31). The target tube reachability problem arises within several contexts in control system design, including model predictive control, which is described in Section 6.5 (for a recent discussion, see the paper by Mayne [MayOl)). For a survey of the associated issues, including extensions to continuous-time systems and additional references, see BIanchini [Blag9].

There has been considerable recent research on minimax formulations of general optimization problems under uncertainty, such as linear programming problems. This approach, which is known as robust optimization, is also based on a set-membership description of the uncertainty; for some representative works, see Ben-Tal and Nemirovski [BeN98], [BeNOl], and Bertsimas and Sim [BeS03].

## E XEReISES

## 4.1 (Linear-Quadratic Problems with Forecasts)

Consider the linear-quadratic problem first examined in Section 4.1 ( $A k^{\prime} B k$ : known) for the case where at the beginning of period $k$ we have a forecast $\gamma_{k} \mathrm{E}\{\mathrm{I}, 2, \ldots, n\}$ consisting of an accurate prediction that $w k$ will be selected in accordance with a particular probability distribution $P_{k / Y k}$ (cf. Section 1.4). The vectors $w k$ need not have zero mean under the distribution $P_{k / Y k}$. Show that the optimal control law is of the form

$$
\mu_{k}\left(x_{k}, y_{k}\right)=-\left(B_{k}^{\prime} K_{k+1} B_{k}+R_{k}\right)^{-1} B_{k}^{\prime} K_{k+1}\left(A_{k} x_{k}+E\left\{w_{k} \mid y_{k}\right\}\right)+\alpha_{k}
$$

where the matrices $K k$ are given by the Riccati equation (4.3) and (4.4) and $\alpha_{k}$ are appropriate vectors.

Consider a scalar linear system

$$
x_{k+1}=a_{k} x_{k}+b_{k} u_{k}+w_{k}, \quad k=0,1, \ldots, N-1
$$

where $a k, b k E R$, and each $w k$ is a Gaussian random variable with zero mean and variance $\sigma^{2}$. We assume no control constraints and independent disturbances. Show that the control law $\left\{\mu_{0}^{*}, \mu_{1}^{*}, \ldots, \mu_{N}^{*}-\mathrm{I}\right\}$ that minimizes the cost function

$$
E\left\{\exp \left[x_{N}^{2}+\sum_{k=0}^{N-1}\left(x_{k}^{2}+r u_{k}^{2}\right)\right]\right\}, \quad r>0
$$

is linear in the state variable, assuming that the optimal cost is finite for every xo. Show by example that the Gaussian assumption is essential for the result to hold. (For analyses of multidimensional versions of this exercise, see Jacobson [Jac73], Whittle [Whi82], [Whi90], and Basar [BasOO].)

## 4.3

Consider an inventory problem similar to the problem of Section 4.2 (zero fixed cost). The only difference is that at the beginning of each period $k$ the decision maker, in addition to knowing the current inventory level $X k$, receives an accurate forecast that the demand $w k$ will be selected in accordance with one out of two possible probability distributions $P_{\ell}, P_{s}$ (large demand, small demand). The a priori probability of a large demand forecast is known (d. Section 1.4).
(a) Obtain the optimal ordering policy for the case of a single-period problem.
(b) Extend the result to the N -period case.
(c) Extend the result to the case of any finite number of possible distributions.

## 4.4

Consider the inventory problem of Section 4.2 (zero fixed cost), where the purchase costs $c k, k=0,1, \ldots, N-1$, are not initially known, but instead they are independent random variables with a priori known probability distributions. The exact value of the cost $c k$, however, becomes known to the decision maker at the beginning of the kth period, so that the inventory purchasing decision at time $k$ is made with exact knowledge of the cost $c k$. Characterize the optimal ordering policy assuming that $p$ is greater than all possible values of $c k$.

## 4.5

Consider the inventory problem of Section 4.2 for the case where the cost has the general form

$$
E\left\{\sum_{k=0}^{N} r_{k}\left(x_{k}\right)\right\}
$$

The functions $r k$ are convex and differentiable and

$$
\lim _{x \rightarrow-\infty} \frac{d r k(X)}{d x}=- \text { oo, } \quad \lim _{x \rightarrow \infty} \frac{d r_{k}(x)}{d x}=\infty, \quad k=O, \ldots, N
$$

(a) Assume that the fixed cost is zero. Write the DP algorithm for this problem and show that the optimal ordering policy has the same form as the one derived in Section 4.2.
(b) Suppose there is a one-period time lag between the order and the delivery of inventory; that is, the system equation is of the form

$$
x_{k+1}=x_{k}+u_{k-1}-w_{k}, \quad k=0,1, \ldots, N-1,
$$

where $U-I$ is given. Reformulate the problem so that it has the form of the problem of part (a). Hint: Make a change of variables $Y k=X k+U k-I$.

## 4.6 (Inventory Control for Nonzero Fixed Cost)

Consider the inventory problem of Section 4.2 (nonzero fixed cost) under the assumption that unfilled demand at each stage is not backlogged but rather is lost; that is, the system equation is $X k+1=\max (\mathrm{O}, X k+U k-W k)$ instead of $X k+1=X k+U k-W k$. Complete the details of the following argument, which shows that a multiperiod ( $s, S$ ) policy is optimal.
Abbreviated Proof: (due to S. Shreve) Let $I N(X)=0$ and for all $k$

$$
\begin{gathered}
G_{k}(y)=c y+E\left\{h \max \left(0, y-w_{k}\right)+p \max \left(0, w_{k}-y\right)+J_{k+1}\left(\max \left(0, y-w_{k}\right)\right)\right\} \\
J_{k}(x)=-c x+\min _{u \geq 0}\left[K \delta(u)+G_{k}(x+u)\right]
\end{gathered}
$$

where $\delta(0)=0, \delta(u)=1$ for $u>0$. The result will follow if we can show that $G k$ is K-convex, continuous, and $G k(Y) \rightarrow 00$ as jyl $\rightarrow 00$. The difficult part is to show K-convexity, since K-convexity of $G k-I$ does not imply K-convexity of $E\{J k+l(\max (O, y-w))\}$. It will be sufficient to show that K-convexity of $G_{k+1}$ implies K-convexity of

$$
\begin{equation*}
H(y)=\operatorname{pmax}(O,-y)+J k+1(\max (\mathrm{O}, y)), \tag{4.70}
\end{equation*}
$$

or equivalently that

$$
\left.\mathrm{K}+H(y+z) \geq H(y)+z\left(\underline{H(y)}-\frac{H(y}{\hbar}=\underline{b}\right)\right), \quad z \geq 0, b>0, Y \in \Re .
$$

By K-convexity of $G k+1$ we have for appropriate scalars $S k+l$ and $S k+1$ such that $G k+I(S k+l)=\min _{y} G_{k+1}(y)$ and $\mathrm{K}+G k+l(S k+l)=G k+l(S k+I)$ :

$$
J_{k+1}(x)= \begin{cases}K+G_{k+1}\left(S_{k+1}\right)-c x & \text { if } x<S k+l,  \tag{4.72}\\ G_{k+1}(x)-c x & \text { if } x \geq S k+l,\end{cases}
$$

and $J k+1$ is K-convex by the theory of Section 4.2.
Case 1: $0 \leq y-b<y \leq y+z$. In this region, Eq. (4.71) follows from K-convexity of $J k+l$.

Sec. 4.7 Notes, numued, and Exercises
Case 2: $y-b<y \leq y+z \leq 0$. In this region, $H$ is linear and hence K-convex.

Case 3: $y-b<y \leq 0 \leq y+z$. In this region, Eq. (4.71) may be written [in view of Eq. (6.1)] as

$$
K+J_{k+1}(y+z) \geq J_{k+1}(0)-p(y+z)
$$

We will show that

$$
\begin{equation*}
K+J_{k+1}(z) \geq J_{k+1}(0)=p z, \quad z \geq 0 . \tag{4.73}
\end{equation*}
$$

If $0<s_{k+1} \leq z$, then using Eq. (4.72) and the fact $p>\mathrm{c}$, we have

$$
K+J_{k+1}(z)=K-c z+G_{k+1}(z) \geq K-p z+G_{k+1}\left(S_{k+1}\right)=J_{k+1}(0)=p z
$$

If $0 \leq z \leq S k+l$, then using Eq. (4.72) and the fact $p>\mathrm{c}$, we have

$$
K+J_{k+1}(z)=2 K-c z+G_{k+1}\left(S_{k+1}\right) \geq K-p z+G_{k+1}\left(S_{k+1}\right)=J_{k+1}(0)-p z .
$$

If $S k+1 \leq 0 \leq z$, then using Eq. (4.72), the fact $p>\mathrm{c}$, and part (iv) of the lemma in Section 4.2, we have

$$
K+J_{k+1}(z)=K-c z+G_{k+1}(z) \geq G_{k+1}(0)-p z=J_{k+1}(0)-p z .
$$

Thus Eq. (4.73) is proved and Eq. (4.71) follows for the case under consideration.
Case 4: $y-b<0<y \leq y+z$. Then $0<Y<b$. If

$$
\begin{equation*}
\left.\underline{H(y)}=\underline{H(O)}>\underline{H(y)}=\frac{H(y}{b}=\underline{b}\right), \tag{4.74}
\end{equation*}
$$

then since $H$ agrees with $J_{k+1}$ on $[0,00)$ and $J_{k+1}$ is K-convex,

$$
\begin{aligned}
K+H(y+z) & \geq H(y)+z \underline{(H(Y)} \frac{-}{\bar{y}} \underline{H(O))} \\
& \left.\geq H(y)+z\left(\underline{H(y)}-\frac{H(y}{\bar{b}}-\underline{b}\right)\right)
\end{aligned}
$$

where the last step follows from Eq. (4.74). If

$$
\underline{H(y)}=\underline{y}=\underline{H(O)}<\underline{H(y)}=\frac{H(y}{b}=\underline{b},
$$

then we have

$$
H(y)-H(O)<\frac{\boldsymbol{y}}{\boldsymbol{b}}(H(y)-H(y-b))=\frac{\boldsymbol{y}}{\boldsymbol{b}}(H(y)-H(O)+p(y-b))
$$

It follows that

$$
\left(1-\frac{y}{b}\right)(H(y)-H(0))<\left(\frac{y}{b}\right) p(y-b)=-p y\left(1-\frac{y}{b}\right)
$$

and since $b>y$,

$$
\begin{equation*}
H(y)-H(O)<-p y \tag{4.75}
\end{equation*}
$$

Now we have, using the definition of $H$, Eqs. (4.73) and (4.75),

$$
\begin{aligned}
H(y)+\underline{z H(Y)}-\frac{H(y}{\bar{b}}-\underline{b)} & =H(y)+z \underline{(H(O)}=\frac{p y}{-\frac{H(0)}{b} \pm \underline{p r y}}-\underline{b »)} \\
& =H(y)-p z \\
& <H(O)-p(y+z) \\
& \leq K+H(y+z)
\end{aligned}
$$

Hence Eq. (4.73) is proved for this case as well. Q.E.D.

## 4.7

Consider the inventory problem of Section 4.2 (zero fixed cost) with the difference that successive demands are correlated and satisfy a relation of the form

$$
w_{k}=e_{k}-\gamma e_{k-1}, \quad k=0,1, \ldots
$$

where $\gamma$ is a given scalar, ek are independent random variables, and e-1 is given.
(a) Show that this problem can be converted into an inventory problem with independent demands. Hint: Given wo, wi, ..., $W k-I$, we can determine ek-l in view of the relation

$$
e k-I=\gamma^{k} \mathrm{e}-1+\sum_{\mathrm{i}=\mathrm{O}}^{\mathrm{k}-\mathrm{I}} \gamma^{\mathrm{i}} w_{k-I-i}
$$

Define $z k=x k+\gamma e_{k-1}$ as a new state variable.
(b) Show that the same is true when in addition there is a one-period delay in the delivery of inventory (cf. Exercise 4.5).
4.8

Consider the inventory problem of Section 4.2 (zero fixed cost), the only difference being that there is an upper bound $\bar{b}$ and a lower bound $\underline{b}$ to the allowable values of the stock $X k$. This imposes the additional constraint on $U k$

$$
\underline{b}+d \leq u_{k}+x_{k} \leq \bar{b}_{1}
$$

where $d_{-}>$is the maximum value that the demand $w k$ can take (we assume $\underline{b}+d<\bar{b})$. Show that an optimal policy $\left\{\mu_{0}^{*}, \ldots, \mu_{N}^{*}-\mathrm{I}\right\}$ is of the form

$$
\mu_{k}^{*}\left(x_{k}\right)= \begin{cases}S_{k} & \text { if } x_{k}<S_{k} \\ 0 & \text { if } x_{k} \geq S_{k}\end{cases}
$$

where $S_{0}, S_{1}, \ldots, S_{N-1}$ are some scalars

## 4.9

Consider the inventory problem of Section 4.2 (nonzero fixed cost) with the difference that demand is deterministic and must be met at each time period (i.e., the shortage cost per unit is 00 ). Show that it is optimal to order a positive amount at period $k$ if and only if the stock $x_{k}$ is insufficient to meet the demand $w k$. Furthermore, when a positive amount is ordered, it should bring up stock to a level that will satisfy demand for an integral number of periods.

### 4.10 [VeW5], [Tsi84b]

Consider the inventory control problem of Section 4.2 (zero fixed cost) with the only difference that the orders $U k$ are constrained to be nonnegative integers. Let $J_{k}$ be the optimal cost-to-go function. Show that:
(a) $J k$ is continuous.
(b) $J k(X+1)-J k(X)$ is a nondecreasing function of $x$.
(c) There exists a sequence $\{S k\}$ of numbers such that the policy given by

$$
\mu_{k}\left(x_{k}\right)= \begin{cases}n & \text { if } s_{k}-n \leq x k<s_{k} \quad n+1, \quad n=1,2, \ldots \\ 0 & \text { if } x k \geq s_{k}\end{cases}
$$

is optimal.

### 4.11 (Capacity Expansion Problem)

Consider a problem of expanding over $N$ time periods the capacity of a production facility. Let us denote by $x k$ the production capacity at the beginning of the kth period and by $u k \geq$ the addition to capacity during the kth period. Thus capacity evolves according to

$$
x_{k+1}=x_{k}+u_{k}, \quad k=O, l, \ldots, N-1
$$

The demand at the kth period is denoted $w k$ and has a known probability distribution that does not depend on either $x k$ or $U k$. Also, successive demands are assumed to be independent and bounded. We denote:
$C k(U k)$ : expansion cost associated with adding capacity $U k$,
$P k(X k+U k-W k):$ penalty associated with capacity $X k+U k$ and demand Wk,
$S(X N)$ : salvage value of final capacity $X N$.
Thus the cost function has the form

$$
\underset{\substack{w_{k} \\ k=0,1, \ldots, N-1}}{E}\left\{-S\left(x_{N}\right)+\sum_{k=0}^{N-1}\left(C_{k}\left(u_{k}\right)+P_{k}\left(x_{k}+u_{k}-w_{k}\right)\right)\right\}
$$

(a) Derive the DP algorithm for this problem
(b) Assume that $S$ is a concave function with $\lim _{\mathrm{x}-+\mathrm{oo}} d S(x) j d x=0, P k$ are convex functions, and the expansion cost $C_{k}$ is of the form

$$
C k(U)=\{K+C k U \quad \text { if } u>0
$$

## 0

$$
\hat{\mathrm{I}}_{U}=\hat{0}
$$

where $K \geq 0, C k>$ for all $k$. Show that the optimal policy is of the $(8, S)$ type assuming $C k Y+E\{P k(Y-W k)\} \rightarrow 00$ as $|y| \rightarrow 00$.

### 4.12

We want to use a machine to produce a certain item in quantities that meet as closely as possible a known (nonrandom) sequence of demands $d k$ over $N$ periods. The machine can be in one of two states: good $(G)$ or bad $(B)$. The state of the machine is perfectly observed and evolves from one period to the next according to

$$
P(G \mid G)=\lambda_{G}, \quad P(B \mid G)=1-\lambda_{G}, \quad P(B \mid B)=\lambda_{B}, \quad P(G \mid B)=1-\lambda_{B}
$$ where $\lambda_{G}$ and $\lambda_{B}$ are given probabilities. Let $X k$ be the stock at the beginning of the $k$ th period. If the machine is in good state at period $k$, it can produce $U k$, where $U k \mathrm{E}[0, \bar{u}]$, and the stock evolves according to

$$
X k+l=X k+U k \quad d k
$$

otherwise the stock evolves according to

$$
X k+l=X k-d k
$$

There is a cost $g(X k)$ for having stock $X k$ in period $k$, and the terminal cost is also $g(X N)$. We assume that the cost per stage 9 is a convex function such that $g(x) \rightarrow$ oo as |Xi $\rightarrow 00$. The objective is to find a production policy that minimizes the total expected cost.
(a) Prove inductively a convexity property of the cost-to-go functions, and show that for each $k$ there is a target stock level $S k+l$ such that if the machine is in the good state, it is optimal to produce $u_{k}^{*} \mathrm{E}[0, \bar{u}]$ that will bring $x_{k+1}$ as close as possible to $S k+l$.
(b) Generalize part (a) for the case where each demand $d k$ is random and takes values in an interval $[0, \bar{d}]$ with given probability distribution. The stock and the state of the machine are still perfectly observable.

### 4.13 (A Gambling Problem)

A gambler enters a game whereby he may at time $k$ stake any amount $U k \geq$ that does not exceed his current fortune $X k$ (defined to be his initial capital plus his gain or minus his loss thus far). He wins his stake back and as much more with probability $p$, where $\frac{1}{2}<p<1$, and he loses his stake with probability ( $1-p$ ). Show that the gambling strategy that maximizes $E\{\ln x N\}$, where $X N$ denotes his fortune after $N$ plays, is to stake at each time $k$ an amount $U k=(2 p-I) x k$. Hint: The problem is related to the portfolio problem of Section 4.3.

### 4.14

Consider the dynamic portfolio problem of Section 4.3 for the case where at each period $k$ there is a forecast that the rates of return of the risky assets for that period will be selected in accordance with a particular probability distribution as in Section 1.4. Show that a partially myopic policy is optimaL

### 4.15

Consider a problem involving the linear system

$$
x_{k+1}=A_{k} x_{k}+B_{k} u_{k}, \quad k=0,1, \ldots, N-1,
$$

where the $n \times n$ matrices $A k$ are given, and the $n \times \mathrm{m}$ matrices $B k$ are random and independent with given probability distributions that do not depend on $X k$, $U k$. The problem is to find a policy that maximizes $E\left\{U\left(C^{\prime} X N\right)\right\}$, where c is a given n-dimensional vector. We assume that $U$ is a concave twice continuously differentiable utility function satisfying for all $Y$

$$
\begin{aligned}
& U^{\prime}(y) \\
- & \underline{U^{\prime \prime}(y)}
\end{aligned}=a+b y,
$$

and that the control is unconstrained. Show that the optimal policy consists of linear functions of the current state. Hint: Reduce the problem to a onedimensional problem and use the results of Section 4.3.

### 4.16

Suppose that a person wants to sell a house and an offer comes at the beginning of each day. We assume that successive offers are independent and an offer is $W_{j}$ with probability $P_{j}, j=1, \ldots, n$, where $W_{j}$ are given nonnegative scalars. Any offer not immediately accepted is not lost but may be accepted at any later date. Also, a maintenance cost c is incurred for each day that the house remains unsold. The objective is to maximize the price at which the house is sold minus the maintenance costs. Consider the problem when there is a deadline to sell the house within $N$ days and characterize the optimal policy.

### 4.17

Assume that we have $X_{o}$ items of a certain type that we want to sell over a period of $N$ days. At each day we may sell at most one item. At the kth day, knowing the current number $X k$ of remaining unsold items, we can set the selling price $U k$ of a unit item to a nonnegative number of our choice; then, the probability $\lambda_{k}\left(u_{k}\right)$ of selling an item on the kth day depends on $U k$ as follows:

$$
\lambda_{k}\left(u_{k}\right)=\alpha e^{-u_{k}}
$$

where $\alpha$ is a given scalar with ${ }^{\mathrm{O}}<\alpha \leq 1$. The objective is to find the optimal price setting policy so as to maximize the total expected revenue over $N$ days.
(a) Assuming that, for all $k$, the cost-to-go function $J k(X k)$ is monotonically nondecreasing as a function of $X k$, prove that for $X k>0$, the optimal prices have the form

$$
\mu_{k}^{*}\left(x_{k}\right)=1+J_{k+1}\left(x_{k}\right)-J_{k+1}\left(x_{k}-1\right)
$$

and that

$$
J_{k}\left(x_{k}\right)=\alpha e^{-\mu_{k}^{*}\left(x_{k}\right)}+J_{k+1}\left(x_{k}\right)
$$

(b) Prove simultaneously by induction that, for all $k$, the cost-to-go function $J k(X k)$ is indeed monotonically nondecreasing as a function of $X k$, that the optimal price $\mu_{k}^{*}\left(x_{k}\right)$ is monotonically nonincreasing as a function of $X k$, and that $J k(X k)$ is given in closed form by

$$
J_{k}\left(x_{k}\right)= \begin{cases}(N-k) \alpha e^{-1} & \text { if } X k \geq N-k \\ \sum_{i=k}^{N-x_{k}} \alpha e^{-\mu_{i}^{*}\left(x_{k}\right)}+x_{k} \alpha e^{-1} & \text { if }<x_{k}<N-k, \\ 0 & \text { if } X k=0 .\end{cases}
$$

### 4.18 (Optimal Termination of Sampling) (www)

This is a classical problem, which when appropriately paraphrased, is known as the job selection, or as the secretary selection, or as the spouse selection problem. A collection of $N \geq 2$ objects is observed randomly and sequentially one at a time. The observer may either select the current object observed, in which case the selection process is terminated, or reject the object and proceed to observe the next. The observer can rank each object relative to those already observed, and the objective is to maximize the probability of selecting the "best" object according to some criterion. It is assumed that no two objects can be judged to be equal. Let $r^{*}$ be the smallest positive integer $r$ such that

$$
\frac{1}{-1+} \quad \frac{1}{-2}+\ldots+{ }_{r}^{1} \leq 1
$$

Show that an optimal policy requires that the first $r^{*}$ objects be observed. If the $r^{*}$ th object has rank 1 relative to the others already observed, it should be selected; otherwise, the observation process should be continued until an object of rank 1 relative to those already observed is found. Hint: We assume that, if the rth object has rank 1 relative to the previous $(r-1)$ objects, then the probability that it is best is $r I N$. For $k \geq r^{*}$, let $\operatorname{Jk}(O)$ be the maximal probability of finding the best object assuming $k$ objects have been selected and the kth object is not best relative to the previous $(k-1)$ objects. Show that

$$
J k(O)=\frac{k}{N}\left(\frac{1}{N-1}+\ldots+\frac{1}{k}\right)
$$

A driver is looking for parking on the way to his destination. Each parking place is free with probability $p$ independently of whether other parking places are free or not. The driver cannot observe whether a parking place is free until he reaches it. If he parks $k$ places from his destination, he incurs a cost $k$. If he reaches the destination without having parked the cost is C .
(a) Let $P_{k}$ be the minimal expected cost if he is $k$ parking places from his destination, where $F_{O}=C$. Show that

$$
F_{k}=p \min \left(k, F_{k-1}\right)+q F_{k-1}, \quad k=1,2, \ldots,
$$

where $q=1-p$.
(b) Show that an optimal policy is of the form: never park if $k \geq k^{*}$, but take the first free place if $k<k^{*}$, where $k$ is the number of parking places from the destination and $k^{*}$ is the smallest integer i satisfying $q^{i-1}<(p C+q)-l$.

### 4.20 [Whi82]

A person may go hunting for a certain type of animal on a given day or stay home. When the animal population is $x$, the probability of capturing one animal is $p(x)$, a known increasing function, and the probability of capturing more than one is zero. A captured animal is worth one unit and a day of hunting costs $c$ units. Assume that $x$ does not change due to deaths or births, that the hunter knows $x$ at all times, that the horizon is finite, and that the terminal reward is zero. Show that it is optimal to hunt only when $p(x) \geq c$.

### 4.21

Consider the scalar linear system $X k+l=a X k+b U k$, where $a$ and $b$ are known At each period $k$ we have the option of using a control $U k$ and incurring a cost $q x_{k}^{2}+r u_{k}^{2}$, or else stopping and incurring a stopping cost $t x_{k}^{2}$. If we have not stopped by period $N$, the terminal cost is the stopping cost $t x_{N}^{2}$. We assume that $q \geq 0, r>0, t>0$. Show that there is a threshold value for $t$ below which immediate stopping is optimal at every initial state, and above which continuing at every state $X k$ and period $k$ is optimal

### 4.22

Consider a situation involving a blackmailer and his victim. In each period the blackmailer has a choice of: a) Accepting a lump sum payment of $R$ from the victim and promising not to blackmail again. b) Demanding a payment of $u$, where $u E[0,1]$. If blackmailed, the victim will either: 1) Comply with the demand and pay $u$ to the blackmailer. This happens with probability $1-u$. 2) Refuse to pay and denounce the blackmailer to the police. This happens with probability $u$. Once known to the police, the blackmailer cannot ask for any more
money. The blackmailer wants to maximize the expected amount of money he gets over $N$ periods by optimal choice of the payment demands $U k$. (Note that there is no additional penalty for being denounced to the police.) Write a DP algorithm and find the optimal policy.

### 4.23 [Whi82]

The Greek mythological hero Theseus is trapped in King Minos' Labyrinth maze. He can tryon each day one of $N$ passages. If he enters passage $i$ he will escape with probability $p i$, he will be killed with probability $q i$, and he will determine that the passage is a dead end with probability $\left(\begin{array}{l}1 \\ P i\end{array}-q i\right)$, in which case he will return to the point from which he started. Use an interchange argument to show that trying passages in order of decreasing Pi/qi maximizes the probability of escape within $N$ days.

### 4.24 (Hardy's Theorem)

Let $\{a l, \ldots, a n\}$ and $\{b l, \ldots, b n\}$ be monotonically nondecreasing sequences of numbers. Let us associate with each $\mathrm{i}=1, \ldots, n$ a distinct index $j i$, and consider the expression $\sum_{i=1}^{n} a i b_{j i}$. Use an interchange argument to show that this expression is maximized when $j i=\mathrm{i}$ for all i , and is minimized when $j i=n-\mathrm{i}+1$ for all i .

### 4.25

A busy professor has to complete $N$ projects. Each project $k$ has a deadline $d k$ and the time it takes the professor to complete it is $t k$. The professor can work on only one project at a time and must complete it before moving on to a new project. For a given order of completion of the projects, denote by $c k$ the time of completion of project $k$, Le.,

$$
c_{k}=t_{k}+\underbrace{}_{\substack{\text { projects i } \\ \text { completed before } k}}
$$

The professor wants to order the projects so as to minimize the maximum tardiness, given by $\operatorname{maxkE}\{1, \ldots, \mathrm{~N}\} \max \left(\mathrm{O}, C k-d_{k}\right)$. Use an interchange argument to show that it is optimal to complete the projects in the order of their deadlines (do the project with the closest deadline first).

### 4.26

Assume that we have two gold mines, Anaconda and Bonanza, and a gold-mining machine. Let $X A$ and $X B$ be the current amounts of gold in Anaconda and Bonanza, respectively ( $X_{A}$ and $X_{B}$ are integer). When the machine is used in Anaconda or Bonanza, there is a probability $P$ that $\mathfrak{r}_{A X A} 1$ (or $\mathfrak{r}_{r B X B} 1$, respectively)
of the gold will be mined without damaging the machine, and a probability $1-P$ that the machine will be damaged beyond repair and no gold will be mined. We assume that $0<r A<1$ and $0<r B<1$. We want to find a policy that selects the mine in which to use the machine at each period so as to maximize the total expected amount of gold mined.
(a) Use an interchange argument to show that it is optimal to mine Anaconda if and only if $r A X A \geq r B X B$.
(b) Solve the problem for $X A=2, x B=4, r A=0.4, r B=0.6, P=0.9$.

### 4.27

Consider the quiz contest problem of Example 5.1, where is an order constraint that each question i may be answered only after a given number $k_{i}$ of other questions have been answered. Use an interchange argument to show that an optimal list can be constructed by ordering the questions in decreasing order of $P_{i} R_{i} /(I-P i)$ and by sequentially answering the top question in the list out of those that are available (have not yet been answered and satisfy the order constraints).

### 4.28

0
Consider the quiz contest problem of Example 5.1, where there is a cost $F_{i} \geq$ for failing to answer question i correctly (in addition to losing the reward $R i$ ),
(a) Use an interchange argument to show that it is optimal to answer the questions in order of decreasing (PiRi-(1-Pi)Fi)/(1-Pi)'
(b) Solve the variant of the problem where there is an option to stop answering questions.

### 4.29

Consider the quiz contest problem of Example 5.1, where there is a maximum number of questions that can be answered, which is smaller than the number of questions that are available.
(a) Show that it is not necessarily optimal to answer the questions in order of decreasing PiRi/(l-Pi)' Hint: Try the case where only one out of two available questions can be answered.
(b) Give a simple algorithm to solve the problem where the number of available questions is one more than the maximum number of questions that can be answered.

### 4.30 (Reachability of One-Dimensional Linear Systems)

Generalize the analysis of Example 4.6.1 for the case of the one-dimensional linear system

$$
x_{k+1}=a x_{k}+b u_{k}+w_{k}, \quad k=0, \ldots, N-1
$$

where $x_{k}$ should be kept within an interval $[-\alpha, \alpha]$, using controls from an interval $[\beta, \beta]$, and in spite of the effects of the disturbances that can take values from the interval,]. Derive an algorithm to generate the effective target sets, and characterize the set of initial states from which reachability of the target tube is guaranteed. What happens to this set as $N \rightarrow \mathbf{o o}$ ?

### 4.31 (Reachability of Ellipsoidal Tubes [BeR71b], [Ber72a]) (www)

Consider the linear system

$$
x_{k+1}=A_{k} x_{k}+B_{k} u_{k}+w_{k}
$$

where the controls $u_{k}$ and the disturbances $w_{k}$ must belong to the ellipsoids

$$
U_{k}=\left\{x \mid u^{\prime} R_{k} u \leq 1\right\}, \quad W_{k}=\left\{x \mid w^{\prime} D_{k} w \leq 1\right\}
$$

where $R_{k}$ and $D_{k}$ are given positive definite symmetric matrices.
(a) Focus on a single period $k$, and consider the problem of finding an ellipsoid

$$
\underline{x}=\left\{x \mid x^{\prime} K x \leq \mathrm{I}\right\}
$$

where $K$ is a positive definite symmetric matrix, such that $\underline{X}$ is contained in the intersection of the following two sets: (1) an ellipsoid $\left\{x \mid x^{\prime} \Xi x \leq \mathrm{I}\right\}$, where $\Xi$ is a positive definite symmetric matrix, and (2) the set of all states $x$ such that there exists a $U \mathrm{E} U_{k}$ with the property that for all $W \mathrm{E} W_{k}$, we have $A k X+B k U+w E X$, where

$$
X=\left\{x \mid x^{\prime} \Psi x \leq \mathrm{I}\right\}
$$

and $\Psi$ is a given positive definite symmetric matrix. Show that if for some scalar $\beta \mathrm{E}(0,1)$, the matrix

$$
F^{-1}=(1-\beta)\left(\Psi^{-1}-\beta^{-1} D_{k}^{-1}\right)
$$

is well-defined as a positive definite matrix, an appropriate matrix $K$ is given by

$$
K=A_{k}^{\prime}\left(F^{-1}+B_{k} R_{k}^{-1} B_{k}^{\prime}\right)^{-1} A_{k}+\Xi
$$

Furthermore, the linear control law

$$
\mu(x)=-\left(R_{k}+B_{k}^{\prime} F B_{k}\right)^{-1} B_{k}^{\prime} F A_{k} x
$$

satisfies the constraint $\mu(x) \mathrm{E} U_{k}$ for all $x \mathrm{E} \underline{X}$ and achieves reachability of $X$ if $x$ E $\underline{X}$, i.e., $\boldsymbol{\mu}$ is such that $A k x+B_{k} \mu(x)+W \mathrm{E} X$ for all $x$ E $X$ and $w$ E Wk. Hint: Use the fact that the vector sum of two ellipsoids $\left\{x \mid x^{\prime} E l x \leq \mathrm{I}\right\}$ and $\left\{x \mid x^{\prime} E_{2} x \leq \mathrm{I}\right\}$ (with $E l$ and $E_{2}$ positive definite symmetric) is contained in the ellipsoid $\left\{x \mid x^{\prime} E x \leq \mathrm{I}\right\}$, where

$$
E^{-1}=\beta^{-1} E_{1}^{-1}+(1-\beta)^{-1} E_{2}^{-1}
$$

and $\beta$ is any scalar with $\quad \begin{gathered}\mathrm{O} \\ <\beta<1\end{gathered}$.
(b) Consider an ellipsoidal target tube $\left\{\hat{X}_{0}, \hat{X}_{1}, \ldots, \hat{X}_{N}\right\}$, where

$$
\hat{X}_{k}=\left\{x \mid x^{\prime} \Xi_{k} x \leq 1\right\}
$$

and the $\Xi_{k}$ are given positive definite symmetric matrices. Let the matrix sequences $\{F k\}$ and $\left\{K_{k}\right\}$ be generated by the algorithm

$$
\begin{gathered}
K_{N}=\Xi_{N}, \\
F_{k+1}^{-1}=\left(1-\beta_{k}\right)\left(K_{k+1}^{-1}-\beta_{k}^{-1} \mathrm{D} ; ; 1\right), \quad k=0,1_{0} \ldots, N-1, \\
K_{k}=A_{k}^{\prime}\left(\mathrm{F}_{\mathrm{k}+1}+B_{k} R_{k} l B_{k}^{\prime}\right)-l A_{k}+\Xi_{k}, \quad k=, 1, \ldots, N-1,
\end{gathered}
$$

where $\beta_{k}$ are scalars satisfying $<\beta_{k}<1$. Use the procedure of part (a) to show that a linear control law of the form

$$
\mu_{k}\left(x_{k}\right)=-\left(R_{k}+B_{k}^{\prime} F_{k+1} B_{k}\right)^{-1} B_{k}^{\prime} F_{k+1} A_{k} x_{k}, \quad k=O, l, \ldots, N-1
$$

achieves reachability of the target tube, provided the matrices $F_{k}$ are welldefined as positive definite matrices and $X o$ satisfies $x_{0}^{\prime} K_{0}^{-1} x_{0} \leq 1$.
(c) Suppose that the matrices $A k, B_{k}, R k, D_{k}$, and $\Xi_{k}$ do not depend on $k$, and that the algebraic matrix equation

$$
K=A^{\prime}\left((1-\beta)\left(K^{-1}-\beta^{-1} D^{-1}\right)+B R^{-1} B^{\prime}\right) \quad+\Xi
$$

has a positive definite solution $\underline{K}$ for some $\beta E(0,1)$ for which the matrix

$$
F_{-} 1=\left(1-\underline{(3)\left(K_{-}{ }^{\prime}\right.} \quad \beta^{-1} D_{-}^{1}\right)
$$

is well defined as a positive definite matrix. Show that if the initial state belongs to the set $\underline{X}=\left\{x \mid x^{\prime} \underline{K} x \leq I\right\}$, then all subsequent states will belong to $\underline{X}$ when the stationary linear control law

$$
\mu(x)=-\left(R+B^{\prime} F B\right)-l B^{\prime} F A x
$$

is used.

### 4.32 (Pursuit-Evasion Games and Reachability [BeR71b])

Consider the linear system

$$
X k+l=A_{k} x_{k}+B_{k} u_{k}+G k V k, \quad k=0,1, \ldots, N-1
$$

where the controls $U k$ and $v k$ are selected by two antagonistic players from sets $U k$ and $V k$, respectively, with exact knowledge of $x k$ (but without knowledge of the other player's choice at time $k$ ). The player selecting $U k$ aims to bring the state of the system within some given set $X$ at some time $k=1, \ldots, N$, while the player selecting $v_{k}$ aims to keep the state of the system outside the set $X$ at all times $k=1, \ldots, N$. Relate this problem to the problem of reachability of a target tube, and characterize the sets of initial conditions $X o$ starting from which the two players are guaranteed to achieve their objective with suitable choice of their control laws.

A famous but somewhat vain opera singer is scheduled to sing on $N$ successive nights. If she is satisfied with her performance on a given night $k$ (which happens with probability $p$, independently of the previous history) she will sing on the following night (i.e., night $k+1$ ). If she is not satisfied, however, she sulks and declares that she will not sing further. In this case, the only way to placate her into performing on the following night is for the opera director to send her an expensive gift, costing G dollars, which successfully placates her with probabiiity $q$ (independently of the previous history). If the gift does not placate her, the missed performance costs the opera house C dollars. The opera director may send a gift on any night, regardless of the success he has had with gifts on previous nights. The objective is to find a policy for when to send a gift and when not to, that minimizes the total cost from the $N$ nights.
(a) Write a DP algorithm for solving the problem, and characterize as best as you can the optimal policy.
(b) Repeat part (a) for the case where the probability $q$ is not constant, but rather is a decreasing function of the current stage.

### 4.34

An enterprising but somewhat foolish graduate student has invested the tuition for next semester in the stock market. As a result, he currently possesses a certain amount of stock that he/she must sell by registration day, which is $N$ days away. The stock must be sold in its entirety on a single day, and will then be deposited in a bank where it will earn interest at a daily rate $r$. The value of the stock on day $k$ is denoted by $x k$ and it evolves according to

$$
x_{k+1}=\lambda x_{k}+w_{k}, \quad x_{0}: \text { given },
$$

where $\lambda$ is a scalar with $0<\lambda<1$, and $w k$ is a random variable taking one of a finite number of positive values. We assume that $w_{o}, \ldots, w n-1$ are independent and identically distributed. The student wants to maximize the expected value of the money he/she has on registration day.
(a) Write a DP algorithm for solving the problem, and characterize as best as you can the optimal policy.
(b) Assume that the student has the option of selling only a portion of his stock on a given day. What if anything would he/she do different?

# Problems with Imperfect State 

## Information

We have assumed so far that the controller has access to the exact value of the current state, but this assumption is often unrealistic. For example, some state variables may be inaccessible, the sensors used for measuring them may be inaccurate, or the cost of obtaining the exact value of the state may be prohibitive. We model situations of this type by assuming that at each stage the controller receives some observations about the value of the current state, which may be corrupted by stochastic uncertainty.

Problems where the controller uses observations of this type in place of the state are called problems of imperfect state information, and are the subject of this chapter. We will find that even though there are DP algorithms for imperfect information problems, these algorithms are far more computationally intensive than in the perfect information case. For this reason, in the absence of an analytical solution, imperfect information problems are typically solved suboptimallyin practice. On the other hand, we will see that conceptually, imperfect state information problems are no different than the perfect state information problems we have been studying so far. In fact by various reformulations, we can reduce an imperfect state information problem to one with perfect state information. We will study two different reductions of this type, which will yield two different DP algorithms. The first reduction is the subject of the next section, while the second reduction will be given in Section 5.4.

### 5.1 REDUCTION TO THE PERFECT INFORMATION CASE

We first formulate the imperfect state information counterpart of the basic problem.

## Basic Problem with Imperfect State Information

Consider the basic problem of Section 1.2 where the controller, instead of having perfect knowledge of the state, has access to observations $Z k$ of the form

$$
Z o=h o(x o, v a), \quad z_{k}=h_{k}\left(x_{k}, u_{k-1}, v_{k}\right), \quad k=1,2, \ldots, N-1 .
$$

The observation $Z k$ belongs to a given observation space $Z k$. The random observation disturbance $V k$ belongs to a given space $V k$ and is characterized by a given probability distribution

$$
P_{V k}(. \mid X k, \cdots, X o, U k-l, \cdot ., U o, W k-l, \cdots, W o, V k-l, \cdots, v a) \text {, }
$$

which depends on the current state and the past states, controls, and disturbances.

The initial state $X O$ is also random and characterized by a given probability distribution $P_{x o}{ }^{\prime}$ The probability distribution $P_{W k}(. I X k, U k)$ of $W k$ is given, and it may depend explicitly on $X k$ and $U k$ but not on the prior disturbances $W 0, \ldots, W k-l, V o, . \cdot, V k-l$. The control $U k$ is constrained to take values from a given nonempty subset $U k$ of the control space $O k$. It is assumed that this subset does not depend on $X k$.

Let us denote by $I k$ the information available to the controller at time $k$ and call it the information vector. We have

$$
\begin{align*}
I k & =(z o, Z 1, \cdot ., Z k, U 0, \mathrm{U} 1, . ., U k-l), \quad k=1,2, \ldots, N-1  \tag{5.1}\\
{[0} & =\text { Zoo }
\end{align*}
$$

We consider the class of policies consisting of a sequence of functions $\pi=\left\{\mu_{0}, \mu_{1}, \ldots, \mu_{N-1}\right\}$, where each function, uk maps the information vector $I_{k}$ into the control space $O k$ and

$$
\mu_{k}\left(I_{k}\right) \in U_{k}, \quad \text { for all } I k, \quad k=0,1, \ldots, N-1
$$

Such policies are called admissible. We want to find an admissible policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots, \mu_{N^{-}}{ }_{1}\right\}$ that minimizes the cost function

$$
J_{\pi}=\underset{\substack{x_{0}, w_{k}, v_{k} \\ k=0, \ldots, N-1}}{E}\left\{g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(x_{k}, \mu_{k}\left(I_{k}\right), w_{k}\right)\right\}
$$

subject to the system equation

$$
x_{k+1}=f_{k}\left(x_{k}, \mu_{k}\left(I_{k}\right), w_{k}\right), \quad k=0,1, \ldots, N-1
$$

and the measurement equation

$$
\begin{aligned}
& Z o=h o(x o, v o), \\
& Z k=h_{k}\left(x_{k}, \mu_{k-1}\left(I_{k-1}\right), v_{k}\right), \quad k=1,2, \ldots, N-1 .
\end{aligned}
$$

Note the difference from the perfect state information case. Whereas before we were trying to find a rule that would specify the control $U k$ to be applied for each state $X k$ and time $k$, now we are looking for a rule that gives the control to be applied for every possible information vector $I_{k}$ (or state of information), that is, for every sequence of observations received and controls employed up to time $k$.

## Example 5.1.1 (Multiaccess Communication)

Consider a collection of transmitting stations sharing a common channel, for example, a set of ground stations communicating with a satellite at a common frequency. The stations are synchronized to transmit packets of
data at integer times. Each packet requires one time unit (also called a slot) for transmission. The total number $a k$ of packet arrivals during slot $k$ is independent of prior arrivals and has a given probability distribution. The stations do not know the backlog $X k$ at the beginning of the kth slot (the number of packets waiting to be transmitted). Packet transmissions are scheduled using a strategy (known as slotted Aloha) whereby each packet residing in the system at the beginning of the kth slot is transmitted during the kth slot with probability $U k$ (common for all packets). If two or more packets are transmitted simultaneously, they collide and have to rejoin the backlog for retransmission at a later slot. However, the stations can observe the channel and determine whether in anyone slot there was a collision (two or more packets), a success (one packet), or an idle (no packets). These observations provide information about the state of the system (the backlog $X k$ ) and can be used to select appropriately the control (the transmission probability $U k$ ). The objective is to keep the backlog small, so we assume a cost per stage $g k(X k)$, which is a monotonically increasing function of $X k$

The state of the system here is the backlog $X k$ and evolves according to the equation

$$
X k+l=X k+a k \quad t k
$$

where $a k$ is the number of new arrivals and $t k$ is the number of packets successfully transmitted during slot $k$. Both $a k$ and $t k$ may be viewed as disturbances, and the distribution of $t k$ depends on the state $X k$ and the control $U k$. It can be seen that $t k=1$ (a success) with probability $x k u k(l$ $U k) X k k^{1}$, and $t k=0$ (idle or collision) otherwise [the probability of anyone of the $X k$ waiting packets being transmitted, while all the other packets are not transmitted, is $u k(l-U k) X k^{-}{ }^{1}{ }^{1}$.

If we had perfect state information (i.e., the backlog $X k$ were known at the beginning of slot $k$ ), the optimal policy would be to select the value of $U k$ that maximizes $x k u k\left(l-u_{k}\right)^{x_{k}-1}$, which is the success probability.t By setting the derivative of this probability to zero, we find the optimal (perfect state information) policy to be

$$
\mu_{k}\left(x_{k}\right)=\frac{1}{x_{k}}, \quad \text { for all } X k \geq 1
$$

$t$ For a more detailed derivation, note that the DP algorithm for the perfect state information problem is

$$
\begin{aligned}
J k(X k)=g k(X k)+\min _{0 \leq u_{k} \leq 1} \underset{a k}{E}\{ & P\left(X k^{\prime} U k\right) J k+l(X k+a k-1) \\
& +(1-p(X k, U k)) J k+l(X k+a k)\}
\end{aligned}
$$

where $p(X k, U k)$ is the success probability $x k u k(l-U k) X k^{-1}$. Since the cost per stage $g k(X k)$ is an increasing function of the backlog $X k$, it is clear that each cost-to-go function $J k(X k)$ is an increasing function of $X k$ (this can also be proved by induction). Thus $\mathrm{Jk}+1(X k+a k) \geq \mathrm{J}_{\mathrm{k}}+1(X k+a k-1)$ for all $X k$ and $a k$, based on which the DP algorithm implies that the optimal $U k$ maximizes $p(X k, U k)$ over [O,IJ.

In practice, however, $X k$ is not known (imperfect state information), and the optimal control must be chosen on the basis of the available observations (i.e., the entire channel history of successes, idles, and collisions). These observations relate to the backlog history (the past states) and the past transmission probabilities (the past controls), but are corrupted by stochastic uncertainty. Mathematically, we may write an equation $Z k+l=V k+l$, where $Z k+l$ is the observation obtained at the end of the kth slot, and the random variable $V k+l$ yields an idle with probability $\left(1-u_{k}\right)^{x_{k}}$, a success with probability $x k u k(l-U k) X k-{ }^{-}$, and a collision otherwise.

It can be seen that this is a problem that fits the given imperfect state information framework. Unfortunately, the optimal solution to this problem is very complicated and for all practical purposes cannot be computed. A suboptimal solution will be discussed in Section 6.1.

## Reformulation as a Perfect State Information Problem

We now show how to effect the reduction from imperfect to perfect state information. As in the discussion of state augmentation in Section 1.4, it is intuitively clear that we should define a new system whose state at time $k$ is the set of all variables the knowledge of which can be of benefit to the controller when making the kth decision. Thus a first candidate as the state of the new system is the information vector $I_{\boldsymbol{k}}$. Indeed we will show that this choice is appropriate.

We have by the definition of the information vector ref. Eq. (5.1)]

$$
\begin{equation*}
I_{k+1}=\left(I_{k}, Z k+l, U k\right), \quad k=0,1, \ldots, N-2, \quad 10=Z o o \tag{5.2}
\end{equation*}
$$

These equations can be viewed as describing the evolution of a system of the same nature as the one considered in the basic problem of Section 1.2. The state of the system is $I_{k}$, the control is $U k$, and $Z k+l$ can be viewed as a random disturbance. Furthermore, we have

$$
P\left(Z k+l \mid I_{k}, u_{k}\right)=P\left(Z k+l l I_{k}, u_{k}, z_{0}, z_{1}, \ldots, Z k\right)
$$

since $Z o, Z 1, \ldots, Z k$ are part of the information vector $1 k$. Thus the probability distribution of $Z k+l$ depends explicitly only on the state $I_{k}$ and control $U k$ of the new system (5.2) and not on the prior "disturbances" $Z k, \cdots, Z o o$ By writing

$$
E\{g k(X k, U k, W k)\}=E\left\{\underset{x k, w k}{E}\left\{9 k\left(X k, u_{k}, W k\right) 11_{\mathrm{k}, U}, U k\right\}\right\}
$$

we can similarly reformulate the cost function in terms of the variables of the new system. The cost per stage as a function of the new state $I_{k}$ and the control $U k$ is

$$
\begin{equation*}
\tilde{g}_{k}\left(I_{k}, u_{k}\right)=\underset{X k, W k}{E}\left\{9 k(X k, U k, W k) \mid I_{k}, u_{k}\right\} \tag{5.3}
\end{equation*}
$$

Thus the basic problem with imperfect state information has been reformulated as a problem with perfect state information that involves the system (5.2) and the cost per stage (5.3). By writing the DP algorithm for this latter problem and substituting the expressions (5.2) and (5.3), we obtain

$$
\begin{equation*}
+9 N-\mathrm{I}(X N-I, U, N-I, W N-I) \mid I N-I, U N-I\}] \tag{5.4}
\end{equation*}
$$

and for $k=0,1, \ldots, N-2$,
$J k(I k)=\min _{u k E U k}\left[{ }_{x k, w_{k}, Z k+l}^{E}\left\{9 k\left(X k^{\prime} u_{k}, W k\right)+J k+1\left(I k, Z k+1, u_{k}\right) 11 k, u_{k}\right\}\right]$.
These equations constitute one possible DP algorithm for the imperfect state information problem. An optimal policy $\left\{\mu_{0}^{*}, \mu_{1}^{*}, \ldots, \mu_{N}^{*}-\mathrm{I}\right\}$ is obtained by first minimizing in the right-hand side of the DP Eq. (5.4) for every possible value of the information vector $I N_{-1}$ to obtain $\mu_{N^{-1}}^{*}\left(I N_{-1}\right)$. Simultaneously, $\mathrm{JN}_{\mathrm{N}} \mathbf{1}\left(I N_{-1}\right)$ is computed and used in the computation of $I N-2(I N-2)$ via the minimization in the DP Eq. (5.5), which is carried out for every possible value of $I N-2$ ' Proceeding similarly, $I N-3(I N-3)$ and $\mu_{N-3}^{*}$ and so on are obtained, until $\operatorname{Jo}(I o)=\operatorname{Jo}(z o)$ is computed. The optimal cost $\mathrm{J}^{*}$ is then given by

$$
\mathrm{J}^{*}=E_{z_{Q}}\{J o(z o)\}
$$

## Machine Repair Example

A machine can be in one of two states denoted $P$ and $\bar{P}$. State $P$ corresponds to a machine in proper condition (good state) and state $\bar{P}$ to a machine in improper condition (bad state). If the machine is operated for one time period, it stays in state $P$ with probability $\frac{2}{3}$ if it started in $P$, and it stays in state $\bar{P}$ with probability 1 if it started in $\bar{P}$. The machine is operated for a total of three time periods and starts in state $P$. At the end of the first and second time periods the machine is inspected and there are two possible inspection outcomes denoted $G$ (probably good state) and $B$ (probably bad state). If the machine is in the good state $P$, the inspection outcome is $G$ with probability $\frac{3}{4}$; if the machine is in the bad state $\bar{P}$, the inspection outcome is $B$ with probability $\frac{3}{4}$ :

$$
P(G l x=P)=\frac{3}{4}, \quad P(B \mid x \quad P)=\frac{1}{4}
$$

$$
P(G \mid x=\bar{P})=\frac{1}{4}, \quad P(B \mid x=\bar{P})=\frac{3}{4}
$$

see Fig. 5.1.1. After each inspection one of two possible actions can be taken:

C: Continue operation of the machine.
$S$ : Stop the machine, determine its state through an accurate diagnostic test, and if it is in the bad state $\bar{P}$ bring it back to the good state $P$.

At each period there is a cost of 2 and 0 units for starting the period with a machine in the bad state $\bar{P}$ and the good state $P$, respectively. The cost for taking the stop-and-repair action $S$ is 1 unit and the terminal cost is 0 .


Figure 5.1.1 State transition diagram and probabilities of inspection outcomes in the machine repair example.

The problem is to determine the policy that minimizes the expected costs over the three time periods. In other words, we want to find the optimal action after the result of the first inspection is known, and after the results of the first and second inspections, as well as the action taken after the first inspection, are known.

It can be seen that this example falls within the general framework of the problem of this section. The state space consists of the two states $P$ and $\bar{P}$,

$$
\text { state space }=\{P, \bar{P}\}
$$

and the control space consists of the two actions

$$
\text { control space }=\{C, S\} .
$$

The system evolution may be described by introducing a system equation

$$
X k+1=W k, \quad k=0,1,
$$

where for $k=0,1$, the probability distribution of $W k$ is given by

$$
P\left(W k=P \mid X k=P, u_{k}=\mathrm{C}\right)=\frac{2}{3}, \quad P\left(W k=\bar{P} \mid X k=\quad u_{k}=\mathrm{C}\right)=\frac{1}{3},
$$

$$
\begin{array}{ll}
P(W k \quad P \mid X k \quad U k=\mathrm{C}) \quad 0, & P(W k=\bar{P} \mid x k=\bar{P}, U k=\mathrm{C})=1 \\
P(W k=P \mid X k=P, U k=S)=\frac{2}{3}, & P(W k=\bar{P} \mid X k=P, U k=S)=\frac{1}{3} \\
P(W k \quad P I X k=\quad u k=S)=\frac{2}{3}, & P(W k=\bar{P} \mid X k=\bar{P}, U k=S)=\frac{1}{3}
\end{array}
$$

We denote by $X 0, \mathrm{Xl}, \mathrm{X} 2$ the state of the machine at the end of the first, second, and third time period, respectively. Also we denote by Uo the action taken after the first inspection (end of first time period) and by Ul the action taken after the second inspection (end of second time period). The probability distribution of $X o$ is

$$
P\left(\begin{array}{ll}
X o & P
\end{array}\right)=\frac{2}{3}, \quad P(x o=\bar{P})=\frac{1}{3}
$$

Note that we do not have perfect state information, since the inspections do not reveal the state of the machine with certainty. Rather the result of each inspection may be viewed as a measurement of the system state of the form

$$
z_{k}=v_{k}, \quad k \quad 0,1
$$

where for $k=0,1$, the probability distribution of $V k$ is given by

$$
\begin{aligned}
& P(V k=\mathrm{G} \mid X k=P)=\frac{3}{4}, \quad P(V k=B \quad \mid X k=P)=\frac{1}{4} \\
& P(V k=\mathrm{G} \mid X k=\bar{P})=\frac{1}{4}, \quad P(V k=B \quad \mid X k=\bar{P})=\frac{3}{4}
\end{aligned}
$$

The cost resulting from a sequence of states $X o, \mathrm{Xl}$ and actions $U 0, \mathrm{Ul}$ is

$$
g(X o, u o)+g(X l, \mathrm{Ul})
$$

where

$$
g(P, \mathrm{C}) \quad 0, \quad g(P, S)=1, \quad g(\bar{P}, \mathrm{C})=2, \quad g(\bar{P}, S)=1
$$

The information vector at times and 1 is

$$
10=Z o, \quad I_{1}=(z o, \mathrm{Zl}, u o)
$$

and we seek functions $\mu_{0}(I o), \mu 1\left(I_{1}\right)$ that minimize

$$
\begin{aligned}
& \underset{\substack{\boldsymbol{w}_{0}, v_{I} \\
v_{0},{ }_{2}}}{\boldsymbol{E}} \text { wI }\left\{g\left(x o, \mu_{0}\left(I_{0}\right)\right)+g\left(X l^{\prime} \mu_{1}\left(I_{1}\right)\right)\right\} \\
& =\underset{\substack{x Q, W Q, W I \\
v Q, \mathrm{vI}}}{E}\left\{g\left(x_{0}, \mu_{0}\left(z_{0}\right)\right)+g\left(x_{1}, \mu_{1}\left(z_{0}, z_{1}, \mu_{0}\left(z_{0}\right)\right)\right)\right\} .
\end{aligned}
$$

We now apply the DP algorithm. It involves taking the minimum over the two possible actions, C and S , and it has the form

$$
\begin{gathered}
J_{k}\left(I_{k}\right)=\min \left[P\left(X k=P 11_{k}, C\right) g(P, \mathrm{C})+P\left(X k=\bar{P} \mid I_{k}, C\right) g(P, \mathrm{C})\right. \\
\quad+\underset{Z k+I}{E}\left\{J_{k+1}\left(I_{k}, \mathrm{C}, Z k+l\right) \mid I_{k}, C\right\}, \\
P\left(X k=P \mid I_{k}, S\right) g(P, S)+P\left(X k=\bar{P} \mid I_{k}, S\right) g(\bar{P}, S) \\
\left.\quad+\underset{Z k+I}{E}\left\{J_{k+1}\left(I_{k}, S, z_{k+1}\right) \mid I_{k}, S\right\}\right]
\end{gathered}
$$

where $k=0,1$, and the terminal condition is $J_{2}\left(I_{2}\right)=0$.
Last Stage: We use Eq. (5.4) to compute $J l\left(I_{1}\right)$ for each of the eight possible information vectors // (zo, zl,uo). As indicated by the above DP algorithm, for each of these vectors, we shall compute the expected cost of the possible actions, $\mathrm{Ul}=\mathrm{C}$ and $\mathrm{Ul}=S$, and select as optimal the action with the smallest cost. We have

$$
\text { cost of } \mathrm{C}=2 \cdot P\left(X I=\bar{P} \mid I_{1}\right), \quad \text { cost } \text { of } S=1
$$

and therefore

$$
J_{1}\left(I_{1}\right)=\min \left[2 P\left(X I=\bar{P} \mid I_{1}\right), 1\right] .
$$

The probabilities $P\left(X l=\bar{P} \mid I_{1}\right)$ can be computed by using Bayes' rule and the problem data. Some of the details will be omitted. We have:
(1) For $I I=(G, G, S)$

$$
\begin{aligned}
P(\mathrm{Xl}=\bar{P} \mid G, \mathrm{G} S)= & \frac{P(X I=\bar{P}, \mathrm{G}, G I S)}{P(G, G I S)} \\
& \frac{\frac{1}{3} \frac{1}{4}\left(\frac{2}{3} \cdot \frac{3}{4}+\frac{1}{3} \cdot \frac{1}{4}\right)}{\left(\frac{2}{3} \cdot \frac{3}{4}+\frac{1}{3} \cdot \frac{1}{4}\right)^{2}}
\end{aligned}
$$

Hence

$$
J l(G, G, S)=\frac{2}{7}, \quad \mu_{1}^{*}(G, \mathrm{G}, S)=\mathrm{C}
$$

(2) For $I_{1}=(B, G, S)$

$$
P(X I=\bar{P} I B, G, S)=P(X I=\bar{P} I G, G, S)=\frac{1}{7}
$$

$$
J l(B, G, S)=\frac{2}{7}, \quad \mu_{1}^{*}(B, \mathrm{G}, S)=\mathrm{C}
$$

(3) For $I_{1}=(G, E, S)$

$$
\begin{aligned}
& P\left(_{X l}=\bar{P} 1 G B S\right)=\frac{P(X l=\bar{P}, G, E!S)}{P(G, E I S)} \\
&=\left(\frac{1}{3} \cdot \frac{3}{4} \cdot\left(\frac{2}{3} \cdot \frac{3}{4}+\frac{1}{3} \cdot \frac{1}{4}\right)\right. \\
&\left.\frac{1}{4}+\frac{1}{3} \cdot \frac{3}{4}\right)\left(\frac{2}{3} \cdot \frac{3}{4}+\frac{1}{3} \cdot \frac{1}{4}\right)-\frac{\mathbf{3}}{5} \\
& J_{l}(G, E, S)=1, \quad \mu_{1}^{*}(G, E, S)=S .
\end{aligned}
$$

(4) For $I_{1}=(E, B, S)$

$$
\bar{P} \mid B, E, S)=P(X|=\mathrm{P}| G, E, S)=\frac{3}{5}
$$

$$
J 1(B, B, S)=1, \quad \mu_{1}^{*}(B, B, S)=S
$$

(5) For $I_{1}=(G, G, C)$

$$
\begin{gathered}
\left.\left.P\left({ }_{X l}=\bar{P} \mid G G, C\right)=\frac{P(X l=\bar{P}, G, G}{P(G, G T C} \right\rvert\, \mathrm{C}\right) \\
\\
J I(G, G, \mathrm{C})=\frac{1}{5}, \quad \mu_{1}^{*}(G, G, C)=C
\end{gathered}
$$

(6) For $I_{1}=(B, G, C)$

$$
\begin{aligned}
P(X I & =\bar{P} \mid E, G, C)=\underline{21} \underline{2}^{\prime} \\
J 1(E, G, C) & =\underline{22} \underline{23}^{\prime} \quad \mu_{1}^{*}(B, G, C)=C .
\end{aligned}
$$

(7) For $I_{1}=(G, E, C)$

$$
P(X I=-\bar{P} \mid G, E, C)=\stackrel{9}{13}
$$

$$
J l(G, E, \mathrm{C})=1, \quad \mu_{1}^{*}(G, B, \text { C }) \quad S
$$

(8) For $I_{1}=(B, E, C)$

$$
\begin{gathered}
P(X 1=\overline{-} P I \quad E, B, C)=\underline{33} \\
J_{l}(E, E, C)=1, \quad \mu_{1}^{*}(B, E, C)=S .
\end{gathered}
$$

Summarizing the results for the last stage, the optimal policy is to continue $(U 1=C)$ if the result of the last inspection was $G$, and to stop $(U 1=S)$ if the result of the last inspection was $B$.

First Stage: Here we use the DP Eq. (5.5) to compute $J o(l o)$ for each of the two possible information vectors $10=(G), l_{0}=(B)$. We have

$$
\begin{aligned}
\text { cost of } \mathrm{C}= & 2 P(x o=\bar{P} \mid 10, \mathrm{C})+\underset{\mathrm{Zl}}{E}\{J l(10, \mathrm{Zl}, \mathrm{C}) \mid 10, \mathrm{C}\} \\
= & 2 P(x o=\bar{P} \mid 10, \mathrm{C})+P(\mathrm{Zl}=G \mid 10, C) J l(10, G, \mathrm{C}) \\
& +P(\mathrm{Zl}=B \mid 10, C) J l(10, B, C),
\end{aligned}
$$

cost of $\left.S=1+\underset{\mathrm{zl}}{E\{ } J l\left(I_{o}, Z l, S\right) 11_{0}, S\right\}$

$$
=1+P(Z l=G \mid 10, S) J_{1}\left(I_{0}, G, S\right)+P(Z l=E \quad 110, S) J_{l}(10, B, S),
$$

and

$$
\begin{gathered}
J o(1 o)=\min [2 P(X O=\bar{P} \mathrm{I} 10, \mathrm{C})+\underset{\mathrm{zl}}{E}\{J l(10, \mathrm{Zl}, \mathrm{C}) \mid 10, C\}, \\
1+\underset{\mathrm{Zl}}{E}\{J l(10, \mathrm{Zl}, \mathrm{~S}) \mid 10, S\}]
\end{gathered}
$$

(1) For $10=(G)$ : Direct calculation yields

$$
\begin{gathered}
P(Z l=G \mid G, C)=\frac{15}{28}, \quad P(Z l=B \quad \mid G, C)=\frac{13}{28}, \\
P(Z l=G \mid G, 8)=\underline{7}, \quad P(Z l=B I G, S)=\underline{12^{\prime}} \\
P(x o=-P I G, C)=\frac{1}{7},
\end{gathered}
$$

and hence

$$
\begin{gathered}
J_{0}(G)=\min \left[2 \cdot \frac{1}{7}+\frac{15}{28} J_{1}(G, G, C)+\frac{13}{28} J_{1}(G, B, C)\right. \\
\left.1+\frac{7}{12} J_{1}(G, G, S)+\frac{5}{12} J_{1}(G, B, S)\right]
\end{gathered}
$$

Using the values of $J l$ obtained in the previous stage

$$
\begin{gathered}
J o(G)=\min \left[2 \cdot \frac{1}{7}+\frac{15}{28} \cdot \frac{2}{5}+\frac{13}{2} \dot{8}^{1}, 1+\frac{-7}{12} \cdot \frac{2}{7}+\frac{5}{12} \cdot 1^{]}\right. \\
=\min \underline{[27} \underline{28^{\prime}} \underline{19} \underline{12}=\underline{28^{\prime}} \\
J o(G)=\underline{27} \quad \mu_{0}^{\prime}(G)=\mathrm{C} .
\end{gathered}
$$

(2) For $10=(B)$ : Direct calculation yields

$$
\begin{gathered}
P(Z|=G| B, C)=\frac{23}{\underline{60}} \quad P(Z|=B| B, C)=\underline{37} \\
P(Z|=G| B, S)=\underline{7}, \underline{12^{\prime}} \\
P(x o=\bar{P} \mid B, C)=\frac{3}{5}
\end{gathered}
$$

and

$$
\begin{gathered}
J_{0}(B)=\min \left[2 \cdot \frac{3}{5}+\frac{23}{60} J_{1}(B, G, C)+\frac{37}{60} J_{1}(B, B, C)\right. \\
\left.1+\frac{7}{12} J_{1}(B, G, S)+\frac{5}{12} J_{1}(B, B, S)\right]
\end{gathered}
$$

Using the values of $J_{l}$ obtained in the previous state

$$
\begin{aligned}
& J_{\mathrm{o}}(B)=\min \frac{[131}{60}, \frac{19]}{12}=\frac{19}{12^{\prime}} \\
& J o(B)=12^{\prime} \quad \mu_{0}^{*}(B)=S
\end{aligned}
$$

Summarizing, the optimal policy for both stages is to continue if the result of the latest inspection is $G$, and to stop and repair otherwise.

The optimal cost is

$$
J^{*}=P(G) J o(G)+P(B) J o(B)
$$

We can verify that $P(G)=\frac{72}{12}$ and $P(B)=\frac{5}{12}$, so that

$$
J^{*}=\frac{7}{12} \cdot \frac{27}{28}+\frac{5}{12} \cdot \frac{19}{12}=\frac{176}{144}
$$

In the above example, the computation of the optimal policy and the optimal cost by means of the DP algorithm (5.4) and (5.5) was possible because the problem was very simple. It is easy to see that for a more complex problem, the computational requirements of the DP algorithm can be prohibitive, particularly if the number of possible information vectors $I_{k}$ is large (or infinite). Unfortunately, even if the control and observation spaces are simple (one-dimensional or finite), the space of the information vector $I_{k}$ may have large dimension. This makes the application of the algorithm very difficult or computationally impossible in many cases. However, there are some problems where an analytical solution is possible, and the next two sections deal with such problems.

### 5.2 LINEAR SYSTEMS AND QUADRATIC COST

We will show how the DP algorithm of the preceding section can be used to solve the imperfect state information analog of the linear system/quadratic cost problem of Section 4.1. We have the same linear system

$$
x_{k+1}=A_{k} x_{k}+B_{k} u_{k}+w_{k}, \quad k=0,1, \ldots, N-1
$$

and quadratic cost

$$
E\left\{x_{N}^{\prime} Q_{N} x_{N}+\sum_{k=0}^{N-1}\left(x_{k}^{\prime} Q_{k} x_{k}+u_{k}^{\prime} R_{k} u_{k}\right)\right\}
$$

but now the controller does not have access to the current state. Instead it receives at the beginning of each period $k$ an observation of the form

$$
z_{k}=C_{k} x_{k}+v_{k}, \quad k=0,1, \ldots, N-1
$$

where $Z k \mathrm{E} \Re^{s}, C k$ is a given $s$ x $n$ matrix, and $v k \mathrm{E} \Re^{s}$ is an observation noise vector with given probability distribution. Furthermore, the vectors $V k$ are independent, and independent from $w k$ and $X 0$ as welL We make the same assumptions as in Section 4.1 concerning the input disturbances $w k$, i.e., that they are independent, zero mean, and that they have finite variance. The system matrices $A_{k}, B_{k}$ are known; there is no analytical solution of the imperfect information counterpart of the model with random system matrices considered in Section 4.1.

From the DP Eq. (5.4) we have

$$
\begin{aligned}
I N-I(I N-I)=\min _{U N-1}[ & \underset{X N-I, W N-I}{E}\left\{x_{N-1}^{\prime} Q_{N-1} x_{N-1}+u_{N-1}^{\prime} R_{N-1} u_{N-1}\right. \\
& +(A N-1 X N-I+B N-I U N-I+W N-I) \\
& . Q N(A N-1 X N-I+B N-I U N-I+W N-I) \text { IIN-I }\}]
\end{aligned}
$$

Since $E\{W N-I \mid I N-I\}=E\left\{w_{N-1}\right\}=0$, this expression can be written as

$$
\begin{align*}
J_{N-1}\left(I_{N-1}\right)= & \underset{x_{N-1}}{E}\left\{x_{N-1}^{\prime}\left(A_{N-1}^{\prime} Q_{N} A_{N-1}+Q_{N-1}\right) x_{N-1} \mid I_{N-1}\right\} \\
& +\underset{W N-1}{E}\left\{w_{N-1}^{\prime} Q_{N} w_{N-1}\right\} \\
& +\min _{U N-1}\left[u_{N-1}^{\prime}\left(B_{N-1}^{\prime} Q_{N} B_{N-1}+R_{N-1}\right) u_{N-1}\right. \\
& \left.\quad+2 E\left\{X N-| | I_{N-1}\right\}^{\prime} A_{N-1}^{\prime} Q N B N-1 U N-I\right]^{\prime} \tag{5.6}
\end{align*}
$$

The minimization yields the optimal policy for the last stage:

$$
\begin{aligned}
& =\mu_{N-1}^{*}\left(I_{N-1}\right) \\
& =-\left(B_{N-1}^{\prime} Q_{N} B_{N-1}+R_{N-1}\right)^{-1} B_{N-1}^{\prime} Q_{N} A_{N-1} E\left\{x_{N-1} \mid I_{N-1}\right\}^{2}
\end{aligned}
$$

and upon substitution in Eq. (5.6), we obtain

$$
\begin{aligned}
J_{N-1}\left(I_{N-1}\right)= & \underset{X N-1}{E}\left\{x_{N-1}^{\prime} K N-I X N-l \mid I_{N-1}\right\} \\
& +\underset{X N-1}{E}\left\{\left(X N-l-E\left\{X N-l \mid I_{N-1}\right\}\right)^{\prime}\right. \\
& \left.. P N-l\left(X N-l-E\left\{X N-l \mid I_{N-1}\right\}\right) \mid I N-I\right\} \\
& +\underset{W N-1}{E}\left\{w_{N-1}^{\prime} Q_{N} w_{N-1}\right\}
\end{aligned}
$$

where the matrices $K N-l$ and $P N-l$ are given by

$$
\begin{aligned}
& P N-l=A_{N-1}^{\prime} Q_{N} B_{N-1}\left(R_{N-1}+B_{N-1}^{\prime} Q_{N} B_{N-1}\right)^{-1} B_{N-1}^{\prime} Q_{N} A_{N-1} \\
& K N-l \quad A_{N-1}^{\prime} Q_{N} A_{N-1}-P N-l+\mathrm{QN}-1
\end{aligned}
$$

Note that the optimal policy (5.6) is identical to its perfect state information counterpart except that $X N-l$ is replaced by its conditional expectation $E\left\{X N--l \mid I_{N-1}\right\}$. Note also that the cost-to-go $I N-. l(I N-l)$ exhibits a corresponding similarity to its perfect state information counterpart except that $\mathrm{J} \mathrm{N}-1$ (IN-1) contains an additional middle term, which is in effect a penalty for estimation error.

Now the DP equation for period $N-2$ is
$J_{N-2}\left(I_{N-2}\right)=\min _{u_{N-2}}\left[{\underset{x}{N-2}, w_{N-2}, z_{N-1}}_{E}\left\{x_{N-2}^{\prime} Q_{N-2} x_{N-2}+u_{N-2}^{\prime} R_{N-2} u_{N-2}\right.\right.$

$$
+\mathrm{IN}--1(\mathrm{IN}-1) \mid I N-2, U N-2\}]
$$

$$
=E\left\{x_{N-2}^{\prime} Q_{N-2} x_{N-2} \mid I N-2\right\}
$$

$$
+\min _{U N-2}\left[u_{N-2}^{\prime} R_{N-2} u_{N-2}+E\left\{x_{N-1}^{\prime} K_{N-1} x_{N-1} \mid I N-2, u N-2\right\}\right]
$$

$$
+E\left\{(X N-I-E\{X N-l \mid I N-I\})^{\prime}\right.
$$

$$
P N-l(X N-l-E\{X N-l l I N-I\}) \mid I N-2, U N-2\}
$$

$$
\begin{equation*}
+\underset{W N-1}{E}\left\{w_{N-1}^{\prime} Q_{N} w_{N-1}\right\} \tag{5.8}
\end{equation*}
$$

Note that we have excluded the next to last term from the minimization with respect to $U N-2$. We have done so since this term turns out to be independent of $U N-2$. To show this fact, we need the following lemma.

The lemma says essentially that the quality of estimation as expressed by the statistics of the error $x_{k}-E\left\{x_{k} \mid I_{k}\right\}$ cannot be influenced by the choice of control. This is due to the linearity of both the system and the measurement equation. In particular, $x_{k}$ and $E\left\{x_{k} \mid I_{k}\right\}$ contain the same linear terms in (uo, ... ,
which cancel each other out.

Lemma 5.2.1: For every $k$, there is a function $M_{k}$ such that we have

$$
x_{k}-E\left\{x_{k} \mid I_{k}\right\}=M_{k}\left(x_{0}, w_{0}, \ldots, w_{k-1}, v_{0}, \ldots, v_{k}\right),
$$

independently of the policy being used.

Proof: Fix a policy and consider the following two systems. In the first system there is control as determined by the policy,

$$
+B_{k} u_{k}+w_{k}, \quad z_{k}=C_{k} x_{k}+v_{k}
$$

while in the second system there is no control,

$$
\bar{x}_{k+1}=A_{k} \bar{x}_{k}+\bar{w}_{k}, \quad \bar{z}_{k}=C_{k} \bar{x}_{k}+\bar{v}_{k}
$$

We consider the evolution of these two systems when their initial conditions are identical,

$$
X o=\bar{x}_{0},
$$

and when their system disturbance and observation noise vectors are also identical,

$$
W k=\bar{w}_{k}, \quad v_{k}=\bar{v}_{k}, \quad k=0,1, \ldots, N-1 .
$$

Consider the vectors

$$
\begin{array}{rlrl}
Z^{k} & =(z o, & , Z k)^{\prime}, & \bar{Z}^{k}=\left(\bar{z}_{0},\right. \\
W^{k} & =(w o, & \left., \bar{z}_{k}\right)^{\prime}, \\
& , W k)^{\prime}, & V^{k}=(v o, & , V k)^{\prime},
\end{array} \quad U^{k}=\left(u_{0}, \ldots, u_{k}\right)^{\prime} .
$$

Linearity implies the existence of matrices $F k^{\prime} G k$, and $H k$ such that

$$
\begin{aligned}
& X k=F k X O+G k U k-l+H k W^{k-1} \\
& \bar{x}_{k}=F k X O+H k W k \cdot-1
\end{aligned}
$$

Since the vector $U k-l=\left(u 0, \ldots, u_{k-1}\right)^{\prime}$ is part of the information vector $I_{k}$, we have $U k-l=E\{U k-l \mid I k\}$, so

$$
\begin{aligned}
& E\{X k \mid I k\}=F k E\left\{x o \mid I_{k}\right\}+G k U k-1+H k E\left\{W k-l \mid I_{k}\right\}, \\
& \left.E\left\{\bar{x}_{k} \mid I k\right\}=F k E\left\{X O \mid I_{k}\right\}+\quad \mid I_{k}\right\} .
\end{aligned}
$$

We thus obtain

$$
x_{k}-E\left\{x_{k} \mid I_{k}\right\}=\bar{x}_{k}-E\left\{\bar{x}_{k} \mid I_{k}\right\}
$$

From the equations for $z k$ and $\bar{z}_{k}$, we see that

$$
\bar{Z}^{k}=Z^{k}-R_{k} U^{k-1}=S_{k} W^{k-1}+T_{k} V^{k}
$$

where $R k^{\prime} S k$, and $T k$ are some matrices of appropriate dimension. Thus, the information provided by $I k=(Z k, U k-I)$ regarding $\bar{x}_{k}$ is summarized in $\quad$, and we have $E\left\{\bar{x}_{k} \mid I k\right\}=E\left\{\bar{x}_{k} \mid \bar{Z}^{k}\right\}$, so that

$$
x_{k}-E\left\{x_{k} \mid I_{k}\right\}=\bar{x}_{k}-E\left\{\bar{x}_{k} \mid \bar{Z}^{k}\right\}
$$

The function $M k$ given by

$$
M_{k}\left(x_{0}, w_{0}, \ldots, w_{k-1}, v_{0}, \ldots, v_{k}\right) \quad \bar{x}_{k}-E\left\{\bar{x}_{k} \mid \bar{Z}^{k}\right\}
$$

serves the purpose stated in the lemma. Q.E.D.
We can now justify excluding the term
$E\left\{\left(x_{N-1}-E\left\{x_{N-1} \mid I_{N-1}\right\}\right)^{\prime} P_{N-1}\left(x_{N-1}-E\left\{x_{N-1} \mid I_{N-1}\right\}\right) \mid I_{N-2}, u_{N-2}\right\}$
from the minimization in Eq. (5.8), as being independent of $U N-2^{\prime}$ Indeed, by using the lemma, we see that

$$
X N-1-E\left\{X N-1 \mid I_{N-1}\right\}=\xi_{N-1}
$$

where $\xi_{N-1}$ is a function of $X o, W o, . \cdot ., W N-2, v a, . \cdot ., V N-1$. Since $\xi_{N-1}$ is independent of $U N-2$, the conditional expectation of $\xi_{N^{-1} P N-1}^{\prime} \xi_{N-1}$ satisfies

$$
E\left\{\xi_{N-1}^{\prime} P_{N-1} \xi_{N-1} \mid I_{N-2}, u_{N-2}\right\}=E\left\{\xi_{N-1}^{\prime} P_{N-1} \xi_{N-1} \mid I_{N-2}\right\}
$$

Returning now to our problem, the minimization in Eq. (5.8) yields, using an argument similar to the one for the last stage,

$$
\begin{aligned}
& u_{N-2}^{*}=\mu_{N-2}^{*}\left(I_{N-2}\right) \\
& \quad=-\left(R N-2+B_{N-2}^{\prime} K_{N-1} B_{N-2}\right)^{-1} B_{N-2}^{\prime} K_{N-1} A_{N-2} E\left\{x_{N-2} \mid I N-2\right\}
\end{aligned}
$$

We can proceed similarly to obtain the optimal policy for every stage:

$$
\begin{equation*}
\mu_{k}^{*}\left(I_{k}\right)=L_{k} E\left\{x_{k} \mid I_{k}\right\} \tag{5.9}
\end{equation*}
$$

where the matrix $L k$ is given by

$$
L_{k}=-\left(R_{k}+B_{k}^{\prime} K_{k+1} B_{k}\right)^{-1} B_{k}^{\prime} K_{k+1} A_{k}
$$

with the matrices $K_{k}$ given recursively by the Riccati equation

$$
\begin{gathered}
K_{N}=Q_{N} \\
P_{k}=A_{k}^{\prime} K_{k+1} B_{k}\left(R_{k}+B_{k}^{\prime} K_{k+1} B_{k}\right)^{-1} B_{k}^{\prime} K_{k+1} A_{k} \\
K k=A_{k}^{\prime} K_{k+1} A_{k}-P_{k}+Q k
\end{gathered}
$$

The key step in this derivation is that at stage $k$ of the DP algorithm, the minimization over $U k$ that defines $J_{k}\left(I_{k}\right)$ involves the additional terms

$$
E\left\{\left(x_{s}-E\left\{x_{s} \mid I_{s}\right\}\right)^{\prime} P_{s}\left(x_{s}-E\left\{x_{s} \mid I_{s}\right\}\right) \mid I_{k}, u_{k}\right\}
$$

where $s=k+1, \ldots, N-1$. By using the argument of the proof of the earlier lemma, it can be seen that none of these terms depends on $U k$ so that the presence of these terms does not affect the minimization in the DP algorithm. As a result, the optimal policy is the same as the one for the perfect information case, except that the state $X k$ is replaced by its conditional expectation $E\left\{X k \mid I_{k}\right\}$.

It is interesting to note that the optimal controller can be decomposed into the two parts shown in Fig. 5.2.1:
(a) An estimator, which uses the data to generate the conditional expectation $E\left\{X k \mid I_{k}\right\}$.
(b) An actuator, which multiplies $E\{X k \mid I k\}$ by the gain matrix and applies the control input $u_{k}=\operatorname{LkE}\left\{X k \mid I_{k}\right\}$.
Furthermore, the gain matrix $L_{k}$ is independent of the statistics of the problem and is the same as the one that would be used if we were faced with the deterministic problem, where $W k$ and $x Q$ would be fixed and equal to their expected values. On the other hand, as shown in Appendix the estimate $\hat{x}$ of a random vector $x$ given some information (random vector) $I$, which minimizes the mean squared error $E x\left\{l l x-\hat{x} \|^{2} \mid I\right\}$ is precisely the conditional expectation $E\{x \mid I\}$ (expand the quadratic form and set to zero the derivative with respect to $\hat{x}$ ). Thus the estimator portion of the optimal controller is an optimal solution of the problem of estimating the state Xk assuming no control takes place, while the actuator portion is an optimal solution of the control problem assuming perfect state information prevails. This property, which shows that the two portions of the optimal controller can be designed independently as optimal solutions of an estimation and a control problem, has been called the separation theorem for


Figure 5.2.1 Structure of the optimal controller for the linear-quadratic problem. It consists of an estimator, which generates the conditional expectation $E\left\{X k \mid I_{k}\right\}$, and an actuator, which multiplies $E\left\{X k \mid I_{k}\right\}$ by the gain matrix $L k$
linear systems and quadratic cost and occupies a central position in modern automatic control theory.

Another interesting observation is that the optimal controller applies at each time $k$ the control that would be applied when faced with the deterministic problem of minimizing the cost-to-go

$$
x_{N}^{\prime} Q_{N} x_{N}+\sum_{i=k}^{\mathbf{N}-\mathbf{I}}\left(x_{i}^{\prime} Q_{i} x_{i}+u_{i}^{\prime} R_{i} u_{i}\right)
$$

and the input disturbances $W k, w_{k+1}, \ldots, W N-l$ and current state $X k$ were known and fixed at their conditional expected values, which are zero and $E\left\{x k \mid I_{k}\right\}$, respectively. This is another manifestation of the certainty equivalence principle, which was referred to in Section 4.1. A similar result holds in the case of correlated disturbances; see Exercise 5.1.

Inlp.lernent, atllon Aspects - Steady-State Controller
As explained in the perfect information case, the linear form of the actuator portion of the optimal policy is particularly attractive for implementation. In the imperfect information case, however, we need to implement an estimator that produces the conditional expectation

$$
\hat{x}_{k}=E\left\{x_{k} \mid I_{k}\right\},
$$

and this is not easy in general. Fortunately, in the important special case, where the disturbances $w_{k}, v_{k}$, and the initial state Xo are Gaussian random vectors, a convenient implementation of the estimator is possible by means of the Kalman filtering algorithm, which is developed in Appendix E. This algorithm is organized recursively so that to produce $\hat{x}_{k+1}$ at time $k+1$,
only the most recent measurement $Z k+l$ is needed, together with $U k$. In particular, we have for all $k=0, \ldots, N-1$,
$\hat{x}_{k+1}=A_{k} \hat{x}_{k}+B_{k} u_{k}+\Sigma_{k+1 \mid k+1} C_{k+1}^{\prime} N_{k+1}^{-1}\left(z_{k+1}-C_{k+1}\left(A_{k} \hat{x}_{k}+B_{k} u_{k}\right)\right)$,
and

$$
\hat{x}_{0}=E\left\{x_{0}\right\}+\Sigma_{0 \mid 0} C_{0}^{\prime} N_{0}^{-1}\left(z_{0}-C_{0} E\left\{x_{0}\right\}\right)
$$

where the matrices $\Sigma_{k \mid k}$ are precomputable and are given recursively by

$$
\Sigma_{k+1 \mid k+1}=\Sigma_{k+1 \mid k}-\Sigma_{k+1 \mid k} C_{k+1}^{\prime}\left(C_{k+1} \Sigma_{k+1 \mid k} C_{k+1}^{\prime}+N_{k+1}\right)^{-1} C_{k+1} \Sigma_{k+1 \mid k}
$$

$$
\Sigma_{k+1 \mid k}=A_{k} \Sigma_{k \mid k} A_{k}^{\prime}+M_{k}, \quad k=0,1, \ldots, N-1
$$

with

$$
\Sigma_{0 \mid 0}=S-S C b(\operatorname{CoSCb}+N O)-l \operatorname{COS}
$$

In these equations, $M k^{\prime} N k^{\prime}$ and $S$ are the covariance matrices of $W k, V k$, and $X o$, respectively, and we assume that $W k$ and $V k$ have zero mean; that is

$$
\begin{gathered}
E\left\{w_{k}\right\}=E\left\{v_{k}\right\}=0, \\
M_{k}=E\left\{w_{k} w_{k}^{\prime}\right\}, \quad N_{k}=E\left\{v_{k} v_{k}^{\prime}\right\}, \\
S=E\{(x o \quad E\{x o\})(x o-E\{X O\}) /\}^{\prime}
\end{gathered}
$$

In addition, the matrices $N k$ are assumed to be positive definite.
Consider now the case where the system and measurement equations, and the disturbance statistics are stationary. We can then drop subscripts from the system matrices. Assume that the pair $(A, B)$ is controllable and that the matrix $Q$ can be written as $Q=F^{\prime} F$, where $F$ is a matrix such that the pair $(A, F)$ is observable. By the theory of Section 4.1, if the horizon tends to infinity, the optimal controller tends to the steady-state policy

$$
\begin{equation*}
\mu^{*}\left(I_{k}\right)=L \hat{x}_{k} \tag{5.10}
\end{equation*}
$$

where

$$
\begin{equation*}
L=-(R+B I K B)-I B I K A, \tag{5.11}
\end{equation*}
$$

and $K$ is the unique positive semidefinite symmetric solution of the algebraic Riccati equation

$$
K=A^{\prime}(K-K B(R+B I K B)-I B I K) A+Q .
$$

By a similar argument, it can be shown (see Appendix E) that $\hat{x}_{k}$ can be generated in the limit as $k \rightarrow$ oo by a steady-state Kalman filtering algorithm:

$$
\begin{equation*}
\hat{x}_{k+1}=(A+B L) \hat{x}_{k}+\bar{\Sigma} C^{\prime} N^{-1}\left(z_{k+1}-C(A+B L) \hat{x}_{k}\right), \tag{5.12}
\end{equation*}
$$

where $\bar{\Sigma}$ is given by

$$
\bar{\Sigma}=\Sigma-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1} C \Sigma
$$

and $\Sigma$ is the unique positive semidefinite symmetric solution of the Riccati equation

$$
\Sigma=A\left(\Sigma-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1} C \Sigma\right) A^{\prime}+M
$$

The assumptions required for this are that the pair $(A, \mathrm{C})$ is observable and that the matrix $M$ can be written as $M=D D^{\prime}$, where $D$ is a matrix such that the pair $(A, D)$ is controllable. The steady-state controller of Eqs. (5.10)-(5.12) is particularly attractive for practical implementation. Furthermore, as shown in Appendix E, it results in a stable closed-loop system, under the preceding controllability and observability assumptions.

### 5.3 MINIMUM VARIANCE CONTROL OF LINEAR SYSTEMS

We have considered so far the control of linear systems in state variable form as in the previous section. However, linear systems are often modeled by means of an input-output equation, which is more economical in the number of parameters needed to describe the system dynamics. In this section we consider single-input, single-output, linear, time-invariant systems, and a special type of quadratic cost function. The resulting optimal policy is particularly simple and has found wide application. We first introduce some of the basic facts regarding linear systems in input-output form. Detailed discussions may be found in the books by Astrom and Wittenmark [AsW84]' [AsW90], Goodwin and Sin [GoS84], and Whittle [Whi63].

We consider a single-input single-output discrete-time linear system, which is specified by an equation of the form

$$
\begin{equation*}
Y k+a l Y k-l+\ldots+a m Y k-m=b O U k+b_{1} U k-l+\ldots+b m u k-m \tag{5.13}
\end{equation*}
$$

where $a i, b_{i}$ are given scalars. The scalar sequences $\{U k \mid k=0, \pm 1, \pm 2, \ldots\}$, $\{Y k \mid k=0, \pm 1, \pm 2, \ldots\}$ are viewed as the input and output of the system, respectively. Note that we allow time to extend to -00 as well as +00 ; this will be useful for describing generic properties of the system relating © stability. We will later revert to our usual convention of starting at time and proceeding forward.

It is convenient to describe this type of system by means of the backward shift operator, denoted $s$, which when operating on a sequence $\{X k \mid k=0, \quad \pm 2, \ldots\}$ shifts its index by one unit; that is,

$$
s\left(x_{k}\right)=x_{k-1}, \quad k=0, \pm 1, \pm 2, \ldots
$$

We denote by $s r$ the operator resulting from $r$ successive applications of $s$ :

$$
\begin{equation*}
s^{r}\left(x_{k}\right)=x_{k-r}, \quad k=0, \pm 1 \tag{5.14}
\end{equation*}
$$

We also write for simplicity $\operatorname{Sr} X k=X k-r$. The forward shift operator, denoted $s-l$, is the inverse of $s$ and is defined by

$$
s^{-1}\left(x_{k}\right)=x_{k+1}, \quad k=0, \pm 1, \pm 2, \ldots
$$

Thus the notation (5.14) holds for all integers $r$. We can form linear combinations of operators of the form $s r$. Thus, for example, the operator $\left(s+2 s^{2}\right)$ is defined by

$$
\left(s+2 s^{2}\right)\left(x_{k}\right)=x_{k-1}+2 x_{k-2}, \quad k=0, \pm 1, \pm 2, \ldots
$$

With this notation, Eq. (5.13) can be written as

$$
A(s) y_{k}=B(s) u_{k}
$$

where $A(s), B(s)$ are the operators

$$
\begin{aligned}
& A(s)=1+a_{1} s+\cdots+a_{m} s^{m} \\
& B(s)=b o+b_{l} s+., .+b m s^{m}
\end{aligned}
$$

Sometimes it is convenient to write the equation $A(S) Y k=B(S) U k$ as

$$
y_{k}=\frac{B(s)}{A(s)} u_{k}
$$

or

$$
\begin{aligned}
& A(s) \\
& \bar{B}(s)
\end{aligned}
$$

The meaning of both equations is that the sequences $\{Y k\}$ and $\left\{u_{k}\right\}$ are related via $A(S) Y k=B(S) U k^{\prime}$ There is a certain ambiguity here in that, for a fixed $\{U k\}$, the equation $A(S) Y k=B(S) U k$ has an infinite number of solutions in $\{Y k\}^{\prime}$ For example, the equation

$$
y_{k}+a y_{k-1}=u_{k}
$$

0
for $U k \equiv$ has as solutions all sequences of the form $Y k=f 3(-a) k$, where $\beta$ is any scalar; the solution becomes unique only after some boundary condition for the sequence $\{Y k\}$ is specified. As will be discussed shortly, however, for stable systems and for a bounded sequence $\left\{u_{k}\right\}$ there is a unique solution $\{Y k\}$ that is bounded. It is this solution that will be denoted by $(B(s) j A(s)) U k$ in what follows. The reader who is familiar with linear
dynamic system theory will note that $B(s) / A(s)$ can be viewed as a transfer function involving z -transforms.

We now introduce some terminology. When the sequences $\{Y k\}$ and $\{U k\}$ satisfy $A(S) Y k=B(S) U k$, we say that $Y k$ is obtained by passing $U k$ through the filter $B(s) / A(s)$. This comes from engineering terminology, where linear time-invariant systems are commonly referred to as filters. We also refer to the equation $A(S) Y k=B(S) U k$ as the filter $B(s) / A(s)$.

A filter $B(s) / A(s)$ is said to be stable if the polynomial $A(s)$ has all its (complex) roots strictly outside the unit circle of the complex plane; that is, $I p l>1$ for all complex $p$ satisfying $A(p)=0$. A stable filter $B(s) / A(s)$ has the following two properties:
(a) Every solution $\{Y k\}$ of

$$
A(S) Y k=0
$$

satisfies limk $-+=Y k=0$; that is, the output $Y k$ tends to zero if the input sequence $\{U k\}$ is identically zero.
(b) For every bounded sequence $\left\{\bar{u}_{k}\right\}$, the equation

$$
A(s) y_{k}=B(s) \bar{u}_{k}
$$

has a unique solution $\left\{\bar{y}_{k}\right\}$ within the class of bounded sequences. Furthermore, every solution $\{Y k\}$ (possibly unbounded) of the equation satisfies

$$
\lim _{k \rightarrow \infty}\left(y_{k}-\bar{y}_{k}\right)=0 .
$$

For example, consider the system

$$
Y k-0.5 Y k-l=U k .
$$

Given the bounded input sequence $\bar{u}_{k}=\{\ldots, 1,1,1, \ldots\}$, the set of all solutions is given by

$$
Y^{k}=2+\frac{\beta}{2^{k}}
$$

where ( 3 is a scalar, but of these the only bounded solution is $\bar{y}_{k}=$ $\{\ldots, 2,2,2, \ldots\}$. The solution $\left\{\bar{y}_{k}\right\}$ can thus be naturally associated with the input sequence $\{U k\}$; it is also known as the forced response of the system due to the input $\{U k\}^{\prime}$

## ARMAX Models - Reduction to State Space Form

We now consider a linear system with output $Y k$, which is driven by two inputs: a random noise input Ek, and a control input $U k$. It has the form

$$
\begin{align*}
Y k+a l Y k-l+\ldots+a m Y k-m= & b l U k-l+\ldots+b m u k-m \\
& +\epsilon_{k}+\text { ClEk- } 1+\ldots+C m E k-m, \tag{5.15}
\end{align*}
$$

and it is known as an ARMAX model (AutoRegressive, Moving Average, with eXogenous input). We assume throughout that the random variables $\epsilon_{k}$ are mutually independent. We can write the model in the shorthand form

$$
A(s) y_{k}=B(s) u_{k}+C(s) \epsilon_{k},
$$

where the polynomials $A(s), B(s)$, and $C(s)$ are given by

$$
\begin{aligned}
& A(s)=1+a l S+\ldots+a m s^{m} \\
& B(s)=b l s+\ldots+b m s^{m}, \\
& \mathrm{C}(s)=1+\mathrm{Cl} S+\ldots+\mathrm{C}_{\mathrm{m}} s m .
\end{aligned}
$$

The ARMAX model is very common and its derivation is outlined in Appendix F, where it is shown that without loss of generality we can assume that $\mathrm{C}(s)$ has no roots strictly inside the unit circle. For much of the analysis in subsequent sections, it will be necessary to exclude the critical case where $\mathrm{C}(s)$ has roots on the unit circle and assume that $\mathrm{C}(s)$ has all its roots strictly outside the unit circle. This assumption is usually satisfied in practice.

In several situations, analysis and algorithms relating to the ARMAX model are greatly simplified if $C(s)=1$ so that the noise terms $C(S) E k=\epsilon_{k}$ are independent. However, this is typically an unrealistic assumption. To emphasize this point and see how easily the noise can be correlated, suppose that we have a first-order system

$$
x_{k+1}=a x_{k}+w_{k}
$$

where we observe

$$
y_{k}=x_{k}+v_{k}
$$

Then

$$
\begin{aligned}
Y k+l & =X k+l+{ }^{\prime} ' U k+l \\
& =a X k+W k+' U k+l \\
& =a(Y k-' U k)+W k+{ }^{\prime} U k+l,
\end{aligned}
$$

so finally

$$
y_{k+1}=a y_{k}+v_{k+1}-a v_{k}+w_{k}
$$

However, the noise sequence $\left\{' U k+l-a^{\prime} U k+W k\right\}$ is correlated even if $\{' U k\}$ and $\{W k\}$ are individually and mutually independent.

The ARMAX model (5.15) can be put into state space form. The process is based on state augmentation and can perhaps be best understood in terms of an example. Consider the system

$$
\begin{equation*}
y_{k}+a_{1} y_{k-1}+a_{2} y_{k-2}=b_{1} u_{k-1}+b_{2} u_{k-2}+\epsilon_{k}+c_{1} \epsilon_{k-1} \tag{5.16}
\end{equation*}
$$

We have

By setting

$$
\begin{gathered}
X k=\left(\begin{array}{c}
y_{k} \\
y_{k-1} \\
U_{k-k} \\
\epsilon_{k}
\end{array}\right), \quad\left(\begin{array}{c}
\epsilon_{k+1} \\
0 \\
0 \\
\epsilon_{k+1}
\end{array}\right) \\
A=\left(\begin{array}{cccc}
-a_{1} & -a_{2} & b_{2} & c_{1} \\
1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0
\end{array}\right), \quad B=\left(\begin{array}{c}
b_{1} \\
0 \\
1 \\
0
\end{array}\right)
\end{gathered}
$$

we can write Eq. (5.17) as

$$
X k+l=A X k+B U k+w_{k}
$$

where $\left\{w_{k}\right\}$ is a stationary independent process. We have arrived at this state space model through state augmentation. Notice that the state $X k$ includes $\epsilon_{k}$. Thus if the controller is assumed to know at time $k$ only the present and past outputs $Y k, y_{k-1}, \ldots$, and past controls $U k-I, U k-2, \ldots$ (but not $\epsilon_{k} \epsilon_{k-1}, \ldots$ ), we are faced with a model of imperfect state information. If $\mathrm{Cl}=0$ in Eq. (5.16) then the state space model can be simplified so that

$$
X k=\left(\begin{array}{c}
y_{k} \\
y_{k-1} \\
u_{k=1}
\end{array}\right)
$$

in which case we have perfect state information. More generally, we have perfect state information in the ARMAX model (5.15) if $\mathrm{b}_{1} \neq 0$ and $\mathrm{CI}=$ $\mathrm{C} 2=\ldots=\mathrm{c}_{\mathrm{m}}=0$.

## Minimum Variance Control: Perfect State Information Case

Consider the perfect state information case of the ARMAX model (5.15):

$$
y_{k}+a_{1} y_{k-1}+\cdots+a_{m} y_{k-m}=b_{1} u_{k-1}+\cdots+b_{m} u_{k-m}+\epsilon_{k}
$$

where $\mathrm{b}_{\wedge} \neq 0$. A problem of interest, known as the minimum variance control problem, is to select $U k$ as a function of the present and past outputs $Y k, y_{k-1}, \ldots$, as well as the past controls $U k-I, U k-2, \ldots$, so as to minimize the cost

$$
E\left\{\sum_{k=1}^{N}\left(y_{k}\right)^{2}\right\}
$$

There are no constraints on $U k$. By transforming the system to state space form, we see that this problem can be reduced to a perfect state information linear-quadratic problem where the state $X k$ is

$$
\left(y_{k}, y_{k-1}, \ldots, y_{k-m+1}, u_{k-1}, \ldots, u_{k-m+1}\right)^{\prime}
$$

The problem is of the same nature as the linear-quadratic problem of Section 4.1 except that the corresponding matrices $R k$ in the quadratic cost function are zero here. Nonetheless, in Section 4.1 we used the invertibility of $R k$ only to ensure that various matrices in the optimal policy and the Riccati equation are invertible. If invertibility of these matrices can be guaranteed by other means, the same analysis applies even if $R k$ is positive semidefinite. This is indeed the case here. An analysis analogous to the one of Section 4.1 shows that the optimal control $u_{k}^{*}$ at time $k$ (given $Y k, Y k-I, \cdots, Y k-m+1$ and $\left.u_{k-1}, \ldots, U k-m+l\right)$ is the same as the one that would be applied if all future disturbances $\epsilon_{k+1}, \ldots, \epsilon_{N}$ were set equal to zero, their expected value (certainty equivalence). It follows that

$$
\mu_{k}^{*}\left(y_{k}, \ldots, y_{k-m+1}, U k-I, \cdots, U k-m+l\right)=\underline{1} \underline{\mathrm{bl}}(a l Y k+\ldots+a m Y k-m+l
$$

$$
\left.-b 2 U k-1 \_\ldots-b m U k-m+l\right)
$$

and $\left\{u_{k}^{*}\right\}$ is generated via the equation

$$
b_{1} u_{k}^{*}+b_{2} u_{k-1}^{*}+\cdots+b_{m} u_{k-m+1}^{*}=a_{1} y_{k}+a_{2} y_{k-1}+\cdots+a_{m} y_{k-m+1}
$$

In other words, $\left\{u_{k}^{*}\right\}$ is generated by passing $\{Y k\}$ through the linear filter $\underline{A(s) j B(s)}$, where

$$
\begin{gathered}
\bar{A}(s)=a l+a 2 S+\ldots+a m s^{m_{-} l}=\mathrm{s}-1(A(s) \\
\bar{B}(s)=\mathrm{b}_{1}+b 2 S+\ldots+b m s^{m_{-} l}=s-l B(s)
\end{gathered}
$$

as shown in Fig. 5.3.1. The resulting closed-loop system is

$$
\begin{equation*}
Y k=\epsilon_{k} \tag{5.18}
\end{equation*}
$$

and the associated cost is

$$
N E\left\{\left(\epsilon_{k}\right)^{2}\right\}
$$

Notice that the optimal policy, called minimum variance control law, is time invariant and does not depend on the horizon $N$.

Whereas the optimal closed-loop system as given by Eq. (5.18) is clearly stable, we can anticipate serious difficulties if the filter $\bar{A}(s) j \underline{B}(s)$ in the feedback loop is unstable. For if $\underline{B}(s)$ has some roots inside the unit circle, then the sequence $\{U k\}$ will tend to be unbounded. This is illustrated by the following example.


Figure 5.3.1 Minimum variance control with perfect state information. Structure of the optimal closed-loop system, where $A(8)=1+a l S+\ldots+a m .8 \mathrm{~m} ., B(s)=$ $b_{1} 8+\ldots+b m .8 m$., $\bar{A}(s)=S-1(A(8)-1)$, and $\vec{B}(s)=8-1 B(s)$.

工یun2pic 5.3.1 (An Optimal but Unstable Controller)
Consider the system

$$
Y k+Y k-1=U k-1-2 U k-2+\epsilon_{k} .
$$

The optimal policy is

$$
u_{k}=2 u_{k-1}+y_{k}
$$

and the optimal closed-loop system is

$$
y_{k}=\epsilon_{k}
$$

which is a stable system. On the other hand, the last two equations yield

$$
u_{k}=2 u_{k-1}+\epsilon_{k}
$$

Thus, $U k$ is generated by an unstable system, and in fact it is given by

$$
u_{k}=\sum_{n=0}^{k} 2^{n} \epsilon_{k-n}
$$

Therefore, even though the output $Y k$ stays bounded, the control $U k$ typically becomes unbounded.

For another view of the same difficulty, suppose that the coefficients $\mathrm{b}_{1}, \ldots, b_{m}$ of $\underline{B}(s)$ are slightly different from the ones of the true system.

We will show that if the feedback filter $\bar{A}(s) / \bar{B}(s)$ is unstable, then the closed-loop system is also unstable in the sense that both $u_{k}$ and $Y k$ become unbounded with probability one.

Assume that the system is governed by

$$
\begin{equation*}
A^{0}(s) y_{k}=B^{0}(s) u_{k}+\epsilon_{k} \tag{5.19}
\end{equation*}
$$

while the policy is calculated under the assumption that the system model is

$$
A(S) Y k=B(s) u_{k}+\epsilon_{k},
$$

where the coefficients of $A(s)$ and $E(s)$ differ slightly from those of $A O(s)$, $E O(s)$. Define $\bar{A}^{0}(s), \bar{B} \mathbf{O}(s)$ by

$$
\begin{gathered}
1+s \bar{A}^{0}(s)=A O(s), \\
s \bar{B}^{\mathbf{O}}(s)=E O(s) .
\end{gathered}
$$

Note that $\bar{A}^{0}(s)=\bar{A}(s)$ and $\bar{B}^{\mathbf{O}}(s)=\bar{B}(s)$ if $A O(s)=A(s), B O(s)=B(s)$. By multiplying Eq. (5.19) with $\bar{B}(s)$ and by using the relation defining the optimal policy

$$
\bar{B}(s) u_{k}=\bar{A}(s) y_{k}
$$

we obtain

$$
\underline{B(s) A O(s) Y k}=B O(s) A(s) Y k+\underline{B(S) E k^{\prime}}
$$

If the coefficients of $\bar{A}^{0}(s)$ and $\bar{B}^{\mathbf{O}}(s)$ are close to those of $\bar{A}(s), B(s)$, then the roots of the polynomial

$$
\bar{B}(s)+s\left(\bar{B}(s) \bar{A}^{0}(s)-\bar{B}^{0}(s) \bar{A}(s)\right)
$$

are close to the roots of $\underline{B}(s)$. Thus the closed-loop system is stable only if the roots of $\underline{B(s)}$ are outside the unit circle, or equivalently, if and only ifthe filter $\bar{A}(s) \dot{j} \underline{B}(s)$ is stable. If our model is exact, the closed-loop system will be stable due to what is commonly referred to as a pole-zero cancellation. However, the slightest modeling discrepancy will induce instability.

The conclusion from the preceding analysis is that the minimum variance control law is advisable only if it can be realized through a stable filter $[B(s)$ has roots outside the unit circle]. Even if $\underline{B(s)}$ has its roots outside the unit circle, but some of these roots are near the unit circle, the performance of the minimum variance policy can be very sensitive to variations in the parameters of the polynomials $A(s)$ and $E(s)$. One way to overcome this sensitivity is to change the cost to

$$
E\left\{\sum_{k=1}^{N}\left(\left(y_{k}\right)^{2}+R\left(u_{k-1}\right)^{2}\right)\right\}
$$

where $R$ is some positive parameter. This requires solution via the Riccati equation as in Section 4.1. For a detailed derivation, see Astrom [Ast83].

In some problems, the system equation includes an additional external input sequence $\{V k\}$, the values of which can be measured by the controller as they occur. In particular, consider the scalar system

$$
\begin{aligned}
Y k+a l Y k-1+\ldots+a m Y k-m= & \mathrm{b}_{1} u_{k-1}+\ldots+b m U k-m \\
& +\mathrm{d}_{1} v_{k-1}+\ldots+d m V k-m+\epsilon_{k}
\end{aligned}
$$

where each value $V k$ becomes known to the controller without error at time k. The minimum variance controller then takes the form

$$
\begin{gathered}
\mu_{k}^{*}\left(y_{k}, \ldots, Y k-m+I, U k-I, \cdots, U k-m+l, V k, \cdots, V k-m+l\right) \\
=\frac{1}{\mathrm{bl}}\left(a l Y k+\ldots+a m Y k-m+1-\mathrm{d}_{1} V k-\ldots d m V k-m+1\right. \\
\left.-b 2 U_{k-I^{\prime \prime}}-b m U k-m+I\right)
\end{gathered}
$$

and the optimal controls $\left\{u_{k}^{*}\right\}$ are generated by

$$
\bar{B}(s) u_{k}^{*}=\bar{A}(s) y_{k}-\bar{D}(s) v_{k},
$$

where

$$
\begin{array}{ll}
\bar{A}(s)=\mathrm{a} 1+a 2 S+" ' & +a_{m} s^{m-1} \\
\bar{B}(s)=\mathrm{b} 1+b 2 S+ & +b m s^{m_{-}} l \\
\underline{D(s)}=\mathrm{d}_{1}+d 2 S+ & +d m s^{m_{-}} l
\end{array}
$$

The closed-loop system is again $Y k=\epsilon_{k}$, but for practical purposes it is stable only if $B(s)$ has its roots outside the unit circle. The process whereby external inputs are measured and used for control is commonly referred to as feedforward control.

## Imperfect State Information Case

Consider now the general ARMAX model

$$
\begin{aligned}
Y k+a l Y k-1+\ldots+a m Y k-m= & b_{M} u_{k-M}+\ldots+b m U k-m \\
& +\epsilon_{k}+\text { CIEk-1 }
\end{aligned}
$$

or, equivalently,

$$
A(s) y_{k}=B(s) u_{k}+C(s) \epsilon_{k}
$$

where

$$
\begin{aligned}
& A(s)=1+a l s+\ldots,+a m s^{m} \\
& B(s)=b M S^{M}+\ldots+b m s^{m} \\
& C(s)=l+C I S+\cdots+c m s^{m}
\end{aligned}
$$

We assume the following;
(1) $b M \neq 0$ and $1 \leq N l \leq m$.
(2) $\{\mathrm{Ek}\}$ is a zero mean, independent, stationary process.
(3) The polynomial $C(s)$ has all its roots outside the unit circle. (As explained in Appendix F, this assumption is not overly restrictive.)
The controller knows at each time $k$ the past inputs and outputs. Thus the information vector at time $k$ is

$$
I_{k}=\left(y_{k}, y_{k-1}, \ldots, y_{-m+1}, u_{k-1}, u_{k-2}, \ldots, u_{-m+M}\right)
$$

(We include in the information vector the control inputs U-I, $\ldots, U-m+M$. If control starts at time 0 , these inputs will be zero.) There are no constraints on $U k$. The problem is to find a policy $\left\{\mu_{0}\left(I_{0}\right), \ldots, \mu_{N-1}(\right.$ IN-I) $\}$ that minimizes

$$
E\left\{\sum_{k=1}^{N}\left(y_{k}\right)^{2}\right\}
$$

By using state augmentation, we can cast this problem into the framework of the linear-quadratic problem of Section 5.2. The corresponding linear system in state space format involves a state $X_{k}$ given by

$$
x_{k}=\left(y_{k+M-1}, \ldots, y_{k+M-m}, u_{k-1}, \ldots, u_{k+M-m}, \epsilon_{k+M-1}, \ldots, \epsilon_{k+M-m}\right)
$$

Because $Y k+M-I$,.. "Yk+l and $\epsilon_{k+M-1}, \ldots$, Ek+M-m are unknown to the controller, we are faced with a problem of imperfect state information.

An analysis analogous to the one of Section 5.2 shows that certainty equivalence holds; that is, the optimal control $u_{k}^{*}$ at time $k$ given $I_{k}$ is the same as the one that would be applied in the deterministic problem where the current state

$$
x_{k}=\left(y_{k+M-1}, \ldots, y_{k+M-m}, u_{k-1}, \ldots, u_{k+M-m}, \epsilon_{k+M-1}, \ldots, \epsilon_{k+M-m}\right)
$$

is set equal to its conditional expected value given $I_{k}$, and the future disturbances $E k+M, \ldots, \epsilon_{N}$ are set equal to zero (their expected value).

Thus the optimal control $u_{k}^{*}=\mu_{k}^{*}\left(I_{k}\right)$ is obtained by solving for $U k$ the equation
$E\left\{Y k+M \mid U k, I_{k}\right\}=E\{Y k+M \mid Y k, Y k-l, " \prime \quad \mathrm{Y}-\mathrm{m}+1, U k, U k-I, \cdots, U-m+M\}$ $=0$.

This leads to the problem of calculating $E\{Y k+M 11 k, U k\}$, known as the forecasting or prediction problem, which is important in its own right. We first treat the easier case where there is no delay $(M=1)$ and then discuss the more general case where the delay can be positive.

## Forecasting for ARMAX Models - No $\quad(M=1)$

Assume that $M=1$. We would like to generate an equation for the forecast $E\left\{Y k+l \backslash I_{k}, U k\right\}$, and then determine the optimal control $u_{k}^{*}=\mu_{k}^{*}\left(I_{k}\right)$ by setting this forecast to zero. Let us introduce an auxiliary sequence $\{Z k\}$ via the equation

$$
Z k=Y k-\epsilon_{k}
$$

A key fact is that, since $\left\{\epsilon_{k}\right\}$ is an independent, zero-mean sequence, we have

$$
E\left\{Z k+l \mid I_{k}, U k\right\}=E\{Y k+l 11 k, U k\}
$$

We can thus obtain the desired forecast of $Y k+l$ by forecasting $Z k+l$ instead. We can then obtain the optimal control $u_{k}^{*}$ by setting $E\left\{Z k+1, l k, u_{k}^{*}\right\} \quad 0$.

By using the definition $Z k=Y k-\epsilon_{k}$ to express $Y k$ in terms of $Z k$ in the ARMAX model equation for $M=1$, we obtain

$$
\begin{equation*}
Z k+l+C l Z k+\ldots+C m Z k-m+1=\mathrm{bl} U k+\ldots+b m U k-m+l+W k \tag{5.20}
\end{equation*}
$$

where

$$
W k=\left(\mathrm{Cl}-a_{1}\right) y_{k}+\ldots+\left(c_{m}-a m\right) Y k-m+l .
$$

We note that $W k$ is perfectly observable by the controller; however, the scalars $Z k, \ldots, Z k-m+1$ are not known to the controller because the initial conditions $Z 0, \ldots, Z 1-m$ of the system (5.20) are unknown. Nonetheless, the system (5.20) is stable, since the roots of the polynomial $C(s)$ have been assumed to be outside the unit circle. As a result, the initial conditions do not matter asymptotically. In other words, if we generate a sequence $\left\{\hat{y}_{k}\right\}$ using the system (5.20) and zero initial conditions, Le.,

$$
\hat{y}_{k+1}+c_{1} \hat{y}_{k}+\cdots+c_{m} \hat{y}_{k-m+1}=b_{1} u_{k}+\cdots+b_{m} u_{k-m+1}+w_{k}
$$

with

$$
\hat{y}_{0}=0, \quad \hat{y}_{-1}=0, \quad \hat{y}_{1-m}=0,
$$

then we will have

$$
\lim _{k \rightarrow \infty}\left(\hat{y}_{k}-Z k\right)=0
$$

Thus, $\hat{y}_{k+1}$ is an asymptotically accurate approximation to the optimal forecast $E\left\{Y k+l \mid I_{k},{ }^{\prime} U k\right\}$.

## Minimum Variance Control: Imperfect State Information and No Delay

Based on the earlier discussion, an asymptotically accurate approximation to the minimum variance policy is obtained by setting $U k$ to the value that makes $\hat{y}_{k+1}=0$; that is, by solving for $U k$ the equation

$$
\hat{y}_{k+1}+c_{1} \hat{y}_{k}+\cdots+c_{m} \hat{y}_{k-m+1}=b_{1} u_{k}+\cdots+b_{m} u_{k-m+1}+w_{k} .
$$

If this policy is followed, however, the earlier forecasts $\hat{y}_{k}, \ldots$, will be equal to zero. Thus the (approximate) minimum variance policy is given by

$$
\begin{aligned}
U k= & \frac{1}{\mathrm{~b}_{1}}\left(W k-b_{2} u k-l-\ldots-b m U k-m+l\right) \\
& =\frac{1}{\mathrm{~b}_{1}^{\prime}}\left((a l-C l) Y k+\ldots+\left(a_{m}-C m\right) Y k-m-l\right. \\
& \left.-b_{2} u k-l-\ldots \quad b m U k-m+l\right)^{\prime}
\end{aligned}
$$

By substituting this policy in the ARMAX model

$$
\begin{aligned}
Y k+l+a l Y k+\ldots+a m Y k-m+l= & \mathrm{bl} U k+\ldots+b m U k-m+l \\
& +\epsilon_{k+1}+c_{1} \epsilon_{k}+\ldots+c_{m} \epsilon_{k-m+1}
\end{aligned}
$$

we see that the closed-loop system becomes

$$
y_{k+1}-\epsilon_{k+1}+c_{1}\left(y_{k}-\epsilon_{k}\right)+\cdots+c_{m}\left(y_{k-m+1}-\epsilon_{k-m+1}\right)=0
$$

or equivalently $C(S)\left(Y k-\epsilon_{k}\right)=0$. Since $C(s)$ has its roots outside the unit circle, this is a stable system, and we have

$$
Y k=\epsilon_{k}+\gamma(k)
$$

where $\gamma(k) \rightarrow 0$ as $k \rightarrow 00$.

## Forecasting: The General Case

Consider now the general case where the delay $M$ can be greater than 1 . The forecasting problem can still be nicely solved by using a certain trick to transform the ARMAX equation into a more convenient form. To this end, we first obtain polynomials $F(s)$ and $G(s)$ of the form

$$
\begin{aligned}
& F(s)=1+f_{1} s+\ldots+f_{M-1} s^{M-1} \\
& G(s)=g o+g l S+\ldots+\mathrm{gm}_{-1} \mathrm{sm}_{-}{ }^{1}
\end{aligned}
$$

which satisfy

$$
\begin{equation*}
C(s)=A(s) F(s)+s M G(s) . \tag{5.21}
\end{equation*}
$$

The coefficients of $F(S)$ and $G(S)$ are uniquely determined from those of $C(S)$ and $A(S)$ by equating coefficients of both sides of the relation

$$
\begin{aligned}
1+C l S+\ldots+c_{m} s^{m}= & \left(1+a l S+\ldots+a_{m} s^{m}\right)\left(1+f_{1} s+\ldots+f M-l S^{\mathrm{M}-1}\right) \\
& +\operatorname{sM}\left(g O+g l S+\ldots+\mathrm{gm}_{-} 1 \mathrm{Sm}-1\right)
\end{aligned}
$$

## Example 5.3.2

Let $\mathrm{m}=3$ and $M=2$. Then the preceding equation takes the form

$$
1+c_{1} s+c_{2} s^{2}+c_{3} s^{3}=\left(1+a_{1} s+a_{2} s^{2}+a_{3} s^{3}\right)\left(1+f_{1} s\right)+s^{2}\left(g_{0}+g_{1} s+g_{2} s^{2}\right)
$$

and by equating coefficients we have

$$
c_{1}=a_{1}+f_{1}, \quad c_{2}=a_{2}+a_{1} f_{1}+g_{0}, \quad c_{3}=a_{3}+a_{2} f_{1}+g_{1}, \quad a_{3} f_{1}+g_{2}=0
$$

from which $f_{1}, g o, g l$, and $g 2$ are uniquely determined.
The ARMAX model can be written as

$$
\begin{equation*}
A(s) y_{k+M}=\bar{B}(s) u_{k}+C(s) \epsilon_{k+M}, \tag{5.22}
\end{equation*}
$$

where

$$
\bar{B}(s)=s-M B(s)=b M+b M+{ }_{l} S+\ldots+b m s m-M
$$

lVIultiplying both sides of Eq. (5.22) with $F(s)$, we have

$$
F(s) A(s) Y k+M=\underline{F(s) B(s) U k}+F(s) C(S) c k+M,
$$

and using Eq. (5.21) to express $F(s) A(s)$ as $C(s)-s M G(s)$, we obtain

$$
(C(s)-s M G(s)) Y k+M=F(s) \bar{B}(s) u_{k}+F(s) C(S) c k+M,
$$

or equivalently

$$
\begin{equation*}
C(s)(Y k+M-F(S) c k+M)=\underline{F(s) B(s) U k}+G(S) Y k^{\prime} \tag{5.23}
\end{equation*}
$$

Let us now introduce the auxiliary sequence $\{Z k\}$ via the equation $z_{k+M}=y_{k+M}-F(s) \epsilon_{k+M}=y_{k+M}-\epsilon_{k+M}-f_{1} \epsilon_{k+M-1}-\cdots-f_{M-1} \epsilon_{k+1}$.

Note that when $M=1$, we have $F(s)=1$ and $Z k=Y k-\epsilon_{k}$, so $\{Z k\}$ is the same sequence as the one introduced earlier for the case of no delay. Again, since $\{c k\}$ is an independent, zero-mean sequence, by taking expectations in the definition $Z k+M=Y k+M-F(S) c k+M$, we obtain

$$
E\left\{z_{k+M} \mid I_{k}, u_{k}\right\}=E\left\{y_{k+M} \mid I_{k}, u_{k}\right\}
$$

and we can obtain the desired forecast of $Y k+M$ by forecasting $Z k+M$ in its place. Furthermore, by Eq. (5.23), $z_{k+M}$ is written as

$$
C(s) z_{k+M}=w_{k}
$$

or

$$
\begin{equation*}
z_{k+M}+c_{1} z_{k+M-1}+\cdots+c_{m} z_{k+M-m}=w_{k} \tag{5.24}
\end{equation*}
$$

where

$$
\begin{equation*}
W k=F(s) \bar{B}(s) u_{k}+G(S) Y k^{\prime} \tag{5.25}
\end{equation*}
$$

Since the scalar $W k$ of Eq. (5.25) is available at time $k$ (i.e., it is determined from $I_{k}$ and $U k$ ), the system (5.24) can serve as a basis for forecasting $Z k+M$. We would be able to predict exactly $Z k+M$ and use it as a forecast of $Y k+M$ if we knew appropriate initial conditions with which to start the equation (5.24) that generates it. We don't know such initial conditions, but because this equation represents a stable system, the choice of initial conditions does not matter asymptotically, as we proceed to explain more formally.

We consider the sequence $\hat{y}_{k+M}$ generated by

$$
\hat{y}_{k+M}+C I Y k+M-l+\ldots+C m Y k+M-m=W k
$$

with initial condition

$$
\begin{equation*}
\hat{y}_{M-1}=\hat{y}_{M-2}=\ldots=\hat{y}_{M-m}=0, \tag{5.26}
\end{equation*}
$$

and we claim that the forecast $E\left\{Z k+M 11_{k}\right\}$ can be approximated by $\hat{y}_{k+M}$. To see this, note that from Eqs. (5.24) to (5.26) we have

$$
Z k+M=\hat{y}_{k+M}+\left(\gamma_{1}(k) z_{M-1}+\ldots+i m(k) Z M-m\right)
$$

and

$$
E\{Z k+M \mid I k, u k\}=\hat{y}_{k+M}+\sum_{\mathrm{i}=1}^{\mathrm{m}} i i(k) E\{Z M-i \mid \quad U k\},
$$

where $\gamma_{\mathbf{1}}(k), \ldots, \operatorname{im}(k)$ are appropriate scalars depending on $k$. Since $C(s)$ has all its roots outside the unit circle, we have (compare with the discussion on stability earlier in this section)

$$
\lim _{k \rightarrow \infty} \gamma_{1}(k)=\lim _{k \rightarrow \infty} \gamma_{2}(k)=\ldots=\lim _{k \rightarrow \infty} \gamma_{m}(k)=0
$$

It follows that, for large values of $k$,

$$
\hat{y}_{k+M} \simeq E\left\{Z k+M \mid I_{k}, U k\right\}=E\left\{Y k+M \mid I_{k}, U k\right\} .
$$

(More precisely, we have $\left|\hat{y}_{k+M}-E\left\{Y k+M \mid I_{k}, u_{k}\right\}\right| \rightarrow 0$ as $k \rightarrow o 0$, where the convergence is in the mean-square sense.)

In conclusion, an asymptotically accurate approximation to the optimal forecast $E\{Y k+M 11 k, U k\}$ is given by $\hat{y}_{k+M}$ and is generated by the equation

$$
\begin{equation*}
\hat{y}_{k+M}+C I Y k+M-l+\ldots+C m Y k+M-m=F(s) \bar{B}(s) u_{k}+G(S) Y k \tag{5.27}
\end{equation*}
$$

with the initial condition

$$
\begin{equation*}
\hat{y}_{M-1}=\hat{y}_{M-2}=\ldots=\hat{y}_{M-m}=0 . \tag{5.28}
\end{equation*}
$$

## Minimum Variance Control: The General Case

Based on the earlier discussion, the minimum variance policy is obtained by solving for $U k$ the equation $E\left\{Y k+M \mid I_{k}, U k\right\}=0$. Thus an asymptotically accurate approximation is obtained by setting $U k$ to the value that makes $Y k+M=0$, that is, by solving for $U k$ the equation [ef. Eqs. (5.27) and (5.28)]

$$
\underline{F(s) B(s) U k}+G(S) Y k=c_{1} \hat{y}_{k+M-1}+\ldots+C m Y k+M-m
$$

If this policy is followed, however, the earlier forecasts $Y k+M_{-1}, \ldots, \hat{y}_{k+M-m}$ will be equal to zero. Thus the (approximate) minimum variance policy is given by

$$
\begin{equation*}
\underline{F(s) B(s) U k}+G(S) Y k=0 ; \tag{5.29}
\end{equation*}
$$

that is, $u_{k}^{*}$ is generated by passing $Y k$ through the linear filter

$$
-G(s) / F(s) B(s),
$$

as shown in Fig. 5.3.2.


Figure 5.3.2 Minimum variance control with imperfect state information. Structure of the optimal closed-loop system.

From Eqs. (5.23) and (5.29), we obtain the equation for the closedloop system

$$
C(S)(Y k+M \quad F(S) E k+M)=0
$$

Since $C(s)$ has its roots outside the unit circle, we obtain

$$
Y k+M=F(S) E k+M+\gamma(k)
$$

where $\gamma(k) \rightarrow 0$ as $k \rightarrow$ oo. So asymptotically, the closed-loop system takes the form

$$
Y k=\mathrm{Ek}+f_{1} \epsilon_{k-1}+\ldots+f M-I E k-M+l .
$$

Let us consider now the stability properties of the closed-loop system when the true system parameters differ slightly from those of the assumed model. Let the true system be described by

$$
\begin{equation*}
A^{0}(s) y_{k}=s^{M} \bar{B}^{0}(s) u_{k}+C^{0}(s) \epsilon_{k} \tag{5.30}
\end{equation*}
$$

while $U k$ is given by the minimum variance policy

$$
\begin{equation*}
\underline{F(s) B(s) U k}+G(S) Y k=0 \tag{5.31}
\end{equation*}
$$

where

$$
C(s)=A(s) F(s)+s M G(s)
$$

Operating on Eq. (5.30) with $\underline{F(s) B(s)}$ and using Eq. (5.31), we obtain

$$
F(s) \bar{B}(s) A O(s) Y k=-s M \bar{B}(s) G(S) Y k+F(s) \overline{B( }(s) C O(S) E k^{\prime}
$$

Combining the last two equations and collecting terms, we have

$$
\{F(s) B(s) A O(s)+(C(s)-\underline{A(s) F(s)) B o(s)}\} Y k=\underline{F(s) B(s) C O(S) E k}
$$

or

$$
\left\{\bar{B}^{0}(s) C(s)+F(s)\left(\bar{B}(s) A^{0}(s)-A(s) \bar{B}^{0}(s)\right)\right\} y_{k}=F(s) \bar{B}(s) C^{0}(s) \epsilon_{k}
$$

If the coefficients of $A O(s), \bar{B}^{\mathbf{O}}(s)$, and $C O(s)$ are near those of $A(s), \bar{B}(s)$, and $C(s)$, then the poles of the closed-loop system will be near the roots of $\underline{B(s) C(s)}$. Thus the closed-loop system will be in effect stable only if the roots of $\underline{B}(s)$ are strictly outside the unit circle, similar to the perfect state information case examined earlier.

### 5.4 SUFFICIENT STATISTICS

The main difficulty with the DP algorithm for imperfect state information problems is that it is carried out over a state space of expanding dimension. As a new measurement is added at each stage $k$, the dimension of the state (the information vector $I_{k}$ ) increases accordingly. This motivates an effort to reduce the data that are truly necessary for control purposes. In other words, it is of interest to look for quantities known as sufficient statistics,
which ideally would be of smaller dimension than $I_{k}$ and yet summarize all the essential content of $I_{k}$ as far as control is concerned.

Consider the DP algorithm (5.4) and (5.5), restated here for convenience:

$$
\begin{align*}
J_{N-1}\left(I_{N-1}\right)=\min _{\prime U N-I E U N-l}[ & \underset{x_{N-1}, w_{N-1}}{E}\{g N(h v-I(X N-I, U N-I, W N-I)) \\
& \left.\left.+g N-I(X N-I, U N-I, W N-I) \mid I_{N-1}, u_{N-1}\right\}\right] \tag{5.32}
\end{align*}
$$

$J k(I k)=\min _{u k E U k}\left[\frac{E}{X k, W k, Z k+l}\left\{\operatorname{lgk}(X k, U k, W k)+J k+1\left(I_{k}, z_{k+1}, U k\right) \mid I_{k}, U k\right\}\right]$.
Suppose that we can find a function $\operatorname{Sk}(I k)$ of the information vector, such that a minimizing control in Eqs. (5.32) and (5.33) depends on $I k$ via $S_{k}\left(I_{k}\right)$. By this we mean that the minimization in the right-hand side of the DP algorithm (5.32) and (5.33) can be written in terms of some function $H k$ as

$$
\min _{u_{k} \in U_{k}} H_{k}\left(S_{k}\left(I_{k}\right), u_{k}\right)
$$

Such a function $S k$ is called a sufficient statistic. Its salient feature is that an optimal policy obtained by the preceding minimization can be written as

$$
\mu_{k}^{*}\left(I_{k}\right)=\bar{\mu}_{k}\left(S_{k}\left(I_{k}\right)\right)
$$

where $\bar{\mu}_{k}$ is an appropriate function. Thus, if the sufficient statistic is characterized by a set of fewer numbers than the information vector $I_{k}$, it may be easier to implement the policy in the form $U k=\bar{\mu}_{k}\left(S_{k}\left(I_{k}\right)\right)$ and take advantage of the resulting data reduction.

### 5.4.1 The Conditional State Distribution

There are many different functions that can serve as sufficient statistics. The identity function $S_{k}\left(I_{k}\right)=I k$ is certainly one of them. In this section, we will derive another important sufficient statistic: the conditional probability distribution $P X k l l k$ of the state $X k$, given the information vector $I_{k}$. An extra assumption is necessary for this, namely that the probability distribution of the observation disturbance $V k+1$ depends explicitly only on the immediately preceding state, control, and system disturbance $X k, U k, w_{k}$, and not on Xk-I, ..., Xo, Uk-I, ... ,Uo, Wk-I ...,Wo, Vk-I,•.. ,Va. Under this assumption, we will show that for all $k$ and $I_{k}$, we have

$$
\begin{equation*}
J_{k}\left(I_{k}\right)=\min _{u_{k} \in U_{k}} H_{k}\left(P_{x_{k} \mid I_{k}}, u_{k}\right)=\bar{J}_{k}\left(P_{x_{k} \mid I_{k}}\right) \tag{5.34}
\end{equation*}
$$

where $H_{k}$ and $\bar{J}_{k}$ are appropriate functions.

To this end, we will use an important fact that relates to state estimation of discrete-time stochastic systems: the conditional distribution $P_{x k} l_{l k}$ can be generated recursively by an equation of the form

$$
\begin{equation*}
P_{x_{k+1} \mid I_{k+1}}=\Phi_{k}\left(P_{x_{k} \mid I_{k}}, u_{k}, z_{k+1}\right) \tag{5.35}
\end{equation*}
$$

where $\Phi_{k}$ is some function that can be determined from the data of the problem. Let us postpone a justification of this for the moment, and accept it for the purpose of the following discussion.

We note that to perform the minimization in Eq. (5.32), it is sufficient to know the distribution $P X N-I I I N--1$ together with the distribution $P_{w_{N-1} \mid x_{N-1}, u_{N-1}}$, which is part of the problem data. Thus, the minimization in the right-hand side of Eq. (5.32) is of the form
$J_{N-1}\left(I_{N-1}\right)=\min _{u_{N-1} \in U_{N-1}} H_{N-1}\left(P_{x_{N-1} \mid I_{N-1}}, u_{N-1}\right)=\bar{J}_{N-1}\left(P_{x_{N-1} \mid I_{N-1}}\right)$
for appropriate functions $H_{N-}$, and $\bar{J} N-I$.
We now use induction, i.e., we assume that

$$
\begin{equation*}
J_{k+1}\left(I_{k+1}\right)=\min _{u_{k+1} \in U_{k+1}} \quad\left(P_{x_{k+1} \mid I_{k+1}}, u_{k+1}\right)=\bar{J}_{k+1}\left(P_{x_{k+1} \mid I_{k+1}}\right) \tag{5.36}
\end{equation*}
$$

for appropriate functions $H_{k+1}$ and $l k+l$, and we show that

$$
\begin{equation*}
J_{k}\left(I_{k}\right)=\min _{u_{k} \in U_{k}} H_{k}\left(P_{x_{k} \mid I_{k}}, u_{k}\right)=\bar{J}_{k}\left(P_{x_{k} \mid I_{k}}\right) \tag{5.37}
\end{equation*}
$$

for appropriate functions $H_{k}$ and $J_{k}$.
Indeed, using Eqs. (5.35) and (5.36), the DP equation (5.33) is written as

$$
\begin{equation*}
J_{k}\left(I_{k}\right)=\min _{U k E \eta_{k}} E\left\{9 k(X k, U k, W k)+\bar{J}_{k+1}\left(\Phi_{k}\left(P_{x_{k} \mid I_{k}}, u_{k}, z_{k+1}\right)\right) \mid I k, u k\right\}^{\prime} \tag{5.38}
\end{equation*}
$$

In order to calculate the expression being minimized over $U k$ above, we need, in addition to $P x k l_{l k}{ }^{\prime}$ the joint distribution

$$
P\left(x_{k}, w_{k}, z_{k+1} \mid I_{k}, u_{k}\right)
$$

or, equivalently,

$$
P\left(x_{k}, w_{k}, h_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right), u_{k}, v_{k+1}\right) \mid I_{k}, u_{k}\right)
$$

This distribution can be expressed in terms of $P_{X k} I l k^{\prime}$ the given distributions

$$
P\left(w_{k} \mid x_{k}, u_{k}\right), \quad P\left(v_{k+1} \mid f_{k}\left(x_{k}, u_{k}, w_{k}\right), u_{k}, w_{k}\right)
$$

and the system equation $X k+l=f_{k}\left(x_{k}, u k, W k\right)$. Therefore the expression minimized over $U k$ in Eq. (5.38) can be written as a function of PXkllk and ' $U k$, and the DP equation (5.33) can be written as

$$
J_{k}\left(I_{k}\right)=\min _{u_{k} \in U_{k}} H_{k}\left(P_{x_{k} \mid I_{k}}, u_{k}\right)
$$

for a suitable function $H k$. Thus the induction is complete and it follows that the distribution PXkllk is a sufficient statistic.

Note that if the conditional distribution $P_{X k} l l k$ is uniquely determined by another expression $S_{k}\left(I_{k}\right)$, i.e.,

$$
P_{x_{k} \mid I_{k}}=G_{k}\left(S_{k}\left(I_{k}\right)\right)
$$

for an appropriate function $G k$, then $S_{k}\left(I_{k}\right)$ is also a sufficient statistic. Thus, for example, if we can show that PXkllk is a Gaussian distribution, then the mean and the covariance matrix corresponding to PXkllk form a
sufficient statistic.

Regardless of its computational value, the representation of the optimal policy as a sequence of functions of the conditional probability distri-

$$
\mu_{k}\left(I_{k}\right)=\bar{\mu}_{k}\left(P_{x_{k} \mid I_{k}}\right), \quad k=0,1, \ldots, N-1
$$

is conceptually very usefuL It provides a decomposition of the optimal controller in two parts:
(a) An estimator, which uses at time $k$ the measurement $Z k$ and the control $U k-l$ to generate the probability distribution $P_{X k}^{11 k}$
(b) An actuator, which generates a control input to the system as a function of the probability distribution $P_{X k} l l k$ (Fig. 5.4.1).
This interpretation has formed the basis for various suboptimal control schemes that separate the controller a priori into an estimator and an actuator and attempt to design each part in a manner that seems "reasonable." Schemes of this type will be discussed in Chapter 6.

The

## -vmenuluans State Distribution Recursion

There remains to justify the recursion

$$
\begin{equation*}
P_{x_{k+1} \mid I_{k+1}}=\Phi_{k}\left(P_{x_{k} \mid I_{k}}, u_{k}, z_{k+1}\right) \tag{5.39}
\end{equation*}
$$

Let us first give an example.


Figure 5.4.1 Conceptual separatIOn of the optimal controller into an estimator and an actuator.

## Example 5.4.1 (Search Problem)

In a classical problem of search, one $h$ s to decide at each period whether to
 reveals It WIt pro a $h^{1} \quad$ t $\quad 0$ values' el ${ }^{\text {th }}{ }_{e r}$ a treasure is present in the Site
 or it is not. $\mathrm{T}_{\mathrm{d}}$ e the con observation $Z k+1 \mathrm{t}_{\mathrm{a}} \mathrm{k}_{\mathrm{es}}$ two values, treasure found or not folternd, searife ff the SHe-IS not search ${ }^{\text {d }}$, the value of $z k+1$ is Irrelevant.

## Denote

$P k$ : probability a treasure IS present at the beginning of period $k$.
This probability evolves according to the equation

$P k+I^{-} \quad$| $P k$ | if the site is not searched at time $k$, |
| :--- | :--- |
| $\underbrace{P k}_{P k(I-(3)+1}$ | if the site is searched and a treasure is found, |
| if the site is searched but no treasure IS found. |  |

The second relatiOn holds becaus the treasure is removed after a successful
 to the kth period probabIhty of a treasure being of an unsuccessful search). The preunsuccessful divided by the probabilit $\begin{aligned} & \text { a } \\ & \text { ceding equation defines the desIred recurSIOn or }\end{aligned}$ of the state and is a special case of Eq. (5.39).

The general form of the recursion

$$
P_{x k+1} I I_{k+1}=\Phi_{k}\left(P_{X k} \| / k^{\prime} U k, Z k+l\right)
$$

is developed in ExerClse 5 . ${ }^{7}$ for the, case where the state, control, observa-
tion, and disturbance spaces are ill ${ }^{\text {ta }}$ s. ${ }^{\text {s. }}$.
are the real line and all random variables involved possess probability density functions, the conditional density $p\left(X k+l \mid I_{k+1}\right)$ is generated from $p(X k 11 k), U k$, and $Z k+l$ by means of the equation

$$
\begin{aligned}
p\left(X k+l \mid I_{k+1}\right)= & p\left(X k+l \mid I_{k}, U k, Z k+l\right) \\
& \frac{p\left(X k+l, Z k+l \left\lvert\, \frac{\left.I_{k}, u_{k}\right)}{p(Z k+l \mid l k, U k)}\right.\right.}{} \\
= & \frac{p\left(X k+l \mid I_{k}, U k\right) p\left(Z k+l \mid I_{k}, U k, X k+l\right)}{\int_{-\infty}^{\infty} p\left(X k+l \mid I_{k}, U k\right) p\left(Z k+l \mid 1 k, U k, x_{k+1}\right) d x_{k+1}}
\end{aligned}
$$

In this equation all the probability densities appearing in the right-hand side may be expressed in terms of $p\left(X k \quad 11_{k}\right), U k$, and $Z k+l$. In particular, the density $p\left(X k+l \mid I_{k}, U k\right)$ may be expressed through $p(X k 11 k), U k$, and the system equation $X k+l=f k(X k, U k, W k)$ using the given density $p(W k, X k, U k)$ and the relation

$$
p\left(W k \mid I_{k}, U k\right)=\int_{-\infty}^{\infty} p\left(X k \mid I_{k}\right) p\left(w_{k} \mid X k, U k\right) d x k
$$

Similarly, the density $p\left(Z k+l, I_{k}, U k, X k+l\right)$ is expressed through the measurement equation $Z k+l=h_{k+1}(X k+l, U k, V k+l)$ using the densities

$$
p(X k \quad 11 k), \quad p(W k \mid X k, U k), \quad p(V k+l \mid X k, U k, W k) .
$$

By substituting these expressions in the equation for $p(X k+l I I k+l)$, we obtain a dynamic system equation for the conditional state distribution of the desired form. Other similar examples will be given in subsequent sections. A mathematically rigorous substantiation of the recursion $P_{X k}+1 I_{I} k+l$ $\Phi_{k}\left(P x k l l k^{\prime} U k, Z k+l\right)$ can be found in Bertsekas and Shreve [BeS78].

## Alternative Perfect State Information Reduction

Finally, let us formally rewrite the DP algorithm in terms of the sufficient statistic Pxkl ${ }_{l k}{ }^{\prime}$ Using Eqs. (5.35), (5.37), and (5.38), we have for $k<N$ - 1

$$
\begin{align*}
\bar{J}_{k(P X k l l k)=} \min _{u k E U k}\left[{\underset{x k, w k, z k+l}{E}\{ }^{E} \begin{array}{rl} 
& \\
& +X k, U k, W k) \\
& \left.\left.+\bar{J}_{k+l}\left(\Phi_{k}\left(P_{x_{k} \mid I_{k}}, U k, Z k+l\right)\right) \mid I_{k}, U k\right\}\right]
\end{array}\right.
\end{align*}
$$

In the case where $k=N-1$, we have

$$
\begin{aligned}
& \bar{J}_{N-1}\left(P_{x_{N-1} \mid I_{N-1}}\right) \\
& \\
& \qquad \min _{N-I E U N-l} {\left[X N_{-} I E, W N_{-} l\left\{g_{N}\left(f_{N-1}\left(x_{N-1}, u_{N-1}, w_{N-1}\right)\right)\right.\right.} \\
&\left.\left.+g_{N-1}\left(x_{N-1}, u_{N-1}, w_{N-1}\right) \mid I N-l, U N-l\right\}\right]
\end{aligned}
$$

This DP algorithm yields the optimal cost as

$$
J^{*}=\underset{z_{0}}{E}\left\{\bar{J}_{0}\left(P_{x_{0} \mid z_{0}}\right)\right\}
$$

where $\bar{J}_{0}$ is obtained by the last step, and the probability distribution of $z o$ is obtained from the measurement equation $Z o=h o(x o, v o)$ and the distributions of $X o$ and $v o$.

By observing the form of Eq. (5.40), we note that it has the standard DP structure, except that Pxkl $l_{k}$ plays the role of the "state." Indeed the role of the "system" is played by the recursive estimator of Pxkl $1 k$,

$$
P_{x_{k+1} \mid I_{k+1}}=\Phi_{k}\left(P_{x_{k} \mid I_{k}}, u_{k}, z_{k+1}\right)
$$

and this system fits the framework of the basic problem (the role of control is played by $U k$ and the role of the disturbance is played by $Z k+l$ ). Furthermore, the controller can calculate (at least in principle) the state $P X k \mid I_{k}$ of this system at time $k$, so perfect state information prevails. Thus the alternate DP algorithm (5.40)-(5.41) may be viewed as the DP algorithm of the perfect state information problem that involves the above system, whose state is $P x k l_{I k}$, and an appropriately reformulated cost function. In the absence of perfect knowledge of the state, the controller can be viewed as controlling the "probabilistic state" Pxkl ${ }_{l k}$ so as to minimize the expected cost-to-go conditioned on the information $I_{k}$ available.

## Example 5.4.1 (Continued)

Let us write the DP algorithm (5.40) for the search problem of Example 5.4.1, assuming that the treasure's worth is $V$, that each search costs $C$, and that once we decide not to search at a particular time, then we cannot search at future times. The algorithm takes the form

$$
\bar{J}_{k}\left(p_{k}\right)=\max \left[0,-C+p_{k} \beta V+\left(1-p_{k} \beta\right) \bar{J}_{k+1}\left(\frac{p_{k}(1-\beta)}{p_{k}(1-\beta)+1-p_{k}}\right)\right]
$$

with $\bar{J} N(P N)=0$. From this algorithm, it is straightforward to show by induction that the functions $] \mathrm{k}$ satisfy $] k(P k)=0$ for all $P k \leq C /(f 3 V)$, and that it is optimal to search at period $k$ if and only if

$$
C \leq p_{k} \beta V
$$

Thus, it is optimal to search if and only if the expected reward from the next search is greater or equal to the cost of the search.

### 5.4.2 Finite-State Systems

We will now consider systems that are stationary finite-state Markov chains, in which case the conditional probability distribution PXklik is characterized by a finite set of numbers. The states are denoted $1,2, \ldots, n$. When a control $u$ is applied, the system moves from state $i$ to state $j$ with probability Pij $(U)$. The control $U$ is chosen from a finite set $U$. Following a state transition, an observation is made by the controller. There is a finite number of possible observation outcomes, and the probability of each depends on the current state and the preceding control. The information available to the controller at stage $k$ is the information vector

$$
I_{k}=\left(z_{1}, \ldots, z_{k}, u_{0}, \ldots, u_{k-1}\right)
$$

where for all $\mathrm{i}, \mathrm{Zi}$ and $U i$ are the observation and control at stage i, respectively. Following the observation $Z k$, a control ' $U k$ is chosen by the controller, and a cost $g\left(X k^{\prime} U k\right)$ is incurred, where $X k$ is the current (hidden) state. The terminal cost for being at state $x$ at the end of the $N$ stages is denoted $G(x)$. We wish to minimize the expected value of the sum of costs incurred over the IV stages.

As discussed in Section 5.4.1, one can reformulate the problem into a problem of perfect state information: the objective is to control the column vector of conditional probabilities

$$
p_{k}=\left(p_{k}^{1}, \ldots, p_{k}^{n}\right)^{\prime}
$$

where

$$
p_{k}^{i}=P\left(x_{k}=i \mid I_{k}\right), \quad \mathrm{i}=1, \ldots, n o
$$

We refer to $P k$ as the belief state. It evolves according to an equation of the form

$$
P k+1=\Phi\left(p_{k}, U k, Z k+1\right)
$$

where the function $\Phi$ represents an estimator, as discussed in Section 5.4.1. The initial belief state $P o$ is given.

The corresponding DP algorithm [see Eqs. (5.40) and (5.41)] has the form

$$
\begin{equation*}
\min _{u_{k} \in U}\left[p_{k}^{\prime} g\left(u_{k}\right)+E_{z_{k+1}}\left\{\bar{J}_{k+1}\left(\Phi\left(p_{k}, u_{k}, z_{k+1}\right)\right) \mid p_{k}, u_{k}\right\}\right] \tag{5.42}
\end{equation*}
$$

where $g(U k)$ is the column vector with components $g(l, U k), \ldots, g(n, U k)$, and $p_{k}^{\prime} g\left(u_{k}\right)$, the expected stage cost, is the inner product of the vectors $P k$ and $g\left(u_{k}\right)$. The algorithm starts at stage IV, with

$$
\bar{J}_{N}\left(p_{N}\right)=p_{N}^{\prime} G
$$

where $G$ is the column vector with components the terminal costs $G(i)$, $\mathrm{i}=1, \ldots, n$, and proceeds backwards. Note that in this DP algorithm, the conditional distribution of $Z k+1$ given $P k$ and $U k$ can be computed using the transition probabilities Pij $(u)$, and the known conditional distribution of $Z k+l$ given $X k+l$ and $U k$. In particular, we have for any possible observation value $z$,

$$
P\left(Z k+1=\left.Z\right|_{P k, U k)}=\sum_{\mathrm{i}=1}^{n} p_{k}^{i} \underset{\substack{n \\ j=1}}{n} i j(U k) P(Z k+1=Z \mid X k+1=j, U k)^{\prime}\right.
$$

It turns out that the cost-to-go functions $\bar{J}_{k}$ in the DP algorithm are piecewise linear and concave. The demonstration of this fact is straightforward, but tedious, and is outlined in Exercise 5.7. A consequence of the piecewise linearity property is that $\bar{J}_{k}$ can be characterized by a finite set of scalars. Still, however, for fixed $k$, the number of these scalars can increase fast with IV, and there may be no computationally efficient way to solve the problem (see Papadimitriou and Tsitsiklis [PaT87J). We will not discuss here any special procedures for computing (see Lovejoy [Lov91a], [Lov91b], and Smallwood and Sondik [SmS73], [Son71J). Instead we will demonstrate the DP algorithm by means of examples.

## Example 5.4.2 (Machine Repair)

In the two-state machine repair example of Section 5.1, let us denote

$$
P I=P\left(X l=\bar{P} \mid I_{1}\right), \quad p o=P(x o=\bar{P} \mid 10)^{\prime}
$$

The equation relating $P l, P O, U n, \mathrm{Zl}$ is written as

$$
P I=\Phi_{0}(p o, U n, Z l) .
$$

One may verify by straightforward calculation that $\Phi_{0}$ is given by

$$
p_{1}=\Phi_{0}\left(p_{0}, u_{0}, z_{1}\right)=\left\{\begin{array}{lll}
\frac{1}{7} & \text { if } U_{0}=5, & Z l=\mathrm{G} \\
\frac{3}{5} & \text { if } U_{o}=5, & Z I=B \\
\frac{1+2 p_{0}}{7-4 p_{0}} & \text { if } u 1=\mathrm{C}, & Z I=\mathrm{G} \\
\frac{3+6 p_{0}}{5+4 p_{0}} & \text { if } U_{0}=\mathrm{C}, & Z I=B
\end{array}\right.
$$

The DP algorithm (5.42) may be written in terms of Po, PI, and $\Phi_{0}$ above as

$$
l l(P 1)=\min [2 \mathrm{pl}, \mathrm{IJ},
$$

$$
\begin{aligned}
l o(p o)=\min [2 P O & +P(Z I=\mathrm{G} \mid P o, C)] l\left(\Phi_{0}(p o, \mathrm{C}, \mathrm{G})\right) \\
& +P(Z I=B \mid p o, C)] l\left(\Phi_{0}(p o, \mathrm{C}, B)\right) \\
1+ & P(Z l=\mathrm{G} \mid P o, 5)] 1\left(\Phi_{0}(p o, 5, G)\right) \\
& +p(Z l=B \mid p o, 5)] 1\left(\Phi_{0}(p o, 5, B)\right) \mathrm{J}
\end{aligned}
$$

The probabilities entering in the second equation may be expressed in terms of $P o$ by straightforward calculation as

$$
\begin{aligned}
\mathrm{P}(Z I=G \mid p o, C)=\frac{7-4 p o}{12}, & P(Z I=B \mid P o, C)=\frac{5+4 p o}{12}, \\
P(Z I=G \mid p o, S)=\underline{7}, & P(Z I=B \mid p o, S)=\frac{5}{12}-
\end{aligned}
$$

Using these values we have

$$
\begin{gathered}
1 \mathrm{o}(P o)=\min \left[2_{P O}+\underline{7}_{72}^{-4 P O]_{1}} \frac{(1}{7}+\frac{\left.2 P_{0}\right)}{4 p o}+\underline{5}+\frac{4 p o]_{1}}{12} \frac{(3}{5}+\frac{6 P o)}{\mp p o}\right. \\
\left.1+\frac{7}{12} \bar{J}_{1}\left(\frac{1}{7}\right)+\frac{5}{12} J_{1}\left(\frac{3}{5}\right)\right] .
\end{gathered}
$$

By minimization in the equation defining] 1 (PI), we obtain an optimal policy for the last stage

$$
\begin{cases}C & \text { if } p_{1} \leq \frac{1}{2} \\ S & \text { if } p_{1}>\frac{1}{2}\end{cases}
$$

Also by substitution of $] 1$ (PI) and by carrying out the straightforward calculation, we obtain

$$
\bar{J} o(p o)=\begin{array}{ll}
\frac{19}{12} & \text { if } \frac{3}{8} \leq P o \leq 1, \\
& \text { if } \leq p o \leq \frac{3}{8},
\end{array}
$$

and an optimal policy for the first stage:

$$
\mu_{0}^{*}(P o)= \begin{cases}C & \text { if } P o \leq \frac{3}{8} \\ & \text { if } P o>\frac{3}{8}\end{cases}
$$

Note that

$$
\begin{array}{cc}
P(x o=\bar{P} \mid Z o=G)=\frac{1}{7}, & P(x o=\bar{P} \mid Z o=B)=\frac{3}{5}, \\
P(z o=G)=\underline{7}, & P(z o=B)=\frac{5}{12},
\end{array}
$$

so that the formula

$$
J^{\star} \quad \underset{z_{0}}{E}\left\{\bar{J}_{0}\left(P_{x_{0} \mid z_{0}}\right)\right\}=\frac{7}{12} \bar{J}_{0}\left(\frac{1}{7}\right)+\frac{5}{12} \bar{J}_{0}\left(\frac{3}{5}\right)=\frac{176}{144}
$$

yields the same optimal cost as the one obtained in Section 5.1 by means of the DP algorithm (5.4) and (5.5).

## Example 5.4.3 (A Problem of Instruction)

Consider a problem of instruction where the objective is to teach a student a certain simple item. At the beginning of each period, the student may be in one of two possible states:
$L$ : Item learned.
$L$ : Item not learned.
At the beginning of each period, the instructor must make one of two decisions:
$T$ : Terminate the instruction.
$\bar{T}$ : Continue the instruction for one period and then conduct a test that indicates whether the student has learned the item.

The test has two possible outcomes:
$R$ : Student gives a correct answer.
$R$ : Student gives an incorrect answer.
The transition probabilities from one state to the next if instruction takes place are given by

$$
\begin{array}{ll}
P(X k+l \quad L \mid X k=L)=1, & P\left(x_{k+1}=\bar{L} \mid X k=L\right)=0, \\
P(X k+l=L \mid X k=\bar{L})=t, & P(X k+l=\bar{L} \mid X k=\bar{L})=1-t,
\end{array}
$$

where $t$ is a given scalar with ${ }^{\mathrm{O}}<t<1$.
The outcome of the test depends probabilistically on the state of knowl-. edge of the student as follows:

$$
\begin{array}{ll}
P(Z k=R \mid X k=L)=1, & P(Z k=\bar{R} \mid X k=L)=0, \\
P(Z k=R \mid X k=\bar{L})=r, & P(Z k=\bar{R} \mid X k=\bar{L})=1-r,
\end{array}
$$

$$
\mathrm{O}_{<}
$$

where $r$ is a given scalar with $<r<1$. The probabilistic structure of the problem is illustrated in Fig. 5.4.2.


Figure 5.4.2 Probabilistic structure of the instruction problem.

The cost of instruction and testing is I per period. The cost of terminating instruction is 0 or $C>0$ if the student has learned or has not learned the item, respectively. The objective is to find an optimal instruction-termination policy for each period $k$ as a function of the test information accrued up to that period, assuming that there is a maximum of $N$ periods of instruction.

It is straightforward to reformulate this problem into the framework of the basic problem with imperfect state information and to conclude that the decision whether to terminate or continue instruction at period $k$ should depend on the conditional probability that the student has learned the item given the test results so far. This probability is denoted

$$
P k=P\left(x_{k} \mid I_{k}\right)=P\left(X k=L \mid Z o, z_{1}, \ldots, Z k\right)
$$

In addition, we can use the DP algorithm (5.40) and (5.41) defined over the space of the sufficient statistic $P k$ to obtain an optimal policy. However, rather than proceeding with this elaborate reformulation, we prefer to argue and obtain this DP algorithm directly.

Concerning the evolution of the conditional probability $P k$ (assuming instruction occurs), we have by Bayes' rule

$$
p_{k+1}=P\left(x_{k+1}=L \mid z_{0}, \ldots, z_{k+1}\right)=\frac{P\left(x_{k+1}=L, z_{k+1} \mid z_{0}, \ldots, z_{k}\right)}{P\left(z_{k+1} \mid z_{0}, \ldots, z_{k}\right)}
$$

where

$$
\begin{aligned}
P(Z k+1 & \mid Z o, \cdot ., Z k) \\
= & P(X k+1=L \mid Z 0, \ldots, Z k) P(Z k+1 \mid Z o, \ldots, Z k, X k+1=L) \\
& \quad+P(X k+1=\bar{L} \mid Z 0, \ldots, Z k) P(Z k+l \mid Z o, \ldots, Z k, X k+1=\bar{L})
\end{aligned}
$$

From the probabilistic descriptions given, we have
$P\left(z_{k+1} \mid z_{0}, \ldots, z_{k}, x_{k+1}=L\right)=P\left(z_{k+1} \mid x_{k+1}=L\right)= \begin{cases}1 & \text { if } Z k+1=R, \\ 0 & \text { if } Z k+1=\bar{R} .\end{cases}$

$$
\begin{aligned}
P(Z k+1 \mid Z 0, \ldots, Z k, X k+1=\bar{L}) & =P\left(z_{k+1} \mid X k+1=\bar{L}\right) \\
& =\left\{\begin{aligned}
r & \text { if } Z k+1=1=\frac{R}{R} \\
1-r & \text { if } Z k+1=\bar{R}
\end{aligned}\right. \\
P(X k+1=L \mid Z o, \quad, Z k) & =P k+(1-p k) t, \\
P(X k+1=\bar{L} \mid Z o, \quad, Z k) & =(1-P k)(l-t) .
\end{aligned}
$$

Combining these equations, we obtain

$$
p_{k+1}=\Phi\left(p_{k}, z_{k+1}\right)= \begin{cases}\frac{p_{k}+\left(1-p_{k}\right) t}{p_{k}+\left(1-p_{k}\right) t+\left(1-p_{k}\right)(1-t) r} & \text { if } Z k+1=R \\ 0 & \text { if } z_{k+1}=\bar{R}\end{cases}
$$

or equivalently

$$
p_{k+1}=\Phi\left(p_{k}, z_{k+1}\right)= \begin{cases}\frac{1-(1-t)\left(1-p_{k}\right)}{1-(1-t)(1-r)\left(1-p_{k}\right)} & \text { if } Z k+1=R  \tag{5.43}\\ 0 & \text { if } Z k+1=\bar{R}\end{cases}
$$

This equation is a special case of the general recursive update equation (5.39) for the conditional probability of the state. A cursory examination of Eq. (5.43) shows that, as expected, the conditional probability $p_{k+1}$ that the student has learned the item increases with every correct answer and drops to zero with every incorrect answer.

We now derive the DP algorithm for the problem. At the end of the $N t h$ period, assuming instruction has continued to that period, the expected cost is

$$
\vec{J}_{N}\left(p_{N}\right)=\left(1-p_{N}\right) C
$$

At the end of period $N-1$, the instructor has calculated the conditional probability $P N-I$ that the student has learned the item and wishes to decide whether to terminate instruction and incur an expected cost $(1-P N-I) C$ or continue the instruction and incur an expected cost $I+E\{ ] N(p N)\}$. This leads to the following equation for the optimal expected cost-to-go:

$$
] N-I(P N-I)=\min [(1-P N-I) C, 1+(1-t)(l-P N-I) C]
$$

The term (1-- PN-I)C is the cost of terminating instruction, while the term $(1-t)(l-P N-I)$ is the probability that the student still has not learned the item following an additional period of instruction.

Similarly, the algorithm is written for every stage $k$ by replacing $N$ by $k+1$ :

$$
\bar{J}_{k}\left(p_{k}\right)=\min \left[\left(1-p_{k}\right) C, I+\underset{z_{k+1}}{E}\left\{\bar{J}_{k+1}\left(\Phi\left(p_{k}, z_{k+1}\right)\right)\right\}\right]
$$

Now using expression (5.43) for the function $\Phi$ and the probabilities

$$
\begin{aligned}
& P(Z k+l=\widetilde{R} \mid P k)=(1-t)(l-r)(l-P k) \\
& P(Z k+l=R \mid P k)=1-(1-t)(l-r)(1-P k),
\end{aligned}
$$

we have

$$
\bar{J}_{k}\left(p_{k}\right)=\min \left[\left(1-p_{k}\right) C, I+A_{k}\left(p_{k}\right)\right]
$$

where

$$
\begin{aligned}
A k(P k)= & P\left(Z k+1=R \mid I_{k}\right) \bar{J}_{k+1}\left(\Phi\left(p_{k}, R\right)\right) \\
& +P\left(Z k+1=\bar{R} \mid I_{k}\right) \bar{J}_{k+1}\left(\Phi\left(p_{k}, \bar{R}\right)\right)
\end{aligned}
$$

or, equivalently, using Eq. (5.43),

$$
+(1-t)(l-r)(l-P k)] k+I(O)
$$

As shown in Fig. 5.4.3, if $1+(1-t) C \leq \mathrm{C}$ or, equivalently, if


Figure 5.4.3 Determining the optimal instruction policy in the last period.
then there exists a scalar $\alpha_{N-1}$ with optimal policy for the last period:

$$
\begin{aligned}
\text { continue instruction } & \text { if } P N-I \leq \alpha_{N-1} \\
\text { terminate instruction } & \text { if } P N-I>\alpha_{N-1}
\end{aligned}
$$

In the opposite case, where $I \geq t e$, the cost of instruction is so high relative to the cost of not learning that instructing the student is never optimal.

It may be shown by induction (Exercise 5.8) that if $I<t e$, the functions $A^{k}(p)$ are concave and piecewise linear for each $k$ and satisfy, for all $k$,

$$
A_{k}(1)=0
$$

Furthermore, they satisfy for all $k$,

$$
\begin{gathered}
A_{k}(p) \geq A_{k}\left(p^{\prime}\right), \quad \text { for } \quad{ }^{\circ} \leq P<p^{\prime} \leq 1, \\
A k-1(p) \leq A k(p) \leq A k+1(p), \quad \text { for all } P \mathrm{E}[0,1] .
\end{gathered}
$$

Thus the functions. $(1-P k) C$ and $1+A k(P k)$ intersect at a single point, and from the DP algonthm (5.44), we obtain that the optimal policy for each penod is determined by the unique scalars $a k$, which are such that

$$
\left(1-\alpha_{k}\right) C=I+A_{k}\left(\alpha_{k}\right), \quad k=0,1, \ldots, N-1
$$

An optimal policy for period $k$ is given by

| continue instruction | if $P k \leq \alpha_{k}$, |
| ---: | :--- |
| terminate instruction | if $P k>\alpha_{k}$, |

Since the functions $A_{k}(p)$ are monotonically nondecreasing with respect to $k$, it follows from Fig. 5.4.4 that

$$
\alpha_{N-1} \leq \alpha_{N-2} \leq \ldots \leq \alpha_{k} \leq \ldots \leq 1-\frac{1}{\mathrm{C}^{\prime}}
$$

and therefore the sequence $\left\{\alpha_{k}\right\}$ converges to some scalar $\bar{\alpha}$ as $k \rightarrow-\mathbf{o o}$. Thus, as the horizon gets longer, the optimal policy (at least for the initial stages) can be approximated by the stationary policy

$$
\begin{align*}
\text { continue instruction } & i f p k \leq \bar{\alpha}_{k} \\
\text { terminate instruction } & \text { ifpk }>\bar{\alpha}_{k} \tag{5.45}
\end{align*}
$$



Figure 5.4.4 Demonstrating that the instruction thresholds are decreasing with time.

It turns out that this stationary policy has a convenient implementation that does not require the calculation of the conditional probability at each stage. From Eq. (5.43), we see that $P k+l$ increases over $P k$ if a correct answer $R$ is given, and drops to zero if an incorrect answer $\bar{R}$ is given. For $m=1,2, \ldots 0^{\prime}$ define recursively the probability $\pi_{m}$ of getting $m$ successive correct answers following an incorrect answer:

$$
\pi_{1}=\Phi(0, R), \quad \pi_{2}=\Phi\left(\pi_{1}, R\right), \quad, \pi_{k+1}=\Phi\left(\pi_{k}, R\right), . .
$$

Let $n$ be the smallest integer for which $\pi n>\overline{\boldsymbol{\alpha}}$. Then the stationary policy (5.45) can be implemented by terminating instruction if and only if $n$ successive correct answers have been received.

## Example 5.4.4 (Sequential Hywumesis Testing)

Let us consider a hypothesis testing problem that is typical of statistical sequential analysis. A decision maker can make observations, at a cost C each, relating to two hypotheses. Given a new observation, he can either accept one of the hypotheses or delay the decision for one more period, pay the cost C , and obtain a new observation. At issue is trading off the cost of observation with the higher probability of accepting the wrong hypothesis.

Let $Z 0, \mathrm{zl}, \ldots, \mathrm{ZN}-l$ be the sequence of observations. We assume that they are independent, identically distributed random variables taking values on a finite set $Z$. Suppose we know that the probability distribution of the $Z k^{\prime} S$ is either $f o$ or $f_{1}$ and that we are trying to decide on one of these. Here, for any element $z \mathrm{E} Z, f o(z)$ and $f_{1}(z)$ denote the probabilities of $z$ when $f o$ and $f_{1}$ are the true distributions, respectively. At time $k$ after observing $Z 0, \ldots, Z k$, we may either stop observing and accept either $f o$ or $f_{1}$, or we may take an additional observation at a cost $\mathrm{C}>0$. If we stop observing and make a choice, then we incur zero cost if our choice is correct, and costs $L_{o}$ and $L_{1}$ if we choose incorrectly $f o$ and $f_{1}$, respectively. We are given the a priori probability $P$ that the true distribution is $f o$, and we assume that at most $N$ observations are possible.

It can be seen that we are faced with an imperfect state information problem involving the two states:
$x^{0}:$ true distribution is $f o$,
$x_{1}$ : true distribution is $f_{1}$

The alternate DP algorithm (5.40) and (5.41) is defined over the interval [ 0,1 ] of possible values of the conditional probability

$$
p_{k}=P\left(x_{k}=x^{0} \mid z_{0}, \ldots, z_{k}\right) .
$$

Similar to the previous section, we will obtain this algorithm directly.
The conditional probability $P k$ is generated decursively according to the $^{2}$ following equation [assuming $f o(z)>0, f_{1}(z)>$ for all $z \mathrm{E} \mathrm{Z]:}$

$$
\begin{gather*}
P o=\frac{p f o(z o)}{p f_{0}\left(z_{0}\right)^{\prime}+(1-p) f_{1}\left(z_{0}\right)},  \tag{5.46}\\
P k+l=\frac{p_{k} f_{0}\left(z_{k+1}\right)}{P k f o(z k+l)+\left(1-p_{k}\right) f_{1}(Z k+l)}, \quad k=0,1, \ldots, N-2, \tag{5.47}
\end{gather*}
$$

where P is the a priori probability that the true distribution is $f 0$. The optimal expected cost for the last period is

$$
\begin{equation*}
\bar{J}_{N-1}\left(p_{N-1}\right)=\min \left[\left(1-p_{N-1}\right) L_{0}, p_{N-1} L_{1}\right] \tag{5.48}
\end{equation*}
$$

where $\left(1-P_{N}-1\right) L o$ is the expected cost for accepting $f o$ and $P N-1 L_{1}$ is the expected cost for accepting $f_{1}$. Taking into account Eqs. (5.46) and (5.47),
we obtain the optimal expected cost-to-go for the kth period as

$$
\begin{aligned}
\bar{J}_{k}\left(p_{k}\right)=\min & {\left[\left(1-p_{k}\right) L_{0}, p_{k} L_{1}\right.} \\
& \left.C+\underset{z_{k+1}}{E}\left\{\bar{J}_{k+1}\left(\frac{p_{k} f_{0}\left(z_{k+1}\right)}{p_{k} f_{0}\left(z_{k+1}\right)+\left(1-p_{k}\right) f_{1}\left(z_{k+1}\right)}\right)\right\}\right]
\end{aligned}
$$

where the expectation over $Z k+l$ is taken with respect to the probability distribution

$$
p\left(z_{k+1}\right)=p_{k} f_{0}\left(z_{k+1}\right)+\left(1-p_{k}\right) f_{1}\left(z_{k+1}\right)
$$

Equivalently, for $k=0,1, \ldots, N-2$,

$$
\begin{equation*}
\bar{J}_{k}\left(p_{k}\right)=\min \left[\left(1-p_{k}\right) L_{0}, p_{k} L_{1}, C+A_{k}\left(p_{k}\right)\right] \tag{5.49}
\end{equation*}
$$

where

$$
\begin{equation*}
A_{k}\left(p_{k}\right)=\underset{z_{k+1}}{E}\left\{\bar{J}_{k+1}\left(\frac{p_{k} f_{0}\left(z_{k+1}\right)}{p_{k} f_{0}\left(z_{k+1}\right)+\left(1-p_{k}\right) f_{1}\left(z_{k+1}\right)}\right)\right\} \tag{5.50}
\end{equation*}
$$

An optimal policy for the last period (see Fig. 5.4.5) is obtained from the minimization indicated in Eq. (5.48):

$$
\begin{array}{ll}
\text { accept } f o & \text { if } P N-l \geq \gamma \\
\text { accept } f_{1} & \text { if } P N-l<\gamma
\end{array}
$$

where $\gamma$ is determined from the relation ( $1-\gamma) L_{0}=\gamma L_{1}$ or equivalently

$$
\gamma=\frac{L o}{L_{o}+L_{1}} .
$$



Figure 5.4.5 Determining the optimal policy for the last period.

We now prove that the functions $A k:[0,1] \rightarrow R$ of Eq. (5.50) are concave, and satisfy for all $k$ and $P \mathrm{E}[0,1]$

$$
\begin{gather*}
A_{k}(0)=A_{k}(1)=0  \tag{5.51}\\
A k-1(p) \leq A k(p) \tag{5.52}
\end{gather*}
$$

Indeed, we have for all $P E[0,1]$

$$
\bar{J}_{N-2}(p) \leq \min \left[(1-p) L_{0}, p L_{0}\right]=\bar{J}_{N-1}(p)
$$

By making use of the stationarity of the system and the monotonicity property of DP (Exercise 1.23 in Chapter 1), we obtain

$$
\bar{J}_{k}(p) \leq \bar{J}_{k+1}(p)
$$

for all $k$ and $P E[0,1]$. Using Eq. (5.50), we obtain $A_{k}+l(p) \leq A k(p)$ for all $k$ and $p \mathrm{E}[0,1]$.

To prove concavity of $A k$ in view of Eqs. (5.48) and (5.49), it is sufficient to show that concavity of $] k+1$ implies concavity of $A_{k}$ through relation (5.50). Indeed, assume that $\bar{J}_{k+1}$ is concave over [0,1]. Let $Z l, Z 2, \ldots, z n$ denote the elements of the observation space $Z$. We have from Eq. (5.50) that

$$
A_{k}(p)=\sum_{i=1}^{n}\left(p f_{0}\left(z^{i}\right)+(1-p) f_{1}\left(z^{i}\right)\right) \bar{J}_{k+1}\left(\frac{p f_{0}\left(z^{i}\right)}{p f_{0}\left(z^{i}\right)+(1-p) f_{1}\left(z^{i}\right)}\right)
$$

Hence it is sufficient to show that concavity of $] k+1$ implies concavity of each of the functions

$$
H_{i}(p)=\left(p f_{0}\left(z^{i}\right)+(1-p) f_{1}\left(z^{i}\right)\right) \bar{J}_{k+1}\left(\frac{p f_{0}\left(z^{i}\right)}{p f_{0}\left(z^{i}\right)+(1-p) f_{1}\left(z^{i}\right)}\right)
$$

To show concavity of $H i$, we must show that for every $\lambda \mathrm{E}[0,1], P 1$, $P 2$ E [0,1] we have

$$
\lambda H_{i}\left(p_{1}\right)+(1-\lambda) H_{i}\left(p_{2}\right) \leq H_{i}\left(\lambda p_{1}+(1-\lambda) p_{2}\right)
$$

Using the notation

$$
\xi_{1}=p_{1} f_{0}\left(z^{i}\right)+\left(1-p_{1}\right) f_{1}\left(z^{i}\right), \quad \xi_{2}=p_{2} f_{0}\left(z^{i}\right)+\left(1-p_{2}\right) f_{1}\left(z^{i}\right)
$$

this inequality is equivalent to

$$
\begin{aligned}
& \lambda \xi_{1}+\frac{\lambda \xi_{1}}{(1-\lambda) \xi_{2}} \bar{J}_{k+1} \frac{(P d a(z i))}{\xi_{1}}+\xi_{1}{ }^{\frac{(1}{+}-\frac{\lambda) \xi_{2}}{(1-\lambda) \xi_{2}}} \bar{J}_{k+1} \frac{(P 2 f a(Z i))}{\xi_{2}} \\
& \leq \bar{J}_{k+1}\left(\frac{\left(\lambda p_{1}+(\mathbb{1}-\lambda) p_{2}\right) f_{0}\left(z^{i}\right)}{\lambda \xi_{1}+(1-\lambda) \xi_{2}}\right)
\end{aligned}
$$

This relation, however, is implied by the concavity of $J_{k+1}$.
Using Eqs. (5.51) and (5.52), we obtain (see Fig. 5.4.6) that if

$$
C+A_{N-2}\left(\frac{L_{0}}{L_{0}+L_{1}}\right)<\frac{L_{0} L_{1}}{L_{0}+L_{1}}
$$

then an optimal policy for each period $k$ is of the form

$$
\begin{aligned}
& \text { accept } f a \text { if } P k \geq \alpha_{k} \\
& \text { accept } f_{1} \text { if } P k \leq \beta_{k}
\end{aligned}
$$

continue the observations if $\beta_{k}<P k<\alpha_{k}$,
where the scalars $\alpha_{k}, \beta_{k}$ are determined from the relations

$$
\begin{aligned}
\beta_{k} L_{1} & =\mathrm{C}+A k(13 k) \\
\left(1-\alpha_{k}\right) L_{0} & =\mathrm{C}+A_{k}\left(\alpha_{k}\right)
\end{aligned}
$$

Furthermore, we have

$$
\begin{gathered}
\cdots \leq \alpha_{k+1} \leq \alpha_{k} \leq \alpha_{k-1} \leq \cdots \leq 1-\frac{C}{T} \\
\ldots \geq \beta_{k+1} \geq \beta_{k} \geq \beta_{k-1} \geq \ldots \geq \underline{C}
\end{gathered}
$$

Hence as $N \rightarrow$ oo the sequences $\left\{\alpha_{N}-i_{i}\right\},\left\{\beta_{N}-{ }^{\prime} L\right\}$ converge to scalars $\bar{\alpha}, \bar{\beta}$, respectively, and the optimal policy is approximated by the stationary policy

$$
\begin{align*}
& \text { accept } f a \text { if } P k \geq \bar{\alpha} \\
& \text { accept } f_{1} \text { if } P k \leq \beta \tag{5.55}
\end{align*}
$$

continue the observations if $\beta<P k<\bar{\alpha}$.
Now the conditional probability $P k$ is given by

$$
\begin{equation*}
p_{k}=\frac{p f_{0}\left(z_{0}\right) \cdots f_{0}\left(z_{k}\right)}{p f_{0}\left(z_{0}\right) \cdots f_{0}\left(z_{k}\right)+(1-p) f_{1}\left(z_{0}\right) \cdots f_{1}\left(z_{k}\right)} \tag{5.56}
\end{equation*}
$$

where $P$ is the a priori probability that $f a$ is the true hypothesis. Using Eq. (5.56), the stationary policy (5.53)-(5.55) can be written in the form

$$
\begin{align*}
& \text { accept } f a \text { if } R k \geq A  \tag{5.57}\\
& \text { accept } f_{1} \text { if } R k \leq B \tag{5.58}
\end{align*}
$$

continue the observations if $B<R k<A$,
where

$$
A=\frac{(1-p) \bar{\alpha}}{p(1-\bar{\alpha})}
$$


p

Figure 5.4.6 Determining the optimal hypothesis testing policy.

$$
B=\frac{(1-p) \bar{\beta}}{p(l-\bar{\beta})}
$$

and

$$
R k=\frac{!O(z o)}{f_{1}\left(z_{0}\right)} \cdots \frac{!O(Z k)}{f_{1}\left(z_{k}\right)}
$$

Note that $R k$ can be easily generated by means of the recursive equation

$$
R_{k+1}=\frac{f_{0}\left(z_{k+1}\right)}{f_{1}\left(z_{k+1}\right)} R_{k}
$$

The policy (5.57)-(5.59) is known as the sequential probability ratio test, an.d was among the first formal methods studied in statistical sequential analySIS by Wald [WaI47]. The optimality of this policy for the infinite horizon version of the problem will be shown in Vol. II, Chapter 3.

### 5.5 NOTES, SOURCES, AND EXERCISES

For literature on linear-quadratic problems with imperfect state information, see the references quoted for Section 4.1 and the survey paper by Witsenhausen [vVit71]. The Kalman filtering algorithm is discussed in many textbooks, such as Anderson and Moore [AnM79], Ljung and Soderstrom [LjS83]. For linear-quadratic problems with Gaussian uncertainties and observation cost in the spirit of Exercise 5.6, see Aoki and Li [AoL69]. Exercise $\overline{5} .1$, which indicates the form of the certainty equivalence principle when the random disturbances are correlated, is based on an unpublished report by the author [Ber70]. The minimum variance approach IS described in Astrom and Wittenmark [AsW84], and Whittle [Whi63].

Problems with exponential cost functions are discussed in James, Baras, and Elliott [JBE94], and Fernandez-Gaucherand and Markus [FeM94]

The idea of data reduction via a sufficient statistic gained wide attention following the paper by Striebel [Str65]; see also Shiryaev [Shi64], [Shi66]. For the analog of the sufficient statistic idea in sequential minimax problems with set membership description of the uncertainty, see Bertsekas and Rhodes [BeR73].

Sufficient statistics have been used for analysis of finite-state problems with imperfect state information by Eckles [Eck68], and by Smallwood and Sondik [SmS73], [Son71]. The proof of piecewise linearity of the cost-to-go functions, and an algorithm for their computation are given by Smallwood and Sondik [SmS73], [Son71]. For further material on finite-state problems with imperfect state information, see Arapostathis et. aL [ABF93], Lovejoy [Lov91a], [Lov91b], and White and Scherer [WhS89].

The instruction model described in Example 5.4.3 has been considered (with some variations) by a number of authors such as Atkinson, Bower and Crothers [ABC65], Groen and Atkinson [GrA66], Karush and Dear [KaD66], and Smallwood [Sma71].

For a discussion of the sequential probability ratio test (cf. Example 5.4.4) and related subjects, see Chernoff [Che72], DeGroot [DeG70], [Whi82], and the references quoted therein. The treatment given here stems from Arrow, Blackwell, and Girshick [ABG49].

## EXERCISES

5.1 (Linear Quadratic Problems - Correlated Distmrbances)

Consider the linear system and measurement equation of Section 5.2 and consider the problem of finding a policy $\left\{\mu_{0}^{*}\left(I_{0}\right), \ldots, \mu_{N-1}^{*}\left(I_{N-1}\right)\right\}$ that minimizes the quadratic cost

$$
E\left\{x_{N}^{\prime} Q x_{N}+\sum_{k=0}^{N-1} u_{k}^{\prime} R_{k} u_{k}\right\} .
$$

Assume, however, that the random vectors $X o, W a, \ldots, W_{N-1}, v a, \cdot \cdot, V_{N}-1$ are correlated and have a given joint probability distribution, and finite first and second moments. Show that the optimal policy is given by

$$
\mu_{k}^{*}\left(I_{k}\right)=L_{k} E\left\{y_{k} \mid I_{k}\right\},
$$

where the gain matrices $L k$ are obtained from the algorithm

$$
L k=-\left(B_{k}^{\prime} K_{k+1} B_{k}+R k\right)-l B_{k}^{\prime} K_{k+1} A_{k}
$$

$K N=Q$,
$K k \quad A_{k}^{\prime}\left(K_{k+1}-K_{k+1} B_{k}\left(B_{k}^{\prime} K_{k+1} B_{k}+R_{k}\right)^{-1} B_{k}^{\prime} K_{k+1}\right) A_{k}$,
and the vectors $Y k$ are given by

$$
Y k=X k+A_{k}^{-1} w_{k}+A_{k}^{-1} A_{k+1}^{-1} W k+l+\ldots+A_{k}^{-1} \ldots A_{N-1}^{-1} W N-l
$$

(assuming the matrices $A_{O}, \mathrm{~A}_{1}, \ldots, A_{N}-1$ are invertible). Hint: Show that the cost can be written as

$$
E\left\{y_{0}^{\prime} K_{0} y_{0}+\sum_{k=0}^{N-1}\left(u_{k}-L_{k} y_{k}\right)^{\prime} P_{k}\left(u_{k}-L_{k} y_{k}\right)\right\}
$$

where

$$
P_{k}=B_{k}^{\prime} K_{k+1}^{-} B_{k}+R_{k}
$$

## 5.2

Consider the scalar system

$$
\begin{aligned}
x_{k+1} & =x_{k}+u_{k}+w_{k} \\
z_{k} & =x_{k}+v_{k}
\end{aligned}
$$

where we assume that the initial condition $X o$, and the disturbances $W k$ and $V k$ are all independent. Let the cost be

$$
E\left\{x_{N}^{2}+\sum_{k=0}^{N-1}\left(x_{k}^{2}+u_{k}^{2}\right)\right\}
$$

and let the given probability distributions be

$$
\begin{gathered}
p(x o=2)=\frac{1}{2}, \quad P(W k=1)=\frac{1}{2}, \quad p\left(V k=\frac{1}{4}\right)=\frac{1}{2} \\
p(x o=-2)=\frac{1}{2}, \quad p(W k=-1)=\frac{1}{2}, \quad P\left(V k=-\frac{1}{4}\right)=\frac{1}{2} .
\end{gathered}
$$

(a) Determine the optimal policy. Hint: For this problem, $E\left\{X k \mid I_{k}\right\}$ can be determined from $E\left\{X k-l \mid I_{k-1}\right\}, U k-l$, and $Z k$.
(b) Determine the policy that is identical to the optimal except that it uses a linear least squares estimator of $x k$ given $I k$ in place of $E\{x k I l k\}$ (see Appendix $\mathrm{E}-$ this policy can be shown to be optimal within the class of all policies that are linear functions of the measurements).
(c) Determine the asymptotic form of the policies in parts (a) and (b) as $N \rightarrow$ $\infty$.
5.3

A linear system with Gaussian disturbances and Gaussian initial state

$$
x_{k+1}=A x_{k}+B u_{k}+w_{k},
$$

is to be controlled so as to minimize a quadratic cost similar to that in Section 5.2. The difference is that the controller has the option of choosing at each time $k$ one of two types of measurements for the next stage $(k+1)$ :

$$
\begin{aligned}
\text { first type: } & Z k+l= \\
\text { second type: } & Z k+l=C 2 X k+l+
\end{aligned}
$$

Here $\mathrm{C}^{1}$ and $C^{2}$ are given matrices of appropriate dimension, and $\left\{v_{k}^{1}\right\}$ and $\left\{v_{k}^{2}\right\}$ are zero-mean, independent, random sequences with given finite covariances that do not depend on $X o$ and $\{W k\}^{\prime}$ There is a cost gl (or $g 2$ ) each time a measurement of type 1 (or type 2 ) is taken. The problem is to find the optimal control and measurement selection policy that minimizes the expected value of the sum of the quadratic cost

$$
x_{N}^{\prime} Q x_{N}+\sum_{k=0}^{N-l}\left(x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}\right)
$$

and the total measurement cost. Assume for convenience that $N=2$ and that the first measurement $Z o$ is of type 1 . Show that the optimal measurement selection at $k=0$ and $k \quad 1$ does not depend on the value of the information vectors 10 and $1_{1}$, and can be determined a priori. Describe the nature of the optimal policy

## 5.4

Consider a scalar single-input, single-output system given by

$$
\begin{aligned}
Y k+a l Y k-l+\ldots+a m Y k-m= & b M U k-M+\ldots+b m U k-m \\
& +\epsilon_{k}+c_{1} \epsilon_{k-1}+\ldots+C m t k-m+V k-n
\end{aligned}
$$

where $1 \leq M \leq \mathrm{m}, 0 \leq n \leq \mathrm{m}$, and $V k$ is generated by an equation of the form

$$
v_{k}+d_{1} v_{k-1}+\cdots+d_{m} v_{k-m}=\xi_{k}+\epsilon_{1} \xi_{k-1}+\cdots+\epsilon_{m} \xi_{k-m}
$$

and the polynomials $\left(1+C l S+"+c m S^{m}\right),\left(l+d l s+\cdots+d m s^{m}\right)$, and $(1+\mathrm{tlS}+$ $\ldots+\epsilon_{m} s^{m}$ ) have roots strictly outside the unit circle. The value of the scalar $V k$ is observed by the controller at time $k$ together with $Y k$. The sequences $\left\{\epsilon_{k}\right\}$ and $\left\{\xi_{k}\right\}$ are zero mean independent identically distributed with finite variances. Find an easily implementable approximation to the minimum variance controller minimizing $E\left\{\sum_{k=0}^{N}\left(y_{k}\right)^{2}\right\}$. Discuss the stability properties of the closed-loop system.

## (Finite-State Systems Imperfect State Information)

(a) Within the framework of the basic problem with imperfect state information, consider the case where the system and the observations are linear:

$$
\begin{gathered}
X k+l=A k X k+B k U k+W k \\
Z k=C k X k+V k
\end{gathered}
$$

The initial state $X o$ and the disturbances $W k$ and $V k$ are assumed Gaussian and mutually independent. Their covariances are given, and $W k$ and $V k$ have zero mean. Show that $E\{x o \mid I o\}, \ldots, E\{X N-l \mid I N-I\}$ constitute a sufficient statistic for this problem.
(b) Use the result of part (a) to obtain an optimal policy for the special case of the single-stage problem involving the scalar system and observation

$$
\begin{aligned}
& \mathrm{x} 1=X o+u o \\
& Z o=X o+V a
\end{aligned}
$$

and the cost function $E\{l x k l\}$.
(c) Generalize part (b) for the case of the scalar system

$$
\begin{gathered}
x_{k+1}=a x_{k}+u_{k} \\
Z k=C X k+V k
\end{gathered}
$$

and the cost function $E\left\{\sum_{k=1}^{N}\left|x_{k}\right|\right\}$. The scalars $a$ and c are given. Note: You may find useful the following "differentiation of an integral" formula:

$$
\begin{aligned}
& \frac{d}{d y}_{\alpha(y)}^{j^{\prime}(3(Y)} f(y, \xi) d \xi= J\left(3(y) \frac{d I f(y, \xi)}{d \psi} d \xi\right. \\
& \alpha(y) \\
&+f\left(Y,(3(y)) \frac{d \beta(y)}{d y}-f(y, \alpha(y)) \frac{d \alpha(y)}{d y}\right.
\end{aligned}
$$

## 5.6

Consider a machine that can be in one of two states, good or bad. Suppose that the machine produces an item at the end of each period. The item produced is either good or bad depending on whether the machine is in a good or bad state at the beginning of the corresponding period, respectively. We suppose that once the machine is in a bad state it remains in that state until it is replaced. If the machine is in a good state at the beginning of a certain period, then with probability $t$ it will be in the bad state at the end of the period. Once an item is produced, we may inspect the item at a cost $l$ or not inspect. If an inspected item is found to be bad, the machine is replaced with a machine in good state at a cost $R$. The cost for producing a bad item is $\mathrm{C}>0$. Write a DP algorithm for obtaining an optimal inspection policy assuming a machine initially in good state and a horizon of N periods. Solve the problem for $t=0.2, I=1, R=3$, C $\quad 2$, and $\mathrm{N}=8$. (The optimal policy is to inspect at the end of the third period and not inspect in any other period.)

Consider a system that at any time can be in anyone of a finite number of states $1,2, \ldots, n$. When a control $U$ is applied, the system moves from state i to state j with probability $P i j(u)$. The control $U$ is chosen from a finite collection $u^{1}, u^{2}, \ldots \bullet, u^{m}$. Following each state transition, an observation is made by the controller. There is a finite number of possible observation outcomes $z 1, z 2, \ldots, z q$. The probability of occurrence of $z^{\theta}$, given that the current state is j and the preceding control was $u$, is denoted by $r_{j}(\boldsymbol{U}, \boldsymbol{\theta}), \boldsymbol{\theta}=1, \ldots, q$.
(a) Consider the column vector of conditional probabilities

$$
P_{k}=\left[p_{k}^{1}, \ldots, p_{k}^{n}\right]^{\prime}
$$

where

$$
p_{k}^{j}=P(X k=\mathrm{j} \quad \mathrm{I} Z o, \cdots, Z k, U o, \cdot \cdot, U k-I), \quad \mathrm{j}=1, \ldots, n
$$

and show that it can be updated according to

$$
p_{k+1}^{j}=\frac{\sum_{i=1}^{n} p_{k}^{i} p_{i j}\left(u_{k}\right) r_{j}\left(u_{k}, z_{k+1}\right)}{\sum_{s=1}^{n} \sum_{i=1}^{n} p_{k}^{i} p_{i s}\left(u_{k}\right) r_{s}\left(u_{k}, z_{k+1}\right)}, \quad \mathrm{j}=1, \ldots, n .
$$

Write this equation in the compact form

$$
P k+l=\frac{[r(U k, Z k+l)]}{r\left(U k^{\prime} z_{k+1}\right)^{\prime} P} \frac{\left[P(U k)^{\prime} P k\right]}{U k)^{\prime} P k}
$$

where prime denotes transposition and
$P(U k)$ is the $n \times n$ transition probability matrix with ijth element $\operatorname{Pij}(U k)$,
$r\left(u k^{\prime} Z k+l\right)$ is the column vector with $\mathbf{j}$ th coordinate $r_{j}\left(u_{k}, Z k+l\right)$,
$\left[P(U k)^{\prime} P k\right]_{j}$ is the j th coordinate of the vector $P(U k)^{\prime} P k$,
$[r(U k, Z k+l)]^{*}\left[P(U k)^{\prime} P_{k}\right]$ is the vector whose j th coordinate is the scalar $r j(u k, Z k+l)\left[P(U k)^{\prime} P k\right]_{j}$.
(b) Assume that there is a cost for each stage $k$ denoted $9 k(i, u, j)$ and associated with the control $U$ and a transition from i to j . There is no terminal cost. Consider the problem of finding a policy that minimizes the sum of expected costs per stage over N periods. Show that the corresponding DP algorithm is given by

$$
\begin{aligned}
& \bar{J} N-l(P N-. l)= \\
& \min _{u E\left\{u l, \ldots, u^{m}\right\}} P_{N-1}^{\prime} G_{N-1}(u) \\
& \bar{J}_{k}\left(P_{k}\right)= \min _{u \in\left\{u^{1}, \ldots, u^{m}\right\}}\left[P_{k}^{\prime} G_{k}(u)\right. \\
&\left.+\sum_{\theta=1}^{q} r(u, \theta)^{\prime} P(u)^{\prime} P_{k} \bar{J}_{k+1}\left(\frac{[r(u, \theta)] *\left[P(u)^{\prime} P_{k}\right]}{r(u, \theta)^{\prime} P(u)^{\prime} P_{k}}\right)\right]
\end{aligned}
$$

where $\mathrm{G}_{k}\left({ }^{\prime} u\right)$ is the vector of expected kth stage costs given by

$$
G_{k}(u)=\left(\begin{array}{c}
\sum_{j=1}^{n} p_{1 j}(u) g_{k}(1, u, j) \\
\vdots \\
\sum_{j=1}^{n} p_{n j}(u) g_{k}(n, u, j)
\end{array}\right)
$$

(c) Show that, for all $k, J k$ when viewed as a function defined on the set of vectors with nonnegative coordinates, is positively homogeneous; that is,

$$
\bar{J}_{k}\left(\lambda P_{k}\right)=\lambda \bar{J}_{k}\left(P_{k}\right)
$$

for all $\lambda>0$. Use this fact to write the DP algorithm in the simpler form

$$
\bar{J}_{k}\left(P_{k}\right)=\min _{u \in\left\{u^{1}, \ldots, u^{m}\right\}}\left[P_{k}^{\prime} G_{k}(u)+\sum_{\theta=1}^{q} \bar{J}_{k+1}\left([r(u, \theta)] *\left[P(u)^{\prime} P_{k}\right]\right)\right]
$$

(d) Show by induction that, for all $k, \bar{J}_{k}$ has the form

$$
\bar{J}_{k}\left(P_{k}\right)=\min \left[P_{k}^{\prime} \alpha_{k}^{1}, \ldots, P_{k}^{\prime} \alpha_{k}^{m_{k}}\right]
$$

where $\alpha_{k}^{1}, \ldots, \alpha_{k}^{m_{k}}$ are some vectors in $\Re^{n}$.

## 5.8

Consider the functions $] k(P k)$ in the instruction problem of Example 5.4.3. Show inductively that each of these functions is piecewise linear, concave, and of the form

$$
\left.\bar{J}_{k} \Phi k=\operatorname{mm}^{\left[\alpha_{k}^{\prime}\right.}+\beta_{k}^{1} p_{k}, \alpha_{k}+\beta_{k}^{2} P k, \ldots, \alpha_{k}^{m} k+\beta_{k}^{m_{k}} P k\right]
$$

where $\alpha_{k}^{1}, \ldots, \alpha_{k}^{m_{k}}, \beta_{k}^{1}, \ldots, \beta_{k}^{m_{k}}$ are suitable scalars.

## 5.9

A person is offered $N$ free plays to be distributed as he pleases between two slot machines A and B. Machine A pays $\alpha$ dollars with known probability sand nothing with probability $(1-s)$. Machine B pays $\beta$ dollars with probability $P$ and nothing with probability $(1-p)$. The person does not know $P$ but instead has an a priori probability distribution $F(p)$ for $p$. The problem is to find a playing policy that maximizes expected profit. Let $(m+n)$ denote the number of plays in machine B after $k$ free plays $(\mathrm{m}+n \leq k)$, and let $m$ denote the number of successes and $n$ the number of failures. Show that a DP algorithm for this problem is given for $\mathrm{m}+n \leq k$ by
$] N-l(m, n)=\max [s \alpha, p(m, n) f J]$,

$$
\begin{aligned}
J k(m, n)=\max & {[s(\alpha+J k+l(m, n))+(1-S)] k+l(m, n) } \\
& p(m, n)(\beta+] \mathrm{k}+1(\mathrm{~m}+1, n))+(1-p(m, n))] \mathrm{k}+1(\mathrm{~m}, n+1)]
\end{aligned}
$$

where

$$
p(m, n)=\frac{\mathrm{pm}+1(1-p) n d F(p)}{p^{m}(1-p) n d F(p)}
$$

Solve the problem for $N=6, \alpha=\beta=1, s=0.6, d F(p) / d p=1$ for $0 \leq p \leq \mathrm{L}$ [The answer is to play machine $B$ for the following pairs $(m, n):(0,0),(1,0)$, $(2,0),(3,0),(4,0),(5,0),(2,1),(3,1),(4,1)$. Otherwise, machine A should be played.]

### 5.10

A person is offered 2 to 1 odds in a coin-tossing game where he wins whenever a tail occurs. However, he suspects that the coin is biased and has an a priori probability distribution $F(p)$ for the probability $p$ that a head occurs at each toss. The problem is to find an optimal policy of deciding whether to continue or stop participating in the game given the outcomes of the game so far. A maximum of $N$ tossings is allowed. Indicate how such a policy can be found by means of DP.

### 5.11

Consider the ARMAX model of Section 5.3 where instead of $E\left\{\sum_{k=1}^{N}\left(y_{k}\right)^{2}\right\}$, the cost is

$$
E\left\{\sum_{k=1}^{N}\left(y_{k}-\bar{y}\right)^{2}\right\}
$$

where $\bar{y}$ is a given scalar. Generalize the minimum variance policy for this case.

### 5.12

Consider the ARMAX model

$$
Y k+a Y k-l=U k-M+\mathrm{Ek},
$$

where $M \geq \mathrm{L}$ Show that the minimum variance controller is

$$
\mu_{k}\left(I_{k}\right)=a U k-l-a^{2} u k_{-} 2+.,-\left(\_l\right) M-l_{a} M-l_{u k-} M+l-(-1)^{M} a M Y k
$$

that the resulting closed-loop system is

$$
Y k=E k-a \epsilon_{k-1}+a^{2} E k \_2-' "+(-l) M-l a M-l E k \_M+l,
$$

and that the long-term output variance is

$$
E\{(Y k) 2\}=\frac{]-a^{2 M}}{1-a_{2}} E\{(\mathrm{Ek}) 2\}
$$

Discuss the qualitative difference between the cases $\mid a l<1$ and $l a l>1$, and relate it to the stability properties of the uncontrolled system $Y k+a Y k-l=\mathrm{Ek}$ and the size of the delay $M$.

### 5.13 (Linear-Quadratic Problems with Disturbance Estimation)

Consider the linear-quadratic problem discussed in Section $4.1\left(\boldsymbol{A}_{k}, \boldsymbol{B}_{k}\right.$ : known). The state $X k$ is perfectly observed at each stage, and the demands $W k$ are independent, identically distributed random vectors. However, the (common) distribution of the $W k$ is unknown. Instead it is known that this distribution is one out of two given distributions $F_{1}$ and $F_{2}$, and that tow a priori probability that $\mathrm{F}_{1}$ is the correct distribution is a given scalar $q$, with $<q<1$. For convenience, assume that $W k$ can take a finite number of values under each of $F_{1}$ and $F_{2}$.
(a) Formulate this as an imperfect state information problem, and identify the state, control, system disturbance, observation, and observation disturbance.
(b) Show that $(X k, q k)$, where

$$
q k=\mathrm{P}\left(\text { distribution is } \mathrm{F}_{\mathrm{l}} \mathrm{I} W o, \cdots, W k-I\right),
$$

is a suitable sufficient statistic, and write a corresponding DP algorithm.
(c) Show that the optimal control law is of the form

$$
\mu_{k}\left(x_{k}, q_{k}\right)=-\left(B_{k}^{\prime} K_{k+1} B_{k}+R_{k}\right)^{-1} B_{k}^{\prime} K_{k+1} A_{k} x_{k}+c_{k}\left(q_{k}\right),
$$

where the matrices $K_{k}$ are given by the Riccati equation, and $C k(q k)$ are appropriate functions of $q k$. Hint: Show that the cost-to-go function has the form

$$
J_{k}\left(x_{k}, q k\right)=x_{k}^{\prime} K_{k} x_{k}+a k(q k)^{\prime} X k+b k(q k),
$$

where $a_{k}\left(q_{k}\right)$ and $b k(q k)$ are appropriate functions of $q k$.

### 5.14 (Asset Selling Problem with Offer Estimation)

Consider the asset selling problem of Section 4.4. The offers $W k$ are independent and identically distributed. However, the (common) distribution of the $W k$ is unknown. Instead it is known that this distribution is one out of two given distributions PI and $F_{2}$, and that thoa priori probability that FI is the correct distribution is a given scalar $q$, with $<q<1$.
(a) Formulate this as an imperfect state information problem, and identify the state, control, system disturbance, observation, and observation disturbance.
(b) Show that $(X k, q k)$, where

$$
q k=\mathrm{P}\left(\text { distribution is } \mathrm{F}_{\mathrm{I}} \mathrm{I} W o, \cdots, W k-I\right)
$$

is a suitable sufficient statistic, write a corresponding DP algorithm, and derive the form of the optimal selling policy.

Consider an inventory control problem where stock evolves according to

$$
x_{k+1}=x_{k}+u_{k}-w_{k}
$$

and the cost of stage $k$ is

$$
C U k+h \max (O, W k-X k-U k)+\operatorname{pmax}(O, X k+U k-W k)
$$

where $\mathrm{c}, h$, and $p$ are positive scalars with $p>c$. There is no terminal cost. The stock $X k$ is perfectly observed at each stage. The demands $W k$ are independent, identically distributed, nonnegative random variables. However, the (common) distribution of the $W k$ is unknown. Instead it is known that this distribution is one out of two given distributions FI and $F_{2}$, and that tho priori probability that $\mathrm{F}_{1}$ is the correct distribution is a given scalar $q$, with $<q<1$.
(a) Formulate this as an imperfect state information problem, and identify the state, control, system disturbance, observation, and observation disturbance.
(b) Write a DP algorithm in terms of a suitable sufficient statistic.
(c) Characterize as best as you can the optimal policy.

Consider the search problem of Example 5.4.1 for different values of the search horizon $N$.
(a) Show that for any value of the a priori probability $P o$ that is strictly less than 1 , there is a threshold value of $N$, call it $\underline{N}$, such that the optimal reward function $J o(P o)$ is independent of $N$ as long as $N>\underline{N}$.
(b) For $N$ greater than the threshold $\underline{N}$ of part (a) and for a given value of $p o$, give a method to calculate the value of $J o(p o)$ that does not use the DP algorithm.
(c) Suppose that there are two sites that can be searched, instead of one. The sites may contain a treasure of corresponding values VI and $V^{2}$ (independently of each other), and the probabilities of a successful search are $\beta_{1}$ and $l 32$, respectively. After finding a treasure in one site, one may continue searching for the treasure in the other site (but of course each search costs C). Write a DP algorithm involving the probabilities $p_{k}^{1}$ and $p_{k}^{2}$ that a treasure is present at sites 1 and 2 , respectively.
(d) Under the assumptions of part (c), show that for any values of the a priori probabilities $p_{0}^{1}, p_{0}^{2}$, there is a threshold value of $N$, call it $N$, such that the optimal cost-to-go function $J_{0}\left(p_{0}^{1}, p_{0}^{2}\right)$ is independent of $N$ as long as $N>\underline{N}$. Find the optimal search policy if $N>\underline{N}$.

## Approximate

## Dynamic Programming

| Contents |  |
| :---: | :---: |
| 6.1. Certainty Equivalent and Adaptive Control | p. 283 |
| 6.1.1. Caution, Probing, and Dual Control | p. 289 |
| 6.1.2. Two-Phase Control and Identifiability | p. 291 |
| 6.1.3. Certainty Equivalent Control and Identifiability | p. 293 |
| 6.1.4. Self-Tuning Regulators | p. 298 |
| 6.2. Open-Loop Feedback Control. | p. 300 |
| 6.3. Limited Lookahead Policies | p. 304 |
| 6.3.1. Performance Bounds for Limited Lookahead Policies | p. 305 |
| 6.3.2. Computational Issues in Limited Lookahead | p. 310 |
| 6.3.3. Problem Approximation - Enforced Decomposition | p. 312 |
| 6.3.4. Aggregation - | p. 319 |
| 6.3.5. Parametric Cost-to-Go Approximation | p. 325 |
| 6.4. Rollout Algorithms | p. 335 |
| 6.4.1. Discrete Deterministic Problems. | p. 342 |
| 6.4.2. $Q$-Factors Evaluated by Simulation | p. 361 |
| 6.4.3. Q-Factor Approximation | p. 363 |
| 6.5. Model Predictive Control and Related Methods | p. 366 |
| 6.5.1. Rolling Horizon Approximations | p. 367 |
| 6.5.2. Stability Issues in Model Predictive Control | p. 369 |
| 6.5.3. Restricted Structure Policies | p. 376 |
| 6.6. Additional Topics in Approximate DP | p. 382 |
| 6.6.1. Discretization. | p. 382 |
| 6.6.2. Other Approximation Approaches | p. 384 |
| 6.7. Notes, Sources, and Exercises. | p. 386 |

We have seen that it is sometimes possible to use the DP algorithm to obtain an optimal policy in closed form. However, this tends to be the exception. In most cases a numerical solution is necessary. The associated computational requirements are often overwhelming, and for many problems a complete solution of the problem by DP is impossible. To a great extent, the reason lies in what Bellman has called the "curse of dimensionality." This refers to an exponential increase of the required computation as the problem's size increases.

Consider for example a problem where the state, control, and disturbance spaces are the Euclidean spaces $\Re^{n}, \Re^{m}$, and $\Re^{r}$, respectively. In a straightforward numerical approach, these spaces are discretized. Taking $d$ discretization points per state axis results in a state space grid with $d^{n}$ points. For each of these points the minimization in the right-hand side of the DP equation must be carried out numerically, which involves comparison of as many as $d^{m}$ numbers. To calculate each of these numbers, one must calculate an expected value over the disturbance, which is the weighted sum of as many as $d^{\prime}$ numbers. Finally, the calculations must be done for each of the $N$ stages. Thus as a first approximation, the number of computational operations is at least of the order of $N d^{n}$ and can be of the order of $N d^{n}+^{m} t^{r}$. It follows that for perfect state information problems with Euclidean state and control spaces, DP can be applied numerically only if the dimensions of the spaces are relatively small. Based on the analysis of the preceding chapter, we can also conclude that for problems of imperfect state information the situation is hopeless, except for very simple or very special cases.

In the real world, there is an additional aspect of optimal control problems that can have a profound impact on the feasibility of DP as a practical solution method. In particular, there are many circumstances where the structure of the given problem is known well in advance, but some of the problem data, such as various system parameters, may be unknown until shortly before control is needed, thus -seriously constraining the amount of time available for the DP computation. Usually this occurs as a result of one or both of the following situations:
(a) A family of problems is addressed, rather than a single problem, and we do not get to know the exact problem to be solved until shortly before the control process begins. As an example, consider a problem of planning the daily route of a utility vehicle within a street network so that it passes through a number of points where it must perform some service. The street network and the vehicle characteristics may be known well in advance, but the service points may vary from day to day, and may not become known until shortly before the vehicle begins its route. This example is typical of situations, where the same problem must periodically be solved with small variations in its data. Yet, if DP is to be used, the solution of one instance of the problem
may not help appreciably in solving a different instance.
(b) The problem data changes as the system is being controlled. As an example, consider the route planning example in case (a) above, and assume that new service points to be visited arise as the vehicle is on its way. It is possible in principle to model these data changes in terms of stochastic disturbances, but then we may end up with a problem that is too complicated for analysis or solution by DP. A frequently employed alternative is to use on-line replanning, whereby the problem is resolved on-line with the new data, as soon as these data become available, and control continues with a policy that corresponds to the new data.
A common feature of the above situations, which can seriously impact the solution, is that there may be stringent time constraints for the computation of the controls. This may substantially exacerbate the "curse of dimensionality" problem mentioned above.

As indicated by the above discussion, in practice one often has to settle for a suboptimal control scheme that strikes a reasonable balance between convenient implementation and adequate performance. In this chapter we discuss some general approaches for suboptimal control, which are based on approximations to the DP algorithm. We begin with two general schemes to simplify the DP computation, certainty equivalent control (Section 6.1), which replaces the stochastic quantities of the problem by deterministic nominal values, and open-loop-feedback control (Section 6.2), which ignores in part the availability of information in the future. These two schemes set the stage for limited lookahead control, which together with its many variations (Sections 6.3-6.5), is one of the principal approaches for suboptimal control. We also discuss adaptive control in the context of certainty equivalent control. This discussion is not used in subsequent developments, so the reader may skip Sections 6.1.1-6.1.4 if desired.

### 6.1 CERTAINTY EQUIVALENT AND ADAPTIVE CONTROL

The certainty equivalent controller (CEC) is a suboptimal control scheme that is inspired by linear-quadratic control theory. It applies at each stage the control that would be optimal if the uncertain quantities were fixed at some "typical" values; that is, it acts as if a form of the certainty equivalence principle were holding.

The advantage of the CEC is that it replaces the DP algorithm with what is often a much less demanding computation: the solution of a deterministic optimal control problem at each stage. This problem yields an optimal control sequence, the first component of which is used at the current stage, while the remaining components are discarded. The main attractive characteristic of the CEe is its ability to deal with stochastic
and even imperfect information problems by using the mature and effective methodology of deterministic optimal control.

We describe the CEC for the general problem with imperfect state information of Section 5.1. As can be expected, the implementation is considerably simpler if the controller has perfect state information. Suppose that we have an "estimator" that uses the information vector $I_{k}$ to produce a "typical" value $\bar{x}_{k}\left(I_{k}\right)$ of the state. Assume also that for every state-control pair $(X k, U k)$ we have selected a "typical" value of the disturbance, which we denote by $\bar{w}_{k}\left(x_{k}, U k\right)^{\prime}$ For example, if the state spaces and disturbance spaces are convex subsets of Euclidean spaces, the expected values

$$
\bar{x}_{k}\left(I_{k}\right)=E\left\{x_{k} \mid I_{k}\right\}, \quad \bar{w}_{k}\left(x_{k}, u_{k}\right)=E\left\{u_{k} \mid x_{k}, u_{k}\right\}
$$

can serve as typical values.
The control input $\bar{\mu}_{k}\left(I_{k}\right)$ applied by the CEC at each time $k$ is determined by the following rule:
(1) Given the information vector $I_{k}$, compute the state estimate $x k(I k)$.
(2) Find a control sequence $\left\{\bar{u}_{k}, \bar{u}_{k+1}, \ldots, \bar{u}_{N-1}\right\}$ that solves the deterministic problem obtained by fixing the uncertain quantities $X k$ and $W k, \ldots, W N-I$ at their typical values:

$$
\begin{aligned}
& \text { N-I } \\
& \operatorname{minimize} g N(X N)+\underset{i=k}{ } g i(X i, U i, W i(X i, U i)) \\
& i=k
\end{aligned}
$$

subject to the initial condition $X k=\bar{x}_{k}\left(I_{k}\right)$ and the constraints
Ui $\mathrm{E} \quad x_{i+1}=f_{i}\left(x_{i}, u_{i}, \bar{w}_{i}\left(x_{i}, u_{i}\right)\right), \quad \mathrm{i}=k, k+1, \ldots, N-1$.
(3) Use as control the first element in the control sequence found:

$$
\bar{\mu}_{k}\left(I_{k}\right)=\bar{u}_{k}
$$

Note that step (1) is unnecessary if we have perfect state information; in this case we simply use the known value of the $X k$. The deterministic optimization problem in step (2) must be solved at each time $k$, once the initial state $\bar{x}_{k}\left(I_{k}\right)$ becomes known by means of an estimation (or perfect observation) procedure. A total of $N$ such problems must be solved by the CEC at every system run. In many cases of interest, these deterministic problems can be solved by powerful numerical methods such as conjugate gradient, Newton's method, augmented Lagrangian, and sequential quadratic programming methods; see e.g. Luenberger [Lue84] or Bertsekas [Ber99]. Furthermore, the implementation of the CEC requires no storage of the type required for the optimal feedback controller.

An alternative to solving $N$ optimal control problems in an "on-line" fashion is to solve these problems a priori. This is accomplished by computing an optimal feedback controller for the deterministic optimal control problem obtained from the original problem by replacing all uncertain quantities by their typical values. It is easy to verify, based on the equivalence of open-loop and feedback implementation of optimal controllers for deterministic problems, that the implementation of the CEC given earlier is equivalent to the following.

Let $\left\{\mu_{0}^{d}\left(x_{0}\right), \ldots, \mu_{N-1}^{d}(X N-I)\right\}$ be an optimal controller obtained from the DP algorithm for the deterministic problem

$$
\begin{array}{ll}
\operatorname{minimize} & g N(X N)+\sum_{k=0}^{\mathbf{N}-\mathbf{I}} g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), \bar{w}_{k}\left(x_{k}, u_{k}\right)\right) \\
\text { subject to } X k+1=f_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), \bar{w}_{k}\left(x_{k}, u_{k}\right)\right), \quad \mu_{k}\left(x_{k}\right) \text { E } U k, \quad k \geq 0 .
\end{array}
$$

Then the control input $\bar{\mu}_{k}\left(I_{k}\right)$ applied by the CEC at time $k$ is given by

$$
\bar{\mu}_{k}\left(I_{k}\right)=\mu_{k}^{d}\left(\bar{x}_{k}\left(I_{k}\right)\right)
$$

as shown in Fig. 6.1.1.


Figure 6.1.1 Structure of the certainty equivalent controller when implemented in feedback form.

In other words, an equivalent alternative implementation of the CEC consists of finding a feedback controller $\left\{\mu_{0}^{d}, \mu_{1}^{d}, \ldots, \mu_{N-1}^{d}\right\}$ that is optimal for a corresponding deterministic problem, and subsequently using this controller for control of the uncertain system [modulo substitution of the state $X k$ by its estimate $\left.\bar{x}_{k}\left(I_{k}\right)\right]$. Either one of the definitions given for the $C E C$ can serve as a basis for its implementation. Depending on the nature of the problem, one method may be preferable to the other.

The CEC approach often performs well in practice and yields nearoptimal policies. In fact, for the linear-quadratic problems of Sections 4.1
and 5.2, the CEC is identical to the optimal controller (certainty equivalence principle). It is possible, however, that a CEC performs strictly worse than the optimal open-loop controller (see Exercise 6.2).

In what follows in this section, we will discuss a few variants of the CEC, and we will then focus on one particular type of methodology, adaptive control of systems with unknown parameters.

## cortonter Equivalent Control with Heuristics

Even though the CEC approach simplifies a great deal the computations, it still requires the solution of a deterministic optimal control problem at each stage. This problem may be difficult, and a more convenient approach may be to solve it suboptimally using a heuristic algorithm. To simplify notation, let us assume perfect state information [the ideas to be discussed can also be applied to imperfect state information problems, by substituting $X k$ with its estimate $x k(I k)]$. Then, in this approach, given $X k$, we use some (easily implementable) heuristic to find a suboptimal control sequence $\left\{\bar{u}_{k}, U k+l, " " \bar{u}_{N-1}\right\}$ for the problem

$$
\operatorname{minimize} g N(X N)+\sum_{i=k}^{N-l} g i\left(X i,{ }^{\prime} U i, \bar{w}_{i}\left(x_{i}, U i\right)\right)
$$

subject to

$$
u_{i} \in U_{i}\left(x_{i}\right), \quad x_{i+1}=f_{i}\left(x_{i}, u_{i}, \bar{w}_{i}\left(x_{i}, u_{i}\right)\right), \quad \mathrm{i}=k, k+1, \ldots, \mathrm{~N}-1 .
$$

We then use $\bar{u}_{k}$ as the control for stage $k$.
An important enhancement of this idea is to use minimization over the first control $U k$ and to use the heuristic only for the remaining stages $k+1, \ldots, N-1$. To implement this variant of the CEC, we must apply at time $k$ a control $\bar{u}_{k}$ that minimizes over $U k \mathrm{E} U_{k}\left(x_{k}\right)$ the expression

$$
\begin{equation*}
g_{k}\left(x_{k}, u_{k}, \bar{w}_{k}\left(x_{k}, u_{k}\right)\right)+H_{k+1}\left(f_{k}\left(x_{k}, u_{k}, \bar{w}_{k}\left(x_{k}, u_{k}\right)\right)\right), \tag{6.1}
\end{equation*}
$$

where $H k+l$ is the cost-to-go function corresponding to the heuristic, i.e., $H_{k+l}(X k+l)$ is the cost incurred over the remaining stages $k+1, \ldots, \mathrm{~N}-\mathbf{1}$ starting from a state $X k+l$, using the heuristic, and assuming that the future disturbances will be equal to their typical values $\bar{w}_{i}\left(x_{i}, U i\right)^{\prime}$ Note that for any next-stage state $X k+l$, it is not necessary to have a closedform expression for the heuristic cost-to-go $H k+l(X k+l)$. Instead we can generate this cost by running the system forward from $X k+l$ and accumulating the corresponding single-stage costs. Since the heuristic must be run for each possible value of the control $U k$ to calculate the costs
$\left(f_{k}\left(X k^{\prime} U k, \bar{w}_{k}\left(x_{k},{ }^{\prime} U k\right)\right)\right)$ needed in the minimization, it is necessary to discretize the control constraint set if it is not already finite.

Note that the general structure of the preceding variant of the $C E C$ is similar to the one of standard DP. It involves minimization of the expression (6.1), which is the sum of a current stage cost and a cost-to-go starting from the next state. The difference with DP is that the optimal cost-to-go $J_{k+1}^{*} 1(X k+1)$ is replaced by the heuristic cost $H k+1(X k+1)$, and the disturbance $W k$ is replaced by its typical value $\bar{w}_{k}\left(x_{k}, U k\right)$ (so that there is no need to take expectation over $W k$ ). We thus encounter for the first time an important suboptimal control idea, based on an approximation to the DP algorithm: minimizing at each stage $k$ the sum of approximations to the current stage cost and the optimal cost-to-go. This idea is central in other types of suboptimal control such as the limited lookahead, rollout, and model predictive control approaches, which will be discussed in Sections 6.3-6.5.

## Partially Stochastic Certainty Equivalent Control

In the preceding descriptions of the $C E C$ all future disturbances are fixed at their typical values. A useful variation for some imperfect state information problems is to take into account the stochastic nature of these disturbances, and to treat the problem as one of perfect state information, using an estimate $\bar{x}_{k}\left(I_{k}\right)$ of $X k$ as if it were exact. Thus, if $\left\{\mu_{0}^{p}\left(x_{0}\right), \ldots, \mu_{N-1}^{p}(X N-l)\right\}$ is an optimal policy obtained from the DP algorithm for the stochastic perfect state information problem
minimize $E\left\{9 N(X N)+\sum_{k=0}^{N-1} g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right\}$
subject to $X k+l=f k\left(X k^{\prime} \mu_{k}\left(x_{k}\right), W k\right), \quad \mu_{k}\left(x_{k}\right) \mathrm{E} U k, \quad k=0, \ldots, \mathrm{~N}-1$, then the control input $\bar{\mu}_{k}\left(I_{k}\right)$ applied by this variant of CEC at time $k$ is given by

$$
\bar{\mu}_{k}\left(I_{k}\right)=\mu_{k}^{p}\left(\bar{x}_{k}\left(I_{k}\right)\right) .
$$

Generally, there are several variants of the CEC, where the stochastic uncertainty about some of the unknown quantities is explicitly dealt with, while all other unknown quantities are replaced by estimates obtained in a variety of ways. Let us provide some examples.

## Example 6.1.1 (Multiaccess Communication)

Consider the slotted Aloha system described in Example 5.1.1. It is very difficult to obtain an optimal policy for this problem, primarily because there is no simple characterization of the conditional distribution of the state (the system backlog), given the channel transmission history. We therefore resort to a suboptimal policy. As discussed in Section 5.1, the perfect state information version of the problem admits a simple optimal policy:

$$
\mu_{k}\left(x_{k}\right)=\frac{1}{x_{k}}, \quad \text { for all } x_{k} \geq 1
$$

As a result, there is a natural partially stochastic CEC,

$$
\bar{\mu}_{k}\left(I_{k}\right)=\min \left[1, \frac{1}{\bar{x}_{k}\left(I_{k}\right)}\right]
$$

where $\bar{x}_{k}\left(I_{k}\right)$ is an estimate of the current packet backlog based on the entire past channel history of successes, idles, and collisions (which is $I k$ ). Recursive estimators for generating $\bar{x}_{k}\left(I_{k}\right)$ are discussed by Mikhailov [Mik79], Hajek and van Loon [HaL82], Tsitsiklis [Tsi87], and Bertsekas and Gallager [BeG92].

## Example 6.1.2 (Finite-State Systems with Imperfect

 State Information)Consider the case where the system is a finite-state Markov chain under imperfect state information. The partially stochastic CEC approach is to solve the corresponding problem of perfect state information, and then use the controller thus obtained for control of the imperfectly observed system, modulo substitution of the exact state by an estimate obtained via the Viterbi algorithm described in Section 2.2.2. In particular, suppose that $\left\{\mu_{0}^{p}, \ldots, \mu_{N}^{p}-\mathrm{I}\right\}$ is an optimal policy for the corresponding problem where the state is perfectly observed. Then the partially stochastic CEC, given the information vector $I k$, uses the Viterbi algorithm to obtain (in real time) an estimate $\bar{x}\left(I_{k}\right)$ of the current state $X k$, and applies the control

$$
\bar{\mu}_{k}\left(I_{k}\right)=\mu_{k}^{p}\left(\bar{x}_{k}\left(I_{k}\right)\right)
$$

## Hxample 6.1.3 (Systems with Unknown Parameters)

We have been dealing so far with systems having a known system equation. In practice, however, there are many cases where the system parameters are not known exactly or change over time. One possible approach is to estimate the unknown parameters from input-output records of the system by using system identification techniques. This is a broad and important methodology, for which we refer to textbooks such as Kumar and Varaiya [KuV86], Ljung and Soderstrom [LjS83], and Ljung [Lju86]. However, system identification can be time consuming, and thus difficult to apply in an on-line control context Furthermore, the estimation must be repeated if the parameters change.

The alternative is to formulate the stochastic control problem so that unknown parameters are dealt with directly. It can be shown that problems involving unknown system parameters can be embedded within the framework of our basic problem with imperfect state information by using state augmentation. Indeed, let the system equation be of the form

$$
x_{k+1}=f_{k}\left(x_{k}, \theta, u_{k}, w_{k}\right)
$$

where $\theta$ is a vector of unknown parameters with a given a priori probability distribution. We introduce an additional state variable $Y k \quad \theta$ and obtain a system equation of the form

$$
\begin{gathered}
X k+1) \\
\left(\begin{array}{c}
X k+1
\end{array}\right) \\
(f k(X k, Y k, U k, W k)) \\
Y k
\end{gathered}
$$

This equation can be written compactly as

$$
\tilde{x}_{k+1}=\tilde{f}_{k}\left(\tilde{x}_{k}, U k, W k\right),
$$

where $\tilde{x}_{k}=(X k, Y k)$ is the new state, and $] k$ is an appropriate function. The initial state is

$$
\tilde{x}_{0}=(x o, e)
$$

With a suitable reformulation of the cost function, the resulting problem becomes one that fits our usual framework.

Unfortunately, however, since $Y k$ (i.e., e) is unobservable, we are faced with a problem of imperfect state information even if the controller knows the state $X k$ exactly. Thus, typically an optimal solution cannot be found. Nonetheless, the partially stochastic CEC approach is often convenient. In particular, suppose that for a fixed parameter vector $\ell$, we can compute the corresponding optimal policy

$$
\left\{\mu_{0}^{*}(l a, e), \ldots, \mu_{N-1}^{*}(I N-1, e)\right\}
$$

this is true for example if for a fixed $\ell$, the problem is linear-quadratic of the type considered in Sections 4.1 and 5.2. Then a partially stochastic CEC takes the form

$$
\bar{\mu}_{k}\left(I_{k}\right)=\mu_{k}^{*}\left(I_{k}, \hat{\theta}_{k}\right)
$$

where $\hat{\theta}_{k}$ is some estimate of $\theta$ based on the information vector $I_{k}$. Thus, in this approach, the system is identified while it is being controlled. However, the estimates of the unknown parameters are used as if they were exact.

The approach of the preceding example is one of the principal methods of adaptive control, that is, control that adapts itself to changing values of system parameters. In the remainder of this section, we discuss some of the associated issues. Because adaptive control is somewhat disjoint from other material in the chapter, the reader may skip directly to Section 6.2.

### 6.1.1 Caution, Probing, and Dual Control

Suboptimal control is often guided by the qualitative nature of optimal control. It is therefore important to try to understand some of the characteristic features of the latter in the case where some of the system parameters are unknown. One of these is the need for balance between "caution" (the need for conservatism in applying control, since the system is not fully known), and "probing" (the need for aggressiveness in applying control, in order to excite the system enough to be able to identify it). These notions cannot be easily quantified, but often manifest themselves in specific control schemes. The following example provides some orientation; see also Bar-Shalom [Bar81].

## wxample 6.1.4 [Kum83]

Consider the linear scalar system

$$
x_{k+1}=x_{k}+b u_{k}+w_{k}, \quad k=0,1, \ldots, N-1,
$$

and the quadratic terminal cost $E\left\{\left(x_{\mathrm{N}}\right)^{2}\right\}$. Here everything is as in Section 4.1 (perfect state information) except that the control coefficient $b$ is unknown. Instead, it is known that the a priori probability distribution of $b$ is Gaussian with mean and variance

$$
\bar{b}=E\{b\}>0, \quad \sigma_{b}^{2}=E\left\{(b-\bar{b})^{2}\right\} .
$$

Furthermore, $w k$ is zero mean Gaussian with variance $\sigma_{w}^{2}$ for each $k$.
Consider first the case where $N \quad 1$, so the cost is calculated to be

The minimum over $U o$ is attained at

$$
\text { Uo } \quad-\overline{\bar{b}}^{2}+\frac{\bar{b}}{\bar{b}^{2}} \boldsymbol{X} \boldsymbol{\sigma}
$$

and the optimal cost is verified by straightforward calculation to be

$$
\frac{\sigma_{b}^{2}}{\bar{b}^{2}+\sigma_{b}^{2}} x_{0}^{2}+\sigma_{w}^{2}
$$

Therefore, the optimal control here is cautious in that the optimum luol decreases as the uncertainty in $b$ (Le., $\sigma_{b}^{2}$ ) increases.

Consider next the case where $N=2$. The optimal cost-to-go at stage 1 is obtained by the preceding calculation:

$$
\begin{equation*}
J_{1}\left(I_{1}\right)=\frac{\sigma_{b}^{2}(1)}{(b(l))^{\frac{2}{2}}+\sigma_{b}^{2}(1)^{\mathrm{x} 1}+\sigma_{w}^{2}} \tag{6.2}
\end{equation*}
$$

where $I_{1}=(x o, U o, \mathrm{x} 1)$ is the information vector and

$$
\bar{b}(1)=E\left\{b \mid I_{1}\right\}, \quad \sigma_{b}^{2}(1)=E\left\{(b-\bar{b}(1))^{2} \mid I_{1}\right\}
$$

Let us focus on the term $\sigma_{b}^{2}(1)$ in the expression (6.2) for $J_{1}\left(I_{1}\right)$. We can obtain $\sigma_{b}^{2}(1)$ from the equation $\mathrm{x} 1=X o+b u o+w o$ (which we view as a noise-corrupted measurement of b) and least-squares estimation theory (see Appendix E). The formula for $\sigma_{b}^{2}(1)$ will be of no further use to us, so we just state it without going into the calculation:

$$
\sigma_{b}^{2}(1)=\frac{\sigma_{b}^{2} \sigma_{w}^{2}}{u_{0}^{2} \sigma_{b}^{2}+\sigma_{w}^{2}}
$$

The salient feature of this equation is that $\sigma_{b}^{2}(1)$ is affected by the control $U o$. Basically, if luol is small, the measurement $\mathrm{x}_{1}=X o+b u o+w o$ is dominated by wo and the "signal-to-noise ratio" is small. Thus to achieve small error variance $\sigma_{b}^{2}(1)$ [which is desirable in view of Eq. (6.2)], we must apply a control $U o$ that is large in absolute value. A choice of large control to enhance parameter identification is often referred to as probing. On the other hand, if luo I is large, $|\mathrm{X}| \mid$ will also be large, and this is not desirable in view of Eq. (6.2). Therefore, in choosing $U o$ we must strike a balance between caution (choosing a small value to keep xı reasonably small) and probing (choosing a large value to improve the signal-to-noise ratio and enhance estimation of $b$ ).

The tradeoff between the control objective and the parameter estimation objective is commonly referred to as dual control.

### 6.1.2 Two-Phase Control and

$\qquad$
An apparently reasonable form of suboptimal control in the presence of unknown parameters (d. Example 6.1.3) is to separate the control process into two phases, a parameter identification phase and a control phase. In the first phase the unknown parameters are identified, while the control takes no account of the interim results of identification. The final parameter estimates from the first phase are then used to implement an optimal control law in the second phase. This alternation of identification and control phases may be repeated several times during any system run in order to take into account subsequent changes of the parameters.

One drawback of this approach is that information gathered during the identification phase is not used to adjust the control law until the beginning of the second phase. Furthermore, it is not always easy to determine when to terminate one phase and start the other.

A second difficulty, of a more fundamental nature, is due to the fact that the control process may make some of the unknown parameters invisible to the identification process. This is the problem of parameter identifiability, discussed by Ljung [Lju86], which is best explained by means of an example.

## Example 6.1.5

Consider the scalar system

$$
x_{k+1}=a x_{k}+b u_{k}+w_{k}, \quad k=0,1, \ldots, N-1
$$

with the quadratic cost

$$
E\left\{\sum_{k=1}^{N}\left(x_{k}\right)^{2}\right\}
$$

We assume perfect state information, so if the parameters $a$ and $b$ are known, this is a minimum variance control problem (d. Section 5.3), and the optimal
control law is

$$
\mu_{k}^{*}\left(x_{k}\right)=-\frac{a}{b} x_{k}
$$

Assume now that the parameters $a$ and $b$ are unknown, and consider the two-phase method. During the first phase the control law

$$
\begin{equation*}
\tilde{\mu}_{k}\left(x_{k}\right)=\gamma x_{k} \tag{6.3}
\end{equation*}
$$

is used ( $\gamma$ is some scalar; for example, $\gamma=-\bar{a} / \bar{b}$, where $\bar{a}$ and $\bar{b}$ are a priori estimates of $a$ and $b$, respectively). At the end of the first phase, the control law is changed to

$$
\bar{\mu}_{k}\left(x_{k}\right)=-\frac{\hat{a}}{\hat{b}} x_{k}
$$

where $\hat{a}$ and $\hat{b}$ are the estimates obtained from the identification process. However, with the control law (6.3), the closed-loop system is

$$
x_{k+1}=(a+b \gamma) x_{k}+W k
$$

so the identification process can at best identify the value of $(a+b \gamma)$ but not the values of both $a$ and $b$. In other words, the identification process cannot discriminate between pairs of values ( $a^{\prime} \mathrm{b}_{1}$ ) and ( $a 2, b_{2}$ ) such that al $+b_{1} \gamma=a 2+b_{2} \gamma$. Therefore, $a$ and $b$ are not identifiable when feedback control of the form (6.3) is applied.

One way to correct the difficulty is to add an additional known input $\delta_{k}$ to the control law (6.3); that is, use

$$
\tilde{\mu}_{k}\left(x_{k}\right)=\gamma x_{k}+\delta_{k}
$$

Then the closed-loop system becomes

$$
X k+l=(a+b \gamma) x_{k}+b \delta_{k}+W k
$$

and the knowledge of $\{X k\}$ and $\left\{\delta_{k}\right\}$ makes it possible to identify $\left(a+b_{l}\right)$ and $b$. Given $\gamma$, one can then obtain estimates of $a$ and $b$. Actually, to guarantee this in a more general context where the system is of higher dimension, the sequence $\left\{\delta_{k}\right\}$ must satisfy certain conditions: it must be "persistently exciting" (see for example Ljung and Soderstrom [LjS83] for further explanation of this concept).

A second possibility to bypass the identifiability problem is to change the structure of the system by artificially introducing a one-unit delay in the control feedback. Thus, instead of considering control laws of the form $\tilde{\mu}_{k}\left(x_{k}\right)=\gamma x_{k}$, as in Eq. (6.3), we consider controls of the form

$$
U k=\hat{\mu}_{k}\left(x_{k-1}\right)=\gamma x_{k-1}
$$

The closed-loop system then becomes

$$
x_{k+1}=a x_{k}+b \gamma x_{k-1}+w_{k}
$$

and given $\gamma$, it is possible to identify both parameters $a$ and $b$. This technique can be generalized for systems of arbitrary order, but artificially introducing a control delay makes the system less responsive to control.

### 6.1.3 Certainty Equivalent Control and ниенидианиниу

At the opposite extreme of the two-phase method we have the certainty equivalent control approach, where the parameter estimates are incorporated into the control law as they are generated, and they are treated as if they were true values. In terms of the system

$$
x_{k+1}=f_{k}\left(x_{k}, \theta, u_{k}, w_{k}\right)
$$

considered in Example 6.1.3, suppose that, for each possible value of $\theta$, the control law $\pi^{*}(\theta)=\left\{\mu_{0}^{*}(', \theta), \ldots, \mu_{N-1}^{*}\left({ }^{\prime}, \theta\right)\right\}$ is optimal with respect to a certain cost $J_{\pi}\left(x_{0}, \theta\right)$. Then the (suboptimal) control used at time $k$ is

$$
\hat{\mu}_{k}\left(I_{k}\right)=\mu_{k}^{*}\left(x_{k}, \hat{\theta}_{k}\right)
$$

where $\hat{\theta}_{k}$ is an estimate of $\theta$ based on the information

$$
I_{k}=\left\{x_{0}, x_{1}, \ldots, x_{k}, u_{0}, u_{1}, \ldots, u_{k-1}\right\}
$$

available at time $k$; for example,

$$
\hat{\theta}_{k}=E\left\{\theta \mid I_{k}\right\}
$$

or, more likely in practice, an estimate obtained via an on-line system identification method (see [KuV86], [LjS83], [Lju86]).

One would hope that when the horizon is very long, the parameter estimates $\hat{\theta}_{k}$ will converge to the true value $\theta$, so the certainty equivalent controller will become asymptotically optimal. Unfortunately, we will see that difficulties related to identifiability arise here as well.

Suppose for simplicity that the system is stationary with a priori known transition probabilities $P\left\{X k+\left.1\right|_{X k}, U k, \theta\right\}$ and that the control law used is also stationary:

$$
\hat{\mu}_{k}\left(I_{k}\right)=\mu^{*}\left(x_{k}, \hat{\theta}_{k}\right), \quad k=0,1, \ldots
$$

There are three systems of interest here (cf. Fig. 6.1.2):
(a) The system (perhaps falsely) believed by the controller to be true, which evolves probabilistically according to

$$
P\left\{x_{k+1} \mid x_{k}, \mu^{*}\left(x_{k}, \hat{\theta}_{k}\right), \hat{\theta}_{k}\right\}
$$

(b) The true closed-loop system, which evolves probabilistically according to

$$
P\left\{x_{k+1} \mid x_{k}, \mu^{*}\left(x_{k}, \hat{\theta}_{k}\right), \theta\right\}
$$



Figure 6.1.2 The three systems involved in certainty equivalent control, where $\theta$ is the true parameter and $\hat{\theta}_{k}$ is the parameter estimate at time $k$. Loss of optimality occurs when the true system differs asymptotically from the optimal closed-loop system. If the parameter estimates converge to some value $\theta$, the true system typically becomes asymptotically equal to the system believed to be true. However, the parameter estimates need not converge, and even if they do, both systems may be different asymptotically from the optimal.
(c) The optimal closed-loop system that corresponds to the true value of the parameter, which evolves probabilistically according to

$$
P\left\{x_{k+1} \mid x_{k}, \mu^{*}\left(x_{k}, \theta\right), \theta\right\}
$$

For asymptotic optimality, we would like the last two systems to be equal asymptotically. This will certainly be true if $\hat{\theta}_{k} \rightarrow()$. However, it is quite possible that either
(1) $\hat{\theta}_{k}$ does not converge to anything, or that
(2) $\hat{\theta}_{k}$ converges to a parameter $\hat{\theta} \neq 0$.

There is not much we can say about the first case, so we concentrate on the second. To see how the parameter estimates can converge to a wrong value, assume that for some $\hat{\theta} \neq \theta$ and all $X k+l, X k$, we have

$$
\begin{equation*}
P\left\{x_{k+1} \mid x_{k}, \mu^{*}\left(x_{k}, \hat{\theta}\right), \hat{\theta}\right\}=P\left\{x_{k+1} \mid x_{k}, \mu^{*}\left(x_{k}, \hat{\theta}\right), \theta\right\} . \tag{6.4}
\end{equation*}
$$

In words, there is a false value of parameter for which the system under closed-loop control looks exactly as if the false value were true. Then, if the controller estimates at some time the parameter to be $\hat{\theta}$, subsequent data will tend to reinforce this erroneous estimate. As a result, a situation may develop where the identification procedure locks onto a wrong parameter value, regardless of how long information is collected. This is a difficulty with identifiability of the type discussed earlier in connection with twophase controL

On the other hand, if the parameter estimates converge to some (possibly wrong) value, we can argue intuitively that the first two systems (believed and true) typically become equal in the limit as $k \rightarrow$ oo, since, generally, parameter estimate convergence in identification methods implies that the data obtained are asymptotically consistent with the view of the system one has based on the current estimates. However, the believed and true systems mayor may not become asymptotically equal to the optimal closed-loop system. We first present two examples that illustrate how, even when the parameter estimates converge, the true closed-loop system can differ asymptotically from the optimal, thereby resulting in a certainty equivalent controller that is strictly suboptimal. We then discuss the special case of the self-tuning regulator for ARMAX models with unknown parameters, where, remarkably, it turns out that all three of the above systems are typically equal in the limit, even though the parameter estimates typically converge to false values.

## Example 6.1.6 [BoV79]

Consider a two-state system with two controls $u^{1}$ and $u^{2}$. The transition probabilities depend on the control applied as well as a parameter $B$, which is known to take one of two values $B^{*}$ and $\hat{\theta}$. They are as shown in Fig. 6.1.3. There is zero cost for a transition from state 1 to itself and a unit cost for all other transitions. Therefore, the optimal control at state 1 is the one that maximizes the probability of the state remaining at 1 . Assume that the true parameter is $B^{*}$ and that

$$
p_{11}\left(u^{1}, \hat{\theta}\right)>p_{11}\left(u^{2}, \hat{\theta}\right), \quad p_{11}\left(u^{1}, \theta^{*}\right)<p_{11}\left(u^{2}, \theta^{*}\right)
$$

Then the optimal control is $u^{2}$, but if the controller thinks that the true parameter is $\hat{\theta}$, it will apply $\mathbf{u}^{1}$. Suppose also that

$$
p_{11}\left(u^{1}, \hat{\theta}\right)=p_{11}\left(u^{1}, \theta^{*}\right)
$$

Then, under $u^{1}$ the system looks identical for both values of the parameter, so if the controller estimates the parameter to be $\hat{\theta}$ and applies $u^{1}$, subsequent data will tend to reinforce the controller's belief that the true parameter is indeed $\hat{\theta}$.

More precisely, suppose that we estimate $\theta$ by selecting at each time $k$ the value that maximizes

$$
P\left\{8 I_{\left.I_{k}\right\}}=\frac{P\left\{I_{k}\right.}{P} \frac{18\} P(B)}{\left(I_{k}\right)},\right.
$$

where $P(8)$ is the a priori probability that the true parameter is $\theta$ (this is a popular estimation method). Then if $P(\hat{\theta})>P\left(8^{*}\right)$, it can be seen, by using induction, that at each time $k$, the controller will estimate falsely $\theta$ to be $\hat{\theta}$ and apply the incorrect control $u^{1}$. To avoid the difficulty illustrated in this example, it has been suggested to occasionally deviate from the certainty


Transition probabilities for $\boldsymbol{\theta}=8^{\star}$
(true parameter value)
$p_{11}\left(u^{1}\right)=0.5$


Transition probabilities for $\theta=\hat{\theta}$ (false parameter value)

Figure 6.1.3 Transition probabilities for the two-state system of Example 6.1.6. Under the nonoptimal control $\mathrm{u}^{1}$, the system looks identical under the true and the false values of the parameter $\theta$.
equivalent control, applying other controls that enhance the identification of the unknown parameter (see Doshi and Shreve [DoS80], and Kumar and Lin [KuL82]). For example, by making sure that the control $u^{2}$ is used infrequently but infinitely often, we can guarantee that the correct parameter value will be identified by the preceding estimation scheme.

## Example 6.1.7 [Kum83]

Consider the linear scalar system

$$
x_{k+1}=a x_{k}+b u_{k}+w_{k}
$$

where we know that the parameters are either $(a, b)=(1,1)$ or $(a, b)=$ $(0,-1)$. The sequence $\{W k\}$ is independent, stationary, zero mean, and Gaussian. The cost is quadratic of the form

$$
\sum_{k=O}^{\mathbf{N - I}}\left(\left(x_{k}\right)^{2}+2(U k) 2\right),
$$

where $N$ is very large, so the stationary form of the optimal control law is used (see Section
equation to be

$$
\mu^{*}\left(x_{k}\right)= \begin{cases}-\frac{x_{k}}{2} & \text { if }(a, b)=(1,1) \\ 0 & \text { if }(a, b)=(0,-1)\end{cases}
$$

To estimate ( $a, b$ ), we use a least-squares identification method. The value of the least-squares criterion at time $k$ is given by

$$
\begin{align*}
& V k(l, 1)=\sum_{i=0}^{\mathrm{k}-1}\left(x_{i+1}-X i-U i\right) 2, \quad \text { for }(a, b)=(1,1),  \tag{6.5}\\
& V k(O,-1)=\sum_{i=0}^{\mathrm{k}-1}\left(x_{i+1}+U i\right) 2, \quad \text { for }(a, b)=(0,-1) \tag{6.6}
\end{align*}
$$

The control applied at time $k$ is

$$
u_{k}=\tilde{\mu}_{k}\left(I_{k}\right)= \begin{cases}-\frac{x_{k}}{2} & \text { if } \operatorname{Vk}(l, 1)<\operatorname{Vk}(O,-1), \\ 0 & \text { if } \operatorname{Vk}(1,1)>\operatorname{Vk}(0,-1)\end{cases}
$$

Suppose the true parameters are $\theta=(0,-1)$. Then the true system evolves according to

$$
\begin{equation*}
x_{k+1}=-u_{k}+w_{k} \tag{6.7}
\end{equation*}
$$

If at time $k$ the controller estimates incorrectly the parameters to be $\hat{\theta}=(1,1)$, because $V_{k}(\hat{\theta})<V k(B)$, the control applied will be $U k=-X k / 2$ and the true closed-loop system will evolve according to

$$
\begin{equation*}
x_{k+1}=\frac{x_{k}}{2}+w_{k} \tag{6.8}
\end{equation*}
$$

On the other hand, the controller thinks (given the estimate $\hat{\theta}$ ) that the closedloop system will evolve according to

$$
\begin{equation*}
x_{k+1}=x_{k}+u_{k}+w_{k}=x_{k}-\frac{x_{k}}{2}+w_{k}=\frac{x_{k}}{2}+w_{k} \tag{6.9}
\end{equation*}
$$

so from Eqs. (6.7) and (6.8) we see that under the control law $U k=-X k / 2$, the closed-loop system evolves identically for both the true and the false values of the parameters [ef. Eq. (6.4)].

To see what can go wrong, note that if $V k(\hat{\theta})<V k(B)$ for some $k$ we will have, from Eqs. (6.5)-(6.9),

$$
x_{k+1}+u_{k}=x_{k+1}-x_{k}-u_{k}
$$

so from Eqs. (6.5) and (6.6) we obtain

$$
V_{k+1}(\hat{\theta})<V_{k+1}(\theta)
$$

Therefore, if $V_{l}(\hat{\theta})<V_{l}(8)$, the least-squares identification method will yield the wrong estimate $\hat{\theta}$ for eveTy $k$. To see that this can happen with positive probability, note that, since the true system is $x_{k+1}=-\boldsymbol{U k}+w k$, we have

$$
\begin{aligned}
& V_{l}(\hat{8})=(\mathrm{Xl}-X o-u o)^{2}=(w o-X o-2 u o)^{2} \\
& V_{l}(8)=(\mathrm{Xl}+U O) 2=w_{0}^{2}
\end{aligned}
$$

Therefore, the inequality $V_{l}(\hat{\theta})<V_{l}(8)$ is equivalent to

$$
(x o+2 U O) 2<2 w o(x o+2 u o)
$$

which will occur with positive probability since wo is Gaussian.
The preceding examples illustrate that loss of identifiability is a serious problem that frequently arises in the context of certainty equivalent control.

### 6.1.4 Self-Tuning Regulators

We described earlier the nature of the identifiability issue in certainty equivalent control: under closed-loop control, incorrect parameter estimates can make the system behave as if these estimates were correct [cf. Eq. (6.4)]. As a result, the identification scheme may lock onto false parameter values. This is not necessarily bad, however, since it may happen that the control law implemented on the basis of the false parameter values is near optimal. Indeed, through a fortuitous coincidence, it turns out that in the practically important minimum variance control formulation (Section 5.3), when the parameter estimates converge, they typically converge to false values, but the resulting control law typically converges to the optimal. We can get an idea about this phenomenon by means of an example.

## Example 6.1.8

Consider the simplest ARMAX model:

$$
y_{k+1}+a y_{k}=b u_{k}+\epsilon_{k+1}
$$

The minimum variance control law when $a$ and $b$ are known is

$$
u_{k}=\mu_{k}\left(I_{k}\right)=\frac{a}{b} y_{k}
$$

Suppose now that $a$ and $b$ are not known but are identified on-line by means of some scheme. The control applied is

$$
\begin{equation*}
u_{k}=\frac{\hat{a}_{k}}{\hat{b}_{k}} y_{k} \tag{6.10}
\end{equation*}
$$

where $\hat{a}_{k}$ and $\hat{b}_{k}$ are the estimates obtained at time $k$. Then the difficulty with identifiability occurs when
where $\hat{a}$ and $\hat{b}$ are such that the true dosed-loop system given by

$$
y_{k+1}+a y_{k}=\frac{b \hat{a}}{\hat{b}} y_{k}+\epsilon_{k+1}
$$

coincides with the closed-loop system that the controller thinks is true on the basis of the estimates $\hat{a}$ and $\hat{b}$. This latter system is

$$
y_{k+1}=\epsilon_{k+1}
$$

For these two systems to be identical, we must have
which means that the control law (6.10) asymptotically becomes optimal despite the fact that the asymptotic estimates $\hat{a}$ and $\hat{b}$ may be incorrect.

Example 6.1 .8 can be extended to the general ARMAX model of Section 5.3 with no delay:

If the parameter estimates converge (regardless of the identification method used and regardless of whether the limit values are correct), then a minimum variance controller thinks that the closed-loop system is asymptotically

$$
Y k=\epsilon_{k}
$$

Furthermore, parameter estimate convergence intuitively means that the true closed-loop system is also asymptotically $Y k=E k$, and this is clearly the optimal closed-loop system. Results of this type have been proved in the literature in connection with several popular methods for parameter estimation. In fact, surprisingly, in some of these results, the model adopted by the controller is allowed to be incorrect to some extent.

One issue that we have not discussed is whether the parameter estimates indeed converge. A complete analysis of this issue is quite difficult. We refer to the survey paper by Kumar [Kum85], and the textbooks by Goodwin and Sin [GoS84], Kumar and Varaiya [KuV86], and Astrom and vVittenmark [AsW90] for a discussion and sources on this subject. However, extensive simulations have shown that with proper implementation, these estimates typically converge for the type of systems likely to arise in many applications.

### 6.2 OPEN-LOOP FEEDBACK CONTROL

Generally, in a problem with imperfect state information, the performance of the optimal policy improves when extra information is available. However, the use of this information may make the DP calculation of the optimal policy intractable. This motivates an approximation, based on a more tractable computation that in part ignores the availability of extra information.

Let us consider the imperfect state information problem under the assumption of Section 5.4.1, which guarantees that the conditional state distribution is a sufficient statistic, i.e., that the probability distribution of the observation disturbance $V k+l$ depends explicitly only on the immediately preceding state, control, and system disturbance $X k, U k, W k$, and not on $X k-l, \ldots, X o, U k-I, \ldots, U O, W k-l \ldots, W o, V k-l, \cdots, V a$.

We introduce a suboptimal policy known as the open-loop feedback controller (OLFC), which uses the current information vector $I_{k}$ to determine $P_{x} k I l k^{\prime}$ However, it calculates the control $U k$ as if no further measurements will be received, by using an open-loop optimization over the future evolution of the system. In particular, $U k$ is determined as follows:
(1) Given the information vector $I_{k}$, compute the conditional probability distribution PXkllk (in the case of perfect state information, where $I_{k}$ includes $X k$, this step is unnecessary).
(2) Find a control sequence $\left\{\bar{u}_{k}, U k+l, " \prime \bar{u}_{N-1}\right\}$ that solves the openloop problem of minimizing

$$
E\left\{g_{N}\left(x_{N}\right)+\sum_{i=k}^{N-1} g_{i}\left(x_{i}, u_{i}, w_{i}\right) \mid I_{k}\right\}
$$

subject to the constraints

$$
x_{i+1}=f_{i}\left(x_{i}, u_{i}, w_{i}\right), \quad U i \mathrm{E} U i, \quad \mathrm{i}=k, k+1, \ldots, N-1 .
$$

(3) Apply the control input

$$
\bar{\mu}_{k}\left(I_{k}\right)=\bar{u}_{k} .
$$

Thus the OLFC uses at time $k$ the new measurement $z k$ to calculate the conditional probability distribution PXkllk' However, it selects the control input as if future measurements will be disregarded.

Similar to the CEC, the OLFC requires the solution of $N$ optimal control problems. Each problem may again be solved by deterministic optimal control techniques. The computations, however, may be more complicated than those for the CEe, since now the cost involves an expectation with
respect to the uncertain quantities. The main difficulty in the implementation of the OLFC is the computation of PXkllk' In many cases one cannot compute $P_{x k l} l_{l k}$ exactly, in which case some "reasonable" approximation scheme must be used. Of course, if we have perfect state information, this difficulty does not arise.

In any suboptimal control scheme, one would like to be assured that measurements are advantageously used. By this we mean that the scheme performs at least as well as any open-loop policy that uses a sequence of controls that is independent of the values of the measurements received. An optimal open-loop policy can be obtained by finding a sequence $\left\{u_{0}^{*}, u_{1}^{*}, \ldots, u_{N}^{*}-I\right\}$ that minimizes

$$
\bar{J}\left(u_{0}, u_{1}, \ldots, u_{N-1}\right)=E\left\{g_{N}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(x_{k}, u_{k}, w_{k}\right)\right\}
$$

subject to the constraints

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}, w_{k}\right), \quad u_{k} \in U_{k}, \quad k=0,1, \ldots, N-1
$$

A nice property of the OLFC is that it performs at least as well as an optimal open-loop policy, as shown by the following proposition. By contrast, the CEC does not have this property (for a one-stage problem, the optimal open-loop controller and the OLFC are both optimal, but the CEC may be strictly suboptimal; see also Exercise 6.2).

Propositi 6.2.1: The cost $J_{\bar{\pi}}$ corresponding to an OLFC satisfies

$$
\leq J_{0}^{*}
$$

(6.11)
where $J_{0}^{*}$ is the cost corresponding to an optimal open-loop policy.

Proof: We assume throughout the proof that all expected values appearing are well defined and finite, and that the minimum in the following Eq. (6.14) is attained for every $I_{k}$. Let $\bar{\pi}=\left\{\bar{\mu}_{0}, \bar{\mu}_{1}, \ldots \bar{\mu}_{N-1}\right\}$ be the OLFC. Its cost is given by

$$
\begin{equation*}
J_{\bar{\pi}}=\underset{z Q}{E}\{J o(I o)\}=\underset{z Q}{E}\{J o(z o)\} \tag{6.12}
\end{equation*}
$$

where $\bar{J}_{0}$ is obtained from the recursive algorithm

$$
\begin{aligned}
I N-l(I N-l)= & \underset{x N-1, W N-1}{E}\left\{g N\left(f N-l \quad, \bar{\mu}_{N-1}\left(I_{N-1}\right), w_{N-1}\right)\right) \\
& \left.+g_{N-1}\left(x_{N-1}, \bar{\mu}_{N-1}\left(I_{N-1}\right), w_{N-1}\right) \mid I N-I\right\}
\end{aligned}
$$

$$
\begin{aligned}
& \underset{x_{k}, w_{k}, v_{k+1}}{E}\left\{g_{k}\left(x_{k}, \bar{\mu}_{k}\left(I_{k}\right), w_{k}\right)\right. \\
& \left.+\bar{J}_{k+1}\left(I_{k}, h_{k+1}\left(f_{k}\left(x_{k}, \bar{\mu}_{k}\left(I_{k}\right), w_{k}\right), \bar{\mu}_{k}\left(I_{k}\right), v_{k+1}\right), \bar{\mu}_{k}\left(I_{k}\right)\right) \mid I_{k}\right\} \\
& k=0, \ldots, N-1
\end{aligned}
$$

(6.13)
where $h k$ is the function involved in the measurement equation as in the basic problem with imperfect state information of Section 5.1.

Consider the functions $J_{k}^{c}\left(I_{k}\right), k=0,1, \ldots, N-1$, defined by

$$
\begin{equation*}
J_{k}^{c}\left(I_{k}\right)=\min _{\substack{u_{i} \in U_{i} \\ i=k, \ldots, N-1}} E\left\{9 N(X N)+\sum_{i=k}^{N-1} 9 i\left(X i, U i, \text { Wi) } \mid I_{k}\right\}\right. \tag{6.14}
\end{equation*}
$$

The minimization problem in this equation is the one that must be solved at time $k$ in order to calculate the control $\bar{\mu}_{k}\left(I_{k}\right)$ of the OLFC. Clearly, $J_{k}^{c}\left(I_{k}\right)$ can be interpreted as the optimal open-loop cost from time $k$ to time $N$ when the current information vector is $I_{k}$. It can be seen that

$$
\begin{equation*}
\underset{Z 0}{E}\{J 8(z o)\} \leq J_{0}^{*} \tag{6.15}
\end{equation*}
$$

since $J_{0}^{*}$ is the minimum over ${ }^{\prime} U 0, \ldots, U N_{-1}$ of the total expected cost and can be written as

$$
\min _{U O O^{\prime \prime \prime}, U N-l} \underset{Z_{0}}{E}\{E\{\operatorname{cost} \mid z o\}\}
$$

while $E z_{o}\{J 8(z o)\}$ can be written as

$$
\underset{z_{0}}{E}\left\{\min _{u_{0}, \ldots, u_{N-1}} E\{\cos t \mid z, 0\}\right\}
$$

(we generally have $\mathrm{E}\{\min [].) \leq \min [E\{\cdot\}]$ ). We will prove that

$$
\begin{equation*}
\bar{J}_{k}\left(I_{k}\right) \leq J_{k}^{c}\left(I_{k}\right), \quad \text { for all } I_{k} \text { and } k \tag{6.16}
\end{equation*}
$$

Then from Eqs. (6.12), (6.15), and (6.16), it will follow that

$$
J_{\bar{\pi}} \leq J_{0}^{*}
$$

which is the relation to be proved. We show Eq. (6.16) by induction.
By the definition of the OLFC and Eq. (6.14), we have

$$
\bar{J}_{N-1}\left(I_{N-1}\right)=J_{N-1}^{c}\left(I_{N-1}\right), \quad \text { for all } I N_{-1}
$$

and hence Eq. (6.16) holds for $k=N-1$. Assume that

$$
\begin{equation*}
\bar{J}_{k+1}\left(I_{k+1}\right) \leq J_{k+1}^{c}\left(I_{k+1}\right), \quad \text { for all } I_{k+1} \tag{6.17}
\end{equation*}
$$

Then from Eqs. (6.13), (6.14), and (6.17), we have

$$
\begin{aligned}
& \bar{J}_{k}\left(I_{k}\right)={\underset{x_{k}, w_{k}, v_{k+1}}{E}\{9 k(X k, T i k(I k), W k), ~(1)} \\
& \left.+\bar{J}_{k+1}\left(I_{k}, h_{k+1}\left(f_{k}\left(x_{k}, \bar{\mu}_{k}\left(I_{k}\right), w_{k}\right), \bar{\mu}_{k}\left(I_{k}\right), v_{k+1}\right), \bar{\mu}_{k}\left(I_{k}\right)\right) \mid I_{k}\right\} \\
& \leq \underset{X k, W k, V k+l}{E}\left\{9 k\left(X k, \bar{\mu}_{k}\left(I_{k}\right), W k\right)\right. \\
& \left.+J_{k+1}^{c}\left(I k, h k+l\left(f_{k}\left(X k^{\prime} \bar{\mu}_{k}\left(I_{k}\right), W k\right), \bar{\mu}_{k}\left(I_{k}\right), V k+l\right), \bar{\mu}_{k}\left(I_{k}\right)\right) \mid I_{k}\right\} \\
& \underset{x_{k}, w_{k}, v_{k+1}}{E}\left\{\min _{\substack{\boldsymbol{U} i \boldsymbol{E} \boldsymbol{U i} \\
i=k+1, \ldots, N-1}}^{\underset{\substack{X k+l, W i \\
x i+l=f i(x i, u i, w i)}}{E}\left\{g_{k}\left(x_{k}, \bar{\mu}_{k}\left(I_{k}\right), w_{k}\right)\right.}\right. \\
& \left.\left.+\sum_{i=k+l}^{N-l} g i(X i, U i, W i)+g N(X N) \mid I_{k+1}\right\} \mid I_{k}\right\} \\
& \min _{\substack{u i E U i \\
\mathrm{i}=\mathrm{k}+1, \ldots, \mathrm{~N}-1}}^{\substack{X k, W k, W i \\
\mathrm{Xi}+1=f i(X i, U i, W i) \\
\mathrm{i}=\mathrm{k}+1, \ldots, N-1}} \mid \boldsymbol{E}, \ldots\left(g_{N}\left(x_{N}\right)+g_{k}\left(x_{k}, \bar{\mu}_{k}\left(I_{k}\right), w_{k}\right)\right. \\
& X k+l=f_{k}\left(x_{k}, \bar{\mu}_{k}\left(I_{k}\right), w_{k}\right) \\
& \left.+\sum_{i=k+l}^{N-l} g i(X i, U i, W i) \mid I_{k}\right\} \\
& =J_{k}^{c}\left(I_{k}\right) \text {. }
\end{aligned}
$$

The second inequality follows by interchanging expectation and minimization (since we generally have $\mathrm{E}\{\min [\cdot]) \leq \min [E\{\cdot\}]$ ) and by "integrating out" $V k+l$. The last equality follows from the definition of OLFC. Thus Eq. (6.16) is proved for all $k$ and the desired result is shown. Q.E.D.

The preceding proposition shows that the OLFC uses the measurements advantageously even though it selects at each period the present control input as if no further measurements will be taken in the future. It is worth noting that by Eq. (6.16), $J_{k}^{c}\left(I_{k}\right)$, which is the calculated openloop optimal cost from time $k$ to time $N$, provides a readily obtainable performance bound for the OLFC.

## Partial Open-Loop Feedback Control

A form of suboptimal control that is intermediate between the optimal feedback controller and the OLFC is provided by a generalization of the OLFC called the partial open-loop feedback controller (POLFC for short). This controller uses past measurements to compute $P_{x l_{l k}}$ ' but calculates the control input on the basis that some (but not necessarily all) of the
measurements will in fact be taken in the future, and the remaining measurements will not be taken.

This method often allows one to deal with those measurements that are troublesome and complicate the solution. As an example consider an inventory problem such as the one considered in Section 4.2, where forecasts in the form of a probability distribution of each of the future demands become available over time. A reasonable form of POLFC calculates at each stage an optimal ( $s, S$ ) policy based on the current forecast of future demands and follows this policy until a new forecast becomes available. When this happens, the current policy is abandoned in favor of a new one that is calculated on the basis of the new probability distribution of future demands, etc. Thus the complications due to the forecasts are bypassed at the expense of suboptimality of the policy obtained.

We note that an analog of Prop. 6.2.1 can be shown for the POLFC
Bertsekas [Ber76)). In fact the corresponding error bound is potentially much better than the bound (6.11), reflecting the fact that the POLFC takes into account the future availability of some of the measurements.

We will discuss further the idea of ignoring a portion of the information for the purpose of obtaining a tractable suboptimal policy in Section 6.5.3. There we will generalize the OLFC and the POLFC by embedding them within a more general suboptimal scheme.

### 6.3 LIMITED LOOKAHEAD POLICIES

An effective way to reduce the computation required by DP is to truncate the time horizon and use at each stage a decision based on lookahead of a small number of stages. The simplest possibility is to use a one-step lookahead pol'icy whereby at stage $k$ and state $X k$ one uses the control $\bar{\mu}_{k}\left(x_{k}\right)$, which attains the minimum in the expression

$$
\begin{equation*}
\min _{u k E U k(X k)} E\{9 k(X k, U k, W k)+J k+l(f k(X k, U k, W k))\}, \tag{6.18}
\end{equation*}
$$

where $\mathrm{J}_{\mathrm{k}+1}$ is some approximation of the true cost-to-go function $\mathrm{Jk}+1$, with $\mathbf{J}_{N} \quad g N$. Similarly, a two-step lookahead policy applies at time $k$ and state $X k$, the control $\bar{\mu}_{k}\left(x_{k}\right)$ attaining the minimum in the preceding equation, where now $J_{k}+l$ is obtained itself on the basis of a one-step lookahead approximation. In other words, for all possible states $X k+l$ that can be generated via the system equation starting from $X k$,

$$
X k+l=f_{k}\left(x_{k}, U k, W k\right)
$$

we have

$$
\begin{aligned}
\left(x_{k+1}\right)= & \min _{\prime U k+l E U k+l(X k+l)} E\left\{9 k+l \quad u_{k+1}, W k+l\right) \\
& +J k+2(f k+l(X k+l, U k+l, W k+l))\}
\end{aligned}
$$

where $J_{k}+2$ is some approximation of the cost-to-go function $J_{k}+{ }_{2}$. Policies with lookahead of more than two stages are similarly defined.

Note that the limited lookahead approach can be used equally well when the horizon is infinite. One simply uses as the terminal cost-to-go function an approximation to the optimal cost of the infinite horizon problem that starts at the end of the lookahead. Thus the following discussion, with a few straightforward modifications, applies to infinite horizon problems as well.

Given the approximations $J_{k}$ to the optimal costs-to-go, the computational savings of the limited lookahead approach are evident. For a one-step lookahead policy, only a single minimization problem has to be solved per stage, while in a two-step policy the corresponding number of minimization problems is one plus the number of all possible next states $X k+l$ that can be generated from the current state $X k$.

However, even with readily available cost-to-go approximations $J_{k}$, the minimization over $U k E U k(X k)$ in the calculation of the one-step lookahead control [cf. Eq. (6.18)] may involve substantial computation. In a variant of the method that aims at reducing this computation, the minimization is done over a subset

$$
\bar{U}_{k}\left(x_{k}\right) \subset U_{k}\left(x_{k}\right) .
$$

Thus, the control $\bar{\mu}_{k}\left(x_{k}\right)$ used in this variant is one that attains the minimum in the expression

$$
\begin{equation*}
\min _{u_{k} \in \bar{U}_{k}\left(x_{k}\right)} E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\} \tag{6.19}
\end{equation*}
$$

A practical example of this approach is when by using some heuristic or approximate optimization, we identify a subset $\bar{U} k(X k)$ of promising controls, and to save computation, we restrict attention to this subset in the one-step lookahead minimization.
6.3.1 Performance Bounds for Limited Lookahead Policies

Let us denote by $J k(X k)$ the expected cost-to-go incurred by a limited lookahead policy $\left\{\bar{\mu}_{0}, \bar{\mu}_{1}, \ldots, \bar{\mu}_{N-1}\right\}$ starting from state $X k$ at time $k[J k(X k)$ should be distinguished from $J k(X k)$, the approximation of the cost-to-go that is used to compute the limited lookahead policy via the minimization in Eq. (6.19)). It is generally difficult to evaluate analytically the functions $\vec{J}_{k}$, even when the functions $J_{k}$ are readily available. We thus aim to obtain some estimates of $]_{k}(x k)$. The following proposition gives a condition under which the one-step lookahead policy achieves a cost $\bar{J}_{k}\left(X_{k}\right)$ that is better than the approximation $\operatorname{Jk}(X k)$. The proposition also provides a readily computable upper bound to $J k(X k)$.
$\underline{\mathbb{P}_{\mathbf{1}}}=$ 6.3.1: Assume that for all $X k$ and $k$, we have

$$
\begin{equation*}
\min _{u_{k} \in \bar{U}_{k}\left(x_{k}\right)} E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\} \leq \tilde{J}_{k}\left(x_{k}\right) \tag{6.20}
\end{equation*}
$$

Then the cost-to-go functions $\bar{J}_{k}$ corresponding to a one-step lookahead policy that uses $J_{k}$ and $\bar{U}_{k}\left(x_{k}\right)$ [cf. Eq. (6.19)] satisfy for all $X k$ and $k$

$$
\begin{equation*}
\bar{J}_{k}\left(x_{k}\right) \leq \min _{u_{k} \in \bar{U}_{k}\left(x_{k}\right)} E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\} \tag{6.21}
\end{equation*}
$$

Proof: For $k=0, \ldots, N-1$, denote

$$
\begin{equation*}
\hat{J}_{k}\left(x_{k}\right)=\min _{u_{k} \in \bar{U}_{k}\left(x_{k}\right)} E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\widetilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\} \tag{6.22}
\end{equation*}
$$

and let $I N=g N$. We must show that for all $X k$ and $k$, we have

$$
\bar{J}_{k}\left(x_{k}\right) \leq \hat{J}_{k}\left(x_{k}\right)
$$

We use backwards induction on $k$. In particular, we have $J N(X N)$ $\hat{J}_{N}\left(x_{N}\right)=\tilde{I} N(X N)=g N(X N)$ for all XN. Assuming that $J_{k^{+}}{ }_{1}(X k+l) \leq$ $J k+l\left(x_{k+1}\right)$ for all $X k+l$, we have

$$
\begin{aligned}
J k(X k) & E\left\{g k\left(X k^{\prime} \bar{\mu}_{k}\left(x_{k}\right), W k\right)+J k+_{l}\left(f_{k}\left(X k^{\prime} \bar{\mu}_{k}\left(x_{k}\right), W k\right)\right)\right\} \\
\leq & E\left\{g\left(X k^{\prime} \bar{\mu}_{k}(X k), W k\right)+J k+_{l}\left(f k\left(X k^{\prime} \bar{\mu}_{k}\left(x_{k}\right), W k\right)\right)\right\} \\
\leq & E\left\{g\left(X k^{\prime} \bar{\mu}_{k}\left(x_{k}\right), W k\right)+J_{k+1}\left(f_{k}\left(X k^{\prime} \bar{\mu}_{k}\left(x_{k}\right), W k\right)\right)\right\} \\
& J k(X k),
\end{aligned}
$$

for all $X k$. The first equality above follows from the DP algorithm that defines the costs-to-go $\bar{J}_{k}$ of the limited lookahead policy, while the first inequality follows from the induction hypothesis, and the second inequality follows from the assumption (6.20). This completes the induction proof. Q.E.D.

Note that by Eq. (6.21), the value $J k(X k)$ of Eq. (6.22), which is the calculated one-step lookahead cost from state $X k$ at time $k$, provides a readily obtainable performance bound for the cost-to-go $\bar{J} k(x k)$ of the one-step lookahead policy. Furthermore, using also the assumption (6.20), we obtain for all $X k$ and $k$,

$$
\bar{J}_{k}\left(x_{k}\right) \leq \tilde{J}_{k}\left(x_{k}\right)
$$

i.e., the cost-to-go of the one-step lookahead policy is no greater than the lookahead approximation on which it is based. The critical assumption (6.20) in Prop. 6.3.1 can be verified in a few interesting special cases, as indicated by the following examples.

## Example 6.3.1 (Rollout Algorithm)

Suppose that $\mathrm{J} k(X k)$ is the cost-to-go of some given (suboptimal) heuristic policy $\pi=\left\{\mu_{0}, \ldots, \mu_{N-1}\right\}$ and that the set $\bar{U}_{k}\left(x_{k}\right)$ contains the control $\mu_{k}\left(x_{k}\right)$ for all $X k$ and $k$. The resulting one-step lookahead algorithm is called the rollout algorithm and will be discussed extensively in Section 6.4. From the DP algorithm (restricted to the given policy $\pi$ ), we have

$$
\tilde{J}_{k}\left(x_{k}\right)=E\left\{g_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right)\right\},
$$

which in view of the assumption $\mu_{k}\left(x_{k}\right) \mathrm{E} \bar{U}_{k}\left(x_{k}\right)$, yields

$$
J k(X k) \geq \min _{u k E U k(x k)} E\{g k(X k, U k, W k)+J k+l(J k(X k, U k, W k))\}
$$

Thus, the assumption of Prop. 6.3.1 is satisfied, and it follows that the rollout algorithm performs better than the heuristic on which it is based, starting from any state and stage.

Example 6.3.2 (Rollout Algorithm with Multiple Heuristics)
Consider a scheme that is similar to the one of the preceding example, except that $\mathbf{J}_{k}(x k)$ is the minimum of the cost-to-go functions corresponding to m heuristics, Le.,

$$
\tilde{J}_{k}\left(x_{k}\right)=\min \left\{J_{\pi_{1}, k}\left(x_{k}\right), \ldots, J_{\pi_{m}, k}\left(x_{k}\right)\right\}
$$

where for each $\mathbf{j}, J_{\pi_{i}}, k(X k)$ is the cost-to-go of a policy $\pi_{j}=\left\{\mu_{j, 0}, \ldots, \mu_{j, N-1}\right\}$, starting from state $J k$ at stage $k$. From the DP algorithm, we have, for all $\mathbf{j}$,

$$
J_{\pi_{j}, k}\left(x_{k}\right)=E\left\{g_{k}\left(x_{k}, \mu_{j, k}\left(x_{k}\right), w_{k}\right)+J_{\pi_{j}, k+1}\left(f_{k}\left(x_{k}, \mu_{j, k}\left(x_{k}\right), w_{k}\right)\right)\right\}
$$

from which, using the definition of $\boldsymbol{J}_{k}$, it follows that

$$
\begin{aligned}
J_{\pi_{j}, k}\left(x_{k}\right) & \geq E\left\{g_{k}\left(x_{k}, \mu_{j, k}\left(x_{k}\right), w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, \mu_{j, k}\left(x_{k}\right), w_{k}\right)\right)\right\} \\
& \geq \min _{u k E U k(x k)} E\{g k(X k, U k, W k)+J k+l(J k(X k, U k, W k))\}^{\prime}
\end{aligned}
$$

Taking the minimum of the left-hand side over $\mathbf{j}$, we obtain

$$
J k(X k) \geq \min _{u k E U k(x k)} E\left\{g k(X k, U k, W k)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}
$$

Thus, Prop. 6.3.1 implies that the one-step lookahead algorithm based on the heuristic algorithms' costs-to-go $J_{\pi_{1}, k}\left(x_{k}\right), \ldots, J_{\pi_{m}, k}\left(x_{k}\right)$ performs better than all these heuristics, starting from any state and stage.

Generally, the cost-to-go approximation functions $J_{k}$ need not satisfy the assumption (6.20) of Prop. 6.3.1. The following proposition does not require this assumption. It is useful in some contexts, including the case where the minimization involved in the calculation in the one-steplookahead policy is not exact.
$\underset{\tilde{x}}{\operatorname{Pr}} \quad$ 6.3.2: Let $\tilde{J} k, k=0,1, \ldots, N$, be functions of $X k$ with $\tilde{I N}(X N)=g N(X N)$ for all $X N$, and let $\pi=\left\{\bar{\mu}_{0}, \bar{\mu}_{1}, \ldots, \bar{\mu}_{N-1}\right\}$ be a policy such that for all $X k$ and $k$, we have

$$
\begin{equation*}
E\left\{9 k(X k, P k(X k), W k)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, \bar{\mu}_{k}\left(x_{k}\right), w_{k}\right)\right)\right\} \leq J k(X k)+\delta_{k} \tag{6.23}
\end{equation*}
$$

where $\delta_{0}, \delta_{1}, . . " \delta_{N-1}$ are some scalars. Then for all $X k$ and $k$, we have

$$
J_{\pi, k}\left(x_{k}\right) \leq J k(X k)+\sum_{i=k}^{\mathbf{N}-\mathbf{I}} \delta_{i}
$$

where $J_{\pi, k}\left(x_{k}\right)$ is the cost-to-go of $\pi$ starting from state $X k$ at stage $k$.

Proof: We use backwards induction on $k$. In particular, we have $J_{\pi, N}\left(x_{N}\right)=$ $\tilde{I} N(X N)=g N(X N)$ for all $X N$. Assuming that

$$
J_{\pi, k+1}\left(x_{k+1}\right) \leq \tilde{J}_{k+1}\left(x_{k+1}\right)+\sum_{\mathrm{i}=\mathrm{k}+1}^{\mathrm{N}-\mathbf{1}} \delta_{i}
$$

for all $X k+1$, we have

$$
\begin{aligned}
J_{\pi, k}\left(x_{k}\right) & =E\left\{9 k(X k, P k(X k), W k)+J_{\pi, k+1}\left(f_{k}\left(x_{k}, \bar{\mu}_{k}\left(x_{k}\right), w_{k}\right)\right)\right\} \\
& \leq E\{g(X k, P k(X k), W k)+J k+l(f k(X k, 71 k(X k), W k))\}+\sum_{i=k+1}^{N-1} \delta_{i} \\
& \leq J k(X k)+\delta_{k}+\sum_{\mathrm{i}=\mathrm{k}+1}^{N-1} \delta_{i},
\end{aligned}
$$

for, all $X k$. The first equality above follows from the DP algorithm that defines the costs-to-go $J_{\pi, k}$ of $\pi$, while the first inequality follows from the induction hypothesis, and the second inequality follows from the assumptIOn (6.23). This completes the induction proof. Q.E.D.

## Example 6.3.3 (Certainty Equivalent Control)

Consider the CEC for the case of a perfect state information problem, where each disturbance $w_{k}$ is fixed at a nominal value $\bar{w}_{k}, k=0, \ldots, N-\mathbf{1}$, which is independent of $X k$ and $u_{k}$. Consider the optimal value of the problem solved by the CEC at state $X k$ and stage $k$,

$$
\tilde{J}_{k}\left(x_{k}\right)=\min _{\substack{x_{i+1}=f_{i}\left(x_{i}, u_{i}, \bar{w}_{i}\right) \\ u_{i} \in U_{i}\left(x_{i}\right),=1=k, \ldots, N-1}}\left[g_{N}\left(x_{N}\right)+\sum_{i=k}^{N-1} g_{i}\left(x_{i}, u_{i}, \bar{w}_{i}\right)\right],
$$

and let $\mathrm{J} N(X N)=g N(X N)$ for all $X N$. Recall that the CEC applies the control $\bar{\mu}_{k}\left(x_{k}\right)=\bar{u}_{k}$ after finding an optimal control sequence $\left\{\bar{u}_{k}, \ldots, \bar{u}_{N-1}\right\}$ for the deterministic problem in the right-hand side above. Note also that the following DP equation

$$
\tilde{J}_{k}\left(x_{k}\right)=\min _{u_{k} \in U_{k}\left(x_{k}\right)}\left[g_{k}\left(x_{k}, u_{k}, \bar{w}_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, \bar{w}_{k}\right)\right)\right]
$$

holds, and that the control $\bar{u}_{k}$ applied by the CEC minimizes in the right-hand side.

Let us now apply Prop. 6.3.2 to derive a performance bound for the CEC. We have for all $X k$ and $k$,

$$
\begin{aligned}
J k(X k) & =g_{k}\left(x_{k}, \bar{\mu}_{k}\left(x_{k}\right), \bar{w}_{k}\right)+\mathbf{J}_{\mathrm{k}+1}\left(J k\left(x_{k}, \bar{\mu}_{k}\left(x_{k}\right), \bar{w}_{k}\right)\right) \\
& =E\left\{9(X k, \bar{\mu} k(x k), w k)+\mathbf{J} k+1\left(J k\left(X k, \bar{\mu}_{k}(X k), w_{k}\right)\right)\right\}-\gamma_{k}(x k)
\end{aligned}
$$

where $r k(X k)$ is defined by

$$
\begin{aligned}
\gamma_{k}(X k)= & E\left\{g\left(X k^{\prime} \bar{\mu}_{k}(X k), w_{k}\right)+\mathbf{J}_{\mathrm{k}+1}\left(J k\left(X k^{\prime} \bar{\mu}_{k}\left(x_{k}\right), w_{k}\right)\right)\right\} \\
& -g_{k}\left(x_{k}, \bar{\mu}_{k}\left(x_{k}\right), \bar{w}_{k}\right)-\bar{J}_{k+1}\left(f_{k}\left(x_{k}, \bar{\mu}_{k}\left(x_{k}\right), \bar{w}_{k}\right)\right) .
\end{aligned}
$$

It follows that

$$
E\left\{g\left(x_{k}, \bar{\mu}_{k}\left(x_{k}\right), w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, \bar{\mu}_{k}\left(x_{k}\right), w_{k}\right)\right)\right\} \leq \tilde{J}_{k}\left(x_{k}\right)+\delta_{k},
$$

where

$$
\delta_{k}=\operatorname{maX}_{X k} \operatorname{ra} k(X k)
$$

and by Prop. 6.3.2, we obtain the following bound for the cost-to-go function $J k(X k)$ of the CEC:

$$
J k(X k) \leq \tilde{J}_{k}\left(x_{k}\right)+\sum_{i=k}^{\mathrm{N}-1} \delta_{i} .
$$

The preceding performance bound is helpful when it can be shown that $\delta_{k} \leq 0$ for all $k$, in which case we have $\bar{J}_{k}\left(x_{k}\right) \leq J k(X k)$ for all $X k$ and $k$. This is true for example if for all $X k$ and $U k$, we have

$$
E\left\{g\left(x_{k}, u_{k}, w_{k}\right)\right\} \leq g_{k}\left(x_{k}, u_{k}, \bar{w}_{k}\right)
$$

$$
E\left\{\bar{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\} \leq \bar{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, \bar{w}_{k}\right)\right)
$$

The most common way to assert that inequalities of this type hold is via some kind of concavity assumptions; for example, the inequalities hold if the state, control, and disturbance spaces are Euclidean spaces, $\bar{w}_{k}$ is the expected value ofwk, and the functions $g\left(X k, U k,{ }^{\prime}\right)$ and $J k+l\left(f k\left(X k, U k,{ }^{\prime}\right)\right)^{\prime}$ viewed as functions of $W k$, are concave (this is known as Jensen's inequality, and at least in the case where $W k$ takes a finite number of values, follows easily from the definition of concavity). It can be shown that the concavity conditions just described are guaranteed if the system is linear with respect to $X k$ and $W k$, the cost functions $g k$ are concave with respect to $X k$ and $W k$ for each fixed $U k$, the terminal cost function $g N$ is concave, and the control constraint sets $U k$ do not depend on $X k$.

### 6.3.2 Computational Issues in Limited Lookahead

We now discuss the computation of the cost-to-go approximations and the corresponding minimization of the one-step lookahead costs.

## Minimization Using Nonlinear Programming

One approach to obtain the control $\bar{\mu}_{k}\left(x_{k}\right)$ used by the one-step lookahead policy is to exhaustively calculate and compare the one-step lookahead costs of all the controls in the set $\underline{U} k(X K)$. In some cases, there is a more efficient alternative, which is to solve a suitable nonlinear programming problem. In particular, if the control space is the Euclidean space $\Re^{m}$, then for a one-step lookahead control calculation, we are faced with a minimization over a subset of $\Re^{m}$, which may be approached by continuous optimization/nonlinear programming techniques.

It turns out that even a multistage lookahead control calculation can be approached by nonlinear programming. In particular, assume that the disturbance can take a finite number of values, say $r$. Then, it can be shown that for a given initial state, an I-stage perfect state information problem (which corresponds to an l-step lookahead control calculation) can be formulated as a nonlinear programming problem with $\mathrm{m}\left(1+r^{l},{ }^{1}\right)$ variables. We illustrate this by means of an important example where $I=2$ and then discuss the general case.

## wixample 6.3.4 (Two-Stage Stochastic Programming)

Here we want to find an optimal two-stage decision rule for the following situation: In the first stage we will choose a vector $U o$ from a subset $U o$ C $\Re^{m}$ with cost $g o$ (uo). Then an uncertain event represented by a random variable $W$ will occur, where $w$ will take one of the values wi, ..., $w^{\top}$ with corresponding probabilities pi, $\ldots, p T$. We will know the value $w^{j}$ once it occurs, and we must
then choose a vector $u_{1}^{J}$ from a subset $U l\left(u o, w^{j}\right) \mathrm{C} \Re^{m}$ at a cost $g l\left(u_{1}^{J}, w^{j}\right)$. The objective is to minimize the expected cost

$$
g o(u o)+\sum_{\mathbf{j}=\mathbf{1}} p^{j} g_{1}\left(u_{1}^{j}, w^{j}\right)
$$

subject to

$$
U o \mathrm{E} U o, \quad u_{1}^{j} \in U_{1}\left(u_{0}, w^{j}\right), \quad j=1, \ldots, r
$$

This is a nonlinear programming problem of dimension $r n(l+r)$ (the optimization variables are $\left.u o, u_{1}^{1}, \ldots, u_{1}^{r}\right)$. It can also be viewed as a two-stage perfect state information problem, where $\mathrm{x} 1=$ wo is the state equation, wo can take the values wi, ... $w^{T}$ with probabilities $\mathbf{p I}, \ldots, p T$, the cost of the first stage is $g o(u o)$, and the cost of the second stage is $g l(\mathrm{X}, \mathrm{U})$.

The preceding example can be generalized. Consider the basic problem of Chapter 1 for the case where there are only two stages $(l=2)$ and the disturbances Wo and WI can independently take one of the $r$ values $W I, \ldots, w^{T}$ with corresponding probabilities $p I, \ldots, p T$. The optimal cost function $J o(x o)$ is given by the two-stage DP algorithm

$$
\begin{aligned}
J o(x o)= & \min _{u o E U o(x o)}\left[\sum_{j=1}^{r} p^{j}\{g o(x o, u o, w J)\right. \\
& +\underset{u_{1}^{j} E U l\left(f O\left(X O, u o, w^{j}\right)\right)}{ } \min _{i=1}\left[\sum _ { i = 1 } ^ { r } p i \left\{g l\left(f o\left(x o, u o, w^{J}\right), u_{1}^{j}, W i\right)\right.\right. \\
& \left.\left.\left.\left.+g_{2}\left(f_{1}\left(f_{0}\left(x_{0}, u_{0}, w^{j}\right), u_{1}^{j}, w^{i}\right)\right)\right\}\right]\right\}\right]
\end{aligned}
$$

This DP algorithm is equivalent to solving the nonlinear programming problem
$\operatorname{minimize} \sum_{j=1}^{r} p^{j}\left\{90\left(X O U o, w^{j}\right)+\sum_{i=1}^{r} p^{i}\left\{91\left(J o\left(x o, u o, w^{j}\right), u_{1}^{j}, w i\right)\right.\right.$ $\left.\left.+92\left(f_{1}\left(J o\left(x o, u o, w^{j}\right), u_{1}^{J}, w i\right)\right)\right\}\right\}$
subject to Uo E $U o(x o), \quad u_{1}^{j} \mathrm{E} U l\left(f o\left(x o, u o, w^{J}\right)\right), \quad j=1, \ldots, \mathrm{r}$.
If the controls $U o$ and $U \prime$ are elements of $\Re^{m}$, the number of variables in the above problem is $\mathrm{m}(1+r)$. More generally, for an I-stage perfect state information problem a similar reformulation as a nonlinear programming problem requires $\mathrm{m}\left(1+r^{l} .{ }^{1}\right)$ variables. Thus if the number of lookahead stages is relatively small, a nonlinear programming approach may be the preferred option in calculating suboptimal limited lookahead policies.

## Choosing the Approximate Cost-to-Go

A key issue in implementing a limited lookahead policy is the selection of the cost-to-go approximation at the final step. It may appear important at first that the true cost-to-go function should be approximated well over the range of relevant states; however, this is not necessarily true. What is important is that the cost-to-go differentials (or relative values) be approximated well; that is, for an l-step lookahead policy it is important to have

$$
\tilde{J}_{k+l}(x)-\tilde{J}_{k+l}\left(x^{\prime}\right) \approx J_{k+l}(x)-J_{k+l}\left(x^{\prime}\right)
$$

for any two states $x$ and $x^{\prime}$ that can be generated 1 steps ahead from the current state. For example, if equality were to hold above for all $x, x^{\prime}$, then $J k+l(X)$ and $J k+l(X)$ would differ by the same constant for each relevant $x$ and the l-step lookahead policy would be optimal.

The manner in which the cost-to-go approximation is selected depends very much on the problem being solved. There is a wide variety of possibilities here. We will discuss three such approaches:
(a) Problem Approximation: The idea here is to approximate the optimal cost-to-go with some cost derived from a related but simpler problem (for example the optimal cost-to-go of that problem). This possibility is discussed and is illustrated with examples in Sections 6.3.3 and 6.3.4.
(b) Pammetric Cost-to-Go Approximation: The idea here is to approximate the optimal cost-to-go with a function of a suitable parametric form, whose parameters are tuned by some heuristic or systematic scheme. This possibility is discussed in Section 6.3.5 and is illustrated using the computer chess paradigm. Additional methods of this type are discussed in Vol. II.
(c) Rollout Approach: Here the optimal cost-to-go is approximated by the cost of some suboptimal policy, which is calculated either analytically, or more commonly, by simulation. Generally, if a reasonably good suboptimal policy is known (e.g., a certainty equivalent or open-loopfeedback controller, or some other problem-dependent heuristic), it can be used to obtain a cost-to-go approximation. This approach is also particularly well-suited for deterministic and combinatorial problems. It is discussed at length in Section 6.4.

### 6.3.3 Problem

uxamanum Enforced Decomposition
An often convenient approach for cost-to-go approximation is based on solution of a simpler problem that is tractable computationally or analytically. Here is an illustrative example, involving a convenient modification of the probabilistic structure of the problem.

## Example 6.3 .5

Consider the problem of an unscrupulous innkeeper who charges one of $m$ different rates $\mathrm{rl}, \ldots, r_{m}$ for a room as the day progresses, depending on whether he has many or few vacancies, so as to maximize his expected total income during the day (Exercise 1.25 in Chapter 1). A quote of a rate $r_{i}$ is accepted with probability $p_{i}$ and is rejected with probability $1-p_{i}$, in which case the customer departs, never to return during that day. When the number $y$ of customers that will ask for a room during the rest of the day (including the customer currently asking for a room) is known and the number of vacancies is $x$, the optimal expected income $J(x, y)$ of the innkeeper is given by the DP recursion

$$
J(x, y)=\max _{i=1, \ldots, m}\left[p_{i}\left(r_{i}+J(x-1, y-1)\right)+\left(1-p_{i}\right) J(x, y-1)\right]
$$

for all $x \geq 1$ and $y \geq 1$, with initial conditions

$$
J(x, O)=J(O, y)=0, \quad \text { for all } x \text { and } y
$$

On the other hand, when the innkeeper does not know $y$ at the times of decision, but instead only has a probability distribution for $y$, it can be seen that the problem becomes a difficult imperfect state information problem. Yet a reasonable one-step lookahead policy is based on approximating the optimal cost-to-go of subsequent decisions with $J(x-1, \bar{y}-1)$ or $J(x, \bar{y}-1)$, where the function $J$ is calculated by the above recursion and $\bar{y}$ is the closest integer to the expected value of $y$. In particular, according to this one-step lookahead policy, when the innkeeper has a number of vacancies $x \geq 1$, he quotes to the current customer the rate that maximizes $p_{i}\left(r_{i}+J(x-1, \bar{y}-1)-J(x, \bar{y}-1)\right)$.

The preceding example is based on replacing the problem uncertainty (the random variable $y$ ) with a "certainty equivalent" (the scalar $\overline{\boldsymbol{y}}$ ). The next example describes a generalization of this type of approximation, based on simplifying the stochastic structure of the problem.

## Example 6.3.6 (Approximation Using Scenarios)

One possibility to approximate the optimal cost-to-go is to use certainty equivalence, in the spirit of Section 6.1. In particular, for a given state $x_{k+1}$ at time $k+1$, we fix the remaining disturbances at some nominal values $\bar{w}_{k+1}, \ldots, \bar{w}_{N-1}$, and we compute an optimal control trajectory starting from $X k+l$ at time $k+1$. The corresponding cost, denoted by $J k+l(X k+l)$, is used to approximate the optimal cost-to-go $J k+l(X k+l)$ for the purpose of computing the corresponding one-step lookahead policy. Thus to compute the one-step lookahead control at state $X k$, we need to solve a deterministic optimal control problem from all possible next states $f_{k}\left(x_{k}, ' U k, W k\right)$ and to evaluate the corresponding optimal cost $\mathrm{J}_{\mathrm{k}+1}(J k(X k, U k, W k))$ based on the nominal values of the uncertainty.

A simpler but less effective variant of this approach is to compute $J k+l(X k+l)$ as the cost-to-go of a given heuristic (rather than optimal) policy
for the deterministic problem that corresponds to the nominal values of the uncertainty and the starting state $X k+1$. The advantage of using certainty equivalence here is that the potentially costly calculation of the expected value of the cost is replaced by a single state-control trajectory calculation.

The certainty equivalent approximation involves a single nominal trajectory of the remaining uncertainty. To strengthen this approach, it is natural to consider multiple trajectories of the uncertainty, called scenarios, and to construct an approximation to the optimal cost-to-go that involves, for every one of the scenarios, the cost of either an optimal or a given heuristic policy. Mathematically, we assume that we have a method, which at a given state $X k+l$, generates $I v l$ uncertainty sequences

$$
w^{m}\left(x_{k+1}\right)=\left(w_{k+1}^{m}, \ldots, w_{N-1}^{m}\right), \quad \mathrm{m}=1, \ldots, M
$$

These are the scenarios considered at state $x k+1$. The cost $J k+1(X k+1)$ is approximated by

$$
\begin{equation*}
J k+l(X k+l, r)=r o+\sum_{m=l}^{M} r_{m} C_{m}\left(x_{k+1}\right) \tag{6.24}
\end{equation*}
$$

where $r \quad(r o, r l, \ldots, r M)$ is a vector of parameters, and $\mathrm{C}_{\mathrm{m}}\left(x_{k+1}\right)$ is the cost corresponding to an occurrence of the scenario $W^{m}(X k+1)$, when starting from state $X k+l$ and using either an optimal or a given heuristic policy.

The parameters ro, $r l, \ldots, r \mathrm{M}$ may depend on the time index, and in more sophisticated schemes, they may depend on some characteristics of the state (see our subsequent discussion of feature-based architectures in Section 6.3.5). We may interpret the parameter $r_{m}$ as an "aggregate weight" that encodes the aggregate effect on the cost-to-go function of uncertainty sequences that are similar to the scenario $W^{m}(X k+1)$. Note that, if ro $=0$, the approximation (6.24) may also be viewed as a calculation by limited simulation, based on just the $N l l$ scenarios $W m(X k+l)$, and using the weights $r_{m}$ as "aggregate probabilities." One difficulty with this approach is that we have to choose the parameters (ro, rll, ... $r M$ ). For this, we may either use some heuristic scheme based on trial and error, or some of the more systematic schemes of neuro-dynamic programming, discussed in Vol. II.

We finally mention a variation of the scenario-based approximation method, whereby only a portion of the future uncertain quantities are fixed at nominal scenario values, while the remaining uncertain quantities are explicitly viewed as random. The cost of scenario $m$ at state $X k+l$ is now a random variable, and the quantity $C m(X k+l)$ used in Eq. (6.24) should be the expected cost of this random variable. This variation is appropriate and makes practical sense as long as the computation of the corresponding expected scenario costs $\mathrm{C}_{\mathrm{m}}\left(x_{k+1}\right)$ is convenient.

## Enforced Hecomposition of Weakly Coupled Systems

The simplification/approximation approach is often well-suited for problems that involve a number of subsystems that may be coupled through
the system equation, or the cost function, or the control constraints, but the degree of coupling is "relatively weak." It is difficult to define precisely what constitutes "weak coupling," but in specific problem contexts, usually this type of structure is easily recognized. For such problems it is often sensible to introduce approximations by artificially decoupling the subsystems in some way, thereby creating either a simpler problem or a simpler cost calculation, where subsystems can be dealt with in isolation. There are a number of different ways to effect this type of artificial decomposition, and the best approach is often problem-dependent.

As an example consider a deterministic problem, where the control $U k$ at time $k$ consists of $r n$ components, $U k=\left\{u_{k}^{1}, \ldots, u_{k}^{m}\right\}$, with $u_{k}^{i}$ corresponding to the ith subsystem. Then to compute a cost-to-go approximation at a given state $x k$, one may try a one-subsystem-at-a-time approach: first optimize over the control sequence $\left\{u_{k}^{1}, U_{k+}^{1} l^{\prime} \ldots, U_{N}^{1}-\mathrm{I}\right\}$ of the first subsystem, while keeping the controls of the remaining subsystems at some nominal values, then minimize over the controls of the second subsystem, while keeping the controls of the first subsystem at the "optimal" values just computed and the controls of subsystems $3, \ldots, r n$ to the nominal values, and continue in this manner. There are several possible variations, for example to make the order in which the subsystems are considered subject to optimization as well. Let us illustrate this approach by means of an example.

## Example 6.3.7 (Vehicle Routing)

Consider m vehicles that move along the arcs of a given graph. Each node of the graph has a given "value" and the first vehicle that will pass through the node will collect its value, while vehicles that pass subsequently through the node do not collect any value. This may serve as a model of a situation where there are various valuable tasks to be performed at the nodes of a transportation network, and each task can be performed at most once and by a single vehicle. We assume that each vehicle starts at a given node and after at most a given number of arc moves, it must return to some other given node. The problem is to find a route for each vehicle satisfying these constraints, so that the total value collected by the vehicles is maximized.

This is a difficult combinatorial problem that in principle can be approached by DP. In particular, we can view as the state the set of current positions of the vehicles together with the list of nodes that have been traversed by some vehicle in the past, and have thus "lost" their value. Unfortunately, the number of these states is enormous (it increases exponentially with the number of nodes and the number of vehicles). The version of the problem involving a single vehicle, while still difficult in principle, can often be solved in reasonable time either exactly by DP or fairly accurately using a heuristic. Thus a one step-lookahead policy suggests itself, with the value-to-go approximation obtained by solving single vehicle problems.

In particular, in a one step-lookahead scheme, at a given time $k$ and from a given state we consider all possible kth moves by the vehicles, and at
the resulting states we approximate the optimal value-to-go with the value corresponding to a suboptimal set of paths. These paths are obtained as follows: we fix an order of the vehicles and we calculate a path for the first vehicle, assuming the other vehicles do not move. (This is done either optimally by DP, or nearly optimally using some heuristic.) Then we calculate a path for the second vehicle in the order, taking into account the value collected by the first vehicle, and we similarly continue: for each vehicle, we calculate in the given order a path, taking into account the value collected by the preceding vehicles. We end up with a set of paths that have a certain total value associated with them, and which correspond to the particular order for considering the vehicles. We can also calculate other sets of paths and their corresponding total values, for other orders (possibly all orders) for considering the vehicles. We then use as the value-to-go approximation at the given state the maximal value over all the sets of paths computed from that state.

Another context where enforced decomposition may be an attractive possibility, is when the subsystems are coupled only through the disturbance. In particular, consider $m$ subsystems of the form

$$
x_{k+1}^{i}=f^{i}\left(x_{k}^{i}, u_{k}^{i}, w_{k}^{i}\right), \quad \mathrm{i}=1, \ldots, m .
$$

Here the ith subsystem has its own state $x_{k}^{i}$, control $u_{k}^{i}$, and cost per stage $g^{i}\left(x_{k}^{2}, u_{k}^{2}, w_{k}^{2}\right)$, but the probability distribution of $w_{k}^{i}$ depends on the full state

$$
x_{k}=\left(x_{k}^{1}, \ldots, x_{k}^{m}\right)
$$

A natural form of suboptimal control is to solve at each stage $k$ and for each $i$, the ith subsystem optimization problem where the probability distribution of each of the future disturbances $w_{k+1}^{i}, \ldots w_{N-1}^{i}$ is fixed at some distribution that depends only on the corresponding "local" states $x_{k+1}^{i}, \ldots, x_{N^{-1}}^{i}$. This distribution may be derived on the basis of some nominal values $\bar{x}_{k+1}^{j}, \ldots, \bar{x}_{N^{-1}}^{j} j \neq \mathrm{i}$, of the future states of the other subsystems, and these nominal values may in turn depend on the full current state $x k$. The first control $\bar{u}_{k}^{i}$ in the optimal policy thus obtained is applied at the ith subsystem in stage $k$, and the remaining portion of this policy is discarded.

Let us also discuss in some detail an example of subsystem decomposition where the coupling comes through the control constraint.

## Example 6.3.8 (Flexible Manufacturing)

Flexible manufacturing systems (FMS) provide a popular approach for increasing productivity in manufacturing small batches of related parts. There are several workstations in an FMS, and each is capable of carrying out a variety of operations. This allows the simultaneous manufacturing of more than one part type, reduces idle time, and allows production to continue even when a workstation is out of service because of failure or maintenance.

Consider a work center in which $n$ part types are produced. Denote $u_{k}^{i}$ : the amount of part i produced in period $k$.
$d_{k}^{i}$ : a known demand for part i in period $k$.
$x_{k}^{i}$ : the cumulative difference of amount of part i produced and demanded up to period $k$.
Let us denote also by $x k, U k, d_{k}$ the n -dimensional vectors with coordinates $x_{k}^{i}, u_{k}^{i}, d_{k}^{i}$, respectively. We then have

$$
\begin{equation*}
x_{k+1}=x_{k}+u_{k}-d_{k} \tag{6.25}
\end{equation*}
$$

The work center consists of a number of workstations that fail and get repaired in random fashion, thereby affecting the productive capacity of the system (i.e., the constraints on $U k$ )' Roughly, our problem is to schedule part production so that $x k$ is kept around zero to the extent possible.

The productive capacity of the system depends on a random variable $\alpha_{k}$ that reflects the status of the workstations. In particular, we assume that the production vector $u_{k}$ must belong to a constraint set $U\left(\alpha_{k}\right)$. We model the evolution of $\alpha_{k}$ by a Markov chain with known transition probabilities $P\left(\alpha_{k+1} \mid \alpha_{k}\right)$. In practice, these probabilities must be estimated from individual station failure and repair rates, but we will not go into the matter further. Note also that in practice these probabilities may depend on $U k$. This dependence is ignored for the purpose of development of a cost-to-go approximation. It may be taken into account when the actual suboptimal control is computed.

We select as system state the pair ( $X k, \alpha_{k}$ ), where $x k$ evolves according to Eq. (6.25) and $\alpha_{k}$ evolves according to the Markov chain described earlier. The problem is to find for every state $\left(X k^{\prime} \alpha_{k}\right)$ a production vector $u_{k} \mathrm{E} U\left(\alpha_{k}\right)$ such that a cost function of the form

$$
J_{\pi}\left(x_{0}\right)=E\left\{\sum_{k=0}^{N-1} \sum_{i=1}^{n} g^{i}\left(x_{k}^{i}\right)\right\}
$$

is minimized. The function $g i$ expresses the desire to keep the current backlog or surplus of part $i$ near zero. Two examples are $g i\left(X^{i}\right)=\beta_{i}\left|x^{i}\right|$ and $g i\left(X^{i}\right)=$ $\beta_{i}\left(x^{i}\right)^{2}$, where $\beta_{i}>0$.

The DP algorithm for this problem is

$$
\begin{align*}
J k\left(X k, \alpha_{k}\right) & \sum_{\mathrm{i}=1} g^{i}\left(x_{k}^{i}\right)  \tag{6.26}\\
& +\min _{u_{k} \in U\left(\alpha_{k}\right) \alpha_{k+1}} E\left\{J k+l\left(X k+U k \quad d_{k}, \alpha_{k+1}\right) \mid \alpha_{k}\right\} .
\end{align*}
$$

If there is only one part type ( $n=1$ ), the optimal policy can be fairly easily determined from this algorithm (see Exercise 6.7). However, in general, the algorithm requires a prohibitive amount of calculation for an FMS of realistic size (say for $n>10$ part types). We thus consider a one-step lookahead policy


Figure 6.3.1 Inner and outer approximations of the production capacity constraint set by hypercubes in the flexible manufacturing example.
with the cost-to-go $J k+l$ replaced by an approximation $\tilde{J}_{k+1}$ that exploits the nearly separable structure of our problem.

In particular, we note that the problem can to a large extent be decomposed with respect to individual part types. Indeed, the system equation (6.25) and the cost per stage are decoupled, and the only coupling between parts comes from the constraint $U k$ E $U\left(\alpha_{k}\right)$. Suppose we approximate $U\left(\alpha_{k}\right)$ by inner and outer approximating hypercubes $\underline{U}\left(\alpha_{k}\right)$ and $\bar{U}\left(\alpha_{k}\right)$ of the form

$$
\begin{gathered}
\underline{U}\left(\alpha_{k}\right)=\left\{u_{k}^{i} \mid 0 \leq u_{k}^{i} \leq \underline{B}_{i}\left(\alpha_{k}\right)\right\} \\
\widetilde{U}\left(\alpha_{k}\right)=\left\{u_{k}^{i} \mid 0 \leq u_{k}^{i} \leq \bar{B}_{i}\left(\alpha_{k}\right)\right\} \\
\underline{U}\left(\alpha_{k}\right) \text { С } U\left(\alpha_{k}\right) \text { С } \bar{U}\left(\alpha_{k}\right)
\end{gathered}
$$

as shown in Fig. 6.3.1. If $U\left(\alpha_{k}\right)$ is replaced for each $\alpha_{k}$ by either $\bar{U}\left(\alpha_{k}\right)$ or $\underline{U}\left(\alpha_{k}\right)$, then the problem is decomposed completely with respect to part types. For every part i the DP algorithm for the outer approximation is given by

$$
\begin{align*}
\bar{J}_{k}^{2}\left(x_{k}^{i}, \alpha_{k}\right)= & g^{i}\left(x_{k}^{i}\right) \\
& +\min _{0 \leq u_{k}^{i} \leq \bar{B}_{i}\left(\alpha_{k}\right)} E\left\{\begin{array}{c}
k+1
\end{array}\right. \tag{6.27}
\end{align*}
$$

and for the inner approximation it is given by

$$
\begin{align*}
\left.\alpha_{k}\right)= & g^{i}\left(x_{k}^{i}\right) \\
& +\min _{0 \leq u_{k}^{i} \leq \underline{B}_{i}\left(\alpha_{k}\right) \alpha_{k+1}} E\left\{\underline{J}_{k+1}^{i}\left(x_{k}^{i}+u_{k}^{i}-d_{k}^{i}, \alpha_{k+1}\right) \mid \alpha_{k}\right\} . \tag{6.28}
\end{align*}
$$

Furthermore, since $\underline{U}\left(\alpha_{k}\right) \subset U\left(\alpha_{k}\right) \subset \bar{U}\left(\alpha_{k}\right)$, the cost-to-go functions $\bar{J}_{k}^{2}$ and $\underline{J}_{k}^{i}$ provide lower and upper bounds to the true cost-to-go function $J_{k}$,

$$
\sum_{i=1}^{n} \bar{J}_{k}^{i}\left(x_{k}^{i}, \alpha_{k}\right) \leq J k\left(X k, \alpha_{k}\right) \leq \sum_{\mathbf{i}=1} \underline{J}_{k}^{2}\left(x_{k}^{i}, \alpha_{k}\right)
$$

and can be used to construct approximations to $J k$ that are suitable for a one-step lookahead policy. A simple possibility is to adopt the averaging approximation

$$
\tilde{J}_{k}\left(x_{k}, \alpha_{k}\right)=\frac{1}{2} \sum_{i=1}^{n}\left(\bar{J}_{k}^{i}\left(x_{k}^{i}, \alpha_{k}\right)+\underline{J}_{k}^{i}\left(x_{k}^{i}, \alpha_{k}\right)\right)
$$

and use at state $\left(X k, \alpha_{k}\right)$ the control $\tilde{u}_{k}$ that minimizes [ef. Eq. (6.26)]
$E\left\{\sum_{i=1}^{n}\left(\bar{J}_{k+1}^{i}\left(x_{k}^{i}+u_{k}^{i}-d_{k}^{i}, \alpha_{k+1}\right)+\underline{J}_{k+1}^{i}\left(x_{k}^{i}+u_{k}^{i}-\alpha_{k+1}\right)\right) \mid \alpha_{k}\right\}$
over all $U k$ E $U\left(\alpha_{k}\right)$. Multiple upper bound approximations, based on multiple choices of $\underline{B}_{i}\left(\alpha_{k}\right)$, can also be used.

To implement this scheme, it is necessary to compute and store the approximate cost-to-go functions $\bar{J}_{k}^{i}$ and $\underline{J}_{k}^{i}$ in tables, so that they can be used in the real-time computation of the suboptimal control via the minimization of expression (6.29). The corresponding calculations [cf. the DP algorithms (6.27) and (6.28)] are nontrivial, but they can be carried out off-line, and in any case they are much less than what would be required to compute the optimal controller. The feasibility and the benefits of the overall approach have been demonstrated by simulation in the thesis by Kimemia [Kim82], on which this example is based. See also Kimemia, Gershwin, and Bertsekas [KGB82], and Tsitsiklis [Tsi84a].

For some other examples of decomposition approaches, see Wu and Bertsekas [WuB99], which deals with admission control in cellular communication networks, and Meuleau et. al. [MHK98], which deals with problems of resource allocation

## 6.3 .4

An alternative method for constructing a simpler and more tractable problem is based on reducing the number of states by "combining" many of them together into aggregate states. This results in an aggregate problem, with fewer states, which may be solvable by exact DP methods. The optimal cost-to-go functions of the aggregate problem is then used to construct a one-step-lookahead cost approximation for the original problem. The precise form of the aggregate problem may depend on intuition and/or heuristic reasoning, based on our understanding of the original problem.

In this subsection, we will discuss various aggregation methods, starting with the case of a finite-state problem. We will focus on defining the transition probabilities and costs of the aggregate problem, and to simplify notation, we suppress the time indexing in what follows. We generically denote:
$I, \bar{I}$ : The set of states of the original system at the current and the next stage, respectively.

Pij (u): The transition probability of the original system from state i E I to state $j \in \bar{I}$ under control $u$.
$g(\mathrm{i}, u, j)$ : The transition cost of the original system from state i E $I$ to state $j$ E $\bar{I}$ under control $u$.
$S, \bar{S}:$ The set of states of the aggregate system at the current and the next stage, respectively.
$r_{s t}(u)$ : The transition probability of the aggregate system from state $s \mathrm{E} S$ to state $t \mathrm{E} \bar{S}$ under control $u$.
$h(s, u)$ : The expected transition cost of the aggregate system from state $s$ E $S$ under control $u$.
For simplicity, we assume that the control constraint set $U(i)$ is the same for all states i E $I$. This common control constraint set, denoted by $U$, is chosen as the control constraint set at all states $s$ E $S$ of the aggregate problem.

There are several types of aggregation methods, which bring to bear intuition about the problem's structure in different ways. All these methods are based on two (somewhat arbitrary) choices of probabilities, which relate the original system states with the aggregate states:
(1) For each aggregate state $s$ E $S$ and original system state i E $I$, we specify the disaggregation pmbability qsi (we have $\sum_{i \in I} q s i=1$ for each $s$ E S). Roughly, qsi may be interpreted as the "degree to which $s$ is represented by i."
(2) For each original system state $j \mathrm{E} \bar{I}$ and aggregate state $t \mathrm{E} \bar{S}$, we specify the aggregation pmbability Wjt (we have $\sum_{t \in \bar{S}} W j t \quad 1$ for each $j \mathrm{E} \bar{I}$ ). Roughly, Wjt may be interpreted as the "degree of membership of $j$ in the aggregate state $t . "$

Note that in general, the disaggregation and aggregation probabilities may change at each stage (since the state space may change at each stage). On the other hand, for a stationary problem, where state and control spaces, system equation, and cost per stage that are the same for all stages, the disaggregation and aggregation probabilities will ordinarily also be the same for all stages.

As an illustration consider the following example of aggregation.

## 6.3 .9 (Hard Aggregation)

We are given a partition of the original system state spaces $I$ and $\bar{I}$ into subsets of states (each state belongs to one and only one subset). We view each subset as an aggregate state. This corresponds to aggregation probabilities

$$
W j t=1 \quad \text { if state } j \mathrm{E} \bar{I} \text { belongs to aggregate state/subset } t \mathrm{E} \bar{S},
$$

and (assuming all states that belong to aggregate state/subset $s$ are "equally representative") disaggregation probabilities

$$
q_{s i}=1 / n_{s} \quad \text { if state i E } I \text { belongs to aggregate state/subset } s \mathrm{E} S \text {, }
$$

where $n_{s}$ is the number of states of $s$.
Given the disaggregation and aggregation probabilities, qsi and Wjt , and the original transition probabilities $\operatorname{Pij}(u)$, we envisage an aggregate system where state transitions occur as follows:
(i) From aggregate state $s$, generate state i according to $q s i$.
(ii) Generate a transition from i to j according to $\operatorname{Pij}(U)$, with cost $g(i, u, j)$.
(iii) From state $j$, generate aggregate state $t$ according to $W j t$.

Then, the transition probability from aggregate state $s$ to aggregate state $t$ under $u$, and the corresponding expected transition cost, are given by

$$
\begin{gathered}
T s t(U)=\sum_{i E I} \sum_{j \in \bar{I}} q s i P i j(U) W j t, \\
h(s, u)=\sum_{i E I} \sum_{j \in \bar{I}} q s i P i j(U) g(i, u, j) .
\end{gathered}
$$

These transition probabilities and costs define the aggregate problem. After solving for the optimal costs-to-go $J(t), t \mathrm{E} \bar{S}$, of the aggregate problem, the costs of the original problem are approximated by

$$
\begin{equation*}
J(j)=\sum_{t \in \bar{S}} W j t J(t), \quad j \in I \tag{6.30}
\end{equation*}
$$

As an illustration, for the preceding hard aggregation Example 6.3.9, the aggregate system transition process works as follows: Starting from an aggregate state/subset $s$, we generate with equal probability a state i in $s$, then a next state j E $\bar{I}$ according to the transition probabilities $\operatorname{Pij}(u)$, and then we declare as next aggregate state the subset $t$ to which j belongs. The corresponding transition probability and expected transition cost are

$$
r_{s t}(u)=\frac{1}{n^{s}}{ }_{i E s}^{\mathrm{jEt}} \quad \operatorname{Pij}(U),
$$



Figure 6.3.1 Disaggregation and aggregation probabilities in the hard and soft aggregation Examples 6.3.9 and 6.3.10. The difference is that in soft aggregation, some of the aggregation probabilities are strictly between 0 and 1 , i.e., a state of the original system may be associated with multiple aggregate states.

$$
h(s, u)=\frac{1}{n_{\mathrm{s}}} \sum_{\mathrm{iEs}} \sum_{j \in \bar{I}} P \cdot i j(u) g(i, u, j)
$$

After computing the optimal costs $J(t)$ of the aggregate problem, we use Eq. (6.30) to obtain the approximate cost function $J(j)$, which for this hard aggregation example is piecewise constant, with all states $j E \bar{S}$ that belong to the same aggregate state/subset $t$ having the same value of $J(j)$.

### 6.3.10 (Soft Aggregation)

In hard aggregation, the aggregate states/subsets are disjoint, and each original system state is associated with a single aggregate state. A generalization is to allow the aggregate states/subsets to overlap, with the aggregation probabilities $W_{j t}$ quantifying the "degree of membership" of $j$ in the aggregate state $t$. Thus, an original system state $j$ may be a member of multiple aggregate states/subsets $t$, and if this is so, the aggregation probabilities $W_{j t}$ will be positive but less than 1 for all $t$ that contain $j$ (we should still have $w_{j t}=1$ ); see Fig. 6.3.I.
example, assume that we are dealing with a queue that has space for 100 customers, and that the state is the number of spaces occupied at a given time. Suppose that we introduce four aggregate states: "nearly empty"
( $0-10$ customers), "lightly loaded" (11-50 customers), "heavily loaded" (51-90 customers), and "nearly full" (91-100 customers). Then it makes sense to use soft aggregation, so that a state with close to 50 customers is classified neither as "lightly loaded" nor as "heavily loaded," but is viewed instead as associated with both of these aggregate states, to some degree.

It can be seen from Eq. (6.30), that in soft aggregation, the approximate cost-to-go function $\mathbf{J}$ is not piecewise constant, as in the case of hard aggregation, and varies "smoothly" along the boundaries separating aggregate states. This is because original system states that belong to multiple aggregate states/subsets have approximate cost-to-go that is a convex combination of the costs-to-go of these aggregate states.

The choices of aggregate states, and aggregation and disaggregation probabilities are often guided by insight into the problem's structure. The following is an example.

## Example 6.3.11 (Admission Control in a Service Facility)

Consider a facility that serves $m$ types of customers. At each time period, a single customer of each type arrives, and requires a level of service per unit time that is random and has a given distribution (which depends on the customer type). The facility must decide which of the arriving customers to admit at each time period. The total service capacity of the facility is given, and must at all times be no less than the sum of service levels of the customers currently within the facility. An admitted customer of a given type has a given probability of leaving the facility at each time period, independently of his required. level of service and of how long he has already been in the facility. An admitted customer also pays to the facility a given amount per period, which is proportional to his required level of service (with the constant of proportionality depending on the customer type). The objective here is to maximize the expected revenue of the facility over a given horizon. Thus the tradeoff is to provide selective preference to high-paying customer types, or alternatively expressed, to avoid filling up the facility with low-paying long-staying customers, thereby potentially blocking out some high-paying customers.

In a problem of this type, the system state, just prior to making an admission decision, is the entire list of customers of each type within the facility as well as their service levels (together with the list of service levels of the customers that have just arrived). There is clearly an extremely large number of states. Intuition suggests here that it is appropriate to aggregate all customers of a given type, and represent them with their total service level. Thus the aggregate state in this approach is the list of total required service level of each type within the facility (together with the list of service levels of the customers that have just arrived - an uncontrollable state component, d. Section 1.4). Clearly it is much easier to deal with such a space of aggregate states in a DP context.

The choice of aggregation probabilities here is clear, since any original system state maps naturally into a unique aggregate state. The rationale for specifying the disaggregation probabilities, while somewhat arbitrary, is
guided by intuition into the problem structure. Given an aggregate state (a total level of service for each customer type within the facility), we must generate through the disaggregation probabilities, a "representative" list of customers of each type. It is not difficult to devise reasonable heuristic methods for doing so. The disaggregation and aggregation probabilities, together with the transition probabilities of the original system, specify the transition probabilities and the expected cost per stage of the aggregate problem.

Some simplification techniques, aimed at reducing the complexity of the DP computation, can be interpreted in terms of aggregation. The following is an example.

## 

A technique often used to reduce the computational requirements of DP is to select a small collection $S$ of states from $I$ and a small collection $\bar{S}$ of states from $\bar{I}$, and define an aggregate problem whose states are those in Sand $\bar{S}$. The aggregate problem is then solved by DP and its optimal costs are used to define approximate costs for all states in I and $\bar{I}$. This process, is known as coarse grid approximation, and is motivated by the case where the original state spaces $I$ and $\bar{I}$ are dense grids of points obtained by discretization of continuous state spaces, while the collections $S$ and $\bar{S}$ represent coarser subgrids.

The aggregate problem may be formalized by specifying the disaggregation probabilities to associate the states of $S$ with themselves (since $S \subset 1$ ):

$$
q s s=1, \quad q s i=0 \text { if } \mathrm{i} \neq s, \quad s \mathrm{E} S .
$$

The aggregation probabilities are used to represent each state in $\bar{I}$ as a probabilistic mixture of states in $\bar{S}$, respectively, possibly using some geometrical attribute of the state space

The aggregation methodology also applies to problems with an infinite number of states. The only difference is that for each aggregate state, the disaggregation probabilities are replaced by a disaggregation distribution over the original system's state space. Among other possibilities, this type of aggregation provides methods for discretizing continuous state space problems, as illustrated by the following example.

## Example 6.3.13 (Discretization of Continuous State Spaces)

Assume for simplicity a stationary problem, where the state space of the original problem is a bounded region of a Euclidean space. The idea here is to discretize this state space using some finite $\operatorname{grid}\left\{X l, \ldots, x^{M}\right\}$, and then to express each nongrid state by a linear interpolation of nearby grid states. By this we mean that the grid states $x l, \ldots, x^{M}$ are suitably selected within the state space, and each nongrid state $x$ is expressed as

$$
x=\sum_{r n=l}^{M} w m(x) x^{m}
$$

for some nonnegative weights $w m(x)$, which add to 1 and are chosen on the basis of some geometric considerations. We view the weights $w m(x)$ as aggregation probabilities, and we specify the disaggregation probabilities to associate the grid states with themselves, i.e.,

$$
q_{x} m_{x} m_{2}=1, \quad \text { for all } \mathrm{m}
$$

similar to the coarse grid approach of Example 6.3.12.
The aggregation and disaggregation probabilities just given specify the aggregate problem, which has a finite state space, the set $\left\{X l, \ldots, x^{M}\right\}$, and can be solved by DP to obtain the corresponding optimal costs-to-go $\hat{J}_{k}\left(x^{m}\right)$, $\mathrm{m}=1, \ldots, M$, for each stage $k$. Then the cost-to-go of each nongrid state $x$ at stage $k$ is approximated by

$$
J k(x)=\sum_{\mathrm{m}=1}^{\mathrm{M}} w m(x) J k\left(x^{m}\right)
$$

We finally note that one may address the solution of the aggregate problem itself by some suboptimal method, thereby introducing an additional layer of approximation in the solution of the original problem.

### 6.3.5 Parametric Cost-to-Go Approximation

The idea here is to select from within a parametric class of functions, some cost-to-go approximations $J_{k}$, which will be used in a limited lookahead scheme in place of the optimal cost-to-go functions $J_{k}$. Such parametric classes of functions are called approximation architectures, and are generically denoted by $J(x, r)$, where $x$ is the current state and $r=\left(\mathrm{rl}, \ldots, r_{m}\right)$ is a vector of "tunable" scalar parameters, also called weights (to simplify notation, we suppress the time indexing in what follows). adjusting the weights, one can change the "shape" of the approximation so that it is reasonably close to the true optimal cost-to-go function.

There is an extensive methodology for the selection of the weights. The simplest and often tried approach is to do some form of semi-exhaustive or semi-random search in the space of weight vectors and adopt the weights that result in best performance of the associated one-step lookahead controller. Other more systematic approaches are based on various forms of cost-to-go evaluation and least squares fit. We will discuss such approaches briefly here and more extensively in Vol. II in the context of the methodology of neuro-dynamic programming; see also the books by Bertsekas and Tsitsiklis [BeT96], and Sutton and Barto [SuB98].

There is also a large variety of approximation architectures, involving for example polynomials, neural networks, wavelets, various types of basis functions, etc. We provide a brief discussion of architectures based on extraction of features, and we refer to the specialized literature (e.g., Bertsekas and Tsitsiklis [BeT96], Bishop [Bis95], Haykin [Hay99], Sutton and Barto [SuB98]) for detailed discussions.

## pproximation Architectures Based on Feature Extraction

Clearly, for the success of the cost function approximation approach, it is very Important to select a class of functions $J(x, r)$ that is suitable for the problem at hand. One particularly interesting type of cost approximatiOn is provided by feature extraction, a process that maps the state $x$ into some other vector $y(x)$, called the feature vector associated with state $x$. The vector $y(x)$ consists of scalar components $\mathrm{Yl}(x), \ldots, Y m(x)$, called feat'ures. These features are usually handcrafted based on whatever human intelligence, insight, or experience is available', and are meant to capture the most important aspects of the current state $x$. A feature-based cost approximation has the form

$$
J(x, r)=J(y(x), r),
$$

where $r$ is a parameter vector. Thus, the cost approximation depends on the state $x$ through its feature vector $y(x)$ (see Fig. 6.3.2).


Figure 6.3.2 Using a feature extraction mapping to generate an input to a cos approxImator

The idea is t.hat the cost-to-go function J to be approximated may be a highly. complicated nonllnear map and it is sensible to try to break Its complexIty mto smaller, less complex pieces. Ideally, the features will encode much of the nonlinearity that is inherent in $\mathbf{J}$, and the approximation may be qUİte accurate without requiring a complicated function $J$. For example, with a well-chosen feature vector $y(x)$, a good approximation to the cost-to-go is often provided by linearly weighting the features, i.e.,

$$
J(x, r)=J(y(x), r)=\sum_{i=1}^{m} r_{i} y_{i}(x),
$$

where $r_{1}, \ldots, r_{\mathrm{m}}$ is a set of tunable scalar weights.
Note that the use of a feature vector implicitly involves the grouping of states mto the subsets that share the same feature vector, i.e., the subsets

$$
S v=\{x I Y(x)=v\},
$$

where $v$ is possible value of $y(x)$. These subsets form a partition of the state space, and the approximate cost-to-go function $J(y(x), r)$ is piecewise
constant with respect to this partition; that is, it assigns the same cost-to-go value $\mathbf{J}(v, r)$ to all states in the set $S v$. This suggests that a feature extraction mapping is well-chosen if states that have the same feature vector have roughly similar optimal cost-to-go.

A feature-based architecture is also related to the aggregation methodology of Section 6.3.4. In particular, suppose that we introduce $m$ aggregate states, $1, \ldots, \mathrm{~m}$, and associated aggregation and disaggregation probabilities. Let $r_{i}$ be the optimal cost-to-go associated with aggregate state i. Then, the aggregation methodology yields the linear parametric approximation

$$
J(x, r)=\sum_{i=1}^{m} \operatorname{TiYi}(X),
$$

where $Y i(x)$ is the aggregation probability associated with state $X$ and aggregate state i. Thus, within this context, the aggregation probability $Y i(X)$ may be viewed as a feature, which, roughly speaking, specifies the "degree of membership of $x$ to aggregate state i."

We now illustrate the preceding concepts with a detailed discussion of computer chess, where feature-based approximation architectures play an important role.

## Computer Chess

Chess-playing computer programs are one of the more visible successes of artificial intelligence. Their underlying methodology relies provides an interesting case study in the use of approximate DP. It involves the idea of limited lookahead, but also illustrates some DP ideas that we have not had much opportunity to look at in detail. These are the idea of forward depthfirst search, an important memory-saving technique that was discussed in Section 2.3 in the context of label correcting methods, and the idea of alpha-beta pruning, which is an effective method for reducing the amount of computation needed to find optimal strategies in competitive games.

The fundamental paper on which all computer chess programs are based was written by one of the most illustrious modern-day applied mathematicians, C. Shannon [Sha50]. It was argued by Shannon that whether the starting chess position is a win, loss, or draw is a question that can be answered in principle, but the answer will probably never be known. He estimated that, based on the chess rule requiring at least one pawn advance or capture within every 50 moves (otherwise a draw is declared), there are on the order of $10^{120}$ different possible sequences of moves in a chess game. He concluded that to examine these and select the best initial move for White would require $10^{90}$ years of a "fast" computer's time (fast here relates to the standards of the '50s, but the number $10^{90}$ is overwhelming even by today's standards). As an alternative, Shannon proposed a limited lookahead of a few moves and eval'uating the end positions by means of a
scoring junction. The scoring function may involve, for example, the calculation of a numerical value for each of a set of major features of a position (such as material balance, mobility, pawn structure, and other positional factors), together with a method to combine these numerical values into a single score. Thus, we may view a scoring function as a feature-based architecture for evaluating a chess position/state.

Consider first a one-move lookahead strategy for selecting the first move in a given position $P$. Let $M_{1}, \ldots, N I_{r}$ be all the legal moves that can be made in position $P$ by the side to move. Denote the resulting positions by $M_{1} P, \ldots, N I r P$, and let $S\left(M_{1} P\right), \ldots, S(M r P)$ be the corresponding scores (the convention here is that White is favored in positions with high score, while Black is favored in positions with low score). Then the move selected by White (Black) in position $P$ is the move with maximum (minimum) score. This is known as the backed-up score of $P$ and is given by

$$
B S(P)=\left\{\begin{array}{rll}
\max \{S(M I P), & \left., S\left(M_{r} P\right)\right\} & \text { if White is to move in } P, \\
\min \left\{S\left(M_{1} P\right),\right. & , S(M r P)\} & \text { if Black is to move in } P .
\end{array}\right.
$$

This process is illustrated in Fig. 6.3.3.


Figure 6.3.3 A one-move lookahead tree. If White moves at position $P$, the best move is $M_{1}$ and the backed-up score is +5 . If Black moves in position $P$, the best move is $M_{3}$, and the backed-up score of Pis -3 .

Consider next a two-move lookahead strategy in a given position $P$. Assume for concreteness that vVhite moves, and let the legal moves be $M l, \ldots, N I r$ and the corresponding positions be $M 1 P, \ldots, M r P$. Then in each of the positions $M i P, \mathrm{i}=1, \ldots, r$, apply the one-move lookahead strategy with Black to move. This gives a best move and a backed-up score $B S(N I i P)$ for Black in each of the positions $M i P, \mathrm{i}=1, \ldots, r$. Finally, based on the backed-up scores $B S(M 1 P), \ldots, B S(M r P)$, apply a one-move lookahead strategy for White, thereby obtaining the best move at position $P$ and a backed-up score for position $P$ of


Figure 6.3.4 A two-move lookahead tree with 'White to move. The backed-up scores are shown in parentheses. The best initial move is $M_{3}$ and the principal continuation is $\left(M_{3}, \text { M32 }\right)^{\prime}$


Figure 6.3.5 A four-move lookahead tree with White to move. The backed-up scores are shown in parentheses. The best initial move is NII. The principal continuation is heavily shaded.

The sequence of best moves is known as the principal continuation. The process is illustrated in 6.3.4. It is clear that Shannon's method as just described can be generalized for an arbitrary number of lookahead moves (see Fig. 6.3.5).

Generally, to evaluate the best move at a given position and the corresponding backed-up score using lookahead of $n$ moves, one can use the following DP-like procedure:


Figure 6.3.6 Traversing a tree in depth-first fashion. The numbers indicate the order in which the scores of the terminal positions and the backed-up scores of the intermediate positions are evaluated.


Figure 6.3.7 Storage requirements of the depth-first version of the minimax algorithm for the tree of Fig. 6.3.6. At the time that the terminal position marked by a checkmark is scored, only the solid-line moves are stored in memory. The dotted-line moves have been generated and purged from memory. The broken-line moves have not been generated as yet. The memory required grows linearly with the depth of the lookahead.

1. Evaluate the scores of all possible positions that are $n$ moves ahead from the given position $P$.
2. Using the scores of the terminal positions just evaluated, calculate the backed-up scores of all possible positions that are $n-1$ moves ahead from $P$.
3. For $k=1, \ldots, n-1$, using the backed-up scores of all possible po-sitions that are $n-k$ moves ahead from $P$, calculate the backed-up scores of all positions that are $n-k-1$ moves ahead from $P$.
The above procedure requires a lot of memory storage even for a modest number of lookahead moves. Shannon pointed out that with an alternative but mathematically equivalent calculation, the amount of memory required grows only linearly with the depth of lookahead, thereby allowing chess programs to operate in limited-memory microprocessor systems. This is accomplished by searching the tree of moves in depth-first fashion, and by generating new moves only when needed, as illustrated in Figs. 6.3.6 and 6.3.7. It is only necessary to store the one move sequence under current examination together with one list of legal moves at each level of the search tree. The precise algorithm is described by the following routine, which calls itself recursively.

## Minimax Algorithm

To determine the backed-up score $B S(n)$ of position $n$, do the following:

1. If $n$ is a terminal position return its score. Otherwise:
2. Generate the list of legal moves at position $n$ and let the corresponding positions be $n \mathrm{I}, \ldots, n_{r}$. Set the tentative backed-up score $T B S(n)$ of position $n$ to oo if it is vVhite's turn to move at $n$ and to -oo if it is Black's turn to move at $n$.
3. $\operatorname{For} \mathrm{i}=1, \ldots, r$, do:
a. Determine the backed-up score $B S(n i)$ of position $n i$.
b. If it is White's move at position $n$, set
$T B S(n) \quad \max _{\{ }\{T B S(n), B S(n i)\} ;$
if it is Black's move at position $n$, set

$$
T B S(n):=\min \{T B S(n), B S(n i)\} .
$$

4. Return $B S(n)=T B S(n)$.

The idea of solving one-step lookahead problems with a terminal cost (or backed-up score) that summarizes future costs is of course central in the DP algorithm. Indeed, it can be seen that the minimax algorithm just described is nothing but the DP algorithm for minimax problems (see Section 1.6). Here, positions and moves can be identified with states and controls, respectively, there are only terminal costs (the scores of the terminal positions), and the backed-up score of a position is nothing but the optimal cost-to-go at the corresponding state.

The minimax algorithm is also known as the type A strategy. Shannon argued that with this strategy, one could not expect a computer to seriously challenge human players of even moderate strength. In a typical chess position there are around 30 to 35 legal moves. It follows that for an $n$-move lookahead there will be around $30^{n}$ to $35^{n}$ terminal positions to be scored. Thus the number of terminal positions grows exponentially with the size of lookahead, practically limiting $n$ to being of the order of 10 with present computers. Unfortunately, in some chess positions it is essential to look a substantially larger number of moves ahead. In particular, in dynamic positions involving many captures and countercaptures, the necessary size of lookahead can be very large.

These considerations led Shannon to consider another strategy, called type $B$, whereby the depth of the search tree is variable. He suggested that at each position the computer give all legal moves a preliminary examination and discard those that are "obviously bad." A scoring function together with some heuristic strategy can be used for this purpose. Similarly, he suggested that some positions that are dynamic, such as those involving many captures or checkmate threats, be explored further beyond the nominal depth of the search.

Chess-playing computer programs typically use a combination of Shannon's type $A$ and $B$ strategies. These programs use scoring functions, the forms of which have evolved by trial-and-error, but they also use sophisticated heuristics to evaluate dynamic terminal positions in detail. In particular, an effective algorithm, known as swapofJ, is used to quickly analyze long sequences of captures and countercaptures, thereby making it possible to score realistically complex, dynamic positions (see Levy [Lev84] for a description). One may view such heuristics as either defining a sophisticated scoring function, or as implementing a type $B$ strategy.

The efficiency of the minimax algorithm can be substantially improved by using the alpha-beta pruning procedure (denoted $\alpha-\beta$ for short), which can be used to forego some calculations involving positions that cannot affect the selection of the best move. To understand the $\alpha-\beta$ procedure, consider a chess player pondering the next move at position $P$. Suppose that the player has already exhaustively analyzed one relatively good move $\mathrm{M}_{1}$ with corresponding score $B S(M 1 P)$ and proceeds to examine the next move M2. Suppose that as the opponent's replies are examined, a particularly strong response is found, which assures that the score of $M 2$ will be
worse than that of $M_{1}$. Such a response, called a refutation of move $M_{2}$, makes further consideration of move $M 2$ unnecessary (i.e., the portion of the search tree that descends from move $M_{2}$ can be discarded). An example is shown in Fig. 6.3.8.


Figure 6.3.8 The $\alpha-\beta$ procedure. White has evaluated move $M_{1}$ to have backedup score $(+1)$, and starts evaluating move $M 2$. The first reply of Black is a refutation of $M 2$, since it leads to a temporary score of -2 , less than the backedup score of MI. Since the backed-up score of M2 will be -2 or less, M2 will be inferior to $l V I l$. Therefore, it is not necessary to evaluate move $M_{2}$ further.

The $\alpha-\beta$ procedure can be generalized to trees of arbitrary or irregular depth and can be incorporated very simply into the minimax algorithm. Generally, if in the process of updating the backed-up score of a given position (step 3b) this score crosses a certain bound, then no further calculation is needed regarding that position. The cutoff bounds are adjusted dynamically as follows:

1. The cutoff bound in position $n$, where Black has to move, is denoted $\alpha$ and equals the highest current score of all ancestor positions of $n$ where White has to move. The exploration of position $n$ can be terminated as soon as its temporary backed-up score equals or falls below $\alpha$.
2. The cutoff bound in position $n$, where White has to move, is denoted $\beta$ and equals the lowest current value of all ancestor positions of $n$ where Black has the move. The exploration of position $n$ can be terminated as soon as its temporary backed-up score rises above (3.
The process is illustrated in Fig. 6.3.9. It can be shown that the backed-up score and optimal move at the starting position are unafJected by the incorporation of the $\alpha-\beta$ procedure in the minimax algorithm. We leave the verification of this fact to the reader (Exercise 6.8). It can also be seen that the $\alpha-\beta$ procedure will be more efJective if the best moves in each position are explored first. This tends to keep the $\alpha$ bounds high and the $\beta$ bounds low, thus saving a maximum amount of calculation. Current
chess programs use sophisticated techniques for ordering moves so as to maximize the effectiveness of the $\alpha-\beta$ procedure. We discuss briefly two of these techniques: iterative deepening and the killer heuristic.


Figure 6.3.9 The $\alpha-\beta$ procedure applied to the tree of Fig. 6.3.5. For example, the $/ 3$-cutoff in position Pl is due to the fact that its temporary score ( +20 ) exceeds its current $/ 3$-bound $(+16)$. The $\alpha$-cutoffs in positions $P 2, P a$, and $P 4$ are due to the fact that the corresponding temporary scores, $+8,+11$, and +11 , have fallen below the current $\alpha$-bound, which is +16 , the current temporary score in position $P$.

Iterative deepening, in its pure form, consists of first conducting a search based on lookahead of one move; then carrying out (from scratch) a search based on lookahead of two moves; then carrying out a search based on lookahead of three moves and so on. This process is 'continued either up to a fixed level of lookahead or until some limit on computation time is exceeded. At each iteration associated with a certain level of lookahead, one obtains a best move at the starting position, which is examined first in the subsequent iteration that requires one extra move of lookahead. This enhances the power of the $\alpha-\beta$ procedure, thereby more than making up for the extra computation involved in doing a short lookahead search before doing a longer one. (Actually, given that the number of terminal positions increases on the average by a factor of the order of 30 with each additional level of lookahead, the extra computation is relatively small.) An additional benefit of this method is that a best move is maintained throughout the search and can be produced at any time as needed. This comes in handy in commercial programs that incorporate a feature whereby the computer is forced to move either upon exhausting a given time allocation or upon command by a human opponent. An improvement of the method is to
obtain a thoroughly sorted list of moves at the starting position via a one-move lookahead, and then use the improved ordering in subsequent iterations to enhance the performance of the $\alpha-\beta$ procedure.

The killer heuristic is similar to iterative deepening in that it aims at examining first the most powerful moves at each position, thereby enhancing the pruning power of the $\alpha-\beta$ procedure. To understand the idea, suppose that in some position, White selects the first move $M_{1}$ from a candidate list $\left\{M_{1}, M_{2}, M_{3}, \ldots\right\}$, and upon examining Black's responses to $M^{1}$ finds that a particular move, which we will refer to as the killer move, is by far Black's best. Then it is often true that the killer move is also Black's best response to the second and subsequent moves $M_{2}, M_{3}, \ldots$ in White's list. It is therefore a good idea from the point of view of $\alpha-\beta$ pruning to consider the killer move first as a potential response to the remaining moves . $M 2, M_{3}, \ldots$ Of course, this does not always work as hoped, in which case it is advisable to change the killer move depending on subsequent results of the computation. In fact, some programs maintain lists of more than one killer move at each level of lookahead.

The $\alpha-\beta$ procedure is safe in the sense that searching a game tree with it and without it will produce the same result. Some computer chess programs use more drastic tree-pruning procedures, which usually require less computation for a given level of lookahead, but may miss on occasion the strongest move. There is some debate at present regarding the merits of such procedures. The books by Levy [Lev84] and Newborn [New75] consider this subject, and provide a broader discussion of the limitations of computer chess programs. A fascinating account of the development of a checkers computer program that implements many of the ideas discussed here is given by Schaeffer [Sch97].

### 6.4 ROLLOUT ALGORITHMS

We now discuss a specific type of cost-to-go approximation within the context of a limited lookahead scheme. Recall that in the one-step lookahead method, at stage $k$ and state $X k$ we use the control $\bar{\mu}_{k}(X k)$ that attains the minimum in the expression

$$
\min _{u_{k} \in U_{k}\left(x_{k}\right)} E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\},
$$

where $J_{k}+_{l}$ is some approximation of the true cost-to-go function $\mathrm{Jk}+\mathrm{l}^{\prime}$ In the rollout algorithm, the approximating function $\mathrm{Jk}+1$ is the cost-to-go of some known heuristic/suboptimal policy $\pi=\left\{\mu_{0}, . . \cdot, \mu_{N-1}\right\}$, called base policy (see also Example 6.3.1). The policy thus obtained is called the rollout policy based on $\pi$. Thus the rollout policy is a one-step lookahead
policy, with the optimal cost-to-go appTOximated by the cost-to-go of the base policy.

The process of starting from some suboptimal policy and generating another policy using the one-step lookahead process described above is also called policy improvement. This process will be discussed in Section 7.2 and in Vol. II in the context of the policy iteration method, which is a primary method for solving infinite horizon problems.

Note that it is also possible to define rollout policies that make use of multistep (say, l-step) lookahead. Here we assign to every state $x$ that can be reached in $I$ steps, the exact cost-to-go of the base policy, as computed by Monte Carlo simulation of several trajectories that start at $x$. Clearly, such multistep lookahead involves much more on-line computation, but it may yield better performance than its single-step counterpart. In what follows, we concentrate on rollout policies with single-step lookahead.

The viability of a rollout policy depends on how much time is available to choose the control following the transition to state $x$ and on how expensive the Monte Carlo evaluation of the expected value

$$
E\left\{g k(X k, U k, W k)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}
$$

is. In particular, it must be possible to perform the Monte Carlo simulations and calculate the rollout control within the real-time constraints of the problem. If the problem is deterministic, a single simulation trajectory suffices, and the calculations are greatly simplified, but in general, the computational overhead can be substantial.

It is possible, however, to speed up the calculation of the rollout policy if we are willing to accept some potential performance degradation. For example, we may use an approximation $J_{k}+_{1}$ of $J_{k}+_{1}$ to identify a few promising controls through a minimization of the form

$$
\left.\operatorname{mUnin}_{U k}{\underset{k}{X k}}\right) E\left\{g k(X k, U k, W k)+J_{k+1(f k(x k, u k, w k))\},}\right.
$$

and then restrict attention to these controls, using fairly accurate Monte Carlo simulation. In particular, the required values of $J_{k}+_{1}$ may be obtained by performing approximately the Monte Carlo simulation, using a limited number of representative trajectories. Adaptive variants of this approach are also possible, whereby we use some heuristics to adjust the accuracy of the Monte Carlo simulation depending on the results of the computation.

Generally, it is important to use as base policy one whose expected cost-to-go is conveniently calculated. The following is an example.

## Example 6.4.1 (The Quiz Problem)

Consider the quiz problem of Example 4.5.1, where a person is given a list of $N$ questions and can answer these questions in any order he/she chooses.

Question i will be answered correctly with probability pi, and the person will then receive a reward Vi. At the first incorrect answer, the quiz terminates and the person is allowed to keep his or her previous rewards. The problem is to choose the ordering of questions so as to maximize the total expected reward.

We saw that the optimal sequence can be obtained using an interchange argument: questions should be answered in decreasing order of the "index of preference" $p_{i} v_{i} /(1-P i)^{\prime}$ We refer to this as the index policy. Unfortunately, with only minor changes in the structure of the problem, the index policy need not be optimal. Examples of difficult variations of the problem may involve one or more of the following characteristics:
(a) A limit on the maximum number of questions that can be answered, which is smaller than the number of questions $N$. To see that the index policy is not optimal anymore, consider the case where there are two questions, only one of which may be answered. Then it is optimal to answer the question that offers the maximum expected reward PiVi.
(b) A time window for each question, which constrains the set of time slots when each question may be answered. Time windows may also be combined with the option to refuse answering a question at a given period, when either no question is available during the period, or answering any one of the available questions involves excessive risk.
(c) Precedence constraints, whereby the set of questions that can be answered in a given time slot depends on the immediately preceding question, and possibly on some earlier answered questions.
(d) Sequence-dependent rewards, whereby the reward from answering correctly a given question depends on the immediately preceding question, and possibly on some questions answered earlier.
Nonetheless, even when the index policy is not optimal, it can conveniently be used as a base policy for the rollout algorithm. The reason is that at a given state, the index policy together with its expected reward can be easily calculated. In particular, each feasible question order $\ldots, i N$ ) has expected reward equal to

$$
p_{i_{1}}\left(v_{i_{1}}+p_{i_{2}}\left(v_{i_{2}}+p_{i_{3}}\left(\cdots+p_{i_{N}} v_{i_{N}}\right) \cdots\right)\right)
$$

Thus the rollout algorithm based on the index heuristic operates as follows: at a state where a given subset of questions have already been answered, we consider the set of questions $J$ that are eligible to be answered next. For each question $\boldsymbol{j}$ E $J$, we consider a sequence of questions that starts with $\boldsymbol{j}$ and continues with the remaining questions chosen according to the index rule. We compute the expected reward of the sequence, denoted $R(j)$, using the above formula. Then among the questions $\boldsymbol{j} \mathrm{E} J$, we choose to answer next the one with maximal $R(j)$. A computational study of rollout algorithms for the quiz problem and some variations, using several methods for approximating the cost-to-go of the base policy, is given by Bertsekas and Castanon [BeC99].

## Cost morovement with a Rollout Algorithm

Rollout policies have a nice property: in their pure form, they always result in improved performance over the corresponding base policy. This is essentially a consequence of Prop. 6.3.1 (see Example 6.3.1), but for the purpose of convenient reference, we adapt the proof of that proposition to the rollout context. Let $] k(X k)$ and $H k(X k)$ be the costs-to-go of the rollout and the base policies, respectively, starting from a state $X k$ at time $k$. We will show that $] k(X k) \leq H k(X k)$ for all $X k$ and $k$, so that the rollout policy $\bar{\pi}$ is an improved policy over the base policy $\pi$. We have] $N(X N)=$ $H N(X N)=g N(X N)$ for all $X N$. Assuming that $] k+l(X k+l) \leq H_{k+1}\left(x_{k+1}\right)$ for all $x k+1$, we have

$$
\begin{aligned}
J k(X k) & E\left\{g k\left(X k^{\prime} \bar{\mu}_{k}(X k), W k\right)+\bar{J} k+l\left(f k\left(X k^{\prime} \bar{\mu}_{k}\left(x_{k}\right), W k\right)\right)\right\} \\
\leq & \left.E\left\{g k\left(X k^{\prime} \bar{\mu}_{k}\left(x_{k}\right), W k\right)+H k+l U k\left(X k^{\prime} \bar{\mu}_{k}\left(x_{k}\right), W k\right)\right)\right\} \\
\leq & E\left\{g k\left(X k^{\prime} \mu_{k}\left(x_{k}\right), W k\right)+H k+l\left(f k\left(X k^{\prime} \mu_{k}\left(x_{k}\right), W k\right)\right)\right\} \\
= & H k(X k),
\end{aligned}
$$

for all $X k$. The first inequality above follows from the induction hypothesis, the second inequality follows from the definition of the rollout policy, and the first and second equalities follow from the DP algorithm that defines the costs-to-go of the rollout and the base policies, respectively. This completes the induction proof that $\bar{\pi}$ is an improved policy over $\pi$.

Empirically, it has been observed that the rollout policy typically produces considerable (and often dramatic) cost improvement over the base policy. However, there is no solid theoretical support for this observation. The following example provides some insight into the nature of cost improvement.

## Example 6.4.2 (The Breakthrough Problem)

Consider a binary tree with $N$ stages as shown in Fig. 6.4.1. Stage $k$ of the tree has $2^{k}$ nodes, with the node of stage 0 called root and the nodes of stage $N$ called leaves. There are two types of tree arcs: free and blocked. A free (or blocked) arc can (cannot, respectively) be traversed in the direction from the root to the leaves. The objective is to break through the graph with a sequence of free arcs (a free path) starting from the root, and ending at one of the leaves.

One may use DP to discover a free path (if one exists) by starting from the last stage and by proceeding backwards to the root node. The kth step of the algorithm determines for each node of stage $N-k$ whether there is a free path from that node to some leaf node, by using the results of the preceding step. The amount of calculation at the kth step is $O\left(2^{N}-{ }^{k}\right)$. Adding the calculations for the $N$ stages, we see that the total amount of calculation is $O\left(N 2^{N}\right)$, so it increases exponentially with the number of stages. For this reason it is interesting to consider heuristics that require calculation that is


Figure 6.4.1 Binary tree for the breakthrough problem. Each arc is either free or is blocked (crossed out in the figure). The problem is to find a path from the root to one of the leaves, which is free (such as the one shown with thick lines).
linear or polynomial in $N$, but may sometimes fail to determine a free path, even when a free path exists.

Thus, one may suboptimally use a greedy algorithm, which starts at the root node, selects a free outgoing arc (if one is available), and tries to construct a free path by adding successively nodes to the path. Generally, at the current node, if one of the outgoing arcs is free and the other is blocked, the greedy algorithm selects the free arc. Otherwise, it selects one of the two outgoing arcs according to some fixed rule that depends only on the current node (and not on the status of other arcs). Clearly, the greedy algorithm may fail to find a free path even if such a path exists, as can be seen from Fig. 6.4.1. On the other hand the amount of computation associated with the greedy algorithm is $O(N)$, which is much faster than the $O\left(N 2^{N}\right.$ ) computation of the DP algorithm. Thus we may view the greedy algorithm as a fast heuristic, which is suboptimal in the sense that there are problem instances where it fails while the DP algorithm succeeds.

Let us also consider the rollout algorithm that uses the greedy algorithm as the base heuristic. This algorithm starts at the root and tries to construct a free path by exploring alternative paths constructed by the greedy algorithm. At the current node, it proceeds according to the following two Cases:
(a) If at least one of the two outgoing arcs of the current node is blocked, the rollout algorithm adds to the current path the arc that the greedy algorithm would select at the current node.
(b) If both outgoing arcs of the current node are free, the rollout algorithm considers the two end nodes of these arcs, and from each of them it runs the greedy algorithm. If the greedy algorithm succeeds in finding a free path that starts from at least one of these nodes, the rollout algorithm stops with a free path having been found; otherwise, the rollout algorithm moves to the node that the greedy algorithm would select at the current node.

Thus, when both outgoing arcs are free, the rollout algorithm explores further the suitability of these arcs, as in case (b) above. Because of this additional discriminatory capability, the rollout algorithm always does at least as well as the greedy (it always finds a free path when the greedy algorithm does, and it also finds a free path in some cases where the greedy algorithm does not). This is consistent with our earlier discussion of the generic cost improvement property of the rollout algorithm over the base heuristic. On the other hand, the rollout algorithm applies the greedy heuristic as many as $2 N$ times, so that it requires $O\left(N^{2}\right)$ amount of computation - this is intermediate between the $O(N)$ computation of the greedy and the $O\left(N 2^{N}\right)$ computation of the DP algorithm.

Let us now calculate the probabilities that the algorithms will find a free path given a randomly chosen breakthrough problem. In particular, we generate the graph of the problem randomly, by selecting each of its arcs to be free with probability $p$, independently of the other arcs. We then calculate the corresponding probabilities of success for the greedy and the rollout algorithms.

The probability $G k$ that the greedy algorithm will find a free path in a graph of $k$ stages is the probability of a "success" in each of the $k$ stages, where a success is counted whenever at least one of the two arcs involved is free, an event of probability $1-(1-p) 2$ or $p(2-p)$. Thus we have, using the independence of the blocked/unblocked status of the arcs,

$$
G_{k}=(p(2-p))^{k}
$$

The probability $R k$ that the rollout algorithm will find a free path in a graph of $k$ stages can be calculated by means of a recursion, as we now show. At a given node no with $k$ stages to go, consider the path ( $n o, n 1, \ldots, n k$ ) generated by the greedy algorithm, and let ( $n o, \mathrm{nl}$ ) and ( $n o, n_{1}^{\prime}$ ) denote the incident arcs of node no. Let PI denote the probability that exactly one of the $\operatorname{arcs}(n o, \mathrm{nl})$ and ( $n o, n_{1}^{\prime}$ ) is free, so

$$
P I=2 p(1-p),
$$

and let $P 2$ denote the probability that both $\operatorname{arcs}(n o, \mathrm{nl})$ and (no, $\left.n_{1}^{\prime}\right)$ are free, so

$$
P_{2}=p^{2} .
$$

To calculate the probability $R k$ of the event that the rollout algorithm succeeds in finding a free path, we partition this event into the following four mutually exclusive events, and we calculate their probabilities:
(1) Event El. Exactly one of the $\operatorname{arcs}(n o, \mathrm{nl})$ and (no, $n_{1}^{\prime}$ ) is free [by necessity (no, nl) since it is chosen by the greedy algorithm] and the rollout algorithm finds a free path starting from nl. The probability of this event is $p e E l)=P l R k-l$.
(2) Event E2. Both arcs (no, nl) and (no, $n_{1}^{\prime}$ ) are free and the greedy algorithm finds a free path starting from nl . The probability of this event is $P(E 2)=P 2 G k-l$.
(3) Event $\mathrm{E}_{\mathrm{s}}$. Both the arcs (no, nl ) and (no, $n_{1}^{\prime}$ ) are free, and the greedy algorithm does not find a free path starting from nl but finds a free path from $n_{1}^{\prime}$. The probability of this event is peEs $)=P 2(1-G k-l) G k-l$.

Event $E_{4}$. Both the $\operatorname{arcs}(n o, n l)$ and ( $n o, n_{1}^{\prime}$ ) are free, the greedy algorithm does not find a free path starting from either nl or $n_{1}^{\prime}$, but the rollout algorithm finds a free path from nl. The probability of this event is $P_{2}\left(l-G_{k-1}\right)^{2} H k-l$, where $H k-l$ is the conditional probability that the rollout finds a free path from nl given that the greedy does not find a free path from nl. We have

$$
R_{k-1}=G_{k-1}+\left(1-G_{k-1}\right) H_{k-1}
$$

so ( $1-G k-l) H k-l=R k-l-G k-l$, and the probability of the event $E 4$ is

$$
P\left(E_{4}\right)=P_{2}\left(1-G_{k-1}\right)^{2} H_{k-1}=P_{2}\left(1-G_{k-1}\right)\left(R_{k-1}-G_{k-1}\right)
$$

Thus, by adding the probabilities of the above mutually exclusive and collectively exhaustive events, we have

$$
\begin{aligned}
R k & =p e E l)+P(E 2)+p e E s)+P(E 4) \\
& =P l R k-l+P 2(G k-l+(1-G k-l) G k-l+(1-G k-l)(R k-l-G k-l)) \\
& =(P I+P 2(l-G k-l)) R k-l+P 2 G k-l .
\end{aligned}
$$

From this, by substituting the expressions $P I=2 p(l-p)$ and $P 2=p 2$, we obtain

$$
\begin{aligned}
R k & =\left(2 p(1-p)+p 2\left(1 \_G k-l\right)\right) R k-l+p 2 G k_{-} l \\
& =p(2-p) R_{k}-l+p^{2} G k_{-} l(1-R k-l),
\end{aligned}
$$

with the initial condition $R o=1$. Since $\lim _{k \rightarrow \infty} G k=$ and $p(2-p)<1$, it follows from the above equation that $\lim _{k \rightarrow \infty} R k=0$. Furthermore, by dividing with $G_{k}=p(2-p) G k-l$, we have

$$
\frac{R_{k}}{G_{k}}=\frac{R_{k-1}}{G_{k-1}}+\frac{p}{2-p}\left(1-R_{k-1}\right)
$$

so since $\lim _{k \rightarrow \infty} R k=0$, we obtain for large $N$

$$
\frac{R_{N}}{G_{N}}=O\left(N \frac{p}{2-p}\right)
$$

Thus, asymptotically, the rollout algorithm requires $O(N)$ times more computation, but has an $O(N)$ times larger probability of finding a free path than the greedy algorithm. This type of tradeoff appears to be quailitatively typical: the rollout algorithm achieves a substantial performance improvement over the base heuristic at the expense of extra computation that is equal to the computation time of the base heuristic times a factor that is a low order polynomial of the problem size.

## whinnianumar Issues in Rollout Algorithms

We now consider in more detail implementation issues and specific properties of rollout algorithms in a variety of settings. To compute the rollout control $\bar{\mu}_{k}\left(x_{k}\right)$, we need for all $U k E \operatorname{Ek}(X k)$ the value of

$$
Q k(X k, U k)=E\left\{g k(X k, U k, W k)+H_{k+l}(f k(X k, U k, W k))\right\},
$$

known as the $Q$-factor of $(X k, U k)$ at time $k$. Alternatively, for the computation of $\bar{\mu}_{k}\left(x_{k}\right)$ we need the value of the cost-to-go

$$
H k+l(f k(X k, U k, W k))
$$

of the base policy at all possible next states $f_{k}\left(x_{k}, U k, W k\right)$, from which we can compute the required Q -factors.

We will focus on the case where a closed form expression of the $Q$ factor is not available. We assume instead that we can simulate the system under the base policy $\pi$, and in particular, that we can generate sample system trajectories and corresponding costs consistently with the probabilistic data of the problem. We will consider several cases and possibilities, we will point out their advantages and drawbacks, and we will discuss the contexts within which they are most appropriate. These cases are:
(1) The deterministic problem case, where $W k$ takes a single known value at each stage. We provide an extensive discussion of this case, focusing not only on traditional deterministic optimal control problems, but also on quite general combinatorial optimization problems, for which the rollout approach has proved convenient and effective.
(2) The stochastic problem case with Q-factors e'valuated by Monte-Carlo simulation. Here, once we are at state $X k$, the Q -factors $Q k(X k, U k)$ are evaluated on-line by Monte-Carlo simulation, for all $U k \mathrm{E} U k(X k)$.
(3) The stochastic problem case with Q-factors approximated in some way. One possibility is to use a certainty equivalence approximation, where the problem is genuinely stochastic, but the values $H k(X k)$ are approximated by the cost-to-go of $\pi$ that would be incurred if the system were replaced by a suitable deterministic system from state $X k$ and time $k$ onward (assumed certainty equivalence). There are also other possibilities based on the use of an approximation architecture and some form of least squares.
We examine each of these three cases in the next three subsections.

### 6.4.1 Discrete Deterministic Problems

Let us assume that the problem is deterministic, i.e., that $W k$ can take only one value at each stage $k$. Then, starting from state $X k$ at stage $k$, the
base policy $\pi$ produces deterministic sequences of states and controls $\left\{U k, \ldots, u_{N-1}\right\}$ such that

$$
x_{i+1}=f\left(x_{i}, u_{i}\right), \quad \mathrm{i}=k, \ldots, N-1
$$

and a cost

$$
g_{k}\left(x_{k}, u_{k}\right)+\cdots+g_{N-1}\left(x_{N-1}, u_{N-1}\right)+g_{N}\left(x_{N}\right) .
$$

Thus the Q-factor

$$
Q_{k}\left(x_{k}, u_{k}\right)=g_{k}\left(x_{k}, u_{k}\right)+H_{k+1}\left(f_{k}\left(x_{k}, u_{k}\right)\right)
$$

can be obtained by running $\pi$ starting from state $f k(X k, U k)$ and time $k+1$, and recording the corresponding cost $H k+l\left(f_{k}\left(x_{k}, U k\right)\right)$. The rollout control $\bar{\mu}_{k}\left(x_{k}\right)$ is obtained by calculating in this manner the Q -factors $Q k(X k, U k)$ for all $U k \mathrm{E} U k(X k)$, and setting

$$
\bar{\mu}_{k}\left(x_{k}\right)=\arg \min _{u_{k} \in U_{k}\left(x_{k}\right)} Q_{k}\left(x_{k}, u_{k}\right)
$$

Aside for being convenient for the deterministic special case of the basic problem of Chapter 1, this rollout method can be adapted for general discrete or combinatorial optimization problems that do not necessarily have the strong sequential character of the basic problem. For such problems the rollout approach provides a convenient and broadly applicable suboptimal solution method that goes beyond and indeed enhances the common types of heuristics, such as greedy algorithms, local search, genetic algorithms, tabu search, and others.

To illustrate the ideas involved, let us consider the problem

$$
\begin{align*}
& \operatorname{minimize} G(u) \\
& \text { subject to } U \mathrm{E} U \tag{6.31}
\end{align*}
$$

where $U$ is a finite set of feasible solutions and $G(u)$ is a cost function. We assume that each solution $U$ has $N$ components; that is, it has the form $U=$ (UI $U 2, \ldots, U N$ ), where $N$ is a positive integer. Under this assumption, we can view the problem as a sequential decision problem, where the components UI, $\cdots, U N$ are selected one-at-a-time. An n-tuple (UI, ... , un) consisting of the first $n$ components of a solution is called an $n$-solution. We associate $n$-solutions with the nth stage of a DP problem. In particular, for $n=1, \ldots, N$, the states of the nth stage are of the form (Ul,." ,Un). The initial state is a dummy (artificial) state. From this state we may move to any state $\left(u_{1}\right)$, with UI belonging to the set
$U \iota=\left\{\tilde{u}_{1} \mid\right.$ there exists a solution of the form $\left.\left(\tilde{u}_{1}, \tilde{u}_{2}, \ldots, \tilde{u}_{N}\right) \mathrm{E} U\right\}$.

Thus $U_{l}$ is the set of choices of Ul that are consistent with feasibility.
More generally, from a state of the form (Ul, ... Un-I), we may move to any state of the form ( $\mathrm{Ul}, \ldots, U n-I, U n$ ), with Un belonging to the set

$$
\begin{array}{r}
U n(U I, \ldots, U n-I)=\left\{\tilde{u}_{n} I\right. \text { there exists a solution of the form } \\
 \tag{6.32}\\
\left.\left(\mathrm{Ul}, \ldots, u_{n-1}, \tilde{u}_{n}, \ldots, \tilde{u}_{N}\right) \mathrm{E} U\right\} .
\end{array}
$$

The choices available at state $(\mathrm{Ul}, \ldots, \mathrm{Un}-1)$ are $U n \mathrm{E} U n(U I, \ldots, U n-I)^{\prime}$ These are the choices of $u_{n}$ that are consistent with the preceding choices $\mathrm{U}, \ldots, U n-I$, and are also consistent with feasibility. The terminal states correspond to the N -solutions ( $\mathrm{Ul}, \ldots, U N$ ), and the only nonzero cost is the terminal cost $G\left(u_{1}, \ldots, U N\right)$.

Let $\mathrm{J}^{*}(\mathrm{Ul}, \ldots, U n)$ denote the optimal cost starting from the n -solution (Ul, $\ldots$, Un), i.e., the optimal cost of the problem over solutions whose first $n$ components are constrained to be equal to $U i, \mathrm{i}=1, \ldots, n$, respectively. If we knew the optimal cost-to-go function $\mathrm{J}^{*}(\mathrm{Ul}, \ldots, U n)$, we could construct an optimal solution by a sequence of N single component minimizations. In particular, an optimal solution $\left(u_{1}^{*}, \ldots, u_{N}^{*}\right)$ could be obtained through the algorithm

$$
u_{i}^{*}=\arg \min _{u_{i} \in U_{i}\left(u_{1}^{*}, \ldots, u_{i-1}^{*}\right)} J *\left(u_{1}^{*}, \ldots, u_{i-1}^{*}, u_{i}\right), \quad \mathrm{i}=1, \ldots, N .
$$

Unfortunately, this is seldom viable, because of the prohibitive computation required to obtain $\mathrm{J}^{*}(\mathrm{Ul}, \ldots, U n)$.

Assume now that we have a heuristic, which starting from an $n$ solution ( $U I, "$ "Un), produces an $N$-solution ( $u_{1}, \ldots, U n \prime U n+l \ldots, U N$ ) whose cost is denoted by $H(\mathrm{Ul}, \ldots$, Un). Such a heuristic may be viewed as a base policy for the problem, in the sense that given the current state (Ul, $\ldots, U n$ ) it generates the next decision $U n+l$ as the first component of the remaining portion $\left(U_{n}+1, \ldots, U N\right)$ of the solution. Let us consider the corresponding rollout algorithm. It can be seen that this algorithm selects as the first solution component

$$
\bar{u}_{1}=\arg \min _{\mathrm{U} 1 \mathrm{EU} 1} H(U I),
$$

and operates sequentially, for $n=1, \quad, \mathrm{~N}-1$, as follows:
Given a partial solution $\left(\bar{u}_{1}, \quad, \bar{u}_{n}\right)$, it runs the heuristic starting from the partial solutions $\left(\bar{u}_{1}, \quad, \bar{u}_{n}, U n+l\right)$ corresponding to all the possible next solution components $U n+l \mathrm{E} U n+l\left(\bar{u}_{1}, \ldots, \bar{u}_{n}\right)$, and selects as next component

$$
\bar{u}_{n+1}=\arg \min _{U n+l \in U_{n+1}\left(\bar{u}_{1}, \ldots, \bar{u}_{n}\right)} H\left(\bar{u}_{1}, \ldots \bar{u}_{n}, u_{n+1}\right) .
$$

In order to analyze most economically the preceding algorithm and its variants, we will embed it within a more general and flexible framework for
discrete optimization. To this end, we introduce a graph search problem, which contains as special cases broad classes of discrete/integer optimization problems, and will serve as the context of our methodology. We will then describe and analyze a basic form of a one-step lookahead algorithm, we will discuss some of its variations, we will illustrate it by means of some examples, and we will discuss its connection with the DP context.

As we explain later (see the end of Section 6.4.1), the algorithm to be introduced is not quite a rollout algorithm in the sense discussed so far, because, strictly speaking, it does not use the cost of a heuristic policy as a one-step lookahead cost approximation, except under a special assumption (the sequential consistency assumption, to be described later). The basic idea of the algorithm is, however, very close to a rollout: it is a one-step lookahead policy with cost approximation derived from a heuristic. Thus, with a stretch of terminology, we will call this algorithm "rollout" as well.

## The Basic Rollout Algorithm for Discrete

Let us introduce a graph search problem that will serve as a general model for discrete optimization. We are given a graph with set of nodes $N$, set of $\operatorname{arcs} A$, and a special node $s$, which we call the origin. The arcs are directed in the sense that arc ( $\mathbf{i}, \mathrm{j}$ ) is distinct from $\operatorname{arc}(j, i)$. We are also given a subset $\underline{N}$ of nodes, called destinations, and a cost $g(\mathrm{i})$ for each destination i . The destination nodes are terminal in the sense that they have no outgoing arcs. For simplicity, we assume that the node and arc sets, $N$ and $A$ contain a finite number of elements. However, the following analysis and discussion applies, with minor modifications in language, to the case of a countably infinite number of nodes and a finite set of outgoing arcs from each node. We want to find a path that starts at the origin $s$, ends at one of the destination nodes i E $\underline{N}$, and is such that the cost $g(i)$ is minimized.

In the context of the discrete optimization problem (6.31), nodes i correspond to n-tuples (Ul, , Un) consisting of the first $n$ components of a solution, where $n=1, \quad \mathrm{~N}$. Arcs lead from nodes of the form (U1"", Un-1) to nodes of the form (Ul, $\ldots, u_{n-1}, u_{n}$ ), and there is an arc for each Un of the form (6.32). An interesting property of this special case is that its associated graph is acyclic.

In our terminology, a path is a sequence of arcs
(il,i2), (i2,i3),"" (im-1,im),
all of which are oriented in the forward direction. The nodes il and im are called the start node and the end node of the path, respectively. For convenience, and without loss of generality, t we will assume that given an

[^1]ordered pair of nodes ( $\mathrm{i}, j$ ), there is at most one arc with start node i and end node $j$, which (if it exists) will be denoted by (i, $j$ ). In this way, a path consisting of arcs (iI, i2), (i2, i3), , (im-1, $i_{m}$ ) is unambiguously specified as the sequence of nodes ( $\mathrm{iI}, \mathrm{i} 2, \quad, i_{m}$ ).

Let us assume that we have a heuristic path construction algorithm, denoted $\mathcal{H}$, which given a non-destination node i $\notin \underline{N}$, constructs a path
iI, $\left.\ldots, i_{m}, \bar{i}\right)$ starting at i and ending at one of the destination nodes
Implicit in this assumption is that for every non-destination node, there exists at least one path starting at that node and ending at some destination node. We refer to the algorithm $\mathcal{H}$ as the base heuristic, since we will use this algorithm as the basic building block for constructing the rollout algorithm to be introduced shortly.

The end node $\bar{i}$ of the path constructed by the base heuristic $\mathcal{H}$ is completely specified by the start node i. We call $\overline{\boldsymbol{i}}$ the projection of i under $\mathcal{H}$, and we denote it by $p(i)$. We denote the corresponding cost by $H(i)$,

$$
H(i)=g(p(i))
$$

The projection of a destination node is the node itself by convention, so that for all $\mathrm{i} \mathrm{E} \underline{\underline{\mathcal{N}}}$ we have $\mathrm{i}=p(i)$ and $H(i)=g(i)$. Note that while the base heuristic $\mathcal{H}$ will generally yield a suboptimal solution, the path that it constructs may involve a fairly sophisticated suboptimization. For example, $\mathcal{H}$ may construct several paths ending at destination nodes according to some heuristics, and then select the path that yields minimal cost.

One possibility for suboptimal solution of the problem is to start at the origin $s$ and use the base heuristic $\mathcal{H}$ to obtain the projection $p(s)$. We instead propose to use $\mathcal{H}$ to construct a path to a destination node sequentially. At the typical step of the sequence, we consider all downstream neighbors $j$ of a node $i$, we run the base heuristic $\mathcal{H}$ starting from each of these neighbors, and obtain the corresponding projections and costs. We then move to the neighbor that gives the best projection. This sequential version of $\mathcal{H}$ is called the rollout algorithm based on $\mathcal{H}$, and is denoted by $\mathcal{R H}$.

To formally describe the rollout algorithm, let $N$ (i) denote the set of downstream neighbors of node $i$,

$$
N(i) \quad\{j \mid(i, j) \text { is an } \operatorname{arc}\} .
$$

Note that $N(i)$ is nonempty for every non-destination node i , since by assumption there exists a path starting at $i$ and ending at its projection $p(i)$. The rollout algorithm $\mathcal{R H}$ starts with the origin node $s$. At the typical step, given a node sequence $\left(s, i_{1}, \ldots, i_{m}\right)$, where $\mathrm{i}_{\mathrm{m}}$ is not a destination, $\mathcal{R H}$ adds to the sequence a node $i m+1$ such that

$$
\begin{equation*}
=\arg \min _{j E N\left(i_{r n}\right)} H(j) . \tag{6.33}
\end{equation*}
$$



Figure 6.4.2 Illustration of the rollout algorithm. After m steps of the algorithm, we have the path $\left(8, \mathrm{il}, \ldots, i_{r n}\right)$. To extend this path at the next step, we generate the set $\mathrm{N}\left(i_{m}\right)$ of neighbors of the terminal node $i_{m}$, and select from this set the neighbor that has the best projection, i.e.

$$
\operatorname{irn}+1=\arg \min _{j E N\left(i_{r n}\right)} H(j)=\arg \min _{j E N\left(i_{r n}\right)} g(p(j)) .
$$

If im +1 is a destination node, $\mathcal{R H}$ terminates. Otherwise, the process is repeated with the sequence ( $s, \mathrm{il}, \ldots, \quad i_{m+1}$ ) replacing ( $s, \mathrm{il}, \ldots, i_{m}$ ); see Fig. 6.4.2.

Note that once $\mathcal{R H}$ has terminated with a path ( $s, i 1, \ldots, i_{m}$ ), we will have obtained the projection $p(i k)$ of each of the nodes $i k^{\prime} k=1, \ldots, \mathrm{~m}$. The best of these projections yields a cost

$$
\min _{k=I, \ldots, m} H(i k)=\min _{k=I, \ldots, m} g(p(i k))^{\prime}
$$

and the projection that corresponds to the minimum above may be taken as the final (suboptimal) solution produced by the rollout algorithm. We may also compare the above minimal cost with the cost $9(p(s))$ of the projection $p(s)$ of the origin, and use $p(s)$ as the final solution if it produces a smaller cost. This will ensure that the rollout algorithm will produce a solution that is no worse than the one produced by the base heuristic.

## Example 6.4.3 (Traveling Salesman Problem)

Let us consider the traveling salesman problem, whereby a salesman wants to find a minimum mileage/cost tour that visits each of iV given cities exactly once and returns to the city he started from. We associate a node with each city $\mathrm{i}=1, \ldots, \mathrm{iV}$, and we introduce an arc $(i, j)$ with traversal cost aij for each ordered pair of nodes $i$ and $j$. Note that we assume that the graph is complete; that is, there exists an arc for each ordered pair of nodes. There is no loss of generality in doing so because we can assign a very high cost aij to
an $\operatorname{arc}(i, j)$ that is precluded from participation in the solution. The problem is to find a cycle that goes through all the nodes exactly once and whose sum of arc costs is minimum.

There are many heuristic approaches for solving the traveling salesman problem. For illustration purposes, let us restrict attention to the simple nearest neighbor heuristic. Here, we start from a path consisting of just a single node $i_{1}$ and at each iteration, we enlarge the path with a node that does not close a cycle and minimizes the cost of the enlargement. In particular, after $k$ iterations, we have a path $\{\mathrm{iI}, \ldots, i k\}$ consisting of distinct nodes, and at the next iteration, we add an arc ( $i k, \mathrm{ik}+1$ ) that minimizes aiki over all arcs ( $i k$, i) with $\mathrm{i} \neq \mathrm{i} 1, \ldots, i k$. After $N \mathbf{- 1}$ iterations, all nodes are included in the path, which is then converted to a tour by adding the final arc (iN, i1).

We can formulate the traveling salesman problem as a graph search problem as follows: There is a chosen starting city, say il corresponding to the origin of the graph search problem. Each node of the graph search problem corresponds to a path (i1,i2, .., ik), where $\mathrm{i} 1, \mathrm{i} 2, \ldots, \mathrm{i} k$ are distinct cities. The neighbor nodes of the path (i1, i2, .., ik) are paths of the form (i1, i2, ,ik, which correspond to adding one more unvisited city $i k+1 \neq$ $\mathrm{i} 1, \mathrm{i} 2, \quad, i k$ at the end of the path. The destinations are the cycles of the form (i1, i2, ,iN), and the cost of a destination in the graph search problem is the cost of the corresponding cycle. Thus a path from the origin to a destination in the graph search problem corresponds to constructing a cycle in $N-1$ arc addition steps, and at the end incurring the cost of the cycle.

Let us now use as base heuristic the nearest neighbor method. The corresponding rollout algorithm operates as follows: After $k$ iterations, we have a path $\{\mathrm{i} 1, \ldots, i k\}$ consisting of distinct nodes. At the next iteration, we run the nearest neighbor heuristic starting from each of the paths of the form $\{\mathrm{i} 1, \ldots, i k, \mathrm{i}\}$ where $\mathrm{i} \neq \mathrm{i} 1, \ldots, i k$, and obtain a corresponding cycle. We then select as next node $i k+1$ of the path the node $i$ that corresponds to the best cycle thus obtained.

## Termination and Sequential Consistency

We say that the rollout algorithm $\mathcal{R} \mathcal{H}$ is terminating if it is guaranteed to terminate finitely starting from any node. Contrary to the base heuristic $\mathcal{H}$, which by definition, has the property that it yields a path terminating at a destination starting from any node, the rollout algorithm $\mathcal{R} \mathcal{H}$ need not have this property in the absence of additional conditions. The termination question can usually be resolved quite easily, and we will now discuss a few different methods by which this can be done.

One important case where $\mathcal{R} \mathcal{H}$ is terminating is when the graph is acyclic, since then the nodes of the path generated by $\mathcal{R} \mathcal{H}$ cannot be repeated within the path, and their number is bounded by the number of nodes in $N$. As a first step towards developing another case where $\mathcal{R} \mathcal{H}$ is terminating, we introduce the following definition, which will also set the stage for further analysis of the properties of $\mathcal{R} \mathcal{H}$.

Definition 6.4.1: We say that the base heuristic $\mathcal{H}$ is sequentially consistent if for every node i , it has the following property: If $\mathcal{H}$ generates the path (i, iI, $\ldots, i_{m}, \bar{i}$ ) when it starts at i , it generates the path $\left(i_{1}, \ldots, i_{m}, \bar{i}\right)$ when it starts at the node il.

Thus $\mathcal{H}$ is sequentially consistent if all the nodes of a path that it generates have the same projection. There are many examples of sequentially consistent algorithms that are used as heuristics in combinatorial optimization, including the following.

## Example 6.4.4 (Greedy Algorithms as Base Heuristics)

Suppose that we have a function $F$, which for each node i, provides a scalar estimate $F(\mathrm{i})$ of the optimal cost starting from i , that is, the minimal cost $g(\bar{i})$ that can be obtained with a path that starts at i and ends at one of the destination nodes $\overline{\boldsymbol{i}} \mathrm{E} \quad$ Then $F$ can be used to define a base heuristic, called the greedy algorithm with respect to $F$, as follows:

The greedy algorithm starts at a node i with the (degenerate) path that consists of just node i. At the typical step, given a path (i, iI, $\ldots, i_{r n}$ ), where $i_{m}$ is not a destination, the algorithm adds to the path a node $i_{m+1}$ such that

$$
\begin{equation*}
i_{m+1}=\arg \min _{j E N\left(i_{r n}\right)} F(j) \tag{6.34}
\end{equation*}
$$

In the case where $i_{m+1}$ is a destination, the algorithm terminates with the path ( $\mathrm{i}, \mathrm{i}_{1}, \ldots, i_{m}, \mathrm{i}_{\mathrm{rn}}+1$ ). Otherwise, the process is repeated with the path (i, il, $\ldots, i_{m}$, irn +1 ) replacing (i, $\mathrm{il}, \ldots, i_{m}$ ).

An example of a greedy algorithm is the nearest neighbor heuristic for the traveling salesman problem (cf. Example 6.4.3). Recall from that example that nodes of the graph search problem correspond to paths (sequences of distinct cities), and a transition to a neighbor node corresponds to adding one more unvisited city to the end of the current path. The function $F$ in the nearest neighbor heuristic specifies the cost of the addition of the new city.

It is also interesting to note that by viewing $F$ as a cost-to-go approximation, we may consider the greedy algorithm to be a special type of one-step lookahead policy. Furthermore, if $F(j)$ is chosen to be the cost obtained by some base heuristic starting from $j$, then the greedy algorithm becomes the corresponding rollout algorithm. Thus, it may be said that the rollout algorithm is a special case of a greedy algorithm. However, the particular choice of $F$ used in the rollout algorithm is responsible for special properties that are not shared by other types of greedy algorithms.

Let us denote by $\mathcal{H}$ the greedy algorithm described above and assume that it terminates starting from every node (this has to be verified independently). Let us also assume that whenever there is a tie in the minimization of Eq. (6.34), $\mathcal{H}$ resolves the tie in a manner that is fixed and independent of the starting node i of the path, e.g., by resolving the tie in favor of the
numerically smallest node $j$ that attains the minimum in Eq. (6.34). Then it can be seen that $\mathcal{H}$ is sequentially consistent, since by construction, every node on a path generated by $\mathcal{H}$ has the same projection.

For a sequentially consistent base heuristic $\mathcal{H}$, we will assume a restriction in the way the rollout algorithm $\mathcal{R} \mathcal{H}$ resolves ties in selecting the next node on its path; this restriction will guarantee that $\mathcal{R} \mathcal{H}$ is terminating. In particular, suppose that after $m$ steps, $\mathcal{R} \mathcal{H}$ has produced the node sequence $\left(8, i l, \ldots, i_{m}\right)$, and that the path generated by $\mathcal{H}$ starting from $i_{m}$ is $(\mathrm{im}, \mathrm{im}+1, \cdot l m+2, \ldots, \bar{i})$. Suppose that among the neighbor set $N\left(i_{m}\right)$, the node attains the minimum in the selection test

$$
\begin{equation*}
\min _{j E N\left(i_{r n}\right)} H(j) \tag{6.35}
\end{equation*}
$$

but there are also some other nodes, in addition to im+l, that attain this minimum. Then, we require that the tie is broken in favor of $\mathrm{im}+1$, i.e., that the next node added to the current sequence $\left(8, i l, \ldots, i_{m}\right)$ is $i m+l$. Under this convention for tie-breaking, we show in the following proposition that the rollout algorithm $\mathcal{R} \mathcal{H}$ terminates at a destination and yields a cost that is no larger than the cost yielded by the base heuristic $\mathcal{H}$. t
6.4.1: Let the base heuristic $\mathcal{H}$ be sequentially consistent. Then the rollout algorithm $\mathcal{R} \mathcal{H}$ is terminating. Furthermore, if (il, ..., $i_{\tilde{m}}$ ) is the path generated by $\mathcal{R} \mathcal{H}$ starting from a non-destination node il and ending at a destination node $i_{\tilde{m}}$, the cost of $\mathcal{R} \mathcal{H}$ starting from $\boldsymbol{i}_{1}$ is less or equal to the cost of $\mathcal{H}$ starting from $i_{1}$. In particular, we have

$$
\begin{equation*}
H\left(i_{1}\right) \geq H\left(i_{2}\right) \geq \cdots \geq H\left(i_{\tilde{m}-1}\right) \geq H\left(i_{\bar{m}}\right) \tag{6.36}
\end{equation*}
$$

[^2]Furthermore, for all $m=1, \ldots, \tilde{m}$,

$$
\begin{equation*}
H(i m)=\min \left\{H(i l), \min _{\mathrm{jEN}(\mathrm{il})} H(j), \ldots, \min _{\mathrm{jEN}\left(\mathrm{i}_{\mathrm{rn}-1}\right)} H(j)\right\} \tag{6.37}
\end{equation*}
$$

Proof: Let im and $\mathrm{im}+1$ be two successive nodes generated by $\mathcal{R H}$, and let (im, $i_{m+1}^{\prime}, i_{m+2}^{\prime}, \ldots, \bar{i}_{m}$ ) be the path generated by $\mathcal{H}$ starting from im, where $\bar{i}_{m}$ is the projection of $\mathrm{i}_{\mathrm{m}}$. Then, since $\mathcal{H}$ is sequentially consistent, we have

$$
H\left(i_{m}\right)=H\left(i_{m+1}^{\prime}\right)=g\left(\bar{i}_{m}\right) .
$$

Furthermore, since $i_{m+1}^{\prime} \mathrm{E} N\left(i_{m}\right)$, we have using the definition of $\mathcal{R H}$ ref. Eq. (6.33)]

$$
H\left(i_{m+1}^{\prime}\right) \geq \min _{j \in N\left(i_{m}\right)} H(j)=H(i m+l)
$$

Combining the last two relations, we obtain

$$
\begin{equation*}
H\left(i_{m}\right) \geq H\left(i_{m+1}\right) \quad \min _{j E N\left(i_{r n}\right)} H(j) \tag{6.38}
\end{equation*}
$$

To show that $\mathcal{R H}$ is terminating, note that in view of Eq. (6.38), either $H(i m)>H(i m+l)$, or else $H(i m)=H(i m+l)$. In the latter case, in view of the convention for breaking ties that occur in Eq. (6.35) and the sequential consistency of $\mathcal{H}$, the path generated by $\mathcal{H}$ starting from im+1 is the tail portion of the path generated by $\mathcal{H}$ starting from $i_{m}$, and has one arc less. Thus the number of nodes generated by $\mathcal{R H}$ between successive times that the inequality $H\left(i_{m}\right)>H\left(i_{m+1}\right)$ holds is finite. On the other hand, the inequality $H\left(i_{m}\right)>H(i m+l)$ can occur only a finite number of times, since the number of destination nodes is finite, and the destination node of the path generated by $\mathcal{H}$ starting from $\mathrm{i}_{\mathrm{m}}$ cannot be repeated if the inequality $H(i m)>H(i m+l)$ holds. Therefore, $\mathcal{R} \mathcal{H}$ is terminating.

If (il, $\ldots, i_{\bar{m}}$ ) is the path generated by $\mathcal{R} \mathcal{H}$, the relation (6.38) implies the desired relations (6.36) and (6.37). Q.E.D.

Proposition 6.4 .1 shows that in the sequentially consistent case, the rollout algorithm $\mathcal{R} \mathcal{H}$ has an important "automatic cost sorting" property, whereby it follows the best path generated by the base heuristic $\mathcal{H}$. In particular, when $\mathcal{R H}$ generates a path (il, $\ldots, i_{\tilde{m}}$ ), it does so by using $\mathcal{H}$ to generate a collection of other paths and corresponding projections starting from all the successor nodes of the intermediate nodes il, $\ldots, i_{\check{m}-1}$. However, (ii, $\ldots, i_{\tilde{m}}$ ) is guaranteed to be the best among this path collection and $i_{\bar{m}}$ has minimal cost among all generated projections [ef. Eq. (6.37)]. Of course this does not guarantee that the path generated by $\mathcal{R H}$ will be a
near-optimal path, because the collection of paths generated by $\mathcal{H}$ may be "poor." Still, the property whereby $\mathcal{R} \mathcal{H}$ at all times follows the best path found so far is intuitively reassuring.

The following example illustrates the preceding concepts.

## Example 6.4.5 (One-Dimensional Walk)

Consider a person who walks on a straight line and at each time period takes either a unit step to the left or a unit step to the right. There is a cost function assigning cost $g(i)$ to each integer i. Given an integer starting point on the line, the person wants to minimize the cost of the point where he will end up after a given and fixed number $N$ of steps

We can formulate this problem as a graph search problem of the type discussed in the preceding section. In particular, without loss of generality, let us assume that the starting point is the origin, so that the person's position after $n$ steps will be some integer in the interval $[-n, n]$. The nodes of the graph are identified with pairs $(k, m)$, where $k$ is the number of steps taken so far $(k=1, \ldots, N)$ and $m$ is the person's position ( $m E[-k, k]$ ). A node $(k, \mathrm{~m})$ with $k<N$ has two outgoing arcs with end nodes $(k+1, \mathrm{~m}-1)$ (corresponding to a left step) and $(k+1, \mathrm{~m}+1)$ (corresponding to a right step). The starting state is $(0,0)$ and the terminating states are of the form $(N, m)$, where m is of the form $N-2 l$ and $l \mathrm{E}[0, N]$ is the number of left steps taken.

Let the base heuristic $\mathcal{H}$ be defined as the algorithm, which, starting at a node $(k, m)$, takes $N-k$ successive steps to the right and terminates at the node $(N, \mathrm{~m}+\mathrm{N}-k)$. Note that $\mathcal{H}$ is sequentially consistent. The rollout algorithm $\mathcal{R H}$, at node ( $k, \mathrm{~m}$ ) compares the cost of the destination node $(N, \mathrm{~m}+\mathrm{N}-k)$ (corresponding to taking a step to the right and then following $\mathcal{H})$ and the cost of the destination node $(N, \mathrm{~m}+N-k-2)$ (corresponding to taking a step to the left and then following $\mathcal{H})$.

Let us say that an integer i $\mathrm{E}[-N+2, N-2]$ is a local minimum if $g(i-2) \geq g(i)$ and $g(i) \leq g(i+2)$. Let us also say that N (or $-N$ ) is a local minimum if $g(N-2) \leq g(N)$ [or $g(-N) \leq g(-N+2)$, respectively]. Then it can be seen that starting from the origin $(0,0), \mathcal{R H}$ obtains the local minimum that is closest to N, (see Fig. 6.4.3). This is no worse (and typically better) than the integer $N$ obtained by $\mathcal{H}$. Note that if 9 has a unique local minimum in the set of integers in the range $[-N, N]$, the minimum must also be global, and it will be found by $\mathcal{R H}$. This example illustrates how $\mathcal{R H}$ may exhibit "intelligence" that is totally lacking from $\mathcal{H}$, and is in agreement with the result of Prop. 6.4.1


Figure 6.4.3 Illustration of the path generated by the rollout algorithm $\mathcal{R} \mathcal{H}$ in Example 6.4.5. The algorithm keeps moving to the left up to the time where the base heuristic $\mathcal{H}$ generates two destinations $(N, \bar{i})$ and $(N, \bar{i}-\cdot 2)$ with $g(\bar{i}) \leq$ $g(\bar{i}-2)$. Then it continues to move to the right ending at the destination ( $N, 7$ ), which corresponds to the local minimum closest to $N$.

Definition 6.4.2: We say that the base heuristic $\mathcal{H}$ is sequentially improving if for every non-destination node i , we have

$$
\begin{equation*}
H(i) \geq \min _{\mathrm{jEN}(\mathrm{i})} H(j) \tag{6.39}
\end{equation*}
$$

It can be seen that a sequentially consistent $\mathcal{H}$ is also sequentially improving, since sequential consistency implies that $H$ (i) is equal to one of the values $H(j)$, j E $N(i)$. We have the following generalization of Prop. 6.4.1, which also bears a relation with the general cost estimate for one-step lookahead policies of Prop. 6.3.1.

Proposition 6.4.2: Let the base heuristic $\mathcal{H}$ be sequentially improving, and assume that the rollout algorithm $\mathcal{R H}$ is terminating. Let (il,"" $i_{\tilde{m}}$ ) be the path generated by $\mathcal{R H}$ starting from a nondestination node il and ending at a destination node $i_{\tilde{m}}$. Then the
cost of $\mathcal{R} \mathcal{H}$ starting from il is less or equal to the cost of $\mathcal{H}$ starting from il. In particular, we have for all $m=1, \ldots, \tilde{m}$,

$$
\begin{equation*}
H\left(i_{m}\right)=\min \left\{H\left(i_{l}\right), \min _{\mathrm{jEN}(\mathrm{il})} H(j), \ldots, \min _{\mathrm{jEN}\left(\mathrm{i}_{\mathrm{m}}-1\right)} H(j)\right\} \tag{6.40}
\end{equation*}
$$

Proof: For each $m=1, \ldots, \tilde{m}$ 1, we have

$$
H\left(i_{m}\right) \geq \min _{\mathrm{jEN}\left(\mathrm{i}_{\mathrm{m}}\right)} H(j)
$$

by the sequential improvement assumption, while we have

$$
\min _{j \in N\left(i_{m}\right)} H(j)=H\left(i_{m}+l\right)
$$

by the definition of the rollout algorithm. These two relations imply Eq. (6.40). Since the cost of $\mathcal{R} \mathcal{H}$ starting from il is $H\left(i_{\tilde{m}}\right)$, the result follows. Q.E.D.

## Example 6.4.6 ne-Dimensional Walk Continued)

Consider the one-dimensional walk problem of Example 6.4.5, and let $\mathcal{H}$ be defined as the algorithm that, starting at a node ( $k, m$ ), compares the cost $g(m+N-k)$ (corresponding to taking all of the remaining $N-k$ steps to the right) and the cost $9(\mathrm{~m}-N+k)$ (corresponding to taking all of the remaining $N \quad k$ steps to the left), and accordingly moves to node

$$
(N, m+N-k) \quad \text { if } \quad g(m+N-k) \leq g(m-N+k)
$$

or to node

$$
(N, m \quad N+k) \quad \text { if } \quad g(m-N+k)<g(m+N-k) .
$$

It can be seen that $\mathcal{H}$ is not sequentially consistent, but is instead sequentially improving. Using Eq. (6.40), it follows that starting from the origin $(0,0)$, $\mathcal{R} \mathcal{H}$ obtains the global minimum of 9 in the interval $[-N, N]$, while $\mathcal{H}$ obtains the better of the two points $-\mathrm{Nand} N$.

Proposition 6.4.2 actually follows from a general equation for the cost of the path generated by the rollout algorithm, which holds for any base heuristic (not necessarily one that is sequentially improving). This is given in the following proposition, which is related to Prop. 6.3.2.

Hroposition 6.4.3: Assume that the rollout algorithm $\mathcal{R H}$ is terminating. Let (il, ... $i_{\tilde{m}}$ ) be the path generated by $\mathcal{R} \mathcal{H}$ starting from a non-destination node il and ending at a destination node $i_{\tilde{m}}$. Then the cost of $\mathcal{R} \mathcal{H}$ starting from $\boldsymbol{i}_{1}$ is equal to

$$
H\left(i_{1}\right)+\delta_{i_{1}}+\cdots+\delta_{i_{\tilde{m}-1}}
$$

where for every non-destination node $i$, we denote

$$
\delta_{i}=\min _{\mathrm{jEN}(\mathrm{i})} H(j) \quad H(i)
$$

Proof: We have by the definition of the rollout algorithm

$$
H\left(i_{m}\right)+\delta_{i_{m}}=\min _{j \in N\left(i_{m}\right)} H(j)=H\left(i_{m+1}\right), \quad \mathrm{m}=1, \ldots, \tilde{m}-1
$$

By adding these equations over $m$, we obtain

$$
\begin{equation*}
H\left(i_{1}\right)+\delta_{i_{1}}+\cdots+\delta_{i_{\tilde{m}-1}}=H\left(i_{\tilde{m}}\right) \tag{I}
\end{equation*}
$$

Since the cost of $\mathcal{R} \mathcal{H}$ starting from il is $H\left(i_{\tilde{m}}\right)$, the result follows.
If the base heuristic is sequentially improving, we have $\delta_{i} \leq 0$ for all non-destination nodes i, so it follows from Prop. 6.4.3 that the cost of the rollout algorithm is less or equal to the cost of the base heuristic (cf. Prop. 6.4.2).

## The Fortified Rollout Algorithm

We now describe a variant of the rollout algorithm that implicitly uses a sequentially improving base heuristic, so that it has the cost improvement property of Prop. 6.4.2. This variant, called the fortified rollout algorithm, and denoted by $\mathcal{R} \mathcal{H}$, starts at the origin $s$, and after m steps, maintains, in addition to the current sequence of nodes ( $3, \mathrm{il}, \ldots, i_{m}$ ), a path

$$
\left.P\left(\imath_{m}^{\prime}\right)=\kappa_{m}^{\prime \prime}, t_{m+1}^{f}, \ldots, i_{k}^{f}\right)
$$

ending at a destination $i_{k}^{\prime}$. Roughly speaking, the path $P\left(\mathrm{i}_{\mathrm{m}}\right)$ is the tail portion of the best path found after the first m steps of the algorithm, in the sense that the destination $i_{k}^{\prime}$ has minimal cost over all the projections of nodes calculated thus far.

In particular, initially $P(s)$ is the path generated by the base heuristic $\mathcal{H}$ starting from $s$. At the typical step of the fortified rollout algorithm $\mathcal{R} \overline{\mathcal{H}}$,
we have a node sequence $\left(8, i_{1}, \ldots, i_{m}\right)$, where $i_{m}$ is not a destination, and the path $P(i m)=\left(i_{m}, i_{m+1}^{\prime}, \cdots, i_{k}^{\prime}\right)$. Then, if

$$
\begin{equation*}
\min _{j \in N\left(i_{m}\right)} H(j)<g\left(i_{k}^{\prime}\right) \tag{6.41}
\end{equation*}
$$

$\underline{\mathcal{R H}}$ adds to the node sequence $\left(s, i 1, \ldots, i_{m}\right)$ the node

$$
\mathrm{im}+1=\arg \min _{\mathrm{jEN}\left(\mathrm{i}_{\mathrm{m}}\right)} H(j),
$$

and sets $P(i m+l)$ to the path generated by $\mathcal{H}$, starting from im +1 . On the other hand, if

$$
\begin{equation*}
\min _{\mathrm{jEN}(\mathrm{im})} H(j) \geq g\left(i_{k}^{\prime}\right) \tag{6.42}
\end{equation*}
$$

$\mathcal{R H}$ adds to the node sequence $\left(s, i l, \ldots, i_{m}\right)$ the node

$$
i_{m+1}=i_{m+1}^{\prime}
$$

and sets $P\left(i_{m}+l\right)$ to the path (im+1 $\left.i_{m+2}^{\prime}, \ldots i_{k}^{\prime}\right)$. If $\mathrm{im}+1$ is a destination, $\mathcal{R} \mathcal{H}$ terminates, and otherwise $\underline{\mathcal{R} \mathcal{H}}$ repeats the process with ( $8, \mathrm{iI}, \ldots, \mathrm{im}+1$ ) replacing $\left(8, i_{1}, \ldots, i m\right)$, and $P(i m+l)$ replacing $P(i m)$, respectively.

The idea behind the construction of $\underline{\mathcal{R} \mathcal{H}}$ is to follow the path $P\left(i_{m}\right)$ unless a path of lower cost is discovered through Eq. (6.41). We can show that $\mathcal{R H}$ may be viewed as the rollout algorithm $\mathcal{R H}$ corresponding to a modified version of $\mathcal{H}$, called fortified $\mathcal{H}$, and denoted $\underline{\mathcal{H}}$. This algorithm is applied to a slightly modified version of the original problem, which involves an additional downstream neighbor for each node $\boldsymbol{i}_{\mathrm{m}}$ that is generated in the course of the algorithm $\underline{\mathcal{H} \mathcal{H}}$ and for which the condition (6.42) holds. For every such node $i_{m}$, the additional neighbor is a copy of $i_{m+1}^{\prime}$, and the path generated by $\underline{\mathcal{H}}$ starting from this copy is $\left(i_{m+1}^{\prime}, \ldots i_{k}^{\prime}\right)$. From every other node, the path generated by $\underline{\mathcal{H}}$ is the same as the path generated by $\mathcal{H}$.

It can be seen that $\underline{\mathcal{H}}$ is sequentially improving, so that $\underline{\mathcal{R H}}$ is terminating and has the automatic cost sorting property of Prop. 6.4.2; that is,

$$
H\left(i_{m}\right)=\min \left\{H\left(i_{1}\right), \min _{\mathrm{jEN}(\mathrm{il})} H(j), \ldots, \min _{\mathrm{jEN}(\mathrm{im}-1)} H(j)\right\}
$$

The above property can also be easily verified directly, using the definition of $\underline{\mathcal{R} \mathcal{H}}$. Finally, it can be seen that when $\mathcal{H}$ is sequentially consistent, the rollout algorithm $\mathcal{R H}$ and its fortified version $\mathcal{R} \mathcal{H}$ coincide.

## Using Multipn Path Construction Algorithms

In many problems, several promising path construction heuristics may be available. It is then possible to use all of these heuristics in the rollout framework. In particular, let us assume that we have $K$ algorithms $\mathcal{H}_{1}, \ldots, \mathcal{H}_{K}$. The kth of these algorithms, given a non-destination node i , produces a path $\left(i, \mathrm{il}, \ldots, i_{m}, \bar{i}\right)$ that ends at a destination node $\bar{i}$, and the corresponding cost is denoted by $H k(i)=g(\bar{i})$. We can incorporate the $K$ algorithms in a generalized version of the rollout algorithm, which uses the minimal cost

$$
\begin{equation*}
H(i)=\min _{k=1, \ldots, K} H k(i) \tag{6.43}
\end{equation*}
$$

in place of the cost obtained by anyone of the $K$ algorithms $\mathcal{H}_{1}, \ldots, \mathcal{H}_{K}$.
In particular, the algorithm starts with the origin node 8. At the typical step, given a node sequence ( $8, \mathrm{il}, " \mathrm{"} i_{m}$ ), where im is not a destination, the algorithm adds to the sequence a node im+l such that

$$
\mathrm{im}+1=\arg \min _{\mathrm{jEN}(\mathrm{im})} H(j)
$$

If im+l is a destination node, the algorithm terminates, and otherwise the process is repeated with the sequence $\left(8, \mathrm{il}, \ldots, i_{m}, i_{m+1}\right)$ replacing ( $8, \mathrm{il}, \ldots, i_{m}$ ).

An interesting property, which can be readily verified by using the definitions, is that if all the algorithms $\mathcal{H}_{1}, \ldots, \mathcal{H}_{K}$ are sequentially improving, the same is true for $\mathcal{H}$. This is consistent with the analysis of Example 6.3.2.

The fortified version of the rollout algorithm $\underline{\mathcal{R} \mathcal{H}}$ easily generalizes for the case of Eq. (6.43), by defining the path generated starting from a node $i$ as the path generated by the path construction algorithm, which attains the minimum in Eq. (6.43).

In an alternative version of the rollout algorithm that uses multiple path construction heuristics, the results of the $K$ algorithms $\mathcal{H}_{1}, \ldots, \mathcal{H}_{K}$ are weighted with some fixed scalar weights $r k$ to compute $H(i)$ for use in Eq. (6.33):

$$
\begin{equation*}
H(i)=\sum_{\mathrm{k}=1}^{K} r k H k(i) \tag{6.44}
\end{equation*}
$$

The weights $T k$ may be adjusted by trial and error. An alternative and more sophisticated possibility, is to use weights that depend on the node i and which are obtained by training using the neuro-dynamic programming methodology described in Vol. II.

## Extension for Intermediate Arc Costs

Let us consider a variant of the graph search problem where in addition to the terminal cost $g(i)$, there is a cost $\mathrm{c}(\mathrm{i}, j)$ for a path to traverse an arc
$(i, j)$. Within this context, the cost of a path $\left(i_{1}, i 2, \ldots, i n\right)$ that starts at il and ends at a destination node in is redefined to be

$$
\begin{equation*}
g\left(i_{n}\right)+\sum_{\mathrm{k}=1}^{\mathrm{n}-1} C\left(i k^{\prime} \mathrm{i}_{\mathrm{k}+1}\right) . \tag{6.45}
\end{equation*}
$$

Note that when the cost $g(\mathrm{i})$ is zero for all destination nodes i , this is the problem of finding a shortest path from the origin node 8 to one of the destination nodes, with $c(i, j)$ viewed as the length of arc $(i, j)$. We have seen in Chapter 2 that there are efficient algorithms for solving this problem. However, here we are interested in problems where the number of nodes is very large, and the use of the shortest path algorithms of Chapter 2 is impractical.

One way to transform the problem with arc costs into one involving a terminal cost only is to redefine the graph of the problem so that nodes correspond to sequences of nodes in the original problem graph. Thus if we have arrived at node $i k$ using path (il, $\ldots, i k$ ), the choice of $i k+1$ as the next node is viewed as a transition from (il, $\ldots, i k$ ) to (il, $\ldots, i k^{\prime} \mathrm{ik}+1$ ). Both nodes (il, $\ldots, i k)$ and ( $i_{1}, \ldots, i k, i_{k+1}$ ) are viewed as nodes of a redefined graph. Furthermore, in this redefined graph, a destination node has the form (iI, i2, ... in), where in is a destination node of the original graph, and has a cost given by (6.45).

After the details are worked out, we see that to recover our earlier algorithms and analysis, we need to modify the cost of the heuristic algorithm $\mathcal{H}$ as follows: If the path (il, .., in) is generated by $\mathcal{H}$ starting at iI, then

$$
H(i l)=g\left(i_{n}\right)+{ }_{\mathrm{k}=1}^{\mathbf{n}-\mathbf{1}} c\left(i_{k}, \mathrm{i} \mathrm{k}+1\right)
$$

Furthermore, the rollout algorithm $\mathcal{R H}$ at node $\mathrm{i}_{\mathrm{m}}$ selects as next node the node

$$
\mathrm{im}+1=\arg \min _{j \in N\left(i_{m}\right)}[c(i m, j)+H(j) J
$$

[ef. (6.33)]. The definition of a sequentially consistent algorithm remains unchanged. Furthermore, Prop. 6.4.1 remains unchanged except that Eqs. (6.36) and (6.37) are modified to read
$H\left(i_{k}\right) \geq c\left(i_{k}, i_{k+1}\right)+H\left(i_{k+1}\right)=\min _{j \in N\left(i_{k}\right)}\left[c\left(i_{k}, j\right)+H(j)\right], \quad k=1, \ldots, m-1$.
A sequentially improving algorithm should now be characterized by the property

$$
H\left(i_{k}\right) \geq c\left(i_{k}, \quad+H\left(i_{k+1}\right)\right.
$$

where is the next node on the path generated by $\mathcal{H}$ starting from $i k$. Furthermore, Prop. 6.4.2 remains unchanged, except that (6.40) is modified to read

$$
H\left(i_{k}\right) \geq \min _{j \in N\left(i_{k}\right)}\left[c\left(i_{k}, j\right)+H(j)\right], \quad k=1, \ldots, \mathrm{~m} \quad 1 .
$$

Finally, the criterion $\operatorname{minjEN}\left(\mathrm{i}_{\mathrm{m}}\right) H(j)<g\left(i_{k}^{\prime}\right)$ [ef. Eq. (6.41)] used in the fortified rollout algorithm, given the sequence $\left(8, \mathrm{il}, \ldots, i_{m}\right)$, where $\mathrm{i}_{\mathrm{m}} \notin \overline{\mathcal{N}}$, and the path $P\left(i_{m}\right)=\left(i_{m}, i_{m+1}^{\prime} 1^{\prime} \ldots, i_{k}^{\prime}\right)$, should be replaced by

$$
\min _{j E N\left(i^{m}\right)}\left[c(i m, j)+H(j) J<g\left(i_{k}^{\prime}\right)+c\left(i_{m}, i_{m+1}^{\prime}\right)+{ }_{l=m+l}^{\mathrm{k}-1} c\left(i_{l}^{\prime}, i_{l+1}^{\prime}\right)\right.
$$

Rollout Algorithms with mensumen Lookahead
We may incorporate multistep lookahead into the rollout framework. To describe the case of 2-step lookahead, suppose that after m steps of the rollout algorithm, we have the current node sequence ( $8, \mathrm{il}, \ldots, i_{m}$ ). We then consider the set of all 2 -step-ahead neighbors of $i_{m}$, defined as

$$
\begin{aligned}
& N 2\left(i_{m}\right)=\left\{j \mathrm{ENI} j \mathrm{E} N\left(i_{m}\right) \text { and } j \mathrm{E} \underline{N},\right. \\
& \\
& \text { or } \left.j \mathrm{E} N(n) \text { for some } n \mathrm{E} N\left(i_{m}\right)\right\} .
\end{aligned}
$$

We run the base heuristic $\mathcal{H}$ starting from each $j \mathrm{E} N 2\left(i_{m}\right)$ and we find the node $\bar{j} \mathrm{E} N 2\left(i_{m}\right)$ that has projection of minimum cost. Let $\mathrm{i}_{\mathrm{m}}+\mathrm{I} \mathrm{E} N\left(i_{m}\right)$ be the node next to $i_{m}$ on the (one- or two-arc) path from im to $\bar{j}$. If im +1 is a destination node, the algorithm terminates. Otherwise, the process is repeated with the sequence ( $8, \ldots, i_{m}, \mathrm{im}+1$ ) replacing ( $s$, $\qquad$ $\left.\ldots ., i_{m}\right)$.
Note that a fortified version of the rollout algorithm described above is possible along the lines described earlier. Also, it is possible to eliminate from the set $N 2\left(i_{m}\right)$ some of the 2 -step neighbors of $i_{m}$ that are judged less promising according to some heuristic criterion, in order to limit the number of applications of the base heuristic. This may be viewed as selective depth lookahead. Finally, the extension of the algorithm to lookahead of more than two steps is straightforward: we simply replace the 2 -stepahead neighbor set $N 2$ ( $i_{m}$ ) with a suitably defined k -step ahead neighbor set $N k\left(i_{m}\right)$.

## nterpretation in Terms of DP

Let us now reinterpret the graph-based rollout algorithm within the context of deterministic DP. We will aim to view the base heuristic as a suboptimal
policy, and to view the rollout algorithm as a policy obtained by a process of policy improvement, provided the base heuristic is sequentially consistent.

To this end, we cast the graph search problem as a sequential decision problem, where each node corresponds to a state of a dynamic system. At each non-destination node/state $i$, a node $j$ must be selected from the set of neighbors $N(\mathrm{i})$; then if $j$ is a destination, the process terminates with cost $g(j)$, and otherwise the process is repeated with $j$ becoming the new state. The DP algorithm calculates for every node i, the minimal cost that can be achieved starting from $i$, that is, the smallest value of $g(\bar{i})$ that can be obtained using paths that start from $i$ and end at destination nodes $\bar{i}$. This value, denoted.$I^{*}(i)$, is the optimal cost-to-go starting at node i. Once $J^{*}(i)$ is computed for all nodes i, an optimal path (iI, i2, $\ldots, i_{m}$ ) can be constructed starting from any initial node/state il by successively generating nodes using the relation

$$
\begin{equation*}
i_{k+1}=\arg \min _{j \in N\left(i_{k}\right)} J *(j), \quad k=1, \ldots, \mathrm{~m}-1 \tag{6.46}
\end{equation*}
$$

up to the point where a destination node $i_{m}$ is encountered. $t$
A base heuristic $\mathcal{H}$ defines a policy $\pi$, i.e., an assignment of a successor node to any non-destination node. However, starting from a given node i, the cost of $\pi$ need not be equal to $H(i)$ because if a path (iI, i2, i3, $\ldots, i_{m}$ ) is generated by $\mathcal{H}$ starting from node iI, it is not necessarily true that the path ( $\mathrm{i} 2, i 3, \ldots, i_{m}$ ) is generated by the base heuristic starting from $i 2$. Thus the successor node chosen at node i2 by policy $\pi$ may be different than the one used in the calculation of $H(i l)$. On the other hand, if $\mathcal{H}$ is sequentially consistent, the cost of policy $\pi$ starting from a node $i$ is $H$ (i), since sequential consistency implies that the path that the base heuristic generates starting at the successor node is part of the path it generates at the predecessor node. Thus the cost improvement property of the rollout algorithm in the sequentially consistent case also follows from the cost improvement property shown earlier in the DP context.

Generally, we can view the rollout algorithm $\mathcal{R} \mathcal{H}$ as a one-step lookahead policy that uses $H(j)$ as a cost-to-go approximation from state $j$. In some cases, $H(j)$ is the cost of some policy (in the DP sense), such as for example when $\mathcal{H}$ is sequentially consistent, as explained above. In general, however, this need not be so, in which case we can view $H(j)$ as a convenient cost-to-go approximation that is derived from the base heuristic. Still, the rollout algorithm $\mathcal{R H}$ may improve on the cost of the base heuristic (e.g., when $\mathcal{H}$ is sequentially improving, d. Prop. 6.4.2) just as a general one-step lookahead policy may improve on the corresponding one-step lookahead cost approximation (d. Prop. 6.3.1).
t We assume here that there are no termination/cycling difficulties of the type illustrated in the footnote following Example 6.4.4.

### 6.4.2 Q-Factors Evaluated by Simulation

We now consider a stochastic problem and some computational issues regarding the implementation of the rollout policy based on a given heuristic policy. A conceptually straightforward approach to compute the rollout control at a given state $X k$ and time $k$ is to use Monte-Carlo simulation. To implement this algorithm, we consider all possible controls $U k$ E $U k(X k)$ and we generate a "large" number of simulated trajectories of the system starting from $X k$, using $U k$ as the first control, and using the policy $\pi$ thereafter. Thus a simulated trajectory has the form

$$
x_{i+1}=f_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), w_{i}\right), \quad \mathrm{i}=k+1, \ldots, N-1
$$

where the first generated state is

$$
x_{k+1}=f_{k}\left(x_{k}, u_{k}, w_{k}\right)
$$

and each of the disturbances $W k, \ldots, W N-l$ is an independent random sample from the given distribution. The costs corresponding to these trajectories are averaged to compute an approximation $\tilde{Q}_{k}\left(x_{k}, U k\right)$ to the Q -factor

$$
E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+J_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}
$$

Here, $\tilde{Q}_{k}\left(x_{k}, U k\right)$ is an approximation to $Q k(X k, U k)$ because of the simulation error resulting from the use of a limited number of trajectories. The approximation becomes increasingly accurate as the number of simulated trajectories increases. Once the approximate Q-factor $\tilde{Q}_{k}\left(x_{k}, U k\right)$ corresponding to each control $U \mathrm{E} U k(X k)$ is computed, we can obtain the (approximate) rollout control $\bar{\mu}_{k}\left(x_{k}\right)$ by the minimization

$$
\bar{\mu}_{k}\left(x_{k}\right)=\arg \min _{u_{k} \in U_{k}\left(x_{k}\right)} \tilde{Q}_{k}\left(x_{k}, u_{k}\right)
$$

There is a serious flaw with this approach, due to the simulation error involved in the calculation of the Q-factors. In particular, for the calculation of $\bar{\mu}_{k}\left(x_{k}\right)$ to be accurate, the Q-factor differences

$$
Q_{k}\left(x_{k}, u_{k}\right)-Q_{k}\left(x_{k}, \hat{u}_{k}\right)
$$

must be computed accurately for all pairs of controls $U k$ and $\hat{u}_{k}$, so that these controls can be accurately compared. On the other hand, the simulation/approximation errors in the computation of the individual Q-factors $Q k(X k, U k)$ may be magnified through the preceding differencing operation.

An alternative approach is to approximate by simulation the Q-factor difference $Q k(X k, U k) \quad Q k\left(X k, \hat{u}_{k}\right)$ by sampling the difference

$$
C_{k}\left(x_{k}, u_{k}, \mathbf{w}_{k}\right)-C_{k}\left(x_{k}, \hat{u}_{k}, \mathbf{w}_{k}\right),
$$

where $w_{k}=\left(w_{k}, w_{k+1}, \ldots, w_{N-1}\right)$ and
$C k(X k, U k, W k) \quad g N(X N)+g d x k, U k, W k)+{ }_{\mathrm{i}=\mathrm{k}+\mathrm{I}}^{\mathrm{N}-\mathrm{I}} g_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), w_{i}\right)$.

This approximation may be far more accurate than the one obtained by differencing independent samples of $C k\left(X k, u_{k}, W k\right)$ and $C k\left(X k, \hat{u}_{k}, W k\right)$. Indeed, by introducing the zero mean sample errors

$$
D_{k}\left(x_{k}, u_{k}, \mathbf{w}_{k}\right)=C_{k}\left(x_{k}, u_{k}, \mathbf{w}_{k}\right)-Q_{k}\left(x_{k}, u_{k}\right)
$$

it can be seen that the variance of the error in estimating $Q k(X k, U k)$ $Q k\left(X k, \hat{u}_{k}\right)$ with the former method will be smaller than with the latter method if and only if

$$
\begin{aligned}
E_{W k}, \hat{\mathbf{w}}_{k}\left\{I D k\left(X k^{\prime} U k, W k\right)\right. & \left.-D k\left(X k, \hat{u}_{k}, \hat{\mathbf{w}}_{k}\right)_{1}{ }^{2}\right\} \\
& >E_{W k}\left\{I D k\left(X k^{\prime} U k, W k\right)-D k\left(X k, \hat{u}_{k},\right.\right.
\end{aligned}
$$

or equivalently

$$
\begin{equation*}
E\left\{D_{k}\left(x_{k}, u_{k}, w_{k}\right) D_{k}\left(x_{k}, \hat{u}_{k}, w_{k}\right)\right\}>0 \tag{6.47}
\end{equation*}
$$

i.e., if and only if the correlation between the errors $D k(X k, U k, W k)$ and $D k\left(X k, \hat{u}_{k}, W k\right)$ is positive. A little thought should convince the reader that this property is likely to hold in many types of problems. Roughly speaking, the relation (6.47) holds if changes in the value of $U k$ (at the first stage) have little effect on the value of the error $D k(X k, U k, W k)$ relative to the effect induced by the randomness of $W k$. In particular, suppose that there exists a scalar $\gamma<1$ such that, for all $X k, U k$, and $\hat{u}_{k}$, there holds

$$
\begin{equation*}
E\left\{I D k(x k, U k, W k)-D k\left(X k, \hat{u}_{k}, \mathrm{wk}\right) 1^{2}\right\} \leq \gamma E\left\{\operatorname{IDk}\left(X k^{\prime} U k, \mathrm{wk}\right) 1^{2}\right\} \tag{6.48}
\end{equation*}
$$

Then we have

$$
\begin{aligned}
& D k\left(X k, U k, w_{k}\right) D_{k}\left(x_{k}, \hat{u}_{k}, W k\right) \\
&= I D k(X k, U k, \\
&+D k(X k, U k, W k)\left(D k\left(X k, \hat{u}_{k}, W k\right)-D k(X k, U k, W k)\right) \\
& \geq I D k(X k, U k, W k)^{2} \\
&-I D k(X k, U k, W k)_{1} . \operatorname{IDk}\left(X k, \hat{u}_{k}, W k\right)-D k(X k, U k, W k) \mid
\end{aligned}
$$

from which we obtain, using also (6.48),

$$
\begin{aligned}
& E\left\{D_{k}\left(x_{k}, u_{k}, w k\right) D k\left(X k, \hat{u}_{k}, \quad\right\}\right. \\
& \geq E\{\operatorname{IDk}(X k, U k, \\
&-E\left\{\operatorname{IDk}\left(X k, U k, w_{k}\right)\left|\cdot \operatorname{ID} k(X k, \quad W k)-D k\left(X k, u_{k}, W k\right)\right|\right\} \\
& \geq E\left\{I D k\left(X k^{\prime} U k, w k\right) 1^{2}\right\}-\frac{1}{2} E\left\{\mid D_{k}\left(x_{k}, U k,\right.\right. \\
&-\frac{1}{2} E\left\{I D k\left(X k, \hat{u}_{k}, W k\right)-D k\left(X k, U k, \quad{ }^{2}\right\}\right. \\
& \geq \frac{1-}{2} \quad\left\{I D k(X k, U k, W k) 1^{2},\right.
\end{aligned}
$$

Thus, under the assumption (6.48) and the assumption

$$
E\left\{\left|D_{k}\left(x_{k}, u_{k}, \mathbf{w}_{k}\right)\right|^{2}\right\}>0
$$

the condition (6.47) holds and guarantees that by averaging cost difference samples rather than differencing (independently obtained) averages of cost samples, the simulation error variance decreases.

### 6.4.3 Q-Faetor Approximation

Let us now consider the case of a stochastic problem and various possibilities for approximating the costs-to-go $H k(X k), k=1, \ldots, \mathrm{~N} \quad 1$, of the base policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots, \mu_{N-1}\right\}$, rather than calculating them by Monte-Carlo simulation. For example, in a certainty equivalence approach, given a state $X k$ at time $k$, we fix the remaining disturbances at some "typical" values $\bar{w}_{k+1}, \ldots, \bar{w}_{N-1}$, and we approximate the true Q-factor

$$
Q_{k}\left(x_{k}, u_{k}\right)=E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+H_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}
$$

with

$$
\tilde{Q}_{k}\left(x_{k}, u_{k}\right)=E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{H}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}
$$

where $\quad\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)$ is obtained by

$$
\left(f k\left(X k, u_{k}, W k\right)\right)=g N(X N)+g_{i=k+\mathbf{l}}^{\mathbf{N}-\mathbf{I}} g_{i}\left(\bar{x}_{i}, \mu_{i}\left(x_{i}\right), \bar{w}_{i}\right)
$$

the initial state is

$$
\bar{x}_{k+1}=f_{k}\left(x_{k}, u_{k}, w_{k}\right)
$$

and the intermediate states are given by

$$
\bar{x}_{i+1}=f_{i}\left(\bar{x}_{i}, \mu_{i}\left(x_{i}\right), \bar{w}_{i}\right), \quad \mathrm{i}=k+1, \ldots, N-1
$$

Thus, in this approach, the rollout control is approximated by

$$
\tilde{\mu}_{k}\left(x_{k}\right)=\arg \min _{u_{k} \in U_{k}\left(x_{k}\right)} \tilde{Q}_{k}\left(x_{k}, u_{k}\right)
$$

Note that the approximate cost-to-go $\tilde{H}_{k+1}(X k+1)$ represents an approximation of the true cost-to-go $H k+l(X k+1)$ of the base policy based on a single sample (the nominal disturbances $W k+1,{ }^{\prime \prime \prime}, \bar{w}_{N-1}$ ). A potentially more accurate approximation is obtained using multiple nominal disturbance sequences and averaging the corresponding costs with appropriate nominal probabilities, similar to the scenario approximation approach of Example 6.3.6.

Let us also mention another approach for approximation of the cost-to-go $H_{k+1}$ of the base policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots, \mu_{N-1}\right\}$, using an approximation architecture. Here we calculate (possibly approximate) values for the cost-to__go of the base policy at a finite set of state-time pairs, and then we select the weights through a "least-squares fit" of these values.

In particular, suppose that we have calculated the correct value of the cost-to-go $H_{N-1}\left(x^{\cdot i}\right)$ at the next-to-last stage for $s$ states $x i, \mathrm{i}=1, \ldots, s$, through the DP formula

$$
\begin{aligned}
H_{N-1}\left(x_{N-1}\right)=E & \left\{g_{N-1}\left(x_{N-1}, \mu_{N-1}\left(x_{N-1}\right), w_{N-1}\right)\right. \\
& \left.+g N\left(i N-I\left(X N-1, \mu_{N-1}(X N-1), W N-1\right)\right)\right\}
\end{aligned}
$$

and the given terminal cost function $g N$. We can then approximate the entire function $H N-1(X N-1)$ by a function of some given form

$$
\tilde{H}_{N-1}\left(x_{N-1}, r_{N-1}\right)
$$

where $r_{N-1}$ is a vector of weights, which can be obtained by solving the problem

$$
\begin{equation*}
\min _{r} \sum_{i=1}^{s} H_{N-1}(x i)-\tilde{H}_{\mathrm{N}-1}(X i, r)_{1}^{2} \tag{6.50}
\end{equation*}
$$

For example if
is specified to be a linear function of $m$ features $Y l(x), \ldots, Y m(X)$,

$$
(x, r)=\sum_{j=l}^{\boldsymbol{m}} \operatorname{IjYj}(X)
$$

the least squares problem (6.50) is

$$
\left.\min _{r} \sum_{i=1}^{s}!H N_{-} l\left(X^{i}\right) \quad \sum_{j=1}^{m} r_{j} y_{j}\left(x^{i}\right)\right|^{2}
$$

This is a linear least squares problem that can be solved in closed form (its cost function is convex quadratic in the vector $r$ ).

Note that this approximation procedure can be enhanced if we have additional information on the true cost-to-go function $H_{N-1}\left(X_{N-1}\right)$. For example, if we know that $H_{N-} 1(X N-1) \geq 0$ for all $X N-l$, we may first compute the approximation $\tilde{H}_{N-1}\left(X_{N-1}, r_{N-1}\right)$ by solving the least-squares problem (6.50) and then replace this approximation by

$$
\max \left\{0, H N-1\left(X N-1, r_{N-1}\right)\right\}
$$

Once an approximating function $\tilde{H}_{N-1}\left(X_{N-1}, r_{N-1}\right)$ for the next-tolast stage has been obtained, it can be used to similarly obtain an approximating function $\tilde{H}_{N-2}\left(x_{N-2}, r_{N-2}\right)$. In particular, (approximate) cost-to-go function values $H N \_2\left(X^{\prime i}\right)$ are obtained for $s$ states $x^{\prime i}, \mathrm{i}=1, \ldots, s$, through the (approximate) DP formula

$$
\begin{aligned}
H N-2(X N-2)=E & \left\{g_{N-2}\left(x_{N-2}, \mu_{N-2}\left(x_{N-2}\right), w_{N-2}\right)\right. \\
& \left.+\tilde{H}_{N-1}\left(i N-2\left(X N-2, \mu_{N-2}\left(x_{N-2}\right), W N-2\right), I N-I\right)\right\}
\end{aligned}
$$

These values are used to approximate the cost-to-go function $\mathrm{H}_{N-} 2(X N-2)$ by a function of some given form

$$
f I N-2(x N-2, ' N-2),
$$

where $r_{N-2}$ is a vector of parameters, which is obtained by solving the problem

$$
\min _{r} \sum_{i=1}^{s} \mid \hat{H}_{N-2}\left(x^{i}\right)-H N_{-} 2\left(X^{i}, r\right)_{1}{ }^{2}
$$

The process can be similarly continued to obtain
$r_{k}$ ) up to $k=0$ by solving for each $k$ the problem

$$
\begin{equation*}
\min _{r} \sum_{i=1}^{s}\left|\hat{H}_{k}\left(x^{i}\right)-\tilde{H}_{k}\left(x^{i}, r\right)\right|^{2} \tag{6.51}
\end{equation*}
$$

Given the approximations $H O(X O, ' O), \ldots, \tilde{H}_{N-1}$
, IN-I) to the cost-to-go of the base policy, one may obtain a suboptimal policy by using at state-time pair $\left(X_{k}, k\right)$ the one-step lookahead control
$\bar{\mu}_{k}\left(x_{k}\right)=\arg \min _{U k E U k(X k)} E\left\{g k(X k, U k, W k)+\quad\left(i k\left(X k^{\prime} U k, W k\right), r_{k+1}\right)\right\}$.
This control must be calculated on-line once the state $X k$ at time $k$ becomes known.

### 6.5 MODEL PREDICTIVE CONTROL AND RELATED METHODS

In many control problems where the objective is to keep the state of a system near some desired point, the linear-quadratic models of Sections 4.1 and 5.3 are not satisfactory. There are two main reasons for this:
(a) The system may be nonlinear, and using for control purposes a model that is linearized around the desired point may be inappropriate.
(b) There may be control and/or state constraints, which are not handled adequately through a quadratic penalty on state and control. The reason may be special structure of the problem dictating that, for efficiency purposes, the system should often be operated at the boundary of its constraints. The solution obtained from a linear-quadratic model is not suitable for this, because the quadratic penalty on state and control tends to "blur" the boundaries of the constraints.
These inadequacies of the linear-quadratic model have motivated a form of suboptimal control, called model predictive control (MPC), which combines elements of several ideas that we have discussed so far: certainty equivalent control, multistage lookahead, and rollout algorithms. We will focus primarily on the most common form of MPC, where the system is either deterministic, or else it is stochastic, but it is replaced with a deterministic version by using typical values in place of all uncertain quantities, as in the certainty equivalent control approach. At each stage, a (deterministic) optimal control problem is solved over a fixed length horizon, starting from the current state. The first component of the corresponding optimal policy is then used as the control of the current stage, while the remaining components are discarded. The process is then repeated at the next stage, once the next state is revealed. We will also briefly discuss a version of MPC where there is uncertainty with a set-membership description.

The primary objective in MPC, aside from fulfilling the state and control constraints of the problem, is to obtain a stable closed-loop system. Note here that we may only be able to guarantee the stability of the deterministic model that forms the basis for the calculations of the MPC. This is consistent with a common practice in control theory: design a stable controller for a deterministic model of the system, and expect that it will provide some form of stability in a realistic stochastic environment as well.

In Section 6.5.2, we will discuss the mechanism by which stability is achieved by MPC under some reasonable conditions. We will first discuss in the next subsection some issues of multistage lookahead, which are relevant to MPC, but are also important in a broader context. In Section 6.5.3, we provide a general unifying framework for suboptimal control, which includes as special cases several approaches discussed in this chapter, OLFC, rollout, and MPC, and captures the mathematical essence of their attractive properties.

### 6.5.1 Rolling Horizon Approximati

Let us consider the l-step lookahead policy when the cost-to-go approximation is just zero. With this policy, at each stage we apply a control that would be optimal if the remaining horizon length were $/$ and there were no terminal cost. Thus at the typical stage $k$ we ignore the costs incurred in stages $k+I+1$ and beyond, and accordingly we neglect the corresponding long-range effects of our action. We call this the rolling horizon approach. In a variant of this approach, following the $l$ steps of lookahead, we use a cost-to-go approximation that is equal to the terminal cost function $g N$. This is essential if $g N$ is significant relative to the costs per stage accumulated over / stages.

We may also use a rolling horizon approach for infinite horizon problems. Then the length of the horizon of the problem solved at each stage stays the same at all stages. As a result, for a time-invariant system and cost per stage, the rolling horizon approach produces a stationary policy (the controls applied at the same state but in different stages are the same). This is a generic characteristic of infinite horizon control, as we have seen in the context of linear-quadratic problems (see also the discussion of Vol. II).

Naturally, a policy obtained using a rolling horizon is typically not optimal. One is tempted to conjecture that if the size of the lookahead $l$ is increased, then the performance of the rolling horizon policy is improved. This, however, need not be true as the following example shows.

## Example 6.5.1

This is an oversimplified problem, which, however, demonstrates the basic pitfall of the rolling horizon approach.


Figure 6.5.1 The problem of Example 6.5.1.

Consider a deterministic setting where at the initial state there are two possible controls, say 1 and 2 (see Fig. 6.5.1). At all other states there is only one available control, so a policy consists of just the initial choice between controls 1 and 2. Suppose that (based on the cost of the subsequent $N$ stages) control 1 is optimal. Suppose also that if control 2 is chosen, an "unfavorable" (high cost) state results after / transitions, followed by a "particularly favorable" state, which is then followed by other "unfavorable" states. Then, in contrast with the $l$-step lookahead policy, the $(l+\mathrm{I})$-step lookahead policy may view the inferior control 2 as being better, because it may be "fooled" by the presence of the "particularly favorable" state $l+1$ transitions ahead.

The rolling horizon approach is also interesting in the context of a rollout algorithm, where we need to calculate the cost-to-go of the base policy at various states. It is possible to use a rolling horizon approximation in the calculation of this cost-to-go. Thus, from the given state, we calculate the cost of the base policy over a fixed number of stages, rather than over the entire remaining horizon. This can result in significant computational savings. Furthermore, there may also be an improvement in the performance of the rollout policy if a rolling horizon approximation is used. One reason is the phenomenon illustrated in the preceding example. In fact, because of the suboptimality of the base policy, this phenomenon can get exaggerated, as shown in the following example.

## Example 6.5.2

Consider an $N$-stage stopping problem where at each stage we may either stop with a stopping gest equal to 0 , or continue at a certain cost that is either $-\epsilon$ or 1, where $<\epsilon<\operatorname{liN}$ (see Fig. 6.5.2). Let the first state with continuation cost equal to 1 be state m . Then the optimal policy is to stop after m steps at state m . The corresponding optimal cost is $-m \epsilon$. It can also be seen that an $l$-step rolling horizon approach with the cost evaluated optimally over the $/$ steps (rather than suboptimally using a base heuristic) is optimal.

Consider now the rollout policy where the base heuristic is to continue at every state (except the last where stopping is mandatory). It can be seen that this policy will stop at the initial state at a cost of 0 , since it will evaluate the continuation action as having positive cost, in view of the fact $\mathbf{1 - N \epsilon > 0 ,}$ and will thus prefer the stopping action. However, the rollout policy that uses a rolling horizon of $/$ stages, with $/ \leq \mathrm{m}$, will continue up to the first $\mathrm{m}-/+1$ stages, thus compiling a cost of $-(m-I+1) \mathrm{E}$. Thus, as the length $/$ of the rolling horizon becomes shorter, the performance of the rollout policy improves!

Another example of a rollout algorithm whose performance can be improved by using a rolling horizon approximation is the breakthrough problem of Example 6.4.2. In this case, the evolution of the system under


Figure 6.5.2 The problem of Example 6.5.2.
a rollout algorithm, where the greedy heuristic is evaluated using an l-step rolling horizon approximation, can be modeled using a Markov chain with $I+1$ states (see Exercise 6.18 ). Using this Markov chain, it is possible to ascertain that for a problem with a large number of steps the length of the rolling horizon that maximizes the breakthrough probability approaches an optimal value that is essentially independent of $N$.

### 6.5.2 Stability Issues in Mode! Predictive Control

As mentioned earlier, model predictive control was initially motivated by the desire to introduce nonlinearities, and control and/or state constraints into the linear-quadratic framework, and obtain a suboptimal but stable closed-loop system. With this in mind, we will describe MPC for the case of a stationary, possibly nonlinear, deterministic system, where state and control belong to some Euclidean spaces. The system is

$$
X k+l=f(X k, U k), \quad k=0,1, \ldots,
$$

and the cost per stage is quadratic:

$$
x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}, \quad k=0,1, \ldots
$$

where $Q$ and $R$ are positive definite symmetric matrices. We impose state and control constraints

$$
X k \text { EX, } \quad u_{k} \in U\left(x_{k}\right), \quad \mathbf{k}=\mathbf{O}, 1, \ldots,
$$

and we assume that the set $X$ contains the origin of the corresponding Euclidean space. Furthermore, if the system iof at the origin, it can be kept there at no cost with control equal to 0 , i.e., $\mathrm{E} U(O)$ and $f(O, O)=0$. We want to derive a stationary feedback controller that applies control $\bar{\mu}(x)$ at state $x$, and is such that, for all initial states $X_{0} E$ the state of the closed-loop system

$$
X k+l=f\left(X k^{\prime} \bar{\mu}\left(x_{k}\right)\right)
$$

satisfies the state and control constraints, and the total cost over an infinite number of stages is finite:

$$
\begin{equation*}
{ }_{k=O}^{\text {оо }}\left(x_{k}^{\prime} Q x_{k}+\bar{\mu}\left(x_{k}\right)^{\prime} R \bar{\mu}\left(x_{k}\right)\right)<00 . \tag{6.52}
\end{equation*}
$$

Note that because of the positive definiteness of $Q$ and $R$, the feedback controller $\bar{\mu}$ is stable in the sense that $X k \rightarrow 0$ and $\bar{\mu}\left(x_{k}\right) \rightarrow 0$ for all initial states Xo EX. (In the case of a linear system, the assumption of positive definiteness of $Q$ may be relaxed to positive semidefiniteness, together with an observability assumption of the type introduced in Section 4.1 and Prop. 4.1.1.)

In order for such a controller to exist, it is evidently sufficient [in view of the assumption that $1(0,0)=0$ ] that there exists a positive integer m such that for every init'ial state Xo $\mathbf{E X}$, one can find a sequence of controls ' $U k, k=0,1, \ldots, \mathrm{~m}-1$, which drive to 0 the state $X_{m}$ of the system at time m , while keeping all the preceding states $X l, X 2, \ldots, X m-l$ w'ith'in X and satisfy'ing the control constraints Uo E $U(x o), \ldots, \mathrm{Um}-1 \mathrm{E}$ $U(X m-l)$. We refer to this as the constrained controllability assumption (d. the corresponding assumption of Section 4.1). In practical applications, this assumption can often be checked easily. Alternatively, the state and control constraints can be constructed in a way that the assumption is satisfied; the methodology of reachability of target tubes, discussed in Section 4.6.2, can be used for this purpose.

Let us now describe a form of IVIPC under the preceding assumption. At each stage $k$ and state $X k$ EX, it solves an m-stage deterministic optimal control problem involving the same quadratic cost and the requirement that the state after m stages be exactly equal to 0 . This is the problem of minimizing

$$
\sum_{i=k}^{\mathrm{k}+\mathrm{m}-1}\left(x_{i}^{\prime} Q x_{i}+u_{i}^{\prime} R u_{i}\right)
$$

subject to the system equation constraints

$$
x_{i+1}=f\left(x_{i}, u_{i}\right), \quad \mathrm{i}=k, k+1, \ldots, k+\mathrm{m}-1,
$$

the state and control constraints

$$
X^{\prime \prime} \mathrm{E} \quad u_{i} \in U\left(x_{i}\right), \quad \mathrm{i}=k, k+1, \ldots, k+\mathrm{m}-1,
$$

and the terminal state constraint

$$
X k+m=0
$$

By the constrained controllability assumption, this problem has a feasible solution. Let $\left\{\bar{u}_{k}, \bar{u}_{k+1}, \ldots \bar{u}_{k+m-1}\right\}$ be a corresponding optimal control sequence. The MPC applies at stage $k$ the first component of this sequence,

$$
\bar{\mu}\left(x_{k}\right)=\bar{u}_{k},
$$

and discards the remaining components.

## Example 6.5.3

Consider the scalar linear system

$$
x_{k+1}=x_{k}+u_{k},
$$

and the state and control constraints

$$
x_{k} \operatorname{EX}=\{x| | x \mid \leq \operatorname{lo} 5\}, \quad u_{k} \in U\left(x_{k}\right)=\{u| | u \mid \leq 1\} .
$$

Let also $\mathrm{Q}=R=1$. We select $m=2$. For this value of m , the constrained controllability assumption is satisfied.

When at state $x_{k} \mathrm{EX}$, the MPC minimizes the two-stage cost

$$
x_{k}^{2}+u_{k}^{2}+(X k+U k) 2+u_{k+1}^{2},
$$

subject to the control constraints

$$
\left|u_{k}\right| \leq 1, \quad\left|u_{k+1}\right| \leq 1,
$$

and the state constraints

$$
X k+l \mathrm{EX}, \quad x_{k+2}=x_{k}+u_{k}+u_{k+1}=0
$$

It is easily verified that this minimization yields $U k+l-\left(x_{k}+u_{k}\right)$ and $U k=-(2 / 3) X k$. Thus the MPC takes the form

$$
\bar{\mu}\left(x_{k}\right)=-\frac{2}{3} x_{k}
$$

and the dosed-loop system is

$$
x_{k+1}=\frac{1}{3} x_{k}, \quad k=0,1, \ldots
$$

Jote that while the dosed-loop system is stable, its state is never driven to if started from $x o \neq 0$.

We now show that the MPC satisfies the stability condition (6.52). Let $X o, U o, X l$, U1,. .. be the state and control sequence generated by MPC, so that

$$
u_{k}=\bar{\mu}\left(x_{k}\right), \quad x_{k+1}=f\left(x_{k}, \bar{\mu}\left(x_{k}\right)\right), \quad k=0,1, \ldots
$$

Denote $J(x)$ the optimal cost of the $m$-stage problem solved by MPC when at a state $x$ EX. Let also $j(x)$ be the optimal cost starting at $x$ of a corresponding ( $m-1$ )-stage problem, i.e., the optimal value of the quadratic cost

$$
\sum_{k=0}^{m-2}\left(x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}\right)
$$

where $X_{0}=x$, subject to the constraints

$$
u_{k} \in \bar{U}\left(x_{k}\right), \quad k=0,1, \ldots, m-2
$$

and

$$
\text { Xm-1 } 0 .
$$

[For states $x$ E X for which this problem does not have a feasible solution, we write $J\left(x_{\varnothing}=00\right.$.] Since having one less stage in our disposal to drive the state to cannot decrease the optimal cost, we have for all $x \mathrm{E}$ X

$$
\begin{equation*}
J(x) \leq J(x) \tag{6.53}
\end{equation*}
$$

From the definitions of $\hat{J}$ and $J$, we have for all $k$,
$\min _{u E U(x)}\left[x_{k}^{\prime} Q x_{k}+u^{\prime} R u+J(f(x k, u))\right]=x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}+\tilde{J}\left(x_{k+1}\right)=J(X k)$,
so using Eq. (6.53), we obtain

$$
x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}+J(X k+l) \leq J(X k), \quad k=0,1, \ldots
$$

Adding this equation for all $k$ in a range $[0, K]$, where $K=0,1, \ldots$, we obtain

$$
J(X K+l)+\sum_{k=0}^{K}\left(x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}\right) \leq J(x o)
$$

Since $\hat{J}\left(x_{K+1}\right) \geq 0$, it follows that

$$
\begin{equation*}
\sum_{k=0}^{K}\left(x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}\right) \leq J(x o), \quad K=O, I, \ldots, \tag{6.55}
\end{equation*}
$$

and taking the limit as $K \rightarrow 00$,

$$
\sum_{k=0}^{\infty}\left(x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}\right) \leq J(x o)<00 .
$$

This shows the stability condition (6.52).
We note that the one-step lookahead funct'ion $\mathbf{J}$ implicitly used by lIIIPC [cf. Eg. (6.54)\} is the cost-to-go function of a certain policy. This is the policy that drives to 0 the state after $m-1$ stages and keeps the state at 0 thereafter, while observing the state and control constraints $X k$ E X and $U k E U(X k)$, and minimizing the quadratic cost. Thus, we can also view IVIPC as a rollout algorithm with the policy just described viewed as the base heuristic. In fact the stability property of MPC is a special case of the cost improvement property of rollout algorithms, which in the
case of a quadratic cost, implies that if the base heuristic results in a stable closed-loop system) the same is true for the corresponding rollout algorithm.

Regarding the choice of the horizon length $m$ for the IVIPC calculations, note that if the constrained controllability assumption is satisfied for some value of $m$, it is also satisfied for all larger values of $m$. Furthermore, it can be seen that the m -stage cost $J(x)$, which by Eq. (6.55), is an upper bound to the cost of MPC, cannot increase with $m$. This argues for a larger value of m . On the other hand, the optimal control problem solved at each stage by the IVIPC becomes larger and hence more difficult as $m$ increases. Thus, the horizon length is usually chosen on the basis of some experimentation: first use target tube reachability methods (cf. Section 4.6.2) to ensure that $m$ is large enough for the constrained controllability assumption to hold with a target tube that is sufficiently large for the practical problem at hand, and then by further experimentation to ensure overall satisfactory performance.

The MPC scheme that we have described is just the starting point for a broad methodology with many variations, which often relate to the suboptimal control methods that we have discussed so far in this chapter. For example, in the problem solved by MPC at each stage, instead of the requirement of driving the system state to 0 in $m$ steps, one may use a large penalty for the state being nonzero after $m$ steps. Then, the preceding analysis goes through, as long as the terminal penalty is chosen so that Eq. (6.53) is satisfied. In another variant one may use a nonquadratic cost function, which is everywhere positive except at $(x, \varphi \boldsymbol{D}=(0,0)$. In still another variant, instead of aiming to drive the state to after $m$ steps, one aims to reach a sufficiently small neighborhood of the origin, within which a stabilizing controller, designed by other methods, may be used. This variant is also well-suited for taking into account disturbances described by set membership, as we now proceed to explain.

## MPC with Set-Membership Disturbances

To extend the MPC methodology to the case where there are disturbances $W k$ in the system equation

$$
x_{k+1}=f\left(x_{k}, u_{k}, w_{k}\right)
$$

we must first modify the stability objective. The reason is that in the presence of disturbances, the stability condition (6.52) is impossible to meet. A reasonable alternative is to introduce a set-membership constraint $W k \mathrm{E} W(X k, U k)$ for the disturbance and a target set $T$ for the state, and to require that the controller specified by MPC drives the state to $T$ with finite quadratic cost.

To formulate the MPC, we assume that $T$ C and that once the system state enters $T$, we will use some control law $\tilde{\mu}$ that keeps the state
within $T$ for all possible values of the disturbances, i.e.,

$$
\begin{equation*}
f(x, \tilde{\mu}(x), w) \mathrm{E} T, \quad \text { for all } x \in T, w \in W(x, \tilde{\mu}(x)) \tag{6.56}
\end{equation*}
$$

The detailed methodology by which such a target set $T$ and control law $\tilde{\mu}$ are obtained is outside our scope. We refer to the discussion of reachability of target tubes in Section 4.6 .2 for orientation into this problem and references; see also Exercise 4.31, and Exercises 3.21 and 3.22 of Vol. II. We view $T$ essentially as a cost-free and absorbing state, similar to our view of the origin in the earlier deterministic context. Consistent with this interpretation, we introduce the stage cost function

$$
\begin{array}{ll}
g(x, u)=\left\{X 01 Q x+u^{\prime} R u\right. & \text { if } x \notin T, \\
& \text { if } x \mathrm{E} T
\end{array}
$$

The MPC is now defined as follows: At each stage $k$ and state $x k$ EX with $x k \notin T$, it solves the $m$-stage minimax control problem of finding a policy $\hat{\mu}_{k}, \hat{\mu}_{k+1}, \ldots, \hat{\mu}_{k+m-1}$ that minimizes

$$
\max _{\substack{W i E W(X i, P,(X i », i=k, k+l, \ldots, k+m-1}} \sum_{i=k}^{\mathbf{k}+\mathbf{m}-1} g\left(x_{i}, \mu\left(x_{i}\right)\right),
$$

subject to the system equation constraints

$$
x_{i+1}=f\left(x_{i}, u_{i}, w_{i}\right), \quad k, k+1, \ldots, k+\mathrm{m}-1
$$

the control and state constraints

$$
X i \mathrm{E} \quad U i \mathrm{E} U(X i), \quad \mathrm{i}=k, k+1, \ldots, k+\mathrm{m}-1,
$$

and the terminal state constraint

$$
X i \mathrm{E} T, \quad \text { for some } \mathrm{i} \mathrm{E}[k+1, k+\mathrm{m}] .
$$

These constraints must be satisfied for all disturbance sequences satisfying

$$
w_{i} \in W\left(x_{i}, \hat{\mu}\left(x_{i}\right)\right), \quad \mathrm{i}=k, k+1, \ldots, k^{\prime}+\mathrm{m}-1
$$

The MPC applies at stage $k$ the first component of the policy $\hat{\mu}_{k}, \hat{\mu}_{k+1}$, $\ldots, \hat{\mu}_{k+m-1}$ thus obtained,

$$
\bar{\mu}\left(x_{k}\right)=\hat{\mu}_{k}\left(x_{k}\right)
$$

and discards the remaining components. For states $x$ within the target set $T$, the MPC applies the control $\tilde{\mu}(x)$ that keeps the state within $T$, as per Eq. (6.56), at no further cost $[\bar{\mu}(x)=\tilde{\mu}(x)$ for $x \in T]$.

We make a constrained controllability assumption, namely that the problem solved at each stage by MPC has a feasible solution for all $x k$ E $X$ with $x k \notin T$ (this assumption can be checked using the target tube reachability methods of Section 4.6.2). Note that this problem is a potentially difficult minimax control problem, which generally must be solved by DP (cf. the algorithm of Section 1.6).

## Example 6.5.4

This example is a version of the preceding one, modified to account for the presence of disturbances. We consider the scalar linear system

$$
x_{k+1}=x_{k}+u_{k}+w_{k}
$$

and the state and contml constraints

$$
x k \mathrm{E} X=\{x \text { Ilx1 } \leq 1.5\}, \quad u_{k} \in U\left(x_{k}\right)=\{u| | u \mid \leq 1\}
$$

and assume that the disturbances satisfy

$$
w_{k} \in W\left(x_{k}, u_{k}\right)=\{w| | w \mid \leq 0.2\}
$$

We select $m=2$, and it can be verified that for the target set

$$
T=\{x \mathrm{Ilx1} \leq 0.2\}
$$

the constrained controllability assumption is satisfied, and the condition (6.56) is also satisfied using some control law $\tilde{\mu}$ namely $\tilde{\mu}(x)=-x$.

The associated 2 -stage minimax control problem to be solved at each stage by MPC requires a DP solution. At the last stage, assuming $x \notin T$, the DP algorithm calculates
$J(x)=$

$$
\min _{\substack{|u| \leq 1,|x+u+\bar{w}| \leq 0.2 \text { for all }|\bar{w}| \leq 0.2}}^{\left.\max _{|w| \leq 0.2}\left(x^{2}+u^{2}\right)\right] .}
$$

This is a straightforward minimization. It is feasible if and only if Ixl $\leq 1$ and it yields a minimizing policy for the last stage:

$$
\hat{\mu}_{1}(x)=-x, \quad \text { for all } x \notin T \text { with Ixl } \leq 1,
$$

and a cost-to-go

$$
\tilde{J}(x)=2 x^{2}, \quad \text { for all } x \notin T \text { with } \text { Ixl } \leq 1 .
$$

At the first stage, the DP algorithm calculates

$$
\left.\min _{\substack{|u| \leq 1,|x+u+\bar{w}| \leq 1 \text { for all }|\bar{w}| \leq 0.2}} \max _{||w| \leq 0.2}\left(X^{2}+U^{2}+J(x+u+w)\right)\right]
$$

Of, since the maximum over $w$ is attained for $w \quad 0.2 \operatorname{sgn}(x+u)$,

$$
\min _{\substack{|u| \leq 1, 2 \operatorname{sgn}(x+u) \mid \leq 1}}\left[x^{2}+u^{2}+2(x+u+0.2 \operatorname{sgn}(x+u)) 2\right]
$$

or

$$
\min _{\substack{|u| \leq 1,|x+u| \leq 0.8}}\left[x^{2}+u^{2}+2\left(x^{2}+\quad+2 x u+0.41 x+u \mid+0.04\right)\right]
$$

This minimization is again straightforward, and yields the MPC

$$
\begin{array}{rll}
\bar{\mu}(x)=\left\{-\min \left[x, \frac{2}{3}(x+0.2)\right]\right. & \text { if } x \mathrm{E}(0.2,1.5] \\
& \min \left[-x,-\frac{2}{3}(x-0.2)\right] & \text { if } x \mathrm{E}[-1.5,-0.2)
\end{array}
$$

This piecewise linear form of the MPC should be compared with the corresponding linear form, $\bar{\mu}(x)=-(2 / 3) x$, of Example 6.5 .3 , in which there are no disturbances.

The stability analysis of JIIIPC (in the modified sense of reaching the target set $T$ with finite quadratic cost, for all possible disturbance values) is similar to the one given earlier in the absence of disturbances. It is also possible to view JIIIPC in the presence of disturbances as a special case of a rollout algorithm, suitably modified to take account of the set-membership description of the disturbances. The details of this analysis are sketched in Exercise 6.21.

### 6.5.3 Restricted Structure Policies

We will now introduce a general unifying suboptimal control scheme that contains as special cases several of the control schemes we have discussed: OLFC, POLFC, rollout, and MPC. The idea is to simplify the problem by selectively restricting the information and/or the controls available to the controller, thereby obtaining a restricted but more tractable problem structure, which can be used conveniently in a one-step lookahead context.

An example of such a structure is one where fewer observations are obtained, or one where the control constraint set is restricted to a single or a small number of given controls at each state. Generally, a restricted structure is associated with a problem where the optimal cost achievable is less favorable than in the given problem; this will be made specific in what follows. At each stage, we compute a policy that solves an optimal control problem involving the remaining stages and the restricted problem structure. The control applied at the given stage is the first component of the restricted policy thus obtained.

An example of a suboptimal control approach that uses a restricted structure is the OLFC, where one uses the information available at a given stage as the starting point for an open-loop computation (where future observations are ignored). Another example is the rollout algorithm, where at a given stage one restricts the controls available at future stages to be those applied by some heuristic policy. Still another example is MPC, which
under some conditions may be viewed as a form of rollout algorithm, as discussed in the preceding subsection.

For a problem with N stages, implementation of the suboptimal scheme to be discussed requires the solution of a problem involving the restricted structure at each stage. The horizon of this problem starts at the current stage, call it $k$, and extends up to the final stage $N$. This solution yields a control $u_{k}$ for stage $k$ and a policy for the remaining stages $k+1, \ldots, N-1$ (which must obey the constraints of the restricted structure). The control $U k$ is used at the current stage, while the policy for the remaining stages $k+1, \ldots, \mathbf{N}-1$ is discarded. The process is repeated at the next stage $k+1$, using the additional information obtained between stages $k$ and $k+1$; this is similar to CEC, OLFC, multistage lookahead, and MPC.

Similarly, for an infinite horizon model, implementation of the suboptimal scheme requires, at each stage $k$, the solution of a problem involving the restricted structure and a (rolling) horizon of fixed length. The solution yields a control $U k$ for stage $k$ and a policy for each of the remaining stages. The control $U k$ is then used at stage $k$, and the policy for the remaining stages is discarded. For simplicity in what follows, we will focus attention to the finite horizon case, but the analysis applies, with minor modifications, to infinite horizon cases as well.

Our main result is that the performance of the suboptimal control scheme is no worse than the one of the restricted problem, i.e., the problem corresponding to the restricted structure. This result unifies and generalizes our analysis for open-loop-feedback control (which is known to improve the cost of the optimal open-loop policy, cf. Section 6.2), for the rollout algorithm (which is known to improve the cost of the corresponding heuristic policy, cf. Section 6.4), and for model predictive control (where under some reasonable assumptions, stability of the suboptimal closed-loop control scheme is guaranteed, ef. Section 6.5.2).

For simplicity, we focus on the imperfect state information framework for stationary finite-state Markov chains with N stages (ef. Section 5.4.2); the ideas apply to much more general problems with perfect and imperfect state information, as well problems with an infinite horizon. We assume that the system state is one of a finite number of states denoted $1,2, \ldots, n$. When a control $u$ is applied, the system moves from state $i$ to state $j$ with probability $P i j(U)$. The control $u$ is chosen from a finite set $U$. Following a state transition, an observation is made by the controller. There is a finite number of possible observation outcomes, and the probability of each depends on the current state and the preceding control. The information available to the controller at stage $k$ is the information vector

$$
I_{k}=\left(z_{1}, \ldots, z_{k}, u_{0}, \ldots, u_{k-1}\right)
$$

where for all i, $z_{i}$ and $U_{i}$ are the observation and control at stage i , respectively. Following the observation $z k$, a control $U k$ is chosen by the controller,
and a cost $g\left(X k^{\prime} u_{k}\right)$ is incurred, where $X k$ is the current (hidden) state. The terminal cost for being at state $x$ at the end of the $N$ stages is denoted $G(x)$. We wish to minimize the expected value of the sum of costs incurred over the $N$ stages.

As discussed in Section 5.4, we can reformulate the problem into a problem of perfect state information where the objective is to control the column vector of conditional probabilities

$$
p_{k}=\left(p_{k}^{1}, \ldots, p_{k}^{n}\right)^{\prime}
$$

with

$$
p_{k}^{j}=P\left(X k \quad j \mid I_{k}\right), \quad j=1, \ldots, n o
$$

We refer to $P k$ as the belief state, and we note that it evolves according to an equation of the form

$$
p_{k+1}=\Phi\left(p_{k}, u_{k}\right.
$$

The function $\Phi$ represents an estimator, as discussed in Section 5.4. The initial belief state po is given.

The corresponding DP algorithm was given in Section 5.4, and has the form

$$
J_{k}\left(p_{k}\right)=\min _{u_{k} \in U}\left[p_{k}^{\prime} g\left(u_{k}\right)+E_{z_{k+1}}\left\{J_{k+1}\left(\Phi\left(p_{k}, u_{k}, z_{k+1}\right)\right) \mid p_{k}, u_{k}\right\}\right]
$$

where $g\left(u_{k}\right)$ is the column vector with components $g\left(l, u_{k}\right), \ldots, g\left(n, u_{k}\right)$, and $p_{k}^{\prime} g\left(u_{k}\right)$, the expected stage cost, is the inner product of the vectors $P k$ and $g\left(u_{k}\right)$. The algorithm starts at stage $N$, with

$$
J_{N}\left(p_{N}\right)=p_{N}^{\prime} G
$$

where $G$ is the column vector with components $G\left(1^{\prime}\right), \ldots, G(n)$, and proceeds backwards.

We will also consider another control structure, where the information vector is

$$
\bar{I}_{k}=\left(\bar{z}_{1}, \ldots, \bar{z}_{k}, u_{0}, \ldots, u_{k-1}\right), \quad k=0, \ldots, \boldsymbol{N}-\mathbf{1},
$$

with $\bar{z}_{i}$ being some observation for each i (possibly different from Zi ), and the control constraint set at each $P k$ is a given set $U(P k)^{\prime}$ The probability distribution of $\bar{z}_{k}$ given $X k$ and $u_{k-1}$ is known, and may be different than the one of Zk. Also $\bar{U}\left(p_{k}\right)$ may be different than $U$ [in what follows, we will assume that $\bar{U}(P k)$ is a subset of $U J$.

We introduce a suboptimal policy, which at stage $k$, and starting with the current belief state $P k$, applies a control $\bar{\mu}_{k}\left(p_{k}\right)$ E $U$, based on the assumption that the future observations and control constraints will be
according to the restricted structure. More specifically, this policy chooses the control at the typical stage $k$ and state $x k$ as follows:

## Restricted Structure <br> $\square$ At stage $k$ and state $X k$, apply the

 control$$
\bar{\mu}_{k}\left(p_{k}\right)=u_{k},
$$

where

$$
\left(u_{k}, \hat{\mu}_{k+1}\left(\bar{z}_{k+1}, u_{k}\right), \ldots, \hat{\mu}_{N-1}\left(\bar{z}_{k+1}, \ldots, \bar{z}_{N-1}, u_{k}, \ldots, u_{N-2}\right)\right)
$$

is a policy that attains the optimal cost achievable from stage $k$ onward with knowledge of $P k$ and with access to the future observations $\bar{z}_{k+1}, \ldots, \bar{z}_{N-1}$ (in addition to the future controls), and subject to the constraints

$$
u_{k} \mathrm{E} U, \quad \mu_{k+1}\left(p_{k+1}\right) \in \bar{U}\left(p_{k+1}\right), \ldots, \mu_{N-1}\left(p_{N-1}\right) \in \bar{U}\left(p_{N-1}\right)
$$

Let $J k(P k)$ be the cost-to-go, starting at belief state $P k$ at stage $k$, of the restricted structure policy $\left\{\bar{\mu}_{0}, \ldots{ }^{\prime} \bar{\mu}_{N-1}\right\}$ just described. This is given by the DP algorithm
$J k(P k)=\quad\left(\bar{\mu}_{k}\left(p_{k}\right)\right)+E_{Z k+l}\{ ] \mathbf{K}+\left|\left(\Phi\left(p_{k}, \bar{\mu}_{k}\left(p_{k}\right), Z k+l\right)\right)\right|_{\left.P k, \bar{\mu}_{k}\left(p_{k}\right)\right\}}$ (6.57)
for all $P k$ and $k$, with the terminal condition $] N(P N)=p_{N}^{\prime} G$ for all $P N$.
Let us also denote by $J_{k}^{r}\left(p_{k}\right)$ the optimal cost-to-go of the restricted problem, i.e., the one where the observations and control constraints of the restricted structure are used exclusively. This is the optimal cost achievable, starting at belief state $P k$ at stage $k$, using the observations $\bar{z}_{i}, \mathrm{i}=k+1, \ldots, N \quad 1$, and subject to the constraints

$$
u_{k} \in \bar{U}\left(p_{k}\right), \quad \mu_{k+1}\left(p_{k+1}\right) \in \bar{U}\left(p_{k+1}\right), \ldots, \mu_{N-1}\left(p_{N-1}\right) \in \bar{U}\left(p_{N-1}\right)
$$

We will show, under certain assumptions to be introduced shortly, that

$$
\bar{J}_{k}\left(p_{k}\right) \leq J_{k}^{r}\left(p_{k}\right), \quad \forall p_{k}, k=O, \ldots, N \quad 1,
$$

and we will also obtain a readily computable upper bound to $] k(P k)$. To this end, for a given belief vector $P k$ and control $u_{k} \mathrm{E} \quad$ we consider three optimal costs-to-go corresponding to three different patterns of availability of information and control restriction over the remaining stages $k+1, \ldots, N-1$. We denote:
$Q k\left(P k, u_{k}\right)$ : The cost achievable from stage $k$ onward starting with $P k$, applying $u_{k}$ at stage $k$, and optimally choosing each future control $u_{i}, \mathrm{i}=k+1, \ldots, N-1$, with knowledge of $P k$, the observations $z_{k+1}, \ldots, Z i$ and the controls $u_{k}, \ldots, u_{i-1}$, and subject to the constraint $u_{i} \mathrm{E} U$.
$Q_{k}^{c}\left(p_{k},{ }^{\prime} U k\right)$ : The cost achievable from stage $k$ onward starting with $P k$, applying $u_{k}$ at stage $k$, and optimally choosing each future control $u_{i}, i \quad k+1, \ldots, N-1$, with knowledge of $P k$, the observations $, \ldots, \bar{z}_{i}$, and the controls $u_{k}, \ldots, u_{i-1}$, and subject to the constraint $U i \mathrm{E} U(P i)$. Note that this definition is equivalent to

$$
\begin{equation*}
Q_{k}^{c}\left(p_{k}, \bar{\mu}_{k}\left(p_{k}\right)\right)=\min _{U k E U} Q_{k}^{c}\left(p_{k}, u_{k}\right) \tag{6.58}
\end{equation*}
$$

where $\bar{\mu}_{k}\left(p_{k}\right)$ is the control applied by the restricted structure policy just described.
$\hat{Q}_{k}^{c}\left(p_{k}, u_{k}\right)$ : The cost achievable from stage $k$ onward starting with $P k$, applying $u_{k}$ at stage $k$, optimally choosing the control $u_{k+1}$ with knowledge of $P k$, the observatioun $Z k+l$, and the control $u_{k}$, subject to the constraint ' $U k+l \mathrm{E} U$, and optimally choosing each of the remain-
controls $u_{i}, \mathrm{i}=k+2, \ldots, N-1$, with knowledge of $P k$, the observations $Z k+l, Z k+2,{ }^{\prime \prime} \bar{z}_{i}$, and the controls $u_{k}, \ldots, u_{i-1}$, and subject to the constraints $u_{i} \mathrm{E} \underline{U}(P i)$.
Thus, the difference between $Q_{k}^{c}\left(p_{k}, u_{k}\right)$ and $Q k\left(P k, u_{k}\right)$ is due to the difference in the control constraint and the information available to the controller at all future stages $k+1, \ldots, N-1[U(P k+l), \ldots, \underline{U(P N-l)}$ versus $U$, and $\quad, \ldots, \bar{z}_{N-1}$ versus $Z k+l$,"" $Z N-l$, respectively]. The difference between $Q_{k}^{c}\left(p_{k}, u_{k}\right)$ and $\hat{Q}_{k}^{c}\left(p_{k}, u_{k}\right)$ is due to the difference in the control constraint and the information available to the controller at the single stage $k+1\left[U(P k+l)\right.$ versus $U$, and $\bar{z}_{k+1}$ versus $Z k+l$, respectively]. Our key assumptions are that

$$
\begin{equation*}
\bar{U}(P k) \subset U, \quad \forall P k, \quad k=0, \ldots, N-1, \tag{6.59}
\end{equation*}
$$

$Q k\left(P k, u_{k}\right) \leq \hat{Q}_{k}^{c}\left(p_{k}, u_{k}\right) \leq Q_{k}^{c}\left(p_{k}, u_{k}\right), \quad \forall P k, u_{k} \mathrm{E} U, k=0, \ldots, N-1$.
Roughly, this means that the control constraint $U(P k)$ is more stringent than $U$, and the observations $Z k+l, " " \bar{z}_{N-1}$ are "weaker" (no more valuable in terms of improving the cost) than the observations $Z k+l, \ldots, Z N-l$. Consequently, if Eqs. (6.59) and (6.60) hold, we may interpret a controller that uses in part the observations $\bar{z}_{k}$ and the control constraints $\underline{U(P k), ~ i n ~}$ place of $Z k$ and $U$, respectively, as "handicapped" or "restricted."

Let us denote:
$J k(P k)$ : The optimal cost-to-go of the original problem, starting at belief state $P k$ at stage $\boldsymbol{k}$. This is given by

$$
\begin{equation*}
J k(P k)=\min _{U k E U} Q k\left(P k, u_{k}\right) \tag{6.61}
\end{equation*}
$$

$J_{k}^{c}\left(p_{k}\right)$ : The optimal cost achievable, starting at belief state $P k$ at stage $k$, using the observations $\bar{z}_{i}, \mathrm{i}=k+1, \ldots, N-1$, and subject to the constraints
$u_{k} \mathrm{E} U, \quad \mu_{k+1}\left(p_{k+1}\right)$ E $\underline{U(P k+l), \ldots, \mu_{N-1}\left(p_{N-1}\right) \text { E } \bar{U}\left(p_{N-1}\right) . ~ . ~ . ~ . ~}$
This is given by

$$
\begin{equation*}
J_{k}^{c}\left(p_{k}\right)=\min _{U k E U} Q_{k}^{c}\left(p_{k}, u_{k}\right) \tag{6.62}
\end{equation*}
$$

and it is the cost that is computed when solving the optimization problem of stage $k$ in the restricted structure policy scheme. Note that we have for all $P k$,

$$
\begin{equation*}
J_{k}^{r}\left(p_{k}\right)=\min _{u k E \bar{U}(P k)} Q_{k}^{c}\left(p_{k}, u_{k}\right) \geq \min _{U k} E U \text { Q } Q_{k}^{c}\left(p_{k}, u_{k}\right) \quad J_{k}^{c}\left(p_{k}\right) \tag{6.63}
\end{equation*}
$$

where the inequality holds in view of the assumption $U(P k) \subset U$.
Our main result is the following:

Proposition 6.5.1: Under the assumptions (6.59) and (6.60), there holds

$$
\leq \bar{J}_{k}\left(p_{k}\right) \leq J_{k}^{c}\left(p_{k}\right) \leq J_{k}^{r}\left(p_{k}\right), \quad \forall P k, \quad k=0, \ldots \quad-1
$$

Proof: The inequality $J k(P k) \leq J k(P k)$ is evident, since $J k(P k)$ is the optimal cost-to-go over a class of policies that includes the restricted structure policy $\left\{\bar{\mu}_{0}, \ldots, \bar{\mu}_{N-1}\right\}$. Also the inequality $J_{k}^{c}\left(p_{k}\right) \leq J_{k}^{r}\left(p_{k}\right)$ follows from the definitions; see Eq. (6.63). We prove the remaining inequality $J k(P k) \leq J_{k}^{c}\left(p_{k}\right)$ by induction on $k$.

We have $I N(P N)=I N(P N)=0$ for allpN. Assume that for allpk +1 , we have

$$
\bar{J}_{k+1}\left(p_{k+1}\right) \leq J_{k+1}^{c}\left(p_{k+1}\right) .
$$

Then, for all $P k$,

$$
\begin{aligned}
J k(P k)= & \left.p_{k}^{\prime} g\left(\bar{\mu}_{k}\left(p_{k}\right)\right)+E_{Z k+l}\{ ] \mathrm{k}+1\left(\Phi\left(P k^{\prime} \bar{\mu}_{k}\left(p_{k}\right), Z k+l\right)\right) \mid P k, \bar{\mu}_{k}\left(p_{k}\right)\right\} \\
\leq & \left.p_{k}^{\prime} g\left(\bar{\mu}_{k}\left(p_{k}\right)\right)+E_{Z k+\prime \mid} J_{k+1}^{c}\left(\Phi\left(p k^{\prime} \bar{\mu}_{k}\left(p_{k}\right), Z k+l\right)\right) \mid P k, \bar{\mu}_{k}\left(p_{k}\right)\right\} \\
= & p_{k}^{\prime} g\left(\bar{\mu}_{k}\left(p_{k}\right)\right) \\
& +E Z_{k}+{ }_{l} \mid \min _{U k+l} E U \\
= & \left.Q_{k+1}^{c}\left(\Phi\left(p_{k}, \bar{\mu}_{k}\left(p_{k}\right), z_{k+1}\right), u_{k+1}\right) \mid p_{k}, \bar{\mu}_{k}\left(p_{k}\right)\right\} \\
\leq & Q_{k}^{c}\left(p_{k}, \bar{\mu}_{k}\left(p_{k}\right)\right) \\
= & J_{k}^{c}\left(p_{k}\right)
\end{aligned}
$$

where the first equality holds by Eq. (6.57), the first inequality holds by the induction hypothesis, the second equality holds by Eq. (6.62), the third equality holds by the definition of $\hat{Q}_{k}^{c}$, the second inequality holds by the assumption (6.60), and the last equality holds from the definition (6.58) of the restricted structure policy. The induction is complete. Q.E.D.

The main conclusion from the proposition is that the performance of the restricted structure policy $\left\{\bar{\mu}_{0}, \ldots \bar{\mu}_{N-1}\right\}$ is no worse than the performance associated with the restricted control structure. Furthermore, at each stage $k$, the value $J_{k}^{c}\left(p_{k}\right)$, which is obtained as a byproduct of the online computation of the control $\bar{\mu}_{k}\left(p_{k}\right)$, is an upper bound to the cost-to-go $J k(P k)$ of the suboptimal policy. This is consistent with Prop. 6.2.1, which shows the cost improvement property of the OLFC, and Prop. 6.3.1, which is the basis for the cost improvement property of the rollout algorithm and the stability property of MPC.

### 6.6 ADDITIONAL TOPICS IN APPROXIMATE DP

We close this chapter with a brief discussion of a few additional topics relating to approximate DP. We first address some of the discretization issues that arise when continuous state and control spaces are approximated by discrete spaces for DP computation purposes. We then describe some alternative suboptimal control approaches.

### 6.6.1 Discretization

An important practical issue is how to deal computationally with problems involving nondiscrete state and control spaces. In particular, problems with continuous state, control, or disturbance spaces must be discretized in order to execute the DP algorithm. Here each of the continuous spaces of the problem is replaced by a space with a finite number of elements, and the system equation is appropriately modified. Thus the resulting approximating problem involves a finite number of states, and a set of transition probabilities between these states. Once the discretization is done, the DP algorithm is executed to yield the optimal cost-to-go function and an optimal policy for the discrete approximating problem. The optimal cost function and/or the optimal policy for the discrete problem may then be extended to an approximate cost function or a suboptimal policy for the original continuous problem through some form of interpolation. We have already seen an example of such a process in the context of aggregation (ef. Example 6.3.13).

A prerequisite for success of this type of discretization is consistency. By this we mean that the optimal cost of the original problem should be
achieved in the limit as the discretization becomes finer and finer. Consistency is typically guaranteed if there is a "sufficient amount of continuity" in the problem; for example, if the cost-to-go functions and the optimal policy of the original problem are continuous functions of the state. This in turn can be guaranteed through appropriate continuity assumptions on the original problem data (see the references given in Section 6.7).

Continuity of the cost-to-go functions may be sufficient to guarantee consistency, even if the optimal policy is discontinuous in the state. What may happen here is that for some states there may be a large discrepancy between the optimal policy of the continuous problem and the optimal policy of its discretized version, but this discrepancy may occur over a portion of the state space that diminishes as the discretization becomes finer. As an example consider the inventory control problem of Section 4.2 with nonzero fixed cost. We obtained an optimal policy of the $(s, S)$ type

$$
\mu_{k}^{*}\left(x_{k}\right)= \begin{cases}S_{k}-x_{k} & \text { if } X k<S k \\ 0 & \text { if } X k \geq S k\end{cases}
$$

which is discontinuous at the lower threshold $S k$. The optimal policy obtained from the discretized problem may not approximate well the optimal around the point of discontinuity $S k$, but it is intuitively clear that the discrepancy has a diminishing effect on the optimal cost as the discretization becomes finer.

If the original problem is defined in continuous time, then the time must also be discretized, to obtain a discrete-time approximating problem. The issue of consistency becomes now considerably more complex, because the time discretization affects not only the system equation but also the control constraint set. In particular, the control constraint set may change considerably as we pass to the appropriate discrete-time approximation. As an example, consider the two-dimensional system

$$
\dot{x}_{1}(t)=u_{1}(t), \quad \dot{x}_{2}(t)=u_{2}(t)
$$

with the control constraint

$$
U l(t) \mathrm{E}\{--\mathrm{I}, \mathrm{I}\}, \quad U 2(t) \mathrm{E}\{-\mathrm{I}, \mathrm{I}\} .
$$

It can be seen then that the state at time $t+\Delta t$ can be anywhere within the square centered at $x(t)$ with side of length $2 \Delta t$ (note that the effect of any control in the interval $[-1,1]$ can be obtained in the continuous-time system by "chattering" between the +1 and -1 controls). Thus, given $\Delta t$, the appropriate discrete-time approximation of the control constraint set should involve a discretized version of the entire unit square, the convex hull of the control constraint set of the continuous-time problem. An example that illustrates some of the pitfalls associated with the discretization process is given in Exercise 6.10.

A general method to address the discretization issues of continuoustime/space optimal control is the aggregation/discretization approach described in Example 6.3.13. The idea is to discretize, in addition to time, the state space using some finite grid, and then to approximate the cost-togo of nongrid states by linear interpolation of the cost-to-go values of the nearby grid states. Thus, the grid states $\mathrm{x} 1, \ldots, x^{M}$ are suitably selected within the state space, and each nongrid state $x$ is expressed as

$$
x=\sum_{\mathrm{m}=1}^{\mathrm{M}} w m(x) x^{m}
$$

for some nonnegative weights $w m(x)$, which add to 1 . When this is worked out (cf. Example 6.3.13), one ends up with a stochastic optimal control problem having as states the finite number of grid states, and transition probabilities that are determined from the weights $w^{m}(x)$ above. If the original continuous-time optimal control problem has fixed terminal time, the resulting stochastic control approximation has finite horizon. If the terminal time of the original problem is free and subject to optimization, the stochastic control approximation is of the stochastic shortest path type to be discussed in Section 7.2.

Suce t.ne costs-to-go $\mathcal{J}_{k}\left(x^{i \cdot i}\right)$ vf the grid states in the ctochasuc approximating problem are computed, the cost-to-gO each nongrid state $x$ at stage $k$ is approximated by

$$
J k(x)=\sum_{\mathrm{m}=1}^{\mathrm{M}} w m(x) J k\left(x^{m}\right) .
$$

We refer to the papers by Gonzalez and Rofman [GoR85], and by Falcone [Fa187] for an account of this approach, and to the survey paper by Kushner [Kus90], and the monograph by Kushner and Dupuis [KuD92] for a detailed analysis of the associated consistency issues.

An important special case is the continuous-space shortest path problem, described in Exercise 6.10. For the corresponding stochastic shortest path problem, a finitely terminating adaptation of the Dijkstra shortest path algorithm has been developed by Tsitsiklis [Tsi95]; see Exercises 2.10 and 2.11 in Chapter 2 of Vol. II. Other related works are the papers by Bertsekas, Guerriero, and lVlusmanno [BGM95], and Polymenakos, Bertsekas, and Tsitsiklis [PBT98], which develop continuous space versions of label correcting algorithms, such as the Small-Label-First algorithm discussed in Section 2.3.1.

### 6.6.2 Other Approximation Approacl

We mention briefly three additional approaches for using approximations. In the first approach, the optimal cost-to-go functions $J k(X k)$ are approximated with functions $J k(X k, T k)$, where $T 0, T l, \ldots, T N-l$ are unknown parameter vectors, which are chosen to minimize some form of error in the

DP equations; for example by solving the problem

$$
\min _{r_{0}, \ldots, r_{N-1}} \sum_{(x k, k) E \tilde{S}} \mid \tilde{J}_{k}\left(x_{k}, T k\right)
$$

$$
-\min _{U k E U k(X k)} \underset{W k}{E}\{9 k(X k, U k, W k)+J k+1(f k(X k, U k, W k), T k+l)\}
$$

where $\tilde{S}$ is a suitably chosen subset of "representative" state-time pairs. The above minimization can be attempted using some type of gradient method. Note that there is some difficulty in doing so because the cost function of Eq. (6.64) may be nondifferentiable for some values of $r$. However, there are adaptations of gradient methods that work with nondifferentiable cost functions, and for which we refer to the specialized literature. One possibility is to replace the nondifferentiable term

$$
\min _{u_{k} \in U_{k}\left(x_{k}\right)} \underset{w_{k}}{E}\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right), r_{k+1}\right)\right\}
$$

by a smooth approximation (see Bertsekas [Ber82b], eh. 3). The approach of equations will also be discussed in more detail within an infinite horizon context (see Vol. II, Section 2.3).

In the second approach, optimal policies are directly approximated. In particular, suppose that the control space is a Euclidean space, and that we obtain, for a finite number of states $i=1, \ldots, m$, the minimizing controls

$$
\hat{\mu}_{k}\left(x^{i}\right)=\arg \min _{u_{k} \in U_{k}\left(x^{i}\right)} E\left\{g_{k}\left(x^{i}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x^{i}, u_{k}, w_{k}\right), r_{k+1}\right)\right\}
$$

We can then approximate the optimal policy $\mu_{k}\left(x_{k}\right)$ by a function of some given form

$$
\tilde{\mu}_{k}\left(x_{k}, S k\right)
$$

where $s_{k}$ is a vector of parameters obtained by solving the problem

$$
\begin{equation*}
\min _{s k} \sum_{i=1}^{m} \| \hat{\mu}_{k}\left(x^{i}\right)-\tilde{\mu}_{k}\left(x^{i}, s k\right) 111^{2} \tag{6.65}
\end{equation*}
$$

In the case of deterministic optimal control problems, we can take advantage of the equivalence between open-loop and feedback control to carry out the approximation process more efficiently. In particular, for such problems we may select a representative finite subset of initial states, and generate an optimal open-loop trajectory starting from each of these states. (Gradient-based methods can often be used for this purpose.) Each
of these trajectories yields a sequence of pairs $(x k, J k(x k))$ and a sequence of pairs ( $X k^{\prime} \mu_{k}\left(x_{k}\right)$ ), which can be used in the least-squares approximation procedures discussed above. In particular, we can use the exact values $J_{k}\left(x^{i}\right)$ and $\mu_{k}\left(x^{i}\right)$ obtained from the optimal open-loop trajectories in place of $J k\left(X^{i}\right)$ and $\hat{\mu}_{k}\left(x^{i}\right)$, respectively, in the least-squares problems of Eqs. (6.51) and (6.65).

In the third approach, sometimes called optimization in policy space, we parameterize the set of policies by a vector $S=(s o, S l, \ldots, S N-1)$ and we optimize the corresponding cost over this vector. In particular, we consider policies of the form

$$
\pi(s)=\left\{\tilde{\mu}_{0}\left(x_{0}, s_{0}\right), \ldots, \tilde{\mu}_{N-1}\left(x_{N-1}, s_{N-1}\right)\right\}
$$

where the $\tilde{\mu}_{k}(\cdot, \cdot)$ are functions of a given form. We then minimize over $s$ the expected cost

$$
E\left\{J_{\pi(s)}(x o)\right\}
$$

where $J_{\pi(s)}(x o)$ is the cost of the policy $\pi(s)$ starting from the initial state $X o$, and the expected value is taken with respect to a suitable probability distribution of $x o$. One of the difficulties associated with this approach is that optimization of $E\left\{J_{\pi(s)}(x o)\right\}$ over $s$ may be Lime-consuming, because it may require some brute force search, local search, or random search method. Sometimes, it is possible to use a gradient-based approach for optimizing $E\left\{J_{\pi(s)}(x o)\right\}$ over $s$, but this can be time-consuming as well.

In an important special case of this approach, the parameterization of the policies is indirect through a parameterization of an approximate cost-to-go function. In particular, for a given parameter vector $s \quad(s o, \ldots$, we define

$$
\tilde{\mu}_{k}\left(x_{k}, s_{k}\right)=\arg \min _{u_{k} \in U_{k}\left(x_{k}\right)} E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right), s_{k}\right)\right\}
$$

where $J_{k}+_{1}\left(^{\prime},.\right)$ is a function of a given form. For example, $J_{k}+_{l}$ may represent a linear feature-based architecture, where $S k$ is a vector of adjustable scalar weights multiplying corresponding features of states $X k+1$ (d. Section 6.3.5). Note that the policies

$$
\pi(s)=\left\{\tilde{\mu}_{0}\left(x_{0}, s o\right), \ldots, \tilde{\mu}_{N-1}\left(X N_{-1}, S N_{-1}\right)\right\}
$$

form a class of one-step lookahead policies parametrized by $s$. By optimizing over $s$ the corresponding expected cost $E\left\{J_{\pi(s)}(x o)\right\}$, we end up with a one-step lookahead policy that is optimal within this class.

## AND EXERCISES

Many schemes for suboptimal control have been discussed in this chapter, and it may be helpful to summarize them here. Most of these schemes are
based on one-step lookahead, whereby we apply at stage $k$ and state $X k$ the control $\bar{\mu}_{k}\left(x_{k}\right)$ that minimizes over $U k \mathrm{E} U k(X k)$

$$
E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}
$$

where $J_{k}+1$ is a suitable cost-to-go approximating function; in some cases, the control constraint set and/or the expected cost per stage are also approximated. The principal distinction between alternative approaches is the method for calculating $J_{k}+l_{1}$. There are several possibilities (and variations thereoff), the principal of which are:
(a) Explicit cost-to-go approximation. Here $J_{k+1}$ is computed off-line in one of a number of ways.
(1) By solving a related problem, obtained for example by aggregation or enforced decomposition, and by deriving $J_{k+1}$ from the optimal cost-to-go of that problem.
(2) By introducing a parametric approximation architecture, possibly using features. The parameters of the architecture are tuned by some form of heuristic or systematic method.
(b) Implicit cost-to-go approximation. Here the values of $\boldsymbol{J}_{k+1}$ at the states $f k(x k, U k, W k)$ are computed on-line as needed, by using an open-loop computation (optimal or suboptimal/heuristic, with or without a rolling horizon). We focused on a few possibilities, all which were interpreted under the unifying framework of restricted structure policies in Section 6.5.3:
(1) Open-loop-feedback control, where an optimal open-loop computation is used, starting from the state $X k$ (in the case of perfect state information) or the conditional probability distribution of the state (in the case of imperfect state information).
(2) Rollout, where the cost-to-go of a suboptimal/heuristic policy is used as $J_{k}+1$. This cost is computed as necessary by on-line simulation (which in some variants may be approximate and/or use a rolling horizon).
(3) Model predictive control, where an optimal control computation is used in conjunction with a rolling horizon. This computation is deterministic, possibly based on a simplification of the original problem via certainty equivalence, but there is also a minimax variant that implicitly involves reachability of target tube computations.
A few important variations of the preceding schemes should be mentioned. The first is the use of multistep lookahead, which aims to improve the performance of one-step lookahead, at the expense of increased on-line
computation. The second is the use of certainty equivalence, which simplifies the off-line and on-line computations by replacing the current and future unknown disturbances $w k, \ldots, W N-l$ with nominal values. A third variation, which applies to problems of imperfect state information, is to use one of the preceding schemes with the unknown state $x k$ replaced by some estimate.

While the idea of one-step lookahead is old, it has gained a lot of credibility recently, thanks to extensive research on approximate dynamic programming and wide acceptance of model predictive control in practical applications. With experience and research, the relative merits of different approaches have been clarified to some extent, and it is now understood that some schemes possess desirable theoretical performance guarantees, while others do not. In particular, in this chapter, we have discussed qualitative and/or quantitative performance guarantees for open-loop feedback control (cf. Prop. 6.2.1 and the discussion of Section 6.5.3), rollout amples 6.3.1 and 6.3.2), and model predictive control (the stability guarantee discussed in Section 6.5.2). The performance guarantee for certainty equivalent control (cf. Prop. 6.3.2 and Example 6.3.3) is weaker, and indeed for some stochastic problems, certainty equivalent control may be outperformed by open-loop control (see Exercise 6.2) For addItional theoretical analysis on performance bounds. Witsenhausen [Wit69], [Wit70]. Despite the recent progress in theory and practical experience, the methodology for perfoImance analysis of suboptimal control schemes is not very satisfactory at present, and the validation of a suboptimal policy by simulation is often essential in practice. This is true of all approaches described in this chapter, including ones that are not based on one-step lookahead, such as approximation in policy space (cf. Section 6.6.2).

Excellent surveys of adaptive control, which contain many other references, are given by Astrom [Ast83] and Kumar [Kum85]. Self-tuning regulators received wide attention following the paper by Astrom and Wittenmark [AsW73]. For textbook treatments of adaptive control, see Astrom and Wittenmark [AsW94], Goodwin and Sin [GoS84], Hernandez-Lerma [Her89], Ioannou and Sun [IoS96], Krstic, Kanellakopoulos, and Kokotovic [KKK95], Kumar and Varaiya [KuV86], Sastry, Bodson, and Bartram [SBB89], and Slotine and Li [SLL91].

Open-loop feedback control was suggested by Dreyfus [Dre65]. Its superiority over open-loop control (cf. Prop. 6.2.1) was established by the author in the context of minimax control [Ber72b]. A generalization of this result is given by White and Harrington [WhH80]. The POLFC was proposed in Bertsekas [Ber76].

Stochastic programming problems have been discussed in detail in the literature the texts by Birge and Louveaux [BiL97], Kall and Wallace [KaW94], and Prekopa [Pre95]). The connections between stochastic programming and stochastic optimal control have been highlighted by Varaiya and Wets [VaW89].

There is a long history of limited lookahead approximations in specificapplication contexts. The performance bounds for limited lookahead policies, given in Section 6.3.1 and Exercises 6.11-6.15 are new.

The main idea of rollout algorithms, obtaining an improved policy starting from some other suboptimal policy using a one-time policy improvement, has appeared in several DP application contexts. In the context of game-playing computer programs, it has been proposed by Abramson [Abr90] and by Tesauro [TeG96]. The name "rollout" was coined by Tesauro in specific reference to rolling the dice in the game of backgammon. In Tesauro's proposal, a given backgammon position is evaluated by "rolling out" many games starting from that position, using a simulator, and the results are averaged to provide a "score" for the position. The internet contains a lot of material on computer backgammon and the use of rollout, in some cases in conjunction with multistep lookahead and cost-to-go approximation.

The application of rollout algorithms to discrete optimization problems has its origin in the neuro-dynamic programming work of the author and J. Tsitsiklis [BeT96], and has been further formalized by Bertsekas, Tsitsiklis, and Wu [BTW97], Bertsekas [Ber97], and Bertsekas and Castanon [BeC99]. The analysis of the breakthrough problem (Example 6.4.2) is based on unpublished joint work of the author with D. Castanon and J. Tsitsiklis. An analysis of the optimal policy and some suboptimal policies for this nrobiem is given by Pearl [Pea84]. A discussion of rollout algorithms as applied to network optimization problems may be found in the author's network optimization book [Ber98a]. The technique for variance reduction in the calculation of Q -factor differences (Section 6.4.2) is from Bertsekas [Ber97].

For work on rollout algorithms, see Christodouleas [Chr97], Secomandi [SecOO], [Sec01], [Sec03], Bertsimas and Demir [BeD02], Ferris and Voelker [FeV02], [FeV04], McGovern, Moss, and Barto [MMB02], Savagaonkar, Givan, and Chong [SGC02], Bertsimas and Popescu [BeP03], Guerriero and Mancini [GuM03], Tu and Pattipati [TuP03], Wu, Chong, and Givan [WCG03], Chang, Givan, and Chong [CGC04], IVleloni, Pacciarelli, and Pranzo [MPP04], and Yan, Diaconis, Rusmevichientong, and Van Roy [YDR05]. These works discuss a broad variety of applications and case studies, and generally report positive computational experience.

The model predictive control approach has become popular in a variety of control system design contexts, and particularly in chemical process control, where meeting explicit control and state constraints is an important practical issue. Over time, there has been increasing awareness of the connection with the problem of reachability of target tubes, set-membership descriptions of uncertainty, and minimax control (see the discussion of Section 4.6). The stability analysis given here is based on the work of Keerthi and Gilbert [KeG88]. For extensive surveys of the field, see Morari and Lee [MoL99], and Mayne et. al. [IVIRROO], who give many references. For
related textbooks, see Camacho and Bordons [CaB04], and Maciejowski [Mac02]. The connection with rollout algorithms and one-time policy iteration reported in Section 6.5.2 is new. The material of Section 6.5.3 on the unifying suboptimal control framework based on restricted structure policies is also new.

The computational requirements for solving stochastic optimal control problems are discussed from the point of view of computational complexity in the survey by Blondel and Tsitsiklis [BITOO], who give several additional references; see also Rust [Rus97]. For consistency analyses of various discretization and approximation procedures for discrete-time stochastic optimal control problems, see Bertsekas [Ber75], [Ber76a], Chow and Tsitsiklis [ChT89], [ChT91], Fox [Fox71], and Whitt [Whi78], [Whi79]. A discretization method that takes advantage of the special structure of finite-state imperfect state information problems was first given by Lovejoy [Lov91a]; see also the survey [Lov91b]. For more recent work, based on the aggregation/discretization approach described in Example 6.3.13, see Yu and Bertsekas [YuB04]. The discretization issues of continuous-time/space optimal control problems have been the subject of considerable research; see Gonzalez and Rofman [GoR85], Falcone [Fa187], Kushner [Kus90], and Kushner and Dupuis [KuD92], which give additional sources.

There have been important algorithmic developments for certain types of continuous space shortest path problems. A finitely terminating adaptation of the label setting (Dijkstra) method has been developed by Tsitsiklis [Tsi95]. This method was rediscovered later, under the name "fast marching method," by Sethian [Set99a], [Set99b], who discusses several other related methods and many applications, as well as by Helmsen et al. [HPC96]. Efficient analogs of label correcting algorithms for continuous space shortest path problems were developed by Bertsekas, Guerriero, and Musmanno [BGIVI95], and Polymenakos, Bertsekas, and Tsitsiklis [PBT98].

## EXERCISES

6.1

Consider a problem with perfect state information involving the $n$-dimensional linear system of Section 4.1:
and a cost function of the form

$$
\underset{\substack{w_{k} \\ k=0, \ldots, N-1}}{E}\left\{g_{N}\left(c^{\prime} x_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(u_{k}\right)\right\}
$$

where $c E R n$ is a given vector. Show that the DP algorithm for this problem can be carried out over a one-dimensional state space.

## 6.2

Argue that for a one-stage problem, the optimal open-loop controller and the OLFC are both optimal. Construct an example where the CEC may be strictly suboptimal. Also work out the following two-stage example, due to [ThW•66], which involves the following two-dimensional linear system with scalar control and disturbance:

$$
x_{k+1}=x_{k}+b u_{k}+d w_{k}, \quad k=0,1
$$

where $b=(1,0) /$ and $d=(1 / 2, \sqrt{2} / 2)^{\prime}$. The initial state is $X o=0$. The controls $U o$ and Ul are unconstrained. The disturbances $W o$ and wi are independent random variables and each takes the values 1 and -1 with equal probability $1 / 2$. Perfect state information prevails. The cost is

$$
\underset{w_{0}, w_{1}}{E}\left\{\left\|x_{2}\right\|\right\}
$$

where $11 \cdot 11$ denotes the usual Euclidean norm. Show that the CEC with nominal values $\bar{w}_{0}=\bar{w}_{1}=0$ has worse performance than the optimal open-loop controller. In particular, show that the optimal open-loop cost and the optimal closed-loop cost are both $\sqrt{3} / 2$, but the cost corresponding to the CEC is 1 .

## 6.3

Consider a two-stage problem with perfect state information involving the scalar system

$$
X o=1, \quad \mathrm{x}_{1}=X o+u o+w o, \quad x_{2}=f\left(x_{1}, u_{1}\right)
$$

The control constraints are $U o, u_{1} \mathrm{E}\{\mathrm{O},-\mathrm{I}\}$. The random variable Wo takes the values 1 and -1 with equal probability $1 / 2$. The function $f$ is defined by

$$
\begin{gathered}
f(l, O)=f(l,-1)=f(-1,0)=f(-1,-1)=0.5 \\
f(2,0)=0, \quad f(2,-1)=2, \quad f(O,-1)=0.6, \quad f(O, O)=2
\end{gathered}
$$

The cost function is
(a) Show that one possible OLFC for this problem is

$$
\bar{\mu}_{0}\left(x_{0}\right)=-1, \quad \bar{\mu}_{1}\left(x_{1}\right)= \begin{cases}0 & \text { if } x_{1}= \pm 1,2 \\ -1 & \text { if } x_{1}=0\end{cases}
$$

and the resulting cost is 0.5 .
(b) Show that one possible CEC for this problem is

$$
\bar{\mu}_{0}\left(x_{0}\right)=0, \quad \bar{\mu}_{1}\left(x_{1}\right)= \begin{cases}0 & \text { if } x_{\prime}= \pm 1,2 \\ -1 & \text { if } x_{\prime}=0\end{cases}
$$

and the resulting cost is 0.3 . Show also that this CEC is an optimal feedback controller.

## 6.4

Consider the system and cost function of Exercise 6.3 but with the difference that

$$
!(\mathrm{O},-1)=0 .
$$

(a) Show that the controller of part (a) of Exercise 6.3 is both an OLFC and CRC, and that the corresponding cost is 0.5
(b) Assume that the control constraint set fin the first stage is $\{0\}$ rather than $\{0,-I\}$. Show that the controller of part (b) uí Bxercise 6.3 is both an OLFC and a CEC, and that the corresponding cost is U . Note: This problem illustrates a pathology that occurs generically in suboptimal control; that is, if the control constraint set is restricted, the performance of a suboptimal scheme may be improved. To see this, consider a problem and a suboptimal control scheme that is not optimal for the problem. Let $\pi^{*}=\left\{\mu_{0}^{*}, \ldots, \mu_{N-1}^{*}\right\}$ be an optimal policy. Restrict the control constraint set so that only the optimal control $\mu_{k}^{*}\left(x_{k}\right)$ is allowed at state $X k$. Then the cost attained by the suboptimal control scheme will be improved.
6.5

Consider the ARMAX model

$$
Y k+l+a Y k=b U k+E \mathrm{k}+1+C E k
$$

where the parameters $a, b$, and c are unknown. The controller hypothesizes a model of the form

$$
Y k+l+a Y k=U k+\mathrm{Ek}+1
$$

and uses at each $k$ the minimum variance/certainty equivalent control

$$
u_{k}^{*}=\hat{a}_{k} y_{k},
$$

where $\hat{a}_{k}$ is the least-squares estimate of $a$ obtained as

$$
\hat{a}_{k} \quad \arg \min _{a} \sum_{n=l}^{k}\left(y_{n}+a y_{n-1}-u_{n-1}^{*}\right)^{2} .
$$

Write a computer program to test the hypothesis that the sequence $\left\{\hat{a}_{k}\right\}$ converges to the optimal value, which is $(\mathrm{c} a) / b$. Experiment with values lal $<1$ and $|a|>1$

## 6.6 (Semilinear Systems)

Consider the basic problem for semilinear systems (Exercise 1.13 in Chapter 1) Show that the OLFC, and the CEC, with nominal values of the disturbances equal to their expected values, are optimal for this problem.

## 6.7

Consider the production control problem of Example 6.3 .8 for the case where there is only one part type $(n=1)$, and assume that the cost per stage is a convex function 9 with $\lim _{|x| \rightarrow 1} g(x)=00$.
(a) Show that the cost-to-go function $\operatorname{Jk}\left(X k, \alpha_{k}\right)$ is convex as a function of $X k$ for each value of $\alpha_{k}$.
(b) Show that for each $k$ and $\alpha_{k}$, there is a target value $\bar{x}_{k+1}$ such that for each $X k$ it is optimal to choose the control $U k$ E $U_{k}\left(\alpha_{k}\right)$ that brings $X k+l=$ $X k+U k-d k$ as close as possible to $\bar{x}_{k+1}$.
6.8

Provide a careful argument showing that searching a chess position with and without $\alpha-\beta$ pruning will give the same result.

## 6.9

In a version of the game of Nim, two players start with a stack of five pennies and take turns removing one, two, or three pennies from the stack. The player who removes the last penny loses. Construct the game tree and verify that the second player to move will win with optimal play

### 6.10 Lontmuous Space Shortest Path Problems)

Consider the two-dimensional system

$$
\dot{x}_{1}(t)=u_{1}(t), \quad \dot{x}_{2}(t)=u_{2}(t)
$$

with the control constraint $\|\cup(t)\|=1$. We want to find a state trajectory that starts at a given point $x(O)$, ends at another given point $x(T)$, and minimizes

$$
\int_{0}^{T} r(x(t)) d t
$$

The function $r(\cdot)$ is nonnegative and continuous, and the final time $T$ is subject to optimization. Suppose we discretize the plane with a mesh of size $\Delta$ that passes through $x(O)$ and $x(T)$, and we introduce a shortest path problem of going from
$X(O)$ to $x(T)$ using moves of the following type: from each mesh point $\bar{x}=\left(\bar{x}_{1}, \bar{x}_{2}\right)$ we can go to each of the mesh points $\left(\bar{x}_{1}+\Delta, \bar{x}_{2}\right),\left(\bar{x}_{1}-\Delta, \bar{x}_{2}\right),\left(\bar{x}_{1}, \bar{x}_{2}+\Delta\right)$, and
$\bar{x}_{2}-\Delta$ ), at a cost $r(\bar{x}) \Delta$. Show by example that this is a bad discretization of the original problem in the sense that the shortest distance need not approach the optimal cost of the original problem as $\Delta \rightarrow 0$.

### 6.11 Iscretization of Convex Problems)

Consider a problem with state space 5 , where 5 is a convex subset of $\Re^{n}$. Suppose that $\hat{S}=\{Y l, \ldots, Y M\}$ is a finite subset of 5 such that 5 is the convex hull of $\hat{S}$, and consider a one-step lookahead policy based on approximate cost-to-go functions $\mathbf{J}_{0}, \mathbf{J}_{1}, \ldots, \mathbf{J}_{N}$ defined as follows:

$$
\tilde{J}_{N}(x)=g_{N}(x), \quad \forall x \in 5
$$

and for $k=1, \ldots, N-1$,
$J k(x)=\min \left\{\sum_{i=1}^{M} \lambda_{i} \hat{J}_{k}\left(y_{i}\right) \mid \sum_{i=1}^{M} \lambda_{i} y_{i}=x, \sum_{i=1}^{M} \lambda_{i}=1, \lambda_{i} \geq 0, \mathrm{i}=1, \ldots, M\right\}^{\prime}$
where $\mathrm{J}_{k}(x)$ is defined by

$$
J k(x)=\min _{U E U k(x)} E\left\{g_{k}\left(x, u, w_{k}\right)+J k+l(f k(X, U, W k))\right\}
$$

Thus $\mathbf{J} k$ is obtained from $\mathbf{J} k+l$ as a "grid-based" convex piecewise linear approximation to $\mathbf{J}_{k}$ based on the $M$ values

$$
\mathrm{J}_{k}(Y l), \ldots, \mathrm{J}_{k}(Y l V I)
$$

Assume that the cost functions $g k$ and the system functions $f k$ are such that the function $\mathrm{j} k$ is real-valued and convex over 5 whenever $\mathbf{J}_{\mathrm{k}+1}$ is real-valued and convex over 5. Use Prop. 6.3.1 to show that the cost-to-go functions $J_{k}$ corresponding to the one-step lookahead policy satisfy for all $x \mathrm{E} 5$

$$
J k(x) \leq J k(x) \leq J k(x), \quad k=0,1, \ldots, \boldsymbol{N}-1 .
$$

### 6.12 (One-Step Lookahead with Cost per Stage and Constraint

 Approximations)Consider a one-step lookahead policy as in Section 6.3, where $\mathbf{J} k(x k)$ is chosen to be the optimal cost-to-go of a different problem where the costs-per-stage and control constraint sets are $\tilde{g}_{k}(X k, U k, W k)$ and $U k(X k)$, respectively, [rather than

$$
U k, W k) \text { and } U k(X k)] \text {. Assume that for all } k, X k, U k, W k \text {, we have }
$$

$$
g k(X k, U k, W k) \leq g k(X k, U k, W k), \quad U k(X k) \mathrm{c} \bar{U}_{k}\left(x_{k}\right)
$$

Use Prop. 6.3.1 to show that the costs-to-go $\overline{\boldsymbol{J}}_{k}$ of the one-step lookahead policy satisfy

$$
J k(X k) \leq J k(X k)
$$

for all $X k$ and $k$. Extend this result for the case where $\tilde{g}_{k}$ satisfies instead

$$
g k(X k, U k, W k) \leq g k(X k, U k, W k)+\delta_{k}
$$

where $\delta_{k}$ are some scalars that depend only on $k$.

### 6.13 (One-Step Lookahead/Rollout for Shortest Paths)

Consider a graph with nodes $1, \quad, N$, and the problem of finding a shortest path from each of the nodes $1, \quad, N-1$ to node $N$ with respect to a given set of arc lengths aij. We assume that all cycles have positive length. Let $F(i)$, $\mathrm{i}=1, \ldots, N$, be some given scalars with $F(N)=0$, and denote

$$
\begin{equation*}
\hat{F}(i)=\min _{j \in J_{i}}\left[a_{i j}+F(j)\right], \quad \mathrm{i}=1, \ldots, \boldsymbol{N}-\mathbf{1} \tag{6.66}
\end{equation*}
$$

where for each $i, J_{i}$ is a nonempty subset of the set of neighbor nodes $\{j \mid$ $(i, j)$ is an arc $\}$.
(a) Assume that $\hat{F}(i) \leq F(i)$ for all i $=1, \ldots, N-1$. Let $j(i)$ attain the minimum in Eq. (6.66) and consider the graph consisting of the $N \mathbf{- 1}$ arcs (i, $\mathrm{j}(\mathrm{i})), \mathrm{i}=1, \ldots, N-1$. Show that this graph contains no cycles and for each $\mathrm{i}=1, \ldots, N-1$, it contains a unique path Pi starting at i and ending at $N$. Show that the length of $P i$ is less or equal to $\hat{F}(i)$.
(b) Would the conclusion of part (a) hold if the cycles of the original graph are assumed to have nonnegative (rather than positive) length?
(c) Let $F(i)$ be the length of some given path from node i to node $N$ with $F(N)=0$, and assume that for the first arc of $\bar{P}_{i}$, say $(i, j i)$, we have ji E $J_{i}$. Assume further that

$$
F(i) \geq a_{i j_{i}}+F\left(j_{i}\right)
$$

[this is satisfied with equality if $\bar{P}_{i}$ consists of arc (i, $j i$ ) followed by path $\bar{P}_{j}$, which is true if the paths $\bar{P}_{i}$ form a tree rooted at the destination $N$; for example if the paths were obtained by solving some related shortest path problem]. Show that $\hat{F}(i) \leq F(i)$ for all $\mathrm{i}=1, \ldots, \mathrm{~N}-1$ -
(d) Assume that $\mathrm{J}_{\mathrm{i}}=\{\mathrm{j} \mid(i, j)$ is an arc $\}$. Let Pi be the paths obtained as in part (a) when the scalars $F(i)$ are generated as in part (c). Interpret Pi as the result of a rollout algorithm that uses an appropriate heuristic, and show that for each i, the length of $P i$ is less or equal to the length of $\bar{P}_{i}$.
(e) Assume that $\mathbf{J}_{\mathbf{i}}=\{\mathrm{j} \mid(i, j)$ is an arc $\}$. Let us view the scalars $F(i)$ as the node labels of a label correcting method. This method starts with labels $F(i)=$ oo for all $\mathrm{i} \neq \mathrm{N}$ and $F(N)=0$, and at each step sets

$$
F(i)=\left\{_{\left\{j /(i, j) \min _{l S \text { an arc }\}}[a i j+\mathrm{F}(\mathrm{j}) \mathrm{J} .4 .\right.}\right.
$$

for some node $\mathrm{i} \neq N$ for which the above equality is violated (the method terminates if this equality holds for all $\mathrm{i} \neq N$ ). Show that in the course of this method, the labels $F(i)$ satisfy the assumptions of part (c) at all times (at or before termination) for which $F(i)<$ oo for all i.

### 6.14 Performance Bounds for Two-Step Lookahead Policies)

Consider a two-step lookahead policy as in Section 6.3, and assume that for all $x_{k}$ and $k$, we have

$$
\hat{J}_{k}\left(x_{k}\right) \leq \tilde{J}_{k}\left(x_{k}\right)
$$

where $j N \quad g N$ and for $k=0, \ldots, N-1$,

$$
\hat{J}_{k}\left(x_{k}\right)=\min _{u_{k} \in \bar{U}_{k}\left(x_{k}\right)} E\left\{g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right\}
$$

Consider the cost-to-go functions] $k$ corresponding to the two-step lookahead policy that uses $\mathrm{J} k$ and $\bar{U} k(x k)$. Show that for all $x k$ and $k$, we have

$$
\bar{J}_{k}\left(x_{k}\right) \leq J_{k}^{+}\left(x_{k}\right) \leq \hat{J}_{k}\left(x_{k}\right) \leq \tilde{J}_{k}\left(x_{k}\right)
$$

where $J_{k}^{++}$is the function obtained by two DP iterations starting from $J_{k+2}$ :

$$
J_{k}^{+}\left(x_{k}\right)=\min _{u_{k} \in \bar{U}_{k}\left(x_{k}\right)} E\left\{g _ { k } \left(x_{k}, u_{k} \quad \quad\left(\tilde { N } _ { k } \left(\Psi_{k}, u_{k},\right.\right.\right.\right.
$$

0.15 (Pollomt Algorithms with Errors)

Consider the graph search problem of Section 6.4.1 and let $\mathcal{H}$ be a sequentially improving base heuristic. Suppose that we generate a path (iI, ..., $i_{\bar{m}}$ ) according to

$$
\mathrm{im}+1=\arg \min _{j E N\left(i_{m}\right)} H(j), \quad \mathrm{m}=1, \ldots, m-1,
$$

where differs from the cost $H(j)$ of the base heuristic by the error

$$
e(j)=\hat{H}(j)-H(j)
$$

(a) Assuming that $\mathrm{I} e(j) \mathrm{I} \leq \epsilon$ for all $j$, show that the cost of the generated path is less than or equal to $\left.H\left(i_{1}\right)+\underline{2(m} 1\right) \epsilon$. Hint: Use the relation

$$
H\left(i_{m+1}\right)-\epsilon \leq \hat{H}\left(i_{m+1}\right)=\min _{j E N\left(i_{m}\right)} \hat{H}(j) \leq \min _{j E N\left(i_{m}\right)} H(j)+\epsilon \leq H\left(i_{m}\right)+\epsilon
$$

(b) Modify the estimate of part (a) for the case where we have $0 \leq e(j) \leq \epsilon$ for all $j$, and for the case where we have $-\epsilon \leq e(j) \leq 0$ for all $j$.
(c) Consider the case where $\mathcal{H}$ is optimal so that $H(j)=J^{*}(j)$, and derive a bound on the difference between the cost of the generated path and the optimal cost starting from $\mathrm{i}_{1}$.

### 6.16 (Breakthrough Problem with Random Heuristic)

Consider the breakthrough problem of Example 6.4 .2 with the difference that instead of the greedy heuristic, we use the random heuristic, which at a given node selects one of the two outgoing arcs with equal probability. Denote by

$$
D_{k}=p^{k}
$$

the probability of success of the random heuristic in a graph of $k$ stages, and by $R k$ the probability of success of the corresponding rollout algorithm. Show that for all $k$

$$
R_{k}=p(2-p) R_{k-1}+p^{2} D_{k-1}\left(1-R_{k-1}\right)
$$

and that

$$
\frac{R_{k}}{D_{k}}=(2-p) \quad+p(l-R k-I)
$$

Conclude that $R k / D k$ increases exponentially with $k$.
6.17

ConsIder the breakthrough problem of Example 6.4.2 with the difference that there are three outgoing arcs from each node instead of two. Each arc is free with probability $p$, independently of other arcs. Derive an equation for the ratio $R k / G k$, where $G k$ is the probability of success of the greedy heuristic for a $k$ stage problem, and $R_{k}$ is the probability of success of the corresponding rollout algorithm. Verify that the results of Example 6.4 .2 still hold in a qualitative sense, and that $R k / G_{k}$ increases linearly with $k$.

### 6.18 (Breakthrough Problem with a

Horizon

Consider the breakthrough problem of Example 6.4.2 and consider a rolling horizon-type of rollout algorithm that uses a greedy base heuristic with l-step lookahead. This is the same algorithm as the one described in Example 6.4.2, except that if both outgoing arcs of the current node at stage $k$ are free, the rollout algorithm considers the two end nodes of these arcs, and from each of them it runs the greedy algorithm for $\min \{l, N-k-I\}$ steps. Consider a Markov chain with $I+1$ states, where states $i=0, \ldots, I-1$ correspond to the path generated by the greedy algorithm being blocked after i arcs. State $I$ corresponds to the path generated by the greedy algorithm being unblocked after $/$ arcs.
(a) Derive the transition probabilities for this Markov chain so that it models the operation of the rollout algorithm.
(b) Use computer simulation to generate the probability of a breakthrough, and to demonstrate that for large values of $N$, the optimal value of $l$ is roughly constant and much smaller than $N$ (this can also be justified analytically, by using properties of Markov chains).

### 6.19 (Rollout for Constrained

Consider the deterministic constrained DP problem involving the system

$$
X k+l=!k(X k, U k),
$$

where we want to minimize the cost function

$$
g_{N}^{1}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}^{1}\left(x_{k}, U k\right)
$$

subject to the constraints

$$
g_{N}^{m}\left(x_{N}\right)+\sum_{k=0}^{N-1} g_{k}^{m}\left(x_{k}, u_{k}\right) \leq b^{m}, \quad m=2, \ldots, M
$$

cf. Section 2.3.4. We assume that each state $X k$ takes values in a finite set and each controi $\boldsymbol{u}_{\boldsymbol{k}} \quad$ values in a finite constraint set $U k(X k)$ that depends on $X k$. We describe an extension of the rolivue involving some base heuristic, which is feasible in the sense that when started from tue given initial state $X o$, it produces a state/control trajectory that satisfies the constranits of the problem.

Consider a rollout algorithm, which at stage $k$, maintains a partial state/control trajectory

$$
T k=(x o, U o, X l, \cdots, U k-I, X k)
$$

that starts at the given initial state $X o$, and is such that $x_{i+1}=!\mathbf{i}(X i, U i)$ and $U i \in U_{i}\left(x_{i}\right)$ for all $\mathrm{i}=0,1, \ldots, k-1$. For such a trajectory, let $\operatorname{Cm}(X k)$ be the corresponding values of constraint functions

$$
C m(X k)=\sum_{\mathrm{i}=\mathrm{O}}^{k-l} g_{i}^{m}\left(x_{i}, U i\right), \quad m=2, \ldots, M
$$

For each $U k \in U_{k}\left(x_{k}\right)$, let $X k+l=!k(X k, U k)$ be the next state, and let $J(X k+l)$ and $6^{m}(X k+l)$ be the cost-to-go and values of constraint functions of the base heuristic starting from $X k+l$.

The algorithm starts with the partial trajectory To that consists of just the initial state $x o$. For each $k=0, \ldots, N-1$, and given the current trajectory $T k$, it forms the subset of controls $U k \in U k(X k)$ that together with the corresponding states $X k+l=!k(X k, U k)$ satisfy

$$
C m(X k)+g_{k}^{m}\left(x_{k}, U k\right)+6^{m}(X k+l) \leq b^{m}, \quad \mathrm{~m}=2, \ldots, M .
$$

The algorithm selects from this set a control $U k$ and corresponding state $X k+l$ such that

$$
g_{k}^{1}\left(x_{k}, U k\right)+\tilde{J}\left(x_{k+1}\right)
$$

is minimum, and then it forms the trajectory $T k+l$ by adding $(U k, X k+l)$ to $T k$. Formulate analogs of the assumptions of sequential consistency and sequential improvement of Section 6.4.1, under which the algorithm is guaranteed to generate a feasible state/control trajectory that has no greater cost than the cost associated with the base heuristic. Note: For a description and analysis of a generalized version of this algorithm, see the author's report "Rollout Algorithms for Constrained Dynamic Programming," LIDS Report 2646, MIT, April 2005.

### 6.20 (Rollout for Minimax Problems)

Consider the minimax DP problem, as described in Section 1.6, and a one-step lookahead policy based on lookahead functions $J_{I}, \ldots, J N$, with $I N=g N$. This is the policy obtained by minimizing at state $X k$ the expression

$$
\max _{w_{k} \in W_{k}\left(x_{k}, u_{k}\right)}\left[g_{k}\left(x_{k}, u_{k}, w_{k}\right)+\tilde{J}_{k+1}\left(f_{k}\left(x_{k}, u_{k}, w_{k}\right)\right)\right]
$$

over $U k \in U k(X k)$.
(a) State and prove analogs of Props. 6.3.1 and 6.3.2.
(b) Consider a rollout algorithm where $\mathbf{J}_{k}$ are the cost-to-go functions corresponding to some base heuristic. Show that the cost-to-go functions] $k$ of the rollout algorithm satisfy $\bar{J} k(x k) \leq \mathbf{J} k(x k)$ for all $X k$ and $k$.

### 6.21 (MPC with Disturbances)

Consider the MPC framework of Section 6.5.2, including disturbances with setmembership description. Let $\bar{\mu}$ be the policy obtained from MPC.
(a) Use the constrained controllability assumption to show that $\vec{\mu}$ attains reachability of the target tube $\{X, X, \ldots\}$ in the sense that

$$
f(x, \bar{\mu}(x), w) \mathbf{E X}, \quad \text { for all } x \in \mathbf{X} \text { and } w \in \quad \bar{\mu}(x))
$$

(b) Consider any sequence $\{x o, U o, X l, U I, \ldots\}$ generated by MPC [i.e., Xo EX, $X o \notin T, U k=\bar{\mu}\left(x_{k}\right), X k+l=!(X k, U k, W k)$, and $\left.W k \in W(X k, U k)\right]$. Show that

$$
\left(x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}\right) \leq \hat{J}\left(x_{0}\right)<\infty
$$

$k=0$
where $K_{T}$ is the smallest integer $k$ such that $X k \in T$ (with $K T=$ oo if $x_{k} \notin T$ for all $k$ ), and $\mathbf{J}(x)$ is the optimal cost starting at state $x \in X$ of the m-stage minimax control problem solved by MPC. Hint: Argue as in the case where there are no disturbances. Consider an optimal control problem that is similar to the one solved at each stage by MPC, but has one stage less. In particular, given $x \in X$ with $x \notin T$, consider the minimax control problem of finding a policy $\hat{\mu}_{0}, \hat{\mu}_{1}, \ldots, \hat{\mu}_{m-2}$ that minimizes

$$
\max _{\substack{w_{i} \in W\left(x_{i}, \mu\left(x_{i}\right)\right), i=0,1, \ldots, m-2}} \sum_{i=0}^{m-2} g\left(x_{i}, \hat{\mu}_{i}\left(x_{i}\right)\right)
$$

subject to the system equation constraints

$$
x_{i+1}=f\left(x_{i}, \hat{\mu}_{i}\left(x_{i}\right), w_{i}\right), \quad \mathrm{i}=0,1, \ldots, \mathrm{~m}-2
$$

the control and state constraints

$$
x_{i} \mathrm{EX}, \quad \hat{\mu}_{i}\left(x_{i}\right) \in U\left(x_{i}\right), \quad \mathrm{i}=0,1, \ldots, m-2,
$$

and the terminal state constraint

$$
x_{i} \in T, \quad \text { for some i } \mathrm{E}[1, m-1] \text {. }
$$

These constraints must be satisfied for all disturbance sequences with

$$
w_{i} \in W\left(x_{i}, \hat{\mu}_{i}\left(x_{i}\right)\right), \quad 0,1, \ldots, m-2
$$

Let $\mathbf{J}(x o)$ be the corresponding optimal value, and define $\mathbf{J}(x o)=0$ for Xo E $T$, and $J(x o)=00$ for all $X o \notin T$ for which the problem has no feasible solution. Show that the control $\bar{\mu}(x)$ applied by MPC at a state $x \in X$ with $x \notin T$, minimizes over $u \mathrm{E} U(x)$

$$
\max _{w E W(x, u)}\left[X l Q x+u^{\prime} R u+J(J(x, u, W))\right]
$$

and use the fact $J(x) \leq J(x)$ to show that for all $X \in X$ with $x \notin T$, we have

$$
\max _{w E W(x, u)}\left[x^{\prime} Q x+\bar{\mu}(x)^{\prime} R \bar{\mu}(x)+\hat{J}(f(x, \bar{\mu}(x), w))\right] \leq J(x)
$$

Conclude that for all $k$ such that $X k E X$ with $X k \notin T$, we have

$$
x_{k}^{\prime} Q x_{k}+u_{k}^{\prime} R u_{k}+\hat{J}\left(x_{k+1}\right) \leq \hat{J}\left(x_{k}\right)
$$

where $J(X k+l)={ }^{\mathrm{O}}$ if $x_{k+1} \mathrm{E} T$. Add over $k=0,1, \ldots, K T-\mathrm{I}$.
(c) Show that under MPC, the state $X k$ of the system must belong to $T$ for all sufficiently large $k$, provided that minxEx, $x_{x \notin T} x^{\prime} Q x>0$. Use Example 6.5.3 to show the need for this assumption.
(d) Interpret the policy $\bar{\mu}$ produced by MPC as a rollout policy with an appropriate base heuristic. Hint: View $\bar{\mu}$ as a one-step lookahead policy with one-step lookahead approximation function equal to $J$, defined in the hint to part (b).

## Introduction

## Infinite Horizon Problems

## Contents

7.1. An Overview ..... p. 4027.2. Stochastic Shortest Path Problemsp. 405
7.3. Discounted Problemsp. 417
7.4. Average Cost per Stage Problems ..... p. 421p. 435
7.6. Notes, Sources, and Exercises.

In this chapter, we provide an introduction to infinite horizon problems. These problems differ from those considered so far in two respects:
(a) The number of stages is infinite.
(b) The system is stationary, Le., the system equation, the cost per stage, and the random disturbance statistics do not change from one stage to the next.
The assumption of an infinite number of stages is never satisfied in practice, but is a reasonable approximation for problems involving a finite but very large number of stages. The assumption of stationarity is often satisfied in practice, and in other cases it approximates well a situation where the system parameters vary slowly with time.

Infinite horizon problems are interesting because their analysis is elegant and insightful, and the implementation of optimal policies is often

> ivi exampie,
are typically stationary, i.e., the optimal rule for choosing controls does not change from one stage to the next.

On the other hand, infinite horizon problems generally require more sophisticated analysis than their finite horizon counterparts, because of the need to analyze limiting behavior as the horizon tends to infinity. This analysis is often nontrivial and at times reveals surprising possibilities. Our treatment will be limited to finite-state problems. A far more detailed development, together with applications from a variety of fields can be found in Vol. II of this work.

### 7.1 AN OVERVIE"V

There are four principal classes of infinite horizon problems. In the first three classes, we try to minimize the total cost over an infinite number of stages, given by

$$
J_{\pi}\left(x_{0}\right)=\lim _{N \rightarrow \infty} \underset{\substack{w_{k} \\ k=0,1, \ldots}}{E}\left\{\sum_{k=0}^{N-1} \alpha^{k} g\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right\}
$$

$J_{\pi}\left(x_{0}\right)$ denotes the cost associated with an initial state $X o$ and a policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots\right\}$, and $\alpha$ is a positive scalar with $0<\alpha \leq 1$, called the discount factor. The meaning of $\alpha<1$ is that future costs matter to us less than the same costs incurred at the present time. As an example, think of kth period dollars depreciated to initial period dollars by a factor of $(1 \quad$ where $r$ is a rate of interest; here $\alpha=1 /(1+r)$. An important concern in total cost problems is that the limit in the definition of $J_{\pi}\left(x_{0}\right)$ be finite. In the first two of the following classes of problems, this is guaranteed
through various assumptions on the problem structure and the discount factor. In the third class, the analysis is adjusted to deal with infinite cost for some of the policies. In the fourth class, this sum need not be finite for any policy, and for this reason, the cost is appropriately redefined.
(a) Stochastic shortest path problems. Here, $\alpha=1$ but there is a special cost-free termination state; once the system reaches that state it remains there at no further cost. We will assume a problem structure such that termination is inevitable (this assumption will be relaxed somewhat in Chapter 2 of Vol. II). Thus the horizon is in effect finite, but its length is random and may be affected by the policy being used. These problems will be considered in the next section and their analysis will provide the foundation for the analysis of the other types of problems considered in this chapter.
(b) Discounted problems with bounded cost per stage. $\quad \alpha<1$ and the absolute cost per stage $\lg (x, u, w) I$ is bounded from above by some constant Nl; this makes the cost $\mathrm{J}_{\pi}(x o)$ well defined because it is the infinite sum of a sequence of numbers that are bounded in absolute value by the decreasing geometric progression $\left\{\alpha^{k} M\right\}$. We will consider these problems in Section 7.3.
(c) D'iscounted and undiscounted problems with unbounded cost per stage. Here the discount factor $\alpha$ mayor may not be less than 1 , and the cost per stage may be unbounded. These problems require a complicated analysis because the possibility of infinite cost for some of the policies is explicitly dealt with. We will not consider these problems here; see Chapter 3 of Vol. II.
(d) Average cost per problems. Minimization ofthe total cost $J_{\pi}\left(x_{0}\right)$ makes sense only if $J_{\pi}\left(x_{0}\right)$ is finite for at least some admissible policies $\pi$ and some initial states $x o$. Frequently, however, it turns out that $J_{\pi}\left(x_{0}\right)=00$ for every policy $\pi$ and initial state Xo (think of the case where $\alpha=1$, and the cost for every state and control is positive). It turns out that in many such problems the average cost per stage, defined by

$$
\left.\lim _{N \rightarrow \infty} \frac{1}{N} \underset{\substack{W k \\ \mathrm{k}=\mathrm{O}, 1, \ldots}}{E} \sum_{k=0}^{N-1} g\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right\}
$$

is well defined and finite. We will consider some of these problems in Section 7.4.

## A Preview of Infinite Horizon Results

There are several analytical and computational issues regarding infinite horizon problems. Many of these revolve around the relation between the
optimal cost-to-go function $J^{*}$ of the infinite horizon problem and the optimal cost-to-go functions of the corresponding N -stage problems. In particular, consider the case $\alpha=1$ and let $I N(x)$ denote the optimal cost of the problem involving N stages, initial state $x$, cost per stage $g(x, u, w)$, and zero terminal cost. The optimal N -stage cost is generated after N iterations of the DP algorithm

$$
\begin{equation*}
J k+l(X)=\min _{u E U(x)} E\{g(x, u, w)+J k(J(X, U, w))\}, \quad k=O, l, \ldots \tag{7.1}
\end{equation*}
$$

starting from the initial condition $\operatorname{Jo}(x)=0$ for all $x$ (note here that we have reversed the time indexing to suit our purposes). Since the infinite horizon cost of a given policy is, by definition, the limit of the corresponding N -stage costs as $\mathrm{N} \rightarrow 00$, it is natural to speculate that:
(1) The optimal infinite horizon cost is the limit of the corresponding N -stage optimal costs as $\mathrm{N} \rightarrow 00$; that is,

$$
\begin{equation*}
J^{*}(x)=\lim _{1_{1}^{11+\infty 0}} I N(X) \tag{7.2}
\end{equation*}
$$

for all states $x$. This relation is extremely valuable computationally and analytically, and, fortunately, it typically holds. In particular, it holds for the models of the next tivu sections [categories (a) and (b) above]. However, there are some unusual exceptions for problems in category above, and this illustrates that infinite horizon problems should be approached with some care. This issue is discussed in more detail in Vol. II.
(2) The following limiting form of the DP algorithm should hold for all states $x$,

$$
J^{*}(x)=\min _{u E U(x)} E\left\{g(x, u, w)+J^{*}(J(x, u, w))\right\}
$$

as suggested by Eqs. (7.1) and (7.2). This is not really an algorithm, but rather a system of equations (one equation per state), which has as solution the costs-to-go of all the states. It can also be viewed as a functional equation for the cost-to-go function $J^{*}$, and it is called Bellman's equation. Fortunately again, an appropriate form of this equation holds for every type of infinite horizon problem of interest.
(3) If $\mu(x)$ attains the minimum in the right-hand side of Bellman's equation for each $x$, then the policy $\{\mu \mu \ldots\}$ should be optimal. This is true for most infinite horizon problems of interest and in particular, for all the models discussed in this chapter.
Most of the analysis of infinite horizon problems revolves around the above three issues and also around the issue of efficient computation of $J^{*}$ and an optimal policy. In the next three sections we will provide a discussion of these issues for some of the simpler infinite horizon problems, all of which involve a finite state space.

## Total Cost Problem Formulation

Throughout this chapter we assume a controlled finite-state discrete-time dynamic system whereby, at state $i$, the use of a control $u$ specifies the transition probability Pij $(u)$ to the next state $j$. Here the state $i$ is an element of a finite state space, and the control $u$ is constrained to take values in a given finite constraint set $U(i)$, which may depend on the current state i. As discussed in Section 1.1, the underlying system equation is

$$
x_{k+1}=w_{k},
$$

where $w_{k}$ is the disturbance. We will generally suppress $w_{k}$ from the cost to simplify notation. Thus we will assume a kth stage cost $g\left(x_{k}, u_{k}\right)$ for using control $u_{k}$ at state $x_{k}$. This amounts to averaging the cost per stage over all successor states in our calculations, which makes no essential difference in the subsequent analysis. Thus, if $\tilde{g}(\mathrm{i}, u, j)$ is the cost of using $u$ at state i and moving to state $j$, we use as cost per stage the expected cost $g(i, u)$ given by

$$
g(i, u)=\sum_{j} p_{i j}(u) \tilde{g}(i, u, j)
$$

The total expected cost associated with an initial state i and a policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots\right\}$ is

$$
J_{\pi}(i)=\lim _{N \rightarrow \infty} E\left\{\sum_{k=0}^{N-1} \alpha^{k} g\left(x_{k}, \mu_{k}\left(x_{k}\right)\right) \mid x_{0}=i\right\},
$$

where $\alpha$ is a discount factor with $0<\alpha \leq 1$. In the following two sections, we will impose assumptions that guarantee the existence of the above limit. The optimal cost from state i , that is, the minimum of $J_{\pi}(i)$ over all admissible $\pi$, is denoted by $J^{*}$ (i). A stationary policy is an admissible policy of the form $\pi=\{\mu \mu \cdot .$.$\} , and its corresponding cost function is denoted$ by $J_{\mu}(i)$. For brevity, we refer to $\{\mu \mu \ldots\}$ as the stationary policy $\mu$. We say that $\mu$ is optimal if

$$
J_{\mu}(i)=J^{*}(i)=\min _{\pi} J_{\pi}(i), \quad \text { for all states i. }
$$

### 7.2 STOCHASTIC SHORTEST PATH PROBLEMS

Here, we assume that there is no discounting ( $\alpha=1$ ), and to make the cost meaningful, we assume that there is a special cost-free termination state $t$. Once the system reaches that state, it remains there at no further cost, i.e., $\operatorname{Ptt}(u)=1$ and $g(t, u)=0$ for all $u \mathrm{E} U(t)$. We denote by $1, \ldots, n$ the states other than the termination state $t$.

We are interested in problems where reaching the termination state is inevitable, at least under an optimal policy. Thus, the essence of the problem is to reach the termination state with minimum expected cost. We call this problem the stochastic shortest path problem. The deterministic shortest path problem is obtained as the special case where for each statecontrol pair ( $i, u$ ), the transition probability $P_{i j}(u)$ is equal to 1 for a unique state $j$ that depends on (i, $u$ ). The reader may also verify that the finite horizon problem of Chapter 1 can be obtained as a special case by viewing as state the pair ( $X k, k$ ) (see also Section 3.6 of Vol. II).

Certain conditions are required to guarantee that, at least under an optimal policy, termination occurs with probability 1 . We will make the following assumption that guarantees eventual termination under all policies:

Assumption 7.2.1: There exists an integer $m$ such that regardless of the policy used and the initial state, there is positive probability that the termination state will be reached after no more that m stages; that is, for all admissible policies it we have

$$
\begin{equation*}
\rho_{\pi}=\max _{i=1} P\left\{x_{m} \neq\left. t\right|_{X_{o}}=\mathrm{i}, \text { it }\right\}<1 . \tag{7.3}
\end{equation*}
$$

We note, however, that the results to be presented are valid under more general circumstances.t Furthermore, it can be shown that if there cive on integer $m$ with the property of Assumption 7.2.1, then there also exists an integer less or equal to $n$ property 7.12). Thus, we can always use $m=n$ in Assumption 7.2.1, if no smaller value of $m$ is

[^3]known. Let
$$
p=\max _{\pi} \rho_{\pi}
$$

Note that $\rho_{\pi}$ depends only on the first m components of the policy $\pi$. Furthermore, since the number of controls available at each state is finite, the number of distinct m -stage policies is also finite. It follows that there can be only a finite number of distinct values of $\rho_{\pi}$ so that

$$
p<1 .
$$

We therefore have for any 'it and any initial state i

$$
\begin{aligned}
\left.P_{\{\times 2 m} \neq t \mid X Q=\mathrm{i}, \pi\right\}= & P\left\{X 2 m \neq t \mid x_{m} \neq t, x Q=\mathrm{i}, \pi\right\} \\
& . P\left\{x_{m} \neq t \mid X Q=i, i T\right\} \\
\leq & p 2 .
\end{aligned}
$$

More generally, for each admissible policy it, the probability of not reaching the termination state after km stages diminishes like $p k$ regardless of the initial state, that is,

$$
\begin{equation*}
P\left\{x_{k m} \neq t \mid x_{0}=i, \pi\right\} \leq \rho^{k}, \quad \mathrm{i}=1, \ldots, n o \tag{7.4}
\end{equation*}
$$

Thue ine limit defining the associated total cost vector $J_{\pi}$ exists and is nnite, since the expected cost incurred in the $m$ periods between km and $(k+I) m \quad 1$ is bounded in absolute value by

$$
m p k \max _{\substack{i=1, \ldots, n \\ u \in U(i)}} \operatorname{Ig}(i, u) \mid
$$

In particular, we have

$$
\begin{equation*}
\left|J_{\pi}(i)\right| \leq \operatorname{mon}_{k=Q} m p k \max _{\substack{i=1, \ldots(n) \\ u \in \cup(i)}} \operatorname{Ig}(i, u)\left|=\frac{m}{1-P} \max _{\substack{i=1, \ldots, n \\ u \in U(i)}} \operatorname{Ig}(i, u)\right| \cdot \tag{7.5}
\end{equation*}
$$

The results of the following proposition are basic and are typical of many infinite horizon problems. The key idea of the proof is that the "tail" of the cost series,

$$
\underset{k=m K}{E\left\{g\left(X k^{\prime} \mu_{k}\left(x_{k}\right)\right)\right\}}
$$

vanishes as $K$ increases to oo, since the probability that $x m K \neq t$ decreases like $p K$ [ef. Eq. (7.4)].

Proposition 7.2.1 Under Assumption 7.2.1, the following hold for the stochastic shortest path problem:
(a) Given any initial conditions $J o(l), \ldots, J o(n)$, the sequence $J k(i)$ generated by the iteration

$$
\begin{equation*}
J_{k+1}(i)=\min _{u \in U(i)}\left[g(i, u)+\sum_{j=1}^{n} p_{i j}(u) J_{k}(j)\right], \quad \mathrm{i}=1, \ldots, n \tag{7.6}
\end{equation*}
$$

converges to the optimal cost $J^{*}(i)$ for each i. [Note that, by reversing the time index this iteration can be viewed as the DP algorithm for a finite horizon problem with terminal cost function equal to $J o$. In fact, $J k(i)$ is the optimal cost starting from state i of a k-stage problem with cost per stage given by 9 and terminal cost at the end of the $k$ stages given by Jo.]
(b) The optimal costs $J^{*}(l), \ldots, J^{*}(n)$ satisfy Bellman's equation,

$$
\begin{equation*}
\underset{\substack{i, \ldots \\ u E U(i)}}{ } 19(i, U)+\left.\sum_{j=1}^{n} p_{i j}(u) J^{*}(j)\right|^{\prime}, \quad i=1, \ldots, n \tag{7.7}
\end{equation*}
$$

and in fact they are the unique solution of this equation.
(c) For any stationary policy $\mu$, the costs $J_{\mu}(1), \ldots, J_{\mu}(n)$ are the unique solution of the equation

$$
J_{\mu}(i)=g(i, \mu(i))+{ }_{j=l}^{n} p_{i j}(\mu(i)) J_{\mu}(j), \quad \mathrm{i}=1, \ldots, n o
$$

Furthermore, given any initial conditions $J o(l), \ldots, J o(n)$, the sequence $J k(i)$ generated by the DP iteration

$$
J k+l(\mathrm{i}) \quad g(\mathrm{i}, \mu(\mathrm{i}))+\sum_{j=l}^{n} p_{i j}(\mu(\mathrm{i})) J k(j), \quad \mathrm{i}=1, \ldots, n,
$$

converges to the cost $J_{\mu}(i)$ for each i.
(d) A stationary policy $\mu$ is optimal if and only if for every state i , $\mu(i)$ attains the minimum in Bellman's equation (7.7).

Proof: (a) For every positive integer $K$, initial state $x o$, and policy $\pi=$ $\left\{\mu_{0}, \mu_{1}, \ldots\right\}$, we break down the cost $J_{\pi}\left(x_{0}\right)$ into the portions incurred over
the first $m K$ stages and over the remaining stages

$$
\begin{aligned}
J_{\pi}\left(x_{0}\right)= & \lim _{N \rightarrow \infty} E\left\{\sum_{k=0}^{N-1} g\left(x_{k}, \mu_{k}\left(x_{k}\right)\right)\right\} \\
= & E\left\{\sum_{k=0}^{m K-1} g\left(x_{k}, \mu_{k}\left(x_{k}\right)\right)\right\} \\
& +\lim _{N \rightarrow \infty} E\left\{\sum_{k=m K}^{N-1} g\left(x_{k}, \mu_{k}\left(x_{k}\right)\right)\right\} .
\end{aligned}
$$

Let $M$ denote the following upper bound on the cost of an m-stage cycle, assuming termination does not occur during the cycle,

$$
N l=\mathrm{m} \max _{\substack{i=1, \ldots, n \\ u E U(i)}} 19(\mathrm{i}, u) l .
$$

The expected cost during the $K t h \mathrm{~m}$-stage cycle [stages $K m$ to ( $K+1$ )m-1] is upper bounded by $N 1 l_{p K}$ [ef. Eqs. (7.4) and (7.5)], so that

$$
\left|\lim _{N \rightarrow \infty} E\left\{\sum_{k=m K}^{N-1} g\left(x_{k}, \mu_{k}\left(x_{k}\right)\right)\right\}\right| \leq M \sum_{k=K}^{\infty} \rho^{k}=\frac{\rho^{K} M}{1-\rho} .
$$

Also, denoting $J o(t)=0$, let us view $J o$ as a terminal cost function and bound its expected value under $\pi$ after $m K$ stages. We have

$$
\begin{aligned}
!E\{J O(X m K)\}! & =\mid \sum_{i=1}^{n} P\left(x_{m K}=\mathrm{i} \mid \mathrm{XO}, 7 \mathrm{r}\right) \mathrm{JO}(\mathrm{i}) 1 \\
& \leq\left(\sum_{i=1}^{n} P\left(X m K=\mathrm{i} \mid \mathrm{I}_{\mathrm{Xo}}, \pi\right)\right) \max _{i=1, \ldots, n} \mid J_{0}(i) \mathrm{I} \\
& \leq p \max _{i=1, \ldots, n} \mathrm{jJo}(\mathrm{i}) \mathrm{I},
\end{aligned}
$$

since the probability that $X m K \neq t$ is less or equal to $p K$ for any policy. Combining the preceding relations, we obtain

$$
\begin{align*}
-\rho^{K} \max _{i=1, \ldots, n}\left|J_{0}(i)\right| & +J_{\pi}\left(x_{0}\right)-\frac{\rho^{K} M}{1-\rho} \\
& \leq E\left\{J_{0}\left(x_{m K}\right)+\sum_{k=0}^{m K-1} g\left(x_{k}, \mu_{k}\left(x_{k}\right)\right)\right\}  \tag{7.8}\\
& \leq \rho^{K} \max _{i=1, \ldots, n}\left|J_{0}(i)\right|+J_{\pi}\left(x_{0}\right)+\frac{\rho^{K} M}{1-\rho}
\end{align*}
$$

Note that the expected value in the middle term of the above inequalities is the $m K$-stage cost of policy $\pi$ starting from state $X 0$, with a terminal cost $J_{0}\left(x_{m K}\right)$; the minimum of this cost over all $\pi$ is equal to the value $\operatorname{JmK}(X O)$, which is generated by the DP recursion (7.6) after $m K$ iterations. Thus, by taking the minimum over $\pi$ in Eq. (7.8), we obtain for all $X_{0}$ and $K$,

$$
\begin{align*}
\_p K_{i=1, \ldots, n} \max \left|J_{0}(i)\right|+J^{*}(x o) & -\frac{p K M}{1} \frac{M}{p} \\
& \leq J m K(X O)  \tag{7.9}\\
& \leq p K_{i=1, \ldots, n}\left|J_{0}(i)\right|+J^{*}(x o)+\frac{p K M}{1-\rho}
\end{align*}
$$

and by taking the limit as $K \rightarrow 00$, we obtain $\lim _{K \rightarrow \infty} J m K(X O)=J^{*}(x o)$ for all $x$. Since

$$
\left|J_{m K+q}\left(x_{0}\right)-J_{m K}\left(x_{0}\right)\right| \leq \rho^{K} M, \quad q=1, \ldots, m
$$

we see that $\operatorname{limK}_{\rightarrow \infty} \operatorname{JmK}+q(x 0)$ is the same for all $q=1, \ldots, m$, so that we have $\lim _{k \rightarrow \infty} J k(X O)=J^{*}(x o)$.
(b) By taking the limit as $k \rightarrow$ oo in the DP iteration (7.6) and using the result of part (a), we see that $J^{*}(1), \ldots, J^{*}(n)$ satisfy Bellman's equation. To show uniqueness, observe that if $\mathbf{J}(1), \ldots, J(n)$ satisfy Bellman's equation, then the DP iteration (7.6) starting from $\mathrm{J}(1), \ldots, J(n)$ just replicates $\mathbf{J}(1), \ldots, J(n)$. It follows from the convergence result of part (a) that $J(i) \quad J^{*}(i)$ for all i.
(c) Given the stationary policy $\mu$, we can consider a modified stochastic shortest path problem, which is the same as the original except that the control constraint set contains only one element for each state $i$, the control $\mu(i)$; that is, the control constraint set is $\tilde{U}(i)=\{\mu(i)\}$ instead of $U(i)$. From part (b) we then obtain that $J_{\mu}(1), \ldots, J_{\mu}(n)$ solve uniquely Bellman's equation for this modified problem, that is,

$$
J_{\mu}(i)=g(i, \mu(i))+{ }_{j=1}^{n} p_{i j}(\mu(i)) J_{\mu}(j), \quad i=1, \ldots, n,
$$

and from part (a) it follows that the corresponding DP iteration converges to $J_{\mu}(i)$.
(d) We have that $\mu(i)$ attains the minimum in Eq. (7.7) if and only if we have

$$
J^{*}(i)=\min _{u E U(i)}[9 U, u)+\underset{j=1}{n}
$$

$n$

Part (c) and the above equation imply that $J_{\mu}(i)=J^{*}(i)$ for all i. Conversely, if $J_{\mu}(\mathrm{i})=J^{*}(\mathrm{i})$ for all i, parts (b) and (c) imply the above equation. Q.E.D.

Example 7.2.1 (Minimizing Expected Time to Termination)
The case where

$$
g(i, u)=1, \quad i=1, \ldots, n, \quad u \in U(i)
$$

corresponds to a problem where the objective is to terminate as fast as possible on the average, while the corresponding optimal cost $J^{\star}(\mathrm{i})$ is the minimum expected time to termination starting from state i. Under our assumptions, the costs $J^{*}(i)$ uniquely solve Bellman's equation, which has the form

$$
J^{*}(i)=\min _{u E U(i)}\left[1+t P i j(U) J_{j=1}^{*}(j)\right], \quad i=1, \ldots, n o
$$

In the special case where there is only one control at each state, $J^{*}(\mathrm{i})$ represents the mean first passage time from i to $t$ (see Appendix D). These times, denoted $m i$, are the unique solution of the equations

$$
\mathrm{mi}=1+\sum_{j=1} p_{i j} m_{j}, \quad \mathrm{i}=1, \ldots, n o
$$

## Example 7.2.2

A spider and a fly move along a straight line at times $k=0,1, \ldots$ The initial positions of the fly and the spider are integer. At each time period, the fly moves one unit to the left with probability $P$, one unit to the right with probability $P$, and stays where it is with probability $1-2 p$. The spider, knows the position of the fly at the beginning of each period, and will always move one unit towards the fly if its distance from the fly is more that one unit. If the spider is one unit away from the fly, it will either move one unit towards the fly or stay where it is. If the spider and the fly land in the same position at the end of a period, then the spider captures the fly and the process terminates. The spider's objective is to capture the fly in minimum expected time.

We view as state the distance between spider and fly. Then the problem can be formulated as a stochastic shortest path problem with states $0,1, \ldots, n$, where $n$ is the initial distance. State 0 is the termination state where the spider captures the fly. Let us denote $\operatorname{Plj}(M)$ and $P l j(M)$ the transition probabilities from state 1 to state $j$ if the spider moves and does not move, respectively, and let us denote by $P i j$ the transition probabilities from a state $i \geq 2$. We have

$$
\text { Pii } \quad P, \quad P i(i-l)=1-2 p, \quad P i(i-2)=P, \quad \mathrm{i} \geq 2,
$$

$$
\operatorname{Pll}(I N I)=2 p, \quad P l O(I v I)=1-2 p,
$$

$$
p_{12}(M)=p, \quad \underline{P l l(N I)}=1-2 p, \quad \underline{P l O(N I)}=p,
$$

with all other transition probabilities being 0 .
For states $\mathrm{i} \geq 2$, Bellman's equation is written as

$$
\begin{equation*}
J^{*}(i)=1+p J^{*}(i)+(1-2 p) J^{*}(i-1)+p J^{*}(i-2), \quad i \geq 2 \tag{7.10}
\end{equation*}
$$

where $J^{*}(0)=0$ by definition. The only state where the spider has a choice is when it is one unit away from the fly, and for that state Bellman's equation is given by

$$
\begin{equation*}
\left.J^{*}(l)=1+\min [2 p]^{*}(1), p J^{*}(2)+(1-2 p) J^{*}(1)\right] \tag{7.11}
\end{equation*}
$$

where the first and the second expression within the bracket above are associated with the spider moving and not respectively. By writing Eq. (7.10) for $\mathrm{i} \quad$ 2. we obtain

$$
\left.J^{*}(2)=1+p\right]^{*}(2)+(1-2 p) J^{*}(1)
$$

from which

$$
\begin{equation*}
J^{*}(2)=\frac{1}{1}_{p}+\left(1-\frac{2 p) J^{*}(1)}{1-P}\right. \tag{7.12}
\end{equation*}
$$

Substituting this expression in Eq. (7.11), we obtain

$$
J^{*}(1)=1 \quad \min \left[2 p r(I), \quad \frac{1}{=} \frac{p}{=}+\frac{p(1-2 p) J^{*}(1)}{1-p}+(1-2 p) J^{*}(1)\right]
$$

or equivalently,

$$
J^{*}(1)=1+\min \left[2 p J^{*}(1), \quad-p+\left(1 \quad \frac{2 p)}{1-p} \frac{J^{*}}{p}(1)\right]\right.
$$

To solve the above equation, we consider the two cases where the first expression within the bracket is larger and is smaller than the second expression. Thus we solve for $J^{*}(l)$ in the two cases where

$$
\begin{gather*}
J^{*}(1)=1+2 p J^{*}(1)  \tag{7.13}\\
2 p]^{*}(1) \leq-p+\frac{(1-p) J^{*}}{\square-1} \frac{11}{\square} \tag{7.14}
\end{gather*}
$$

and

$$
\begin{aligned}
& J^{*}(l)=1+{ }_{\overline{1}}-P_{-}+\left(1 \frac{2 p) J^{*}(1)}{1-p}\right. \\
& 2 p J^{*}(1) \geq p^{+}+\begin{array}{c}
(1-2 p) J^{*}(1) \\
1-p
\end{array}
\end{aligned}
$$

The solution ofEq. (7.13) is seen to be $J^{*}(l)=1 /(1-2 p)$, and by substitution in Eq. (7.14), we find that this solution is valid when

$$
1-2 p<\frac{p}{1-p}+\frac{1}{1-p}
$$

or equivalently (after some calculation), $p \leq 1 / 3$. Thus for $p \leq 1 / 3$, it is optimal for the spider to move when it is one unit away from the fly.

Similarly, the solution of Eq. (7.15) is seen to be $J^{*}(1) \quad l / p$, and by substitution in Eq. (7.14), we find that this solution is valid when

$$
2>\frac{p}{1-p}+\frac{1-2 p}{p(1-p)}
$$

or equivalently (after some calculation), $p \geq 1 / 3$. Thus, for $p \geq$ it is optimal for the spider not to move when it is one unit away from the fly.

The minimal expected number of steps for capture when the spider is one unit away from the fly was calculated earlier to be

$$
\begin{array}{cl}
J^{*}(l)=\{1 /(1-2 P) & \text { if } p \leq 1 / 3 \\
l / p & \text { if } p \geq 1 / 3
\end{array}
$$

Given the value of $J^{*}(l)$, we can calculate from Eq. (7.12) the minimal expected number of steps for capture when two units away, $J *(2)$, and we can then obtain the remaining values $J^{*}(\mathrm{i}), \mathrm{i}=3, \ldots, n$, from Eq. (7.10).

## Value Iteration and Error Bounds

The DP iteration

$$
\begin{equation*}
J_{k+1}(i)=\min _{u \in U(i)}\left[g(i, u)+\sum_{j=1}^{n} p_{i j}(u) J_{k}(j)\right], \quad \mathrm{i}=1, \ldots, n \tag{7.16}
\end{equation*}
$$

is called value iteration and is a principal method for calculating the optimal cost function $J^{*}$. Generally, value iteration requires an infinite number of iterations, although there are important special cases where it terminates finitely (see Vol. Section 2.2). Note that from (7.9) we obtain that the error

$$
I J m K(i)-J^{*}(i)!
$$

is bounded by a constant multiple of $p K$.
The value iteration algorithm can sometimes be strengthened with the use of some error bounds. In particular, it can be shown (see Exercise 7.13) that for all $k$ and $j$, we have

$$
J k+l(j)+\left(N^{*}(j)-1\right) c_{k} \leq J^{*}(j) \leq J_{\mu^{k}}(j) \leq J k+l(j)+(N k(j)-1) \bar{c}_{k}
$$

(7.17)
where $\mu^{k}$ is such that $\mu^{k}(\mathbf{i})$ attains the minimum in the kth value iteration (7.16) for all $i$, and
$N^{*}(j)$ : The average number of stages to reach $t$ starting from $j$ and using some optimal stationary policy,
$N k(j)$ : The average number of stages to reach $t$ starting from $j$ and using the stationary policy $\mu^{k}$,

$$
\underline{c}_{k}=\min _{i=1, \ldots, n}\left[J_{k+1}(i)-J_{k}(i)\right], \quad \bar{c}_{k}=\max _{i=1, \ldots, n}[J k+1(i)-J k(i)]
$$

Unfortunately, the values $N^{*}(j)$ and $N k(j)$ are easily computed or approximated only in the presence of special problem structure (see for example the next section). Despite this fact, the bounds (7.17) often provide a useful guideline for stopping the value iteration algorithm while being assured that $J k$ approximates $J^{*}$ with sufficient accuracy.

## Iteration

There is an alternative to value iteration, which always terminates finitely. This algorithm is called policy iteration and operates as follows: we start with a stationary policy $\mu^{0}$, and we generate a sequence of new policies $\mu^{1}, \mu^{2}, \ldots$ Given the policy $\mu^{k}$, we perform a policy evaluation step, that computes $J_{\mu^{k}}(\mathrm{i}), \mathrm{i}=1, \ldots, n$, as the solution of the (linear) system of equations

$$
\begin{equation*}
J(i)=g\left(i, \mu^{k}(i)\right)+\sum_{j=1}^{n} p_{i j}\left(\mu^{k}(i)\right) J(j), \quad \mathrm{i}=1, \ldots, n, \tag{7.18}
\end{equation*}
$$

in the $n$ unknowns $J(I), \ldots, J(n)$ [d. Prop. 7.2.1(c)]. We then perform a policy improvement step, which computes a new policy $\mu^{k+1}$ as

$$
\begin{equation*}
\mu^{k+1}(i)=\arg \min _{u \in U(i)}\left[g(i, u)+\sum_{j=1}^{n} p_{i j}(u) J_{\mu^{k}}(j)\right], \quad \mathrm{i}=1, \ldots, n o \tag{7.19}
\end{equation*}
$$

The process is repeated with $\mu^{k+1}$ used in place of $\mu^{k}$, unless we have $J_{\mu^{k+1}}(i)=J_{\mu^{k}}(\mathrm{i})$ for all i, in which case the algorithm terminates with the policy $\mu^{k}$. The following proposition establishes the validity of policy iteration.
$\underline{\mathbf{P r}_{~}^{x}}$ - $\quad$ 7.2.2 Under Assumption 7.2.1, the policy iteration algorithm for the stochastic shortest path problem generates an improving sequence of policies [that is, $J_{\mu^{k+1}}(i) \leq J_{\mu^{k}}(i)$ for all i and $k$ ] and terminates with an optimal policy.

Proof: For any $k$, consider the sequence generated by the recursion

$$
I N+1(i)=g\left(i, \mu^{k+1}(i)\right)+\sum_{j=I}^{n} p_{i j}\left(\mu^{k+1}(i)\right) J_{N}(j), \quad \mathrm{i}=1, \ldots, n
$$

where $N=0,1, \ldots$, and

$$
J o(i)=J_{\mu^{k}}(i), \quad i=1, \ldots, n o
$$

From Eqs. (7.18) and (7.19), we have

$$
\begin{aligned}
J o(i) & =g\left(i, \mu^{k}(i)\right)+\sum_{j=1}^{n} p_{i j}\left(\mu^{k}(i)\right) J_{0}(j) \\
& \geq g\left(i, \mu^{k+1}(i)\right)+\sum_{j=1}^{n} p_{i j}\left(\mu^{k+1}(i)\right) J_{0}(j) \\
& =J_{1}(i)
\end{aligned}
$$

for all i. By using the above inequality we obtain (compare with the monotonicity property of DP, Exercise 1.23 in Chapter 1)

$$
\begin{aligned}
J 1(i) & =g\left(i, \mu^{k+1}(i)\right)+{ }_{j=l}^{n} p_{i j}\left(\mu^{k+1}(i)\right) J_{0}(j) \\
& \geq g\left(i, \mu^{k+1}(i)\right)+\sum_{j=I}^{n} p_{i j}\left(\mu^{k+1}(i)\right) J_{1}(j) \\
& =J_{2}(i)
\end{aligned}
$$

for all $i$, and by continuing similarly we have

$$
\begin{equation*}
J_{0}(i) \geq J_{1}(i) \geq \cdots \geq J_{N}(i) \geq J_{N+1}(i) \geq \cdots, \quad \mathrm{i}=1, \ldots, n \tag{7.20}
\end{equation*}
$$

Since by Prop. 7.2.1(c), IN(i) $\rightarrow J_{\mu^{k+1}}\left(\right.$ i), we obtain $J o(i) \geq J_{\mu^{k+1}}($ i) or

$$
J_{\mu^{k}}(i) \geq J_{\mu^{k+1}}(i), \quad \mathrm{i}=1, \ldots, n, \quad k=0,1, \ldots
$$

Thus the sequence of generated policies is improving, and since the number of stationary policies is finite, we must after a finite number of iterations, say $k+1$, obtain $J_{\mu^{k}}$ (i) $=J_{\mu^{k+1}}$ (i) for all i. Then we will have equality holding throughout in Eq. (7.20), which means that

$$
J_{\mu^{k}}(i)=\min _{u \in U(i)}\left[g(i, u)+\sum_{j=1}^{n} p_{i j}(u) J_{\mu^{k}}(j)\right], \quad i=1, \ldots, n o
$$

Thus the costs $J_{\mu^{k}}(1), \ldots, J_{\mu^{k}}(n)$ solve Bellman's equation, and by Prop. 7.2.1(b), it follows that $J_{\mu^{k}}(i)=J^{*}(i)$ and that $\mu^{k}$ is optimal. Q.E.D.

The linear system of equations (7.18) of the policy evaluation step can be solved by standard methods such as Gaussian elimination, but when the number of states is large, this is cumbersome and time-consuming. A typically more efficient alternative is to approximate the policy evaluation step with a few value iterations aimed at solving the corresponding system (7.18). One can show that the policy iteration method that uses such approximate policy evaluation yields in the limit the optimal costs and an optimal stationary policy, even if we evaluate each policy using an arbitrary positive number of value iterations (see Vol. II, Section 1.3).

Another possibility for approximating the policy evaluation step is to use simulation, and this is a key idea in the rollout algorithm, discussed in Section 6.4. Simulation also plays an important role in the neuro-dynamic programming methodology, discussed in Ch. 2 of Vol. II. In particular, when the number of states is large, one can try to approximate the cost-to-go function $J_{\mu^{k}}$ by simulating a large number of trajectories under the policy $\mu^{k}$, and perform some form of least fit nf using an approximation architecture Section 6.3.4). These are a number of variations of this Which are discussed in more detail in Vol. II and in the research monograph by Bertsekas and Tsitsiklis [BeT96].

## Linear

Suppose that we use value iteration to generate a sequence of vectors $J k=\left(J k(l), \ldots, J_{k}(n)\right)$ starting with an initial condition vector $J o=$ (Jo(1), $\ldots, J o(n))$ such that

$$
J_{0}(i) \leq \min _{U \in \bar{U}(i)}\left[9\left(i^{\prime} U\right)+\sum_{j \equiv 1}^{n} p_{i j}(u) J_{0}(j)\right] \quad \mathrm{i}=1, \ldots, n o
$$

Then we will have $J_{k}(\mathbf{i}) \leq J_{k+1}$ (i) for all $k$ and i (the monotonicity property of Exercise 1.23 in Chapter 1). It follows from Prop. 7.2.1(a) that we will also have $J o(i) \leq J^{*}(i)$ for all i. Thus $J^{*}$ is the "largest" $J$ that satisfies the constraint

$$
J(i) \leq g(i, u)+\sum_{j=1}^{n} p_{i j}(u) J(j), \quad \text { for all } \cdot \mathrm{i}=1, \ldots, n \text { and } u \in U(i)
$$

In particular, $J^{*}(l), \ldots, J^{*}(n)$ solve the linear program of maximizing $\sum_{i=1}^{n} J(i)$ subject to the above constraint Fig. 7.2.1). Unfortunately, for large $n$ the dimension of this program can be very large and its solution can be impractical, particularly in the absence of special structure.
$J(2)$
$J(2)=g\left(2, u^{2}\right)+p_{21}\left(u^{2}\right) u(1)+p_{z z}\left(u^{2}\right)(2)$


Figure 7.2.1 Linear program associated with a two-state stochastic shortest path problem. The constraint set is shaded, and the objective to maximize is $. l(1)+. l(2)$. Note that because we have $. l(i) \leq J^{*}(i)$ for all i and vectors.$l$ in the constraint set, the vector $.1^{*}$ maximizes any linear cost function of the form $\sum_{i=1}^{n} f 3 i . l(i)$ where $f 3 i \geq 0$ for all i. If $f 3 i>0$ for all i , then $.1^{*}$ is the unique optimal solution of the corresponding linear program.

### 7.3 DISCOUNTED PROBLEMS

We now consider a discounted problem, where there is a discount factor $\alpha<1$. We will show that this problem can be converted to a stochastic shortest path problem for which the analysis of the preceding section holds. To see this, let $\mathrm{i}=1, \ldots, n$ be the states, and consider an associated stochastic shortest path problem involving the states $1, \ldots, n$ plus an extra termination state $t$, with state transitions and costs obtained as follows: From a state $\mathrm{i} \neq t$, when control $U$ is applied, a cost $g(\mathrm{i}, u)$ is incurred, and the next state is $j$ with probability $\alpha p_{i j}(u)$ and $t$ with probability $1-\alpha$; see Fig. 7.3.1. Note that Assumption 7.2.1 of the preceding section is satisfied for the associated stochastic shortest path problem.

Suppose now that we use the same policy in the discounted problem and in the associated stochastic shortest path problem. Then, as long as termination has not occurred, the state evolution in the two problems is governed by the same transition probabilities. Furthermore, the expected cost of the kth stage of the associated shortest path problem is $g\left(X k^{\prime} \mu_{k}\left(x_{k}\right)\right)$ multiplied by the probability that state $t$ has not yet been reached, which is $\alpha^{k}$. This is also the expected cost of the kth stage for the discounted problem. We conclude that the cost of any policy starting from a given state, is the same for the original discounted problem and for the associated stochastic shortest path problem. Furthermore, value iteration produces identical iterates for the two problems. We can thus apply


Figure 7.3.1 Transition probabilities for an a-discounted problem and its associated stochastic shortest path problem. In the latter problem, the probability that the state is not $t$ after $k$ stages is $\alpha^{k}$. The expected cost at each state $\mathrm{i}=1, \ldots, n$ is $g(i, u)$ for both problems, but it must be multiplied by $\alpha^{k}$ because of discounting (in the discounted case) or because it is incurred with probability $\alpha^{k}$ when termination has not yet been reached (in the stochastic shortest path case).
the results the preceding sectlUll ou tha latter problem and obtain the following:

Proposition 7.3.1 The following hold for the discounted problem:
(a) The value iteration algorithm

$$
\begin{equation*}
\min _{u \in U(i)}\left[g(i, u)+\alpha \sum_{j=1}^{n} p_{i j}(u) J_{k}(j)\right] \quad i=1, \ldots, n, \tag{7.21}
\end{equation*}
$$

converges the optimal costs $(i), i=1, \ldots, n$, starting from arbitrary initial conditions Jo(I), $\ldots, \mathrm{Jo}(n)$.
(b) The optimal costs $J^{*}(I), \ldots, J^{*}(n)$ of the discounted problem satisfy Bellman's equation,

$$
\begin{equation*}
\mathrm{J}^{*}(\mathrm{i})=\min _{u \in U(i)}\left[g\left(i^{\prime} u\right)+\alpha \sum_{j=1}^{n} p_{i j}(u) J^{*}(j)\right] \quad \mathrm{i}=1, \ldots, n, \tag{7.22}
\end{equation*}
$$

and in fact they are the unique solution of this equation.
(c) For any stationary policy $\mu$, the costs $J_{\mu}(1), \ldots, J_{\mu}(n)$ are the unique solution of the equation

$$
J_{\mu}(\mathrm{i})=g(i, \mu(i))+\alpha \sum_{j=1}^{n} p_{i j}(\mu(i)) J_{\mu}(j), \quad \mathrm{i}=1, \ldots, n o
$$

Furthermore, given any initial conditions $\mathrm{Jo}(\mathrm{I}), \ldots, \mathrm{Jo}(\mathrm{n})$, the sequence $J k$ (i) generated by the DP iteration

$$
J k+l(i)=g(i, \mu(i))+\alpha \sum_{j=l}^{n} p_{i j}(\mu(i)) J_{k}(j), \quad i=1, \ldots, n,
$$

converges to the cost $J_{\mu}(\mathrm{i})$ for each i.
(d) A stationary policy $\mu$ is optimal if and only if for every state i, $\mu(i)$ attains the minimum in Bellman's equation (7.22).
(e) The policy iteration algorithm given by
$\mu^{k+1}(i)=\arg \min _{u \in U(i)}\left[g(i, u)+\alpha \sum_{j=1}^{n} p_{i j}(u) J_{\mu^{k}}(j)\right], i=1, \ldots, n$,
generates an improving sequence of policies and terminates with an optimal policy.

Proof: Parts (a)-(d) and part
by applying parts (a)-(d) of Prop. 7.2.1, and $\quad 1.2 .2$, respectively, to the associated stochastic shui ${ }^{+2}$ est path proidem described above. Q.E.D.

Bellman's equation (7.22) has a familiar DP interpretation. At state i, the optimal cost $\mathrm{J}^{*}$ (i) is the minimum over all controls of the sum of the expected current stage cost and the expected optimal cost of all future stages. The former cost is $g(i, u)$. The latter cost is $\mathrm{J}^{*}(j)$, but since this cost starts accumulating after one stage, it is discounted by multiplication with $\alpha$.

As in the case of stochastic shortest path problems [see Eq. (7.9) and the discussion following the proof of Prop. 7.2.1], we can show that the error

$$
\operatorname{IJk}(i)-J^{*}(i) 1
$$

is bounded by a constant times $\alpha^{k}$. Furthermore, the error bounds (7.17) become

$$
\begin{equation*}
J_{k+1}(j)+\frac{\alpha}{1-\alpha} \underline{c}_{k} \leq J^{*}(j) \leq J_{\mu^{k}}(j) \leq J_{k+1}(j)+\frac{\alpha}{1-\alpha} \bar{c}_{k}, \tag{7.23}
\end{equation*}
$$

where $\mu^{k}$ is such that $\mu^{k}$ (i) attains the minimum in the kth value iteration (7.21) for all i , and

$$
\underline{c}_{k}=\min _{i=1, \ldots, n}\left[J_{k+1}(i)-J_{k}(i)\right], \quad \bar{c}_{k}=\max _{i=1, \ldots, n}\left[J_{k+1}(i)-J_{k}(i)\right]
$$

since for the associated stochastic shortest path problem it can be shown that for every policy and starting state, the expected number of stages to reach the termination state $t$ is $1 /(1-\alpha)$, so that the terms $N^{*}(j)-1$ and $N k(j) \quad 1$ appearing in Eq. (7.17) are equal to $\alpha /(1-\alpha)$. We note also that there are a number of additional enhancements to the value iteration algorithm for the discounted problem (see Section 1.3 of Vol. II). There are also discounted cost variants of the approximate policy iteration and linear programming approaches discussed for stochastic shortest path problems.

## Example 7.3.1 (Asset Selling)

Consider an infinite horizon version of the asset selling example of Section 4.4, assuming the set of possible offers is finite. Here, if accepted, the amount $x_{k}$ offered in period $k$, will be invested at a rate of interest $r$. By depreciating the sale amount to period 0 dollars, we view $(1+r)-k x k$ as the reward for selling the asset in period $k$ at a price $x k$, where $r>0$ is the rate of interest. Then we have a total discounted reward problem with discount factor $\alpha=$ $1 /(1 r)$. The analysis of the present section is applicable, and the optimal value function $J^{*}$ is the unique solution of Bellman's equation

$$
J^{*}(x)=\max \begin{gathered}
{\left[\begin{array}{c}
E\left\{\begin{array}{l}
J^{*}(w) \\
\underline{1} \pm \underline{r}
\end{array}\right]
\end{array}, . ~\right.}
\end{gathered}
$$

(see Section
The optimal reward function is characterized by the critical number

$$
\bar{\alpha}=\frac{E\left\{\frac{\left.J^{*}(w)\right\}}{l+r},\right.}{}
$$

which can be calculated as in Section 4.4. An optimal policy is to sell if and only if the current offer $x_{k}$ is greater than or equal to $\bar{\alpha}$.
for the states $\mathrm{i}=0,1, \ldots, n-1$, and takes the form

$$
\begin{equation*}
J^{*}(n) \quad K+\alpha(1 \quad p) J^{*}(O)+\alpha p J^{*}(1) \tag{7.25}
\end{equation*}
$$

for state $n$. The first expression within brackets in Eq. (7.24) corresponds to processing the i unfilled orders, while the second expression corresponds to leaving the orders unfilled for one more period. When the maximum $n$ of unfilled orders is reached, the orders must necessarily be processed, as indicated by Eq. (7.25).

To solve the problem, we observe that the optimal cost $J^{*}$ (i) is monotonically nondecreasing in i. This is intuitively clear, and can be rigorously proved by using the value iteration method. In particular, we can show by using the (finite horizon) DP algorithm that the k -stage optimal cost functions $J k(i)$ are monotonically nondecreasing in $i$ for all $k$ (Exercise 7.7), and then argue that the optimal infinite horizon cost function $J^{*}$ (i) is also monotonically nondecreasing in i , since

$$
J^{*}(i)=\lim _{k \rightarrow \infty} J_{k}(i)
$$

by Prop. 7.3.1(a). Given that $J^{*}(i)$ is monotonically nondecreasing in i , from Eq. (7.24) we have that if processing a batch of $m$ orders is optimal, that is,

$$
K+\alpha(1-p) J^{*}(O)+\alpha p J^{*}(1) \leq \mathrm{em}+\alpha(1-p) J^{*}(m)+\alpha p J^{*}(m+1),
$$

then processing a batch of $\mathrm{m}+1$ orders is also optimal. Therefore a threshold policy, that is, a policy that processes the orders if their number exceeds some threshold integer $\mathrm{m}^{*}$, is optimal.

We leave it as Exercise 7.8 for the reader to verify that if we start the policy iteration algorithm with a threshold policy, every subsequently generated policy will be a threshold policy. Since there are $n+1$ distinct threshold policies, and the sequence of generated policies is improving, it follows that the policy iteration algorithm will yield an optimal policy after at most $n$ iterations.

### 7.4 AVERAGE COST PER STAGE PROBLEMS

The methodology of the last two sections applies mainly to problems where the optimal total expected cost is finite either because of discounting or because of a cost-free termination state that the system eventually enters. In many situations, however, discounting is inappropriate and there is no natural cost-free termination state. In such situations it is often meaningful to optimize the average cost per stage starting from a state $i$, which is defined by

$$
J_{\pi}(i)=\lim _{N \rightarrow \infty} \frac{1}{N} E\left\{\sum_{k=0}^{N-1} g\left(x_{k}, \mu_{k}\left(x_{k}\right)\right) \mid x_{0}=i\right\}
$$

Let us first argue heuristically that for most problerns of interest the average cost per stage of a policy and the optimal average cost per stage are independent of the initial state.

To this end we note that the average cost per stage of a policy primarily expresses cost incurred in the long term. Costs incurred in the early stages do not matter since their contribution to the average cost per stage is reduced to zero as $N \rightarrow 00$; that is,

$$
\begin{equation*}
\lim _{N \rightarrow \infty} \frac{1}{N} E\left\{\sum_{k=\emptyset}^{K} g\left(X k^{\prime} \mu_{k}\left(x_{k}\right)\right)\right\}=0 \tag{7.26}
\end{equation*}
$$

for any fixed $K$. Consider now a stationary policy $\mu$ and two states $i$ and $j$ such that the system will, under $\mu$, eventually reach $j$ with probability 1 starting from i. Then intuitively, it is clear that the average costs per stage starting from $i$ and from $j$ cannot be different, since the costs incurred in the process of reaching $j$ from $i$ do not contribute essentially to the average cost per stage. More precisely, let $K_{i j}(\mu)$ be the first passage time from i to $j$ under $\mu$, that is, the first index $k$ for which $X k=j$ starting from $X Q=\mathrm{i}$ under $\mu$ - . $\quad$ нen Llle average cose pei rorresponding ou initial condition $X Q=i$ can be expressed as

$$
\begin{aligned}
J_{\mu}(i)= & \lim _{N \rightarrow \infty} 1\left\{\sum_{k=0}^{K_{i j}(\mu)-1} g\left(X k^{\prime} \mu\left(x_{k}\right)\right)\right\} \\
& +\lim _{N \rightarrow \infty} \frac{1}{N} E\left\{\sum_{k=K_{i j}(\mu)}^{N} g\left(X k^{\prime} \mu\left(x_{k}\right)\right)\right\}
\end{aligned}
$$

If $E\left\{K^{\prime} i j(\mu)\right\}<$ oo (which is equivalent to assuming that the system eventually reaches $j$ starting from i with probability 1 ; see Appendix D ), then it can be seen that the first limit is zero [cf. Eq. (7.26)], while the second limit is equal to $J_{\mu}(j)$. Therefore,

$$
J_{\mu}(i)=J_{\mu}(j), \quad \text { for all } i, j \text { with } E\left\{K_{i j}(\mu)\right\}<00
$$

The preceding argument suggests that the optimal cost $J^{*}$ (i) should also be independent of the initial state $i$ under normal circumstances. To see this, assume that for any two states $i$ and $j$, there exists a stationary policy $\mu$ (dependent on $i$ and $j$ )such that $j$ can be reached from $i$ with probability 1 under $\mu$. Then it is impossible that

$$
\mathrm{J}^{*}(j)<\mathrm{J}^{*}(i)
$$

since when starting from i we can adopt the policy $\mu$ up to the time when $j$ is first reached and then switch to a policy that is optimal when starting from $j$, thereby achieving an average cost starting from $i$ that is equal to $J^{*}(\mathrm{j})$. Indeed, it can be shown that

$$
\mathrm{J}^{*}(\mathrm{i})=J^{*}(j), \quad \text { for all } i, j=1, \ldots, n
$$

under the preceding assumption (see Vol. Section 4.2).

## The Associated Stochastic Shortest Path Problem

The results of this section can be proved under a variety of different assumptions (see Chapter 4 in Vol. II). we will make the following assumption, which will allow us to use the stochastic shortest path analysis of Section 7.2.
upti 7.4.1: One of the states, by convention state $n$, is such that for some integer $\mathrm{m}>0$, and for all initial states and all policies, $n$ is visited with positive probability at least once within the first m stages.

Assumption 7.4.1 can be shown to be equivalent to the assumption that the special state $n$ is recurrent in the Markov chain corresponding to each stationary policy (see Appendix D for the definition of a recurrent state, and Exercise 2.3 of Chapter 2 in Vol. II for a proof of this equivalence).

Under Assumption 7.4.1 we will make an important connection of the average cost problem with an associated stochastic shortest path problem. To motivate this connection, consider a sequence of generated states, and divide it into "independent" cycles marked by successive visits to the state $n$. The first includes the transitions from the initial state to the first visit to state $n$, and the kth cycle, $k=2,3, \ldots$, includes the transitions from the $(k-1)$ th to the kth visit to state $n$. Each of the cycles can be viewed as a state trajectory of a corresponding stochastic shortest path problem with the termination state being essentially $n$.

More precisely, this problem is obtained by leaving unchanged all transition probabilities $p_{i j}(u)$ for $j \neq n$, by setting all transition probabilities $p_{i n}(u)$ to 0 , and by introducing an artificial termination state $t$ to which we move from each state $i$ with probability $p_{i n}(u)$; see Fig. 7.4.1. Note that Assumption 7.4.1 is equivalent to the Assumption 7.2.1 of Section 7.2 under which the results of Section 7.2 on stochastic shortest path problems were shown.

We have specified the probabilistic structure of the stochastic shortest path problem so that its state trajectories replicate the state trajectories of a single cycle of the average cost problem. We will next argue that if we fix the expected stage cost incurred at state i to be

$$
g(i, u)-\lambda^{*}
$$

where $\lambda^{*}$ is the optimal average cost per stage starting from the special state $n$, then the associated stochastic shortest path problem becomes essentially equivalent to the original average cost per stage problem. Furthermore, Bellman's equation for the associated stochastic shortest path problem can


Figure 7.4.1 Transition probabilities for an average cost problem and its associated stochastic shortest path problem. The latter problem is obtained by introducing, in addition to $1, \ldots, n$, an artificial termination state $t$. The corresponding transition probabilities are obtained from the transition probabilities of the original average cost problem as follows: the probabilities of transition from the states $\mathrm{i} \neq t$ to state $t$ are set equal to $\operatorname{Pin}(u)$, the probabilities of transition from all states to state $n$ are set to 0 , and all other transition probabilities are left unchanged.
be viewed as Bellman's equation for the
armane onat per stage problem.

For \& Heuristic argument of why this is so, note that under all stationary policies there will be an infinite number of cycles marked by successive visits to state $n$. From this, it can be conjectured (and it can also be shown, as will be seen later) that the average cost problem is equivalent to a minimum cycle cost problem. This is the problem of finding a stationary policy $\mu$ that minimizes the average cycle cost

$$
\frac{C_{n n}(\mu)}{N_{n n}(\mu)},
$$

where for a fixed $\mu$,
$C_{n n}(\mu)$ : expected cost starting from $n$ up to the first return to $n$,
$N_{n n}(\mu)$ : expected number of stages to return to $n$ starting from $n$.
The intuitive idea here is that the ratio $C_{n n}(\mu) / N_{n n}(\mu)$ is equal to the average cost of $\mu, \mathrm{t}$ so the optimal average cost $\lambda^{*}$ is equal to the optimal cycle cost. Therefore, we have

$$
\begin{equation*}
C_{n n}(\mu)-N_{n n}(\mu) \lambda^{*} \geq 0, \quad \text { for all } \mu, \tag{7.27}
\end{equation*}
$$

[^4]with equality holding if $\mu$ is optimal. Thus, to attain an optimal $\mu$, we must minimize over $\mu$ the expression $C_{n n}(\mu)-N_{n n}(\mu) \lambda^{*}$, which is the expected cost of $\mu$ starting from $n$ in the associated stochastic shortest path problem with stage costs
$$
g(i, u)-\lambda^{*}, \quad i=1, \ldots, n o
$$

Let us denote by $h^{*}$ (i) the optimal cost of this stochastic shortest path problem when starting at the nontermination states $i=1, \ldots, n$. Then by Prop. 7.2.1(b), $h^{*}(I), \ldots, h^{*}(n)$ solve uniquely the corresponding Bellman's equation, which has the form

$$
\begin{equation*}
h^{*}(i)=\min _{u E \cup(i)} g(i, u)-\lambda^{*}+\sum_{J=1}^{l} p_{i j}(u) h^{*}(j)^{l}, \quad i=1, \ldots, n, \tag{7.28}
\end{equation*}
$$

since in the stochastic shortest path problem, the transition probability from i to $j \neq n$ is $\operatorname{Pij}(u)$ and the transition probability from $i$ to $n$ is zero under all $u$. If $\mu^{*}$ is a stationary policy that minimizes the cycle cost, then this policy must satisfy,

$$
C_{n n}\left(\mu^{*}\right)-N_{n n}\left(\mu^{*}\right) \lambda^{*}=0,
$$

and from Eq. (7.27), this policy must also be optimal for the associated stochastic shortest path problem. Thus, we must have

$$
h^{*}(n)=C_{n n}\left(\mu^{*}\right)-N_{n n}\left(\mu^{*}\right) \lambda^{*}=0 .
$$

By using this equation, we can now write Bellman's equation (7.28) as

$$
\begin{equation*}
\lambda^{*}+h^{*}(i)=\min _{u \in U(i)}\left[g(i, u)+t P_{j=1} i j(U) h^{*}(j)\right], \quad i=1, \ldots, n . \tag{7.29}
\end{equation*}
$$

Equation (7.29), which is really Bellman's equation for the associated stochastic shortest path problem, will be viewed as Bellman's equation for the average cost per stage problem. The preceding argument indicates that this equation has a unique solution as long as we impose the constraint $h^{*}(n)=0$. Furthermore, by minimization of its right-hand side we should obtain an optimal stationary policy. We will now prove these facts formally.
$N_{l}, N 2, \ldots, N_{m}$ are the corresponding numbers of stages of these cycles, we have

$$
\lambda_{\mu}=\lim _{m \rightarrow \infty} \frac{\sum_{k=1}^{m} C_{k}}{\sum_{k=1}^{m} N_{k}}=\lim _{m \rightarrow \infty} \frac{\sum_{k=1}^{m} C_{k}}{m} \cdot \lim \quad \frac{m}{\frac{m}{=1} N_{k}}=C_{n n}(\mu) \cdot \frac{1}{N_{n n}(\mu)}
$$

(with probability one).

## Bellman's Equation

The following proposition provides the main results regarding Bellman's equation.

Proposition 7.4.1 Under Assumption 7.4.1 the following hold for the average cost per stage problem:
(a) The optimal average cost $\lambda^{*}$ is the same for all initial states and together with some vector $h^{\star} \quad\left\{h^{*}(1), \ldots, h^{\star}(n)\right\}$ satisfies Bellman's equation
$\lambda^{*}+h^{*}(i)=\min _{u \in U(i)}\left[9\left(i^{\prime} u\right)+\sum_{j=1}^{n} p_{i j}(u) h^{*}(j)\right]$
Furthermore, it $\mu_{(6)} a^{+t}$ ans the minimum in the above equation for all $i$, the stationary policy $\mu$ io .. In addition, out of all vectors $h^{*}$ satisfying this equation, there is a vector for which $h^{*}(n)=0$.
(b) If a scalar $\lambda$ and a vector $h=\{h(I), \ldots, h(n)\}$ satisfy Bellman's equation, then $\lambda$ is the average optimal cost per stage for each initial state
(c) Given a stationary policy $\mu$ with corresponding average cost per stage $\lambda_{\mu}$ there is a unique vector $h_{\mu}=\left\{h_{\mu}(1), \ldots, h_{\mu}(n)\right\}$ such that $h_{\mu}(n)=0$ and

$$
\lambda_{\mu}+h_{\mu}(i)=g(i, \mu(i))+{ }_{j=l}^{n} p_{i j}(\mu(i)) h_{\mu}(j), \quad \mathrm{i}=1, \ldots, n o
$$

Proof: (a) Let us denote

$$
\begin{equation*}
\tilde{\lambda}=\operatorname{m}_{\mu} \dot{m}_{n} \frac{C_{n n}(\mu)}{N_{n n}(\mu)} \tag{7.31}
\end{equation*}
$$

where $C_{n n}(\mu)$ and $N_{n n}(\mu)$ have been defined earlier, and the minimum is taken over the finite set of all stationary policies. Note that $C_{n n}(\mu)$ and $N_{n n}(\mu)$ are finite in view of Assumption 7.4.1 and the results of Section 7.2. Then we have

$$
C_{n n}(\mu)-N_{n n}(\mu) \tilde{\lambda} \geq 0,
$$

with equality holding for all $\mu$ that attain the minimum in (7.31). Consider the associated stochastic shortest path problem when the expected
stage cost incurred at state $i$ is

$$
g(i, u) \quad \tilde{\lambda} .
$$

Then by Prop. 7.2.1(b), the costs $h^{*}(I), \ldots, h^{*}(n)$ solve uniquely the corresponding Bellman's equation

$$
\begin{equation*}
\left.h^{*}(i)=\min _{u E \cup(i)} g(i, u)-\tilde{\lambda}+\sum_{j=1}^{n-1} p_{i j}(u) h^{*}(j)\right], \tag{7.32}
\end{equation*}
$$

since the transition probability from $i$ to $n$ is zero in the associated stochastic shortest path problem. An optimal stationary policy must minimize the cost $C_{n n}(\mu)-N_{n n}(\mu) \tilde{\lambda}$ and reduce it to zero [in view of Eq. (7.31)], so we see that

$$
\begin{equation*}
h^{*}(n)=0 . \tag{7.33}
\end{equation*}
$$

Thus, Eq. (7.32) is written as

$$
\begin{equation*}
\tilde{\lambda}+h^{*}(i)=\min _{u \in U(i)}\left[9\left(i^{\prime} u\right)+t P j j(U) h^{*}(j)\right], \quad i=1, \ldots, n . \tag{7.34}
\end{equation*}
$$

We will show that this relation implies that $\tilde{\lambda}=\lambda^{*}$.
let $\pi=\left\{\mu_{0}, \mu_{1}, \quad\right\}$ be any admissible policy, let $N$ be a positive integer, and for all $k=0, \quad, N-1$, define $J_{k}(\mathrm{i})$ using the following recursion

$$
=h^{*}(i), \quad \mathrm{i}=1, \ldots, n,
$$

$$
\begin{equation*}
\text { (i) }=g\left(i, \mu_{N-k-1}(i)\right)+\sum_{j=1}^{n} p_{i j}\left(\mu_{N-k-1}(i)\right) J_{k}(j), \quad \mathrm{i}=1, \ldots, n o \tag{7.35}
\end{equation*}
$$

Note that $J_{N}(i)$ is the N -stage cost of $\pi$ when the starting state is i and the terminal cost function is $h^{*}$. From (7.34), we have

$$
\tilde{\lambda}+h^{*}(i) \leq g\left(i, \mu_{N-1}(i)\right)+\sum_{j=1}^{n} p_{i j}\left(\mu_{N-1}(i)\right) h^{*}(j), \quad i=1, \ldots, n,
$$

or equivalently, using Eq. (7.35) for $k=0$ and the definition of $J 0$,

$$
\tilde{\lambda}+J_{0}(i) \leq J_{1}(i), \quad \mathrm{i}=1, \ldots, n o
$$

Using this relation, we have

$$
\begin{aligned}
g\left(i, \mu_{N-2}(i)\right)+\tilde{\lambda} & +\sum_{j=1}^{n} p_{i j}\left(\mu_{N-2}(i)\right) J_{0}(j) \\
& \leq g\left(i, \mu_{N-2}(i)\right)+{ }_{j=I}^{n}\left(\mu_{N-2}(i)\right) J_{1}(j), \quad \text { i }=1, \ldots, n .
\end{aligned}
$$

Eq. (7.34) and the definition $J o(j)=h^{*}(j)$, the left-hand side of the above inequality is no less than $2 \tilde{\lambda}+h^{*}(i)$, while by Eq. (7.35), the righthand side is equal to $J 2$ (i). We thus obtain

$$
2 \tilde{\lambda}+h^{*}(i) \leq J_{2}(i), \quad i=1, \ldots, n o
$$

By repeating this argument several times, we obtain

$$
k \tilde{\lambda}+h^{*}(i) \leq J_{k}(i), \quad k=0, \ldots, N, \quad \mathrm{i}=1, \ldots, n
$$

and in particular, for $k=N$,

$$
\begin{equation*}
\tilde{\lambda}+\frac{h^{*}(i)}{N} \leq \frac{1}{N} J_{N}(i), \quad \text { i } \quad 1, \ldots, n o \tag{7.36}
\end{equation*}
$$

Furthermore, equality holds in the above relation if $\mu_{k}(i)$ attains the minimum in Eq. (7.34) for all i and $k$.

Let us now take the limit as $N \rightarrow \infty$ in Eq. (7.36). The left-hand side tends to $\tilde{\lambda}$. We claim that the right-hand side tends to $J_{\pi}(i)$, the average cost per stage of $\pi$ starting at state $i$. The reason is that from the definition (7.35), $I N(i)$ is the $N$-stage cost of $\pi$ starting at $i$, when the terminal cost function is $h^{*}$; when we take the limit of $(I / N) J N(i)$, the dependence on the terminal cost function $h^{*}$ disappears. Thus, by taking the limit as $N \rightarrow \infty$ in Eq. (7.36), we obtain

$$
\tilde{\lambda} \leq J_{\pi}(i), \quad i=1, \ldots, n
$$

for all admissible $\pi$, with equality if $\pi$ is a stationary policy $\mu$ such that $\mu(\mathrm{i})$ attains the minimum in Eq. (7.34) for all i and $k$. It follows that

$$
\tilde{\lambda}=\min _{\pi} J_{\pi}(i)=\lambda^{*}, \quad \mathrm{i}=1, \ldots, n
$$

and from Eq. (7.34), we obtain the desired Eq. (7.30).
Finally, Eqs. (7.33) and (7.34) are equivalent to' Bellman's equation (7.32) for the associated stochastic shortest path problem. Since the solution to the latter equation is unique, the same is true for the solution of Eqs. (7.33) and (7.34).
(b) The proof of this part is obtained by using the argument of the proof of part (a) following Eq. (7.34).
(c) The proof of this part is obtained by specializing part (a) to the case where the constraint set at each state $\mathbf{i}$ is $U(i)=\{\mu(i)\}$.

An examination of the preceding proof show that the unique vector $h^{*}$ in Bellman's equation (7.30) for which $h^{*}(n)=$ is the optimal cost vector for the associated stochastic shortest path problem when the expected stage cost at state $i$ is

$$
g(i, u)-\lambda^{*}
$$

[cf. (7.32)]. Consequently, $h^{*}(i)$ has the interpretation of a relative or differential cost; it is the minimum of the difference between the expected cost to reach $n$ from $i$ for the first time and the cost that would be incurred if the cost per stage was the average $\lambda^{*}$. We note that the relation between optimal policies of the stochastic shortest path and the average cost problems is clarified in Exercise 7.15.

We finally mention that Prop. 7.4.1 can be shown under considerably weaker conditions (see Section 4.2 of Vol. II). In particular, Prop. 7.4.1 can be shown assuming that all stationary policies have a single recurrent class, even if their corresponding recurrent classes do not have state $n$ in common. The proof, however, requires a more sophisticated use of the connection with an associated stochastic shortest path problem. Proposition 7.4.1 can also be shown assuming that for every pair of states $i$ and $j$, there exists a stationary policy under which there is positive probability of reaching $j$ starting from i. In this case, however, an associated stochastic shortest path problem cannot be defined and the corresponding connection with the average cost per stage problem cannot be made. The analysis of Chapter 4 of Vol. II relies on another connection that exists between the average cost per stage problem and the discounted cost problem, but to establish this connection and to fully explore its ramifications, a much more sophisticated analysis is required.

## Example 7.4.1

Consider the average cost version of the manufacturer's problem of Example 7.3.2. Here, state 0 plays the role of the special state $n$ in Assumption 7.4.1. Bellman's equation takes the form
$\lambda^{*}+h^{*}(i)=\min \left[\mathrm{K}+(1-p) h^{*}(O)+p h^{*}(I), \mathrm{ci}+(1-p) h^{*}(i)+p h^{*}(i+1) \mathrm{J}\right.$, (7.37)
for the states $\mathrm{i}=0,1, \ldots, n-1$, and takes the form

$$
\lambda^{*}+h^{*}(n)=K+(1-p) h^{*}(O)+p h^{*}(I)
$$

for state $n$. The first expression within brackets in Eq. (7.37) corresponds to processing the i unfilled orders, while the second expression corresponds to leaving the orders unfilled for one more period. The optimal policy is to process i unfilled orders if

$$
K+(1-p) h^{*}(O)+p h^{*}(I) \leq \mathrm{ci}+(1-p) h^{*}(i)+p h^{*}(i+1)
$$

If we view $h^{*}(i), \mathrm{i}=1, \ldots, n$, as differential costs associated with an optimal policy, it is intuitively clear that $h^{*}(\mathrm{i})$ is monotonically nondecreasing with i [this can also be proved by interpreting $h^{*}(i)$ as optimal costs-to-go for the associate stochastic shortest path problem, or by using analysis based on the theory presented in Vol. II, Section 4.2]. As in Example 7.3.2, the monotonicity property of $h^{*}(i)$ implies that a threshold policy is optimal.

## Value Iteration

The most natural version of the value iteration method for the average cost problem is simply to select arbitrarily a terminal cost function, say $J o$, and to generate successively the corresponding optimal k-stage costs $J_{k}(i)$, $k=1,2, \ldots$ This can be done by executing the DP algorithm starting with $J o$, that is, by using the recursion

$$
\begin{equation*}
J_{k+1}(i)=\min _{u \in U(i)}\left[g(i, u)+\sum_{j=1}^{n} p_{i j}(u) J_{k}(j)\right] \quad \mathrm{i}=1, \ldots, n o \tag{7.38}
\end{equation*}
$$

It is natural to expect that the ratios $J k(i) / k$ should converge to the optimal average cost per stage $\Lambda^{*}$ as $k \rightarrow o o$, i.e.,

$$
\lim _{k \rightarrow \infty} \frac{J k(i)}{k}=\lambda^{*} .
$$

To show this, let us define the recursion

$$
\min _{u E U(i)}\left[9(i, u)+\sum_{f=1}^{n} p_{i j}(u) J_{k}^{*}(j)\right], \quad i=1, \ldots, n
$$

with the initial condition

$$
J_{0}^{*}(i)=h^{*}(i), \quad i=1, \ldots, n,
$$

where $h^{*}$ is a differential cost vector satisfying Bellman's equation

$$
\begin{equation*}
\lambda^{*}+h^{*}(i)=\min _{u E U(i)}\left[9\left(i^{\prime} U\right)+\sum_{j=1}^{n} p_{i j}(u) h^{*}(j)\right], \quad \mathrm{i}=1, \ldots, n o \tag{7.39}
\end{equation*}
$$

Using this equation, it can be shown by induction that for all $k$ we have

$$
J_{k}^{*}(i)=k \lambda^{*}+h^{*}(i), \quad \mathrm{i}=1, \ldots, n o
$$

On the other hand, it can be seen that for all $k$,

$$
J_{k}^{*}(i) \mid \leq \max _{J=l, \ldots, n} I J o(j)-h^{*}(j) l, \quad \mathrm{i}=1, \ldots, n o
$$

The reason is that $J k$ (i) and $J_{k}^{*}$ (i) are optimal costs for two k-stage problems that differ only in the corresponding terminal cost functions, which are Jo and $h^{*}$, respectively. From the preceding two equations, we see that for all $k$,

$$
J_{k}(\mathrm{i})-k \lambda^{*}\left|\leq \max _{J=l, \ldots, n}\right| J^{\circ}(j)-\left.h^{\star}(j)\left|+\max _{J=, \ldots, n}\right| h^{*}(j)\right|^{\prime} \quad \mathrm{i}=1, \ldots, n o
$$

so that $J k(i) / k$ converges to $\lambda^{*}$ at the rate of a constant divided by $k$. Note that the above proof shows that $J k(i) / k$ converges to $\lambda^{*}$ under any conditions that guarantee that Bellman's equation (7.39) holds for some vector $h^{*}$.

The value iteration method just described is simple and straightforward, but has two drawbacks. First, since typically some of the components of $J k$ diverge to oo or -oo, direct calculation of limk-+oo $J k(\mathrm{i}) / k$ is numerically cumbersome. Second, this method will not provide us with a corresponding differential cost vector $h^{*}$. We can bypass both difficulties by subtracting the same constant from all components of the vector $J k$, so that the difference, call it $h k^{\prime}$ remains bounded. In particular, we can consider the algorithm

$$
\begin{equation*}
h_{k}(i)=J_{k}(i)-J_{k}(s), \quad \mathrm{i}=1, \ldots, n o \tag{7.40}
\end{equation*}
$$

where $s$ is some fixed state. By using Eq. (7.38), we then obtain

$$
\begin{aligned}
h k+l(i)= & J k+l(i)-J k+l(S) \\
= & \min _{U E(i)}\left[9\left(i^{\prime} u\right)+\sum_{j=1}^{n} p_{i j}(u) J_{k}(j)\right], \\
& \min _{u \in U(s)}\left[g(S, U)+\sum_{j=1}^{n} p_{s j}(u) J_{k}(j)\right],
\end{aligned}
$$

from which in view of the relation $h k(j)=J k(j)-J k(S)$, we have

$$
\begin{align*}
h k+1(i)= & \min _{u \in U(i)}\left[9\left(i^{\prime} U\right)+\sum_{j=1}^{n} p_{i j}(u) h_{k}(j)\right] \\
& -\min _{u E U(s)}\left[g\left(S^{\prime} u\right)+\sum_{j=1}^{n} p_{s j}(u) h_{k}(j)\right] \quad \mathrm{i}=1, \ldots, n . \tag{7.41}
\end{align*}
$$

The above algorithm, known as relative value iteration, is mathematically equivalent to the value iteration method (7.38) that generates $J k(i)$. The iterates generated by the two methods merely differ by a constant [d. Eq. (7.40)], and the minimization problems involved in the corresponding iterations of the two methods are mathematically equivalent. However, under Assumption 7.4.1, it can be shown that the iterates $h k(i)$ generated by the relative value iteration method are bounded, while this is typically not true for the value iteration method.

It can be seen that if the relative value iteration (7.41) converges to some vector $h$, then we have

$$
\lambda+h(i)=\min _{u \in U(i)}\left[g(i, u)+\sum_{j=1}^{n} p_{i j}(u) h(j)\right]
$$

where

$$
\lambda=\min _{u \in U(s)}\left[9(8, u)+\sum_{j=1}^{n} p_{s j}(U) h(j)\right]
$$

By Prop. 7.4.1(b), this implies that $\lambda$ is the optimal average cost per stage for all initial states, and $h$ is an associated differential cost vector. Unfortunately, the convergence of the relative value iteration is not guaranteed under Assumption 7.4.1 (see Exercise 7.14 for a counterexample). A stronger assumption is required. It turns out, however, that there is a simple variant of the relative value iteration for which convergence is guaranteed under Assumption 7.4.1. This variant is given by

$$
(i)=\left(\begin{array}{ll}
1 & T) h k(i)+\min _{u \in U(i)}\left[g\left(i^{\prime} u\right)+\tau \sum_{j=1}^{n} p_{i j}(u) h_{k}(j)\right] \tag{7.42}
\end{array}\right]
$$

$$
\left.\min _{u E U(s)} 19(5, U)+T J L .=n 1 P S j(U) h k(j)\right], \quad i=/, \ldots, n
$$

where $T$ is a scalar such that $0<T<1$
1 we obtain
the relative sutiation (7 . The convergence proof of this algorithm is somewhat complicated. It can be found in Section 4.3 of Vol. II.

## Iteration

It is possible to use a policy iteration algorithm for the average cost problem. This algorithm operates similar to the policy iteration algorithms of the preceding sections: given a stationary policy, we obtain an improved policy by means of a minimization process, and continue until no further improvement is possible. In particular, at the typical step of the algorithm, we have a stationary policy $\mu^{k}$. We then perform a policy evaluation step; that is, we obtain corresponding average and differential costs $\lambda^{k}$ and $h k(i)$ satisfying

$$
\begin{aligned}
& \lambda^{k}+h k(i)=g\left(i, \mu^{k}(i)\right)+ \sum_{j=l}^{n} p_{i j}\left(\mu^{k}(i)\right) h^{k}(j), \quad \mathrm{i}=1, \ldots, n, \\
& h k(n)=0 .
\end{aligned}
$$

We subsequently perform a policy improvement step; that is, we find a stationary policy $\mu^{k+1}$, where for all i, $\mu^{k+1}$ (i) is such that

$$
\begin{aligned}
g\left(i, \mu^{k+1}(i)\right)+ & \sum_{j=1}^{n} p_{i j}\left(\mu^{k+1}(i)\right) h^{k}(j) \\
& \min _{u \in U(i)}\left[9(i, u)+\sum_{j=1}^{n} \operatorname{Pij}(U) h k(j)\right]
\end{aligned}
$$

If $\lambda^{k+1}=M k$ and $h^{k}{ }_{+}^{1}(\mathrm{i})=h^{k}(\mathrm{i})$ for all i, the algorithm terminates; otherwise, the process is repeated with $\mu^{k+1}$ replacing $\mu^{k}$.

To prove that the policy iteration algorithm terminates, it is sufficient that each iteration makes some irreversible progress towards optimality, since there are finitely many stationary policies. The type of irreversible progress that we can demonstrate is described in the following proposition, which also shows that an optimal policy is obtained upon termination.

Proposition 7.4.2 Under Assumption 7.4.1, in the policy iteration algorithm, for each $k$ we either have

$$
\lambda^{k+1}<\lambda^{k}
$$

or else we have

$$
\lambda^{k+1}=\lambda^{k}, \quad h^{k+1}(i) \leq h^{k}(i), \quad \mathrm{i}=1, \ldots, n o
$$

Furthermore, the algorithm terminates, and the policies $\mu^{k}$ and $\mu^{k+1}$ obtained upon termination are optimal.

Proof: To simplify notation, denote $\mu^{k}=\mu, \mu^{k+1}=\bar{\mu}, \lambda^{k}=\lambda, \lambda^{k+1} \quad \lambda$, $h k(i)=h(i), h^{k+1}(i)=\quad$ Define for $\mathrm{N}=1,2, \ldots$

$$
h N(i)=g(i, \bar{\mu}(i))+\sum_{j=1}^{n} p_{i j}(\bar{\mu}(i)) h_{N-1}(j), \quad i=1, \ldots, n
$$

where

$$
h o(i)=h(i), \quad i=1, \ldots, n o
$$

Note that $h N(i)$ is the $N$-stage cost of policy $\bar{\mu}$ starting from i when the terminal cost function is $h$. Thus we have

$$
\begin{equation*}
\bar{\lambda}=J_{\bar{\mu}}(i)=\lim _{N \rightarrow \infty} \underline{N} \underline{N} N(i), \quad i=1, \ldots, n, \tag{7.43}
\end{equation*}
$$

since the contribution of the terminal cost to $(1 / N) h_{N}$ (i) vanishes when $\mathrm{N} \rightarrow$ oo. By the definition of $\bar{\mu}$ and Prop. 7.4.1(c), we have for all i

$$
\begin{aligned}
h 1(i) & =g(i, \bar{\mu}(i))+\sum_{j=1}^{n} p_{i j}(\bar{\mu}(i)) h_{0}(j) \\
& \left.\leq g(i, \mu(i))+\sum_{j=1}^{n} p_{i j}(\mu(i)) h_{0}(j)\right) \\
& =\lambda+h o(i) .
\end{aligned}
$$

From the above equation, we also obtain

$$
\begin{aligned}
h_{2}(i)= & g(i, \bar{\mu}(i))+\sum_{j=l}^{n} p_{i j}(\bar{\mu}(i)) h_{1}(j) \\
\leq & g(i, \bar{\mu}(i))+\sum_{j=l}^{n} p_{i j}(\bar{\mu}(i))(\lambda+h o(j)) \\
& \lambda+9(\mathrm{i}, \bar{\mu}(\mathrm{i}))+\sum_{j=l}^{n} P i j(\bar{\mu}(\mathrm{i})) h o(j) \\
\leq & \lambda+g(i, \mu(i))+\sum_{j=l}^{n} p_{i j}(\mu(i)) h_{0}(j) \\
= & 2 \lambda+h o(i)
\end{aligned}
$$

and by proceeding similarly, we see that for all i and N we have

$$
h_{N}(i) \leq N \lambda+h_{0}(i)
$$

Thus,

$$
\frac{1}{1 \mathrm{~V}} h_{-i v}(i) \leq \lambda+\underline{1} \underline{N h o(i)}
$$

and by $\quad$ the limit as $\mathrm{N} \rightarrow 0$ oo and using Eq. (7.43), we obtain $\bar{\lambda} \leq \lambda$.
If $\bar{\lambda}=\lambda$, then it is seen that the iteration that produces $\mu^{k+1}$ is a policy improvement step for the associated stochastic shortest path problem with cost per stage

$$
g(i, u)-\lambda
$$

Furthermore, $h(\mathrm{i})$ and $\bar{h}(\mathrm{i})$ are the optimal costs starting from $i$ and corresponding to $\mu$ and $\bar{\mu}$, respectively, in this associated stochastic shortest path problem. Thus, by Prop. 7.2.2, we must have $\bar{h}(\mathrm{i}) \leq h(\mathrm{i})$ for all i.

In view of the improvement properties just shown, no policy can be repeated without termination of the algorithm. Since there are only a finite number of policies, it follows that the algorithm will terminate. Let us now show that when the algorithm terminates with $\bar{\lambda}=\lambda$ and $\bar{h}(\mathrm{i})=h(\mathrm{i})$ for all $i$, the policies $\bar{\mu}$ and $\mu$ are optimal. Indeed, upon termination we have for all $i$

$$
\begin{aligned}
\lambda+h(\mathrm{i}) & =\bar{\lambda}+\bar{h}(\mathrm{i}) \\
& =g(i, \bar{\mu}(i))+\sum_{j=l}^{n} p_{i j}(\bar{\mu}(i)) \bar{h}(j) \\
& =9(\mathrm{i}, \bar{\mu}(\mathrm{i}))+\sum_{j=l}^{n} P i j(\bar{\mu}(\mathrm{i})) h(j) \\
& =\min _{u E U(i)}\left[9\left(i^{\prime} U\right)+t_{j=l} P j j(U) h(j)\right]
\end{aligned}
$$

Thus $\lambda$ and $h$ satisfy Bellman's equation, and by Prop. 7.4.1(b), $\lambda$ must be equal to the optimal average cost. Furthermore, $\bar{\mu}(i)$ attains the minimum in the right-hand side of Bellman's equation, so by Prop. 7.4.1(a), $\bar{\mu}$ is optimal. Since we also have for all i

$$
\lambda+h(i)=g(i, \mu(i))+{ }_{j=l}^{n} p_{i j}(\mu(i)) h(j)
$$

the same is true for $\mu$. Q.E.D.
We note that policy iteration can be shown to terminate with an optimal stationary policy under less restrictive conditions than Assumption 7.4.1 (see Vol. II, Section 4.3).

### 7.5 SEMI-MARKOV PROBLEMS

have considered so far problems where the cost per stage does not depend on the time required for transition from one state to the next. Such problems have a natural discrete-time representation. On the other hand, there are situations where controls are applied at discrete times but cost is continuously accumulated. Furthermore, the time between successive control choices is variable; it may be random or it may depend on the current state and the choice of control. For example, in queueing systems state transitions correspond to arrivals or departures of customers, and the corresponding times of transition are random. In this section, we discuss continuous-time, infinite horizon problems with a finite number of states. We will provide a fairly straightforward extension of our earlier infinite horizon analysis for discrete-time problems.

We assume that there are $n$ states, denoted by $1, \ldots, n$, and that state transitions and control selections take place at discrete times, but the length of the time interval from one transition to the next is random. The state and control at any time $t$ are denoted by $x(t)$ and $u(t)$, respectively, and stay constant between transitions. We use the following notation:
$t k$ : The time of occurrence of the kth transition. By convention, we denote $t o=0$.

$$
\begin{aligned}
& X k=X(t k): \text { We have } x(t)=X k \text { for } t k \leq t<t k+l^{\prime} \\
& U k=U(t k): \text { We have } u(t)=U k \text { for } t k \leq t<t k+1
\end{aligned}
$$

In place of transition probabilities, we have transition distributions Qij( $r, u$ ), which for a given pair $(i, u)$, specify the joint distribution of the transition interval and the next state:

$$
Q_{i j}(\tau, u)=P\left\{t_{k+1}-t_{k} \leq \tau, x_{k+1}=j \mid x_{k}=i, u_{k}=u\right\}
$$

Note that the transition distributions specify the ordinary transition probabilities via

$$
P^{\prime} i j(U)=P\{X k+l=j \mid X k=\mathrm{i}, U k=U\}=\lim _{\tau \rightarrow \infty} Q i j(7, U)
$$

Note also that the conditional cumulative distribution function (CDF) of $\tau$ given $\mathrm{i}, j, 11$, is

$$
-t k \leq \tau \mid X k=i \quad x k+l=j, U k=U\}=\underset{P i j}{\operatorname{Qij}(7, l l,)}
$$

[assuming that $\operatorname{Pij}(u)>0]$. I $\quad Q i j(\tau, u)$ can be viewed as a "scaled CDF", i.e., a CDF multiplied by $p_{i j}(\dot{u})$ (see Fig. 7.4.2).

(a)


Figure 7.4.2 Illustration of the transition distributions $Q_{i j}(T, u)$ and the conditional CDF of $\tau$. Figures (a) and (b) correspond to the cases where $T$ is a continuous and a discrete random variable, respectively.

An advantage of working with transition distributions $\operatorname{Qij}(7,11$,$) is that$ they can be used to model discrete, continuous, and mixed distributions for the transition time $\tau$. Generally, expected values of functions of $\tau$ can be written as integrals involving the differential of $Q i j$ with respect to $\tau$, denoted $\operatorname{dQij}(7, u)$. For example, the conditional expected value of $\tau$ given $\mathrm{i}, j$, and $u$ is written using the conditional CDF (7.44) as

$$
\begin{equation*}
\mathrm{E}\{71 \overline{i, j,-}, u\}=\int_{Q}^{\infty} \tau \underset{p_{i j}}{\operatorname{dij}}-u(J) \tag{7.45}
\end{equation*}
$$

If $Q_{i j}(\tau, u)$ is continuous and piecewise differentiable with respect to $\tau$, its partial derivative

$$
\left.q_{i j}(\tau, u)=\frac{d Q i j(7,}{d \tau} \underline{u}\right)
$$

can be viewed as a "scaled" density function for $\tau$. Then, $d Q i j(\tau, u)$ may be replaced by $q i j(\tau, u) d 7$, and expected values of functions of $\tau$ can be written in terms of $q i j(7, u)$. For example, Eq. (7.45) is written as

$$
E\{T \mid \bar{i}, \bar{j}, u\}=\int_{\Omega}^{\infty} \tau_{p_{i j} \subseteq u^{-}}^{q i j(7, U)} d \tau
$$

If $Q i j(\tau, u)$ is discontinuous and "staircase-like," then $\tau$ is a discrete random variable, and expected values of functions of $\tau$ can be written as summations.

We will assume that for each state i and control $u \mathrm{E} U(i)$, the expected transition time, denoted $\bar{\tau}_{i}(u)$, is nonzero and finite:

$$
\begin{equation*}
0<\bar{\tau}_{i}(u)<\infty \tag{7.46}
\end{equation*}
$$

In view of Eq. (7.45), this expected transition time is given by

$$
\bar{\tau}_{i}(u)=\sum_{j=1}^{n} p_{i j}(u) E\{\tau \mid i, j, u\}=\sum_{j=1}^{n} \int_{0}^{\infty} \tau d Q_{i j}(\tau, u)
$$

Optimal control problems involving continuous-time Markov chains of the type described above are called semi-Markov problems. The reason is that, for a given policy, while at a transition time $t k$ the future of the system probabilistically depends only on the current state, at other times it may depend in addition on the time elapsed since the preceding transition. In fact, if we were to allow the control to depend continuously on the time $t$ (rather than restricting the choice of control to just the transition times $t k$ ), we would obtain a problem where there is genuine benefit to the controller for knowing the time elapsed since the preceding transition. We would then have to include this elapsed time as of the state, and we would obtain a difficult (infinite state space) problem. This type of complication is avoided in our formulation by restricting the control to change only at the transition times $t k$.

We note, however, that there is a special case where the future of the system depends only on its current state at all times, and there is no benefit in allowing the control to depend continuously on the time elapsed since the preceding transition. This is the case where the transition distributions are exponential, of the form

$$
Q_{i j}(\tau, u)=p_{i j}(u)\left(1-e^{-\nu_{i}(u) \tau}\right)
$$

where $\operatorname{Pij}(u)$ are transition probabilities, and $V i(u)$ are given positive scalars, called the transition rates at the corresponding states $i$. In this case, if the system is in state i and control $u$ is applied, the next state will be $j$ with probability $P i j(u)$, and the time interval between the transition to state $i$ and the transition to the next state is exponentially distributed with parameter $V i(11$,$) ; that is,$
$P\{$ transition time interval $>\tau$ Ii, $u\}=e-V i(u) T$.
The exponential distribution has the so called memoryless property, which in our context implies that, for any time t between the transition times $t k$
and , the additional time $t k+1-t$ needed to effect the next transition is independent of the time $t-t k$ that the system has been in the current state. To see this, use the following generic calculation

$$
\begin{aligned}
P\left\{\tau>r 1+r 2 \mid T>r_{1}\right\}= & P\{T>r 1+r 2\} \\
& \left.\frac{P\{\tau>r 1}{}\right\} \\
& \frac{e-v(r l+r 2)}{e^{-\nu r_{1}}} \\
= & e^{-\nu r_{2}} \\
& -P\left\{\tau>r_{2}\right\}
\end{aligned}
$$

where $T 1=t-t k, T 2=t k+1-\mathrm{t}$, and $\nu$ is the trensition rate. Thus, when the transition distributions are exponential, the state evolves in continuous time as a Markov process, but this need not be true i.r a more general distribution.

We assume that for given state $i$ and control $u \mathrm{E} U(i)$, the cost that is incurred in a small time interval $d t$ is $g(i, u) d t$. Thus, we mä view $g(i, u)$ as cost peT unit time. Based on this generic cost structure, we whe consider analogs of the discounted and average cost per stage problems of the preceding sections.

Discounted Problems
Here the cost function has the form

$$
\lim _{T \rightarrow \infty} E\left\{\int_{0}^{T} e-\left(3 \operatorname{tg}\left(x(t)^{\prime} U(t)\right) d t\right\}\right.
$$

where ,6 is a given positive discount parameter. Since the cost per unit time, $9(x(t), u(t))$, remains constant between transitions, the expected cost of a single transition from state $i$ under control $u$ is given by

$$
\begin{aligned}
G(i, u) & =\mathrm{E}\left\{\int_{0}^{\tau} e^{-\beta t} g(\mathrm{i}, U) d t\right\} \\
& =g(i, u) \mathrm{E}\left\{\int_{0}^{\tau} e^{-\beta t} d t\right\} \\
& =g(i, u) \quad\left\{E r\left\{\int_{0}^{\tau} e^{-\beta t} d t \mid j\right\}\right\} \\
& =g(\mathrm{i}, u) \sum_{\mathrm{j}=1}^{n} \operatorname{Pij}(u) \int_{0}^{\circ}\left(\int_{0}^{\tau} e-(3 t d t) \frac{d Q_{i j}(\tau,}{-p_{i j}(u)} \underline{u)}\right.
\end{aligned}
$$

or equivalently, since $\int_{0}^{\tau} e-\left(3 t d t=(1)_{-} e-(3 T) /(3\right.$,

$$
\begin{equation*}
G(i, u)=g(i, u) \sum_{j=I} \quad 0 \quad \operatorname{loo} 1_{-} e-(3 T(T, U) \tag{7.47}
\end{equation*}
$$

The cost of an admissible policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots\right\}$ starting from state $i$ is given by

$$
J_{\pi}(i)=\lim _{N \rightarrow \infty} \sum_{k=0}^{N-1} E\left\{\int_{t_{k}}^{t_{k+1}} e^{-\beta t} g\left(x_{k}, \mu_{k}\left(x_{k}\right)\right) d t \mid x_{0}=i\right\}
$$

This cost can be broken down into the sum of the expected cost of the first transition, which is $G\left(1, \mu_{0}(i)\right)$, plus the expected cost-to-go starting from the next state, discounted by the factor $e^{-\beta \tau}$, where $\tau$ is the (random) time when the first transition occurs:

$$
J_{\pi}(i)=G\left(i, \mu_{0}(i)\right)+E\left\{e^{-\beta \tau} J_{\pi_{1}}(j) \mid x_{0}=i, u_{0}=\mu_{0}(i)\right\}
$$

The last term in the above equation can be calculated as

$$
\begin{aligned}
E\left\{e-\left(3 T J_{\pi_{1}}(j)\right.\right. & \left.\mid X_{0}=i, U_{0}=\mu_{0}(\mathrm{i})\right\} \\
& =E\left\{E\left\{e-(3 T \mid j\} J_{\pi_{1}}(\mathrm{j}) \mid X_{0}=\mathrm{i}, U_{0}=\mu_{0}(\mathrm{i})\right\}\right. \\
& =\sum_{J=1}^{n} p_{i j}\left(\mu_{0}(\mathrm{i})\right)\left(\int_{0}^{\infty} e^{-\beta \tau} \frac{\left.d Q, ; \frac{(\tau,}{p_{i j}} \frac{\left.\mu_{0}(i)\right)}{\left(\mu_{0}(i)\right)}\right) J_{1 T 1}(j)}{}\right. \\
& =\sum_{J=I}^{n} m_{i j}\left(\mu_{0}(i)\right) J_{\pi_{1}}(j)
\end{aligned}
$$

where for any $U E U(\mathrm{i}), \operatorname{mij}^{( }(u)$ is given by

$$
\begin{equation*}
m_{i j}(u)=\int_{0}^{\infty} e^{-\beta \tau} d Q_{i j}(\tau, u) \tag{7.49}
\end{equation*}
$$

Thus, combining Eqs. (7.47)-(7.49), we see that $J_{I T}$ (i) can be written as

$$
\begin{equation*}
J_{I T}(i)=G\left(i, \mu_{0}(i)\right)+\sum_{j=I}^{n} m_{i j}\left(\mu_{0}(i)\right) J_{\pi_{1}}(j) \tag{7.50}
\end{equation*}
$$

which is similar to the corresponding equation for discounted discrete-time problems [mij $\left(\mu_{0}(\mathrm{i})\right)$ replaces $\left.\alpha p_{i j}\left(\mu_{0}(\mathrm{i})\right)\right]$.

In analogy with the discrete-time case, we may associate with a stochastic shortest path problem involving an artificial termination state t . Under control $u$, from state $i$ the system moves to state $j$ with probability $m i j(u)$ and to the termination state $t$ with probability

$$
1-\sum_{j=1}^{n} m i j(u)
$$

The assumption of a positive expected transition time [d. Eq. (7.46)] implies that

$$
{ }_{j=1}^{n} \operatorname{mij}(u)<1, \quad \text { for all } i, u \mathrm{E} U(i)
$$

so that the Assumption 7.2.1, which is required for the validity of our stochastic shortest path analysis of Section 7.2, is satisfied. By using an essentially identical approach to the one of Section 7.3 , we can derive analogs of all the discounted cost results of Prop. 7.3.1. In particular, the optimal cost function $\mathrm{J}^{*}$ is the unique solution of Bellman's equation

$$
J^{*}(i)=\min _{u \in U(i)}\left[G(i, u)+\sum_{j=1}^{n} m_{i j}(u) J^{*}(j)\right] .
$$

In addition, there are analogs of the computational methods of Section 7.3, including value iteration, policy iteration, and linear programming. What is happening here is that essentially we have the equivalent of a discretetime discounted problem where the discount factor depends on $i$ and $u$.

We finally note that in some problems, in addition to the cost per unit time $g$, there is extra (instantaneous) one-stage cost $\hat{g}(i, u)$ that is incurred the time the control $u$ is chosen at state $i$, and is independent of the length of the transition interval. In this case, Bellman's equation takes the form

$$
\begin{equation*}
J^{*}(i)=\min _{u \in U(i)}\left[\hat{g}(i, u)+G(i, u)+{ }_{j=1}^{n} \operatorname{mij}(u) J^{*}(j)\right] \tag{7.51}
\end{equation*}
$$

and the various computational methods are appropriately adjusted. Another problem variation arises when the cost $g$ depends on the next state $j$. Here, once the system goes into state $i$, a control $u \in U(i)$ is selected, the next state is determined to be $j$ with probability $P_{i j}(u)$, and the cost incurred is $g(i, u, j)$. In this case, $\mathrm{G}(i, u)$ should be defined by

$$
G(i, u)={ }_{j=1}^{n} \int_{0}^{\infty} g(i, u, j)^{1-e^{-\beta \tau}} d \operatorname{Qij}(T, U),
$$

[cf. Eq. (7 and the preceding development goes through without modification.
7.5.1

Consider the manufacturer's problem of Example 7.3.2, with the only difference that the times between the arrivals of successive orders are uniformly distributed in a given interval [0,
and $c$ is the cost per unit time of an
unfilled order. Let $F$ and $N F$ denote the choices of filling and not filling the orders, respectively. The transition distributions are

$$
Q_{i j}(\tau, F)= \begin{cases}\min \left[1, \frac{\tau}{\tau_{\max }}\right] & \text { if } j=1, \\ 0 & \text { otherwise },\end{cases}
$$

and

$$
Q_{i j}(\tau, N F)= \begin{cases}\min \left[1, \frac{\tau}{\tau_{\max }}\right] & \text { if } j=i+1, \\ 0 & \text { otherwise }\end{cases}
$$

The one-stage expected cost $G$ of Eq. (7.47) is given by

$$
G(i, F)=0, \quad G(i, N F)=\gamma c i
$$

where

$$
\gamma=\left\lvert\, \begin{array}{lc}
T \max & 1-e^{-\beta \tau} \\
0 & \beta \tau_{\max }
\end{array} d T\right.
$$

The scalars ffiij of Eq. (7.49) that are nonzero are

$$
m_{i 1}(F)=m_{i(i+1)}(N F)=\alpha
$$

where

$$
\alpha=\int_{0}^{\operatorname{Tmax}} \frac{e^{-\beta \tau}}{\operatorname{Tmax}} d \tau=------
$$

Bellman's equation has the form [ef. Eq. (7.51)]

$$
\mathrm{J}(\mathrm{i})=\min [\mathrm{K}+\alpha J(1), \gamma c i+\alpha J(i+1) \mathbf{J}, \quad i=1,2, \ldots
$$

As in Example 7.3.2, we can conclude that there exists a threshold i* such that it is optimal to fill the orders if and only if their number i exceeds $\mathrm{i}^{*}$.

## Average Cost Problems

A natural cost function for the continuous-time average cost problem would be

$$
\begin{equation*}
\lim _{T \rightarrow \infty} 1 \quad\left\{\int_{0}^{T} g(x(t), u(t)) d t\right\} \tag{7.52}
\end{equation*}
$$

However, we will use instead the cost function

$$
\begin{equation*}
\lim _{N \rightarrow \infty} \frac{1}{E\left\{t_{N}\right\}} E\left\{\int_{0}^{t_{N}} g(x(t), u(t)) d t\right\} \tag{7.53}
\end{equation*}
$$

where $t N$ is the completion time of the Nth transition. This cost function is also reasonable and turns out to be analytically convenient. We note, however, that the cost functions (7.52) and (7.53) are equivalent under the conditions of the subsequent analysis, although a rigorous justification of
this is beyond our scope (see Ross [Ros70], p. 52 and p. 160 for related discussion).

For each pair $(i, u)$, we denote by $G(i, u)$ the one-stage expected cost corresponding to state $i$ and control $u$. We have

$$
G(i, u)=g(i, U) T i(U)
$$

where $\bar{\tau}_{i}(u)$ is the expected value of the transition time corresponding to $(i, u)$. the cost per unit time 9 depends on the next state $j$, the expected transition cost $G(i, u)$ should be defined by

$$
G(i, u)={ }_{j=1}^{n} \int_{\mathrm{Q}}^{\infty} g(i, u, j) \operatorname{TdQij}(T, u)
$$

and the following analysis and results go through without modification.] The cost function of an admissible policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots\right\}$ is given by
$J_{\pi}(i)=\lim _{N \rightarrow \infty} \frac{1}{E\left\{t_{N} \mid x_{0}=i, \pi\right\}} E\left\{\sum_{k=0}^{N-1} \int_{t_{k}}^{t_{k+1}} g\left(x_{k}, \mu_{k}\left(x_{k}\right)\right) d t \mid x_{0}=i\right\}$.
will see that character of the solution of the problem is determined by the structure of tife embedded Markov chain, which is the controlled discrete-time :Markov chain whose transition probabilities are

$$
p_{i j}(u)=\lim _{\tau \rightarrow \infty} Q_{i j}(\tau, u)
$$

In particular, assuming that the embedded Markov chain satisfies Assumption 7.4.1 of Section 7.4, we can show that the costs $J$ * (i) are independent of i.

It turns out that Bellman's equation for average cost semi-Markov problems takes the form

$$
h(i)=\min _{u \in U(i)}\left[G(i, u)-\lambda \bar{\tau}_{i}(u)+\sum_{j=1}^{n} p_{i j}(u) h(j)\right]
$$

As a special case, when $\bar{\tau}_{i}(u)=1$ for all $(i, u)$, we obtain the corresponding Bellman's equation for discrete-time problems, given in Section 7.4. We motivate the above form of Bellman's equation with the stochastic shortest path argument that we used in Section 7.4. We consider a sequence of generated states, and divide it into cycles marked by successive visits to the special state $n$. Each of the cycles can be viewed as a state trajectory of a corresponding stochastic shortest path problem with the termination state being essentially $n$, as in Section 7.4.

We next conjecture that the average cost problem is equivalent to the minimum cycle cost problem of finding a stationary policy $\mu$ that minimizes the average cycle cost

$$
\frac{C_{n n}(\mu)}{T_{n n}(\mu)}
$$

where for a fixed $\mu$,
$C n n(\mu)$ : expected cost starting from $n$ up to the first return to $n$,

$$
T n n(\mu): \text { expected time to return to } n \text { starting from } n
$$

An intuitive conjecture is that the optimal average cost $\lambda^{*}$ is equal to the optimal cycle cost, so it satisfies

$$
\begin{equation*}
C_{n n}(\mu)-N_{n n}(\mu) \lambda^{*} \geq 0, \quad \text { for all } \mu \tag{7.54}
\end{equation*}
$$

with equality holding if $\mu$ is optimal. Thus, to attain an optimal $\mu$, we must minimize over $\mu$ the expression $C_{n n}(\mu)-T_{n n}(\mu) \lambda^{*}$, which is the expected cost of $\mu$ starting from $n$ in the associated stochastic shortest path problem with stage costs

$$
G(i, u)-\lambda^{*} \bar{\tau}_{i}(\mu(i)), \quad \mathrm{i}=1, \ldots, n o
$$

Let us denote by $h^{*}(\mathrm{i})$ the optimal cost of this stochastic shortest path problem when starting at state i. Then $h^{*}(I), \ldots, h^{*}(n)$ solve uniquely the corresponding Bellman's equation, which has the form

$$
\begin{equation*}
h^{*}(i)=\min _{u \in U(i)}\left[G(i, u)-\lambda^{*} \bar{\tau}_{i}(u)+\sum_{j=1}^{n-1} p_{i j}(u) h^{*}(j)\right], \quad i=1, \ldots, n o \tag{7.55}
\end{equation*}
$$

If $\mu^{*}$ is an optimal stationary policy, then this policy must satisfy

$$
C_{n n}\left(\mu^{*}\right)-N_{n n}\left(\mu^{*}\right) \lambda^{*}=0
$$

and from (7.54), this policy must also be optimal for the associated stochastic shortest path problem. Thus, we must have

$$
h^{*}(n)=C_{n n}\left(\mu^{*}\right)-N_{n n}\left(\mu^{*}\right) \lambda^{*}=0
$$

By using this equation, we can now write Bellman's equation (7.55) as

$$
h^{*}(i)=\min _{u \in U(i)}\left[G\left(i^{\prime} U\right)-\lambda^{*} \bar{\tau}_{i}(u)+t P i j(U) h^{*}(j)\right], \quad i=1, \ldots, n o
$$

(7.56)

If there is an "instantaneous" one-stage cost $\hat{g}(i, u)$, the term
u) should be replaced by $\hat{g}(i, u)+G(i, u)$ in this equation.

Given the correct form of Bellman's equation and the connection with the associated stochastic shortest path problem, it is possible to essentially repeat the proof of Prop. 7.4.1 and to obtain analogous results to those for the discrete-time case.

Example 7.5.2
Consider the average cost version of the manufacturer's problem of Example 7.5.1. Here we have

$$
\bar{\tau}_{i}(F)=\bar{\tau}_{i}(N F)=\frac{\bar{\tau}_{\max }}{n}
$$

$$
G(i, F)=K, \quad G(i, N F)=\mathrm{CiT}^{\max }
$$

where $F$ and $N F$ denote the decisions to fill and not fill the orders, respectively. Bellman's equation (7.56) takes the form

$$
h^{*}(i)=\min \left[K-\lambda^{*} \frac{\bar{\tau}_{\max }}{2}+h^{*}(1), c i \frac{\bar{\tau}_{\max }}{2}-\lambda^{*} \frac{\bar{\tau}_{\max }}{2}+h^{*}(i+1)\right]
$$

We leave it as an exercise for the reader to show that there exists a threshold $i^{*}$ such that it is optimal to fill the orders if and only if i exceeds $i^{*}$.

## Example 7.5.3 [LiR71]

Consider a person providing a certain type of service to customers. Potential customers arrive according to a Poisson process with rate $r$; that is, the customer's interarrival times are independent and exponentially distributed with parameter $r$. Each customer offers one of $n$ pairs ( $m^{\prime} i^{\prime} T_{i}$ ), i $=1, \ldots, n$, where $m i$ is the amount of money offered for the service and $T_{i}$ is the average amount of time that will be required to perform the service. Successive offers are independent and offer $\left(m i^{\prime} T_{i}\right)$ occurs with probability $p i$, where $\sum_{i=1}^{n} P i=1$. An offer may be rejected, in which case the customer leaves, or may be accepted in which case all offers that arrive while the customer is being served are lost. The problem is to determine the acceptance-rejection policy that maximizes the service provider's average income per unit time.

Let us denote by $i$ the state corresponding to the offer ( $m i^{\prime} T i$ ), and let $A$ and $R$ denote the accept and reject decision, respectively. We have

$$
\begin{array}{cc}
=T_{i}+\frac{1}{r}, \quad \bar{\tau}_{i}(R)=\frac{1}{r}, \quad G(i, & G(i, R)=0 \\
p_{i j}(A)=p_{i j}(R)=p_{j} &
\end{array}
$$

Bellman's equation is given by

$$
h^{*}(i)=\min \left[-m_{i}-\lambda^{*}\left(T_{i}+\frac{1}{r}\right)+\sum_{j=1}^{n} p_{j} h^{*}(j),-\lambda^{*} \frac{1}{r}+\sum_{j=1}^{n} p_{j} h^{*}(j)\right]
$$

It follows that an optimal policy is to

$$
\text { accept offer (i, } \quad \text { if and only if } \quad \frac{m_{i}}{T_{i}} \geq-\lambda^{*}
$$

where $-\lambda^{*}$ is the optimal average income per unit time.

### 7.6 NOTES, SOU^uル. AND EXERCISES

This chapter is only an introduction to infinite horizon problems. There is an extensive theory for these problems with interesting mathematical and computational content. Volume II provides a comprehensive treatment and gives many references to the literature.

The presentation in this chapter is original in that it uses the stochastic shortest path problem as the starting point for the analysis of the other problems. This line of development not only explains intuitively the connections between the various types of problems, but also leads to new solution methods. For example, an alternative value iteration algorithm for the average cost problem, based on the connection with the stochastic shortest path problem, is given in Bertsekas [Ber98b], and in Section 4.3 of Vol. II. On the other hand, there are also important results for undiscounted and average cost problems that cannot be obtained through the connection with the stochastic shortest path problem. Some of these alternative lines of analysis are pursued in Vol. II.

Semi-Markov problems were introduced by Jewell [Jew63] and were also discussed by Ross [Ros70]. Volume II contains a broader exposition of Semi-Markov problems, and applications to queueing and related systems.

## EXERCISES

## 7.1

A tennis player has a Fast serve and a Slow serve, denoted $F$ and $S$, respectively. The probability of $F$ (or $S$ ) landing in bounds is $P F$ (or $p s$, respectively). The probability of winning the point assuming the serve landed in bounds is $q F$ (or $q s$, respectively). We assume that $P F<P S$ and $q F>q s$. The problem is to find the serve to be used at each possible scoring situation during a single game in order to maximize the probability of winning that game.
(a) Formulate this as a stochastic shortest path problem, argue that Assumption 7.2 .1 of Section 7.2 holds, and write Bellman's equation.
(b) Computer assignment: Assume that $q F \quad 0.6$, $q s=0.4$, and $P S=0.95$. Use value iteration to calculate and plot (in increments of 0.05 ) the probability of the server winning a game with optimal serve selection as a function of $P F$.

A quarterback can choose between running and passing the ball on any given play. The number of yards gained by running is integer and is Poisson distributed with parameter $\lambda_{r}$. A pass is incomplete with probability $p$, is intercepted with probability $q$, and is completed with probability $1-p-q$. When completed, a pass gains an integer number of
is Poisson distributed with parameter $\lambda_{p}$. We assume that the probability of scoring a touchdur on a single play starting i yards from the goal is equal to the probability of gainis? a number of yards greater than or equal to i. We assume also that yardage cannot lost on any play and that there are no penalties. The ball is turned over to the ther team on a fourth down or when an interception occurs.
(a) Formulate the problem as a stochastic shortest path problem, argue that Assumption 7.2.1 of Section 7.2 holds, and write Bellman's equation.
(b) Computer assignment: Use value iteration to compute the quarterback's play-selection policy that maximizes the probability of scoring a touchdown on any single drive for $\lambda_{r}=3, \lambda_{p}=10, p=0.4$, and $q=0.05$.

## 7.3

A computer manufacturer can be in one of two states. In state 1 his product sells well, while in state 2 his product sells poorly. While in state 1 he can advertise his product in which case the one-stage reward is 4 units, and the transition probabilities are Pll $=0.8$ and $P 12=0.2$. If in state 1 , he does not advertise, the reward is 6 units and the transition probabilities are $\mathrm{P} 11=\mathrm{P} 12=0.5$. While in state 2 , he can do research to improve his product, in which case the onestage reward is -5 units, and the transition probabilities are $P 21=0.7$ and $P 22=0.3$. If in state 2 he does not do research, the reward is -3 , and the transition probabilities are $P_{P 21}=0.4$ and ${ }_{P 22}=0.6$. Consider the infinite horizon, discounted version of this problem.
(a) Show that when the discount factor $\alpha$ is sufficiently small, the computer manufacturer should follow the "shortsighted" policy of not advertising (not doing research) while in state 1 (state 2 ). By contrast, when $\alpha$ is sufficiently close to unity, he should follow the "farsighted" policy of advertising (doing research) while in state 1 (state 2 ).
(b) For $\alpha \quad 0.9$ calculate the optimal policy using policy iteration.
(c) For $\alpha=0.99$, use a computer to solve the problem by value iteration, with and without the error bounds (7.23).

## 7.4

An energetic salesman works every day of the week. He can work in only one of two towns $A$ and B on each day. For each day he works in town $A$ (or B ) his expected reward is $r_{A}$ (or $r_{B}$, respectively). The cost for changing towns is c . Assume that $\mathrm{c}>r_{A}>r_{B}$ and that there is a discount factor $\alpha<1$.
(a) Show that for $\alpha$ sufficiently small, the optimal policy is to stay in the town he starts in, and that for $\alpha$ sufficiently close to 1 , the optimal policy is to move to town $A$ (if not starting there) and stay in $A$ for all subsequent times.
(b) Solve the problem for $\mathrm{c}=3, r_{A}=2, r_{B}=1$, and $\alpha=0.9$ using policy iteration.
(c) Use a computer to solve the problem of part (b) by value iteration, with and without the error bounds (7.23).

## 7.5

A person has an umbrella that she takes from home to office and vice versa. There is a probability $p$ of rain at the time she leaves home or office independently of earlier weather. If the umbrella is in the place where she is and it rains, she tanes the umbrella to go to the other place (this involves no cost). If there is no umbrelie and it rains, there is a cost $W$ for getting wet. If the umbrella is in the place where she is but it does not rain, she may take the umbrella to go to the other place (thi involves an inconvenience cost $V$ ) or she may leave the umbrella behind (this involves no cost). Costs are discounted at a factor $\alpha<1$.
(a) Formulate this as an infinite horizon total cost discounted problem. Hint: Try to use as few states as possible.
(b) Characterize the optimal molicy in Lesi as you can.
7.6

For the tennis player's problem (Exercise 7.1), show that it is optimal (regardless of score) to use $F$ on both serves if

$$
\left(p_{F} q_{F}\right) /\left(p_{S} q_{S}\right)>1,
$$

to use $S$ on both serves if

$$
(P F q F) /(P S q S)<1+P F-P s
$$

and to use $F$ on the first serve and $S$ on the second otherwise.

## 7.7

Consider the value iteration method for the Example 7.3.2:

$$
\begin{aligned}
J k+1(i)=\min \quad & +\alpha(1-p) J k(O)+\operatorname{apJ} J(l), \\
& c i+a(1-p) J k(i)+\operatorname{apJ} J(i+1)], \quad O, l, \ldots, n \quad l,
\end{aligned}
$$

$$
(n)=K+\alpha(1-p) J k(O)+\operatorname{apJ} J(l)
$$

where $J o(i)=0$ for all $i$. Show by induction that $J k(i)$ is monotonically nondecreasing in i.

Consider the policy iteration algorithm for the problem of Example 7.3.2.
(a) Show that if we start the algorithm with a threshold policy, every subsequently generated policy will be a threshold policy. Note: This requires a careful argument.
(b) Carry out the algorithm for the case $\mathrm{c}=1, K=5, n=10, P=0.5$, $\alpha=0.9$, and an initial policy that always processes the unfilled orders.

## 7.9

Solve the average cost version ( $\alpha=1$ ) of the computer manufacturer's problem by using value iteration and by using policy iteration (Exercise 7.3).

### 7.10

An unemployed worker receives a job offer at each time period, which she may accept or reject. The offered salary takes one of $n$ possible values $\mathrm{WI}, \ldots, w^{n}$ with given probabilities, independently of preceding offers. If she accepts the offer, she must keep the job for the rest of her life at the same salary level. If she rejects the offer, she receives unemployment compensation $\rho^{f}$, whe current period ana is eligible to accept future offers.
mcome is discounted by a factor
$\alpha<1$.
(a) Show is a threshold $\bar{w}$ such that it is optimal to accept an offer
and only if its salary is larger than $\bar{w}$, and characterize $\bar{w}$. and only if its salary is larger than $\bar{w}$, and characterize $\bar{w}$.
(b) Consider the variant of the problem where there is a given probability $P i$ that the worker will be fired from her job at anyone period if her salary is Wi. Show that the result of part (a) holds in the case where $P i$ is the same for all i. Analyze the case where $P i$ depends on i.

### 7.11

Do part (b) of Exercise 7.10 for the case where income is not discounted and the worker maximizes her average income per period.

Show that one can always take $m=n$ in Assumption 7.2.1. Hint: For any $\pi$ and i, let $S k$ (oi) be the set of states that are reachable with positive probability from i under $\pi$ in $k$ stages or less. Show that under Assumption 7.2.1, we cannot have $S k(i) \quad S k+1$ (i) while $t \neq S k(\mathrm{i})$.

### 7.13

Show the error bounds (7.17). These bounds constitute a generalization to the stochastic shortest path problem of the bounds (7.23) for the discounted problem, which have a long history, starting with the work of McQueen [McQ66]. Hint: Complete the details of the following argument. Let $\mu^{k}$ (i) attain the minimum in the value iteration (7.16) for all i. Then, in vector form, we have

$$
J k+1=g k+P k J k,
$$

where $J k$ and $g k$ are the vectors with components $J k(i), i=1, \ldots, n$, and $g k\left(i, \mu^{k}(i)\right), i=1, \ldots, n$, respectively, and $P k$ is the matrix whose components are the transition probabilities $P_{i j}\left(\mu^{k}(\mathrm{i})\right)$. Also from Bellman's equation, we have

$$
J^{*} \leq g k+P k J^{*}
$$

where the vector inequality above is meant to hold separately for each component. Let $\mathrm{e}=(1, \ldots, 1)^{\prime}$. Using the above two relations, we have

$$
\begin{equation*}
J^{*}-J_{k} \leq J^{*}-J_{k+1}+\bar{c}_{k} e \leq P_{k}\left(J^{*}-J_{k}\right)+\bar{c}_{k} e \tag{7.57}
\end{equation*}
$$

Multiplying this relation with $P k$ and adding $\bar{c}_{k} e$, we obtain

$$
P_{k}\left(J^{*}-J k\right)+\bar{c}_{k} e \leq P_{k}^{2}\left(J^{*}-J k\right)+\bar{c}_{k}(I+P k) e
$$

Similarly continuing, we have for all $r \geq 1$

$$
J^{*}-J k+1+\bar{c}_{k} e \leq P_{k}^{r}\left(J^{*}-J k\right)+\bar{c}_{k}\left(I+P k+\ldots+P_{k}^{r-1}\right) e
$$

For $s=1,2, \ldots$, the ith component of the vector $P_{k}^{s} e$ is equal to the probability $\left.P_{\left\{x_{s}\right.} \neq \mathrm{t} \mid X_{0}=\mathrm{i}, \mu^{k}\right\}$ that $t$ has not been reached after $s$ stages starting from i anu' 'Ising the stationary

Thus, Assumption 7.2.í İmplíes that $\lim _{r \rightarrow \infty} P_{k}^{r}=0$, wilie we have

$$
\lim _{r \rightarrow \infty}\left(I+P k+\ldots+P_{k}^{r-1}\right) e=N^{k}
$$

where $N^{k}$ is the vector $\left(N^{k}(I), \ldots, N k(n)\right)^{\prime}$. Combining the above two relations, we obtain

$$
J^{*} \leq J_{k+1}+\bar{c}_{k}\left(N^{k}-e\right)
$$

proving the desired upper bound.
The lower bound is proved similarly, by using in place of $\mu^{k}$, an optimal stationary policy $\mu^{*}$. In particular, in place of Eq. (7.57), we can show that

$$
J k-J^{*} \leq J k+1-J^{*}-\underline{c}_{k} e \leq P^{*}\left(J k-J^{*}\right)-\underline{c}_{k} e
$$

where $P^{*}$ is the matrix with elements $P^{\prime} i j\left(\mu^{*}(i)\right)$. We similarly obtain for all $r \geq 1$

$$
\mathrm{Jk}+1-J^{*}-\underline{c}_{k} e \leq\left(D^{*} r\left(J k-J^{*}\right)-c_{k}\left(I+P^{*}+\ldots+\left(D^{*} r^{l}\right) \mathrm{e}\right.\right.
$$

from which $\mathrm{Jk}+1+{\underset{c}{c}}\left(N^{*} \quad e\right) \leq J^{*}$, where $N^{*}$ is the vector $\left(N^{*}(1), \ldots, N^{*}(n)\right)^{\prime}$.

Apply the relative value iteration algorithm (7.41) for the case where there are two states and only one control per state. The transition probabilities are $\mathrm{Pll}=\epsilon$, $P 12=1-\epsilon, P 21=1-\epsilon$, and $P 22=\epsilon$, where $0 \leq \epsilon<1$. Show that if $0<\epsilon$ the algorithm converges, but if $\epsilon=0$, the algorithm may not converge. Show also that the variation (7.42) converges when $\boldsymbol{\epsilon}=0$.

### 1.15

Consider the average cost problem and its associated stochastic shortest path problem when the expected cost incurred at state i is $9(i, u)-\lambda^{*}$.
(a) Use Prop. 7.2.1(d) and Prop. 7.4.1(a) to show that if a stationary policy is optimal for the latter problem it is also optimal for the former.
(b) Show by example that the reverse of part (a) need not be true.

### 7.16

Consider a problem of operating a machine that can be in anyone of $n$ states, denoted $1,2, \ldots, n$. We denote by $g(i)$ the operating cost per period when the machine is in state $i$, and we assume that

$$
g(l) \leq g(2) \leq \ldots \leq g(n)
$$

The implication here is that state $i$ is better than state $i+1$, and state 1 corresponds to a machine being in the best condition. The transition probabilities during one period of operation satisfy

$$
\begin{gathered}
P^{\prime} i(i+l)>0 \quad \text { if } \mathrm{i}<n, \\
P i j=0 \quad \text { if } \mathrm{j} \neq \mathrm{i}, \mathrm{j} \neq \mathrm{i}+1
\end{gathered}
$$

We assume that at the start of each period we know the state of the machine and we must choose one of the following two options:
(1) Let the machine operate one more period in the state it currently is.
(2) Repair the machine and bring it to the best state 1 at a cost $R$.

We assume that the machine, once repaired, is guaranteed to stay in state 1 for one period. In subsequent periods, it may deteriorate to states $\mathrm{j}>1$.
(a) Assume an infinite horizon and a discount factor $\alpha E(0,1)$, and show that there is an optimal policy which is a threshold policy; that is, it takes the form

$$
\text { replace if and only if } \mathrm{i} \geq \mathrm{i}^{*}
$$

where $i^{*}$ is some integer.
(b) Show that the policy iteration method, when started with a threshold policy, generates a sequence of threshold policies.

### 1.11

Consider a person providing a certain type of service to customers. The person receives at the beginning of each time period with probability $P i$ a proposal by a customer of type i , where $\mathrm{i}=1,2, \ldots, n$, who offers an amount of money Mi . We assume that $\sum_{i=1}^{n} P i=1$. The person may reject the offer, in which case the customer leaves and the person remains idle during that period, or the person may accept the offer in which case the person spends some random amount of time with that customer. In particular, we assume that the probability that the type i customer will leave after $k$ periods ( $k=1,2, \ldots$ ), given that the customer has already stayed with the person for $k-1$ periods is a given scalar $f 3 \mathrm{E}(0,1)$. The problem is to determine an acceptance-rejection policy that maximizes
(a) Formulate the person's problem as an average cost per stage problem, and show that the optimal cost is independent of the initial state.
(b) Show that there exists a scalar $\lambda$ and an optimal policy that accepts the offer of a type $i$ customer if and only if
where $T_{i}$ is the expected time spent with the type i customer.

### 7.18

A person has an asset to sell for whe she receives offers that take one $n$ values $s_{j}, \mathrm{j}=1, \ldots, n$. The times between successive offers are random, identically distributed, and independent of preceding times. Let $Q j(T)$ be the probability that the time between successive offers is less or equal to $T$ and the next offer is $s j$. Find the offer acceptance policy that maximizes $E\left\{\alpha^{T} s\right\}$, where $T$ is the time of sale, $s$ is the sale price, and $\alpha E(0,1)$ is a discount factor.

### 7.19

An unemployed worker receives job offers, which she may accept or reject. The times between successive offers are independent and exponentially distributed with parameter $r$. The offered salary (per unit time) takes one of $n$ possible values $w i, \mathrm{i}=1, \ldots, n$, with given probabilities $p i$, independently of preceding offers. If she accepts an offer at salary $w i$, she keeps the job for a random amount of time that has expected value $t_{i}$. If she rejects the offer, she receives unemployment compensation c (per unit time) and is eligible to accept future offers. Solve the problem of maximizing the worker's average income per unit time.

Consider a computing system where the interarrival times of the jobs are independent and exponentially distributed with parameter $r$. A job may be rejected, in which case the job is lost, or may be accepted in which case all jobs that arrive while the job is being processed are lost. There are $n$ types of jobs. Each arriving job is of type i with probability $p i$, independently of earlier jobs, and if processed, is worth a fixed positive benefit $b_{i}(\mathrm{i}=1, \ldots, n)$. Jobs of type i require an average amount of time $T i$ to complete processing. The problem is to determine the acceptance-rejection policy that maximizes the system's average benefit per unit time.
(a) Argue that the analysis of Example 7.5.3 applies for this problem.
(b) Calculate the optimal average benefit per unit time $\lambda^{*}$ in terms of the given quantities for the case where there are only two job types.
(c) Suppose that the time to process a job of type $i$ is exponentially distributed with mean $T_{i}$. Assume further that the system can process up to a given number $m>1$ of jobs simultaneously (rather than just one). Formulate the problem as an average benefit per unit time semi-Markov problem, and write Bellman's equation for the case where $\mathrm{m}=2$. Why do we need the exponential distribution assumption?

### 7.21

Formulate a semi-Markov version of the stochastic shortest path problem of Section 7.2. The cost function has the form

$$
\lim _{T \rightarrow \infty} E\left\{\prod_{0}^{I} 9(X(t), U(t)) d t\right\}
$$

and there is a cost-free and absorbing state. Use the transition distributions $Q^{i j}$ to formulate an assumption that is analogous to Assumption 7.2.1. Under this assumption, state and justify a result that parallels Prop. 7.2.1.

### 7.22

A treasure hunter has obtained a lease to search a site that contains $n$ treasures, and wants to find a searching policy that maximizes his expected gain over an infinite number of days. At each day, knowing the current number of treasures not yet found, he may decide to continue searching for more treasures at a cost c per day, or to permanently stop searching. If he searches on a day when there are $i$ treasures on the site, he finds $\mathrm{m} \mathrm{E}[0, \mathrm{i}]$ treasures with given probability $p\left(m_{\mathrm{I}} \mathrm{i}\right)$, where we assume that $p\left(O_{1} \mathrm{i}\right)<1$ for all $\mathrm{i} \geq 1$, and that the expected number of treasures found,
decreases monotonically with i. Each found treasure is worth 1 unit.
(a) Formulate the problem as an infinite horizon DP problem.
(b) Write Bellman's equation. How do you know that this equation holds and has a unique solution?
(c) Start policy iteration with the policy that never searches. How many policy iterations does it take to find an optimal policy, and what is that optimal policy?

The latest slot machine model has three arms, labeled 1,2 , and 3 . A single play with arm i , where $\mathrm{i}=1,2,3$, costs Ci dollars, and has two possible outcomes: a "win," which occurs with probability pi, and a "loss," which occurs with probability 1 - Pi. The slot machine pays you $m$ dollars each time you complete a sequence of three successive "wins," with each win obtained using a different arm.
(a) Consider the problem of finding the arm-playing order that minimizes the expected cost if you are restricted to stop at the first time the machine pays you. Formulate this problem as a stochastic shortest path problem where arm-playing orders are identified with stationary policies, and write Bellman's equation for each stationary policy.
(b) Show that the expected cost of the arm-playing order ABC is

$$
c_{A}+\frac{p_{A} c_{B} \neq p_{A} p_{B} c_{C}-p_{A} p_{B} p_{C} m}{1-P A P B P C}
$$

Show that it is optimal to play the arms in order of decreasing ci/(1-Pi)'
(c) Consider the problem of finding the arm-playing order that minimizes the average expected cost per play, assuming you play infinitely many times. Formulate this problem as an average cost per stage problem, where armplaying orders are identified with stationary policies, and write Bellman's equation for each stationary policy.
(d) Show that the expected cost per play of the arm-playing order ABC is

$$
\underline{C A} \pm \frac{P A C B}{ \pm P A P B C C} \frac{p A p B P c m}{1+P A P B}
$$

Is it possible that the optimal playing order is different than the one of part (b)? If this is so, how do you explain it?

A person has a house that he rents at a fixed amount $R$ per time period. At the beginning of each period $k$, the person receives an offer $w k$ to sell the house. The amount $w k$ takes one of m given values $\mathrm{w}_{\mathrm{I}}, \ldots, w^{m}$, with corresponding positive probabilities $q 1, \ldots, q m$, independently of preceding offers. The person, at the
beginning of each period, must decide whether to accept the current offer or to decline the offer and continue to rent the house.

Upon selling the house, the sale amount, call it $w$, is immediately reinvested in some way so that it yields at time $k$ a random amount $Y k W$, where $Y k$ takes one of 8 given values $y^{1}, \ldots, Y^{s}$. The value of $Y k$ evolves according to a Markov chain with a single recurrent class and given transition probabilities

$$
p_{i j}=P\left(y_{k+1}=y^{j} \mid y_{k}=y^{i}\right), \quad \mathrm{i}, j=1, \ldots, 8 .
$$

(a) Suppose that at time $k$ the house is sold when $Y k$ is equal to Vi. Let

$$
\bar{y}(i)=\lim _{N \rightarrow \infty} \frac{1}{N} E\left\{\sum_{l=k}^{k+N-1} y_{l} \mid y_{k}=y^{i}\right\}
$$

be the average future yield per unit time. Show that $\bar{y}(\mathrm{i})$ is equal to a common value $\bar{y}$, independent of $i$, and derive a Bellman-type equation for this value.
(b) Suppose that the person's objective is to maximize the average monetary benefit per time .. Aryue that an optImal policy is to wait _.avi tne maximum possible offer $\bar{w}=\max \left\{w^{1}, \ldots, w^{m}\right\}$ is received and then sell the house, assuming that $R / \bar{y} \leq \bar{w}$. Given this result, discuss whether the average cost formulation is satisfactory for this problem.
(c) Suppose that the person's objective is to maximize the total discounted monetary benefit over an infinite horizon, with a discount factor $\alpha<\mathrm{I}$. Show that for each $\mathrm{i}=1, \ldots, 8$, there is a threshold $t(i)$ such that it is optimal to sell the house at period $k$ when $Y k=y^{i}$ and the current offer is larger than $\mathrm{t}(\mathrm{i})$.

### 7.25

You have just bought your first car, which raises the issue of where to park it. At the beginning of each day you may either park it in a garage, which costs $G$ per day, or on the street for free. However, in the latter case, you run the risk of getting a parking ticket, which costs $T$, with probability $P j$, where $j$ is the number of consecutive days that the car has been parked on the street (e.g., on the first day you park on the street, you have probability Pl of getting a ticket, on the second successive day you park on the street, you have probability $P 2$, etc). Assume that $P_{j}$ is monotonically nondecreasing in $\mathbf{j}$, and that you may receive at most one ticket per day when parked on the street. Assume also that there exists an integer $m$ such that $p m T>G$.
(a) Formulate this as an infinite horizon discounted cost problem with finite state space and write the corresponding Bellman's equation.
(b) Characterize as best as you can the optimal policy.
(c) Let $n$ be the total number of states. Show how to use policy iteration so that it terminates after no more than $n$ iterations. Hint: Use threshold policies as in Problem 7.8.
(d) Formulate the infinite horizon average cost version of this problem with finite state space and write the corresponding Bellman's equation. State an assumption under which Bellman's equation holds.

### 7.26

An engineer has invented a better mouse trap and is interested in selling it for the right price. At the beginning of each period, he receives a sale offer that takes one of the values $81, . ., 8 n$ with corresponding probabilities $P 1, \ldots, p n$, independently of prior offers. If he accepts the offer he retires from engineering. If he refuses the offer, he may accept subsequent offers but he also runs the risk that a competitor will invent an even better mouse trap, rendering his own unsaleable; this happens with probability $\{3>0$ at each time period, independently of earlier time periods. While he is overtaken by the competitor, at each time period, he may choose to retire from engineering, or he may choose to invest an amount $v \geq 0$, in which case he has a probability $\gamma$ to improve his mouse trap, overtake his competitor, and start receiving offers as earlier. The problem is to determine the engineer's strategy to maximize his discounted expected payoff (minus investment cost), assuming a discount factor $\alpha<1$.
(a) Formulate the problem as an infinite horizon discounted cost problem and write the corresponding Bellman's equation.
(b) Characterize as best as you can an optimal policy.
(c) Assume that there is no discount factor. Does the problem make sense as an average cost per stage problem?
(d) Assume that there is no discount factor and that the investment cost $v$ is equal to 0 . Does the problem make sense as a stochastic shortest path problem, and what is then the optimal policy?

### 7.27 Eilminatmg Self-Transitions)

Consider a stochastic shortest path problem (SSP) with termination state $t$, the nontermination states $1, \ldots, n$, transition probabilities $\operatorname{Pij}(U)$, and expected costs per stage $g(i, u)$. Let Assumption 7.2.1 hold.
(a) Modify the costs and transition probabilities as follows:

$$
\begin{gathered}
\tilde{g}(i, u)=\frac{g(i, u)}{1-P i i(U)}, \quad i=l, \ldots, n, U E U(i), \\
\tilde{p}_{i j}(u)=\left\{\begin{array}{ll}
0 & \text { if } j=\mathrm{i}, \\
\overline{\mathrm{I}} & \text { if } j \neq \mathrm{i},
\end{array} \quad \mathrm{i}=1, \ldots, n, j=1, \ldots, n, t, \quad \cup \mathrm{E} U(i),\right.
\end{gathered}
$$

to obtain another SSP without self-transitions. Show that the modified SSP is equivalent to the original in the sense that its stationary policies and optimal policies have the same cost functions. What is the interpretation of the transitions of the modified SSP in terms of transitions of the original?
(b) Fix a policy $\mu$. Let $J k$ and $\mathbf{J} k$ be the sequences of cost vectors generated by value iteration (for the fixed policy) in the original and the modified SSP, respectively, starting from the same initial vector $J_{O}$. Show that value iteration is faster for the modified SSP in the sense that if $J o \leq J_{1}$, then $J k \leq J k \leq J^{*}$ for all $k$, and if $J o \geq J_{1}$, then $J k \geq J k \geq J^{*}$ for all $k$.

### 7.28 Cost Problems with Nonnegative Costs)

This is a theoretical problem whose purpose is to provide some additional analysis for undiscounted cost problems, including an extension of the results of Section 7.2 for stochastic shortest problems. The idea is to use the analysis of Section 7.3 for discounted problems to curive the basic results for total undiscounted cost problems under the assumption that the stage costs are nonnegative and the optimal costs are finite. These results apply, ang others, to some stochastic shortest path problems where not all stationary policies are proper and Assumption 7.2.1 is violated

Consider a controlled Markov chain with states $i=1, \ldots, n$, controls $u$ chosen from a finite constraint set $U(i)$ for each state $i$, and transition probabilities $P i j(U)$. (The states may include a cost-free and absorbing termination state, but this is not relevant for the following analysis.) The cost of the kth stage at state i when control $u$ is applied has the form

$$
\alpha^{n} g(i, u) ; \quad \quad, \ldots, \quad \cup \mathrm{E} U(i)
$$

where $\alpha$ is a scalar from $(0,1]$. Our key assumption is that

$$
0 \leq g(i, u), \quad \mathrm{i}=1, \ldots, n, u \mathrm{E} U(i)
$$

For any policy $\pi$, let $J_{\pi, \alpha}$ be the cost function for the a-discounted problem ( $\alpha<1$ ), and let $J_{\pi}$ be the cost function for the problem where $\alpha=1$. Note that for $\alpha \quad 1$, we may have $J_{\pi}(i)=$ oo for some $\pi$ and i. However, the limit defining $J_{\pi}(i)$ exists either as a real number or oo, thanks to the assumption $0 \leq g(i, u)$ for all $i$ and $u$. Let $J_{\alpha}^{*}(i)$ and $J^{*}(i)$ be the optimal costs starting from i , when $\alpha<1$ and $\alpha=1$, respectively. We assume that

$$
J^{*}(\mathrm{i})<00, \quad \mathrm{i}=1, \ldots, n,
$$

(this is true in particular for the case of a stochastic shortest path problem if there exists a proper stationary policy, Le., a policy under which there is a positive transition probability path from every state to the termination state).
(a) Show that for all $\alpha<1$, we have

$$
0 \leq J_{\alpha}^{*}(i) \leq J^{*}(\mathrm{i}), \quad \mathrm{i}=1, \ldots, n o
$$

(b) Show that for any admissible policy $\pi$, we have

$$
\lim _{\alpha \uparrow 1} J_{\pi, \alpha}(\mathrm{i})=J_{\pi}(i), \quad \mathrm{i}=1, \ldots, n o
$$

Furthermore,

$$
\left.\lim _{\alpha \uparrow 1} J_{\alpha}^{*}(i)=r e i\right), \quad \mathrm{i}=1, \ldots, n o
$$

Hint: To show the first equality, note that for any $\alpha<1, N$, and $\pi$ $\left\{\mu_{0}, \mu_{1}, \ldots.\right\}$, we have

$$
J_{\pi}(i) \geq J_{\pi, \alpha}(i) \geq \sum_{k=O}^{N-1} \alpha^{k} E\left\{g\left(i k^{\prime} \mu_{k}\left(i_{k}\right)\right) \mathrm{l}_{\mathrm{o}}=\mathrm{i}, \pi\right\} .
$$

Take the limit as $\alpha \rightarrow 1$ and then take the limit as $N \rightarrow 00$. For the second equality, consider a stationary policy $\mu$ and a sequence $\left\{a_{m}\right\} C(0,1)$ with $\alpha_{m} \rightarrow 1$ such that $J_{\mu, \alpha_{m}}=J_{\alpha_{m}}^{*}$ for all m
(c) Use Bellman's equation for $\alpha<1$, to show that $J^{*}$ satisfies Bellman's equation for $\alpha=1$ :

$$
r(i)=\min _{u \in U(i)}\left[9\left(i^{\prime} U\right)+\sum_{j=l}^{n} \operatorname{Pij}(u) J^{*}(j)\right] \quad i=1, \ldots, n o
$$

(d) Let J be such that $0 \leq J(i)<0$ for all i. Show that if

$$
\left.J(i) \geq \min _{u \in U(i)}[g(j, u)+t P i j(U)](j)\right], \quad \mathrm{i}=1, \ldots, n,
$$

then $J\left(\right.$ i) $\geq J^{*}$ (i) for all i. Show also that if for some stationary policy $\mu$, we have

$$
J(i) \geq g(i, \mu(i))+\sum_{j=l} p_{i j}(\mu(i)) \tilde{J}(j), \quad i=1, \ldots, n
$$

then $J(i) \geq J_{\mu}(i)$ for all i. Hint: Argue that

$$
\left.J(i) \geq \min _{u \in U(i)}\left[9\left(i^{\prime} u\right)+\alpha \sum_{j=l}^{n} P i j(U) J U\right)\right] \quad \mathrm{i}=1, \ldots, n
$$

use value iteration to show that $\mathbf{J} \geq J_{\alpha}^{*}$, and take the limit as $\alpha \rightarrow \mathbf{I}$.
(e) For $\alpha=1$, show that if

$$
\left.\mu^{*}(i)=\arg \min _{u \in U(i)}\left[9\left(i^{\prime} U\right)+t P i j(U) J^{*} U\right)\right], \quad \mathrm{i}=1, \ldots, n,
$$

$$
)=1
$$

then $\mu^{*}$ is optimal. Hint: Use part (d) with $J=J^{*}$
(f) For $\alpha=1$, show that for the value iteration method, given by

$$
J_{k+1}(i)=\min _{u \in U(i)}\left[g(i, u)+\sum_{j=1}^{n} p_{i j}(u) J_{k}(j)\right], \quad \mathrm{i}=1, \ldots, n
$$

we have $J_{k}(i) \rightarrow J^{*}(i)$, $\mathrm{i}=1, \ldots, n$, assuming that

$$
0 \leq \mathrm{Jo}(\mathrm{i}) \leq \mathrm{J}^{*}(\mathrm{i}), \quad \mathrm{i}=1, \ldots, \text { no }
$$

Give examples showing what may happen when this last assumptiol is violated. Hint: Prove the result by first assuming that $J_{0}$ is the zeru function.
(g) Show that the set of states $Z=\left\{\mathrm{i} \mid J^{*}(\mathrm{i})=0\right\}$ is nonempty. Furthermore, under an optimal stationary policy $\mu^{*}$, the set of states $Z$ is cost-free and absorbing, Le., $g\left(i, \mu^{*}(i)\right)=0$ and $p_{i j}\left(\mu^{*}(i)\right)=0$ for all i.E $Z$ and $\mathrm{j} \notin Z$. In addition, is proper in the sense that for every state $\imath \notin Z$, under $\mu^{*}$, there is a positive probability path that starts at i and ends at a state of Z.

## APPENDIX

## Mathematical Review

The purpose of this appendix is to provide a list of mathematical definitions, notations, and results that are used firequently in the text. For detailed expositions, the reader may consult textbooks such as Hoffman and Kunze [HoK71], Royden [Roy88], Rudin [Rud75], and Strang [Str76].

## A.I SETS

If $x$ is a member of the set $S$, we write $\boldsymbol{x} \boldsymbol{E S}$. We write $x \notin S$ if $x$ is not a member of $S$. A set $S$ may be specified by listing its elements within braces. For example, by writing $S=\left\{x t, x 2, \ldots, x_{n}\right)$ we mean that the set 8 consists of the elements $x l, x 2, \ldots, x_{n}$. A set $S$ may also be specified in the generic form

$$
8=\{x \mid x \text { satisfies } P\}
$$

as the set of elements satisfying property $P$. For example,

$$
S=\{x \mid x: \text { real }, 0 \leq x \leq \mathrm{I}\}
$$

denotes the set of all real numbers $x$ satisfying $0 \leq x \leq 1$.
The union of two sets 8 and $T$ is denoted by BUT and the intersection of 8 and $T$ is denoted by 8 nT . The union and the intersection of a sequence of sets $81,8_{2}, \ldots$
$\ldots$ are denoted by $\cup_{k=1}^{\infty} S_{k}$ and $n k=18 k$, respectively. If 8 is a subset of $T$ (i.e., if every element of $S$ is also an element of $T$ ), we write 8 c $T$ or $T \supset 8$.

Finite and Countable Sets

A set $S$ is said to be finite if it consists of a finite number of elements. It is said to be countable if there exists a one-to-one function from $S$ into the set of nonnegative integers. Thus, according to our definition, a finite set is also countable but not conversely. A countable set $S$ that is not finite may be represented by listing its elements $x o, X l, X 2, \ldots$ (i.e., $S=\{x o, X l, X 2, " '\})^{\prime}$ A countable union of countable sets is countable, that is, if $A=\{a o, a l, \ldots\}$ is a countable set and $S a o^{\prime} S a l^{\prime} \ldots$ are each countable sets, then $U k=o S a k$ is also a countable set.

## Sets of Real Numbers

If $a$ and $b$ are real numbers or $+00,-00$, we denote by $[a, b]$ the set of numbers $X$ satisfying $a \leq X \leq b$ (including the possibility $X=+00$ or $x=-00$ ). A rounded, instead of square, bracket denotes strict inequality in the definition. Thus ( $a, b],[a, b)$, and $(a, b)$ denote the set of all $x$ satisfying $a<x \leq b, a \leq x<b$, and $a<x<b$, respectively.

If $S$ is a set of real numbers that is bounded above, then there is a smallest real number $y$ such that $x \leq y$ for all $x E S$. This number is called the least upper buand or supremum of $S$ and is denoted by $\sup \{x \mathrm{I} X E S\}$ or
$\lfloor X E S\}$. (This is somewhat inconsistent with normal mathematical usage, where the use of max in place of sup indicates that the supremum is attained by some element of S.) Similarly, the greatest real number $z$ such that $z \leq x$ for all $\boldsymbol{x} \boldsymbol{E} \boldsymbol{S}$ is called the greatest lower bound or infimum of $S$ and is denoted by $\inf \{x \mid X E S\}$ or $\min \{x \mid X E S\}$. If $S$ is unbounded above, we write $\sup \{x \mid X E S\}=+00$, and if it is unbounded below, we write $\inf \{\mathrm{x} \mid X E S\}=-00$. If $S$ is the empty set, then by convention we write $\inf \{\mathrm{x} \mid X E S\}=+00$ and $\sup \{\mathrm{x} I X E S\}=-00$.

The set of all n-tuples $x=\left(X l, \ldots, x_{n}\right)$ of real numbers constitutes the $n$-dimensional Euclidean space, denoted by $\Re^{n}$. The elements of $\Re^{n}$ are referred to as n-dimensional vectors or simply vectors when confusion cannot arise. The one-dimensional Euclidean space $\Re^{1}$ consists of all the real numbers and is denoted by $\Re$. Vectors in $\Re^{n}$ can be added by adding their corresponding components. They can be multiplied by a scalar by multiplication of each component by a scalar. The inner product of two vectors $x=\left(X l, \ldots, x_{n}\right)$ and $y=(Y l, " \prime Y n)$ is denoted by $x^{\prime} y$ and is equal to

XiYi. The norm of a vector $x=\left(X l, \ldots, x_{n}\right) \mathrm{E} \Re^{n}$ is denoted by Ixll and is equal to $\left(x^{\prime} x\right)^{1 / 2}=$ $\left.x_{i}^{2}\right)^{1 / 2}$.

A set of vectors $a l, a 2, \ldots, a k$ is said to be linearly dependent if there exist scalars $\lambda_{1}, \lambda_{2}, \ldots, \lambda_{k}$, not all zero, such that

$$
\lambda_{1} a_{1}+\cdots+\lambda_{k} a_{k}=0 .
$$

If no such set of scalars exists, the vectors are said to be linearly independent.

## A. 3 MATRICES

An $m \times n$ matrix is a rectangular array of numbers, referred to as elements or components, which are arranged in m rows and $n$ columns. If $\mathrm{m}=n$ the matrix is said to be square. The element in the ith row and jth column of a matrix $A$ is denoted by a subscript $i j$, such as $a i j$, in which case we write $A=[a i j]$. The $n \times n$ identity matrix, denoted by $I$, is the matrix with elements $a i j=0$ for $\mathrm{i} \neq j$ and aii $=1$, for $\mathrm{i}=1, \ldots, n$. The sum of two $\mathrm{m} \times n$ matrices $A$ and $B$ is written as $A+B$ and is the matrix whose elements are the sum of the corresponding elements in $A$ and $B$. The
of a matrix $A$ and a scalar $\lambda$, written as $\lambda A$ or $A \lambda$, is obtained by multiplying each element of $A$ by $\lambda$. The product $A B$ of an $\mathrm{m} \times n$ matrix $A$ and an $n \times p$ matrix $B$ is the $\mathrm{m} \times p$ matrix C with elements $C_{i j}=\sum_{k=1}^{n}$ aikbkj. If $b$ is an n-dimensional column vector and $A$ is an $\mathrm{m} \times n$ matrix, then $A b$ is an m-dimensional column vector.

The transpose of an $\mathrm{m} \times n$ matrix $A$ is the $n \times \mathrm{m}$ matrix $A^{\prime}$ with elements $a_{i j}^{\prime}=a j i$. The elements of a given row (or column) of $A$ constitute a vector called a row vector (or column vector, respectively) of $A$. A square matrix $A$ is symmetric if $A^{\prime}=A$. An $n \times n$ matrix $A$ is called nonsingular or invertible if there is an $n \times n$ matrix called the inverse of $A$ and denoted by $A^{-1}$, such that $A^{-1} A=I=A A^{-1}$, where $I$ is the $n \times n$ identity matrix. An $n \times n$ matrix is nonsingular if and only if its $n$ row vectors are linearly independent or, equivalently, if its $n$ column vectors are linearly independent. Thus, an $n \times n$ matrix $A$ is nonsingular if and only if the relation $A v=0$, where $v \mathrm{E} \Re^{n}$, implies that $v=0$.

## Rank of a Matrix

The rank of a matrix $A$ is equal to the maximum number of linearly independent row vectors of $A$. It is also equal to the maximum number of linearly independent column vectors. Thus, the rank of an $\mathrm{m} \times n$ matrix is at most equal to the minimum of the dimensions m and $n$. An $\mathrm{m} \times n$ matrix is said to be of full rank if its rank is maximal, that is, if its rank is equal to the minimum of $m$ and $n$. A square matrix is of full rank if and only if it is nonsingular.

## Eigenvalues

Given a square $n \times n$ matrix $A$, the determinant of the matrix $\gamma I-$ where I is the $n \times n$ identity matrix and $\gamma$ is a scalar, is an nth degree polynomial. The $n$ roots of this polynomial are called the eigenvalues of $A$. Thus, $\gamma$ is an eigenvalue of $A$ if and only if the matrix $\gamma I-A$ is singular, or equivalently, if and only if there exists a nonzero vector $v$ such that $A v=\gamma v$. Such a vector $v$ is called an eigenvector corresponding to $\gamma$. The eigenvalues and eigenvectors of $A$ can be complex even if $A$ is real. ${ }_{A}$ matrix $A$ is singular if and only if it has an eigenvalue that is equal to zero. If $A$ is nonsingular, then the eigenvalues of A-I are the reciprocals of the eigenvalues of $A$. The eigenvalues of $A$ and $A I$ coincide.

If $\gamma_{1}, \ldots, \gamma_{n}$ are the eigenvalues of $A$, then the eigenvalues of $\mathrm{c} 1+$ where c is a scalar and I is the identity matrix, are $\mathrm{c}+\gamma_{1}, \ldots, \mathrm{c}+\gamma_{n}$. The eigenvalues of where $k$ is any positive integer, are equal to $\gamma_{1}^{k}, \cdots, \gamma_{n}^{k}$. From this it follows that $\lim _{k \rightarrow 0} A k=$ aif and only if all the eigenvalues of A lie strictly within the unit circle of the complex plane. Furthermore, if the latter condition holds, the iteration

$$
x_{k+1}=A x_{k}+b
$$

where $b$ is a given vector, converges to

$$
\bar{x}=(\mathrm{I}-A)-\mathrm{lb},
$$

which is the unique solution of the equation $x=A x+b$.
If all the eigenvalues of $A$ are distinct, then their number is exactly $n$, and there exists a set of corresponding linearly independent eigenvectors. In this case, if $\gamma_{1}, \ldots, \gamma_{n}$ are the eigenvalues and $\mathrm{VI}_{1}, \ldots, v_{n}$ are such eigenvectors, every vector $x \mathrm{E} \Re^{n}$ can be decomposed as

$$
X=\sum_{i=1}^{n} \xi_{i} v_{i}
$$

where $\xi_{i}$ are some unique (possibly complex) numbers. Furthermore, we have for all positive integers $k$,

$$
A k x=\gamma_{\mathbf{i}=1}^{n} \gamma_{i}^{k} \xi_{i} v_{i}
$$

If $A$ is a transition probability matrix, that is, all the elements of $A$ are nonnegative and the sum of the elements of each of its rows is equal to 1 , then all the eigenvalues of $A$ lie within the unit circle of the complex plane. Furthermore, 1 is an eigenvalue of $A$ and the unit vector $(1,1, \ldots, 1)$ is a corresponding eigenvector.

## Positive Definite and Semidefinite Svmmetric Matrices

A square symmetric $n \times n$ matrix $A$ is said to be positive se.midefinite if $x \mid A x \geq$ a for all $x \mathrm{E} \Re^{n}$. It is said to be positive definite If $x \mid A x>$ a for all nonzero $x \mathrm{E} \Re^{n}$. The matrix $A$ is said to be negative semidefinite (definite) if $-A$ is positive semidefinite (definite). In this book, the notions of positive definiteness and semidefimteness WIll be used only in connectIOn with symmetric matrices.
-
A positive definite symmetric matrix is invertible and Its inverse IS also positive definite symmetric. Also, an invertible positive semïdefin.lte symmetric matrix is positive definite. Analogous results hold for negative definite and semidefinite symmetric matrices. If $A$ and Bare $n \times n$ pOSItIve semidefinite (definite) symmetric matrices, then the matrix $\lambda A+\mu B$ is also positive semidefinite (definite) symmetric for .all $\lambda \geq \mathrm{a}$. and $\mu \geq \mathrm{a}$. If $A$. IS an $n \mathrm{x} n$ positive semidefinite symmetric matnx and $C$ IS an m x. $n$ matrix, then the matrix $C A C /$ is positive semidefinite symmetric. If $A$ IS positive definite symmetric, and $C$ has rank m (equivalently, $\mathrm{m} \leq$ nand $C$ has full rank), then $C A C I$ is positive definite symmetric.

An $n \times n$ positive definite symmetric matnx $A$ can be wntten as $C C I$ where $C$ is a square invertible matrix. If $A$ is positive semidefinite symmetric and its rank is m , then it can be written as $C C^{\prime}$, where $C$ is an $n \times \mathrm{m}$ matrix of full rank.

A symmetric $n \times n$ matrix $A$ has real eigenvalues and a set of $n$ real linearly indepeadent eigenvectors, which are orthogonal (the imner product of any pair is a). If $A$ is positive semidefinite (definite) symmetric, ItS cigenvalues are nonnegative (respectively, positive).

## Partitioned Matrices

It is often convenient to partition a matrix into submatrices. For example, the matrix

$$
A=\left(\begin{array}{ll}
a_{11} & a_{14} \\
a_{21} & a_{24} \\
a_{31} & a_{34}
\end{array}\right)
$$

may be partitioned into

$$
A=\left(\begin{array}{ll}
\mathrm{A}^{\mathrm{A}} 11 & A 12 \\
\mathrm{~A}_{2 \mathrm{I}} & A 22
\end{array}\right)
$$

where

$$
\begin{array}{rlrl}
\mathrm{All}=\left(\begin{array}{ll}
\text { all } & a 12
\end{array}\right), & \mathrm{A}_{12} & =\left(\begin{array}{ll}
a 13 & \mathrm{a} 14
\end{array}\right), \\
& =\left(\begin{array}{ll}
a_{31}^{21} & a 22 \\
a 31 & a 32
\end{array}\right), & & =\left(\begin{array}{ll}
a^{23} & a_{33}^{24}
\end{array}\right) .
\end{array}
$$

We separate the components of a partitioned matrix by a space, as in (BC), or by a comma, as in $(B, C)$. The transpose of the partitioned matrix $A$ is

$$
A^{\prime}=\left(\begin{array}{ll}
A_{11}^{\prime} & A_{21}^{\prime} \\
A_{12}^{\prime} & A_{22}^{\prime}
\end{array}\right.
$$

Partitioned matrices may be multiplied just as nonpartitioned matrices, provided the dimensions involved in the partitions are compatible. Thus if

$$
A \quad B=\left(B_{B} l l\right.
$$

then

$$
\begin{gathered}
\left.A B=\begin{array}{ll}
(A l l B l l+A 12 B 21 & A l l B 12+A l 2 B 22
\end{array}\right), \\
A 21 B l l+A 22 B 21
\end{gathered} \text { A2lB12+A22B22}, ~, ~
$$

provided the dimensions of the submatrices are such that the preceding products $A i j B j k, i, j, k=1,2$ can be formed.

## Matrix Inversion Formulas

Let $A$ and $B$ be square invertible matrices, and let $C$ be a matrix of appropriate dimension. Then, if all the following inverses exist, we have

$$
(A+C B C /)-l=A-I-A-1 C(B-1+0 \cdot A-1 C)-l C^{\prime} A-I
$$

The equation can be verified by multiplying the right-hand side by

$$
A+O B O^{\prime}
$$

and showing that the product is the identity matrix.
Consider a partitioned matrix $1 \backslash 1$ of the form

$$
M=\left(\begin{array}{ll}
A & B \\
C & D
\end{array}\right)
$$

Then we have

$$
M^{-1}=\left(\begin{array}{cc}
-D-Q_{C Q} & -Q B D-1
\end{array}\right)
$$

where

$$
Q=(A-B D-1 C)-1,
$$

provided all the inverses exist. The proof is obtained by multiplying $M$ with the expression given for $M-1$ and verifying that the product yields the identity matrix.

## A. 4 ANALYSIS

## Convergence of Sequences

A sequence of vectors $x o, X l, \ldots, X k, \boldsymbol{\mathcal { O }}$ in $\Re^{n}$, denoted by $\left\{x_{k}\right\}$, is said to converge to a limit $X$ if $j j x k-X i i l \rightarrow$ as $k \rightarrow 00$ (i.e., if, given any $\epsilon>0$, there is an integer $N$ such that for all $k \geq N$ we have $I j x k-\mathrm{XII}<\mathrm{E}$ ). If $\{X k\}$ converges to $x$, we write $X k \rightarrow x$ or $\lim _{k \rightarrow \infty} X k=x$. We have $A X k+B Y k \rightarrow A x+B y$ if $X k \rightarrow x, Y k \rightarrow Y$, and $A, B$ are matrices of appropriate dimension.

A vector $x$ is said to be a limit point of a sequence $\{x k\}$ if there is a subsequence of $\{x k\}$ that converges to $x$, that is, if there is an infinite subset $\mathcal{K}$ of the nonnegative integers such that for any $\epsilon>0$, there is an integer $N$ such that for all $k \mathrm{E} \mathcal{K}$ with $k \geq N$ we have $I l x k-x i i<\mathrm{E}$.

A sequence of real numbers $\{r k\}$, which is monotonically nondecreasing (nonincreasing), that is, satisfies $r k \leq r k+l$ for all $k$, must either converge to a real number or be unbounded above (below). In the latter case we write $\lim _{k \rightarrow \infty} r k=00(-00)$. Given any bounded sequence of real numbers $\{r k\}$, we may consider the sequence $\{S k\}$, where $S k=\sup \{\mathrm{riji} \geq k\}$. Since this sequence is monotonically nonincreasing and bounded, it must have a limit. This limit is called the limit superior of $\{r k\}$ and is denoted by $\lim \sup _{k \rightarrow \infty} r k$. The limit inferior of $\{r k\}$ is similarly defined and is denoted by $\liminf _{k \rightarrow \infty} r k$. If $\{r k\}$ is unbounded above, we write $\lim \sup _{k \rightarrow \infty} r k \quad 00$, and if it is unbounded below, we write $\liminf _{k \rightarrow \infty} r k=-00$. We also use this notation if $r k \mathrm{E}[-00,00]$ for all $k$.

## Open, Closed, and Compact Sets

A subset $S$ of $\Re^{n}$ is said to be open if for every vector $\boldsymbol{x} \boldsymbol{E} \boldsymbol{S}$ one can find an $\epsilon>$ such that $\{z\|\|-x i i<\mathrm{E}\}$ C $S$. A set $S$ is closed if and only if every convergent sequence $\{X k\}$ with elements in $S$ converges to a point that also belongs to $S$. A set $S$ is said to be compact if and only if it is both closed and bounded (i.e., it is closed and for some $M>0$ we have $j \mid x \| \leq M$ for all $X E S)$. A set $S$ is compact if and only if every sequence $\{X k\}$ with elements in $S$ has at least one limit point that belongs to $S$. Another important fact is that if $S o, S l, \ldots, S k, \ldots$ is a sequence of nonempty compact sets in $\Re^{n}$ such that $S k \supset S k+1$ for all $k$, then the intersection $\cap_{k=0}^{\infty} S_{k}$ is a nonempty and compact set.

## Continuous Functions

A function $f$ mapping a set $S l$ into a set $S 2$ is denoted by $f: S l \rightarrow S 2$. A function $f: \Re^{n} \rightarrow \Re^{m}$ is said to be continuous if for all $x, f(x k) \rightarrow f(x)$ whenever $X k \rightarrow x$. Equivalently, $f$ is continuous if, given $x$ E $\Re^{n}$ and $\epsilon>0$,
there is a $\delta>0$ such that whenever $\|y-x\|<\delta$, we have $\operatorname{Ilf}(y)-f(x) \| l<\epsilon$. The function

$$
\left(a_{1} f_{1}+a_{2} f_{2}\right)(\cdot)=a_{1} f_{1}(\cdot)+a_{2} f_{2}(\cdot)
$$

is continuous for any two scalars $a I, a 2$ and any two continuous functions $f_{1}, f_{2}: \Re^{n} \rightarrow \Re^{m}$. If 51, 53 are any sets and $f_{1}: 51 \rightarrow S_{2}, f_{2}: 52 \rightarrow 53$ are functions, the function $f_{2} f_{1}: 51 \rightarrow 53$ defined by $\left(f_{2} \cdot f_{1}\right)(x)=$ $f_{2}\left(f_{1}(x)\right)$ is called of $f_{1}$ and $f_{2}$. If $f_{1}: \Re^{n} \rightarrow \Re^{m}$ and $f_{2} \cdot \Re^{m} \rightarrow \Re^{p}$ are continuous, then $f_{2} \cdot f_{1}$ is also continuous.

Derivatives
Let $f: \Re^{n} \mapsto \Re$ be some function. For a fixed $x$ E $\Re^{n}$, the first partial derivative of $f$ at the point $x$ with respect to the ith coordinate is defined by

$$
\frac{8 f(x)}{\partial x_{i}}=\lim _{\alpha \rightarrow 0} \frac{f(x}{ \pm} \frac{\left.\alpha e_{i}\right)}{\alpha}=f(x)
$$

where ei is the 'ith unit vector, and we assume that the above limit exists. If the partial derivatives with respect to all coordinates exist, $f$ is called differentiable at $x$ and its gradient at $x$ is defined to be the column vector

$$
\left.\nabla f(x)=\frac{\underline{a f(X)})}{\underline{\vdots}} \begin{array}{c}
a f(X) \\
\underline{a X n}
\end{array}\right)
$$

The function $f$ is called differentiable if it is differentiable at every $x \mathrm{E} \Re^{n}$. If $\nabla f(x)$ exists for every $x$ and is a continuous function of $x, f$ is said to be continuously differentiable. Such a function admits, for every fixed $x$, the first order expansion

$$
f(x+y)=f(x)+y^{\prime} \nabla f(x)+0(\|Y Y\|)
$$

where $0(\| y \mid D$ is a function of $y$ with the property limllyll $\mathrm{O} O(\|y\|) /\|y\|=0$.
A vector-valued function $f: \Re^{n} \mapsto \Re^{m}$ is called differentiable (respectively, continuously differentiable) if each component $f i$ of $f$ is differentiable (respectively, continuously differentiable). The gradient matrix of $f$, denoted by $\nabla f(x)$, is the $n \times m$ matrix whose ith column is the gradient $\nabla f i(X)$ of $f i$. Thus,

$$
\nabla f(x)=\left[\nabla f_{1}(x) \cdots \nabla f_{m}(x)\right]
$$

The transpose of $\nabla f$ is the Jacobian of $f$; it is the matrix whose ijth entry is equal to the partial derivative $\partial f i l a x j$.

If the gradient $\nabla f(x)$ is itself a differentiable function, then $f$ is said to be twice differentiable. We denote by $\nabla^{2} f(x)$ the Hessian matrix of $f$ at $x$, that is, the matrix

$$
\nabla^{2} f(x)=\left[\frac{\partial^{2}}{\partial x^{i} \partial x^{j}}\left(x^{\prime}\right)\right]
$$

the elements of which are the second partial derivatives of $f$ at $x$.
Let $f: \Re^{k} \mapsto \Re^{m}$ and $9: \Re^{m} \mapsto \Re^{n}$ be continuously differentiable functions, and let $h(x)=g(J(x))$. The chain rule for differentiation states that

$$
\nabla h(x)=Y f(x) y g(f(x)), \quad \text { for all } x \mathrm{E} \Re^{k} .
$$

For example, if $A$ and $B$ are given matrices, then if $h(x)=$ we have $\nabla h(x)=A^{\prime}$ and if $h(x)=A B x$, we have $\nabla h(x)=B^{\prime} A^{\prime}$.

## A. 5 CONVEX SETS AND FUNCTIONS

A subset 0 of $\Re^{n}$ is said to be convex if for every $x, y E O$ and every scalar $\alpha$ with $0 \leq \alpha \leq 1$, we have $\alpha x+(1-\alpha) y \mathrm{E} \mathrm{O}$. In words, 0 is convex if the line segment connecting any two points in 0 belongs to O . A function $f: C \rightarrow \quad$ defined over a convex subset 0 of $\Re^{n}$, is said to be convex if for every $x, y \leq 0$ aud every scalar $\alpha$ with $0 \leq \alpha \leq 1$ we have

$$
f(\alpha x+(1-\alpha) y) \leq \alpha f(x)+(1-\alpha) f(y)
$$

The function $f$ is said to be concave if $(-f)$ is convex, or equivalently if for every $x, y E O$ and every scalar $\alpha$ with $0 \leq \alpha \leq 1$ we have

$$
f(\alpha x+(1-\alpha) y) \geq \alpha f(x)+(1-\alpha) f(y)
$$

If $f: 0 \rightarrow \Re$ is convex, then the sets $\Gamma_{\lambda}=\{x \mid x \in 0, f(x) \leq \lambda\}$ are convex for every scalar $\lambda$. An important property is that a real-valued convex function defined over $\Re^{n}$ is continuous.

If $f_{1}, f_{2}, \ldots, f_{m}$ are convex functions defined over a convex subset 0 of $\Re^{n}$ and $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{m}$ are nonnegative scalars, then the function $\alpha_{1} f_{1}+\ldots+\alpha_{m} f_{m}$ is also convex over O . If $f: \Re^{m} \rightarrow \Re$ is convex, $A$ is an $\mathrm{m} \times n$ matrix, and $b$ is a vector in $\Re^{m}$, the function $9: \Re^{n} \rightarrow \Re$ defined by $9(x)=f(A x+b)$ is also convex. If $f: \Re^{n} \rightarrow \Re$ is convex, then the function $g(x)=E w\{f(x+w)\}$, where $w$ is a random vector in convex function provided the expected value is finite for every $x \mathrm{E} \Re^{n}$.

For functions $f: \Re^{n} \rightarrow \Re$ that are differentiable, there are alternative characterizations of convexity. Thus, $f$ is convex if and only if

$$
f(y) \geq f(x)+\nabla f(x)^{\prime}(Y-x), \quad \text { for all } x, y \mathrm{E} \Re^{n} .
$$

If $f$ is twice continuously differentiable, then $f$ is convex if and only if $\nabla^{2} f(x)$ is a positive semidefinite symmetric matrix for every $x$ E $\Re^{n}$.

For accounts of convexity and its applications in optimization, see Bertsekas [BN003) and Rockafellar [Roc70).

## APPENDIX

## Optimization Theory

The purpose of this appendix is to provide a few definitions and results of deterministic optimization. For detailed expositions, which include both convex and nonconvex problems, see textbooks such as Bertsekas [Ber9g], [BN003], Luenberger [Lue84], and Rockafellar [Roc70].

## . 1 OPTIMAL SOLUTIONS

Given a set $S$, a real-valued function $j: S \rightarrow \Re$, and a subset $X e S$, the optimization problem

$$
\begin{align*}
& \operatorname{minimize} j(x)  \tag{B.1}\\
& \text { subject to } x \mathrm{E} X,
\end{align*}
$$

is to find an element $x^{*} \mathrm{E} X$ (called a minimizing element or an optimal solution) such that

$$
j\left(x^{*}\right) \leq j(x), \quad \text { for all } x \mathrm{EX}
$$

For any minimizing element $x^{*}$, we write

$$
x^{*}=\arg \min _{x E X} j(x)
$$

Note that a minimizing element need not exist. For example, the scalar functions $j(x)=x$ and $j(x)=e X$ have no minimizing elements over the set of real numbers. The first function decreases without bound to -oo as $x$ tends toward -oo, while the second decreases toward 0 as $x$ tends toward -oo but always takes positive values. Given the range of values that $j(x)$ takes as $x$ ranges over $X$, that is, the set of real numbers

$$
\{j(x) \mid x \mathrm{EX}\},
$$

there are two possibilities:

1. The set $\{j(x)$ I $x \in X\}$ is unbounded below (Le., contains arbitrarily small real numbers) in which case we write

$$
\min \left\{\left.\mathrm{j}(\mathrm{x})\right|_{x \in X} X\right\}=-\mathrm{oo} \quad \text { or } \quad \min _{x E X} j(x)=-\mathrm{oo}
$$

2. The set $\{j(x) \mid x \in X\}$ is bounded below; that is, there exists a scalar M such that $N l \leq j(x)$ for all $x$ EX. The greatest lower bound of $\{j(x) \mid x \in X\}$ is also denoted by

$$
\min \{\mathrm{j}(\mathrm{x}) \mid x \mathrm{E} X\} \quad \text { or } \quad \min _{x E X} j(x)
$$

In either case we call minxEx $j(x)$ the optimal value of problem (B.1). A maximization problem of the form

```
maximize j (x)
subject to }x\in\mathbb{E}
```

may be converted to the minimization problem

$$
\begin{aligned}
& \operatorname{minimize}-j(x) \\
& \text { subject to } x \mathrm{E} X
\end{aligned}
$$

in the sense that both problems have the same optimal solutions, and the optimal value of one is equal to minus the optimal value of the other. The optimal value for the maximization problem is denoted by maxxEX $j(x)$.

Existence of at least one optimal solution in problem (B.l) is guaranteed if $j: \Re^{n} \rightarrow \Re$ is a continuous function and $X$ is a compact subset of $\Re^{n}$. This is the Weierstrass theorem. By a related result, existence of an optimal solution is guaranteed if $j: \Re^{n} \rightarrow \Re$ is a continuous function, $X$ is closed, and $j(x) \rightarrow 0$ if Ixll $\rightarrow 00$.

## B. 2 OPTIMALITY CONDITIONS

Optimality conditions are available when $J$ is a differentiable function on $\Re^{n}$ and $X$ is a convex subset of $\Re^{n}$ (possibly $X \quad \Re^{n}$ ). In particular, if $x^{*}$ is an optimal solution of problem (B.1), $J: \Re^{n} \rightarrow \Re$ is a continuously differentiable function on $\Re^{n}$, and $X$ is convex, we have

$$
\begin{equation*}
v J\left(x^{*}\right)^{\prime}\left(x-x^{*}\right) \geq 0, \quad \text { for all } x \mathrm{E} X \tag{B.2}
\end{equation*}
$$

where $\nabla J\left(x^{*}\right)$ denotes the gradient of $J$ at $X^{*}$. When $X=\Re^{n}$ (i.e., the minimization is unconstrained), the necessary condition (B.2) is equivalent to

$$
\begin{equation*}
\nabla J\left(x^{*}\right)=0 . \tag{B.3}
\end{equation*}
$$

When $J$ is twice continuously differentiable and $X=\Re^{n}$, an additional necessary condition is that the Hessian matrix $\nabla^{2} J\left(x^{*}\right)$ be positive semidefin'ite at $x^{*}$. An important fact is that iJ J: $\Re^{n} \rightarrow \Re$ is a convex Junction and $X$ is convex) then Eq. (B.2) is both a necessary and a sufficient condition Jor optimality of $x^{*}$.

Other types of , AJllmtions deal with the case winere the constraint set ${ }^{V}$ iunsists of equality and inequality constraints, i.e., problems ur the form

```
minimize \(J(x)\)
subject to \(h l(x)=0, \ldots, h m(x)=0, g l(X) \leq 0, \ldots, g r(X) \leq 0\),
```

where $J$, hi, gj are continuously differentiable functions from $\Re^{n}$ to $\Re$.
We say that the vectors $\lambda^{*}=\left(\lambda_{1}^{*}, \ldots, \lambda_{m}^{*}\right)$ and $\mu^{*}=\left(\mu_{1}^{*}, \ldots, \mu_{r}^{*}\right)$ are Lagrange multiplier vectors corresponding to a local minimum $x^{*}$ if they satisfy the following conditions:

$$
\begin{gathered}
v \boldsymbol{J}\left(x^{*}\right)+\sum_{\mathrm{i}=1}^{\mathrm{m}} \lambda_{i}^{*} \nabla h_{i}\left(x^{*}\right)+\sum_{j=l}^{r} \mu_{j}^{*} \nabla g_{j}\left(x^{*}\right)=0, \\
\mu_{j}^{*} \geq 0, \quad \text { for all } j=1, \ldots, r, \\
\mu_{j}^{*}=0, \quad \text { for all } j \notin A\left(x^{*}\right),
\end{gathered}
$$

where is the index set of inequality constraints that are active at $x^{*}$ :

$$
A\left(x^{*}\right)=\left\{j \mid g j\left(x^{*}\right)=0\right\} .
$$

Lagrange multiplier theory revolves around conditions under which Lagrange multiplier vectors are guaranteed to exist for a given local minimum $X^{*}$. Such conditions are known as constraint qualifications. Some of the most useful ones are the following:

CQ1: The equality constraint gradients $v h i(x *), i=1, \ldots, m$, and the active inequality constraint gradients $\nabla g j\left(x^{*}\right), j \mathrm{E} \quad$ are linearly independent.
CQ2: The equality constraint gradients $v h i\left(x^{*}\right), \mathrm{i}=1, \ldots, m$, are linearly independent, and there exists ay $E \Re^{n}$ such that

0
0
$v h i\left(x^{*}\right)^{\prime} y=$ for all $\mathrm{i}=1, \ldots, \mathrm{~m}, \quad \operatorname{vgj}\left(x^{*}\right)^{\prime} y<\quad$ for all $j \mathrm{E} A\left(x^{*}\right)$.
CQ3: The functions $h i$ are linear and the functions $g j$ are concave.
CQ4: The functions $h i$ are linear, the functions $g j$ are convex, and there exists ayE $\Re^{n}$ such that

$$
g j(Y)<0, \quad \text { for all } j=1, \ldots, r
$$

Each of the above constraint qualifications implies the existence of at least one Lagrange multiplier vector associated with $x^{*}$ (unique in the case of CQ1); see e.g., [Ber99] for a detailed account.

## B. 3 MINIMIZATION OF QUADRATIC FORMS

Let $J: \Re^{n} \rightarrow \Re$ be a quadratic form

$$
J(x)=\frac{1}{2} x^{\prime} Q x+b^{\prime} x
$$

where $Q$ is a symmetric $n \times n$ matrix and $b \mathrm{E} \Re^{n}$. Its gradient is given by

$$
\nabla f(x)=Q x+b
$$

The function $J$ is convex if and only if $Q$ is positive semidefinite. If $Q$ is positive definite, then $J$ is convex and $Q$ is invertible, so by (B.3), a vector $x^{*}$ minimizes $J$ if and only if

$$
v J\left(x^{*}\right) \quad Q x^{*}+b=0
$$

or equivalently

$$
x^{*}=-Q-1 b .
$$

## APPENDIX

## Probability Theory

This appendix lists selectively some of the basic probabilistic notions that we will be using. Its main purpose is to familiarize the reader with some of our terminology. It is not meant to be exhaustive, and the reader should consult textbooks such as Ash [Ash70], Feller [Fe168], Papoulis [Pap65], Ross [Ros85], Stirzaker [Sti94], and Bertsekas and Tsitsiklis [BeT02] for detailed accounts. For fairly accessible treatments of measure theoretic probability, see Adams and Guillemin [AdG86], and Ash [Ash72]

## C. 1

## PROBABILITY SPACES

A probability space consists of
(a) A set $\Omega$.
(b) A collection $\mathcal{F}$ of subsets of $\Omega$, called events, which includes $\Omega$ and has the following properties:
(1) If $A$ is an event, then the complement $\bar{A}=\{w \in \Omega \mid w \notin A\}$ is also an event. (The complement of $\Omega$ is the empty set and is considered to be an event.)
(2) If $\quad \ldots, A_{k}, \ldots$ are events, then $\bigcup_{k=1}^{\infty} A_{k}$ is also an event.
(3) If $\ldots$ are events, then $\cap_{k=1}^{\infty} A_{k}$ is also an event.
(c) A function $P(\cdot)$ assigning to each event $A$ a real number $P(A)$, called the probability oj the event and satisfying:
(1) $P(A) \geq 0$ for every event $A$.
(2) $P(\Omega)=1$.
(3) $P(A I \cup A 2)=P(A I)+P(A 2)$ for every pair of disjoint events AI, A 2.
(4) $P\left(\cup_{k=1}^{\infty} A k\right)=\sum_{k=1}^{\infty} P(A k)$ for every sequence of mutually disjoint events $A I, A 2, \ldots, A k^{\prime} \ldots$

The function $P$ is referred to as a probability measure
Convention for Finite and Countable Probability Spaces
The case of a probability space where the set $\Omega$ is a countable (possibly finite) set is encountered frequently in this book. When we specify that $\Omega$ is finite or countable, we implicitly assume that the associated collection of events is the collection of all subsets of $\Omega$ (including $\Omega$ and the empty set). Then, if $\Omega$ is a finite set, $\Omega=\left\{W I, W 2, \quad w_{n}\right\}$, the probability space is specified by the probabilities $p_{1}, p_{2}, \quad, P n$, where $P i$ denotes the probability of the event consisting of just $\omega_{i}$. Similarly, if $\Omega=$ $\{W I, W 2, \ldots, W k, \ldots\}$, the probability space is specified by the corresponding probabilities $P I, P 2, \ldots, P k, \cdots$ In either case we refer to $P 2, \ldots, P n)$ or $(P I, P 2, \ldots, P k, \ldots)$ as a probability distribution over $\Omega$.

## C. 2 RANDOM VARIABLES

A random var'iable on a probability space $(\Omega, \mathcal{F}, P)$ is a function $x: \Omega \rightarrow \Re$ such that for every scalar $\lambda$ the set

$$
\{\omega \in \Omega \mid x(\omega) \leq \lambda\}
$$

is an event (i.e., belongs to the collection $F$ ). An n-dimensional random vector $x=\left(\mathrm{Xl}, X 2, \ldots, x_{n}\right)$ is an n-tuple of random variables $\mathrm{Xl}, X 2, \ldots, x_{n}$, each defined on the same probability space.

We define the distribution junction $F: \Re \rightarrow \Re$ [or cumulative distribution junction (CDF for short)] of a random variable $x$ by

$$
F(z)=p(\{w \in \Omega \mid x(w) \leq z\}) ;
$$

that is, $F(z)$ is the probability that the random variable takes a value less than or equal to $z$. We define the distribution function $F: \Re^{n} \rightarrow \Re$ of a random vector $X=\left(\mathrm{Xl}, X 2, \ldots, x_{n}\right)$ by
$F(Z I, Z 2, \cdots, Z n)=p(\{w \in \Omega \mid X I(W) \leq Z 1, X 2(W) \leq Z 2, \cdots, x n(w) \leq Z n\})$.

Given the distribution function of a random vector $x=\left(\mathrm{Xl}, \ldots, x_{n}\right)$, the (marginal) distribution function of each random variable $x_{i}$ is obtained from

$$
F_{i}\left(z_{i}\right)=\lim _{z_{j} \rightarrow \infty, j \neq i} F\left(z_{1}, z_{2}, \ldots, z_{n}\right)
$$

The random variables $\mathrm{X}, \ldots, X_{n}$ are said to be independent if

$$
F\left(z_{1}, z_{2}, \ldots, z_{n}\right)=F_{1}\left(z_{1}\right) F_{2}\left(z_{2}\right) \cdots F_{n}\left(z_{n}\right)
$$

for all scalars $\mathrm{Z}, \ldots, \mathrm{Zn}$.
The expected value oj a random variable $x$ with distribution function $F$ is defined by

$$
E\{x\}=\int_{-\infty}^{\infty} z d F(z)
$$

provided the integral is well-defined. The expected value oj a random vector $x \quad\left(\mathrm{Xl}, \ldots, x_{n}\right)$ is the vector

$$
E\{x\}=\left(E\left\{x_{1}\right\}, E\left\{x_{2}\right\}, E\left\{x_{\infty}\right\}\right.
$$

The covariuluce matrix of a random vector $x=\left(\mathrm{Xl}, \ldots, x_{n}\right)$ with expected value $E\{x\}=\left(\bar{x}_{1}, \ldots, \bar{x}_{n}\right)$ is defined to be the $n \mathrm{x} n$ positive semidefinite symmetric matrix

$$
\left(\begin{array}{cc}
E\left\{\left(x_{1}-\bar{x}_{1}\right)^{2}\right\} & E\{(X l \\
\vdots & \left.\left.\bar{x}_{1}\right)\left(x_{n}-\bar{x}_{n}\right)\right\} \\
E\left\{\left(x_{n}-\bar{x}_{n}\right)\left(x_{1}-\bar{x}_{1}\right)\right\} & E\left\{\left(X_{n}-\bar{x}_{n}\right)^{2}\right\}
\end{array}\right)
$$

provided the expected values are well-defined.
Two random vectors $x$ and yare said to be uncorrelated if

$$
E\left\{(X-E\{x\})(y-E\{y\})^{\prime}\right\}=0
$$

where $(x-E\{x\})$ is viewed as a column vector and $(y-E\{y\})^{\prime}$ is viewed as a row vector.

The random vector $x=\left(\mathrm{Xl}, \ldots, x_{n}\right)$ is said to be characterized by a probability density junction $f: \Re^{n} \rightarrow \Re$ if

$$
\left., z_{2}, \ldots, z_{n}\right)=\int_{-\infty}^{z_{1}} \int_{-\infty}^{z_{2}} \ldots \int_{-\infty}^{z_{n}} f\left(y_{1}, \ldots, y_{n}\right) d y_{1} \ldots d y_{n}
$$

for every $Z 1, \ldots, Z n$.

### 0.3 CONDITIONAL PROBABILITY

We restrict ourselves to the case where the underlying probability space $\Omega$ is a countable (possibly finite) set and the set of events is the set of all subsets of $\Omega$.

Given two events $A$ and $B$, we define the conditional pTObability oj $B$ given $A$ by

$$
P(B \mid A)= \begin{cases}\boldsymbol{O}(A n B) & \text { if } P(A)>0 \\ \{-\underline{p e A}) & \text { if } P(A)=0\end{cases}
$$

We also use the notation $P\{B \mid A\}$ in place of $P(B \mid A)$. If $B_{1}, B 2, \ldots$ are a countable (possibly finite) collection of mutually exclusive and exhaustive events (i.e., the sets $B i$ are disjoint and their union is $\Omega$ ) and $A$ is an event, then we have

$$
P(A)=\sum P\left(A \cap B_{i}\right)
$$

From the two preceding relations, we obtain the total pTObability theorem:

$$
P(A)=\sum_{i} P(B i) P(A \mid B i)
$$

We thus obtain for every $k$,

$$
\frac{P\left(A \cap B_{k}\right)}{P(A)} \quad \frac{P(B k) P(A \mid B k)}{\sum_{i} P(B i) P(A \mid B i)}
$$

assuming that $P(A)>0$. This relation is referred to as Bayes' rule.
Consider now two random vectors $x$ and $y$ taking values in $\Re^{n}$ and $\Re^{m}$, respectively [i.e., $x(w) \mathrm{E} \Re^{n}, y(w) \mathrm{E} \Re^{m}$ for all $w \in \Omega$ ]. Given two subsets $X$ and $Y$ of $\Re^{n}$ and $\Re^{m}$, respectively, we denote

$$
P(X \mid Y)=P(\{w \mid x(w) \in X\} \mid\{w \mid y(w) \in Y\})
$$

For a fixed vector $v E \Re^{n}$, we define the conditional distribution junction of $x$ given $v$ by

$$
F(z \mid v)=P(\{w \mid x(w) \leq \mathrm{z}\} \mid\{w \mid y(w)=v\})
$$

and the conditional expectation of $x$ given $v$ by

$$
E\{x \mid v\}=\int_{\Re^{n}}^{\boldsymbol{r}} z d F(z \mid v)
$$

assuming that the integral is well-defined. Note that $E\{x \mid v\}$ is a function mapping $v$ into $\Re^{n}$.

Finally, let us provide Bayes' rule for random vectors. If $w I, W 2, \ldots$ are the elements of $\Omega$, denote

$$
z_{i}=x\left(\omega_{i}\right), \quad v_{i}=y\left(\omega_{i}\right), \quad \mathrm{i}=1,2, \ldots
$$

Also, for any vectors $z \mathrm{E} \Re^{n}, v \mathrm{E} \Re^{m}$, let us denote

$$
P(z) \quad p(\{w l x(w)=z\}), \quad P(v)=p(\{w \mid y(w)=v\}) .
$$

have $P(z)=0$ if $z \neq z i, \mathrm{i}=1,2, \ldots$, and $P(v)=0$ if $v \neq v i, \mathrm{i}=1,2, \ldots$ Denote aisi

$$
\begin{aligned}
& P(z l v)=p(\{w l x(w)=z\} \mid\{w l y(w)=v\}) \\
& P(v \mid z)=P(\{w \mid y(\omega)=. I\{w l x(w)=z\})
\end{aligned}
$$

Then, for all $k=1,2, \ldots$, Bayes' rule yields

$$
P\left(z_{k} \mid v\right)= \begin{cases}\sum_{i}^{P\left(z_{k}\right) P\left(v \mid z_{k}\right)} P & \text { if } P(v)>0 \\ 0 & \text { if } P(v)=0\end{cases}
$$

## APPENDIX

## Finite-State Markov Chains

This appendix provides some of the basic probabilistic notions related to stationary Markov chains with a finite number of states. For detailed presentations, see Ash [Ash70], Bertsekas and Tsitsiklis [BeT02], Chung [Chu60], Gallager Kemeny and Snell [KeS60], and Ross [Ros85].

## D. 1 STATIONARY MARKOV CHAINS

A square $n \times n$ matrix [ $P_{i j}$ ] is said to be a stochastic matrix if all its elements are nonnegative, that is, Pij $\geq 0, i, j=1, \ldots, n$, and the sum of the elements of each of its rows is equal to 1 , that is, $\sum_{j=1}^{n} P i j=1$ for all i = 1, ..., по

Suppose we are given a stochastic $n \mathrm{x} n$ matrix $P$ together with a finite set of states $8=\{I, \ldots, n\}$. The pair $(8, P)$ will be referred to as a stationary finite-state Markov chain. We associate with $(8, P)$ a process whereby an initial state xe E 8 is chosen in accordance with some initial probability distribution

$$
\mathrm{rQ}=\left(r_{0}^{1}, r_{0}^{2}, \ldots, r_{0}^{n}\right)
$$

Subsequently, a transition is made from state xQ to a new state xı E 8 in accordance with a probability distribution specified by $P$ as follows. The
probability that the new state will be $j$ is equal to $P i j$ whenever the initial state is i, i.e.,

$$
P(X l=j \mid \times Q=\mathrm{i})=P i j, \quad i, j=1, \ldots, n .
$$

Similarly, subsequent transitions produce states $x 2, x 3, \ldots$ in accordance with

$$
\begin{equation*}
P(X k+l=j \mid X k=i)=P i j, \quad i, j=1, \ldots, n . \tag{D.1}
\end{equation*}
$$

The probability that after the kth transition the state $X k$ will be $j$, given that the initial state xa is $i$, is denoted by

$$
\begin{equation*}
p_{i j}^{k}=P\left(X k=\left.j\right|_{\mathrm{XQ}}=\mathrm{i}\right), \quad i, j=1, \ldots, n \tag{D.2}
\end{equation*}
$$

A straightforward calculation shows that these probabilities are equal to the elements of the matrix $p k$ ( $P$ raised to the kth power), in the sense that $p_{i j}^{k}$ is the element in the ith row and $\mathbf{j}$ th column of $p k$ :

$$
\begin{equation*}
p k=\left[p_{i j}\right] . \tag{D.3}
\end{equation*}
$$

Given the initial probability distribution $P Q$ of the state $X Q$ (viewed as a row vector in $\Re^{n}$ ), the probability distribution of the state $X k$ after $k$ transitions

$$
r_{k}=\left(r_{k}^{1}, r_{k}^{2}, \ldots, r_{k}^{n}\right)
$$

(viewed again as a cuiv retion) is aime hy

$$
\begin{equation*}
r_{k}=r_{0} P^{k}, \quad k=1,2, \ldots \tag{DA}
\end{equation*}
$$

This relation follows from Eqs. (D.2) and (D.3) once we write

$$
r_{k}^{J}={ }_{\mathrm{i}=1}^{n} P\left(X k=\left.j\right|_{\mathrm{XQ}}=i\right) r_{0}^{i}=\sum_{\mathrm{i}=1}^{n} p_{i j}^{k} r_{0}^{i}
$$

## D. 2 CLASSIFICATION OF STATES

Given a stationary finite-state Markov chain ( $8, P$ ), we say that two states i and $j$ communicate if there exist two positive integers $\mathrm{k}_{1}$ and $k 2$ such that $p_{i j}^{k_{1}}>0$ and $p_{j i}^{k_{2}}>0$. In words, states i and $j$ communicate if one can be reached from the other with positive probability.

Let $\tilde{S}$ C 8 be a subset of states such that:

1. All states in $\tilde{S}$ communicate.
2. If $i$ E $\tilde{S}$ and $j \notin \tilde{S}$, then $p_{i j}^{k}=0$ for all $k$.

Then we say that $\tilde{S}$ forms a recurrent class of states.
If 8 forms by itself a recurrent class (i.e., all states communicate with each other), then we say that the Markov chain is irreducible. It is possible that there exist several recurrent classes. It can also be proved that at least one recurrent class must exist. A state that belongs to some recurrent class is called recurrent; otherwise it is called transient. We have

$$
\lim _{k \rightarrow \infty} p_{i i}^{k}=0 \quad \text { if and only if } \quad \mathrm{i} \text { is transient. }
$$

In other words, if the process starts at a transient state, the probability of returning to the same state after $k$ transitions diminishes to zero as $k$ tends to infinity.

The definitions imply that if the process starts within a recurrent class, it stays within that class. If it starts at a transient state, it eventually (with probability one) enters a recurrent class after a number of transitions, and subsequently remains there.

## D. 3 LIMITING PROBABILITIES

An important property of any stochastic matrix $P$ is that the matrix $P^{*}$ defined by
exists [in the sense that the sequences of the elements of (liN) $\sum_{k=0}^{N-1} p k$ converge to the corresponding elements of $\left.P^{\star}\right]$. A proof of this is given in Prop. A. 1 of Appendix A in Vol. II. The elements $p_{i j}^{*}$ of $P^{*}$ satisfy

$$
p_{i j}^{*} \geq 0, \quad \sum_{\mathbf{j}=1}^{n} p_{i j}^{*}=1, \quad i, j=1, \ldots, n o
$$

Thus, $P^{*}$ is a stochastic matrix.
Note that the ( $\mathrm{i}, \mathrm{j}$ ) th element of the matrix $p k$ is the probability that the state will be $j$ after $k$ transitions starting from state $i$. With this in mind, it can be seen from the definition (D.5) that $p_{i j}^{*}$ can be interpreted as the long term fraction of time that the state is $j$ given that the initial state is i. This suggests that for any two states i and if in the same recurrent class we have $p_{i j}^{*}=p_{i i^{\prime}}^{*} \mathbf{j}^{\prime}$ and this can indeed be proved. In particular, if a Markov chain is irreducible, the matrix $P^{*}$ has identical rows. Also, if $j$ is a transient state, we have

$$
p_{i j}^{*}=0, \quad \text { for all } \mathrm{i}=1, \ldots, n
$$

so the columns of the matrix $P^{*}$ corresponding to transient states consist of zeroes.

## FIRST PASSAGE TIMES

Let us denote by $q_{i j}^{k}$ the probability that the state will be $j$ for the first time after exactly $k \geq 1$ transitions given that the initial state is $i$, that is,

$$
q_{i j}^{k}=P\left(X k=j, X_{m} \neq j, 1 \leq \mathrm{m}<k \mid X Q=\mathrm{i}\right) .
$$

Denote also, for fixed $\boldsymbol{i}$ and $j$,

$$
K_{i j}=\min \{\mathrm{k} \geq 1 \mid X k=j, X Q=\mathrm{i}\}
$$

Then Kij, called the first passage time from $i$ to $j$, may be viewea co a random variable. We have, for every $k=1,2, \ldots$,

$$
P\left(K_{i j}=k\right)=q_{i j}^{k}
$$

and we write

$$
P(K i j=00)=P(X k \neq j, k=1,2, \ldots \mid X Q=\mathrm{i})=1-\sum_{k=1}^{\infty} q_{i j}^{k} .
$$

Note that it is possible that $\sum_{k=1}^{\infty} q_{i j}^{k}<1$. This will occur, for example, if $j$ cannot be reached from i , in which case $q_{i j}^{k}=0$ for all $k=1,2, \ldots$ The mean first passage time from $\boldsymbol{i}$ to $j$ is the expected value of $K i j$ :

$$
E\left\{K_{i j}\right\}= \begin{cases}\sum_{k=1}^{\infty} k q_{i j}^{k} & \text { if } \sum_{k=1}^{k} q_{i j}^{k}=1 \\ \infty & \text { if } \sum_{k=1}^{\infty} q_{i j}^{k}<1\end{cases}
$$

It may be proved that if $i$ and $j$ belong to the same recurrent class then

$$
E\left\{K_{i j}\right\}<\infty
$$

In fact if there is only one recurrent class and $t$ is a state of that class, the mean first passage times $E\{$ Kit $\}$ are the unique solution of the following linear system of equations
see Example 7.2.1. If $i$ and $j$ belong to two different recurrent classes, then $E\left\{K^{\prime} i j\right\} E\left\{K_{j i}\right\} \quad$ oo. If $\boldsymbol{i}$ belongs to a recurrent class and $j$ is transient, we have $E\{K i j\}=0$.

## APPENDIX

## Kalman Filtering

In this appendix we present the basic principles of least-squares estimation and their application in estimating the state of a linear discrete-time dynamic system using measurements that are linear in the state variables.

Fundamentally, the problem is the following. There are two random vectors $x$ and $y$, which are related through their joint probability distribution so that the value of one provides information about the value of the other. We get to know the value of $y$, and we want to estimate the value of $x$ so that the average squared error between $x$ and its estimate is minimized. A related problem is to find the best estimate of $x$ within the class of all estimates that are linear in the measured vector $y$. We will specialize these problems to a case where there is an underlying linear dynamic system. In particular, we will estimate the state of the system using measurements that are obtained sequentially in time. By exploiting the special structure of the problem, the computation of the state estimate can be organized conveniently in a recursive algorithm the Kalman filter.

## E. 1 LEAST-SQUARES ESTIMATION

Consider two jointly distributed random vectors $x$ and $y$ taking values in $\Re^{n}$ and $\Re^{m}$, respectively. We view $y$ as a measurement that provides some information about $x$. Thus, while prior to knowing your estimate of $x$ may
have been the expected value $E\{x\}$, once the value of $y$ is known, we want to form an updated estimate $x(y)$ of the value $x$. This updated estimate depends, of course, on the value of $y$, so we are interested in a rule that gives us the estimate for each possible value of $y$, i.e., we are interested in a function $x\left({ }^{\prime}\right)$, where $x(y)$ is the estimate of $x$ given $y$. Such a function $x(\cdot): \Re^{m} \rightarrow \Re^{m}$ is called an estimator. We are seeking an estimator that is optimal in some sense and the criterion we shall employ is based on minimization of

$$
\begin{equation*}
\underset{x, Y}{E}\{l l x-x(y) 112\} . \tag{E.l}
\end{equation*}
$$

 more, throughout the appendix, we assume that all encuuntered expected values are finite.

An estimator that minimizes the expected squared error above over all $x(\cdot): \Re^{n} \rightarrow \Re^{m}$ is called a least-squares estimator and is denoted by $x^{*}(\cdot)$. Since

$$
\underset{x, y}{E}\left\{\|x-x(y)\|^{2}\right\}=\underset{y}{E}\left\{\frac{E}{x}\left\{\|x-x(y)\|^{2} \mid y\right\}\right\}
$$

it is clear that $x^{*}($.$) is a least-squares estimator if x^{*}(y)$ minimizes the conditional expectation in the right-hand side above for every y $\mathrm{E} \Re^{m}$, that is,

$$
\begin{equation*}
\underset{x}{E}\left\{\| \mathrm{X}-x^{*}(y) 112 \mid y\right\}=\min _{z \in \Re^{m}} E_{x}\{l l x-z|12| y\}, \quad \text { for all } y \mathrm{E} \Re^{m} \tag{E.2}
\end{equation*}
$$

By carrying out this minimization, we obtain the following proposition.

Prof -. E.1: The least-squares estimator $x^{*}(\cdot)$ is given by

$$
\begin{equation*}
x^{*}(y)=\underset{x}{E}\{x \mid y\}, \quad \text { for all } y \mathrm{E} \Re^{m} \tag{E.3}
\end{equation*}
$$

Proof: We have for every fixed $z \mathrm{E} \Re^{n}$

$$
\underset{x}{E}\{l l x-z \mid 12, y\}=\underset{x}{E}\left\{\| \mathbf{X}_{\|} 2_{1} y\right\}-2 z^{\prime} \underset{x}{E}\{x, y\}+\| z 112
$$

setting to zero the derivative with respect to $z$, we see that the above expression is minimized by $z=E x\{x \mid y\}$, and the result follows. Q.E.D.

## E. 2 LINEAR LEAST-SQUARES ESTIMATION

The least-squares estimator $\operatorname{Ex}\{x \mid y\}$ may be a complicated nonlinear function of $y$. As a result its practical calculation may be difficult. This motivates finding optimal estimators within the restricted class of linear estimators, i.e., estimators of the form

$$
\begin{equation*}
x(y)=A y+b, \tag{EA}
\end{equation*}
$$

where $A$ is an $n \mathrm{x}$ m matrix and $b$ is an n -dimensional vector. An estimator

$$
\hat{x}(y)=\hat{A} y+\hat{b}
$$

where $\hat{A}$ and $\hat{b}$ minimize

$$
\underset{x, Y}{E}\left\{l l x-A y-b \mid \|^{2}\right\}
$$

over all $n \mathrm{x} \mathrm{m}$ matrices $A$ and vectors $b \mathrm{E} \Re^{n}$ is called a linear least-squares estimator.

In the special case where $x$ and yare jointly Gaussian random vectors it turns out that the conditional expectation $\operatorname{Ex}\{x \mid y\}$ is a linear function of $y$ (plUS a constant vector), and as a result, a linear least-squares estimator is also a least-squares estimator. This is shown in the next proposition.

Proposition E.2: If $x$, yare jointly Gaussian random vectors, then the least-squares estimate $\operatorname{Ex}\{x \mid y\}$ of $x$ given $y$ is linear in $y$.

Proof: Consider the random vector $z \mathrm{E} \Re^{n+m}$

$$
z=\binom{x}{y}
$$

and assume that $z$ is Gaussian with mean

$$
\begin{equation*}
\bar{z}=E\{z\}=\binom{E\{x\}}{E\{y\}} \tag{E.5}
\end{equation*}
$$

and covariance matrix

$$
\begin{aligned}
& -\left(\begin{array}{ll}
\Sigma_{x x} & \\
\Sigma_{y x} & \Sigma_{y y}
\end{array}\right) .
\end{aligned}
$$

To our proof we assume that $\Sigma$ is a positive definite symmetric matrix so that it possesses an inverse; the result, however, holds without this assumption. We recall that if $z$ is Gaussian, its probability density function is of the form

$$
p(z)=p(x, y)=c e^{-\frac{1}{2}(z-\bar{z})^{\prime} \Sigma^{-1}(z-\bar{z})},
$$

where

$$
\mathrm{c}=(2 \pi)^{-(n+m) / 2}(\operatorname{det} \Sigma)^{-1 / 2}
$$

and $\operatorname{det} \Sigma$ denotes the determinant of $\Sigma$. Similarly the probability density functions of $x$ and yare of the form

$$
\begin{aligned}
& p(x)=c_{1} e^{-\frac{1}{2}(x-\bar{x})^{\prime} \sum_{x x}^{-1}(x-\bar{x})} \\
& p(y)=c_{2} e^{-\frac{1}{2}(y-\bar{y})^{\prime} \sum_{y y}^{-1}(y-\bar{y})}
\end{aligned}
$$

where Cl and C 2 are appropriate constants. By Bayes' rule the conditional probability density function of $x$ given $y$ is

$$
\begin{equation*}
p(x \mid y)=\frac{p(x, y)}{p(y)}=\frac{c}{c_{2}} e^{-\frac{1}{2}\left((z-\bar{z})^{\prime} \Sigma^{-1}(z-\bar{z})-(y-\bar{y})^{\prime} \Sigma_{y y}^{-1}(y-\bar{y})\right)} . \tag{E.7}
\end{equation*}
$$

It can now be seen that there exist a positive definite symmetric $n \times n$ matrix $D$, an $n \mathrm{x}$ m matrix $A$, a vector $b \mathrm{E} \Re^{n}$, and a scalar $s$ such that
$(z-\bar{z})^{\prime} \Sigma^{-1}\left(\begin{array}{ll}z & \bar{z}\end{array}\right)-(y-\bar{y})^{\prime} \Sigma_{y y}^{-1}(y-\bar{y})=(x-A y-b) I D-1(x-A y-b)+s$.
This is because by substitution of the expressions for $\bar{z}$ alld $\Sigma$ of Eqs. (E.5) and (E.6), the left-hand side of Eq. (E.8) becomes a quadratic form in $x$ and $y$, which can be put in the form indicated in the right-hand side of Eq. (E.8). In fact, by computing the inverse of $\Sigma$ using the partitioned matrix inversion formula (Appendix A) it can be verified that $A, b, D$, and $s$ in Eq. (E.8) have the form
$A=\Sigma_{x y} \Sigma_{y y}^{-1}, \quad b=\bar{x}-\Sigma_{x y} \Sigma_{y y}^{-1} \bar{y}, \quad D=\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{y x}, \quad s=0$.
Now it follows from Eqs. (E.8) and (E.7) that the conditional expectation $E x\{x \mid y\}$ is of the form $A y+b$, where $A$ is some $n \times m$ matrix and $b \mathrm{E} \Re^{n}$. Q.E.D.

We now turn to the characterization of the linear least-squares estimator.
$\underline{\text { Prof }} \quad \mathbb{E} .3:$ Let $x, y$ be random vectors taking values in $\Re^{n}$ and $\Re^{m}$, respectively, with given joint probability distribution. The expected values and covariance matrices of $x$, yare denoted by

$$
E\{x\}=\bar{x} \quad E\{y\}=\bar{y},
$$

$$
\begin{array}{ll}
E\left\{(x-\bar{x})(x-\bar{x})^{\prime}\right\}=\Sigma_{x x}, & E\left\{(y-\bar{y})(y-\bar{y})^{\prime}\right\}= \\
E\left\{(x-\bar{x})(y-\bar{y})^{\prime}\right\}=\Sigma_{x y}, & E\left\{(y-\bar{y})(x-\bar{x})^{\prime}\right\}=\Sigma_{x y}^{\prime} \tag{E.II}
\end{array}
$$

and we assume that $\Sigma_{y y}$ is invertible. Then the linear least-squares estimator of $x$ given $y$ is

$$
\hat{x}(y)=\bar{x}+\Sigma_{x y} \Sigma_{y y}^{-1}(y-\bar{y}) .
$$

The corresponding error covariance matrix is given by

$$
\begin{equation*}
\underset{x, y}{E}\{(x-\hat{x}(y))(x-X(y)) /\}=\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{y x} \tag{E.13}
\end{equation*}
$$

Proof: The linear least-squares estimator is defined as

$$
\hat{x}(y)=\hat{A} y+\hat{b}
$$

where $\quad \hat{b}$ minimize the function $\left.f(A, b)=E x, y_{f} \| \mathrm{X}-A y-b l^{2}\right\}$ over $A$ and $b$. Taking the derivatives of $f(A, b)$ with respect to $A$ and $b$ and setting them to zero, we obtain the two conditions

$$
\begin{gather*}
O=\left.\frac{\partial f}{\partial A}\right|_{\hat{A}, \hat{b}}=2 \underset{x, y}{E}\{(\hat{b}+\hat{A} y-X) y /\},  \tag{E.14}\\
0=\left.\frac{\partial f}{\hat{a} b}\right|_{\hat{A}, \hat{b}}=2 \underset{x, y}{E}\{\hat{b}+\hat{A} y-x\} \tag{E.15}
\end{gather*}
$$

The second condition yields

$$
\begin{equation*}
\hat{b}=\bar{x}-\hat{A} \bar{y} \tag{E.16}
\end{equation*}
$$

and by substitution in the first, we obtain

$$
\begin{equation*}
\underset{x, Y}{E}\left\{y(A(y-\bar{y})-(x-\bar{x}))^{\prime}\right\}=0 \tag{E.17}
\end{equation*}
$$

We have

$$
\underset{x, y}{E}\{\hat{A}(y-\bar{y})-(x-\bar{x})\}^{\prime}=0
$$

so that

$$
\begin{equation*}
\bar{y} \underset{x, y}{E}\{\hat{A}(y-\bar{y})-(x-\bar{x})\}!=0 \tag{E.18}
\end{equation*}
$$

subtracting Eq. (E.18) from Eq. (E.17), we obtain

$$
\underset{x, y}{E}\left\{(y-\bar{y})(\hat{A}(y-\bar{y})-(x-\bar{x}))^{\prime}\right\}=0
$$

Equivalently,

$$
\Sigma_{y y} \hat{A}^{\prime}-\Sigma_{y x}=0
$$

from which

$$
\begin{equation*}
\hat{A}=\Sigma_{y, n}^{\prime} \Sigma_{y y}^{-1}=\Sigma_{x y} \Sigma_{y y}^{-1} \tag{E.19}
\end{equation*}
$$

Using the expressions ana q) for $\hat{b}$ and respectively, we obtain

$$
\hat{x}(y)-\hat{A} y+\hat{b}=\bar{x}+\quad \bar{y})
$$

Which was to be proved. The desired Eq. (E.13) for ths error covariance follows upon substitution of the expression for $\hat{x}(y)$ obtaineQ a ${ }^{\text {hone. }}$ Q.E.D.

We list some of the properties of the least-squares estimator as corve laries.

Corollary E.3.1: The linear least-squares estimator is unbiased, i.e.,

$$
E_{y}\{\hat{x}(y)\}=\bar{x}
$$

Proof: This follows from Eq. (E.12). Q.E.D.

Corollary E.3.2: The estimation error $x-\hat{x}(y)$ is uncorrelated with both $y$ and $\hat{x}(y)$, i.e.,

$$
\begin{gathered}
\underset{x, y}{E}\{y(x-x(y))!\}=0, \\
\underset{x, y}{E}\{x(y)(x-x(y))!\}=0 .
\end{gathered}
$$

Proof: The first equality is Eq. (E.14). The second equality can be written as

$$
\underset{x, y}{E}\left\{(\hat{A} y+\hat{b})(x-\hat{x}(y))^{\prime}\right\}=0
$$

and follows from the first equality and Cor. E.3.1. Q.E.D.
Corollary E.3.2 is known as the orthogonal projection principle. It states a property that characterizes the linear least-squares estimate and forms the basis for an alternative treatment of least-squares estimation as a problem of projection in a Hilbert space of random variables (see Luenberger [Lue69]).

Corollary E.3.3: Consider in addition to $x$ and $y$, the random vector $z$ defined by

$$
z=C x
$$

where $C$ is a given $p \mathrm{xm}$ matrix. Then the linear least-squares estimate of $z$ given $y$ is

$$
\hat{z}(y)=C \hat{x}(y),
$$

and the corresponding error covariance matrix is given by

$$
\underset{z, y}{E}\left\{(z-\hat{z}(y))(z-\hat{z}(y))^{\prime}\right\}=C \underset{x, y}{E}\{(x-\hat{x}(y))(x-x(y))!\} C!
$$

Proof: We have $E\{z\}=\bar{z}=C \bar{x}$ and

$$
\begin{aligned}
&\left.\quad(z-\bar{z})(z-\bar{z})^{\prime}\right\} \quad C \Sigma_{x x} C^{\prime} \\
& \Sigma_{z y}= \underset{z, y}{E}\left\{(z-\bar{z})(y-\bar{y})^{\prime}\right\}=C \Sigma_{x y}, \\
& \Sigma_{y z}= \Sigma_{z y}^{\prime}=\Sigma_{y x} C^{\prime} .
\end{aligned}
$$

By Prop. E. 3 we have

$$
\begin{aligned}
& \hat{z}(y)=\bar{z}+\Sigma_{z y} \Sigma_{y y}^{-1}(y-\bar{y})= C \bar{x}+C \Sigma_{x y} \Sigma_{y y}^{-1}(y-\bar{y})=C \hat{x}(y), \\
&{\underset{x, y}{E}\{(z-\hat{z}(y))(z-i C Y))!\}=} \begin{aligned}
& E_{x} \\
& C\left(\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{y x}\right) \mathrm{C}! \\
&= \mathrm{C}_{x, Y}\left\{(x-x(y))(x-x(y))^{\prime}\right\} C^{\prime \prime}
\end{aligned}
\end{aligned}
$$

Corollary 雨.3.4: Consider in addition to $x$ and $y$, an additional random vector $z$ of the form

$$
\begin{equation*}
z=C y+u \tag{E.20}
\end{equation*}
$$

where C is a given $p \mathrm{x} \mathrm{m}$ matrix of $\operatorname{rank} p$ and $u$ is a given vector in $\Re^{p}$. Then the linear least-squares estimate $\hat{x}(z)$ of $x$ given $z$ is

$$
\begin{equation*}
\hat{x}(z)=\bar{x}+\quad \quad^{\prime}\left(C \Sigma_{y y} C^{\prime}\right)^{-1}(z-C \bar{y}-u) \tag{E.21}
\end{equation*}
$$

and the corresponding error covariance matrix is

$$
\begin{equation*}
\underset{x, z}{E}\left\{(x-\hat{x}(z))(x-\hat{x}(z))^{\prime}\right\}=\quad-\Sigma_{x y} C^{\prime}\left(C^{\prime} \Sigma_{y_{s}} C^{\prime}\right)^{-1} C \Sigma_{y x} \tag{E.22}
\end{equation*}
$$

Proof: We have

$$
\begin{gather*}
\bar{z}=E\{z\}=C \bar{y}+u  \tag{E.23a}\\
\Sigma_{z z}=E\left\{(z \quad \bar{z})(z-\bar{z})^{\prime}\right\}=C \Sigma_{y y} C^{\prime}  \tag{E.23b}\\
\Sigma_{z x}=E\left\{(z-\bar{z})(x-\bar{x})^{\prime}\right\}=C \Sigma_{y x} \tag{E.23c}
\end{gather*}
$$

$$
\begin{equation*}
\left.(x \quad \bar{x})(z-\bar{z})^{\prime}\right\}=\Sigma_{x y} C^{\prime} \tag{E.23~d}
\end{equation*}
$$

From Prop. E. 3 we have

$$
\begin{gather*}
\hat{x}(z)=\bar{x}+\Sigma_{x z} \Sigma_{z z}^{-1}(z-\bar{z})  \tag{E.24a}\\
\underset{x, z}{E}\left\{(x-\hat{x}(z))(x-\hat{x}(z))^{\prime}\right\}=\Sigma_{x x}-\Sigma_{x z} \Sigma_{z z}^{-1} \Sigma_{z x} \tag{E.24b}
\end{gather*}
$$

where $\Sigma_{z z}=C \Sigma_{y y} C^{\prime}$ has an inverse, since $\Sigma_{y y}$ is invertible and $C$ has rank p. substituting the relations (E.23) into Eqs. (E.24a) and (E.24b) the result follows. Q.E.D.

Frequently we want to estimate a vector of parameters $x \mathrm{E} \Re^{n}$ given a measurement vector $z E \Re^{m}$ of the form $z=C x+v$, where $C$ is a given $\mathrm{m} \times n$ matrix, and $v \mathrm{E} \Re^{m}$ is a random measurement error vector. The following corollary gives the linear least-squares estimate $\hat{x}(z)$ and its error covariance.

## Coronary $\mathrm{E} .3 .5:$ Let

$$
z=C x+v
$$

where C is a given mxn matrix, and the random vectors $x \mathrm{E} \Re^{n}$ and $v E \Re^{m}$ are uncorrelated. Denote

$$
\begin{array}{ll}
E\{x\}=\bar{x}, & E\left\{(x \quad \bar{x})(x-\bar{x})^{\prime}\right\}=\Sigma_{x x} \\
E\{v\}=\bar{v}, & E\left\{(v-\bar{v})(v-\bar{v})^{\prime}\right\}=\Sigma_{v v}
\end{array}
$$

and assume further that $\Sigma_{v v}$ is a positive definite matrix. Then

$$
\begin{gathered}
\hat{x}(z)=\bar{x}+\Sigma_{x x} C^{\prime}\left(C \Sigma_{x x} C^{\prime}+\Sigma_{v v}\right)^{-1}(z-C \bar{x}-\bar{v}) \\
\underset{x, v}{E}\left\{(x-\bar{x}(z))(x-\bar{x}(z))^{\prime}\right\}=\Sigma_{x x}-\Sigma_{x x} C^{\prime}\left(C \Sigma_{x x} C^{\prime}+\Sigma_{v v}\right)^{-1} C \Sigma_{x x}
\end{gathered}
$$

## Proof: Define

$$
y=\left(X^{\prime} V /\right)^{\prime}, \quad \bar{y}=\left(\begin{array}{ll}
\bar{x}^{\prime} & \bar{v}^{\prime}
\end{array}\right)^{\prime}, \quad O=\left(\begin{array}{ll}
C & 1
\end{array}\right)
$$

Then we have $z=\tilde{C} y$, and Cor. E.3.3,

$$
\begin{gathered}
\hat{x}(z)=\left(\begin{array}{ll}
1 & 0
\end{array}\right) \hat{y}(z) \\
E\left\{(x-\hat{x}(z))\left(x-\quad '^{\prime}\right\}=\left(\begin{array}{ll}
1 & 0
\end{array}\right) E\{(y-y(z))(y-y(Z)) /\}\binom{I}{0}\right.
\end{gathered}
$$

where $\hat{y}(z)$ is the linear least-squares estimate of $y$ given $z$. By applying Cor. E.3.4 with $u=0$ and $x=y$ we obtain

$$
\begin{gathered}
\hat{y}(z)=\bar{y}+\Sigma_{y y} \tilde{C}^{\prime}\left(\tilde{C} \Sigma_{y y} \tilde{C}^{\prime}\right)^{-1}(z-\tilde{C} \bar{y}) \\
E\{(y-\hat{y}(z))(y-y(Z)) /\}=\Sigma_{y y}-\Sigma_{y y} \tilde{C}^{\prime}\left(\tilde{C} \Sigma_{y y} \tilde{C}^{\prime}\right)^{-1} \tilde{C} \Sigma_{y y}
\end{gathered}
$$

By using the equations

$$
\Sigma_{y y}=\left(\begin{array}{cc}
\Sigma_{x x} & 0 \\
0 & \Sigma_{v v}
\end{array}\right), \quad O=\left(\begin{array}{ll}
C & 1
\end{array}\right)
$$

and by carrying out the straightforward calculation the result follows. Q.E.D.

The next two corollaries deal with least-squares estimates involving multiple measurement vectors that are obtained sequentially. In particular,
the corollaries show how to modify an existing least-squares estiulate $\hat{x}(y)$ to obtain $\hat{x}(y, z)$ once an additional vector $z$ becomes known. This is a central operation in Kalman filtering.

Corollary E.3.6: Consider in addition to $x$ and $y$, an additional random vector $z$ taking values in $\Re^{p}$, which is uncorrelated with $y$. Then the linear least-squares estimate $\hat{x}(y, z)$ of $x$ given $y$ and $z$ [i.e., given the composite vector $(y, z)]$ has the form

$$
\begin{equation*}
\hat{x}(y, z)=\hat{x}(y)+\hat{x}(z)-\bar{x} \tag{E.25}
\end{equation*}
$$

where $\hat{x}(y)$ and $\hat{x}(z)$ are the linear least-squares estimates of $x$ given $y$ and given $z$, respectively. Furthermore,

$$
\underset{x, y, z}{E}\left\{(x-\hat{x}(y, z))(x-\hat{x}(y, z))^{\prime}\right\}=\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{y x}-\Sigma_{x z} \Sigma_{z z}^{-1} \Sigma_{z x}
$$

where

$$
\begin{gathered}
\underset{\boldsymbol{x}, \boldsymbol{z}}{E}\left\{(x-\bar{x})(z-\bar{z})^{\prime}\right\}, \quad \Sigma_{z x}=\underset{X, z}{E}\left\{(z-\bar{z})(x-\bar{x})^{\prime}\right\}, \\
\Sigma_{z z}=\underset{z}{E}\left\{(z-\bar{z})(z-\bar{z})^{\prime}\right\}, \quad \bar{z}=\underset{z}{E}\{z\},
\end{gathered}
$$

and it is assumed that $\Sigma_{z z}$ is invertible.

Proof: Let

$$
w=\binom{y}{z}, \quad \bar{w}=\binom{\bar{y}}{\bar{z}}
$$

(E.12) we have

$$
\begin{equation*}
\hat{x}(w)=\bar{x}+\Sigma_{x w} \Sigma_{w w}^{-1}(w-\bar{w}) \tag{E.27}
\end{equation*}
$$

## Furthermore

$$
\Sigma_{x w}=\left[\Sigma_{x y}, \Sigma_{x z}\right],
$$

and since $y$ and $z$ are uncorrelated, we have

$$
\Sigma_{w w}=\left(\begin{array}{cc}
\Sigma_{y y} & 0 \\
0 & \Sigma_{z z}
\end{array}\right)
$$

Substituting the above expressions in Eq. (E.27), we obtain

$$
\hat{x}(w)=\bar{x}+\quad\left(\begin{array}{ll}
y & \bar{y})+\Sigma_{x z} \Sigma_{z z}^{-1}(z-\bar{z})=\hat{x}(y)+\hat{x}(z)-\bar{x}, ~
\end{array}\right.
$$

and Eq.
is proved. The proof of Eq. (E.26) is similar by using the relations above and the covariance Eq. (E.13). Q.E.D.

Corollary E.3.7: Let $z$ be as in the preceding corollary and assume that $y$ and $z$ are not necessarily uncorrelated, that is, we may have

$$
\Sigma_{y z}=\Sigma_{z y}^{\prime}=\underset{y, z}{E}\left\{(y-\bar{y})(z-\bar{z})^{\prime}\right\} \neq 0
$$

Then

$$
\begin{equation*}
\hat{x}(y, z)=\hat{x}(y)+\hat{x}(z-\hat{z}(y))-\bar{x} \tag{E.28}
\end{equation*}
$$

where $\hat{x}(z-\hat{z}(y))$ denotes the linear least-squares estimate of $x$ given the random vector $z-z(y)$ and $\hat{z}(y)$ is the linear least-squares estimate of $z$ given $y$. Furthermore,

$$
\begin{equation*}
\underset{x, Y, z}{E}\left\{(x-\hat{x}(y, z))(x-\hat{x}(y, z))^{\prime}\right\}=\Sigma_{x x}-\Sigma_{x y} \Sigma_{y y}^{-1} \Sigma_{y x}-\hat{\Sigma}_{x z} \hat{\Sigma}_{z z}^{-1} \hat{\Sigma}_{z x} \tag{E.29}
\end{equation*}
$$

where

$$
\begin{aligned}
& \hat{\Sigma}_{x z}=\underset{x, y, z}{E}\left\{(x-\bar{x})(z-z(y))^{\prime}\right\} \\
& \hat{\Sigma}_{z z}=\underset{y, z}{E}\left\{(z-z(y))(z-z(y))^{\prime}\right\} \\
& \hat{\Sigma}_{z x}=\underset{x, Y, z}{E}\left\{(z-z(y))(x-\bar{x})^{\prime}\right\}
\end{aligned}
$$

Proof: It can be seen that, since $\hat{z}(y)$ is a linear function of $y$, the linear least-squares estimate of $x$ given $y$ and $z$ is the same as the linear leastsquares estimate of $x$ given $y$ and $z-\hat{z}(y)$. By Cor. E.3.2 the random vectors $y$ and $z-\hat{z}(y)$ are uncorrelated. Given this observation the result follows by application of the preceding corollary.

## STATE ESTIMATION - THE KALMAN FILTER

Consider now a linear dynamic system of the type considered in Section 5.2 but without a control vector $(U k \equiv 0)$

$$
\begin{equation*}
x_{k+1}=A_{k} x_{k}+w_{k}, \quad k=0,1, \ldots, \mathbf{N}-1, \tag{E.30}
\end{equation*}
$$

where $x_{k} \mathrm{E} \Re^{n}$ and $w_{k} \mathrm{E} \Re^{n}$ denote the state and random disturbance vectors, respectively, and the matrices $A k$ are known. Consider also the
measurement vaumusu

$$
\begin{equation*}
z_{k}=C_{k} x_{k}+v_{k}, \quad k=0,1, \ldots, \mathbf{N}-1, \tag{E.31}
\end{equation*}
$$

where $Z k E \Re^{s}$ and $V k E \Re^{s}$ are the observation and observation noise vectors, respectively.

We assume that $X o$, $W o, \ldots, W N-l, V o, \ldots, V N-l$ are independent random vectors with given probability distributions and that

$$
\begin{equation*}
E\left\{w_{k}\right\}=E\left\{v_{k}\right\}=0, \quad k=0,1, \ldots, \mathrm{~N}-1 \tag{E.32}
\end{equation*}
$$

We use the notation

$$
\begin{equation*}
S=E\left\{\left(x_{0}-E\{x o\}\right)(x o-E\{x o\})^{\prime}\right\}, \quad M k=E\left\{w_{k} w_{k}^{\prime}\right\}, \quad N k=E\left\{v_{k} v_{k}^{\prime}\right\} \tag{E.33}
\end{equation*}
$$

and we assume that $N k$ is positive definite for all $k$.

## A Nonrecursive Least-Squares Estimate

We first give a straightforward but somewhat tedious method to derive the linear least-squares estimate of $X k+l$ or $X k$ given the values of $Z o, Z l, \cdots, Z k$. Let us denote

$$
Z_{k}=\left(z_{0}^{\prime}, z_{1}^{\prime}, \ldots, z_{k}^{\prime}\right)^{\prime}, \quad r_{k-1}=\left(x_{0}^{\prime}, w_{0}^{\prime}, w_{1}^{\prime}, \ldots, w_{k-1}^{\prime}\right)^{\prime}
$$

In this method, we first find the linear least-squares estimate of $r k-l$ given $Z k$, and we then obtain the linear least-squares estimate of $X k$ given $Z k$ after expressing $X k$ as a linear function of $r k-l$.

For each i with $0 \leq \mathrm{i} \leq k$ we have, by using the system equation,

$$
x_{i+1}=L_{i} r_{i}
$$

where is the $n \times(n(i+1))$ matrix

## I) .

As a result we may write

$$
Z_{k}=\Phi_{k-1} r_{k-1}+V_{k}
$$

where

$$
V_{k}=\left(v_{0}^{\prime}, v_{1}^{\prime}, \ldots, v_{k}^{\prime}\right)^{\prime}
$$

and
is an $s(k+1) \times(n k)$ matrix of the form


We can thus use Cor. E.3.5, the equations above, and the data of the problem to compute

$$
\hat{r}_{k-1}\left(Z_{k}\right) \quad \text { and } \quad E\left\{\left(r_{k-1}-\hat{r}_{k-1}\left(Z_{k}\right)\right)\left(r_{k-1}-\hat{r}_{k-1}\left(Z_{k}\right)\right)^{\prime}\right\}
$$

Let us denote the linear least-squares estimates of $X k+l$ and $X k$ given $Z k$ by $\hat{x}_{k+1 \mid k}$ and $\hat{x}_{k \mid k}$, respectively. We can now obtain $\hat{x}_{k \mid k}=\hat{x}_{k}\left(Z_{k}\right)$ and the corresponding error covariance matrix by using Cor. E.3.3, that is,

$$
\hat{x}_{k \mid k}=L_{k-1} \hat{r}_{k-1}\left(Z_{k}\right)
$$

$$
\begin{aligned}
& E\left\{(X k-x k l k)(X k-x k l k)^{\prime}\right\} \\
& \quad=L k-l E\left\{(r k-l-T k-l(Z k))\left(r k-l-\hat{r}_{k-1}\left(Z_{k}\right)\right)^{\prime}\right\} L_{k-1}^{\prime}
\end{aligned}
$$

These equations may in turn be used to yield and the corresponding error covariance again via Cor. E.3.3.

## The Kalman Filtering Algorithm

The preceding method for obtaining the least-squares estimate of $X k$ is cumbersome when the number of measurements is large. Fortunately, the sequential structure of the problem can be exploited and the computations can be organized conveniently, as first proposed by Kalman [Ka160]. The main attractive feature of the Kalman filtering algorithm is that the estimate $\hat{x}_{k+1 \mid k}$ can be obtained by means of a simple equation that involves the previous estimate $\hat{x}_{k \mid k-1}$ and the new measurement $Z k$ but does not involve any of the past measurements $Z 0, Z l, \ldots, Z k-l$.

Suppose that we have computed the estimate $\hat{x}_{k \mid k-1}$ together with the covariance matrix

$$
\Sigma_{k \mid k-1}=E\left\{\left(x_{k}-\hat{x}_{k \mid k-1}\right)\left(x_{k}-\hat{x}_{k \mid k-1}\right)^{\prime}\right\}
$$

At time $k$ we receive the additional measurement

$$
z_{k}=C_{k} x_{k}+v_{k}
$$

We may use now Cor. E.3.7 to compute the linear least-squares estimate of $X k$ given $Z k-l=\left(z b, z_{1}^{\prime}, \ldots, z_{k-1}^{\prime}\right)^{\prime}$ and $Z k$. This estimate is denoted by $\hat{x}_{k \mid k}$ and, by Cor. E.3.7, it is given by

$$
\begin{equation*}
\hat{x}_{k \mid k}=\hat{x}_{k \mid k-1}+\hat{x}_{k}\left(z_{k}-\hat{z}_{k}\left(Z_{k-1}\right)\right)-E\left\{x_{k}\right\}, \tag{E.35}
\end{equation*}
$$

where $Z k(Z k-l)$ denotes the linear least-squares estimate of $Z k$ given $Z k-l$ and $\hat{x}_{k}(Z k-Z k(Z k-l))$ denotes the linear least-squares estimate of $X k$ given $\left(Z k-\hat{z}_{k}\left(Z_{k-1}\right)\right)$.

We now calculate the term $\hat{x}_{k}\left(z_{k}-Z k(Z k-l)\right)$ in Eq.
We have by Eqs. (E.31), (E.32), and Cor. E.3.3,

$$
\begin{equation*}
\hat{z}_{k}\left(Z_{k-1}\right)=C_{k} \hat{x}_{k \mid k-1} . \tag{E.36}
\end{equation*}
$$

Also we use Cor. E.3.3 to obtain

$$
\begin{equation*}
E\left\{\left(Z k-\hat{z}_{k}\left(Z_{k-1}\right)\right)(Z k-Z k(Z k-l))^{\prime}\right\}=C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N k^{\prime} \tag{E.37}
\end{equation*}
$$

$E\left\{X k\left(Z k-\hat{z}_{k}\left(Z_{k-1}\right)\right)^{\prime}\right\}$

$$
\begin{aligned}
& =E\left\{X k(C k(X k-X k l k-l))^{\prime}\right\}+E\left\{x_{k} v_{k}^{\prime}\right\} \\
& =E\left\{\left(\begin{array}{ll}
X k & \hat{x}_{k \mid k-1}
\end{array}\right)\left(x_{k}-\hat{x}_{k \mid k-1}\right)^{\prime}\right\} C_{k}^{\prime}+E\left\{\hat{x}_{k \mid k-1}\left(X k-\hat{x}_{k \mid k-1}\right)^{\prime}\right\} C_{k}^{\prime}
\end{aligned}
$$

The last term in the right-hand side above is zero by Cor. E.3.2, so by using Eq. (E.34) we have

$$
\begin{equation*}
E\left\{X k(Z k \quad Z k(Z k-l))^{\prime}\right\} \quad \Sigma_{k \mid k-1} C_{k}^{\prime} \tag{E.38}
\end{equation*}
$$

Using Eqs. (E.36)-(E.38) in Prop. E.3, we obtain

$$
\left(z_{k}-\hat{z}_{k}\left(Z_{k-1}\right)\right)=E\left\{x_{k}\right\}+\Sigma_{k \mid k-1} C_{k}^{\prime}\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N k\right)-l(Z k-C k X k l k-1)
$$

and Eq. (E.35) is written as

$$
\begin{equation*}
\hat{x}_{k \mid k}=x k l k-l+\Sigma_{k \mid k-1} C_{k}^{\prime}\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N k\right)-l(Z k-C k X k l k-1)^{\prime} \tag{E.39}
\end{equation*}
$$

By using Cor. E.3.3 we also have

$$
\begin{equation*}
\hat{x}_{k+1 \mid k}=A_{k} \hat{x}_{k \mid k} . \tag{EAO}
\end{equation*}
$$

Concerning the covariance matrix $\Sigma_{k+1 \mid k}$, we have from Eqs. (E.30), (E.32), and Cor. E.3.3,

$$
\begin{equation*}
\Sigma_{k+1 \mid k}=A_{k} \Sigma_{k \mid k} A_{k}^{\prime}+M_{k} \tag{EA1}
\end{equation*}
$$

where

$$
\Sigma_{k \mid k}=E\left\{\left(x_{k}-\hat{x}_{k \mid k}\right)\left(x_{k}-\hat{x}_{k \mid k}\right)^{\prime}\right\} .
$$

The error covariance matrix may be computed via Cor. E.3.7 similar to $\hat{x}_{k \mid k}$ [ef. Eq. (E.35)]. Thus, we have from Eqs. (E.29), (E.37),

$$
\begin{equation*}
\Sigma_{k \mid k}=\Sigma_{k \mid k-1}-\Sigma_{k \mid k-1} C_{k}^{\prime}\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N_{k}\right)^{-1} C_{k} \Sigma_{k \mid k-1} \tag{E.42}
\end{equation*}
$$

Equations (E.39)-(EA2) with the initial conditions [ef. Eq. (E.33)]

$$
\begin{equation*}
\hat{x}_{0 \mid-1}=E\{x o\}, \quad \Sigma_{0 \mid-1}=5 \tag{EA3}
\end{equation*}
$$

constitute the Kalman jiltering algorithm. This algorithm recursively generates the linear least-squares estimates $\hat{x}_{k+1 \mid k}$ or $\hat{x}_{k \mid k}$ together with the associated error covariance matrices $\Sigma_{k+1 \mid k}$ or $\Sigma_{k \mid k}$. In particular, given $\Sigma_{k \mid k-1}$ and $\hat{x}_{k \mid k-1}$, Eqs. (E.39) and (E.42) yield $\Sigma_{k \mid k}$ and $\hat{x}_{k \mid k}$, and then Eqs. (EA1) and (EAO) yield $\Sigma_{k+1 \mid k}$ and $\hat{x}_{k+1 \mid k}$.

An alternative expression for Eq. (E.39) is

$$
\begin{equation*}
\hat{x}_{k \mid k}=A_{k-1} \hat{x}_{k-1 \mid k-1}+\Sigma_{k \mid k} C_{k}^{\prime} N_{k}^{-1}\left(z_{k}-C_{k} A_{k-1} \hat{x}_{k-1 \mid k-1}\right), \tag{E.44}
\end{equation*}
$$

which can be obtained from Eqs. (E.39) and (EAO) by using the following equality

$$
\begin{equation*}
\Sigma_{k \mid k} C_{k}^{\prime} N_{k}^{-1}=\Sigma_{k \mid k-1} C_{k}^{\prime}\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N k\right)-l \tag{EA5}
\end{equation*}
$$

This equality may be verified by using Eq. (EA2) to write

$$
\Sigma_{k \mid k} C_{k}^{\prime} N_{k}^{-1}=\left(\Sigma_{k \mid k-1}-\Sigma_{k \mid k-1} C_{k}^{\prime}\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N_{k}\right)^{-1} C_{k} \Sigma_{k \mid k-1}\right) C_{k}^{\prime} N_{k}^{-1}
$$

$$
=\Sigma_{k \mid k-1} C_{k}^{\prime}\left(N_{k}^{-1}-\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N_{k}\right)^{-1} C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime} N_{k}^{-1}\right)
$$

and then use in the above formula the following calculation

$$
\begin{aligned}
N_{k}^{-1} & =\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N k\right)-l\left(C_{k} \Sigma_{k \mid \bar{k}-1} C_{k}^{\prime}+N_{k}\right) N_{k}^{-1} \\
& =\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N_{k}\right)^{-1}\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime} N_{k}^{-1}+1\right) .
\end{aligned}
$$

When the system equation contains a control vector $U k$,

$$
X k+l=A k X k+B_{k} u_{k}+W k, \quad k=0,1, \ldots, N-1,
$$

it is straightforward to show that Eq. (EA4) takes the form

$$
\begin{align*}
\hat{x}_{k \mid k}= & A k-1 X_{k}-l l k-1+\text { Bk-1 Uk-1 }  \tag{EA6}\\
& +\Sigma_{k \mid k} C_{k}^{\prime} N_{k}^{-1}\left(z_{k}-\text { CkAk-1Xk-1Ik-1 }-C_{k} B_{k-1} u_{k-1}\right)
\end{align*}
$$

where $\hat{x}_{k} I k$ is the linear least-squares estimate of $x k$ given $Z 0, Z l, \ldots, Z k$ and $U 0, \mathrm{U} 1, \cdots, U k-I^{\prime}$ The equations (EA1)-(EA3) that generate $\Sigma_{k \mid k}$ remain unchanged.

## Sec. E. 4

## Steady-State Kalman Filtering Algorithm

Finally we note that Eqs. (EA1) and (EA2) yield
$\Sigma_{k+1 \mid k}=A k\left(\Sigma_{k \mid k-1}-\Sigma_{k \mid k-1} C_{k}^{\prime}\left(C_{k} \Sigma_{k \mid k-1} C_{k}^{\prime}+N_{k}\right)^{-1} C_{k} \Sigma_{k \mid k-1}\right) A_{k}^{\prime}+M k^{\prime}$
with the initial condition $\Sigma_{0 \mid-1}=S$. This equation is a matrix Riccati equation of the type considered in Section 4.1. Thus when $A k, C k, N k^{\prime}$ and $M_{k}$ are constant matrices,

$$
A_{k}=A, \quad C_{k}=C, \quad N_{k}=N, \quad M_{k}=M, \quad k=0,1, \ldots, N-1,
$$

we have by invoking the proposition proved there, that $\Sigma_{k+1 \mid k}$ tends to a positive definite symmetric matrix $\Sigma$ that solves the algebraic Riccati equation

$$
\Sigma=A\left(\Sigma-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1} C \Sigma\right) A^{\prime}+N I l
$$

assuming observability of the pair $(A, C)$ and controllability of the pair $(A, D)$, where $\mathrm{M}=D D^{\prime}$. Under the same conditions, we have $\Sigma_{k \mid k} \rightarrow \bar{\Sigma}$, where from Eq. (EA2),

$$
=\Sigma-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1} C \Sigma
$$

We may then write the Kalman filter recursion [ef. Eq. (EA4)] in the asymptotic form

$$
\hat{x}_{k \mid k}=A \hat{x}_{k-1 \mid k-1}+\bar{\Sigma} C^{\prime} N^{-1}\left(z_{k}-C A \hat{x}_{k-1 \mid k-1}\right)
$$

This estimator is simple and convenient for implementation.

## STABILITY ASPECTS

Let us consider now the stability properties of the steady-state form of the Kalman filter. From Eqs. (E.39) and (EAO), we have

$$
\begin{equation*}
\hat{x}_{k+1 \mid k}=A \hat{x}_{k \mid k-1}+A \Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1}\left(z_{k}-C \hat{x}_{k \mid k-1}\right) \tag{EA9}
\end{equation*}
$$

Let $e k$ denote the "one-step prediction" error

$$
e_{k}=x_{k}-\hat{x}_{k \mid k-1} .
$$

By using Eq. (E.49), the system equation

$$
x_{k+1}=A x_{k}+w_{k}
$$

and the measurement equation

$$
z_{k}=C x_{k}+v_{k}
$$

we obtain

$$
e_{k+1}=\left(A-A \Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1} C\right) e_{k}+w_{k}-A \Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1} v_{k}
$$

From the practical point of view it is important that the error equation (E.50) represents a stable system, that is, the matrix

$$
\begin{equation*}
A-A \Sigma C^{\prime}\left(C \Sigma C^{\prime}+\mathrm{N}\right)-1 C \tag{E.51}
\end{equation*}
$$

has eigenvalues strictly within the unit circle. This, however, follows by Prop. 4.4.1 of Section 4.1 under the observability and controllability assumptions given earlier, since $\Sigma$ is the unique positive semidefinite symmetric solution of the algebraic Riccati equation

$$
\Sigma=A\left(\Sigma-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1} C \Sigma\right) A^{\prime}+M
$$

Actually this proposition yields that the transpose of the matrix (E.51) has eigenvalues strictly within the unit circle, but this is sufficient for our purposes since the eigenvalues of a matrix are the same as those of its transpose.

Let us consider also the stability properties of the equation governing the estimation error

$$
\tilde{e}_{k}=X k-\hat{x}_{k \mid k}
$$

We have by a straightforward calculation

$$
\tilde{e}_{k}=\left(I-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)-l C\right) e k-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1} v_{k}
$$

By multiplying both sides of Eq. (E.50) by I $\Sigma C^{\prime}\left(C \Sigma C^{\prime}+1 \mathrm{~V}\right)-1 C$ and by using Eq. (E.52), we obtain

$$
\begin{aligned}
\tilde{e}_{k+1} & +\Sigma C^{\prime}\left(C \Sigma C^{\prime}+\mathrm{N}\right)-1 \mathrm{vk}+1 \\
& =\left(A-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)-l C A\right)\left(\tilde{e}_{k}+\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)-l V k\right) \\
& +\left(\mathrm{I}-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+\mathrm{N}\right)-1 \mathrm{C}\right)\left(W k-A \Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)--l v k\right),
\end{aligned}
$$

or equivalently

$$
\begin{aligned}
\tilde{e}_{k+1}= & \left(A-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)-l C A\right) e k \\
& +\left(\mathrm{I}-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)-l C\right) W k-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+\mathrm{N}\right)-\mathrm{lVk}+1 .
\end{aligned}
$$

Since the matrix has eigenvalues strictly within the unit circle, the sequence $\left\{e_{k}\right\}$ generated by Eq. (E.50) tends to zero whenever the vectors $w k$ and $v k$ are identically zero for all $k$. Hence, by Eq. (E.52), the same is true for the sequence $\left\{\tilde{e}_{k}\right\}$. It follows from Eq. (E.53) that the matrix

$$
A-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)-I C A
$$

(E.54)
has eigenvalues strictly within the unit circle, and the estimation error sequence $\left\{\tilde{e}_{k}\right\}$ is generated by a stable system.

Let us finally consider the stability properties of the 2 n -dimensional system of equations with state vector $\left(x_{k}^{\prime}, \hat{x}_{k}^{\prime}\right)$ :

$$
\begin{gather*}
x_{k+1}=A x_{k}+B L \hat{x}_{k}  \tag{E.55}\\
\hat{x}_{k+1}=\bar{\Sigma} C^{\prime} N^{-1} C A x_{k}+\left(A+B L-\bar{\Sigma} C^{\prime} N^{-1} C A\right) \hat{x}_{k} \tag{E.56}
\end{gather*}
$$

This is the steady-state, asymptotically optimal closed-loop system that was encountered at the end of Section 5.2.

We assume that the appropriate observability and controllability assumptions stated there are in effect. By using the equation

$$
\bar{\Sigma} C^{\prime} N^{-1}=\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1}
$$

shown earlier, we obtain from Eqs. (E.55) and (E.56) that

$$
\left(x_{k+1}-\hat{x}_{k+1}\right)=\left(A-\Sigma C^{\prime}\left(C \Sigma C^{\prime}+N\right)^{-1} C A\right)\left(x_{k}-\hat{x}_{k}\right)
$$

Since we have proved that the matrix (E.54) has eigenvalues strictly within the unit circle, it follows that

$$
\begin{equation*}
\lim _{k \rightarrow \infty}\left(X k+l \quad \hat{x}_{k+1}\right)=0 \tag{E.57}
\end{equation*}
$$

for arbitrary initial states $X o$ and $\hat{x}_{0}$. From Eq. (E.55) we obtain

$$
\begin{equation*}
X k+l=(A+B L) X k+B L\left(\hat{x}_{k}-X k\right) \tag{E.58}
\end{equation*}
$$

Since in accordance with the theory of Section 4.1 the matrix $(A+B L)$ has eigenvalues strictly within the unit circle, it follows from Eqs. (E.57) and (E.58) that we have

$$
\begin{equation*}
\lim _{k \rightarrow \infty} X k=0 \tag{E.59}
\end{equation*}
$$

and hence from Eq. (E.57),

$$
\begin{equation*}
\lim _{k \rightarrow \infty} \hat{x}_{k}=0 \tag{E.60}
\end{equation*}
$$

Since the equations above hold for arbitrary initial states $X 0$ and $\hat{x}_{0}$ it follows that the system defined by Eqs. (E.55) and (E.56) is stable.

## E. 5 GAUSS-MARKOV ESTIMATORS

Suppose that we want to estimate a vector $x \mathrm{E} \Re^{n}$ given a measurement vector $z \mathrm{E} \Re^{m}$ that is related to $x$ by

$$
\begin{equation*}
z=C x+v \tag{E.61}
\end{equation*}
$$

where C is a given $\mathrm{m} \times n$ matrix with rank m , and $v$ is a random measurement error vector. Let us assume that $v$ is uncorrelated with $x$, and has a known mean and a positive definite covariance matrix

$$
\begin{equation*}
E\{v\}=v, \quad E\left\{(v-\bar{v})(v-\bar{v})^{\prime}\right\}=\Sigma_{v v} . \tag{E.62}
\end{equation*}
$$

If the a priori probability distribution of $x$ is known, we can obtain a linear least-squares estimate of $x$ given $z$ by using the theory of Section E. 2 (d. Cor. E.3.5). In many cases, however, the probability distribution of $x$ is unknown. In such cases we can use the Gauss-Markov estimator, which is optimal within the class of linear estimators that satisfy certain restrictions, as described below.

Let us consider an estimator of the form

$$
\hat{x}(z)=\hat{A}(z-\bar{v})
$$

where $\hat{A}$ minimizes

$$
\begin{equation*}
\left.f(A)=\underset{X, Z}{E}\{\| X-\quad-\bar{v}) \|^{2}\right\} \tag{E.63}
\end{equation*}
$$

over all $n \mathrm{x}$ m matrices $A$. Since $x$ and $v$ are uncorrelated, we have using Eqs. (E.61)-(E.63)

$$
\begin{aligned}
f(A) & \left.=E_{x, v}\{\| X-A C x-\quad-\bar{v}) \|^{2}\right\} \\
& \left.=E_{x}\{I I U-A C) x I^{2}\right\}+E_{v}\{\| A(v-v) 112\}
\end{aligned}
$$

where $I$ is the $n \mathrm{x} n$ identity matrix. Since $f(A)$ depends on the unknown statistics of $x$, we see that the optimal matrix $\hat{A}$ also depends on these statistics. We can circumvent this difficulty by requiring that
$A C=I$.
Then our problem becomes

$$
\begin{aligned}
& \text { minimize } \quad \underset{v}{E}\left\{I I A(v-\bar{v}) \|^{2}\right\} \\
& \text { subject to } A C=I .
\end{aligned}
$$

Note that the requirement $A C=I$ is not only convenient analytically, but also makes sense conceptually. In particular, it is equivalent to requiring that the estimator $x(z)=-\bar{v})$ be unbiased in the sense that

$$
E\{x(z)\}=E\{x\}=\bar{x}, \quad \text { for all } \bar{x} \mathrm{E} \Re^{n} .
$$

This can be seen by writing

$$
E\{x(z)\}=E\{A(C x+v-\bar{v})\}=A C E\{x\}=A C \bar{x}=\bar{x} .
$$

To derive the optimal solution $\hat{A}$ of problem (E.64), let $a_{i}^{\prime}$ denote the ith row of $A$. We have

$$
\begin{aligned}
& \|A(v-\bar{v})\|^{2}=(v-\bar{v})^{\prime}\left(\begin{array}{ll}
a_{1} & \left.a_{n}\right)\left(\begin{array}{c}
a_{1}^{\prime} \\
\vdots \\
a_{n}^{\prime}
\end{array}\right)(v-\bar{v}), ~
\end{array}\right. \\
& =\sum_{\mathrm{i}=1}^{n}(v-\bar{v})^{\prime} a_{i} a_{i}^{\prime}(v-\bar{v}) \\
& { }^{n} a_{i}^{\prime}(v-\bar{v})(v-\bar{v})^{\prime} a_{i} . \\
& i=1
\end{aligned}
$$

Hence, the minimization problem (E.64) can also be written as

$$
\begin{aligned}
& \operatorname{minimize} \quad \sum_{i=1}^{n} a_{i}^{\prime} \Sigma_{v v} a_{i} \\
& \text { subject to } \quad \mathrm{C}^{\prime} a i=e i, \quad \mathrm{i}=1, \ldots, n,
\end{aligned}
$$

where $e_{i}$ is the ith column of the identity matrix. The minimization can be carried out separately for each i, yielding

$$
\hat{x}_{i}=\Sigma_{v v}^{-1} C\left(C^{\prime} \Sigma_{v v} C\right)^{-1} e_{i}, \quad 1, \ldots ; n
$$

and finally

$$
\hat{A}=\left(C^{\prime} \Sigma_{v v}^{-1} C\right)^{-1} C^{\prime} \Sigma_{v v}^{-1}
$$

Thus, the Gauss-Markov estimator is given by

$$
\begin{equation*}
\hat{x}(z)=\left(C^{\prime} \Sigma_{v v}^{-1} C\right)^{-1} C^{\prime} \Sigma_{v v}^{-1}(z-\bar{v}) . \tag{E.65}
\end{equation*}
$$

Let us also calculate the corresponding error covariance matrix. We have

$$
\begin{aligned}
E\left\{(x-\hat{x}(z))(x-x(z))^{\prime}\right\} & =E\left\{(x-\hat{A}(z-\bar{v}))(x-\hat{A}(z-\bar{v}))^{\prime}\right\} \\
& =E\left\{\hat{A}(v-\bar{v})(v-\bar{v})^{\prime} \hat{A}^{\prime}\right\} \\
& =\hat{A} \Sigma_{v v} \hat{A}^{\prime} \\
& =\left(C^{\prime} \Sigma_{v v}^{-1} C\right)^{-1} C^{\prime} \Sigma_{v v}^{-1} \Sigma_{v v} \Sigma_{v v}^{-1} C\left(C^{\prime} \Sigma_{v v}^{-1} C\right)^{-1}
\end{aligned}
$$

and finally

$$
\begin{equation*}
E\left\{(x-\hat{x}(z))(x-\hat{x}(z))^{\prime}\right\}=\left(C^{\prime} \Sigma_{v v}^{-1} C\right)^{-1} . \tag{E.66}
\end{equation*}
$$

Finally, let us compare the Gauss-1\IIarkov estimator with the linear least-squares estimator of Cor. E.3.5. Assuming that $\Sigma_{x x}$ is invertible, a straightforward calculation shows that the latter estimator can be written as

$$
\begin{equation*}
\hat{x}(z)=\bar{x}+\left(\Sigma_{x x}^{-1}+C^{\prime} \Sigma_{v v}^{-1} C\right)^{-1} C^{\prime} \Sigma_{v v}^{-1}(z-C \bar{x}-\bar{v}) \tag{E.67}
\end{equation*}
$$

By comparing Eqs. (E.65) and (E.67), we see that the Gauss-Markov estio mator is obtained from the linear least-squares estimator by setting $\bar{x}=$ and $\Sigma_{x x}^{-1}=0$, i.e., a zero mean and infinite covariance for the unknown random variable $x$. Thus, the Gauss-Markov estimator may be viewed as a limiting form of the linear least-squares estimator. The error covariance matrix (E.66) of the Gauss-Markov estimator is similarly related with the error covariance matrix of the linear least-squares estimator.

## E. 6 DETERMINISTIC LEAST-SQUARES ESTIMATION

Suppose again that we want to estimate a vector $x \mathrm{E} \Re^{n}$ given a measurement vector $z \mathrm{E} \Re^{m}$ that is related to $x$ by

$$
z=C x+v
$$

where C is a known $\mathrm{m} x \mathrm{n}$ matrix of rank m . However, we know nothing about the probability distribution of $x$ and $v$, and thus we can't use a statistically-based estimator. Then it is reasonable to select as our estimate the vector $\hat{x}$ that minimizes

$$
f(x)=\operatorname{liz}-\mathrm{CX} \|^{2},
$$

that is, the estimate that fits best the data in a least-squares sense. We denote this estimate by $\hat{x}(z)$.

By setting to zero the gradient of $f$ at $\hat{x}(z)$, we obtain

$$
\left.\nabla f\right|_{\hat{x}(z)}=2 C^{\prime}(C x(z)-z)=0
$$

from which

$$
\begin{equation*}
\hat{x}(z)=\left(C^{\prime} C\right)-l C^{\prime} Z \tag{E.68}
\end{equation*}
$$

An interesting observation is that the estimate (E.68) is the same as the Gauss-Markov estimate given by Eq. (E.65), provided the measurement error has zero mean and covariance matrix equal to the identity, i.e., $\bar{v}=0$, $\Sigma_{v v}=I$. In fact, if instead of liz - CXll ${ }^{2}$, we minimize

$$
(z-\bar{v}-C x)^{\prime} \Sigma_{v v}^{-1}(z-\bar{v}-C x)
$$

then the deterministic least-squares estimate obtained is identical to the Gauss-Markov estimate. If instead of liz - Cxl1 ${ }^{2}$ we minimize

$$
(x-\bar{x})^{\prime} \Sigma_{x x}^{-1}(z-\bar{x})+(z-\bar{v}-C x)^{\prime} \Sigma_{v v}^{-1}(z-\bar{v}-C x),
$$

then the estimate obtained is identical to the linear least-squares estimate given by Eq. (E.67). Thus, we arrive at the interesting conclusion that the estimators obtained earlier on the basis of a stochastic optimization framework can also be obtained by minimization of a deterministic measure of fitness of estimated parameters to the data at hand.

## APPENDIX

Modeling
Stochastic Linear
Systems

In this appendix we show how controlled linear time-invariant systems with stochastic inputs can be represented by the ARMAX model used in Section 5.3.

## F. 1 LINEAR SYSTEMS WITH STOCHASTIC INPUTS

Consider a linear system with output $\{Y k\}$, control input $\left\{u_{k}\right\}$, and an additional zero-mean random input $\{W k\}$. We assume that $\{W k\}$ is a stationary (up to second order) stochastic process. That is, $\{W k\}$ is a sequence of random variables satisfying, for all i, $k=0, \pm 1, \pm 2, \ldots$,

$$
E\left\{w_{k}\right\}=0, \quad E\left\{w_{0} w_{i}\right\}=E\left\{w_{k} w_{k+i}\right\}<\infty
$$

(All references to stationary processes in this section are meant in the limited sense just described.) By linearity, $Y k$ is the sum of one sequence $\left\{y_{k}^{1}\right\}$ due to the presence of $\{U k\}$ and another sequence $\left\{y_{k}^{2}\right\}$ due to the presence of $\{W k\}$ :

$$
\begin{equation*}
Y k=y_{k}^{1}+y_{k}^{2} . \tag{F.l}
\end{equation*}
$$

We assume that $y_{k}^{1}$ and $y_{k}^{2}$ are generated by some filters $B_{l}(\mathrm{~s}) / \boldsymbol{A l}(\mathrm{s})$ and $B_{2}(s) / A_{2}(s)$, respectively:

$$
\begin{equation*}
A_{1}(s) y_{k}^{1}=B_{1}(s) u_{k} \tag{F.2a}
\end{equation*}
$$

$$
\begin{equation*}
A_{2}(s) y_{k}^{2}=B 2(S) W k \tag{F.2b}
\end{equation*}
$$

Operating on Eqs. (F.2a) and (F.2b) with $A 2(S)$ and $A 1(s)$, respectively, adding, and using Eq. (F.1), we obtain

$$
\begin{equation*}
A(S) Y k=\underline{B(S) U k}+V k, \tag{F.3}
\end{equation*}
$$

where $\bar{A}(s)=A_{l}(s) A_{2}(s), \bar{B}(s)=A 2(S) B l(s)$, and $\{V k\}$, given by

$$
\begin{equation*}
V k=A l(S) B 2(S) W k, \tag{FA}
\end{equation*}
$$

is cu zero generally correlated, stationary stochastic process.
We are interested in the case where $U k$ is a control input applied after $Y k$ has occurred and has been observed, so that in Eq. (F.2a) we have $B l(O)=0$. Then, we may assume that the polyicmials $\bar{A}(s)$ and $\underline{B(s)}$ have the form

$$
\bar{A}(s)=1+\bar{a}_{1} s+\cdots+\bar{a}_{m_{0}} s^{m_{0}}, \quad \bar{B}(s)=\bar{b}_{1} s+\cdots+\bar{b}_{m_{0}} s^{m_{0}}
$$

for some scalars $\bar{a}_{i}$ and $\bar{b}_{i}$, and some positive integer $m o$.
To summarize, we have constructed a model of the form

$$
\bar{A}(s) y_{k}=\bar{B}(s) u_{k}+v_{k},
$$

where $\bar{A}(s)$ and $\underline{B}(s)$ are polynomials of the preceding form and $\{V k\}$ is some zero-mean, correlated, stationary stochastic process. We now need to model further the sequence $\{V k\}$.

## PROCESSES WITH RATIONAL SPECTRUM

Given a zero-mean, stationary scalar process $\left\{v_{k}\right\}$, denote by $V(k)$ the autocorrelation function

$$
V(k)=E\left\{v_{i} v_{i+k}\right\}, \quad k=0, \pm 1, \pm 2, \ldots
$$

We say that $\{V k\}$ has rational spectrum if the transform of $\{V(k)\}$ defined by

$$
S_{v}(\lambda)=\sum_{k=-\infty}^{\mathrm{o}} V(k) e^{-j k \lambda}
$$

exists for $\lambda E[-1 T, 1 T]$ and can be expressed as

$$
\begin{equation*}
S_{v}(\lambda)=\sigma^{2} \frac{\left|C\left(e^{i \lambda}\right)\right|^{2}}{|n \cdot i \lambda| 1 \eta}, \quad \lambda \in[-\pi, \pi] \tag{F.5}
\end{equation*}
$$

where $\sigma$ is a scalar, $C(z)$ and are some polynomials with real coefficients

$$
\begin{align*}
& \mathrm{C}(z)=1+\mathrm{cl} z+\ldots+c m z^{m}  \tag{F.6a}\\
& D(z)=1+d_{1} z+\cdots+d_{m} z^{m} \tag{F.6b}
\end{align*}
$$

and $D(z)$ has no roots on the unit circle $\{z\|z\|=I\}$.
The following facts are of interest:
(a) If $\{V k\}$ is an uncorrelated process with $V(O)=\sigma^{2}, V(k)=0$ for $k \neq 0$, then

$$
S_{v}(\lambda)=\sigma^{2}, \quad \lambda \in[-\pi, \pi]
$$

and clearly $\{V k\}$ has rational spectrum.
(b) If $\{V k\}$ has rational spectrum $S v$ given by Eq. (F.5), then $S v$ can be written as

$$
S_{v}(\lambda)=\tilde{\sigma}^{2} \frac{\left|\tilde{C}\left(e^{j \lambda}\right)\right|^{2}}{\left|\tilde{D}\left(e^{j \lambda}\right)\right|^{2}}, \quad \lambda \in[-\pi, \pi]
$$

where $\tilde{\sigma}$ is a scalar and $\tilde{C}(z), D(z)$ are unique real polynomials of the form

$$
\begin{aligned}
& \tilde{C}(z)=1+\tilde{c}_{1} z+\cdots+\tilde{c}_{m} z^{m} \\
& \tilde{D}(z)=1+\tilde{d}_{1} z+\cdots+\tilde{d}_{m} z^{m}
\end{aligned}
$$

such that:
(1) $\tilde{C}(z)$ has all its roots outside or on the unit circle, and if $C(z)$ has no roots on the unit circle, then the same is true for $\tilde{C}(z)$.
(2) $D(z)$ has all roots strictly outside the unit circle.

These facts are seen by noting that if $p \neq 0$ is a root of $D(z)$, then $\left|D\left(e^{j \lambda}\right)\right|^{2}=D\left(e^{j \lambda}\right) D\left(e^{-j \lambda}\right)$ contains a factor

$$
\left(1-\rho^{-1} e^{j \lambda}\right)\left(1-\rho^{-1} e^{-j \lambda}\right)=\rho^{-2}\left(\rho-e^{j \lambda}\right)\left(\rho-e^{-j \lambda}\right)
$$

A little reflection shows that the roots of $\tilde{D}(z)$ should be $p$ or $\rho^{-1}$ depending on whether $p$ is outside or inside the unit circle. Similarly, the roots of $\tilde{C}(z)$ are obtained from the roots of $C(z)$. Thus the polynomials $\tilde{C}(z)$ and $\tilde{D}(z)$ as well as $\tilde{\sigma}^{2}$ can be uniquely determined. We may thus assume without loss of generality that $C(z)$ and $D(z)$ in Eq. (F.5) have no roots inside the unit circle.
There is a fundamental result here that relates to the realization of processes with rational spectrum. The proof is hard; see for example, Ash and Gardner [AsG75, pp. 75-76].
$\mathbb{P}^{\text {roposition }} \mathrm{F} .1:$ If $\left\{v_{k}\right\}$ is a zero-mean, stationary stochastic process with rational spectrum

$$
S_{v}(\lambda)=\sigma^{2} \frac{\left|C\left(e^{j \lambda}\right)\right|^{2}}{\left|D\left(e^{j \lambda}\right)\right|^{2}}, \quad \lambda \in[-\pi, \pi]
$$

where the polynomials $O(s)$ and $D(s)$ are given by

$$
O(S)=1+C I S+\ldots+c m s^{m}, \quad D(s)=1+d_{1} s+\cdots+d_{m} s^{m}
$$

and are assumed (without loss of generality) to have no roots inside the unit circle, then there exists a zero-mean, uncorrelated stationary process $\{\mathrm{Ek}\}$ with $E\left\{\epsilon_{k}^{2}\right\}=\sigma^{2}$ such that for all $k$

$$
v_{k}+d_{1} v_{k-1}+\cdots+d_{m} v_{k-m}=\epsilon_{k}+c_{1} \epsilon_{k-1}+\cdots+c_{m} \epsilon_{k-m}
$$

## THE ARMAX MODEL

Let us now return to the problem of representation of a linear system with stochastic inputs. We had arrived at the model

$$
\begin{equation*}
\bar{A}(s) y_{k}=\bar{B}(s) u_{k}+v_{k} \tag{F.7}
\end{equation*}
$$

If the zero-mean stationary process $\{V k\}$ has rational spectrum, the preceding analysis and proposition show that there exists a zero-mean, uncorrelated stationary process $\{E k\}$ satisfying

$$
D(S) V k=O(S) E k,
$$

where $O(s)$ and $D(s)$ are polynomials, and $O(s)$ has no roots inside the unit circle. Operating on both sides of Eq. (F.7) with $D(s)$ and using the relation $D(S) V k$

$$
O(S) E k \text {, we obtain }
$$

$$
\begin{equation*}
A(S) Y k=B(S) U k+O(S) E k \tag{F.8}
\end{equation*}
$$

where $A(s) \quad D(s) A(s)$ and $B(s)=\underline{D(s) B(s) . ~ S i n c e ~} \bar{A}(0)=1, \underline{B(O)}=0$, we can write Eq. (F.8) as

$$
Y k+a_{\mathrm{i}=1}^{\mathrm{m}} a^{\prime} i Y k^{\prime} ' i=\sum_{\mathrm{i}=1}^{\mathrm{m},} b i U k-i+\mathrm{Ek}+\sum_{\mathrm{i}=1}^{\mathrm{m}} C i E k-i
$$

for some integer $m$ and scalars $a^{\prime}, b i, C i, \mathrm{i}=1, \ldots, \mathrm{~m}$. This is the ARMAX model that we have used in Section 5.3.

## APPENDIX

## Formulating Problems

Under Uncertainty

\section*{Decision

## 正

## 正

In this appendix we discuss various approaches for formulating problems of decision under uncertainty. After a brief discussion of the min-max approach, we focus on the expected utility approach, and we show how this approach can be theoretically justified even if the decision maker is sensitive to the "variability" or "risk" associated with the results of different decisions.

## G. 1 THE PROBLEM OF DECISION UNDER UNCERTAINTY

A decision problem in one of its simplest and most abstract forms consists of three nonempty sets $\mathcal{D}, N$, and 0 , a function $f: \mathcal{D} \times N \mapsto 0$, and a complete and transitive relation $\preceq$ on O . Here
$\mathcal{D}$ is the set of possible decisions,
$N$ indexes the uncertainty in the problem and may be called the set of "states of nature,"
$\mathcal{O}$ is the set of outcomes of the decision problem,
$f$ is the function that determines which outcome will result from a given decision and state of nature, i.e., if decision $d \mathrm{E} \mathcal{D}$ is selected and state of nature $n \mathrm{EN}$ prevails, then the outcome $f(d, n) \mathrm{E} \mathcal{O}$ occurs,
$\prec$ is a relation that determines our preference among the outcomes.t Thus, for $01,02 \mathrm{EO}$, by $01 \preceq 02$ we mean that outcome 02 is at least as preferable as outcome 01 . By completeness of the relation, we mean that every two elements of $\mathcal{O}$ are related, i.e., given any $01,02 \mathrm{EO}$, there are three possibilities: either $01 \preceq \mathrm{O} 2$ but not $\mathrm{O}_{2} \preceq 01$, or $02 \preceq 01$ but not $01 \preceq 02$, or both $01 \preceq 02$ and $02 \preceq 01_{1}$. By transitivity we mean that $01 \preceq 02$ and $02 \preceq 03$ implies $01 \preceq 03$ for any three elements $01,02,03 \mathrm{EO}$.

## Example G. 1

Consider an individual that may bet $\$ 1$ on the toss of a coin or not bet at all. If he bets and guesses correctly, he wins $\$ 1$ and if he does not guess correctly, he loses $\$ 1$. Here $\mathcal{D}$ consists of three elements

$$
\mathcal{D}=\{\text { bet on heads, bet on tails, not bet }\},
$$

$N$ consists of two elements

$$
N=\{\text { heads, tails }\}
$$

and 0 consists of three elements, the three possible final fortunes of the player

$$
\mathcal{O}=\{\$ \mathrm{O}, \$ 1, \$ 2\}
$$

The preference relation on $\mathcal{O}$ is the natural one, Le., $\preceq 1, ~ \preceq 2,1 \preceq 2$, and the values of the function $f$ are given by

$$
\begin{array}{r}
f(H, H)=\$ 2, \quad f(T, H)=\$ 0, \quad f(\text { not bet }, H)=\$ 1, \\
f(H, T)=\$ 0, \quad f(T, T)=\$ 2, \quad f(\text { not bet }, T)=\$ 1 .
\end{array}
$$

Now the relative order by which we rank outcomes is usually clear in any given situation. On the other hand, for the decision problem to be completely formulated, we need a ranking among decisions that is consistent in a well-defined sense with our ranking of outcomes. Furthermore, to facilitate a mathematical or computational analysis, this ranking should be determined by a numerical function $F$ that maps the set of decisions $\mathcal{D}$ to the set of real numbers $\Re$ and is such that
$d_{1} \preceq d_{2}$ if and only if $F\left(d_{1}\right) \leq F\left(d_{2}\right)$,
for all $d_{1}, d 2 \mathrm{E} \mathcal{D}$,
$t$ The symbol $\prec$ in this appendix will be used (somewhat loosely) to denote a preference relation within either the set of outcomes or the set of decisions. The precise meaning should be clear from the context, and hopefully the use of the same symbol to denote different preference relations will create no confusion.
where the notation $d_{l} \preceq d 2$ implies that the decision is at least as preferable as the decision $d l$.

It is by no means clear how one should go about determining and characterizing a ranking among decisions. For example, in the gambling example above, different people will have different preferences as to accepting or refusing the gamble. In fact, the method by which one goes from a ranking of outcomes to a ranking of decisions is a central issue in decision theory. There are a number of approaches and viewpoints, and we now proceed to discuss some of these.

## Payoff Functions, Dominant, and $N$ oninferior Decisions

Let us consider the case where it is possible to assign to each element of $\mathcal{O}$ a real number in a way that the order between elements of 0 agrees with the usual order of the corresponding numbers. In particular, we assume that there exists a real-valued function $G: \mathcal{O} \mapsto \Re$ with the property
$G\left(O_{1}\right) \leq G\left(O_{2}\right)$ if and only if $O_{1} \preceq O_{2}, \quad$ for all $01,02 \mathrm{E} \mathcal{O}$. (G.2)
Such a $G$ does not always exist (see Exercise G.2). However, its existence can be guaranteed under quite general assumptions. In particular, one may show that it exists if $\mathcal{O}$ is a countable set. Also if $G$ exists, it is far from unique, since if $\Phi$ is any monotonically increasing function $\Phi: \Re \mapsto \Re$, the composite function $\Phi \cdot G$ [defined by $(\Phi \cdot G)(O)=\Phi(G(O))]$ has the same property (C.2) as $G$. For instance, in the example given earlier, a function $G:\{0,1,2\} \mapsto \Re$ satisfies Eq. (C.2) if and only if $G(O)<G(1)<G(2)$ and there is an infinity of such functions.

For any choice of $G$ satisfying Eq. (G.2), we define the function $J$ : $\mathcal{D} \times \mathcal{N} \mapsto \Re$ by means of

$$
J(d, n)=G(f(d, n))
$$

and call it a payofffunction.
Given a payoff function $J$, it is possible to obtain a complete ranking of decisions by means of a numerical function in the special case of certainty (the case where the set $\mathcal{N}$ of states of nature consists of a single element $\bar{n})$. By defining

$$
F(d)=J(d, \bar{n})
$$

we have
$d l \preceq d 2$ if and only if $F\left(d_{1}\right) \leq F(d 2)$ if and only if $f\left(d_{1}, \bar{n}\right) \preceq f(d 2, \bar{n})$,
and the numerical function $F$ defines a complete ranking of decisions.

In the case where there is uncertainty, i.e., when $N$ contains more than one element, the order on $\mathcal{O}$ induces only a partial order on $\mathcal{D}$ by means of the relations

$$
\begin{align*}
\mathrm{d}_{1} \preceq d_{2} & \text { if and only if } F\left(d_{1}\right) \leq F\left(d_{2}\right), \text { for all } n \mathrm{EN} \\
& \text { if and only if } f\left(d l^{\prime} n\right) \preceq f\left(d_{2}, n\right), \quad \text { for all } n \mathrm{EN.} \tag{G.3}
\end{align*}
$$

In this partial order, it is not necessary that every two elements of $\mathcal{D}$ be related, i.e., for some $d, d^{\prime} E \mathcal{D}$ we may have neither $d \preceq d^{\prime}$ nor $d^{\prime} \preceq d$. If, however, for two decisions $d 1, d 2$ E $\mathcal{D}$ we have $\mathrm{d}_{1} \preceq \mathrm{~d} 2$ in the sense of Eq. (G.3), then we can conclude that $d 2$ is at least as preferable as d1 since the resulting outcome $f(d 2, n)$ is at least as preferable as $f(d 1, n)$, regardless of the state of nature $n$ that will occur.

A decision $d^{*} E \mathcal{D}$ is called a dominant decision if

$$
\mathrm{d} \preceq d^{*}, \quad \text { for all } \mathrm{d} E \mathcal{D}
$$

where $\preceq$ is understood in the sense of the partial order defined by Eq. (G.3). Naturally such a decision need not exist, but if it does exist, then it may be viewed as optimal. Unfortunately. in most problems of interest to an analyst there exists uummant decision. For instance, hiis in in in the gambline $\Xi_{\text {xample }}$ G.1, as the reader can easily verify. In fact no two Jwions are related in the sense of Eq. (G.3) for this example.

In the absence of a dominant decision, one can consider the set $\mathcal{D}_{m} \mathrm{C}$ $\mathcal{D}$ of all noninferior decisions, where $\mathrm{d}_{\mathrm{m}} \mathrm{E} \mathcal{D}_{m}$ if for every $\mathrm{d} \mathrm{E} \mathcal{D}$ the relation $\mathrm{d}_{\mathrm{m}} \preceq d$ implies $\mathrm{d} \preceq \mathrm{d}_{\mathrm{m}}$ in the sense of the partial order defined by Eq. (G.3). In terms of a payoff function J, noninferior decisions may be characterized by
$\mathrm{d}_{\mathrm{m}} \in \mathcal{D}_{m}$ if and only if there is no $\mathrm{d} E \mathcal{D}$ such that

$$
\begin{aligned}
& J\left(d_{m}, n\right) \leq J(d, n) \text { for all } n \mathrm{EN} \text { and } \\
& J\left(d_{m}, n\right)<J(d, n) \text { for some } n \mathrm{EN} .
\end{aligned}
$$

Clearly it makes sense to consider only the decisions in $\mathcal{D}_{m}$ as candidates for optimality since any decision that is not in $\mathcal{D}_{m}$ is dominated by one that belongs to $\mathcal{D}_{m}$. Furthermore, it may be proved that the set $\mathcal{D}_{m}$ is nonempty when the set $\mathcal{D}$ is a finite set, so that at least for this case there exists at least one noninferior decision. However, in practice the set $\mathcal{D}_{m}$ of noninferior decisions often is either difficult to determine or contains too many elements. For instance, in the gambling example given earlier, the reader may verify that every decision is noninferior.

Whenever the partial order ofEq. (G.3) fails to produce a satisfactory ranking among decisions, one must turn to other approaches to formulate the decision problem. The approaches that we will examine assume a notion of a generalized outcome of a decision and introduce a complete order on the set of these generalized outcomes that is consistent with the original order on the set of outcomes $\mathcal{O}$. The complete order on the set of generalized outcomes in turn induces a complete order on the set of decisions.

## The Min-Max Approach

In the min-max (or max-min) approach we take the point of view that the generalized outcome of a decision $d$ is the set of all possible outcomes resulting from d :

$$
f(d, N)=\{\mathrm{O} \mathrm{E} \mathcal{O} \text { Ithere exists } n \mathrm{EN} \text { with } f(d, n)=\mathrm{O}\}
$$

In addition, we adopt a pessimistic attitude and rank the sets $f(d, N)$ on the basis of their worst possible element. In particular, we introduce a complete order on the set of all subsets of $\mathcal{O}$ by means of the relation

$$
\mathcal{O}_{1} \preceq \mathcal{O}_{2} \text { if and only if } \inf _{O \in \mathcal{O}_{1}} G(O) \leq \inf _{O \in \mathcal{O}_{2}} G(0), \text { for all } \mathcal{O}_{1}, \mathcal{O}_{2} \subset \mathcal{O}
$$

where $\mathcal{O}_{1}, \mathcal{O}_{2}$ is any pair of subsets of $\mathcal{O}$, and $G$ is a numerical function consistent with the order on $\mathcal{O}$ in accordance with Eq. (G.2). From Eq. (G.4) we have a complete order on the set of decisions $\mathcal{D}$ by means of

$$
\begin{aligned}
\mathrm{d}_{1} \preceq \mathrm{~d} 2 & \text { if and only if } f(d 1, N) \preceq f(d 2, N) \\
& \text { if and only if } \inf _{n E N} G(f(d 1, n)) \leq \inf _{n E N} G(f(d 2, n)),
\end{aligned}
$$

or in terms of a payoff function J ,

$$
\mathrm{d} 1 \preceq \mathrm{~d} 2 \text { if and only if } \inf _{n E N} J(d 1, n) \leq \inf _{n E N} J(d 2, n) .
$$

Thus, by using the min-max approach, the decision problem is formulated concretely in that it reduces to maximizing over $\mathcal{D}$ the numerical function

$$
F(d)=\inf _{n E N} J(d, n)
$$

Furthermore, it can be easily shown that the elements of $\mathcal{D}$ that maximize $F(d)$ above will not change if $J$ is replaced by $\Phi \cdot J$, where $\Phi: \Re \mapsto \Re$ is any monotonically increasing function. Nonetheless, the min-max approach is pessimistic in nature and will often produce an unduly conservative decision. Characteristically, in the gambling example the optimal decision according to the min-max approach is to refuse the gamble.

We next discuss another approach for formulating decision problems. This approach is quantifying the likelihood of various states of nature through probabilities.

## G. 2 EXPECTED UTILITY THEORY AND RISK

In many decision problems under uncertainty we have additional information about the mechanism by which states of nature occur. In particular,
we are often in a position to know that these states occur in accordance with a given probabilistic mechanism, which may depend on the decision $d$ adopted. To be specific, assume for convenience that the set of states of nature $N$ is either a finite set or a countable sett and that for every decision $d \mathrm{E} V$ we know that states of nature occur according to a given probability law $P(\cdot \mid d)$ defined on $N$. Now each decision $d \mathrm{E} V$ specifies the probability of each outcome via the function $f(d,$.$) and the relation$

$$
P_{d}(O)=P(\{n \mid f(d, n)=O\} \mid d), \quad \text { for all } 0 \mathrm{E} \mathcal{O}
$$

In this relation, $P_{d}(0)$ denotes the probability that the outcome 0 will occur when the decision $d$ is adopted. One may view the probability law $P_{d}$ associated with each $d \mathrm{E} \mathcal{D}$ as a "probabilistic outcome" (or "generalized outcome" to use the term of the preceding section) corresponding to $d$, since $P_{d}$ specifies the probabilistic mechanism by which outcomes occur once $d$ is selected. We shall also use the term lottery $\ddagger$ for a probability law on the set of outcomes. In the example G.l, the decision "bet on heads" has as a outcome the probability law (or lottery) $0,1 / 2$ )
the set of outcomes $0=\{\$ 0, \$ 1, \$ 2\}$. The decision "bet on tails" has the same generalized outcome, while the decision "not bet" has as a generalized outcome the probability law $(0,1,0)$.

The basic idea of the expected utility approach is the following. We already have a complete ranking of the outcomes, i.e., the elements of O. If we had a complete ranking of all lotteries on the set of outcomes (presumably consistent with the original ranking on 0 in the sense that if the outcome $0_{1}$ is preferable to the outcome $\mathrm{O}_{2}$, then the lottery assigning probability 1 to 01 is preferable to the lottery assigning probability 1 to outcome 02 ), then we could in turn obtain a complete ranking of all decisions in $\mathcal{D}$. This is true simply because we could rank any two decisions $d_{1}, d 2 \mathrm{E} \mathcal{D}$ according to the relative order of their corresponding lotteries $P_{d 1}, P_{d 2}$, i.e., by means of the relation

$$
d_{1} \preceq d_{2} \text { if and only if } P_{d_{1}} \preceq P_{d_{2}}
$$

The fundamental premise of the expected utility approach is to assume at the outset that the decision maker has a complete ranking of all lotteries on the set of outcomes, i.e., the decision maker is in a position to express his preference between any two probability laws on the set of

[^5]outcomes. This in turn settles the question of ranking decisions in view of the preceding relation. Furthermore, if there exists a numerical function $G$ by means of which preferences on the set of lotteries can be expressed,
$$
P d \prime \preceq P d 2 \text { if and only if } G\left(P_{d_{1}}\right) \leq G(P d 2)^{\prime}
$$
then decisions can be ranked by means of a numerical function $F$,
$$
d l \preceq d_{2} \text { if and only if } F\left(d_{1}\right) \leq F\left(d_{2}\right)
$$
where $F(d)=G\left(P_{d}\right)$ for all $d \mathrm{E} \mathcal{D}$.
The. aspect of this formulation that is analytically very appealing, however, is that the ordering of decisions can be expressed not only by a function $G$ as above, but also by means of an essentially unique numerical function called the utility function. This function, denoted maps the space of outcomes into the set of real numbers and satisfies
\[

$$
\begin{align*}
& d 1 \preceq d 2 \text { if and only if } P_{d \prime} \preceq P_{d 2} \\
&  \tag{G.5}\\
& \text { if and only if } E\{U(f(d 1, n)) \mid d 1\} \leq E_{\{ }\left\{U(f(d 2, n)) \mid d_{2},\right.
\end{align*}
$$
\]

where the expectations are taken with respect to the corresponding probabihty law $P(\cdot \mid d)$ on $N$. The problem of selecting an optimal decision is thus reduced to the problem of maximizing over $\mathcal{D}$ the expected value of the numerical function $U$.

To clarify the problem formulation based on the approach of this section and to illustrate the advantages resulting from the introduction of a utility function, let us consider an example.

## exampie G. 2

Consider a problem of allocating one unit of capital between two investment opportunities A and B. Opportunity A yields $\$ 1.5$ per dollar invested with certainty, while opportunity B yields $\$ 1$ per dollar invested with probability $1 / 2$ and $\$ 3$ per dollar invested with probability $1 / 2$. the problem is to decide on the fractions $d$ and $(1-d)$ of the capital to be invested in opportunities A and B , respectively, where $0 \leq d \leq \mathrm{I}$.

In terms of the framework of the decision problem of Section the set of decisions $\mathcal{D}$ consists of the interval [ 0,1 ], i.e., the set of values that the fraction $d$ invested in A can take. The set of states of nature $N$ consists of two elements $n \mathrm{I}, \mathrm{n} 2$, where $n \mathrm{I}$ : B yields $\$ 1$ per dollar invested, and $n 2$ : B yields $\$ 3$ per dollar invested. The set of outcomes $\mathcal{O}$ may be taken to be the interval 3], which is the set of possible final fortunes of the investor resulting from all possible decisions and states of nature. The function $f$ that determines the outcome corresponding to any decision $d$ and state of nature $n$ is given by

$$
f\left(d_{, n}\right)= \begin{cases}1.5 d+(1-d) & \text { if } n=\mathrm{nl} \\ l .5 d+3(1-d) & \text { if } n=n 2 .\end{cases}
$$

The preference relation on the set of outcomes is the natural one, i.e., a final fortune 01 is at least as preferable as a final fortune $\mathrm{O}_{2}$ if 01 is numerically greater than or equal to $\mathrm{O}_{2}$ (i.e., $02 \preceq 0_{1}$ if $\mathrm{O}_{2} \leq 0_{1}$ ),

Let us note that since $B$ has a higher expected rate of return, the decision that maximizes expected value of profit is to invest exclusively in opportunity $\mathrm{B} \quad\left(d^{*}=0\right)$. On the other hand the optimal decision based on the max-min approach is to invest exclusively in A $\left(d^{*}=1\right)$ since in this approach one maximizes profit based on the assumption that the most unfavorable state of nature will occur. Mathematically this can be verified by noting that $d^{*}=1$ maximizes over $[0,1]$ the function $F(d)$ given by

$$
F(d) \quad \min \{1.5 d+(1-d), \quad 1.5 d+3(1-d)\}
$$

Note that the approach of maximizing expected profit and the max-min approach lead to very different decisions. Yet it is safe to assume that many decision makers would settle on a decision that differs from both decisions mentioned above and that invests a positive fraction of the capital in both opportunities A and B.

Now in the expected utility approach, the fundamental assumption is that the decision maker has a complete ranking of all lotteries on the set of outcomes. In other words, given any two probability distributions on the interval of final fortunes $[0,3]$, the decision maker can express his preference between the two, in the sense that he can point out the probability distribution in accordance with which he would rather have his final fortune selected. Now the probability distribution on the set of final fortunes corresponding to a decision $d$ is the one that assigns probability $1 / 2$ to $(1.5 d+(1-d))$ and probability $1 / 2$ to $(1.5 d+3(1-d))$. According to the expected utility approach, a decision $d$ is optimal if its corresponding probability distribution is at least as preferable as all other probability distributions of the type described above. It should be clear, however, that a mathematical formulation of the corresponding optimization problem is very cumbersome since it is difficult to visualize or conjecture the form of a numerical function by means of which these probability distributions can be ranked. On the other hand, let us assume that a utility function $U$ satisfying Eq. (G.5) .exists (and it does exist under mild assumptions, as will be indicated shortly). Then an optimal decision is one that solves the problems

$$
\begin{array}{ll}
\operatorname{maximize} & E\{U(J(d, n))\} \\
\text { subject to } & 0 \leq d \leq 1
\end{array}
$$

Substituting the problem data, we have

$$
E\{U(J(d, n))\}=\frac{1}{2}(U(1.5 d+(1-d))+U(1.5 d+3(1-d)))
$$

so the maximization problem is conveniently formulated.
As an example, let us assume that the decision maker's utility function is quadratic of the form

$$
U(O)=\alpha O-O^{2}
$$

where $\alpha$ is some scalar. We require that $6<\alpha$ so that $U(0)$ is increasing in the interval $[0,3]$. This is necessary for the original preference relation on the set of outcomes to be consistent with the one specified by the utility function. Solution of the maximization problem above yields the optimal decision $d^{*}$, where

$$
d^{*}= \begin{cases}0_{(8-\alpha) / 5} & \text { if } 8 \leq \alpha \\ \text { if } 6<\alpha<8\end{cases}
$$

Note that for $6<\alpha<8$, a positive fraction of the capital is invested in opportunity A even though it offers a return that is less than the average return of $B$.

It should be noted, of course, that different decision makers faced with the same decision problem may have different utility functions, so that before the problem can be numerically solved, the form of the utility function must be specified. This can be done experimentally if necessary (see Exercise G.3). However, the importance of the notion of a utility function satisfying Eq. (G.5) lies primarily with the fact that under relatively mild assumptions, it exists and can serve as the starting point of analysis of the decision problem. The reason is that important conclusions about optimal decisions can often be obtained based on either incomplete knowledge of the utility function or fairly general assumptions on its form.

We provide below the theorem of existence of a utility function for the case where the set of outcomes $\mathcal{O}$ is a finite set. For more general cases, see the book by Fishburn [Fis70].

Consider the set $\mathcal{O}$ of outcomes and assume that it is a finite set, $\mathcal{O}=\{\mathrm{Ol}, 02, \ldots, \mathrm{ON}\}$. Let $\mathcal{P}$ be the set of all probability laws $P=$ ( $p l, p 2, ' ", P N$ ) on $\mathcal{O}$, where $P i$ is the probability of outcome $\mathrm{Oi}, \mathrm{i}=$ $1, \ldots, N$. For any $P I, P 2 \in \mathcal{P}, P I=\left(p i, \cdots, p_{N}^{1}\right), P 2=\left(p_{1}^{2}, \ldots, p_{N}^{2}\right)$, and any $\alpha \mathrm{E}[0,1]$, we use the notation

$$
\alpha P_{1}+(1-\alpha) P_{2}=\left(\alpha p_{1}^{1}+(1-\alpha) p_{1}^{2}, \ldots, \alpha p_{N}^{1}+(1-\alpha) p_{N}^{2}\right)
$$

Let us make the following assumptions:
A.I There exists a complete and transitive relation $\preceq$ on $P$. (For any $P I, P 2$ E $P$, we write $P I \sim P 2$ if $P I \preceq P 2$ and $P 2 \preceq P I$, and we write $P I \prec P 2$ if $P I \preceq P 2$ but not $P 2 \preceq P I$.)
A. 2 If $P I \sim P 2$, then for all $\alpha \mathrm{E}[0,1]$ and all $P E P$

$$
\alpha P_{1}+(1-\alpha) P \sim \alpha P_{2}+(1-\alpha) P
$$

A. 3 If $P I \prec P 2$, then for all $\alpha \mathrm{E}(0,1]$ and all $P E P$

$$
\alpha P_{1}+(1-\alpha) P \prec \alpha P_{2}+(1-\alpha) P .
$$

A. 4 If $P I \prec P 2 \prec P 3$, there exists an $\alpha E(0,1)$ such that

$$
\alpha P_{1}+(1-\alpha) P_{3} \sim
$$

Before proving the expected utility theorem, let us briefly discuss the above assumptions. It is convenient for interpretation purposes to view each of the outcomes $0_{1}, 02, \ldots, O N$ as a monetary prize. Consider any probability law $(P I, P 2, \ldots, P N)$ on the set of outcomes. Imagine a pointer that spins in the center of a circle divided into $N$ regions, and assume that it spins in a way that when it stops it is equally likely to be pointing in any direction. The region associated with each prize $\mathrm{Oi}, \mathrm{i}=1, \ldots, N$, occupies a fraction $P i$ of the circumference of the circle. Then we associate with $P$ the game (or lottery) whereby we spin the wheel and win the prize correspondizes to the fithin which the pointer stops. Now given any two probability laws $P /$ and $P 2$ and a scalar $\alpha E \quad 1]$ we can associate with the probability law

$$
\alpha P_{1}+(1-\alpha) P_{2}
$$

the following game. A pointer is spun in the center of a circle divided in two regions, say 1 and 2, occupying respective fractions $\alpha$ and ( $1-\alpha$ ) of its circumference. Depending on whether the pointer stops in region 1 or region 2, the game corresponding to $P 1$ and $P 2$ is played and a prize is won accordingly.

Assumption A. 1 requires that we are able to state our preference between games such as the above, which correspond to any two probability laws $P I$ and $P 2$. Furthermore, our preference relation must be transitive, I.e., if $P 1 \preceq P 2$ and $P 2 \preceq P 3$, then $P I \preceq P 3$. This is the basic assumption, which forms the core of the expected utility approach. Assumptions A. 2 and A. 3 have obvious interpretations and both seem reasonable. Assumption A. 4 is a continuity assumption requiring that if $P I \prec P 2 \prec P 3$, one is indifferent to the game associated with $P_{2}$ and a game whose outcome decides with respective probabilities $\alpha$ and ( $1-\alpha$ ) whether the game associated with or $P_{3}$ will be played. This assumption is inconsistent with a worst-case viewpoint whereby one ranks lotteries according to the worst outcome that can occur with positive probability, and has been the subject of some controversy. For example, consider the extreme situation where there are three outcomes $0_{1}$ death, $\mathrm{O}_{2}=$ receive nothing, and $0_{1}$ receive $\$ 1$. Then it appears reasonable that any probability law that assigns a positive probability to $0{ }_{1}$ (death) cannot be preferable or equivalent to any probability law that assigns a zoo probability to $0_{1}$. Yet Assumption $A A$ requires that for some $\alpha$ with $<\alpha<1$ we are indifferent between the status quo and a game whereby we receive $\$ 1$ with probability $(1-\alpha)$ and die with probability $\alpha$. On the other hand, it is possible to argue that if the probability of death $\alpha$ is extremely close to zero, then this might actually be the case.

The following theorem is the central result of the expected utility theory. It states that a preference relation on the set of all lotteries, which satisfies Assumptions A.I-AA can be characterized numerically by means of an essentially unique function, the utility function. Note that this result
concerns an arbitrary preference relation on lotteries on the set of outcomes and is thus completely decoupled from any decision problem that one may be considering.

Proposition G.1: Under Assumptions A.I-AA, there exists a realvalued function $U: \mathcal{O} \mapsto \Re$, called utility function, such that for all $P I, P_{2} \in P$,

$$
P_{1} \preceq P_{2} \text { if and only if } E_{P_{1}}\{U(O)\} \leq E_{P_{2}}\{U(O)\}
$$

where we denote by $E p\{\cdot\}$ the expected value with respect to a probability law $P$. Furthermore, $U$ is unique up to a positive linear transformation, i.e., if $U^{*}$ is another function with the above property, there exists a positive scalar 31 and a scalar 32 , such that

$$
U^{*}(O)=31 U(0)+32, \quad \text { for all } 0 \mathrm{O}
$$

Proof: We first show the following statement:
S If $P I \prec P_{3}$, and $P 2$ is such that $P I \preceq P 2 \preceq P 3$, then there exists a unique scalar $\alpha \mathrm{E}[0,1]$ such that

$$
\begin{equation*}
\alpha P_{1}+(1-\alpha) P_{3} \sim P_{2} \tag{G.7}
\end{equation*}
$$

Furthermore, if $P_{2}^{\prime}$ is such that $P I \preceq P 2 \preceq P_{2}^{\prime} \preceq P 3$ and $\alpha^{\prime}$ corresponds to $P_{2}^{\prime}$ as in Eq. (G.7), then $\alpha \geq \alpha^{\prime}$.
Indeed if $P I \sim P 2 \prec P_{3}$, then $\alpha=1$ is the unique scalar satisfying Eq. (G.7), since if for some $\alpha E[0,1)$ we had

$$
\alpha P_{1}+(1-\alpha) P_{3} \sim P_{2} \sim \alpha P_{1}+(1-\alpha) P_{2}
$$

then Assuroption A. 3 would be contradicted. Similarly, if $P 1 \prec P 2 \sim P 3$, then $\alpha=$ is the unique scalar satisfying Eq. (G.7). Assume now that $P 1 \prec P 2 \prec P 3$. Then by Assumption there exists an $\alpha_{1} \mathrm{E}(0,1)$ satisfying Eq. (G.7). Assume that $\alpha_{1}$ is not unique and there exists another scalar $\alpha_{2} \mathrm{E}(0,1)$ such that Eq. (G.7) is satisfied, i.e.,

$$
\begin{equation*}
\alpha_{1} P_{1}+\left(1-\alpha_{1}\right) P_{3} \sim P_{2} \sim \alpha_{2} P_{1}+\left(1-\alpha_{2}\right) P_{3} \tag{G.8}
\end{equation*}
$$

Let us assume that ${ }^{\mathrm{O}}<\alpha_{1}<\alpha_{2}<1$. Then we have

$$
\begin{equation*}
P_{3}=\frac{\alpha_{2}-\alpha_{1}}{1-\alpha_{1}} P_{3}+\frac{1-\alpha_{2}}{1-\alpha_{1}} \tag{G.g}
\end{equation*}
$$

$$
\begin{equation*}
\alpha_{2} P_{1}+\left(1-\alpha_{2}\right) P_{3}=\alpha_{1} P_{1}+(1-\alpha)\left\{\frac{\alpha_{2}}{11}-\alpha_{1}+\frac{1-\alpha_{2}}{1-\alpha_{1}} P_{3}\right\} \tag{G.10}
\end{equation*}
$$

Since $P I \prec P 3$, we have by Assumption A. 3 and Eq. (G.g)

$$
\frac{\alpha_{2}-\alpha_{1}}{1-\alpha_{1}} P_{1}+\frac{1-\alpha_{2}}{1-\alpha_{1}} \quad \prec \frac{\alpha_{2}-\alpha_{1}}{1-\alpha_{1}} P_{1}+\frac{1-\alpha_{2}}{1-\alpha_{1}} P_{3}=P_{3}
$$

Again, using Assumption A. 3 and Eq. (G.10), we have

$$
\alpha_{2} P_{1}+\left(1-\alpha_{2}\right) P_{3} \prec \alpha_{1} P_{1}+\left(1-\alpha_{1}\right) P_{3}
$$

However, this contradicts Eq. (G.8) and hence the uniqueness of the scalar $\alpha$ in Eq. (G.7) is proved.

To show that $P I \preceq P 2 \preceq P_{2}^{\prime} \preceq P 3$ implies $\alpha \geq \alpha^{\prime}$, assume the contrary, i.e., $\alpha<\alpha^{\prime}$. Then we have, using Assumption A.3,

$$
\begin{aligned}
P_{2}^{\prime} & \sim \alpha^{\prime} P I+\left(1-\alpha^{\prime}\right) P_{3} \\
& =\left(1-\alpha+\alpha^{\prime}\right)\left\{\begin{array}{l}
\alpha \\
1-\alpha+\alpha^{\prime}
\end{array}{ }^{1-\alpha l+\alpha^{\prime}} \begin{array}{c}
1-\alpha+\alpha^{\prime}
\end{array} P_{3}\right\}+\left(\alpha^{\prime}-\alpha\right) P_{1} \\
& \prec\left(1-\alpha+\alpha^{\prime}\right)\left\{\begin{array}{l}
1-\alpha^{\prime} \\
1-\alpha+\alpha^{\prime} \\
1-\alpha+\alpha^{\prime}
\end{array}\right\}+\left(\alpha^{\prime}-\alpha\right) P_{1} \\
& =\alpha P_{1}+\left(\begin{array}{ll}
1-\alpha) P_{3}
\end{array}\right.
\end{aligned}
$$

Hence $P_{2}^{\prime} \prec P 2$, which contradicts the assumption $P 2 \preceq P_{2}^{\prime}$. It follows that $\alpha \geq \alpha^{\prime}$ and statement S is proved.

Now consider the probability laws

$$
\bar{P}_{1}=(1,0, \ldots, 0), \quad \bar{P}_{2}=(0,1, \ldots, 0), \quad \ldots, \quad \bar{P}_{N}=(0,0, \ldots, 1)
$$

Assume without loss of generality that $\bar{P}_{1} \preceq \bar{P}_{2} \preceq \ldots \preceq \bar{P}_{N}$ and assume further that $P_{1} \prec \bar{P}_{\mathrm{N}}$ (if $\bar{P}_{1} \sim \bar{P}_{2}, \sim \ldots \sim \overline{\bar{P}}_{\mathrm{N}}$, the proof of the proposition is trivial). Let be any scalars with $\mathrm{Al}<\quad$ and define

$$
U\left(O_{1}\right)=\quad U\left(O_{N}\right)=A_{N}
$$

Let $\alpha_{i}, i=1, \ldots, n$, be the unique scalar $\alpha_{i} \mathrm{E}[0,1]$ such that

$$
\begin{equation*}
\alpha_{i} \bar{P}_{1}+\left(1-\alpha_{i}\right) \bar{P}_{N} \sim \quad i=1, \ldots \tag{G.11}
\end{equation*}
$$

and define

$$
\begin{equation*}
U\left(O_{i}\right)=A_{i}=\alpha_{i} A_{1}+(1-\quad i=1, \ldots, N \tag{G.12}
\end{equation*}
$$

We shall prove that the function $U: \mathcal{O} \mapsto \Re$ defined above has the desired property (G.6). Indeed for any probability law $P=(P I, \ldots, P N)$, it is easily
seen that $\bar{P}_{1} \prec P \prec \bar{P}_{\mathbf{N}}$, and thus we can define $\alpha(P)$ to be the unique scalar in $[0,1]$ such that

$$
\begin{equation*}
\alpha(P) \bar{P}_{1}+(1-\alpha(P)) \bar{P}_{N} \sim P \tag{G.13}
\end{equation*}
$$

From statement S we obtain for all $P, P^{\prime}$

$$
\begin{equation*}
P \preceq P^{\prime} \text { if and only if } \alpha(P) \geq \alpha\left(P^{\prime}\right) \tag{G.14}
\end{equation*}
$$

Now from Eq. (G.11), we have

$$
\begin{align*}
P= & \sum_{i=1}^{N} p_{i} \bar{P}_{i} \\
\sim & \sum_{i=1}^{N} p_{i}\left(\alpha_{i} \bar{P}_{i}+\left(1-\alpha_{i}\right) \bar{P}_{N}\right)  \tag{G.15}\\
& \sum_{i=1}^{N} p_{i} \alpha_{i} \bar{P}_{1}+\left(1-\sum_{i=1}^{N} p_{i} \alpha_{i}\right) \bar{P}_{N} .
\end{align*}
$$

Comparing Eqs. (G.13) and (G.15), we obtain

$$
\alpha(P)=\sum_{i=1}^{\mathrm{N}} p_{i} \alpha_{i}
$$

and from Eq. (G.14),

$$
\begin{equation*}
P I \preceq P 2 \text { if and only if } \sum_{\mathrm{i}=1}^{\mathrm{N}} p_{i}^{1} \alpha_{i} \geq \sum_{\mathrm{i}=1}^{\mathrm{N}} p_{i}^{2} \alpha_{i} . \tag{G.16}
\end{equation*}
$$

From Eq. (G.12), we have $\alpha_{i}=(A N-A i) /\left(A_{N}-A I\right)$, and substituting in Eq. (G.16), we obtain

$$
P I \preceq P 2 \text { if and only if } \sum_{\mathrm{i}=1}^{\mathrm{N}} p_{i}^{1} A i \leq \sum_{\mathrm{i}=1}^{\mathrm{N}} p_{i}^{2} A i,
$$

which is equivalent to the desired relation (G.6).
There remains to show that the function $U$ defined by Eq. (G.12) is unique up to a positive linear transformation. Indeed if $U^{*}$ were another utility function satisfying Eq. (G.6), then by denoting $1, \ldots, \mathrm{~N}$, we would have from Eqs. (G.11) and (G.6)

$$
U^{*}\left(O_{i}\right)=\alpha_{i} U^{*}\left(O_{1}\right)+\left(1-\alpha_{i}\right) U^{*}\left(O_{N}\right)
$$

It follows that

$$
\alpha_{i}=\frac{A_{N}-A_{i}}{A_{N}-A_{1}} \quad \frac{A_{N}^{*}}{A_{N}^{*}}-\frac{A_{i}^{*}}{\boldsymbol{A} i^{\prime}}
$$

from which

$$
A_{i}^{*}=A_{N}^{A_{N}^{*}-A^{*}} A_{i} A_{i}+A_{\mathrm{N}}^{*} \quad \frac{A_{N}\left(A_{N}^{*}=\frac{\left.A_{1}^{*}\right)}{A N-\boldsymbol{A} l}\right.}{\text { In }}
$$

This proves the theorem. Q.E.D.
Returning now to the decision problem, once we assume the existence of a preference relation on the set of lotteries that is characterized by a utility function, we can rank decisions as follows: Given the probability law $P(\cdot \mid d)$ on the set of states of nature $N$, every decision $d \mathrm{E} \mathcal{D}$ induces a probability law (or lottery) $P d$ on the set of outcomes $\mathcal{O}$. Under the assumptions of the expected utility theorem, there exists a utility function $U: \mathcal{O} \mapsto \Re$ such that for any $d_{1}, d_{2} \mathrm{E} \mathcal{D}$

$$
P_{d_{1}} \preceq P_{d_{2}} \text { if and only if } E_{P_{d_{1}}}\{U(O)\} \leq E_{P_{d_{2}}}\{U(O)\}
$$

We have, however,

$$
E_{P_{d}}\{U(O)\}=E\{U(f(d, n)) \mid d\}, \quad \text { for all } d E \mathcal{D}
$$

where the expectation on the left is taken with respect to $P d$ and the expectation on the right is taken with respect to the probability law $P(\cdot \mid d)$ on $N$. Hence

$$
P_{d_{1}} \preceq P_{d_{2}} \text { if and only if } E\left\{U\left(f\left(d_{1}, n\right)\right) \mid d_{1}\right\} \leq E\left\{U\left(f\left(d_{2}, n\right)\right) \mid d_{2}\right\} .
$$

By ranking decisions $d E \mathcal{D}$ in accordance with the ranking of the corresponding $P_{d}$, i.e.,

$$
\begin{aligned}
\mathrm{d}_{1} \preceq d 2 & \text { if and only if } P d l \preceq P d 2 \\
& \text { if and only if } E\left\{U(f(d I, n)) \mid \mathrm{d}_{1\}} \leq E\{U(f(d 2, n)) \mid d 2\}^{\prime}\right.
\end{aligned}
$$

we obtain a complete order on the set $\mathcal{D}$ induced by the utility function $U$. The optimal decision is found by maximization of the numerical function $F: \mathcal{D} \mapsto \Re$, where

$$
F(d)=E\{U(f(d, n)) \mid d\}
$$

and the decision problem is formulated in a way that is amenable to mathematical analysis.

The N otion of Risk

Consider a decision maker possessing a utility function $U$ defined over an interval $X$ of real numbers. We say that the decision maker is risk averse if

$$
\begin{equation*}
E_{P}\{U(x)\} \leq U\left(E_{P}\{x\}\right) \tag{G.17}
\end{equation*}
$$

for every probability distribution $P$ on $X$ for which the expected value above is finite. In other words, a decision maker is risk averse if he always prefers the expected value of the lottery over the lottery itself. Such behavior characterizes most decision makers. One may show that risk aversion is equivalent to concavity of the utility function (see Appendix A for the definition and properties of concave and convex functions.) On the other hand, we say that the decision maker is risk preferring if the opposite inequality holds in Eq. (G.17), which is the case of a convex utility function. A gambler playing an unbiased roulette and receiving no reward or pleasure from gambling per se is a typical example of a risk preferring decision maker. Finally, a decision maker having a linear utility function is said to be risk neutral.

The notion of risk is important since it captures a basic attribute of the attitudes of the decision maker and often characterizes significant aspects of his behavior. An important and widely accepted measure of risk has been proposed by Pratt [Pra64]. He introduced the function

$$
\begin{equation*}
r(x)=-\frac{U^{\prime \prime}(x)}{U^{\prime}(x)} \tag{G.18}
\end{equation*}
$$

where $U^{\prime}$ and $U^{\prime \prime}$ denote the first and second derivative of $U$, and it is assumed that $U^{\prime}(x) \neq 0$ for all $x$. This function, called the index of absolute risk aversion, measures locally (at the point $x$ ) the risk aversion of the decision maker. It can be interpreted as follows.

Let $x$ be a gamble over the set of real numbers (Le., a random variable) with given distribution and expected value $\bar{x}=E\{x\}$. Let us denote by $y$ the amount of insurance the decision maker is willing to pay in order to avoid the gamble $x$, and instead receive the expected value $\bar{x}$ of the gamble. In other words, $y$ is such that

$$
\begin{equation*}
U(\bar{x}-y)=E\{U(x)\} \tag{G.19}
\end{equation*}
$$

Intuitively, $y$ provides a natural measure of risk aversion. Using a Taylor series expansion around $\bar{x}$, we have

$$
\begin{equation*}
U(\bar{x}-y)=U(\bar{x})-y U^{\prime}(\bar{x})+o(y) \tag{G.20}
\end{equation*}
$$

where by $O(y)$ we denote a quantity that is negligible compared with the scalar $\alpha$ provided $\alpha$ is close to zero, i.e., $\lim _{\alpha \rightarrow 0}(o(\alpha) / a)=0$. Also we have

$$
\begin{align*}
E\{U(x)\} & =E\left\{U(\bar{x})+(x-\bar{x}) U^{\prime}(\bar{x})+\frac{1}{2}(x-\bar{x})^{2} U^{\prime \prime}(\bar{x})+o\left((x-\bar{x})^{2}\right)\right\} \\
& =U(\bar{x})+\frac{1}{2} \sigma^{2} U^{\prime \prime}(\bar{x})+E\left\{o\left((x-\bar{x})^{2}\right)\right\} \tag{G.21}
\end{align*}
$$

where $\sigma^{2}$ is the variance of $x$. From Eqs. (G.19)-(G.21), we have

$$
y U^{\prime}(\bar{x})=-\frac{1}{2} \sigma^{2} U^{\prime \prime}(\bar{x})+O(y)+E\left\{o\left((x-\bar{x})^{2}\right)\right\}
$$

From this equation and Eq. (G.18) it follows that the amount of insurance or risk premium $y$ that the decision maker is willing to pay is proportional (up to first order) to the index of absolute risk aversion $r(\bar{x})$ at the mean $\bar{x}$ of the gamble, thus justifying the use of $r$ as a measure of local risk aversion. Notice that in the investment Example G.2, we have $r(y)=2 /(\alpha-2 y)$, so $r(y)$ tends to decrease as $\alpha$ increases. This fact is reflected in the optimal investment, where an increasing fraction of the capital is invested in the risky asset as $\alpha$ is increased.

The index $r(x)$ often plays an important role in the analysis of behavior of decision makers. It is generally accepted that for most decision makers, $r(x)$ is a decreasing or at least nonincreasing function of $x$, i.e., the decision maker more readily accepts risk as his wealth is increased. On the other hand, for the quadratic utility function $U(x)=-\frac{1}{2} x^{2}+b x+c$, the index $r(x)$ is equal to $(b-x)-1$ and is an increasing function of $x$ (for $x<b)$. For this reason the quadratic utility function is often considered inappropriate or at least accepted with reservation in economics applications, despite the analytical simplifications resulting from its use.

## Example G. 3

An individual with given initial wealth $\alpha$ wishes to invest part of it in a risky asset offering a rate of return $e$, and the rest in a secure asset offering rate of return $s>0$. We assume that $s$ is known with certainty while $e$ is a random variable with known probability distribution $P$. If $x$ is the amount invested in the risky asset, then the final wealth of the decision maker is given by

$$
y=s(\alpha-x)+e x=s \alpha+(e-s) x
$$

The decision to be made by the individual is to choose $x$ so as to maximize

$$
J(x)=E\{U(y)\}=E\{U(s \alpha+(e-s) x)\}
$$

subject to the constraint $x \geq 0$. We assume that $U$ is a concave, monotonically increasing, twice continuously differentiable function with negative second derivative, and with index of absolute risk aversion

$$
r(y)=-\underline{U I I(y)} .
$$

We also assume that the probability distribution of e is such that all expected values appearing below are finite, and furthermore we assume that the utility function $U$ is such that the maximization problem has a solution (the necessary and sufficient conditions for this have an interesting economic interpretation, which is discussed in Bertsekas [Ber74]).

Now given $\alpha$, the amount $x^{*}$ to be invested in the risky asset is determined from the necessary conditions

$$
\begin{equation*}
\frac{d J\left(x^{*}\right)}{d x}=E\left\{(\mathrm{e}-s) U^{\prime}\left(s \alpha+(\mathrm{e}-s) x^{*}\right)\right\}=0, \quad \text { if } x^{*}>0 \tag{G.22}
\end{equation*}
$$

$$
\frac{d J\left(x^{*}\right)}{d x}<, \quad \text { if } x^{*}=0
$$

Now since $U^{\prime}$ is everyohere positive it follows that if $E\{(e-s)\}>0$, then we cannot have $x^{*}=$ since

$$
\frac{d J(0)}{d x}=E\{(e-s)\} U^{I}(s \alpha)>0
$$

Hence $E\{(\mathrm{e}-s)\}>{ }^{\mathrm{O}}$ implies $x^{*}>0$, or in words, a positive amount will be invested in the risky asset if its expected rate of return is greater than the rate of return of the secure asset.

Assume now that $E\{(\mathrm{e}-s)\}>$ and denote by $x^{*}(\alpha)$ the amount invested in the risky asset when the initial wealth is $\alpha$. We would like to investigate the effects of changes in initial wealth $\alpha$ on the amount $x^{*}(\alpha)$ invested. By differentiating Eq. (G.22) with respect to $\alpha$ we obtain

$$
E\left\{(e-s) U^{\prime \prime}\left(s \alpha+(e-s) x^{*}(\alpha)\right)\left(s \alpha+(e \quad s)\left(d x^{*}(\alpha) / d \alpha\right)\right)\right\}=0
$$

from which

$$
\frac{d x^{*}(\alpha)}{d \alpha}=-\frac{E\left\{(e-s) U I I\left(s a+(e-s) x^{*}(a)\right)\right\}}{E\left\{(e-s)^{2} U^{\prime \prime}\left(s \alpha+(e-s) x^{*}(a)\right)\right\}}
$$

Since the denominator is always negative and the constant $s$ is positive, the sign of $d x^{*}(\alpha) j d \alpha$ is the same as the sign of

$$
E\left\{(e-s) U^{\prime \prime}\left(s \alpha+(e-s) x^{*}(a)\right)\right\}
$$

which using the definition of the index of absolute risk aversion $r(y)$ is equal to

$$
f(\alpha)=E\left\{(e-s) U^{I}\left(s \alpha+(e-s) x^{*}(\alpha)\right) r\left(s \alpha+(e-s) x^{*}(\alpha)\right)\right\}
$$

Now assume that $r(y)$ is monotonically decreasing, i.e.,

$$
r\left(y_{1}\right)>r\left(y_{2}\right) \quad \text { if } \quad \mathrm{Yl}<Y 2
$$

Then we have

$$
(e-s) r\left(s \alpha+(e-s) x^{*}(\alpha)\right) \leq(e-r) r(s \alpha)
$$

with strict inequality if $\mathrm{e} \neq s$, and from the preceding relations, we obtain

$$
\begin{aligned}
f(\alpha) & >-r(s \alpha) E\left\{(\mathrm{e}-s) U^{\prime}\left(s \alpha+(\mathrm{e}-s) x^{*}(\alpha)\right)\right\} \\
& =-r(s \alpha) d J\left(x^{*}(\alpha)\right) \\
& =0 .
\end{aligned}
$$

Thus we have $f(\alpha)>0$ and hence $d x^{*}(\alpha) / d \alpha>0$ if $r(y)$ is monotonically decreasing. Similarly, we obtain $d x^{*}(\alpha) j d \alpha<0$ if $r(y)$ is monotonically increasing. In words, the individual, given more wealth, will invest more (less) in the risky asset if his utility function has decreasing (increasing) index of absolute risk aversion. Aside from its intrinsic value, this result illustrates the important role of the index of risk aversion in shaping significant aspects of a decision maker's behavior.

## STOCHASTIC OPTIMAL CONTROL PROBLEMS

The class of decision problems considered so far in this appendix is very broad. In this book we are interested in a subclass of decision problems that involves a dynamic system. Such systems have an input-output description and furthermore in such systems, inputs are selected sequentially after observing past outputs. This allows the possibility of feedback. Let us first give an abstract description of these problems.

Let us consider a system characterized by three sets $U, W$, and $Y$, and a function $S: U \times W \mapsto Y$. We call $U$ the input set, $W$ the uncertainty set, $Y$ the output set, and $S$ the system function. Thus an input $u \mathrm{E} U$ and an uncertain quantity $w E W$ produce an output $y=S(u, w)$ through the system function $S$ (see Fig. G.1). Implicit here is the assumption that the choice of the input $u$ is somehow controlled by a decision maker or device to be designed, while $w$ is chosen by nature according to some mechanism, probabilistic or not.


Figure G. 1 Structure of an uncertain system: $u$ is the input, $w$ is the uncertain state of nature, $y$ is the output, and $S$ is the system function.

In many problems that evolve in time, the input is a time function or sequence, and there may be a possibility of observing the output $y$ as it evolves in time. Naturally, this output may provide some information about the uncertain quantity $w$, which may be fruitfully taken into account in choosing the input $u$ by means of a feedback mechanism.

Let us say that a function $\pi: Y \mapsto U$ is a feedback controller (otherwise called policy or decision function) for the system if for each $w$ E $W$ the equation

$$
u=\pi(S(u, w))
$$

has a unique solution (dependent on $w$ ) for $u$. Thus for any fixed $w$, a feedback controller $\pi$ generates a unique input $u$ and hence a unique output $y$ (see Fig. G.2). In any practical situation, the class of admissible feedback controllers is further restricted by causality (present inputs should not depend on future outputs), and possibly other constraints.


Figure G. 2 Structure of a feedback controller $\boldsymbol{\pi}$. We require that for each $w$ E $W$ the equation

$$
u=\pi(S(u, w))
$$

has a unique solution in $u$.

Given the system ( $U, \quad Y, S$ ) and a set of admissible controllers it is possible to formulate a decision problem in accordance with the theory of the previous section. We take II as the decision set and $W$ as the set of states of nature. We take as the set of outcomes the Cartesian product of $W$, and $Y$, i.e.,

$$
O=(U \times W \times Y)
$$

Now a feedback controller $\pi$ E II and a state of nature $w E W$ generate a unique outcome ( $u, w, y$ ), where $u$ is the unique solution of the equation $u=\pi(S(u, w))$ and $y=S(u, w)$. Thus we may write $(u, w, y)=f(\pi, w)$, where $f$ is some function determined by the system function $S$.

If G is a numerical function ordering our preferences on $0, \mathrm{~J}$ is the corresponding payoff function for the decision problem above, and a maxmin viewpoint is adopted, then the problem becomes one of finding $\pi$ E II that maximizes

$$
\min _{w \in W} J(\pi, w)=\min _{w \in W} G(u, w, y)
$$

where $u$ and yare expressed in terms of $\pi$ and $w$ by means of $u=\pi(S(u, w))$ and $y=S(u, w)$ (min here denotes least upper bound over the corresponding set $W$ ).

If $w$ is selected in accordance with a known probabilistic mechanism, i.e., a given probability law that may depend on $\pi$, and the function Sand the elements of II satisfy suitable (measurability) assumptions, then it is possible to use a utility function $U$ to formulate the decision problem as one of finding $\pi$ E II that maximizes

$$
F(\pi)=E\{U(u, n, w)\}
$$

$\sim$ wand $Y$ are expressed in terms of $\pi$ and $w$ by means of $u=\pi(S(u, w))$ andy $=S(u, w)$.

While in the formulations just given, we have reduced the problem to one of decision under certainty [the problem of maximizing over II the numerical function $F(\pi)$ ], this is not an easy problem. The reason is that due to the feedback possibility, the set II is a set offunctions (of the system output). This renders inapplicable deterministic optimization techniques, such as those based on linear and nonlinear programming, or Pontryagin's minimum principle. Dynamic programming offers some possibility of analysis by decomposing the problem of minimizing $F(\pi)$ into a sequence of much simpler optimization problems that are solved backwards in time, as discussed in Chapter 1.

Finally, let us indicate how to convert the basic problem of Section 1.2 into the general form given in this section. Referring to the discrete time dynamic system

$$
\begin{equation*}
x_{k=1}=f_{k}\left(x_{k}, u_{k}, w_{k}\right), \quad k=0,1, \ldots, N-1, \tag{G.23}
\end{equation*}
$$

introduced in Section 1.2, the system input is the control sequence $u=$ $\{U O, U I, \ldots, U N--I\}$, the uncertainty is $w=\left\{W O, W l, \ldots, w_{N-1}\right\}$ (perhaps together with the initial state $X_{0}$, if $X_{0}$ is uncertain), the output is the state sequence $y=\{X O, \mathrm{x} 1, \ldots, X N\}$, and the system function is determined in the obvious manner from the system equation (G.23). The class II of admissible feedback controllers is the set of sequences of functions $\pi=$ $\left\{\mu_{0}, \mu_{1}, \ldots, \mu_{N-1}\right\}$, where $\mu_{k}$ is a function that depends on the output $y$ exclusively through the state $x k$. Furthermore, $\mu_{k}$ must satisfy constraints such as $\mu_{k}\left(x_{k}\right) \mathrm{E} U k(X k)$ for all $x_{k}$ and $k$.

## EXERCISES

## G. 1

Show that there exists a function $\mathrm{G}: 0 \mapsto \Re$ satisfying relation (G.2) provided the set 0 is countable. Show also that if the set of decisions is finite, there exists at least one noninferior decision.

## G. 2

Let $0=[-1,1]$. Define an order on $\mathcal{O}$ by means of

$$
O_{1}<O_{2} \text { if and only if }\left|O_{1}\right|<\left|O_{2}\right| \text { or } O_{1}<O_{2}=\left|O_{1}\right|
$$

Show that there exists no real-valued function G on $\mathcal{O}$ such that

$$
0_{1} \prec \mathrm{O}_{2} \text { if and only if } G(O I)<G\left(02, \quad \text { for all } 0_{1}, \mathrm{O}_{2} \mathrm{EO} .\right.
$$

Hint: Assume the contrary and associate with every $0 \mathrm{E}(0,1)$ a rational number $r(0)$ such that

$$
\mathrm{G}(-\mathrm{O})<r(O)<G(O)
$$

Show that if $0_{1} \neq \mathrm{O}_{2}$, then $r(01) \neq r(02)^{\prime}$

## G. 3 Expermental Measurement of UdUty)

Consider an individual faced with a decision problem with a finite collection of outcomes $0_{1}, \mathrm{O}_{2}, \ldots, O N$. Assume that the individual has a preference relation over the set of lotteries on the set of outcomes satisfying Assumptions A.I-AA of the expected utility theorem, and hence a utility function over the set of outcomes exists. Suppose also that $0_{1} \preceq 02 \preceq \ldots \preceq O N$ and furthermore that $0_{1} \prec O N$.
(a) Show that the following method will dermine a utility function. Define $U(01)=0, U(O N)=1$. Let Pi with $\leq P i \leq 1$ be the probability for which one is indifferent between the lottery $\left\{\left(\mathrm{I}-\mathrm{P}^{\prime} i\right), 0, \ldots, 0, \mathrm{Pi}\right\}$ and Oi occurring with certainty. Then let $U(O i)=P i$. Try the procedure on yourself for $\mathrm{Oi}=100 \mathrm{i}$ with $\mathrm{i}=0,1, \ldots, 10$.
(b) Show that the following procedure will also yield a utility function. Determine $U(O N-1)$ as in (a), but set

$$
U\left(O_{N-2}\right)=\tilde{p}_{N-2} U\left(O_{N-1}\right)
$$

where $\tilde{p}_{N-2}$ is the probability for which one is indifferent between the lottery $\left\{\left(\mathrm{I}-\tilde{p}_{N-2}\right), 0, \ldots, 0, P N-2,0\right\}$ and $O N-2$ occurring with certainty.

Similarly, set $U(O i)=\tilde{p}_{i} U\left(O_{i+1}\right)$, where $\tilde{p}_{i}$ is the appropriate probability. Again try this procedure on yourself for $\mathrm{Oi}=100 \mathrm{i}$ with $\mathrm{i}=0,1, \ldots, 10$, and compare the results with the ones obtained in part (a).
(c) Devise a third procedure whereby the utilities $U(O l), U(O z)$ are specified initially and $U(O i), \mathrm{i}=3, \ldots, N$, is obtained from $U(O i-2), U(O i-l)$ through a comparison of the type considered above. Again try this procedure on yourself for $\mathrm{Oi}=100 \mathrm{i}$ with $\mathrm{i}=0,1, \ldots, 10$.

## G. 4

Suppose that two persons, A and B, want to make a bet. Person A will pay $\$ 1$ to person B if a certain event occurs and person B will pay $x$ dollars to person A if the event does not $\boldsymbol{O}^{\text {ccur }}$. Person A believes that the probability of the event occurring $P A$ with $<P A<1$, while person $B$ believes that this probability is $P B$ with $<P B<1$. Suppose that the utility functions $U_{A}$ and $U B$ of persons $A$ and $B$ are strictly increasing functions of monetary gain. Let $\alpha$, (3 be such that

$$
U_{A}(\alpha)=\frac{U_{A}(0)-p_{A} U_{A}(-1)}{1-p_{A}}, \quad U B\left(-(3)=\frac{U B(O)}{\perp}=\frac{P B U B(l) .}{p_{B}}\right.
$$

Show that if $\alpha<\beta$. .auly value of $x$ between $\alpha$ and (3 is a mutually satisfactory bet.

## References

[ABC65] Atkinson, R. C., Bower, G. H., and Crothers, E. J., 1965. An Introduction to Mathematical Learning Theory, Wiley, N. Y.
[ABF93] Arapostathis, A., Borkar, V., Fernandez-Gaucherand, E., Ghosh, M., and Marcus, S., 1993. "Discrete-Time Controlled Markov Processes with Average Cost Criterion: A Survey," SIAM J. on Control and Optimization, Vol. 31, pp. 282-344.
[ABG49] Arrow, K. J., Blackwell, D., and Girshick, M. A., 1949. "Bayes and Minimax Solutions of Sequential Design Problems," Econometrica, Vol. 17, pp. 213-244
[AGK77] Athans, M" Ku, R., and Gershwin, S. B., 1977. "The Uncertainty Threshold Principle," IEEE Trans. on Automatic Control, Vol. AC-22, pp. 491-495.
[AHM51] Arrow, K. J., Harris, T., and Marschack, J., 1951. "Optimal Inventory Policy," Econometrica, Vol. 19, pp. 250-272.
[AKS58] Arrow, K. J., Karlin, S., and Scarf, H., 1958. Studies in the Mathematical Theory of Inventory and Production, Stanford Univ. Press, Stanford, CA.
[Abr90] Abramson, B., 1990. "Expected-Outcome: A General Model of Static Evaluation," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 12, pp. 182-193.
[AdG86] Adams, M., and Guillemin, V., 1986. Measure Theory and Probability, Wadsworth and Brooks, Monterey, CA.
[AsG75] Ash, R. B., and Gardner, M. F., 1975. Topics in Stochastic Processes, Academic Press, N. Y.
[AnM79] Anderson, B. D.O., and Moore, J. B., 1979. Optimal Filtering, Prentice-Hall, Englewood Cliffs, N. J.
[AoL69] Aoki, M., and Li, M. T., 1969. "Optimal Discrete-Time Control Systems with Cost for Observation," IEEE Trans. Automatic Control, Vol. AC-14, pp. 165-175.
[AsW73] Astrom, K. J., and Wittenmark, B., 1973. "On Self-Tuning Regulators," Automatica, Vol. 9, pp. 185-199.
[AsW84] Astrom, K. J., and vVittenmark, B., 1984. Computer Controlled Systems, Prentice-Hall, Englewood Cliffs, N. J.
[AsW94] Astrom, K. J., and Wittenmark, B., 1994. Adaptive Control, (2nd Ed.), Prentice-Hall, Englewood Cliffs, N. J.
[Ash70] Ash, R. B., 1970. Basic Probability Theory, Wiley, N. Y.
[Ash72] Ash, R. B., 1972. Real Analysis and Probability, Academic Press, N. Y.
[Ast83] Astrom, K. J., 1983. "Theory and Applications of Adaptive Control A Survey," Automatica, Vol. 19, pp. 471-486.
[AtF66] Athans, M., and Falb, P., 1966. Optimal Control, McGraw-Hill, N. Y.
[BGM95] Bertsekas, D. P., Guerriero, F., and Musmanno, R., 1995. "Parallel Shortest Path Methods for Globally Optimal Trajectories," High Performance Computing: Technology, Methods, and Applications, (J. Dongarra et al., Eds.), Elsevier.
[BGM96] Bertsekas, D. P., Guerriero, F., and Musmanno, R., 1996. "Parallel Label Corrrecting Methods for Shortest Paths," J. Optimization Theory Appl., Vol. 88, 1996, pp. 297-320.
[BMS99] Boltyanski, V., Martini, H., and Soltan, V., 1999. Geometric Methods and Optmizurion Problems, Kluwer, Boston.
[BN003] Bertsekas, D. r., .;+h A. Nedic, A., and A. E. Ozdaglar, 2003. Convex Analysis and Optimization, Athena Scientuic Belmont, MA.
[BTW97] Bertsekas, D. P., Tsitsiklis, J. N., ana. C., 1997. "Rollout Algorithms for Combinatorial Optimization," Heuristics, Vol. 3, pp. 245-2 22.
[BaB95] Basar, T., and Bernhard, P., 1995. H-oo Optimal Control and Relaud Minimax Design Problems: A Dynamic Game Approach, Birkhuser, Boston, MA.
[Bar81] Bar-Shalom, Y, 1981. "Stochastic Dynamic Programming: Caution and Probing," IEEE Trans. on Automatic Control, Vol. AC-26, pp. 1184-1195.
[Bas91] Basar, T., 1991. "Optimum Performance Levels for Minimax Filters, Predictors, and Smoothers," Systems and Control Letters, Vol. 16, pp. 309-317.
[BasOO] Basar, T., 2000. "Risk-Averse Designs: From Exponential Cost to Stochastic Games," In T. E. Djaferis and 1. C. Schick, (Eds.), System Theory: Modeling, Analysis and Control, Kluwer, Boston, pp. 131-144.
[BeC99] Bertsekas, D. P., and Castanon, D. A., 1999. "Rollout Algorithms for Stochastic Scheduling Problems," Heuristics, Vol. 5, pp. 89-108.
[BeC04] Bertsekas, D. P., and Castanon, D. A., 2004. Unpublished Collaboration.
[BeD62] Bellman, R, and Dreyfus, S., 1962. Applied Dynamic Programming, Princeton Univ. Press, Princeton, N. J.
[BeD02] Bertsimas, D., and Demir, R, 2002. "An Approximate Dynamic Programming Approach to Multi-Dimensional Knapsack Problems," Management Science, Vol. 4, pp. 550-565.
[BeG92] Bertsekas, D. P., and Gallager, R G., 1992. Data Networks (2nd Edition), Prentice-Hall, Englewood Cliffs, N. J.
[BeN98] Ben-Tal, A., and Nemirovski, A., 1998. "Robust Convex Optimization," Math. of Operations Research, Vol. 23, pp. 769-805.
[BeNOl] Ben-Tal, A., and Nemirovski, A., 2001. Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications, SIAM, Phila., PA
[BeP03] Bertsimas, D., and Popescu, 1., 2003. "Revenue Management in a Dynamic Network Environment,"Transportation Science, Vol. 37, pp. 257-277.
[BeR71a] Bertsekas, D. P., and Rhodes, 1. B., 1971. "Recursive State Estimation for a Set-Membership Description of the Uncertainty," IEEE Trans. Automatic Control, Vol. AC-16, pp. 117-128.
[BeR71b] Bertsekas, D. P., and Rhodes, 1. B., 1971. "On the Minimax Reachability of Target Sets and Target Tubes," Automatica, Vol. 7, pp. 233-247.
[BeR73] Bertsekas, D. P., and Rhodes, 1. B., 1973. "Sufficiently Informative Functions and the Minimax Feedback Control of Uncertain Dynamic Systems," IEEE Trans. Automatic Control, Vol. AC-18, pp. 117-124.
[BeS78] Bertsekas, D. P., and Shreve, S. E., 1978. Stochastic Optimal Control: The Discrete Time Case, Academic Press, N. Y.; republished by Athena Scientific, Belmont, MA, 1996; can be downloaded from the author's website.
[BeS03] Bertsimas, D., and Sim, M., 2003. "Robust Discrete Optimization and Network Flows," Math. Programming, Series B, Vol. 98, pp. 49-71.
[BeT89] Bertsekas, D. P., and Tsitsiklis, J. N., 1989. Parallel and Distributed Computation: Numerical Methods, Prentice-Hall, Englewood Cliffs, N. J.; republished by Athena Scientific, Belmont, MA, 1997.
[BeT91] Bertsekas, D. P., and Tsitsiklis, J. N., 1991. "An Analysis of Stochastic Shortest Path Problems," Math. Operations Res., Vol. 16, pp. 580-595.
[BeT96] Bertsekas, D. P., and Tsitsiklis, J. N., 1996. Neuro-Dynamic Programming Athena Scientific, Belmont, MA.
[BeT97] Bertsimas, D., and Tsitsiklis, J. N., 1997. Introduction to Linear Optimization Athena Scientific, Belmont, MA.
[BeI57] Bellman, R., 1957. Dynamic Programming, Princeton University Press, Princeton, N. J.
[Ber70] Bertsekes, D. P., 1970. "On the Separation Theorem for Linear Systems, Quadratic Criteria, and Correlated Noise," Unpublished Report, Electronic Systems Lab., Massachusetts Institute of Technology.
[Ber71] Bertsekas, D. P., 1971. "Control of Uncertain Systems With a Set-Membership Description of the Uncertainty," Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA (available in scanned form from the author's www site).
[Ber72a] Bertsekas, D. P., 1972. "Infinite Time Reachability of State Space Regions by Using Feedback Control," IEEE Trans. Automatic Control, Vol. AC-17, pp. 604-613.
[Ber72b] Bertsekas, D. P., 1972. "On the Solution of Some Minimax Control Problems" Proc. 1972 IEEE Decision and Control Conf., New Orleans, LA.
[Ber74] Bertsekas, D. P., 1974. "Necessary and Sufficient Conditions for Existence of an Optimal Portfolio," J. Econ. Theory, Vol. 8, pp. 235-247.
[Ber75] Bertsekas, D. P., 1975. "Convergence of Discretization Procedures in Dynamic Programming," IEEE Trans. Automatic Control, Vol. AC-20, pp. 415-419.
[Ber76] Bertsekas, D. P., 1976. Dynamic Programming and Stochastic Control, Academic Press, N. Y.
[Ber82a] Bertsekas, D. P., 1982. "Distributed Dynamic Programming," IEEE Trans. Automatic Control, Vol. AC-27, pp. 610-616.
[Ber82b] Bertsekas, D. P., 1982. Constrained Optimization and Lagrange Multiplier Methods, Academic Press, N. Y.; republished by Athena Scientific, Belmont, MA, 1996.
[Ber93] Bertsekas, D. P., 1993. "A Simple and Fast Label Correcting Algorithm for Shortest Paths," Networks, Vol. 23, pp. 703-709.
[Ber97] Bertsekas, D. P., 1997. "Differential Training of Rollout Policies," Proc. of the

35th Allerton Conference on Communication, Control, and Computing, Allerton Park, Ill.
[Ber98a] Bertsekas, D. P., 1998. Network Optimization: Continuous and Discrete Models, Athena Scientific, Belmont, MA.
[Ber98b] Bertsekas, D. P., 1998. "A New Value Iteration Method for the Average Cost Dynamic Programmming Problem," SIAM J. on Control and Optimization, Vol. 36, pp. 742-759.
[Ber99] Bertsekas, D. P., 1999. Nonlinear Programming, (2nd Ed.), Athena Scientific, Belmont, MA
[BiL97] Birge, J. R, and Louveaux, 1997. Introduction to Stochastic Programming, Springer-Verlag, New York, N. Y.
[Pux) 5 ] Pishon C. M. 1995. Neural Networks for Pattern Recognition, Oxford University Press, N. Y.
[BITOO] Blondel, V. D., and Tsitsiklis, J. N., 2000. "A buirev of Computational Complexity Results in Systems and Control," Automatica, Vol. 36, pp. 1249-1274.
[Bla99] BIanchini, F., 1999. "Set Invariance in Control A Survey," A.-ltomatica, Vol 35, pp. 1747-1768.
[BoV79] Borkar, V., and Varaiya, P. P., 1979. "Adaptive Control of Markov Chains, I: Finite Parameter Set," IEEE Trans. Automatic Control, Vol. AC-24, pp. 953-958.
[CGC04] Chang, H. S., Givan, R L., and Chong, E. K. P., 2004. "Parallel Rollout for Online Solution of Partially Observable Markov Decision Processes," Discrete Event Dynamic Systems, Vol. 14, pp. 309-341.
[CaB04] Camacho, E. F., and Bordons, C., 2004. Model Predictive Control, 2nd Edition, Springer-Verlag, New York, N. Y.
[ChT89] Chow, C.-S., and Tsitsiklis, J. N., 1989. "The Complexity of Dynamic Programming," Journal of Complexity, Vol. 5, pp. 466-488.
[ChT91] Chow, C.-S., and Tsitsiklis, J. N., 1991. "An Optimal One-Way Multigrid Algorithm for Discrete-Time Stochastic Control," IEEE Trans. on Automatic Control, Vol. AC-36, 1991, pp. 898-914.
[Che72] Chernoff, H., 1972. "Sequential Analysis and Optimal Design," Regional Conference Series in Applied Mathematics, SIAM, Philadelphia, PA.
[Chr97] Christodouleas, J. D., 1997. "Solution Methods for Multiprocessor Network Scheduling Problems with Application to Railroad Operations," Ph.D. Thesis, Operations Research Center, Massachusetts Institute of Technology.
(Chu60] Chung, K. L., 1960. Markov Chains with Stationary Transition Probabilities, Springer-Verlag, N. Y.
[CoL55] Coddington, E. A., and Levinson, N., 1955. Theory of Ordinary Differential Equations, McGraw-Hill, N. Y.
[DeG70] DeGroot, M. H., 1970. Optimal Statistical Decisions, McGraw-Hill, N. Y.
[DeP84] Deo, N., and Pang, C., 1984. "Shortest Path Problems: Taxonomy and Annotation," Networks, Vol. 14, pp. 275-323.
[DeI89] Deller, J. R, 1989. "Set Membership Identification in Digital Signal Processing," IEEE ASSP Magazine, Oct., pp. 4-20.
[DoS80] Doshi, B., and Shreve, S., 1980. "Strong Consistency of a Modified Maximum

Likelihood Estimator for Controlled Markov Chains," J. of Applied Probability, Vol. 17, pp. 726-734.
[Dre65] Dreyfus, S. D., 1965. Dynamic Programming and the Calculus of Variations Academic Press, N. Y.
[Dre69] Dreyfus, S. D., 1969. "An Appraisal of Some Shortest-Path Algorithms," Operations Research, Vol. 17, pp. 395-412.
[Eck68] Eckles, J. E., 1968. "Optimum Maintenance with Incomplete Information," Operations Res., Vol. 16, pp. 1058-1067.
[Elm78] Elmaghraby, S. E., 1978. Activity Networks: Project Planning and Control by Network Models, Wiley-Interscience, N. Y.
[FaI87] Falcone, M., 1987. "A Numerical Approach to the Infinite Horizon Problem of Deterministic Control Theory," Appl. Math. Opt., Vol. 15, pp. 1-13.
[FeM94] Fernandez-Gaucherand, E., and Markus, S. 1., 1994. "Risk Sensitive Optimal Control of Hidden Markov Models," Proc. 33rd IEEE Conf. Dec. Control, Lake Buena Vista, Fla.
[FeV02] Ferris, M. C., and Voelker, M. M., 2002. "Neuro-Dynamic Programming for Radiation Treatment Planning," Numerical Analysis Group Research Report NA-02/06, Oxford University Computing Laboratory, Oxford University.
[FeV04] Ferris, M. C., and Voelker, M. M., 2004. "Fractionation in Radiation Treatmen Planning," Mathematical Programming B, Vol. 102, pp. 387-413.
[Fel68] Feller, W., 1968. An Introduction to Probability Theory and its Applications Wiley, N. Y.
[Fis70] Fishburn, P. C., 1970. Utility Theory for Decision Making, Wiley, N. Y.
[For56] Ford, L. R, Jr., 1956. "Network Flow Theory," Report P-923, The Rand Corporation, Santa Monica, CA.
[For73] Forney, G. D., 1973. "The Viterbi Algorithm," Proc. IEEE, Vol. 61, pp. 268-278. [Fox71] Fox, B. L., 1971. "Finite State Approximations to Denumerable State Dynamic Progams," J. Math. Anal. Appl., Vol. 34, pp. 665-670.
[GaP88] Gallo, G., and Pallottino, S., 1988. "Shortest Path Algorithms," Annals of Operations Research, Vol. 7, pp. 3-79.
[GaI99] Gallager, R G., 1999. Discrete Stochastic Processes, Kluwer, Boston.
[GoR85] Gonzalez, R, and Rofman, E., 1985. "On Deterministic Control Problems: An Approximation Procedure for the Optimal Cost, Parts I, II," SIAM J. Control Optimization, Vol. 23, pp. 242-285.
[GoS84] Goodwin, G. C., and Sin, K. S. S., 1984. Adaptive Filtering, Prediction, and Control, Prentice-Hall, Englewood Cliffs, N. J.
[GrA66] Groen, G. J., and Atkinson, R C., 1966. "Models for Optimizing the Learning Process," Psychol. BulL, Vol. 66, pp. 309-320
[GuF63] Gunckel, T. L., and Franklin, G. R, 1963. "A General Solution for Linear Sampled-Data Control," Trans. ASME Ser. D. J. Basic Engrg., Vol. 85, pp. 197-201.
[GuMOl] Guerriero, F., and Musmanno, R, 2001. "Label Correcting Methods to Solve Multicriteria Shortest Path Problems," J. Optimization Theory Appl., Vol. 111, pp. 589-613.
[GuM03] Guerriero, F., and Mancini, M., 2003. "A Cooperative Parallel Rollout Algorithm for the Sequential Ordering Problem," Parallel Computing, Vol. 29, pp. 663-677. [HMS55] Holt, C. C., Modigliani, F., and Simon, H. A., 1955. "A Linear Decision Rule for Production and Employment Scheduling," Management Sci., Vol. 2, pp. 1-30.
[HPC96] Helmsen, J., Puckett, E. G., Colella, P., and Dorr, M., 1996. "Two New Methods for Simulating Photolithography Development," SPIE, Vol. 2726, pp. 253-261.
[HaL82] Hajek, B., and van Loon, T., 1982. "Decentralized Dynamic Control of a Mul tiaccess Broadcast Channel," IEEE Trans. Automatic Control, Vol. AC-27, pp. 559-569
[Hak70] Hakansson, N. H., 1970. "Optimal Investment and Consumption Strategies under Risk for a Class of Utility Functions," Econometrica, Vol. 38, pp. 587-607.
[Hak71] Hakansson, N. H., 1971. "On Myopic Portfolio Policies, With and Without Serial Correlation of Yields," The Journal of Business of the University of Chicago, Vol. 44 pp. 324-334.
[Han80] Hansen, P., 1980. "Bicriterion Path Problems," in Multiple-Criteria Decision Making: Theory and Applications, Edited by G. Fandel and T. Gal, Springer Verlag, Heidelberg, Germany, pp. 109-127.
[Hay98] Haykin, S., 1998. Neural Networks: A Comprehensive Foundation, (2nd Ed.), McMillan, N. Y.
[Hes66] Hestenes, M. R, 1966. Calculus of Variations and Ontimal Control Theory, Wiley, N. Y.
[Her89] Hernand= werma, 0., 1989. Adaptive Markov Control Processes, SpringerVerlac, -. Y .
[HoK71] Hoffman, K., and Kunze, R, 1971. Linear Algebra, Prentice-Hall, Englewood Cliffs, N. J.
[IEE71] IEEE Trans. Automatic Control, 1971. Special Issue on Linear-Quadratic Gaussian Problem, Vol. AC-16.
[IoS96] Ioannou, P. A., and Sun, J., 1996. Robust Adaptive Control, Prentice-Hall, Englewood Cliffs, N. J.
[JBE94] James, M. R, Baras, J. S., and Elliott, R J., 1994. "Risk-Sensitive Control and Dynamic Games for Partially Observed Discrete-Time Nonlinear Systems," IEEE Trans. on Automatic Control, Vol. AC-39, pp. 780-792.
[Jac73] Jacobson, D. H., 1973. "Optimal Stochastic Linear Systems With Exponential Performance Criteria and their Relation to Deterministic Differential Games," IEEE Trans. Automatic Control, Vol. AC-18, pp. 124-131.
[Jaf84] Jaffe, J. M., 1984. "Algorithms for Finding Paths with Multiple Constraints," Networks, Vol. 14, pp. 95-116.
[Jew63] Jewell, W., 1963. "Markov Renewal Programming I and II," Operations Research, Vol. 2, pp. 938-971.
[JoT61] Joseph, P. D., and Tou, J. T., 1961. "On Linear Control Theory," AlEE Trans., Vol. 80 (II), pp. 193-196.
 Control Design, J. Wiley, N. Y.
[KGB82] Kimemia, J., Gershwin, S. B., and Bertsekas, D. P., 1982. "Computation of Production Control Policies by a Dynamic Programming Technique," in Analysis and

Optimization of Systems, A. Bensoussan and J. L. Lions (eds.), Springer-Verlag, N. Y. pp. 243-269.
[KLB92] Kosut, R. L., Lau, M. K., and Boyd, S. P., 1992. "Set-Membership Identification of Systems with Parametric and Nonparametric Uncertainty," IEEE Trans. on Automatic Control, Vol. AC-37, pp. 929-941.
[KaD66] Karush, W., and Dear, E. E., 1966. "Optimal Stimulus Presentation Strategy for a Stimulus Sampling Model of Learning," J. Math. Psychology, Vol. 3, pp. 15-47.
[KaK58] Kalman, R E., and Koepcke, R W., 1958. "Optimal Synthesis of Linear Sampling Control Systems Using Generalized Performance Indexes," Trans. ASME, Vol. 80, pp. 1820-1826
[KaW94] Kall, P., and Wallace, S. W., 1994. Stochastic Programming, Wiley, Chichester, UK.
[Ka160] Kalman, R E., 1960. "A New Approach to Linear Filtering and Prediction Problems," Trans. ASME Ser. D. J. Basic Engrg., Vol. 82, pp. 35-45.
[KeS60] Kemeny, J. G., and Snell, J. L., 1960. Finite Markov Chains, Van NostrandReinhold, N. Y.
[KeG88] Keerthi, S. S., and Gilbert, E. G., 1988. "Optimal, Infinite Horizon Feedback Laws for a General Class of Constrained Discete Time Systems: Stability and MovingHorizon Approximations," J. Optimization Theory Appl., Vo. 57, pp. 265-293
[Kim82] Kimemia, J., 1982. "Hierarchical Control of Production in Flexible Manufacturing Systems," Ph.D. Thesis, Dep. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology.
[KuA77] Ku, $\mathrm{N}_{\mathrm{n} .,}$, and Athans, M., 1977. "Further Results on the Uncertainty Threshold Principle," IEEE Trans. on Automatic Control, Vol. AC-22, pp. "866-868.
[KuD92] Kushner, H. J., and Dupuis, P. G., 1992. Numerical Methods for Stochastic Control Problems in Continuous Time, Springer-Verlag, N. Y.
[KuL82] Kumar, P. R, and Lin, W., 1982. "Optimal Adaptive Controllers for Unknown Markov Chains," IEEE Trans. Automatic Control, Vol. AC-27, pp. 765-774.
[KuV86] Kumar, P. R, and Varaiya, P. P., 1986. Stochastic Systems: Estimation, Identification, and Adaptive Control, Prentice-Hall, Englewood Cliffs, N. J.
[KuV97] Kurzhanski, A., and Valyi, 1., 1997. Ellipsoidal Calculus for Estimation and Control, Birkhauser, Boston, MA.
[Kum83] Kumar, P. R, 1983. "Optimal Adaptive Control of Linear-Quadratic-Gaussian Systems," SIAM J. on Control and Optimization, Vol. 21, pp. 163-178.
[Kum85] Kumar, P. R, 1985. "A Survey of Some Results in Stochastic Adaptive Control," SIAM J. on Control and Optimization, Vol. 23, pp. 329-380.
[Kus90] Kushner, H. J., 1990. "Numerical Methods for Continuous Control Problems in Continuous Time," SIAM J. on Control and Optimization, VoL 28, pp. 999-1048.
[Las85] Lasserre, J. B., 1985. "A Mixed Forward-Backward Dynamic Programming Method Using Parallel Computation," J. Optimization Theory Appl., VoL 45, pp. 165168.
[Lev84] Levy, D., 1984. The Chess Computer Handbook, B. T. Batsford Ltd., London
[LiR71] Lippman, S. A., and Ross, S. M., 1971. "The Streetwalker's Dilemma: A JobShop Model," SIAM J. of Appl. Math., VoL 20, pp. 336-342.
[LjS83] Ljung, L., and Soderstrom, T., 1983. Theory and Practice of Recursive Identification, MIT Press, Cambridge, MA.
[Lju86] Ljung, L., 1986. System Identification: Theory for the User, Prentice-Hall, Englewood Cliffs, N. J.
[Lov91a] Lovejoy, W. S., 1991. "Computationally Feasible Bounds for Partially Observed Markov Decision Processes," Operations Research, Vol. 39, pp. 162175.
[Lov91b] Lovejoy, W. S., 1991. "A Survey of Algorithmic Methods for Partially Observed Markov Decision Processes," Annals of Operations Research, Vol. 18, pp. 47-66
[Lue69] Luenberger, D. G., 1969. Optimization by Vector Space Methods, Wiley, N. Y. [Lue84] Luenberger, D. G., 1984. Linear and Nonlinear Programming, Addi-son...Wesley Reading, MA.
[MHKYoj N., Hauskrecht, M., Kim, K.-E., Peshkin, L., Kaelbling, L. K., and Dean, T., 1998. "Solving Vely _ ?rge Weakly Coupled Markov Decision Processes," Proc. of the Fifteenth National Conference on rtificial Intelligence, Madison, WI, pp. 165-172. [MMB02] McGovern, A., Moss, E., and Barto, 2., 2002. "Building a Basic Building Block Scheduler Using Reinforcement Learning and Rollouts," Machine Learning, Vol. 49, pp. 141-160.
[MPP04] Meloni, C., Pacciarelli, D., and Pranzo, M., 2004. "A Rollout Metaheuristic for Job Shop Scheduling Problems," Annals of Operations Research, Vol. 131, pp. 215-235.
[MRROO] Mayne, D. Q., Rawlings, J. B., Rao, C. V., and Scokaert, P. O. M., 2000. "Constrained Model Predictive Control: Stability and Optimality," Automatica, Vol. 36, pp. 789-814.
[Mac02] Maciejowski, J. M., 2002. Predictive Control with Constraints, Addison-Wesley, Reading, MA.

Martins, E. Q. V., 1984. "On a Multicriteria Shortest Path Problem," European J. of Operational Research, Vol. 16, pp. 236-245.
[MayOl] Mayne, D. Q., 2001. "Control of Constrained Dynamic Systems," European Journal of Control, Vol. 7, pp. 87-99.
[McQ66] MacQueen, J., 1966. "A Modified Dynamic Programming Method for Markovian Decision Problems," J. Math. Anal. Appl., Vol. 14, pp. 38-43.
[Mik79] Mikhailov, V. A., 1979. Methods of Random Multiple Access, Candidate Engineering Thesis, Moscow Institute of Physics and Technology, Moscow.
[MoL99] Morari, M., and Lee, J. H., 1999. "Model Predictive Control: Past, Present, and Future," Computers and Chemical Engineering, Vol. 23, pp. 667-682.
[Mos68] Mossin, J., 1968. "Optimal Multi-Period Portfolio Policies," J. Business, Vol. 41, pp. 215-229.
[NeW88] Nemhauser, G. L., and Wolsey, L. A., 1988. Integer and Combinatorial Opti mization, Wiley, N. Y
[New75] Newborn, M., 1975. Computer Chess, Academic Press, N. Y.
[Nic66] Nicholson, T., 1966. "Finding the Shortest Route Between Two Points in a Network," The Computer Journal, Vol. 9, pp. 275-280.
[Nil71] Nilsson, N. J., 1971. Problem-Solving Methods in Artificial Intelligence, McGraw Hill, N. Y.
[Ni180] Nilsson, N. J., 1971. Principles of Artificial Intelligence, Morgan-Kaufmann, San Mateo, Ca .
[PBG65] Pontryagin, L. S., Boltyanski, V., Gamkrelidze, R, and Mishchenko, E., 1965 The Mathematical Theory of Optimal Processes, Interscience Publishers, Inc., N. Y.
[PBT98] Polymenakos, L. C., Bertsekas, D. P., and Tsitsiklis, J. N., 1998. "Efficient Algorithms for Continuous-Space Shortest Path Problems," IEEE Trans. on Automatic Control, Vol. 43, pp. 278-283
[PaS82] Papadimitriou, C. H., and Steiglitz, K., 1982. Combinatorial Optimization: Algorithms and Complexity, Prentice-Hall, Englewood Cliffs, N. J.
[PaT87] Papadimitriou, C. H., and Tsitsiklis, J. N., 1987. "The Complexity of Markov Decision Processes," Math. Operations Res., Vol. 12, pp. 441-450.
[Pap74] Pape, V., 1974. "Implementation and Efficiency of Moore Algorithms for the Shortest Path Problem," Math. Progr., Vol. 7, pp. 212-222.
[Pap65] Papoulis, A., 1965. Probability, Random Variables and Stochastic Processes McGraw-Hill, N. Y.
[Pea84] Pearl, J., 1984. Heuristics, Addison-Wesley, Reading, MA
[Pic90] Picone, J., 1990. "Continuous Speech Recognition Using Hidden Markov Modth, " IEEE ASSP Magazine, July Issue, pp. 26-41.
[Pin95] Pinedo, M., 1995. Scheuuling: Theory, Algorithms, and Systems, Prentice-Hall, Englewood Cliffs, N. J.
[Pra64] Pratt, J. W., 1964. "Risk Aversion in the Small and in the Large," Econometrica, Vol. 32, pp. 300-307
[Pre95] Prekopa, A., 1995. Stochastic Programming, Kluwer, Boston
[PrS94] Proakis, J. G., and Salehi, M., 1994. Communication Systems Engineering, Prentice-Hall, Englewood Cliffs, N. J.
[Rab89] Rabiner, L. R, 1989. "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition," Proc. of the IEEE, Vol. 77, pp. 257-286.
[Roc70] Rockafellar, R T., 1970. Convex Analysis, Princeton University Press, Princeton, N. J.
[Ros70] Ross, S. M., 1970. Applied Probability Models with Optimization Applications, Holden-Day, San Francisco, CA
[Ros83] Ross, S. M., 1983. Introduction to Stochastic Dynamic Programming, Academic Press, N. Y.
[Ros85] Ross, S. M., 1985. Probability Models, Academic Press, Orlando, Fla.
[Roy88] Royden, H. L., 1988. Principles of Mathematical Analysis, (3rd Ed.), McGrawHill, N. Y.
[Rud76] Rudin, W., 1976. Real Analysis, (3rd Ed.), McGraw-Hill, N. Y.
[Rus97] Rust, J., 1997. "Using Randomization to Break the Curse of Dimensionality," Econometrica, Vol. 65, pp. 487-516.
[SBB89] Sastry, S., Bodson, M., and Bartram, J. F., 1989. Adaptive Control: Stability, Convergence, and Robustness, Prentice-Hall, Englewood Cliffs, N. J.
[SGC02] Savagaonkar, U., Givan, R, and Chong, E. K. P., 2002. "Sampling Techniques for Zero-Sum, Discounted Markov Games," in Proc. 40th Allerton Conference on Com-
munication, Control and Computing, Monticello, Ill.
[Sam69] Samuelson, P. A., 1969. "Lifetime Portfolio Selection by Dynamic Stochastic Programming," Review of Economics and Statistics, Vol. 51, pp. 239-246.
[Sar87] Sargent, T. J., 1987. Dynamic Macroeconomic Theory, Harvard Univ. Press, Cambridge, MA.
[Sca60] Scarf, H., 1960. "The Optimality of ( $s$, S) Policies for the Dynamic Inventory Problem," Proceedings of the 1st Stanford Symposium on Mathematical Methods in the Social Sciences, Stanford University Press, Stanford, CA.
[Sch68] Schweppe, F. C., 1968. "Recursive State Estimation; Unknown but Bounded Errors and System Inputs," IEEE Trans. Automatic Control, Vol. AC-13.
[Sch74] Schweppe, F. C., 1974. Uncertain Dynamic Systems, Academic Press, N. Y.
[Sch97] Schaeffer, J., 1997. One Jump Ahead, Springer-Verlag, N. Y.
[SecOO] Secomandi, N., 2000. "Comparing Neuro-Dynamic Programming Algorithms for the Vehicle Routing Problem with Stochastic Demands," Computers and Operations Resem Vol. 27, pp. 1201-1225.
[SecOl] Secomandi, $\mathrm{I} .$, - "A Rollout Policy for the Vehicle Routing Problem with Stochastic Demands," Operations Restar_in, 49. pp. 796-802.
[Sec03] Secomandi, N., 2003. "Analysis of a Rollout Approach to Sequentin Problems with Stochastic Routing Applications," J. of Heuristics, Vol. 9, pp. 321-352.

Sethian, J. A., 1999. Level Set Methods and Fast Marching Methods Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science, Cambridge University Press, N. Y.
[Set99b] Sethian, J. A., 1999. "Fast Marching Methods," SIAM Review, Vol. 41, pp. 199-235.
[Sha50] Shannon, C., 1950. "Programming a Digital Computer for Playing Chess," Phil. Mag., Vol. 41, pp. 356-375.
[Shi64] Shiryaev, A. N., 1964. "On Markov Sufficient Statistics in Non-Additive Bayes Problems of Sequential Analysis," Theory of Probability and Applications, Vol. 9, pp. 604-618.
[Shi66] Shiryaev, A. N., 1966. "On the Theory of Decision Functions and Control by an Observation Process with Incomplete Data," Selected Translations in Math. Statistics and Probability, Vol. 6, pp. 162-188.

Shreve, S. E., 1981. "A Note on Optimal Switching Between Two Activities," Naval Research Logistics Quarterly, Vol. 28, pp. 185-190.
[Sim56] Simon, H. A., 1956. "Dynamic Programming Under Uncertainty with a Quadratic Criterion Function," Econometrica, Vol. 24, pp. 74-81.
[SkI88] Sklar, B., 1988. Digital Communications: Fundamentals and Applications, Pren-tice-Hall, Englewood Cliffs, N. J.
[SIL91] Slotine, J.-J. E., and Li, W., Applied Nonlinear Control, Prentice-Hall, Englewood Cliffs, N. J.
[SmS73] Smallwood, R D., and Sondik, E. J., 1973. "The Optimal Control of Partially Observable Markov Processes Over a Finite Horizon," Operations Res., Vol. 11, pp. 1071-1088.
[Sma71] Smallwood, RD., 1971. "The Analysis of Economic Teaching Strategies for a Simple Learning Model," J. Math. Psychology, Vol. 8, pp. 285-301.
[Son71] Sondik, E. J., 1971. "The Optimal Control of Partially Observable Markov Processes," Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, Stanford, CA.
[StW91] Stewart, B. S., and White, C. C., 1991. "Multiobjective $A^{*}$ " J. ACM, Vol. 38, pp. 775-814.
[Sti94] Stirzaker, D., 1994. Elementary Probability, Cambridge University Press Cambridge.
[StL89] Stokey, N. L., and Lucas, R E., 1989. Recursive Methods in Economic Dynamics, Harvard University Press, Cambridge, MA.
[Str65] Striebel, C. T., 1965. "Sufficient Statistics in the Optimal Control of Stochastic Systems," J. Math. Anal. Appl., Vol. 12, pp. 576-592.
[Str76] Strang, G., 1976. Linear Algebra and its Applications, Academic Press, N. Y.
[SuB98] Sutton, R, and Barto, A. G., 1998. Reinforcement Learning, MIT Press, CambrIdge, MA.
[TeG96] Tesauro, G., and Galperin, G. R, 1996. "On-Line Policy Improvement Using Monte Carlo Search," presented at the 1996 Neural Information Processing Systems Conference, Denver, CO ; also in M. Mozer et al. (eds.), Advances in Neural Information Processing Systems 9, MIT Press (1997).

Thau, F. E., and Witsenhausen, H. S., 1966. "A Comparison of Closed-Loop and Open-Loop Optimum Systems," IEEE Trans. Automatic Control, Vol. AC-11, pp. 619-621.
[The54] Theil, H., 1954. "Econometric Models and Welfare Maximization," Weltwirtsch. Arch., Vol. 72, pp. 60-83.
[Tsi84a] Tsitsiklis, J. N., 1984. "Convexity and Characterization of Optimal Policies in a Dynamic Routing Problem," J. Optimization Theory Appl., Vol. 44, pp. 105-136.
[Tsi84b] Tsitsiklis, J. N., 1984. "Periodic Review Inventory Systems with Continuous Demand and Discrete Order Sizes," Management ScL, Vol. 30, pp. 1250-1254.
[Tsi87] Tsitsiklis, J. N., 1987. "Analysis of a Multiaccess Control Scheme," IEEE Trans. Automatic Control, Vol. AC-32, pp. 1017-1020.
[Tsi95] Tsitsiklis, J. N., 1995. "Efficient Algorithms for Globally Optimal Trajectories," IEEE Trans. Automatic Control, Vol. AC-40, pp. 1528-1538.
[TuP03] Tu, F., and Pattipati, K. R, 2003. "Rollout Strategies for Sequential Fault Diagnosis," IEEE Trans. on Systems, Man and Cybernetics, Part A, pp. 86-99.
[YuB04] Yu, H., and Bertsekas, D. P., 2004. "Discretized Approximations for POMDP with Average Cost," Proc. of 20th Conference on Uncertainty in Artificial Intelligence, Banff, Canada.
[VaW89] Varaiya, P., and Wets, R J-B., 1989. "Stochastic Dynamic Optimization Approaches and Computation," Mathematical Programming: State of the Art, M. Iri and K. Tanabe (eds.), Kluwer, Boston, pp. 309-332.
[Vei65] Veinott, A. F., Jr., 1965. "The Optimal Inventory Policy for Batch Ordering," Operations Res., Vol. 13, pp. 424-432.
[Vei66] Veinott, A. F., Jr., 1966. "The Status of Mathematical Inventory Theory," Management Sci., Vol. 12, pp. 745-777.
[Vin74] Vincke, P., 1974. "Problemes Multicriteres," Cahiers du Centre d' Etudes de Recherche Operationnelle, Vol. 16, pp. 425-439.
[Vit67] Viterbi, A. J., 1967. "Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm," IEEE Trans. on Info. Theory, Vol. IT-13, pp. 260-269.
[WCG03] Wu, G., Chong, E. K. P., and Givan, R. L., 2003. "Congestion Control Using Policy Rollout," Proc. 2nd IEEE CDC, Maui, Hawaii, pp. 4825-4830.
[Wa147] Wald, A., 1947. Sequential Analysis, Wiley, N. Y.
[WeP80] Weiss, G., and Pinedo, M., 1980. "Scheduling Tasks with Exponential Service Times on Nonidentical Processors to Minimize Various Cost Functions," J. Appl. Prob., Vol. 17, pp. 187-202.
[vVhH80] White, C. C., and Harrington, D. P., 1980. "Application of Jensen's Inequality to Adaptive Suboptimal Design," J. Optimization Theory Appl., Vol. 32, pp. 89-99.
[WhS89] White, C. C., and Scherer, W. T., 1989. "Solution Procedures for Partially Observed Markov Decision Processes," Operations Res., Vol. 30, pp. 791-797.
[Whi63] Whittle, P., 1963. Prediction and Regulation by Linear Least-Square Methods, English Universities Press, London.
[vVhi69] White, D. J., 1969. Dynamic Programming, Holden-Day, San Francisco, CA.
[Whi78] Whitt, W., 1978. "Approximations of Dynamic Programs I," Math. Operations Res., Vol. 3, pp. 231-243.
[Whi79] Whitt, W., 1979. "Approximations of Dynamic Programs II," Math. Operations Res., Vol. 4, pp. 179-185.
[vVhi82] 'Whittle, P., 1982. Optimization Over Time, Wiley, N. Y., Vol. 1, 1982, Vol. 2, 1983.
[vVhi90] Whittle, P., 1990. Risk-Sensitive Optimal Control, Wiley, N. Y
[Wit66] vVitsenhausen, H. S., 1966. "Minimax Control of Uncertain Systems," Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA.
[Wit68] vVitsenhausen, H. S., 1968. "Sets of Possible States of Linear Systems Given Perturbed Observations," IEEE Trans. Automatic Control, Vol. AC-13.
[Wit69] Witsenhausen, H. S., 1969. "Inequalities for the Performance of Suboptimal Uncertain Systems," Automatica, Vol. 5, pp. 507-512.
[vVit70] Witsenhausen, H. S., 1970. "On Performance Bounds for Uncertain Systems," SIAM J. on Control, Vol. 8, pp. 55-89.
[Wit71] vVitsenhausen, H. S., 1971. "Separation of Estimation and Control for DiscreteTime Systems," Proc. IEEE, Vol. 59, pp. 1557-1566.
[WoI98] Wolsey, L. A., 1998. Integer Programming, Wiley, N. Y.
[vVuB99] Wu, C. C., and Bertsekas, D. P., 1999. "Distributed Power Control Algorithms for Wireless Networks," unpublished report, available from the author's www site.
[YDR05] Yan, X., Diaconis, P., Rusmevichientong, P., and Van Roy, B., 2005. "Solitaire: Man Versus Machine," Advances in Neural Information Processing Systems, Vol. 17, to appear.

INDEX

A
A* algorithm, 87, 95
ARMAX model, 238, 503, 506
Adaptive control, 289
Adjoint equation, 118, 130
Admissible policy, 13, 219
Aggregation, 319
Aggregation probabilities, 320
Alpha-beta pruning, 332
Asset selling, 176, 278, 420
Asynchronous algorithms, 102
Augmentation of state, 35
Autoregressive process, 239
Average cost problem, 403, 421, 441

## B

Backward shift operator, 236
Basic problem, 12, 218
Bayes' rule, 475
Bellman, 51
Bellman's equation, 404, 408, 418, 426, 440, 443
Best-first search, 86
Brachistochrone problem, 133
Branch-and-bound algorithm, 88
Breadth-first search, 84
Breakthrough problem, 338, 368, 388, 396, 397

## C

CEC, 283
Calculus of variations, 108, 120
Capacity expansion, 207
Caution, 289
Certainty equivalence principle, 28, 161
Certainty equivalent control, 283, 293,
309
Chess, 11, 15, 32, 327
Closed-loop control, 4
Closed set, 465
Coarse grid, 324
Communicating states, 478
Compact set, 465
Composition of functions, 466
Concave function, 467
Conditional probability, 475
Constrained DP, 95, 397
Constrained shortest path, 91

Constrained controllability, 370, 374
Constraint feasibility problem, 91
Constraint qualification, 470
Continuous function, 465
Continuously differentiable, 466
Controllability, 152, 370
Control law, 13
Control trajectory, 106
Convergence of sequences, 465
Convex function, 467
Convex set, 467
Convolutional coding, 74
Correlated disturbances, 37, 181, 271
Cost-to-go function, 24
Countable set, 460
Covariance matrix, 474
Critical path analysis, 68
Cumulative distribution function, 473

## D

D'Esopo-Pape method, 86
DP algorithm, 18, 222, 256
DP algorithm proof, 23, 44, 48
Data networks, 97, 102
Decision function, 525
Delays, 35
Depth-first search, 85, 327
Detectability, 159
Differential cost, 429
Dijkstra's algorithm, 86, 99, 101, 384
389
Disaggregation probabilities, 320
Discounted cost, 52, 403, 417, 438
Discretization, 324, 382
Distributed computation, 102
Distribution function, 473
Disturbance, 13
Dominant decision, 510
Dual control, 289
E
Eigenvalue, 462
Eigenvector, 462
Euclidean space, 460
Event, 472
Existence of optimal solutions, 469
Expected value, 474

Exponential cost function, 53, 202
Exponential distribution, 437
F
Fast marching method, 389
Feature extraction, 326
Feature vectors, 326
Feedback controller, 525
First passage time, 411, 480
Flexible manufacturing, 316
Forecasts, 38, 175, 195, 202, 304
Fortified rollout algorithm, 355
Forward DP algorithm, 66
Forward shift operator, 237
Four queens problem, 77
Full rank, 461

## G

## Gambling, 208

Gauss-Markov estimator, 499
Gaussian random vector, 234, 273, 483, 484
Gradient, 466
Gradient matrix, 466
Greedy algorithm, 339, 349

## H

Hamilton-Jacobi-Bellman equation, 109 Hamiltonian function, 118
Hard aggregation, 321
Hidden Markov model, 70
Hypothesis testing, 266 I

Identifiability, 291, 293
Improper policy, 406
Independent random variables, 474
Infimum, 460
Information vector, 219
Inner product, 460
Interchange argument, 186, 189 Inventory control, 3, 21, 162, 204 Investment problems, 60, 170 Irreducible Markov chain, 479 Isoperimetric problem, 144
Iterative deepening, 334

## K

K-convexity, 166
Kalman filter, 192, 234, 481, 49
Killer heuristic, 334

> L

L'Hopital's problem, 145

LLL strategy 86, 97, 98
Label correcting method, 81, 384, 389
Label setting method, 86, 384, 389
Limit inferior, 465
Limit point, 465
Limit superior, 465
Limited lookahead policy, 304
Linear independence, 461
Linear programming, 416
Linear quadratic problems, 27, 114, 122,
148, 202, 229, 241, 270

## M

Markov chains, 477
Mean first passage time, 411, 480
Memoryless property, 437
Minimax algorithm, 331
Minimum principle, 119, 129
Minimum-time problem, 136
Minimum variance control, 236, 296
Minimax control problems, 46, 197,332,
374, 388, 398
Model predictive control, 366, 369, 376,

## 388, 398

Monotonicity property, 59
Moving average process, 239
Multiaccess communication, 219, 287
Multiobjective DP, 93
Multiobjective shortest path, 92
Multiplicative cost, 54
Multistep lookahead, 304, 359
Myopic policy, 175

## N

Noninferior decision, 92, 510
Noninferior solution, 92
Norm, 460
O
OLFC, 300
Observability, 152
One-step-Iookahead rule, 184
Open-loop control, 4, 301, 376
Open-loop feedback control, 300, 376
Open set, 465
Optimal cost function, 14
Optimal value function, 14
Optimality conditions, 470
Optimality principle, 18, 93
Optimization in policy space, 386

## P

POLFC, 303
Partially myopic policy, 175

Partial open-loop feedback control, 303 376
Payoff function, 50
Pole-zero cancellation, 243
Policy, 13, 525
Policy iteration, 336, 414, 419, 432
Pontryagin minimum principle, 115, 119
Portfolio analysis, 170
Positive definite matrix, 463
Positive semidefinite matrix, 463
Principle of optimality, 18
Probability density function, 474
Probability distribution, 473
Probability space, 472
Probing, 289
Proper policy, 406
Pursuit-evasion game, 215

Q
Q-factor, 342, 361, 363
Quadratic cost, 27, 114, 122, 148, 229
240, 369
Queueing control, 10, 34

## R

Random variable, 473
Rank of a matrix, 461
Rational spectrum, 504
Reachability, 197,201,214,215,370,374, 388
Recurrent state, 479
Relative cost, 429
Relative value iteration, 431
Replacement problems, 8, 34
Riccati equation, 114, 151
Riccati equation convergence, 153
Risk, 17, 53, 521
Rolling horizon, 367
Rollout algorithm, 307, 335, 372, 376

## S

SLF method, 86, 97, 98
Scenarios, 313
Scheduling problems, 7, 19, 186
Self-tuning regulator, 298
Sequential consistency, 349
Sequential improvement, 353
Semi-Markov problems, 435
Semilinear systems, 55, 391
Separation theorem, 233
Sequential hypothesis testing, 266
Sequential probability ratio, 270
Set-membership estimation, 191
Set-membership models, 191, 373, 388

Shortest path problem, 65, 384, 389, 406
Singular problem, 139
Slotted Aloha, 220
Soft aggregation, 321
Speech recognition, 73
Stabilizability, 159
Stable filter, 238
Stable system, 153
State trajectory, 106
Stationary policy, 405
Stochastic matrix, 477
Stochastic programming, 310
Stochastic shortest paths, 384, 403, 405
Stopping problems, 176
Sufficient statistic, 252
Supremum, 460
Symmetric matrix, 461

## T

Terminating process, 53
Terminating rollout algorithm, 348
Tetris, 41
Time lags, 35
Total probability theorem, 475
Transient state, 479
Transition probabilities, 478
Transition rate, 437
Transpose, 461
Traveling salesman problem, 78, 347, 349

## u

Uncertainty threshold principle, 160
Uncontrollable state components, 39
Uncorrelated random variables, 474
Unknown-but-bounded disturbances, 197
Utility function, 171, 513, 517, 526
Utility theory, 511

## V

Value iteration, 413, 418, 430, 445
Value of information, 14
Vehicle routing, 315
Viterbi algorithm, 73, 288

## W

Weierstrass theorem, 469


[^0]:    $\dagger$ For the benefit of the mathematically oriented reader we note that in the preceding equation, "min" denotes the greatest lower bound (or infimum) of the set of numbers $\left\{J_{\pi-}(x o) \mid \pi \mathrm{E} I \mathrm{I}\right\}$. A notation more in line with normal mathematical usage would be to write $J^{*}\left(x_{0}\right)=\inf _{\pi \in \Pi} J_{\pi}\left(x_{0}\right)$. However (as discussed in Appendix B), we find it convenient to use "min" in place of "inf" even when the infimum is not attained. It is less distracting, and it will not lead to any confusion.

[^1]:    $t$ In the case where there are multiple arcs connecting a node pair, we can merge all these arcs to a single arc, since the set of destination nodes that can be reached from any non-destination node will not be affected.

[^2]:    t For an example where this convention for tie-breaking is not observed and as a consequence $\mathcal{R} \mathcal{H}$ does not terminate, assume that there is a single destination $d$ and that all other nodes are arranged in a cycle. Each non-destination node i has two outgoing arcs: one arc that belongs to the cycle, and another arc which is (i, $d$ ). Let $\mathcal{H}$ be the (sequentially consistent) base heuristic that starting from a node $\mathrm{i} \neq d$, generates the path ( $\mathrm{i}, \mathrm{d}$ ). When the terminal node of the path is node $i$, the rollout algorithm $\mathcal{R} \mathcal{H}$ compares the two neighbors of $i$, which are $d$ and the node next to i on the cycle, call it $j$. Both neighbors have $d$ as their projection, so there is tie in Eq. (6.35). It can be seen that if $\mathcal{R} \mathcal{H}$ breaks ties in favor of the neighbor $j$ that lies on the cycle, then $\mathcal{R} \mathcal{H}$ continually repeats the cycle and never terminates.

[^3]:    $t$ Let us call a stationary policy $\pi$ proper if the condition (7.3) is satisfied for some $m$, and call $\pi$ improper otherwise. It can be shown that Assumption 7.2.1 is equivalent to the seemingly weaker assumption that all stationary policies are proper (see Vol. II, Exercise 2.3). However, the results of Prop. 7.2.1 can also be shown under the genuinely weaker assumption that there exists at least one proper policy, and furthermore, every improper policy results in infinite expected cost from at least one initial state (see Bertsekas and Tsitsiklis [BeT89], [BeT91], or Vol. II, Chapter 2). These assumptions, when specialized to deterministic shortest path problems, are similar to the ones we used in Chapter 2. They imply that there is at least one path to the destination from every starting node and that all cycles have positive cost. Still another set of assumptions under which the results of Prop. 7.2.1 hold is described in Exercise 7.28, where again improper policies are allowed, but the stage costs $g(\mathrm{i}, u)$ are assumed nonnegative, and the optimal costs $J^{*}(\mathrm{i})$ are assumed finite.

[^4]:    t For a heuristic argument, let $\lambda_{\mu}$ be the average cost per stage corresponding to a stationary policy $\mu$ and consider a trajectory of the system under $\mu$ starting from state $n$. If $\mathrm{C}_{1}, C_{2}, \ldots \mathrm{em}$ are the costs incurred in the first m cycles, and

[^5]:    t If $N$ is not countable, it is necessary to introduce a probability space structure on Nand $\mathcal{O}$ as in Appendix C. Furthermore, it is necessary that the function $f(d,$.$) satisfy certain (measurability) assumptions.$
    $\ddagger$ The term "lottery" is associated with conceptually convenient device of viewing outcomes as prizes of some sort and viewing a fixed probabilistic mechanism for winning a prize as a lottery.

