

*Dynamic Programming
and Optimal Control*

Volume II

THIRD EDITION

Dimitri P. Bertsekas

Massachusetts Institute of Technology

WWW site for book information and orders

<http://www.athenasc.com>



Athena Scientific, Belmont, Massachusetts

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

Email: info@athenasc.com
WWW: <http://www.athenasc.com>

Cover Design: Ann Gallager, www.gallagerdesign.com

© 2007, 2001, 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 information storage and retrieval) without permission in writing from
the publisher.

Publisher's Cataloging-in-Publication Data

Bertsekas, Dimitri P.
Dynamic Programming and Optimal Control
Includes Bibliography and Index
1. Mathematical Optimization. 2. Dynamic Programming. I. Title.
QA402.5 .B465 2007 519.703 01-75941

ISBN 1-886529-30-2 (Vol. II)
ISBN 1-886529-26-4 (Vol. I)
ISBN 1-886529-08-6 (Two-volume set – latest editions)

Contents

1. Infinite Horizon – Discounted Problems

1.1. Minimization of Total Cost – Introduction	p. 3
1.1.1. The Finite-Horizon DP Algorithm	p. 5
1.1.2. Shorthand Notation and Monotonicity	p. 6
1.1.3. A Preview of Infinite Horizon Results	p. 8
1.1.4. Randomized and History-Dependent Policies	p. 10
1.2. Discounted Problems with Bounded Cost per Stage	p. 12
1.3. Finite-State Systems – Computational Methods	p. 19
1.3.1. Value Iteration and Error Bounds	p. 22
1.3.2. Variants of Value Iteration	p. 30
1.3.3. Policy Iteration	p. 38
1.3.4. Linear Programming	p. 51
1.3.5. Limited Lookahead and Rollout Policies	p. 53
1.4. The Role of Contraction Mappings	p. 57
1.4.1. Sup-Norm Contractions	p. 57
1.4.2. m -Stage Sup-Norm Contractions	p. 62
1.4.3. Discounted Problems - Unbounded Cost per Stage	p. 63
1.5. Scheduling and Multiarmed Bandit Problems	p. 66
1.6. Notes, Sources, and Exercises	p. 78

2. Stochastic Shortest Path Problems

2.1. Problem Formulation	p. 94
2.2. Bellman’s Equation	p. 97
2.3. Value Iteration	p. 105
2.4. Policy Iteration	p. 108
2.5. Countable State Problems	p. 112
2.6. Notes, Sources, and Exercises	p. 114

3. Undiscounted Problems

3.1. Unbounded Costs per Stage	p. 142
--	--------

3.2. Linear Systems and Quadratic Cost	p. 140
3.3. Inventory Control	p. 142
3.4. Optimal Stopping	p. 145
3.5. Optimal Gambling Strategies	p. 150
3.6. Nonstationary and Periodic Problems	p. 157
3.7. Notes, Sources, and Exercises	p. 162
4. Average Cost per Stage Problems	
4.1. Finite-Spaces Average Cost Models	p. 174
4.1.1. Relation with the Discounted Cost Problem	p. 178
4.1.2. Blackwell Optimal Policies	p. 184
4.1.3. Optimality Equations	p. 194
4.2. Conditions for Equal Average Cost for all Initial States	p. 198
4.3. Value Iteration	p. 203
4.3.1. Single-Chain Value Iteration	p. 206
4.3.2. Multi-Chain Value Iteration	p. 222
4.4. Policy Iteration	p. 229
4.4.1. Single-Chain Policy Iteration	p. 229
4.4.2. Multi-Chain Policy Iteration	p. 234
4.5 Linear Programming	p. 239
4.6. Infinite Space Problems	p. 245
4.6.1. A Sufficient Condition for Optimality	p. 253
4.6.2. Finite State Space and Infinite Control Space	p. 254
4.6.3. Countable States – Vanishing Discount Approach	p. 264
4.6.4. Countable States – Contraction Approach	p. 267
4.6.5. Linear Systems with Quadratic Cost	p. 271
4.7. Notes, Sources, and Exercises	p. 273
5. Continuous-Time Problems	
5.1. Uniformization	p. 286
5.2. Queueing Applications	p. 293
5.3. Semi-Markov Problems	p. 304
5.4. Notes, Sources, and Exercises	p. 315
6. Approximate Dynamic Programming	
6.1. Cost Approximation	p. 325
6.2. Approximate Policy Iteration – Direct Policy Evaluation	p. 329
6.2.1. Gradient Methods for Direct Policy Evaluation	p. 332
6.2.2. TD(λ)	p. 336
6.2.3. Optimistic Policy Iteration	p. 337
6.2.4. Approximate Policy Iteration Based on Q -Factors	p. 338
6.3. Indirect Methods for Policy Evaluation	p. 340
6.3.1. Policy Evaluation by Projected Value Iteration	p. 341

6.3.2. Least Squares Policy Evaluation (LSPE)	p. 346
6.3.3. PVI(λ) and LSPE(λ)	p. 348
6.3.4. The LSTD(λ) Algorithm	p. 355
6.3.5. The TD(λ) Algorithm	p. 357
6.3.6. Summary and Examples	p. 359
6.4. <i>Q</i> -Learning	p. 362
6.4.1. <i>Q</i> -Factor Approximations	p. 364
6.4.2. <i>Q</i> -Learning for Optimal Stopping Problems	p. 366
6.5. Stochastic Shortest Path Problems	p. 369
6.6. Average Cost Problems	p. 374
6.6.1. Approximate Policy Evaluation	p. 374
6.6.2. Approximate Policy Iteration	p. 385
6.6.3. <i>Q</i> -Learning for Average Cost Problems	p. 388
6.7. Approximation in Policy Space	p. 391
6.8. Notes, Sources, and Exercises	p. 398

Appendix A: Measure-Theoretic Issues in Dynamic Programming
p. 403

A.1. A Two-Stage Example	p. 403
A.2. Resolution of the Measurability Issues	p. 408

References p. 417

Index p. 439

CONTENTS OF VOLUME I

1. The Dynamic Programming Algorithm

- 1.1. Introduction
- 1.2. The Basic Problem
- 1.3. The Dynamic Programming Algorithm
- 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

2. Deterministic Systems and the Shortest Path Problem

- 2.1. Finite-State Systems and Shortest Paths
- 2.2. Some Shortest Path Applications
 - 2.2.1. Critical Path Analysis
 - 2.2.2. Hidden Markov Models and the Viterbi Algorithm
- 2.3. Shortest Path Algorithms
 - 2.3.1. Label Correcting Methods
 - 2.3.2. Label Correcting Variations - A^* Algorithm
 - 2.3.3. Branch-and-Bound
 - 2.3.4. Constrained and Multiobjective Problems
- 2.4. Notes, Sources, and Exercises

3. Deterministic Continuous-Time Optimal Control

- 3.1. Continuous-Time Optimal Control
- 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. The Minimum Principle for Discrete-Time Problems
- 3.4. Extensions of the Minimum Principle
 - 3.4.1. Fixed Terminal State
 - 3.4.2. Free Initial State
 - 3.4.3. Free Terminal Time
 - 3.4.4. Time-Varying System and Cost
 - 3.4.5. Singular Problems
- 3.5. Notes, Sources, and Exercises

4. Problems with Perfect State Information

- 4.1. Linear Systems and Quadratic Cost
- 4.2. Inventory Control
- 4.3. Dynamic Portfolio Analysis
- 4.4. Optimal Stopping Problems
- 4.5. Scheduling and the Interchange Argument

- 4.6. Set-Membership Description of Uncertainty
 - 4.6.1. Set-Membership Estimation
 - 4.6.2. Control with Unknown-but-Bounded Disturbances
- 4.7. Notes, Sources, and Exercises

5. Problems with Imperfect State Information

- 5.1. Reduction to the Perfect Information Case
- 5.2. Linear Systems and Quadratic Cost
- 5.3. Minimum Variance Control of Linear Systems
- 5.4. Sufficient Statistics and Finite-State Markov Chains
 - 5.4.1. The Conditional State Distribution
 - 5.4.2. Finite-State Systems
- 5.5. Notes, Sources, and Exercises

6. Suboptimal and Adaptive Control

- 6.1. Certainty Equivalent and Adaptive Control
 - 6.1.1. Caution, Probing, and Dual Control
 - 6.1.2. Two-Phase Control and Identifiability
 - 6.1.3. Certainty Equivalent Control and Identifiability
 - 6.1.4. Self-Tuning Regulators
- 6.2. Open-Loop Feedback Control
- 6.3. Limited Lookahead Policies and Applications
 - 6.3.1. Performance Bounds for Limited Lookahead Policies
 - 6.3.2. Computational Issues in Limited Lookahead
 - 6.3.3. Problem Approximation - Enforced Decomposition
 - 6.3.4. Aggregation
 - 6.3.5. Parametric Cost-to-Go Approximation
- 6.4. Rollout Algorithms
 - 6.4.1. Discrete Deterministic Problems
 - 6.4.2. Q -Factors Evaluated by Simulation
 - 6.4.3. Q -Factor Approximation
- 6.5. Model Predictive Control and Related Methods
 - 6.5.1. Rolling Horizon Approximations
 - 6.5.2. Stability Issues in Model Predictive Control
 - 6.5.3. Restricted Structure Policies
- 6.6. Additional Topics in Approximate DP
 - 6.6.1. Discretization
 - 6.6.2. Other Approximation Approaches
- 6.7. Notes, Sources, and Exercises

7. Introduction to Infinite Horizon Problems

- 7.1. An Overview
- 7.2. Stochastic Shortest Path Problems
- 7.3. Discounted Problems
- 7.4. Average Cost Problems

- 7.5. Semi-Markov Problems
- 7.6. Notes, Sources, and Exercises

Appendix A: Mathematical Review

- A.1. Sets
- A.2. Euclidean Space
- A.3. Matrices
- A.4. Analysis
- A.5. Convex Sets and Functions

Appendix B: On Optimization Theory

- B.1. Optimal Solutions
- B.2. Optimality Conditions
- B.3. Minimization of Quadratic Forms

Appendix C: On Probability Theory

- C.1. Probability Spaces
- C.2. Random Variables
- C.3. Conditional Probability

Appendix D: On Finite-State Markov Chains

- D.1. Stationary Markov Chains
- D.2. Classification of States
- D.3. Limiting Probabilities
- D.4. First Passage Times

Appendix E: Least-Squares Estimation and Kalman Filtering

- E.1. Least-Squares Estimation
- E.2. Linear Least-Squares Estimation
- E.3. State Estimation – Kalman Filter
- E.4. Stability Aspects
- E.5. Gauss-Markov Estimators
- E.6. Deterministic Least-Squares Estimation

Appendix F: Modeling of Stochastic Linear Systems

- F.1. Linear Systems with Stochastic Inputs
- F.2. Processes with Rational Spectrum
- F.3. The ARMAX Model

Appendix G: Formulating Problems of Decision Under Uncertainty

- G.1. The Problem of Decision Under Uncertainty
- G.2. Expected Utility Theory and Risk
- G.3. Stochastic Optimal Control Problems

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.I.T.), where he is currently McAfee Professor of Engineering.

His research spans several fields, including optimization, control, large-scale computation, and data communication networks, and is closely tied to his teaching and book authoring activities. He has written numerous 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.

ATHENA SCIENTIFIC
OPTIMIZATION AND COMPUTATION 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, 2007, ISBN 1-886529-08-6, 920 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 Rockafellar, 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. Constrained Optimization and Lagrange Multiplier Methods, by Dimitri P. Bertsekas, 1996, ISBN 1-886529-04-3, 410 pages
11. Stochastic Optimal Control: The Discrete-Time Case, by Dimitri P. Bertsekas and Steven E. Shreve, 1996, ISBN 1-886529-03-5, 330 pages

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 Maximum Principle, developed from a dynamic programming viewpoint. I have also added a fairly extensive exposition of simulation-based 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 Polymenakos, 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,† 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 Tsitsiklis 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
bertsekas@lids.mit.edu
<http://web.mit.edu/dimitrib/www/home.html>

† Note added in the 2nd edition: This monograph was republished by Athena Scientific in 1996.


Preface to the Second Edition

This second edition of Vol. II should be viewed as a relatively minor revision of the original. The coverage was expanded in a few areas as follows:

- (a) In Chapter 1, material was added on variants of the policy iteration method.
- (b) In Chapter 2, the material on neuro-dynamic programming methods was updated and expanded to reflect some recent developments.
- (c) In Chapter 4, material was added on some new value iteration methods.
- (d) In Chapter 5, the material on semi-Markov problems was revised, with a significant portion simplified and shifted to Volume I.

There are also a few miscellaneous additions and improvements scattered throughout the text. 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 

I would like to express my thanks to the many colleagues who contributed suggestions for improvement of the second edition.

Dimitri P. Bertsekas
bertsekas@lids.mit.edu
<http://web.mit.edu/dimitrib/www/home.html>

June, 2001

Preface to the Third Edition

This is a major revision of the 2nd edition, and contains a substantial amount of new material, as well as a reorganization of old material. The length of the text has increased by more than 50%, and more than half of the old material has been restructured and/or revised. Most of the added material is in four areas.

- (a) The coverage of the average cost problem of Chapter 4 has greatly increased in scope and depth. In particular, there is now a full analysis of multi-chain problems, as well as a more extensive analysis of infinite-spaces problems (Section 4.6).
- (b) The material on approximate dynamic programming has been collected in Chapter 6. It has been greatly expanded to include new research, thereby supplementing the 1996 book “Neuro-Dynamic Programming” by John Tsitsiklis and the author.
- (c) Contraction mappings and their role in various analyses have been highlighted in new material on infinite state space problems (Sections 1.4, 2.5, and 4.6), and in their use in the approximate dynamic programming material of Chapter 6.
- (d) The mathematical measure-theoretic issues that must be addressed for a rigorous theory of stochastic dynamic programming have been illustrated and summarized in an appendix for the benefit of the mathematically oriented reader.

Also some exercises were added and a few sections were revised while preserving their essential content.

I would like to express my thanks to many colleagues who contributed valuable suggestions. I am particularly thankful to Ciamac Moallemi, Steven Shreve, John Tsitsiklis, and Ben Van Roy, who reviewed some of the new material and each contributed several substantial suggestions. I wish to thank especially Janey Yu for her extraordinary help. Janey read with great care and keen eye large parts of the book, contributed analysis, and a large number of incisive, substantive comments, and also collaborated with me in research that was included in Chapter 6.

Dimitri P. Bertsekas
<http://web.mit.edu/dimitrib/www/home.html>
Fall 2006