Optimization over Time. Dynamic Programming and Stochastic Control. Volume 1.
PETER WHITTLE
John Wiley, U.K./U.S.A., 1982. 317pp. £19.50
ISBN 0 471 10120 6

This volume, the first of two, is a welcome addition to dynamic programming literature. Although not claiming to be encyclopaedic, the total work is a thorough survey of most of the areas of current interest in optimization over time. This area has seen some remarkable developments over the last few years, such as the Gittens index for multi-armed bandits, variations on the linear dynamic systems or the work on partially observable systems, and this is the first textbook to incorporate them. Perhaps textbook is a misnomer for, although there are exercises, they teach new concepts rather than reinforce those within the text, and the overall impression is more of a reference book incorporating the important results of the subjects.

This volume is split into three parts. The first half of the book deals with deterministic dynamic programming, which introduces the ideas of the optimality equations for such problems, as well as examples in allocation, scheduling, inventory, C.P.A. and trajectory problems. It describes the linear dynamic models, introduces the ideas of stability, controllability and observability for them and ends with two chapters introducing the maximum principle.

The second part deals with stochastic dynamic programming models which are really Markov decision processes with a finite time horizon. It uses the general results obtained earlier to solve the secretary problem, replacement problems, portfolio selection, gambling and inventory problems. Then comes a very lucid account of the index approach to multi-armed bandit processes and its applications. This was the highlight of the book for me, as the author's own approach reveals the reasonableness and the strength of the results.

The last part of the book returns to the linear dynamic model with quadratic cost function and Gaussian noise. It shows the certainty equivalence (i.e. taking the expected state value in the deterministic optimal control as the optimal stochastic control) and introduces the Kalman filter technique. The theory is extended to risk-sensitive criteria, and the spectral theory approach to the problem is also discussed.

The author displays his expertise in the area to the full. The proofs of the theorems are slick, the examples are illuminating and the pitfalls in the theory carefully pointed out. It should be on the bookshelf of everyone interested in optimal control. Yet, it is not an easy read. The area itself makes notation complicated, and by dealing first in generalities the author adds to the problem. Also there is not one numerical example in the text and, though algebraic examples give more general results, they also emphasise the theoretical tone of the book. I also did not like the references ahead to results in Volume Two, which, until its appearance, leaves one grasping thin air.

Having said that, this book is a very significant contribution to the subject and is a unifying account of some of the most important areas in dynamic optimization. It does this with care and understanding and must become a standard reference for the area. Roll on Volume Two.

L.C. THOMAS