Continuous Time Markov Decision Processes Theory And Applications

This book presents classical Markov Decision Processes (MDP) for real-life applications and optimization. MDP allows users to develop and formally support approximate and simple decision rules, and this book showcases state-of-the-art applications in which MDP was key to the solution approach. The book is divided into six parts. Part 1 is devoted to the state-of-the-art theoretical foundation of MDP, including approximate methods such as policy improvement, successive approximation and infinite state spaces as well as an instructive chapter on Approximate Dynamic Programming. It then continues with five parts of specific and non-exhaustive application areas. Part 2 covers MDP healthcare applications, which includes different screening procedures, appointment scheduling, ambulance scheduling and blood management. Part 3 explores MDP modeling within transportation. This ranges from public to private transportation, from airports and traffic lights to car parking or charging your electric car . Part 4 contains three chapters that illustrates the structure of approximate policies for production or manufacturing structures. In Part 5, communications is highlighted as an important application area for MDP. It includes Gittins indices, down-to-earth call centers and wireless sensor networks. Finally Part 6 is dedicated to financial modeling, offering an instructive review to account for financial portfolios and derivatives under proportional transactional costs. The MDP applications in this book illustrate a variety of both standard and non-standard aspects of MDP modeling and its practical use. This book should appeal to readers for practitioning, academic research and educational purposes, with a background in, among others, operations research, mathematics, computer science, and industrial engineering.

This volume collects papers, based on invited talks given at the IMA workshop in Modeling, Stochastic Control, Optimization, and Related Applications, held at the Institute for Mathematics and Its Applications, University of Minnesota, during May and June, 2018. There were four week-long workshops during the conference. They are (1) stochastic control, computation methods, and applications, (2) queueing theory and networked systems, (3) ecological and biological applications, and (4) finance and economics applications. For broader impacts, researchers from different fields covering both theoretically oriented and application intensive areas were invited to participate in the conference. It brought together researchers from multi-disciplinary communities in applied mathematics, applied probability, engineering, biology, ecology, and networked science, to review, and substantially update most recent progress. As an archive, this volume presents some of the highlights of the workshops, and collect papers covering a broad range of topics.

Put together by two top researchers in the Far East, this text examines Markov Decision Processes - also called stochastic dynamic programming - and their applications in the optimal control of discrete event systems, optimal replacement, and optimal allocations in sequential online auctions. This dynamic new book offers fresh applications of MDPs in areas such as the control of discrete event systems and the optimal allocations in sequential online auctions.

Continuous Time Markov Decision Models with Variable Information

Continuous-Time Markov Chains and Applications

Selected Topics on Continuous-time Controlled Markov Chains and Markov Games

Markov Decision Processes with Their Applications

Two-Time-Scale Methods and Applications

This book provides a unified approach for the study of constrained Markov decision processes with a finite state space and unbounded costs. Unlike the single controller case considered in many other books, the author considers a single controller with several objectives, such as minimizing delays and loss, probabilities, and maximization of throughputs. It is desirable to design a controller that minimizes one cost objective, subject to inequality constraints on other cost objectives. This framework describes dynamic decision problems arising frequently in many engineering fields. A thorough overview of these applications is presented in the introduction. The book is then divided into three sections that build upon each other. The first part explains the theory for the finite state space. The author characterizes the set of achievable expected occupation measures as well as performance vectors, and identifies simple classes of policies among which optimal policies exist. This allows the reduction of the original dynamic into a linear program. A Lagranian approach is then used to derive the dual linear program using dynamic programming techniques. In the second part, these results are extended to the infinite state space and action spaces. The author provides two frameworks: the case where costs are bounded below and the contracting framework. The third part builds upon the results of the first two part and examines asymptotical results of the convergence of both the value and the policies in the time horizon and in the discount factor. Finally, several state truncation algorithms that enable the approximation of the solution of the original control problem via finite linear programs are given.

Dynamic power management is a design methodology aiming at controlling performance and power levels of digital circuits and systems, with the goal of extending the autonomous operation time of battery-powered systems, providing graceful performance degradation when supply energy is limited, and adapting power dissipation to satisfy environmental constraints. Dynamic Power Management: Design Techniques and CAD Tools addresses design techniques and computer-aided design solutions for power management. Different approaches are presented and organized in an order related to their applicability to control-units, macro-blocks, digital circuits and electronic systems, respectively. All approaches are based on the principle of exploiting idleness of circuits, systems, or portions thereof. They involve both the detection of idleness conditions and the freezing of power-consuming activities in the idle components. The book also describes some approaches to system-level power management, including Microsoft's OnNow architecture and the `Advanced Configuration and Power Management' standard proposed by Intel, Microsoft and Toshiba. These approaches migrate power management to the software layer running on hardware platforms, thus providing a flexible and self-configurable solution to adapting the power/performance tradeoff to the needs of mobile (and fixed) computing and communication. Dynamic Power Management: Design Techniques and CAD Tools is of interest to researchers and developers of computer-aided design tools for integrated circuits and systems, as well as to system designers.

A path-breaking account of Markov decision processes-theory and computation This book's clear presentation of theory, numerous chapter-end problems, and development of a unified method for the computation of optimal policies in both discrete and continuous time make it an excellent course text for graduate students and advanced undergraduates. Its comprehensive coverage of important recent advances in stochastic dynamic programming makes it a valuable working resource for operations research professionals, management scientists, engineers, and others. Stochastic Dynamic Programming and the Control of Queueing Systems presents the theory of optimization under the finite horizon, infinite horizon discounted, and average cost criteria. It then shows how optimal rules of operation (policies) for each criterion may be numerically determined. A great wealth of examples from the application area of the control of queueing systems is presented. Nine numerical programs for the computation of optimal policies are fully explicated. The Pascal source code for the programs is available for viewing and downloading on the Wiley Web site at www.wiley.com/products/subject/mathematics. The site contains a link to the author's own Web site and is also a place where readers may discuss developments on the programs or other aspects of the material. The source files are also available via ftp at ftp://ftp.wiley.com/public/sci_tech_med/stochastic Stochastic Dynamic Programming and the Control of Queueing Systems features: * Path-breaking advances in Markov decision process techniques, brought together for the first time in book form * A theorem/proof format (proofs may be omitted without loss of continuity) * Development of a unified method for the computation of optimal rules of system operation * Numerous examples drawn mainly from the control of queueing systems * Detailed discussions of nine numerical programs * Helpful chapter-end problems * Appendices with complete treatment of background material

Finite State Continuous Time Markov Decision Processes with an Infinite Planning Horizon

Time-bounded Reachability in Continuous-time Markov Decision Processes

Initially Stationary [epsilon]-optimal Policies in Continuous Time Markov Decision Chains

Markov Decision Processes

Markov Decision Processes with Applications to Finance

This book covers formulation, algorithms, and structural results of partially observed Markov decision processes, whilst linking theory to real-world applications in controlled sensing. Computations are kept to a minimum, enabling students and researchers in engineering, operations research, and economics to understand the methods and determine the structure of their optimal solution.

This book concerns continuous-time controlled Markov chains and Markov games. The former, which are also known as continuous-time Markov decision processes, form a class of stochastic control problems in which a single decision-maker has a wish to optimize a given objective function. In contrast, there are two or more decision-makers (or players, or controllers) trying to optimize its own objective function in a Markov game. Both decision-making processes appear in a large number of applications in economics, operations research, engineering, and computer science among other areas. The main features of the control and game models studied in the book are the continuous time variable, the denumerable state space, and that the control (or action) sets are Borel spaces. Moreover, the transition and reward rates of the dynamical system may be unbounded. The authors are interested in some aspects of controlled Markov chains and Markov games such as characterizing the optimal reward functions, and determining optimal policies for each of the optimality criteria studied here. The main focus is on advanced optimality criteria (such as, bias, variance, sensitive discount, and Blackwell optimality), though they also deal with the basic optimality criteria (discounted and average reward). A particular emphasis is made on the application of the results presented in this book. One of the main concerns is to propose assumptions on the control and game models that are easily verifiable (and verified) in practice. Moreover, algorithmic and computational issues are also analyzed. In particular, the authors propose approximation results that allow precise numerical approximations of the solution to some problems of practical interest. Applications to population models and epidemic processes are also shown.

Covering formulation, algorithms, and structural results, and linking theory to real-world applications in controlled sensing (including social learning, adaptive radars and sequential detection), this book focuses on the conceptual foundations of partially observed Markov decision processes (POMDPs). It emphasizes structural results in stochastic dynamic programming, enabling graduate students and researchers in engineering, operations research, and economics to understand the underlying unifying themes without getting weighed down by mathematical technicalities. Bringing together research from across the literature, the book provides an introduction to nonlinear filtering followed by a systematic development of stochastic dynamic programming, lattice programming and reinforcement learning for POMDPs. Questions addressed in the book include: when does a POMDP have a threshold optimal policy? When are myopic policies optimal? How do local and global decision makers interact in adaptive decision making in multi-agent social learning where there is herding and data incest? And how can sophisticated radars and sensors adapt their sensing in real time?

Markov Processes for Stochastic Modeling

Partially Observed Markov Decision Processes

Continuous Time Control of Markov Processes on an Arbitrary State Space

Continuous Time Markov Processes

Discrete-Time Markov Chains

This book concerns continuous-time controlled Markov chains, also known as continuous-time Markov decision processes. They form a class of stochastic control problems in which a single decision-maker wishes to optimize a given objective function. This book is also concerned with Markov games, where two decision-makers (or players) try to optimize their own objective function. Both decision-making processes appear in a large number of applications in economics, operations research, engineering, and computer science, among other areas. An extensive, self-contained, up-to-date analysis of basic optimality criteria (such as discounted and average reward), and advanced optimality criteria (e.g., bias, overtaking, sensitive discount, and Blackwell optimality) is presented. A particular emphasis is made on the application of the results herein: algorithmic and computational issues are discussed, and applications to population models and epidemic processes are shown. This book is addressed to students and researchers in the fields of stochastic control and stochastic games. Moreover, it could be of interest also to undergraduate and beginning graduate students because the reader is not supposed to have a high mathematical background: a working knowledge of calculus, linear algebra, probability, and continuous-time Markov chains should suffice to understand the contents of the book.

The system considered may be in one of n states at any point in time; its probability law is a Markov process that depends on the policy (control) chosen. The return to the system over a given planning horizon is the integral over that horizon of a return rate that depends on both the policy and the sample path of the process. The objective is to find a policy that maximizes the expected discounted return as the planning horizon tends to infinity. The case where the discount factor goes to zero is also considered. In all cases it is shown that there is a stationary policy that is optimal, and an algorithm is given to obtain that policy. (Author).

The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence. Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics. Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

Stochastic Dynamic Programming and the Control of Queueing Systems

A Two-Time-Scale Approach

Continuous Time Markov Decision Processes with Discounted Returns

Further Topics on Discrete-Time Markov Control Processes

Semi-Markov and Decision Processes

Markov processes are processes that have limited memory. In particular, their dependence on the past is only through the previous state. They are used to model the behavior of many systems including communications systems, transportation networks, image segmentation and analysis, biological systems and DNA sequence analysis, random atomic motion and diffusion in physics, social mobility, population studies, epidemiology, animal and insect migration, queueing systems, resource management, dams, financial engineering, actuarial science, and decision systems. Covering a wide range of areas of application of Markov processes, this second edition is revised to highlight the most important aspects as well as the most recent trends and applications of Markov processes. The author spent over 16 years in the industry before returning to academia, and he has applied many of the principles covered in this book in multiple research projects. Therefore, this is an applications-oriented book that also includes enough theory to provide a solid ground in the subject for the reader. Presents both the theory and applications of the different aspects of Markov processes Includes numerous solved examples as well as detailed diagrams that make it easier to understand the principle being presented Discusses different applications of hidden Markov models, such as DNA sequence analysis and speech analysis.

This dissertation studies the problem of the control of Markov processes with continuous time parameter and arbitrary state space. The economic criteria used are (a) the expected discounted return over an infinite horizon and (b) the expected average return over an infinite horizon. An extensive theory is available to treat the control problems in which either the time parameter or the state space is discrete. This dissertation extends the available theory to the general case of continuous time parameter and arbitrary state space. An application to the control of the arrival process in an M/G/1 queue is included. (Modified author abstract).

Devoted to a systematic exposition of some recent developments in the theory of discrete-time Markov control processes, the text is mainly confined to MCPs with Borel state and control spaces. Although the book follows on from the author's earlier work, an important feature of this volume is that it is self-contained and can thus be read independently of the first. The control model studied is sufficiently general to include virtually all the usual discrete-time stochastic control models that appear in applications to engineering, economics, mathematical population processes, operations research, and management science.

An Introduction

Discrete Stochastic Dynamic Programming

Modeling, Stochastic Control, Optimization, and Applications

Finite State Continuous-time Markov Decision Processes

Dynamic Probabilistic Systems

The system considered may be in one of n states at any point in time and its probability law is a Markov process which depends on the policy (control) chosen. The return to the system over a given planning horizon is the integral (over that horizon) of a return rate which depends on both the policy and the sample path of the process. The objective is to find a policy which maximizes the expected return over the given planning horizon. A necessary and sufficient condition for optimality is obtained, and a constructive proof is given that there is a piecewise constant policy which is optimal. A bound on the number of switches (points where the piecewise constant policy jumps) is obtained for the case where there are two states. (Author).

Continuous-time Markov decision processes (MDPs), also known as controlled Markov chains, are used for modeling decision-making problems that arise in operations research (for instance, inventory, manufacturing, and queueing systems), computer science, communications engineering, control of populations (such as fisheries and epidemics), and management science, among many other fields. This volume provides a unified, systematic, self-contained presentation of recent developments on the theory and applications of continuous-time MDPs. The MDPs in this volume include most of the cases that arise in applications, because they allow unbounded transition and reward/cost rates. Much of the material appears for the first time in book form.

This book focuses on two-time-scale Markov chains in discrete time. Our motivation stems from existing and emerging applications in optimization and control of complex systems in manufacturing, wireless communication, and ?nancial engineering. Much of our e?ort in this book is devoted to designing system models arising from various applications, analyzing them via analytic and probabilistic techniques, and developing feasible compu- tionalschemes. Ourmainconcernistoreducetheinherentsystemcompl- ity. Although each of the applications has its own distinct characteristics, all of them are closely related through the modeling of uncertainty due to jump or switching random processes. Oneofthesalientfeaturesofthisbookistheuseofmulti-timescalesin Markovprocessesandtheirapplications. Intuitively,notallpartsorcom- nents of a large-scale system evolve at the same rate. Some of them change rapidly and others vary slowly. The di?erent rates of variations allow us to reduce complexity via decomposition and aggregation. It would be ideal if we could divide a large system into its smallest irreducible subsystems completely separable from one another and treat each subsystem indep- dently. However, this is often infeasible in reality due to various physical constraints and other considerations. Thus, we have to deal with situations in which the systems are only nearly decomposable in the sense that there are weak links among the irreducible subsystems, which dictate the oc- sional regime changes of the system. An e?ective way to treat such near decomposability is time-scale separation. That is, we set up the systems as if there were two time scales, fast vs. slow. xii Preface Followingthetime-scaleseparation,weusesingularperturbationmeth- ology to treat the underlying systems.

Markov Decision Processes in Practice

Selected Topics on Continuous-Time Controlled Markov Chains and Markov Games

An Equivalence Between Continuous and Discrete Time Markov Decision Processes

Continuous-time Markov Decision Processes

Borel Space Models and General Control Strategies

This book concerns continuous-time controlled Markov chains, also known as continuous-time Markov decision processes. They form a class of stochastic control problems in which a single decision-maker wishes to optimize a given objective function. This book is also concerned with Markov games, where two decision-makers (or players) try to optimize their own objective function. Both decision-making processes appear in a large number of applications in economics, operations research, engineering, and computer science, among other areas. An extensive, self-contained, up-to-date analysis of basic optimality criteria (such as discounted and average reward), and advanced optimality criteria (e.g., bias, overtaking, sensitive discount, and Blackwell optimality) is presented. A particular emphasis is made on the application of the results herein: algorithmic and computational issues are discussed, and applications to population models and epidemic processes are shown. This book is addressed to students and researchers in the fields of stochastic control and stochastic games. Moreover, it could be of interest also to undergraduate and beginning graduate students because the reader is not supposed to have a high mathematical background: a working knowledge of calculus, linear algebra, probability, and continuous-time Markov chains should suffice to understand the contents of the book. Contents:IntroductionControlled Markov ChainsBasic Optimality CriteriaPolicy Iteration and Approximation TheoremsOvertaking, Bias, and Variance OptimalitySensitive Discount OptimalityBlackwell OptimalityConstrained Controlled Markov ChainsApplicationsZero-Sum Markov GamesBias and Overtaking Equilibria for Markov Games Readership: Graduate students and researchers in the fields of stochastic control and stochastic analysis. Keywords:Markov Decision Processes;Continuous-Time Controlled Markov Chains;Stochastic Dynamic Programming;Stochastic GamesKey Features:This book presents a reader-friendly, extensive, self-contained, and up-to-date analysis of advanced optimality criteria for continuous-time controlled Markov chains and Markov games. Most of the material herein is quite recent (it has been published in high-impact journals during the last five years) and it appears in book form for the first timeThis book introduces approximation theorems which, in particular, allow the reader to obtain numerical approximations of the solution to several control problems of practical interest. To the best of our knowledge, this is the first time that such computational issues are studied for denumerable state continuous-time controlled Markov chains. Hence, the book has an adequate balance between, on the one hand, theoretical results and, on the other hand, applications and computational issuesThe books that analyze continuous-time controlled Markov chains usually restrict themselves to the case of bounded transition and reward rates, which can be reduced to discrete-time models by using the uniformization technique. In our case, however, the transition and the reward rates might be unbounded, and so the uniformization technique cannot be used. By the way, let us mention that in models of practical interest the transition and the reward rates are, typically, unboundedReviews:“The book contains a large number of recent research results on CMCs and Markov games and puts them in perspective. It is written in a very conscious manner, contains detailed proofs of all main results, as well as extensive bibliographic remarks. The

book is a very valuable piece of work for researchers on continuous-time CMCs and Markov games."Zentralblatt MATH

The theory of Markov decision processes focuses on controlled Markov chains in discrete time. The authors establish the theory for general state and action spaces and at the same time show its application by means of numerous examples, mostly taken from the fields of finance and operations research. By using a structural approach many technicalities (concerning measure theory) are avoided. They cover problems with finite and infinite horizons, as well as partially observable Markov decision processes, piecewise deterministic Markov decision processes and stopping problems. The book presents Markov decision processes in action and includes various state-of-the-art applications with a particular view towards finance. It is useful for upper-level undergraduates, Master's students and researchers in both applied probability and finance, and provides exercises (without solutions).

This book gives a systematic treatment of singularly perturbed systems that naturally arise in control and optimization, queueing networks, manufacturing systems, and financial engineering. It presents results on asymptotic expansions of solutions of Komogorov forward and backward equations, properties of functional occupation measures, exponential upper bounds, and functional limit results for Markov chains with weak and strong interactions. To bridge the gap between theory and applications, a large portion of the book is devoted to applications in controlled dynamic systems, production planning, and numerical methods for controlled Markovian systems with large-scale and complex structures in the real-world problems. This second edition has been updated throughout and includes two new chapters on asymptotic expansions of solutions for backward equations and hybrid LQG problems. The chapters on analytic and probabilistic properties of two-time-scale Markov chains have been almost completely rewritten and the notation has been streamlined and simplified. This book is written for applied mathematicians, engineers, operations researchers, and applied scientists. Selected material from the book can also be used for a one semester advanced graduate-level course in applied probability and stochastic processes.

Continuous-Time Markov Decision Processes

Finite Stage Continuous Time Markov Decision Processes with an Infinite Planning Horizon

Bisimulation and Logical Preservation for Continuous-time Markov Decision Processes

Design Techniques and CAD Tools

Finite State Continuous-time Markov Decision Processes, with Applications to a Class of Optimization Problems in Queueing Theory

This book offers a systematic and rigorous treatment of continuous-time Markov decision processes, covering both theory and possible applications to queueing systems, epidemiology, finance, and other fields. Unlike most books on the subject, much attention is paid to problems with functional constraints and the realizability of strategies. Three major methods of investigations are presented, based on dynamic programming, linear programming, and reduction to discrete-time problems. Although the main focus is on models with total (discounted or undiscounted) cost criteria, models with average cost criteria and with impulsive controls are also discussed in depth. The book is self-contained. A separate chapter is devoted to Markov pure jump processes and the appendices collect the requisite background on real analysis and applied probability. All the statements in the main text are proved in detail. Researchers and graduate students in applied probability, operational research, statistics and engineering will find this monograph interesting, useful and valuable.

The Wiley-Interscience Paperback Series consists of selected booksthat have been made more accessible to consumers in an effort toincrease global appeal and general circulation. With these newunabridged softcover volumes, Wiley hopes to extend the lives ofthese works by making them available to future generations ofstatisticians, mathematicians, and scientists. "This text is unique in bringing together so many resultshitherto found only in part in other texts and papers. . . . Thetext is fairly self-contained, inclusive of some basic mathematicalresults needed, and provides a rich diet of examples, applications,and exercises. The bibliographical material at the end of eachchapter is excellent, not only from a historical perspective, butbecause it is valuable for researchers in acquiring a goodperspective of the MDP research potential." —Zentralblatt fur Mathematik ". . . it is of great value to advanced-level students,researchers, and professional practitioners of this field to havenow a complete volume (with more than 600 pages) devoted to thistopic. . . . Markov Decision Processes: Discrete Stochastic DynamicProgramming represents an up-to-date, unified, and rigoroustreatment of theoretical and computational aspects of discrete-timeMarkov decision processes." —Journal of the American Statistical Association

A continuous time Markov decision process with bounded sojourn parameters (the parameters of the exponential sojourn times in the states) is shown to be equivalent to a simpler discrete time Markov decision process. This is for both the discounted and average cost criteria on an infinite horizon. This result follows from (1) an equivalence of certain functionals of equivalent Markov processes, and (2) the property that a continuous time jump Markov process with bounded sojourn parameters is equal in distribution to a similar Markov process whose sojourn parameters are all the same. (Author).

Reinforcement Learning, second edition

Continuous Average Control of Piecewise Deterministic Markov Processes

Markov Decision Processes with Continuous Time Parameter

Dynamic Power Management

Theory and Applications

This invaluable book provides approximately eighty examples illustrating the theory of controlled discrete-time Markov processes. Except for applications of the theory to real-life problems like stock exchange, queues, gambling, optimal search etc, the main attention is paid to counter-intuitive, unexpected properties of optimization problems. Such examples illustrate the importance of conditions imposed in the theorems on Markov Decision Processes. Many of the examples are based upon examples published earlier in journal articles or textbooks while several other examples are new. The aim was to collect them together in one reference book which should be considered as a complement to existing monographs on Markov decision processes.The book is self-contained and unified in presentation.The main theoretical statements and constructions are provided, and particular examples can be read independently of others. Examples in Markov Decision Processes is an essential source of reference for mathematicians and all those who apply the optimal control theory to practical purposes. When studying or using mathematical methods, the researcher must understand what can happen if some of the conditions imposed in rigorous theorems are not satisfied. Many examples confirming the importance of such conditions were published in different journal articles which are often difficult to find. This book brings together examples based upon such sources, along with several new ones. In addition, it indicates the areas where Markov decision processes can be used. Active researchers can refer to this book on applicability of mathematical methods and theorems. It is also suitable reading for graduate and research students where they will better understand the theory.

An integrated work in two volumes, this text teaches readers to formulate, analyze, and evaluate Markov models. The first volume treats the basic process; the second, semi-Markov and decision processes. 1971 edition.

The intent of this book is to present recent results in the control theory for the long run average continuous control problem of piecewise deterministic Markov processes (PDMPs). The book focuses mainly on the long run average cost criteria and extends to the PDMPs some well-known techniques related to discrete-time and continuous-time Markov decision processes, including the so-called ``average inequality approach", ``vanishing discount technique" and ``policy iteration algorithm". We believe that what is unique about our approach is that, by using the special features of the PDMPs, we trace a parallel with the general theory for discrete-time Markov Decision Processes rather than the continuous-time case. The two main reasons for doing that is to use the powerful tools developed in the discrete-time framework and to avoid working with the infinitesimal generator associated to a PDMP, which in most cases has its domain of definition difficult to be characterized. Although the book is mainly intended to be a theoretically oriented text, it also contains some motivational examples. The book is targeted primarily for advanced students and practitioners of control theory. The book will be a valuable source for experts in the field of Markov decision processes. Moreover, the book should be suitable for certain advanced courses or seminars. As background, one needs an acquaintance with the theory of Markov decision processes and some knowledge of stochastic processes and modern analysis.

Technical Report, No.15

Examples In Markov Decision Processes

With Applications to a Class of Optimization Problems in Queuing Theory

Constrained Markov Decision Processes

Markov processes are among the most important stochastic processes for both theory and applications. This book develops the general theory of these processes, and applies this theory to various special examples. The initial chapter is devoted to the most important classical example - one dimensional Brownian motion. This, together with a chapter on continuous time Markov chains, provides the motivation for the general setup based on semigroups and generators. Chapters on stochastic calculus and probabilistic potential theory give an introduction to some of the key areas of application of Brownian motion and its relatives. A chapter on interacting particle systems treats a more recently developed class of Markov processes that have as their origin problems in physics and biology. This is a textbook for a graduate course that can follow one that covers basic probabilistic limit theorems and discrete time processes.

Continuous-Time Markov Decision ProcessesTheory and ApplicationsSpringer Science & Business Media

Finite State Continuous Time Markov Decision Processes with a Finite Planning Horizon

From Filtering to Controlled Sensing