TABLE OF CONTENTS

Introduction

As it is well known, only very few stochastic control problems with partial observation admit an analytic solution; it is therefore of great importance to have efficient numerical methods for their solution. The exact computation of strictly optimal controls may be very difficult, if not impossible, even in discrete time. Our goal with this monograph is therefore to present a detailed description of various possibilities to determine nearly optimal controls for partially observed stochastic control problems in discrete time, where optimality is in the sense of minimizing the expectation of a given cost functional (objective function).

By nearly optimal controls we mean a family of controls with the property that, given any $\varepsilon > 0$, there is a control in the family for which the objective function takes a value that comes within ε of the optimal one. Constructing nearly optimal controls is important for applications : while strictly optimal controls may not always exist and, if they exist, they may be very difficult to determine explicitly, nearly optimal controls exist by the very definition of optimality and, from a practical point of view, they may be just as acceptable as strictly optimal controls. It may also happen that optimal controls have a very complicated structure, difficult to implement in practice, while nearly optimal controls may turn out to be rather simple; in fact, various nearly optimal controls exhibited in this monograph are piecewise constant.

Limiting ourselves to the discrete time case allows us to keep the presentation within a reasonable size and to present the conceptual aspects of the problem without excessive technicalities. On the other hand the discrete time case is also important in practical applications : it includes the partially observed Markov decision problems that have wide applicability, especially in Operations Research and Management Science (see e.g. [25], [40]); it includes also various engineering problems, particularly those when the observations, and therefore also the controls, are taken at discrete time points. Finally, nearly optimal controls for discrete time problems may be used to construct nearly optimal controls for continuous time problems, of which the discrete time problems are an approximation. An approach to this effect can be found in [3] for finite horizon problems and in [29] for infinite horizon problems with discounting. We strictly treat the more comprehensive partially observed case using a setting as general as possible. For this case the associated filtering problem becomes relevant, allowing the partially observed problem to be transformed into an equivalent complete observation problem where the filter becomes the new state (the so-called *separated problem*, see e.g. [2] in the context ot continuous time). This monograph therefore does not only contain results for control problems, but also for controlled filtering problems, in particular on approximations and existence of invariant measures. Although the partially observed case is a generalization of the completely observed one, the methods are rather different so that this monograph has little overlapping with other studies concerning exclusively the completely observed case (for approximation studies see e.g. [21],[22]).

In order to consider a method to be such that it allows the actual construction of nearly optimal controls, we shall require that it reduces the given problem to the solution of an associated problem, where all quantities involved take only a finite number of possible values. Accordingly, in this monograph we shall generally stop the investigations every time we reach a step where the only problem left is that of determining an optimal solution among a finite number of possible ones (an exception is section 4.5.4, where we present a complete computational analysis of an example for the ergodic cost case). This last step problem can in principle always be solved and efficient procedures to this effect are available from the literature, such as various global optimization methods (see e.g. [33],[41]). The major tools used in this monograph to obtain approximating problems with a finite number of possible solutions are either based on discretization-like methods leading to finite-state Markov chains, or on considering approximations of a given trasition kernel by kernels separated in the variables. By the very definition of a family of nearly optimal controls as recalled above, it is clear that the problem of determining such a family is intimately related to approximations of stochastic control with partial observations. In fact, the main tool used in this monograph is the following : approximate the given problem by a family of problems, for each of which an optimal control can actually be computed (in line with the foregoing, each problem in the family should thus admit only a finite number of possible solutions); then show that, if an approximating problem is sufficiently close in a suitable sense to the original one, its optimal control can be extended to become a nearly optimal control for the original problem (given $\varepsilon > 0$, there exists an approximating problem such that the optimal solution for the latter becomes ε-optimal in the original one). The methods we are going to describe are mainly related to work performed by the authors and various coauthors over the past years and for which the results have to a large extent already appeared in the literature (see [3],[10],[30],[31],[36],[37]). In the monograph we present these results in a logical order, we also add various new results derived in order to fill in the gaps still left open, and complement our results with work by other authors. The main material is presented in three stages corresponding to the three chapters from 2 to 4.

The first stage (chapter 2) concerns an approach that appears to be most efficient for the *finite horizon* case and that is based on measure transformation, for which, under a suitable reference probability measure, the observations become i.i.d. random variables, independent of the state process. As a consequence, one can have the same admissible controls in the original and the approximating problems and thus compare the cost functions (expressed as expectations with respect to the same reference probability measure) of the original and the approximating problems corresponding to a same control law, the main objective in this first stage being an approximation that is uniform in the control. This method is thus particularly appropriate for our purposes, but measure transformation requires certain regularity assumptions on the model that may not always be satisfied.

The second stage (chapter 3) is mainly concerned with the *infinite horizon with discounting* case and here we present an approach without the use of measure transformation and thus without the ensuing major benefit that allows to have the same admissible control laws in the original and the approximating problems. As a consequence, we have to proceed along two steps : on a first step, after showing that we may restrict ourselves to Markov controls obtained as functions of the current filter values, we determine nearly optimal control laws (control functions) which, when applied to the true filter values, yield nearly optimal controls. For the actual construction of the nearly optimal control laws we approximate the original problem by simpler ones, for which the associated filter process, based on fictitious discrete-valued observations, takes values in a finite-dimensional space of measures thus making the construction feasible. Extending then suitably these functions to the space of measures where the original filtering process takes its values, we show that they become the desired nearly optimal control functions for the original problem. Since the true (infinite-dimensional) filter values can generally not be computed in practice, on a second step and under additional assumptions guaranteeing the continuity of the nearly optimal control laws, we construct a real finite-dimensional approximating filter process and show that the nearly optimal control laws, when applied to the approximating filter values, still lead to nearly optimal controls. It is worth pointing out that the real approximating filter process can unfortunately not also be used for the construction of the nearly optimal control laws; in fact, a major theoretical problem arising with this real approximating filter is that it is not Markov, a difficulty that we overcome by considering pairs of processes, each pair consisting of the real approximating filter and a "true" filter. In connection with this second stage notice also that, for bounded cost functions, a nearly optimal control law obtained for the finite horizon case can always be extended to become nearly optimal for the infinite horizon case with discounting. On the other hand, the method presented in this second stage and that does not require the regularity assumptions to perform measure transformation, can also be applied to the finite horizon case and it becomes important in the following third stage that concerns the infinite horizon ergodic case. In fact, while under measure transformation one can work with unnnormalized filters, without measure transformation one has to use the more complicated normalized filters that on the other hand however have the advantage of being uniformly bounded measures and this is useful in the ergodic case.

For the third stage (chapter 4) that concerns the *infinite horizon ergodic cost* case (infinite horizon problem with long run average cost criterion), the basic approach follows the lines of that of chapter 3, namely without the use of measure transformation and thus proceeding along two steps. While the second step concerning the construction of a real approximating filter parallels that of chapter 3, for the first step, that concerns the construction of nearly optimal control functions, we present two possible approaches. In the first approach we give various conditions under which, when the discount factor is close to one, a nearly optimal control function computed for the infinite horizon discounted problem is nearly optimal also for the infinite horizon ergodic cost problem (*discounted cost approximation*). The second approach is specific to the ergodic case: using continuous control functions, the filter process itself becomes Markov-Feller, so that under some additional assumptions there exist invariant measures for the filter process. The ergodic cost functional can then be expressed as integral of the cost function with respect to an invariant measure. In order to construct nearly optimal control functions, results on approximation, convergence as well as uniqueness of invariant measures for controlled filtering processes become then of great importance. The general problem of uniqueness is hard. We thus study two situations, where for the invariant measure it is not only possible to show uniqueness, but also to obtain an explicit representation and this is obviously very useful for approximation purposes. The two situations just mentioned are the following: the first one is when the state process is completely observed on a given recurrent subset of the state space and partially outside; the second concerns all those cases when the filter process admits an embedded i.i.d. process. This latter case arises typically when the filter process returns periodically, with bounded average time, to a same measure so that, when costs are bounded, a strong Law of Large Numbers applies.

To convince the reader that our approach is not only theoretical, we intended to present in chapter 4, more precisely in section 4.5.4, a complete computational analysis of a given (nonlinear) problem with an ergodic cost functional. The corresponding programs were worked by dr A. Zemla from IMPAN in Warsaw. The computations were performed along two possible variants: one fully exploiting our approximation approach, the other being partly based on Monte Carlo simulations. In both cases the search for the optimal solution is based on a simulated annealing algorithm. The interested readers may obtain the programs via FTP by contacting the second author for the instructions.

The monograph is intended to be as much as possible self-contained. In chapter 1 we therefore summarize background material concerning properties of the filtering process associated with a partially observed stochastic control problem, as well as basic facts related to measure transformation techniques. Given the generality of our setting, this background material contains partially new results.

Finally, in an Appendix we discuss bibliographical references trying to present in a hopefully comprehensive way the connections between the material presented in the monograph and the existing literature. We apologize for possible omissions.

The monograph may be used as textbook for courses in subjects like Control, Optimization and Applied Stochastic Analysis; furthermore, it may serve not only researchers but also practitioners in areas such as Control Engineering, Operation Research and Management Science, Applied Stastistics and Decision Theory. No particular backgrounds is required, except some basic notions from Probability Theory. Introductory notions from Control (as in [4]) and or Markov Decision Theory (as e.g. [17],[24],[40]) may be useful, but are not required.

This monograph is not only the result of our own work, but also of the interaction with other scientists, mainly those with whom we had a chance to cooperate (see References) and we are grateful to them for all their suggestions and advice. In this context we would also like to mention Professors H. J. Kushner and J. Zabczyk who, although they did not contribute directly to this monograph, had a strong influence on our scientific formation in particular on control and approximations.

Great help was given to us by dr Adam Zemła from IMPAN, who not only wrote the programs, but also tested extensively our methods in particular those for the ergodic control cost problem. The numerical experience that he thus gained allowed him to give us useful advice concerning the methods themselves and to enhance our intuition.

We would probably not have begun working on this monograph, had we not been invited to do so by the Applied Mathematics Committee (CAM) of the Italian National Research Council (CNR) that is editing this series of Applied Mathematics Monographs. Our sincere thanks for the encouragement go therefore to all its members, in particular to Professors G.F. Capriz and F. Giannessi who follow more closely the Monographs series. We particularly appreciate the extremely careful typing of our manuscript by Mrs. Joanna Zemła, secretary in chief at IMPAN and the prompt handling of our monograph by the publisher Giardini from Pisa.

To work on this monograph it was very important to have had the opportunity to meet periodically. Our sincere thanks go therefore also to the host Institutions and funding Agencies and Organizations. The host Institutions were the Institute of Mathematics (IMPAN) of the Polish Academy of Sciences in Warsaw and the Mathematics Department of the University and the Laboratory LADSEB of the Italian National Research Council (CNR) in Padova. The funding for our mutual visits came from GNAFA-CNR through its visiting professors program, from KBN grant 2 2043 91 02, from the project "40% Processi Stocastici e Calcolo Stocastico" of the Italian Ministry for the Universities and Scientific Research, and finally from the exchange program CNR-PAN.

List of basic symbols

$A(E)$: denotes the set of analytic subsets of E;
$\mathcal{B}(E)$: the set of Borel subsets of E;
$b\mathcal{B}(E)$: the set of real valued, bounded Borel functions on E;
$C(E)$: the set of real valued, bounded continuous functions on E;
$P(E)$: the set of probability measures on E, endowed with the weak convergence topology;
$\mathcal{B}(P(E))$: the set of Borel subsets of $P(E)$;
$b\mathcal{B}(P(E))$: the set of real valued, bounded Borel functions on $P(E)$;
$C(P(E))$: the set of real valued, bounded continuous functions on $P(E)$;
$A(P(E),U)$: the set of analytic functions from $P(E)$ in U;
$\mathcal{B}(P(E),U)$: the set of Borel functions from $P(E)$ in U;
$C(P(E),U)$: the set of continuous functions from $P(E)$ in U;
$\nu(f)$ for $\nu \in P(E)$ and $f \in b\mathcal{B}(E)$:

$$\nu(f) = \int_E f(x)\nu(dx)$$

$P(\nu, f)$ for $\nu \in P(E)$ and $f \in b\mathcal{B}(E)$:

$$P(\nu, f) = \int_E \int_E f(z)P(x, dz)\nu(dx)$$

List of assumptions

Assumptions	*Sections*
$(A1)-(A4)$	1.2
$(A5)$	3.2
$(A6),(A7)$	4.2.1
$(A8)-(A10)$	4.3.1
$(A8'),(A9')$	4.3.1
$(A11)$	4.4.3
$(A11')$	4.5.5
$(A12)$	4.5.2
$(B1),(B2)$	2.2.1
$(B1')$	2.3.1
$(B3)-(B5)$	2.2.2
$(B6)-(B8)$	2.3.1
$(B9)$	3.3.2
$(B10)$	$3.3.3.b$
$(B11)$	4.6
$(C1)-(C4)$	2.2.1
$(C5)-C6)$	2.2.2
$(C7)-(C10)$	2.3.2
$(C11)$	3.4.2
$(C12)$	3.5.2
$(C13)$	$4.5.2.b$
$(D1)-(D3)$	3.3.1
$(D4)-(D7)$	4.5.1
$(D8)-(D9)$	$4.5.1.b$
$(E1.1)-(E1.4)$	4.4.2
$(E2.1)-(E2.4)$	4.4.2