Preface

This book deals with numerical methods for control and optimal control problems for nonlinear continuous-time stochastic systems with delays. It is an extension to the model with delays of the Markov chain approximation methods of [58]. For the nondelay problem, these methods are a widely used and powerful class of numerical approximations of optimal costs or other functionals of controlled or uncontrolled stochastic processes in continuous time and have significant applications to deterministic problems. A comprehensive development is in [58].[1]

There are numerous sources of delays in the modeling of realistic physical and biological systems. Many examples arise in communications and queueing, due to the finite speed of signal transmission, the nonnegligible time required to traverse long communications distances, or the time required to go through a queue [90]. Other examples arise because of mechanical transportation delays as, for example in hydraulic control systems, delays due to noninstantaneous human responses or chemical reactions, or delays due to visco-elastic effects in materials. The books [44, 45] contains many concrete examples in mechanics, physics and control, as well as in biology and medecine. These examples are for the most part uncontrolled and deterministic. But many of them would be more realistic if noise were added. Many examples, together with a great deal of information on deterministic delay systems are in [77]. The excellent reference [46] contains a thorough development of the problems of optimal control of deterministic and stochastic delay systems up to its original date of publication (1992), with many examples from biology, mechanics, and elsewhere, as well as a discussion of approximation in policy space algorithms for approximating the optimal cost and control. Other examples can be found in [17, 39, 68, 77, 78]. Examples arise in biological systems due to the time

[1] This book is concerned with optimization and control problems, and with the computation of the expected values of system functionals of interest. Methods for the pathwise numerical solution of the delay equation itself for deterministic and stochastic models are discussed in [2, 6, 32, 47, 48, 69, 81].

delay in the body's adaptive loops, the finite speed of blood flow, or the time required for enzyme or other chemical reactions to occur (see, e.g., [4, Chapter 2]). Models of ecological interactions have been a main source of dynamical models with delays, and applications to financial mathematics are beginning to appear [11]. Very little information is available concerning solutions when the models are nonlinear and stochastic, and numerical methods should be a main source of such information. The reference [75] is concerned with discrete-time approximations to determinstic control problems governed by differential inclusions.

There is a huge literature on control problems for delay systems for the linear model (deterministic or stochastic) with a quadratic cost criterion, and many good computational methods have been developed. Some of the approximations have been done in the spectral domain, based on finite-order rational approximations to the transfer function of the system. Others work with the state-space formulation, where the key issue is the finite-dimensional approximation of the Ricatti equation, which is often done via an approximation to the semigroup of the system. A selection of the available results can be found in [3, 15, 22, 26, 28, 35, 36, 37, 38, 40, 63, 66, 72, 95] and in their references. Although these techniques and algorithms have been very useful for the linear problem, it is not clear (as for the problem without delays), how to adapt them to the nonlinear models that are of concern to us. For this reason, we confine attention to analogs of the approaches that have been found to be very useful for the general no-delay problem, namely the Markov chain approximation method.

The models of the systems of concern in the book are diffusion and reflected diffusion processes, and the results can be extended to cover jump-diffusions. The control might be "ordinary" in the sense that it is a bounded measurable function, or it might be impulsive, or what is known as a "singular" control. All of the usual cost functionals are covered; the discounted cost, stopping on reaching a boundary, optimal stopping, ergodic, etc. Any or all of the path, control, boundary reflection process, or driving Wiener process, might appear in delayed from. Examples where the boundary reflection process might be delayed occur in communications/queueing models, where there is a communications delay. (See Section 1.2 for an example.) If a buffer overflows (corresponding to a lost packet), a signal is sent to the source, which receives it after a delay, and then adjusts its rate of transmission accordingly. The buffer overflow is a component of the boundary reflection process. Models with delays of such boundary reflection terms have not been treated previously.

For the nondelay problem, the approach of the Markov chain approximation method starts by approximating the original controlled process by a controlled Markov chain on a finite state space. The approximation parameter is denoted by h and it might be vector-valued. The original cost functional is also approximated so that it is suitable for the chain. The approximating chain must satisfy a simple condition called "local consistency." This is quite unrestrictive and means simply that from a local point of view and for small

h, the conditional mean and covariance of the changes in state of the chain are proportional to the local mean drift and covariance of the original process, modulo small errors. Many straightforward ways of getting the approximating chains are discussed in [58], where it is seen that the approach is very flexible. The approximation yields a control problem that is close to the original, which gives the method intuitive content that can be exploited for the construction of effective algorithms. After getting the approximating chain, one solves the Bellman equation for the optimal cost (or simply the equation for the value function of interest if there is no control), and proves that the solution converges to the desired optimal cost or value function as h goes to zero. One tries to choose the approximation so that the associated control or optimal control problem can be solved with a reasonable amount of computation and that the approximation errors are acceptable.

The proofs of convergence of the Markov chain approximation method as $h \to 0$ are purely probabilistic. We always work with the processes. No tools from PDE theory or classical numerical analysis are used. The idea behind the proof can be described as follows. For the optimal control problem, starting with the approximating chain with its optimal control, one gets a suitable continuous-time interpolation, and shows that in the sense of weak or distributional convergence, there is a convergent subsequence whose limit is an optimally controlled process of the original diffusion type, and with the original cost function and boundary data. The mathematical basis is the theory of weak convergence of probability measures, and this powerful theory provides a unifying approach for all of the problems of interest. The development in this book depends heavily on the results and methods in [58]. We try to be as self-contained as possible, and do review all of the essential ideas, but it would be beneficial to be familiar with the basic ideas in that source before reading this book.

The probabilistic nature of the methods of process approximation and of the mathematical proofs of convergence allows us to use our physical intuition concerning the original problem in all phases of the development. This gives us great flexibility in the details of the approximation and in the construction of algorithms. These advantages will carry over to the problem with delays. In fact, the probabilistic approach to the approximation and convergence is particularly important when there are delays, since virtually nothing is known about the analytical properties of the associated (infinite-dimensional) Bellman equations for nonlinear problems.

When doing numerical work on general nonlinear systems, it is most convenient if the system is bounded. Many types of systems are *a priori* bounded, owing to the physical constraints on the state variables. For example, systems arising in communications or in approximations to queueing models might be bounded due to the boundedness of the buffers and the possible rates of transmission. Other systems are intrinsically bounded due to saturation effects. Models of many communications and queueing systems involve boundaries on the state space that are reflecting, where the reflection directions are determined by the internal routing of the data in the system [56].

There are two standard ways of bounding a state space if it is not already bounded due to the physical constraints imposed on the model. One might stop the process, with an associated stopping cost, if it attempts to leave a prespecified region. Or one might confine it to a given region via a reflecting boundary (the latter method is common in ergodic cost problems). Both approaches are dealt with. If the boundary is added for numerical purposes, then one might have to experiment with it to assure that it is large enough so that it does not materially affect the quantities of main interest. For simplicity, we confine attention to the diffusion model, with the noise variance not being controlled. The methods can be extended to cover jump diffusions and controlled variance and jumps; One adapts the methods that are used in [58] for such problems analogously to the way that the methods for the covered problems are adapted.

For models without delays, the system state takes values in a subset of some finite-dimensional Euclidean space, and the control is a functional of the current state. For models with delays, the state space must take the path of the delayed quantities (over the delay intervals) into account, and this makes the problem infinite-dimensional. So a major issue in adapting the Markov chain approximation method to models with delays concerns suitable "finite" approximations to the "memory segments" so that a reasonable numerical method can be devised, and much attention is given to this problem.

The methods of approximation that are developed are natural and seem to be quite promising. They deal with issues of approximation that are fundamental. However to date there has been little numerical experience, and considerable further work is required. Yet, judging from the experience with no-delay problems, the methods that are developed are very likely to be the foundation of useful algorithms. There are many additional difficulties to be overcome before effective numerical methods for nonlinear stochastic delay equations become a reality. One is rarely interested in an optimal control. Since the model of interest is often not known precisely, or the implementation of an optimal control might be difficult, what is desired is an understanding of the structure of the control, and how it can be approximated. For the no-delay case in low dimensions this is facilitated by being able to visualize the control via graphical methods. This would be a considerable challenge when there are delays.

Numerical optimization methods are often used as a means of exploring the possible tradeoffs among competing criteria. One solves the optimization problem repeatedly, varying the weights of the various components of interest, to see how a decrease in the value of one component affects the values of the other components, under conditions of optimality, as in [59]. Such information can be invaluable to the system designer, even if optimality is not sought for its own sake.

Next the contents of the various chapters of the book will be described.

Outline of the book. Suppose that the effect of the control action is delayed by an amount $\bar{\theta}$. This can cause serious instabilities. To effectively control in such a case, in determining the current control action one must take into account the control actions that were made in the recent past but whose effects have not yet been seen by the controller, those up to $\bar{\theta}$ units of time back from the present time. Chapter 1 contains some simple examples that dramatically illustrate this point. It also describes the class of examples for which there is a state transformation that reduces the problem to one in finitely many dimensions. The narrowness of this class makes numerical methods all the more important.

Chapter 2 is a summary of the main results that will be needed from the theory of weak convergence of a sequence of random processes, and of the so-called martingale problem for characterizing the limit of a weakly convergent sequence. The theory of weak convergence is an extension to a sequence of random processes of the theory of convergence in probability of a sequence of random variables, and is a fundamental tool for approximation and limit theorems. The primary processes of concern in the proofs of convergence are continuous-time interpolations of the approximating chains, and we will need to show that they have limits that are (in fact, optimal) controlled diffusions. Weak convergence theory, together with the methods of the so-called martingale problem for characterizing the limit procesess as the the desired diffusions, provides the essential tools. With their use, the proofs of convergence are purely probabilistic. For the no-delay case this probabilistic approach to the proofs of convergence of numerical algorithms is the most powerful and flexible. For the delay case, there does not seem to be any alternative since, as noted above, the Bellman equation is infinite-dimensional and virtually nothing is known about it.

Chapter 3 describes the controlled dynamical system models that will be of main interest. The subject of delay equations is vast, whether deterministic, stochastic, or controlled or not; for example, see [27, 39, 44, 45, 51, 57, 68, 73, 74, 77, 78]. The behavior can be quire bizarre, as seen in the examples in [6, 74]. The numerical approximations that are of interest require that the path take values in some compact subset G of a Euclidean space, and this motivates the models. The process can be either stopped on first (if ever) reaching the boundary of G, or else be prevented from leaving G by a boundary reflection process, both being standard models in applications. The stochastic differential equations with path and control delays are reviewed for the cases where the process is either reflected from a boundary or not. Relaxed controls, which are very helpful when dealing with approximation and convergence in control problems, and the Girsanov transformation, an approach to constructing control systems from uncontrolled systems, are discussed. The Girsanov transformation method will be crucial in dealing with the ergodic cost problem in Chapter 5. When the control and possibly the reflection process and/or the driving Wiener process appear in a delayed form, the most direct approaches to the numerical approximation could require an impossibly large memory. One promising way of alleviating this is discussed in Chapter 9 and a dynamical model that is particularly useful for that approach is introduced in Chapter 3.

The existence of an optimal control is also shown. The proof of this fact is important because it is a template for the proofs of convergence of the system and numerical approximations in subsequent chapters. Proofs of the existence of solutions to uncontrolled stochastic delay equations of the diffusion type (without reflecting boundaries) and some of their properties can be found in [39, 68, 73, 74]. For the singular control problem, the definition of the model and the existence of an optimal control are dealt with via a very useful "time transformation" method, which is necessary owing to the possibly wild nature of the associated paths and controls.

Numerical methods involve working with approximations of the original problem whether there are delays or not. The design and success of a numerical approximation is dependent upon the sensitivity of the original model to perturbations in its structure since the numerical algorithm itself is an approximation to the original model. This issue of sensitivity is a particularly acute problem when there are delays owing to the great sensitivity of many such models to parameter variations. See, for example the examples in [74]. One must always be aware of this issue of sensitivity in constructing a numerical approximation. Nevertheless it is important to simplify the original dynamical model as much as possible without sacrificing the essential aspects of the results. Fortunately, for many problems of interest, approximations that are useful for numerical purposes can be obtained.

The key difference between the problem with and without delays is that the state space for the problem with delays involves the "memory segments" of the components whose delayed values appear in the dynamics. The first step in the construction of a numerical approximation involves approximating the original dynamical system. In our case, this entails approximating the delays and dynamics so that the resulting model is simpler, and ultimately finite-dimensional. Chapter 4 is is devoted to a set of model simplifications that have considerable promise when the path or path and/or control are delayed. A variety of approximations are presented, eventually leading to finite-dimensional forms that will be used as the basis of numerical algorithms in Chapters 7–9. To help validate the approximations, simulations that compare the paths of the original and approximated system are presented, and it is seen that the approximations can be quite good.

Delay equations might have rapidly time-varying terms, even rapidly varying delays. This complicates the numerical problem. But, under suitable conditions, there are limit and approximation theorems that allow us to replace the system by a simpler "averaged" one and some such results are presented at the end of Chapter 4.

Chapter 5 is concerned with the average cost per unit time (ergodic cost) problem for nondegenerate reflected diffusion models, where only the path is delayed. The aim is to prepare ourselves for the needs of the numerical algorithms for this case. Hence the issues of model complexity and simplification that were of concern in Chapter 4 are also of concern here. There are only a few results on the ergodic theory for general delay equations. Some, dealing with the problems of existence and convergence of the distributions to invariant measures are [12, 18, 83, 86]. Since they are not quite adequate for the needs of the numerical and approximation problems for the systems of interest, the necessary results are developed, using methods based on the Girsanov transformation and the Doeblin condition, and to the extent possible following the procedures laid out in [56, Chapter 4]. Of particular interest is the demonstration that the various model approximations developed in Chapter 4 can also be used for the ergodic cost problem.

The Markov chain approximation method for the model with no delays is outlined in Chapter 6. We review of the key parts of [58] that will be needed in the sequel. All of the usual process models and cost functions can be handled. For efficiency, the development and analysis of the numerical algorithms in the following chapters is organized to take advantage of the results in [58], wherever possible, and it would be helpful if the reader has some familiarity with that reference. The notation will be slightly different from that in the references [31, 50, 58], since we wish to adapt or simplify it for the particular purposes of this book. The basic and unrestrictive local consistency condition, methods of approximation, continuous time interpolations, and the discounted, singular, impulsive control, and ergodic cost function are covered. The numerical algorithms are based on the finite-state Markov chain approximation. But the convergence proofs are based on continuous-time interpolations of the approximating chains. These interpolations are used for the convergence proofs only and not for the numerical algorithms.

Owing to the local consistency condition, the dynamical system that is represented by a continuous-time interpolation of the chain "resembles" the original controlled diffusion process. Thus we would expect that the optimal cost or the values of the functionals of interest would be close to those for the diffusion. This is quantified by the convergence theorems. There are two (asymptotically equivalent) methods of getting the approximating chains that are of interest, called the "explicit" and "implicit" methods. They differ in the way that the time variable is treated, and each can be obtained from the other. The first method was the basic approach for the nondelay problem. The second method will play a useful role in reducing the memory requirements when there are delays.

The adaptation of the methods of the Markov chain approximation method to the models with delays begins in Chapter 7 and is continued in Chapter 8. It is shown in Chapter 7 that any method of constructing the approximating chain for the no-delay problem can be readily adapted to the delay problem, with the transition probabilities taking the delays into account. The only change in the local consistency condition is the use of the "memory segment" arguments in the drift and diffusion functions.

The algorithms in Chapters 7–9 are well motivated and seem to be quite reasonable. But since the subject is in its infancy, what is presented should be taken as a first step, and will hopefully motivate further work. When constructing a numerical approximation algorithm, there are two main issues that must be kept in mind. The algorithm must be numerically feasible and it must be such that there is a proof of convergence as the approximating parameter goes to zero. These issues inform the structure of the development.

We start the development in Chapter 7 by working with numerical approximations to the original model. Then we turn our attention to the various approximations to the original model that were developed in Chapter 4, with an eye to the feasibility of their numerical approximations, taking the two main issues cited above into consideration. It will be seen that variations of the implicit approximation method of Chapter 6 can be advantageous in dealing with the memory problem. The continuous-time interpolations that are used for the convergence proofs are somewhat more complicated that those for the no-delay case, owing to the need to represent the "memory segment" argument in a way that is convenient for use in the proofs of convergence.

The development is continued in Chapter 8, where classes of numerical approximations that we call the periodic and periodic-Erlang are given. The chapter also contains the proofs of convergence for the algorithms in both chapters. Where possible, the proofs follow the general lines that were used for the no-delay case in [58]. As noted above, one interpolates the chain to a continuous-time process in a suitable manner, shows that the Bellman equation for the interpolation is the same as for the chain, and then that the interpolated processes converge to an optimal diffusion as the approximating parameter goes to zero.

The methods of Chapters 7 and 8 are promising if only the path is delayed or if the control is delayed but the control-value space has only a few points. The memory requirements can become onerous if the reflection process and/or the Wiener process also appear in delayed form, or if the control-value space has more than a few points. Chapter 9 takes an alternative approach that reduces the memory requirements for general nonlinear stochastic problems where the control and reflection terms, as well as the path variables, are delayed. The approach was suggested by the work in [94] for linear deterministic system with a quadratic cost function. Effectively, the delay equation is replaced by a type of stochastic wave equation with no delays, and its numerical solution yields the optimal costs and controls for the original model. The representation is equivalent to the original problem in that any solution to one yields a solution to the other. The details of the appropriate Markov chain approximation are given and the convergence theorem is proved. Theoretically, with the use of appropriate numerical approximations, the dimension of the required memory vector is much reduced, although there is little practical numerical experience as yet.

Because of the large sizes of the state spaces that arise in the numerical approximations, it would seem that the topic is well suited for one of the various approaches that are known as approximate dynamic programming, or even linear programming with suitably sampled constraints. See., e.g., the references in [30, 87, 82]. The success of such approaches usually depends on detailed insight into the "physics" of the problem, so that the approximation can be tailored apprpriately. At this time, it is not at all clear how to use such approaches for our problem, but one must always seek approaches that simplify the problem while yielding meaningful results.

Numbering and cross referencing. Cross reference numbering *within* a chapter does not include the chapter number. For example, within Chapter 5, Equation 4 of Section 3 of Chapter 5 is called Equation (3.4), and Subsection 6 of Section 3 of Chapter 5 is called Subsection 3.6, with the analogous usage for Theorem, Figure, and Assumption. Cross references *between* chapters do include the chapter number. For example, in Chapter 5, a reference to Equation 4 of Section 3 of Chapter 2 is called Equation (2.3.4), and Subsection 6 of Section 3 of Chapter 2 is called Subsection 2.3.6, with the analogous usage for Theorem, Figure, and Assumption.

A glossary of the more frequently used symbols appears at the end of the book.

The author gratefully acknowledges the long-term support of the Army Research Office and the National Science Foundation on numerical methods in stochastic control.

Providence, Rhode Island, USA Harold J. Kushner
February 2008

http://www.springer.com/978-0-8176-4534-2

Numerical Methods for Controlled Stochastic Delay Systems
Kushner, H.
2008, XX, 282 p. 37 illus., Hardcover
ISBN: 978-0-8176-4534-2
A product of Birkhäuser Basel