Automatica 50 (2014) 303–334

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Survey paper

A partial history of the early development of continuous-time nonlinear stochastic systems theory ☆

Harold J. Kushner

Applied Mathematics Department, Brown University, Providence, RI 02912, USA

ARTICLE INFO

Article history:
Received 25 July 2013
Received in revised form
19 October 2013
Accepted 25 October 2013
Available online 27 December 2013

Keywords:
Stochastic optimal control
Viscosity solutions
Dynamic programming principle
Nisio semigroup
Existence of optimal controls
Existence under partial observations
Compactness methods
Wide bandwidth noise
Approximations to systems
Nonlinear filtering
Stochastic stability
Stochastic maximum principle
Ergodic control
Stochastic stability
Filtering and control with wide bandwidth noise

ABSTRACT

This article is a survey of the early development of selected areas in nonlinear continuous-time stochastic control. Key developments in optimal control and the dynamic programming principle, existence of optimal controls under complete and partial observations, nonlinear filtering, stochastic stability, the stochastic maximum principle and ergodic control are discussed. Issues concerning wide bandwidth noise for stability, modeling, filtering and ergodic control are dealt with. The focus is on the earlier work, but many important topics are omitted for lack of space.

1. Introduction

The late 1950s throughout the mid 1970s were a period of renaissance in control theory. The classical theory, largely based on linear systems, Fourier and Laplace transform methods, and stability based on Bode and Nyquist plots, and Routh–Hurwitz criteria, was very successful, and provided the foundations for the great successes in control during WWII. They continued to be important and would be significantly developed later, with or without conjunction with state space methods, particularly for multi-input–multi-output linear systems, as seen in the references on linear systems in Levine (1995). But it was difficult to adapt them to the multi-input–output linear and nonlinear systems in and outside of the traditional control areas that were of increasing interest. For example, applications arising in trajectory optimization, system identification and adaptive control, stochastic networks, applications in operations research, etc. Techniques based on differential equations models had been used earlier as in Minorsky (1922), but fell out of favor due to their difficulty and the intuitive nature and successes of the frequency domain methods. In the stochastic realm, the main methods were based on adaptations of Wiener filtering theory and statistical detection methods. The Wiener theory was based on stationary processes, while most applications had time-varying data and were of interest over a bounded time interval. It was well known that the only processes for which the Wiener–Hopf equation could be solved were Gaussian processes with a finite-order rational spectral density, and that in the Gaussian case such processes could be represented in terms of a finite-order Gauss–Markov process, but limited use was made of state-space representations.

☆ Partially supported by ARO contract FA9550-09-1-0378. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form under the direction of Editor John Baillieul.

E-mail address: hjk@dam.brown.edu.

0005-1098/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.automatica.2013.10.013

Three of the fundamental developments[1] that changed the course of control theory in the middle of the 20th century were the maximum principle of[2] Pontryagin, Boltyanski, Gamkrelidze, and Mischenko (1962), giving necessary conditions for optimality and leading to numerical methods; the dynamic programming (DP) of Bellman (1952, 1954, 1957), giving sufficient conditions for optimality; and the work of Kalman (1960a,b) and Kalman and Bucy (1961), solving fundamental problems in the control, characterization, and filtering of linear (Gauss–Markov) systems, emphasizing the importance of controllability and observability, essentially introducing modern linear control systems theory. The excitement created by these works cannot be overstated. One can often find prior works that in some sense "anticipate" major future developments. The methods used by Pontryagin and colleagues were influenced by earlier work in the calculus of variations such as McShane (1939), particularly his use of convex cones for the admissible perturbations of the terminal state constraint. Dynamic programming had much in common with Caratheodory's work, as noted in Kalman (1960a) and Pesch and Bulirsch (1994). But the timing and elegance, as well as the wealth of immediate applications, were crucial to their influence.

In the late 1950s, in a series of RAND Corporation reports, Peter Swerling (a pioneer in radar signal analysis) did develop a recursive form of the filter for the case where the model was a deterministic linear system, and the observations were linear in the state and white noise corrupted. The aim was to get a minimum mean square estimator. These works partly anticipated those of Kalman on filtering. Key aspects of Kalman's work were the use of the concepts of conditional distribution and of observability and controllability in the analysis of the asymptotic properties of the Riccati equation, and establishing the connection with optimal control under a quadratic cost criterion. These developments helped move the emphasis to optimization and filtering problems for dynamical systems, as opposed to the stationarity and Fourier/Laplace transform methods, and they became the dominant approaches. One could say that the early successes of the maximum principle, the theory of linear filtering and control, and dynamic programming, gave control theory powerful reasons for its existence and for the strong support that it ultimately received. Of course, the effects of the Soviet success with Sputnik, the consequent US investment in the space program, and the efforts of people such as Solomon Lefschetz in convincing the funding agencies of the importance of dynamical systems, cannot be ignored. At that time Itô equation models for stochastic problems were barely heard of. Stochastic control was a completely new area. One could also mention Norbert Wiener's Cybernetics. This had little direct influence on control theory, but was important in that it impressed a wider public with the role of feedback, particularly in biological systems.

This survey concerns a few of the early developments in stochastic control theory. The field is vast and the number of contributors large, and we will focus on some of the developments in a few of the various areas with which the author has been involved. The main focus will be on optimal control and numerical methods for controlled diffusion processes (Section 2), existence of optimal controls (Section 4), issues concerning white vs. wide-bandwidth noise, and the related diffusion and cost/control approximations (Sections 5, 6, 7, 9), nonlinear filtering (Section 6), stochastic stability (Section 7), the stochastic maximum principle (Section 8), and control with ergodic cost criteria (Section 9), a few of the concerns among those who were trying to put the subject on solid ground.

This list omits many important topics which had their roots and key developments about the same era, and were critical in the future development of the subject. They include, but are not limited to, the following list, which in itself gives an idea of the current great breadth and broad view of stochastic control. Optimal stopping, originating in the sequential decision analysis of Wald (1950), was perhaps the first stochastic control problem involving dynamic models, later fully developed by Shiryaev (1978) via probability methods, and in Bensoussan and Lions (1978) and Bensoussan (1982b) from the point of view of variational inequalities. Other omitted important topics include linear filtering and control, Markov chain methods (Puterman, 1994), risk-sensitive and robust control (Basar & Bernhard, 1991; Whittle, 1990), singular control (Fleming & Soner, 2006, Chapter 8; Shreve, 1986), multiscale and singularly perturbed problems (Bensoussan, 1988; Borkar & Gaitsgory, 2007; Kushner, 1990), stochastic approximation (Benveniste, Metivier, and Priouret, 1990; Kushner & Clark, 1978; Kushner & Yin, 2003; Yin, 2002), stochastic games (Shapley, 1953), adaptive control and identification, large deviations methods, control of systems with delays, algebraic-geometric methods, control of stochastic partial differential equations, value function approximation methods, and others.

The emphasis is on some of the earlier work, although to complete a line of thought mention might be made to later works. For simplicity in the presentation, our assumptions are usually more restrictive than those in the references. We can only mention a few of the many authors and contributions, and we apologize to those whose work is not adequately reported on, but hopefully the enthusiasm and explosion of activities in many directions will be clear. Keep in mind that this article is not in any way a survey of the latest results nor an encyclopedia.

Since our focus is on some largely theoretical topics, it might give a skewed view of the broad range of what was being done. We need to emphasize that there was also an enormous activity in other areas that are critical for applications, such as the early and influential work of Aström (1970) and the work reported in Ljung and Söderström (1983) which showed the possibilities for adaptive control and systems modeling and identification, and in areas such as Monte Carlo optimization, neural networks and pattern recognition that also relate to systems theory.

Throughout the article, the state space will be the Euclidean r-dimensional space $\mathbb{R}^r$.

2. Stochastic optimal control

The Bellman equation and the principle of optimality for the optimal value function for diffusion-type processes are fundamental tools for characterizing the value and for analysis and approximations. Yet there are subtleties and their validity is not a-priori obvious. A major concern was to put them on solid ground. Some of the main approaches and issues will be covered in this section. To set the framework, we start with the familiar form for discrete-time Markov chains, and then move on to diffusions, both nondegenerate and degenerate. The principle of optimality will always be referred to as the DP (dynamic programming) principle. Fleming and Soner (2006) contains many historical remarks on the subjects of this section.

2.1. Dynamic programming (DP)

A Markov chain model. Consider a controlled Markov chain $\{X_n\}$ on a state space B, with controls $\{u_n\}$ taking values in a set U, and a one-step transition function $P(z|x, \alpha)$, where α is the current control value, and a cost function[3] $W_n(x, u^n) = E_x^{u^n} \sum_{i=n}^{N} k(X_i, u_i)$,

[1] The brief survey of Marcus (1994) gives a good overview of the dynamism of that period.

[2] Often called the minimum principle, since it is frequently used in minimization problems.

[3] $E_x^{u^n}$ is the expectation under the probability defined by initial state x and that control sequence u^n is to be used.