Automatica 70 (2016) 74–85

Contents lists available at ScienceDirect

Automatica

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

Numerical methods for optimal harvesting strategies in random environments under partial observations ☆

Ky Tran, G. Yin

Department of Mathematics, Wayne State University, Detroit, MI 48202, United States

ARTICLE INFO

Article history:
Received 13 July 2015
Received in revised form
30 December 2015
Accepted 8 March 2016
Available online 15 April 2016

Keywords:
Harvesting problem
Controlled regime-switching diffusion
Singular control
Partial observation
Wonham filter

ABSTRACT

This work is concerned with optimal harvesting problems in random environments. In contrast to the existing literature, the Markov chain is hidden and can only be observed in a Gaussian white noise in our work. We first use the Wonham filter to estimate the Markov chain from the observable evolution of the given process so as to convert the original problem to a completely observable one. Then we treat the resulting optimal control problem. Because the problem is virtually impossible to solve in closed form, our main effort is devoted to developing numerical approximation algorithms. To approximate the value function and optimal strategies, Markov chain approximation methods are used to construct a discrete-time controlled Markov chain. Convergence of the algorithm is proved by weak convergence method and suitable scaling. A numerical example is provided to demonstrate the results.

1. Introduction

This work focuses on optimal harvesting problems for ecosystems formulated by stochastic differential equations with regime switching represented by a continuous-time Markov chain. The problem belongs to the class of singular stochastic control problems motivated by the establishment of ecologically, environmentally, and economically reasonable wildlife management and harvesting policies. Recently, there has been a resurgent interest in determining the optimal harvesting strategies in the presence of stochastic fluctuations. Radner and Shepp (1996) derived the optimal strategy of a model for corporate strategy. Alvarez and Shepp (1998) studied the optimal harvesting plan for the stochastic Verhulst–Pearl logistic model. All the aforementioned works dealt with species living in an environment with a fixed configuration. Recently, Song, Stockbridge, and Zhu (2011) and Tran and Yin (2015) considered singular control problems in random environments modeled by a Markov chain, in which Song et al. (2011) dealt with a single species and Tran and Yin (2015) treated multiple species with interactions.

Suppose that there is a single species $X(t)$ whose growth is subject to the usual fluctuations as well as the abrupt changes of a random environment. Harvesting strategies are introduced to derive financial benefit as well as to control the growth of the population. Let $Z(t)$ denote the total amount harvested from the species up to time t. The goal is to find a harvesting strategy $Z(t)$ that maximizes the expected total discounted income from harvesting, up to the time when the population falls to a given threshold (e.g., extinction), which has the following economic interpretation. Let $X(t)$ be the value at time t of asset/security/investment and $Z(t)$ represent the total amount paid in dividends up to time t. Then $\mathbb{R}_+ = (0, \infty)$ can be regarded as the solvency set, and (13) becomes the problem of finding the optimal stream of dividends from the collection of assets until the time of bankruptcy; see Asmussen and Taksar (1997), Sotomayor and Cadenillas (2011), and Choulli, Taksar, and Zhou (2003).

Harvesting may occur instantaneously, so it results in a singular stochastic control problem in the sense that the optimal harvesting strategy $Z(t)$ may not be absolutely continuous with respect to the Lebesgue measure of the time variable. For instance, if the discounted value and noise intensity are sufficiently large, driving the population to extinction instantly or chattering harvesting strategies might be optimal or near-optimal; see Alvarez and Shepp (1998), and Tran and Yin (2015). In contrast to regular stochastic control problems, in which the displacement of the state due to control is differentiable in time, the harvesting problem considered in this work allows the displacement to be discontinuous. To find the value function and the harvesting strategy, one usually solves a so-called Hamilton–Jacobi–Bellman (HJB) equation. However, for singular control problems with

☆ This research was supported in part by the Air Force Office of Scientific Research under FA9550-15-1-0131. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Valery Ugrinovskii under the direction of Editor Ian R. Petersen.

E-mail addresses: kytran@wayne.edu (K. Tran), gyin@math.wayne.edu (G. Yin).

http://dx.doi.org/10.1016/j.automatica.2016.03.025
0005-1098/

regime switching, the HJB equation is in fact a coupled system of nonlinear quasi-variational inequalities. A closed-form solution is virtually impossible to obtain. The Markov chain approximation methodology developed by Kushner and Dupuis (1992) becomes a viable alternative. As pointed out in Kushner and Dupuis (1992), probabilistic approach using the Markov chain approximation method for controlled diffusions has the following advantages. First, the Markov chain approximation method allows one to use physical insights derived from the dynamics of the controlled diffusion in obtaining a suitable approximation scheme. Second, the Markov chain approximation method does not require much regularity of the controlled processes (solutions of the controlled stochastic differential equations) nor does it rely on the properties of the associated HJB equations. Though it is well recognized the need of developing numerical approximation methods for singular control problems, the results are still scare. For singular controlled diffusions without regime switching, Budhiraja and Ross (2007), and Kushner and Martins (1991) are two of the representative works that carry out a convergence analysis using weak convergence and relaxed control formulation for singular control problems in the setting of Itô diffusions. Recently, some works have been devoted to numerical methods for singular controls with regime switching. Jin, Yin, and Zhu (2012) developed numerical algorithms for finding optimal dividend payout and reinsurance policies under a generalized singular control formulation. A numerical algorithm for optimal dividend payment and investment strategies of regime-switching jump diffusion models with capital injections was then introduced in Jin, Yang, and Yin (2013).

In our work, we focus on the harvesting problem for a partially observed system with a hidden Markov chain. So far, the work on numerical solutions has mostly concentrated on the case the Markov chain is observable. In reality, the environment (Markov chain) can often be only observed with noise. That is, at any given instance, the exact state of residency of the Markov chain is not known. Thus, we cannot see $\alpha(t)$ directly but only have noise-corrupted observation in the form of $\alpha(t)$ plus noise. An effective way to handle control problems of such partially observed systems is to convert them to completely observed ones, which can be done by using a Wonham filter (see, for example, Wonham, 1964). In the literature, the Wonham filters have been used widely to investigate control problems with partial observations; see Tran and Yin (2014) and Yang, Yin, and Zhang (2015) for applications in engineering, finance, and ecology. Note that our work is different from some problems in math finance, in which the $X(t)$ is taken to be the observation. Here we have another observation process in the form of the Markov chain observed in white noise. Such formulations appear in many networked control problems, ecological systems, and cyber–physical systems. Compared to the aforementioned works on numerical methods for singular control problems, in the current work, we take a step towards more useful and realistic model where the Markov chain is unobservable. Although main ideas developed are crucial to the analysis of the current paper, there are key differences in the model that make our analysis more delicate. Using a Wonham filter, we convert the partially observed system into a fully observed controlled diffusion. We then design approximation procedures for the optimal strategies and the value function. We need to use a couple of step sizes $h = (h_1, h_2)$. The parameter $h_1 > 0$ is a discretization parameter for state variables, and $h_2 > 0$ is the step size for time variable. In the actual computing, the computations are involved due to the presence of the Wonham filter.

In contrast to the existing results, our new contributions in this paper are as follows. (i) We use Wonham's filter to formulate the harvesting problem in random environments when the Markov chain is only observable in white Gaussian noise. (ii) We convert the partially observed system to a fully observed system by replacing the unknown Markovian states by their posterior probability estimates. (iii) We develop numerical approximation schemes based on the Markov chain approximation method. Although Markov chain approximation techniques have been used extensively in various control problems, the work on combination of such method for a singular control problem with partial observation seems to be scarce to the best of our knowledge.

The rest of the paper is organized as follows. Section 2 begins with the problem formulation. Section 3 presents the numerical algorithm based on the Markov chain approximation method. Section 4 establishes the convergence of the algorithm. Finally, the paper is concluded with a numerical example for illustration together with further remarks in the last section.

2. Formulation

For $i = 1, \dots, r$, let $X^i(t)$ be the population size of the ith species in the ecosystem at time t and denote $X(t) = \big(X^1(t), \dots, X^r(t)\big)' \in \mathbb{R}^r$ (with z' denoting the transpose of $z \in \mathbb{R}^{r_1 \times r_2}$ with $r_1, r_2 \geq 1$). Suppose that species $X^i(t)$ live in random environments. In addition to the random fluctuations of the population, we also assume that the growth of the species is subject to abrupt changes within a finite number of configurations of the environment. For simplicity, we assume that the switching among different environments is memoryless and that the waiting time for the next switch is exponentially distributed. In fact, this phenomenon is frequently observed in nature; see Slatkin (1978) and Yin and Zhu (2010). Thus we can model the random environments and other random factors in the ecological system by a continuous-time Markov chain $\alpha(t)$ taking values in $\mathcal{M} = \{1, 2, \dots, m\}$ with the generator given by $Q = (q^{ij}) \in \mathbb{R}^{m\times m}$. Assume throughout the paper that both the Markov chain $\alpha(t)$ and the r-dimensional standard Wiener process $w(\cdot) = \big(w^1(\cdot), \dots, w^r(\cdot)\big)'$ are defined on a complete filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}(t), P)$, where $\{\mathcal{F}(t)\}$ is a filtration satisfying the usual condition (i.e., right continuous, increasing, and $\mathcal{F}(0)$ containing all the null sets).

In an effort to capture the salient feature that continuous dynamics and discrete events coexist in the ecosystem, we model the evolution in the absence of harvesting by the stochastic differential equation

$$\begin{aligned} &dX(t) = b(X(t), \alpha(t))dt + \sigma(X(t), \alpha(t))dw(t), \\ &X(0) = x_0 \in \mathbb{R}^r_+, \quad \alpha(0) = \alpha_0 \in \mathcal{M}, \end{aligned} \tag{1}$$

where $b(\cdot) : \mathbb{R}^r \times \mathcal{M} \mapsto \mathbb{R}^r$, $\sigma(\cdot) : \mathbb{R}^r \times \mathcal{M} \mapsto \mathbb{R}^{r\times r}$ are suitable functions. Furthermore, we assume that the Brownian motion $w(\cdot)$ and the Markov chain $\alpha(\cdot)$ are independent, a commonly used assumption in the literature. We attempt to answer the question: Can we solve optimal harvesting problems if the Markov chain is hidden and we can only treat a partially observed system? In particular, we cannot see $\alpha(t)$ directly but only have noise-corrupted observation in the form of $\alpha(t)$ plus noise. That is, we can observe the following process

$$dy(t) = g(\alpha(t))dt + \sigma_0 dB(t), \quad y(0) = 0, \tag{2}$$

where σ_0 is a positive constant, $g : \mathcal{M} \mapsto \mathbb{R}$ is a one-to-one function, $B(t)$ is a one-dimensional standard Brownian motion being independent of $w(t)$.

To proceed, we denote by $\mathbb{1}_E$ the indicator function of the event E, and use the following notation. For $j = 1, \dots, m$,

$$\begin{aligned} &p^j(t) := \mathbb{1}_{\{\alpha(t)=j\}}, \\ &\varphi^j(t) := P\left(\alpha(t) = j | y(s), 0 \leq s \leq t\right). \end{aligned} \tag{3}$$

Since $\varphi^j(t)$ is the probability vector conditioned on the observation $\sigma\{y(s), 0 \leq s \leq t\}$, $\varphi^j(t) \geq 0$ and $\sum_{j=1}^m \varphi^j(t) = 1$.