Analysis of convergence and derivation of a rate of convergence for optimal stopping

ANR-09-SEGI-004

Final report of team CQFD on task 3.1.1

AGENCE NATIONALE DE LA RECHERCHE ANR

INRIA centre de recherche BORDEAUX – SUD-OUEST

The aim of this task is to propose a computational method for optimal stopping of a piecewise deterministic Markov process, analyze the convergence of our scheme and derive bounds for the convergence rate. The partners involved are INRIA CQFD and Astrium. People involved are B. de Saporta (CQFD), F. Dufour (CQFD) and C. Elegbede (Astrium), with the collaboration of K. Gonzalez (CQFD) and H. Zhang (CQFD).

1 Context

Piecewise-deterministic Markov processes (PDMP's) have been introduced in the literature by M.H.A. Davis [11] as a general class of stochastic models. PDMP's are a family of Markov processes involving deterministic motion punctuated by random jumps. The motion of the PDMP $\{X(t)\}$ depends on three local characteristics, namely the flow ϕ, the jump rate λ and the transition measure Q, which specifies the post-jump location. Starting from x the motion of the process follows the flow $\phi(x,t)$ until the first jump time T_1 which occurs either spontaneously in a Poisson-like fashion with rate $\lambda(\phi(x,t))$ or when the flow $\phi(x,t)$ hits the boundary of the state-space. In either case the location of the process at the jump time T_1: $X(T_1) = Z_1$ is selected by the transition measure $Q(\phi(x,T_1),\cdot)$. Starting from Z_1, we now select the next interjump time $T_2 - T_1$ and postjump location $X(T_2) = Z_2$. This gives a piecewise deterministic trajectory for $\{X(t)\}$ with jump times $\{T_k\}$ and post jump locations $\{Z_k\}$ which follows the flow ϕ between two jumps. A suitable choice of the state space and the local characteristics ϕ, λ, and Q provides stochastic models covering a great number of problems of operations research [11].

The optimal stopping problem consists in finding the best time to stop the process in order to optimize the expectation of a reward function g of the process at that time. The best possible performance is called the value function of the problem, and a stopping time reaching this optimal performance is called an optimal stopping time.

Optimal stopping problems have been studied for PDMP's in [8, 10, 11, 13, 15, 17]. In [15] the author defines an operator related to the first jump time of the process, and shows that the value function of the optimal stopping problem is a fixed point for this operator. The basic assumption in this case is that the final cost function is continuous along trajectories, and it is shown that the value function will also have this property. In [13, 17] the authors adopt some stronger continuity assumptions and boundary conditions to show that the value function of the optimal stopping problem satisfies some variational inequalities, related to integro-differential equations. In [11], M.H.A. Davis assumes that the value function is bounded and locally Lipschitz along trajectories to show that the variational inequalities are necessary and sufficient to characterize the value function of the optimal stopping problem. In [10], the authors weakened the continuity assumptions of [11, 13, 17]. A paper related to our work is [8] by O.L.V. Costa and M.H.A. Davis. It is the only one presenting a computational technique for solving the optimal stopping problem for a PDMP based on a discretization of the state space similar to the one proposed by H. J. Kushner in [16]. In particular, the authors in [8] derive a convergence result for the approximation scheme but no estimation of the rate of convergence is derived.

2 Approach

Our approach is based on quantization. Quantization methods have been developed recently in numerical probability, nonlinear filtering or optimal stochastic control with applications in finance [6, 7, 18, 19, 20, 21]. Roughly speaking, the approach developed in [6, 7, 21] for studying the optimal stopping problem for a continuous-time diffusion process $\{Y(t)\}$ is based on a time-discretization scheme to obtain a discrete-time Markov chain $\{\overline{Y}_k\}$. It is shown that the original continuous-time optimization problem can be converted to an auxiliary optimal stopping problem associated with the discrete-time Markov chain $\{\overline{Y}_k\}$. Under some suitable assumptions, a rate of converge of the auxiliary value function to the original one can be derived. Then, in order to address the optimal stopping problem of the discrete-time Markov chain, a twofold computational method is proposed. The first step consists in approximating the Markov chain by a quantized process. There exists an extensive literature on quantization methods for random variables and processes, see [14, 18, 21] and references therein. The second step is to approximate the conditional expectations which are used to compute the backward dynamic programming formula by the conditional expectation related to the quantized process. This procedure leads to a tractable formula called a *quantization tree algorithm* (see Proposition 4 in [6] or section 4.1 in [21]). Providing the cost function and the Markov kernel are Lipschitz, some bounds and rates of convergence are obtained (see for example section 2.2.2 in [6]).

As regards PDMP's, it was shown in [15] that the value function of the optimal stopping problem can be calculated by iterating a functional operator, labeled L, which is related to a continuous-time maximization and a discrete-time dynamic programming formula. Thus, in order to approximate the value function of the optimal stopping problem of a PDMP $\{X(t)\}$, a natural approach would have been to follow the same lines as in [6, 7, 21]. However their method cannot be directly applied to our problem for two main reasons related to the specificities of PDMP's.

First, PDMP's are in essence discontinuous at random times. Therefore, as pointed out in [15], it will be problematic to convert the original optimization problem into an optimal stopping problem associated to a time discretization of $\{X(t)\}$ with nice convergence properties. In particular, it appears ill-advised to propose as in [6] a fixed-step time-discretization scheme $\{X(k\Delta)\}$ of the original process $\{X(t)\}$. Besides, another important intricacy concerns the transition semigroup $\{P_t\}_{t\in\mathbb{R}_+}$ of $\{X(t)\}$. On the one hand, it cannot be explicitly calculated from the local characteristics (ϕ, λ, Q) of the PDMP (see [9, 12]). Consequently, it will be complicated to express the Markov kernel P_Δ associated to the Markov chain $\{X(k\Delta)\}$. On the other hand, the Markov chain $\{X(k\Delta)\}$ is in general not even a Feller chain (see [11, pages 76-77]), therefore it will be hard to ensure it is K-Lipschitz (see Definition 1 in [6]).

Second, the other main difference stems from the fact that the function appearing in the backward dynamic programming formula associated to L and the reward function g is not continuous even if some strong regularity assumptions are made on g. Consequently, the approach developed in [6, 7, 21] has to be refined since it can only handle conditional expectations of Lipschitz-continuous functions.

3 Results

By using the special structure of PDMP's, we are able to overcome both the above-mentionned obstacles. Indeed, associated to the PDMP $\{X(t)\}$, there exists a natural embedded discrete-time Markov chain $\{\Theta_k\}$ with $\Theta_k = (Z_k, S_k)$ where S_k is given by the inter-arrival time $T_k - T_{k-1}$ and Z_k the post-jump locations. The main operator L can be expressed using the chain $\{\Theta_k\}$ and a continuous-time maximization. We first convert the continuous-time maximization of operator L into a discrete-time maximization by using a path-dependent time-discretization scheme. This enables us to approximate the value function by the solution of a backward dynamic programming equation in discrete-time involving conditional expectation of the Markov chain $\{\Theta_k\}$. Then, a natural approximation of this optimization problem is obtained by replacing $\{\Theta_k\}$ by its quantized approximation. It must be pointed out that this optimization problem is related to the calculation of conditional expectations of indicator functions of the Markov chain $\{\Theta_k\}$. As said above, it is not straightforward to obtain convergence results as in [6, 7, 21]. We deal successfully with indicator functions by showing that the event on which the discontinuity actually occurs is of small enough probability. This enables us to provide rate of convergence for the approximation scheme.

In addition and more importantly, this numerical approximation scheme enables us to propose a computable stopping rule which also is an ϵ-optimal stopping time of the original stopping problem. Indeed, for any $\epsilon > 0$ one can construct a stopping time, labeled τ, such that

$$V(x) - \epsilon \leq \mathbf{E}_x\Big[g(X(\tau))\Big] \leq V(x)$$

where $V(x)$ is the value function associated to the original stopping problem when the process starts from point x. Our computational approach is attractive in the sense that it does not require any additional calculations. Moreover, we can characterize how far it is from optimal in terms of the value function. In [6, section 2.2.3, Proposition 6], another criteria for the approximation of the optimal stopping time has been proposed. In the context of PDMP's, it must be noticed that an optimal stopping time does not generally exists as shown in [15, section 2].

4 Applications

The numerical procedure described above has been first tested on a simple academic model of PDMP. Then, we have successfully implemented it to solve a maintenance problem submitted by Astrium. We consider a metallic structure subject to corrosion and want to determine the best time to intervene before the failure.

5 Dissemination of results

The theoretical part of this work with rigorous proofs and the academic example is published in an international peer-reviewed journal *Annals of Applied Probability* [2]. The detailed study of the corrosion model proposed by Astrium is given in the proceeding of *Lambda-Mu 17* conference, a French peer-reviewed conference on reliability and safety, and co-signed by CQFD and Astrium [3]. It is also submitted for publication in an international peer-reviewed journal [4]. B. de Saporta and F. Dufour were invited to present these results at the workshop *Modern trends in controlled stochastic processes* in Liverpool [1]. F. Dufour was invited to present these results at SPA2010 conference in Osaka [5].

References

Publications of the Fautocoes team

[1] de Saporta, B., Dufour, F. Numerical method for optimal stopping of piecewise deterministic Markov processes. In *Modern trends in controlled stochastic processes*, A. Piunovskiy, Ed., Luniver press, 2010.

[2] de Saporta, B., Dufour, F. and Gonzalez, K. Numerical method for optimal stopping of piecewise deterministic Markov processes. *Annals of Applied Probability 20*, (2010), 1607–1637.

[3] de Saporta, B., Dufour, F., Zhang, H. and Elegbede, C. Arrêt optimal pour la maintenance prédictive In *Lambda-mu 17* (La Rochelle, 2010).

[4] de Saporta, B., Dufour, F., Zhang, H. and Elegbede, C. Optimal stopping of piecewise deterministic Markov processes. *Submitted in january 2011.*

[5] Dufour, F. Arrêt optimal pour la maintenance prédictive In *Stochastic Processes and their Applications* (Osaka, 2010).

References in the text

[6] Bally, V., and Pagès, G. A quantization algorithm for solving multi-dimensional discrete-time optimal stopping problems. *Bernoulli 9*, 6 (2003), 1003–1049.

[7] Bally, V., Pagès, G., and Printems, J. A quantization tree method for pricing and hedging multidimensional American options. *Math. Finance 15*, 1 (2005), 119–168.

[8] Costa, O. L. V., and Davis, M. H. A. Approximations for optimal stopping of a piecewise-deterministic process. *Math. Control Signals Systems 1*, 2 (1988), 123–146.

[9] Costa, O. L. V., and Dufour, F. Stability and ergodicity of piecewise deterministic Markov processes. *SIAM J. Control Optim. 47*, 2 (2008), 1053–1077.

[10] Costa, O. L. V., Raymundo, C. A. B., and Dufour, F. Optimal stopping with continuous control of piecewise deterministic Markov processes. *Stochastics Stochastics Rep. 70*, 1-2 (2000), 41–73.

[11] Davis, M. H. A. *Markov models and optimization*, vol. 49 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London, 1993.

[12] Dufour, F., and Costa, O. L. V. Stability of piecewise-deterministic Markov processes. *SIAM J. Control Optim. 37*, 5 (1999), 1483–1502 (electronic).

[13] Gątarek, D. On first-order quasi-variational inequalities with integral terms. *Appl. Math. Optim. 24*, 1 (1991), 85–98.

[14] Gray, R. M., and Neuhoff, D. L. Quantization. *IEEE Trans. Inform. Theory 44*, 6 (1998), 2325–2383. Information theory: 1948–1998.

[15] Gugerli, U. S. Optimal stopping of a piecewise-deterministic Markov process. *Stochastics 19*, 4 (1986), 221–236.

[16] Kushner, H. J. *Probability methods for approximations in stochastic control and for elliptic equations*. Academic Press [Harcourt Brace Jovanovich Publishers], New York, 1977. Mathematics in Science and Engineering, Vol. 129.

[17] Lenhart, S., and Liao, Y. C. Integro-differential equations associated with optimal stopping time of a piecewise-deterministic process. *Stochastics 15*, 3 (1985), 183–207.

[18] Pagès, G. A space quantization method for numerical integration. *J. Comput. Appl. Math. 89*, 1 (1998), 1–38.

[19] Pagès, G., and Pham, H. Optimal quantization methods for nonlinear filtering with discrete-time observations. *Bernoulli 11*, 5 (2005), 893–932.

[20] Pagès, G., Pham, H., and Printems, J. An optimal Markovian quantization algorithm for multi-dimensional stochastic control problems. *Stoch. Dyn. 4*, 4 (2004), 501–545.

[21] Pagès, G., Pham, H., and Printems, J. Optimal quantization methods and applications to numerical problems in finance. In *Handbook of computational and numerical methods in finance*. Birkhäuser Boston, Boston, MA, 2004, pp. 253–297.