Approximation of the value function of an impulse control problem for Piecewise Deterministic Markov Processes

Benoîte de Saporta François Dufour Huilong Zhang

INRIA Bordeaux Sud-Ouest
University of Bordeaux
funded by ANR-09-SEGI-004

August 30 2011 – IFAC 11 – Milano

Definition of piecewise deterministic Markov processes

Davis (80's)

General class of non-diffusion dynamic stochastic hybrid models: deterministic motion punctuated by random jumps.

Applications

Engineering systems, operations research, management science, economics, dependability and safety,. . .

Dynamics

Hybrid process $X_t = (m_t, y_t)$

- discrete mode $m_t \in \{1, 2, \dots, p\}$
- Euclidean state variable $y_t \in \mathbb{R}^n$

Local characteristics for each mode m

- E_m open subset of $\mathbb{R}^d$, ∂E_m its boundary and $\overline{E}_m$ its closure
- Flow ϕ_m: $\mathbb{R}^d \times \mathbb{R} \to \mathbb{R}^d$ deterministic motion between jumps, one-parameter group of homeomorphisms
- Intensity λ_m: $\overline{E}_m \to \mathbb{R}_+$ intensity of random jumps
- Markov kernel Q_m on $(\overline{E}_m, \mathcal{B}(\overline{E}_m))$ selects the post-jump location

Iterative construction

Starting point

$$X_0 = Z_0 = (m, y)$$

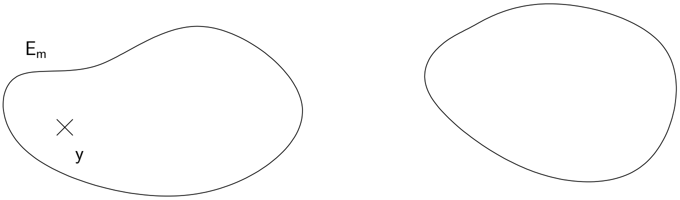

Iterative construction

X_t follows the deterministic flow until the first jump time $T_1 = S_1$

$$X_t = \big(m, \phi_m(y,t)\big), \quad t < T_1$$

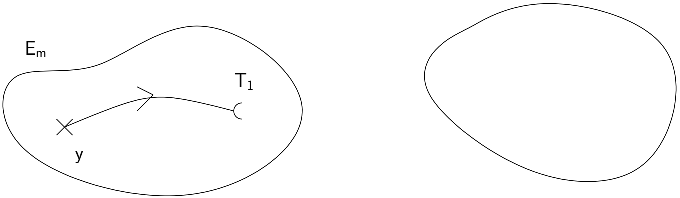

Iterative construction

Post-jump location $Z_1 = (M_1, Y_1)$ selected by

$$Q_m(\phi_m(y, T_1), \cdot)$$

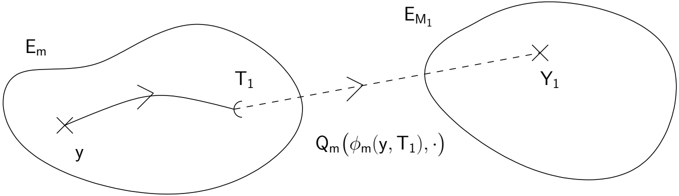

Iterative construction

X_t follows the flow until the next jump time $T_2 = T_1 + S_2$

$$X_{T_1+t} = \big(M_1, \phi_{M_1}(Y_1, t)\big), \quad t < S_2$$

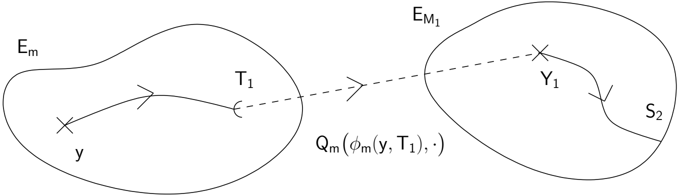

Iterative construction

Post-jump location $Z_2 = (M_2, Y_2)$ selected by

$$Q_{M_1}(\phi_{M_1}(Y_1, S_2), \cdot) \ldots$$

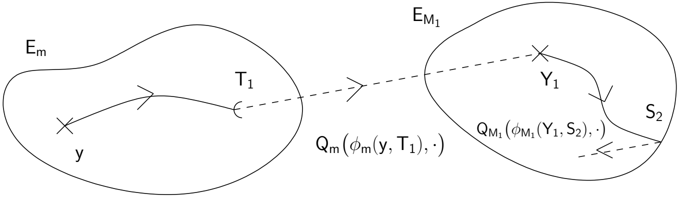

Embedded Markov chain

$\{X_t\}$ strong Markov process (M.H.A. Davis)

Natural embedded Markov chain

- Z_0 starting point, $S_0 = 0$, $S_1 = T_1$
- Z_n new mode and location after n-th jump, $S_n = T_n - T_{n-1}$, time between two jumps

Proposition

(Z_n, S_n) is a discrete-time Markov chain
Only source of randomness of the PDMP

Impulse control

Impulse control

Choose

- intervention times
- new starting points for the process after the interventions

in order to minimize a cost function

Application: maintenance of a complex system

Machine subject to failure of its components. Choose

- the intervention dates to perform a maintenance
- the nature of the maintenance: full of partial reparation

Mathematical definition

Strategy $\mathcal{S} = (\tau_n, R_n)_{n \geq 1}$

- τ_n intervention times
- R_n new positions after intervention

Value function

$$\mathcal{J}^{\mathcal{S}}(x) = E_x^{\mathcal{S}}\left[\int_0^{\infty} e^{-\alpha s} f(Y_s)ds + \sum_{i=1}^{\infty} e^{-\alpha\tau_i} c(Y_{\tau_i}, Y_{\tau_i^+})\right]$$

$$\mathcal{V}(x) = \inf_{\mathcal{S}\in\mathbb{S}} \mathcal{J}^{\mathcal{S}}(x)$$

- f, c cost functions, α discount factor
- Y_t controlled process, $\mathbb{S}$ set of admissible strategies

Dynamic programming

Costa, Davis, 1988

For any function $g \geq$ cost of the no-impulse strategy

- $v_0 = g$
- $v_n = \mathcal{L}(v_{n-1})$

$$v_n(x) \xrightarrow[n\to\infty]{} \mathcal{V}(x)$$

Dynamic programming operator

$$
\begin{aligned}
\mathcal{L}(w)(x) &= L(Mw, w)(x) \\
&= \Bigg(\inf_{t \le t^*(x)} \mathbb{E}_x \Big[F(x,t) + \mathrm{e}^{-\alpha S_1} w(Z_1) \mathbf{1}_{\{S_1 < t \wedge t^*(x)\}} \\
&\quad + \mathrm{e}^{-\alpha t \wedge t^*(x)} Mw(\phi(x, t \wedge t^*(x)) \mathbf{1}_{\{S_1 \ge t \wedge t^*(x)\}} \Big] \Bigg) \\
&\quad \wedge \mathbb{E}_x \Big[F(x, t^*(x)) + \mathrm{e}^{-\alpha S_1} w(Z_1) \Big]
\end{aligned}
$$

with

$$
\begin{aligned}
F(x,t) &= \int_0^{t \wedge t^*(x)} \mathrm{e}^{-\alpha s - \Lambda(x,s)} f\big(\phi(x,s)\big)\, ds \\
Mw(x) &= \inf_{y \in \mathbb{U}} \{ c(x,y) + w(y) \}
\end{aligned}
$$

Our aim

Propose a numerical method

- to compute an approximation of the value function
- with error bounds

Main difficulty

Discretization of the dynamic programming operator

Our approach

Discretize the underlying Markov chain (Z_n, S_n)

Quantization

Quantization of a random variable X

Approximate X by $\widehat{X}$ taking finitely many values such that $\|X - \widehat{X}\|_p$ is minimum

- Find a finite weighted grid Γ with $|\Gamma| = K$
- Set $\widehat{X} = p_\Gamma(X)$ closest neighbor projection

Algorithms

There exist algorithms providing

- Γ
- law of $\widehat{X}$
- transition probabilities for quantization of Markov chains

Example: $\mathcal{N}(0, I_2)$:

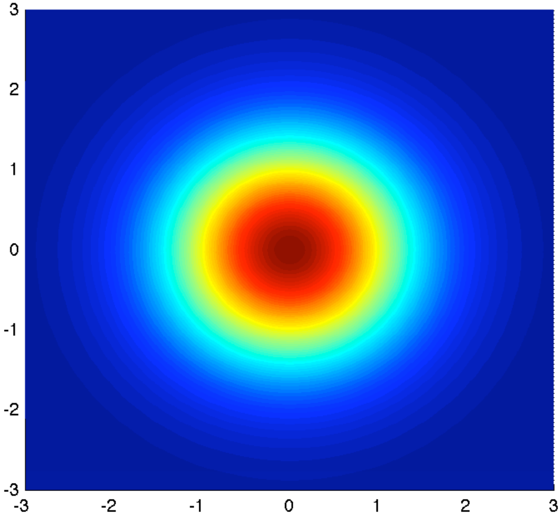

Example: $\mathcal{N}(0, I_2)$:

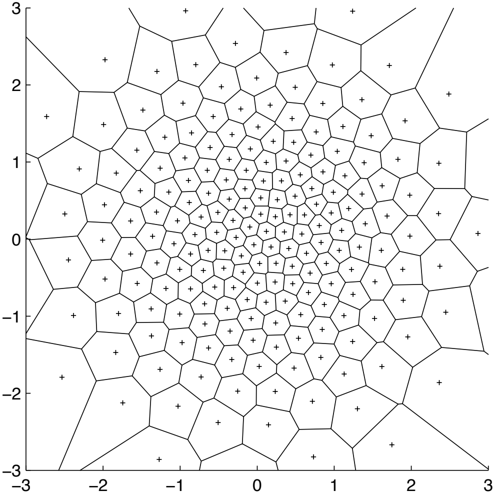

Example: $\mathcal{N}(0, I_2)$:

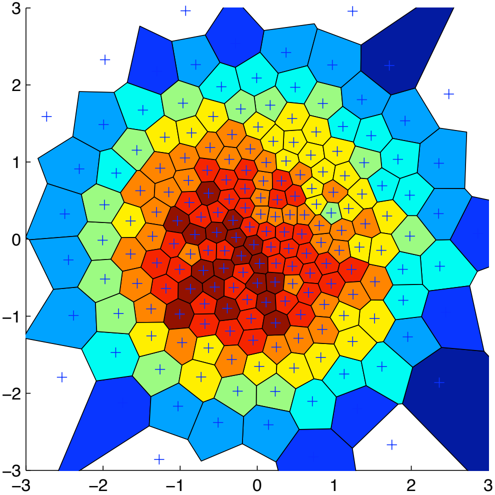

Horizon and control set

- finite set $\mathbb{U}$ of new starting points
- select horizon N such that $v_N(x) - \mathcal{V}(x)$ small enough

$$\longrightarrow \text{numerical approximation of } v_N(x)$$

Main idea

Replace the dynamic programming iteration of functions by an iteration of random variables

Backward dynamic programming

Change of notation

For a well chosen function g and large enough N

- $v_N = g$
- $v_n = \mathcal{L}(v_{n+1})$

$$v_0(x) \simeq \mathcal{V}(x)$$

Dynamic programming

Markov property

$$
\begin{aligned}
v_n(Z_n) &= L(Mv_{n+1}, v_{n+1})(Z_n) \\
&= \Bigg(\inf_{t \le t^*(Z_n)} \mathbb{E}\Big[F(Z_n, t) + e^{-\alpha S_{n+1}} v_{n+1}(Z_{n+1}) \mathbf{1}_{\{S_{n+1} < t \wedge t^*(Z_n)\}} \\
&\quad + e^{-\alpha t \wedge t^*(Z_n)} Mv_{n+1}\big(\phi(Z_n, t \wedge t^*(Z_n))\big) \mathbf{1}_{\{S_{n+1} \ge t \wedge t^*(Z_n)\}} \mid Z_n \Big] \\
&\quad \wedge \mathbb{E}\Big[F(Z_n, t^*(Z_n)) + e^{-\alpha S_{n+1}} v_{n+1}(Z_{n+1}) \mid Z_n \Big]
\end{aligned}
$$

with

$$
\begin{aligned}
F(x,t) &= \int_0^{t \wedge t^*(x)} e^{-\alpha s - \Lambda(x,s)} f\big(\phi(x,s)\big)\, ds \\
Mv_{n+1}(x) &= \inf_{y \in \mathbb{U}} \big\{ c(x,y) + v_{n+1}(y) \big\}
\end{aligned}
$$

Recurrence on random variables

$v_n(Z_n)$ expression of $v_{n+1}(Z_{n+1})$, Z_n, S_{n+1}
$+\ v_{n+1}(y)$ for all y in $\mathbb{U}$

Numerical scheme

- first compute recursively $\widetilde{v}_n(y)$ approximation of $v_n(y)$ for all y in $\mathbb{U}$
- then compute recursively $\widehat{v}_n(\widehat{Z}_n)$ approximation of $v_n(Z_n)$

Discretization

In the expression of operator L replace

- inf by min over a discretized grid
- Z_n, Z_{n+1}, S_{n+1} by their quantized approximation starting from $Z_0 \in \mathbb{U}$

Recurrence on random variables

$v_n(Z_n)$ expression of $v_{n+1}(Z_{n+1})$, Z_n, S_{n+1}
$+ \; v_{n+1}(y)$ for all y in $\mathbb{U}$

Numerical scheme

- first compute recursively $\widetilde{v}_n(y)$ approximation of $v_n(y)$ for all y in $\mathbb{U}$
- then compute recursively $\widehat{v}_n(\widehat{Z}_n)$ approximation of $v_n(Z_n)$

Discretization

In the expression of operator L replace

- inf by min over a discretized grid
- Z_n, Z_{n+1}, S_{n+1} by their quantized approximation starting from $Z_0 = x$

Properties of the numerical scheme

- the quantized process has no Markov property $\longrightarrow$ a different approximation of L for each time step and each starting point
- Under Lipschitz regularity assumption, convergence of the scheme with errors bounds depending on
  - the time discretization step inf $\longrightarrow$ min
  - the quantization error $(Z_n, S_n) \rightarrow (\widehat{Z}_n, \widehat{S}_n)$

Numerical scheme in practice

Draw two series of quantization grids for (Z_n, S_n)

- $(\widehat{Z}_n, \widehat{S}_n)$ starting uniformely on $\mathbb{U}$
- $(\widehat{Z}_n, \widehat{S}_n)$ starting from x

Compute the triangular scheme (ex: $N = 3$)

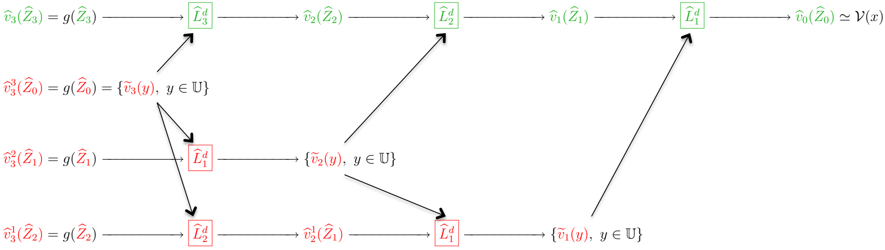

Simple example

object moving on $[0;1[$ with constant speed

Local characteristics

- $\phi(x,t) = x + t$
- $\lambda(x) = 3x$: as the object comes closer to 1 the probablity to jump increases
- $Q(x,\cdot)$ uniform law on $[0;1/2]$

Control parameters

- discount factor $\alpha = 2$
- running cost $f(x) = 1 - x$, constant intervention cost $c(x,y) = 0.08$
- control set $\mathbb{U} = \{k/50, 0 \leq k < 50\}$

Trajectories

Examples of trajectories for $X_0 = 0$ up to the 10-th jump

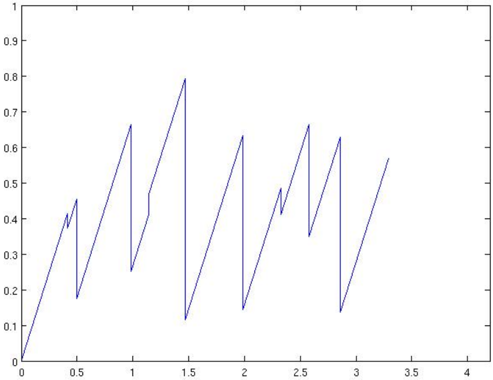

Trajectories

Examples of trajectories for $X_0 = 0$ up to the 10-th jump

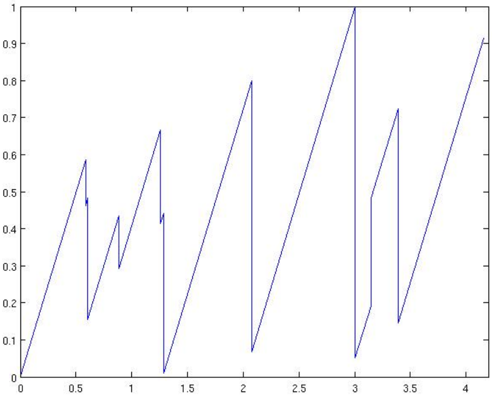

Approximation of the value function

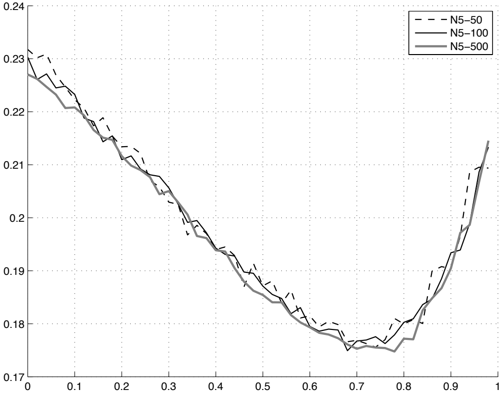

horizon $N = 5$

Approximation of the value function

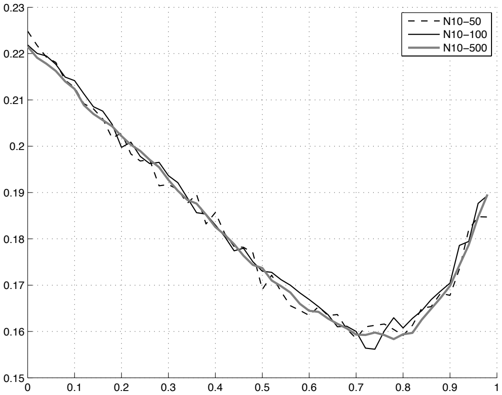

horizon $N = 10$

Approximation of the value function

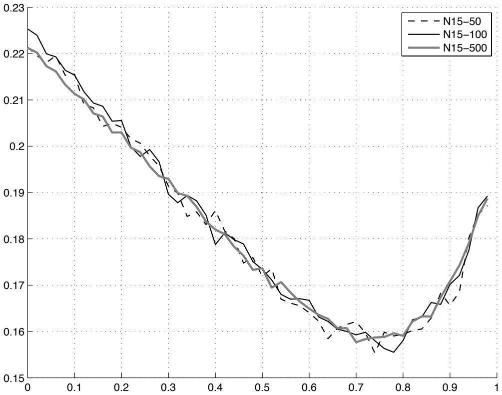

horizon $N = 15$

Choice of the computation horizon

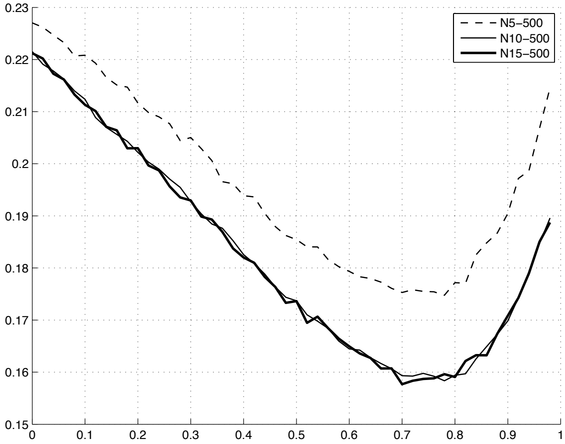

Perspectives

- try our method on industrial examples
- propose a numerical scheme to approximate an ε-optimal strategy

  $\longrightarrow$ open post-doc positions