Numerical method for optimal stopping of PDMP and maintenance optimization

Camille Baysse
Thales Optronique

Benoîte de Saporta
Univ. Bordeaux

Anne Gégout-Petit
Univ. Lorraine

International Workshop on Applied Probability

Antalya, 17 June 2014

Outline

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, internet traffic, neurosciences, biology, dependability and safety. . .

Dynamics

Hybrid process $X_t = (m_t, y_t)$

- discrete mode $m_t \in \{1, 2, \ldots, p\}$
- 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 post-jump location

Two types of jumps

- $t^*(m,y)$ deterministic exit time starting from (m,y)

$$t^*(m,y) = \inf\{t > 0 : \phi_m(y,t) \in \partial E_m\}$$

- law of the first jump time T_1 starting from (m,y)

$$\mathbb{P}_{(m,y)}(T_1 > t) = \begin{cases} \mathrm{e}^{-\int_0^t \lambda_m\left(\phi_m(y,s)\right) ds} & \text{if} \quad t < t^*(m,y) \\ 0 & \text{if} \quad t \geq t^*(m,y) \end{cases}$$

Remark

T_1 has a density on $[0, t^*(m,y)[$ but has an atom at $t^*(m,y)$

$$\mathbb{P}_{(m,y)}(T_1 = t^*(m,y)) > 0$$

Iterative construction

Starting point

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

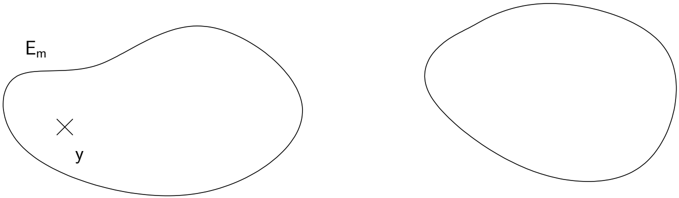

Iterative construction

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

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

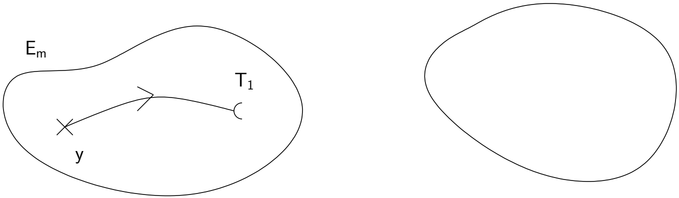

Iterative construction

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

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

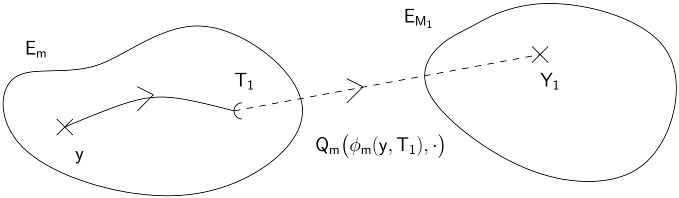

Iterative construction

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

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

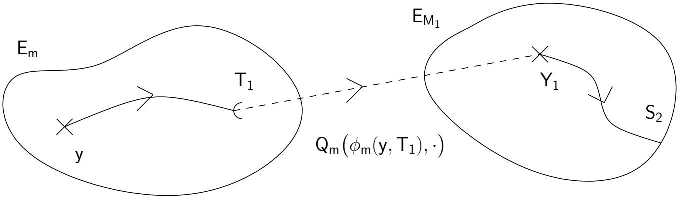

Iterative construction

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

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

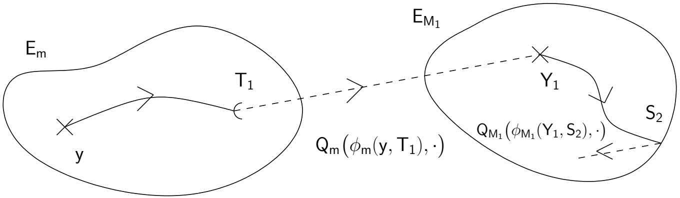

Embedded Markov chain

$\{X_t\}$ strong Markov process [Davis 93]

Natural embedded Markov chain

- Z_0 starting point, $T_0 = 0$, $S_0 = 0$
- Z_n new mode and location after n-th jump
  T_n date of n-th jump, $S_n = T_n - T_{n-1}$

Important property

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

Numerical methods for PDMP

Fact

- numerous application domains
- numerous theoretical results
  [Davis 93], [Jacobsen 06], [Costa Dufour 13]
- processes easy to simulate if explicit flow
- very few dedicated numerical methods in the literature for optimal control
  [Costa Davis 88, 89]

Aim of the talk

Propose a new numerical method

- adapted to the specificities of PDMPs
- with proofs and rate of convergence
- implementable in practice

To solve approximately the optimal stopping problem for PDMPs

Outline

Problem setting

- Reward function g
- Random time horizon: N-th jump time T_N
- $\mathcal{M}_N$ set of all stopping times $\tau \leq T_N$

Optimal stopping problem

- compute the value function

$$V(x) = \sup_{\tau \in \mathcal{M}_N} \mathbb{E}_x[g(X_\tau)]$$

- find an (ε-)optimal stopping time τ^* that reaches $V(x)(-\varepsilon)$

Recursion for value functions

[Gugerli 1986]

Dynamic programming

- $v_N = g$
- $v_n = L(v_{n+1}, g)$ pour $n \leq N-1$

$$v_0(x) = \sup_{\tau \in \mathcal{M}_N} \mathbb{E}_x[g(X_\tau)] = V(x)$$

$$\begin{aligned} L(w,g)(x) \\ = \quad & \sup_{u \leq t^*(Z_n)} \Big\{ \mathbb{E}\Big[w(Z_{n+1}) \mathbb{1}_{\{S_{n+1} < u\}} + g\big(\phi(Z_n, u)\big) \mathbb{1}_{\{S_{n+1} \geq u\}} \;\Big|\; Z_n = x \Big] \Big\} \\ & \vee\, \mathbb{E}\big[w(Z_{n+1}) \;\big|\; Z_n = x \big] \end{aligned}$$

Recursion for random variables

Dynamic programming

- $v_N(Z_N) = g(Z_N)$
- $v_n(Z_n) = L(v_{n+1}, g)(Z_n)$ pour $n \leq N-1$

$$v_0(Z_0) = \sup_{\tau \in \mathcal{M}_N} \mathbb{E}_x[g(X_\tau)]$$

$$\begin{aligned} v_n(Z_n) &= L(v_{n+1}, g)(Z_n) \\ &= \sup_{u \leq t^*(Z_n)} \left\{ \mathbb{E}\left[v_{n+1}(Z_{n+1}) \mathbb{1}_{\{S_{n+1} < u\}} + g\big(\phi(Z_n, u)\big) \mathbb{1}_{\{S_{n+1} \geq u\}} \;\middle|\; Z_n \right] \right\} \\ &\quad \vee \mathbb{E}\left[v_{n+1}(Z_{n+1}) \;\middle|\; Z_n \right] \end{aligned}$$

Strategy

- discretize the chain (Z_n, S_n) using quantization
- replace (Z_n, S_n) by its approximation $(\widehat{Z}_n, \widehat{S}_n)$ in L
  $\longrightarrow$ computable approximation
- study convergence, derive error bounds
  - indicator functions in the dynamic programming equation
    $\longrightarrow$ be careful with the time grids

Quantization

[Pagès 98], [Pagès, Pham, Printems 04]...

Quantization of a random variable $X \in L^p(\mathbb{R}^d)$

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

Asymptotic properties

If $E[|X|^{p+\eta}] < +\infty$ for some $\eta > 0$ then

$$\lim_{K\to\infty} K^{1/d} \min_{|\Gamma|\leq K} \|X - \widehat{X}^\Gamma\|_p = C$$

Algorithms

There exist algorithms providing

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

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

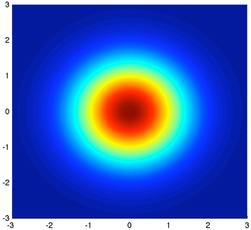

Algorithms

There exist algorithms providing

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

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

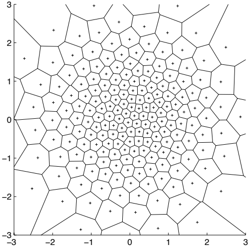

Algorithms

There exist algorithms providing

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

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

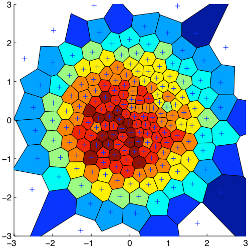

Grids construction

Model ⟶ simulator of trajectories ⟶ grids

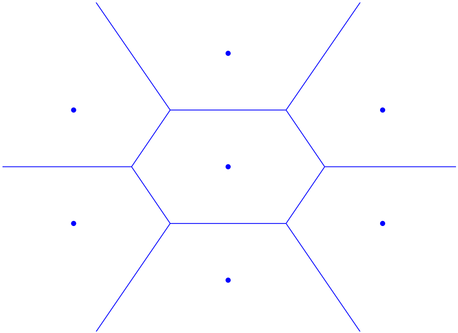

Grids construction

Model $\longrightarrow$ simulator of trajectories $\longrightarrow$ grids

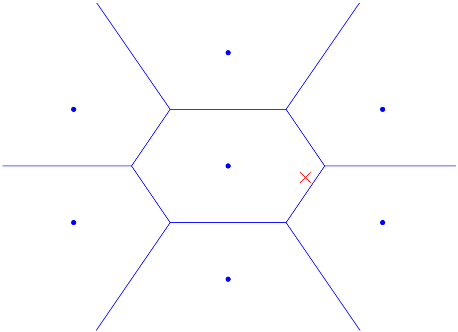

Grids construction

Model ⟶ simulator of trajectories ⟶ grids

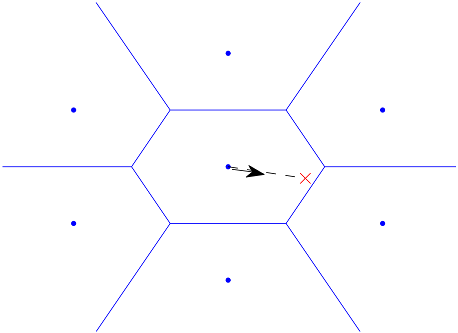

Grids construction

Model $\longrightarrow$ simulator of trajectories $\longrightarrow$ grids

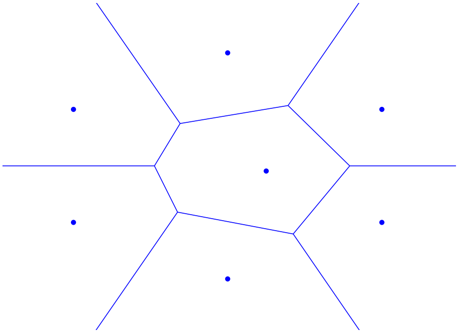

Assets and drawbacks of quantization

Assets

- a simulator of the target law is enough to build the grids
- automatic construction of grids
- convergence rate for $\mathbb{E}[|f(X) - f(\widehat{X})|]$ if f lipschitz

Drawbacks

- computation time
- curse of dimension
- open questions of convergence of the algorithms

Discretization

Approximation of the value function

- $\widehat{v}_N(\widehat{Z}_N) = g(\widehat{Z}_N)$
- $\widehat{v}_n(\widehat{Z}_n) = \widehat{L}_d^n(\widehat{v}_{n+1}, g)(\widehat{Z}_n)$ for $n \leq N-1$

$$
\begin{aligned}
&\widehat{L}_d^n(v_{n+1}, g)(\widehat{Z}_n) \\
&= \max_{u \in G(\widehat{Z}_n)} \Big\{ \mathbb{E}\Big[v(\widehat{Z}_{n+1}) \mathbf{1}_{\{\widehat{S}_{n+1} < u\}} + g\big(\phi \widehat{Z}_n, u)\big) \mathbf{1}_{\{\widehat{S}_{n+1} \geq u\}} \;\Big|\; \widehat{Z}_n \Big] \Big\} \\
&\quad \vee \mathbb{E}\big[v(\widehat{Z}_{n+1}) \;\big|\; \widehat{Z}_n \big]
\end{aligned}
$$

Convergence

[dS, Dufour, Gonzalez, Ann. Appl. Proba. 2010]

Theroem

Lipschitz regularity assumptions on ϕ, λ, Q, t^* and g

$$|v_0(x) - \widehat{v}_0(x)| \leq C\sqrt{EQ}$$

C explicit constant,
EQ quantization error

$\sqrt{\cdot}$ due to the indicator functions

ϵ-optimal stopping time

Computable stopping rule $\hat{\tau}$

- explicit iterative construction
- no extra computation
- true stopping time for the original process (X_t)

Theorem

Same assumptions

$$|v_0(x) - \mathbb{E}_x[g(X_{\hat{\tau}})]| \leq C_1 EV + C_2 \sqrt{EQ}$$

C_1, C_2 explicit constants, EV value function error, EQ quantization error

Provides another approximation of the value function via Monte Carlo simulations

Outline

Application to maintenance optimization

- $X_t = (m_t, y_t)$ state of a machine at time t
- T_n failure of some components

Maintenance optimization

Find an optimal balance between

- changing the components too early/often
- do nothing until total breakdown

Industrial problem from Thales optronique

Compute an optimized maintenance date for an equipment subject to different kinds of failures

Air conditioning unit

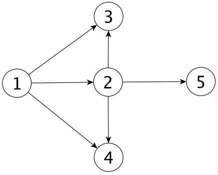

- State 1: stable state
- State 2: degraded ball bearing
- State 3: failed electrovalve
- State 4: electronic failure
- State 5: failed ball bearing

PDMP model

Transition rates

- degraded ball bearing and failed electorovalve: Weibull distributions ⇒ time dependent intensity
- electronic and ball bearing failures: exponential distribution

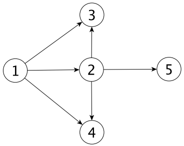

- State 1: stable state
- State 2: degraded ball bearing
- State 3: failed electrovalve
- State 4: electronic failure
- State 5: failed ball bearing

PDMP model

- discrete mode $m_t \in \{1, 2, 3, 4, 5\}$
- Euclidean state variable $y_t = t$ working time

Trajectories without maintenance

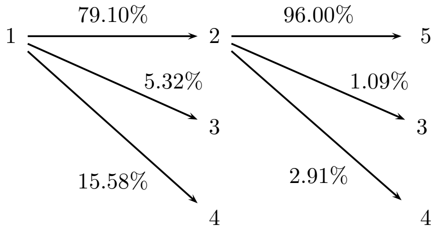

- 1: stable state
- 2: degraded ball bearing
- 3: failed electrovalve
- 4: electronic failure
- 5: failed ball bearing

Reward function

$$g(m,t) = \frac{t}{p(m)}$$

- $p(1) = 6$ price of maintenance in stable mode
- $p(2) = 6$ price of maintenance in degraded ball bearing mode
- $p(3) = 5$ price of repair of electrovalve failure
- $p(4) = 3.5$ price of repair of electronic failure
- $p(5) = 12$ price of repair of ball bearing failure

Maintenance optimisation

- Better to start a maintenance in mode 2 than wait for total failure mode 5
- Failure modes 3 and 4 cost less than maintenance

Average performance without maintenance: 342.72

Trajectories with maintenance

Without maintenance

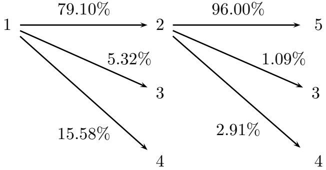

Average performance : 342.72

With maintenance

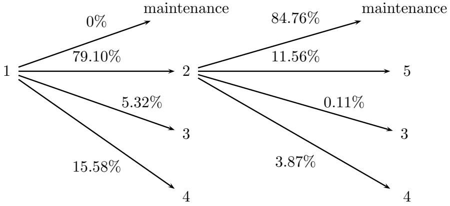

Average performance : 592.47

Conclusion and perspectives

Assets and drawbacks of the numerical method

- practical method
- computation time on line/off line
- curse of dimension

Perspectives

- Optimal policy for impulse control Thales Optronique
- Numerical methods for MDP Airbus, DCNS

Thank you