Numerical method for control of piecewise deterministic Markov processes

Benoîte de Saporta, François Dufour

Adrien Brandejsky, Karen Gonzalez, Huilong Zhang

University of Bordeaux and Inria

University of Bordeaux
Inria CQFD
France

ICMC – USP 4 June 2013

Outline

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, internet traffic, neurosciences, biology, 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 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

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

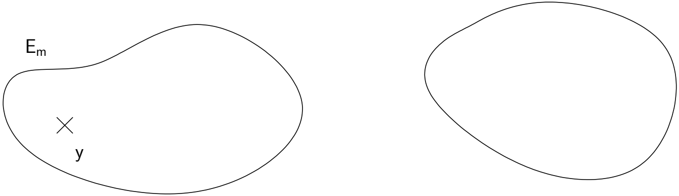

Iterative construction

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

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

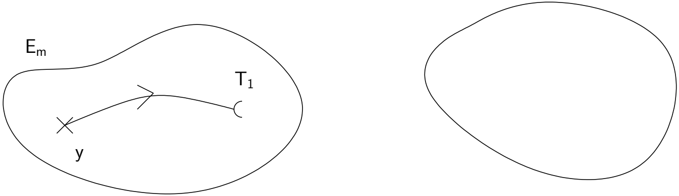

Iterative construction

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

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

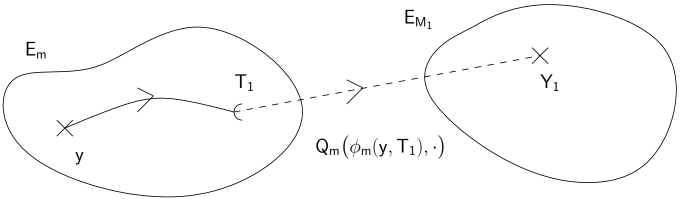

Iterative construction

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

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

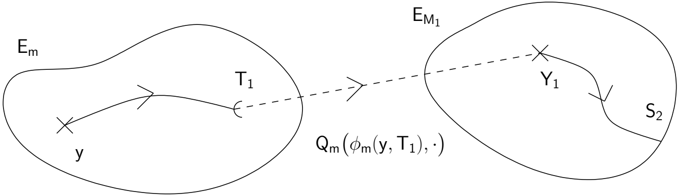

Iterative construction

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

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

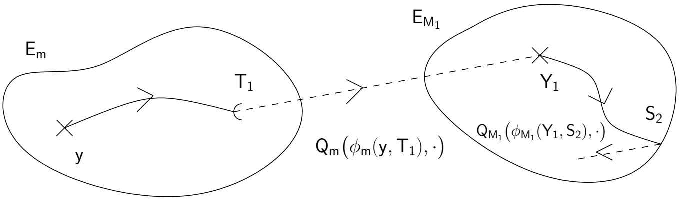

Embedded Markov chain

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

Natural embedded Markov chain

- Z_0 starting point, $S_0 = 0$, $S_1 = T_1$
- 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

Industrial example from Astrium Space Transportation

Material subject to corrosion and randomly exposed to different stressing ambiences

Corrosion problem

Material subject to corrosion

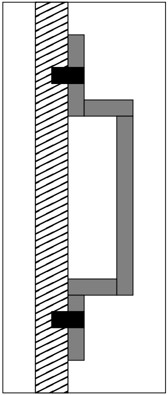

- support structure for other equipments
- small size: one point of measure
- long service life $\rightarrow$ monitor the thickness loss due to corrosion

Usage profile

Material subject to corrosion

Usage profile

Storage in 3 successive differents environments for random times

1. workshop
2. submarine in operation
3. submarine in dry dock

Strong safety requirements

$\Downarrow$

Monitor the thickness loss

Degradation process

- Deterministic succession of environments : $1 \to 2 \to 3 \to 1 \cdots$

Degradation process

- Deterministic succession of environments : $1 \to 2 \to 3 \to 1 \cdots$
- Random time in environment i with distribution $Exp(\lambda_i)$

Degradation process

- Deterministic succession of environments : $1 \to 2 \to 3 \to 1 \cdots$
- Random time in environment i with distribution $Exp(\lambda_i)$
- Initial protection against corrosion active for random time with Weibull distribution

Degradation process

- Deterministic succession of environments : $1 \to 2 \to 3 \to 1 \cdots$
- Random time in environment i with distribution $Exp(\lambda_i)$
- Initial protection against corrosion active for random time with Weibull distribution
- Equation of thickness loss in environment i

$$d_t = \rho_i \Big(t - \eta_i + \eta_i exp(-t/\eta_i) \Big)$$

  - ρ_i random corrosion rate in environment i with uniform distribution
  - η_i deterministic time in environnement i.

Degradation process

- Deterministic succession of environments : $1 \to 2 \to 3 \to 1 \cdots$
- Random time in environment i with distribution $Exp(\lambda_i)$
- Initial protection against corrosion active for random time with Weibull distribution
- Equation of thickness loss in environment i

$$d_t = \rho_i \Big(t - \eta_i + \eta_i exp(-t/\eta_i) \Big)$$

  - ρ_i random corrosion rate in environment i with uniform distribution
  - η_i deterministic time in environnement i.

Inefficient material if $d_t \geq 0.2mm$

Simulated trajectories

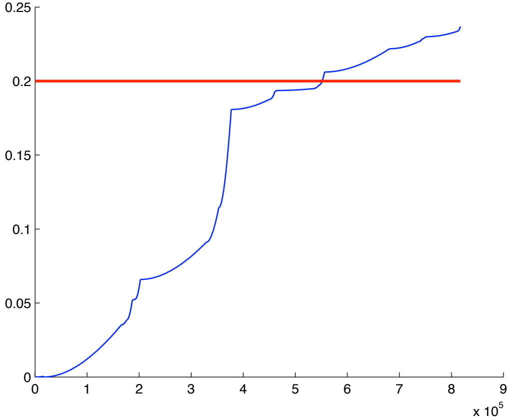

Simulated trajectories

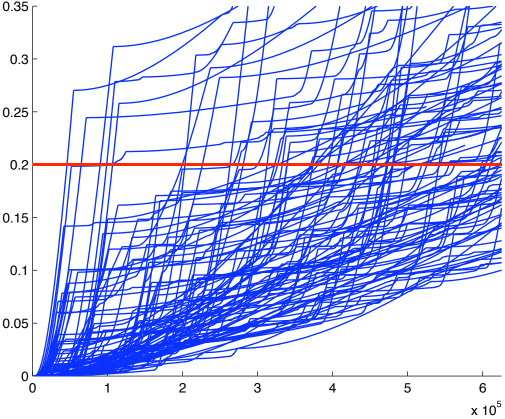

Outline

Optimal stopping

- Reward function g
- Time horizon N-th jump T_N

Optimal stopping problem

- compute the value function

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

$\mathcal{M}$ set of all stopping times $\tau \leq T_N$ for the natural filtration of the process (X_t)

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

Application to maintenance optimization

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

Optimal stopping

Find an optimal balance between

- changing the components too early/often
- no maintenance leading to a total breakdown

Numerical methods

Fact

Very few numerical methods available for PDMP's
[Costa Davis 88, 89]

Our aim

Propose numerical methods

- suitable for PDMP's
- with proof of convergence and convergence rate
- that can work in practice

Aim of the talk

Objective

Propose a numerical method

- to evaluate the value function
- to compute an ε-optimal stopping rule

with error bounds

Strategy

- Use dynamic programming
- Use the embedded Markov chain (Z_n, S_n)
- Adapt the methodology developed for diffusion processes based on quantization

Outline

Numerical method for diffusion processes

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

Y_t continuous-time Markov diffusion process

Numerical method for diffusion processes

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

Y_t continuous-time Markov diffusion process

1. time discretization (Euler scheme) $Y_k = Y_{k\Delta t}$ discrete-time Markov chain with continuous state space

Numerical method for diffusion processes

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

Y_t continuous-time Markov diffusion process

1. time discretization (Euler scheme) $Y_k = Y_{k\Delta t}$ discrete-time Markov chain with continuous state space
2. quantization replace Y_k by a random variable $\widehat{Y}_k$ taking values in a finite state space

Numerical method for diffusion processes

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

Y_t continuous-time Markov diffusion process

1. time discretization (Euler scheme) $Y_k = Y_{k\Delta t}$ discrete-time Markov chain with continuous state space
2. quantization replace Y_k by a random variable $\widehat{Y}_k$ taking values in a finite state space
3. replace the conditional expectations in the dynamic programming equation by finite sums

Numerical method for diffusion processes

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

Y_t continuous-time Markov diffusion process

1. time discretization (Euler scheme) $Y_k = Y_{k\Delta t}$ discrete-time Markov chain with continuous state space
2. quantization replace Y_k by a random variable $\widehat{Y}_k$ taking values in a finite state space
3. replace the conditional expectations in the dynamic programming equation by finite sums

Lipschitz-continuity conditions $\Longrightarrow$ convergence rate of the approximated value function to the original one

Specificities of PDMP's

- jumps at random times
- indicator functions in the dynamic programming equation

Solution

- use the embedded Markov chain (Z_n, S_n)
- be careful with the time grids

Quantization

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

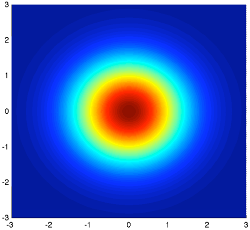

Algorithms

There exist algorithms providing

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

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

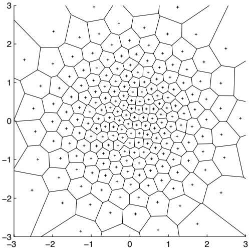

Algorithms

There exist algorithms providing

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

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

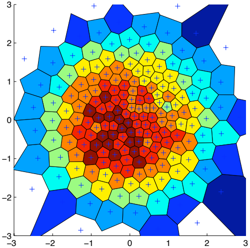

Grids construction

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

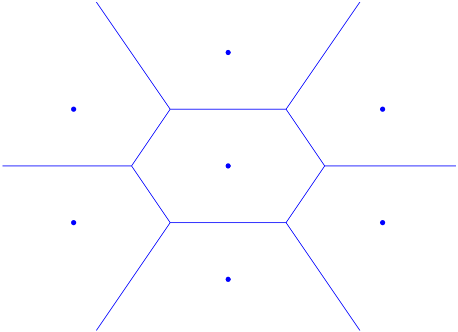

Grids construction

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

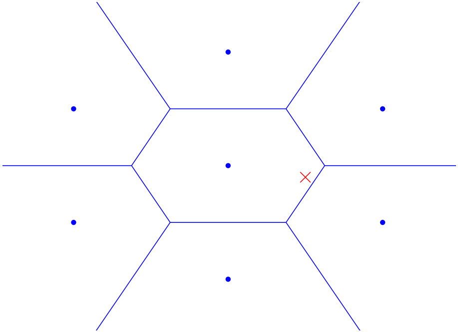

Grids construction

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

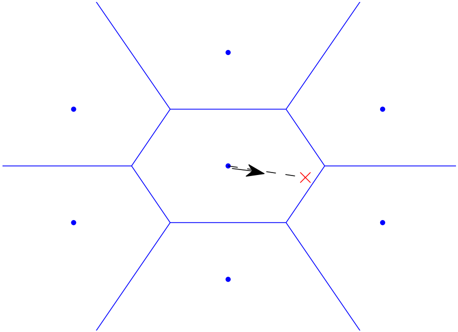

Grids construction

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

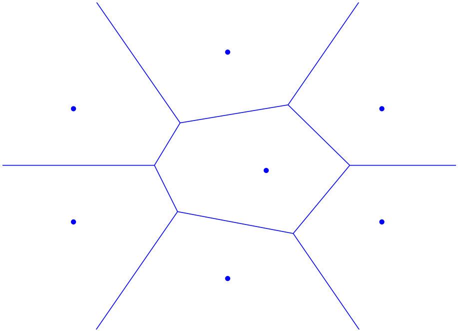

Assets and drawbacks of quantization

Assets

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

Drawbacks

- computation time
- curse of dimension
- open questions o conn thevergence of the algorithms

Optimal stopping for PDMP's

[Gugerli, 1986]

Dynamic programming equation

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

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

$$\begin{aligned} & L(w,g)(x) \\ = \quad & \sup_{t \geq 0} \Big[\int_0^{t \wedge t^*(x)} \lambda Q w(\phi(x,s)) \mathrm{e}^{-\Lambda(x,s)} ds + g\big(\phi(x, t \wedge t^*(x))\big) \mathrm{e}^{-\Lambda(x, t \wedge t^*(x))} \Big] \\ & \vee \int_0^{t^*(x)} \lambda Q w(\phi(x,s)) \mathrm{e}^{-\Lambda(x,s)} ds + Qw(\phi(x, t^*(x))) \mathrm{e}^{-\Lambda(x, t^*(x))} \end{aligned}$$

Probabilistic interpretation

Interpretation of operator L

$$
\begin{aligned}
L(w,g)(x) & \\
= \quad & \sup_{t\geq 0}\Big[\int_0^{t\wedge t^*(x)}\lambda Qw(\phi(x,s))\mathrm{e}^{-\Lambda(x,s)}ds + g\big(\phi(x,t\wedge t^*(x))\big)\mathrm{e}^{-\Lambda(x,t\wedge t^*(x))}\Big] \\
& \vee \int_0^{t^*(x)}\lambda Qw(\phi(x,s))\mathrm{e}^{-\Lambda(x,s)}ds + Qw(\phi(x,t^*(x)))\mathrm{e}^{-\Lambda(x,t^*(x))}
\end{aligned}
$$

Probabilistic interpretation

Interpretation of operator L

$$
\begin{aligned}
& L(w,g)(x) \\
& = \sup_{u \leq t^*(x)} \Big\{ \mathbb{E}_x \Big[w(Z_1) \mathbf{1}_{\{S_1 < u\}} + g\big(\phi(x,u)\big) \mathbf{1}_{\{S_1 \geq u\}} \Big] \Big\} \\
& \quad \vee \mathbb{E}_x \big[w(Z_1) \big]
\end{aligned}
$$

Probabilistic interpretation

Interpretation of operator L

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

Recursive computation

Backward dynamic programming equation

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

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

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

Discretization

Approximation of the value function

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

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

Discretization

Time discretization

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

$$L_d(v_{n+1}, g)(Z_n)$$

$$= \max_{u \in G(Z_n)} \Big\{ \mathbb{E}\Big[v(Z_{n+1}) \mathbf{1}_{\{S_{n+1} < u\}} + g\big(\phi(Z_n, u)\big) \mathbf{1}_{\{S_{n+1} \geq u\}} \;\Big|\; Z_n \Big] \Big\}$$

$$\vee\ \mathbb{E}\big[v(Z_{n+1}) \;\big|\; Z_n \big]$$

Discretization

Quantization

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

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

Discretization

Approximation of the value function

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

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

Convergence rate

[An. Appl. Proba. 2010]

Theorem

Lipschitz 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

Optimal stopping time

Optimal stopping time : τ^*

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

Existence?

Optimal stopping time

Optimal stopping time : τ^*

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

Existence?

ϵ-optimal stopping time : $\hat{\tau}$

$$v_0(x) - \epsilon \leq \mathbb{E}_x[g(X_{\hat{\tau}})] \leq v_0(x)$$

Optimal stopping time

Optimal stopping time : τ^*

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

Existence?

ϵ-optimal stopping time : $\hat{\tau}$

$$v_0(x) - \epsilon \leq \mathbb{E}_x[g(X_{\hat{\tau}})] \leq v_0(x)$$

Proposition of a computable stopping rule $\hat{\tau}$

- explicit iterative construction
- no extra computation

Optimal stopping time

Optimal stopping time : τ^*

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

Existence?

ϵ-optimal stopping time : $\hat{\tau}$

$$v_0(x) - \epsilon \leq \mathbb{E}_x[g(X_{\hat{\tau}})] \leq v_0(x)$$

Proposition of a computable stopping rule $\hat{\tau}$

- explicit iterative construction
- no extra computation
- true stopping time in $\mathcal{M}$

Optimality

[An. Appl. Proba. 2010]

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

Maintenance policy for the corrosion model

One intervention before ⇒ structure as good as new

Maintenance optimization : balance between

- early intervention unnecessary and costly
- late intervention dangerous and costly

Margin optimization

During conception

- ensure with 95% confidence that no maintenance will be required before a given date

Reward function

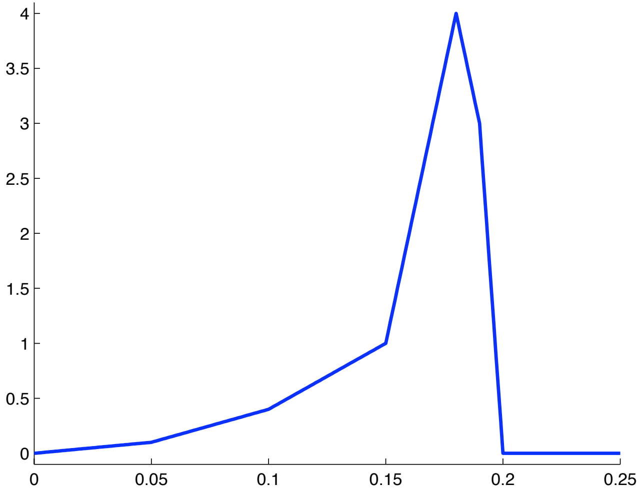

Iterative stopping rule

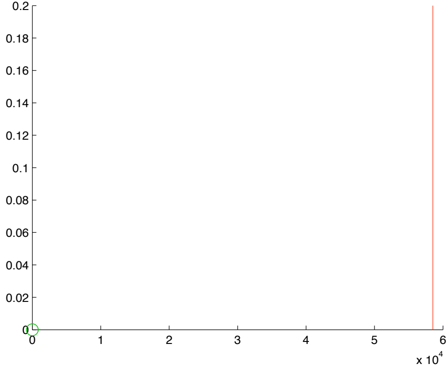

Iterative stopping rule

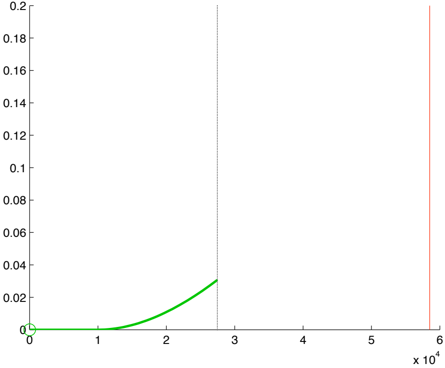

Iterative stopping rule

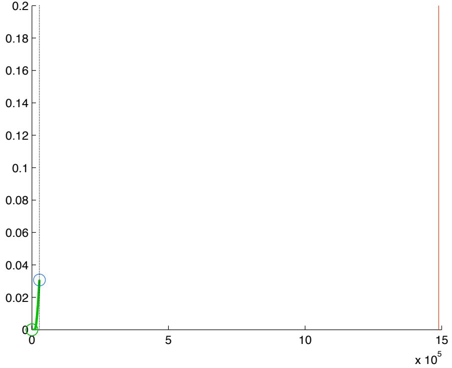

Iterative stopping rule

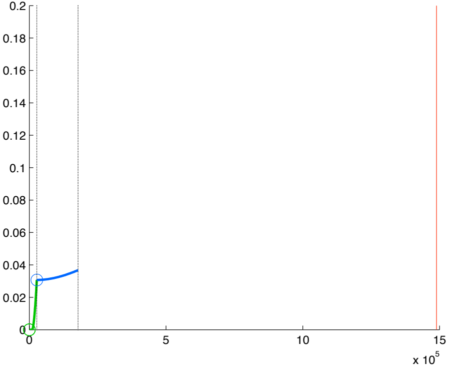

Iterative stopping rule

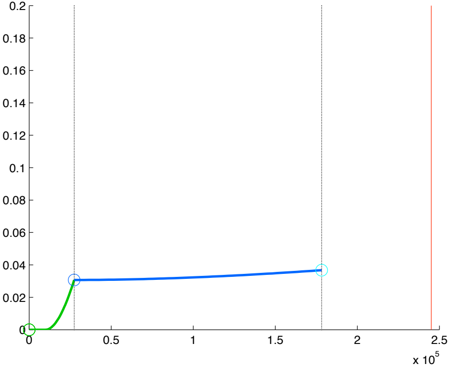

Iterative stopping rule

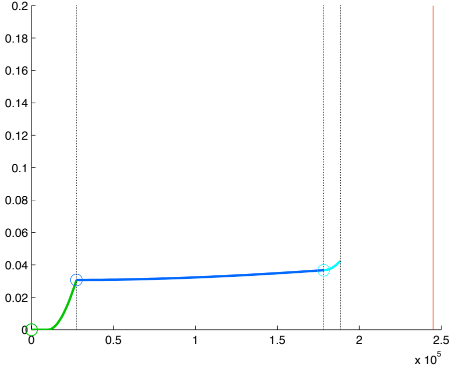

Iterative stopping rule

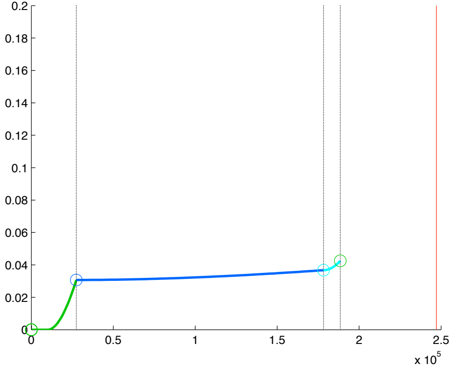

Iterative stopping rule

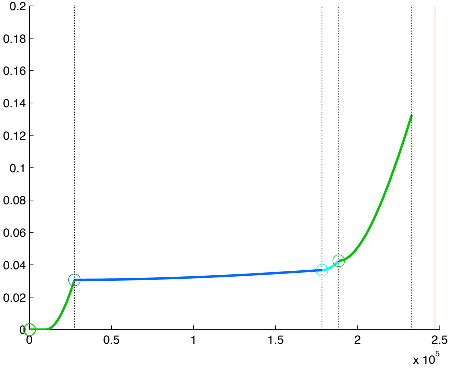

Iterative stopping rule

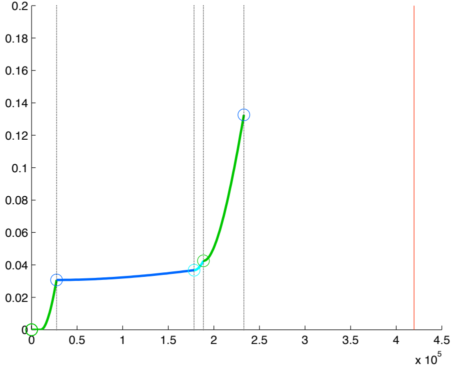

Iterative stopping rule

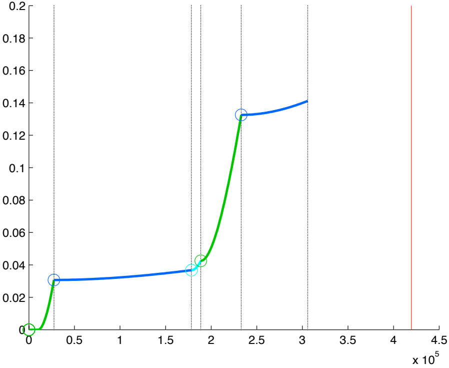

Iterative stopping rule

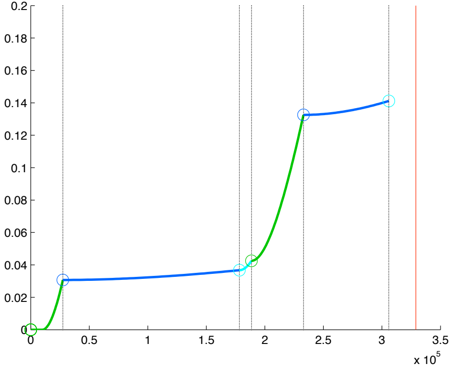

Iterative stopping rule

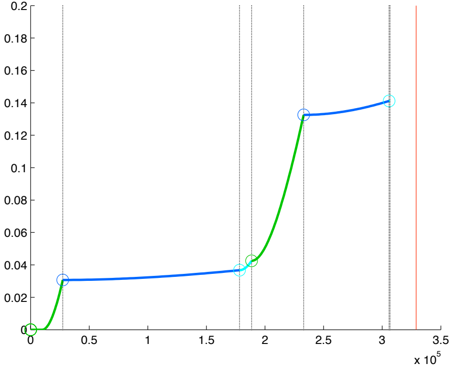

Iterative stopping rule

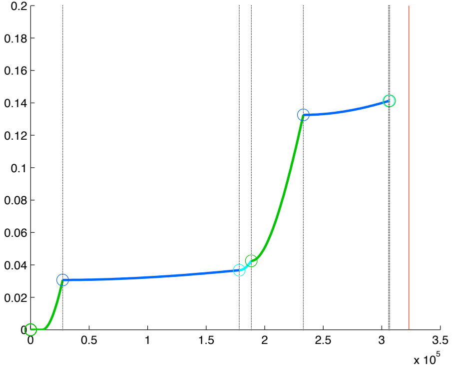

Iterative stopping rule

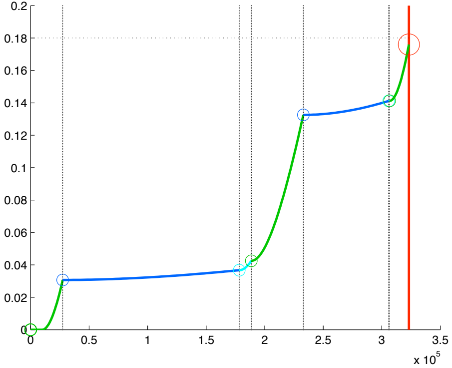

Margin optimization

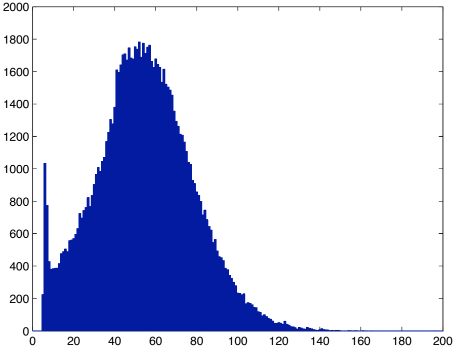

Date	Probability
5 ans	0.0002
10 ans	0.0304
15 ans	0.0524
20 ans	0.0793
40 ans	0.2647
60 ans	0.6048
80 ans	0.8670
100 ans	0.9691
150 ans	0.9997

Computation of the value function

Numerical results (true value : 4)

Number of points in quantization grids	Approximated value function	Monte Carlo value function
10	2.48	0.94
50	2.70	1.84
100	2.94	2.10
200	3.09	2.63
500	3.39	3.15
1000	3.56	3.43
2000	3.70	3.60
5000	3.82	3.73
8000	3.86	3.75

Comparison with theoretical optimal stopping time

[J. Risk Reliability 2012]
Optimal stopping time $\tau^* = \inf\{t \ : \ d_t \geq 0.02\}$

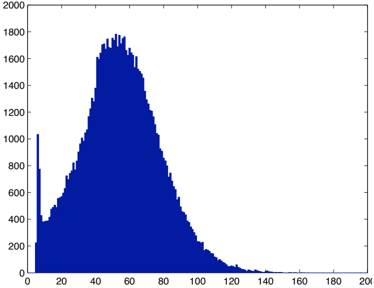

our stopping rule

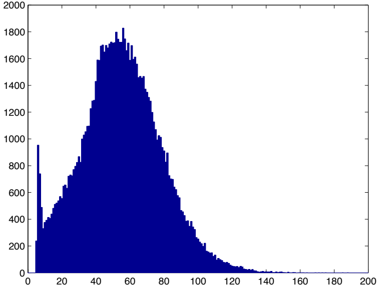

theoretical stopping time

Outline

Partial observations

Only a noisy observation of (X_t) is available

Optimal stopping problem

- compute the value function

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

  $\mathcal{M}$ set of all stopping times $\tau \leq T_N$ for natural filtration $(\mathcal{F}_t^{\mathcal{O}})$ of an observation process
- find an (ε-)optimal stopping time τ^* that reaches $V(x)(-\varepsilon)$

Observation process

- S_n perfectly observed
- Z_n observed through a noise

$$Y_n = \phi(Z_n) + W_n$$

- continuous time observation process

$$Y_t = \sum_{n=0}^{\infty} \mathbb{1}_{[T_n, T_{n+1}[}(t) Y_n$$

- filtration $\mathcal{F}_t^{\mathcal{O}} = \sigma(Y_s, s \leq t)$

From partial observations to complete observations

Methodology

- introduce the filter process $\Pi_t = \mathbb{E}[X_t \mid \mathcal{F}_t^{\mathcal{O}}]$
- transform the initial problem into a completely observed one for the filter process

Main drawbacks

- infinite dimension of the filter
- the new optimal stopping problem involves a process that is not a PDMP

Optimal stopping under partial observation

[Stoch. Proc. Appl. 2013]

Main steps

- Study the filter process $\Pi_n^i = \mathbb{P}[Z_n = x_i \mid \mathcal{F}_{T_n}^{\mathcal{O}}]$
  - recursive computation $\Pi_n = \Psi(\Pi_{n-1}, Y_n, S_n)$
  - Markov property
- Derive the dynamic programming equation for the new completely observed problem
- Replace (Π_n, S_n) by its quantized approximation
- Same kind of results for the convergence with rate

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 or partial repair

Simplified mathematical definition

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

- τ_n intervention times
- R_n new positions after intervention

Value function

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

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

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

Results

[Automatica 2012]

- Convergence results for the value function only
- more involved numerical scheme as the recurrence not he value functions does not yield an autonomous recurrence on random variables

References

[Stoch. Proc. Appl. 2013] A. Brandejsky, B. de Saporta, and F. Dufour. Optimal stopping for partially observed piecewise-deterministic markov processes. Stoch. Proc. Appl. 2013.

[Costa Davis 88] O. L. V. Costa and M. H. A. Davis. Approximations for optimal stopping of a piecewise-deterministic process. Math. Control Signals Systems 1988

[Costa Davis 89] O. L. V. Costa and M. H. A. Davis. Impulse control of piecewise-deterministic processes. Math. Control Signals Systems 1989

[Davis 93] M. H. A. Davis. Markov models and optimization, volume 49 of Monographs on Statistics and Applied Probability. Chapman & Hall, London, 1993.

[An. Apl. Proba. 2010] B. de Saporta, F. Dufour, and K. Gonzalez. Numerical method for optimal stopping of piecewise deterministic Markov processes. Ann. Appl. Proba. 2010.

[J. Risk Reliability 2012] B. de Saporta, F. Dufour, H. Zhang, and C. Elegbede. Optimal stopping for the predictive maintenance of a structure subject to corrosion. J. Risk and Reliability 2012.

[Automatica 2012] B. de Saporta and F. Dufour. Numerical method for impulse control of piecewise deterministic markov processes. Automatica 2012.

[Gugerli 86] U. S. Gugerli. Optimal stopping of a piecewise-deterministic Markov process. Stochastics 1986.

[Pagès 98] G. Pagès. A space quantization method for numerical integration. J. Comput. Appl. Math. 1998.

[Pagès, Pham, Printems 04] G. Pagès, H. Pham, and J. Printems. An optimal Markovian quantization algorithm for multi-dimensional stochastic control problems. Stoch. Dyn. 2004.