Numerical Approximation of Optimal Strategies for Impulse Control of Piecewise Deterministic Markov Processes
Application to Maintenance Optimisation

Benoîte de Saporta, François Dufour, Huilong Zhang

Univ. Montpellier, Bordeaux INP, Univ. Bordeaux

Outline

Maintenance optimization

Equipments

- with several components
- subject to random degradation and failures

Maintenance optimization problem: find some optimal balance between

- repairing/changing components too often
- do nothing and wait for the total failure of the system

Optimize some criterion

- minimize a cost: repair, maintenance, unavailability penalty, failure penalty, ...
- maximize a reward: availability, production, ...

Maintenance optimization

Equipments

- with several components
- subject to random degradation and failures

Maintenance optimization problem: find some optimal balance between

- repairing/changing components too often
- do nothing and wait for the total failure of the system

Optimize some criterion

- minimize a cost: repair, maintenance, unavailability penalty, failure penalty, ...
- maximize a reward: availability, production, ...

Impulse control problem

Impulse control

Select

- intervention dates
- new starting point for the process at interventions

to minimize a cost function

Piecewise deterministic Markov processes

General class of non-diffusion dynamic stochastic hybrid models: deterministic motion punctuated by random jumps.
[CD 89], [Davis 93], [dSDZ 14], ...

Piecewise deterministic Markov processes

Starting point

$$X_0 = (m, x)$$

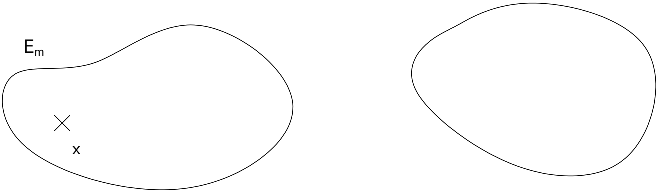

Piecewise deterministic Markov processes

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

$$X_t = \big(m, \phi_m(x,t)\big), \quad \mathbb{P}_{(m,x)}(T_1 > t) = \mathrm{e}^{-\int_0^t \lambda_m\left(\phi_m(x,s)\right) ds}$$

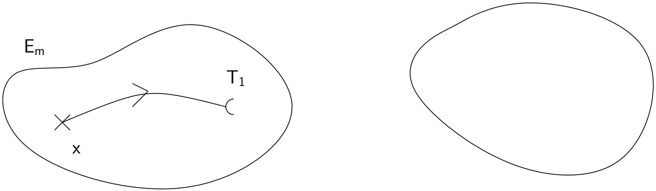

Piecewise deterministic Markov processes

Post-jump location (m_1, x_{T_1}) selected by the Markov kernel

$$Q_m\big(\phi_m(x, T_1), \cdot\big)$$

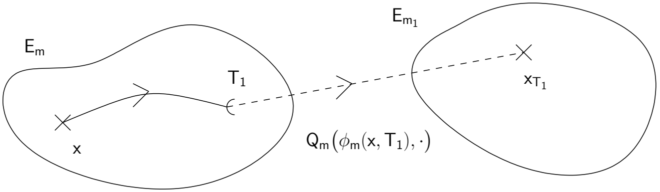

Piecewise deterministic Markov processes

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

$$X_{T_1+t} = \left(m_1, \phi_{m_1}(x_{T_1}, t)\right), \quad t < S_2$$

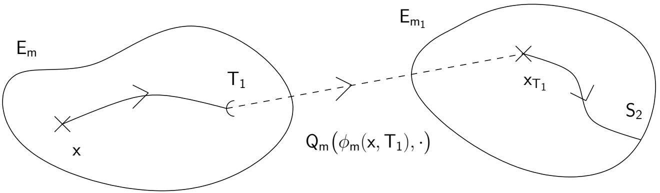

Piecewise deterministic Markov processes

Post-jump location (m_2, x_{T_2}) selected by Markov kernel

$$Q_{m_1}\left(\phi_{m_1}(x_{T_1}, S_2), \cdot\right) \ldots$$

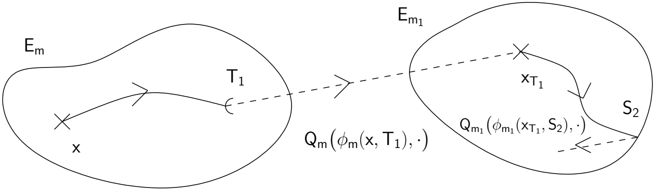

Embedded Markov chain

$\{X_t\}$ 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

Mathematical definition of impulse control

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

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

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

dS, Dufour, Geeraert, 2017

Construction of ϵ-optimal strategies based on the dynamic programming operator

Dynamic programming

Jump-or-intervention operator

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

with

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

Approximation scheme - Value function

Based on time-dependent discretizations of the state space of (Z_n, S_n)

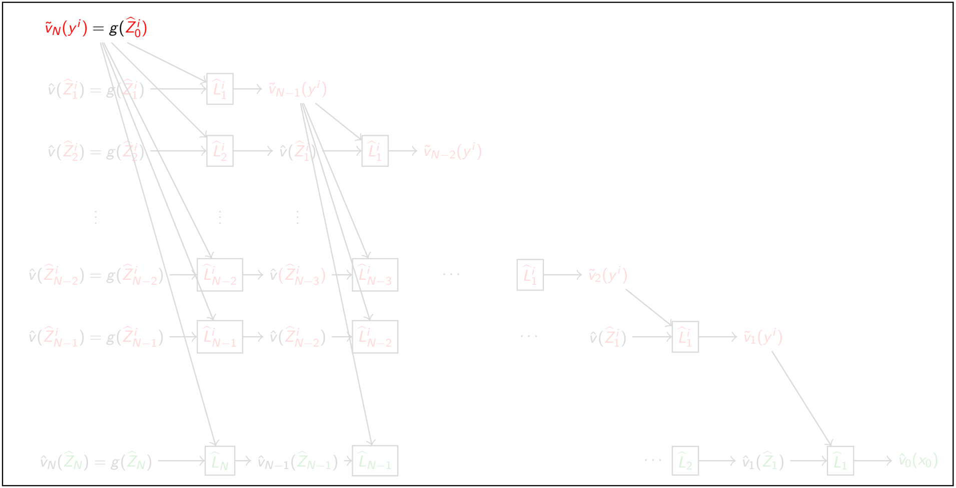

Approximation scheme - Value function

Based on time-dependent discretizations of the state space of (Z_n, S_n)

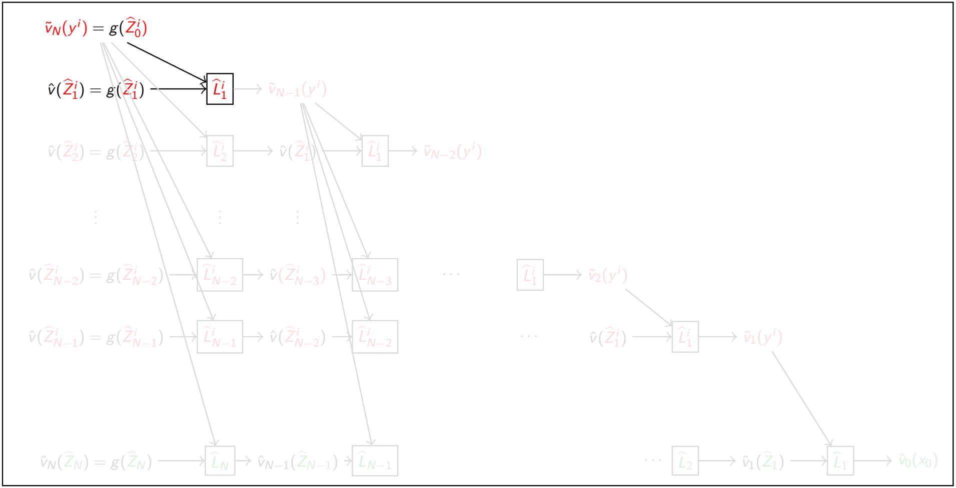

Approximation scheme - Value function

Based on time-dependent discretizations of the state space of (Z_n, S_n)

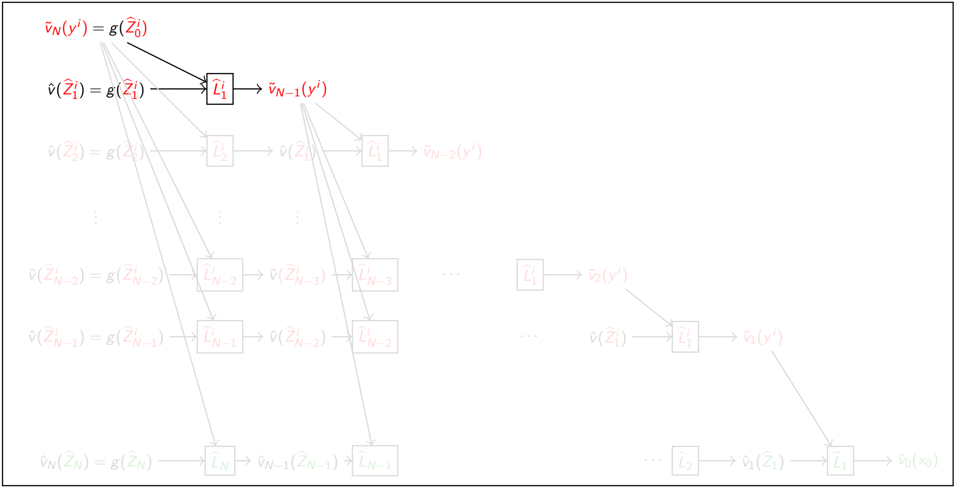

Approximation scheme - Value function

Based on time-dependent discretizations of the state space of (Z_n, S_n)

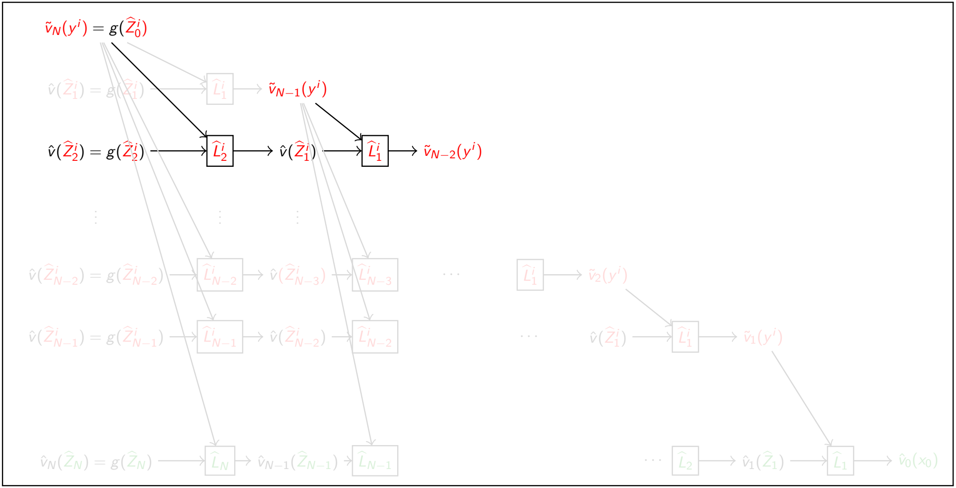

Approximation scheme - Value function

Based on time-dependent discretizations of the state space of (Z_n, S_n)

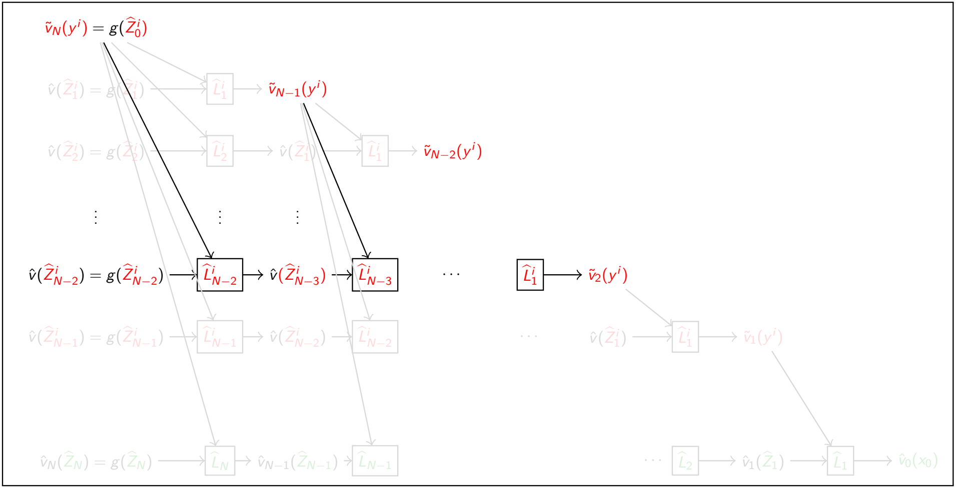

Approximation scheme - Value function

Based on time-dependent discretizations of the state space of (Z_n, S_n)

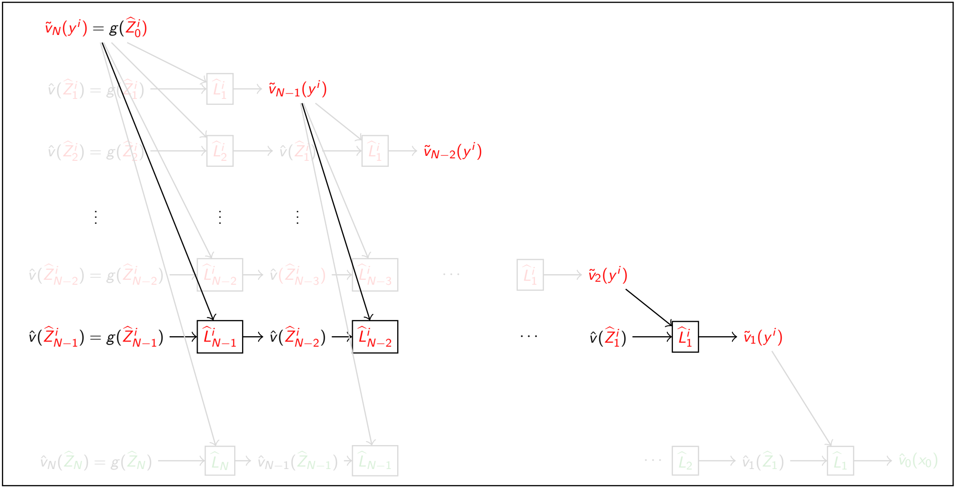

Approximation scheme - Value function

Based on time-dependent discretizations of the state space of (Z_n, S_n)

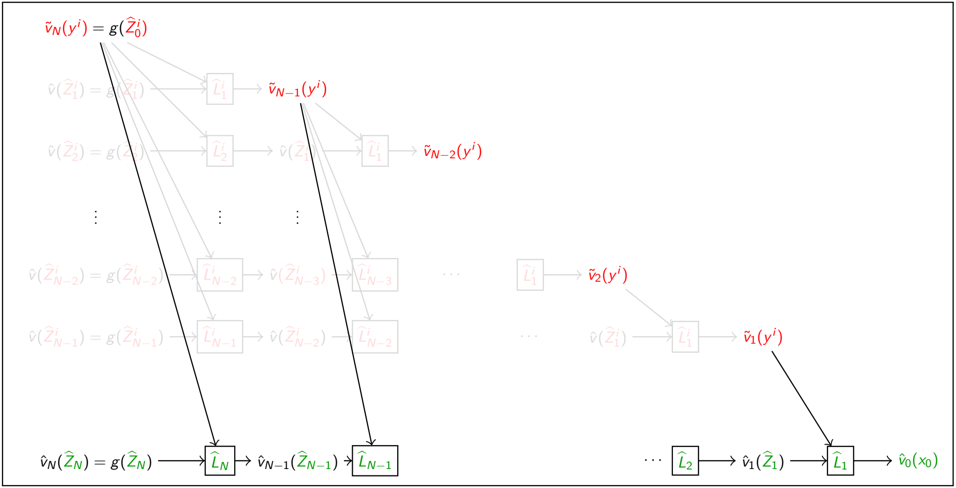

Approximation scheme - ϵ-optimal strategy

Based on time-dependent discretizations of the state space of (Z_n, S_n)

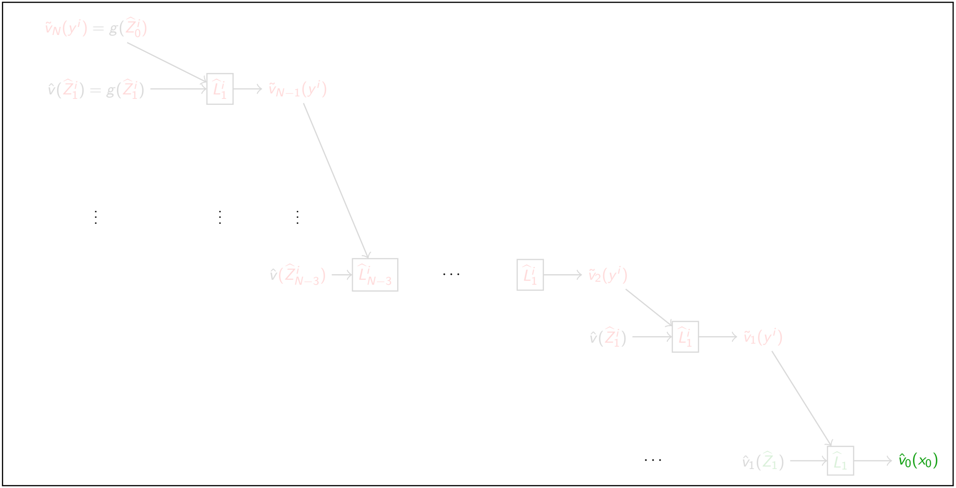

Approximation scheme - ϵ-optimal strategy

Based on time-dependent discretizations of the state space of (Z_n, S_n)

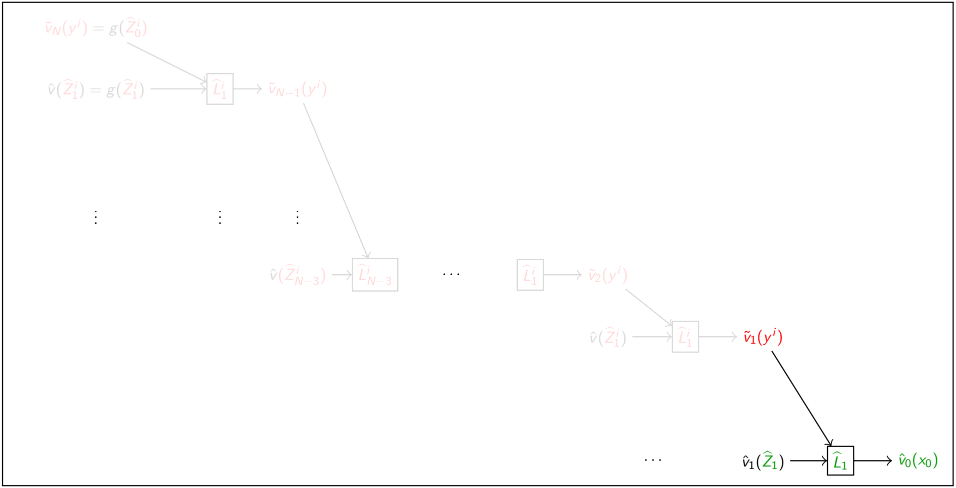

Approximation scheme - ϵ-optimal strategy

Based on time-dependent discretizations of the state space of (Z_n, S_n)

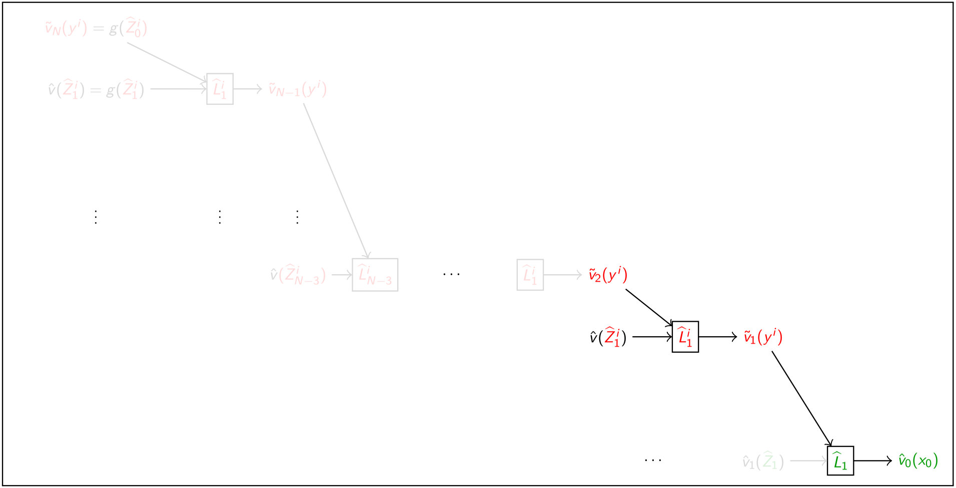

Approximation scheme - ϵ-optimal strategy

Based on time-dependent discretizations of the state space of (Z_n, S_n)

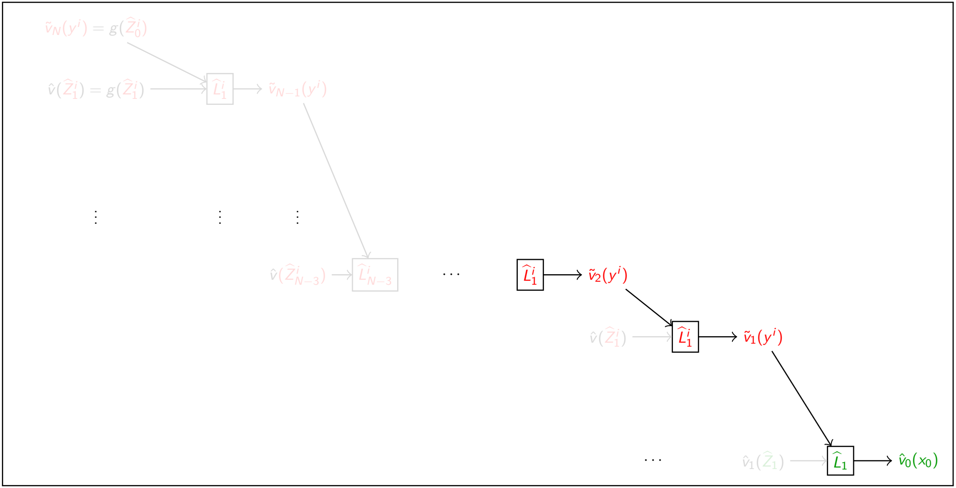

Approximation scheme - ϵ-optimal strategy

Based on time-dependent discretizations of the state space of (Z_n, S_n)

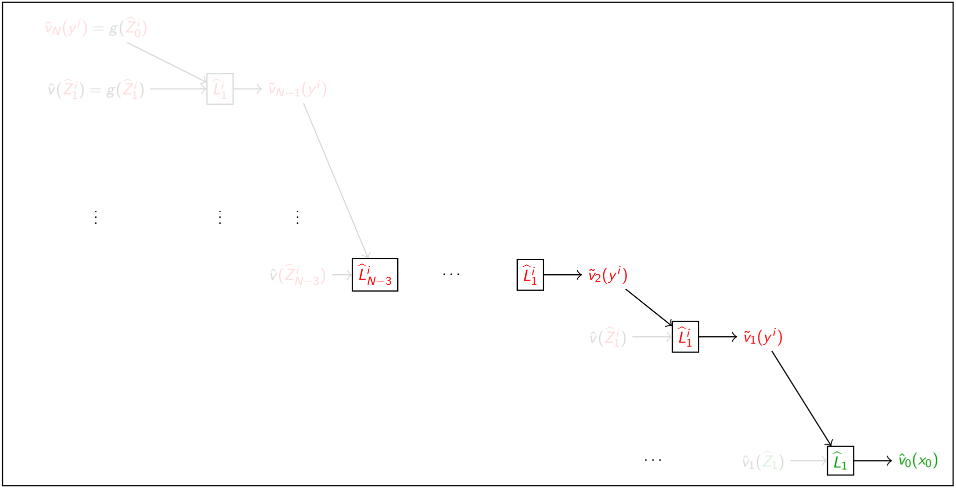

Approximation scheme - ϵ-optimal strategy

Based on time-dependent discretizations of the state space of (Z_n, S_n)

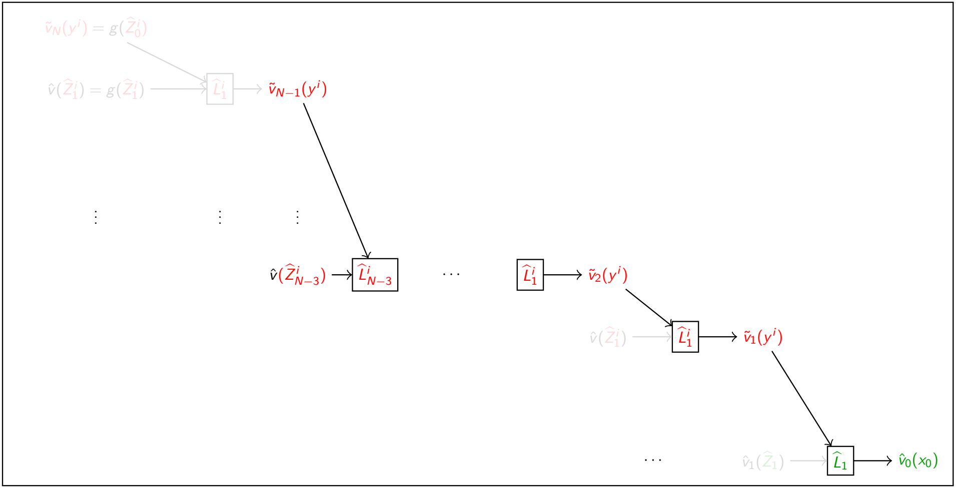

Approximation scheme - ϵ-optimal strategy

Based on time-dependent discretizations of the state space of (Z_n, S_n)

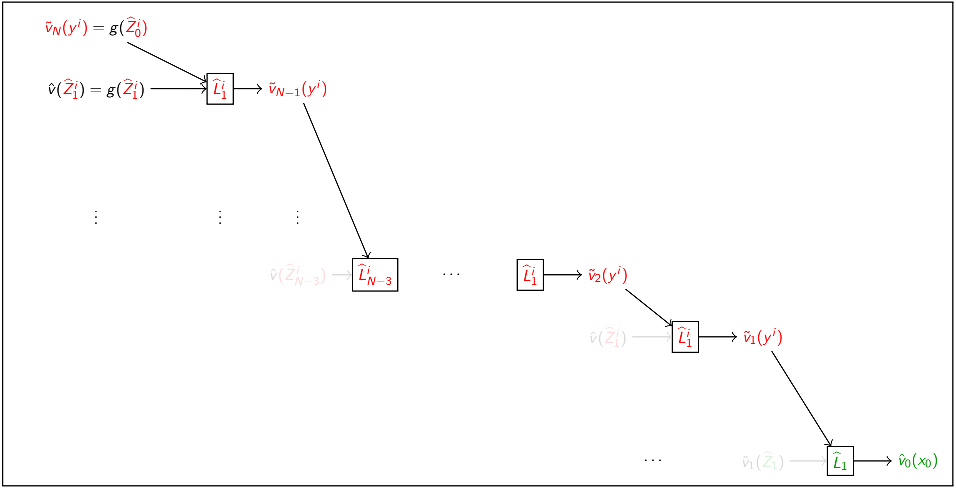

Approximation scheme - ϵ-optimal strategy

Based on time-dependent discretizations of the state space of (Z_n, S_n)

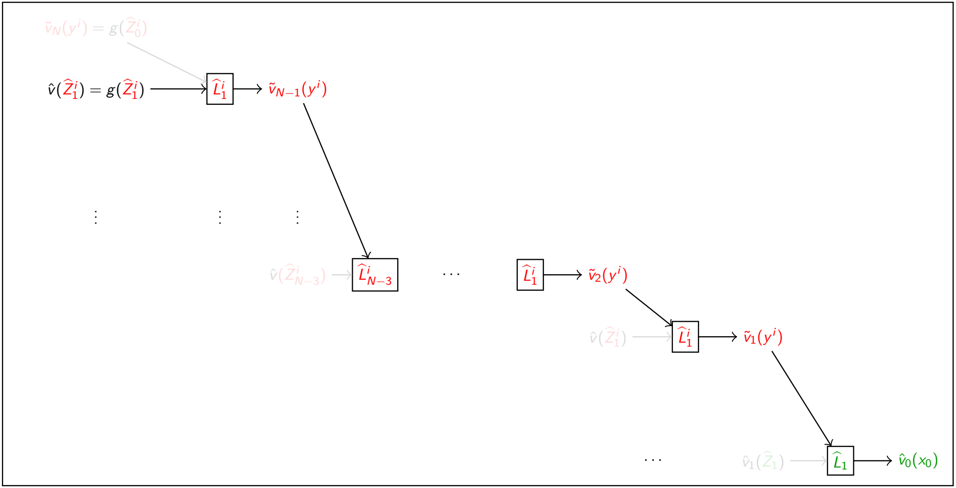

Equipment model

Typical model with 4 components

- Component 1: 2 states – stable $\xrightarrow{\text{Exponential}}$ failed
- Component 2: 2 states – stable $\xrightarrow{\text{Weibull}}$ failed
- Components 3 and 4: 3 states

stable $\xrightarrow{\text{Weibull}}$ degraded $\xrightarrow{\text{Exponential}}$ failed

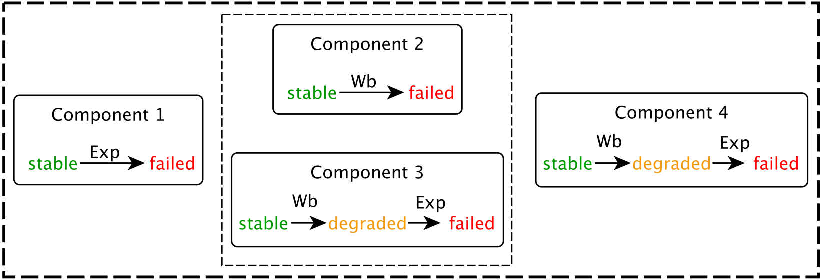

Maintenance operations

Possible maintenance operations

- All components, all states: do nothing
- Components 1 and 2, all states: change
- Components 3 and 4: change in all states, repair only in stable or degraded states

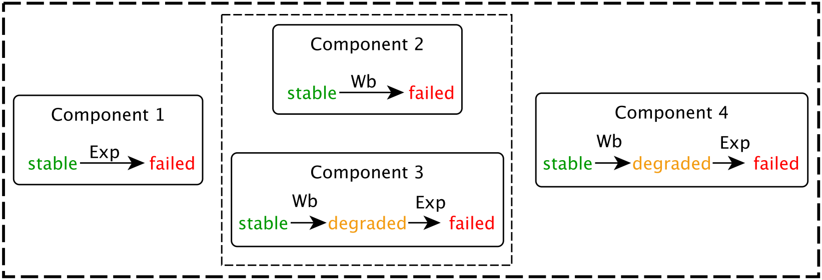

Criterion to optimize

Minimize the maintenance + unavailability costs

- unavailability cost proportional to time spend in failed state
- fixed cost for going to the workshop + repair < change costs

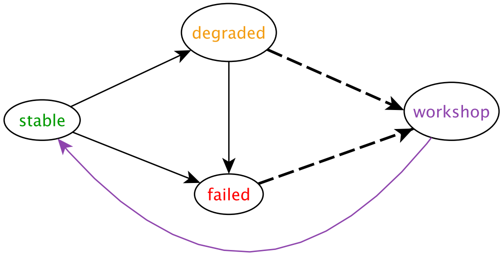

PDMP model of the equipment

- Euclidean variables: 5 time variables
  - functioning time of components 2, 3 and 4
  - calendar time
  - time spent in the workshop
- Discrete variables: 225 modes
  - state of the components / maintenance operations

Parameters to tune

- Number of points in the control grid (underlying continuous model)
- Number of point in the quantization grids for (Z_n, S_n)
- Approximation horizon N such that $v_N(x) - \mathcal{V}(x)$ small enough $\simeq$ allowed number of jumps + interventions
- bounding function g
- Time discretization step for inf

Step 1: Exact simulation of the PDMP

Implementation of an exact simulator for reference strategies to serve as benchmark

- Strategy 1: do nothing
- Strategy 2: send equipment to workshop 1 day after failure, change all degraded components, change all failed ones
- Strategy 3: send equipment to workshop 1 day after degradation, change all degraded components, change all failed ones

Strategy	1	2	3
Mean cost	19952	11389	8477

Step 1: Exact simulation of the PDMP

Implementation of an exact simulator for reference strategies to serve as benchmark

- Strategy 1: do nothing
- Strategy 2: send equipment to workshop 1 day after failure, change all degraded components, change all failed ones
- Strategy 3: send equipment to workshop 1 day after degradation, change all degraded components, change all failed ones

Strategy	1	2	3
Mean cost	19952	11389	8477

Step 2 and 3: Discretisation of the control set $\mathbb{U}$ and te embedded Markov chain

Finite control set $\mathbb{U}$

$\Longrightarrow$ discretize the functioning times at interventions

$\Longrightarrow$ project the real times on the grid feasibly

Compromise between precision and computation time

Tests on strategy 3

Grid	Number of points	relative error
$3 \times 3 \times 3 \times 5$	246	0.10344
$4 \times 4 \times 4 \times 5$	331	0.0241
$5 \times 5 \times 5 \times 5$	592	0.0062
$3 \times 3 \times 3 \times 11$	615	0.0341
$4 \times 4 \times 4 \times 11$	923	0.0819
$5 \times 5 \times 5 \times 11$	1855	0.0186
$6 \times 6 \times 6 \times 11$	2110	0.0066
$7 \times 7 \times 7 \times 11$	2617	0.0071
$8 \times 8 \times 8 \times 11$	3359	0.0066
$3 \times 3 \times 3 \times 21$	1230	0.0034
$4 \times 4 \times 4 \times 21$	1899	0.0170
$5 \times 5 \times 5 \times 21$	2960	0.0095
$6 \times 6 \times 6 \times 21$	4220	0.0065
$7 \times 7 \times 7 \times 21$	5536	0.0059
$8 \times 8 \times 8 \times 21$	7111	0.0047

Step 4: Calibrating N the number of allowed jumps + interventions

Horizon N (number of iterations)

- 5 for Strategy 1
- up to 30 for Strategy 2 (mean 6)
- up to 25 for Strategy 3 (mean 6)

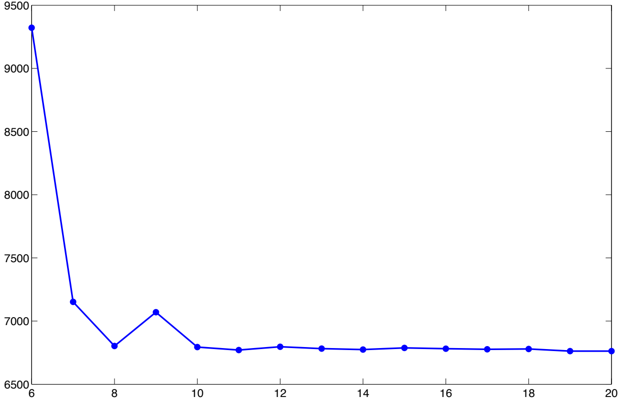

Step 5: Approximation of the value function

Strategy 1	Strategy 2	Strategy 3	Approx. Value function
19952	11389	8477	7076

- relative gain of 19.8% vs Strategy 5

Step 6: Optimally controlled trajectories

Strategy 1	Strategy 2	Strategy 3	Approx. Value function	Optimally controlled traj.
19952	11389	8477	7076	6733

- numerical validation of the algorithm with various starting points: consistent results

Conclusion

Numerical method to derive a feasible ϵ-optimal strategy

- rigorously validated [dSD 12, dSDG 17]
- with general error bounds for the approximation of the value function
- numerically demanding but viable in low dimensional examples

References

[CD 89] O. Costa, M. Davis *Impulse control of piecewise-deterministic processes*
[Davis 93] M. Davis, *Markov models and optimization*
[dSD 12] B. de Saporta, F. Dufour *Numerical method for impulse control of piecewise deterministic Markov processes*
[dSDG 17] B. de Saporta, F. Dufour, A. Geeraert *Optimal strategies for impulse control of piecewise deterministic Markov processes*
[dSDZ 14] B. de Saporta, F. Dufour, H. Zhang *Numerical methods for simulation and optimization of PDMPs: application to reliability*
[P 98] G. Pagès *A space quantization method for numerical integration*
[PPP 04] G. Pagès, H. Pham, J. Printems *An optimal Markovian quantization algorithm for multi-dimensional stochastic control problems*