Stochastic Gradient Method: Applications

February 03, 2015

Lecture Outline

A Basic Two-Stage Recourse Problem

We consider the management of a water reservoir. Water is drawn from the reservoir by way of random consumers. In order to ensure the reservoir supply, 2 decisions are taken at successive time steps.

- A first supply decision q_1 is taken without any knowledge of the effective consumption, the associated cost being equal to $\frac{1}{2}\, c_1 \big(q_1\big)^2$, with $c_1 > 0$.
- Once the consumption d has been observed (realization of a r.v. $\boldsymbol{D}$ defined over a probability space $(\Omega, \mathcal{A}, \mathbb{P})$), a second supply decision q_2 is taken in order to maintain the reservoir at its initial level, that is, $q_2 = d - q_1$, the cost associated to this second decision being equal to $\frac{1}{2}\, c_2 \big(q_2\big)^2$, with $c_2 > c_1 > 0$.

The problem is to minimize the expected cost of operation.

Mathematical Formulation and Solution

Problem Formulation

- q_1 is a deterministic decision variable,
- whereas q_2 is the realization of a random variable $\boldsymbol{Q}_2$.

$$\min_{(q_1, \boldsymbol{Q}_2)} \frac{1}{2} c_1 \big(q_1\big)^2 + \frac{1}{2} \mathbb{E}\Big(c_2 \big(\boldsymbol{Q}_2\big)^2\Big) \quad \text{s.t.} \quad q_1 + \boldsymbol{Q}_2 = \boldsymbol{D} \ .$$

Equivalent Problem

$$\min_{q_1 \in \mathbb{R}} \frac{1}{2} \mathbb{E}\Big(c_1 \big(q_1\big)^2 + c_2 \big(\boldsymbol{D} - q_1\big)^2\Big)$$

Analytical solution: $q_1^\sharp = \dfrac{c_2}{c_1 + c_2} \mathbb{E}\big(\boldsymbol{D}\big)$.

Stochastic Gradient Algorithm

$$\boldsymbol{Q}_1^{(k+1)} = \boldsymbol{Q}_1^{(k)} - \frac{1}{k}\Big(\big(c_1 + c_2\big)\boldsymbol{Q}_1^{(k)} - c_2\boldsymbol{D}^{(k+1)}\Big) .$$

Algorithm (initialization)

```
//
// Random generator
//
   rand('normal'); rand('seed',123);
//
// Random consumption
//
   moy = 10.; ect = 5.;
//
// Criterion
//
   c1 = 3.; c2 = 1.;
//
// Initialization
//
   x = [ ]; y = [ ];
```

Algorithm (iterations)

```
//
// Algorithm
//
   qk = 0.;
   for k = 1:100
      dk = moy + (ect*rand(1));
      gk = ((c1+c2)*qk) - (c2*dk);
      ek = 1/k;
      qk = qk - (ek*gk);
      x = [x ; k]; y = [y ; qk];
   end
//
// Trajectory plot
//
   plot2d(x,y);
   xtitle('Stochastic Gradient ','Iter.','Q1');
```

A Realization of the Algorithm

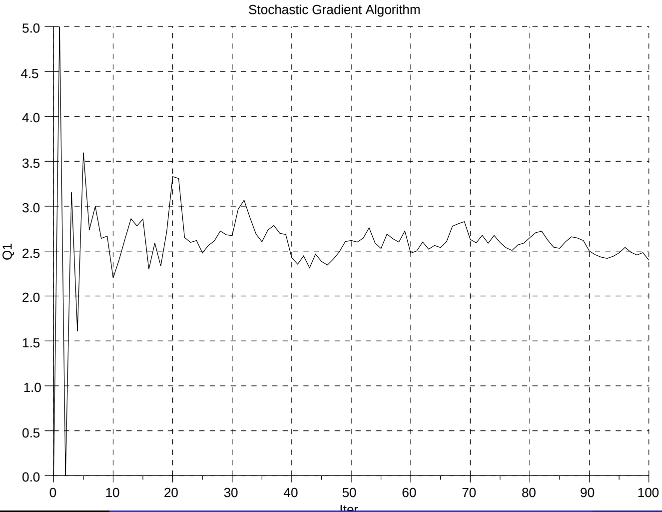

More Realizations...

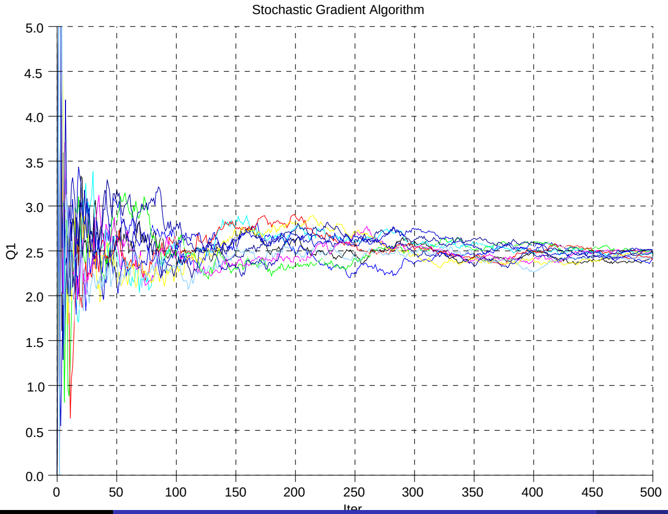

Slight Modification of the problem

As in the basic two-stage recourse problem,

- a first supply decision q_1 is taken without any knowledge of the effective consumption, the associated cost being equal to $\frac{1}{2}\, c_1 \big(q_1\big)^2$,
- a second supply decision q_2 is taken once the consumption d has been observed (realization of a r.v. $\boldsymbol{D}$), the cost of this second decision being equal to $\frac{1}{2}\, c_2 \big(q_2\big)^2$.

The difference between supply and demand is penalized thanks to an additional cost term $\frac{1}{2}\, c_3 \big(q_1 + q_2 - d\big)^2$. The new problem is :

$$\min_{(q_1, \boldsymbol{Q}_2)} \frac{1}{2} \mathbb{E}\Big(c_1 \big(q_1\big)^2 + c_2 \big(\boldsymbol{Q}_2\big)^2 + c_3 \big(q_1 + \boldsymbol{Q}_2 - \boldsymbol{D}\big)^2 \Big) .$$

Question: how to solve it using a stochastic gradient algorithm?

Trade-off Between Investment and Operation (1)

A company owns N production units and has to meet a given (non stochastic) demand d.

- For each unit i, the decision maker first takes an investment decision $u_i \in \mathbb{R}$, the associated cost being $I_i(u_i)$.
- Then a discrete disturbance $w_i \in \{w_{i,a}, w_{i,b}, w_{i,c}\}$ occurs.
- Knowing all noises, the decision maker selects for each unit i an operating point $v_i \in \mathbb{R}$, which leads to a cost $c_i(v_i, w_i)$ and a production $e_i(v_i, w_i)$.

The goal is to minimize the overall expected cost, subject to the following constraints:

- investment limitation: $\Theta(u_1, \ldots, u_N) \leq 0$,
- operating limitation: $v_i \leq \varphi_i(u_i) \; , \quad i = 1 \ldots, N$,
- demand satisfaction: $\sum_{i=1}^{N} e_i(v_i, w_i) - d = 0$.

Trade-off Between Investment and Operation (2)

Questions.

1. Write down the optimization problem. Is it possible to apply the stochastic gradient algorithm in a straightforward manner?
2. Extract the optimization subproblem obtained when both the investment $u = (u_1, \ldots, u_N)$ and the noise $w = (w_1, \ldots, w_N)$ are fixed. The value of this subproblem is denoted $f^\sharp(u, w)$.
   - Discuss the resolution of this subproblem.
   - Give assumptions in order to have a "nice" function $f^\sharp$.
   - Compute the partial derivatives of $f^\sharp$ w.r.t. u.
3. Reformulate the initial optimization problem using this new function $f^\sharp$ and apply the stochastic gradient algorithm in the two following cases:
   - the investment limitation reduces to $u_i \in [\underline{u}_i, \overline{u}_i]$, $i = 1, \ldots, N$,
   - the investment limitation has the form $\Theta(u_1, \ldots, u_N) \leq 0$.

The Financial Problem

The price of an option with payoff $\psi(\boldsymbol{S}_t, 0 \leq t \leq T)$ is given by

$$P = \mathbb{E}\Big(\mathrm{e}^{-rT}\psi(\boldsymbol{S}_t, 0 \leq t \leq T)\Big) ,$$

where the dynamics of the underlying n-dimensional asset $\boldsymbol{S}$ is described by the following stochastic differential equation

$$\mathrm{d}\boldsymbol{S}_t = \boldsymbol{S}_t\big(r\mathrm{d}t + \sigma(t, \boldsymbol{S}_t)\mathrm{d}\boldsymbol{W}_t\big) , \quad \boldsymbol{S}_0 = x ,$$

r being the interest rate and $\sigma(t, y)$ being the volatility function.

Discretization

Most of the time, the exact solution is not available. To overcome the difficulty, one considers a discretized approximation of $\boldsymbol{S}$ (by using an Euler's scheme), so that the price P is approximated by

$$\widehat{P} = \mathbb{E}\Big(\mathrm{e}^{-rT}\psi(\widehat{\boldsymbol{S}}_{t_1},\ldots,\widehat{\boldsymbol{S}}_{t_d})\Big) .$$

In such cases, the discretized function can be expressed in terms of the Brownian increments, or equivalently using a random normal vector. A compact form for the discretized price is

$$\widehat{P} = \mathbb{E}\big(\phi(\boldsymbol{G})\big) ,$$

where $\boldsymbol{G}$ is a $n \times d$-dimensional Gaussian vector with identity covariance matrix and zero-mean.

A Parameterized Change of Variable

The problem is now to compute $\widehat{P} = \mathbb{E}\big(\phi(\boldsymbol{G})\big)$ using Monte Carlo simulations. By a change of variables, we obtain that

$$\widehat{P} = \mathbb{E}\Big(\phi(\boldsymbol{G} + \theta)\mathrm{e}^{-\langle \theta \,, \boldsymbol{G} \rangle - \frac{\|\theta\|^2}{2}}\Big) \ ,$$

for any $\theta \in \mathbb{R}^{n \times d}$. Let us denote by $\widehat{V}(\theta)$ the associated variance:

$$\begin{aligned} \widehat{V}(\theta) &= \mathbb{E}\Big(\phi(\boldsymbol{G} + \theta)^2\mathrm{e}^{-2\langle \theta \,, \boldsymbol{G} \rangle - \|\theta\|^2}\Big) - \mathbb{E}\Big(\phi(\boldsymbol{G})\Big)^2 \ , \\ &= \mathbb{E}\Big(\phi(\boldsymbol{G})^2\mathrm{e}^{-\langle \theta \,, \boldsymbol{G} \rangle + \frac{\|\theta\|^2}{2}}\Big) - \mathbb{E}\Big(\phi(\boldsymbol{G})\Big)^2 \ . \end{aligned}$$

The last expression shows that function $\widehat{V}$ is strictly convex and differentiable without any specific assumptions on ϕ. Moreover,

$$\nabla \widehat{V}(\theta) = \mathbb{E}\Big((\theta - \boldsymbol{G})\phi(\boldsymbol{G})^2\mathrm{e}^{-\langle \theta \,, \boldsymbol{G} \rangle + \frac{\|\theta\|^2}{2}}\Big) \ ,$$

so that the gradient of $\widehat{V}$ does not involves any derivative of ϕ.

Variance Minimization

The goal is to compute the parameter θ such that the variance $\widehat{V}(\theta)$ related to $\widehat{P}$ is as small as possible, in order to optimize the convergence speed of the Monte Carlo simulations:

$$\min_{\theta\in\mathbb{R}^d} \mathbb{E}\Big(\phi(\boldsymbol{G})^2 \mathrm{e}^{-\langle\theta\,,\boldsymbol{G}\rangle+\frac{\|\theta\|^2}{2}}\Big)\ .$$

This optimization problem is well suited for being solved by the stochastic gradient algorithm:

$$\boldsymbol{\theta}^{(k+1)} = \boldsymbol{\theta}^{(k)} - \epsilon^{(k)}\big(\boldsymbol{\theta}^{(k)} - \boldsymbol{G}^{(k+1)}\big)\phi\big(\boldsymbol{G}^{(k+1)}\big)^2 \mathrm{e}^{-\langle\boldsymbol{\theta}^{(k)}\,,\boldsymbol{G}^{(k+1)}\rangle+\frac{\|\boldsymbol{\theta}^{(k)}\|^2}{2}}\ ,$$

and its unique solution is denoted $\theta^\sharp$.

Non adaptive Algorithm

1. Using a N-sample of $\boldsymbol{G}$, obtain an approximation $\boldsymbol{\theta}^{(N)}$ of $\theta^\sharp$ by N iterations of the stochastic gradient algorithm.

2. Once $\boldsymbol{\theta}^{(N)}$ has been obtained, use the standard Monte Carlo method to compute an approximation of the price $\widehat{P}$ by using another N-sample of $\boldsymbol{G}$:

$$\widehat{\boldsymbol{P}}^{(N)} = \frac{1}{N}\sum_{k=1}^{N} \phi(\boldsymbol{G}^{(N+k)} + \boldsymbol{\theta}^{(N)})\mathrm{e}^{-\langle \boldsymbol{\theta}^{(N)}, \boldsymbol{G}^{(N+k)}\rangle - \frac{\|\boldsymbol{\theta}^{(N)}\|^2}{2}} .$$

This algorithm requires $2N$ evaluations of the function ϕ, whereas a crude Monte Carlo method evaluates ϕ only N times. The non adaptive algorithm is efficient as soon as $\widehat{V}(\theta^\sharp) \leq \widehat{V}(0)/2$.

Adaptive Algorithm

Combine the 2 previous algorithms, and compute simultaneously approximations of $\theta^\sharp$ and $\widehat{P}$ by using the same N-sample of $\boldsymbol{G}$:[10]

$$\boldsymbol{\theta}^{(k+1)} = \boldsymbol{\theta}^{(k)} - \epsilon^{(k)}\big(\boldsymbol{\theta}^{(k)} - \boldsymbol{G}^{(k+1)}\big)\phi\big(\boldsymbol{G}^{(k+1)}\big)^2 \mathrm{e}^{-\langle \boldsymbol{\theta}^{(k)}, \boldsymbol{G}^{(k+1)}\rangle + \frac{\|\boldsymbol{\theta}^{(k)}\|^2}{2}} ,$$

$$\widehat{\boldsymbol{P}}^{(k+1)} = \widehat{\boldsymbol{P}}^{(k)} - \frac{1}{k+1}\Big(\boldsymbol{P}^{(k)} - \phi\big(\boldsymbol{G}^{(k+1)} + \boldsymbol{\theta}^{(k)}\big)\,\mathrm{e}^{-\langle \boldsymbol{\theta}^{(k)}, \boldsymbol{G}^{(k+1)}\rangle - \frac{\|\boldsymbol{\theta}^{(k)}\|^2}{2}}\Big) .$$

A Central Limit Theorem is available for this algorithm:

$$\sqrt{N}\Big(\widehat{\boldsymbol{P}}^{(N)} - \widehat{P}\Big) \xrightarrow{\mathcal{D}} \mathcal{N}\Big(0, \widehat{V}(\theta^\sharp)\Big) .$$

[10] the last relation being the recursive form of:

$$\widehat{\boldsymbol{P}}^{(N)} = \frac{1}{N}\sum_{k=1}^{N} \phi\big(\boldsymbol{G}^{(k+1)} + \boldsymbol{\theta}^{(k)}\big)\,\mathrm{e}^{-\langle \boldsymbol{\theta}^{(k)}, \boldsymbol{G}^{(k+1)}\rangle - \frac{\|\boldsymbol{\theta}^{(k)}\|^2}{2}} .$$

Mission to Mars

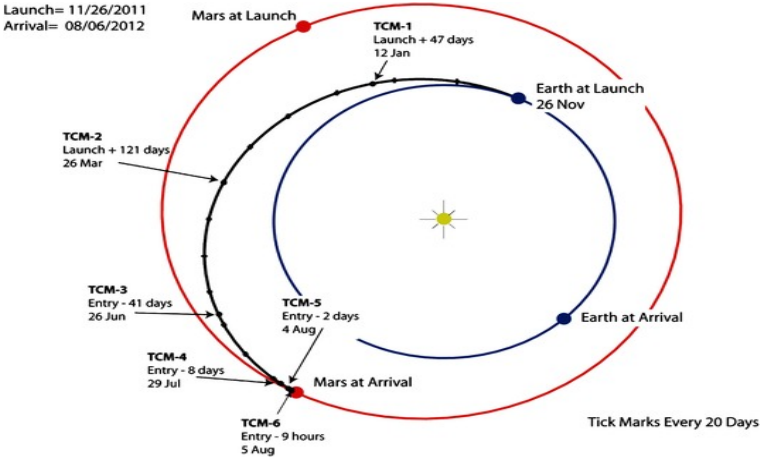

Satellite Model

$$\frac{\mathrm{d}r}{\mathrm{d}t} = v \, , \quad \frac{\mathrm{d}v}{\mathrm{d}t} = -\mu \frac{r}{\|r\|^3} + \frac{F}{m}\kappa \, , \tag{8a}$$

$$\frac{\mathrm{d}m}{\mathrm{d}t} = -\frac{T}{g_0 I_{\mathrm{sp}}}\delta \, . \tag{8b}$$

(8a): 6-dimensional state vector (position r and velocity v).
(8b): 1-dimensional state vector (mass m including fuel).

κ involves the direction cosines of the thrust and the on-off switch δ of the engine (3 controls), and $\mu, F, T, g_0, I_{\mathrm{sp}}$ are constants.

The deterministic control problem is to drive the satellite from the initial condition at t_{i} to a known final position r_{f} and velocity v_{f} at t_{f} (given) while minimizing fuel consumption $m(t_{\mathrm{i}}) - m(t_{\mathrm{f}})$.

Deterministic Optimization Problem

Using *equinoctial coordinates* for the position and velocity

$\rightsquigarrow$ state vector $x \in \mathbb{R}^7$,

and *cartesian coordinates* for the thrust of the engine

$\rightsquigarrow$ control vector $u \in \mathbb{R}^3$,

the deterministic optimization problem writes as follows:

$$\min_{u(\cdot)} K\big(x(t_\mathrm{f})\big)$$

subject to:

$$x(t_\mathrm{i}) = x_\mathrm{i} \;, \quad \overset{\bullet}{x}(t) = f\big(x(t), u(t)\big) \;,$$

$$\|u(t)\| \leq 1 \quad \forall t \in [t_\mathrm{i}, t_\mathrm{f}] \;,$$

$$C\big(x(t_\mathrm{f})\big) = 0 \;.$$

Engine Failure

- Sometimes, the engine may fail to work when needed: the satellite drifts away from the deterministic optimal trajectory. After the engine control is recovered, it is not always possible to drive the satellite to the final target at t_f.
- By anticipating such possible failures and by modifying the trajectory followed before any such failure occurs, one may increase the possibility of eventually reaching the target.
- But such a deviation from the deterministic optimal trajectory results in a deterioration of the economic performance.
- The problem is thus to balance the increased probability of eventually reaching the target despite possible failures against the expected economic performance, that is, to quantify the price of safety one is ready to pay for.

Stochastic Formulation (1)

A failure is modeled using two random variables:

- $\mathbf{t}_{\rm p}$: random initial time of the failure,
- $\mathbf{t}_{\rm d}$: random duration of the failure.

For any realization $(t_{\rm p}^{\xi}, t_{\rm d}^{\xi})$ of a failure:

1. $u(\cdot)$ denotes the control used prior to the failure
   $\rightsquigarrow$ u is defined over $[t_{\rm i}, t_{\rm f}]$ but implemented over $[t_{\rm i}, t_{\rm p}^{\xi}]$
   and corresponds to an open-loop control,
2. the control during the failure is 0 over $[t_{\rm p}^{\xi}, t_{\rm p}^{\xi} + t_{\rm d}^{\xi}]$,
3. $v^{\xi}(\cdot)$ denotes the control used after the failure
   $\rightsquigarrow$ v^{ξ} is defined over $[t_{\rm p}^{\xi} + t_{\rm d}^{\xi}, t_{\rm f}]$ (if nonempty)
   and corresponds to a closed-loop strategy $\mathbf{V}$.

The satellite dynamics in the stochastic formulation writes:

$$x^{\xi}(t_{\rm i}) = x_{\rm i} \ , \quad \dot{x}^{\xi}(t) = f^{\xi}\big(x^{\xi}(t), u(t), v^{\xi}(t)\big) \ .$$

Stochastic Formulation (2)

The problem is to minimize the expected cost (fuel consumption)

- w.r.t. the open-loop control u and the closed-loop strategy $\mathbf{V}$,
- the probability to hit the target at t_{f} being at least equal to p.

$$\min_{u(\cdot)} \min_{\mathbf{V}(\cdot)} \mathbb{E}\Big(K\big(x^{\xi}(t_{\mathrm{f}})\big)\Big)$$

$$\min_{u(\cdot)} \min_{\mathbf{V}(\cdot)} \mathbb{E}\Big(K\big(x^{\xi}(t_{\mathrm{f}})\big) \;\Big|\; C\big(x^{\xi}(t_{\mathrm{f}})\big) = 0\Big)$$

subject to:

$$x^{\xi}(t_{\mathrm{i}}) = x_{\mathrm{i}} \;, \quad \dot{x}^{\xi}(t) = f^{\xi}\big(x^{\xi}(t), u(t), v^{\xi}(t)\big) \;,$$

$$\|u(t)\| \leq 1 \quad \forall t \in [t_{\mathrm{i}}, t_{\mathrm{f}}] \;, \quad \|v^{\xi}(t)\| \leq 1 \quad \forall t \in [t_{\mathrm{p}}^{\xi} + t_{\mathrm{d}}^{\xi}, t_{\mathrm{f}}] \;,$$

$$\mathbb{P}\Big(C\big(x^{\xi}(t_{\mathrm{f}})\big) = 0\Big) \geq p \;.$$

Indicator Function

Consider the real-valued indicator function:

$$\mathbb{I}(y) = \begin{cases} 1 & \text{if } y = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Then

$$\mathbb{P}\Big(C\big(x^{\xi}(t_{\mathrm{f}})\big) = 0\Big) = \mathbb{E}\Big(\mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big) ,$$

and

$$\mathbb{E}\Big(K\big(x^{\xi}(t_{\mathrm{f}})\big) \;\Big|\; C\big(x^{\xi}(t_{\mathrm{f}})\big) = 0\Big) = \frac{\mathbb{E}\Big(K\big(x^{\xi}(t_{\mathrm{f}})\big) \times \mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big)}{\mathbb{E}\Big(\mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big)} .$$

Problem Reformulation

The problem is (shortly) reformulated as

$$\min_{u(\cdot)} \min_{\mathbf{V}(\cdot)} \frac{\mathbb{E}\Big(K\big(x^{\xi}(t_{\mathrm{f}})\big) \times \mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big)}{\mathbb{E}\Big(\mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big)}$$

$$\text{s.t.} \quad \mathbb{E}\Big(\mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big) \geq p \, .$$

Such a formulation is however not well-suited for a numerical implementation (e.g. Arrow-Hurwicz algorithm), because

a ratio of expectations is not an expectation!

An Useful Lemma

Using compact notation, the previous problem writes:

$$\min_{\mathbf{u}} \frac{J(\mathbf{u})}{\Theta(\mathbf{u})} \quad \text{s.t.} \quad \Theta(\mathbf{u}) \geq p \,, \tag{9}$$

in which J and Θ assume positive values.

1. If $\mathbf{u}^\sharp$ is a solution of (9) and if $\Theta(\mathbf{u}^\sharp) = p$, then $\mathbf{u}^\sharp$ is also a solution of

$$\min_{\mathbf{u}} J(\mathbf{u}) \quad \text{s.t.} \quad \Theta(\mathbf{u}) \geq p \,. \tag{10}$$

2. Conversely, if $\mathbf{u}^\sharp$ is a solution of (10), and if an optimal Kuhn-Tucker multiplier $\beta^\sharp$ satisfies the condition

$$\beta^\sharp \geq \frac{J(\mathbf{u}^\sharp)}{\Theta(\mathbf{u}^\sharp)} \,,$$

then $\mathbf{u}^\sharp$ is also a solution of (9).

Back to a Cost in Expectation

Using the previous lemma, we aim at solving a problem in which the cost and the constraint functions correspond to expectations:

$$\begin{aligned} &\min_{u(\cdot)} \min_{\mathbf{V}(\cdot)} \mathbb{E}\Big(K\big(x^{\xi}(t_{\mathrm{f}})\big) \times \mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big) \\ \text{s.t.} \quad &\mathbb{E}\Big(\mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big) \geq p\,, \end{aligned}$$

or equivalently (Interchange Theorem [R&W, 1998]):

$$\begin{aligned} &\min_{u(\cdot)} \mathbb{E}\Big(\min_{v^{\xi}(\cdot)} K\big(x^{\xi}(t_{\mathrm{f}})\big) \times \mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big) \\ \text{s.t.} \quad &\mathbb{E}\Big(\mathbb{I}\big(\big\|C\big(x^{\xi}(t_{\mathrm{f}})\big)\big\|\big)\Big) \geq p\,. \end{aligned}$$

Lagrangian Formulation

$$\min_{u(\cdot)} \mathbb{E}\Big(\min_{v^\xi(\cdot)} K\big(x^\xi(t_\mathrm{f})\big) \times \mathbb{I}\big(\big\| C\big(x^\xi(t_\mathrm{f})\big)\big\|\big)\Big)$$

$$\text{s.t.} \quad p - \mathbb{E}\Big(\mathbb{I}\big(\big\| C\big(x^\xi(t_\mathrm{f})\big)\big\|\big)\Big) \leq 0 \qquad \leftrightsquigarrow \qquad \mu$$

Assume there exists a saddle point for the associated Lagrangian. In order to solve

$$\max_{\mu \geq 0} \; \min_{u(\cdot)} \left\{ \mu\, p + \mathbb{E}\Big(\underbrace{\min_{v^\xi(\cdot)} \big(K\big(x^\xi(t_\mathrm{f})\big) - \mu\big) \times \mathbb{I}\big(\big\| C\big(x^\xi(t_\mathrm{f})\big)\big\|\big)}_{W(u,\mu,\xi)} \Big) \right\} .$$

that is,

$$\max_{\mu \geq 0} \; \min_{u(\cdot)} \Big\{ \mu\, p \;+\; \mathbb{E}\big(W(u,\mu,\xi)\big) \Big\} ,$$

we use an adapted Arrow-Hurwicz algorithm ([Culioli, 1994]).

Algorithm Overview

Arrow-Hurwicz algorithm

At iteration k,

1. draw a failure $\xi^k = (t_{\mathrm{p}}^{\xi^k}, t_{\mathrm{d}}^{\xi^k})$ according to its probability law,
2. compute the gradient of W w.r.t. u and update $u(\cdot)$:
$$u^{k+1} = \Pi_{\mathcal{B}}\Big(u^k - \varepsilon^k \nabla_u W(u^k, \mu^k, \xi^k)\Big) ,$$
3. compute the gradient of W w.r.t. μ and update μ:
$$\mu^{k+1} = \max\Big(0, \mu^k + \rho^k\big(p + \nabla_\mu W(u^{k+1}, \mu^k, \xi^k)\big)\Big) .$$

First Difficulty: $\mathbb{I}$ Is Not a Smooth Function

At every iteration k, we must evaluate function W as well as its derivatives w.r.t. $u(.)$ and μ. But W is not differentiable!

$$\mathbb{I}(y) = \begin{cases} 1 & \text{if } y = 0, \\ 0 & \text{otherwise,} \end{cases} \quad \rightsquigarrow \quad \mathbb{I}_r(y) = \begin{cases} \left(1 - \frac{y^2}{r^2}\right)^2 & \text{if } y \in [-r, r], \\ 0 & \text{otherwise.} \end{cases}$$

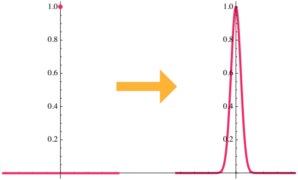

There are specific rules to drive r to 0 as the iteration number k goes to infinity in order to obtain the best asymptotic Mean Quadratic Error of the gradient estimates ([Andrieu et al., 2007]).

Second Difficulty: Solving the Inner Problem

The approximated closed-loop problem to solve at each iteration is:

$$W_{r^k}(u^k, \xi^k, \mu^k) = \min_{v^\xi(\cdot)} \Big\{ \big(K\big(x^\xi(t_\mathrm{f})\big) - \mu^k\big) \times \mathbb{I}_{r^k}\big(\big\|C\big(x^\xi(t_\mathrm{f})\big)\big\|\big) \Big\} .$$

In this setting, we have to check if the target is reached up to r^k. Different cases have to be considered:

1. the target can be reached accurately,
2. the target can be reached up to r^k only,
3. the target cannot be reached up to r^k.

Note that if reaching the target is possible but too expensive (that is, if $K\big(x^\xi(t_\mathrm{f})\big) \geq \mu^k$), the best thing to do is to stop the engine!

In practice, the solution of the approximated problem is derived from the resolution of two standard optimal control problems...

Parameters Tuning

Gradient step length:

$$\varepsilon^k = \frac{a}{b+k} \qquad , \qquad \rho^k = \frac{c}{d+k} ,$$

$\rightsquigarrow$ usual for a stochastic gradient algorithm.

Smoothing parameter:

$$r^k = \frac{\alpha}{\beta + k^{\frac{1}{3}}} ,$$

$\rightsquigarrow$ MQE reduced by a factor 1000 in about 100.000 iterations.

Example: Interplanetary Mission

- $t_{\mathrm{i}} = 0.70$ and $t_{\mathrm{f}} = 8.70$ (normalized units),
- $\mathbf{t}_{\mathrm{p}}$: exponential distribution: $\mathbb{P}\big(\mathbf{t}_{\mathrm{p}} \geq t_{\mathrm{f}}\big) \approx 0.58 = \pi_{\mathrm{f}}$,
- $\mathbf{t}_{\mathrm{d}}$: exponential distribution: $\mathbb{P}\big(0.035 \leq \mathbf{t}_{\mathrm{d}} \leq 0.125\big) \approx 0.80$.

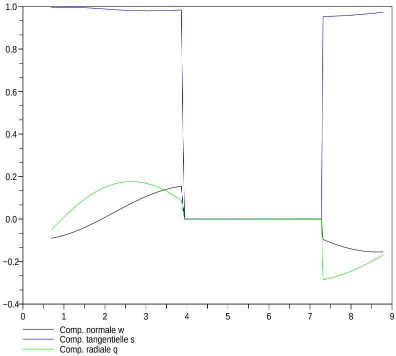

The deterministic optimal control has a "bang–off–bang" shape.

Along the optimal trajectory, the probability to recover a failure is: $p^{\mathrm{det}} \approx 0.94$.

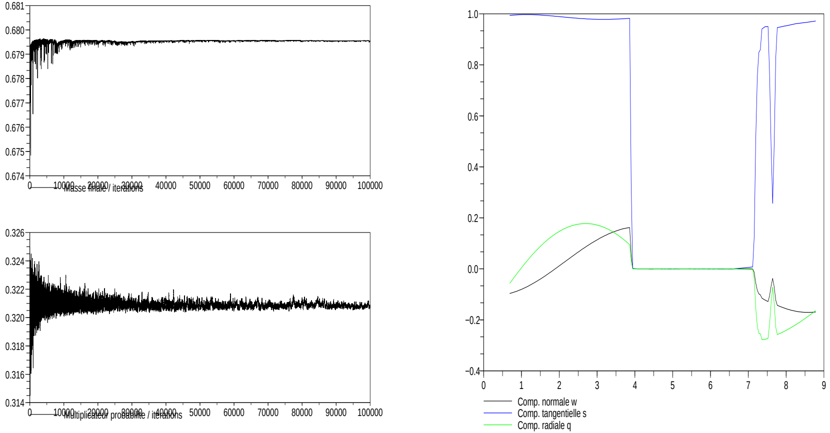

Figure: Probability level $p = 0.750$

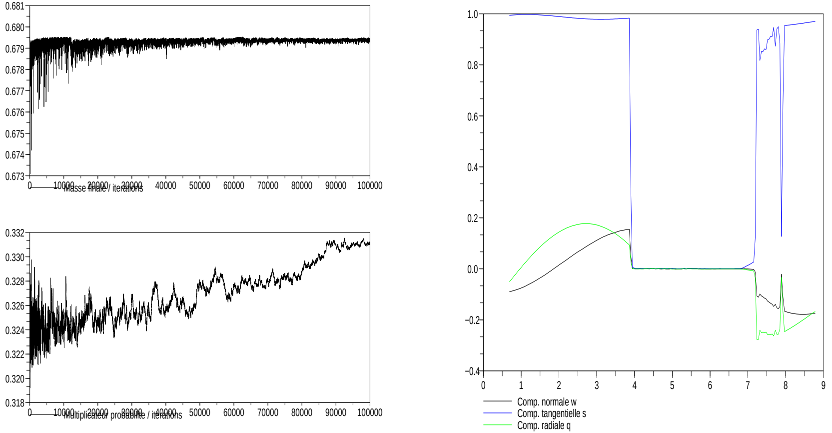

Figure: Probability level $p = 0.960$

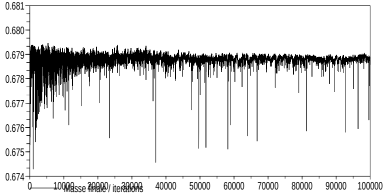

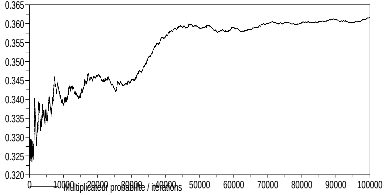

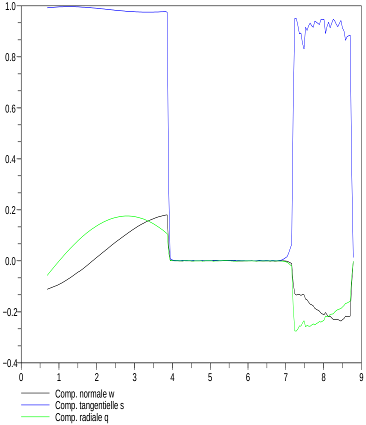

Figure: Probability level $p = 0.990$

The Price of Safety. . .

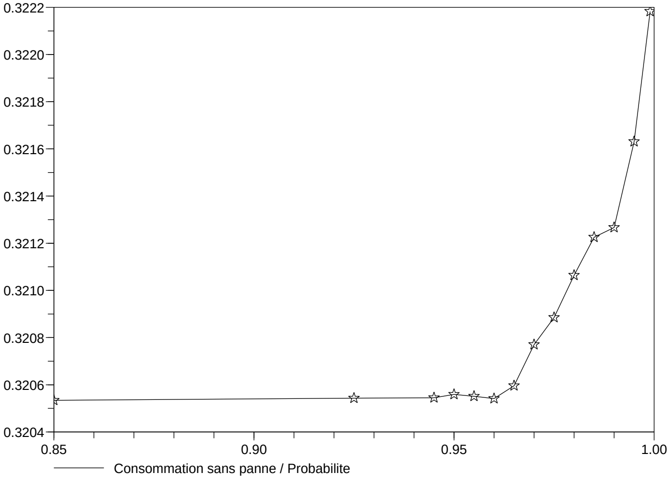

H. Robbins and S. Monro.
A stochastic approximation method.
Annals of Mathematical Statistics, 22:400–407, 1951.

B. T. Polyak and A. B. Juditsky.
Acceleration of stochastic approximation by averaging.
SIAM Journal on Control and Optimization, 30(4):838–855, 1992.

M. Duflo.
Algorithmes stochastiques.
Springer Verlag, 1996.

H. J. Kushner and D. S. Clark.
Stochastic Approximation Methods for Constrained and Unconstrained Systems.
Springer Verlag, 1978.

J.-C. Culioli and G. Cohen.
Decomposition-coordination algorithms in stochastic optimization.
SIAM Journal on Control and Optimization, 28(6):1372–1403, 1990.

J.-C. Culioli and G. Cohen.
Optimisation stochastique sous contraintes en espérance.
CRAS, t.320, Série I, pp. 75-758, 1995.

Laetitia Andrieu, Guy Cohen and Felisa J. Vázquez-Abad.
Gradient-based simulation optimization under probability constraints.
European Journal of Operational Research, 212, 345-351, 2011.

R. T. Rockafellar and R. J.-B. Wets.
Variational Analysis, volume 317 of *A series of Comprehensive Studies in Mathematics*.
Springer Verlag, 1998.