A Stochastic Min-plus Algorithm for Deterministic Optimal Control

Marianne Akian(1), Jean-Philippe Chancelier(2) and Benoît Tran(2)

(1) École Polytechnique, CMAP and INRIA

(2) École Nationale des Ponts et Chaussées, CERMICS

In short

In her thesis, Zheng Qu [1] exploited the min-plus linearity of the Bellman operators associated with a continuous time linear-quadratic optimal control problem, which involves both continuous and discrete controls, in order to build a stochastic algorithm for such a deterministic framework.
Here, we describe an adaptation of her work in discrete time and show that this adaptation converges almost surely to the value functions after a finite number of steps.

I - Introduction

We will denote by $\mathbb{X} := \mathbb{R}^n$ the space of states, by $\mathbb{U} := \mathbb{R}^m$ the space of continuous controls and let $\mathbb{V}$ be a finite set of discrete controls. We want to solve the following deterministic multistage optimal control problem involving mixed continuous and discrete controls

$$\begin{aligned} &\min_{\substack{x \in \mathbb{X}^T \\ (u,v) \in (\mathbb{U}\times\mathbb{V})^{T-1}}} \sum_{t=0}^{T-1} c_t^{v_t}(x_t, u_t) + \psi(x_T) \\ &s.t. \begin{cases} x_{t+1} = f_t^{v_t}(x_t, u_t) \\ x_0 \in \mathbb{X} \text{ is given,} \end{cases} \end{aligned} \tag{1}$$

II - Assumptions

For every time $t \in [\![0, T-1]\!]$ and for every discrete control $v \in \mathbb{V}$:

- The dynamic $f_t^v : \mathbb{X} \times \mathbb{U} \to \mathbb{X}$ is linear.
- The cost function $c_t^v : \mathbb{X} \times \mathbb{U} \to \mathbb{R} \cup \{+\infty\}$ is (jointly) convex.
- The final cost function $\psi : \mathbb{X} \to \mathbb{R} \cup \{+\infty\}$ is convex.

III - Dynamic Programming

Following the Dynamic Programming principle, the solution of Problem (1) is given by the solution of the *Bellman Equation* :

$$\begin{cases} V_T := \psi \\ V_t := \inf\limits_{(u,v)\in\mathbb{U}\times\mathbb{V}} c_t^v(\cdot, u) + V_{t+1}\left(f_t^v(\cdot, u)\right). \end{cases}$$

We now introduce the *Bellman Operator* $\mathcal{B}_t$, for every functional $\phi : \mathbb{X} \to \mathbb{R}$ we define :

$$\mathcal{B}_t(\phi) := \inf_{(v,u)\in\mathbb{V}\times\mathbb{U}} c_t^v(\cdot, u) + \phi\left(f_t^v(\cdot, u)\right).$$

Thus we can rewrite the dynamic programming principle as :

$$\begin{cases} V_T = \psi \\ V_t : x \in \mathbb{X} \mapsto \mathcal{B}_t(V_{t+1})(x). \end{cases}$$

IV - Bellman Operators

For a general switched deterministic optimal control problem, given a fixed discrete control $v \in \mathbb{V}$ the Bellman Operators

$$\mathcal{B}_t^v(\phi) := \inf_{u\in\mathbb{U}} c_t^v(\cdot, u) + \phi\left(f_t^v(\cdot, u)\right).$$

are *min-plus linear* and preserve the set of the convex functions.

- If ϕ_1 and ϕ_2 are two functionals over $\mathbb{X}$, then we have :
$$\inf_{\phi\in\{\phi_1,\phi_2\}} \mathcal{B}_t^v(\phi) = \mathcal{B}_t^v\left(\inf_{\phi\in\{\phi_1,\phi_2\}} \phi\right).$$
- If λ is a constant function and ϕ is a functional on $\mathbb{X}$, then we have :
$$\mathcal{B}_t^v(\lambda + \phi) = \lambda + \mathcal{B}_t^v(\phi).$$
- If $\phi : \mathbb{X} \to \mathbb{R}$ is a convex function then $\mathcal{B}_t^v(\phi)$ is also a convex function.

Important result : (sur)approximations by convex functions

For every time $t \in [\![0, T]\!]$, there exists a finite non-empty set Q_t of convex functions such that $V_t = \inf_{q\in Q_t} q$.

V - Main observation

We have that

$$\mathrm{Card}(Q_t) = \mathrm{Card}(Q_{t+1}) \cdot \mathrm{Card}(\mathbb{V}).$$

Thus the cardinality of Q_t grows exponentially as t decreases to 0.
Hence, naive computations of the value function lead to untractable algorithms.
We describe a tractable algorithm which approximates the value function V_t by a function V_t^k which will be defined as a minimum over a finite subset Q_t^k of Q_t :

$$V_t^k := \inf_{q\in Q_t^k} q \geq \inf_{q\in Q_t} q =: V_t.$$

VI - Stochastic min-plus algorithm

We set for every time $t \in [\![0, T]\!]$ and for every $k \geq 0$, $V_t^k := \inf_{q\in Q_t^k} q$ where $Q_t^k \subset Q_t$. For every $t \in [\![0, T-1]\!]$ we fix a compact K_t.

- For $k = 0$, define for every $t \in [\![0, T]\!]$, $Q_t^0 := \emptyset$.
- For every $k \geq 1$, **input** : Q_t^k.
- Draw x_t^{k+1} from a probability law μ_t whose support is K_t.
- If $\mathcal{B}_t\left(V_{t+1}^{k+1}\right)\left(x_t^{k+1}\right) < V_t^k\left(x_t^{k+1}\right)$ then, for some
$$q \in \operatorname*{arg\,min}_{q'\in Q_{t+1}^{k+1}} \mathcal{B}_t(q')\left(x_t^{k+1}\right),$$
output : $Q_t^k \cup \{q\}$. Otherwise, **output** : Q_t^k.

VII - Computational remark

This algorithm needs to compute $\mathcal{B}_t(\phi)$ where ϕ is convex. It works at least when ϕ is convex and quadratic by solving an algebraic Riccati Equation. Moreover the image $\mathcal{B}_t^v(\phi)$ is also convex and quadratic.

Convergence result

For every $t \in [\![0, T]\!]$ the sequence of surapproximations $\left(V_t^k\right)_{k\in\mathbb{N}}$ is equal to V_t on K_t for k large enough. That is, almost surely there exists a rank k^* such that for every $k \geq k^*$ and $x \in K_t$ we have

$$V_t(x) = V_t^k(x).$$

VIII - Future work

- Extend the current work when the supports change with k.
- Extend to the stochastic case.
- Numerical experiments.

References

[1] Zheng Qu.
Nonlinear Perron-Frobenius Theory and Max-plus Numerical Methods for Hamilton-Jacobi Equations.
PhD thesis, Ecole Polytechnique X, October 2013.

École des Ponts ParisTech

CMAP