poema

Polynomial Optimization, Efficiency through Moments and Algebra

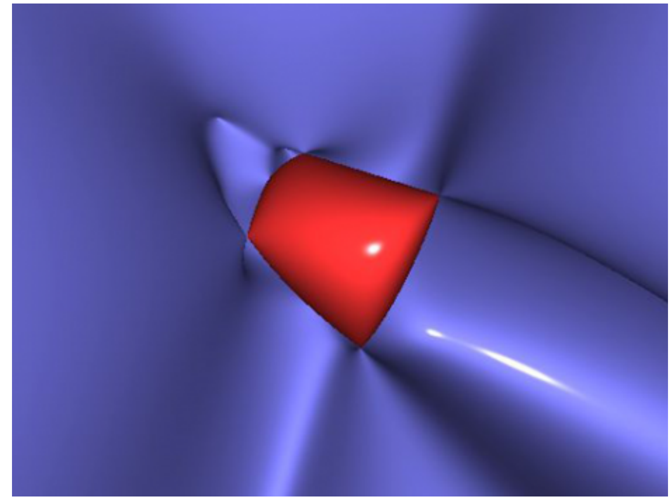

Online learning weeks of the POEMA network

Didier Henrion

Polynomial optimal control

10 June 2020

Recall the main steps of the moment-SOS aka Lasserre hierarchy.

Given a **nonlinear nonconvex** problem:

1. Reformulate it as a **linear** problem (at the price of enlarging or changing the space of solutions);

2. Solve approximately the linear problem with a hierarchy of tractable **convex** relaxations (of increasing size);

3. Ensure **convergence**: either the original problem is solved at a finite relaxation size, or its solution is approximated with increasing quality.

At each step, **conic duality** is an essential ingredient.

POC

A **polynomial optimal control** (POC) problem is a time-varying extension of a polynomial optimization problem

$$\begin{array}{rcll} v^*(t_0, x_0) & := & \inf_u & \int_{t_0}^{T} l(x_t, u_t)dt + l_T(x_T) \\ & & \text{s.t.} & \dot{x}_t = f(x_t, u_t),\ x_{t_0} = x_0 \\ & & & x_t \in X,\ u_t \in U,\ \forall t \in [t_0, T] \\ & & & x_T \in X_T \end{array}$$

All the given data f, l, l_T are polynomial
and the given sets X, X_T, U are semi-algebraic

Terminal time T can be either given or free

The function v^* of the initial data t_0, x_0 is the **value function**

Why is the value function important ?

From value function to optimal control

From the value function v^* we can derive an optimal control

$$u_t^* \in \arg\min_u \{l(x_t, u) + \text{grad}\, v^*(t, x_t) \cdot f(x_t, u)\}$$

by solving an **optimization** problem

Then we can **verify** optimality

$$l(x_t, u_t^*) + \frac{\partial v^*(t, x_t)}{\partial t} + \text{grad}\, v^*(t, x_t) \cdot f(x_t, u_t^*) = 0$$

HJB PDE

The value function solves the Hamilton-Jacobi-Bellman (HJB) equation, a nonlinear first-order partial differential equation (PDE)

$$\frac{\partial v(t,x)}{\partial t} + h(t, \operatorname{grad} v(t,x)) = 0$$
$$v(T,.) = l_T$$

with Hamiltonian conjugate to the Lagrangian

$$h(t,p) := \inf_u \{l(x,u) + p \cdot f(x,u)\}$$

If $u \mapsto f(x,u)$ is affine and $u \mapsto l(x,u)$ is quadratic convex the HJB PDE is a classical Riccati equation

In general this PDE does not have a regular solution, and a notion of weak solution (viscosity solution) must be defined

What is the geometry of the value function ?

Eikonal equation of geometric optics

Minimum time to reach the boundary of a given set $X \subset \mathbb{R}^n$ with velocity bounded in $U := \{u \in \mathbb{R}^m : \|u\|_2 \leq 1\}$

$$\begin{array}{rcll} v^*(x_0) & := & \inf_u & \int_0^T dt \\ & & \text{s.t.} & \dot{x}_t = u_t, \ x_0 \text{ given} \\ & & & x_t \in X, \ u_t \in U, \ \forall t \in [0,T] \\ & & & x_T \in X_T := \partial X \end{array}$$

Here, the value function does not depend on time

Let $n = m = 1$ and $X = U = [-1, 1]$, what is $x \mapsto v^*(x)$?

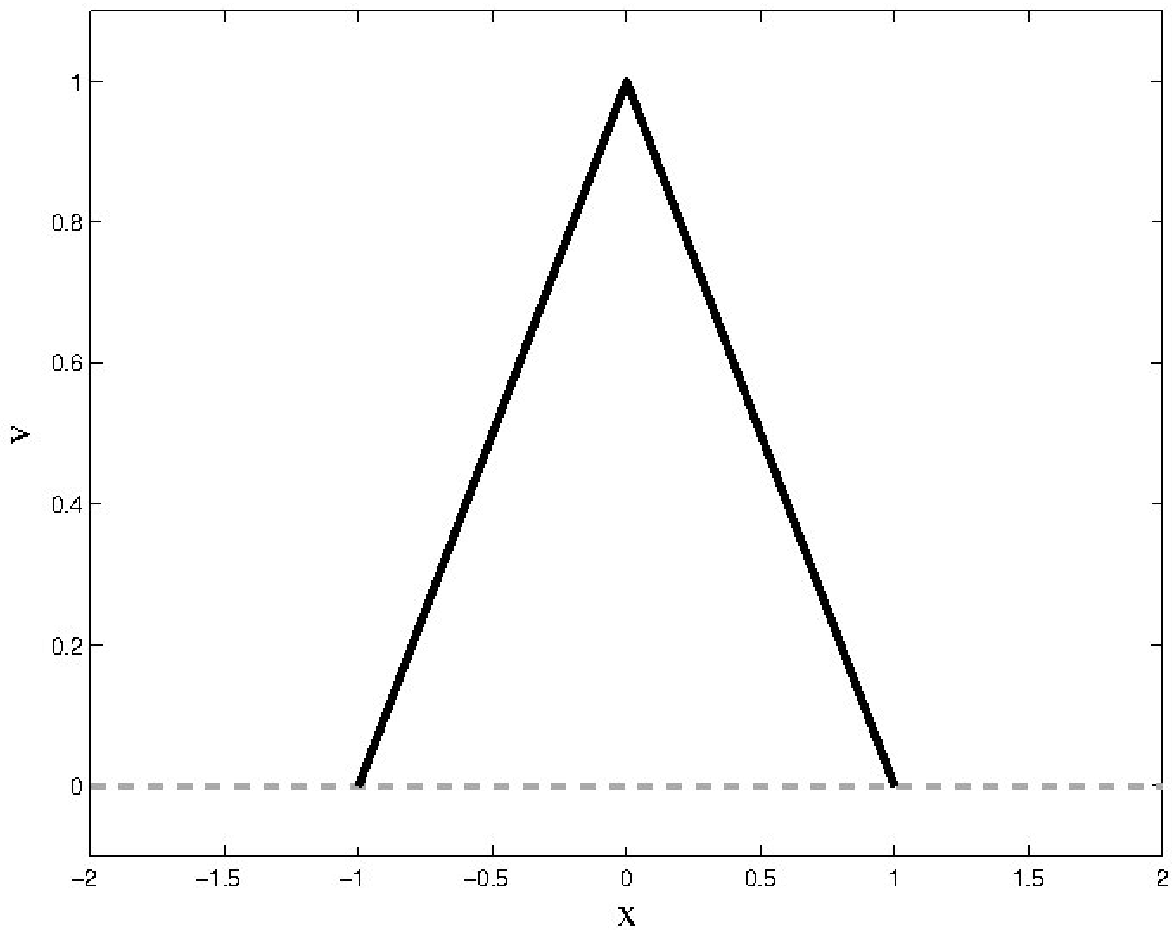

v
1
0.8
0.6
0.4
0.2
0
−2
−1.5
−1
−0.5
0
0.5
1
1.5
2
x

Eikonal HJB

Lagrangian $l(x,u) = 1$, vector field $f(x,u) = u$
Hamiltonian $h(x,p) := \inf_{\|u\|_2 \le 1}\{l(x,u) + p \cdot f(x,u)\} = 1 - \|p\|_2^2$

$$\begin{aligned} 1 - \|\text{grad}\, v\|_2^2 = 0 \quad & \text{on } X \\ v = 0 \quad & \text{on } \partial X \end{aligned}$$

Amongst the many functions vanishing on ∂X and with unit gradient on X there is a unique solution in the viscosity sense

It is nondifferentiable at the origin

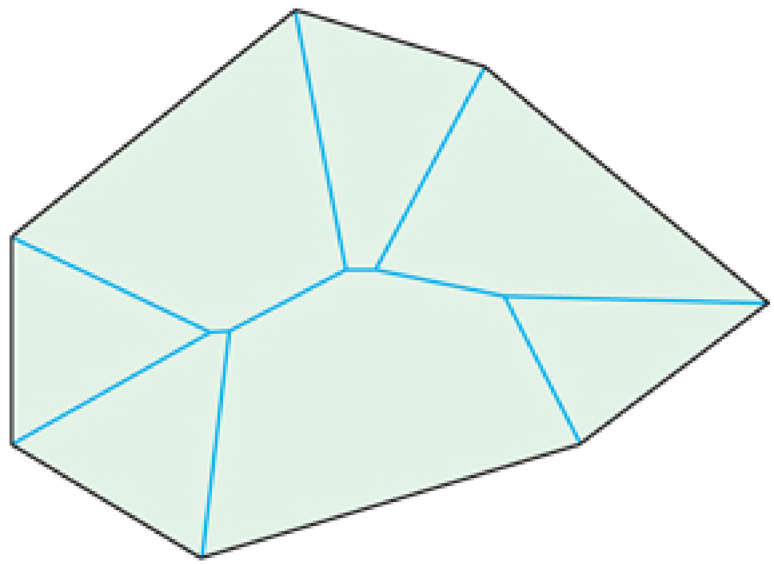

Cut locus = set of nondifferentiability points
of eikonal function of a polytope X

The value function can be complicated

Wheeled robot value function

Minimum time problem for Dubins or Reeds-Shepp car of nonholonomic robotics

$$
\begin{aligned}
\dot{x}_t &= u_t \cos\theta_t \\
\dot{y}_t &= u_t \sin\theta_t \\
\dot{\theta}_t &= v_t
\end{aligned}
$$

equivalent to Brockett's integrator

$$
\begin{aligned}
\dot{x}_t &= u_t \\
\dot{y}_t &= v_t \\
\dot{z}_t &= y_t u_t - x_t v_t
\end{aligned}
$$

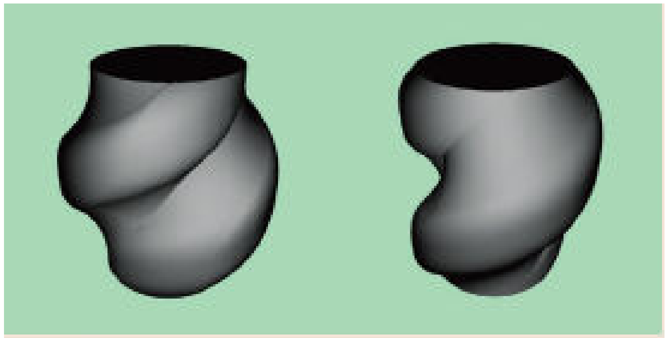

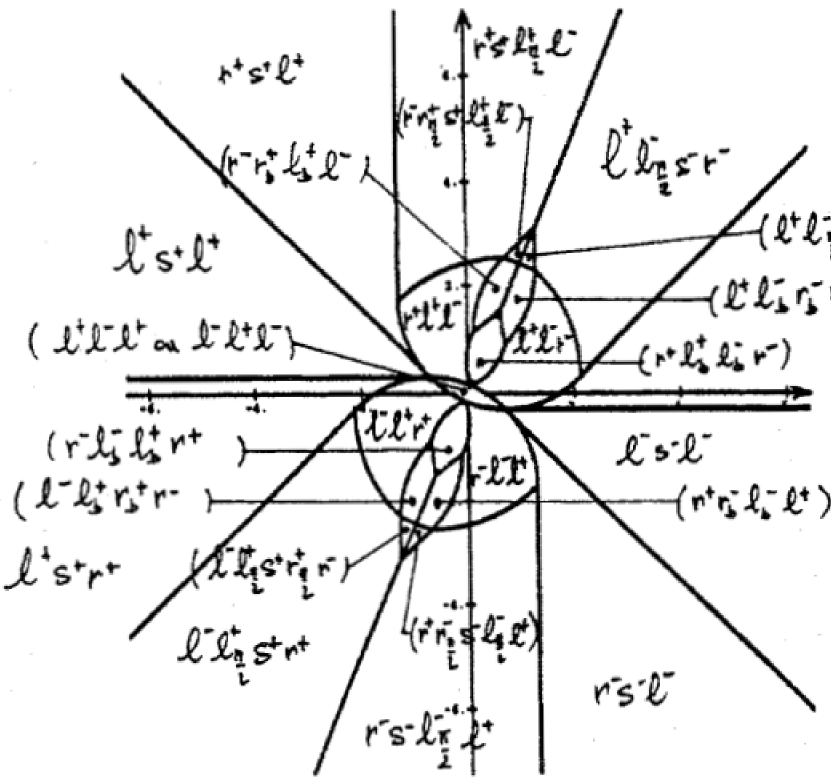

$\theta = -\frac{\pi}{4}$

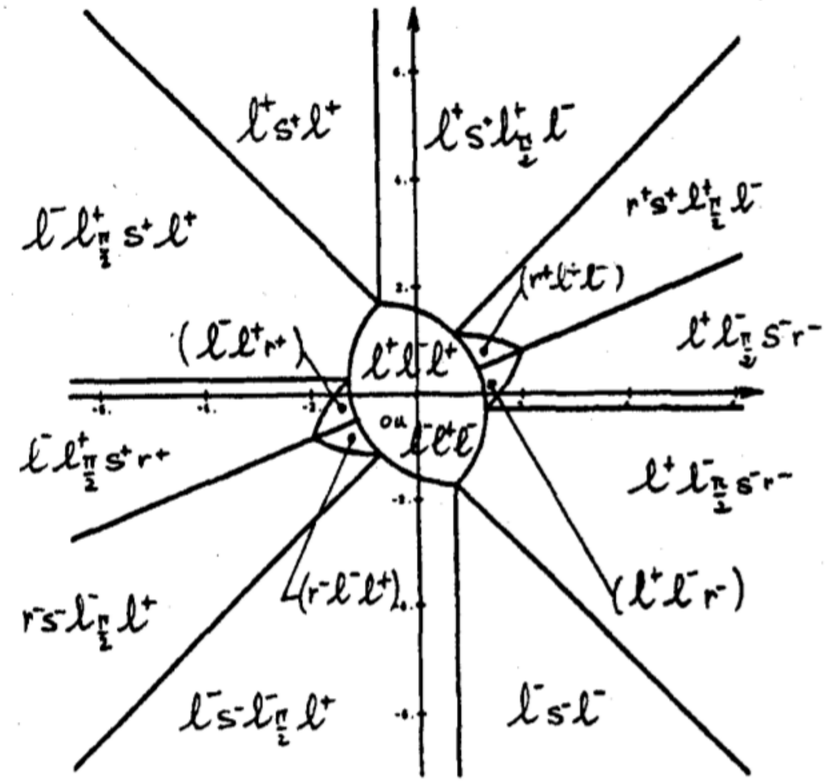

$\theta = -\frac{3\pi}{4}$

The value function is not smooth

.. just lower semicontinuous

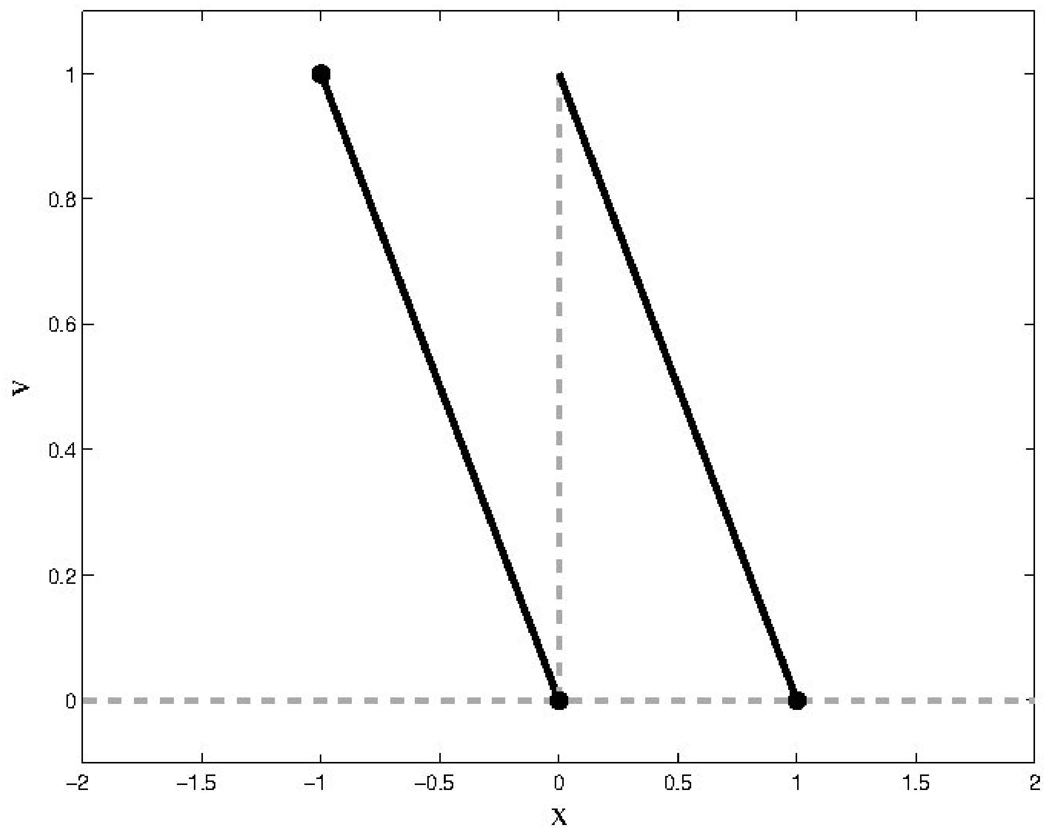

Minimum time for $X = [-1, 1]$, $U = [0, 1]$, $X_T = \{0, 1\}$

In general there is no optimal control !

No optimal control

Bolza problem

$$\begin{aligned} v^* &= \inf_u \int_0^1 (x_t^2 + (u_t^2 - 1)^2) dt \\ & \text{s.t.} \quad \dot{x}_t = u_t,\ x_0 = 0 \\ & \quad\quad x_t \in X := [-1,1],\ u_t \in U := [-1,1]\ \forall t \in [0,1] \end{aligned}$$

The cost is **nonconvex** in the control

Let us construct a minimizing sequence...

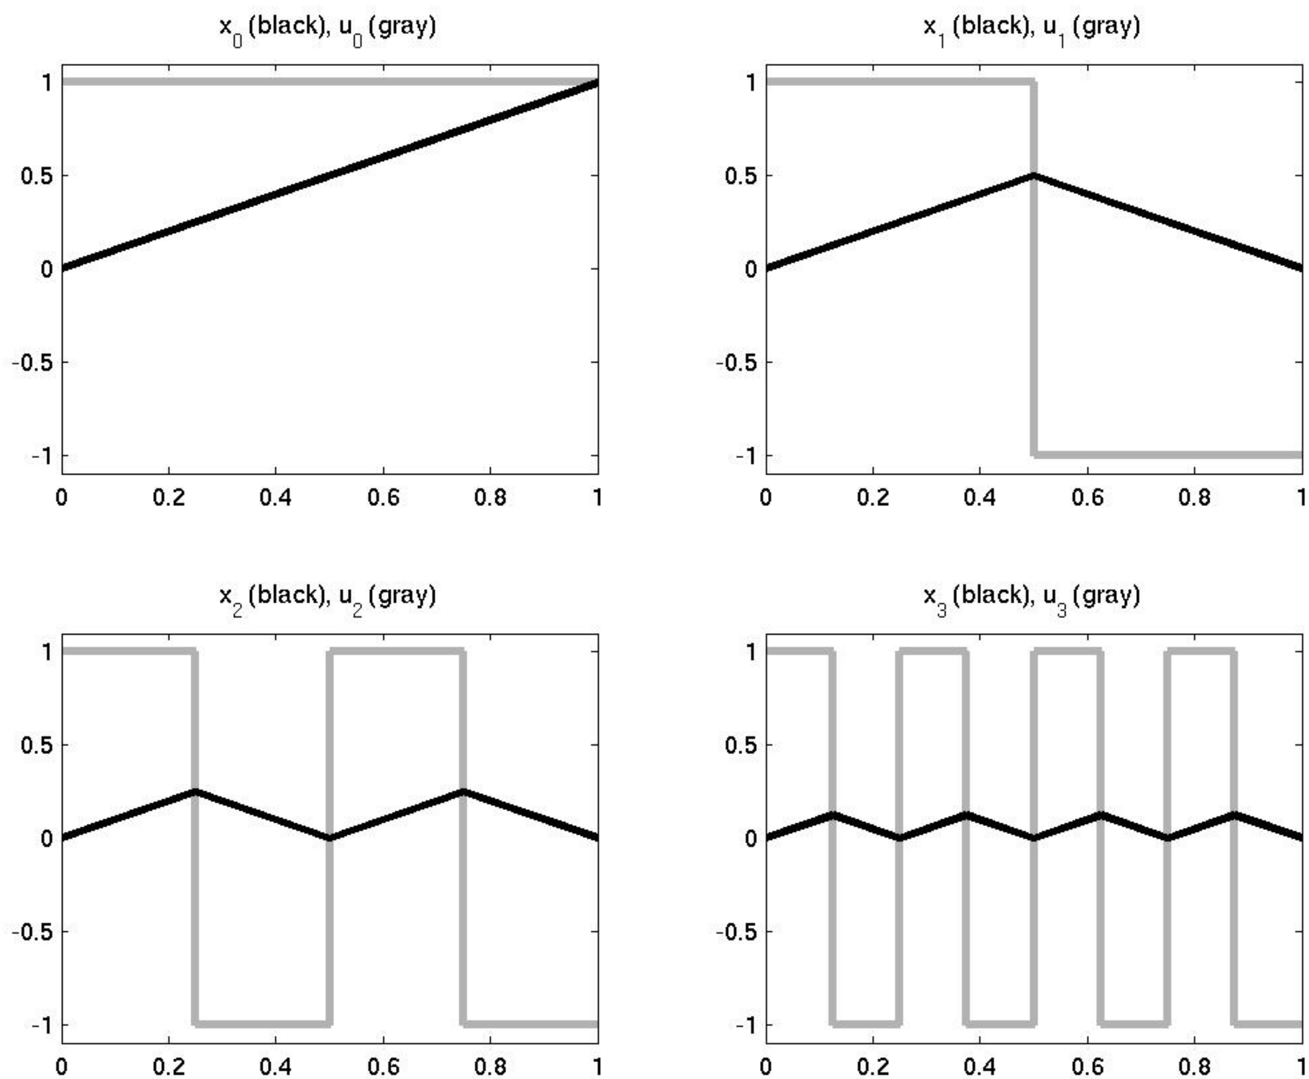

x_0 (black), u_0 (gray)
x_1 (black), u_1 (gray)
x_2 (black), u_2 (gray)
x_3 (black), u_3 (gray)
1
0.5
0
-0.5
-1
0
0.2
0.4
0.6
0.8
1

The infimum $v^* = 0$ is **not** attained in the space of measurable functions of time

$$t \mapsto u_t \in U$$

so let us enlarge the space of allowable controls

We proceed as previously for POP and MPI set approximation

Instead of classical controls let us consider **relaxed controls**

$$t \mapsto \omega_t(du) = \omega(du|t) \in \mathscr{P}(U)$$

as time-dependent probability measures on U

The controlled ordinary differential equation (ODE)

$$\dot{x}_t = f(x_t, u_t), \quad u_t \in U$$

becomes a relaxed controlled ODE

$$\dot{x}_t = \int_U f(x_t, u)\omega_t(du)$$

Classical controls correspond to $\omega_t(du) = \delta_{u_t}(du)$

Relaxed controls capture limit behavior such as e.g. oscillations

$$\lim_{r\to\infty} \int_{t_0}^{T} v(u_{rt})dt = \int_{t_0}^{T} \int_U v(u)\omega_t(du)dt, \quad \forall v \in C(U)$$

Exercise 3.1: What is the limit $\omega_t(du)$ if $u_r(t) = \sin(2\pi rt)$?

The classical Bolza problem

$$
\begin{array}{rcl}
v^* & = & \inf_u \int_0^1 (x_t^2 + (u_t^2 - 1)^2) dt \\
 & & \text{s.t.} \quad \dot{x}_t = u_t, \; x_0 = 0 \\
 & & \qquad x_t \in [-1, 1], \; u_t \in [-1, 1] \;\; \forall t \in [0, 1]
\end{array}
$$

becomes the relaxed Bolza problem

$$
\begin{array}{rcl}
v_R^* & = & \inf_{\omega_t} \int_0^1 \int_U (x_t^2 + (u^2 - 1)^2) \, \omega_t(du) dt \\
 & & \text{s.t.} \quad \dot{x}_t = \int_U u \, \omega_t(du), \; x_0 = 0 \\
 & & \qquad x_t \in [-1, 1], \; \omega_t \in \mathscr{P}([-1, 1]) \;\; \forall t \in [0, 1]
\end{array}
$$

Exercise 3.2: Prove that there is no relaxation gap: $v^* = v_R^*$ and that the relaxed infimum is attained at $\omega_t^* = \frac{1}{2}(\delta_{-1} + \delta_{+1})$.

Let's relax

The classical POC problem

$$\begin{array}{rcll} v^*(t_0,x_0) & := & \inf_u & \int_{t_0}^{T} l(x_t,u_t)dt + l_T(x_T) \\ & & \text{s.t.} & \dot{x}_t = f(x_t,u_t),\ x_{t_0} = x_0 \\ & & & x_t \in X,\ u_t \in U,\ \forall t \in [t_0,T] \\ & & & x_T \in X_T \end{array}$$

becomes a relaxed POC problem

$$\begin{array}{rcll} v_R^*(t_0,x_0) & := & \min_{\omega_t} & \int_{t_0}^{T} \int_U l(x_t,u)\omega_t(du)dt + l_T(x_T) \\ & & \text{s.t.} & \dot{x}_t = \int_U f(x_t,u)\omega_t(du),\ x_{t_0} = x_0 \\ & & & x_t \in X,\ \omega_t \in \mathscr{P}(U),\ \forall t \in [t_0,T] \\ & & & x_T \in X_T \end{array}$$

and under reasonable convexity assumptions, it can be shown that there is **no relaxation gap**: $v_R^* = v^*$

Not relaxed enough

The POC problem is still nonlinear in the state trajectory

For a given initial state x_0 and a given relaxed control ω_t, let us introduce the **occupation measure**

$$d\mu(t,x,u|x_0) := dt\,\omega_t(du)\delta_{x_t}(dx|x_0,u)$$

corresponding to the trajectory x_t

Integration of an observable along time trajectory becomes

$$\int_{t_0}^{T} v(x_t)dt = \int_{t_0}^{T}\int_X\int_U v(x)d\mu(t,x,u|x_0) = \langle v,\mu\rangle, \quad \forall v \in C(X)$$

Given $v \in C^1([t_0, T] \times X)$, notice that

$$\begin{aligned} v(T, x_T) - v(t_0, x_0) &= \int_{t_0}^{T} dv(t, x_t) \\ &= \int_{t_0}^{T} \dot{v}(t, x_t) dt \\ &= \int_{t_0}^{T} \mathscr{L} v(t, x_t) dt \end{aligned}$$

where $\mathscr{L}$ is an operator from $C^1([t_0, T] \times X)$ to $C([t_0, T] \times X \times U)$:

$$v \mapsto \mathscr{L} v := \frac{\partial v}{\partial t} + \text{grad}\, v \cdot f$$

The relation

$$v(T, x_T) - v(t_0, x_0) = \int_{t_0}^{T} \mathscr{L} v(t, x_t) dt$$

can be written as follows

$$\int_X v\mu_T - \int_X v\mu_0 = \int_{t_0}^{T} \int_X \int_U \mathscr{L} v\mu$$

after defining $\mu_0 := \delta_{t_0}(dt)\delta_{x_0}(dx)$ and $\mu_T := \delta_T(dt)\delta_{x_T}(dx)$

This is the **Liouville equation**

$$\mathscr{L}'\mu + \mu_T = \mu_0$$

involving the adjoint operator

$$\mu \mapsto \mathscr{L}'\mu := \frac{\partial \mu}{\partial t} + \mathrm{div}(f\mu)$$

For a given initial state x_0 and a given control u_t,
we replace the **nonlinear** Cauchy problem

$$\dot{x}_t = f(x_t, u_t), \quad x_{t_0} = x_0$$

with the **linear** Liouville problem

$$\frac{\partial \mu}{\partial t} + \mathrm{div}(f\mu) + \mu_T = \delta_{t_0}\delta_{x_0}$$

Lemma: There is a unique solution to the Liouville equation, whose state conditional is concentrated on the solution to the Cauchy problem: $d\mu(t,x,u|x_0) = dt\, \delta_{u_t}(du)\delta_{x_t}(dx|x_0,u)$.

Also true when replacing the classical control u_t with a relaxed control ω_t, in which case $d\mu(t,x,u|x_0) = dt\, \omega_t(du)\delta_{x_t}(dx|x_0,u)$

Step 1 - Linear reformulation

The original POC problem

$$\begin{array}{rcll} v^*(t_0,x_0) & := & \inf_u & \int_{t_0}^T l(x_t,u_t)dt + l_T(x_T) \\ & & \text{s.t.} & \dot{x}_t = f(x_t,u_t),\ x_{t_0} = x_0 \\ & & & x_t \in X,\ u_t \in U,\ \forall t \in [t_0,T] \\ & & & x_T \in X_T \end{array}$$

becomes a **linear** problem (LP)

$$\begin{array}{rcll} p^*(t_0,x_0) & := & \min_{\mu,\mu_T} & \int l\mu + \int l_T\mu_T \\ & & \text{s.t.} & \frac{\partial \mu}{\partial t} + \text{div}(f\mu) + \mu_T = \delta_{t_0}\delta_{x_0} \end{array}$$

on measures $\mu \in C([t_0,T] \times X \times U)'_+$, $\mu_T \in C(\{T\} \times X_T)'_+$

Lemma: $p^* = v^*$

The primal measure LP

$$\begin{array}{rcll} p^*(t_0,x_0) & := & \min_{\mu,\mu_T} & \int l\mu + \int l_T\mu_T \\ & & \text{s.t.} & \frac{\partial\mu}{\partial t} + \text{div}(f\mu) + \mu_T = \delta_{t_0}\delta_{x_0} \\ & & & \mu \in C([t_0,T]\times X\times U)'_+,\ \ \mu_T \in C(\{T\}\times X_T)'_+ \end{array}$$

has a dual LP

$$\begin{array}{rcll} d^*(t_0,x_0) & := & \sup_v & v(t_0,x_0) \\ & & \text{s.t.} & l + \frac{\partial v}{\partial t} + \text{grad}\, v\cdot f \in C([t_0,T]\times X\times U)_+ \\ & & & l_T - v(T,.) \in C(\{T\}\times X_T)_+ \end{array}$$

on functions $v \in C^1([t_0,T]\times X)$

Exercise 3.3: Derive the dual and prove that there is no duality gap: $p^* = d^*$.

Step 2 - Convex hierarchy

To solve the primal LP

$$
\begin{array}{rcll}
p^*(t_0,x_0) & := & \min_{\mu,\mu_T} & \int l\mu + \int l_T\mu_T \\
& & \text{s.t.} & \frac{\partial\mu}{\partial t} + \text{div}(f\mu) + \mu_T = \delta_{t_0}\delta_{x_0} \\
& & & \mu \in C([t_0,T]\times X\times U)'_+, \ \ \mu_T \in C(\{T\}\times X_T)'_+
\end{array}
$$

and dual LP

$$
\begin{array}{rcll}
d^*(t_0,x_0) & := & \sup_v & v(t_0,x_0) \\
& & \text{s.t.} & l + \frac{\partial v}{\partial t} + \text{grad}\, v\cdot f \in C([t_0,T]\times X\times U)_+ \\
& & & l_T - v(T,.) \in C(\{T\}\times X_T)_+
\end{array}
$$

with X, X_T bounded basic semialgebraic and l, l_T, f polynomial we can readily use the moment-SOS hierarchy

We replace $C(.)_+$ with $Q(.)_{r,r}$ for increasing relaxation order r and we get sequences p^*_r and d^*_r as well as pseudo-moments and polynomials v_r in $\mathbb{R}[x]_r$

Step 3 - Convergence

Dual LP

$$\begin{array}{rcll} d^*(t_0,x_0) & := & \sup_v & v(t_0,x_0) \\ & & \text{s.t.} & l + \frac{\partial v}{\partial t} + \operatorname{grad} v \cdot f \in C([t_0,T] \times X \times U)_+ \\ & & & l_T - v(T,.) \in C(\{T\} \times X_T)_+ \end{array}$$

Lemma (lower bound on the value function): it holds $v^* \geq v$ on $[t_0,T] \times X$ for every admissible v

Exercise 3.4: Prove this by combining the dual inequalities evaluated on an admissible trajectory.

Lemma (maximizing sequence): there exists an admissible sequence $(v_r)_{r\in\mathbb{N}}$ such that $\lim_{r\to\infty} v_r(t_0,x_0) = v^*(t_0,x_0)$.

This later lemma follows from strong LP duality. It means that at the limit the graph of v_r touches the graph v from below

Theorem (uniform convergence along trajectories): For any admissible $(v_r)_{r\in\mathbb{N}}$ and any optimal trajectory $(x_t)_{t\in[t_0,T]}$ it holds

$$0 \leq v^*(t, x_t) - v_r(t, x_t) \leq v^*(t_0, x_0) - v_r(t_0, x_0) \underset{r\to\infty}{\longrightarrow} 0.$$

In other words, the gap between the value function and its lower bound decreases along optimal trajectories

The result holds for any optimal trajectory and hence for all of them simultaneously

In the Liouville equation

$$\frac{\partial \mu}{\partial t} + \mathrm{div}(f\mu) + \mu_T = \delta_{t_0}\delta_{x_0}$$

instead of a Dirac right hand side we can use a general probability measure $\xi_0 \in \mathscr{P}(X)$ supported on a **set of initial conditions**

$$\frac{\partial \mu}{\partial t} + \mathrm{div}(f\mu) + \mu_T = \delta_{t_0}\xi_0 =: \mu_0$$

Equivalently, instead of using the occupation measure

$$d\mu(t,x,u|x_0) := dt\, \omega_t(du)\delta_{x_t}(dx|x_0,u)$$

we use the **averaged** occupation measure

$$d\mu(t,x,u) := \int_X d\mu(t,x,u|x_0)d\xi_0(x_0)$$

The value function also becomes averaged

$$\bar{v}^*(\mu_0) := \int_X v^*(t_0, x_0)\xi_0(x_0)$$

and matches the primal LP averaged value

$$\begin{array}{rcll} \bar{p}^*(\mu_0) & := & \min_{\mu,\mu_T} & \langle l, \mu\rangle + \langle l_T, \mu_T\rangle \\ & & \text{s.t.} & \frac{\partial \mu}{\partial t} + \mathrm{div}(f\mu) + \mu_T = \mu_0 \\ & & & \mu \in C([t_0, T] \times X \times U)'_+, \;\; \mu_T \in C(\{T\} \times X_T)'_+ \end{array}$$

and the dual LP averaged value

$$\begin{array}{rcll} \bar{d}^*(\mu_0) & := & \sup_v & \langle v, \mu_0\rangle \\ & & \text{s.t.} & l + \frac{\partial v}{\partial t} + \mathrm{grad}\, v \cdot f \in C([t_0, T] \times X \times U)_+ \\ & & & l_T - v(T, .) \in C(\{T\} \times X_T)_+ \end{array}$$

For any solution (μ, μ_T) of the primal LP, there are measures $\xi_t \in \mathscr{P}(X)$, $\omega_{t,x} \in \mathscr{P}(U)$ such that $d\mu(t,x,u) = dt\, \xi_t(dx)\omega_{t,x}(du)$, $d\mu_0(t,x) = \delta_{t_0}(dx)\xi_{t_0}(dx)$, $d\mu_T(t,x) = \delta_T(dx)\xi_T(dx)$

Let $(x_t)_{t\in[t_0,T]}$ be the trajectory starting at $x_{t_0} = x_0$

The map $F_t : x_0 \mapsto x_t$ is called the flow of the controlled ODE, and $\xi_t = F_{t\#}\xi_0$ is the image measure of the initial distribution

Theorem (uniform convergence): For any maximizing dual sequence $(v_r)_{r\in\mathbb{N}}$ and any $t \in [t_0, T]$ it holds

$$0 \leq \int_X (v^*(t,x) - v_r(t,x))\xi_t(dx) \leq \int_X (v^*(t_0,x_0) - v_r(t_0,x_0))\xi_0(dx) \underset{r\to\infty}{\longrightarrow} 0$$

First example: turnpike control

Turnpike control

$$v^*(t_0, x_0) :=$$

$$\begin{aligned} \inf_u \quad & \int_{t_0}^{2} (x_t + u_t)\,dt \\ \text{s.t.} \quad & \dot{x}_t = 1 + x_t - x_t u_t,\ x_{t_0} = x_0 \\ & x_t \in [-3, 3],\ u_t \in [0, 3] \end{aligned}$$

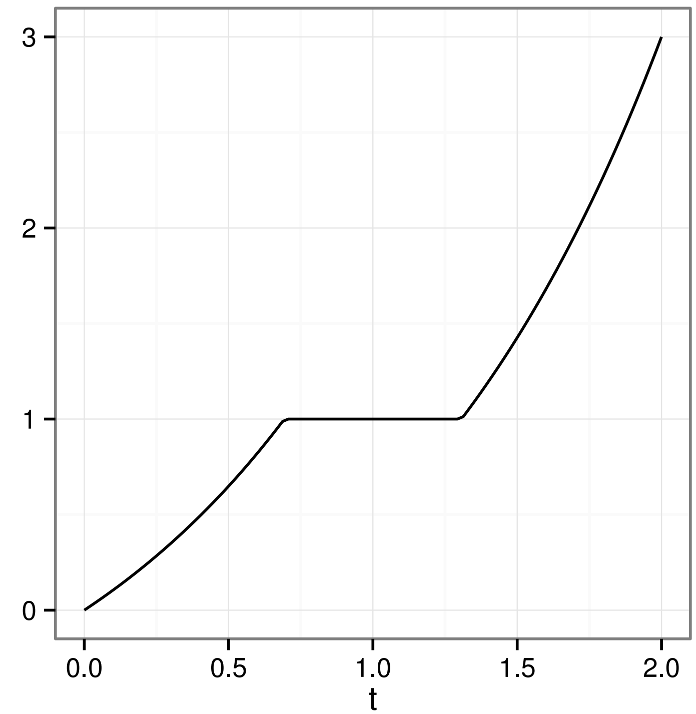

optimal trajectory x_t starting at $(t_0, x_0) = (0, 0)$

Turnpike control

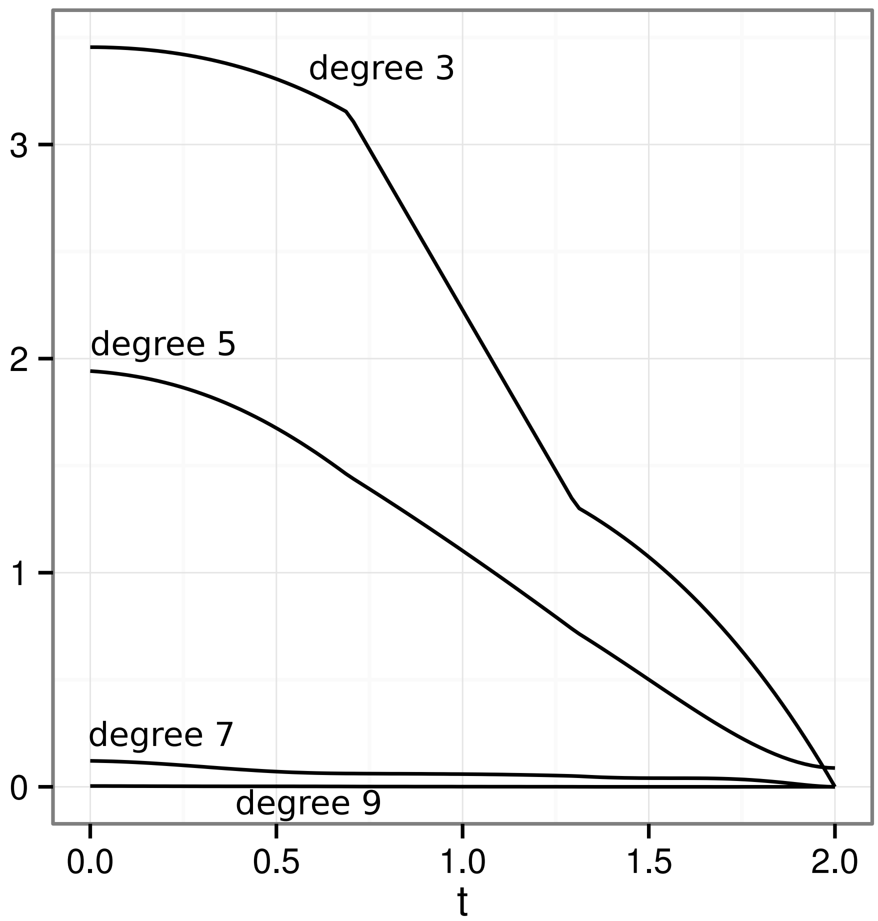

Differences $t \mapsto v^*(t, x_t) - v_r(t, x_t)$ between the actual value function and its poly. approx. of deg. $r = 3, 5, 7, 9$ along the optimal trajectory starting at $(t_0, x_0) = (0, 0)$

Observe convergence along this trajectory, as well as time decrease of the difference

Second example: LQR set control

LQR set control

$$v^*(t_0, x_0) :=$$

$$\begin{array}{ll} \inf_u & \int_{t_0}^{1} (10x_t^2 + u_t^2) dt \\ \text{s.t.} & \dot{x}_t = x_t + u_t,\ x_{t_0} = x_0 \end{array}$$

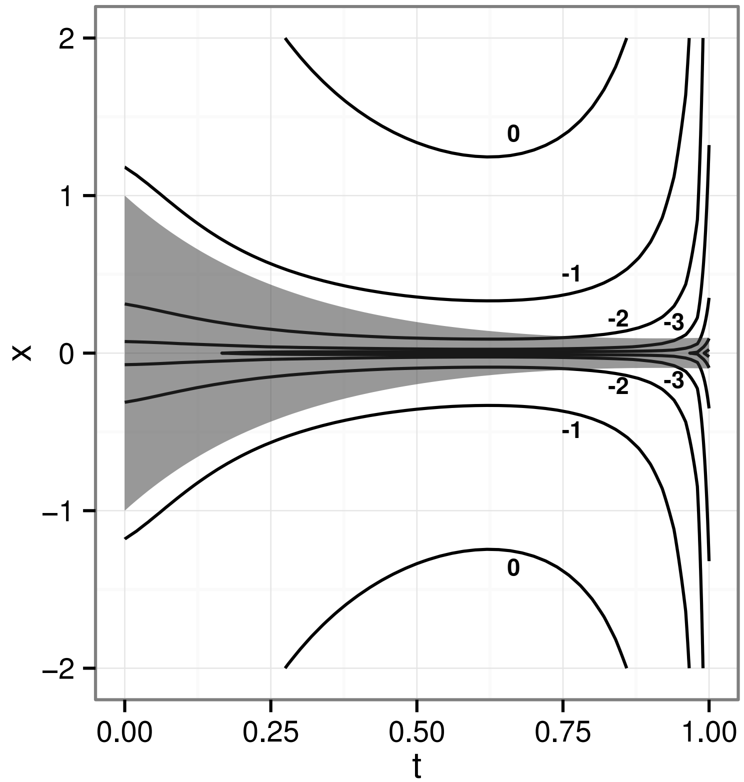

Contour lines of $(t, x) \mapsto \log(v^*(t, x) - v_6(t, x))$ with v_6 poly. approx. of deg. 6 to actual value function v^* obtained by transporting the Lebesgue measure on $[-1, 1]$

Recall the main steps of the moment-SOS aka Lasserre hierarchy.

Given a **nonlinear nonconvex** problem:

1. Reformulate it as a **linear** problem (at the price of enlarging or changing the space of solutions);

2. Solve approximately the linear problem with a hierarchy of tractable **convex** relaxations (of increasing size);

3. Ensure **convergence**: either the original problem is solved at a finite relaxation size, or its solution is approximated with increasing quality.

At each step, **conic duality** is an essential ingredient.

Thank you for your attention !