Didier Henrion

Polynomial optimization and optimal control

CAS MINES ParisTech 2021

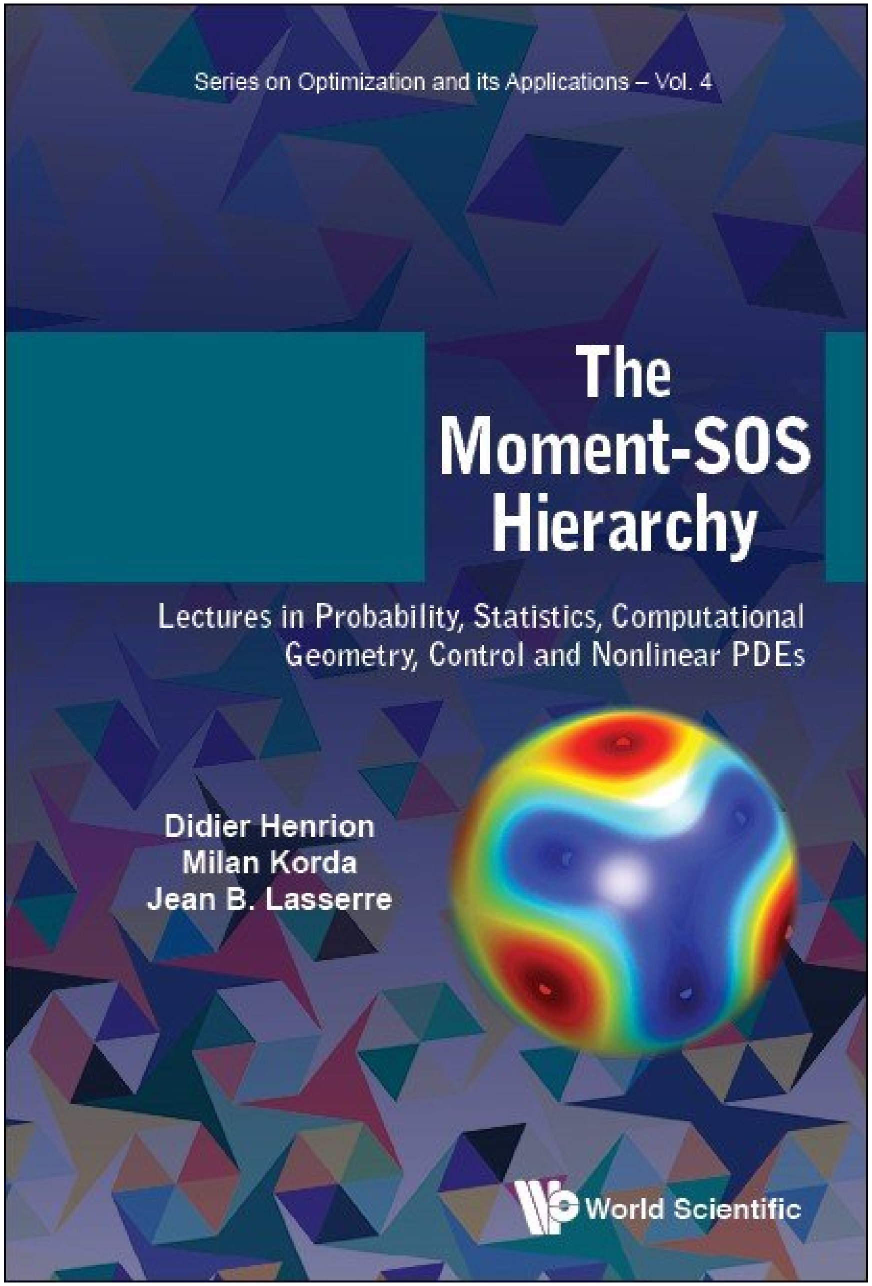

Poema

Polynomial Optimization, Efficiency through Moments and Algebra

poema-network.eu

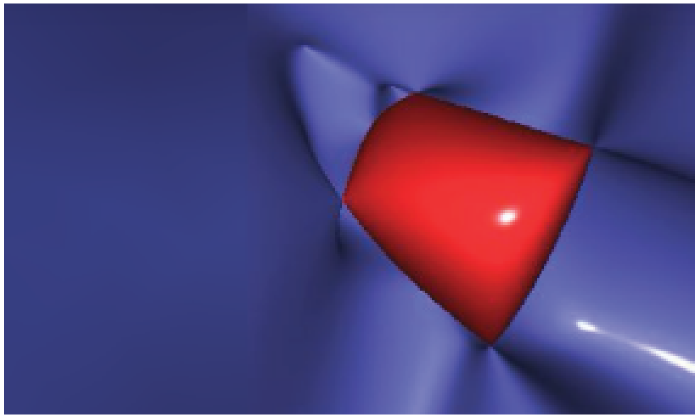

Moment-SOS aka Lasserre hierarchy

Nonlinear nonconvex problem reformulated as infinite-dimensional **linear** optimization problem

Solved approximately with a family of **convex** (semidefinite) relaxations of increasing size indexed by relaxation order $r \in \mathbb{N}$

Based on the **duality** between the cone of positive polynomials and moments and their sum of squares (SOS) and linear matrix inequality (LMI) approximations

Approximate solutions to the nonlinear nonconvex problem can be **extracted** from the solutions of the convex relaxations

1. Polynomial optimization

POP

Given multivariate real polynomials $p, p_1, \ldots, p_k$, solve **globally**

$$\begin{array}{rcll} v^* & := & \min_x & p(x) \\ & & \text{s.t.} & x \in X := \{x \in \mathbb{R}^n : p_k(x) \geq 0,\ k = 1, \ldots, m\} \end{array}$$

where X is bounded and $p \in \mathbb{R}[x]_d$ has degree d

Equivalently

$$\begin{array}{rcll} v^* & := & \max_{v \in \mathbb{R}} & v \\ & & \text{s.t.} & p - v \in P(X) \end{array}$$

where $P(X)$ is the convex cone of positive polynomials on X

However this cone is difficult to manipulate directly

Inner approximations

Since $X := \{x \in \mathbb{R}^n : p_k(x) \geq 0,\ k = 1, \ldots, m\}$ is bounded, we can assume that $p_1(x) = R^2 - \sum_{i=1}^n x_i^2$ for R large enough

Let $p_0(x) := 1$ and for $r \geq d$ define the convex cone

$$Q(X)_r := \{p \in \mathbb{R}[x]_d : p = \sum_{k=0} \underbrace{s_k p_k}_{\in \mathbb{R}[x]_r},\ s_k \text{ SOS}\}$$

Observe that $Q(X)_r \subset Q(X)_{r+1} \subset P(X)$

Theorem [Putinar 1993]: $\overline{Q(X)_\infty} = P(X)$

In words, every positive polynomial on a compact semialgebraic set can be approximated arbitrary well by SOS polynomials

Testing whether a polynomial is SOS reduces to **semidefinite programming** (SDP)

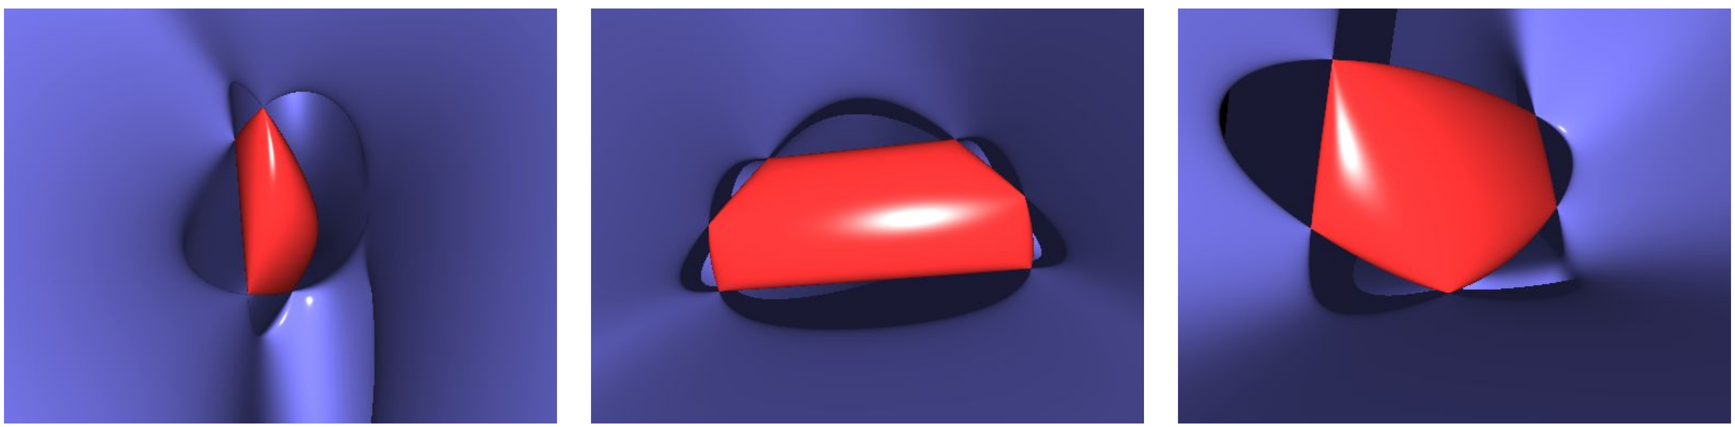

Semidefinite programs can be solved efficiently with primal-dual interior-point methods

SOS hierarchy

Since $Q(X)_r \subset Q(X)_{r+1} \subset P(X)$ we have a hierarchy of SDP problems of increasing size

$$\begin{array}{rcll} v_r^* & := & \max_{v\in\mathbb{R}} & v \\ & & \text{s.t.} & p - v \in Q(X)_r \end{array}$$

yielding a converging monotone sequence of lower bounds

$$v_r^* \leq v_{r+1}^* \leq \cdots \leq v_\infty^* = v^*$$

At a given r^* we want to detect if the bound is **exact**: $v^*_{r^*} = v^*$

For that **convex duality** is essential [Lasserre 2001]

Primal formulation on **positive measures**

$$
\begin{array}{rcl}
v^* & = & \min_\mu \quad \int p(x) d\mu(x) \\
 & & \text{s.t.} \quad \underbrace{\begin{array}{l} \int d\mu(x) = 1 \\ \mu \in C(X)'_+ \end{array}}_{\mu \in \text{Prob}(X)}
\end{array}
$$

with dual on **positive continuous functions**

$$
\begin{array}{rcll}
v^* & = & \max_{v \in \mathbb{R}} & v \\
 & & \text{s.t.} & p - v \in C(X)_+
\end{array}
$$

At a given r^* we want to detect if the bound is **exact**: $v^*_{r^*} = v^*$

For that **convex duality** is essential [Lasserre 2001]

Primal formulation on **positive measures** and **moments**

$$\begin{array}{rcll} v^* & = & \min_\mu & \int p(x)d\mu(x) \\ & & \text{s.t.} & \underbrace{\int d\mu(x) = 1 \quad \mu \in C(X)'_+}_{\mu \in \text{Prob}(X)} \end{array}$$

$$\begin{array}{rcll} v^* & = & \min_y & \sum_a p_a y_a \\ & & \text{s.t.} & y_0 = 1 \\ & & & y \in P(X)' \end{array}$$

with dual on **positive continuous functions** and **polynomials**

$$\begin{array}{rcll} v^* & = & \max_{v\in\mathbb{R}} & v \\ & & \text{s.t.} & p - v \in C(X)_+ \end{array}$$

$$\begin{array}{rcll} v^* & = & \max_{v\in\mathbb{R}} & v \\ & & \text{s.t.} & p - v \in P(X) \end{array}$$

Moments

Let $(b_a(x))_{a \in \mathbb{N}^n_d}$ denote a basis of vector space $\mathbb{R}[x]_d$ indexed in $\mathbb{N}^n_d := \{a \in \mathbb{N}^n : \sum_{k=1}^n a_k \leq d\}$ of cardinality $\binom{n+d}{n}$

The polynomial p can then be written as

$$p(x) = \sum_{a \in \mathbb{N}^n_d} p_a b_a(x)$$

and the objective function can be written as

$$\int p(x) d\mu(x) = \sum_{a \in \mathbb{N}^n_d} p_a y_a$$

which is a linear function of the **moments** of measure μ

$$y_a = \int_X b_a(x) d\mu(x)$$

Moment-SOS hierarchy

So we have a primal moment hierarchy

$$\begin{array}{rcll} v_r^* & = & \min_y & \sum_a p_a y_a \\ & & \text{s.t.} & y_0 = 1 \\ & & & y \in Q(X)'_r \end{array}$$

with explicit LMI relaxations of the cone of moments on X (called **pseudo-moments**, or pseudo-expectations) whose dual is the SOS hierarchy

$$\begin{array}{rcll} v_r^* & := & \max_{v \in \mathbb{R}} & v \\ & & \text{s.t.} & p - v \in Q(X)_r \end{array}$$

In the primal hierarchy, global optimality is ensured whenever y are moments of the Dirac measure at a global optimum x^*

... or more generally, whenever y are moments of a measure **concentrated** on global optima $X^* := \{x \in X : p(x) = v^*\}$

Extracting global optimizers

To certify exactness, we can post-process the solution of the primal SDP and check the rank of the so-called **moment matrix**

$$M_r(y) := \left(\int b_{a_r}(x) b_{a_c}(x) d\mu(x) \right)_{a_r, a_c \in \mathbb{N}^n_r}$$

If the rank of $M_r(y)$ does not increase when r increases, then the moment relaxation is exact [Curto & Fialkow 1991]

Global solutions extracted by linear algebra, as implemented in our Matlab interface GloptiPoly [H & Lasserre 2003]

Exactness at finite relaxation order is generic [Nie 2014]

Approximating global optimizers

Since the moment matrix is positive semidefinite, it holds

$$M_d(y) = PEP'$$

where P is an orthonormal matrix whose columns are denoted p_i and E is a diagonal matrix of eigenvalues $e_{i+1} \geq e_i \geq 0$

Each column p_i is the vector of coefficients in basis $b(x)$ of a polynomial $p_i(x)$, so that

$$p_i' M_d(y) p_i = \int p_i^2(x) d\mu(x) = e_i$$

Let $r \in \mathbb{N}$ and define the **Christoffel-Darboux** polynomial SOS

$$p_{\text{CD}}(x) := \sum_{i=1}^{r} p_i^2(x)$$

Given $\beta > 0$, let $\gamma := \sum_{i=1}^{r} e_i/\beta$ so that $\mu(\{x : p_{\text{CD}}(x) \leq \gamma\}) \geq 1-\beta$

Hence the measure is concentrated on small sublevel sets of the Christoffel-Darboux polynomial [Lasserre & Pauwels 2019]

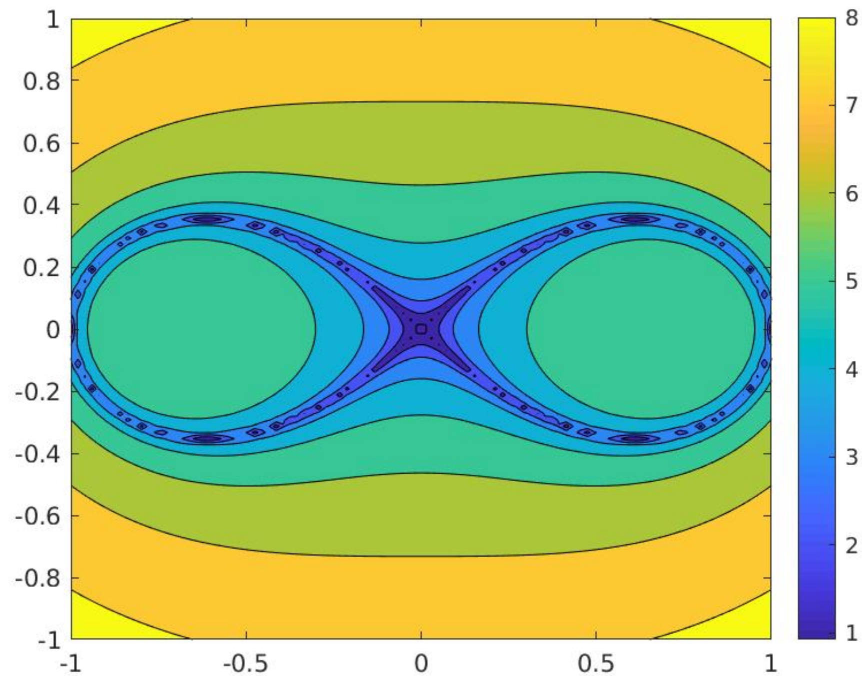

Moment matrix of order 4 and size 15 for the POP $\min_{x\in\mathbb{R}^2}(x_1^2+x_2^2)^2-x_1^2+x_2^2$

2. Polynomial optimal control

POC

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

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

All the given data f, l, l_T are polynomial
and the given sets X, X_T, U are compact 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**

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)

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

with Hamiltonian conjugate to the Lagrangian

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

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

The value function can be discontinuous and complicated

And there are additional difficulties...

No optimal control

Bolza problem

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

Note that the cost is **nonconvex** in the control

Let us construct a minimizing sequence...

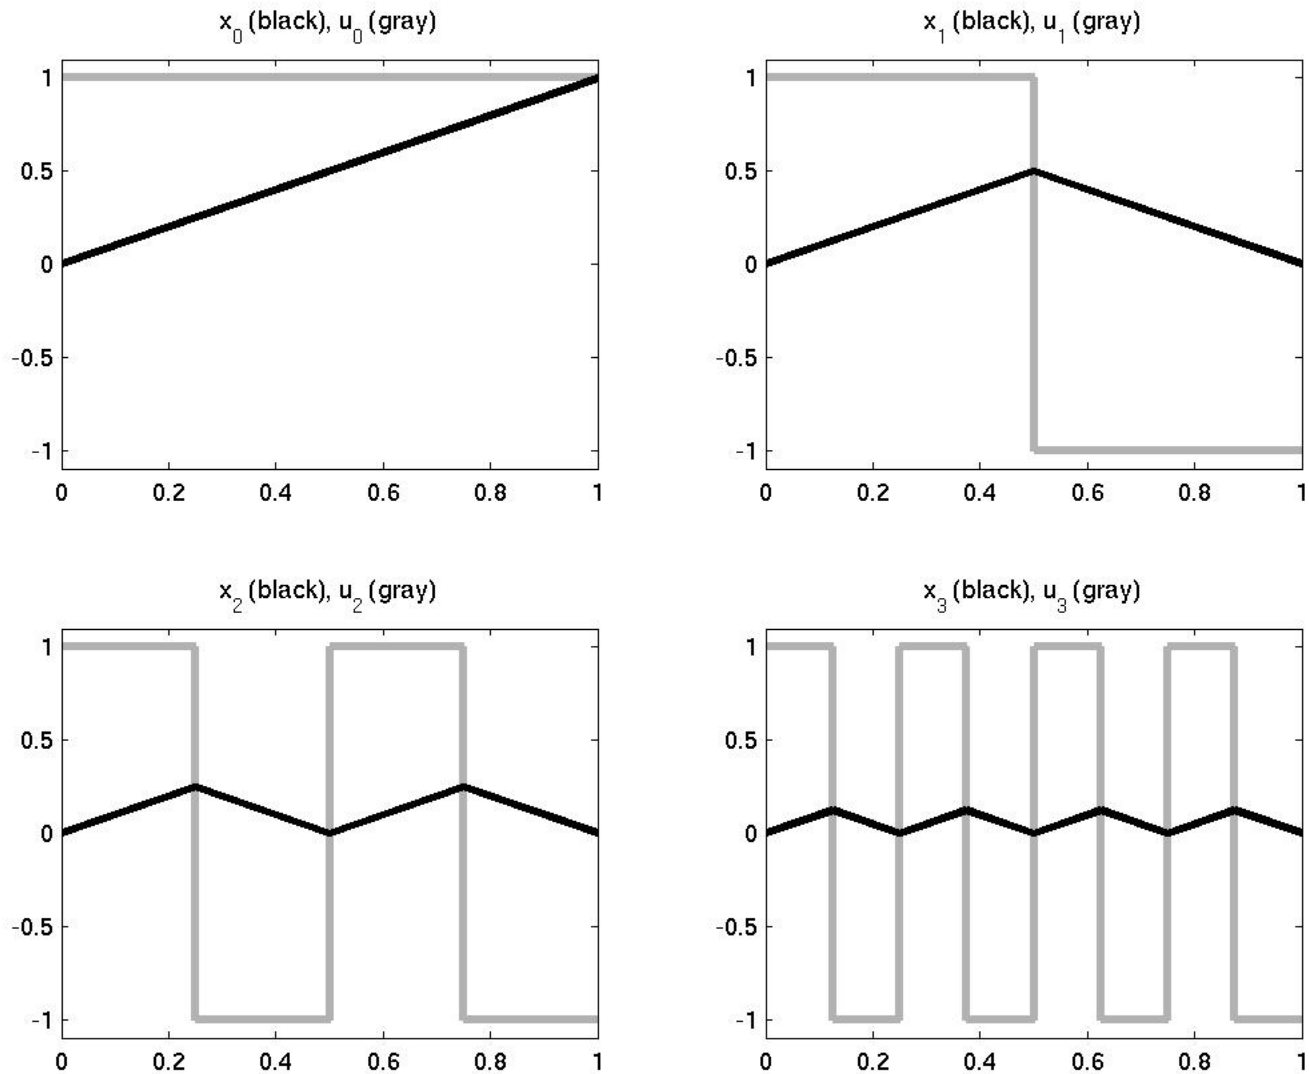

x0 (black), u0 (gray)
x1 (black), u1 (gray)
x2 (black), u2 (gray)
x3 (black), u3 (gray)

The infimum $v^*(0,0) = 0$ 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

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

$$t \mapsto \omega_t(du) = \omega(du|t) \in \mathrm{Prob}(U)$$

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

$$\dot{x}_t = \int_U f(x_t, u)\omega_t(du), \quad \omega_t \in \mathsf{Prob}(U)$$

or equivalently a convex differential inclusion

$$\dot{x}_t \in \mathsf{Conv}\{f(x_t, u) : u \in U\}$$

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

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

The classical Bolza problem

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

becomes the relaxed Bolza problem

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

There is no relaxation gap: $v^*(0,0) = v_R^*(0,0) = 0$
and the relaxed infimum is attained at $\omega_t^* = \frac{1}{2}(\delta_{-1} + \delta_{+1})$

Let's relax

The classical POC problem

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

becomes a relaxed POC problem

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

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

Not relaxed enough

The POC problem is now linear in the relaxed control ω_t but it remains nonlinear in the state x

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

The occupation measure $\mu := dt\,\omega_t\,\delta_{x_t}$ and the terminal measure $\mu_T := \delta_{(T,x_T)}$ solve the **Liouville equation**

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

which should be understood in the weak sense, i.e.

$$\int v\mu_T = v(t_0,x_0) + \int \left(\frac{\partial v}{\partial t} + \mathrm{grad}\, v \cdot f\right)\mu$$

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

The non-linear ODE $\dot{x} = f(x,u)$ has been relaxed to a **linear** transport PDE on measures

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)

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

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

It can be shown that there is **no relaxation gap**: $p^* = v^*$

The primal measure LP

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

has a dual LP

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

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

It can be shown that there is **no duality gap**: $p^* = d^*$

Convergence guarantees

Dual LP

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

Lemma: for every admissible v it holds $v^* \geq v$ on $[t_0,T]\times X$

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

Subsolutions to HJB PDE [Lasserre, H, Prieur, Trélat. SICON 2008]

Theorem [H & Pauwels 2017]**:** For any admissible $(v_r)_{r\in\mathbb{N}}$ and 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 words, the gap between the value function and its lower bound decreases **uniformly in time** along optimal trajectories

Convergence in space is pointwise but we can do better...

In the Liouville equation

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

instead of a Dirac right hand side we can use a general probability measure $\xi_0 \in \mathrm{Prob}(X)$ 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

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

Given an initial condition x_0 and a relaxed control ω_t, let $(x_t)_{t\in[t_0,T]}$ be the solution to the controlled ODE

Let ξ_t denote the image measure of ξ_0 through the flow map $x_0 \mapsto x_t$, such that $\xi_t(A) := \xi_0(\{x_0 : F_t(x_0) \in A\})$ for all $A \subset X$

The averaged occupation measure writes $d\mu(t, x, u) = dt\omega_t(du)\xi_t(dx)$

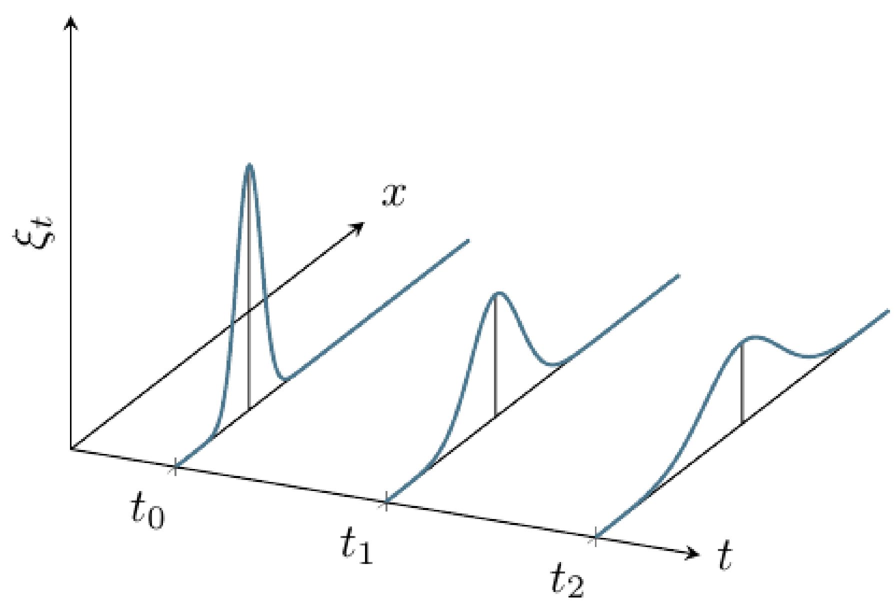

The value function also becomes averaged

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

and it matches the primal LP averaged value

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

and the dual LP averaged value

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

Theorem [H & Pauwels 2017]**:** For any maximizing sequence $(v_r)_{r\in\mathbb{N}}$ it holds

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

Hence by transporting a probability measure ξ_0,
we have $L_1(\xi_0)$ **convergence** to the value function

Now let's compute

To solve the primal LP

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

and dual LP

$$\begin{array}{ll} \sup_v & \int v\mu_0 \\ \text{s.t.} & l + \frac{\partial v}{\partial t} + \mathrm{grad}\, v\cdot f \in C([t_0,T]\times X\times U)_+ \\ & l_T - v(T,.) \in C(\{T\}\times X_T)_+ \end{array}$$

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 $Q(.)_r$ for increasing relaxation order r and at the price of solving **SDP problems** of increasing size we get **pseudo-moments** and **polynomials** v_r in $\mathbb{R}[x]_r$

Numerical examples

Turnpike control

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

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

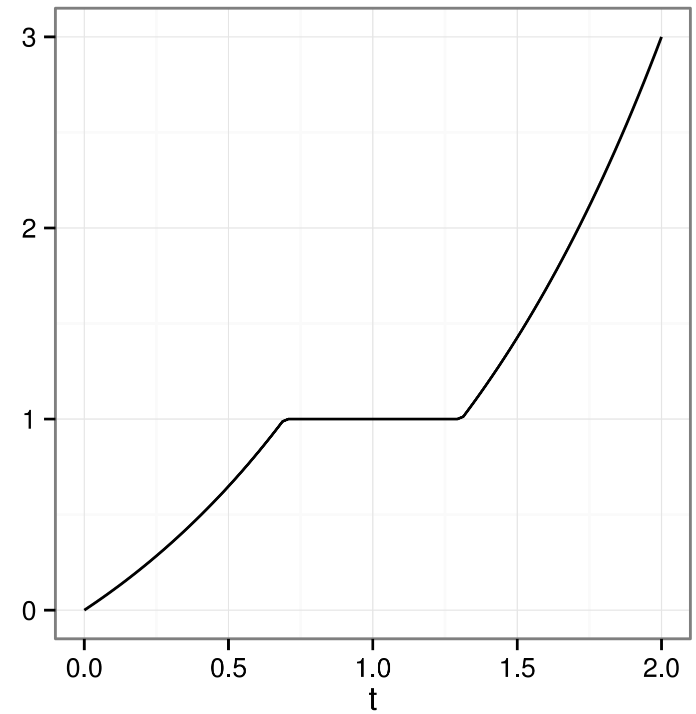

optimal trajectory x_t starting at $(t_0, x_0) = (0, 0)$

Turnpike control

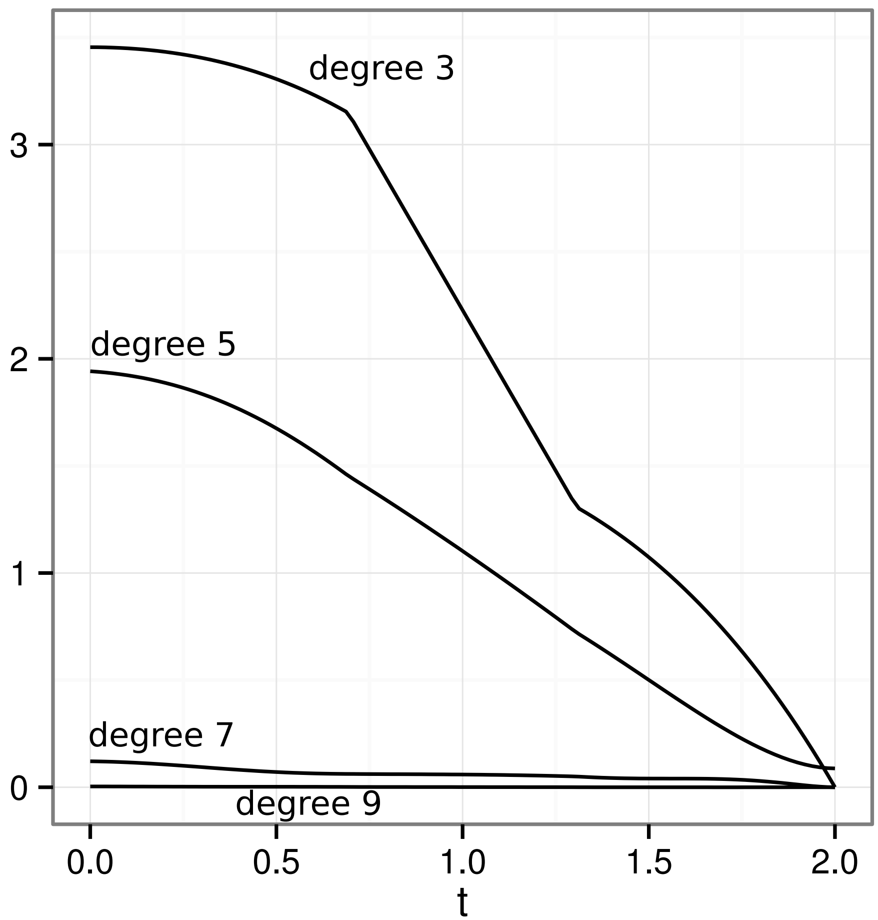

Differences $t \mapsto v^*(t, x_t) - v_r(t, x_t)$ between the actual value function and its poly. approx. of deg. $r = 3, 5, 7, 9$ along the optimal trajectory starting at $(t_0, x_0) = (0, 0)$

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

LQR set control

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

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

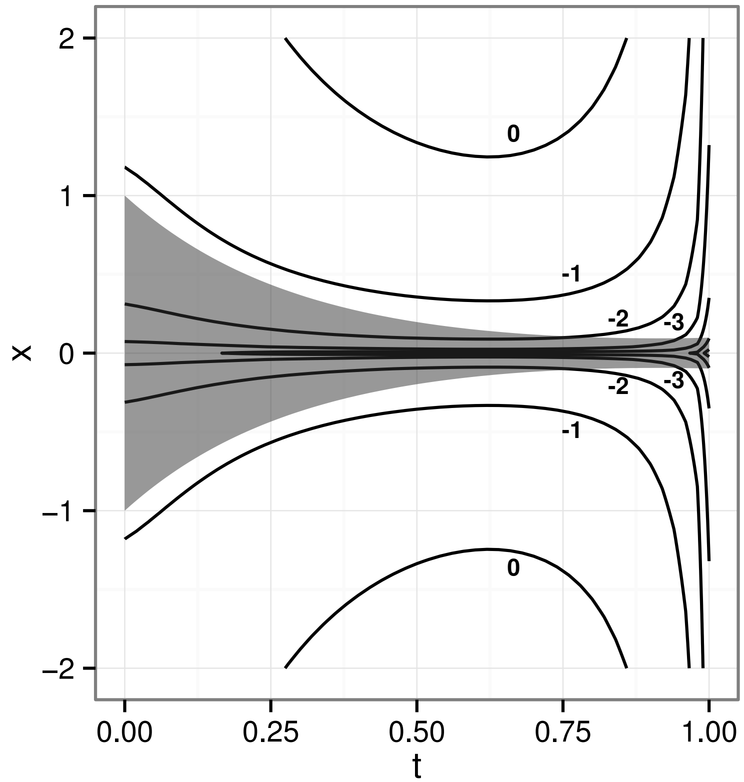

Contour lines of $(t, x) \mapsto \log(v^*(t, x) - v_6(t, x))$ with v_6 poly. approx. of deg. 6 to actual value function v^* obtained by transporting the Lebesgue measure on $[-1, 1]$

Minimum time double integrator with state constraints

With the moment-SOS hierarchy, we compute the pseudo-moments of degree 8 of the occupation measure, and we construct the moment matrices of size 45 of the control and state marginals

For each time we minimize the respective Christoffel-Darboux polynomial (from left to right: control, first and second state, red curves to be compared with the analytic solutions in black)

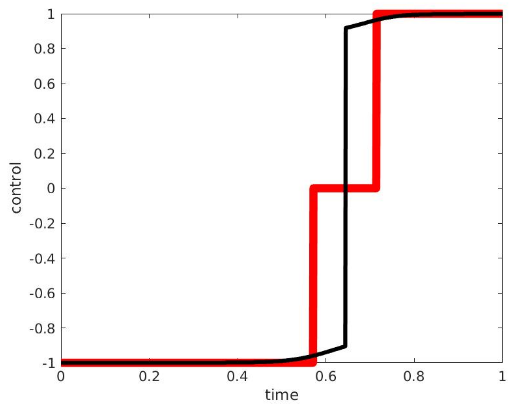

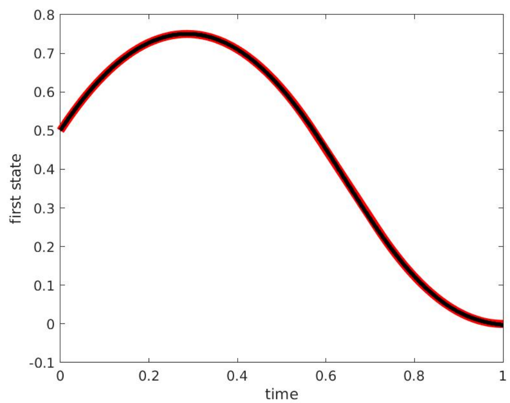

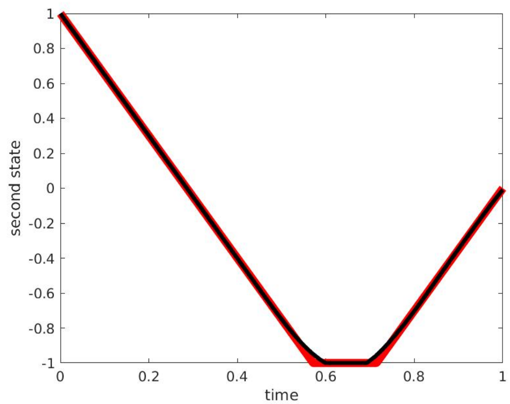

Take-home messages

Polynomial optimization (POP) and optimal control (POC) can be solved approximately with the **moment-SOS hierarchy**

Non-linear non-convex problems reformulated as primal **linear** problems on probability measures or occupation measures

Dual linear SOS problems give **bounds** on the optimal value with convergence guarantees

From the primal solutions we can **certify** global optimality (linear algebra on the moment matrix) and/or **extract** approximate solutions (Christoffel-Darboux polynomial)

Current research directions

Exploit various kinds of sparsity to improve scability of the moment-SOS hierarchy

From optimal control of ODEs to SDEs and PDEs

Occupation measures on infinite-dimensional spaces

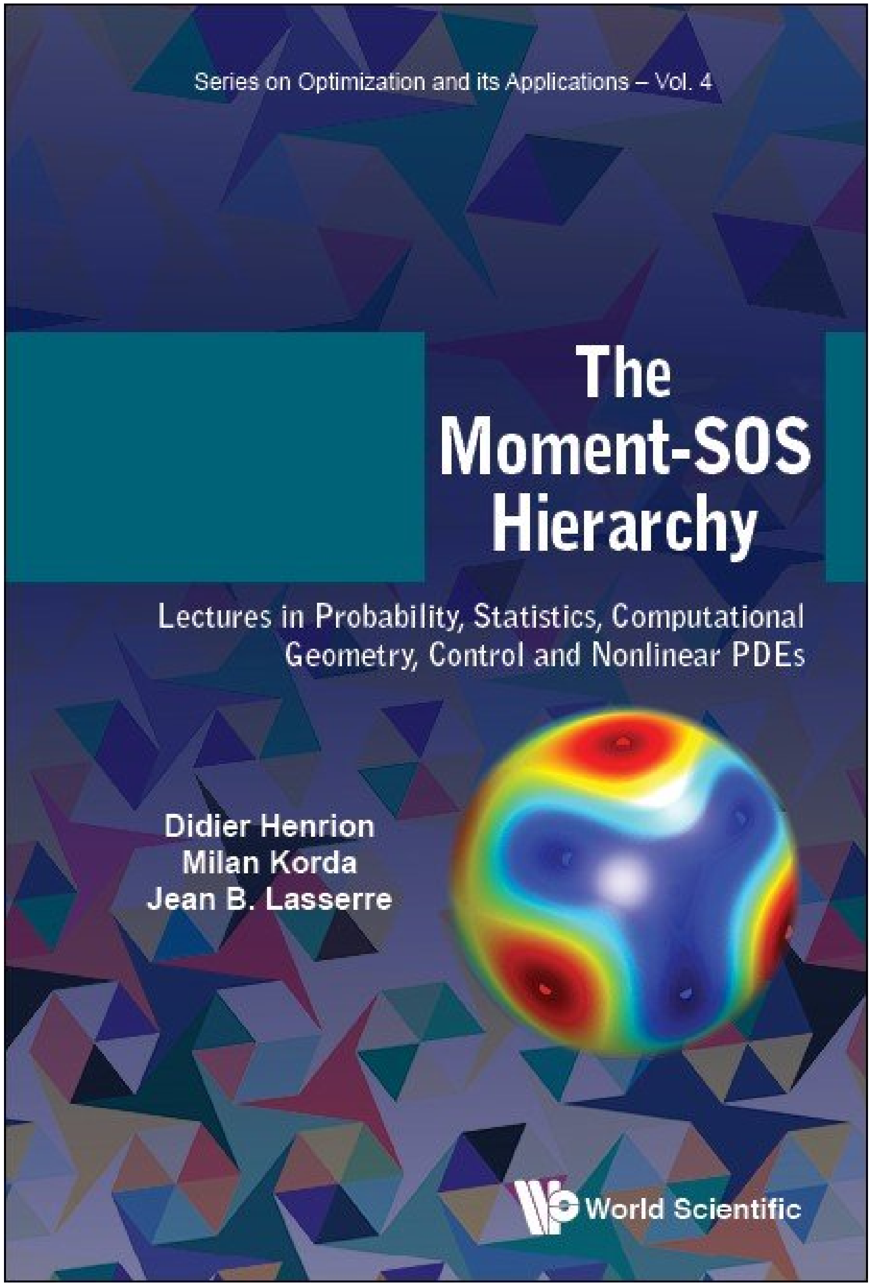

Poema

Polynomial Optimization, Efficiency through Moments and Algebra

poema-network.eu