Spatial Decomposition/Coordination Methods for Stochastic Optimal Control Problems

Practical aspects and theoretical questions

P. Carpentier, J-Ph. Chancelier,
M. De Lara, V. Leclère

École des Ponts ParisTech

16 December 2014

Large scale storage systems stand as powerful motivation

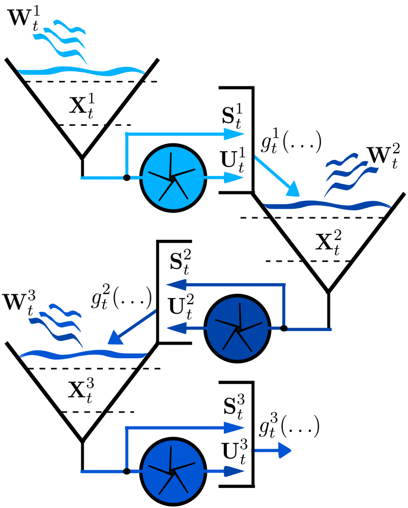

- The Optimization and Systems team was created in 2000 at École des Ponts ParisTech with emphasis on stochastic optimization
- Since 2011, we witness a growing demand from (small and large) energy firms for stochastic optimization, fueled by a deep and fast transformation of power systems

Something is changing in power systems

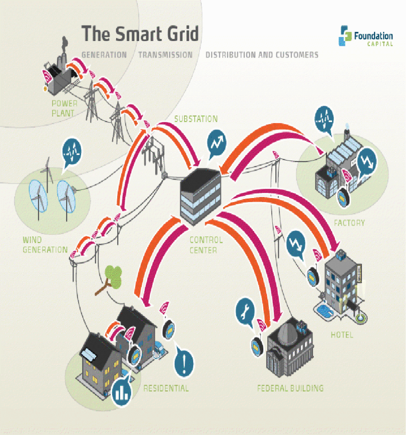

- Renewable energies penetration, telecommunication technologies and markets remold power systems and challenge optimization
- More renewable energies $\rightarrow$ more unpredictability + more variability $\rightarrow$
  - more storage $\rightarrow$ more dynamic optimization, optimal control
  - more stochastic optimization

hence, stochastic optimal control (SOC)

Lecture outline

1. Decomposition and coordination
   - A bird's eye view of decomposition methods
   - A brief insight into Progressive Hedging
   - Spatial decomposition methods in the deterministic case
   - The stochastic case raises specific obstacles
2. Dual approximate dynamic programming (DADP)
   - Problem statement
   - DADP principle and implementation
   - Numerical results on a small size problem
3. Theoretical questions
   - Existence of a saddle point
   - Convergence of the Uzawa algorithm
   - Convergence w.r.t. information
4. Summary and research agenda

A long-term effort in our group

1976 A. Benveniste, P. Bernhard, G. Cohen, "On the decomposition of stochastic control problems", *IRIA-Laboria research report*, No. 187, 1976.

1996 P. Carpentier, G. Cohen, J.-C. Culioli, A. Renaud, "Stochastic optimization of unit commitment: a new decomposition framework", *IEEE Transactions on Power Systems*, Vol. 11, No. 2, 1996.

2006 C. Strugarek, "Approches variationnelles et autres contributions en optimisation stochastique", *Thèse de l'ENPC*, mai 2006.

2010 K. Barty, P. Carpentier, P. Girardeau, "Decomposition of large-scale stochastic optimal control problems", *RAIRO Operations Research*, Vol. 44, No. 3, 2010.

2014 V. Leclère, "Contributions to decomposition methods in stochastic optimization", *Thèse de l'Université Paris-Est*, juin 2014.

Let us fix problem and notations

$$\min_{\mathbf{u},\mathbf{x}} \overbrace{\mathbb{E}}^{\text{"risk-neutral"}} \left(\sum_{i=1}^{N} \left(\sum_{t=0}^{T-1} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + K^i(\mathbf{x}_T^i) \right) \right)$$

subject to dynamics constraints

$$\underbrace{\mathbf{x}_{t+1}^i}_{\text{state}} = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \underbrace{\mathbf{w}_{t+1}}_{\text{uncertainty}}) , \quad \mathbf{x}_0^i = f_{-1}^i(\mathbf{w}_0)$$

to measurability constraints on the control $\mathbf{u}_t^i$

$$\mathbf{u}_t^i \preceq \sigma(\mathbf{w}_0, \dots, \mathbf{w}_t) \iff \mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_0, \dots, \mathbf{w}_t \right)$$

and to instantaneous coupling constraints

$$\sum_{i=1}^{N} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

(The letter u stands for the Russian word for control: *upravlenie*)

Couplings for stochastic problems

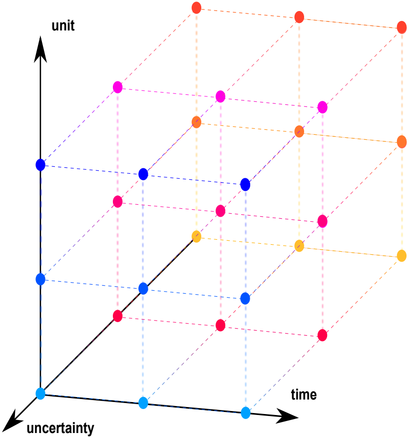

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

Couplings for stochastic problems: in time

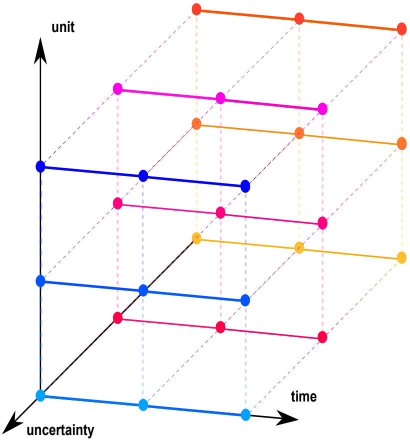

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

Couplings for stochastic problems: in uncertainty

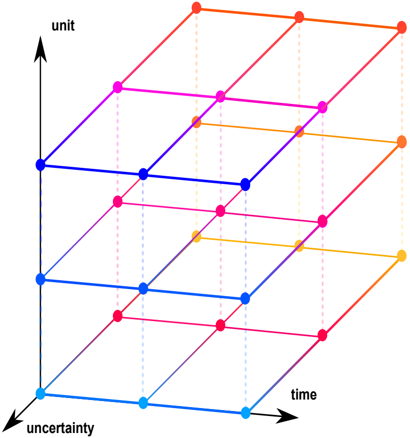

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_0, \ldots, \mathbf{w}_t\right)$$

Couplings for stochastic problems: in space

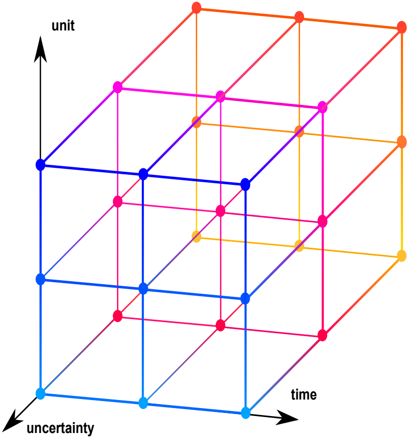

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_0, \ldots, \mathbf{w}_t\right)$$

$$\sum_{i} \Theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Can we decouple stochastic problems?

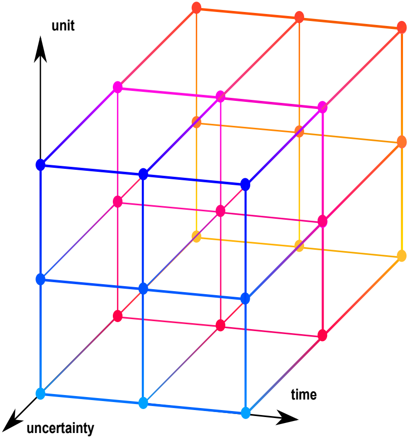

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t. } \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_0, \ldots, \mathbf{w}_t\right)$$

$$\sum_{i} \Theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Decompositions for stochastic problems: in time

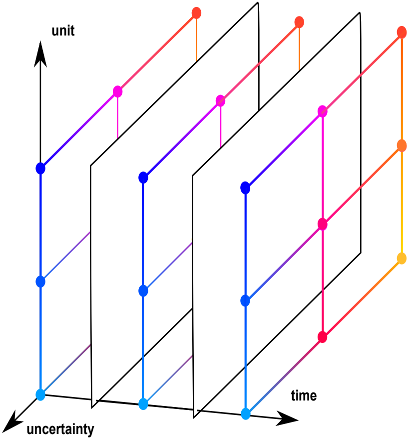

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_0, \dots, \mathbf{w}_t\right)$$

$$\sum_{i} \Theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Dynamic Programming
Bellman (56)

Decompositions for stochastic problems: in uncertainty

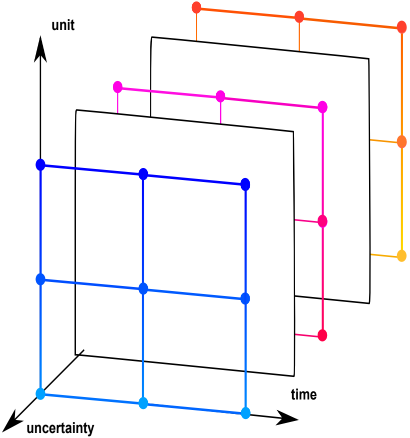

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_0, \ldots, \mathbf{w}_t\right)$$

$$\sum_{i} \Theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Progressive Hedging

Rockafellar - Wets (91)

Decompositions for stochastic problems: in space

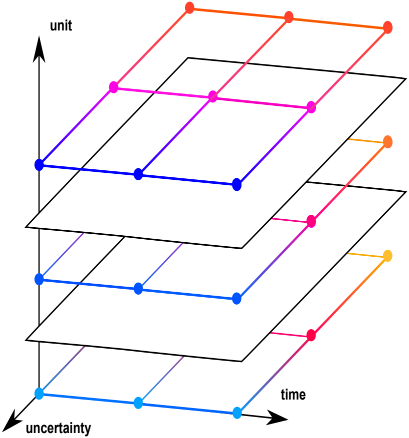

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\text{s.t.} \quad \mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_0, \dots, \mathbf{w}_t\right)$$

$$\sum_{i} \Theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Dual Approximate
Dynamic Programming

Outline of the presentation

Non-anticipativity constraints are linear

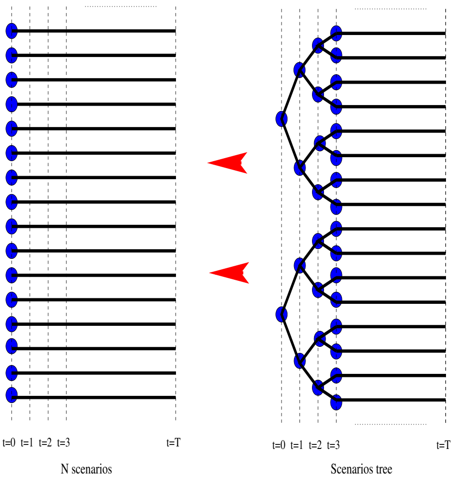

- From tree to scenarios (comb)
- Equivalent formulations of the non-anticipativity constraints
  - pairwise equalities
  - all equal to their mathematical expectation
- Linear structure

$$\mathbf{u}_t = \mathbb{E}\left(\mathbf{u}_t \;\middle|\; \mathbf{w}_0, \ldots, \mathbf{w}_t\right)$$

Progressive Hedging stands as a scenario decomposition method

We dualize the non-anticipativity constraints

- When the criterion is strongly convex,
  we use a Lagrangian relaxation (algorithm "à la Uzawa")
  to obtain a scenario decomposition
- When the criterion is linear,
  Rockafellar - Wets (91) propose to use an augmented Lagrangian,
  and obtain the Progressive Hedging algorithm

Decomposition and coordination

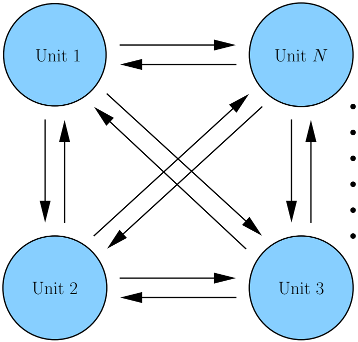

Interconnected units

- The system to be optimized consists of interconnected subsystems
- We want to use this structure to formulate optimization subproblems of reasonable complexity
- But the presence of interactions requires a level of coordination
- Coordination iteratively provides a local model of the interactions for each subproblem
- We expect to obtain the solution of the overall problem by concatenation of the solutions of the subproblems

Example: the "flower model"

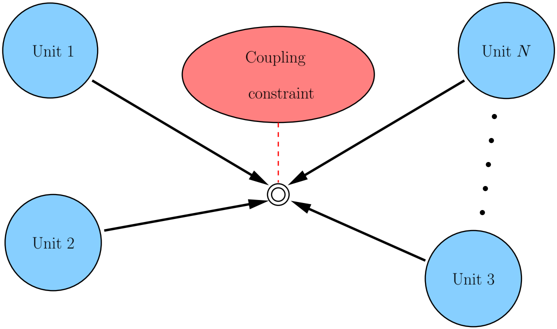

$$\min_u \sum_{i=1}^N J_i(u_i)$$

$$\text{s.t.} \quad \sum_{i=1}^N \theta_i(u_i) = 0$$

Unit Commitment Problem

Intuition of spatial decomposition

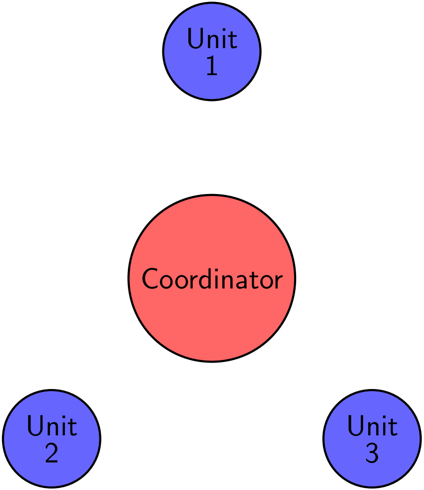

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ
  - the units send their optimal decision u_i
  - the coordinator compares total production $\sum_{i=1}^{N} \theta_i(u_i)$ and demand, and then updates the price accordingly
  - and so on...

Intuition of spatial decomposition

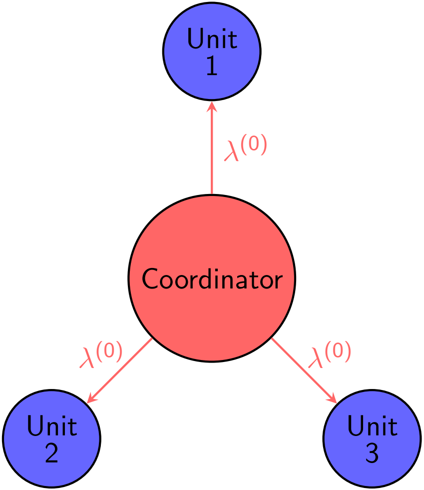

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ
  - the units send their optimal decision u_i
  - the coordinator compares total production $\sum_{i=1}^{N} \theta_i(u_i)$ and demand, and then updates the price accordingly
  - and so on...

Intuition of spatial decomposition

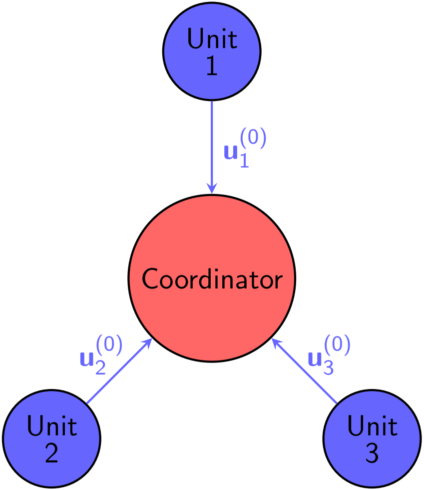

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ
  - the units send their optimal decision $\mathbf{u}_i$
  - the coordinator compares total production $\sum_{i=1}^{N} \theta_i(u_i)$ and demand, and then updates the price accordingly
  - and so on...

Intuition of spatial decomposition

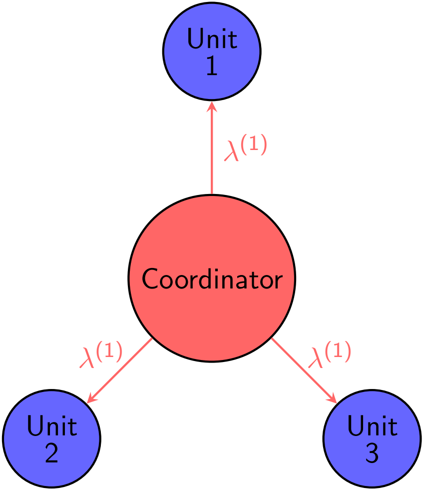

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ
  - the units send their optimal decision $\mathbf{u}_i$
  - the coordinator compares total production $\sum_{i=1}^{N} \theta_i(u_i)$ and demand, and then updates the price accordingly
  - and so on...

Intuition of spatial decomposition

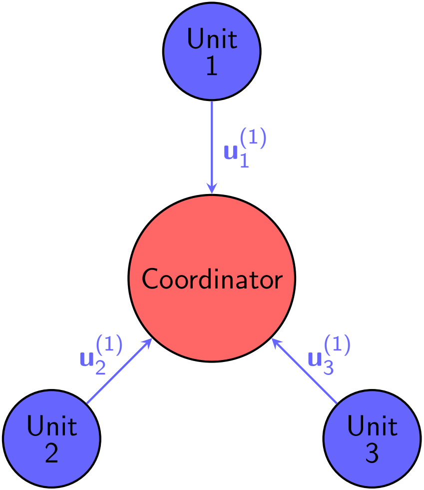

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ
  - the units send their optimal decision $\mathbf{u}_i$
  - the coordinator compares total production $\sum_{i=1}^{N} \theta_i(u_i)$ and demand, and then updates the price accordingly
  - and so on...

Intuition of spatial decomposition

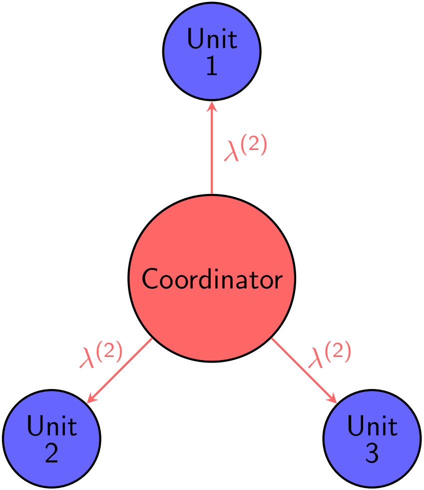

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ
  - the units send their optimal decision $\mathbf{u}_i$
  - the coordinator compares total production $\sum_{i=1}^{N} \theta_i(u_i)$ and demand, and then updates the price accordingly
  - and so on...

Intuition of spatial decomposition

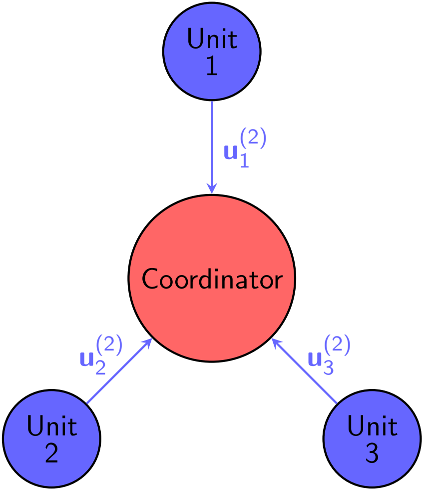

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ
  - the units send their optimal decision $\mathbf{u}_i$
  - the coordinator compares total production $\sum_{i=1}^{N} \theta_i(u_i)$ and demand, and then updates the price accordingly
  - and so on...

Price decomposition relies on dualization

$$\min_{u_i \in \mathcal{U}_i, i=1\ldots N} \sum_{i=1}^{N} J_i(u_i) \quad \text{subject to} \quad \sum_{i=1}^{N} \theta_i(u_i) = 0$$

1. Form the Lagrangian and assume that a saddle point exists

$$\max_{\lambda \in \mathcal{V}} \min_{u_i \in \mathcal{U}_i, i=1\ldots N} \sum_{i=1}^{N} \Big(J_i(u_i) + \langle \lambda \, , \theta_i(u_i) \rangle \Big)$$

2. Solve this problem by the dual gradient algorithm "à la Uzawa"

$$u_i^{(k+1)} \in \arg\min_{u_i \in \mathcal{U}_i} J_i(u_i) + \left\langle \lambda^{(k)} \, , \theta_i(u_i) \right\rangle \, , \quad i = 1 \ldots , N$$

$$\lambda^{(k+1)} = \lambda^{(k)} + \rho \sum_{i=1}^{N} \theta_i(u_i^{(k+1)})$$

Remarks on decomposition methods

- The theory is available for infinite dimensional Hilbert spaces, and thus applies in the stochastic framework, that is, when the $\mathcal{U}_i$ are spaces of random variables
- The minimization algorithm used for solving the subproblems is not specified in the decomposition process
- New variables $\lambda^{(k)}$ appear in the subproblems arising at iteration k of the optimization process

$$\min_{u_i \in \mathcal{U}_i} J_i(u_i) + \left\langle \lambda^{(k)} , \theta_i(u_i) \right\rangle$$

- These variables are fixed when solving the subproblems, and do not cause any difficulty, at least in the deterministic case

Price decomposition applies to various couplings

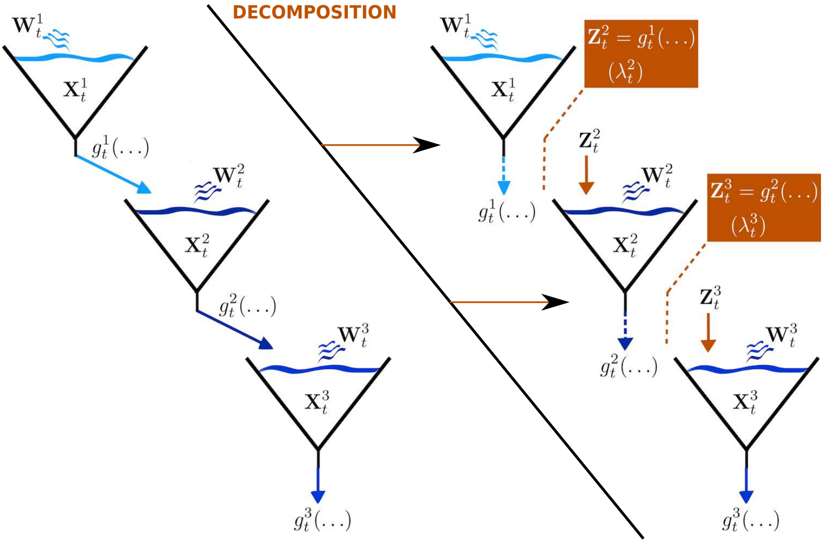

Stochastic optimal control (SOC) problem formulation

Consider the following SOC problem

$$\min_{\mathbf{u},\mathbf{x}} \ \mathbb{E}\bigg(\sum_{i=1}^{N} \Big(\sum_{t=0}^{T-1} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + K^i(\mathbf{x}_T^i) \Big) \bigg)$$

subject to the constraints

$$\begin{aligned} \mathbf{x}_0^i &= f_{-1}^i(\mathbf{w}_0) \ , & i = 1 \dots N \\ \mathbf{x}_{t+1}^i &= f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) \ , \qquad t = 0 \dots T-1 \ , & i = 1 \dots N \end{aligned}$$

$$\mathbf{u}_t^i \preceq \mathcal{F}_t = \sigma(\mathbf{w}_0, \dots, \mathbf{w}_t) \ , \ \ t = 0 \dots T-1 \ , \ \ i = 1 \dots N$$

$$\sum_{i=1}^{N} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0 \ , \qquad t = 0 \dots T-1$$

Stochastic optimal control (SOC) problem formulation

Consider the following SOC problem

$$\min_{\mathbf{u},\mathbf{x}} \sum_{i=1}^{N} \left(\mathbb{E}\Big(\sum_{t=0}^{T-1} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + K^i(\mathbf{x}_T^i) \Big) \right)$$

subject to the constraints

$$\begin{aligned}
\mathbf{x}_0^i &= f_{-1}^i(\mathbf{w}_0) , & i = 1 \dots N \\
\mathbf{x}_{t+1}^i &= f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) , \quad & t = 0 \dots T-1 , \ i = 1 \dots N
\end{aligned}$$

$$\mathbf{u}_t^i \preceq \mathcal{F}_t = \sigma(\mathbf{w}_0, \dots, \mathbf{w}_t) , \ t = 0 \dots T-1 , \ i = 1 \dots N$$

$$\sum_{i=1}^{N} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0 , \qquad t = 0 \dots T-1$$

Dynamic programming yields centralized controls

- As we want to solve this SOC problem using dynamic programming (**DP**), we suppose to be in the Markovian setting, that is, $\mathbf{w}_0, \ldots, \mathbf{w}_T$ are a white noise
- The system is made of N interconnected subsystems, with the control $\mathbf{u}_t^i$ and the state $\mathbf{x}_t^i$ of subsystem i at time t
- The optimal control $\mathbf{u}_t^i$ of subsystem i is a function of the whole system state $\left(\mathbf{x}_t^1, \ldots, \mathbf{x}_t^N\right)$

$$\mathbf{u}_t^i = \gamma_t^i\big(\mathbf{x}_t^1, \ldots, \mathbf{x}_t^N\big)$$

Naive decomposition should lead to decentralized feedbacks

$$\mathbf{u}_t^i = \widehat{\gamma}_t^i(\mathbf{x}_t^i)$$

which are, in most cases, far from being optimal...

Straightforward decomposition of dynamic programming?

The crucial point is that the optimal feedback of a subsystem a priori depends on the state of all other subsystems, so that using a decomposition scheme by subsystems is not obvious...

As far as we have to deal with dynamic programming, the central concern for decomposition/coordination purpose boils down to

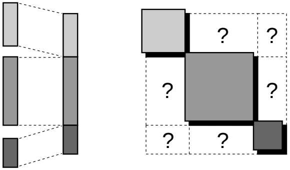

- how to decompose a feedback γ_t w.r.t. its domain $\mathbb{X}_t$ rather than its range $\mathbb{U}_t$?

  And the answer is
- impossible in the general case!

Price decomposition and dynamic programming

When applying price decomposition to the problem by dualizing the (almost sure) coupling constraint $\sum_i \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$,
multipliers $\mathbf{\Lambda}_t^{(k)}$ appear in the subproblems arising at iteration k

$$\min_{\mathbf{u}^i, \mathbf{x}^i} \mathbb{E}\Big(\sum_t L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + \mathbf{\Lambda}_t^{(k)} \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \Big)$$

- The variables $\mathbf{\Lambda}_t^{(k)}$ are fixed random variables, so that the random process $\mathbf{\Lambda}^{(k)}$ acts as an additional input noise in the subproblems
- But this process may be correlated in time, so that the white noise assumption has no reason to be fulfilled
- DP cannot be applied in a straightforward manner!

Question: how to handle the coordination instruments $\mathbf{\Lambda}_t^{(k)}$ to obtain (an approximation of) the overall optimum?

Optimization problem

The SOC problem under consideration reads

$$\min_{\mathbf{u},\mathbf{x}} \ \mathbb{E}\bigg(\sum_{i=1}^{N} \Big(\sum_{t=0}^{T-1} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + K^i(\mathbf{x}_T^i) \Big) \bigg)$$

subject to dynamics constraints

$$\begin{aligned} \mathbf{x}_0^i &= f_{-1}^i(\mathbf{w}_0) \\ \mathbf{x}_{t+1}^i &= f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) \end{aligned}$$

to measurability constraints

$$\mathbf{u}_t^i \preceq \sigma(\mathbf{w}_0, \dots, \mathbf{w}_t)$$

and to instantaneous coupling constraints

$$\sum_{i=1}^{N} \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) = 0$$

Constraints to be **dualized**

Assumptions

Assumption 1 (White noise)

Noises $\mathbf{w}_0, \dots, \mathbf{w}_T$ are independent over time

Hence dynamic programming applies: there is no optimality loss to look after the controls $\mathbf{u}_t^i$ as functions of the state at time t

Assumption 2 (Constraint qualification)

A saddle point of the Lagrangian $\mathcal{L}$ exists (more on that later)

$$\mathcal{L}(\mathbf{x}, \mathbf{u}, \mathbf{\Lambda}) = \mathbb{E}\Bigg(\sum_{i=1}^{N} \bigg(\sum_{t=0}^{T-1} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + K^i(\mathbf{x}_T^i) + \sum_{t=0}^{T-1} \mathbf{\Lambda}_t \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \bigg) \Bigg)$$

where the $\mathbf{\Lambda}_t$ are $\sigma(\mathbf{w}_0, \dots, \mathbf{w}_t)$-measurable random variables

Assumption 3 (Dual gradient algorithm)

Uzawa algorithm applies. . . (more on that later)

Uzawa algorithm

At iteration k of the algorithm,

1. Solve Subproblem i, $i = 1, \dots, N$, with $\mathbf{\Lambda}^{(k)}$ fixed

$$\min_{\mathbf{u}^i, \mathbf{x}^i} \mathbb{E}\left(\sum_{t=0}^{T-1} \left(L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + \mathbf{\Lambda}_t^{(k)} \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \right) + K^i(\mathbf{x}_T^i) \right)$$

subject to

$$\mathbf{x}_{t+1}^i = f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$
$$\mathbf{u}_t^i \preceq \sigma(\mathbf{w}_0, \dots, \mathbf{w}_t)$$

whose solution is denoted $\left(\mathbf{u}^{i,(k+1)}, \mathbf{x}^{i,(k+1)}\right)$

2. Update the multipliers $\mathbf{\Lambda}_t$

$$\mathbf{\Lambda}_t^{(k+1)} = \mathbf{\Lambda}_t^{(k)} + \rho_t \left(\sum_{i=1}^{N} \theta_t^i\left(\mathbf{x}_t^{i,(k+1)}, \mathbf{u}_t^{i,(k+1)}\right) \right)$$

Structure of a subproblem

- Subproblem i reads

$$\min_{\mathbf{u}^i,\mathbf{x}^i} \mathbb{E}\left(\sum_{t=0}^{T-1} \left(L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \underbrace{\mathbf{w}_{t+1}}_{\text{white}}) + \underbrace{\mathbf{\Lambda}_t^{(k)}}_{???} \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \right) \right)$$

subject to

$$\begin{aligned} \mathbf{x}_{t+1}^i &= f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) \\ \mathbf{u}_t^i &\preceq \sigma(\mathbf{w}_0, \ldots, \mathbf{w}_t) \end{aligned}$$

- Without some knowledge of the process $\mathbf{\Lambda}^{(k)}$ (we just know that $\mathbf{\Lambda}_t^{(k)} \preceq (\mathbf{w}_0, \ldots, \mathbf{w}_t)$), the informational state of this subproblem i at time t cannot be summarized by the physical state $\mathbf{x}_t^i$

We outline the main idea in DADP

- The core idea of DADP is to replace
  the multiplier $\boldsymbol{\Lambda}_t^{(k)}$ at iteration k
  by its conditional expectation $\mathbb{E}(\boldsymbol{\Lambda}_t^{(k)} \mid \mathbf{y}_t)$
- where we introduce a new adapted "information" process
  $\mathbf{y} = \big(\mathbf{y}_0, \ldots, \mathbf{y}_{T-1}\big)$
- (More on the interpretation later)

Let us go on with our "trick"

- Using this idea, we replace Subproblem i by

$$\min_{\mathbf{u}^i,\mathbf{x}^i} \mathbb{E}\bigg(\sum_{t=0}^{T-1} \Big(L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) + \mathbb{E}(\mathbf{\Lambda}_t^{(k)} \mid \mathbf{y}_t) \cdot \theta_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i) \Big) + K^i(\mathbf{x}_T^i) \bigg)$$

subject to

$$\begin{aligned} \mathbf{x}_{t+1}^i &= f_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1}) \\ \mathbf{u}_t^i &\preceq \sigma(\mathbf{w}_0, \dots, \mathbf{w}_t) \end{aligned}$$

- The conditional expectation $\mathbb{E}(\mathbf{\Lambda}_t^{(k)} \mid \mathbf{y}_t)$ is an (updated) function of the variable $\mathbf{y}_t$
- so that Subproblem i involves the two noises processes $\mathbf{w}$ and $\mathbf{y}$

If $\mathbf{y}$ follows a dynamical equation, DP applies

We obtain a dynamic programming equation by subsystem

Assuming a non-controlled dynamics $\mathbf{y}^i_{t+1} = h^i_t(\mathbf{y}_t, \mathbf{w}_{t+1})$ for the information process $\mathbf{y}$, the DP equation writes

$$
\begin{aligned}
V^i_T(x,y) &= K^i(x) \\
V^i_t(x,y) &= \min_u \mathbb{E}\Big(L^i_t(x,u,\mathbf{w}_{t+1}) \\
&\qquad\qquad + \mathbb{E}\big(\boldsymbol{\Lambda}^{(k)}_t \mid \mathbf{y}_t = y\big) \cdot \theta^i_t(x,u) \\
&\qquad\qquad\qquad + V^i_{t+1}\big(\mathbf{x}^i_{t+1}, \mathbf{y}_{t+1}\big)\Big)
\end{aligned}
$$

subject to the extended dynamics

$$
\begin{aligned}
\mathbf{x}^i_{t+1} &= f^i_t(x,u,\mathbf{w}_{t+1}) \\
\mathbf{y}^i_{t+1} &= h^i_t(y,\mathbf{w}_{t+1})
\end{aligned}
$$

What have we done?

- Trick: DADP as an approximation of the optimal multiplier

$$\lambda_t \quad \rightsquigarrow \quad \mathbb{E}\big(\lambda_t \mid \mathbf{y}_t\big)$$

- Interpretation: DADP as a decision-rule approach in the dual

$$\max_{\lambda} \min_{\mathbf{u}} L\big(\lambda, \mathbf{u}\big) \quad \rightsquigarrow \quad \max_{\lambda_t \preceq \mathbf{y}_t} \min_{\mathbf{u}} L\big(\lambda, \mathbf{u}\big)$$

- Interpretation: DADP as a constraint relaxation in the primal

$$\sum_{i=1}^{n} \theta_t^i\big(\mathbf{x}_t^i, \mathbf{u}_t^i\big) = 0 \quad \rightsquigarrow \quad \mathbb{E}\Big(\sum_{i=1}^{n} \theta_t^i\big(\mathbf{x}_t^i, \mathbf{u}_t^i\big) \;\Big|\; \mathbf{y}_t\Big) = 0$$

A bunch of practical questions remains open

⋆ How to choose the information variables $\mathbf{y}_t$?

- Perfect memory: $\mathbf{y}_t = \big(\mathbf{w}_0, \dots, \mathbf{w}_t\big)$
- Minimal information: $\mathbf{y}_t \equiv \mathrm{cste}$
- Static information: $\mathbf{y}_t = h_t\big(\mathbf{w}_t\big)$
- Dynamic information: $\mathbf{y}_{t+1} = h_t\big(\mathbf{y}_t, \mathbf{w}_{t+1}\big)$

⋆ How to obtain a feasible solution from the relaxed problem?

- Use an appropriate heuristic!

⋆ How to accelerate the gradient algorithm?

- Augmented Lagrangian
- More sophisticated gradient methods

⋆ How to handle more complex structures than the flower model?

We consider 3 dams in a row, amenable to DP

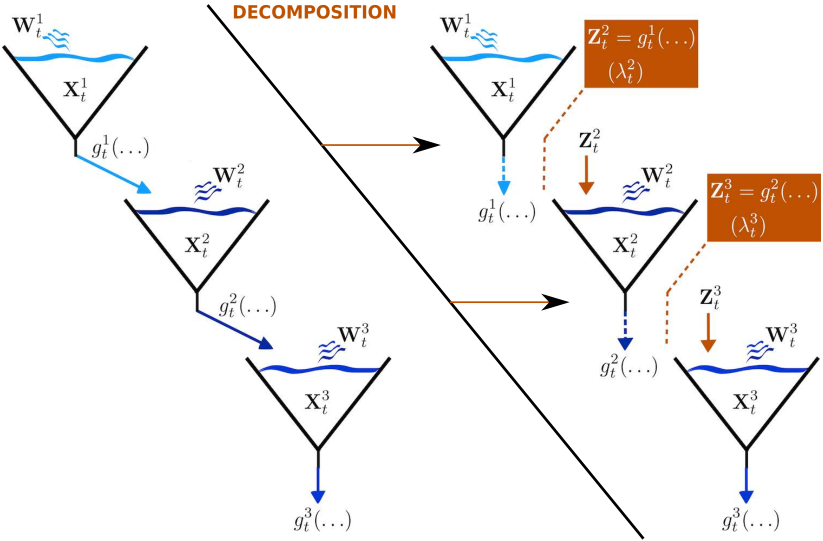

Problem specification

- We consider a 3 dam problem, over 12 time steps
- We relax each constraint with a given information process $\mathbf{y}^i$ that depends on the constraint
- All random variable are discrete (noise, control, state)

- We test the following information processes

  Constant information: equivalent to replace each a.s. constraint by the expected constraint

  Part of noise: the information process depends on the constraint and is the inflow of the above dam $\mathbf{y}^i_t = \mathbf{w}^{i-1}_t$

  Phantom state: the information process mimicks the optimal trajectory of the state of the first dam (by statistical regression over the known optimal trajectory in this case)

Numerical results are encouraging

	DADP - $\mathbb{E}$	DADP - $\mathbf{w}^{i-1}$	DADP - dyn.	DP
Nb of iterations	165	170	25	1
Time (min)	2	3	67	41
Lower Bound	-1.386×10^6	-1.379×10^6	-1.373×10^6	
Final Value	-1.335×10^6	-1.321×10^6	-1.344×10^6	-1.366×10^6
Loss	-2.3%	-3.3%	-1.6%	ref.

$\rightsquigarrow$ *PhD thesis of J.-C. Alais*

Summing up DADP

- Choose an information process $\mathbf{y}$ following $\mathbf{y}_{t+1} = \widetilde{f}_t\big(\mathbf{y}_t, \mathbf{w}_{t+1}\big)$
- Relax the almost sure coupling constraint into a conditional expectation
- Then apply a price decomposition scheme to the relaxed problem
- The subproblems can be solved by dynamic programming with the modest state $\big(\mathbf{x}_t^i, \mathbf{y}_t\big)$
- In the theoretical part, we give
  - Conditions for the existence of an L^1 multiplier
  - Convergence of the algorithm (fixed information process)
  - Consistency result (family of information process)

What are the issues to consider?

- We treat the coupling constraints in a stochastic optimization problem by duality methods
- Uzawa algorithm is a dual method which is naturally described in an Hilbert space, but we cannot guarantee the existence of an optimal multiplier in the space $\mathrm{L}^2\big(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n\big)$!
- Consequently, we extend the algorithm to the non-reflexive Banach space $\mathrm{L}^\infty\big(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n\big)$, by giving a set of conditions ensuring the existence of a $\mathrm{L}^1\big(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n\big)$ optimal multiplier,
  and by providing a convergence result of the algorithm
- We also have to deal with the approximation induced by the information variable: we give an epi-convergence result related to such an approximation

$\rightsquigarrow$ *PhD thesis of V. Leclère*

Abstract formulation of the problem

We consider the following abstract optimization problem

$$(\mathcal{P}) \qquad \min_{\mathbf{u}\in\mathcal{U}^{\mathrm{ad}}} J(\mathbf{u}) \quad \text{s.t.} \quad \Theta(\mathbf{u}) \in -C$$

where $\mathcal{U}$ and $\mathcal{V}$ are two Banach spaces, and

- $J : \mathcal{U} \to \overline{\mathbb{R}}$ is the objective function
- $\mathcal{U}^{\mathrm{ad}}$ is the admissible set
- $\Theta : \mathcal{U} \to \mathcal{V}$ is the constraint function to be dualized
- $C \subset \mathcal{V}$ is the cone of constraint

Here, $\mathcal{U}$ is a space of random variables, and J is defined by

$$J(\mathbf{u}) = \mathbb{E}\big(j(\mathbf{u}, \mathbf{w})\big)$$

1. Decomposition and coordination
   - A bird's eye view of decomposition methods
   - A brief insight into Progressive Hedging
   - Spatial decomposition methods in the deterministic case
   - The stochastic case raises specific obstacles

2. Dual approximate dynamic programming (DADP)
   - Problem statement
   - DADP principle and implementation
   - Numerical results on a small size problem

3. Theoretical questions
   - Existence of a saddle point
   - Convergence of the Uzawa algorithm
   - Convergence w.r.t. information

4. Summary and research agenda

Standard duality in L^2 spaces (I)

Assume that $\mathcal{U} = \mathrm{L}^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n)$ and $\mathcal{V} = \mathrm{L}^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m)$

The standard sufficient constraint qualification condition

$$0 \in \mathrm{ri}\Big(\Theta\big(\mathcal{U}^{\mathrm{ad}} \cap \mathrm{dom}(J)\big) + C\Big)$$

is scarcely satisfied in such a stochastic setting

Proposition 1

If the σ-algebra $\mathcal{F}$ is not finite modulo $\mathbb{P}$,[a]
then for any subset $U^{\mathrm{ad}} \subset \mathbb{R}^n$ that is not an affine subspace, the set

$$\mathcal{U}^{\mathrm{ad}} = \Big\{\mathbf{u} \in \mathrm{L}^p(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n) \mid \mathbf{u} \in U^{\mathrm{ad}} \quad \mathbb{P}-a.s.\Big\}$$

has an empty relative interior in L^p, for any $p < +\infty$

[a] If the σ-algebra is finite modulo $\mathbb{P}$, $\mathcal{U}$ and $\mathcal{V}$ are finite dimensional spaces

Standard duality in L^2 spaces (II)

Consider the following optimization problem
(a variation on a linear example given by R. Wets)

$$\begin{aligned}
\inf_{u_0,\mathbf{u}_1} \quad & u_0^2 + \mathbb{E}\big((\mathbf{u}_1+\alpha)^2\big) \\
\text{s.t.} \quad & u_0 \geq a \\
& \mathbf{u}_1 \geq 0 \\
& u_0 - \mathbf{u}_1 \geq \mathbf{w} && \text{to be dualized}
\end{aligned}$$

where $\mathbf{w}$ is a random variable uniform on $[1,2]$

For $a < 2$, we exhibit a maximizing sequence of multipliers for the dual problem that does not converge in L^2.
(We are in the so-called *non relatively complete recourse* case, that is, the case where the constraints on $\mathbf{u}_1$ induce a stronger constraint on u_0)

The optimal multiplier is not in L^2, but in $(\mathrm{L}^\infty)^\star$

Constraint qualification in $\left(\mathrm{L}^\infty, \mathrm{L}^1\right)$

From now on, we assume that

$$\mathcal{U} = \mathrm{L}^\infty\left(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n\right)$$
$$\mathcal{V} = \mathrm{L}^\infty\left(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m\right)$$
$$C = \{0\}$$

where the σ-algebra $\mathcal{F}$ is not finite modulo $\mathbb{P}$

We consider the pairing $\left(\mathrm{L}^\infty, \mathrm{L}^1\right)$ with the following topologies:

- $\sigma\left(\mathrm{L}^\infty, \mathrm{L}^1\right)$: weak* topology on L^∞ (coarsest topology such that all the L^1-linear forms are continuous),
- $\tau\left(\mathrm{L}^\infty, \mathrm{L}^1\right)$: Mackey-topology on L^∞ (finest topology such that the continuous linear forms are only the L^1-linear forms)

Weak* closedness of linear subspaces of L^∞

Proposition 2

Let $\Theta : \mathrm{L}^\infty(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n) \to \mathrm{L}^\infty(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m)$ *be a linear operator, and assume that there exists a linear operator* $\Theta^\dagger : \mathrm{L}^1(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m) \to \mathrm{L}^1(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n)$ *such that:*

$$\langle \mathbf{v} , \Theta(\mathbf{u}) \rangle = \langle \Theta^\dagger(\mathbf{v}) , \mathbf{u} \rangle , \ \forall \mathbf{u}, \ \forall \mathbf{v}$$

Then the linear operator Θ *is weak** *continuous*

Applications

- $\Theta(\mathbf{u}) = \mathbf{u} - \mathbb{E}(\mathbf{u} \mid \mathcal{B})$: non-anticipativity constraints
- $\Theta(\mathbf{u}) = A\mathbf{u}$ with $A \in \mathcal{M}_{m,n}(\mathbb{R})$: finite number of constraints

A duality theorem

$$(\mathcal{P}) \qquad \min_{\mathbf{u}\in\mathcal{U}} J(\mathbf{u}) \quad \text{s.t.} \quad \Theta(\mathbf{u}) = 0$$

with $J(\mathbf{u}) = \mathbb{E}\big(j(\mathbf{u},\mathbf{w})\big)$

Theorem 1

Assume that j is a convex normal integrand, that J is continuous in the Mackey topology at some point $\mathbf{u}_0$ such that $\Theta(\mathbf{u}_0) = 0$, and that Θ is weak* continuous on $\mathrm{L}^\infty\big(\Omega,\mathcal{F},\mathbb{P};\mathbb{R}^n\big)$
Then, $\mathbf{u}^\star \in \mathcal{U}$ is an optimal solution of Problem $\big(\mathcal{P}\big)$ if and only if there exists $\lambda^\star \in \mathrm{L}^1\big(\Omega,\mathcal{F},\mathbb{P};\mathbb{R}^m\big)$ such that

- $\mathbf{u}^\star \in \arg\min_{\mathbf{u}\in\mathcal{U}} \mathbb{E}\Big(j(\mathbf{u},\mathbf{w}) + \lambda^\star \cdot \Theta(\mathbf{u})\Big)$
- $\Theta(\mathbf{u}^\star) = 0$

Extension of a result given by R. Wets for non-anticipativity constraints

Uzawa algorithm

$$(\mathcal{P}) \qquad \min_{\mathbf{u}\in\mathcal{U}} J(\mathbf{u}) \quad \text{s.t.} \quad \Theta(\mathbf{u}) = 0$$

with $J(\mathbf{u}) = \mathbb{E}\big(j(\mathbf{u},\mathbf{w})\big)$

The standard Uzawa algorithm makes sense

$$\mathbf{u}^{(k+1)} \in \operatorname*{arg\,min}_{\mathbf{u}\in\mathcal{U}^{\mathrm{ad}}} J(\mathbf{u}) + \left\langle \lambda^{(k)} , \Theta(\mathbf{u}) \right\rangle$$

$$\lambda^{(k+1)} = \underbrace{\lambda^{(k)}}_{\text{dual}} + \rho \underbrace{\Theta\big(\mathbf{u}^{(k+1)}\big)}_{\text{primal}}$$

Note that all the multipliers $\lambda^{(k)}$ *belong to* $\mathrm{L}^\infty\big(\Omega,\mathcal{F},\mathbb{P};\mathbb{R}^m\big)$, *as soon as the initial multiplier* $\lambda^{(0)} \in \mathrm{L}^\infty\big(\Omega,\mathcal{F},\mathbb{P};\mathbb{R}^m\big)$

Convergence result

Theorem 2

Assume that

1. $J : \mathcal{U} \to \overline{\mathbb{R}}$ *is proper, weak* l.s.c., differentiable and* a*-convex*
2. $\Theta : \mathcal{U} \to \mathcal{V}$ *is affine, weak* continuous and* κ*-Lipschitz*
3. $\mathcal{U}^{\mathrm{ad}}$ *is weak* closed and convex,*
4. *an admissible* $\mathbf{u}_0 \in \mathrm{dom}\, J \cap \Theta^{-1}(0) \cap \mathcal{U}^{\mathrm{ad}}$ *exists*
5. *an optimal* L^1*-multiplier to the constraint* $\Theta\big(\mathbf{u}\big) = 0$ *exists*
6. *the step* ρ *is such that* $0 < \rho < \frac{2a}{\kappa}$

Then, there exists a subsequence $\big\{\mathbf{u}^{(n_k)}\big\}_{k \in \mathbb{N}}$ *of the sequence generated by the Uzawa algorithm converging in* L^∞ *toward the optimal solution* $\mathbf{u}^\star$ *of the primal problem*

Discussion

- The result is not as good as expected (convergence of a subsequence)
- Improvements and extensions (inequality constraint) needed
- The Mackey-continuity assumption forbids
  the use of extended functions
  - In order to deal with almost sure bound constraints, we can turn towards the work of T. Rockafellar and R. Wets
  - In a series of 4 papers (stochastic convex programming), they have detailed the duality theory on two-stage and multistage problems, with the focus on non-anticipativity constraints
  - These papers require
    - a strict feasability assumption
    - a relatively complete recourse assumption

Relaxed problems

Following the interpretation of DADP in terms of a relaxation of the original problem, and given a sequence $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ of subfields of the σ-field $\mathcal{F}$, we replace the abstract problem

$$(\mathcal{P}) \qquad \min_{\mathbf{u}\in\mathcal{U}} J(\mathbf{u}) \quad \text{s.t.} \quad \Theta(\mathbf{u}) = 0$$

by the sequence of approximated problems:

$$(\mathcal{P}_n) \qquad \min_{\mathbf{u}\in\mathcal{U}} J(\mathbf{u}) \quad \text{s.t.} \quad \mathbb{E}\big(\Theta(\mathbf{u}) \bigm| \mathcal{F}_n\big) = 0$$

We assume the Kudo convergence of $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ toward $\mathcal{F}$:

$$\mathcal{F}_n \longrightarrow \mathcal{F} \iff \mathbb{E}\big(\mathbf{z} \bigm| \mathcal{F}_n\big) \xrightarrow{\mathrm{L}^1} \mathbb{E}\big(\mathbf{z} \bigm| \mathcal{F}\big) \,, \ \forall \mathbf{z} \in \mathrm{L}^1(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R})$$

Convergence result

Theorem 3

Assume that

- $\mathcal{U}$ *is a topological space*
- $\mathcal{V} = \mathrm{L}^p(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^m)$ *with* $p \in [1, +\infty)$
- J *and* Θ *are continuous operators*
- $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ *Kudo converges toward* $\mathcal{F}$

Then the sequence $\{\widetilde{J}_n\}_{n\in\mathbb{N}}$ *epi-converges toward* $\widetilde{J}$*, with*

$$\widetilde{J}_n = \begin{cases} J(\mathbf{u}) & \text{if } \mathbf{u} \text{ satisfies the constraints of } (\mathcal{P}_n) \\ +\infty & \text{otherwise} \end{cases}$$

Summing up theoretical questions

- Conditions for the existence of an L^1 multiplier
- Convergence of the algorithm (fixed information process)
- Consistency result (family of information process)

Discussing DADP

- DADP (Dual Approximate Dynamic Programming) is a method to design stochastic price signals allowing decentralized agents to act as a team
- Hence, DADP is especially adapted to tackle large-scale stochastic optimal control problems, such as those found in energy management
- A host of theoretical and practical questions remains open
- We would like to test DADP on "network models" (smart grids) extending the works already made on "flower models" (unit commitment problem) and on "chained models" (hydraulic valley management)

Let us move to broader stochastic optimization challenges

- Stochastic optimization requires to make risk attitudes explicit
  - robust, worst case, risk measures, in probability, almost surely, etc.
- Stochastic dynamic optimization requires to make online information explicit
  - State-based functional approach
  - Scenario-based measurability approach

Numerical walls

- in dynamic programming, the bottleneck is the dimension of the state
- in stochastic programming, the bottleneck is the number of stages

Here is our research agenda for stochastic decomposition

- Combining different decomposition methods
  - time: dynamic programming
  - scenario: progressive hedging
  - space: dual approximate dynamic programming
- Designing risk criterion compatible with decomposition (time-consistent dynamic risk measures)
- Mixing decomposition with analytical properties (convexity, linearity) on costs, constraints and dynamics functions