Decomposition Methods in Multistage Stochastic Optimization

Michel DE LARA
CERMICS, École des Ponts ParisTech, France
and
Pierre Carpentier, ENSTA ParisTech, France
Jean-Philippe Chancelier, École des Ponts ParisTech, France
Vincent Leclère, École des Ponts ParisTech, France

CERMICS, France

October 28, 2014

Outline of the presentation

1. Framing stochastic optimization problems
   - Working out a toy example
   - Expliciting risk attitudes
   - Handling online information
2. Solving stochastic optimization problems by decomposition methods
   - A bird's eye view of decomposition methods
   - Spatial decomposition methods in the deterministic case
   - The stochastic case raises specific obstacles
3. Summary and research agenda

Let us work out a toy example of economic dispatch as a cost-minimization problem under supply-demand balance

- Production: consider two energy production units
  - a "cheap" limited one with which we can produce quantity q_0, with $0 \leq q_0 \leq q_0^\sharp$, at cost $c_0 q_0$
  - an "expensive" unlimited one with which we can produce quantity q_1, with $0 \leq q_1$, at cost $c_1 q_1$, with $c_1 > c_0$
- Consumption: the demand is $D \geq 0$
- Balance: ensuring at least the demand

$$D \leq q_0 + q_1$$

- Optimization: total costs minimization

$$\min_{q_0, q_1} \underbrace{c_0 q_0 + c_1 q_1}_{\text{total costs}}$$

When the demand D is deterministic, the optimization problem is well posed

- The deterministic demand D is a single number, and we minimize

$$\min_{q_0, q_1} c_0 q_0 + c_1 q_1$$

$$\text{under the constraints} \quad \begin{array}{rl} 0 & \leq q_0 \leq q_0^\sharp \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \end{array}$$

- The solution is $q_0^\star = \min\{q_0^\sharp, D\}$, $\quad q_1^\star = [D - q_0^\sharp]_+$, that is,
  - if the demand D is below the capacity $q_0^\sharp$ of the "cheap" energy source

$$D \leq q_0^\sharp \Rightarrow q_0^\star = D \,, \quad q_1^\star = 0$$

  - if the demand D is above the capacity $q_0^\sharp$ of the "cheap" energy source, you have to have recourse to the "expensive" source

$$D > q_0^\sharp \Rightarrow q_0^\star = q_0^\sharp \,, \quad q_1^\star = D - q_0^\sharp$$

- Now, what happens when the demand D is no longer deterministic?

What happens if we replace the uncertain value D of the demand by its mean $\overline{D}$ in the deterministic solution?

- If we suppose that the demand D is a random variable $D : \Omega \to \mathbb{R}_+$, with mathematical expectation $\mathbb{E}(D) = \overline{D}$
- and that we propose the "deterministic solution"

$$q_0^{(\overline{D})} = \min\{q_0^\sharp, \overline{D}\} \ , \ \ q_1^{(\overline{D})} = [\overline{D} - q_0^\sharp]_+$$

- we cannot assure the inequality

$$\underbrace{D(\omega)}_{\text{uncertain}} \leq \underbrace{q_0 + q_1}_{\text{deterministic}} \ , \ \forall \omega \in \Omega$$

because $\sup_{\omega \in \Omega} D(\omega) > \overline{D} = q_0^{(\overline{D})} + q_1^{(\overline{D})}$

- Are there better solutions among the deterministic ones?

When the demand D is bounded above, the robust optimization problem has a solution

- In the robust optimization problem, we minimize

$$\min_{q_0, q_1} c_0 q_0 + c_1 q_1$$

$$\text{under the constraints} \qquad \begin{array}{rll} 0 & \leq q_0 \leq q_0^\sharp & \\ 0 & \leq q_1 & \\ D(\omega) & \leq q_0 + q_1 & \forall \omega \in \Omega \end{array}$$

- When $D^\sharp = \sup_{\omega \in \Omega} D(\omega) < +\infty$, the solution is $q_0^\star = \min\{q_0^\sharp, D^\sharp\}$, $\quad q_1^\star = [D^\sharp - q_0^\sharp]_+$
- Now, the total cost $c_0 q_0^\star + c_1 q_1^\star$ is an increasing function of the upper bound $D^\sharp$ of the demand
- Is it not too costly to optimize under the worst-case situation?

What happens if we solve the problem demand value by demand value?

- If we solve the problem for each possible value $d = D(\omega)$ of the random variable D, when $\omega \in \Omega$, we obtain a collection of "solutions"

$$q_0^{(d)} = \min\{q_0^\sharp, d\}\,, \quad q_1^{(d)} = [d - q_0^\sharp]_+$$

- Now, we face an informational issue
  - if the demand D is observed before selecting the quantities q_0 and q_1, this collection of "solutions" is optimal in many understandings
  - whereas, on the contrary, how can we glue together those "solutions" to cook up quantities q_0 or q_1 that do not depend upon the unknown quantities d?
- When the demand D is not observed, we do not know :-(
  and this is a big issue with that so-called scenarios method
- Therefore, we can remain with a feasability issue

To overcome the above difficulties, we turn to stochastic optimization

- We suppose that the demand D is a random variable, and minimize

$$\min_{q_0, q_1} \mathbb{E}[c_0 q_0 + c_1 q_1]$$

$$\text{under the constraints} \quad \begin{array}{ll} 0 & \leq q_0 \leq q_0^{\sharp} \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \\ q_1 & \text{depends upon } D \end{array}$$

and we emphasize two issues, new with respect to the deterministic case

- expliciting online information issue:
  the decision q_1 depends upon the random variable D
- expliciting risk attitudes:
  we aggregate the total costs with respect to all possible values
  by taking the expectation $\mathbb{E}[c_0 q_0 + c_1 q_1]$

Turning to stochastic optimization forces one to specify online information

- We suppose that the demand D is a random variable, and minimize

$$\min_{q_0, q_1} \mathbb{E}[c_0 q_0 + c_1 q_1]$$

$$\text{under the constraints} \quad \begin{array}{rl} 0 & \leq q_0 \leq q_0^{\sharp} \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \\ q_1 & \text{depends upon } D \end{array}$$

- specifying that the decision q_1 depends upon the random variable D, whereas q_0 does not, forces to consider two stages and a so-called non-anticipativity constraint (more on that later)
  - first stage: q_0 does not depend upon the random variable D
  - second stage: q_1 depends upon the random variable D

Turning to stochastic optimization forces one to specify risk attitudes

▷ We suppose that the demand D is a random variable, and minimize

$$\min_{q_0, q_1} \mathbb{E}[c_0 q_0 + c_1 q_1]$$

$$\text{under the constraints} \quad \begin{array}{ll} 0 & \leq q_0 \leq q_0^{\sharp} \\ 0 & \leq q_1 \\ D & \leq q_0 + q_1 \\ q_1 & \text{depends upon } D \end{array}$$

▷ Now that q_1 depends upon the random variable D, it is also a random variable, and so is the total cost $c_0 q_0 + c_1 q_1$; therefore, we have to aggregate the total costs with respect to all possible values, and we chose to do it by taking the expectation $\mathbb{E}[c_0 q_0 + c_1 q_1]$

In the uncertain framework, two additional questions must be answered with respect to the deterministic case

Question (expliciting risk attitudes)

How are the uncertainties taken into account
in the payoff criterion and in the constraints?

Question (expliciting available online information)

Upon which online information are decisions made?

Outline of the presentation

The output of a stochastic optimization problem is a random variable. How can we rank random variables?

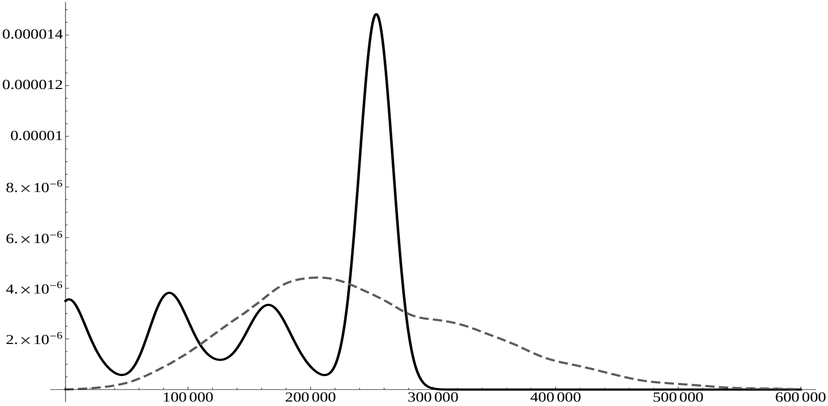

How are the uncertainties taken into account in the payoff criterion and in the constraints?

In a probabilistic setting, where uncertainties are random variables, a classical answer is

- to take the mathematical expectation of the payoff (risk-neutral approach)

$$\mathbb{E}(\text{payoff})$$

- and to satisfy all (physical) constrainsts almost surely that is, practically, for all possible issues of the uncertainties (robust approach)

$$\mathbb{P}(\text{constrainsts}) = 1$$

But there are many other ways to handle risk: robust, worst case, risk measures, in probability, almost surely, by penalization, etc.

Outline of the presentation

1. Framing stochastic optimization problems
   - Working out a toy example
   - Expliciting risk attitudes
   - Handling online information
2. Solving stochastic optimization problems by decomposition methods
   - A bird's eye view of decomposition methods
   - Spatial decomposition methods in the deterministic case
   - The stochastic case raises specific obstacles
3. Summary and research agenda

Upon which online information are decisions made?

We navigate between two stumbling blocks: rigidity and wizardry

- On the one hand, it is suboptimal to restrict oneself, as in the deterministic case, to open-loop controls depending only upon time, thereby ignoring the available information at the moment of making a decision
- On the other hand, it is impossible to suppose that we know in advance what will happen for all times: clairvoyance is impossible as well as look-ahead solutions

The in-between is $\boxed{\text{non-anticipativity constraint}}$

There are two ways to express the non-anticipativity constraint

Denote the uncertainties at time t by w_t, and the control by u_t

- ▷ Functional approach
  The control u_t may be looked after under the form

$$u_t = \phi_t\big(\underbrace{w_{t_0}, \ldots, w_{t-1}}_{\text{past}}\big)$$

  where ϕ_t is a function, called policy, strategy or decision rule

- ▷ Algebraic approach
  When uncertainties are considered as random variables (measurable mappings), the above formula for u_t expresses the measurability of the control variable u_t with respect to the past uncertainties, also written as

$$\sigma(u_t) \subset \sigma\big(\underbrace{w_{t_0}, \ldots, w_{t-1}}_{\text{past}}\big)$$

What is a solution at time t?

- In deterministic control, the solution u_t at time t is a single number
- In stochastic control, the solution u_t at time t is a random variable expressed
  - either as $u_t = \phi_t\big(w_{t_0}, \dots, w_{t-1}\big)$, where $\phi_t : \mathbb{W}^{t-t_0} \to \mathbb{R}$
  - or as $u_t : \Omega \to \mathbb{R}$ with measurability constraint $\sigma(u_t) \subset \sigma\big(w_{t_0}, \dots, w_{t-1}\big)$
- Now, as time t goes on, the domain of the function ϕ_t expands, and so do the conditions $\sigma(u_t) \subset \sigma\big(w_{t_0}, \dots, w_{t-1}\big)$
- Therefore, for numerical reasons, the information $\big(w_{t_0}, \dots, w_{t-1}\big)$ has to be compressed or approximated

There are two classical ways to compress information

▷ State-based functional approach

In the special case of the Markovian framework with $\big(w_{t_0}, \dots, w_T\big)$ white noise, there is no loss of optimality to look for solutions as

$$u_t = \psi_t \underbrace{\left(x_t\right)}_{\text{state}} \quad \text{where} \quad \underbrace{x_t \in \mathbb{X}}_{\text{fixed space}} \; , \quad \underbrace{x_{t+1} = f_t(x_t, u_t, w_t)}_{\text{dynamical equation}}$$

▷ Scenario-based measurability approach

- Scenarios are approximated by a finite family $\big(w^s_{t_0}, \dots, w^s_T\big)$, $s \in S$
- Solutions $q^s_{i,t}$ are indexed by $s \in S$ with the constraint that if two scenarios coincide up to time t, so must do the controls at time t

$$\big(w^s_{t_0}, \dots, w^s_{t-1}\big) = \big(w^{s'}_{t_0}, \dots, w^{s'}_{t-1}\big) \Rightarrow q^s_{i,t} = q^{s'}_{i,t}$$

- In the case of the scenario tree approach, the scenarios $\big(w^s_{t_0}, \dots, w^s_T\big)$, $s \in S$, are organized in a tree, and controls $q^n_{i,t}$ are indexed by nodes n on the tree

More on what is a solution at time t State-based approach $u_t = \psi_t(x_t)$

- The mapping ψ_t can be computed in advance (that is, at initial time t_0) and evaluated at time t on the available online information at that time t
  - either exactly (for example, by dynamic programming)
  - or approximately (for example, among linear decision rules) because the computational burden of finding *any* function is heavy
- The value $u_t = \psi_t(x_t)$ can be computed at time t
  - either exactly by solving a proper optimization problem, which raises issues of dynamic consistency
  - or approximately (for example, by assuming that controls from time t on are open-loop)

More on what is a solution at time t
Scenario-based approach

- An optimal "solution" can be computed scenario by scenario, with the problem that we obtain solutions such that

$$\left(w^s_{t_0}, \dots, w^s_{t-1}\right) = \left(w^{s'}_{t_0}, \dots, w^{s'}_{t-1}\right) \text{ and } u^s_t \neq u^{s'}_t$$

- Optimal solutions can be computed scenario by scenario and then merged (for example, by progressive hedging) to be forced to satisfy

$$\left(w^s_{t_0}, \dots, w^s_{t-1}\right) = \left(w^{s'}_{t_0}, \dots, w^{s'}_{t-1}\right) \Rightarrow u^s_t = u^{s'}_t$$

- The value u_t can be computed at time t depending on $\left(w^s_{t_0}, \dots, w^s_{t-1}\right)$
  - either exactly by solving a proper optimization problem, which raises issues of dynamic consistency
  - or approximately (for example, by a sequence of two-stages problems)

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.

Couplings for stochastic problems

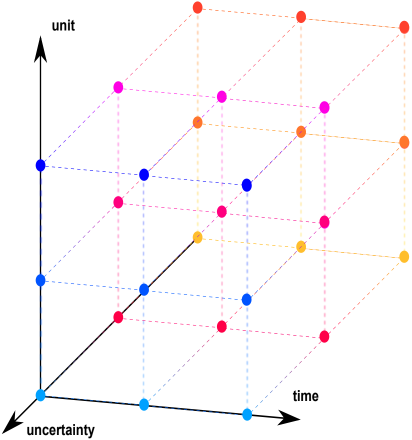

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

Couplings for stochastic problems: in time

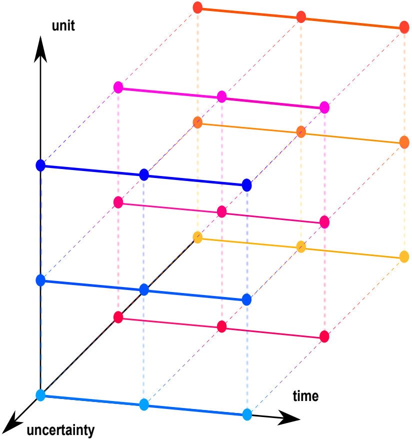

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

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

Couplings for stochastic problems: in uncertainty

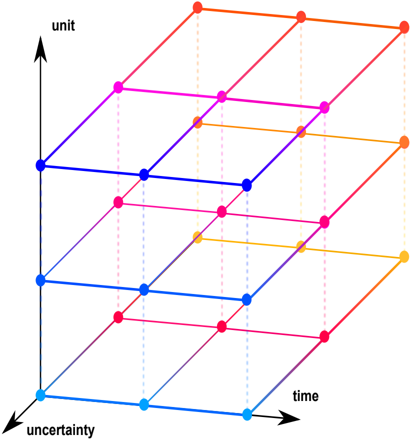

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

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

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \dots, \mathbf{w}_t\right)$$

Couplings for stochastic problems: in space

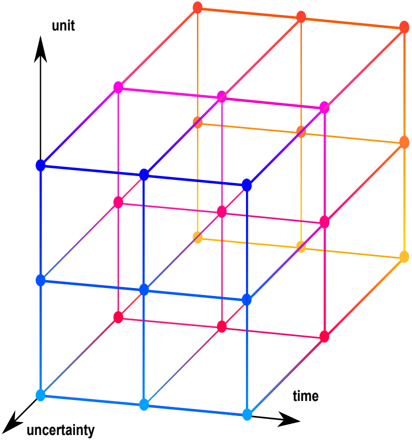

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

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

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \dots, \mathbf{w}_t\right)$$

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

Can we decouple stochastic problems?

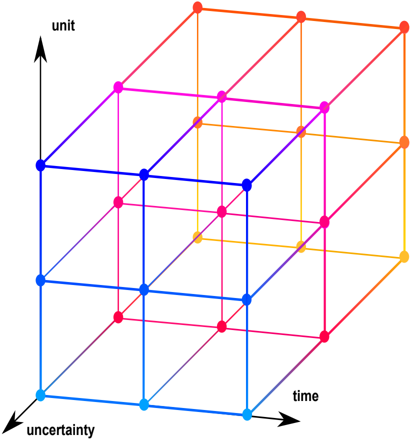

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

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

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \ldots, \mathbf{w}_t\right)$$

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

Decompositions for stochastic problems: in time

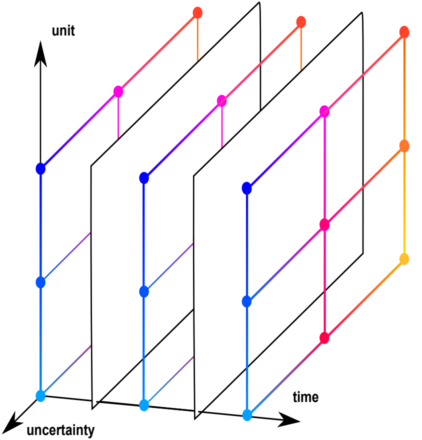

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

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

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \ldots, \mathbf{w}_t\right)$$

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

Dynamic Programming

Bellman (56)

Decompositions for stochastic problems: in uncertainty

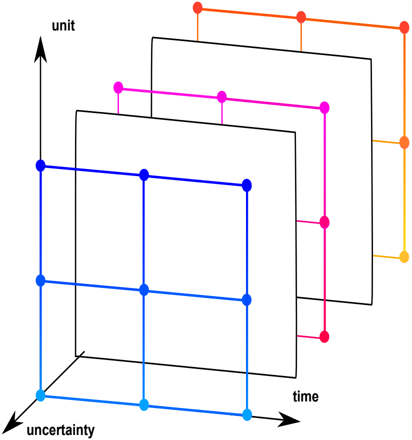

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_\omega L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

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

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \,\middle|\, \mathbf{w}_1, \ldots, \mathbf{w}_t\right)$$

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

Progressive Hedging

Rockafellar - Wets (91)

Decompositions for stochastic problems: in space

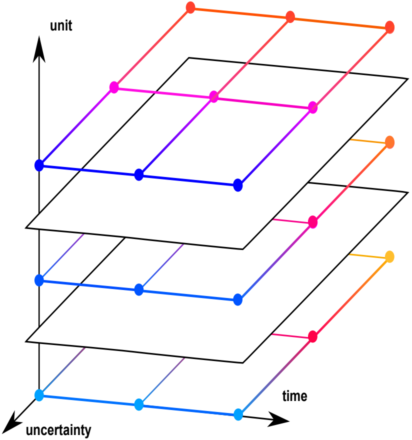

$$\min \sum_{\omega} \sum_{i} \sum_{t} \pi_{\omega} L_t^i(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_{t+1})$$

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

$$\mathbf{u}_t^i = \mathbb{E}\left(\mathbf{u}_t^i \;\middle|\; \mathbf{w}_1, \ldots, \mathbf{w}_t\right)$$

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

Dual Approximate
Dynamic Programming

Decomposition and coordination

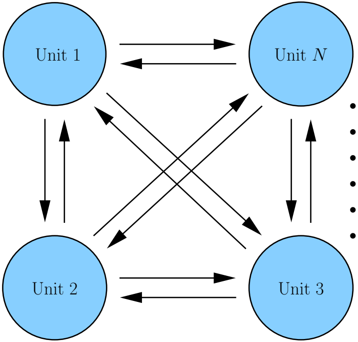

Interconnected units

- 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”

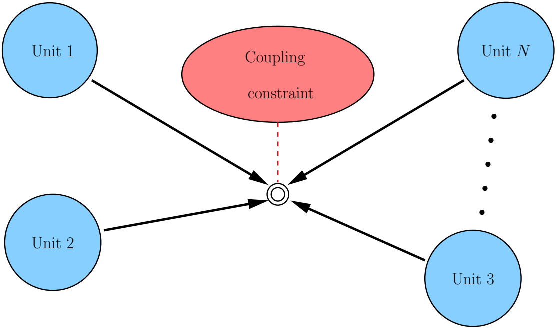

$$\min_{u} \sum_{i=1}^{N} J_i(u_i)$$

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

Unit Commitment Problem

Intuition of spatial decomposition

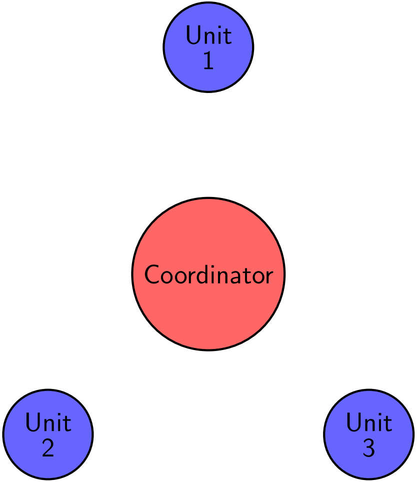

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition

Intuition of spatial decomposition

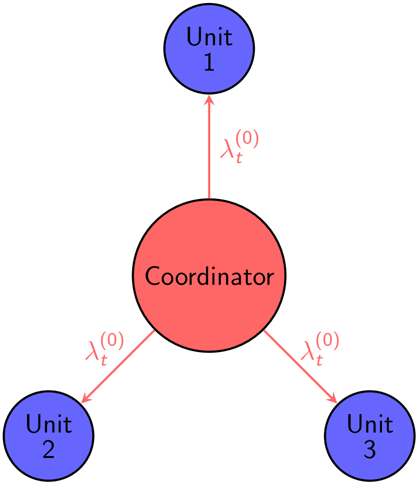

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t

Intuition of spatial decomposition

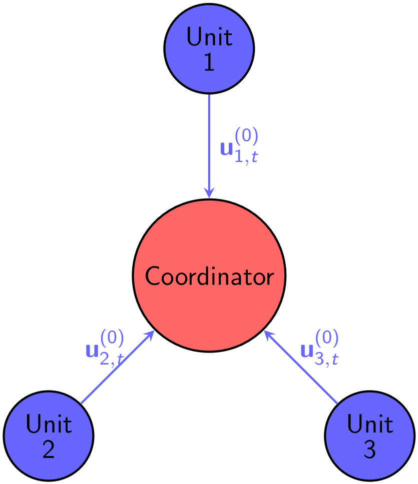

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}_t^{(i)}$

Intuition of spatial decomposition

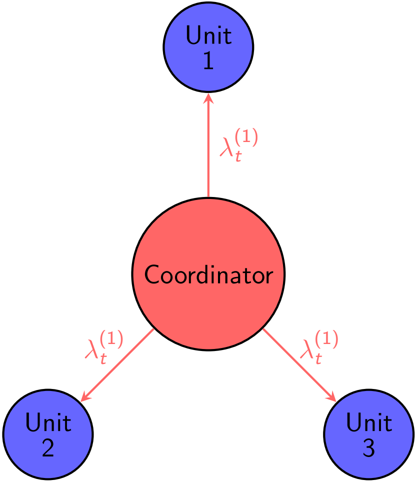

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}_t^{(i)}$
  - the coordinator compares total production and demand, and then updates the price

Intuition of spatial decomposition

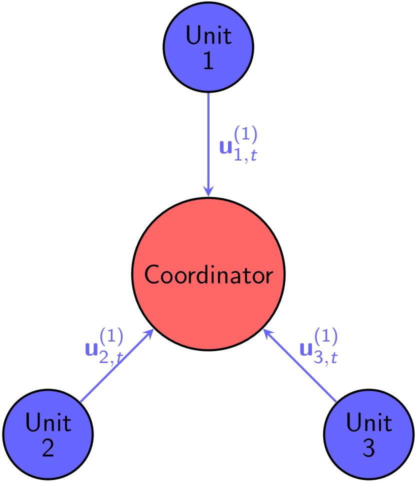

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}_t^{(i)}$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Intuition of spatial decomposition

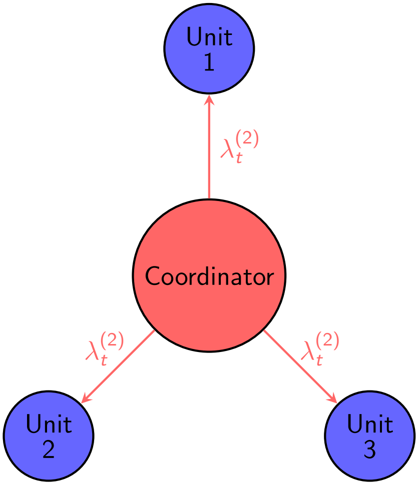

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}_t^{(i)}$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Intuition of spatial decomposition

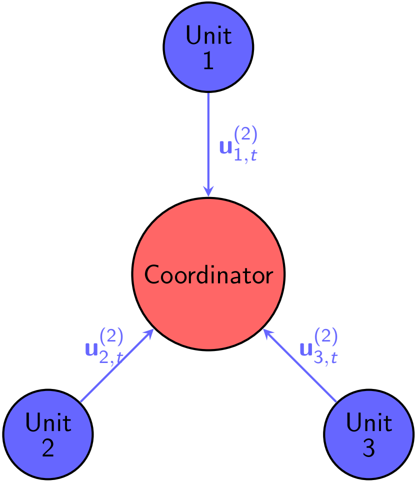

- Purpose: satisfy a demand with N production units, at minimal cost
- Price decomposition
  - the coordinator sets a price λ_t
  - the units send their production $\mathbf{u}^{(i)}_t$
  - the coordinator compares total production and demand, and then updates the price
  - and so on...

Price decomposition relies on dualization

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

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

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

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

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

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

Price decomposition relies on dualization

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

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

$$\max_{\lambda \in \mathcal{V}} \min_{u \in \mathcal{U}} \sum_{i=1}^{N} \Big(J_i(u_i) + \left\langle \lambda \,, \theta_i(u_i) \right\rangle \Big) - \left\langle \lambda \,, \theta \right\rangle$$

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

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

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

Remarks on decomposition methods

- The theory is available for infinite dimensional Hilbert spaces, and thus applies in the stochastic framework, that is, when $\mathcal{U}$ is a space 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

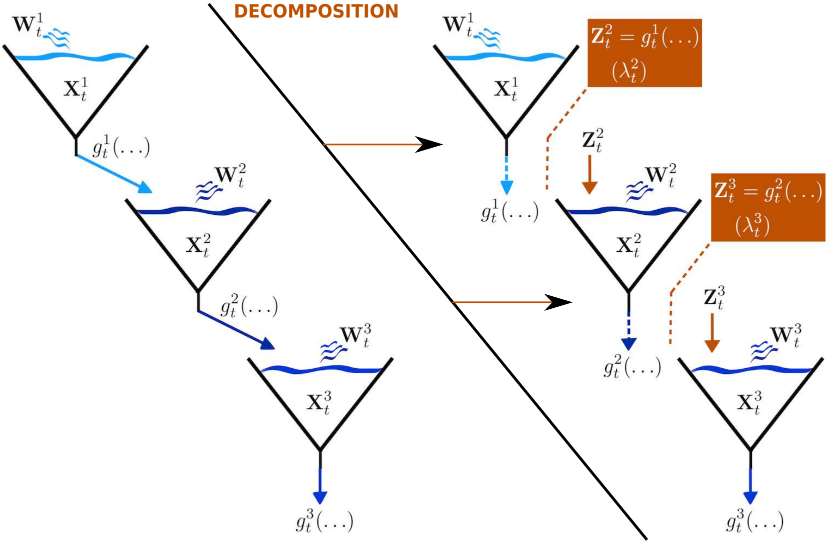

Stochastic optimal control (SOC) problem formulation

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

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

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

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

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

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

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

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

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
- 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 $\left(\mathbf{x}_t^1, \dots, \mathbf{x}_t^N\right)$

$$\mathbf{u}_t^i = \gamma_t^i\left(\mathbf{x}_t^1, \dots, \mathbf{x}_t^N\right)$$

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

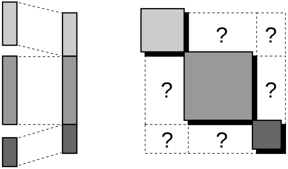

- ▷ how to decompose a feedback γ_t w.r.t. its domain $\mathbb{X}_t$ rather than its range $\mathbb{U}_t$?

  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 $\mathbf{\Lambda}_t^{(k)}$'s 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
- ▷ Dynamic Programming cannot be applied in a straightforward manner!

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

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 $\mathbf{\Lambda}_t^{(k)}$'s 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
- ▷ Dynamic Programming cannot be applied in a straightforward manner!

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

Framing, framing, framing

- Stochastic optimization requires to make risk attitudes explicit
- Stochastic dynamic requires to make online information explicit
- These explicitations raise a bunch of issues, because of the
  - many ways to represent risk (criterion, constraints)
  - many information structures
  - tremendous numerical obstacles to overcome

Handling risk and online information

- Risk
  - robust, worst case, risk measures, in probability, almost surely, by penalization
- Online information
  - State-based functional approach
  - Scenario-based measurability approach

Numerical walls

- in dynamic programming, the bottleneck is the dimension of the state (no more than 3)
- in stochastic programming, the bottleneck is the number of stages (no more than 2)

Decomposing and mixing different decompositions

- Decomposition with respect to
  - time: dynamic programming
  - scenario: progressive hedging
  - space: dual approximate dynamic programming
- Research agenda
  - designing risk criterion compatible with decomposition
  - combining different decomposition methods
  - mixing with analytical properties (convexity, linearity)