A fair comparison of two stochastic optimization algorithms

Benchmarking MPC vs SDDP

April 10, 2017

Why using stochastic optimization?

We aim to tackle uncertainties in Energy Management System.

Problem: we do not know in advance the uncertainties, common in the management of energetical systems:

- Electrical demands
- Hot water demands
- Outdoor temperature
- Wind's speed
- Solar irradiation
- etc.

Introducing the problem

Here, we focus on the management of a domestic microgrid

Sensitivity analysis w.r.t two uncertainties:

- Electrical demands
- Solar irradiation

Introducing the problem

Here, we focus on the management of a domestic microgrid

Sensitivity analysis w.r.t two uncertainties:

- Electrical demands
- Solar irradiation

Introducing the problem

Here, we focus on the management of a domestic microgrid

Sensitivity analysis w.r.t two uncertainties:

- Electrical demands
- Solar irradiation

We compare two classes of algorithm:

The Mainstream: Model Predictive Control (MPC)
(use forecasts to predict the future uncertainties)

The Challenger: Stochastic Dual Dynamic Programming (SDDP)
(model uncertainties with discrete probability laws)

Outline

A brief recall of the single house problem

- Physical modelling
- Optimization problem
- Resolution Methods

Handling solar irradiation

- Academic modeling
- Realistic modeling

Framing the optimization problem

We aim to

- Minimize electrical's bill
- Maintain a comfortable temperature inside the house

To achieve these goals, we can

- store electricity in battery;
- store heat in hot water tank.

We control the stocks every 15mn over one day.

We formulate a multistage stochastic programming problem

Microgrid's description

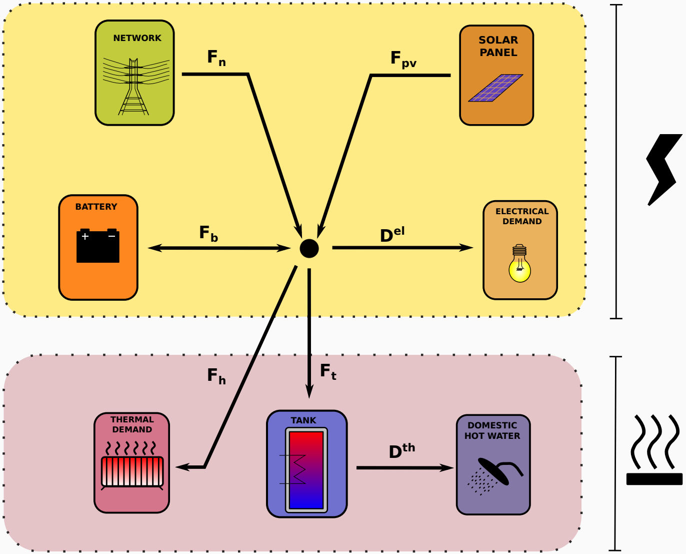

Outline

We introduce states, controls and noises

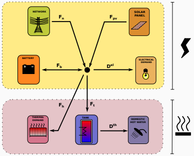

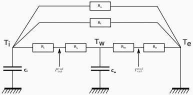

- **Stock variables** $X_t = \left(\mathbf{B}_t, \mathbf{H}_t, \theta^i_t, \theta^w_t\right)$
  - $\mathbf{B}_t$, battery level (kWh)
  - $\mathbf{H}_t$, hot water storage (kWh)
  - θ^i_t, inner temperature ($^\circ C$)
  - θ^w_t, wall's temperature ($^\circ C$)
- **Control variables** $U_t = \left(\mathbf{F}^+_{\mathbf{B},t}, \mathbf{F}^-_{\mathbf{B},t}, \mathbf{F}_{A,t}, \mathbf{F}_{\mathbf{H},t}\right)$
  - $\mathbf{F}^+_{\mathbf{B},t}$, energy stored in the battery
  - $\mathbf{F}^-_{\mathbf{B},t}$, energy taken from the battery
  - $\mathbf{F}_{A,t}$, energy used to heat the hot water tank
  - $\mathbf{F}_{\mathbf{H},t}$, thermal heating
- **Uncertainties** $W_t = \left(\mathbf{D}^E_t, \mathbf{D}^{DHW}_t, \mathbf{\Phi}^s_t\right)$
  - $\mathbf{D}^E_t$, electrical demand (kW)
  - $\mathbf{D}^{DHW}_t$, domestic hot water demand (kW)
  - $\mathbf{\Phi}^s_t$, external radiations (kW)

Discrete time state equations

So we have the four state equations (all linear):

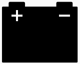

$$\mathbf{B}_{t+1} = \alpha_{\mathbf{B}} \mathbf{B}_t + \Delta T \Big(\rho_c \mathbf{F}^{+}_{\mathbf{B},t} - \frac{1}{\rho_d} \mathbf{F}^{-}_{\mathbf{B},t} \Big)$$

$$\mathbf{H}_{t+1} = \alpha_{\mathbf{H}} \mathbf{H}_t + \Delta T \big[\mathbf{F}_{A,t} - \mathbf{D}^{DHW}_t \big]$$

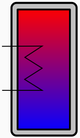

$$\boldsymbol{\theta}^w_{t+1} = \boldsymbol{\theta}^w_t + \frac{\Delta T}{c_m} \left[\frac{\boldsymbol{\theta}^i_t - \boldsymbol{\theta}^w_t}{R_i + R_s} + \frac{\boldsymbol{\theta}^e_t - \boldsymbol{\theta}^w_t}{R_m + R_e} + \gamma \mathbf{F}_{\mathbf{H},t} + \frac{R_i}{R_i + R_s} P^{int}_t + \frac{R_e}{R_e + R_m} \boldsymbol{\Phi}^s_t \right]$$

$$\boldsymbol{\theta}^i_{t+1} = \boldsymbol{\theta}^i_t + \frac{\Delta T}{c_i} \left[\frac{\boldsymbol{\theta}^w_t - \boldsymbol{\theta}^i_t}{R_i + R_s} + \frac{\boldsymbol{\theta}^e_t - \boldsymbol{\theta}^i_t}{R_v} + \frac{\boldsymbol{\theta}^e_t - \boldsymbol{\theta}^i_t}{R_f} + (1-\gamma) \mathbf{F}_{\mathbf{H},t} + \frac{R_s}{R_i + R_s} P^{int}_t \right]$$

which will be denoted:

$$\boxed{X_{t+1} = f_t(X_t, U_t, W_{t+1})}$$

Outline

Prices and temperature setpoints vary along time

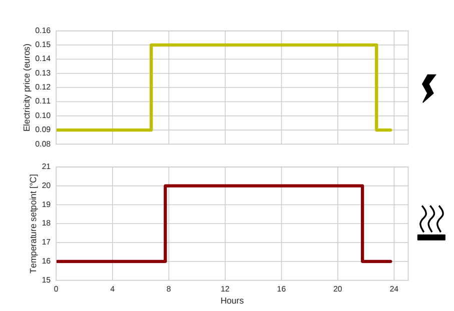

- $T_f = 24\text{h}$, $\Delta T = 15\text{mn}$
- Electricity peak and off-peak hours
- $\pi_t^E = 1.5$ euros/kWh (10x higher than usual)
- Temperature set-point $\bar{\theta}_t^i = 16^\circ C$ or $20^\circ C$

The costs we have to pay

- Cost to import electricity from the network

$$-\underbrace{b_t^E \max\{0, -\mathbf{F}_{NE,t+1}\}}_{\text{selling}} + \underbrace{\pi_t^E \max\{0, \mathbf{F}_{NE,t+1}\}}_{\text{buying}}$$

where we define the recourse variable (electricity balance):

$$\underbrace{\mathbf{F}_{NE,t+1}}_{Network} = \underbrace{D^E_{t+1}}_{Demand} + \underbrace{\mathbf{F}^+_{\mathbf{B},t} - \mathbf{F}^-_{\mathbf{B},t}}_{Battery} + \underbrace{\mathbf{F}_{\mathbf{H},t}}_{Heating} + \underbrace{\mathbf{F}_{A,t}}_{Tank} - \underbrace{\mathbf{F}_{pv,t}}_{Solar\ \ panel}$$

- Virtual Cost of thermal discomfort: $\kappa_{th}(\underbrace{\boldsymbol{\theta}^i_t - \bar{\boldsymbol{\theta}}^i_t}_{\text{deviation from setpoint}})$

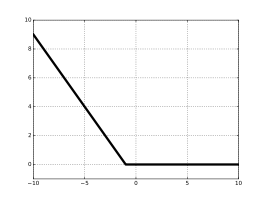

κ_{th}

Piecewise linear cost

Penalize temperature if

below given setpoint

Instantaneous and final costs for a single house

- The instantaneous convex costs are

$$L_t(X_t, U_t, W_{t+1}) = \underbrace{-b_t^E \max\{0, -\mathbf{F}_{NE,t+1}\}}_{buying} + \underbrace{\pi_t^E \max\{0, \mathbf{F}_{NE,t+1}\}}_{selling} + \underbrace{\kappa_{th}(\boldsymbol{\theta}_t^i - \bar{\boldsymbol{\theta}}_t^i)}_{discomfort}$$

- We add a final linear cost

$$K(X_T) = -\pi^{\mathbf{H}}\mathbf{H}_T - \pi^{\mathbf{B}}\mathbf{B}_T$$

to avoid empty stocks at the final horizon T

That gives the following stochastic optimization problem

$$\min_{X,U} \quad J(X,U) = \mathbb{E}\left[\sum_{t=0}^{T-1} \underbrace{L_t(X_t, U_t, W_{t+1})}_{instantaneous\ \ cost} + \underbrace{K(X_T)}_{final\ \ cost}\right]$$

$$\begin{aligned} s.t \quad X_{t+1} = f_t(X_t, U_t, W_{t+1}) & \quad \text{Dynamic} \\ X^\flat \leq X_t \leq X^\sharp & \\ U^\flat \leq U_t \leq U^\sharp & \\ X_0 = X_{ini} & \\ \sigma(U_t) \subset \sigma(W_1, \dots, W_t) & \quad \text{Non-anticipativity} \end{aligned}$$

That gives the following stochastic optimization problem

$$\min_{X,U} \quad J(X,U) = \mathbb{E}\left[\sum_{t=0}^{T-1} \underbrace{L_t(X_t, U_t, W_{t+1})}_{instantaneous\ \ cost} + \underbrace{K(X_T)}_{final\ \ cost}\right]$$

$$\begin{aligned} s.t \quad X_{t+1} &= f_t(X_t, U_t, W_{t+1}) \quad \text{Dynamic} \\ X^\flat \leq X_t &\leq X^\sharp \\ U^\flat \leq U_t &\leq U^\sharp \\ X_0 &= X_{ini} \\ \sigma(U_t) \subset \sigma(W_1, \dots, W_t) & \quad \text{Non-anticipativity} \end{aligned}$$

Because of the non-anticipativity constraint, we can not solve the optimization problem with standard methods (such as stochastic gradient)

Outline

MPC vs SDDP: uncertainties modelling

The two algorithms use optimization scenarios to model the uncertainties:

MPC

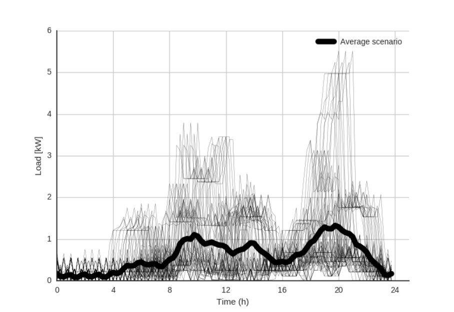

SDDP

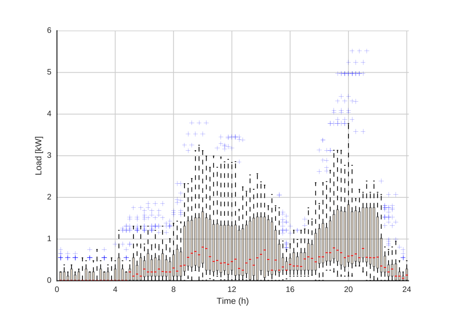

MPC vs SDDP: online resolution

At the beginning of time period $[\tau, \tau + 1]$, do

MPC

- Consider a **rolling horizon** $[\tau, \tau + H[$
- Consider a **deterministic scenario** of demands (forecast) $(\overline{W}_{\tau+1}, \dots, \overline{W}_{\tau+H})$
- Solve the **deterministic optimization** problem

$$\min_{X,U} \left[\sum_{t=\tau}^{\tau+H} L_t(X_t, U_t, \overline{W}_{t+1}) + K(X_{\tau+H}) \right]$$

$$\begin{aligned} s.t. \quad X_\cdot &= (X_\tau, \dots, X_{\tau+H}) \\ U_\cdot &= (U_\tau, \dots, U_{\tau+H-1}) \\ X_{t+1} &= f(X_t, U_t, \overline{W}_{t+1}) \\ X^\flat &\leq X_t \leq X^\sharp \\ U^\flat &\leq U_t \leq U^\sharp \end{aligned}$$

- Get optimal solution $(U_\tau^\#, \dots, U_{\tau+H}^\#)$ over horizon $H = 24h$
- Send only first control $U_\tau^\#$ to assessor, and iterate at time $\tau + 1$

SDDP

- We consider the approximated value functions $\left(\widetilde{V}_t\right)_0^T$

$$\underbrace{\widetilde{V}_t}_{\text{Piecewise affine functions}} \leq V_t$$

- Solve the **stochastic optimization problem**:

$$\min_{u_\tau} \mathbb{E}_{W_{\tau+1}} \Big[L_\tau(X_\tau, u_\tau, W_{\tau+1}) + \widetilde{V}_{\tau+1}\Big(f_\tau(X_\tau, u_\tau, W_{\tau+1})\Big) \Big]$$

$\Rightarrow$ this problem resumes to solve a LP at each timestep

- Get optimal solution $U_\tau^\#$
- Send $U_\tau^\#$ to assessor

A brief recall on Dynamic Programming

Dynamic Programming

μ_t is the probability law of W_t and is being used to estimate expectation and compute **offline** value functions with the backward equation:

$$V_T(x) = K(x)$$

$$V_t(x_t) = \min_{U_t} \mathbb{E}_{\mu_t}\Big[\underbrace{L_t(x_t, U_t, W_{t+1})}_{\text{current cost}} + \underbrace{V_{t+1}\Big(f(x_t, U_t, W_{t+1})\Big)}_{\text{future costs}}\Big]$$

A brief recall on Dynamic Programming

Dynamic Programming

μ_t is the probability law of W_t and is being used to estimate expectation and compute **offline** value functions with the backward equation:

$$V_T(x) = K(x)$$
$$V_t(x_t) = \min_{U_t} \mathbb{E}_{\mu_t}\Big[\underbrace{L_t(x_t, U_t, W_{t+1})}_{\text{current cost}} + \underbrace{V_{t+1}\Big(f(x_t, U_t, W_{t+1})\Big)}_{\text{future costs}} \Big]$$

Stochastic Dual Dynamic Programming

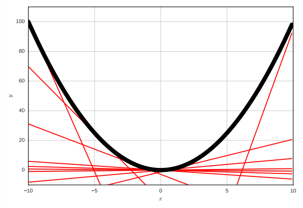

- Convex value functions V_t are approximated as a supremum of a finite set of affine functions
- Affine functions (=cuts) are computed during forward/backward passes, till convergence
- SDDP makes an extensive use of LP solver

$$\widetilde{V}_t(x) = \max_{1 \le k \le K} \{\lambda_t^k x + \beta_t^k\} \le V_t(x)$$

Outline

How to forecast solar irradiation?

We suppose that we have available at midnight a forecast $\hat{\Phi}$, with error bounds $(\varepsilon_0, \dots, \varepsilon_T)$. The realization of Φ_t is equal to

$$\Phi_t = \hat{\Phi}_t \times (1 + \varepsilon_t) \ .$$

How to forecast solar irradiation?

We suppose that we have available at midnight a forecast $\hat{\Phi}$, with error bounds $(\varepsilon_0, \ldots, \varepsilon_T)$. The realization of Φ_t is equal to

$$\Phi_t = \hat{\Phi}_t \times (1 + \varepsilon_t) \ .$$

Objective

We aim to identify the sensitivity of the two algorithms w.r.t the modelling of ε_t

How to forecast solar irradiation?

We suppose that we have available at midnight a forecast $\hat{\Phi}$, with error bounds $(\varepsilon_0, \dots, \varepsilon_T)$. The realization of Φ_t is equal to

$$\Phi_t = \hat{\Phi}_t \times (1 + \varepsilon_t) \; .$$

Objective

We aim to identify the sensitivity of the two algorithms w.r.t the modelling of ε_t

We model the error ε_t as a random variable. Different models are available:

- First with gaussian white noise, supposing that the process $(\varepsilon_0, \dots, \varepsilon_T)$ is time independent,
- Then with an autoregressive process, to have a more accurate modelling of the time dependency

Outline

White noise process

We recall that the irradiation corresponds to a forecast and an error:

$$\Phi_t = \hat{\Phi}_t \times (1 + \varepsilon_t)$$

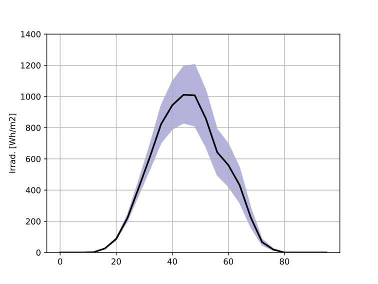

We first consider that for all t, ε_t is Gaussian:

$$\varepsilon_t \sim \mathcal{N}(0, \sigma_t)$$

and that the standard-deviation increases linearly over time

$$\sigma_t = \sigma_0 + (\sigma_T - \sigma_0)\frac{t}{T}$$

Discretizing the probability laws

Numerical optimization requires the discretization of continuous variables.

We use optimal quantization to approximate the continuous gaussian distribution of ε_t with a discrete probability distribution.

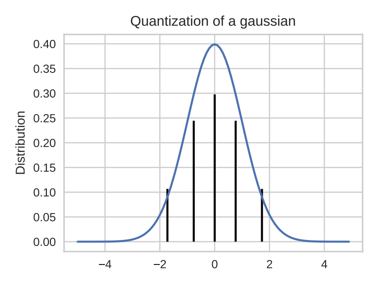

The probability measure of ε_t is approximated as

$$\mu[\varepsilon_t] \approx \sum_{i=1}^{n} \pi_i \delta_{w_i}$$

where π_i is the probability that the event $\varepsilon_t = w_i$ occurs.

Decision-Hazard or Hazard-Decision?

- In Decision-Hazard, the decision $\mathbf{U}_t$ is taken before the realization of the uncertainties $\mathbf{W}_{t+1}$ in $[t, t+1[$.
- In Hazard-Decision, the decision $\mathbf{U}_t$ is taken after the realization of the uncertainties $\mathbf{W}_{t+1}$ in $[t, t+1[$.

Hence irrealistic, Hazard-Decision gives a lower-bound of the Decision-Hazard problem.
(the more information, the better the algorithm is)

Hazard-Decision

In HD, we know the realization $w_{\tau+1}$ of the uncertainties $\mathbf{W}_{t+1}$ during the following interval $[\tau, \tau+1[$.

MPC forecasts:

$$\big(w_{\tau+1}, \mathbb{E}(\mathbf{W}_{\tau+2}), \ldots, \mathbb{E}(\mathbf{W}_T)\big)$$

and solves the deterministic optimization problem.

Hazard-Decision

In HD, we know the realization $w_{\tau+1}$ of the uncertainties $\mathbf{W}_{t+1}$ during the following interval $[\tau, \tau+1[$.

MPC forecasts:

$$\big(w_{\tau+1}, \mathbb{E}(\mathbf{W}_{\tau+2}), \dots, \mathbb{E}(\mathbf{W}_T)\big)$$

and solves the deterministic optimization problem.

SDDP solves the following LP problem:

$$\begin{aligned}
&\min_{u_\tau} \Big[L_\tau(x_\tau, u_\tau, w_{\tau+1}) + \theta \Big] \\
s.t \quad & x_{\tau+1} = f_\tau(x_\tau, u_\tau, w_{\tau+1}) \\
& \theta \geq \big\langle \lambda^c_{\tau+1}, x_{\tau+1} \big\rangle + \beta^c_{\tau+1} \quad \forall c \in \mathbb{C}_{\tau+1}
\end{aligned}$$

where $\mathbb{C}_{\tau+1}$ is the set of cuts uses to approximate the value function $V_{\tau+1}$.

Decision-Hazard

In DH, we know only the probability distribution of the uncertainties $\mathbf{W}_{t+1}$

MPC forecasts:

$$\big(\mathbb{E}(\mathbf{W}_{\tau+1}), \mathbb{E}(\mathbf{W}_{\tau+2}), \ldots, \mathbb{E}(\mathbf{W}_{T})\big)$$

and solves the deterministic optimization problem.

Decision-Hazard

In DH, we know only the probability distribution of the uncertainties $\mathbf{W}_{t+1}$

MPC forecasts:

$$\Big(\mathbb{E}\big(\mathbf{W}_{\tau+1}\big), \mathbb{E}\big(\mathbf{W}_{\tau+2}\big), \ldots, \mathbb{E}\big(\mathbf{W}_{T}\big)\Big)$$

and solves the deterministic optimization problem.

SDDP solves the following LP problem:

$$\begin{aligned} \min_{u_\tau} \ & \sum_{i=1}^{n} \pi_i \Big(L_\tau(x_\tau, u_\tau, w^i_{\tau+1}) + \theta^i \Big) \\ s.t \quad & x^i_{\tau+1} = f_\tau(x_\tau, u_\tau, w^i_{\tau+1}) \quad \forall i \\ & \theta^i \geq \big\langle \lambda^c_{\tau+1}, x^i_{\tau+1} \big\rangle + \beta^c_{\tau+1} \quad \forall i, \ c \in \mathbb{C}_{\tau+1} \end{aligned}$$

where $\mathbb{C}_{\tau+1}$ is the set of cuts uses to approximate the value function $V_{\tau+1}$ and n is the size of the discrete probability law.

Numerical settings

We compare different level of uncertainties, corresponding to different final standard-deviation σ_T.

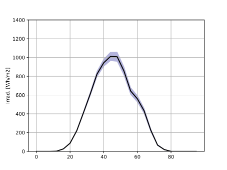

$\sigma_T = 5$ %

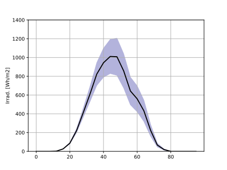

$\sigma_T = 20$ %

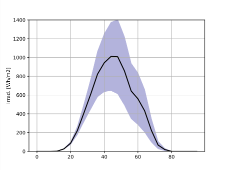

$\sigma_T = 40\%$

Assessment scenarios

We generate n_{assess} scenarios

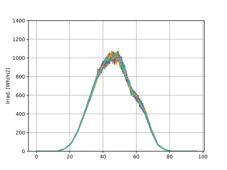

$\sigma_T = 5$ %

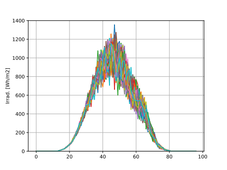

$\sigma_T = 20$ %

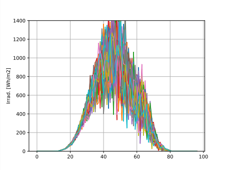

$\sigma_T = 40\%$

And then, let's roll!

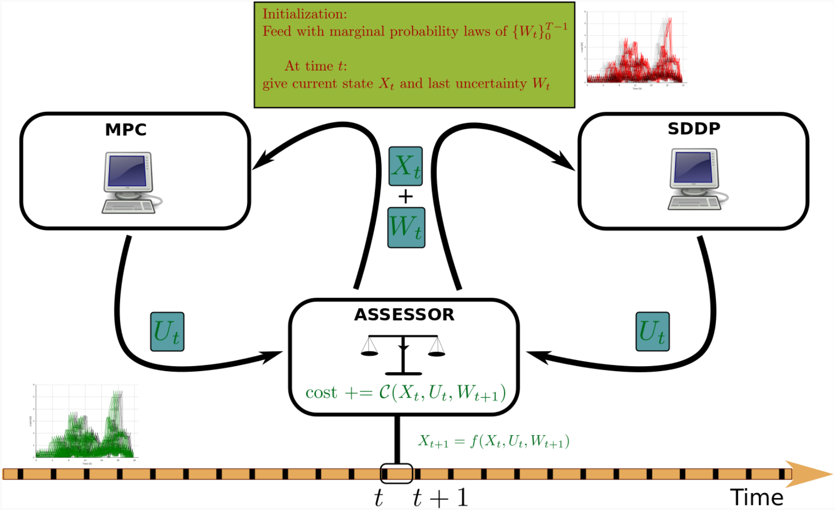

Results

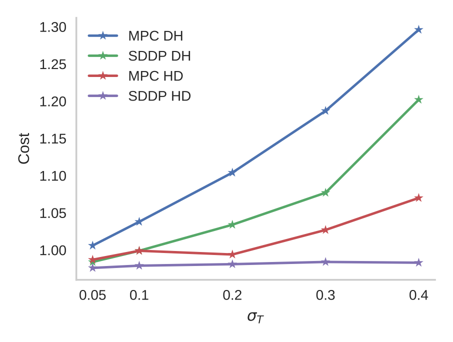

	HD		DH	
σ_T	SDDP	MPC	SDDP	MPC
5 %	0.976	0.987	0.984	1.006
10 %	0.979	0.999	0.984	1.038
20 %	0.981	0.994	1.034	1.104
30 %	0.984	1.027	1.077	1.187
40 %	0.983	1.070	1.202	1.296

Outline

Description

Modelling solar irradiation with white noise is a shortfall.

We rather have to model the process $(\varepsilon_0, \ldots, \varepsilon_T)$ as an ARMA process.

We define the nebulosity as:

$$n_t^s = \frac{\Phi_t}{\Phi_t^{clear}}$$

- Φ_t^{clear} is given by some trigonometric laws (position of the sun in the sky).
- n_t^s can be modelled with an AR process.

Description

Modelling solar irradiation with white noise is a shortfall.

We rather have to model the process $(\varepsilon_0, \ldots, \varepsilon_T)$ as an ARMA process.

We define the nebulosity as:

$$n_t^s = \frac{\Phi_t}{\Phi_t^{clear}}$$

- Φ_t^{clear} is given by some trigonometric laws (position of the sun in the sky).
- n_t^s can be modelled with an AR process.

Still a work in progress! :-)

Conclusion

- The more uncertainties, the better SDDP is towards MPC
- We obtained similar results while tackling electrical and hot water demands
- We have to study more realistic uncertainties, corresponding to real data
- We aim to use decomposition algorithms to tackle bigger problems, with a lot more houses! :-D