Théorie des jeux pour la négociation entre la production d'énergie et les tâches informatique dans un datacenter à énergies renouvelables

Jean-Marc Pierson
Jean-Marc.Pierson@irit.fr

GreenDays, June 25th, 2019

Context

Does using IT technologies have any consequences ?

⇒ IT consumes a huge amount of energy

- sending an email with an attach file consumes as much as one low-power bulb of high power for one hour

⇒ Data Centers in the US consumed 91 billions of kWh in 2013

⇒ Data Centers in Europe consumed 56 billions of kWh in 2013

- increasing the energy efficiency of data-centers
- supplying data-centers with only green energy

Does using IT technologies have any consequences ?

⇒ IT consumes a huge amount of energy

- sending an email with an attach file consumes as much as one low-power bulb of high power for one hour

⇒ Data Centers in the US consumed 91 billions of kWh in 2013

⇒ Data Centers in Europe consumed 56 billions of kWh in 2013

- increasing the energy efficiency of data-centers
- supplying data-centers with only green energy

Data Centers and energy efficiency

Does using IT technologies have any consequences ?

⇒ IT consumes a huge amount of energy
- sending an email with an attach file consumes as much as one low-power bulb of high power for one hour

⇒ Data Centers in the US consumed 91 billions of kWh in 2013

⇒ Data Centers in Europe consumed 56 billions of kWh in 2013

Data Centers reached 4% of the global energy consumption in 2015

- increasing the energy efficiency of data-centers
- supplying data-centers with only green energy

Data Centers and energy efficiency

Does using IT technologies have any consequences ?

⇒ IT consumes a huge amount of energy

- sending an email with an attach file consumes as much as one low-power bulb of high power for one hour

⇒ Data Centers in the US consumed 91 billions of kWh in 2013

⇒ Data Centers in Europe consumed 56 billions of kWh in 2013

Data Centers reached 4% of the global energy consumption in 2015

- increasing the energy efficiency of data-centers
- supplying data-centers with only green energy

DATAZERO : an innovative data-center model

DATAZERO : an innovative data-center model

Adapting the IT load to
the available power
&
Adapting the power to
the incoming IT load

while using a mix of only green energy sources (without grid power usage)

DATAZERO : the big picture

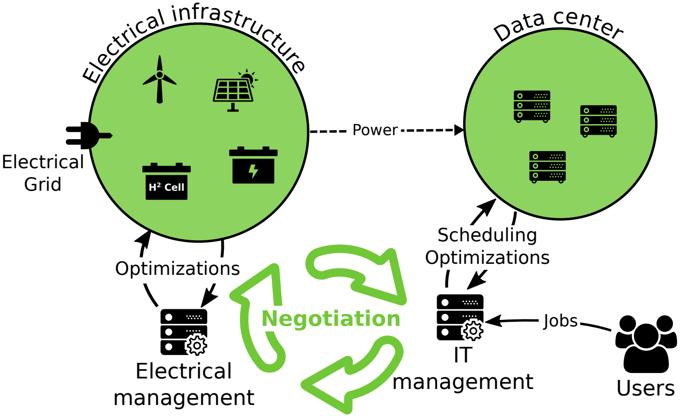

Outline

Negotiation between ITDM and PDM

Context and overview of the problem

Problem statement

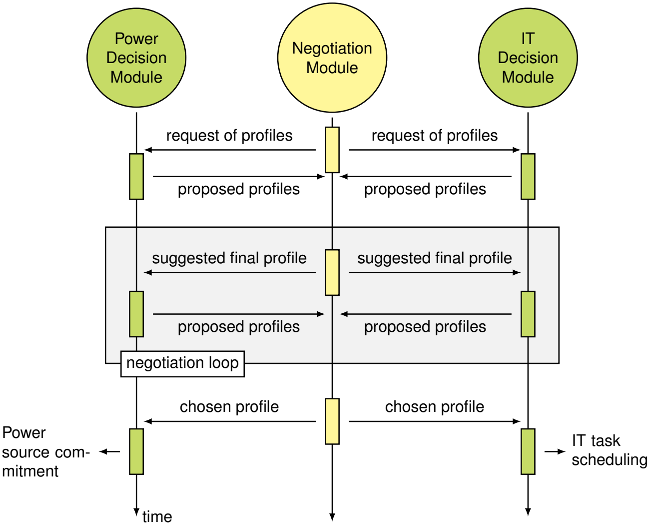

Infrastructure and negotiation

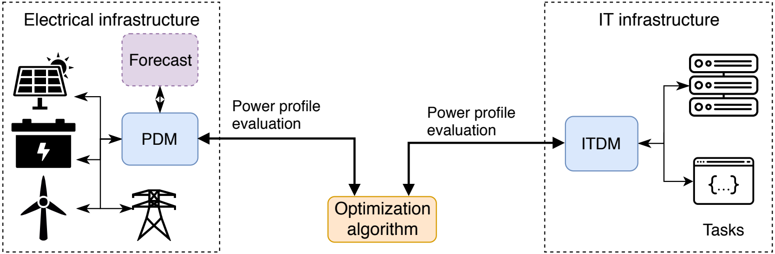

Separated IT and electrical optimizations

- Ability to evaluate power plan impact
- Internal objective (utility)
- Black box functions $\mathbb{R}^T \rightarrow \mathbb{R}$
- Computationally expensive

Infrastructure and negotiation

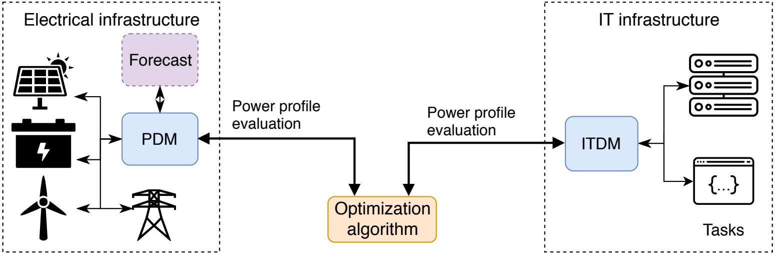

Separated IT and electrical optimizations

- Ability to evaluate power plan impact
- Internal objective (utility)
- Black box functions $\mathbb{R}^T \to \mathbb{R}$
- Computationally expensive

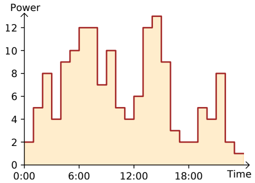

Infrastructure and negotiation

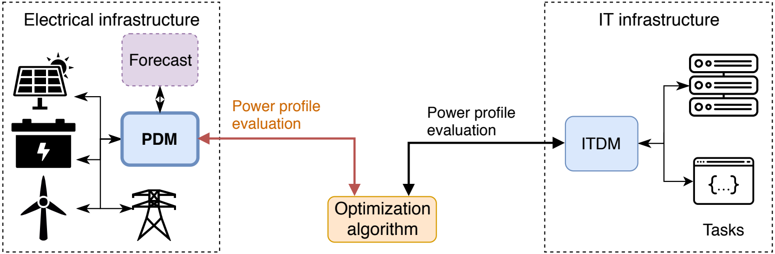

Separated IT and electrical optimizations

- Ability to evaluate power plan impact
- Internal objective (utility)
- Black box functions $\mathbb{R}^T \rightarrow \mathbb{R}$
- Computationally expensive

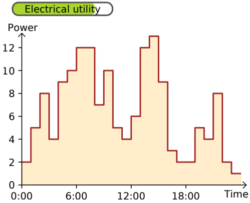

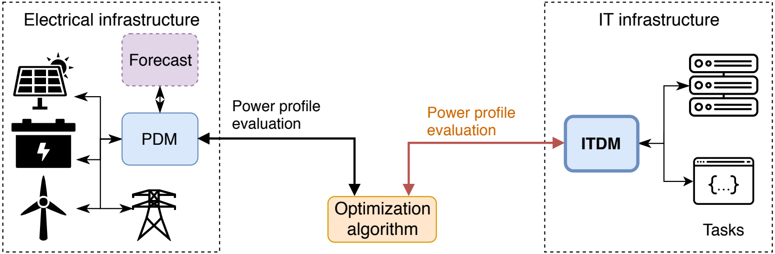

Separated IT and electrical optimizations

- Ability to evaluate power plan impact
- Internal objective (utility)
- Black box functions $\mathbb{R}^T \rightarrow \mathbb{R}$
- Computationally expensive

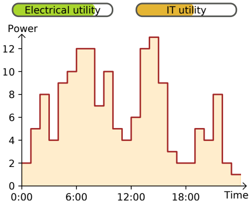

Multi-objective aspect

- Each DM has one or more objectives to satisfy
- Objectives may differ between DM
  - QoS related for ITDM, environmental impact for PDM

Managing different objectives

3 options studied :

- Finding a set of good solutions (set of possible trade-offs) (Pareto-based approach)
- Maximizing the weighted sum of the utilities, under the constraint of a distance between the two resulting profiles (SAN approach)
- Playing a game between the PDM and the ITDM so that each one maximizes its profit (GAN approach)

Multi-objective aspect

- Each DM has one or more objectives to satisfy
- Objectives may differ between DM
  - QoS related for ITDM, environmental impact for PDM

Managing different objectives

3 options studied :

- Finding a set of good solutions (set of possible trade-offs) (Pareto-based approach)
- Maximizing the weighted sum of the utilities, under the constraint of a distance between the two resulting profiles (SAN approach)
- Playing a game between the PDM and the ITDM so that each one maximizes its profit (GAN approach)

Pareto front

Pareto front : Multi-objective optimization and heuristics

- Find Pareto front (best trade-offs)

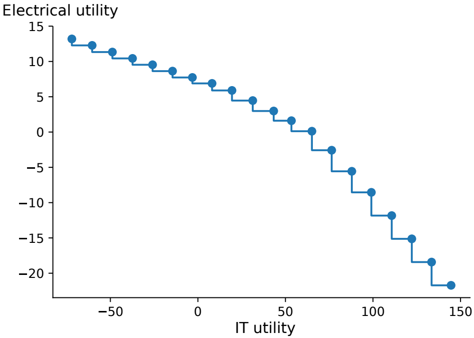

Multi-Objective Evolutionary Algorithms

- Well studied area, various approaches
- Focused on SPEA2 (genetic algorithm). Maximization of the hypervolume of solutions

- Find Pareto front (best trade-offs)

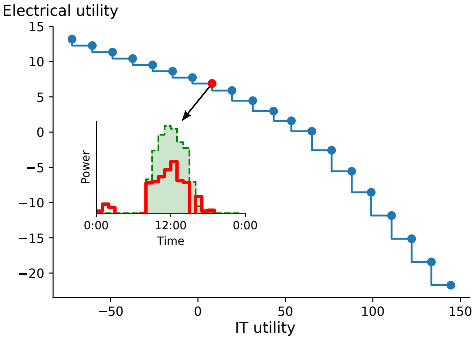

Multi-Objective Evolutionary Algorithms

- Well studied area, various approaches
- Focused on SPEA2 (genetic algorithm). Maximization of the hypervolume of solutions

- Find Pareto front (best trade-offs)

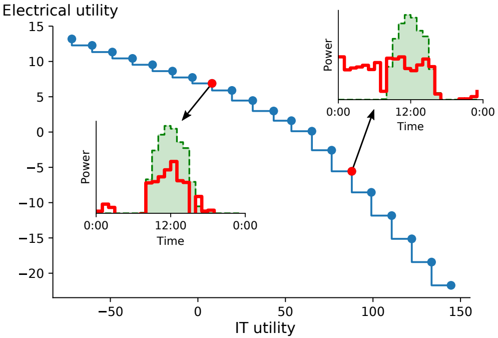

Multi-Objective Evolutionary Algorithms

- Well studied area, various approaches
- Focused on SPEA2 (genetic algorithm). Maximization of the hypervolume of solutions

Pareto front : Multi-objective optimization and heuristics

- Find Pareto front (best trade-offs)

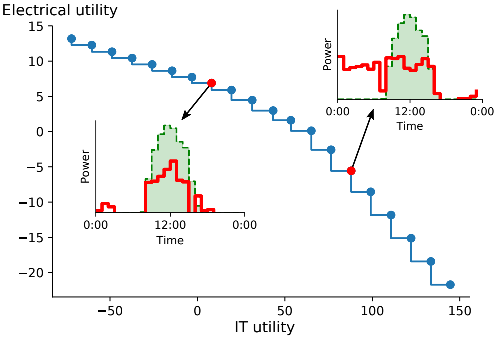

Multi-Objective Evolutionary Algorithms

- Well studied area, various approaches
- Focused on SPEA2 (genetic algorithm). Maximization of the hypervolume of solutions

SAN and GAN : Turn based approaches

Overview of SAN and GAN approaches

Main points

- Both algorithms are based on scheduling
  - DMs generate multiple scheduling solutions
  - Then we find negotiation solution from those scheduling solutions
- Both SAN and GAN negotiates in turn-based strategy
  - When ITDM runs scheduling (to follow PDM), PDM does not, and vice versa
  - We define 2 modes : "Follow PDM" mode (FLW_PD) and "Follow ITDM" mode (FLW_IT)
  - For both SAN and GAN (for the entire of the presentation), the whole system is executed under only 1 mode at a time

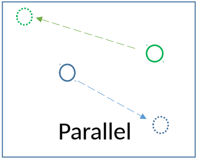

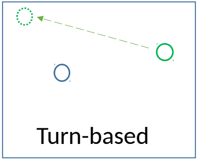

Scheduling Based Negotiation (SAN)

SAN algorithm

Definitions

- The set of ITDM profiles is $\{x_1, x_2, ..., x_m\}$
- The set of PDM profiles is $\{y_1, y_2, ..., y_n\}$
- Depending on each specific context, a profile may also be named "hint" or "candidate"

2 stages

- Stage 1 : Checking for matched pair
  - Decision variable : the pair {PDM profile, ITDM profile} : $\{x \in \{x_1, x_2, ..., x_m\}, y \in \{y_1, y_2, ..., y_n\}\}$
  - Objective : maximize sum of utility

$$\max_{\{x,y\}} (u(x) + u(y)) \tag{1}$$

  - Constraint : $d(x, y) < \epsilon$
    where $d()$ is the distance between x and y
  - If we can't find any matched pair, run Stage 2 : Negotiating.

Stage 2 : Negotiating :

- General mechanism
  - At a time, the whole system is executed under only 1 mode : follow ITDM (FLW_IT) or follow PDM (FLW_PD)
  - NM decides to switch between two modes using *"verify quality of rescheduling"*
  - Repeat until matched pair found
- Two modes :
  - FLW_IT : Follow the ITDM
    - NM sends the ITDM hints to PDM
    - PDM uses an algorithm (e.g. greedy, linear program) to find multiple scheduling solutions as candidates
    - PDM evaluates quality of candidates by *"weighted similarity"* to hints
    - PDM selects candidates with high *"weighted similarity"* as its news hints and sends back to NM
  - FLW_PD : Similar to FLW_IT, following the PDM

Algorithm

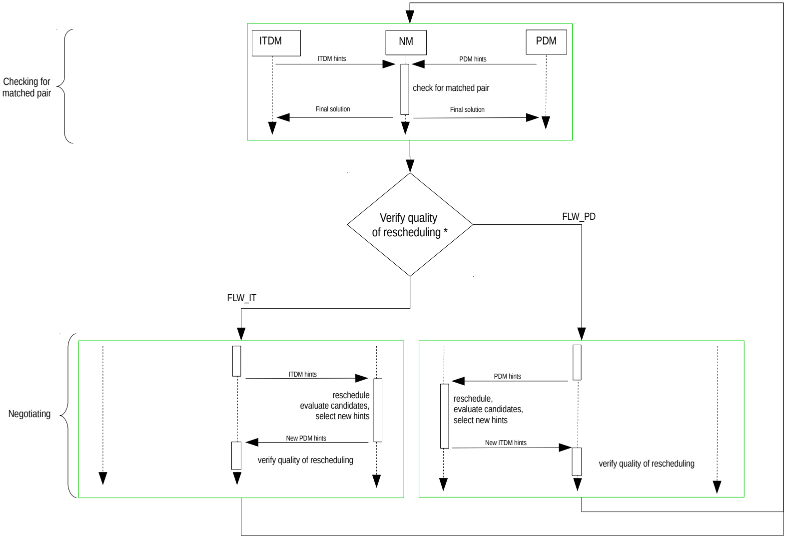

(*) verfify quality of rescheduling: compare “distance between the best ITDM hint and the best PDM hint” before and after rescheduling

Formulation

Formulation : Distance measurement

- Mean Square Error between profile $x = \{x_1, ..., x_T\}$ and $y = \{y_1, ..., y_T\}$:

$$d(x,y) = \frac{1}{T}\sum_{i=1}^{T}(x_i - y_i)^2 \tag{2}$$

- Pearson correlation between them :

$$d(x,y) = \frac{\sqrt{\sum_{i=1}^{T}(x_i - \bar{x})^2}\sqrt{\sum_{i=1}^{T}(y_i - \bar{y})^2}}{\sum_{i=1}^{T}(x_i - \bar{x})(y_i - \bar{y})}, \tag{3}$$

Formulation : Weighted similarity

$$w(x, Y) = \sum_{n=1}^{N} (u(x) + u(y^n)) \frac{1}{d(x, y^n)}, \quad (4)$$

Or, if we want to adjust impact of utility and distance :

$$w(x, Y) = \sum_{n=1}^{N} \left(\alpha(u(x) + u(y^n)) + (1 - \alpha) \frac{1}{d(x, y^n)} \right), \quad (5)$$

where $Y = \{y^1, ..., y^n, ...y^N\}$ and $y^n = \{y_1^n, ..., y_T^n\}$, and the values are normalized to (0,1]

Example

Example

Follow ITDM

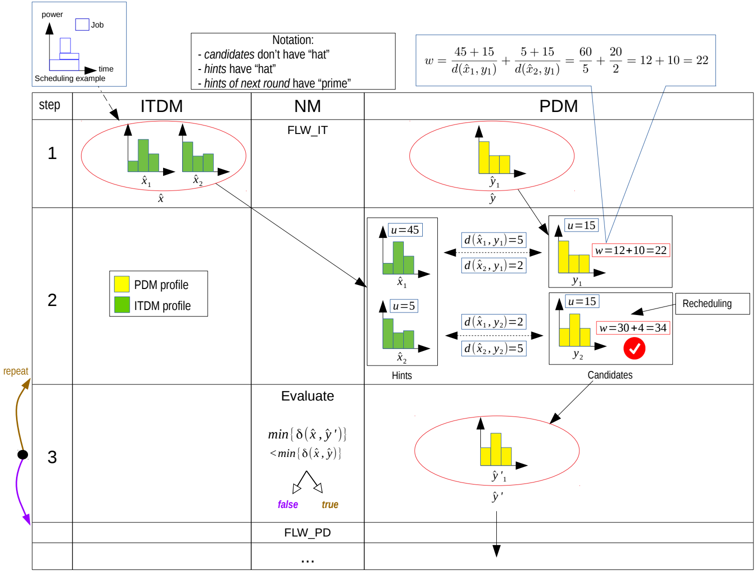

Follow PDM

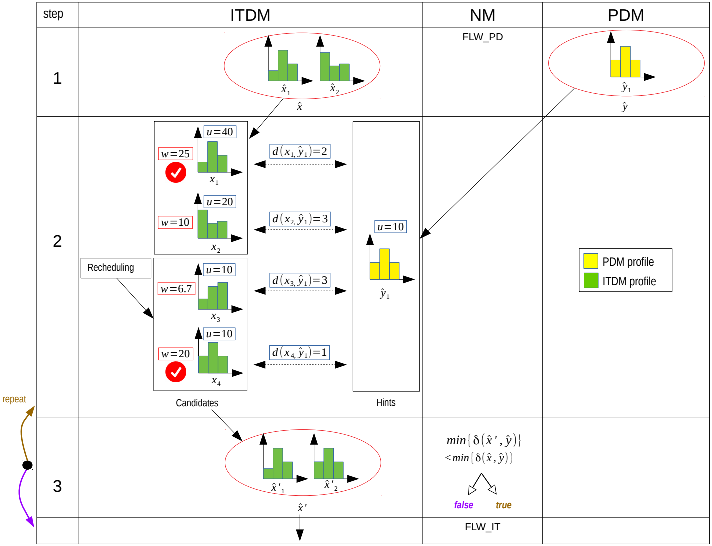

step
ITDM
NM
PDM
1
$\hat{x}_1$
$\hat{x}_2$
$\hat{x}$
FLW_PD
$\hat{y}_1$
$\hat{y}$
2
$u=40$
$w=25$
x_1
$u=20$
$w=10$
x_2
Recheduling
$u=10$
$w=6.7$
x_3
$u=10$
$w=20$
x_4
Candidates
$d(x_1, \hat{y}_1)=2$
$d(x_2, \hat{y}_1)=3$
$d(x_3, \hat{y}_1)=3$
$d(x_4, \hat{y}_1)=1$
$u=10$
$\hat{y}_1$
Hints
PDM profile
ITDM profile
repeat
3
$\hat{x}'_1$
$\hat{x}'_2$
$\hat{x}'$
$min\{\delta(\hat{x}', \hat{y})\}$
$< min\{\delta(\hat{x}, \hat{y})\}$
false
true
FLW_IT

GAME BASED NEGOTIATION (GAN)

Model

Model

Game players

ITDM:

- Objective: (2) (1)
  - max {payment from users – payment to PDM}
- Decision variables:
  - price
  - scheduling

revenue - cost

PDM:

- Objective: (1) (2)
  - max {payment from ITDM – (opex + capex cost}
- Decision variables:
  - purchased power
  - scheduling

Game model

Hybrid model

- Non-cooperative: each player maximizes their own utility
- Cooperative: sometimes a player follows the other's suggestion

Supplier-buyer game diagram

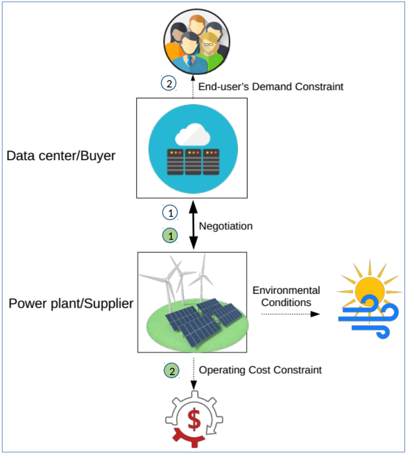

Model

Motivation

- The buying-selling nature of the system
  - Selling : PDM is selling power $\rightarrow$ controls the price
  - Buying : ITDM is buying power $\rightarrow$ decide the order/purchase
- Advantages of pricing
  - Power source availability can be reflected in price
  - Through price, pattern of order reflects pattern of PDM's desirable supply $\rightarrow$ drive demand toward supply

Model

Preliminary

- Players are selfish
  - They try to maximize their utility
  - Each player negotiates just because he foresees some benefit
  - We introduce *incentive pricing mechanism* : each player tries to find offers that are attractive to the other player.
- An unexpected situation may occurred : all players can't foresee their benefit and stop negotiate without reaching any agreement.
  - From the view of the whole system, this situation is unacceptable, no transaction is done, the players obtain zero utility
  - If this situation occurred, we introduce *sacrifice mechanism*, in which the players gradually sacrifice their utility until they reach an agreement.

- T : Time window
- $\hat{x}, x^{IT}, x^{PD}$ are profiles
- $\hat{x} = \{\hat{x}_1, \hat{x}_2, ..., \hat{x}_T\}$, $x^{IT} = \{x_1^{IT}, x_2^{IT}, ..., x_T^{IT}\}$ $x^{PD} = \{x_1^{PD}, x_2^{PD}, ..., x_T^{PD}\}$

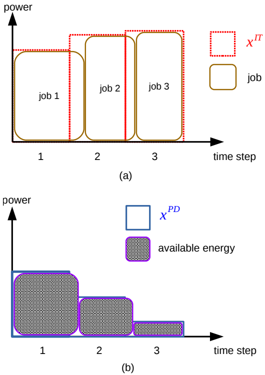

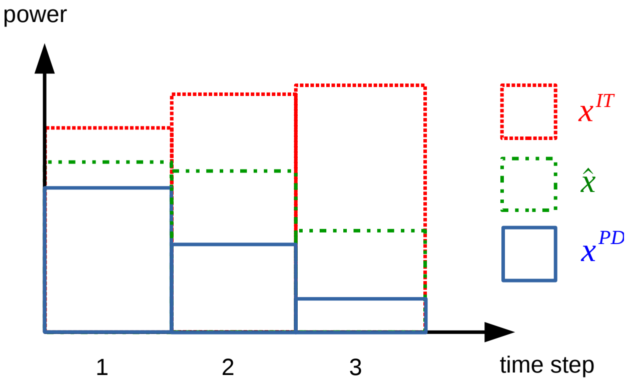

Definition 2

- $\pi = \{\pi_1, \pi_2, \ldots, \pi_T\}$, $\pi^{IT} = \{\pi_1^{IT}, \pi_2^{IT}, \ldots, \pi_T^{IT}\}$, $\pi^{PD} = \{\pi_1^{PD}, \pi_2^{PD}, \ldots, \pi_T^{PD}\}$

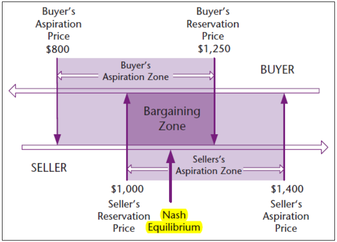

2006__Krause__Bargaining Stances and Outcomes in Buyer–Seller Negotiations: Experimental Results (The Journal of Supply Chain Management)

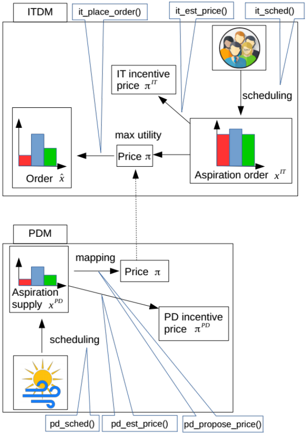

Negotiation Model

Negotiation Model

Turn-based play

- There are 3 variables :
  - ITDM local variable $it\_mod = \{FLW\_IT, FLW\_PD\}$,
  - PDM local variable $pd\_mod = \{FLW\_IT, FLW\_PD\}$,
  - global variable $mod = \{FLW\_IT, FLW\_PD\}$.
- The mode of the system is only depended on *mod*. Two local variables only show the capability of the DMs.
- At a time, the system is executed under only one mode
- However, each player is selfish, he always wants the other player to follow himself, i.e., $it\_mod = FLW\_IT$, $pd\_mod = FLW\_PD$. In this situation, the negotiation can't be processed.
- Therefore
  - if ITDM is also capable of following PDM, we will set $it\_mod = FLW\_PD$
  - if PDM is also capable of following ITDM, we set $pd\_mod = FLW\_IT$

Negotiation solution

- Stopping criteria : $\hat{x}$ approximates both x^{IT} and x^{PD}
- Final solution is the last $\hat{x}$

Incentive pricing mechanism :

- ITDM follows PDM : when x^{PD} and π^{PD} are attractive * :
  it_cost(aspiration supply, PD incentive price) $<$ it_cost(order, price) :

$$c(x^{PD}, \pi^{PD}) < c(\hat{x}, \pi) \quad (6)$$

- PD follows ITDM : when ITDM's offers are attractive :
  pd_revenue(aspiration order, IT incentive price) $>$ pd_revenue(order, price)

$$r(x^{IT}, \pi^{IT}) > r(\hat{x}, \pi) \quad (7)$$

(*) next page

Negotiation Model

(*) How π^{PD} can be attractive to ITDM :

- Given :
  - Definition : x^{PD} is the PDM's desirable supply
  - Then, the cost associated with this supply is lower than the cost associated with other supplies
- As a result
  - the PDM estimates the amount of cost it can reduce when it provides the ITDM with x^{PD}, instead of $\hat{x}$.
  - the PDM computes a π^{PD}such that its total utility increases

$$u(x^{PD}, \pi^{PD}) > u(\hat{x}, \pi) \tag{8}$$

How π^{IT} can be attractive to PDM : similarly

Negotiation Model

Sacrifice mechanism

- Unexpected situation : both ITDM and PDM are not capable of following each other, but an agreement is not reached
- Solution : ITDM gradually increases its *sacrifice variable* α :

$$\alpha \leftarrow \alpha + \gamma \tag{9}$$

  then the incentive pricing mechanism at ITDM becomes :

$$c(x^{PD}, \pi^{PD}) - \alpha < c(\hat{x}, \pi) \tag{10}$$

- Similarly, incentive pricing mechanism at PDM becomes :

$$r(x^{IT}, \pi^{IT}) + \alpha > r(\hat{x}, \pi) \tag{11}$$

Algorithm

Algorithm

Note :

- At a time, the whole system is executed under only one mode : FLW_IT or FLW_PD
- DMs only exchange data that have been updated/modified, other data can be stored and reused

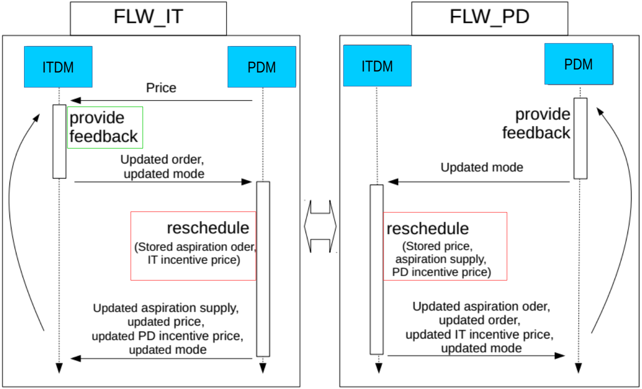

Algorithm

Graphical interpretation of the algorithm

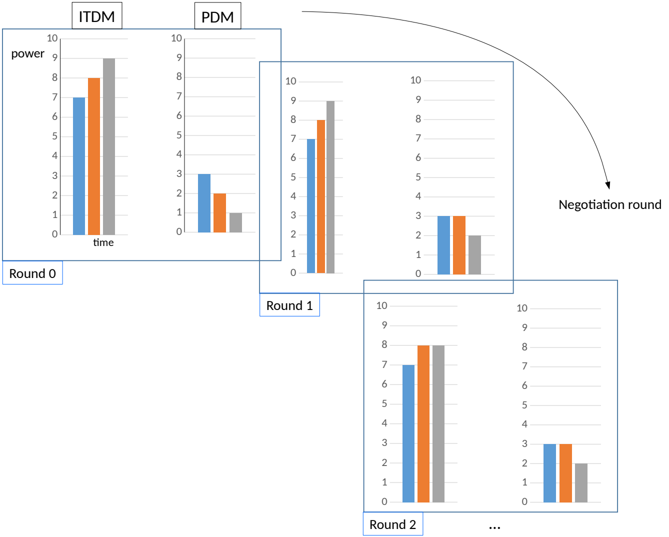

ITDM Formulation

$$u(x,\pi) = r(x) - c(x,\pi) = \sum_{i=1}^{J} u_i(x) - \sum_{k=1}^{T} \pi_k x_k, \tag{12}$$

where

- $r(\cdot)$: revenue of ITDM
- $c(\cdot)$: ITDM's payment to PDM
- J : the number of ITDM's jobs
- $u_i(\cdot)$: the payment of the users' i-th job to ITDM

The payment from users is computed based on the Amazon EC2 pricing

PDM Formulation

PDM Formulation

$$u(x,\pi,S) = r(x,\pi) - c(x,S) = \sum_{i=1}^{T} \pi_i x_i - \sum_{j \in S} c_j^{OP}(x) c_j^{CAP} \quad (13)$$

where

- $c_j^{OP}(\cdot)$: operational rate of the j-th power source component
- $c_j^{CAP}(\cdot)$: the capital cost
- S is the number of utilized power source components

ITDM Procedures

- it_sched() : Generate x^{IT} from the scheduling solution, such that
  - Utility $u(x^{IT}, \pi)$ is maximized
  - Revenue of x^{IT} must be larger than revenue of previous round's aspiration order $\dot{x}^{IT}$
  - And x^{IT} has to be closer to x^{PD} than previous round's aspiration order $\dot{x}^{IT}$

$$x^{IT} = \arg\max_{x} \; u(x, \pi) \tag{14}$$

$$s.t \quad r(x) > r(\dot{x}^{IT}) \tag{15}$$

$$d(x, x^{PD}) < d(\dot{x}^{IT}, x^{PD}) \tag{16}$$

ITDM Procedures

- it_place_order() : ITDM finds $\hat{x}$ that maximizes the ITDM's utility function $u(x, \pi)$

$$\hat{x} = \arg\max_{x} u(x, \pi) \quad (17)$$

- it_est_price() : ITDM estimates new π^{IT} that is more attractive to PDM than previous round's IT incentive price $\dot{\pi}^{IT}$, while keeping the ITDM's total utility non-decreased :

$$p = \pi^{IT} = \dot{\pi}^{IT} \quad (18)$$

$$while \;\; u(x^{IT}, p) \geq u(\hat{x}, \pi) \quad (19)$$

$$\pi^{IT} = p \quad (20)$$

$$p_i = p_i + p_i/N, \qquad i = 1, ..., T \quad (21)$$

where N is an integer, set through experiment parameters

PDM Procedures

- pd_propose_price(x^{PD}) : the price π is generated such that

$$\pi_i = \frac{1}{x_i^{PD}}, \quad i = 1, ..., T \tag{22}$$

where x_i^{PD} is normalized to (0,1]

- pd_est_price() : Similar to it_est_price(), PDM estimates a new π^{PD} that is more attractive to ITDM than previous round's PD incentive price $\dot{\pi}^{PD}$, while keeping PDM's utility non-decreased :

$$p = \pi^{PD} = \dot{\pi}^{PD} \tag{23}$$

$$while \;\; u(x^{PD}, p) \geq u(\hat{x}, \pi) \tag{24}$$

$$\pi^{PD} = p \tag{25}$$

$$p_i = p_i - p_i/N, \qquad i = 1, ..., T \tag{26}$$

PDM Procedures

- pd_sched() : Generate x^{PD} from the scheduling solution, such that
  - Utility $u(x^{PD}, \pi, S)$ is maximized
  - Cost of x^{PD} must be smaller than cost of previous round's aspiration supply $\dot{x}^{PD}$
  - New price π has to be closer to π^{IT} than the previous round's price $\dot{\pi}$

$$x^{PD} = \arg\max_{x} \; u(x, \pi, S) \tag{27}$$

$$s.t \quad c(x) < c(\dot{x}^{PD}) \tag{28}$$

$$d(\pi, \pi^{IT}) < d(\dot{\pi}, \pi^{IT}) \tag{29}$$

where $\pi \leftarrow$ pd_propose_price(x)

Example

Simplified example

round	row	PD-Player	Common Space	IT-Player
1	1	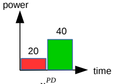 x^{PD}	$\pi^{PD}=\{7,6\}$ $\pi=\{8,6\}$ 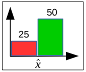 $\hat{x}$ $\pi^{IT}=\{3,3\}$	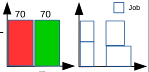 x^{IT} Scheduling solution
	2	Assume $\begin{cases} r(x^{IT},\pi^{IT})+\alpha=0.25 \\ r(\hat{x},\pi)=0.22 \end{cases}$ $r(x^{IT},\pi^{IT})+\alpha>r(\hat{x},\pi)$ → ***true***		Assume $\begin{cases} c(x^{PD},\pi^{PD})-\alpha=0.13 \\ c(\hat{x},\pi)=0.15 \end{cases}$ $c(x^{PD},\pi^{PD})-\alpha<c(\hat{x},\pi)$ → ***true***
	3	*pd_pre = FLW_IT*	***mod = FLW_IT***	*it_pre = FLW_PD*
2	4	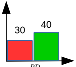 x^{PD}	$\pi^{PD}=\{7,6\}$ $\pi=\{7,6\}$ 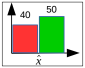 $\hat{x}$ $\pi^{IT}=\{3,3\}$	
	5	Assume $\begin{cases} r(x^{IT},\pi^{IT})+\alpha=0.25 \\ r(\hat{x},\pi)=0.28 \end{cases}$ $r(x^{IT},\pi^{IT})+\alpha>r(\hat{x},\pi)$ → ***false***		Assume $\begin{cases} c(x^{PD},\pi^{PD})-\alpha=0.16 \\ c(\hat{x},\pi)=0.17 \end{cases}$ $c(x^{PD},\pi^{PD})-\alpha<c(\hat{x},\pi)$ → ***true***
	6	*pd_pre = FLW_PD*	***mod = FLW_PD***	*it_pre = FLW_PD*
3	7		$\pi^{PD}=\{7,6\}$ $\pi=\{7,6\}$ 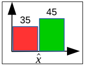 $\hat{x}$ $\pi^{IT}=\{6,5\}$	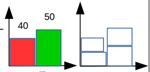 x^{IT} New scheduling solution
	8	Assume $\begin{cases} r(x^{IT},\pi^{IT})+\alpha=0.23 \\ r(\hat{x},\pi)=0.24 \end{cases}$ $r(x^{IT},\pi^{IT})+\alpha>r(\hat{x},\pi)$ → ***false***		Assume $\begin{cases} c(x^{PD},\pi^{PD})-\alpha=0.16 \\ c(\hat{x},\pi)=0.14 \end{cases}$ $c(x^{PD},\pi^{PD})-\alpha<c(\hat{x},\pi)$ → ***false***
repeat	9	*pd_pre = FLW_PD*		*it_pre = FLW_IT*
	10	$d(x^{PD},\hat{x})>\varepsilon$ ***true*** / ***false***		$d(x^{PD},\hat{x})>\varepsilon$ ***true*** / ***false***

Experiment

Setup of Experiment

Real PDM & ITDM

- PDM weather information : 1 month
- Time window : 3 days or 72 hours
- Timestep : 1 hour or 3,600,000 ms
- PDM sizing : ≈1kW
- Run time : ≈10 minutes

Distance

- Calculation : using Pearson correlation
- Distance is not always decreasing because the profiles are evaluated by both utility and distance
- Negotiation results depend a lot on the series of utilities from DMs

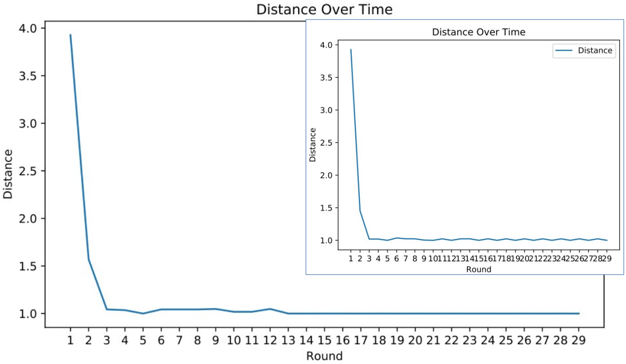

Violation

- Calculation : sum of the amount that the ITDM profiles excesses PDM's profiles
- A significant reason for this result : DMs scheduling algorithms

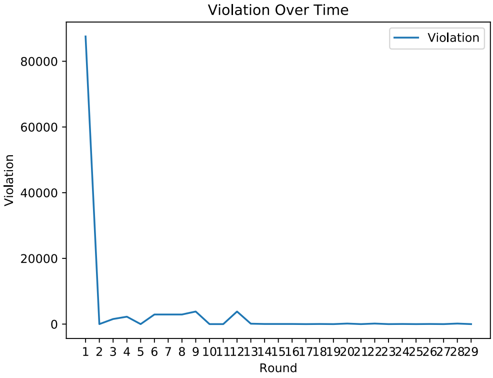

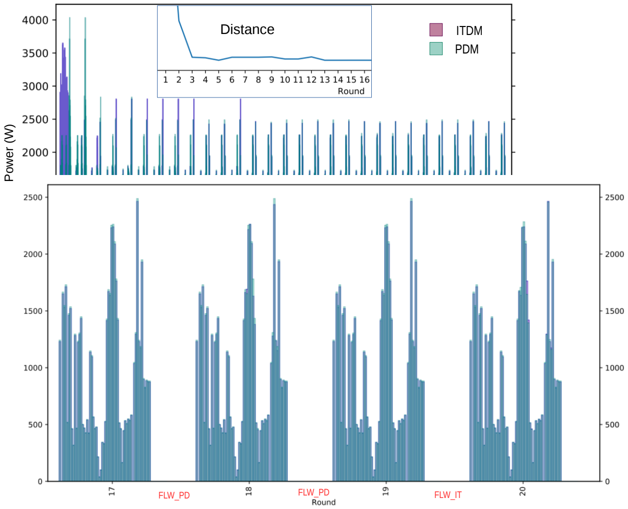
Distance
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
Round
ITDM
PDM
Power (W)
4000
3500
3000
2500
2000
2500
2000
1500
1000
500
0
17
18
19
20
FLW_PD
FLW_PD
FLW_IT
Round

Visit www.datazero.org for more information ! !