A Pareto-based GA for Scheduling HPC Applications on Distributed Cloud Infrastructures

Y. Kessaci, N. Melab et E-G. Talbi
Dolphin Team , Université Lille 1, LIFL-CNRS, INRIA Lille-Nord Europe

Outline

- Motivation
- Cloud Model
- Contribution
- Experimentations
- Perspectives
- Conclusion

Motivation

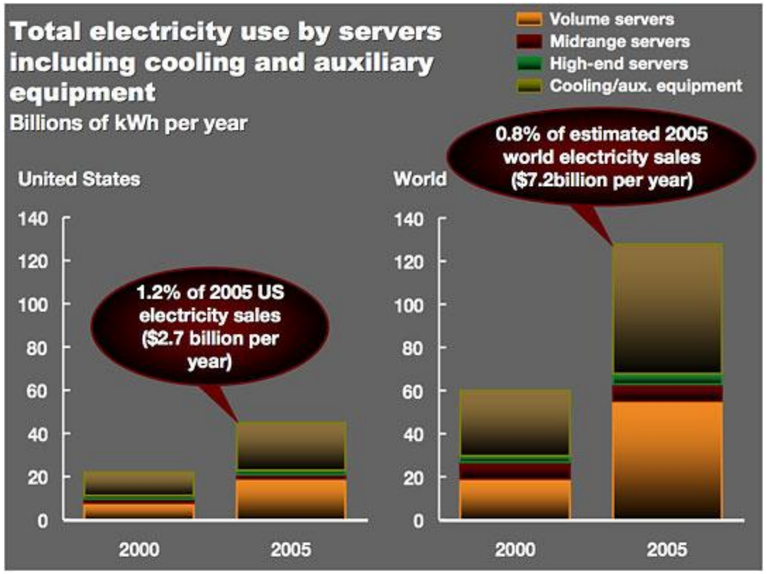

Source: Koomey 2007

➢ The total electricity consumed by servers doubled over the period 2000 to 2005 in worldwide.

➢ Information and Communications Technology (ICT) industry accounts for approximately **2%** of CO2 emissions.

➢ Energy consumption represents more than 42% of the total data center budget.

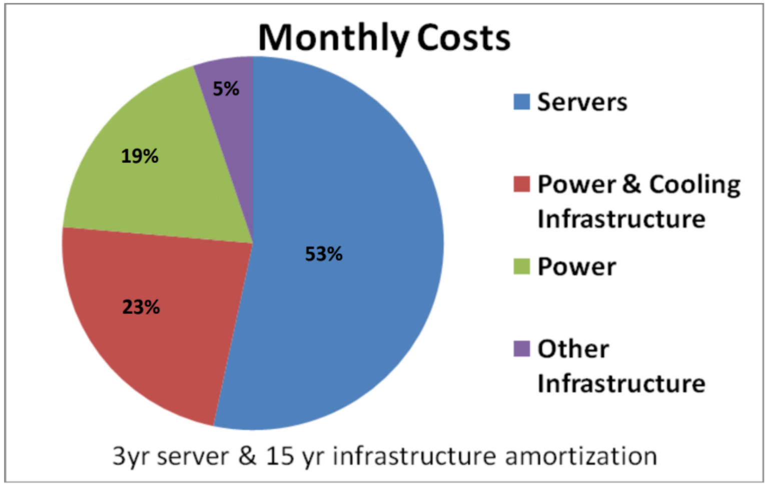

Source: Hamilton 2009

Overview

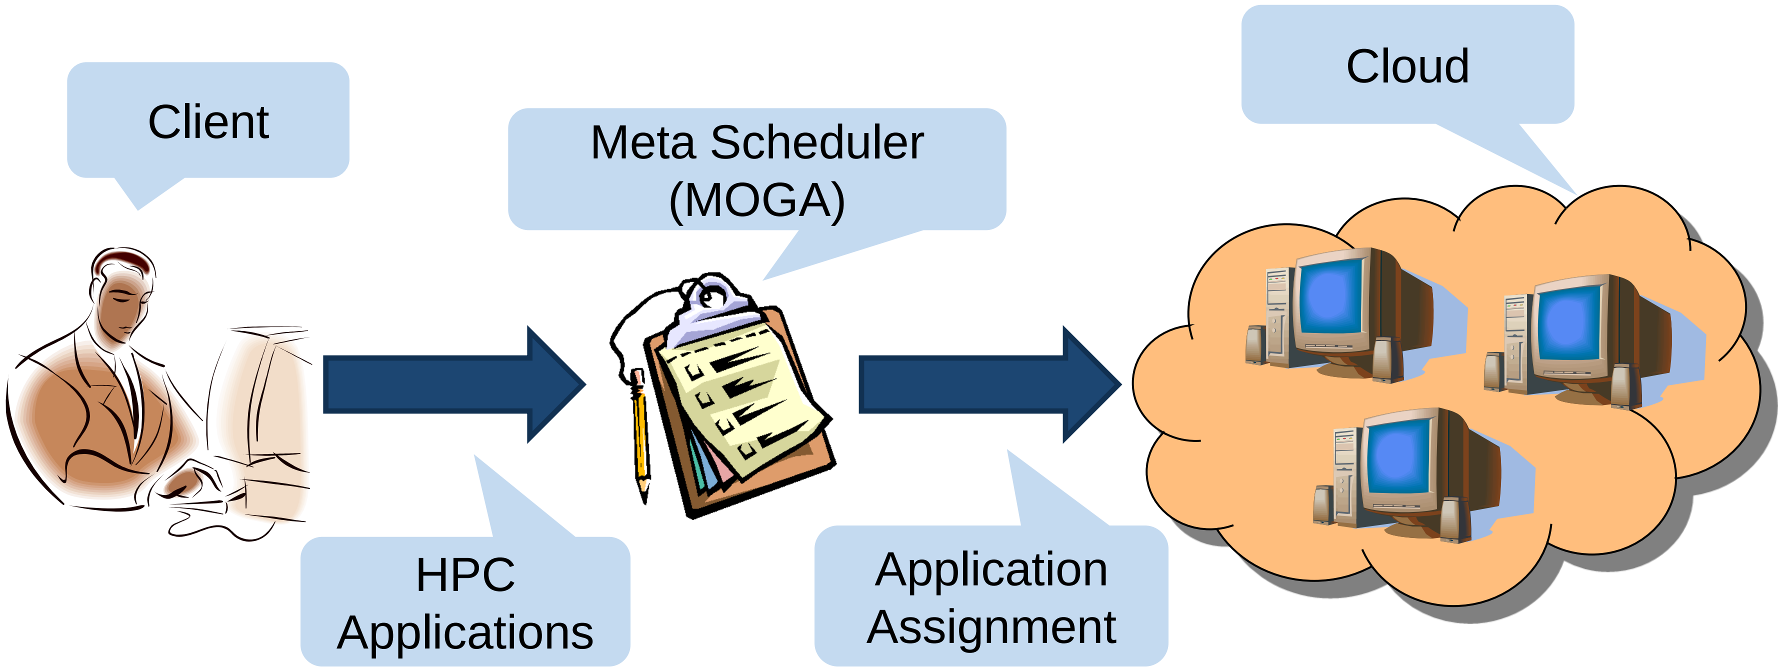
Client
Meta Scheduler (MOGA)
Cloud
HPC Applications
Application Assignment

Infrastructure as a Service (IAAS)

- The cloud provider deals with: the virtualization, server hardware, storage and networks.

- The company deals with: applications, runtimes, SOA integration, databases and server software.

  - ✓ Two-thirds model (Client, Provider).
  - ✓ Economic model based on the consumption of infrastructure.
  - ✓ Payment on demand (pay as you go).

Clouds Federation

- ✓ Geographically distributed cloud
- ✓ 3 continents (America, Europe et Oceania)
- ✓ **Open Cirrus** architecture

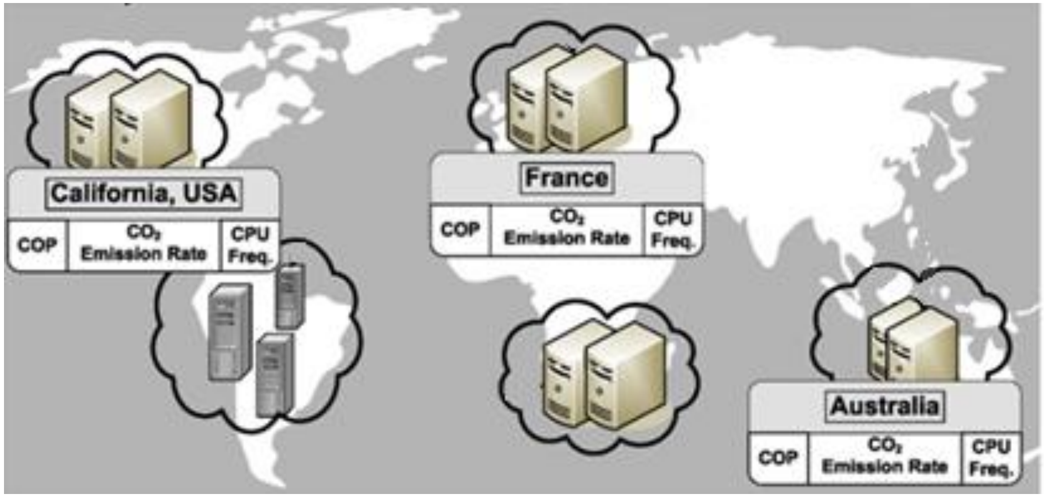

Location	COP Rate	Co2 Rate (kg/kW h)	Electricity price ($/kW h)	Cpu Alpha	Cpu Beta	Max Frequency (Ghz)	Opt Frequency (Ghz)	Nb processor
New York, USA	3.05288430624641	0.389	0.15	65	7.5	1.8	1.630324	2050
Pennsylvania, USA	1.69138630912639	0.574	0.09	75	5	1.8	1.8	2600
California, USA	2.19652511128224	0.275	0.13	60	60	2.4	0.793701	650
Ohio, USA	1.27017081994563	0.817	0.09	75	5.2	2.4	1.93201	540
North Carolina, USA	1.84324268393684	0.563	0.07	90	4.5	3.0	2.154435	600
Texas, USA	1.60832327774260	0.664	0.1	105	6.5	3.0	2.00639	350
France	0.915954867727123	0.083	0.17	90	4.0	3.2	2.240702	200
Australia	3.09966378393583	0.924	0.11	105	4.4	3.2	2.285084	250

CO2 rate ➔ Department Of Energy (DOE)
Electricity price ➔ Energy Information Administration (EIA)
Other parameters ➔ Statistics on other works (Wang et al) or synthetic data generation

Application model : HPC

- The meta-scheduler receives flow of requests (applications) with for each application j the following triple:

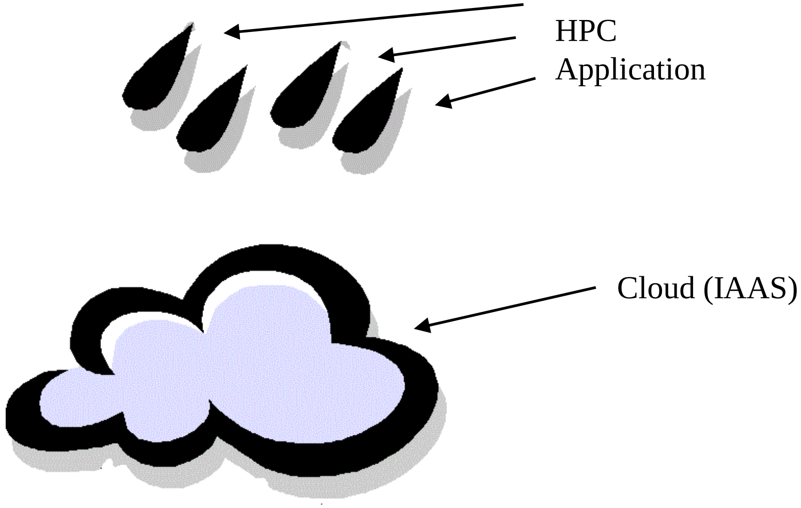

$$(e_j, n_j, d_j)$$

e_j The execution time of the application (booking)
n_j The number of reserved processors
d_j The deadline of the application

Meta-Scheduling

➢ Solution

- Optimal application assignment based on **metrics** and with **constraints** respect

➢ Tool

- Multi-objective genetic algorithm **(MO-GA)**

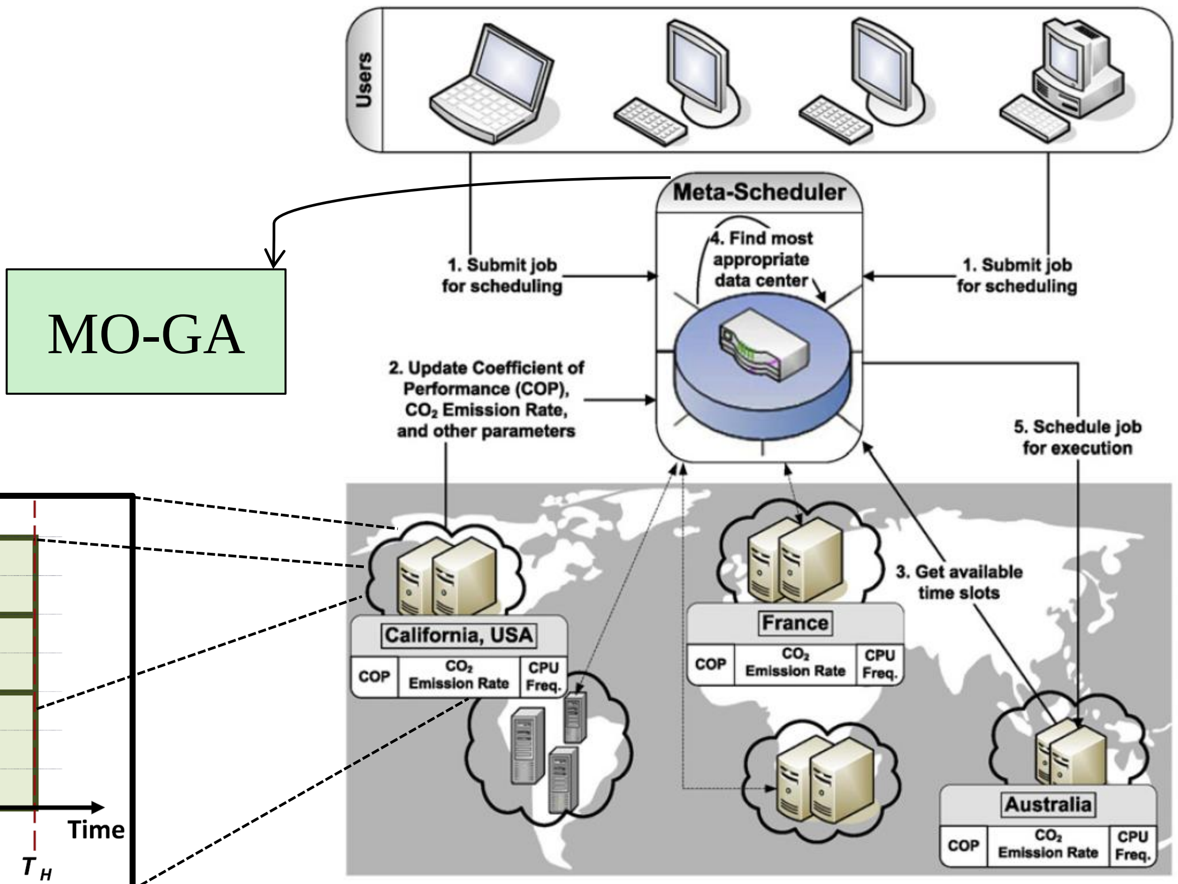

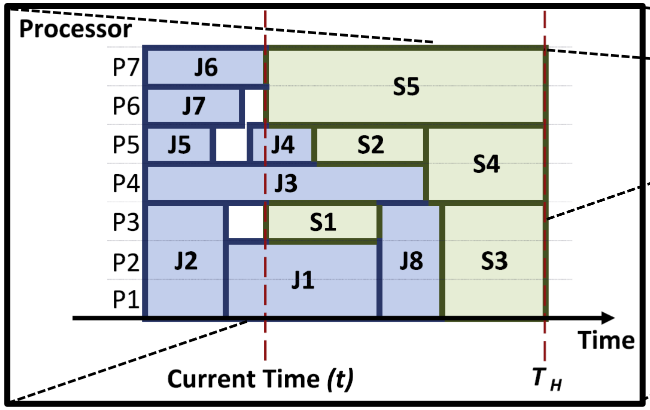

Source: S.K Garg et al (Environment-conscious scheduling of hpc applications on distributed cloud-oriented data centers)

State of the Art

Approach	Greenhouse gas/Energy consumption	HPC workload	Distributed data centers	Energy cost aware scheduling	Market-oriented scheduling	Genetic algorithm scheduling	Pareto optimization
Rizvandi et al.	yes	no	no	no	no	no	no
Lee and Zomaya.	yes	no	no	no	no	no	no
Tesauro et al.	yes	yes	no	no	no	no	no
Orgerie et al.	yes	yes	yes	no	no	no	no
Lee et al.	no	no	no	no	yes	no	no
Burge et al.	yes	yes	no	yes	no	no	no
Yu and Buyya.	no	no	yes	no	yes	yes	no
Garg et al.	no	no	yes	no	yes	yes	no
Garg et al.	yes	yes	yes	yes	yes	no	no
Our work	yes	yes	yes	yes	yes	yes	yes

Problematic

➢ Problem modeling

✓ Scheduling J applications on N clouds (**NP hard**) respecting the constraints…

- … Deadline (Strong constraint): no delays
- … Each application must be affected to one and only one cloud (Loosely Coupled Cloud Federation)
- … The cloud have to provide the total number of processors requested by an application without overriding the deadline

➢ Objectives

✓ Meet maximum client's requests …

- … by optimizing the three objectives (energy, CO2, profit)

Metrics (1)

➢Energy

- For each cloud **minimizing** the energy consumption
  - ✓ Energy needed for the computation (E^c)
  - ✓ Energy required for cooling the cloud (E^h)

$$E_{ij}^c = (\beta_i + \alpha_i f_{ij}^3) \times n_j \times e_{ji}$$

$$E^h = \frac{E^c}{COP} \qquad E_{total} = E^c + E^h$$

Source: Google

➢CO2

- For each cloud **minimizing** the amount of CO2 emission

$$CO_2 = E_{total} \times CO_2 Rate$$

Source: Google

Metrics (2)

➢Profit

➢ Optimizing the profit consists of **maximizing** the provider's earnings at each meta-scheduling

Price charged to the client — Electrical energy price

$$Profit_{ij} = n_j \times e_{ji} \times p^c - (p_i^e \times E_{ij})$$

where: n_j number of processors
e_{ji} execution time of the application
p^c client's price per hour
p_i^e electricity price
E_{ij} consumed energy

Source: Google

Algorithm

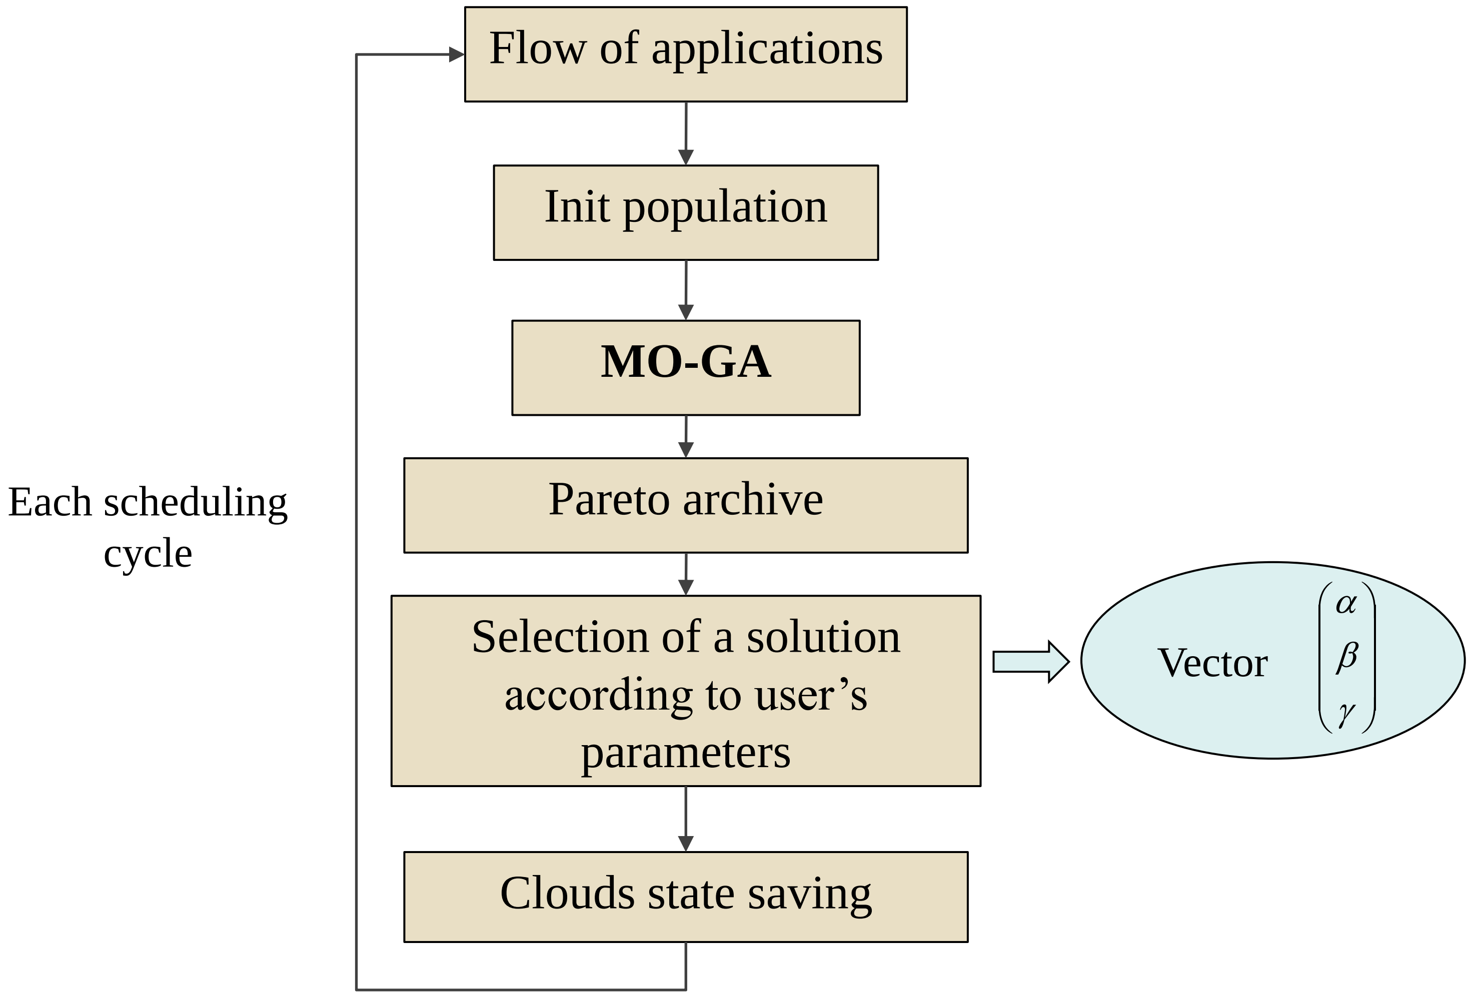

Flow of applications
Init population
MO-GA
Pareto archive
Selection of a solution according to user's parameters
Clouds state saving
Each scheduling cycle
Vector $\begin{pmatrix}\alpha\\ \beta\\ \gamma\end{pmatrix}$

Selection Mechanism

$$v_1 = \begin{pmatrix} \alpha = 0 & f_1 \\ \beta = 1 & f_2 \end{pmatrix}$$

$$v_2 = \begin{pmatrix} \alpha = \frac{\sqrt{2}}{2} & f_1 \\ \beta = \frac{\sqrt{2}}{2} & f_2 \end{pmatrix}$$

$$v_3 = \begin{pmatrix} \alpha = 1 & f_1 \\ \beta = 0 & f_2 \end{pmatrix}$$

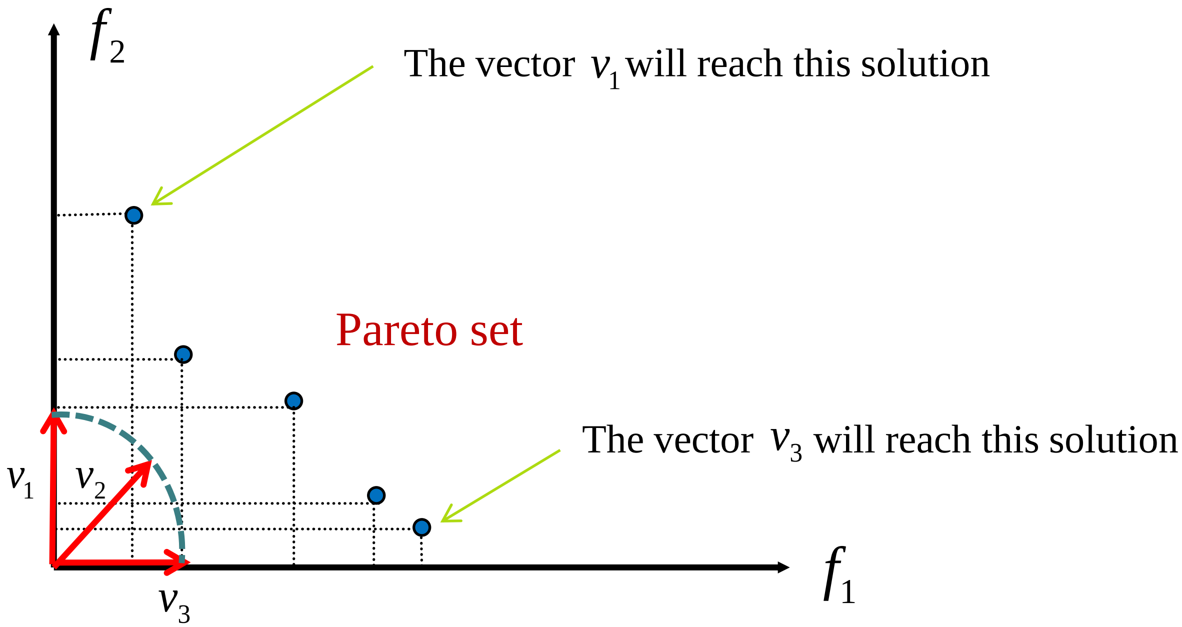

Cloud State Saving

Cloud during the previous state

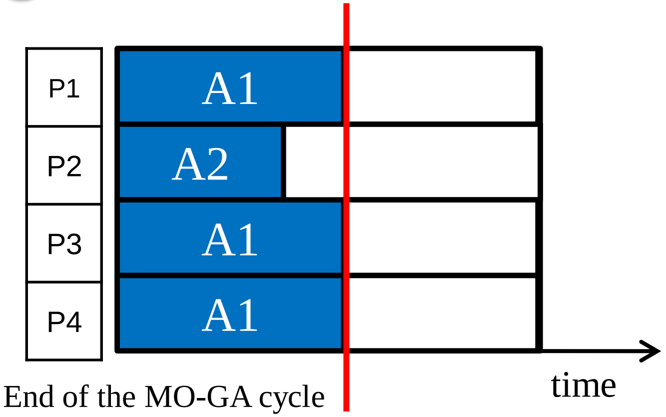

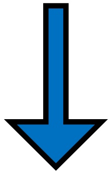

Cloud after updating the state

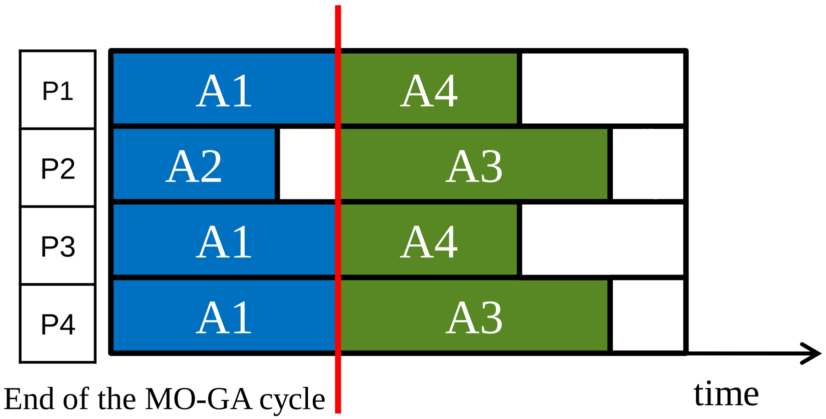

Experimentation

Experiment parameters

- 4 different arrival rates (Low, Medium, High, Very high). Each move from an arrival rate to another represents 10 times more applications arrival during the same period of time.
- Different selection policies for the transition state (Vector).
- Realistic experiments drawn from the archives of Feitelson's Parallel Workload.
  (http://www.cs.huji.ac.il/labs/parallel/workload/l_llnl_thunder/index.html)
- Total number of applications: **128662** (5 months requests).
- The deadlines are generated using a bimodal distribution (80% ”Low Urgency“, 20% ”High Urgency”).
- A LU application has on average a three times larger deadline than HU application.

Comparison

- Problem not treated in the literature with a **multi-objective Pareto approach**.
- Comparison done with a heuristic that maximizes the number of scheduled applications.

Results (1)

- Profit Oriented Selection Vector Results

Vector settings: Arrival rate \ Value for each criterion	Profit: Energy (kW h)	Profit: CO_2 (Kg)	Profit: Profit ($)	Profit: Failed applications	Heuristic: Energy (kW h)	Heuristic: CO_2 (Kg)	Heuristic: Profit ($)	Heuristic: Failed applications
Low	2.067e+06	1.168e+06	**4.803e+06**	2325	1.807e+06	709039	**4.726e+06**	2270
Medium	2.033e+06	1.095e+06	**4.764e+06**	2723	1.833e+06	727748	**4.696e+06**	2635
High	2.214e+06	1.180e+06	**4.681e+06**	3657	2.026e+06	939550	**4.677e+06**	3357
Very high	2.113e+06	1.039e+06	**4.647e+06**	6522	2.036e+06	972912	**4.660e+06**	5168

- None Specific objective oriented selection Vector Results

Vector settings: Arrival rate \ Value for each criterion	Average: Energy (kW h)	Average: CO_2 (Kg)	Average: Profit ($)	Average: Failed applications	Heuristic: Energy (kW h)	Heuristic: CO_2 (Kg)	Heuristic: Profit ($)	Heuristic: Failed applications
Low	1.807e+06	709050	4.726e+06	**2270**	1.807e+06	709039	4.726e+06	2270
Medium	1.813e+06	717417	4.700e+06	**2588**	1.833e+06	727748	4.696e+06	2635
High	1.931e+06	837531	4.683e+06	**3045**	2.026e+06	939550	4.677e+06	3357
Very high	2.036e+06	972912	4.660e+06	**5168**	2.036e+06	972912	4.660e+06	5168

Results (2)

Scheduling Success Rate

- The average application scheduling success rate through the different arrivals is **97.46%**.

- The improvement compared to the heuristic which maximizes the number of scheduled application is **2.67%**.

Criterions Improvement

- The total energy consumption of the distributed cloud have been reduced by up to **4.66%.**

- The green house gas emissions by up to **10.85%.**

- The provider's profit increased by up to **1.62%.**

Inria

Perspectives

➢ Proposing a better energy reduction by using the Dynamic Voltage Scaling (DVS) within each data center.

➢ Introducing delays in the approach with a new pricing model with penalties.

➢ Integrating dynamicity in the meta-scheduler, for a real time applications' reallocation.

➢ Using EGI as a cloud federation in our next model and benefit from the geographical dispersion offered by **Europe**.
… a part in a Cloud Distribution ex: StratusLab.

Source: EGI

StratusLab & MO-GA

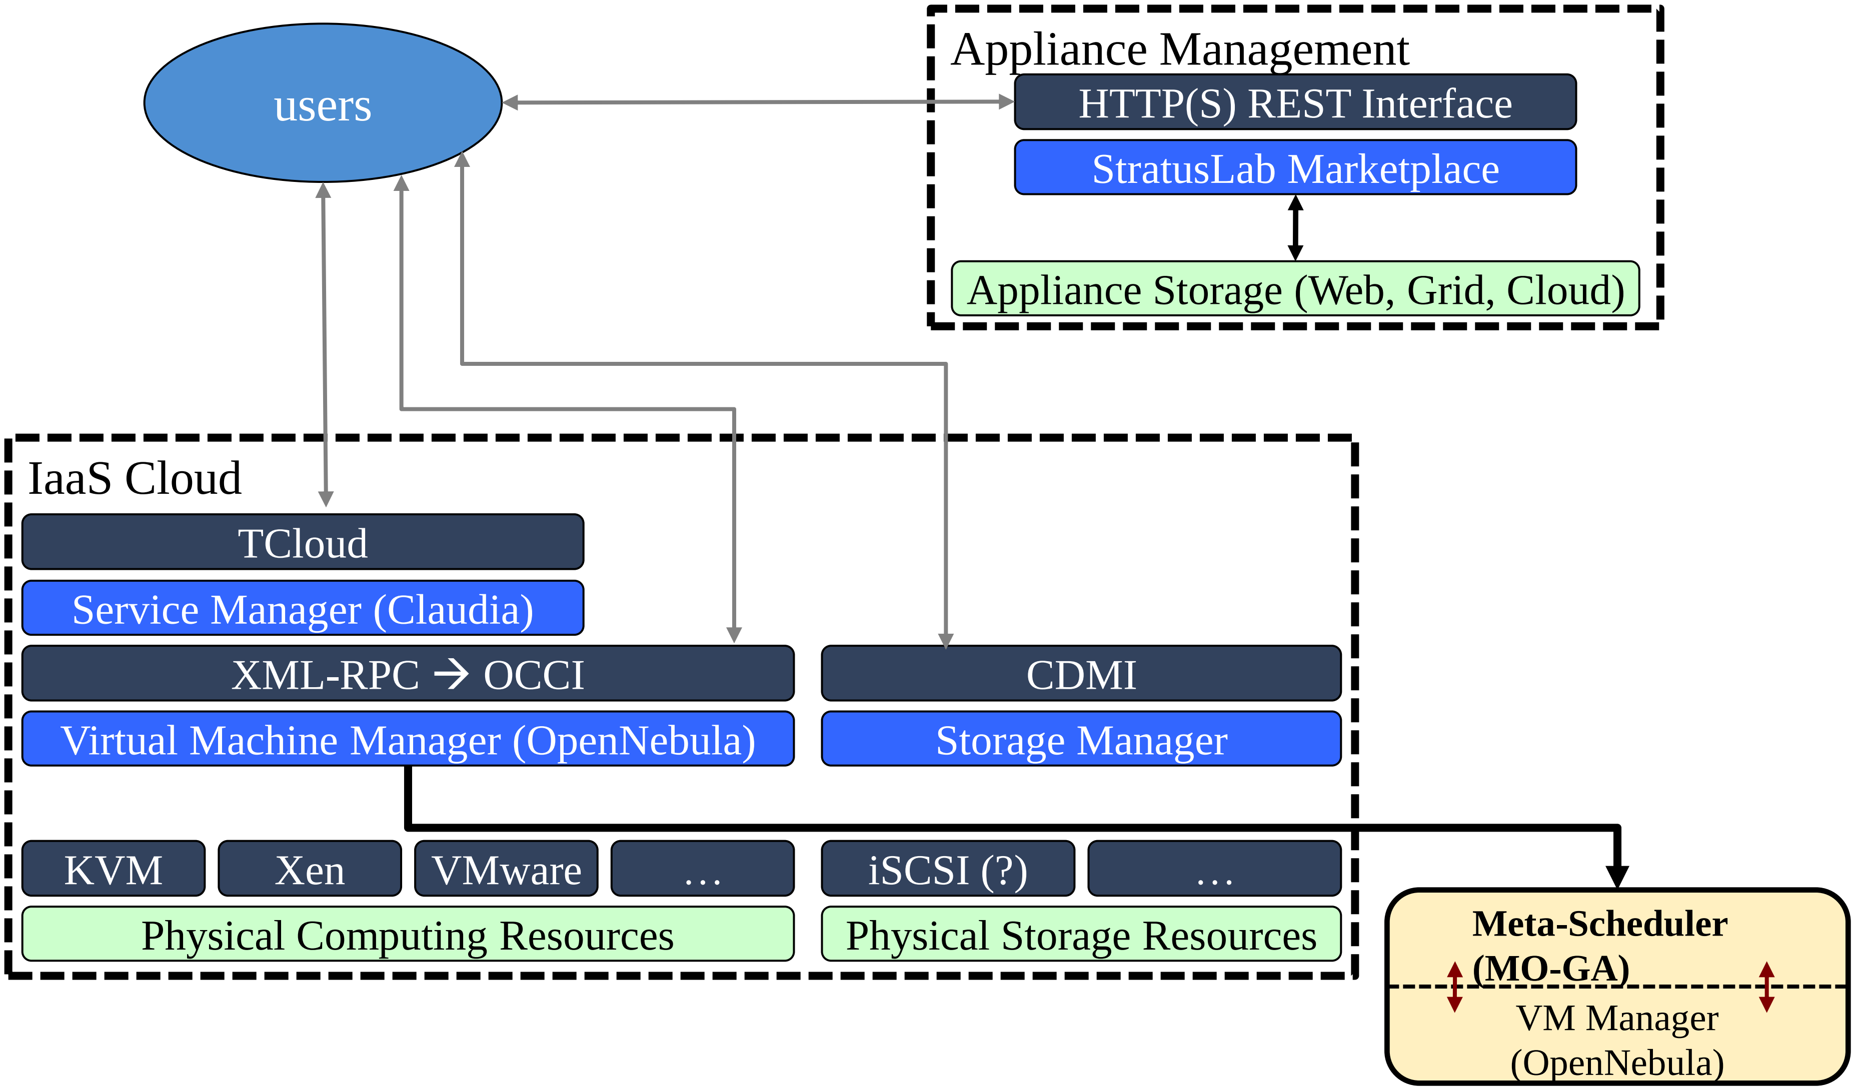

Conclusion

- Geographical dispersion helps to optimize significantly energy consumption and CO_2 emission rates.

- The multi-objective genetic algorithm allows the exploitation of a large number of possible meta-scheduling (solutions) which leads to better results.

- MO-GA gives a higher rate of scheduled application compared with the heuristic.

- MO-GA allows an improvement of all the treated objectives in comparison with the heuristic.

Thank you for your attention

Questions ?