Power-aware Multi-DataCenter Management using Machine Learning

Josep Ll. Berral, Ricard Gavaldà, Jordi Torres

2nd Intl. Workshop on Power-aware Algorithms, Systems, and Architectures 2013

October 1st – Lyon (France)

Context: Energy, Quality of Service and Self-Management

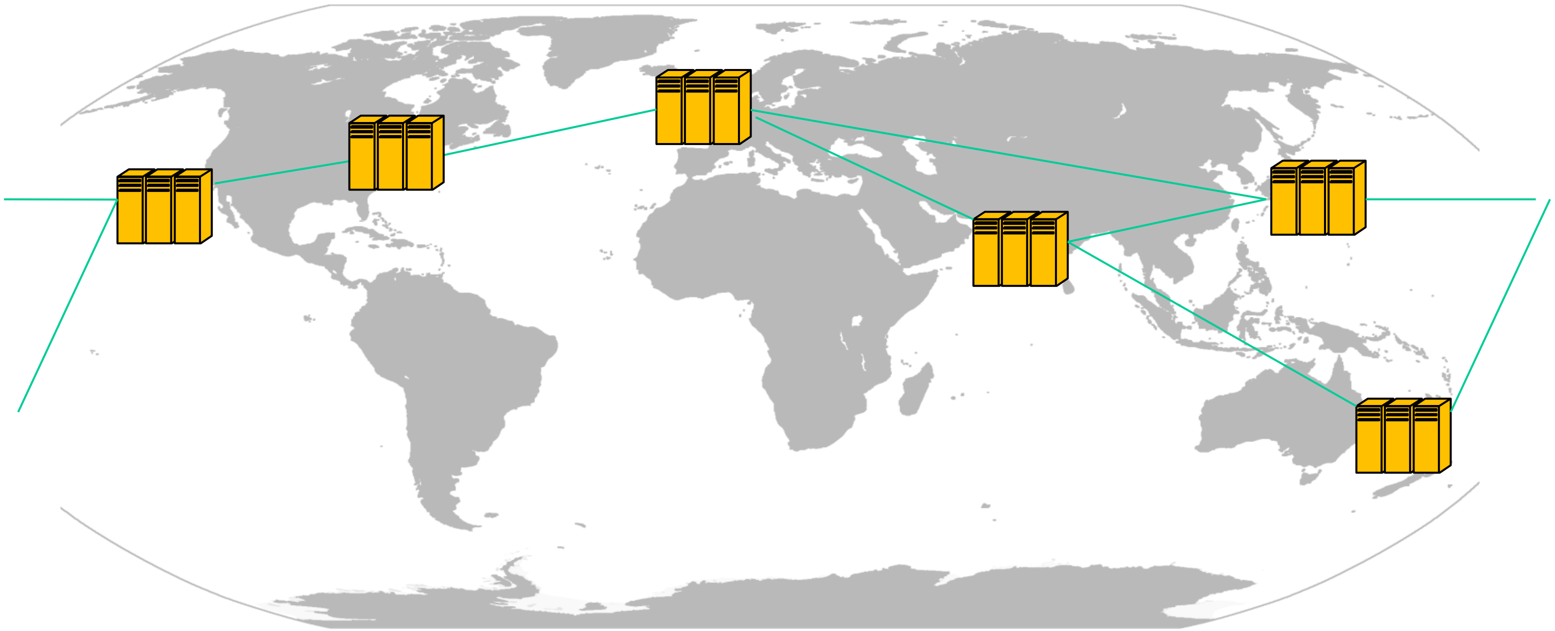

- Scenario: Multi-Datacenter Network
  - Achieve allocation of virtualized web-services
  - … keeping good Quality of Service
  - … reducing energy costs
  - … and doing this “automatically”

Context: Autonomic Computing and Machine Learning

- Keywords:
  - Autonomic Computing (AC): Automation of management
  - Machine Learning (ML): Learning patterns and predict them

- Applying AC to energy control:
  1. Self-management must include energy policies
  2. Optimization mechanisms are becoming more complex
  3. Decision makers can be improved through adaption over time

- Modeling and prediction:
  - Obtain a predictive model from the system from the past
  - ...using minimal expert knowledge

Introduction

- Energy Saving in Cloud Self-management:
  - Apply the well-known consolidation strategy

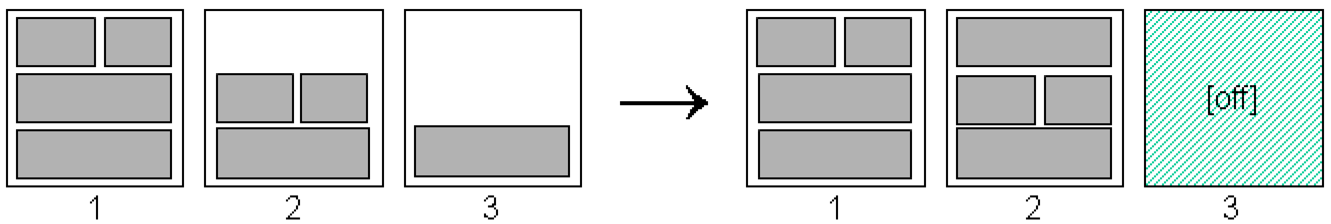

- Challenges:
  - How do we consolidate? Optimal place for a job/VM
  - How much resources used? Required resources for the job/VM
  - Resulting QoS / Energy cost in the new placement?

- Contributions:
  - Apply ML to learn about resource performance
  - On a mathematical model for a multi-DC (Benefit-Cost optimization)
  - Also include elements of geographical location (and their properties)

Multi-DataCenter Business Model

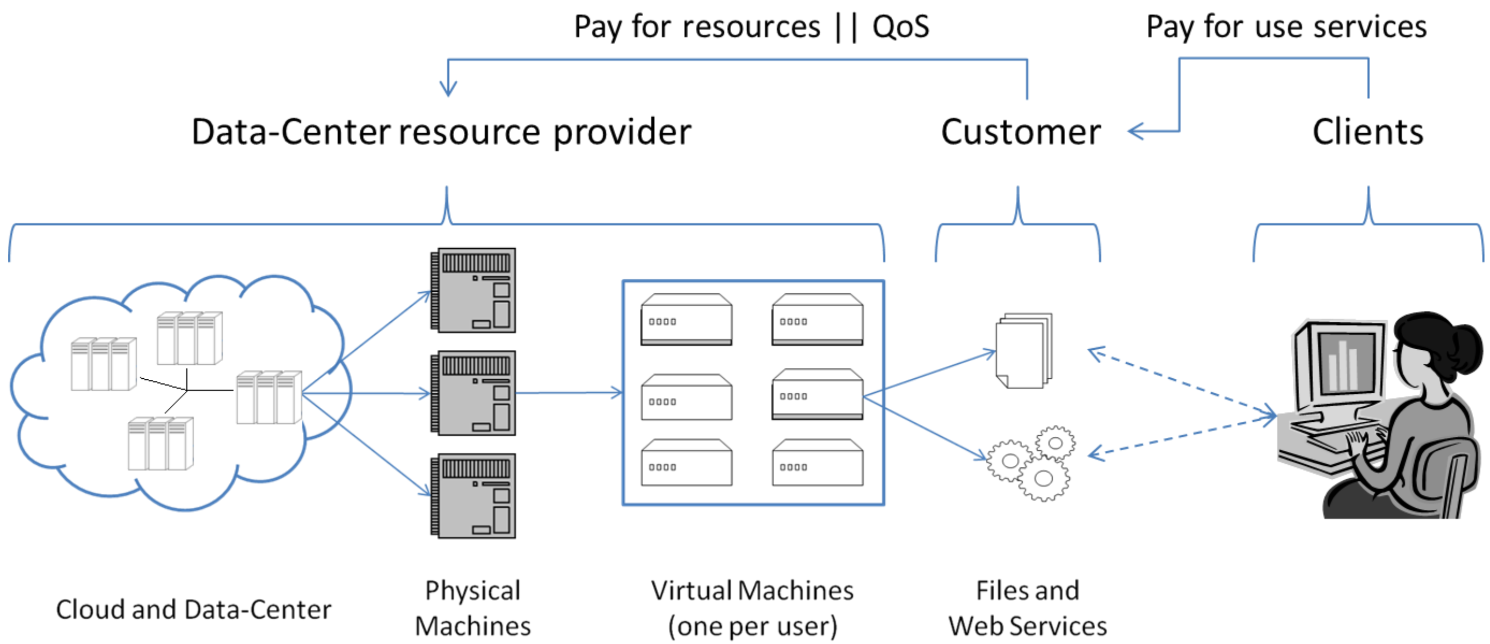

- Specific Case of Study:
  - Transactional jobs, Quality of Service (i.e. "Response Time")
- Problem:
  - As a provider: Schedule properly VMs to PMs

Multi-DataCenter Scenario

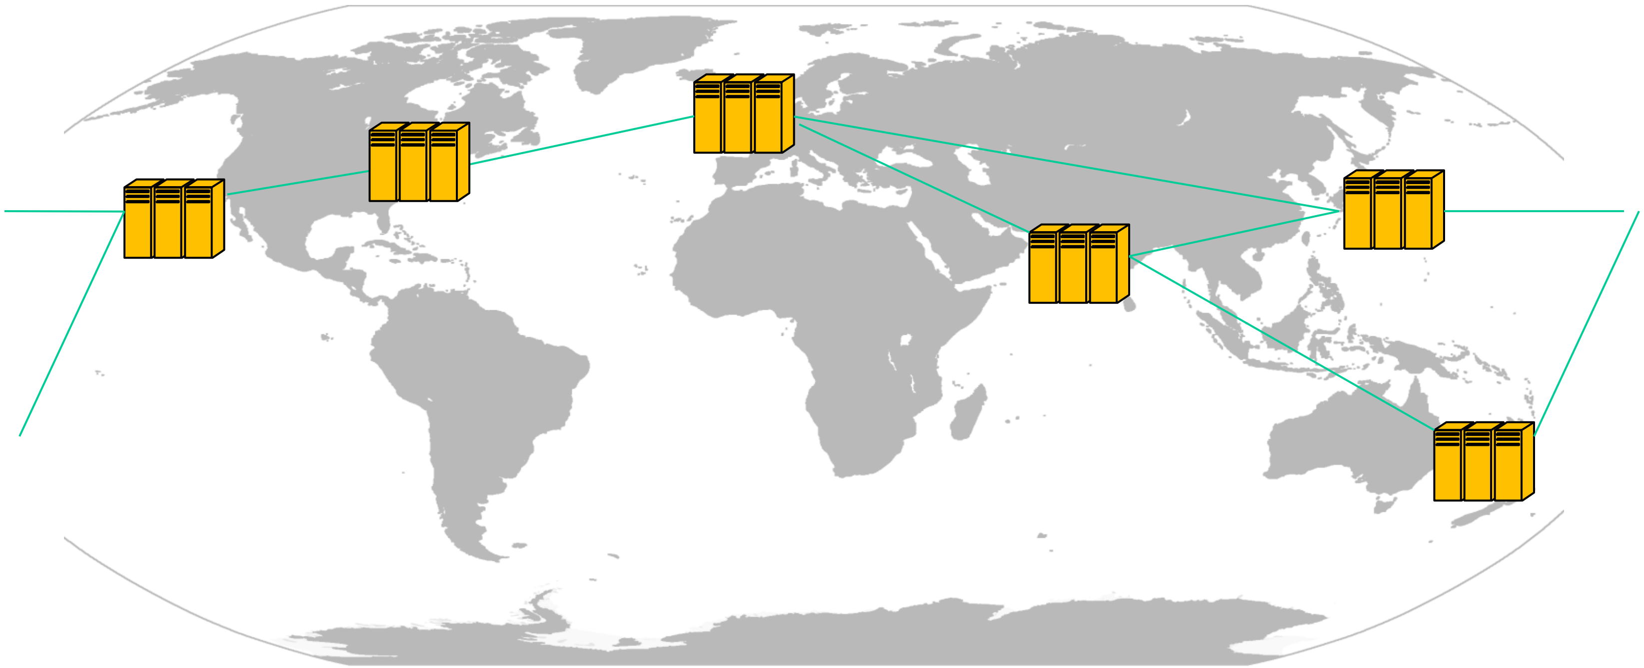

- Network of DataCenters
  - Each location has its own energy prices
  - Each client connects to our DC network through the closest DC
  - Each VM may have clients from around the world
  - Each location clients have different “timetables”

Problem parts…

1. Model the multi-datacenter
   – Create a mathematical model to represent the multi-DataCenter

2. Fit the model to observations:
   – Relevant variables only available *a posteriori*
   – ML creates a model from past examples

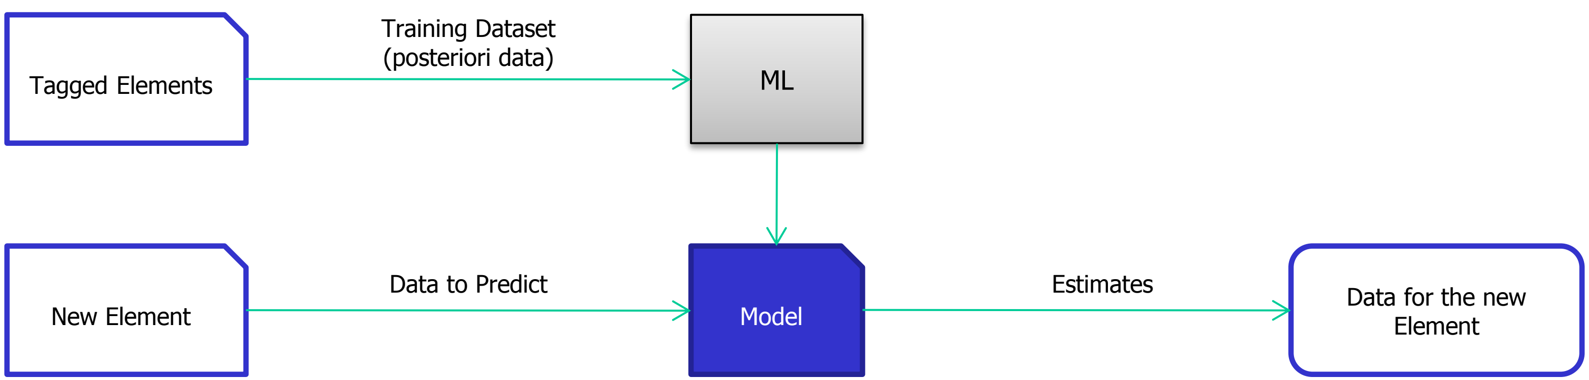

3. Solving the optimization problem

Modeling the Multi-DataCenter

- Mathematical Model: Find VMs $\rightarrow$ (hosts × resources)
  - Profit = Benefits for running VMs – QoS penalties – power costs
  - Outputs: Schedule optimizing profit
  - Constraints: maintaining the consistence of M-DC and operations
- Quality of Service
  - $RT = RT_{process} + RT_{transport}$ ("Latencies")
- Subject to:
  - VM requirements, depending on load
  - Power functions, depending on resources and locations
  - Migration penalties, on distances and VM volumes
  - QoS, depending on resource competence and client distance

Modeling the Multi-DataCenter

- Mathematical Model: Find VMs $\rightarrow$ (hosts × resources)
  - Profit = Benefits for running VMs – QoS penalties – power costs
  - Outputs: Schedule optimizing profit
  - Constraints: maintaining the consistence of M-DC and operations

- Quality of Service
  - $RT = RT_{process} + RT_{transport}$ ("Latencies")

Machine Learning

- Subject to:
  - VM requirements, depending on load
  - Power functions, depending on resources and locations
  - Migration penalties, on distances and VM volumes
  - QoS, depending on resource competence and client distance

Learning and Prediction

- Applying modeling and prediction
  - How much CPU/Mem/IO… will each VM demand?
  - How good will each VM behave?

- Learning on the given scenario
  - Apply ML modeling techniques for VM CPU/MEM/IO
  - Also: learn PM CPU aggregate
  - Also: learn QoS as “RT” or “SLA”

- Benefits:
  - When changing machines, we only need to re-learn ML models
  - We discover the bottlenecks of the system

	ML Method	Correl.	MAE	Err-StDev	Train/Val	Date Range
Predict VM CPU	M5P ($M = 4$)	0.854	$4.41\%_{CPU}$	$4.03\%_{CPU}$	959/648	$[0, 400]$ $\%_{CPU}$
Predict VM MEM	Linear Reg.	0.994	26.85 MB	93.30 MB	959/1324	$[256, 1024]$ MB
Predict VM IN	M5P ($M = 2$)	0.804	1.77 KB	4.01 KB	319/108	$[0, 33]$ KB
Predict VM OUT	M5P ($M = 2$)	0.777	25.55 KB	22.06 KB	319/108	$[0, 141]$ KB
Predict PM CPU	M5P ($M = 4$)	0.909	$14.45\%_{CPU}$	$7.70\%_{CPU}$	477/95	$[25, 400]$ $\%_{CPU}$
Predict VM RT	M5P ($M = 4$)	0.865	0.234 s	1.279 s	1887/364	$[0, 19.35]$ s
Predict VM SLA	K-NN ($K = 4$)	0.985	0.0611	0.0815	1887/364	$[0.0, 1.0]$

Experiments

- Intra-DataCenter comparatives
  - Using approximate algorithms (ordered best-fit):

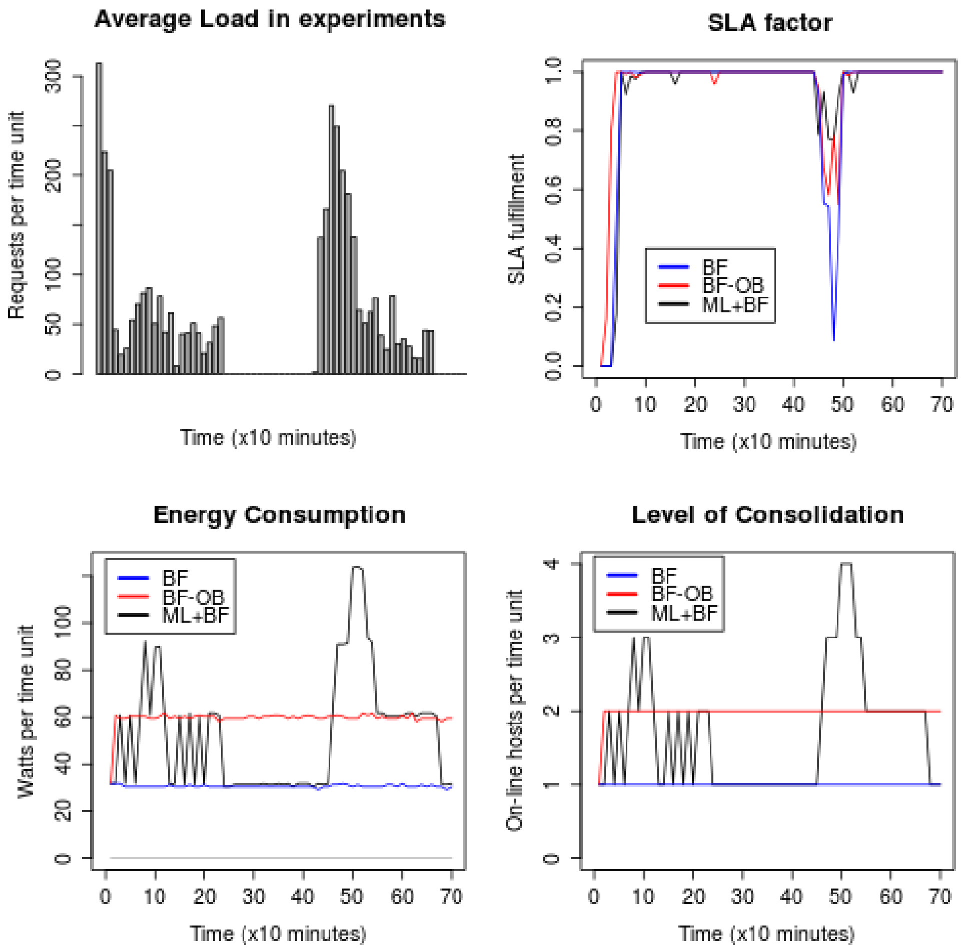

Experiments

- Inter-DataCenter results
  - In a dynamic context, energy savings may increase when consolidating spare VMs
  - Average SLA increases when migration costs are smaller than benefit improvements
  - When no load, VMs are sent to cheapest place to stay parked
  - ML models detect QoS violations better than no ML

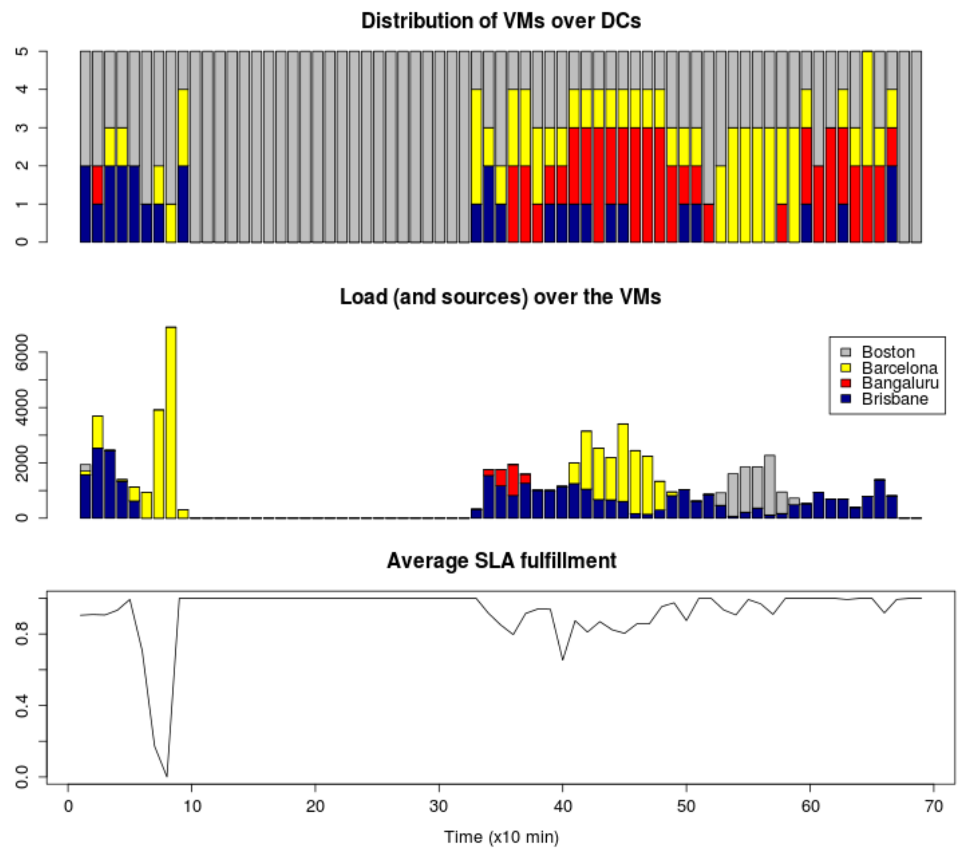

Energy/QoS/Load Trade-offs

- Trade-off between energy consumption and SLA (QoS)

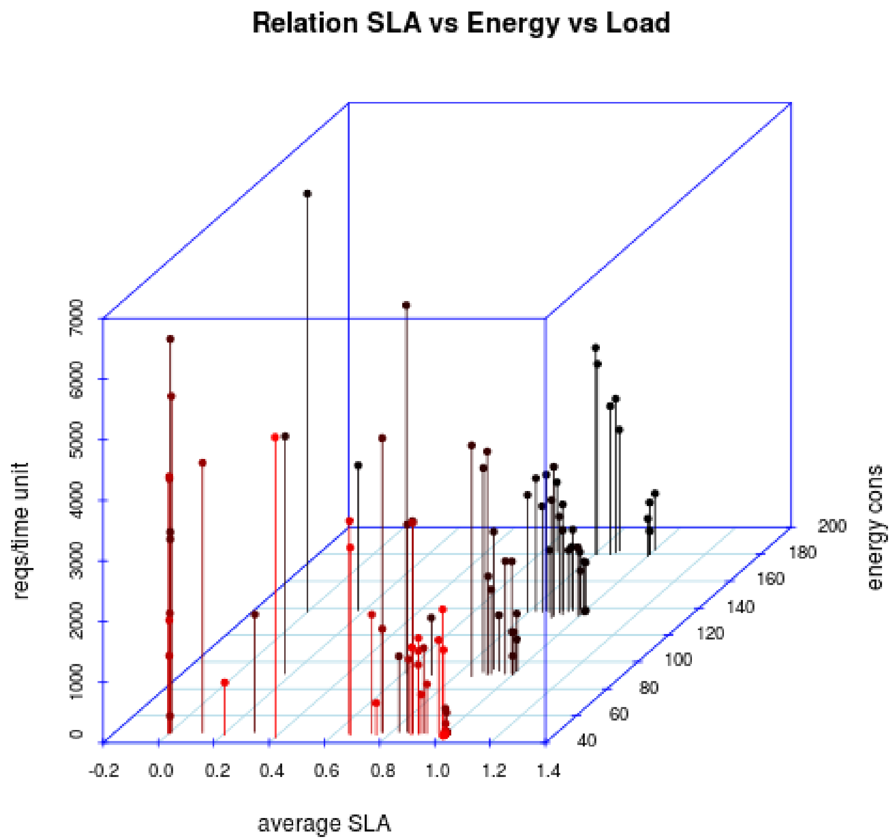

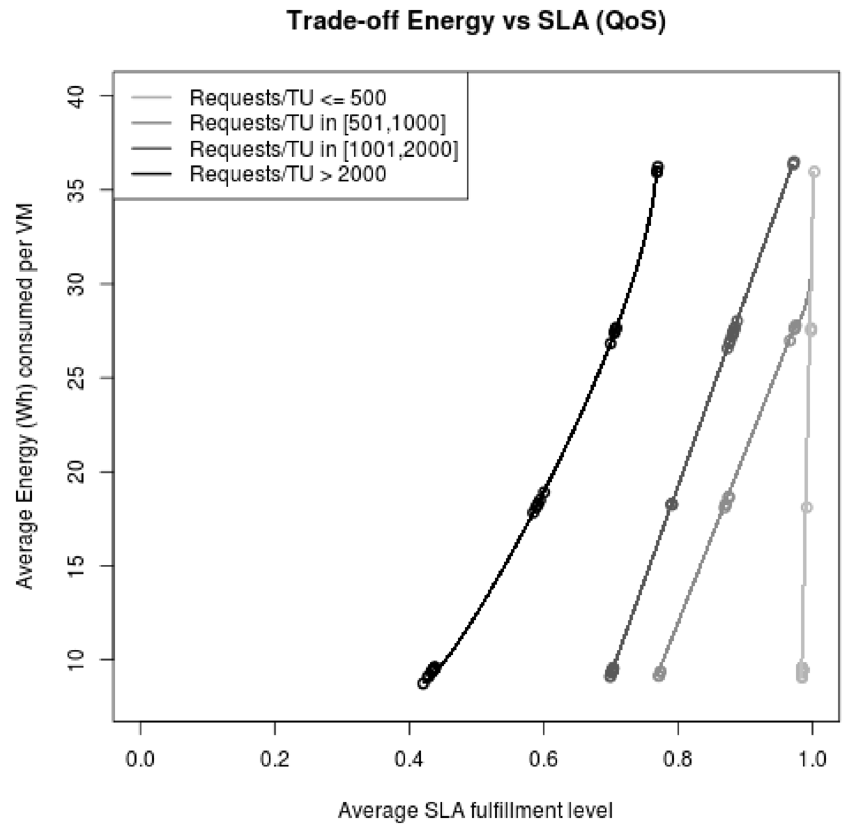

Summary

- Focus the "VMs × PMs" allocation problem :
  - With mathematical modeling on multi-datacenter systems
  - Focused on energy consumption and quality of service
  - Usage of automatic modeling through machine learning
- Contributions:
  - Introduce localization variables to a DC management model
  - Studied learning models on different kind of machines and views of QoS
  - Trade-off between SLA fulfillment and energy for transactional jobs
- Learning and Experimentation Results
  - When having different energy prices, de-location becomes a good option
- Future work:
  - Study new relevant variables to the multi-DC model, and other kind of jobs and web-services

Thank you for your attention

Josep Ll. Berral, Ricard Gavaldà, Jordi Torres

berral@ac.upc.edu

(also jlberral@lsi.upc.edu)