An Energy Aware Framework for Virtual Machine Placement in Cloud Federated Data Centres

Corentin Dupont

Authors: Corentin Dupont (Create-Net); Giovanni Giuliani (HP Italy); Fabien Hermenier (INRIA); Thomas Schulze (Uni Mannheim); Andrey Somov (Create-Net)

"FIT4Green expects to provide at least 20% saving in direct server and network devices energy consumption (with respect to a traditionally managed data centre) and induce an additional 30% saving due to reduced cooling needs"

- ✓ **FIT4Green** seeks energy saving policies for DCs, enhancing the effects inside a federation by an aggressive strategy for reducing the energy consumption in ICT
- ✓ We aim at reducing cost for companies ➔ Strengthening competitive position
- ✓ **FIT4Green** needs to be DC framework agnostic:
  - Demonstrated in Cloud computing, Traditional computing, Super computing and Networking

FIT4 GREEN

Table of Contents

Introduction

The policies seek to:

Consolidate application/services and turn unused servers off.

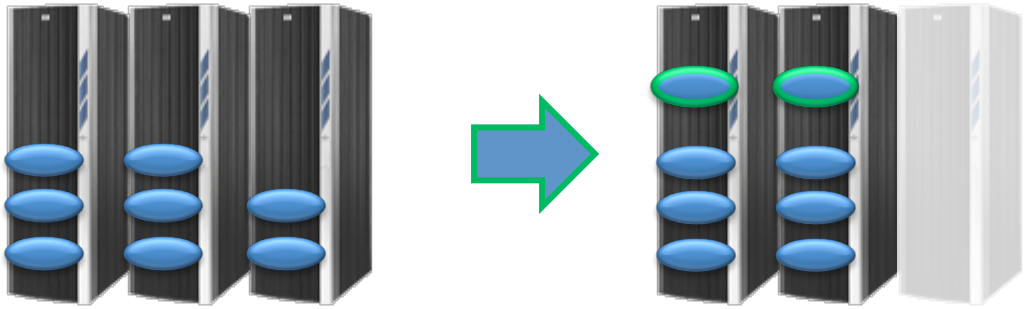

Relocate application/services to efficient servers

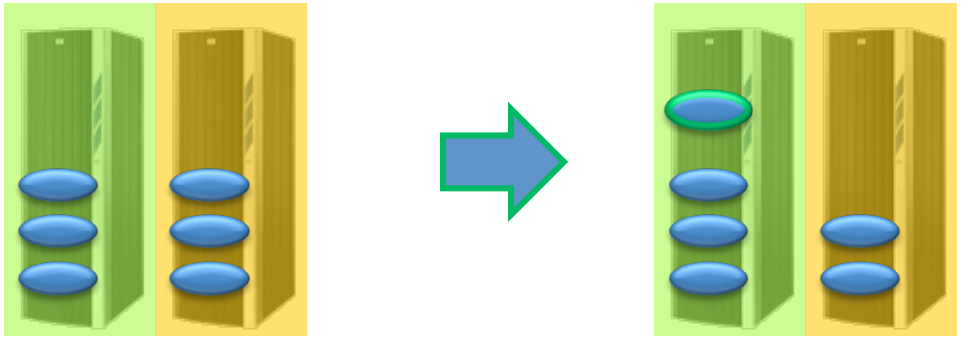

The strategies are ranked through their Energy KPIs

Single allocation

Find the most energy efficient and suitable resource for a **new Workload**.

Global optimization

Rearrange the resources in a way that saves maximum amount of energy or carbon emission.

Requirements

- Flexibility, extensibility
- Deep exploration of the search space

Abstracting out the constraints

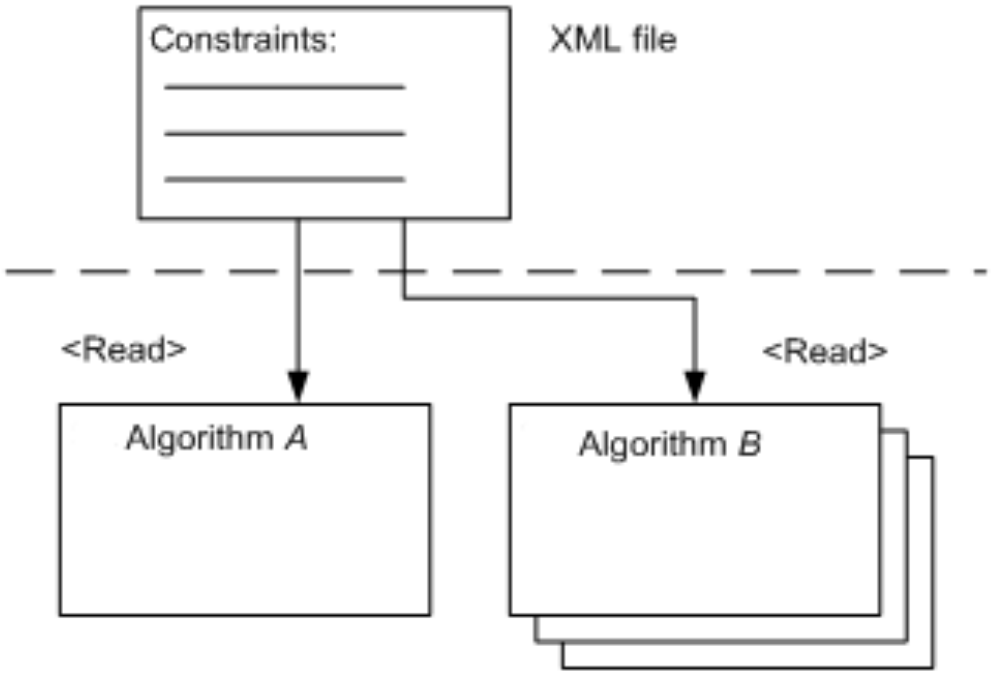

Framework Design

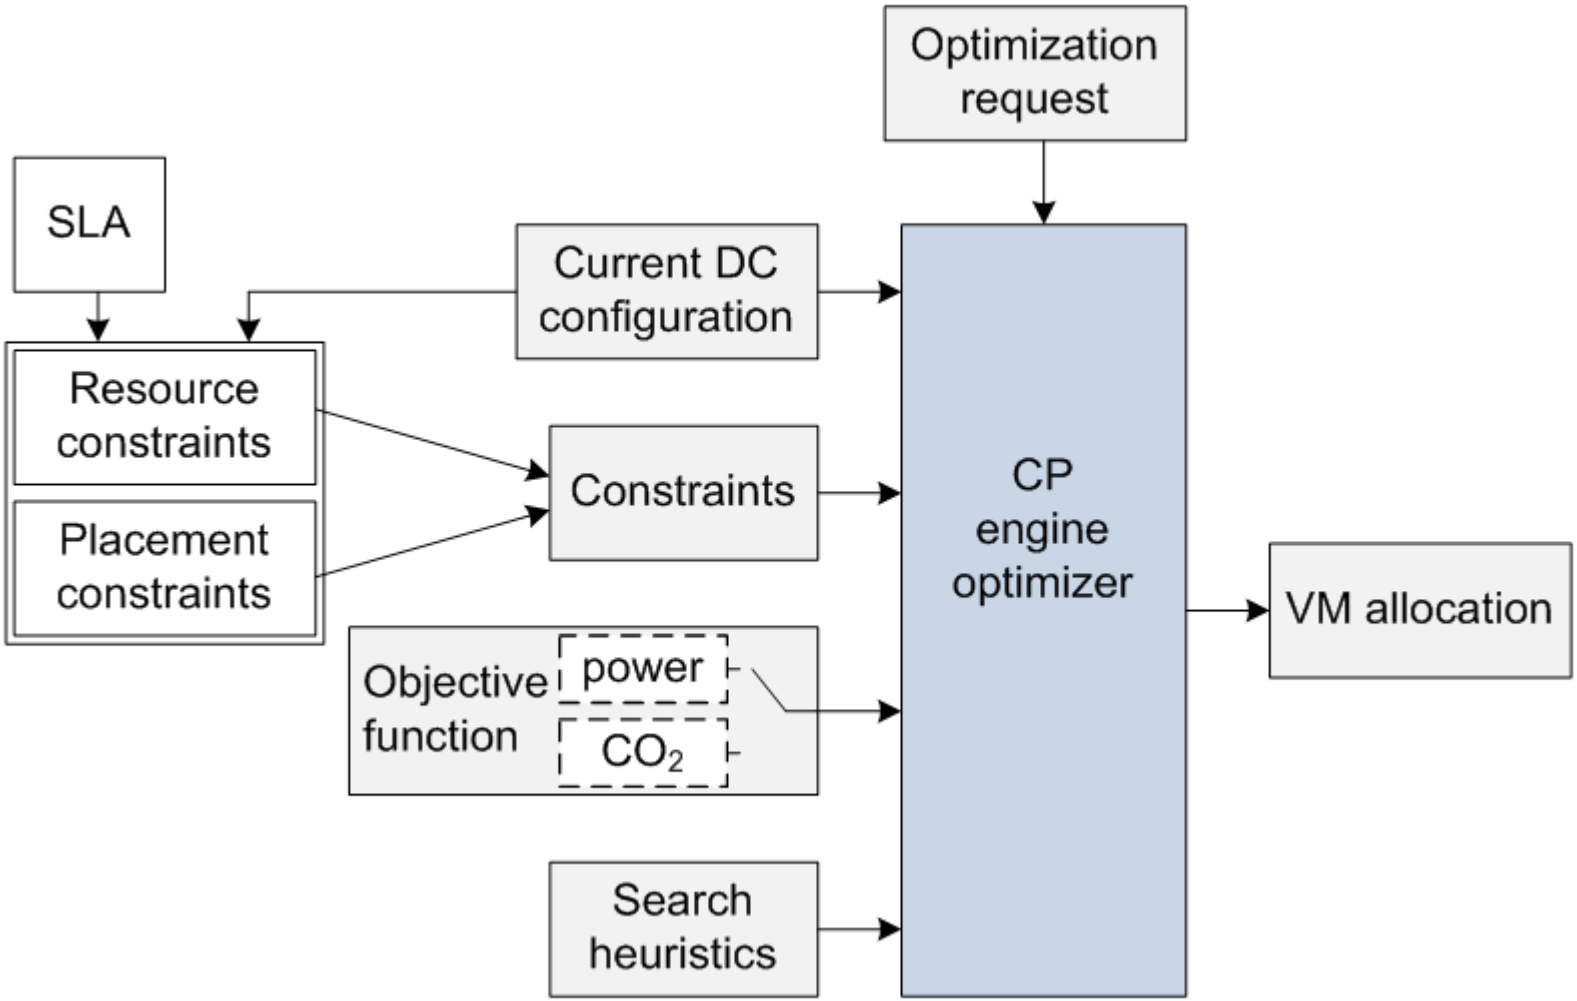

SLA constraints flow

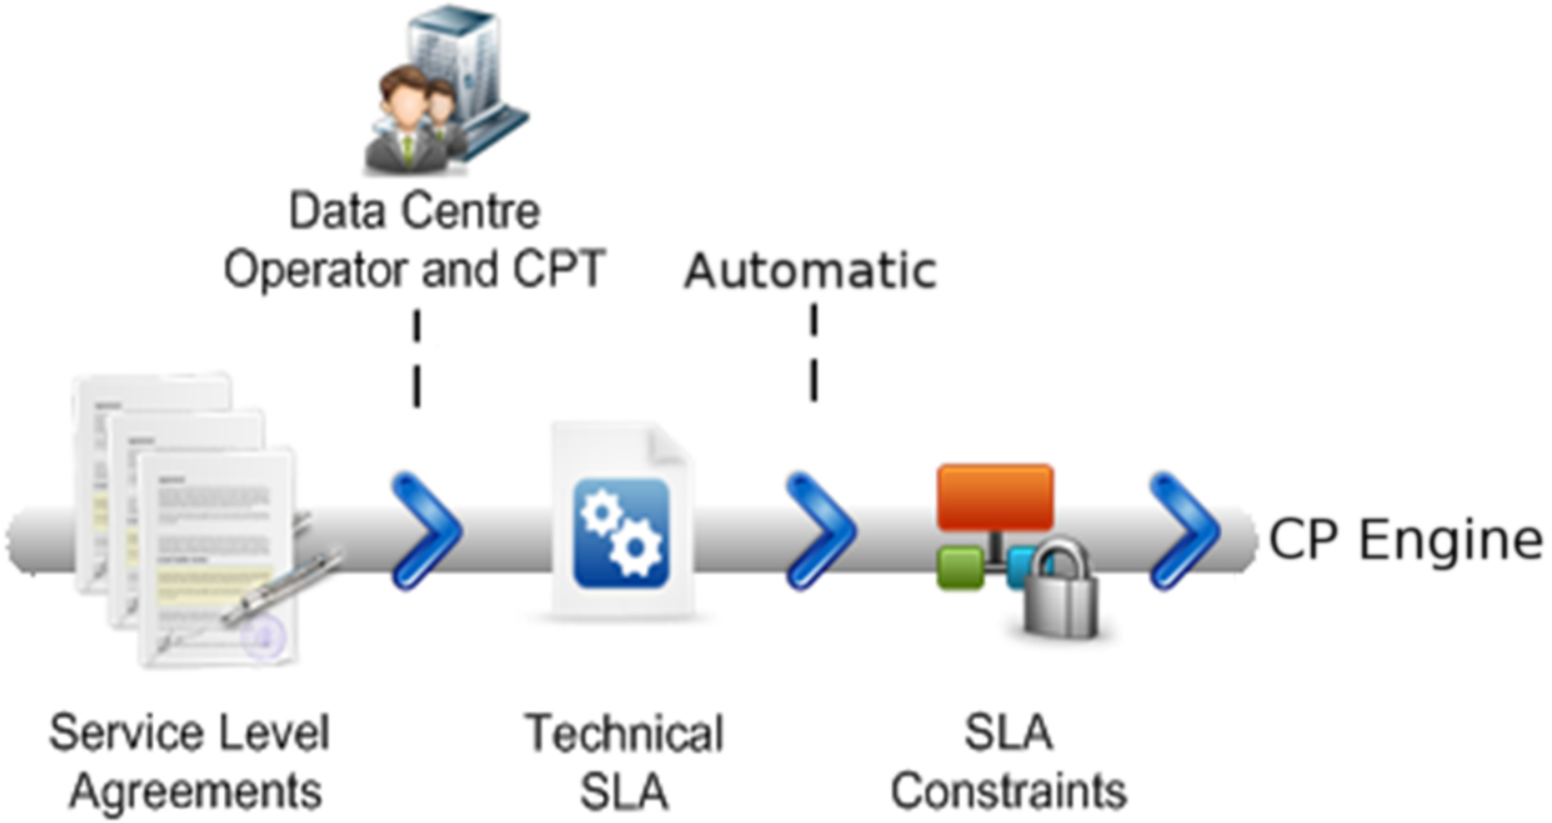

SLA CONSTRAINTS

SLA constraints examples

Category	Constraint	Approach	LoC
Hardware	HDD	Choco + ext. Entropy	121+(25)
	CPUCores	Entropy (‘fence’)	0+(25)
	CPUFreq	Entropy (‘fence’)	0+(25)
	RAM	Choco + ext. Entropy	123+(25)
	GPUCores	Entropy (‘fence’)	0+(25)
	GPUFreq	Entropy (‘fence’)	0+(47)
	RAIDLevel	Entropy (‘fence’)	0+(47)
QoS	MaxCPULoad	Choco + ext. Entropy	90+(25)
	MaxVLoadPerCore	Choco + ext. Entropy	109+(25)
	MaxVCPUPerCore	Choco + ext. Entropy	124+(25)
	Bandwidth	Entropy (‘fence’)	0+(49)
	MaxVMperServer	Entropy (‘capacity’)	0+(25)
Availability	PlannedOutages	Choco + ext. Entropy	Future Work
	Availability	Choco + ext. Entropy	Future Work
Additional Metrics	Dedicated Server	Entropy (‘capacity’)	0 + (25)
	Access	Entropy (‘fence’)	0 + (25)

POWER OBJECTIVE MODEL

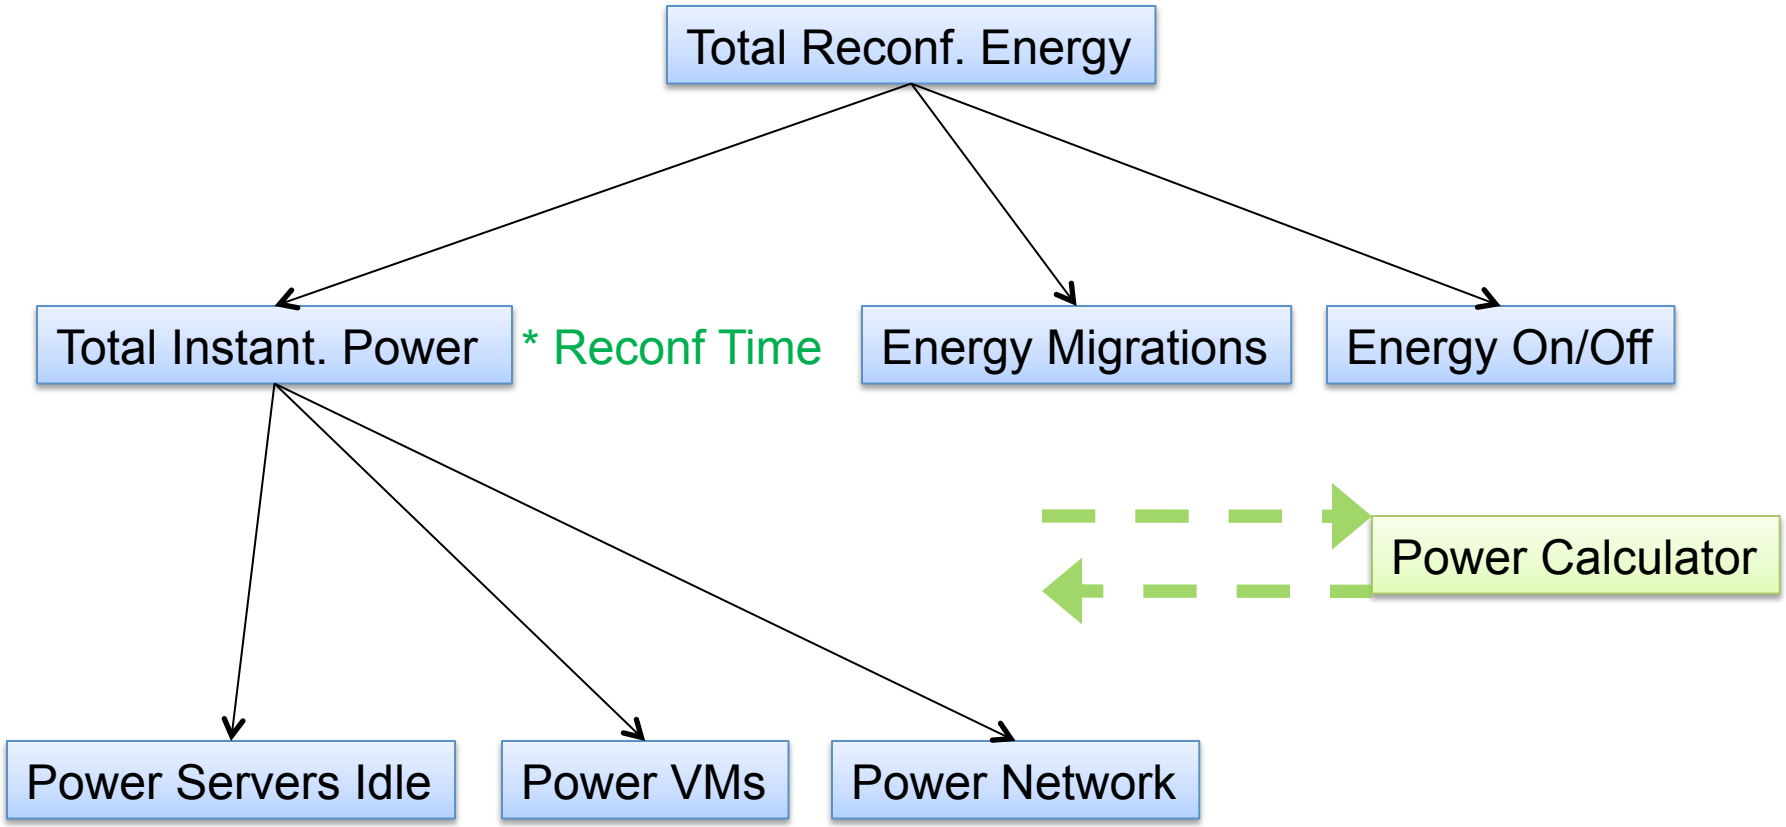

HEURISTICS

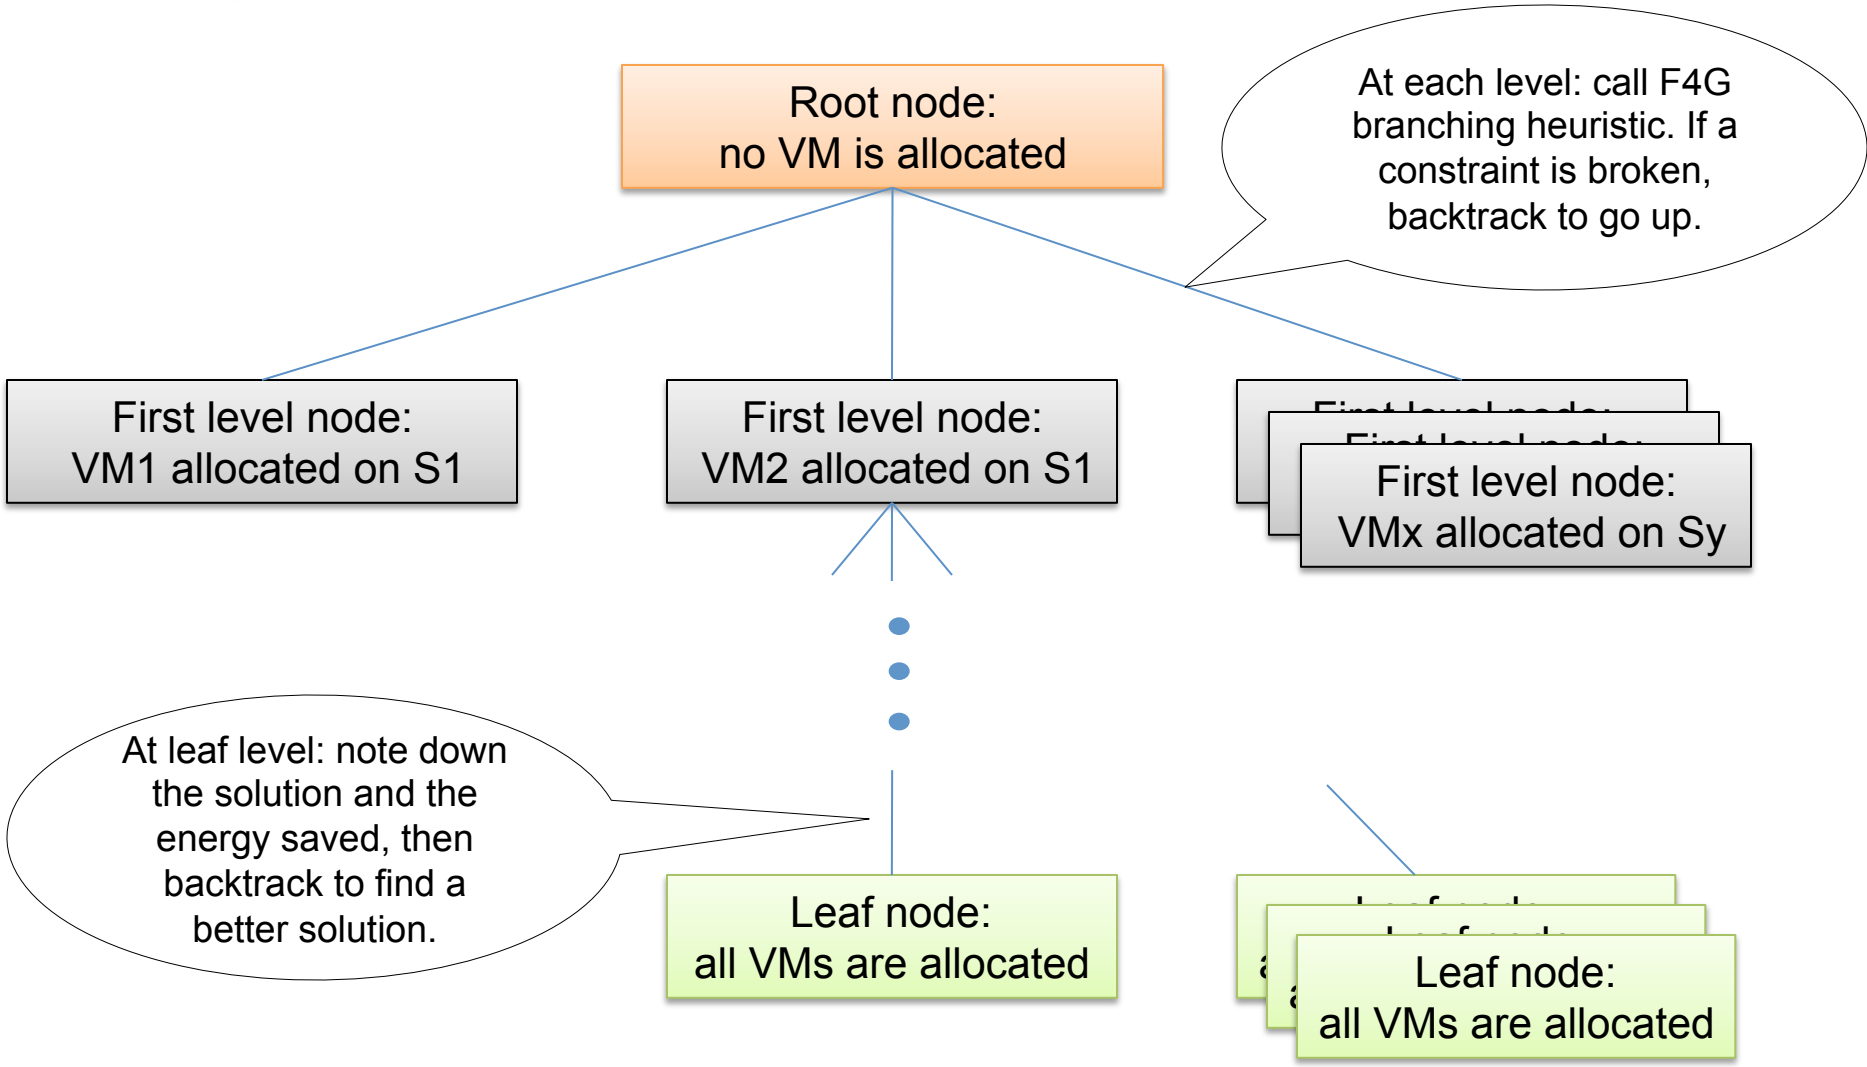
Root node:
no VM is allocated
At each level: call F4G branching heuristic. If a constraint is broken, backtrack to go up.
First level node:
VM1 allocated on S1
First level node:
VM2 allocated on S1
First level node:
VMx allocated on Sy
At leaf level: note down the solution and the energy saved, then backtrack to find a better solution.
Leaf node:
all VMs are allocated
Leaf node:
all VMs are allocated

Composable heuristics

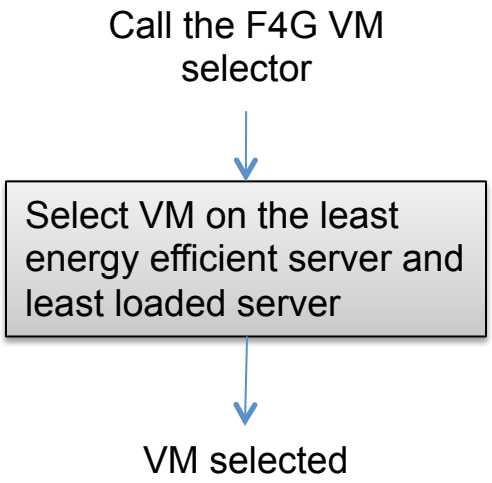

- Candidate VM for migration

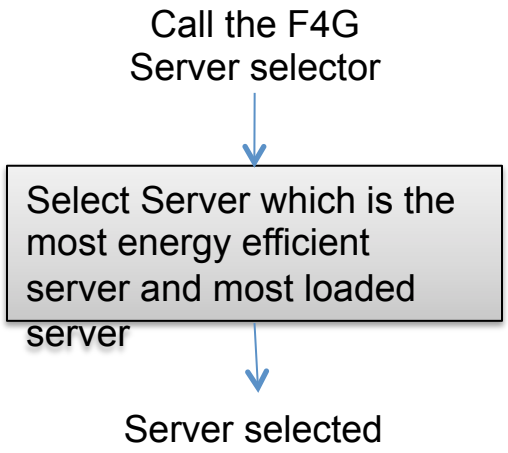

- Target server for migration

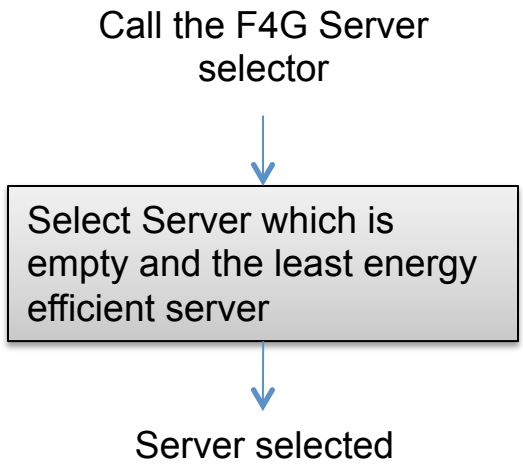

- Candidate Server for extinction

To sum up…

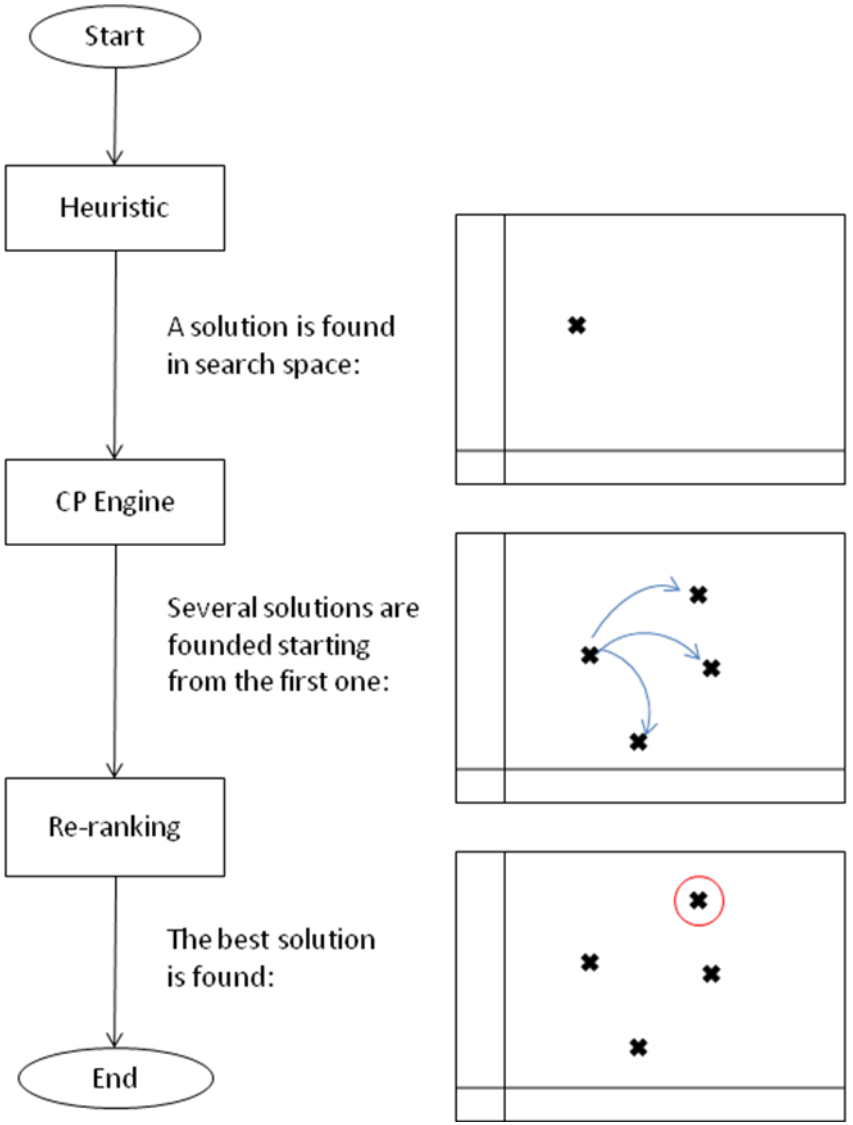
Start
Heuristic
A solution is found in search space:
CP Engine
Several solutions are founded starting from the first one:
Re-ranking
The best solution is found:
End

Experiments on Cloud Testbed

Lab trial ressources

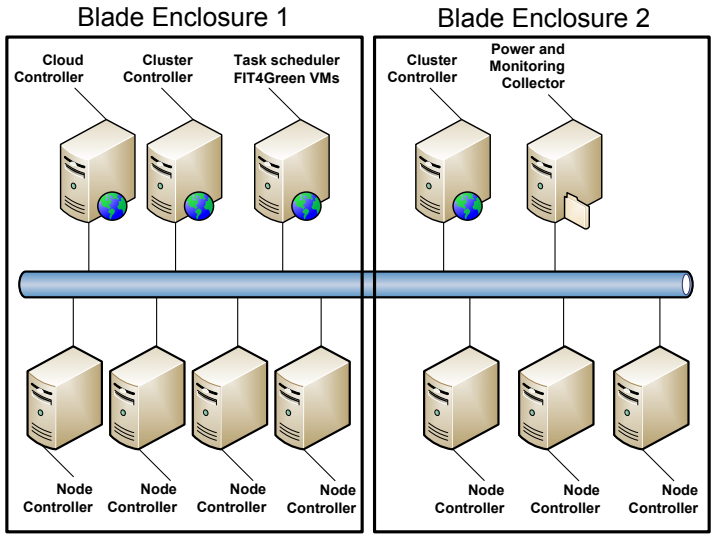

	Enclosure 1	Enclosure 2
Processor model	Intel Xeon E5520	Intel Xeon E5540
CPU frequency	2.27GHz	2.53GHz
Cpu& Cores	Dual cpu – Quad core	Dual cpu – Quad core
RAM	24 GB	24GB

- DC1: 4 BL 460c blades using VMWare ESX v4.0 native hypervisor, 3 blades for Cluster and Cloud Control
- DC2: 3 BL460c blades using VMWare ESX v4.0 native hypervisor, 2 blades for Cluster Control and Power and Monitoring System.

Experiments on Cloud Testbed

Lab trial Workload

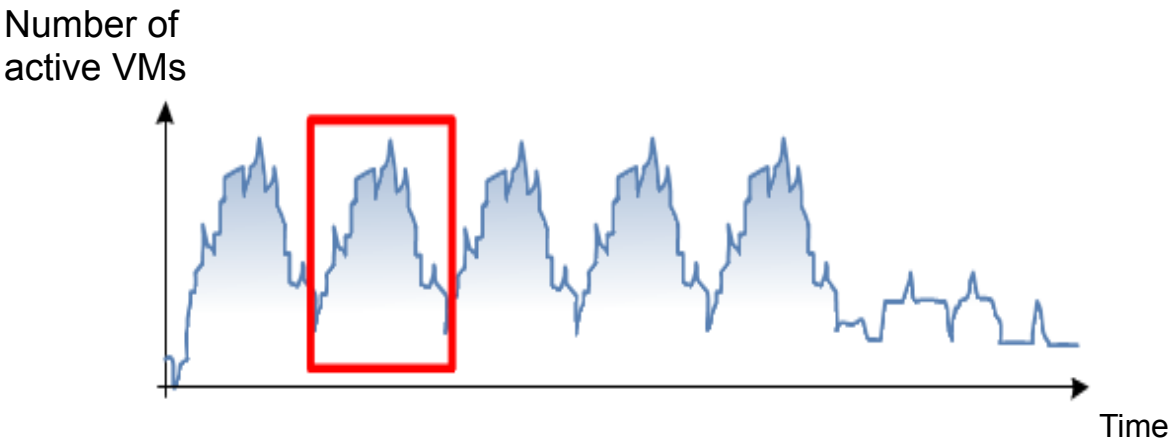

Total number of active virtual machines during full week of work

Active SLAs constraints:

- Max vCPU per core = 2
- Min VM Slot = 3
- Max VM Slot = 6

Experiments on Cloud Testbed

Final test results for the various configurations

Configuration	Data Centre 1	Data Centre 2	Energy for Federation
Without FIT4Green	6350 Wh	4701 Wh	11051 Wh
With FIT4Green Static Allocation	5190 Wh	4009 Wh	9199 Wh Saving 16.7%
With FIT4Green Dynamic Allocation	5068 Wh	3933 Wh	9001 Wh Saving 18.5%
With FIT4Green Optimized Policies	4860 Wh	3785 Wh	8645 Wh Saving 21.7%

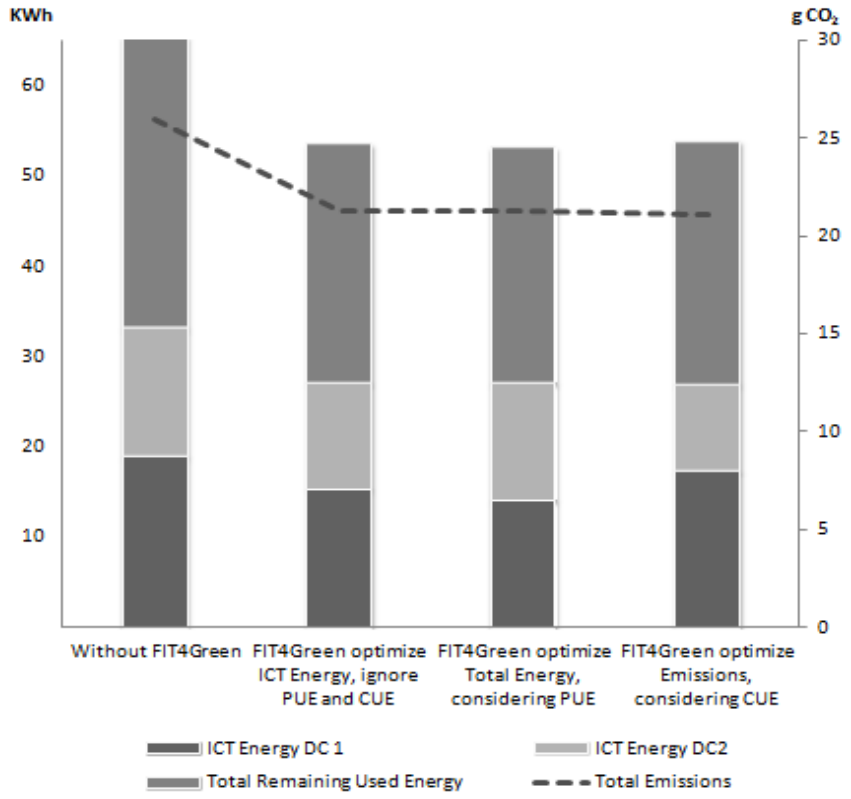

Scalability Evaluation

#	Configuration	Placement constraints activated
1	1 datacenter	none
2	1 datacenter with overbooking factor=2	"MaxVCPUPerCore" constraint set on each server
3	2 federated datacenters	"Fence" constraint set on each VM

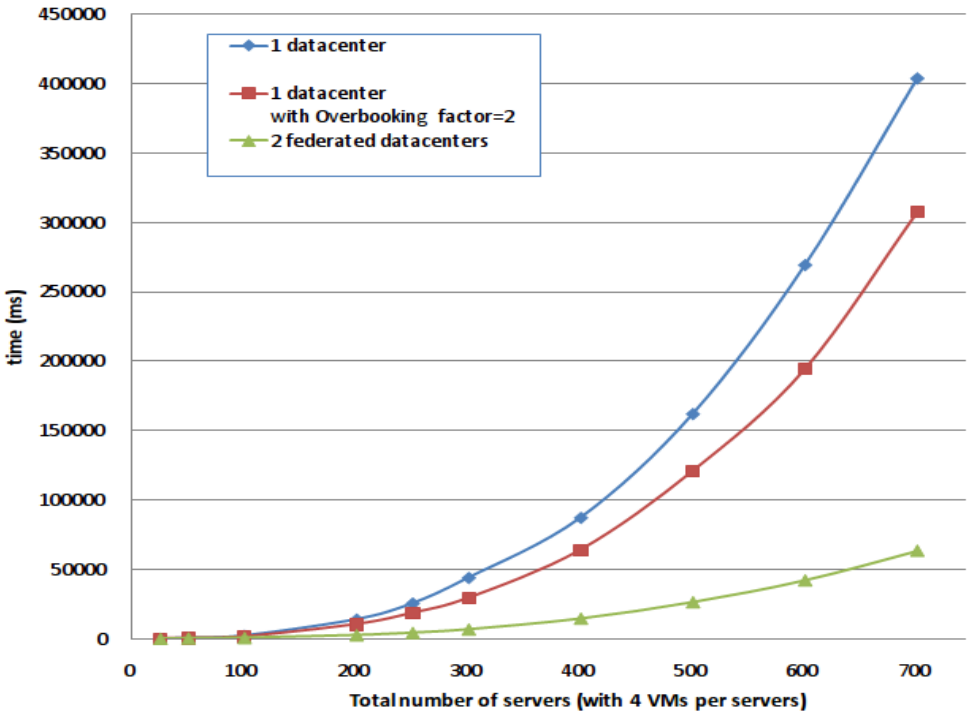

CONCLUSION & FUTURE WORK

- ✓ Energy aware resource allocation in datacenters
- ✓ Flexibility & extensibility
- ✓ Saves up to 18% in HP experiment
- ✓ Scalability with parallel processing

- ✓ Future work:
  - ✓ SLA re-negotiation
  - ✓ Green SLAs

FIT 4 GREEN

Thanks for your attention

Any Question?