Scheduling applications among heterogeneous computing resources to reach energy proportionality

Violaine VILLEBONNET
Laurent LEFEVRE – Inria Avalon ENS Lyon
Jean-Marc PIERSON, Georges DA COSTA, Patricia STOLF – IRIT Toulouse

GreenDays@Toulouse
March 17th 2015

Introduction

- Concept of Energy Proportionality
- Heterogeneous Computing - Example of ARM big.LITTLE
- Generalization at datacenter scale

Concept of Energy Proportionality

« The case for Energy-Proportional Computing »
L. A. Barroso and U. Hölzle, IEEE Computer , 2007

- Average server load between 10 and 50 %
  Most inefficient region
- High idle consumption
  Can be up to 50 % of peak power

→ A perfect proportional curve would bring huge energy savings

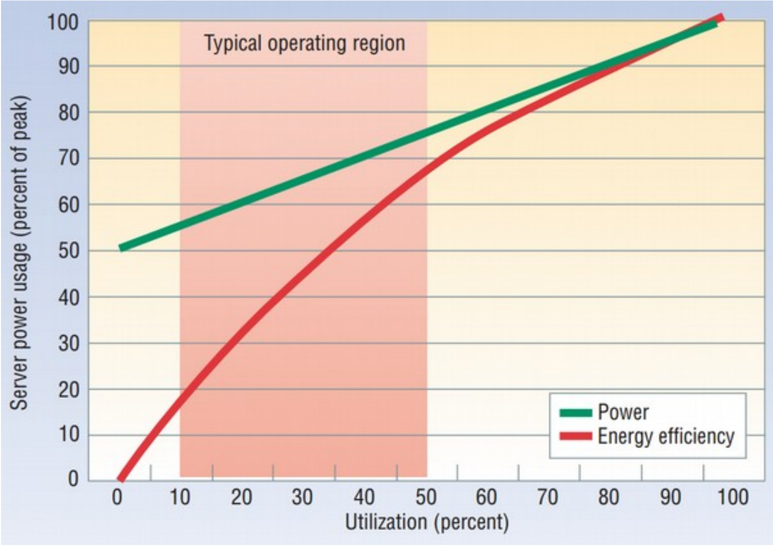

Hardware constructors are making efforts, but existing processors are still far from energy proportionality…

Heterogeneous computing

Energy proportional hardware does not exist yet
→ Heterogeneous computing !

Combination of processors with different characteristics in terms of performance and energy consumption

Example : **ARM big.LITTLE**

to extend battery life of mobile devices

2 processors (4 cores each)
- big : ARM Cortex A15
- LITTLE : ARM Cortex A7

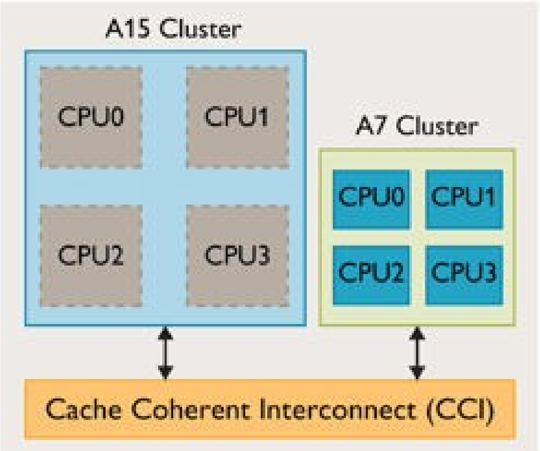

Interconnected by a Cache Coherence System

Generalization at datacenter scale

BIG

MEDIUM

LITTLE

HETEROGENEOUS INFRASTRUCTURE :

- Low power processors → Reduce static costs
- Classical servers → Only used at their most energy efficient region

RESOURCE MANAGEMENT :

Dynamically execute applications on the most suitable architecture : least consuming processor for the needed performances

TECHNICAL CHALLENGES :

- Applications migrations between different Instruction Set Architectures : ARM and x86
- Scheduling decisions : migrations and machines On/Off

Our Approach

- « Big Medium Little » datacenter infrastructure
- BML Architecture profiling
- BML Scheduling Framework

« Big Medium Little » Infrastructure

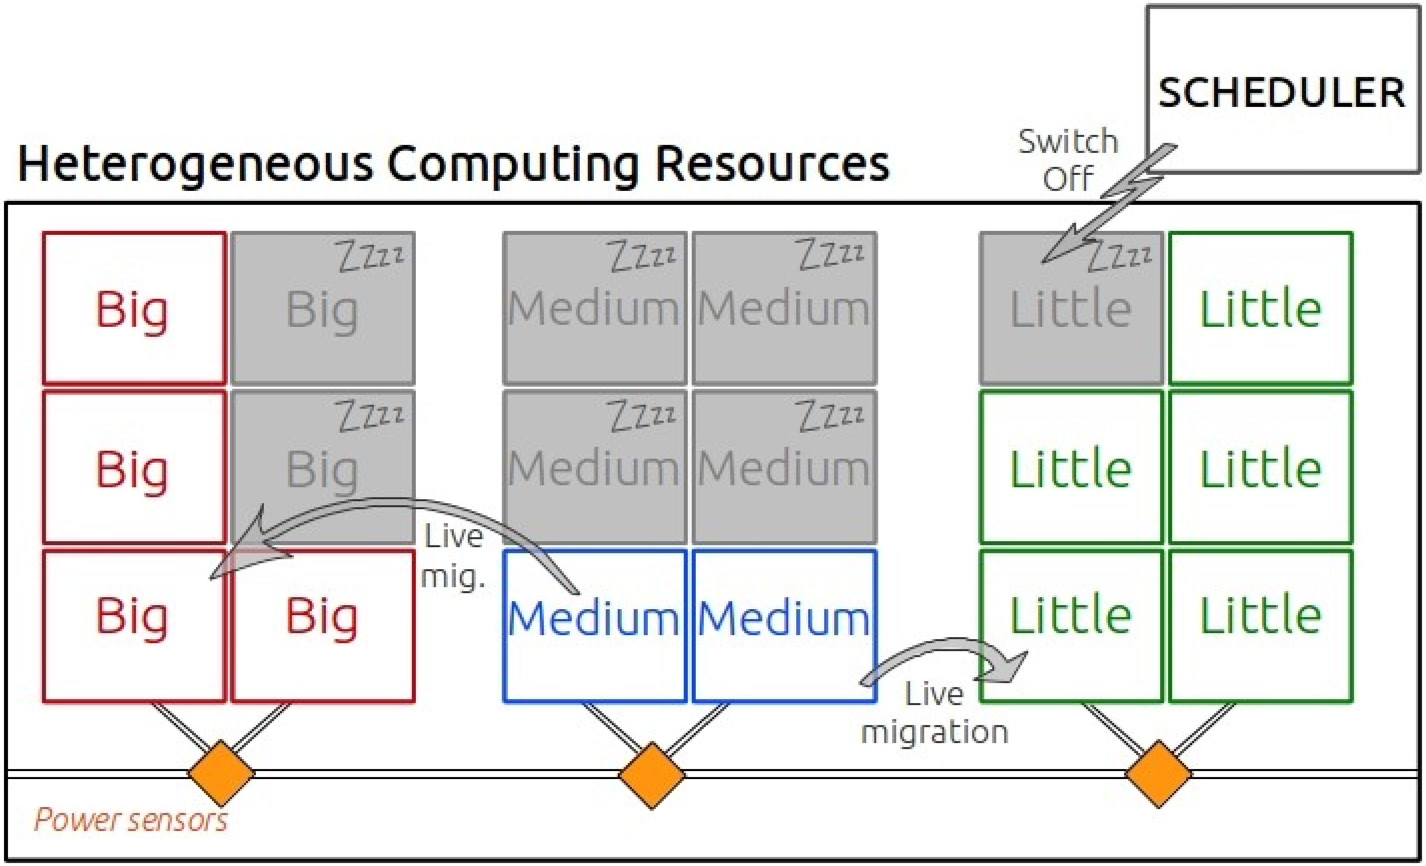

BML Infrastructure - Profiling

BIG

MEDIUM

LITTLE

Power and Performance Profiling

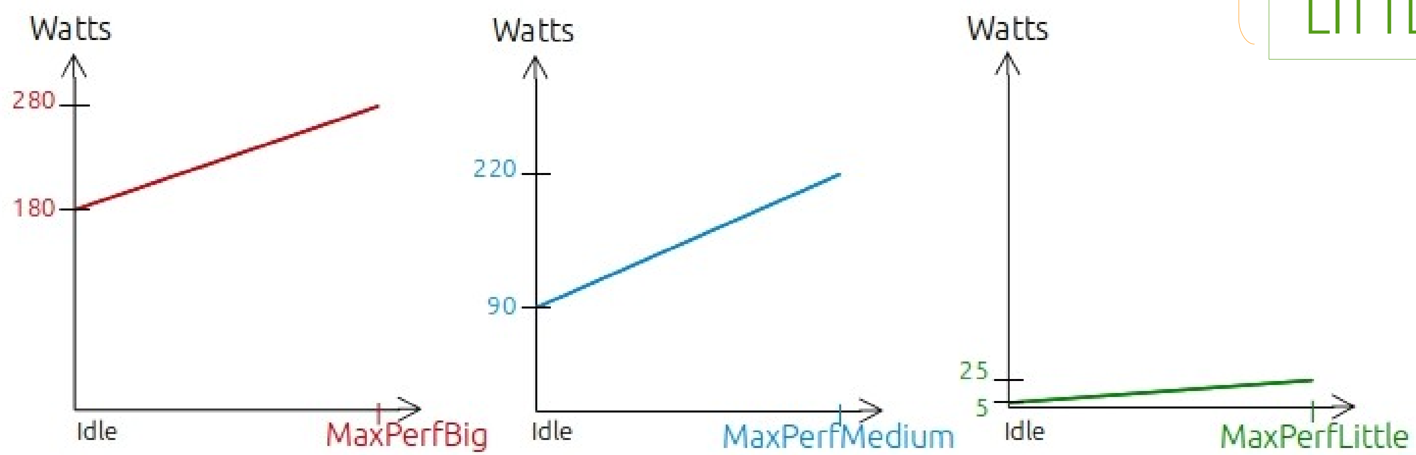

Emulation overhead

	Native performance		Emulated performance	
Big	...	...	...	...
Medium	...	...	...	...
Little	...	...	...	...

Migration time and power overhead

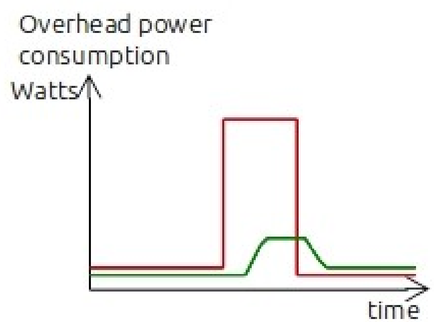

On/Off time and power overhead

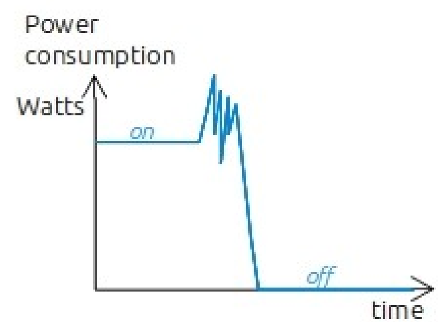

BML Scheduling Framework

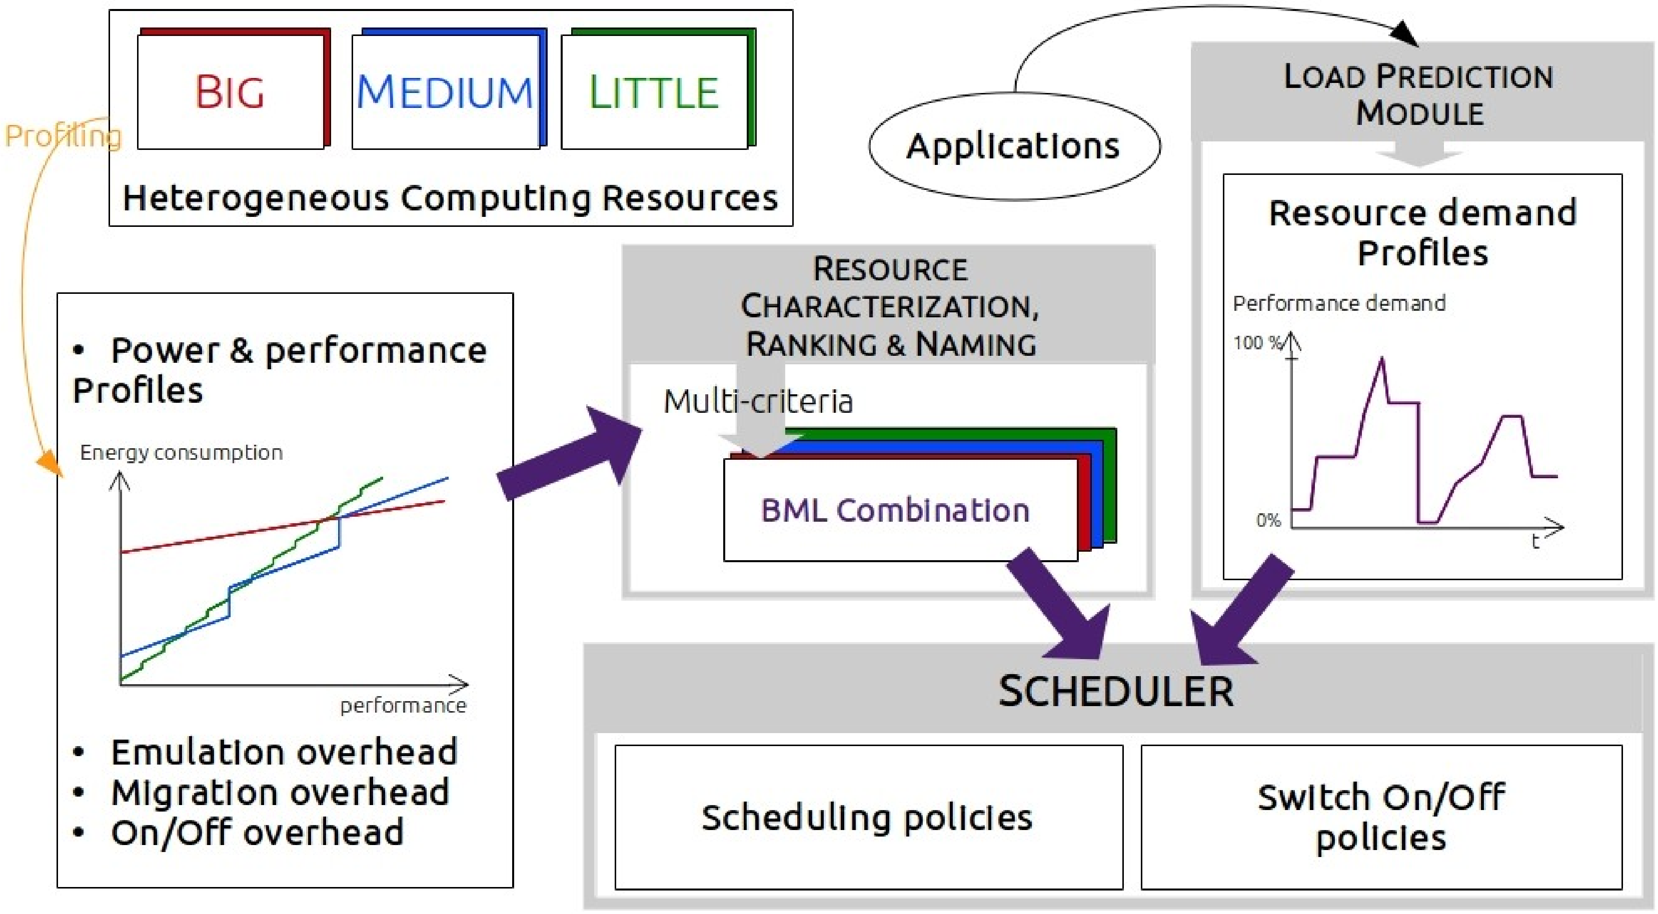

VALIDATION

- Use-case: Stateless web servers
- Experiments for hardware profiling
- Creation of Ideal BML Combination

Use-case for validation : Stateless Web Servers

Chosen Hardware :

Codename	Chromebook	Taurus	Parapluie
Fullname	Samsung Chromebook	Dell PowerEdge R720	HP Proliant DL165 G7
Architecture	ARMv7 32 bits	x86 64 bits	x86 64 bits
CPU	2 × Cortex-A15	2 × Intel Xeon E5-2630	2 × AMD Opteron 6164
Total cores	2	12	24
Power consumption	5 - 25 W	96 - 227 W	180 - 280 W
Release year	2012	2012	2010

Grid'5000

Web server benchmarking :

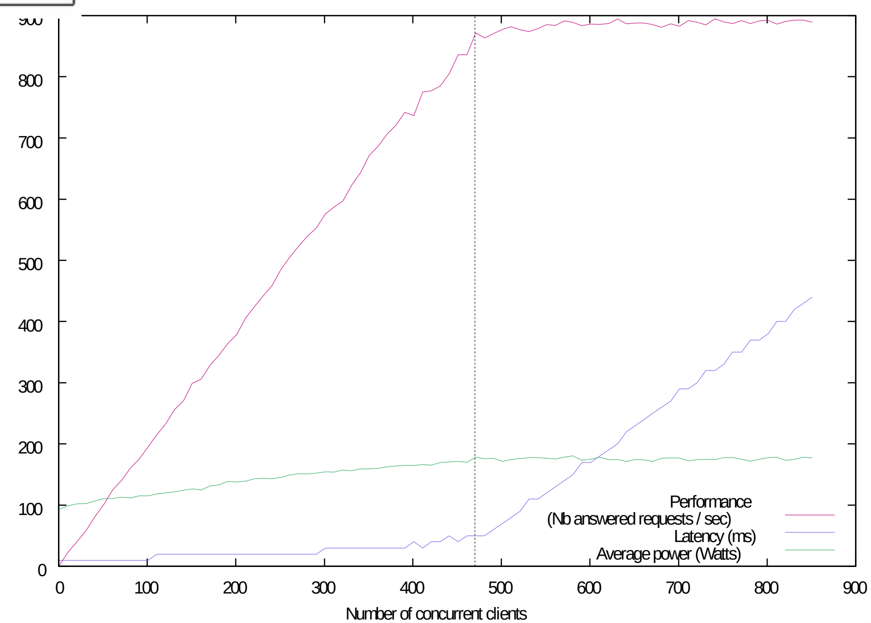

Hardware Profiling Results

Profiling and BML tagging :

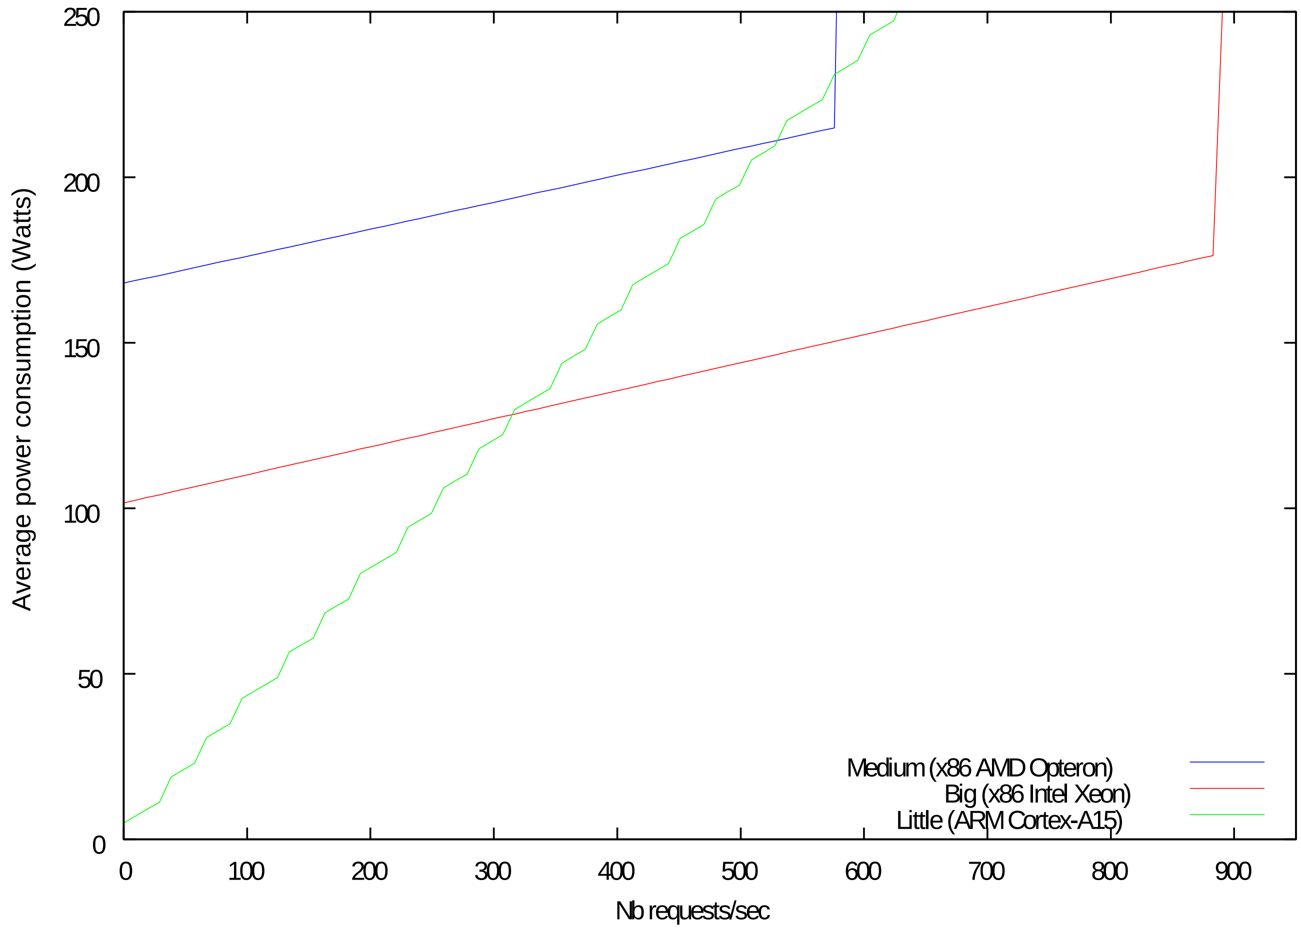

Creating Ideal BML Combination

Combination of least consuming architectures

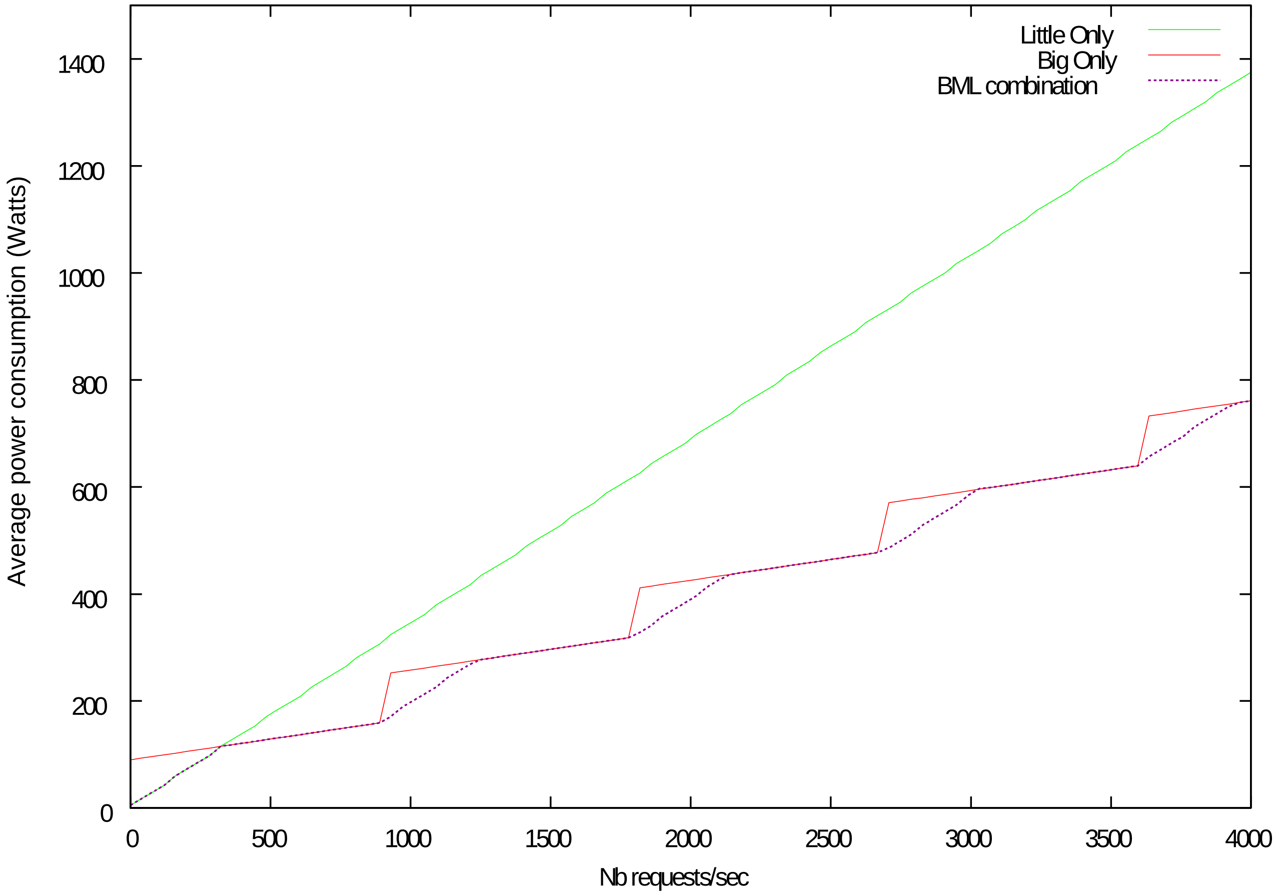

SIMULATION RESULTS

- Scenario: WorldCup 1998 web requests traces
- Gains of BML combination over homogeneous datacenters
- Considering On/Off overheads

Validation with Web Server : Scenario

World Cup 1998 traces

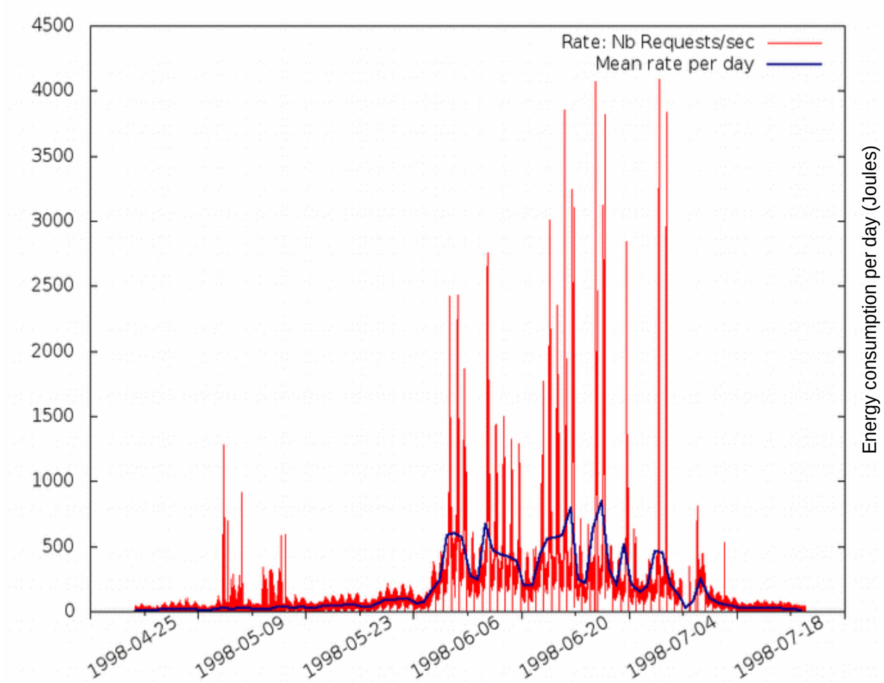

Energy consumption per day

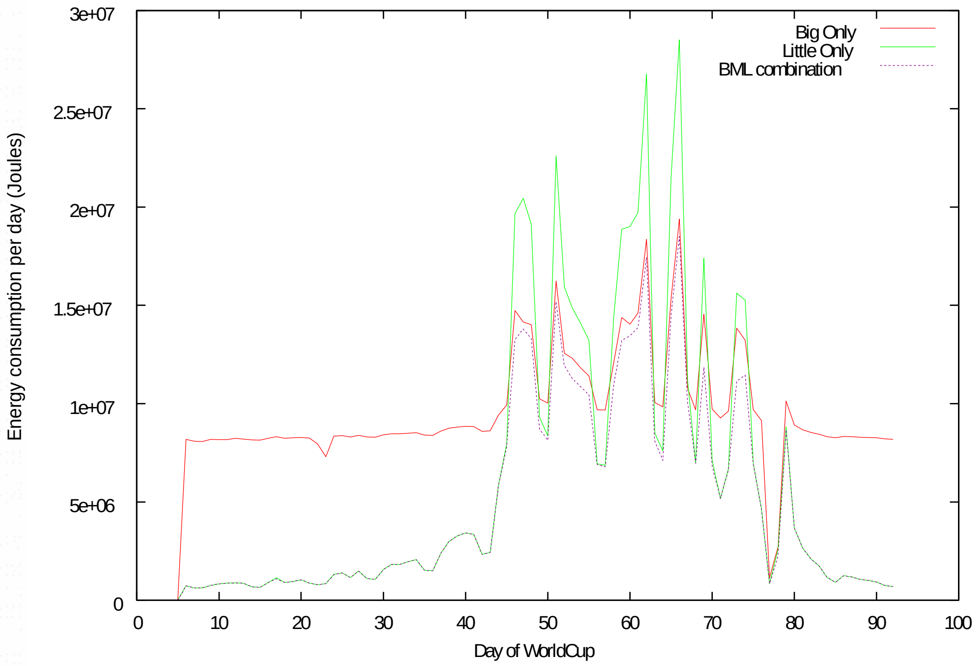

Results: Gains of BML combination

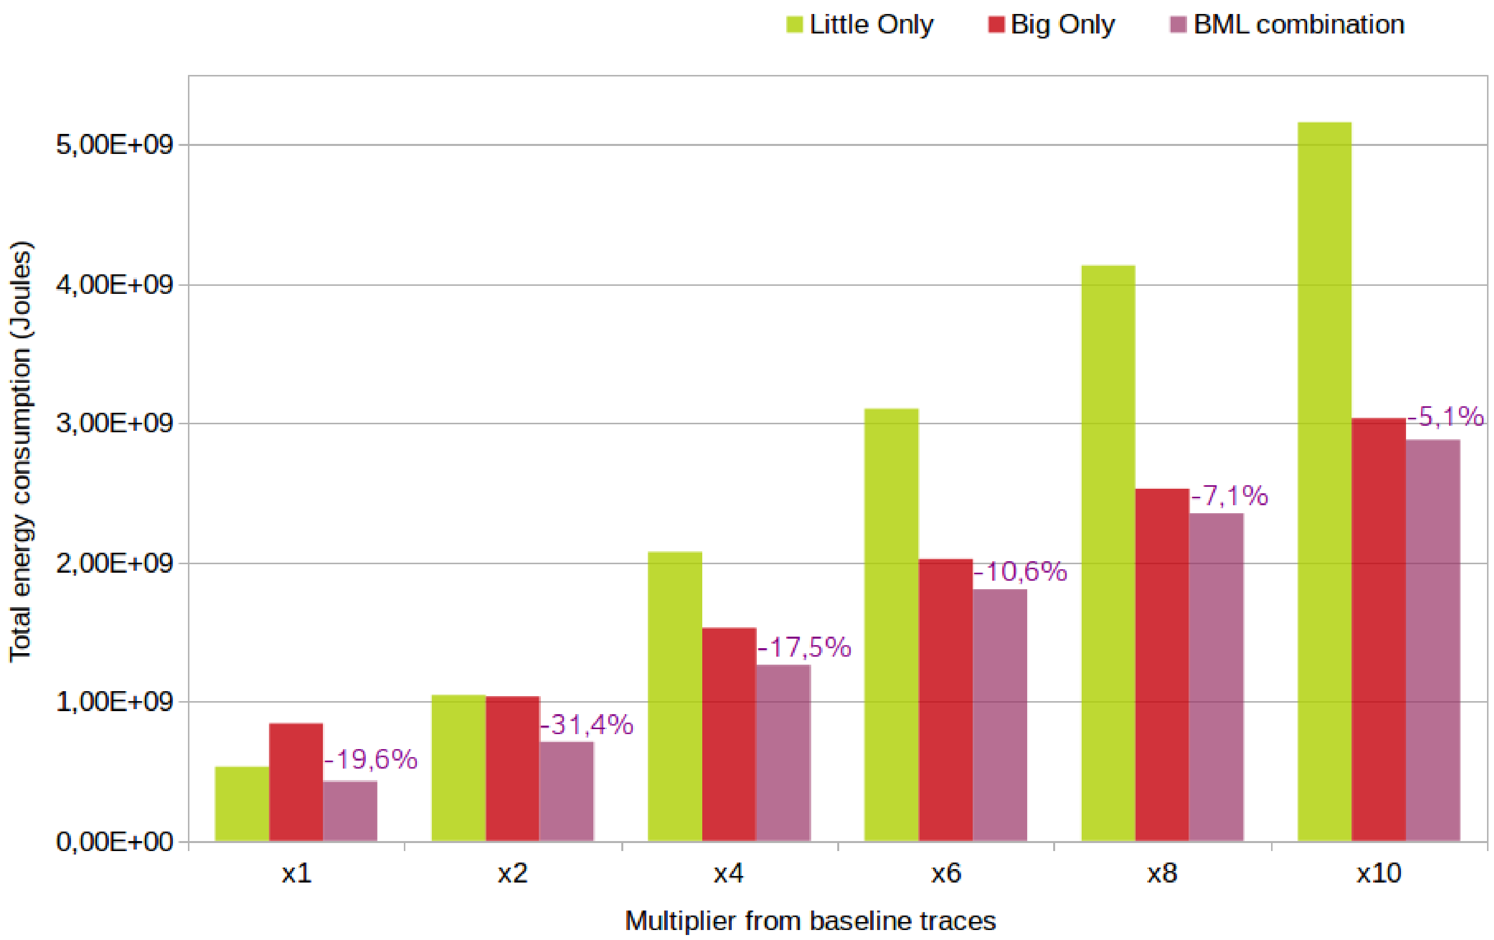

Results: Considering On/Off overheads

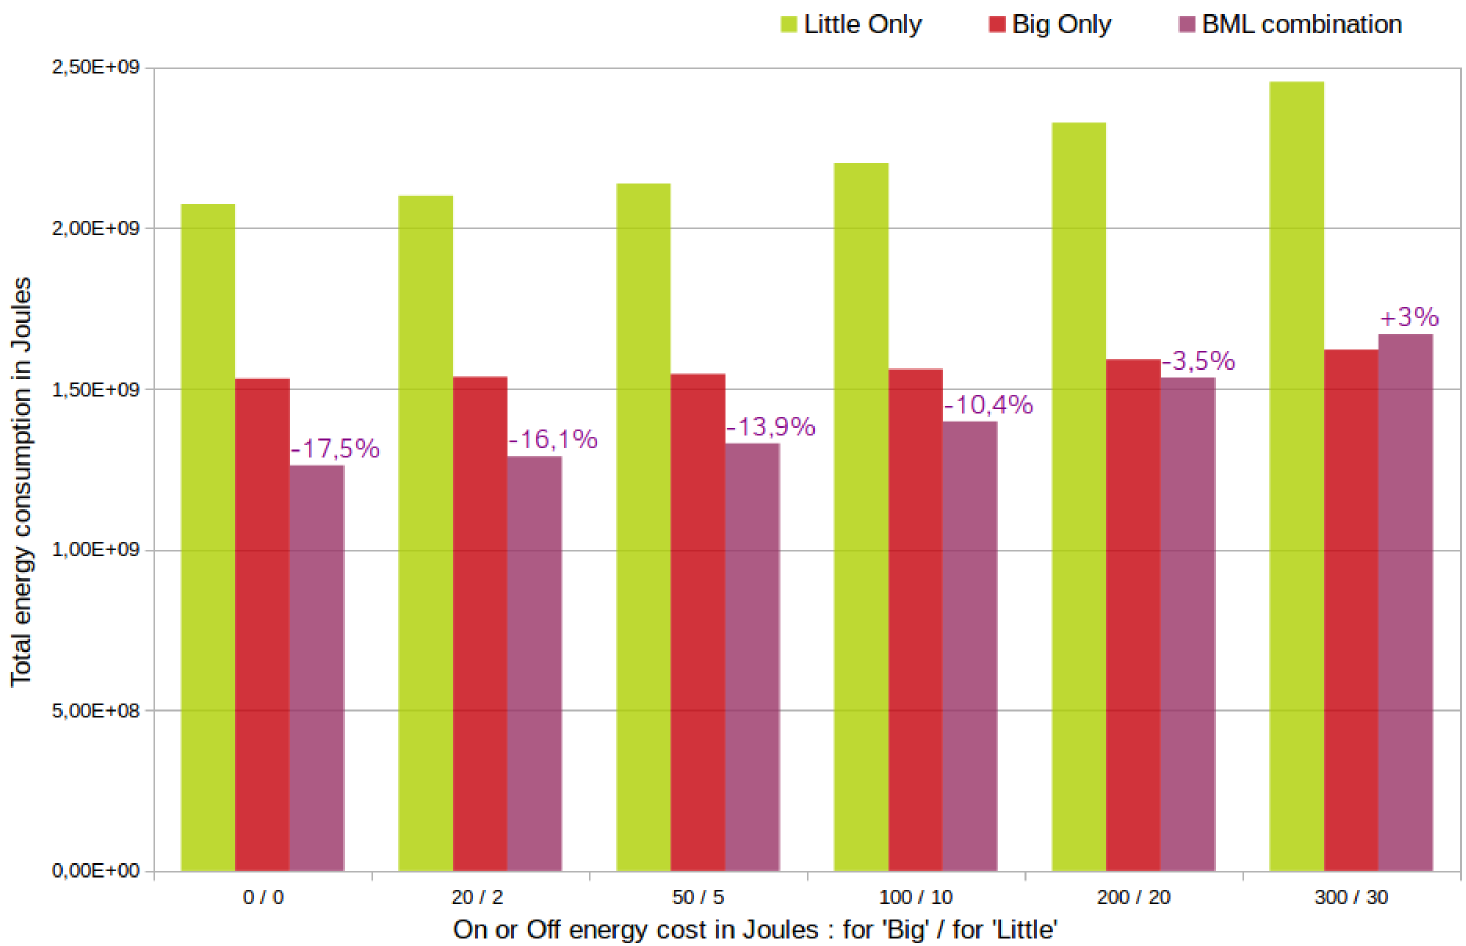

Conclusions, limitations & Future Work

CONCLUSIONS

- BML allows to adapt energy consumption in case of load variability
- High importance of On/Off decisions

FUTURE WORK

- Scheduling over limited datacenter infrastructure
- Take into account Quality of Service
- Enhance migrations & On/Off decisions
- Validation with other use-case applications

CONCLUSIONS

- BML allows to adapt energy consumption in case of load variability
- High importance of On/Off decisions

FUTURE WORK

- Scheduling over limited datacenter infrastructure
- Take into account Quality of Service
- Enhance migrations & On/Off decisions
- Validation with other use-case applications

Thank you!

violaine.villebonnet@inria.fr