



Energy Profiling/ Efficiency / Proportionality → The Green Activities in Hemera



Laurent Lefèvre
laurent.lefevre@inria.fr

Hemera Final Evaluation, Paris, December 17, 2014

INRIA AVALON / LIP
Ecole Normale
Supérieure de Lyon

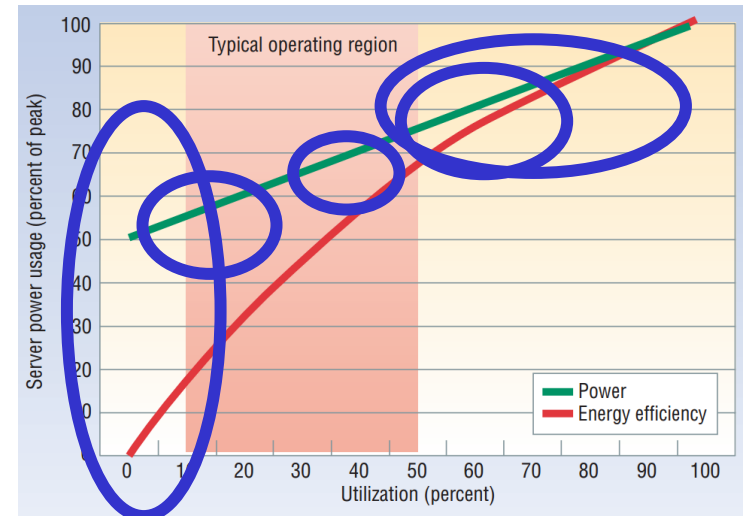
Energy : 1st limiting factor for large scale systems ((hpc)datacenter, clouds, internet)?

- Energy consumption is growing :
Top500 : Nov 2010 : 127 MW – Nov 2014 : 215 MW (not all referenced) -
Green500 : 590 MW (Nov. 14 – all referenced)
- Only usage ! not the full life cycle which is bad : planned obsolescence, rebound effect, design (rare minerals), difficult recycling...
- How to build future infrastructures platforms and make them (more) energy sustainable/responsible ? -
Multi dimension approaches :
hardware, software, usage



Energy@Hemera : 3 selected addressed scientific challenges

- **Energy Profiling and Modeling** : How to profile energy usage of large scale distributed systems ? Focus on physical resources (servers and networks) & virtual machines
- **Energy Efficiency** : Providing energy efficiency to parallel applications ? Applying large scale leverages without apriori knowledge of applications
- **Energy Proportionality** : Dealing with hybrid computing



Luiz André Barroso and Urs Hölzle, « **The case for Energy-Proportional Computing** », *IEEE Computer*, 2007

Green activities during Hemera

Scientific challenge: Energy Profiling of Large Scale Applications -- Laurent Lefèvre (Avalon), Jean-Marc Pierson (Sepia), Jean-Marc Menaud (Ascola)

Working Group: Energy Efficient Large Scale Experimental Distributed Systems -- Laurent Lefèvre (Avalon), Jean-Marc Menaud (Ascola)

Core teams : Ascola (Nantes), Avalon (Lyon), Mescal (Grenoble), Myriads (Rennes), Reso (Lyon), Astre/Sepia (IRIT, Toulouse)

Second circle teams : Adam (Lille), Dolphin (Lille), LRI, Univ Luxembourg, Telecom ParisTech + Industrial partners : Bull, Eaton

Joint Phds : G. Landry Tsafack (Reso-Astre) & V. Villebonnet (Avalon-Sepia)

Joint Postdocs: B. Cornea (Myriads-Avalon), H. Kumar Mehta (Ascola-Sepia)

Energy profiling : The Green Grid5000


Eaton, Schleifenbauer, OmegaWatt,
Dell iDrac6, WattsUp, Zimmer LMG450,...



Name	Protocol(s)/Link(s)	Frequency	Resolution
Eaton	Serial or SNMP/Ethernet	5s	1 W
Schleifenbauer	SNMP/Ethernet	3s	0.1 W
OmegaWatt	"IrDA"/Serial	1s	0.125 W
Dell iDrac6	IPMI/Ethernet	5s	7 W
WattsUp?	Legacy/USB	1s	0.1 W
ZES LMG450	Serial	0.05s	0.01 W



Extending to profiling to low power processors: SBC card (Mescal)




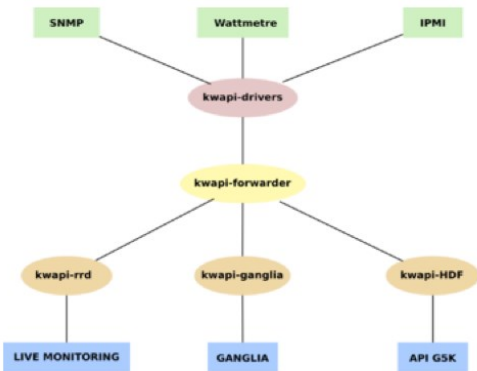
New Kwapi Grid5000 API for Energy profiling in production




Unified way to access energy logs for energy profiling

- Energy consumption per resource
- Energy consumption per user jobs

Energy measurement injection in user application and services



Details



Profiling energy consumption of virtual resources

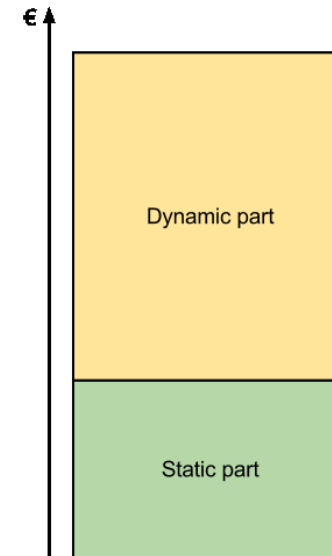
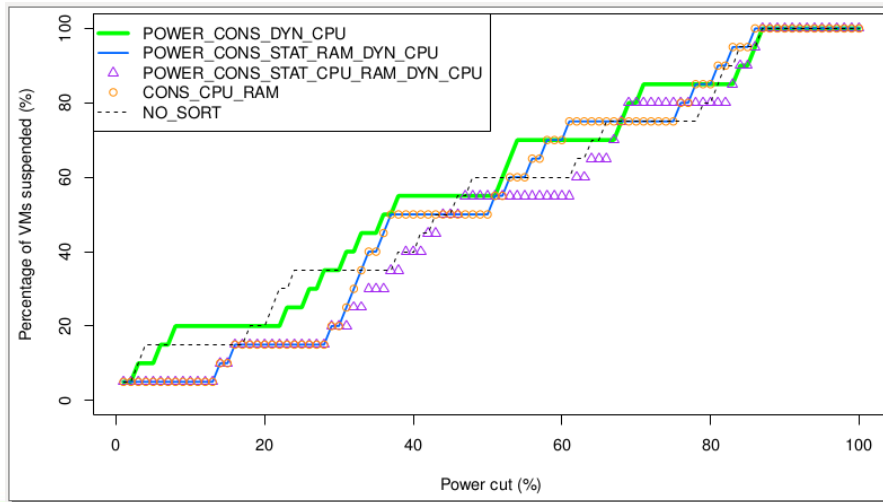



Joint Postdoc : H. Kumar Mehta (Ascola-Sepia 2013-14)
– GreenCom2013

Estimating energy consumption of virtual machines

Focus on idle VMs (in order to manage on/off of physical resources)

Various policies of splitting static&dynamic electrical costs to Vms



Including Energy profile of networks



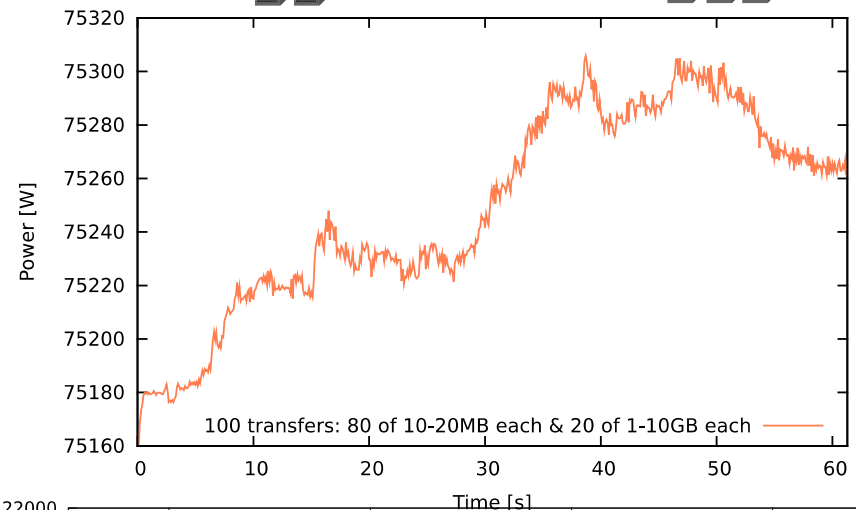

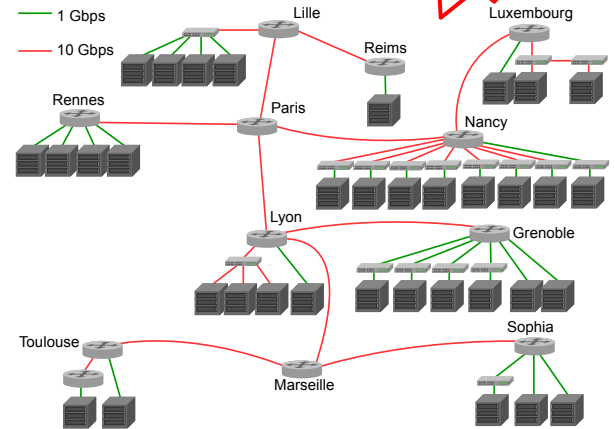


Joint Postdoc : B. Cornea (Myriads-Avalon 2013-2014) – IEEE CloudNet2014

Some infrastructure are too difficult to profile (remote network equipment, end2end architectures...) – simulation is mandatory

The ecofen approach - From NS3 to SimGrid

Ex: full network profiling, adding net energy cost in VM migration, analysis of impact of green levers



Energy efficiency : adapt resources to the need of applications




HPC applications keep growing in complexity : too many bugs in HPC applications already present, adding energy management and considerations won't help :=)

HPC programmers are not yet ready for eco design of applications

Applications can share the same infrastructure : Optimizations made for saving energy considering some applications are likely to impact the performance of others

Instead of looking at applications and service => Focusing on the infrastructure

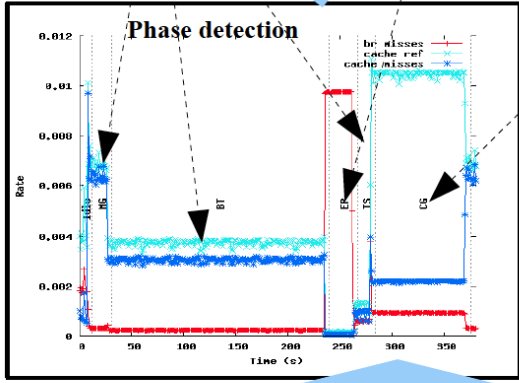
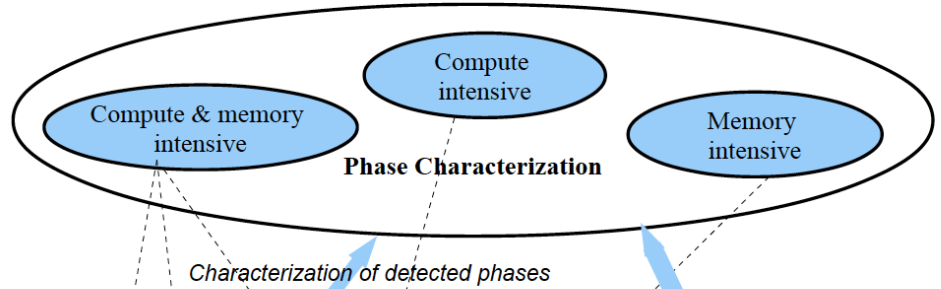
- Detect and characterize system's runtime behaviours/phases
- Optimize each subsystem (storage, memory, interconnect, CPU) accordingly



Joint PhD : G. Tsafack Chetsa
(RESO-Astre 2010-13) – FGCS2013,
PDP2015, Greencom2013,
ICPADS12, SBAC-PAD12

Online analysis without knowledge on applications

- Irregular usage of resources
- Phase detection, characterisation
- Power saving modes deployment
- MREEF framework



get the characteristics of the new vector
If identification successful


Identification and decision making

Match any new EV with a known phase if any and make the appropriate decision

memory intensive: slow down the Processor

Detect new phases

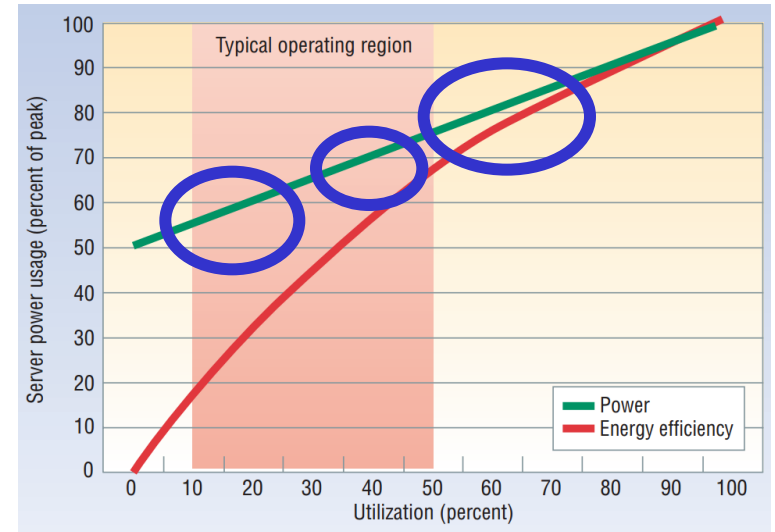
Phase label	Possible reconfiguration decisions
compute intensive	switch off memory banks; send disks to sleep; scale the processor up; put NICs into LPI mode
memory intensive	scale the processor down; decrease disks or send them to sleep; switch on memory banks
mixed	switch on memory banks; scale the processor up send disks to sleep; put NICs into LPI mode
communication intensive	switch off memory banks; scale the processor down switch on disks
I/O intensive	switch on memory banks; scale the processor down; increase disks, increase disks (if needed)



What about missing parts of the curve ?

- Specific conditions of workload
- Gaps between bursts

- Exploiting heterogeneity of processors (flops, watts, flops per watt) to fill the missing parts



The graal of energy proportionnality

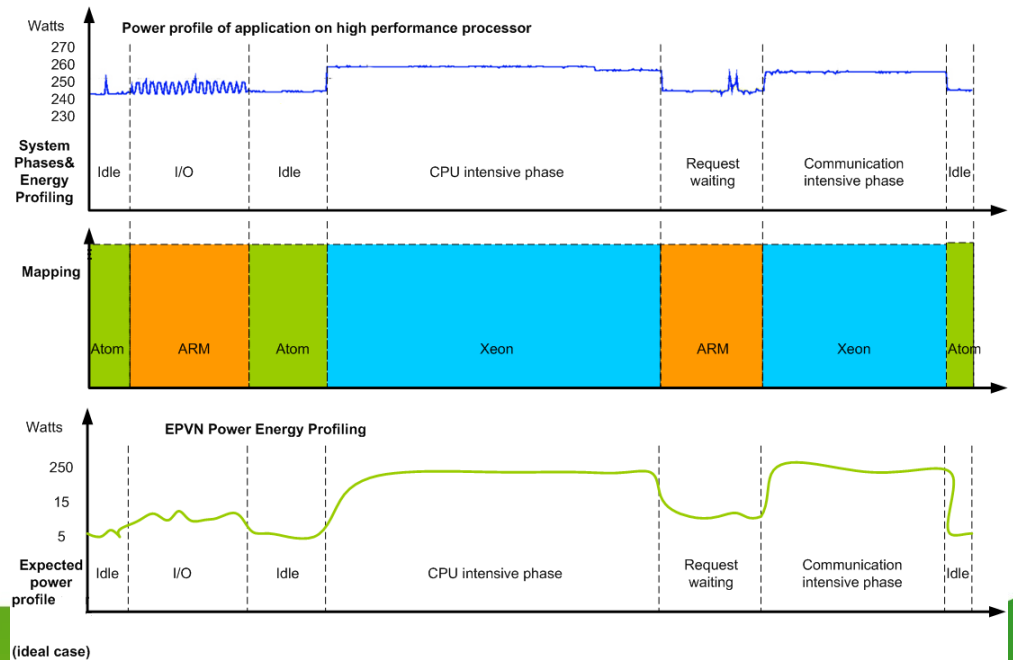
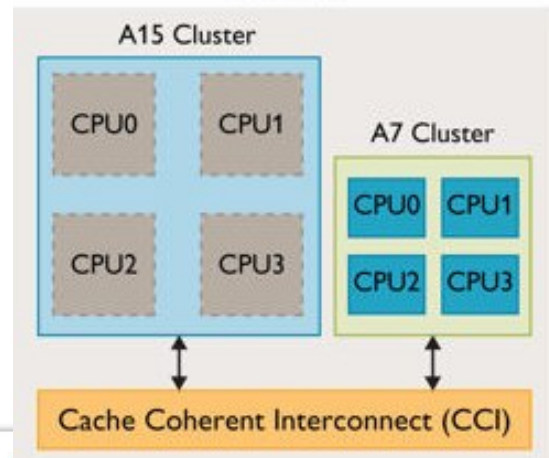



Joint Phd : V. Villebonnet (Sepia-Avalon 2013-16) – IEEE
Sustaincom2014

Dealing with heterogeneous computing for energy
efficiency and proportionality

Coordinating and mapping jobs on multi (heterogeneous)
CPUs – extending BIG.Little concept

Proof of concept : application classes



Energy@Hemera : Animations



- The GreenDays series :

- GreenDays@Rennes (July 2014): "Data deluge, virtual machines and connected objects : Energy Efficiency more than ever needed. Impact on architectures, systems and networks"
- GreenDays@Lille (Nov. 2013) : "HPC & Cloud : Are they green ?"
- GreenDays@Luxembourg (Jan. 2013) : "Energy efficiency: what else / what next?"
- GreenDays@Lyon (Jan. 2012) : "Energy efficiency: how to monitor and impact on applications?"
- GreenDays@Paris (June 2011) : "Energy efficiency: how to apply our models and solutions in production infrastructures?"

- Tutorial series during G5K school : Put some Green in your Grid'5000 experiments!
(2010,2011,2014)

- Sharing traces : ICT Energy Logs repository

A screenshot of the ICT Energy Logs website. The page has a green header with the title "ICT Energy logs" and the subtitle "energy and power consumption traces from large-scale distributed systems". There are navigation links for "Home", "Grid5000 Logs", and "Contact us". The main content area is divided into two columns. The left column has a section titled "Motivation and Goals" with text explaining the site's purpose and the challenges of large-scale distributed systems. The right column has a section titled "latest news" with two entries: "November 9, 2010" and "July 6, 2010".

Home Grid5000 Logs Contact us

ICT Energy logs

energy and power consumption traces from large-scale distributed systems

Motivation and Goals

The goal of this site is to provide (anonymized) workload, energy and power consumption traces from large-scale distributed environments to researchers and to practitioners alike.

Large-scale distributed systems are widely used from DNA sequence analysis to economic forecasting. Driven mostly by application performance, these systems have constantly increased in size and computing power of their resources.

The power supply requirements of these systems have increased in a similar fashion, which has raised concerns about the energy they consume. Power management of resources is then required on the basis of real power consumption measurements.

The goal of this website is to provide a virtual meeting place where practitioners and researchers can exchange energy traces.

Contributing

If you have traces from your large-scale distributed system (data center, grid, cloud, network), either complete or partial, please contact us. We will help you gather the traces, we will anonymize and process the traces, we will publish with your consent the traces, and we will give credit where its due.

latest news

November 9, 2010
Official launch of the ICT Energy Logs for all web page

July 6, 2010
Grid5000 Logs are available !!
Description of the logs can be found [here](#).

Energy@Hemera Summary

Addressing strong scientific challenges which were needing joint teams expertise

Community structuring / sharing of knowledge

Exchange and integration of tools and software components between teams,
prototype tools to “production”

Post Hemera : Green activities continue...

Animation : GreenDays@Toulouse (March 2015)

Joint Proposals : ANR

Inria
INVENTEURS DU MONDE NUMÉRIQUE

laurent.lefevre@inria.fr