PARA DIME

Power Characterization of Servers in Heterogeneous Cloud Environments

Mascha Kurpicz, Anita Sobe, Pascal Felber
Université de Neuchâtel

This project and the research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement n° 318693

Motivation

- Bigger data centers
- More powerful CPUs

- Cloud computing requires more energy than India or Germany
- Goal:

  Reduce energy consumption on multiple levels

PARA DIME

Context

- Heterogeneous hardware within a data center is common
- Multi-cloud scenarios: connecting heterogeneous data centers

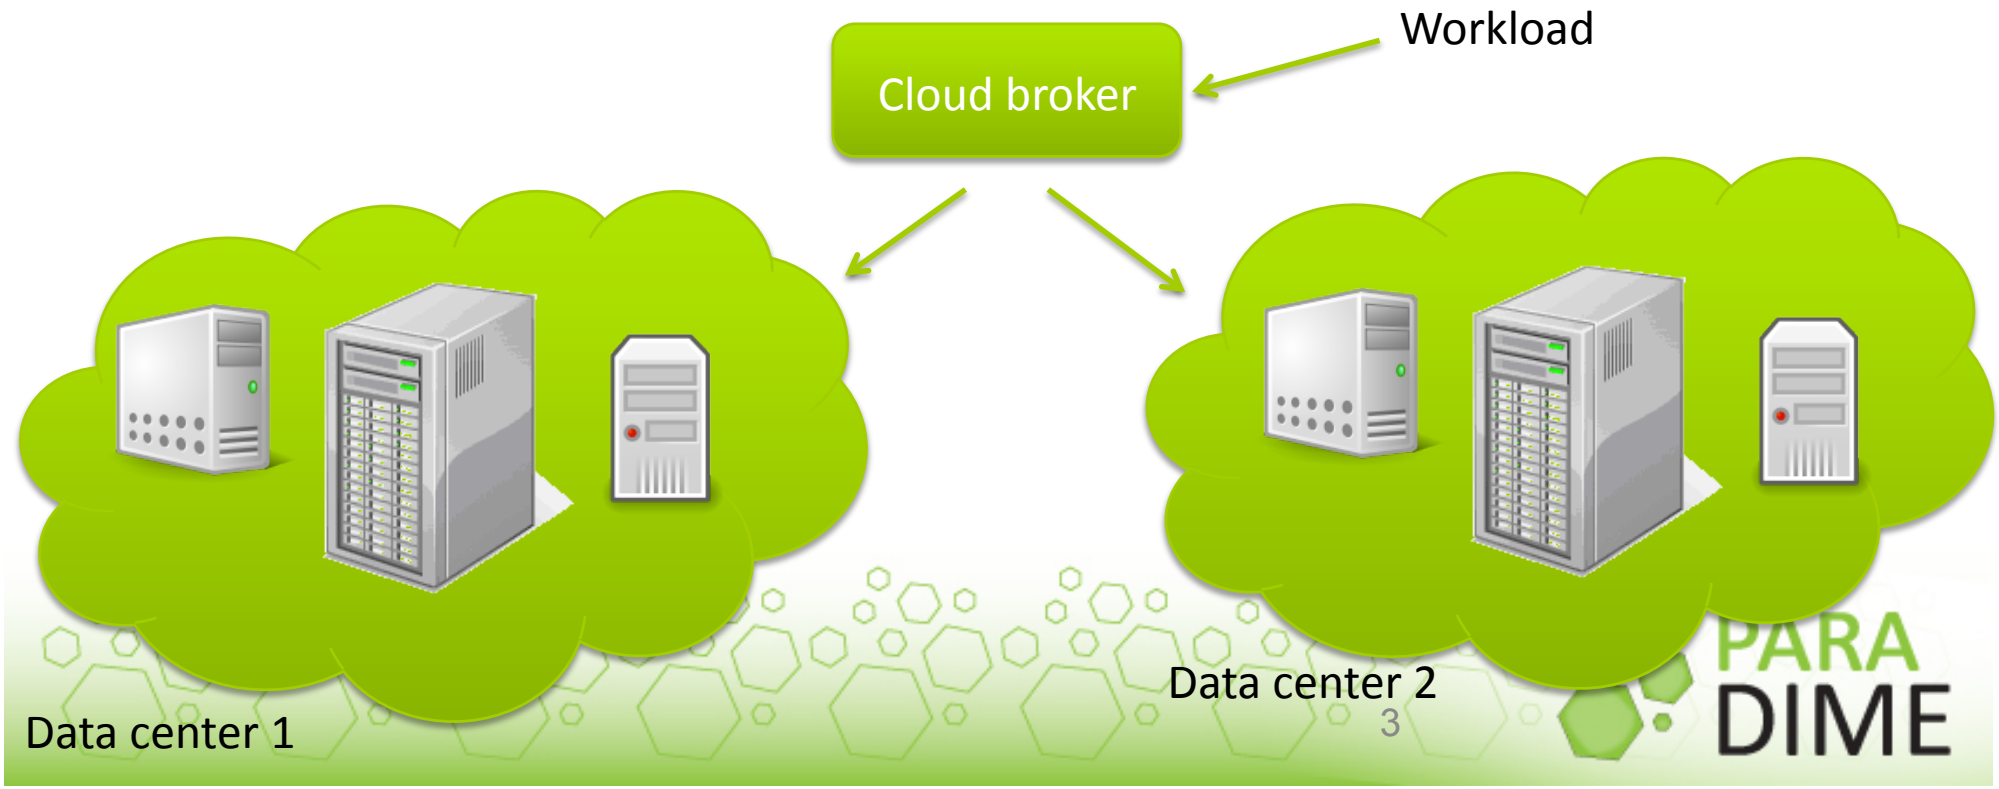

CrossCloud Brokers‘14

- Study about power consumption for different workloads
  - CPU
  - Disk
  - Real-world application
- On heterogeneous hardware

"**Using Power Measurements as a Basis for Workload Placement in Heterogeneous Multi-Cloud Environments**", Kurpicz, M., A. Sobe, And P. Felber, CrossCloudBrokers '14 (co-located to Middleware 2014), Bordeaux, France, ACM, 12/2014.

PARA DIME

Physical power meter

- PowerSpy device from Alciom
- Setup: power every second (Watt)

Bluetooth

PARA DIME

Metrics

$$E = P \times t$$ Joule = Watt x seconds

$$\frac{Perf}{W} = \frac{Throughput}{P}$$ E.g. for disk workload: Read Rate / Watt

PARA DIME

Idle power consumption

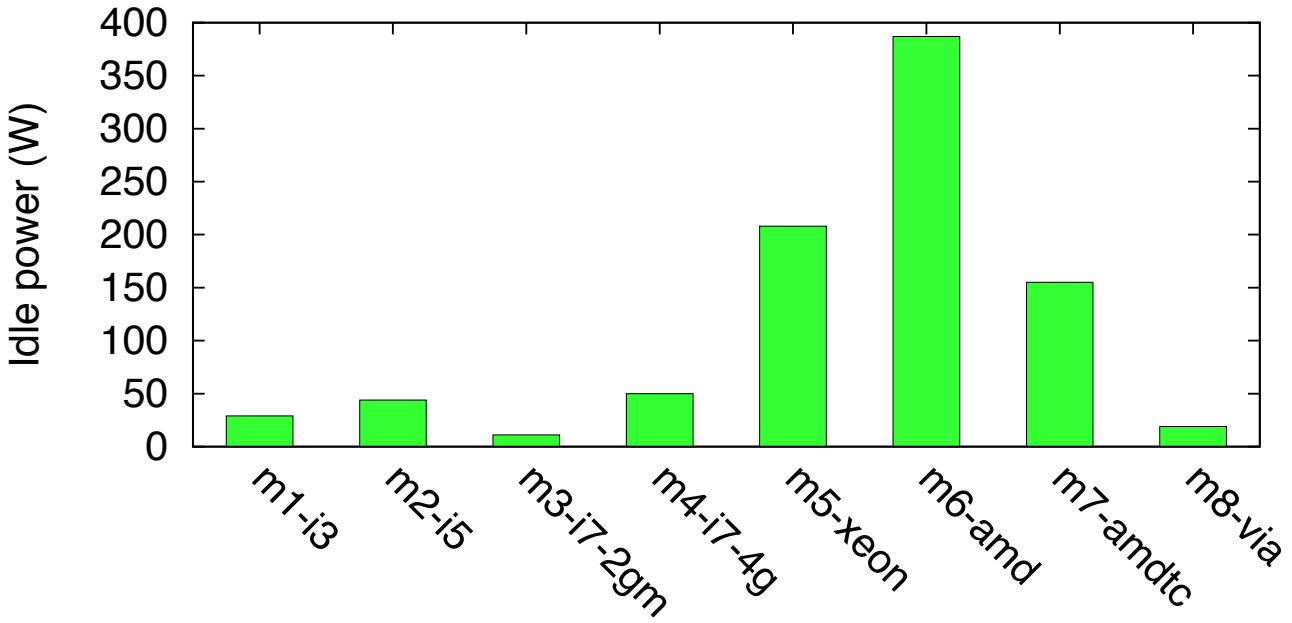

M1-i3	M2-i5	M3-i7-2gm	M4-i7-4g	M5-xeon	M6-amd	M7-amdtc	M8-via
Desktop	Desktop	Mobile	Desktop	Desktop	Server	Desktop	Mobile

Idle power consumption

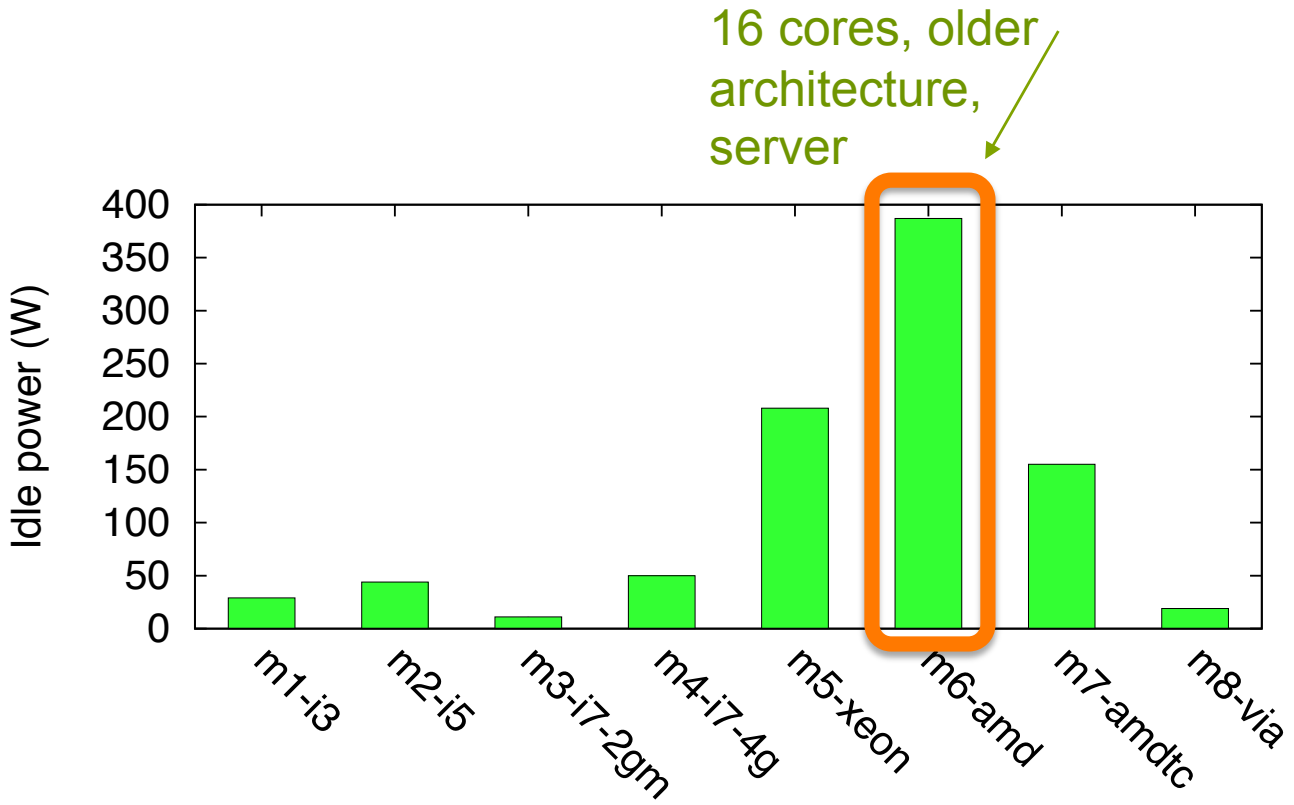

M1-i3	M2-i5	M3-i7-2gm	M4-i7-4g	M5-xeon	M6-amd	M7-amdtc	M8-via
Desktop	Desktop	Mobile	Desktop	Desktop	Server	Desktop	Mobile

CPU workload (factorial)

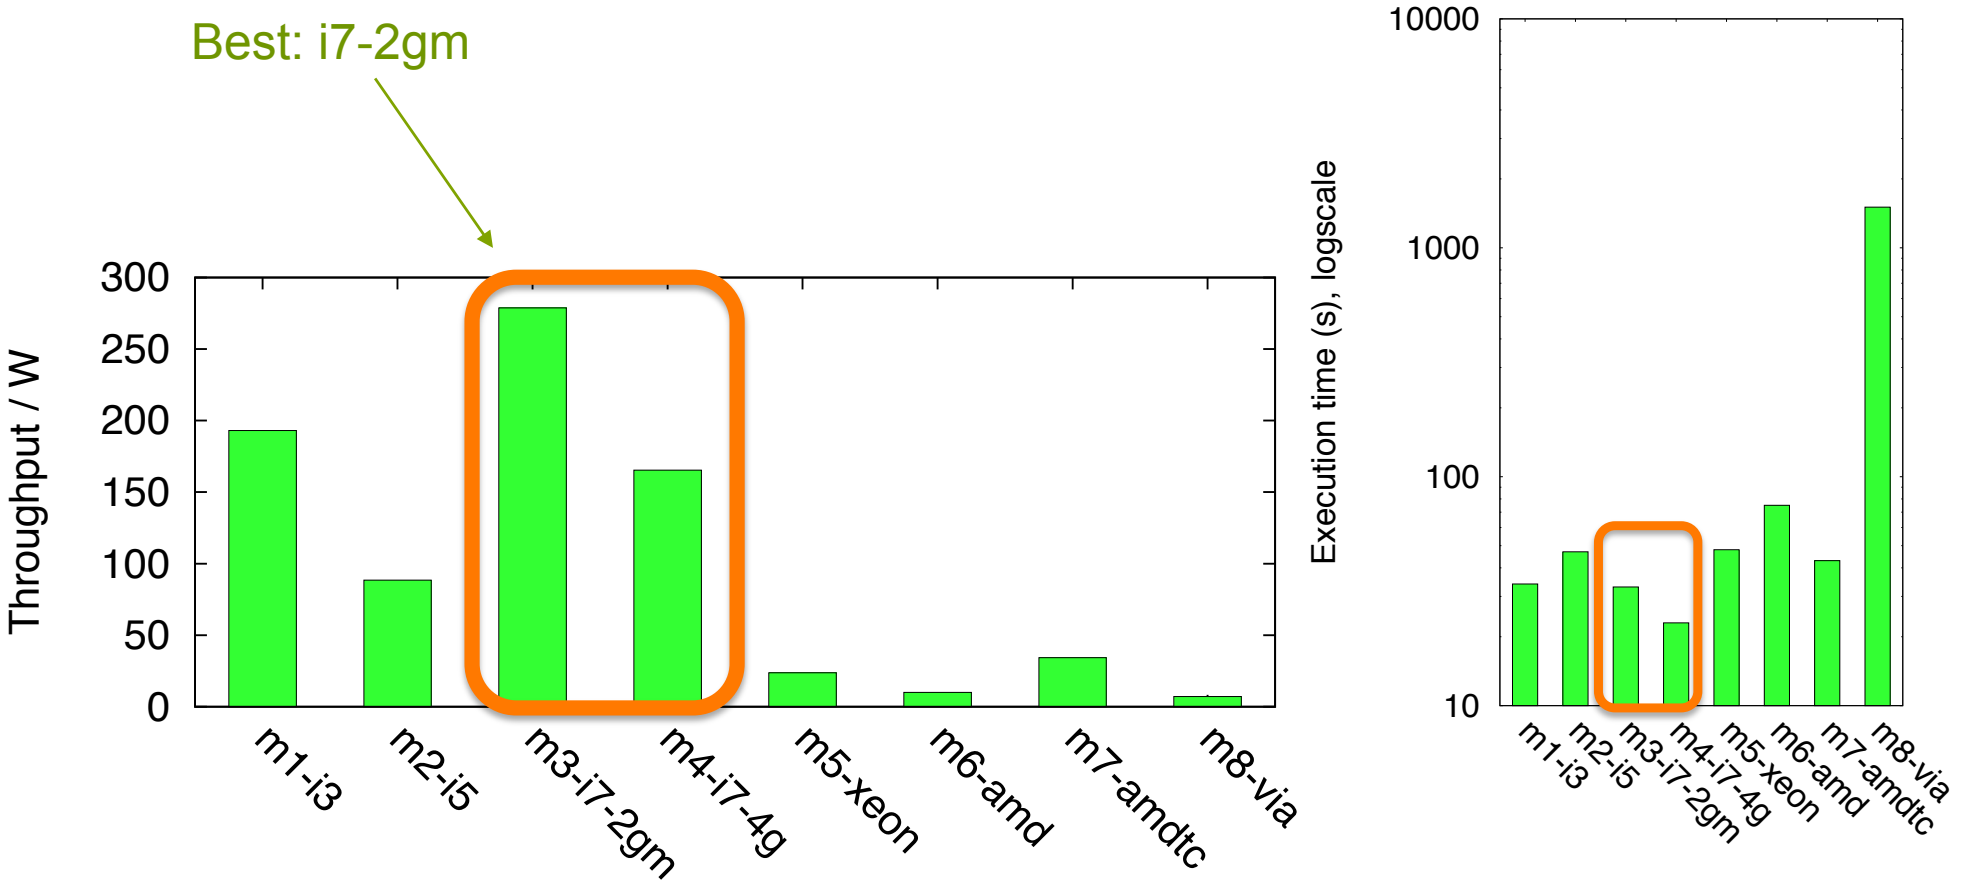

PARA DIME

Disk workload (Bonnie++)

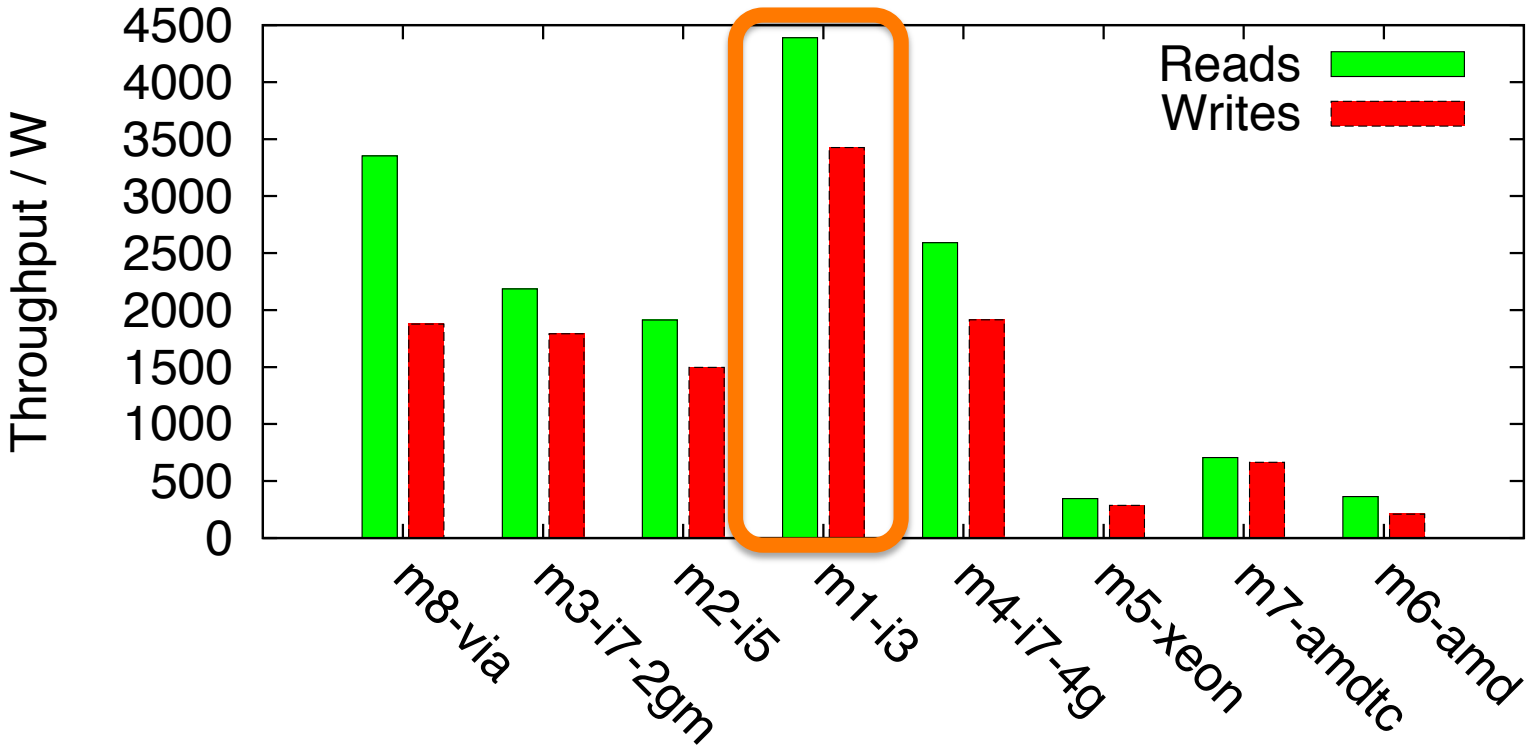

	M8-via	M3-i7-2gm	M2-i5	M1-i3	M4-i7-4g	M5-xeon	M7-amdtc	M6-amd
Type	Mobile	Mobile	Desktop	Desktop	Desktop	Desktop	Desktop	Server
Disk RPM	5400	5400	5900	7200	7200	7200	7200	7200

PARA DIME

Impact on energy-aware scheduling

- Different scheduling possibilities on the same two machines

	M1-i3	M4-i7-4g	Total (J)
Placement 1	5xDisk	5xCPU	14’370
Placement 2	5xCPU	5xDisk	16’110

M1-i3	M4-i7-4g
Desktop	Desktop

PARA DIME

Current work: Job and HW profiles

- HW profile on reference machine
- Extrapolation from one machine to another
- Online job profiling
- Estimation of job energy consumption as input for scheduling decision

TECHNISCHE UNIVERSITÄT DRESDEN

UNIVERSITÉ DE NEUCHÂTEL

PARA DIME

HW profile

- Profile machine m1 as a reference
- CPU (*usr* and *sys*) and *disk*
- Utilization intervals $u_1,\ldots,u_n$

Util	u_1	u_2	...
usr	10W	20W	...
sys	10W	15W	...
disk	3W	5W	...

PARA DIME

Extrapolation for other HW

Util	u_1	u_2	...
usr	10W	20W	...
sys	10W	15W	...
disk	3W	5W	...

...

Util	u_1	u_2	...
usr			...
sys			...
disk			...

PARA DIME

Utilization mapping between machines

- Utilization mapping tables
  - For CPU (*sys* and *usr*)
  - For *disk*

sys(%)			
m1	10	20	…
m2			
m3			
…			

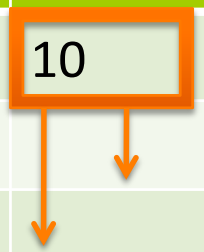

PARA DIME

Online job profiling – machine m1

- On job arrival, monitor part of the job on m1
- Measure CPU and disk utilization
- Look up power consumption in HW profile of m1

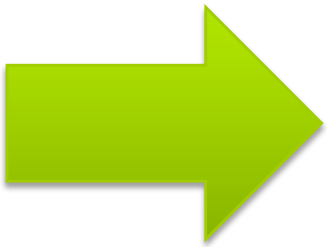

Expected power consumption on machine m1

Online job profiling – other machines

- Utilization mapping for machines $m_2,..,m_n$
- Look up power values in HW table for mapped utilization values

Expected power consumption on machines $m_2,...,m_n$

- Provide table with expected power consumption for the different machines to the scheduler

Workflow

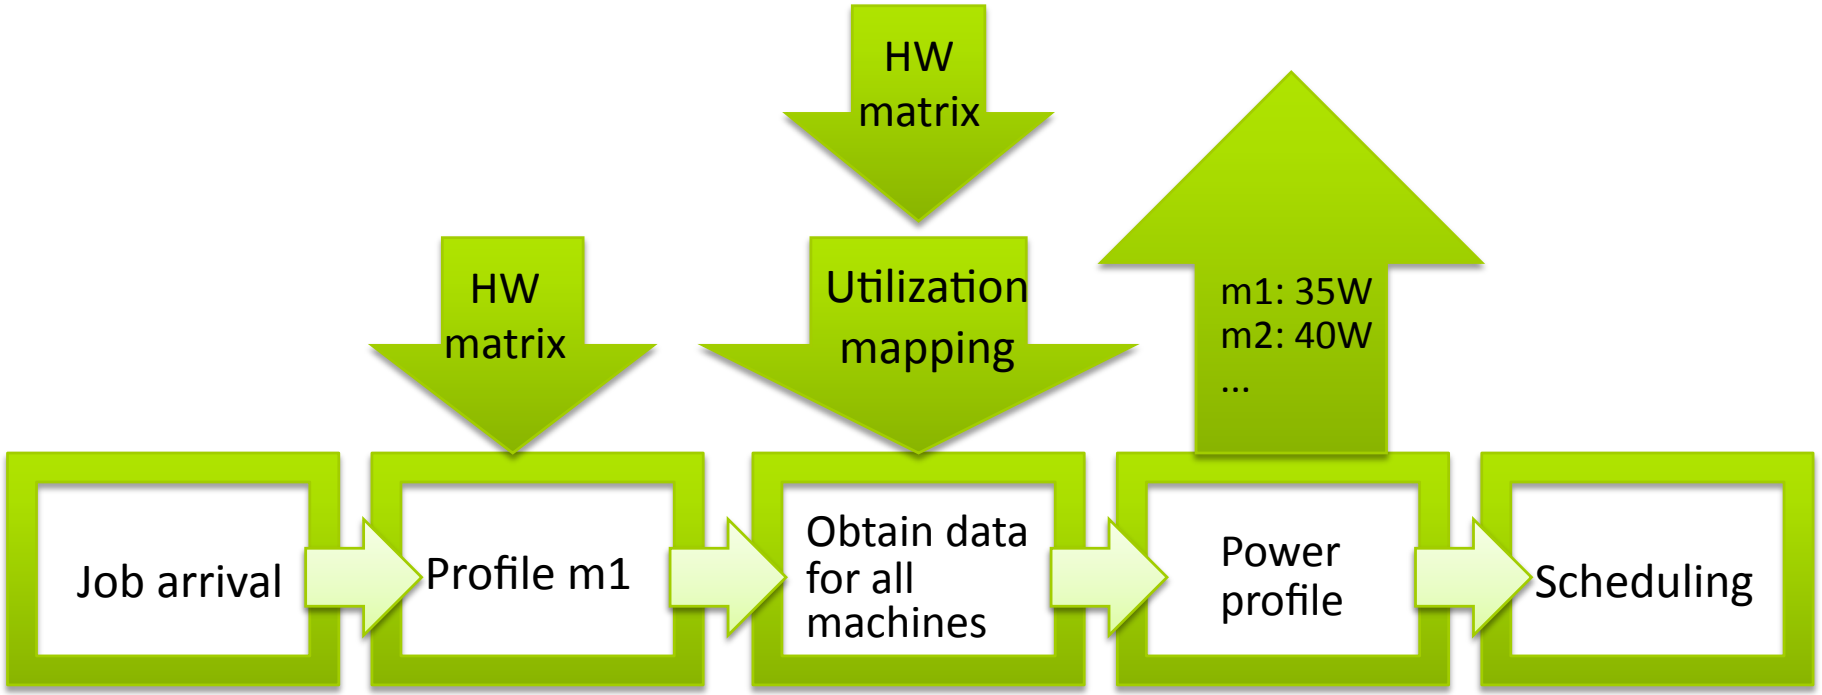
HW matrix
HW matrix
Utilization mapping
m1: 35W
m2: 40W
...
Job arrival
Profile m1
Obtain data for all machines
Power profile
Scheduling

Scheduler

Estimated power consumption on each machine

Estimated execution time

Energy efficient scheduling decision

Open points

- Data locality
- Which subset of the workload to monitor?
- What HW can be covered by the model?
- Exact definition of the mapping functions

Conclusion

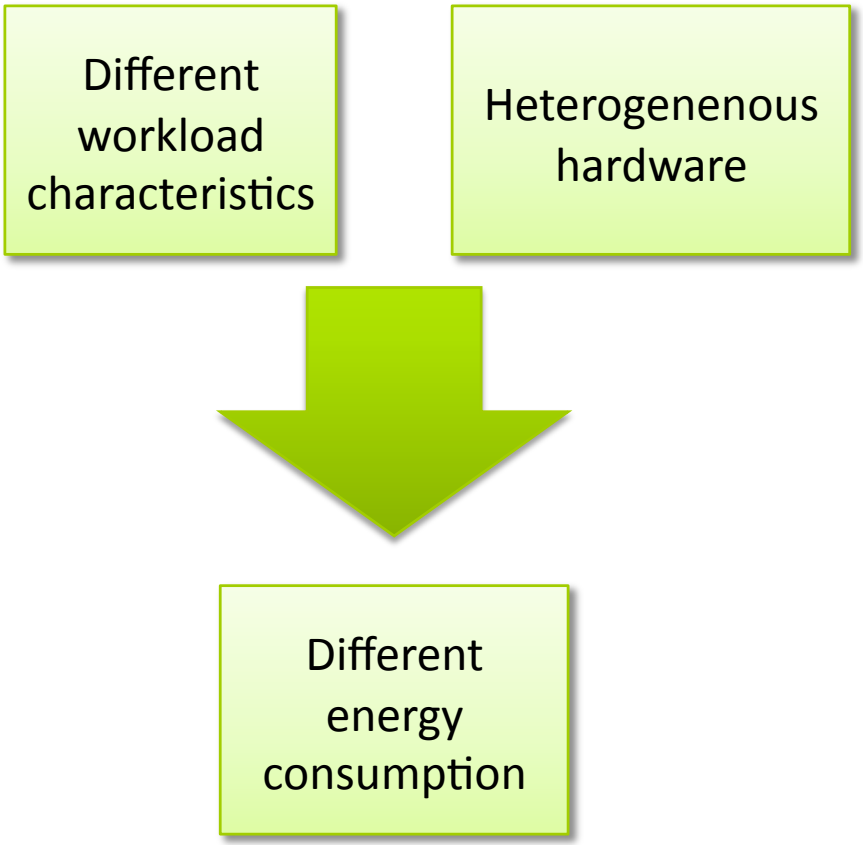
Different workload characteristics
Heterogenenous hardware
Different energy consumption

Conclusion

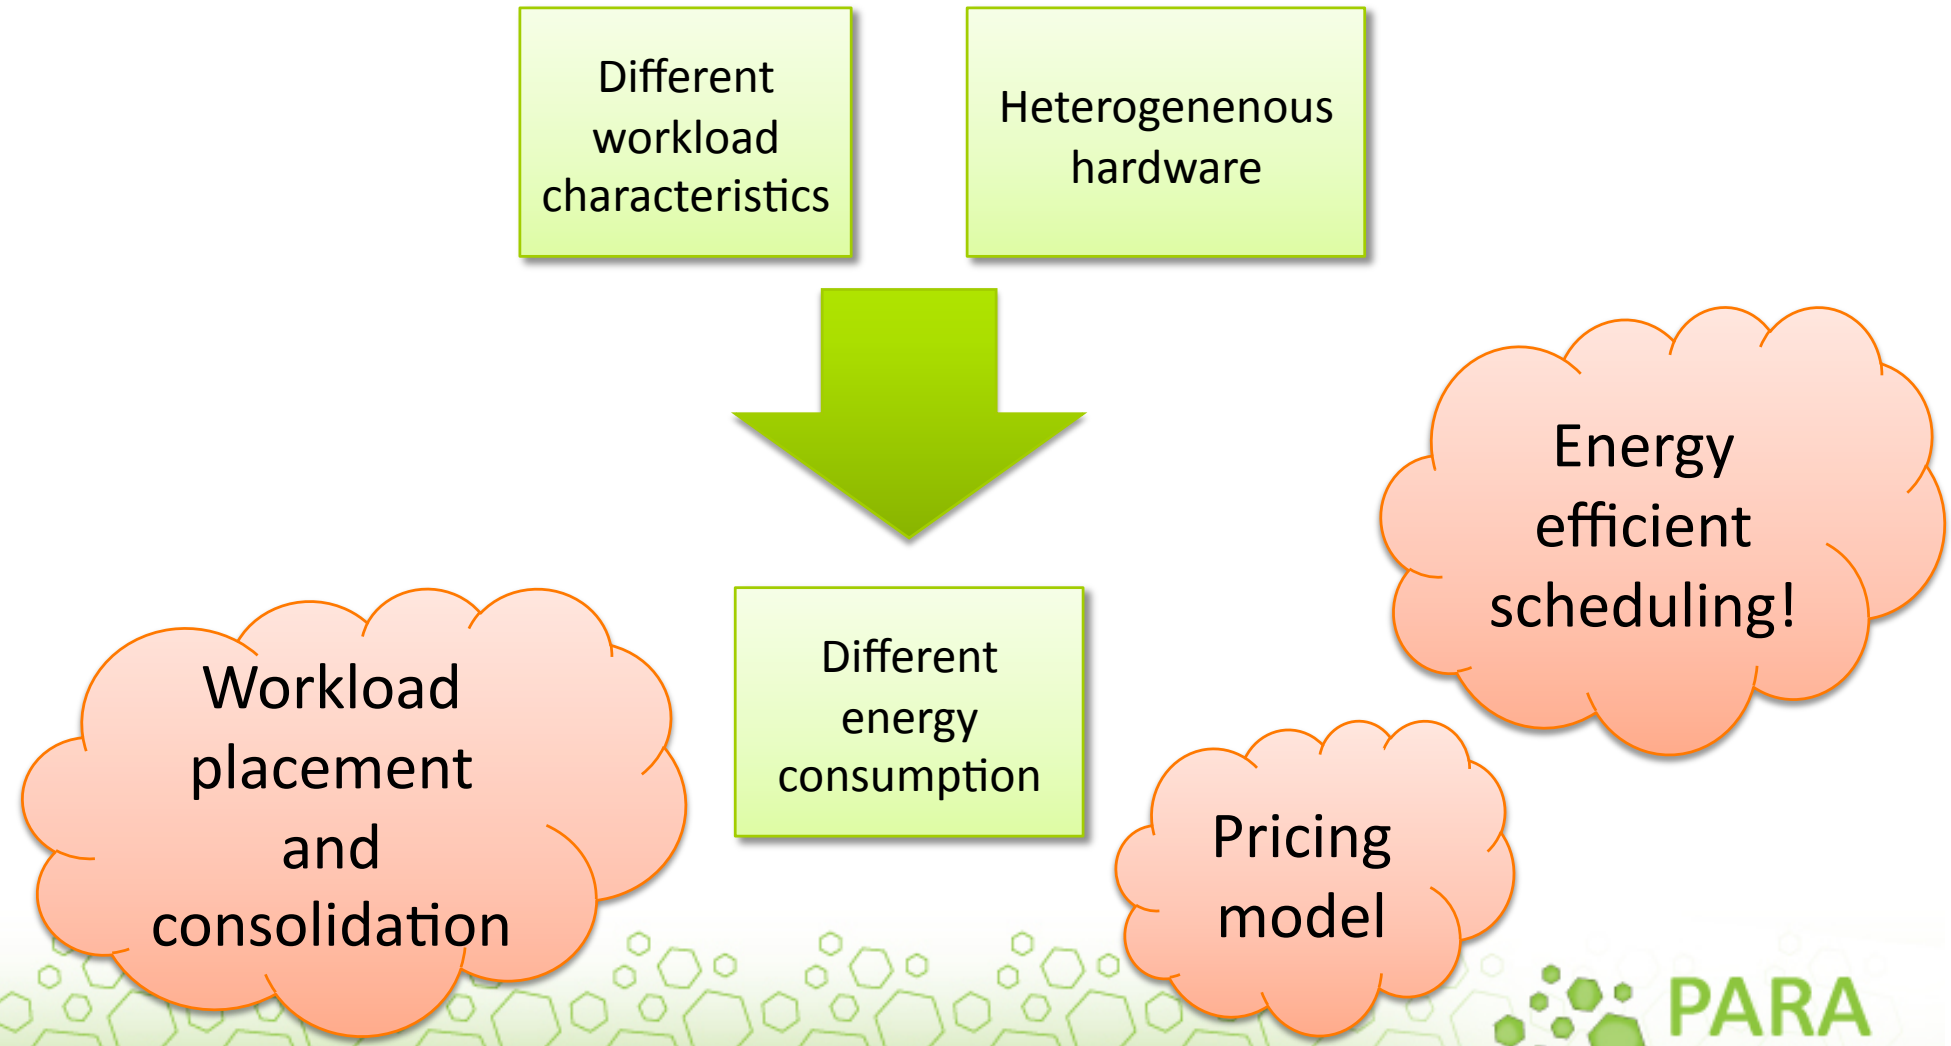

PARA DIME

PARA DIME

Power Characterization of Servers in Heterogeneous Cloud-Environments

Mascha Kurpicz, Université de Neuchâtel

This project and the research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement n° 318693