POZNAŃ SUPERCOMPUTING AND NETWORKING CENTER

Energy efficiency in HPC

Radosław Januszewski
Poznań Supercomputing and Networking Center

The motivation top 500 2009

- Roadrunner
  - 129600 IBM Cell and AMD Opteron cores
- Jaguar
  - 150152 AMD Opteron cores
- Nehalem cluster
  - 26304 Intel Nehalem cores

The motivation

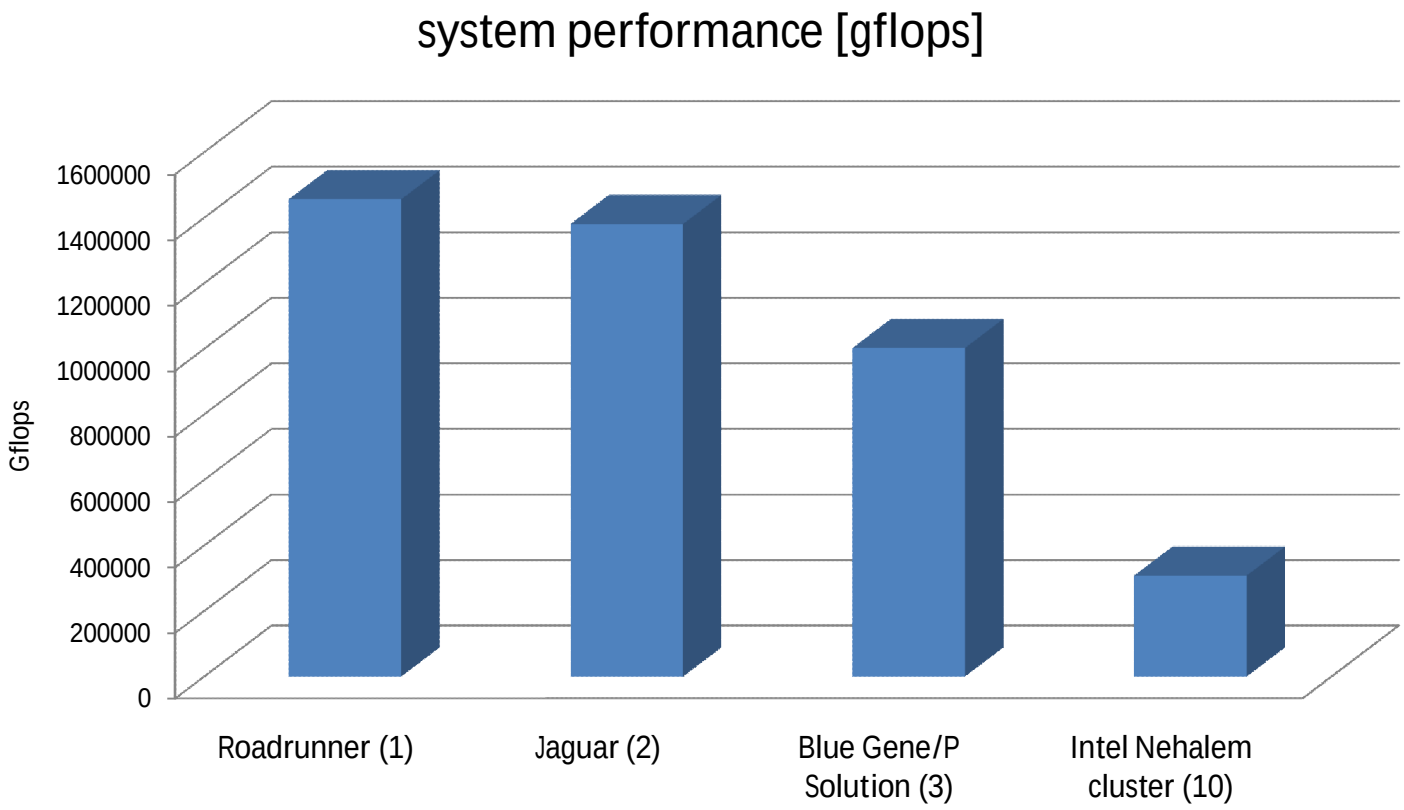

The motivation

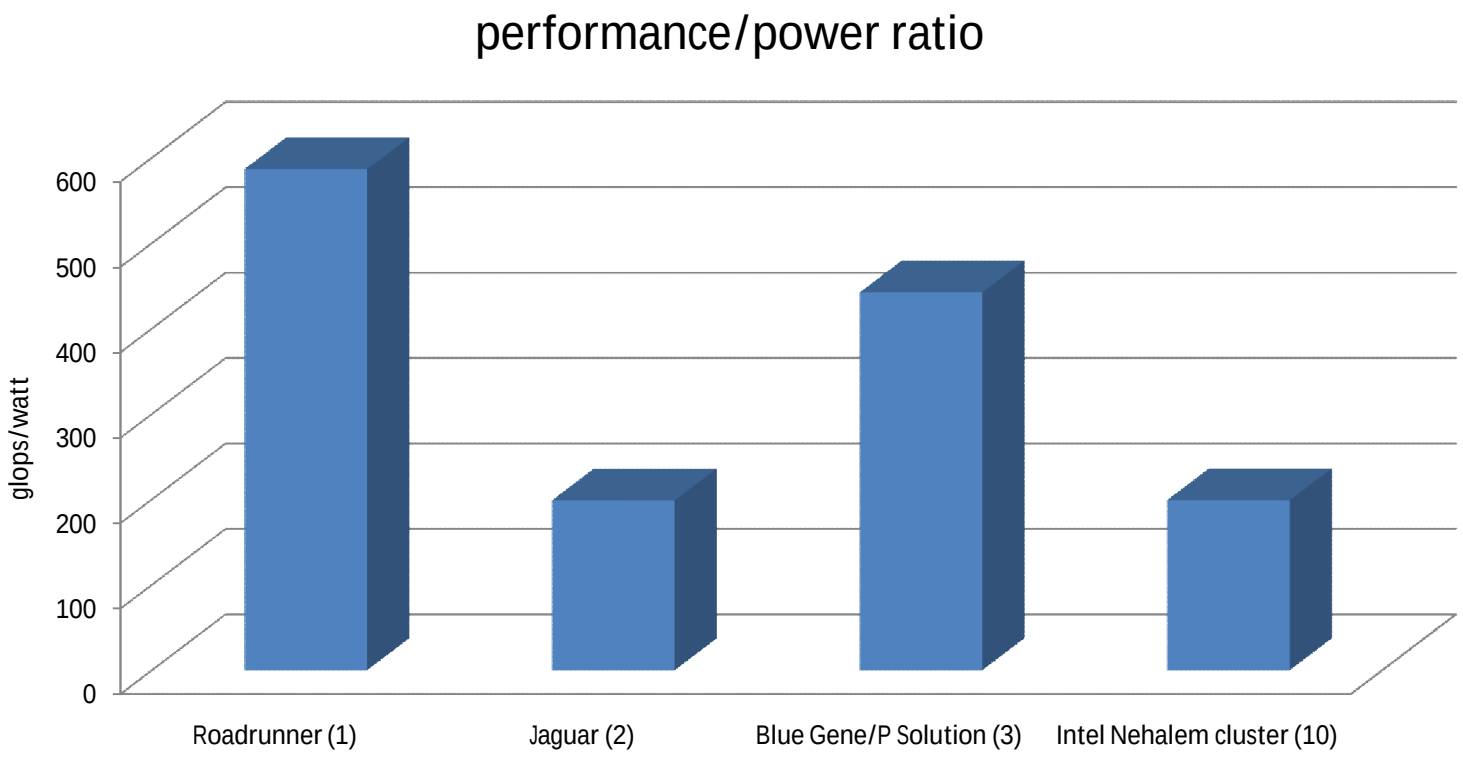


PRACE
What is using all that power?

How it was done

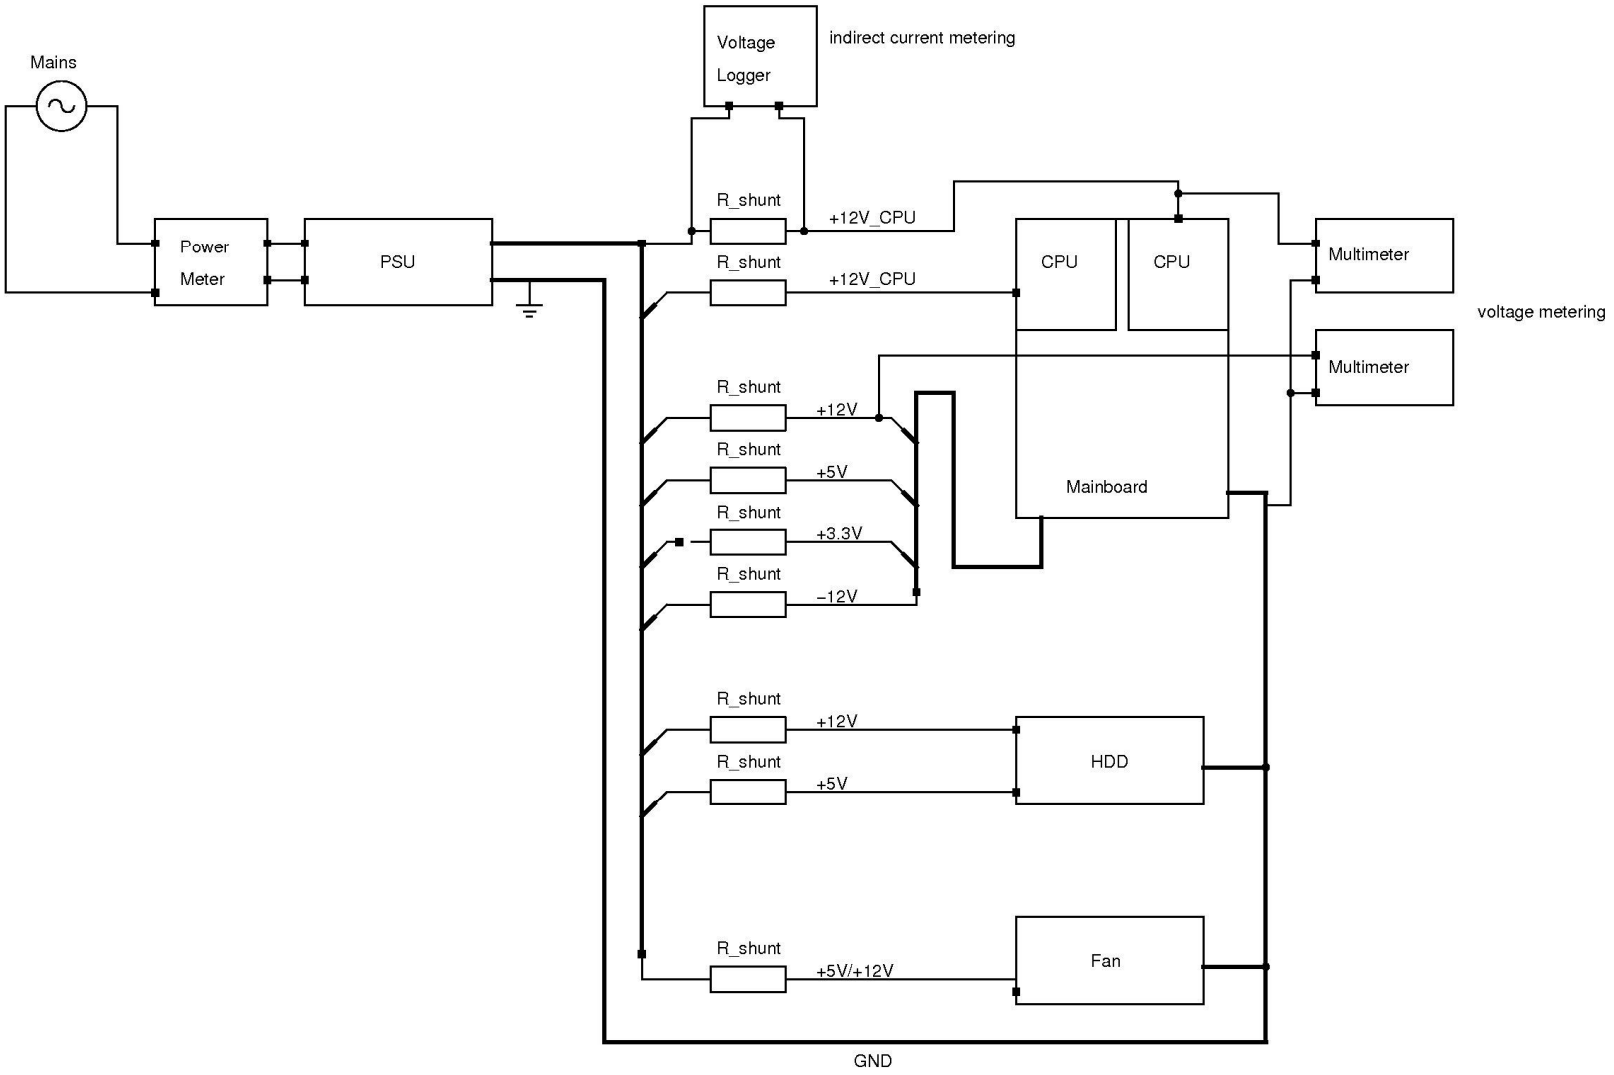

How it was done

- high resolution power meter Lutron DW-6090

- multi-channel data logger Yokogawa XL-100

- Multimeter PICOTEST M3500A

- additional equipment such as current shunt resistors

How it was done

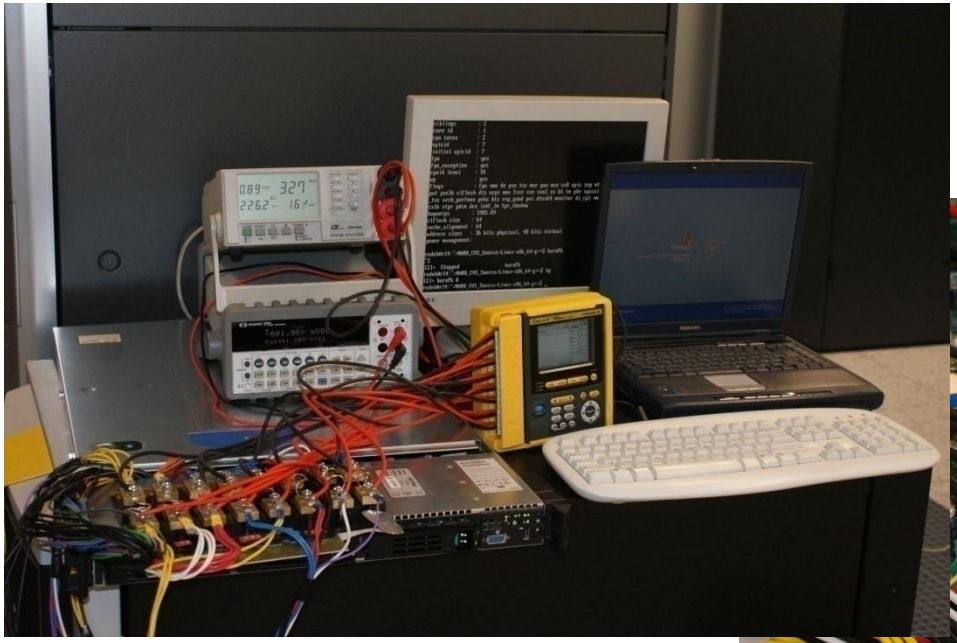

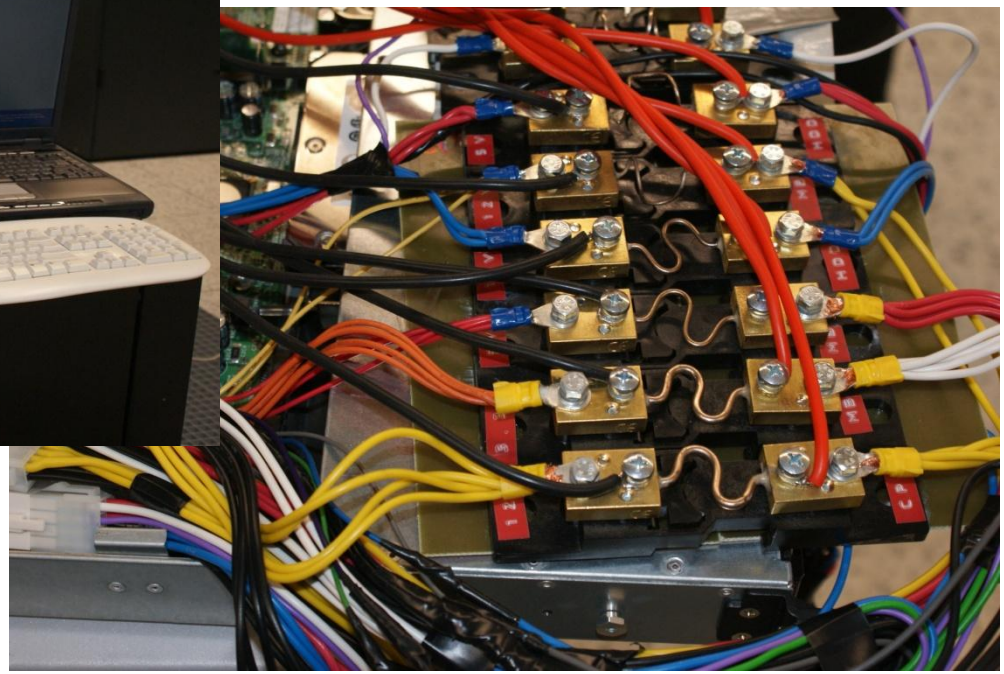

For each server a customized wiring and shunt configuration

The results

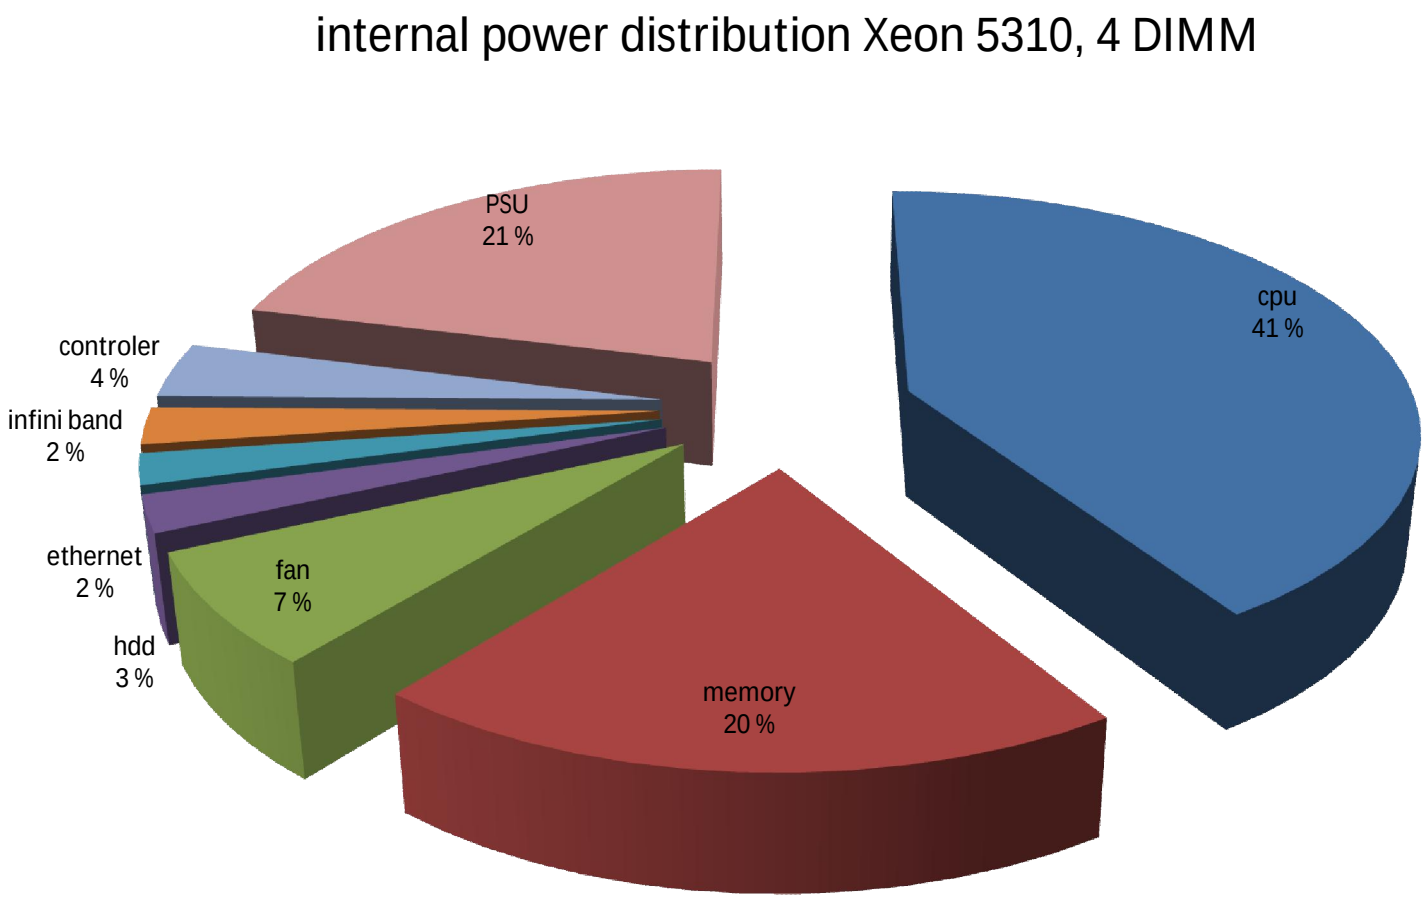
internal power distribution Xeon 5310, 4 DIMM
PSU
21 %
cpu
41 %
controler
4 %
infini band
2 %
ethernet
2 %
hdd
3 %
fan
7 %
memory
20 %

The results

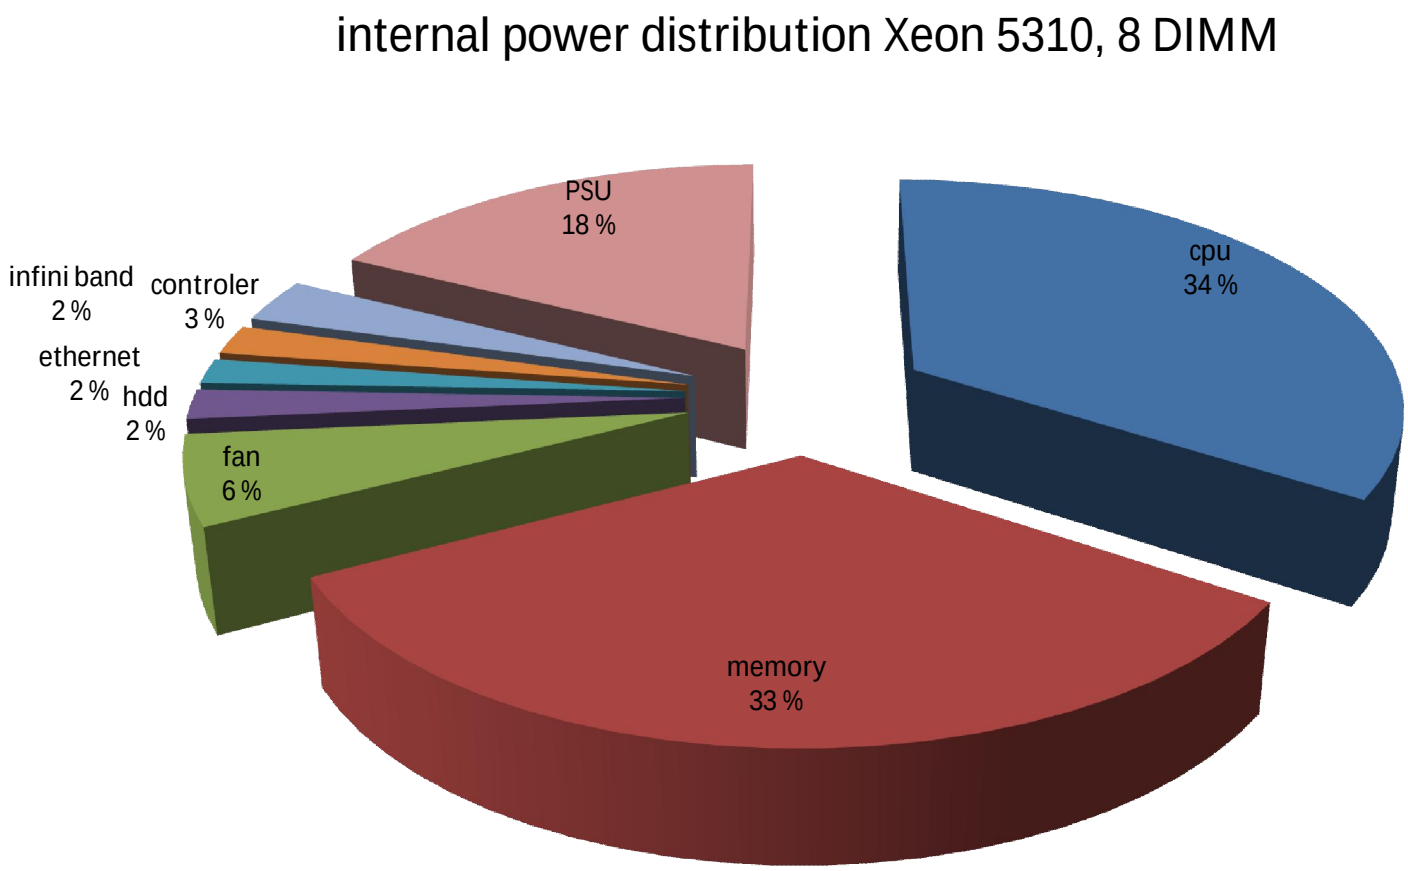
internal power distribution Xeon 5310, 8 DIMM
PSU
18 %
cpu
34 %
infini band
2 %
controler
3 %
ethernet
2 %
hdd
2 %
fan
6 %
memory
33 %

Power distribution (AMD Istanbul blade system)

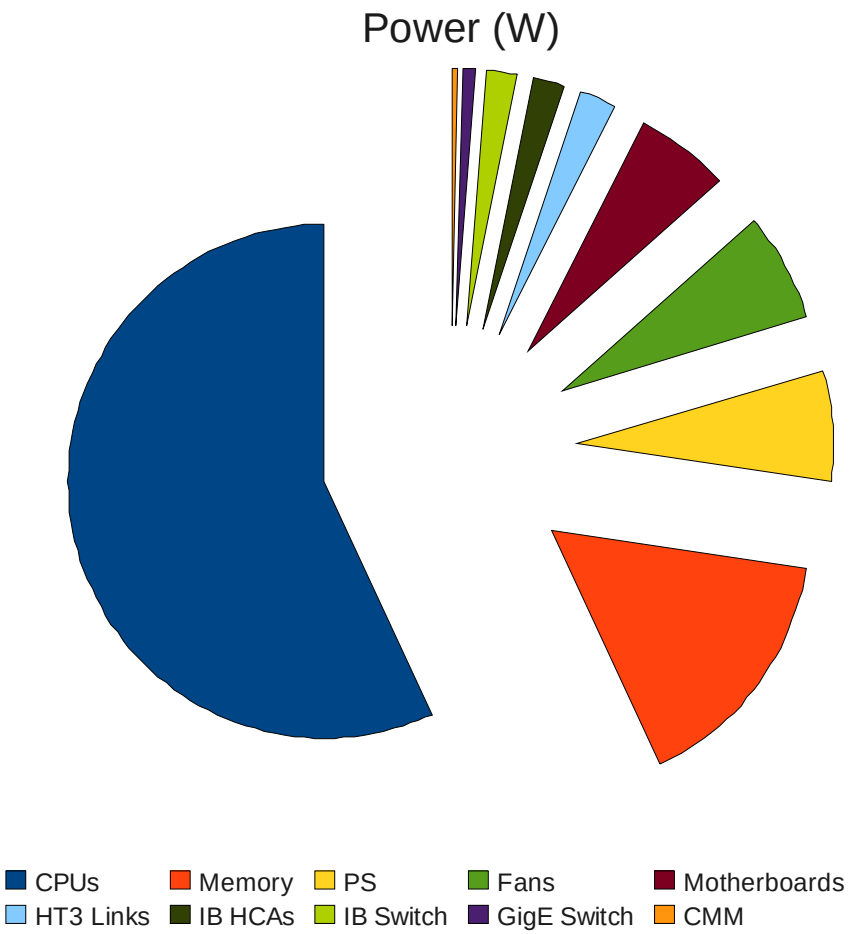

Blade measurements

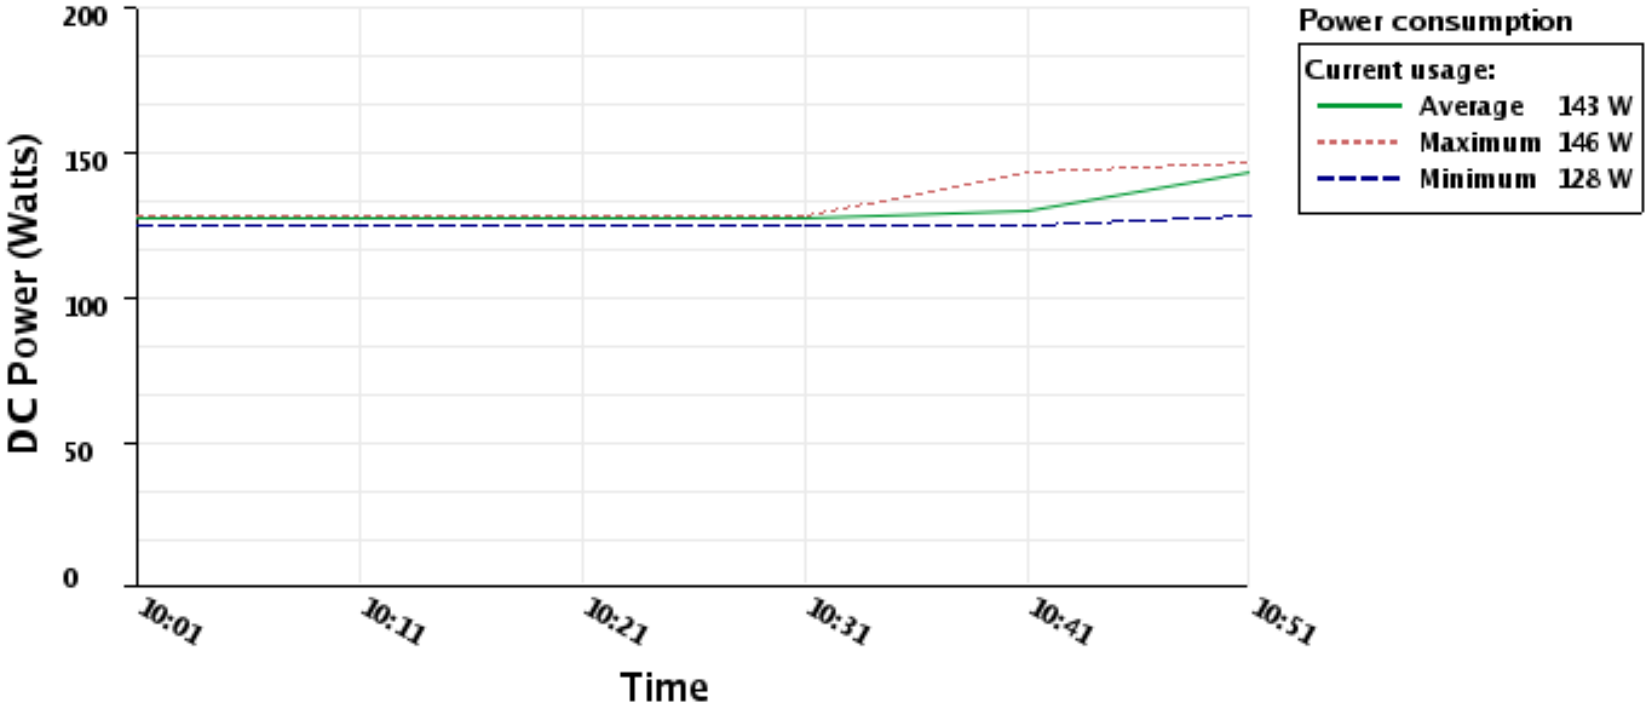

Overall power consumption

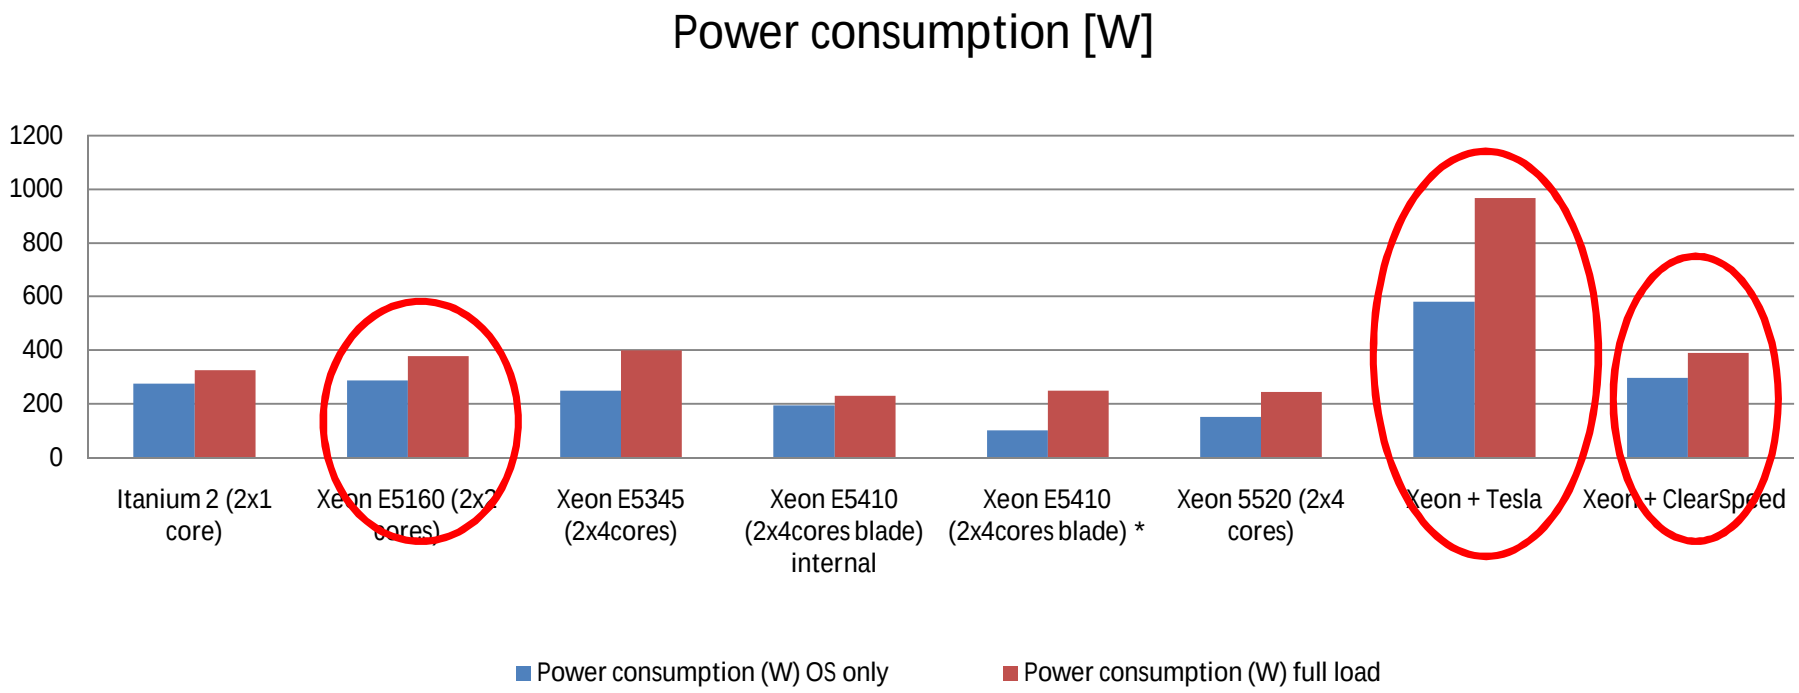

Power efficiency

NAMD appoa2 benchmark

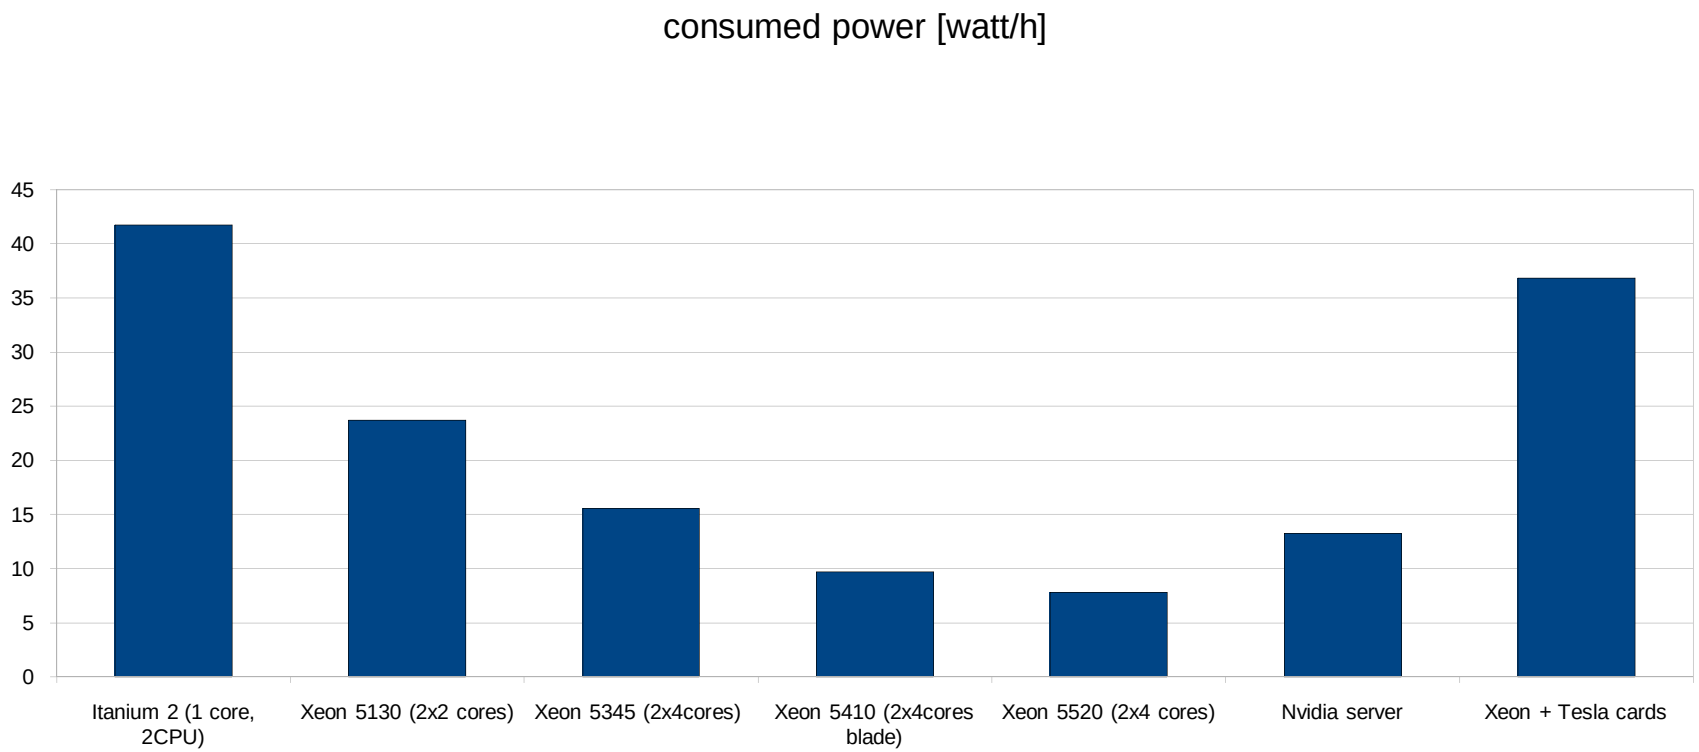

Turn it off...

Switched off machines still consume energy!

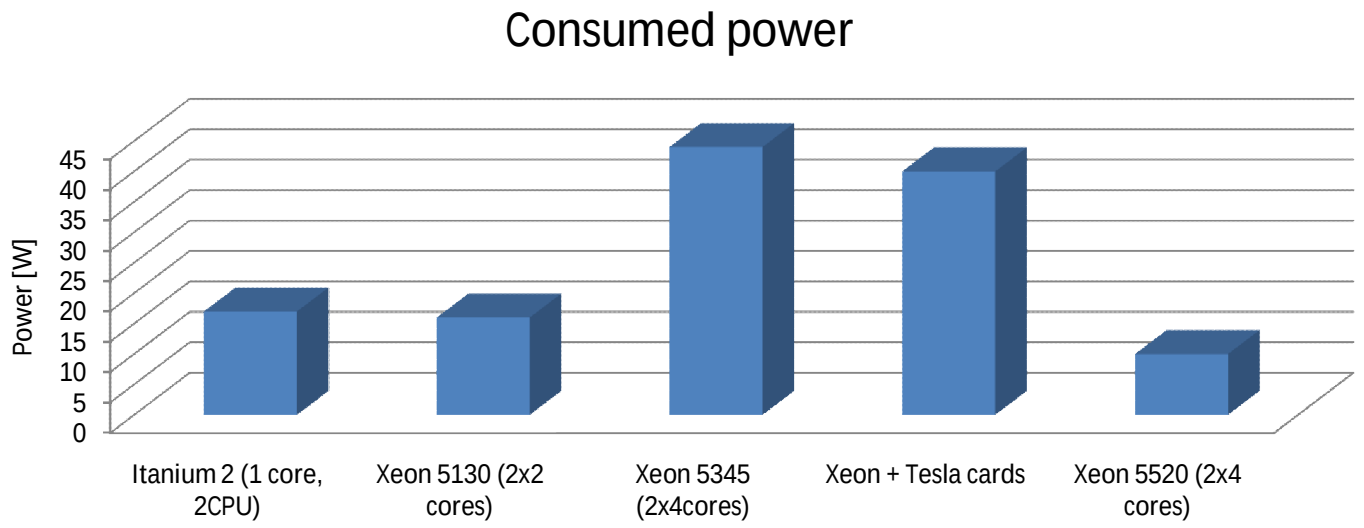

Accelerators and GPU – a new hope?

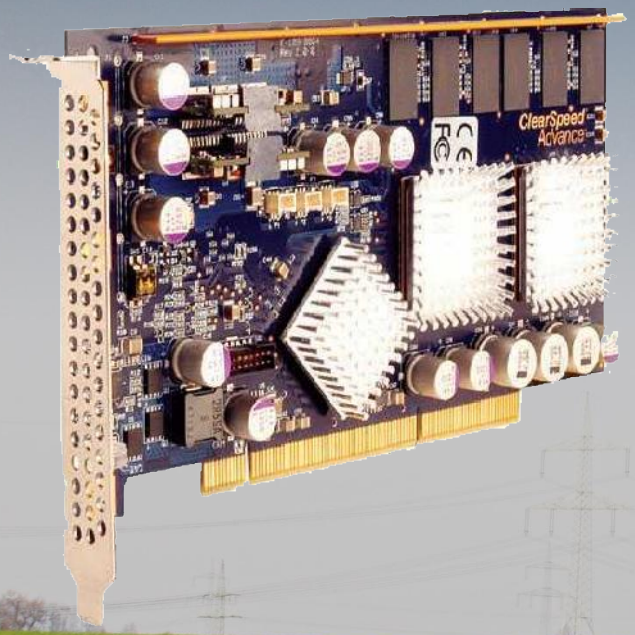

Accelerators

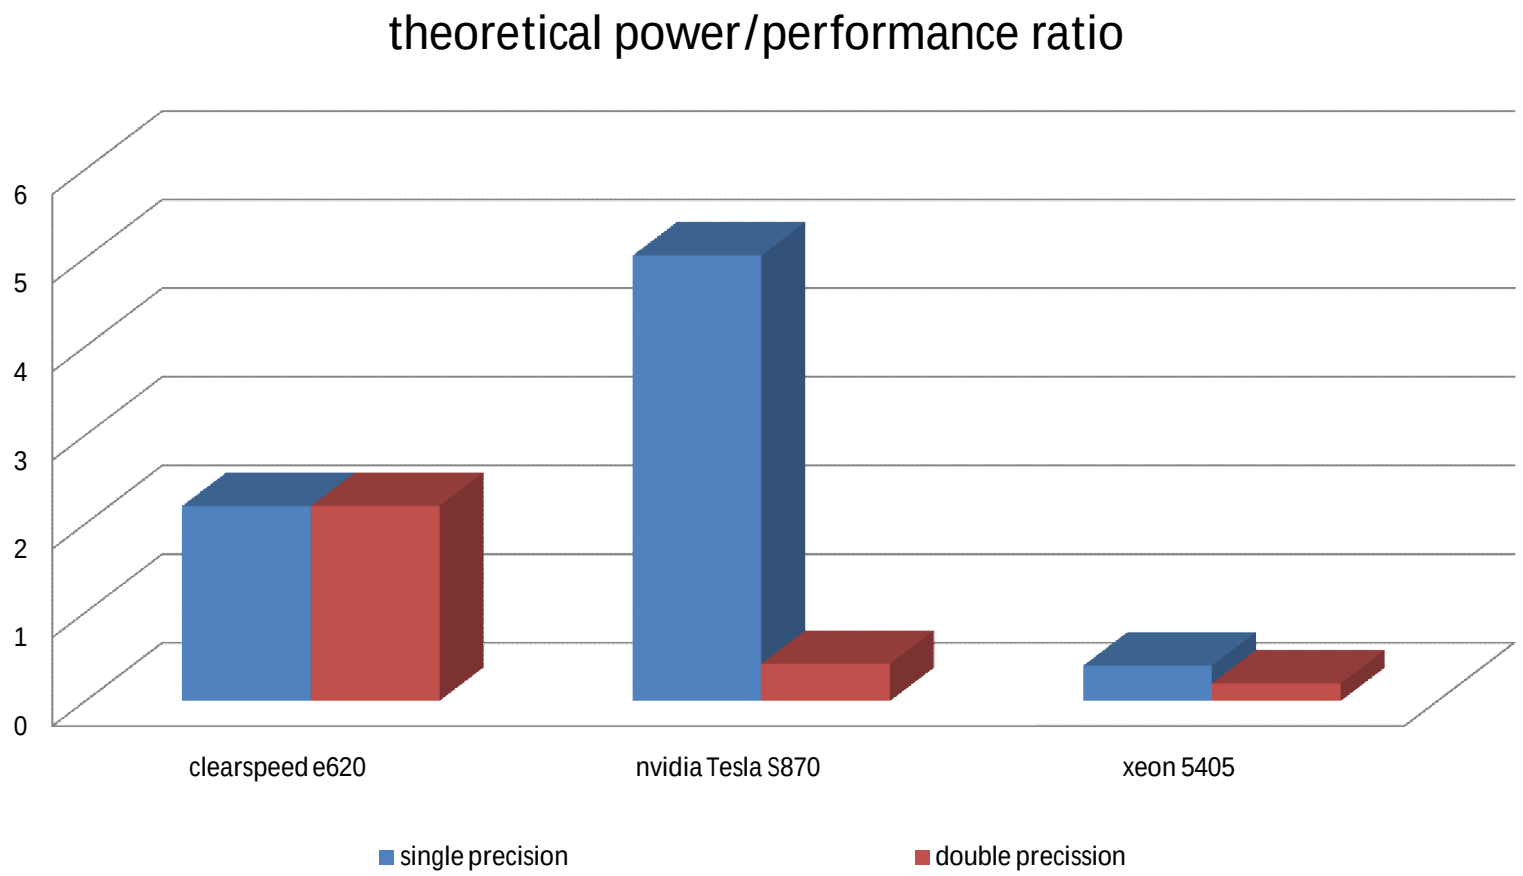

Tesla synthetic benchmarks

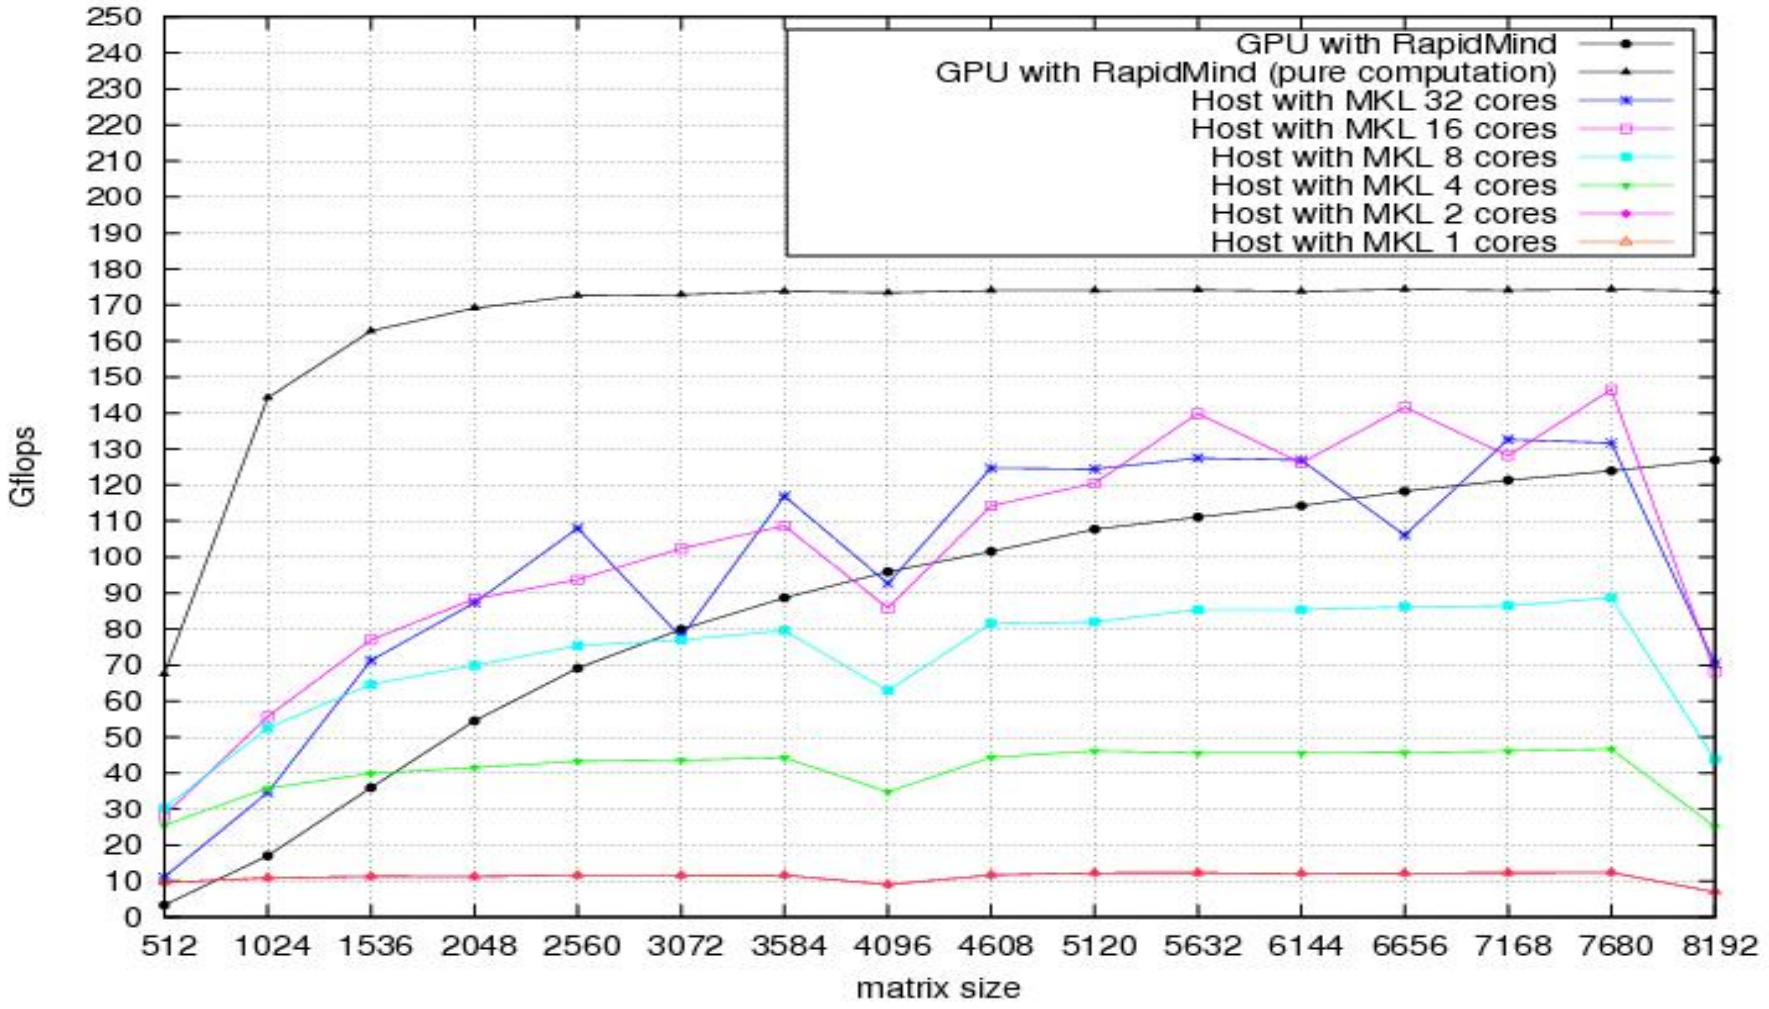

Accelerators
Nvidia

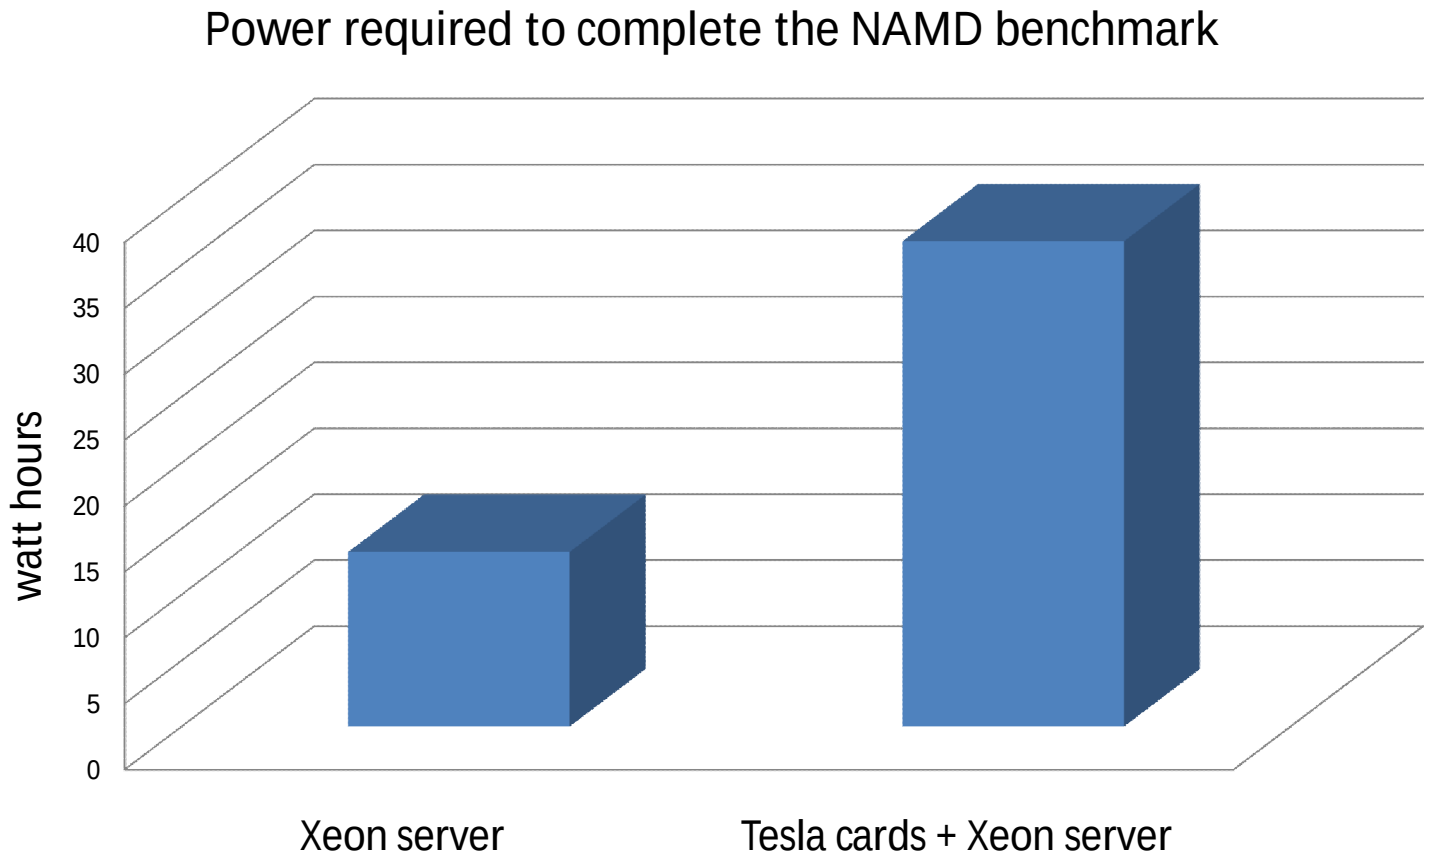

Accelerators
power overhead

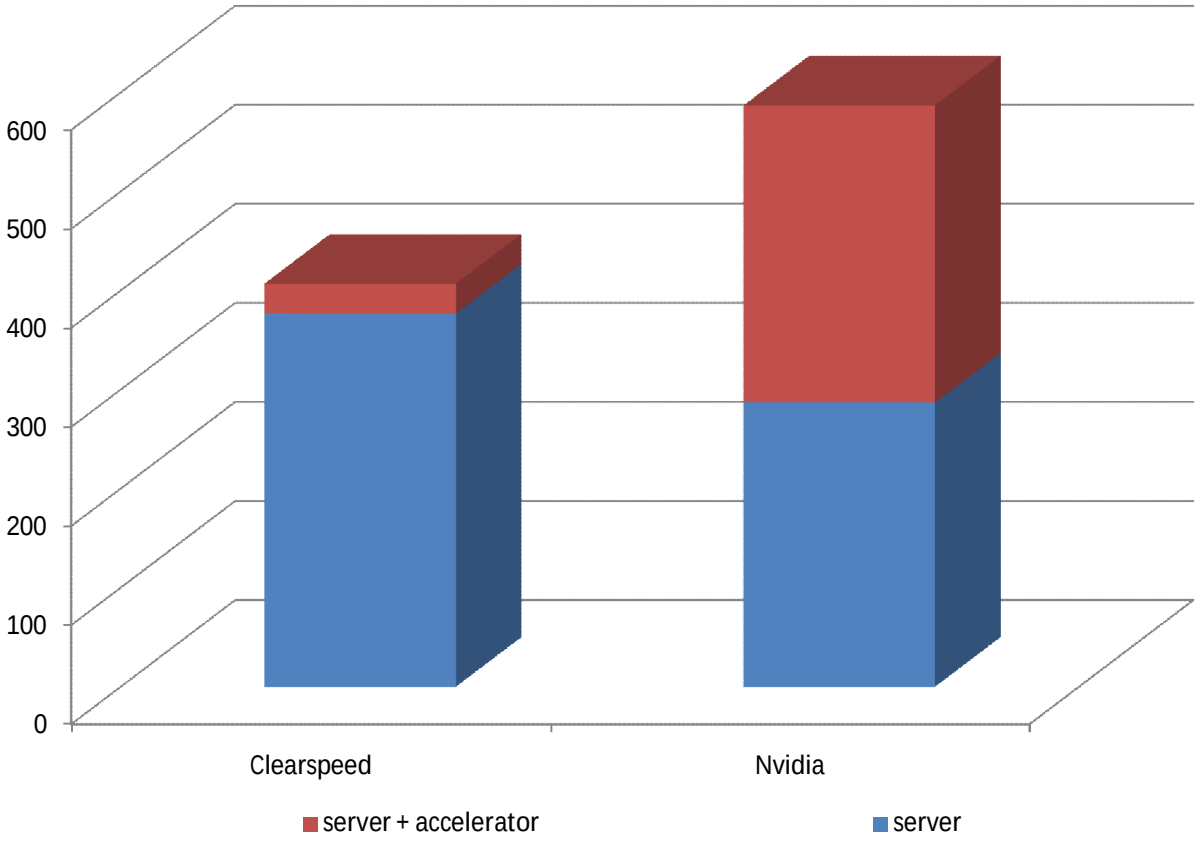

What was measured

Sicortex 1458

- 243 CPUs (MIPS 66Mhz) 2W
- 1458 cores
- Kautz graph inteconnect
- 1,9 TB memory

- 2,1 Tflops
- less than 5 kW consumed power

Sicortex

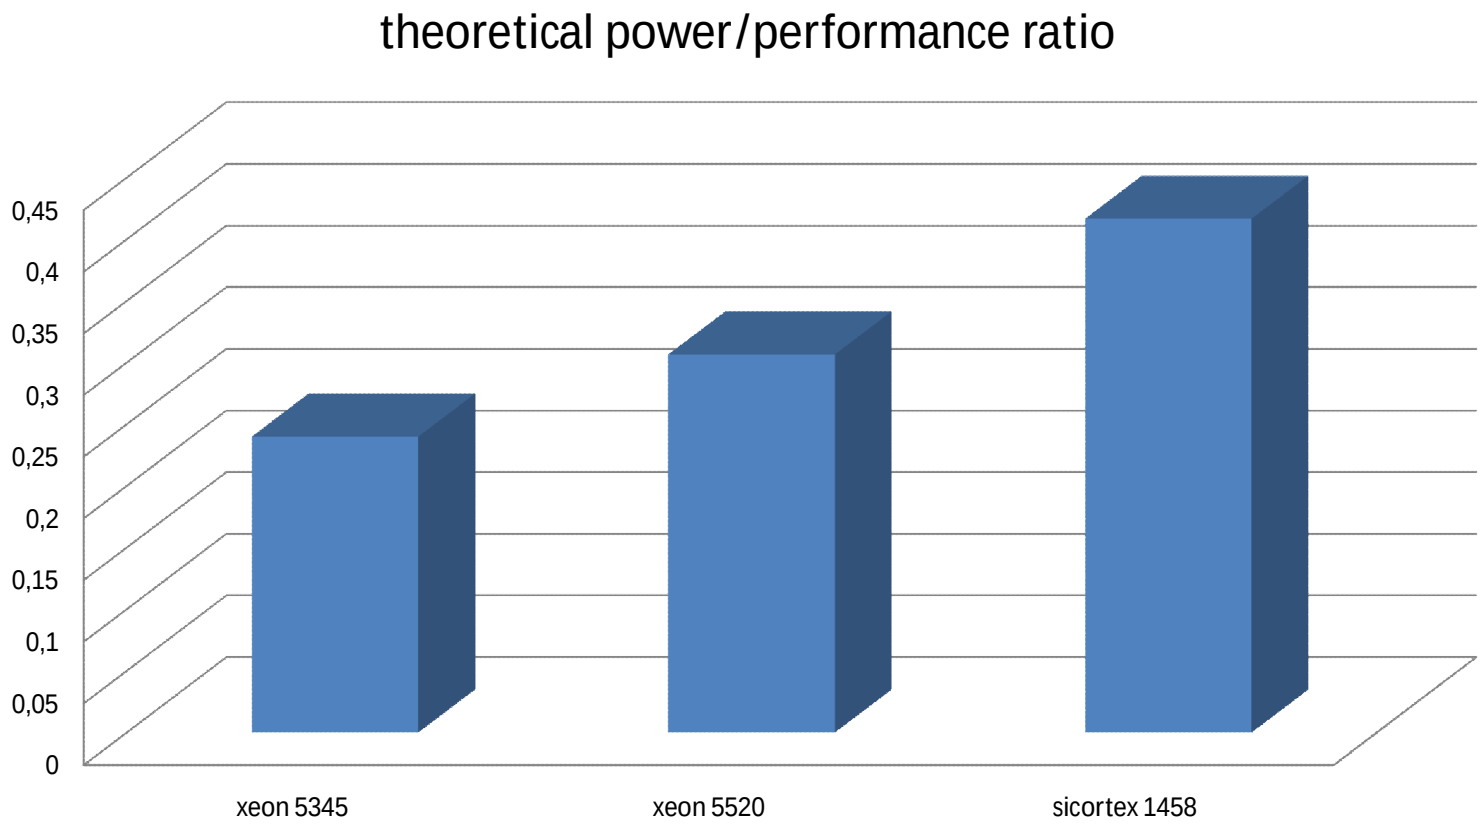
theoretical power/performance ratio
0,45
0,4
0,35
0,3
0,25
0,2
0,15
0,1
0,05
0
xeon 5345
xeon 5520
sicortex 1458

Sicortex

tests

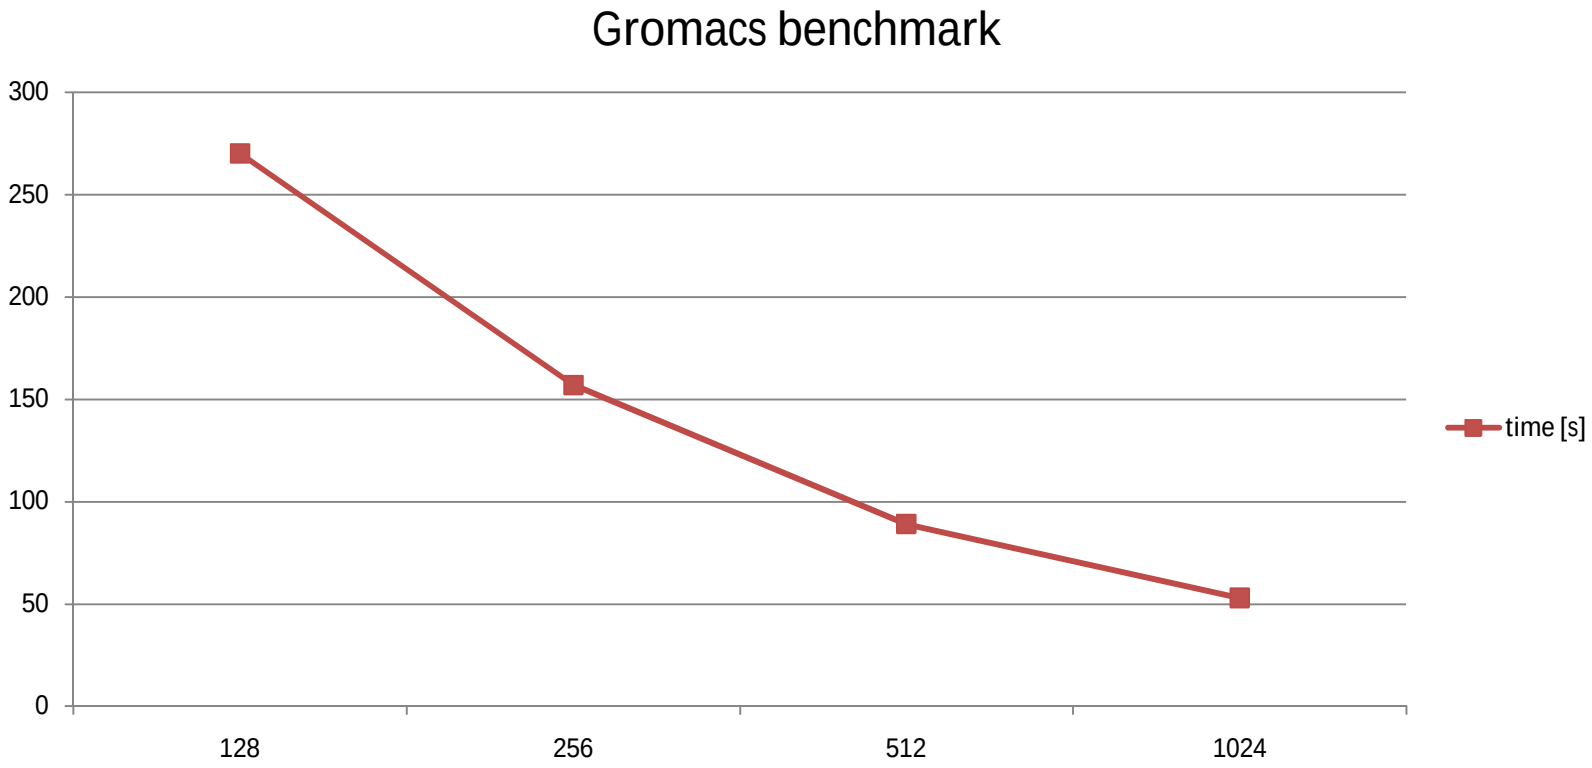

Software is crucial

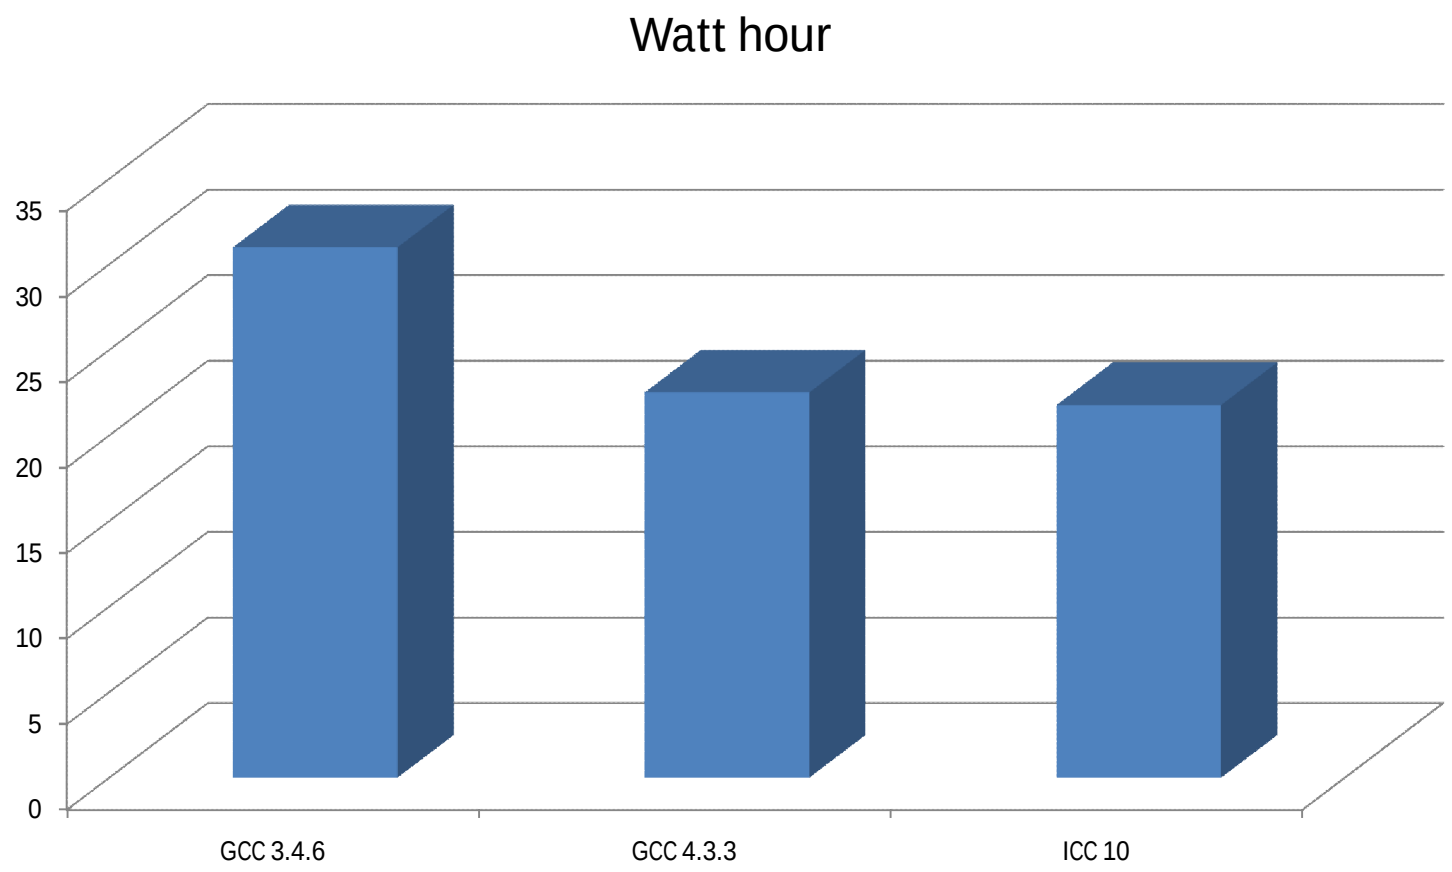

conclusions

- Physical form of the servers can be as crucial as the installed hardware (CPU)
- Being green is at the cost…
- It is hard to have efficient „good for everything” system. There ARE solutions that are better in terms of energy efficiency than x86 (commodity based) servers
- Software support is crucial
  - the accelerators are great but are green only when on 100% load.
  - getting maximum of the hardware is sometimes hard due to lack of appropriate software

The future

- In short time scale power efficiency can be improved by simple steps.
  - technology
  - funding!
- More efficient programs - non-commercial approach...
- new power management policies
  - power aware schedulers?
  - power aware software?

POZNAŃ
SUPERCOMPUTING AND
NETWORKING
CENTER

Thank you!

Radoslaw.januszewski@man.poznan.pl