Proportional Computing

Georges Da Costa

Georges.Da-Costa@irit.fr IRIT, Toulouse University

EXTREMEGREEN

Workshop
May, 13th 2013

Plan

1 Context

2 Proportional computing

3 Experiments

4 Conclusion & perspective

Small is beautiful

What is the status of energy-efficiency in data-centers?

- Currently, lots of work is done on consolidation
- Metric are evaluated at full charge
- For all usages ?
- Not for:
  - Irregular workload
  - Low workload
- Why is it important ?
  - Because we are not in a linear world !
  - Even for really efficient systems (at full charge)

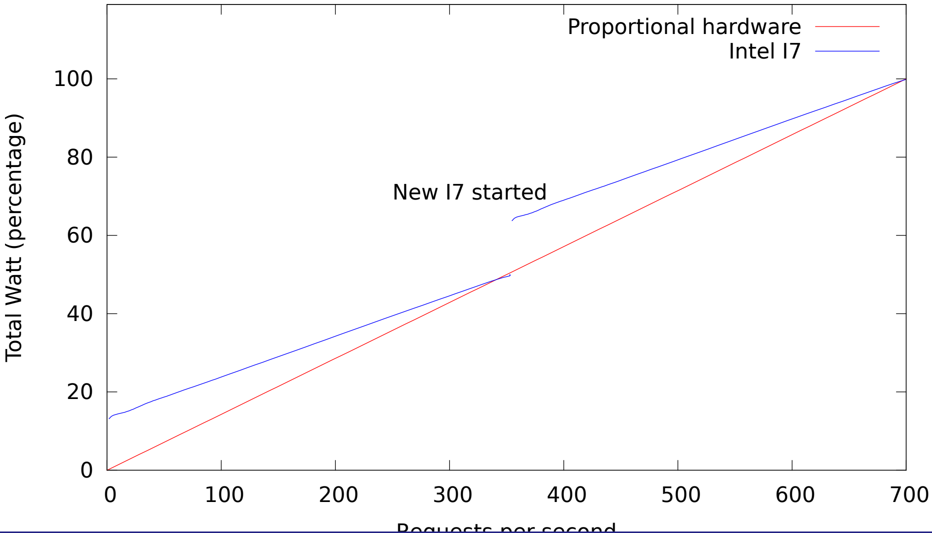
Proportional hardware
Intel I7
New I7 started
Total Watt (percentage)
0
20
40
60
80
100
0
100
200
300
400
500
600
700
Requests per second

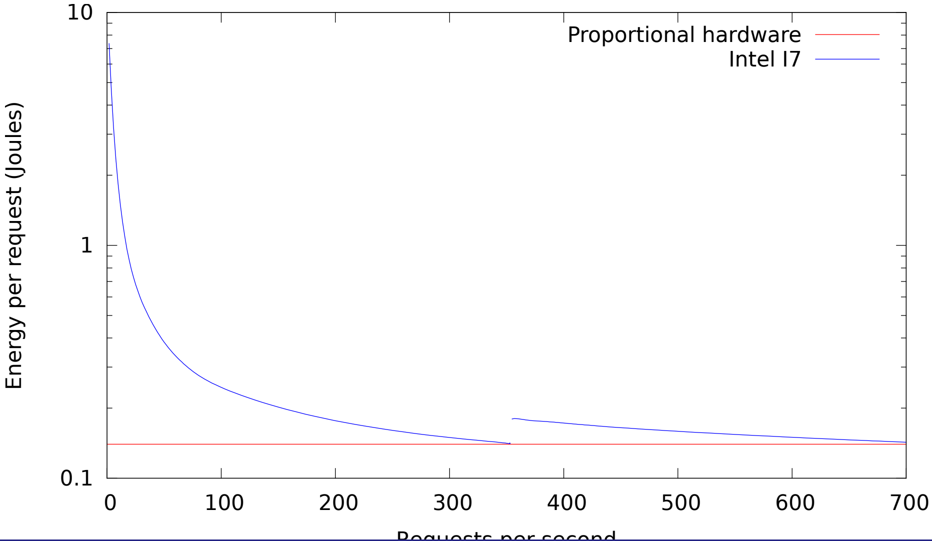
10
1
0.1
Energy per request (Joules)
Proportional hardware
Intel I7
0 100 200 300 400 500 600 700
Requests per second

Current methods

- Current methods:
  - On high load, consolidation.
  - On high number of requests, overhead is spread on lots of nodes
  - But wasted Watts continue to add-up
- What do we want ?
  - proportional computing
  - idle load = 0W

Plan

1 Context

2 Proportional computing

3 Experiments

4 Conclusion & perspective

The big lie of homogeneous systems

- Most current approach : Homogeneous systems
- True in HPC centers
- But data centers are constructed step by step
  - Historical building of DC $\rightarrow$ quite different hardware
  - Position in a data center : distance to cooling infrastructure
- But also, compare:
  - Very energy efficient systems at low load, high overhead
  - Energy-inefficient system at max load, low overhead

State of the art of proportional computing

- Dark Silicon
  - Ongoing research
  - Mostly on processor
  - Switch off unused processors units
- Heterogeneous on-die cores
  - Big.LITTLE ARM : Cortex A7 + Cortex A15, 20μs migration
  - NVIDIA Optimus : CPU-integrated GPU + Full-fleged GPU, 1/5th frame migration
- Same problems
  - Motherboard facilities (bus, network,...) always on and less dynamic

Our idea

Use nodes depending on the real load, not the peak load

Example:

Processor	Watt range	Max request/s	Efficiency (W/r)
Intel I7	11 - 42	353	.12
Intel Atom	8 - 9	34	.26
Raspberry Pi	2.56 - 2.81	5.6	.50

Intuition: Several small node and intermediary nodes to have a multi-scale smooth curve

Zoom on Intel I7

Zoom on Intel Atom

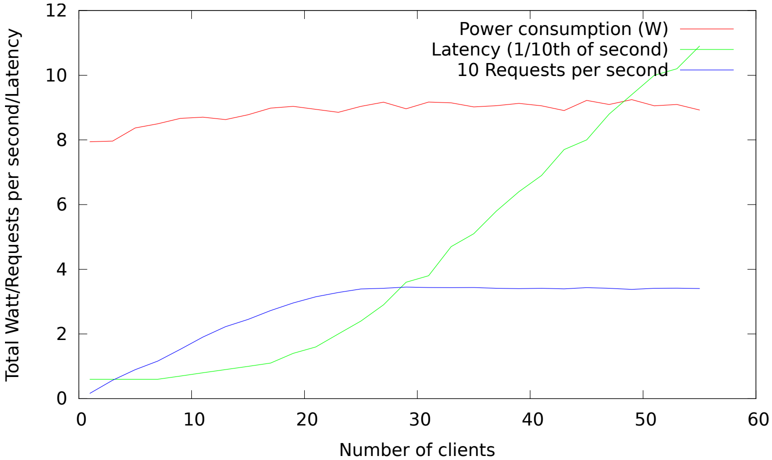

Comparison

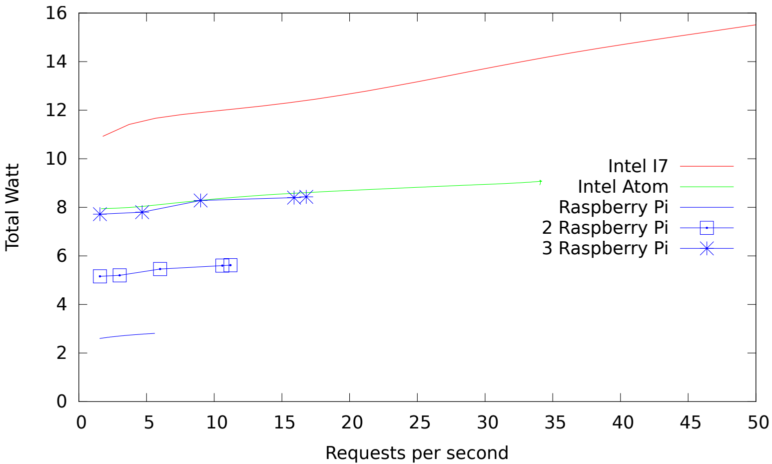
Total Watt
Requests per second
0
2
4
6
8
10
12
14
16
0
5
10
15
20
25
30
35
40
45
50
Intel I7
Intel Atom
Raspberry Pi
2 Raspberry Pi
3 Raspberry Pi

A near-linear behavior

Take the most efficient hardware as function of workload:

0→5 req/s	1 Raspberry Pi
5→10 req/s	2 Raspberry Pi
10→35 req/s	1 Intel Atom
35→40 req/s	1 Intel Atom + 1 Raspberry Pi
40→350 req/s	1 Intel I7

For over 350req/s, use modulo 350, and use as many I7 as necessary

Still far far away...

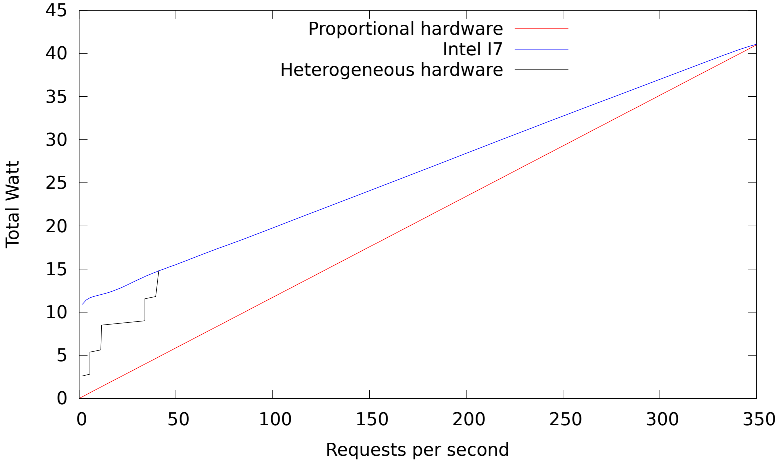

Not so far, efficiency view !

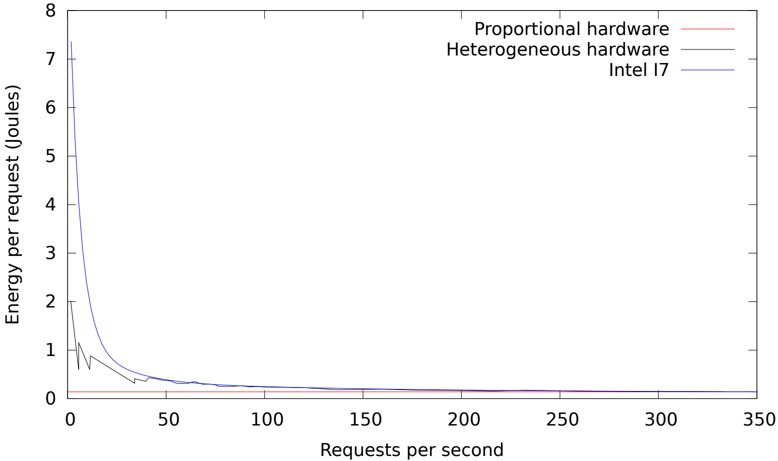

Plan

1 Context

2 Proportional computing

3 Experiments

4 Conclusion & perspective

98 Football World Cup

- Available data
- 92 days of web server access logs
- Workload precise at the second level
- Four geographic locations, three in US, one in France
- Several phases
  - Low phases, first 40 days and last 10 days
  - High phase, during the competition

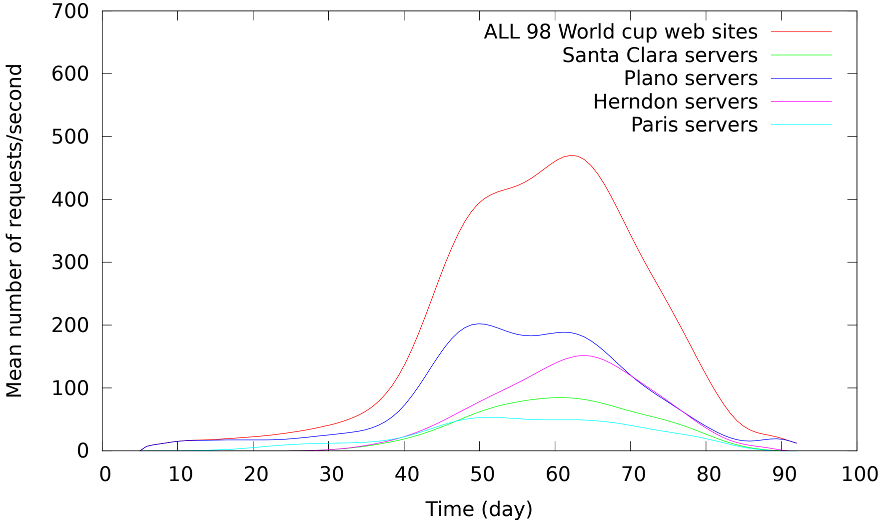

ALL 98 World cup web sites
Santa Clara servers
Plano servers
Herndon servers
Paris servers
Mean number of requests/second
Time (day)
0
100
200
300
400
500
600
700
0
10
20
30
40
50
60
70
80
90
100

Comparison if using one single data-center

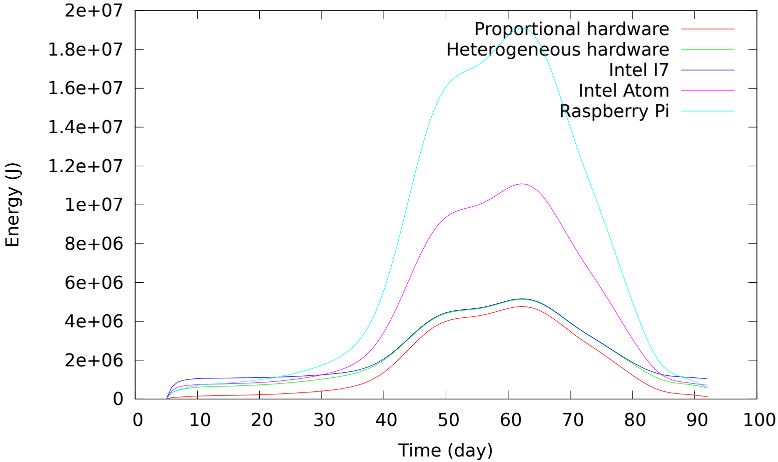

Zoom on the most efficient methods

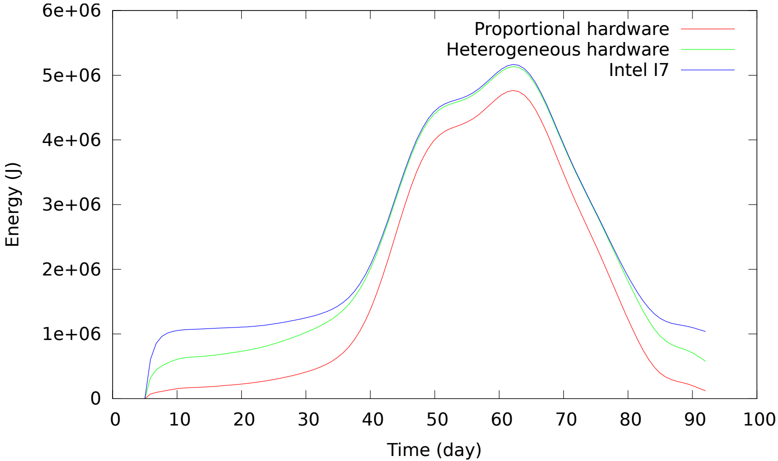

Far reached goal: proportional computing

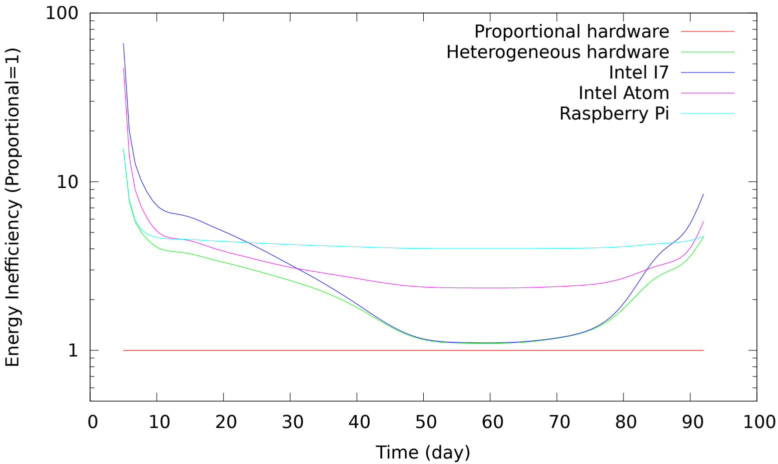

Comparison if using multiple data-centers

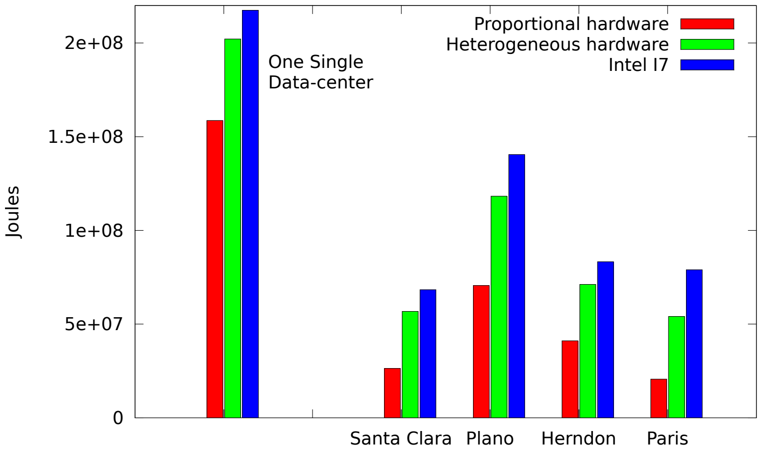

Not limited to Data-Centers

Several studies on personal computer usages:

- Average load of switched-on computers in an university[1]: 5.5%
  - 80% and 60% energy reduction for perfect and heterogeneous systems
- Between 40 and 80% computers of several computer science department are idle but switched-on[2].
  - Worst case for the proposed approach: 30% and 20% energy reduction
- Using the load distribution of a Google data-center[3]
  - Resp. 60% and 5% improvement

[1] P. Domingues, P. Marques, and L. Silva. *Resource usage of windows computer laboratories*. In International Conference Workshops on Parallel Processing (ICPP), 2005.
[2] Anurag Acharya, Guy Edjlali, and Joel H. Saltz. *The utility of exploiting idle workstations for parallel computation*. In SIGMETRICS, 1997
[3] L.A. Barroso and U. Holzle. *The case for energy-proportional computing*. Computer, 2007.

Plan

1 Context

2 Proportional computing

3 Experiments

4 Conclusion & perspective

Conclusion & perspective

- Improves efficiency several types of systems
  - Reduces the need of increasing density $\rightarrow$ reduces cooling cost
- Still inefficient most of the time
- Next steps
  - Increase range of efficiency with different hardware
  - Take into account frequencies
  - Tackle the migration problem (currently only service-oriented)
  - Save the world with Extreme Energy Savings