COST EFFICIENCY VS ENERGY EFFICIENCY

Anna Lepak
Universität Hamburg
Seminar: Energy-Efficient Programming
Wintersemester 2014/2015

TOPIC

- Cost Efficiency vs Energy Efficiency
  - How much money do we have to pay to acquire an HPC platform and to maintain it
  - in consideration of energy efficiency

OUTLINE

1. Introduction
2. HPC Platform
    1. What is it
    2. Why do we need such platforms?
3. Total Cost of Ownership
    1. TCO
    2. Lowering the cost
    3. Power Management
    4. Cooling
    5. Example: Google’s Data Center
    6. Brainware
4. Conclusion

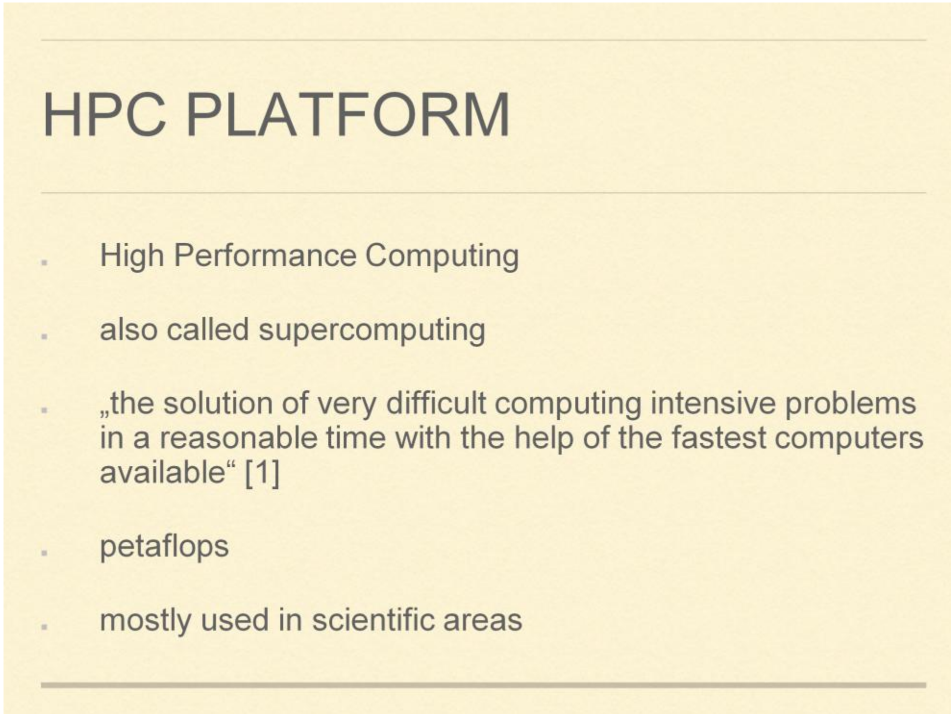

HPC stands for „High Performance Computing“; also called „supercomputing“

- quote: „the solution of very difficult computing intensive problems in a reasonable time with the help of the fastest computers available“
- - > they can calculate really fast: calculation at the speed of a nanosecond
- Measured in FLOPS: “Floating Point Operations Per Second” ; measure for computer performance
- today: PETAFLOPS (10^{15})
- Scientific areas:
- - examples: medicine (epidemics, pandemics), physics, climate research

Not only in scientific areas, but also in modeling new cars and airplanes: less costly to simulate a new model then making new prototypes

TOTAL COST OF OWNERSHIP

- building and maintaining an datacenter is expensive
- -> engineers need to focus on cost efficiency

- Total Cost of Ownership (TCO):
  - the money that is spend during a lifetime of a HPC platform

With higher computing power come more costs
So we need cost efficiency

TOTAL COST OF OWNERSHIP

- Investment costs (to acquire an HPC Platform)
  - hardware (servers, storage, cooling systems, cabling, network, …)
  - software
  - datacenter construction
- Operational costs
  - Energy efficiency
  - Personnel („brainware“ pays off, more to it later)
  - Maintenance

TOTAL COST OF OWNERSHIP

- higher computing performance -> higher energy consumption
- energy costs have become a contributor to TCO
- Green500 list : reflects computing efficiency (not raw computing power)

TOTAL COST OF OWNERSHIP

- Investments worldwide:
  - 2011: $ 27 billion
  - 2012: $ 29 billion
  - 2017: $ 40 billion

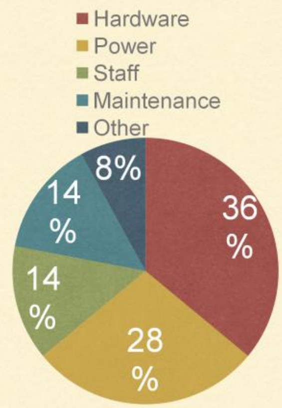

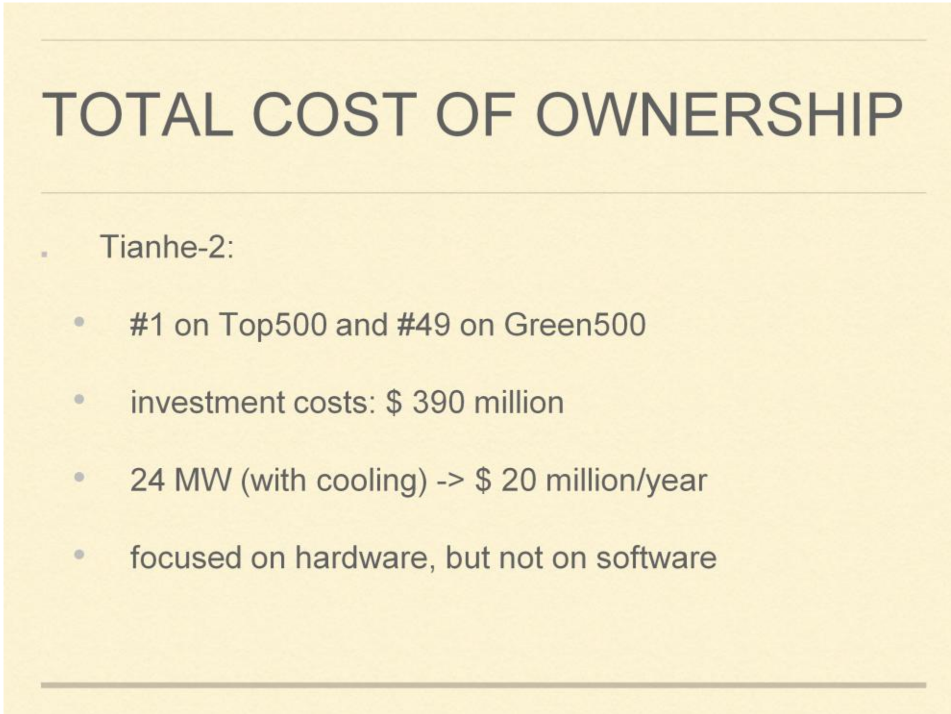

Tianhe-2 is chinese and means: milkyway-2

- #1 on TOP500 list
- #49 on Green500 list
- Built for $390 million by National University of Defence Technology
- 24 MW (with cooling) => $20 million a year ($65000 – $100000 a day)
- 34 petaflops
- World‘s fastest computing system, but isn‘t used to it‘s fullest capacity (lack of good software)
- => too much focus on good hardware, not software

LOWERING THE COSTS

- better planning of the whole project
  - what kind of software do we need
  - is the supercomputer too powerful for the problem/simulation it is designed for?
  - lowering the power consumption -> lowers costs

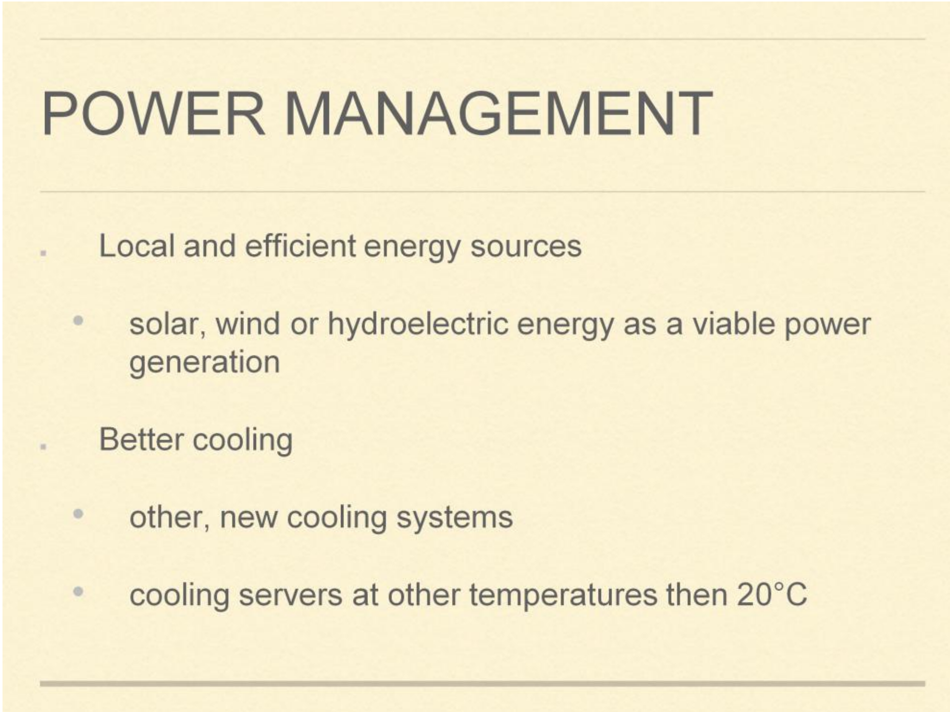

Local and efficient energy sources

- Renewable energy -> saves much money and is economically good
- Not always possible (good climate for datacenter <-> bad climate for renewable energy)

Better cooling:

- Use different, new cooling systems
- Most of energy is spent on cooling
- Most datacenters use cooling which requires much energy
- Free cooling?
- Raising aisle temp tp 27° C
- - usually chilly at 20°C
- BUT: no server or network equipment needs temperatures at 20°C.
- Can still run efficiently, w/o failure at 27°C
- Improves efficiency of chillers
- -> allows higher chilled water temp -> reduced runtime of chillers

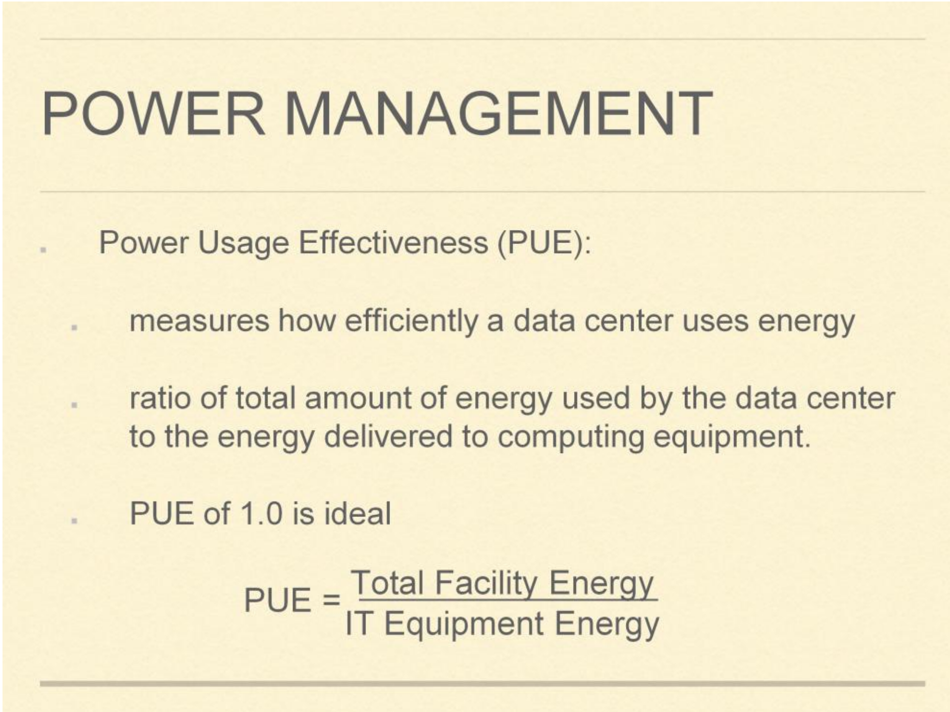

PUE

- Measures efficiency of a data centers energy usage
- Ideal PUE: 1.0
- Average PUE 2 – 2.5

- Cooling energy is not IT equipment energy.
- IT Equipment Energy is only energy spent on the IT Equipment, like servers etc.

COOLING

- cooling takes much energy
- Traditional Cooling: chillers
  - cold water or liquid coolant exchange heat with the hot air
  - the hot liquid has to be cooled down to be reused -> chillers
  - removes heath via a vapor-compression

COOLING

- better way: cooling towers
  - warm water from data center flows down a tower
  - cools down mainly through evaporation
  - cheaper than chillers (free cooling)
  - in colder climates (but not too cold)

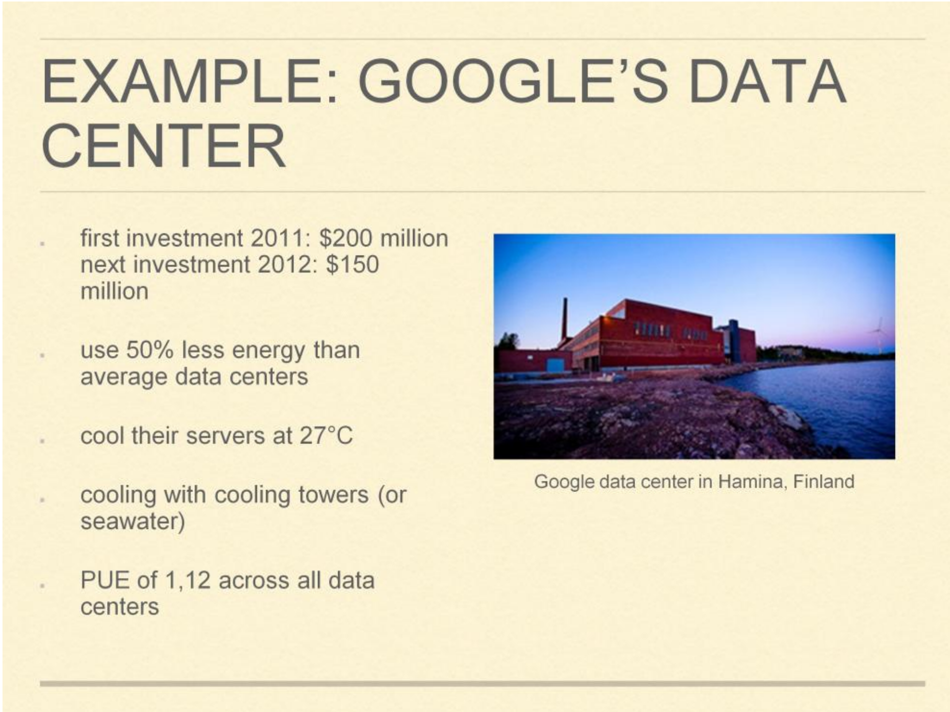

Google‘s data center are a good example for energy efficient HPC. They are always trying new ways to make their data centers even more energy efficient.

- Investements.
- Average PUE across all data centers: 1,12
- But one data center has a PUE of 1,06 (best value)
- „Standby mode“: waiting for task -> use little energy
- Buy only needed hardware
- Raising aisle temperature to 27°C
- Prevent hot air mixing with cold air
- -> curtains to seal off the cold air
- -> appropriate ducting and permanent enclosure
- Cooling system: cooling towers (or sea water)

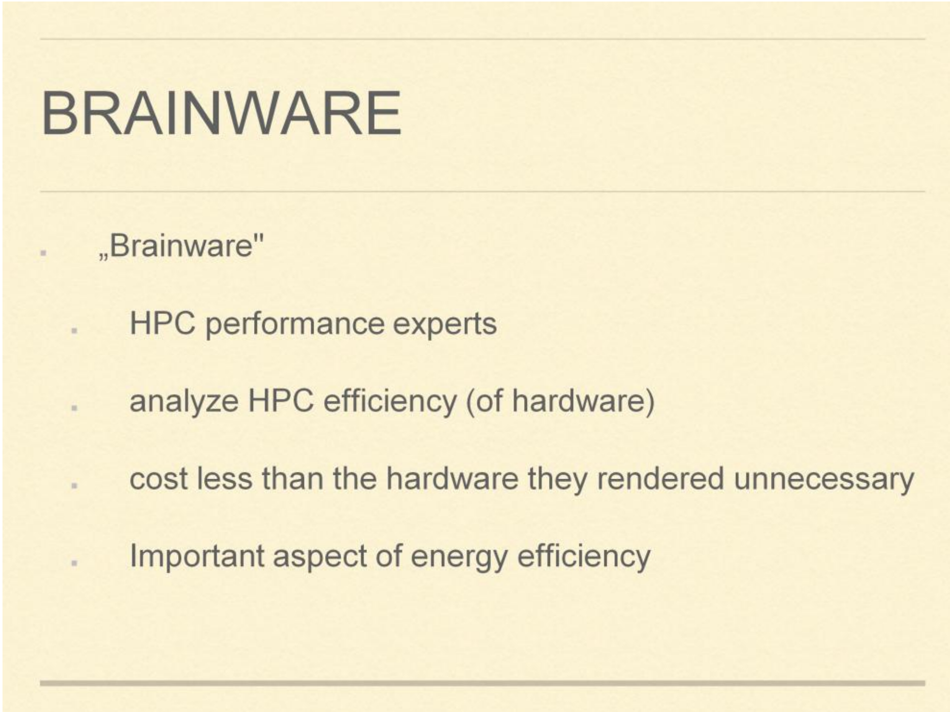

- Trained HPC specialists support performance optimization of user‘s codes.
- analyze efficiency of hardware and software -> make changes to make it more efficient
- Cost less then the unnecessary hardware

BRAINWARE

- Assumptions:
  - 7,5 € Mio for Infrastructure
  - Hardware 2 Mio €
  - 4 years maintenance, then new hardware
  - Power consumption: 850 KW
  - ISV software provided by users

	COST/YEAR	PERCENTAGE
BUILDING	300.000 €	5,46 %
INVESTMENT	2.000.000 €	36,14 %
HARDWARE MAINTENANCE	800.000 €	14,46 %
POWER	1.563.660 €	28,26 %
LINUX	0 €	0,00 %
BATCH SYSTEM	100.000 €	1,81 %
ISV SOFTWARE	0 €	0,00 %
HPC SOFTWARE	50.000 €	0,90 %
STAFF (12 FTE)	720.000 €	13,01 %
TOTAL SUM	5.533.660 €	100 %

Example from RWTH Aachen University:
Wrote an paper about Brainware and used University computing center

1. Building: 7,5 Mio €
   for compute cluster, offices
   25 years
   so 300.000€ each year
2. Hardware
   2 Mio €
   Maintenance: typically 10% of investment price
   maintenance for 4 years
   Why 4? Afterwards more advantegeous to buy new equipment with better performance (Moore‘s Law)
3. Power consumption
   will be operated 24/7
   PUE 1,5
   assumed power consumption of 850 KW per year
   1840€ per KW/Year => 1,56 Mio € per year
4. Software
   employ free software to a large extend
   calculate minor cost for compiler and tools (50000€)
   licenses provided by end users
5. Staff
   12 full time employees
   running the HPC equipment and supporting users
   60000€/year for 1 FTE

BRAINWARE

- It takes 2 months to tune one project
- An expert can handle 5 projects per year
- HPC experts can improve the performance of projects by 5,10 or 20 %
- HPC performence expert can take care of 10 projects at a timeti
- First take care of the „hot spots“ (top projects in order of CPU usage)

BRAINWARE

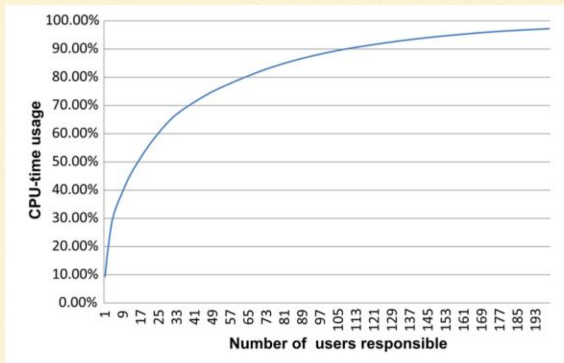

„Brainware for Green HPC" ; Christian Bischof, Dieter an Mey, Christian Iwainsky

BRAINWARE

- Example:
  - 1,5 FTE
  - take care of 15 projects
  - 50 % of CPU usage
  - 10% performance improvement
  - => 0,1 * 0,5 * 5,5 Mio € - 1,5 * 60000 € = 185000 € Savings

BRAINWARE

- Example:
  - 3 FTE
  - 30 projects
  - 60 % CPU usage
  - 0,2 * 0,6 * 5,5 Mio € - 3 * 60000 € = 480000 € Savings

- Brainware pays off

CONCLUSION

- higher computation power -> higher energy consumption -> higher costs
- good planning needed
- different methods to lower the TCO
    - but: making an HPC more energy efficient, makes ist also more cost-efficient

REFERENCES

- [1] „Advances In Computers: Green and Sustainable Computing Part II"; Ali Hurson; 2013
- [2] „The Datacenter as a Computer"
- [3] „Brainware for Green HPC"; Christian Bischof, Dieter an Mey, Christian Iwainsky; Springer Verlag; 2011
- [4] „Cost Model for Planning, Development and Operation of an Data Center": http://www.hpl.hp.com/techreports/2005/HPL-2005-107R1.pdf
- [6] „Google data center in Hamina": http://www.google.com/about/datacenters/inside/locations/hamina/
- [7] „Google data centers; Efficiency: How we do it" http://www.google.com/about/datacenters/efficiency/internal/#water-and-cooling

REFERENCES

- http://www.green500.org/news/green500-list-june-2014
- „Chinas Supercomputing Strategy called out“ http://www.hpcwire.com/2014/07/17/dd/
- „World’s fastest computer Tianhe-2, might get very little use“ http://www.scmp.com/news/china/article/1264529/worlds-fastest-computer-tianhe-2-might-get-very-little-use
- „Determining Total Cost of Ownership for Data Center and Network Room Infrastructure“ http://www.linuxlabs.com/PDF/Data%20Center%20Cost%20of%20Ownership.pdf
- „Tianhe-2“ http://en.wikipedia.org/wiki/Tianhe-2
- HPC Industry Dynamics http://www.intersect360.com/industry/reports.php?id=95