Green Computing: Datacentres

Simin Nadjm-Tehrani
Department of Computer and Information Science (IDA)
Linköping University
Sweden

Many thanks to **Jordi Cucurull**
For earlier versions of this course material

This lecture

- Datacentre design
- Efficiency metrics
- Energy-proportional computing
- Initiatives

- Seminar groups
- Datacentre visit

This lecture

- **Datacentre design**
- Efficiency metrics
- Energy-proportional computing
- Initiatives

Chapter 4 and 5 of the 1st Edition of the e-book on the course web (2009)

- Seminar groups
- Data centre visit

Datacentres

"Datacenters are buildings where multiple servers and communication gear are co-located because of their common environmental requirements and for physical security needs,and ease of maintenance."

Barroso and Hölzle

Luiz André Barroso and Urs Hölzle. The Datacenter as a Computer - An Introduction to the Design of Warehouse-Scale Machines. Morgan & Claypool Publishers, 2009 (120 pages).

Note: a 2013 2nd Edition of this book (156 pages) as also available with more Google-oriented updates. Most of the basic material in the 2009 is still worthwhile to study.

IT Architectural overview

Energy use at a datacentre

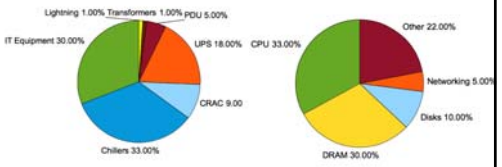

Datacentre overhead Server consumption

From Luiz André Barroso and Urs Hölzle, 2009, Data (Right chart) from one Google facility 2007.

Power and cooling systems

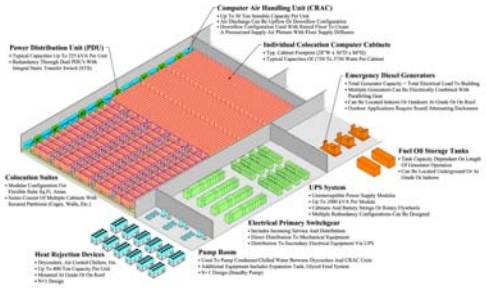

Graphic provided courtesy of DLB Associates

Power system components

- Primary switchgear
  - Breakers for protection
  - Transformers (10-20kV to 110-600V)
- Diesel generators
  - Switched on in case of utility power failure
- Uninterruptible Power Supply (UPS)
  - Batteries for short term energy provision
  - Functionality
    - Switch energy source between mains power (from utility) and power from Diesel generators
    - Sustain system power with batteries during mains power failure
    - Power feed conditioning

Power system components

- Power Distribution Units (PDUs)
  - Break a higher voltage line into several circuits
    - 200-480V line to many 110-220V circuits that feed the servers
    - A ground short (in server or power supply) will only break one circuit
  - Distribute energy to each rack
  - Provide redundancy (A-side, B-side) so that with a power supply failure fast switching can take place

Cooling system

- Computer Room Air Conditioning (CRAC)
  - Blows cold air under the floor plenum
  - Cold air moves to front of server racks (cold aisle)
  - Cold air flows through server racks
  - Warm air is expelled in the back (warm aisle)

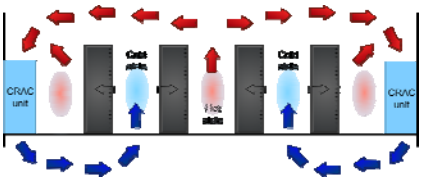

Cooling: operational range

- Datacentre cannot operate without cooling
- CRAC units cool the room's air
  - Liquid coolant is pumped from chillers or cooling towers
  - Coils are kept cool (12-14 ºC) with liquid coolant
  - Warm air is pushed through the coils by fans
  - Cold air (16-20 ºC) is moved to the floor plenum
- Air reaches the servers at 18-22 ºC

Free cooling

- Much more efficient than chillers
- Cool the coolant to much lower temperatures before reaching the chiller

Cooling tower

Glycol-based radiator

Fans to push air from outside

Redundancy and reliability

- Redundancy for fault tolerance and during maintenance
  - Applied to both power and cooling systems
- Configurations: N+1, N+2, 2N

Tier Classification (ANSI/TIA 942)

Type	Availability	Description
Tier I	99.67%	Single path for power & cooling. No redundancy (N)
Tier II	99.74%	Single path for power & cooling. Redundancy (N+1)
Tier III	99.98%	Multiple paths for power & cooling (only one active). Concurrently maintainable. Redundancy (N+1)
Tier IV	99.995%	Multiple active paths for power & cooling. Redundancy in both paths (min (N+1))

W.Pitt Turner IV, J.H.Seader, K.G.Brill. Tier classifications define site infrastructure performance, Uptime Institute, White Paper

Container-based datacentres

Server racks, power distribution and heat exchange inside a container!!

Highly efficient cooling
High server density

Image from Wikimedia by user Sun Microsystem

This lecture

- Datacentre design
- **Efficiency metrics**
- Energy-proportional computing
- Initiatives

- Seminar groups
- Datacentre visit

Sources of efficiency losses

- Power systems

Element	Losses
Transformers	0.50%
UPSs	7-12%
Highly efficient UPSs	3%
Low-voltage power (110-220V) cables	1-3%

- Cooling
  - Fans that move cool and warm air
  - Mix of cool and warm air during long paths
  - Too low temperature selection
    - 25-27 °C better than traditional 20 °C

Energy efficiency metrics

- Power Usage Effectiveness (PUE)
  - Proposed by The Green Grid association
  - Efficiency of the IT support infrastructure

$$PUE = \frac{Total\ facility\ power}{Total\ IT\ equipment\ power}$$

Historic data: PUEs between 1.5 and 2.0

Google, Microsoft, Apple have reported PUEs around 1.1X

C. Beladi. Green Grid datacenter power efficiency metrics: PUE and DCiE. White paper. 2008

Location matters…

- One quarter of Singapore data centers have PUE higher than 3, Japan and Australia claim average PUE of 2.2 and 2.25 respectively (APAC datacenter survey, April 2013)
- Facebook Luleå: 1.08 (Computer Sweden, 7 Feb 2013)

Image: https://www.facebook.com/LuleaDataCenter

Energy efficiency metrics

- PUE has to be used carefully
  - Example: Server fans in IT-PAC Microsoft data centre container module

IT-PAC airflow management is improved

Server fans are no longer required

$$\frac{Facility\ power}{IT\ equipment\ power}$$

PUE increases

Overall energy consumption reduced despite PUE increase

- Other aspects must be taken into account

http://www.datacenterknowledge.com/archives/2011/01/31/microsoft-eliminates-server-fans-despite-pue-hit/

Energy efficiency metrics

- Green Grid's Datacentre Performance Efficiency (DCPE)

$$DCPE = \frac{Amount\ of\ computational\ work}{Total\ energy\ used}$$

- Barroso and Hölze propose:

$$DCPE = \frac{1}{PUE} \times \frac{1}{SPUE} \times \frac{Computation}{Total\ Energy\ to\ Electronic\ Components}$$

- Computing efficiency
- IT equipment efficiency
- Building facilities efficiency

Energy efficiency metrics

- Server Power Usage Effectiveness (SPUE)
  - Efficiency of the IT infrastructure
    - Losses in power supply, voltage regulator modules, and cooling fans

$$SPUE = \frac{Total\ server\ input\ power}{Total\ useful\ power}$$

    - Useful power
      - Consumption of electronics directly involved in computation
      - E.g. motherboard, disks, CPU, DRAM...

Servers have SPUEs between 1.6 and 1.8

State of the art servers should be less than 1.2!!!

Energy efficiency metrics

- Total Power Usage Effectiveness (TPUE)
  - Show efficiency of electromechanical overheads.

$$TPUE = PUE \times SPUE$$

Example: 2.0 PUE
1.6 SPUE

$$TPUE = 2.0 \times 1.6 = 3.2$$

For each productive Watt another 2.2 W consumed!!!

This lecture

- Datacentre design
- Efficiency metrics
- **Energy-proportional computing**
- Initiatives

- Seminar groups
- Datacentre visit

Load and energy efficiency

- Most of the time spent in no energy-efficient load regions

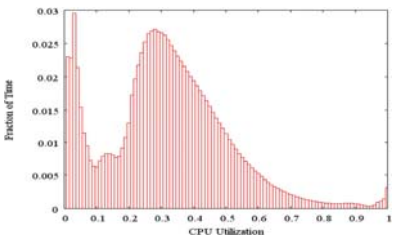

Activity profile of 5000 Google servers in 6 months

Image from Luiz André Barroso and Urs Hölzle. The Datacenter as a Computer 2009.

Load and energy efficiency

- Almost no time in idle state
  - In low load there are several hundreds of queries
  - Load spread out over all available servers
    - Optimised for performance and/or availability

Sustainable approach

Concentrate load in a few servers
Switch off unused servers

But!!!

More software complexity
Energy for data distribution

Energy proportional computing

- Consumption proportional to the load
  - Ideal linear function without constants

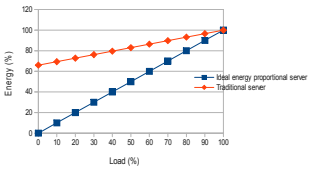

Traditional server data based on data from Barroso and Hölzle 2009.

Energy proportional computing

- Capacity to adapt consumption to load

Device	Dynamic power
CPU	3.0x
Memory	2.0x
Disks	1.3x
Network switches	1.2x

More hardware improvements are required!!!

Meanwhile only option is to switch off hardware!!!

Energy spent: revisited

CPU a larger proportion of energy costs than that reported in 2009!

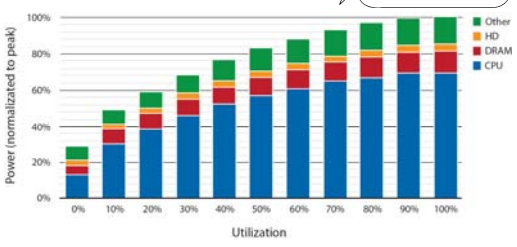

Luiz André Barroso and Urs Hölzle. The Datacenter as a Computer, 2nd Edition 2013

This lecture

- Datacentre design
- Efficiency metrics
- Energy-proportional computing
- **Initiatives**

- Seminar groups
- Datacentre visit

ENERGY STAR

- Future specifications
  - Program Requirements for Computer Servers 2.0
    - Currently first draft available
    - Will include active consumption
      - as opposed to just idle consumption

Industry drive: Microsoft joins Facebook to lead the Open Computer Project on servers Computer Sweden, 30 Jan. 2014

http://blogs.technet.com/b/microsoft_blog/archive/2014/01/27/microsoft-contributes-cloud-server-designs-to-the-open-compute-project.aspx

  - Program Requirements for Datacenter Storage 1.0
    - Will establish maximum consumption of storage products

"Fully-functional storage system that supplies data storage services to clients and devices attached directly or through a network"

Storage product: Definition from Energy Star specification

EU Code of Conduct for Datacentres

- European action to reduce consumption of datacentres
  - Best practices for datacentres operators
    - V3.0 Guide (Feb 2011)
  - Operators register and commit to their application

- Roles of applicant
  - Participant
    - Operator of datacentre or equipment in it
    - Commitments
      - Annual report of energy consumption
      - Implementation of some of the best practices
  - Endorser
    - Support the initiative and participants

http://re.jrc.ec.europa.eu/energyefficiency/html/standby_initiative.htm

EU Code of Conduct for Data Centres

- Scope of the best practices
  - Datacentre utilisation, management and planning
  - IT equipment and services
  - Cooling
  - Datacentre power equipment
  - Other equipment
  - Building
  - Monitoring

- Best practice expected implementation
  - Immediate
  - During software install or update
  - During new IT introduction or replacement
  - During building of datacenter or retrofit
  - Optional

This lecture

- Datacentre design
- Efficiency metrics
- Energy-proportional computing
- Initiatives

- **Seminar groups**
- Datacentre visit

Seminar groups

- Groups are being formed
- We have 62 students that registered in webreg and sent their preferences
- That means we will move around a bit so that 3 groups of 14 (A, B, C) and two groups of 10 (D, E) will be formed

More details in Sem0

- **Important notice:** Absence/dropout after group formation will be reported as a fail in Ladok

Datacentre visit

- We have the opportunity to see the equipment at the company Availo (Mjärdevi) that will show us around their premises
- Based on the past years' experience roughly half of the students are interested to go
- We have booked two occasions:
  - **10/2 at 15.00**
  - **11/2 at 16.00**
- Webreg is open for registering for these visits!

Questions?