

Power/Energy and Cooling in Data Centers

Kai Shen

11/5/2013

CSC 296/576 - Fall 2013

1

Overview

- Power (instantaneous) vs. energy (integral of power over time)
- Power and energy uses in data centers
 - Behind-the-scene costs of big data management and computing
 - Costs include infrastructure setup, operating expenses, and environmental impact (public image)
 - Expensive: a substantial portion of total data center cost
 - Google claimed to have saved over \$1B from energy improvements
- Cooling is very much related to power uses
 - High-power equipments need more cooling
 - Cooling also consumes power
- Data centers on the Columbia river
<http://www.dailywireless.org/2012/04/17/green-power-for-columbia-river-data-centers/>

11/5/2013

CSC 296/576 - Fall 2013

2

Data Center Energy Efficiency

$$\text{Efficiency} = \frac{\text{Useful computer energy}}{\text{Total data center energy}} = \frac{1}{\text{PUE}} \times \frac{1}{\text{SPUE}}$$

- $\frac{1}{\text{PUE}}$ represents facility efficiency – getting power to IT equipment
- $\frac{1}{\text{SPUE}}$ represents server efficiency – getting power to computers

11/5/2013

CSC 296/576 - Fall 2013

3

Facility Efficiency

$$\text{Efficiency} = \frac{\text{Useful computer energy}}{\text{Total data center energy}} = \frac{1}{\text{PUE}} \times \frac{1}{\text{SPUE}}$$

- Power usage effectiveness (PUE)
 - Ratio of total facility energy to IT energy (those consumed by computing and network equipments)
 - Non-IT energy uses on chillers, UPS, other mechanics (powered door), etc.
- Average: ~3.0 (in 2006) \Rightarrow ~1.8 (in 2012)
- Leading data centers at or below 1.2 today
 - Careful air flow handling
 - Elevated temperature
 - Free cooling

11/5/2013

CSC 296/576 - Fall 2013

4

Server Efficiency

$$\text{Efficiency} = \frac{\text{Useful computer energy}}{\text{Total data center energy}} = \frac{1}{\text{PUE}} \times \frac{1}{\text{SPUE}}$$

- Server Power usage effectiveness (SPUE)
 - Ratio of IT energy to useful server energy (used by electronics inside servers)
 - Losses due to power supply costs, voltage regulation, and server cooling fans.
- Efficiency from half to about 80%
 - Figures 4.2, 4.3, 4.4 (pages 52-53)


11/5/2013

CSC 296/576 - Fall 2013

5

Energy Efficiency within a Server

- Ideally the computing/data resources and their configurations are optimized for highest energy efficiency
- Challenging because the optimal system setup is dependent on the workload and the workload is dynamic



- Resources are provisioned for peak load, so many machines are idle at low load periods (but they still consume power)

11/5/2013

CSC 296/576 - Fall 2013

6

Server Consolidation

- Monitor the workload level and shutdown some machines at low load periods
 - Challenge: bring up a machine takes time which may hurt QoS
 - Different power down levels (with different power saving and bringup delay)
- But at odds with design of distributed big data systems
 - Data is spread over the local storage of all machines (for load distribution and replication)
 - All (or many) must be up to ensure data availability

11/5/2013

CSC 296/576 - Fall 2013

7

Energy-Proportional Computing

- Can we keep a machine up (ready for work) but it only consumes power when performing work and the energy use is proportional to its work?
 - So we don't need to worry about powering down/up machines while still be energy-efficient in face of dynamic load.
 - Advocated by Barroso and Holze at Google
- Current computer systems aren't energy-proportional
 - Figure 5.6 (page 79)
 - CPU/RAM are better than other components
 - CPU frequency/voltage scaling, duty-cycle modulation etc.
 - Things are getting better: 60% idle-peak power ratio (2006) ⇒ 40% (2009) ⇒ 20% (2012)
 - We see 7% idle-peak power ratio on a Google Nexus7!

11/5/2013

CSC 296/576 - Fall 2013

8

Software System Optimization

- Adaptive control of the CPU power/speed configuration
 - Knowledge on the application workload allows more aggressive optimization
 - Multicore scheduling (where to place two map tasks on a two-socket, two-core machine)
- Power / data availability tradeoff
 - In a distributed big data system, a reduction of replication degree allows more machines to shutdown at a cost of less availability
 - A restarting replica needs to sync and catch up with lost updates


11/5/2013

CSC 296/576 - Fall 2013

9

Power/Energy Accounting

- Power measurements at servers
- CPU core-level measurement is actually hard
 - But can be modeled using software/hardware event frequencies
- User / request-level power accounting
 - Carbon footprint of a web search?
 - Distribution of request energy usage




11/5/2013

CSC 296/576 - Fall 2013

10

Power Viruses

- Power viruses in cloud systems



- Power virus containment [Shen et al. ASPLOS 2013]

11/5/2013

CSC 296/576 - Fall 2013

11

Conventional vs. Wimpy CPUs

- Conventional CPU (regular CPUs on desktops / servers)
- Wimpy CPU (what you have in a smart phone, worse)
- Why wimpy CPUs?
 - With good parallelism, a larger number of wimpy CPUs can do the job done at a smaller number of conventional CPUs
 - Cluster of wimpy CPUs may save power [Andersen et al. 2009]
- But good parallelism isn't easy to get

11/5/2013

CSC 296/576 - Fall 2013

12

Data Center Cooling

- Cooling efficiency matters
 - Principle: emitted hot air moves away as quickly as possible, without mixing with the cool air
 - Figure 4.6 (page 54)
- Software system management
 - Avoid hot spots
 - Temperature-aware workload placement

11/5/2013

CSC 296/576 - Fall 2013

13

Modular Design and Management

- Container-based data centers
 - Figure 4.15 (page 64)

11/5/2013

CSC 296/576 - Fall 2013

14

Data Centers on Renewable Energy

- Renewable energy
 - Hydro power is commonly used
 - Wind and solar power are still expensive, deployed for environmental concerns (public image) and expected high energy costs in the future
- Principles:
 - Renewable energy has less predictable capacity, energy storage is expensive ⇒ opportunistic use
 - Moving computing is more efficient than moving energy, but moving big data isn't cheap
 - Shift work schedule according to availability of renewable energy

11/5/2013

CSC 296/576 - Fall 2013

15

Really???

- Floating data center in the ocean
 - Wind and waves provide energy
 - Sea water to cool

11/5/2013

CSC 296/576 - Fall 2013

16



Disclaimer

- Preparation of this class was helped by materials in the book “The Datacenter as a Computer” by Luiz A. Barroso, Jimmy Clidaras, and Urs Holzle.