GREEN ICT

Anne-Cécile Orgerie

Lecture, Telecom SudParis
9th October 2020

Who I am

- Full-time researcher at CNRS (about 33,000 people)
- Located in Rennes, France.
- IRISA laboratory (about 1,000 people)
- Myriads team: INRIA, CNRS, University of Rennes, INSA, ENS Rennes (about 30 people)
- Energy efficiency in large-scale distributed systems

http://www.people.irisa.fr/Anne-Cecile.Orgerie

Green Computing?

"Designing, manufacturing, using, and disposing of computers, servers, and associated subsystems -- such as monitors, printers, storage devices, and networking and communications systems -- efficiently and effectively with minimal or no effect on the environment."

Sam Murugesan, "*Harnessing Green IT: Principles and Practices*" IEEE IT Professional, 2008.

- Green use: reduce usage of hazardous materials
- Green design: design compliant with the environment
- Green disposal: recycling e-waste with little impact
- Green manufacturing: new products without hazardous substances

Outline

I. Introduction to (not) green ICT

II. Trails to green ICT from my research point of view

A. Green computing history
B. Data center level
C. Measuring energy consumption
D. Slowing down
E. Switching off unused resources
F. Efficient scheduling
G. Exploiting renewable energy

III. Concluding remarks

Energy for dummies

Physical measure	Unit
Current (I)	Amperes (A)
Voltage (U)	Volts (V)
Power (P)	Watts (W)
Energy (E)	Joules (J) or Wh

P = U x I
E = P x t

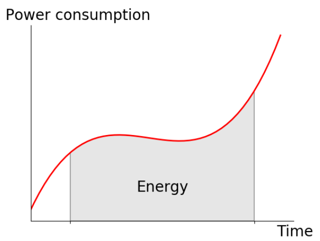

ICT impact ?

- What is ICT carbon impact in comparison with global impact?
- What is carbon impact?
- Which part of the lifecyle of an ICT product has more carbon impact?

https://app.klaxoon.com/join/VNCJ9JX

ICT is not green

ICT is responsible for 2% to 10% of CO2 emissions

M. Mills. The Cloud Begins With Coal - Big Data, Big Networks, Big Infrastructure, and Big Power. Technical report, Digital Power Group, Aug. 2013.

Computing in the 21st century?

The Cloud

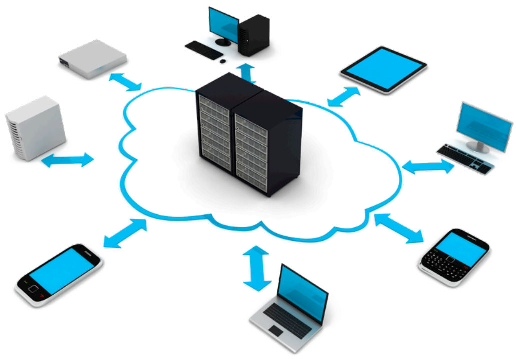

ICT energy consumption

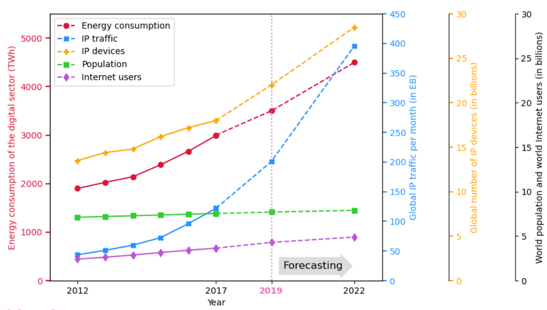

3% of the global energy consumption in 2018.

Rapport Lean ICT : Pour une sobriété Numérique, 2018, https://theshiftproject.org
Cisco Visual Networking Index: Forecast and Methodology [2013—2019].
World Population: Past, Present, and Future *https://www. worldometers.info/world-population* 2019.
International Telecommunication Union, *Measuring the Information Society Report*, 2018.

Resource waste

Servers are used 6% on average.

Source*: Revolutionizing Data Center Energy Efficiency*, McKinsey, July **2008**.

Networks are lightly of unevenly utilized

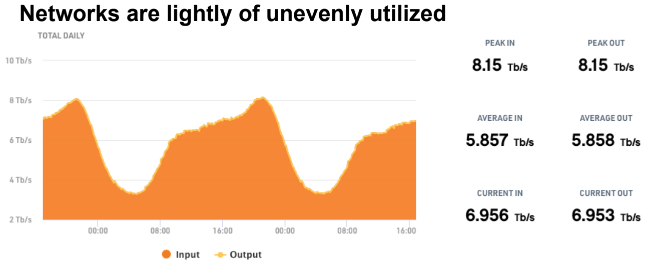

Daily aggregated traffic on AMS-IX(Amsterdam Internet eXchange Point), October 2020.

Cloud computing in 1 slide

Cloud computing: access through networks to on-demand self-service, configurable, shared computing resources.

- Mutualization of services
- Elasticity of infrastructures
- Externalization of data

Jevons Paradox: the increase in efficiency with which a resource is used tends to increase (rather than decrease) the rate of consumption of that resource.

Economies of scale

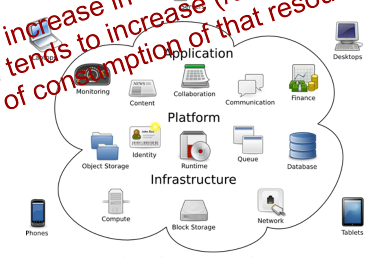

Can we save ICT…

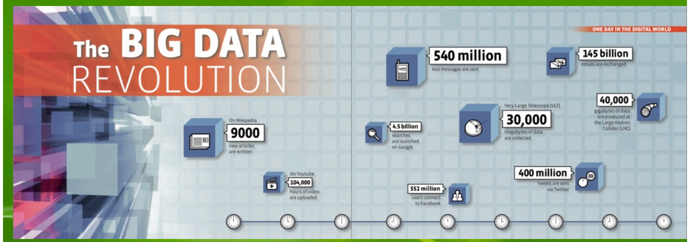

"The Big Data revolution", CNRS Ineternational Magazine, 2013.

… without changing users' habits

Internet of many things

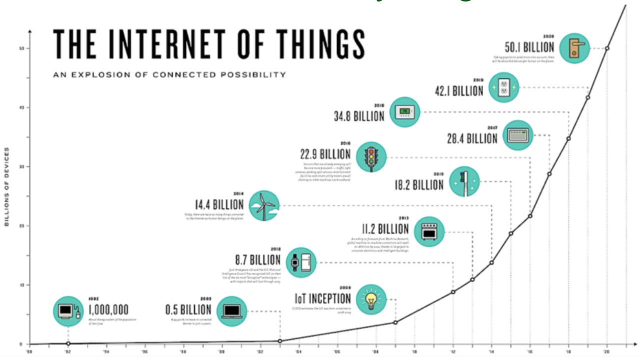

Source: Cisco

Practical Internet of Things

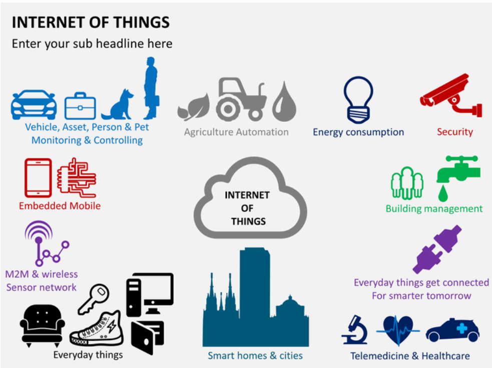

 http://www.supinfo.com/articles/single/4235-internet-of-things

 [Source : https://www.google.fr/about/datacenters]

One Google Data Center (Dalles)

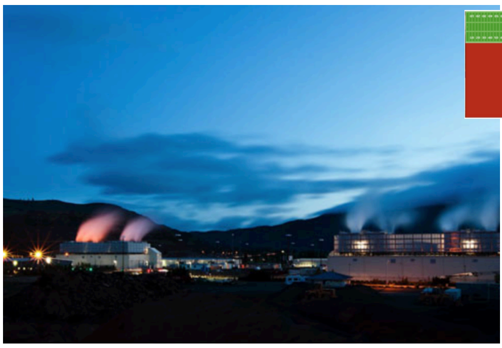

11,5 football fields
100,000 servers

100 MWatts

https://www.google.com/about/datacenters/inside/locations/the-dalles/

UPS to the rescue

Uninterruptible power supply:

- Emergency power system
- Used to protect hardware from power disruption
- Supplies energy stored in batteries, supercapacitors of flywheels (converted into alternating current)

But only for few minutes!

And then?

Engine-generator

OVH example

Roubaix site: ~ 10,000 servers

8 MVA at max: 1,600L/h of oil

Tests: every 2-3 weeks

https://lafibre.info/ovh-datacenter/test-de-groupes/

Other data centers

http://www.datacentermap.com/blog/datacenter-container-55.html

https://news.microsoft.com/features/microsoft-research-project-puts-cloud-in-ocean-for-the-first-time/#sm.0000q5ts4lqgfez110wem1gb0ig5o

Inside the cloud

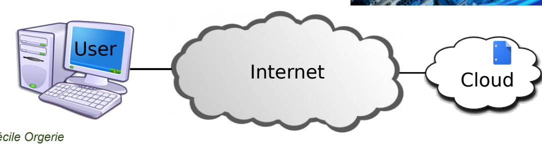

Distribution of ICT energy consumption

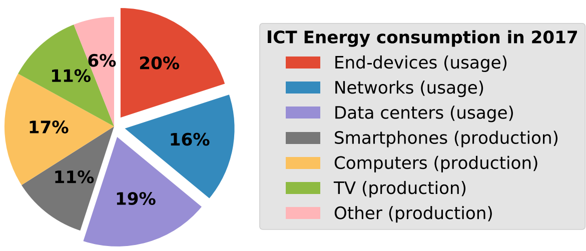

Rapport Lean ICT : Pour une sobriété Numérique, 2018
https://theshiftproject.org

CO2 is a different metric

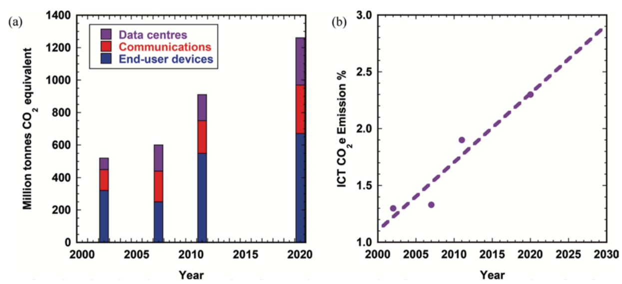

"ICT - Energy Concepts for Energy Efficiency and Sustainability", G. Fagas, L. Gammaitoni, J. Gallagher and D. Paul, InTech (open) book, 2017.

But globally, it provides the same message

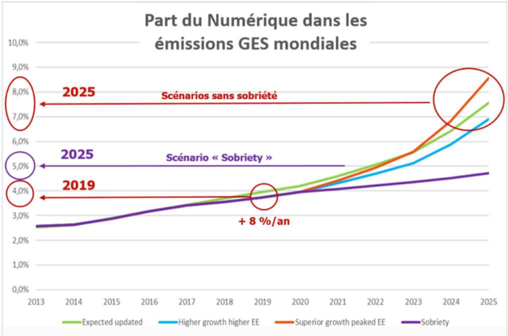

Green computing history

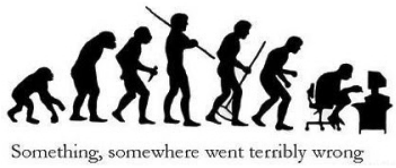

First eco-labels

- Energy Star: international standard for energy efficient consumer products
  - 1992, USA
  - Voluntary labeling program
  - To promote energy-efficient monitors, climate control equipment and other technologies
  - Main result: sleep mode
- TCO certification
  - 1992, Sweden
  - To promote low magnetic and electrical emissions from CRT-based computer displays

Specific eco-labels

- 80 Plus: voluntary certification program to promote efficient energy use in computer power supply units (2004):
  - More than 80% energy efficiency at 20%, 50% and 100% of rated load
  - Power factor of 0.9 or greater at 100% load
  - 80 Plus Titanium : 90% energy efficient
- RoHS: Restriction of Hazardous Substances Directive (2003):
  - Adopted in 2003, effective in 2006 in EU
  - Restricted use of six materials

Outline of my research work

1. Data center level
2. Understanding energy consumption
3. Slowing down
4. Switching off unused resources
5. Consolidation the load
6. Exploiting virtualization capabilities
7. Consuming renewable energy
8. Networking equipment

Data center level

Where is electricity consumed? Data center example

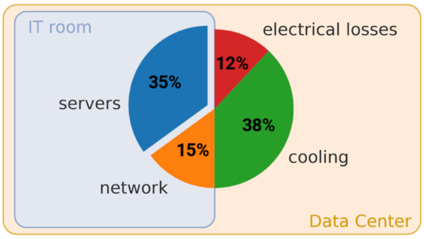

J. Ni and X. Bai, "A Review of Air Conditioning Energy Performance in Data Centers", 2017

Great part consumed by facilities

→ Cooling accounts for 30-50% of the total value

Let's reduce the heat

- Water-based cooling

https://www.datacenterknowledge.com/archives/2012/12/11/defense-department-cool-servers-with-hot-water

Let's reduce the heat

- Oil-based cooling

Green Revolution Cooling, https://www.grcooling.com

Let's reduce the heat

- Free cooling

https://www.google.com/about/datacenters/inside/locations/hamina/

Water-cooled doors

- Hybrid cooling

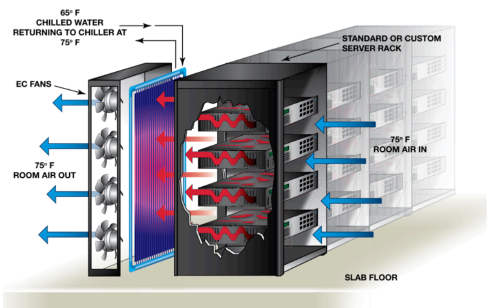

Exploded View of Chilled Door®

https://www.monman.com/motivair-chilled-door-rack-cooling-details

Water required!

Water tanks

Capacity:
900,000 liters

Google data center in South Carolina, US

https://www.google.com/about/datacenters/gallery/index.html#/tech/19

Water pipes

[Source : https://www.google.com/about/datacenters/gallery/index.html#/tech/19]

Thermal management

- Cooling costs a lot, so reduce heat production to reduce energy consumption

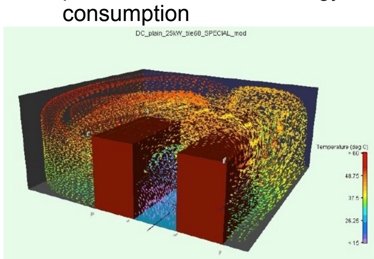

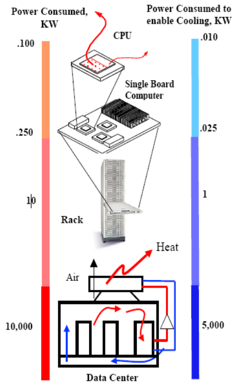

"*Energy Aware Grid: Global Workload Placement based on Energy Efficiency*", HP Technical Report, 2002.
"*Power Provisioning for a Warehouse-sized Computer*", X. Fan, W. Weber and L. Barroso, ISCA, 2007.

Understanding energy consumption

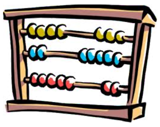

Understanding energy consumption

- Mandatory to optimize the energy consumption of a resource or an application
- Mandatory to simulate or emulate energy consumption
- Other usages:
  - monitoring,
  - forecasting,
  - accounting...

Grid'5000

How to measure energy consumption?

- Power usage per device, per process, per service, per rack?
- Software tools: **powertop**
- Event counters
- Sensors

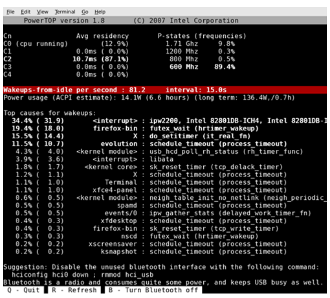

Less Watts. Saving Power with Linux on Intel® Platforms

http://lesswatts.org/projects/powertop/

"*Balancing power consumption in multiprocessor systems*", A. Merkel and F. Bellosa, SIGOPS Oper. Syst. Rev., 2006.

Intel Power Gadget

https://software.intel.com/en-us/articles/intel-power-gadget-20

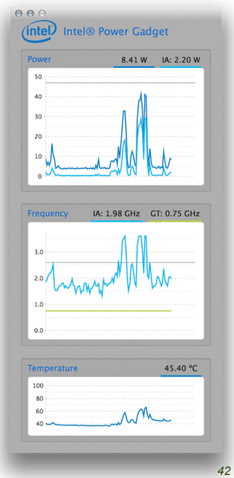

Without wattmeters

PAPI (Performance Application Programming Interface) can read RAPL (Running Average Power Limit) values

- Uses software power model, hardware performance counters, temperature, leakage models and I/O models

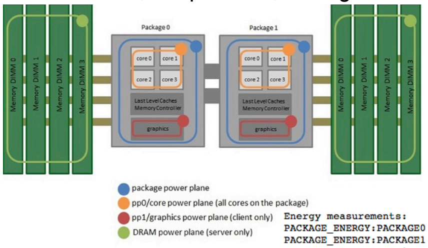


```
Energy measurements:
PACKAGE_ENERGY:PACKAGE0   176.450363J   (Average Power 42.9W)
PACKAGE_ENERGY:PACKAGE1    75.812454J   (Average Power 18.4W)
DRAM_ENERGY:PACKAGE0       11.899246J   (Average Power 2.9W)
DRAM_ENERGY:PACKAGE1        8.341141J   (Average Power 2.0W)
PP0_ENERGY:PACKAGE0       118.029236J   (Average Power 28.7W)
PP0_ENERGY:PACKAGE1        16.759064J   (Average Power 4.1W)
```


With an example: Grid'5000

- French experimental testbed
- 15,000 cores
- 8 sites
- Dedicated Gb network
- Designed for research on large-scale parallel and distributed systems

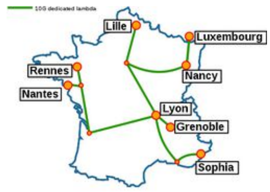

Grid'5000

A monitored site: Lyon

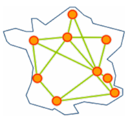

- 78 nodes
- 50 power measurements per node and per second
- Multiple views + logs on demand

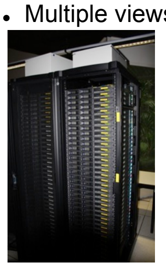

Supercomputers

N°1: Fugaku (TOP500 June 2020)

415 Petaflops, 28.335 MW, **7,299,072 cores,**

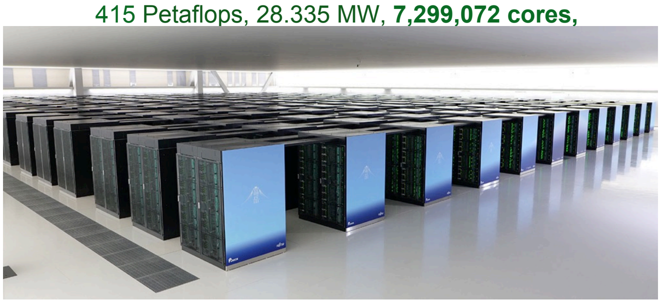

https://spectrum.ieee.org/tech-talk/computing/hardware/japans-fugaku-supercomputer-is-first-in-the-world-to-simultaneously-top-all-high-performance-benchmarks

Typical applications:
- Artificial intelligence
- Disaster-prevention simulations
- …

Top500

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,397,824	143,500.0	200,794.9	9,783
2	**Sierra** - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
3	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
4	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 , NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482
5	**Piz Daint** - Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100 , Cray Inc. Swiss National Supercomputing Centre (CSCS) Switzerland	387,872	21,230.0	27,154.3	2,384

http://www.top500.org, Top500 list, November 2018.

Performance development

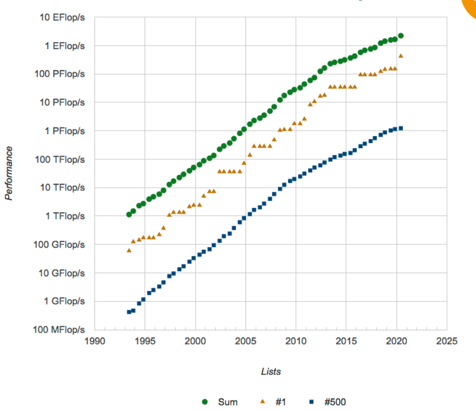

The Green500

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	393	**MN-3** - MN-Core Server, Xeon 8260M 24C 2.4GHz, MN-Core, RoCEv2/MN-Core DirectConnect, Preferred Networks Preferred Networks Japan	2,080	1,621.1	77	21.108
2	7	**Selene** - DGX A100 SuperPOD, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	272,800	27,580.0	1,344	20.518
3	468	**NA-1** - ZettaScaler-2.2, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-SC2 700Mhz, PEZY Computing / Exascaler Inc. PEZY Computing K.K. Japan	1,271,040	1,303.2	80	18.433
4	204	**A64FX prototype** - Fujitsu A64FX, Fujitsu A64FX 48C 2GHz, Tofu interconnect D, Fujitsu Fujitsu Numazu Plant Japan	36,864	1,999.5	118	16.876
5	26	**AiMOS** - IBM Power System AC922, IBM POWER9 20C 3.45GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM Rensselaer Polytechnic Institute Center for Computational Innovations (CCI) United States	130,000	8,339.0	512	16.285

https://www.top500.org/green500/, Green500 list, June 2020.

Energy spent in a server

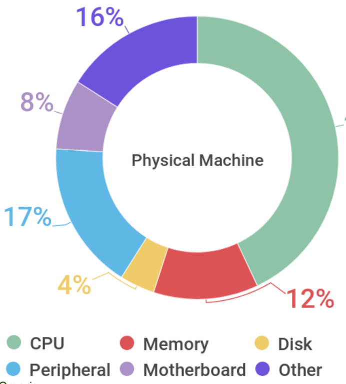

"GreenCloud: A Packet-level Simulator of Energy-aware Cloud Computing Data Centers,", D. Kliazovich, P. Bouvry, and S. U. Khan, The Journal of Supercomputing, vol. 62, no 3, pp 1263–1283, 2012.

Courtesy of David Guyon

Profiling applications

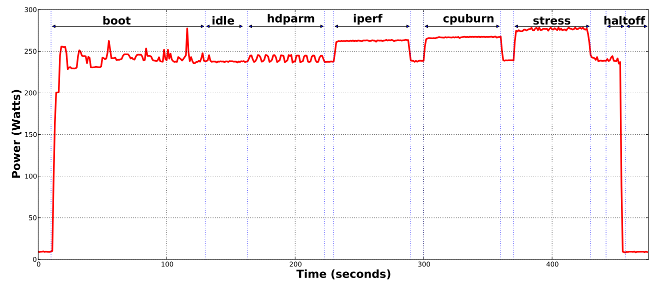

- Running benchmark applications
- Idle power consumption is really **high**.

Variable workload & non-power proportionality

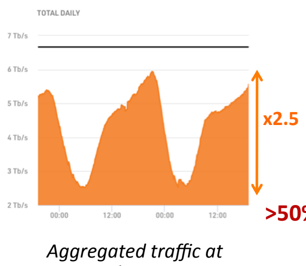

Aggregated traffic at Amsterdam Internet Exchange Point (AMS-IX).

[Source: https://www.ams-ix.net/ams, May 2019]

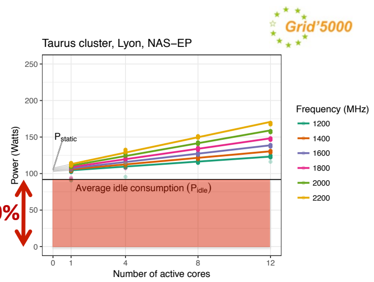

[Cluster 2017]

High idle power consumption.

Ideal server vs. current server

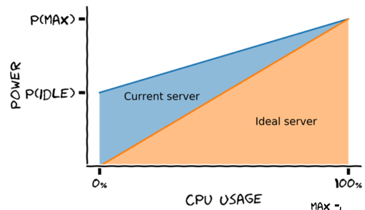

Courtesy of David Guyon

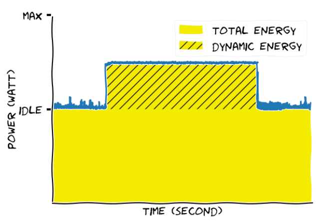

Wasted energy

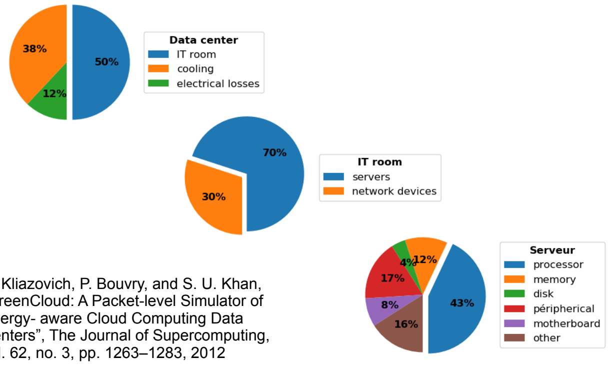

D. Kliazovich, P. Bouvry, and S. U. Khan, "GreenCloud: A Packet-level Simulator of Energy- aware Cloud Computing Data Centers", The Journal of Supercomputing, vol. 62, no. 3, pp. 1263–1283, 2012

Courtesy of David Guyon

How to measure energy efficiency?

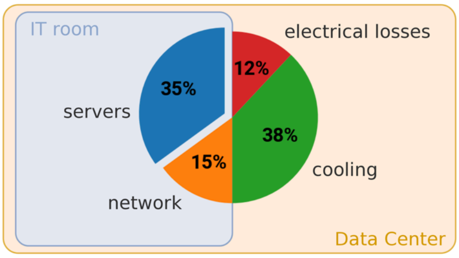

Courtesy of David Guyon

PUE: Power usage effectiveness

$$PUE = \frac{Total\ Facility\ Power}{IT\ Equipment\ Power}$$

"*Green Grid Data Center Power Efficiency Metrics: PUE and DCIE*", Green Grid White Paper, 2008.

the green grid

PUE examples

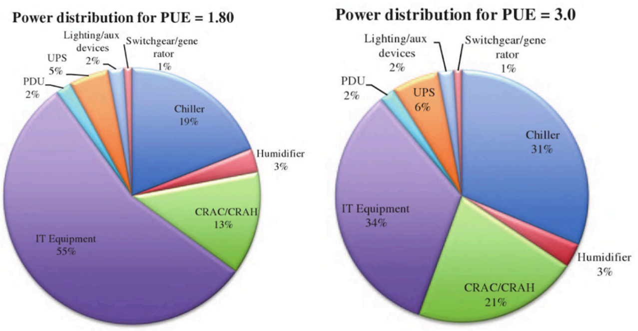

PUE as a selling point: Google case

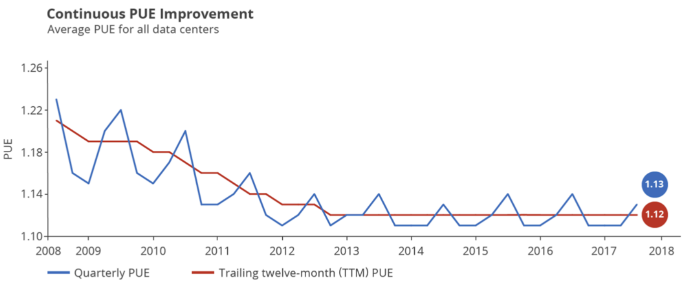

https://www.google.fr/about/datacenters/

Facebook live dashboard

OPEN Compute Project

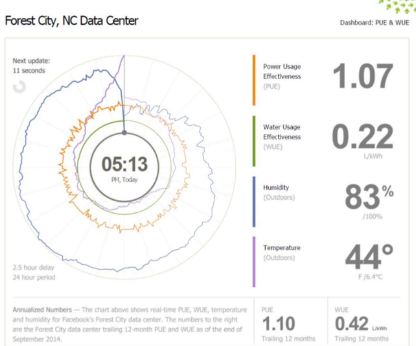

https://www.facebook.com/ForestCityDataCenter/app_288655784601722

Energy cost of an Internet box?

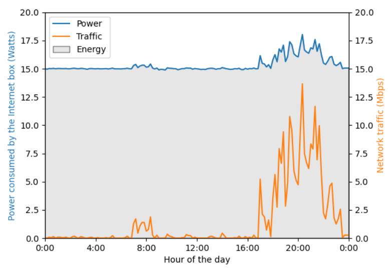

Energy cost in Joules/bit:

$$cost = \frac{energy}{data\_volume}$$

- When the bow is idle, it consumes15 Watts
- At 8am, 15.15 Watts and 0.65 Mbps
- At 8pm, 16.05 Watts and 4.75 Mbps
- At max, 25 Watts and 60 Mbps

Efficiency of an Internet box

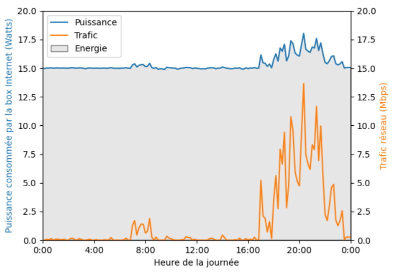

Efficiency metrics	Conditions	µJoules/bit
Instantaneous dynamic cost	At 8am, 0.15 Watts 0.65 Mbps	0.23
Instantaneous overall cost	At 8am, 15.15 Watts 0.65 Mbps	23.3
	At 8pm, 16.05 Watts 4.75 Mbps	3.38
Dayly overall cost	On average, 15.35 Watts 1.55 Mbps	9.9
Overall minimal cost	At max, 25 Watts 60 Mbps	0.42

Two orders of magnitude

Internet energy intensity

Estimates for the energy demand of Internet transmissions.

Study	Method	System boundary: Networking equipment	Optical fibers	End devices	Data for	Energy intensity
Koomey et al. (2004)	Top–down	X	X	X	2000	<136 kWh/GB
Taylor and Koomey (2008)	Top–down	X	X	X	2006	8.8–24.3 kWh/GB
Weber et al. (2010)	Top–down	X	X	X	2008	7 kWh/GB
Pickavet et al. (2008)	Top–down	X	X		2008	1.8 kWh/GB[a]
Lanzisera et al. (2012)	Top–down	X			2008	0.39 kWh/GB[a]
Baliga et al. (2007)	Model-based	X	X		2007	0.7–2.1 kWh/GB[a]
Baliga et al. (2009)	Model-based	X	X		2008	>0.179 kWh/GB
Baliga et al. (2011)	Model-based	X	X		2011 (?)	0.006 kWh/GB
Schien et al. (2012)	Bottom–up	X	X		2009	0.057 kWh/GB[a]
Coroama et al. (2013)	Bottom–up	X	X		2009	<0.2 kWh/GB

[a] Calculated by the authors based on the results provided in the cited study.

"Estimates published over the last decade diverge by up to four orders of magnitude — from 0.0064 kilowatt-hours per gigabyte (kWh/GB) to 136 kWh/GB."

"Assessing Internet energy intensity: A review of methods and results", V. Coroama, L. Hilty, Environmental Impact Assessment Review, 2014.

Standards, Consortiums, Projects

- Energy Star
- Green Grid
- Efficient Servers
- The green 500
- ...

Slowing down

Power modes by component

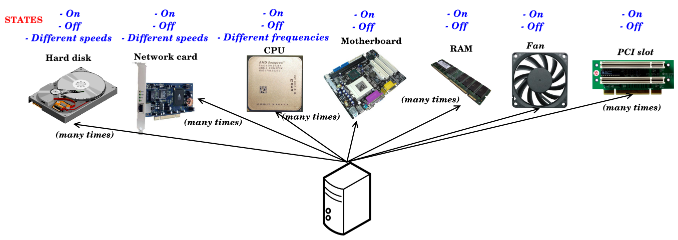

Dynamic adaptation to the load:

- Hard disks: spin-down
- Processors: Dynamic Voltage Frequency Scaling (DVFS)
- Network cards: Adaptive Link Rate

DVFS principle

CPU-Level DPM

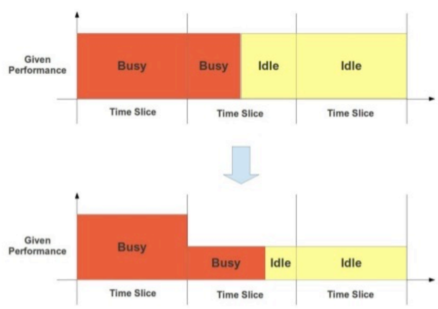

DVFS under load

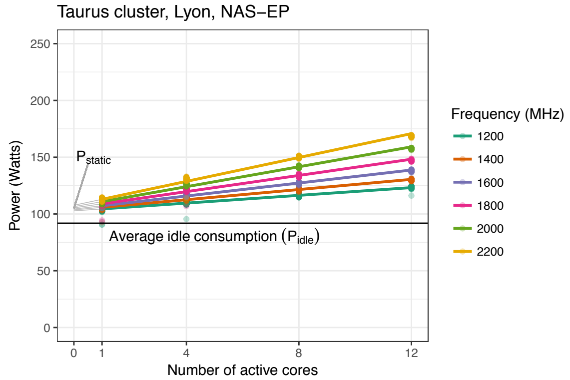

"*Predicting the Energy Consumption of MPI Applications at Scale Using a Single Node*", C. Heinrich, T. Cornebize, A. Degomme, A. Legrand, A. Carpen-Amarie, S. Hunold, A.-C. Orgerie, and M. Quinson, IEEE Cluster 2017.

big.LITTLE architecture

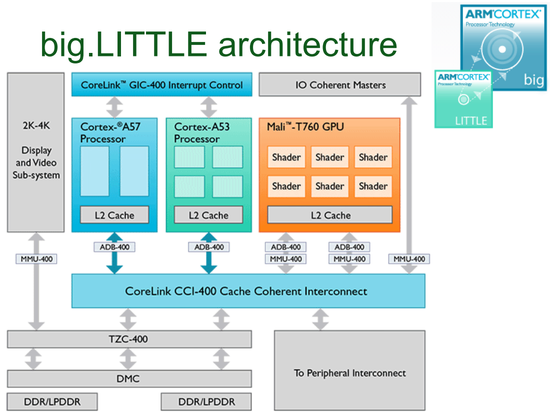

Fig. 3: Block diagram of a big.LITTLE System on Chip design representation

https://www.arm.com/products/processors/technologies/biglittleprocessing.php

Simulating energy consumption of applications

Simulation useful because:

- No need of big infrastructure
- Testing new approaches
- Fair comparison between two methods
- Reproducible

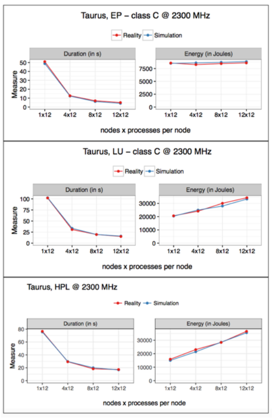

Reproducible results:
https://gitlab.inria.fr/fheinric/paper-simgrid-energy

Simulating energy consumption

http://simgrid.gforge.inria.fr

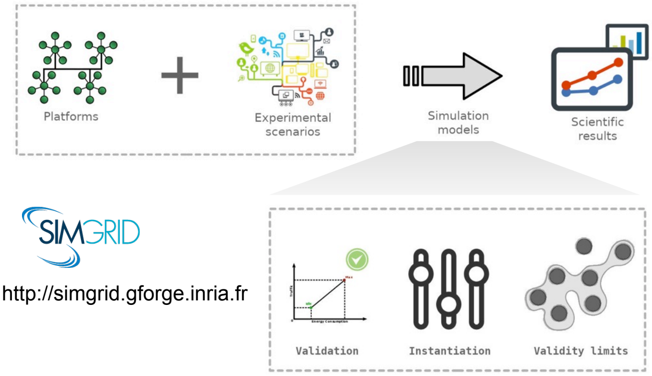

Courtesy of Loic Guegan

Energy consumption of IoT

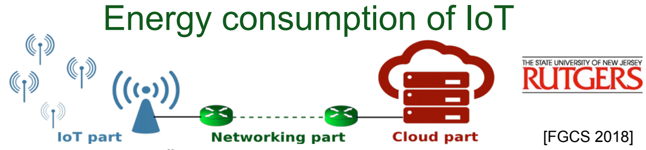

[FGCS 2018]

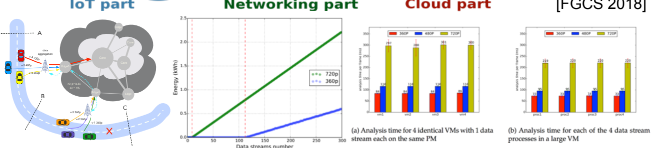

Scenario: Data stream analysis from cameras embedded on vehicles to detect objects on the road

Simulation of the networking part (wireless and wired)

(a) Analysis time for 4 identical VMs with 1 data stream each on the same PM (b) Analysis time for each of the 4 data stream processes in a large VM

Real measurements for the Cloud part: various configurations, image resolutions and application accuracy

Scenario	IoT	Network	Cloud
Edge Cloud	10.96 Watts	0.07 Watts	32.3 Watts
Core Cloud	10.96	0.11 Watts	22.8 Watts

Switching off unused resources

Switching off

Idea: adapt the set of active resources to the load

Issues:

- Does it reduce the life time of resources?
- How to switch resources on again?
- How much time does it take to switch?
- Switch off or sleep?
- Does the middleware consider the resources as dead?

Does a switched off node consume energy?

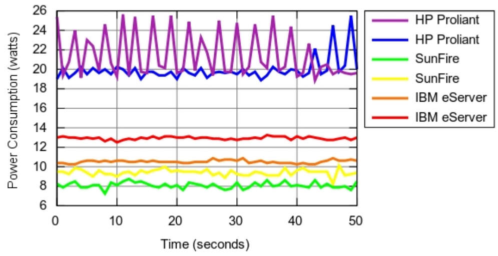

"*Demystifying Energy Consumption in Grids and Clouds* ", A.-C. Orgerie et al., WIPGC, 2010.

When can we switch off?

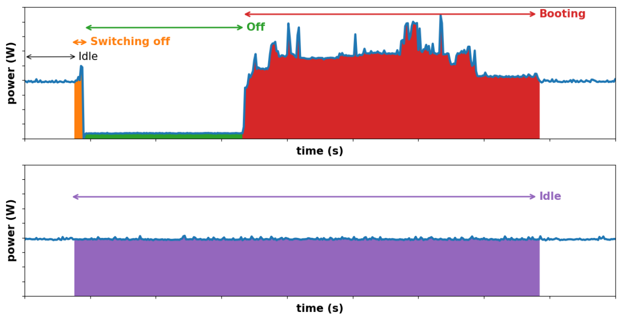

$$T_s = \frac{E_s - P_{OFF}(\delta_{ON\to OFF} + \delta_{OFF\to ON}) + E_{ON\to OFF} + E_{OFF\to ON}}{P_I - P_{OFF}}$$

"*Towards Energy Aware Reservation Infrastructure for Large-Scale Experimental Distributed Systems* ", A.-C. Orgerie et al., Parallel Processing Letters, 2009.

On/off tools

- Suspend to disk (hibernation)
- Suspend to RAM (standby or sleep)
- Wake on LAN
- IPMI (Intelligent Platform Management Interface)
- Presence proxies
- Smart PDU (Power Distribution Unit)

Energy-efficient scheduling to switch off more resources

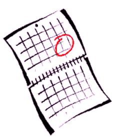

Consolidation

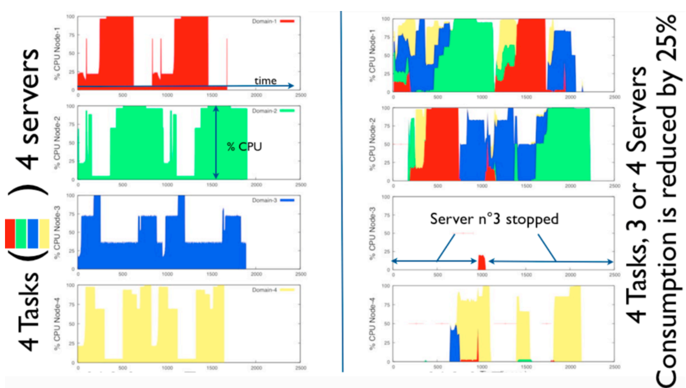

Courtesy of Jean-Marc Menaud

Consolidation in space

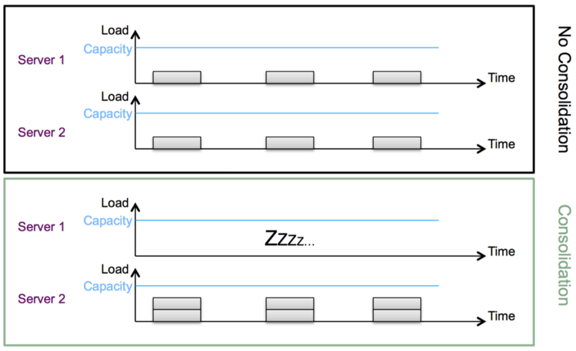

Courtesy of Ekhiotz Jon Vergara

Consolidation in time

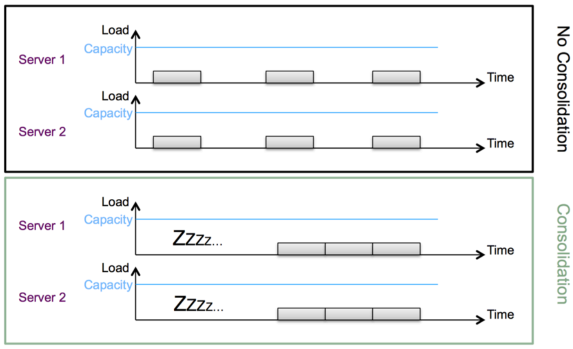

Courtesy of Ekhiotz Jon Vergara

Consolidation in space & time

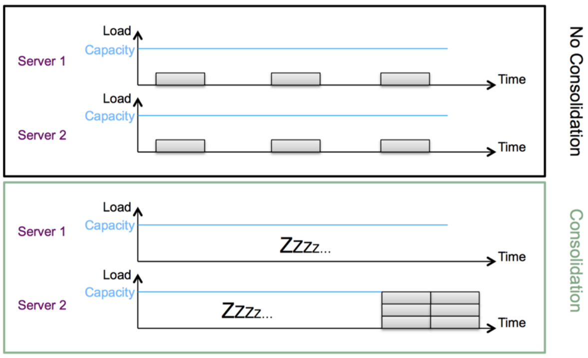

Courtesy of Ekhiotz Jon Vergara

Virtualization

Servers are used 6% on average.

Source: *Revolutionizing Data Center Energy Efficiency*, McKinsey, Juillet 2008.

Virtualization to the rescue.

Exploiting live-migration capabilities

Idea: migrating virtual machines to consolidate the load on the fewer number of physical resources

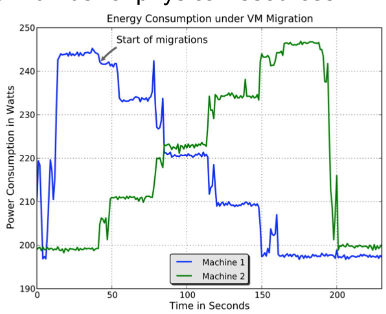

"*Energy Aware Clouds*", A.-C. Orgerie et al., book chapter in *Grids, Clouds and Virtualization*, Springer, 2010.

Infrastructure level & renewable energy

Follow-the-* approaches

- Follow the moon: free cooling with air from outside during cool days, and on hot weather days, computing load is shifted to other data centers
- Follow the sun/wind: use renewable energy sources

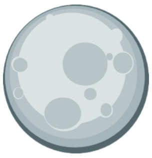

Problem

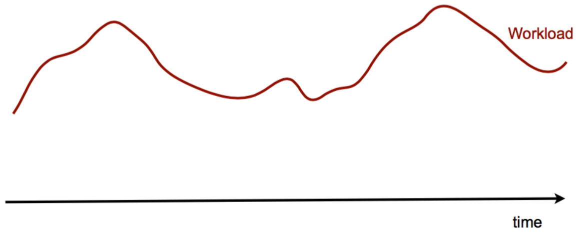

Courtesy of Jean-Marc Menaud

Problem

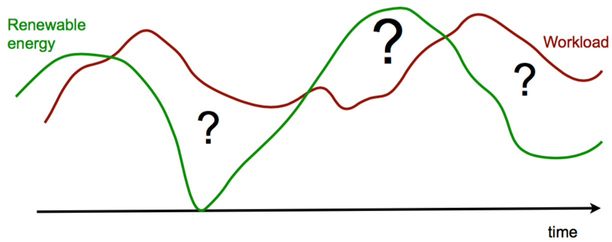

Courtesy of Jean-Marc Menaud

Problem

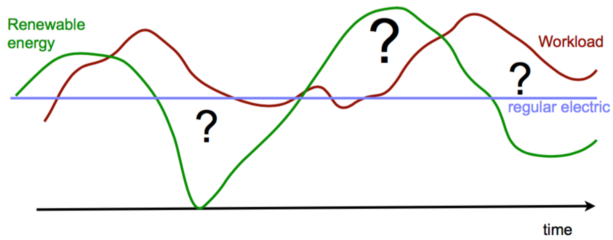

Courtesy of Jean-Marc Menaud

Opportunistic scheduling

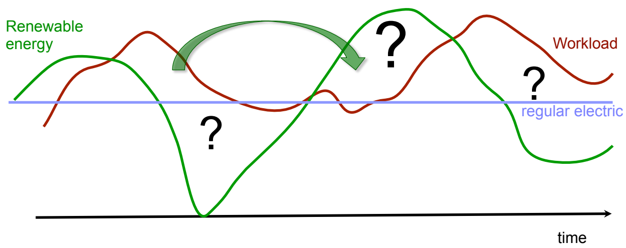

Opportunistic scheduling

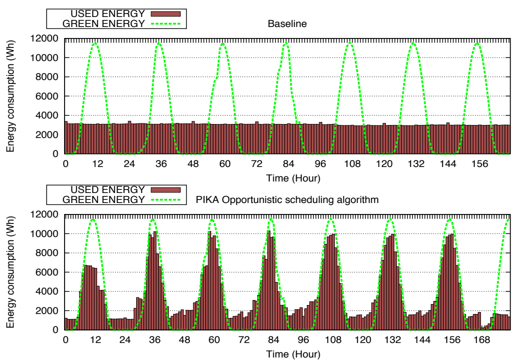

"*Opportunistic Scheduling in Clouds Partially Powered by Green Energy*", Y. Li, A.-C. Orgerie and J.-M. Menaud, GreenCom 2015.

Energy storage

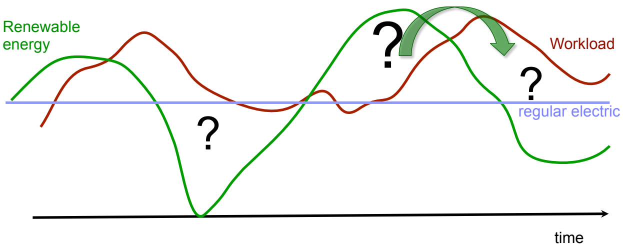

Real infrastructure: Parasol

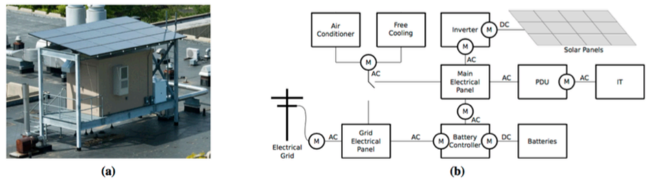

Figure 1. (a) Outside view of Parasol. (b) Parasol's power distribution and monitoring infrastructure.

THE STATE UNIVERSITY OF NEW JERSEY RUTGERS

http://parasol.cs.rutgers.edu

A real example

Apple's North Carolina iCloud data center

40 MW (max) of power

Apple Environmental Responsibility Report, 2017.

Renewable energy

- two 20 MW and one 18 MW solar arrays
- one 10 MW biogas fuel cells
- producing 244 million kWh annually
- daily on-site production: 60-100% of facility's consumption

Around 450 acres (1,800,000 m²) needed for solar farms

Apple Environmental Responsibility Report 2016

Apple's renewable energy use

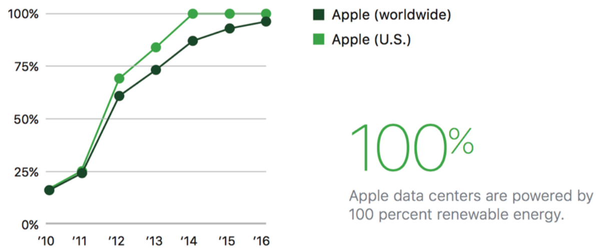

In just six years, Apple's use of renewable energy to power its corporate facilities, retail stores, and data centers worldwide went from 16 percent in 2010 to 96 percent in 2016.

https://images.apple.com/environment/pdf/Apple_Environmental_Responsibility_Report_2017.pdf

How is this possible?

The Truth About Apple's '100% Renewable' Energy Usage

Alex Epstein, Forbes, January 2016.

http://tinyurl.com/yc3r7c73

Apple pays off consumers and other companies to give it 'green credits' for its coal electricity usage

Conclusions

State of the art

Computing resources

Node level
- sleep state (hibernation)
- DVFS
- green softwares
- hardware capabilities

Infrastructure level
- green sources (follow-the-sun)
- thermal management
- workload consolidation
- task scheduling

Virtualization
- virtualization layer (green hypervisors)
- VM migration
- Cloud level

Wasted energy

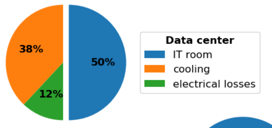

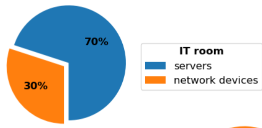

D. Kliazovich, P. Bouvry, and S. U. Khan, “GreenCloud: A Packet-level Simulator of Energy- aware Cloud Computing Data Centers”, The Journal of Supercomputing, vol. 62, no. 3, pp. 1263–1283, 2012

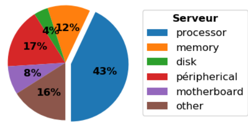

Courtesy of David Guyon

Sources of energy waste

Complex tradeoffs and modularity

System design is full of complex tradeoffs
- Peak vs. average performance
- Peak vs. average load
- General-purpose vs. dedicated
- High vs. best effort availability
- Backward compatibility

System functionality as independent modules
- Modularity and interaction
- System components designed separately (CPU, network interface...)

Courtesy of Simin Nadjm-Tehrani

Performance and general purpose

Optimization for peak performance scenario
- Low average system utilization
- Benchmarks stress worst-case performance workloads
  - Systems optimized for these scenarios

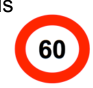

Good performance for a multitude of different applications
- Union of maximum requirements of each application class
- E.g. smartphone vs. MP3 player
- Legacy solutions

Growth and availability

- Overprovisioning to plan for the future
  - Ensure enough capacity
- Redundancy to increase availability

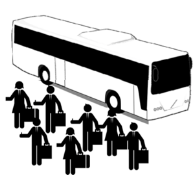

Vs.

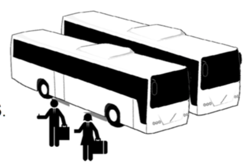

Design process structure

- Hardware and software separately
- Divided system functionality across components
- Layers
  - Local optimizations not optimal for global efficiency
  - E.g. worst-case assumption at each layer

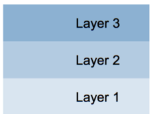

Energy efficiency at design stage

- Replacement with a more power-efficient alternative
- Holistic solutions
  - Look at problem with broad scope
  - Cross-layer interaction
- Optimize energy efficiency for the common case
- Design only for required functionality and requirements

Energy efficiency at runtime

- Trade off some other qualities for energy
- Disable or scale down unused resources
- Combination of multiple tasks in a single energy event
- Spend someone else's power
- Spend power to save power
- Monitor energy consumption to be energy-aware
- Predict resource usage trends
- Control algorithms and policies

To go a bit further

ICT for Green vs. Green ICT

- **ICT for Green**
  - Use ICT technologies to reduce the environmental footprint of other processes and sectors
  - E.g. smart grids, smart buildings, etc.
- **Green ICT**
  - Reduction of the ICT's environmental footprint
  - E.g. energy-aware data centers

The (in)dispensable weather toaster

In 2017: 5 connected devices / person
20 billion devices worldwide.

Forrester Research, "*Connected devices forecast, 2012 to 2017*", white paper, 2013.

The smart frying pan

Assisteo, Tefal, 2017.

In 2017: 5 connected devices / person
20 billion devices worldwide.

Forrester Research, "*Connected devices forecast, 2012 to 2017*", white paper, 2013.

Are we going on the good way?

- New functionalities
- Create new practices and needs
- Multiplication of the devices
- Capability overlap
- Health issues

Complete life-cycle

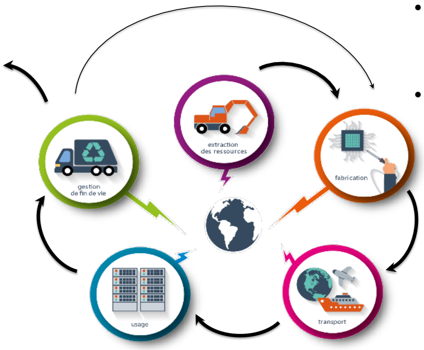

- 1.4 billion smartphones sold in 2015.
- Average life duration of first-hand smartphones < 2 years in 2015.

A. Scarsella, W. Stofega, "Worldwide Smartphone Forecast Update , 2015-2019", IDC report, 2015.

Durability and life cycle

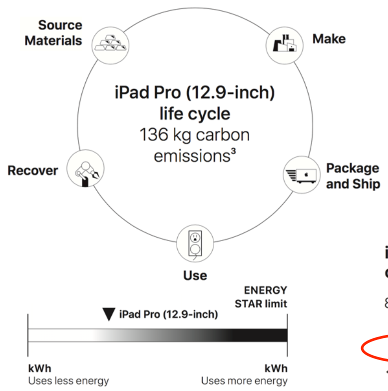

iPad Pro (12.9-inch) life cycle carbon emissions

83%	Production
11%	Transport
6%	Use
<1%	End-of-life processing

Product environmental report, Apple, 2018.

User = person responsible

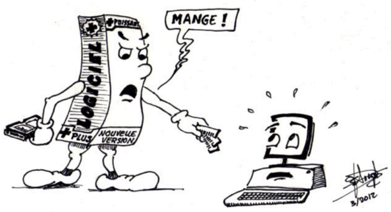

- Bloatware
- Obsolescence

In 2014, on average, 35 applications installed per smartphone, among which: 11 are used every week and 12 are never used.

Harris Interactive, 2015

Video games

Video games purshased on download
Do not fit on DVD any more
Network-hungry

[Source : Sandvine, The Global Internet Phenomena Report, 2018.]

What else?

What you can do

- Completely switch off unused devices
- Remove unused applications
- Erase useless (old) emails, photos, etc.
- Be careful when sending emails (attachments, receivers)
- Be careful when coding (image size, active loops, etc.)
- Look at eco-labels when buying new equipment
- Keep devices longer if they are still working
- Avoid capability overlap
- Stay energy-aware…

Opportunities

- To think differently
- To propose new things
- To build differently
- To design a sustainable future

interstices.info

ADEME

http://ecoinfo.cnrs.fr

Thank you for your attention

http://people.irisa.fr/Anne-Cecile.Orgerie