Eco-design or data centres collapse :

- environmental impact of digital
- 3 scenarii
- GreenIT challenges and role

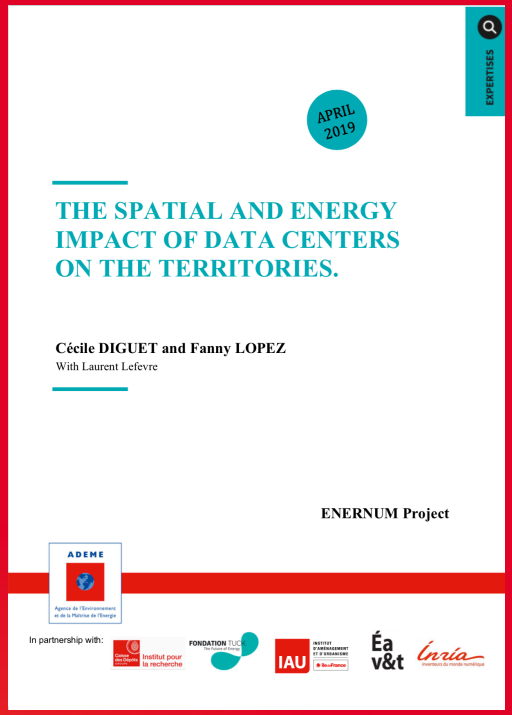

Laurent Lefèvre
laurent.lefevre@inria.fr

INRIA **AVALON / LIP**
Ecole Normale Supérieure de Lyon

Marne la Vallée, June 6, 2019

Is digital crazy (out of control) ?

CES2018 : building human++ of tomorrow : ultra-connected, ultra-observed, ultra-mobile -> ultra digital

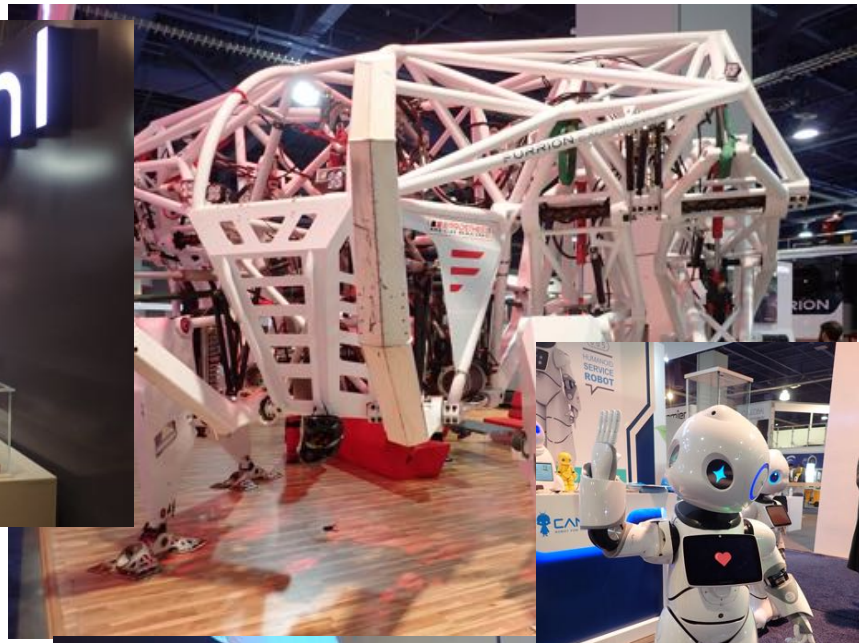

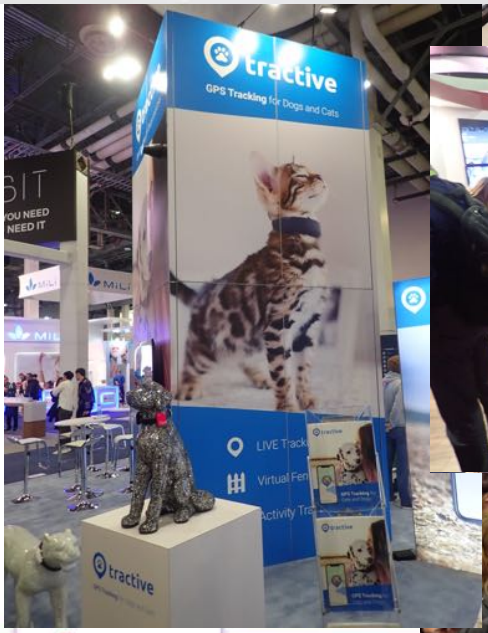

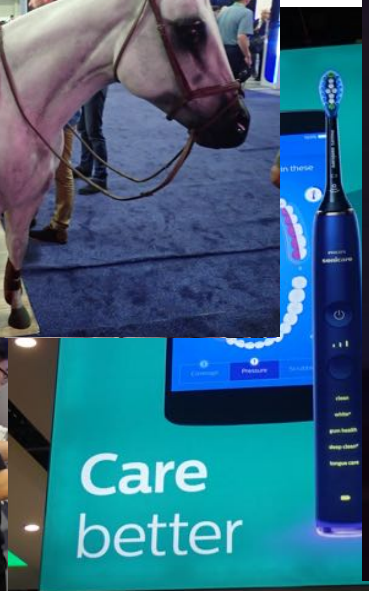

Is digital crazy (out of control) ?

The data/usage problem : the Pope effect !

The problem : the Pope effect !

The problem : the Pope effect !

Such much data to process, store and transport
All these data are taken by datacenters and networks

In one hour….

In France : 1,8 Mt of electronic waste in 2017 : 3.4 tonnes per minute

IoT : 15 B ?
2030 : 50-150 B ?

In 2010… the Cloud/Datacenter was no more virtual…

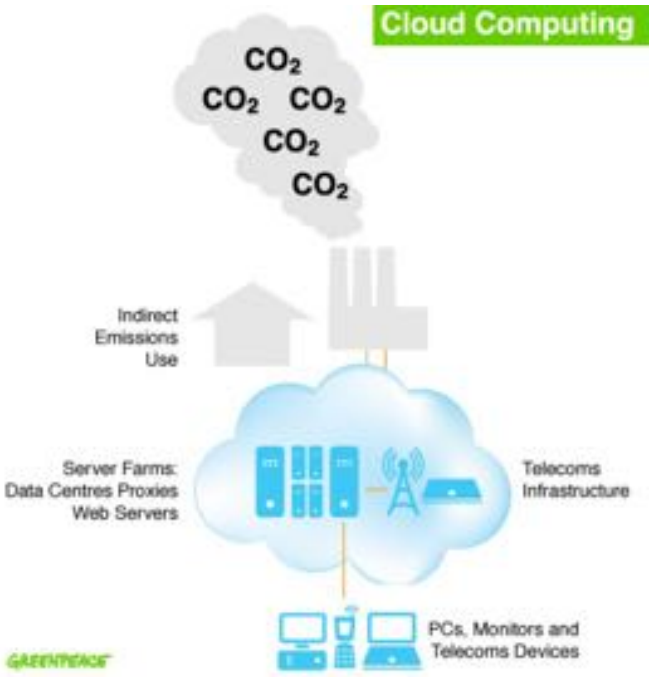

Greenpeace report : Make IT Green : Cloud Computing and its contribution to Climate Change (2010)

But the Cloud should be Green ! … the myth

- Virtualization of computers
- Virtualization of networks
- Virtualization of services
- Cloud is based on improved physical infrastructures and thus benefit from their EE
- Can benefit from renewable energy
- Aggregation / Consolidation effect

The reality : the Cloud faces a lot of issues

Uptime Institute, 2015

30% of US servers are « comatose »

« those that have not delivered information or computing services in six months or more »

- Overprovisioning of computing, storing and networking resources
- Optimizations technics are rarely applied in real clouds
- Providers afraid on resilience/users/QoS impact

Electricity Usage of Digital world

>10% of worlwide electricity

Increase by 8% per year

4% of GHG

Focusing on energy and power

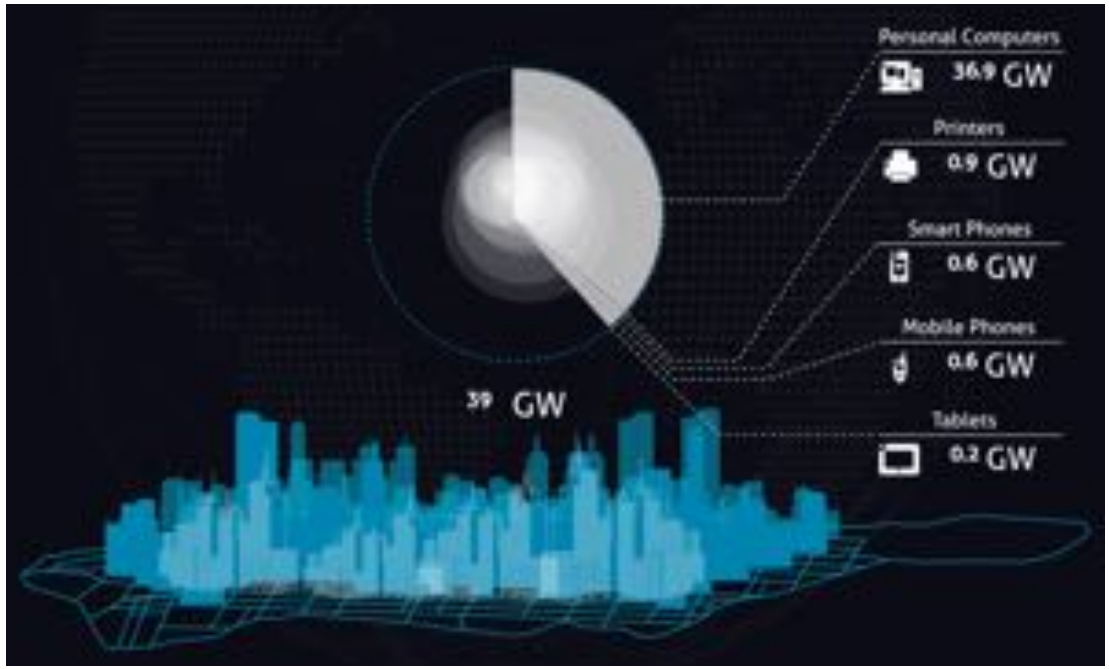

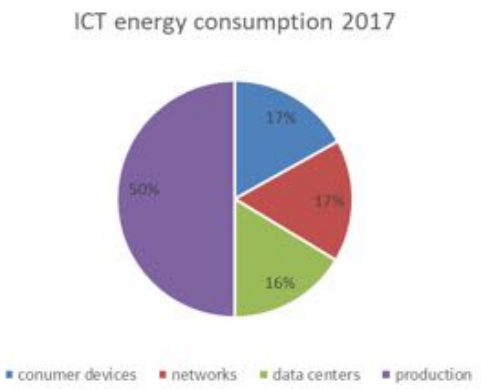

Source : leanICT, shift projet

Impact of ICT LifeCycle : the real full impact…

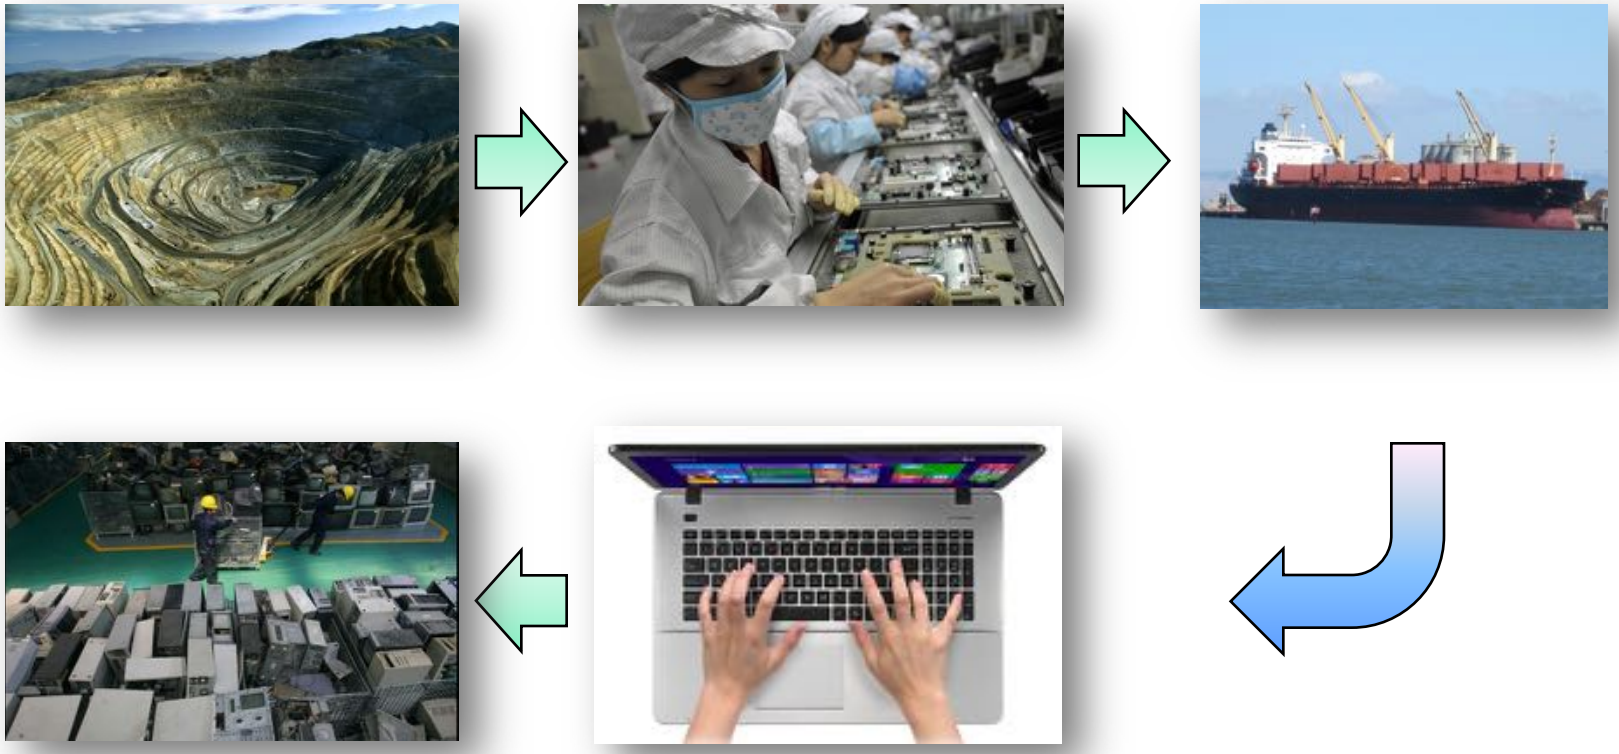

Multiple impacts

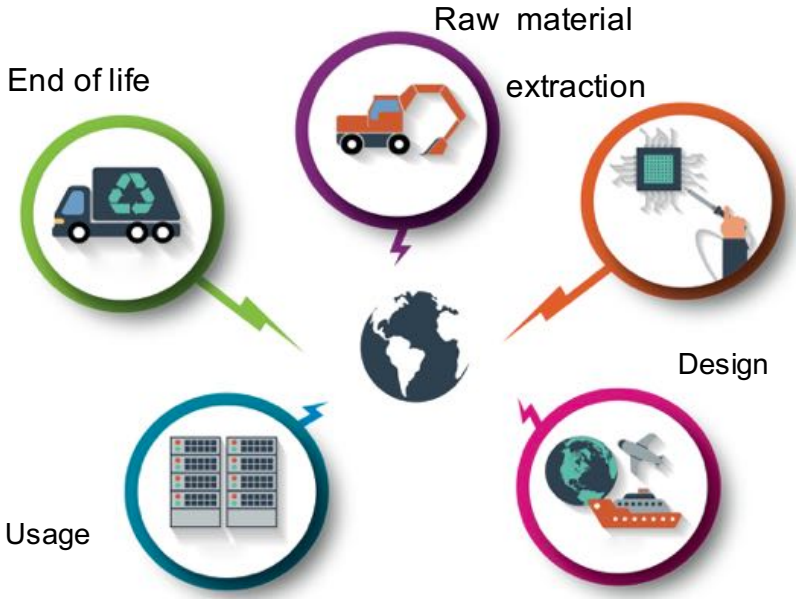

5 mandatory steps

Source : Discutforum

Source : Newsweek

With a lot of impacts and metrics

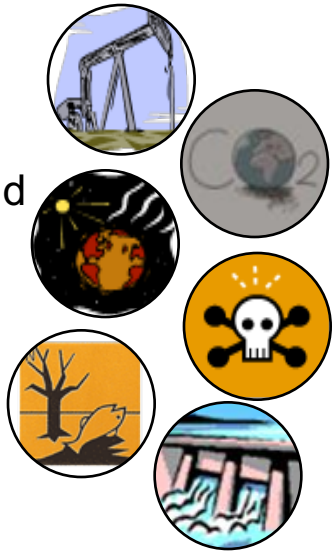

Limited stocks of metals and rare earths
17 out of 60 metals are recycled

Two visions of the smart city

(for a computer scientist)

Smart city is wonderful *(for a computer scientist)*

- Optimization of everything: transport, mobility, energy-> Innovative scenario
- Sensing everything : pollution, noise -> full of data
- Industry 4.0 -> money !
- Happy citizens, happy companies. E-learning !
- More sensors, terminals, screens, datacenters

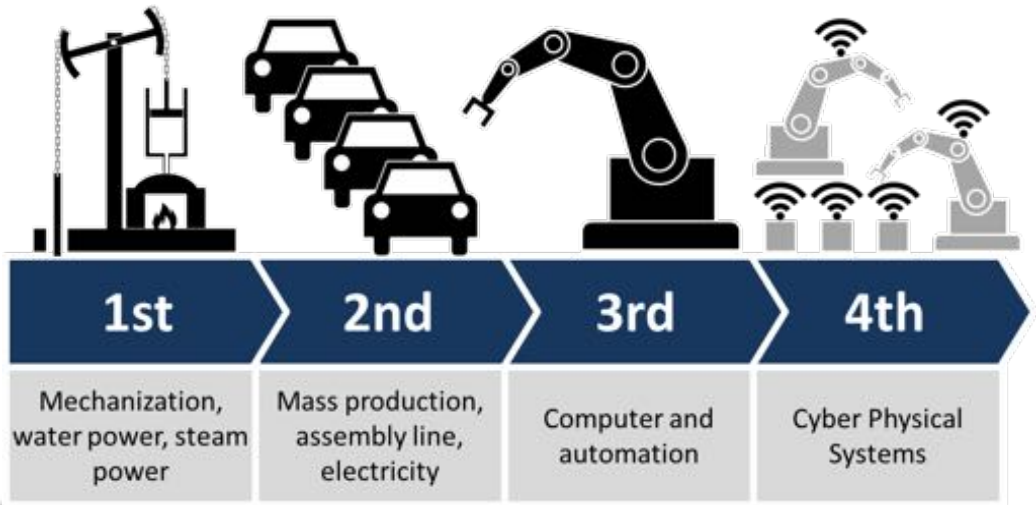

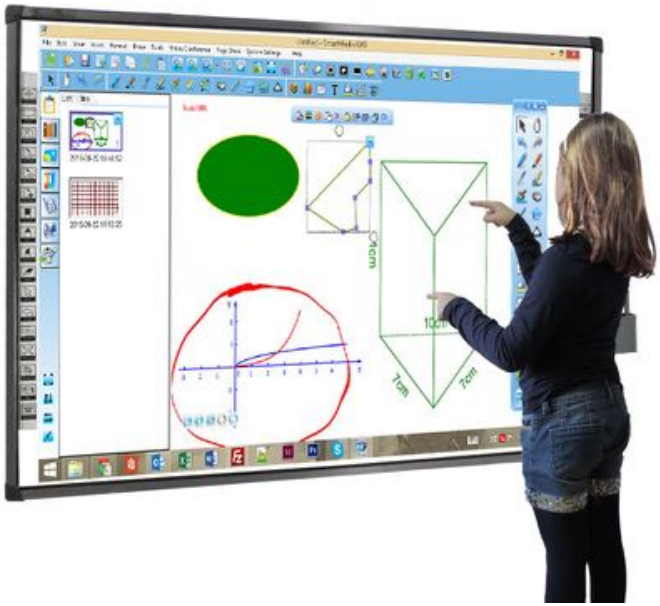

Smart city is awful *(for a computer scientist)*

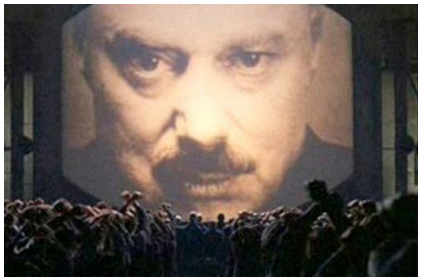

- More usage/pression on (numerical) services, short term experiments, BigData for what ?
- Energy consumption, heterogeneous technologies
- Too much sensors /domotic everywhere, that quickly fail…
- Cocktail of mobile networks > heath impact?
- Freedom of citizens

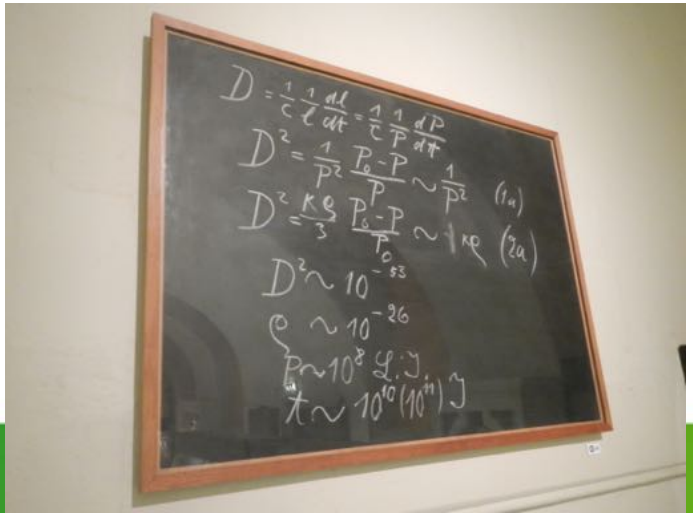

3 scenarii to support smart cities

- More & More : digital growth and ultra-centralization
- More & same : Stabilization of the digital technical system and infrastructural diversity
- Less & Less : digital ultra-decentralization: the end of data centers?

Scenario : More & More : digital growth and ultra-centralization

- More IOT, more services -> more cloud, more big datacenters
- AI everywhere ?
- IT4Green must work

Big Datacentres : ala Google ?

Google :

Few M servers ?

15 datacentres ?

All renewable energy

PUE (Power Usage Effectiveness): 1.12

Bigger than previous

- We can build what we want…
  - **China Telecom- Inner Mongolia Information Park :** 1.2 millions of servers - 3 Billion $ - 150 MW
  - **Bitmain / Antpool** : 60 MW ? for Bitcoin production
  - **Kolos Norway:** 2018-2019 : 70 MW -> 1 GW (100% RE)

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
6	**Trinity** - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect , **Cray Inc.** DOE/NNSA/LANL/SNL United States	979,072	20,158.7	41,461.2	7,578
7	**AI Bridging Cloud Infrastructure (ABCI)** - PRIMERGY CX2570 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, Infiniband EDR , **Fujitsu** National Institute of Advanced Industrial Science and Technology (AIST) Japan	391,680	19,880.0	32,576.6	1,649
8	**SuperMUC-NG** - ThinkSystem SD530, Xeon Platinum 8174 24C 3.1GHz, Intel Omni-Path , **Lenovo** Leibniz Rechenzentrum Germany	305,856	19,476.6	26,873.9	

Top500.org : Nov.2018

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	375	**Shoubu system B** - ZettaScaler-2.2, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-SC2 , **PEZY Computing / Exascaler Inc.** Advanced Center for Computing and Communication, RIKEN Japan	953,280	1,063.3	60	17.604
2	374	**DGX SaturnV Volta** - NVIDIA DGX-1 Volta36, Xeon E5-2698v4 20C 2.2GHz, Infiniband EDR, NVIDIA Tesla V100 , **Nvidia** NVIDIA Corporation United States	22,440	1,070.0	97	15.113
3	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	9,783	14.668
4	7	**AI Bridging Cloud Infrastructure (ABCI)** - PRIMERGY CX2570 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, Infiniband EDR , **Fujitsu** National Institute of Advanced Industrial Science and Technology (AIST) Japan	391,680	19,880.0	1,649	14.423
5	22	**TSUBAME3.0** - SGI ICE XA, IP139-SXM2, Xeon E5-2680v4 14C 2.4GHz, Intel Omni-Path, NVIDIA Tesla P100 SXM2 , **HPE** GSIC Center, Tokyo Institute of Technology Japan	135,828	8,125.0	792	13.704
6	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	7,438	12.723
7	446	**AIST AI Cloud** - NEC 4U-8GPU Server, Xeon E5-2630Lv4 10C 1.8GHz, Infiniband EDR, NVIDIA Tesla P100 SXM2 , **NEC** National Institute of Advanced Industrial Science and Technology Japan	23,400	961.0	76	12.681

Green500.org : Nov.2018

Scenario 2: More & same : Stabilization of the digital technical system and infrastructural diversity: quest for a difficult resilience

- Less hyperscale datacenters
- More alternet
- New connected objects and scenarii of usage : from autonomous cars

to security services -> need of low latency reaction (5G < 10 ms)

- Edge effect to support more services

Edge wave in data centers

Distributing DCs - impact of cooling / heat recycling – Hybrid systems (GPUs, FPGAs, Low Power processors)

Qarnot computing

Defab

Stimergy

Truck @ Delta Power Solutions

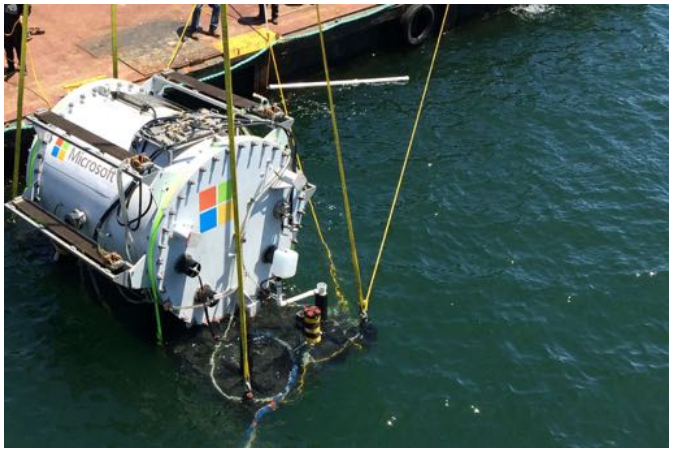

Natclick @ Microsoft

-> Fog : for E2E computing

Scenario 3 : less & less

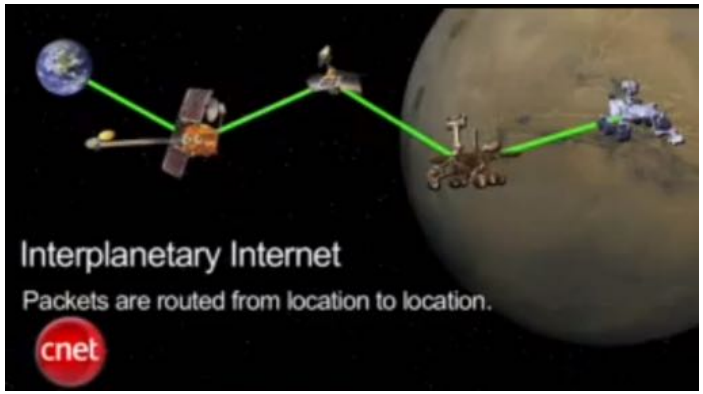

- Energy has become very expensive, scarce
- Best effort Internet, intermittent connectivity, lowtech design
- Improving “locavoracity” of service
- DTN like scenario with little infrastructure
- For long distance services : follow the wind, follow the sun, renewable fist/only

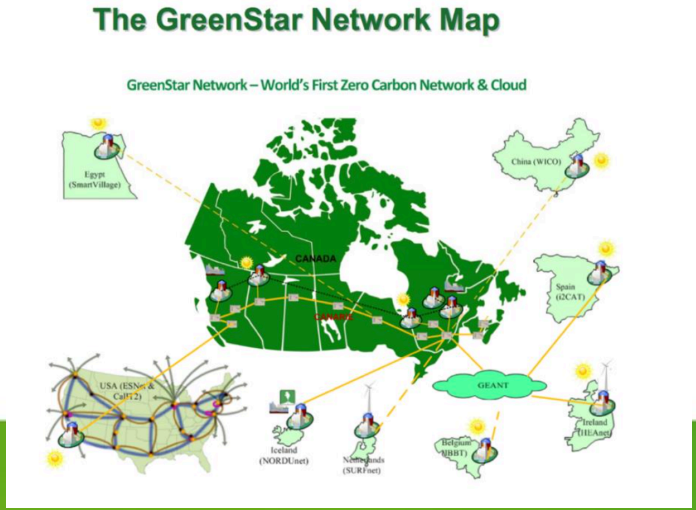

Connecting without infrastructure : DTN

- Ex company making money and providing services with DTN : (First Mille Solution)
- Services :
  - Offline web search
  - Emails
  - Voicemails/ video mails/ SMS

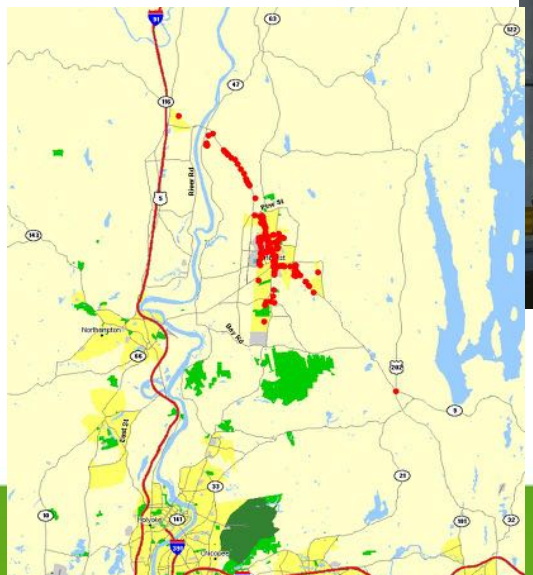

- UMASS / Amherst
- 40 buses
- Bus to bus throughput : 2 Mbits

Need to reduce energy impact of large scale infrastructures -> Energy Efficiency can help

GreenIT challenge and role with target on usage

- **Managing energy leverages**
- **Reaching energy proportionality**
- **Eco-design**

Managing multiple energy leverages

- **Shutdown** : to reduce the number of useless plugged resources
- **Slowdown** : to adapt performances (and energy consumption) to the real needs of applications, services and protocols
- **Optimizing** : to modify applications and services in order to make them greener
- **Consolidating / aggregating** : relocate services and applications on smaller number of physical resources

From green families to a lot of power capabilities

Green Families : Shutdown, Slowdown, Optimizing, Consolidating / aggregating :

- Node Shutdown
- Node Hibernation
- Node Suspend To Ram
- DVFS : Dynamic Voltage and Frequency Scaling
- NTV ; near threshold voltage
- AVX : Advanced Vector Extensions
- Low Power Idle
- Adaptive Link Rate
- Green scheduling policies
- Energy budget aware scheduling
- Power Capping
- Green Programming
- Simple / Double precision computing…

Adressing this leverages coordination is a big research challenge

The grail of Energy proportionality

Luiz André Barroso and Urs Hölzle, « The case for Energy-Proportional Computing », IEEE Computer, 2007

High idle consumption- Can be up to 50 % of peak power
Average server load between 10 and 50 % - Most inefficient region

Servers : Important Idle consumption – weakly proportional

Networks : constant consumption independant of usage

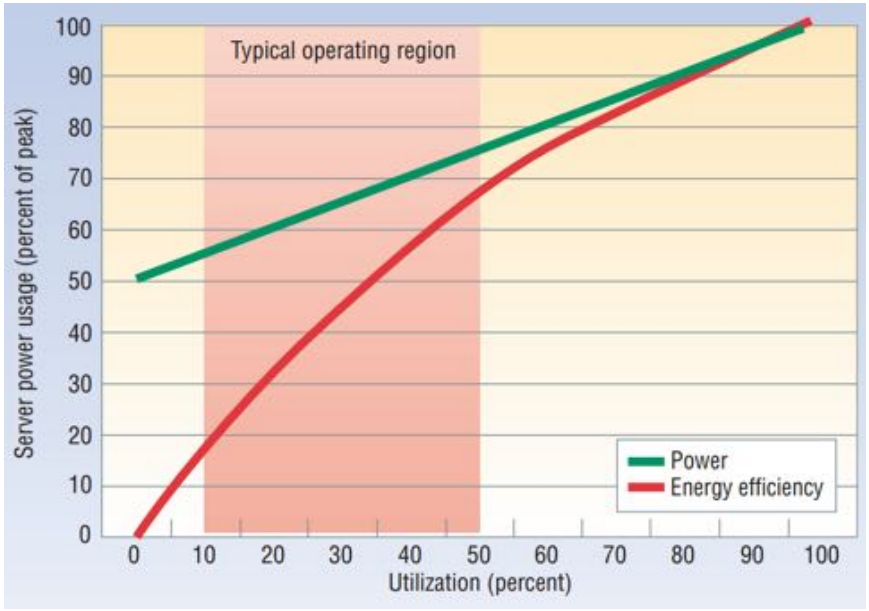

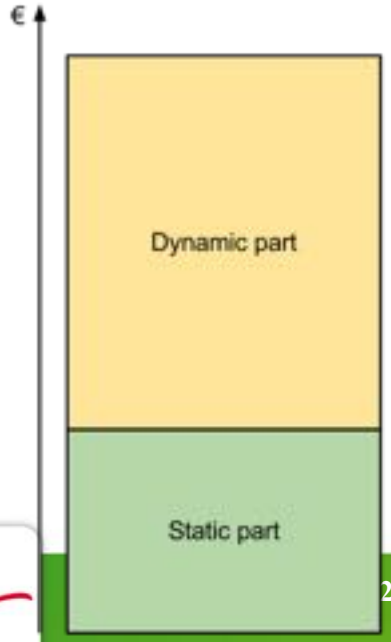

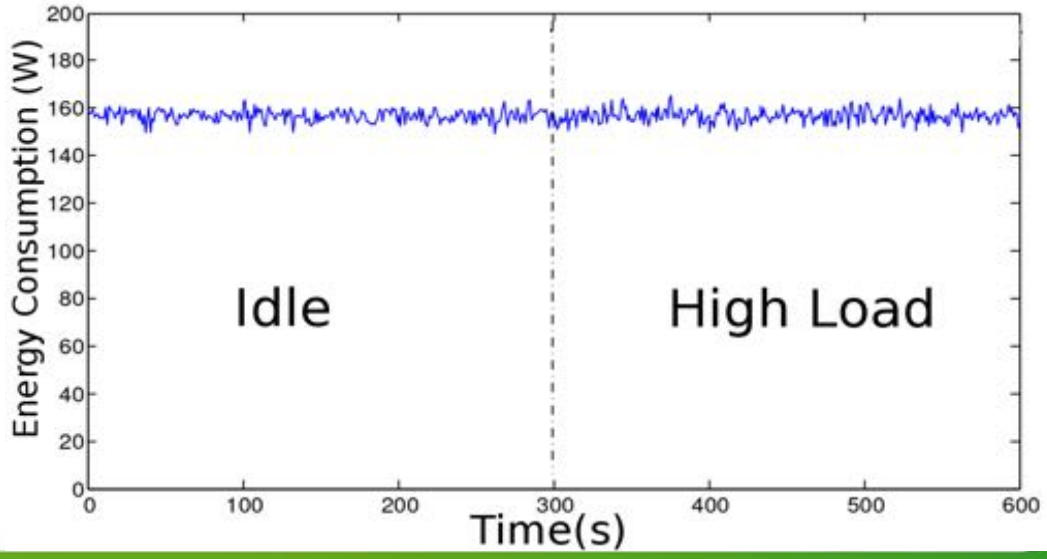

Eco design of software and hardware

Be ready to go slowly

Formula1 : Speed: 260 km/h @ 400 km/h - 75 à 100 liters/100 km

2CV : Speed: 115 km/h <5 liters/100km

Be ready to go lowtech as old computers software design from Paul&Bill and Steve&Steve…

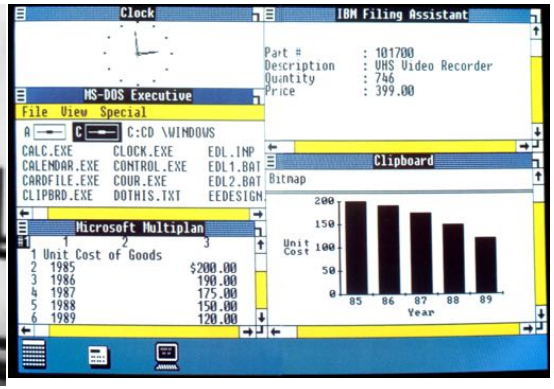

Windows 1 (85) et Mac OS 1 (84) : need of 256 KB of memory

GreenIT is not sufficient, usage must change… Where do we put our personal cursor of acceptance ?

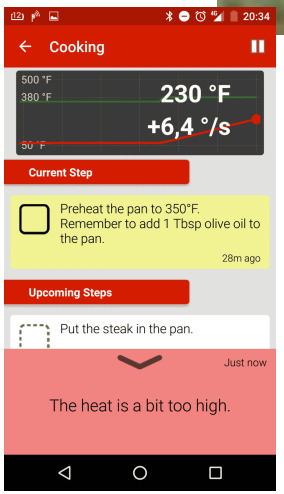

We are all frogs !

We must get out out of the pan (« Syndrome de Stockholm Numérique »)
(Jean-Romain Lhomme)

We could become colibris !

We are all frogs !

We must get out out of the pan (« Syndrome de Stockholm Numérique »)
(Jean-Romain Lhomme)

We have to become more than colibris !

What can we do ?

- Cloud/Datacenters should remain for a few more years (with various forms : from fog to edge to pervasive to invisible to…)
- Future could be a mix of the 3 (or more) scenarii

- Avoiding waste / overprovisioning
- Keeping equipments as much as possible
- Taking care of rebound effects : : improving one step of lifecycle could degrade other step and thus increase energy consumption of lifecycle

- Quantifying IT4Green

- What about compensating carbon : planting trees ? Not clear…
  - A french citizen : 12 t eq CO2 per year – some tree can store 30 Kg of CO2 per year -> need to plant 400 trees and redo it/increase it again every 30-40 years ?

Thanks to : F. Berthoud, E. Drezet, A-C. Orgerie, I. Rais

Laurent.lefevre@inria.fr