GreenFactory : orchestrating power capabilities and energy leverages at large scale for energy efficient infrastructures

Laurent Lefèvre
laurent.lefevre@inria.fr

with I. Rais, A. Benoit and A.-C. Orgerie

CCDSC2016, Dareizé, October 4, 2016

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

Thanks to Jack and Bernard !

An urgent issue !

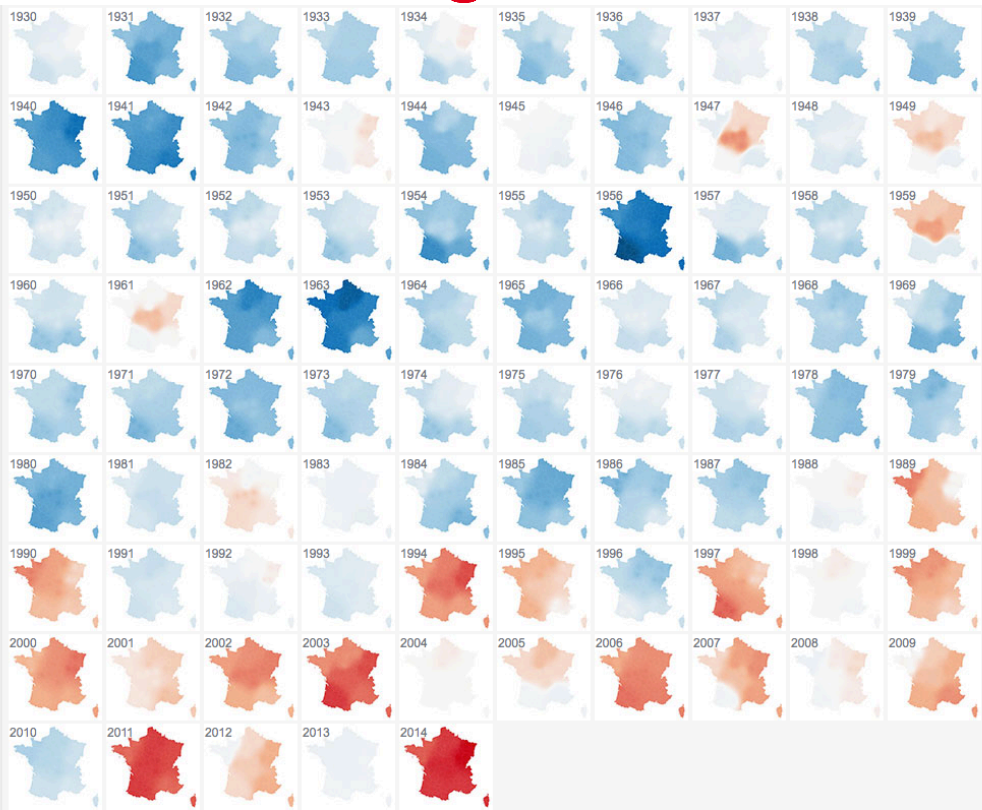

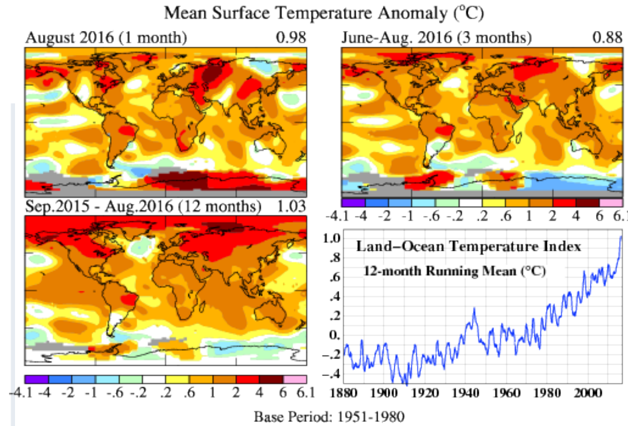

COP21 · CMP11
PARIS 2015
UN CLIMATE CHANGE CONFERENCE

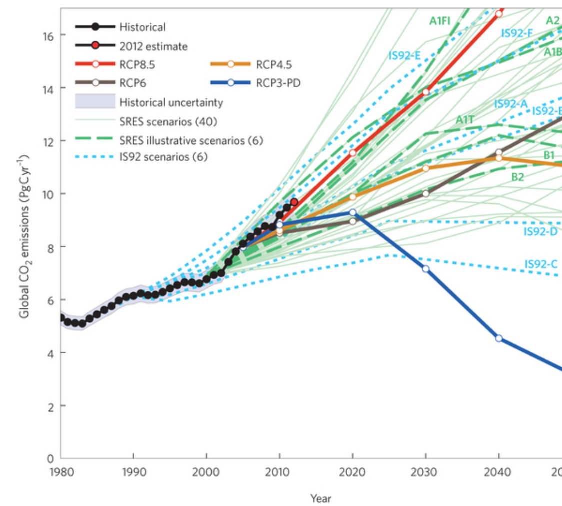

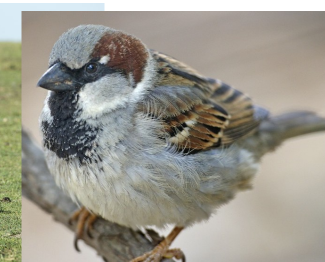

Links with me/ICT ?

Lifecycle of ICT (Info. & Comm. Tech.)

5 mandatory steps

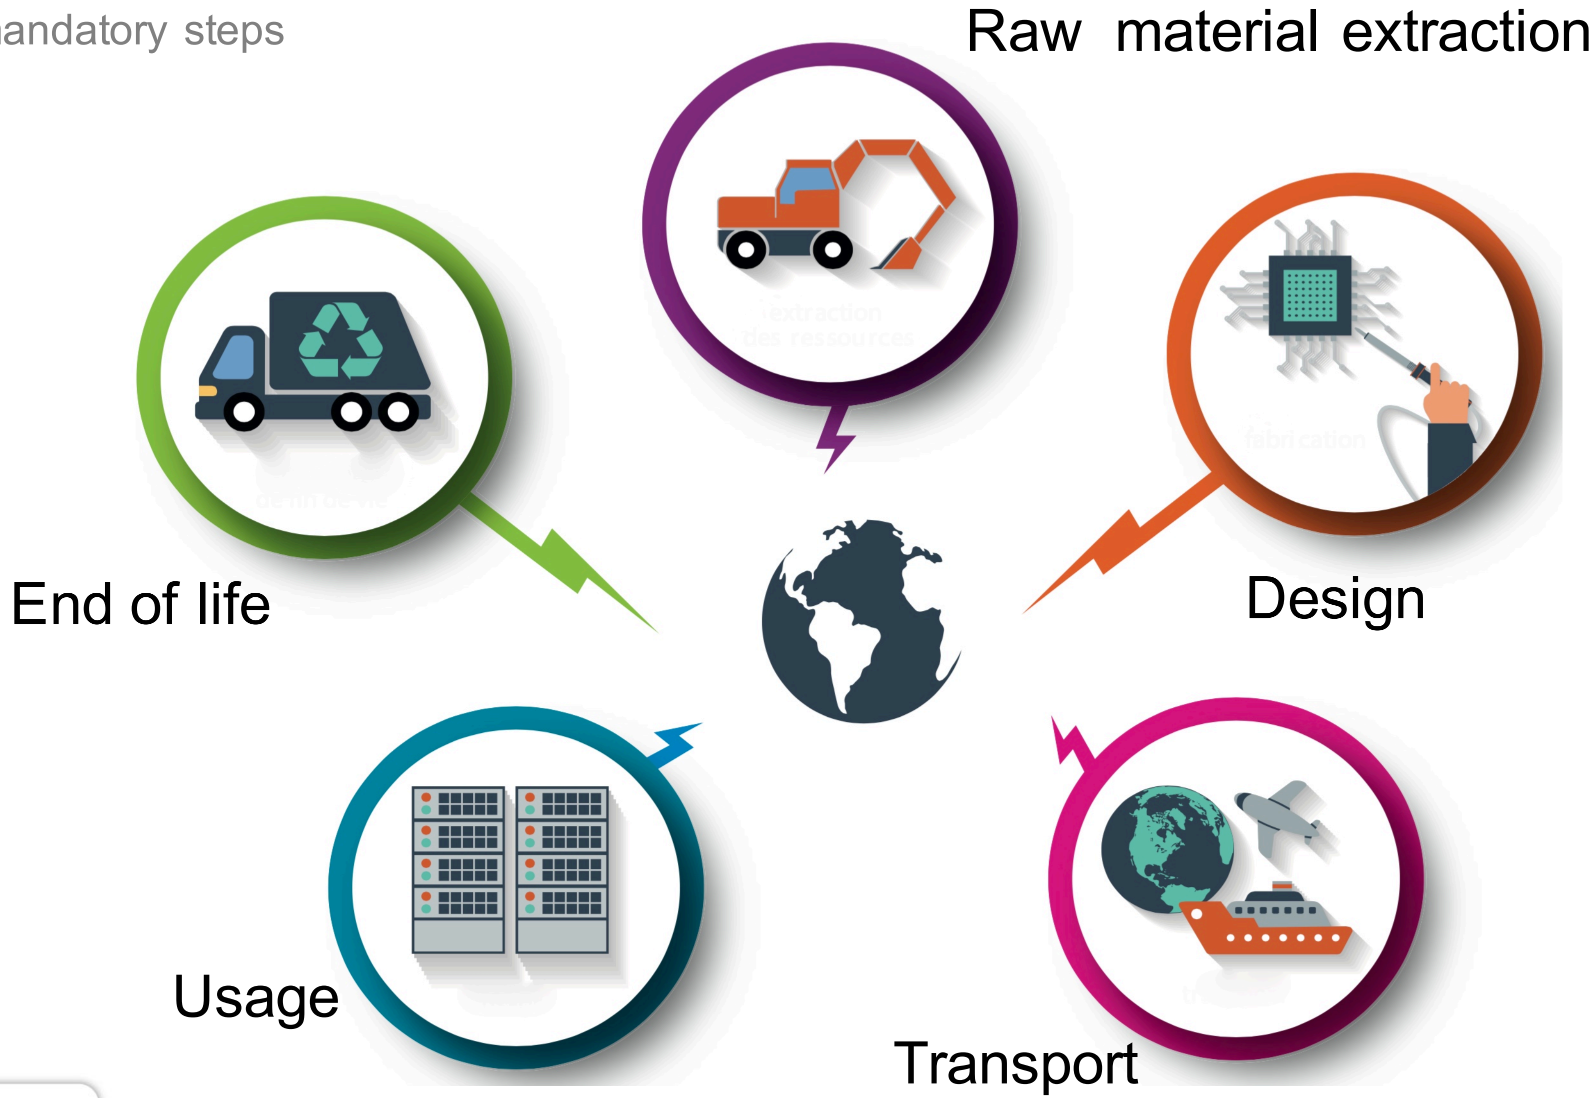

Life Cycle multi criteria impacts / Metrics

Primary energy consumption

Climate changing : emission of greenhouse gases

Ozone layer impact

Human toxicity : in air water, soil with potential risk on humans

Eco / Biodiversity toxicity : emission in air, water and soil with potential risk for aquatic fauna and flora

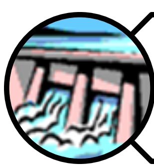

Water consumption

Lifecycle of ICT : Raw Material extraction

ICT need plastic, silica and metals

Extracting metals face Mineralogical barrier !

Ex : Copper : 1930 : 18 kg per ton – 2010 : 8 kg per ton

Less concentration of metal : more energy needed for extraction
10 % of primary energy is used for metal extraction and refinery
Geopolitical tensions / Metals from conflicts zones

Ex : China produces 97% of rare earths / 75 % Gallium / 71% Germanium / restrictions and quotas
Example : indium (coming from zinc) :
- 2007 : estimation of reserve at 6000 ton : prices explode
- New deposits become economically exploitable : prices drop
- Now reserves estimated at 50000 tons

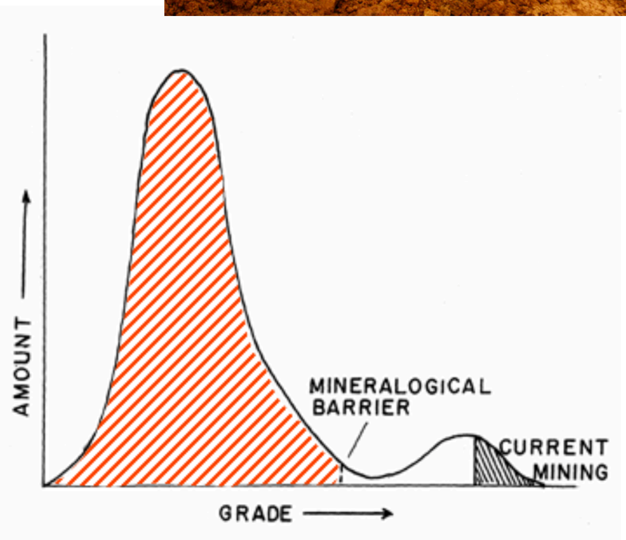

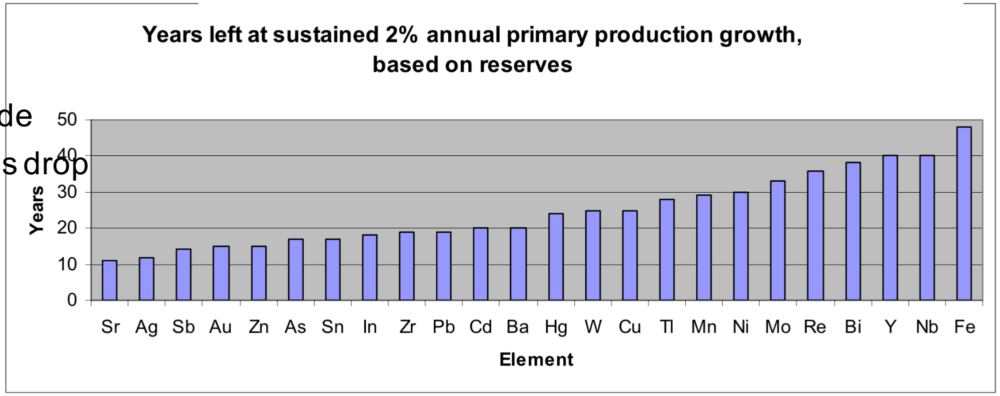

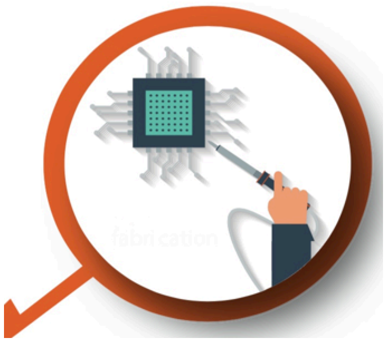

Lifecycle of ICT : Design

In 70's and 80's < 10 metals used for ICT
Since 2000's > 50 metals and rare earths

Better, smaller, faster, resilient…

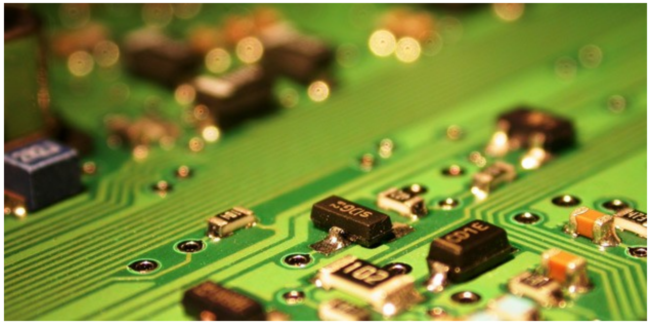

1 H																	2 He
3 Li	4 Be											5 B	6 C	7 N	8 O	9 F	10 Ne
11 Na	12 Mg											13 Al	14 Si	15 P	16 S	17 Cl	18 Ar
19 K	20 Ca	21 Sc	22 Ti	23 V	24 Cr	25 Mn	26 Fe	27 Co	28 Ni	29 Cu	30 Zn	31 Ga	32 Ge	33 As	34 Se	35 Br	36 Kr
37 Rb	38 Sr	39 Y	40 Zr	41 Nb	42 Mo		44 Ru	45 Rh	46 Pd	47 Ag	48 Cd	49 In	50 Sn	51 Sb	52 Te	53 I	54 Xe
55 Cs	56 Ba	*	72 Hf	73 Ta	74 W	75 Re	76 Os	77 Ir	78 Pt	79 Au	80 Hg	81 Tl	82 Pb	83 Bi			
		**															
*Lanthanides (Rare earths)			57 La	58 Ce	59 Pr	60 Nd		62 Sm	63 Eu	64 Gd	65 Tb	66 Dy	67 Ho	68 Er	69 Tm	70 Yb	71 Lu
**Actinides				90 Th		92 U											

Lifecycle of ICT : End of Life

Best European ICT recycler : Umicore (Belgium) : 17 out of 60 metals can be recycled and re injected/sold to the market

Dismantling / refurbishing / recycling ….Is expensive … some countries still export WEEE despite Basel convention

- < 1%
- 1 – 10%
- 10 – 25%
- 25 – 50%
- > 50%

1 H																	2 He
3 Li	4 Be											5 B	6 C	7 N	8 O	9 F	10 Ne
11 Na	12 Mg											13 Al	14 Si	15 P	16 S	17 Cl	18 Ar
19 K	20 Ca	21 Sc	22 Ti	23 V	24 Cr	25 Mn	26 Fe	27 Co	28 Ni	29 Cu	30 Zn	31 Ga	32 Ge	33 As	34 Se	35 Br	36 Kr
37 Rb	38 Sr	39 Y	40 Zr	41 Nb	42 Mo		44 Ru	45 Rh	46 Pd	47 Ag	48 Cd	49 In	50 Sn	51 Sb	52 Te	53 I	54 Xe
55 Cs	56 Ba	*	72 Hf	73 Ta	74 W	75 Re	76 Os	77 Ir	78 Pt	79 Au	80 Hg	81 Tl	82 Pb	83 Bi			
		**															
*Lanthanides (Rare Earths)			57 La	58 Ce	59 Pr	60 Nd		62 Sm	63 Eu	64 Gd	65 Tb	66 Dy	67 Ho	68 Er	69 Tm	70 Yb	71 Lu
**Actinides				90 Th		92 U											

Source : UNEP / Recycling rates of metals 2011

Lifecycle of ICT : Usage

2013… gwatt.net

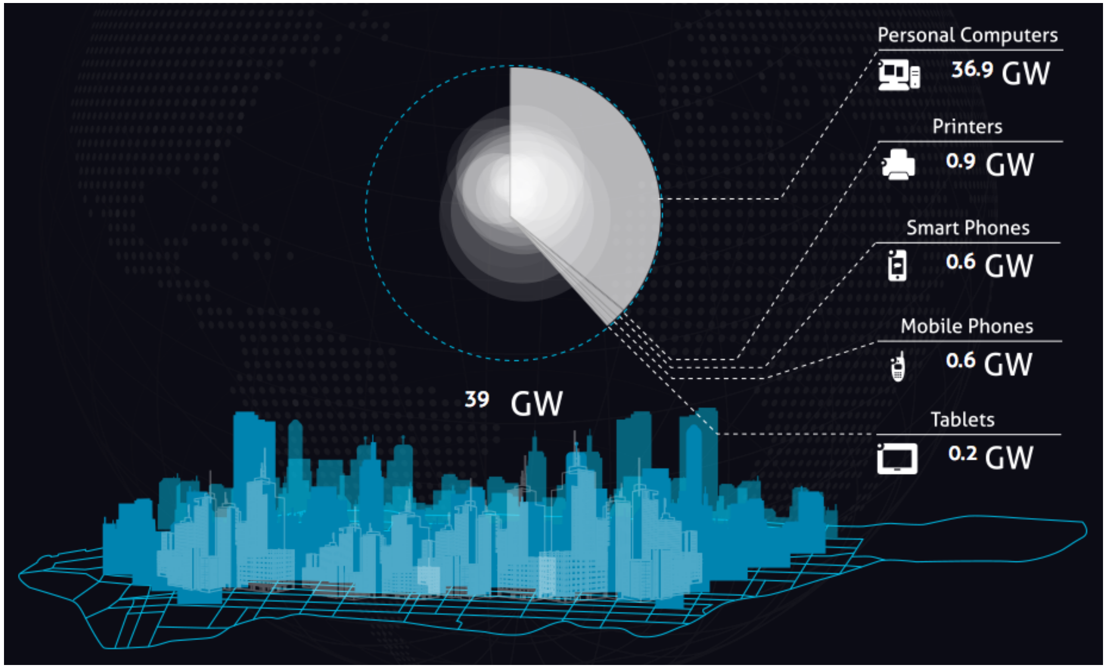

Only considered metrics :

- power
- energy

….watts, watts/h, joules

Lifecycle of ICT : Usage : a lot of electricity or not ?

3 consuming families :
- end users terminals
- network
- datacenters

== 100 GW !

Nuclear reactor (in France) : 0.9 to 1.4 GW – 59 reactors / 19 power plants
Future Hinkley Point C Central : 2 * 1.6 GW (7% of UK electricity- for 2023?)
World : 250 nuclear power plants / 441 nuclear reactor : 382 GW

ICT == 100 nuclear reactors / 5% (10%?) of consumed electricity → 100 billions of euros of electricity per year

Double waves in data centers

Merging / distributing - impact of cooling / heat recycling

Qarnot

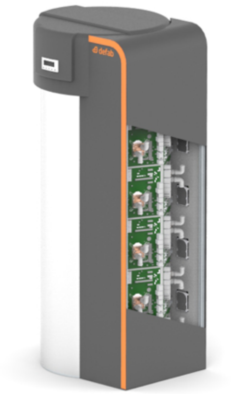

Defab

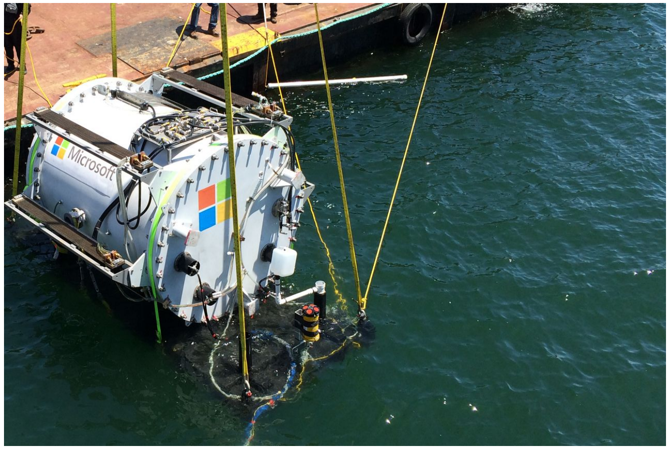

Microsoft

-> Fog/Edge

Infrastructure impacts the DC consumption

Less resources on end users equipment ->

more cloud services -> more data centres

Once infras are optimized .. The IT game can happen … Energy capabilities families during ICT usage..

- **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

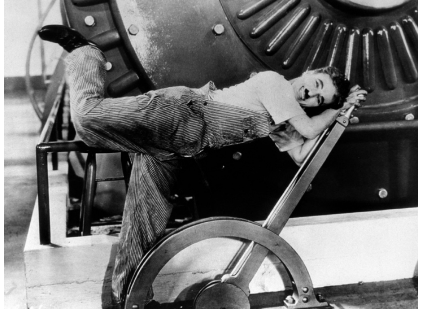

- Node Shutdown
- Node Hibernation
- Node Suspend To Ram
- DVFS : Dynamic Voltage and Frequency Scaling
- NTV ; near thresold voltage
- AVX : Advanced Vector Extensions
- Low Power Idle
- Adaptive Link Rate
- Green scheduling policies
- Energy budget aware scheduling
- Power Capping
- Green Programming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

From green families to a lot of power capabilities

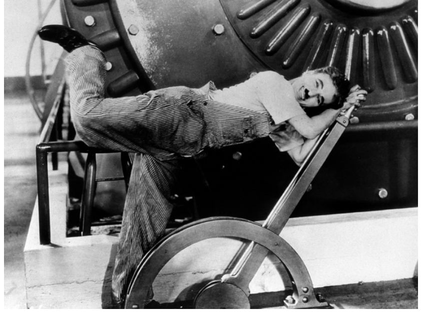

- Node Shutdown
- Node Hibernation
- Node Suspend To Ram
- DVFS : Dynamic Voltage and Frequency Scaling
- NTV ; near threshold voltage
- AVX : Advanced Vector Extensions
- Low Power Idle
- Adaptive Link Rate
- Green scheduling policies
- Energy budget aware scheduling
- Power Capping
- Green Programming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

Energy leverages !

Problem : Most of scientific literature only focus on Mono energy leverage usage …

ex : DVFS : (at least) 241 (IEEE) articles focusing on how to use that technology to reduce energy consumption

The target : Creating GreenFactory an orchestration framework

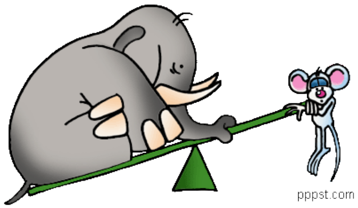

How to orchestrate the usage of energy leverages available through time on modulable IT resources to answer fixed constraints ?

Orchestrate : Multi level (Hardware, Middleware, Application) - Based on what knowledge/prediction ? Static / online

Leverages : the most relevant ones ? Leverages in opposition ?

Green Leverages for who ?

User point of view :

- How could we manage to help him choose the right solution ?
- How to use all these leverages that could help users lower the energy consumption of application ?
- How could we lower or regulate energy consumption using a wise usage of chosen leverages while maintaining good performances ?

Admin /providers :

- How to configure management tools present on a supercomputer (in a energy aware way) ?
- How could we abstract leverages usage to the user ?

Focus on 2 energy leverages :

Shutdown

Components based programming

Shutdown : the “easiest one” ?

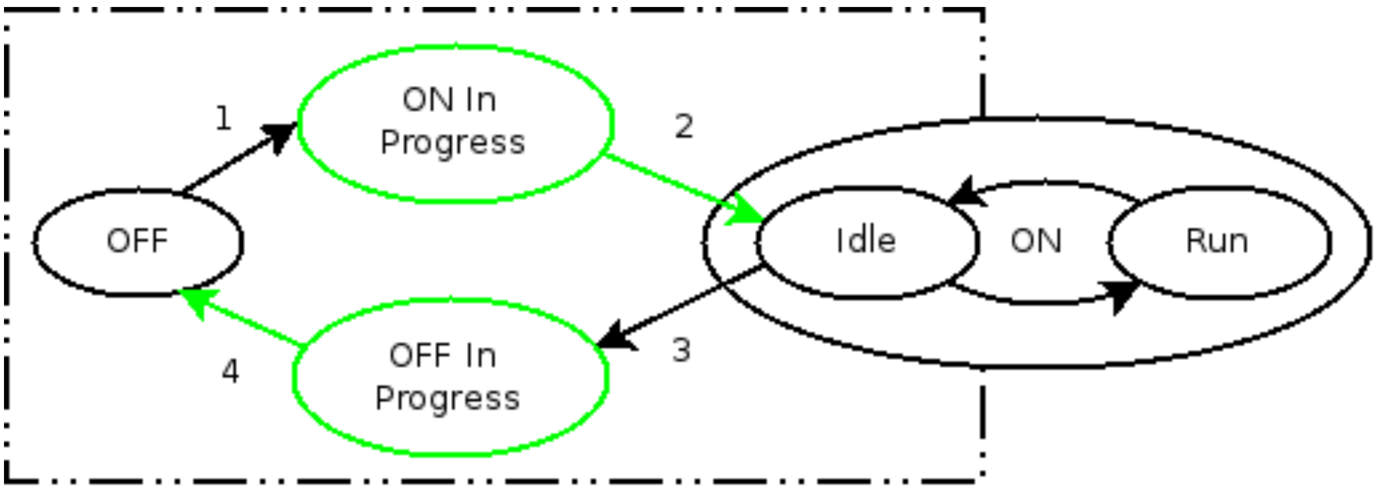

Challenge : agressive on/off are not always the best solution !

- Must be intelligently managed and applied
  - Exploiting gaps between bursts
  - Only switching off if energy gains are possible
  - Impact on equipment resilience / impact on cooling
  - Need to understand usage and to predict them

- Same approaches and methodology :
  - When and what to switch off ? (widely explored in SoA)
    - How to support some significant energy reduction ?

  - When and what to switch on ? (more difficult)
    - Which network parts (in order to avoid congestion) ?
    - Computing servers to avoid over provisioning…

OnOff Models: more complex than expected !

Pattern, power capping and energy aware : iif the overall nodes under supervision are respecting the fixed constraints

Electric Provider aware : iif the modular constraints imposed by the electric provider are respected

Cooling aware : iif the recommended thermal envelop is respected

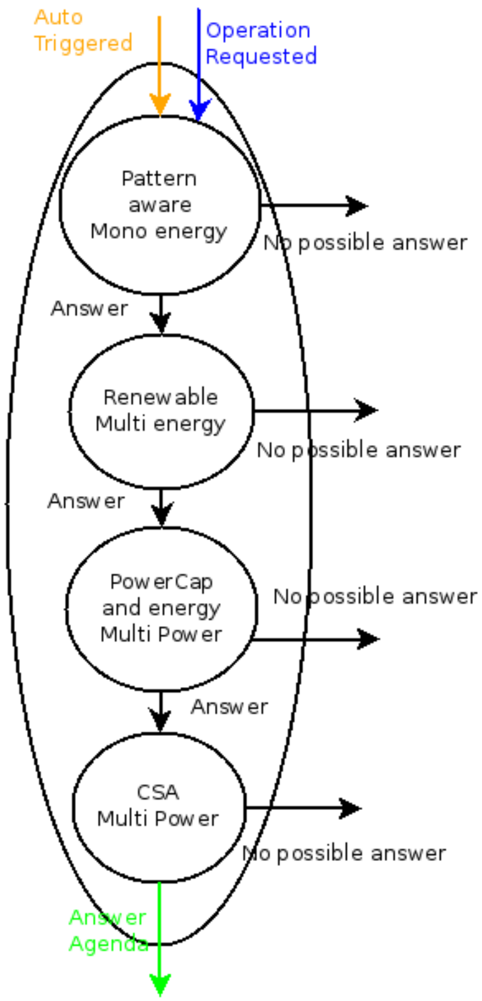

Component based programming : a possible energy leverage ?

Building the required service with EE

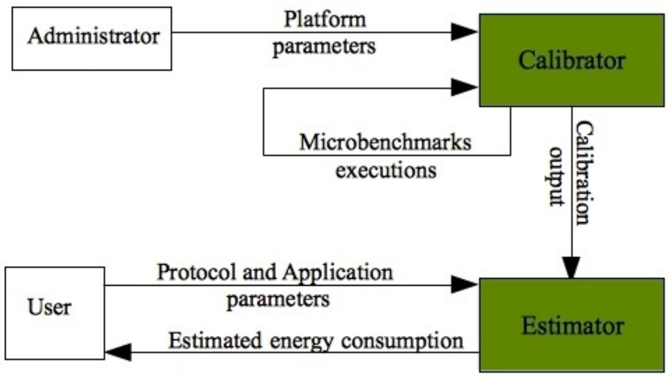

- Considered services : resilience & data broadcasting
- 4 steps: **Service analysis (decomposing), Measurements, Calibration, Estimation**
- Helping users make the right choices depending on context and parameters

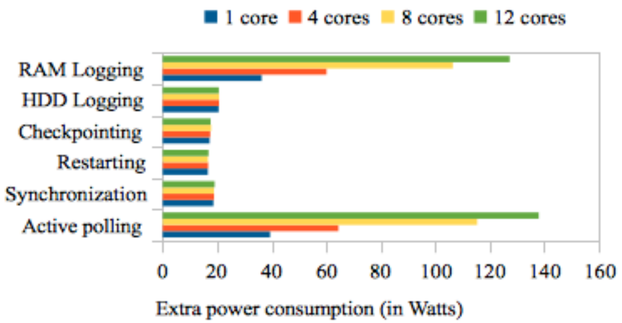

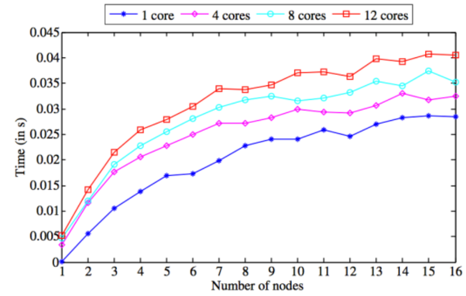

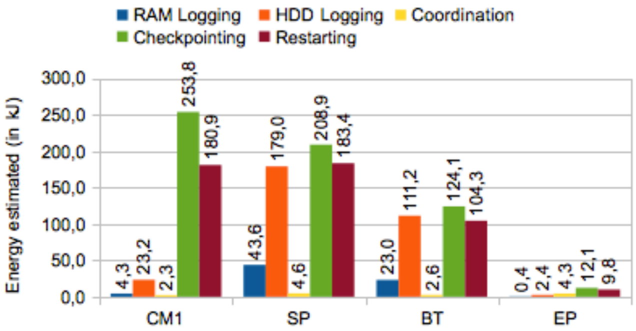

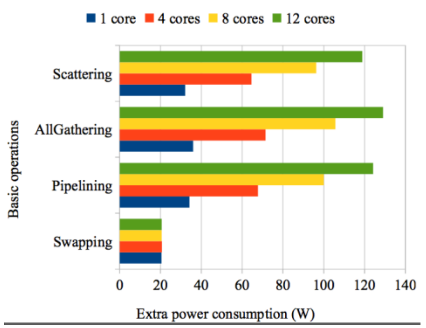

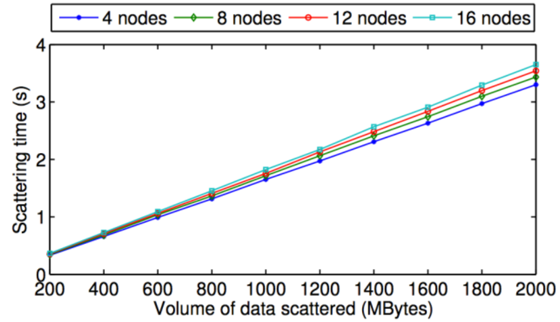

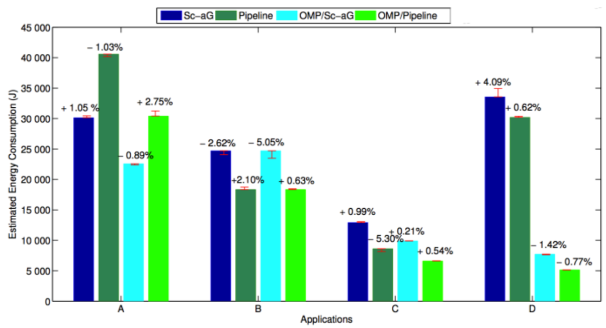

M. Diouri, Olivier Glück, Laurent Lefevre, and Franck Cappello. **"ECOFIT: A Framework to Estimate Energy Consumption of Fault Tolerance Protocols during HPC executions"**, *CCGrid2013, the 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, Delft, the Netherlands, May 13-16, 2013

From Green Programming to composant assembly as a leverage

« Green programming » is an interesting leverage
But very hard to apply in real world

Aim : Ease energy control
Through component-based application
Generating various versions / assembly of components / configuration

An example : The Shallow water application
- Component based application
- Implemented with a component based multi stencil language (MSL)

First results : find the best EE version ?

4 machines

Grid5000

4 versions of assembly

Varying configuration

#proc / # threads

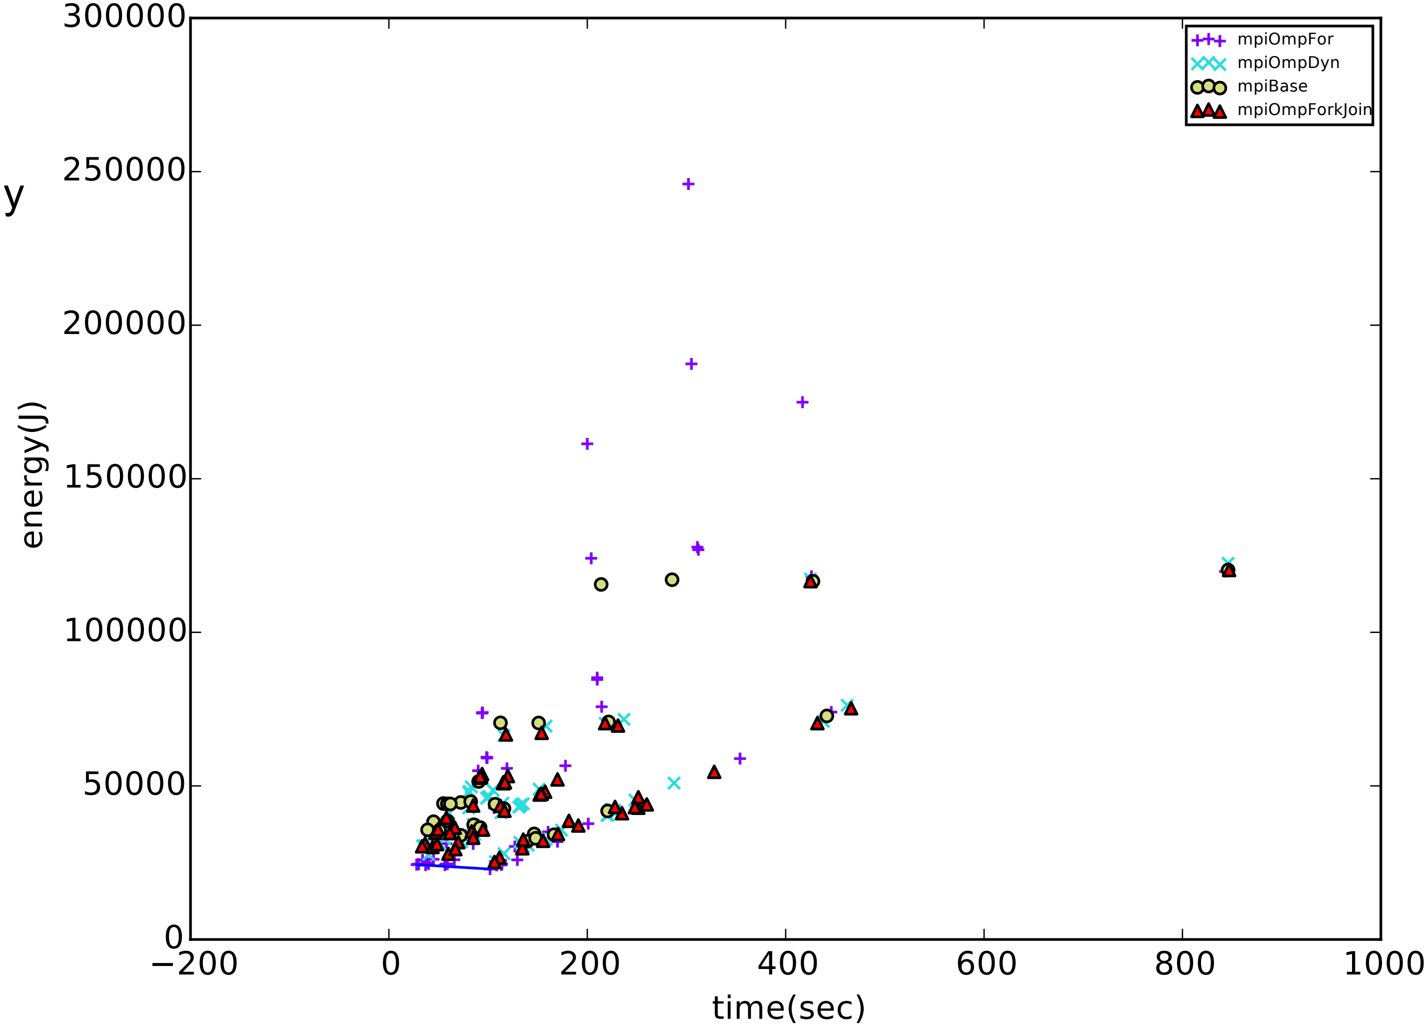

On the Pareto front

Between energy and average powercap

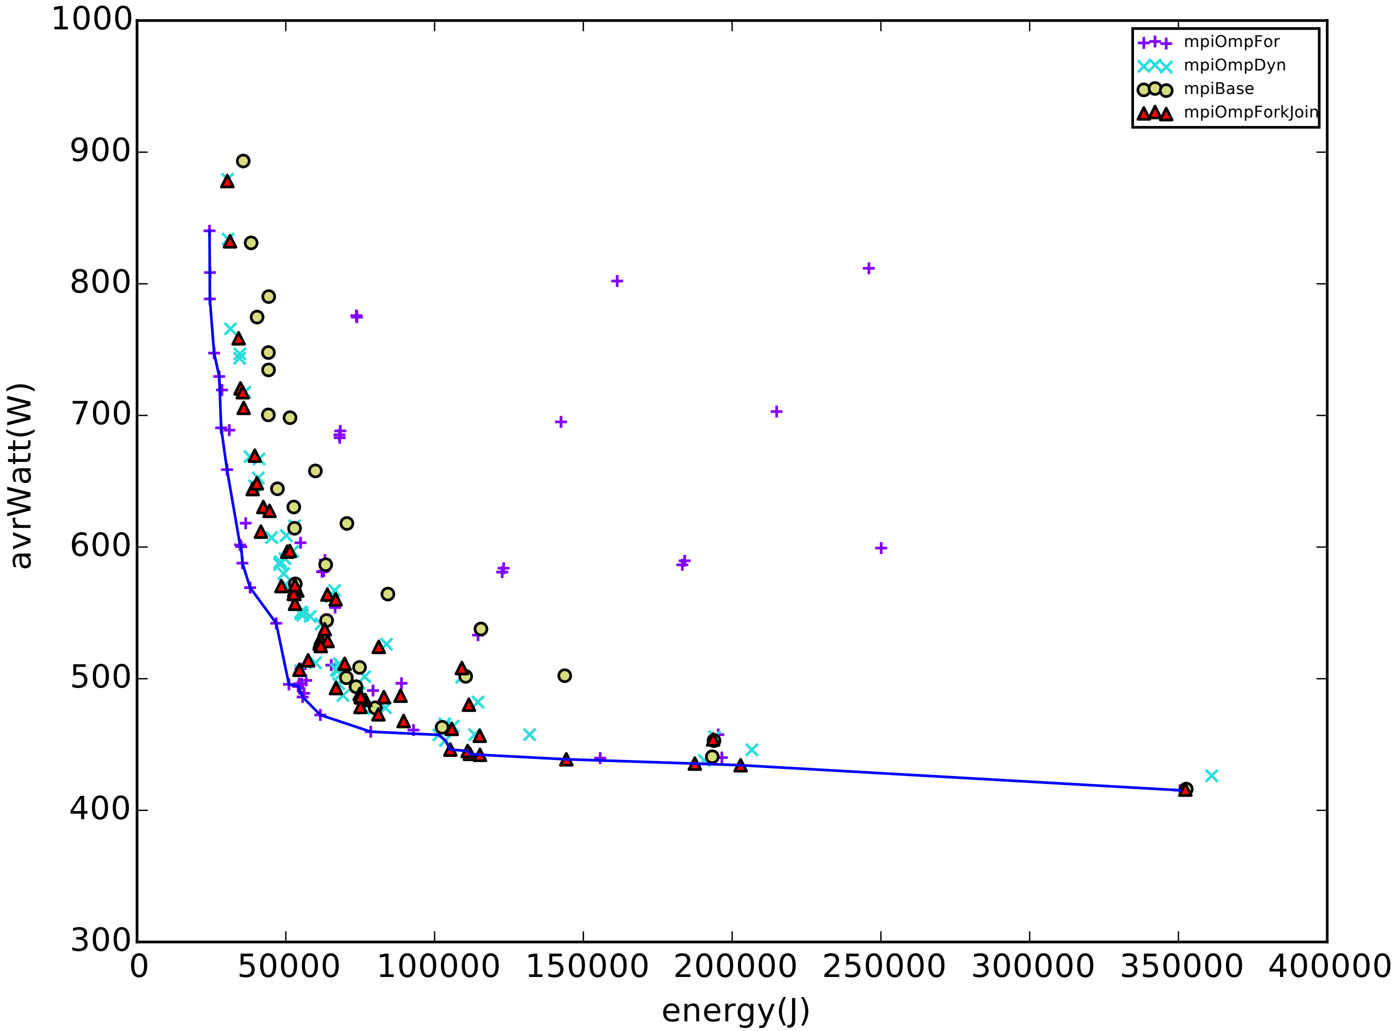

For a power capping scenario

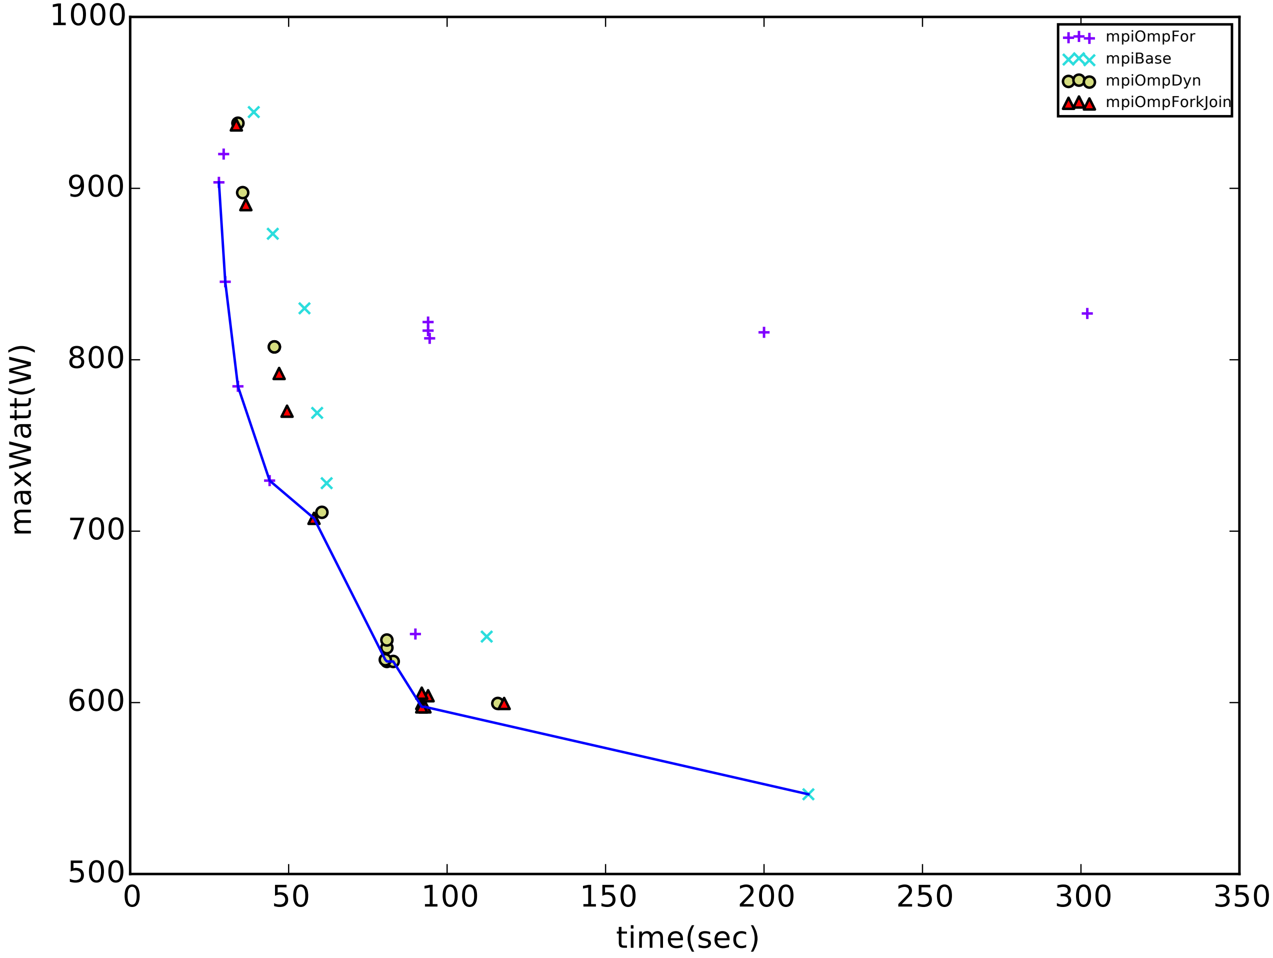

Larger scale

Curie Machine (1024 used procs)

4 versions of assembly

Varying configuration

#proc / # threads

Expensive to explore more assemblies

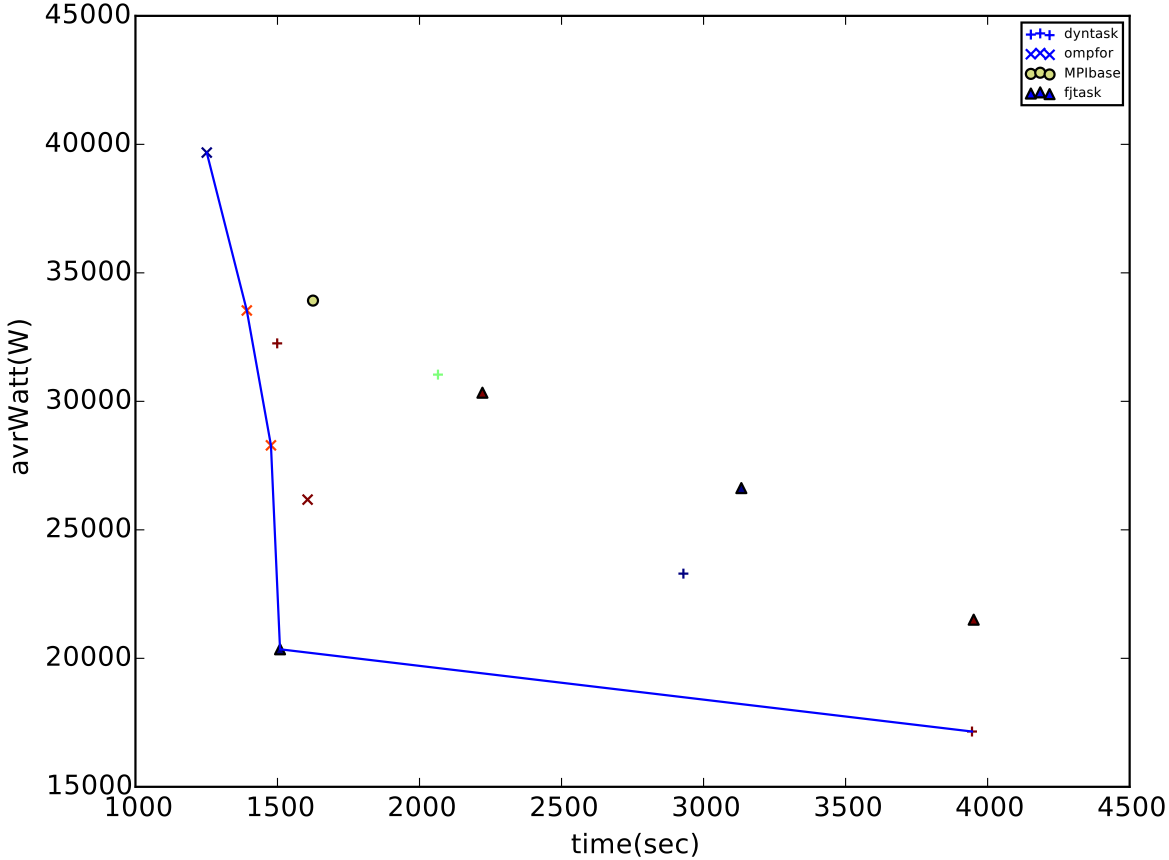

Conclusions

- ICT is not clean and has multiple environmental impacts
- Being clear and realistic : usage is only one part of the problem
- An improved and longer usage of ICT with the help of energy leverages orchestration is mandatory

Thanks to : M. Dias de Asuncao, A. Benoit, F. Berthoud, H. Coullon, M. Diouri, E. Drezet, A.-C. Orgerie, G. Tsafack, C. Perez, I. Rais

INVENTEURS DU MONDE NUMÉRIQUE

Laurent.lefevre@inria.fr