

REDUCING THE ENERGY IMPACT OF VIRTUALIZED DATA CENTERS

Anne-Cécile Orgerie

E3-RSD school
4th October 2018, Dinard



Who I am

- Full-time researcher at CNRS (about 34,000 people)
- Located in Rennes, France.
- IRISA laboratory (about 800 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>

Outline



- I. Energy consumption in data centers**
- II. Measuring power consumption of servers**
- III. Modeling power consumption of servers**
- IV. Concluding remarks**

ICT's electricity consumption

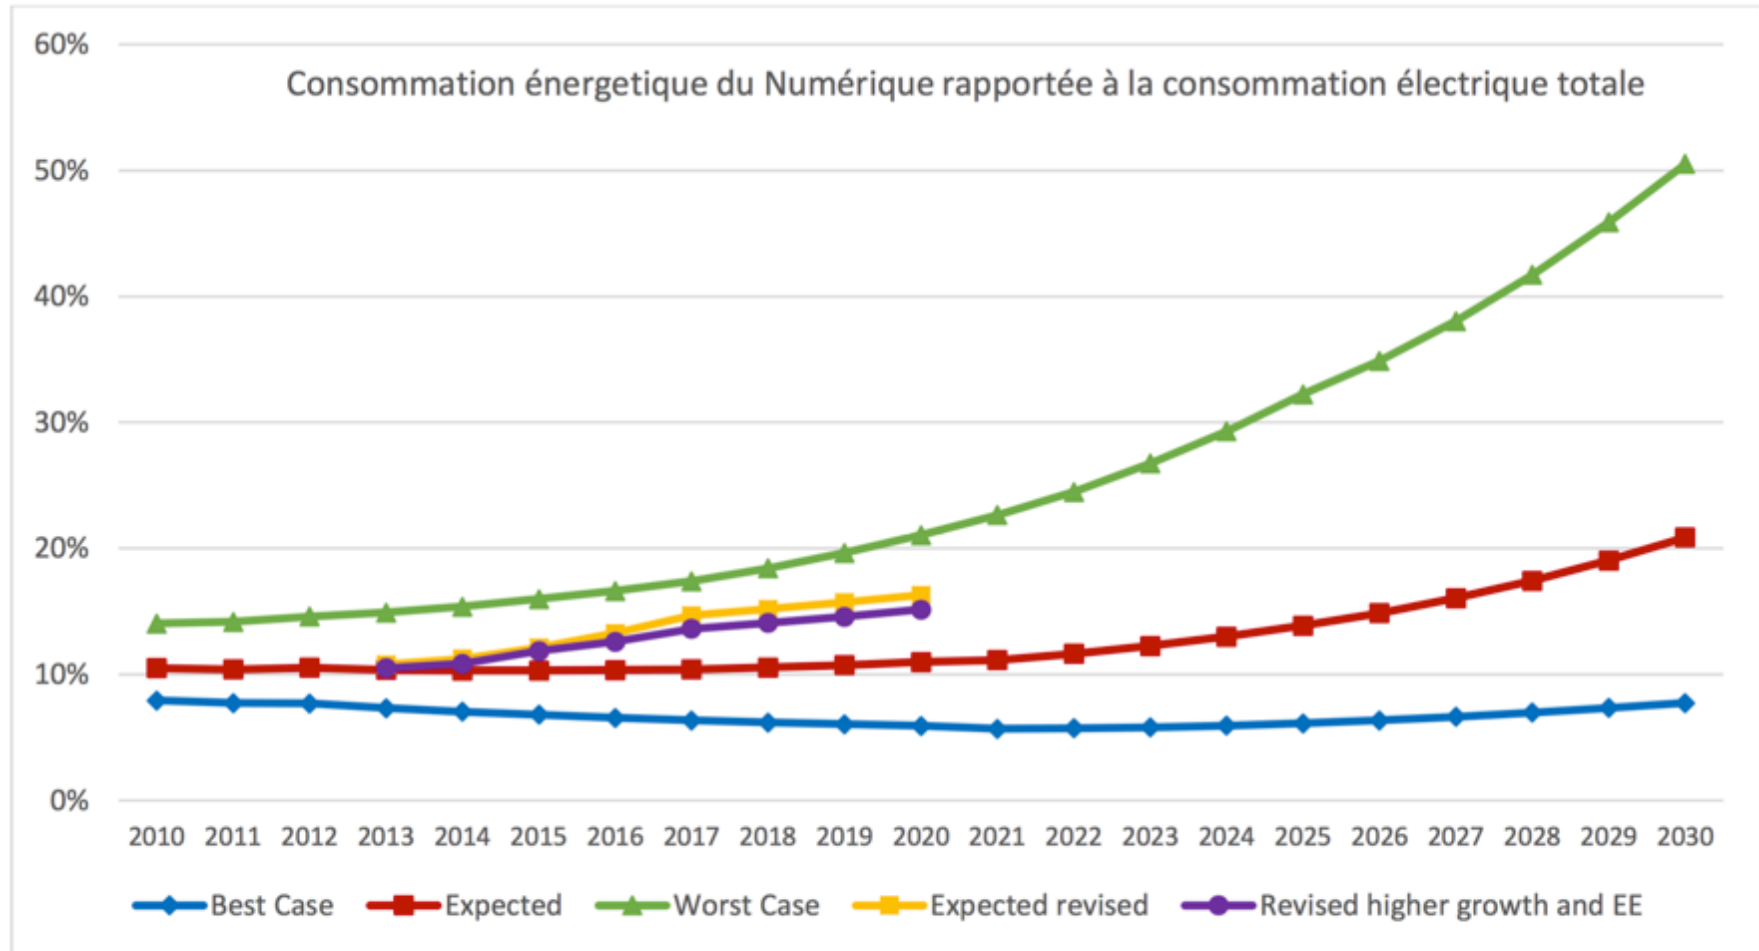
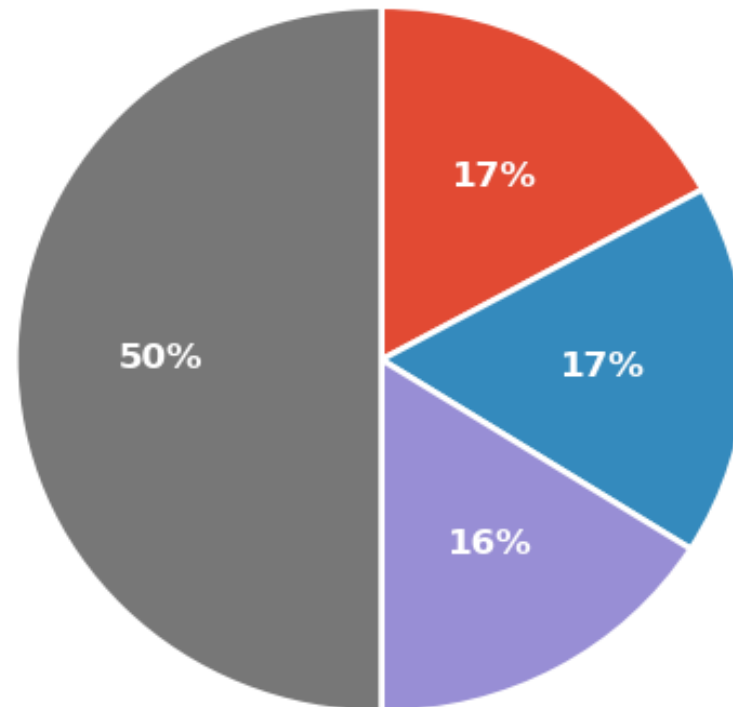


Figure 2 : Évolution 2010-2020 de la consommation énergétique du Numérique rapportée à la consommation électrique mondiale⁷
[Source: calculé par The Shift Project à partir des données publiées par Andrae et Edler (2015)]

14% of the global electricity consumption in 2017.

Distribution of ICT energy consumption

End devices Networks Data centers Production



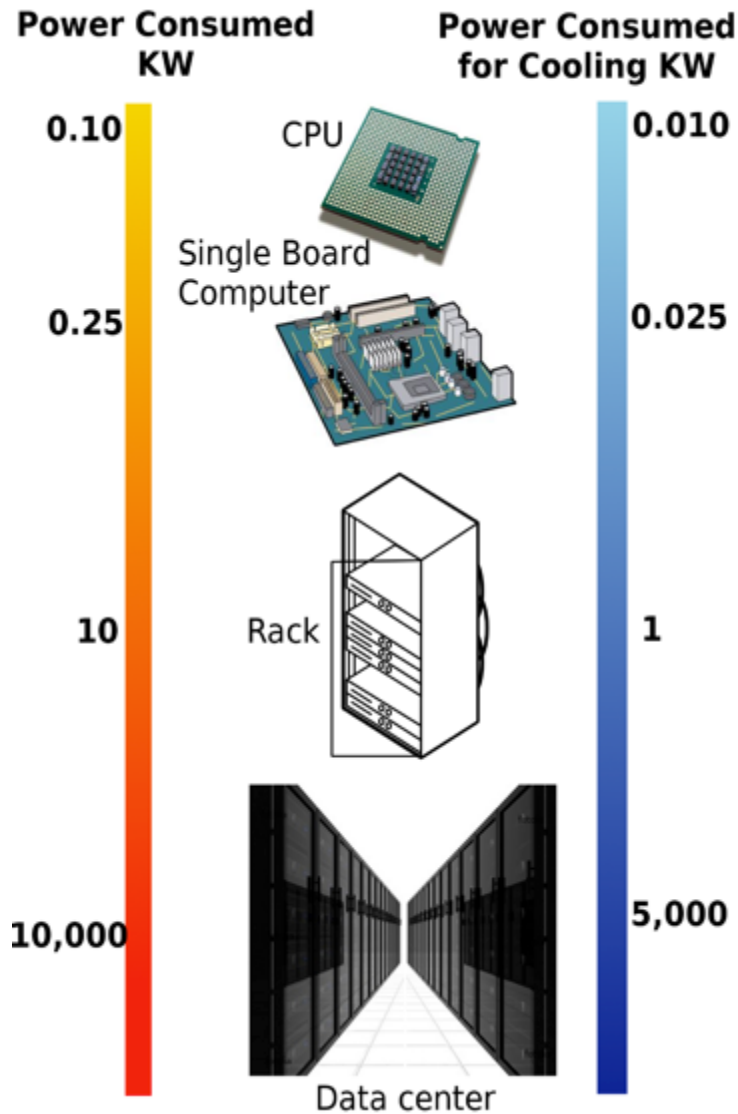
**THE SHIFT
PROJECT**
THE CARBON TRANSITION THINK TANK

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

Anne-Cécile Orgerie



Saving energy



Low power processors (Big.LITTLE)

Multi-core architectures

Energy-efficient dedicated architectures (FPGA, GPU)

Dynamic Voltage Frequency Scaling

Workload consolidation techniques

On/off policies

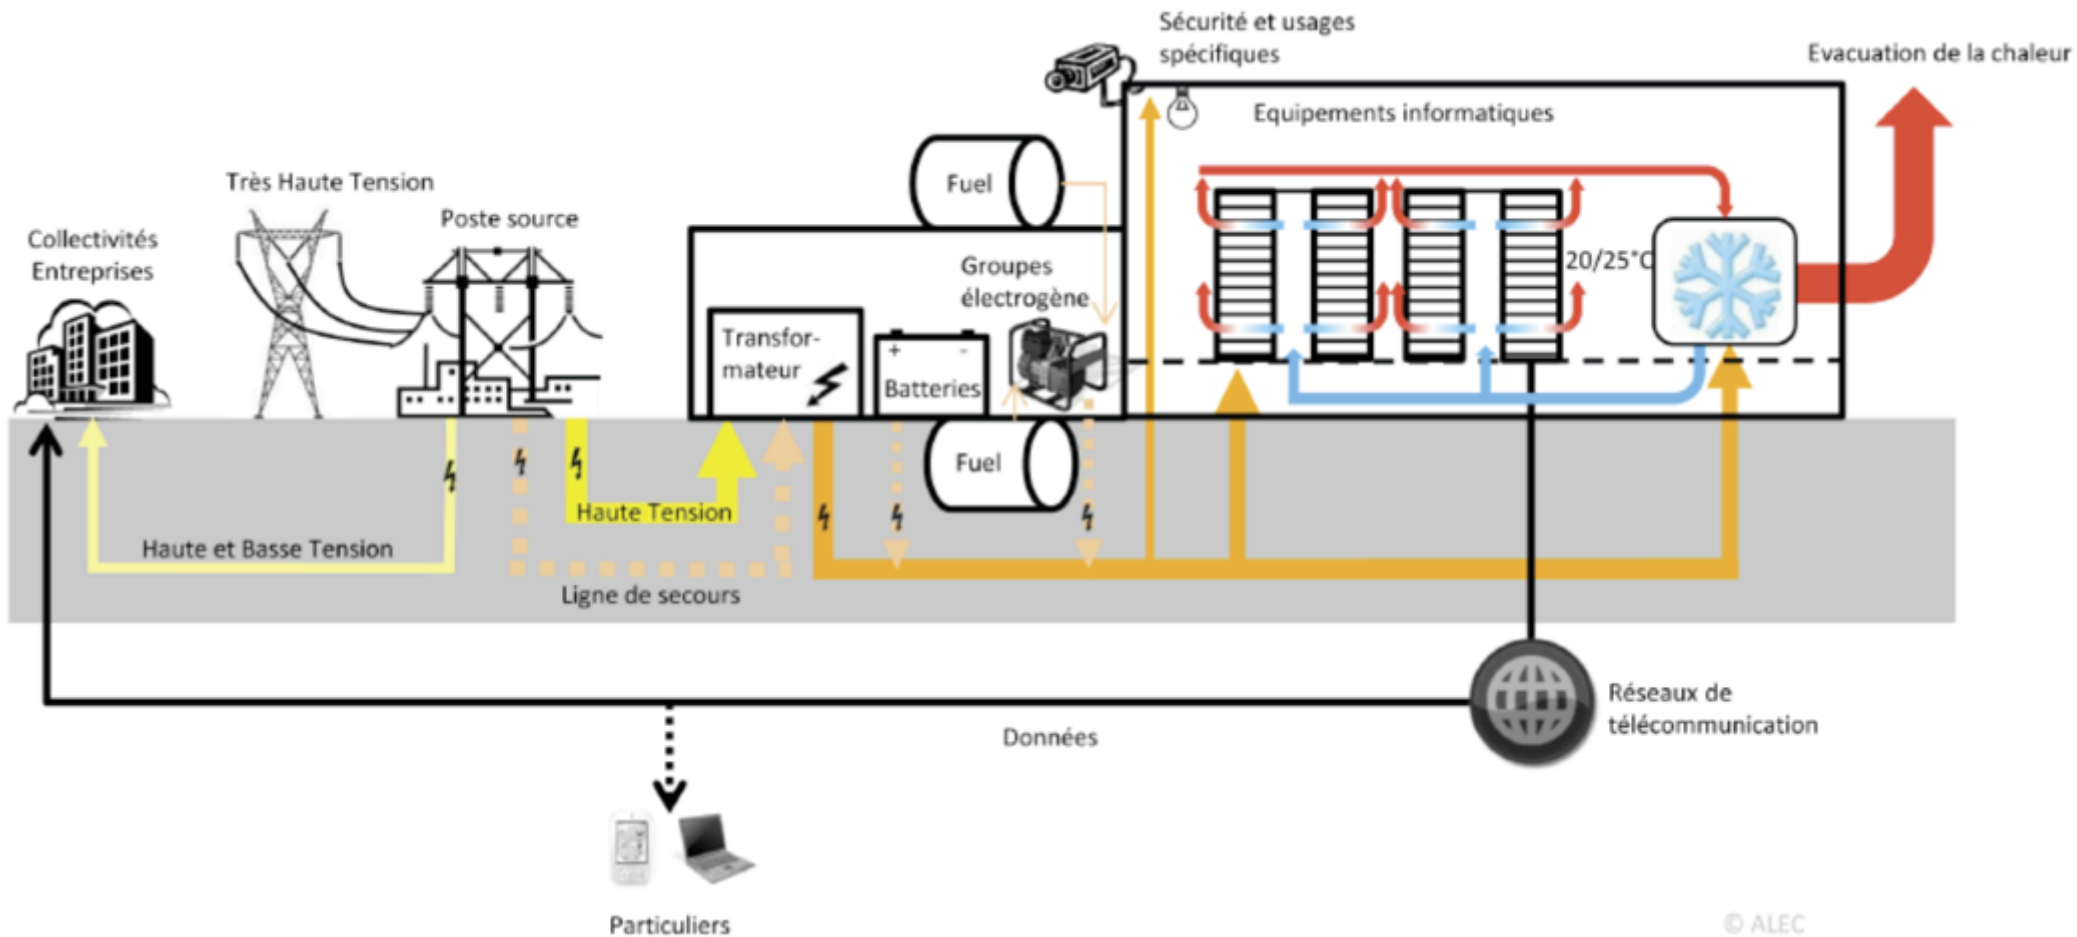
Hot spot management

Workload peak reduction

Dynamic adaptation



Overall data center view



Courtesy of Jean-Marc Menaud

How to measure energy efficiency?

- Benchmarks and metrics
- PUE: Power usage effectiveness

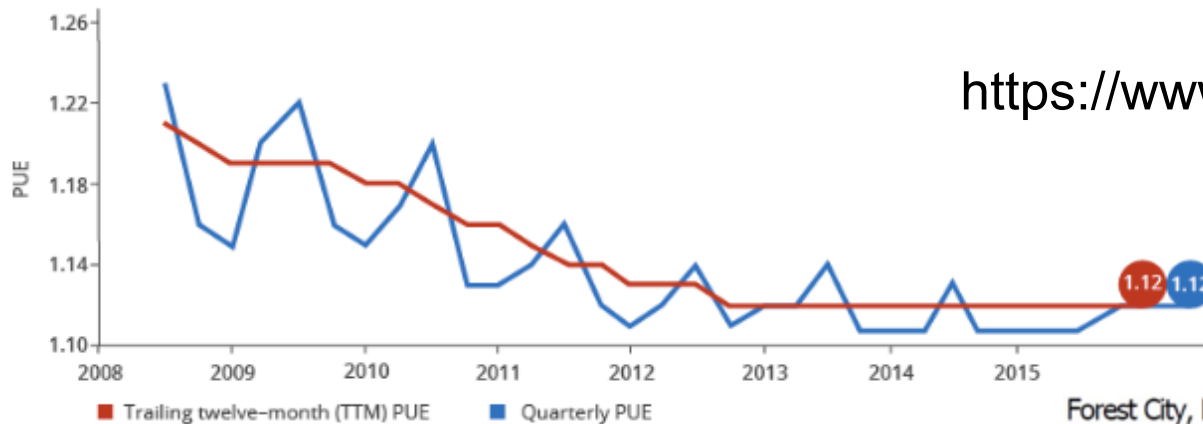
$$PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

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



PUE as a selling point

Continuous PUE Improvement
Average PUE for all data centers

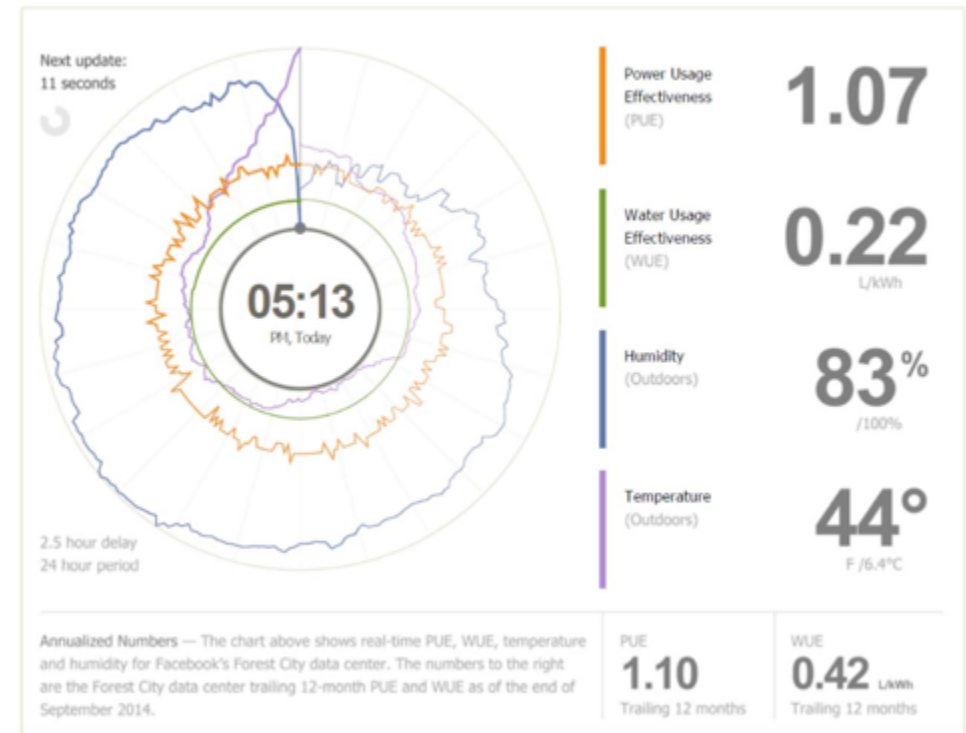


<https://www.google.fr/about/datacenters/>

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

Forest City, NC Data Center

Dashboard: PUE & WUE



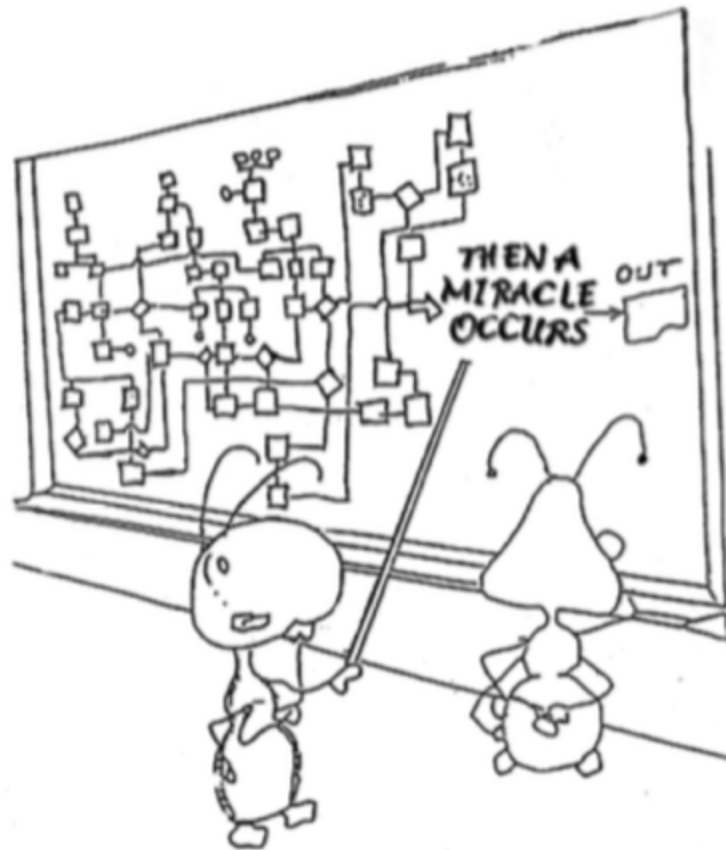
Code of Conduct



- From European Commission
- First release in 2008
- Guidelines, recommendations and best practices (including ASHRAE rules, EnergyStar, etc.)
- Improving understanding of energy demand within the data center
- Raising awareness
- Recommending energy efficient best practice and targets.

<https://ec.europa.eu/jrc/en/energy-efficiency>

Model



Energy models to

- Understand how energy is consumed by hardware
- Estimate energy-performance trade-off
- Guide energy-aware resource allocation policies
- Evaluate energy consumption of applications
- Study what-if scenarios



Energy models of

- Data centers (air cooling, etc.)
- Servers
- Wired network devices
- Cores
- Virtual machines
- Green levers



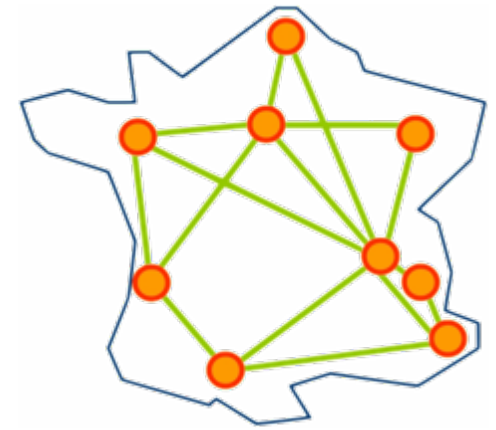
- Multi-criteria
- Multi-scale
- Multi-precision

Models... later!

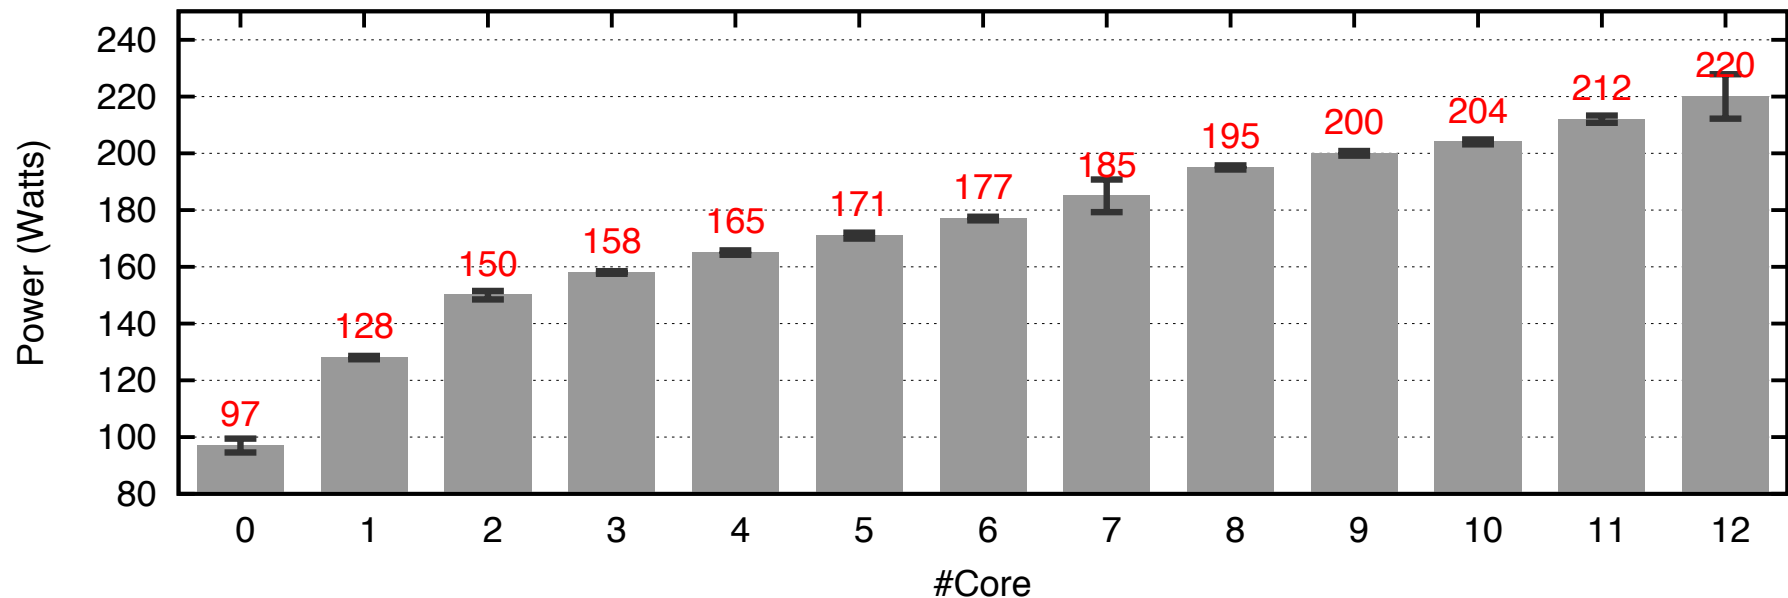
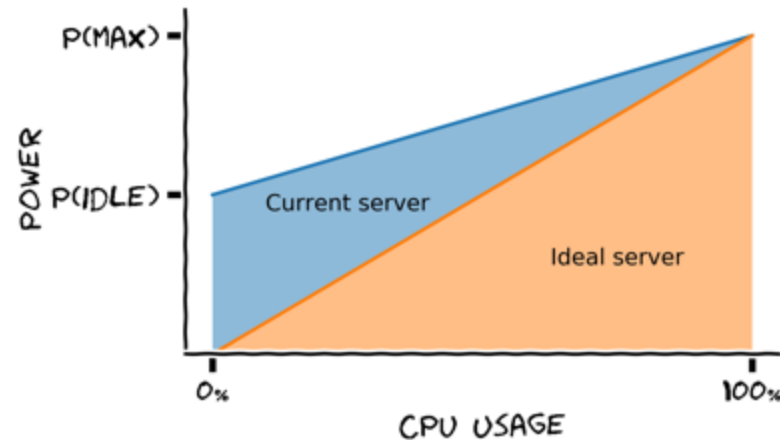
Experiments



Wattmeters

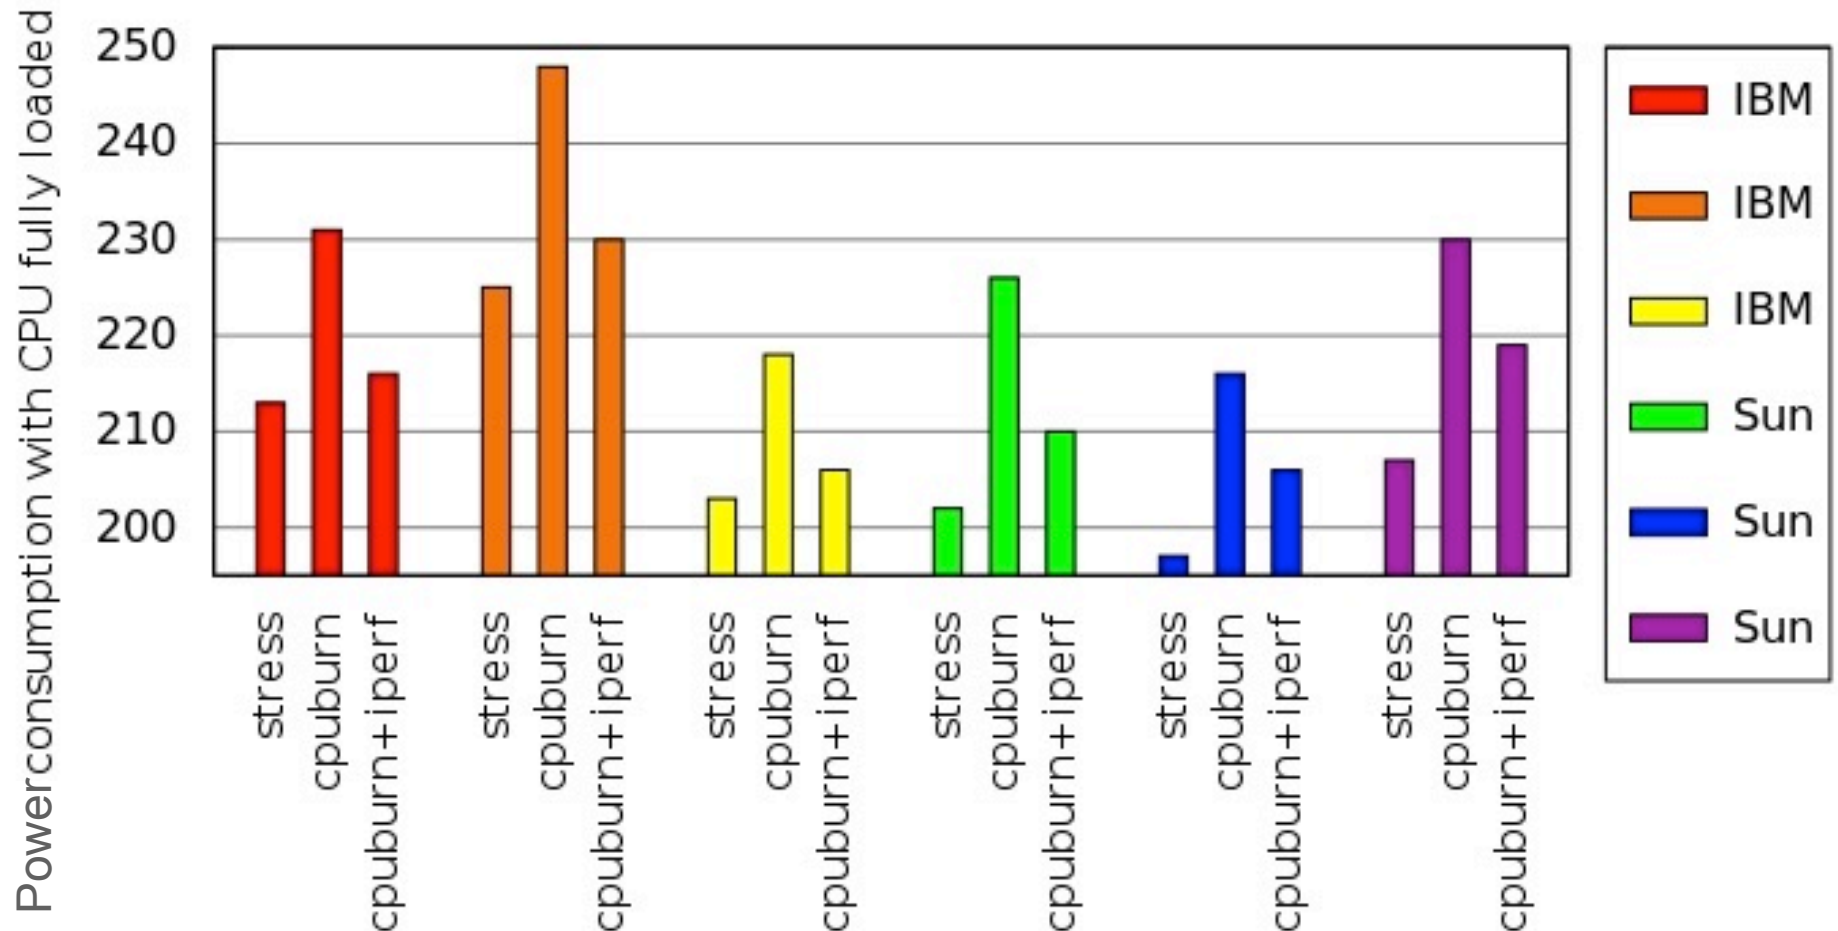


Servers are not power proportional



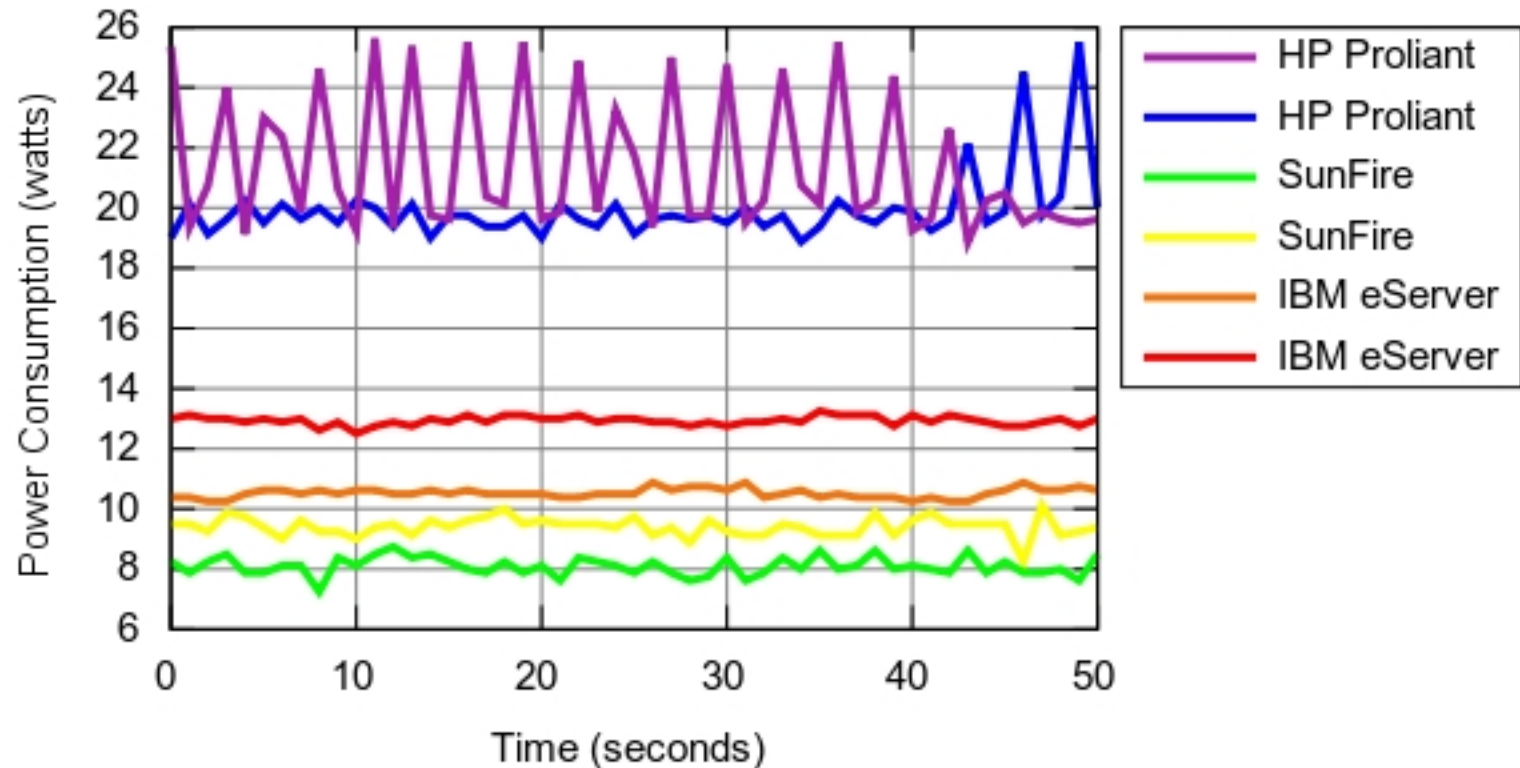
*“Opportunistic Scheduling in Clouds Partially Powered by Green Energy”,
Y. Li, A.-C. Orgerie and J.-M. Menaud, IEEE GreenCom 2015.*

CPU utilization is not the right metric



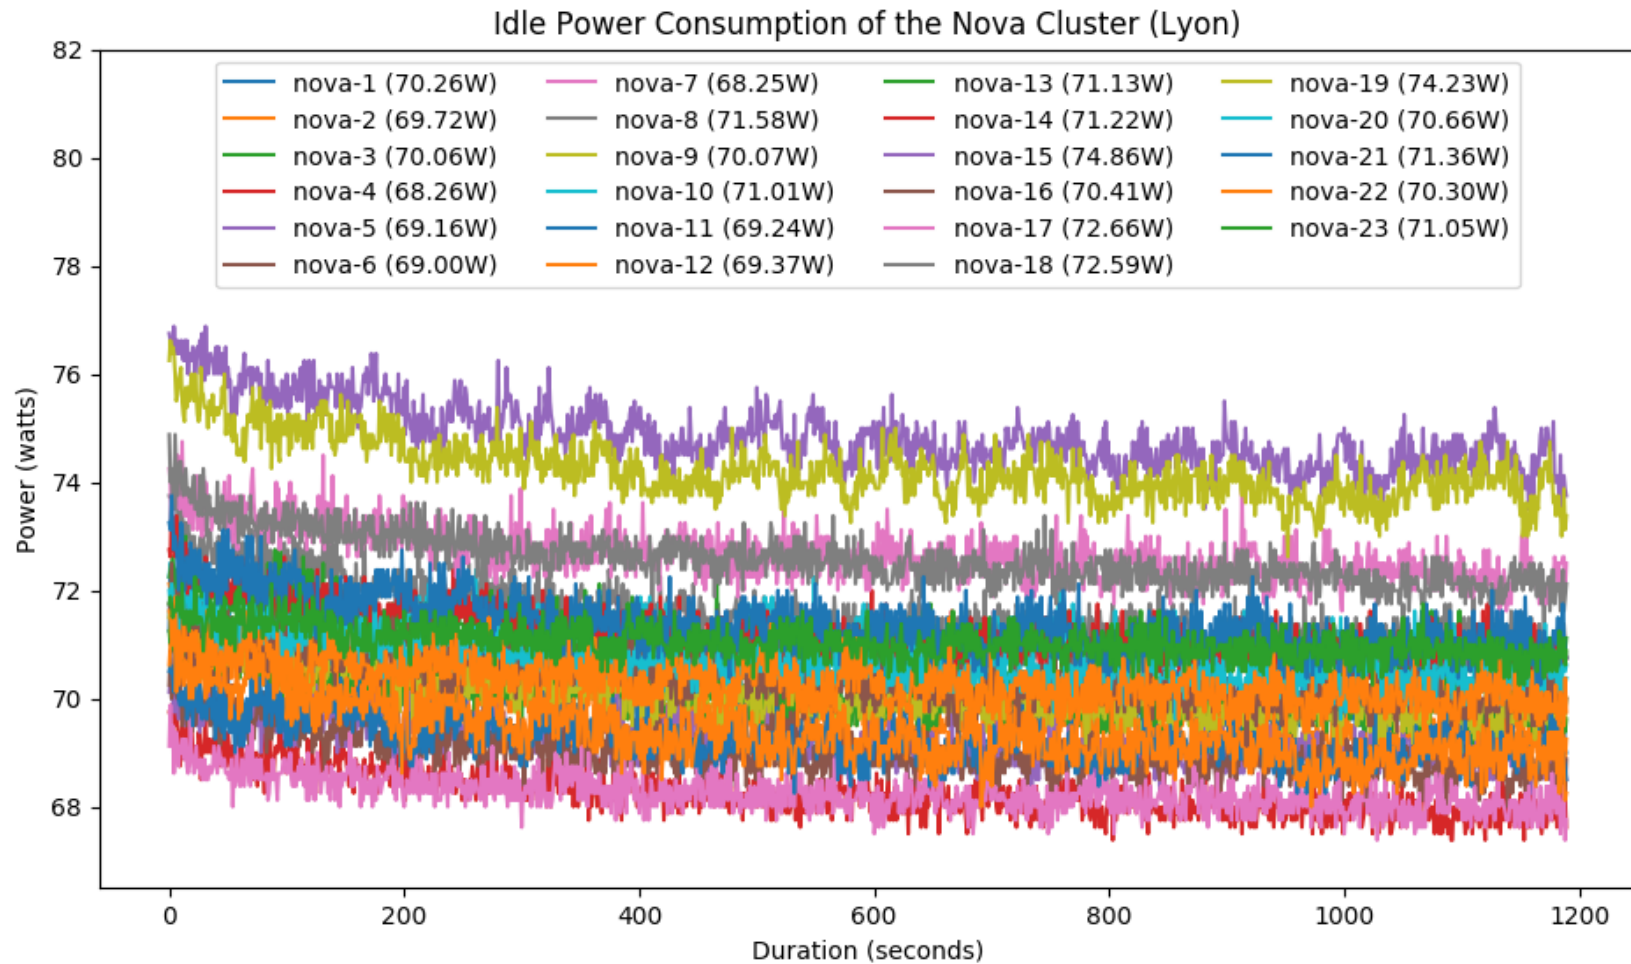
*“Demystifying Energy Consumption in Grids and Clouds”,
A.-C. Orgerie J.-P. Gelas and L. Lefèvre, WIPGC, 2010.*

Switched off servers



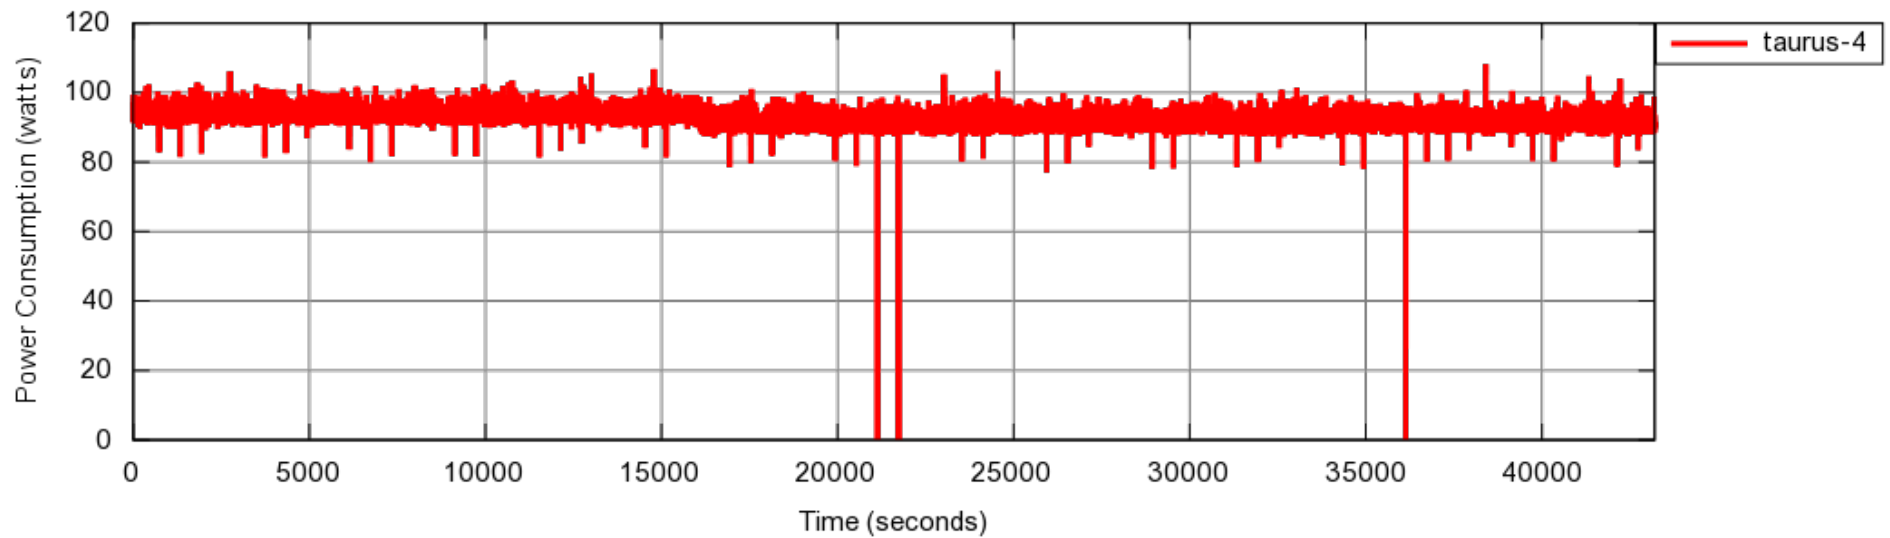
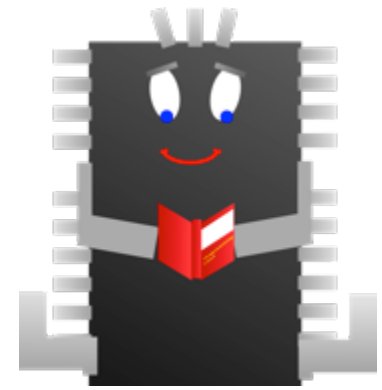
*“Demystifying Energy Consumption in Grids and Clouds”,
A.-C. Orgerie, J.-P. Gelas and L. Lefèvre, WIPGC, 2010.*

Homogeneous servers

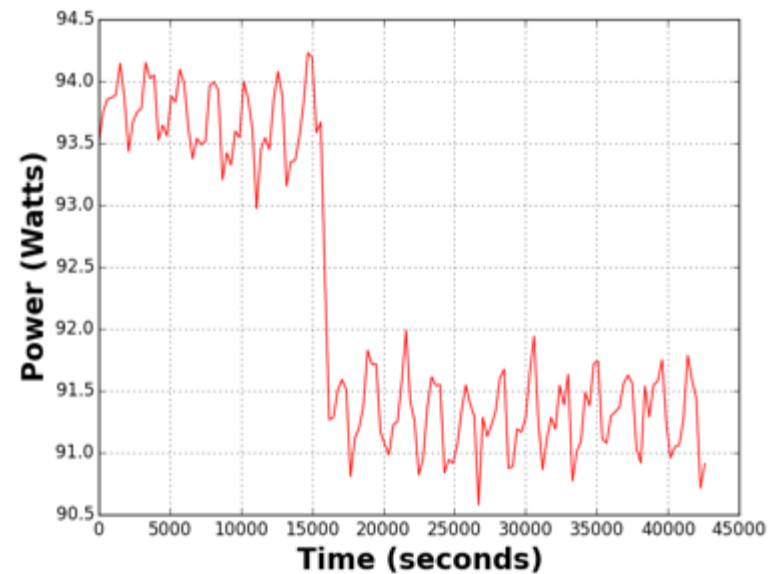
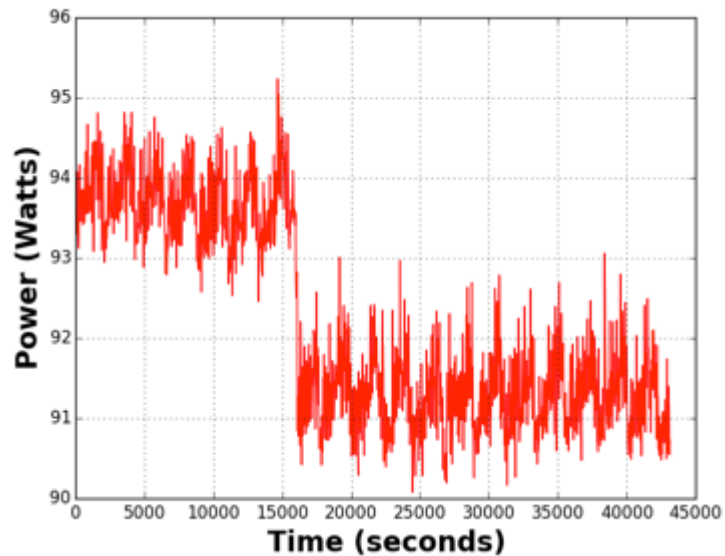
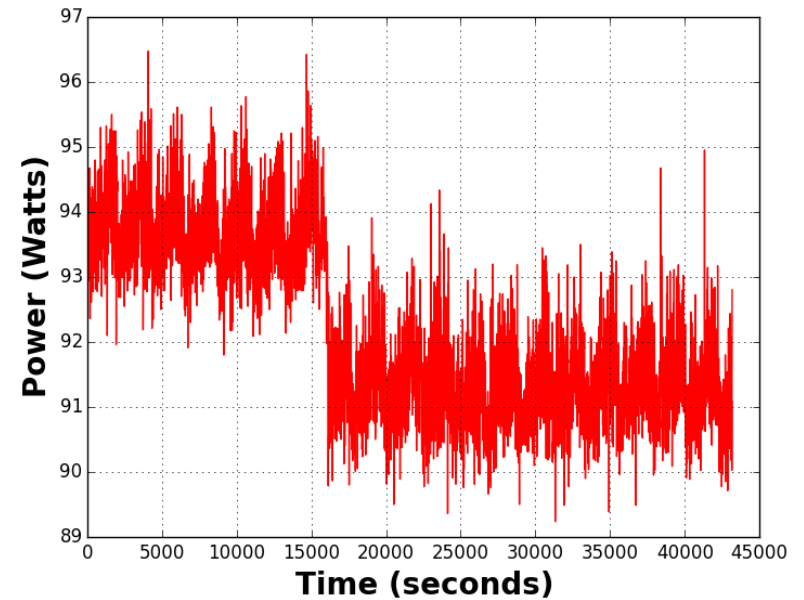
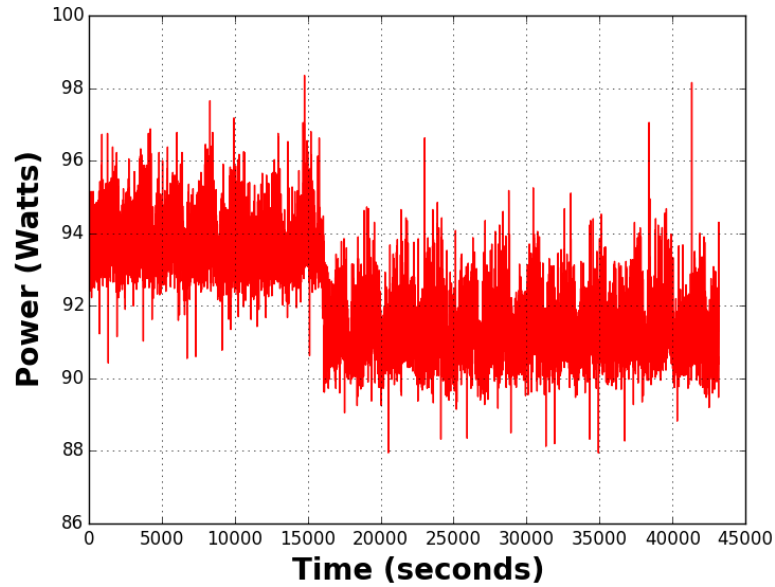


Courtesy of David Guyon

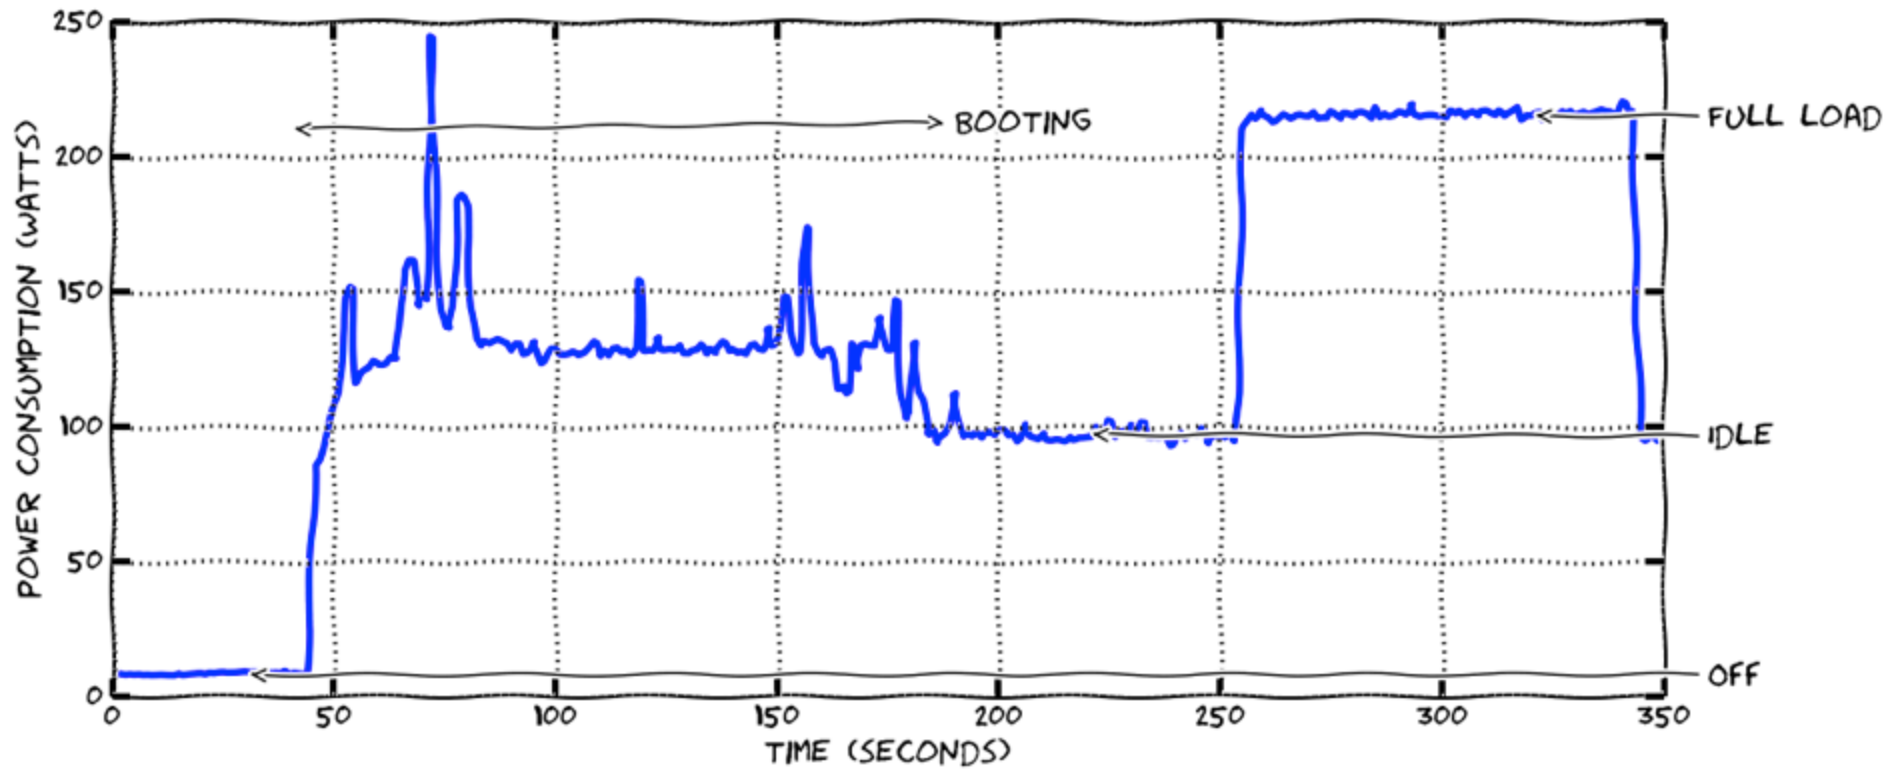
Idle consumption



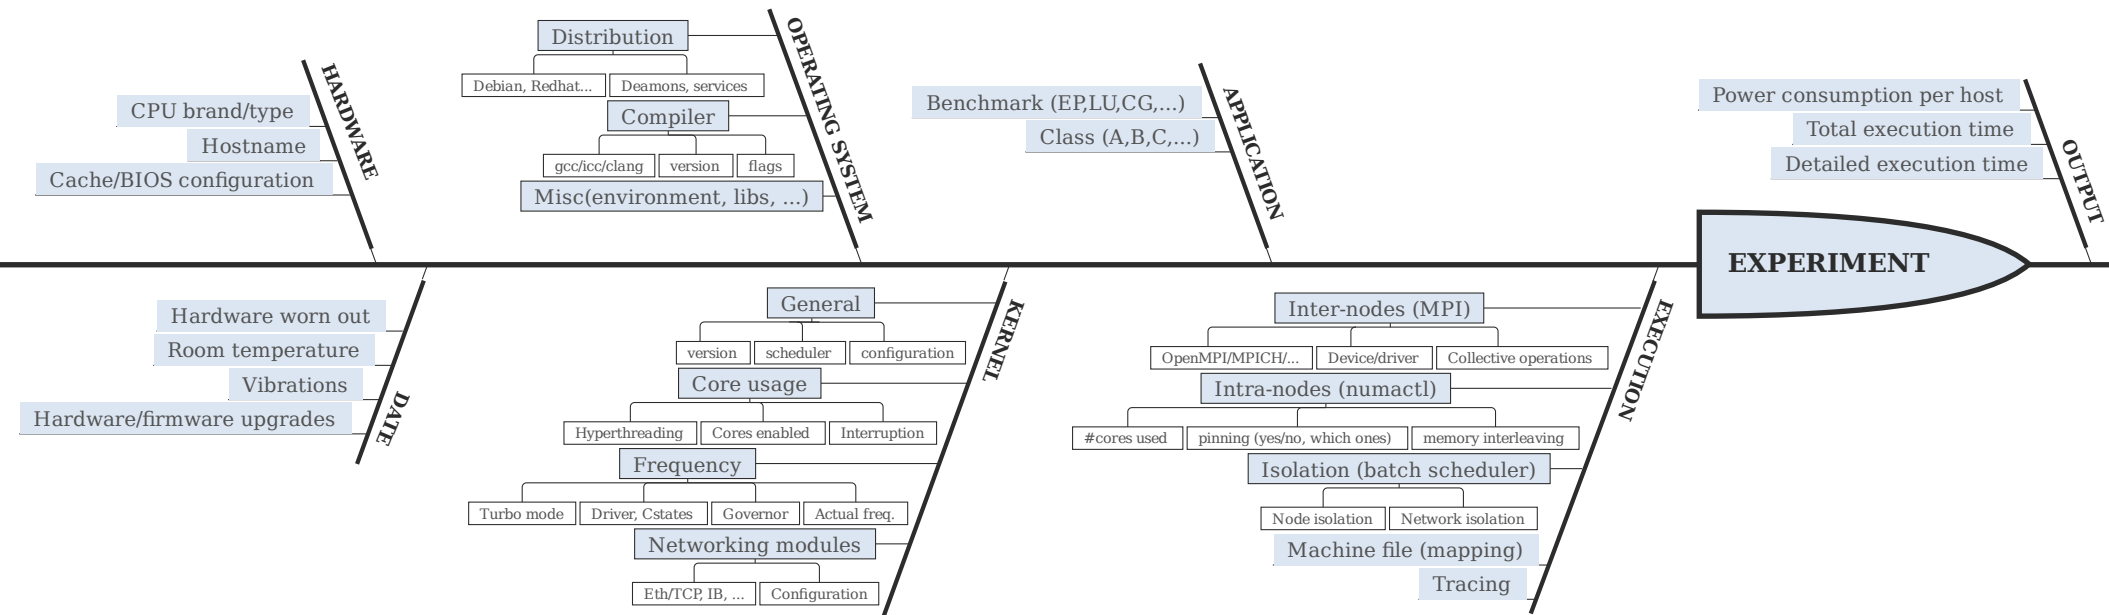
Smoothed idle consumption



Off-On Sequence



Measurement methodology



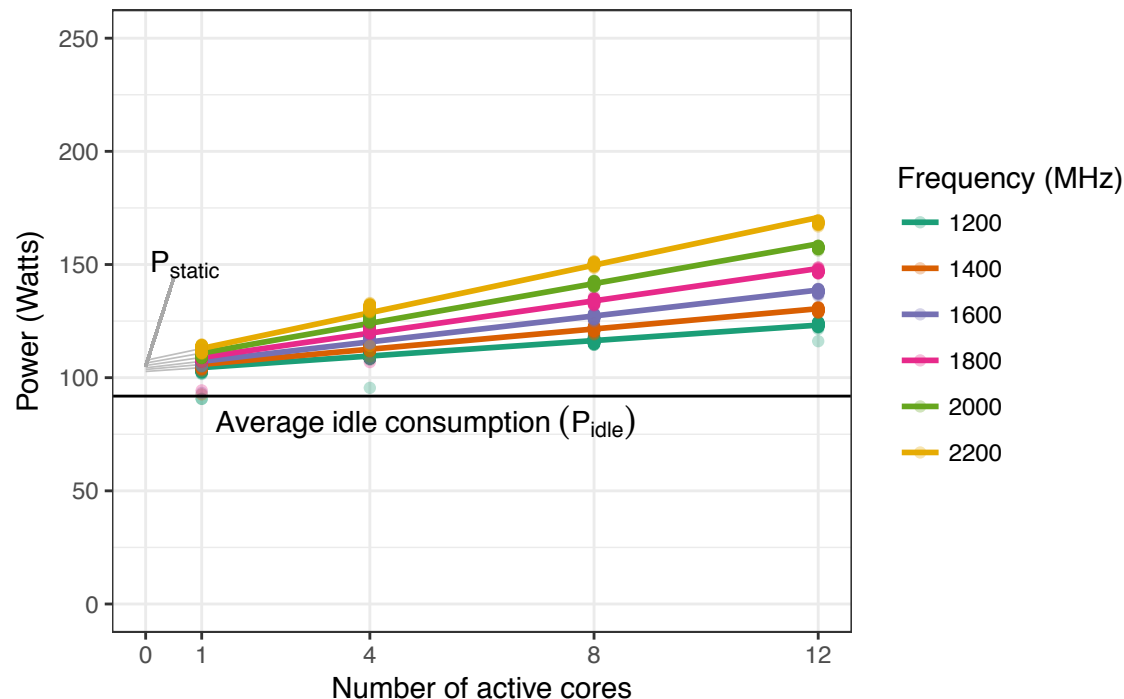
"Predicting the Performance and the Power Consumption of MPI Applications With SimGrid", C. Heinrich, A. Carpen-Amarie, A. Degomme, S. Hunold, A. Legrand, A.-C. Orgerie and M. Quinson, Research Report, 2017.

Validating the models



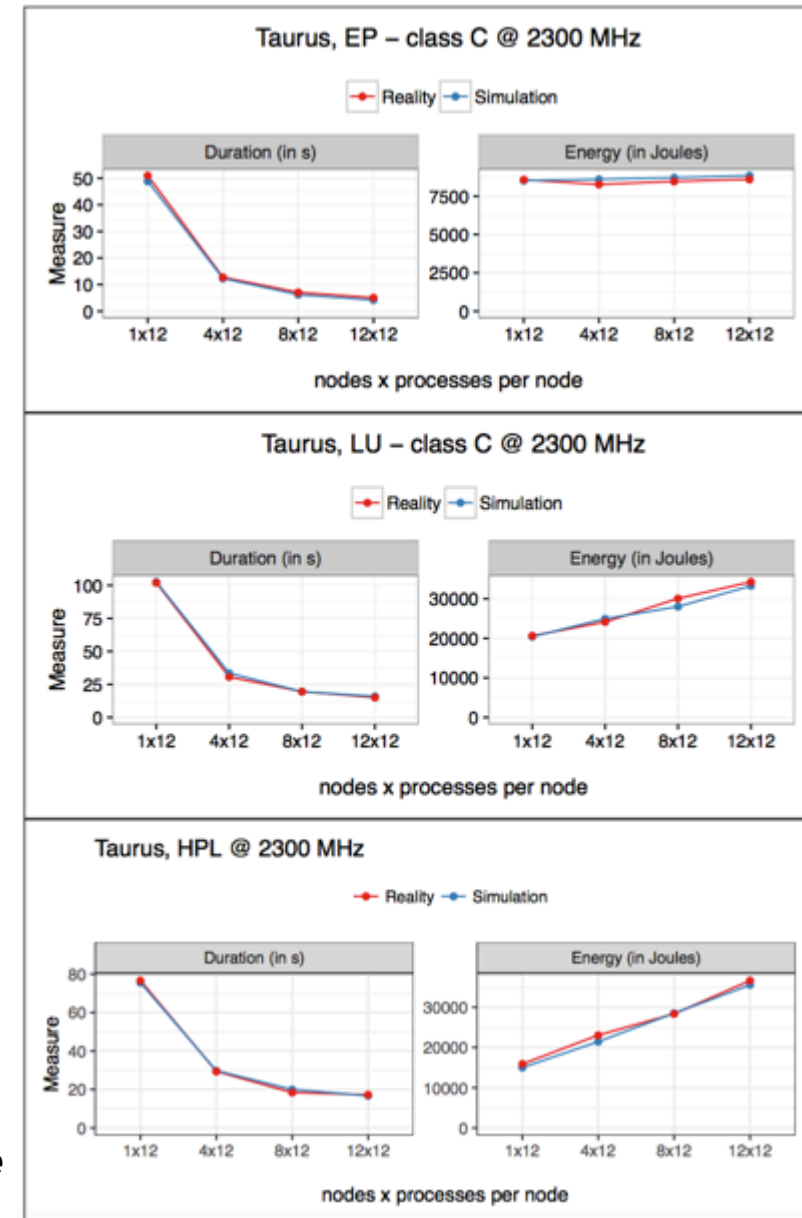
- multicore
- DVFS
- On/off

Taurus cluster, Lyon, NAS-EP

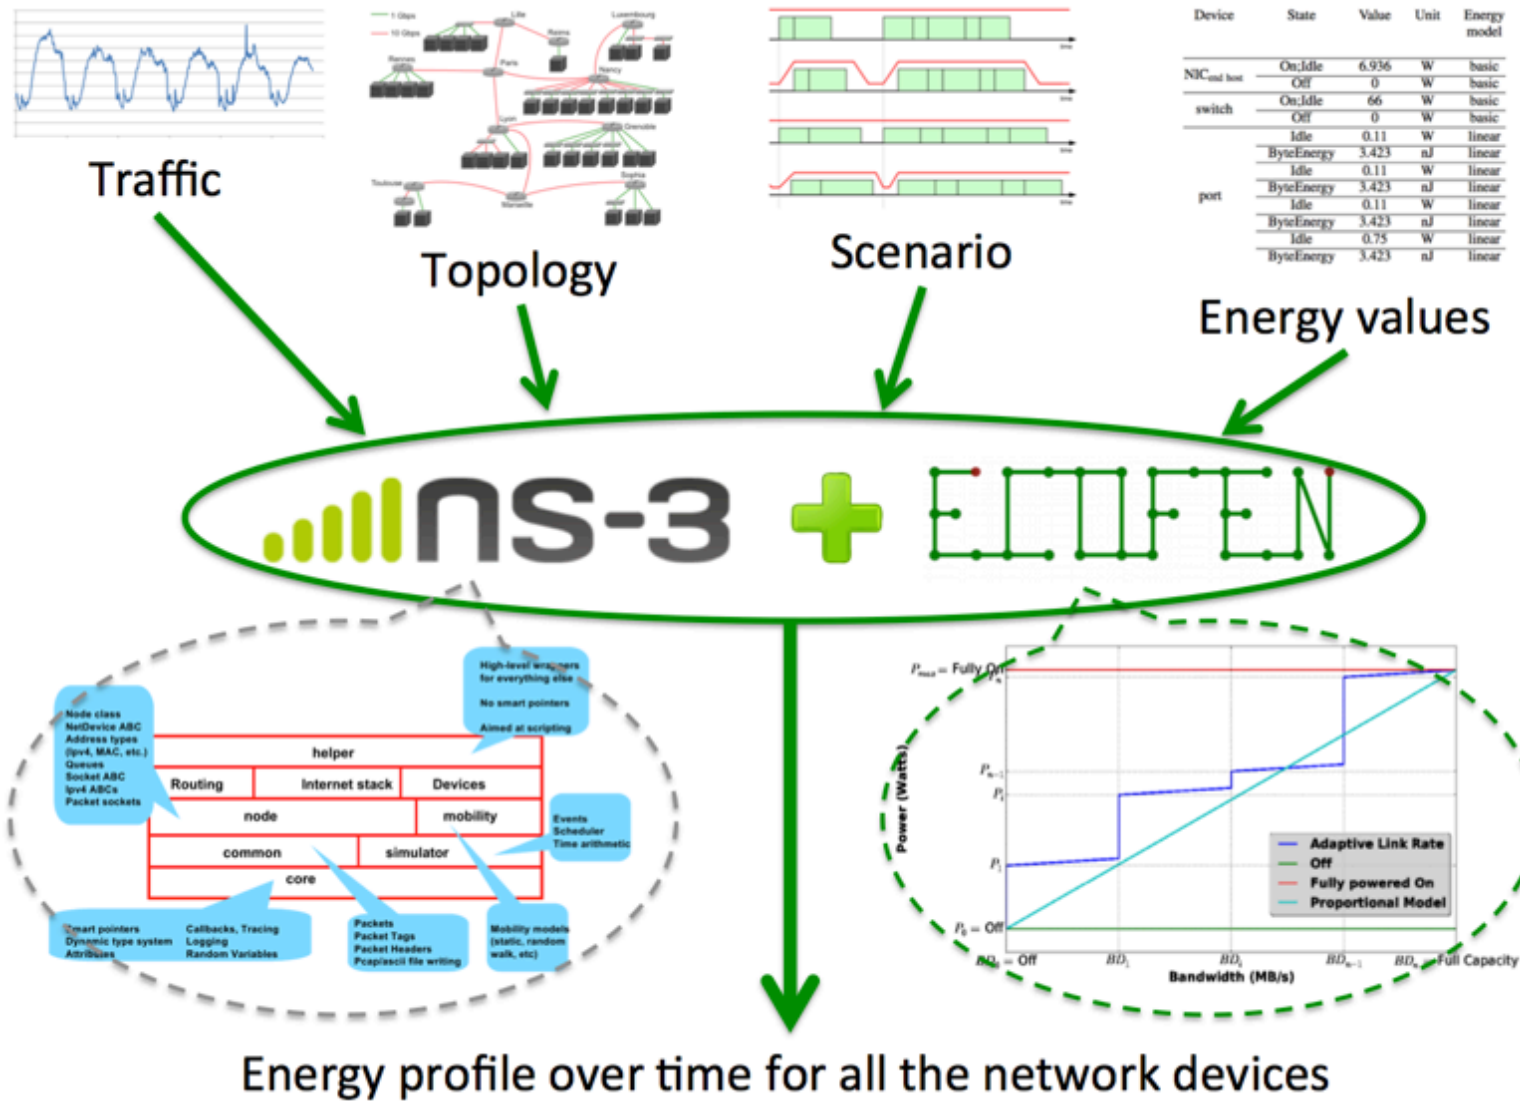


“Predicting the Energy-Consumption of MPI Applications at Scale Using Only a Single Node”, C. Heinrich, T. Cornebize, A. Degomme, A. Legrand, A. Carpen-Amarie, S. Hunold, A.-C. Orgerie, and M. Quinson, IEEE Cluster 2017.

Anne-Cécile Orgerie



Wired networks



lip

i3S
sophia antipolis

"Simulation toolbox for studying energy scenarios in wired networks", A.-C. Orgerie, B. Amersho, T. Haudebourg, M. Quinson, M. Rifai, D. Lopez Pacheco, and L. Lefèvre, CNSM 2018.

Anne-Cécile Orgerie

What is difficult

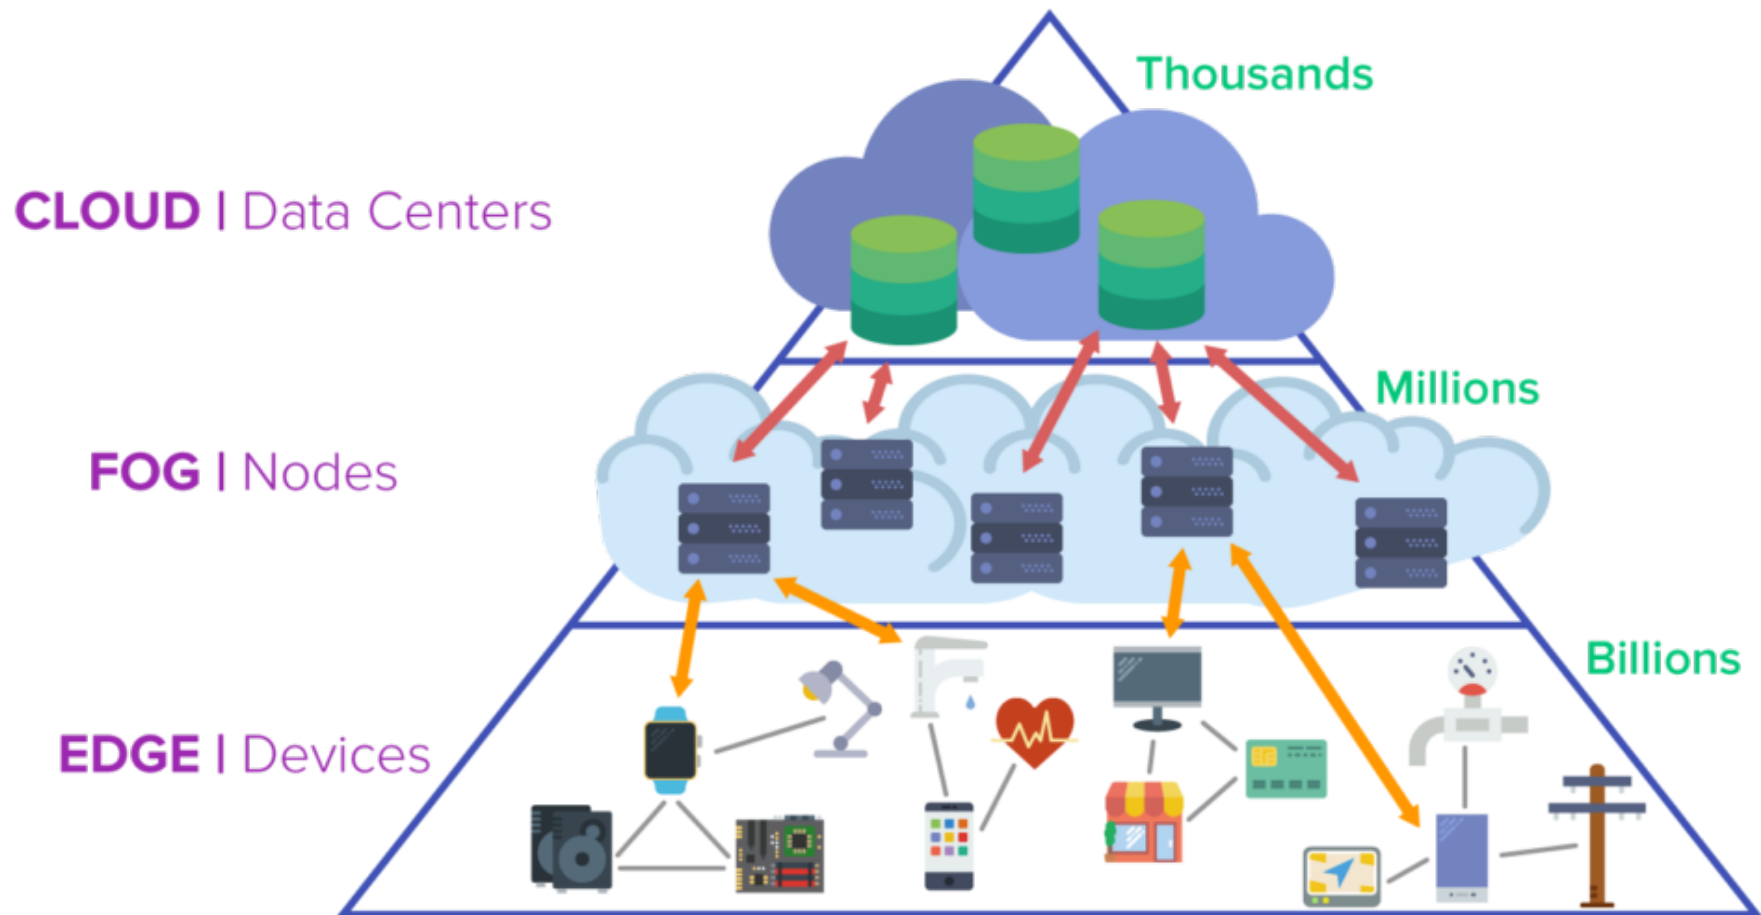


- Instrument realistic infrastructures
- Measure accurately consumption of resources
- Isolate influential factors
- Combine energy models with performance models
- Propose models integrating inherent variability
- Perform campaign measurements
- Publish: either new models or invalidation studies

Scenario

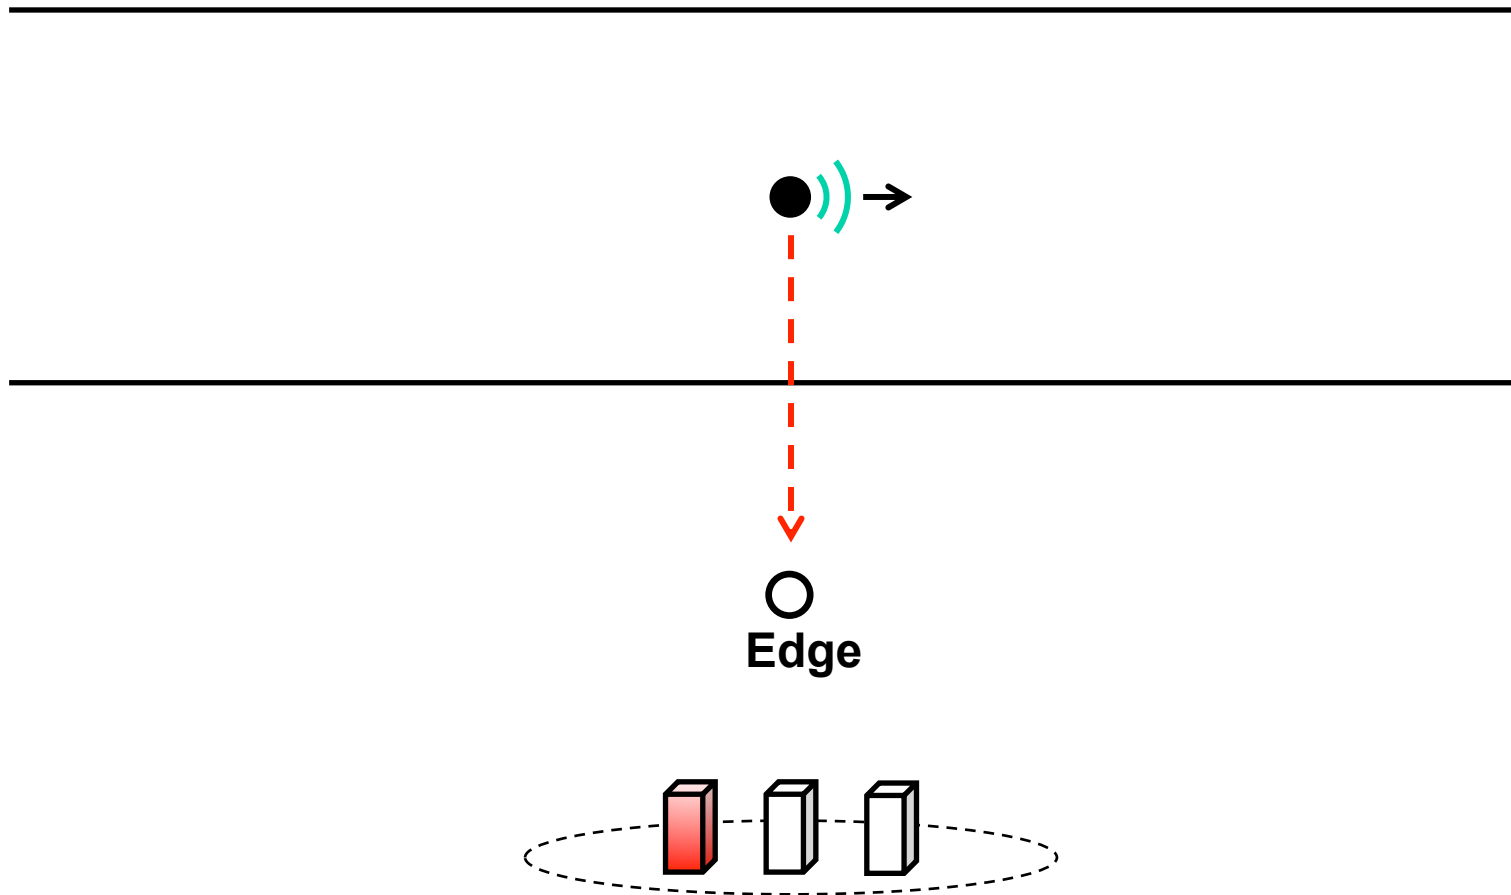


New cloud architectures



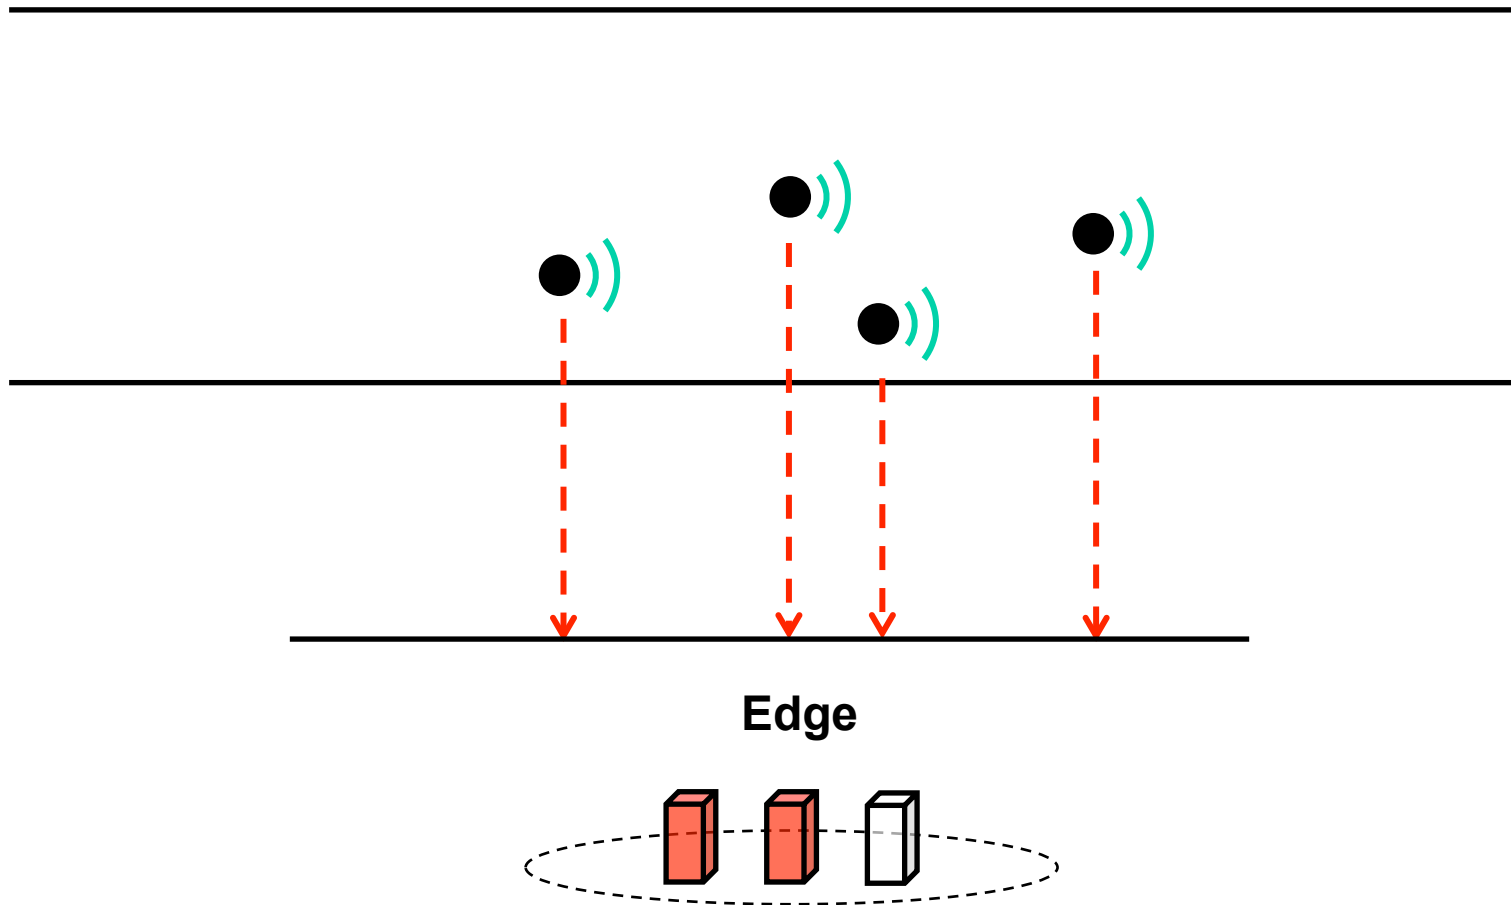
<https://erpinnews.com/fog-computing-vs-edge-computing>

Edge Model



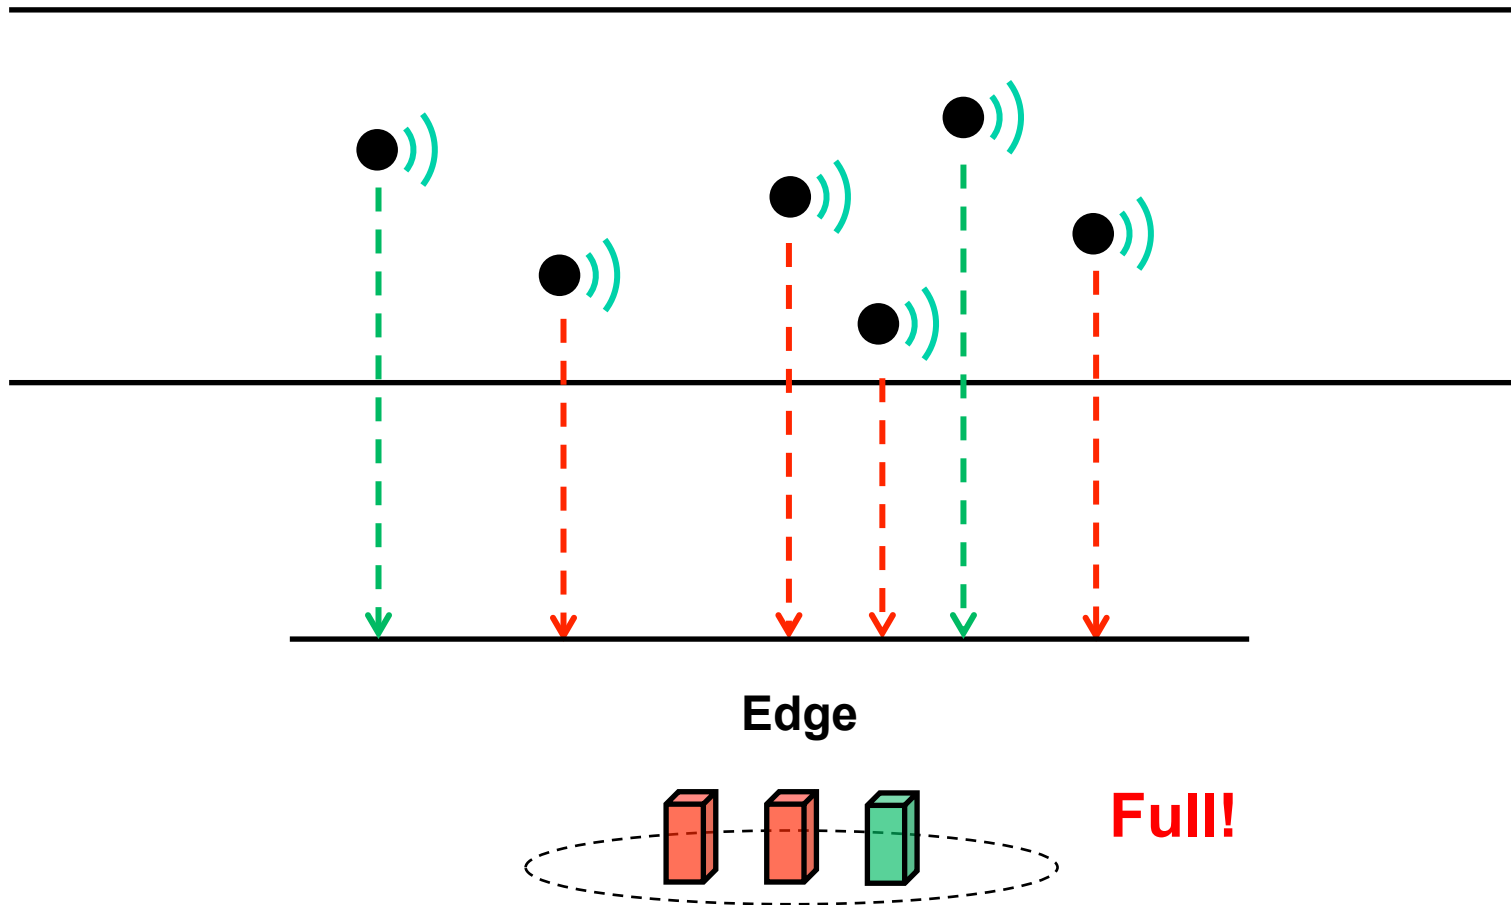
Courtesy of Yunbo Li

Edge Model



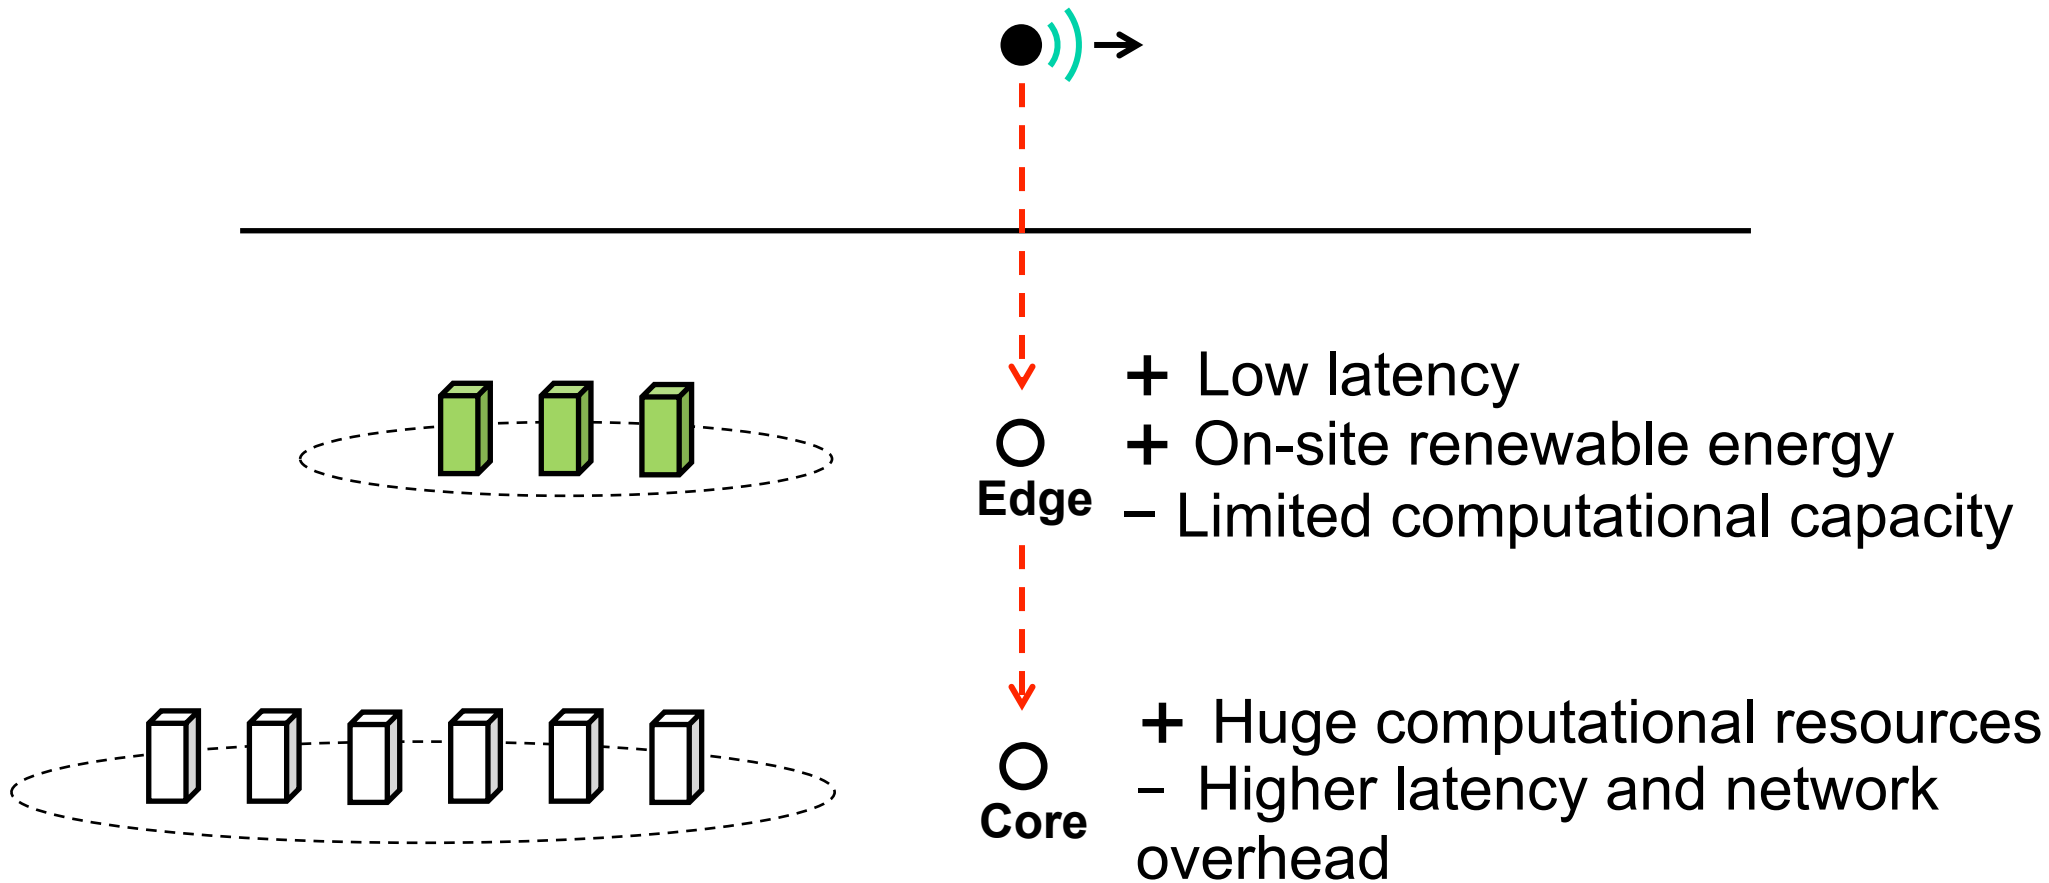
Courtesy of Yunbo Li

Edge Model



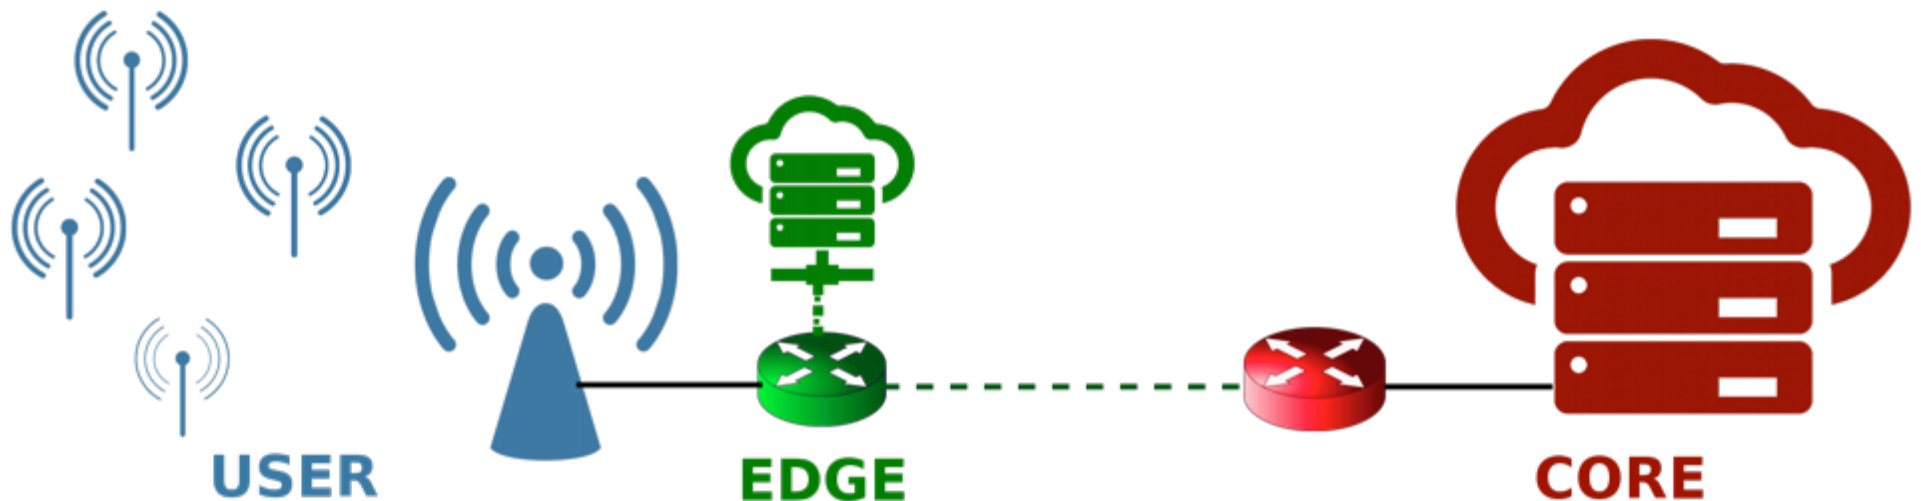
Courtesy of Yunbo Li

Edge-Core Model



Problem

How to decide to compute at the edge or offload at the edge depending on QoS and energy-efficiency for a given IoT application?



"Leveraging Renewable Energy in Edge Clouds for Data Stream Analysis in IoT", Y. Li, A.-C. Orgerie, I. Roderio, M. Parashar and J.-M. Menaud, p 186-195, IEEE/ACM CCGrid 2017.

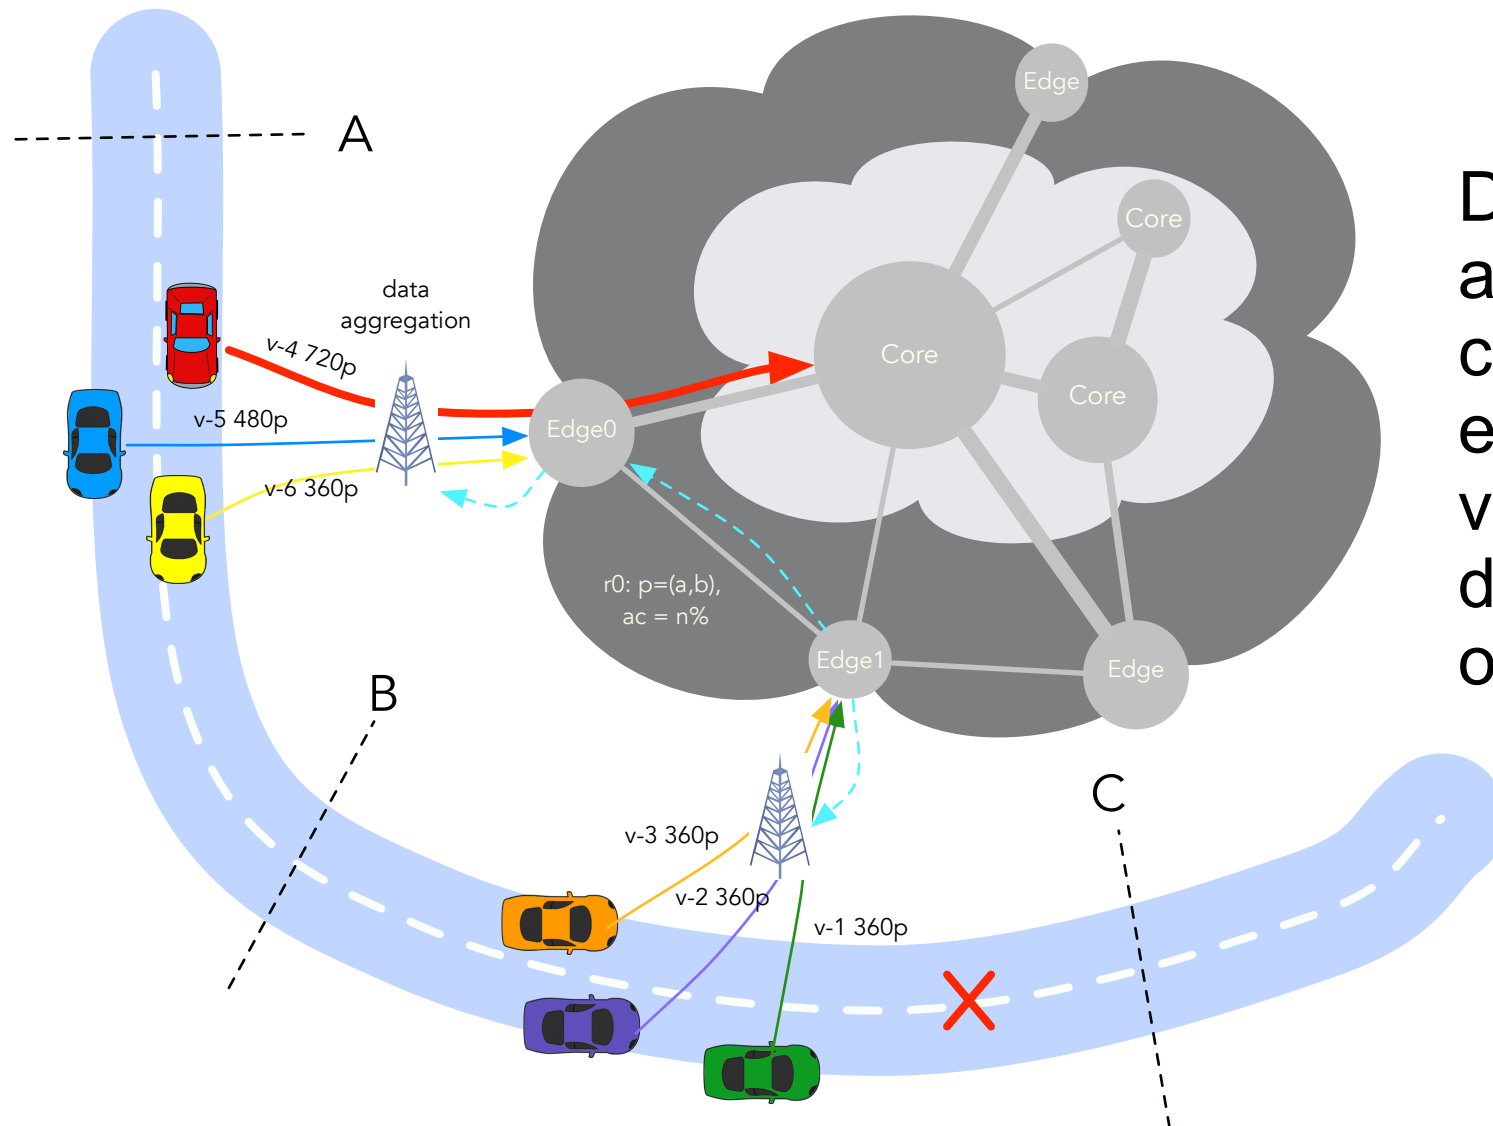
Costs of running on edge/core cloud for a given application

Depends on:

- Application's characteristics (generated traffic)
- Application's required QoS (response time, security, etc.)
- Cloud computing capacities:
 - Resource availability
 - Computing & storage capacities
 - Virtual technology (containers, VM configuration, etc.)
- Network bandwidth
- Renewable energy availability

Performance/energy trade-off

Application-driven approach



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

Evaluation metrics

- Application accuracy (detection probability)
- Service performance (response time)
- Energy consumption
- Green energy consumption



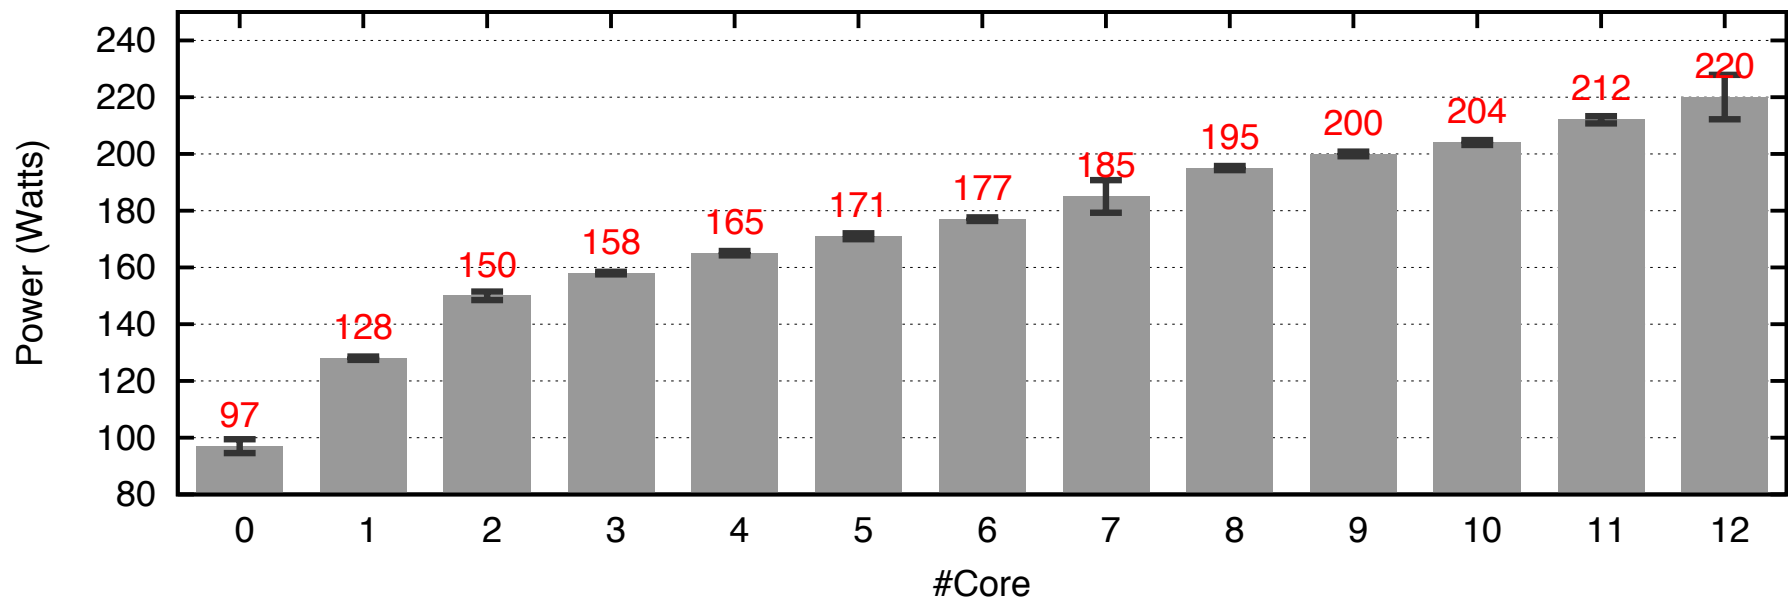
Application details

- Haar classifier (in OpenCV) to analyze video streams for object detection
- Videos encoded in H.264 at 25 fps in 3 resolutions (360p, 480p, 720p)
- Analysis of about 1 frame over 3 (8 fps)
- 5 minutes videos for the experiments

	resolution	bit rate
360p	640 x 360	514 kb/s
480p	720 x 480	706 kb/s
720p	1280 x 720	1176 kb/s

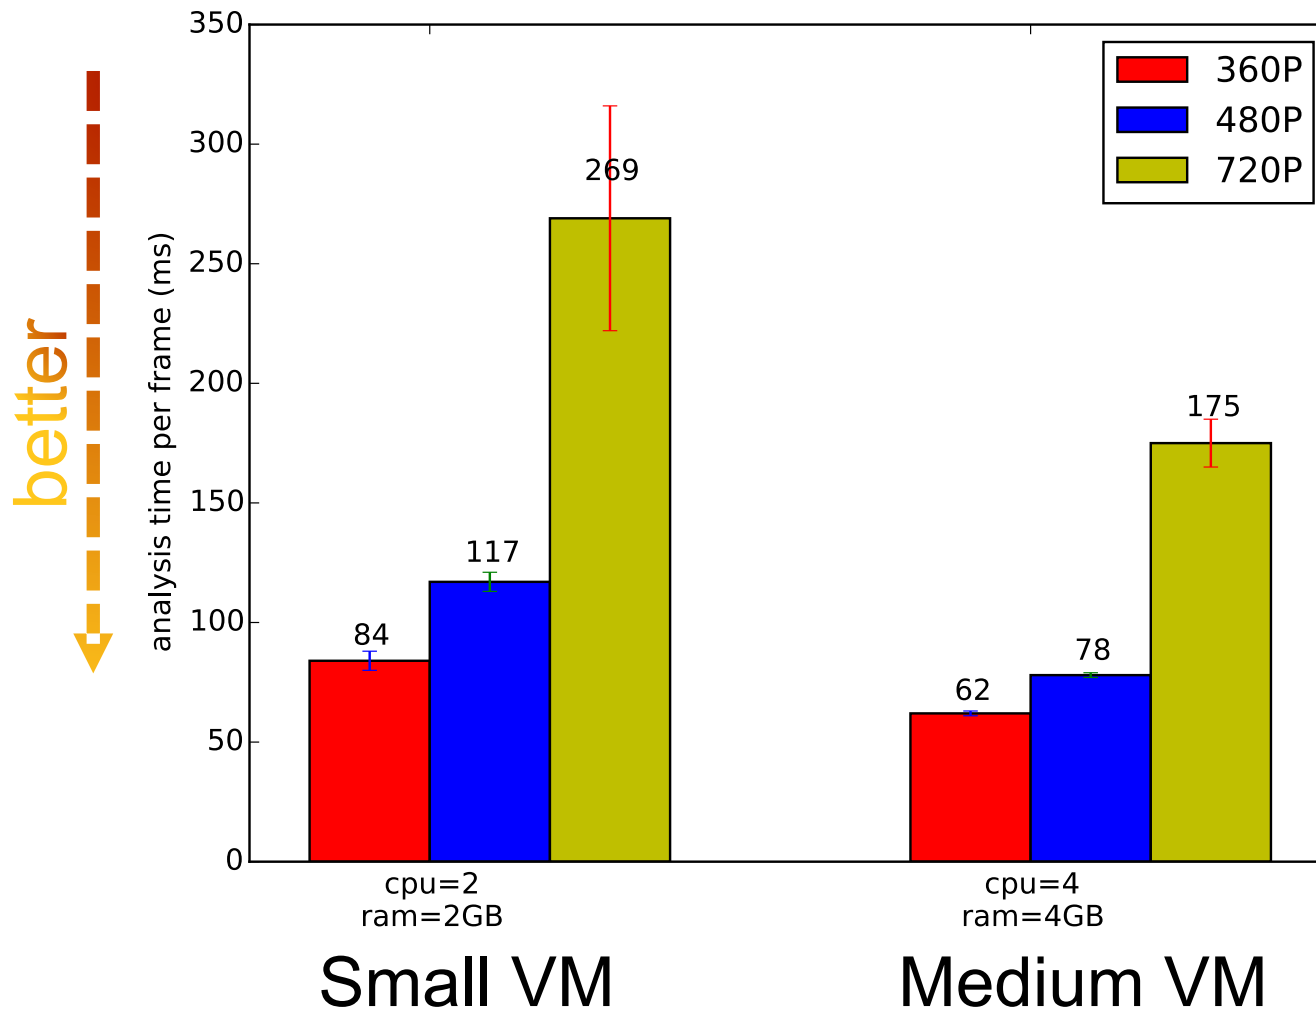
Servers' power profile

- x86 servers with 12 physical cores (2.3 GHz), 32 GB RAM (*Taurus*)
- KVM-based virtualization layer



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

Analysis time on different VM sizes

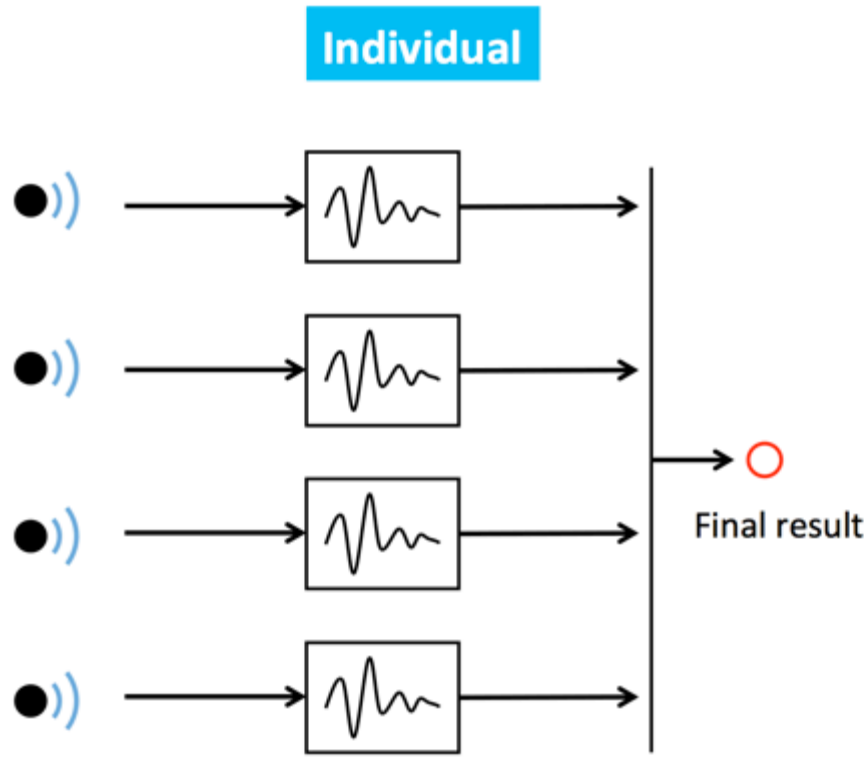


Medium VM
better, but not
linear scalability

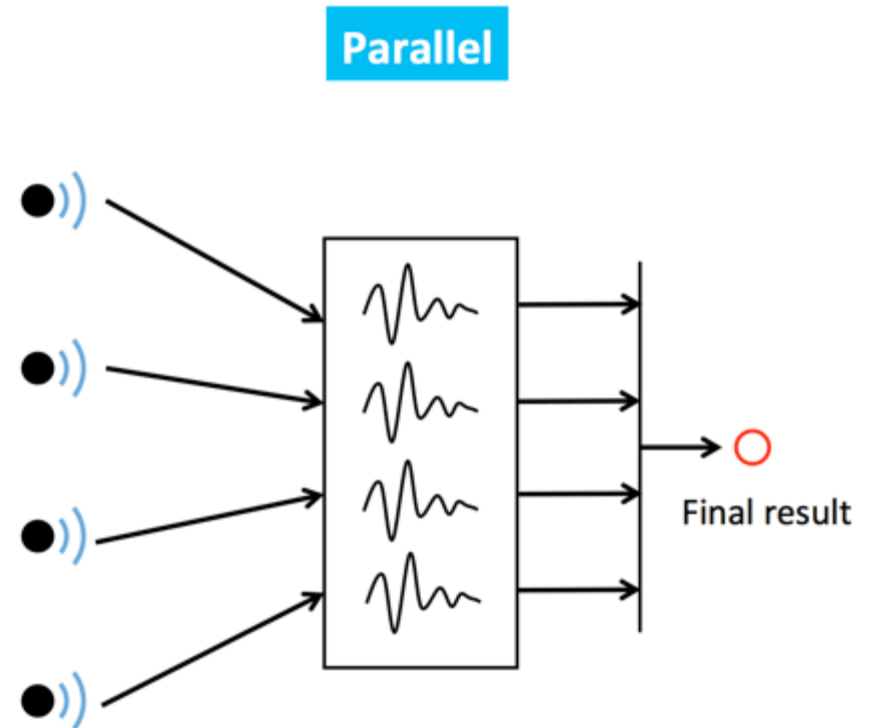
*Depends on
applications'
elasticity*

Real measurements based on 10 runs for each experiment.

Service configuration



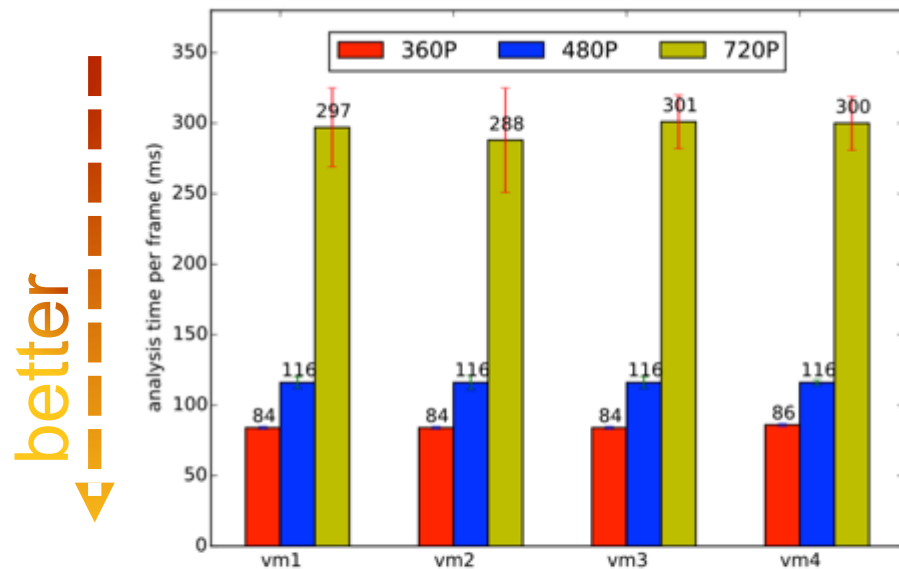
One VM per stream



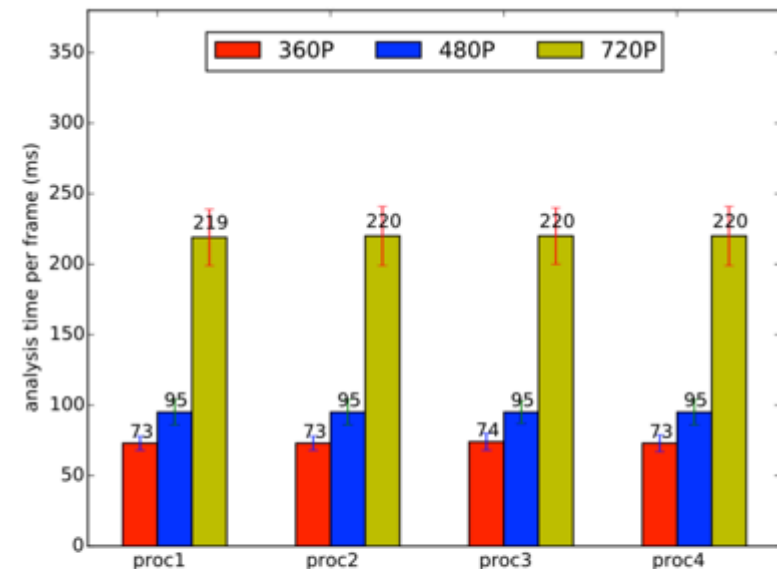
A VM for several streams

Courtesy of Yunbo Li

Several small VMs vs. one large VM



(a) Analysis time for 4 identical VMs with 1 data stream each on the same PM

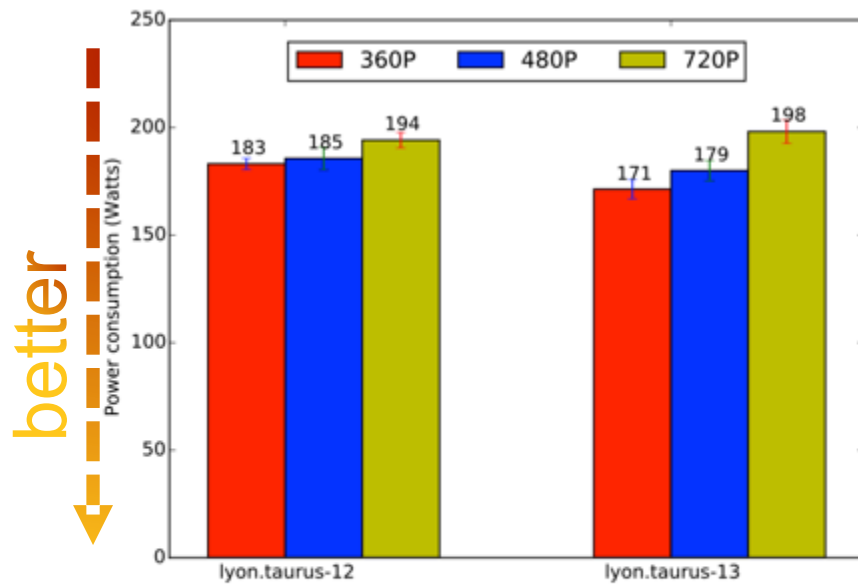


(b) Analysis time for each of the 4 data stream processes in a large VM

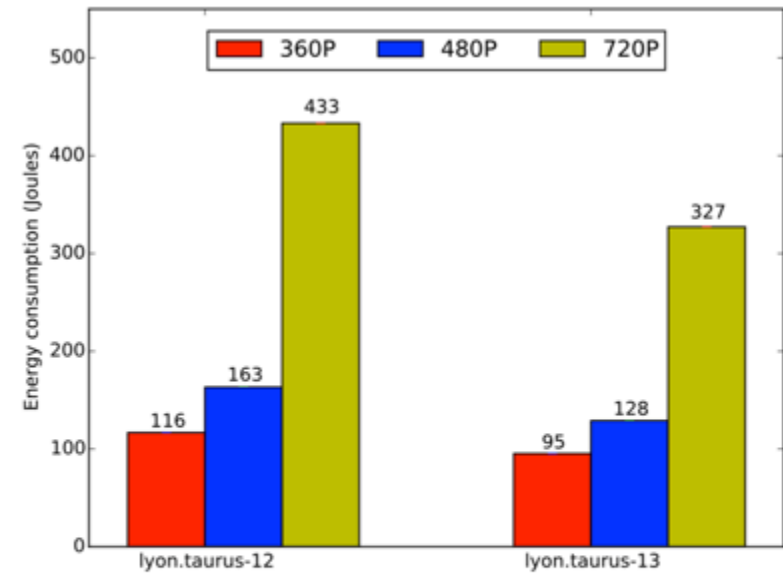
Better performance with one large VM
Large VMs less easy to consolidate, repair, etc.

Depends on application's resource usage

Power and energy consumption



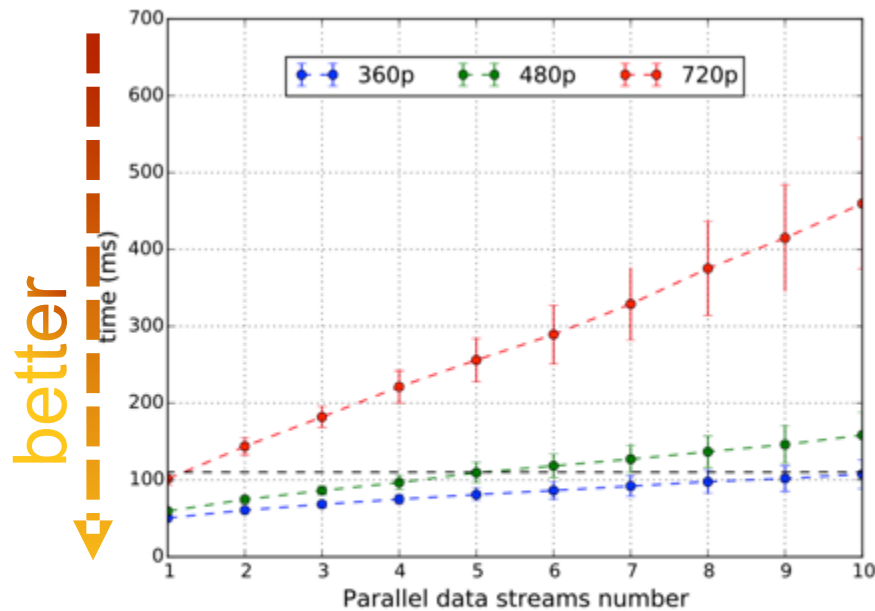
(c) Power consumption for 4 small VMs on Taurus-12 and 1 large VM on Taurus-13 for the same amount of computation



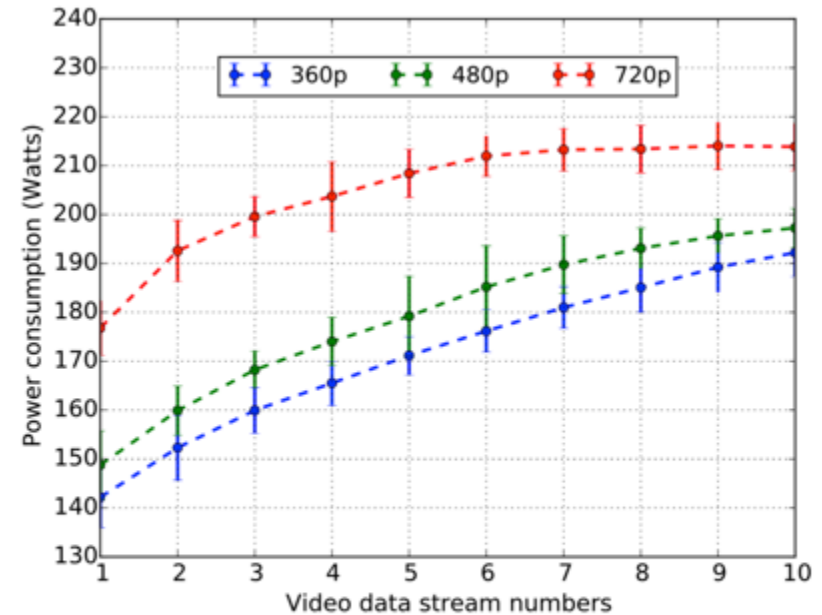
(d) Energy consumption for analyzing a 5 mn video on Taurus-12 with 4 small VMs and on Taurus-13 with 1 large VM

Power consumptions almost equivalent
Better energy consumption with large VM

Consolidation within a single VM



(e) Analysis time with parallel data streams in a large VM

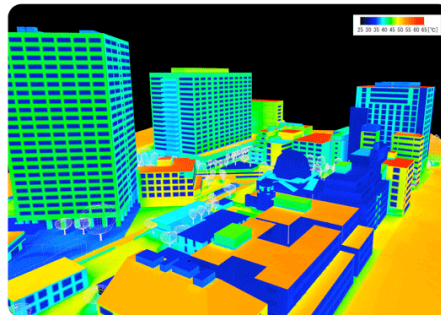


(f) Power consumption with parallel data streams in a large VM

8 frames per second to analyze: 0.125 ms per frame max
A large VM can handle: 11 360p streams, 5 480p streams and 1 720p stream.

Depends on required application accuracy.

Simulations

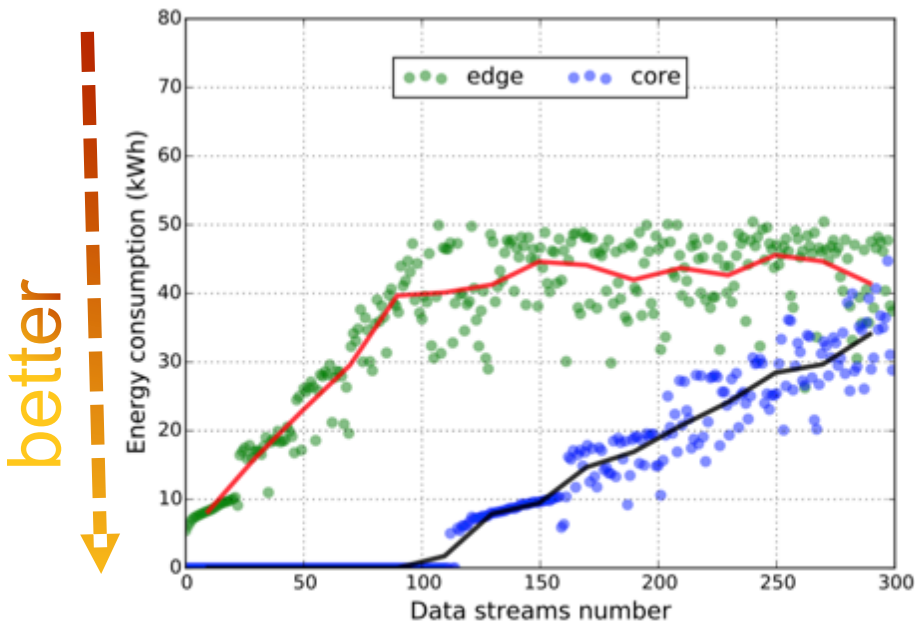


Cloud configurations

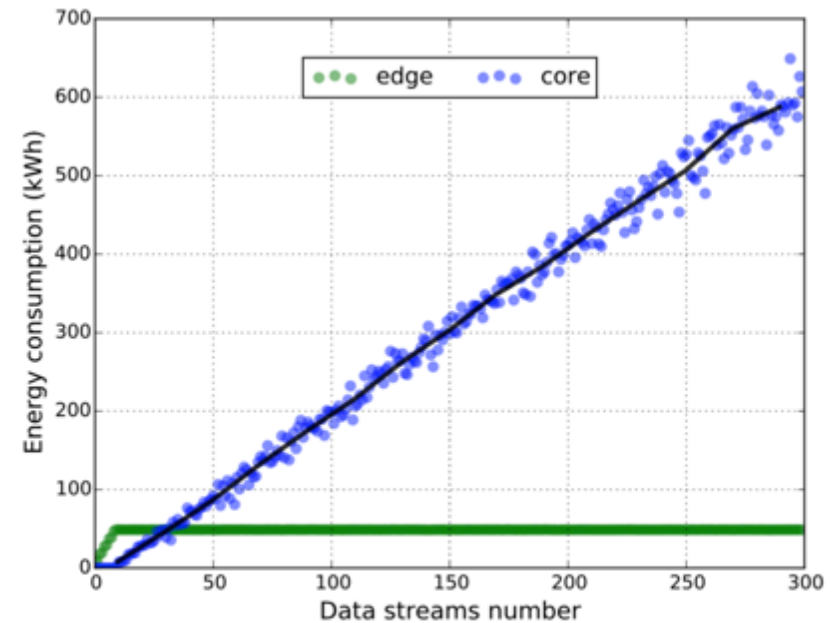


- **Core cloud**
 - 100 servers
 - 100 ms latency with the edge devices
- **Edge cloud**
 - 5 servers
- **Unused resources are switched off.**

Energy consumption at edge & core



(a) Energy consumption with resolution 360p

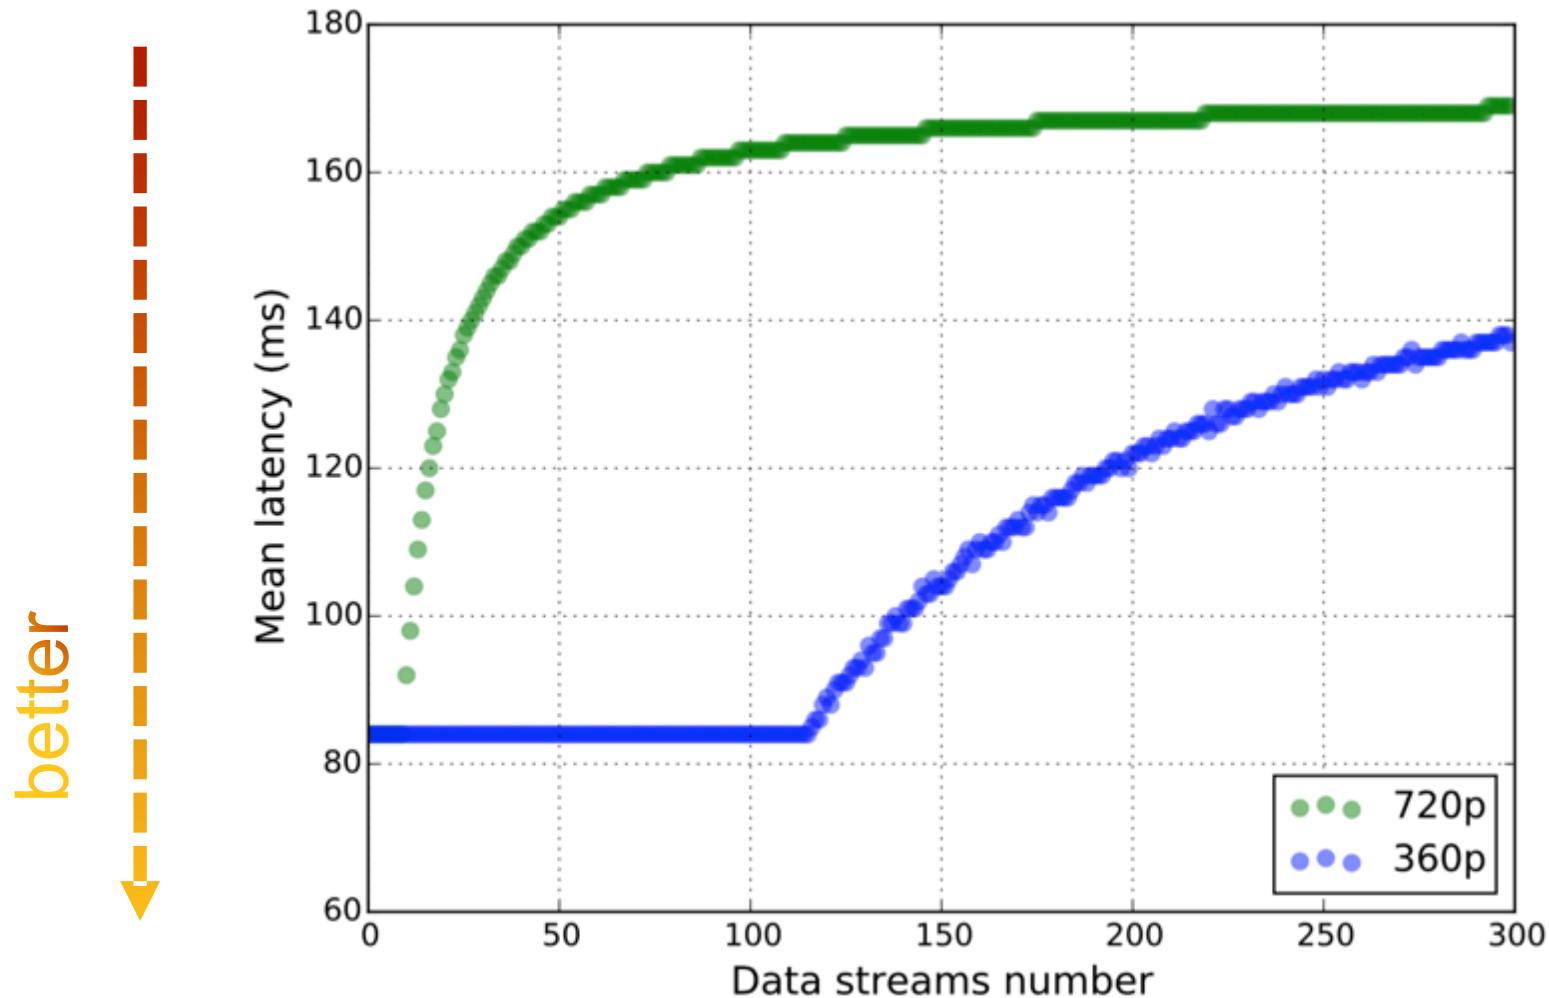


(b) Energy consumption with resolution 720p

Edge can handle: 112 360p data streams and
16 720p data streams.

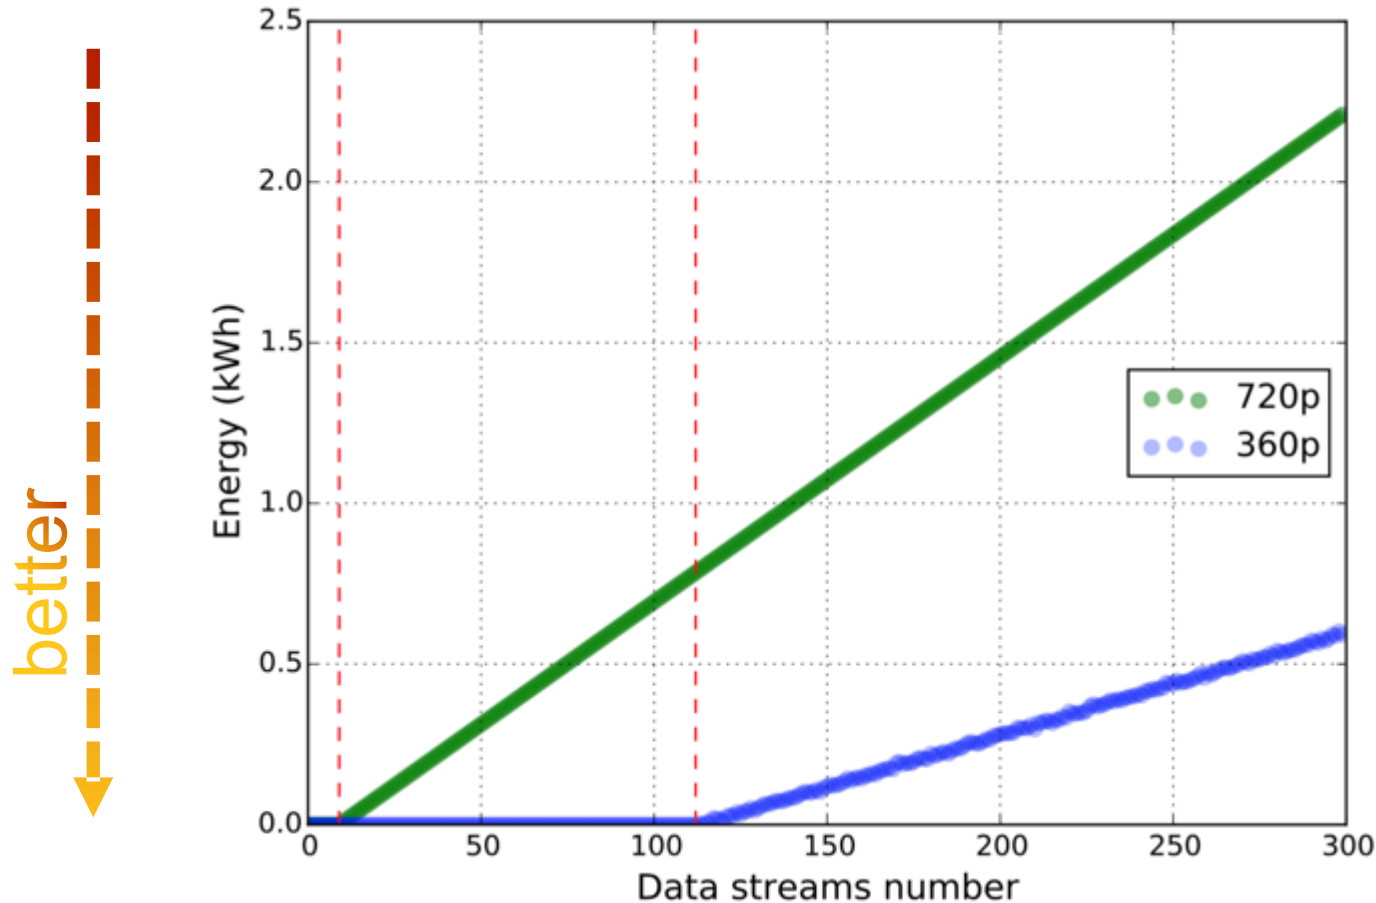
Depends on servers' architecture.

Average application latency



Depends on edge's resources availability

Network energy consumption



Cost per-bit
energy model
for network

Model from: F.Jalali,
K.Hinton, R.Ayre,
T.Alpcan, R.S.Tucker,
“Fog Computing May
Help to Save Energy in
Cloud Computing”,
JSAC 34 (5), 2016.

Depends on application traffic and edge resources

Application accuracy

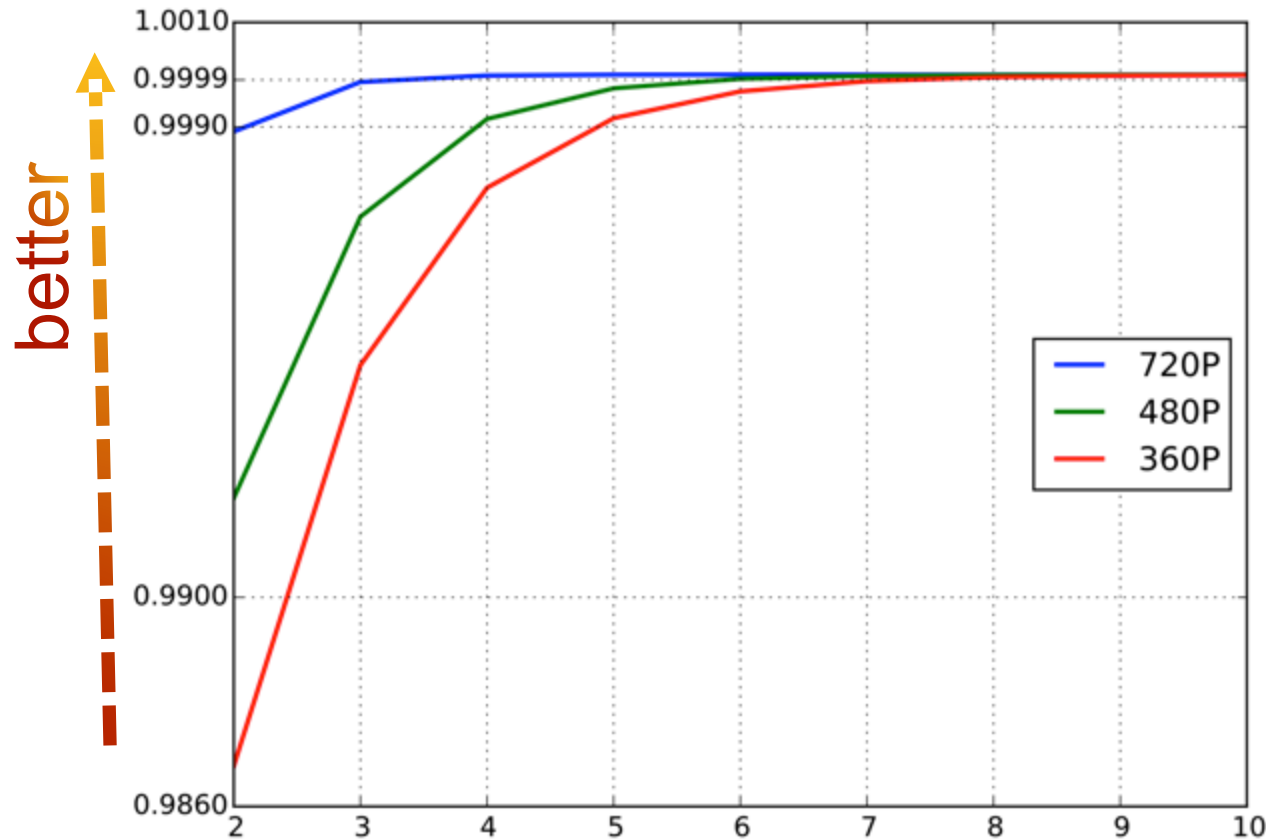
Object detection accuracy

better ↑

Classes	720p	480p	360p
car	96.7%	91%	88.5%
body	97.7%	94.9%	90.7%
dog	96.1%	94.9%	90.7%
total	96.7%	92.3%	87.9%

**Is it better to have 1 car with 720p resolution
or 2 cars with 360p resolutions?**

Reliability



Not feasible in real-time with 1 VM

Feasible in real-time with 1 VM

Number of data streams for reaching required accuracy

# nines	720p	480p	360p
99.9%	3	4	6
99.99%	4	6	8
99.999%	5	7	11

Summary

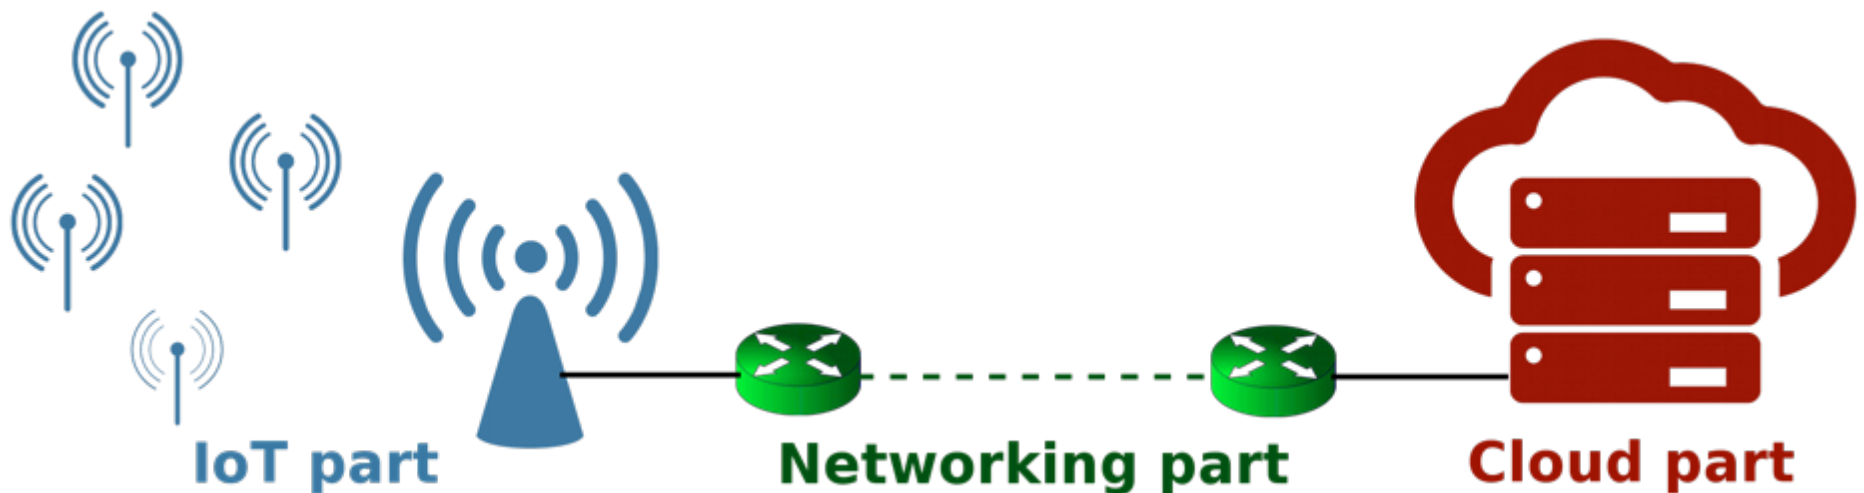
Offloading the data to process video streams at edge:

- Effectively reduces the response time
- Avoids unnecessary data transmission between devices and core Cloud
- Best configuration highly depending on the applications' characteristics



What about the other parts?

Which part consumes the most?



“End-to-end Energy Models for Edge Cloud-based IoT Platforms: Application to Data Stream Analysis in IoT”, Y. Li, A.-C. Orgerie, I. Roderio, B. Lemma Amersho, M. Parashar, J.-M. Menaud, FGCS, vol. 87, p 667-678, 2018.

Parameters of our example

Parameter	Value
Voltage	3.3 V
Idle current	0.273 A
CCA Busy State current	0.273 A
Tx current	0.38 A
Rx current	0.313 A
Channel Switching current	0.273 A
Sleep current	0.033 A

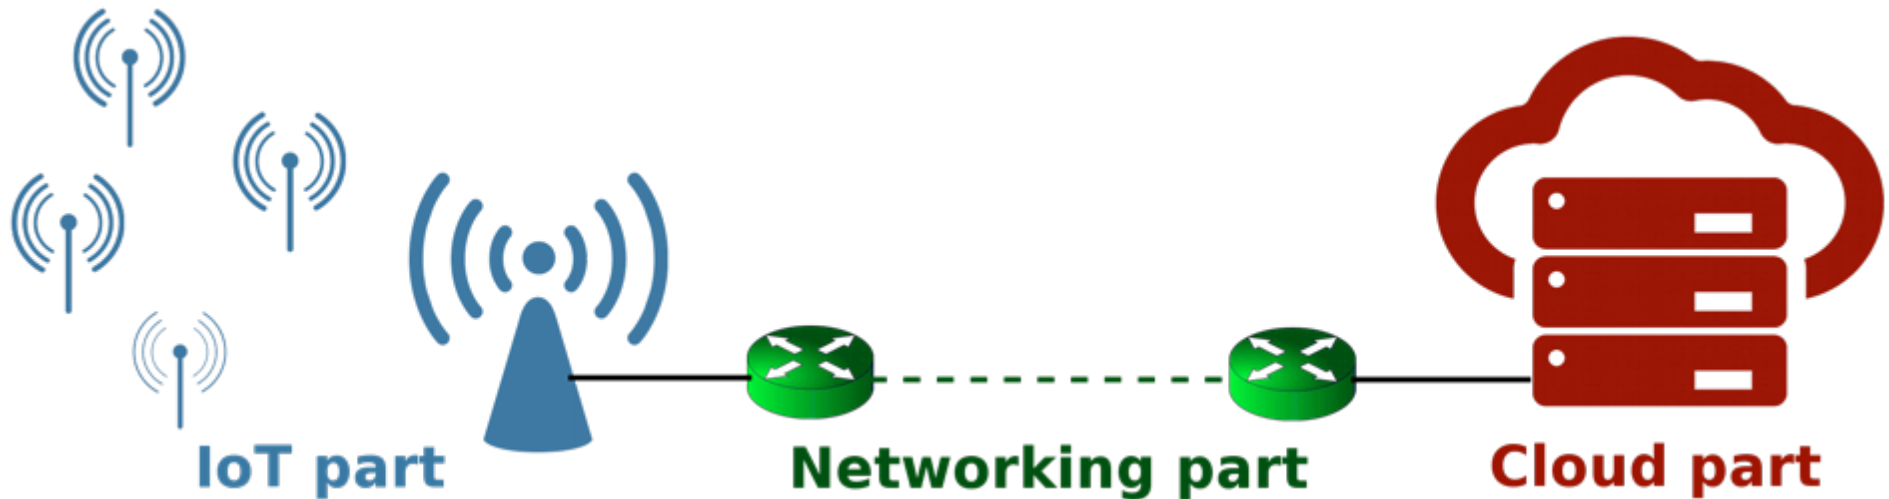
IoT devices (camera)

Network devices

Cloud data centers
PUE = 1.7 for edge
PUE = 1.2 for core

Parameter	Edge router	Core router
Idle consumption	4,095 Watts	11,070 Watts
Max consumption	4,550 Watts	12,300 Watts
Traffic	560 Gbps	4,480 Gbps
Energy	37 nJ/bit	12.6 nJ/bit

Experimental setup



Simulations



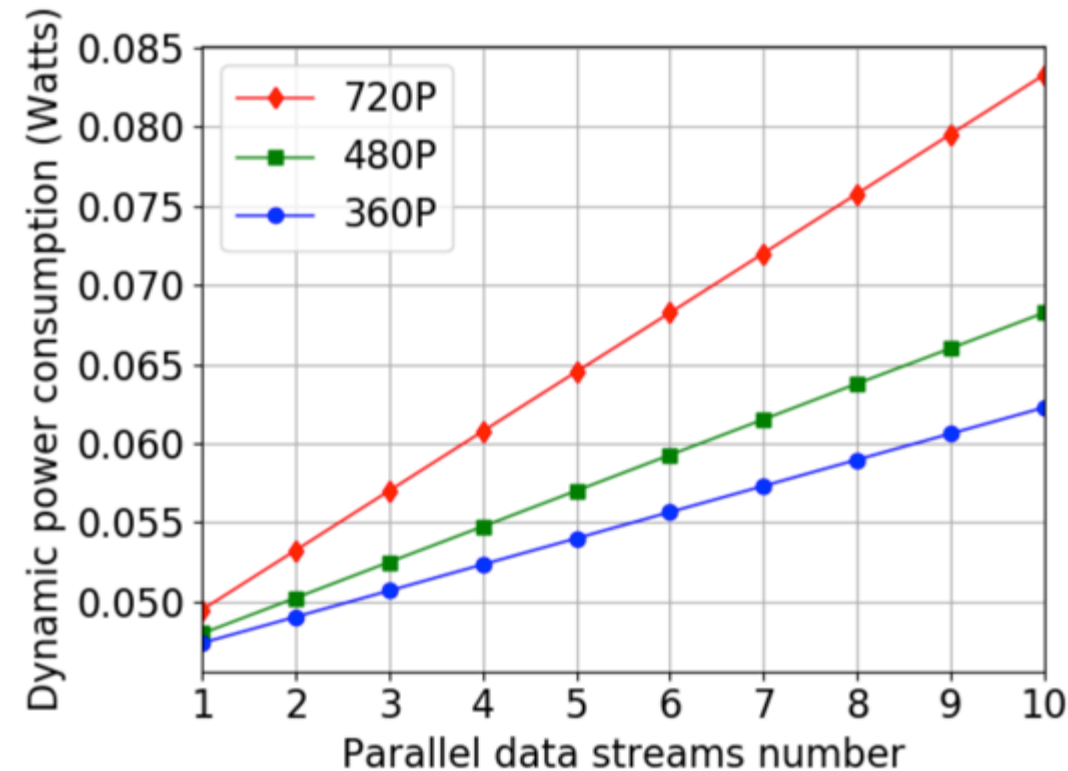
Model from literature



Real measurements



IoT consumption per device

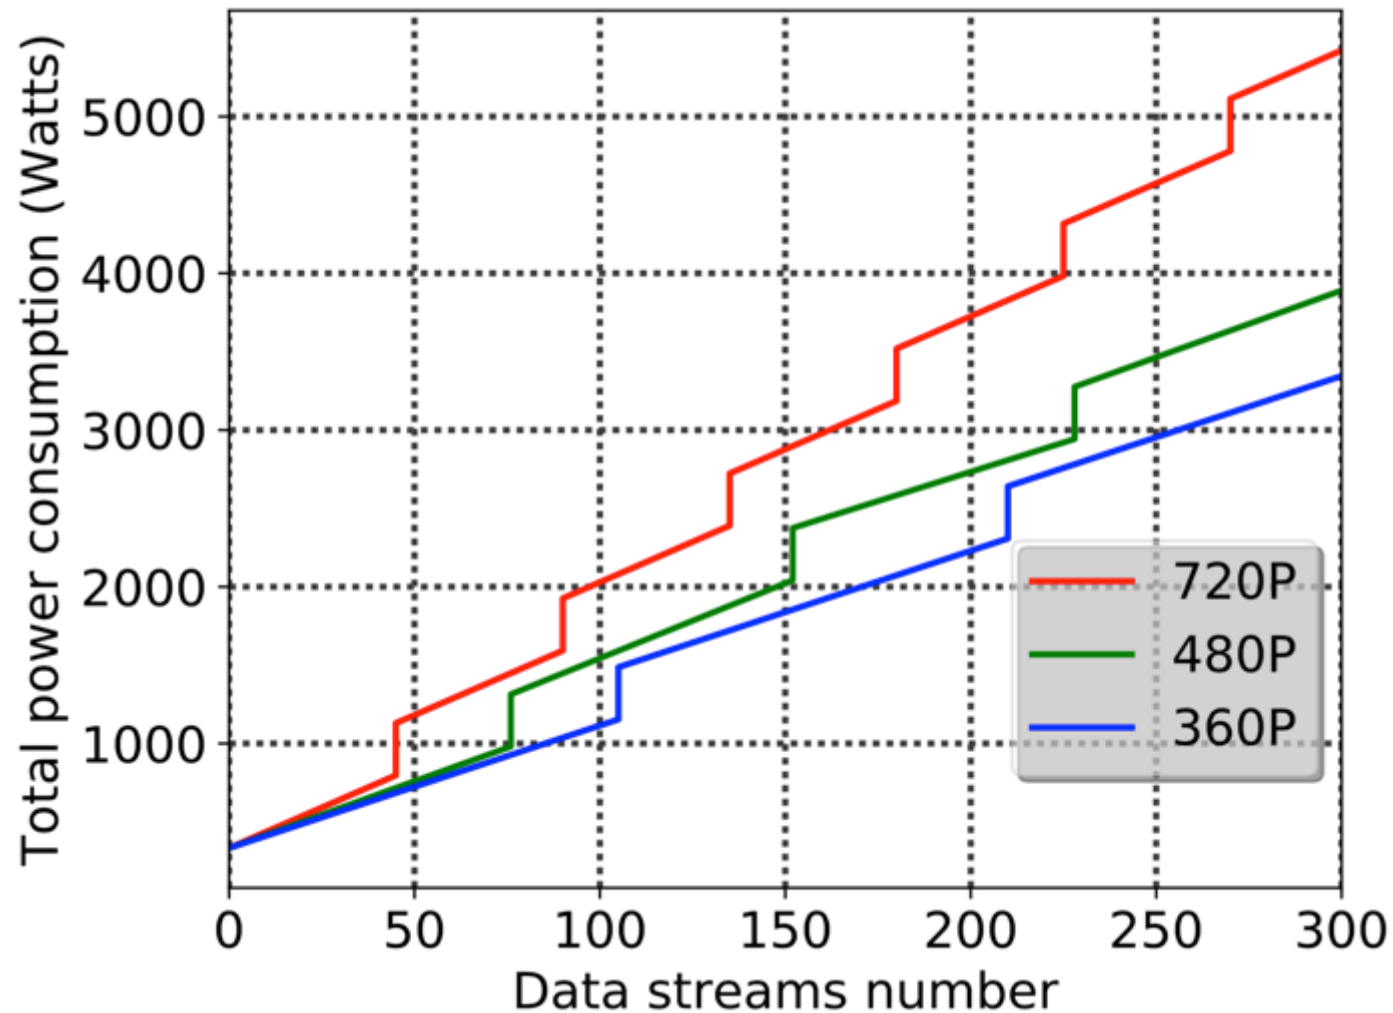


Dynamic power consumption

# devices	360p	480p	720p
1	6.907	6.908	6.909
2	12.869	12.87	12.873
3	18.831	18.832	18.837
4	24.792	24.795	24.801
5	30.754	30.757	30.765
6	36.716	36.719	36.728
7	42.677	42.682	42.692
8	48.639	48.644	48.656
9	54.601	54.606	54.62
10	60.562	60.568	60.583

Overall power consumption

IoT part including access point





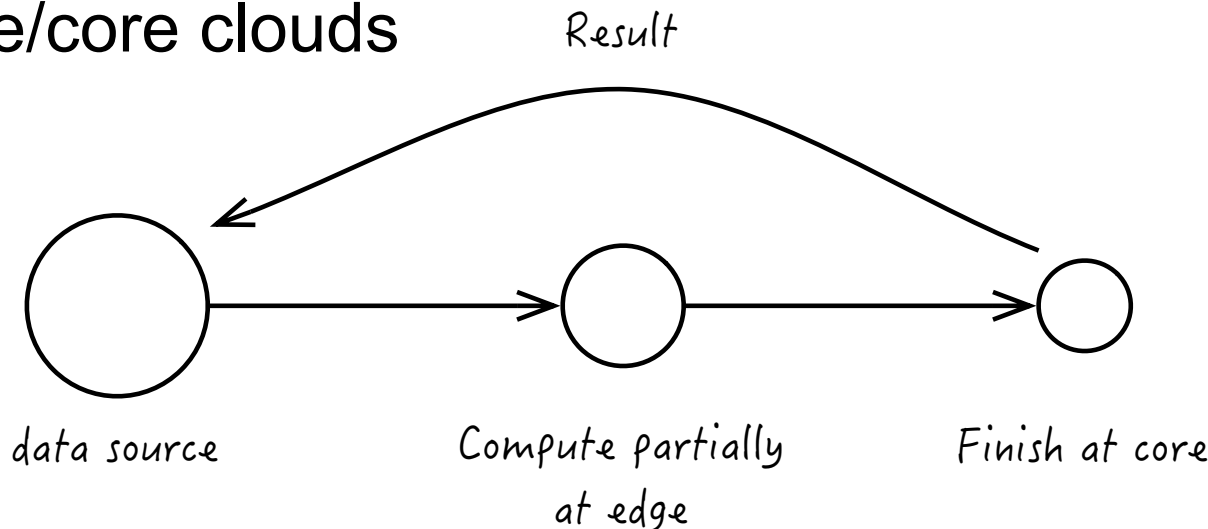
Overall evaluation

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

- Cost per 360p stream for each part
- Consumption when in use
- Not including all infrastructure costs
- **IoT part:** accurate for the given scenario in an ideal case (without loss on the 802.11 network)
- **Network part:** following literature model (based on average Internet traffic, so probably underestimated)
- **Cloud part:** measured, accurate on the given servers

Conclusions

- Typical application: data stream analysis for IoT devices and applications
- Real power and performance measurements on a concrete use-case (for the Cloud part)
- Exploration of possible trade-offs between performance (response time and accuracy) and energy consumption (*green and brown*)
- First step towards energy-aware IoT applications relying on edge/core clouds



Conclusions

- **End-to-end** energy consumption
- Cloud part non negligible
- Started with the study of a given application
- Extending existing simulators with generic validated energy models
- On-going work...
 - Other IoT devices
 - Using other network protocols



The ICT world



The (in)dispensable weather toaster

<https://www.behance.net/gallery/4417181/Jamy-Smart-Toaster>



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

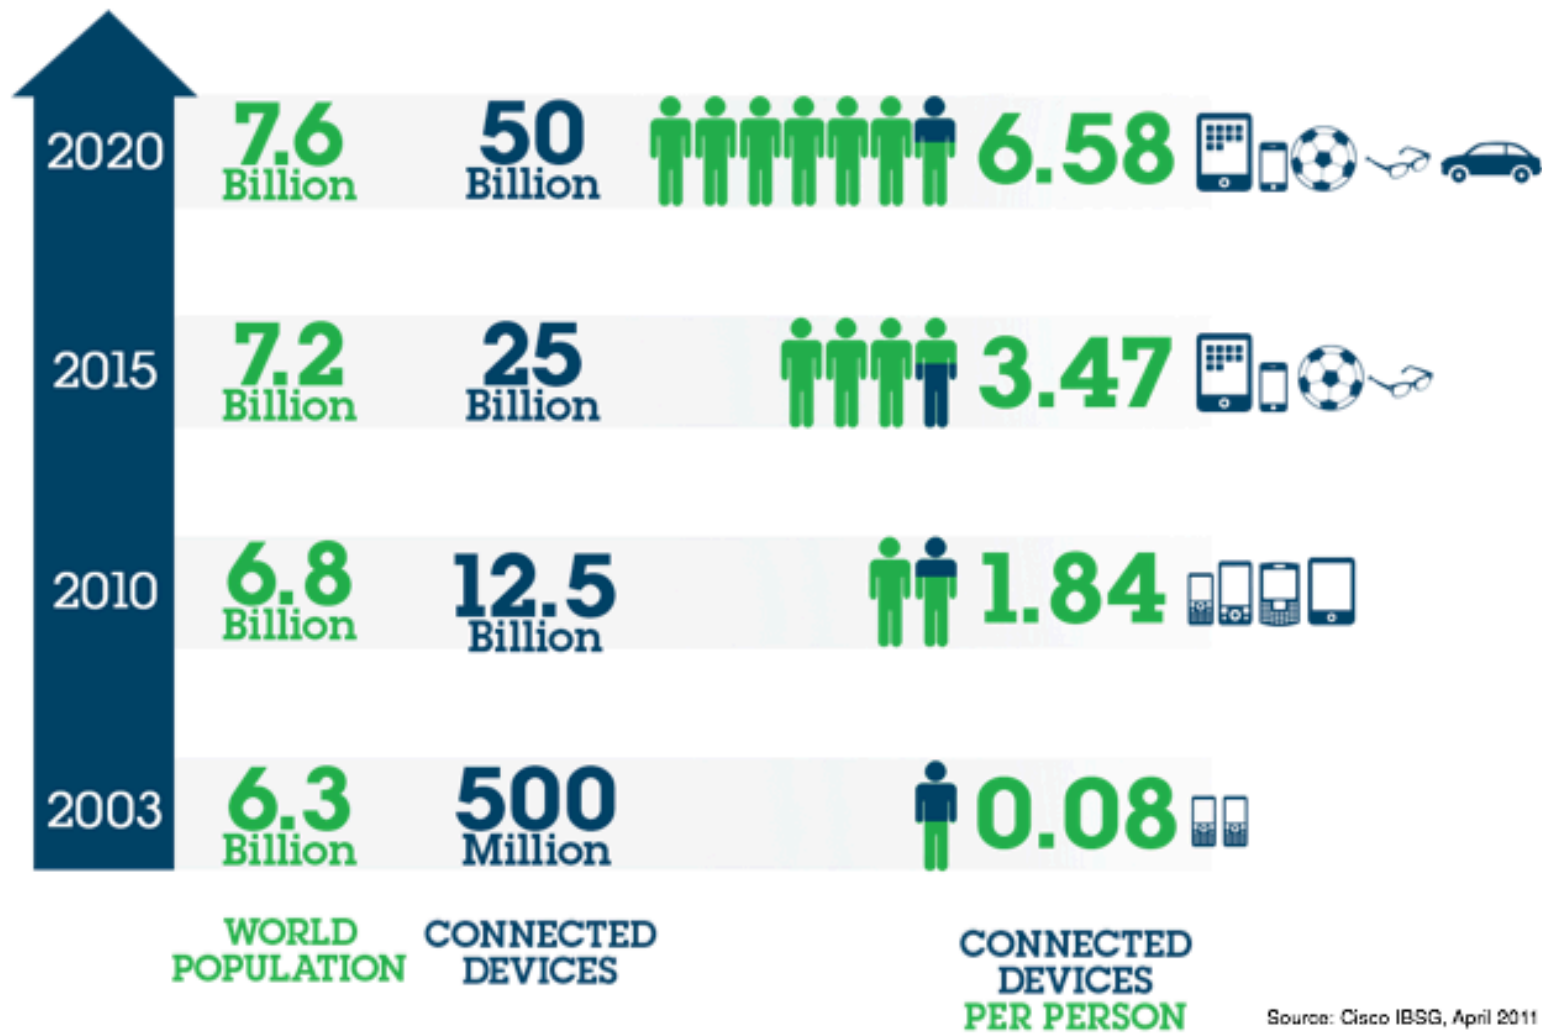
Forrester Research, "Connected devices forecast, 2012 to 2017", white paper, 2013.

The smart frying pan



Assistéo, Tefal, 2017.

Internet of all the Things



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

Are we going on the good way?



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

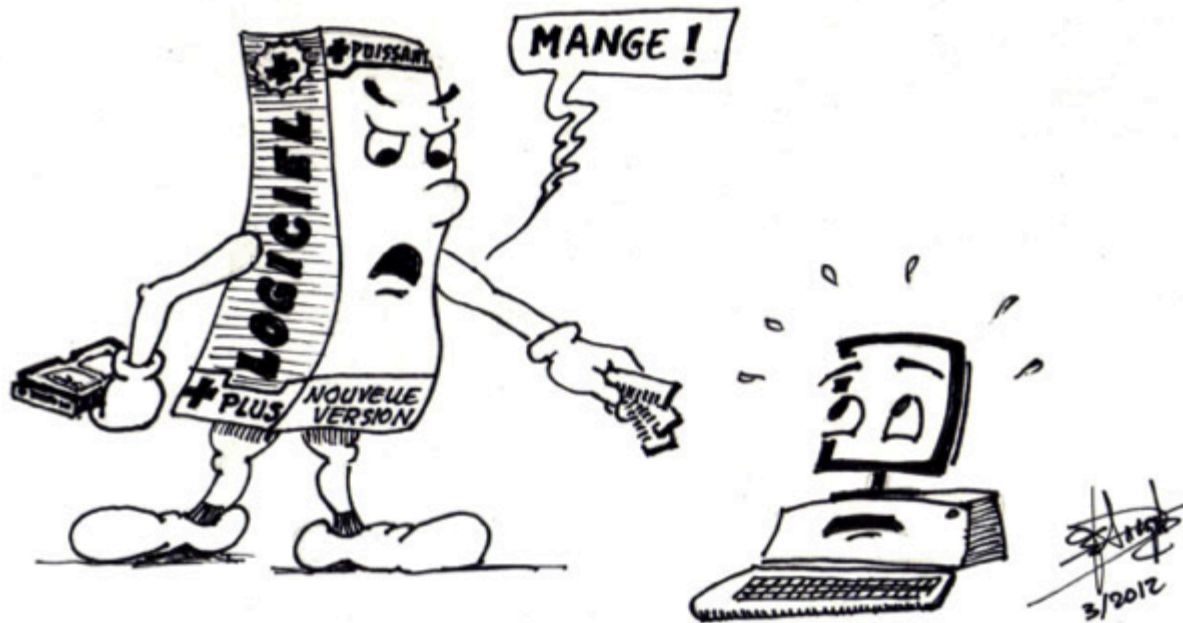
Complete life-cycle



- 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.

User = responsible person



- 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, “*Usages & attitudes vis-à-vis des applications mobiles*”, survey, 2015.

<http://ecoinfo.cnrs.fr>



[Le GDS EcoInfo](#) [Services](#) [Thématiques](#) [Formations](#) [Solutions](#) [Communications](#) [Les confs EcoInfo](#)

GDS EcoInfo

Pour une informatique éco-responsible

Groupement de Service (G.D.S.) EcoInfo : des ingénieurs et des chercheurs (CNRS, INRIA, ParisTech, Université Grenoble Alpes, université de Strasbourg, Université Aix-Marseille, etc....) à votre service pour réduire les impacts écologiques et sociétaux des TIC.



Copyright : © Natalya Yakovleva, iStockphoto / Fotolia / L. Revelin Falcoz - CNRS

conf éco-conception : supports en ligne

Partenaires



Livre



Plaquette



A la une

Thank you for your attention!

