# End-to-end Energy Models for Edge Cloud-based IoT Platforms: Application to Data Stream Analysis in IoT

**Anne-Cécile Orgerie**

GreenDays 2018
3rd July 2018, Toulouse

# Outline

- Context
- Edge clouds
- IoT use-case exploiting renewable energy
- End-to-end energy consumption
- Towards generic energy models…
- Conclusions

# ICT's electricity consumption

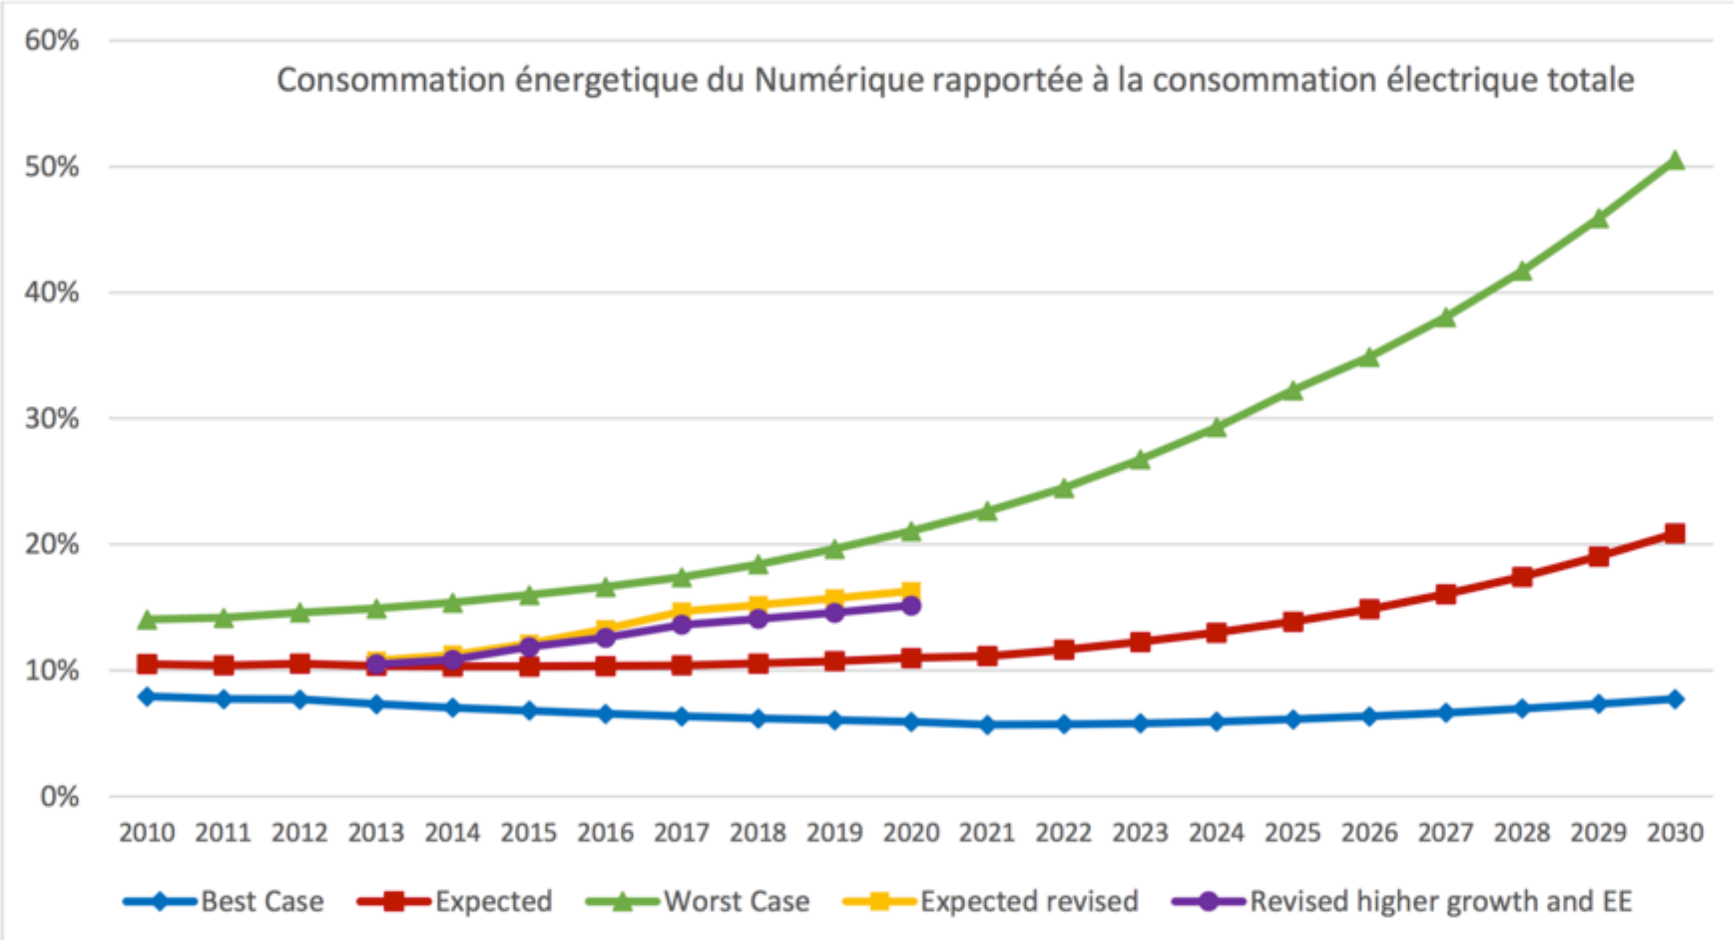


Figure 2 : Évolution 2010-2020 de la consommation énergétique du Numérique rapportée à la consommation électrique mondiale[7]
*[Source: calculé par The Shift Project à partir des données publiées par Andrae et Edler (2015)]*

**14% of the global electricity consumption in 2017.**

# Distribution of ICT energy consumption

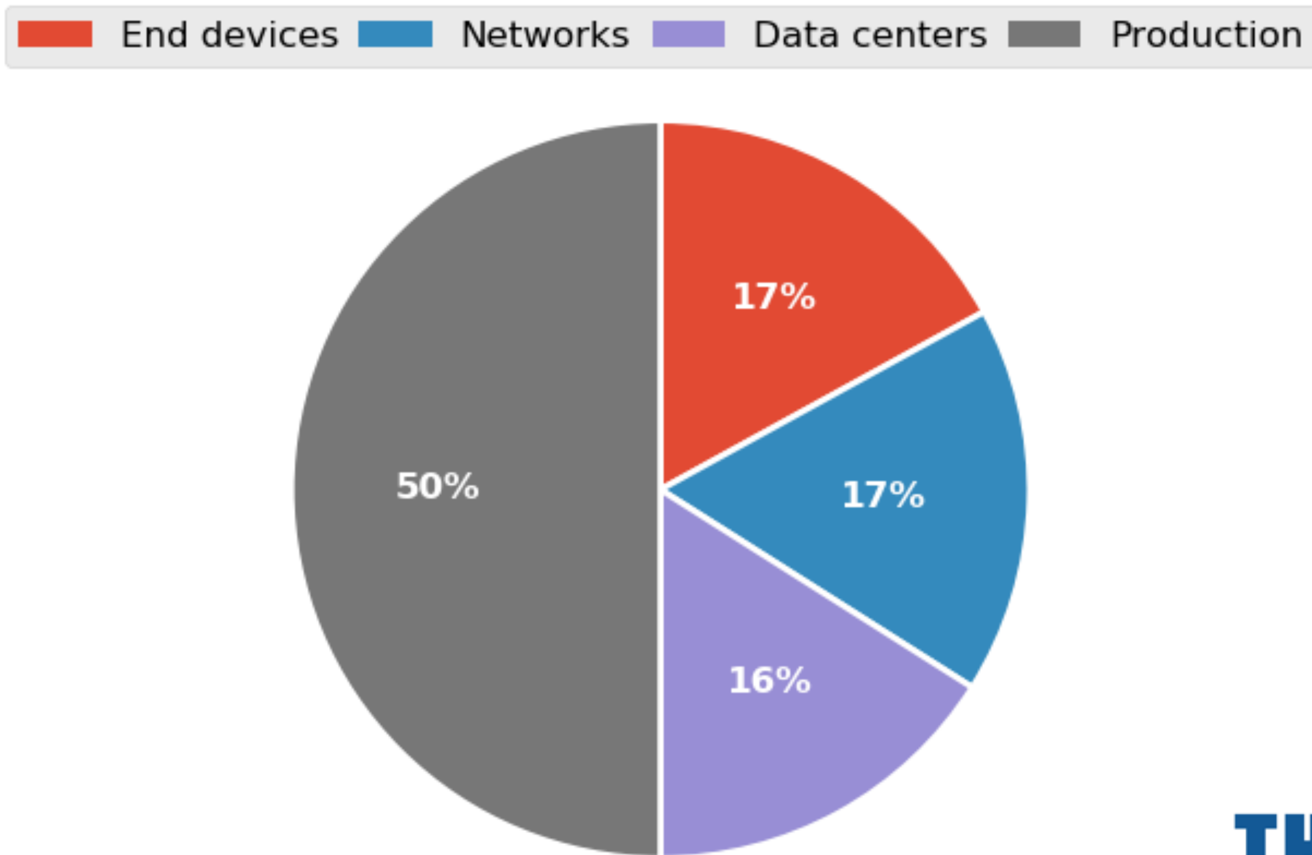


Rapport intermédiaire Lean ICT :
Pour une sobriété Numérique, 2018
https://theshiftproject.org

# IoT & Cloud

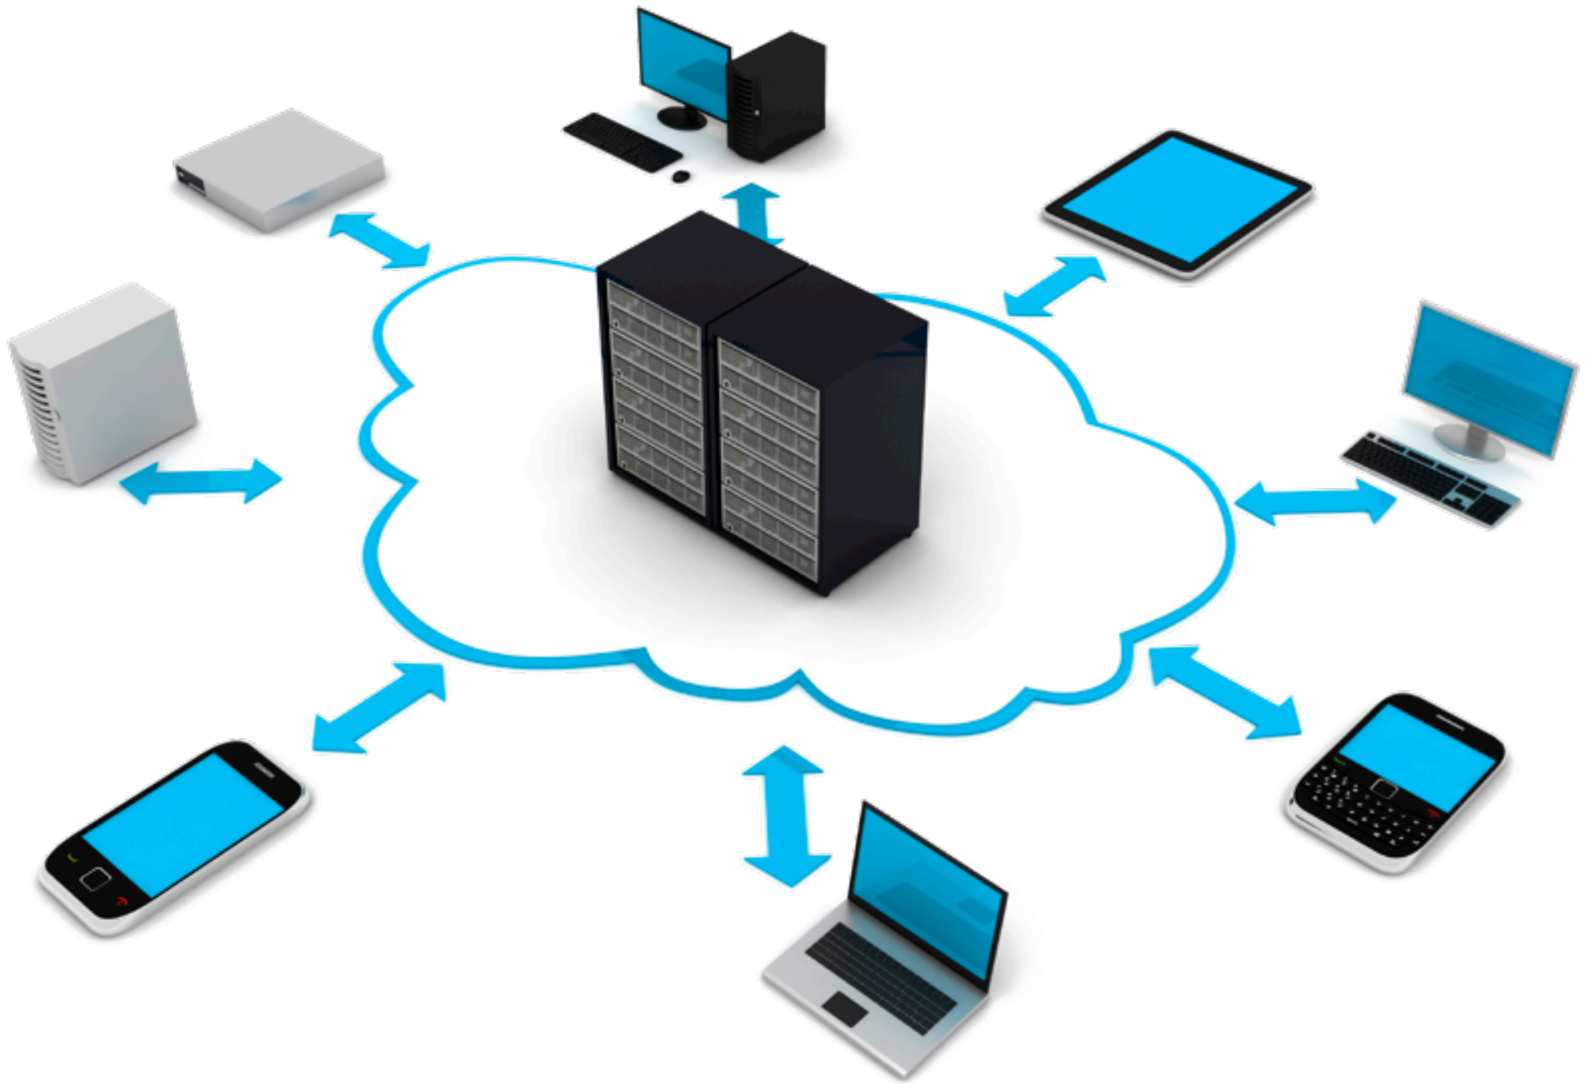

# Cloud computing promise

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

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

↳ **Economies of scale**

**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.

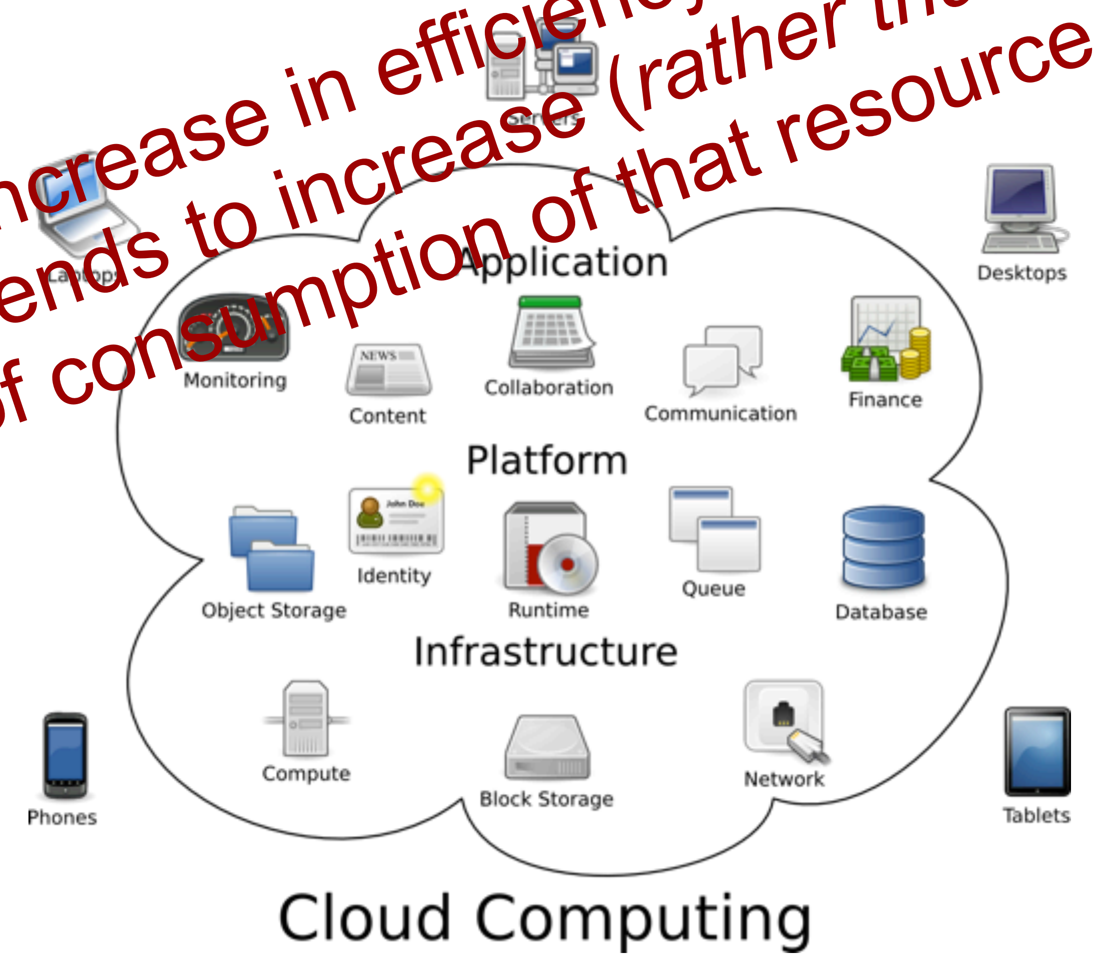

# IoT & edge cloud computing

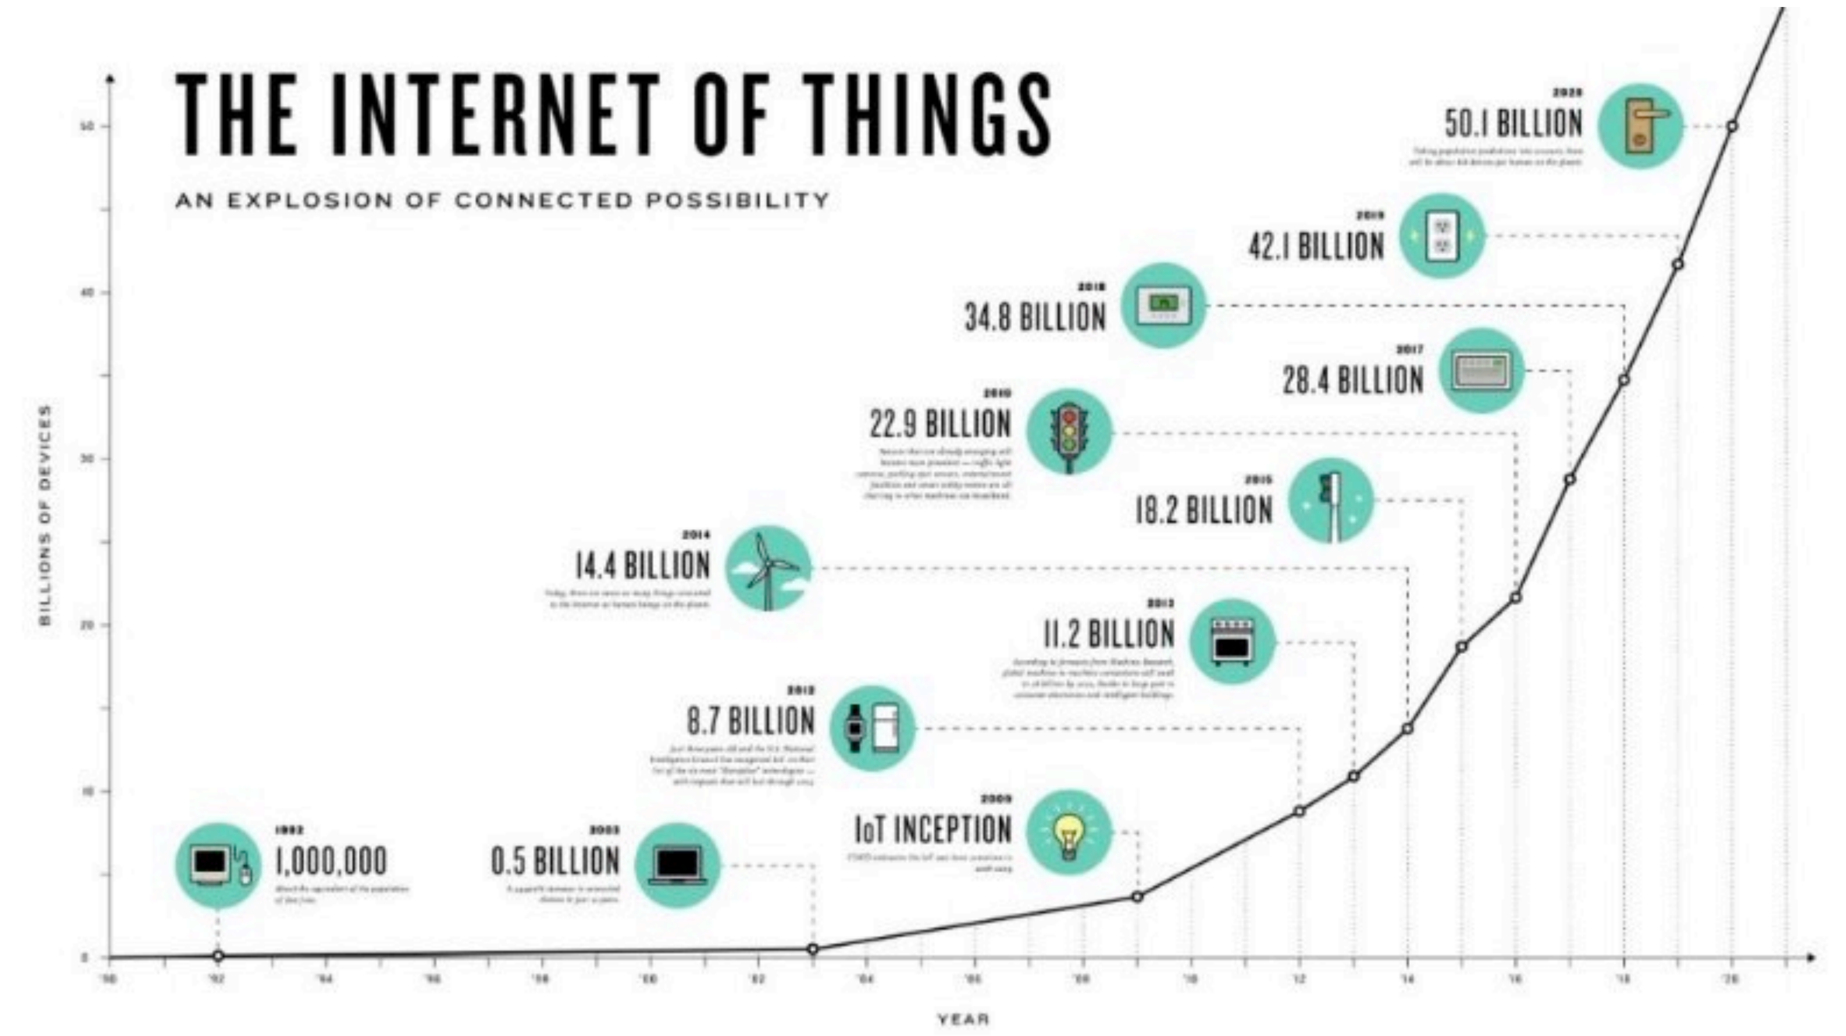

# New cloud architectures

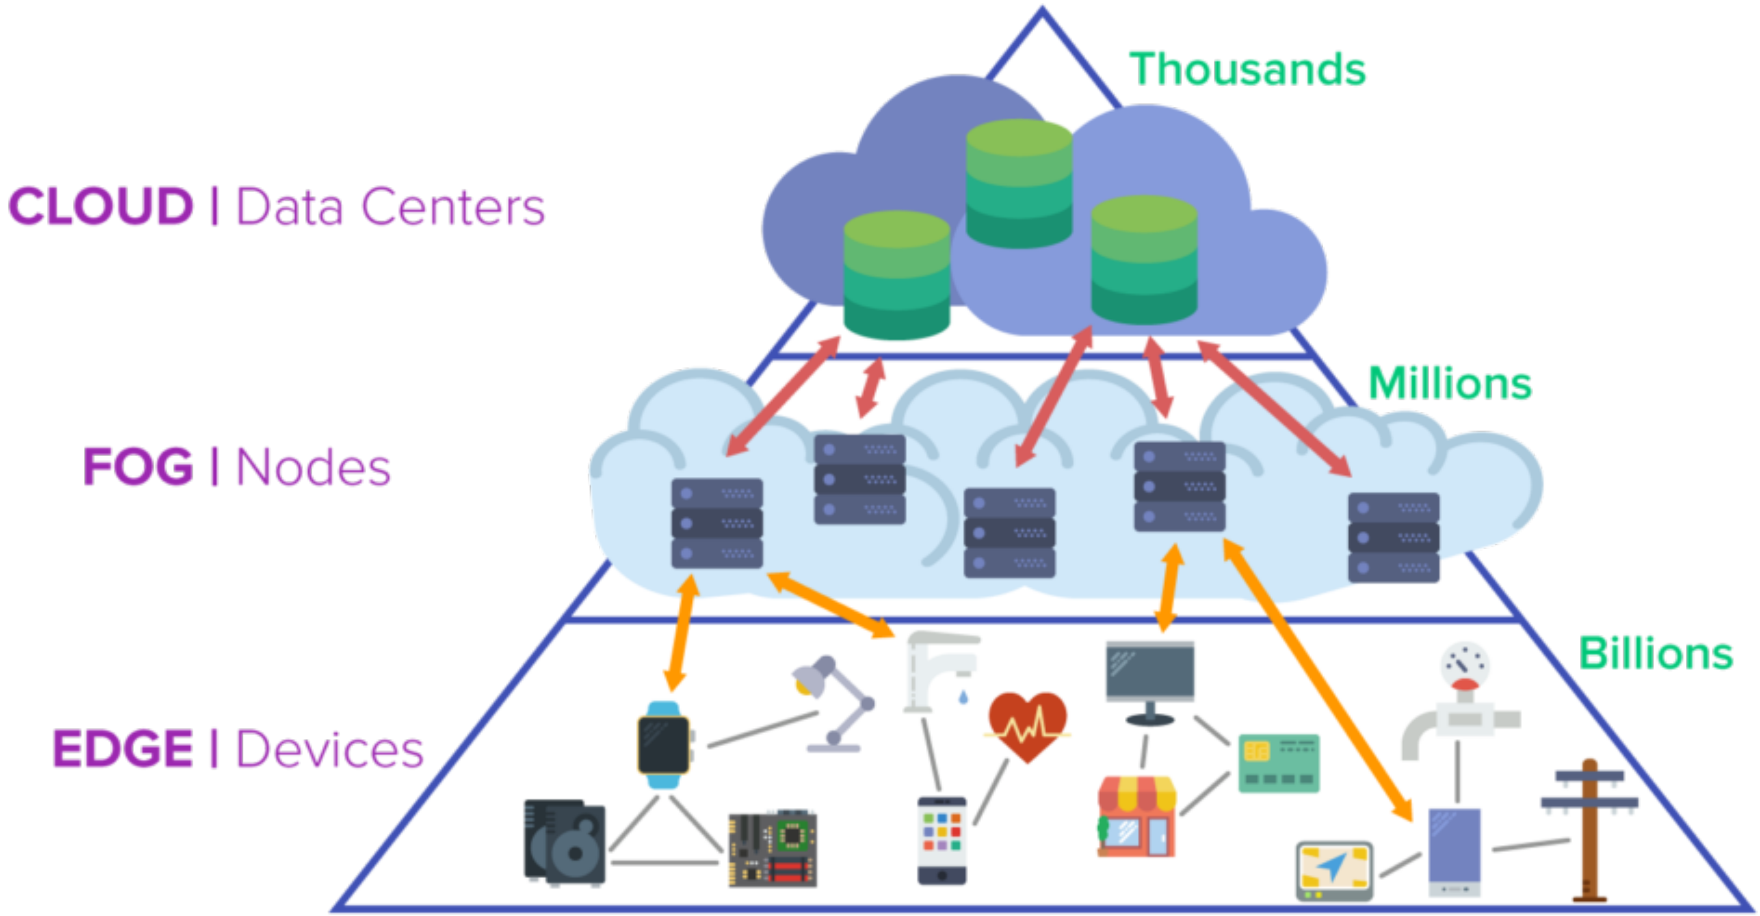


https://erpinnews.com/fog-computing-vs-edge-computing

# Edge Cloud

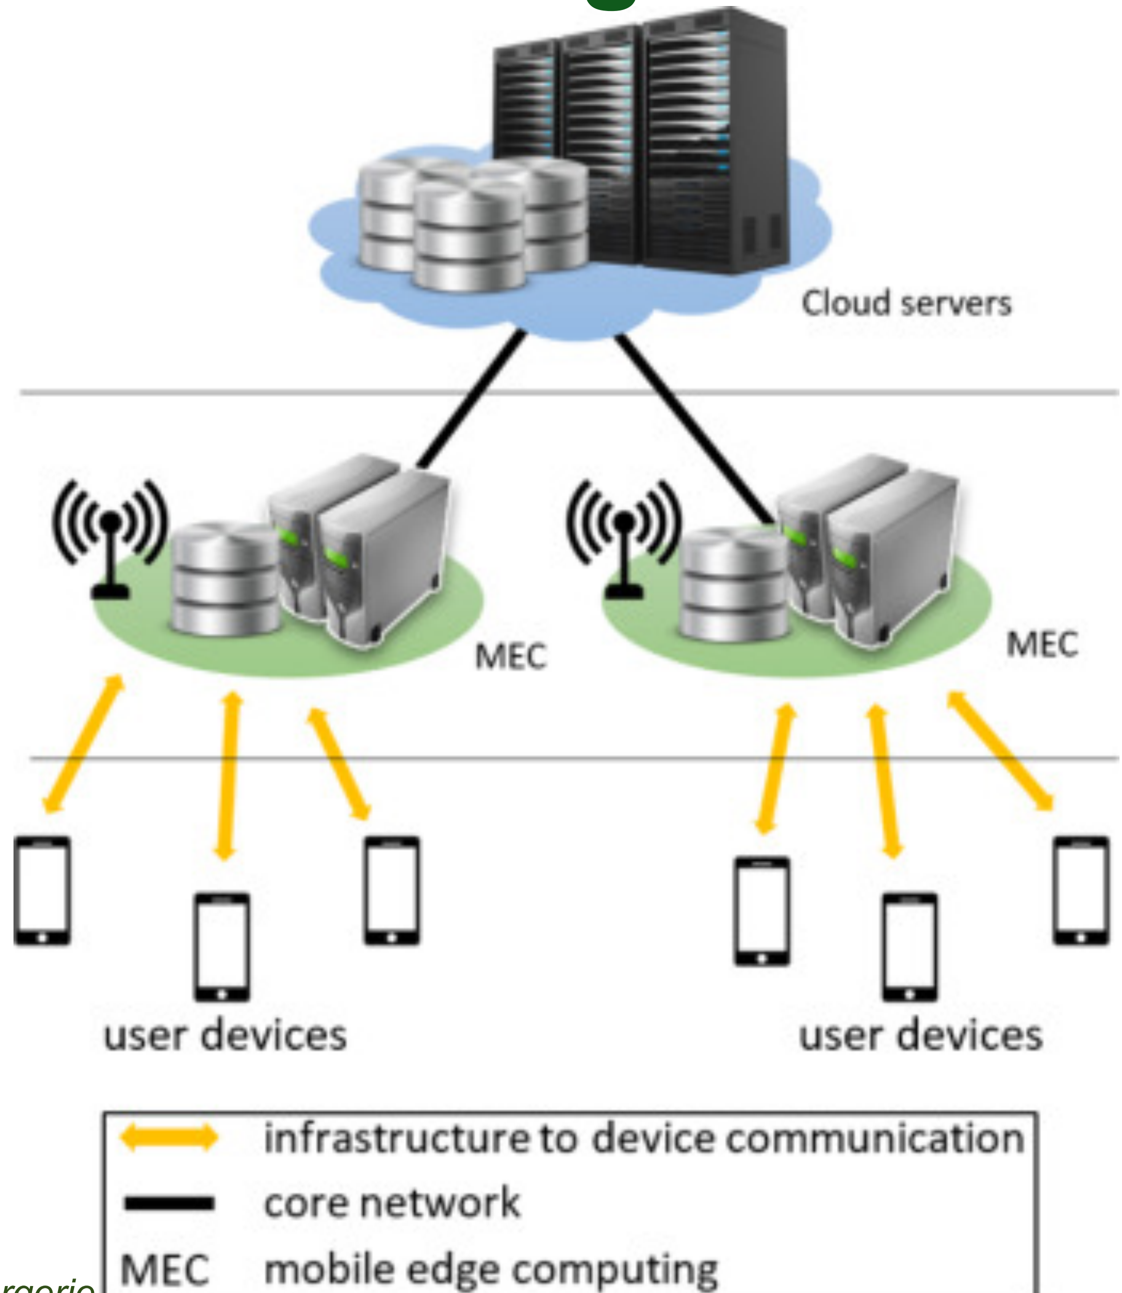


Stability
Availability
Latency

Low latency
Heterogeneity
Low capacity

"*Recovery for overloaded mobile edge computing*", D. Satria et al., FGCS 2017.

# Frightening examples

- Google in 2010 : 900 000 servers, ~ 2 billion kWh
- Facebook in 2012: ~ 532 million kWh
- In 2012, about 509 147 datacenters worldwide, ~ electricity consumption of 30 nuclear power
- In 2017, more than 20 billion devices connected.
- *More than 50 billion in 2023?*

**→ Electricity is a big problem in data centers**

Source : **J. Koomey**, *Growth in Data Center Electricity Use 2005 to 2010*. Analytics Press, août 2011.

https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/

# Using renewable sources

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 m$^2$) needed for solar farms**

# Data center's energy consumption

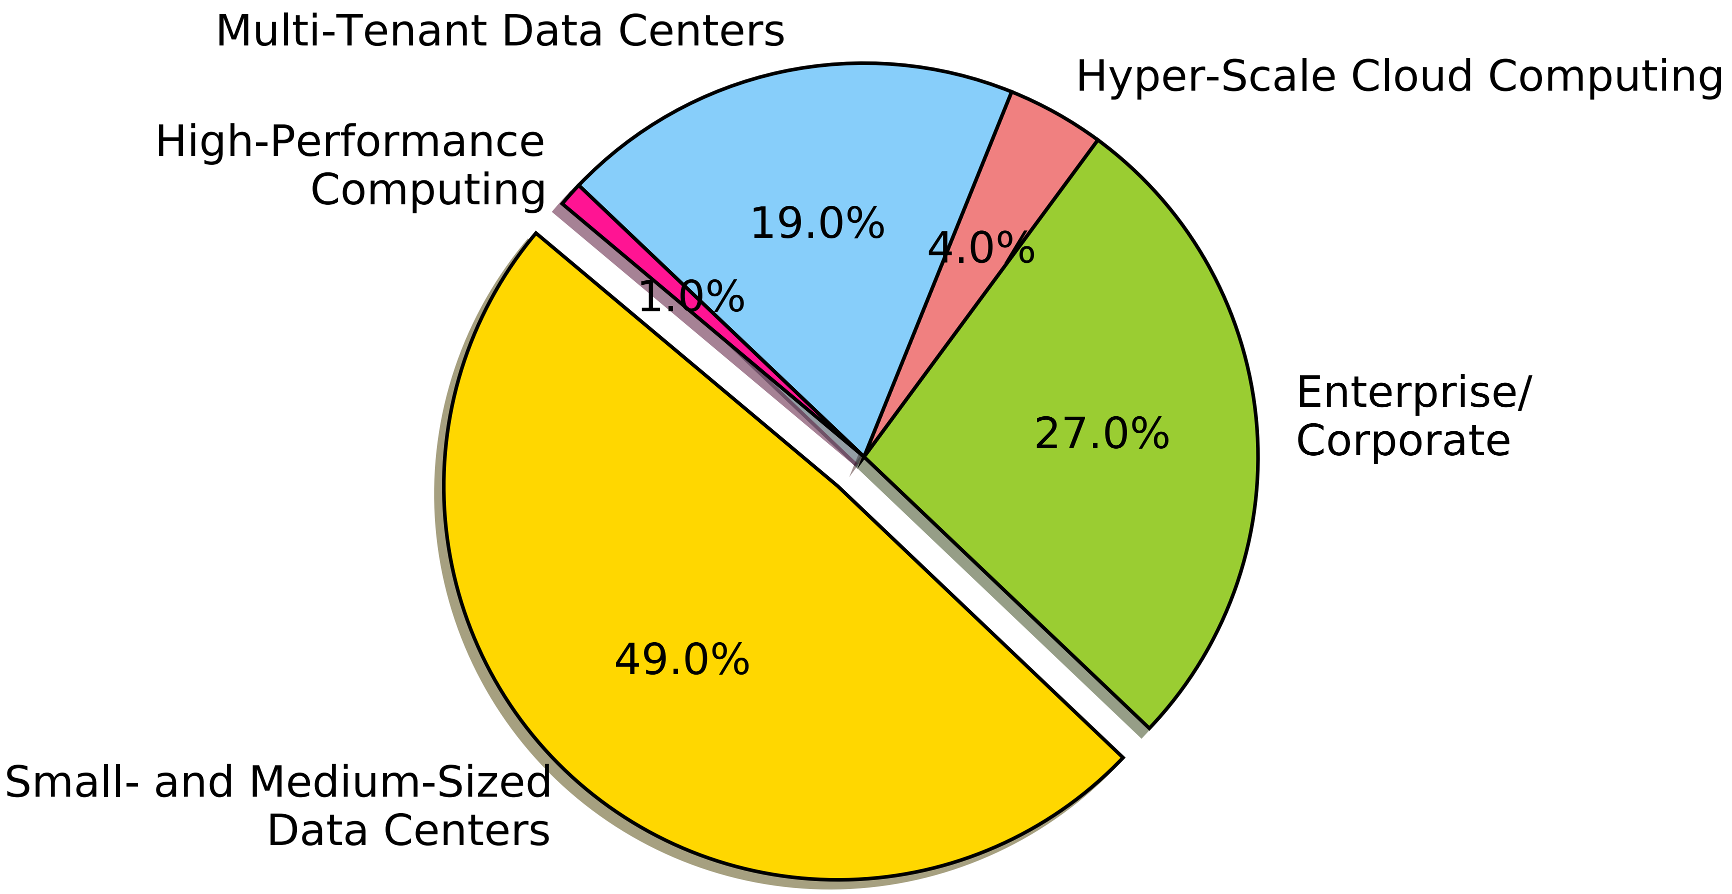


Estimated U.S. data center electricity consumption by market segment (2011)

Source: Data Center Efficiency Assessment, NRDC White paper, 2014.

# Renewable sources to power small DCs

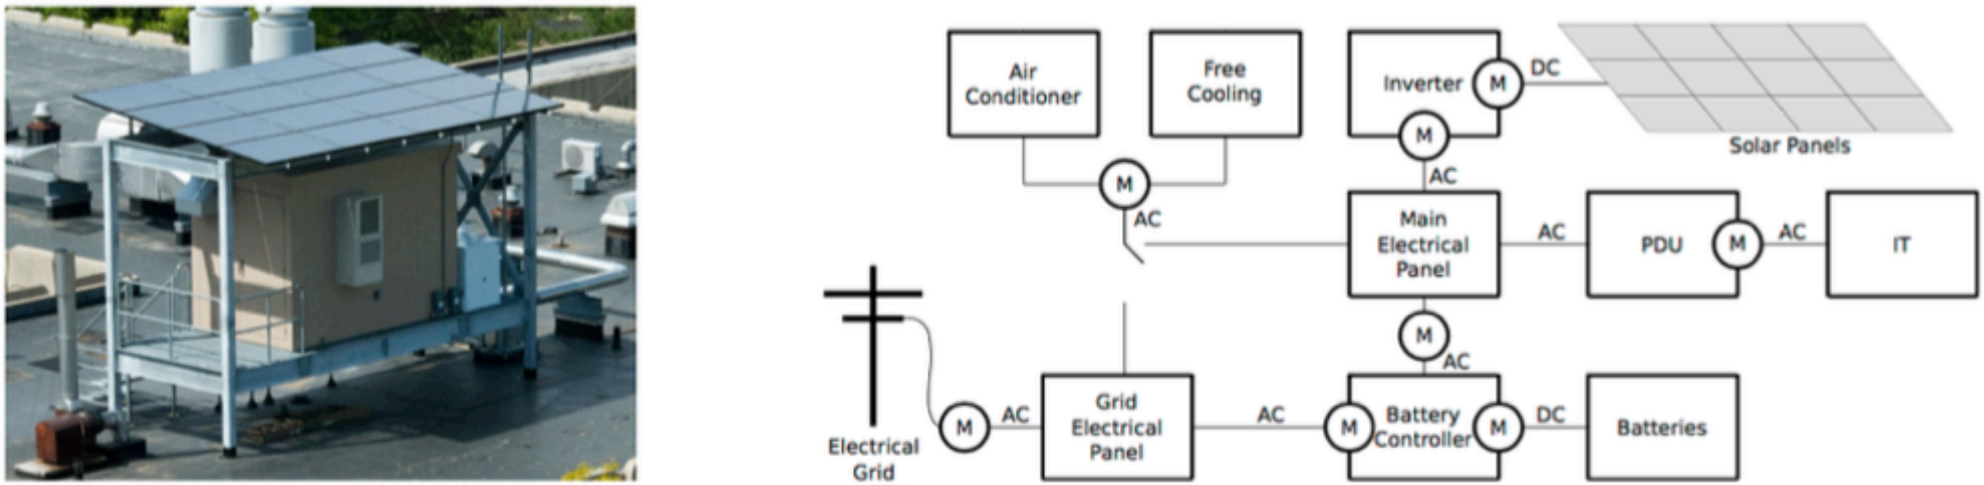


(a) (b)

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

Parasol, 2012, http://parasol.cs.rutgers.edu

THE STATE UNIVERSITY OF NEW JERSEY
RUTGERS

# Demand-side management

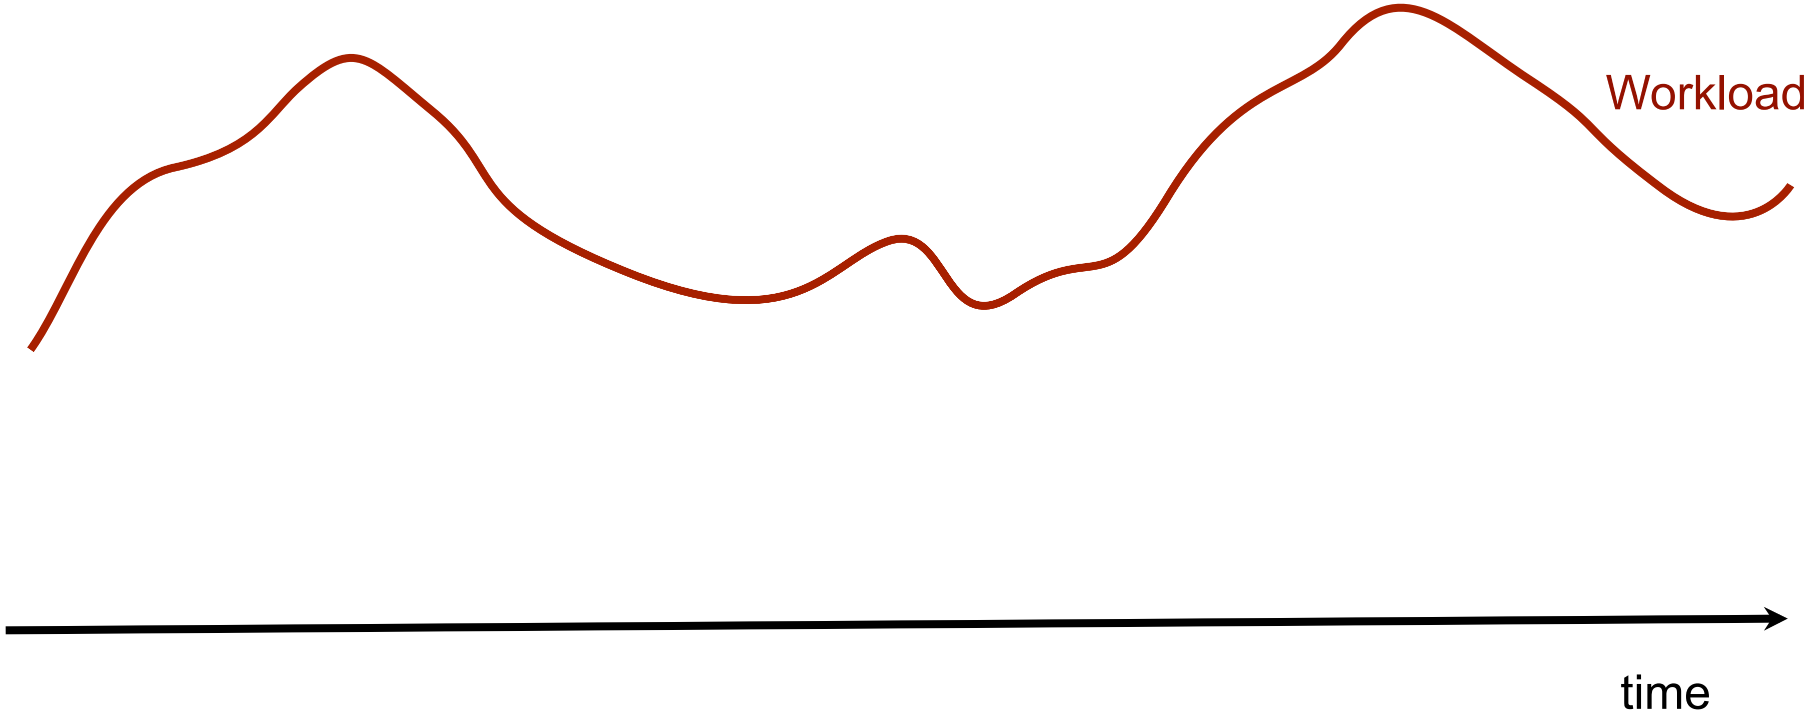


Courtesy of J-M Menaud

# Demand-side management

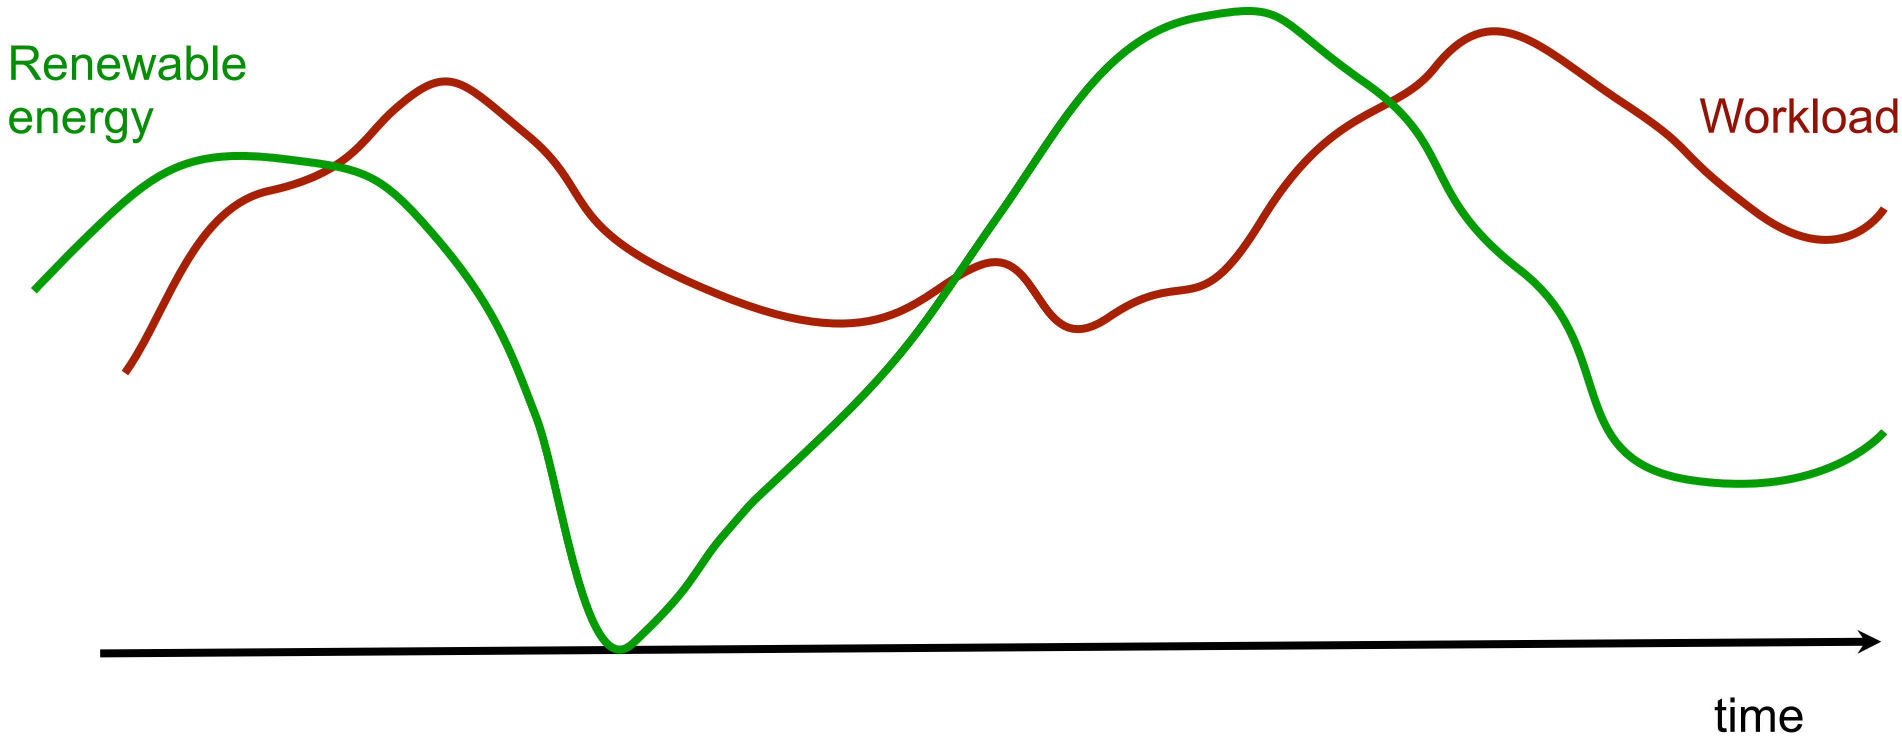


Courtesy of J-M Menaud

# Demand-side management

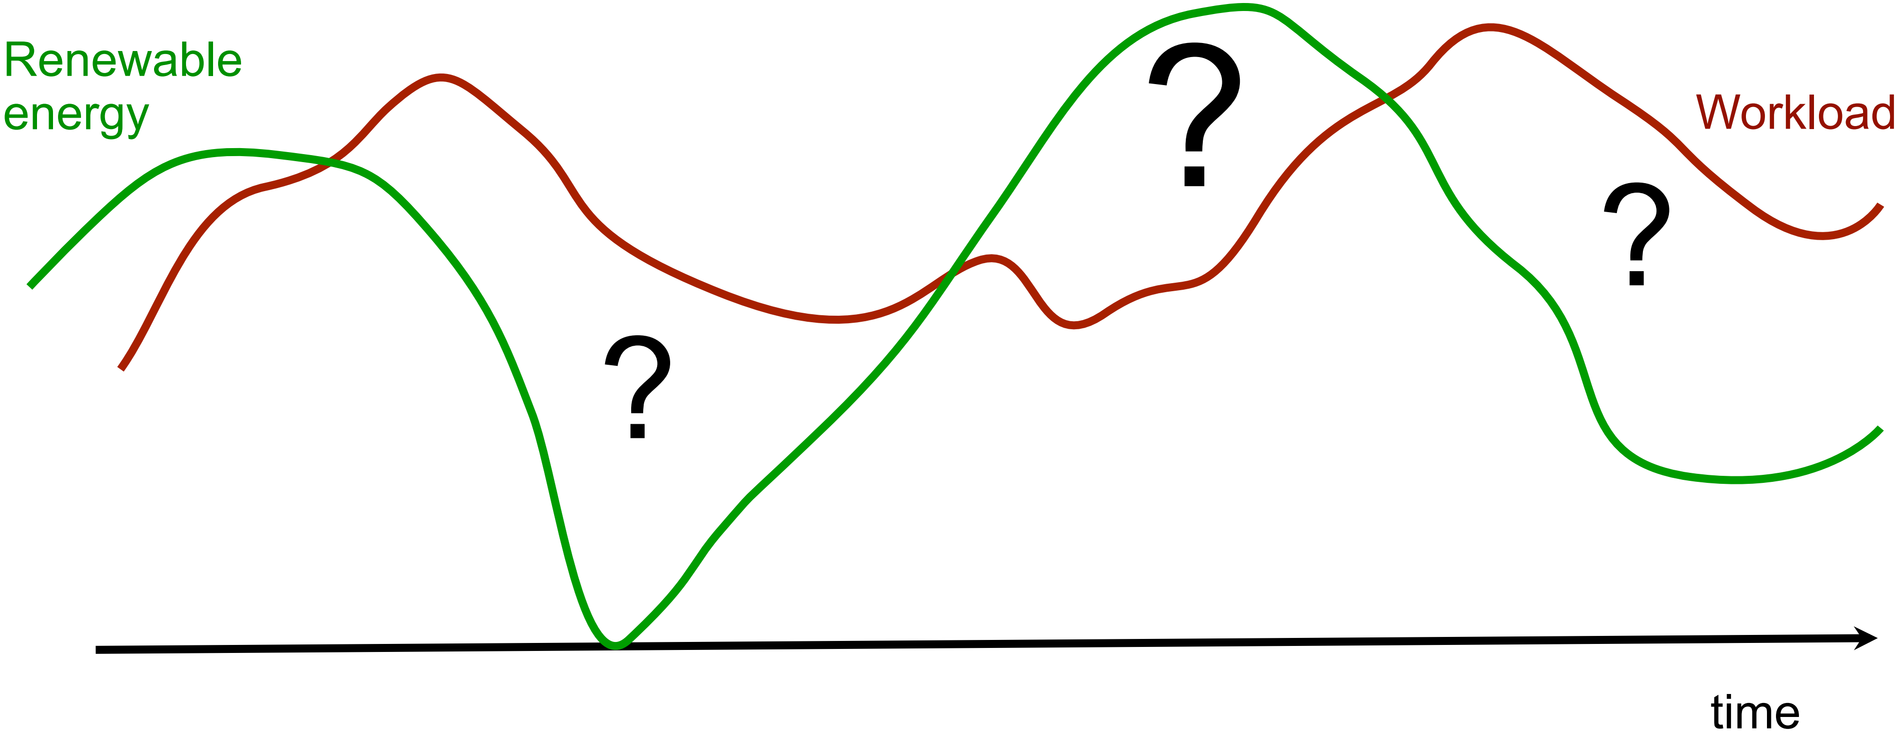


Courtesy of J-M Menaud

# Demand-side management

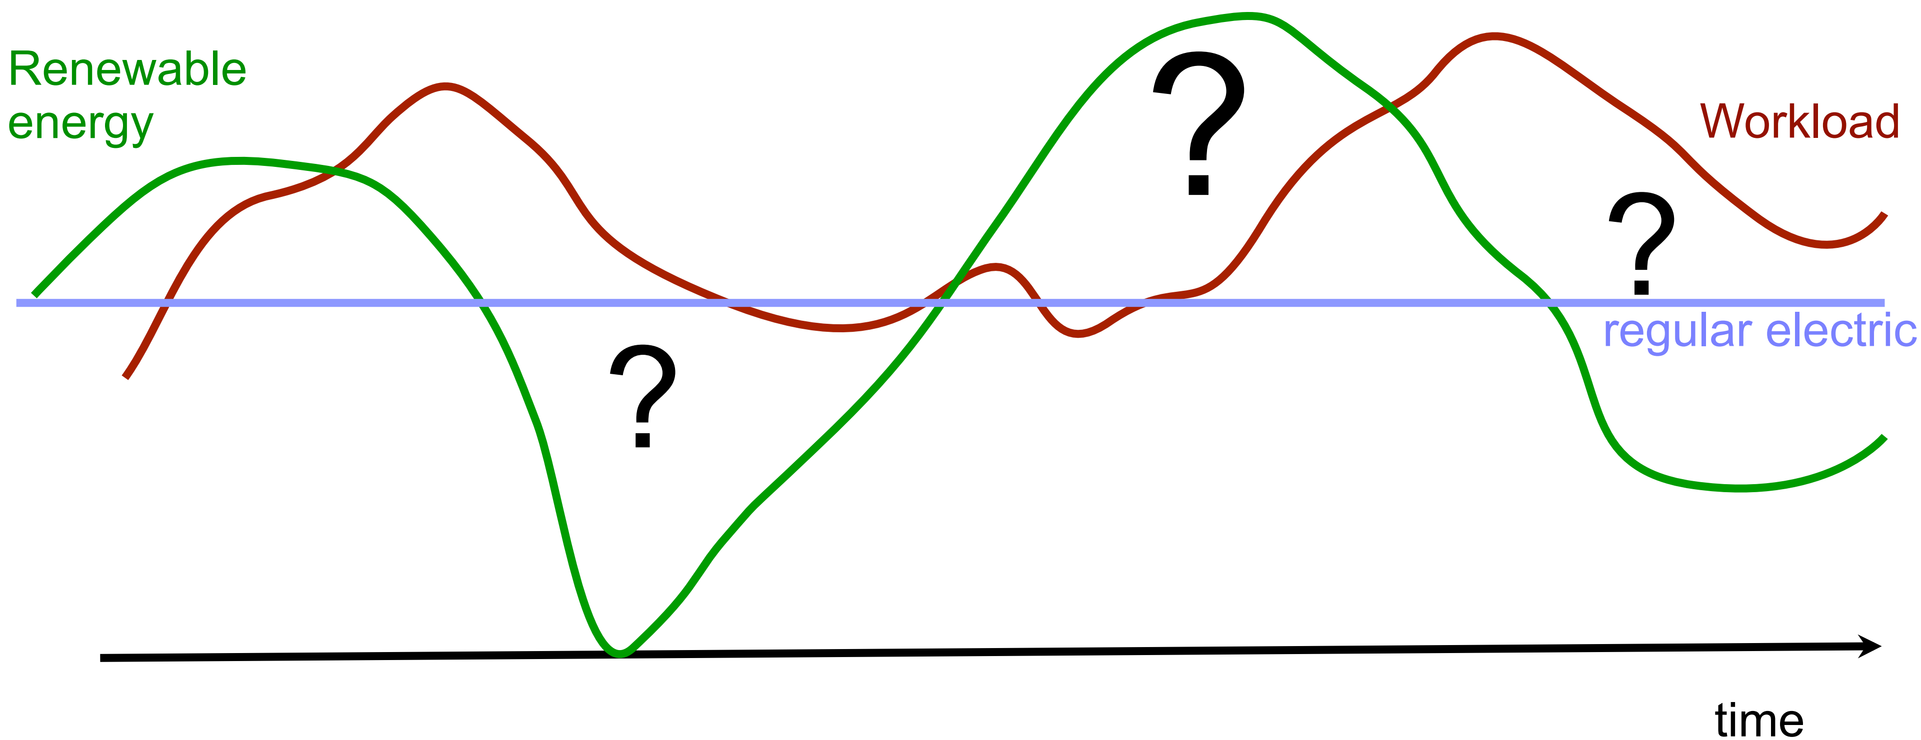


Courtesy of J-M Menaud

# Energy storage

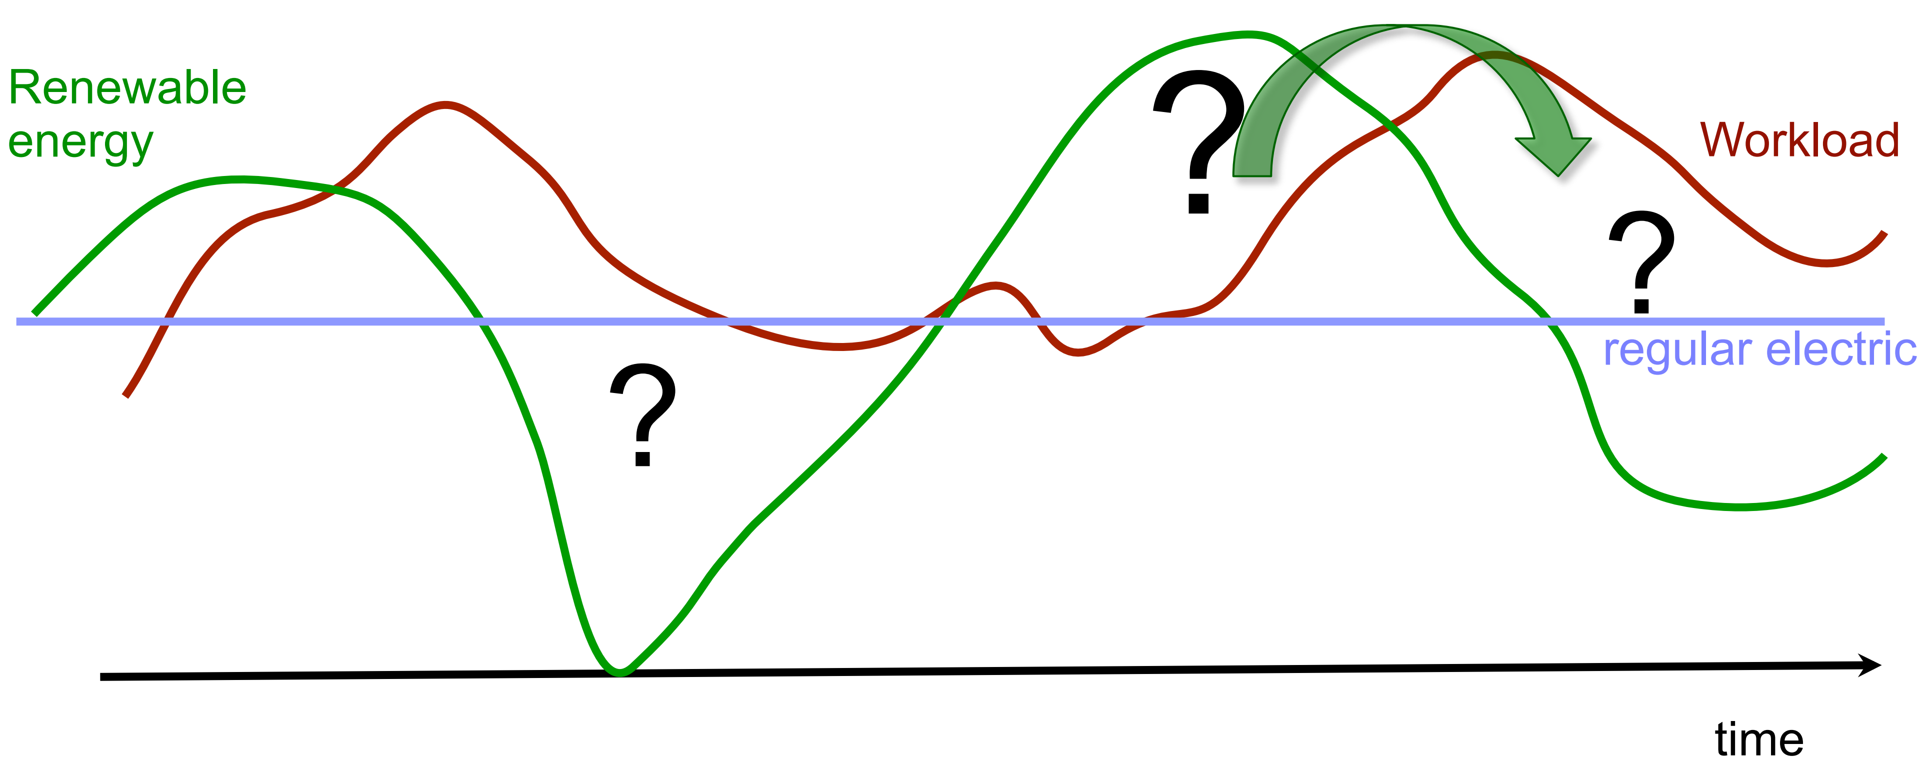


# Edge Model

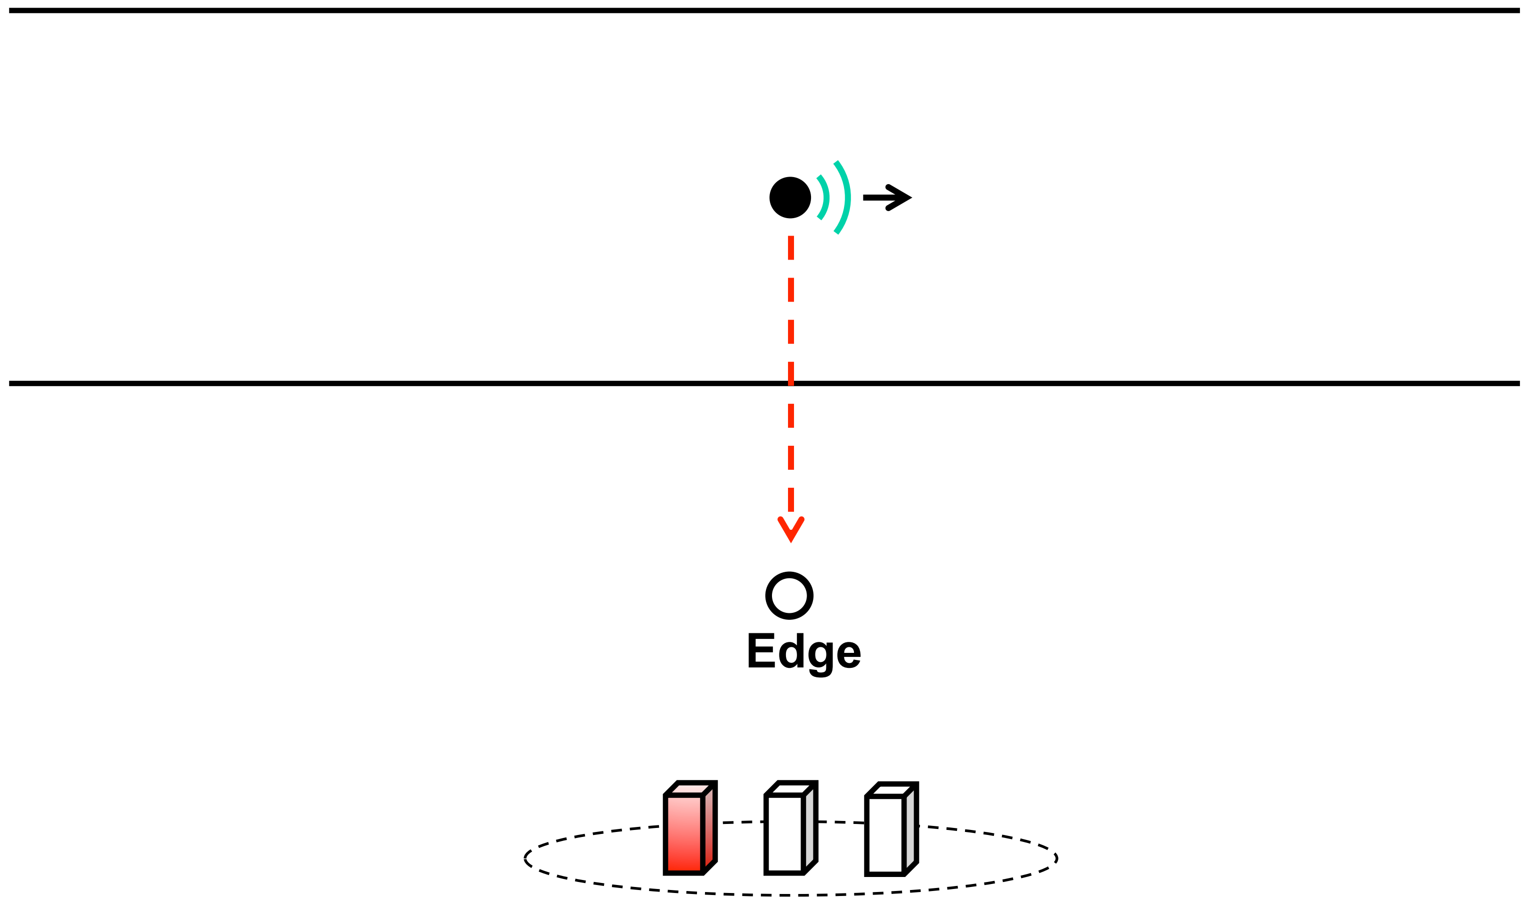


Courtesy of Yunbo Li

# Edge Model

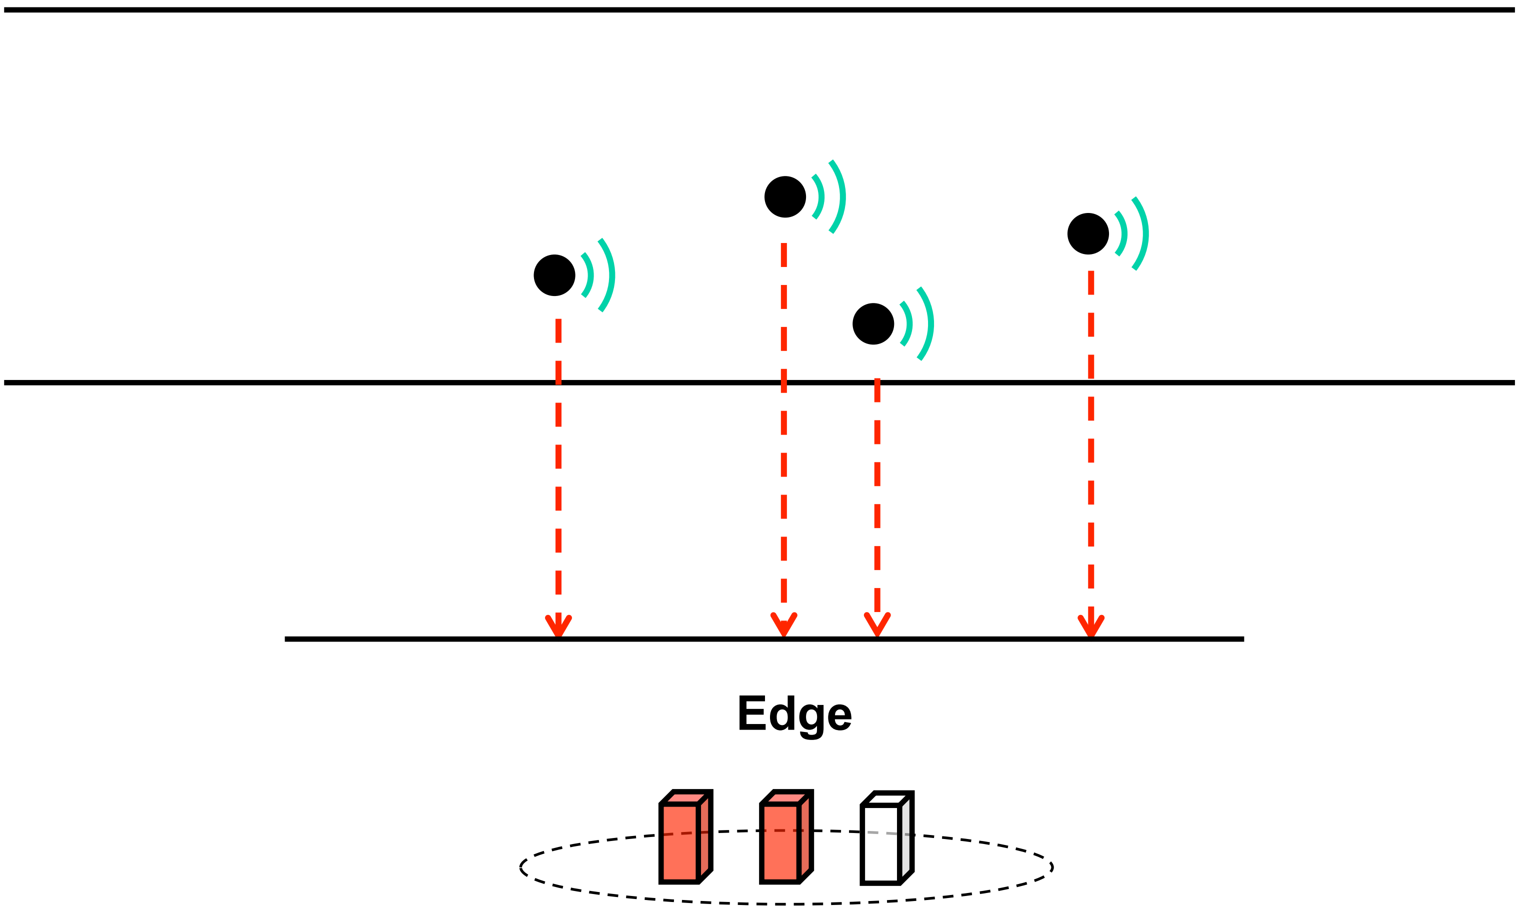


Courtesy of Yunbo Li

# Edge Model

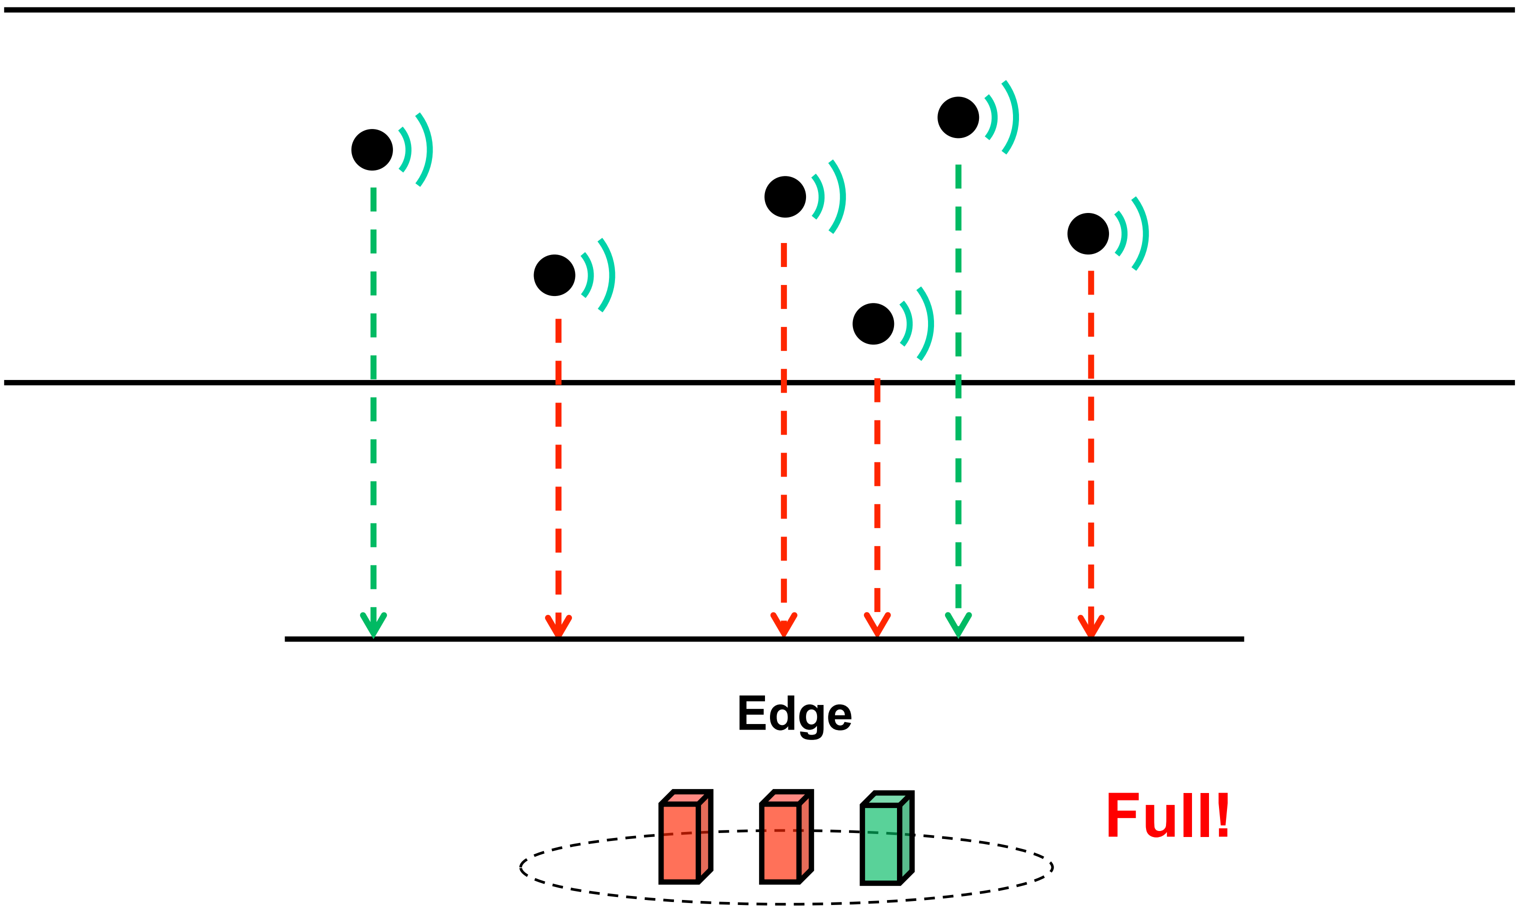


Courtesy of Yunbo Li

# Edge-Core Model

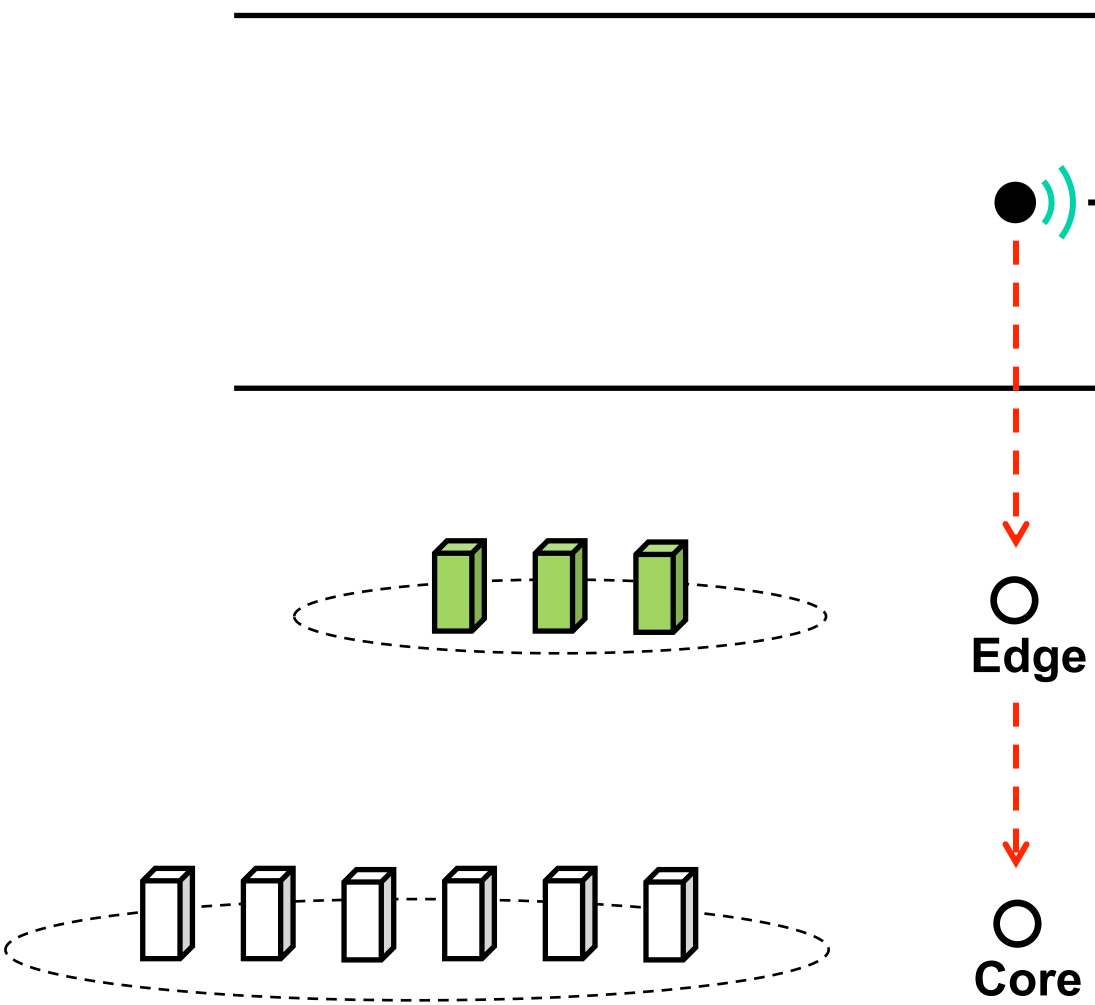


+ Low latency
+ On-site renewable energy
– Limited computational capacity

+ Huge computational resources
– Higher latency and network overhead

# 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?***

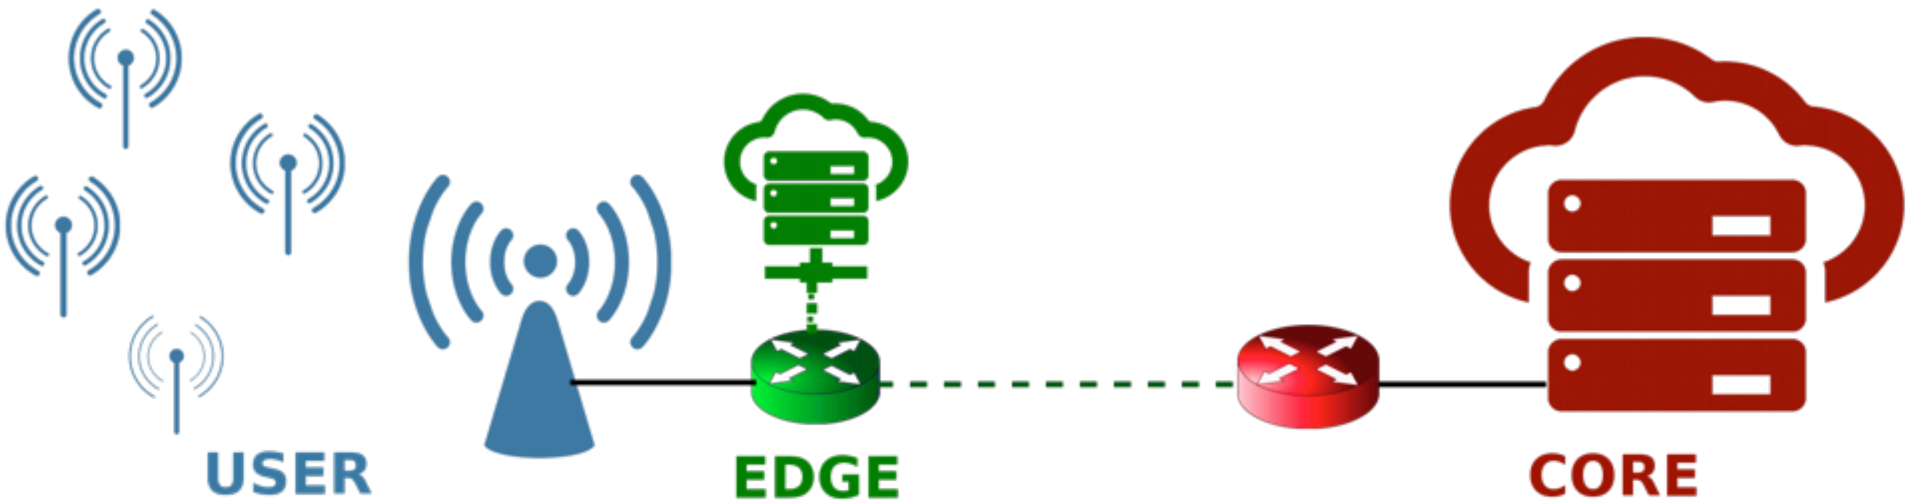


*"Leveraging Renewable Energy in Edge Clouds for Data Stream Analysis in IoT",* Y. Li, A.-C. Orgerie, I. Rodero, M. Parashar and J.-M. Menaud, p 186-195, IEEE/ACM CCGrid 2017.

# 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**

# Model

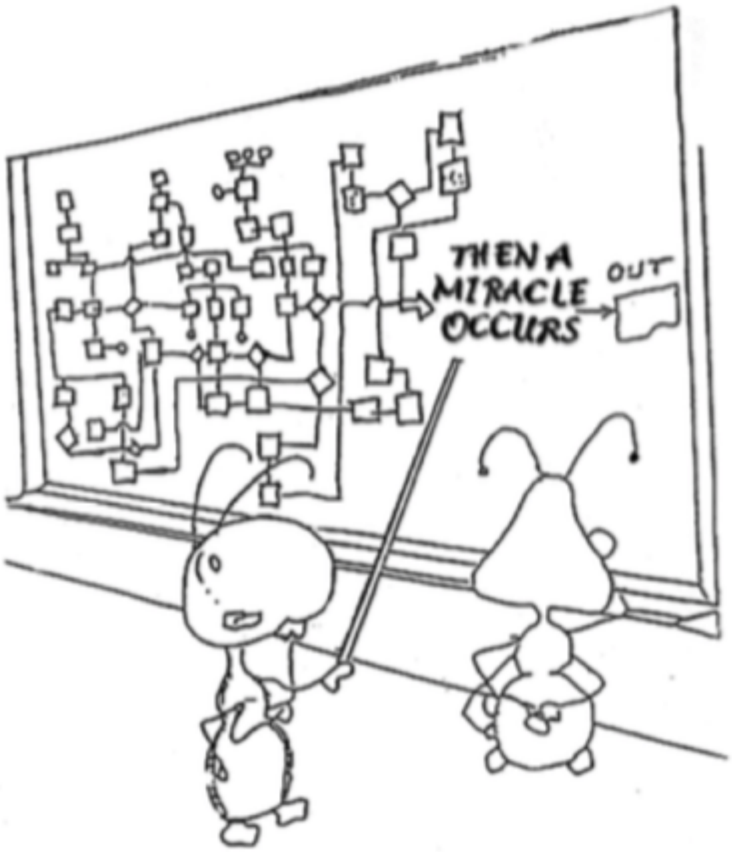


**… later!**

# Scenario

# Application-driven approach

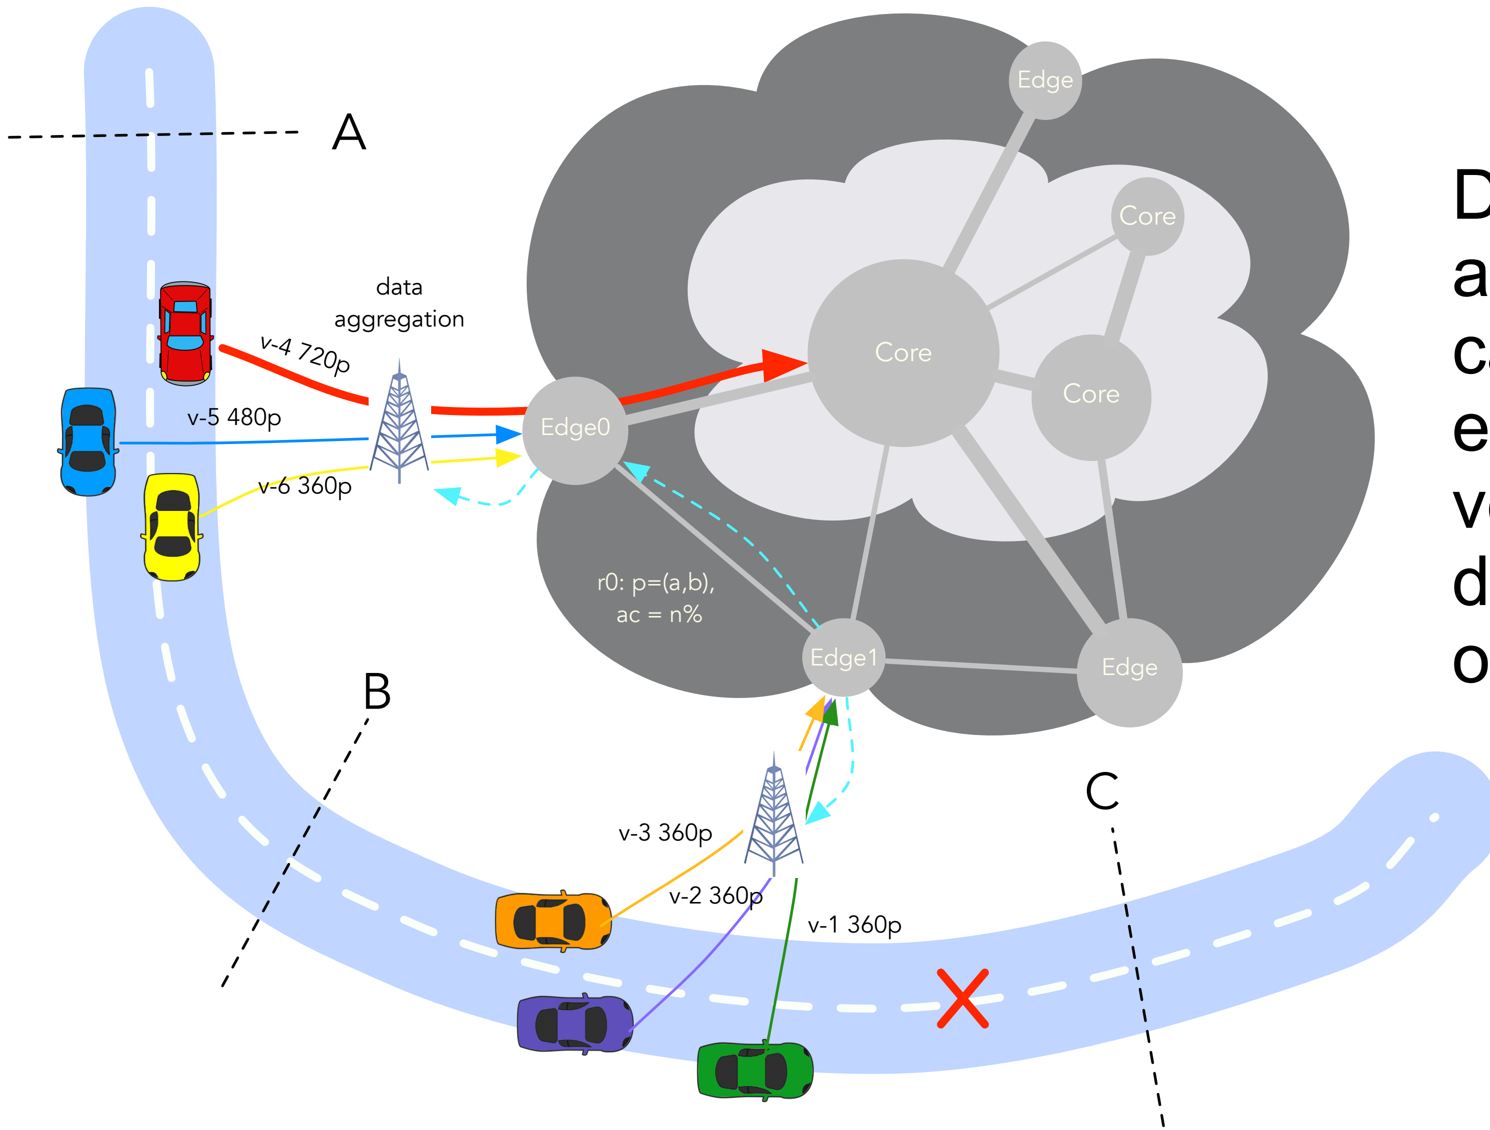


Data stream analysis from cameras embedded on vehicles to detect objects on the road

# 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 |

# Wattmeters

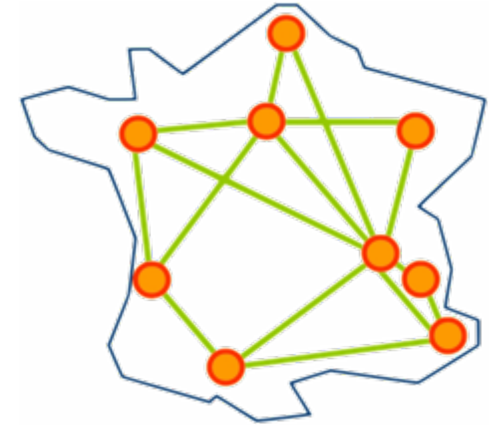


"*The Green Grid'5000: Instrumenting a Grid with Energy Sensors*", M. Dias de Assunção, J.-P. Gelas, L. Lefèvre and A.-C. Orgerie, INGRID 2010.

# Servers' power profile

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

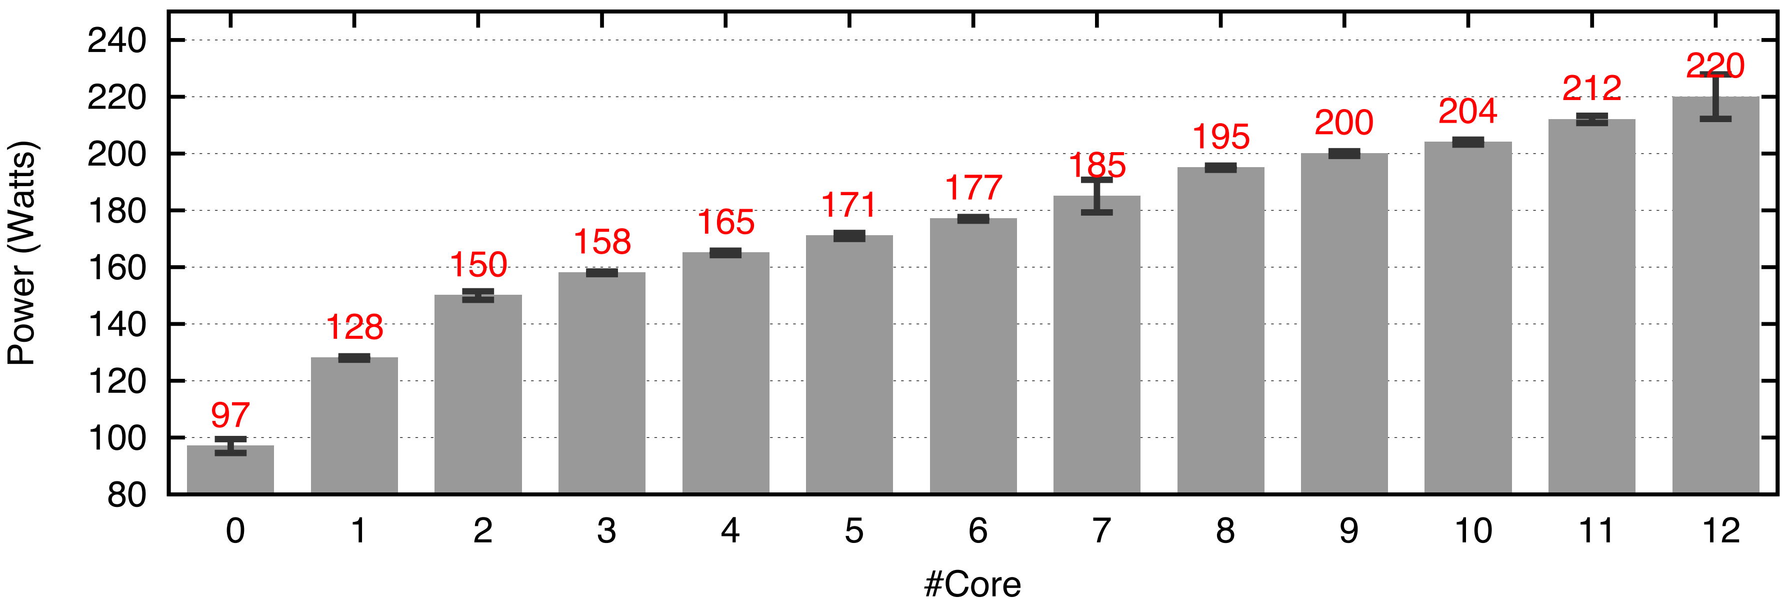


"*Opportunistic Scheduling in Clouds Partially Powered by Green Energy*", Y. Li, A.-C. Orgerie and J.-M. Menaud, IEEE GreenCom 2015.

# Experimental methodology

1. Application benchmarking on real infrastructure

2. Extrapolated results based on simulation

Using:

- Servers monitored by wattmeters

- Photovoltaic panel production traces

# Experiments

# Analysis time on different VM sizes

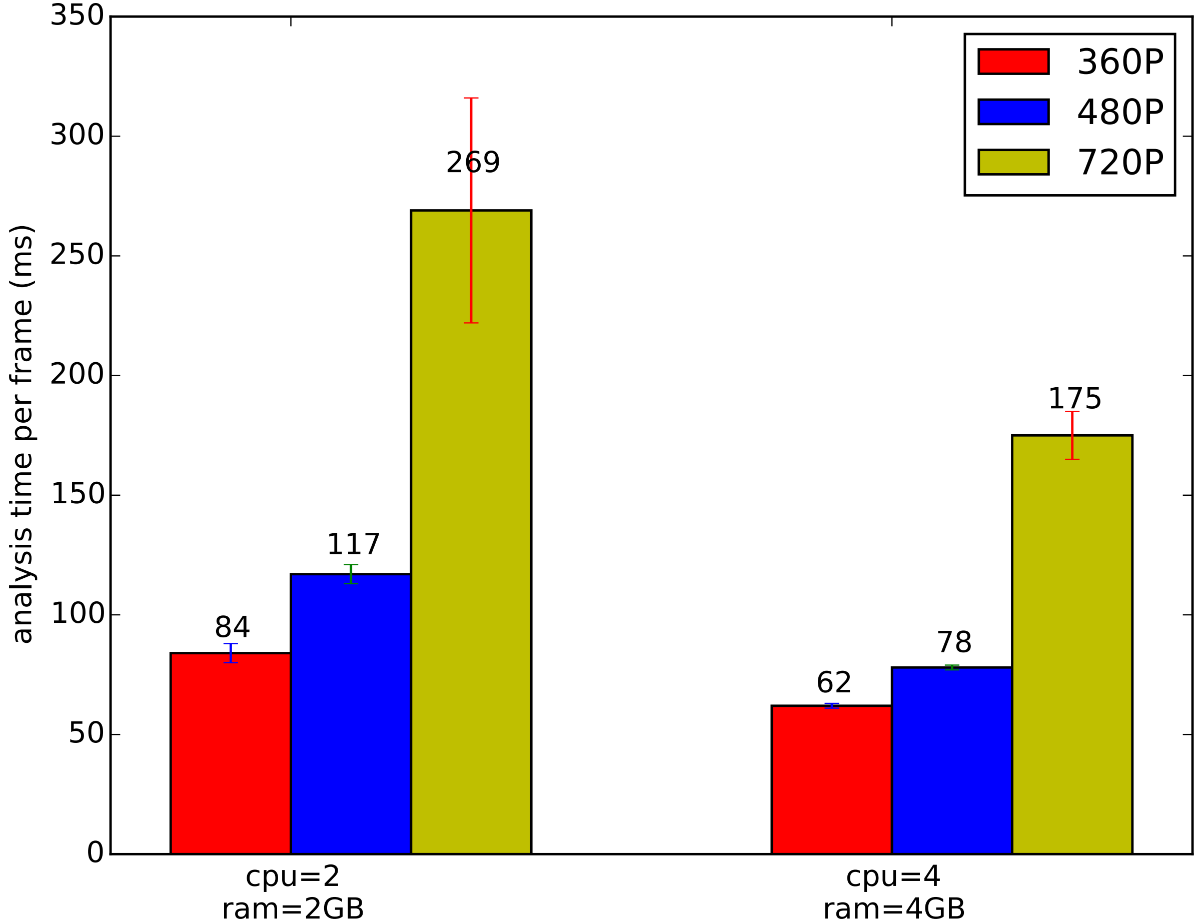


Small VM | Medium VM

Medium VM better, but not linear scalability

*Depends on applications' elasticity*

Real measurements based on 10 runs for each experiment.

# Service configuration

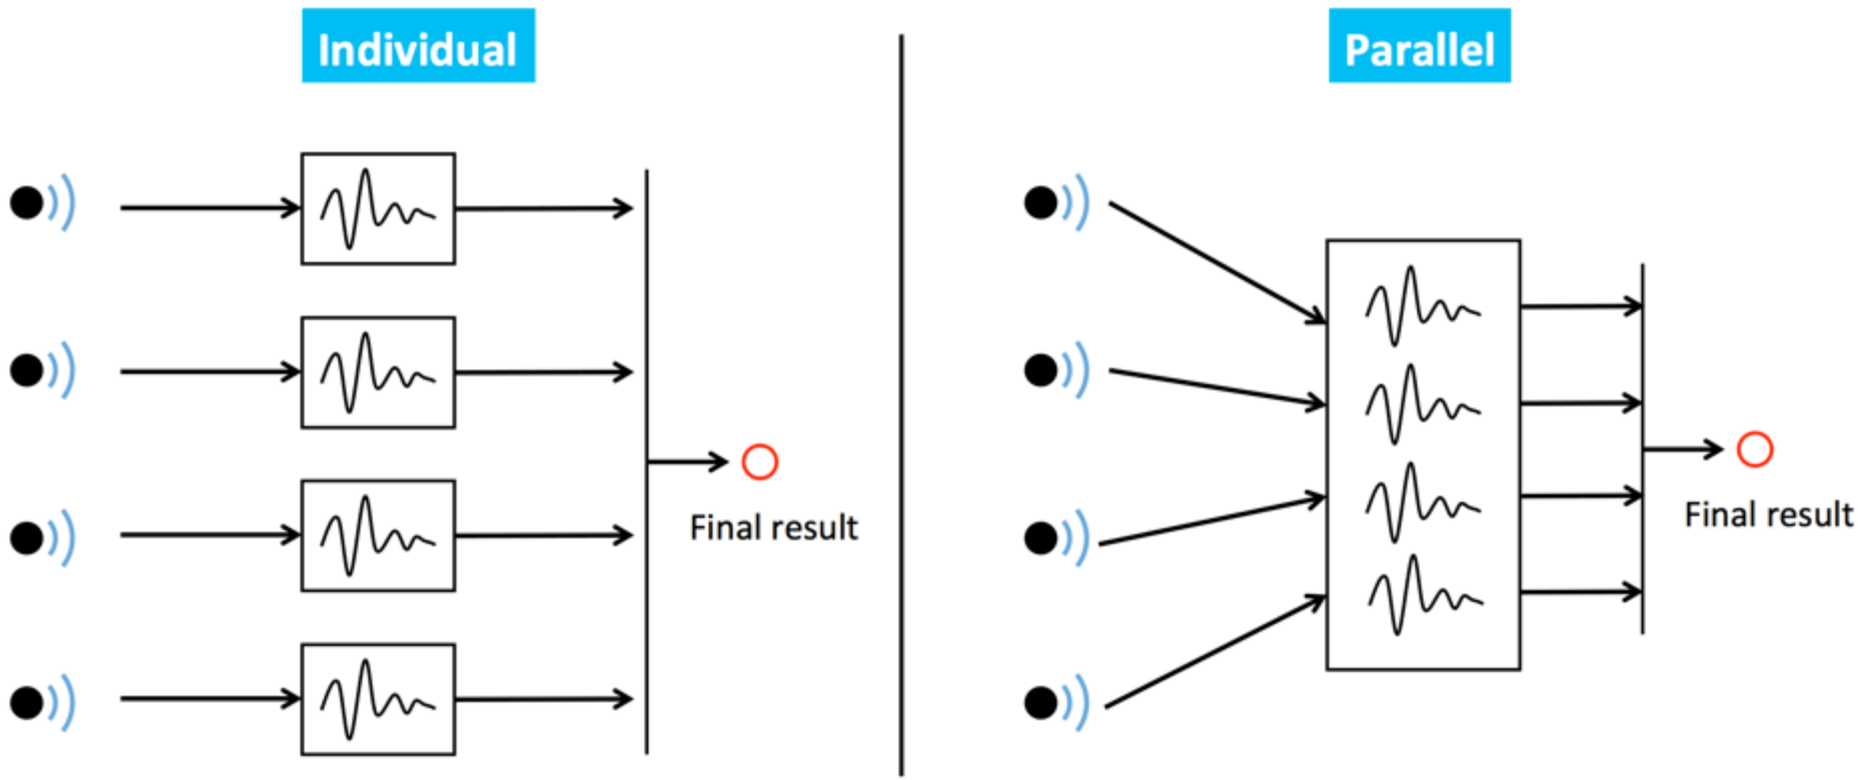


One VM per stream

A VM for several streams

Courtesy of Yunbo Li

# Several small VMs vs. one large VM

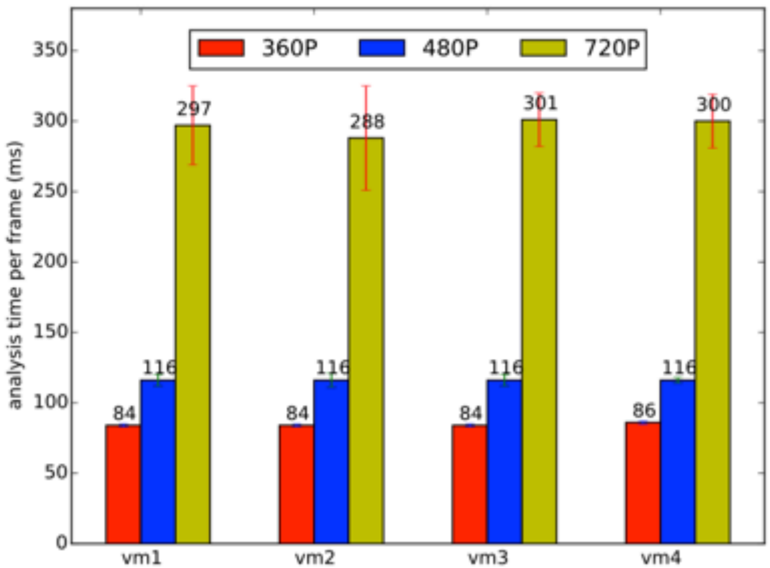


(a) Analysis time for 4 identical VMs with 1 data stream each on the same PM

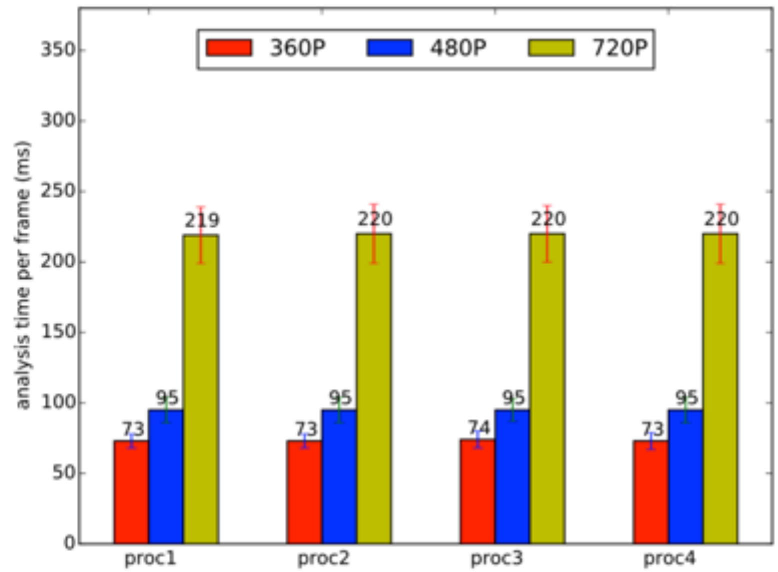


(b) Analysis time for each of the 4 data stream processes in a large VM

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

*Depends on application's resource usage*

# Power and energy consumption

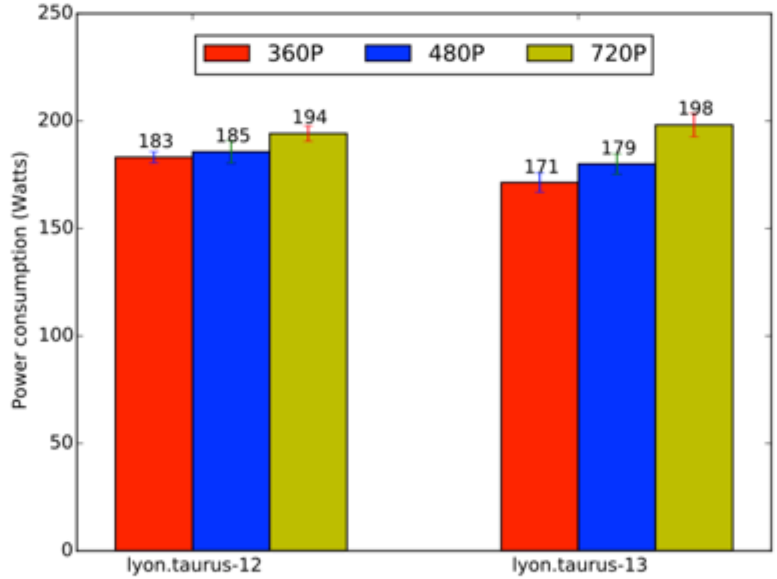


(c) Power consumption for 4 small VMs on Taurus-12 and 1 large VM on Taurus-13 for the same amount of computation

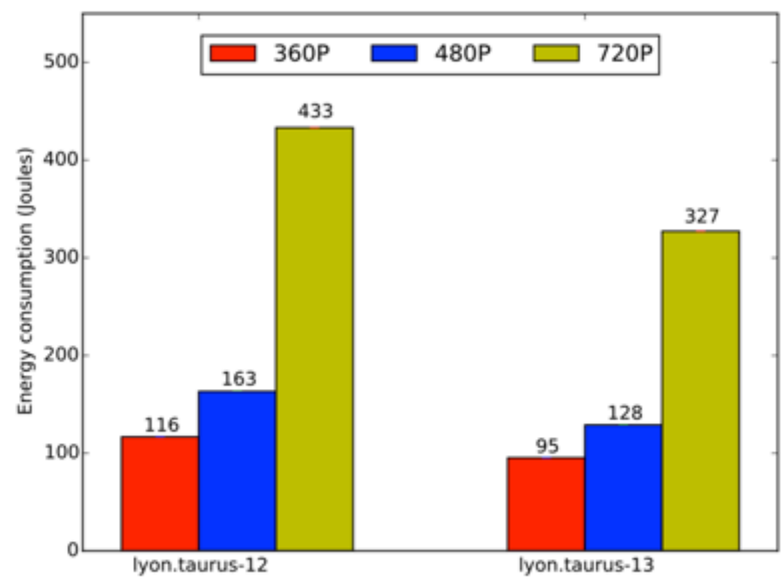


(d) Energy consumption for analyzing a 5 mn video on Taurus-12 with 4 small VMs and on Taurus-13 with 1 large VM

Power consumptions almost equivalent
Better energy consumption with large VM

# Consolidation within a single VM

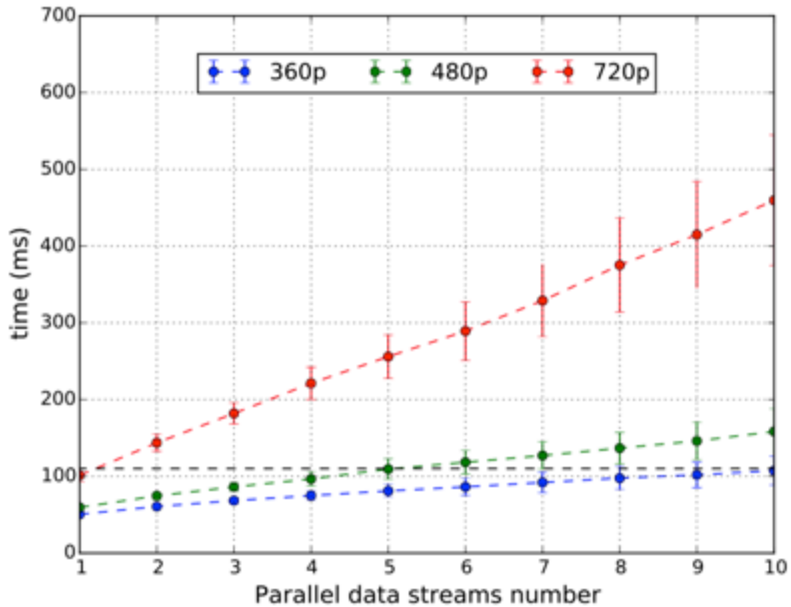


(e) Analysis time with parallel data streams in a large VM

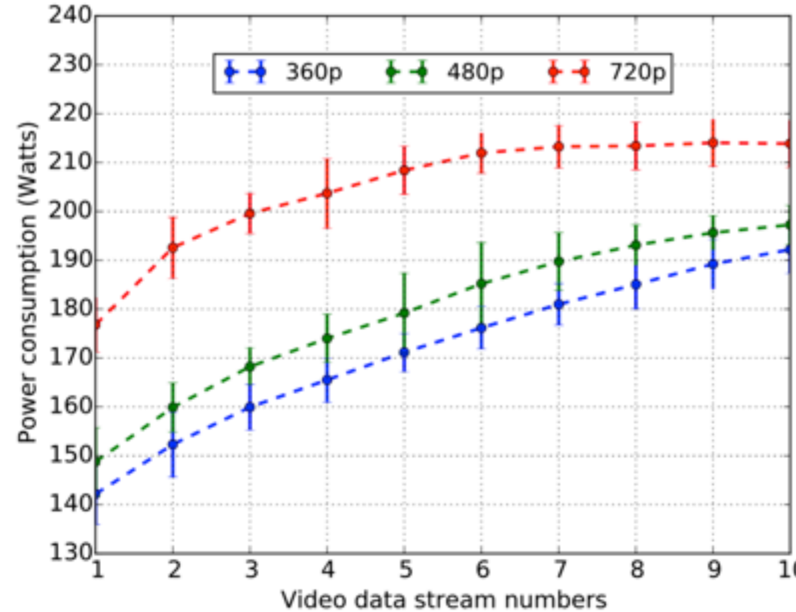


(f) Power consumption with parallel data streams in a large VM

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

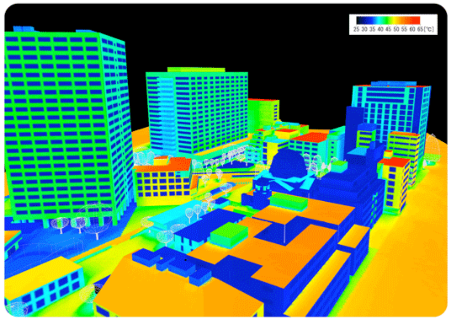

# Cloud configurations

- **Core cloud**
  - 100 servers
  - 100 ms latency with the edge devices
- **Edge cloud**
  - 5 servers
  - Equipped with on-site photovoltaic panels: 17 $m^2$
  - Equipped with batteries: 10 kWh (~130 L for LA)
  - Connected to the regular electric grid
- Unused resources are switched off.

# Energy storage model

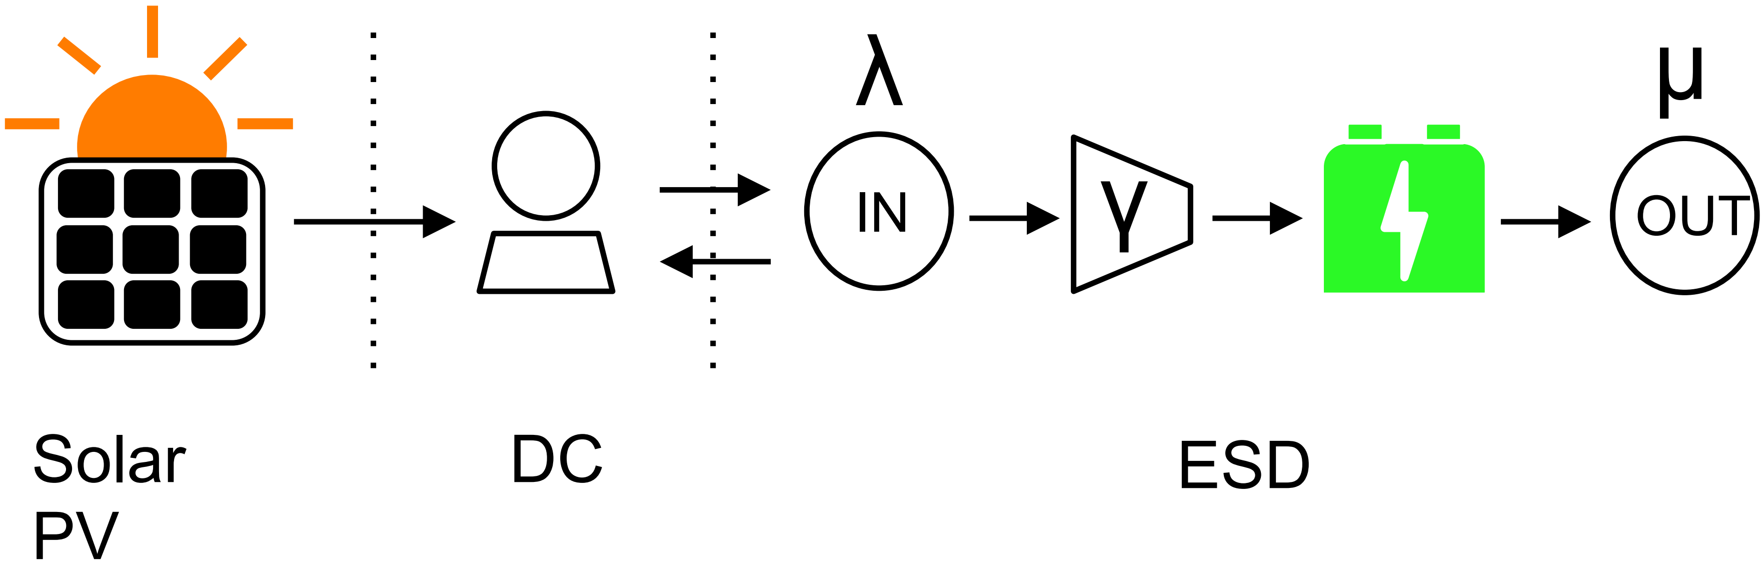


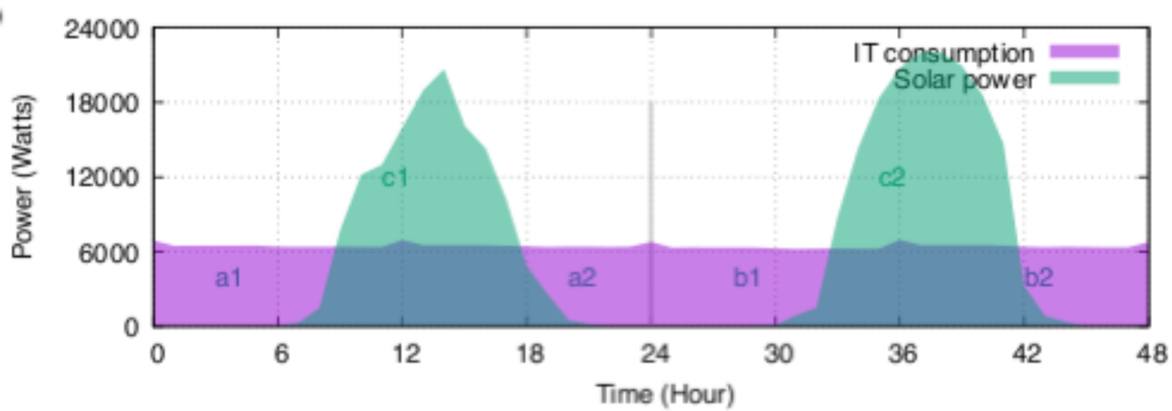


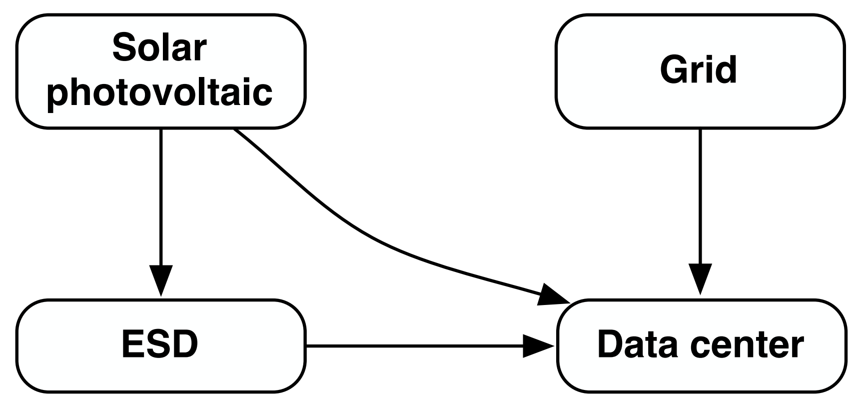


**"***Balancing the use of batteries and opportunistic scheduling policies for maximizing renewable energy consumption in a Cloud data center***",** Y. Li, A.-C. Orgerie and J.-M. Menaud, PDP 2017.

# Energy production: photovoltaic panel real traces

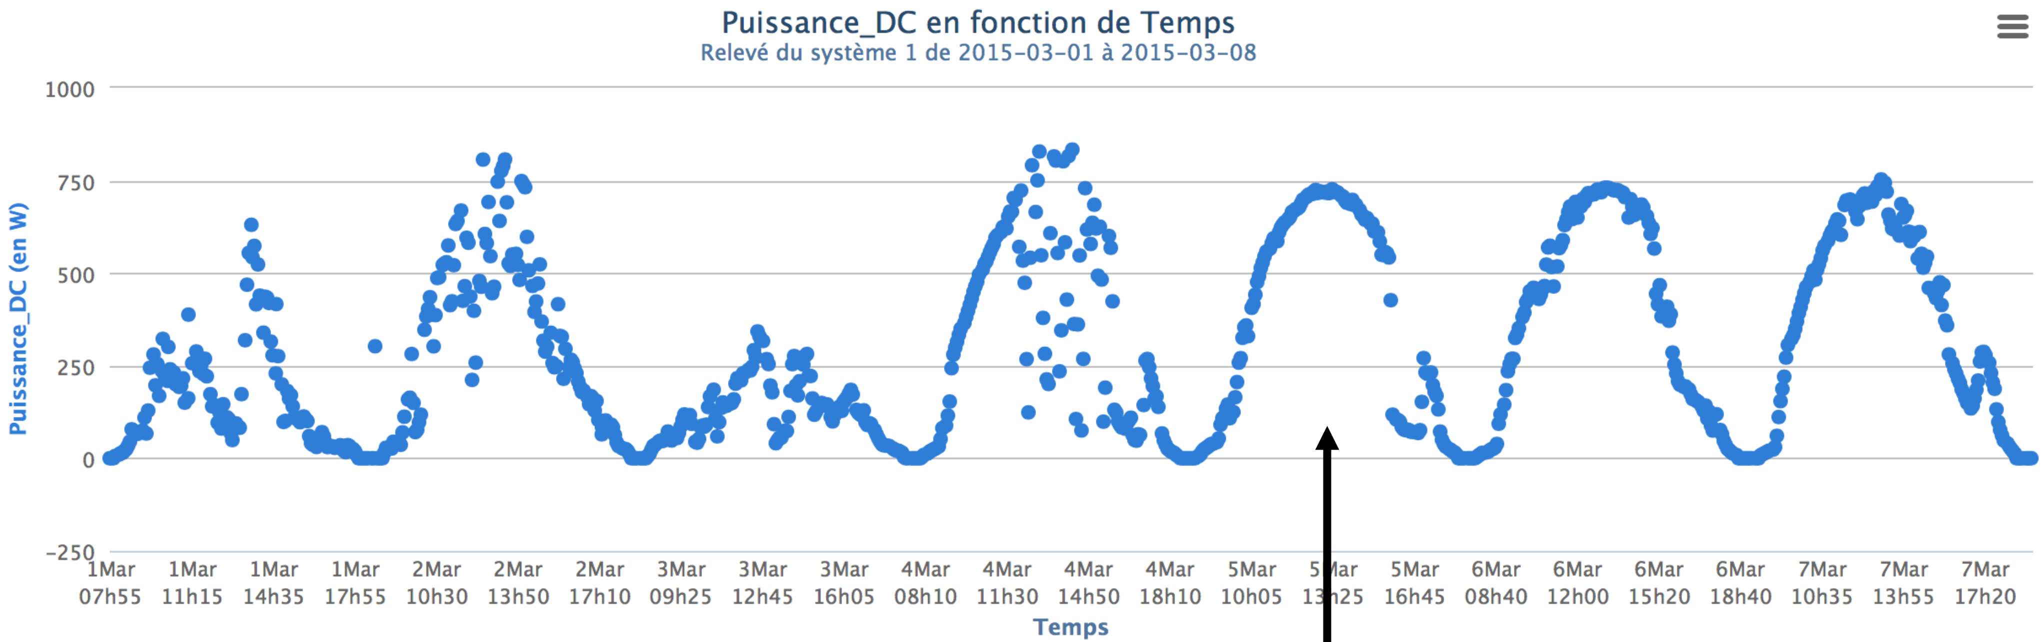


cloudless sunny day


# Energy consumption at edge & core

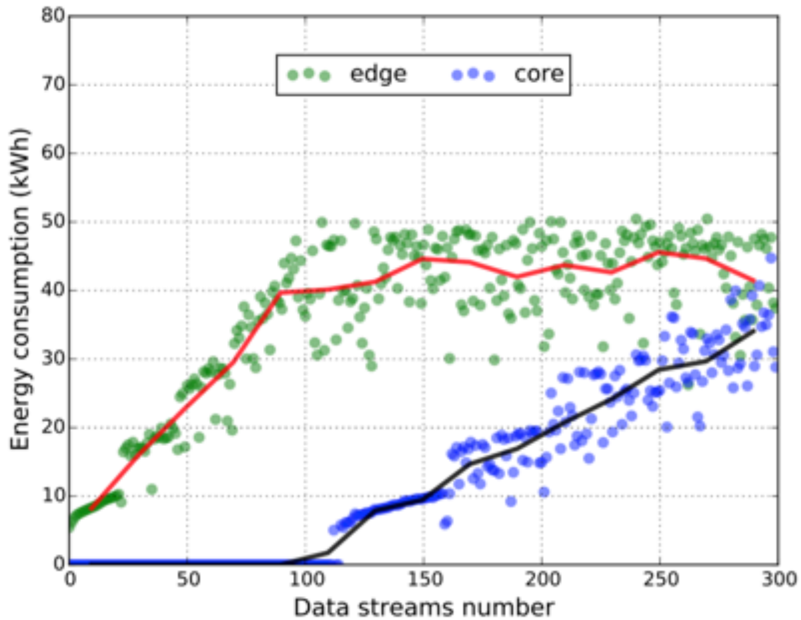

(a) Energy consumption with resolution 360p

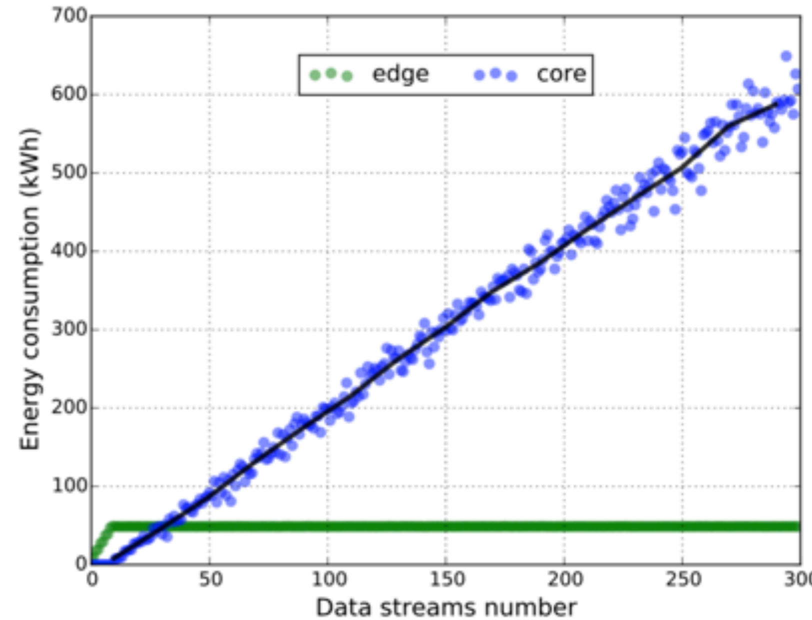

(b) Energy consumption with resolution 720p

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

*Depends on servers' architecture.*

# Brown energy consumption

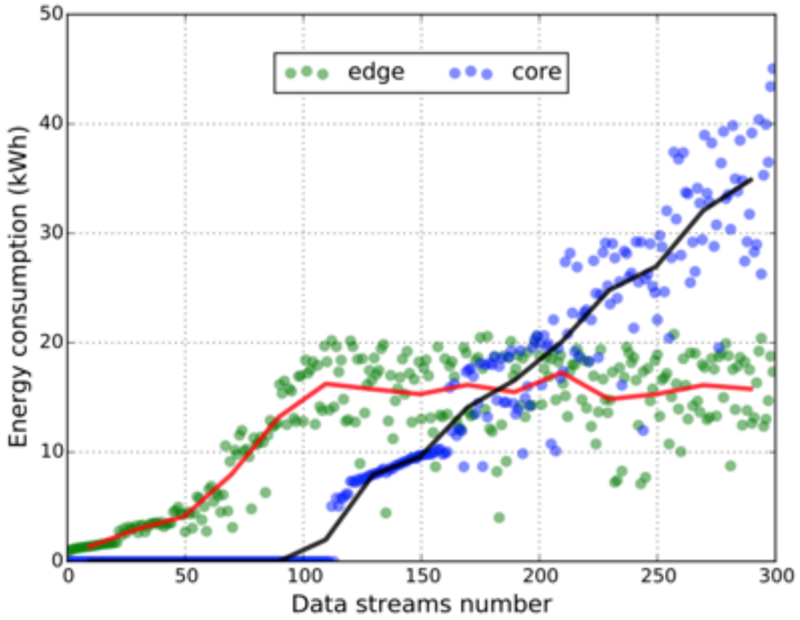


(c) Brown energy consumption with resolution 360p

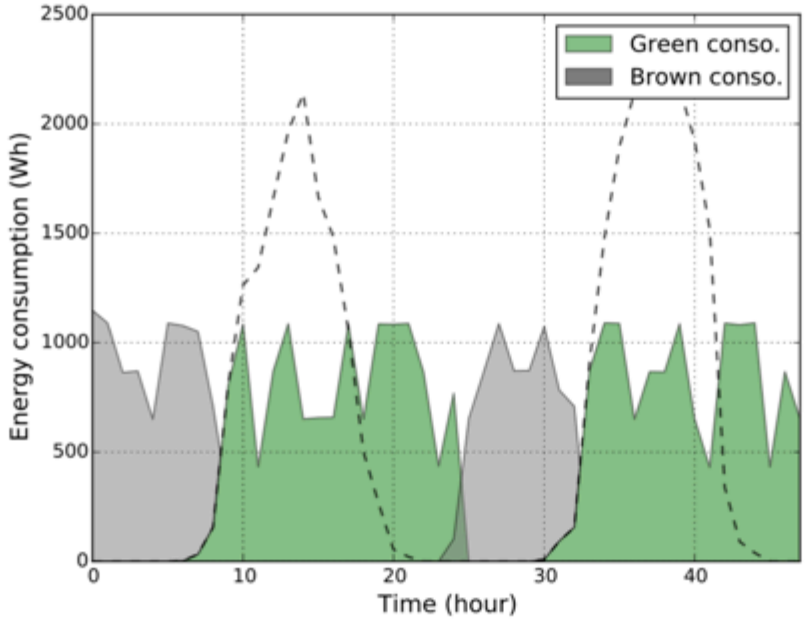


(d) 2 days of energy consumption at edge

Electricity consumption from regular grid is reduced by half with the solar panels and batteries at the edge.

# Average application latency

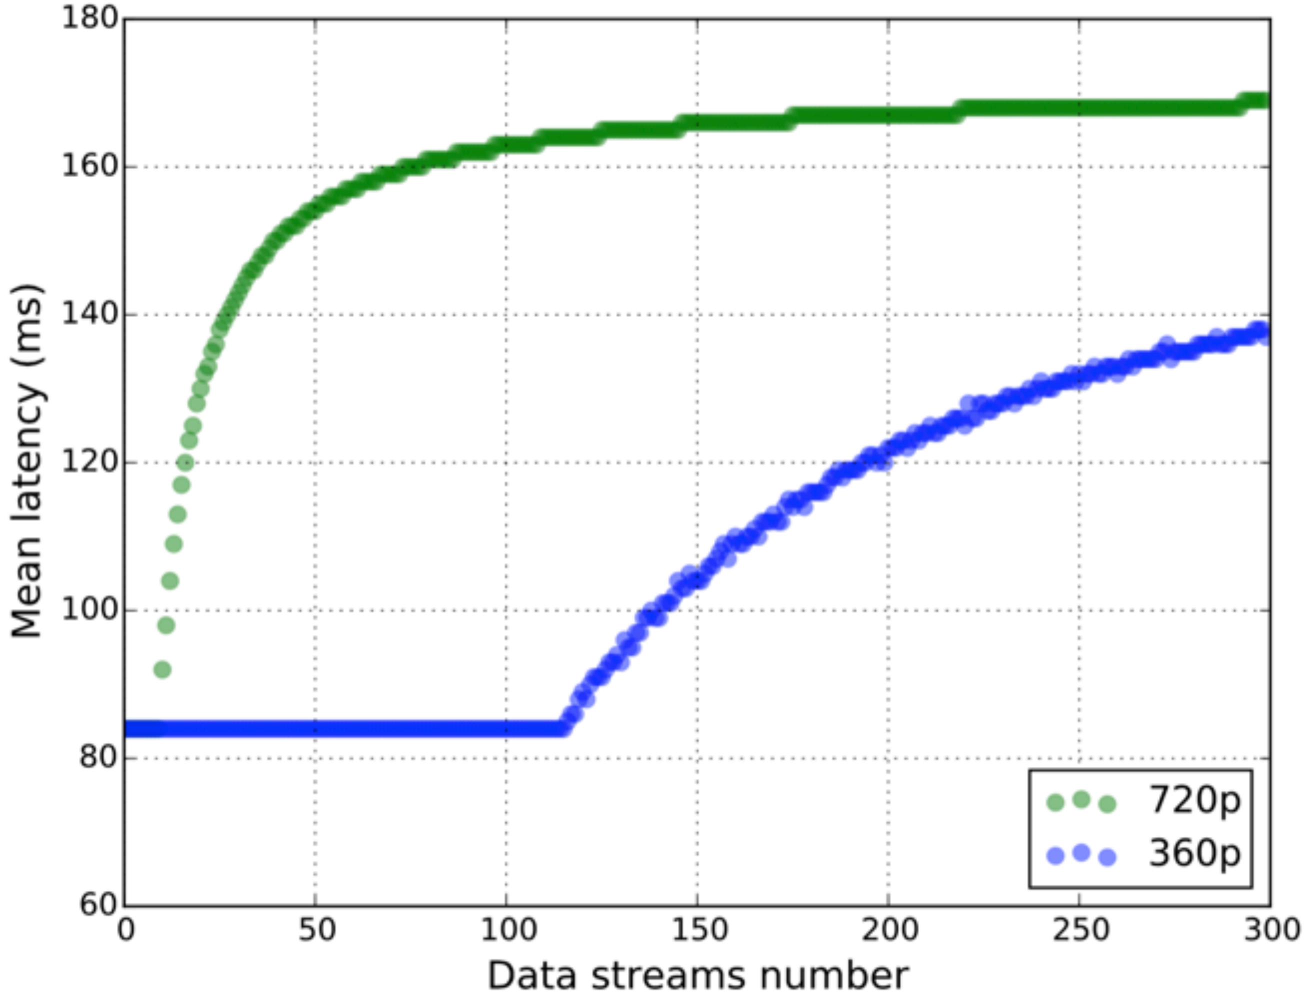


*Depends on edge's resources availability*

# Network energy consumption

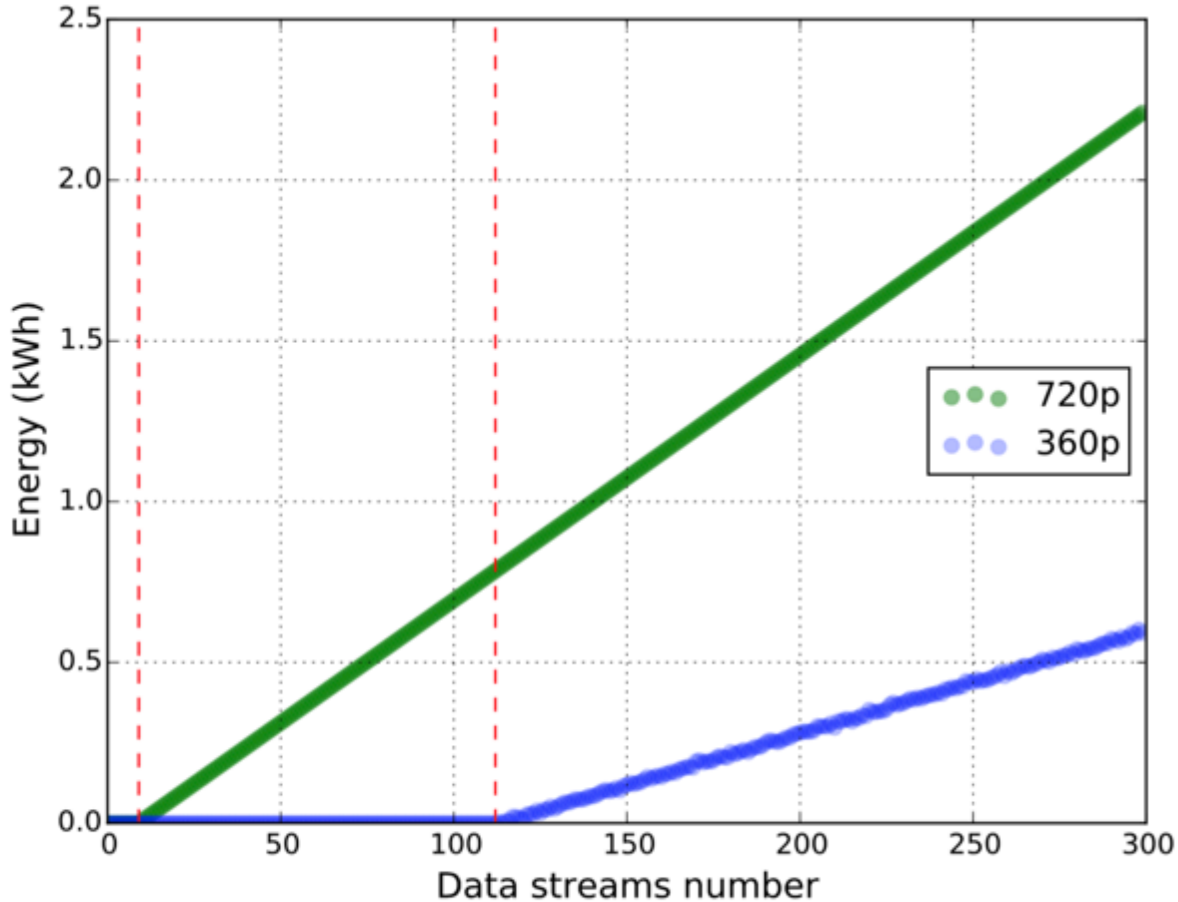


Cost per-bit energy model for network

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

| 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

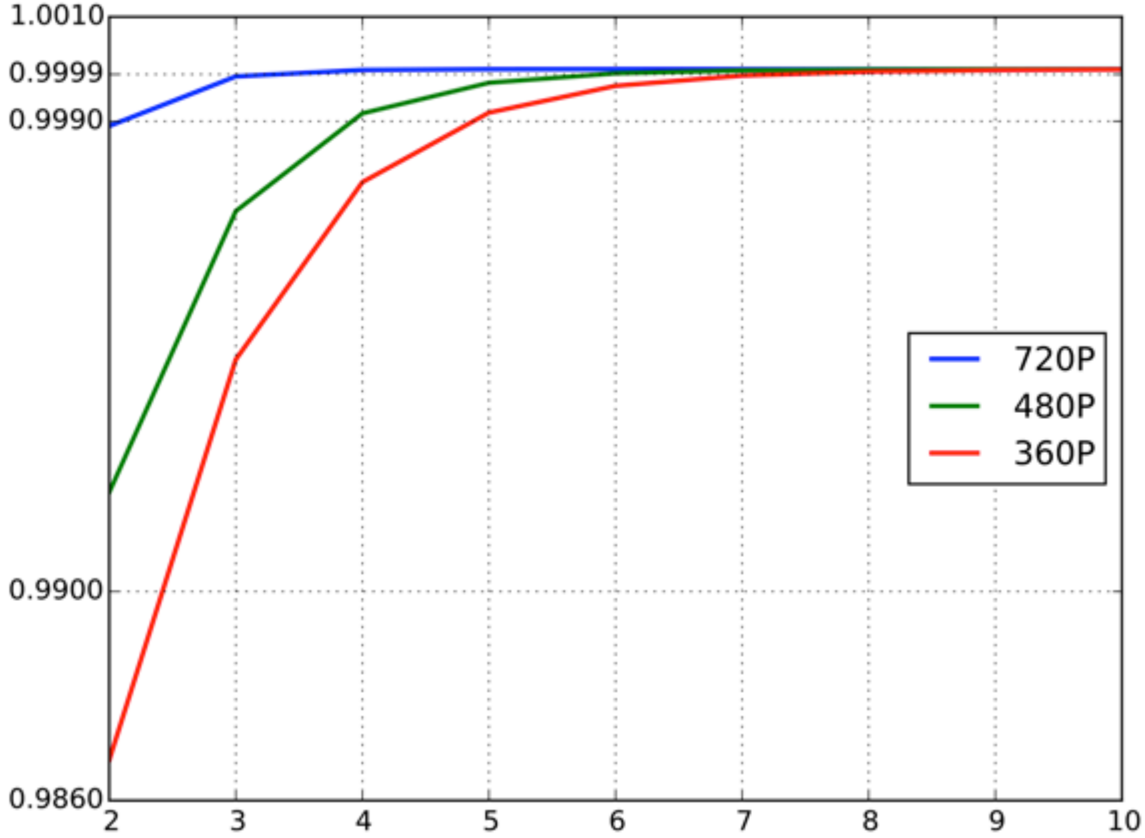


Not feasible in real-time with 1 VM

Feasible in real-time with 1 VM

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

Number of data streams for reaching required accuracy

# Summary

Offloading the data to process video streams at edge:

- Effectively reduces the response time
- Avoids unnecessary data transmission between edge and core
- Extends for instance the battery lifetime of end-user equipment
- On-site renewable energy production and batteries in our scenario can save up to 50% total consumed energy consumed at the edge

# What about the other parts?

## *Which part consumes the most?*

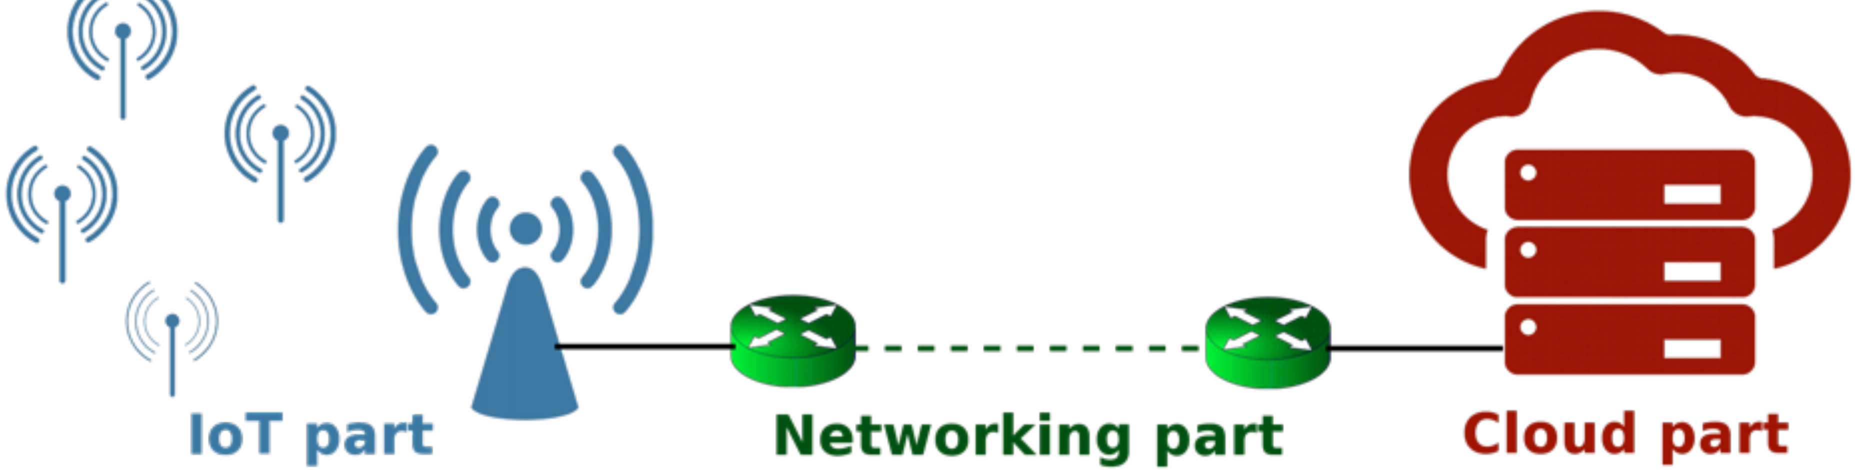


"End-to-end Energy Models for Edge Cloud-based IoT Platforms: Application to Data Stream Analysis in IoT", Y. Li, A.-C. Orgerie, I. Rodero, B. Lemma Amersho, M. Parashar, J.-M. Menaud, FGCS, vol. 87, p 667-678, 2018.

# Parameters of our example

**IoT devices (camera)**

| 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 |

**Network devices**

| 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 |

**Cloud data centers**

PUE = 1.7 for edge

PUE = 1.2 for core

# Experimental setup

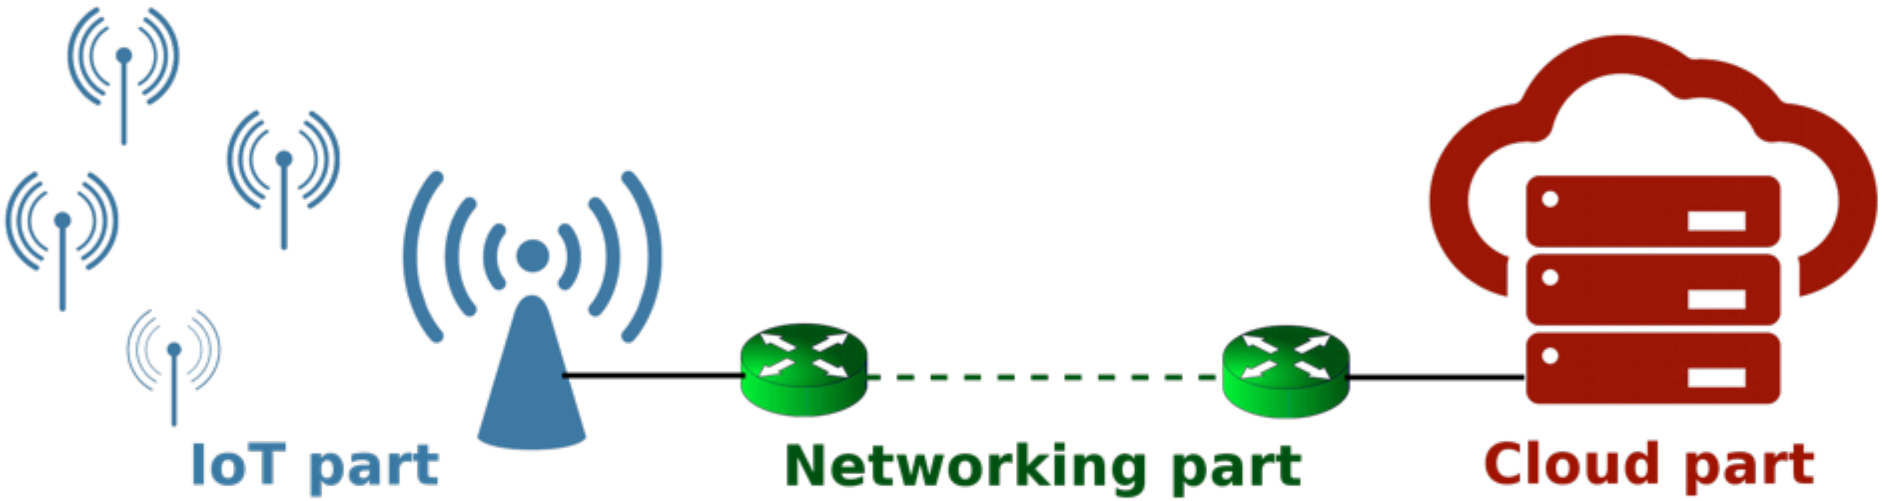


*Simulations*

ns-3

*Model from literature*

*Real measurements*

# IoT consumption per device

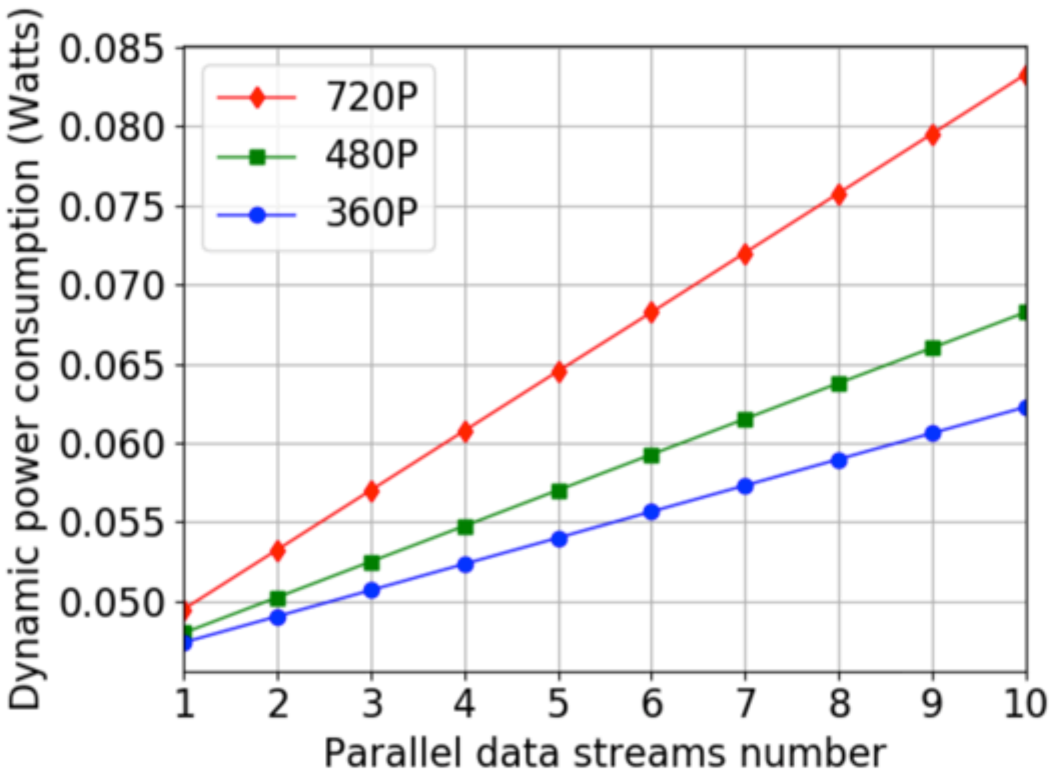


Dynamic power consumption

Overall 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 |

# IoT part including access point

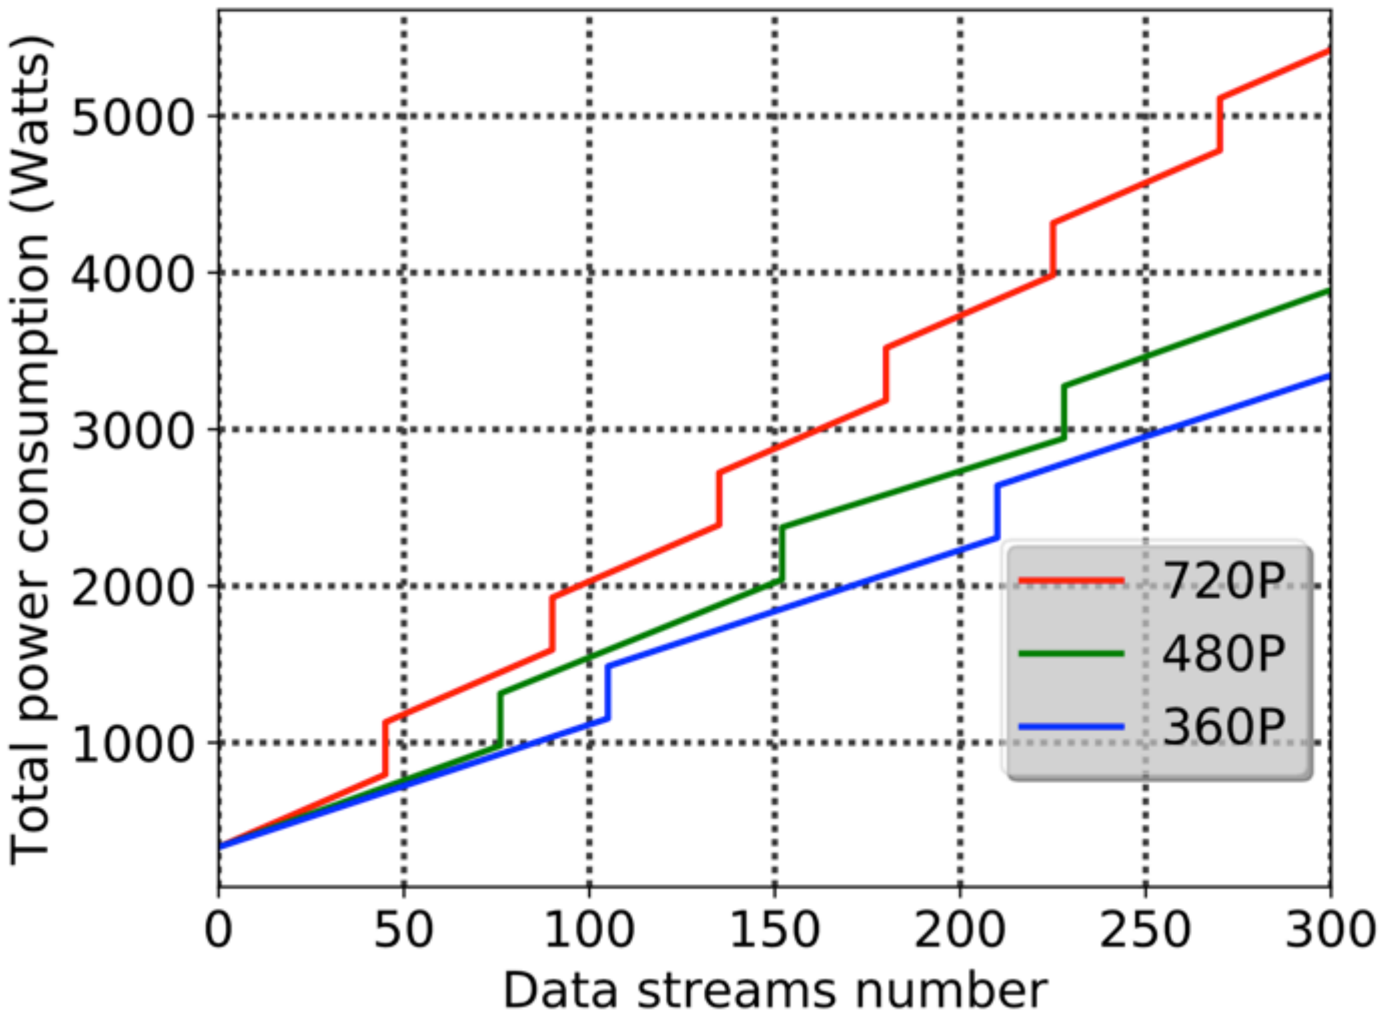

# Overall evaluation

| 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 (probably underestimated)
- **Cloud part**: measured, accurate on the given servers

# Moving jobs between data centers?

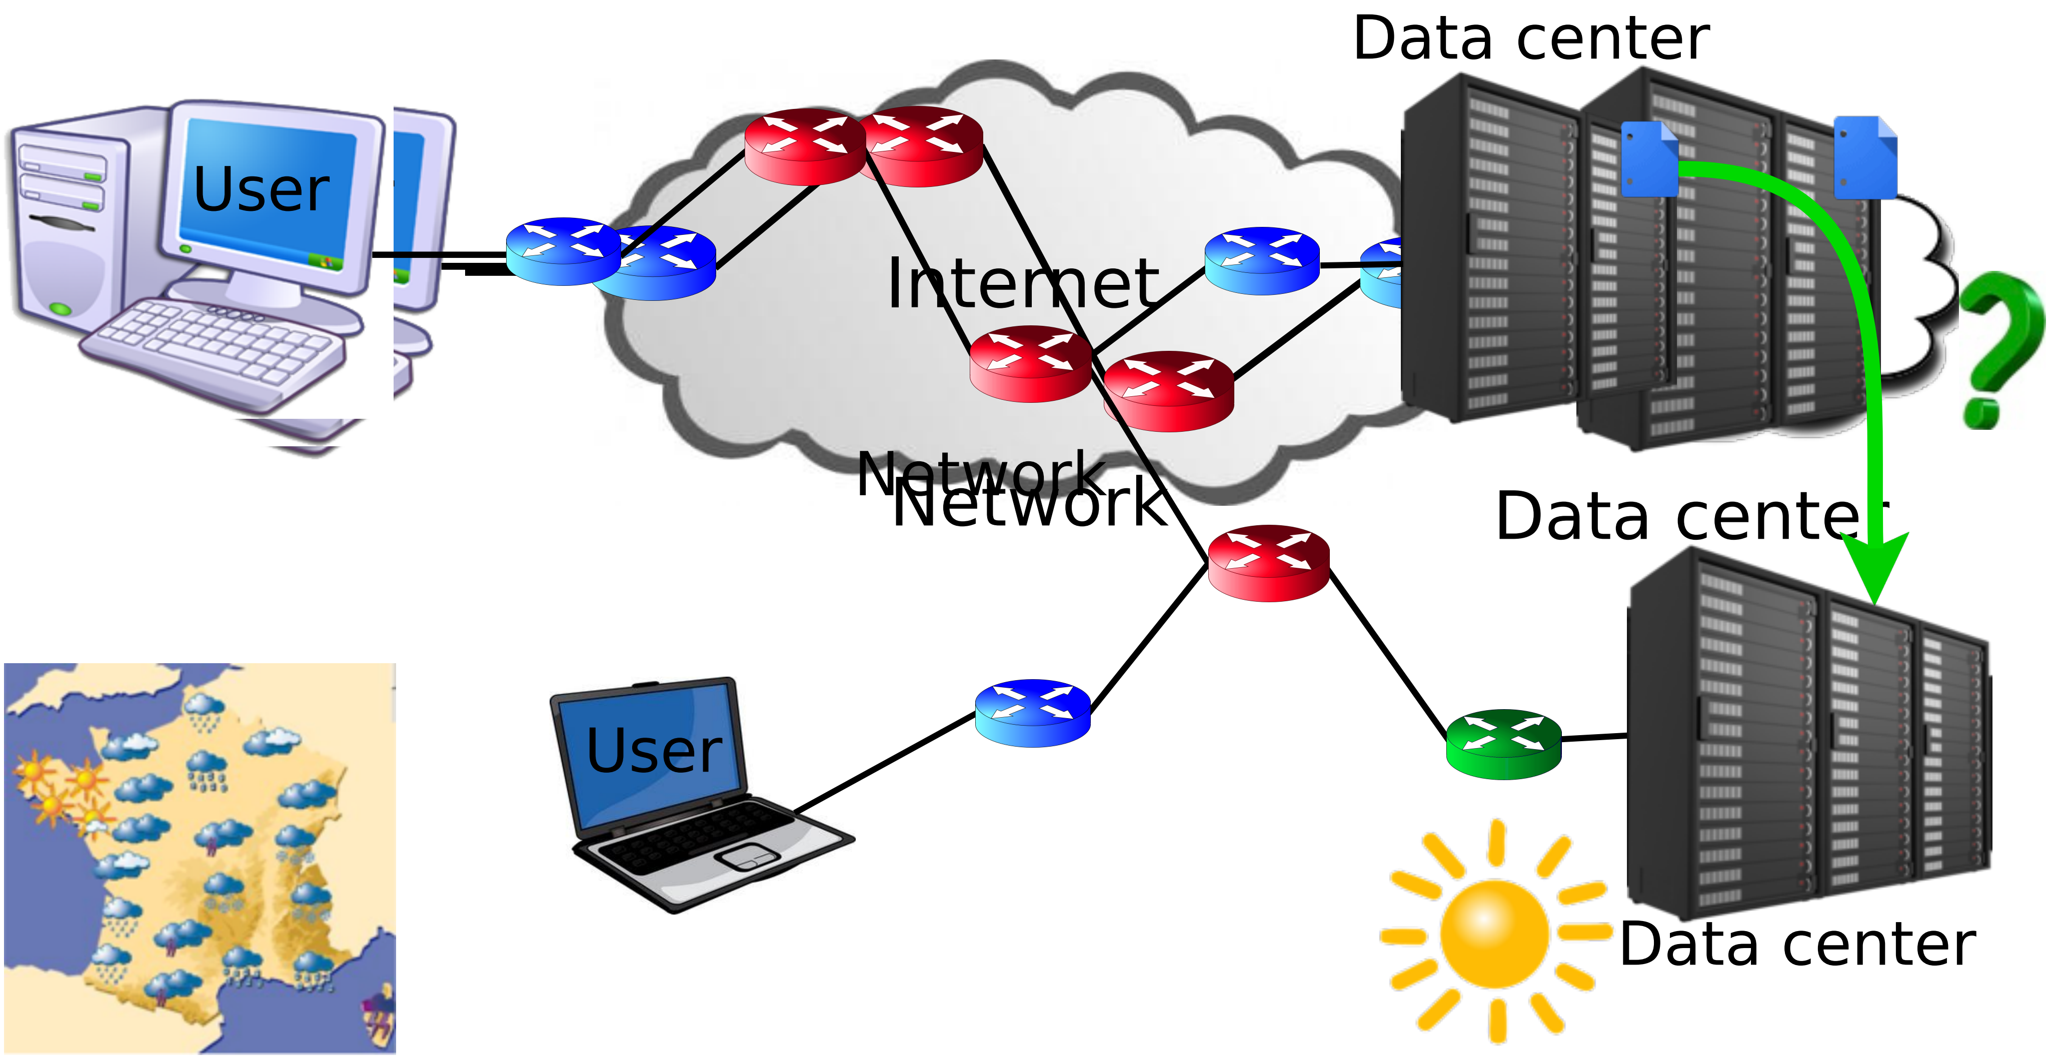

# Distributed Clouds

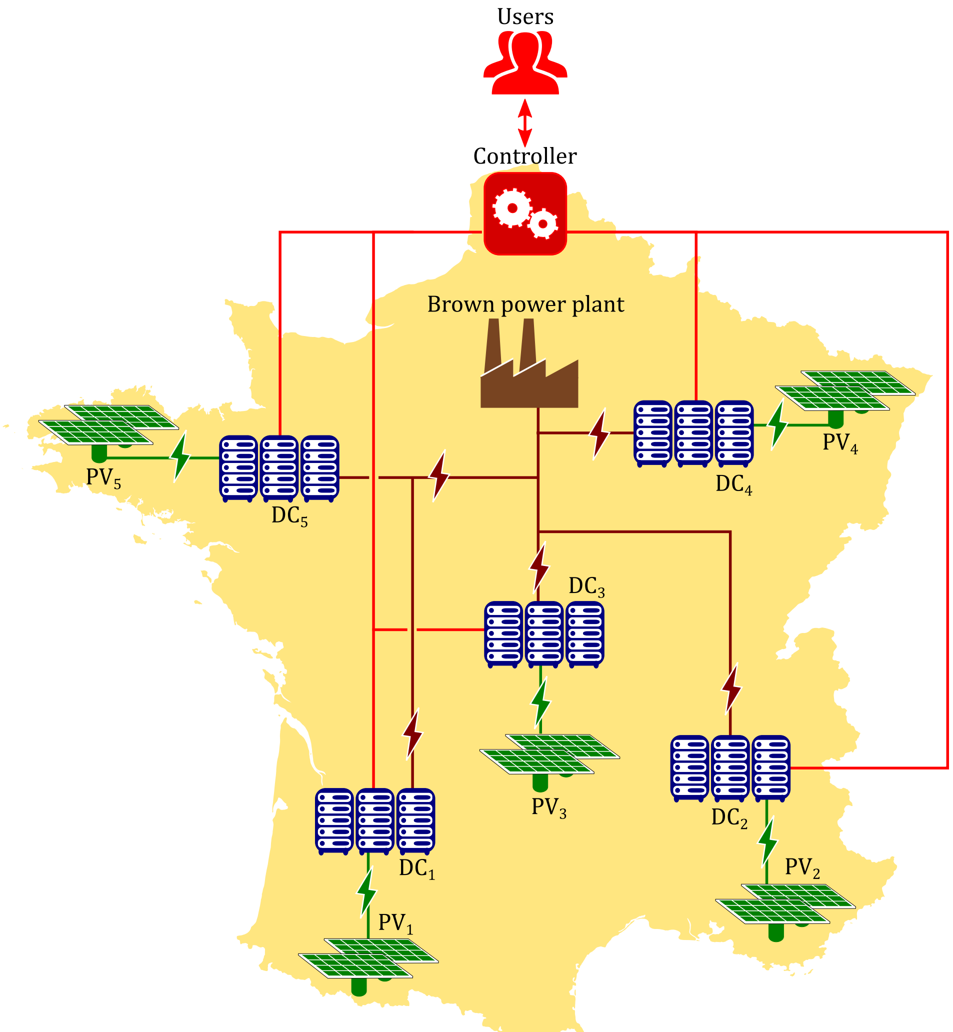


- Distributed small data centers
- Stochastic model of solar electricity production
- Greedy heuristics for resource allocation
- VM migration model
- Simulation of real data center traces

B. Camus, F. Dufossé, A.-C. Orgerie , "A stochastic approach for optimizing green energy consumption in distributed clouds", SMARTGREENS 2017.

# Validating the models

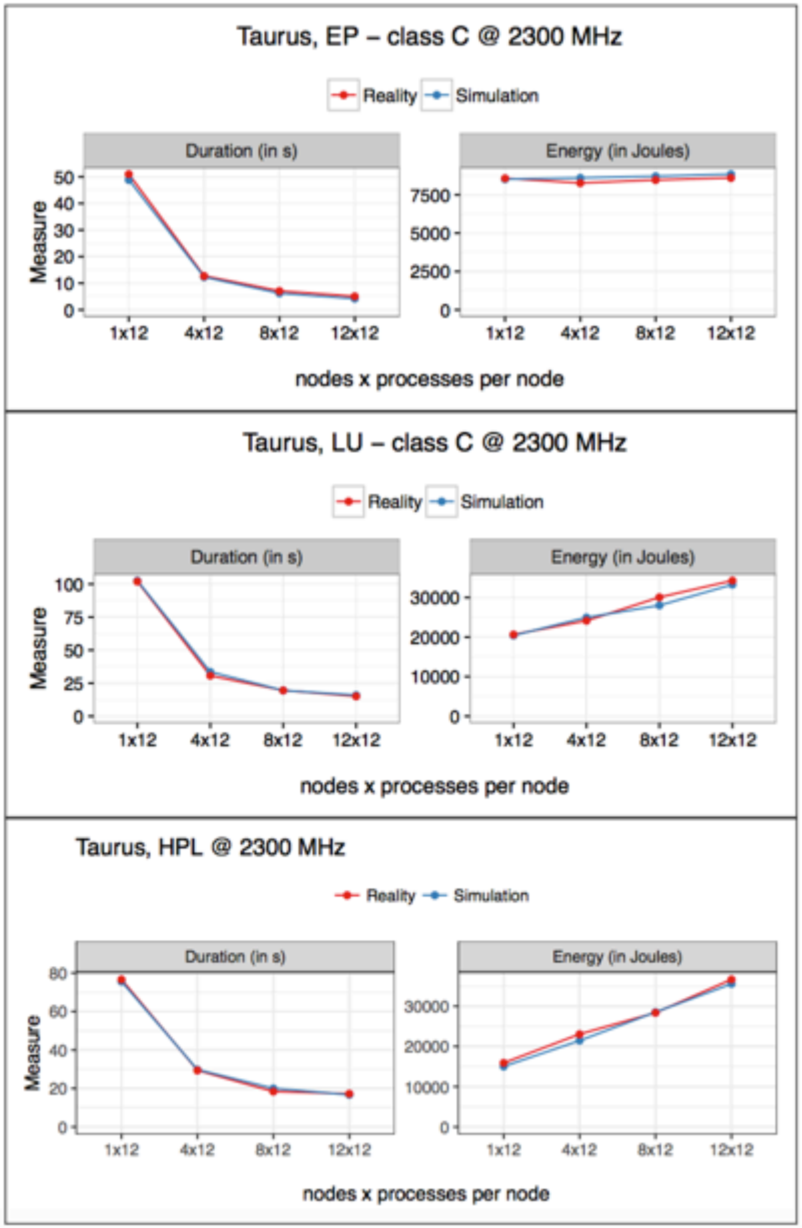


SIMGRID

- DVFS
- On/off

HAC SPECIS

"*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.

# Wired networks

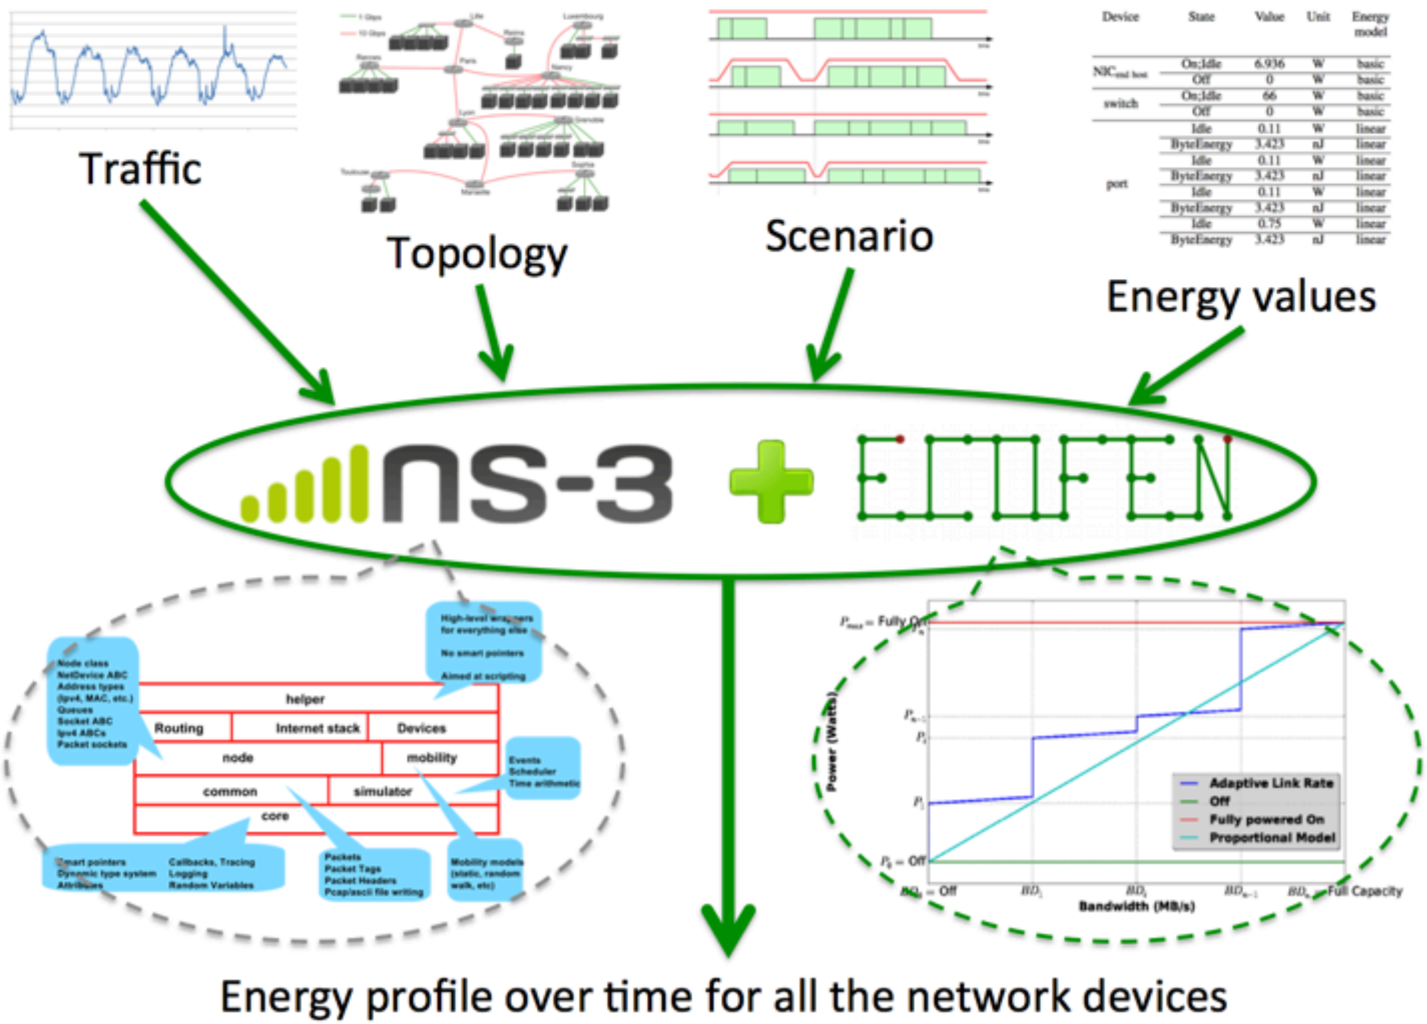


*"Simulation toolbox for studying energy scenarios in wired networks",*
A.-C. Orgerie, B. L. Amersho, T. Haudebourg, M. Quinson, M. Rifai, D. Lopez Pacheco, and L. Lefèvre, CNSM 2018.

# Conclusions

- Typical (future) application: data stream analysis for IoT devices and applications
- Real power and performance measurements on a concrete use-case
- 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

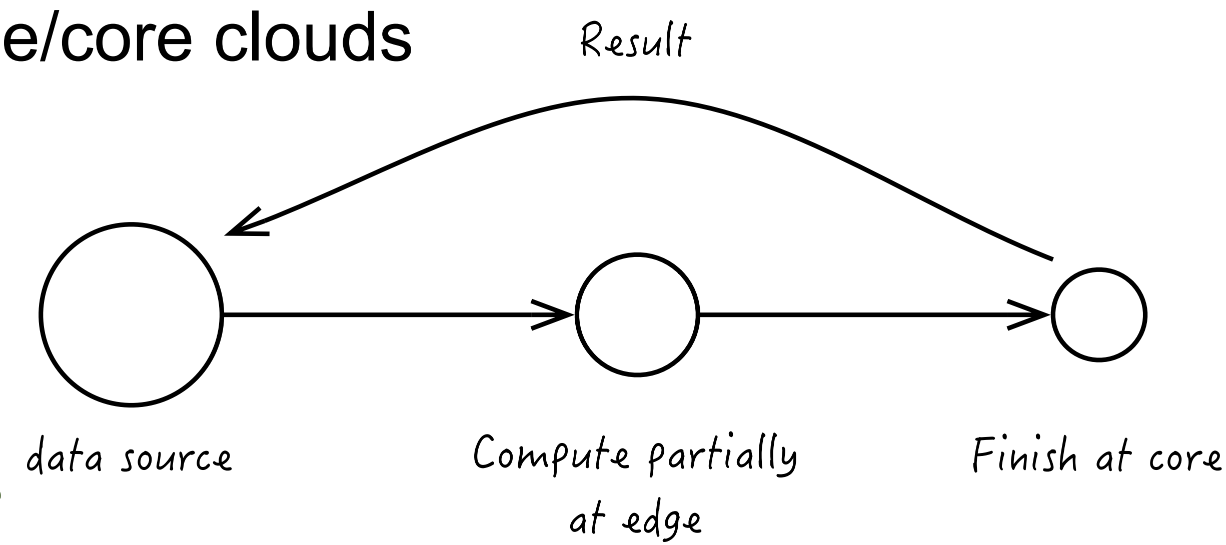

# Conclusions

- **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

# Top of the iceberg

# Are we going on the good way?




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

# The (in)dispensable weather toaster



In 2023: 50 billion IoT devices worldwide.

https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/

# The smart frying pan

Assisteo, Tefal, 2017.

# E3-RSD 2018

## http://e3rsd.irisa.fr

## 1 - 4 October 2018, Dinard, France

*Research school for early stage researchers on Energy Efficiency of Networks and Distributed Systems*

# Thank you for your attention

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