EPOC : Energy Proportional and Opportunistic Computing system

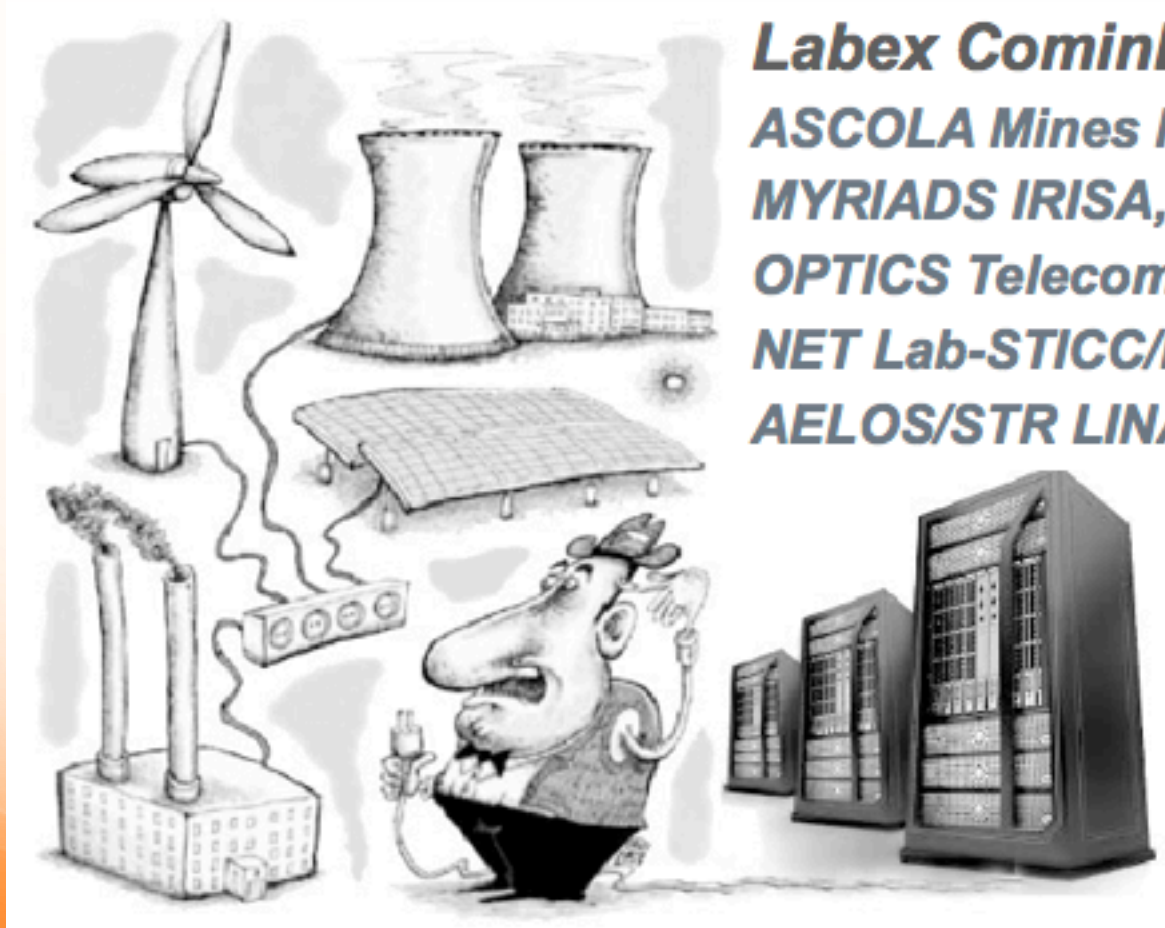

Labex CominLabs

ASCOLA Mines Nantes, INRIA, LINA

MYRIADS IRISA, INRIA

OPTICS Telecom Bretagne

NET Lab-STICC/MOM ENIB

AELOS/STR LINA, IRCCyN

ComInLabs

Project 2013-2017

- **Permanent Staff involved**
  - Jean-Marc Menaud, Thomas Ledoux ASCOLA->STACK,
  - Nicolas Beldiceneau, TASC
  - Philippe Gravey, Michel Morvan, OPTICS,
  - Anne Cecile Orgerie, Jean-Louis Pazat MYRIADS
  - Ammar Sharaiha, Pascal Morel, ENIB
  - Claude Jard, Olivier H. Roux, Didier Lime STR-AELOS
- **PhD involved**
  - Barbara Dumas-Feris (2013-2017)
  - Sabbir Hasan (2013-2017)
  - Yunbo Li (2013-2017)
  - Gilles Madi-Wamba (2014-2017)

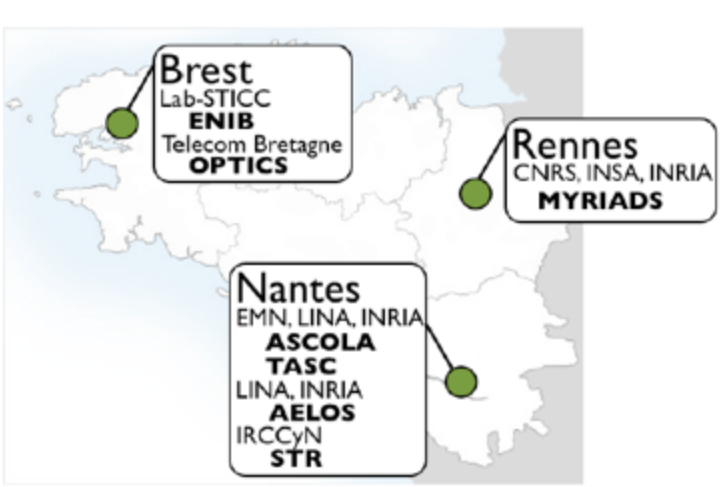

Main Focus

- With the emergence of the Future Internet and the dawning of new IT models such as Fog/Edge/Cloud computing, the usage of data centers, and consequently their power consumption, increase dramatically
- Optimizing the energy consumption of mono-site Cloud DCs connected to the regular electrical grid and to renewable-energy sources

Originality

- The EPOC project addressed the energy consumption for mono-micro-datacenter (no more 50 servers) connected renewable-energy sources.
  - In 2013, this positioning could be considered at the margin because most of the research on energy work focused on large data centers.
  - in 2018, with the emergence of Fog/Edge computing, our original initial positioning was an innovative and precursor to current mainstream research
- Workload-driven vs **Power-driven**

The problem in 3 curves

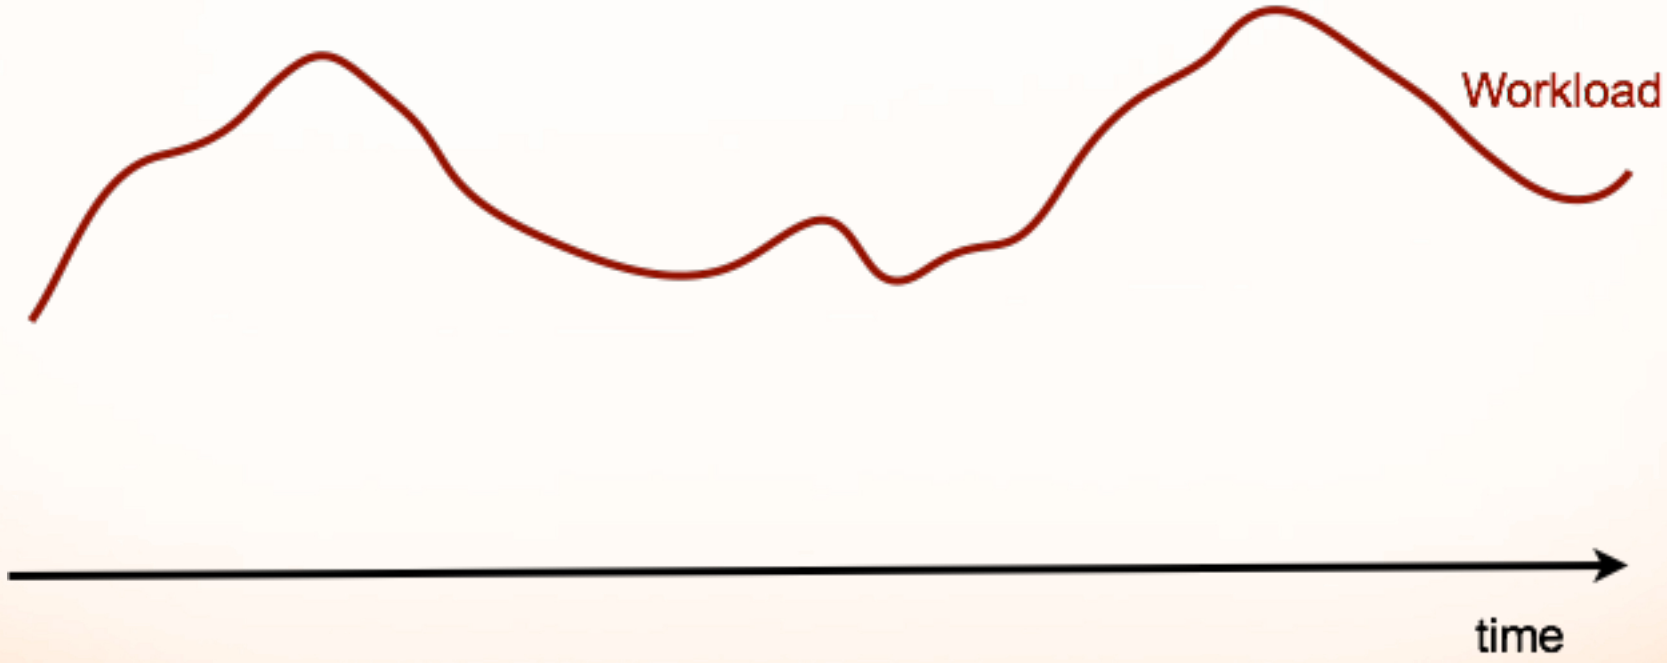

The problem in 3 curves

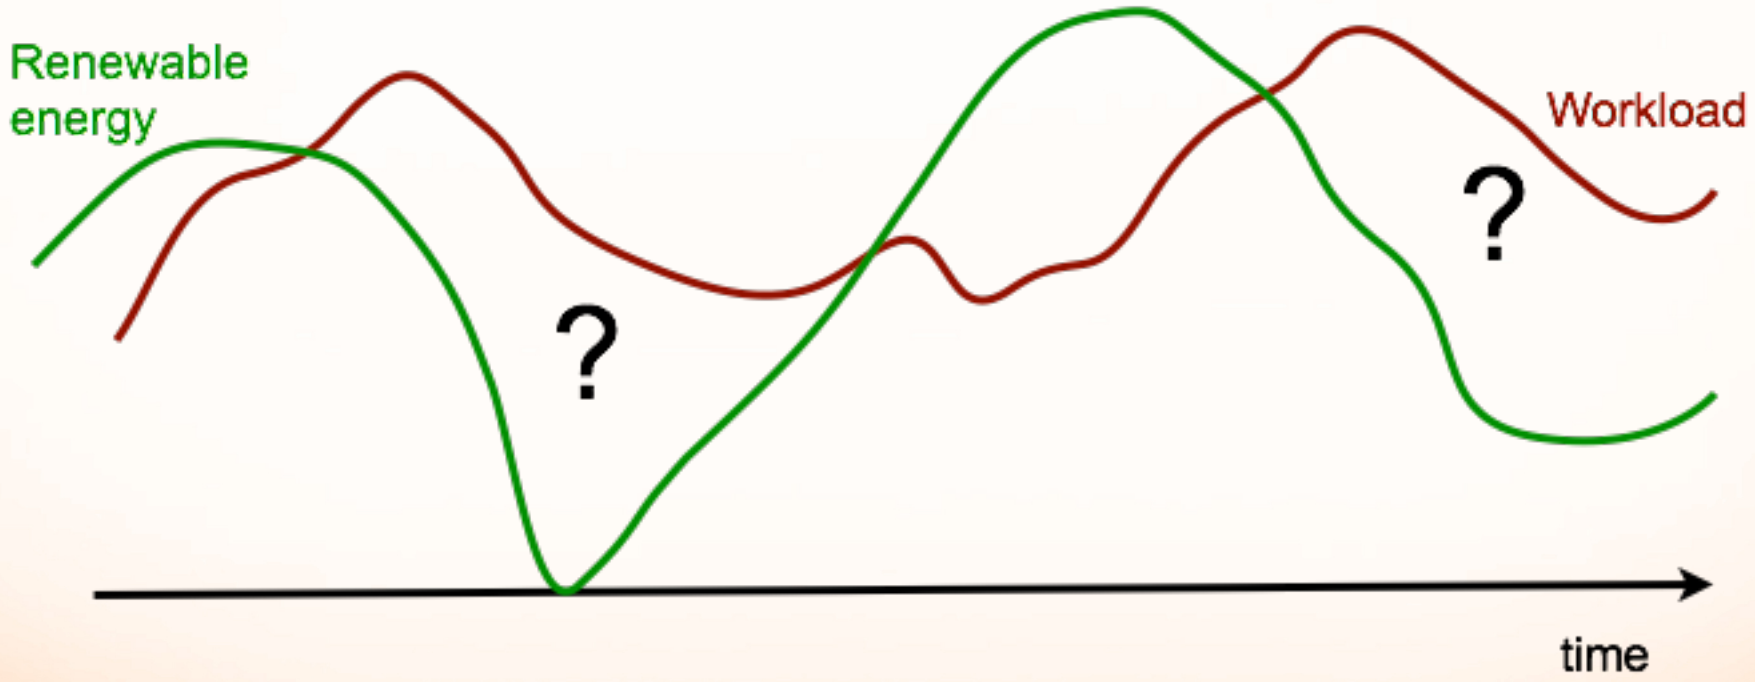

CominLabs

The problem in 3 curves

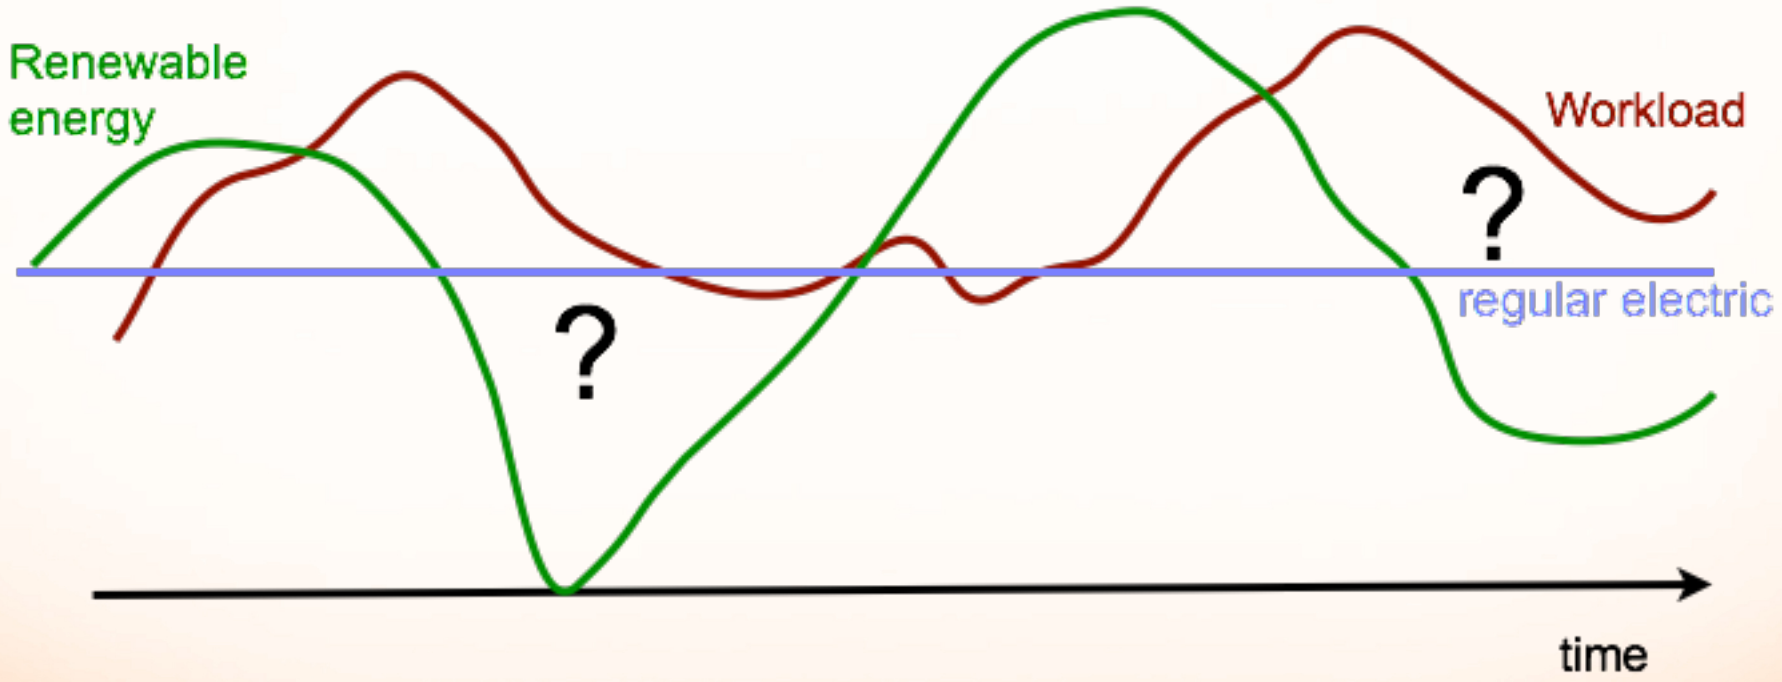

Workload-driven vs **Power-driven**

Our approach

- Address this problem at
  - Hardware level
    - Innovative infrastructure (Hyper-convergence) with low latency and high bandwidth network card.
  - IaaS level
    - Scheduling online and batch job
  - Software level
    - Greening interactive application (with SLA respect)
  - Model level
    - For workload prediction models

Hardware

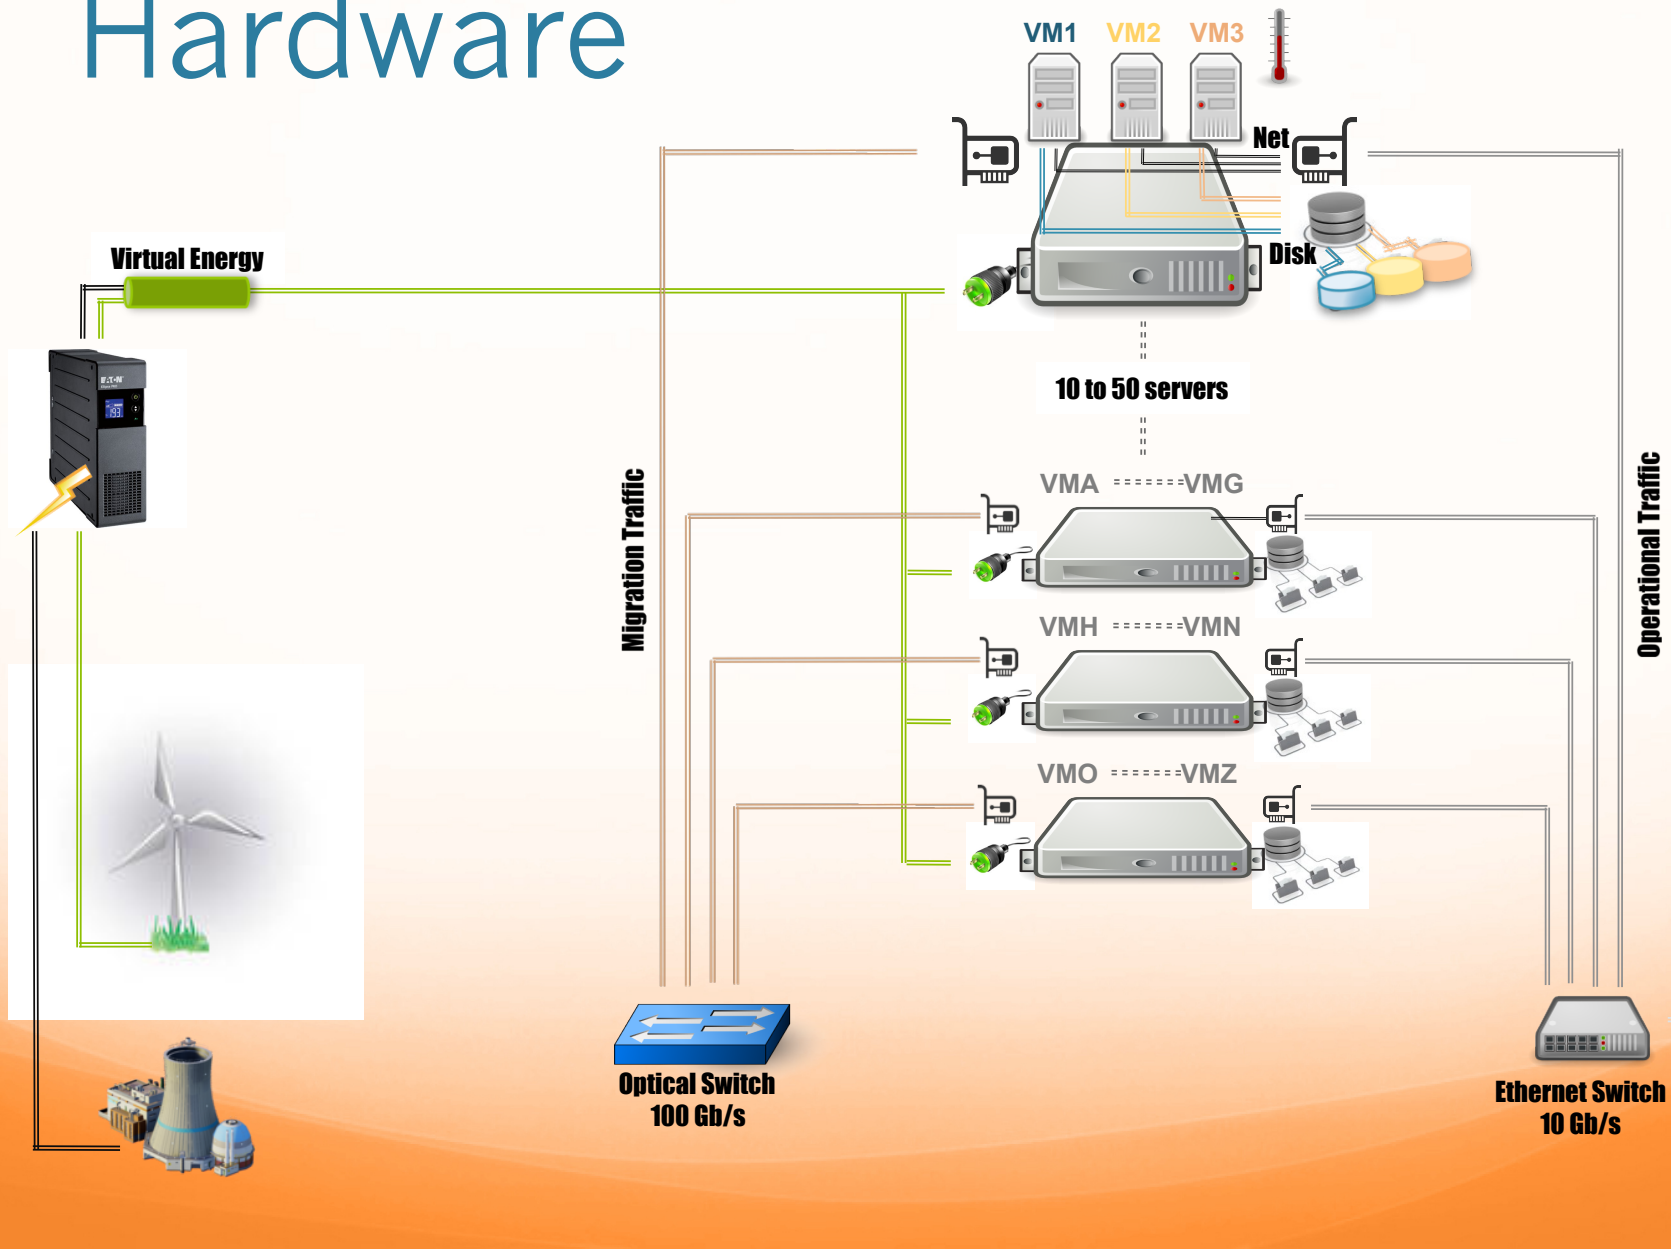
VM1
VM2
VM3
Net
Disk
Virtual Energy
10 to 50 servers
Migration Traffic
Operational Traffic
VMA
VMG
VMH
VMN
VMO
VMZ
Optical Switch
100 Gb/s
Ethernet Switch
10 Gb/s
Internet

Hardware:

support for live migration without NAS/SAN

- No dedicated storage network in order to reduce the consumption but we want turn off some servers
  - Migration of VM RAM and Disk (of about 130G)
  - Network bit rate must be near to 100 Gbit/s…
- Network: A purely passive network architecture based on fast-wavelength-tunable transmitters under the name of POPI.
  - extension called E-POPI that allows to increase the number of connected servers by using several transmission bands

IaaS:
Opportunistic scheduling

- to optimize the usage of renewable energy, one way consists in carefully scheduling the workload to align it with the time-varying renewable energy production
- An Opportunistic scheduling in presence of batch jobs
  - A solution mixing approaches in order to achieve a balance in all aspects (job, energetic costs of turning ON/OFF physical machines..), implying minimizing the renewable energy losses

One result

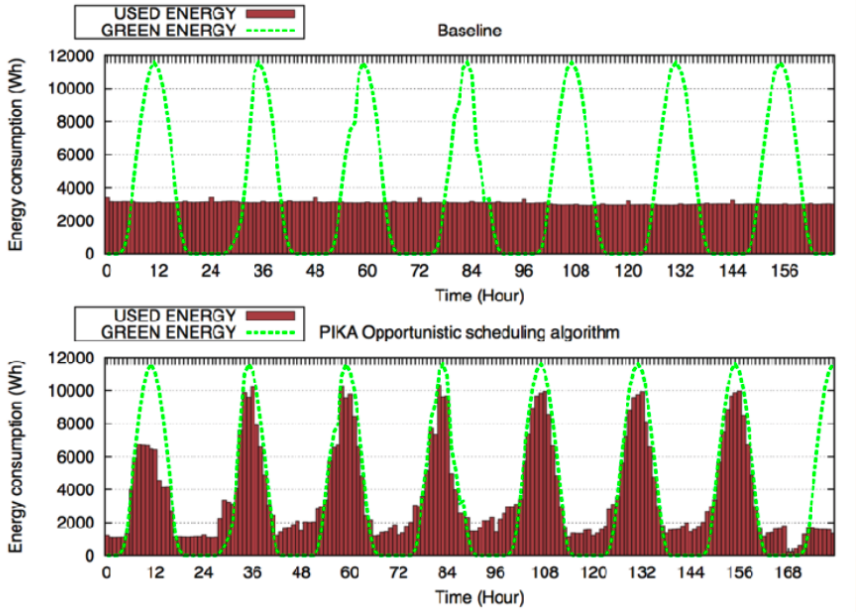

Policy	Total Energy (Wh)	Brown Energy (Wh)	Green Energy (Wh)
Baseline	768,724	442,085	326,639
Baseline + ESD	792,155	280,441	511,714
PIKA	892,458	378,569	513,889
PIKA + ESD	914,944	209,935	705,009

Software:

Greening interactive application

- How to make Cloud application, specially interactive application, adaptive to green energy availability while traditional QoS properties can be at satisfactory level so to lower carbon footprint?
- Virtualization of green energy concept to tackle both the forecasting error and intermittency of green energy.
- A self-adaptive autoscaler architecture to enable smart usage of energy in an interactive application
  - Extend to: actuate at application and at infrastructure level depending on the application behavior and energy availability

One result

- **GPaaScaler**
  - Reduce significant amount of brown energy consumption by 35% compared to a baseline approach when green energy aware adaptation is considered

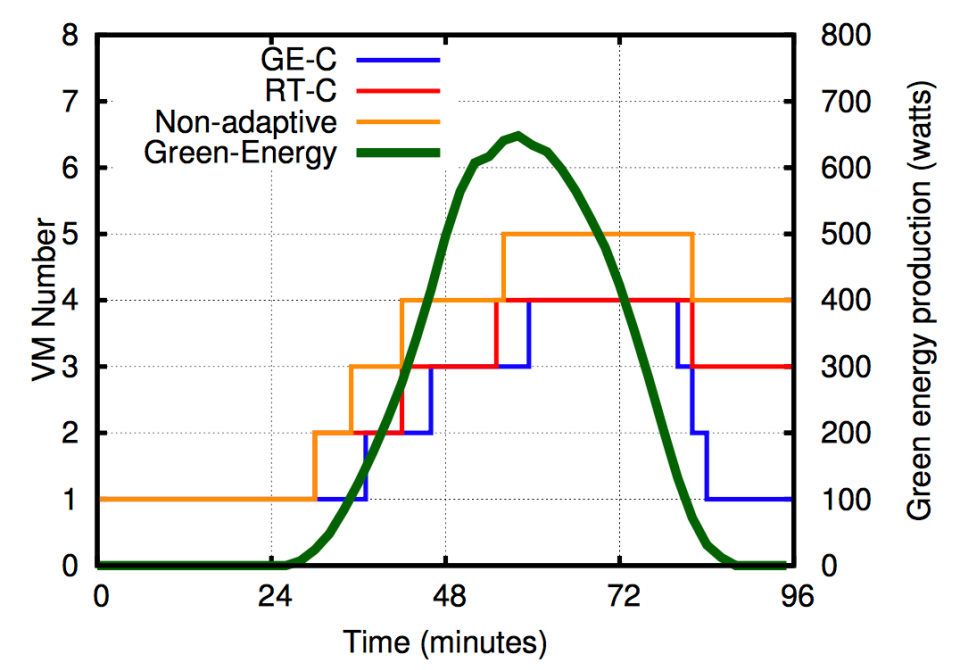

Model:
workload prediction models

- workload prediction models
  - used for dynamic resource provisioning in Cloud environments with respect to renewable/green energy availability
  - used in workload trace generators that can help to extend an experimental dataset in order to test more widely resource optimization heuristics
- Both approaches are shown to be complimentary as neural networks give better prediction results, while constraint programming is more suitable for trace generation

One result

- Prediction of values
  - Implemented with Sictus2, Mac OS 10.10 Yosemite, 16Go, Intel i7 à 2,93GHz
  - Real workload from EasyVirt : 50 PM, 1000 VM, 6 months, values range from **3000 to 30000**
- Average value error for the CP model : **2800**
- Average value error for the NN model : **1600**

Results

- International Scientific production:

2014	2015	2016	2017	Total
1	3	4	11	**19**

- Projects
  - SeDuCe: experimental infrastructure dedicated to the study of data centers with low energy footprint.
  - HYDDA: deployment of applications on heterogeneous platforms and how optimize resources usage
- Collaborations
  - Manish Parashar is Distinguished Professor of Computer Science at Rutgers (USA)

CominLabs

Seduce

Cluster
1000 Cores
Cooling : 35°c

Photovoltaïc Panel
50kWc

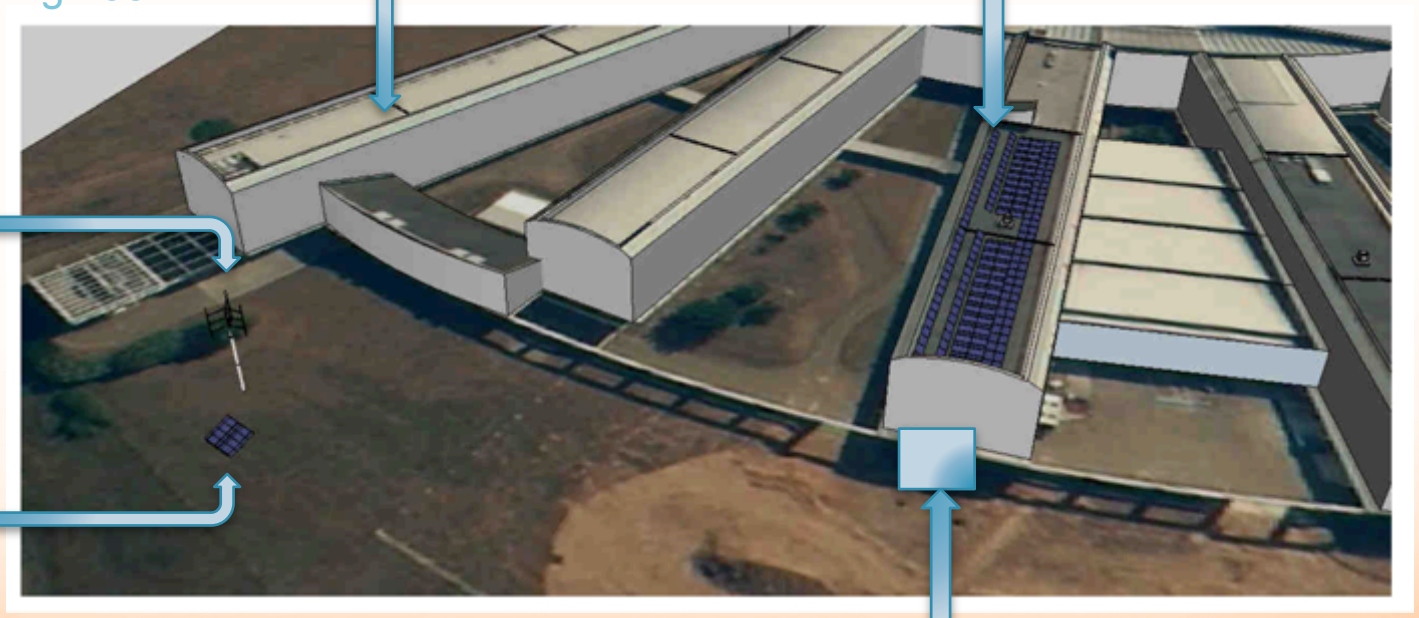

Tracker
3kWc

Tracker
3kWc

Technical local

ComInLabs

Our PhD

- Yunbo Li, Resource allocation in a Cloud partially powered by renewable energy sources
  - post-doc position at Orange Labs, Rennes.
- Md Sabbir Hasan, Smart management of renewable energy in clouds : from infrastructure to application
  - post-doc position at Orange Labs, Grenoble.
- Barbara Dumas-Feris, Réseaux optiques en mode paquet pour les connexions internes à un centre de données
  - post-doc position at IMT-Atlantique
- Gilles Madi Wamba, Combiner programmation par contraintes et apprentissage machine pour construire un modèle éco-énergétique pour petits et moyens data centers"
  - post-doc position at IMT Atlantique