



Energies renouvelables et centres de données

Jean-Marc Menaud

ASCOLA Mines Nantes, INRIA, LINA

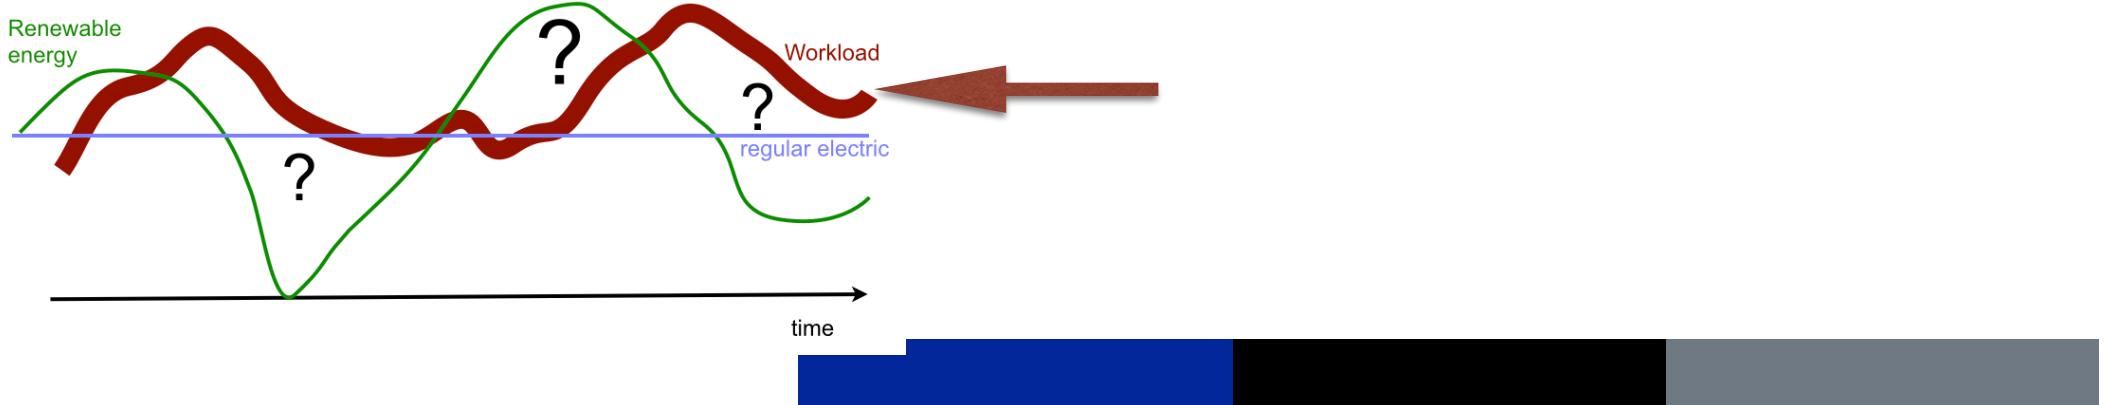
**Ecole Jeunes Chercheurs sur
l’Efficacité Energétique des
Réseaux et Systèmes Distribués**

26 Mai 2016
Mines Nantes



Multi-resources

- **Generalization :**
 - **Server resources**
CPU, RAM, Disk, Net, Energy
 - **Rack Ressources**
Net, cooling, space
 - **Data center resources**
Cooling, Humidity, Noise, Electrical, Phases, UPS, ...
- **How optimize virtualized datacenter with multiple inter-dependent objectives ?**
 - Ex: you can increase room temperature for reducing the cooling energy consumption, but a collateral effect should be done by a fan speedup (and increase all servers power consumptions).
 - How can express relation between cooling and server consumption
Server consumption and noise etc.
- **Multi-resources dynamic consolidation**

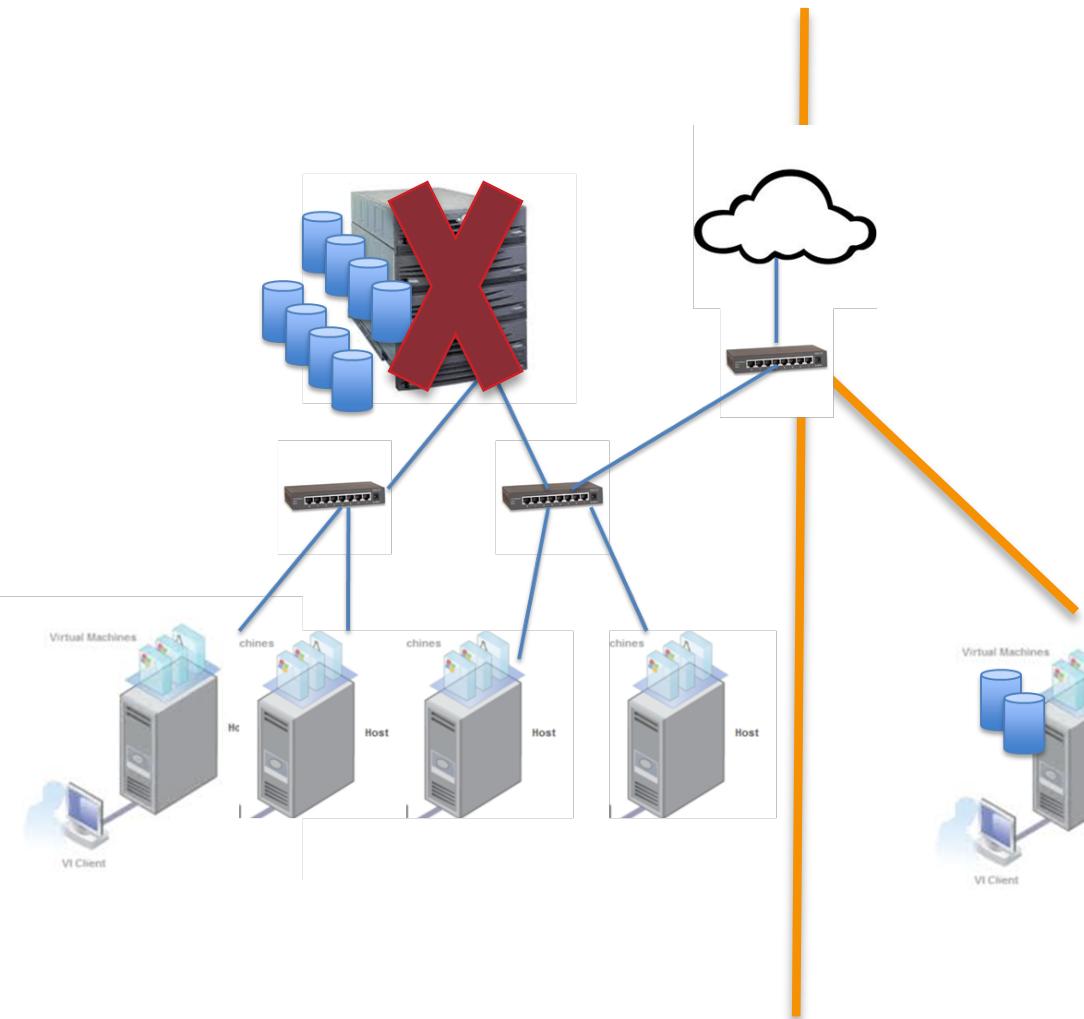


Solutions logicielles pour la maîtrise énergétique des centres de données

Grands principes
Energy driven

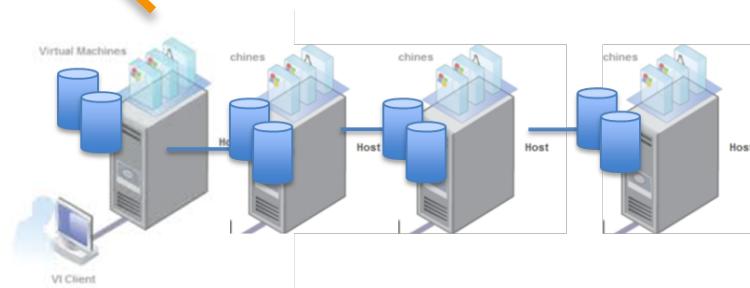
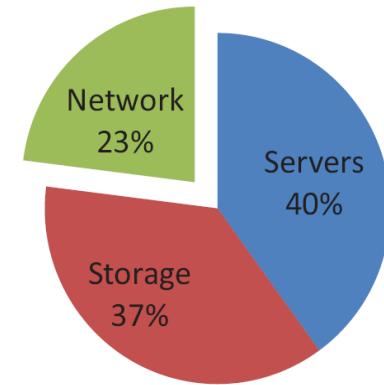
Infrastructure

Standard Solution



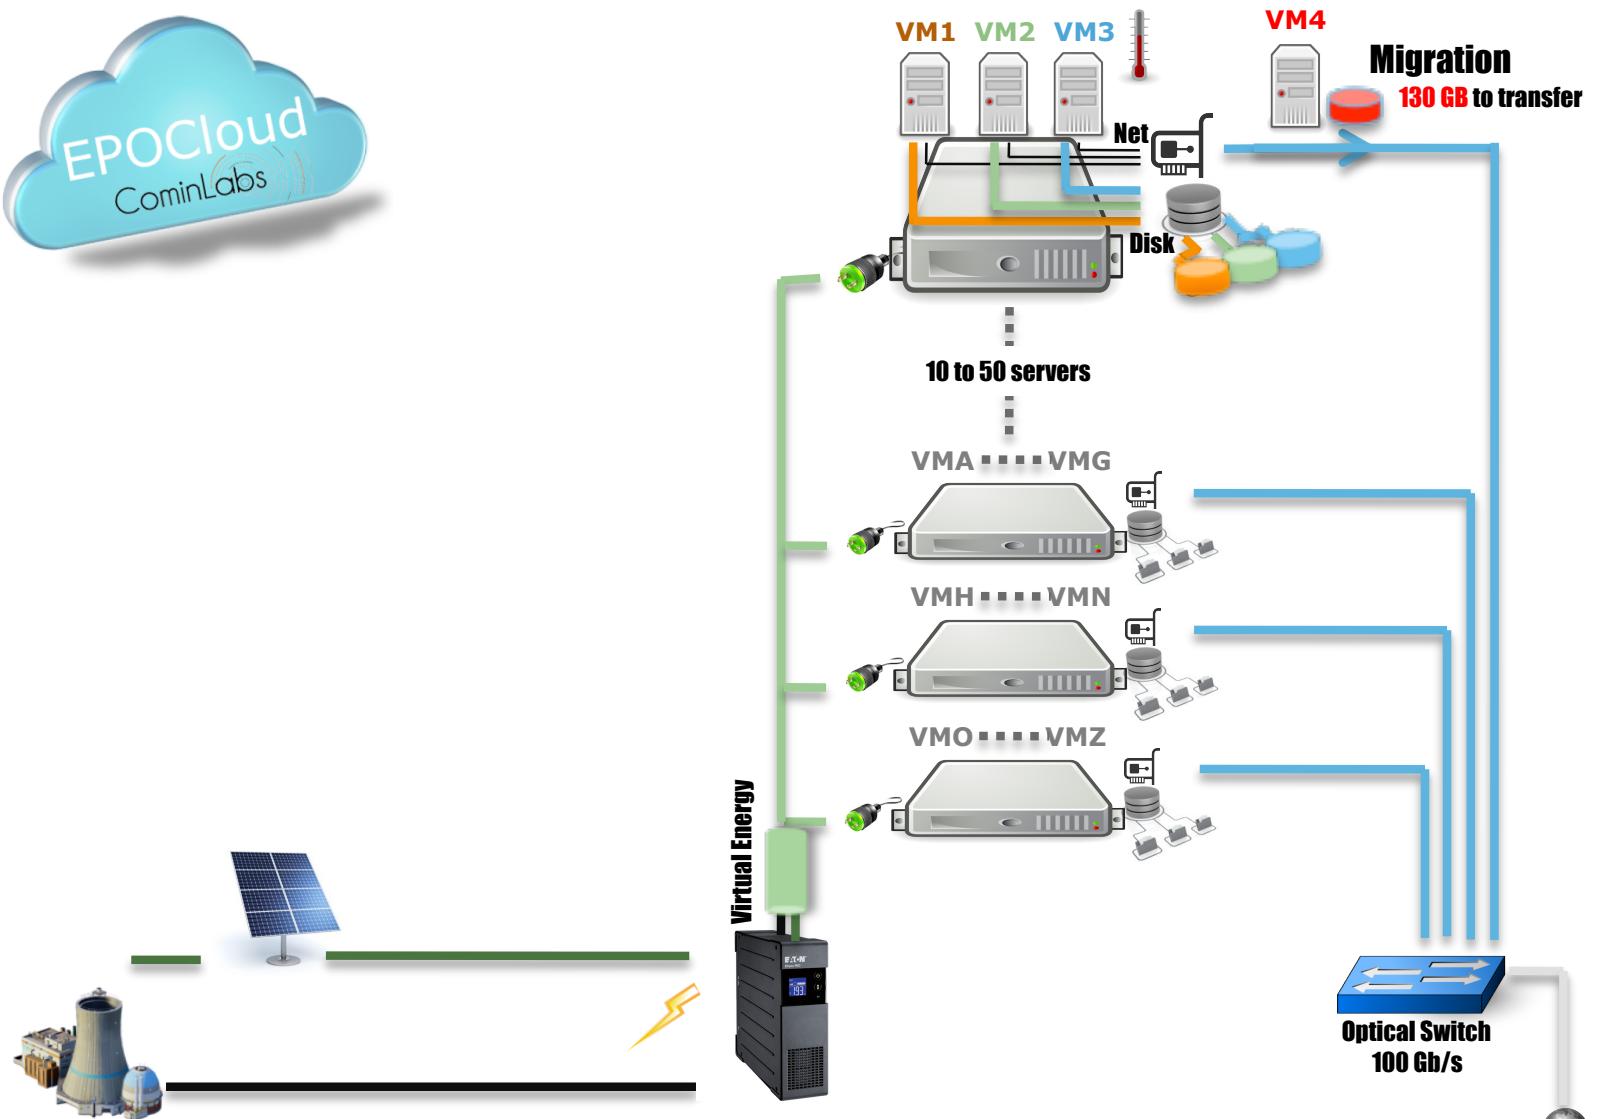
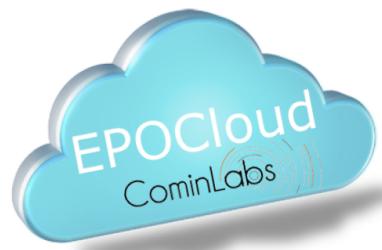
EPOCloud

Power consumption distribution

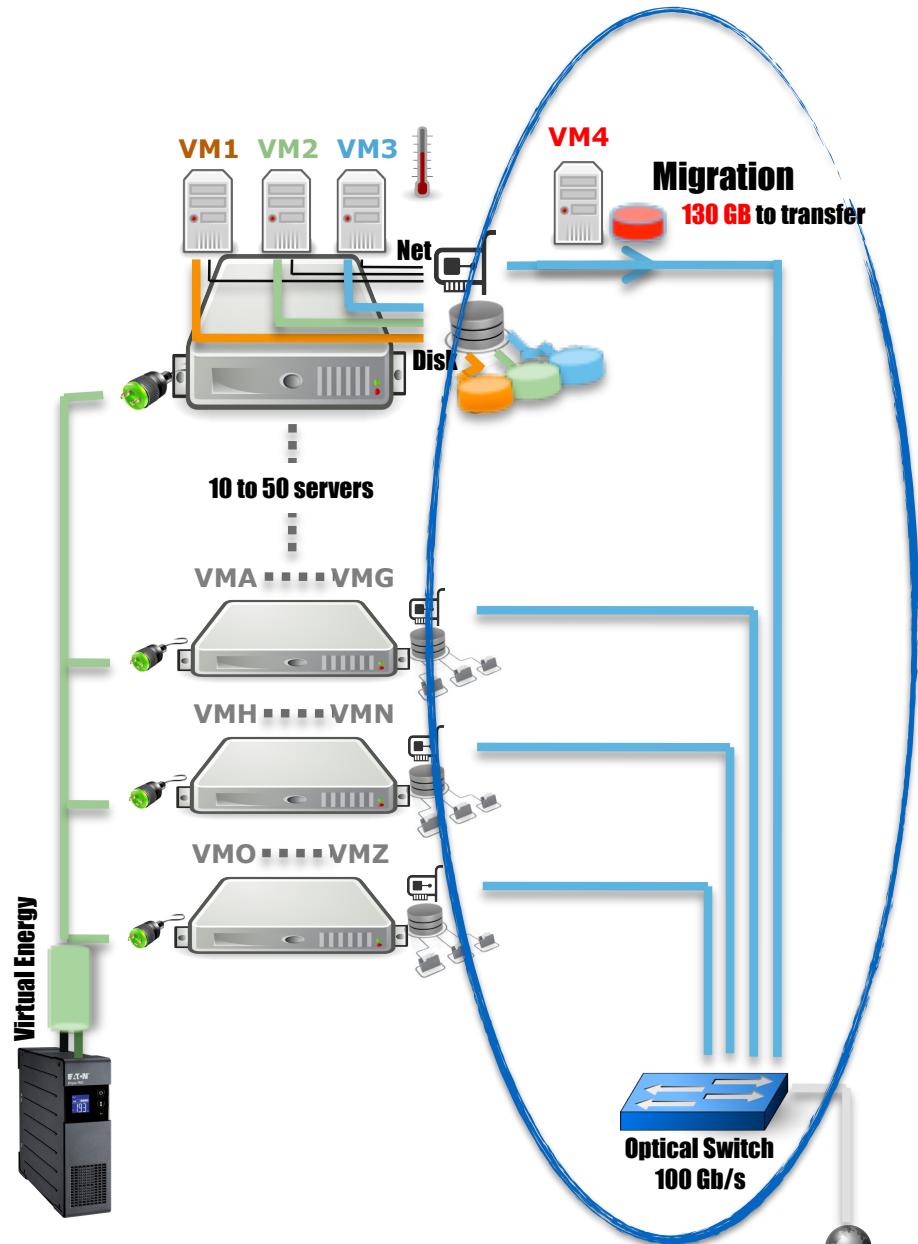
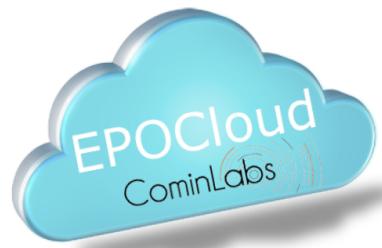


2

Infrastructure

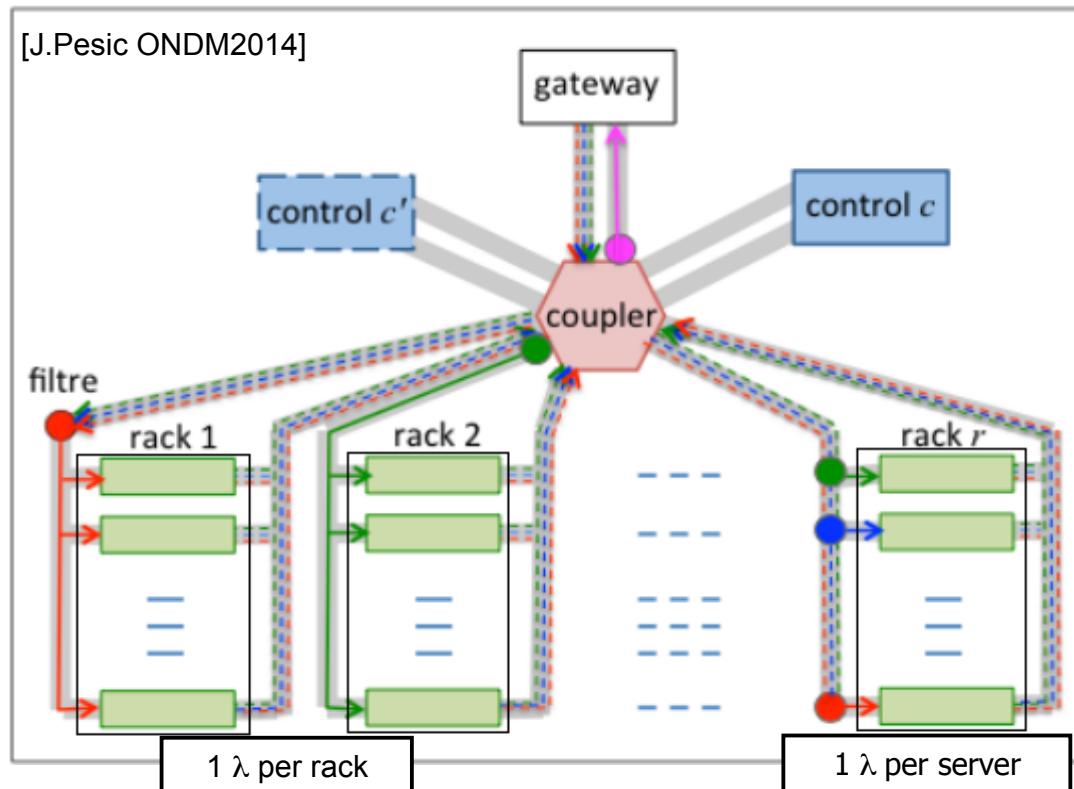


Optical Network



POPI architecture (based on TWIN)

**POPI : Passive Optical Pod Interconnect
(using tunable lasers)**



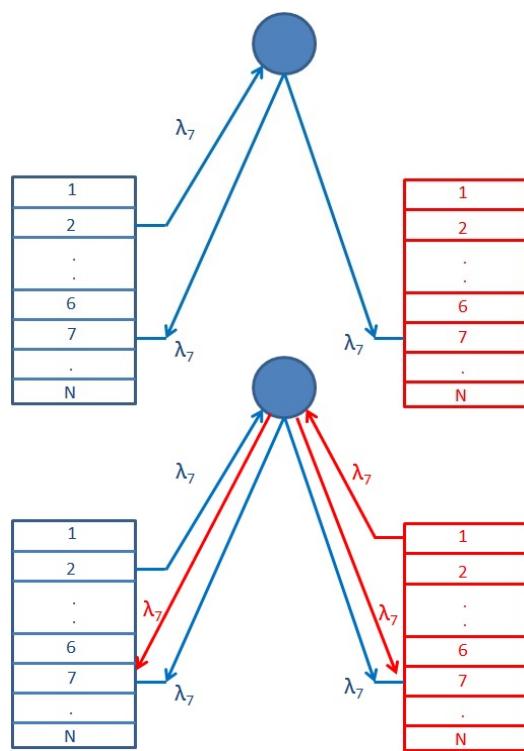
Current POPI limitation is about 60 channels using 10G NRZ format and 50 GHz grid in C-band (1530-1565 nm).

TWIN: Time-domain Wavelength Interleaved Networks

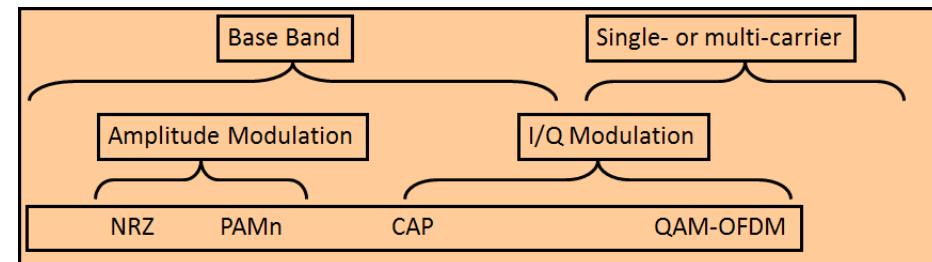
Enhancing POPI architecture

DOUBLE BAND POPI

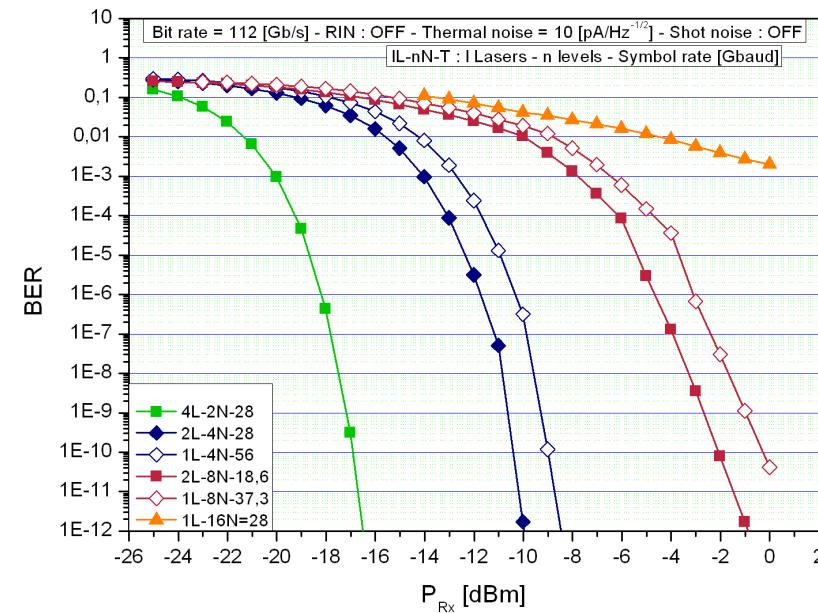
- By using C+L (1565-1625 nm) bands



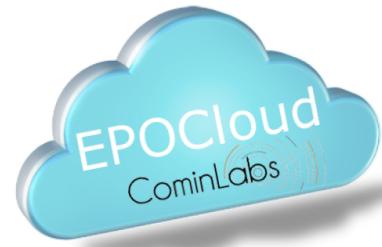
SHORT REACH MODULATION FORMATS



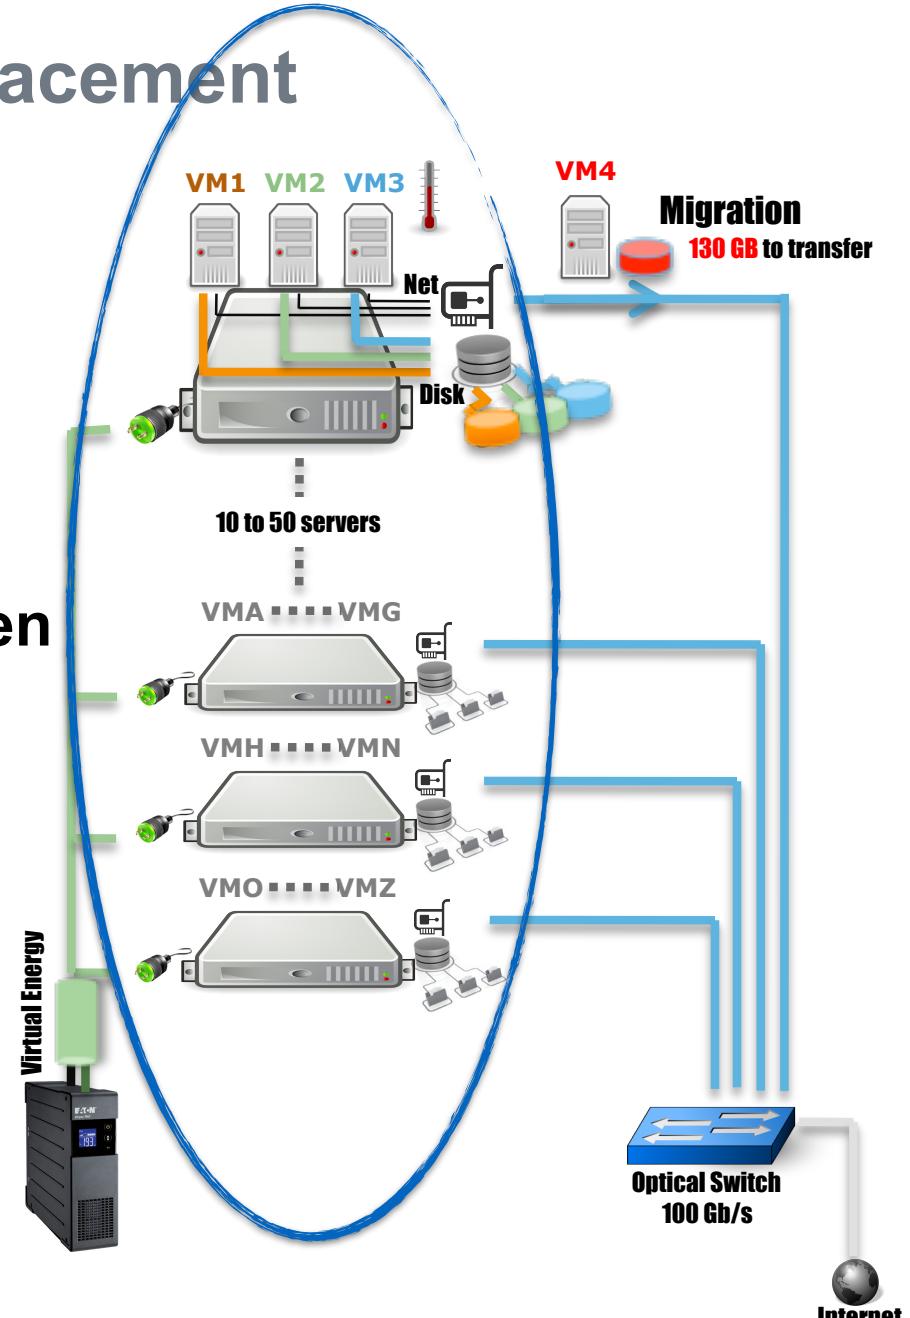
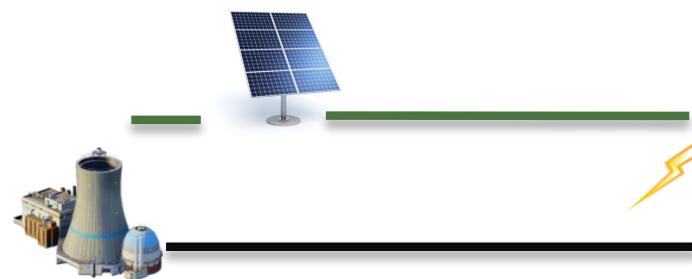
PAM : first VPI simulation results



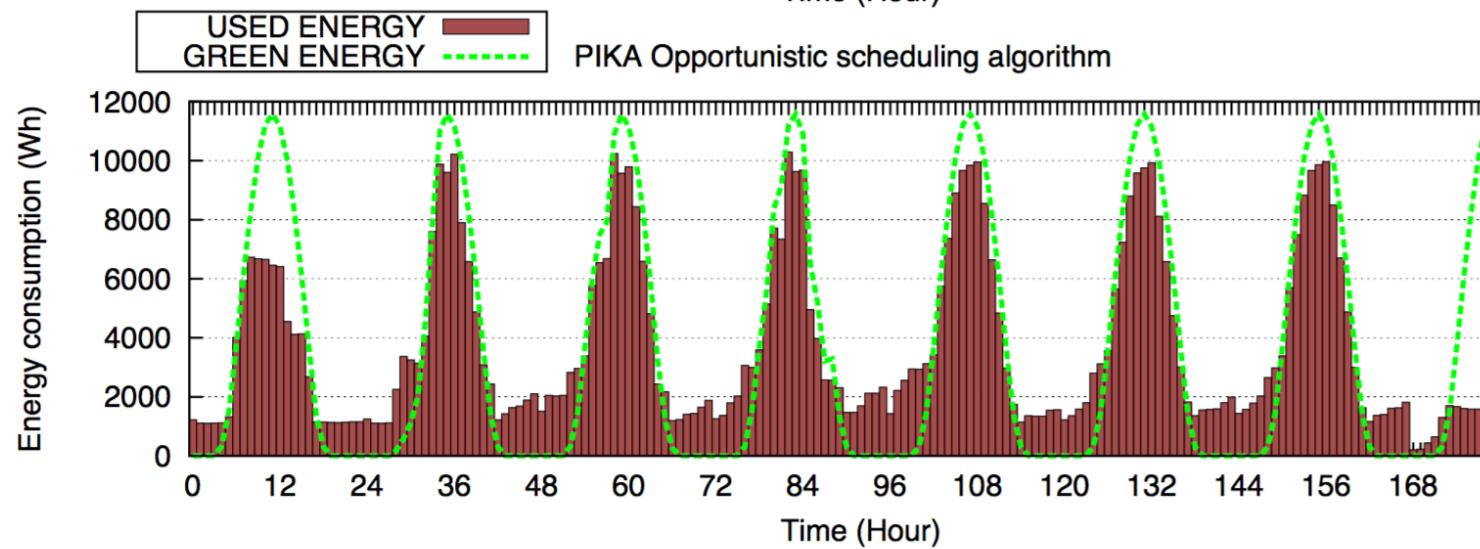
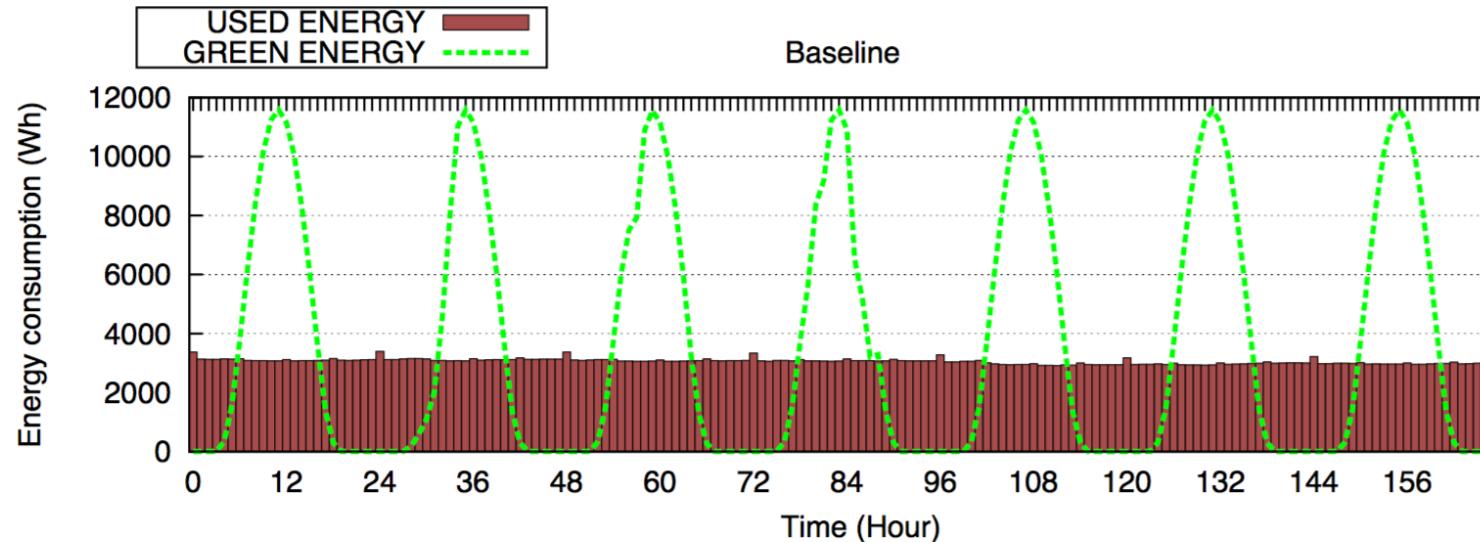
Virtual Machine placement



Workload-driven/Power-driven



Last Result



Job scheduling

- **Where** (e.g. #node/PMs, #cluster, PMs set etc.) The node has been chosen should fit VM's demand
 - Constraints of PMs:
 1. CPU resource (e.g. number of core)
 2. Memory resource

Distinguish between two kinds of job:

- ① web-job - Non-Interruptible
- ② batch-job - Interruptible with deadline

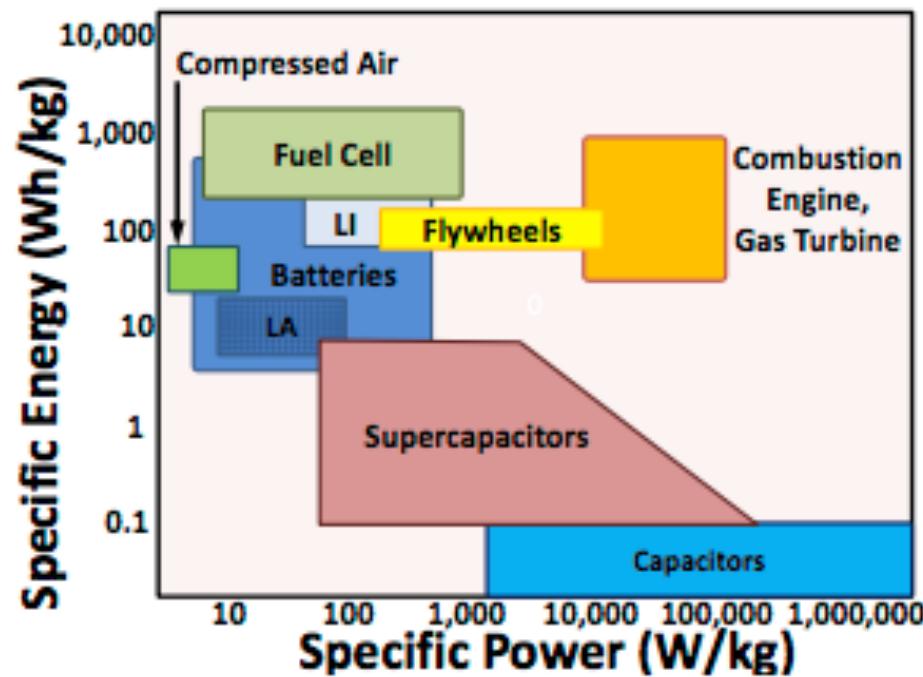
Schedule/Reschedule jobs at the beginning of each time frame:

- **When**
 - Schedule web-job first (Algo-webjob)
 - If there are still *free* resources, then schedule the batch-job (Algo-batchjob)
 - If there is not enough green energy, try to switch off PMs to save energy (Algo-consolidation if necessary)

1 / 2

Algorithm	Total E. C.	Brown E. C.	Renewable E. C.
Baseline	513.633	259.559	254.073
PIKA	676.895	142.957	533.938
	31% ↑	44.9% ↓	110.1% ↑

Et les batteries ?





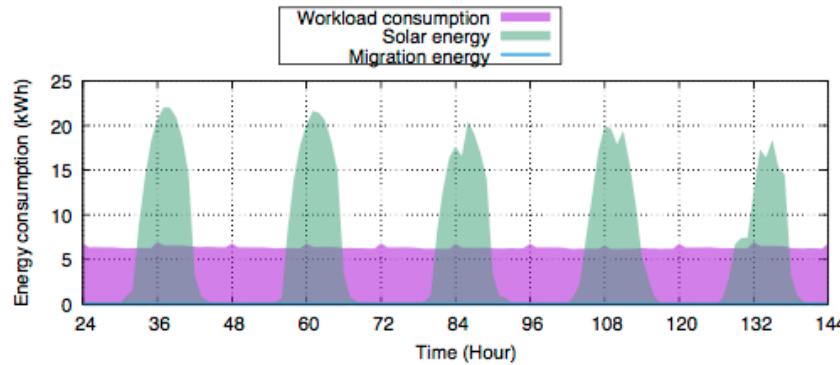
Batteries lead-acid Batteries lithium-ion Ultra/Super-capacitors
Flywheels Comp. Air Energy Storage

ESD	LA	LI	UC	FW	CAES
Energy Cost C_k^{eng} (\$/kWh)	200	525	10000	5000	50
Power Cost C_k^{pow} (\$/kW)	125	175	100	250	600
Energy Density v_k^{eng} (Wh/L)	80	150	30	80	6
Power Density v_k^{pow} (W/L)	128	450	3000	1600	0.5
Discharge:Charge Rate γ_k	10	5	1	1	4
Life Cycle L_{cyc_k} (# discharges x 1000)	2	5	1000	200	15
Max. DoD DoD_k^{max} (%)	80	80	100	100	100
Float Life T_k^{max} (years)	4	8	12	12	12
Energy Efficiency η_k (%)	75%	85%	95%	95%	68%
Self-discharge μ_k per Day	0.3%	0.1%	20%	100%	low
Ramp Time T_k^{ramp} (sec)	0.001	0.001	0.001	0.001	600

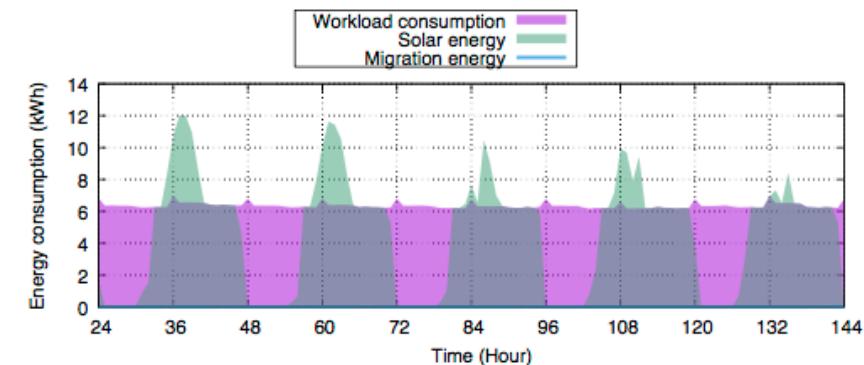
Our last results

Algorithm	Brown energy (kWh)
baseline no bat.	442.085
baseline bat.	280.441
PIKA no bat.	378.568
PIKA bat.	209.935

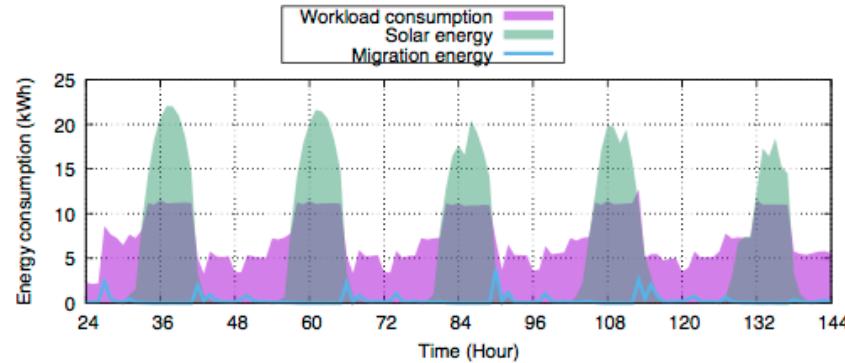
Table II: Brown energy consumption



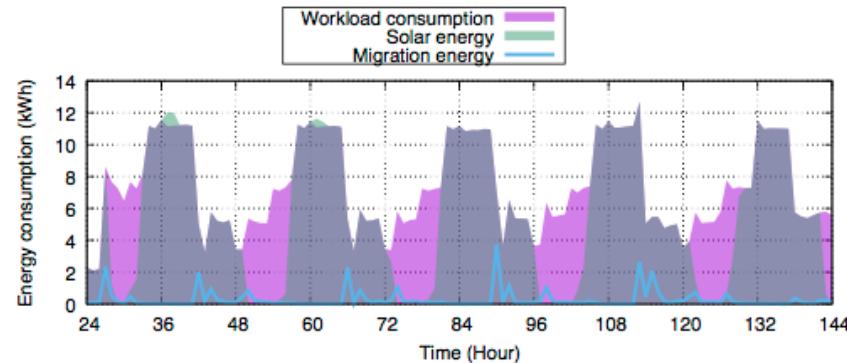
(a) Baseline without ESD



(b) Baseline with ESD



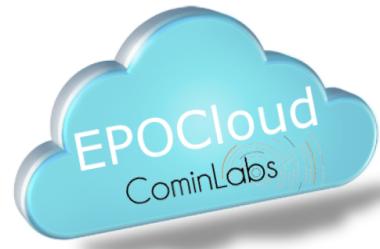
(c) PIKA without ESD



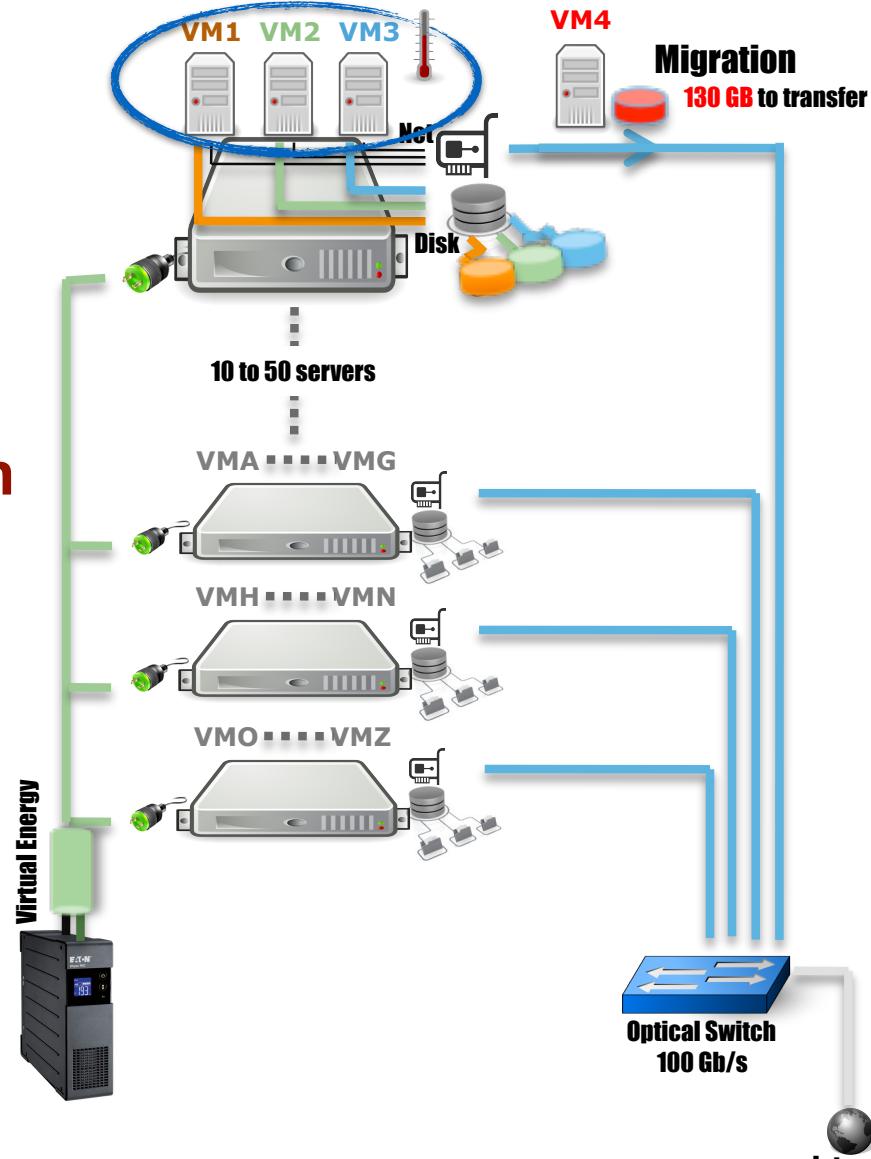
(d) PIKA with ESD

Figure 11: The energy consumption with fixed solar PV panel dimension (160 m^2) and fixed LI battery size (40 kWh)

Virtual Machine placement

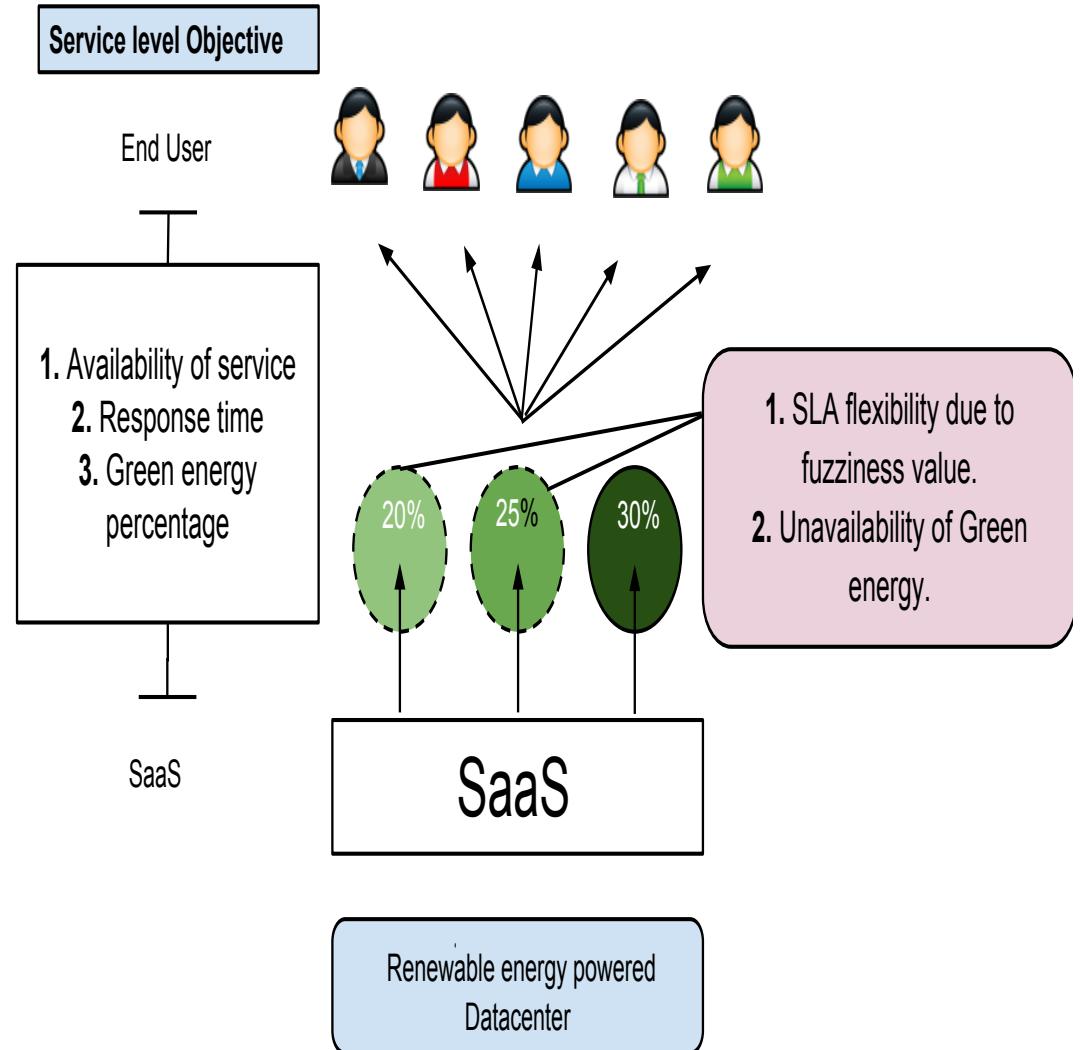


■ SLA, Application degradation



Enforcing Green SLA

- ❖ Green SLA can be contracted based on the usage of green energy for the service.
 - Clients want their application to run with x% (e.g., 30%) green energy on the cloud.
- ❖ Managing Green SLA.
 - Trading Green energy
 - Reducing energy consumption of application

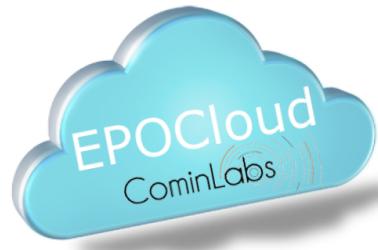




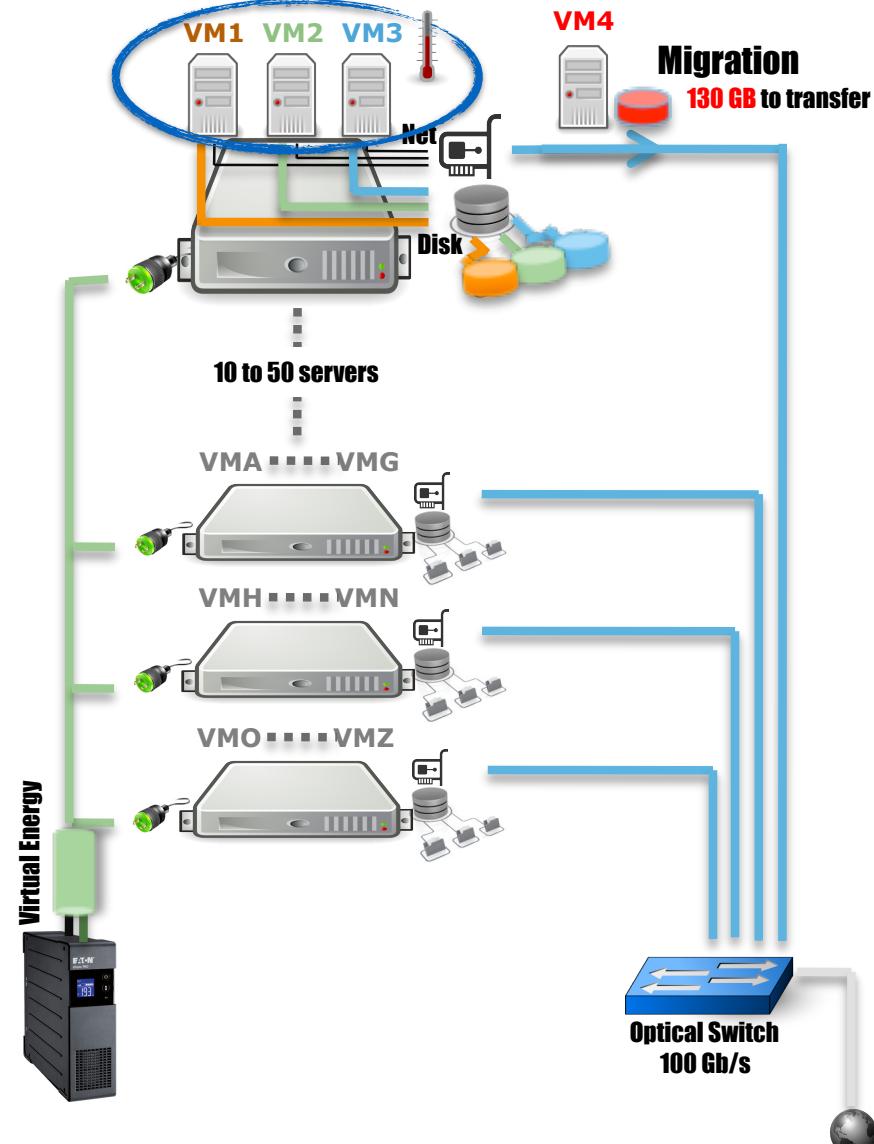
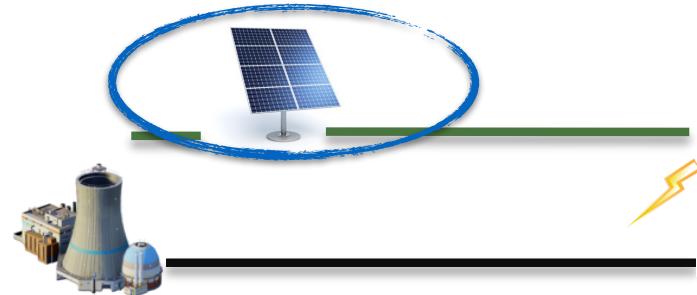
Performance degradation

- ✧ Reducing energy consumption
 - By performance/functionality/feature degradation of Cloud application.
 - Degradation based on the presence and absence of green energy.
- ✧ Performance is related to CPU and Ram usage, thus affecting energy consumption.
- ✧ How to degrade?
 - Algorithm selection depending on time and space complexity.
 - Finding sweet spot frequencies for algorithms.
 - Reducing CPU clock when deadline is not nearest.

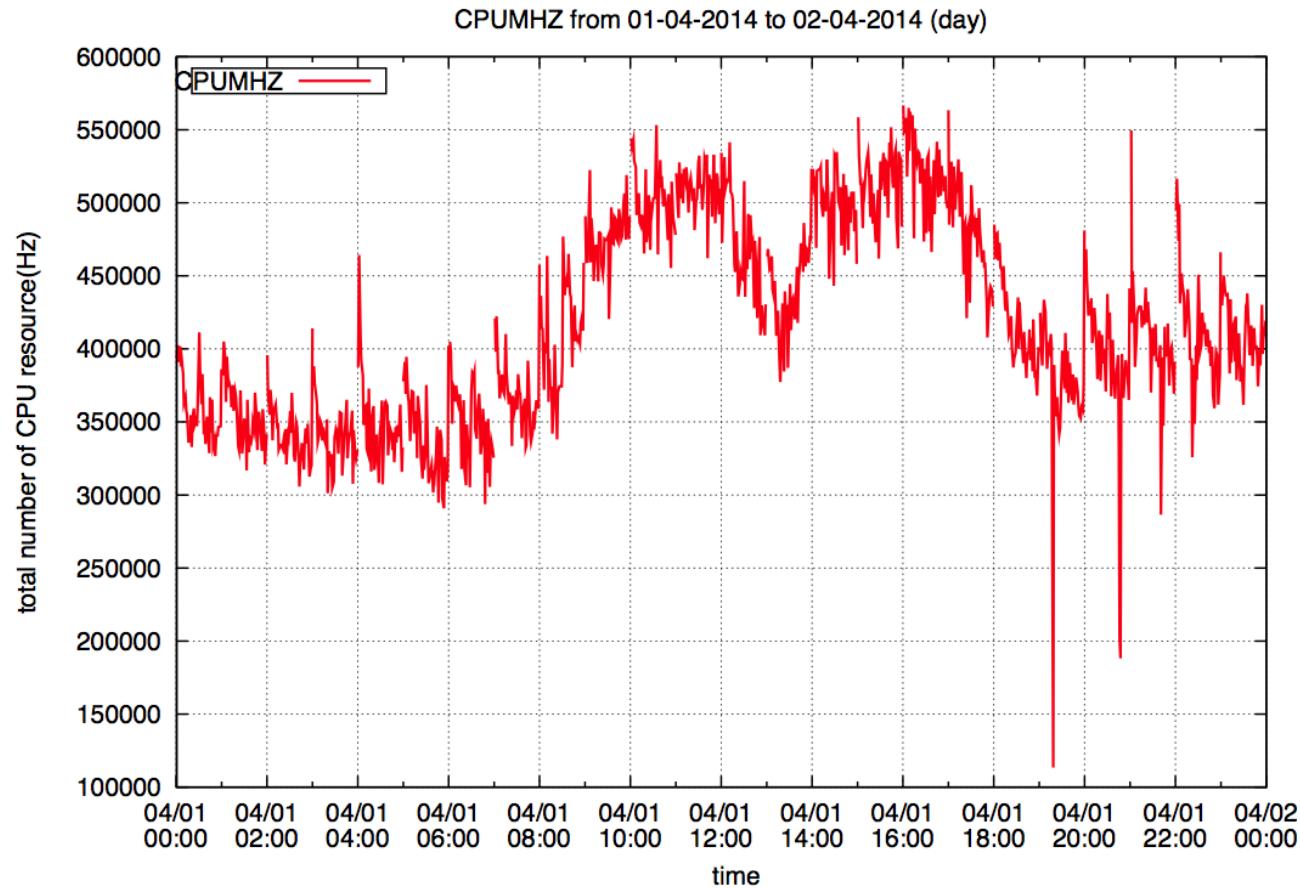
Virtual Machine placement



■ Prédictions

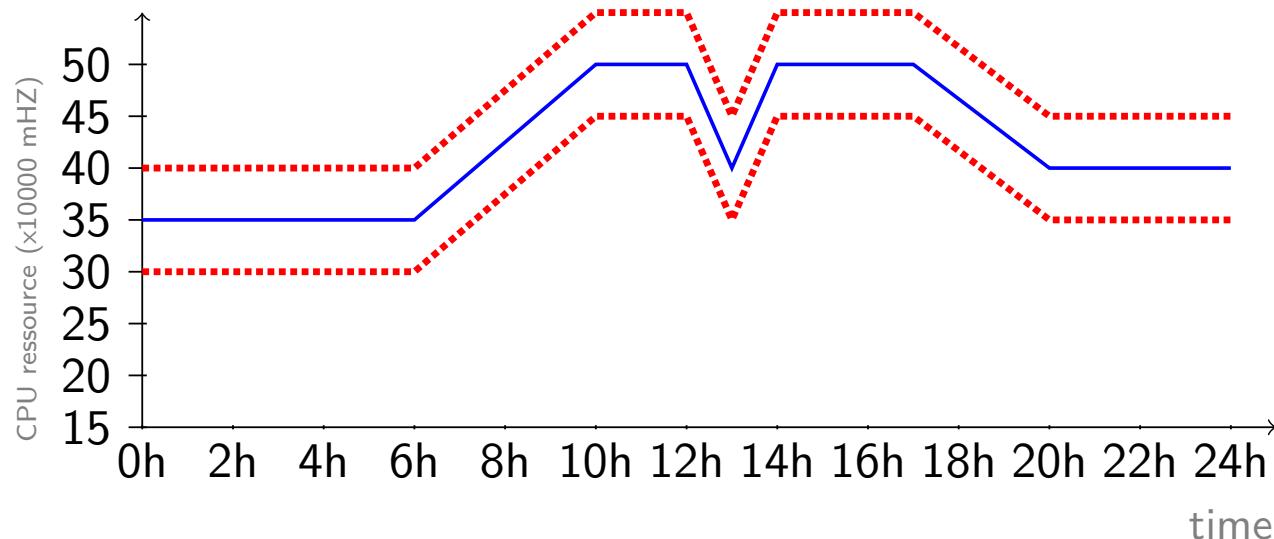


Dynamic model : Server load evolution



11 mars 2015

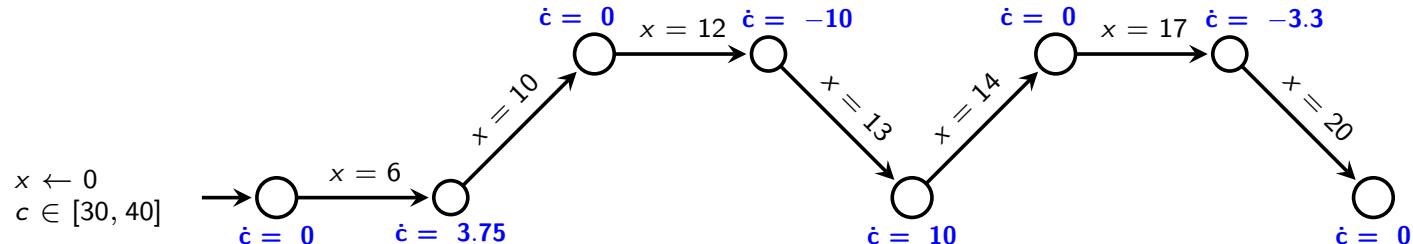
Dynamic model : pattern recognition



- Smoothing the curves.
- Clustering to characterize typical behaviours (working day, weekend, public holiday).
- For each cluster capture and model (set of constraints, e.g. number of peaks) and use it as a curve generator.

20 mars 2015

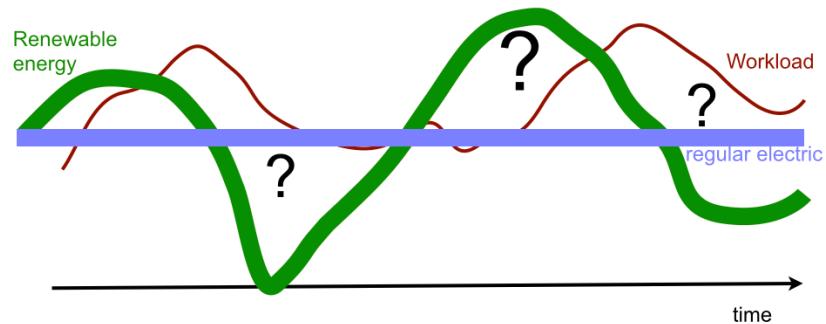
Dynamic model : Prediction model



Prediction automaton

- The states and transitions of the automaton represent the variations in the server load.
- The automaton will further be transformed into a constraint network using a constraint automata.
- Each state of the automaton represent a bin packing problem.
- At a given date t , we solve the system for the corresponding state of the automaton, and eventually few states later : The chosen solution should take into consideration future events.

20 mars 2015



Infrastructures

Pour travaux de recherche



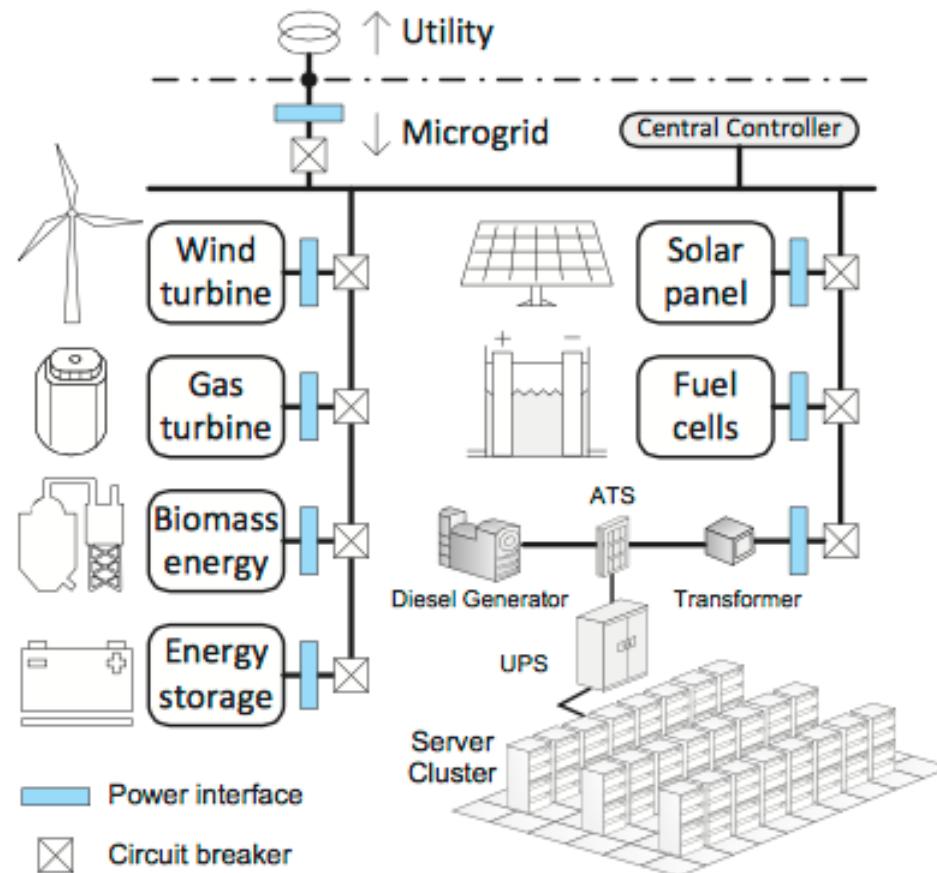
Politique de Google : Energy Broking



Greening our grid
through renewable energy purchases

Google™

Architecture Classique



https://www.usenix.org/system/files/conference/icac14/icac14-paper-li_chao.pdf

Parasol



(a)

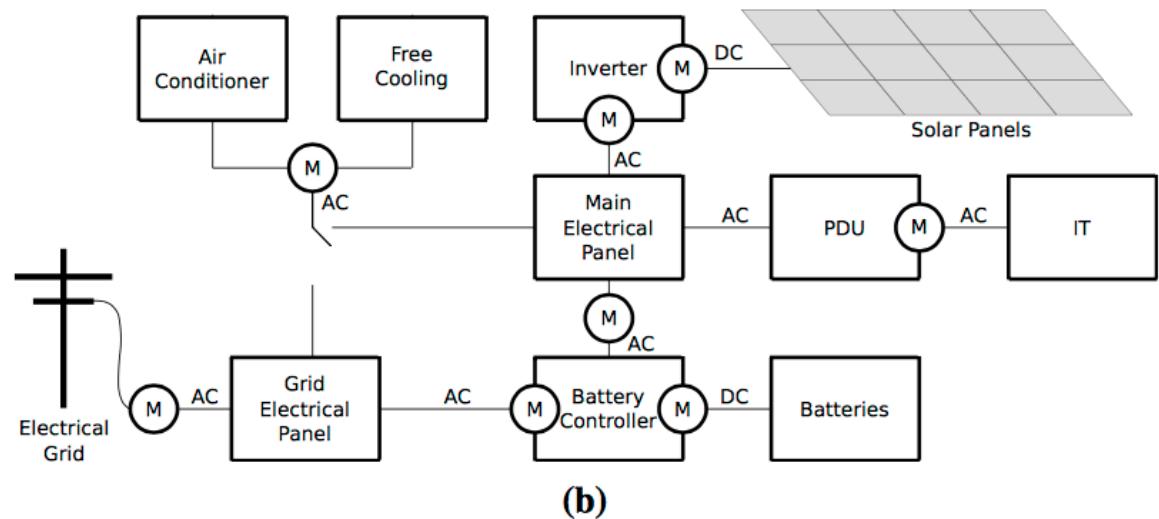


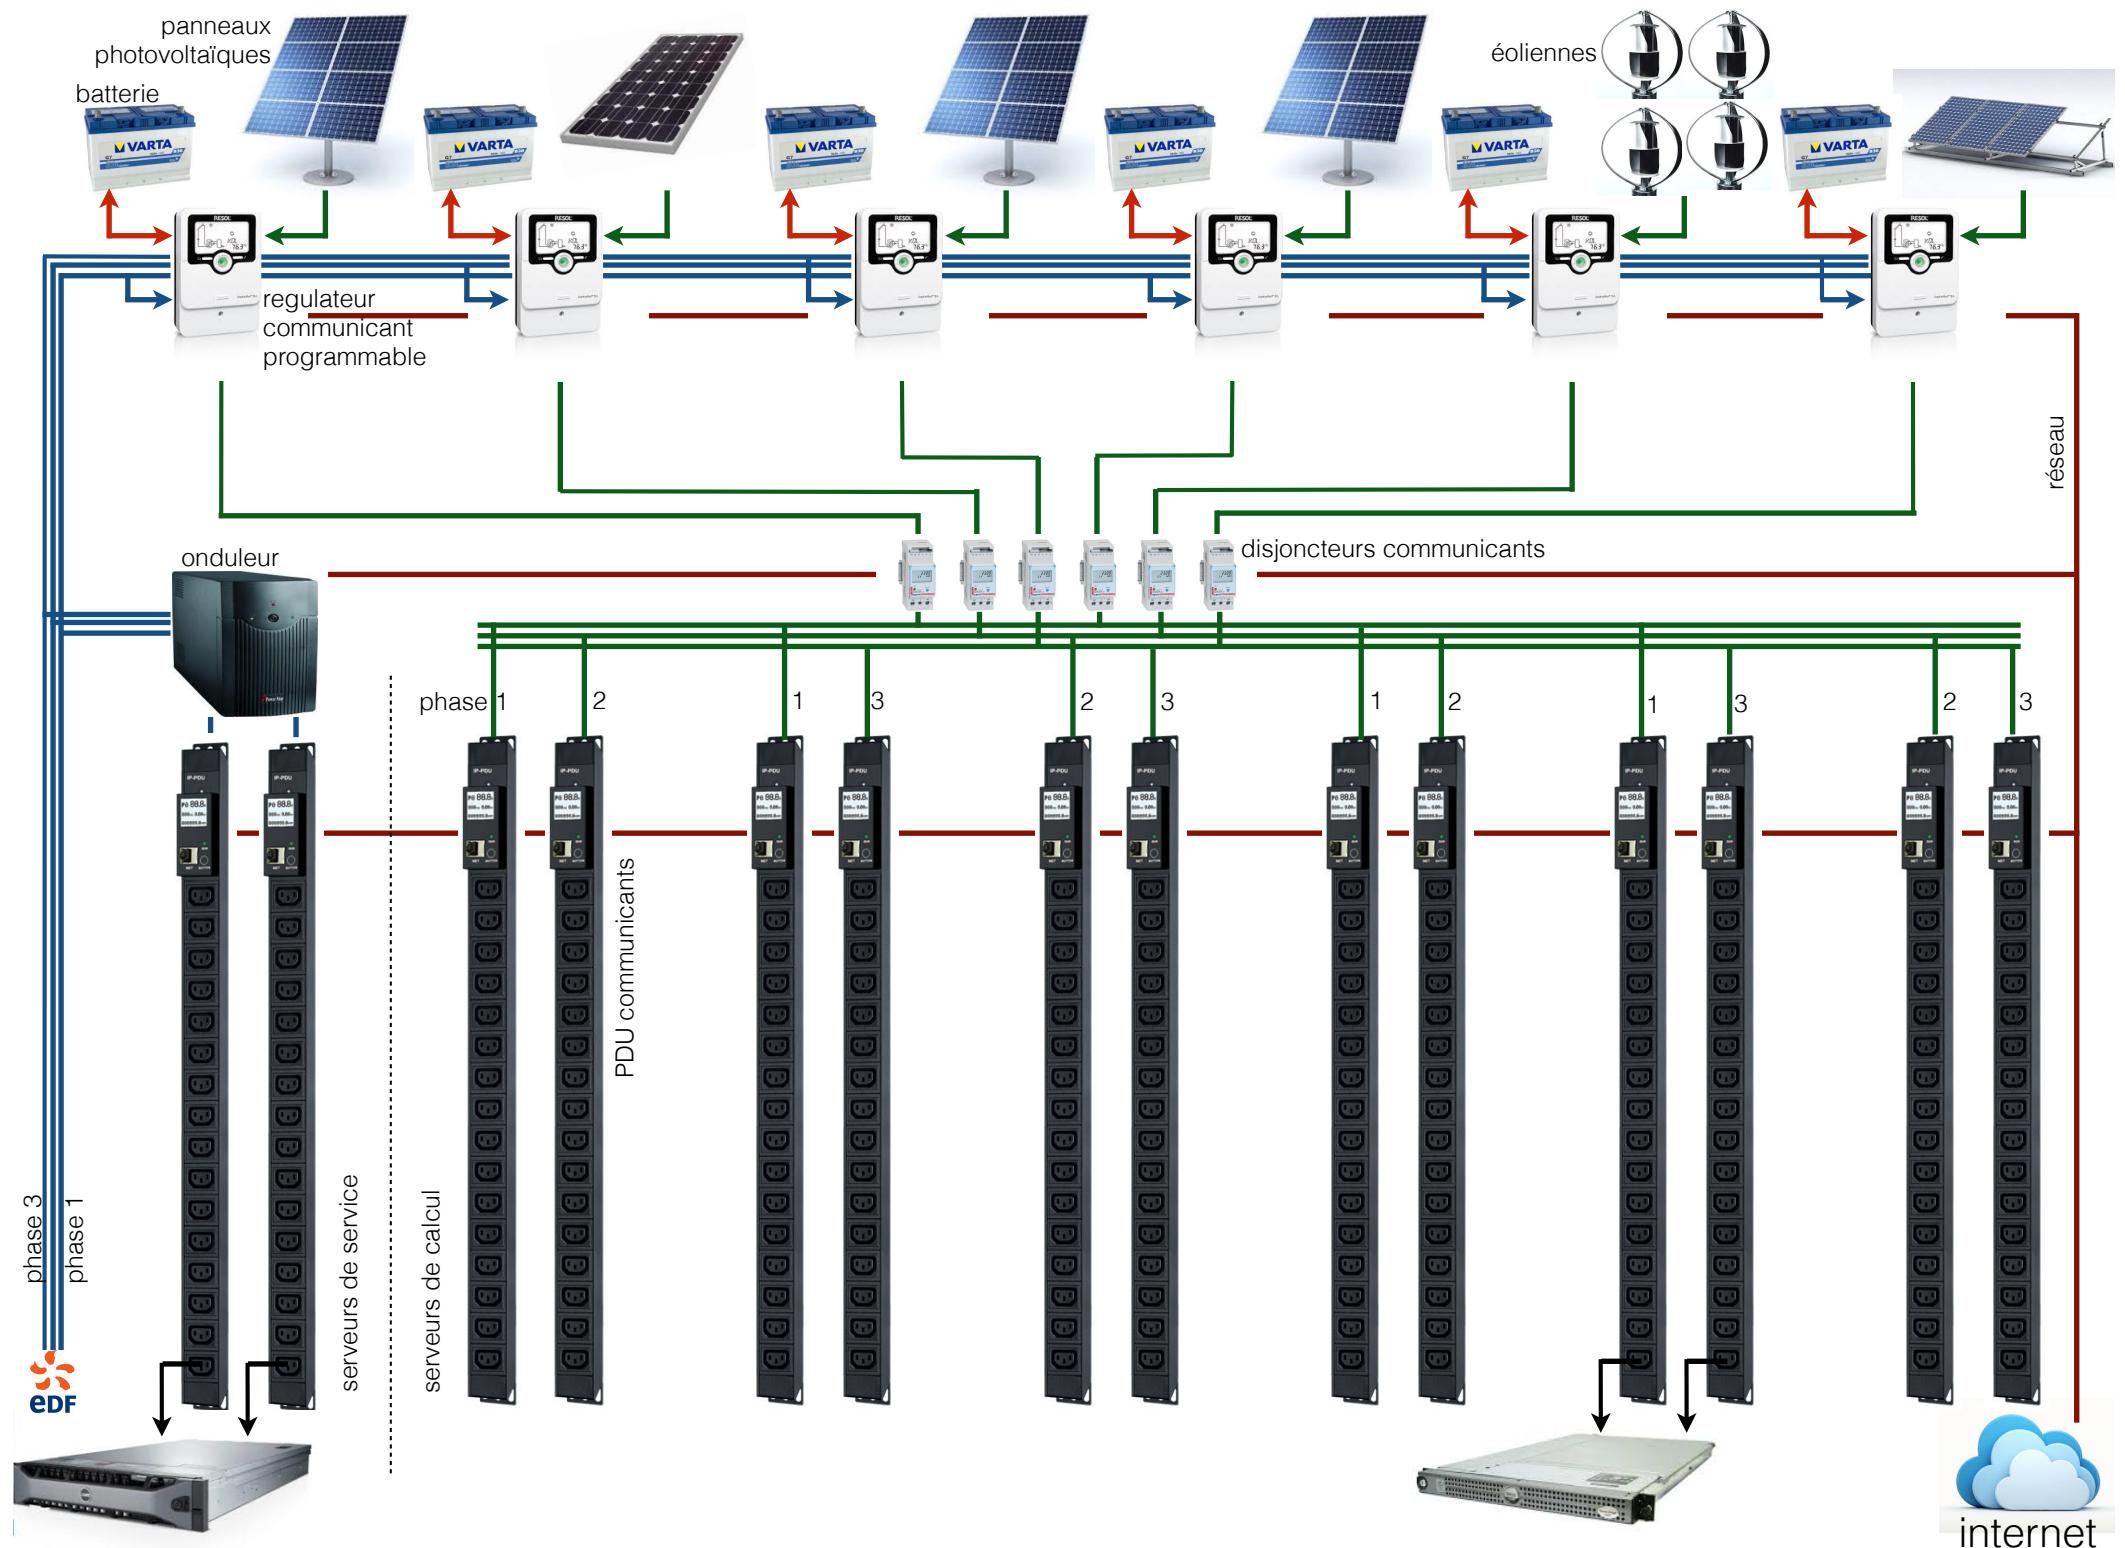
Figure 1. (a) Outside view of Parasol. (b) Parasol's power distribution and monitoring infrastructure.

<http://www.cs.rutgers.edu/~ricardob/papers/top-picks14.pdf>



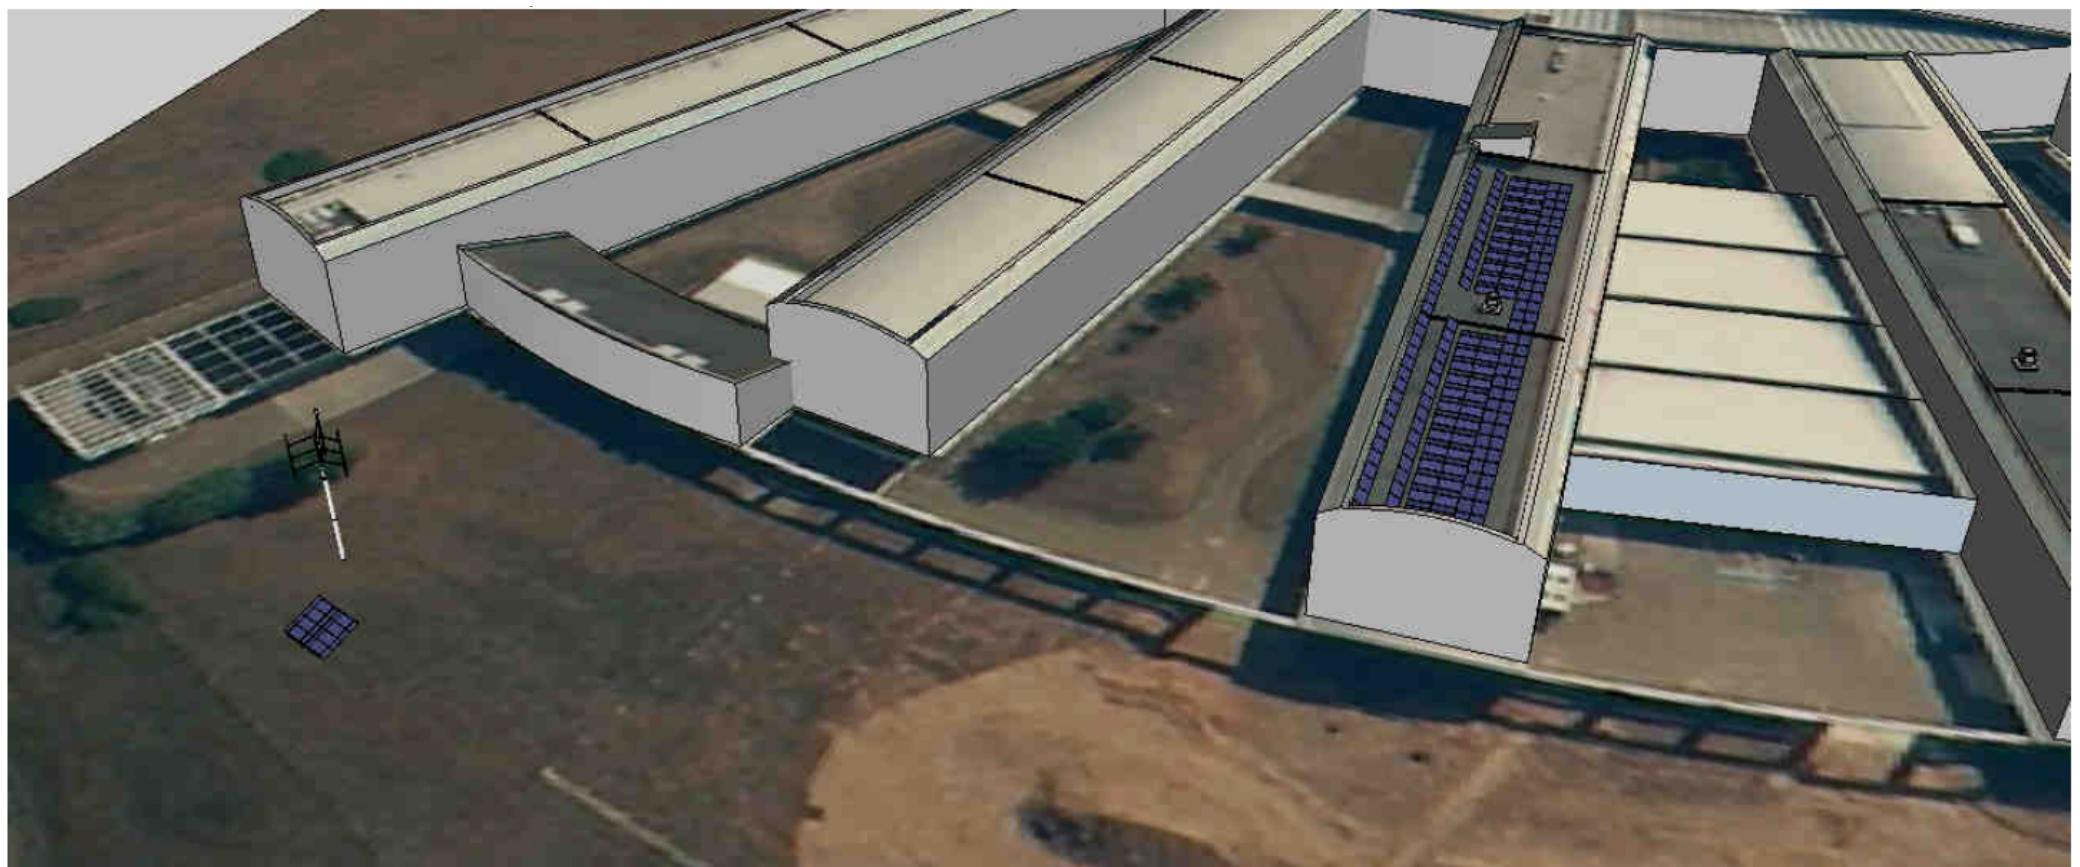
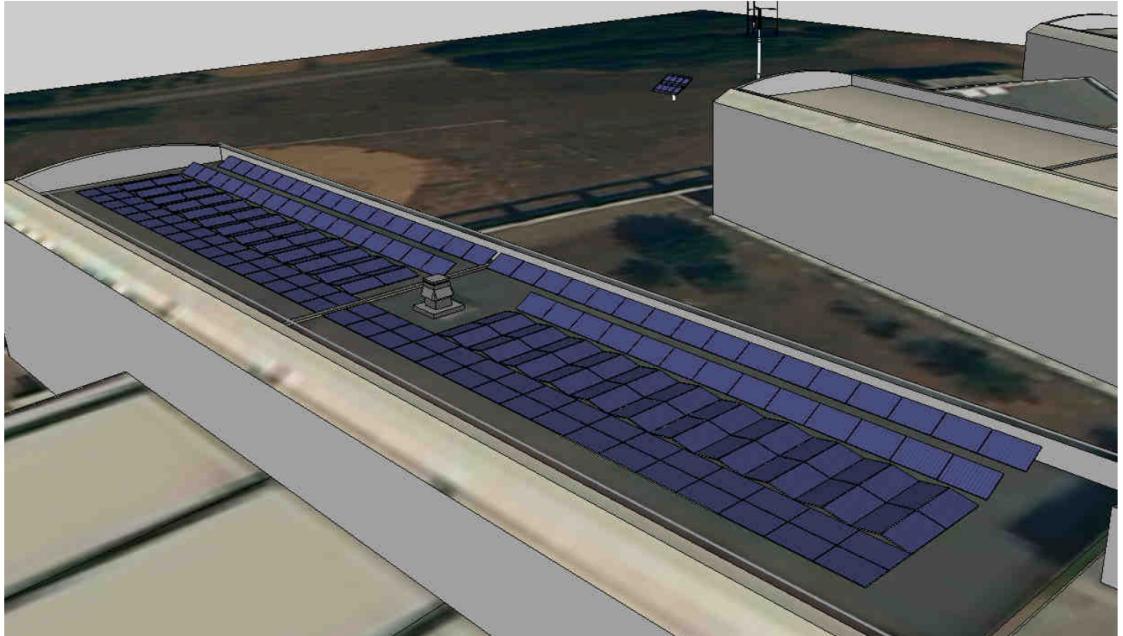
SeDuCe :
Une infrastructure recherche
unique pour les travaux sur les
centre de données a faible
emprunte énergétique

2016-2020





Seduce

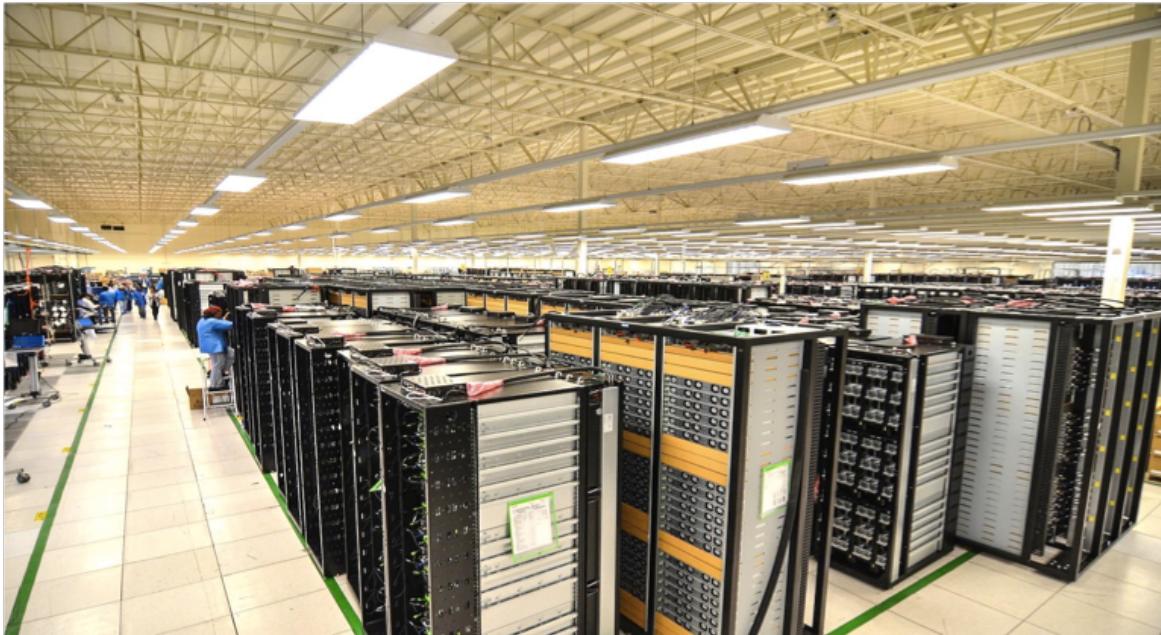


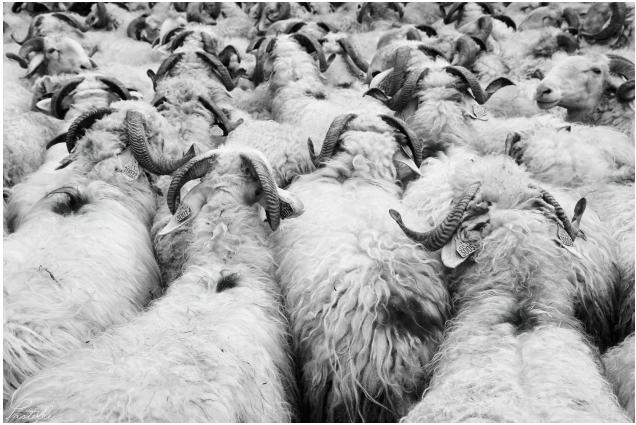


Seduce



Winter Fell Server





2017-...



Tranhumance :

Follow the sun, wind ...



Luc Olivier

Follow-the-renewables

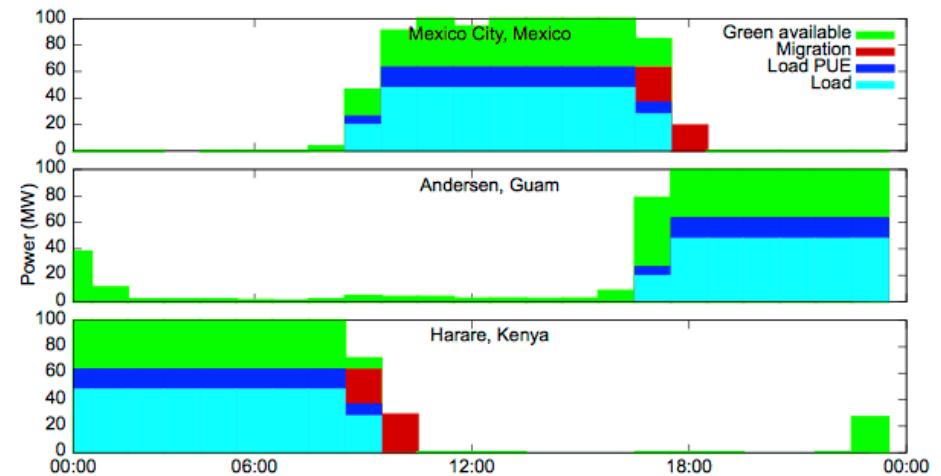
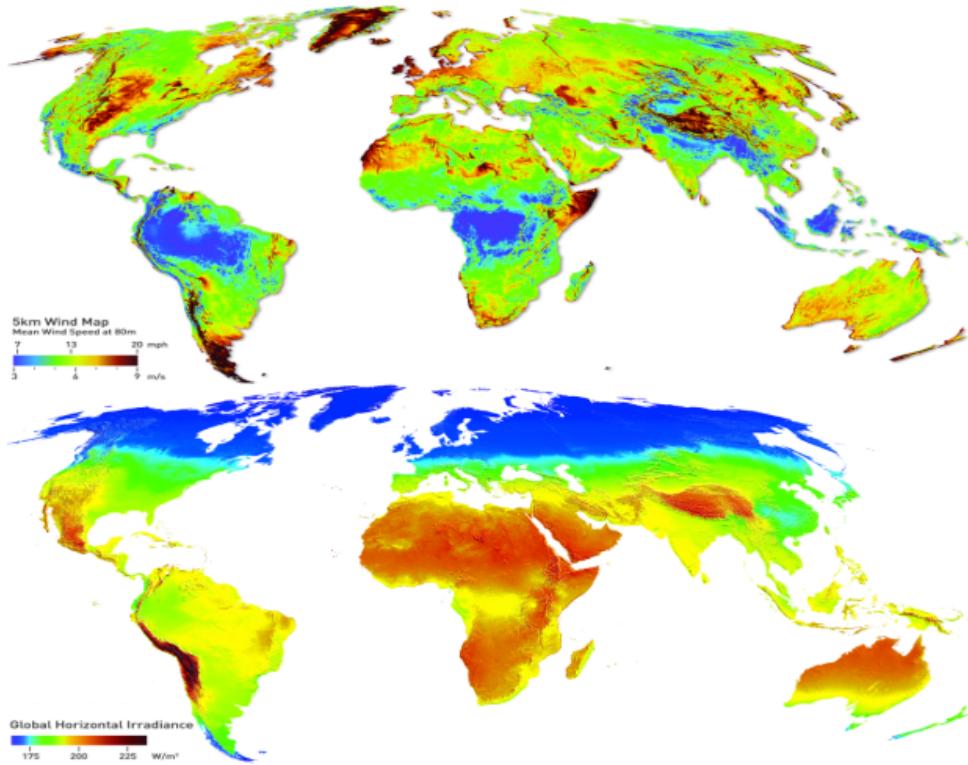


Figure 15. Load distribution to achieve 100% green energy without energy storage. The green area shows the amount of green energy produced, light blue shows the computing load, dark blue shows the PUE energy overheads, and red shows energy used by migrations.

<http://www.cs.rutgers.edu/~ricardob/papers/icdcs14.pdf>



