ANR Datazero

DATAcenter with Zero Emission and RObust management using renewable energy

October 1st, 2015 — December 31st 2019

as green as possible

Jean-Marc.Pierson@irit.fr

PROJET FINANCÉ PAR L'ANR - PROJECT FUNDED BY THE ANR

IRIT CNRS - INPT - UPS - UT1 - UT2J

EATON Powering Business Worldwide

Laplace

femto-st SCIENCES & TECHNOLOGIES

An innovative datacenter model

- Adapting the *IT load* to the Power available

and

- Adapting the *Power* to the incoming IT load

while avoiding unnecessary operations and materials and using a mix of energy sources

Target and Outcome

- Small and middle size datacenters Cloud / Virtualization (1000 m^2, 1 MW)
- Rethink datacenter with alternative power

RETHINK

data ZERO

as green as possible

First, we need an electrical and IT connection…

- Traditionally, a lot of redundancy is present to prevent shortage of power feeds and absorb peak IT demands
- *How to limit the redundancy while keeping safety?*
  - Use only critical redundant material
- How to limit the losses in the electricity conversions and transports?
  - Connect directly Direct Current Busses with onsite electricity production

data ZERO
as green as possible

“Classical” Electrical Schema: Up to N+1 (Green+Grid)

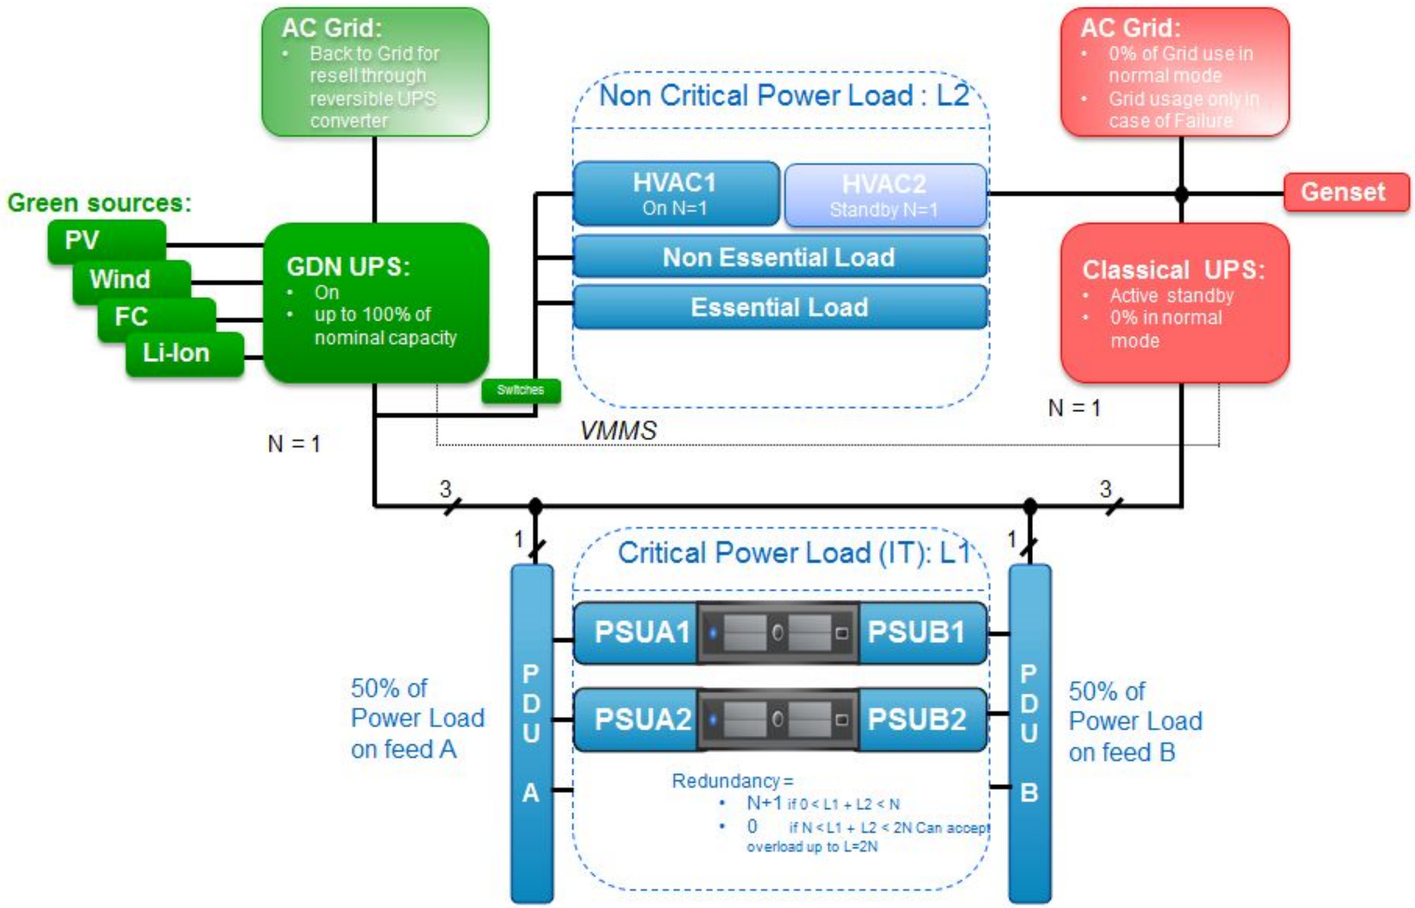

data ZERO
as green as possible

Innovative electrical Schema: Up to 2N Green

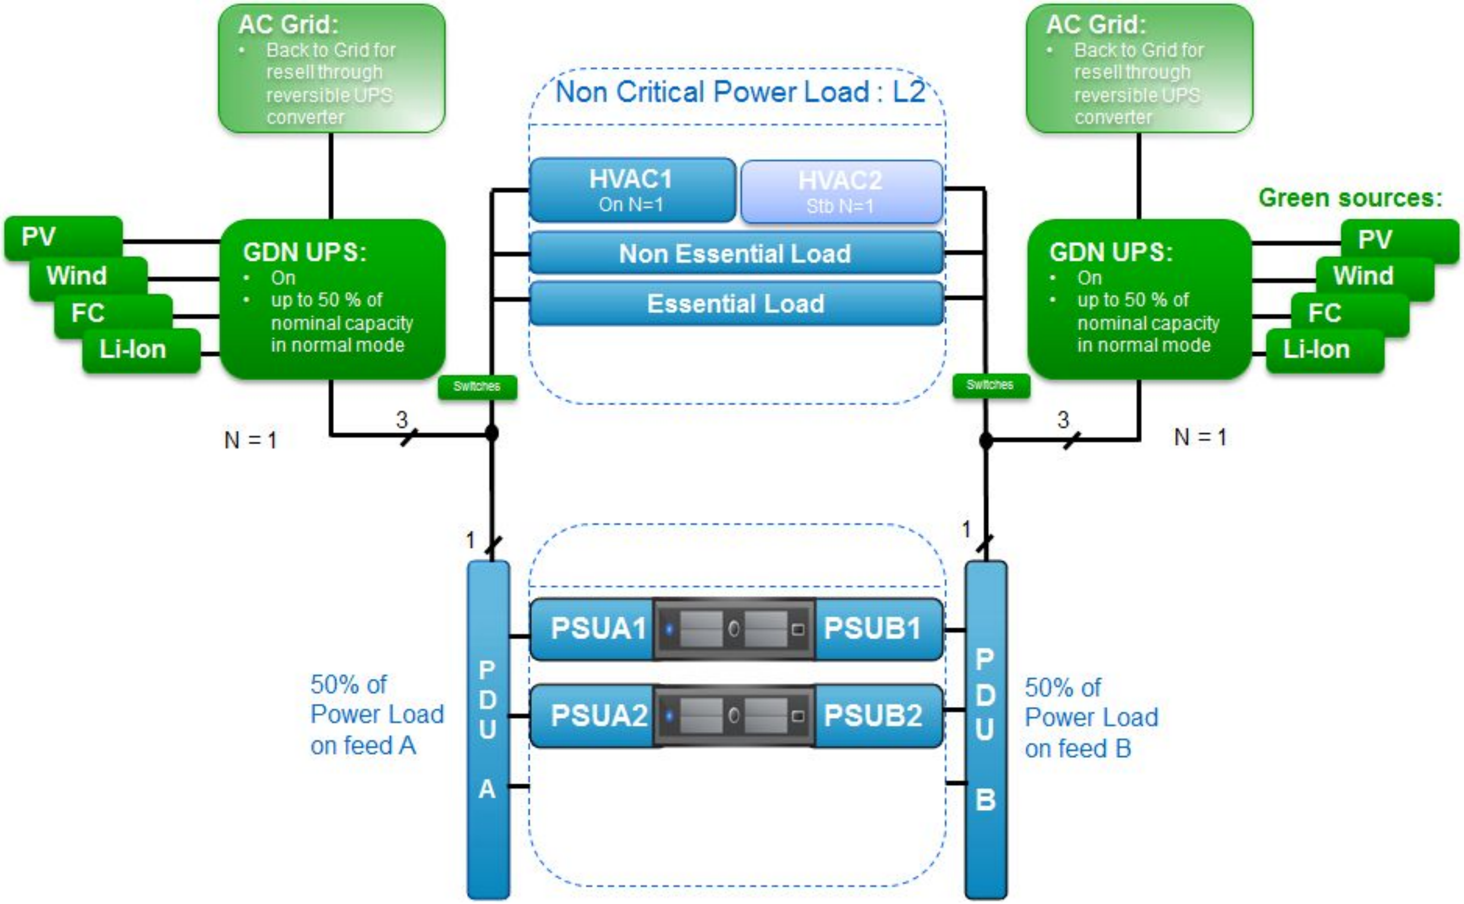

Second, we need a mix of energy sources

- How to ensure a reliable and emissions-free power supply?

- Additional challenges:
  - Sources are diverse, each with their own characteristics
  - How many of each type should we install?
  - Then how to operate them efficiently?
  - How to handle uncertainty from renewable generation?
  - What is the impact of sources ageing?

data ZERO
as green as possible

The components form a microgrid

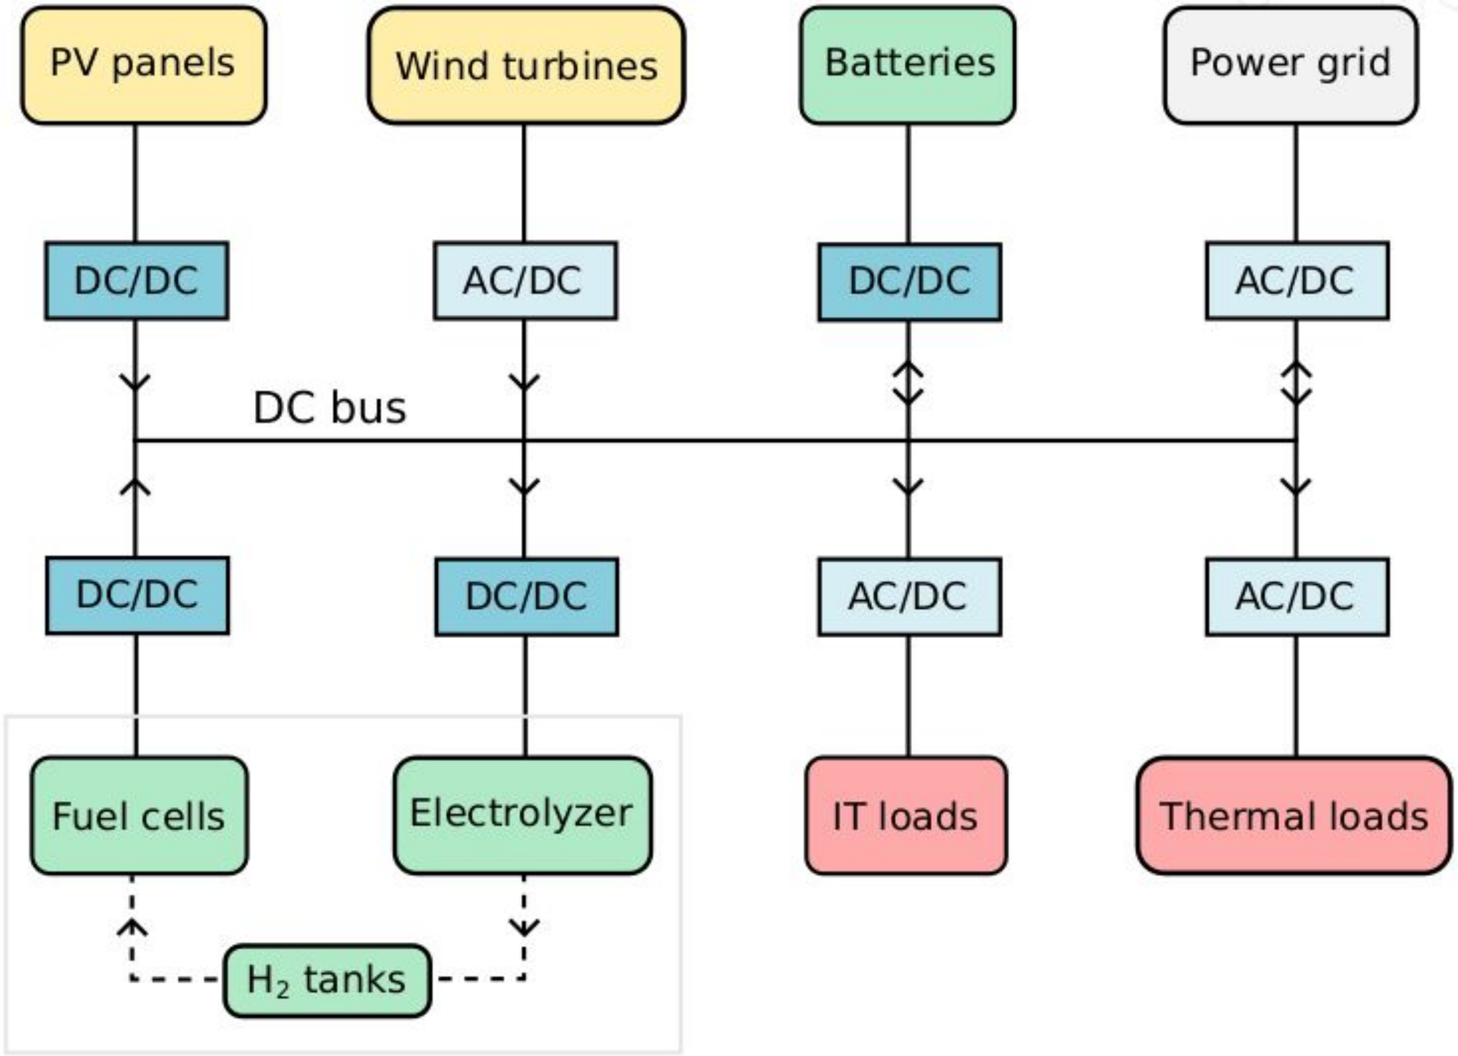

Model example: fuel cell

- Converts hydrogen to electricity
- Performance is degraded with time and use
- Ageing model based on experimental data

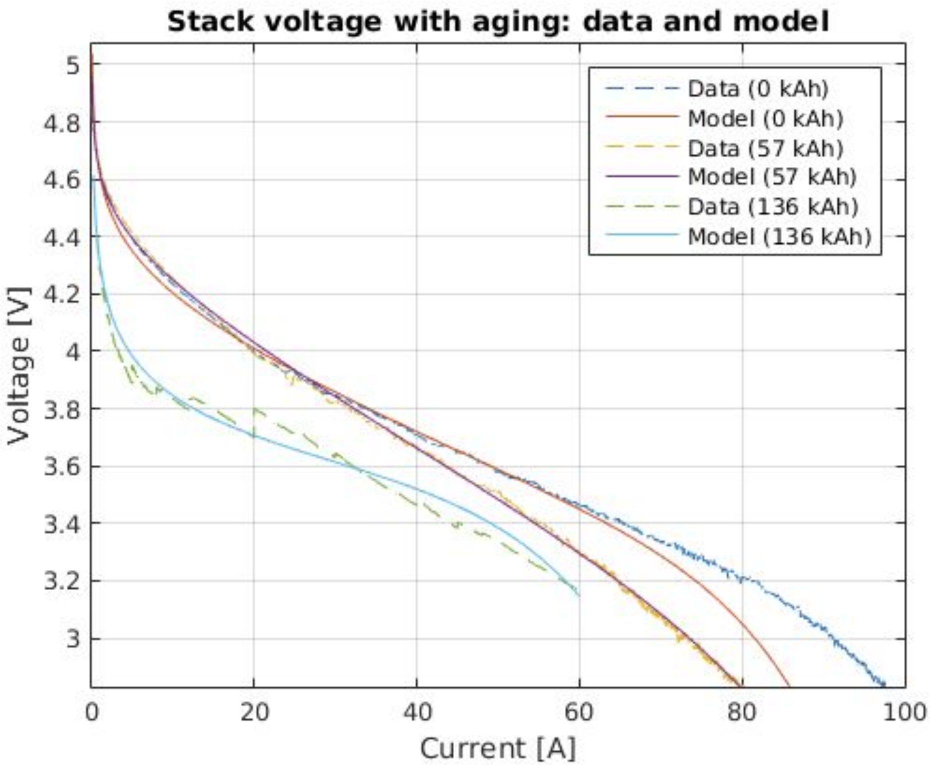

Then, we need computers in a datacenter

- How do they consume power? What is the link between their usage and their electricity consumption ?
- What is the impact of the operating system or the hardware tuning ?
- What is the impact of using Cloud Computing technologies ? What is the costs of hypervisors ? What is the cost of migration of tasks on the servers ?

data ZERO
as green as possible

We need to profile the IT

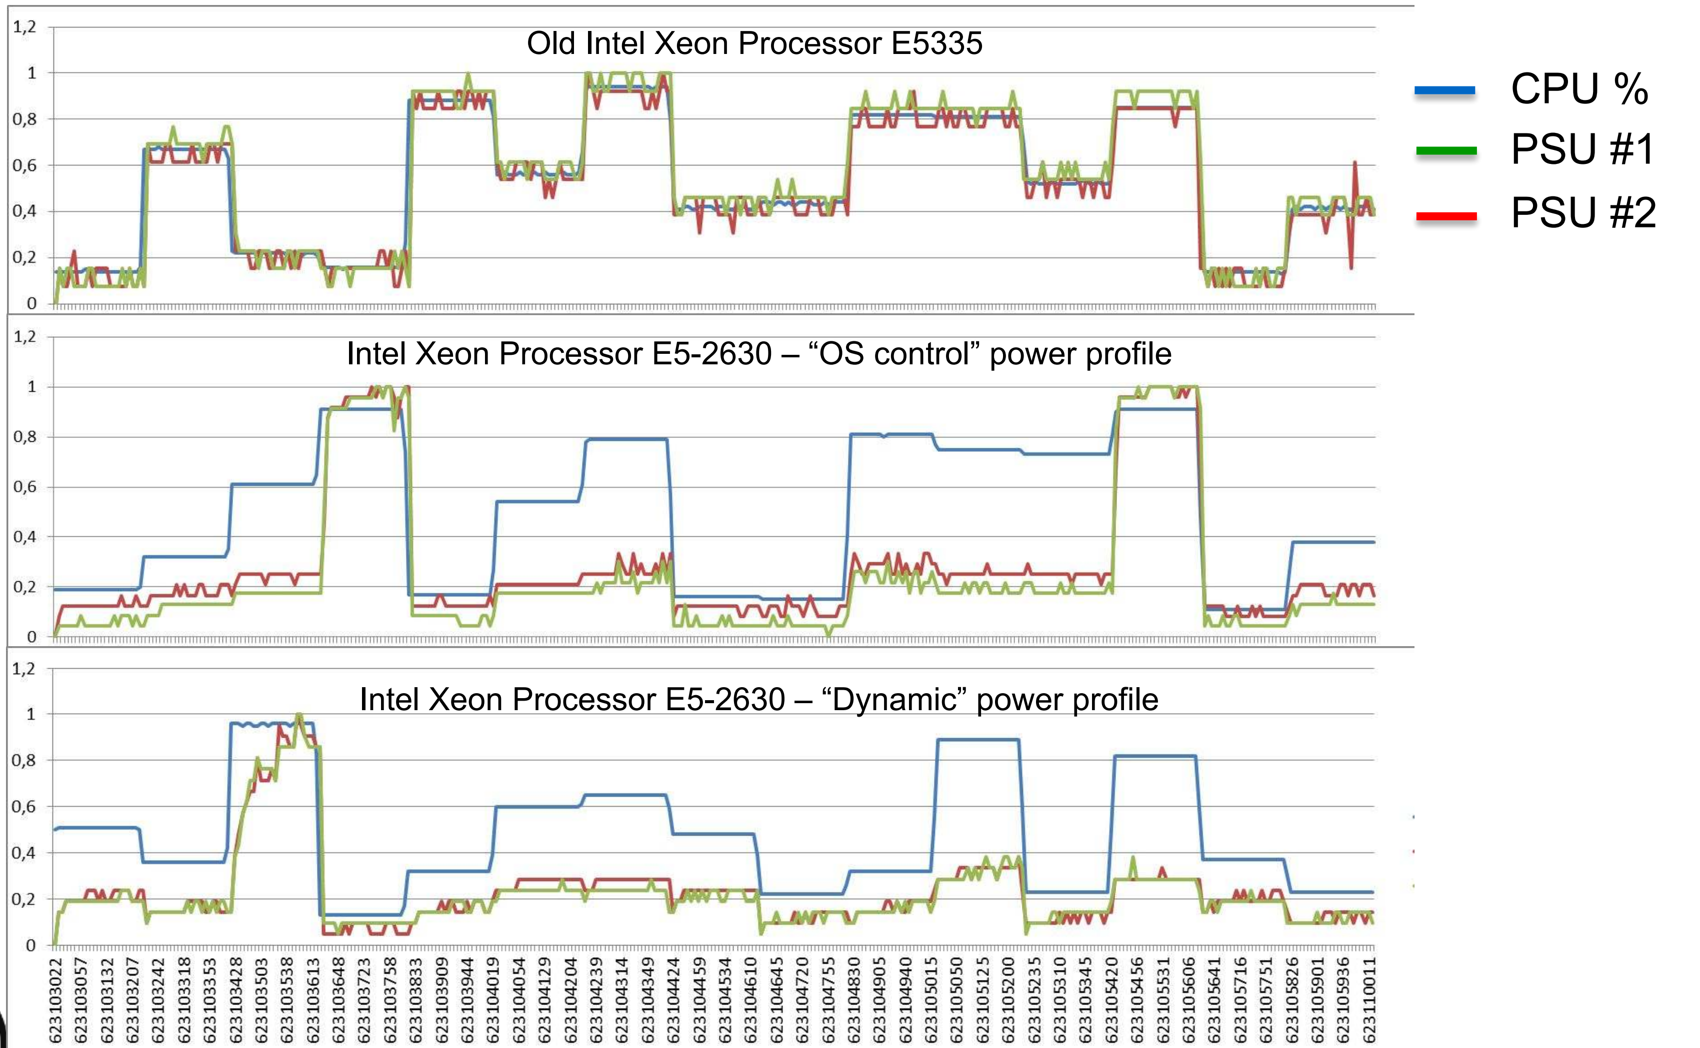

as green as possible

Impact of virtualization and energy saving parameters of servers

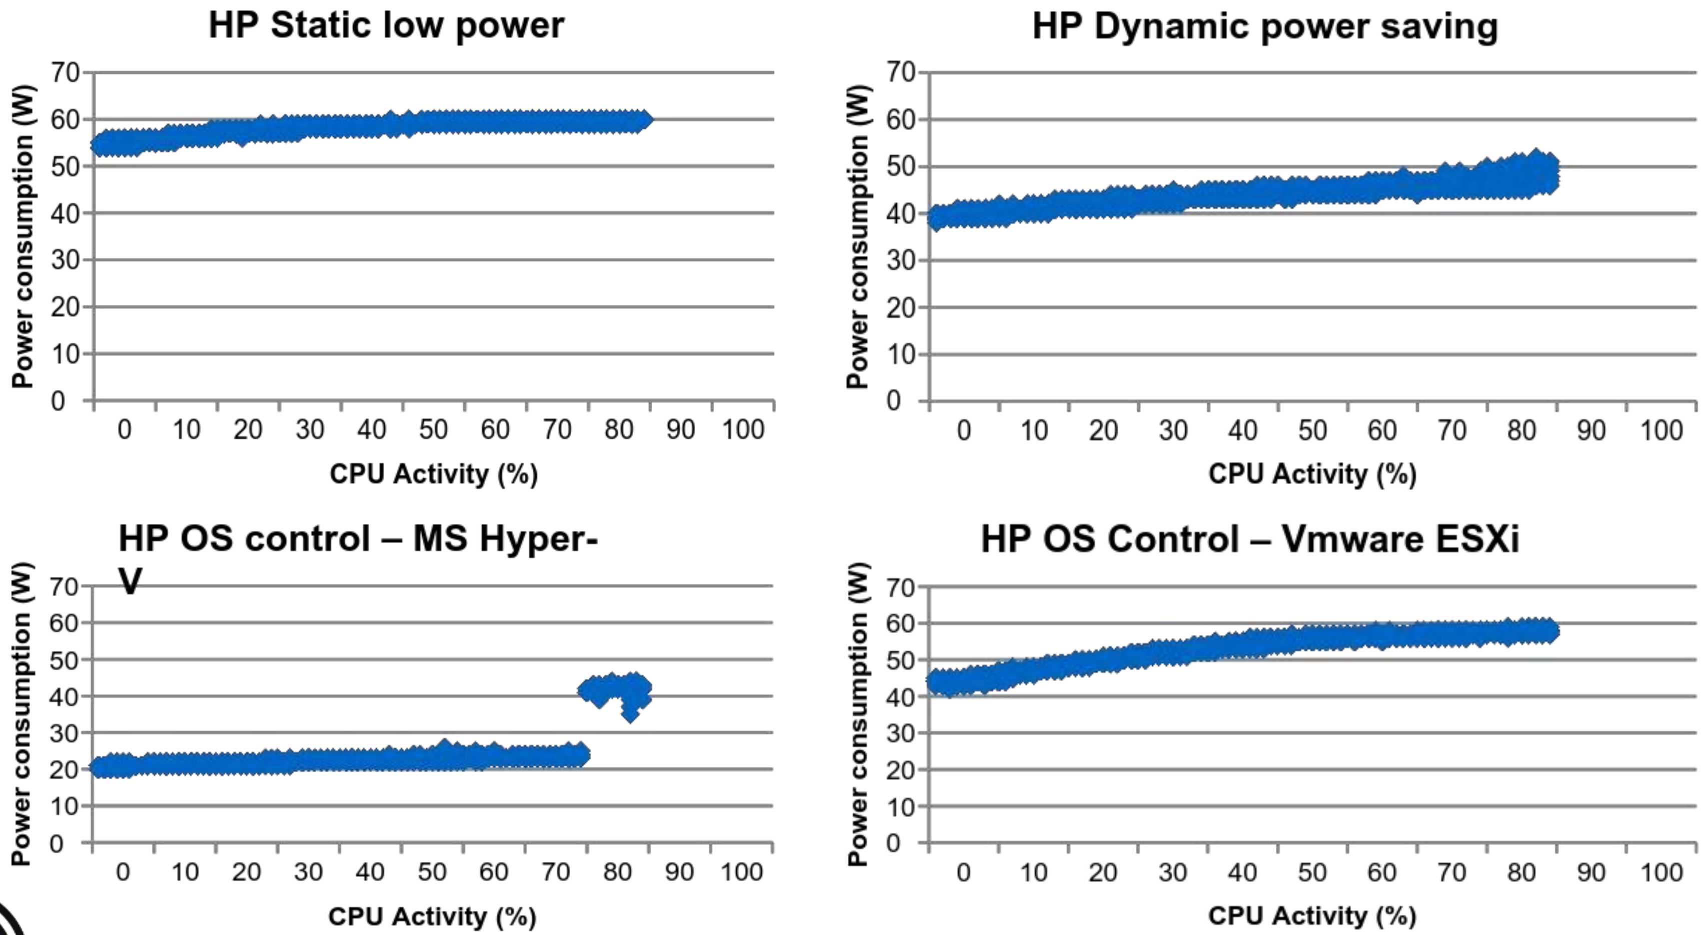

as green as possible

Impact of virtual machine migration

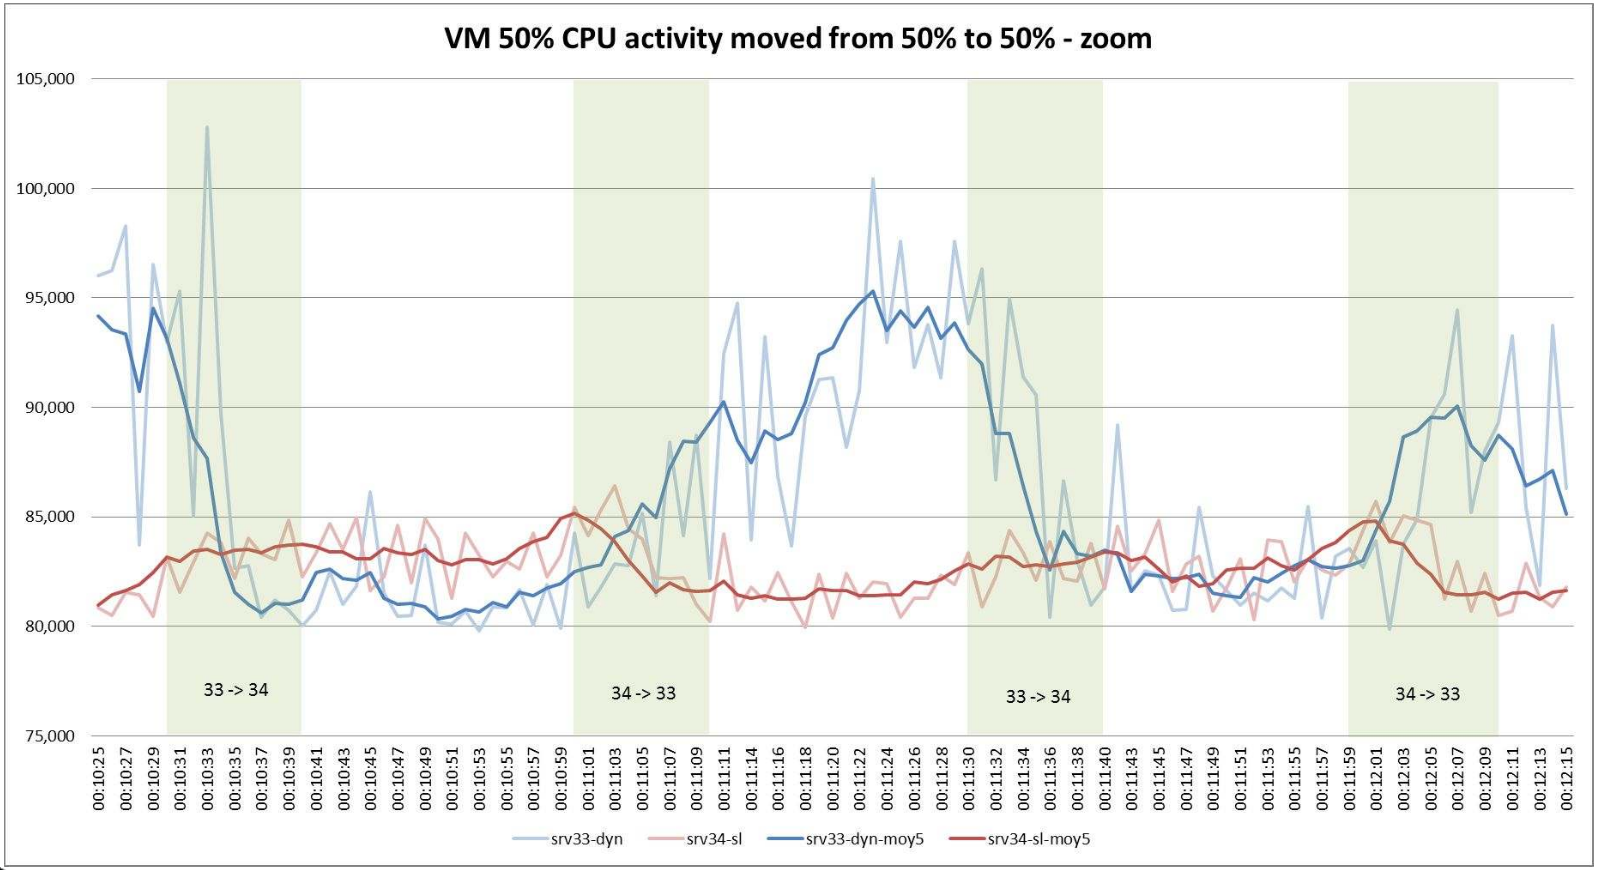

Then, we need applications on these computers!

- Does an application consume power ?
  - No, they consume resources (CPU, memory, ...), that translates to power consumption
- And applications consume power differently depending on their (obtained or targeted) performances

We need to know the applications

An application is modelled by:

- Submission information
  - Time of arrival
  - Requirements
- Resource consumption over time
  - Processor, memory, IO, network
- Policy related informations (scheduling priority, ...)
- Context information
  - Degradation level

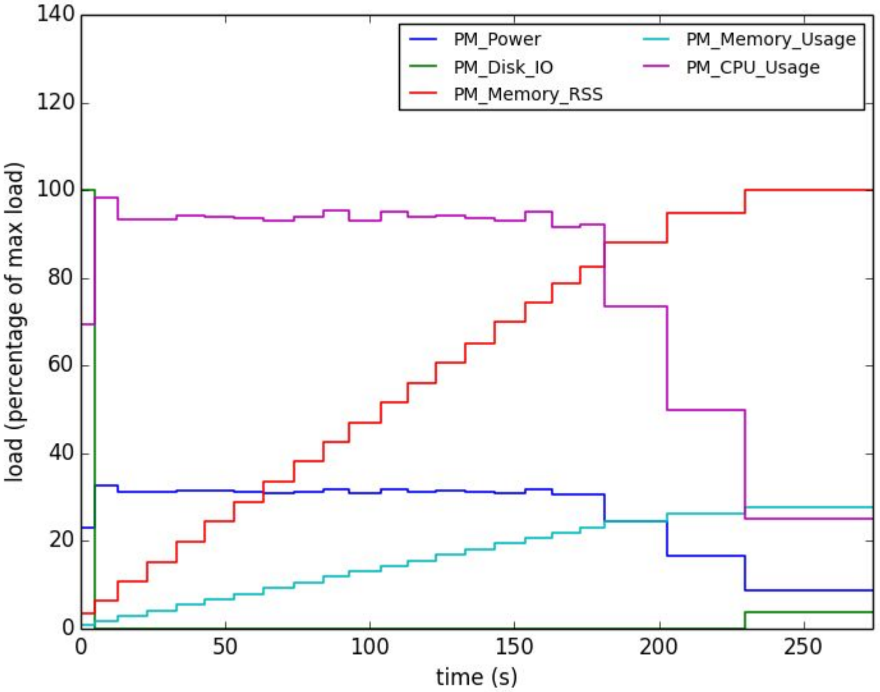

Toolchain to abstract applications

```
    write(fd, a1, 4);
}
close(fd);
unlink("/dev/shm/test_file");
for(i=0; i<5000; i++) {
  usleep(1000);
  a=sqrt((double)i)*sin(a+cos(a))
  a=sqrt((double)i)*sin(a+cos(a))
  a=sqrt((double)i)*sin(a+cos(a))
  a=sqrt((double)i)*sin(a+cos(a))
```

```
#cpumodel Generic Intel
#timestamp userLoad systemLoad vSize RS
1464856255 0.000000 0.000000 11173888 2
1464856256 0.000000 0.000000 11173888 2
1464856257 0.000000 0.000000 11173888 2
1464856258 0.000000 0.000000 11173888 2
1464856259 0.000000 0.000000 11173888 2
1464856260 0.000000 0.000000 11173888 2
1464856261 0.000000 0.000000 11173888 2
1464856262 0.000000 0.000000 11173888 2
```

```
<resourceConsumptionProfile>
    <resourceConsumption>
        <referenceHardware>Generic Intel</re
        <duration>PT11S</duration>
        <behaviour name="userLoad">
            <value>0.000000</value>
        </behaviour>
        <behaviour name="systemLoad">
            <value>0.000000</value>
        </behaviour>
```

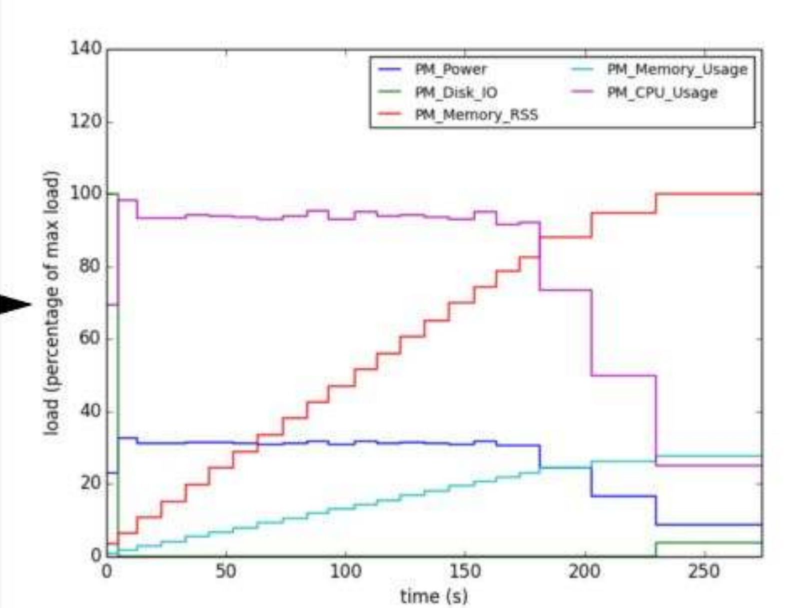

An *application* is a set of profiles depending on context

Multiple profiles

Cloud application service

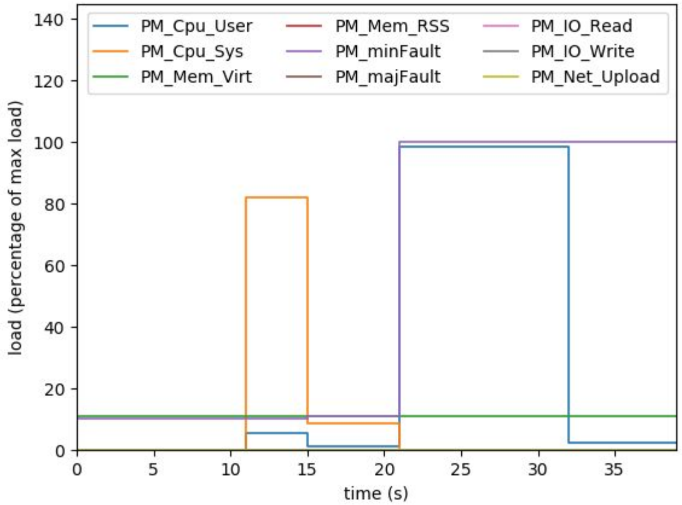

Degraded application on slower hardware

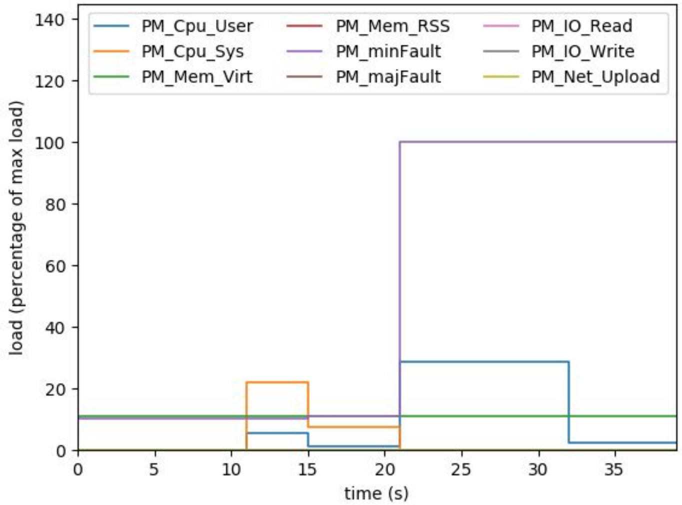

Same application on slower hardware

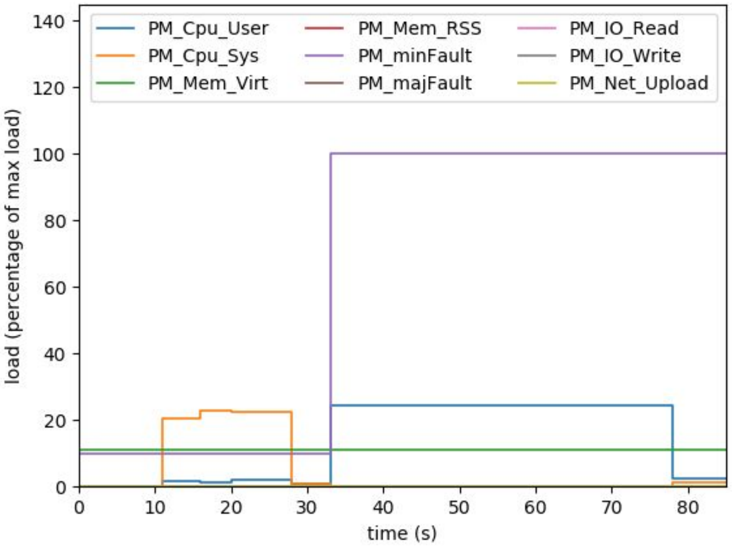

An *application* is a set of profiles depending on context

- Hardware
- Quality of service

Context is included each XML file

<degradationLevel>none</degradationLevel>

We need to negotiate

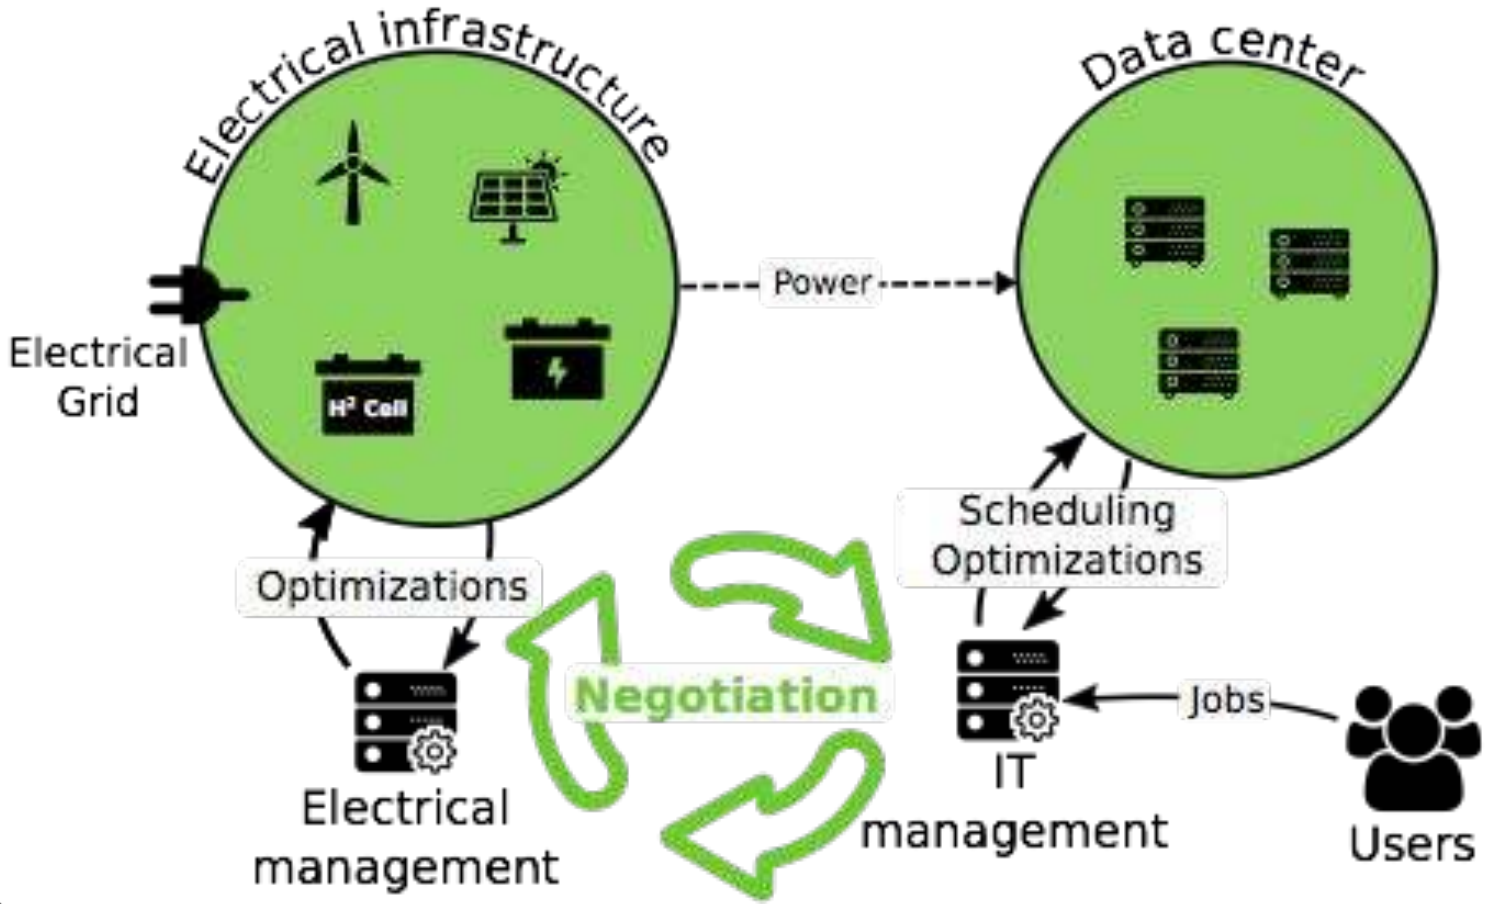

data ZERO
as green as possible

Finally...

- We have electrical connections
- We have some (renewable) power sources
- We have some computers, with their applications
- How to make all these elements coexist ?
- Traditionally: consider all the elements at once, and optimise the global problem !

➽ We believe in an alternative way: the negotiation !

data ZERO
as green as possible

Negotiation Module

Example: 1st approach based on attraction concept:

Representing the interest of a system for a proposal

- interest for the electrical part to produce 2000W from 11:00 am to 1:00 pm?
- interest for the IT part to execute a task at 11:00 am?

Abstracted to a float number in [-1, 1]

✓ - 1 → absolute reject (a very high penalty)
✓ +1 → a very high attraction
✓ 0 → neutral
✓ intermediate values available

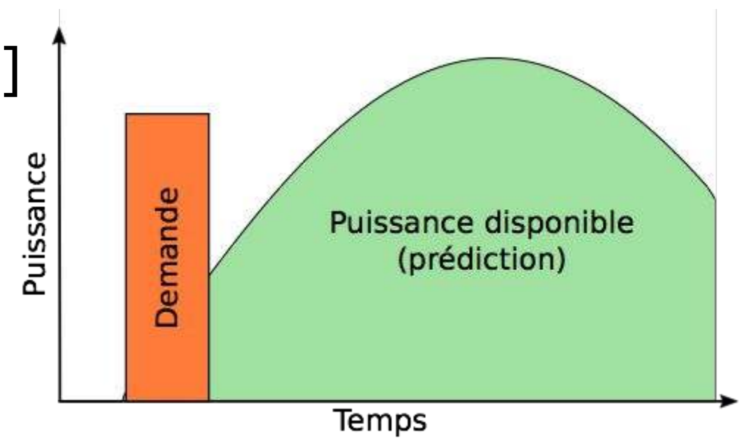

Attraction values depend on the possibilities

Negotiation: IT attraction

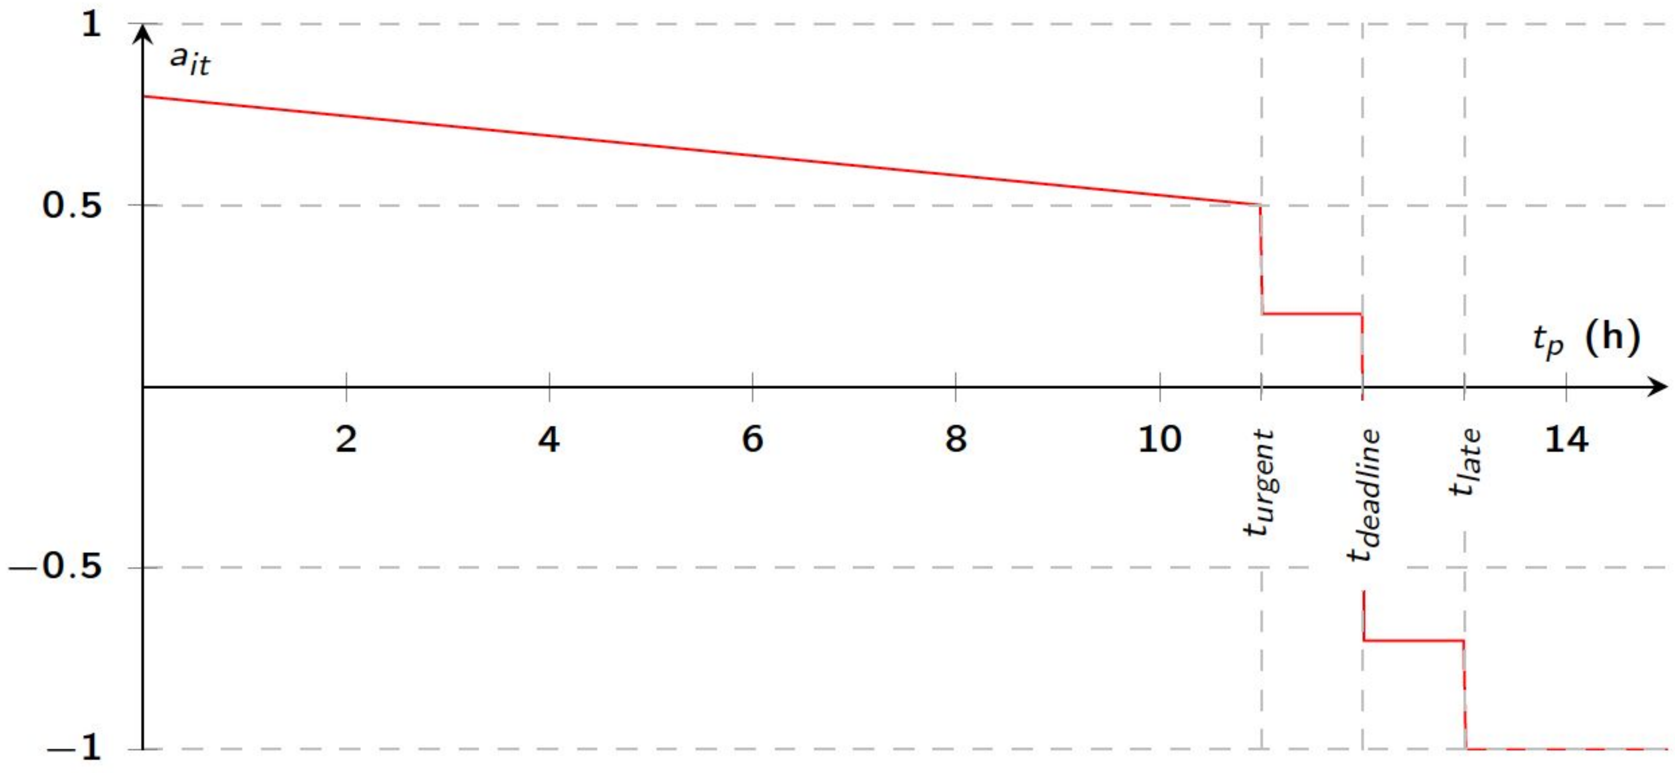

- Flexibility is given with the deadline

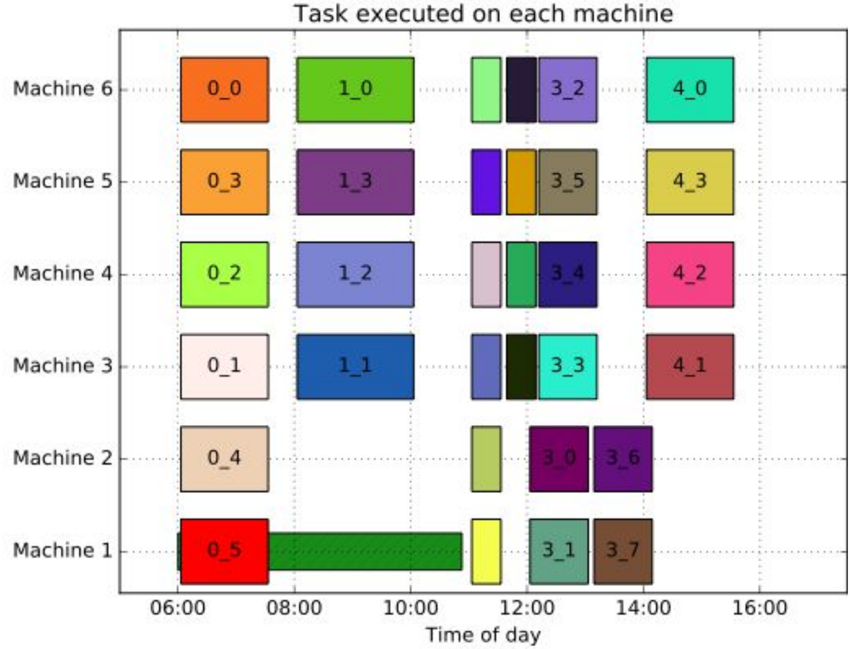
Task executed on each machine
Machine 6
Machine 5
Machine 4
Machine 3
Machine 2
Machine 1
0_0
1_0
3_2
4_0
0_3
1_3
3_5
4_3
0_2
1_2
3_4
4_2
0_1
1_1
3_3
4_1
0_4
3_0
3_6
0_5
3_1
3_7
06:00
08:00
10:00
12:00
14:00
16:00
Time of day

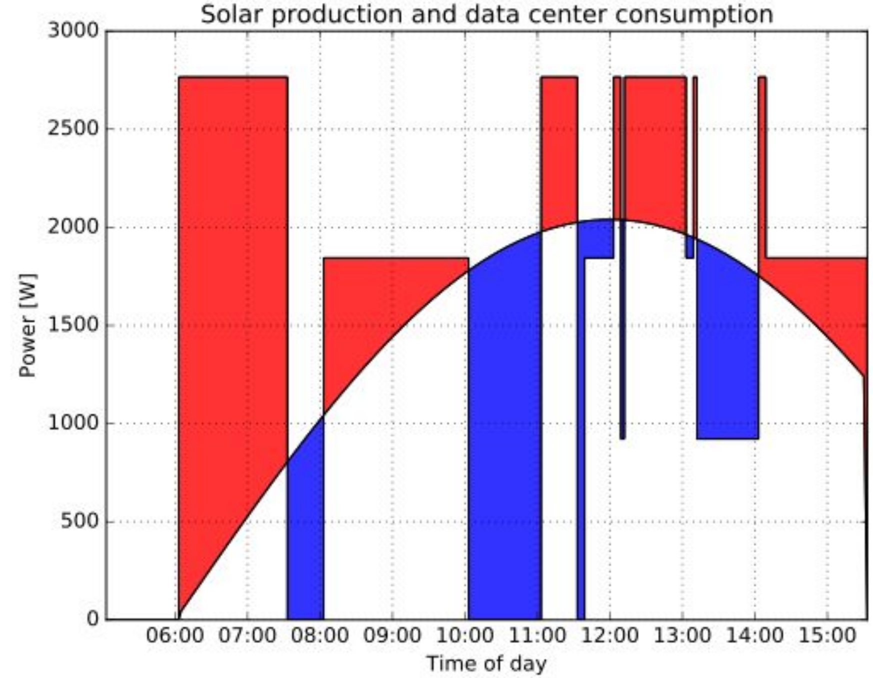
Solar production and data center consumption
Power [W]
3000
2500
2000
1500
1000
500
0
06:00
07:00
08:00
09:00
10:00
11:00
12:00
13:00
14:00
15:00
Time of day

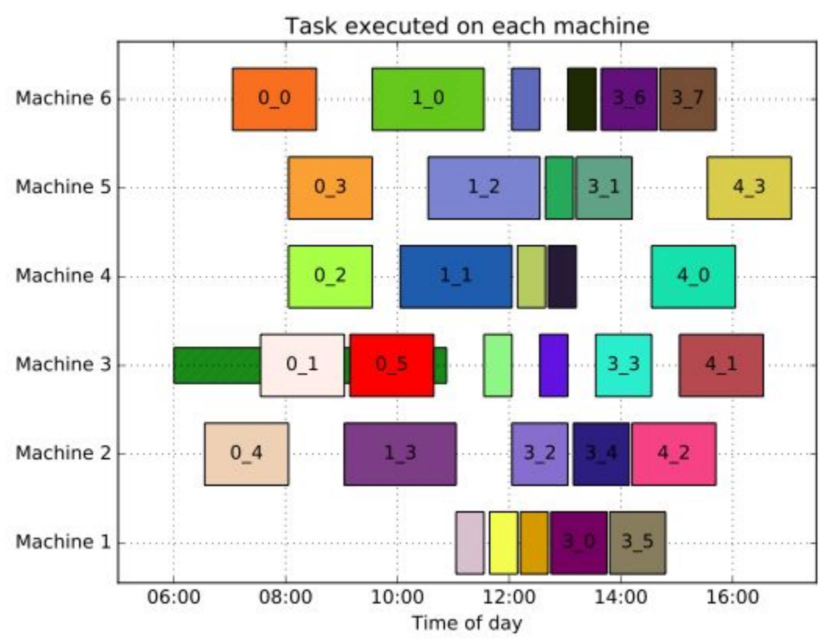
Task executed on each machine
Machine 6
Machine 5
Machine 4
Machine 3
Machine 2
Machine 1
0_0
1_0
3_6
3_7
0_3
1_2
3_1
4_3
0_2
1_1
4_0
0_1
0_5
3_3
4_1
0_4
1_3
3_2
3_4
4_2
3_0
3_5
06:00
08:00
10:00
12:00
14:00
16:00
Time of day

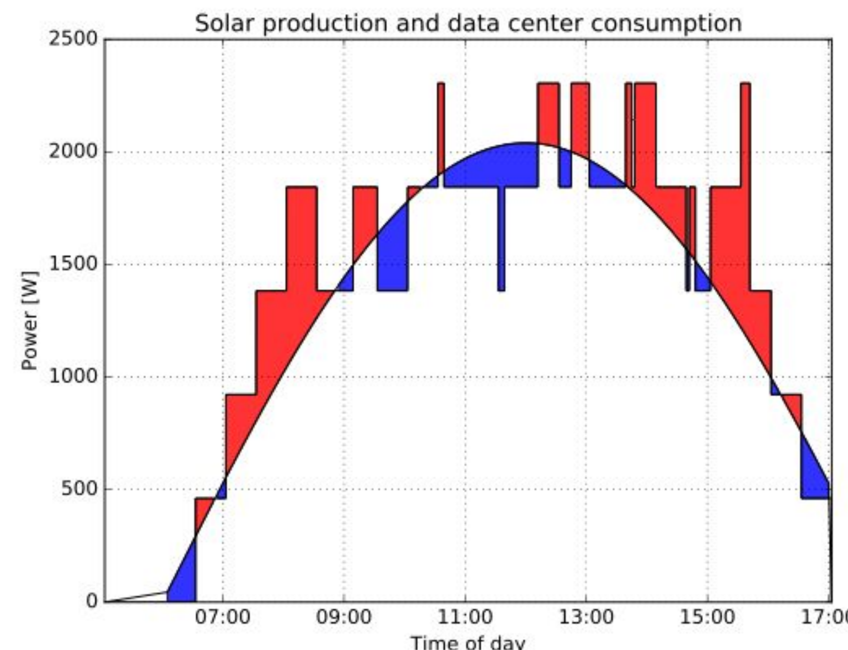
Solar production and data center consumption
Power [W]
2500
2000
1500
1000
500
0
07:00
09:00
11:00
13:00
15:00
17:00
Time of day

But how to assemble and test what we propose?

- We can't build a 1 MW datacenter (with the money we have :))
- But we can make smaller experiments to validate the models (e.g., power production models, servers consumption, ...)
- AND we integrate results within a platform, mixing both simulated parts and real platforms

Middleware, Simulation, Real Hardware

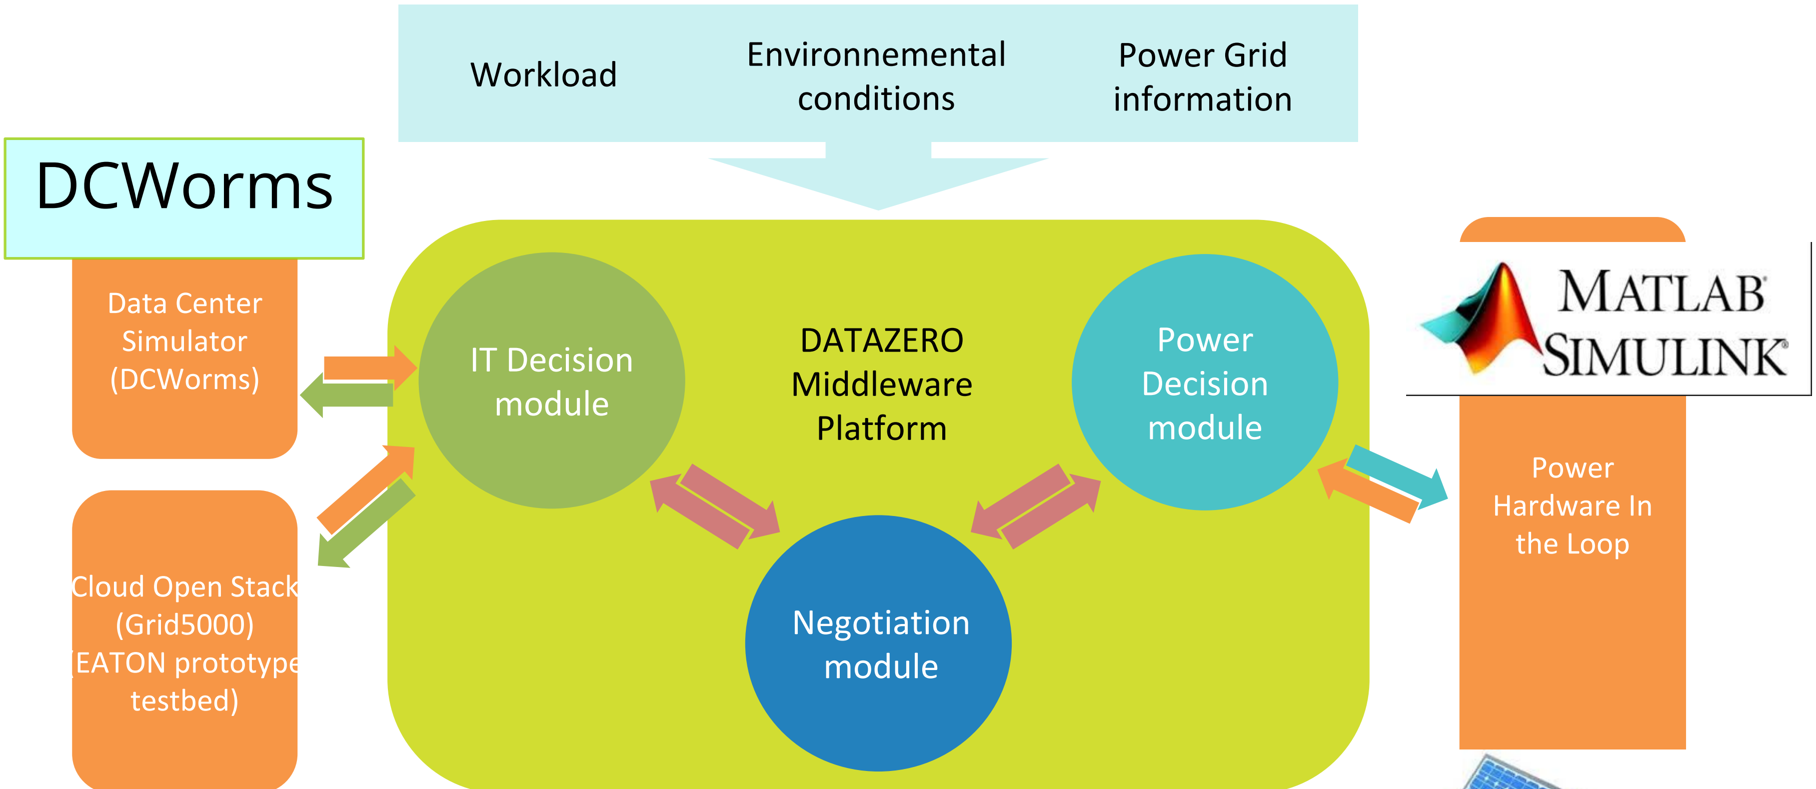

openstack
CLOUD SOFTWARE

As a result, we have a Power Dynamic Data Center with Long Term Green Power Storage

Challenge 1. Making demand and envelope constraints coincide

Challenge 2. Sizing the infrastructure

Challenge 3. Commanding electrical converters

Challenge 4. Scheduling IT load

Challenge 5. Managing thermal burden

Challenge 6. Keeping the complexity at pace

Challenge 7. Developing a simulation toolkit

as green as possible

Conclusion

- An ongoing work
- Still about 1 years to go
- About 25 researchers involved

- Keep in touch!

Jean-Marc.Pierson@irit.fr