Scaling your experiments

Lucas Nussbaum[1]

RESCOM'2017

June 2017

[1] The Grid'5000 part is joint work with S. Delamare, F. Desprez, E. Jeanvoine, A. Lebre, L. Lefevre, D. Margery, P. Morillon, P. Neyron, C. Perez, O. Richard and many others

IEEE Cluster 2014 Madrid 22-26 September

Grid'5000: Running a Large Instrument for Parallel and Distributed Computing Experiments

F. Desprez

INRIA Grenoble Rhône-Alpes, LIP ENS Lyon, Avalon Team

Joint work with G. Antoniu, Y. Georgiou, D. Glesser, A. Lebre, L. Lefèvre, M. Liroz, D. Margery, L. Nussbaum, C. Perez, L. Pouillioux

The data-centers facet of SILECS (A.K.A. Grid'5000)

Frédéric Desprez & Lucas Nussbaum
Grid'5000 Scientific & Technical Directors

2019-04-25

SILECS/Grid'5000

Adrien Lebre, Jan 2022

Validation in (Computer) Science

- Two classical approaches for validation:
  - Formal: equations, proofs, etc.
  - Experimental, on a scientific instrument
- Often a mix of both:
  - In Physics, Chemistry, Biology, etc.
  - In Computer Science

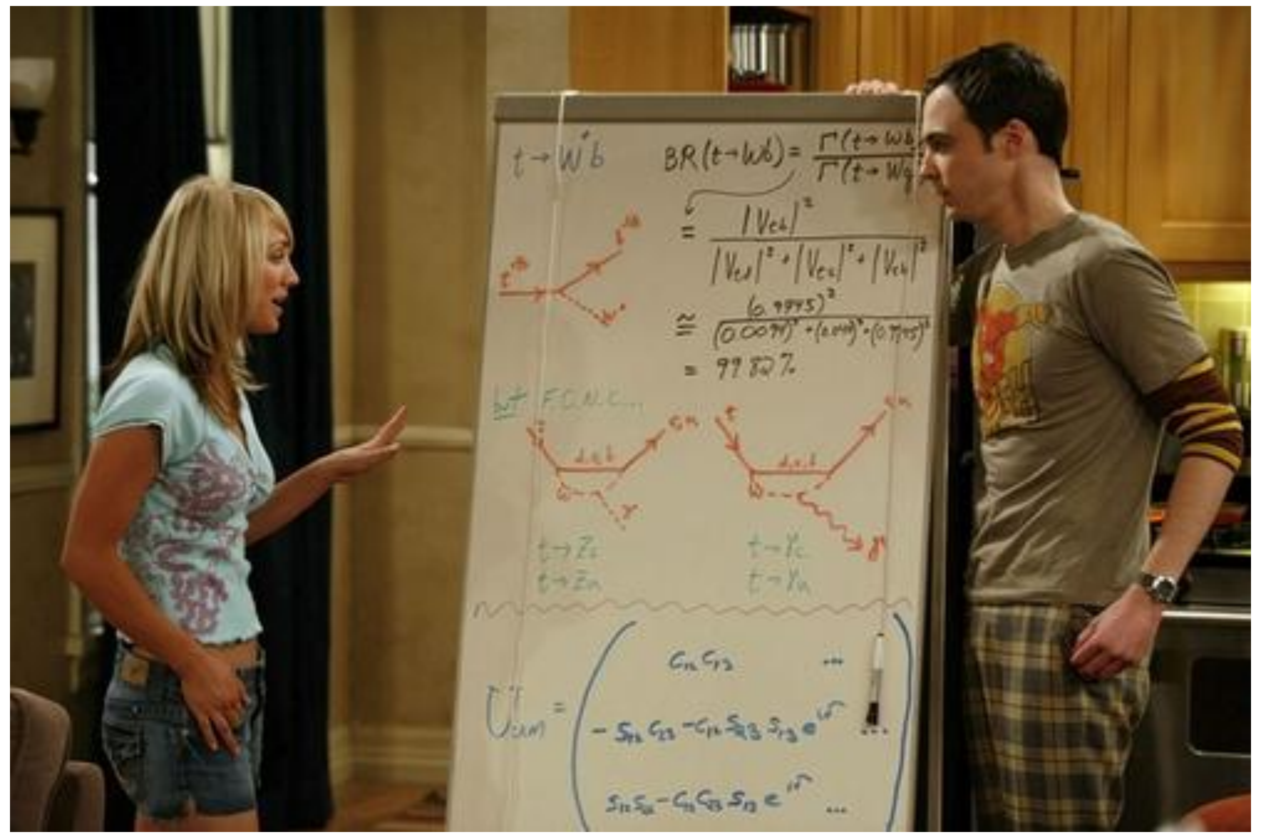

Analytic modeling

Grid'5000

Purely analytical (mathematical) models

- Demonstration of properties (theorem)
- Models need to be tractable: over-simplification?
- Good to understand the basic of the problem
- Most of the time ones still perform a experiments (at least for comparison)

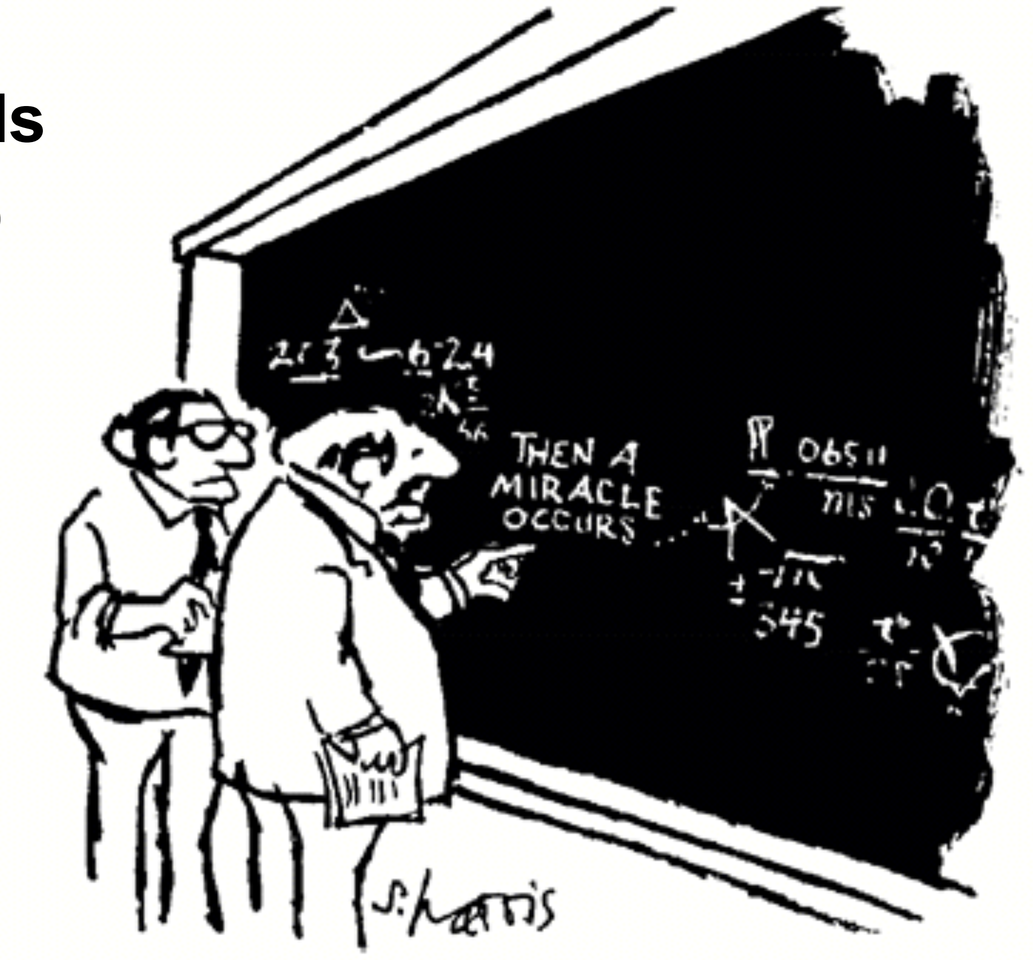

"I THINK YOU SHOULD BE MORE EXPLICIT HERE IN STEP TWO."

For a practical impact (especially in distributed computing): analytic studies are not always sufficient or possible! Micro-services Deathstar nightmare

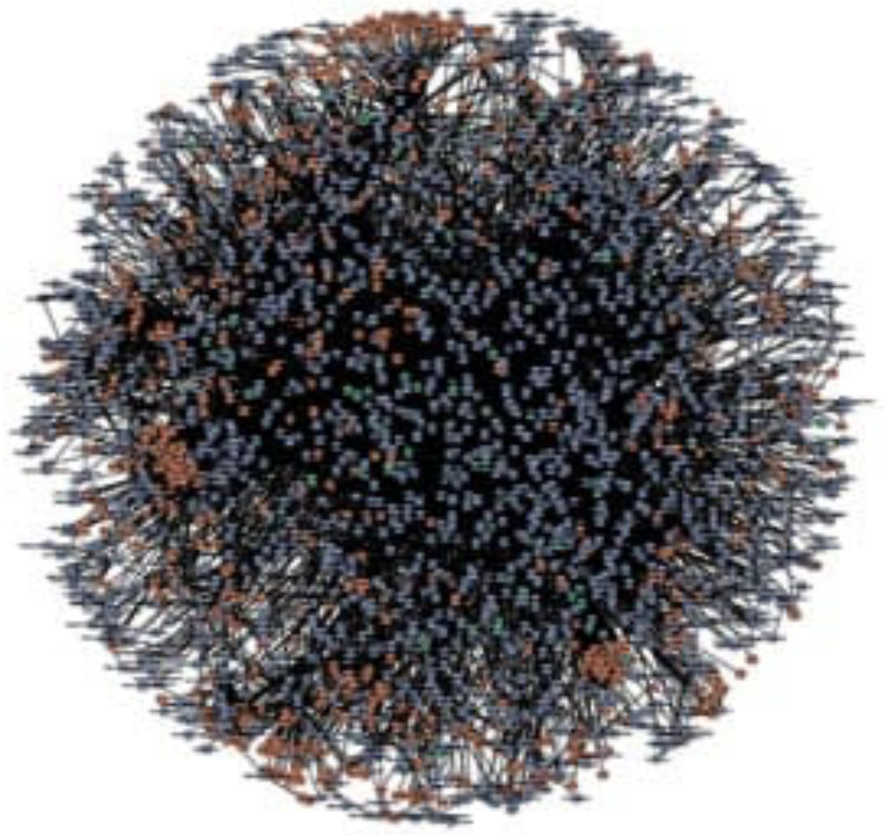

amazon.com

NETFLIX

DC & networking: peculiar fields in CS

- Performance and scalability are central to results
  - But depend greatly on the environment (hardware, network, software stack, etc.)
  - Many contributions are about *fighting* the environment
    - Making the most out of limited, complex and different resources (e.g. memory/storage hierarchy, asynchronous communications)
    - Handling performance imbalance, noise ⇝ asynchronism, load balancing
    - Handling faults ⇝ fault tolerance, recovery mechanisms
    - Hiding complexity ⇝ abstractions: middlewares, runtimes
- Validation of most contributions require experiments
  - Formal validation often intractable or unsuitable
  - Even for more theoretical work ⇝ simulation (SimGrid, CloudSim)
- Experimenting is difficult and time-consuming. . . but often neglected
  - *Everybody is doing it, not so many people are talking about it*

Experimental culture not comparable with other science

Grid'5000

Different studies

- 1994: 400 papers
  - Between 40% and 50% of CS ACM papers requiring experimental validation had none (15% in optical engineering) [Lukovicz et al.]
- 1998: 612 papers
  - "*Too many articles have no experimental validation*" [Zelkowitz and Wallace 98]
- 2009 update
  - *Situation is improving*
- 2007: Survey of simulators used in P2P research
  - Most papers use an unspecified or custom simulator

Computer science not at the same level than some other sciences

- Nobody redo experiments
- Lack of tool and methodologies

Paul Lukowicz et al. **Experimental Evaluation in Computer Science: A Quantitative Study**. In: *J.l of Systems and Software* 28:9-18, 1994
M.V. Zelkowitz and D.R. Wallace. **Experimental models for validating technology**. Computer, 31(5):23-31, May 1998
Marvin V. Zelkowitz. **An update to experimental models for validating computer technology.** In: *J. Syst. Softw.* 82.3:373–376, Mar. 2009
S. Naicken et al. **The state of peer-to-peer simulators and simulations**. In: *SIGCOMM Comput. Commun. Rev.* 37.2:95–98, Mar. 2007

Experimental methodologies

Simulation

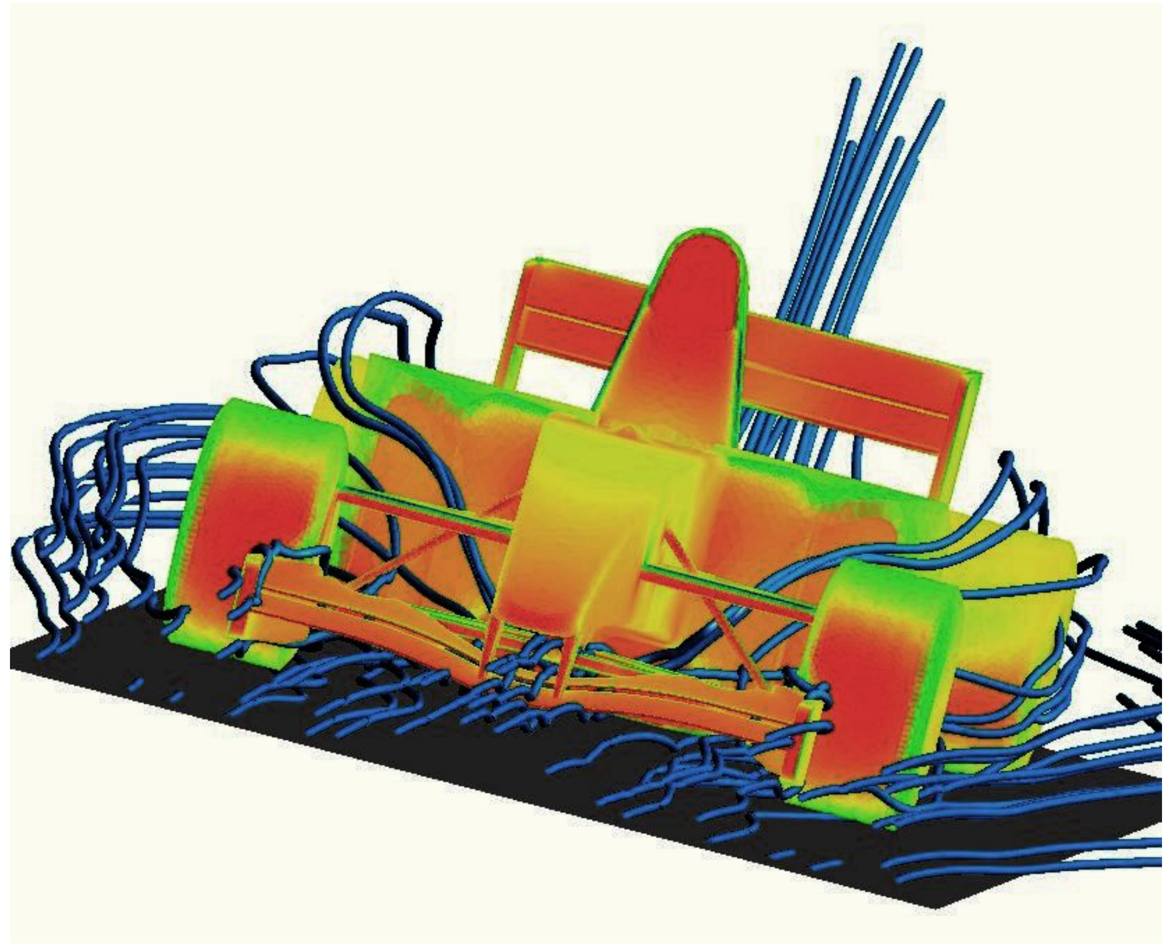

1. **Model** application
2. **Model** environment
3. **Compute** interactions

Real-scale experiments

Execute the **real** application
on **real** machines

Complementary solutions:

- ☺ Work on algorithms
- ☺ More scalable, easier

- ☺ Work with real applications
- ☺ Perceived as more realistic

Experimental methodologies: Emulation

- Take a real system
- Degrade it to make it match experimental conditions

From ideas to applications

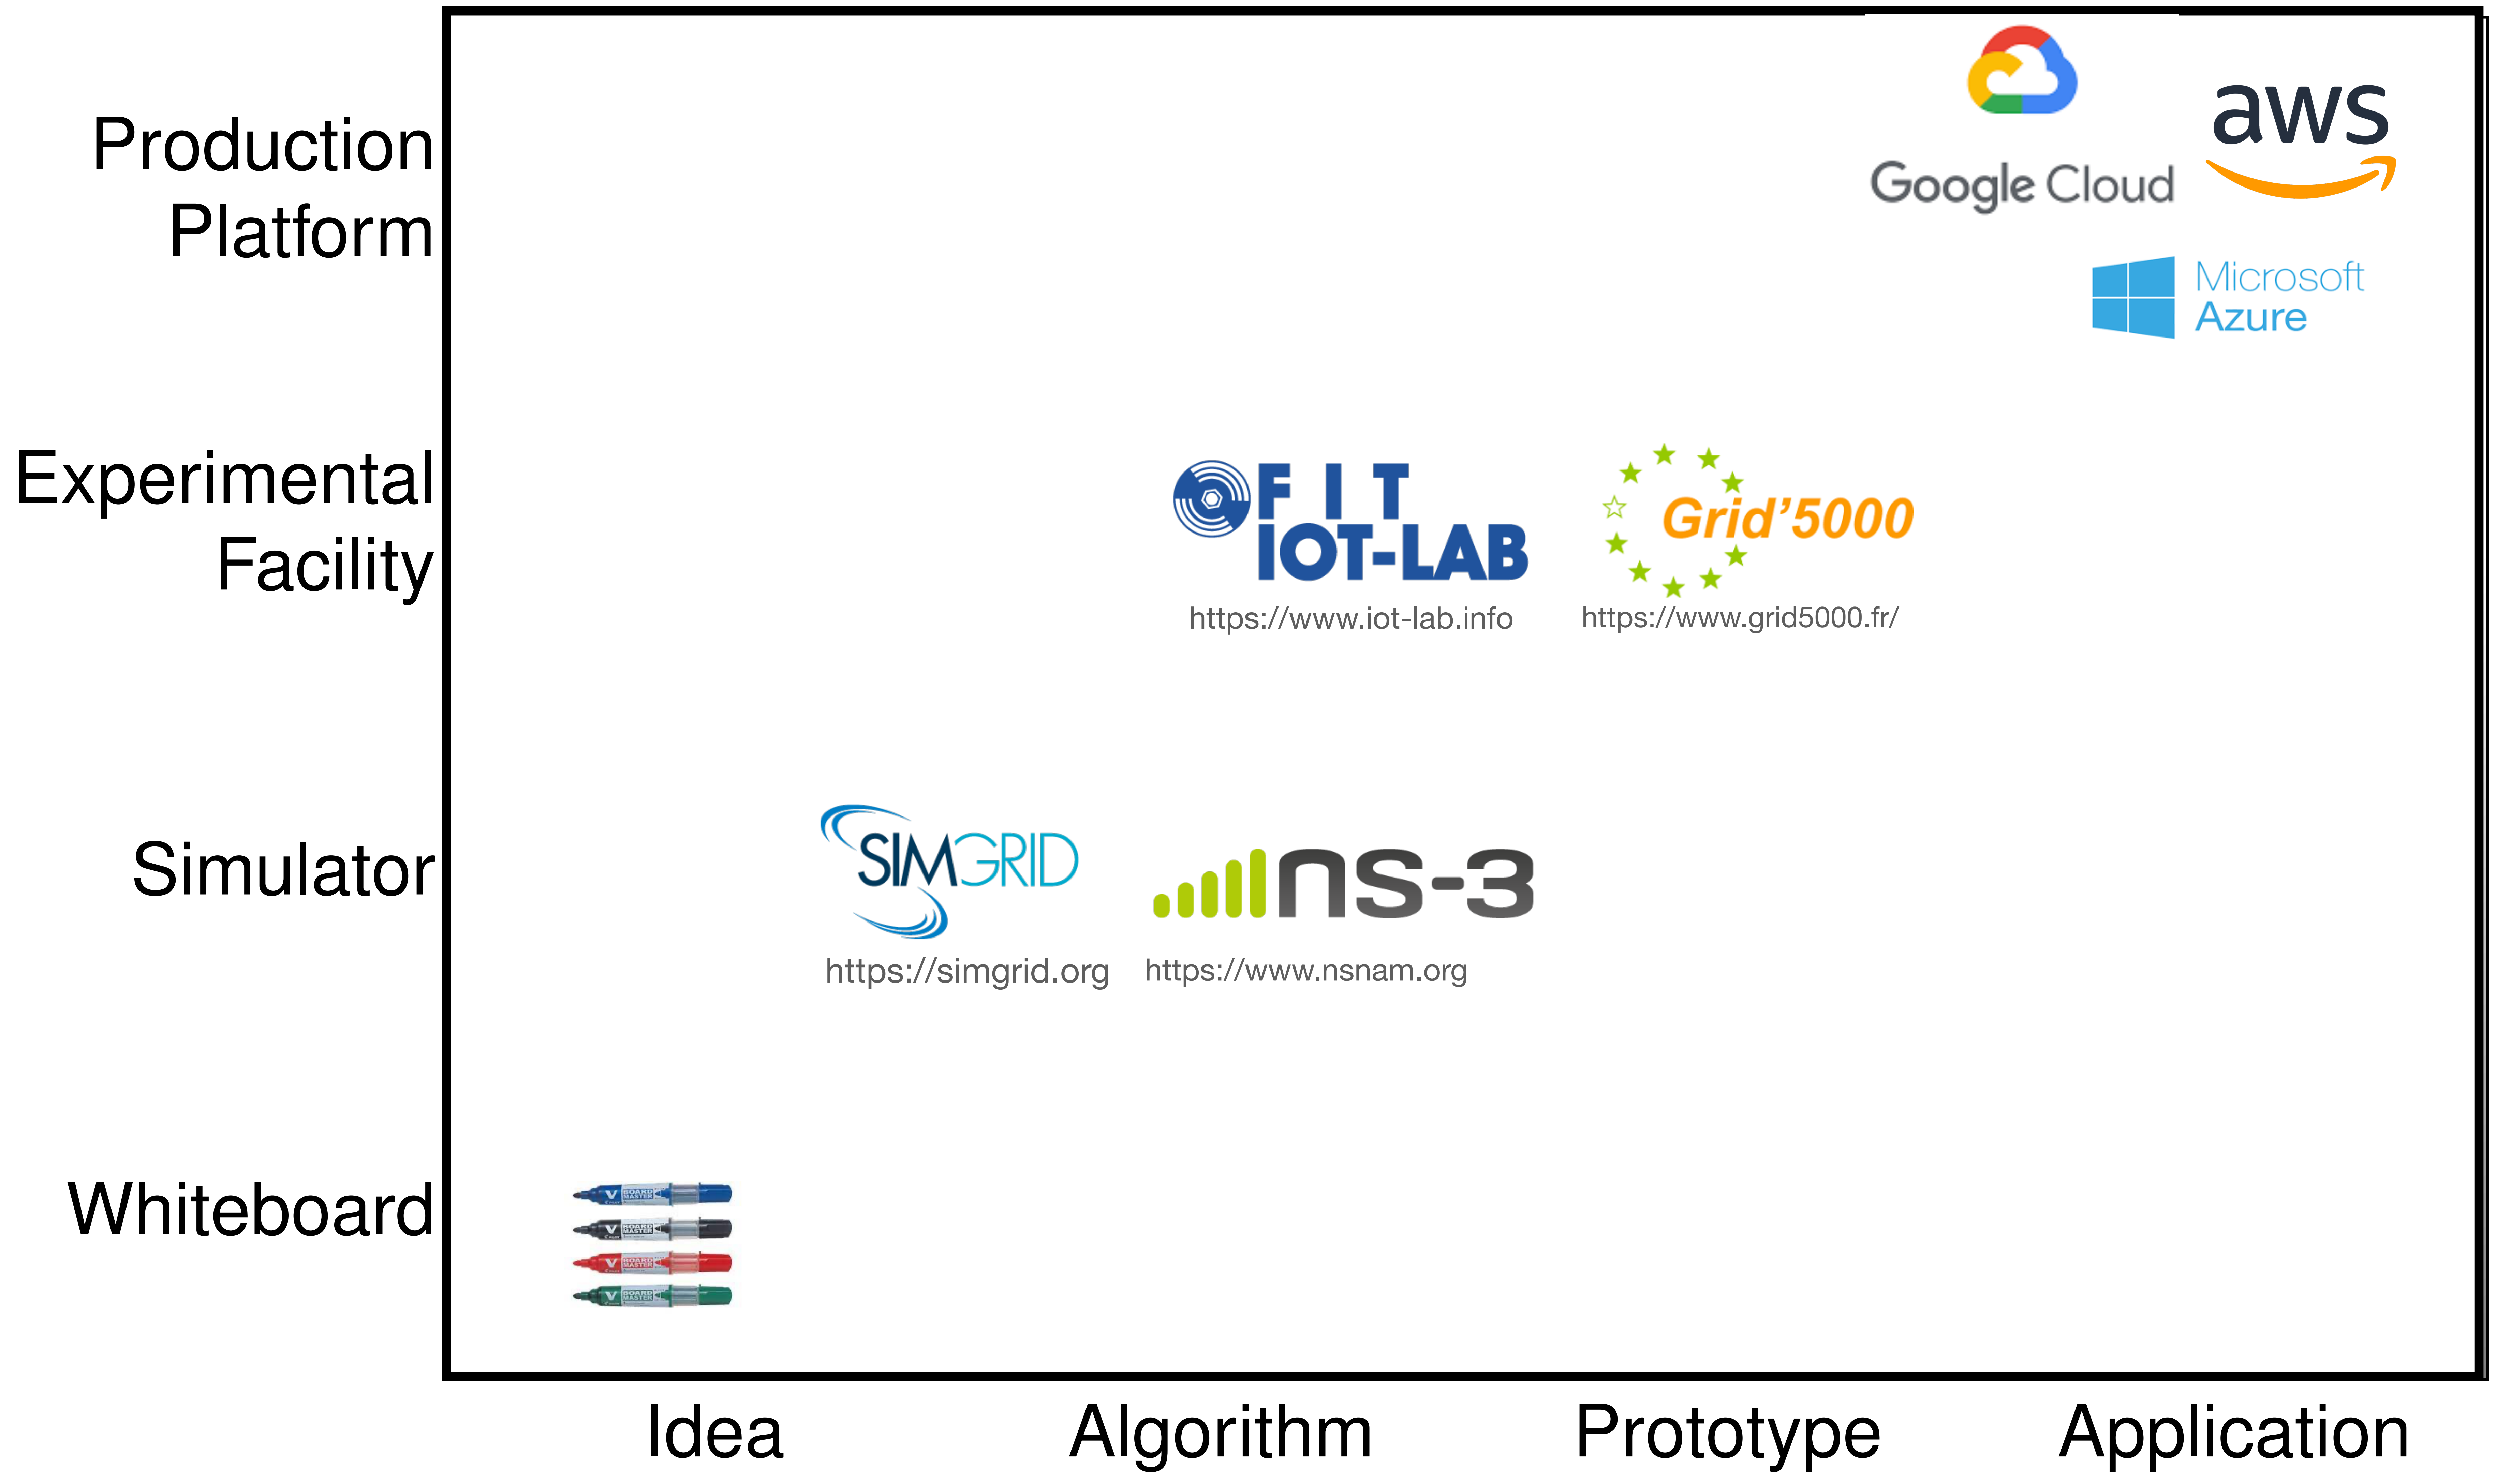

Wireless (WiFi, 4G/LTE, SDR): CorteXlab[12], R2lab

- Sets of customizable wireless nodes in an anechoic chamber
- For experiments on the physical layer

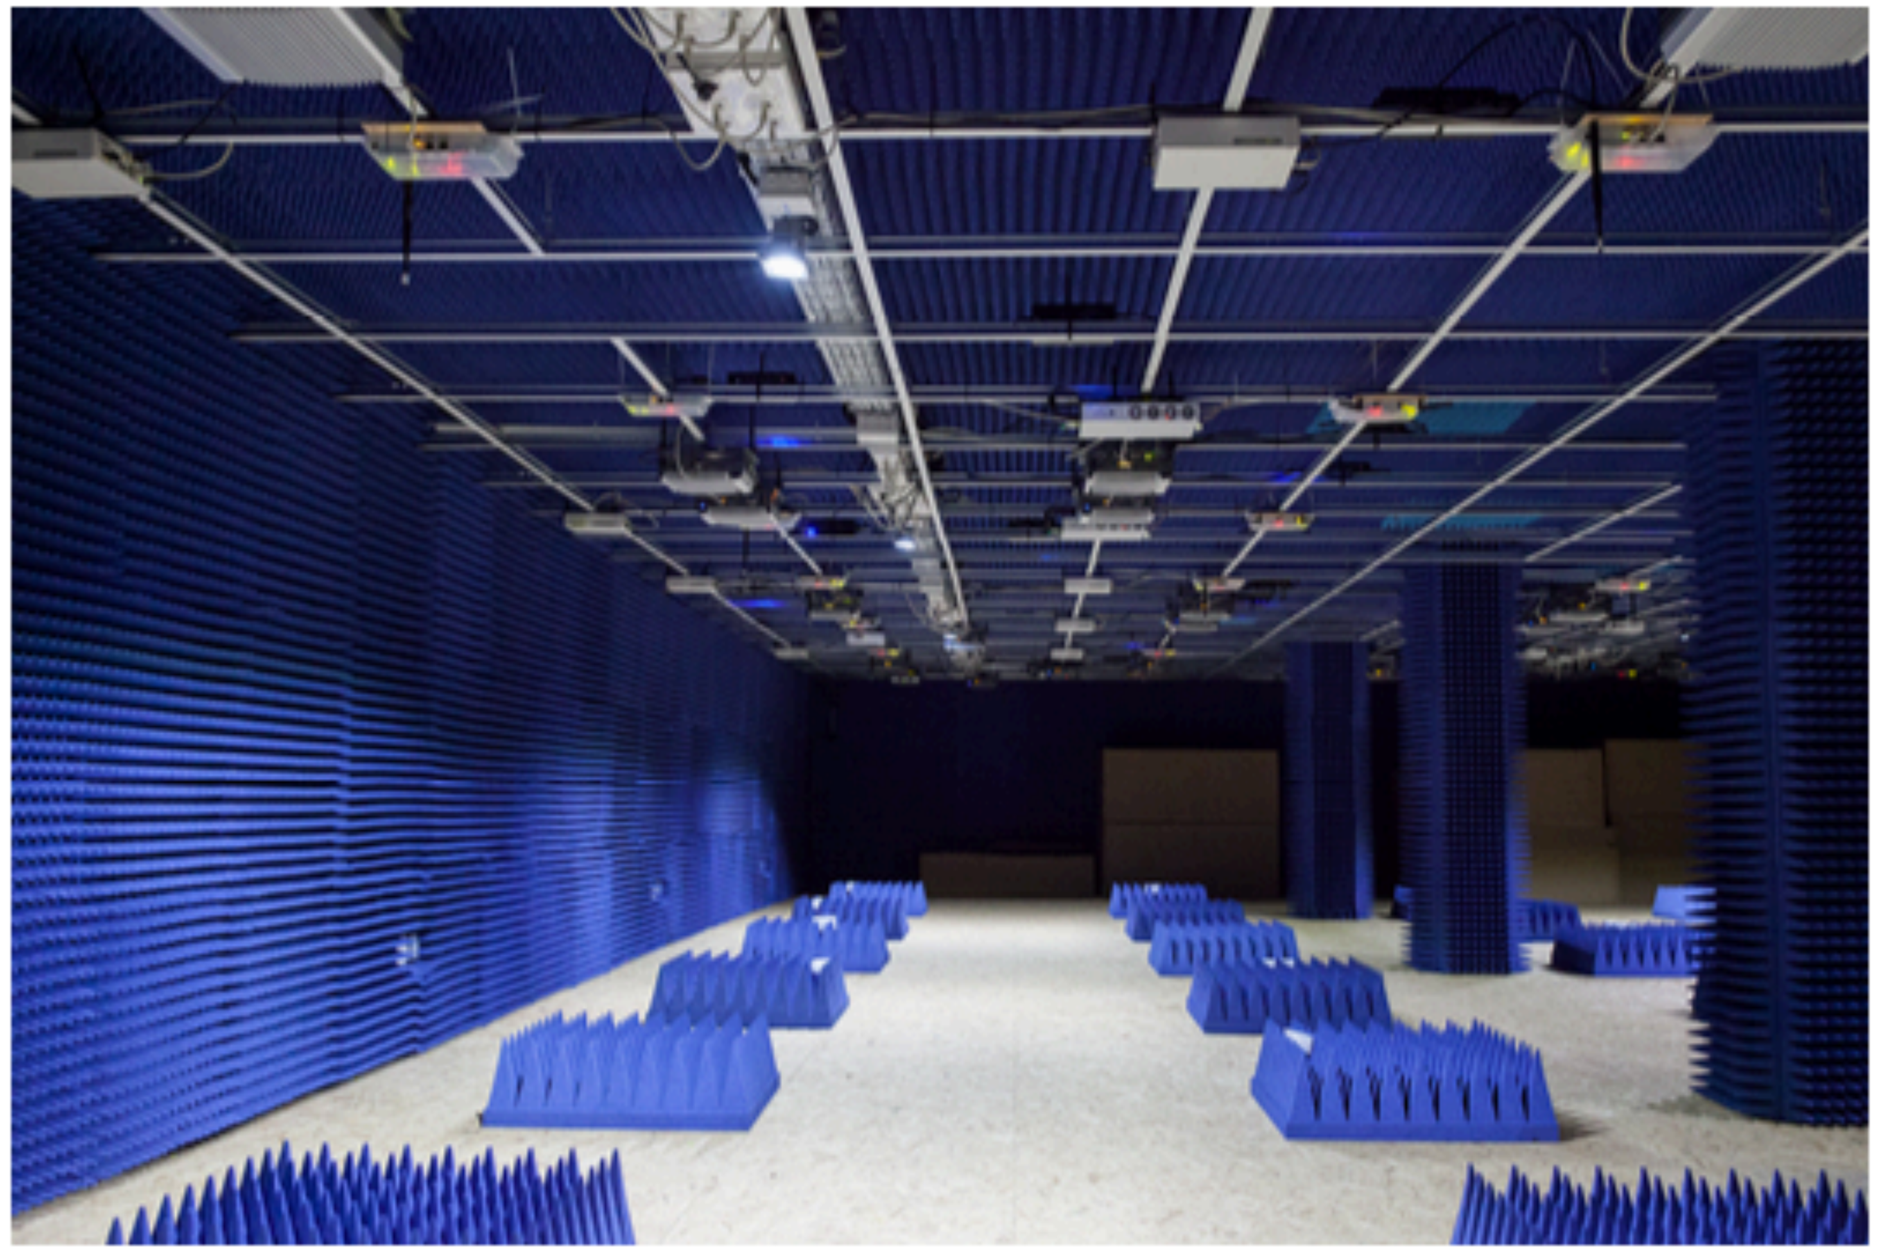

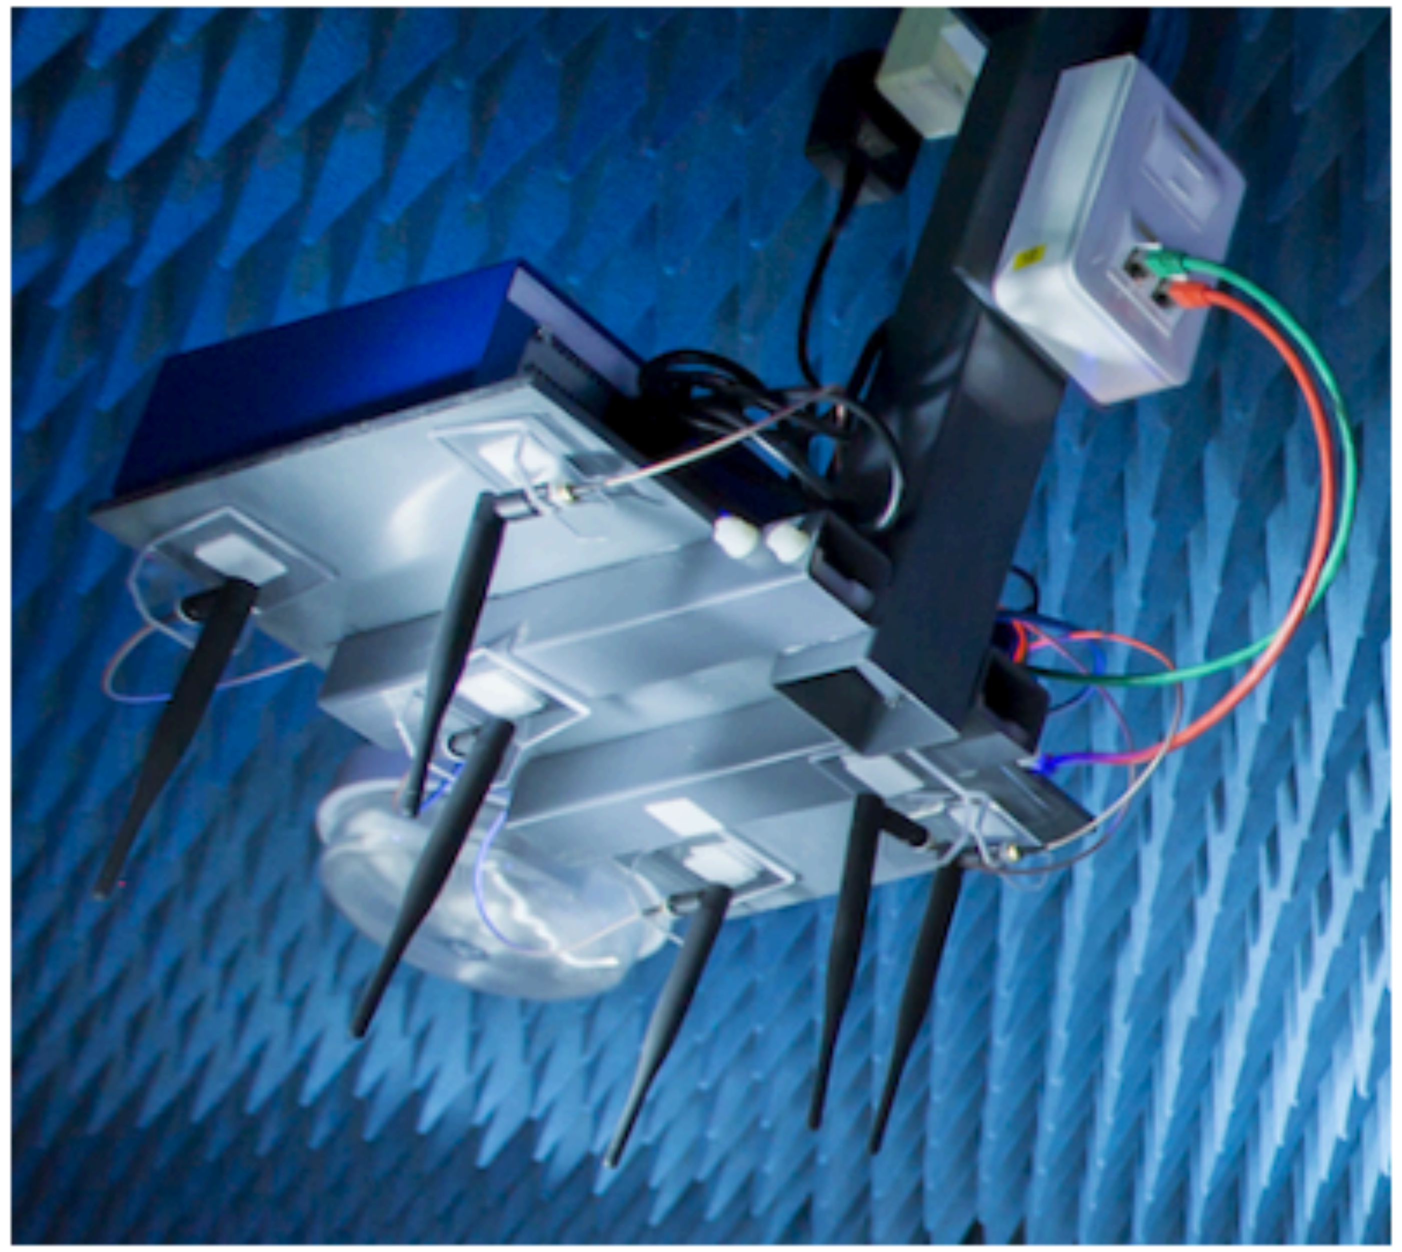

`http://www.cortexlab.fr`
`https://r2lab.inria.fr`

[12] Albdelbassat Massouri et al. "CorteXlab: An Open FPGA-based Facility for Testing SDR & Cognitive Radio Networks in a Reproducible Environment". In: *INFOCOM'2014 Demo/Poster Session*. 2014.

Internet of Things: FIT IoT-Lab[11]

- 2769 wireless sensors (from WSN430 to Cortex A8)
- 7 sites (Grenoble, Lille, Strasbourg, Saclay, Rennes, IMT Paris, Lyon)
- Also mobile robots
- Typical experiment: IoT communication protocols

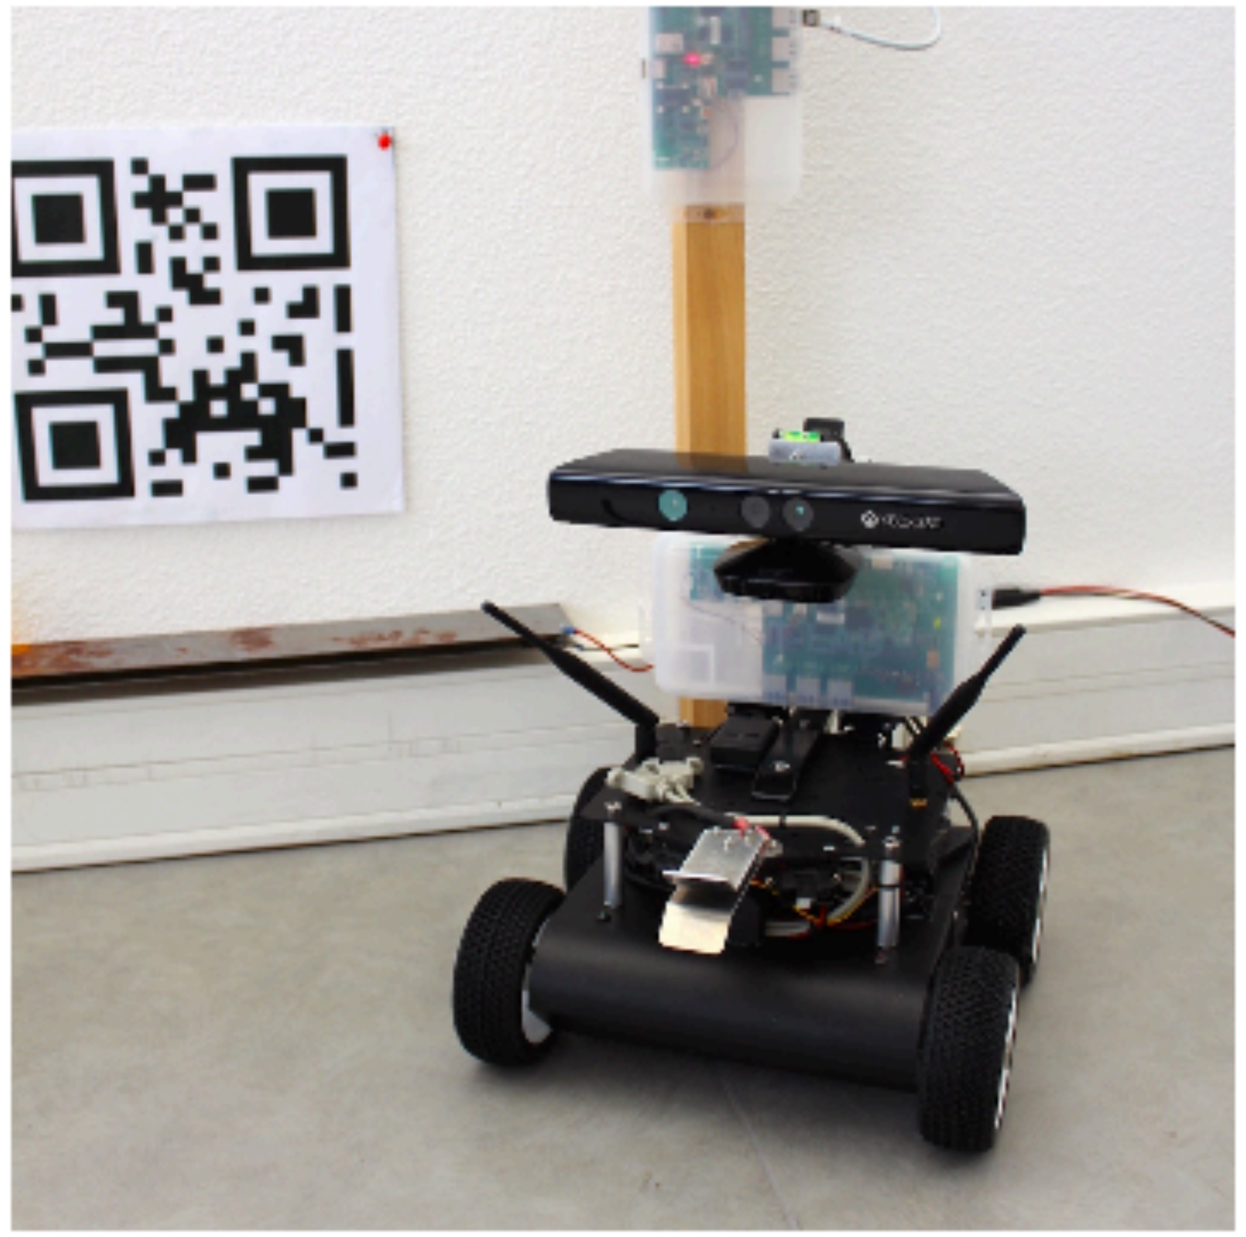

`https://www.iot-lab.info/`

[11] Cedric Adjih et al. "FIT IoT-LAB: A large scale open experimental IoT testbed". In: *IEEE 2nd World Forum on Internet of Things (WF-IoT)*. 2015.

The Grid'5000 testbed

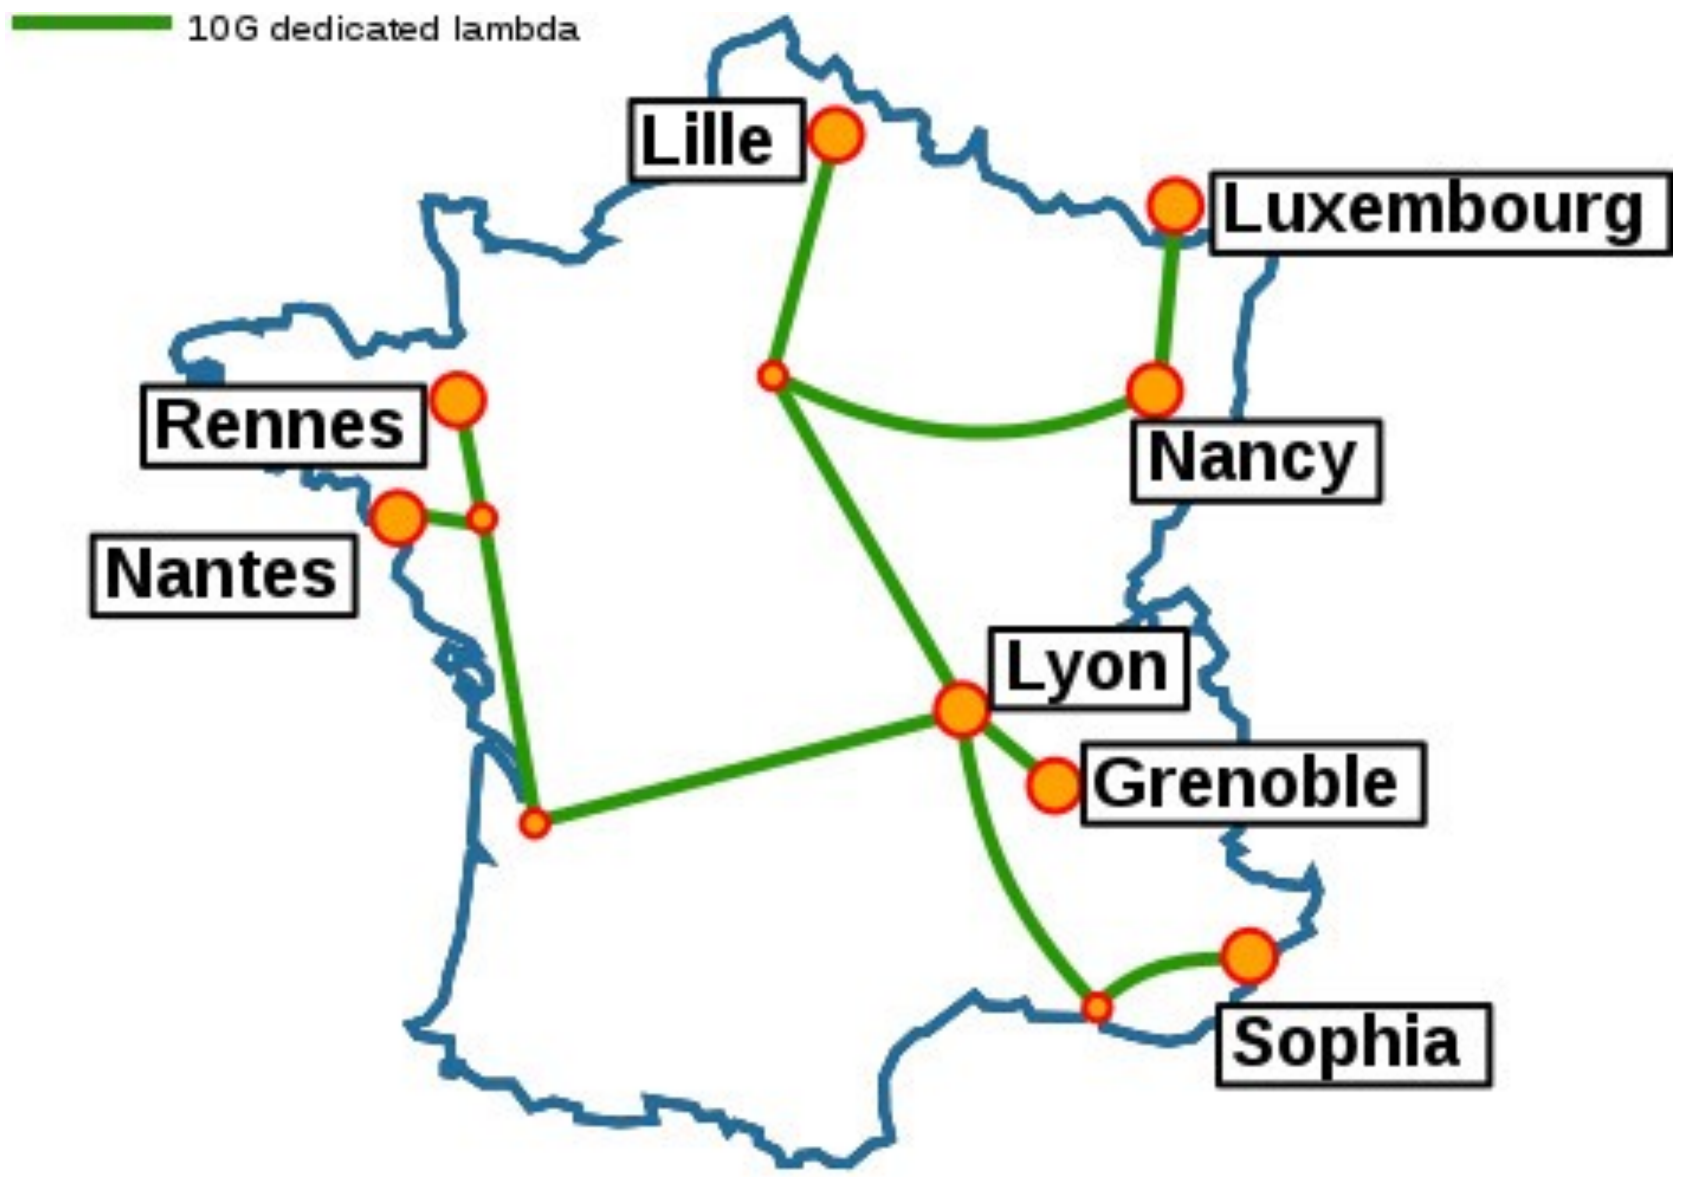

- **A large-scale testbed for distributed computing**
  - 8 sites, 31 clusters, 828 nodes, 12328 cores
  - Dedicated 10-Gbps backbone network
  - 550 users and 120 publications per year
- **A meta-cloud, meta-cluster, meta-data-center**
  - Used by CS researchers in HPC, Clouds, Big Data, Networking, AI
  - To experiment in a fully controllable and observable environment
  - Similar problem space as Chameleon and Cloudlab (US)
  - Design goals
    - Support high-quality, reproducible experiments
    - On a large-scale, distributed, shared infrastructure

Landscape – cloud & experimentation[1]

- Public cloud infrastructures (AWS, Azure, Google Cloud Platform, etc.)
  - ☹ No information/guarantees on placement, multi-tenancy, real performance
- Private clouds: Shared observable infrastructures
  - ☺ Monitoring & measurement
  - ☹ No control over infrastructure settings
  - ⇝ Ability to understand experiment results
- Bare-metal as a service, fully reconfigurable infrastructure (Grid'5000)
  - ☺ Control/alter all layers (virtualization technology, OS, networking)
  - ⇝ *In vitro* Cloud

And the same applies to all other environments (e.g. HPC)

[1] Inspired from a slide by Kate Keahey (Argonne Nat. Lab.)

Some ~~recent~~ results from Grid'5000 users

- Portable Online Prediction of Network Utilization (Inria Bdx + US)
- Energy proportionality on hybrid architectures (LIP/IRISA/Inria)
- Maximally Informative Itemset Mining (Miki) (LIRM/Inria)
- Asynchronous data storage for large-scale simulations (Inria)
- BeBida: Mixing HPC and BigData Workloads (LIG)
- HPC: In Situ Analytics (LIG/Inria)
- Addressing the HPC/Big-Data/IA Convergence
- An Orchestration Syst. for IoT Applications in Fog Environment (LIG/Inria)
- Toward a resource management system for Fog/Edge infrastructures
- Distributed Storage for Fog/Edge infrastructures (LS2N)
- From Network Traffic Measurements to QoE for Internet Video (Inria)

Users and publications

Year	2010	2011	2012	2013	2014	2015	2016	2017	2018
Active users	564	553	592	514	528	458	573	600	564
Publications	154	141	101	134	106	143	122	151	127
PhD & HDR	14	20	9	27	24	30	27	23	22
Usage rate	50%	56%	58%	63%	63%	63%	55%	53%	70%

- 1313 active users over the last 3 years
- 3769 active users since 2003
- 2007 publications that benefited from Grid'5000 in our **HAL collection**[6]
  - Computer Science: 96%, Mathematics: 2.4%, Physics: 2.4%
  - Since 2015: LORIA: 23%, IRISA: 23%, LIG: 19%, LIP: 13%, LS2N: 13%, CRISTAL: 5%, LIRMM: 5%, LIP6: 3%

[6] `https://hal.archives-ouvertes.fr/GRID5000`

Conclusions

- An advanced and established infrastructure for the *data-center* facets of Computer Science
  - Large-scale, distributed
  - Shared (many involved laboratories and institutions)
  - Designed for reconfigurability, observability, reproducible research
- Future: SILECS
  - SILECS Infrastructure for Large-scale Experimental Computer Science
  - On the fondations of Grid'5000 and FIT (IoT-Lab, CorteXLab, R2Lab, etc.)
  - Experiment on a single infrastructure, from edge to cloud

https://www.silecs.net

https://slices-ri.eu