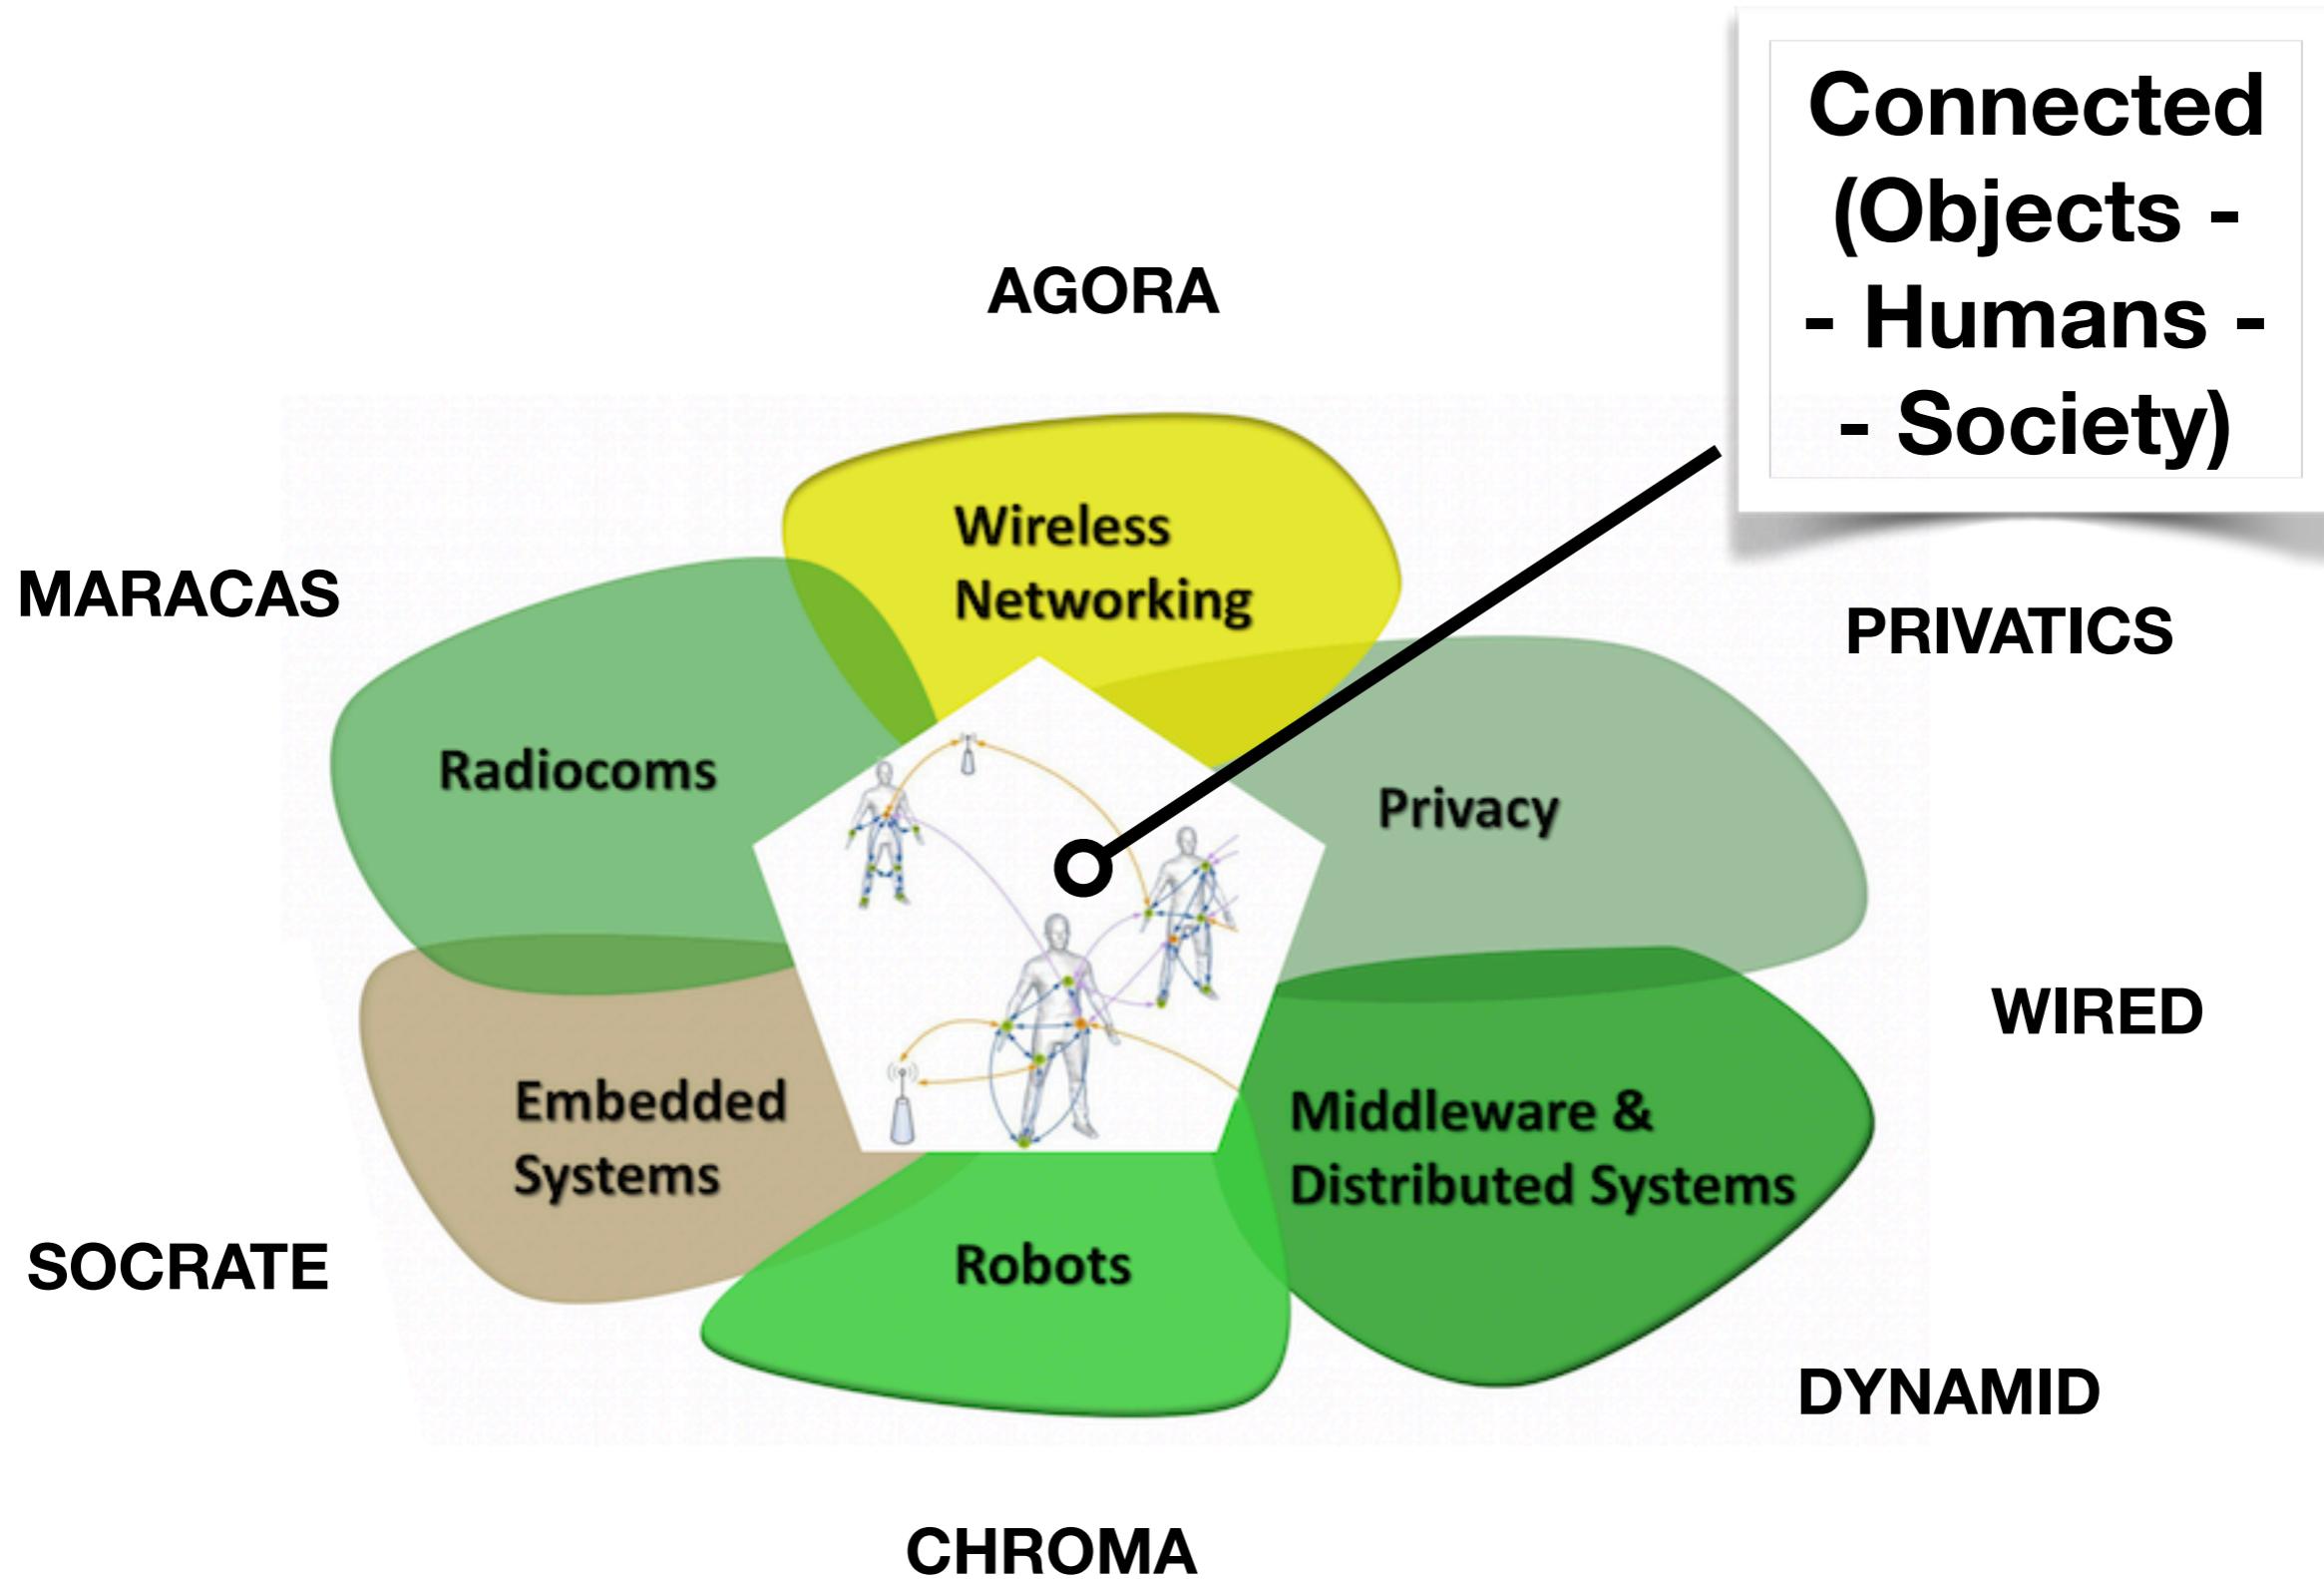


YOUPI: feedbacks, ongoing and future works of a Fog/ Edge Computing platform

Frédéric Le Mouël
Univ. Lyon, INSA Lyon, CITI Lab

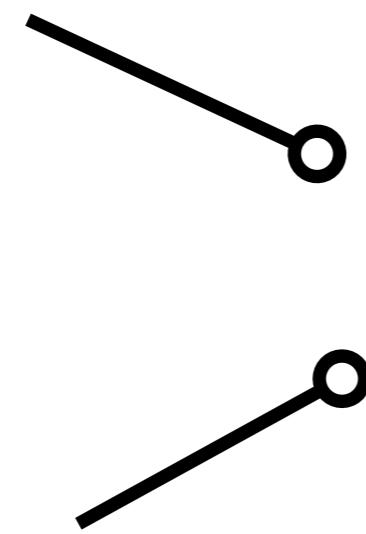


INSA / INRIA CITI Lab



Communications issues for IoT

Deploying dense networks



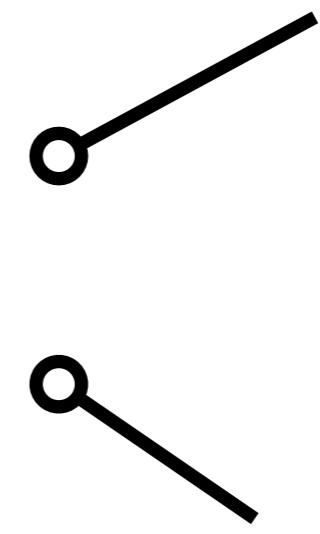
Machine-to-Machine (M2M)
communication protocols

Urban
Wireless
Networks

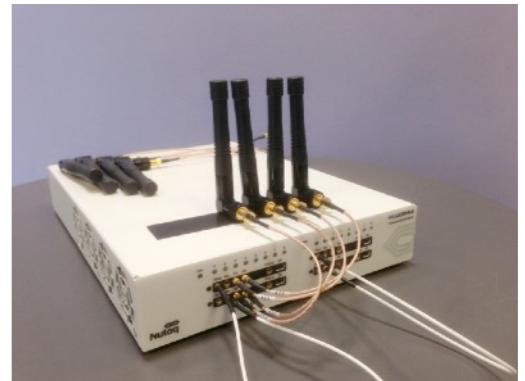


Flexible Radio Front-End

Exploiting the data
carried by the network



Software
and Cognitive Radio



22 USRP + 16 PicoSDR nodes

http://www.cortexlab.fr/?page_id=387

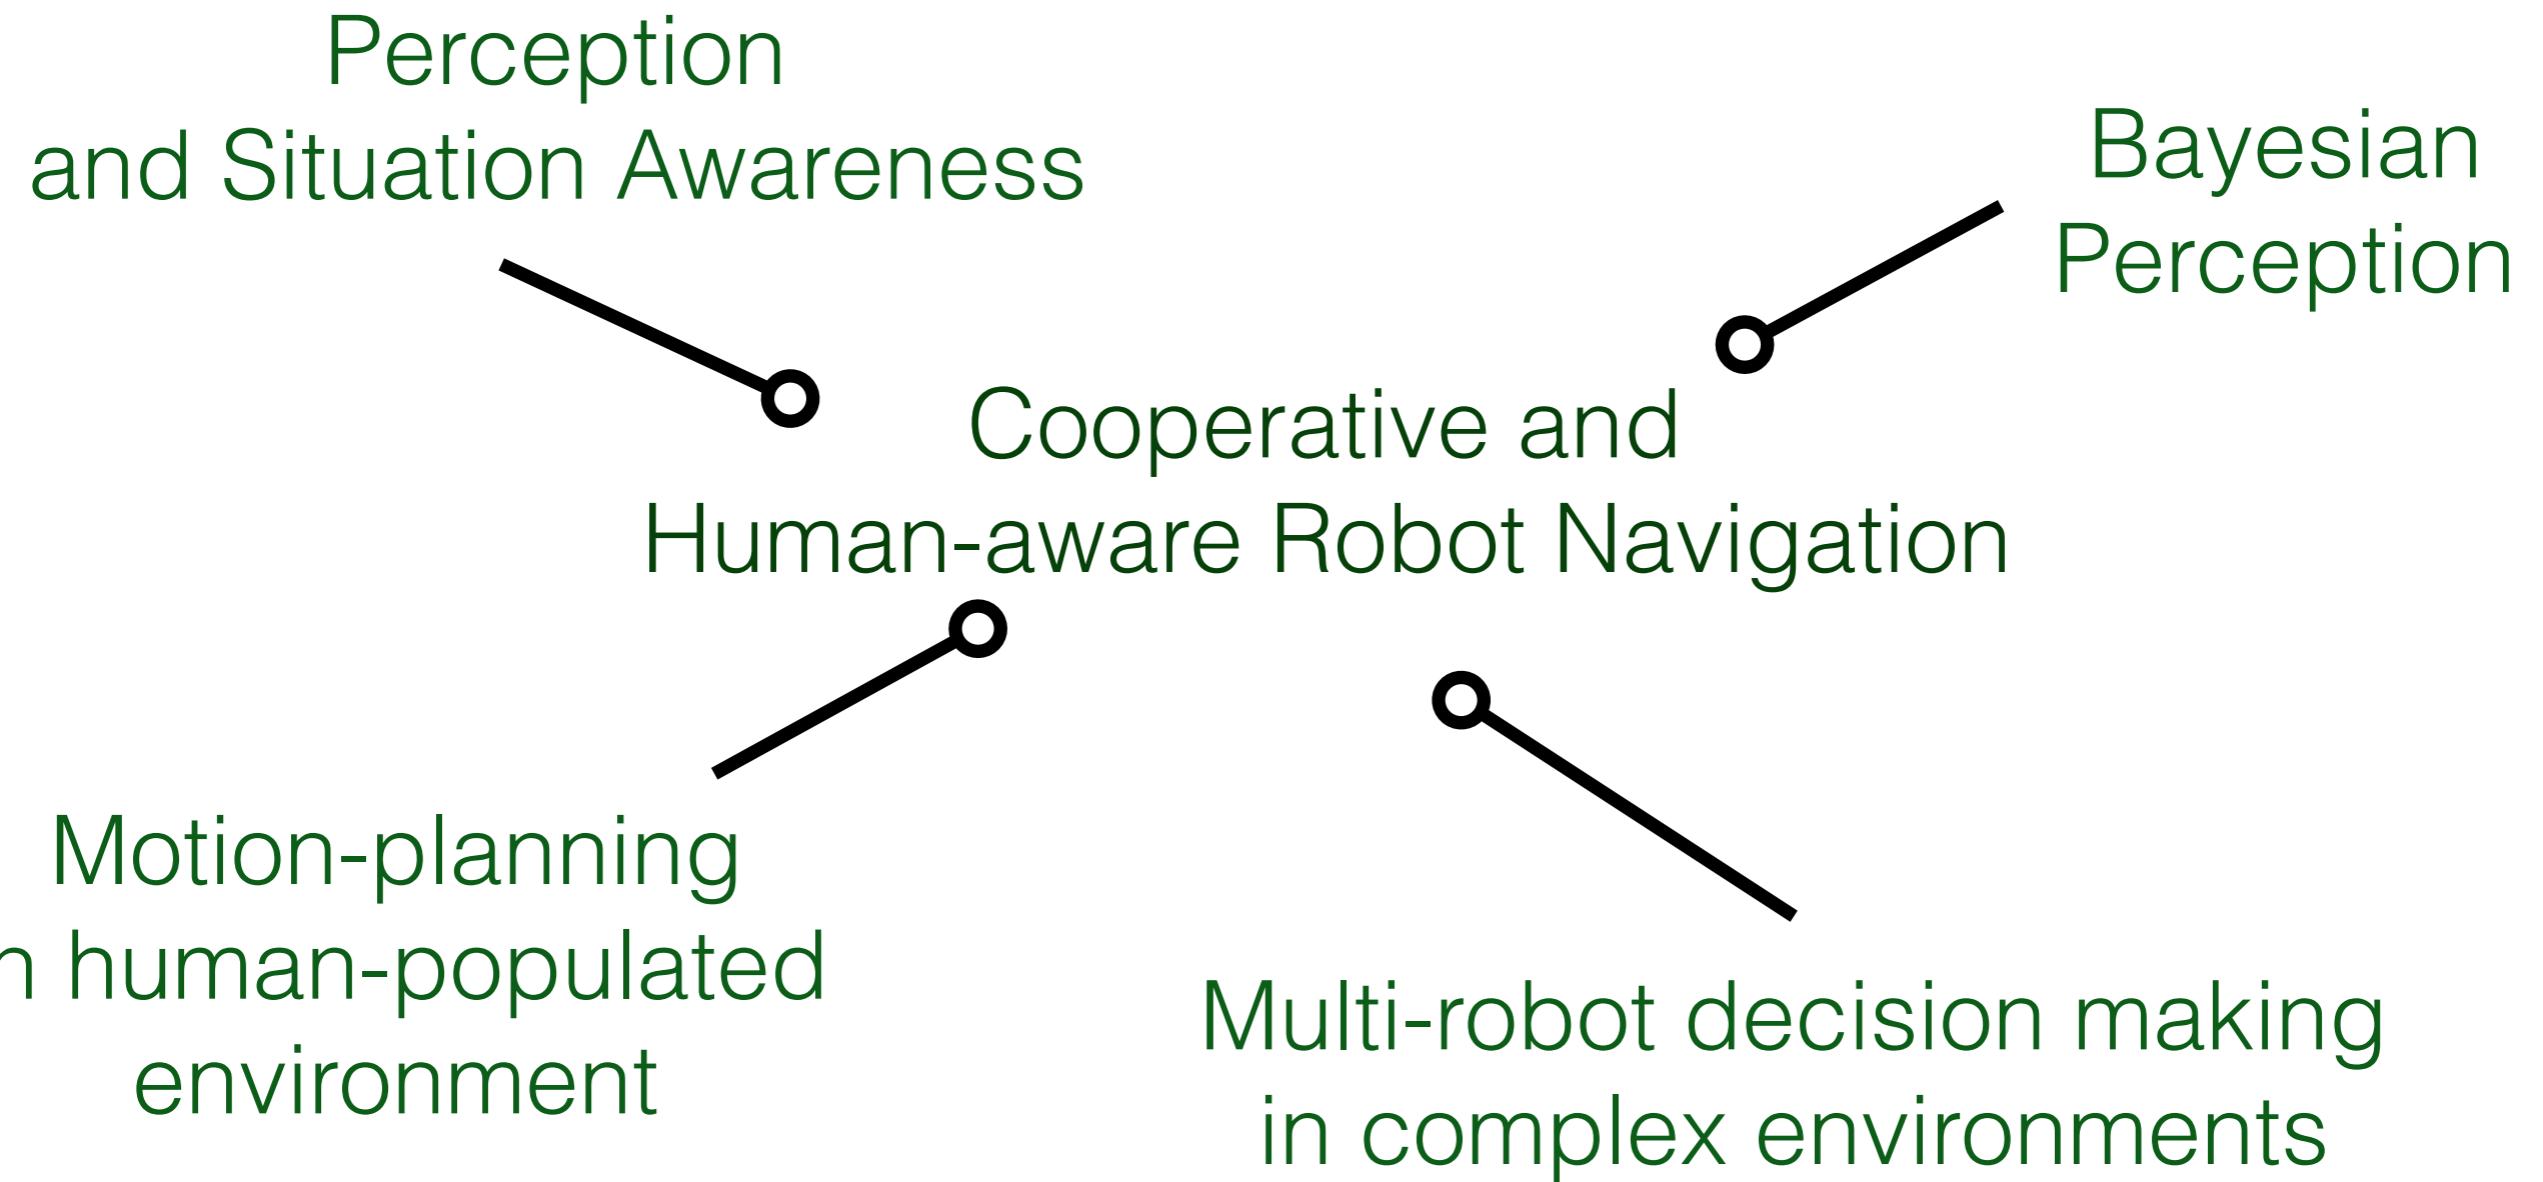
FIT/IoTLab - <https://www.iot-lab.info/deployment/lyon/>



11 A8 nodes

<https://www.iot-lab.info/hardware/a8/>

Cooperation / Mobility issues for Robots



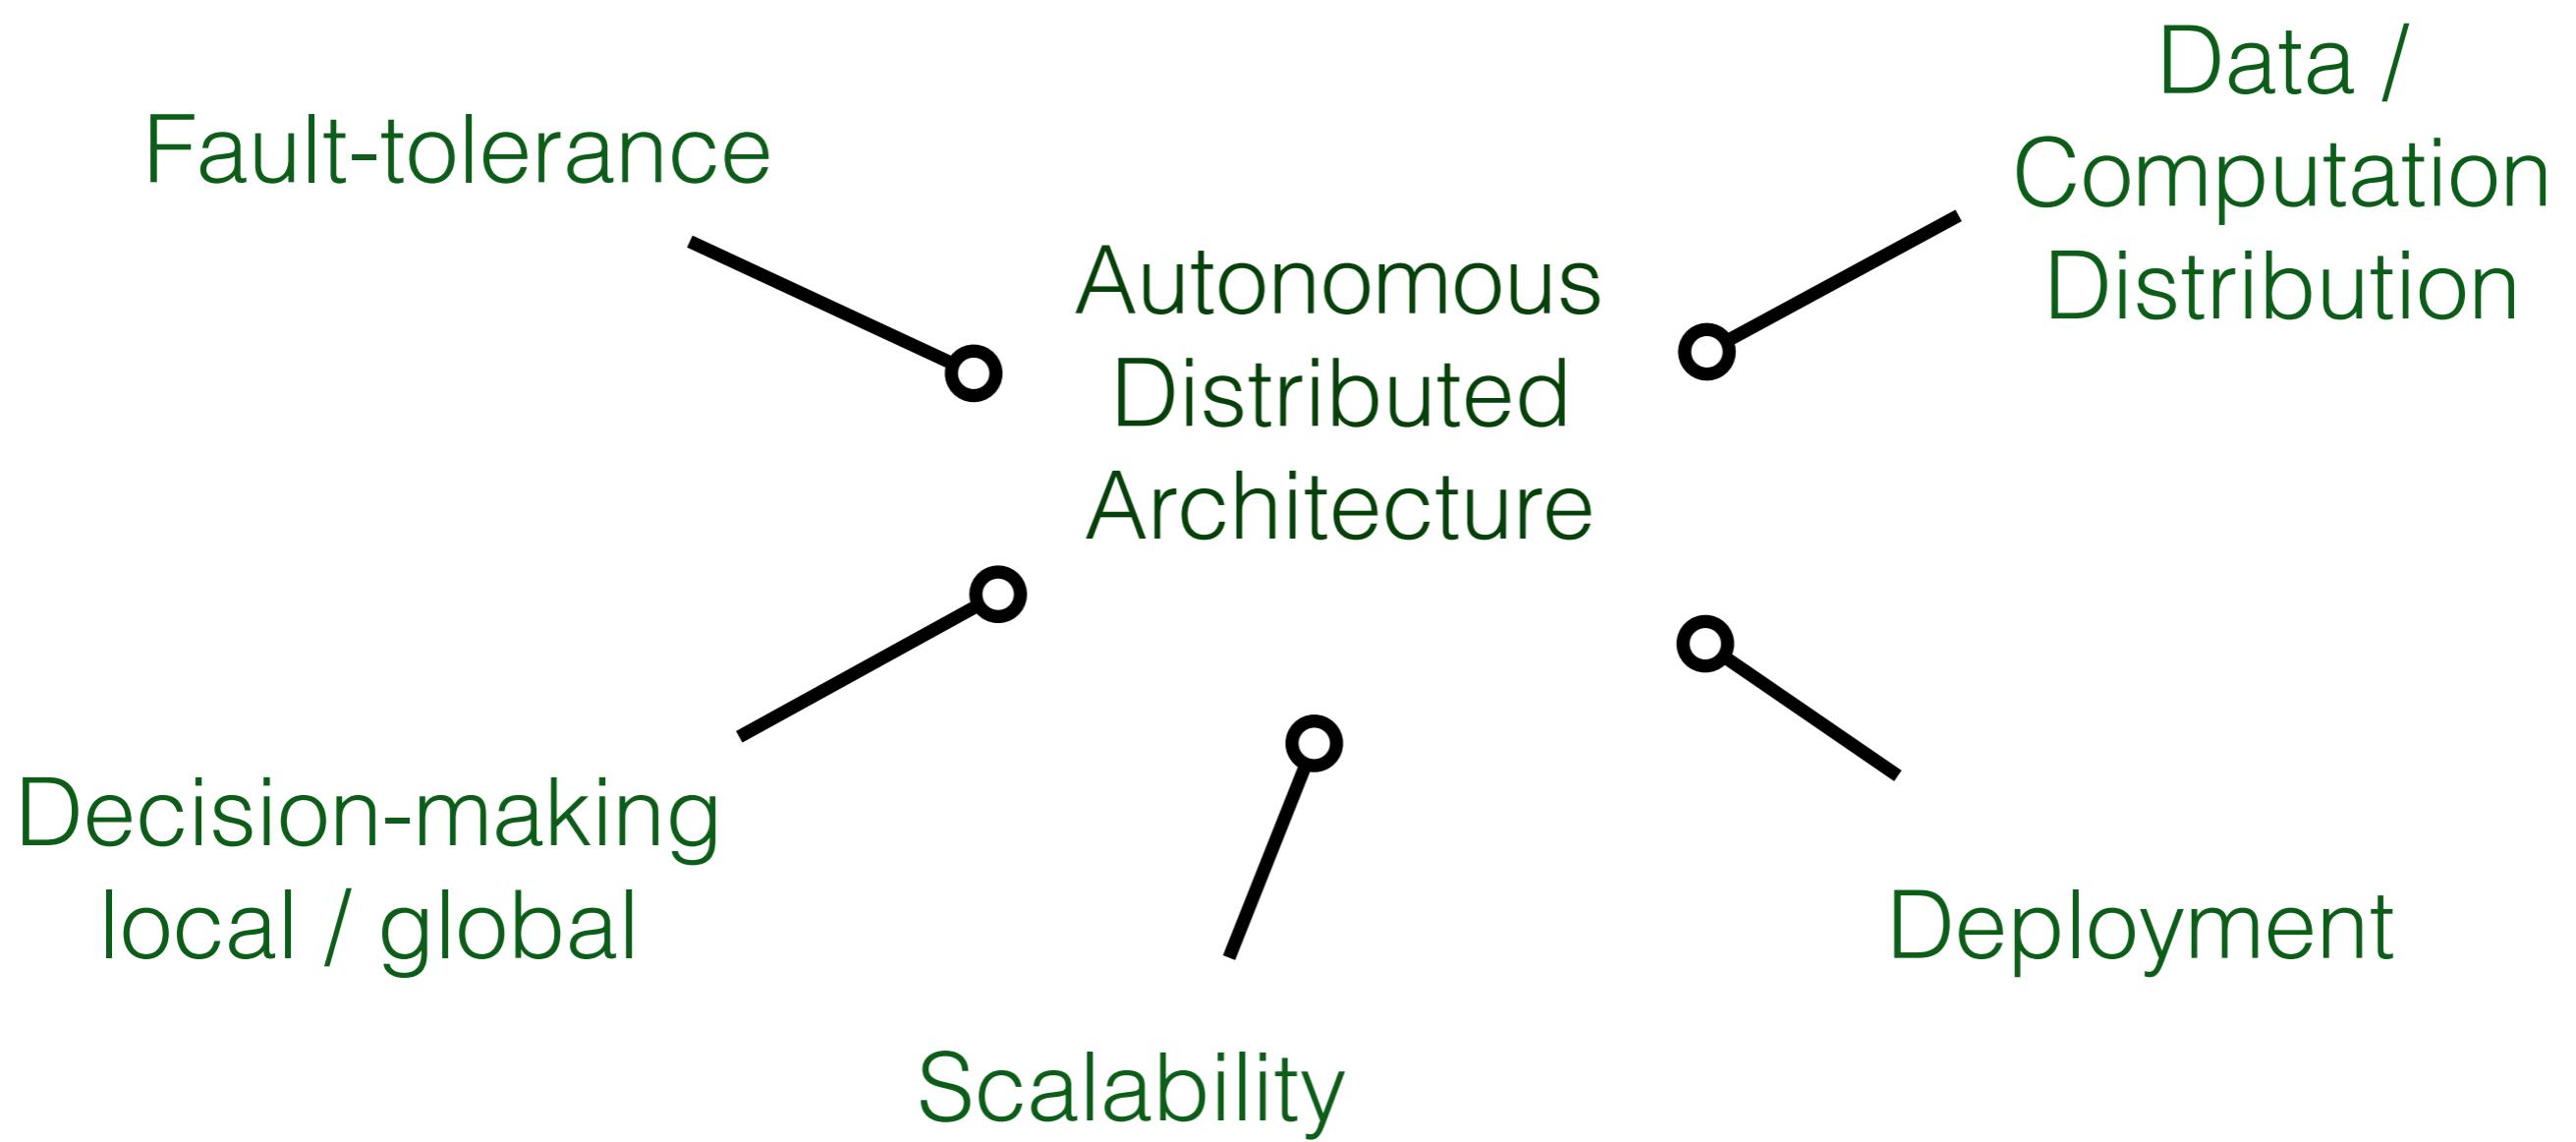
FIT IoTLab - <https://www.iot-lab.info/deployment/lyon/>



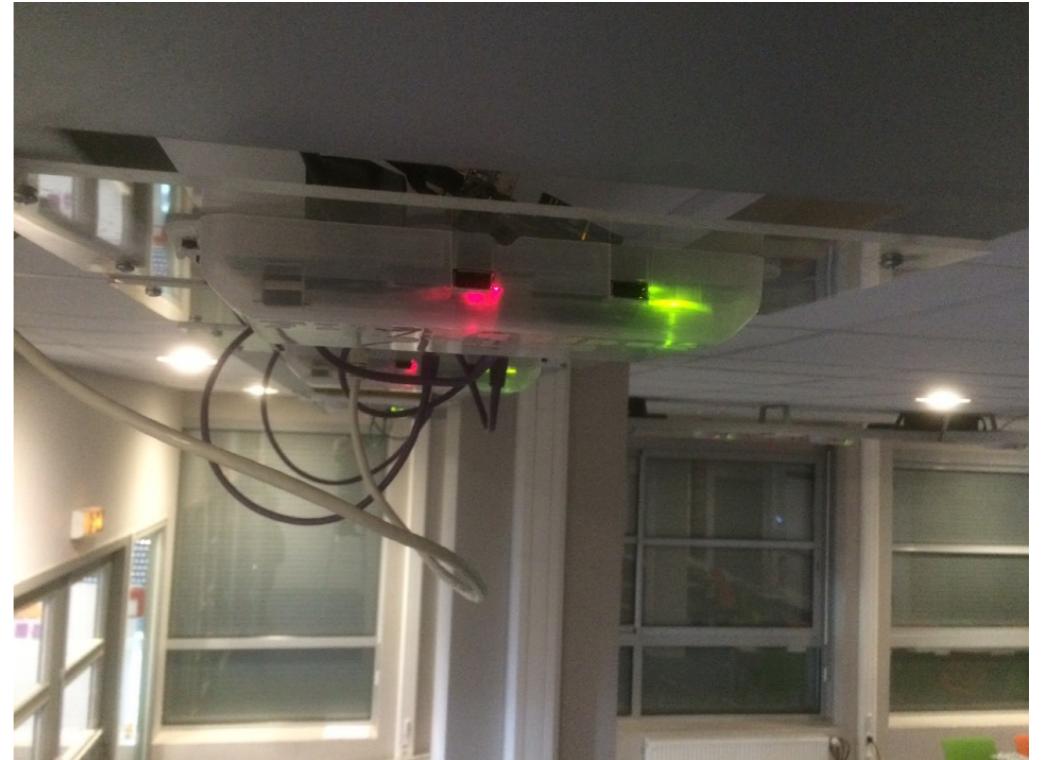
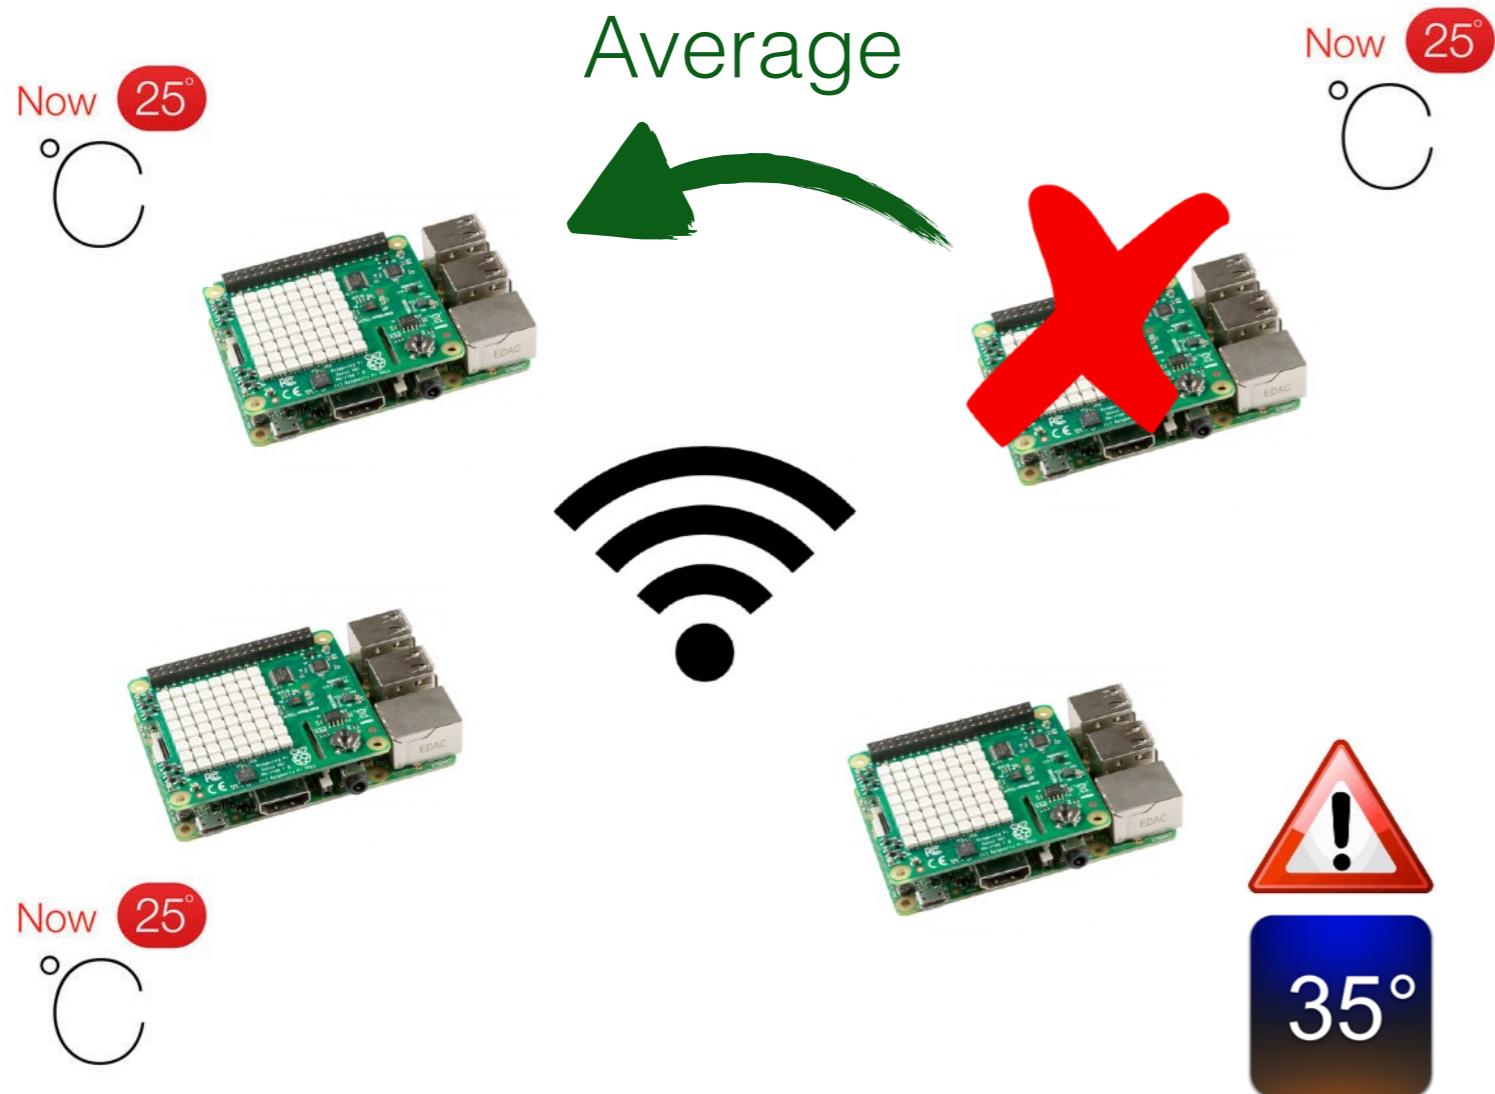
12 TurtleBot 2
Drone Parrot Bibop



Software / middleware issues for IoT



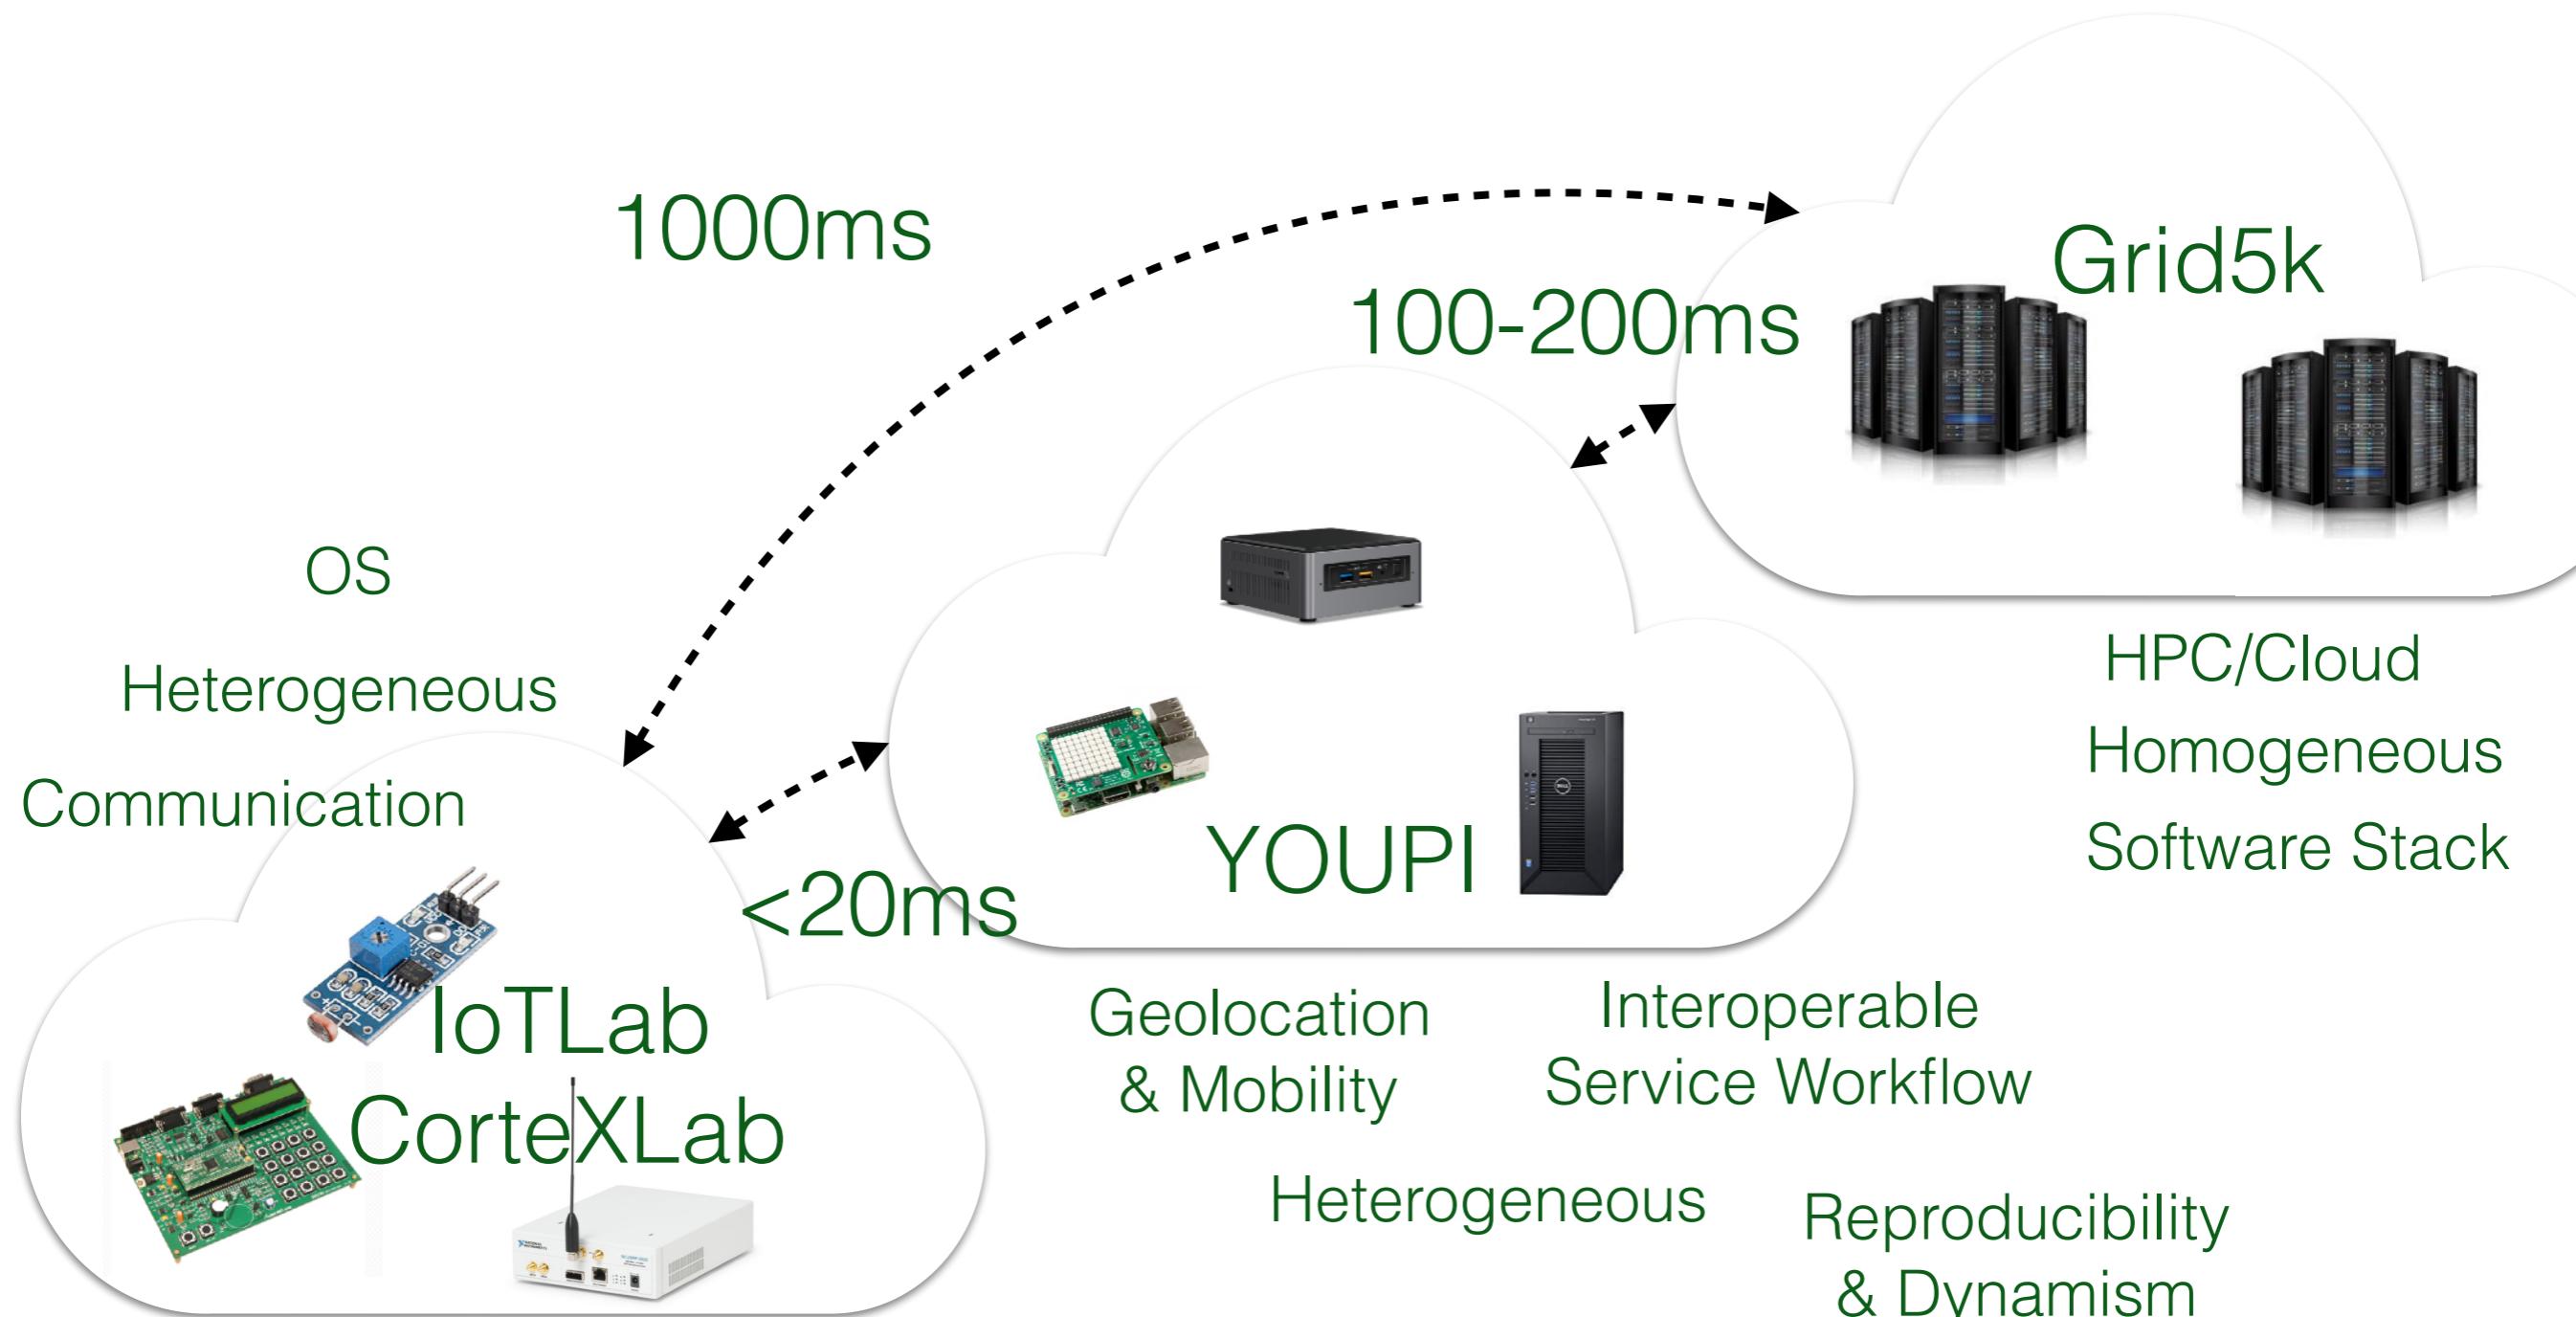
YOUPI - Fog/Edge Computing



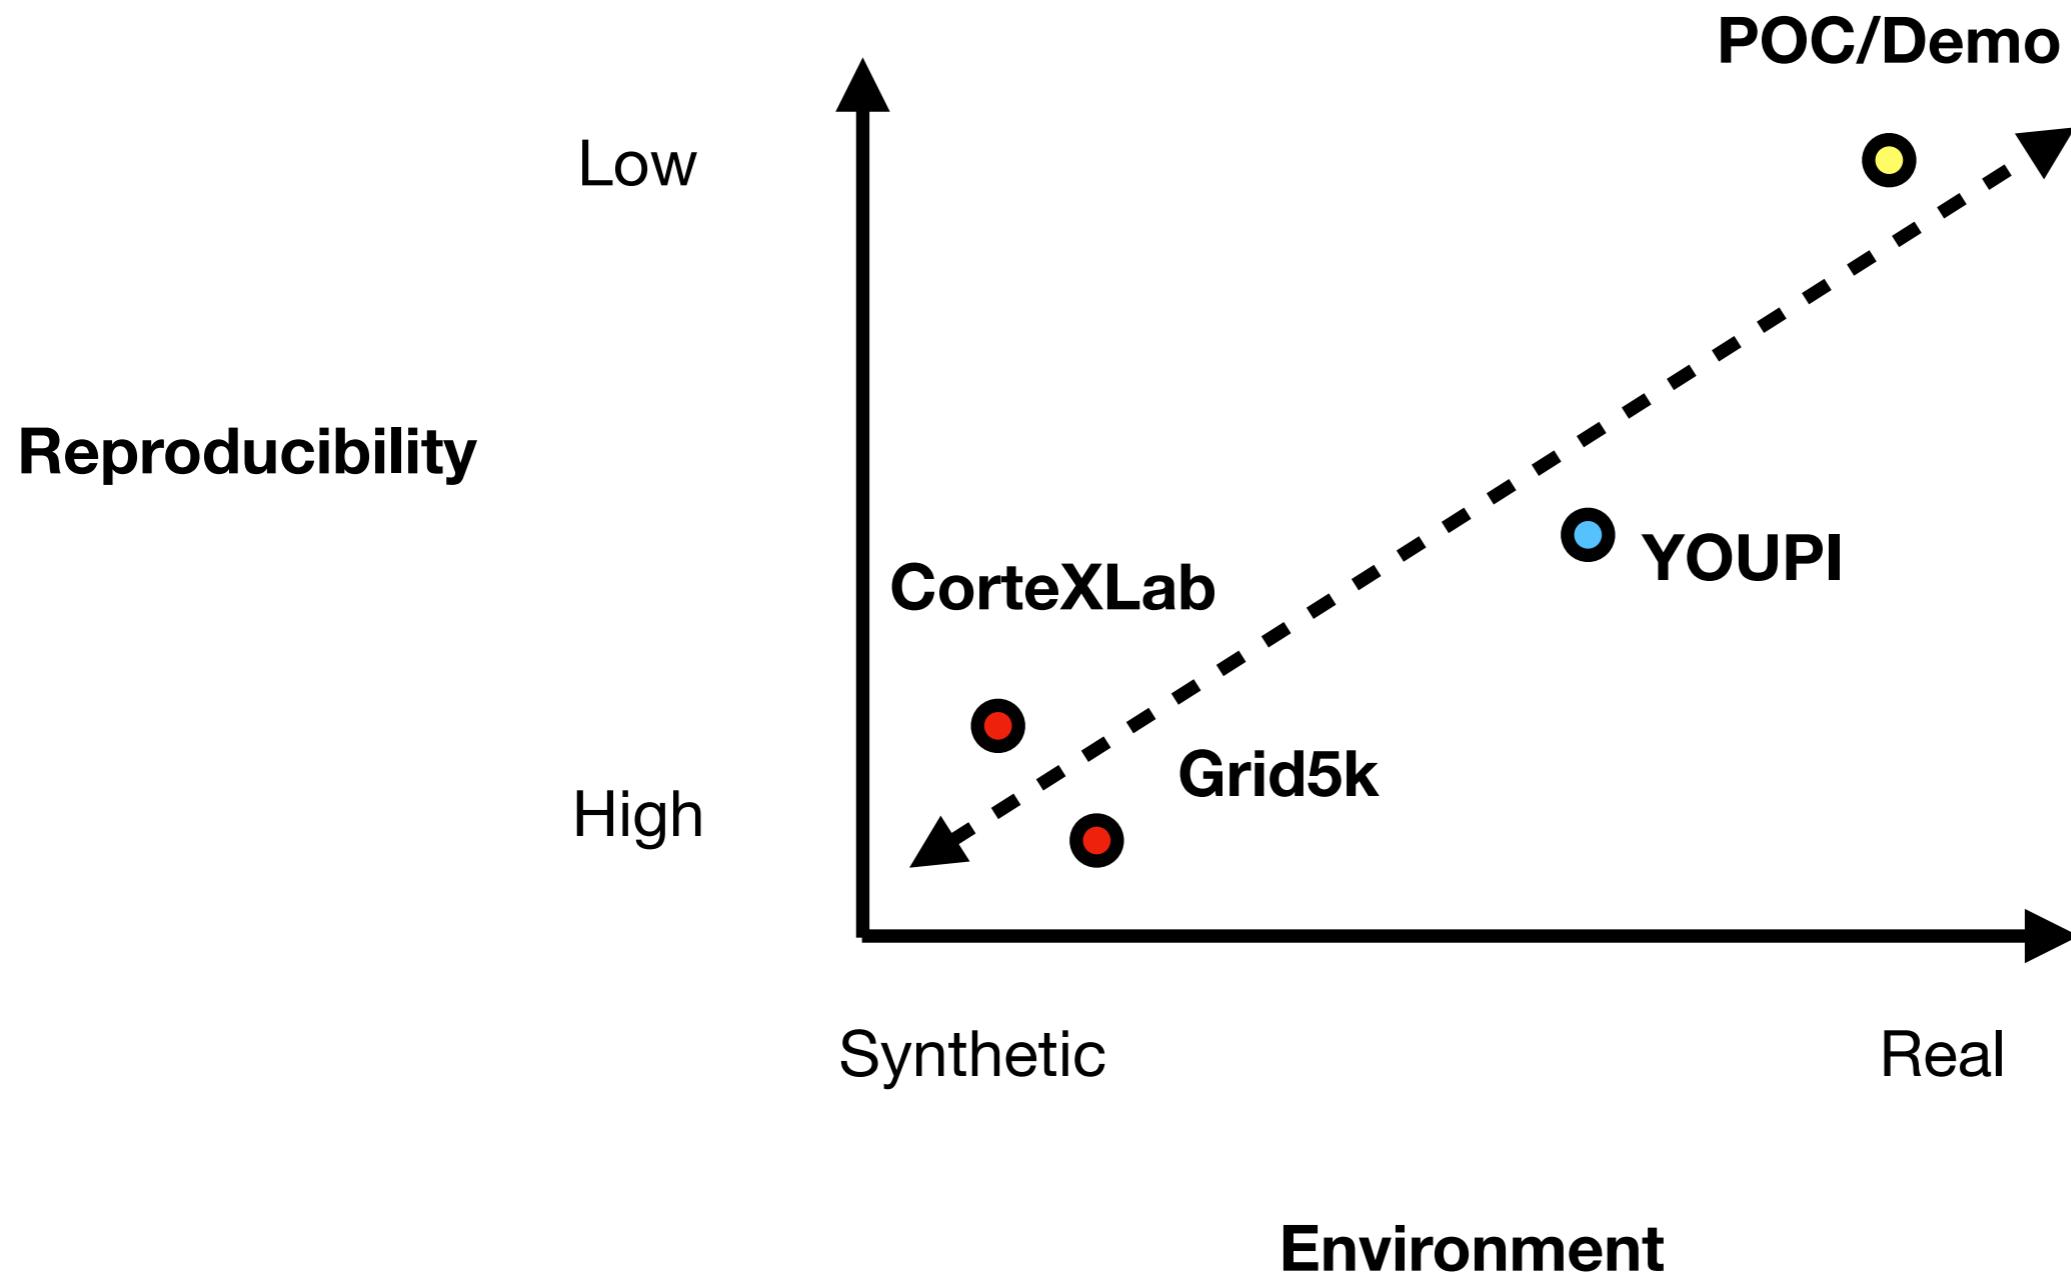
25+ RPI
3 NUC
3 Nvidia



YOUPI vs FIT(IoTLab/ CortexLab)/Grid5000 ?



Reproducibility ?



Goals

- YOUPI: Yet Another *PI platform
 - Off-the-shelf hardware and software
 - Capturing dynamic properties of the environment - not reproducible, but traceable - wireless connections/interferences, mobility, user behaviour
 - Replay/analyse offline

Heterogeneous Hardware

- 15 RPI3 + 10 RPI2 + 2 BanaPI (+ Arduinos)
 - RPI3: **ARM** Quad Core **1.2GHz** Broadcom BCM2837
64bit CPU, 1GB RAM

Heterogeneous Hardware

- 3 NUCs
 - NUC8I7HVK: Intel Kaby Lake G Core i7-8809G **3.1 GHz Quad-Core 8 thread** / GPU Radeon RX Vega M GH Graphics 4 Go HBM2 (24 Compute Units, **1063 MHz**) / **32 Go DDR4-SODIMM 2400 MHz**
 - NUC7I3BNH: Intel Core i3-7100U Dual-Core **2.4 GHz - 4 threads** / GPU Intel HD Graphics 620 (24 Compute Units, **300 MHz**) / **16 Go DDR3-SODIMM 2133 MHz**

Heterogeneous Hardware

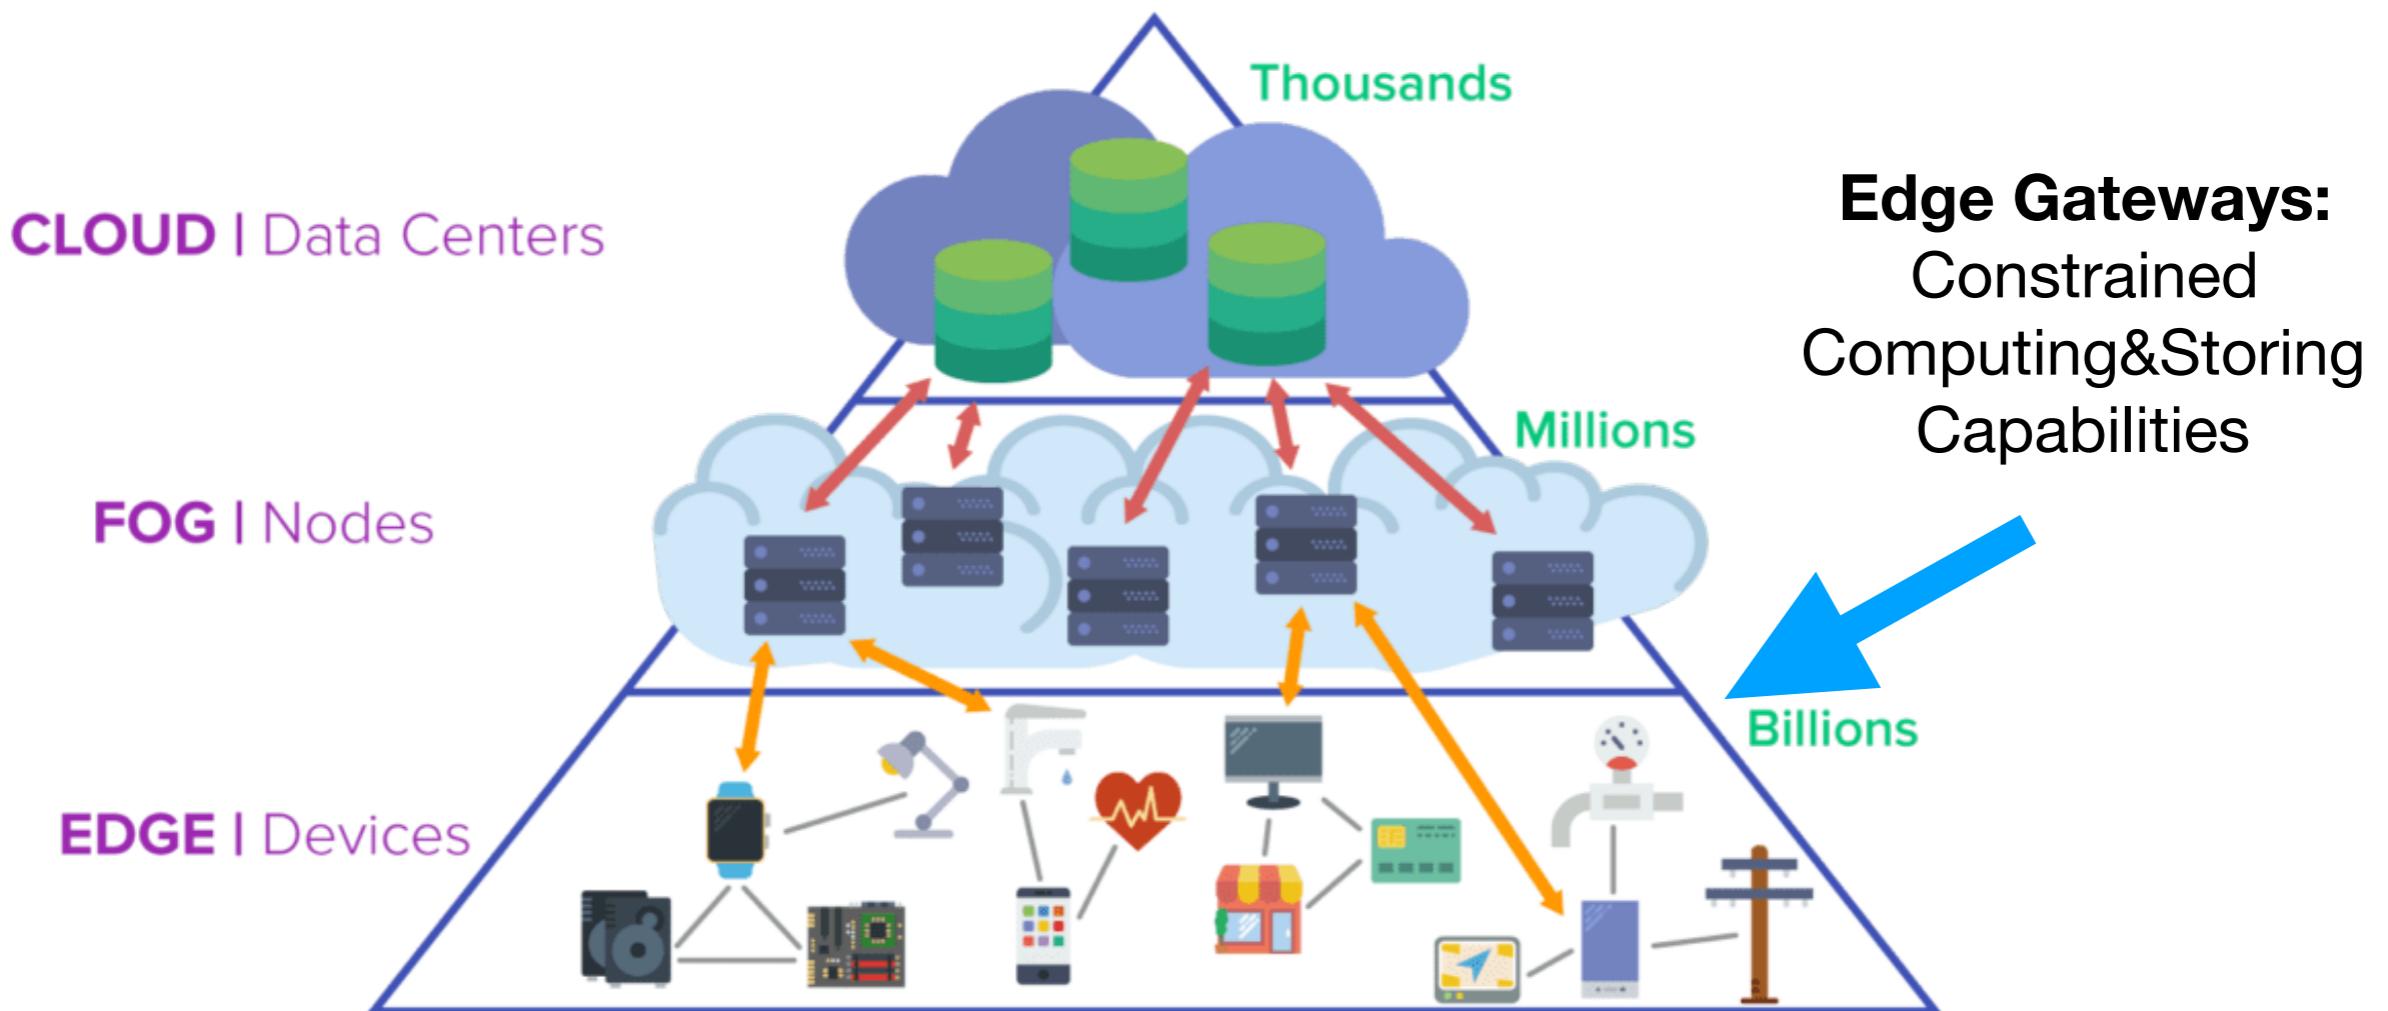
- 3 Nvidias
 - Jetson TX1: **256-core NVIDIA Maxwell™ GPU 998 MHz**
 - Nano: **128-core NVIDIA Maxwell™ GPU 921 MHz**

Heterogeneous Wireless Communication Interfaces

- Bluetooth Low Energy Short
- 6LoWPAN
- 802.11b/g/n/ac Mid Range
- LoRaWAN Long
- (Sigfox)
- + Ethernet Control Plan

Energy Management

- Power-Over-Ethernet Switch / POE splitter
 - Off-the-shell electric plugs: high failure probability
 - Energy monitoring, switch on/off nodes



Software Management

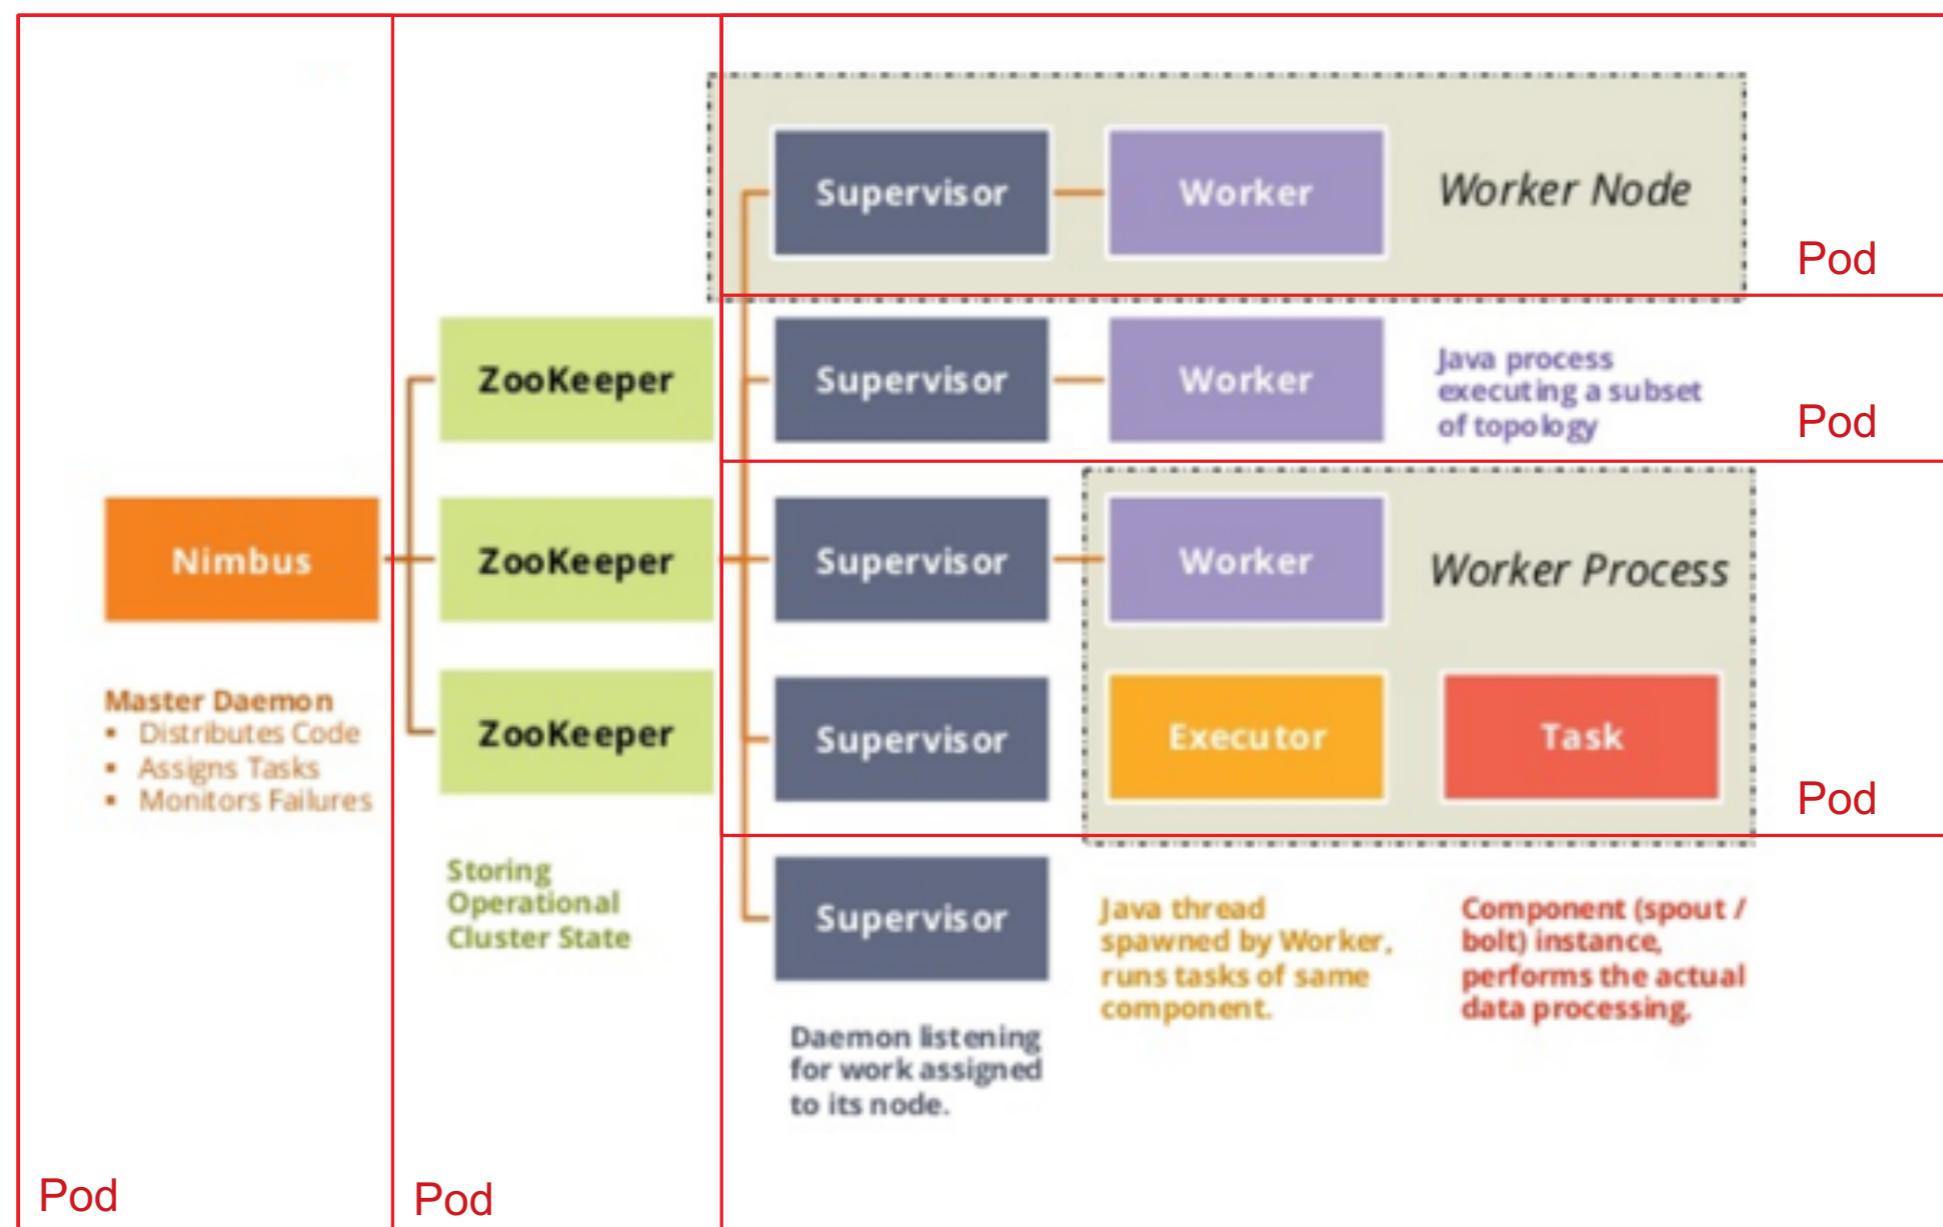
- Image Server
 - Network PXE boot / reflash all nodes / specific configuration
- Off-the-shell containers
 - Docker
 - Kubernetes

2 experiment feedbacks

- Streaming application deployment
- Edge topology capacity

Apache Storm deployment

Joann Felipe Castellanos



Out of memory

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: storm-worker-controller
  labels:
    name: storm-worker
spec:
  replicas: 3
  selector:
    matchLabels:
      name: storm-worker
      uses: nimbus
  template:
    metadata:
      labels:
        name: storm-worker
        uses: nimbus
    spec:
      containers:
        - name: storm-worker
          image: daksher/storm-worker
          resources:
            limits:
              cpu: 200m
              memory: 500Mi
            requests:
              cpu: 100m
              memory: 100Mi
          ports:
            - hostPort: 6700
              containerPort: 6700
            - hostPort: 6701
              containerPort: 6701
            - hostPort: 6702
              containerPort: 6702
            - hostPort: 6703
              containerPort: 6703
```

```
pi@pi1:~ $ kubectl get pods
NAME                               READY   STATUS    RESTARTS   AGE
nimbus-694fc8cbcfc-p7hm7           1/1    Running   0          6h25m
prometheus-deployment-648fc4ff4-ncmjh 1/1    Running   0          8h
prometheus-deployment-648fc4ff4-v9g8j 1/1    Running   1          8h
storm-worker-controller-58c99f8c9d-ptlnv 1/1    Running   0          6h4m
storm-worker-controller-58c99f8c9d-r78bj 1/1    Running   0          6h4m
storm-worker-controller-58c99f8c9d-z2lc6  1/1    Running   0          6h4m
zookeeper-cddf49659-4qg4z           1/1    Running   0          8h
```

**1G RAM : infinite loop of oom-killed,
service relaunch every 10s**

FogZimulateur

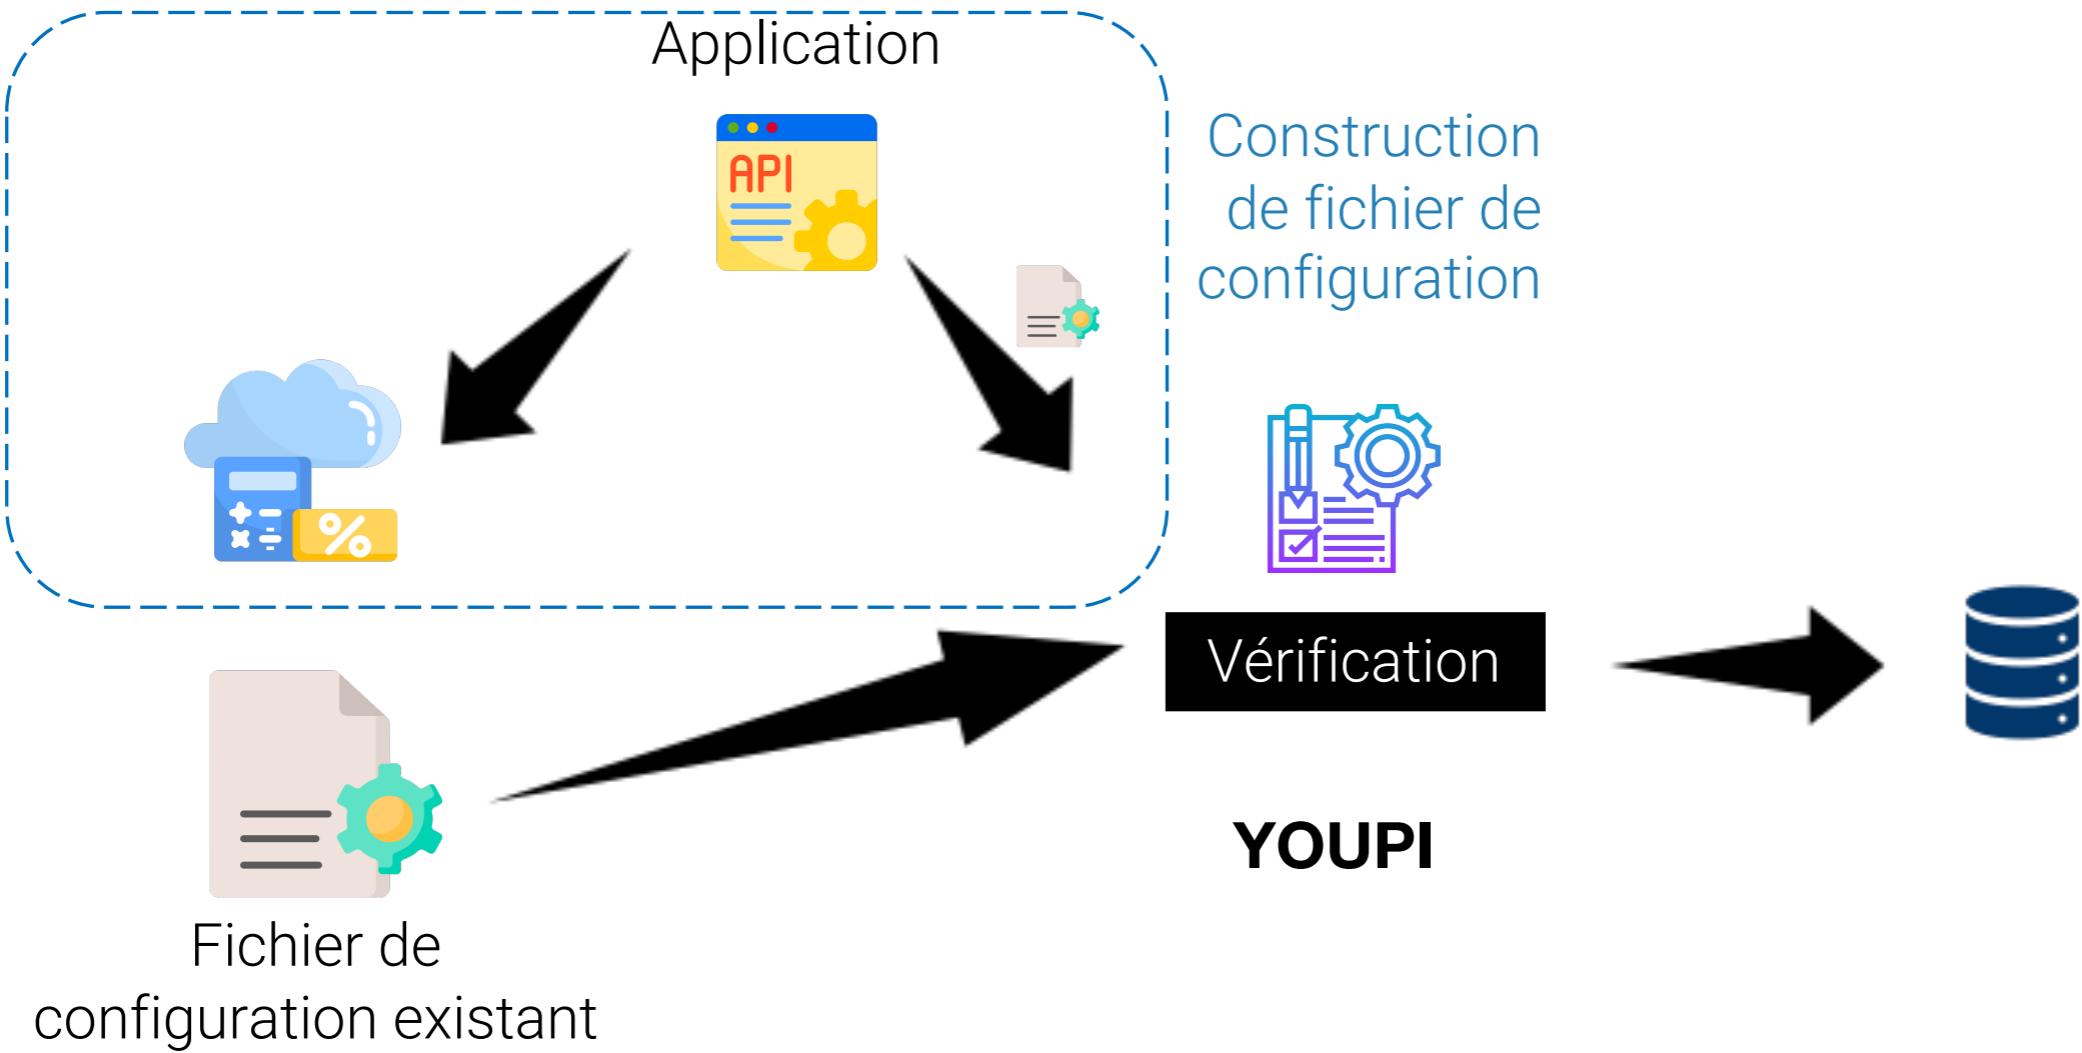
Zakaria Fraoui



Eddy Caron



-  Portail web collaboratif, accessible et multi-utilisateurs
-  Tableau de bord ergonomique et interactif
-  Simulation et prédition des performances des plates-formes fog computing ,
-  Génération des graphes et des statistiques





Cloud
Quantity : 1
Throughput : 660. Kbps



Fog-devices
Quantity : 2
Throughput : 660. Kbps



Sensors
Quantity : 12
Throughput : 720. Kbps



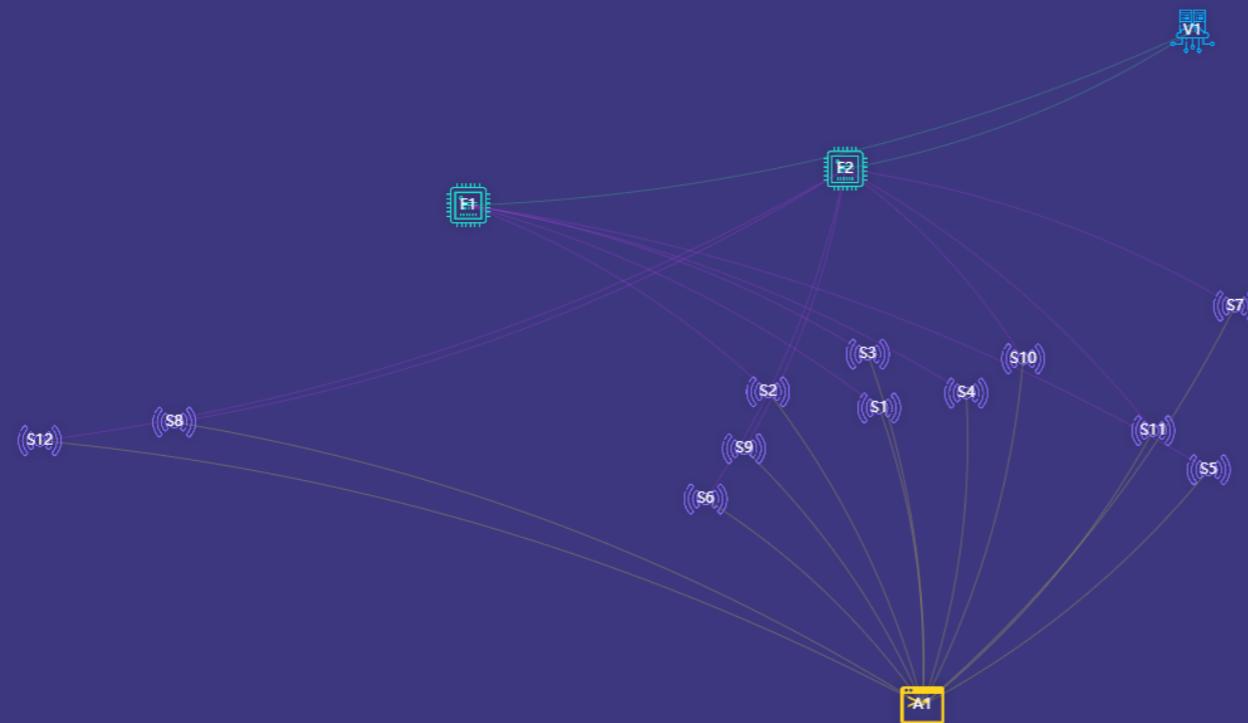
Applications
Quantity : 1
Throughput : 60.0 Kbps

Platform

SIMULATE

Platform Topology

Application IoT Edge Cloud



Throughput Simulation (Kbps)

	Throughput	Change	Index
1	88.80	▲ 0%	1
2	177.6	▲ 88%	2
3	222.0	▲ 44%	3
4	266.4	▲ 44%	4
5	355.2	▲ 88%	5

A1

ON OFF

Description

weather

Throughput

60 Kbps

Hardware Requirements

1 GB

100 Mb

1.2GHz 64bit CPU

ADD

MODIFY

DELETE

statistics

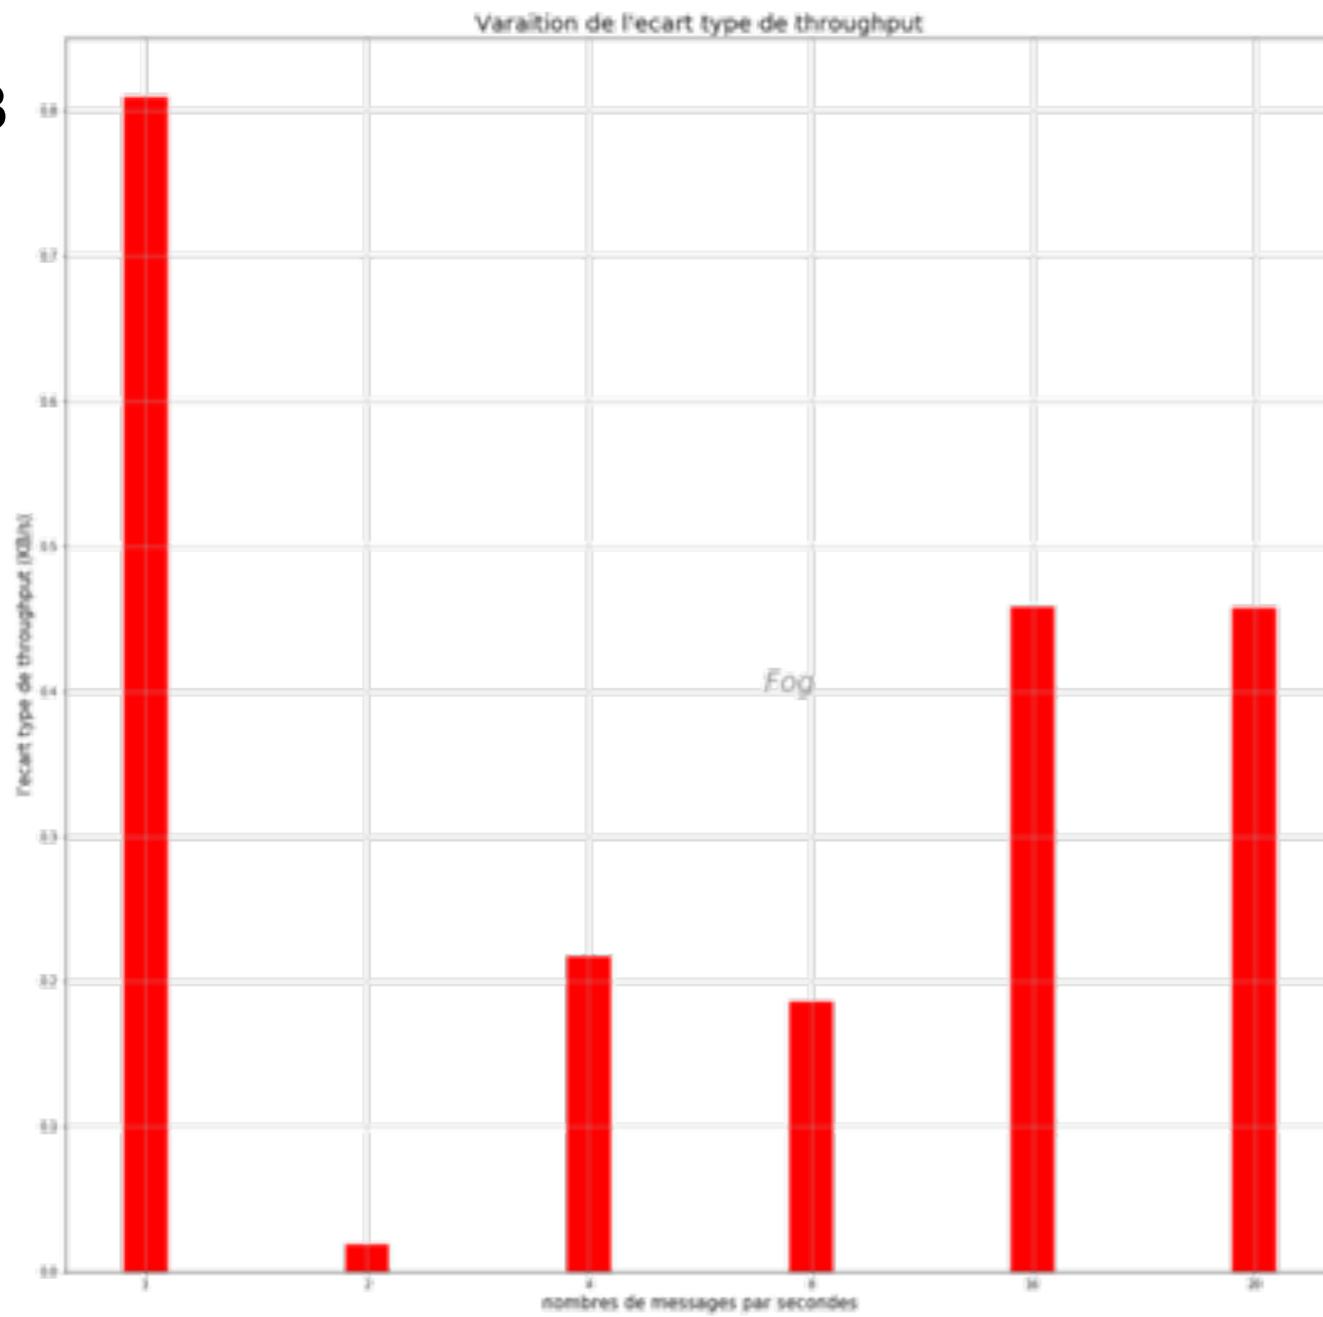
Throughput(Kbps)



Upload a Platform

Choisir un fichier citi3.json

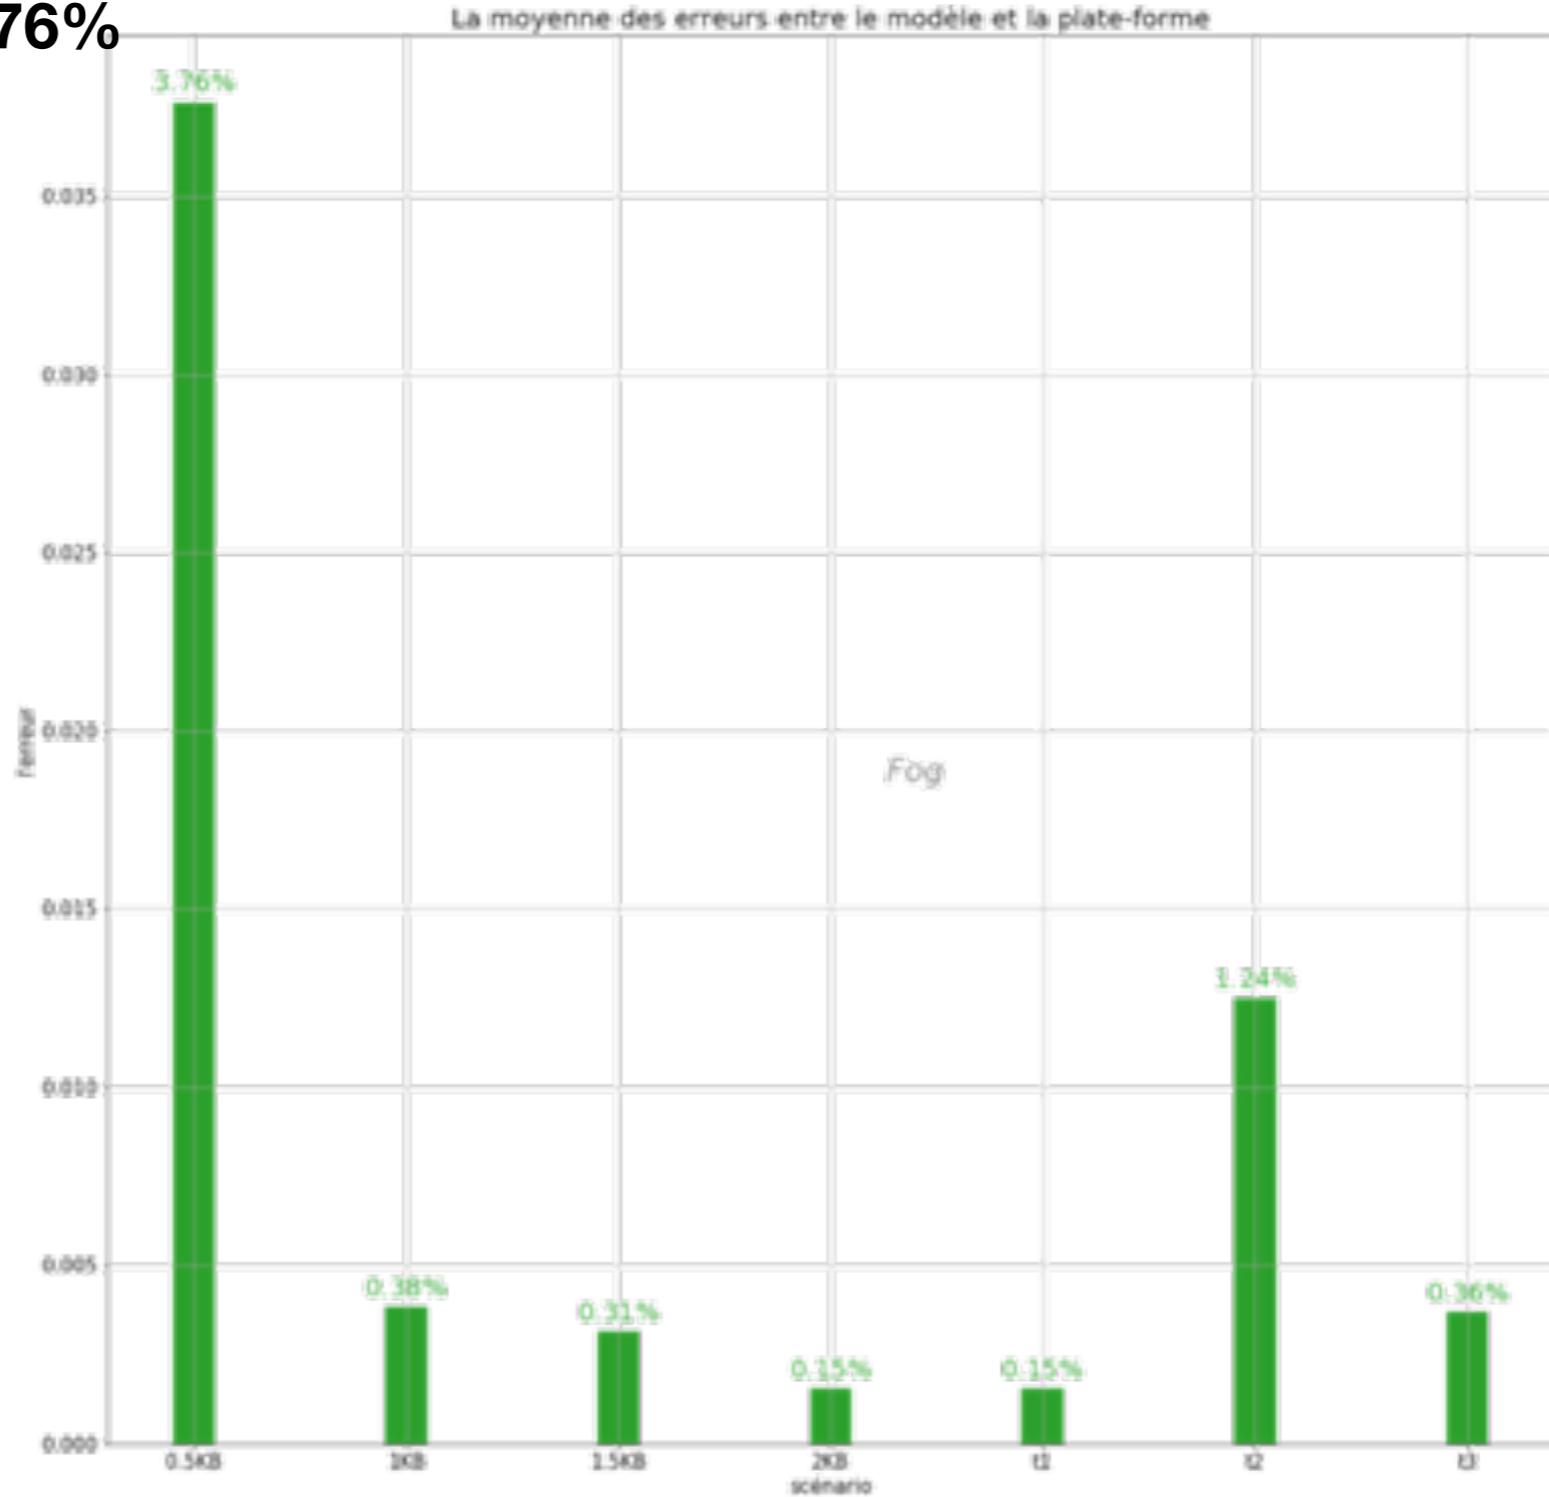
CREATE THE PLATFORM



L'écart type de throughput du fog est inférieur à 1

L'erreur maximum entre le modèle et les résultats de tests est *inférieur* à 4%

3.76%

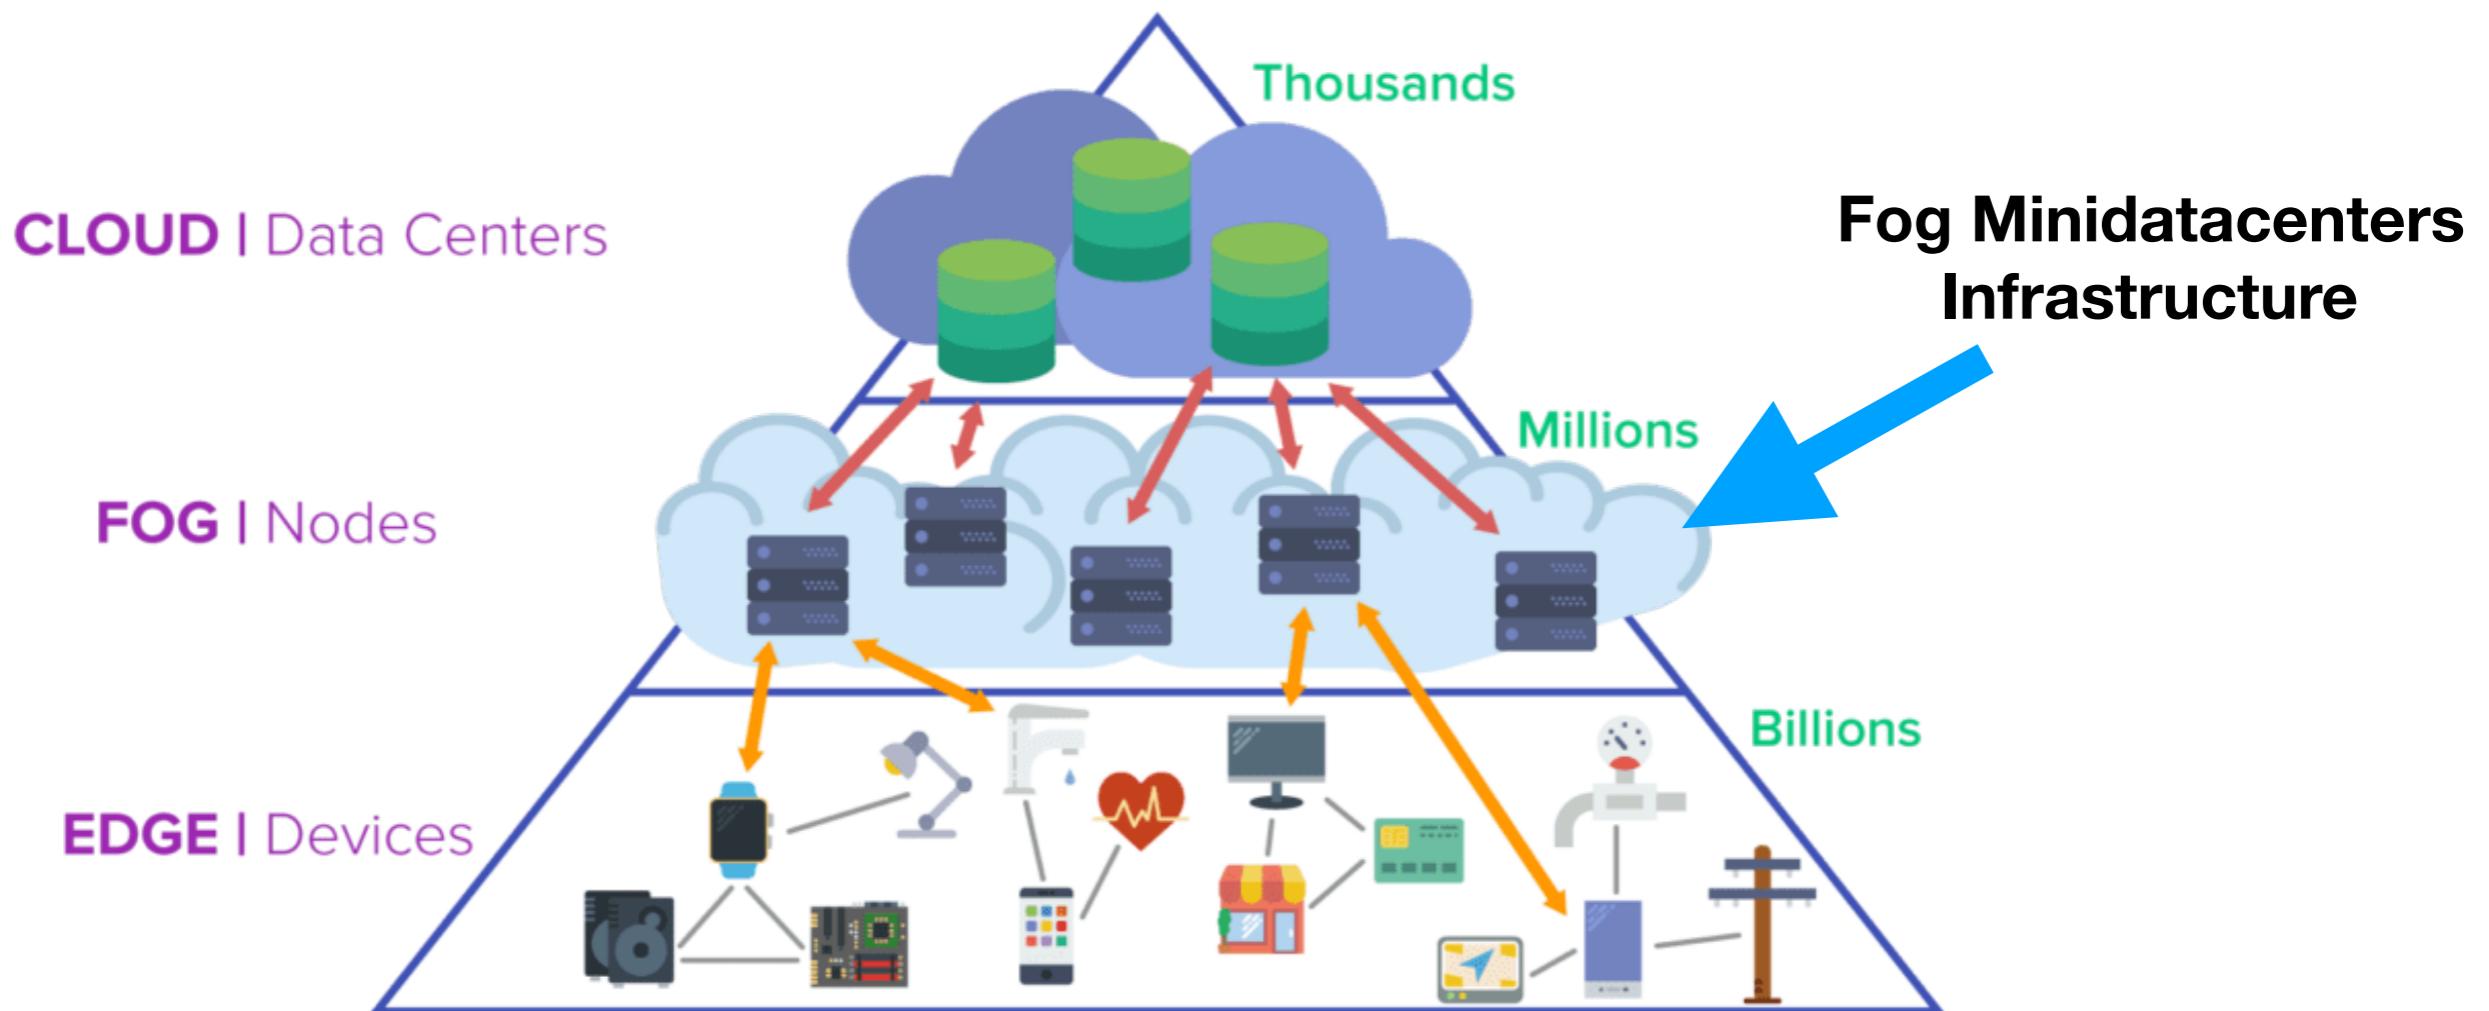


Future research work

- Off-the-shell electronic components: significant failure probability (electric plug, SD cards), normal behaviour
 - Highly-dynamic local fault-tolerance at the Edge? vs Energy?
- VXLAN, tunnelling SSH: encapsulation, not acceptable delay for Edge Computing
 - Efficient cross-domain network management?
- Several dynamic scenarios in 5G (IoT, vehicle, ...)
 - Computing/storing services handoff in multi-proximity clouds?

Future research work

- Fog Computing : SILECS Lyon - INSA & ENS
 - CPER project proposal
 - Interconnecting Gerland/La Doua Campus - City mobility



Discussions ?

- CITI <http://www.citi-lab.fr>
- Dynamid <http://dynamid.citi-lab.fr>
- Chroma <https://team.inria.fr/chroma>
- Agora <https://team.inria.fr/agora>

