ECO-ICT SCHOOL
7-11 Oct 2024

V. Lala, A. Pirovano, J. Radzik

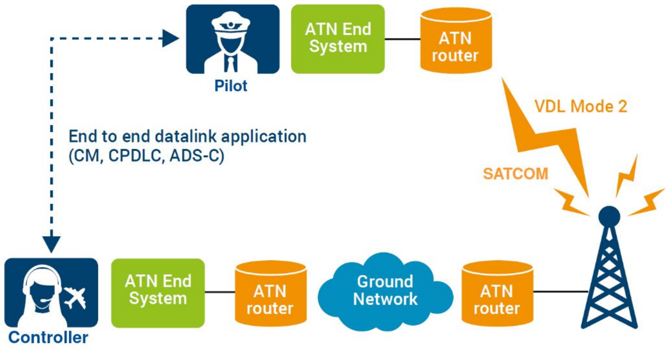


**Aeronautical communications**

Upcoming migration to the ATN/IPS for air-to-ground datalink communications

**TCP/IP Networks**

Emergence of new transport protocols for TCP/IP networks (TCP Cubic, QUIC, MPTCP, ...)

**Which set of end-to-end reliability mechanisms is best suited to datalink communication for the future ATN/IPS?**

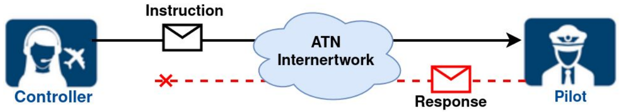


**Where is the reliability 🤔 ?**

# Where I work

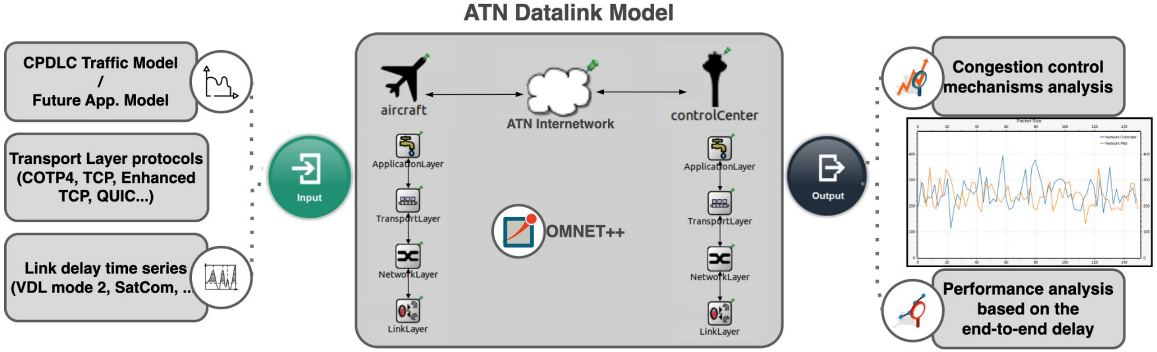


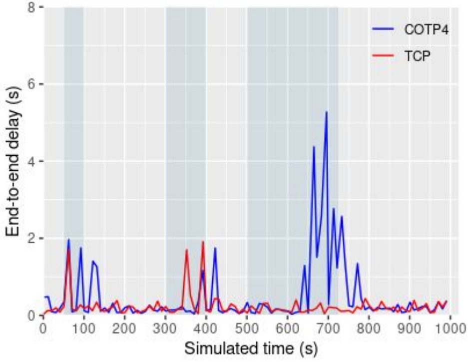


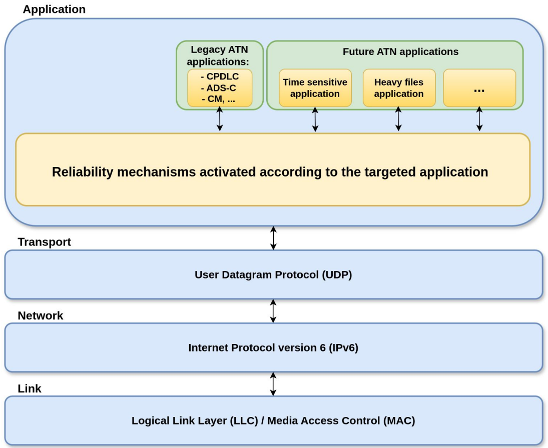


# How I study it

RÉPUBLIQUE FRANÇAISE
ENAC
isae SUPAERO

# Designing Fog Infrastructure for Environmental Monitoring

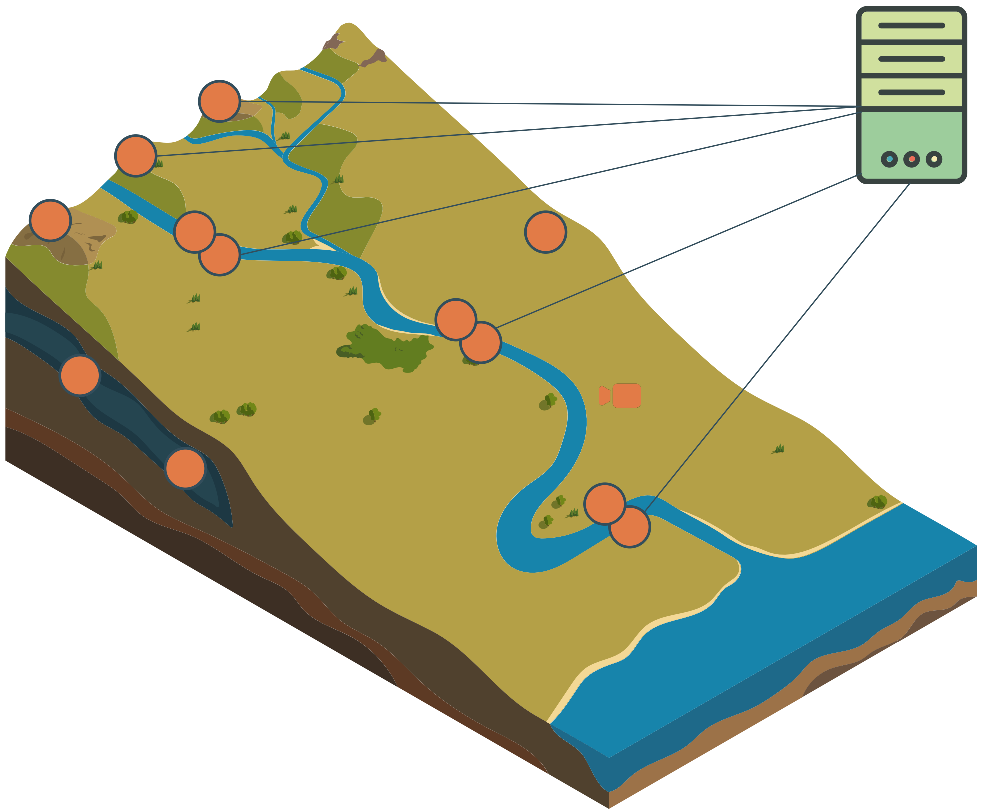

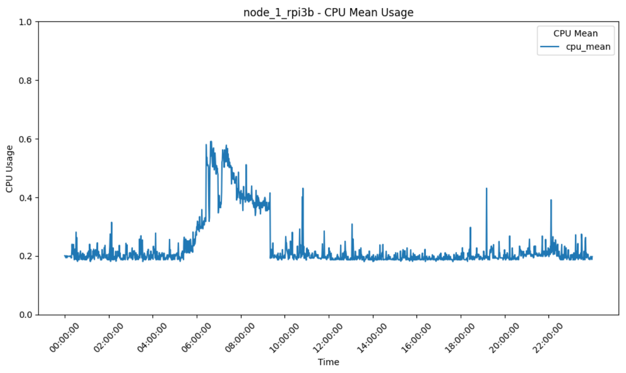


Event driven workload

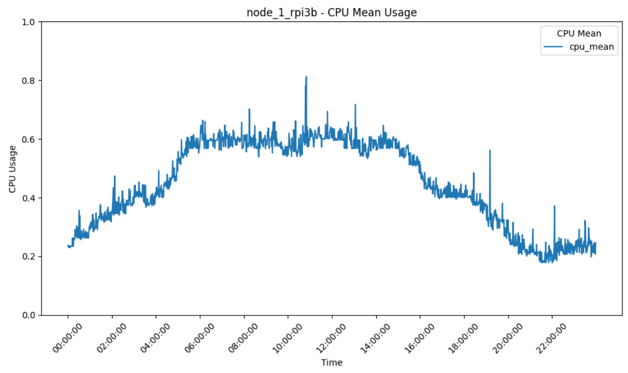


Long term workload

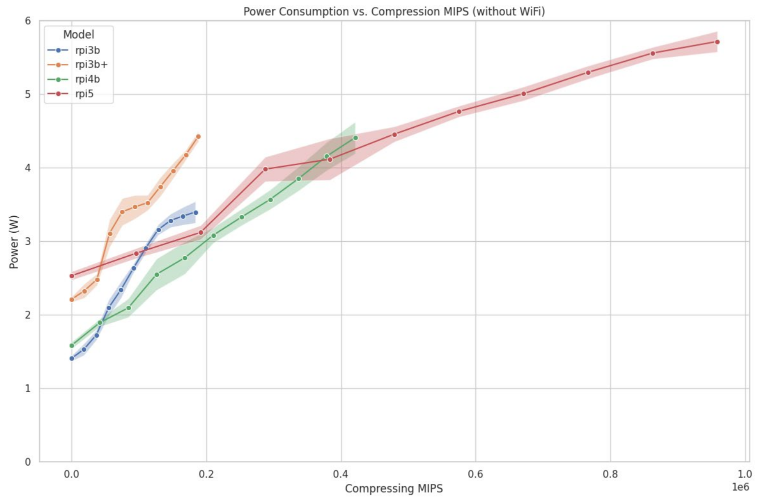


UMR IRISA

Université de Rennes

GEOSCIENCES Rennes

*Ammar KAZEM, email: ammar.kazem@irisa.fr*

# Energy-Efficient Traffic Steering for Slicing 5G via Deep Reinforcement Learning

**André-Cédric BESSALA**, Guilherme IECKER RICARDO, Gentian JAKLLARI

IRIT, Universite de Toulouse, CNRS, Toulouse INP, UT3

**Eco-ICT 2024 school**

## Background

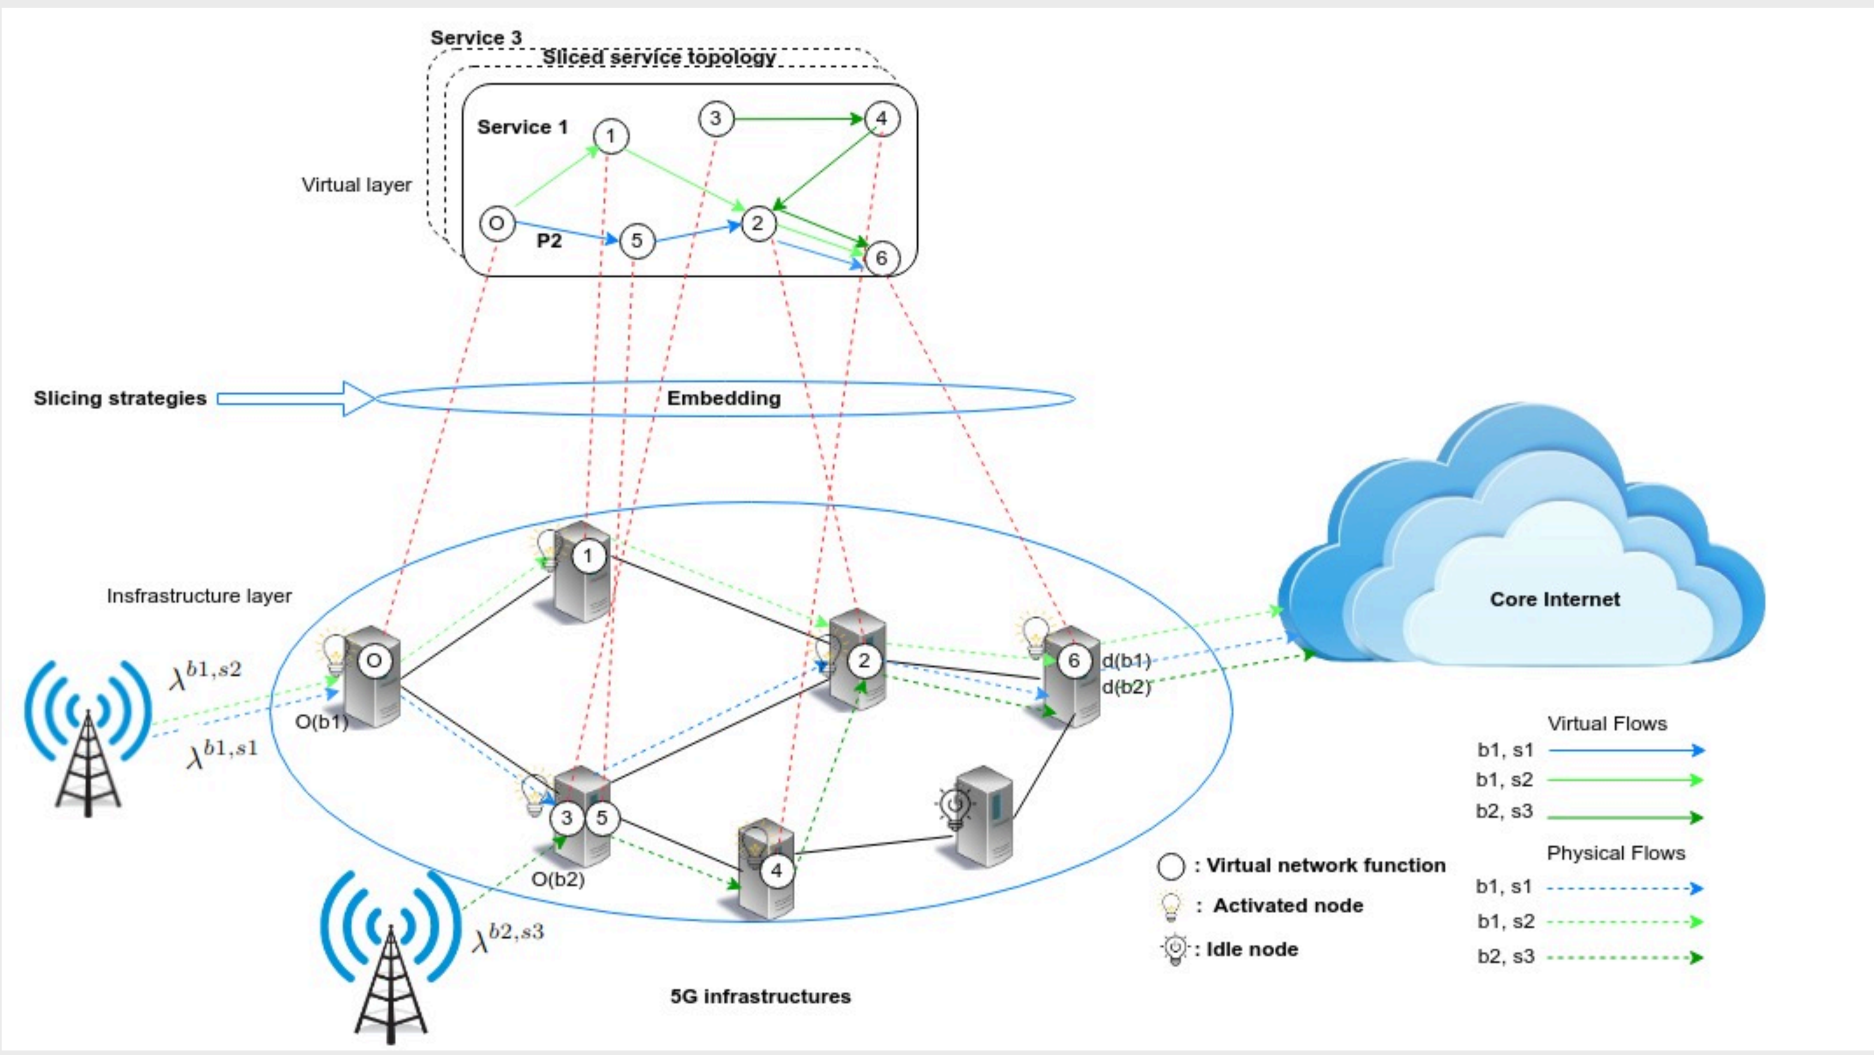


Fig. 1: 5G Infrastructures

5G System :

1. **RAN** Radio Access Network that introduces new radio technologies.
2. **Core network** that relies on softwarization and a cloud-native architecture.
3. **SDN**-enabled transport network that interconnects the RAN and the core network.

**Network slicing is a technology to enables the partitioning of the network infrastructure into multiple tenants networks, each tailored for certain QoS for their users.**

## Problem formulation

Given a network topology and a virtualized slice service pool : How can we guarantee both slice performance and low energy consumption ?

1. We assume that we can optimize energy efficiency in the network with resource allocation and traffic steering.
2. What are our resources: network bandwidth and processing power.
3. Objective: solve energy-efficient 5G slicing placement problem.

1. **Which physical machines will uses to host each VNF of slice service provider request ?**
2. **How to map the VNFs ⇒ how to handle SFC ?**
3. **How to re-routing traffic between VNFs ?**
4. **How to gaurantee low network energy cost for NSPR provisionning ?**

We address in this problem four challenges: energy saving, resource utilization, network traffic and QoS, then the problem can be modeled as NUM (1)(2)(Network Utility Maximization) problem :

$$\mathrm{U}(x,y,z) = \mathrm{E}(x,y,z) - \alpha(C_{ij} - \sum_{b\in B}\sum_{s\in S} \hat{\lambda}^{b,s} \cdot z^{*}_{ij}) \quad (1)$$

## Methodology

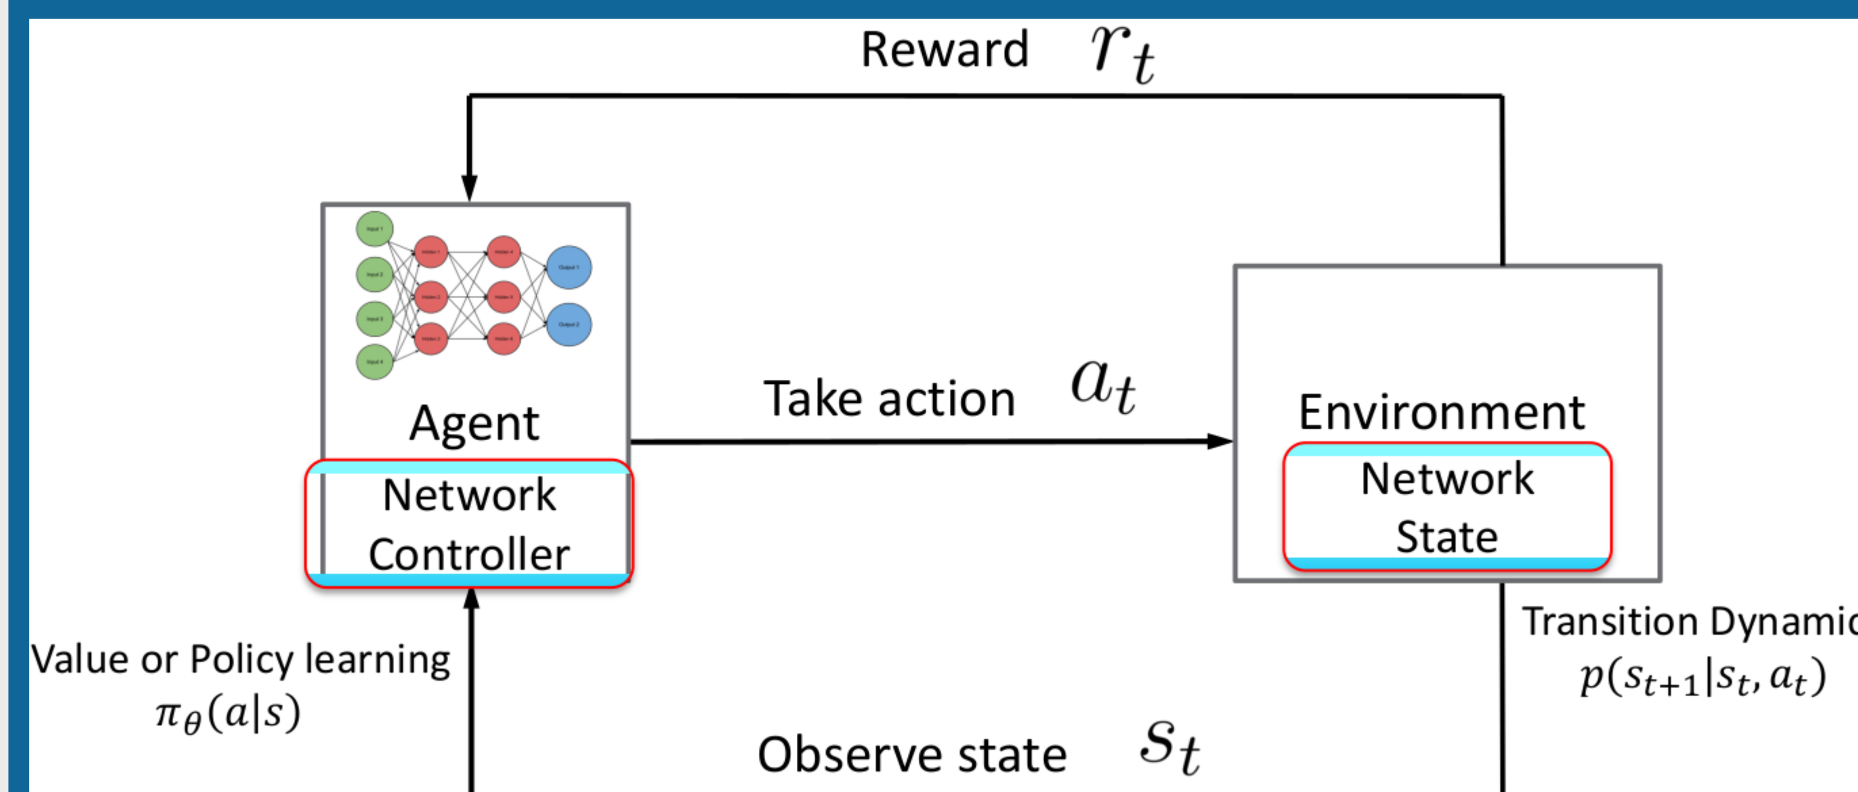


Fig. 2: We propose RL to solve the problem

**MDP of the problem :**

1. State Space : $S = (\hat{x}_i, \hat{y}_i^{b,s}, \hat{z}_{ij}^{b,s}, \hat{\lambda}^{b,s})$
2. Action Space: find the optimal configuration for each slice request;
3. Reward: utility function of the problem;
4. Policy: specifies how the controller chooses its action given the observation;
5. Goal: find optimal policy for the controller that maximizes the expected return.

## Design and Evaluation

- ➢ Working with Network Emulators, Ryu controller, mininet, networkx for simulating network behavior
- ➢ We designing our network states and actions in the OpenAI Gym environment
- ➢ Agent learning algorithm: DQN

- — Reward function and objective function
- — Energy saving
- — Link load
- — Path stretch for energy solution vs shortest path for each request

## References

[1] C. Liu, P. Wu, M. Xu, Y. Yang, and N. Geng, "Scalable deep reinforcement learning-based online routing for multi-type service requirements," *IEEE Transactions on Parallel and Distributed Systems*, vol. 34, no. 8, pp. 2337–2351, 2023.

[2] Z. Xu, J. Tang, J. Meng, W. Zhang, Y. Wang, C. H. Liu, and D. Yang, "Experience-driven networking : A deep reinforcement learning based approach," in *IEEE INFOCOM 2018-IEEE conference on computer communications*. IEEE, 2018, pp. 1871–1879.

# Reliable and cost-efficient Data Placement and Repair in P2P over immutable Data

**Context:**

- Data volume remains a major challenge for distributed systems and big data communities.
- Appealing solution is to utilize the storage resources available on connected devices.
- This requires addressing several issues including node failure, node availability (churn), and how to guarantee data availability and avoid data loss.
- Replication provides high data availability, BUT it incurs high storage overhead and large network transfers.
- Erasure Coding (EC) offers high data availability with minimal storage overhead, countering replication issues.
- Objective: How can EC be efficiently implemented and optimized in P2P storage systems ?

**Completed Work:**

- IPFS (Interplanetary file system) is an open source and is widely used a P2P file sharing system that relies on content addressing,DAGs and distributed hash tables.
- We have implemented Reed-Solomon(RS) erasure codes into the Go implementation of IPFS.

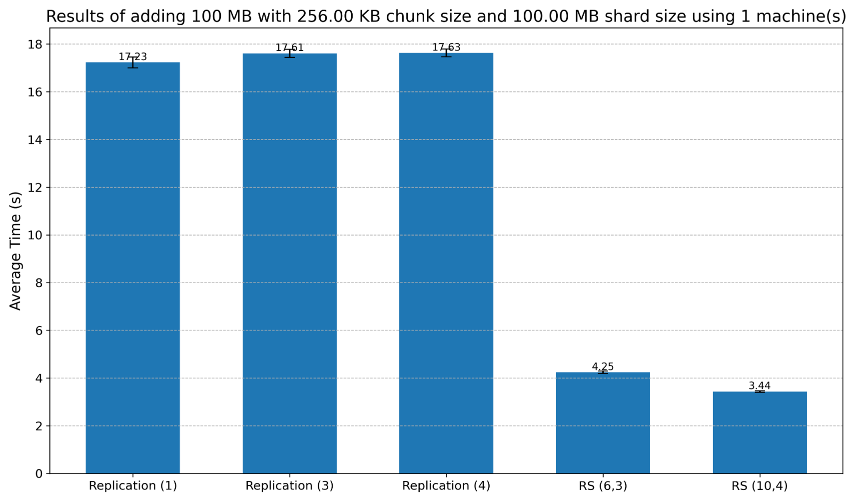


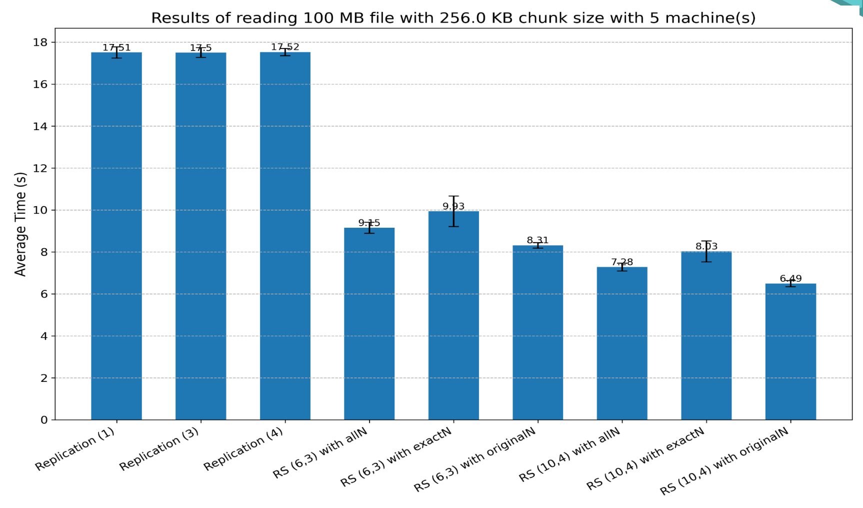


**Current Works:**

Implementing repair jobs under EC and I will work on designing and implementing new data placement strategy.

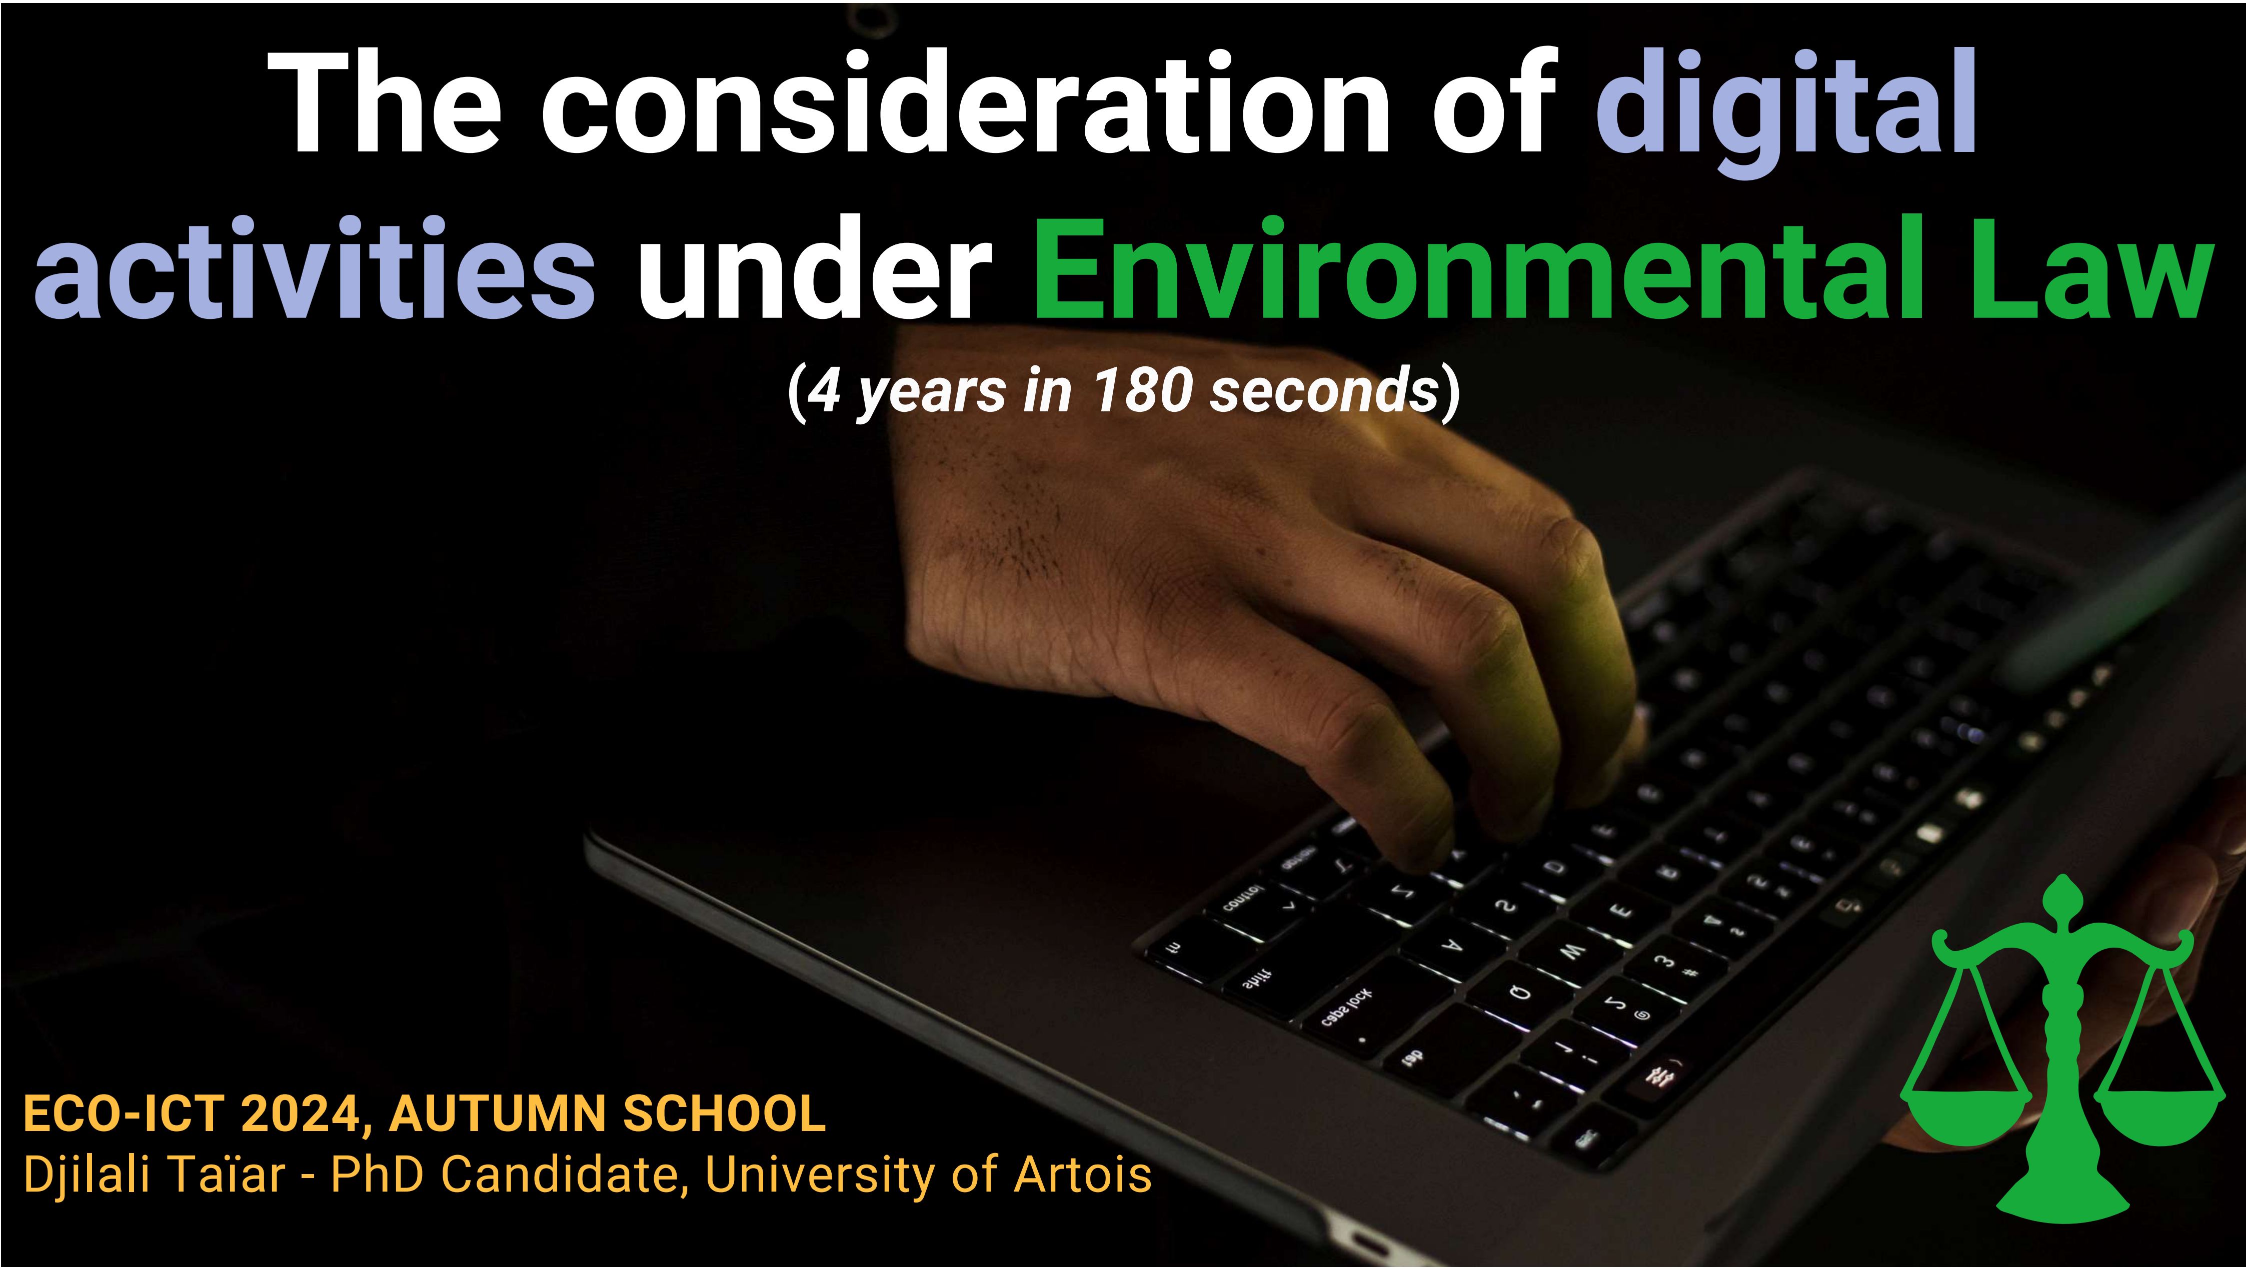
The consideration of digital activities under Environmental Law
(4 years in 180 seconds)
ECO-ICT 2024, AUTUMN SCHOOL
Djilali Taïar - PhD Candidate, University of Artois

Ahmad Mahmod
PhD Student

# Software Defined Low-power Lossy Wireless Networks

ICUBE

## Context

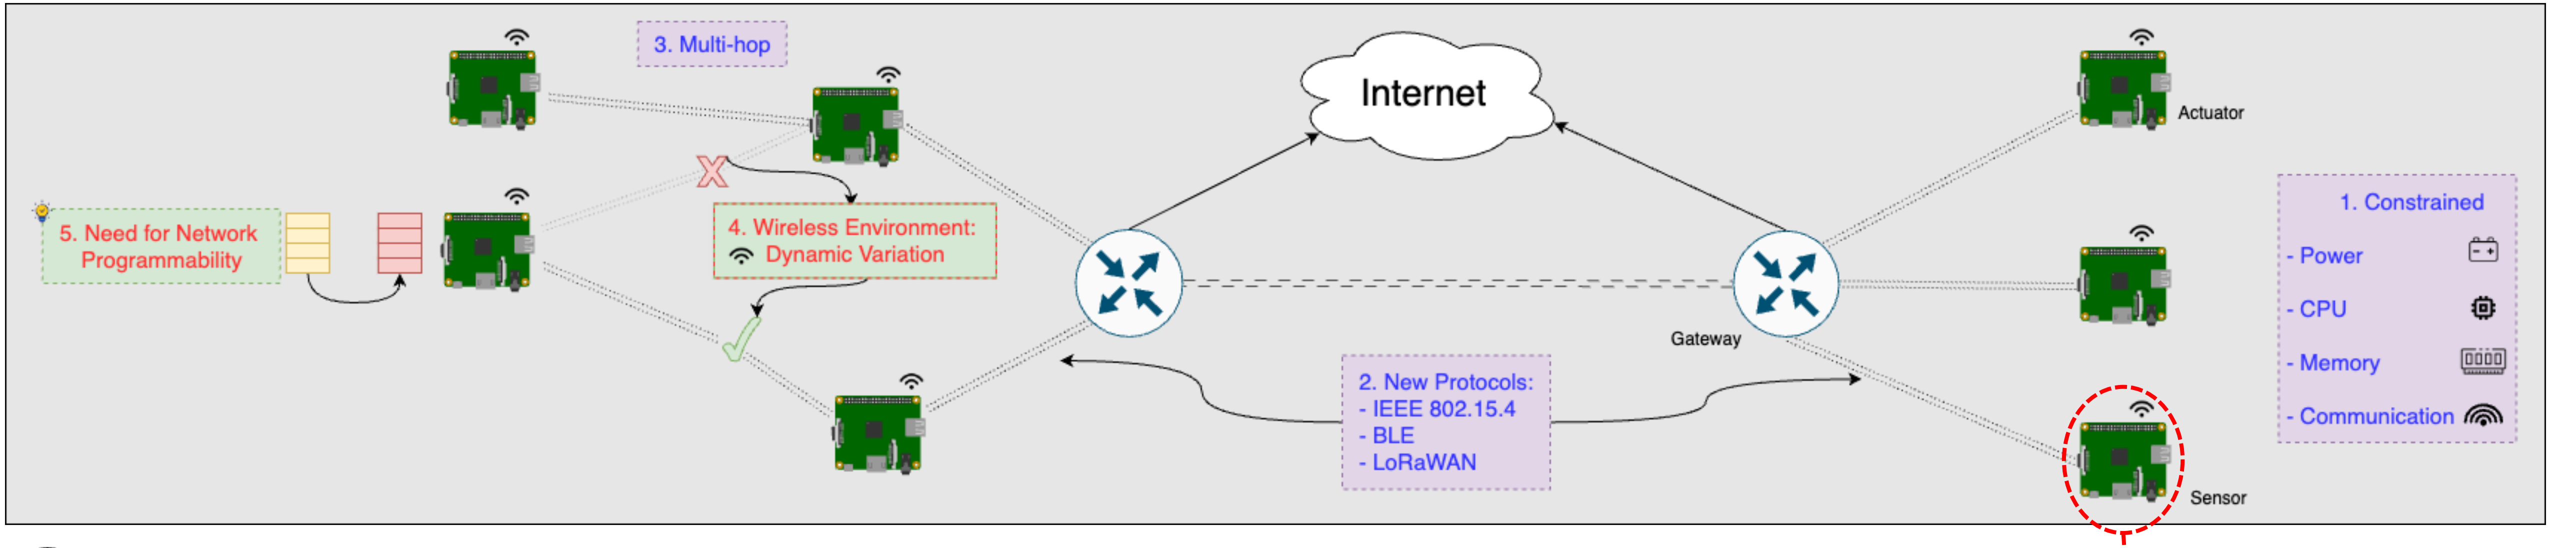


## Solution

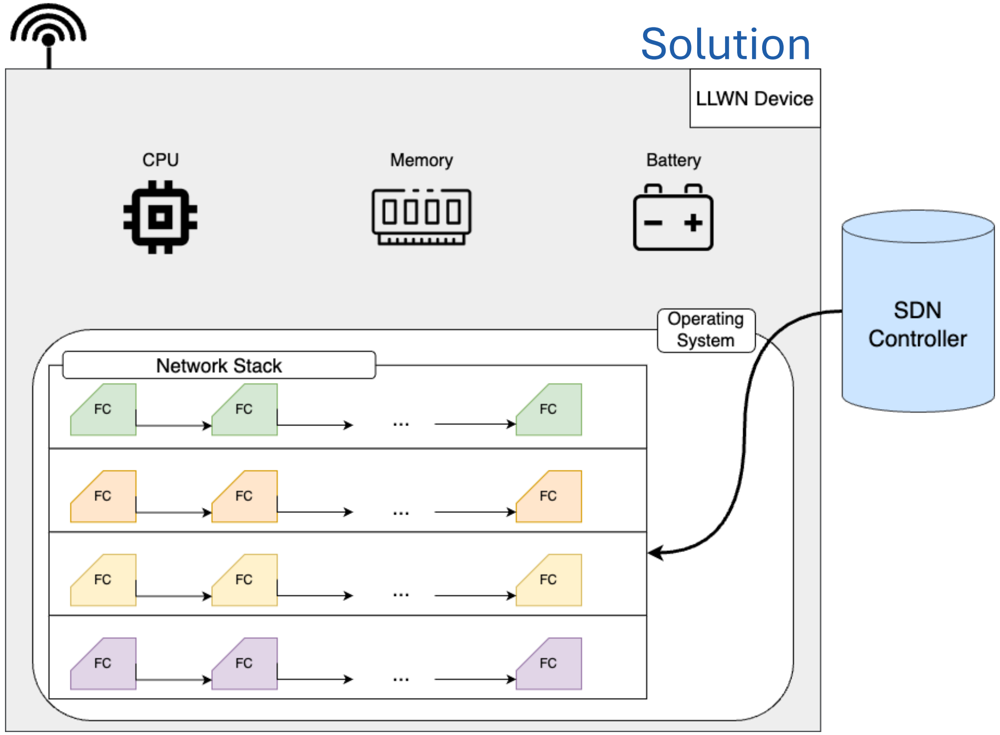


FC

**Femto Container: A lightweight VM for constrained devices**

Programmable LLWN Device

- Reduce power consumption
- Increase Reusability
- Reduce e-waste

## Problem

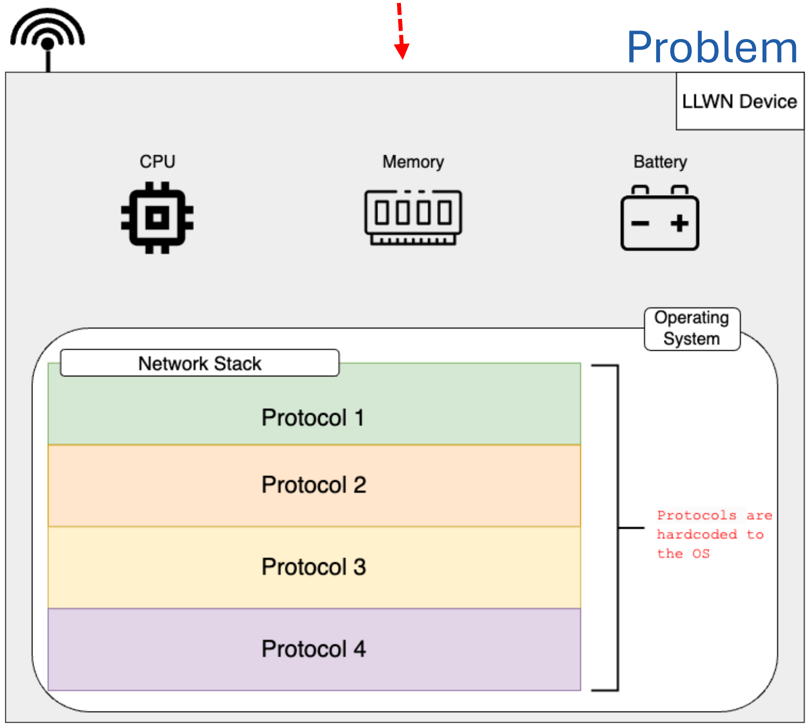

# Heterogeneity-aware resource management in the edge-cloud continuum

**Romain Carlier** (1st year Ph.D. student) and Prof. Etienne Rivière, UCLouvain, Belgium

## CONTEXT AND MOTIVATION

- The edge-cloud continuum is of **increasing popularity** for multiple use cases, and increasingly **heterogeneous**
- Current resource allocation approaches are **single-application-centric**, raising some concerns about the **environmental footprint**

## OBJECTIVES

- Focus on a multi-application view of infrastructure resources to:
  - Maximize the lifespan of existing appliances, raising **endurance**
  - Assign only necessary and sufficient resources, targeting **thriftiness**

## CHALLENGES

1. Taming ISA-heterogeneity
2. Characterizing application profiles
3. Designing metrics and scheduling policies

## APPROACH

**1 Heterogeneity**

Consider **WebAssembly** as a universal runtime for the continuum

WA

**2 Profiles**

scikit learn

Collect **energy/performance** metrics and leverage **similarities** between workloads

PowerAPI

**3 Scheduling**

Use the uncovered profiles to **integrate my objectives** in an edge-cloud **orchestrator**

UCLouvain icteam

# What about **Pablo Leboulanger?**

## BEFORE

- Master : **Fundamental Computer Science**

- Internship : **Modeling Consumption Mobile Network**

  G.Gunnebaud & A.Bugeau

- Voluntary Service : **Development Green Mobility**

https://theshiftproject.org/mondes-virtuels-reseaux/

## NOW

*Minimalist cloud sober in energy and software and material ressources*

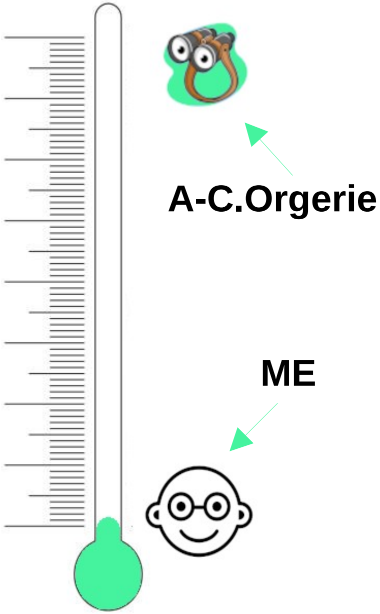


https://carecloud.irisa.fr/project/

pablo.leboulanger@irisa.fr

# Joint Training Design and Network Resource Allocation for Distributed Machine Learning

**PhD Student: Tiago da Silva Barros (Université Côte d’Azur)**

**very large model**

**Very good accuracy**

**medium size model**

**Good accuracy**

**much faster**

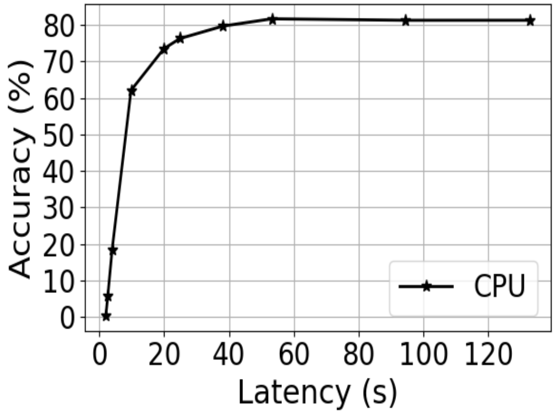


**Reducing model size can help to achieve energy sobriety**

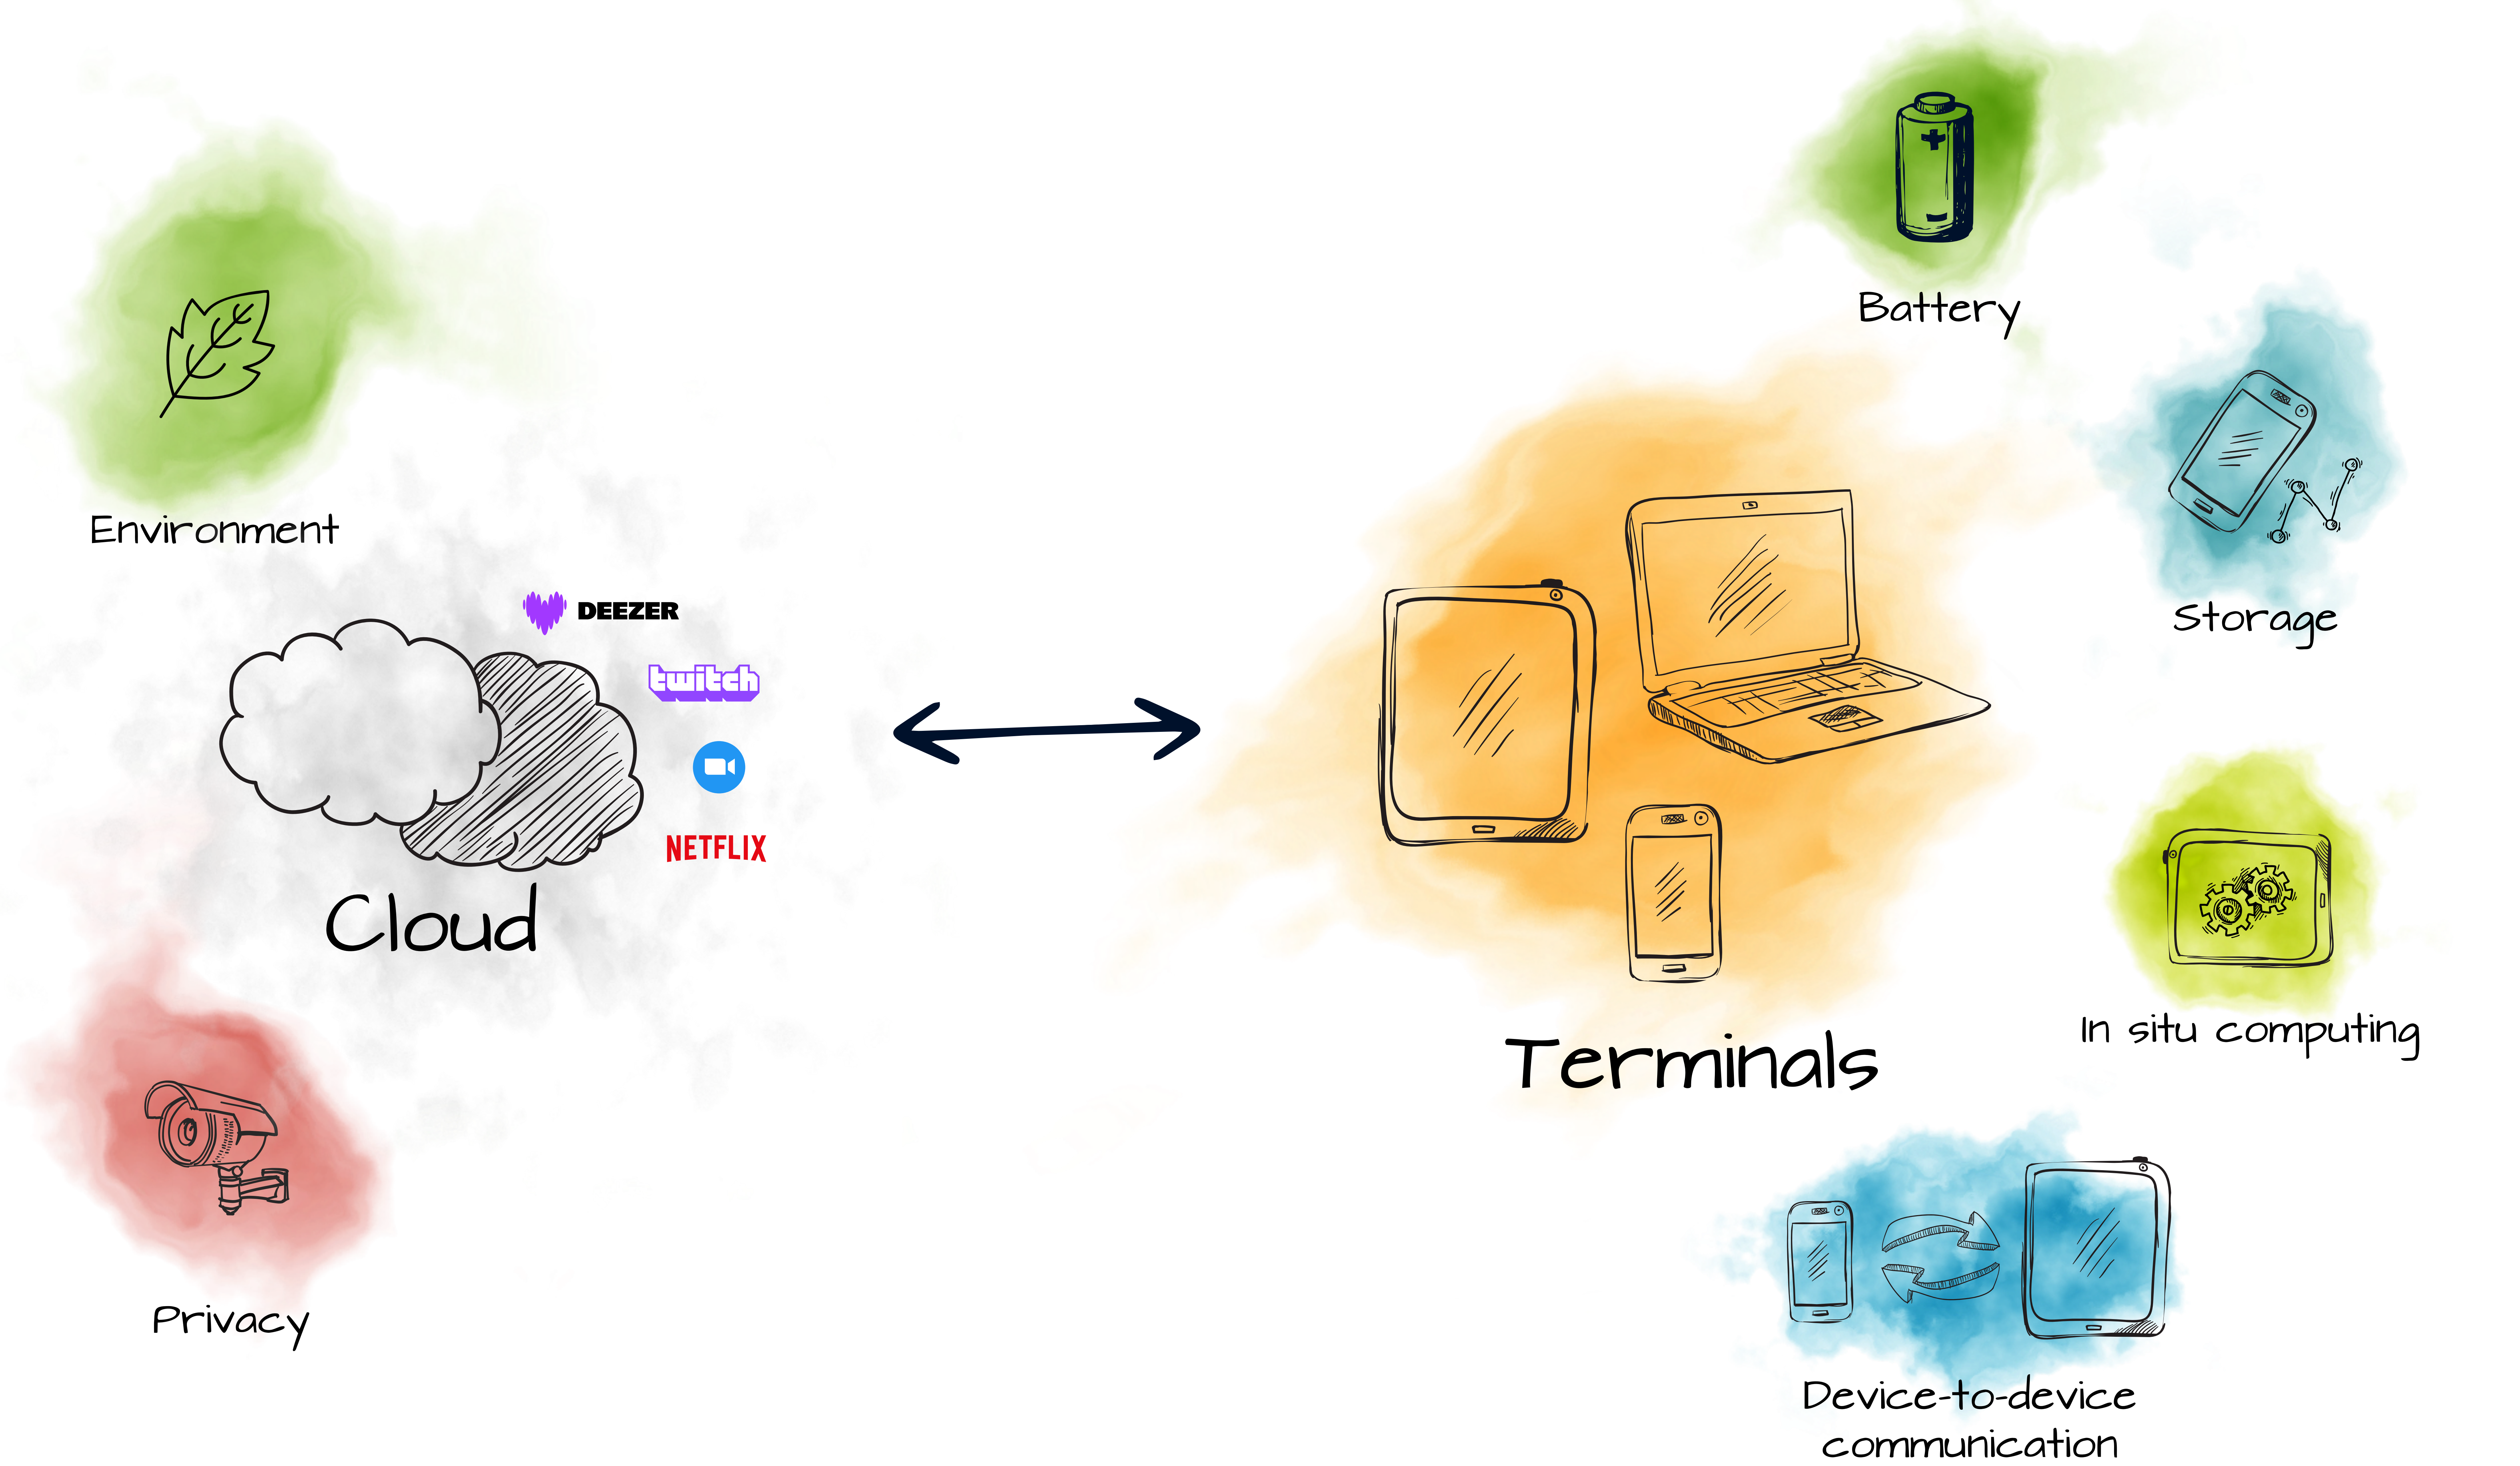
Environment
DEEZER
twitch
NETFLIX
Cloud
Privacy
Battery
Storage
Terminals
In situ computing
Device-to-device communication

# Problem

## 5G in France:

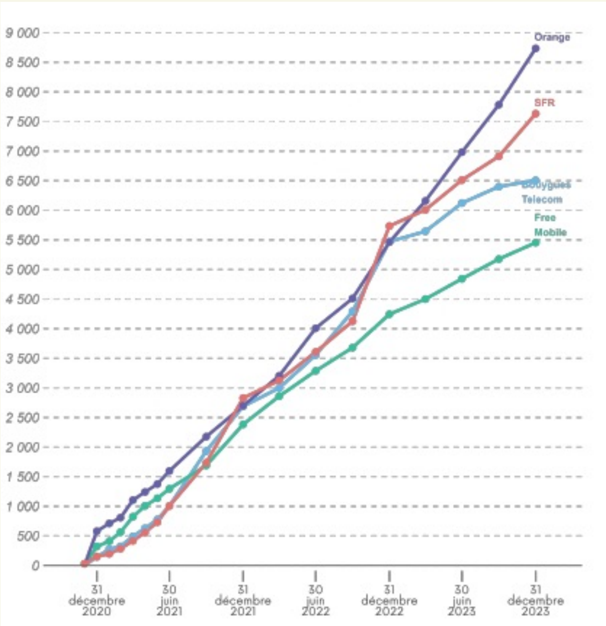


Deployment of 3.5 Ghz 5G sites - ARCEP

## 5G and Sustainability?

# My PhD

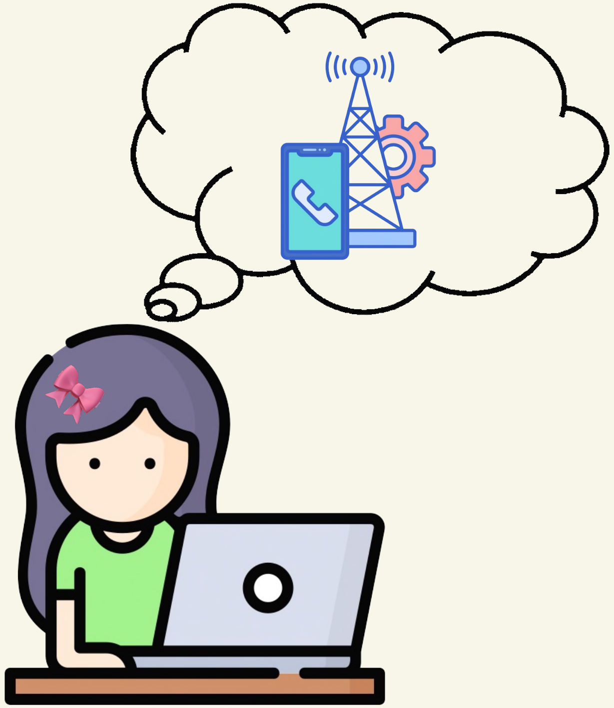

5G is complex
(Massive MIMO,
Beamfoarming ...)

# Goal


* 5G Capacity
* Operational Power Consumption

M. Ghali, HowNet, LIP, ENS Lyon

# MAI'S PHD

**Prof. Georges Da Costa**
*SEPIA team - IRIT*
*University Toulouse III*

**Thesis**: Exploring the tradeoffs between **Energy** and performance of **Federated Learning** algorithms
=> Build a reproducible framework

**Assoc. Prof. Millian Poquet**
*SEPIA team - IRIT*
*University Toulouse III*

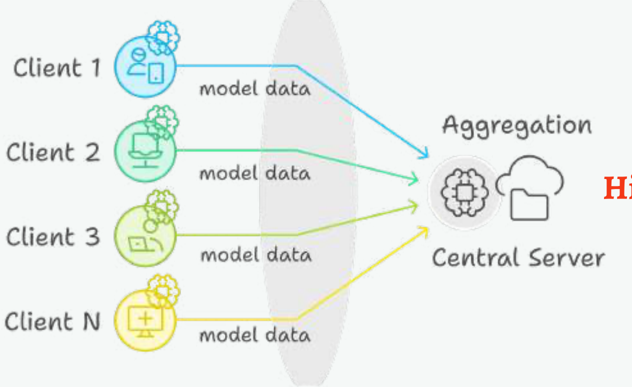


**High Energy Consumption?**

- Computation (CPU, GPU, ...)
- Communication (Data transfer,...)
- Etc (Cooling,...)

**Mai-Huong Do**
***PhD Student***
*SEPIA team - IRIT*
*University Toulouse III*

HuongDM1896.github.io

Huong.Do-Mai@irit.fr

# Electric Vehicles to balance the Grid

Matthieu SILARD
Anne-Cécile ORGERIE
Nicolas MONTAVONT
Georgios Z. PAPADOPOULOS

Smooths consumption by controlling HVAC, water heaters, or EV charging

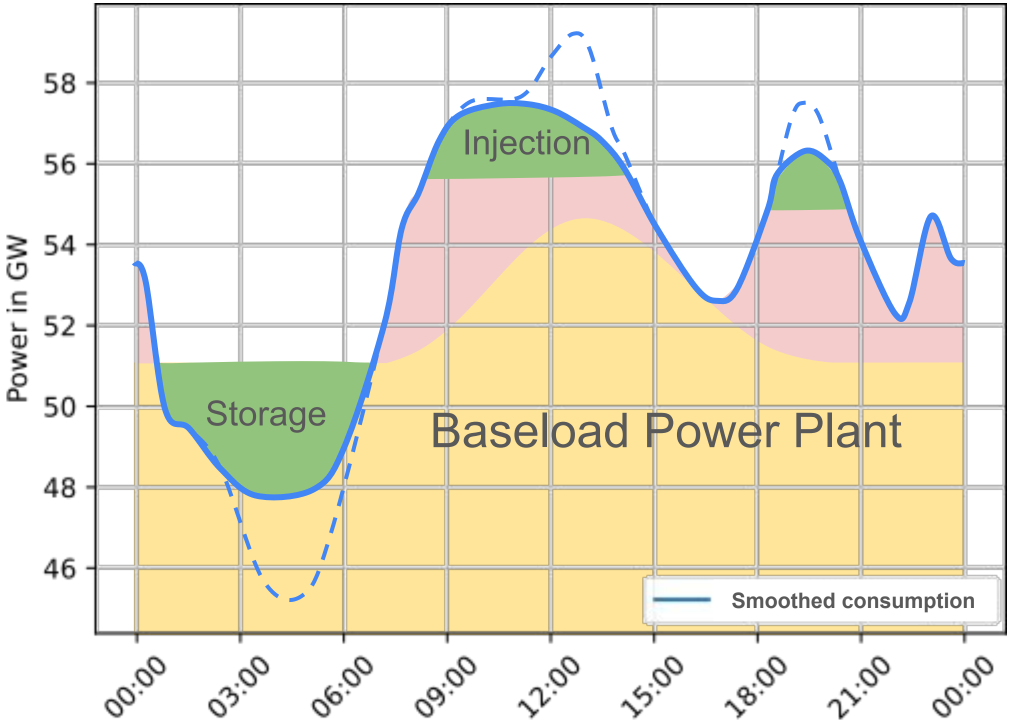


Peaking Power Plant

- Oil
- Gas
- Dams

} costly or limited

Baseload Power Plant

- Coal
- Nuclear

} cheap but no flexibility

+ solar/wind power

Storage/Injection

- Pumped-storage hydroelectricity
- EVs with V2X capabilities

**38GWh potential storage!**

matthieu.silard@irisa.fr

# Speed or energy-efficiency ? That is the question.

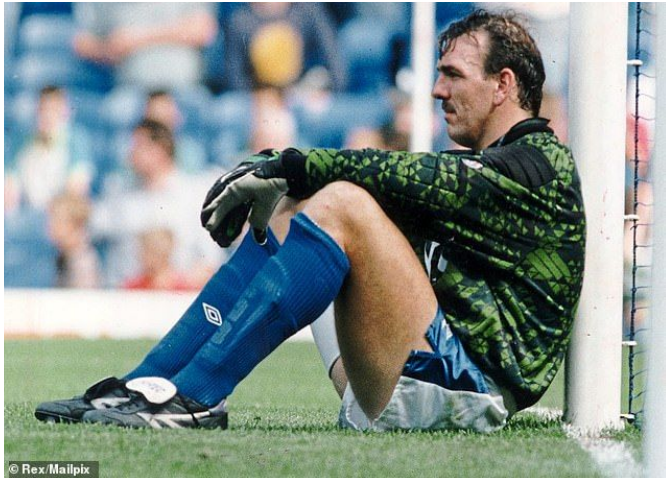


**Energy-consuming**
**Fast response**

**Energy-saving**
**Slow response**

UCLouvain fnrs LA LIBERTÉ DE CHERCHER

Problem

## CLOUD

*APPLICATIONS*

**INCREASE OF ENERGY CONSUMPTION IN CLOUD APPLICATIONS, AND NO ADAPTATION STRATEGY CONCERNING ENERGY EFFICIENCY**

Motivation

# Self-Adaptive Systems

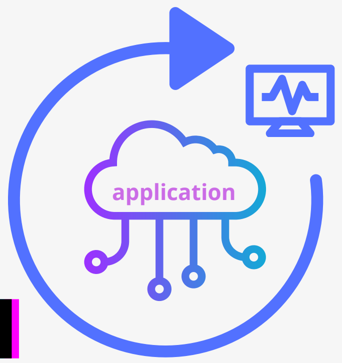


Motivation

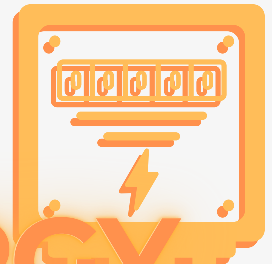

## ENERGY CONSUMPTION

**ENERGY PROFILERS**

*APPLICATION LEVEL*

PowerAPI

Motivation

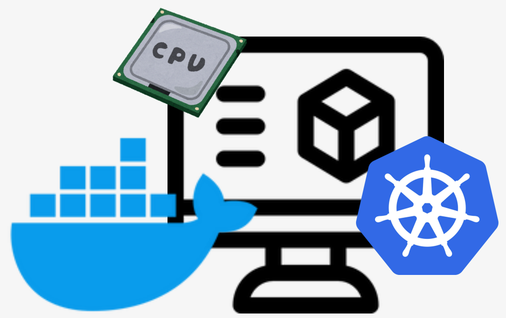


## RESOURCES PROVISIONING

Literature

**REVIEW**

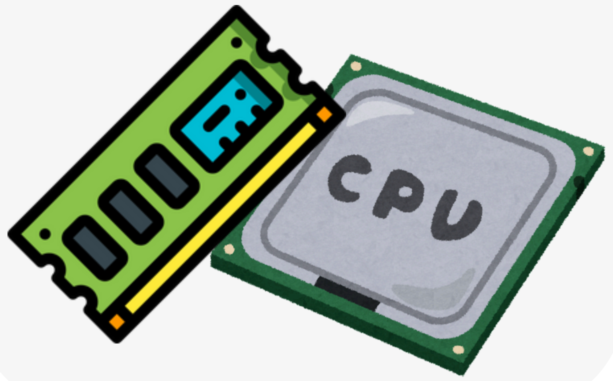


Proposal

# *Self-Adaptive* ENERGY EFFICIENT Architecture

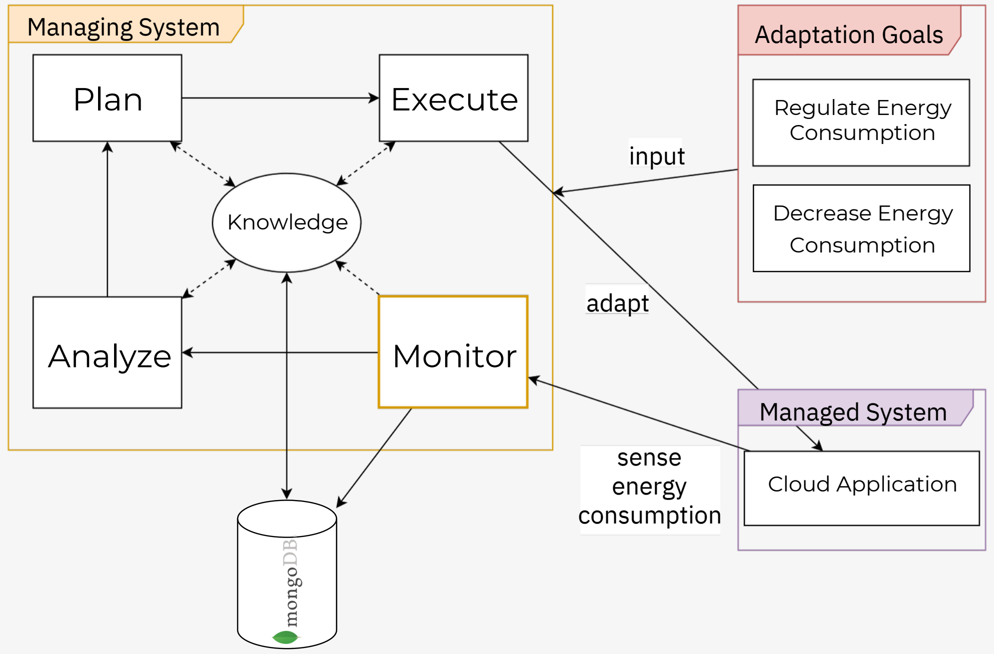



## Henrique DE MEDEIROS

*PhD Student - Started on September 2024*

*Institut Polytechnique de Paris, Télécom SudParis*


TELECOM SudParis · IP PARIS · ensiie Paris | Évry-Courcouronnes · CARECloud · s@movar · E4C INTERDISCIPLINARY CENTER · FRANCE 2030 PROGRAMME DE RECHERCHE CLOUD

# DATA

# Towards sustainable and resilient AI

## Frugal and energy-aware distributed machine learning platforms

[1] Lo, Sin Kit, et al. "Architectural patterns for the design of federated learning systems." Journal of Systems and Software 191 (2022): 111357.

## 1/ Problematic: reduce energy consumption but many configurations

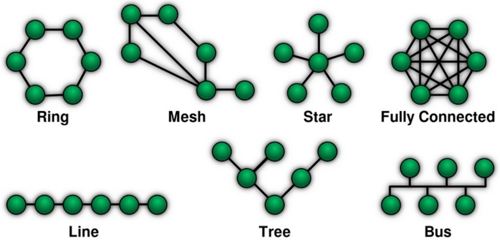


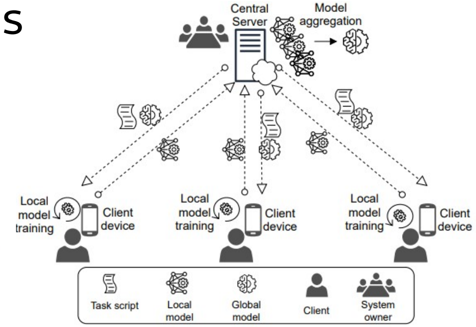


Network topology

Task allocation[1]

## 2/ Methodology:

Discrete simulation:

\+ Fast, flexible, reproductibility

— Calibration required

Experimenting:

\+ Realist scenario, will help calibration

— Harder config and deployment

## 3/ What's done: **simulator**, analysis tool, config optimizer based on **evolution algorithm**

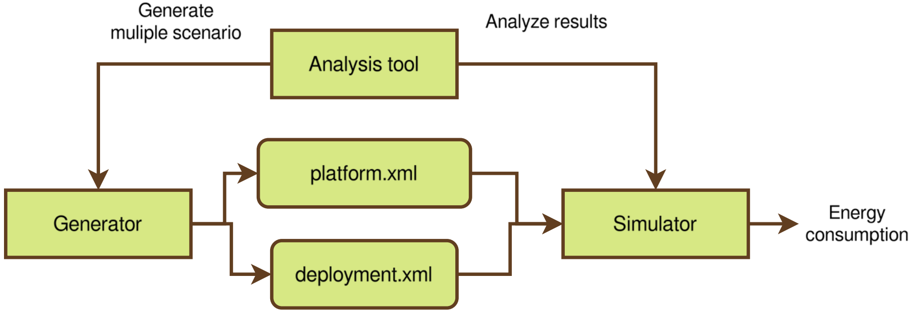


## 4/ Early results:

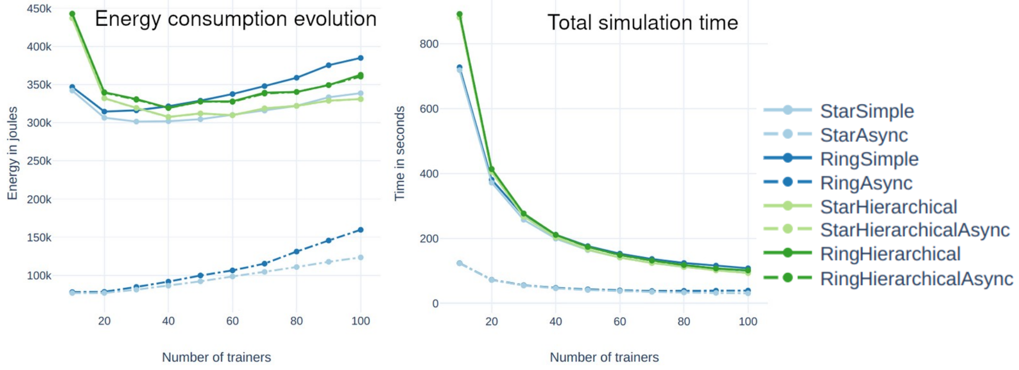

# (Provisional) Title: "Load shifting and multi-platform services for energy efficiency"

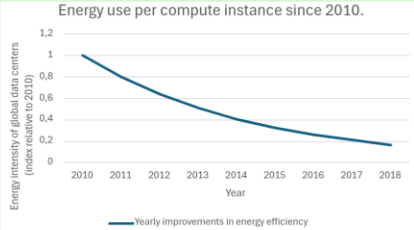


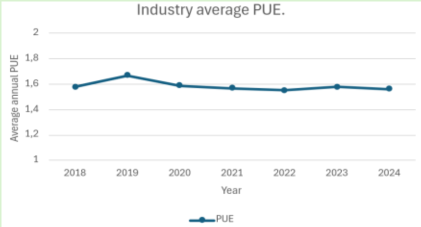


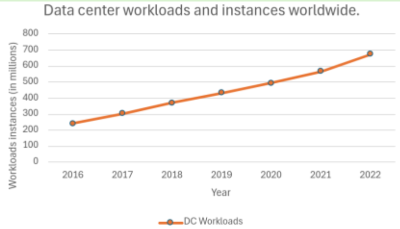


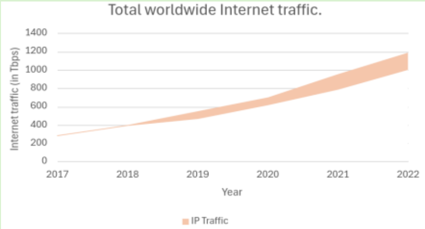


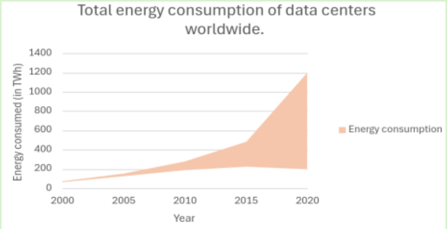


Despite improvement in efficiency, **+** increase in workloads and users **=** leads to more energy consumed !

Focusing on efficiency by workload and instance is not enough to reduce energy consumption and carbon emissions. To reduce the ICT sector's carbon emissions, we propose combining existing shifting techniques to change the time, machine and processes to compute a workload, and aim the least carbon-intensive energy consumed.

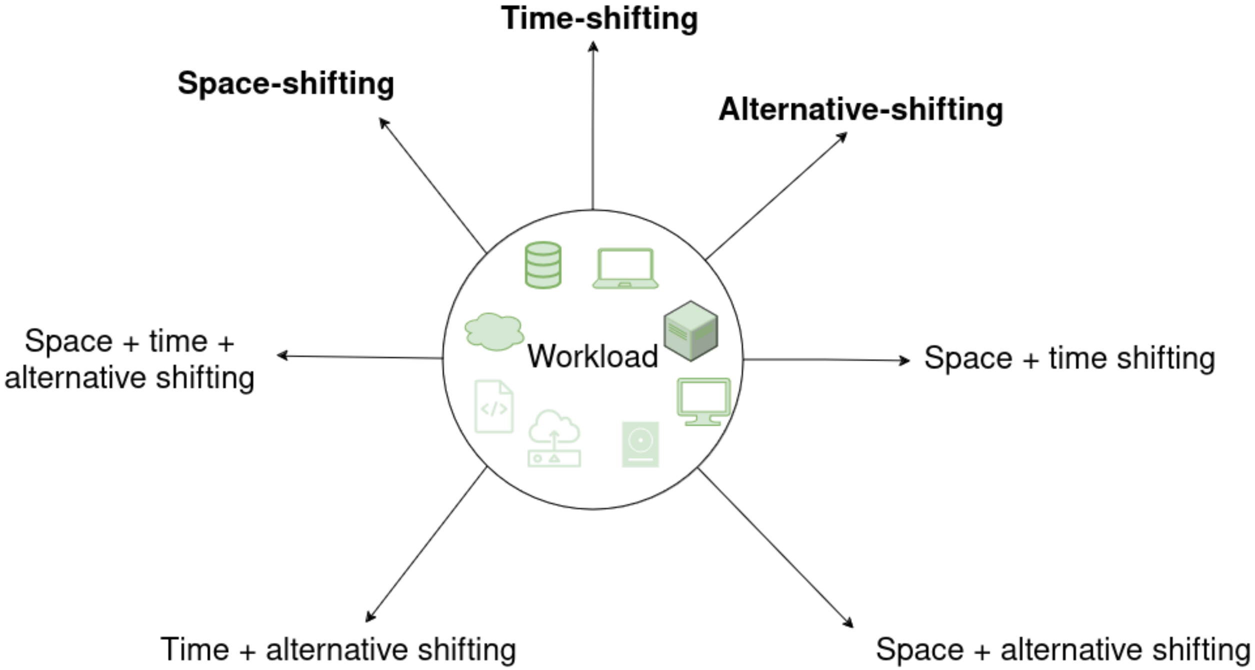


Université de Pau et des Pays de l'Adour

Universidad Zaragoza


Author: **Nicolas Tirel**, directors: **Philippe Roose** (LIUPPA) and **Sergio Ilarri** (I3A, UNIZAR)
Supervised also by **Adel Noureddine** and **Olivier Le Goaër** (LIUPPA)

# Matteo Chancerel

## Magellan team, Rennes

M1

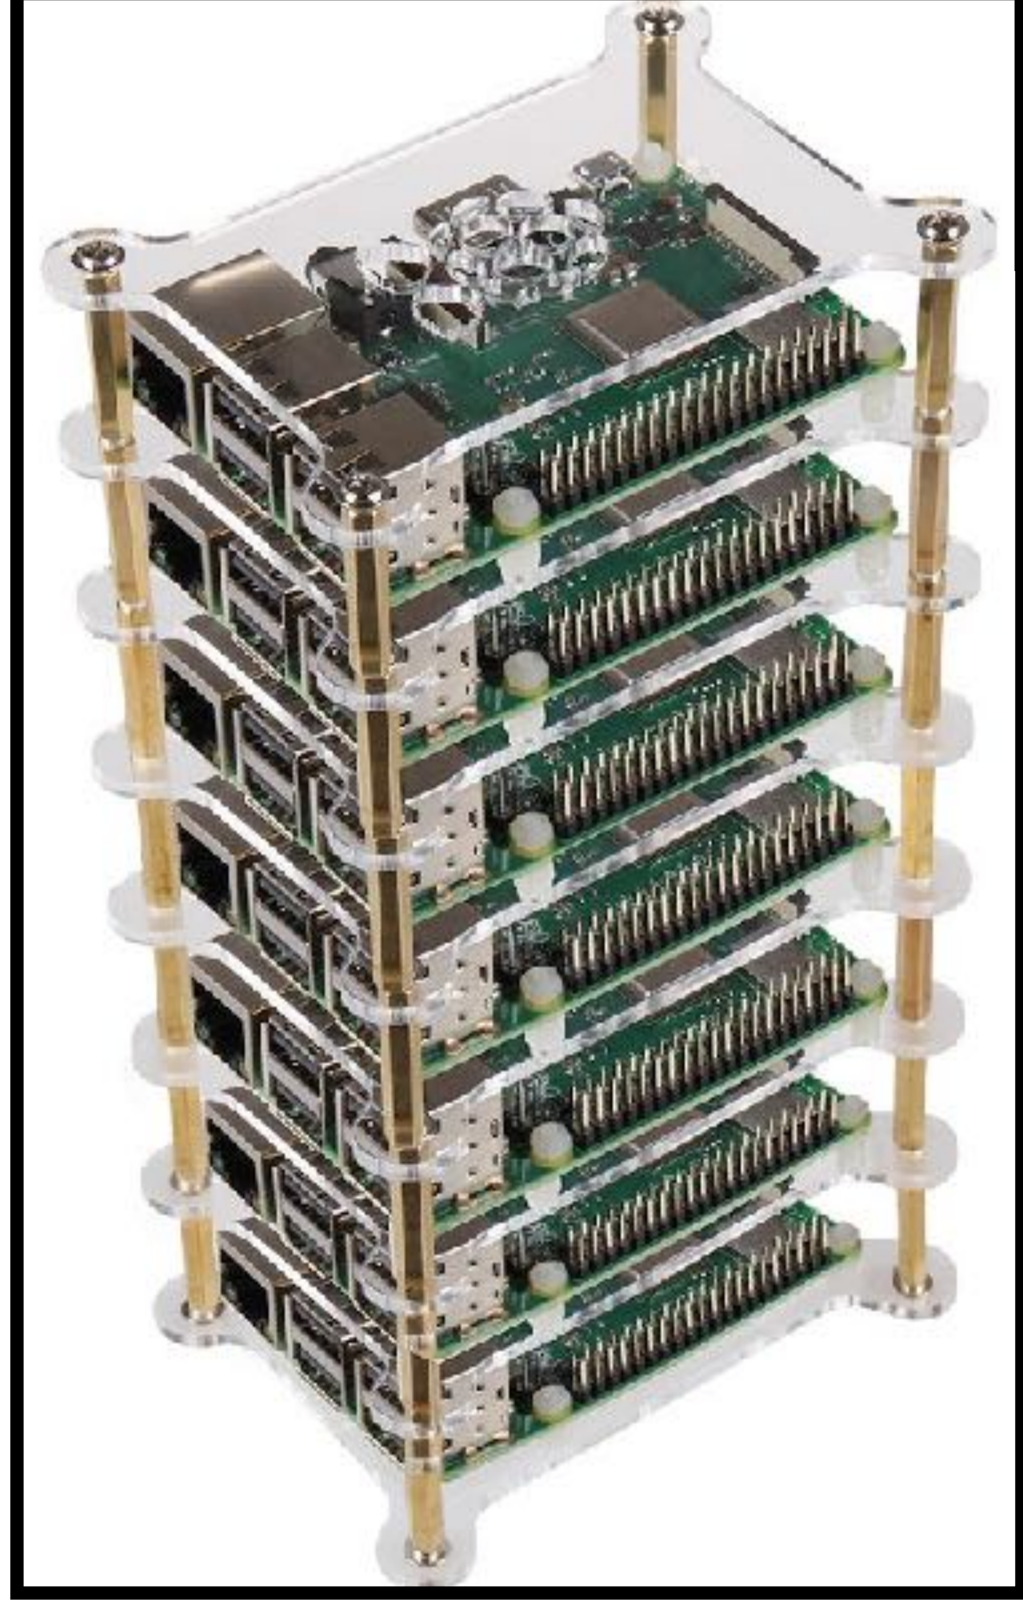

L3

$X = \overline{Orb(x)}$ is the Toeplitz subshift generated by $x$.

| | | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| ... | .a | • | a | • | a | • | a | • | a | • | a | • | a | • | a | • | a | • | ... |
| ... | .a | b | a | • | a | b | a | • | a | b | a | • | a | b | a | • | a | b | ... |
| ... | .a | b | a | a | a | b | a | • | a | b | a | a | a | b | a | • | a | b | ... |
| ... | .a | b | a | a | a | b | a | b | a | b | a | a | a | b | a | • | a | b | ... |
| ... | .a | b | a | a | a | b | a | b | a | b | a | a | a | b | a | a | a | b | ... |

M2

$$\pi = 3 + \cfrac{1}{7 + \cfrac{1}{15 + \cfrac{1}{1 + \cfrac{1}{292 + \cfrac{1}{1 + \cfrac{1}{1 + \cfrac{1}{1 + \cfrac{1}{2 + \cfrac{1}{1 + \cfrac{1}{3 + \cfrac{1}{1 + \cfrac{1}{14 + \cfrac{1}{2 + \cfrac{1}{\dots}}}}}}}}}}}}}}$$

PHD : Optimization of a fully distributed fog powered by renewable energy sources.

# Energy-efficient Microservice-based Software Architectures: Energy Consumption and Dynamic Evolution in Cloud Environments

## MOTIVATION

- Microservices elasticity mitigates resource overprovisioning [Fontana de Nardin et al., 2021]

## PROBLEM

- Tracking energy consumption of a single microservice is not sufficient [Anand et al., 2023]

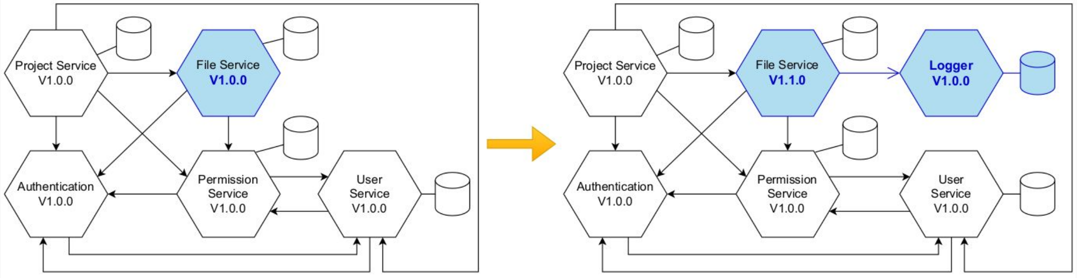


Microservice application evolution [Wang 2022]

**RQ1**: How to take into account energy efficiency when analysing microservice-based architectures?

- Measure of energy consumption <u>at the granularity of a microservice</u> for a given set of workloads
- Provide energy consumption data for <u>planning a future reconfiguration</u> or different versions of microservices

**RQ2**: How to estimate energy consumption of individual end-to-end requests?

- Determine <u>Energy-aware Critical Paths</u> of requests using Distributed Tracing

**RQ3**: How to enhance DT for energy consumption in asynchronous workloads?

- Identify <u>Energy Bottlenecks</u>

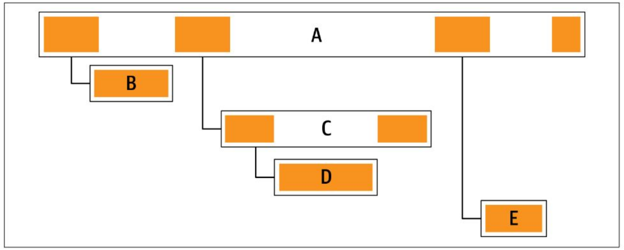


Critical path of a trace [Parker et al., 2020]

**César PERDIGÃO BATISTA**
PhD Student from 01/2024
at SAMOVAR/TSP/IPPARIS

TELECOM SudParis
IP PARIS
CARECloud
FRANCE 2030 PROGRAMME DE RECHERCHE CLOUD

Guillaume RAFFIN - alumet.dev

⚡? Software energy consumption

📈? Performance

EVIDEN LIG

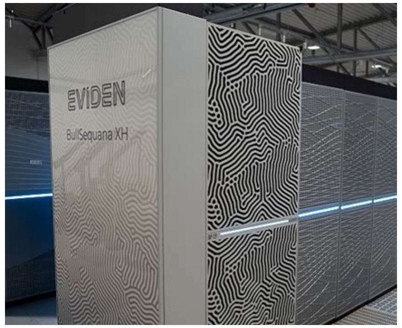


slurm workload manager

# Alumet

**Modular measurement framework and tool**

- High performance
  High frequency
- Plugin system
  => bespoke tools
- Metrics standard
  + flexible attributes
- Written in Rust
  (async inside!)