ORGANISED BY

BDV BIG DATA VALUE ASSOCIATION

HOSTED BY

COSMOTE

DATAWEEK

JOIN.LEARN.SHARE.GET VALUE

Odyssey of AI:

Navigating the Data Seas

27 - 28 May / Athens

IN COLLABORATION WITH

ATC ATHENS TECHNOLOGY CENTER

ΕΠΙΣΕΥ ICCS

core|IC

DEMOKRITOS NATIONAL CENTRE FOR SCIENTIFIC RESEARCH

FORTH INSTITUTE OF CHEMICAL ENGINEERING SCIENCES

NEUROPUBLIC Information Systems & Technologies

data-week.eu
#DataWeek25

DATAWEEK
JOIN.LEARN.SHARE.GET VALUE

ML-Based Autonomic System Management in the Edge-Cloud Continuum

Christos Antonopoulos
Computer Systems Lab - ECE Department
University of Thessaly, Greece

ORGANISED BY BDV BIG DATA VALUE ASSOCIATION
HOSTED BY COSMOTE
IN COLLABORATION WITH ATC CORE|IC FORTH DEMOKRITOS NEUROPUBLIC

CHRISTOS ANTONOPOULOS

Professor in the Department of Electrical and Computer Engineering at the University of Thessaly and Director of its Computer Systems Lab.

Research focus on system software and software optimization for cloud, high-performance, and embedded systems, approximate computing, and edge computing.

Authored 80+ publications (4 best paper awards).

Coordinated and contributed to major national and EU projects and has served as a scientific advisor to CNH Industrial on AI and computational challenges related to smart/precision agriculture.

MLSysOps Overview

HORIZON-CL4-2022-DATA-01-02 - Cognitive Cloud: AI-enabled computing continuum from Cloud to Edge (RIA)

Grant ID: 101092912

12 partners

8 countries

1/1/2023-31/1/2026

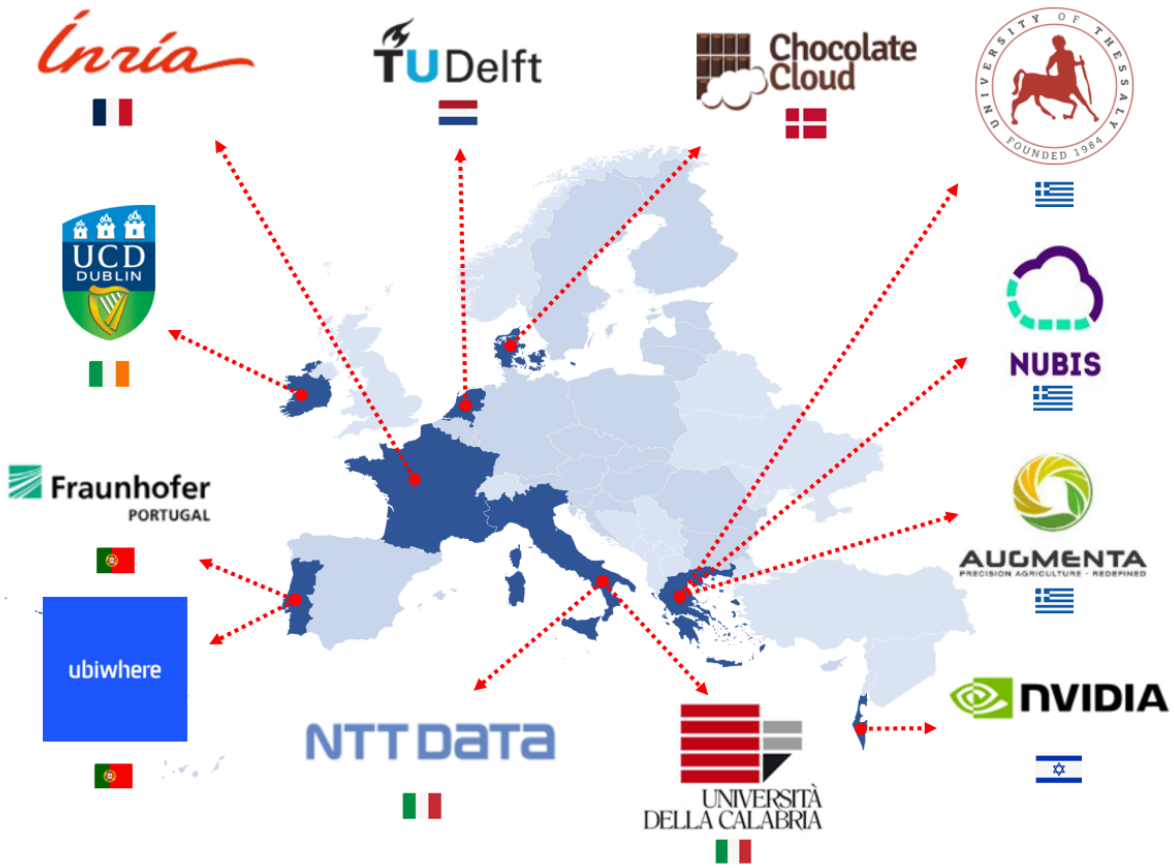

THE TREND

Continuum systems

- Applications are moving outside the cloud
  - Start involving nodes and resources toward or directly at the edge of the Internet…
  - … including powerful but also resource-constrained IoT devices

THE CHALLENGES

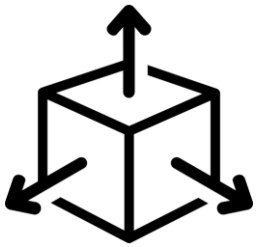

Scale

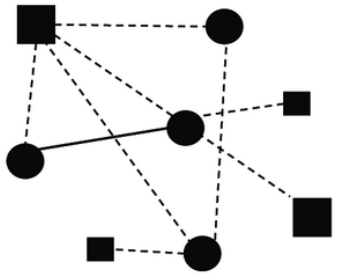

Heterogeneity

Security

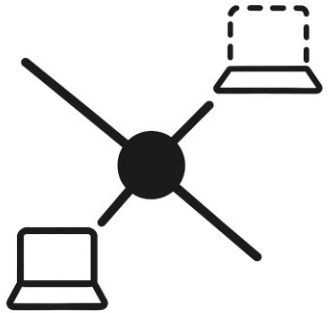

Volatility

Impossible to monitor/manage by a human

DATAWEEK

JOIN.LEARN.SHARE.GET VALUE

ARE AI/ML METHODS FOR AUTONOMIC SYSTEM MANAGEMENT AND CONFIGURATION IN THE CLOUD-EDGE-IOT CONTINUUM A FEASIBLE SOLUTION?

MLSysOps Concept

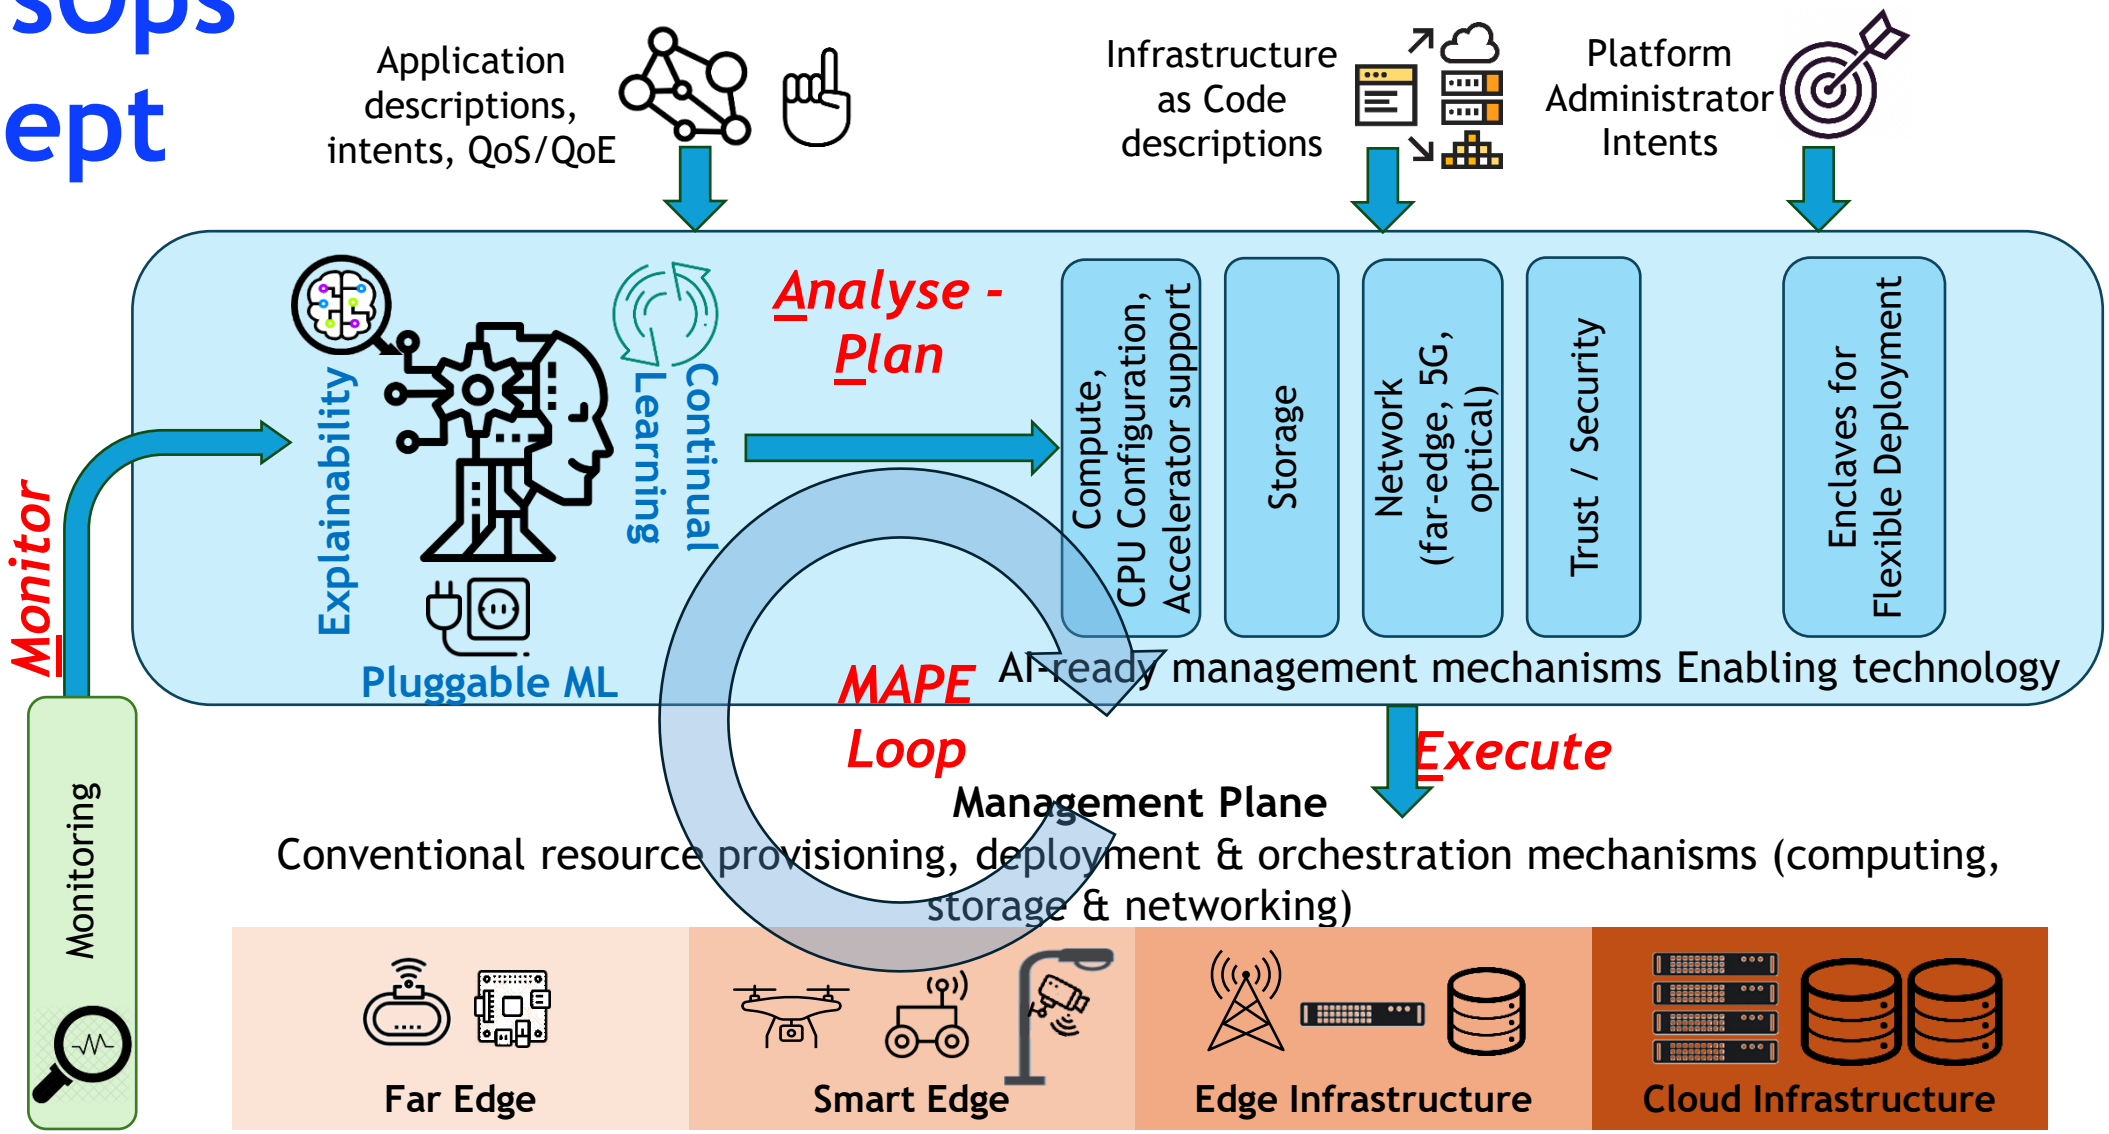

MLSysOps Instantiation

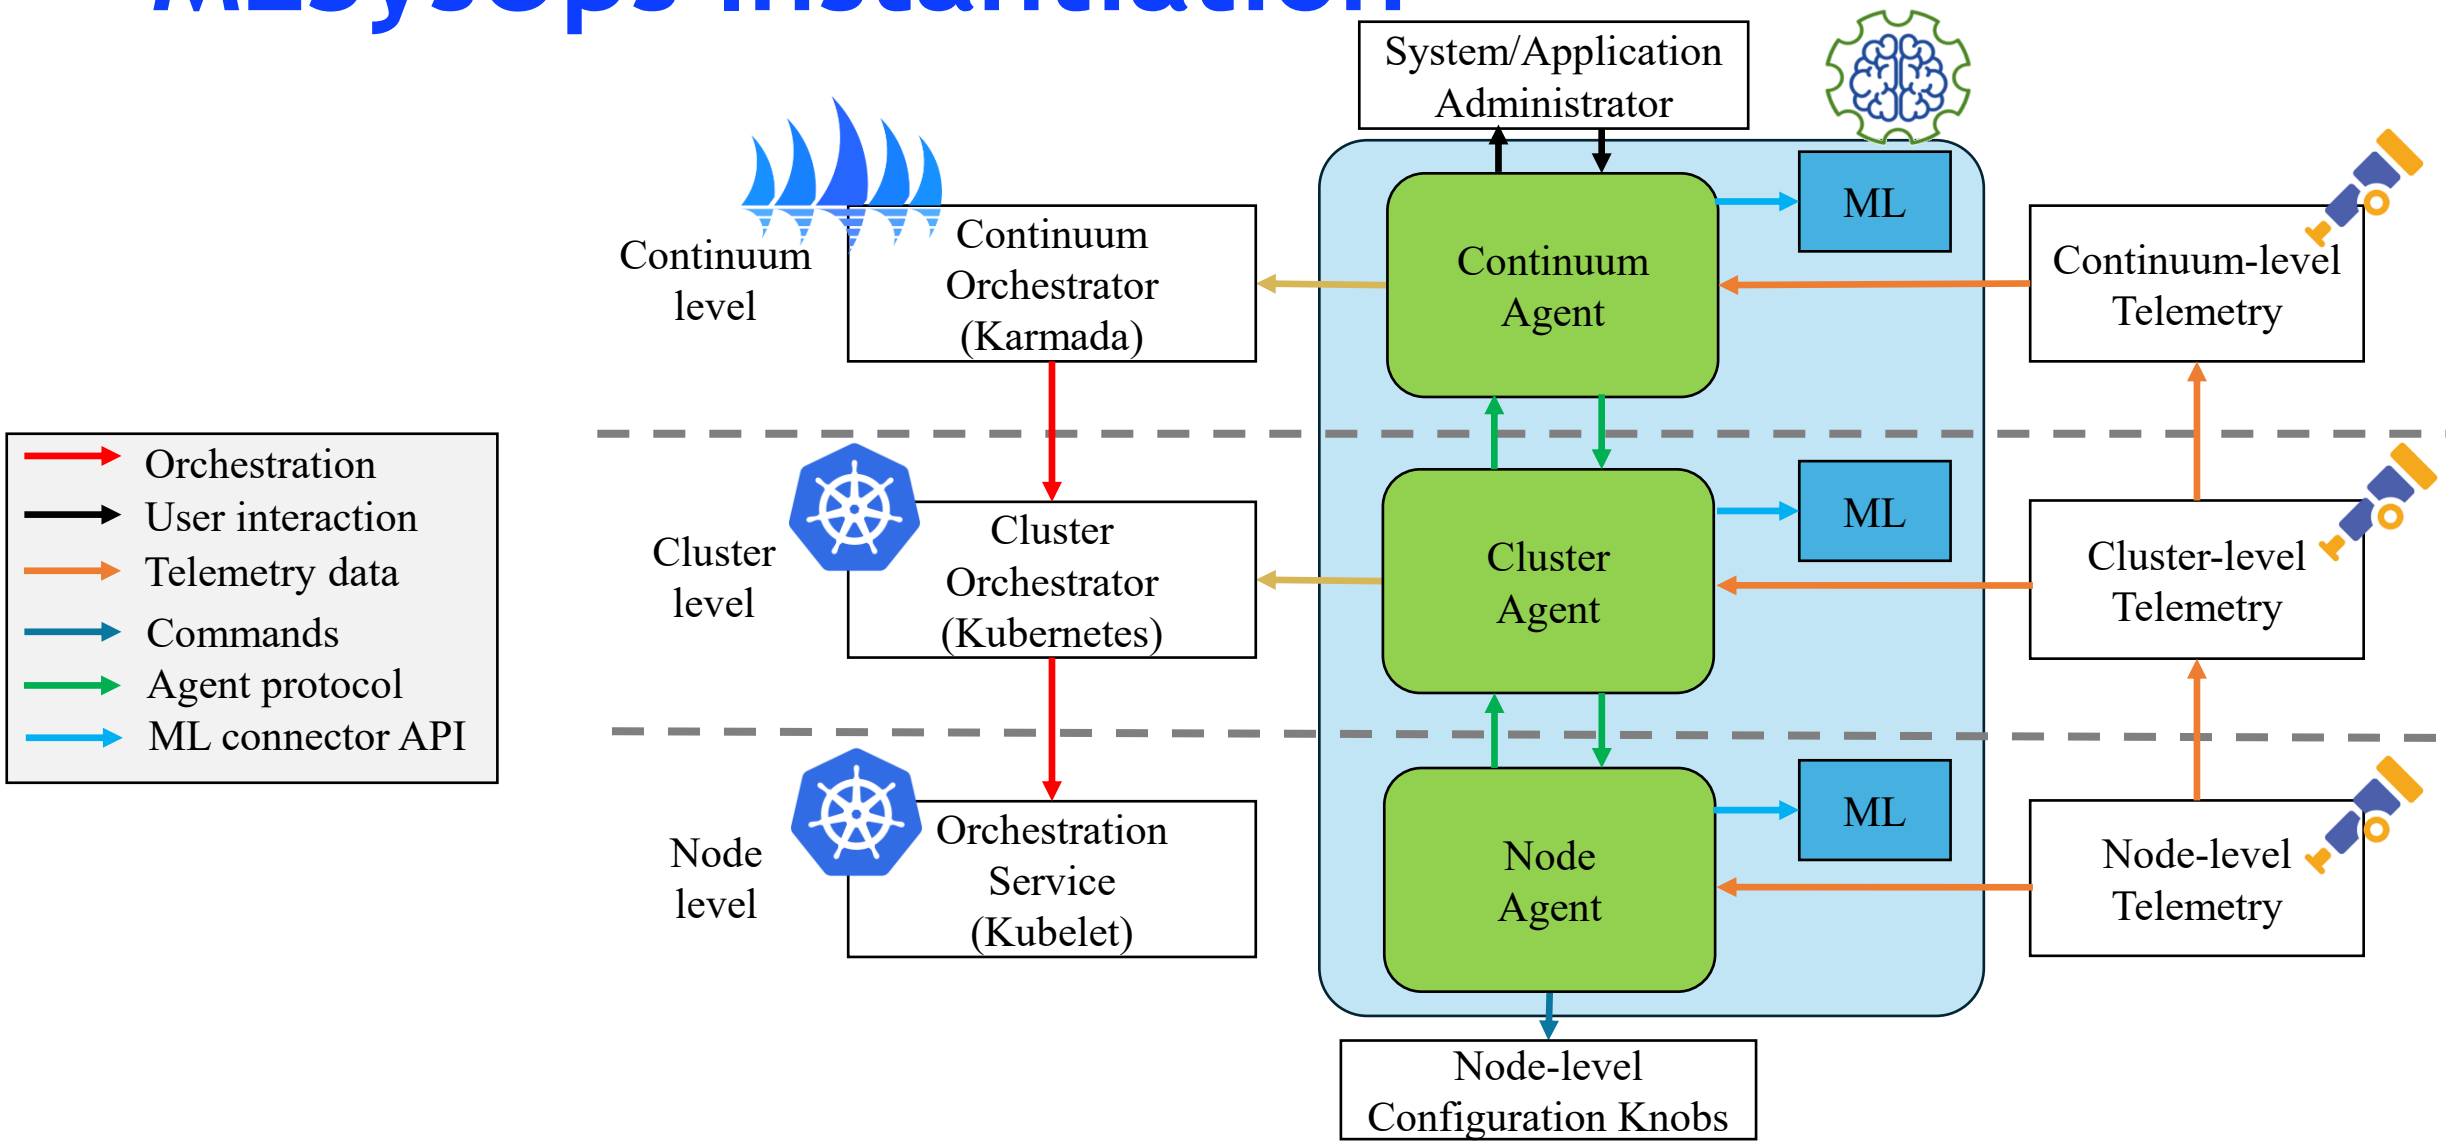

ML Models

- ML models may become outdated; adapt to a changing world/conditions
  - Reinforcement learning (RL)
- Do not train from scratch for each new setup; reuse and adapt pre-trained models
  - Transfer learning (TL)
- Privacy-preserving and distributed/scalable training
  - Federated Learning (FL) & Split Federated Learning (SFL)
- ML models (and their training) can be deployed as special applications

Use Cases

Smart Cities

ubiwhere

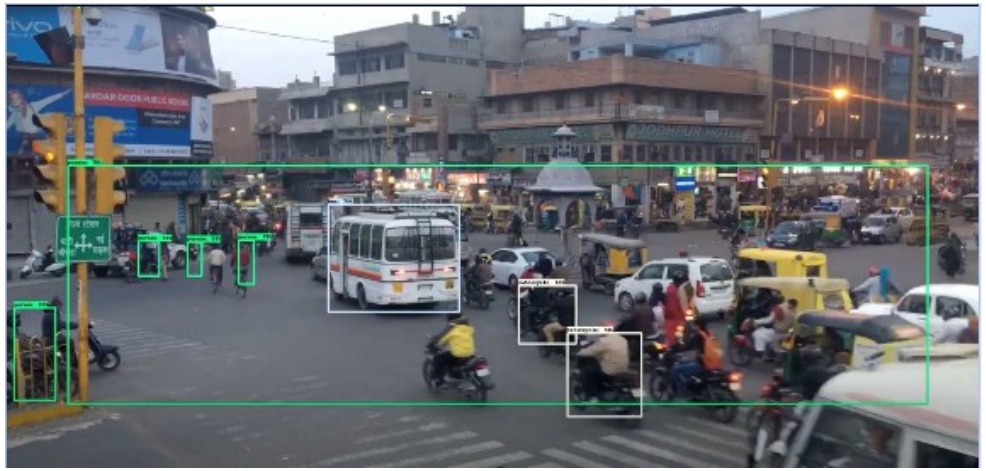

Improve existing system via AI/ML-driven control to manage/configure application modules that can be deployed on smart lampposts and/or to datacenters.

Smart Agriculture

AUGMENTA
PRECISION AGRICULTURE - REDEFINED

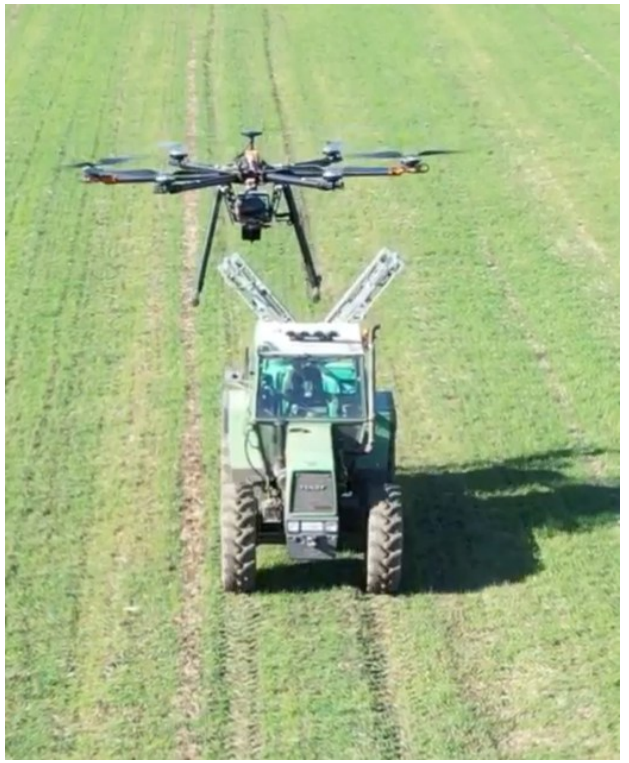

Improve existing system via AI/ML-driven control to jointly manage/configure the image processing application pipeline on a tractor and a companion drone.

ML for System Management

Three success scenarios

Can we do better with ML-based resource management for the Cloud?

Solution Overview

- Node-level ML model
  - Runs locally on each node
  - Predicts future CPU utilization and VM lifetime
  - Drives VM management decisions
- Cluster agent
  - Receives predictor output
  - Applies a heuristic to globally optimize cluster resources

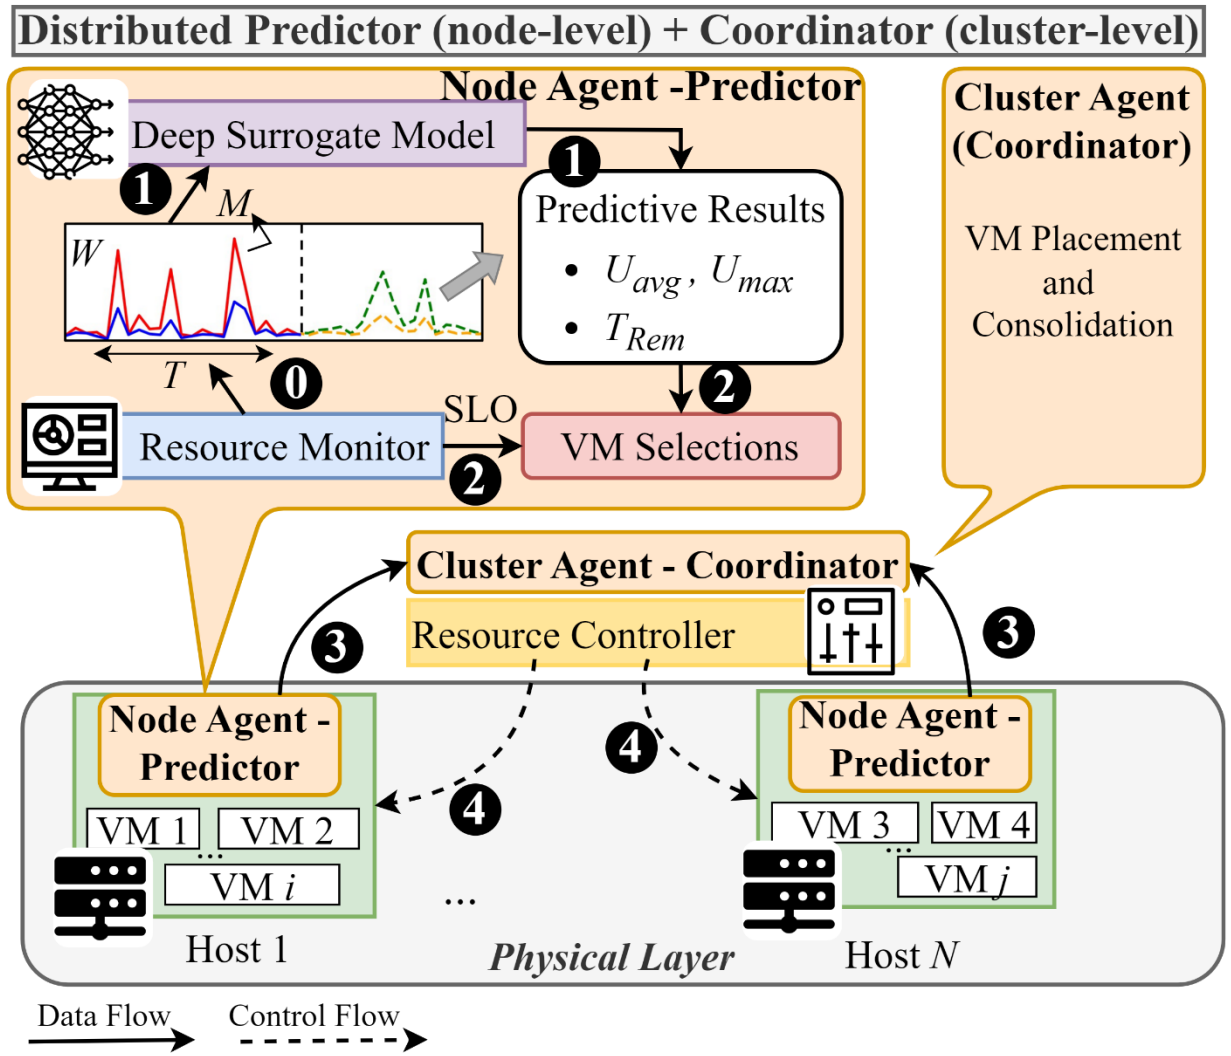

Results

- Migration counts
  - **Decrease of 41.33%** compared with SoTA conventional (non-ML) policy
  - Only 7.5% over oracle (Ground Truth)
- SLO violations
  - **Lower by 34.98%** compared with SoTA conventional policy
  - A mere 1.1% more than oracle (Ground Truth)

But do we have to train from scratch for each target cluster?

Transferring DRL Agents Between Infrastructures

- Trained DRL agents for 4 different cluster sizes (4, 8, 16, 32 nodes)
- Transferred agents between clusters and compared with training from scratch
- Pre-trained agents consistently outperformed agents trained from scratch when comparing performance over the same training duration
  - Performance improved by up to 54%
  - In some cases, pre-trained agents outperformed the converged from-scratch performance in less than 25% of the training time

Transferring Learning Insights

- **Insight 1: Larger** clusters benefit the most from transfer learning
- **Insight 2: Source-target** cluster size **similarity** matters
- **Insight 3: Small-to-large** transfers are more **effective** than the reverse

Can ML improve security at the edge?

Why is device authentication necessary?

📈 IoT explosion with **30+ billion devices** expected in **2025** across smart homes, healthcare, and industrial sectors.

🎯 Unsecured devices become prime targets for sophisticated **cyber attacks**.

- 🔓 Unauthorized access leading to critical **data breaches** and system **compromises**.
- 🕵️ Device spoofing enabling dangerous **Man-in-the-Middle** attacks.
- 🤖 Vulnerable devices being hijacked for large-scale **DDoS botnet attacks**.

⚠️ Increasing complexity of cyber threats targeting **connected devices**.

Solution

- Using multiparametric Physical-Level Authentication (PLA), on top of cryptographic authentication solutions

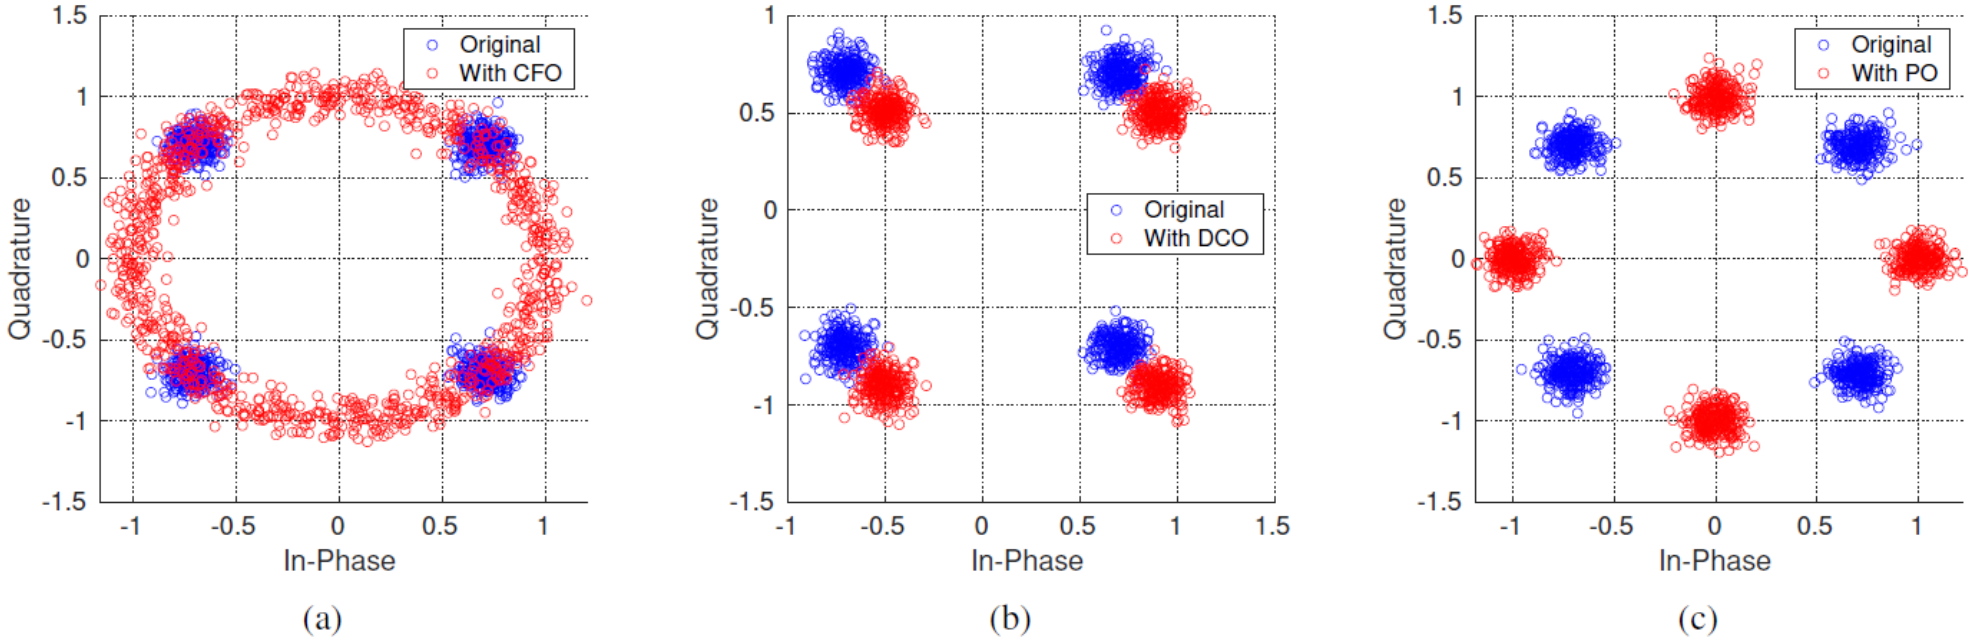

- Based on unique physical characteristics of the communication subsystem.

(a): Carrier Frequency Offset (b): Direct Current (DC) Offset (c): Phase Offset

Setup

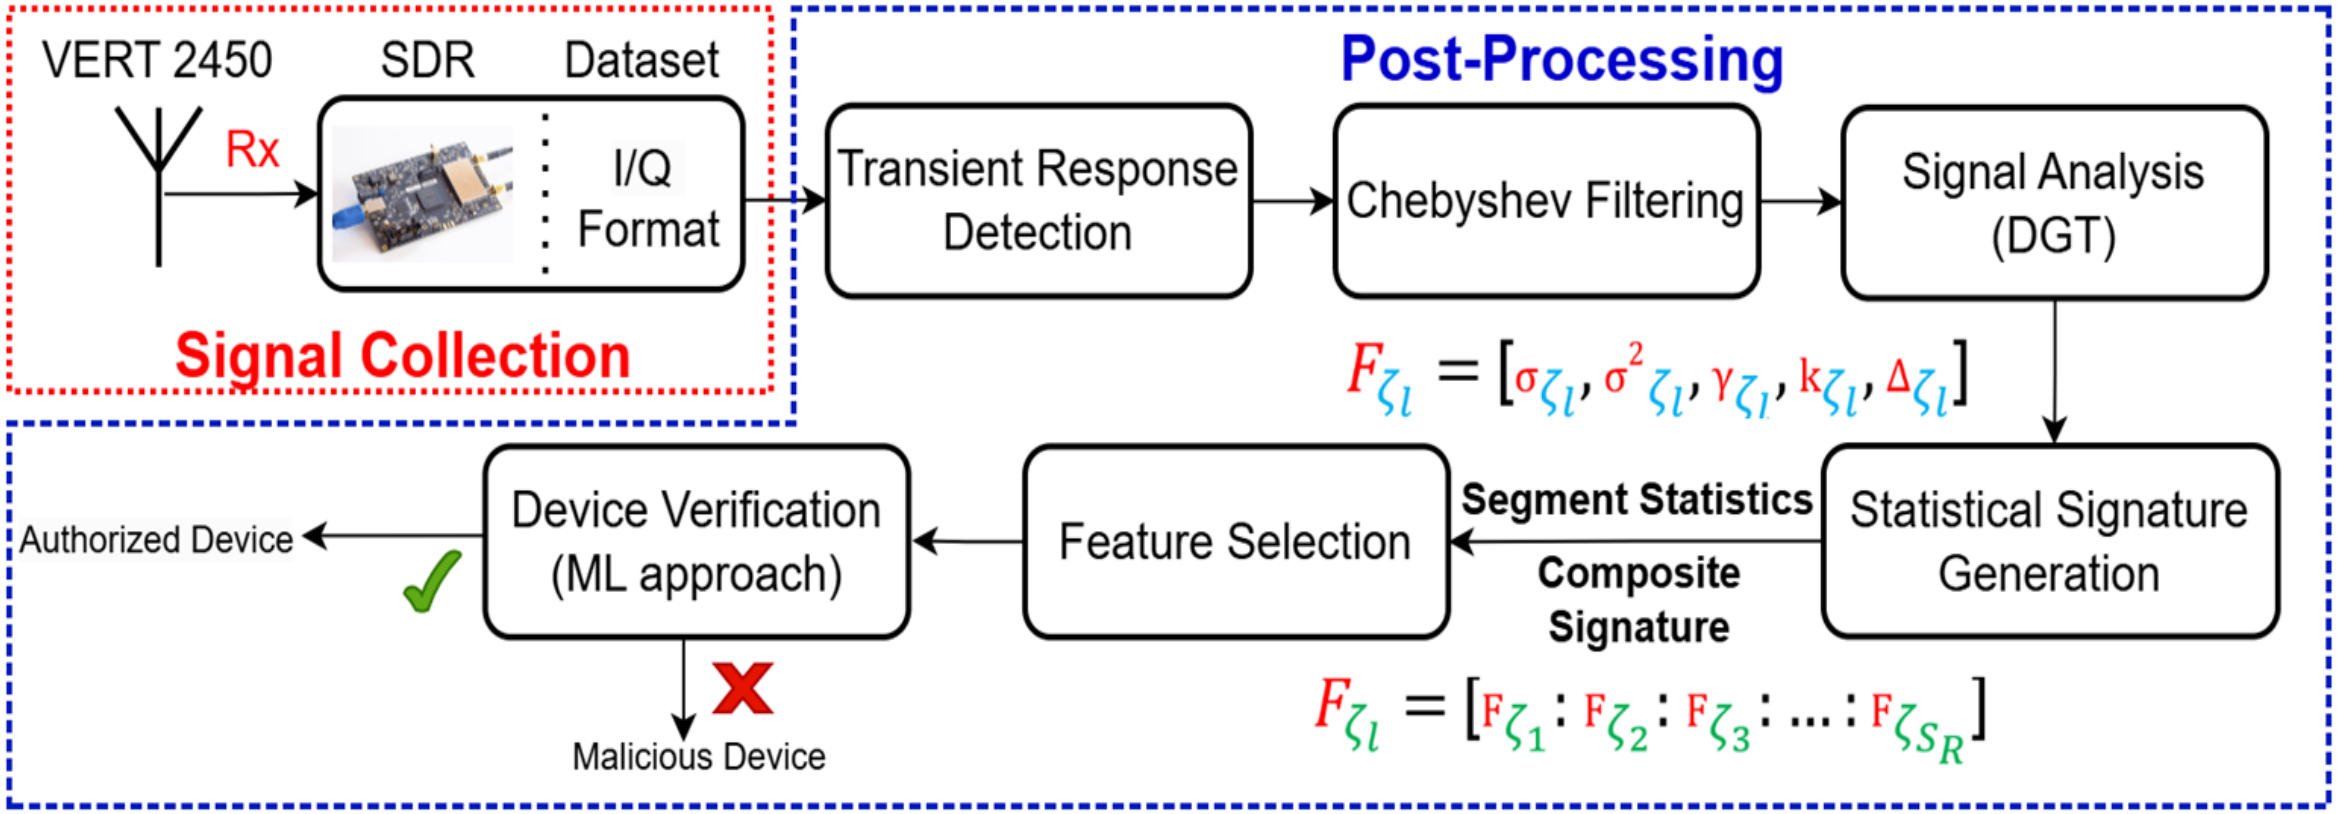
VERT 2450
SDR
Dataset
Rx
I/Q Format
Signal Collection
Post-Processing
Transient Response Detection
Chebyshev Filtering
Signal Analysis (DGT)
$F_{\zeta_l} = [\sigma_{\zeta_l}, \sigma^2_{\zeta_l}, \gamma_{\zeta_l}, k_{\zeta_l}, \Delta_{\zeta_l}]$
Statistical Signature Generation
Segment Statistics
Composite Signature
Feature Selection
Device Verification (ML approach)
Authorized Device
Malicious Device
$F_{\zeta_l} = [F_{\zeta_1} : F_{\zeta_2} : F_{\zeta_3} : \ldots : F_{\zeta_{S_R}}]$

Results

- 95% - 98% malicious devices detection rate.
  - Even with large number and diverse characteristics of malicious devices…
  - … and for varying SNR levels.
- Mean inference time: 3.75ms.
- Energy consumption: < 25.5mJ.

Coming Soon...

Open-Source Release

https://github.com/mlsysops-eu/mlsysops-framework

- Open-source components of the MLSysOps framework (late **6/2025**)
  - Orchestrators
  - Runtimes (sandboxed & generic)
  - Agent templates (continuum, cluster, node)
  - Policy
  - ML Connector

Hackathon & HiPEAC Workshop

- Hackathon organization (9/2025, University of Calabria, Rende, Italy)
  - Familiarize developers with the MLSysOps framework
  - Foster application & ML models development
- Workshop organization (ML4ECS - HiPEAC 2026)
  - https://ml4ecs.e-ce.uth.gr/
  - Edgeless + CODECO + MLSysOps

Thank you!

https://mlsysops.eu

mlsysops-eu/mlsysops-framework

@mlsysops

mlsysops

@mlsysopsproject

Newsletter https://mlsysops.eu/communication/