# P4Pir: In-Network Analysis for Smart IoT Gateways

Mingyuan Zang*, Changgang Zheng§, Lars Dittmann*, and Noa Zilberman§

*Technical University of Denmark, §University of Oxford

* ✉ {minza@dtu.dk, ladit@dtu.dk}, § ✉ {changgang.zheng@eng.ox.ac.uk, noa.zilberman@eng.ox.ac.uk}

# In-Network Classification

## IIsy [1], Planter [2]

- A trained ML model → programmable network devices
- Support >11 ML models (+50 variants): Bayes, SVM, DT, NN, …
- Running on Intel Tofino switch, AMD FPGA, NVIDIA DPU, …
- Line-rate performance

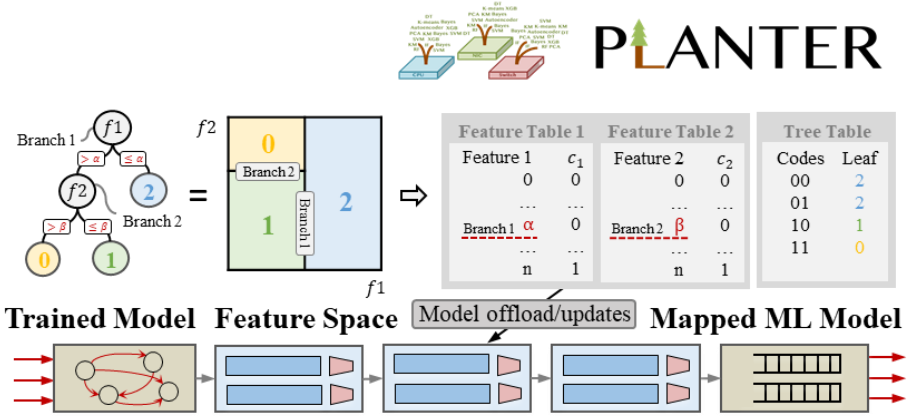


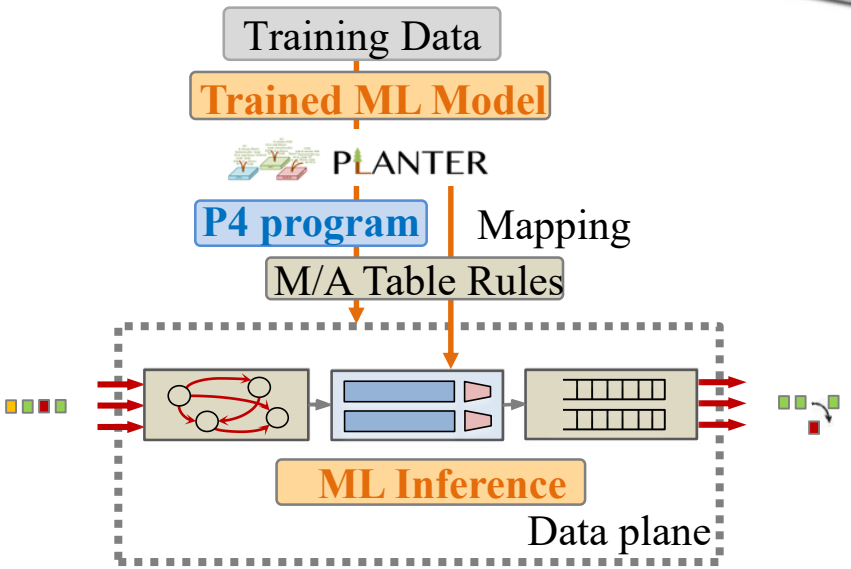



[1] C. Zheng et al., "IIsy: Practical In-Network Classification," arXiv preprint arXiv:2205.08243, 2022

[2] C. Zheng et al., "Automating In-Network Machine Learning," arXiv preprint arXiv:2205.08824, 2022

Can In-Network Classification bring benefits to traffic analysis in the Internet of Things (IoT) networks?

# IoT Network – What's different?

## 5G Requirements

Extremely low latency requirements

e.g. Process automation (latency<50ms)[1]

## IoT Security

Emerging attack variants

→ Threat to network infrastructure

## Distributed Devices

Limited computing resources

→ Lack of security measures

→ ***Fast spreading threats with changing patterns***

Typical Solution: Cloud-based services are limited in fast reaction

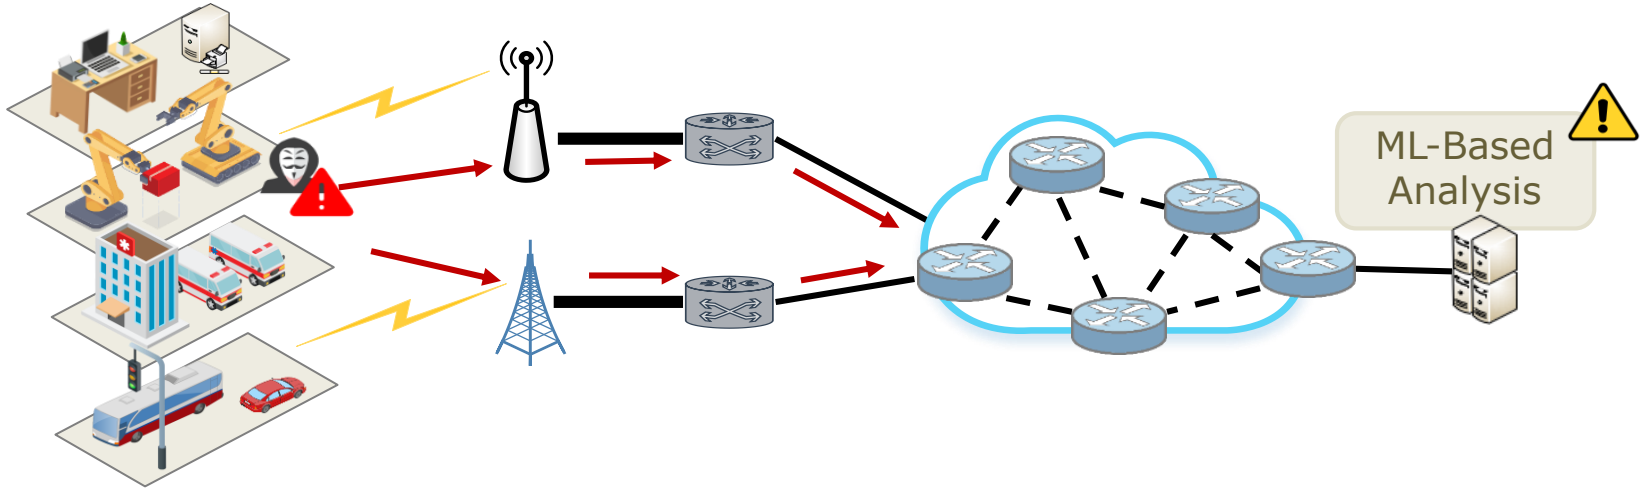


[1] "Service requirements for the 5G system (3GPP TS 22.261 version 16.14.0 Release 16)," 3GPP, Standard, Apr. 2021.

# IoT Network – What's different?

## 5G Requirements

Extremely low latency requirements

e.g. Process automation (latency<50ms)[1]

## IoT Security

Emerging attack variants

→ Threat to network infrastructure

## Distributed Devices

Limited computing resources

→ Lack of security measures

→ ***Fast spreading threats with changing patterns***

Our Solution: Offload ML capability to network edge devices for fast mitigation

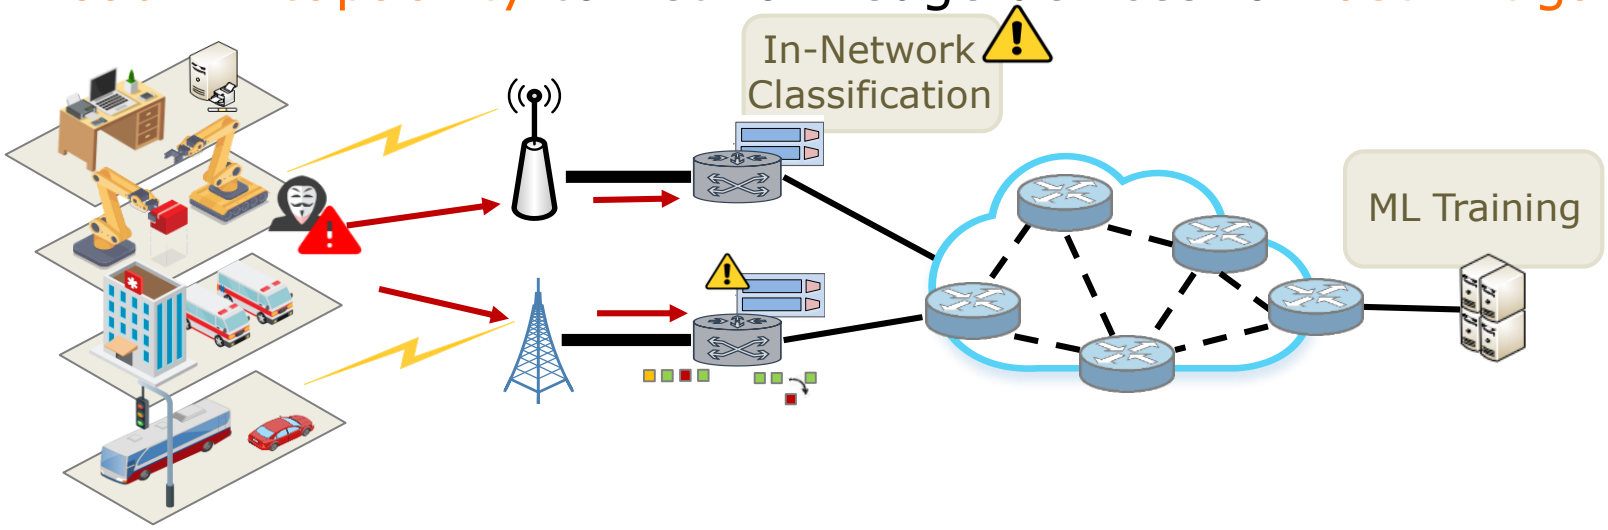


[1] "Service requirements for the 5G system (3GPP TS 22.261 version 16.14.0 Release 16)," 3GPP, Standard, Apr. 2021.

# Design Challenges

- Cheap IoT gateway devices?
- Continuous defense? (24×7 security operation)
- Distributed deployment?

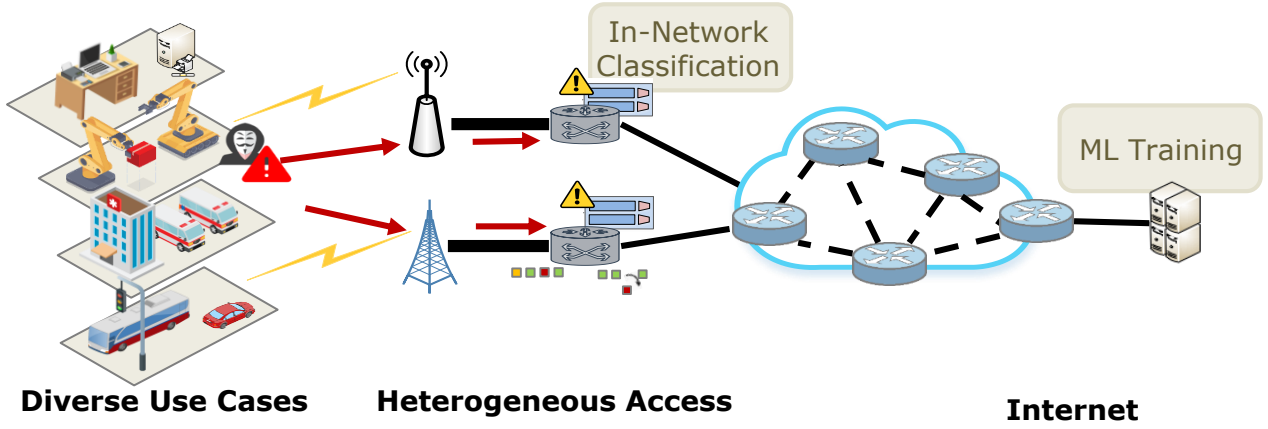

# Challenge I: Continuous Defense

## Proposed Design *P4Pir*

**Solution: Real-time traffic analysis with in-network classification**

- In-network classification on cheap device (P4Pi – P4 in RPi)
- Tree-based ML for lightweight deployment
- Runtime ML updates for continuous defense

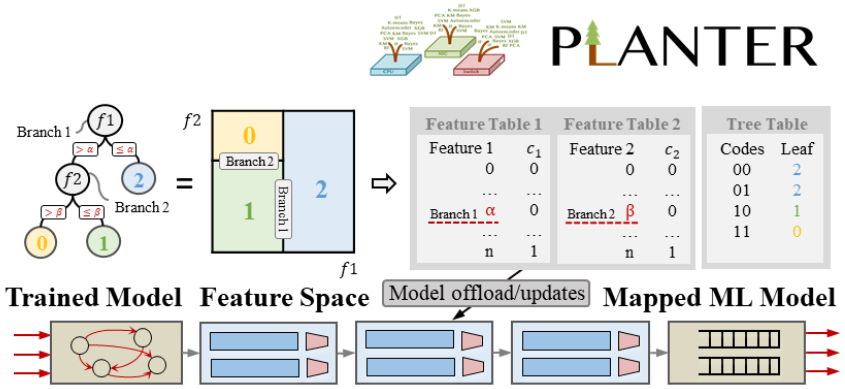


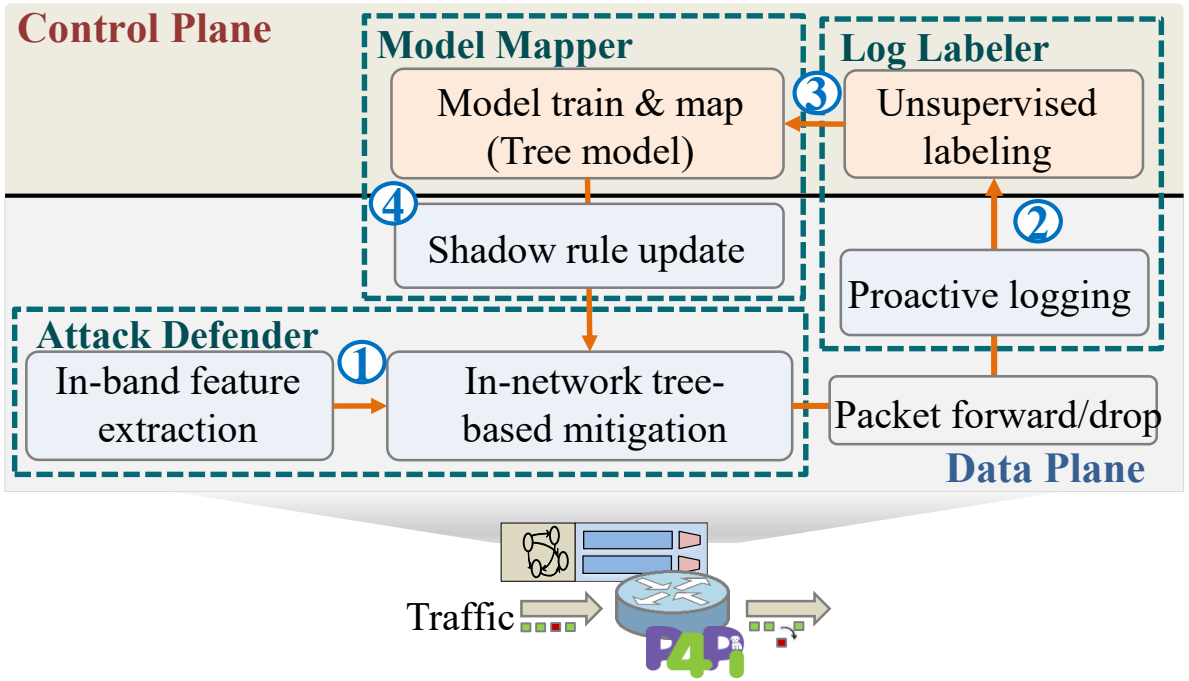


M. Zang et al., "Towards Continuous Threat Defense: In-Network Traffic Analysis for IoT Gateways," IEEE Internet of Things Journal, 2023

# Challenge I: Continuous Defense

## Proposed Design *P4Pir*

**Solution: Runtime reconfiguration for in-network model**

- Digest-based logging
- Proactive labeling & retraining
- Hitless shadow M/A rule (classification rule) updates
- Avoid function recompilation/forwarding disruption

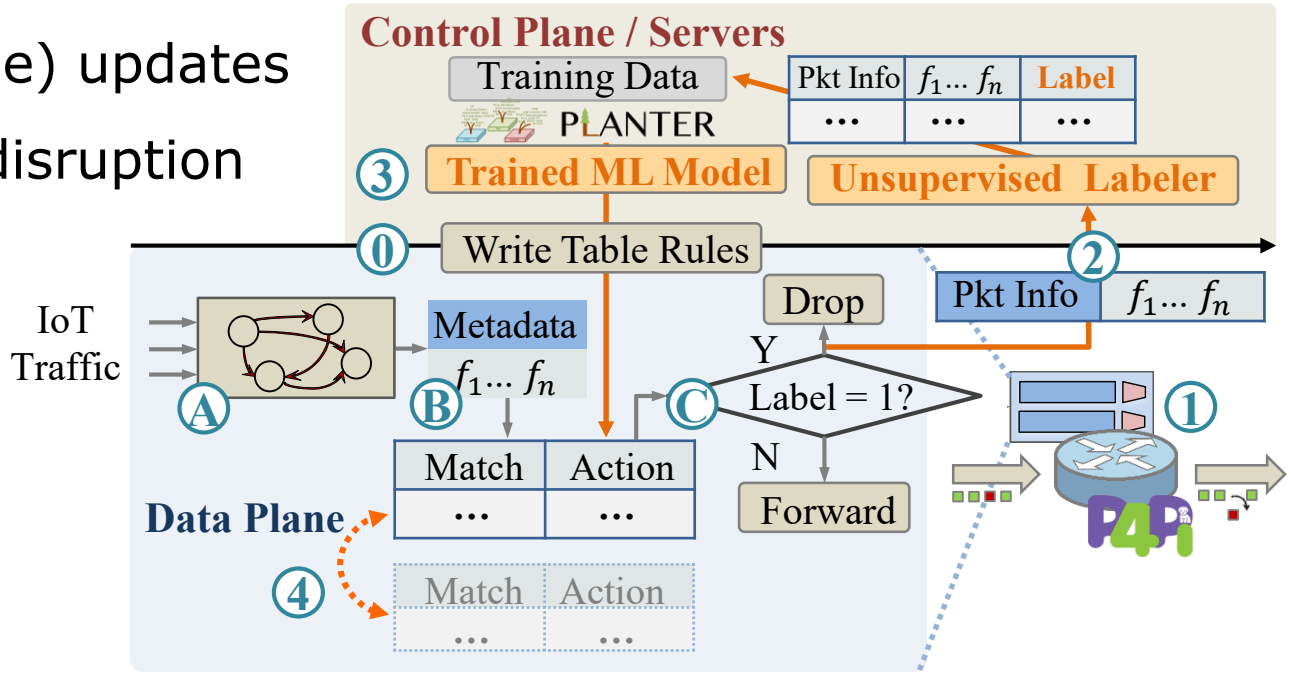


M. Zang et al., "Towards Continuous Threat Defense: In-Network Traffic Analysis for IoT Gateways," IEEE Internet of Things Journal, 2023

# Challenge I: Continuous Defense

## Proposed Design *P4Pir*

### Prototype

- P4Pi [1] (Raspberry Pi 4 Model B), Dell EMC Edge Gateway 5200

### Performance

- >30% accuracy ↑, real-time mitigation, negligible jitter, 8% ↑ on CPU utilization

TABLE III
DETECTION ACCURACY ON DATASET CICIDS 2017.

| | | SCAN | SCAN→DOS | | SCAN→BOT* | |
|---|---|---|---|---|---|---|
| | | Init | Base | P4Pir | Base | P4Pir |
| DT | ACC | 0.987 | 0.604 | 0.932 | 0.900 | 0.923 |
| | F1 | 0.984 | 0.568 | 0.868 | 0.776 | 0.820 |
| RF | ACC | 0.989 | 0.731 | 0.942 | 0.987 | 0.989 |
| | F1 | 0.985 | 0.027 | 0.869 | 0.964 | 0.987 |

* Init - Initial state, Base - Baseline from static in-network inference model, SCAN - port scanning attack, DoS - DDoS LOIT attack, BOT - Botnet ARES attack. "→" indicates the change in attack pattern from the initial state to an emerging attack

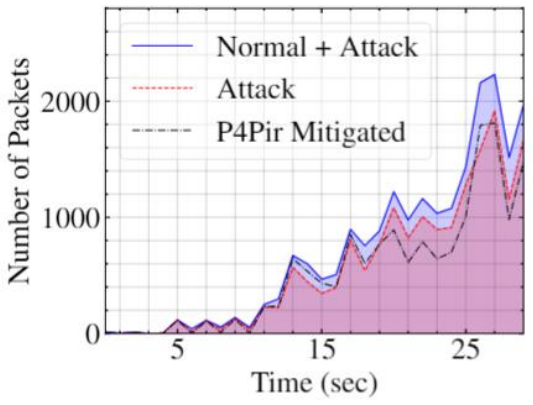


(a) Mitigation performance.

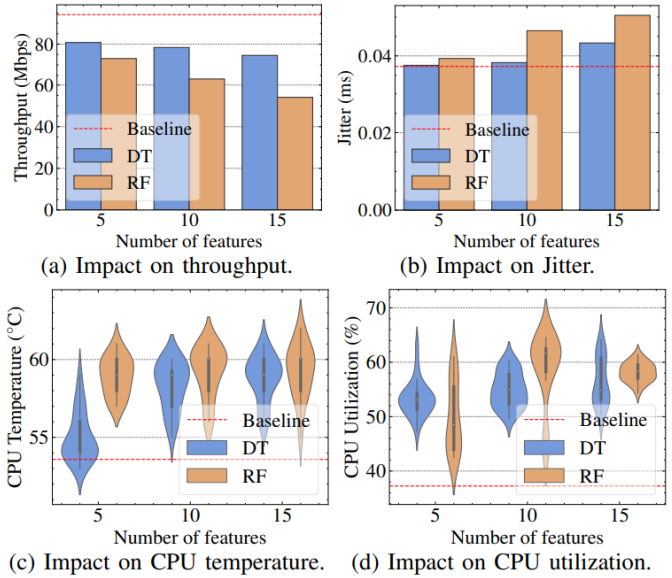


(a) Impact on throughput. (b) Impact on Jitter.
(c) Impact on CPU temperature. (d) Impact on CPU utilization.

[1] S. Laki et al., "P4Pi: P4 on Raspberry Pi for networking education," SIGCOMM Comput. Commun. Rev., vol. 51, no. 3, p. 17–21, 2021

# Challenge II: Distributed Deployment

## Proposed Design *FLIP4*

**Solution: Federated learning-based in-network classification:**

- In-network classification in distributed IoT gateways
- Federated ML training & updates
- Privacy-preserving model parameters sharing

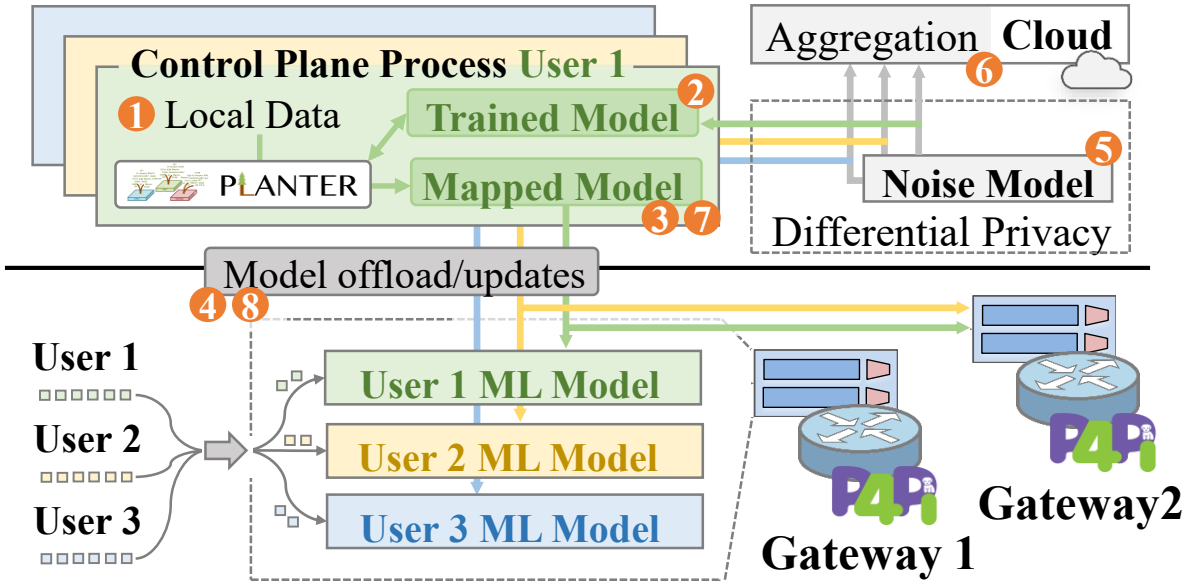


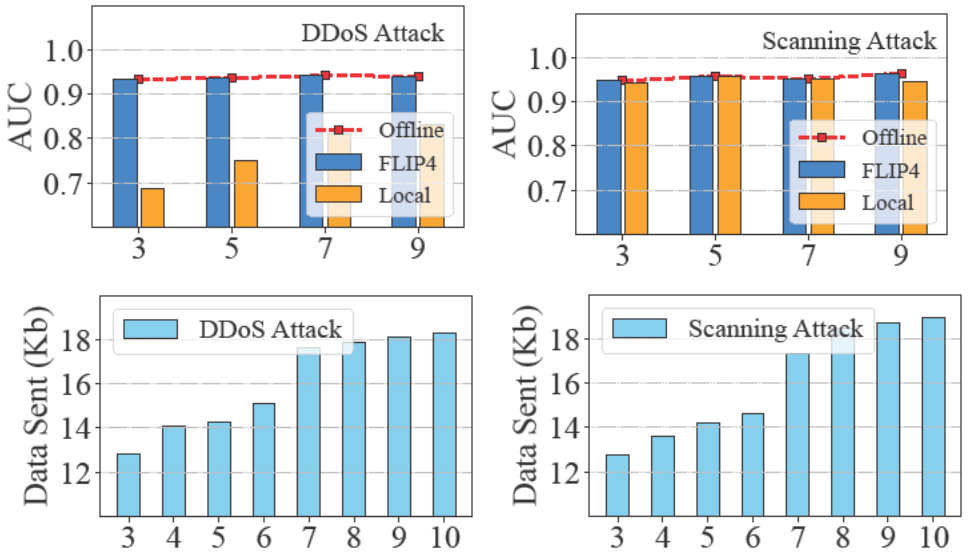


M. Zang et al., "Federated Learning-Based In-Network Traffic Analysis on IoT Edge," IFIP Networking 2023 - Sec4IoT, 2023

# Summary

In-network classification brings benefits to IoT scenarios:

- Feasible on cheap IoT gateway devices
- Swift analysis and reaction to detected incidents
- Scalable to distributed devices

**Further work:**

- Optimized resources?
- More services?

Open-source codes:

# List of Papers


[1] C. Zheng et al., “IIsy: Practical In-Network Classification,” arXiv:2205.08243, 2022

[2] C. Zheng et al., “Automating In-Network Machine Learning,” arXiv:2205.08824, 2022

[3] M. Zang et al., “P4Pir: In-Network Analysis for smart IoT gateways,” Proceedings of the SIGCOMM '22 Poster and Demo Sessions, 2022

[4] M. Zang et al., “Towards Continuous Threat Defense: In-Network Traffic Analysis for IoT Gateways,” IEEE Internet of Things Journal, 2023

[5] M. Zang et al., “Federated Learning-Based In-Network Traffic Analysis on IoT Edge,” IFIP Networking - Sec4IoT, 2023

[6] X. Hong et al., “LOBIN: In-Network Machine Learning for Limit Order Books,” IEEE HPSR, 2023



*We acknowledge the support from VMware, EU Horizon SMARTEDGE (101092908), Otto Mønsted Foundation, Nordic University Hub on Industrial IoT (HI2OT) by NordForsk.*

*We thank Radostin Stoyanov (Oxford), Damu Ding (Oxford), Eder Ollora Zaballa (DTU), Tomasz Koziak (DTU) for help with experimental setups and valuable discussions and feedbacks.*