Towards Real-Time Intrusion Detection in P4-Programmable 5G User Plane Functions

<u>Aristide Tanyi-Jong Akem</u>, Marco Fiore

IMDEA Networks Institute, Madrid, Spain

Euro'P4 2024

Outline

- In-Network Machine Learning
- P4-based 5G UPF
- Intrusion detection in the UPF
- Conclusion

Outline

- In-Network Machine Learning
- P4-based 5G UPF
- Intrusion detection in the UPF
- Conclusion

In-network machine learning

- Machine Learning (ML) is playing a key role in network automation
- In Software-Defined Networking (SDN), these models run in the control plane

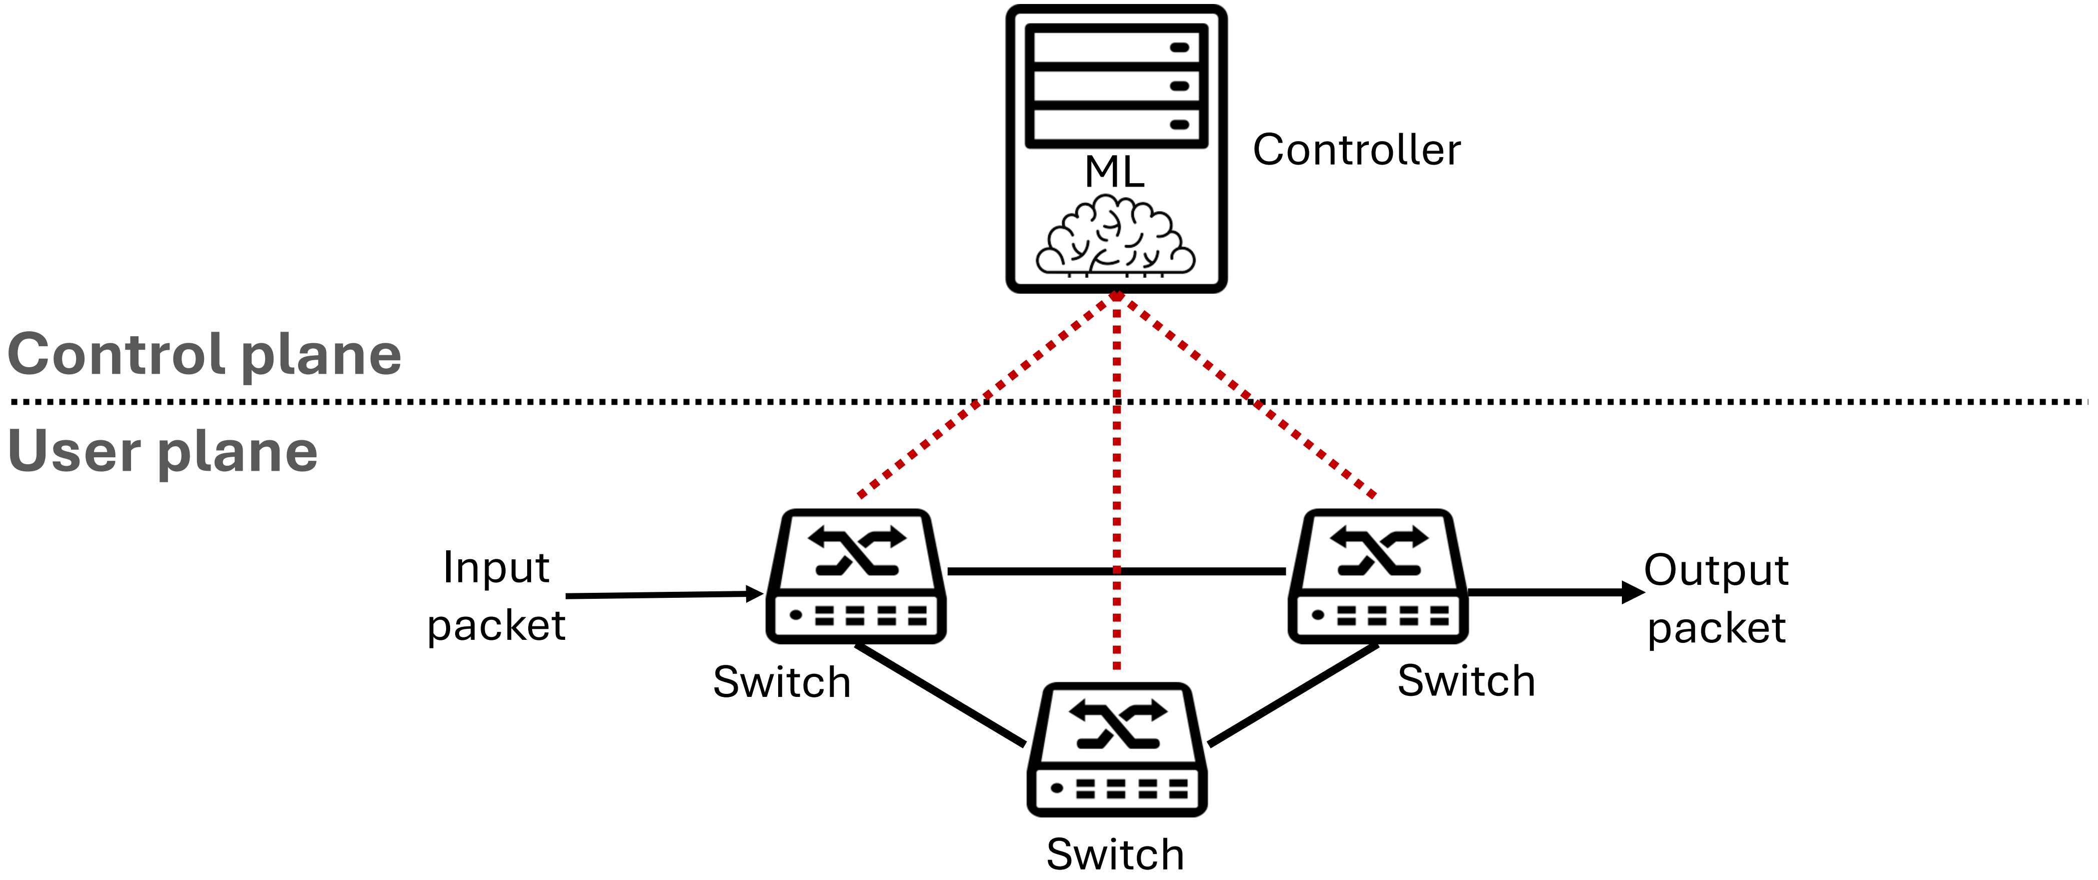

In-network machine learning

- Control plane ML requires back-and-forth communication with the user plane
- This induces ms-level delays which are undesirable in low-latency applications

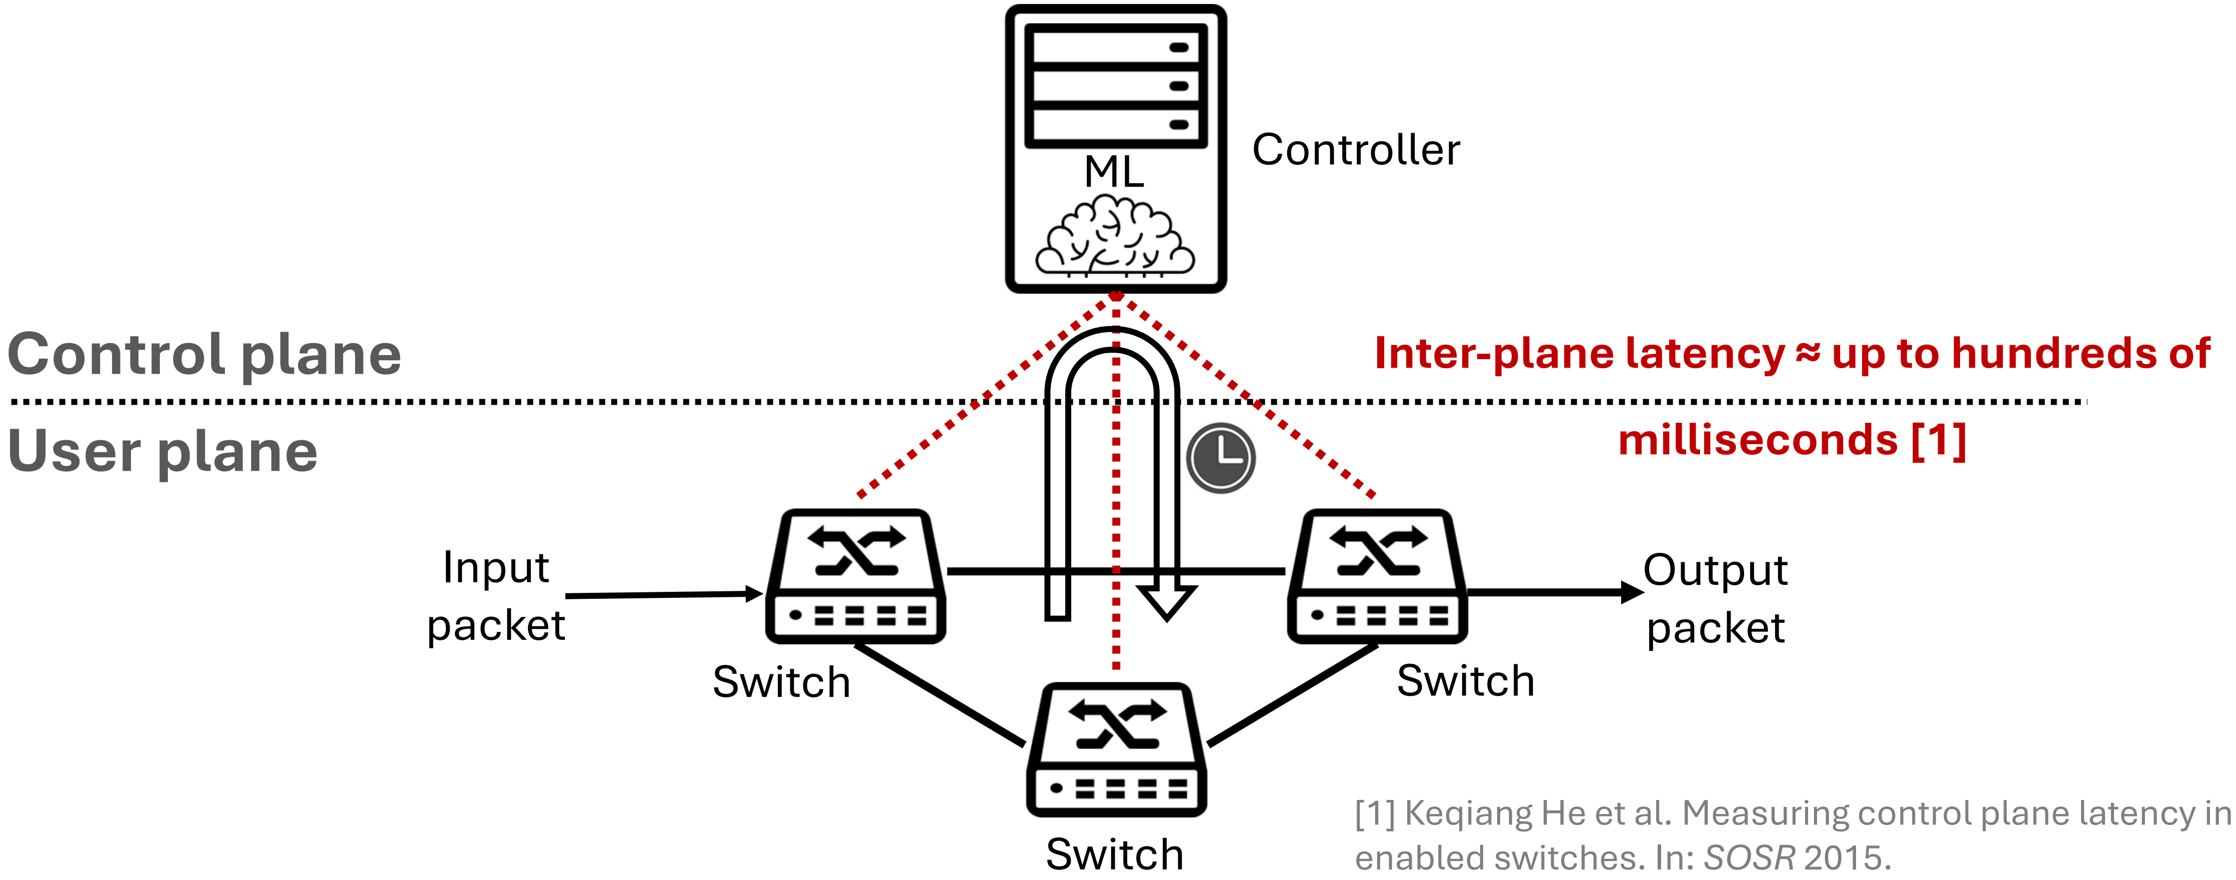

[1] Keqiang He et al. Measuring control plane latency in SDN-enabled switches. In: *SOSR* 2015.

In-network machine learning

The advent of programmable switches and domain-specific languages like P4 has made it possible to deploy trained ML models into the user plane

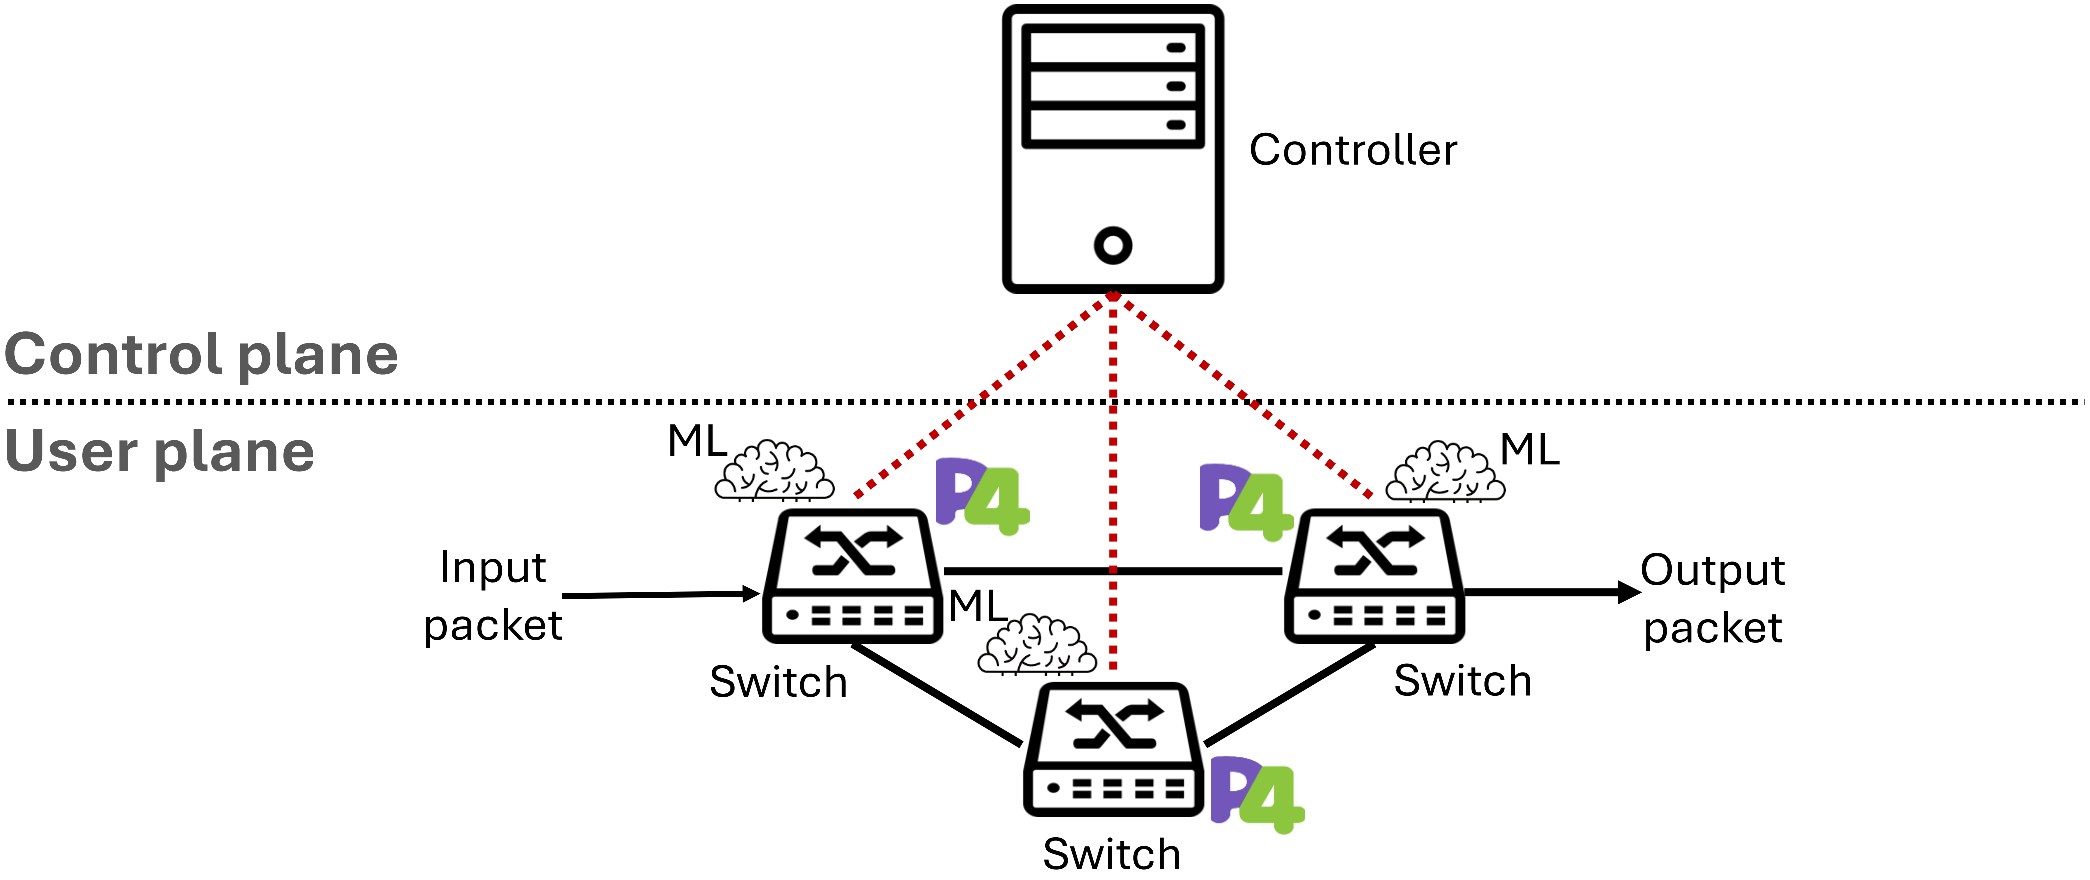

In-network machine learning

The advent of programmable switches and domain-specific languages like P4 has made it possible to deploy trained ML models into the user plane

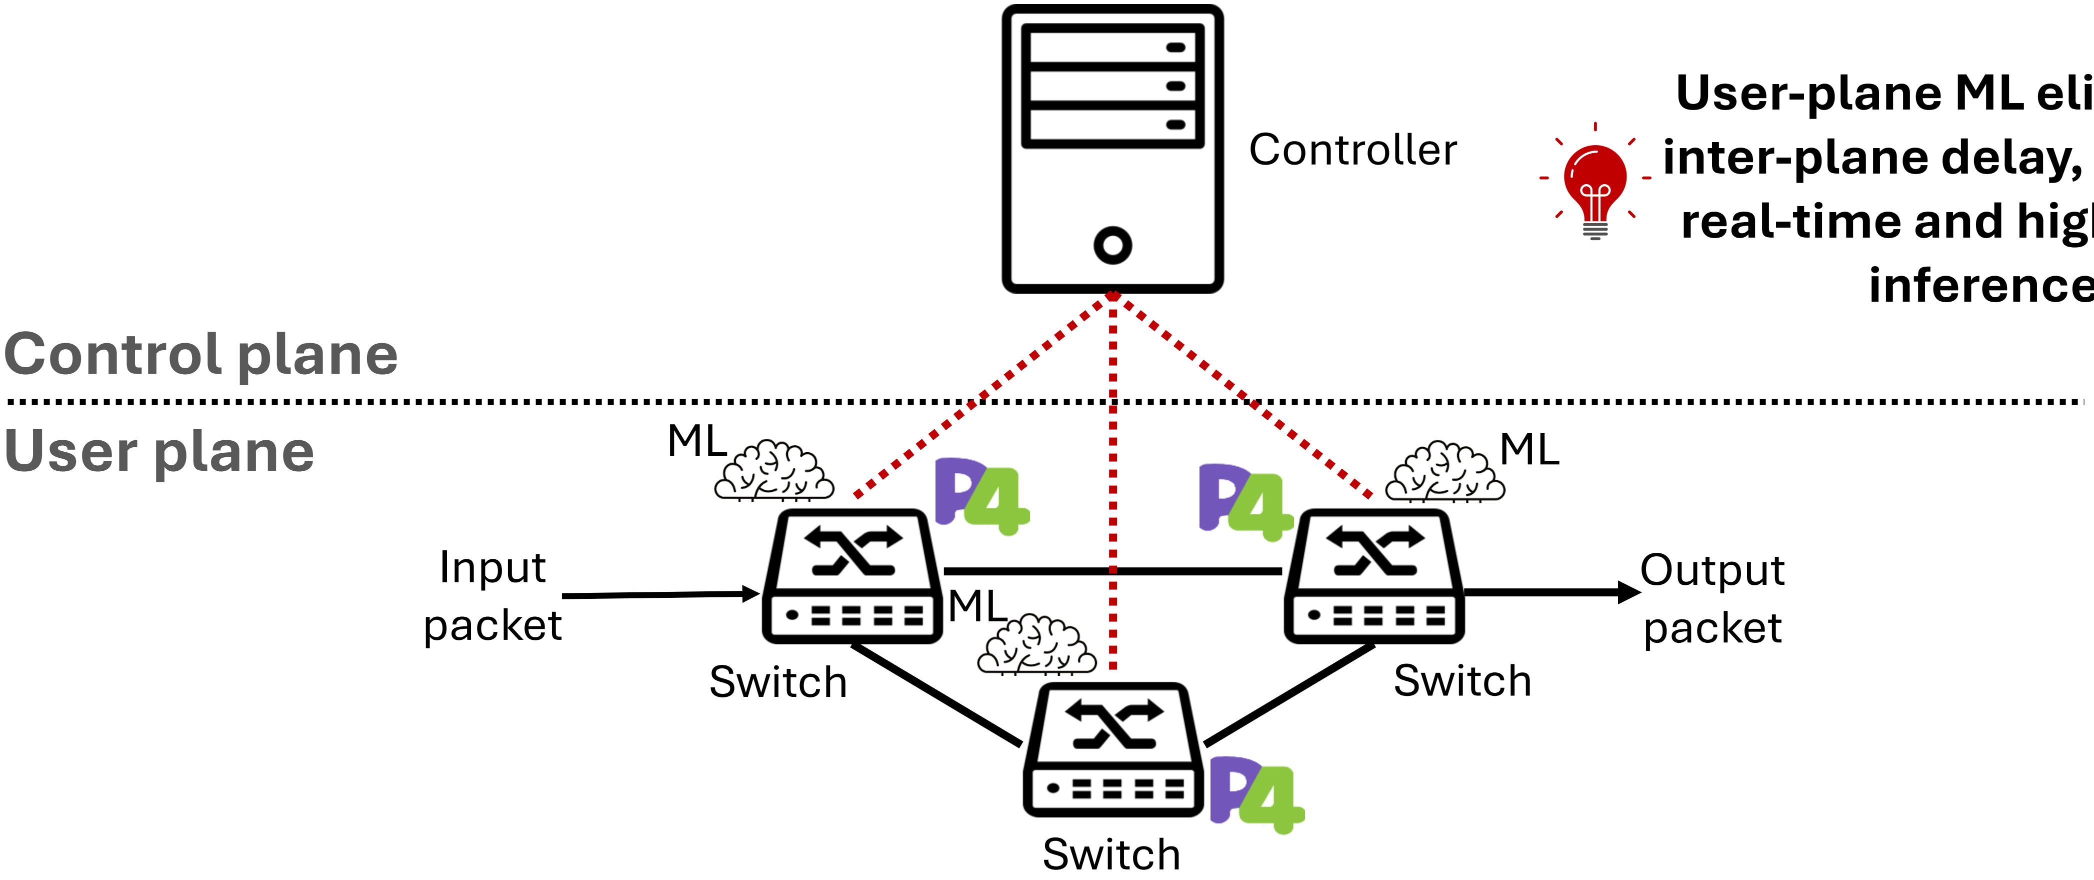

User-plane ML eliminates inter-plane delay, enabling real-time and high-speed inference

In-network machine learning

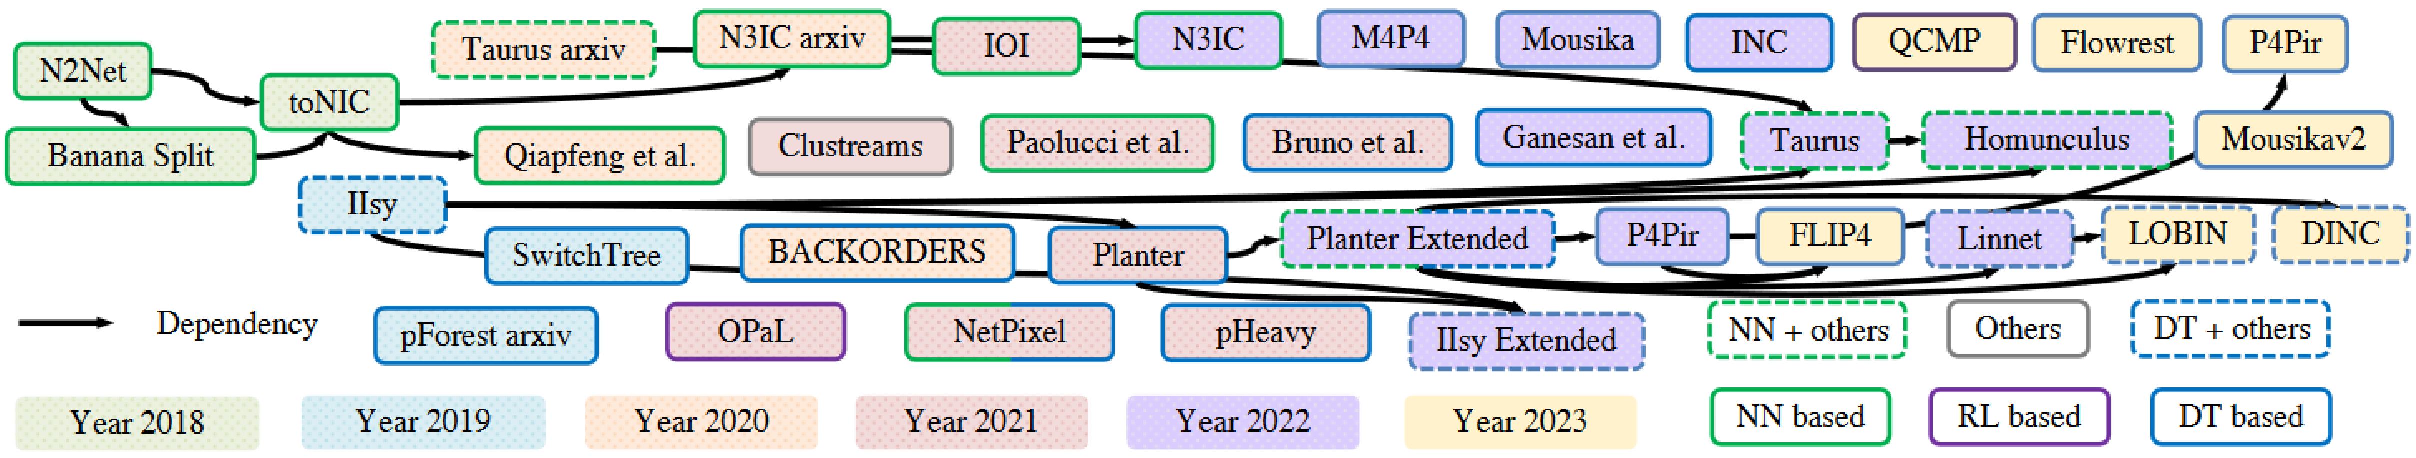

C. Zheng, X. Hong, D. Ding, S. Vargaftik, Y. Ben-Itzhak and N. Zilberman, "In-Network Machine Learning Using Programmable Network Devices: A Survey," in *IEEE Communications Surveys & Tutorials*, vol. 26, no. 2, 2024.

Outline

- In-Network Machine Learning
- P4-based 5G UPF
- Intrusion detection in the UPF
- Conclusion

5G User Plane Function (UPF)

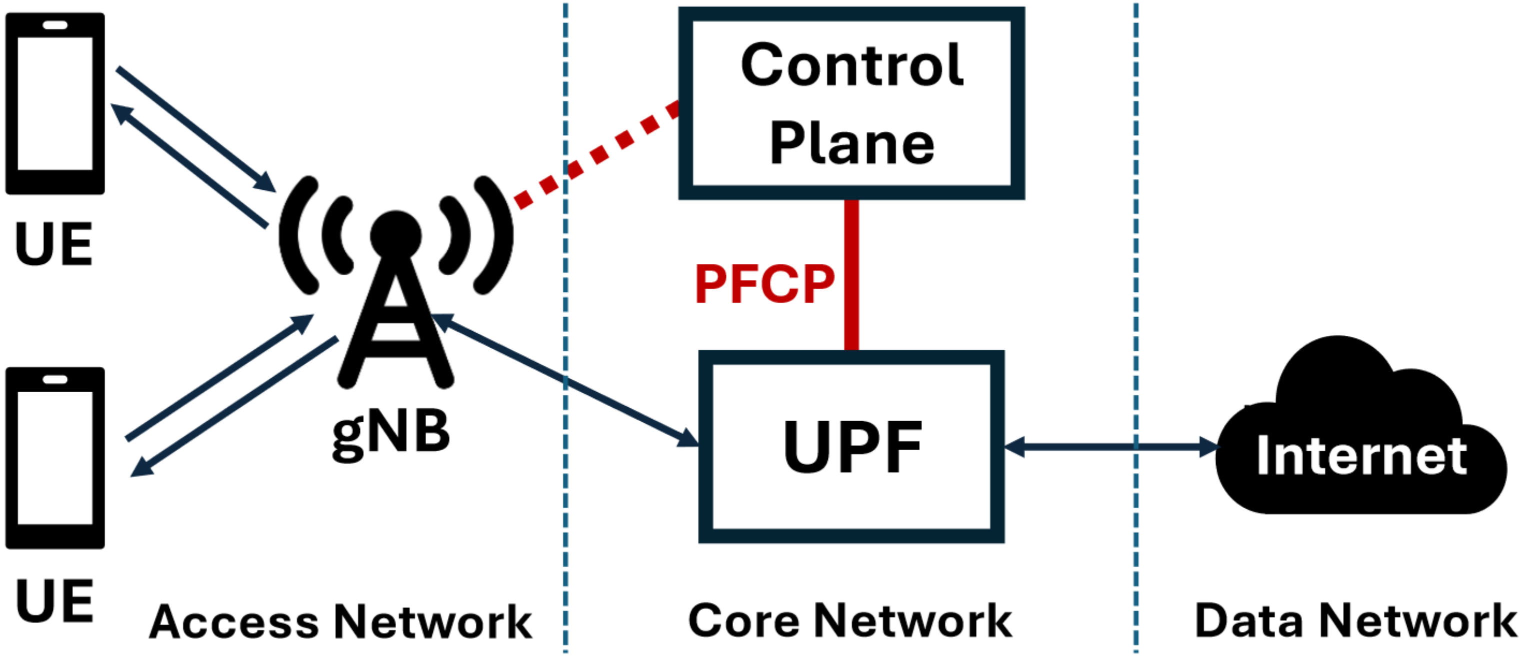

P4-based 5G User Plane Function (UPF)

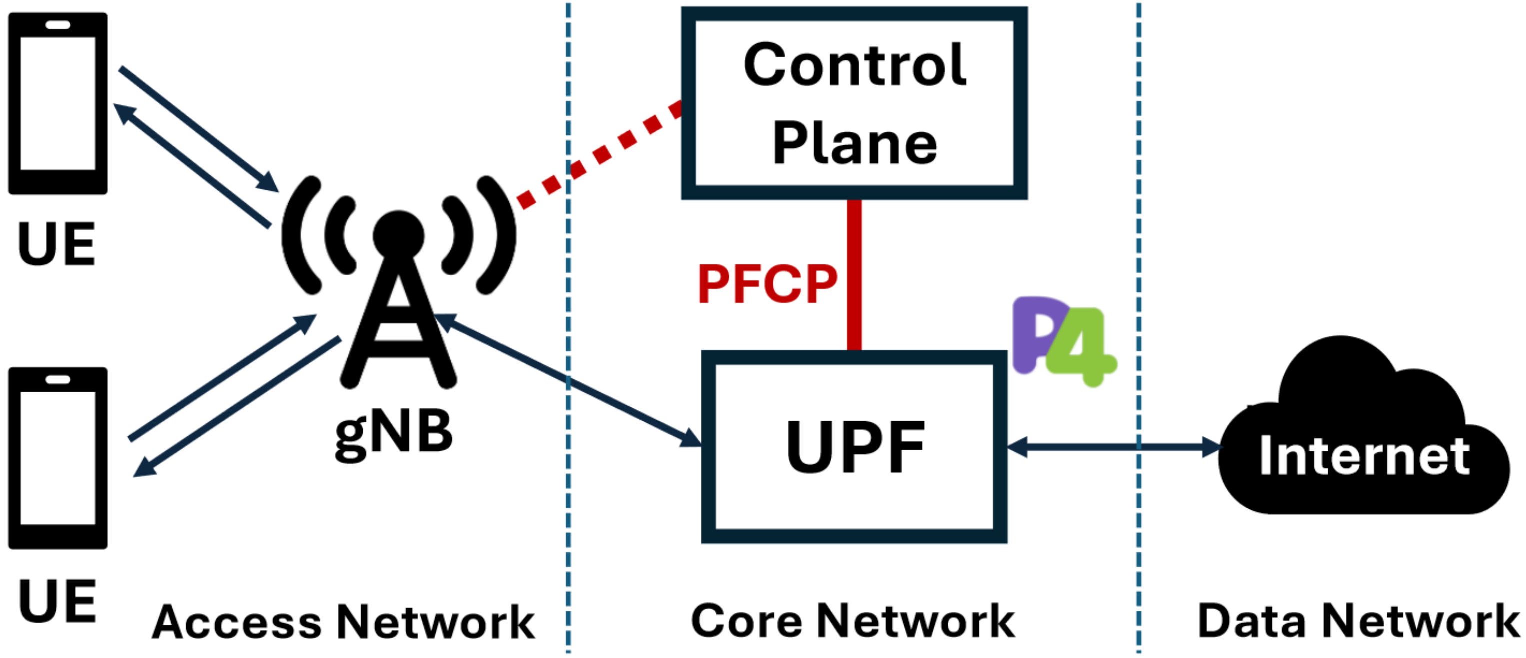

[1] R. MacDavid et al. A P4-based 5G user plane function. In SOSR. ACM, 2021.

[2] A. Bose et al. AccelUPF: accelerating the 5G user plane using programmable hardware. In SOSR. ACM, 2022.

UPF packet processing flow

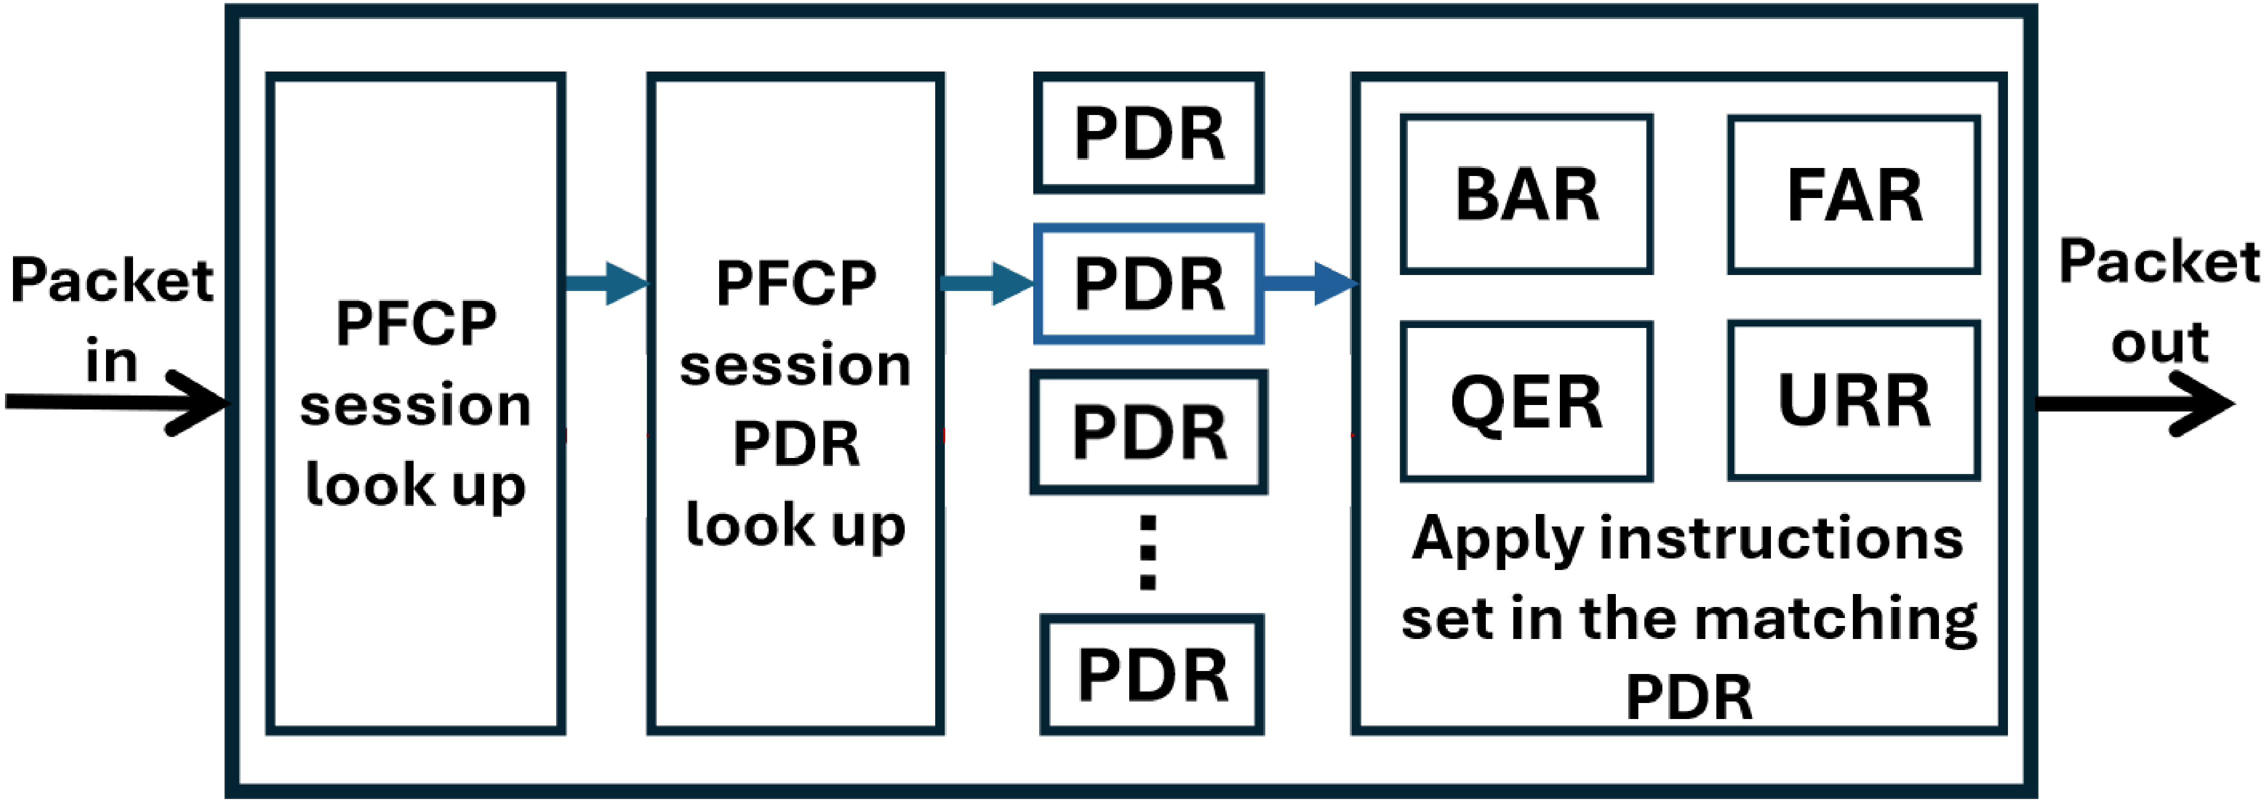

UPF packet processing flow

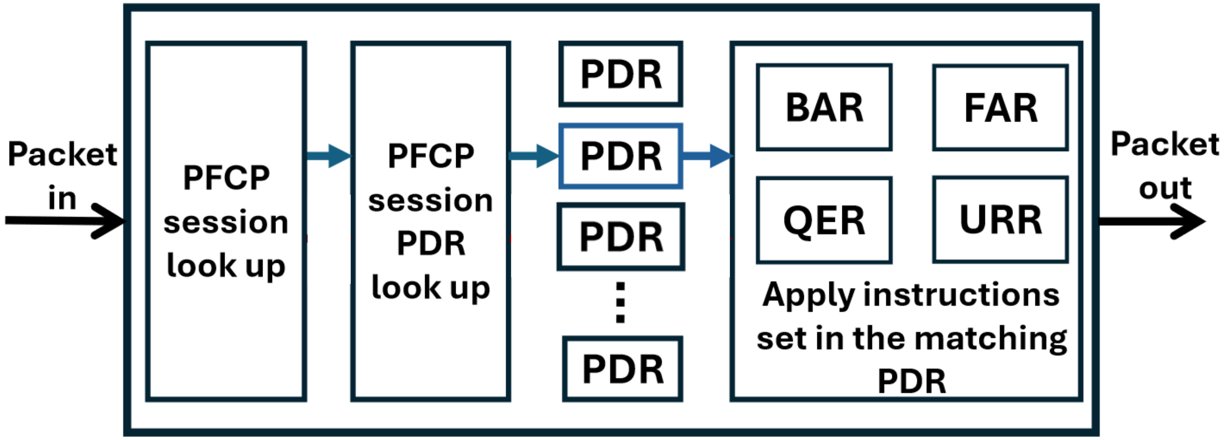

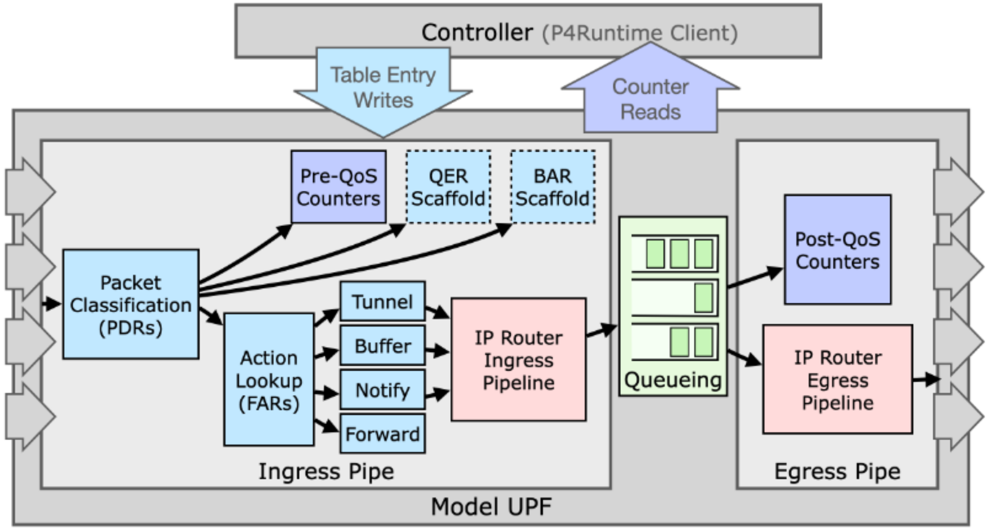

[1] R. MacDavid et al. A P4-based 5G user plane function. In SOSR. ACM, 2021.

Outline

- In-Network Machine Learning
- P4-based 5G UPF
- **Intrusion detection in the UPF**
- Conclusion

Integrating ML into the P4-based 5G UPF

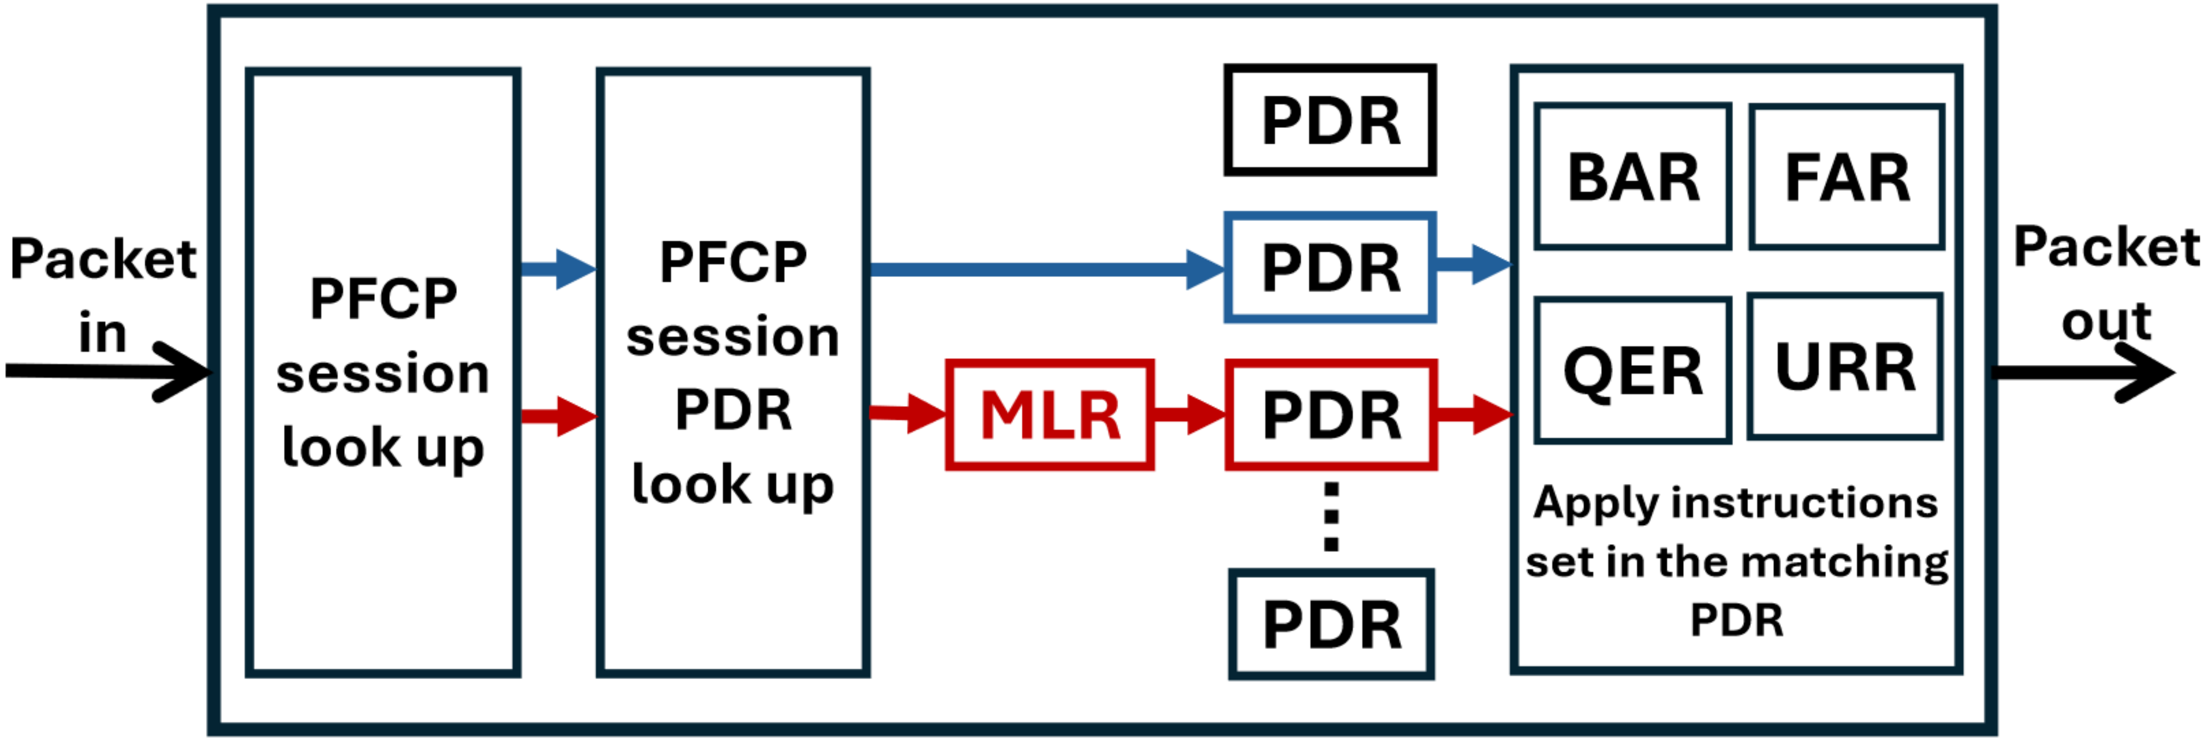

Integrating ML into the P4-based 5G UPF

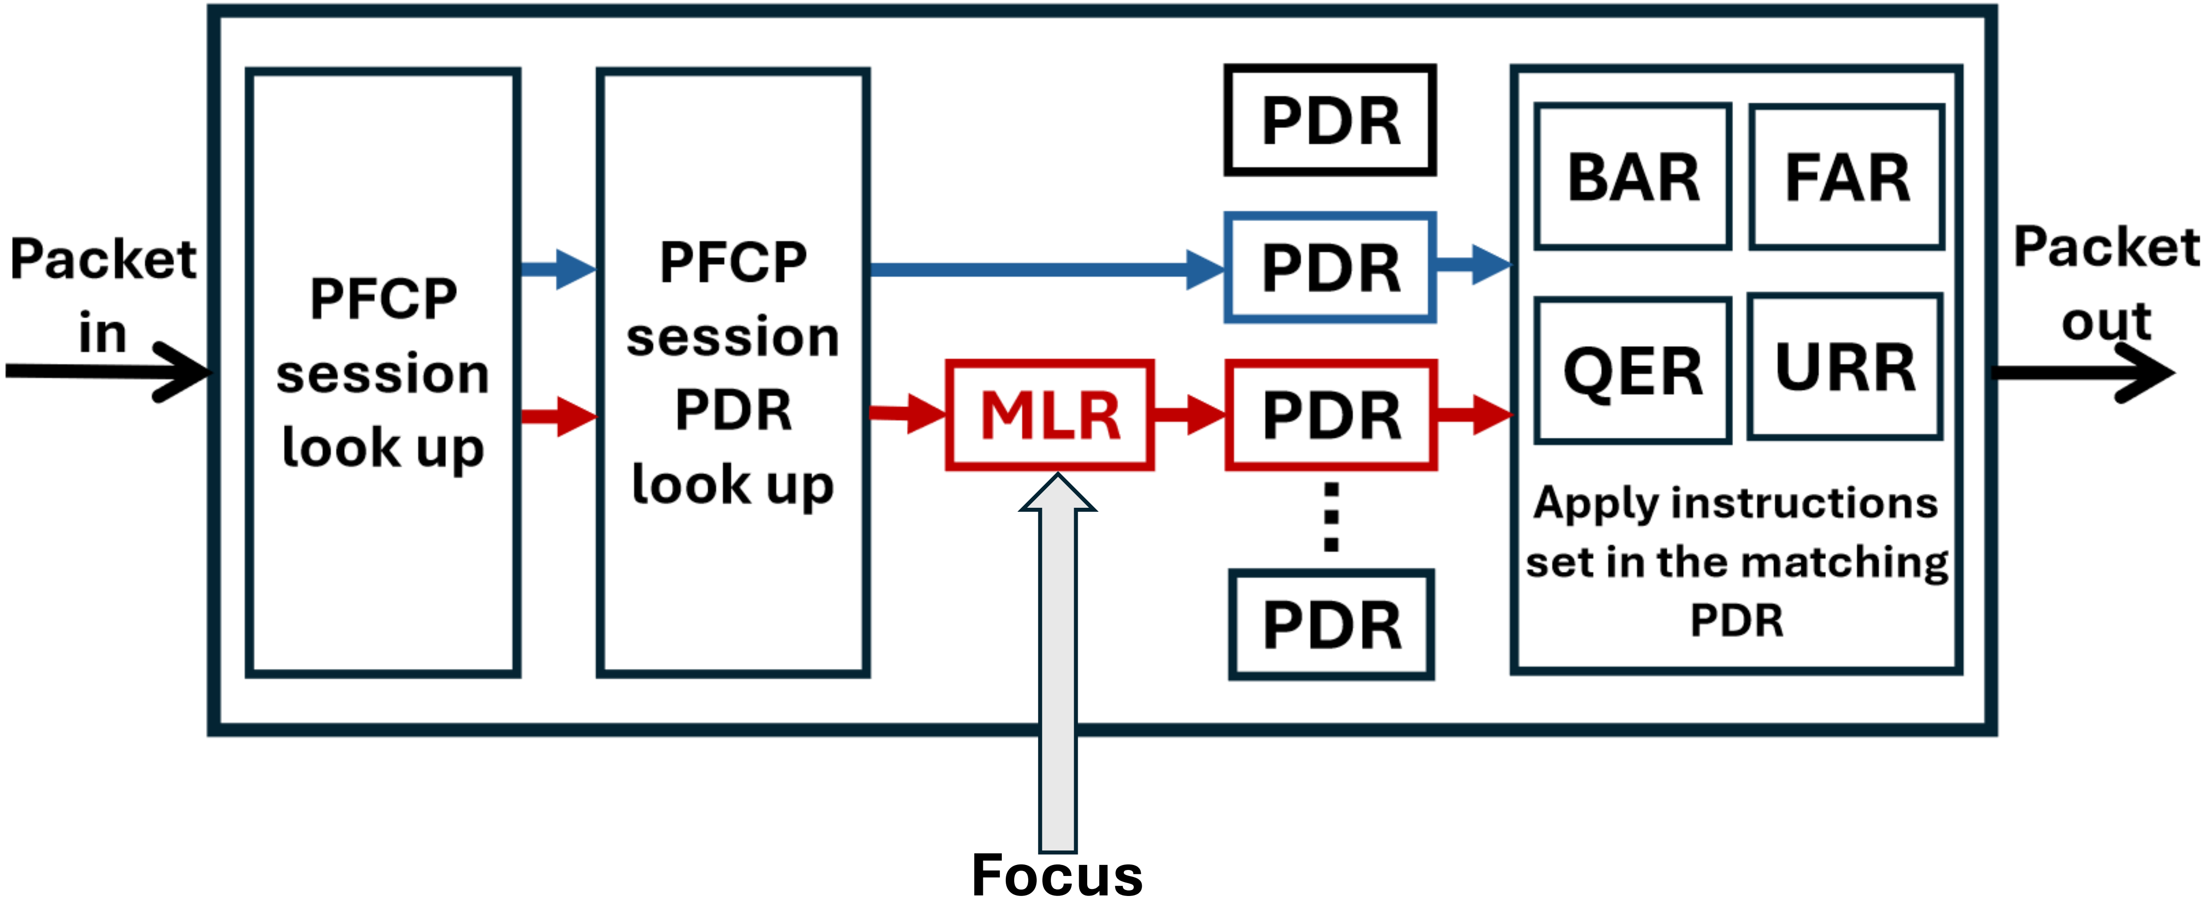

Intrusion detection in the UPF with In-switch ML

Tree-based models are most suitable for in-switch ML

- Their simple logical structure makes them easy to map to switch pipelines [1]

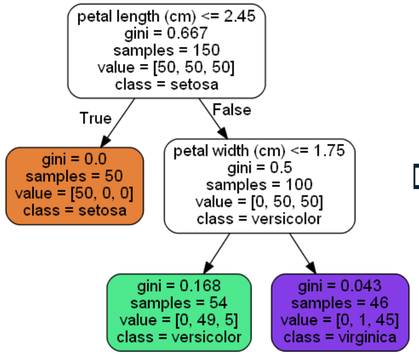

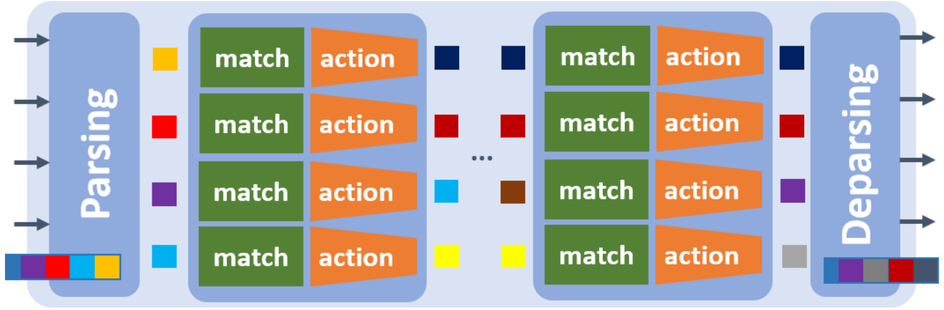

- They still outperform deep learning on tabular data [2]

[1] Zhaoqi Xiong et al. Do Switches Dream of Machine Learning? Toward In-Network Classification. In HotNets. ACM, 2019.

[2] Léo Grinsztajn, et al. Why do tree-based models still outperform deep learning on typical tabular data? In NeurIPS, 2022.

In-switch ML inference workflow

Dataset (.pcap) → Feature Extraction and computation (Tshark) → Model Training (Python) → Model translation to M/A table entries (Python)

ML server

Control plane

User plane

In-switch ML inference workflow

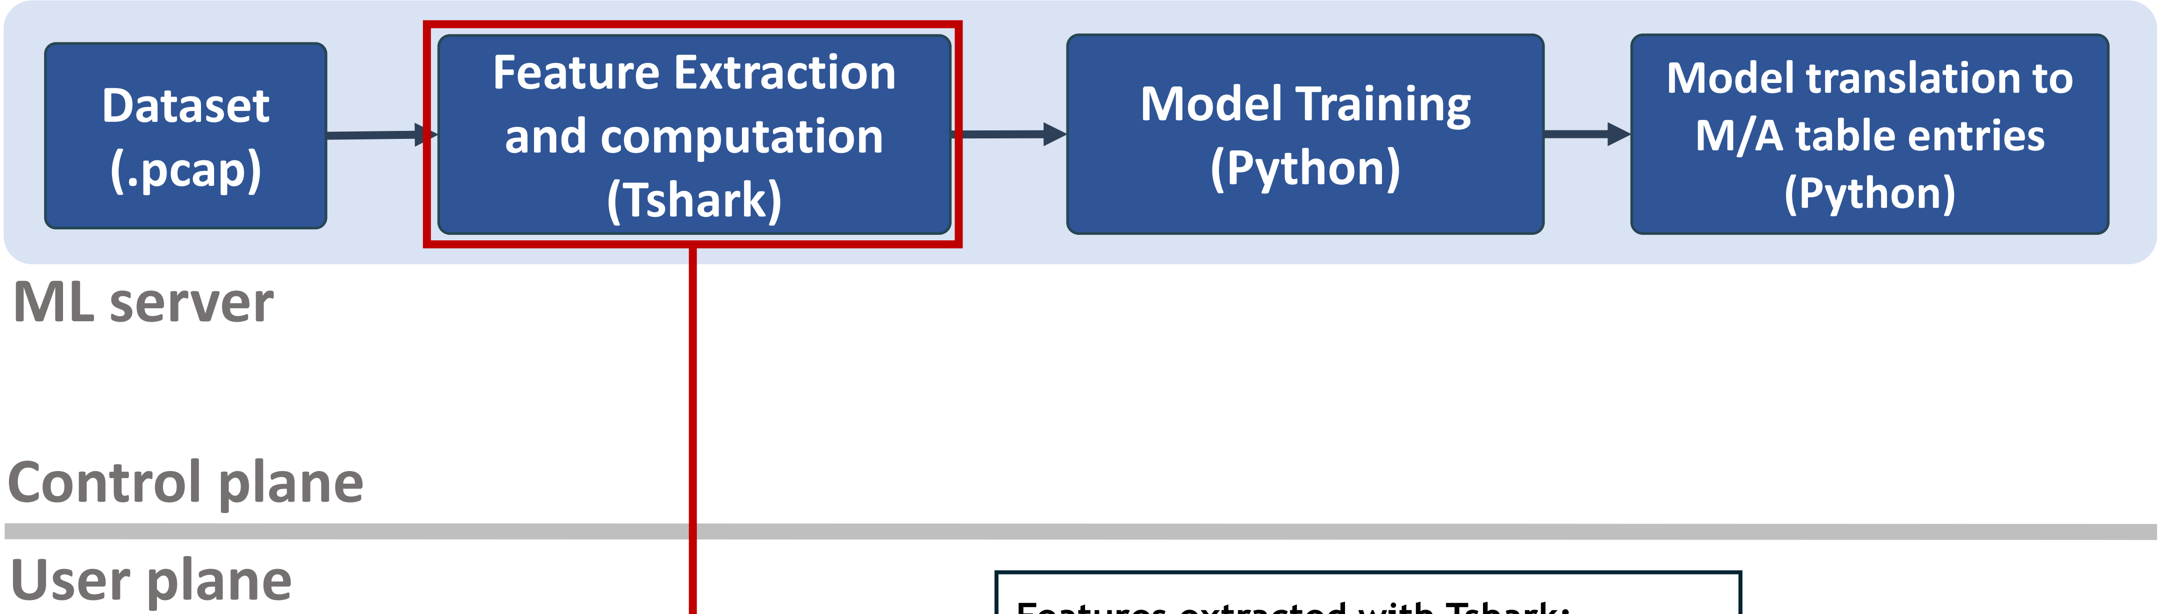

Features extracted with Tshark:
- Packet length,
- Source port & destination port,
- Protocol,
- TCP flags (SYN, ACK, FIN, PSH, RST),
- TCP header length,
- TCP window size,
- UDP length,
- Time-to-live (TTL).

In-switch ML inference workflow

Dataset (.pcap) → Feature Extraction and computation (Tshark) → Model Training (Python) → Model translation to M/A table entries (Python)

ML server

Control plane

User plane

Model training:
- Scikit-Learn Python libraries

Feature selection:
- Importance as expressed by the Mean Decrease in Impurity (MDI)

Hyperparameters:
- Number of trees (for RF)
- *Max tree depth, etc.*

In-switch ML inference workflow

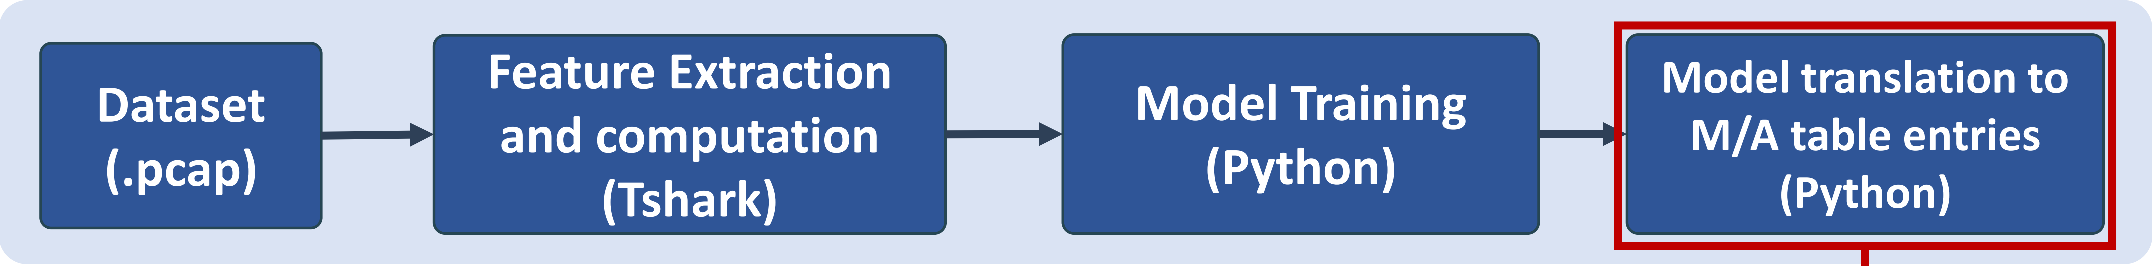

ML server

Control plane

User plane

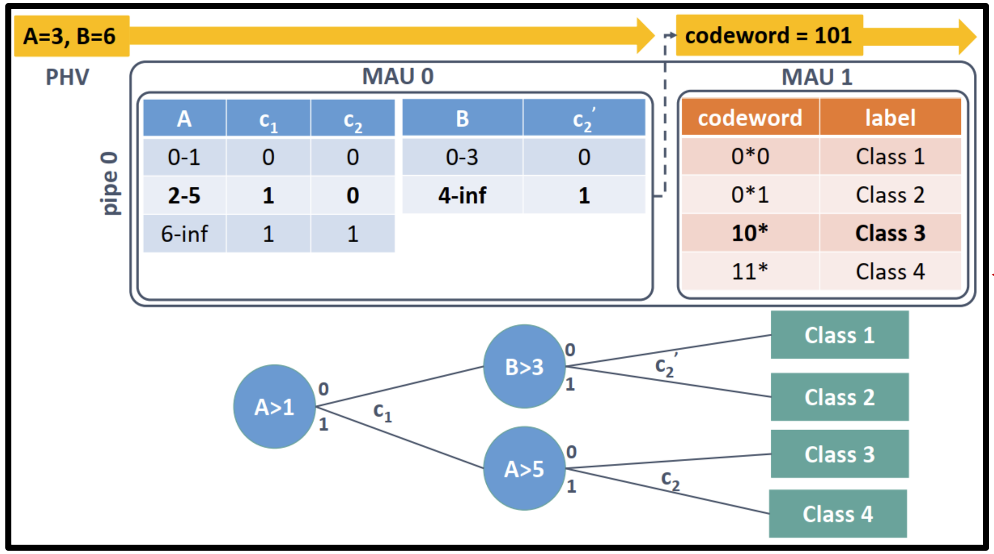

[1] C. Zheng and N. Zilberman. **Planter: Seeding trees within switches**. In SIGCOMM Poster Session. ACM, 2021.

In-switch ML inference workflow

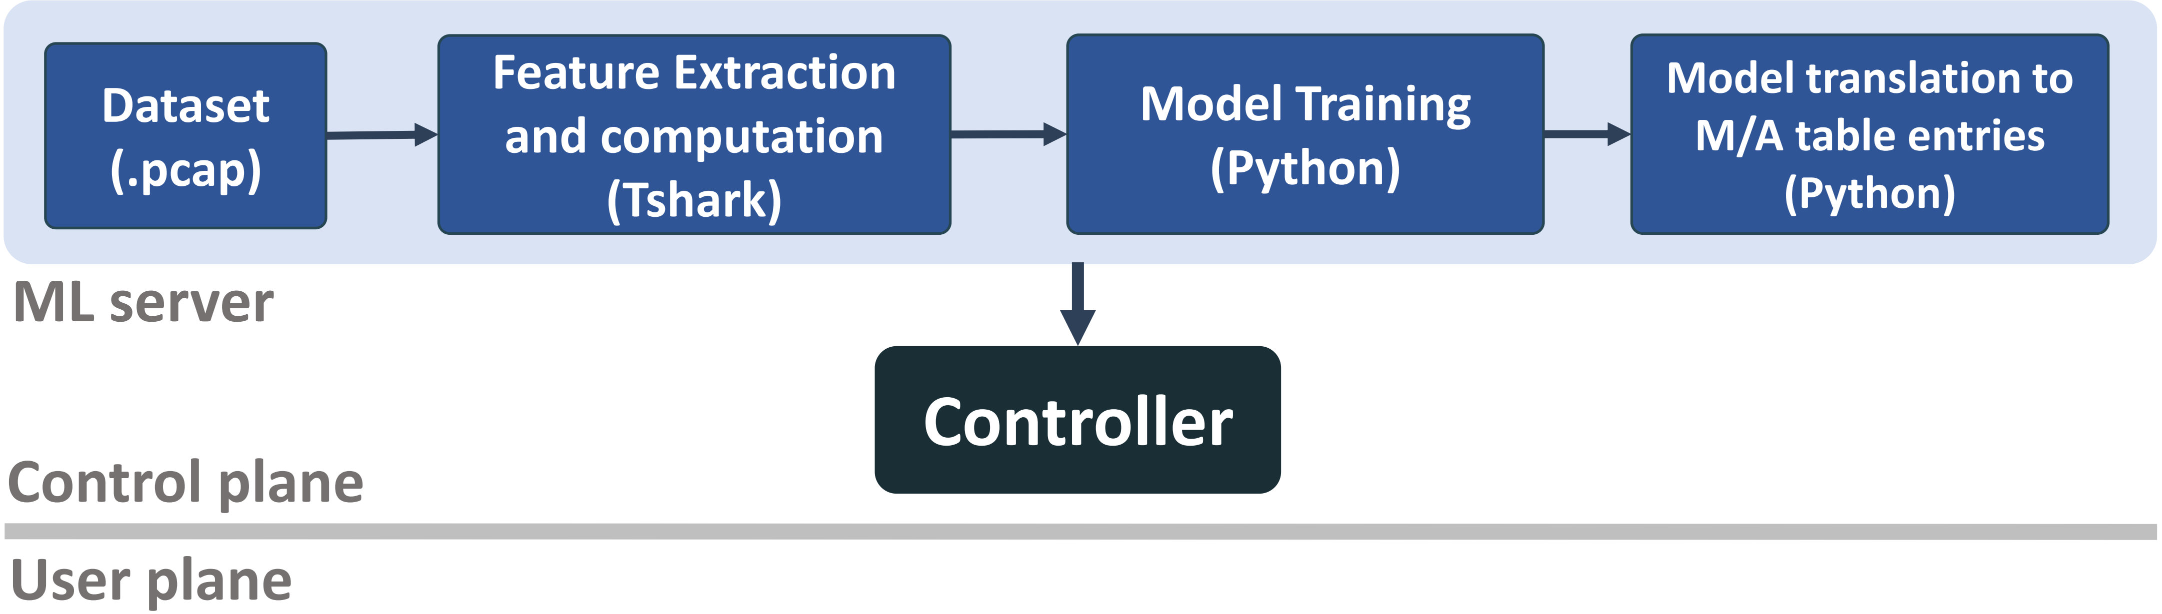

In-switch ML inference workflow

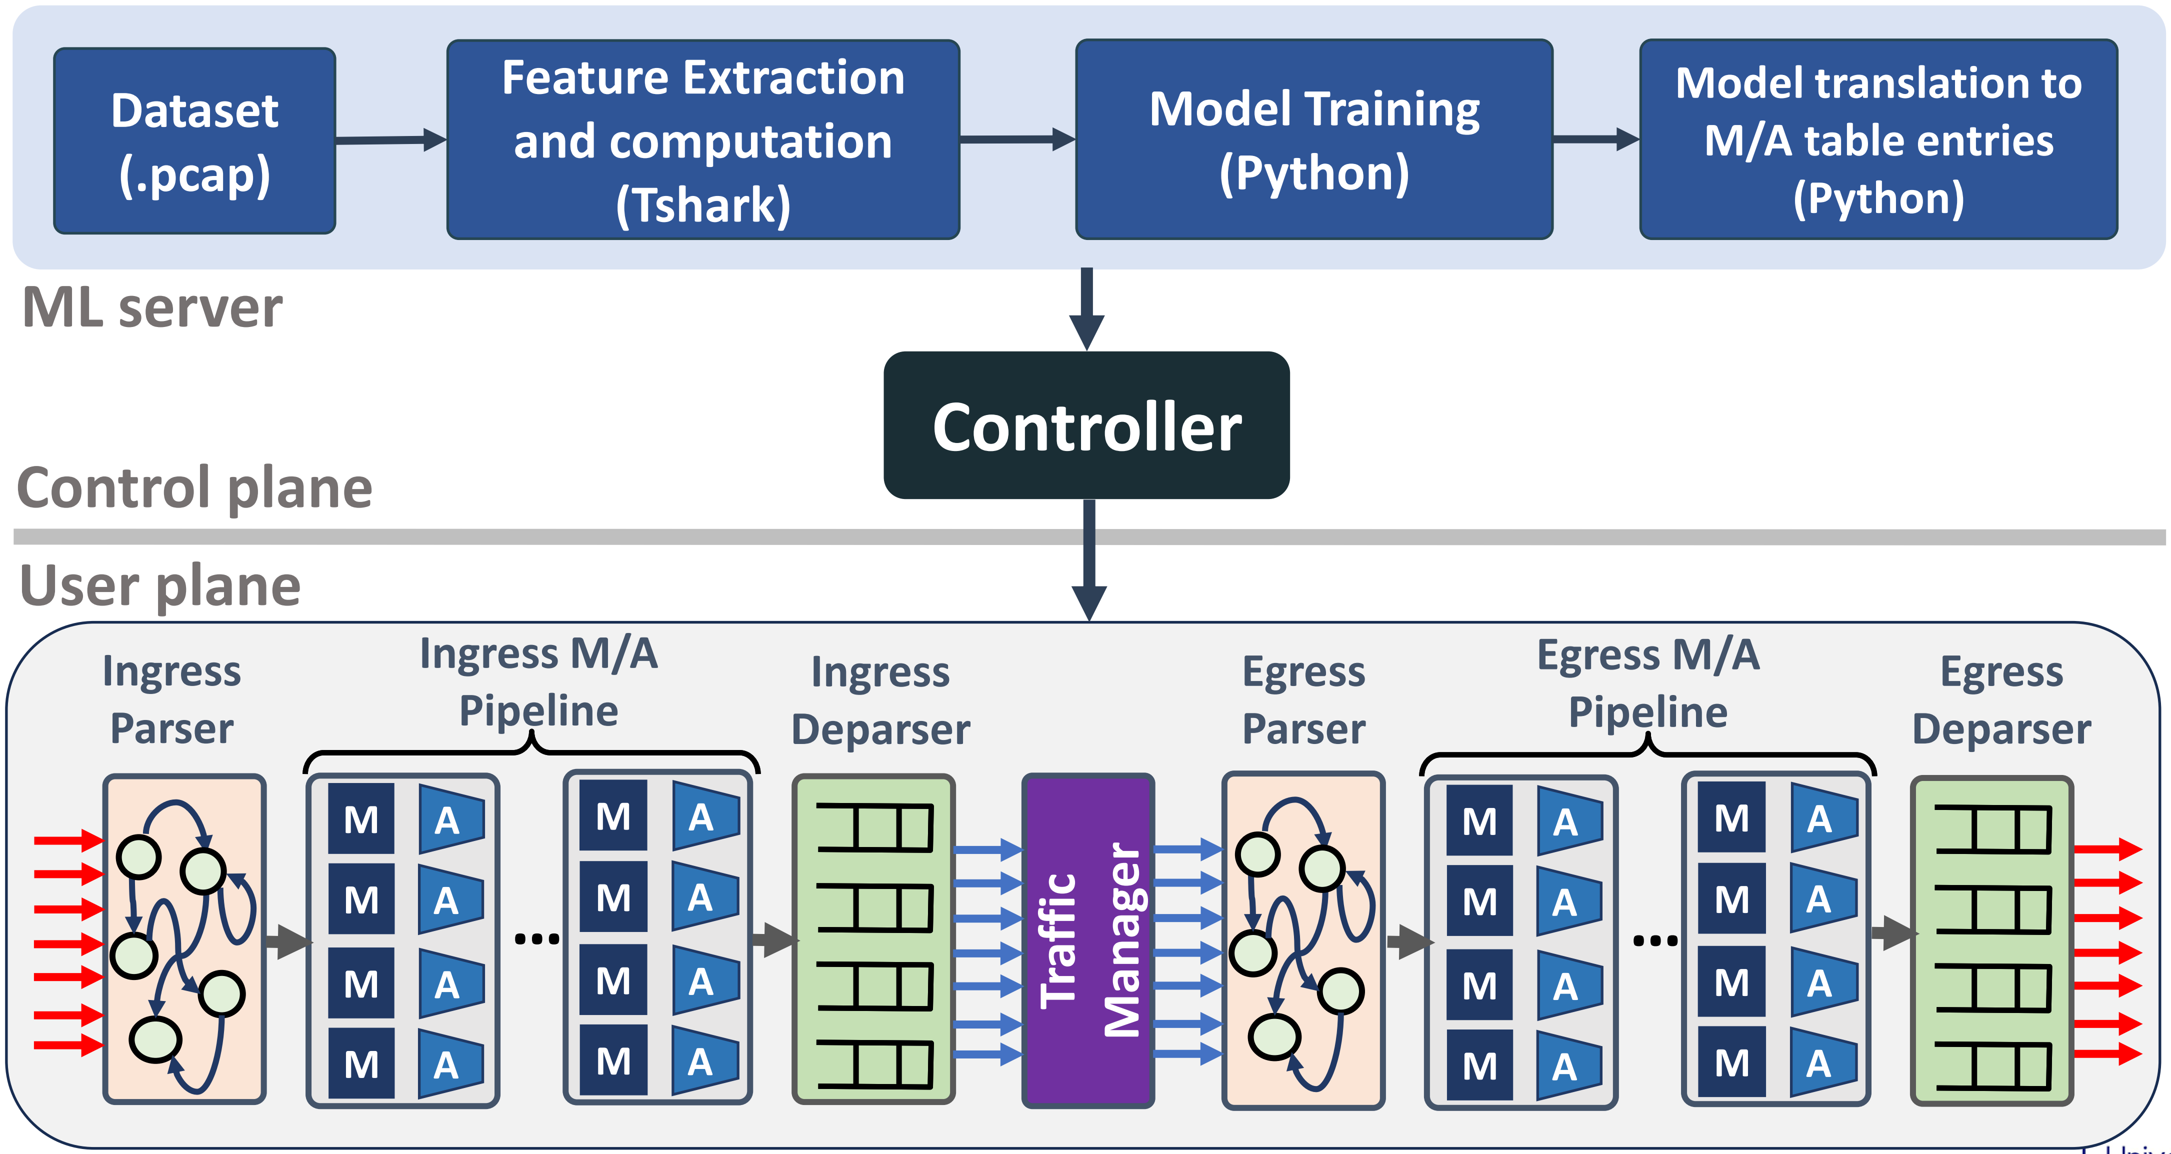

In-switch model implementation

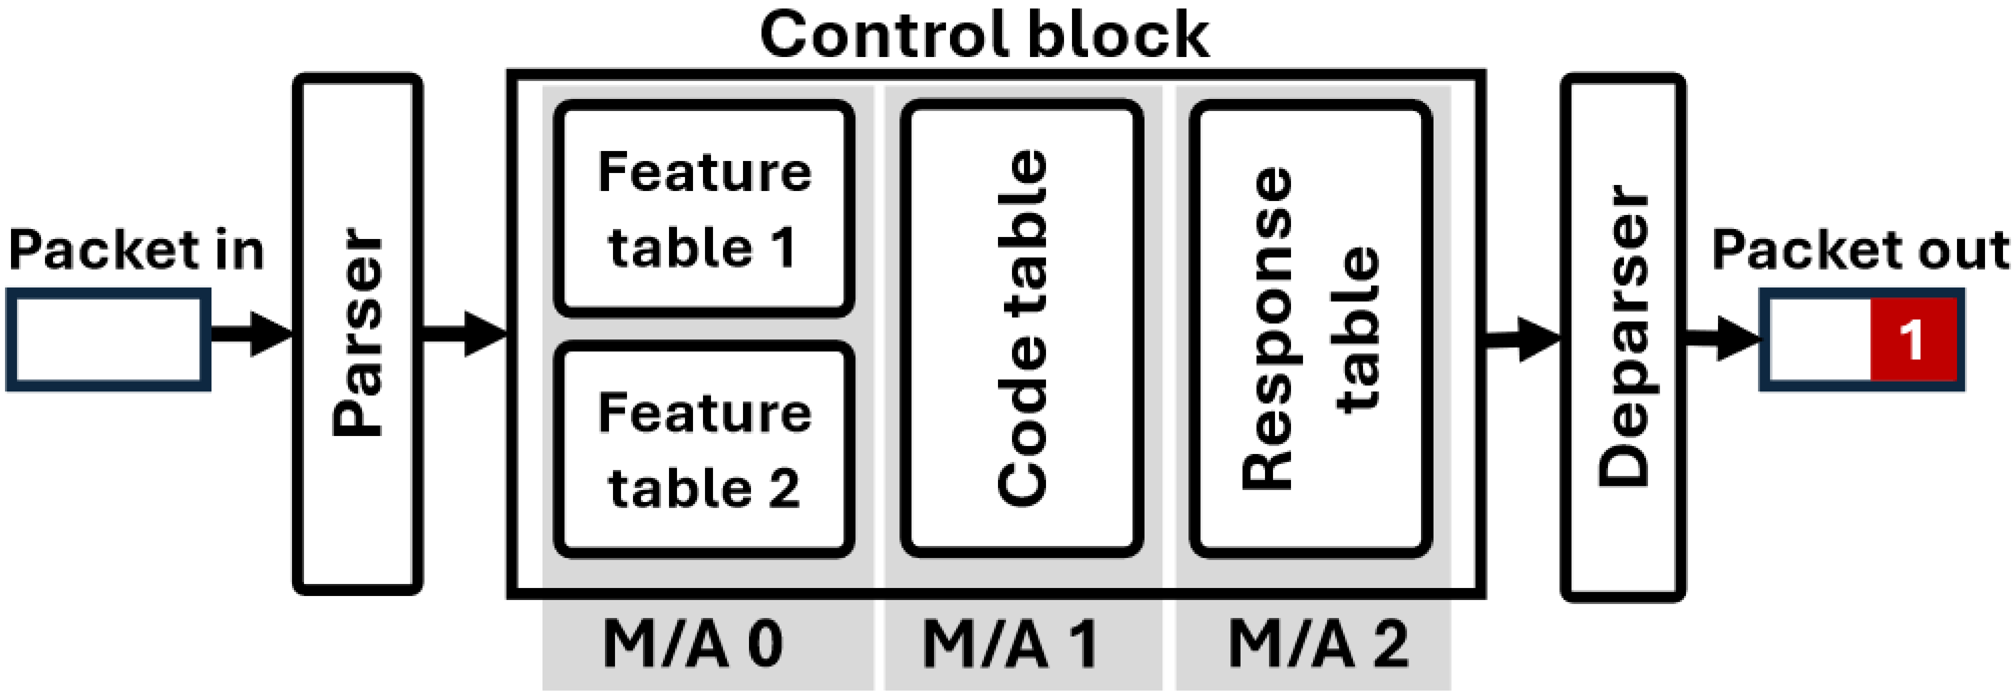

Use case: intrusion detection in a 5G test network

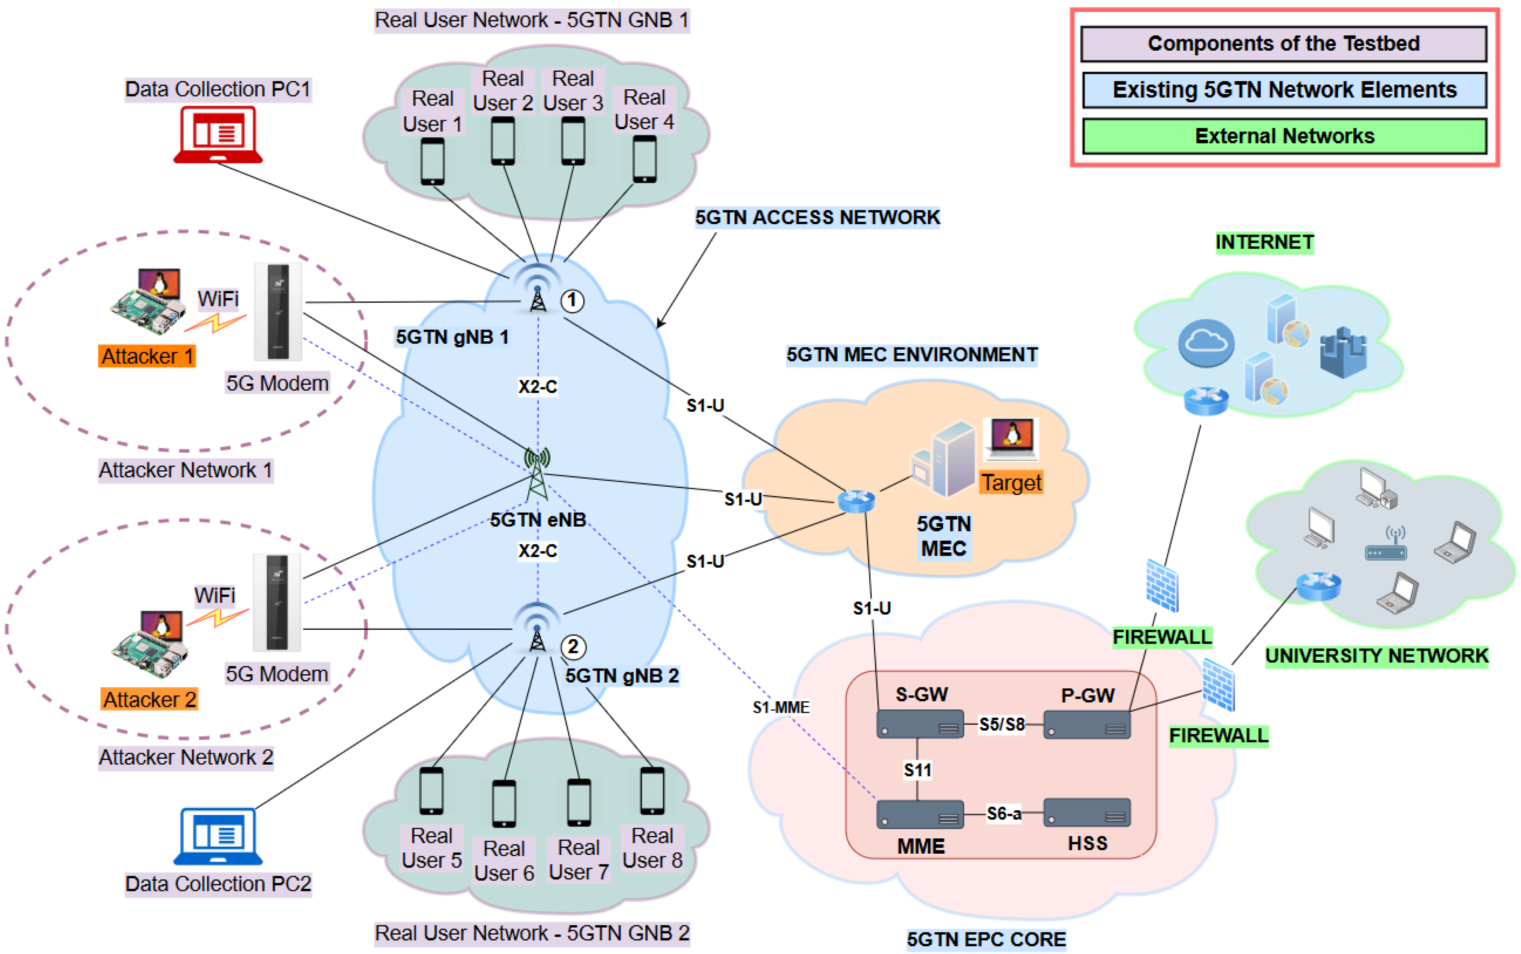

Dataset: 5G-NIDD

- Attackers target a server deployed in the 5GTN MEC
- **8 Attacks:** DoS and port scans
- **DoS attacks:** ICMP Flood, UDP Flood, SYN Flood, HTTP Flood, and Slowrate DoS
- **Port scans:** SYN Scan, TCP Connect Scan, and UDP Scan
- **Task:** Detect and separate the 8 attacks from benign traffic

Use case: intrusion detection in a 5G test network

***Model*:** Decision tree with maximum depth 37, using 10 features

***Features*:** packet length, TTL, destination port, source port, TCP window size, TCP Push flag, TCP header length, TCP Reset flag, TCP Fin flag, and UDP length

***Metrics*:** TPR, FPR, TNR, FNR, F1 Score; macro & weighted averages

Results – Classification accuracy

Class	F1 Score	TPR	FPR	TNR	FNR
Benign	99.993%	99.986%	0.000%	100.000%	0.014%
Slowrate DoS	90.507%	87.145%	0.663%	99.337%	12.855%
SYN Flood	100.000%	100.000%	0.000%	100.000%	0.000%
UDP Scan	99.727%	99.586%	0.000%	100.000%	0.414%
ICMP Flood	99.397%	100.000%	0.000%	100.000%	0.000%
TCP Connect Scan	99.589%	99.224%	0.000%	100.000%	0.776%
HTTP Flood	93.874%	96.292%	1.674%	98.326%	3.708%
SYN Scan	99.698%	99.891%	0.002%	99.998%	0.109%
UDP Flood	100.000%	100.000%	0.000%	100.000%	0.000%
Macro Avg	**98.087%**	**98.014%**	**0.260%**	**99.740%**	**1.986%**
Weighted Avg	**97.985%**	**97.998%**	**0.338%**	**99.662%**	**2.002%**

Use case: intrusion detection in a 5G test network

Results – Classification accuracy

Class	F1 Score	TPR	FPR	TNR	FNR
Benign	99.993%	99.986%	0.000%	100.000%	0.014%
Slowrate DoS	90.507%	87.145%	0.663%	99.337%	12.855%
SYN Flood	100.000%	100.000%	0.000%	100.000%	0.000%
UDP Scan	99.727%	99.586%	0.000%	100.000%	0.414%
ICMP Flood	99.397%	100.000%	0.000%	100.000%	0.000%
TCP Connect Scan	99.589%	99.224%	0.000%	100.000%	0.776%
HTTP Flood	93.874%	96.292%	1.674%	98.326%	3.708%
SYN Scan	99.698%	99.891%	0.002%	99.998%	0.109%
UDP Flood	100.000%	100.000%	0.000%	100.000%	0.000%
Macro Avg	**98.087%**	**98.014%**	**0.260%**	**99.740%**	**1.986%**
Weighted Avg	**97.985%**	**97.998%**	**0.338%**	**99.662%**	**2.002%**

Use case: intrusion detection in a 5G test network

Results – Resource consumption

SRAM	TCAM	Ternary Match Input Xbar	VLIW	Action Data Bus Bytes	Logical Table ID
1.40%	8.00%	11.50%	2.10%	4.60%	5.70%

Outline

- In-Network Machine Learning
- P4-based 5G UPF
- Intrusion detection in the UPF
- **Conclusion**

Conclusions & future work

- We built on successes in in-network ML and hardware accelerated 5G UPFs to enable high-speed network intrusion detection

- Evaluation on a 5G test network dataset shows how an in-switch model can achieve high classification accuracy with low resource usage

- Future work will focus on a full integration of ML inference into the UPF and an experimental evaluation in complete 5G network setting

Thank you!

aristide.akem@imdea.org

This work has been supported by:

Project **PCI2022-133013 (ECOMOME)**, funded by MICIU/AEI/10.13039/501100011033 and the European Union “NextGenerationEU”/PRTR.

Funded by the Horizon Europe programme of the European Union under Grant Agreement No. 101139270 (ORIGAMI)

Project **ORIGAMI**, funded by the Horizon Europe programme of the European Union, under Grant Agreement 101139270.