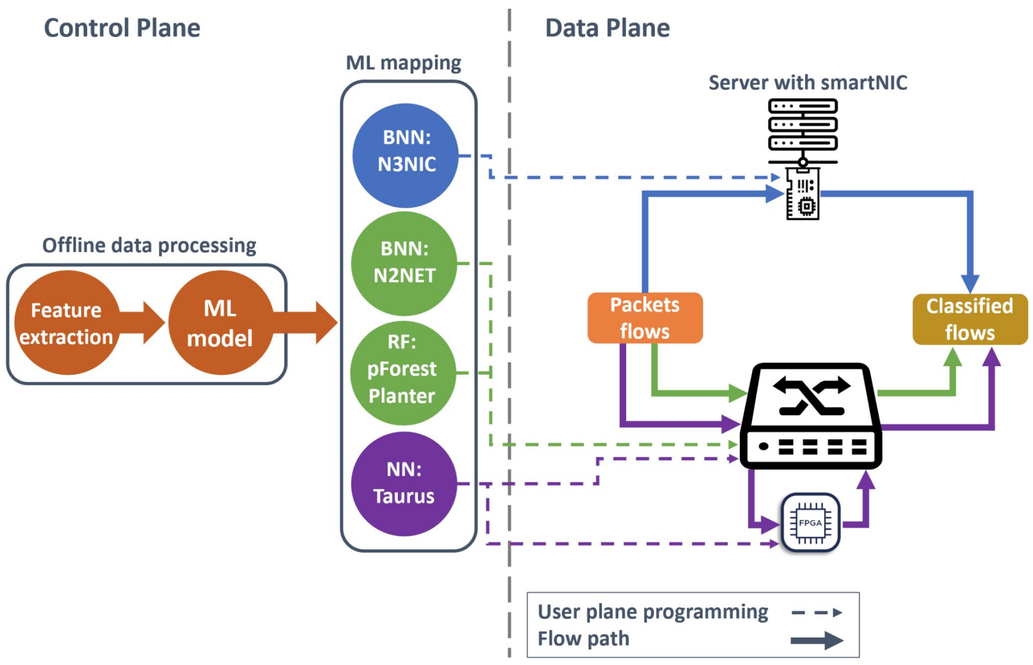
Implementation and Scalability Evaluation of Random Forests for In-Switch Inference

Aristide Akem[1,2], Michele Gucciardo[1] and Marco Fiore[1]
[1]IMDEA Networks Institute, [2]Universidad Carlos III de Madrid

Introduction

We comparatively evaluate state-of-the-art solutions for in-switch machine learning inference. We demonstrate that random forest (RF) models attain accuracies on par with those of approaches based on neural networks, which are also less amenable to in-switch operation. Next, we implement the top two solutions for in-switch random forest representation using a unified framework that we propose to ensure their fair comparison. We then verify their performance, resource consumption and scalability with respect to the capabilities of production-grade programmable switches.

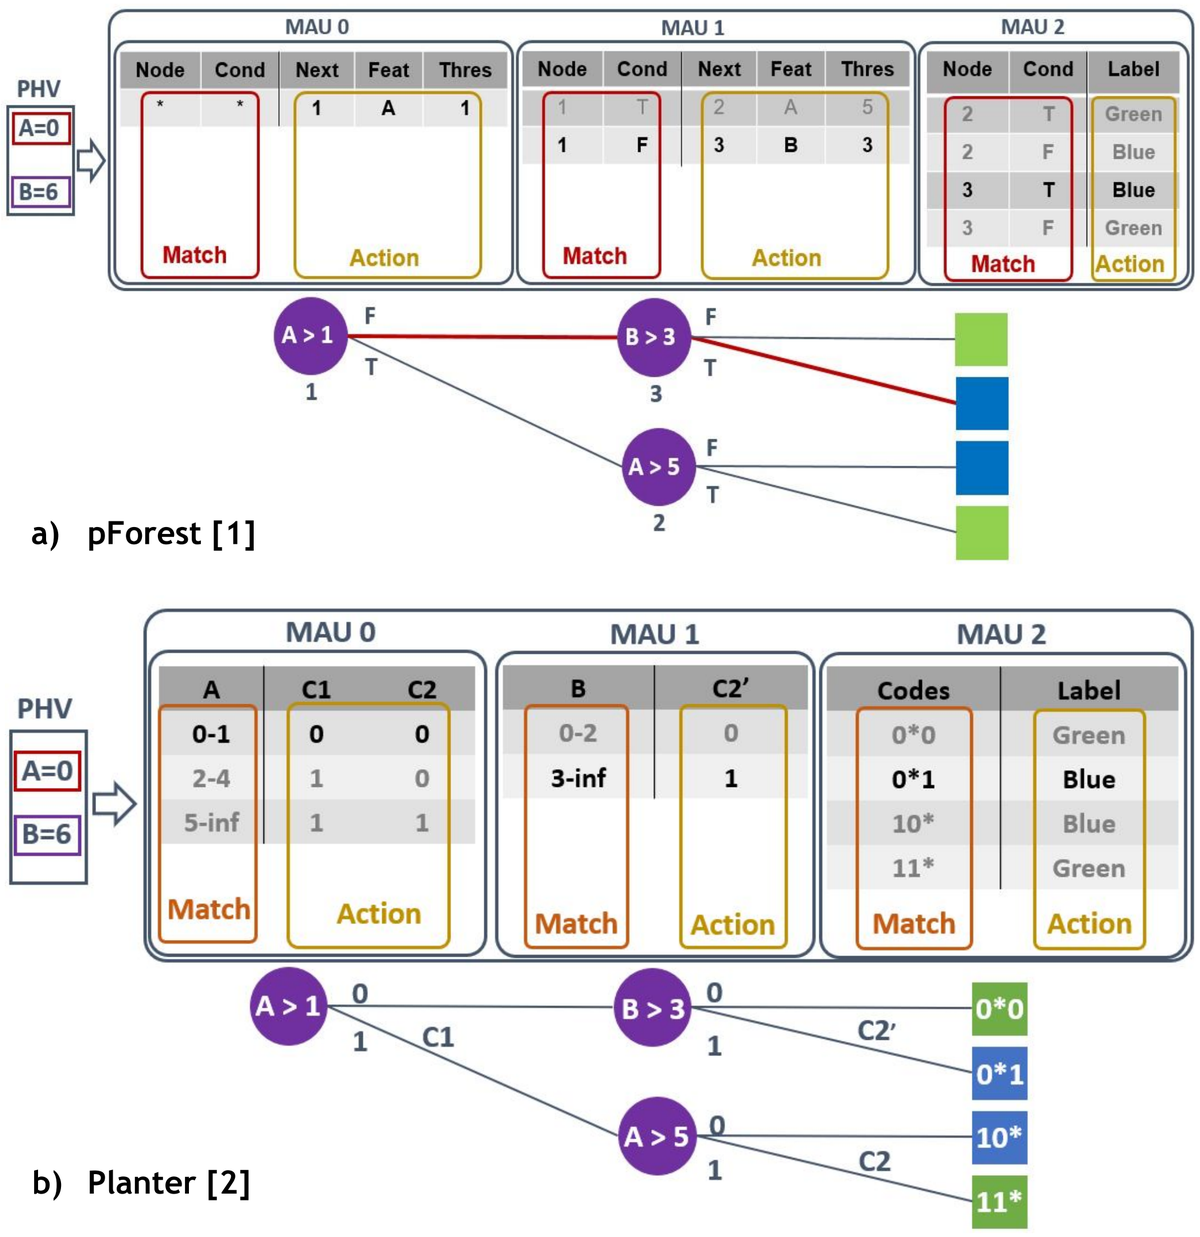
Proposed Approaches to In-Switch Inference



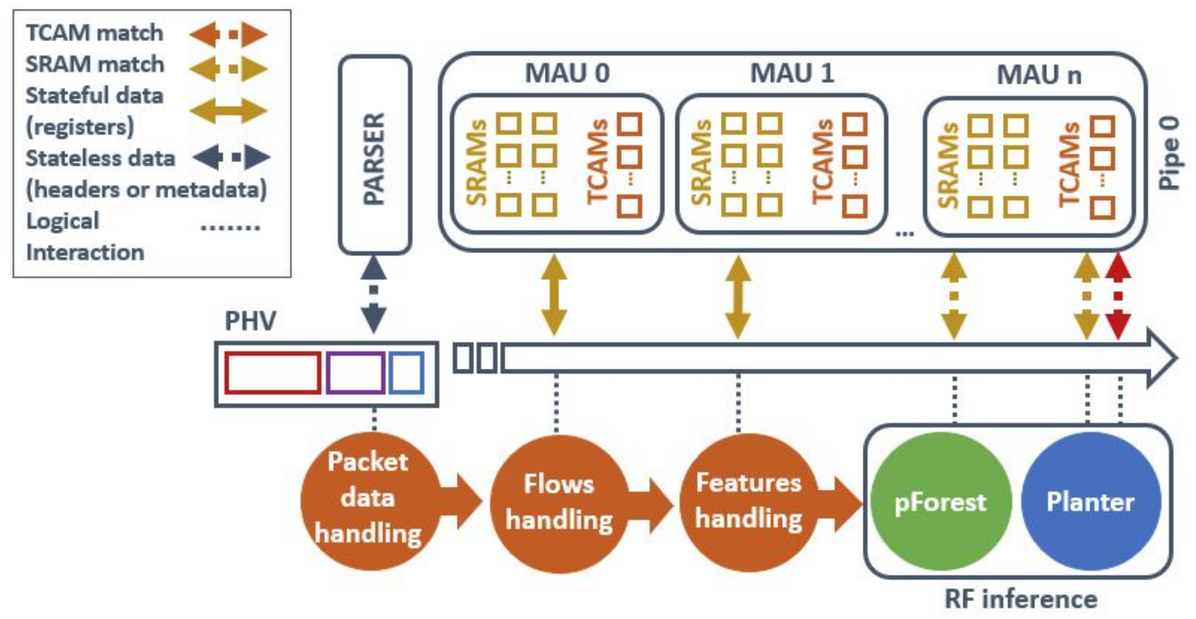
Comparison of Machine Learning Models

Use Case	Dataset	Model	Accuracy	Precision	Recall	F1-Score
Traffic Classification	UNIBS 2009	DT(9)	**99.737**	**96.345**	**97.406**	**96.851**
		RF (9,5)	99.521	95.077	95.107	95.078
		NN (64,32,8)	99.223	95.261	89.116	91.129
		BNN (64,32,8)	89.747	90.479	89.746	87.169
Anomaly Detection	CICIDS 2017	DT(10)	99.744	99.799	99.590	99.694
		RF (10, 5)	**99.762**	**99.805**	**99.628**	**99.716**
		NN(128,64,2)	99.259	99.263	99.259	99.260
		BNN(128,64,2)	89.807	87.305	92.622	88.802

Approaches to in-Switch RF Implementation



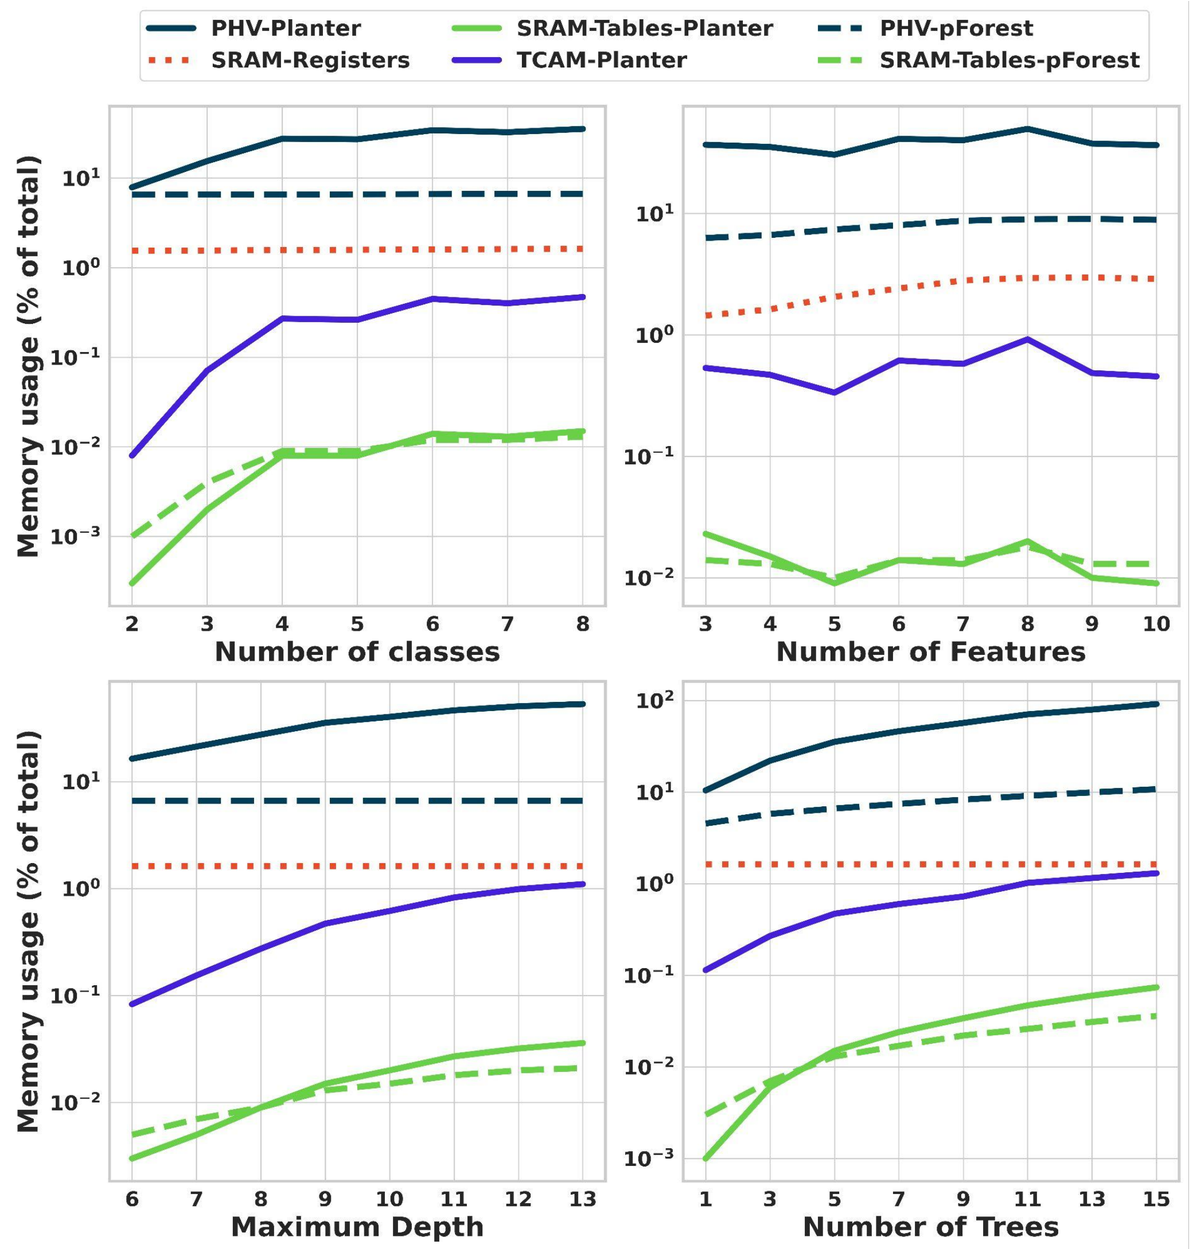
Unified Framework for RF Implementation



Results from In-Switch Inference with RFs

Use Case	Dataset	Approach	Accuracy	Precision	Recall	F1-Score
Traffic Classification	UNIBS 2009	RF (9,5) pForest	99.787	92.231	96.353	93.884
		RF (9,5) Planter	99.714	91.965	95.259	93.270
Anomaly Detection	CICIDS 2017	RF (4, 5) pForest	98.918	98.861	97.656	98.736
		RF (4, 5) Planter	98.965	98.870	97.797	98.791

Scalability Analysis



Conclusion

Our results show that although we can already obtain satisfactory results for simple use cases of traffic classification and anomaly detection with in-switch machine learning, there are critical aspects of specific resource usage which have to be further investigated and optimized. These constitute directions for future research on data plane inference.

References

[1] Coralie Busse-Grawitz *et al*. 2019. pForest: In-network inference with random forests. *arXiv:1909.05680*

[2] Changgang Zheng *et al*. 2021. Planter: Seeding trees within switches. *SIGCOMM '21 Demos and Posters*. pp. 12–14.

© Institute IMDEA Networks

This work has has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement no. 101017109 "DAEMON" and the Marie Skłodowska-Curie grant agreement no. 860239 "BANYAN"