

# POSMAC: Powering Up In-Network AR/CG Traffic Classification with Online Learning

Alireza Shirmarz, Fábio Luciano Verdi  
*Dept. of Computer Science*  
*Universidade Federal de São Carlos (UFSCar)*  
Sorocaba, Brazil  
ashirmarz@ufscar.br, verdi@ufscar.br

Suneet Kumar Singh, Christian Esteve Rothenberg  
*Dept. of Computer Engineering*  
*Universidade Estadual de Campinas (UNICAMP)*  
Campinas, Brazil  
ssingh@dca.fee.unicamp.br, chesteve@dca.fee.unicamp.br

**Abstract**—In this demonstration, we showcase POSMAC<sup>1</sup>, a platform designed to deploy Decision Tree (DT) and Random Forest (RF) models on the NVIDIA DOCA DPU, equipped with an ARM processor, for real-time network traffic classification. Developed specifically for Augmented Reality (AR) and Cloud Gaming (CG) traffic classification, POSMAC streamlines model evaluation, and generalization while optimizing throughput to closely match line rates.

**Index Terms**—ML model deployment, Online traffic classification, AR/CG, retraining platform.

## I. INTRODUCTION

Network traffic classification is crucial in applications like AR and CG, where encryption and dynamic port usage obscure traditional identifiers [1]. In our previous work, we utilized RTP features to achieve a 94.8% accuracy in classifying traffic [2]. To further enhance model accuracy, we developed POSMAC<sup>2</sup>, a platform that refines models by integrating new PCAP data and employing transfer learning techniques. This pre-deployment process, essential for avoiding the costs of deploying unoptimized models, outputs a robust and well-generalized model. POSMAC, leveraging NVIDIA DOCA on BlueField-3 DPU in a Docker environment, produces models that are efficient and cost-effective for deployment on BlueField-3 and other dataplanes, as supported by [3]–[6].

## II. ARCHITECTURE

POSMAC is a Docker-based platform consisting of four main components: (1) the Pcap Pool, (2) Traffic Classifier (TC), (3) Application Servers (APS), and (4) Online Trainer (OT), as illustrated in Fig 1. Each component is labeled and marked with green circles, connected by directed grey arrows indicating communication direction. Communication occurs at the datalink layer using MAC addresses, with dedicated networks ensuring exclusive connectivity between components (shown by the grey arrows). The online learning data flows, depicted by blue, red, and black dotted lines, indicate data transfer direction across the platform. Detailed descriptions of each component follow in subsequent sections.

<sup>1</sup>Platform of Optimization & Deployment of the Online Self-Trainer Model for AR/CG Traffic Classification.

<sup>2</sup><https://github.com/dcomp-leris/POSMAC>.

### A. PCAP Pool

The PCAP Pool is a key component of the POSMAC platform, simulating incoming network traffic to the server. It stores pre-collected PCAP files from various applications, with traffic categorized into three labels: ‘AR’, ‘CG’, and ‘other’, representing AR, CG, and other UDP-based applications like video conferencing and live streaming. After preprocessing with TCPReplay to adjust data rates, traffic is routed from the PCAP Pool to the TC component, where it is efficiently classified. The PCAP Pool can simulate traffic or be replaced by real user-generated traffic, enabling testing under diverse frame rates, resolutions, and network bandwidths.

### B. Traffic Classifier (TC)

The Traffic Classifier component within the POSMAC platform utilizes pre-trained models, such as DT or RF, initially deployed on the Data Processing Unit (DPU) BlueField 3.0, which features an ARM processor. These models are iteratively trained and improved using the OT block. This component receives network packets and routes them to the appropriate server based on the classified application type. It directly connects to three distinct application servers—each representing a specific class (AR, CG, or other applications)—to ensure targeted traffic handling. The TC comprises three key operations: **(1) Extracting Features (EXT)**: This function extracts relevant features from the packets. **(2) Network Traffic Classifier (NTC)**: Utilizes the RF or DT model to classify network flows using the extracted features. **(3) Forwarding Packets (FP)**: Routes packets to the corresponding application server based on their classified type. Additionally, this component maintains a direct link to the OT, periodically sending batched feature data to facilitate continuous model training and refinement.

### C. Application Servers (APS)

In the POSMAC platform, the Application Servers component deploys servers with heterogeneous computing resources, which may be physical, virtual, or containerized. Each server type connects directly to both the TC and OT components. The platform designates one container per application class—specifically AR, CG, and others. Each application server, such as the AR server, receives network flows and verifies their

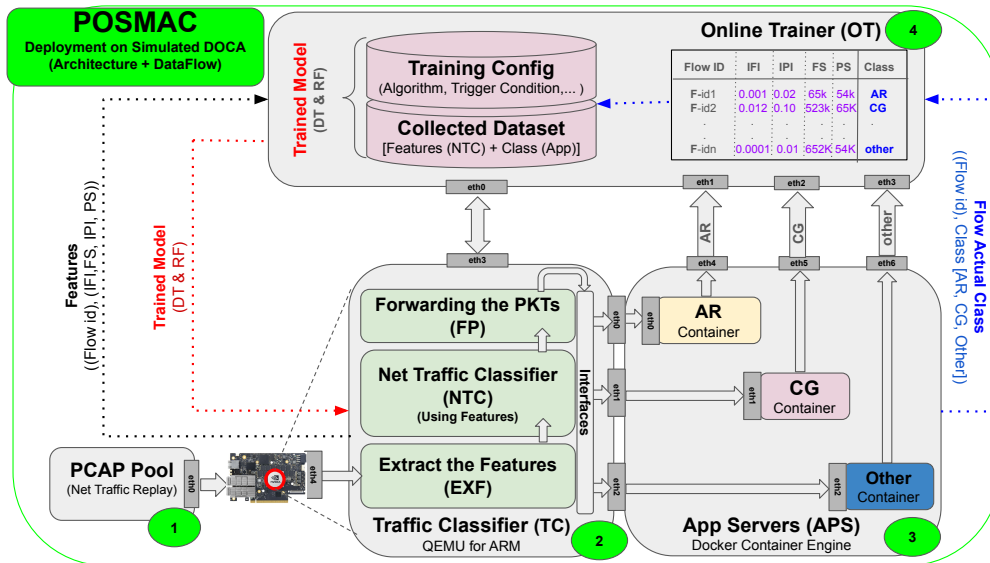


Fig. 1. POSMAC Architecture: Dotted lines illustrate online learning data flows — black for feature extraction, blue for class labeling, and red for transferring the retrained ML model.

correct routing based on the server IP configurations set in the PCAP Pool. It also labels the flows accordingly.  $Flow_{IDS}$ , which are unique identifiers, along with their corresponding labels (AR, CG, or other), are then transmitted to the OT. This connection to the OT is facilitated directly through a designated NIC, ensuring efficient communication and data handling within the platform.

#### D. Online Trainer (OT)

The OT is a crucial component that periodically trains the model, maintaining direct connections to both the TC and the Apps. It collects datasets from the TC’s EXT function, which includes Packet Size (PS), Inter Packet Interval (IPI), Frame Size (FS), and Inter Frame Interval (IFI), each tagged with a unique  $flow_{ID}$ . These datasets are further enriched with labels provided by the Application Servers, corresponding to each  $flow_{ID}$ . The OT is programmed with configurable rules that dictate the timing and conditions under which training occurs. Once the model is updated and enhanced with new data, it is serialized and transmitted back to the TC, where it is loaded and deployed according to predefined configuration rules. This cyclical process ensures that the classification model remains accurate and up-to-date with the latest network traffic patterns.

### III. CURRENT DEPLOYMENT & OUTPUT

Our initial deployment utilizes DT and RF models to classify network traffic into categories such as AR, CG, and others, based on metrics like PS, IPI, FS, and IFI. Using the DOCA simulator for BlueField 3.0 DPU, our models initially achieve 95.7% accuracy, which approaches nearly 100% through transfer learning as the model is retrained with new data. The improvement rate depends on the retraining strategy in OT. Our platform not only provides real-time reports on accuracy and resource usage but also enables detailed analysis of TC performance, including packet processing rates and system

requirements. The results are carefully logged to optimize classification accuracy and throughput. In the demo, we will showcase POSMAC components by analyzing the accuracy, throughput, and latency of the TC as it processes traffic from various applications using pre-collected PCAP files. Our next phase is to deploy these models on BlueField 3.0 hardware to validate and enhance performance in real-world conditions. The primary challenge in deploying online learning models on hardware is obtaining accurate, continuously updated class labels in real-world environments. However, the platform uses online learning to enhance and generalize models before hardware deployment, though this real-world loop poses the challenge of evolving labels.

#### ACKNOWLEDGMENT

This work was supported by Ericsson Telecomunicações Ltda., and by the Sao Paulo Research Foundation (FAPESP), grant 2021/00199-8, CPE SMARTNESS .

#### REFERENCES

- [1] A. Azab, M. Khasawneh, S. Alrabaae, K.-K. R. Choo, and M. Sarsour, “Network traffic classification: Techniques, datasets, and challenges,” *Digital Communications and Networks*, 2022.
- [2] A. Shirmarz, F. L. Verdi, S. K. Singh, and C. E. Rothenberg, “From pixels to packets: Traffic classification of augmented reality and cloud gaming,” in *2024 IEEE 10th International Conference on Network Softwarization (NetSoft)*, 2024, pp. 195–203.
- [3] S. U. Jafri, S. Rao, V. Shrivastav, and M. Tawarmalani, “Leo: Online ml-based traffic classification at multi-terabit line rate.” *NSDI*, 2024.
- [4] C. Zheng, Z. Xiong, T. T. Bui, S. Kaupmees, R. Bensoussane, A. Bernabeu, S. Vargaftik, Y. Ben-Itzhak, and N. Zilberman, “Iisy: Hybrid in-network classification using programmable switches,” *IEEE/ACM Transactions on Networking*, pp. 1–16, 2024.
- [5] A. T.-J. Akem, B. Büttin, M. Gucciardo, M. Fiore *et al.*, “Jewel: Resource-efficient joint packet and flow level inference in programmable switches,” in *IEEE ICC*, 2024.
- [6] T. Swamy, A. Zulfiqar, L. Nardi, M. Shahbaz, and K. Olukotun, “Homunculus: Auto-generating efficient data-plane ml pipelines for datacenter networks,” in *Proceedings of ACM ASPLOS*.