Advanced Threat Defense with In-Network Traffic Analysis for IoT Gateways

Mingyuan Zang[§], Changgang Zheng[†], Lars Dittmann[§], and Noa Zilberman[†]
[§]Technical University of Denmark, [†]University of Oxford
[§]{minza, ladit}@dtu.dk, [†]{name.surname}@eng.ox.ac.uk

Introduction IoT devices are widely and dynamically deployed in diverse use cases, leading to a surge of security threats that were previously overlooked. Timely mitigation services are entailed to prevent network-wide damage caused by fierce attacks, which is crucial to secure ultra-reliable low-latency communications (URLLC) in 5G. IoT gateways with both cellular and Ethernet interface connecting to devices are expected to provide first-line of defense. Traditional firewall functions are insufficient for analyzing traffic with multiple protocols (e.g. Modbus/messaging protocols). Machine learning (ML) aids in advanced traffic analysis, but efficient deployment remains a barrier. Cloud-based ML offers complex analysis in mobile networks but is limited in timely mitigation. Offloading ML inference to IoT gateways can reduce mitigation time, but is limited in providing flexible, low-overhead, and continuous traffic parsing and analysis to handle emerging threats. Programmable Data Planes (PDP) and in-network computing offer opportunities for fast and flexible in-network ML-based packet processing [3]. Nonetheless, it is challenging to apply them to IoT gateways with limited resources to cope with continuous learning on emerging threats for easy maintenance.

Our design We present P4Pir, an in-network traffic analysis solution integrated within IoT gateways. We utilize the flexibility of both data plane and control plane provided by PDP to enable continuous learning and consistent ML updates. By introducing run-time traffic parsing and shadow model updates at the IoT gateway, P4Pir achieves real-time multi-protocol data collection, in-network ML-based attack mitigation, and hitless runtime ML updates.

Figure 1 depicts P4Pir's workflow. Step ① shows a typical workflow of in-network ML-based detection [3], where a trained ML model (e.g. Decision Tree) is mapped to a set of table rules on PDP for inference to analyze incoming traffic. On top of it, P4Pir utilizes the programmability to parse features from multiple IoT protocol headers (e.g., messaging protocols, Modbus) in step ② A & B and identifies suspicious traffic from incoming traffic based on the table rules in step ② C. While benign traffic is being forwarded, suspicious traffic will be dropped immediately and logged to the control plane by encapsulating the extracted features in digests [2]. Based on digests, P4Pir uses Isolation Forest algorithm for unsupervised labeling and retrains the current model to learn

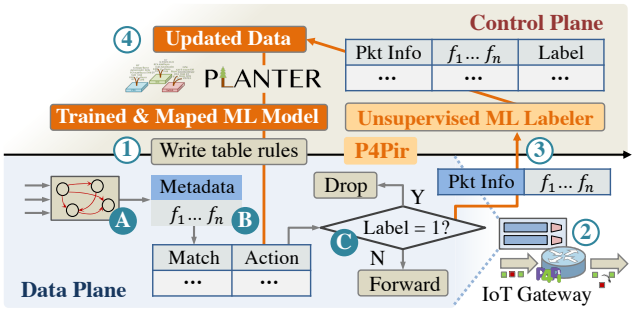

Figure 1: P4Pir workflow for IoT gateway.

from the labeled digests for new traffic patterns, shown in steps ③. A set of new table rules is then generated to map the new model's parameters, as in step ④. A shadow update scheme is proposed to seamlessly insert new table rules. With this workflow, P4Pir can learn from newly incoming traffic and mitigate abnormal traffic continuously.

Performance P4Pir prototype was developed on P4Pi [1], using Raspberry Pi 4 Model B with 8GB RAM, and bmv2 with v1model architecture. We replayed two public datasets *EDGE-IIOTSET* and *CIC-IDS2017* for model evaluation. Results demonstrate that P4Pir can scale and detect emerging attacks by retraining and updating in-network models (>30% accuracy enhancement). The proposed shadow update scheme has a negligible impact on network performance (e.g. only cause 14% throughput reduction and negligible jitter increase). P4Pir consumes insignificant CPU resources (e.g. 8% increment on CPU utilization rate).

Conclusion We presented P4Pir, an in-network ML-based analysis solution to defense emerging threats for IoT gateways. P4Pir can be deployed in 5G IoT gateway to timely mitigate malicious traffic in low-latency communications. *Future work* will focus on distributed deployment of P4Pir.

REFERENCES

[1] Sándor Laki, Radostin Stoyanov, Dávid Kis, Robert Soulé, Péter Vörös, and Noa Zilberman. 2021. P4Pi: P4 on Raspberry Pi for Networking Education. *SIGCOMM Comput. Commun. Rev.* 51, 3 (2021).
[2] Mingyuan Zang, Changgang Zheng, Radostin Stoyanov, Lars Dittmann, and Noa Zilberman. 2022. P4Pir: In-Network Analysis for Smart IoT Gateways. In *SIGCOMM '22 Poster and Demo Sessions*. ACM, 46–48.
[3] Changgang Zheng, Mingyuan Zang, Xinpeng Hong, Riyad Bensoussane, Shay Vargaftik, Yaniv Ben-Itzhak, and Noa Zilberman. 2022. Automating In-Network Machine Learning. arXiv:2205.08824 [cs.NI]