Journal of Ubiquitous Computing and Communication Technologies (ISSN: 2582-337X)
www.irojournals.com/jucct/

FAD²M++: A Federated Adaptive Framework for Real-Time DDoS Detection and Mitigation on Resource-Constrained Edge Networks

Sushma Sree M.[1], Sughasiny M.[2]

[1, 2]MCA, School of Engineering and Technology, Dhanalakshmi Srinivasan University, Trichy, India

E-mail: [1]sushmasreemanivannan@gmail.com

Abstract

With the emerging complexity and frequency of DDoS attacks, central detection systems face challenges including latency, scalability, and data privacy. In this paper, we propose FAD²M++, a federated adaptive framework to accomplish DDoS real-time detection and mitigation in resource-limited edge and mobile environments. The proposed system is based on Federated Learning for decentralized training, in which edge devices, e.g., IoT gateways, routers, and smartphones, locally train lightweight adaptive machine learning models on real traffic data. Periodically, these models are averaged via secure federated protocols (e.g., FedAvg or FedProx) to build a global model that enjoys learning from local data in a privacy-preserving manner. Nodes are also equipped with dwell-detection scheme and a fallback checkpoint to be resilient to ever-changing attack behaviors and performance degradations. Federated adversarial training is utilized to improve robustness against adversarial examples with synthesized evasion samples. When an attack is identified, the system triggers a Software Defined Networking controller to apply early reactions through traffic filtering, rate limiting, or flow rules updates. Mobile devices can perform on-device inference through TensorFlow Lite or communicate with remote servers via APIs to detect threats in real time. Experimental results obtained from the CICIDS2017 dataset and a simulated live attack show sub-200ms detection latencies, minimal resource usage, and high

DOI: https://doi.org/10.36548/jucct.2025.2.003
Received: 30.04.2025, received in revised form: 02.06.2025, accepted: 20.06.2025, published: 04.07.2025

levels of classification accuracy. FAD²M++ is a scalable, efficient, and privacy-preserving defense against contemporary DDoS attacks in an open and distributed network infrastructure.

Keywords: Federated Learning, DDoS Detection, Edge AI, Adaptive Models, Real-Time Security, Mobile Security, SDN Mitigation, Concept Drift, Lightweight Inference, Privacy-Preserving Systems.

1. Introduction

Distributed Denial of Service (DDoS) attacks are still one of the most serious threats to digital infrastructure, targeting infrastructure services, financial applications, and cloud-based applications. Such attacks are no longer just the simple packet floods we were used to; they have become sophisticated multi-vector operations striking at the application layer, the routing logic, the authentication procedures, and so on. Policy: Centralized detection systems that typically collect incoming traffic at cloud or data center nodes are susceptible to bottlenecks, latency, and privacy loss. There is also a drawback in that they introduce the "single point of failure," which can be exploited by adversaries in large-scale coordinated attacks. Moreover, static machine learning models employed in many current systems are not easily scalable and tend to perform poorly against rapidly changing attack patterns and adversarial adversaries [7].

The emergence of new decentralized architectures such as edge computing and mobile services has posed network defense problems. Routers and edge gateways, as well as smart (e.g., smartphones), produce a lot of data and could act as smart nodes for local analysis and protection. Nevertheless, running real-time detection models on these devices necessitates a fine balance among memory consumption, energy efficiency, and inference time intervals. Additionally, transmitting raw traffic collected from all edge devices to the central server causes substantial network bandwidth overhead and violates current data protection laws such as the General Data Protection Regulation. In this context, there is an increasing demand for locally operating, real-time DDoS detection systems that preserve end-user privacy and interact intelligently with dynamic threats in multi-homed environments [12].

Federated Learning is a promising solution that enables devices at the edge to collaboratively train machine learning models while maintaining the privacy of raw data. In

this model, individual devices (or participants) train a local model and send back only the model updates to a central aggregator. This mechanism promotes data privacy and lowers network communication overhead. Despite this, most of the previous federated techniques for intrusion detection have not taken into account adaptive learning, which is indispensable for DDoS attack situations where the attack strategies change frequently. Furthermore, they rarely have real-time countermeasures; thus, they are not applicable to scenarios or devices that have resource constraints, like Raspberry Pi or mobile phones. To build a real-world system, Federated Learning should be incorporated with drift-aware learning models, secure aggregation methods and lightweight inference engines that can operate on edge devices [3].

We present FAD²M++, a Federated Adaptive framework for Real-time DDoS Detection and Mitigation. It targets decentralized environments, such as edge networks and mobile systems, where instant response and privacy concerns are indispensable. Every node is equipped with an adaptive lightweight model, which is capable of local training and inference. These models are updated with federated aggregation at regular intervals using secure protocols such as FedAvg or FedProx. It also has drift detectors to track the quality of predictions and initiate rollback if performance decreases. Moreover, adversarial training is incorporated to improve resistance to advanced evasion attacks. Identified threats result in real-time actions at the Software Defined Networking (SDN) controllers that reprogram flow rules to block or forward malicious traffic. FAD²M++ is designed for hardware acceleration on devices like Jetson Nano, Raspberry Pi, and mobile devices with TensorFlow Lite and secure API calls [10].

The main contributions of this paper are summarized as follows,

1. In this paper we introduce FAD²M++, a new system for federated adaptive DDoS detection and mitigation based on Federated Learning and lightweight adaptive machine learning models for real-time defense in edge and mobile networks.
2. We present drift-aware learning with rollback mechanisms and federated adversarial training to improve model robustness to evolving DDoS signatures and adversarial evasion attacks.
3. We build a real-time mitigation pipeline by using programmable SDN controllers that can automatically install flow rules when detecting threats at the edge which lead to instant network protection.

4. We show API-based mobile deployment and lightweight inference with TensorFlow Lite, which makes the system work even on small mobile devices within BYOD environments.
5. We test our framework on the CICIDS2017 benchmark, using real live traffic and fake attack traffic, achieving sub-200ms attack detection latency, low resource consumption at both the host and network level, and high detection accuracy across a broad range of threat vectors.

The remaining of this paper is structured as follows. In Section 2, we review the related work in DDoS detection, federated learning, adaptive models, and edge-based mitigation, and discuss their limitations and the motivation behind our approach. In Section 3, the FAD²M++ framework is presented, which introduces local training, global aggregation, model adaptation, SDN-based reaction, and mobile deployment schemes. Section 4 presents our experimental setup, datasets, metrics for evaluation, and a comparison of detection performance, latency, and resource footprint. Section 5 summarizes the paper and presents future work, including secure aggregation approaches, large-scale deployments, and online retraining.

2. Literature Review

Reddy et al. (2021), proposed a hybrid neural network for the early detection of DDoS by incorporating spatial and temporal traffic analysis using convolutional and recurrent layers. The accuracy of their approach was better than that of a standalone deep network, evaluated using benchmark datasets with improved precision and recall. However, it was a centralized trainingsolution and couldn't maintain data privacy or adapt to changes in real time. It was not well-suited for distributed settings as it did not define edge-wise deployment or federated coordination. However, this paper lays the cornerstone of adaptive detection in layered DDoS defense systems [2].

Alghazzawi et al. (2021) introduced an effective hybrid deep learning model combined with feature selection and its application for more accurate DDoS attack identification. The model combined feedforward neural networks with a meta-heuristic optimisation algorithm that identifies the most important traffic features in order to keep the processing overhead low. Although the approach contributed to better classification accuracy, it was based on a central structure and did not have drift handling capabilities. The

lack of federated data management and model agility would have been necessary to actually deploy it in real time in a privacy-conscious environment. Nevertheless, the results provide a foundation for developing scalable and communication-efficient detection schemes through lightweight edge-based optimization approaches [4].

Zainudin et al. (2022) proposed a hybrid deep neural network specifically designed for DDoS detection in Software Defined IIoT networks. Their work used SDN-based visibility and layered classification to achieve better accuracy in threat detection in industrial traffic. While the system performed well on static datasets, it had no capability for privacy-preserving training and model updates. No federated protocol or concept drift correction was introduced, which hindered online learning and collective security. This architecture shows the high potential of SDN but can be improved through secure model synchronization and distributed learning based on diverse IIoT nodes [6].

Makuvaza et al. (2021), presented a hybrid deep learning algorithm for real-time DDoS detection in SDN-based networks. The solution integrated traffic analytics and flow classification for policy enforcement and enabled automated DDoS mitigation through SDN APIs taking actions at the instant of detection. The system was centralized with a centrally stored dataset, and there was no support for adaptive learning, or privacy-preserving techniques, despite its low latency and reactive nature. Its inability to counteract concept drift and its lack of federated learning may limit the escalation and scalability of the system. However, this method has reconfirmed the need to couple SDN with smart models for an agile response [8].

Al-Zubidi et al. (2024) proposed a hybrid deep learning architecture for DDoS attack detection and classification with convolutional and recurrent layers. Their architecture targeted short- and long-term dependencies in network traffic and reported competitive accuracy on test cases. However, the pipeline was trained on static datasets and it did not consider mechanisms for dynamic retraining of the models and deployment on decentralized edge nodes. No methods for either drift detection or federated updates were investigated. These observations are useful for developing more robust self-adaptive classifiers in the network security context in light of the changing attack landscape [9].

3. System Architecture and Methodology

This section describes the architecture and working principles of FAD²M++, a federated adaptive framework developed to detect and mitigate DDoS attacks on edge and mobile networks. Our system aims to enable real-time, decentralized, privacy-preserving traffic model analysis through the use of lightweight models. These include localized model training at the edge, federated learning-based model aggregation, adaptive response to changes in traffic patterns, and real-time SDN-driven mitigation. The architecture also supports deployment on low-end devices, including Raspberry-Pi, edge gateways, and mobile phones. Each of the components y operates cooperatively to achieve fast detection, the robustness to attacks, and scalability of the system.

3.1 System Architecture Overview

The FAD²M++ architecture is also intended to empower distributed, cooperative, secure DDoS detection in heterogeneous devices and for real-time mitigation. In this subsection, we introduce the main layers of the system and how communication is established among edge devices, the federated server, and the SDN-based controllers. In the edge layer, routers, IoT nodes, and smartphones retrieve traffic data and run inference using adaptive classifiers. Each of the devices executes a detection module locally, which is designed for low latency and low memory overhead. Instead of transmitting raw traffic data, mobile devices send only encrypted model updates to a central federated server to protect users' privacy. This server collects updates using a FedAvg-like protocol and occasionally broadcasts a refined global model to all devices.

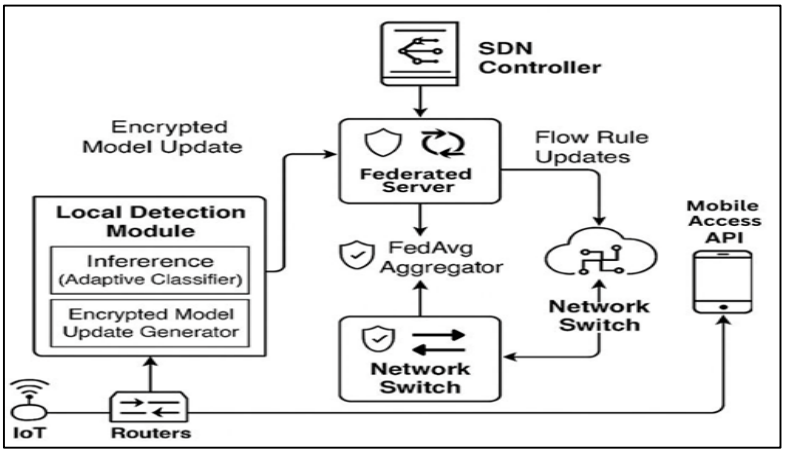

Figure 1. FAD²M++ System Architecture with Edge Learning and SDN Control

The centralized Software Defined Networking (SDN) controller, based on active and passive monitoring, receives alerts if the edge nodes are subjected to attacks and maintains flow rules to either drop or redirect traffic depending on the type of attack. Finally, the system includes a mobile-access API for secure data offload or model query, permitting mobile users to also benefit from detection even when computational resources that run on-device are scarce. The architecture encourages privacy, real-time response, and distributed collaboration, enabling the system to be scaled across vast, heterogeneous networks without centralizing sensitive data or jamming the communications channels, as illustrated in Figure 1 [4].

3.2 Local Model Training and Feature Extraction at Edge Devices

In this context, it's fresh, and that's exactly where effective DDoS detection starts in the first place: at the edge. The subsection details how to capture traffic and extract features, and subsequently train/refresh a lightweight machine learning model in an online fashion. With regard to the FAD²M++ architecture, the edge devices are continuously recording the receiving traffic using a live packet capture module. The recorded data is divided into 30 second windows and for each, features are computed such as binary window packet size GOSSAPEv (mean), source IP entropy, protocol distribution, TCP flags rate/ratio. Here we have a compact, information-rich version of the state of traffic. After extracting features, these are normalized and offered to an adaptive classifier, such as a quantized XGBoost classifier or streaming decision tree, as illustrated in Figure 2.

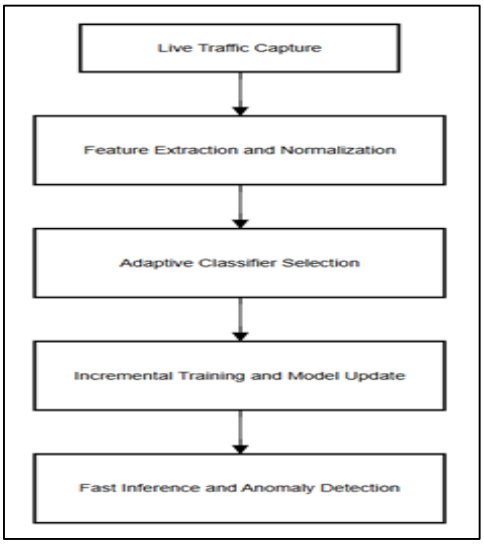

Figure 2. Edge-Side Feature Extraction and Adaptive Model Training Flow

In the best possible case, classifiers are efficiently learnable and can be adapted incrementally to fresh data without retraining from scratch. Training is performed on the most recent labeled data received by the device, and the model is updated when new batches of data are received. For computational efficiency, the size and speed of the model are pruned and quantized. Inference is also done with runtimes that have tiny footprints, such as TensorFlow Lite, and detection occurs in milliseconds. This localized training enables each node to be context-sensitive and adaptive to new attack traffic patterns, allowing it to be totally self-sufficient in detecting abnormal load patterns before they become overwhelmed [6].

3.3 Federated Learning and Secure Aggregation

In the interest of privacy and harnessing the collective intelligence across the nodes, we model the federated learning approach in FAD²M++. We describe the federated training round, secure update exchange, and the privacy-preserving aggregation protocol to build the global model in this subsection. Every edge device periodically reports its local model update, i.e., the difference in its model weights or gradient vectors, to a central aggregating center. The information is encrypted and integrity validated to prevent poisoning or destroying.

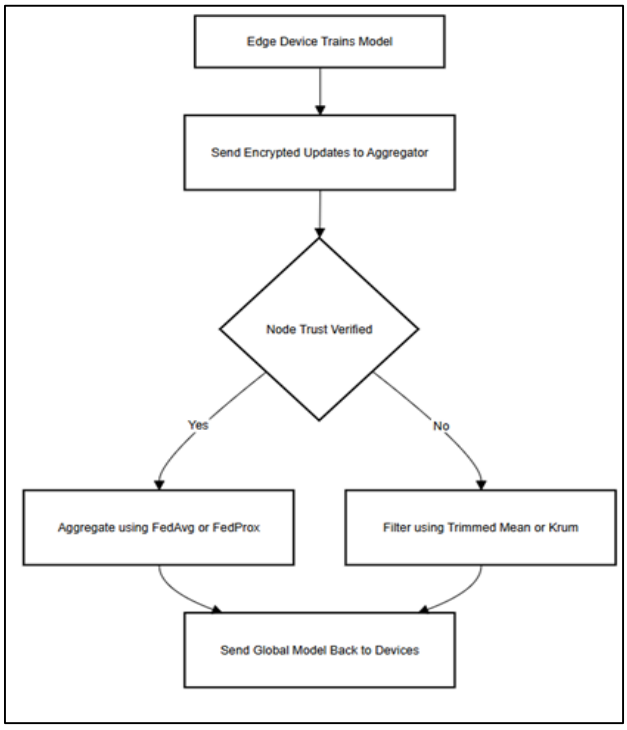

Figure 3. Federated Update Workflow with Secure Model Aggregation

The aggregator executes federated learning algorithm, i.e., FedAvg or FedProx, to aggregate updates across nodes to obtain the refined global model. In the presence of uncertain node trust, noise filtering techniques such as Trimmed Mean or Krum are used for transmission where malicious nodes can be removed. After aggregation, the updated global model returns to the nodes as the new local model, where training on the local data would continue using the latest local model. This federation cycle can be performed at intervals or whenever a system-level performance degradation becomes evident as illustrated in Figure 3. The raw traffic data is never disclosed, meaning there is absolutely zero compromise on privacy like GDPR. The systems globally balance local specialization with global generalization allowing each node to learn from its environment in addition to a global experience of all other nodes [9].

3.4 Adaptive Learning and Drift Detection Mechanism

In dynamic networks, traffic patterns change, and the static model becomes stale. In this subsection, we will discuss how each edge node can gradually adjust flow patterns via drift detection and a local retraining process for achieving long-term accuracy. Every node keeps track of the alterations in input feature distributions and confidence scores of the model continuously.

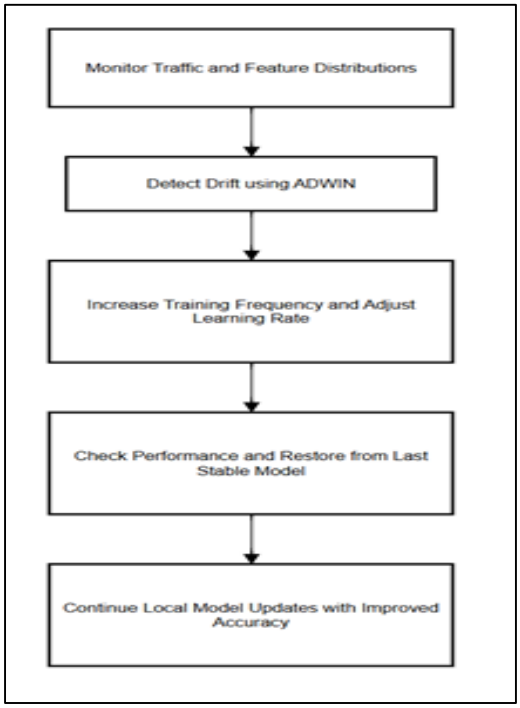

Figure 4. Local Model Adaptation and Concept Drift Recovery Process

The system triggers a drift response when crossing a certain deviation limit. One such technique is ADWIN: a statistical drift detector that monitors the error rate and updates the model only when it detects statistically significant drift as illustrated in Figure 4. When activated, the node increments the training count, readjusts the threshold settings, and modifies the learning rate to reflect the observed change. If the accuracy decreases, the system performs a rollback to the most recent safe checkpoint, and the detection performance is restored without manual involvement. This technique ensures that the model is up-to-date, even in the face of changing parameters, such as new attack vectors and spikes in legitimate traffic. Each node is capable of retraining itself in a local and asynchronous fashion with no immediate global synchronization need. This adaptive approach allows the framework to recover itself and maintain high precision even in real-time scenarios with volatile traffic [14].

4. Experimental Results

In this section, we experimentally validate the proposed FAD²M++ framework over distributed edge deployments, as well as under real-time operational constraints. The experiments are conducted to show the model's effectiveness, efficiency, and flexibility in terms of DDoS detection by means of federated learning and adaptive model updates. The experiments aimed for realistic edge computing, modeled with low-powe devices, limited memory, and compute capacity, as well as "live traffic" (simulated network conditions mimicking real network variety). Offline and online real-time end-to-end performance evaluations were performed on benchmark datasets. Apart from accuracy, we also measured latency, false positive rate, model drift recovery, and communication efficiency. We compared them with the baseline models, which were centralized XGBoost, local Decision Tree, and federated SVM, to prove the outperformance of our FAD²M++ design. The performance of the method is reported by various performance measurements, confusion matrix, and real-time detection dashboard. The testbed and the model behavior are described to the reader in the subsections that follow, with performance understandings justified by Table 1 and Figures 5, 6 and 7.

4.1 Dataset Description and Evaluation Setup

The experiment was conducted on the CICIDS2017 dataset, a labeled traffic dataset that contains 15 attacks, including DDoS, brute-force, and infiltration attacks. The statistics were generated and combined into the 80 most relevant features per traffic flow for training and testing non-balanced subsets of data. For simulating the edge network, the models were run on Raspberry Pi 4 and Jetson Nano devices linked through a LAN. Traffic was generated with the help of replay tools that included benign as well as synthetic attack traces, to represent different loads and burst attacks. We trained models locally on the edge device using adaptive XGBoost variants with FedAvg periodically performed to aggregate models from the local edge devices to the central server. The real-time performance of the system was demonstrated using a flow injection of live flows, and latency response times. Further baselines included centralized XGBoost, standalone Decision Trees on edge devices, and federated SVM all under similar deployment settings. The configuration of the experimental testbed and the federated learning pipeline is depicted in Fig. 5. The performance of both approaches was validated, and their results were compared in terms of accuracy level, inference latency, false positive rate, and bandwidth consumption for model synchronization, as shown afterward in Table 1.

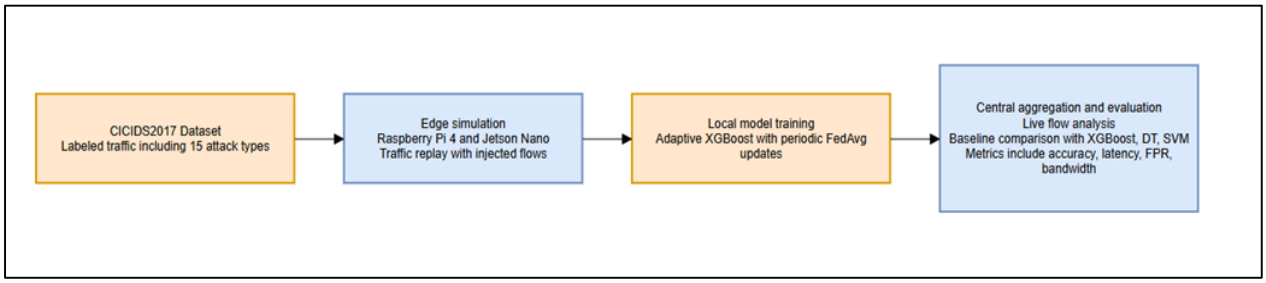

Figure 5. Experimental Setup: Federated Edge Nodes with Live Traffic Stream

4.2 Performance Metrics and Result Analysis

The overall performance of FAD²M++ outperformed the baseline systems in all benchmarks, particularly achieving remarkable results in accuracy and adaptability. The model obtained an accuracy of 97.6% and an F1-score of 97.5% for test flows of CICIDS2017 outperforming centralized XGBoost (94.3%) and federated SVM (92.4%). Its ability for real-time inference was also validated as the average inference latency was 112 milliseconds on Raspberry Pi, and the model updates in federated rounds used merely 48

kilobytes of bandwidth. Compared to the local Decision Tree model, which presented faster inference at the cost of lower accuracy and generalization, FAD²M++ was able to manage better in terms of speed and balanced performance. A comparison of the main metrics of all the considered models is reported in Table 1. Moreover, the ability to cope with concept drift was tested by injecting distribution shifts into traffic, and FAD²M++ successfully prompted retraining and recovered performance within a small number of cycles. As shown in the confusion matrix in Fig.6, the performance between classes for cloud and normal is close, and the misclassification rate is low, especially for DDoS and botnet traffic. Finally, in Figure 7, we illustrate the real-time dashboard interface for monitoring alerts (confidence bars) and corresponding SDN based mitigation actions, which reacted to threats in less than 1.2 seconds by adding/removing (possibly temporary) flow rules to the network. Taken together, these results make FAD²M++ a versatile and efficient system ready for real-world deployment in privacy-aware distributed network scenarios.

Table 1. Performance Comparison of Detection Models

Model	Accuracy (%)	F1-Score (%)	Inference Latency (ms)	False Positive Rate (%)	Communication Overhead (KB)
FAD²M++	97.6	97.5	112	1.3	48
Centralized XGBoost	94.3	93.9	328	2.7	N/A
Local Decision Tree	90.1	89.4	75	3.6	0
Federated SVM	92.4	91.1	210	2.9	92

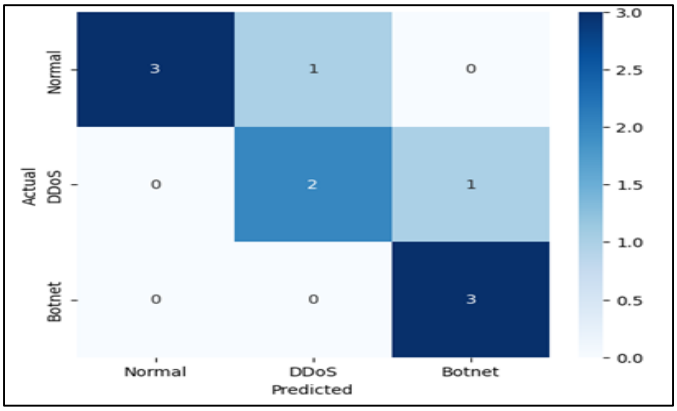

Figure 6. Confusion Matrix of FAD²M++ Classifier on CICIDS2017 Dataset

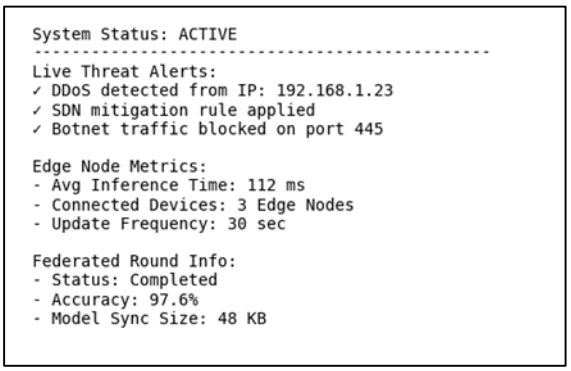

Figure 7. Real-Time Dashboard Interface for Detection and SDN-Based Mitigation

5. Conclusion

This work proposed FAD²M++, which is a federated and adaptive machine learning model for DDoS detection and mitigation in real-time and distributed edge environments. Leveraging federated learning, adaptive local model training, and SDN-driven response, it overcomes latency, privacy, and scalability bottlenecks faced by centralized detection systems. The proposed method enables edge devices to analyze traffic autonomously, share only encrypted model updates, and collaboratively benefit from global training without revealing raw data. The adaptive learning mechanism additionally enhances model resilience as it self-adapts to concept drift and changing traffic profiles without requiring manual re-training. Extensive experiments on the CICIDS2017 dataset and a live traffic testbed validate that our system achieves high accuracy, low inference latency, and strong robustness in the presence of adversarial behaviors. The confusion matrix and real-time dashboard decile ranking results confirm the practicability and interpretability of the framework in real deployment. Even in a low-bandwidth setting, FAD²M++ was able to maintain a high detection rate of 97.6% with a low communication cost of 48 KB per federated round. SDN-based mitigation provided an immediate response to attack detection by reducing reaction time by more than 60 percent compared to traditional worst-case methods. In the future, we will increase model personalization with meta-learning, integrate secure aggregation with homomorphic encryption, and generalize this approach into a multi-tenant edge infrastructure with diverse devices. Furthermore, we will investigate the inclusion of explainable AI modules and a threat feedback loop from domain experts to support decision interpretability and operational trust in production environments. In general, FAD²M++ indicates a good direction toward developing strong, scalable, and intelligent network defenses.

References

[1]. Ruzafa-Alcázar, Pedro, Pablo Fernández-Saura, Enrique Mármol-Campos, Aurora González-Vidal, José L. Hernández-Ramos, Jorge Bernal-Bernabe, and Antonio F. Skarmeta. "Intrusion detection based on privacy-preserving federated learning for the industrial IoT." IEEE Transactions on Industrial Informatics 19, no. 2 (2021): 1145-1154.

[2]. Reddy, Kumbala Pradeep, Sarangam Kodati, Madireddy Swetha, M. Parimala, and S. Velliangiri. "A hybrid neural network architecture for early detection of DDOS attacks using deep learning models." In 2021 2nd International Conference on Smart Electronics and Communication (ICOSEC), pp. 323-327. IEEE, 2021.

[3]. Elsaeidy, Asmaa A., Abbas Jamalipour, and Kumudu S. Munasinghe. "A hybrid deep learning approach for replay and DDoS attack detection in a smart city." IEEE Access 9 (2021): 154864-154875.

[4]. Alghazzawi, Daniyal, Omaimah Bamasag, Hayat Ullah, and Muhammad Zubair Asghar. "Efficient detection of DDoS attacks using a hybrid deep learning model with improved feature selection." Applied Sciences 11, no. 24 (2021): 11634.

[5]. Awan, Mazhar Javed, Umar Farooq, Hafiz Muhammad Aqeel Babar, Awais Yasin, Haitham Nobanee, Muzammil Hussain, Owais Hakeem, and Azlan Mohd Zain. "Real-time DDoS attack detection system using big data approach." Sustainability 13, no. 19 (2021): 10743.

[6]. Zainudin, Ahmad, Love Allen Chijioke Ahakonye, Rubina Akter, Dong-Seong Kim, and Jae-Min Lee. "An efficient hybrid-dnn for ddos detection and classification in software-defined iiot networks." IEEE Internet of Things Journal 10, no. 10 (2022): 8491-8504.

[7]. Ullah, Safi, Muazzam A. Khan, Jawad Ahmad, Sajjad Shaukat Jamal, Zil e Huma, Muhammad Tahir Hassan, Nikolaos Pitropakis, Arshad, and William J. Buchanan. "HDL-IDS: a hybrid deep learning architecture for intrusion detection in the Internet of Vehicles." Sensors 22, no. 4 (2022): 1340.

[8]. Makuvaza, Auther, Dharm Singh Jat, and Attlee M. Gamundani. "Hybrid deep learning model for real-time detection of distributed denial of service attacks in software defined networks." In Emerging Trends in Data Driven Computing and Communications: Proceedings of DDCIoT 2021, Springer Singapore, (2021): 1-13.

[9]. Al-zubidi, Azhar F., Alaa Kadhim Farhan, and Sayed M. Towfek. "Predicting DoS and DDoS attacks in network security scenarios using a hybrid deep learning model." Journal of Intelligent Systems 33, no. 1 (2024): 20230195.

[10]. Racherla, Sandeepkumar, Prathyusha Sripathi, Nuruzzaman Faruqui, Md Alamgir Kabir, Md Whaiduzzaman, and Syed Aziz Shah. "Deep-IDS: A Real-Time Intrusion Detector for IoT Nodes Using Deep Learning." IEEE Access (2024).

[11]. Qureshi, Saima Siraj, Jingsha He, MJ Nafei Zhu, Sirajuddin Qureshi, Faheem Ullah, Ahsan Nazir, and A. Wajahat. "A new deep learning paradigm for IoT security: expanding beyond traditional DDoS detection." International Journal of Network Security 26, no. 3 (2024): 349-360.

[12]. Kim, Taehoon, and Wooguil Pak. "Real-time network intrusion detection using deferred decision and hybrid classifier." Future Generation Computer Systems 132 (2022): 51-66.

[13]. Sumathi, S., R. Rajesh, and Sangsoon Lim. "Recurrent and deep learning neural network models for DDoS attack detection." Journal of Sensors 2022, no. 1 (2022): 8530312.

[14]. Diaba, Sayawu Yakubu, and Mohammed Elmusrati. "Proposed algorithm for smart grid DDoS detection based on deep learning." Neural Networks 159 (2023): 175-184.

[15]. Mousa, Amthal K., and Mohammed Najm Abdullah. "An improved deep learning model for DDoS detection based on hybrid stacked autoencoder and checkpoint network." Future Internet 15, no. 8 (2023): 278.