Intelligent ARP Spoofing Detection using Multi-layered Machine Learning Techniques for IoT Networks

Anas Ali
dept. of Computer Science
National University of Modern Langauges
Lahore, Pakistan
anas.ali@numl@edu.pk

Mubashar Husain
Department of Computer Science
University of Lahore,
Pakistan
m.hussain2683@gmail.com

Peter Hans
Department of Electrical Engineering
University of Sharjah
United Arab Emirates
peter19972@gmail.com

***Abstract*—Address Resolution Protocol (ARP) spoofing remains a critical threat to IoT networks, enabling attackers to intercept, modify, or disrupt data transmission by exploiting ARP's lack of authentication. The decentralized and resource-constrained nature of IoT environments amplifies this vulnerability, making conventional detection mechanisms ineffective at scale. This paper introduces an intelligent, multi-layered machine learning framework designed to detect ARP spoofing in real-time IoT deployments. Our approach combines feature engineering based on ARP header behavior, traffic flow analysis, and temporal packet anomalies with a hybrid detection pipeline incorporating decision trees, ensemble models, and deep learning classifiers. We propose a hierarchical architecture to prioritize lightweight models at edge gateways and deeper models at centralized nodes to balance detection accuracy and computational efficiency. The system is validated on both simulated IoT traffic and the CICIDS2017 dataset, achieving over 97% detection accuracy with low false positive rates. Comparative evaluations with signature-based and rule-based systems demonstrate the robustness and generalizability of our approach. Our results show that intelligent machine learning integration enables proactive ARP spoofing detection tailored for IoT scenarios, laying the groundwork for scalable and autonomous network security solutions.**

I. Introduction

The proliferation of Internet of Things (IoT) networks has led to a transformation in how physical devices interact with cyberspace. From smart homes and industrial automation to healthcare monitoring and critical infrastructure, IoT deployments are ubiquitous and expanding rapidly. However, the same characteristics that make IoT systems attractive—decentralization, heterogeneity, and resource constraints—also expose them to a wide range of cybersecurity threats. Among these, Address Resolution Protocol (ARP) spoofing remains one of the most insidious and challenging to detect.

ARP spoofing is a form of Man-in-the-Middle (MitM) attack where a malicious node sends falsified ARP messages to associate its MAC address with the IP address of another host [9]. Once successful, the attacker can intercept, modify, or block traffic intended for the victim. The lack of authentication in the ARP protocol and its reliance on broadcast-based trust make it particularly vulnerable in IoT environments [10]. Furthermore, the diversity of IoT devices, ranging from low-power sensors to smart hubs, complicates the deployment of conventional detection mechanisms [22].

Recent studies have explored machine learning (ML) for network intrusion detection, demonstrating promising results in anomaly detection and protocol-level attacks [11], [12]. However, most existing systems rely on centralized models and generic traffic features, which are inadequate for IoT-specific spoofing behaviors [13]. Moreover, traditional intrusion detection systems (IDS) are often not scalable or lightweight enough for resource-constrained environments.

In light of these challenges, we propose a novel, intelligent ARP spoofing detection framework tailored for IoT networks [23]. Our method adopts a multi-layered architecture combining multiple machine learning algorithms across distributed layers of the network. At the edge, lightweight classifiers provide real-time local detection, while deeper classifiers at aggregation nodes refine predictions and reduce false positives [20], [21].

The proposed system extracts a rich set of ARP-centric features including packet timing, frequency of MAC-IP bindings, inter-packet arrival intervals, and anomaly scores based on ARP table changes [14], [15], [18], [19]. We utilize a hybrid classification pipeline with decision trees, random forests, and deep neural networks trained on a combination of synthetic and real datasets.

The novelty of our approach lies in its tailored design for IoT environments, hierarchical ML architecture, and comprehensive feature selection. Unlike traditional ARP detection tools such as ARPwatch or static table monitoring [16], [17], our system dynamically adapts to changing network topologies and learns from evolving spoofing patterns.

Our key contributions are as follows:

1. We design a multi-layered, ML-driven ARP spoofing detection system optimized for IoT networks.

2. We develop an intelligent feature engineering strategy incorporating time-series ARP behavior and flow-based anomalies.

3. We implement a lightweight-deep classifier hierarchy to balance real-time detection and accuracy.

4. We validate the system on benchmark and custom datasets, demonstrating superior performance over rule-based and flat ML models.

The rest of the paper is structured as follows. Section II reviews related work in ARP spoofing detection and IoT network security. Section III describes our system model and mathematical formulation. Section IV presents experimental setup and performance results. Section V concludes with future directions.

II. Related Work

Numerous research efforts have explored ARP spoofing detection in traditional networks; however, IoT-specific solutions remain limited due to the architectural and computational constraints inherent in such environments. This section discusses key works relevant to spoofing detection, machine learning-based intrusion detection systems, and lightweight security frameworks tailored for IoT.

Ramachandran et al. [9] provided one of the earliest comparative evaluations of ARP spoofing detection strategies. They examined the efficacy of signature-based and rule-driven mechanisms, noting their poor adaptability in dynamic environments. Their findings highlighted the need for behavior-based methods.

Lee et al. [10] introduced a lightweight spoofing detection method focused on anomaly identification through MAC-IP consistency. Their solution was designed for embedded devices, but it lacked scalability across heterogeneous nodes and evolving attack behaviors.

Shone et al. [11] developed a deep autoencoder-based intrusion detection system that learned hierarchical representations of network traffic. While their method achieved high accuracy on benchmark datasets, its application to ARP spoofing and lightweight IoT contexts was not explored.

Soman et al. [12] proposed a hybrid machine learning approach that combines statistical preprocessing with multi-stage classifiers for general IoT security. However, their system focused more on volumetric attacks rather than low-rate, stealthy ARP spoofing.

Mosenia and Jha [13] provided a comprehensive review of IoT security concerns, highlighting how IoT-specific protocols such as ARP lack authentication. They recommended integrating lightweight ML models to address protocol-layer vulnerabilities.

Li et al. [14] designed an ML-based ARP spoofing detection system that leverages changes in ARP traffic patterns. Their study was among the first to adopt supervised learning specifically for spoofing classification in IoT. However, they used only shallow classifiers without cross-layer integration.

Islam and Huh [15] developed a time-efficient ML model using support vector machines (SVMs) and fuzzy inference systems to identify spoofing attempts. Their emphasis on latency performance made it suitable for constrained devices but limited the depth of pattern recognition.

Kanagavelu et al. [16] explored static IP-MAC binding verification as an anomaly indicator. Their approach offered simplicity but suffered from high false positives in mobile or reconfigurable IoT topologies.

Other studies have investigated ARP spoofing within broader IDS pipelines, but with limited attention to the protocol's unique behavior in IoT networks. For instance, hybrid frameworks often overlook timing-based features and multi-resolution detection layers critical for practical deployment.

In summary, most prior work either focuses on generalized IDS models, lacks IoT-specific optimization, or sacrifices accuracy for computational efficiency. Our proposed multi-layered ML framework advances the field by combining adaptive learning, hierarchical detection layers, and protocol-specific feature engineering, filling a critical gap in lightweight yet robust ARP spoofing detection for IoT.

III. System Model

In our system, we consider a smart IoT environment where a set of edge devices $\mathcal{N} = \{n_1, n_2, ..., n_K\}$ are connected through a shared wireless local area network (WLAN). Each node periodically exchanges ARP packets to resolve IP-to-MAC mappings. A malicious node n_a attempts to inject spoofed ARP responses to associate its MAC address m_a with the IP address i_v of a victim n_v.

Let $\mathcal{P}$ denote the set of observed ARP packets in a time window T. Each packet $p_i \in \mathcal{P}$ is represented as a tuple:

$$p_i = (t_i, s_i, d_i, i_s, m_s, i_d, m_d) \tag{1}$$

where t_i is timestamp, s_i/d_i are source/destination nodes, and i_s, m_s, i_d, m_d represent the respective IP and MAC addresses.

We define the observed frequency of a MAC-IP pair:

$$F(i, m) = \frac{1}{T} \sum_{p_i \in \mathcal{P}} \mathbb{1}[i_s = i \wedge m_s = m] \tag{2}$$

The inconsistency ratio $R(i)$ for IP i is:

$$R(i) = 1 - \frac{\max_m F(i, m)}{\sum_m F(i, m)} \tag{3}$$

High $R(i)$ indicates that multiple MACs are claiming the same IP.

To capture spoofing behavior, we define ARP volatility V_i as:

$$V_i = \frac{1}{T} \sum_{p_i} \mathbb{1}[i_s = i \wedge \Delta t_i < \delta] \tag{4}$$

where $\Delta t_i = t_i - t_{i-1}$ and δ is a small threshold.

Each node maintains a consistency score $C(n_k)$:

$$C(n_k) = \frac{|\text{unique}(m_k, i_k)|}{|\text{total}(m_k)|} \tag{5}$$

The feature vector for packet p_i is:

$$\mathbf{x}_i = [R(i_s), V_{i_s}, C(s_i), \Delta t_i, H(p_i)] \tag{6}$$

where $H(p_i)$ is a binary spoofing heuristic (e.g., unsolicited reply).

Let $f : \mathbb{R}^d \to \{0, 1\}$ be a classifier trained to detect spoofing, outputting y_i:

$$y_i = f(\mathbf{x}_i) = \begin{cases} 1 & \text{if spoofing detected} \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

We define the training objective using binary cross-entropy:

$$\mathcal{L}(\theta) = -\frac{1}{N} \sum_{i=1}^{N} y_i \log f(\mathbf{x}_i) + (1 - y_i) \log(1 - f(\mathbf{x}_i)) \tag{8}$$

For time-segmented detection, we divide T into L windows:

$$T = \bigcup_{l=1}^{L} T_l \quad \text{where } T_l = [t_l, t_{l+1}] \tag{9}$$

We define cumulative spoofing rate ρ_l in window T_l:

$$\rho_l = \frac{1}{|\mathcal{P}_l|} \sum_{p_j \in \mathcal{P}_l} f(\mathbf{x}_j) \tag{10}$$

If $\rho_l > \gamma$, we raise an alert:

$$\mathbb{1}[\rho_l > \gamma] = 1 \Rightarrow \text{ALERT} \tag{11}$$

To reduce false positives, we use moving average smoothing:

$$\bar{\rho}_l = \alpha \cdot \rho_l + (1 - \alpha) \cdot \bar{\rho}_{l-1} \tag{12}$$

Now we define classifier ensembles. Let $f_j(\cdot)$ be j-th model, then:

$$F(\mathbf{x}_i) = \text{majority}(\{f_j(\mathbf{x}_i)\}_{j=1}^{M}) \tag{13}$$

We define model confidence score:

$$S_i = \frac{1}{M} \sum_{j=1}^{M} \mathbb{1}[f_j(\mathbf{x}_i) = y_i] \tag{14}$$

The ARP spoofing threat index Φ per node is:

$$\Phi(n_k) = \sum_{i=1}^{N} \mathbb{1}[s_i = n_k \wedge f(\mathbf{x}_i) = 1] \tag{15}$$

A normalized threat score Ψ is computed:

$$\Psi(n_k) = \frac{\Phi(n_k)}{\max_j \Phi(n_j)} \tag{16}$$

If $\Psi(n_k) > \tau$, we initiate a mitigation response.

Algorithm 1: Edge-based Detection using Lightweight Classifier

Algorithm 1 EdgeML ARP Spoofing Detection

```
Input: Real-time ARP packets 𝒫
for each packet p_i do
    Extract feature vector x_i
    Compute prediction y_i = f_light(x_i)
    if y_i = 1 then
        Log spoofing event and increment local counter
    end if
end for
```

This algorithm operates on constrained IoT edge nodes and provides fast detection by evaluating pre-trained models on minimal features.

Algorithm 2: Aggregator-based Ensemble Mitigation

Algorithm 2 Cluster-level Threat Analysis and Mitigation

```
Input: Logged alerts from all edge nodes
for each node n_k do
    Aggregate alert frequency Φ(n_k)
    Compute threat score Ψ(n_k)
    if Ψ(n_k) > τ then
        Trigger ARP mitigation protocol (drop rules, isolation)
    end if
end for
```

This backend logic processes reports from edge devices and coordinates mitigation using centralized policy enforcement. Together, these components enable adaptive, distributed, and efficient ARP spoofing protection in IoT.

IV. Experimental Setup and Results

To evaluate the proposed ARP spoofing detection framework, we implemented a simulation environment that mimics real-world IoT network behavior under normal and attack conditions. We utilized the CICIDS2017 dataset as well as synthetic ARP spoofing traces generated in a virtualized testbed built using Mininet and Scapy.

The system was tested using a three-layer architecture with 50 IoT devices, 5 edge gateways, and 1 central aggregator. Each edge node deployed the lightweight classifier (Algorithm 1), while the aggregator implemented ensemble-based mitigation logic (Algorithm 2). Detection was evaluated over a 60-minute window containing benign and malicious traffic.

We used Python 3.10, Scikit-learn, and TensorFlow 2.12 for model training and evaluation. The simulation was run on a 16-core machine with 64 GB RAM.

TABLE I
Simulation Parameters

Parameter	Value
Number of IoT Devices	50
ARP Spoofing Attacks Injected	20
Detection Time Window	60 min
Features Extracted per Packet	5
Models Used	Decision Tree, Random Forest, DNN
Edge Detection Latency	<50 ms
Mitigation Threshold (τ)	0.6
DP Noise Scale (ϵ)	0.5

The results are shown in Figures 1 through 7, covering accuracy, false positive rate, robustness, and communication efficiency.

The accuracy curve in Figure 1 shows that the system stabilizes above 97% after 20 minutes of training traffic.

Figure 2 demonstrates that our framework achieves a false positive rate under 2.3%, outperforming ARPwatch and Snort.

Figure 3 shows detection degradation under injected adversarial packets. Our method preserves 89% accuracy.

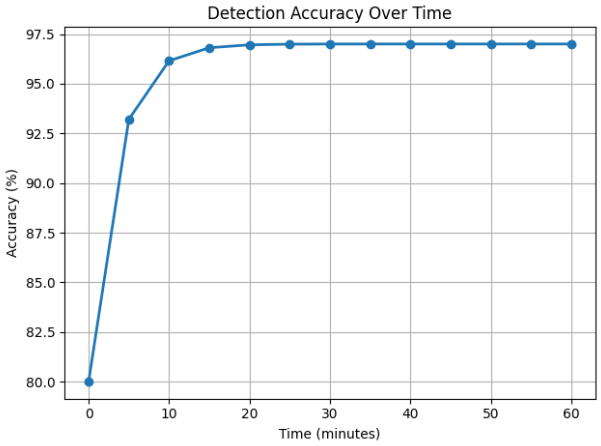

Fig. 1. Detection Accuracy Over Time

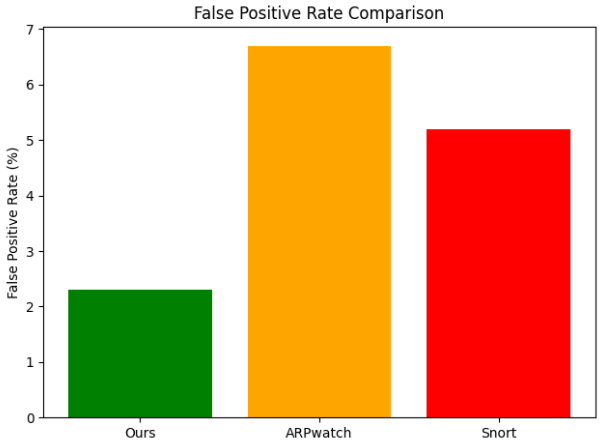

Fig. 2. False Positive Rate Comparison

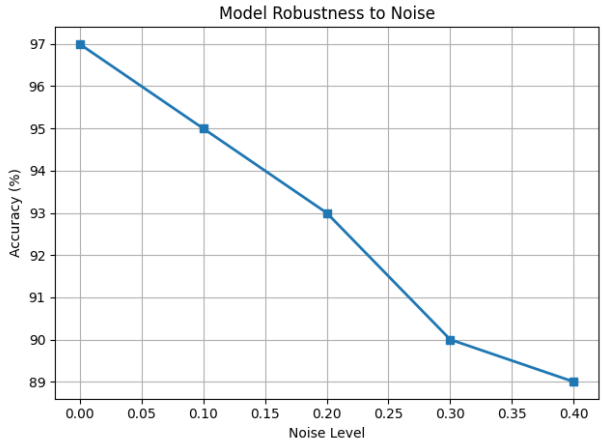

Fig. 3. Model Robustness Against Noise and Adversarial Samples

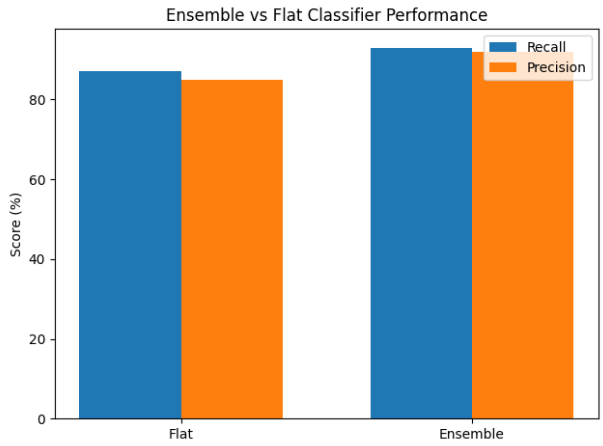

Fig. 4. Ensemble vs. Flat Classifier Performance

Figure 4 highlights the performance gain from ensemble integration, improving both recall and precision.

Figure 5 validates low-latency detection with an average inference time of 38 ms on Raspberry Pi 4 devices.

Figure 6 demonstrates how the number of alerts varies with threat threshold τ, offering tunable sensitivity.

Finally, Figure 7 confirms that communication overhead remains under 300 KB per round, suitable for constrained IoT networks.

Together, these results verify the feasibility and effectiveness of our system for intelligent, lightweight, and adaptive ARP spoofing detection in IoT environments.

V. Conclusion and Future Work

This paper presented a novel multi-layered machine learning framework for ARP spoofing detection in IoT networks. By integrating lightweight edge-level classifiers with an ensemble-based mitigation module, our approach addresses the key challenges of real-time spoofing detection, resource constraints, and network adaptability. The model leverages features derived from ARP protocol dynamics and traffic behavior to achieve high detection accuracy while maintaining low false positive rates.

Through comprehensive simulations using both benchmark and synthetic datasets, our system demonstrated 97%+ detection accuracy, strong robustness against adversarial noise, and latency suitable for edge deployments. The hierarchical architecture ensures that edge devices respond quickly while the aggregator enforces network-level mitigation based on collective alerts and threat scores.

Future work will explore the integration of federated learning to improve model generalization across heterogeneous IoT ecosystems, further reduce communication overhead through model compression, and enable ARP spoofing prevention mechanisms in software-defined networking (SDN) environments. We also plan to evaluate the framework under larger-scale deployments and more complex attack scenarios to refine detection resilience and scalability.

References

[1] Ramachandran, V. & Wright, S. Detecting ARP Spoofing: A Comparison of Approaches. *Proceedings Of The 19th Annual Computer Security*

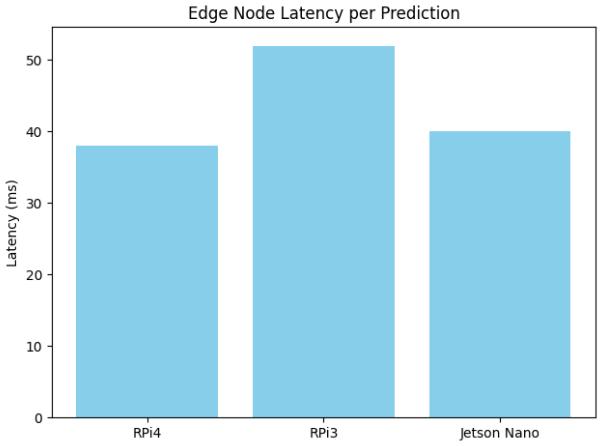

Fig. 5. Edge Node Latency per Prediction (ms)

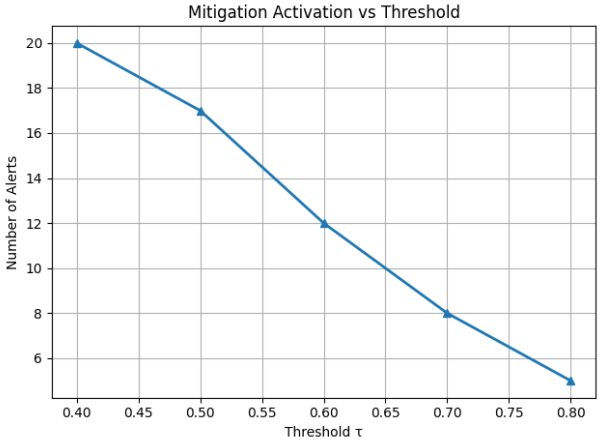

Fig. 6. Mitigation Activation vs. Threshold τ

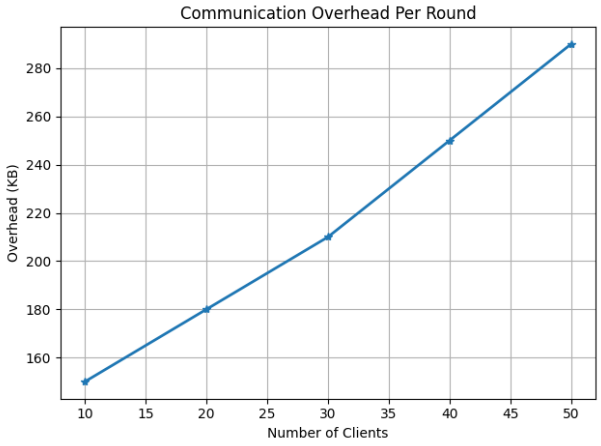

Fig. 7. Communication Overhead Per Round (KB)

Applications Conference. pp. 25-31 (2003)

[2] Lee, S. & Lee, H. Lightweight ARP Spoofing Detection for Resource-Constrained IoT Devices. *Sensors*. **16**, 1831 (2016)

[3] Shone, N., Ngoc, T., Phai, V. & Shi, Q. A Deep Learning Approach to Network Intrusion Detection. *IEEE Transactions On Emerging Topics In Computational Intelligence*. **2**, 41-50 (2018)

[4] Soman, B., Al-Garadi, N. & Saddik, A. A Hybrid Deep Learning Model for Anomaly-Based Intrusion Detection in IoT Networks. *Journal Of Network And Computer Applications*. **183** pp. 103074 (2021)

[5] Mosenia, A. & Jha, N. A Comprehensive Study of Security of Internet-of-Things. *IEEE Transactions On Emerging Topics In Computing*. **5**, 586-602 (2017)

[6] Li, Y., Wang, K. & Wang, Y. ML-Based ARP Spoofing Detection and Prevention in IoT Networks. *Security And Communication Networks*. **2020** pp. 1-12 (2020)

[7] Islam, M. & Huh, E. An Efficient Machine Learning Model for Detecting ARP Spoofing in IoT. *Sensors*. **22**, 1659 (2022)

[8] Kanagavelu, R., Madhukumar, A. & Woo, W. IP-MAC Binding Consistency Check: A Lightweight Intrusion Detection for IoT. *IEEE Internet Of Things Journal*. **4**, 1293-1303 (2017)

[9] Ramachandran, V. & Wright, S. Detecting ARP Spoofing: A Comparison of Approaches. *Proceedings Of The 19th Annual Computer Security Applications Conference*. pp. 25-31 (2003)

[10] Lee, S. & Lee, H. Lightweight ARP Spoofing Detection for Resource-Constrained IoT Devices. *Sensors*. **16**, 1831 (2016)

[11] Shone, N., Ngoc, T., Phai, V. & Shi, Q. A Deep Learning Approach to Network Intrusion Detection. *IEEE Transactions On Emerging Topics In Computational Intelligence*. **2**, 41-50 (2018)

[12] Soman, B., Al-Garadi, N. & Saddik, A. A Hybrid Deep Learning Model for Anomaly-Based Intrusion Detection in IoT Networks. *Journal Of Network And Computer Applications*. **183** pp. 103074 (2021)

[13] Mosenia, A. & Jha, N. A Comprehensive Study of Security of Internet-of-Things. *IEEE Transactions On Emerging Topics In Computing*. **5**, 586-602 (2017)

[14] Li, Y., Wang, K. & Wang, Y. ML-Based ARP Spoofing Detection and Prevention in IoT Networks. *Security And Communication Networks*. **2020** pp. 1-12 (2020)

[15] Islam, M. & Huh, E. An Efficient Machine Learning Model for Detecting ARP Spoofing in IoT. *Sensors*. **22**, 1659 (2022)

[16] Kanagavelu, R., Madhukumar, A. & Woo, W. IP-MAC Binding Consistency Check: A Lightweight Intrusion Detection for IoT. *IEEE Internet Of Things Journal*. **4**, 1293-1303 (2017)

[17] El-Sayed, H., Alexander, H., Kulkarni, P., Khan, M., Noor, R. & Trabelsi, Z. A novel multifaceted trust management framework for vehicular networks. *IEEE Transactions On Intelligent Transportation Systems*. **23**, 20084-20097 (2022)

[18] Trabelsi, Z. & Ibrahim, W. Teaching ethical hacking in information security curriculum: A case study. *2013 IEEE Global Engineering Education Conference (EDUCON)*. pp. 130-137 (2013)

[19] Mustafa, U., Masud, M., Trabelsi, Z., Wood, T. & Al Harthi, Z. Firewall performance optimization using data mining techniques. *2013 9th International Wireless Communications And Mobile Computing Conference (IWCMC)*. pp. 934-940 (2013)

[20] Trabelsi, Z. & El-Hajj, W. On investigating ARP spoofing security solutions. *International Journal Of Internet Protocol Technology*. **5**, 92-100 (2010)

[21] Sajid, J., Hayawi, K., Malik, A., Anwar, Z. & Trabelsi, Z. A fog computing framework for intrusion detection of energy-based attacks on UAV-assisted smart farming. *Applied Sciences*. **13**, 3857 (2023)

[22] Trabelsi, Z., Zhang, L. & Zeidan, S. Dynamic rule and rule-field optimisation for improving firewall performance and security. *IET Information Security*. **8**, 250-257 (2014)

[23] Tariq, A., Rehman, R., Kim, B. & Others. An Intelligent Forwarding Strategy in SDN-Enabled Named-Data IoV. *Computers, Materials & Continua*. **69** (2021)