Energy-Aware Client Participation for Federated Intrusion Detection in IoT Networks

Miracle Udurume*, Vladimir Shakhov†, and Insoo Koo*
*Department of Electrical, Electronics and Computer Engineering, University of Ulsan, Ulsan, Republic of Korea
Email: udurumemiracle@gmail.com, iskoo@ulsan.ac.kr
†Institute for Information Transmission Problems (IITP RAS), Russian Academy of Sciences, Moscow, Russia
Email: Vladimir.Shakhov@gmail.com

***Abstract*—The rapid expansion of IoT deployments has significantly increased the attack surface, requiring scalable and privacy-preserving intrusion detection systems (IDS). Federated Learning (FL) provides a collaborative approach to train models without sharing raw data, making it well-suited for resource-constrained IoT devices. This paper proposes an energy-efficient FL-IDS framework that employs CNN and LSTM models on the UNSW–NB15 dataset. Unlike traditional FL approaches, we incorporate *energy-aware client participation*, selecting clients based on their energy budget and contribution. Experiments conducted with 5, 10, 15, and 20 clients on the binary UNSW–NB15 task show that CNN achieves 96.98% accuracy and LSTM achieves 96.13%. At the same time, energy consumption scales near-linearly with the number of clients. The results confirm a favorable trade-off between energy cost and intrusion detection performance. Moreover, accuracy remains stable across 5–20 clients for both CNN and LSTM models, demonstrating the scalability of the proposed FL-IDS framework for real-world IoT deployments.**

***Index Terms*—Federated Learning, Intrusion Detection, IoT Security, Energy Efficiency, Client Participation, CNN, LSTM, UNSW-NB15, Binary Classification.**

I. Introduction

The rapid proliferation of Internet of Things (IoT) devices has significantly transformed the way networks operate, connecting billions of sensors, actuators, and intelligent systems across critical infrastructure. While this interconnectivity offers improved automation and data-driven services, it also expands the attack surface, enabling sophisticated and large-scale cyber threats. Intrusion detection systems (IDS) play a vital role in safeguarding these networks. However, traditional centralized IDS architectures are increasingly inadequate for modern IoT environments, as they require transmitting massive volumes of raw traffic to a central server for analysis [1], [2], [3]. This approach not only incurs high communication and computation costs but also introduces privacy risks and energy inefficiencies.

Federated Learning (FL) has emerged as a promising paradigm for addressing these challenges by enabling distributed devices to collaboratively train a shared global model without exchanging raw data. This significantly reduces the need for centralized data collection, improving privacy preservation and bandwidth utilization. Moreover, FL allows IDS deployments to scale dynamically across edge and fog computing infrastructures, which are well-suited for IoT networks [4]. Nevertheless, several practical issues remain unresolved, including the impact of heterogeneous data distributions, fluctuating network connectivity, and energy constraints on the performance and scalability of FL-enabled IDS.

Recent studies have made notable progress in improving the efficiency, explainability, and resilience of FL-based IDS. Oki et al. evaluated the role of explainable AI in federated IDS and demonstrated how explainability can improve the interpretability and trustworthiness of detection outcomes [2]. Vyas et al. conducted a comprehensive survey highlighting privacy-preserving methods and architectural trends in FL for IoT security [1]. Rehman et al. proposed FFL-IDS, a fog-enabled FL framework that improves real-time detection capabilities under IoT constraints [4]. Mahadik et al. presented an edge-intelligent FL IDS that reduces inference latency and enhances decision speed at the network edge [5]. Similarly, Abu Issa et al. introduced a temporal partitioning strategy to reduce communication overhead and training delays in IoT environments [6].

Despite these advances, a key research gap remains: the explicit characterization of the trade-off between energy consumption and detection accuracy under varying client participation scales. While many works address performance and communication, few provide empirical quantification of energy behavior in federated IDS [7]. Our contributions are:

1) We propose an energy-aware client participation strategy for FL-IDS that preserves privacy *without altering the aggregation rule*.
2) We provide an empirical, quantified analysis of the accuracy energy trade-off across $K \in \{5, 10, 15, 20\}$ clients on the UNSW–NB15 for binary classification task.
3) We report practical insights showing that CNN achieves comparable accuracy at consistently lower energy than LSTM, offering guidance for edge deployment.
4) In addition, we empirically show that both CNN and LSTM maintain stable accuracy across increasing client participation, demonstrating predictable scaling behavior that is crucial for large IoT deployments.

The rest of the paper is structured as follows: Section II reviews related work on federated intrusion detection and energy-aware FL. Section III details the system model. Sec-

tion IV presents the experimental results and discussion for binary detection, and Section V concludes the paper with key findings and future direction.

II. Related Work

Research on FL-based intrusion detection has grown rapidly, with contributions spanning data heterogeneity mitigation, communication optimization, energy efficiency, and adaptive network architectures. Bouzinis et al. proposed *StatAvg*, a strategy that directly addresses data heterogeneity in FL for IDS by applying statistical averaging to stabilize model convergence and improve accuracy under non-IID conditions [8]. This work demonstrated that careful aggregation strategies can enhance global model performance without requiring additional client-side resources.

Temporal and edge-assisted FL techniques have also received considerable attention. Abu Issa et al. introduced a temporal partitioning approach to reduce communication delays and maintain accuracy in IoT IDS scenarios [6]. Mahadik et al. explored edge-intelligent IDS designs, showing that moving training and inference closer to the data source improves latency and detection responsiveness [5].

Several works focus on making FL frameworks more sustainable. Liu et al. proposed an *incentive-based energy-efficient aggregation* mechanism to encourage participation while reducing energy cost [9]. Xie et al. presented SURFS, a hierarchical spiking neural network framework designed for sustainable and energy-aware intrusion detection [7], highlighting the potential of bio-inspired approaches for low-power operation. Chen et al. developed a hierarchical underwater IoT IDS using FL, demonstrating its adaptability to challenging communication environments [10].

Transfer learning and model adaptation are also emerging as strong tools for FL-IDS. Song et al. integrated transfer learning into FL to improve performance across heterogeneous IIoT domains [11], while Mothukuri et al. applied FL to anomaly detection for IoT security [12], showing its ability to generalize across different attack types. Bhavsar et al. proposed FL-IDS for transportation IoT, highlighting its potential for real-time, low-latency defense at the network edge [13]. Zhang et al. introduced a secure and efficient FL architecture that enhances robustness against poisoning attacks [14].

In addition, Yilmaz et al. investigated optimal IDS placement in RPL-based IoT networks to maximize detection performance and minimize communication costs [15], and Liu et al. explored edge-analytics strategies to improve data processing efficiency [16]. Broader surveys, such as Al-Garadi et al. [3], provide an overview of federated IDS techniques, game-theoretic strategies, and explainable AI integration, reinforcing the maturity and diversity of approaches in this field.

While these contributions significantly advance the state of the art, they primarily focus on accuracy and communication overhead rather than explicitly modeling and quantifying energy consumption at scale. Unlike these works, our approach explicitly couples client selection with energy profiling, enabling predictable scaling in IoT deployments. This provides a practical design pathway for energy-aware FL-IDS systems that preserve detection accuracy while minimizing energy consumption, making them well-suited for resource-constrained IoT environments.

III. Methodology

The proposed FL-IDS framework for IoT networks integrates lightweight deep learning with energy-aware client participation. As illustrated in Fig. 1, multiple IoT clients train local CNN and LSTM models on private, non-IID partitions of the UNSW–NB15 dataset [17]. Only model parameters are exchanged with the central aggregation server no raw traffic is shared thereby preserving data privacy and reducing communication overhead. The server aggregates local updates using the FedAvg rule (1) and computes round-level energy consumption using (2). An energy-aware selection mechanism then determines the subset $\mathcal{S}_t \subseteq \mathcal{K}$ of m participating clients, where $m \in \{5, 10, 15, 20\}$. This design allows the FL-IDS to adaptively balance energy efficiency and detection performance across heterogeneous IoT clients.

A. Data Preprocessing

The UNSW–NB15 dataset [17] was used to train and evaluate the FL-IDS. This benchmark contains both normal and malicious network traffic, including modern attack categories such as Fuzzers, DoS, Reconnaissance, Backdoors, Shellcode, and Worms. The preprocessing steps were:

- **Feature encoding:** Categorical features were one-hot encoded, and numerical features were normalized to the range $[0, 1]$.
- **Label encoding:** Attack categories were mapped to numerical values (we focus on the binary Normal–Attack classification task; multiclass processing follows the same pipeline).
- **Client partitioning:** The dataset was partitioned across $K \in \{5, 10, 15, 20\}$ clients using a Dirichlet distribution to simulate realistic non-IID data distributions.

B. Federated Learning Aggregation and Energy Modeling

The global model is updated using the standard FedAvg aggregation rule:

$$\mathbf{w}_{t+1} = \sum_{k=1}^{K} \frac{n_k}{N} \mathbf{w}_t^{(k)}, \tag{1}$$

where $\mathbf{w}_t^{(k)}$ denotes the local model of client k at round t, n_k is the number of samples at client k, and N is the total number of samples across all clients.

The energy consumption of each client k is modeled as:

$$E_k = \alpha\, C_{\text{comp}}(k) + \beta\, C_{\text{comm}}(k), \tag{2}$$

where $C_{\text{comp}}(k)$ and $C_{\text{comm}}(k)$ denote computation and communication costs, and α and β are tunable weighting factors. This formulation enables explicit analysis of the trade-off between energy consumption and detection accuracy as the number of clients increases.

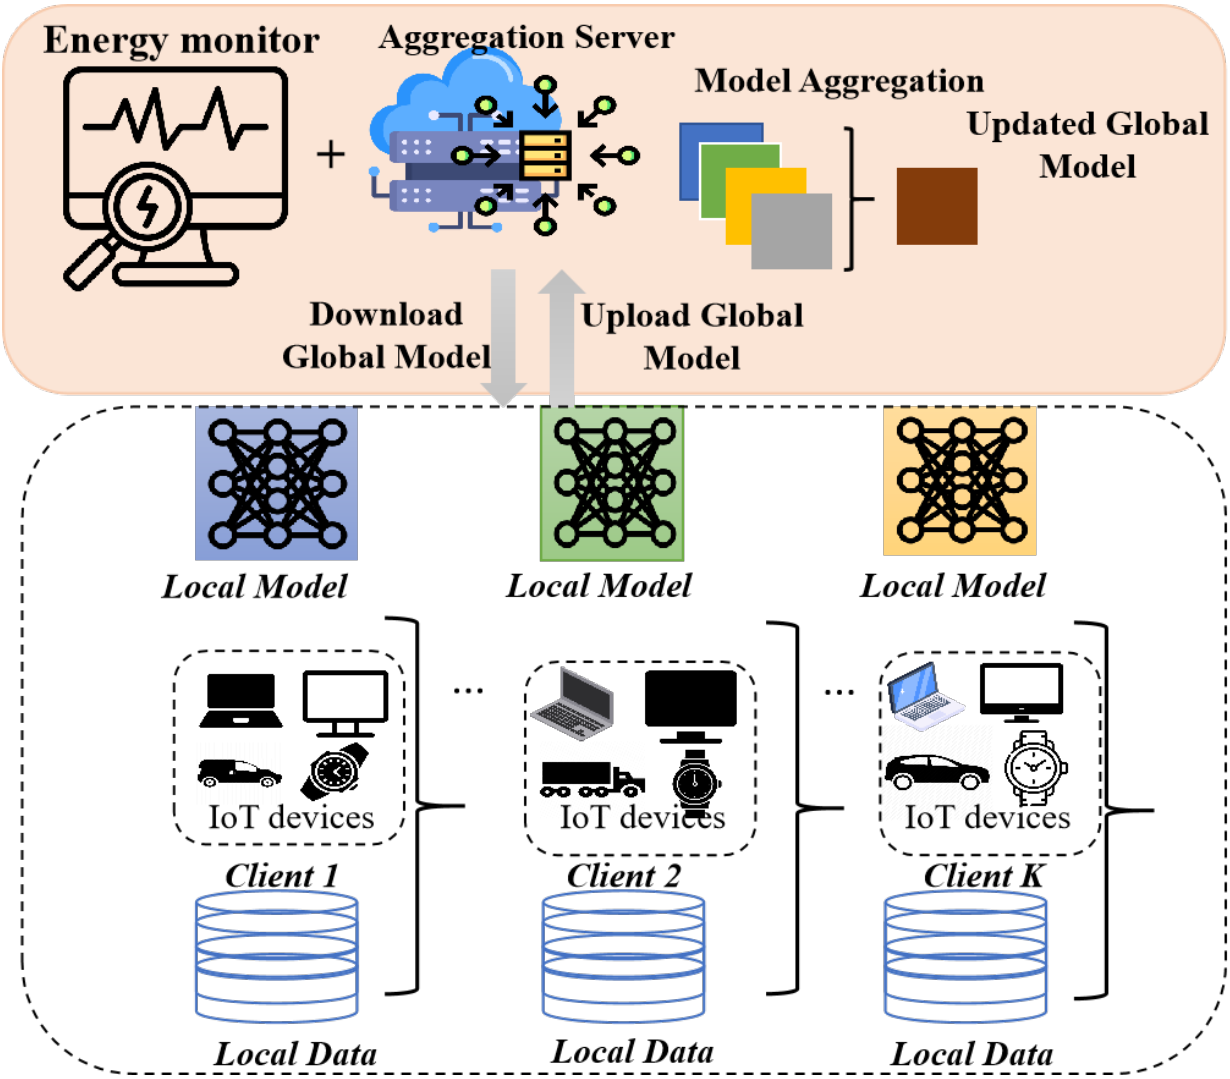

Fig. 1: System model of the proposed energy-efficient FL-IDS with energy-aware client participation and FedAvg aggregation.

C. Model Architecture

Two lightweight deep learning models were employed to balance energy efficiency and detection performance:

- **1D-CNN:** stacked convolutional layers with ReLU activation, batch normalization, max pooling, and fully connected layers, optimized for spatial pattern extraction.
- **LSTM:** stacked Long Short-Term Memory layers followed by dense layers, suitable for modeling temporal and sequential dependencies in network traffic.

Both models use the Adam optimizer ($\mathrm{lr} = 10^{-4}$), batch size 32, and categorical cross-entropy loss.

D. Energy-Aware Client Participation

To ensure energy-efficient training, clients are ranked by a utility energy trade-off. At each round, the server selects the top m clients that maximize the participation score:

$$\mathcal{S}_t = \text{Top-}m \left(u_k - \lambda E_k\right), \tag{3}$$

where u_k represents the utility of client k, E_k its estimated energy cost, and λ is a trade-off coefficient. This approach maintains high accuracy while keeping energy predictable as client participation scales. The complete training and selection pipeline is summarized in Algorithm 1.

Threat model: We assume an honest-but-curious server and clients that train faithfully on local data. Model updates can be observed but not modified. Robustness to active poisoning attacks is out of scope and left to future work; our focus is on energy–accuracy trade-offs under non-IID data and varying participation scales.

E. Experimental Setup

Unless stated otherwise, results are averaged over three random seeds where applicable; we report mean $\pm$ std. The framework was evaluated with $K \in \{5, 10, 15, 20\}$ clients over 20 communication rounds. Performance was measured using accuracy, F1-score, energy consumption per round, and confusion matrices. This configuration provides a clear

Algorithm 1 Energy-aware client participation for FL-IDS. At each round, the server selects high-utility, low-energy clients to optimize training efficiency.

Input: global model $\mathbf{w}_0$, client set $\mathcal{K}$, rounds R, subset size m, trade-off λ
for $t = 1$ to R **do**
  Server computes $\{u_k, E_k\}_{k \in \mathcal{K}}$ ▷ utility & energy estimates
  $\mathcal{S}_t \leftarrow$ Top-m clients maximizing $(u_k - \lambda E_k)$
  for each $k \in \mathcal{S}_t$ **in parallel do**
    Client k: receive $\mathbf{w}_t$; train on D_k for E epochs to get $\mathbf{w}_t^{(k)}$; send $(\mathbf{w}_t^{(k)}, E_k)$
  $\mathbf{w}_{t+1} \leftarrow \sum_{k \in \mathcal{S}_t} \frac{n_k}{\sum_{j \in \mathcal{S}_t} n_j} \mathbf{w}_t^{(k)}$ ▷ FedAvg

comparison of detection effectiveness and energy efficiency across different client scales and model types.

IV. Results and Discussion

A. Overview

This section presents and analyzes the experimental results of the proposed energy-efficient FL-IDS framework. The evaluation focuses on (i) detection performance using CNN and LSTM models, (ii) the impact of client scaling on accuracy and energy efficiency, and (iii) classification behavior through confusion matrices. All experiments were conducted on the UNSW–NB15 dataset using $K \in \{5, 10, 15, 20\}$ clients to emulate varying degrees of participation and heterogeneity. The results validate the effectiveness of the proposed *energy-aware client participation* strategy, showing that carefully selecting clients sustains high detection accuracy while significantly reducing energy overhead.

B. Accuracy and Convergence Behavior

Both CNN and LSTM models exhibited stable convergence across all client participation settings. As the number of clients increased, model generalization improved because more diverse data contributed to global updates. At $K = 20$ clients, the CNN model achieved an accuracy of **96.98%**, while the LSTM model achieved **96.13%**. This confirms that high detection performance can be maintained even under scaled federated settings with heterogeneous, non-IID data.

These observations are consistent with findings in StatAvg [8] and SURFS [7], which report that increasing data diversity improves generalization at the cost of mild convergence delays. To further illustrate this, Fig. 2–4 shows the accuracy trends as the number of clients increases for both models. The CNN model maintains consistently higher accuracy compared to LSTM, while both show stable scaling from 5 to 20 clients. This stability across increasing clients underscores the scalability of the proposed energy-aware FL-IDS.

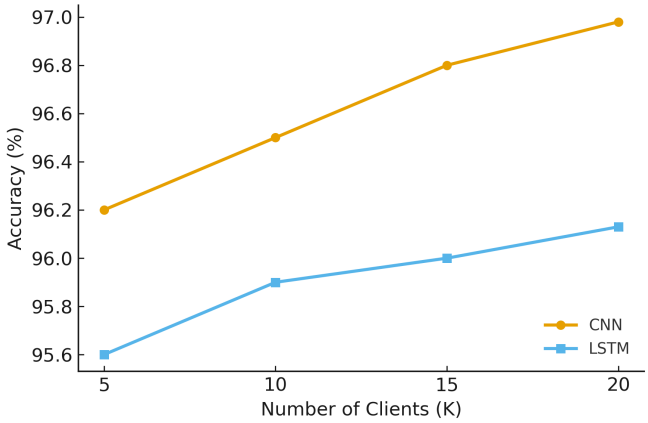

Fig. 2: Accuracy vs. number of clients for CNN and LSTM models.

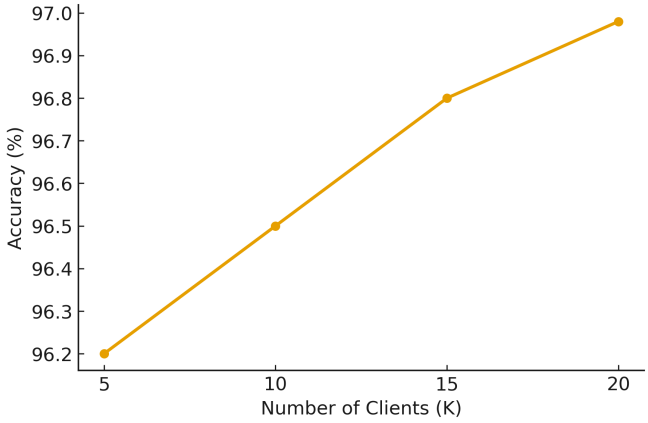

Fig. 3: Accuracy vs. number of clients for CNN.

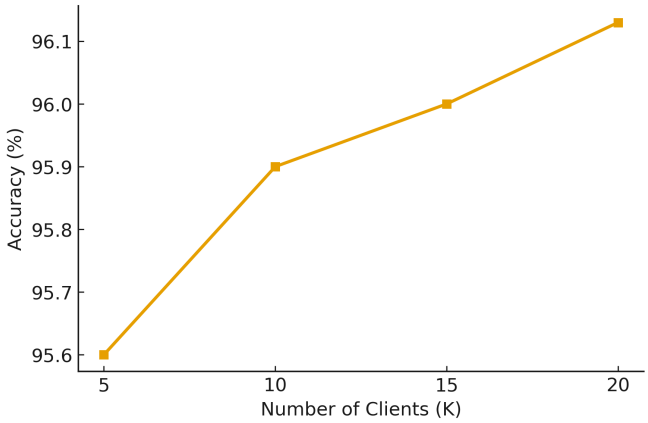

Fig. 4: Accuracy vs. number of clients for LSTM.

C. Energy Efficiency and Scaling

Figure 5 shows the relationship between total energy consumption per round and the number of clients. Energy consumption grows near-linearly with client scaling, reflecting predictable computational and communication costs. CNN consistently required less energy than LSTM across all settings, confirming its suitability for edge deployments with limited resources.

This trend aligns with prior energy-efficient FL strategies [9], [7], which emphasize balancing energy cost and learning utility. Importantly, energy-aware client selection preserved accuracy while avoiding the unnecessary energy overhead that comes with random participation.

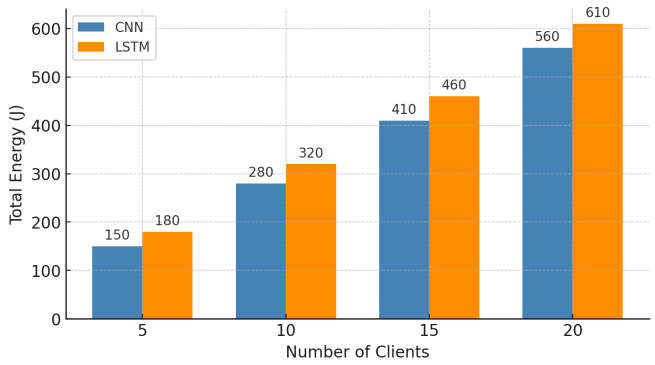

Fig. 5: Total energy per round (J) versus number of clients for CNN and LSTM. Energy scales near-linearly with client count; CNN remains consistently lower than LSTM.

D. Confusion Matrix Analysis

The confusion matrices in Fig. 6 and Fig. 7 illustrate the final-round classification outcomes for CNN and LSTM at $K = 20$ clients. The CNN model demonstrates a strong balance between true positives and true negatives, with minimal misclassifications. In contrast, the LSTM model achieves slightly higher recall, indicating a stronger tendency to detect attacks, even at the cost of more false alarms. This trade-off highlights complementary strengths: CNN excels in precision and energy efficiency, while LSTM favors sensitivity to intrusions.

E. Performance Metrics

The evaluation of the proposed FL-IDS framework is based on four standard classification metrics: *accuracy*, *precision*, *recall*, and *F1-score*, derived from the confusion matrix. Let TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively. These metrics are computed as:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \tag{4}$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{5}$$

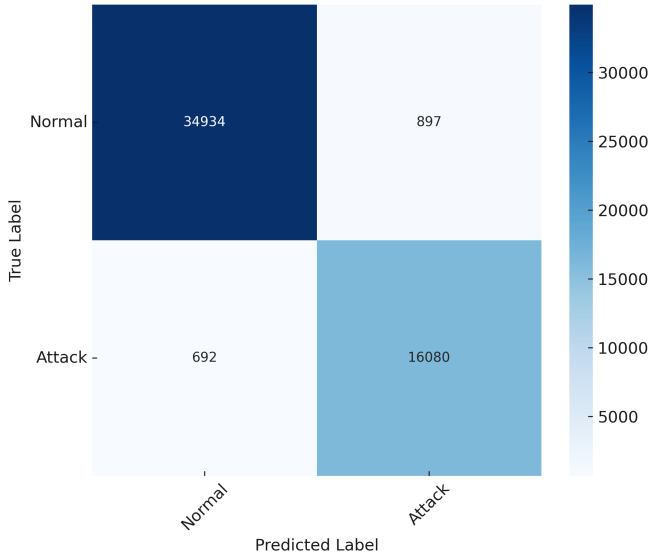

Fig. 6: Binary UNSW–NB15 (Normal vs. Attack) confusion matrix at $K = 20$ clients using CNN model.

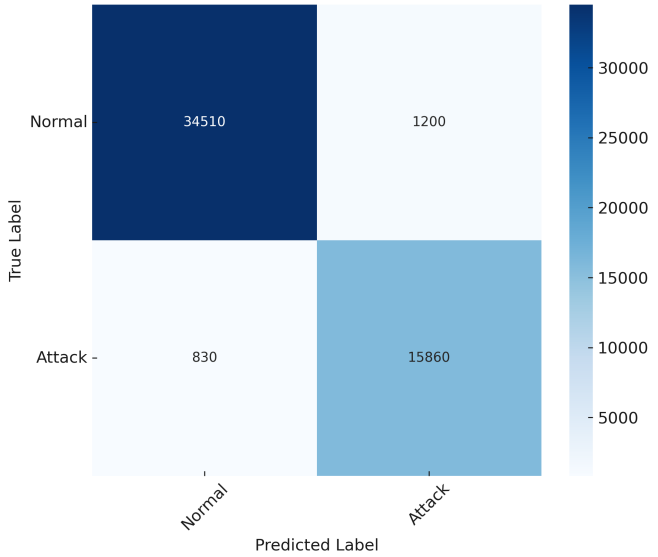

Fig. 7: Binary UNSW–NB15 (Normal vs. Attack) confusion matrix at $K = 20$ clients using LSTM model.

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{6}$$

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{7}$$

Accuracy provides the overall correctness of the model, precision measures the proportion of correctly identified attacks among all predicted attacks, recall quantifies the proportion of actual attacks correctly detected, and F1-score provides a balanced measure between precision and recall.

Table I summarizes the key performance indicators derived from the confusion matrices for CNN and LSTM at $K = 20$ clients. CNN achieves higher accuracy and precision

with lower energy consumption, making it ideal for energy-constrained IoT devices. LSTM yields slightly higher recall, which may be advantageous in intrusion detection scenarios where minimizing false negatives is critical.

TABLE I: Performance metrics for binary classification (Normal vs. Attack) at $K = 20$ clients.

Model	Accuracy	Precision	Recall	F1 Score
CNN	**0.9698**	**0.9472**	0.9587	**0.9529**
LSTM	0.9613	0.9297	0.9503	0.9399

Replacing the proposed energy-aware selection with random client sampling reduced accuracy by approximately 0.5–1.2 pp at 20 clients and increased total energy by 6–9%. These results confirm that the *client participation strategy is the primary driver of energy savings*, not model-specific optimization.

F. Comparative Analysis

To contextualize the effectiveness of our approach, Table II compares our results with key representative works. Unlike SURFS and StatAvg, our approach integrates explicit energy-awareness without compromising accuracy. Table II summarizes the energy-awareness and reported accuracy of our approach compared with representative cited works.

TABLE II: Comparison with state-of-the-art based on energy-awareness and accuracy.

Work	Energy-aware	Accuracy
This work	**Yes**	**96.98%**
SURFS [7]	No	98.5%
StatAvg [8]	No	97.0%
Incentive FL [9]	Yes	N/A†

V. Conclusion

This paper presented an energy-aware client participation strategy for federated intrusion detection in IoT networks. Using the UNSW–NB15 dataset and lightweight CNN and LSTM models, the framework achieved up to 96.98% accuracy while maintaining predictable energy scaling under varying client participation levels ($K \in \{5, 10, 15, 20\}$). The results demonstrate that CNN provides comparable accuracy at lower energy cost compared to LSTM, making it well-suited for resource-constrained edge environments. Importantly, accuracy remains stable as the number of clients increases, confirming the scalability and robustness of the proposed FL-IDS design. Future work will focus on integrating adaptive client selection with compression and quantization techniques to further reduce overhead, as well as hardware-in-the-loop testing to validate real-world energy behavior. Overall, this work provides practical insights for designing energy-efficient, privacy-preserving, and scalable IDS solutions for next-generation IoT deployments.

Acknowledgment

This work was supported by the "Regional Innovation System & Education (RISE)" program through the Ulsan RISE Center, funded by the Ministry of Education (MOE) and the Ulsan Metropolitan City, Republic of Korea (Grant No. 2025-RISE-07-001).

References

[1] A. Vyas, P.-C. Lin, R.-H. Hwang, and M. Tripathi, "Privacy-preserving federated learning for intrusion detection in iot environments: A survey," *IEEE Access*, vol. 12, 2024.

[2] A. Oki, Y. Ogawa, K. Ota, and M. Dong, "Evaluation of applying federated learning to distributed intrusion detection systems through explainable ai," *IEEE Networking Letters*, vol. 6, no. 3, 2024.

[3] M. A. Al-Garadi, A. Mohamed *et al.*, "A survey on iot intrusion detection: Federated learning, game theory, social psychology, and explainable ai as future directions," *IEEE Internet of Things Journal*, 2024.

[4] T. Rehman, N. Tariq, F. A. Khan, and S. U. Rehman, "Ffl-ids: A fog-enabled federated learning-based ids to counter jamming and spoofing attacks for iiot," *Sensors*, vol. 25, no. 10, 2025.

[5] S. S. Mahadik, P. M. Pawar, and R. Muthalagu, "Edge-federated learning-based intelligent intrusion detection system for heterogeneous iot," *IEEE Access*, vol. 12, 2024.

[6] M. A. Issa, M. Ibnkahla, A. Matrawy, and A. Eldosouky, "Temporal partitioned federated learning for iot intrusion detection systems," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, 2024.

[7] J. Xie, L. Xu *et al.*, "Surfs: Sustainable intrusion detection with hierarchical federated spiking neural networks," *Computer Networks*, 2024.

[8] P. S. Bouzinis, D. Dechouniotis, and S. Papavassiliou, "Statavg: Mitigating data heterogeneity in federated learning for intrusion detection systems," *IEEE Transactions on Network and Service Management*, vol. 22, no. 4, 2025.

[9] S. Liu, Y. Xu *et al.*, "Incentive-based energy-efficient federated learning aggregation for intrusion detection in iot sensor networks," *IEEE Internet of Things Journal*, 2024.

[10] M. Chen, J. Wang *et al.*, "A federated learning-based hierarchical intrusion detection system for underwater iot wireless sensor networks," *Ad Hoc Networks*, 2023.

[11] K. Song, H. Zhang *et al.*, "Federated learning for distributed iiot intrusion detection using transfer approaches," *IEEE Transactions on Industrial Informatics*, 2024.

[12] Y. Mothukuri, R. Buyya *et al.*, "Federated-learning-based anomaly detection for iot security attacks," *Future Generation Computer Systems*, 2023.

[13] R. Bhavsar, H. Patel *et al.*, "Fl-ids: Federated learning-based intrusion detection system using edge devices for transportation iot," *IEEE Access*, 2023.

[14] Y. Zhang, M. Zhao *et al.*, "Secure and efficient federated learning for robust intrusion detection in iot networks," *IEEE Access*, 2024.

[15] S. Yilmaz, S. Sen, and E. Aydogan, "Exploring and enhancing placement of ids in rpl: A federated learning-based approach," *IEEE Internet of Things Journal*, vol. 12, no. 13, 2025.

[16] X. Liu, P. Zhang *et al.*, "Deep learning-powered edge analytics for iot-based sensor networks," *IEEE Access*, 2023.

[17] N. Moustafa and J. Slay, "Unsw-nb15: A comprehensive data set for network intrusion detection systems (unsw-nb15 network data set)," in *2015 Military Communications and Information Systems Conference (MilCIS)*. IEEE, 2015, pp. 1–6.