3RD WORKSHOP ON MACHINE LEARNING IN NETWORKING (MaLeNe) PROCEEDINGS

SEPTEMBER 1, 2025

CO-LOCATED WITH
THE 6TH INTERNATIONAL CONFERENCE ON
NETWORKED SYSTEMS (NETSYS 2025)
ILMENAU, GERMANY

Decentralized Federated Learning for Intrusion Detection in 5G Networks: An Asynchronous Consensus-driven Multi-Agent Approach

Nasim Nezhadsistani[1], Francisco Enguix[2], Carlos Carrascosa[2], Burkhard Stiller[1]
[1]Communication Systems Group, Department of Informatics, University of Zürich, Switzerland
Email: {nezhadsistani, stiller}@ifi.uzh.ch
[2]Valencian Research Institute for Artificial Intelligence (VRAIN), Universitat Politècnica de València (UPV),
Camino de Vera s/n, 46022, Valencia, Spain
Email: fraenan@upv.es, carrasco@dsic.upv.es

***Abstract*—5G networks expand the wireless infrastructure attack surface while concurrently limiting data streams through in-your-face privacy regulations. Centralized traffic aggregation-based traditional intrusion-detection pipelines, thus, are plagued with bottlenecks, sole points of failure, and regulatory insurrection. This paper uses PA-CoL, a Parallel Asynchronous Consensus-based Learning framework in which multi-agent base stations train a shared deep neural model without ever exporting raw traffic records. Each agent performs local mini-batches on 5G-NIDD flow features and intermittently averages parameters with a randomly chosen neighbor; no parameter server or global synchrony is required. Experiments across three overlay graphs (Complete, Ring, Small-World), two data Types (IID and non-IID), and three federation sizes (5, 8, 10 agents) show that the newly developed scheme reaches F1 ≥ 0.99 under IID data and F1 ≥ 0.93 under severe non-IID skew. These results indicate that a lightweight, peer-to-peer consensus can deliver carrier-grade intrusion detection for privacy-sensitive 5G edge clouds and can pave the way toward 6G self-defending networks.**

***Index Terms*—5G Network, Federated Learning (FL) , Multi-Agent Systems (MAS), Intrusion Detection.**

I. Introduction

Fifth-generation (5G) mobile networks are a wireless communications paradigm shift that provides greater bandwidth, ultra-reliable low-latency (URLLC), and massive machine-type communications (mMTC). This results in new capabilities like autonomous vehicles, smart manufacturing, and remote medicine. But with these capabilities comes the price tag of an exponentially increased attack surface. Use of encrypted network traffic, pervasive use of network slicing, and the inherent dynamic nature of 5G topologies introduce high degrees of complexity for threat detection and security monitoring. Conventional security analytics, having a dependency on centralized data aggregation and correlation, become less effective in such cases [1]. Centralized solutions have critical drawbacks, including bottlenecks, single points of failure, and lack of end-to-end visibility in highly segmented virtualized environments [4], [10].

An additional level of intricacy arises from stricter data sovereignty and privacy regulations, such as the General Data Protection Regulation (GDPR) within the European Union. These regulations mandate that organizations maintain sensitive user data within local or national borders, effectively excluding unregulated trafficking of raw data to distant cloud servers for processing. Security products, therefore, relying on packet capture and aggregation of all traffic across the network, are not only technologically infeasible but also most likely to be non-compliant with global privacy standards. There is therefore a pressing need for threat detection and response capabilities without losing data locality in the security architectures [28], [29], [31].

Federated Learning (FL) [2] proves to be a viable paradigm to overcome privacy and locality of data concerns. FL facilitates training machine learning models in a decentralized manner through enabling local nodes or edge devices to hold their data locally and send model updates (e.g., parameters or gradients) to a coordination agent. This paradigm reduces risk to privacy and restricts exposure of sensitive data to an absolute minimum. But the traditional FL architecture usually depends on a parameter server in the middle to collect model updates and scatter the updated global model. This server-based architecture creates new threats, such as communication bottlenecks and single points of failure, which are undesirable in the case of large-scale, mission-critical 5G networks. In addition, centralized FL presumes synchronized and consistent participation of all the nodes that, in reality, cannot be ensured due to heterogeneity and dynamic connectivity of mobile devices [9], [16], [18].

Decentralized Federated Learning (DFL) has been suggested to bypass the downsides of centralized servers. Peer-to-peer cooperation in DFL substitutes the central aggregation node such that nodes are able to train models jointly in an unsupervised way without a coordinator node. The distributed approach improves robustness and scalability but typically comes with the cost of rigorous synchronization among cooperating nodes. 5G real-world scenarios are proven to be challenging in device capability, network state, and participation rates shifting unpredictably [5]. Multi-Agent Systems (MAS) [3] offer the natural paradigm for distributed decision-making and intelligence for advanced environments. Autonomous agents and local sensing and knowledge-directed agents interact with each other in MAS to achieve common goals. Translated to 5G security, MAS enable localized discovery and quick response to emerging threats, allowing the requirements of adaptive, distributed defenses [32].

To take advantage of the strong points of DFL and MAS and handle synchrony and heterogeneity problems, this paper uses Parallel Asynchronous Consensus-based Learning (PA-CoL) [8]. PA-CoL facilitates the convergence of the models of decentralized agents without the need for global synchronization. Agents instead update and exchange their models asynchronously with neighbors in such a way that the network converges as an aggregate under delays, network partition, or stragglers. This asynchrony is particularly suited for variable conditions under 5G networks. The key contributions of this paper are as follows:

1) This paper constructs an entirely distributed intrusion detection system (IDS) using PA-CoL over a MAS overlay specially designed for the distinctive nature and needs of 5G networks.
2) This paper applies the developed framework using the *PyTorch* deep learning framework and performs an extensive experiment on the 5G-NIDD dataset [1], with specific emphasis on non-IID (non-identically and independently distributed) data scenarios.
3) This paper compares our detection performance, communication cost, and scalability with conventional FL benchmarks and shows its performance benefits in real-world 5G scenarios.

The remainder of the paper is organized as follows. Section II provides related work on intrusion detection for 5G networks, FL, decentralized training algorithms, and multi-agent systems. Section III offers the proposed framework architecture, such as system model, dataset preprocessing, learning algorithm, and comparisons with baseline methods. Section IV provides a description of the experimental setup and results focusing on performance measures, communication efficiency, and resilience with non-IID scenarios. Finally, Section V summarizes the paper and presents future directions of work.

II. Related Work

A. Federated Learning for Network Security

FL has become an attractive paradigm for cooperative, privacy-aware intrusion detection. The initial *FedAvg* algorithm [2] proved that distributed devices could perform local model learning on their own data in parallel and then exchange model updates with a central server to collect. This work motivated diverse research on privacy-preserving machine learning for network security. The latest developments also involve the use of advanced models like BERT, which has made remarkable discoveries of up to 97.8% accuracy for intrusion detection based on datasets [6]. However, most recent FL solutions still rely on a central server to handle the learning process. This reliance is accompanied by a cluster of weaknesses: the master server is a point of contention, a target for denial-of-service or poisoning attacks, and a single point of failure, which compromises the resilience and scalability needed for 5G network security.

B. Multi-Agent Systems in 5G Security

MAS provides an adaptive and agile answer to distributed security in 5G networks. In MAS, autonomous heterogeneous agents with local knowledge and goals interact and collaborate to detect, respond to, and pursue threats. This distributed approach is most applicable to the dynamic, partitioned, and large-scale environment of 5G. The combination of MAS with FL was even investigated recently to further improve system resilience and fault tolerance to allow agents the capability to collaboratively train and improve security models in the absence of a central controller [7]. Although these are promising directions, existing research in these directions has been mainly theoretical models or simulations, with extremely few being based on real 5G datasets like 5G-NIDD. Therefore, there is an urgent need for experimental investigations that assess the efficiency and usability of MAS-FL hybrids in real-world 5G deployment environments.

C. From Centralized to Decentralized FL in 5G IDS

Early 5G intrusion detection research used the standard *FedAvg* framework, where edge devices provide gradient updates to a parameter server; the framework attains high accuracy on IIoT data but suffers from a single point of failure and is still susceptible to model-poisoning attacks [12], [14]. Server-side bandwidth constraints also reduce scalability when thousands of gNBs transmit traffic features every few seconds. A number of studies find that server-centrism is antithetical to the URLLC aspirations of 5G and support peer-to-peer aggregation instead [11], [13]. There exists a taxonomy of decentralized forms with security in focus as well [5]. They collectively induce a turn towards topologically fully decentralized or at least hierarchically organized training structures.

D. Asynchronous Decentralized FL

Weighted-average consensus on an overlay graph eliminates the server bottleneck, while the synchronous one continues to assume uniform compute speed [27]. Multilayer consensus and dynamic-average consensus enable all agents to update whenever resources become available and thus eliminate stragglers [16], [18]. The pull-based protocol generalizes the concept further by allowing nodes to pull the latest neighbor model on demand [9], while the latency-compensated scheduler rewinds stale gradients [17]. All such systems, though, broadcast each update across the network and thereby flood low-power radios.

E. Security, Privacy, and Verifiability

Consensus ledger verifiable computation prevents tampering but adds additional latency [15]; end-to-end correctness proofs add even more integrity [21]. Update-level trust scoring and the proof-of-data system, which is cryptographic, prevents decentralized FL poisoning but has not yet been tested with real 5G traffic [19], [20]. Communication-saving methods like locally differentially private updates indicate privacy and efficiency can be pursued in tandem [25].

F. Parallel Asynchronous Consensus-based Learning

Consensus-Based Learning approaches make fusion into an iterated averaging problem with no kind of global coordinator [30]. The most recent evolution, PA-CoL, features an explicit coalition layer [8]. Agents sharing similar semantics or geographically co-located agents in PA-CoL organize intra-coalition clusters that reach consensus in frequent rounds; models aggregated from each coalition leader engage in the slower inter-coalition exchange. Experiments across a non-IID image data set reduce total bytes transferred by about 35 % without affecting accuracy compared to single-coalition baselines, and message-direction analysis validates that most traffic is local between coalitions [8]. Since 5G slices natively map to trust domains, we generalize the coalition concept to slice-aware security monitoring [8], [33].

G. Graph and Self-Supervised Models for Network Data

Graph neural networks improve representation of host-to-host relationships and protocol hierarchies: Edge-based GraphSAGE in [26], DIGNN-A in [22], and feature-rearranged (GNNs in [23] all see double-digit F1 gains over multilayer perceptrons. Self-supervised pretext tasks like TS-IDS in [24] see comparable large gains when labelled data is limited. Nevertheless, parameter counts are an order of magnitude greater than for lightweight MLPs, increasing the communication overhead decentralized FL already needs to cope with. Compression methods like the autoencoder pipeline in [29] and hierarchical clustering in [30] mitigate this expense in part.

III. Experimental Design

A. Data Set and Pre-processing

The evaluation relies on the public 5G-NIDD dataset [1], which captures network traffic across nine behaviour classes. Each record carries forty-eight normalised numerical features. Rows with missing values less than $0.1\,\%$ of the file are discarded to avoid introducing noise. A stratified hold-out split reserves $20\,\%$ of the cleaned data for testing so every class preserves its prior. Feature scales are already aligned; therefore, no extra standardisation is applied.

B. Neural Architecture and Local Optimisation

Both agents have the same feed-forward MLP depicted in Fig. 1. There are two hidden widths 128 layers with ReLU activations and dropout rate 0.3 to prevent over-fitting. The output layer provides nine logits. Training uses the Adam optimiser with a step size of 0.001 and categorical cross-entropy loss. A training iteration is a single forward-and-backward pass over a mini-batch of 32 samples.

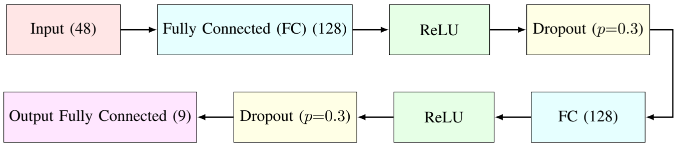

Fig. 1: Architecture of the MLP model with two hidden layers of size 128 and dropout.

C. Data Partitioning Strategy

To examine the impact of statistical bias, two placement modes are considered. *IID mode*: The fold employed for training is shuffled and divided equally into shards, and each agent gets an equal amount of samples and class distribution. *Non-IID mode*: sampling indices are sampled from Dirichlet concentration $\alpha = 0.3$. Biases samples class proportions such that some agents have labels unobservable elsewhere, simulating true edge drift.

D. Round Structure and Consensus Rule

Agents transfer information along static undirected graphs of five, eight, or ten nodes. The considered topologies are complete, small-world (rewiring probability 0.3), and ring. The largest node degree $\deg_{\max}$ fixes the mixing rate

$$\epsilon = \frac{1}{\deg_{\max}} \tag{1}$$

During each round, an agent first performs one local epoch, then selects one neighbor at random and updates its weights via

$$\theta_i \leftarrow (1-\epsilon)\,\theta_i + \epsilon\,\theta_j, \tag{2}$$

where θ_i and θ_j are the agent's parameter vector and the selected neighbor's parameter vector, respectively. There is one such exchange in each round. The protocol is executed for 500 rounds for each $\{$*graph, partition mode*$\}$.

E. Hardware and Software stack

All experiments were conducted on a computing environment equipped with an Intel Core i7 13650HX CPU, featuring 20 logical processors, and 32 GB of RAM. The graphics processing was managed by an NVIDIA GeForce RTX 4070 Laptop GPU. The system used an NVMe SSD storage solution, ensuring rapid data read and write speeds. The operating system for this setup was Windows 11 Pro 64-bit.

IV. Experimental Results

Table I reports the maximum mean-over-agent performance metrics obtained by the parallel Asynchronous Consensus-based Learning framework. This paper evaluates three communication topologies (Complete, Ring, and Small-World), two data distributions (IID and non-IID), and three network sizes (5, 8, and 10 agents). Under IID conditions, different topologies perform almost the well as they can and show very few differences in all metrics and network sizes. This is because, with IID, each agent can access data that is equally representative, which leads to steady and consistent convergence no matter what the network structure is. For example, accuracies and F1 scores often get close to 0.995, which confirms that topology does not matter much in this ideal situation. On the other hand, the non-IID setting (with Dirichlet skew) is more realistic and difficult, as agents have very different and biased local data distributions. In this case, the topology choice becomes a very important design consideration that impacts convergence and the final result.

Under non-IID conditions, the *Complete* topology exhibits significant performance degradation. For example, with 8 agents, the accuracy drops dramatically to 0.6973, and the F1 score decreases to 0.5908. This highlights the vulnerability of fully connected graphs to bias amplification, as global averaging may allow agents with extreme local distributions to disproportionately influence the consensus model. Increasing the number of agents from 8 to 10 slightly improves performance (accuracy rising to 0.8558), suggesting that increased diversity among agents can partially counteract individual biases, but overall, the Complete topology remains the most sensitive to non-IID data.

The ring topology shows greater strength in non-IID setups. By restricting each agent's interactions to close neighbors, it lessens the spread of strong local biases. As shown in Table I, the ring topology keeps high accuracy and F1 scores (for example, an accuracy of 0.9677 and an F1 score of 0.9370 with eight agents), which points to more consistent and dependable model behavior. Also, the ring topology's behavior stays even and less jumpy across varied network sizes, showing it can scale with data differences.

The small-world topology has middle-of-the-road behavior in non-IID settings. While it gains from shortcuts that allow quicker information spread than a simple ring, it also struggles with the impact of highly linked hub nodes, which can make local biases worse if these hubs have slanted data. This is clear in the fairly low precision and F1 scores (like a precision of 0.8341 and an F1 score of 0.7919 with eight agents), which shows a give-and-take between quicker mixing and flexibility to differences.

In considering how scaling impacts performance, merely adding more agents does not assure better results in non-IID setups. For instance, the Complete topology improves somewhat with ten agents. But the Ring and Small-World topologies are more stable, without big drops in performance. This suggests that just increasing the number of agents isn't enough to fix data differences.

The network structure must be carefully thought out. From a real-world use point, these results point to using sparse, structured overlays like the Ring topology when data is inherently non-IID. Situations include edge networks in 5G and distributed intrusion detection systems. Keeping high F1 scores (above 0.93) even with few communication links per agent gives a good balance. This balance is between detection quality and communication costs, which is key for privacy and places with limited bandwidth.

TABLE I: Maximum of the mean-over-agents metrics.

Nodes	Topology	Dirichlet (non-IID)			IID		
		Accuracy	Precision	F1	Accuracy	Precision	F1
5	Complete	0.9687	0.9571	0.9545	0.9959	0.9929	0.9916
	Ring	0.9621	0.9591	0.9383	**0.9963**	**0.9943**	**0.9937**
	Small-World	0.9410	0.8642	0.7990	0.9959	0.9929	0.9918
8	Complete	0.6973	0.6647	0.5908	0.9943	**0.9883**	**0.9845**
	Ring	0.9677	0.9519	0.9370	**0.9949**	0.9875	0.9833
	Small-World	0.7191	0.8341	0.7919	0.9945	0.9882	0.9838
10	Complete	0.8558	0.6964	0.6406	0.9934	0.9859	0.9822
	Ring	0.9547	0.9300	0.9066	0.9944	**0.9869**	**0.9831**
	Small-World	0.8938	0.7209	0.6881	**0.9948**	0.9862	0.9828

Fig.2 (a)-(i) illustrates the per-round evolution of *accuracy*, *F1 score*, and *precision* for networks consisting of 10, 8, and 5 agents, each evaluated across six configurations. These configurations combine three communication topologies (Complete, Ring, and Small-World) with two data distributions (IID and Dirichlet-skewed, $\alpha = 0.3$).

In the IID setting (indicated by blue, orange, and green curves), all topologies achieve rapid and smooth convergence. Accuracy, precision, and F1 score all exceed 0.95 within approximately 10 to 20 communication rounds, regardless of the number of agents. The Complete topology exhibits the steepest initial improvement, typically saturating in as few as 5 rounds. Ring and Small-World topologies follow slightly slower but still very close trajectories. This behavior confirms that when data is balanced and fully representative at each agent, the effect of communication topology becomes negligible. Thus, these IID results act merely as an upper bound on possible performance.

Under non-IID conditions (Dirichlet, $\alpha = 0.3$), the convergence patterns change dramatically and reveal the true impact of topology. The Ring topology (purple curves) shows the most robust and stable performance. For example, in the 8-agent configuration, accuracy surpasses 0.95 by around round 40, and F1 score exceeds 0.90 by about round 60, with minimal oscillations in subsequent rounds. This stability arises because each agent exchanges information only with local neighbors, which helps prevent global consensus from being dominated by agents with highly skewed local data.

The Complete topology (red curves) shows high initial instability with non-IID data. Initial F1 scores may fall below 0.60 in the first 50 rounds, and the trajectory has clear oscillations before gradually stabilizing. Even after 500 rounds, final F1 scores only stabilize around 0.75 to 0.78, highlighting the vulnerability of global averaging to local bias amplification. Nevertheless, as the number of agents increases, the added diversity among peers slightly mitigates these effects, as observed by a final accuracy increase from approximately 0.70 (8 agents) to 0.86 (10 agents).

The Small-World topology (brown curves) presents an intermediate behavior. Its shortcut edges enable a faster initial rise, for instance, with 10 agents, accuracy approaches 0.90 already by round 30. However, these same shortcut links can allow certain highly connected "hub" agents to propagate biased or noisy updates more broadly. This leads to convergence settling at intermediate final scores (e.g., around 0.78 accuracy for 8

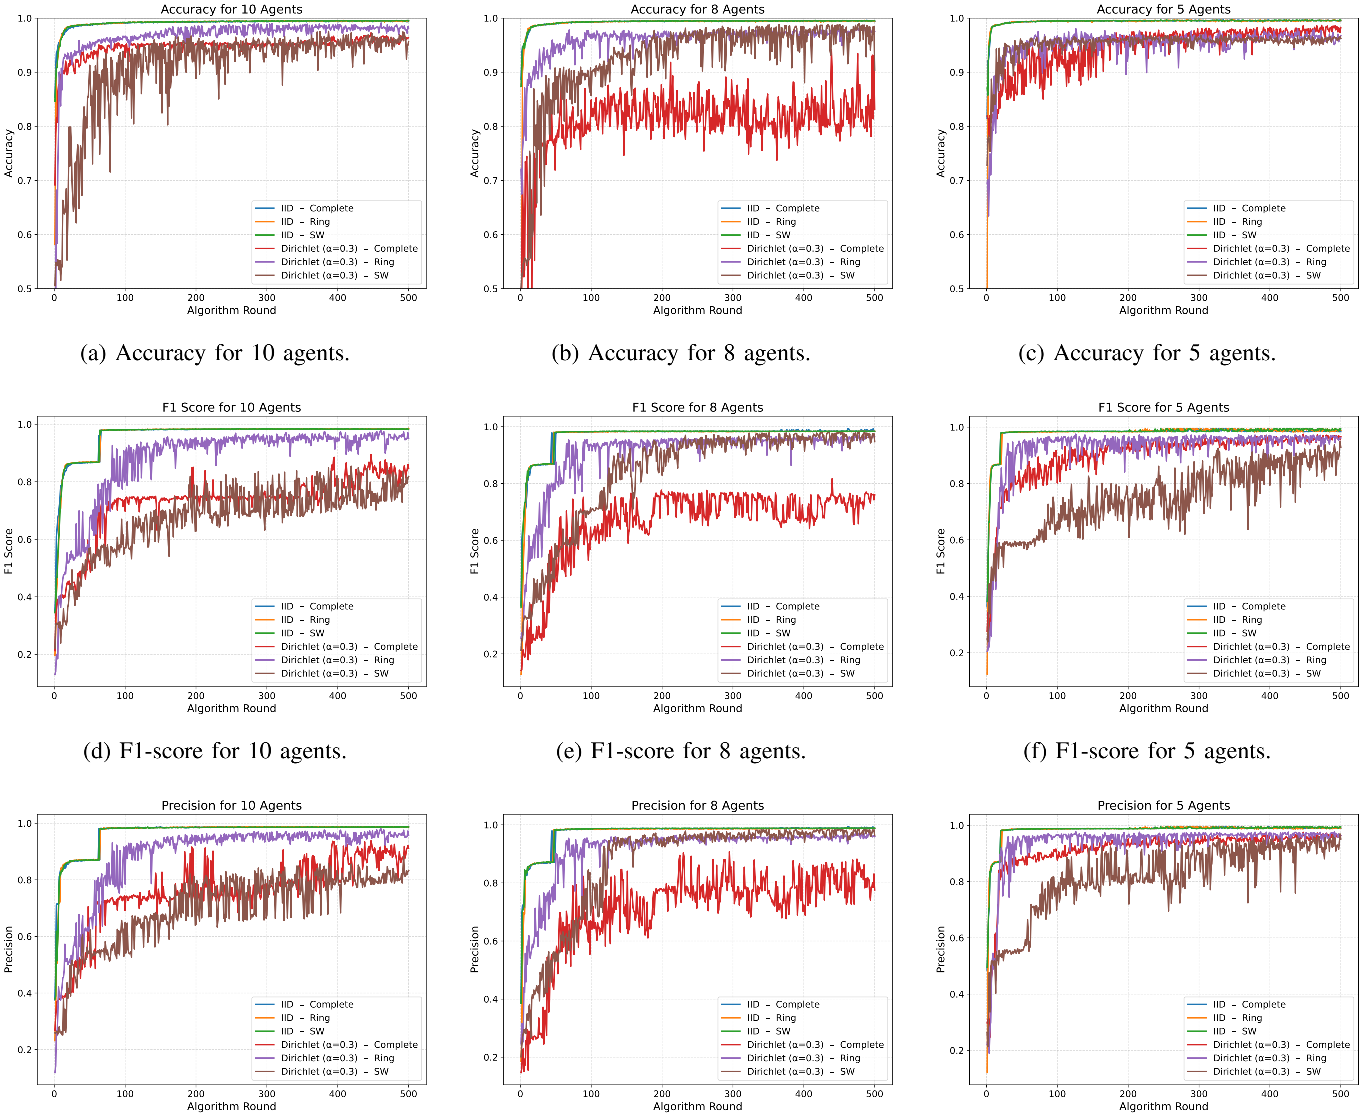

(a) Accuracy for 10 agents.

(b) Accuracy for 8 agents.

(c) Accuracy for 5 agents.

(d) F1-score for 10 agents.

(e) F1-score for 8 agents.

(f) F1-score for 5 agents.

(g) Precision for 10 agents.

(h) Precision for 8 agents.

(i) Precision for 5 agents.

Fig. 2: Accuracy, F1-score, and precision; across all agent configurations.

agents), reflecting a trade-off between accelerated mixing and robustness against local skew. When considering the effect of network size, increasing the number of agents from 5 to 10 slightly increases early-stage variance, especially in the Complete topology, where more potential sources of skew exist. However, in the Complete graph, this added diversity can also support moderate recovery in later stages. Meanwhile, the Ring topology stays very consistent no matter the network size, which points to strong resilience and scalability. Small-world topologies only gain a little from more agents.

From an engineering perspective, these observations highlight that in practical distributed environments such as 5G edge networks, where non-IID data is common and agents often operate with limited communication bandwidth, sparse and structured overlays like the Ring topology provide the most favorable balance between convergence stability and communication overhead. The Ring topology achieves high final F1 scores (above 0.93), stable learning curves, and minimal susceptibility to local data skew, even with minimal connectivity. In contrast, higher-density topologies like Complete or Small-World can be justified only under ultra-tight convergence requirements of application-level latency demands at favorable (IID) conditions. In addition, the insensitivity of terminal performance to the number of agents under IID conditions indicates that neighborhood edge clusters can scale relatively adaptively with minimal hyperparameter readjustment, a desirable aspect in dynamic and fast-expanding 5G FL scenarios.

V. Summary and Future Work

This paper proposes a distributed intrusion detection system for 5G networks, based on the synergetic combination of FL and MAS with the Asynchronous Consensus-based Learning protocol. The proposed solution meets the new requirements of 5G environments in terms of data heterogeneity, latency, and privacy-preserving and scalable threat detection. Our system proves excellent performance in various aspects through rigorous experimentation on the 5G-NIDD dataset. It obtains outstanding detection performance, lowers communication overhead, and retains robustness even in the case of non-IID data distribution and asynchronous update conditions. These results verify the feasibility of employing decentralized collaborative intelligence at the edge of the network to protect next-generation 5G infrastructures. Future work will explore integrating graph neural networks for richer modelling of inter-slice and inter-node traffic relationships and adding trusted execution environments to harden local training and update integrity.

VI. Acknowledgment

This work was supported partially by (a) the University of Zürich UZH, Switzerland, (b) the Horizon Europe Framework Program's project NATWORK, Grant Agreement No. 101139285, funded by the Swiss State Secretariat for Education, Research, and Innovation SERI, under Contract No. 23.00642, and (c) the grant PID2021-123673OB-C31, PRE2022-101563, funded by MICIU/AEI/10.13039/501100011033 and by "ERDF A way of making Europe".

References

[1] S. Samarakoon, Y. Siriwardhana, P. Porambage, M. Liyanage, S.-Y. Chang, J. Kim, J. Kim, and M. Ylianttila, "5G-NIDD: A Comprehensive Network Intrusion Detection Dataset Generated over 5G Wireless Network," *arXiv preprint arXiv:2212.01298*, 2022.

[2] H. B. McMahan and D. Ramage, "Federated Learning: Collaborative Machine Learning without Centralized Training Data," *Google AI Blog*, 6 Apr. 2017. [Online]. Available: https://ai.googleblog.com/2017/04/federated-learning-collaborative.html

[3] M. Wooldridge and N. R. Jennings, "Intelligent Agents: Theory and Practice," *The Knowledge Engineering Review*, vol. 10, no. 2, pp. 115–152, 1995.

[4] C. Carrascosa, A. Pico, M. M. Matagne, M. Rebollo, and J. A. Rincón, "Asynchronous Consensus for Multi-Agent Systems and Its Application to Federated Learning," *Engineering Applications of Artificial Intelligence*, vol. 135, Art. 108840, 2024.

[5] E. Hallaji, R. Razavi-Far, M. Saif, B. Wang, and Q. Yang, "Decentralized Federated Learning: A Survey on Security and Privacy," *arXiv preprint arXiv:2401.17319*, 2024.

[6] F. Adjewa, M. Esseghir, and L. Merghem-Boulahia, "Efficient Federated Intrusion Detection in 5G Ecosystem Using Optimized BERT-Based Model," *arXiv preprint arXiv:2409.19390*, 2024.

[7] N. Latif, W. Ma, and H. B. Ahmad, "Advancements in Securing Federated Learning with IDS: A Comprehensive Review of Neural Networks and Feature Engineering Techniques for Malicious Client Detection," *Artificial Intelligence Review*, vol. 58, Art. 91, Jan. 2025.

[8] F. Enguix, J. A. Rincón, and C. Carrascosa, "Introducing Coalitions to Improve the Performance of Federated Learning Consensus-Based Algorithms (ACoL)," in *Proc. Int. Conf. Practical Applications of Agents and Multi-Agent Systems (PAAMS)*, CCIS 2149, pp. 28–39, 2025.

[9] B. Wang, Z. Tian, J. Ma, W. Zhang, W. She, and W. Liu, "A decentralized asynchronous federated learning framework for edge devices," *Future Generation Computer Systems*, vol. 166, Art. 107683, 2025.

[10] S. Chennoufi, G. Blanc, H. Jmila, and C. Kiennert, "SoK: Federated Learning based Network Intrusion Detection in 5G: Context, State of the Art and Challenges," in *Proc. 19th Int. Conf. Availability, Reliability and Security (ARES)*, Vienna, Austria, 2024.

[11] J. Wu, F. Dong, H. Leung, Z. Zhu, and J. Zhou, "Topology-aware Federated Learning in Edge Computing: A Comprehensive Survey," *ACM Computing Surveys*, vol. 56, no. 10, pp. 1–41, 2024.

[12] A. Karunamurthy, K. Vijayan, P. R. Kshirsagar, and K. T. Tan, "An optimal federated learning-based intrusion detection for IoT environment," *Scientific Reports*, vol. 15, Art. 8696, 2025.

[13] A. Belenguer, J. A. Pascual, and J. Navaridas, "A review of federated learning applications in intrusion detection systems," *Computer Networks*, vol. 258, Art. 111023, 2025.

[14] S. M. S. Bukhari *et al.*, "Enhancing cybersecurity in Edge IIoT networks: An asynchronous federated learning approach with a deep hybrid detection model," *Internet of Things*, vol. 27, Art. 101252, 2024.

[15] F. Zhang, Y. Zhang, S. Ji, and Z. Han, "Secure and decentralized federated learning framework with non-IID data based on blockchain," *Heliyon*, vol. 10, no. 5, e27176, 2024.

[16] M. Rebollo and C. Carrascosa, "Multilayered Asynchronous Consensus-Based Federated Learning (MACoFL)," in *Intelligent Data Engineering and Automated Learning – IDEAL 2024*, LNCS 14452, pp. 386–396.

[17] Y. Xu, Z. Ma, H. Xu, S. Chen, and J. Liu, "FedLC: Accelerating asynchronous federated learning in edge computing," *IEEE Trans. Mobile Computing*, vol. 23, pp. 5327–5343, 2024.

[18] Z. Chen, D. Li, J. Zhu, and S. Zhang, "DACFL: Dynamic average consensus-based federated learning in decentralized sensor networks," *Engineering Applications of Artificial Intelligence*, vol. 135, Art. 108840, 2024.

[19] Z. Alsulaimawi, "Enhancing security in federated learning through adaptive consensus-based model update validation," *arXiv:2403.04803*, 2024.

[20] H. Liu, F. Zhu, and L. Cheng, "Proof-of-Data: A consensus protocol for collaborative intelligence," *arXiv:2501.02971*, 2025.

[21] X. Zhao, A. Wu, Y. Pei, Y.-C. Liang, and D. Niyato, "End-to-end verifiable decentralized federated learning," *arXiv:2404.12623*, 2024.

[22] J. Liu and M. Guo, "DIGNN-A: Real-time network intrusion detection with integrated neural networks based on dynamic graph," *Computers, Materials & Continua*, vol. 82, no. 1, pp. 817–842, 2025.

[23] H.-D. Le and M. Park, "Enhancing multi-class attack detection in graph neural networks through feature rearrangement," *Electronics*, vol. 13, no. 12, Art. 2404, 2024.

[24] H. Nguyen and R. Kashef, "TS-IDS: Traffic-aware self-supervised learning for IoT network intrusion detection," *Knowledge-Based Systems*, vol. 279, Art. 110966, 2023.

[25] L. Li, X. Zhang, Y. Wang, and C. Chen, "Locally differentially private online federated learning with limited communication," *arXiv:2411.18752*, 2024.

[26] W. W. Lo *et al.*, "E-GraphSAGE: A graph neural network-based intrusion detection system for IoT," in *Proc. IEEE/IFIP NOMS*, 2022, pp. 1–9.

[27] A. Giuseppi, S. Manfredi, and A. Pietrabissa, "A weighted average consensus approach for decentralized federated learning," *Machine Intelligence Research*, vol. 19, no. 4, pp. 319–330, 2022.

[28] M. Chahoud, S. Otoum, and A. Mourad, "On the feasibility of federated learning towards on-demand client deployment at the edge," *Information Processing & Management*, vol. 60, no. 1, Art. 103150, 2023.

[29] A. S. M. Tayeen *et al.*, "CAFNet: Compressed autoencoder-based federated network for anomaly detection," in *Proc. IEEE MILCOM*, 2023, pp. 325–330.

[30] X. Sáez-de-Cámara *et al.*, "Clustered federated learning architecture for network anomaly detection in large-scale heterogeneous IoT networks," *arXiv:2303.15986*, 2023.

[31] M. Nivaashini *et al.*, "FEDDBN-IDS: Federated deep belief network-based wireless network intrusion detection system," *EURASIP J. Information Security*, vol. 2024, Art. 22.

[32] E. Gelenbe, B. C. Gül, and M. Nakıp, "DISFIDA: Distributed self-supervised federated intrusion detection algorithm with online learning for health Internet of Things and Internet of Vehicles," *Internet of Things*, vol. 28, Art. 101340, 2024.

[33] P. Kumar, J. Liu, A. S. Md Tayeen, S. Misra, H. Cao, J. Harikumar, and O. Perez, "FLNET2023: Realistic Network Intrusion Detection Dataset for Federated Learning," in *Proc. IEEE Military Communications Conference (MILCOM)*, 2023, pp. 345–350.