Mitigating Communications Threats in Decentralized Federated Learning through Moving Target Defense

Enrique Tomás Martínez Beltrán[1*], Pedro Miguel Sánchez Sánchez[1],
Sergio López Bernal[1], Gérôme Bovet[2], Manuel Gil Pérez[1],
Gregorio Martínez Pérez[1], Alberto Huertas Celdrán[3]

[1]Department of Information and Communications Engineering, University of Murcia, 30100, Spain.
[2]Cyber-Defence Campus, Armasuisse Science and Technology, 3602, Thun, Switzerland.
[3]Communication Systems Group, Department of Informatics (IFI), University of Zurich, 8050, Zürich, Switzerland.

*Corresponding author(s). E-mail(s): enriquetomas@um.es;
Contributing authors: pedromiguel.sanchez@um.es; slopez@um.es;
gerome.bovet@armasuisse.ch; mgilperez@um.es; gregorio@um.es; huertas@ifi.uzh.ch;

Abstract

The rise of Decentralized Federated Learning (DFL) has enabled the training of machine learning models across federated participants, fostering decentralized model aggregation and reducing dependence on a server. However, this approach introduces unique communication security challenges that have yet to be thoroughly addressed in the literature. These challenges primarily originate from the decentralized nature of the aggregation process, the varied roles and responsibilities of the participants, and the absence of a central authority to oversee and mitigate threats. Addressing these challenges, this paper first delineates a comprehensive threat model focused on DFL communications. In response to these identified risks, this work introduces a security module to counter communication-based attacks for DFL platforms. The module combines security techniques such as symmetric and asymmetric encryption with Moving Target Defense (MTD) techniques, including random neighbor selection and IP/port switching. The security module is implemented in a DFL platform, Fedstellar, allowing the deployment and monitoring of the federation. A DFL scenario with physical and virtual deployments have been executed, encompassing three security configurations: (i) a baseline without security, (ii) an encrypted configuration, and (iii) a configuration integrating both encryption and MTD techniques. The effectiveness of the security module is validated through experiments with the MNIST dataset and eclipse attacks. The results showed an average F1 score of 95%, with the most secure configuration resulting in CPU usage peaking at 68% ($\pm$9%) in virtual deployments and network traffic reaching 480.8 MB ($\pm$18 MB), effectively mitigating risks associated with eavesdropping or eclipse attacks.

Keywords: Decentralized Federated Learning, Decentralized Network, Cyberattack Mitigation, Moving Target Defense

1

1 Introduction

The rise of the Internet of Things (IoT) has significantly reshaped the digital landscape, defining an era marked by unprecedentedly interconnected devices. IoT devices produce vast volumes of data every second, spanning various sectors, from healthcare and manufacturing to transportation and home automation. Traditionally, Machine Learning (ML) techniques have been employed to derive meaningful insights from these large datasets. However, these techniques often involve the centralized aggregation of data, a process that raises serious concerns about data privacy, data sovereignty, and overhead [1].

A novel ML approach, known as Federated Learning (FL), has emerged in response to these challenges. FL can train models locally on multiple edge devices, each holding local data samples. This eliminates the need to share raw data, thereby preserving data privacy. Advancing this concept, Decentralized Federated Learning (DFL) represents a paradigm shift within FL [2]. DFL strengthens decentralization by enabling the aggregation of models across multiple nodes, thereby substantially reducing reliance on a centralized server. This advancement not only preserves privacy but also enhances system scalability, robustness, and efficiency, making it particularly suitable for distributed IoT applications. DFL integrates several key processes: (1) participants train models on their edge devices using local data, preserving data privacy; (2) nodes then directly exchange the parameters of their models in pairs, which favors a decentralized network structure; and (3) each node integrates these shared parameters into their local models, resulting in an aggregated and refined model that benefits from the diverse data insights from across the network. This innovative approach addresses single points of failures, trust dependencies, and server node bottlenecks inherent in traditional FL. DFL also eliminates the need for a central server by broadening the model aggregation to multiple nodes. Additionally, DFL employs asynchronous communications, a departure from traditional FL. This feature enables individual nodes to communicate their updates independently of others, contributing to system resilience and ensuring the continued learning process even if some nodes

encounter delays or disconnections [3]. The application of DFL to wireless networks has been motivated by the resilience offered by its asynchronous communication, which is crucial in environments with intermittent and unpredictable connectivity [4]. Specifically, such traits make DFL highly applicable for Unmanned Aerial Vehicle (UAV) networks, where constant and reliable communication is often challenged by diverse factors such as terrain and weather conditions, hence enhancing their cooperative missions [5].

Despite the substantial benefits of DFL, it also introduces new challenges. This approach poses different types of sensitive information necessary for the federation, such as the network topology, the roles of the participants, and communication patterns that can be exploited. Besides, in DFL environments where all participants are connected, the absence of a central authority to manage potential threats raises significant security and privacy concerns [6]. With each participant sharing equal threat exposure, adversarial and communication-based attacks become significant concerns. Adversarial attacks can misguide the learning process by manipulating training data or leveraging the shared model updates to infer sensitive information about the other participants. At the same time, communication-based threats can disrupt the model aggregation process or lead to security breaches and privacy infringements [7]. Addressing these challenges could benefit from adopting a dynamic approach like Moving Target Defense (MTD) [8]. MTD is a security concept that continuously alters attack surfaces to confuse and mislead adversaries, making it difficult for them to launch successful attacks. The potential integration of MTD with encryption in DFL offers a novel approach to enhancing security, particularly in the face of unique challenges in decentralized architectures. This strategy is particularly relevant in DFL, where the decentralized nature of data exchange and interaction presents distinct challenges not adequately addressed by traditional security methods. Combining dynamic MTD techniques with strong encryption proposes an advanced defense against vulnerabilities and threats unique to these systems. Moreover, the literature has not extensively addressed specific attacks within DFL environments, highlighting the need for this innovative integration. Such an

approach underscores the need for innovative solutions tailored to environments with distributed architectures. In recognition of the risks in DFL, and with a special emphasis on communication-based attacks that leverage the inherent decentralization of DFL, this paper presents the following contributions:

- Create a threat model, identifying and understanding the sensitive information vulnerable to threats affecting the communications in DFL, such as eavesdropping, Man in the Middle (MitM), and eclipse attacks.
- Develop an advanced security module for DFL platforms providing secure data exchanges through encryption and dynamic proactive defense using MTD. This module mitigates the threats identified in the comprehensive threat model of DFL, ensuring efficient system operation despite the integrated security measures.
- Implement and deploy the security module within a real-world DFL framework, Fedstellar, integrating it into the frontend, controller, and core components of the platform to enhance the overall security of the DFL approach. Furthermore, this work implements a dual DFL environment using the Fedstellar platform. The initial deployment comprises a physical network of eight heterogeneous devices. Additionally, a virtual deployment with 50 participants facilitates a comprehensive and scalable evaluation of DFL performance. Three security configurations are assessed in both setups: a baseline with no security, a configuration with encryption, and a configuration integrating both encryption and MTD techniques.
- Conduct an in-depth experimental evaluation of the proposed security module using a real-world topology with diverse connections and participants, leveraging the widely used MNIST dataset and a custom implementation of an eclipse attack. The evaluation across both physical and virtual deployments reveals an average F1 score of 95%, which ascends to 98.9% in the absence of security measures. Implementing secure configurations, particularly those utilizing encryption and MTD, leads to an increase in CPU usage, reaching up to 68% ($\pm$9%) in virtual

environments. In addition, the network traffic peaks at 480.8 MB ($\pm$18 MB), while the RAM usage also experiences a moderate rise, with a maximum of 35.9% ($\pm$1.5%) noted in the physical deployment under encryption and MTD settings.

The remainder of this paper is organized as follows: Section 2 provides an in-depth overview of the literature on DFL and its associated security challenges. Section 3 introduces the proposed threat model, highlighting the unique security issues that DFL environments face. Section 4 presents a detailed description of the proposed security module, elucidating its key components and their functionality. Section 5 outlines the experimental setup and evaluation methodology, paving the way for a rigorous assessment of the effectiveness of the security module. Section 6 presents a comprehensive discussion of the results, and Section 7 concludes the paper with a summary of the key findings and an exploration of potential avenues for future research.

2 Related Work

This section gives the insights required to understand the concepts used in the following sections and reviews the main works in the literature associated with the present one.

2.1 Privacy and security in DFL

The promise of DFL as a tool for collaborative learning in heterogeneous and geographically distributed settings continues to drive robust research into its inherent security implications. A comprehensive understanding of its potential threats and appropriate countermeasures enhances cooperative learning practices. Several ground-breaking research efforts have focused on integrating trust within a DFL context. In this regard, Gholami et al. [9] proposed an approach that integrates trust as a metric within a DFL context. Their method used a comprehensive mathematical framework to quantify and aggregate the trustworthiness of individual agents. In parallel, Mothukuri et al. [10] addressed anomaly detection in Internet of Things (IoT) networks by leveraging the distributed nature of FL. They proposed a FL methodology that optimized anomaly detection by aggregating updates from diverse sources.

Their approach hinged on using gated recurrent units (GRUs) in federated training rounds to maximize the accuracy of the overall ML model. Complementing these advancements, Li [11] took an innovative leap by proposing a Trustiness-based Hierarchical Decentralized FL (TH-DFL) framework. It employs a Security Robust Aggregation (SRA) rule to ensure privacy and robustness even in the face of malicious nodes. The TH-DFL framework strikes an optimal balance between privacy and robustness, especially as the group size fluctuates, and exhibits superior resilience against varying forms of attacks.

Security concerns related to jamming attacks have also been extensively studied, especially in wireless networks implementing DFL. Shi et al. [3] shed light on the susceptibility of DFL to these attacks, proposing crucial countermeasures. Their algorithms identify and target pivotal network links for attack prevention and optimal placement of jammers to disrupt the federation process. Their findings point to the urgency for sophisticated defense mechanisms in DFL architectures. Further contributing to the body of knowledge on security threats in DFL, Chen et al. [12] proposed a method called Decentralized FL Historical Gradient (DFedHG). DFedHG utilizes historical gradients to differentiate between regular, untrusted, and malicious users in a DFL environment. This unique solution strengthens the defense against potential threats in DFL systems, accentuating the necessity for sturdy security frameworks.

Securing wireless networks while implementing DFL is a topic of intensive research. Wang [13] introduced a method to ensure the security and efficiency of FL in Wireless Computing Power Networks (WCPNs). Their research presents a secure and decentralized FL solution based on blockchain for WCPN, which allows nodes to freely participate or leave the WCPN federated training without authorization and security threats. This approach uses a blockchain with a proof-of-accuracy (PoAcc) consensus scheme and an evolutionary game-based incentive scheme to ensure the consistency and security of FL in WCPN. On the other hand, Salama [4] proposed a method for Decentralized FL over Slotted ALOHA Wireless Mesh Networking. The approach offers an efficient solution for ML model training without a central server, reducing communication costs and increasing convergence speed. This paper demonstrates how network topologies can impact the performance of ML models, and their results indicate significant promise for DFL in Internet of Things (IoT) systems.

2.2 Security-based DFL solutions

Innovative approaches toward enhancing data protection and secure communication within DFL environments have also seen considerable development. For instance, the FusionFedBlock solution, proposed by Singh et al. [14], merges the strengths of blockchain and DFL to ensure privacy in Industry 5.0. A distributed hash table (DHT) guarantees secure decentralized storage at the cloud layer, while blockchain miners facilitate data verification. FL-SEC, introduced by Qu et al. [15], stands as a breakthrough framework that addresses potential information leakage due to inference attacks, threats of poisoning attacks via falsified data, and high consumption of communication resources. This model uses a custom incentive mechanism and an enhanced sign gradient descent method to protect the privacy of model parameters and significantly reduce communication resource consumption. Contributing further to privacy preservation and trustworthiness in DFL, Wang [16] proposed PTDFL, an efficient and novel DFL scheme. This scheme integrates a gradient encryption algorithm to protect data privacy, employs concise proof for the correctness of the gradients, and uses a local aggregation strategy to ensure that the aggregated result is trustworthy. The unique feature of PTDFL is its support for data owners joining in and dropping out during the entire DFL task.

In the enterprise domain, Arakapis et al. [17] introduced P4L, a private peer-to-peer learning system. As an asynchronous collaborative learning scheme, P4L allows users to participate in the learning process without depending on a centralized infrastructure. It ensures the confidentiality and utility of shared gradients employing strong cryptographic primitives. Also, it maintains resilience to user dropout and fault tolerance, highlighting the practical applicability and effectiveness of decentralized learning solutions in real-world settings. Finally, on the frontier of sixth-generation (6G) networks, Ridhawi et al. [18] proposed a decentralized zero-trust framework for digital twins. By integrating the zero-trust

architecture into digital twin-enabled networks with DFL, they ensured the security, privacy, and authenticity of physical and digital devices. Their approach addresses the challenges of cooperation between devices and network components in a 6G environment, demonstrating the pivotal role of DFL in next-generation networks.

3 Communications Threat Model in DFL

The threat model primarily focuses on the communication aspects of DFL, presuming the co-existence of trusted participants who abide by network protocols and malicious participants who pose multilayered threats. The threat landscape in the communication channels of a DFL environment is complex, with malicious entities potentially playing passive or active roles. Passive malicious entities might eavesdrop on network communications, surreptitiously gaining access to sensitive information such as model parameters, aggregated gradients, or participants' metadata. In contrast, active malicious entities could actively interfere with network operations, manipulate data, introduce false information, or disrupt communication channels. These threats can originate from internal and external sources, with internal threats emerging from compromised or malicious network participants and external threats from entities outside the DFL topology.

As detailed in Table 1, a malicious participant in a DFL environment can extract a wide range of sensitive information, each bearing unique implications implications. A notable example is the extraction of model parameters, such as weights and biases from each neural network layer, which encapsulate the learned knowledge of the model. Although methods like Homomorphic Encryption or Differential Privacy may impede or obscure this extraction, the underlying threat parallels that in FL. Unauthorized access to these parameters could allow a malicious entity to reconstruct the learning model, leading to significant data privacy violations and potentially exposing critical insights. Additionally, the network's topology provides valuable information about its structure and interactions, offering adversaries insights that could facilitate more targeted attacks.

Additionally, the assigned roles within a DFL network can provide an adversary with a detailed understanding of the functional distribution and control mechanisms. Unlike in FL vanilla, where all clients primarily hold the same role, this aspect of DFL architecture can aid an attacker in identifying which nodes to target for maximum disruption. Moreover, performance metrics and resource usage data could expose system vulnerabilities regarding performance and resource allocation strategies. An attacker might infer these metrics from the patterns and volume of network communications [19]. Information about participant activity periods and the underlying model architecture could prove invaluable for an attacker. By analyzing communication timings and frequencies, an attacker might discern when specific nodes are most active, providing insights into the operational rhythms of the network. A deep understanding of the model architecture, obtained through careful observation of network interactions and data exchanges, might expose the structure and operational logic of the model, thereby revealing potential weaknesses for exploitation. Finally, understanding communication patterns could prove beneficial for a malicious entity. By examining the frequency and nature of participant interactions, an attacker could identify critical patterns, anticipate behaviors, and potentially impersonate trusted nodes to gain unauthorized access or disrupt the network.

Numerous potential security threats can compromise the confidentiality, integrity, and availability of federated data and models. These threats primarily arise from the inherent vulnerabilities presented by the decentralization of learning processes and model sharing without the control of a central authority. The following communications threats have been identified (see Table 2):

- *TH1. Eavesdropping.* In a DFL setting, an adversary could covertly monitor network communications or infiltrate a participant node to gain unauthorized access to sensitive data. This data could include model parameters, network topology, and participant roles. The adversary could then leverage this information to disrupt the federated process or impersonate a legitimate participant. This threat often persists undetected due to its

Table 1: Information accessible to a malicious participant in DFL

Information	Description
Model Parameters	Each layer l_i in a model M with n layers has weight $w_i \in \mathbb{R}^{d_i \times d_{i-1}}$ and bias $b_i \in \mathbb{R}^{d_i}$, where d_i is the number of neurons in layer i. The parameters of M are the collection $\{w_i, b_i\}_{i=1}^n$.
Topology	The graph of the network $G(V, E)$, where V is the set of vertices (participants) and E is the set of edges (connections). If $V = \{v_1, v_2, ..., v_n\}$ and $E = \{(v_i, v_j)|v_i, v_j \in V, i \neq j\}$, the topology is fully connected.
Roles	Each participant $p_i \in V$ has a role $r_i \in \{\text{idle, trainer, aggregator, proxy}\}$. This can be mathematically represented by a function $R : V \to \{\text{idle, trainer, aggregator, proxy}\}$, where $R(p_i) = r_i$.
Metrics	Performance of the model (e.g., accuracy, precision, recall, F1 score) and resource usage (CPU, RAM, network) of the nodes. For resources, let R be the resource, U_R the usage, and C_R the capacity. The usage rate is $R_{rate} = \frac{U_R}{C_R}$.
Activity Periods	If $T = \{t_1, t_2, ..., t_n\}$ represent the set of all time intervals and $A = \{a_1, a_2, ..., a_k\} \subseteq T$ the active intervals, then the activity ratio is $A_{ratio} = \frac{\sum_{i=1}^k a_i}{\sum_{i=1}^n t_i}$.
Model Architecture	A feedforward neural network with n layers can be represented as a sequence of function compositions $f(x) = f_n(f_{n-1}(...f_2(f_1(x))))$, where $f_i(x) = \sigma(w_i \cdot x + b_i)$ is the operation for layer i, and σ is the activation function.
Communication Patterns	If $M = \{m_{ij}\}$ is the set of all messages sent from participant i to participant j, the frequency of communication between these participants can be quantified as $F_{ij} = \frac{|m_{ij}|}{\sum_{i,j} |m_{ij}|}$, where $|m_{ij}|$ is the number of messages exchanged.

Table 2: Attacks, goals, and information at risk in DFL

Attack	Goal	Information at Risk
Eavesdropping	Extract sensitive information to undermine integrity and security of the federated participants [!!]	• Model Parameters • Topology • Roles
MitM	Manipulate information or insert malicious data to disrupt federation operations [!!!]	• Communication Patterns • Roles
Network Mapping	Know the network structure to launch more targeted future attacks on the federation [!]	• Topology • Model Architecture
Eclipse Attacks	Isolate a node or group of nodes to extract information or disrupt DFL communications [!!!]	• Activity Periods • Topology • Roles • Communication Patterns

! Low importance, !! High importance, !!! Critical

covert nature, leading to prolonged periods of sensitive data leakage.

- *TH2. MitM.* It involves an attacker intercepting and potentially manipulating the communication between two participant nodes. This enables the attacker to alter exchanged model parameters, introduce spurious data, or eavesdrop on the exchanged information, posing significant challenges to the integrity of the federated process.

- *TH3. Network Mapping.* It aims to understand the structure of the federated network and the roles of participant nodes. By gaining this knowledge, attackers can predict and interfere with network operations, facilitating more targeted and potentially detrimental exploits.

- *TH4. Eclipse.* This attack in DFL seeks to isolate a specific node or a group of nodes from the rest of the network. This isolation distorts the affected nodes' perception

of the network state, causing them to act based on inaccurate information and potentially paving the way for additional security breaches.

In light of the identified threats, a comprehensive security module for DFL must account for these potential attack vectors and implement countermeasures to ensure robust operation and resilience against attacks. Crucially, achieving this goal involves striking a careful balance between enhancing security and managing the additional computational and network overhead that security measures may introduce.

4 Security Module

This section details the proposed security module, particularly examining its integration within a novel DFL platform and how it fortifies the network against a broad spectrum of cyber threats.

4.1 Overview

The security module comprises a set of cybersecurity strategies designed to safeguard the complex exchange of data and models in DFL. The distinctive features of DFL, such as decentralized aggregation, asynchronous communication, limited visibility to near neighbors, and participant independence, necessitate nuanced and versatile security measures. The limited visibility of DFL nodes, usually only to immediate neighbors, restricts the broader network anomaly detection. Participant independence complicates maintaining a secure environment as nodes decide when to commence model training or aggregation. This proposal responds to the growing need for advanced security mechanisms within the field of DFL, considering the diversity and sensitivity of data involved in these systems. This module employs sophisticated encryption methods and MTD techniques, making it highly adaptable to various DFL platforms:

- *Encryption.* Using a combination of symmetric and asymmetric encryption, the module ensures secure model exchanges and efficient key management. This strategy guarantees data confidentiality and provides robust protection against potential breaches.
- *MTD Techniques.* These techniques, which include Neighbor Selection and IP/port

switching, create a dynamic and unpredictable defensive layer within the system. By continuously changing communication pathways and nodes, these techniques make it increasingly difficult for potential attackers to gain a foothold in the system.

4.2 Security Components

The components of the security module comprise encryption techniques and MTD strategies. The encryption techniques, designed to ensure data confidentiality during the model exchange, combine the efficiency of symmetric encryption for data protection with the secure key management of asymmetric encryption. MTD techniques, such as Neighbor Selection and IP/port switching strategies, add a dynamic and shifting defensive layer to the system. These techniques introduce unpredictability and fluidity by continuously altering network communication pathways, making the system difficult for potential attackers to decipher due to the increased complexity and resource requirements for successful attacks. The integration of these components in a federated participant cycle within a DFL environment is depicted in Algorithm 1. This algorithm combines the elements of encryption and MTD, effectively creating a robust security layer within the DFL infrastructure.

4.2.1 Communications Encryption

The integrity and confidentiality of the information exchanged among participants during the federation is a fundamental requirement in secure DFL systems. This security is achieved by combining symmetric and asymmetric encryption techniques, forming a comprehensive, multi-layered security infrastructure.

The first layer of this security architecture employs symmetric encryption. This method is computationally efficient and uses a single key for data encryption and decryption. The Advanced Encryption Standard (AES) algorithm, provided by the *pycryptodome* library, is utilized for this layer. Known for its robust security and broad acceptance, the AES algorithm is an ideal choice, especially considering the resource constraints often present in many devices.

Algorithm 1 Federated participant cycle with Encryption and MTD Techniques in DFL

Require: R: local round, α: learning rate, λ: regularization parameter, S_j: socket to neighbor j, D: local dataset, E_{sym} / E_{asym}: symmetric/asymmetric encryption function, D_{sym} / D_{asym}: symmetric/asymmetric decryption function, MTD_{IP}: IP/port MTD function, MTD_{N}: neighbor selection MTD function
1: **procedure** $MTD_{\text{N}}(N_{\text{all}}, n)$ ▷ **Neighbor Selection**
2:    Initialize an empty list N
3:    **while** $|N| < n$ **do**
4:      Select a neighbor i from N_{all} uniformly at random
5:      **if** $i \notin N$ **then**
6:        Add i to N
7:      **end if**
8:    **end while**
9:    **return** N
10: **end procedure**
11: **procedure** $MTD_{\text{IP}}(config)$ ▷ **IP/Port Switch**
12:    Fetch a list of available IP and ports: $IP_{\text{avail}}, P_{\text{avail}}$
13:    Select a new IP and port from $IP_{\text{avail}}, P_{\text{avail}}$ uniformly at random
14:    Update $config$ with the new IP address and port
15:    **return** $config$
16: **end procedure**
17: $D_{\text{Train}}, D_{\text{Test}} \leftarrow split(D)$
18: **for** r in R **do**
19:    $\theta \leftarrow Initialize()$ ▷ **Initialize Local Model**
20:    **for** each (x, y) in D_{Train} **do**
21:      $\theta \leftarrow \theta - \alpha(\nabla_\theta J(\theta, x, y) + \lambda\theta)$ ▷ **Train**
22:    **end for**
23:    $N \leftarrow MTD_{\text{N}}(N_{\text{all}})$
24:    **for** j in N **do**   ▷ **Send**
25:      $\theta_{\text{enc}} \leftarrow E_{\text{sym}}(\theta, K_{\text{sym}})$
26:      $K_{\text{sym\_enc}} \leftarrow E_{\text{asym}}(K_{\text{sym}}, K_{j_{\text{pub}}})$
27:      Send $\theta_{\text{enc}}, K_{\text{sym\_enc}}$ to j via S_j
28:    **end for**
29:    **while** $not\ Timeout$ **do**
30:      **for** j in N **do**   ▷ **Receive**
31:        $RP_{j_{\text{enc}}}, K_{j_{\text{sym\_enc}}} \leftarrow$ Receive from j via S_j
32:        $K_{j_{\text{sym}}} \leftarrow D_{\text{asym}}(K_{j_{\text{sym\_enc}}}, K_{\text{priv}})$
33:        $RP_j \leftarrow D_{\text{sym}}(RP_{j_{\text{enc}}}, K_{j_{\text{sym}}})$
34:      **end for**
35:    **end while**
36:    $\theta \leftarrow \frac{1}{|N|+1}(\theta + \sum_{j \in N} RP_j)$ ▷ **Aggregate** (FedAvg)
37:    Update Local Model with θ
38: **end for**
39: **for** each (x, y) in D_{Test} **do**
40:    $y_{pred} \leftarrow Predict(\theta, x)$ ▷ **Test**
41:    $L \leftarrow \frac{1}{|D_{\text{Test}}|} \sum_{i=1}^{|D_{\text{Test}}|} l(y_i, y_{pred_i})$ ▷ **Compute Loss**
42: **end for**
43: Send metrics to controller   ▷ **Report Metrics**
44: $MTD_{\text{IP}}(config) \rightarrow config$

The second layer of the security architecture employs asymmetric encryption. This technique provides an additional layer of security by using a pair of keys: a public key for encryption and a private key for decryption. The RSA algorithm, also provided by the *pycryptodome* library, is used for this layer. RSA eliminates risks associated with key sharing in symmetric encryption and ensures a secure channel for key exchange, protecting the symmetric keys used in the AES algorithm. Key distribution and management are central to this interconnected system, facilitated by the controller, which acts as a secure Key Distribution Center (KDC). Upon deployment, each node is authenticated by the controller (see Section 4.3) and issued digital certificates. This process underpins the trust and integrity of the public keys disseminated within the network. Moreover, the controller dynamically manages public key updates, scheduling regular key renewals in line with security protocols to swiftly address potential vulnerabilities.

4.2.2 MTD Techniques

The MTD techniques serve to obfuscate and alter the attack surface dynamically, posing a significant challenge for attackers attempting to exploit system vulnerabilities. The proposed security module incorporates two MTD techniques: Neighbor Selection and IP/port switching.

The Neighbor Selection MTD technique minimizes network topology exposure to potential attackers. This technique can protect the nodes from targeted attacks by dynamically altering their communication partners in each learning cycle. By continually shifting the communication patterns in the network, the likelihood of an attacker successfully predicting or manipulating these patterns is significantly reduced. The random selection of neighbors is implemented using Python's built-in random library, ensuring unbiased and unpredictable selections for each cycle. The process for the Neighbor Selection MTD is fairly straightforward. In each federated round, a node randomly selects a subset of neighbors from all available participants (see Algorithm 1). This selection scheme is implemented using the socket library of Python, which provides low-level networking capabilities suitable for various network protocols, including TCP/IP, common in wired and wireless communications. The socket-based communication scheme offers reliability and flexibility, which are vital in a dynamic DFL environment.

The second technique is IP/port switching MTD, adding another layer of security. This method involves routinely changing the IP addresses and ports used by the federated nodes, further complicating the predictability of the attack surface. An attacker finds it difficult to sustain a prolonged attack on a specific node.

In the proposed security module, IP/port switching is implemented by regularly selecting a new IP address and port from a pool of available ones. This selection is automated and randomized using the built-in capabilities of Python for network configuration. By dynamically altering the IP addresses and ports, the technique disrupts potential attackers' ability to predict the communication structure or execute targeted attacks.

Both techniques need to ensure uninterrupted and secure communication amid IP and port changes employing a rendezvous mechanism. To achieve this, the system implements a predictive notification mechanism. Before a node switches its network configuration, it broadcasts its neighbors an encrypted notification containing the new connection details. This notification, encrypted with the network's standard encryption protocols, allows each recipient node to update its records before the change. This decentralized approach eliminates the need for a real-time directory service and instead relies on the timely dissemination of IP/port updates directly between nodes. As a result, even when an IP address or port changes, the communicating nodes can independently reconcile the new configurations, thereby maintaining uninterrupted and secure connections. This method adheres to the principles of a decentralized network and reinforces the security infrastructure, ensuring the network resolution process remains robust against potential vulnerabilities.

Building on the elaboration of the implemented security techniques, it is essential to understand their effectiveness, as depicted in Table 3. Encryption protects against eavesdropping, MitM, and eclipse attacks by protecting data during transmission. As a complement, MTD offers robust defenses against attacks such as Network Mapping or eclipse attacks.

Table 3: Potential mitigations for attacks in DFL

Security Components	Attacks			
	Eavesdropping	MitM	Network Mapping	Eclipse
Encryption	✓	✓	✗	✓
MTD	✗	✓	✓	✓

4.3 Fedstellar Platform

Fedstellar is an innovative platform that facilitates the training of FL models across a wide array of physical and virtual devices [19]. The platform is a hub for developing, deploying, and managing federated applications and provides a standardized approach for executing these processes. The architecture of Fedstellar is composed of three fundamental elements:

- *Frontend*. A user-centric interface that offers easy experiment setup and real-time monitoring, thus ensuring an intuitive user experience.
- *Controller*. A central command unit orchestrates operations across the platform, ensuring seamless inter-module communication and efficient task execution.
- *Core*. This critical component, deployed on each participating device, is responsible for vital functions such as model training and communication.

These components establish a robust and resilient architecture that provides sophisticated tools and metrics for federation management. This enables high transparency and efficiency in monitoring the learning process. Moreover, the platform contains extensible modules offering data storage, asynchronous capabilities, and effective model training and communication mechanisms.

The security module is integrated into the Fedstellar platform to demonstrate the proposed effectiveness and compatibility of the module. As depicted in Figure 1, the *security module* is a pivotal functionality of the core component responsible for managing secure communications across the platform. Its integration into the core ensures robust protection for the vast and complex communication exchanges characteristic of DFL. To support the overall security structure, enhancements have also been made to the frontend and the controller components of the Fedstellar platform.

In this sense, the frontend encompasses the *security definition* feature, enabling users to set and manage their security parameters conveniently. Conversely, the controller implements *security measures*, a provision that efficiently manages and enforces the established security settings in real time. Also, it incorporates a participant authentication process based on JSON Web Tokens (JWT) during network deployment,

conducted under encrypted communication (see Section 4.2.1). Upon joining the network, each node requests a token from the controller by providing its credentials. The controller validates these credentials and issues a JWT, which the node then uses for all subsequent communications within the network. This token, containing encrypted identity and permission information, ensures that only authenticated nodes participate in the network, enhancing security and preventing unauthorized access. The tokens have a limited lifespan, requiring nodes to periodically re-authenticate, thus maintaining ongoing network integrity.

The integration of the security module maintains compatibility through its design, which leverages threaded processing for non-blocking operations and event passing between modules for effective communication. These provisions ensure that the addition of the module does not disrupt the existing functionalities of the platform but rather harmonizes with them, augmenting the capability of Fedstellar to efficiently manage diverse federations comprising various devices, network topologies, and algorithms.

5 Validation Scenario

The validation scenario of this study emulates an edge computing setting, which evaluates the performance of the proposed security module in a DFL environment. The validation was conducted in two distinct deployments: physical and virtual. The physical deployment, as detailed in Table 4, encompasses a federation of eight physical devices: five Raspberry Pi 4 units and three Rock64 units. These devices are interconnected via a random network topology within the private local network. This topology is designed to mimic dynamic real-world environments, where connections between devices vary, offering insights into the federated process under fluctuating network conditions. The Raspberry Pi 4 units, armed with a 1.5GHz quad-core 64-bit ARM Cortex-A72 CPU and 2GB of RAM, present a delicate balance between size, cost-effectiveness, and computational prowess, thereby rendering them a suitable choice for simulating edge nodes. The remaining three devices, Rock64 boards, enhance the system's heterogeneity by contributing slightly lower processing capabilities, characterized by a 64-bit

ARM Cortex-A53 with a 1.5 GHz clock speed and up to 2GB RAM. To showcase the scalability of the solution, the experiment incorporates a virtual deployment comprising 50 Docker containers. Each container is configured to replicate the processing power and memory capacity of the physical devices. This expanded configuration offers a comprehensive testbed for evaluating the scalability and security module in a more complex DFL network. The physical and virtual deployment is conducted on the Fedstellar platform, specifically designed to facilitate FL experiments. Within this platform, each participating node employs the LeNet5 neural network architecture. The choice of LeNet5 is strategic, given its relatively simple structure that allows for quick training and inference, thus suitable for DFL across devices with varying computational capabilities. The MNIST dataset is utilized to train and validate the federated models. Comprising 70,000 handwritten digits, MNIST provides a balanced and comprehensive dataset for benchmarking classification models.

Table 4: Validation scenario using physical and virtual deployment

Characteristic	Description
Participants	① Physical deployment
	• 5 Raspberry Pi 4
	• 3 Rock64
	② Virtual deployment
	• 50 Docker containers
DFL Platform	Fedstellar [19]
Federation Architecture	DFL
Network Topology	Random
Federated Model	LeNet5
Dataset	MNIST [20]
Security Configuration	① Baseline
	② Encryption
	③ Encryption and MTD
Attack	Eclipse attack:
	• One external attacker
	• One target participant

The security of the federation is assessed under three different configurations, providing an expansive view of its security posture under varied conditions. Initially, the federation functions with ① a baseline with no security measures and no malicious attack for subsequent security comparisons. Following this, the federation incorporates ② encryption techniques, forming its primary line of defense. Finally, the system operates with ③ both encryption and MTD techniques, following

Fig. 1: Overall architecture of Fedstellar and the security module

the design of the proposed security module. To assess the resiliency of the security configuration against cybersecurity threats, the validation scenario simulates an eclipse attack, a significant threat in decentralized networks [21, 22]. The choice of this attack is motivated by the number of security measures it requires, as shown in Table 3. The successful mitigation of this multi-faceted attack in the validation scenario implies a high probability of successful defense against other potential attacks, as enumerated in Table 2. Fig. 2 shows the steps of the eclipse attack deployed: (i) involves isolating a chosen node, (ii) seizing control over its communications, and (iii) extracting valuable information. The implementation of the eclipse attack, as detailed in Algorithm 2, involved several technical considerations, particularly in network communication and manipulation. Initially, it required configuring two nodes to act as compromise participants. These nodes were set up using advanced socket programming techniques, allowing them to establish and hijack communication channels with the target node. By manipulating the routing tables and utilizing custom-built scripts, the attacking nodes were able to redirect traffic, effectively isolating the target node from the rest of the network. Following the steps outlined in Fig. 2, these nodes then took over

the communication channels of the isolated node, using packet-sniffing tools and protocol spoofing to simulate data extraction processes.

Fig. 2: Shematic representation of eclipse attack deployed in the validation

6 Results

This section assesses the security module performance focused on performance indicators such as the F_1 *score* for federated models, the percentage of CPU and RAM usage, network traffic quantified in megabytes (MB), and model convergence time. Fig. 3 and Fig. 4 show the performance

Algorithm 2 Implementation of eclipse attack in DFL

Require: N: Set of all nodes in the network, T: Target node, A: Attacker nodes
Ensure: Isolation and control over the target node T
1: Initialize the network with nodes in N
2: Select target node T from N
3: Initialize attackers in A
4: **for** each node n in N **do**
5:      **if** $n \in A$ **then**   ▷ Node Identification
6:          Begin monitoring communications of T
7:      **end if**
8: **end for**
9: **for** each communication link of T **do**   ▷ Node Isolation
10:      Attacker nodes in A intercept and block communications
11: **end for**
12: **for** each outbound communication from T **do**
13:      Redirect to attacker nodes in A ▷ Seizing Control
14: **end for**
15: **while** T is isolated **do**   ▷ Information Extraction
16:      Extract and analyze data from communications in T
17:      Attacker nodes mimic the legitimate network behavior
18: **end while**
19: **return** Success if T remains isolated and controlled

---

indicators in the physical and virtual deployment, respectively.

The diagram depicted in Fig. 3a demonstrates the average F_1 *score* for the federated models in a physical deployment, under three separate security configurations: baseline without security techniques and malicious attacks on the network, encryption, and encryption combined with MTD techniques to deal with attacks. All three configurations exhibit a consistent growth pattern in the early stages of the federation process ($\approx$10 minutes). The baseline configuration continues upward, achieving an F_1 *score* of 97%. This indicates the potential for high performance when security overheads are absent. However, when examining the configurations that include security measures, there is a slight decline in the F_1 *score*. In the encryption configuration, the F_1 *score* peaks at 94%, while in the combined encryption and MTD setting, the F_1 *score* fluctuates between 92.5%. Similarly, Fig. 4a shows the results in a virtual deployment. In this case, the growth is rapid in the first 6 minutes, with the baseline configuration reaching 98.9%. The encryption configuration achieves an F1 score of 95.5%, while the encryption with MTD configuration attains 93.8%. The variations observed throughout the federation process are likely due to the occasional computational overhead of the security mechanisms during the processing and transmission of data.

A more granular view of the performance in terms of CPU usage is provided by Fig. 3b and

Fig. 4b for physical and virtual deployments, respectively. In the physical deployment, the baseline CPU usage is 54.6% on average, reflecting the computational load of the training process. With the introduction of encryption, there is an increase in CPU usage to 60.9% due to the additional tasks of encrypting and decrypting data. These requirements further escalate when encryption is combined with MTD, leading to an average CPU usage of 63.2%, attributed to managing dynamic communication routes. In contrast, the virtual deployment exhibits a different pattern, as shown in Fig. 4b. The baseline configuration uses about 62.4% of the CPU, which increases to 66.1% with encryption, reflecting the computational overhead in a virtualized environment. Incorporating both encryption and MTD causes a CPU usage rise up to 68%. These figures highlight the increased resource demands in the virtual deployment, particularly when the number of participants increased.

For RAM usage in both physical and virtual deployments, a discernible trend is evident, as highlighted in Fig. 3c and Fig. 4c. In the physical deployment, the baseline configuration exhibits a lower average of 31.9%, reflecting the lower computational footprint when security measures are absent. However, including encryption mechanisms results in a slight increase in RAM usage due to the additional memory demands of the encryption process. Specifically, the encryption configuration averages 33.8%, and when the MTD technique is added alongside encryption, the average RAM usage augments to 35.9%. In contrast, the virtual deployment shows a different usage pattern. The baseline configuration in the virtual environment uses 27% of the RAM, which is lower than in the physical deployment. This increases to 29.5% with encryption and further to 31% when both encryption and MTD are implemented. This is attributable to the additional memory required for managing dynamic communication routes under MTD. Despite the marginal increase, it underscores the added resource requirements induced by security features.

Furthermore, network traffic, as depicted in Fig. **??** for the physical deployment and Fig. 4d for the virtual deployment, provides critical insights into the performance impacts of different security configurations. In the physical setup, the baseline configuration remains modest, averaging around

(a) Model (F_1 score) (b) CPU usage (%) (c) RAM usage (%) (d) Network usage (MB)

Fig. 3: Performance of Fedstellar in a physical deployment with eight participants using MNIST during 60 minutes

(a) Model (F_1 score) (b) CPU usage (%) (c) RAM usage (%) (d) Network usage (MB)

Fig. 4: Performance of Fedstellar in a virtual deployment with 50 participants using MNIST during 60 minutes

110.2 MB. However, the integration of security mechanisms leads to an increase in network usage. The encryption configuration generates an average of 185.2 MB of network traffic, while the encryption with MTD configuration pushes the average even higher, reaching 226 MB. In contrast, the virtual deployment, which involves a larger number of participants, exhibits higher network usage across all configurations. The baseline configuration shows an average network traffic of 429.9 MB, which rises to 448.4 MB with the implementation of encryption and further to 480.8 MB when encryption is combined with MTD techniques.

Moreover, the detailed network metrics, as outlined in Table 5, further elucidate the impacts of these security configurations on network performance. These metrics include (i) throughput, measuring data transmission efficiency; (ii) latency, indicating the communication speed; (iii) packet loss, reflecting data transmission reliability; and (iv) control overhead, representing the network cost due to security management. In physical deployments, the results show a slight decrease in throughput from 92 Mbps in the baseline to 85 Mbps with encryption and MTD, coupled with a gradual increase in latency and packet loss. Conversely, in virtual deployments, the throughput remains consistent across security settings, although lower than in physical setups, indicating a potential bottleneck in virtual environments. Interestingly, latency remains lower in virtual deployments compared to physical ones, possibly due to optimized routing in virtualized networks. However, packet loss and control overhead show a marked increase with more complex security configurations, emphasizing the additional network strain introduced by these security measures.

As illustrated by Table 6, securing DFL systems with encryption and MTD techniques introduces notable computational and network overheads, evident in physical and virtual deployments. While these security measures increase CPU, RAM, and network usage, with virtual deployments showing higher resource utilization

13

Table 5: Network metrics under different security settings in DFL

Deployment	Security Setting	Throughput (Mbps)	Latency (ms)	Packet Loss (%)	Control Overhead (%)
Physical	Baseline	92	61	0.2	3.5
	Encryption	87	63	0.5	4.7
	Encryption and MTD	85	64	1.1	5.9
Virtual	Baseline	85	52	0.6	4.3
	Encryption	81	52	0.6	7.1
	Encryption and MTD	81	53	1.3	7.8

due to a larger number of participants, they are essential for protecting against data breaches and cyberattacks. This study highlights the critical balance in DFL between ensuring high predictive accuracy and adhering to rigorous security protocols, offering a comprehensive view of the performance trade-offs inherent in implementing robust security configurations in real-world scenarios.

7 Conclusion

This work formulated a threat model for DFL communications, providing a detailed understanding of potential security vulnerabilities and sensitive information that could be exposed during interactions between participating nodes. In response to these challenges, an innovative security module was developed for DFL communications. It incorporates robust defensive mechanisms, including symmetric and asymmetric encryption methods and MTD techniques, tailored to the unique structure and requirements of DFL. This security module was deployed within a real-world DFL framework called Fedstellar to evaluate its efficacy and practicality. The validation scenario was conducted through two distinct deployments. The first involved a random topology of eight physical devices engaged in solving an ML task using the MNIST dataset and facing a custom implementation of eclipse attacks. Complementing this, a second deployment was executed in a virtual environment with 50 participants, expanding the scope and scale of the validation to a more extensive network scenario. Both deployments allowed the module to be rigorously evaluated under three security configurations: baseline without security and malicious attacks, encryption, and a composite of encryption and MTD. The assessments validated the performance of the proposed module across both physical and virtual deployments, demonstrating an average F1 score of approximately 93% with an acceptable increase in system overhead. The peak values observed in the physical deployment for CPU usage, network traffic, and RAM usage were 63% ($\pm 7\%$), 226 MB (± 15 MB), and 35.9% ($\pm 1.5\%$), respectively. In the virtual deployment, these metrics slightly increased due to the larger scale of operation, reaching 68% ($\pm 9\%$) for CPU usage, 480.8 MB (± 18 MB) for network traffic, and 31% ($\pm 1.7\%$) for RAM usage. These results demonstrate the efficiency and practicality of the security module in diverse DFL applications, accommodating various deployment scales and complexities.

Future research could consider developing and integrating new security techniques into the current security module to enhance the resilience of DFL environments further. Researchers might assess these enhancements across dynamic network topologies and more participant devices to better understand their efficacy in real-world, large-scale applications. Additionally, simulations with a wider variety of potential attacks would provide valuable insights into the robustness of these defensive methods under diverse threat scenarios. These advancements could significantly contribute to achieving secure, efficient, and scalable deployment of DFL.

Declarations

Table 6: Security settings, information protection, and performance in DFL. *PD:* Physical Deployment, *VD:* Virtual Deployment

Security Settings	Information Protected	Performance Metrics *			
		F1 Score (%)	**CPU (%)**	**RAM (%)**	**Network (MB)**
Baseline	N/A	*PD:* 97 ±0.02	*PD:* 54.4 ±8	*PD:* 32 ±2.3	*PD:* 110.1 ±12
		VD: 98.9 ±0.01	*VD:* 62.4 ±12	*VD:* 27 ±1.2	*VD:* 429.9 ±8
Encryption	• Model Parameters	*PD:* 94 ±0.9	*PD:* 60.7 ±7	*PD:* 33.9 ±2.41	*PD:* 185.1 ±21
	• Roles	*VD:* 95.5 ±0.8	*VD:* 66.1 ±9	*VD:* 29.5 ±1.8	*VD:* 448.4 ±20
	• Communication Patterns				
Encryption and MTD	• Model Parameters	*PD:* 92.5 ±1.1	*PD:* 63 ±7	*PD:* 35.9 ±1.5	*PD:* 226 ±15
	• Roles	*VD:* 93.8 ±0.7	*VD:* 68 ±9	*VD:* 31 ±1.7	*VD:* 480.8 ±18
	• Communication Patterns				
	• Topology				
	• Activity Periods				

* Average values for each participant.

Defense Procurement (armasuisse) with the DEFENDIS and CyberForce projects (CYD-C-2020003), and *(d)* the University of Zürich UZH.

• Availability of data and materials: Data sharing does not apply to this article as no datasets were generated during the current study.

References

[1] Reinsel, D., Gantz, J., Rydnin, J.: The Digitization of the World From Edge to Core (2018). https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

[2] Martínez Beltrán, E.T., Quiles Pérez, M., Sánchez Sánchez, P.M., López Bernal, S., Bovet, G., Gil Pérez, M., Martínez Pérez, G., Huertas Celdrán, A.: Decentralized Federated Learning: Fundamentals, State of the Art, Frameworks, Trends, and Challenges. IEEE Communications Surveys & Tutorials **25**(4), 2983–3013 (2023) https://doi.org/10.1109/COMST.2023.3315746

[3] Shi, Y., Liu, Y., Sun, Y., Lin, Z., Shen, L., Wang, X., Tao, D.: Towards more suitable personalization in federated learning via decentralized partial model training. arXiv preprint arXiv:2305.15157 (2023)

[4] Salama, A., Stergioulis, A., Hayajneh, A.M., Zaidi, S.A.R., McLernon, D., Robertson, I.: Decentralized federated learning over slotted aloha wireless mesh networking. IEEE Access **11**, 18326–18342 (2023) https://doi.org/10.1109/ACCESS.2023.3246924

[5] Xiao, Y., Ye, Y., Huang, S., Hao, L., Ma, Z., Xiao, M., Mumtaz, S., Dobre, O.A.: Fully decentralized federated learning-based onboard mission for uav swarm system. IEEE Communications Letters **25**(10), 3296–3300 (2021) https://doi.org/10.1109/LCOMM.2021.3095362

[6] Perales Gómez, Á.L., Martínez Beltrán, E.T., Sánchez Sánchez, P.M., Huertas Celdrán, A.: TemporalFED: Detecting Cyberattacks in Industrial Time-Series Data Using Decentralized Federated Learning. arXiv preprint arXiv:2308.03554 (2023)

[7] Mothukuri, V., Parizi, R.M., Pouriyeh, S., Huang, Y., Dehghantanha, A., Srivastava, G.: A survey on security and privacy of federated learning. Future Generation Computer Systems **115**, 619–640 (2021) https://doi.org/10.1016/j.future.2020.10.007

[8] Etxezarreta, X., Garitano, I., Iturbe, M., Zurutuza, U.: Low delay network attributes randomization to proactively mitigate reconnaissance attacks in industrial control systems. Wireless Networks (2023) https://doi.org/10.1007/s11276-022-03212-5

[9] Gholami, A., Torkzaban, N., Baras, J.S.: Trusted decentralized federated learning. In: IEEE 19th Annual Consumer Communications & Networking Conference, pp. 1–6 (2022). https://doi.org/10.1109/CCNC49033.2022.9700624

[10] Mothukuri, V., Parizi, R.M., Pouriyeh, S., Dehghantanha, A., Choo, K.-K.R.: FabricFL: Blockchain-in-the-loop federated learning for trusted decentralized systems. IEEE Systems Journal **16**(3), 3711–3722 (2022) https://doi.org/10.1109/JSYST.2021.3124513

[11] Li, Y., Wang, X., Sun, R., Xie, X., Ying, S., Ren, S.: Trustiness-based hierarchical decentralized federated learning. Knowledge-Based Systems, 110763 (2023) https://doi.org/10.1016/j.knosys.2023.110763

[12] Chen, Y., Liang, L., Gao, W.: Non trust detection of decentralized federated learning based on historical gradient. Engineering Applications of Artificial Intelligence **120**, 105888 (2023) https://doi.org/10.1016/j.engappai.2023.105888

[13] Wang, P., Sun, W., Zhang, H., Ma, W., Zhang, Y.: Distributed and secure federated learning for wireless computing power networks. IEEE Transactions on Vehicular Technology, 1–13 (2023) https://doi.org/10.1109/TVT.2023.3247859

[14] Singh, S.K., Yang, L.T., Park, J.H.: Fusionfedblock: Fusion of blockchain and federated learning to preserve privacy in industry 5.0. Information Fusion **90**, 233–240 (2023) https://doi.org/10.1016/j.inffus.2022.09.027

[15] Qu, Y., Xu, C., Gao, L., Xiang, Y., Yu, S.: FL-SEC: Privacy-preserving decentralized federated learning using SignSGD for the Internet of Artificially Intelligent Things. IEEE Internet of Things Magazine **5**(1), 85–90 (2022) https://doi.org/10.1109/IOTM.001.2100173

[16] Wang, L., Zhao, X., Lu, Z., Wang, L., Zhang, S.: Enhancing privacy preservation and trustworthiness for decentralized federated learning. Information Sciences **628**, 449–468 (2023) https://doi.org/10.1016/j.ins.2023.01.130

[17] Arapakis, I., Papadopoulos, P., Katevas, K., Perino, D.: P4l: Privacy preserving peer-to-peer learning for infrastructureless setups. arXiv preprint arXiv:2302.13438 (2023)

[18] Ridhawi, I.A., Otoum, S., Aloqaily, M.: Decentralized zero-trust framework for digital twin-based 6g. arXiv preprint arXiv:2302.03107 (2023)

[19] Martínez Beltrán, E.T., Perales Gómez, Á.L., Feng, C., Sánchez Sánchez, P.M., López Bernal, S., Bovet, G., Gil Pérez, M., Martínez Pérez, G., Huertas Celdrán, A.: Fedstellar: A Platform for Decentralized Federated Learning. Expert Systems with Applications **242**, 122861 (2024) https://doi.org/10.1016/j.eswa.2023.122861

[20] Deng, L.: The MNIST database of handwritten digit images for machine learning research. IEEE Signal Processing Magazine **29**(6), 141–142 (2012) https://doi.org/10.1109/MSP.2012.2211477

[21] Alangot, B., Reijsbergen, D., Venugopalan, S., Szalachowski, P., Yeo, K.S.: Decentralized and lightweight approach to detect eclipse attacks on proof of work blockchains. IEEE Transactions on Network and Service Management **18**(2), 1659–1672 (2021) https://doi.org/10.1109/TNSM.2021.3069502

[22] Niu, B., Chen, Y., Wang, Z., Li, F., Wang, B., Li, H.: Eclipse: Preserving differential location privacy against long-term observation attacks. IEEE Transactions on Mobile Computing **21**(1), 125–138 (2022) https://doi.org/10.1109/TMC.2020.3000730