



PQ-FLCom: Post-quantum secure communication for industrial federated learning

Aniello Castiglione ^a, Vincenzo Loia ^a, Michele Nappi ^b, Chiara Pero ^c,
Matteo Polsinelli ^{a,*}

^a Department of Management & Innovation Systems, University of Salerno, Fisciano, 84084, Italy

^b Department of Computer Science, University of Salerno, Fisciano, 84084, Italy

^c Department of Human Sciences, Link Campus University, Rome, 00165, Italy

ARTICLE INFO

Keywords:

Post-quantum digital signatures
Post-quantum cryptography
Federated learning
ML-DSA-87
Kyber KEM

ABSTRACT

Industrial innovations, driven by technologies such as the Internet of Things (IoT), cloud computing, and Artificial Intelligence (AI), have revolutionized various operational domains, including supply chain management, equipment monitoring, predictive maintenance, and quality control. Deep Learning (DL) has become instrumental in addressing complex, data-driven challenges within industrial settings, relying on large datasets collected from heterogeneous sources. However, centralizing these datasets on cloud platforms poses significant privacy and security risks. Federated Learning (FL) offers a promising solution by enabling distributed model training across multiple nodes while preventing the exchange of raw data. Despite its advantages, FL introduces new vulnerabilities, especially related to the security of communication channels between participating entities. Traditional cryptographic mechanisms, such as digital signatures and model encryption, can mitigate these risks, but the emergence of quantum computing threatens the robustness of conventional solutions. This work explores the integration of Post-Quantum Cryptography (PQC) into FL to enhance security without incurring significant performance degradation. A modular FL architecture with three security levels is proposed: (i) an unprotected baseline; (ii) authenticated communication using ML-DSA digital signatures; and (iii) full protection combining ML-DSA and ML-KEM-512-based key encapsulation. Implemented within the Flower framework, the architecture is evaluated under simulated MitM attacks. Experimental results demonstrate that PQC-enhanced schemes effectively mitigate quantum-resistant threats while maintaining acceptable computational overhead, thereby ensuring model integrity and data confidentiality.

1. Introduction

Industrial innovations have transformed various operational aspects [1], from supply chain management and equipment monitoring to predictive maintenance and quality control, enabled by technologies such as Internet of Things (IoT), cloud computing, and Artificial Intelligence (AI) [2]. These advancements extend beyond traditional manufacturing, influencing diverse sectors such as healthcare [3], automotive systems for autonomous driving and real-time diagnostics [4], and agriculture for precision farming and crop monitoring [5]. Deep Learning (DL) has become a key tool in addressing intricate, data-driven problems in various industrial sectors. Effective implementation of DL solutions often requires gathering vast amounts of data from multiple sources, such as IoT sensors, but centralizing these datasets on cloud platforms introduces notable privacy and security concerns [6].

* Corresponding author.

E-mail address: mpolsinelli@unisa.it (M. Polsinelli).

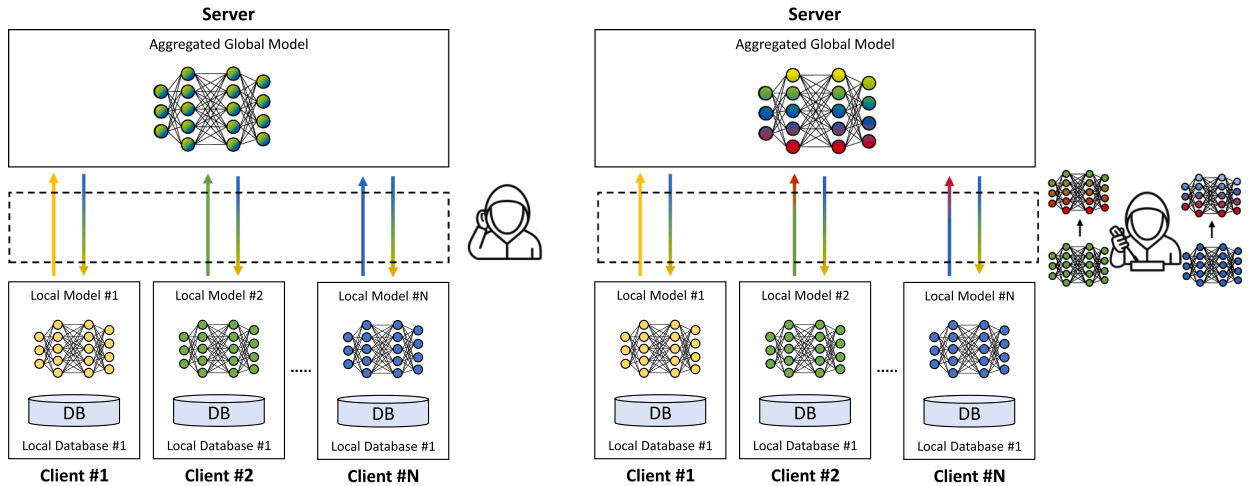


Fig. 1. Overview of the FL architecture, in which multiple clients collaboratively train a shared global model through a central aggregation server. Potential threats include eavesdropping on exchanged model updates and malicious manipulation of transmitted parameters.

Privacy and scalability issues can be mitigated through distributed training across multiple industrial nodes, which leads to Federated Learning (FL) [7,8]. FL is a framework that enhances data privacy by allowing multiple data owners to collaboratively train a shared model without exchanging sensitive raw data. Each client trains the model locally and shares only model updates with a central aggregator. FL is particularly well-suited for distributed IoT and Industrial IoT (IIoT) applications, where data are often generated across geographically distributed devices and systems. This facilitates the development of intelligent, real-time industrial systems that are both privacy-aware and performance-optimized, even in environments with limited connectivity or strict compliance requirements [9].

Recent advances in FL have enabled several IoT and Industrial IoT (IIoT) environments characterized by distributed data generation across sensors, industrial machines, edge devices, and connected vehicles. In Industry 4.0 scenarios, FL supports privacy-preserving collaborative intelligence for applications such as predictive maintenance, industrial process monitoring, fault detection, and quality inspection, enabling distributed model training directly at the factory edge. FL is also increasingly adopted in vehicular networks and intelligent transportation systems (ITSs), where connected vehicles, roadside units (RSUs), and traffic infrastructures collaboratively support applications such as traffic prediction, autonomous driving assistance, collision avoidance, and route optimization without centralized data collection. Beyond industrial and transportation contexts, FL-based IoT architectures are also being explored in smart city and edge intelligence scenarios operating under bandwidth, latency, and privacy constraints. In these distributed environments, secure and quantum-resilient communication mechanisms become critical for protecting the integrity and confidentiality of exchanged FL model updates across potentially untrusted networks [10,11]. However, the distributed nature of FL introduces new security and privacy challenges [12]. Since multiple participants are involved, it becomes harder to detect and defend against malicious behavior. As a result, these threats can be even more severe than those typically encountered in traditional, centralized ML systems [13].

An important aspect to consider in FL is the source of potential attacks. These attacks can originate from clients, the central server, or the communication channel [14]. In particular, adversaries may target the communication channel through Man-in-the-Middle (MitM) attacks by intercepting model updates exchanged between clients and the server [14] or can act like a malicious participant [15] (as shown in Fig. 1). An attacker can eavesdrop on model updates [13] to passively observe gradients or parameters. This can lead to inference attacks, where private data is reconstructed from observed updates. If the attacker modifies updates in transit, they can poison the model [16], introducing backdoors or degrading performance [17]. By spoofing a legitimate client or server, the attacker may inject malicious updates or collect sensitive information [18]. They could also replay old updates, causing the model to regress or behave undesirably.

To enhance security, techniques such as digital signatures for model authentication and model encryption are commonly employed. However, the advent of quantum computing threatens to undermine many of these cryptographic safeguards. In response to these emerging threats, significant efforts are underway to develop quantum-resistant cryptographic schemes. The National Institute of Standards and Technology (NIST), for example, is leading the development of Post-Quantum Cryptography (PQC) algorithms [19], with a particular emphasis on key encapsulation mechanisms (KEMs) and digital signatures resilient to quantum attacks. These advancements have broad applicability [20–22] and are particularly well-suited to the requirements of FL systems.

In this work, a modular FL architecture enhanced with PQC is proposed to address security vulnerabilities posed by quantum computing. The study introduces three progressively secure communication levels: an unprotected baseline (Security Level 0), authenticated communication using post-quantum digital signatures (Security Level 1), and full protection combining digital signatures and post-quantum key encapsulation and authenticated symmetric encryption (Security Level 2).

Experimental evaluations conducted under simulated MitM attacks demonstrate that the PQC-enhanced architecture significantly improves security, ensuring model integrity and data confidentiality without imposing excessive computational overhead. This study

focuses on MitM attacks, as they directly exploit cryptographic vulnerabilities that are expected to become critical in the post-quantum era [23,24]. This threat model was selected due to its direct relevance to communication integrity, which is one of the main security aspects addressed through PQC [25]. Other adversarial scenarios in industrial environments, although highly relevant to real-world deployments, require broader system-level modelling and are therefore considered outside the scope of this PQC-oriented investigation.

The main contributions can be summarized as follows:

1. A modular FL architecture integrating PQC to mitigate quantum-resistant security threats in industrial contexts.
2. Definition of three progressive security levels that balance Confidentiality, Integrity, and Availability (CIA).
3. Implementation within the Flower framework, enabling reproducible experiments and modular integration of PQC primitives.
4. Comprehensive experimental evaluation under simulated MitM attacks demonstrating the security-performance trade-off of the proposed approach.

The remainder of this paper is organized as follows. [Section 2](#) reviews the related literature on security in FL and PQC. [Section 3](#) introduces the proposed PQC-enhanced FL architecture, describing the three progressive security levels and their corresponding mechanisms. [Section 4](#) provides implementation details, including cryptographic libraries, the Flower framework, and the experimental configuration. [Section 5](#) presents and discusses the experimental results, emphasizing the robustness of the proposed model against MitM attacks. [Section 6](#) analyzes the security trade-offs in industrial applications, identifying suitable use cases for each protection level. Finally, [Section 7](#) concludes the paper and outlines future research directions.

2. Literature review

In FL architectures, clients communicate with a central server (or directly with each other) throughout the iterative learning process, exchanging model parameters multiple times. For this reason, it is essential to secure communication channels in order to mitigate threats such as MitM attacks.

The initial FL frameworks were primarily focused on improving efficiency and adaptability for industrial applications. However, security concerns were typically addressed only through the inherent properties of FL, which has proven to be insufficient. For example, in [26], the authors propose an FL framework for anomaly detection in IIoT [27] that uses a CNN-LSTM model to accurately detect anomalies from time-series data. Experiments on real-world datasets show that the framework improves detection accuracy and reduces communication overhead by approximately 50%. With the same objective, in [28], the authors introduced a Federated Transfer Learning Framework for Cross-Domain Prediction (FTL-CDP), designed to improve cross-domain prediction in smart manufacturing environments. This approach enables efficient and accurate learning across heterogeneous applications while maintaining data privacy. In addition to domain-specific optimization efforts, the complexity of building and deploying federated learning infrastructures across heterogeneous environments has also been investigated. For instance, Docker-enabled FL frameworks have been proposed to simplify distributed deployment and data stream processing across heterogeneous clients, leveraging containerization and lightweight communication protocols for scalable execution in IoT and HPC scenarios [29]. Over time, it became clear that more robust and dedicated security measures were necessary, particularly to ensure authenticity by verifying the identity of the signer and confirming that the transmitted content had not been altered. In [30], the authors propose a decentralized FL framework that uses industrial blockchain and threshold digital signatures to securely train machine learning models across space-air-ground-sea networks. The framework introduces a threshold digital signature scheme (TDSS) based on the SM-2 algorithm, enabling multiple distributed nodes to collaboratively sign and verify model updates without relying on a central server. This approach enhances data privacy, system security, and resilience against single points of failure, making it suitable for large-scale and sensitive industrial applications.

Even if authenticity is ensured, an attacker could still intercept transmitted model updates and access weights or gradients sent by the client, potentially extracting sensitive information. Therefore, encryption is necessary when confidentiality is also required, as it ensures that only authorized parties can access and read the content, even if it is intercepted. In [31], the authors introduce privacy and model security in FL (PSFL), a framework designed for Industry 4.0 that ensures both data privacy and model security. PSFL addresses vulnerabilities in traditional federated and blockchain-based FL systems, particularly against poisoning and inference attacks. It combines a multiparty cross-validation mechanism to detect and exclude malicious contributions with a lightweight signcryption algorithm that provides both encryption and authentication for secure model transmission. Similarly, in [2], the authors proposed Privacy-Enhanced Federated Learning (PEFL), a framework designed for industrial AI that enhances both privacy and efficiency. Unlike traditional FL approaches, PEFL uses a dual-layer encryption scheme and distributed differential privacy to protect local data and shared parameters, even under participant collusion. In [32], the authors propose SecFL, a secure FL framework for predicting defects in sheet metal forming under variability. It addresses privacy and security challenges in collaborative model training, especially for small and medium manufacturing companies. SecFL combines traditional FL with encryption and digital signatures to protect both data and model weights. Each client encrypts its data using a symmetric key, which is then encrypted with the server's public key to form a digital envelope. A digital signature is added to ensure authenticity and detect tampering or data poisoning. Beyond industrial manufacturing contexts, FL has been systematically analyzed in digital healthcare systems, where strict privacy regulations and the sensitivity of clinical data impose stringent security requirements [33]. In such environments, even indirect information leakage through model updates may represent a critical vulnerability, further motivating the need for robust communication-level protections.

In [34], the authors introduce Beskar, a framework designed to ensure end-to-end privacy in FL. Beskar integrates post-quantum secure aggregation with differential privacy (DP) to protect user data, model updates, and deployed models against a wide range of adversaries. The framework establishes a comprehensive threat model that encompasses attacks occurring during both training

Table 1

Schematic recap of the three security levels of the proposed PQ-FLCom architecture. The table summarizes the CIA properties guaranteed by each level, the post-quantum cryptographic primitives adopted to enforce them, the underlying post-quantum security assumptions, and the representative application scenarios. The corresponding computational and communication overhead is experimentally evaluated in [Section 5](#).

Aspect	Security Level 0	Security Level 1	Security Level 2
CIA coverage	Availability	Availability + Integrity	Full CIA
Authentication	–	ML-DSA-87	ML-DSA-87
Integrity	–	SHA3-512 + ML-DSA-87	SHA3-512 + ML-DSA-87
Confidentiality	–	–	ML-KEM-512 + AES-256-GCM

and after deployment, while optimizing cryptographic operations for efficiency on resource-constrained devices. Experimental results demonstrate that Beskar achieves strong privacy guarantees and quantum resistance with significantly reduced computational and communication overhead compared to existing FL security protocols. In contrast, the focus of the present work is on securing the communication channel within potentially hostile network environments. The primary objective is to mitigate risks posed by external adversaries seeking to intercept or alter data in transit, such as MitM attacks. To achieve this, PQC is employed for both authentication and encryption, thereby ensuring the integrity and confidentiality of the communication link.

This contribution is justified as it provides a practical, adaptable, and context-specific solution for a critical application domain. While Beskar advances the theoretical frontier of comprehensive privacy in FL, the presented approach offers a pragmatic framework for securing FL systems in the unique and demanding conditions of industrial environments. The two studies are thus complementary contributions to the overarching goal of achieving secure and privacy-preserving DL.

3. Proposed architecture

In modern cryptographic systems, the security of communications against various types of cyber threats is a critical concern. One potential threat is the MitM attack, which occurs when an attacker targets the communication channel between two legitimate parties. By exploiting vulnerabilities in protocols, encryption mechanisms, or other aspects of secure communication [35], the attacker may intercept, observe, modify, drop, or replay the messages exchanged between clients and the server without their knowledge. In the context of FL, this threat is particularly relevant because model updates are repeatedly exchanged during the training process and may contain sensitive information or directly affect the evolution of the global model.

The threat model considered in this work is therefore focused on a network-level MitM adversary acting on the communication channel between FL clients and the central aggregation server. The adversary is assumed to be able to observe transmitted model updates, alter their content in transit, or replace them with manipulated messages before aggregation. However, the attacker is not assumed to compromise the private signing keys of honest clients, the private key of the server, or the local training process executed by honest participants. Moreover, the proposed architecture assumes that the public keys of the server (used for ML-KEM-512 encapsulation) and of the clients (used for ML-DSA-87 signature verification) are distributed authentically during an initial bootstrap phase, for instance through a trusted PKI or a trust-on-first-use mechanism. The security of such bootstrap is considered outside the scope of this study. Under these assumptions, the objective of the proposed architecture is not to address all possible FL security threats, such as poisoning performed by authenticated malicious clients, but to protect the communication layer against unauthorized interception and manipulation of transmitted model updates. Attacks generated by compromised clients that legitimately sign malicious updates require complementary defenses, such as robust aggregation, anomaly detection, or Byzantine-resilient FL mechanisms, and are outside the main scope of this communication-oriented study. In the experimental evaluation ([Section 5](#)), the active tampering capability is concretely simulated through Gaussian and sign-flipping perturbations applied to model updates after signature generation, while confidentiality against passive eavesdropping is guaranteed conceptually by the post-quantum security of ML-KEM-512 combined with AES-256-GCM and is therefore not directly tested through inference attacks.

Digital signature algorithms and encryption mechanisms can mitigate these communication-level risks by providing authenticity, integrity, and confidentiality. However, the rise of quantum computing threatens the long-term security of several conventional public-key cryptographic solutions, creating the need for PQC mechanisms suitable for FL scenarios. At the same time, cryptographic protection introduces computational and communication overhead; therefore, its integration should be modular and adaptable to the requirements of the target industrial application.

As a starting point for defining such requirements, the Confidentiality, Integrity, and Availability (CIA) triad serves as a foundational model in information security, representing the core objectives for protecting information systems [36]. Confidentiality focuses on preventing unauthorized data access, integrity ensures the accuracy and authenticity of information, and availability guarantees timely and reliable access to system resources. Traditionally, industrial environments have prioritized availability first, followed by integrity, with confidentiality often receiving less attention [36]. However, with the growing adoption of Internet-connected systems in Industry 4.0, all three aspects must be considered when designing secure and reliable FL infrastructures.

Achieving stronger confidentiality and integrity guarantees inevitably introduces additional computational and communication costs, which may negatively affect availability, especially in resource-constrained or real-time industrial systems. For this reason, the proposed architecture adopts a three-level security model, illustrated in [Fig. 2](#), in which protection mechanisms are progressively

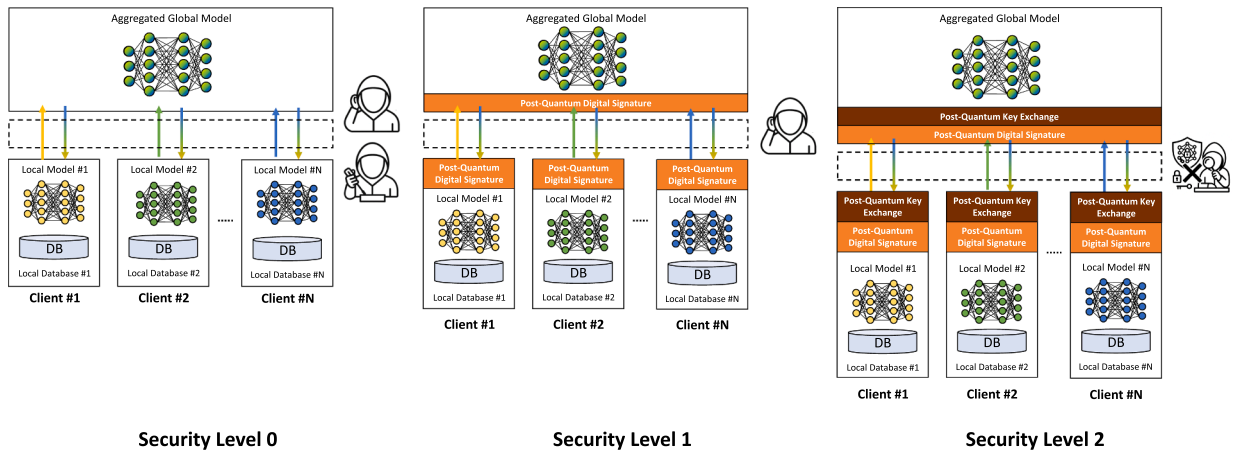


Fig. 2. Overview of the three-tier security model for FL communications. Security Level 0 provides no protection, Security Level 1 introduces post-quantum digital signatures for authentication, while Security Level 2 combines post-quantum signatures and encryption to ensure confidentiality, integrity, and authenticity.

introduced according to the required trade-off between security and overhead. Starting from an unprotected baseline, the framework moves toward increasingly stronger post-quantum communication protection.

At Security Level 0, no cryptographic protection is applied. Model updates are transmitted in plaintext and without authentication, making the communication vulnerable to both passive eavesdropping and active manipulation. This level serves as a reference baseline for quantifying the impact of MitM attacks and for evaluating the benefits introduced by the subsequent security levels.

At Security Level 1, each client signs its serialized local model update using ML-DSA-87 before transmission. The server verifies the corresponding signature before including the received update in the aggregation process. Since the signature is computed over the transmitted model update, any unauthorized modification performed by a MitM adversary after signature generation changes the verified content and causes the signature verification to fail. Therefore, under the assumption that the adversary cannot forge ML-DSA-87 signatures or access the private signing keys of honest clients, this level provides authenticity and integrity of model updates exchanged between clients and the server. Although the update remains visible to a passive observer, unauthorized tampering can be detected and the corrupted update can be excluded from aggregation.

At Security Level 2, confidentiality is added to the authenticity and integrity guarantees provided by Security Level 1. In this configuration, the client signs the model update and protects its transmission by establishing a shared symmetric key with the server through ML-KEM-512. The serialized model update is then encrypted using AES-256-GCM before being transmitted. Upon reception, the server decapsulates the shared key, decrypts the payload, and verifies the ML-DSA-87 signature before aggregation. This level prevents passive adversaries from accessing the plaintext model update and prevents active adversaries from modifying the transmitted content without detection, either through authenticated decryption failure or signature verification failure. AES-256-GCM is adopted to provide authenticated symmetric encryption, while ML-KEM-512 provides post-quantum key establishment.

Each successive level therefore introduces additional protection for the FL communication channel: Security Level 0 provides no protection and is used as a baseline; Security Level 1 protects the aggregation process against in-transit manipulation through post-quantum signature verification; and Security Level 2 additionally protects the confidentiality of model updates against eavesdropping. Replay attacks are not the primary focus of the experimental evaluation; however, the same architecture can be extended by including communication-round identifiers, client identifiers, timestamps, or nonces in the signed payload, thereby binding each update to a specific FL round and preventing the reuse of previously valid messages. [Table 1](#) provides a compact schematic recap of the three security levels, and the overall methodology of the proposed framework is summarized in [Algorithm 1](#).

The proposed architecture adopts a modular separation between the FL workflow and the cryptographic pipeline, enabling each security level to be activated independently of the underlying learning protocol. The three configurations progressively introduce post-quantum primitives selected according to their NIST standardization status and their balanced trade-off between security guarantees and computational efficiency. The implementation details, experimental setup, and corresponding evaluation are presented in the following sections.

4. Implementation details

4.1. LibOQS library

The `liboqs` library developed as part of the Open Quantum Safe (OQS) project provides a comprehensive and modular collection of post-quantum cryptographic primitives, including digital signature schemes and Key Encapsulation Mechanisms (KEMs) submitted to the NIST PQC standardization process. The library is implemented in C for performance optimization and includes a Python interface (`liboqs-python` [37]), enabling seamless integration into Python-based experimental pipelines. In this work, `liboqs` and its Python wrapper were employed to secure communication within an FL architecture based on the Flower framework (see [Section 4.2](#)). All

Algorithm 1 PQC FL.

Require: Local datasets $\{D_1, D_2, \dots, D_N\}$, global model W_{global} , number of rounds R , security level $SL \in \{0, 1, 2\}$

Ensure: Securely aggregated global model W_{global}

- 1: Initialize W_{global} on server
- 2: Initialize local models W_i on clients
- 3: **for** round = 1 to R **do**
- 4: **for** each client $i \in C_{sel}$ **in parallel do**
- 5: $W_i^{local} \leftarrow \text{LocalTraining}(W_{global}, D_i)$ ▷ FL workflow
- 6: **if** $SL = 0$ **then**
- 7: Send W_i^{local} to server
- 8: **else if** $SL = 1$ **then**
- 9: $h \leftarrow \text{SHA3-512}(\text{Serialize}(W_i^{local}))$ ▷ Cryptographic pipeline
- 10: $\sigma \leftarrow \text{ML-DSA-87.Sign}(h, sk_i)$
- 11: Send $(W_i^{local}, \sigma, pk_i)$ to server
- 12: **else if** $SL = 2$ **then**
- 13: $h \leftarrow \text{SHA3-512}(\text{Serialize}(W_i^{local}))$ ▷ Cryptographic pipeline
- 14: $\sigma \leftarrow \text{ML-DSA-87.Sign}(h, sk_i)$
- 15: $(ct, K_{shared}) \leftarrow \text{ML-KEM.Encap}(pk_{srv})$
- 16: $(E(W_i^{local}), tag, nonce) \leftarrow \text{AES-256-GCM.Encrypt}(W_i^{local}, K_{shared})$
- 17: Send $(E(W_i^{local}), tag, nonce, \sigma, ct, pk_i)$ to server
- 18: **end if**
- 19: **end for**
- 20: $ValidModels \leftarrow \emptyset$
- 21: **for** each received update from client i **do**
- 22: **if** $SL = 0$ **then**
- 23: Add W_i^{local} to $ValidModels$
- 24: **else if** $SL = 1$ **then**
- 25: $h' \leftarrow \text{SHA3-512}(\text{Serialize}(W_i^{local}))$ ▷ Cryptographic pipeline
- 26: **if** $\text{ML-DSA-87.Verify}(h', \sigma, pk_i)$ **then**
- 27: Add W_i^{local} to $ValidModels$
- 28: **end if**
- 29: **else if** $SL = 2$ **then**
- 30: $K_{shared} \leftarrow \text{ML-KEM.Decap}(ct, sk_{srv})$ ▷ Cryptographic pipeline
- 31: $W_i^{dec} \leftarrow \text{AES-256-GCM.Decrypt}(E(W_i^{local}), tag, nonce, K_{shared})$
- 32: $h' \leftarrow \text{SHA3-512}(\text{Serialize}(W_i^{dec}))$
- 33: **if** $\text{ML-DSA-87.Verify}(h', \sigma, pk_i)$ **then**
- 34: Add W_i^{dec} to $ValidModels$
- 35: **end if**
- 36: **end if**
- 37: **end for**
- 38: $W_{global} \leftarrow \text{FedAvg}(ValidModels)$ ▷ FL workflow
- 39: Broadcast W_{global} to all clients
- 40: **end for**
- 41: **return** W_{global}

cryptographic operations, including key generation, encapsulation and decapsulation, SHA3-512 hashing, signature generation and verification, and AES-GCM encryption, were executed locally via `liboqs-python`. Two distinct configurations were implemented:

- Security Level 1 - Post-quantum digital signature. Model updates generated by each client are signed using the ML-DSA-87 algorithm, a lattice-based post-quantum digital signature scheme provided by `liboqs`. Prior to signing, the model parameters are serialized and hashed using the SHA3-512 function to produce a fixed-length digest. To ensure consistency in the security level across all cryptographic primitives, SHA3-512 was adopted, thereby aligning the hashing function with the 256-bit post-quantum security level guaranteed by ML-DSA-87. This digest is then signed and transmitted along with the public key and the model update. Upon reception, the server recomputes the hash and verifies the signature, ensuring the integrity and authenticity of the received data.
- Security Level 2 - Post-quantum digital signature and hybrid encryption. This configuration extends the Level 1 scheme by incorporating confidentiality. Each client signs the model using ML-DSA-87 and establishes a shared symmetric key with the server via ML-KEM-512 (Kyber512), a post-quantum Key Encapsulation Mechanism. The signed model is subsequently encrypted using AES-256 in Galois/Counter Mode (AES-256-GCM). Although AES-256-GCM is based on a classical symmetric encryption scheme, AES-256 is generally regarded as compatible with post-quantum security requirements, since the quadratic speedup theoretically

achievable through Grover’s algorithm would still preserve an effective security level of approximately 128 bits [38]. Consequently, AES-256-GCM is widely adopted in hybrid post-quantum cryptographic architectures in which the symmetric session key is established through a post-quantum secure KEM such as ML-KEM-512.

ML-DSA-87 and ML-KEM-512 were adopted due to their favorable trade-off between post-quantum security guarantees, computational efficiency, and communication overhead. ML-DSA, standardized by the National Institute of Standards and Technology (NIST) in Federal Information Processing Standard (FIPS) 204 [39], is based on the hardness of the Module Learning With Errors (Module-LWE) and Module Short Integer Solution (Module-SIS) problems. Compared with alternative post-quantum digital signature schemes, ML-DSA provides an effective balance among signing efficiency, verification complexity, security guarantees, and implementation practicality, making it particularly suitable for FL environments characterized by iterative communication rounds and repeated cryptographic operations [40]. In particular, the ML-DSA-87 parameter set was selected to provide a high post-quantum security margin while maintaining stable computational performance throughout the experimental evaluation.

For key establishment, the framework employs ML-KEM-512, standardized in FIPS 203 [41] and derived from CRYSTALS-Kyber. ML-KEM relies on the hardness of the Module-LWE problem and represents the NIST-standardized post-quantum key encapsulation mechanism [42]. Among the available parameter sets, ML-KEM-512 was preferred because it provides NIST Security Category 1 protection while ensuring reduced computational overhead and compact public-key and ciphertext sizes.

4.2. Flower: A FL framework

In this study, Flower (Flwr) is employed as the framework for the design and coordination of the proposed FL architecture. Flower is an open-source framework written in Python, specifically designed to support scalable and modular FL systems. Developed within a research context at the University of Oxford, it bridges the gap between simulation and real-world applications across heterogeneous edge environments. One of its main strengths is interoperability with diverse machine learning frameworks and programming languages, allowing each client to operate independently with its preferred local training stack. The system architecture is based on high-level abstractions and customizable components, including client logic, server orchestration, and aggregation strategies. These elements can be extended or overridden to accommodate specific design requirements. Although native support for custom message exchange between server and clients is not included, this functionality can be added by modifying the source code. This approach is adopted in the present study to support advanced cryptographic operations.

Flower is suitable for both simulation and deployment scenarios, addressing challenges related to system heterogeneity, bandwidth limitations, and computational variability. It also facilitates the transition from centralized training pipelines to decentralized configurations, ensuring compatibility with existing machine learning workflows. The framework includes a set of pre-implemented strategies based on state-of-the-art FL algorithms, which can be used as-is or extended as needed. It also incorporates robust monitoring and fault-tolerance mechanisms, such as heartbeat tracking, straggler management, and error recovery that is essential in dynamic environments affected by network and device unreliability. Its modular and framework-agnostic design enables the customization of all stages of the FL pipeline, including client selection and model aggregation, while supporting the integration of legacy systems and heterogeneous software stacks. The client-server architecture ensures scalability through adaptive batching and dynamic resource allocation. Although Secure Aggregation (SA) is not available by default, the framework can be extended with external privacy-preserving techniques such as encryption and differential privacy. Flower is actively maintained and benefits from a vibrant open-source community that contributes documentation, examples, and technical improvements.

4.3. Datasets and FL configuration

The original Flower PyTorch setup includes a basic Convolutional Neural Network (CNN), primarily designed for educational purposes and rapid prototyping. However, preliminary evaluations demonstrated that such a lightweight architecture was insufficient to effectively learn from both the CIFAR-10 dataset and the Tiny ImageNet database in a FL setting. To address this limitation, the baseline model was replaced with a ResNet18 architecture pre-trained on ImageNet in the first case, and with a ResNet50 architecture pre-trained on ImageNet in the second case, which provide greater expressive capacity and improved generalization.

CIFAR-10: The experiments were conducted using the CIFAR-10 dataset, consisting of 60,000 color images (32×32 pixels) evenly distributed across 10 categories. In each simulation run, the dataset was partitioned into 10 client subsets using an independent and identically distributed (i.i.d.) strategy, implemented through the `IdPartitioner` module of the `flwr-datasets` library. Each local dataset was further split into training and validation sets following an 80%/20% ratio. To align with the input requirements of ResNet18, all images were resized to 224×224 pixels and normalized to zero mean and unit variance. Local training on each client was performed using stochastic gradient descent (SGD) with a learning rate of 0.001, momentum of 0.9, and a batch size of 32. Each client executed one local epoch per communication round. The FL process comprised 30 global communication rounds. In each round, 5 out of the 10 clients (i.e., 50%) were randomly selected to participate in training and aggregation. This configuration reflects a realistic FL scenario characterized by partial client availability due to heterogeneous system constraints.

Tiny ImageNet: The experiments were extended to the Tiny ImageNet dataset, which comprises 100,000 color images (64×64 pixels) evenly distributed across 200 object categories, with 500 samples per class. Unlike the CIFAR-10 case, a non-independent and identically distributed (non-i.i.d.) strategy was employed to simulate a more realistic federated setting where client datasets reflect heterogeneous data distributions. Specifically, the `DirichletPartitioner` module of the `flwr-datasets` library was used with concentration parameter $\alpha = 0.5$, resulting in skewed class distributions across clients. Each local subset was further divided into training

and validation sets using an 80%/20% split. Given the increased complexity of the dataset, the backbone architecture was replaced with ResNet50 to provide enhanced representational capacity. All input images were resized to 224×224 pixels and normalized to zero mean and unit variance. Local training on each client was carried out using stochastic gradient descent (SGD) with a learning rate of 0.001, momentum of 0.9, and a batch size of 16. Each client executed one local epoch per communication round. The federated process consisted of 200 global communication rounds. Experiments were conducted by simulating federated networks with 50, 75, and 100 clients.

These configurations were specifically designed to assess scalability and practical deployability in progressively more complex federated environments, combining concurrent client participation, non-i.i.d. data heterogeneity, increased communication overhead, and the higher model complexity introduced by ResNet50. This dual experimental setup highlights the contrast between a smaller-scale and balanced scenario (CIFAR-10) and a larger-scale heterogeneous one (Tiny ImageNet), thereby providing a comprehensive evaluation of the proposed FL framework under different data distributions and system constraints. Overall, the two experimental configurations were intentionally designed to evaluate complementary aspects of the proposed framework. In particular, CIFAR-10 serves as a controlled and lightweight FL baseline for analyzing the impact of the proposed PQC communication layers under relatively stable conditions, whereas Tiny ImageNet is employed to assess scalability, communication overhead, and robustness in larger-scale heterogeneous non-i.i.d. federated environments characterized by increased client concurrency and higher model complexity.

4.4. Simulation of MitM attacks

To assess the vulnerability of FL systems to adversarial interference, this study simulates different types of MitM attacks by extending the standard FedAvg aggregation strategy. The objective is to compromise the integrity of model updates by tampering with the parameters transmitted by selected clients before server-side aggregation. As a first baseline scenario, the core attack mechanism consists of injecting Gaussian noise into the model weights returned by compromised clients. Specifically, each parameter tensor is perturbed as follows:

$$\theta' = \theta + \mathcal{N}(0, 0.1) \quad (1)$$

where θ represents the original model parameters and $\mathcal{N}(0, 0.1)$ denotes Gaussian noise with zero mean and standard deviation equal to 0.1. Gaussian perturbation provides a controlled and reproducible baseline that allows the impact of untargeted interference on model updates to be isolated under reproducible conditions. To complement this baseline with a more directly adversarial scenario, an additional Byzantine-style sign-flipping poisoning attack was also considered. In this scenario, the attacker manipulates transmitted model updates by inverting and amplifying the parameters prior to aggregation according to:

$$\theta' = -\lambda\theta \quad (2)$$

where λ controls the attack intensity. In the experiments, $\lambda = 5$ was adopted. The value $\lambda = 5$ was selected as a moderate attack intensity, following common Byzantine-style poisoning settings in which sign inversion is combined with amplification to increase the disruptive effect while keeping the attack bounded. Unlike Gaussian perturbation, sign-flipping actively drives the global optimization process toward divergence through adversarial update inversion, thus representing a significantly more challenging poisoning scenario for FL aggregation.

Two adversarial configurations are considered:

- **Full-compromise scenario:** All clients participating in a given communication round are assumed to be malicious. Consequently, every model update is deliberately altered prior to aggregation. This represents a worst-case setting in which the server has no capability to verify the trustworthiness of participating clients.
- **Probabilistic-compromise scenario:** A more realistic setting in which only a random subset of clients (30% in each round) is assumed to be compromised. The remaining clients behave honestly and transmit unaltered updates. This configuration enables the evaluation of FL resilience in the presence of partial and intermittent adversarial behavior.

For both configurations, the tampered parameters are re-encoded and submitted to the server without any cryptographic safeguards or integrity verification mechanisms (Level 0). The simulation framework logs the identities of compromised clients and tracks key performance indicators, such as global accuracy and loss, throughout the training process. These adversarial scenarios serve as an unprotected baseline for evaluating the efficacy of subsequent security-enhanced strategies, including post-quantum digital signatures and encryption schemes.

Gaussian noise injection is adopted as a representative weak model poisoning strategy due to its generality, reproducibility, and minimal assumptions about the system. It effectively simulates a non-adaptive MitM adversary who tampers with model updates in transit, without requiring access to the original training data or internal model information. The sign-flipping attack complements this baseline scenario by simulating a stronger active adversary explicitly designed to destabilize FL convergence through adversarial directional manipulation of model updates. This allows for a controlled and interpretable evaluation of the cryptographic defenses under study. The simulated adversary operates exclusively on the communication channel and therefore does not possess valid client-side signing keys. While more sophisticated poisoning approaches (e.g., backdoor or targeted gradient attacks) could also be considered, the adopted perturbation strategies provide a clean and interpretable framework for quantifying robustness degradation under adversarial interference. The discussion of the results is deferred to [Section 5](#).

We differentiate between two primary categories of MitM attacks in the FL context: (i) eavesdropping, where a passive adversary intercepts model updates to infer sensitive information about clients' data via gradient inversion or membership inference attacks;

Table 2
Performance metrics of the FL system trained on CIFAR-10 over 30 communication rounds under different adversarial settings.

Scenario	Accuracy (avg / min / max)	Loss (avg / min / max)
No-attack scenario	0.931 / 0.7948 / 0.9408	0.219 / 0.1880 / 0.6638
Full-compromise scenario	0.273 / 0.1828 / 0.3222	2.020 / 1.8513 / 2.2450
Prob.-compromise scenario (30%)	0.606 / 0.4738 / 0.7840	1.107 / 0.6987 / 1.5017

and (ii) model poisoning, where an active adversary tampers with the transmitted weights or gradients to degrade model accuracy, introduce backdoors, or cause convergence failures. In our experiments, model poisoning is simulated through both additive Gaussian perturbation and sign-flipping manipulation of transmitted model updates, focusing on the attacks' impact on global performance metrics. Eavesdropping attacks, while not explicitly simulated here, are mitigated by the confidentiality guarantees of post-quantum cryptography introduced at Security Level 2.

5. Experimental results

All experiments were conducted on a Windows 11 Pro-64-bit system equipped with an Intel Core™ i9-13950HX CPU (13th Gen, 32 logical processors, 2.2GHz base frequency), 32GB RAM, and an NVIDIA RTX 4000 Ada Generation GPU. CUDA version 12.4 was installed. All cryptographic operations were executed on the CPU, while model training leveraged GPU acceleration when applicable. To systematically assess the vulnerability of an FL system in the absence of cryptographic safeguards, three experimental configurations were evaluated under Security Level 0. These include: (i) a baseline scenario with no adversarial interference; (ii) a full-compromise setting in which all participating clients behave maliciously; and (iii) a probabilistic-compromise scenario involving partial client corruption (30%).

5.1. Results on CIFAR-10

The evaluation on the CIFAR-10 dataset was carried out over 30 global communication rounds, with model performance assessed in terms of global accuracy and training loss. As depicted in Fig. 3, when the FL system operates under nominal conditions without any adversarial manipulation, the global model demonstrates consistent and efficient convergence. The accuracy increases rapidly during the initial communication rounds, surpassing 90% within the first ten iterations and stabilizing at approximately 94% by the end of training. The corresponding training loss exhibits a smooth monotonic decline, converging to a final value of approximately 0.19. These results validate the correctness of the learning pipeline and confirm the effectiveness of the chosen system configuration, including the use of the ResNet18 architecture, appropriate hyperparameters, and independently and identically distributed (i.i.d.) data partitioning. This scenario serves as the reference baseline for evaluating the impact of adversarial interference.

The worst-case configuration involves all clients participating in each round behaving maliciously and injecting Gaussian noise into their local model updates. Under these conditions, the global model fails to converge in any meaningful way. The training accuracy fluctuates erratically between 18% and 33% throughout the training process. The corresponding loss curve is characterized by high variance, oscillating between 1.85 and 2.25, without exhibiting any clear sign of convergence. The degradation in performance is due to the uncontrolled aggregation of poisoned updates. This disrupts the optimization trajectory and results in instability in both accuracy and loss. In the absence of integrity verification mechanisms, the central server is unable to differentiate between benign and malicious contributions, thus allowing corrupted updates to dominate the aggregation process. This scenario underscores the inherent fragility of FL systems when deployed without protective countermeasures in adversarial environments.

A more realistic adversarial scenario is illustrated in Fig. 3, where 30% of clients are randomly selected to behave maliciously at each round. Despite the majority of participants acting honestly, the global model experiences clear degradation in both performance and stability. The training accuracy starts from a relatively high value (78%) but quickly declines and fluctuates within a lower range, between 48.0% and 65.0%, with an average accuracy of approximately 60.6%. This irregular pattern reflects the disruptive impact of intermittent model poisoning. The corresponding training loss also exhibits erratic behavior, with values oscillating between 0.7 and 1.5 throughout the training process. Notably, the loss trend fails to exhibit a consistent decrease, indicating a lack of convergence and ongoing instability in the global optimization process. These findings suggest that even limited adversarial interference, such as the sporadic injection of poisoned updates by a minority of clients, can be sufficient to undermine the training dynamics of the FL system. The inability of the server to detect or mitigate corrupted updates prevents effective aggregation, resulting in persistent fluctuations in accuracy and loss. This underscores the importance of incorporating integrity verification mechanisms and robust aggregation strategies to tolerate partial compromise while preserving model convergence and performance. The numerical results summarized in Table 2 confirm the severe degradation in performance under adversarial conditions.

5.2. Results on Tiny ImageNet

The evaluation on the Tiny ImageNet dataset was conducted to examine the behaviour of the FL pipeline in a larger-scale and more challenging setting. Experiments employed ResNet50 as the backbone model, with training carried out over 200 global communication rounds. Data were distributed among clients according to a non-i.i.d. scheme based on a Dirichlet distribution. To investigate the




Fig. 3. Performance comparison of the FL system on CIFAR-10 over 30 communication rounds under baseline, partial-attack, and full-attack scenarios. The results highlight the degradation in model accuracy and convergence stability in the presence of adversarial behavior.

Table 3

Performance metrics of the FL system trained on Tiny ImageNet over 200 communication rounds under different adversarial settings (75 clients).

Scenario	Accuracy (avg / min / max)	Loss (avg / min / max)
No-attack scenario	0.75 / 0.0089 / 0.763	1.53 / 0.94 / 5.27
Full-compromise scenario	0.0077 / 0.0044 / 0.0112	5.32 / 5.27 / 5.39
Prob.-compromise scenario (30%)	0.16 / 0.0056 / 0.30	3.93 / 2.88 / 5.28

effect of system size and client heterogeneity on learning dynamics, three network scales were simulated, involving 50, 75, and 100 clients, respectively. This configuration more closely reflects realistic federated deployments.

Under nominal conditions (no adversarial interference), the global model required the full 200 rounds to achieve a satisfactory operating point. Specifically, the model attained a global accuracy of approximately 75%, with an average loss of about 1.53. It should be emphasized that, in this experimental setup, the primary objective was not to maximize absolute accuracy, but rather to evaluate the network's behaviour in terms of convergence speed, stability, and sensitivity within a complex and heterogeneous federated scenario. When adversarial interference was introduced, the system exhibited a marked deterioration in performance. In the probabilistic-compromise scenario, the global accuracy dropped to approximately 16%, while the average loss increased to around 3.93. This outcome underscores the severe disruptive impact that intermittent, partial poisoning can exert when data are non-i.i.d. and client contributions are heterogeneous. In the worst-case full-compromise scenario, the global model accuracy decreased even further, stabilizing at approximately 0.77% on average, with the best recorded value not exceeding 1.12%. The corresponding loss values were consistently high and, as expected, convergence was not achieved: the loss curves remained effectively flat and exhibited no consistent downward trend.

A comparison between Tiny ImageNet and CIFAR-10 yields two particularly noteworthy observations: (i) the Tiny ImageNet setting is substantially more challenging due to the larger number of classes and the non-i.i.d. partitioning, and (ii) the simulated system scales (50, 75, and 100 clients) make the experimental conditions more representative of realistic FL deployments. Consequently, the effects of both intermittent and full-scale model poisoning are more severe, and the training dynamics are less stable than in the CIFAR-10 (i.i.d.) case. These findings further underscore the necessity of integrity verification mechanisms and robust aggregation strategies when deploying FL in large-scale, heterogeneous environments.

The evaluation across 50, 75, and 100 concurrent clients was specifically designed as a scalability test of the proposed PQC architecture under progressively larger federated configurations. The three settings yielded consistent behaviour across system scales, confirming the stability of the cryptographic layer with respect to the number of participating clients. For conciseness, the quantitative results reported in Table 3 correspond to the 75-client configuration, which provides a representative balance between system scale and computational feasibility within the considered scalability range.

To further strengthen the adversarial evaluation, an additional Byzantine-style sign-flipping attack was considered alongside the previously adopted Gaussian perturbation baseline. Unlike additive Gaussian noise, the sign-flipping strategy actively inverts and amplifies compromised model updates before aggregation, thereby driving the global optimization process toward divergence. Under a probabilistic-compromise setting involving 30% compromised clients, the Security Level 0 configuration failed to converge. In this scenario, the global accuracy remained close to random-guessing performance throughout training, reaching only 0.31% final accuracy with a final loss of approximately 5.30 after 100 communication rounds. These results confirm that sign-flipping produces a substantially more disruptive effect than simple additive perturbations, particularly in heterogeneous non-i.i.d. federated environments. Conversely, under Security Level 1, the same attack was effectively mitigated through ML-DSA-87 signature verification. Since the adversarial manipulation was applied after client-side signature generation, all tampered updates failed the verification process and were consequently discarded before aggregation. As a result, the FL system preserved a stable convergence trajectory, achieving approximately 74.65% final accuracy with a final loss close to 1.04 after 100 communication rounds.

Table 4
Execution times (in seconds) of cryptographic operations over 30 communication rounds.

Component	Operation	Avg Time (s)	Min (s)	Max (s)
Client	Signature Generation	0.08997	0.07891	0.15270
	Encryption (AES-GCM)	0.09721	0.05948	0.13012
	Key Encapsulation (KEM)	0.00022	0.00018	0.00050
Server	Signature Generation	0.07978	0.07746	0.09151
	Decryption (AES-GCM)	0.64179	0.58971	0.67518
	Key Decapsulation (KEM)	0.00194	0.00162	0.00232
	Signature Verification	4.79133	4.29523	5.22536

5.3. Cryptographic overhead analysis

To quantify the computational overhead introduced by post-quantum cryptographic primitives within the proposed FL framework, we conducted a detailed runtime evaluation in the CIFAR-10 setting, considering 30 communication rounds and operating under the highest protection configuration, Security Level 2. This configuration integrates the ML-DSA-87 digital signature scheme with a hybrid encryption mechanism that combines ML-KEM-512 and AES-GCM.

Although the simulation involves a single client instance, all cryptographic operations were independently executed and timed in each round, yielding statistically meaningful measurements that closely approximate the per-round performance expected in realistic multi-client FL scenarios. The analysis covers core cryptographic procedures, including digital signature generation and verification, symmetric encryption and decryption, and key encapsulation and decapsulation.

As summarized in Table 4, on the client side, digital signature generation and AES-GCM encryption exhibited comparable execution times, averaging 0.08997 and 0.09721 s per round, respectively. The ML-KEM-512 key encapsulation phase proved particularly lightweight, completing in just 0.00022 s. On the server side, digital signature generation required approximately 0.07978 s per round. Decryption of the encrypted model parameters emerged as the most time-consuming task, averaging 0.64179 s, mainly due to the structured and high-dimensional nature of the model data. Most notably, verification of the client's signature was the most computationally intensive operation across the system, with an average execution time of 4.79133 s. This value represents the end-to-end pipeline cost of the verification step, which is dominated by the serialization and SHA3-512 hashing of the full ResNet18 state dictionary rather than by the ML-DSA-87 verify primitive itself. Despite this cost, the runtime remained stable throughout all rounds, with a maximum deviation of less than one second. Overall, these results demonstrate that the overhead introduced by post-quantum cryptographic safeguards is bounded, consistent, and predictable. Moreover, due to the inherently parallel structure of FL systems (where multiple clients operate simultaneously) the per-client cryptographic load is expected to scale efficiently. This validates the practical viability of integrating post-quantum protections into FL pipelines, offering enhanced security with an acceptable computational trade-off.

Finally, an important factor to consider is the latency introduced by the digital signature of DL models during the FL process. In an FL setup with N clients and a central server, each client must transmit its signed model update in every communication round, while the server returns the aggregated signed model. The total overhead in bytes per round, attributed exclusively to the signature mechanism, can be estimated as:

$$\text{Total Extra Bytes per Round} = (N + 1) \times \text{Signature Size} \quad (3)$$

and the Latency per Round (LpR) can be approximated as:

$$\text{LpR (seconds)} = \frac{\text{Total Extra Bytes per Round} \times 8}{\text{Bandwidth (bps)}} \quad (4)$$

Table 5 presents the simulated additional latency due to signature overhead. As shown, at 10 Mbps, all signature-related latencies are minimal.

Table 5

Communication overhead and induced latency of different PQC algorithms under varying network conditions (10 Mbps, 1 Mbps, 0.5 Mbps, and 0.1 Mbps) in a FL setup with 10 clients. The results account for both client-to-server and server-to-client signed model transmissions.

Algorithm	Size per Entity (KB)	Total Overhead (KB)	Latency@10 Mbps (s)	Latency@1 Mbps (s)	Latency@0.5 Mbps (s)	Latency@0.1 Mbps (s)
ML-DSA-87 Signature	4.50	49.5	0.0396	0.396	0.792	3.960
ML-KEM-512 ciphertext	0.75	8.25	0.0066	0.066	0.132	0.660
ML-KEM-512 Key	Public 0.78	8.58	0.0069	0.0686	0.1372	0.686
Total (Signature + KEM)	6.03	66.33	0.0531	0.531	1.062	5.310

Table 6

Comparison between classical ECC-based cryptographic primitives and the adopted PQC schemes. Times report average execution over 100 runs. Sizes report average values.

Primitive	Operation	Avg Time (s)	Public Key Size (bytes)	Signature/Ciphertext Size (bytes)
<i>Digital Signatures</i>				
ECDSA-P256	Signing	0.000045	65	71
ECDSA-P256	Verification	0.000075	65	71
ML-DSA-87	Signing	0.002622	2592	4627
ML-DSA-87	Verification	0.000936	2592	4627
<i>Key Exchange</i>				
ECDH-P256	Exchange	0.000074	65	N/A
ML-KEM-512	Encapsulation	0.000307	800	768
ML-KEM-512	Decapsulation	0.000165	800	768

5.3.0.1. Classical cryptographic baseline comparison. To better contextualize the computational and communication overhead introduced by PQC, we additionally benchmarked representative classical baselines, namely ECDSA-P256 for digital signatures and ECDH-P256 for key exchange. The benchmark was performed in the same experimental environment using 100 measured runs after 10 warm-up executions, evaluating standalone cryptographic operations over a SHA3-512 digest of a 1 KB message. This analysis provides a quantitative reference point for assessing the practical trade-offs associated with the adopted PQC primitives. Table 6 summarizes the obtained results.

The results reveal two main trade-offs. From a computational perspective, classical ECC primitives remain faster: ECDSA-P256 signing is approximately 58× faster than ML-DSA-87, whereas ECDH-P256 is about 4× faster than ML-KEM-512 encapsulation. Nevertheless, all PQC operations remain within the sub-millisecond to low-millisecond range, indicating their practical feasibility even in latency-sensitive FL scenarios. From a communication perspective, ML-DSA-87 entails substantially larger public keys and signatures than ECDSA-P256, while ML-KEM-512 introduces larger public keys and ciphertexts than ECDH-P256. In realistic FL deployments, however, model parameter exchanges typically dominate the communication payload, making the additional PQC overhead acceptable in exchange for stronger quantum-resistant security guarantees.

5.4. Comparison with similar approaches

The landscape of secure FL is rapidly evolving, with a growing focus on protecting against emerging threats, particularly those posed by quantum computing. Table 7 above provides a comparative analysis of our proposed approach against several state-of-the-art solutions, highlighting key differences in their main focus, application domain, and the specific security measures they employ.

A significant portion of the compared works concentrate on integrating PQC to safeguard FL frameworks from future quantum attacks. For instance, Zhang et al. [34] and Yang et al. [43] specifically address the vulnerability of secure aggregation protocols to quantum computers. Similarly, Commey et al. [44] explores the integration of PQC with blockchain technology to enhance security. Our work aligns with this research direction by proposing a modular PQC architecture for industrial systems. Another key theme is the use of privacy-preserving techniques to protect data during the FL process. Several papers [45,46] leverage Differential Privacy (DP) to add noise to the model updates, making it difficult for an adversary to infer sensitive information. In [47] Homomorphic Encryption (HE) is used to perform computations directly on encrypted data. Our approach, while focused on PQC, can be complemented by these privacy techniques to provide a more comprehensive security solution. The application domains of these secure FL frameworks are diverse, ranging from IIoT and healthcare to the financial sector. This highlights the broad applicability of FL and the need for domain-specific security solutions. For example, Appiah et al. [48] focuses on the unique challenges of the healthcare domain. Our work contributes to the security of FL in industrial settings, a critical area given the increasing connectivity of industrial control systems. In addition, Djamaa et al. [49] present FedCoRE, a lightweight FL framework designed for highly resource-constrained IoT environments. Unlike the compared security-focused works, FedCoRE optimizes FL efficiency through quantization, CoAP-based encoding, and architectural compression, but does not address quantum threats. While the authors mention possible future use of privacy-preserving techniques, our approach is complementary, as it embeds post-quantum cryptography at the architectural level, making it better suited for industrial systems requiring strong security guarantees.

In summary, while all the compared papers aim to enhance the security and privacy of FL, they do so with different focuses and techniques. Our proposed approach distinguishes itself by providing a comprehensive, PQC-enhanced FL architecture specifically designed for the security requirements of industrial systems, demonstrating that robust security can be achieved with minimal impact on performance.

6. Security trade-offs in industrial FL

FL in industrial environments must carefully balance CIA to meet diverse operational requirements [50]. Given that stronger security measures typically introduce a non-negligible computational overhead and latency for some tasks, choosing appropriate protection levels becomes critical. This section outlines security trade-offs associated with the proposed multi-tier FL architecture, highlighting suitable industrial scenarios for each security level.

Table 7

Overview of representative studies integrating Homomorphic Encryption (HE), Differential Privacy (DP), and Secure Aggregation (SA) with PQC schemes across different application domains.

Method	Application Domain	Quantum Threat Addressed	Comprehensive PQC	Privacy Technique
[34]	General FL for deep learning (Beskar)	Quantum-vulnerable SE	Kyber KEM & Dilithium signatures within 'Beskar'	Post-quantum SE + DP
[47]	General FL	Not addressed	No	HE
[46]	Intrusion detection in IIoT	Post-quantum key exchange (RLWE) for channel security	RLWE-based post-quantum key exchange (not full-stack PQC)	DP + post-quantum key exchange
[45]	Intrusion detection in IIoT	Not addressed	No	DP
[44]	Healthcare analytics	Quantum vulnerability of cryptographic schemes	Falcon, Dilithium, SPHINCS+	Blockchain for data integrity
[43]	Mobile FL	Quantum-vulnerable SE	Kyber KEM & homomorphic PRG	SE using multi-party computation
[48]	Medical data collaboration	Conventional encryption's quantum vulnerability	Integrates lattice-based digital signatures (CRYSTALS-Dilithium)	FHE-based model protection (HE)
[49]	Constrained IoT FL (AIoT)	Not addressed	No	Communication/computation optimization
Our Method	IFL	PQC threats	Modular PQC (ML-DSA, ML-KEM)	Not the primary focus

6.1. Security Level 0

Security Level 0 is suitable for applications where high availability and minimal latency are critical, and the sensitivity of the shared models is low. Typical examples include real-time process control in Industrial Control Systems (ICS), local SCADA monitoring, and low-power IoT sensor networks [51] particularly when the transmitted data consists of non-sensitive operational information. In smart transportation systems, public traffic flow prediction relies on aggregated traffic volume data collected from numerous vehicles [52]. Given that this data is non-personal and presents a low privacy risk, solutions with minimal overhead are favored to support scalability. Another relevant application area is agriculture, particularly in crop yield prediction, where multiple farms collaboratively train a model using environmental sensor data such as soil moisture, temperature, and rainfall to forecast crop yields [53]. In this case, the data relates to environmental and soil conditions, not individuals or specific businesses, resulting in a low privacy risk.

6.2. Security Level 1

For Security Level 1, maintaining the integrity of the model is essential, while the confidentiality of the model itself is considered less critical. This is particularly true when compared to the confidentiality of the underlying data, which remains protected due to the inherent characteristics of FL [54]. In this scenario, suitable applications include smart grid control signals, inter-PLC communications in factory automation, and remote equipment diagnostics. Another emerging example is FL applied in industrial IoT environments, where edge devices collaboratively train machine learning models. Here, ensuring the integrity and authenticity of local model updates sent to a central aggregator is vital, even if full encryption of the data is not always required.

In smart transportation, Cooperative Adaptive Cruise Control (CACC) uses sensors and Vehicle-to-Vehicle (V2V) communication to let vehicles automatically adjust speed and distance, enabling safer, more efficient, and coordinated driving [55]. In this scenario, ensuring the authenticity and integrity of model updates is critical to protect against spoofed or tampered control data from rogue vehicles. At the same time, low latency is essential for real-time driving responses, making lightweight security measures preferable over heavy encryption that could introduce delays.

6.3. Security Level 2

Security Level 2 applies to all scenarios where both confidentiality and integrity are non-negotiable, and where the overhead introduced by post-quantum digital signatures and cryptography is acceptable, regardless of its impact on system performance.

This level is relevant in situations where local model updates could indirectly reveal sensitive operational patterns or production behaviors where even minimal leakage may compromise valuable business or operational intelligence [56]. Examples include remote industrial control systems operating over public networks, industrial IoT gateways transmitting sensitive production data, cloud-based Manufacturing Execution Systems (MES), and over-the-air software or firmware updates. In the medical field, Security Level 2 is equally critical. Applications such as remote patient monitoring systems, telemedicine platforms, and medical IoT devices that collect patient health data require strong guarantees of confidentiality and integrity [57,58]. Even if raw patient data remains local, FL updates can still risk exposing sensitive information. To address these risks, full encryption and digital signatures are required to protect patient privacy, comply with regulatory frameworks such as the General Data Protection Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA), and ensure the integrity and trustworthiness of model aggregation across distributed healthcare facilities [59].

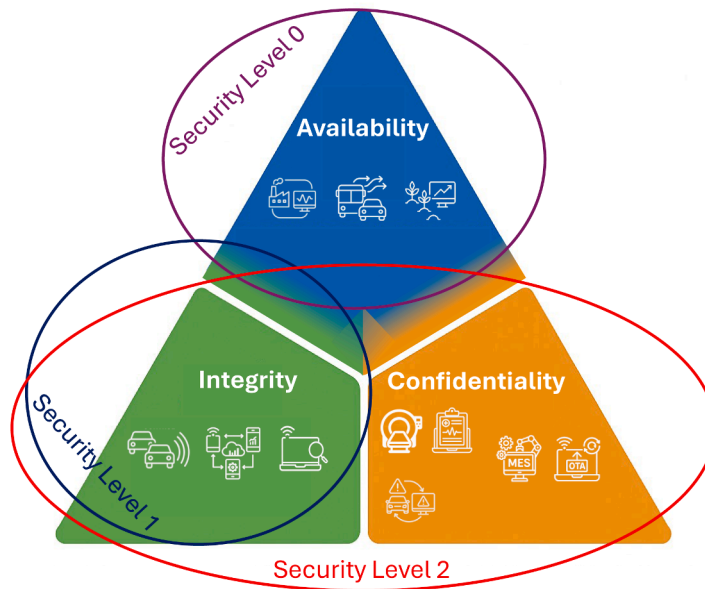


Fig. 4. Security trade-offs in industrial FL, illustrating the relative emphasis placed on CIA at Security Level 2, where both confidentiality and integrity are prioritized, potentially at the cost of reduced availability.

In smart transportation systems, real-time road hazard and accident detection refers to systems that automatically detect and report dangerous conditions or accidents as they occur. In this context, protecting both the confidentiality and integrity of sensitive sensor data from vehicles involved is critical, as compromised or inaccurate data could result in false alarms, missed incidents, or delayed emergency response, potentially putting road users at further risk.

Finally, Fig. 4 presents the CIA model centered on the three fundamental principles: availability, integrity, and confidentiality. These core elements are linked to progressively advanced security levels (Level 0 focuses solely on availability, Level 1 adds integrity, and Level 2 incorporates all three) reflecting the most comprehensive level of protection. The previously discussed case studies corresponding to each security aspect are visually represented within the pyramid.

7. Conclusions

This paper explores the integration of PQC into FL architectures to mitigate emerging security threats, particularly those posed by quantum computing. FL significantly improves data privacy in industrial settings by enabling distributed models to be trained locally, thus avoiding the risks associated with centralized data exposure. Traditional cryptographic methods, effective against classical threats, are insufficient against quantum computing, which has the potential to compromise common security mechanisms such as digital signatures. To address this, the paper proposes a multi-level security framework. The first level, Security Level 0, acts as an unprotected baseline. Security Level 1 incorporates post-quantum digital signatures (ML-DSA-87), ensuring the integrity and authenticity of the models. Security Level 2 combines these signatures with a post-quantum Key Encapsulation Mechanism (Kyber512), offering comprehensive protection by ensuring both confidentiality and integrity. Experimental results indicate that integrating PQC measures significantly enhances security, effectively mitigating MitM attacks without introducing prohibitive computational overhead. However, it was noted that the most computationally demanding operation was signature verification. A critical balance must be maintained among CIA. The proposed security model addresses this by offering varying levels of protection suited to specific operational requirements. Overall, the paper concludes that integrating PQC into FL is not only feasible but also provides robust protection against quantum threats, ensuring the integrity, authenticity, and confidentiality of critical data in industrial FL environments.

Future work also includes extending the experimental evaluation of the proposed framework under more diverse deployment conditions, beyond the simulation-based setting adopted in the present study. In particular, additional analyses are planned to further characterize the behaviour of the cryptographic pipeline under varying hardware configurations and network conditions, with the aim of assessing the practical deployability and overall robustness of the proposed PQC layer across heterogeneous FL environments. Such investigations will complement the current evaluation by providing a broader perspective on the trade-offs between security guarantees, computational efficiency, and communication overhead in realistic operational scenarios.

CRedit authorship contribution statement

Aniello Castiglione: Writing – review & editing, Visualization, Supervision, Conceptualization; **Vincenzo Loia:** Writing – review & editing, Visualization, Supervision, Conceptualization; **Michele Nappi:** Writing – review & editing, Visualization, Supervision,

Conceptualization; **Chiara Pero**: Writing – original draft, Software, Methodology, Investigation, Formal analysis, Conceptualization; **Matteo Polsinelli**: Writing – original draft, Software, Methodology, Investigation, Formal analysis, Conceptualization.

Data availability

Data will be made available on request.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This work was partially supported by the project IDA included in the Spoke 2 - Misinformation and Fakes of the Research and Innovation Program PE00000014, “SEcurity and RIghts in the Cyberspace (SERICS)”, under the National Recovery and Resilience Plan, Mission 4 “Education and Research” - Component 2 “From Research to Enterprise” - Investment 1.3, funded by the European Union - NextGenerationEU.

References

- [1] S. Savazzi, M. Nicoli, M. Bennis, S. Kianoush, L. Barbieri, Opportunities of federated learning in connected, cooperative, and automated industrial systems, *IEEE Commun. Mag.* 59 (2) (2021) 16–21. <https://doi.org/10.1109/MCOM.001.2000200>
- [2] M. Hao, H. Li, X. Luo, G. Xu, H. Yang, S. Liu, Efficient and privacy-enhanced federated learning for industrial artificial intelligence, *IEEE Trans. Ind. Inf.* 16 (10) (2020) 6532–6542. <https://doi.org/10.1109/TII.2019.2945367>
- [3] M. Karatas, L. Eriskin, M. Deveci, D. Pamucar, H. Garg, Big data for healthcare industry 4.0: applications, challenges and future perspectives, *Expert Syst. Appl.* 200 (2022) 116912. <https://doi.org/10.1016/j.eswa.2022.116912>
- [4] C. Llopis-Albert, F. Rubio, F. Valero, Impact of digital transformation on the automotive industry, *Technol. Forecast. Soc. Change* 162 (2021) 120343. <https://doi.org/10.1016/j.techfore.2020.120343>
- [5] R. Abbasi, P. Martinez, R. Ahmad, The digitization of agricultural industry—a systematic literature review on agriculture 4.0, *Smart Agric. Technol.* 2 (2022) 100042. <https://doi.org/10.1016/j.atech.2022.100042>
- [6] H. Ren, H. Li, Y. Dai, K. Yang, X. Lin, Querying in internet of things with privacy preserving: challenges, solutions and opportunities, *IEEE Netw.* 32 (6) (2018) 144–151. <https://doi.org/10.1109/MNET.2018.1700374>
- [7] B. McMahan, E. Moore, D. Ramage, S. Hampson, B.A.y. Arcas, Communication-efficient learning of deep networks from decentralized data, in: *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 54 of *Proceedings of Machine Learning Research*, PMLR, 2017, pp. 1273–1282.
- [8] B. Fakher, M. El Amine Brahmia, I. Bennis, A. Abouaissa, FedWKD: federated learning weighted aggregation with knowledge distillation for IoT forecasting, *Internet Things* 36 (2026) 101849. <https://doi.org/10.1016/j.iot.2025.101849>
- [9] D. Thakur, A. Guzzo, G. Fortino, F. Piccialli, Green federated learning: a new era of green aware AI, *ACM Comput. Surv.* 57 (8) (2025) 1–36. <https://doi.org/10.1145/371836>
- [10] D.C. Nguyen, M. Ding, P.N. Pathirana, A. Seneviratne, J. Li, D. Niyato, H.V. Poor, Federated learning for industrial internet of things in future industries, *IEEE Wirel. Commun.* 28 (6) (2021) 192–199. <https://doi.org/10.1109/MWC.001.2100102>
- [11] D.C. Nguyen, M. Ding, P.N. Pathirana, A. Seneviratne, J. Li, H.V. Poor, Federated learning for internet of things: a comprehensive survey, *IEEE Commun. Surv. Tutor.* 23 (3) (2021) 1622–1658. <https://doi.org/10.1109/COMST.2021.3075439>
- [12] S. Niu, W. Kong, L. Chen, X. Zhou, N. Wang, A homomorphic MAC-based verifiable secure aggregation for federated learning in cloud-edge AIoT, *Comput. Commun.* (2025) 108271. <https://doi.org/10.1016/j.comcom.2025.108271>
- [13] C. Zhang, S. Yang, L. Mao, H. Ning, Anomaly detection and defense techniques in federated learning: a comprehensive review, *Artif. Intell. Rev.* 57 (6) (2024) 150. <https://doi.org/10.1007/s10462-024-10796-1>
- [14] K.N. Kumar, C.K. Mohan, L.R. Cenkeramaddi, The impact of adversarial attacks on federated learning: a survey, *IEEE Trans. Pattern Anal. Mach. Intell.* 46 (5) (2024) 2672–2691. <https://doi.org/10.1109/TPAMI.2023.3322785>
- [15] J. Zhang, H. Zhu, F. Wang, J. Zhao, Q. Xu, H. Li, Security and privacy threats to federated learning: issues, methods, and challenges, *Secur. Commun. Netw.* 2022 (1) (2022) 2886795. <https://doi.org/10.1155/2022/2886795>
- [16] H. Xiong, J. Lv, D. Man, Y. Zhu, T. Liu, H. Wang, C. Xu, W. Yang, A lightweight secret-sharing-based defense against model poisoning attacks in privacy-preserving federated learning, *Comput. Commun.* (2025) 108272. <https://doi.org/10.1016/j.comcom.2025.108272>
- [17] B. Ghimire, D.B. Rawat, Recent advances on federated learning for cybersecurity and cybersecurity for federated learning for internet of things, *IEEE Internet Things J.* 9 (11) (2022) 8229–8249. <https://doi.org/10.1109/JIOT.2022.3150363>
- [18] Y. Lu, X. Huang, Y. Dai, S. Maharjan, Y. Zhang, Blockchain and federated learning for privacy-preserved data sharing in industrial IoT, *IEEE Trans. Ind. Inf.* 16 (6) (2020) 4177–4186. <https://doi.org/10.1109/TII.2019.2942190>
- [19] G. Alagic, G. Alagic, D. Apon, D. Cooper, Q. Dang, T. Dang, J. Kelsey, J. Lichtinger, Y.-K. Liu, C. Miller, et al., Status report on the third round of the NIST post-quantum cryptography standardization process (2022). <https://doi.org/10.6028/NIST.IR.8413>
- [20] F. Wu, B. Zhou, J. Song, L. Xie, Quantum-resistant blockchain and performance analysis, *J. Supercomput.* 81 (3) (2025) 498. <https://doi.org/10.1007/s11227-025-07018-y>
- [21] H.U. Khan, N. Ali, F. Ali, S. Nazir, Transforming future technology with quantum-based IoT, *J. Supercomput.* 80 (15) (2024) 22362–22396. <https://doi.org/10.1007/s11227-024-06251-1>
- [22] U.H. Khan, A. Qamar, R. Khan, M.A. Alawad, F. Alturise, M. Hayat, TrustEdge: a quantum-safe, self-healing framework for federated TinyML in critical ICU monitoring, *Internet Things* 36 (2026) 101855. <https://doi.org/10.1016/j.iot.2025.101855>
- [23] D.J. Bernstein, Post-quantum cryptography, in: *Encyclopedia of Cryptography, Security and Privacy*, Springer, 2025, pp. 1846–1847. https://doi.org/10.1007/978-1-4419-5906-5_386
- [24] M. Mosca, Cybersecurity in an era with quantum computers: will we be ready?, *IEEE Secur. Priv.* 16 (5) (2018) 38–41. <https://doi.org/10.1109/MSP.2018.3761723>
- [25] M. Sosnowski, F. Wiedner, E. Hauser, L. Steger, D. Schoinianakis, S. Gallenmüller, G. Carle, The performance of post-quantum tls 1.3, in: *Companion of the 19th International Conference on Emerging Networking Experiments and Technologies*, 2023, pp. 19–27. <https://doi.org/10.1145/3624354.3630585>
- [26] Y. Liu, N. Kumar, Z. Xiong, W.Y.B. Lim, J. Kang, D. Niyato, Communication-efficient federated learning for anomaly detection in industrial internet of things, in: *GLOBECOM 2020 - 2020 IEEE Global Communications Conference*, 2020, pp. 1–6. <https://doi.org/10.1109/GLOBECOM42002.2020.9348249>

- [27] D. Zhan, W. Zhang, L. Ye, X. Yu, H. Zhang, Z. He, Anomaly detection in industrial control systems based on cross-domain representation learning, *IEEE Trans. Dependable Secure Comput.* 22 (3) (2025) 2505–2518. <https://doi.org/10.1109/TDSC.2024.3520155>
- [28] K.I.-K. Wang, X. Zhou, W. Liang, Z. Yan, J. She, Federated transfer learning based cross-domain prediction for smart manufacturing, *IEEE Trans. Ind. Inf.* 18 (6) (2022) 4088–4096. <https://doi.org/10.1109/TII.2021.3088057>
- [29] A. Nandi, F. Xhafa, R. Kumar, A docker-based federated learning framework design and deployment for multi-modal data stream classification, *Computing* 105 (10) (2023) 2195–2229. <https://doi.org/10.1007/s00607-023-01179-5>
- [30] J. Chen, Z. Wang, G. Srivastava, T.A. Alghamdi, F. Khan, S. Kumari, H. Xiong, Industrial blockchain threshold signatures in federated learning for unified space-air-ground-sea model training, *J. Ind. Integr.* 39 (2024) 100593. <https://doi.org/10.1016/j.jii.2024.100593>
- [31] J. Li, Y. Tian, Z. Zhou, A. Xiang, S. Wang, J. Xiong, PSFL: ensuring data privacy and model security for federated learning, *IEEE Internet Things J.* 11 (15) (2024) 26234–26252. <https://doi.org/10.1109/JIOT.2024.3394168>
- [32] M.A. da Silveira Dib, P. Prates, B. Ribeiro, SecFL - secure federated learning framework for predicting defects in sheet metal forming under variability, *Expert Syst. Appl.* 235 (2024) 121139. <https://doi.org/10.1016/j.eswa.2023.121139>
- [33] F. Xhafa, *Federated Learning for Digital Healthcare Systems*, Elsevier, 2024.
- [34] Y. Zhang, R. Behnia, A.A. Yavuz, R. Ebrahimi, E. Bertino, Efficient full-stack private federated deep learning with post-quantum security, *IEEE Trans. Dependable Secure Comput.* (2025). <https://doi.org/10.1109/TDSC.2025.3568704>
- [35] S. Ansari, S.G. Rajeev, H.S. Chandrashekar, Packet sniffing: a brief introduction, *IEEE Potentials* 21 (5) (2003) 17–19. <https://doi.org/10.1109/MP.2002.1166620>
- [36] K. Tange, M. De Donno, X. Fafoutis, N. Dragoni, A systematic survey of industrial internet of things security: requirements and fog computing opportunities, *IEEE Commun. Surv. Tutor.* 22 (4) (2020) 2489–2520. <https://doi.org/10.1109/COMST.2020.3011208>
- [37] Open Quantum Safe, LibOQS-python: python 3 bindings for LibOQS, 2022. Accessed 18.04.2025, <https://github.com/open-quantum-safe/liboqs-python>.
- [38] M. Grassl, B. Langenberg, M. Roetteler, R. Steinwandt, Applying Grover's algorithm to AES: quantum resource estimates, in: *International Workshop on Post-Quantum Cryptography*, Springer, 2016, pp. 29–43. https://doi.org/10.1007/978-3-319-29360-8_3
- [39] N.I.o. Standards, Technology, Module-Lattice-Based Digital Signature Standard, Technical Report Federal Information Processing Standards Publications (FIPS PUBS) 204, U.S. Department of Commerce, Washington, D.C., 2024. <https://doi.org/10.6028/NIST.FIPS.204>
- [40] A. Castiglione, J.G. Esposito, V. Loia, M. Nappi, C. Pero, M. Polsinelli, Enhancing trust of deep learning models with post-quantum digital signatures, *J. Supercomput.* 81 (11) (2025) 1191. <https://doi.org/10.1007/s11227-025-07669-x>
- [41] N.I.o. Standards, Technology, Module-Lattice-Based Key-Encapsulation Mechanism Standard, Technical Report Federal Information Processing Standards Publication (FIPS PUB) 203, U.S. Department of Commerce, Washington, D.C., 2024. <https://doi.org/10.6028/NIST.FIPS.203>
- [42] R. Avanzi, J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J.M. Schanck, P. Schwabe, G. Seiler, D. Stehlé, et al., *CRYSTALS-Kyber algorithm specifications and supporting documentation*, NIST PQC Round 2 (4) (2019) 1–43.
- [43] S. Yang, Y. Chen, S. Tu, Z. Yang, A post-quantum secure aggregation for federated learning, in: *Proceedings of the 2022 12th International Conference on Communication and Network Security*, 2022, pp. 117–124. <https://doi.org/10.1145/3586102.358612>
- [44] D. Commy, S.G. Hounsinou, G.V. Crosby, Post-quantum secure blockchain-based federated learning framework for healthcare analytics, *IEEE Netw. Lett.* (2025). <https://doi.org/10.1109/LNET.2025.3563434>
- [45] P. Ruzafa-Alcázar, P. Fernández-Saura, E. Mármol-Campos, A. González-Vidal, J.L. Hernández-Ramos, J. Bernal-Bernabe, A.F. Skarmeta, Intrusion detection based on privacy-preserving federated learning for the industrial IoT, *IEEE Trans. Ind. Inf.* 19 (2) (2021) 1145–1154. <https://doi.org/10.1109/TII.2021.3126728>
- [46] O. Friha, M.A. Ferrag, M. Benbouzid, T. Berghout, B. Kantarci, K.-K.R. Choo, 2DF-IDS: decentralized and differentially private federated learning-based intrusion detection system for industrial IoT, *Comput. Secur.* 127 (2023) 103097. <https://doi.org/10.1016/j.cose.2023.103097>
- [47] A. Madi, O. Stan, A. Mayoue, A. Grivet-Sébert, C. Gouy-Pailler, R. Sirdey, A secure federated learning framework using homomorphic encryption and verifiable computing, in: *2021 Reconciling Data Analytics, Automation, Privacy, and Security: A Big Data Challenge (RDAAPS)*, IEEE, 2021, pp. 1–8. <https://doi.org/10.1109/RDAAPS48126.2021.9452005>
- [48] B. Appiah, I. Osei, B.K. Frimpong, D. Commy, K. Owusu-Agyang, G. Assamah, Enhanced federated learning for secure medical data collaboration, *J. Anal. Sci. Technol.* 16 (1) (2025) 13. <https://doi.org/10.1186/s40543-025-00484-2>
- [49] B. Djamaa, H. Yekhlef, M.A. Kouda, A. Bradai, FedCoRE: effective federated learning for constrained RESTful environments in the artificial intelligence of things, *J. Netw. Comput. Appl.* (2025) 104357. <https://doi.org/10.1016/j.jnca.2025.104357>
- [50] Z. Wang, J. Chen, H. Dai, J. Xu, G. Yang, H. Zhou, ACSFL: an adaptive client selection-based federated learning with personalized differential privacy for heterogeneous AIoT environments, *Comput. Commun.* (2025) 108264. <https://doi.org/10.1016/j.comcom.2025.108264>
- [51] I. Ahmad, F. Rodriguez, T. Kumar, J. Suomalainen, S.K. Jagatheesaperumal, S. Walter, M.Z. Asghar, G. Li, N. Papakonstantinou, M. Ylianttila, et al., Communications security in industry X: a survey, *IEEE Open J. Commun. Soc.* 5 (2024) 982–1025. <https://doi.org/10.1109/OJCOMS.2024.3356076>
- [52] Nidhi, J. Grover, Federated learning analysis for vehicular traffic flow prediction: evaluation of learning algorithms and aggregation approaches, *Clust. Comput.* 27 (4) (2024) 5075–5091. <https://doi.org/10.1007/s10586-023-04235-z>
- [53] T. Dey, S. Bera, A. Mukherjee, D. De, R. Buyya, FLYer: federated learning-based crop yield prediction for agriculture 5.0, *IEEE Trans. Artif. Intell.* (2025). <https://doi.org/10.1109/TAI.2025.3534149>
- [54] W. Zhang, X. Li, H. Ma, Z. Luo, X. Li, Federated learning for machinery fault diagnosis with dynamic validation and self-supervision, *Knowl. Based Syst.* 213 (2021) 106679. <https://doi.org/10.1016/j.knosys.2020.106679>
- [55] M.A. Tahiri, A. Rachid, B. Boudmane, I. Mortabit, S. Laaroussi, Toward cooperative adaptive cruise control: a mini-review, in: *2024 International Conference on Circuit, Systems and Communication (ICCS)*, IEEE, 2024, pp. 1–6. <https://doi.org/10.1109/ICCS62074.2024.10617304>
- [56] R. Accorsi, A. Lehmann, N. Lohmann, Information leak detection in business process models: theory, application, and tool support, *Inf. Syst.* 47 (2015) 244–257. <https://doi.org/10.1016/j.is.2013.12.006>
- [57] A. Younesi, E. Oustad, M. Ansari, T. Fahringer, R. Buyya, HealthCare 5.0: an industry 5.0 perspective for next-generation medical systems with synergistic integration of IoT, AI, and 6G, *Internet Things* 35 (2026) 101815. <https://doi.org/10.1016/j.iot.2025.101815>
- [58] G. Ammirata, G.J. Pezzullo, S. Contino, B.D. Martino, R. Pirrone, Federated learning framework for privacy-preserving AI in healthcare, in: *International Conference on Advanced Information Networking and Applications*, Springer, 2025, pp. 316–325. <https://doi.org/10.1109/ICSSAS66150.2025.11081017>
- [59] M. Joshi, A. Pal, M. Sankarasubbu, Federated learning for healthcare domain-pipeline, applications and challenges, *ACM Trans. Comput. Healthc.* 3 (4) (2022) 1–36. <https://doi.org/10.1145/3533708>