Forward-Secrecy-at-Scale for Post-Quantum Federated Learning in Internet of Medical Things

Collins Izuchukwu Okafor [1], Love Allen Chijioke Ahakonye [2], Dong-Seong Kim [1] *, Jae Min Lee [1]
[1] IT-Convergence Engineering, *Kumoh National Institute of Technology*, Gumi, South Korea
* NSLab Co. Ltd., Gumi, South Korea, *Kumoh National Institute of Technology*, Gumi, South Korea
[2] ICT Convergence Research Center, *Kumoh National Institute of Technology*, Gumi, South Korea
(collinsokafor, loveahakonye, dskim, ljmpaul)@kumoh.ac.kr

***Abstract*—Internet of Medical Things (IoMT) deployments increasingly rely on federated learning (FL) to train clinical models without centralizing patient data. Yet three practical risks remain: (i) "harvest-now, decrypt-later" interception that breaks long-term confidentiality once quantum computers mature, (ii) traffic-analysis leakage where packet size and timing fingerprints reveal the participating client, and (iii) poisoning by malicious participants. This study presents *Forward-Secrecy-at-Scale (FS^2)-PQC-FL*, a post-quantum secure aggregation pipeline that adds lightweight per-round forward secrecy via a symmetric ratchet, and metadata-aware traffic shaping to reduce client-identification leakage. FS^2-PQC secure aggregation achieves order-of-magnitude latency gains over CKKS fully homomorphic encryption (up to 38× speedup) while remaining compatible with heterogeneous IoMT devices, and robust aggregation (Trimmed-Mean/Krum) mitigates poisoning-induced accuracy collapse for ECG arrhythmia FL.**

***Index Terms*—Federated learning, Forward secrecy, IoMT, Post-quantum cryptography, Secure aggregation, Traffic analysis.**

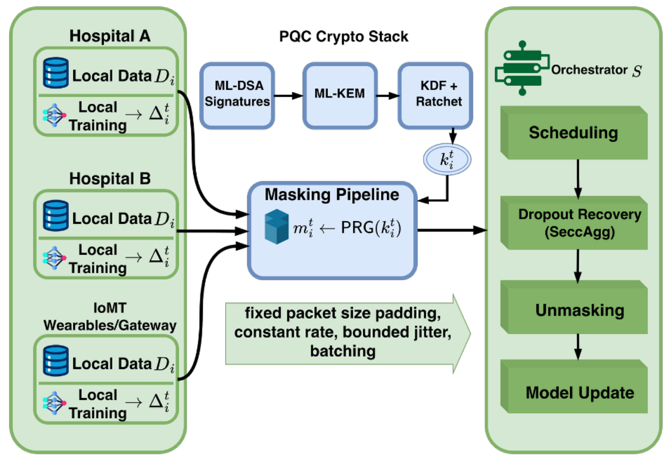

Fig. 1: PFS^2-PQC-FL system architecture

I. Introduction

Federated Learning (FL) facilitates collaborative training across hospitals and Internet of Medical Things (IoMT) devices while maintaining data locality [1]. This is vital for high-volume physiological monitoring under strict privacy constraints [2]. However, data locality alone is insufficient; communication patterns and model updates remain vulnerable to three critical threats: (i) **Quantum Adversaries** utilizing "harvest-now, decrypt-later" (HNDL) attacks [3]; (ii) **Metadata Leakage** via side-channels that reveal clinical behaviors; and (iii) **Inference and Poisoning** attacks that extract private features or degrade diagnostic accuracy [4]. Current Secure Aggregation (SecAgg) protocols often lack *forward secrecy*; a future compromise of long-term keys could expose historical clinical data. While Post-Quantum Cryptography (PQC) is gaining traction in FL, there remains a need for a system-level design that simultaneously ensures quantum resilience, per-round forward secrecy, and robustness to IoMT device dropouts with minimal overhead.To address this, we propose **FS^2-PQC-FL**, a framework for forward-secure, post-quantum FL at scale. Our contributions include:

1) **FS^2-PQC-FL Architecture:** A SecAgg design using ML-KEM and ML-DSA that ratchets fresh per-round keys to ensure forward secrecy.
2) **Robust Privacy:** An aggregation workflow that preserves confidentiality and correctness despite intermittent device participation (dropouts).
3) **Side-Channel Mitigation:** A metadata-aware scheduling strategy using constant-rate and constant-size transmissions.
4) **Performance Evaluation:** Benchmarks against Homomorphic Encryption (CKKS) demonstrating efficiency and robustness under adversarial conditions.

II. System Methodology

Cross-silo federated learning (FL) is used for IoMT deployments, where a central orchestrator S (e.g., a hospital cloud coordinator) coordinates T training rounds across N clients $C_1, \ldots, C_N$. Each client i holds private data $\mathcal{D}_i$ and computes an update Δ_i^t for the global model w^t. The threat model encompasses: (1) a global passive eavesdropper with future quantum capability (HNDL), (2) an observer targeting traffic metadata (size and timing), and (3) up to f Byzantine clients injecting poisoned updates.

Key Establishment (PQC): Clients authenticate control messages via post-quantum signatures (ML-DSA) and derive confidentiality keys through a post-quantum KEM (ML-KEM). To achieve *forward secrecy* without per-round public-

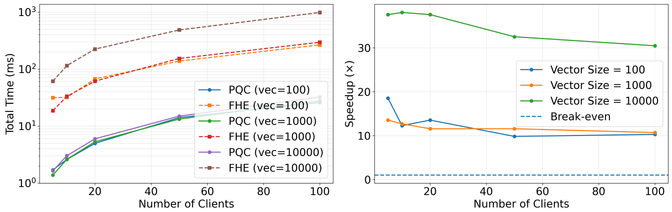

Fig. 2: PQC-SecAgg vs. CKKS-FHE total execution time (log-scale) and derived speedup as the number of clients increases.

key overhead, KEM-derived material is combined with a one-way symmetric ratchet:

$$k_i^t \leftarrow \mathsf{KDF}\big(\mathsf{Decaps}(\mathsf{ct}_i^t),\ r_i^t\big), \quad r_i^{t+1} \leftarrow H(r_i^t). \tag{1}$$

Compromise of current secrets does not reveal past keys provided $H(\cdot)$ is one-way and KDF is a PRF.

Secure Aggregation: Each client transmits a masked update:

$$\widetilde{\Delta}_i^t = \Delta_i^t + m_i^t, \quad m_i^t \leftarrow \mathsf{PRG}(k_i^t). \tag{2}$$

The server aggregates these updates and cancels masks using dropout-tolerant recovery to compute $\sum_i \Delta_i^t$ without exposing individual Δ_i^t, maintaining near-linear cost in N.

Metadata Protection: As payload confidentiality alone does not hide traffic patterns, FS^2-PQC-FL enforces fixed packet sizes through padding, constant inter-round scheduling with bounded jitter, and client-side batching, informed by the leakage trends. FS^2-PQC-FL provides post-quantum confidentiality for session material via ML-KEM and authenticated control messaging via ML-DSA. Computational cost is dominated by KEM decapsulation and symmetric KDF/PRG operations at clients, while the server performs linear-time aggregation with conditional recovery, supporting scalable operation under IoMT constraints. Figure 1 shows the PFS^2-PQC-FL system architecture.

III. Performance Evaluation

A. *PQC-SecAgg vs. CKKS-FHE*

Fig. 2 compares total execution time as clients scale and reports speedup of PQC-SecAgg over CKKS-FHE. PQC-SecAgg stays in the few–tens of milliseconds range at $N \leq 100$, while CKKS incurs two to three orders of magnitude higher latency, yielding $\sim$10×–38× speedups depending on model dimension and cohort size. The gap widens with larger vectors due to ciphertext expansion and the costs of homomorphic arithmetic.

B. *Robustness Under Poisoning*

Fig. 3 evaluates ECG arrhythmia FL under poisoning. FedAvg under attack degrades sharply, while TrimmedMean and Krum recover much of the accuracy, typically maintaining ≈80%–88% in later rounds. This highlights that confidentiality (secure aggregation) must be complemented by robustness for clinical trust. The comparison in Table I shows that the proposed PQC-secure aggregation achieves the lowest runtime (0.28–0.46 ms) compared to recent HE/PQ baselines while maintaining poisoning resilience. In contrast, prior works mainly optimize confidentiality or system efficiency, and robustness against poisoning is not consistently treated as a primary objective across all baselines.

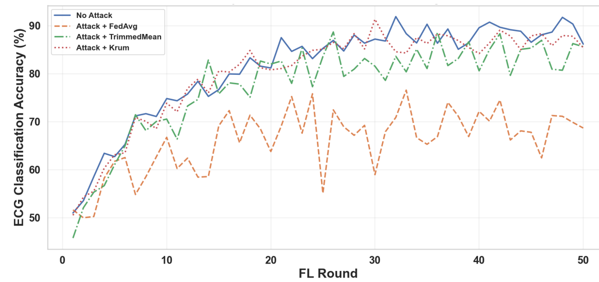

Fig. 3: ECG arrhythmia FL accuracy under poisoning: robust aggregators mitigate degradation compared to FedAvg.

TABLE I: Performance comparison with recent related work.

Study	Technique	Runtime (ms)	Robustness
[1]	CKKS HE	1.85-4.44	Not primary
[5]	Synthetic updates	1.25-1.79	Poisoning-resilient
[4]	PQ cross-silo FL	1.12-2.31	Not primary
This work	PQC-secure agg	0.28–0.46	Poisoning-resilient

IV. Conclusion

FS^2-PQC-FL advances practical post-quantum FL for IoMT by coupling PQC secure aggregation with lightweight forward secrecy and metadata-aware traffic shaping. Compared with CKKS-FHE aggregation, the PQC-SecAgg path achieves consistent order-of-magnitude speedups while remaining compatible with heterogeneous clients. Future work will formalize metadata leakage bounds and integrate adaptive padding that preserves QoS in emergency-care scenarios.

Acknowledgment

This work was partly supported by the Innovative Human Resource Development for Local Intellectualization program through the Institute of IITP grant funded by the Korea government (MSIT) (IITP-2025-RS-2020-II201612, 33%), the Priority Research Centers Program through the NRF funded by the MEST (2018R1A6A1A03024003, 33%), and by the MSIT, Korea, under the ITRC support program (IITP-2025-RS-2024-00438430, 34%).

References

[1] N. Yan, Y. Li, J. Chen, X. Wang, J. Hong, K. He, and W. Wang, "Efficient and Straggler-Resistant Homomorphic Encryption for Heterogeneous Federated Learning," in *IEEE INFOCOM 2024 - IEEE Conference on Computer Communications*, 2024, pp. 791–800.

[2] C. I. Okafor, L. A. C. Ahakonye, D.-S. Kim, and J. M. Lee, "QuantumShield V2X: A Continuous-Variable Quantum Key Distribution Framework for Eavesdropping-Resistant V2X Communications," in *2025 16th International Conference on Information and Communication Technology Convergence (ICTC)*, 2025.

[3] ——, "PureQuantum: Towards A Scalable Blockchain Channel Security in IoT Networks," *Blockchain: Research and Applications*, p. 100372, 2025.

[4] X. Qin and R. Xu, "Efficient Post-Quantum Cross-Silo Federated Learning Based on Key Homomorphic Pseudo-Random Function," *Mathematics*, vol. 13, no. 9, 2025.

[5] M. Fang, S. Nabavirazavi, Z. Liu, W. Sun, S. S. Iyengar, and H. Yang, "Do We Really Need to Design New Byzantine-robust Aggregation Rules?" 2025.