Tamperproof Quantum-Inspired Hierarchical Federated Learning for Side-Channel Security in 6G Resource-Constrained V2X Communications

Simeon Okechukwu Ajakwe*, Dong-Seong Kim
Department of IT Convergence Engineering, Kumoh National Institute of Technology, Gumi, South Korea
*ICT Convergence Research Centre, Kumoh National Institute of Technology, Gumi, South Korea
simeonajlove@gmail.com, dskim@kumoh.ac.kr

***Abstract*—This paper proposes a tamperproof Quantum-Inspired Federated Learning (QiFL) framework for side-channel attack mitigation in resource-constrained V2X networks. Phase 1 demonstrates that a tensor-network-based QiFL model achieves comparable guessing entropy to a conventional CNN while reducing parameters by approximately 95.8% and accelerating hierarchical federated training by 3.6×. Phase 2 integrates cryptographic authentication, anomaly-aware Byzantine-robust aggregation, and quantum-inspired security layers into a hierarchical FL architecture. A qubit (tensor-rank) ablation study shows that a 1-qubit QiFL configuration preserves security-level performance with minimal parameters and runtime, validating its suitability for vehicular edge–cloud deployment.**

***Index Terms*—Quantum-Inspired Federated Learning, Tamperproof Federated Learning, Side-Channel Attack Mitigation, Vehicular Networks, V2X Communications, Hierarchical Federated Learning, Byzantine-Robust Aggregation, Resource-Constrained Edge Devices**

I. INTRODUCTION

Vehicular-to-Everything (V2X) communications are central to intelligent transportation systems, enabling safety-critical applications such as cooperative perception and collision avoidance [1]. However, the cryptographic modules embedded in vehicles and roadside units remain vulnerable to side-channel attacks (SCAs), where attackers exploit power or electromagnetic leakages to recover secret keys. Recent works have shown that deep learning (DL) significantly improves profiling SCAs on datasets such as ASCAD [2], and that data augmentation and architectural tuning further enhance attack success [3]. However, these approaches typically assume centralized training, rely on large convolutional networks with millions of parameters, and ignore the realities of resource-constrained V2X devices and adversarial clients in federated settings [4].

First, centralizing sensitive traces at a single server conflicts with privacy and regulatory requirements for vehicular data [5]. Second, heavyweight DL models are difficult to deploy on on-board units with limited memory, compute, and bandwidth [6]. Third, federated learning (FL) alleviates data centralization, but existing federated learning (FL)–based SCA defenses largely overlook tampering and Byzantine behavior: malicious vehicles can inject poisoned updates, disrupt aggregation, or exfiltrate model information [7].

Motivated by these gaps, this paper proposes a tamperproof Quantum-Inspired Federated Learning (QiFL) framework that leverages quantum-classical artificial intelligence [8], tailored to V2X SCA mitigation. We employ tensor-network–based QiFL models to dramatically reduce parameter counts, embed them into a hierarchical FL architecture (cloud–edge–vehicle) that aligns with V2X infrastructure, and integrate cryptographic authentication, anomaly-aware secure aggregation, and quantum-inspired security layers to counter malicious participants. Thus, the key contributions are threefold:

1) we demonstrate that QiFL achieves approximately 95.8% parameter reduction while preserving SCA-level performance under hierarchical FL.
2) we design a tamperproof FL protocol that remains lightweight yet robust to poisoned or forged updates; and
3) we conduct a qubit (tensor-rank) ablation study showing that a 1-qubit QiFL configuration offers the best trade-off between security performance and resource consumption, validating its suitability for edge–cloud V2X deployment.

In this paper, section II describes the proposed cybercognitive security architecture, section III presents the result discussion and performance evaluation, while section IV concludes the work with future direction.

II. SYSTEM DESIGN MODEL AND ARCHITECTURE

Fig. 1 presents the system design and architecture of the proposed tamperproof quantum-inspired hierarchical federated learning (QiFL) framework for side-channel secure and resource-constrained V2X communications. We describe the hierarchical network model, the federated learning formulation, the side-channel and adversary model, the tamperproof security mechanisms, and the quantum-inspired components.

A. Hierarchical V2X Network Model

We consider a three-tier architecture: (i) a single *Global Aggregation & Security Server* (cloud), (ii) a set of edge servers / RSUs, and (iii) resource-constrained vehicles.

Let

$$\mathcal{E} = \{1, 2, \ldots, E\} \tag{1}$$

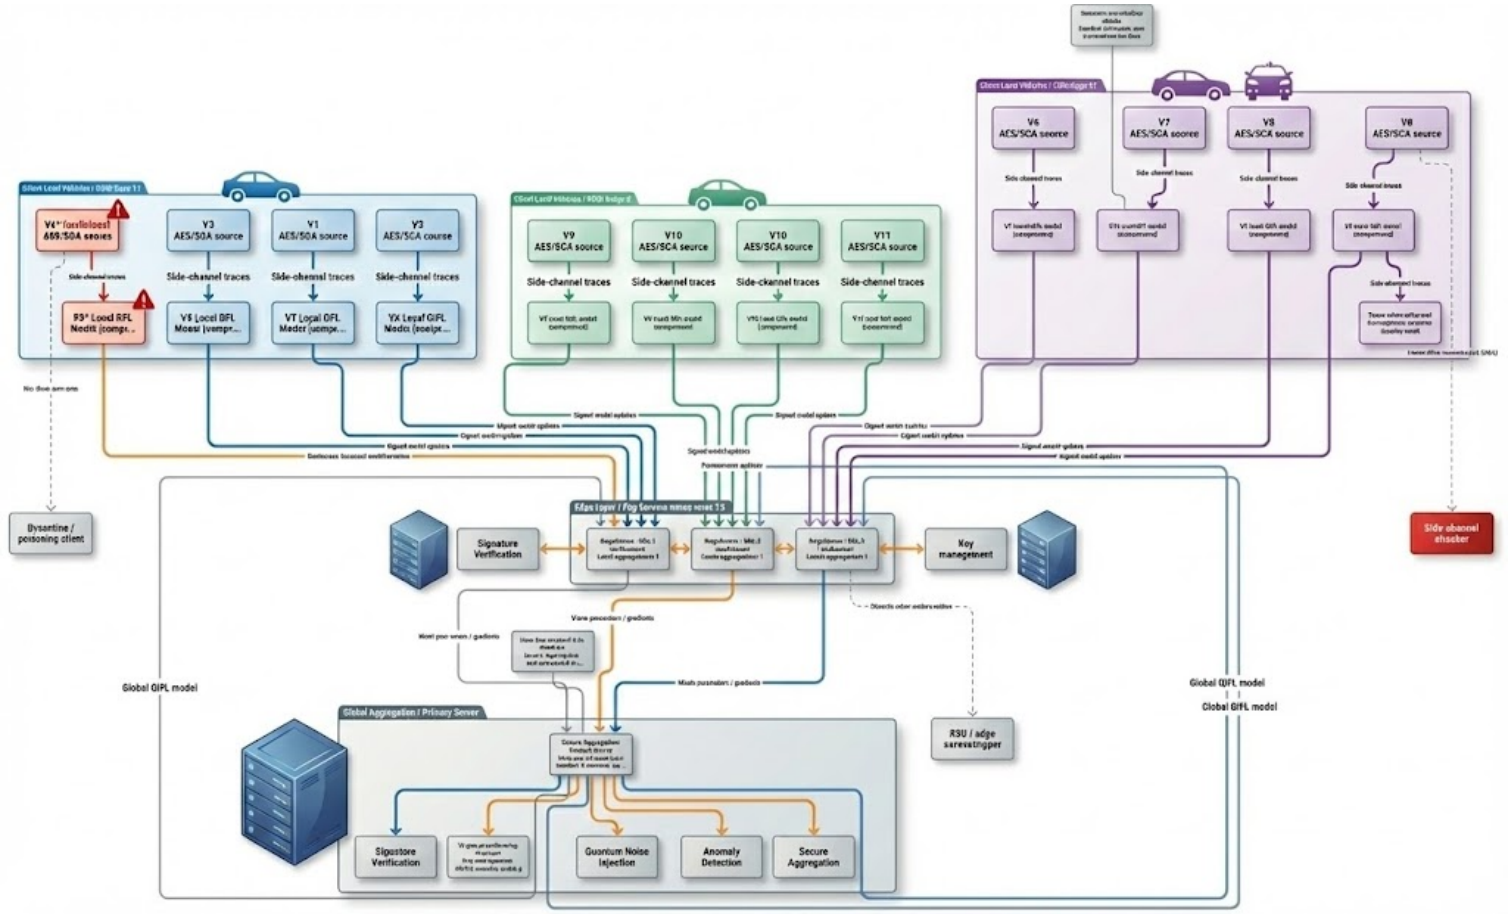

Fig. 1. System architecture of the proposed tamperproof quantum-inspired hierarchical federated learning (QiFL) framework for side-channel secure and resource-constrained V2X communications. The model integrates global tamperproof aggregation, hierarchical edge-level secure learning, vehicle-level compressed QiFL training with side-channel data, and explicit modeling of SCA and Byzantine adversaries.

denote the set of edge servers. Each edge server $e \in \mathcal{E}$ serves a set of vehicles $\mathcal{V}_e = \{1, 2, \ldots, V_e\}$. The global set of vehicles is

$$\mathcal{V} = \bigcup_{e \in \mathcal{E}} \mathcal{V}_e, \qquad V = |\mathcal{V}|. \tag{2}$$

Edge server e maintains a local copy of the global model parameters $\mathbf{w}^{(t)} \in \mathbb{R}^d$ at global round $t \in \{0, 1, \ldots, T\}$. Each vehicle (e, v) holds its own local dataset

$$\mathcal{D}_{e,v} = \{(\mathbf{x}_{e,v}^{(i)}, y_{e,v}^{(i)})\}_{i=1}^{N_{e,v}}, \tag{3}$$

where $\mathbf{x}_{e,v}^{(i)}$ corresponds to a side-channel trace and $y_{e,v}^{(i)}$ is a task-specific label (e.g., key class, leakage profile, or security level).

B. Federated Learning Objective

The global objective of hierarchical federated learning is

$$\min_{\mathbf{w} \in \mathbb{R}^d} F(\mathbf{w}) = \sum_{e \in \mathcal{E}} \alpha_e F_e(\mathbf{w}), \tag{4}$$

where $\alpha_e \geq 0$, $\sum_e \alpha_e = 1$, and $F_e(\mathbf{w})$ denotes the edge-level objective:

$$F_e(\mathbf{w}) = \sum_{v \in \mathcal{V}_e} p_{e,v} f_{e,v}(\mathbf{w}), \tag{5}$$

with $p_{e,v} = \frac{N_{e,v}}{\sum_{u \in \mathcal{V}_e} N_{e,u}}$ and

$$f_{e,v}(\mathbf{w}) = \frac{1}{N_{e,v}} \sum_{i=1}^{N_{e,v}} \ell\big(f_{\mathbf{w}}(\mathbf{x}_{e,v}^{(i)}),\, y_{e,v}^{(i)}\big), \tag{6}$$

where $\ell(\cdot, \cdot)$ is a suitable loss function and $f_{\mathbf{w}}(\cdot)$ denotes the QiFL model.

At global round t, each edge e distributes $\mathbf{w}^{(t)}$ to its vehicles. Vehicle (e, v) performs K local stochastic gradient steps:

$$\mathbf{w}_{e,v}^{(t,k+1)} = \mathbf{w}_{e,v}^{(t,k)} - \eta_{t,k} \nabla f_{e,v}\big(\mathbf{w}_{e,v}^{(t,k)};\, \mathcal{B}_{e,v}^{(t,k)}\big), \tag{7}$$

where $\mathcal{B}_{e,v}^{(t,k)}$ is a mini-batch sampled from $\mathcal{D}_{e,v}$, and the initial condition is $\mathbf{w}_{e,v}^{(t,0)} = \mathbf{w}^{(t)}$.

After local training, the vehicle computes its model update

$$\Delta \mathbf{w}_{e,v}^{(t)} = \mathbf{w}_{e,v}^{(t,K)} - \mathbf{w}^{(t)}. \tag{8}$$

C. Resource-Constrained Communication and Compression

Because vehicles are resource-constrained and V2X links are bandwidth-limited, each update $\Delta \mathbf{w}_{e,v}^{(t)}$ is compressed into a lower-rate representation

$$\widehat{\Delta \mathbf{w}}_{e,v}^{(t)} = \mathcal{C}\big(\Delta \mathbf{w}_{e,v}^{(t)}\big), \tag{9}$$

where $\mathcal{C}(\cdot)$ is a (possibly lossy) compression operator (e.g., quantization, sparsification, or low-rank projection). The effective communication cost is then

$$B_{e,v}^{(t)} = \mathrm{bits}\big(\widehat{\Delta \mathbf{w}}_{e,v}^{(t)}\big), \tag{10}$$

subject to a per-vehicle budget $B_{\max}$:

$$B_{e,v}^{(t)} \le B_{\max}, \qquad \forall e, v, t. \tag{11}$$

D. Side-Channel Leakage Model

Each vehicle (e, v) hosts a cryptographic module (e.g., AES engine) producing side-channel leakages. Let $\mathbf{k}_{e,v}$ denote the secret key and $\mathbf{m}_{e,v}(t)$ the processed message at time t. The observed leakage (e.g., power or EM trace) can be modeled as:

$$\mathbf{L}_{e,v}(t) = g\big(\mathbf{k}_{e,v}, \mathbf{m}_{e,v}(t)\big) + \mathbf{n}_{e,v}(t), \tag{12}$$

where $g(\cdot)$ is a deterministic leakage function and $\mathbf{n}_{e,v}(t)$ is noise.

The local QiFL model is trained to extract security-relevant features from side-channel traces:

$$\mathbf{x}_{e,v}^{(i)} = \mathbf{L}_{e,v}(t_i), \tag{13}$$

$$y_{e,v}^{(i)} \in \mathcal{Y} \tag{14}$$

(e.g., key class, leakage level, or attack success indicator), and the loss ℓ in (4) is evaluated over such pairs.

E. Adversary Model: SCA and Byzantine Clients

1) Side-Channel Attacker: An external side-channel attacker can eavesdrop on $\mathbf{L}_{e,v}(t)$ and attempts to recover $\mathbf{k}_{e,v}$:

$$\widehat{\mathbf{k}}_{e,v} = \mathsf{A}_{\mathrm{SCA}}\big(\{\mathbf{L}_{e,v}(t)\}_t\big), \tag{15}$$

where $\mathsf{A}_{\mathrm{SCA}}$ denotes a generic SCA algorithm (e.g., CPA, DPA, or profiling attacks). The QiFL framework aims to train models that assess and mitigate such risks while keeping the federated updates secure and privacy-preserving.

2) Byzantine / Poisoning Clients: A subset $\mathcal{B}_e \subseteq \mathcal{V}_e$ of vehicles under edge e may be compromised and transmit arbitrary malicious updates $\mathbf{b}_{e,v}^{(t)}$ instead of the honest $\widehat{\Delta \mathbf{w}}_{e,v}^{(t)}$. Formally, the edge receives

$$\widetilde{\Delta \mathbf{w}}_{e,v}^{(t)} = \begin{cases} \widehat{\Delta \mathbf{w}}_{e,v}^{(t)}, & v \notin \mathcal{B}_e, \\ \mathbf{b}_{e,v}^{(t)}, & v \in \mathcal{B}_e. \end{cases} \tag{16}$$

The objective is to design aggregation rules that remain robust under such Byzantine behavior.

F. TamperproofSecurityManager-Cryptographic Signatures

Each vehicle (e, v) shares a secret key $K_{e,v}$ with the TamperproofSecurityManager (TSM) at the global server. The vehicle signs its compressed update using an HMAC:

$$\sigma_{e,v}^{(t)} = \mathrm{HMAC}\big(K_{e,v}, \widehat{\Delta \mathbf{w}}_{e,v}^{(t)} \| t\big), \tag{17}$$

where $\|$ denotes concatenation. The transmitted packet is

$$\mathcal{P}_{e,v}^{(t)} = \big(\widehat{\Delta \mathbf{w}}_{e,v}^{(t)}, \sigma_{e,v}^{(t)}\big). \tag{18}$$

Upon reception, the edge (and/or global server) verifies

$$\mathrm{Valid}\big(\mathcal{P}_{e,v}^{(t)}\big) = \begin{cases} 1, & \text{if } \sigma_{e,v}^{(t)} = \mathrm{HMAC}\big(K_{e,v}, \widehat{\Delta \mathbf{w}}_{e,v}^{(t)} \| t\big), \\ 0, & \text{otherwise.} \end{cases} \tag{19}$$

Only updates with $\mathrm{Valid}(\mathcal{P}_{e,v}^{(t)}) = 1$ are passed to the secure aggregator. The TSM can further maintain a *security score* $s_{e,v}^{\mathrm{TSM}}(t) \in [0, 1]$ for each client, updated over time based on anomaly indicators:

$$s_{e,v}^{\mathrm{TSM}}(t+1) = \phi\Big(s_{e,v}^{\mathrm{TSM}}(t),\, \mathbf{a}_{e,v}^{(t)}\Big), \tag{20}$$

where $\mathbf{a}_{e,v}^{(t)}$ are anomaly features (e.g., deviation from aggregate, historical behavior) and $\phi(\cdot)$ is an update rule (e.g., exponential moving average).

G. SecureAggregator: Byzantine-Resilient Aggregation

At edge server e, the *SecureAggregator* combines the verified updates $\widetilde{\Delta \mathbf{w}}_{e,v}^{(t)}$ into an aggregated edge update $\Delta \mathbf{w}_e^{(t)}$. A general weighted aggregation is

$$\Delta \mathbf{w}_e^{(t)} = \sum_{v \in \mathcal{V}_e} \omega_{e,v}^{(t)}\, \widetilde{\Delta \mathbf{w}}_{e,v}^{(t)}, \tag{21}$$

where $\omega_{e,v}^{(t)} \ge 0$, $\sum_{v \in \mathcal{V}_e} \omega_{e,v}^{(t)} = 1$.

To be resilient against Byzantine updates, we can use coordinate-wise trimmed mean. For each coordinate $j \in \{1, \dots, d\}$, consider the multiset $\mathcal{S}_{e,j}^{(t)} = \{[\widetilde{\Delta \mathbf{w}}_{e,v}^{(t)}]_j : v \in \mathcal{V}_e\}$. Sorting in ascending order gives

$$s_{e,j}^{(t,1)} \le s_{e,j}^{(t,2)} \le \cdots \le s_{e,j}^{(t,|\mathcal{V}_e|)}, \tag{22}$$

and the β-trimmed mean is

$$[\Delta \mathbf{w}_e^{(t)}]_j = \frac{1}{|\mathcal{V}_e| - 2\beta} \sum_{k=\beta+1}^{|\mathcal{V}_e|-\beta} s_{e,j}^{(t,k)}, \tag{23}$$

where β is the trimming parameter.

Alternatively, the weights $\omega_{e,v}^{(t)}$ in (21) can be constructed from security scores (see next subsection), thereby implementing security-weighted aggregation:

$$\omega_{e,v}^{(t)} = \frac{s_{e,v}^{\mathrm{QS}}(t)}{\sum_{u \in \mathcal{V}_e} s_{e,u}^{\mathrm{QS}}(t)}, \tag{24}$$

where $s_{e,v}^{\mathrm{QS}}(t)$ arises from the QuantumSecurityLayer.

The global server performs another secure aggregation over edge-level updates:

$$\Delta \mathbf{w}^{(t)} = \sum_{e \in \mathcal{E}} \alpha_e \Delta \mathbf{w}_e^{(t)}, \tag{25}$$

leading to the global model update:

$$\mathbf{w}^{(t+1)} = \mathbf{w}^{(t)} + \eta_g\big(\Delta \mathbf{w}^{(t)} + \boldsymbol{\xi}^{(t)}\big), \tag{26}$$

where η_g is the global step size and $\boldsymbol{\xi}^{(t)}$ is quantum-inspired noise (Section II-I).

H. QuantumSecurityLayer and Security Scoring

Each local QiFL model embeds a *QuantumSecurityLayer* (QSL), which maps a latent representation to a scalar *security score*.

Let $f_{\mathbf{w}}(\cdot)$ be decomposed as

$$f_{\mathbf{w}}(\mathbf{x}) = g_{\boldsymbol{\theta}_2}\big(h_{\boldsymbol{\theta}_1}(\mathbf{x})\big), \tag{27}$$

where $h_{\boldsymbol{\theta}_1}(\mathbf{x}) \in \mathbb{R}^r$ is an intermediate representation. The QSL transforms $h_{\boldsymbol{\theta}_1}(\mathbf{x})$ into a score:

$$s_{e,v}^{\mathrm{QS}}(t) = \sigma\big(\psi_{\boldsymbol{\theta}_q}(h_{\boldsymbol{\theta}_1}(\mathbf{x}_{e,v}^{(t)}))\big), \tag{28}$$

where $\psi_{\boldsymbol{\theta}_q}(\cdot)$ is a parameterized function inspired by quantum interactions (e.g., entanglement-like mixing) and $\sigma(\cdot)$ is a squashing nonlinearity (e.g., sigmoid) mapping to $[0, 1]$.

A simple instantiation is

$$\psi_{\boldsymbol{\theta}_q}(\mathbf{h}) = \mathbf{h}^\top \mathbf{Q}\mathbf{h} + \mathbf{u}^\top \mathbf{h} + b_q, \tag{29}$$

where $\mathbf{Q} \in \mathbb{R}^{r\times r}$, $\mathbf{u} \in \mathbb{R}^r$, and $b_q \in \mathbb{R}$ are learnable parameters, reminiscent of quadratic forms encountered in quantum Hamiltonians.

The security score (28) can be integrated into the training objective as a regularizer, for example:

$$f_{e,v}^{\mathrm{QiFL}}(\mathbf{w}) = f_{e,v}(\mathbf{w}) + \lambda\, \mathbb{E}_{\mathbf{x}\in\mathcal{D}_{e,v}}\big[\mathcal{R}(s_{e,v}^{\mathrm{QS}}(\mathbf{x}))\big], \tag{30}$$

where $\mathcal{R}(\cdot)$ penalizes low security scores and $\lambda > 0$ is a tradeoff parameter.

I. Quantum Noise Injection

To harden the system against model inversion and reconstruction attacks, the global server injects quantum-inspired noise into aggregated updates. In (26), the noise term $\boldsymbol{\xi}^{(t)}$ can be modeled as

$$\boldsymbol{\xi}^{(t)} \sim \mathcal{N}\big(\mathbf{0},\, \sigma_q^2 \mathbf{I}_d\big), \tag{31}$$

where σ_q^2 is tuned based on the desired privacy level and robustness.

Alternatively, inspired by depolarizing quantum channels, a stochastic mixing with a reference model $\mathbf{w}_{\mathrm{ref}}$ can be used:

$$\mathbf{w}^{(t+1)} = (1-\varepsilon)\left[\mathbf{w}^{(t)} + \eta_g \Delta\mathbf{w}^{(t)}\right] + \varepsilon\, \mathbf{w}_{\mathrm{ref}}, \tag{32}$$

where $\varepsilon \in [0, 1]$ controls the strength of the mixing. This effectively perturbs the learned parameters in a way analogous to quantum depolarization, thereby reducing information leakage through shared updates.

In summary, to ensure tamperproof security against side-channel leakage, the Hierarchical QiFL Protocol allows each global round t to proceed as:

1) **Broadcast:** The global server sends $\mathbf{w}^{(t)}$ to all edge servers, which forward it to their vehicles.
2) **Local Training:** Each vehicle (e, v) performs local QiFL training on $\mathcal{D}_{e,v}$, yielding $\Delta\mathbf{w}_{e,v}^{(t)}$.
3) **Compression & Signing:** Vehicles compute $\widehat{\Delta\mathbf{w}}_{e,v}^{(t)} = \mathcal{C}(\Delta\mathbf{w}_{e,v}^{(t)})$ and $\sigma_{e,v}^{(t)} = \mathrm{HMAC}(K_{e,v}, \widehat{\Delta\mathbf{w}}_{e,v}^{(t)} \| t)$, then send $\mathcal{P}_{e,v}^{(t)}$ to the edge server.
4) **Edge-Level Secure Aggregation:** Edge e verifies signatures, filters invalid/anomalous updates, and aggregates the remaining ones using trimmed mean or security-weighted aggregation (leveraging QSL scores) to obtain $\Delta\mathbf{w}_e^{(t)}$.
5) **Global Tamperproof Aggregation:** The global server collects edge updates, applies SecureAggregator and Anomaly Detection, injects quantum-inspired noise, and updates the global model via (26).
6) **Security Scoring & Key Management:** The TamperproofSecurityManager maintains and updates client security scores and cryptographic keys, mitigating Byzantine and SCA-driven threats over time.

This mathematical formulation captures the key elements of the proposed tamperproof quantum-inspired hierarchical federated learning architecture: global QiFL coordination, edge-level secure aggregation, side-channel-aware vehicle clients, adversary models (SCA and Byzantine), cryptographic tamperproofing, and quantum-inspired security mechanisms.

The experiments use the ASCAD variable-key side-channel dataset [9], containing profiling and attack power traces from AES implementations with associated plaintext and key metadata. Traces are cropped to 1,400 samples, converted to float32, and standardized using StandardScaler (zero mean, unit variance) fitted on profiling traces and applied to attack traces. Labels are derived as the XOR of plaintext and key bytes (256-class S-box output) and, in some analyses, binned into 16 classes. All models and training pipelines are implemented in Python using PyTorch, NumPy, h5py, and scikit-learn.

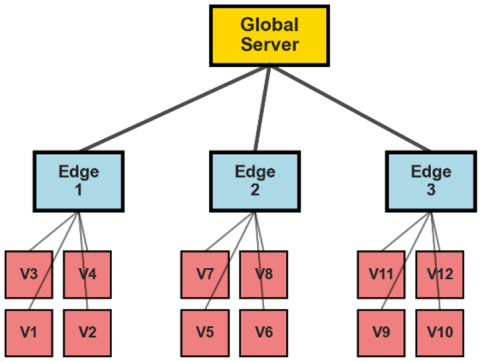

Fig. 2. QiFL configuration highlighting the hierarchies of edge servers in vehicles participating in the training rounds

III. Result and Performance Evaluation

To evaluate the performance of the proposed framework in achieving lightweight tamperproof security, we analyzed the results in two (2) phases: no degradation in SCA performance (without tamperproofing) and tamper-resistance and resilience while maintaining lightweight characteristics. The ASCAD variable-key profiling and attack traces on AES implementations. The task was a 256-class classification of S-box output (per-trace), typical in SCA works. The baseline CNN has three

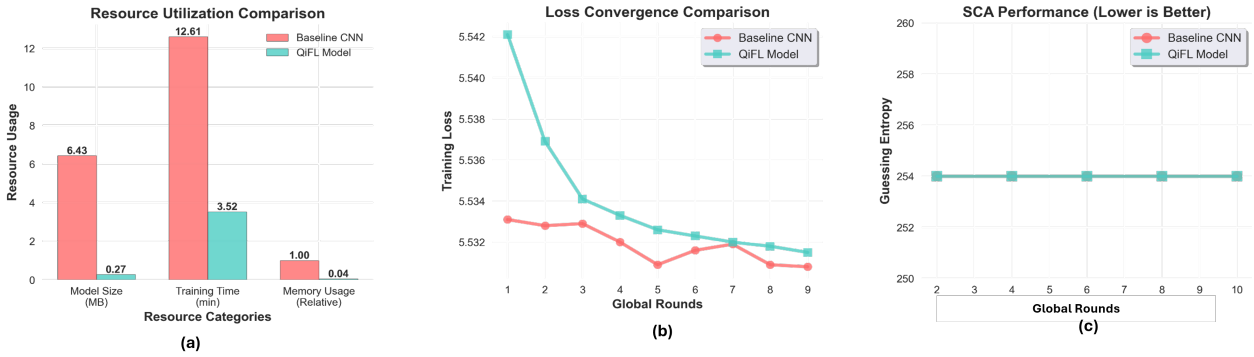

Fig. 3. Model Complexity and Efficiency Performance for Mitigating SCA in V2X Communication

Conv1D blocks and a dense classifier. On the other hand, the QiFL model uses tensor-network layers (Matrix Product State-like decomposition) for parameter compression. It has 3 edge servers (RSUs), each serving 4 vehicular clients (12 total), as seen in Fig. 2.

To train the model, the Hierarchical FL (global–edge–client) has 10 global rounds, 2 edge rounds, and 2 local epochs.

A. QiFL SCA Efficiency and Capacity

Table I captures the model complexity and efficiency of QiFL and the baseline CNN model.

TABLE I
MODEL COMPLEXITY AND EFFICIENCY OF QIFL AND CNN

Model	Parameters	Model Size (float32)	Reduction vs CNN
CNN	12 6862848	≈6.4 MB	-
QiFL	70, 624	≈ 0.27MB	95.8% fewer

From Table I, the QiFL model compresses the baseline by roughly 24× while still being trainable in the hierarchical FL setup. We compared the performance of a baseline centralized CNN with Hierarchical QiFL without tamperproofing on the ASCAD dataset. This supports the claim of resource-constrained suitability for vehicular clients (limited memory, bandwidth, and compute). Also, the communication overhead per FL round is reduced in the same proportion, as model update size scales linearly with parameter count. Also, for the SCA performance (see Fig. 3 (c)), the baseline CNN achieved a final guessing entropy (GE) of 254 and per-trace accuracy of 0.46% while QiFL achieved the same result. GE measures the average rank of the correct cryptographic key among all hypotheses, with higher values indicating stronger resistance to key recovery. In the 256-class AES S-box task, a GE of 254 implies near-maximal key uncertainty, confirming that QiFL preserves side-channel security despite model compression and tamperproof federated learning. In this ASCAD configuration, per-trace accuracy is low, even for the baseline CNN, which is consistent with the difficulty of direct 256-class S-box prediction from noisy traces. Thus, the QiFL achieves identical GE as the baseline, meaning that there is no degradation in side-channel resistance compared to the CNN. Hence, from an SCA perspective, QiFL is as effective as the baseline at modeling the leakage needed for key recovery. Relative to existing CNN-based SCA approaches, this result shows that a quantum-inspired tensor-network architecture can compress the model by ≈96% without sacrificing SCA performance under hierarchical FL. Finally, for model efficiency and suitability in V2X, the baseline CNN had a training time of 12.6 mins while it took QiFL ≈3.5 min, i.e., ≈3.6 × faster (see Fig. 3(a)). The reduction in both parameter count and computation leads to a substantial training speed-up, which is critical for online model updates in dynamic vehicular environments, and lower energy consumption on vehicles and RSUs. This empirically validates that QiFL is better aligned with resource-constrained V2X devices than conventional CNNs while maintaining SCA-level performance (same GE).

B. QiFL Tamper-Proof Security Resilience

To optimize the security resilience of the proposed framework, Phase 2 builds on Phase 1 of the QiFL model and introduces a tamperproof FL framework that combines cryptographic authentication, Byzantine-robust aggregation, anomaly detection, and quantum-inspired security layers as described in section II. The tamperproof framework overheads comprise (i) `TamperproofSecurityManager` per-client HMAC signatures over model parameters and Security scores, combining Signature verification history, Gradient anomaly detection, and Update norm consistency; (ii) `SecureAggregator` with security-weighted, trimmed-mean aggregation, tolerates a fraction of malicious/Byzantine clients and injects small quantum-inspired noise as an additional obfuscation layer; and (iii) `QuantumSecurityLayer` "Entanglement" features and a learned security score that influence the final classifier. We conducted an ablation study on the QiFL model by varying the number of qubits (n = 4, 3, 2, 1) to determine the optimized and best resource-efficient model for V2X communication, as seen in Table III.

As summarized in Table III, the tamperproof QiFL models in Phase 2 contain between 231k and 274k trainable parameters, depending on the qubit (tensor-rank) setting, and a complete tamperproof hierarchical FL run (5 global rounds,

TABLE II
KEY RESULTS COMPARING BASELINE CNN, QIFL, AND TAMPERPROOF QIFL CONFIGURATIONS.

Model / Phase	Setting	Params	Time [min]	GE	Acc. [%]	Performance/Outcome
Baseline CNN (Phase 1)	HFL, 256-class S-box	1,686,848	12.6	254	0.46	Standard CNN, no compression, no tamperproof FL
QiFL (Phase 1)	HFL, 256-class S-box	70,624	3.5	254	0.46	$\approx$ 95.8% parameter reduction, same GE as CNN
Tamperproof QiFL (Phase 2, 1 qubit)	HFL + security, 256-class S-box	231,553	$\approx$ 1.7–1.9	254	0.46	Secure aggregation, signatures, anomaly detection, quantum noise
Centralized QiFL (Stage A, 16-class)	Central, 16-class binned task	55,024	1.38	14 (max 16)	6.38	Simplified task to analyze ML behavior (not core SCA metric)

TABLE III
PHASE 2 TAMPERPROOF QIFL RESOURCE USAGE FOR DIFFERENT QUBIT (TENSOR-RANK) SETTINGS.

#Qubits	Tensor Rank	Parameters	Time [s]	Time [min]
4	16	273,665	114.21	1.90
3	8	249,601	106.68	1.78
2	4	237,569	106.15	1.77
1	2	231,553	100.66	1.68

3 edges × 4 clients, secure aggregation and anomaly detection enabled) completes in approximately 1.7–1.9 minutes on a CPU-only machine. These results confirm that, despite the added costs of key management, anomaly detection, Byzantine-robust aggregation, and quantum-inspired security layers, the overall model and protocol remain lightweight enough for edge–cloud V2X deployments. Across all metrics, the tamperproof QiFL model with 1 qubit is the most resource-efficient security model against SCA in V2X communication, as seen in Fig. 4.

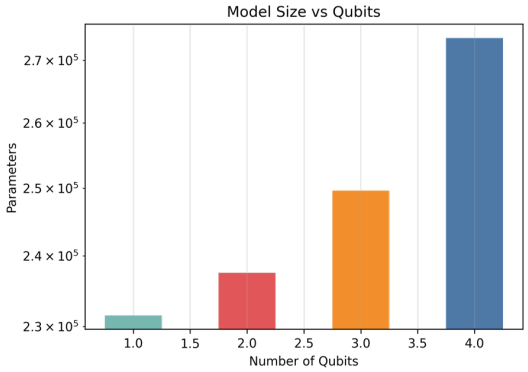

Fig. 4. Tamperproof QiFL ablation study highlighting the number of parameters vs No. of qubits

Finally, Table II summarizes the key differences in performance between the proposed model and baseline CNN in phase 1 and the tamper-proof security-resilient model with increased overhead in phase 2. In phase 1, the QiFL reduces parameters by $\approx$95.8% and training time by $\approx$3.6×, while keeping GE and accuracy identical to the baseline CNN under HFL. In Phase 2, the tamperproof QiFL (1-qubit) maintains the same GE/accuracy regime as Phase 1, even after adding secure aggregation, signatures, anomaly detection, and quantum noise, and runs in 1.7–1.9 minutes per full tamperproof HFL experiment on CPU. The result demonstrates that even under a simplified task (Centralized 16-class run), absolute accuracy remains modest, reinforcing that the main contributions are efficiency and secure FL design, not raw per-trace accuracy.

IV. CONCLUSION AND FUTURE WORK

This work proposed a tamperproof Quantum-Inspired Federated Learning (QiFL) framework for side-channel attack mitigation in resource-constrained V2X networks. By leveraging tensor-network compression and hierarchical federated learning, QiFL achieves over 95% parameter reduction while preserving CNN-level side-channel security, as measured by guessing entropy. The tamperproof extensions remain lightweight and suitable for edge–cloud deployment. While Byzantine-robust aggregation and anomaly detection are integral to the design, explicit empirical evaluation under malicious client attacks is not yet included and will be addressed in future work.

ACKNOWLEDGMENT

This research was supported by the Priority Research Centers Program through the NRF funded by the MEST (2018R1A6A1A03024003) (50%) and by the Institute of Information & Communications Technology Planning & Evaluation (IITP)-Innovative Human Resource Development for Local Intellectualization program grant funded by the Korea government(MSIT)(IITP-2025-RS-2020-II201612) (50%).

REFERENCES

[1] S. O. Ajakwe and D.-S. Kim, "Facets of security and safety problems and paradigms for smart aerial mobility and intelligent logistics," *IET Intelligent Transport Systems*, vol. 18, pp. 2827–2855, 2024.

[2] R. Benadjila, E. Prouff, R. Strullu, E. Cagli, and C. Dumas, "Deep learning for side-channel analysis and introduction to ASCAD database," *Journal of Cryptographic Engineering*, vol. 10, no. 2, pp. 163–188, 2020.

[3] X. Li, Y. Zhang, J. Wang, and H. Liu, "Deep learning-based improved side-channel attacks using data augmentation and gradient penalty," *PLOS ONE*, vol. 19, no. 9, p. e0315340, 2024.

[4] S. O. Ajakwe, K. L. Olabisi, and D.-S. Kim, "Multihop intruder node detection scheme (minds) for secured drones' fanet communication," *IET Intelligent Transport Systems*, vol. 19, no. 1, p. e70080, 2025.

[5] F. Sun, R. R. Brooks, G. Comert, and N. Tusing, "Side-channel security analysis of connected vehicle communications using hidden markov models," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 10, pp. 17 562–17 574, 2022.

[6] M. A. Al-Shareeda, S. Manickam, B. A. Mohammed, Z. G. Al-Mekhlafi, A. Qtaish, A. J. Alzahrani, G. Alshammari, A. A. Sallam, and K. Almekhlafi, "Chebyshev polynomial-based scheme for resisting side-channel attacks in 5g-enabled vehicular networks," *Applied Sciences*, vol. 12, no. 12, p. 5939, 2022.

[7] S. A. A. Hakeem and H. Kim, "Advancing intrusion detection in v2x networks: A comprehensive survey on machine learning, federated learning, and edge ai for v2x security," *IEEE Transactions on Intelligent Transportation Systems*, 2025.

[8] S. O. Ajakwe and D.-S. Kim, "EQAI: Explainable Quantum-Empowered Antispoofing Intelligence for Trustworthy Connected Autonomous Vehicles Communication," *IEEE Internet of Things Journal*, pp. 1–1, 2025.

[9] ANSSI-FR, "Ascad: A side-channel analysis dataset," https://github.com/ANSSI-FR/ASCAD, 2018, accessed: 2025-12-01.