International Journal of Information Security and Privacy
Volume 20 • Issue 1 • January-December 2026

Adaptive Lightweight Federated Learning With Aggregation-Only CKKS for Privacy-Preserving IoT Intrusion Detection

Mahdi Ajdani
http://orcid.org/0000-0002-6969-2901
Islamic Azad University of Qeshm, Iran
Received: August 13th, 2025 | **Accepted:** March 5th, 2026

ABSTRACT

Federated learning (FL) enables collaborative training without sharing raw data, but standard FL exposes client updates and burdens resource-constrained IoT devices. The authors propose AdaptiveCKKS, an FL framework combining aggregation-only CKKS encryption with index-free block sparsification and stochastic quantization. A lightweight controller adaptively selects compression ratio and quantization per device/round based on on-device calibration of bandwidth, CPU, and encryption cost, while CKKS contexts are fixed at enrollment. The server performs ciphertext-only additions, decrypting only the aggregate each round. On BoT-IoT and ToN-IoT datasets, AdaptiveCKKS improves accuracy by 3.2–3.8% over FL and fixed-HE, reduces per-round communication by ~45% and average power by ~39%, and increases resistance to membership inference and gradient inversion attacks. Results are averaged over 10 runs with 95% confidence intervals, and all artifacts are released for reproducibility.

KEYWORDS

Federated Learning, Homomorphic Encryption, CKKS, Secure Aggregation, IoT Security, Intrusion Detection, Communication Compression, Adaptive Optimization, Membership Inference, Gradient Inversion

1. INTRODUCTION

The rapid proliferation of IoT devices has pushed learning to the network edge, where privacy regulations, intermittent connectivity, and tight compute/energy budgets make centralized training impractical. Federated learning (FL) mitigates raw-data sharing, yet two practical hurdles remain in IoT settings: (i) communication overhead under heterogeneous and often asymmetric links, and (ii) privacy risks from gradient/model-update leakage in the presence of honest-but-curious servers or passive eavesdroppers (McMahan et al., 2017). Interactive secure aggregation protocols reduce leakage but introduce extra rounds (Bonawitz et al., 2017), pairwise key management, and non-trivial latency on constrained devices. Homomorphic encryption (HE) is appealing but, in full generality, is often considered too heavy for resource-limited clients.

This work targets a practical middle ground. We pair aggregation-only CKKS—where the server performs ciphertext-only additions and a designated trusted decryptor reveals only the aggregate—with index-free block sparsification and stochastic quantization. A lightweight controller adapts the compression ratio (κ\kappaκ) and quantization level (qqq) per device and per round using short

DOI: 10.4018/IJISP.402007

on-device calibrations of bandwidth, CPU, and encryption cost. Encryption parameters are fixed at enrollment via two deployable CKKS profiles, avoiding re-keying across rounds or clients. Packing utilities align blocks across clients via a shared seed, simplifying slot layout and eliminating per-index metadata.

1.1 Contributions

1. **Practical CKKS secure aggregation.** We instantiate an aggregation-only CKKS setting with fixed contexts, negligible multiplicative depth, and end-to-end encryption of client updates; the server never accesses plaintext updates and only the per-round aggregate is decrypted.
2. **Adaptive, resource-aware controller.** We tune κ\kappaκ and qqq per device/round through fast on-device probes of bandwidth/CPU/encryption cost; error-feedback preserves learning utility under sparsification and quantization without re-keying or extra interaction. (Karimireddy et al., 2019)
3. **Privacy evaluation with standard attacks.** Under an honest-but-curious server and passive eavesdroppers, we report membership-inference AUC and feature-space MSE for gradient inversion, comparing against plaintext aggregation and fixed-HE baselines where aggregates are also decrypted.
4. **Empirical gains on IoT IDS benchmarks.** On BoT-IoT and ToN-IoT, our method improves accuracy and macro-F1 by +2.5–3.5 pp over standard FL and fixed-HE, while reducing per-round communication by ≈40% and average device power by ≈39% (means with 95% CIs over 10 runs). We release CKKS profiles, packing utilities, controller code, and training scripts to support reproducibility on constrained devices.

On BoT-IoT and ToN-IoT, our method improves accuracy and macro-F1 by +3.2–3.8 pp, while reducing per-round communication by ≈45% and average device power by ≈39% (mean ±95% CI over 10 runs).

Finally, we discuss threat/trust assumptions, analyze limitations (e.g., the absence of threshold decryption), and outline how the proposed design can be integrated with interactive secure aggregation or TEEs where appropriate.

2. RELATED WORK

Federated learning for IoT intrusion detection. A growing line of work applies FL to network-traffic and device-behavior modeling in IoT, motivated by regulatory and operational constraints that prevent raw-data centralization. Prior systems typically assume uniform bandwidth or omit on-device cost modeling, which limits deployability on heterogeneous edge fleets; some rely on server-side plaintext aggregation or interactive secure aggregation, leaving either privacy gaps or additional rounds of communication. Experimental studies further analyze federated learning behavior under non-IID data silos and highlight convergence challenges (Singh et al., 2020).

Communication-efficient FL. Gradient compression has been explored via sparsification (e.g., top-k/thresholding with error feedback), quantization (e.g., stochastic and low-bit schemes), and sign-based updates (Aji & Heafield, 2017; Alistarh et al., 2017; Lin et al., 2018; Wen et al., 2017). While these techniques reduce traffic, practical deployments must coordinate indices/metadata across clients and manage non-IID drifts (Hsu et al., 2019; Li et al., 2020). Our design adopts index-free block sparsification with a shared seed for slot alignment and couples it with stochastic quantization and error feedback, which together reduce metadata and preserve utility under heterogeneity.

Secure aggregation protocols. Interactive secure aggregation (e.g., pairwise masking with dropout resilience) protects per-client updates from an honest-but-curious server but incurs extra