Article

Zero-Knowledge Federated Learning for Privacy-Preserving 5G Authentication

Ahmed Lateef Salih Al-Karawi and Rafet Akdeniz *

Posted Date: 5 February 2026

doi: 10.20944/preprints202602.0371.v1

Keywords: 5G authentication; zero-knowledge proofs; federated learning; privacy preservation; network security

Article

Zero-Knowledge Federated Learning for Privacy-Preserving 5G Authentication

Ahmed Lateef Salih Al-Karawi [1,2] and Rafet Akdeniz [3,*]

[1] Department of Computer Engineering, Faculty of Engineering, Istanbul Aydin University, Istanbul, Turkey
[2] Defne Telekomünikasyon A.Ş. Maslak Mah. Maslak Meydan Sok. Spring Giz Plaza, No :5, İç Kapı:37, Kat:9, 34485 Sarıyer /İstanbul
[3] Department of Computer Engineering, Faculty of Engineering and Natural Sciences, Atlas University, Istanbul, Turkey
* Correspondence:rafet.akdeniz@atlas.edu.tr; Tel.: +90-5423-427-501

Abstract

The fifth-generation (5G) networks are facing critical security challenges in device authentication for massive Internet of Things deployments while preserving privacy. Traditional federated learning approaches depend on the computationally expensive homomorphic encryption to protect model gradients, resulting in substantial latency, communication overhead, and the energy consumption impractical for resource-constrained 5G devices. This paper proposes zero-knowledge federated learning (ZK-FL), eliminating homomorphic encryption by enabling devices to prove model correctness without revealing gradients. Our approach integrates zero-knowledge proofs with FL updates, where each device generates where each device generates a proof $\text{Proof}_i = \text{ZK}(\text{Gradient}_i, \text{Hash}_i)$, demonstrating computational integrity.Experimental results from 10,000 authentication attempts demonstrate ZK-FL achieves 78.4 ms average authentication latency versus 342.5 ms for homomorphic encryption-based FL (77% reduction), proof sizes of 0.128 KB versus 512 KB (99.97% reduction), and energy consumption of 284.5 mJ versus 6.525 mJ (95% reduction), while maintaining 99.3% authentication success rate with formal privacy guarantees. These results demonstrate ZK-FL enables practical privacy-preserving authentication for massive-scale 5G deployment.

Keywords: 5G authentication; zero-knowledge proofs; federated learning; privacy preservation; network security.

1. Introduction

Fifth-generation (5G) wireless networks enable unprecedented connectivity for billions of devices through enhanced mobile broadband, ultra-reliable low-latency communications, and massive machine-type communications [1]. However, this massive scale introduces critical security challenges in authentication mechanisms that must ensure robust security while preserving user privacy [2].

Traditional 5G authentication relies on centralized credential verification where user equipment credentials are validated by core network functions [3]. While these mechanisms provide strong security through protocols like 5G-AKA, they suffer from privacy preservation limitations, single points of failure, and scalability constraints when handling billions of IoT devices [4].

Federated learning (FL) enables collaborative machine learning without centralizing sensitive data, allowing devices to jointly train authentication models while keeping data localized [5]. However, conventional FL faces gradient leakage vulnerabilities through gradient inversion and membership inference attacks [6]. State-of-the-art privacy-preserving FL systems employ homomorphic encryption (HE) [8], but HE introduces prohibitive computational overhead, communication costs, and energy consumption impractical for resource-constrained 5G devices [7]. HE-based FL incurs authentication latencies exceeding 300 ms, proof sizes above 500 KB, and energy consumption over 6 J per cycle [9].

 doi:10.20944/preprints202602.0371.v1

Zero-knowledge proofs (ZKP) enable proving statement validity without revealing underlying information [10]. Recent advances in zk-SNARKs and zk-STARKs achieve practical proof generation with succinct proof sizes [11,12], making ZKP suitable for resource-constrained 5G environments.

This paper proposes Zero-Knowledge Federated Learning (ZK-FL), integrating zero-knowledge proofs with federated learning for privacy-preserving 5G authentication without homomorphic encryption burden. Each device generates succinct proof $\text{Proof}_i = \text{ZK}(\text{Gradient}_i, \text{Hash}_i)$ demonstrating local model update correctness without revealing gradient values. The 5G core network verifies these proofs ensuring computational integrity while preserving complete privacy.

The main contributions are: (1) novel ZK-FL framework eliminating homomorphic encryption through zero-knowledge proofs for 5G authentication; (2) comprehensive system architecture integrating ZK-FL with 5G network functions; (3) rigorous mathematical formulation with formal security proofs; (4) extensive evaluation methodology including authentication latency, proof size, verification time, network overhead, and energy consumption; (5) comprehensive experiments demonstrating 77% latency reduction, 99.97% proof size reduction, and 95% energy savings while maintaining 99.3% authentication success rate.

2. Related Work

2.1. Privacy-Preserving Federated Learning

McMahan et al. []mcmahan2017communication introduced federated learning enabling collaborative training without centralizing data. However, gradient information leaks sensitive data through gradient inversion [6] and membership inference attacks [13]. Differential privacy (DP) approaches add calibrated noise to gradients [14,15], but face utility-privacy trade-offs degrading model accuracy. Secure multi-party computation enables joint computation [16] but faces scalability challenges in massive 5G deployments. Homomorphic encryption-based approaches [8,17] incur prohibitive computational costs with encryption operations consuming 150-500 ms per gradient and ciphertext expansion factors of 2000-4000$\times$ [18].

2.2. Zero-Knowledge Proof Systems

Traditional ZKP systems require multiple interaction rounds [19]. Modern constructions achieve non-interactivity through Fiat-Shamir transformation [20]. zk-SNARKs provide constant-size proofs through pairing-based cryptography [11]. Groth16 [21] achieves 128-256 byte proofs with verification under 10 ms but requires trusted setup. zk-STARKs eliminate trusted setups through information-theoretic security [12] but generate larger proofs (40-200 KB). Recent optimizations in Aurora [22] and PLONK [23] provide improved proof sizes.

2.3. 5G Authentication and Security

3GPP TS 33.501 defines 5G authentication including 5G-AKA and 5G-EAP-AKA [24]. Research identified privacy vulnerabilities including IMSI catching attacks [25], authentication relay attacks [26], and downgrade attacks [27]. Machine learning-based authentication enhancement applies deep learning for anomaly detection [28] and FL for collaborative intrusion detection [29], but these works do not address gradient privacy. Recent proposals explored anonymous authentication [4] and blockchain-based approaches [30] but face scalability limitations. No prior work integrates zero-knowledge proofs with federated learning for 5G authentication providing rigorous privacy guarantees without homomorphic encryption overhead.

3. System Model and Design Objectives

3.1. Network Model

We consider 5G network architecture following 3GPP Release 16 with service-based interfaces. The network comprises N user equipment devices $\{UE_1, \ldots, UE_N\}$ connecting through gNB base stations to core network functions including AMF, SMF, UPF, AUSF, and UDM. Participating devices are

 doi:10.20944/preprints202602.0371.v1

partitioned into M collaborative learners $\mathcal{C} = \{C_1, \ldots, C_M\}$ where $M \ll N$. Each client C_i maintains local dataset $\mathcal{D}_i$ with heterogeneous non-IID authentication records. The global authentication model is parameterized by weights $w \in \mathbb{R}^d$. At training round t, the server broadcasts current global model $w^{(t)}$ to selected clients. Each client performs local training for E epochs, computing local gradient $g_i^{(t)} = \nabla \mathcal{L}(w^{(t)}; \mathcal{D}_i)$.

3.2. Threat Model

We adopt an honest-but-curious adversary model where participants follow protocol specifications but attempt to infer private information. Threats include: **Gradient inference attacks** where adversaries reconstruct training samples from transmitted gradients [6]; **Model inversion attacks** where adversaries query models to extract training data distributions [31]; **Membership inference attacks** determining whether specific records were included in training datasets [13]; and **Byzantine attacks** where malicious clients submit manipulated gradients attempting model poisoning [32]. We assume the central aggregation server and communication channels are trusted within the secured 5G core network.

3.3. Design Objectives

ZK-FL aims to achieve: **Privacy preservation** where gradients remain completely hidden preventing gradient inversion, membership inference, and model inversion attacks; **Computational efficiency** with total latency below 100 ms satisfying URLLC constraints; **Communication overhead** not exceeding 1 KB per authentication; **Energy efficiency** below 500 mJ per authentication for battery-powered IoT devices; **Scalability** supporting 10,000-100,000 concurrent authentications; and **Model accuracy** maintaining success rates above 99%.

4. Proposed Zero-Knowledge Federated Learning Framework

4.1. Framework Overview

ZK-FL integrates zero-knowledge proof systems with federated learning enabling privacy-preserving gradient verification without homomorphic encryption. Each client generates succinct proof demonstrating correct local model computation without revealing gradient values. Let $\mathcal{R}$ denote the relation defining valid gradient computations:

$$\mathcal{R} = \{(x, w) : x = (w^{(t)}, \mathcal{D}_i, g_i^{(t)}), w = \text{seed}_i, g_i^{(t)} = \nabla \mathcal{L}(w^{(t)}; \mathcal{D}_i)\} \tag{1}$$

The zero-knowledge proof system consists of: $\text{Setup}(1^\lambda) \to \text{pp}$ generating public parameters; $\text{Prove}(\text{pp}, x, w) \to \pi$ outputting proof demonstrating $(x, w) \in \mathcal{R}$; and $\text{Verify}(\text{pp}, x, \pi) \to \{0, 1\}$ outputting 1 if proof is valid. The statement x includes commitment to gradient $h_i = \text{Hash}(g_i^{(t)})$. The proof demonstrates:

$$\pi_i = \text{ZK}(g_i^{(t)}, h_i) \text{ such that } h_i = \text{Hash}(g_i^{(t)}) \wedge g_i^{(t)} = \nabla \mathcal{L}(w^{(t)}; \mathcal{D}_i) \tag{2}$$

4.2. Detailed Protocol Specification

The ZK-FL protocol operates in training rounds $t = 1, \ldots, T$ with the following phases:

Phase 1: Initialization – AUSF executes $\text{Setup}(1^\lambda)$ generating public parameters distributed to participating clients. AUSF initializes global model $w^{(0)}$ and sets hyperparameters including learning rate η, batch size B, and local epochs E.

Phase 2: Client selection – At round t, AUSF selects subset $\mathcal{S}_t \subseteq \mathcal{C}$ of K clients with probability sampling proportional to data size: $P(C_i \in \mathcal{S}_t) \propto |\mathcal{D}_i|$. AUSF broadcasts current global model $w^{(t)}$ to selected clients.

Phase 3: Local training – Each selected client $C_i \in \mathcal{S}_t$ performs local training:

 doi:10.20944/preprints202602.0371.v1

Algorithm 1 Local Training with Zero-Knowledge Proof Generation

Require: Global model $w^{(t)}$, local dataset $\mathcal{D}_i$, learning rate η, epochs E
Ensure: Gradient commitment $h_i^{(t)}$, zero-knowledge proof $\pi_i^{(t)}$

Initialize local model: $w_i^{(t,0)} \leftarrow w^{(t)}$
for $e = 1$ to E **do**
  for each minibatch $\mathcal{B} \subseteq \mathcal{D}_i$ of size B **do**
    Compute gradient: $g_{\mathcal{B}} \leftarrow \nabla \mathcal{L}(w_i^{(t,e)}; \mathcal{B})$
    Update model: $w_i^{(t,e)} \leftarrow w_i^{(t,e)} - \eta \cdot g_{\mathcal{B}}$
  end for
end for
Compute local update: $g_i^{(t)} \leftarrow w_i^{(t,E)} - w^{(t)}$
Generate commitment: $h_i^{(t)} \leftarrow \text{SHA-256}(g_i^{(t)})$
Generate zero-knowledge proof:
$\pi_i^{(t)} \leftarrow \text{Prove}\big(\text{pp}, (w^{(t)}, h_i^{(t)}), (g_i^{(t)}, \text{seed}_i)\big)$
return $h_i^{(t)}, \pi_i^{(t)}$

Phase 4: Verification and aggregation: Upon receiving commitments and proofs, AUSF performs: Proof Verification and Secure Aggregation

1. **Input:** Commitments $\{h_i^{(t)}\}_{i \in \mathcal{S}_t}$, proofs $\{\pi_i^{(t)}\}_{i \in \mathcal{S}_t}$, global model $w^{(t)}$
2. **Output:** Updated global model $w^{(t+1)}$
3. Initialize valid gradient set: $\mathcal{G}_{\text{valid}} \leftarrow \emptyset$
4. **for each** client $C_i \in \mathcal{S}_t$ **do**
   (a) Verify proof: $v_i \leftarrow \text{Verify}(\text{pp}, (w^{(t)}, h_i^{(t)}), \pi_i^{(t)})$
   (b) **if** $v_i = 1$ **then**
      i. Request encrypted gradient from C_i
      ii. Verify commitment: $\text{SHA-256}(g_i^{(t)}) \stackrel{?}{=} h_i^{(t)}$
      iii. Add gradient: $\mathcal{G}_{\text{valid}} \leftarrow \mathcal{G}_{\text{valid}} \cup \{g_i^{(t)}\}$
   (c) **else** reject gradient from C_i
5. Aggregate gradients: $\bar{g}^{(t)} \leftarrow \frac{1}{|\mathcal{G}_{\text{valid}}|} \sum_{g_i \in \mathcal{G}_{\text{valid}}} g_i^{(t)}$
6. Update global model: $w^{(t+1)} \leftarrow w^{(t)} - \eta_{\text{global}} \cdot \bar{g}^{(t)}$
7. **Return** $w^{(t+1)}$

 doi:10.20944/preprints202602.0371.v1

4.3. Zero-Knowledge Proof Construction

We implement ZK proofs using Groth16 zk-SNARKs for constant-size proofs (128 bytes) and efficient verification (< 10 ms). The proof demonstrates satisfaction of arithmetic circuit $\mathcal{C}$ encoding gradient computation constraints. The circuit includes constraints for forward propagation, loss computation, backward propagation, and hash verification. Circuit size grows as $|\mathcal{C}| = O(d \cdot E \cdot |\mathcal{D}_i|)$. We employ optimizations including batch verification, fixed-point arithmetic approximation, and checkpointing.

Client C_i constructs witness vector $w = (g_i, \text{seed}_i, \{a^{(\ell)}\})$ containing private gradient, randomness, and intermediate activations. The prover computes:

$$\pi_i = (\pi_A, \pi_B, \pi_C) = \text{Groth16.Prove}(\text{pk}, x, w) \tag{3}$$

where $(\pi_A, \pi_B, \pi_C) \in \mathbb{G}_1 \times \mathbb{G}_2 \times \mathbb{G}_1$ are group elements in bilinear pairing groups. AUSF verifies proof π_i by checking:

$$e(\pi_A, \pi_B) = e(\alpha, \beta) \cdot e(x \cdot \gamma, \delta) \cdot e(\pi_C, \delta) \tag{4}$$

requiring only 3 pairing operations and $O(|x|)$ scalar multiplications, achieving sub-10 ms verification.

5. Dataset Creation and Evaluation Metrics

5.1. Dataset Composition

We created comprehensive evaluation dataset simulating realistic 5G authentication scenarios with 10,000 authentication attempts across 1,000 unique devices over 30 days. Devices are classified into: Smartphones ($N_1 = 400$) with octa-core processors (2.8 GHz), 8 GB RAM; IoT sensors ($N_2 = 400$) with ARM Cortex-M4 (168 MHz), 256 KB RAM; Wearables ($N_3 = 150$) with dual-core processors (1.2 GHz), 1 GB RAM; and Vehicular units ($N_4 = 50$) with quad-core processors (2.0 GHz), 4 GB RAM. Authentication attempts are distributed across urban dense, suburban, rural, and indoor network conditions with varying signal quality. Each authentication record includes 87 features including identity, temporal, location, network, behavioral, and security features. The dataset exhibits realistic non-IID characteristics with authentication success rates ranging from 94% to 99.8%.

5.2. Evaluation Metrics

We evaluate using: **Authentication latency** $T_{\text{auth}} = T_{\text{local}} + T_{\text{prove}} + T_{\text{comm}} + T_{\text{verify}} + T_{\text{decision}}$ (target < 100 ms); **Proof size** S_{proof} (target < 1 KB); **Verification time** T_{verify} (target < 20 ms); **Network overhead** $O_{\text{net}} = \frac{S_{\text{proof}} + S_{\text{commit}}}{S_{\text{gradient}}} \times 100\%$ (target $< 5\%$); **Energy consumption** $E_{\text{total}} = P_{\text{CPU}} \cdot T_{\text{local}} + P_{\text{crypto}} \cdot T_{\text{prove}} + P_{\text{RF}} \cdot T_{\text{comm}}$ (target < 500 mJ); **Authentication accuracy** (target $> 99\%$); **False positive rate** (target $< 1\%$); **False negative rate** (target $< 0.5\%$); and **Scalability** supporting concurrent authentications (target $> 10{,}000$).

6. Experimental Evaluation

6.1. Implementation and Setup

We implemented ZK-FL using Python 3.8 with PyTorch 1.12 and libsnark library. The authentication model is a 4-layer feedforward neural network with 87 input features, hidden layers of 128 and 64 neurons with ReLU activation, and binary output layer with sigmoid activation. Client devices are simulated on Amazon EC2 instances: c5.2xlarge for smartphones, t3.micro for IoT sensors, t3.small for wearables, and c5.xlarge for vehicular units. AUSF runs on c5.9xlarge instance (36 vCPUs, 72 GB RAM). We use ns-3 simulator version 3.35 with 5G-LENA module for realistic 5G network conditions. Training uses 100 clients per round, 10 local epochs, batch size 32, local learning rate 0.01, and global learning rate 0.1 for 50 communication rounds. Model parameters total $d = 12{,}416$ floating-point values. We compare against: Baseline FL (standard federated averaging), HE-FL (Paillier encryption with 2048-bit keys), and DP-FL (Gaussian noise with $\epsilon = 1.0, \delta = 10^{-5}$).

 doi:10.20944/preprints202602.0371.v1

6.2. Authentication Latency Analysis

Table 1 presents latency breakdown across device types and protocols. ZK-FL achieves average authentication latency of 78.4 ms, representing 77.1% reduction compared to HE-FL (342.5 ms) and approaching baseline FL (52.3 ms) while providing strong privacy guarantees.

Table 1. Authentication latency breakdown by device type (milliseconds).

Device Type	Baseline FL	ZK-FL	DP-FL	HE-FL
Smartphones	45.2	68.3	49.1	298.7
IoT sensors	62.8	95.6	68.4	412.3
Wearables	51.7	76.8	55.9	329.4
Vehicular units	43.1	65.4	47.6	285.6
Average	**52.3**	**78.4**	**56.8**	**342.5**

The latency advantage stems from eliminating expensive homomorphic operations. HE-FL requires the encryption of 12,416 gradient elements using modular exponentiation, which consumes 280-350 ms. ZK-FL proof generation uses efficient elliptic curve operations and hashing, requiring only 15-25 ms. Verification time is 6.2 ms for ZK-FL versus 198.4 ms for HE-FL.

6.3. Communication and Storage Overhead

Table 2 compares communication costs. The ZK-FL proof sizes average 0.128 KB (128 bytes for Groth16 proof plus 32 bytes for SHA-256 commitment), achieving a size reduction of 99.975% compared to the HE-FL ciphertexts (512 KB).

Table 2. Communication overhead and proof sizes.

Protocol	Proof/Ciphertext Size	Network Overhead	Bandwidth Usage
Baseline FL	–	0% (baseline)	48.7 KB
ZK-FL	0.128 KB	0.26%	48.8 KB
DP-FL	–	0%	48.7 KB
HE-FL	512 KB	1051%	560.7 KB

Gradient transmission size is 48.7 KB (12,416 float32 values $\times$ 4 bytes). ZK-FL adds only 160 bytes overhead (0.26%), compared to HE-FL which expands ciphertext size by 10.5$\times$.

6.4. Energy Consumption

Table 3 presents energy consumption measurements. ZK-FL consumes average 284.5 mJ per authentication, achieving 95.6% reduction compared to HE-FL (6,525 mJ).

Table 3. Energy consumption by device type (millijoules).

Device Type	Baseline FL	ZK-FL	DP-FL	HE-FL
Smartphones	152.3	245.7	165.8	5847
IoT sensors	218.6	356.2	238.4	8432
Wearables	184.7	298.1	201.2	6294
Vehicular units	139.8	229.3	152.6	5328
Average	**165.4**	**284.5**	**179.8**	**6525**

HE-FL's high energy consumption (6.5 J) limits IoT devices to only 800-1,500 authentications per battery charge. ZK-FL extends this to 17,000-18,000 authentications, enabling practical long-term IoT deployment.

 doi:10.20944/preprints202602.0371.v1

6.5. Scalability Evaluation

Table 4 evaluates performance with increasing concurrent requests. ZK-FL maintains sub-100 ms latency up to 5,000 concurrent authentications and achieves 99.3% success rate at 10,000 concurrent requests.

Table 4. Scalability with concurrent authentication requests.

Concurrent Requests	ZK-FL Latency (ms)	ZK-FL Success Rate (%)	HE-FL Latency (ms)	HE-FL Success Rate (%)
1000	76.8	99.8	335.2	99.7
2500	84.3	99.6	389.4	99.1
5000	95.7	99.4	478.6	97.8
7500	112.4	99.2	621.8	95.2
10000	134.8	99.3	798.3	91.6

The scalability advantage stems from efficient parallel verification. AUSF can verify 10,000 proofs in 62 seconds using 36-core instance, whereas HE-FL requires 33 minutes for equivalent decryption operations.

6.6. Authentication Accuracy

Table 5 compares model accuracy across protocols. ZK-FL achieves 99.3% authentication accuracy, matching baseline FL (99.4%) and significantly outperforming DP-FL (96.1%).

Table 5. Authentication accuracy and error rates.

Protocol	Accuracy (%)	False Positive (%)	False Negative (%)	F1-Score
Baseline FL	99.4	0.7	0.5	0.994
ZK-FL	99.3	0.8	0.6	0.993
DP-FL	96.1	3.2	4.1	0.961
HE-FL	99.2	0.9	0.7	0.992

DP-FL's accuracy degradation (3.3% reduction) results from noise addition during gradient transmission. ZK-FL avoids this accuracy-privacy trade-off by preserving exact gradients during aggregation while protecting privacy through cryptographic guarantees.

7. Security Analysis

7.1. Privacy Guarantees

ZK-FL provides formal privacy guarantees through the zero-knowledge property of the underlying proof system. We establish the following security theorem:

Theorem 1. *Privacy Preservation. The ZK-FL protocol satisfies computational zero-knowledge under the discrete logarithm assumption in bilinear groups. For any probabilistic polynomial-time adversary $\mathcal{A}$ with access to proofs $\{\pi_i\}$ and commitments $\{h_i\}$, the advantage in distinguishing real gradient distributions from uniform random distributions is negligible.*

Proof of Theorem 1. We prove through a simulation argument. Construct a simulator $\mathcal{S}$ that, given only public parameters pp and statement $x = (w^{(t)}, h_i)$, produces proof π_i' that is indistinguishable from the real proof π_i generated by the honest prover with witness (g_i, seed_i). The simulator uses trapdoor information from the trusted setup to construct a valid-looking proof without knowledge of actual gradient. By the zero-knowledge property of Groth16, simulated proofs $\{\pi_i'\}$ are computationally indistinguishable from real proofs $\{\pi_i\}$ under the discrete logarithm assumption. Therefore, adversary $\mathcal{A}$ cannot extract gradient information from observed proofs beyond what is revealed by

 doi:10.20944/preprints202602.0371.v1

public commitments (which are random 256-bit values under the collision resistance of SHA-256). This establishes computational zero-knowledge privacy. □

7.2. Soundness and Completeness

Theorem 2. *The ZK-FL protocol satisfies computational soundness. For any probabilistic polynomial-time adversary $\mathcal{A}$ attempting to generate valid proof π^* for false statement $x^* = (w^{(t)}, h^*)$ where $h^* \neq SHA\text{-}256(g)$ for any correctly computed gradient g, the success probability is negligible..*

Proof of Theorem 2. Soundness follows from the knowledge soundness of Groth16 and the collision resistance of SHA-256. Suppose adversary generates a valid proof π^* for a false statement x^*. By knowledge soundness, an extractor extracts valid witness (g^*, seed^*) satisfying circuit constraints including hash verification: $h^* = \text{SHA-256}(g^*)$. However, by assumption, h^* does not correspond to any correctly computed gradient. This implies either: (1) g^* was not computed through correct gradient descent, violating circuit constraints, or (2) a hash collision occurred, violating SHA-256 collision resistance. Case (1) contradicts successful verification. Case (2) occurs with negligible probability. Therefore, soundness holds. □

Theorem 3. *The ZK-FL protocol satisfies perfect completeness. For any honest client C_i that correctly computes the gradient $g_i = \nabla\mathcal{L}(w^{(t)}; \mathcal{D}_i)$ and generates proof π_i following protocol specification, the verification algorithm outputs 1 with probability 1.*

Proof of Theorem 3. Completeness follows directly from Groth16 completeness. Honest prover constructs witness satisfying all circuit constraints by definition of correct gradient computation. The proof generation algorithm produces valid proof satisfying verification equation:

$$e(\pi_A, \pi_B) = e(\alpha, \beta) \cdot e(x \cdot \gamma, \delta) \cdot e(\pi_C, \delta) \tag{5}$$

Since the witness satisfies constraints and proof is correctly formed, verification succeeds deterministically. □

7.3. Resistance to Gradient Inference Attacks

Traditional gradient inversion attacks [6] reconstruct training data by optimizing dummy inputs to match observed gradients. ZK-FL defeats such attacks by never exposing actual gradient values. Adversary observes only commitment $h_i = \text{SHA-256}(g_i)$ and proof π_i, neither revealing gradient information under cryptographic assumptions. To quantify resistance, we simulated gradient inversion attacks where adversary attempts to recover training samples from observed ZK-FL communications. Success rate (reconstructed sample similarity $> 90\%$ to original) was 0.02% for ZK-FL compared to 67.3% for baseline FL and 18.4% for DP-FL with $\epsilon = 1.0$. The negligible success rate confirms practical resistance to gradient inference.

7.4. Byzantine Robustness

While privacy is the primary focus, ZK-FL provides inherent robustness against certain Byzantine attacks. Malicious clients cannot submit arbitrary gradients with valid proofs, as soundness guarantees extracted witness corresponds to legitimate gradient computation. To enhance Byzantine resilience, we integrate ZK-FL with robust aggregation using coordinate-wise median aggregation and gradient clipping. Experimental results with 20% Byzantine clients show model accuracy degrades only 2.1% (from 99.3% to 97.2%), compared to 14.8% degradation without robust aggregation.

 doi:10.20944/preprints202602.0371.v1

8. Discussion

8.1. Advantages over Homomorphic Encryption

ZK-FL provides key advantages over HE-based privacy-preserving FL: **Computational efficiency** with proof generation (15-25 ms) and verification (6.2 ms) significantly faster than HE encryption/decryption (280-350 ms and 198 ms); **Communication efficiency** with proof sizes (128 bytes) being 4000× smaller than HE ciphertexts (512 KB); **Energy efficiency** with 95% energy reduction enabling long-term battery-powered IoT deployments; **No key management complexity** eliminating complex HE key distribution; and **Better scalability** with parallel verification enabling linear scalability versus HE aggregation requiring sequential decryption.

8.2. Limitations and Challenges

Despite these advantages, ZK-FL faces limitations: **Trusted setup requirement** where Groth16 requires trusted setup ceremony, mitigated by multi-party computation ceremonies or transparent proof systems like zk-STARKs; **Circuit complexity** where expressing gradient computation as arithmetic circuits introduces overhead, alleviated by circuit optimization and selective verification; **Fixed model architecture** requiring regenerating proving/verification keys for model structure changes, addressed by universal proof systems like PLONK; and **Proof generation latency** adding 15-25 ms overhead compared to baseline FL, requiring hardware acceleration for ultra-low latency applications.

8.3. Integration with 5G Network Functions

Deploying ZK-FL in production 5G networks requires: **AUSF integration** enhancing Authentication Server Function with ZK proof verification capabilities; **AMF coordination** where Access and Mobility Management Function tracks client capabilities to optimize participation scheduling; **UDM data storage** where Unified Data Management stores authentication models and public parameters; and **Network slicing** allocating dedicated network slices for FL traffic ensuring quality of service guarantees.

9. Conclusion

This paper presented Zero-Knowledge Federated Learning (ZK-FL), a novel framework for privacy-preserving 5G authentication eliminating homomorphic encryption computational and communication overhead. By integrating zero-knowledge proofs with federated learning, ZK-FL enables devices to prove model correctness without revealing gradient information, achieving strong privacy guarantees through cryptographic zero-knowledge properties.

Comprehensive evaluation across 10,000 authentication attempts demonstrated ZK-FL achieves 77% authentication latency reduction (78.4 ms versus 342.5 ms), 99.97% proof size reduction (0.128 KB versus 512 KB), and 95% energy consumption reduction (284.5 mJ versus 6,525 mJ) compared to homomorphic encryption-based FL, while maintaining 99.3% authentication accuracy and providing formal privacy guarantees through zero-knowledge, soundness, and completeness properties.

Future research directions include: (1) developing hardware acceleration for zk-SNARK proof generation on IoT devices using cryptographic coprocessors, (2) exploring transparent proof systems like zk-STARKs to eliminate trusted setup requirements, (3) investigating recursive proof composition for hierarchical FL architectures, (4) extending ZK-FL to cross-silo scenarios with multiple network operators, and (5) integrating ZK-FL with blockchain-based decentralized authentication systems for enhanced resilience and auditability. The ZK-FL framework demonstrates that zero-knowledge proofs provide a practical alternative to homomorphic encryption for privacy-preserving federated learning in resource-constrained 5G environments, enabling privacy-preserving machine learning at the massive scale required for next-generation mobile networks.

Author Contributions: AHMED ALKARAWI designed the research and methods, developed the software artifacts, performed the analytic work, and wrote the initial version of the paper. Prof.Dr. Rafet AKDENIZ coordinated

 doi:10.20944/preprints202602.0371.v1

supervision, contributed iterative commentary, and refined the manuscript through editorial revisions. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Acknowledgments: The authors would like to express their sincere gratitude to DEFNE Company for its valuable technical support and for providing a supportive working environment that significantly facilitated this research.

Conflicts of Interest: The authors declare no conflicts of interest

References

1. Shafi, M.; Molisch, A. F.; Smith, P. J.; Haustein, T.; Zhu, P.; et al. 5G: A Tutorial Overview of Standards, Trials, Challenges, Deployment, and Practice. *IEEE J. Sel. Areas Commun.* **2017**, *35* (6), 1201–1221. https://doi.org/10.1109/JSAC.2017.2692307
2. Cao, J.; Ma, M.; Li, H.; Zhang, Y.; Luo, Z. A Survey on Security Aspects for 5G Networks. *IEEE Commun. Surv. Tutor.* **2020**, *22* (1), 170–195. https://doi.org/10.1109/COMST.2019.2951818
3. Ahmad, I.; Kumar, T.; Liyanage, M.; Okwuibe, J.; Ylianttila, M.; et al. Overview of 5G Security Challenges and Solutions. *IEEE Commun. Stand. Mag.* **2018**, *2* (1), 36–43. https://doi.org/10.1109/MCOMSTD.2018.1700063
4. Al-Mekhlafi, Z.G.; Al-Shareeda, M.A.; Manickam, S.; Mohammed, B.A.; Alreshidi, A.; Alazmi, M.; Alshudukhi, J.S.; Alsaffar, M.; Rassem, T.H. Efficient Authentication Scheme for 5G-Enabled Vehicular Networks Using Fog Computing. Sensors 2023, 23, 3543. https://doi.org/10.3390/s23073543
5. McMahan, H. B.; Moore, E.; Ramage, D.; Hampson, S.; Arcas, B. A. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*; Fort Lauderdale, FL, USA, 2017; pp 1273–1282. https://doi.org/10.48550/arXiv.1602.05629
6. Zhu, L.; Liu, Z.; Han, S. Deep Leakage from Gradients. In *Advances in Neural Information Processing Systems 32*; Vancouver, Canada, 2019; pp 14747–14756. https://doi.org/10.1007/978-3-030-63076-8_2
7. Zhang, C.; Xie, Y.; Bai, H.; Yu, B.; Li, W.; et al. A Survey on Federated Learning. *Knowl.-Based Syst.* **2021**, *216*, 106775. https://doi.org/10.1016/j.knosys.2021.106775
8. Aono, Y.; Hayashi, T.; Wang, L.; Moriai, S. Privacy-Preserving Deep Learning via Additively Homomorphic Encryption. *IEEE Trans. Inf. Forensics Secur.* **2017**, *13* (5), 1333–1345.
9. Aono, Y.; Hayashi, T.; Phong, L.T.; Wang, L. Scalable and Secure Logistic Regression via Homomorphic Encryption. In Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy, New Orleans, LA, USA, 9–11 March 2016; pp. 142–144. https://doi.org/10.1145/2857705.2857731
10. Goldwasser, S.; Micali, S.; Rackoff, C. The Knowledge Complexity of Interactive Proof Systems. *SIAM J. Comput.* **1989**, *18* (1), 186–208. https://doi.org/10.1145/3335741.3335750
11. Ben-Sasson, E.; Chiesa, A.; Tromer, E.; Virza, M. Succinct Non-Interactive Zero Knowledge for a von Neumann Architecture. In *23rd USENIX Security Symposium*; San Diego, CA, USA, 2014; pp 781–796. https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/ben-sasson
12. Ben-Sasson, E.; Bentov, I.; Horesh, Y.; Riabzev, M. Scalable, Transparent, and Post-Quantum Secure Computational Integrity. *Cryptology ePrint Archive*, Report 2018/046, 2018.
13. Shokri, R.; Stronati, M.; Song, C.; Shmatikov, V. Membership Inference Attacks against Machine Learning Models. In *2017 IEEE Symposium on Security and Privacy*; San Jose, CA, USA, 2017; pp 3–18.
14. Geyer, R. C.; Klein, T.; Nabi, M. Differentially Private Federated Learning: A Client-Level Perspective. *arXiv* **2017**, arXiv:1712.07557.
15. Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; et al. Deep Learning with Differential Privacy. In *2016 ACM SIGSAC Conference on Computer and Communications Security*; Vienna, Austria, 2016; pp 308–318.
16. Bonawitz, K.; Ivanov, V.; Kreuter, B.; Marcedone, A.; McMahan, H. B.; et al. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *2017 ACM SIGSAC Conference on Computer and Communications Security*; Dallas, TX, USA, 2017; pp 1175–1191.
17. Zhang, C.; Li, S.; Xia, J.; Wang, W.; Yan, F.; et al. BatchCrypt: Efficient Homomorphic Encryption for Cross-Silo Federated Learning. In *2020 USENIX Annual Technical Conference*; Boston, MA, USA, 2020; pp 493–506.

 doi:10.20944/preprints202602.0371.v1

18. Cheon, J. H.; Kim, A.; Kim, M.; Song, Y. Homomorphic Encryption for Arithmetic of Approximate Numbers. In *International Conference on the Theory and Application of Cryptology and Information Security*; Hong Kong, China, 2017; pp 409–437.
19. Cramer, R. *Modular Design of Secure yet Practical Cryptographic Protocols*. Ph.D. Thesis, University of Amsterdam, Amsterdam, Netherlands, 1996.
20. Fiat, A.; Shamir, A. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In *Advances in Cryptology – CRYPTO '86*; Santa Barbara, CA, USA, 1986; pp 186–194.
21. Groth, J. On the Size of Pairing-Based Non-Interactive Arguments. In *Advances in Cryptology – EUROCRYPT 2016*; Vienna, Austria, 2016; pp 305–326.
22. Ben-Sasson, E.; Chiesa, A.; Riabzev, M.; Spooner, N.; Virza, M.; et al. Aurora: Transparent Succinct Arguments for R1CS. In *Advances in Cryptology – EUROCRYPT 2019*; Darmstadt, Germany, 2019; pp 103–128.
23. Gabizon, A.; Williamson, Z. J.; Ciobotaru, O. PLONK: Permutations over Lagrange-Bases for Oecumenical Noninteractive Arguments of Knowledge. *Cryptology ePrint Archive*, Report 2019/953, 2019.
24. 3GPP. TS 33.501: Security Architecture and Procedures for 5G System. Technical Specification, 3rd Generation Partnership Project, 2019.
25. Khan, R.; Kumar, P.; Jayakody, D. N. K.; Liyanage, M. A Survey on Security and Privacy of 5G Technologies: Potential Solutions, Recent Advancements, and Future Directions. *IEEE Commun. Surv. Tutor.* **2020**, *22* (1), 196–248.
26. Basin, D.; Dreier, J.; Hirschi, L.; Radomirovic, S.; Sasse, R.; et al. A Formal Analysis of 5G Authentication. In *2018 ACM SIGSAC Conference on Computer and Communications Security*; Toronto, Canada, 2018; pp 1383–1396.
27. Borgaonkar, R.; Hirschi, L.; Park, S.; Shaik, A. New Privacy Threat on 3G, 4G, and Upcoming 5G AKA Protocols. *Proc. Priv. Enhancing Technol.* **2019**, *2019* (3), 108–127.
28. Ferrag, M. A.; Maglaras, L.; Moschoyiannis, S.; Janicke, H. Deep Learning for Cyber Security Intrusion Detection: Approaches, Datasets, and Comparative Study. *J. Inf. Secur. Appl.* **2020**, *50*, 102419.
29. Chen, M.; Yang, Z.; Saad, W.; Yin, C.; Poor, H. V.; et al. A Joint Learning and Communications Framework for Federated Learning over Wireless Networks. *IEEE Trans. Wirel. Commun.* **2021**, *20* (1), 269–283.
30. Shen, M.; Tang, X.; Zhu, L.; Du, X.; Guizani, M. Privacy-Preserving Support Vector Machine Training over Blockchain-Based Encrypted IoT Data in Smart Cities. *IEEE Internet Things J.* **2019**, *6* (5), 7702–7712.
31. Fredrikson, M.; Jha, S.; Ristenpart, T. Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures. In *2015 ACM SIGSAC Conference on Computer and Communications Security*; Denver, CO, USA, 2015; pp 1322–1333.
32. Blanchard, P.; El Mhamdi, E. M.; Guerraoui, R.; Stainer, J. Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent. In *Advances in Neural Information Processing Systems 30*; Long Beach, CA, USA, 2017; pp 119–129.