



International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Unified Post-Quantum Cryptography For Cloud-Assisted IoT: Lattice Homomorphic Proxy Re-Encryption With Ghoa-Optimized Parameters and Explainable Security

Bommepalli Narayana Reddy ¹, B.Raja Koti ²

¹Research Scholar, Department of Computer Science Engineering, Gandhi Institute of Technology and Management· Visakhapatnam, Andhra Pradesh, IN

²Assistant Professor, Department of Computer Science Engineering, Gandhi Institute of Technology and Management· Visakhapatnam, Andhra Pradesh, IN

Corresponding author: Bommepalli Narayana Reddy
Email: [nbommepa@gitam.in]

Abstract

Medical and industrial IoT deployments require cryptographic solutions that provide confidentiality for data being computed in the cloud, enable secure delegation to authorized users without disclosing private keys, and remain quantum resilient under the energy and memory constraints of ARM Cortex-class devices. Current methods only partially address these requirements. This proposal arises from the observation that homomorphic encryption schemes can be constructed with a high level of abstraction. The result is a hybrid that selectively supports specific operations through a cryptographic transformation. The Operation homomorphic is categorized into several types. ALPHREN synthesizes four synergistic elements. The multi-objective Ring-LWE parameter optimization problem (which considers security, encryption latency, and memory consumption) is tackled with a Genetic Honey Optimization Algorithm (GHOA). The GHOA converges to around 47 iterations and an optimal parameter set ($n = 1024$, $q = 12289$, $\sigma = 3.19$). The Batch Fan-Vercauteren homomorphic encryption engine can carry out operations on a depth $d \leq 7$ arithmetic circuits whilst allowing for a noise growth that is correctable. One more, a method of one-way proxy re-encryption reliably changes ciphertexts between users without disclosing plaintext. Ultimately, a SHAP (SHapley Additive exPlanations) module interprets the contribution of each parameter to the security-performance trade-off to enable HIPAA and GDPR compliant transparency. We ran the framework on Raspberry Pi 4B and used the datasets MIMIC-III (48,321 clinical records) and UCI IoT Intrusion (625,783 records) for 30 independent trials. ALPHREN was able to generate the keys in 0.43 ms, achieve an encryption throughput of 2.14 MB/s and an energy consumption of 4.7 μ J per operation with a security classification accuracy of 99.1%. Encryption latency was 32.8 $\times$ lower than the closest PQC competitor. Both the Friedman test ($\chi^2(6) = 44.19$, $p < 0.001$) and Bonferroni-corrected Wilcoxon tests showed these results to be statistically significant and of large effect size (Cohen's $d = 1.39$ – 2.11). ALPHREN is the first framework that integrates post-quantum security, homomorphic evaluation, proxy re-encryption and explainable cryptographic parameter optimization on resource-constrained IoT hardware.


Keywords: Ring Learning With Errors; lattice-based cryptography; homomorphic encryption; proxy re-encryption; cloud-assisted IoT; post-quantum cryptography; SHAP explainability; IND-CCA2 security; medical data privacy; meta-heuristic parameter optimization.

1. Introduction

It is a present-day fact that medical IoT deployments are not a future proposition. Clouds have become a common destination for physiological telemetry from cardiac monitors, infusion pumps, continuous glucose sensors, and ventilator controllers for collection, alerting, and clinical decision-making. Within this decade, more than 500 million devices used in healthcare will be connected according to the latest survey-based projections. These architectures are necessary for operation but at the same time give rise to a structural security problem: the cloud intermediaries providing computation must, as per a rigorous threat model, be treated as adversarial [4]. However, the data being processed by these intermediaries is sensitive identifiable patient records subject to HIPAA Section 164.312 and GDPR Article 25 [12]. The sensitive data must remain confidential, and also be auditable for regulatory compliance throughout the pipeline. Classical public-key cryptography has served its purpose well against today's adversaries, but the assumption underlying its security is expiring. The integer factorization problem which forms the hardness assumption of RSA is reduced by Shor's algorithm (i.e., made easier) to polynomial-time on a fault-tolerant quantum computer. Halving the effective bit-security of symmetric schemes like AES is Grover's algorithm. Today, the question is no longer

if large-scale quantum computers will exist, but when will it happen? IBM wants systems with 100,000 qubits by 2033. Google's experiments have been breaking computational records since 2019. Given this, systems with decade-long lifetimes, exactly like deployed medical infrastructure, need cryptographic innovation now. NIST published FIPS 203, 204, and 205 in 2024 to formalize its response to establish lattice-based schemes as NIST's main post-quantum cryptography (PQC) standard [8]. Nonetheless, standardization does not solve deployment: putting PQC on constrained IoT hardware introduces implementation complexity and performance overhead that commodity protocols do not experience.

Lattice-based cryptography the Ring Learning With Errors problem that is Ring-LWE has emerged as the front-runner



for constrained IoT environments due to its $O(n \log n)$ polynomial multiplication complexity via Number Theoretic Transform, compact key representations and a worst-case hardness reduction to Shortest Independent Vectors Problem (SIVP) in ideal lattices [17]. Recent benchmarking efforts by Zhao et al. [21] have resulted in an efficient NTT implementation on an ARM Cortex-A processor running an operating system, achieving key generation in less than a millisecond. Also, Ring-LWE implementations on embedded platforms have matured considerably since the pqm4 benchmarking efforts [20]. Nonetheless, just achieving post-quantum resistance is not enough for cloud-assisted IoT architectures, which impose three additional requirements. The current Ring-LWE deployments do not satisfy these requirements simultaneously. The first requirement is homomorphic computations: all cloud-side analytics must work over encrypted data, without decryption. This protects data sovereignty and enables privacy-preserving machine learning and aggregate statistics. The BFV scheme and its successors like the OpenFHE library provide arithmetic circuit evaluation under encryption, but reported 892 ms per batch latencies on server hardware have always precluded their deployment on IoT-constrained pipelines. The second requirement concerns secure re-encryption delegation. In a multi-party clinical world, data produced by one entity – say, a bedside sensor controlled by a patient's care team – needs to be re-encrypted for consumption by another (the analytics service or specialist clinician) without intermediaries ever

seeing private key material. Proxy re-encryption (PRE), introduced by Ateniese et al. and latter extended to the lattice setting by Zhou et al., is a solution for delegation. Miao et al. show a different application of PRE which is for cloud-IoT data sharing. Regulatory explainability is the third requirement which requires documentation and justification of algorithms according to HIPAA and GDPR. None of the existing lattice-based IoT (Internet of Things) systems provide a mechanism for non-expert operators to audit or interpret the cryptographic parameters governing their deployments. There is an important gap between what regulators expect and what the technology can deliver; this has been the basis of GDPR enforcement actions in medical data contexts since 2021. A systematic review of the literature corroborates that no existing scheme delivers all four requirements post-quantum security, homomorphic evaluation, proxy re-encryption, and parameter explainability within the computational envelope of commodity IoT hardware. Chaudhary et al. performed attribute based encryption for mobile edge IoT which does not have homomorphic capabilities. Khalid et al. [14] propose post-quantum non-enabled blocking authentication. The IoT security architecture has been discussed in detail in paper [15] by Bao et al. However, their focus is only limited to pre-quantum symmetric primitives. Recent surveys by Nejatollahi et al. [18] and Tawalbeh et al. [16] found Ring-LWE the preferred post-quantum candidate for IoT applications, although there are open engineering issues for homomorphic integration. Works addressing HE in the cloud-medical context [28] only operate on server hardware and do not support multi-party delegation. The constructions for lattice PRE by Zhou et al. [31] and Wan et al. [32] do not support homomorphic evaluation. None pay attention to SHAP-based explanation for cryptographic parameters despite a growing pool of XAI literature which is establishing SHAP as the standard attribution mechanism for complex fitness functions [37, 38, 39]. At least to this day and age, no literature addressed the intersection of these four requirements on physical IoT hardware. The introduction of ALPHREN (Adaptive Lattice-Homomorphic Proxy Re-encryption Network) in this document addresses the said user. ALPHREN has Key generation based on Ring-LWE under optimum parameters determined with the use of Genetic Honey Optimization Algorithm (GHOA). It also has a BFV homomorphic evaluation engine supporting circuit of depth $d \leq 7$, a unidirectional proxy re-encryption layer providing IND-PRE-CCA2 security and a SHAP based explainability module for regulatory compliance. The results obtained on Raspberry Pi 4B using the MIMIC-III clinical dataset [40] and UCI IoT Intrusion dataset confirm that the key generation takes place at 0.43 ms at a speed throughput of 2.14 MB/s and energy consumed is 4.7 μ J per operation. Finally, we achieve a security classification accuracy of 99.1%. These results are statistically validated against six competing schemes using the Friedman test ($\chi^2(6) = 39.1$, $p < 0.01$). = 44.19, $p < 0.001$), while Wilcoxon post-hoc tests with Bonferroni correction provided large effect sizes (Cohen's d : 1.39–2.11). The following are the contributions of the present work. The ALPHREN framework provides the first complete system that integrates Ring-LWE encryption, BFV homomorphic computation, proxy re-encryption, and SHAP explainability on commodity IoT hardware. It achieves formal IND-CCA2 security under Fujisaki-Okamoto transform and IND-PRE-CCA2 security under Random Oracle Model. GHOA is a new meta-heuristic that combines the genetic crossover (of Genetic Algorithm) with the honey bee foraging (of ABC algorithm). Results show that GHOA yields optimal Ring-LWE parameters ($n^* = 1024$, $q^* = 12289$, $\sigma^* = 3.19$) in about 47 iterations, while standard genetic algorithm takes about 180. Furthermore, the convergence of GHOA is guaranteed by ergodicity of the Markov chain over the finite parameter lattice. Using SHAP value attributions. The selection of other cryptographic parameters through explanations of the GHOA's multi-objective fitness function is presented here for the first-time quality parameters importance scores n : 0.41, σ : 0.33, q : 0.26 satisfying HIPAA and GDPR audit measures [12]. In conclusion, the empirical methodology applies a statistically rigorous protocol to four datasets, six competitors and two hardware platforms. Effect size reporting is consistent with recent recommendations for randomized algorithm benchmarking [42]. As follows the rest of document is organised Section 2 examines previously published studies about post-quantum crypto systems for IoT devices along with other approaches. The ALPHREN Architecture and Trust model is in Section 3. Section 4 delves into the mathematics. Section 5 describes the GHOA algorithm and its convergence analysis. Section 6 outlines the experimental methods. Sections seven through nine show performance results and SHAP explainability and ablation analysis. Section 10 examines restrictions. Section 11 is the end.

Fig. 1. Overall architecture of the proposed ALPHREN framework, illustrating the end-to-end workflow from IoT data acquisition and Ring-LWE key generation through proxy re-encryption, homomorphic evaluation, secure cloud processing, attack detection, statistical validation, and secure data recovery under a post-quantum security framework.

2. Related Work

2.1 Post-Quantum Cryptography for Constrained IoT Devices

The NIST standardization process [8] as well as implementation research focusing on ARM Cortex (see, e.g., [17,18]) shape the landscape of post-quantum cryptography for IoT devices. In their paper, Sinha et al. [18] survey lightweight post-quantum cryptographic (PQC) schemes for the Internet of Things (IoT). They identify Ring-LWE as the most

promising candidate. It has $O(n \log n)$ multiplication via Number Theoretic Transform (NTT). The public key sizes are 1.5 kilobytes at $n=1024$. Furthermore, it has worst-case hardness reduction to the Shortest Vector Problem in ideal lattices. Pereira et al. [19] demonstrate that their Ring-LWE implementation for ARM Cortex-M4 achieves key generation efficiency of 1.1 ms, which is $2.6\times$ slower than alternatives. The overhead by the presence of our additional GHOA optimization and proxy re-encryption component has a significant cost impact. Howe et al. [20] benchmark Kyber-512 (KEM derived from Module-LWE, standard of NIST) on Raspberry Pi 3 with 0.18 ms in keygen. ALPHREN runs directly on Ring-LWE and incurs 0.43 ms key generation on Raspberry Pi 4B but offers homomorphic evaluation and proxy re-encryption not possible in KEM designs. Khalid and colleagues [21] the techniques they develop for optimizing NTT assembly implementation for ARM processors are directly applied in our implementation. An access control scheme for IoT based on lattice was proposed by Zhang et al. [22]. However, their scheme requires performing $O(n^2)$ matrix operations which cannot be implemented on 8-bit microcontrollers.

2.2 Homomorphic Encryption in Cloud-IoT Architectures

Gentry [23] first constructed fully homomorphic encryption (FHE), followed by leveled FHE by BGV [24], while the integer-based BFV scheme was proposed by [25] and used in ALPHREN. In a paper [9], Li et al. apply BFV to medical data analytics in the cloud. They show that aggregate statistics do not require the patient's data. However, their implementation incurs a latency of 892 ms for every batch of 100 records. In other words, the server hardware they deploy is three orders of magnitude greater than the constraints upon IoT. ALPHREN is limited to homomorphic depth $d \leq 7$, which bounds noise accumulation with sub-millisecond overhead per operation. Boura et al. [26] tackle the complexity of BFV multiplication and obtain the $O(n \log n)$ bound achievable through NTT. TFHE is designed for faster Boolean circuit evaluation on a GPU. Even though it can evaluate gates faster than previous proposals, it cannot evaluate the arithmetic required to extract medical features. It is worth mentioning that previous work utilizes homomorphic encryption for cloud IoT processing [9] and medical data [28] separately, whereas ALPHREN is the first work to combine HE with proxy re-encryption on physical IoT hardware, creating end-to-end encrypted data pipelines from sensor to cloud to authorized receiver.

2.3 Proxy Re-encryption: Classical and Lattice-Based

The idea of proxy re-encryption was first put forward by Blaze et al. [10]; thereafter, it was rigorously formalized by Ateniese et al. [11], which set the foundation of IND-PRE-CCA2 security notion used by ALPHREN. The authors [29] introduce an extensive taxonomy of security notions and constructions for PREs; the IND-PRE-CCA2 property, satisfied by ALPHREN, provides the strongest security guarantee which can be achieved in practice. Green and Ateniese [30] construct the first identity-based PRE scheme to enable attribute-based delegation. ALPHREN follows a simpler design for direct PRE system to reduce the overhead in IoT. Lattice-based PRE is getting growing attention for post-quantum security. NTRURenCrypt proposed by Nuñez et al. [29] relies on NTRU and has competitive performance but allows only single-hop delegation. Shao and Cao [31] build a PRE that is CCA-secure and does not use pairings. Yao et al. [32] propose an attribute-based encryption scheme for IoT whose architecture has similarities with our delegation scheme. ALPHREN enhances this line of work by integrating lattice PRE with homomorphic computation, a synthesis not previously reported in the literature.

2.4 Parameter Optimization for Lattice Cryptography

It is not possible to find a closed form solution to the discrete multi-objective trade-off surface for selecting optimal Ring-LWE parameters [33,34]. The Lattice Estimator [33] of Albrecht et al. has been used by GHOA to empirically estimate the attack cost of the industry-standard BKZ-2.0, the dual attack and the primal (uSVP) attack against the given parameter sets. The GHOA algorithm presented in our work offers an automatic search within the complex dimensioning space proposed by Bernstein and Yang [34] using meta-heuristic tools. The authors of [35] employed a genetic algorithm to optimize the RSA key scheduling, whereas the authors of [36] used a particle swarm optimization to attack the optimal design of S-box. In GHOA, honey bee foraging dynamics are included which, based on empirical evidence, provide faster convergence for hypothetical discrete parameter grids due to scout-forager separation (47 vs. ~ 180 iterations). This allows for exploration while intensifying search on promising regions. To our knowledge, GHOA is the first meta-heuristic applied to the Ring-LWE parameter selection problem. SHAP (SHapley Additive exPlanations) [37], its basis on cooperative game theory, provides theoret...

2.5 Explainability in Cryptographic and Security Systems

SHAP (SHapley Additive exPlanations) [37], rooted in cooperative game theory, provides theoretically grounded feature attribution for machine learning models. Applications to security include intrusion detection explanation [38] and vulnerability prioritization [39]. However, application of SHAP to cryptographic parameter selection where the "model" is a multi-objective fitness function over a discrete parameter space has not been previously reported. ALPHREN's SHAP module fills this gap by treating GHOA's fitness function as the object of explanation, providing non-cryptographers with quantitative guidance about which parameters most strongly influence security-efficiency outcomes. This directly addresses HIPAA and GDPR requirements for algorithmic transparency in medical data processing systems.

3. The ALPHREN Framework: Architecture and Design

3.1 Security Goals and Trust Model

The four formal security properties governing the ALPHREN design are derived from deployment constraints and not arbitrary choices. At least IND-CPA security must be satisfied by local encryption on the IoT node as the cloud provider is able to observe the ciphertexts generated on constrained hardware. The overall end-to-end encryption mechanism is made IND-CCA2 secure with the Fujisaki-Okamoto transform [44] which transforms an IND-CPA-secure lattice scheme into a chosen-ciphertext-secure KEM by hashing the encapsulated key into the randomness a norm construction that adds negligible extra computation. The proxy re-encrypted ciphertexts individually satisfy IND-PRE-CCA2 security which is the strongest achievable notion for unidirectional single-hop re-encryption [11, 29]. This guarantees that even with the proxy relay colluding as a decryption helper, no adversary with access to a re-encryption oracle can distinguish challenge ciphertexts. All three properties are satisfied based on the Ring-LWE hardness assumption. They offer post-quantum security, which no provably post-quantum RSA- and ECC-based alternative can guarantee against polynomial-time quantum adversaries. The threat model is the standard semi-honest proxy assumption, where the cloud server and the proxy relay faithfully execute the appropriate protocol steps, while also trying to passively infer private information from the observed message transcripts. It is a reasonable assumption for commercial clouds. These are covered by contractual data processing agreements (as required under GDPR Article 28) and documented logging environments where activeness deviation from protocol has legal and reputational impact. The threshold re-encryption key distribution will be required to extend to fully malicious proxies in section 10; it is labelled as future work. The IoT node and authorized recipient are fully trusted in their private key domain. The formal adversary is a randomized polynomial-time algorithm endowed with quantum computational power that can request a decryption oracle for any ciphertext other than the challenge ciphertext (as per the IND-CCA2 security game).

3.2 System Architecture

ALPHREN is broken down into four subsystems which are partitioned across three trust domains and the security properties guaranteed by the architecture derive from the boundaries between the domains rather than of the subsystems themselves. In the IoT domain that is fully trusted, the GHOA-driven parameter selection (done once offline) is the only operation requiring access to the secret key. The remaining operations specifically requiring access to the secret key are Ring-LWE key pair generation, local BFV encryption of sensor data, derivation of re-encryption key from recipient's public key, and generation of SHAP audit records. These operations are recorded in the INCM audit logs. Since the INCM never releases key material, the trust boundary of the IoT domain is precisely the ciphertext output interface. All data crossing this boundary are encrypted under pk_A . In the semi-trusted cloud gateway area, the Proxy Re-encryption Relay (PRR) retains re-encryption keys assigned by data owners and performs re-encryption at request. The semi-honest status of PRR is enforced architecturally: It holds $rk_{A \rightarrow B}$ but never sk_A or any plain, so even a passively compromised relay learns nothing more than what can be inferred from the ciphertext distribution. Homomorphic Evaluation Engine occupies the same domain and processes operations on the arithmetic circuit via additions and multiplications with relinearization directly on the ciphertexts received from the PRR. Thus enabling necessary computations for analytics while preserving user's privacy with no decryption step. The Security Classification Module (SCM), residing at the cloud boundary with the recipient domain, maintains intrusion detection classifiers that evaluate encrypted feature vectors. When classification is kept in the encrypted domain, the sovereignty over the data is preserved, meaning that the cloud only learns the classification and not the sensor readings. Protocol Overview

3.3 Protocol Overview

The ALPHREN protocol consists of six consecutive phases, each one establishing the security property needed by the following one. In a manner that guarantees safe initialization of the parameters, secure key management, encrypted computation, controlled delegation of authority and decryption with authorization. The protocol starts by selecting parameters. Before executing any of the cryptographic operations, the Intelligent Network Control Module (INCM) executes the off-line Genetic Honey Optimization Algorithm (GHOA) to determine the optimal set of Ring-LWE parameters (n^*, q^*, σ^*) . The process of optimizing the parameters seeks to achieve 128 bits of post-quantum security while minimizing computational overhead. This process is carried out only once, during system initialization, and does not increase communication session latency at all. After initializing the parameters, protocol enters the key generation phase. Every communicating entity uses the optimized parameters to generate its own public and private key pair. Alice creates the key pair (sk_A, pk_A) in the IoT device, while Bob does the same with (sk_B, pk_B) inside a trusted environment. The secret keys remain local to their respective owners and are never sent nor shared during the protocol. The next step involves the generation of the proxy re-encryption key, which enables delegation without knowledge of the private key. Alice calculates the re-encryption key with her secret key and Bob's public key. $rk_{A \rightarrow B} = RkmCgen(sk_A, pk_B)$. This key is transferred to the Proxy Re-encryption Relay (PRR) through a mutual authenticated channel securely. Re-encryption key is unidirectional by design. In possession of. The presence of $rk_{A \rightarrow B}$ allows the construction of Alice to

Bob ciphertexts by the PRR, but the reverse key $rk_{B \rightarrow A}$ cannot be constructed, i.e., Bob's private key is not made available to the relay. After registering the re-encryption key, Alice encrypts the measured data with her public key. $CT_A = \text{Enc}(pk_A, m)$ where m is the sensor message. The cloud is only sent the resulting ciphertext, so no plaintext ever leaks off the IoT. Even if multiple ciphertexts are gathered or stored, the cloud or relay cannot learn about the underlying data. The encrypted data is then computed during the homomorphic evaluation phase. The Homomorphic Evaluation Engine (HEE) directly performs the required arithmetic over ciphertexts without a required decryption. The assessment generates. $CT_{\text{Result}} = \text{Eval}(f, CT_A)$, which shows the encoded result of $f(m)$. During this computation, the encryption boundary is not broken so that intermediate or final plaintext values are disclosed. The next step is pattern encryption. With the stored re-encryption key, the evaluated ciphertext is transformed by the PRR. $CT_B = \text{ReEnc}(rk_{A \rightarrow B}, CT_{\text{Result}})$. As a result, it will turn a ciphertext that could have originally been decrypted only by Alice into one that can be decrypted only by Bob. Significantly, the transformation alters solely the cryptographic structure of the ciphertext. The relay can perform this operation without being able to learn the message itself or knowing either party's secret key. Ultimately, Bob performs. $m' = \text{Dec}(Sk_B, CT_B)$ to retrieve the output event. When a protocol executes correctly, the output is decrypted will satisfy. $m' = f(m)$




Fig. 2. ALPHREN cryptographic protocol sequence diagram (6 words) The entire six-phase protocol is displayed across five principals: IoT Node (Alice, blue), GHOA Optimizer (purple), Proxy Relay (green), HEE~ / Cloud (red), and Recipient Bob (teal). Arrows show confirmed data flows safety Guarantee IND-CCA2 is maintained throughout, proxy relay does not hold nor access plaintext.

4. Mathematical Modeling

4.1 Ring-LWE Foundations

Let n a power of two and q a prime satisfying the congruence $q \equiv 1 \pmod{2n}$. Instead of a random selection of parameters, this particular congruence ensures there is a primitive $2n$ -th root of unity in $\mathbb{Z}_q$.

The usability of Number Theoretic Transform (NTT) for polynomial multiplications in $(n \log n)$ arithmetic operations is a structural property of the presented framework's architecture. This property acts as the performance engine of the framework. The underlying algebraic platform is defined over the cyclotomic quotient ring $R = \mathbb{Z}[x]/(x^n + 1)$, with its modular counterpart denoted as $R_q = \mathbb{Z}_q[x]/(x^n + 1)$. Elements within R_q are polynomials of degree at most $n - 1$, with coefficients strictly reduced modulo q .

The Ring Learning With Errors (Ring-LWE) distribution $\mathcal{D}_s$, over R^2 is parameterized by a secret polynomial $s \in R$ and an error distribution $\chi = \mathcal{D}_{R, \sigma}$, which represents a discrete Gaussian q distribution over R with a standard deviation of σ . A sample drawn from this distribution takes the

form $(a, b) \in R^2$, where $a \leftarrow R$ is sampled uniformly at random, and the target component is defined

as:

$$b = a \cdot s + e \pmod{q} \text{-----} (1)$$

where $e \leftarrow \chi$ represents a freshly drawn error.

The security of our construction relies on the Ring-LWE hardness assumption, which asserts that samples from $\mathcal{D}_s$ are computationally indistinguishable from a uniform distribution over R^2 , even when faced with a quantum polynomial-time adversary. This assumption benefits from a rigorous worst-case reduction to the Shortest Independent Vectors Problem (SIVP_q) on ideal lattices, as demonstrated by Lyubashevsky, Peikert, and Regev [17]. Consequently, an adversary capable of breaking the Ring-LWE instance could be leveraged to solve SIVP_q instances across all ideal lattices of a corresponding dimension, rather than merely an average-case subset. This worst-case guarantee marks a major security divergence from classical RSA or elliptic-curve assumptions, which depend entirely on average-case hardness conjectures that fail in the presence of quantum computing architectures.

4.2 Key Generation

The KeyGen($1^\lambda, n, q, \sigma$) routine establishes the long-term cryptographic identity of an actor through the following sequence:

1. Sample the secret polynomial $s \leftarrow \mathcal{D}_R$, to serve as the secret key.
2. Sample a public base element $a \leftarrow R_q$ uniformly at random.
3. Draw an initial key-generation error $e_k \leftarrow \mathcal{D}_R$. The key pair (pk,sk) is then constructed as:

$$\text{pk} = (a, b) = (a, -a \cdot s + e_k \pmod{q}) \text{-----} (2)$$

$$\text{sk} = s \text{-----} (3)$$

The polynomial multiplication $a \cdot s$ is evaluated via the NTT to maintain an $(n \log n)$ asymptotic profile. Setting the operational parameters to $n = 1024$ and $q = 12289$, the NTT runs over a 14-bit prime field. Because $q - 1 = 3 \cdot 2^{12}$ is perfectly divisible by $2n = 2048$, the existence of a primitive $2n$ -th root of unity is mathematically guaranteed. Using the primitive root $g = 11 \pmod{12289}$, the specific twiddle factor is determined by:

$$\omega = g^{\frac{q-1}{2n}} \pmod{q} = 11^6 \pmod{12289} = 1479 \text{-----} (4)$$

This parameter configuration yields highly optimized memory footprints. The secret key sk requires exactly:

$$\frac{n \cdot \lceil \log_2 q \rceil}{8} = \frac{1024 \cdot 14}{8} = 1792 \text{ bytes} \text{-----} (5)$$

This packs down to ≈ 1.8 KB (or 2 KB under standard byte-aligned structures). The public key pk spans two such polynomials, totaling 4 KB, which yields an aggregate key material size of 6 KB.

When evaluated alongside classical baselines, these properties remain highly competitive. For instance, while an RSA-2048 public key uses a small public exponent (≈ 3 bytes), its modulus alone demands 256 bytes, and its private key representations routinely exceed 2.5 KB in standard PKCS#8 formatting. The transition to post-quantum security via Ring-LWE therefore carries a minimal and practical storage penalty.

4.3 BFV Encryption and Decryption

To encrypt a message under the Brakerski-Fan-Vercauteren (BFV) paradigm, the message is first mapped to a plaintext polynomial $m_p \in R_q$ by allocating one message bit per coefficient.

The encryption routine $\text{Enc}(\text{pk}, m_p)$ samples transient error polynomials $r, e_1, e_2 \leftarrow \mathcal{D}_{R, \sigma}$ independently, outputting the ciphertext tuple $\text{CT} = (c_0, c_1) \in R^2$ defined by:

$$c_0 = b \cdot r + e_1 + \left\lceil \frac{q}{2} \right\rceil \cdot m_p \pmod{q} \quad (6)$$

$$c_1 = a \cdot r + e_2 \pmod{q} \quad (7)$$

Decryption is performed by evaluating the ciphertext elements against the secret key $\text{sk} = s$:

$$m^* = c_0 + c_1 \cdot s \pmod{q} \quad (8)$$

To trace the structural correctness of this step, we substitute equations (3) and (4) into equation (5), then expand using the definition of b from equation (1):

$$\begin{aligned} m &= (b \cdot r + e_1 + \left\lceil \frac{q}{2} \right\rceil m_p) + (a \cdot r + e_2) \cdot s \\ &= (-a \cdot s + e_k) \cdot r + e_1 + \left\lceil \frac{q}{2} \right\rceil m_p + a \cdot r \cdot s + e_2 \cdot s \end{aligned} \quad (9)$$

$$= \left\lceil \frac{q}{2} \right\rceil m_p + (e_1 + e_2 \cdot s + \overline{e}_k \cdot r) \pmod{q} \quad (10)$$

The aggregate noise term, denoted as $\eta = e_1 + e_2 \cdot s + e_k \cdot r$, explicitly preserves the distinction between the fresh encryption errors (e_1, e_2) and the static key-generation error (e_k) . Successful message extraction via decoding defined by $m = \lfloor m_p \cdot \frac{2}{q} \rfloor$ depends entirely on preserving the

$$\text{infinity-norm bound } \|\eta\|_\infty < \frac{q}{4} \quad m = \lfloor m_p \cdot \frac{2}{q} \rfloor \pmod{2}$$

For our chosen concrete instances where $\sigma = 3.19$ and $n = 1024$, the probability that the noise violates this threshold can be modeled via a standard tail bound argument over a convolved discrete Gaussian. Applying a union bound across all n coefficients, the failure probability per coefficient is bounded by:

$$\Pr[|\eta_i| \geq \frac{q}{4}] \leq 2 \exp\left(-\frac{(q/4)^2}{2\sigma_{\text{eff}}^2}\right) \quad (11)$$

where $\sigma_{\text{eff}} \leq 3\sigma\sqrt{n} \approx 305$ accounts for the expanded variance of the polynomial combinations. Given $q = 12289$, this yields a per-coefficient failure rate below 2^{-67} . Aggregating this across all 1024 independent coefficients ensures that the total decryption failure probability remains safely below $1 - 2^{-57}$, matching the security baselines established in standard literature [17].

4.4 Homomorphic Operations

The BFV engine supports fully homomorphic evaluations over encrypted payloads. Homomorphic addition operates directly on two ciphertexts $\text{CT}_1 = (c_{01}, c_{11})$ and $\text{CT}_2 = (c_{02}, c_{12})$:

$$\text{CT}_{\text{add}} = (c_{01} + c_{02}, c_{11} + c_{12}) \pmod{q} \quad (12)$$

Because the underlying noise expands strictly linearly ($\|\eta_{\text{add}}\|_\infty \leq \|\eta_1\|_\infty + \|\eta_2\|_\infty$), additive operations do not pose a meaningful threat to correctness within reasonable circuit depths.

Homomorphic multiplication, by contrast, requires structural management. Multiplying CT_1 and CT_2 initially yields an expanded, degree-2 ciphertext $\text{CT}_{\text{mul}} = (d_0, d_1, d_2)$ where:

$$d_0 = c_{01} \cdot c_{02}, \quad d_1 = c_{01} \cdot c_{12} + c_{11} \cdot c_{02}, \quad d_2 = c_{11} \cdot c_{12} \pmod{q} \text{----- (13)}$$

To prevent geometric expansion of the ciphertext size during subsequent operations, a relinearization step compresses the tuple back to a standard degree-1 format using a specialized evaluation key $ek = (ek_0, ek_1)$:

$$CT' = (d_0 + ek_0 \cdot d_2, d_1 + ek_1 \cdot d_2) \pmod{q} \text{----- (14)}$$

This operation scales efficiently within $(n \log n)$ time using the NTT.


Each multiplicative layer scales the noise quadratically. Following L consecutive multiplication layers, the maximum noise magnitude is governed by the structural bound:

$$B_L \leq n^L \cdot B^2 \cdot C^{2^L - 1} \text{----- (15)}$$

where B_0 represents the baseline noise from a fresh encryption, and C represents the fixed relinearization noise constant dictated by the parameter choices:

$$C = \frac{8n\sigma^2 \cdot \|ek\|_\infty}{q} \text{----- (16)}$$

Evaluating equation (10) using our real-world parameters ($n = 1024, q = 12289, \sigma = 3.19$) determines the maximum allowable circuit depth. At an evaluation depth of $L = 7$, the noise bound yields $B_7 \approx 1.9 \times 10^3$, sitting safely under the hard ceiling of $q = 3072$. Stepping to $L = 8$ pushes the noise bound to $B_8 \approx 4.1 \times 10^3 > q$, triggering an immediate decryption failure. Thus, our depth-7 restriction is derived directly from these algebraic thresholds.



As a consequence of noise accumulation across these seven sequential levels, the probability of correct decryption shifts from the single-encryption baseline $(1 - 2^{-57})$ down to a circuit-wide assurance of $1 - 2^{-48}$. This remains a statistically sound profile for runtime operations.

4.5 Proxy Re-encryption

The delegation component allows an intermediate proxy to translate ciphertexts from Alice to Bob without learning the underlying data.

The re-encryption key generation routine $RKeyGen(sk_A, pk_B)$ constructs the translation token $rk_{A \rightarrow B}$ as follows:

$$rk_{A \rightarrow B} = s_A + (pk_B) \cdot b_B \pmod{q} \text{----- (17)}$$

where b_B represents the second polynomial element of Bob's public key pk_B

R_q denotes a collision-resistant hash function modeled as a random oracle.

$$= (a_B, b_B), \text{ and } H: R^2 \rightarrow q$$

Given Alice's ciphertext $CT_A = (c_0, c_1)$, the proxy applies this token via the ReEnc routine to generate the modified ciphertext CT_B :

$$CT_B = (c_0 - c_1 \cdot rk_{A \rightarrow B}, c_1 \cdot a_B) \pmod{q} \text{ ----- (18)}$$

Correctness Derivation

Decryption of the transformed ciphertext CT_B under Bob's secret key s_B expands linearly: $Dec(s_B, CT_B) = (c_0 - c_1 \cdot rk_{A \rightarrow B}) + (c_1 \cdot a_B) \cdot s_B \pmod{q}$ ----- (19)

Substituting the explicit definition of the re-encryption token $rk_{A \rightarrow B}$ from equation (11):

$$\begin{aligned} Dec(s_B, CT_B) &= c_0 - c_1 \cdot (s_A + H(pk_B) \cdot b_B) + c_1 \cdot a_B \cdot s_B \\ &= c_0 - c_1 \cdot s_A - c_1 \cdot (pk_B) \cdot b_B + c_1 \cdot a_B \cdot s_B \end{aligned} \text{ ----(21)}$$

Next, expanding Bob's public key term $b_B = -a_B \cdot s_B + e_{k,B}$, resolves the identity:

$$\begin{aligned} Dec(s_B, CT_B) &= c_0 - c_1 \cdot s_A - c_1 \cdot H(pk_B) \cdot (-a_B \cdot s_B + e_{k,B}) + c_1 \cdot a_B \cdot s_B \\ &= (c_0 - c_1 \cdot s_A) + c_1 \cdot a_B \cdot s_B \cdot (1 + (pk_B)) - c_1 \cdot (pk_B) \cdot e_{k,B} \pmod{q} \end{aligned}$$

Because the baseline components $(c_0 - c_1 \cdot s_A)$ preserve the primary plaintext extraction relation

$$\text{----- (20)}$$

q
 $[-]m_p + \eta_A$, the token successfully shifts the decryption capability to Bob's key ring. The residual error components scale within permissible boundaries, ensuring successful decoding with the same bounds established in Section 4.3.

Security Analysis

The proxy re-encryption mechanism enforces strict unidirectionality. The translation token $rk_{A \rightarrow B}$ exposes no mathematical pathway to derive the inverse relation $rk_{B \rightarrow A}$. Attempting to extract the inverse parameter requires isolating Bob's secret key s_B from the public parameter $b_B = -a_B \cdot s_B +$

$e_{k,B}$, an act that reduces directly to solving an instance of the Ring-LWE problem.

Formally, the architecture satisfies IND-PRE-CCA2 security via a sequence of game-based reductions built on the core IND-CCA2 security of the underlying BFV infrastructure and the collision-resistant mapping of H under the Random Oracle Model, adhering to the frameworks established by Ateniese et al. [11] and Nuñez et al. [29]. Any adversary demonstrating a non-negligible advantage in distinguishing re-encrypted ciphertexts under a chosen-ciphertext attack can be directly transformed into an oracle wrapper that breaks the underlying Ring-LWE hardness assumption.

5. Genetic Honey Optimization Algorithm (GHOA) for Ring-LWE Parameter Optimization

5.1 Problem Formulation

The selection of cryptographic parameters for the Ring Learning With Errors (Ring-LWE) scheme can be formulated as a constrained multi-objective optimization problem. The objective is to identify a parameter combination that simultaneously maximizes the security level while minimizing computational latency and memory consumption. Since these objectives exhibit conflicting dependencies on the polynomial dimension, modulus, and error distribution, deriving an analytical optimum is generally intractable. Consequently, a metaheuristic optimization framework is adopted.

The feasible search space is defined as

$$\Omega = \{(n, q, \sigma) \mid n \in \mathcal{N}, q \in \mathcal{P}, \sigma \in \mathcal{S}\}, \text{----- (22)}$$

where

$$\mathcal{N} = \{512, 1024, 2048\}, \mathcal{P} = \{7681, 12289, 40961\}, \text{----- (23)}$$

and

$$\mathcal{S} = \{2.0, 2.5, 3.19, 3.5, 3.9\} \text{ ----- (24)}$$

The selected prime moduli satisfy

$$q \equiv 1 \pmod{2n},$$

which guarantees compatibility with the Number Theoretic Transform (NTT), while the error standard deviations are selected from commonly adopted security ranges for Ring-LWE implementations.

A candidate solution

$$x = (n, q, \sigma) \in \Omega \text{ ----- (25)}$$

is considered feasible only if the following constraints are simultaneously satisfied:

$$C_1: (x) \geq 128 \text{ bits}, C_2: t(x) \leq 1.0 \text{ ms}, C_3: M(x) \leq 64 \text{ KB}, \text{----- (26)}$$

where

- $S(x)$ denotes the estimated quantum security level,
- $t_{enc}(x)$ represents the encryption latency,
- (x) denotes the peak memory requirement. The feasible region is therefore

$$\Omega_f = \{x \in \Omega: C_1 \wedge C_2 \wedge C_3\} \text{ ----- (27)}$$

If any constraint is violated,

$$(x) = 0,$$

thereby excluding infeasible candidates from further consideration.

For feasible solutions, the optimization objective is defined as

$$\max(x), \text{----- (28)}$$

$$x \in \Omega_f$$

The normalization ensures that all objectives contribute on a comparable scale, whereas the weighting coefficients prioritize security while still accounting for computational efficiency and memory utilization.

5.2 Genetic Honey Optimization Algorithm

The proposed Genetic Honey Optimization Algorithm (GHOA) combines the exploitation capability of genetic evolution with the exploration characteristics of honey bee foraging. The hybrid strategy balances intensification around promising solutions and diversification across the discrete search space.

Let

$$(t) = \{(t), (t), \dots, x^{(t)}\} \text{----- (29)}$$

$$1 \qquad \qquad \qquad 2 \qquad \qquad \qquad N$$

denote the population at generation t , where

$$N = 50.$$

Each individual corresponds to one feasible parameter vector

$$(t) = (n_i, q_i, \sigma_i) \text{----- (30)}$$

$$i \qquad i \qquad i \qquad i$$

The optimization process consists of four sequential operators.

Step 1. Fitness Evaluation

The fitness of every candidate is computed as

$$(t)$$

$$f_i = (x_i) \text{----- (31)}$$

Candidates violating any feasibility constraint are immediately assigned

$$f_i = 0.$$

Otherwise,

$$S(x_i)$$

is obtained from the lattice security estimator and substituted into (9).

Step 2. Elite Preservation

The population is sorted according to descending fitness,

$$(1) \geq (2) \geq \dots \geq (N) \text{----- (32)}$$

The elite subset is defined as

$$(t) = \{(1), \dots, x_{(k)}\}, \text{----- (33)}$$

where

$$k = \lfloor 0.2N \rfloor.$$

Elite solutions are copied directly into the next generation,

$$(t+1) \supseteq E(t),$$

thereby preserving the highest-quality individuals.

Step 3. Genetic Recombination

Two elite parents

and

$$x_\alpha = (n_\alpha, q_\alpha, \sigma_\alpha) \text{-----} (34)$$

$$x_\beta = (n_\beta, q_\beta, \sigma_\beta) \text{-----} (35)$$

are randomly selected.

An offspring is generated through component-wise inheritance,

$$x_c = (n_c, q_c, \sigma_c), \text{-----} (36)$$

where

$$n_c \in \{n_\alpha, n_\beta\}, q_c \in \{q_\alpha, q_\beta\}, \sigma_c \in \{\sigma_\alpha, \sigma_\beta\}.$$

This discrete recombination allows promising characteristics from different parents to be inherited simultaneously.

Step 4. Honey Bee Mutation

Each offspring undergoes mutation with probability

$$p_{mut} = 0.15 \text{-----} (37)$$

If mutation occurs, one decision variable is randomly selected,

$$v \in \{n, q, \sigma\}, \text{-----} (38)$$

and updated according to

$$v' = v + \Delta_v, \text{-----} (39)$$

where

$$\Delta_v \in \{-1, +1\} \text{-----} (40)$$

represents movement to an adjacent discrete level.

Any infeasible mutation is rejected and resampled until a feasible neighboring state is obtained. This scout-bee exploration mechanism prevents premature convergence and improves population diversity.

The iterative optimization terminates when

$$\max_i |f^{(i)} - f^{(i-1)}| < 10^{-6} \text{-----} (41)$$

for five consecutive generations or when the maximum iteration limit is reached. The optimal Ring-LWE parameter set is therefore

$$x^* = (n^*, q^*, \sigma^*) = \operatorname{argmax}_{x \in \Omega_f} (x). \text{-----} (42)$$

$x \in \Omega_f$




Fig. 3. GHOA Key Optimization Algorithm Flowchart. The algorithm initializes with 50 candidates from the Ring-LWE parameter grid, evaluates each by the multi-objective fitness function f (security 45%, latency 35%, memory 20%), applies elite selection, genetic crossover, and honey bee mutation, validates lattice security via BKZ-2.0 Lattice Estimator, and checks convergence. The No (not converged) branch loops back for further optimization. Optimal output $n^*=1024$, $q^*=12289$, $\sigma^*=3.19$ is achieved in ~47 iterations (128-bit quantum security, 0.43 ms key generation, 18.4 KB RAM).

5.3 Convergence Analysis

Proposition 1

The proposed GHOA converges to the global optimum over the finite search space Ω with probability one.

Proof

The search space

$$|\Omega| = |\mathcal{N}| \cdot |\mathcal{P}| \cdot |\mathcal{S}| = 3 \times 3 \times 5 = 45 \text{-----} (43)$$

is finite.

Since the mutation probability satisfies

$$p_{mut} > 0,$$

every admissible state has a non-zero probability of being generated after a finite number of transitions. Consequently, the sequence of populations forms an irreducible and aperiodic Markov process over a finite state space.

Furthermore, elite preservation guarantees

$$\max_i^{(t+1)} f_i \geq \max_i^{(t)} f_i \text{-----} (44)$$

indicating that the best fitness value is monotonically non-decreasing throughout the optimization.

Because every feasible solution remains reachable and the best solution cannot be discarded once discovered, the algorithm converges to the global maximizer

$$x^* = \operatorname{argmax}(x) \text{-----} (45)$$

$$x \in \Omega_f$$

6. Experimental Methodology

6.1 Datasets and Pre-processing

Four datasets that are publicly available were chosen that could cover the operation targets of ALPHREN in a complimentary manner. The operation targets are medical data confidentiality, IoT intrusion robustness, and cross-domain generalization. Dataset selection aimed for real-world diversity in record volume and feature dimensionality, to actively stress-test the encryption pipeline in clinically relevant data regimes. The MIMIC-III Clinical Database [40] provides 48,321 de-identified ICU patient records. Each of these records is characterized by 24 physiological features which have been extracted from one-hour observation windows. The physiological features include heart rate, peripheral oxygen saturation (SpO_2), invasive arterial blood pressure, respiratory rate, core temperature, and nine laboratory analytes. In full accordance with the MIMIC data use agreement, all identifiers were removed prior to experimentation. Each feature was zero-centered and had the unit variance imposed. Missing value (overall 3.7%) was forward-filled within the patients' episode. This dataset was the primary benchmark for the throughput of end-to-end encryption, modelling the volume and structure of real clinical IoT gateway traffic. For the evaluation at the network layer, we used the UCI IoT Intrusion Detection Dataset comprising of 625,783 labelled traffic records which we collected from an 11-device smart home testbed. The records include five different types of attacks namely Mirai botnet flooding, ARP spoofing, port scanning, man-in-the-middle injection, and a benign baseline. One-hot encoding was applied to the categorical protocol features. To prevent scale dominance in the classification model, the numerical flow statistics were min-max normalized to $[0, 1]$. To assess the throughput scaling behavior of ALPHREN free from any dataset-specific distributional artifacts, we simulated a synthetic medical IoT stream of 100,000 records using a six-component Gaussian mixture model (GMM). A 5,000-record sample from MIMIC-III was used to assess model fit via minimization of Bayesian information criterion. These results indicate that six components are the parsimonious fit; i.e., BIC reduced monotonically as the number of components went from one to six, and plateaued thereafter. The expectation-maximization method was used to estimate the GMM parameters and the generation fidelity was verified by a two-sample Kolmogorov-Smirnov test on MIMIC-III (all features, $p > 0.12$). The NSL-KDD benchmark [41], with its 125,973 records, is used for the cross-domain validation of attack detection results. This benchmark is partitioned into the standard KDD Train+ and KDD Test+ subsets covering the four canonical attack families (DoS, Probe, R2L, U2R) and benign class. For all the four datasets, training, validation, and holdout test subsets were stratified split by class label in the ratio of 80/10/10. The stratification maintained the same class proportions across the partitions. This process was necessitated for the NSL-KDD dataset. After all, the U2R samples are less than 0.5% of the total records.

6.2 Hardware Platform and Implementation

All cryptographic operations on the IoT side were performed exclusively on a Raspberry Pi 4B running 64-bit Raspberry Pi OS (kernel 5.15.84). The Pi runs on a Broadcom BCM2711, ARM Cortex-A72 quad-core at 1.8 GHz, 4 GB LPDDR4-3200 SDRAM with 15.3 GB/s peak bandwidth. Prior to all measurements, cpufreq-set utility locked CPU frequency scaling to 1.8 GHz in order to avoid thermal throttling from causing non-stationarity issue in timing data. Before benchmarking, a 15-minute warm-up time (with nominal load) was given and for each of the trials, the first five iterations were discarded also to account for cold-cache effect. We offloaded cloud-side homomorphic evaluation to an AWS EC2 c5.large instance (Intel Xeon Platinum 8124M, 2 vCPU, 4 GB RAM). This type of instance offers a deterministic single-tenant compute environment which is necessary for reproducible latency measurement. This consideration excludes burstable instance

families. All cryptographic primitives were implemented in C++17 to Python. The 3.9 bindings are available using pybind11 v2.10.3. The Number Theoretic Transform, developed using ARM NEON SIMD intrinsics following the optimized assembly techniques of Khalid et al. [21], records a 1.8× improvement in throughput as compared to a scalar reference implementation running on the same hardware. Homomorphic operations deployed by BFV library Microsoft SEAL v4.1 were designed with IoT-centric parameter profiles disabling batching slot redundancy to reduce RAM footprint. The Lattice Estimator v0.4.0 was invoked as a Python subprocess for security margin computation. Therefore, GHOA was implemented in C++. The competing scheme implementations maintained their original language and library dependencies wherever possible; that is, where source code was available from the authors' repositories. Specifically, Padma LHE [13] in Python 3.9, Hussien Hash [15] in Python 3.9, and Mahor ECC [16] via OpenSSL 3.0 so as to preserve the conditions under which their published benchmarks were achieved. To the extent that differences between languages could inflate ALPHREN's measured advantage, results are interpreted relative to the post-quantum-capable schemes Padma LHE and HXP, not to the Python implementations for latency comparisons; throughput rankings are stated with this caveat noted. This was done on the same hardware configuration as a hardware-native symmetric reference using. The recorded values are medians over 30 independent trials; 95% confidence intervals are bias-corrected bootstrapped. Given the observed right-skew in the latency distributions, we use 10,000 resamples rather than Gaussian approximation.

6.3 Competing Schemes and Evaluation Baseline

Seven schemes were evaluated under a common benchmarking harness: Padma LHE [13], HXP [14], Hussien Hash [15], Mahor ECC [16], AES-GCM-256 (symmetric reference), Ring-LWE Vanilla (Ring-LWE without GHOA optimization or proxy re-encryption, author-implemented as an ablation anchor), and ALPHREN. For each competing scheme, the parameter set published in the original paper was used without modification; where original parameters were under-specified, the authors' public repositories or correspondence with the original authors resolved ambiguities. Ring-LWE Vanilla was included not as an independent prior-work competitor but as a controlled baseline isolating the marginal contribution of GHOA and the proxy re-encryption layer its results appear in the ablation analysis (Section 9) rather than the primary comparison table. All seven implementations were deployed on the same Raspberry Pi 4B hardware running identically configured operating environments to eliminate platform-induced confounds.

6.4 Statistical Analysis Protocol

The statistical inference strategy follows a two-stage non-parametric protocol, chosen because Shapiro-Wilk normality testing on 30-trial timing distributions rejected normality for five of seven schemes at $\alpha = 0.05$ (W statistics: 0.81–0.93), precluding parametric ANOVA-based inference. Stage 1 applied the Friedman rank-sum test across all seven schemes on each primary metric (key generation time, encryption latency, throughput, memory consumption, energy consumption, security accuracy) as an omnibus test of the null hypothesis that performance is exchangeable across schemes. Stage 2 applied pairwise Wilcoxon signed-rank tests comparing ALPHREN against each of the six external schemes on all primary metrics. Family-wise error rate was controlled by Bonferroni correction at $\alpha = 0.05/6 = 0.0083$ per comparison pair. Effect sizes were quantified as Cohen's $d = (\mu_A - \mu_B)/\sigma_{\text{pooled}}$ [42], where μ and σ denote sample means and pooled standard deviations of the 30-trial distributions. Following established benchmarks [42], $d \geq 0.8$ was classified as a medium effect and $d \geq 1.3$ as a large effect, thresholds selected because they correspond to practically meaningful performance differences in cryptographic engineering contexts not merely statistical ones. A post-hoc sensitivity analysis confirmed that 30 trials provide greater than 90% power to detect the smallest observed effect size ($d = 1.39$) at the Bonferroni-corrected α , satisfying the minimum power requirement for inferential credibility.

7. Results and Discussion

7.1 Key Generation and Encryption Performance

Median performance across all six schemes on Raspberry Pi 4B after 30 independent trials (as per Table 1). ALPHREN enables key generation within 0.43 ms (95% CI: 0.41–0.45 ms); a factor-2.7 faster than Padma LHE (1.17 ms). Also, ALPHREN is faster than ECC-based Mahor (0.89 ms) which lacks post-quantum security. The algorithm implementation of ALPHREN introduces GHOA parameter optimization as well as proxy re-encryption key generation on top of Ring-LWE. This metric reflects the efficiency dividend from the GHOA NTT-compatible parameter selection ($n^*=1024$, $q^*=12289$). Encryption lag deserves dedicated focus. ALPHREN at 0.47 ms (95% CI: 0.45–0.49 ms) is 32.8× faster than Padma LHE (15.40 ms). Padma LHE is based on Python and incurs a heavyweight overhead compared to our ARM NEON NTT [21]. The sub-millisecond encryption aligns easily with clinical IoT sensors' sampling intervals of 1 ms–1 s. Hence, the incorporation of ALPHREN does not lead to detectable pipeline latency in medical monitoring. HXP(8.92 ms) and Mahor ECC(3.74 ms) both surpass the threshold; Hussien Hash(2.10 ms) approaches it with significantly weaker security guarantees. The 2.14 MB/s speed (95% CI: 2.09–2.19 MB/s) used to encrypt roughly 1.2MB of physiological time-series data from an hour's MIMIC-III ICU record does so in roughly 560 ms, well within the 30-second batch upload tolerance of hospital IoT gateways.

The performance of Padma LHE and HXP saturates at 0.39 MB/s, which has a 5.5× throughput gap attributed to the lack of NTT acceleration for their polynomials. AES-GCM-256 obtains 9.81 MB/s, as one would expect for a symmetric cipher. However, it also does not offer post-quantum security or the capability for proxy re-encryption. Thus, it is more of a computational reference than a direct competitor for the multi-party delegation use case. Throughput must be discussed as well as memory usage of 18.4 KB. Padma LHE needed. The Mahor ECC and Gallo 256-bit AES require more than 32KB of SRAM that mid-range Cortex M microcontrollers have. While larger in footprint than that of Hussien Hash (11.3 KB) and AES-GCM-256 (6.8 KB), ALPHREN fits within the hard 64 KB constraint specified in the GHOA optimization (Constraint C3, Section 5.1) and working set in the LPDDR4 cache of Raspberry Pi 4B. The post-quantum-capable schemes have the lowest energy consumption of 4.7 μ J per operation. The measurements below Padma LHE of 2.6 times (12.3 μ J), is suggestive of the latency advantage: fewer clock cycles per operation imply a direct reduction of dynamic power draw under the ARM Cortex-A72's power profile. Table 1 provides a thorough performance comparison based on different schemes in Raspberry Pi 4B and sorted by decreasing thruput. Each scheme was assessed medially over a span of 30 trials with a 95% confidence interval. The presence of * denotes statistical significance when compared with ALPHREN at a Bonferroni-corrected $\alpha = 0.0083$. Throughput and security accuracy measure a better performance while all other metrics signify a better performance when lower.

Scheme	Key Generation (ms)	Encryption Time (ms)	Throughput (MB/s)	Memory (KB)	Energy (μ J)	Security Accuracy (%)
ALPHREN (Proposed)	0.43	0.47	2.14	18.4	4.7	99.1
Padma LHE [13]	1.17	15.40	0.39	34.2	12.3	97.8
HXP [14]	0.78	8.92	0.39	22.1	9.8	96.1
Hussien Hash [15]	0.31	2.10	1.47	11.3	5.9	91.4
Mahor ECC [16]	0.89	3.74	0.83	28.7	11.2	96.4
AES-GCM-256 (Reference)	0.12	0.31	9.81	6.8	1.8	94.3

ALPHREN — Comprehensive Performance Comparison Against State-of-the-Art




Fig. 4. Comprehensive Performance Comparison Results. (a) Key generation (blue, left axis) and encryption time (red, right axis) ALPHREN achieves 0.43 ms and 0.47 ms respectively, optimal among post-quantum schemes. (b) Throughput 2.14 MB/s, 5.5× above Padma LHE and HXP. (c) Security vs. memory scatter ALPHREN achieves the highest security (99.1%) among post-quantum schemes with only moderate memory overhead; the dotted red lines highlight the Pareto frontier position. (d) Energy consumption 4.7 μJ/op lowest among PQ schemes; memory overlay (blue line) confirms the favorable balance.

7.2 Security Classification Results

The UCI IoT Intrusion dataset testing shows ALPHREN achieves security classification accuracy of 99.1% under end-to-end encryption. In terms of per-category results, we see from where the remaining 0.9% comes. The replay attack detection inhibits any replay events within a 5 ms timestamp window, which ended with a 98.7% probability, and for events after that, it was 100%. The impersonation detection was at 100%, the chosen-plaintext attack resistance was at 99.3% probability, the chosen-ciphertext attack resistance was at 99.1% and finally the differential power analysis resistance was at 97.4% as measured through ELMO side-channel simulation [43]. A value of replay detection smaller than 100% is seen not due to a cryptographic reason but rather due to the hardware of the Raspberry Pi 4B. The system clock of the Raspberry Pi has a granularity of 1 ms which cannot properly discern closely spaced replay events. The floor the analysis would reveal going forward would, on the whole, be tied to the upper-bound on the ALPHREN measurement depth. There has to be an answer to why ALPHREN (99.1%) has a 7.7 percentage point greater gap from Hussien Hash [15] (91.4%). Hussien Hash applies an HMAC on the data, providing protection but not anonymity. As a result, the cloud-side classifier can still exploit the plaintext statistical structure in the traffic to identify certain anomalies that remain undetected under IND-CCA2-secure encryption. The difference thus measures a genuine security property difference, not merely a classification difficulty difference: IND-CCA2 encryption actively conceals traffic patterns that weaker authenticators do not conceal. GHOA has measured superiority over the submissions Padma LHE and HXP due to suboptimal parameter and due to the absence of proxy re-encryption which introduces an additional layer of traffic distinguishability in the delegation operations of configurations we evaluated here. The interpretation of our classification results and the formal security proofs we provide in Section 4 goes as follows: First, since empirical accuracy demonstrates the correctness of an implementation, and since our IND-CCA2 reduction (Section 4.3) is under the hardness of Ring-LWE, our IND-CCA2 security then guarantees that no polynomial time quantum algorithm can distinguish ciphertexts.

7.3 Homomorphic and Re-encryption Operation Performance

The proxy re-encryption took 1.87 ms per 32-byte ciphertext block (95% CI: 1.81–1.93 ms), and re-encryption key size is 2.1 KB. To provide context, the closest lattice-based PRE (predicate-enciphered message scheme) to ALPHREN in the literature, NTRURenCrypt [29], achieves a re-encryption overhead of approximately 3.4 ms on similar ARM hardware. Thus, we see that ALPHREN achieves around 1.8× improvement due to a more compact polynomial ring structure of Ring-LWE relative to NTRU's non-ideal lattice. The re-encryption key size of 2.1K bytes fits within a single RAM cache line cluster on a PRR gateway capable of inline caching and not causing eviction overhead during multi-session delegation. Homomorphic addition took 0.09 ms while multiplication needed relinearization. The depth-7 arithmetic circuit over a 24-feature medical record, which assumes a balanced binary multiplication tree (7 HomMul layers) with 6 HomAdd operations per layer (24 parallel input branches processed), finishes in about $0.82 \times 7 + 0.09 \times 42 \approx 9.5$ ms, which is very much tolerable (30 sec / batch) for standard medical predictor pipelines. By imposing a depth-7 restriction, the rate at which noise accumulates is bounded below $q/4$ with probability $1 - 2^{-48}$ (see Section 4.4). Specifically, the circuit vocabulary mean, variance, Fourier coefficient approximations, and threshold comparisons cover the standard physiological feature extraction requirements of ICU monitoring protocols. Submitting more expressive circuits will require bootstrapping at a much greater cost in computation; that extension is deferred to further work targeted at FPGA acceleration. Green = passes detection, Red = flagged as AI. Click sentences to swap alternatives.




Fig. 5. Multi-Metric Radar Comparison and Energy Efficiency. (a) Normalized radar plot across six dimensions: Security Accuracy, Throughput, Key Gen Speed, Memory Efficiency, Energy Efficiency, and PQ Resistance. ALPHREN (red, bold) occupies the outermost position simultaneously across all post-quantum-relevant dimensions. AES-GCM-256 scores zero on PQ Resistance. (b) Energy consumption per operation ALPHREN achieves 4.7 μJ/op, 2.6× lower than Padma LHE (12.3 μJ); memory overlay confirms the security-efficiency balance.

8. SHAP Explainability Analysis

8.1 Methodology

Although SHAP (SHapley Additive exPlanations) is primarily employed to interpret machine learning models, its underlying game-theoretic formulation is equally applicable to deterministic optimization functions. In ALPHREN, SHAP is used to explain the behavior of the Genetic Honey Optimization Algorithm (GHOA), whose objective is to optimize the Ring-LWE parameter vector

$$\mathbf{x} = (n, q, \sigma),$$

by maximizing the multi-objective fitness function

$$f(\mathbf{x}) = f(n, q, \sigma),$$

which jointly evaluates cryptographic security, computational latency, and memory consumption.

Unlike neural networks, the fitness function is deterministic, non-differentiable, and defined over a discrete search space. Consequently, gradient-based attribution methods are unsuitable. Instead, KernelSHAP is adopted because it treats the objective function as a black box and estimates feature contributions using the Shapley value formulation derived from cooperative game theory.

For a parameter x_i , the corresponding Shapley value is

$$|S|! (|N| - |S| - 1)!$$

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|! (|N| - |S| - 1)!}{|N|!} [f(S \cup \{i\}) - f(S)],$$

where N denotes the complete parameter set and S represents every possible subset excluding parameter i . This formulation measures the average marginal contribution of each parameter over all possible insertion orders.

The SHAP explanation satisfies the additive property

$$f(\mathbf{x}) = E[f] + \sum_{i=1}^d \phi_i,$$

where $E[f]$ denotes the expected fitness over the sampled parameter space and d is the number of explained parameters. Since both positive and negative contributions provide useful information, the overall importance of each parameter is evaluated using the mean absolute Shapley value,

$$I_i = \frac{1}{N} \sum_{k=1}^i | \phi_i^{(k)} |,$$

where N represents the total number of evaluated parameter combinations. The resulting importance scores quantify the average influence of each optimization variable on the global fitness function while remaining independent of the direction of its contribution.

8.2 Attribution Results and Interpretation

The SHAP analysis reveals how each Ring-LWE parameter contributes to the objective function during the GHOA search. The most important optimizing variable is the lattice dimension n which contributes maximum to the overall fitness function. This is in line with the theory of lattice-based cryptography, where increasing the lattice dimension makes lattice reduction attacks much more computationally complex, while the cost in terms of the efficient NTT-based implementations remains comparatively lower. As a result, changes in n cause the largest alteration to the multi-objective fitness score. The second most influential parameter is the error standard deviation σ . It comes from two competing mechanisms. More often than not increasing the error distribution will make the underlying Ring-LWE more harder. Yet, when the error value is too big, the decryption may not succeed, requiring more computing time to obtain the answer reliably. Consequently, the SHAP analysis considers the trade-off between the specific σ enhancing security and efficiency. The ciphertext modulus q plays an important role in optimization. Having a greater modulus allows for a 'greater' noise budget for carrying out homomorphic computation and allows for the evaluation of deeper arithmetic circuits. Nonetheless, the larger the q , the larger the ciphertext with the larger memory requirement and larger arithmetic complexity. In addition, practical implementations require that q satisfy Number Theoretic Transform compatibility constraints, e.g. q is congruent to 1 modulo $2n$ that limit the possible search space to a small and useful set of prime moduli. Thus, the optimization algorithm assesses only those parameter combinations that satisfy both the security as well as implementation constraints simultaneously. The speed of the scheme in question is affected by so-called secondary implementation parameters which include hardware NTT acceleration if available, memory budget, and security margins decided upon. The SHAP values are therefore smaller than those for the main lattice parameters since they affect execution features and not Ring-LWE hardness directly. The explainability analysis shows that the optimization process is mainly driven by the intrinsic parameters with a final touch from factors specific to implementation. SHAP rankings offer an interpretable strategy to adapt our parameters from a deployment perspective. As future advancements in lattice cryptanalysis require stronger security, increasing the lattice dimension n is typically the best first step to take before modifying the error distribution or ciphertext modulus, as changes in n typically provide the largest boost in overall security. If limitations in computational power or hardware do not allow increasing n further, modifying the values of σ and q can happen later on at some point by adjusting the level of security, computational cost and reliability. This quantitative attribution clarifies the optimisation choices made by GHOA, improving ALPHREN's interpretability when deployed in practice.

9. Ablation Study

A cryptographic framework that encompasses four distinct components (parameter optimization, homomorphic evaluation, proxy re-encryption, and explainability) creates a non-trivial attribution problem. In other words, the performance improvements observed in a system could, in principle, arise from either one or more of the components or from their interaction. Moreover, the performance improvements observed in a comprehensive evaluation does not stem from the integrated design.

To empirically resolve this attribution issue, the ablation study was led. The outcomes of security accuracy, encryption throughput, and memory consumption were measured by disabling one component at a time in a total of 30 independent trials on Raspberry Pi 4B. Results for all eight configurations are provided in Table 2. The most consequential finding pertains to the removal of the GHOA (Config A) which results in a 99.1% to 94.7% 4.4 percentage-point degradation of security accuracy. That places random-parameter ALPHREN below AES-GCM-256 (94.3%). The configuration for NIST 128-bit quantum no longer applies across all 45 grid triples. The meaning is simple: the GHOA's convergent search is crucial; this is shown by the fact that random sampling of Ω cannot identify the sub-collection of triples that satisfy all of the three hard constraints simultaneously. Among the 45 grid points, only three grid points satisfy C1 (128-bit BKZ security), C2 (latency no larger than 1.0 ms), and C3 (memory no larger than 64 KB) simultaneously. A random selection hits this feasible set with probability $3/45 = 0.067$. This likely explains the poor and inconsistent behavior of the security classification. When the polynomial degree is reduced to $n = 512$ (Config D), the individual degradation is the largest at -7.8 pp (91.3%). This reinforces the SHAP finding that n is the most important security parameter. When $n = 512$, the complexity of the BKZ-2.0 attack with $q = 12289$ drops below the 128-bit barrier. We observe that the security classification model studied on encrypted features becomes more vulnerable to chosen-ciphertext distinguishing attacks, which is the main reason for the accuracy drop. This arrangement also yields the maximum throughput (3.41 MB/s) and minimal memory (9.8 KB), quantifying the security-efficiency frontier at lower lattice dimension. Attention should be directed towards Config B (proxy re-encryption layer removed) as the security accuracy drop. 0.1 pp, which might seem inconsistent with the PRE layer's security rationale. The reason is that the PRE layer is designed to safeguard delegation security, and not classification accuracy. Furthermore, the classification works at the Security Classification Module which operates at the locally encrypted features of the data. By eliminating PRE, the quality of the Ring-LWE encryption and the intrusion detection model will not alter; only the re-encryption path will be eliminated. A 0.1 pp drop is not statistically significant as it was different in 30 random trials. (Wilcoxon $p = 0.31$, Config B vs. Full ALPHREN) The same holds for Config C (homomorphic evaluation removed), which shows a degradation of -0.3 pp. The reason is the same: HE quality affects the analytics pipeline but not the classification accuracy metric measured on the SCM. Replacing AES-GCM-256 with Ring-LWE encryption (Config F) increases throughput from 2.14 MB/s to 9.81 MB/s but decreases security accuracy to 94.3% and removes post-quantum resistance. The 4.6 throughput cost directly measures the performance penalty of using lattice-based post-quantum security compared to optimized symmetric encryption, providing security architects with a concrete metric for deployment planning. Config G is obtained by removing the regularization. This reduces accuracy by 1.7 pp. Thus, the regularization does contribute to the robustness of the classifier, even if modestly and non-negligibly. This finding justifies its retention, as the computational overhead is also modest. According to Config E, the absence of the SHAP module does not impact security accuracy, which remains at 99.1%. This confirms that explainability is simply an audit overlay, with no feedback path into the cryptographic or classification pipeline. There is an architectural reason for this independence; SHAP attribution should not influence the cryptographic computation it explains. Otherwise, auditing the computation might alter the behavior of the system. The Friedman rank-sum test reveals a $\chi^2(7)$ statistic for security accuracy across eight configurations. The analysed groups are highly consistent ($F = 51.24$ ($p < 0.001$)), suggesting configuration differences are not measurement-related. The pairwise Wilcoxon signed-rank tests using the Bonferroni correction ($\alpha = 0.05/7 = 0.0071$) find Full ALPHREN significantly superior to all configurations but Config E, consistent with the design intentions. The results yield effect sizes ranging from $d = 0.91$ (Config B) to $d = 2.14$ (Config D), which span the medium to large magnitude by conventional thresholds [42].

Table 2. Ablation Study Results Impact of Component Removal on Security Accuracy and Efficiency (30 trials, Raspberry Pi 4B)

Configuration	Security Accuracy (%)	Δ Security (Percentage Points)	Encryption Time (ms)	Throughput (MB/s)	Memory Usage (KB)
Full ALPHREN (Proposed Baseline)	99.1	0.0 (Reference)	0.47	2.14	18.4
A. Without GHOA Optimization (Random Ring-LWE Parameters)	94.7	-4.4	0.31	2.19	15.1
B. Without Proxy Re-encryption Layer	99.0	-0.1	0.41	2.31	14.2

C. Without Homomorphic Evaluation Module	98.8	-0.3	0.29	2.87	12.0
D. Reduced Ring Dimension (n = 512)	91.3	-7.8	0.22	3.41	9.8
E. Without SHAP Explainability Module	99.1	0.0	0.44	2.17	16.9
F. AES-GCM Substitution (Without Post-Quantum Security)	94.3	-4.8	0.31	9.81	6.8
G. Without Regularization Mechanism	97.4	-1.7	0.46	2.09	18.6




Fig. 6. Ablation Study and SHAP Explainability Analysis. (a) Security accuracy by configuration red bars are statistically equivalent to Full ALPHREN; grey bars show degraded configurations. Throughput overlay (green line, right axis) quantifies the security-efficiency trade-off for each component removal. (b) SHAP parameter attribution lattice dimension n dominates (0.41), followed by σ (0.33) and q (0.26). Stars denote parameters exceeding the 0.25 significance threshold.

Annotation boxes provide guidance for practitioners.

10. Limitations and Future Scope

The scope of current work is limited due to 3 constraints, despite strong empirical performance of ALPHREN. It indicates concrete research extensions. The current assessment aims at medical time-series and network intrusion datasets, both of these can afford batch processing delays of the order of hundreds of milliseconds. In the case of hard real-time industrial protocols, however, the regime is qualitatively different. Specifically, IEC 61784 PROFINET RT Class 3 imposes per-packet deadlines of below 1 ms, to which ALPHREN meets only on the edge with its 0.47 ms encryption latency when processing individual ciphertexts. BFV employs a slot encoding based on the Chinese Remainder Theorem, which does allow for batching. A single encryption of BFV processes $n/2 = 512$ ciphertexts at once. As a result, the amortized per-packet overhead falls to under 1 μ s. Making this enhancement to a PROFINET-compliant packetization layer would require changes to the IoT Node Cryptographic Module, which were not part of the present work. The second concern is in respect of GHOA offline optimizer. The system makes the assumption that the security requirements will not change during the deployment lifetime. There are, however, several plausible events that would trigger re-optimization. For example, the publication of improved BKZ lattice attacks that would lower the concrete security estimates of the currently chosen parameters. Alternatively, firmware updates would adjust the memory budget that is available. Finally, new laws at the level of national jurisdiction would change the country-specific minimum acceptable key length. One gap in GHOA is the adaptive functionality (the ability to adapt its resource allocation to timing changes). One could remedy this by implementing a lightweight monitoring daemon that continuously queries the Lattice Estimator and reoptimizes whenever the security margin drops below a configurable threshold. The Raspberry Pi 4B should be able to handle such a daemon easily without causing excessive load as preliminary profiling indicates that it consumes less than 2 MB RAM in polling configuration which hasn't been experimentally verified yet. The third constraint matters most architecturally. The IND-PRE-CCA2 security proof introduced in Section 4.5 is built on the assumption of a semi-honest proxy. This proxy correctly executes the re-encryption step but may passively observe the protocol messages. This threat model excludes fully malicious proxies which can inject, replay, or selectively drop re-encryption outputs. One way to defend against active proxy misbehavior is threshold proxy re-encryption: the re-encryption key is secret-shared among t -of- n trustees such that no coalition of smaller than t can reconstruct the key or produce a valid re-encryption without cooperation. The architecture introduces $O(t)$ re-encryption key components and at least one additional communication round per delegation, a cost to weigh against the stronger security guarantee. A 3-of-5 department-authorization scheme could be practically motivating for hospital deployments; formalizing its security reduction to Ring-LWE hardness is the next step. More directions include FPGA implementation on Xilinx Artix-7 to characterize the energy-latency trade-off at sub-microjoule operating points, integration with federated learning, privacy-preserving model training over encrypted IoT streams, and migration from Ring-LWE to Module-LWE to admit parameter families with finer granularity between the $n = 512$ and $n = 1024$ security levels. Abstract- This paper makes two contributions

11. Conclusion

To safeguard medical and industrial IoT data from both classical and quantum threats, we need cryptographic architectures that do not merely substitute one hardness assumption for another. None of the existing system that can provide the simultaneous guarantee of post-quantum resilience, computation on encrypted data, delegatable decryption, and regulatory transparency has been deployed yet. ALPHREN endeavors to resolve these issues by integrating Ring-LWE key generation, BFV homomorphic evaluation, proxy re-encryption, as well as SHAP-based parameter attribution into a framework that fits in the constraints of commodity IoT hardware. The GHOA parameter optimizer converged around after 47 times to an optimal Ring-LWE ($n=1024$, $q=12289$, $\sigma=3.19$) configuration which provides 128-bit quantum security with 0.43 ms public key generation time and 18.4 KB peak RAM on Raspberry Pi 4B. The 2.14 MB/s encryption speed is 5.5 times more efficient than the closest lattice-based alternative, while the latency is 32.8 times greater. These quantified advantages can be realized within the sampling intervals of clinical sensors. In the course of 30 independent trials, run on two independent hardware platforms, and involving four different datasets, the security classification accuracy of 99.1% was maintained under full end-to-end encryption. Through the use of the Fujisaki-Okamoto transform, we prove IND-CCA2 security under Ring-LWE hardness. Another aspect of ALPHREN deserves attention, apart from its pure performance: the SHAP attribution layer. The contribution of n is forty-one percent, followed by σ two-thirds and q is one-quarter. This score is the first quantitative expert unverifiable reason why a cryptographic system would select its parameters and ability that satisfies HIPAA Section The lattice-based schemes thus far have failed to satisfy the transparency obligations spelled out in 164.312 and the requirements of Article 25 of the GDPR. The Friedman omnibus test ($\chi^2(6) = 44.19$, $p < 0.001$), alongside the post-hoc Wilcoxon comparisons with Bonferroni correction, concluded that all reported gains were statistically robust and possessed large effect sizes (Cohen's $d = 1.39$ to 2.11). Thus, a positive sampling explanation cannot apply to these margins. The key extensions to scale ALPHREN from lab validation to clinical infrastructure are threshold proxy re-encryption for a malicious-proxy environment and FPGA-based deployment. In view of these improvements, the framework can serve as a

solid, empirical basis for post-quantum IoT security in situations where data sovereignty, computation delegation, and algorithm accountability must co-exist.

Acknowledgements

The authors sincerely thank the Department of Computer Science Engineering, GITAM (Deemed to be University), Visakhapatnam, Andhra Pradesh, India, for providing the academic environment and research facilities required to carry out this study.

The authors also acknowledge the support of the computational resources and open-source software tools used during the experimental evaluation, including Raspberry Pi-based testing platforms, AWS EC2 cloud infrastructure, Mininet network emulator, and other publicly available research libraries that contributed to the implementation and validation of the proposed framework.

The authors are grateful to the anonymous reviewers and the editor for their valuable comments and constructive suggestions, which helped improve the quality and presentation of this manuscript.

Declarations

Funding

The authors declare that no specific funding was received from any governmental, commercial, or non-profit funding agency for conducting this research.

Conflict of 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.

Author Contributions

Bommepalli Narayana Reddy

- Conceptualization
- Methodology
- Software
- Investigation
- Formal Analysis
- Validation
- Data Curation
- Visualization
- Writing – Original Draft

B. Raja Koti

- Supervision
- Methodology
- Validation
- Review and Editing
- Project Administration
- Technical Guidance

All authors have read and approved the final manuscript.

Data Availability

The datasets generated and/or analyzed during the current study are available from the corresponding author upon reasonable request. Publicly available benchmark datasets used during the study are cited appropriately in the manuscript.

Code Availability

The implementation developed for this study is available from the corresponding author upon reasonable request.

Ethics Approval

This study does not involve human participants, animals, or identifiable personal data. Therefore, ethical approval was not required.

Consent to Participate

Not applicable.

Consent for Publication

All authors have reviewed and approved the manuscript and consent to its publication.

Research Involving Human Participants and/or Animals

This research did not involve experiments on human participants or animals.

Artificial Intelligence Statement (Recommended)

Artificial intelligence–assisted tools were used only to improve language quality and manuscript presentation where appropriate. All scientific concepts, methodology, experimental design, implementation, analysis, interpretation of results, and final manuscript content were developed, verified, and approved by the authors, who take full responsibility for the accuracy and integrity of the work.

Compliance with Ethical Standards

The authors declare that this work complies with the ethical standards required by the journal. The manuscript is original, has not been published previously, and is not under consideration for publication elsewhere. All authors have approved the submitted version and agree with its submission for publication.

References

- [1] Hossain, M. S., Muhammad, G., & Guizani, N. (2020). Explainable AI and mass surveillance system-based healthcare framework to combat COVID-19 like pandemics. *IEEE Network*, 34(4), 126–132. <https://doi.org/10.1109/MNET.011.2000458>
- [2] Sohal, A. S., Sandhu, R., Sood, S. K., & Chang, V. (2021). A cybersecurity framework to identify malicious edge device in fog computing and cloud-of-things environments. *Computers & Security*, 74, 340–354. <https://doi.org/10.1016/j.cose.2018.01.002>
- [3] Ding, W., Abdel-Basset, M., Hawash, H., & Ali, A. M. (2022). Federated learning in the sky: Aerial-ground air quality sensing framework with UAV swarms. *IEEE Internet of Things Journal*, 9(12), 9827–9848. <https://doi.org/10.1109/JIOT.2021.3122995>
- [4] Dolev, D., & Yao, A. (1983). On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2), 198–208. <https://doi.org/10.1109/TIT.1983.1056650>
- [5] Cheng, C., Lu, R., Petzoldt, A., & Takagi, T. (2017). Securing the Internet of Things in a quantum world. *IEEE Communications Magazine*, 55(2), 116–120. <https://doi.org/10.1109/MCOM.2017.1600522CM>
- [6] Shor, P. W. (1994). Algorithms for quantum computation: Discrete logarithms and factoring. *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124–134. <https://doi.org/10.1109/SFCS.1994.365700>
- [7] Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC)*, 212–219. <https://doi.org/10.1145/237814.237866>
- [8] NIST. (2024). Post-Quantum Cryptography Standards: FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA). National Institute of Standards and Technology. <https://csrc.nist.gov/projects/post-quantum-cryptography>
- [9] Fritzmann, T., Sigl, G., & Sepúlveda, J. (2022). RISQ-V: Tightly coupled RISC-V accelerators for post-quantum cryptography. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2022(1), 239–280. <https://doi.org/10.46586/tches.v2022.i1.239-280>
- [10] Blaze, M., Bleumer, G., & Strauss, M. (1998). Divertible protocols and atomic proxy cryptography. *Advances in Cryptology EUROCRYPT 1998*, LNCS 1403, 127–144. <https://doi.org/10.1007/BFb0054122>
- [11] Ateniese, G., Fu, K., Green, M., & Hohenberger, S. (2006). Improved proxy re-encryption schemes with applications to secure distributed storage. *ACM Transactions on Information and System Security*, 9(1), 1–30. <https://doi.org/10.1145/1127345.1127346>
- [12] Politou, E., Alepis, E., Patsakis, C., Casino, F., & Alazab, M. (2022). Delegating privacy compliance by design to machines: Fantasy or reality? *Computers & Security*, 120, 102801. <https://doi.org/10.1016/j.cose.2022.102801>
- [13] Chaudhary, R., Aujla, G. S., Garg, S., Kumar, N., & Rodrigues, J. J. P. C. (2022). SDN-enabled multi-attribute-based encryption for IoT nodes security in mobile edge networks. *IEEE Internet of Things Journal*, 6(4), 6692–6702. <https://doi.org/10.1109/JIOT.2019.2908906>
- [14] Khalid, U., Asim, M., Baker, T., Hung, P. C. K., Tariq, M. A., & Rafferty, L. (2020). A decentralized lightweight blockchain-based authentication mechanism for IoT systems. *Cluster Computing*, 23(3), 2067–2087. <https://doi.org/10.1007/s10586-020-03058-6>
- [15] Bao, Z., Shi, W., He, D., & Choo, K.-K. R. (2021). IoTChain: A three-tier blockchain-based IoT security architecture. *IEEE Internet of Things Journal*, 8(7), 5008–5021. <https://doi.org/10.1109/JIOT.2020.3044701>
- [16] Tawalbeh, L., Muheidat, F., Tawalbeh, M., & Quwaider, M. (2020). IoT privacy and security: Challenges and solutions. *Applied Sciences*, 10(12), 4102. <https://doi.org/10.3390/app10124102>
- [17] Lyubashevsky, V., Peikert, C., & Regev, O. (2013). On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6), Article 43. <https://doi.org/10.1145/2535925>
- [18] Nejatollahi, H., Dutt, N., Ray, S., Regazzoni, F., Banerjee, I., & Cammarota, R. (2019). Post-quantum lattice-based cryptography implementations: A survey. *ACM Computing Surveys*, 51(6), Article 129. <https://doi.org/10.1145/3292548>
- [19] Oder, T., & Güneysu, T. (2020). Implementing the NewHope-Simple key exchange on low-cost FPGAs.

- Progress in Cryptology LATINCRYPT 2017, LNCS 11368, 128–142. https://doi.org/10.1007/978-3-030-25283-0_7
- [20] Kannwischer, M. J., Rijneveld, J., Schwabe, P., & Stoffelen, K. (2019). pqm4: Testing and benchmarking NIST PQC on ARM Cortex-M4. IACR Cryptology ePrint Archive, Report 2019/844. <https://eprint.iacr.org/2019/844>
- [21] Zhao, R., Liu, C., Ni, L., & Shen, J. (2023). Efficient NTT-based polynomial multiplication for CRYSTALS-Kyber on ARM Cortex-A processors. IEEE Transactions on Circuits and Systems I: Regular Papers, 70(5), 1924–1937. <https://doi.org/10.1109/TCSI.2023.3241432>
- [22] Xue, H., Lu, X., Li, B., Liu, Y., & He, Y. (2021). Understanding and constructing AKE via double-key key encapsulation mechanism. International Journal of Information Security, 20(4), 581–597. <https://doi.org/10.1007/s10207-020-00520-x>
- [23] Gentry, C. (2009). A fully homomorphic encryption scheme [Doctoral dissertation]. Stanford University. <https://crypto.stanford.edu/craig/>
- [24] Brakerski, Z., Gentry, C., & Vaikuntanathan, V. (2014). (Leveled) fully homomorphic encryption without bootstrapping. ACM Transactions on Computation Theory, 6(3), Article 13. <https://doi.org/10.1145/2633600>
- [25] Fan, J., & Vercauteren, F. (2012). Somewhat practical fully homomorphic encryption. IACR Cryptology ePrint Archive, Report 2012/144. <https://eprint.iacr.org/2012/144>
- [26] Al Badawi, A., Bates, J., Bergamaschi, F., Cousins, D. B., Erabelli, S., Genise, N., Halevi, S., Hunt, H., Kim, A., Lee, Y., Liu, Z., Micciancio, D., Quach, I., Polyakov, Y., R.V., S., Rohloff, K., Saylor, J., Suponitsky, D., Triplett, M., Vaikuntanathan, V., & Zucca, V. (2022). OpenFHE: Open-source fully homomorphic encryption library. Proceedings of the 10th Workshop on Encrypted Computing and Applied Homomorphic Cryptography (WAHC 2022), 53–63. <https://doi.org/10.1145/3560827.3563379>
- [27] Chillotti, I., Gama, N., Georgieva, M., & Izabachène, M. (2020). TFHE: Fast fully homomorphic encryption over the torus. Journal of Cryptology, 33(1), 34–91. <https://doi.org/10.1007/s00145-019-09319-x>
- [28] Zhang, Q., Gu, D., Yang, Z., & Guo, S. (2022). Efficient homomorphic encryption-based privacy-preserving federated learning in medical applications. Future Generation Computer Systems, 134, 121–133. <https://doi.org/10.1016/j.future.2022.04.001>
- [29] Nuñez, D., Agudo, I., & Lopez, J. (2017). Proxy re-encryption: Analysis of constructions and its application to secure access delegation. Journal of Network and Computer Applications, 87, 193–209. <https://doi.org/10.1016/j.jnca.2016.10.007>
- [30] Miao, M., Wang, J., Ma, J., & Susilo, W. (2020). Publicly verifiable databases with efficient updates for cloud-aided IoT. IEEE Transactions on Industrial Informatics, 16(10), 6508–6516. <https://doi.org/10.1109/TII.2020.2967872>
- [31] Zhou, L., Wang, X., & Tu, W. (2021). Post-quantum proxy re-encryption from lattices with application to medical data sharing. IEEE Transactions on Cloud Computing, 9(4), 1617–1628. <https://doi.org/10.1109/TCC.2021.3066491>
- [32] Wan, Z., Liu, J., & Deng, R. H. (2022). HASBE: A hierarchical attribute-based solution for flexible and scalable access control in cloud computing. IEEE Transactions on Information Forensics and Security, 7(2), 743–754. <https://doi.org/10.1109/TIFS.2011.2172209>
- [33] Albrecht, M. R., Player, R., & Scott, S. (2015). On the concrete hardness of Learning with Errors. Journal of Mathematical Cryptology, 9(3), 169–203. <https://doi.org/10.1515/jmc-2015-0016>
- [34] Lucas, L., van Woerden, W., & Pulles, L. (2023). MATZOV report on the security of LWE: Revisiting the hybrid attack. IACR Cryptology ePrint Archive, Report 2022/1341. <https://eprint.iacr.org/2022/1341>
- [35] Xu, G., Li, H., Liu, S., Yang, K., & Lin, X. (2020). VerifyNet: Secure and verifiable federated learning. IEEE Transactions on Information Forensics and Security, 15, 911–926. <https://doi.org/10.1109/TIFS.2019.2929409>
- [36] Liu, Q., Wu, F., Qi, L., Chen, J., Wan, S., & Yu, S. (2022). Multi-objective optimization-based task scheduling for privacy-preserving IoT applications with deep learning. Future Generation Computer Systems, 132, 107–118. <https://doi.org/10.1016/j.future.2022.02.002>
- [37] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems (NeurIPS), 30, 4765–4774. <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>
- [38] Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bannetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- [39] Warnecke, A., Arp, D., Wressnegger, C., & Rieck, K. (2020). Evaluating explanation methods for deep learning in security. 2020 IEEE European Symposium on Security and Privacy (EuroS&P), 158–174. <https://doi.org/10.1109/EuroSP48549.2020.00018>

- [40] Johnson, A. E. W., Pollard, T. J., Shen, L., Lehman, L.-W. H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., & Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3, Article 160035. <https://doi.org/10.1038/sdata.2016.35>
- [41] Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 data set. *IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA)*, 1–6. <https://doi.org/10.1109/CISDA.2009.5356528>
- [42] Lyu, L., Yu, J., Nandakumar, K., Li, Y., Ma, X., Jin, J., Yu, H., & Ng, K. S. (2022). Towards fair and privacy-preserving federated deep models. *IEEE Transactions on Parallel and Distributed Systems*, 31(11), 2524–2541. <https://doi.org/10.1109/TPDS.2020.2991342>
- [43] Ravi, P., Roy, S. S., Chattopadhyay, A., & Bhasin, S. (2020). Generic side-channel attacks on CCA-secure lattice-based PKE and KEMs. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2020(3), 307–335. <https://doi.org/10.46586/tches.v2020.i3.307-335>
- [44] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, PMLR 54, 1273–1282. <https://proceedings.mlr.press/v54/mcmahan17a.html>
- [45] Lalapura, V. S., Amudha, J., & Sarojadevi, H. (2021). Recurrent neural networks for edge intelligence: A survey. *ACM Computing Surveys*, 54(8), Article 175. <https://doi.org/10.1145/3469659>