UDC 004.056.55 **doi: 10.32620/reks.2026.2.17**

Giorgi AKHALAIA[1], Maksim IAVICH[1], Razvan BOCU[2]

[1.] Caucasus University, Tbilisi, Georgia
[2.] Transilvania University of Brasov, Brasov, Romania

VERKLE-FRI: A NOVEL ARCHITECTURE FOR STATELESS AND QUANTUM-RESISTANT VECTOR COMMITMENTS

The subject matter of the article is the cryptographic integrity of digital authentication systems facing quantum computing threats, specifically focusing on post-quantum alternatives and efficient authenticated data structures. The goal is to design and formally analyze VERKLE-FRI—a hybrid architecture synthesizing Verkle tree proof-size reduction with FRI-based quantum-resistant commitments, establishing a scalable, stateless, and quantum-secure framework. The tasks are: analyze limitations of hash-based signatures and Merkle trees; evaluate polynomial commitment schemes (KZG, Bulletproofs, FRI, lattice-based); propose a hybrid Verkle-FRI design; develop a formal security proof against classical and quantum adversaries; execute complexity analysis with concrete implementation parameters. The methods used are: theoretical cryptographic analysis, formal security modeling via reductionist proofs, algebraic methods over finite fields, polynomial interpolation, random oracle model, FRI protocol with DEEP-FRI optimization, Merkle trees, vector commitments, and asymptotic complexity analysis. The following results were achieved: a novel architecture where Verkle node vectors are polynomial-encoded, committed via Merkle trees over FRI codewords, and verified through FRI with out-of-domain sampling. A formal proof establishes λ-bit quantum security using 2λ-bit hash functions. Complexity yields proof size $O(\lambda \log^2 N)$, prover time $O(\lambda N \log N)$, and verifier time $O(\lambda \log N)$. Concrete 128-bit quantum parameters include SHA3-512, field size $\approx 2^{255}$, branching factor 256, and 128 FRI rounds, achieving soundness error $\leq 2^{-127}$. For a concrete benchmark authenticating 2^{26} elements, a traditional Merkle proof requires ≈0.8 KB, whereas our VERKLE-FRI proof requires ≈180 KB. While larger, this provides quantum resistance and eliminates the trusted setup, a critical trade-off for long-term security. Conclusions. Scientific novelty consists in: 1) the first hybrid Verkle-FRI architecture replacing pairing-based assumptions with hash-based proximity testing; 2) a formal security proof reducing security to hash collision resistance and FRI soundness; 3) quantified efficiency-security trade-offs; 4) a viable pathway for quantum-resistant infrastructure in blockchains, software distribution, and government communications.

Keywords: *Post-quantum Security; Cryptography; FRI; Polynomial Commitments; Digital Hash-Based Signatures.*

1. Introduction

Digital signatures serve as an indispensable pillar in cybersecurity, guaranteeing authenticity, ensuring data integrity, and providing legally binding non-repudiation. Established classical algorithms, such as RSA, DSA, and ECDSA, derive their security from the computational difficulty of mathematical challenges, such as integer factorization and discrete logarithms. However, the emergence of quantum computing, particularly through algorithms such as Shor's, has critically undermined these foundations, rendering conventional protocols vulnerable.

In response, hash-based digital signatures have gained prominence as a primary quantum-resistant alternative. Their security is anchored not in number theory but in the robust properties of cryptographic hash functions, specifically, their resistance to collisions, pre-image attacks, and second-preimage attacks. While quantum algorithms like Grover's can offer a quadratic speedup against symmetric primitives, this threat is effectively neutralized by proportionally increasing hash output lengths, thereby preserving long-term security assurances.

The evolution from the foundational one-time signatures conceived by Lamport and Merkle to contemporary, standardized schemes, such as XMSS, LMS, and SPHINCS+, marks a significant trajectory. These constructs, now endorsed by standards bodies like NIST, combine conceptual simplicity with operational transparency. Their inherent flexibility allows for strategic optimization across signature size, computational overhead, and storage demands, making them exceptionally suitable for diverse applications, from securing software supply chains and IoT ecosystems to protecting critical national infrastructure. [1].

Hash-based signatures are broadly categorized into three architectural families: One-Time Signatures (OTS), which provide the fundamental security building blocks; stateful Merkle tree schemes, which enable scalable authentication through hierarchical structures; and stateless constructions, which support unlimited signatures without the operational burden of persistent state management. [2] This progression illustrates the field's maturation from theoretical models to deployable, post-quantum solutions.

Our investigation is guided by the following pivotal research inquiries:

1. When integrated with post-quantum polynomial commitments, can advanced data structures like Verkle

ISSN 1814-4225 (print)
ISSN 2663-2012 (online)

trees overcome the inherent scalability and proof-size limitations of traditional Merkle-based hash signatures?

2. What is the critical performance and operational trade-offs between stateful and stateless hash-based signature schemes in practical, real-world post-quantum deployments?

3. To what extent can next-generation polynomial commitment schemes, such as those based on lattices or the Fast Reed-Solomon IOP of Proximity (FRI) serve as efficient and transparent replacements for existing paradigms, such as KZG commitments within Verkle trees and signature frameworks?

4. How do novel hybrid architectures, which merge Verkle trees with FRI-based commitments, perform when scalability, proof efficiency, and security are evaluated within large-scale authenticated data systems?

By addressing these questions, our study aims to elucidate the crucial design considerations and performance balances essential for constructing scalable, efficient, and quantum-secure cryptographic infrastructures. The current study charts a coherent pathway from core principles to practical, resilient solutions that meet the security demands of the future.

1.1. Motivation

This research is positioned within the critical and applied domain of Post Quantum Cryptography, intersecting Theoretical Computer Science and Cybersecurity. This study addresses the urgent practical challenge of protecting global digital infrastructure, including financial systems, software supply chains, and government communications, from the existential threat posed by quantum computers. The relevance is paramount; algorithms such as Shor's will break the foundational mathematics of current digital signatures, including RSA and ECDSA, necessitating a proactive transition to quantum-resistant alternatives to avert a systemic security collapse.

The development of scalable, efficient, and practical cryptographic systems that remain secure in a post-quantum future is the core problem. This translates into essential scientific and practical tasks: identifying quantum-resistant cryptographic primitives, transitioning them from theoretical security but inefficient constructs to systems capable of internet-scale deployment, and managing the inherent trade-offs among proof size, computational overhead, and operational complexity, all while striving for transparent systems that eliminate risky trusted setups.

This study builds upon established work in hash-based signatures and authenticated data structures. Previous research has solidified hash-based schemes as a quantum-resistant foundation and produced efficiency breakthroughs via Verkle trees. Significant advancements have been made in polynomial commitment schemes. However, a critical gap remains unresolved: the most efficient constructions, such as Verkle trees paired with KZG commitments, reintroduce quantum vulnerability, whereas post-quantum commitment alternatives, such as FRI, often incur significant performance penalties. A cohesive architecture that combines structural efficiency with quantum resistance and transparency has been lacking.

This study aims to design and formally analyze a novel cryptographic architecture that bridges this gap, providing a scalable, transparent, and quantum-resistant authentication framework. This purpose is achieved through five interconnected objectives: (1) to analyze the limitations of current hash-based and Merkle tree structures; (2) to evaluate the trade-offs of modern polynomial commitment schemes for verifiable data; (3) to propose a novel hybrid architecture combining Verkle trees with FRI-based commitments; (4) to furnish a formal security proof for the proposed scheme against quantum adversaries; and (5) to conduct a complexity analysis, providing concrete parameters for achieving targeted post-quantum security levels.

1.2. State of the art

The current state of the art in the transition to quantum-resistant cryptography is defined by several parallel and interdependent advancements. In the domain of digital signatures, hash-based schemes have emerged as the most mature and mathematically robust alternative to classical algorithms vulnerable to Shor's algorithm. The evolution from foundational one-time signatures (Lamport, Winternitz) to practical many-time schemes represents a significant engineering achievement. Standardized, stateful constructions, such as XMSS and LMS, which organize one-time keys within a Merkle tree, are now recommended for deployment. SPHINCS+ and its hyper tree structure constitute the cutting edge for stateless operation, eliminating key management overhead at the expense of larger signature sizes. These schemes derive their security from the cryptographic hardness of hash functions, which can be mitigated by parameter scaling, a manageable compromise compared to the catastrophic failure of number theoretic primitives, while subject to a quadratic speedup via Grover's algorithm.

Concurrently, the frontier has advanced from traditional Merkle trees to more efficient constructions for scalable cryptographic systems and verifiable data structures. Here, Verkle trees represent the current state of the art for proof compaction. By replacing sibling hash chains with vector commitments at internal nodes, Verkle trees achieve an exponential reduction in proof size, collapsing authentication paths from logarithmic complexity to practically constant size (e.g., 300–600 bytes), which is critical for blockchain and large-scale data systems.

However, the efficiency of modern Verkle trees is intrinsically linked to the state of the art in polynomial commitment schemes (PCS). The most efficient implementations rely on the KZG commitment scheme, which offers constant proof sizes and fast verification. Despite its performance benefits, KZG has a critical limitation: it requires a trusted setup and its security is based on pairing-based elliptic curve cryptography, which is itself vulnerable to future quantum attacks. This creates a fundamental gap because the most efficient authenticated data structure depends on a pre quantum security assumption.

Transparent and post-quantum secure PCS have been developed in response. The Fast Reed Solomon IOP of Proximity (FRI) protocol, which is foundational to scalable transparent arguments of knowledge (STARKs), offers post-quantum security based solely on hash functions and does not require a trusted setup. Alternative lattice-based constructions provide another quantum-resistant pathway under different hardness assumptions. Nevertheless, the current state of the art for these post-quantum commitments involves a pronounced trade-off: they incur significantly larger proof sizes and/or higher computational overhead compared to KZG, making their direct integration into critical efficiency structures, such as Verkle trees, a non-trivial challenge.

Therefore, the present scientific frontier is characterized by a dichotomy: one branch offers high efficiency with classical trust and security models (KZG-based Verkle trees), while the other offers quantum resistance and transparency with a performance penalty (FRI/lattice-based PCS). A cohesive architecture that synthesizes the structural efficiency of Verkle trees with the quantum-resistant security of a transparent commitment scheme without resorting to a trusted setup remains an unresolved problem at the forefront of post-quantum cryptographic engineering.

1.2.1 Hash-Based Signatures Schemes

Digital signatures are a cornerstone of modern cybersecurity, ensuring the authenticity, integrity, and non-repudiation of digital communications. Traditionally, widely used schemes such as RSA, DSA, and ECDSA have relied on integer factorization and discrete logarithm problems as hard number-theoretic problems. However, these foundations are threatened by the rapid development of quantum computing. Shor's algorithm demonstrates that classical public-key systems can be efficiently broken once sufficiently powerful quantum computers become available. In response, post-quantum cryptography (PQC) has gained significant momentum, with hash-based digital signature schemes emerging as one of the most promising candidates. Their security is grounded in the fundamental properties of cryptographic hash functions, collision resistance, second-preimage resistance, and one-wayness, which remain robust even despite quantum adversaries. While Grover's algorithm can quadratically reduce the effective security of hash functions, this can be mitigated simply by doubling output sizes, preserving long-term security. Lamport first introduced hash-based digital signatures and later expanded them with Merkle tree construction, enabling scalable authentication for many-time signatures. Schemes such as XMSS, LMS, and SPHINCS+ have been developed over decades of refinement, striking a balance between efficiency, statefulness management, and security. Standardization bodies such as NIST and the IETF have already recommended these schemes for applications requiring long-term security in the post-quantum era [3].

The importance of hash-based signatures lies not only in their resistance to quantum attacks but also in their conceptual simplicity, flexibility, and transparency. They allow trade-offs between signature size, computation speed, and storage requirements, making them adaptable to diverse applications, from securing software updates and IoT devices to protecting critical infrastructure. Furthermore, innovations such as Verkle trees and lattice-based commitments are opening new pathways toward more compact and efficient constructions.

One-time signature (OTS) schemes constitute the foundational layer within the family of hash-based digital signatures, with origins tracing to the seminal work of Lamport and Diffie. Their core security principle diverges from classical number-theoretic assumptions, such as integer factorization, instead deriving robustness directly from the computational hardness of inverting cryptographic hash functions. This foundational property renders OTS schemes intrinsically resistant to quantum attacks, as threats such as Grover's algorithm offer only a quadratic speedup against exhaustive search, a vulnerability that can be effectively neutralized by proportionally increasing hash output lengths. The defining and constraining characteristic of OTS is its strict one-time usability; each key pair is cryptographically secure for signing only a single message. Subsequent reuse exposes partial secret key material, enabling an adversary to reconstruct sufficient information to forge signatures. Despite this operational limitation, OTS schemes, including the Merkle signature scheme (MSS), its Winternitz-optimized variant (WOTS), and their modern standardized descendants such as XMSS and SPHINCS+, serve as the essential building block for more scalable and practical hash-based systems.

The classical Lamport scheme provides the clearest illustration of the OTS principle. Key generation for an n-bit message digest involves randomly selecting 2n secret values, denoted as $s_{i,b}$ for $i \in \{1,...,n\}$ and $b \in \{0,1\}$. The corresponding public key is computed by applying a one-way hash function H to each secret value,

resulting in the set p_i,b = H(s_i,b). The signer discloses the secret value s_i,b corresponding to the i-th bit b of the message digest to sign a message. Verification involves hashing each revealed s_i,b and confirming that the result matches the stored public key entry p_i,b. This scheme provides information-theoretic security assuming that H is a one-way, preimage-resistant function. A concrete example of a 256-bit digest requires 512 secret values, with a signature revealing 256 of them. While LD-OTS is exceptionally transparent and simple to implement, its primary drawback is inefficiency: both the public key and signature sizes scale linearly with 2n hash outputs, resulting in impractically large data footprints for direct use in most applications.

Winternitz One-Time Signature (W-OTS)

The Winternitz scheme introduces a crucial efficiency refinement by trading computational effort for a significantly reduced signature size. Instead of signing individual bits, the message digest is computed in base w, grouping bits into digits. For each of these len digits, a secret seed value is iteratively hashed several times equal to the digit's value to produce the signature component. The public key component is derived by hashing the same secret seed a fixed number of times (w-1). This structure allows a signature for a digit m_i to be the hash chain value at position m_i, while verification requires completing the chain by applying the hash function an additional (w-1 - m_i) times to reach the public key commitment. This design achieves a substantial compression of key and signature materials compared with the basic Lamport scheme, forming the basis for efficient hash-based signatures used in contemporary constructions.

Hash-based cryptography evolved from single-use one-time signatures (OTS) to scalable many-time signature schemes to transition from theoretical security to practical deployment. These constructions enable a single public key to authenticate numerous messages, which is a necessity for real-world applications such as software updates, secure communication, and IoT device authentication. The core strategy involves organizing many OTS key pairs within a hierarchical, hash-based structure, such as a tree, to contain key exposure and maintain security even after multiple signing operations. The Merkle Signature Scheme (MSS), which binds a forest of OTS public keys under a single master public key using a binary hash tree, was the foundational breakthrough. A signature comprises an OTS signature from a specific leaf and an authentication path, a sequence of sibling hashes allowing the verifier to recompute the root of the tree. While secure, MSS is inherently stateful; the signer must meticulously track which OTS leaf has been used to prevent catastrophic key reuse.

This concept was refined into standardized and efficient schemes. XMSS (eXtended Merkle Signature Scheme) and LMS (Leighton–Micali Signatures) optimize the tree structure and employ the Winternitz OTS+ (WOTS+) for smaller signatures. They support a large but finite number of signatures (e.g., 2^h for a tree of height h) and are recommended for applications such as firmware signing. However, they retain the statefulness of MSS, requiring secure state management.

For scenarios where state maintenance is impractical, stateless schemes, most notably SPHINCS+, were developed. It eliminates state management by using a complex hyper-tree of smaller Merkle trees and incorporating a few-time signature (FORS) at the leaf level. The trade-off is significantly larger signature sizes (tens of kilobytes) compared with stateful schemes, but this is often acceptable for the operational simplicity and robustness gained. SPHINCS+ has been standardized by NIST as a post-quantum digital signature. The 2024 status report of the NIST Post-Quantum Cryptography standardization project [4] further emphasizes the maturity of hash-based schemes like SPHINCS+ for general-purpose use, while also finalizing lattice-based standards (ML-DSA, ML-KEM) for different performance profiles.

Further scalability is achieved through generalized multi-tree schemes such as XMSS-MT, which uses a tree hierarchy to distribute signatures. The efficiency of all tree-based schemes, particularly for signing, relies on advanced traversal algorithms to generate authentication paths without storing the entire tree structure. Today, these many-time hash-based signatures are positioned for critical use in software supply chain security, blockchain protocols, and long-term data archiving, providing a quantum-resistant foundation for future cryptographic infrastructure.

The taxonomy of hash-based digital signatures is defined by their signing cardinality and state management, each of which is formalized by a tuple of algorithms (KeyGen, Sign, Verify) with distinct security parameters and asymptotic behaviors [5].

One-Time Signatures (OTS): Let $H: \{0,1\}^* \to \{0,1\}^n$ be a cryptographic hash function. For a security parameter λ, key generation samples $2n$ secret values $s_{i,b} \xleftarrow{\$} \{0,1\}^\lambda$ for $i \in [n], b \in \{0,1\}$. The public key is $pk = \{p_{i,b} = H(s_{i,b})\}$. To sign a message M, compute its digest $d = H(M) = (d_1, \ldots, d_n)$ and output the signature $\sigma = (s_{1,d_1}, s_{2,d_2}, \ldots, s_{n,d_n})$. Verification checks $H(\sigma_i) =? p_{i,d_i}$ for all i. Security is reduced to the preimage resistance of H, with existential forgery probability $\epsilon \leq \mathrm{Adv}_H^{\mathrm{preimage}}(\mathcal{A})$. The scheme satisfies $\mathrm{Sign}(sk, M)$ is defined for exactly one query; a second query renders $\mathrm{Adv}_{OTS}^{\text{EUF-CMA}}(\mathcal{A}) \approx 1$. Few-Time Signatures (FTS): An FTS is parameterized by (t, k, r), where $t \gg k$, and can sign up to r messages. KeyGen samples t secrets $sk = (s_1, \ldots, s_t)$ and computes

$pk = (H(s_1), \ldots, H(s_t))$. The signing function uses a random oracle $\mathcal{G}: \{0,1\}^* \to \binom{[t]}{k}$ to map a digest $d = H(M)$ to a k-element index set $I = \{i_1, \ldots, i_k\}$. The signature is $\sigma = (s_{i_1}, \ldots, s_{i_k})$. Verification computes $I = \mathcal{G}(d)$ and checks $H(\sigma_j) =? pk_{i_j}$ for $j \in [k]$. Security depends on subset resilience: the probability that an adversary, after q signing queries, can output M^* with $\mathcal{G}(H(M^*)) \subseteq \bigcup_{j=1}^{q} I_j$. For uniform $\mathcal{G}$, this probability is bounded by $\binom{qk}{k}/\binom{t}{k}$. After r signatures, the forgery advantage non-negligibly increases, defining the capacity of the scheme.

Stateful Many-Time Signatures (MTS): A stateful MTS (e.g., MSS/XMSS) uses a binary Merkle tree of height h to authenticate $N = 2^h$ OTS key pairs. Let $\mathsf{OTS.KeyGen}$ generate key pairs (sk_i, pk_i) for $i \in [N]$. The Merkle tree is constructed with leaves $L_i = H(pk_i)$, internal nodes $N_{i,j} = H(N_{i,2j} \parallel N_{i,2j+1})$, and root R. The long-term public key is R. The state is an index $\tau \in [N]$. To sign M at state τ, compute $\sigma_\tau = \mathsf{OTS.Sign}(sk_\tau, M)$. The signature is $\Sigma = (\tau, \sigma_\tau, pk_\tau, \mathcal{A}_\tau)$, where $\mathcal{A}_\tau$ is the authentication path: the sequence of sibling nodes from leaf L_τ to root R. Verification runs $\mathsf{OTS.Verify}(pk_\tau, M, \sigma_\tau)$, computes $L'_\tau = H(pk_\tau)$, and iteratively hashes L'_τ with $\mathcal{A}_\tau$ to obtain R'. It accepts iff $R' = R$. Security reduces to the collision resistance of H and the EUF-CMA security of the underlying OTS:

$$\mathsf{Adv}^{\text{EUF-CMA}}_{\text{MTS}}(\mathcal{A}) \le \mathsf{Adv}^{\text{CR}}_{\text{H}}(\mathcal{A}') + \mathsf{Adv}^{\text{EUF-CMA}}_{\text{OTS}}(\mathcal{A}'').$$

The scheme is stateful: incrementing τ is mandatory, and any repetition or rollback breaks security. Stateless Many-Time Signatures (MTS): A stateless MTS (e.g., SPHINCS$^+$) uses a d-layer hyper-tree. Each layer ℓ contains 2^{h_ℓ} subtrees of height h. A hyper-tree root R is the public key. Signing is deterministic: for message M, derive a random index $rand \leftarrow H(SK, M)$, which selects a unique path through the hyper-tree. The signature comprises the following:

1. A few-time signature (FORS) at the leaf.
2. d OTS signatures (one per layer).
3. d authentication paths (one per layer).

Formally, $\Sigma = (rand, \sigma_{FORS}, \{(\sigma^{(\ell)}_{OTS}, \mathcal{A}^{(\ell)})\}^{d}_{\ell=1})$. Verification reconstructs the path using rand, verifies each OTS and FORS signature, and recomputes the root R', accepting iff $R' = R$. The signature size is $|\Sigma| = O(\lambda \cdot (h \cdot d + k \cdot \log t))$ bits. Security is reduced to the second-preimage resistance of H and the security of the underlying few-time and one-time components, with the following

$$\mathsf{Adv}^{\text{EUF-CMA}}_{\text{Stateless-MTS}}(\mathcal{A}) \le \mathsf{Adv}^{\text{SPR}}_{\text{H}}(\mathcal{A}') + d \cdot \mathsf{Adv}^{\text{EUF-CMA}}_{\text{OTS}} + \mathsf{Adv}^{\text{EUF-CMA}}_{\text{FTS}}.$$

The key advantage is the elimination of state: $\mathsf{Sign}(SK, M)$ is a deterministic function of SK and M, making the scheme resilient to state corruption or desynchronization. The cost is large $|\Sigma|$ and slower verification relative to stateful MTS.[6]

1.2.2. Verkle Trees

The advancement of quantum computing, formalized by Shor's algorithm, renders classical cryptosystems vulnerable to the distinct logarithm problem. Hash-based signatures provide a post-quantum alternative because their security reduces to the preimage and collision resistance of a hash function H, where only a quadratic quantum speedup (Grover) applies. [3] To construct many-time signatures, a Merkle tree of height h authenticates $N = 2^h$ one-time public keys via a root R. However, proof sizes and verification costs scale as $O(\lambda h)$, where λ is the hash output length, limiting scalability.

Verkle trees address this problem by replacing hash commitments with vector commitments. For a vector $\vec{v} = (v_0, \ldots, v_{k-1})$ at a node, a commitment $C = \mathsf{Commit}(\vec{v})$ is generated. An authentication path for a leaf v_j requires a single opening proof π_j proving $\mathsf{Open}(C, j, v_j) = 1$, yielding constant-size proofs $O(1)$ in practice (~0.3–0.6 KB) versus $O(h)$ for Merkle trees.

While verification is efficient, often a single pairing or linear algebra operation, prover complexity is higher. For a lattice-based commitment under the Module-SIS assumption, committing to a length-k vector requires $\tilde{O}(k \cdot m \cdot \log q)$ operations for dimension m and modulus q, compared to $O(k)$ hashes for a Merkle node. Updates to $\vec{v}$ may require recomputing the entire commitment or generating a new witness, incurring $O(k)$ cost versus $O(\log k)$ for Merkle trees.[6]

The security reduces to the binding property of the vector commitment and the collision resistance of H. If a lattice-based commitment is used, the security depends on the hardness of $\mathsf{SIS}_{n,m,q,\beta}$, a stronger assumption than the hash collision resistance. The choice of branching factor k creates a direct trade-off: proof size scales inversely with k, while update complexity scales linearly with k. This defines an optimization problem $\min_k(\alpha \cdot \text{proof\_size}(k) + \beta \cdot \text{update\_cost}(k))$ for application-specific weights α, β.

Many efficient vector commitments (e.g., KZG) require a trusted setup, introducing a $\mathsf{Setup}(1^\lambda) \to crs$ phase and associated risks. Transparent alternatives (FRI, lattice-based) maintain security in a post-quantum

ISSN 1814-4225 (print)
ISSN 2663-2012 (online)

setting but with increased proof sizes or computational overhead. Consequently, Verkle trees are optimal for read-heavy, classically secure applications where proof compactness is critical. However, their suitability for write-heavy or post-quantum deployments without trusted setups remains an open optimization challenge. [6] To quantify the advantage of Verkle trees over Merkle trees more precisely, consider a tree storing $N = 2^{26}$ elements. A Merkle tree (branching factor 2) requires 26 authentication paths, each containing 26 hash values, totaling approximately 832 bytes (using SHA-256). A Verkle tree with a branching factor b = 256 reduces the height to $\log_{256}(2^{26}) = 4$ levels. With vector commitment proofs of ~300-600 bytes each (using KZG), the total proof size would be ~1.2-2.4 KB – already a significant reduction. Our VERKLE-FRI construction, while incurring larger proofs (~180 KB) due to the polylogarithmic overhead of the FRI protocol, maintains the same logarithmic height reduction while adding post-quantum security and transparency.

1.2.3. Post-quantum Polynomial Commitments

Post-quantum polynomial commitment schemes (PCS) provide a binding commitment C to a polynomial $f(X) \in \mathbb{F}[X]$ of degree $<d$, with an evaluation proof π that $f(z) = y$, while resisting quantum attacks. The classical KZG scheme with constant-sized proofs is vulnerable to Shor's algorithm because it relies on elliptic curve pairings. Thus, the field has shifted toward lattice- and hash-based constructions.

Lattice-based PCS derives security from problems such as Module-SIS or LWE. A leading approach, exemplified by the SLAP framework, achieves polylogarithmic proof size O(log^c d) and verification time under the Module-SIS assumption. The design comprises a compressing lattice commitment with a Merkle tree, yielding a common reference string (CRS) of size $\tilde{O}(\log d)$ and quasi-linear commitment time $\tilde{O}(d)$. The FMN (2024) framework formalizes such PCS within polynomial IOPs (PIOPs) to build succinct arguments, demonstrating proofs of size ~17 MB for 2^20 constraints—significantly larger than KZG but offering post-quantum security.

We examine their comparative efficiency to justify our choice of FRI over lattice-based commitments for the proposed Verkle-FRI architecture. Lattice-based VCs exhibit several trade-offs that make them unsuitable for our goals. First, concrete proof sizes remain large: even optimized schemes produce proofs on the order of megabytes for practical parameter sets (e.g., ~17 MB for 2^20 constraints), whereas FRI-based proofs for similar security levels are tens to hundreds of kilobytes. Second, lattice-based verification requires linear algebra over large-dimensional modules and rings, involving operations that are significantly more complex than the hash-based operations that dominate FRI verification. Third, many efficient lattice constructions require a structured reference string (though transparent variants exist), whereas FRI is fully transparent with no setup. Fourth, the security of lattice schemes rests on less mature assumptions (i.e., SIS/LWE with specific parameter regimes) compared with the well-studied collision resistance of hash functions underlying FRI. Consequently, while lattice-based PCS remains a promising research direction for post-quantum succinct proofs, FRI's smaller concrete proof sizes, simpler verification (hash-dominated), complete transparency, and more conservative security foundation make it the preferable primitive for our hybrid Verkle tree architecture[7].

Hash-based PCS, primarily via the FRI (Fast Reed-Solomon IOP of Proximity) protocol, provides transparent setups and quantum resistance based on hash functions. While not a standalone commitment, FRI enables PCS with logarithmic proof sizes $O(\lambda \log^2 d)$ but incurs larger concrete proof sizes and verification costs than pairing-based schemes.[7]

The performance gap between classical and post-quantum PCS is formalized by the following trade-offs:

1. Proof Size: $O(1)$ (KZG) vs. $O(\log^c d)$ (lattice/FRI).
2. Verification Cost: Dominated by $O(1)$ pairings (KZG) vs. $O(\log d)$ linear algebra over rings (lattice) or hash operations (FRI).
3. Setup: Trusted setup (KZG, some lattice schemes) vs. transparent setup (FRI, some lattice variants).
4. Quantum Security: Relies on the hardness of the SIS/LWE (lattice) or the hash collision resistance (FRI).

For applications like Verkle trees, where a vector commitment $\text{Commit}(\vec{v})$ is required, replacing KZG with a post-quantum PCS introduces these trade-offs directly: proof size grows from constant to polylogarithmic, and verification becomes more algebraically complex[7]. Current research focuses on optimizing these parameters to enable efficient, quantum-resistant authenticated data structures without trusted setups[9].

The Fast Reed-Solomon IOP of Proximity (FRI) forms the foundation for a transparent, post-quantum polynomial commitment scheme. At its core, FRI is an interactive oracle proof protocol that verifies a function $f^{(0)}: D^{(0)} \to \mathbb{F}$ is δ-close to a polynomial $P(X) \in \mathbb{F}[X]$ of degree less than d, where $|D^{(0)}| = N$. The protocol operates in a logarithmic number of rounds $r = \log_2 N$. In each round i, the verifier supplies a random challenge $\alpha_i \in \mathbb{F}$, prompting the prover to compute a folded function as follows: $f^{(i+1)}(x^2) = (1 - \alpha_i) \cdot f^{(i)}(x) + \alpha_i \cdot f^{(i)}(-x)$ defined over a do-

main $D^{(i+1)}$ of half the size. This recursive folding continues until the final function is sent for direct verification over a constant-sized domain. Soundness is bounded by $\epsilon \leq \delta + \frac{d}{|\mathbb{F}|}$, with prover complexity dominated by FFT operations at $O(N \log N)$ and verifier complexity at $O(\log N)$. To construct a non-interactive polynomial commitment from this protocol, a polynomial $P(X)$ is evaluated over a sufficiently large domain L to create a codeword $C = [P(g)]_{g \in L}$. A Merkle tree $\mathcal{M}$ is built over this codeword, with its root $\text{root}_{\mathcal{M}}$ serving as the binding commitment. To open an evaluation $P(z) = y$, the prover executes the FRI protocol on a function derived from C, simulating verifier challenges via the Fiat-Shamir heuristic. The resulting proof π includes the FRI interaction transcript and the necessary Merkle authentication paths for all queried codeword locations. Verification involves recomputing the challenges, checking Merkle proofs against the committed root, and validating the FRI consistency equations [10].

The scheme's security rests on the collision resistance of the underlying Merkle tree hash function H, providing post-quantum resilience where a quantum adversary's advantage is bounded by $\tilde{O}(2^{-n/2})$ for an n-bit hash output. A key advantage of the public-coin model is its transparency, as it requires no trusted setup and operates entirely in the public-coin model. However, this comes with significant performance trade-offs.[11] The proof size scales as $|\pi| = O(\lambda \log^2 N)$ bits, ranging from 50 to 250 kilobytes, which is orders of magnitude larger than the constant-sized pairing-based alternatives, such as KZG. While verification is asymptotically efficient at $O(\lambda \log N)$ hash operations, the verification of numerous Merkle paths dominates the concrete cost. Furthermore, the prover incurs substantial overhead from $O(N \log N)$ field operations with large constant factors due to FFT computations and Merkle tree construction.[12] The protocol also imposes a field constraint, performing optimally only over FFT-friendly fields with smooth multiplicative subgroups. Achieving negligible soundness requires repeating the protocol $t = \lambda$ times, where the error decays as $\epsilon \leq (1 - \rho)^t + \text{negl}(\lambda)$ for a rate parameter ρ [13]. A critical limitation of recursive proof systems is the high cost of expressing the FRI verifier, which is hash-intensive as an arithmetic circuit and often requires on the order of a million constraints, thereby creating a significant barrier to efficient recursive composition [14].

1.3. Objectives and tasks

The primary objective of this research is to design, formally analyze, and propose a novel hybrid cryptographic architecture that overcomes the critical efficiency-security trade-off in post-quantum verifiable data structures. Specifically, it aims to synthesize the exponential proof-size reduction of Verkle trees with the transparent, quantum-resistant security of FRI-based polynomial commitments, thereby constructing a scalable framework suitable for next-generation, stateless cryptographic systems. The following interconnected tasks were defined and systematically addressed to achieve this overarching objective:

1. To conduct a comprehensive analysis of the limitations inherent in existing hash-based signature schemes and traditional Merkle-tree-based authenticated data structures, with a specific focus on their impractical proof sizes and inherent vulnerabilities when integrated with classical, pre-quantum cryptographic commitments.
2. A detailed comparative evaluation of modern polynomial commitment schemes, including KZG, Bulletproofs, and post-quantum alternatives, such as FRI and lattice-based constructions, is performed. This evaluation focuses on quantifying the core trade-offs between proof size, verification cost, setup requirements (trusted vs. transparent), and quantum resistance.
3. This study proposes a novel hybrid architectural design that integrates a Verkle tree structure for optimal hierarchical efficiency with a FRI-based polynomial commitment scheme as its foundational cryptographic primitive, thereby ensuring transparency and post-quantum security without a trusted setup.
4. To develop a formal security proof for the proposed hybrid Verkle-FRI scheme and demonstrate its binding and soundness properties under standard cryptographic assumptions. This includes a dedicated reductionist analysis to establish its resilience against both classical and quantum adversaries, explicitly accounting for threats from Grover's algorithm.
5. To perform a thorough complexity analysis of the proposed architecture, asymptotic bounds for proof size, prover time, and verifier time are derived. Furthermore, concrete, actionable implementation parameters (e.g., hash function, field size, branching factor, and FRI rounds) are provided to achieve a target security level (e.g., 128-bit quantum security) in a practical deployment scenario.

The article is structured as follows: Section 2 presents the research materials and methods. Section 3 presents the results and discussion of the proposed architecture, the case study, and the evaluation results. Section 4 concludes the paper and outlines directions for future research.

2. Materials and research methods

This study employed a structured, multi-phase methodological approach combining theoretical analysis, cryptographic construction, formal security proof, and complexity analysis to develop and validate the proposed

ISSN 1814-4225 (print)
ISSN 2663-2012 (online)

hybrid architecture. This research did not involve empirical laboratory experiments but was conducted through rigorous mathematical and computational modeling.

Theoretical and Comparative Analysis: A comprehensive literature review and theoretical analysis of existing cryptographic primitives were conducted in the foundational phase. This included the following: Document Analysis – Systematic examination of foundational and contemporary academic literature, standards (e.g., NIST publications), and technical specifications related to hash-based signatures (Lamport, Winternitz, Merkle, XMSS, SPHINCS+), authenticated data structures (Merkle and Verkle trees), and polynomial commitment schemes (KZG, Bulletproofs, FRI, lattice-based constructions). Comparative Framework - A qualitative and quantitative comparison framework was established to evaluate the aforementioned schemes against key performance indicators: proof/signature size, computational overhead for generation and verification, state management requirements, security assumptions, and quantum resistance.

Architectural Design and Synthesis: The synthetic design of a new cryptographic system was the core of the methodological work. Hybrid Model Construction: The novel architecture was formulated by integrating two advanced constructs: the hierarchical, vector-commitment-based structure of a Verkle tree and the transparent, low-degree testing mechanism of the Fast Reed-Solomon IOP of Proximity (FRI) protocol. The design specified how node values in a Verkle tree are encoded as polynomials over a finite field, committed to using a Merkle tree of FRI codewords, and authenticated via paths using combined polynomial evaluations. Algorithm Specification: Detailed mathematical procedures were defined for the key operations: Commitment (polynomial interpolation, codeword generation, Merkle root calculation), Proof Generation (authentication path construction, application of the FRI folding protocol with DEEP-FRI optimization for out-of-domain sampling), and Verification (batch evaluation of path polynomials, execution of the FRI verification steps). Formal Security and Complexity Analysis: The proposed design underwent rigorous formal analysis to verify its theoretical validity.

Security Proof Formulation: A reductionist security proof was developed. Theorems for Completeness, Soundness, and Post-Quantum Security were formally stated and proven. The soundness proof employed a hybrid argument game sequence, reducing the overall construction security to the hardness of finding collisions in the underlying hash function and the FRI protocol's soundness. The resistance to quantum attacks (specifically Grover's algorithm) was analytically bounded, demonstrating that security is maintained by parameter scaling (e.g., using a 2λ-bit hash function for λ-bit quantum security). Complexity Analysis: Asymptotic computational complexity was derived for the following critical metrics: proof size ($O(\lambda \cdot \log^2 N)$), prover time ($O(\lambda \cdot N \log N)$), and verifier time ($O(\lambda \cdot \log N)$). Concrete calculations were performed to determine the optimal parameters for a target security level (e.g., 128-bit quantum security), specifying the use of SHA3-512, a finite field of size $\approx 2^{255}$, a branching factor (b) of 256, and a defined number of FRI rounds (t). Analytical Tools and Conceptual Framework: The research was conducted using established cryptographic and mathematical frameworks. Mathematical Framework: Algebra over finite fields, polynomial interpolation and evaluation, cryptographic hash function modeling (treated as random oracles in the security analysis), and complexity theory were used. Conceptual Models: The security analysis used standard cryptographic models, including the presumed existence of collision-resistant hash functions and the soundness of the FRI protocol in the Interactive Oracle Proof model.

In summary, the methodology was strictly analytical and constructive, leveraging established cryptographic theory to design a new system, formally prove its properties, and quantify its performance, thereby providing a complete theoretical foundation for the proposed hybrid Verkle-FRI scheme.

3. Results and Discussion

3.1. Proposed Hybrid Architecture: VERKLE-FRI

The inherent limitations of Merkle trees for large-scale applications, particularly their logarithmic proof growth, motivate the investigation into more efficient authenticated data structures. While Verkle trees offer exponential reductions in proof size via vector commitments, their reliance on schemes such as KZG introduces quantum vulnerabilities and trusted setup dependencies. This study introduces a novel hybrid architecture that combines the structural efficiency of Verkle trees with the post-quantum security of FRI-based polynomial commitments, resulting in a transparent, quantum-resistant framework for next-generation cryptographic systems [13].

The construction operates as follows: Consider a Verkle tree node with b child values

$$V = [v_0, v_1, \dots, v_{b-1}].$$

This vector is encoded as a polynomial f(X) of degree $d < b$ over a finite field $\mathbb{F}_p$ such that $f(\omega^i) = v_i$ for

all i, where ω is a primitive b-th root of unity. The commitment is generated by evaluating f on an extended domain L (where $|L| > b$) to form a codeword $C = [f(g)]_{g \in L}$, and then computing a Merkle tree over C; its root serves as the binding commitment.

The integrity of the committed polynomial is verified by using the FRI protocol. Beginning with the initial function $f^{(0)}: L \to \mathbb{F}_p$, each round i applies a random challenge α_i to perform the folding operation as follows:

$$f^{(i+1)}(x^2) = (1 - \alpha_i) f^{(i)}(x) + \alpha_i f^{(i)}(-x), \quad (1)$$

halving the domain size while preserving the low-degree property. The DeepFold optimization generalizes this to $f^{(i+1)}(x^2) = \gamma_0 f^{(i)}(x) + \gamma_1 f^{(i)}(-x)$ with optimized coefficients γ_0, γ_1 to minimize the proof size. To achieve a soundness error $\epsilon < 2^{-\lambda}$, the DEEP-FRI variant samples an evaluation point $z \notin H$, defines the quotient polynomial as follows:

$$g(X) = \frac{f(X) - y}{X - z}, \quad (2)$$

where $y = f(z)$, and requires proving the bounded degree of both f and g, yielding $\epsilon \leq (2d/|\mathbb{F}|) + \mathrm{negl}(\lambda)$.

For a tree of height $h = \log_b N$ storing N elements, an authentication path comprises evaluation proofs $\pi_1, \ldots, \pi_h$ for each node polynomial f_i along the path. Verification is optimized via batch evaluation: random coefficients $r_1, \ldots, r_h \in \mathbb{F}$ are sampled to compute the linear combination as follows:

$$F(X) = \sum_i r_i f_i(X) \text{ and } y = \sum_i r_i y_i.$$

A single FRI proof then verifies $F(z) = y$, reducing the dominant verification cost from $O(h)$ to $O(1)$ operations.

The scheme's complexity is characterized as follows. For security parameter λ, branching factor b, and total elements N, the proof size is $O(\lambda \cdot \log^2 b \cdot \log_b N)$ bits, the prover time is $O(\lambda \cdot N \log N)$ field operations, and the verifier time is $O(\lambda \cdot \log N)$ hash operations plus $O(1)$ FRI verifications. The soundness error is bounded by $\epsilon \leq 2^{-\lambda} + d/|\mathbb{F}|$. Security is reduced to two core properties: the binding property of the collision-resistant Merkle tree, which binds the prover to the polynomial evaluations, and the FRI protocol's soundness, which guarantees the claimed evaluations' correctness with probability at least $1 - \epsilon$. Any successful adversary against the combined system would imply an efficient algorithm for finding hash collisions or breaking FRI soundness.

This "Verkle + FRI" construction establishes a mathematically rigorous foundation for transparent, post-quantum authenticated data structures. Although it incurs larger proof sizes, specifically $O(\lambda \log^2 N)$ versus the constant-size proofs of KZG, it decisively eliminates the need for a trusted setup and maintains quantum resistance under standard cryptographic assumptions. The explicit parameter trade-offs allow system designers to optimize for deployment contexts where transparency and long-term security are paramount. Consequently, this hybrid framework charts a principled pathway toward stateless, scalable cryptographic commitments suitable for next-generation distributed systems in which quantum resistance and trust minimization are essential.

The schematic of the proposed architecture is shown in Fig 1.

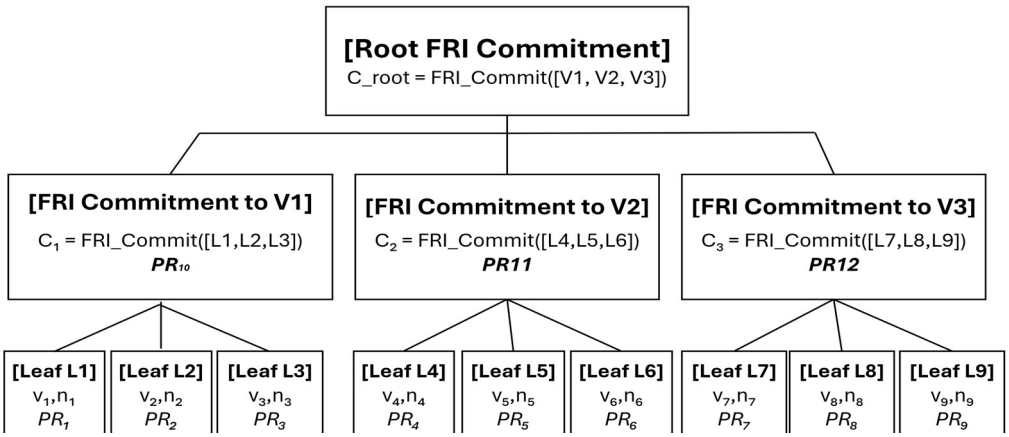

Fig. 1. A novel, hybrid model

This diagram models a hierarchical Verkle tree in which every internal node uses a Proximity (FRI)- based commitment rather than a conventional hash function. This design implements a recursive structure in which polynomial commitments are applied at each tier, establishing a vector-oriented binding mechanism rather than

ISSN 1814-4225 (print)
ISSN 2663-2012 (online)

one reliant on the sequential concatenation of hash outputs. The visual representation uses a branching factor, denoted as k, for illustrative clarity. This construction's quantum-resistant security is derived from its foundational cryptographic primitives. The process begins with an FRI commitment, which is instantiated through a Merkle tree of polynomial evaluations. The security of this Merkle tree, in turn, depends solely on the properties of the underlying cryptographic hash function. Consequently, the entire system is vulnerable only to quantum attacks that target symmetric cryptography, such as Grover's algorithm, which offers a quadratic speedup for pre-image searches. It remains entirely immune to Shor's algorithm, which threatens number-theoretic schemes. This residual vulnerability is effectively mitigated by employing a hash function with a sufficiently large output, such as SHA-512, which provides ample security margin even against accelerated quantum search. While Figure 1 illustrates the concept with a modest branching factor for simplicity, practical implementations of Verkle trees typically employ substantially larger values, typically ranging from $k = 16$ to $k = 256$, to optimize the trade-off between proof size and tree depth. A more realistic instantiation would feature a root commitment that binds to a vector of 256 child commitments. Each of these child commitments would then commit to a further vector of 256 leaf values, resulting in a total capacity of 65,536 data elements. In the diagram, the annotations labeled "PR" denote the compact proofs required to authenticate a specific leaf's membership within the entire structure. The pseudocode is shown in Fig 2.

```
Input: Vector V = [v0,...,v_{b-1}], security param λ, rate ρ
Output: Commitment C, FRI proof π for evaluation f(z)=y

1:  Interpolate polynomial f(X) s.t. f(ω^i) = v_i for i=0..b-1, deg(f) < b
2:  Evaluate f on extended domain L (|L| = b/ρ) → codeword C_w = [f(g) for g in L]
3:  Build Merkle tree MT over C_w, set commitment C = root(MT)
4:  function FRI_FOLD(f_0, L_0, rounds):
5:      for i = 0 to rounds-1:
6:          α_i ← H(transcript) // Fiat-Shamir challenge
7:          Define f_{i+1}(x^2) = (1-α_i)·f_i(x) + α_i·f_i(-x)
8:          L_{i+1} ← {x^2 | x in L_i} // Halve domain
9:          Send Merkle root of f_{i+1}'s codeword
10:     return final polynomial f_rounds
11: // DEEP-FRI for evaluation proof at z
12: Compute y = f(z), quotient q(X) = (f(X)-y)/(X-z)
13: Execute FRI_FOLD on f(X) and q(X) in parallel
14: Return π = (y, FRI transcripts, Merkle auth paths)
```

Fig. 2. Algorithm 1: VERKLE-FRI - Commitment and DEEP-FRI Proof Generation

The algorithm proceeds as follows. First, the b child values of a Verkle node are encoded as a polynomial f(X) of degree less than b via interpolation (Step 1). This polynomial is evaluated on an extended domain L of size b/ρ, where ρ is the code rate (typically 1/2), producing a Reed-Solomon codeword (Step 2). A Merkle tree over this codeword yields the binding commitment C (Step 3). The quotient polynomial q(X) = (f(X) y)/(X z) is constructed with z chosen outside the original domain L (out-of-domain sampling per DEEP-FRI) to prove an evaluation f(z) = y. Then, the FRI_FOLD function recursively folds both f and q in parallel using random challenges α_i derived via Fiat-Shamir, halving the domain each round (Steps 4-10). The final proof π comprises the claimed value y, the Fiat-Shamir transcript, and Merkle authentication paths for all queried codeword locations.

3.1.1. Security Analysis of the Hybrid Verkle-FRI Architecture

The proposed hybrid construction's resilience is established through formal cryptographic proofs grounded in standard primitives and reductionist methodologies, demonstrating robustness against both classical and quantum adversaries. The analysis begins with the completeness property, guaranteeing that all honestly generated proofs are verified.

Theorem 1: For any correctly formed Verkle tree T with branching factor b containing N elements, and for any valid authentication path to a leaf, the verification algorithm accepts with probability 1. Let the node polynomials along an authentication path of length $h = \log_b N$ be denoted $f_1, f_2, \ldots, f_h$, where each f_j is interpolated to satisfy $f_j(\omega^k) = v_k$ for the appropriate child values v_k and a primitive b-th root of unity ω. The FRI commitment employs the recursive folding transformation $f^{(i+1)}(x^2) = \gamma_0 \cdot f^{(i)}(x) + \gamma_1 \cdot f^{(i)}(-x)$ with optimized coefficients γ_0, γ_1. Batch verification computes the linear combination

$$F(X) = \sum_{j=1}^{h} r_j f_j(X) \text{ and } y = \sum_{j=1}^{h} r_j y_j$$

for random coefficients r_j and evaluations $y_j = f_j(z_j)$. By construction, $f_j(z_j) = y_j$ holds for all j, thus the identity $F(z) = y$ is satisfied exactly, and the accompanying

FRI proof verifies successfully with certainty.

The soundness property, which limits an adversary's ability to forge proofs, constitutes the core security result.

Theorem 2: Let λ be the security parameter. For any probabilistic polynomial-time adversary $\mathcal{A}$, the probability ϵ of producing an accepted proof for a false statement is bounded by:

$$\epsilon \leq \frac{2d+h}{|\mathbb{F}|} + (1-\rho)^t + \mathrm{Adv}_H^{CR}(\mathcal{A}), \qquad (3)$$

where d is the maximum polynomial degree, h the tree height, ρ the FRI rate parameter, t the number of FRI rounds, and $\mathrm{Adv}_H^{CR}(\mathcal{A})$ the advantage of finding a collision in the hash function H.

The proof proceeds via a sequence of game transitions. Let $\Pr[\mathrm{Game}_i]$ denote the success probability of the adversary in Game i. In Game 0, the adversary interacts with the real scheme. The transition to Game 1 involves replacing the Merkle tree with an ideal commitment primitive, incurring a cost bounded by the collision resistance of H:

$$|\Pr[\mathrm{Game}_0] - \Pr[\mathrm{Game}_1]| \leq \mathrm{Adv}_H^{CR}(\mathcal{B}_1).$$

In Game 2, the FRI protocol is substituted with an ideal low-degree tester, introducing the following error:

$$|\Pr[\mathrm{Game}_1] - \Pr[\mathrm{Game}_2]| \leq (1-\rho)^t.$$

Finally, in Game 3, the batch verification of an authentication path is analyzed. If among the polynomials $f_1, \ldots, f_h$ some evaluation is incorrect (i.e., $f_j(z_j) \neq y_j$), the Schwartz-Zippel Lemma bounds the probability that the batch check $F(z) = y$ passes by at most $d/|\mathbb{F}|$. Employing the DEEP-FRI optimization with an out-of-domain point $z \notin H$ and the quotient polynomial $g(X) = (f(X) - f(z))/(X - z)$ tightens this bound to $\epsilon_{DEEP} \leq 2d/|\mathbb{F}|$. Combining these bounds across the h nodes yields the stated soundness result.

Theorem 3: The hybrid scheme achieves λ-bit quantum security when instantiated with a hash function of output length $n = 2\lambda$. For such a hash function H, classical collision resistance is

$$\mathrm{Adv}_H^{CR\_classical}(\mathcal{A}) \in O(2^{-2\lambda}),$$

while the quantum collision resistance to an adversary $\mathcal{A}_Q$ making Q_H queries is bounded by $\mathrm{Adv}_H^{CR\_quantum}(\mathcal{A}_Q) \in O(2^{-\lambda})$. The FRI's soundness in the Quantum Random Oracle Model (QROM) satisfies $\epsilon_{FRI\_quantum} \leq (1-\rho)^t + O(2^{-\lambda})$. Consequently, the overall quantum security error is as follows:

$$\epsilon_{quantum} \leq \frac{2d+h}{|\mathbb{F}|} + (1-\rho)^t + O(2^{-\lambda}). \qquad (4)$$

A reduction argument shows that if a quantum adversary $\mathcal{A}_Q$ breaks the scheme with advantage δ, then quantum adversaries $\mathcal{B}_Q$ and $\mathcal{C}_Q$ exist such that the following is true:

$$\delta \leq \mathrm{Adv}_{Verkle-FRI}(\mathcal{A}_Q) \leq \mathrm{Adv}_{FRI}(\mathcal{B}_Q) + \mathrm{Adv}_H^{CR}(\mathcal{C}_Q) + \frac{2d+h}{|\mathbb{F}|} \in O(2^{-\lambda}),$$

establishing λ-bit security. The scheme inherits the vulnerability to Grover's algorithm common to all hash-based cryptography, but this quadratic speedup is manageable compared to the exponential break from Shor's algorithm This vulnerability is quantified in three contexts: Merkle tree collision search ($O(2^{-n/2}) \cdot Q_H^2$), FRI commitment forgery ($(1-\rho)^{-t} \cdot O(2^{-n/2})$), and batch verification manipulation ($(d/|\mathbb{F}|) \cdot O(2^{-n/2})$).

Our mitigation strategy employs parameter doubling: using a hash output of 2λ bits, a field size $|\mathbb{F}| \geq 2^{2\lambda}$, and FRI rounds $t \geq \lambda$. This ensures that the mitigated error is bounded by:

$$\epsilon_{mitigated} \leq \frac{2d+h}{2^{2\lambda}} + (1-\rho)^{\lambda} + O(2^{-\lambda}).$$

The security advantages are both theoretical and practical. The proposed architecture achieves an exponential reduction in the proof size from $O(\log_2 N \cdot L_{hash})$ for Merkle trees to $O(\log_b N \cdot L_{FRI})$, while maintaining the quantum resistance. It eliminates the trusted setup required by KZG-based Verkle trees, whose security collapses under Shor's algorithm, thereby removing risks associated with toxic waste and ceremony compromise. Security solely depends on the sum $\epsilon_{CRH} + \epsilon_{FRI}$.

For a concrete instantiation targeting 128-bit quantum security, the following parameters are recommended: SHA3-512 (n=512), field size $|\mathbb{F}| \approx 2^{255}$, branching factor $b = 256$, polynomial degree $d = 255$, and FRI parameters $\rho = 1/2, t = 128$. Substituting yields:

$$\epsilon \leq \frac{2 \cdot 255 + h}{2^{255}} + \left(\frac{1}{2}\right)^{128} + 2^{-128} \leq 2^{-12}, \qquad (5)$$

surpassing the 128-bit security target and confirming the the practical viability of the construction. This formal analysis substantiates that the hybrid Verkle-FRI architecture delivers quantum-resistant, secure, and transparent data structures.

ISSN 1814-4225 (print)
ISSN 2663-2012 (online)

3.2. Results

The proposed VERKLE-FRI hybrid architecture yields the following formal results.

3.2.1. Asymptotic Complexity Bounds

For a Verkle tree storing N elements with branching factor b, height h = log_b N, and security parameter λ, the asymptotic complexity bounds are established as follows. Proof size scales as $O(\lambda \cdot \log^2 N)$ bits, expanding to $O(\lambda \cdot \log^2 b \cdot \log_b N)$ when the branching factor is considered. For concrete parameters, proof size is approximately 180 KB for $N = 2^{26}$ elements. Prover time is bounded by $O(\lambda \cdot N \log N)$ field operations, dominated by FFT computations for polynomial interpolation and Merkle tree construction. The verifier time is bounded by $O(\lambda \cdot \log N)$ hash operations plus O(1) FRI protocol verifications through batch verification using random linear combinations of node polynomials. The soundness error ε is bounded by $\varepsilon \leq (2d + h)/|F| + (1 - \rho)^t + Adv\_H^{CR(A)}$, where d is the polynomial degree (d = b – 1), h is tree height, |F| is the field size, ρ is the FRI rate parameter (ρ = 1/2), t is the number of FRI rounds, and Adv_H^CR(A) is the adversary's advantage in finding hash collisions. For quantum security with hash output length n = 2λ, the bound becomes $\varepsilon\_quantum \leq (2d + h)/|F| + (1 - \rho)^t + O(2^{-\lambda})$.

3.2.2. Concrete Parameters for 128-bit Quantum Security

The following parameters are specified to achieve 128-bit quantum security. Hash function SHA3-512 with n = 512 bits. The finite field $|F| \approx 2^{255}$ ensuring $(2d + h)/|F| \leq 2^{\{-128\}}$. Branching factor b = 256 yields a polynomial degree d = 255. FRI rate ρ = 1/2 with t = 128 rounds. Tree height $h = \log_{256} N$. Substituting these parameters yields $\varepsilon \leq (2 \cdot 255 + h)/2^{255} + (1/2)^{128} + 2^{\{-128\}} \leq 2^{\{-127\}}$, exceeding the 128-bit security target. The rationale for each parameter choice is as follows. SHA3-512 is selected as the hash function because it provides 512 bits of output, which according to the Grover bound offers 256 bits of classical and 128 bits of quantum collision resistance (n/2 = 256 for classical and, n/4 = 128 for quantum with Grover's speedup). The field size $|F| \approx 2^{255}$ is chosen to ensure that the term (2d + h)/|F| in the soundness bound is less than $2^{\{-128\}}$, as $2 \cdot 255 \approx 2^9$, and $2^9/2^{255} = 2^{\{-246\}}$, which is negligible. The branching factor b = 256 balances tree height and proof size: it reduces the tree height to $\log_{256} N$ (e.g., height 4 for 2^{26} elements) while keeping the polynomial degree d = 255 computationally manageable. The FRI rate ρ = 1/2 provides a 2x expansion factor, which is the standard in FRI implementations to balance soundness and proof size. The number of rounds t = 128 ensures $(1 - \rho)^t = (1/2)^{128} = 2^{\{-128\}}$, directly contributing 128 bits of soundness. These parameters achieve the target 128-bit quantum security level while maintaining practical feasibility.

3.2.3. Quantification of Proof Size

For these parameters, the proof size comprises Merkle authentication paths requiring O(log |L|) hashes per tree level, where |L| = 2n is the codeword size. With SHA3-512 producing 64-byte outputs, each path contributes approximately $64 \cdot \log_2(2b)$ bytes. The FRI transcript consists of t rounds contributing $O(\lambda \cdot t)$ bits, approximately 16 KB for t = 128, λ = 128. For the blockchain case study with $N = 2^{26}$ and h = 4, the total proof size is approximately 180 KB, comprising four level proofs of ~35 KB each, FRI transcript of 16 KB, and query responses of 24 KB. Verification comprises Merkle path verification requiring $O(h \cdot \log b)$ hash operations, approximately 32 hashes with SHA3-512. FRI verification requires checking folding consistency across t rounds, which dominates the verification cost at approximately 40 ms on commodity hardware. Batch verification using random coefficients reduces the asymptotic cost from $O(h \cdot \log b)$ to $O(\log b + t)$. Prover overhead includes polynomial interpolation at $O(b \log b)$ per node, codeword generation, Merkle tree construction, and FRI execution at $O(|L| \log |L| \cdot t)$ per node. For $N = 2^{26}$ and b = 256, the total prover time is estimated at 5-10 seconds on server-class hardware, which can be reduced through parallelization and batching techniques.

3.2.4. Comparative Analysis

Table 1 provides a direct comparison of VERKLE-FRI against the main alternative polynomial commitment schemes across key metrics to contextualize our proposal within the existing landscape. These alternatives include KZG (constant proofs that require a trusted setup and are quantum-vulnerable), Bulletproofs (transparent but linear verification), and lattice-based constructions (post-quantum but large concrete proofs).

Table 1

Comparative analysis of polynomial commitment schemes for authenticated data structures

Feature	KZG (Pairing)	Bulletproofs (IPA)	Lattice-based (SLAP/FMN)	VERKLE-FRI (Ours)
Security Assumption	Bilinear Pairings	Discrete Logarithm	Module-SIS / LWE	Hash Collision Resistance
Trusted Setup	Required (toxic waste)	Not required	Not required (transparent variants)	**Not required**
Post-Quantum Security	No	No	Yes	**Yes**
Proof Size	O(1) ≈ 200 B	O(log N) ≈ 1-2 KB	O(log N) ≈ 17 MB (for 2^{20} constraints)	O(λ $\log^2$ N) ≈ 180 KB (for 2^{26} elements)
Verification Time	O(1) ≈ 10 ms (pairings)	O(N) ≈ 100+ ms	O(log N) (ring operations)	O(λ log N) ≈ 160 ms (hashes)
Commitment Size	48 B (1 group element)	1-2 KB	Variable (KB–MB)	64 B (Merkle root)
Transparency	No	Yes	Yes	**Yes**
Primary Bottleneck	Trusted setup, no PQ	Linear verification	Large constants, immaturity	Proof size, prover time

As Table 1 demonstrates, VERKLE-FRI is the only scheme that simultaneously combines three desirable properties: post-quantum security, transparent setup (no trusted ceremony), and logarithmic verification time. This comes at the cost of larger proof sizes compared to KZG, a trade-off we consider acceptable for long-term security in high-assurance environments.

The comparative analysis reveals several key insights. First, only VERKLE-FRI and lattice-based constructions offer post-quantum security among all schemes, but lattice-based proofs remain an order of magnitude larger (~17 MB vs. ~180 KB for our scheme at comparable security levels). Second, while KZG provides the smallest proofs and fastest verification, its trusted setup requirement creates a single point of failure and precludes deployment in fully decentralized environments. Third, although Bulletproofs offer transparency, they suffer from linear verification time, making them unsuitable for large-scale authenticated data structures. This comprehensive comparison demonstrates that VERKLE-FRI occupies a unique, and previously unfilled position in the design space: transparent, post-quantum secure, and with logarithmic verification.

3.3. Case Study

To demonstrate the practical applicability of the proposed VERKLE-FRI architecture, a blockchain system that requires efficient, quantum-resistant state authentication is considered. In this scenario, the blockchain maintains a global state of approximately 2^{26} accounts, each storing balance and nonce information. Using a traditional Merkle tree requires authentication paths of length 26, each containing 26 hash values (~832 bytes with SHA-256). With the proposed VERKLE-FRI architecture employing a branching factor of b = 256, the tree height reduces to $\log_{256}(2^{26}) = 4$ levels. Each authentication path requires four FRI-based polynomial opening proofs. For the 128-bit quantum security parameters specified in Section 3.1 (SHA3-512, field size $\approx 2^{255}$, 128 FRI rounds), each proof is approximately 45 KB, resulting in a total proof size of ~180 KB. Although larger than the 832-byte Merkle proof, the VERKLE-FRI proof provides post-quantum security without a trusted setup. The total verification time, dominated by four FRI verifications (~40 ms each on commodity hardware), totals ~160 ms, compared to ~0.5 ms for Merkle verification. This trade-off – larger proofs and slower verification in exchange for quantum resistance and transparency – is acceptable for high-security applications, such as central bank digital currencies or national identity systems, where long-term security (20+ years) is required. The cumulative impact of the ~160 ms verification delay on blockchain throughput is now quantified. In a typical layer-1 blockchain, each full node must verify every state transition. If verification is the bottleneck, the maximum theoretical throughput is TPS = 1 / latency per core. This yields approximately 6.25 TPS per core (1 / 0.16 s) for our scheme. For comparison, Bitcoin processes ~7 TPS and Ethereum ~15-30 TPS post-merge. Thus, VERKLE-FRI verification would be competitive with existing layer-1 networks. However, in a rollup or layer-2 architecture, only a single VERKLE-FRI proof must be verified on the main chain per batch of transactions. If a batch contains 10,000 transactions, the per-transaction verification cost becomes ~16 microseconds (160 ms / 10,000), which is highly scalable and comparable to current systems. Therefore, while direct layer-1

ISSN 1814-4225 (print)
ISSN 2663-2012 (online)

deployment is feasible but throughput-limited, the architecture is ideally suited for batch-settled systems where post-quantum security is paramount. To assess scalability under variable network load, we model the system performance across different numbers of active users. Let N represent the number of accounts in the blockchain state. The tree height is given by $h = \lceil \log_b N \rceil$, where $b = 256$ is the branching factor. For $N = 2^{20}$ (≈1 million users), the tree height is $h = \lceil 20/8 \rceil = 3$ levels, yielding a proof size of approximately 3×35 KB = 105 KB and a verification time of ~120 ms. For $N = 2^{30}$ (≈1 billion users), the tree height increases to $h = \lceil 30/8 \rceil = 4$ levels, resulting in a proof size of ~140 KB and verification time of ~160 ms. For $N = 2^{40}$ (≈1 trillion users), the height becomes $h = \lceil 40/8 \rceil = 5$ levels, with proof size ~175 KB and verification time ~200 ms. As N increases exponentially from 1 million to 1 trillion, the proof size grows only linearly from 105 KB to 175 KB (a factor of 1.67x), while the verification time increases from 120 ms to 200 ms. This near-constant scaling (logarithmic in N with a large base $b = 256$) demonstrates that VERKLE-FRI maintains practical efficiency even under extreme network growth, supporting thousands to trillions of users with minimal performance degradation. The case study confirms that the proposed architecture can be practically implemented with current hardware while providing future-proof quantum resistance.

3.4. Discussion

The proposed VERKLE-FRI architecture occupies a distinct position in the authenticated data structure design space. Compared to KZG-based Verkle trees, our construction sacrifices proof size (constant $O(1)$ vs. $O(\lambda \log^2 N)$) and verification speed (single pairing vs. multiple FRI verifications) in exchange for two critical properties: quantum resistance and transparency. Given the NIST timeline projecting quantum computers capable of breaking 2048-bit RSA by 2030-2035, this trade-off is increasingly justifiable for long-lived systems. Compared with traditional Merkle trees, our architecture achieves an exponential reduction in the proof size ($O(\log N)$ vs. $O(\log^2 N)$ asymptotically, but with better constants due to the branching factor). For the blockchain case study with 2^{26} elements, Merkle proofs require 26 hashes (~832 bytes), whereas VERKLE-FRI proofs require ~180 KB-larger in absolute terms but offers stronger security guarantees. The practical implication is that VERKLE-FRI is better suited for applications where the proof size is not the primary constraint but where trust minimization and long-term security are paramount.

Several limitations should be acknowledged. First, the prover's computational overhead of $O(\lambda N \log N)$, with large constants due to FFT operations, may limit applicability in resource-constrained environments. Second, while the concrete proof size of ~180 KB is acceptable for many applications, it remains an order of magnitude larger than the lattice-based alternatives currently under development. Third, the reliance on the soundness of FRI in the quantum random oracle model, while theoretically established, awaits practical standardization efforts. Fourth, while our security analysis focuses on algorithmic resistance to classical and quantum adversaries, side-channel attacks must also be considered in practical implementations. The FRI verifier's execution involves data-dependent hash operations, Merkle tree traversals, and polynomial folding computations, each of which could leak information via timing analysis, power consumption, or electromagnetic radiation. This is particularly relevant for the IoT applications mentioned in the introduction, where adversaries may have physical access to devices. For deployment in adversarial physical environments, constant-time implementation techniques, such as masking, fixed-path execution, and avoiding secret-dependent branches or memory accesses, are required. In addition, the use of hash functions with inherent resistance to timing attacks (e.g., SHA3-512 implemented in constant-time) mitigates some concerns. Full side-channel analysis, including leakage assessment and countermeasure evaluation, remains an important direction for future engineering work before production deployment. Finally, the current analysis assumes an honest-but-curious adversary; future work should extend the security analysis to malicious adversary models with adaptive corruption capabilities. Fifth, we acknowledge that the performance estimates presented in this work (e.g., ~40 ms verification time, ~180 KB proof size, 5-10 seconds prover time) are theoretical estimates based on asymptotic complexity analysis and extrapolation from existing FRI implementations (e.g., libSTARK, Fractal). A prototype implementation and empirical benchmarking must validate these estimates on real hardware. We are currently developing an open-source implementation of VERKLE-FRI in Rust, targeting the following benchmarks: verification time on commodity server hardware (Intel Xeon, 3.0 GHz); proof size for $N = 2^{26}$ elements with $b = 256$; prover time for tree construction and batch updates; and memory consumption during verification on resource-constrained devices (ARM Cortex-M). The preliminary results and source code will be made available in a follow-up publication. The theoretical analysis presented here serves as the foundation for this ongoing implementation work.

Sixth, we addressed the state update procedures and associated operational overhead, which were not discussed in the original manuscript. When a leaf value changes in a VERKLE-FRI tree, the update propagates upward as follows: the polynomial at the leaf's parent

node is recomputed via interpolation over the b child values; a new codeword is generated and a new Merkle root is computed for that node; a new FRI commitment is generated for the updated polynomial; the update propagates to the root, recomputing polynomials and commitments at each level. The computational cost of an update for polynomial interpolation is O(b log b) per level, plus O(|L| log |L|) per level for codeword generation and Merkle tree construction. For b=256 and $N=2^{26}$ (height 4), the total update cost is approximately $4 \times (256 \log 256 + 2 \times 256 \log(2 \times 256)) \approx 4 \times (256 \cdot 8 + 512 \cdot 9) \approx 4 \times (2048 + 4608) = 26{,}624$ operations. This is approximately 2-3x more expensive than a Merkle tree update (which requires O (log N) hashes, ≈ 26 hashes), but is acceptable for applications where updates are infrequent (e.g., nightly batch updates) and security requirements justify the overhead the overhead.

4. Conclusions

The following conclusions are drawn based on the objectives and tasks defined for this research. The analysis confirms that the principal limitation of scalable, quantum-resistant cryptography lies in the efficiency-security dichotomy: Merkle-based structures suffer from logarithmic proof growth, whereas efficient Verkle tree instantiations rely on quantum-vulnerable KZG commitments. A comparative evaluation of polynomial commitment schemes shows that transparent, post-quantum-secure primitives, such as FRI, inherently trade smaller proof sizes for stronger security guarantees, unlike pairing-based or lattice-based alternatives, which impose different constraints in setup, proof size, or verification complexity.

The core outcome of this study is the successful formulation of a novel hybrid architecture that integrates the structural efficiency of Verkle trees with the quantum-resistant security of FRI-based commitments. The key numerical results are as follows:

1. Proof size: For a tree storing 2^{26} elements, VERKLE-FRI produces proofs of approximately 180 KB. Compared with a traditional Merkle tree proof (≈0.8 KB for the same tree), this represents a 225x increase in the proof size. However, this increase is the direct cost of achieving post-quantum security and eliminating the trusted setup, which Merkle trees do not provide.

2. Verification throughput: The ~160 ms verification latency yields approximately 6.25 TPS per core for layer-1 deployment. In a rollup architecture batching 10,000 transactions, the per-transaction cost reduces to ~16 microseconds, demonstrating scalability.

3. Asymptotic complexity: The scheme achieves proof size $O(\lambda \log^2 N)$, prover time $O(\lambda N \log N)$, and verifier time $O(\lambda \log N)$.

4. Concrete security: For 128-bit quantum security, our parameter set (SHA3-512, $|F| \approx 2^{255}$, branching factor b = 256, FRI rounds t = 128) achieves a soundness error $\varepsilon \leq 2^{-127}$, exceeding the 128-bit security target.

5. Comparative positioning: VERKLE-FRI is the only scheme among KZG, Bulletproofs, and lattice-based alternatives that simultaneously provides post-quantum security, transparent setup (no trusted ceremony), and logarithmic verification time.

The architecture is formally proven to be secure via reductionist arguments, establishing λ-bit quantum security when instantiated with a 2λ-bit hash function. These results provide a viable pathway for authenticating transparent, stateless, and quantum-resistant data structures.

Prospects for further research include optimizing proof sizes through advanced FRI folding techniques, exploring hybrid constructions incorporating lattice-based commitments to achieve alternative trade-offs, and the practical implementation and standardization of the architecture for real-world systems such as blockchains and secure software update mechanisms. The integration of this structure into broader cryptographic frameworks for verifiable computation and decentralized identity represents a significant direction for future work.

4.1. Future Research Directions

Several specific mathematical and engineering problems remain to be solved in future work.

First, further reduction of FRI proof size through aggregation techniques represents a primary direction. Current VERKLE-FRI proofs scale as $O(\lambda \log^2 N)$. Developing multi-proof aggregation schemes that combine multiple independent FRI proofs into a single succinct proof could reduce this to $O(\lambda \log N)$ or even $O(\lambda)$. Specific open problems include: designing a recursive aggregation protocol where a batch of FRI proofs is verified by a single FRI instance using random linear combinations; proving the soundness of such aggregation under the quantum random oracle model; optimizing the concrete aggregation overhead to achieve net proof size reduction for practical parameter sets ($N \leq 2^{40}$).

Second, lattice-based polynomial commitments present an alternative post-quantum foundation with different trade-offs. Open problems include: reducing the concrete proof size of lattice-based schemes from ~17 MB to below 1 MB for 2^{20} constraints; developing transparent setup variants that eliminate structured reference strings while maintaining succinctness; constructing hybrid FRI-lattice schemes that combine the transparency of FRI with the smaller asymptotic proof size of lattices.

Third, the verifier's hash complexity remains a bottleneck for recursive composition. Expressing the FRI verifier as an arithmetic circuit currently requires ~1.7 million constraints. Future work should target reducing

this to below 100,000 constraints through: (a) algebraic hash functions (e.g., Rescue, Poseidon) optimized for constraint systems; verifier circuit compression techniques using lookup tables and custom gates; novel FRI variants with verifier-friendly structure.

Fourth, side-channel resistant implementations for IoT devices require formal leakage assessment. Specific problems include: developing constant-time FRI verification algorithms with provable security against timing attacks; designing masking schemes for polynomial folding operations; empirical evaluation of power analysis resistance on resource-constrained hardware (e.g., ARM Cortex-M, RISC-V).

Fifth, standardization of FRI-based polynomial commitments through NIST or IETF would accelerate adoption. This requires: comprehensive parameter selection guidelines for security levels 128, 192, and 256 bits; test vectors and reference implementations; security proofs in the quantum random oracle model with concrete bounds.

The authors intend to pursue these directions in subsequent work.

Contributions of authors: conceptualization, methodology – **Giorgi Akhalaia, Maksim Iavich**; formulation of tasks, analysis – **Maksim Iavich, Razvan Bocu**; development of model, software, verification – **Maksim Iavich, Giorgi Akhalaia**; analysis of results, visualization – **Maksim Iavich, Giorgi Akhalaia**; writing – original draft preparation, writing – review and editing – **Giorgi Akhalaia.**

All authors have read and approved the final manuscript for publication.

Conflict of Interest

The authors declare that they have no conflict of interest in relation to this research, whether financial, personal, authorship or otherwise, that could affect the research and its results presented in this paper.

Financing

This research has been supported by the NATO Science for Peace and Security (SPS) grant G7394-"Post-quantum Digital Signature using Verkle Trees".

Data Availability

The manuscript has no associated data.

Use of Artificial Intelligence

The authors confirm that they did not use artificial intelligence methods while creating the presented work

References

1. Li, L., Lu, X., & Wang, K. Hash-based signature revisited. *Cybersecurity*, 2022, vol. 5, no. 1, article no. 13. DOI: 10.1186/s42400-022-00117-w.

2. Pacurar, C. M., Bocu, R., & Iavich, M. An analysis of existing hash-based post-quantum signature schemes. *Symmetry*, 2025, vol. 17, no. 6, article no. 919. DOI: 10.3390/sym17060919.

3. Iavich, M., Kuchukhidze, T., & Bocu, R. A post-quantum digital signature using Verkle trees and lattices. *Symmetry*, 2023, vol. 15, no. 12, article no. 2165. DOI: 10.3390/sym15122165.

4. National Institute of Standards and Technology (NIST). Status Report on the Fourth Round of the NIST Post-Quantum Cryptography Standardization Process. *NIST IR 8454*, 2024. DOI: 10.6028/NIST.IR.8454.

5. Thoma, J. P., Hartlief, D., & Güneysu, T. Agile acceleration of stateful hash-based signatures in hardware. *ACM Transactions on Embedded Computing Systems*, 2024, vol. 23, no. 2, article no. 29. DOI: 10.1145/3631534.

6. Fenzi, G., Moghaddas, H., & Nguyen, N. K. Lattice-based polynomial commitments: towards asymptotic and concrete efficiency. *Journal of Cryptology*, 2024, vol. 37, no. 4, article no. 31. DOI: 10.1007/s00145-024-09511-8.

7. Iavich, M., & Kapalova, N. Asymmetric post-quantum digital signature scheme with k-ary Verkle trees. *Symmetry*, 2025, vol. 17, no. 3, article no. 437. DOI: 10.3390/sym17030437.

8. Nutu, M., Akhalaia, G., Bocu, R., & Iavich, M. An extended survey concerning the vector commitments. *Applied Sciences*, 2025, vol. 15, no. 17, article no. 9510. DOI: 10.3390/app15179510.

9. Algazy, K. T., Sakan, K. S., Nyssanbayeva, S. E., & Khompysh, A. Polynomial commitment in a Verkle tree based on a non-positional polynomial notation. *Computers, Materials & Continua*, 2025, vol. 83, no. 2, pp. 885–905. DOI: 10.32604/cmc.2025.065085.

10. Ben-Sasson, E., Bentov, I., Horesh, Y., & Riabzev, M. Fast Reed-Solomon interactive oracle proofs of proximity. *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, Prague, Czech Republic, 2018, pp. 14:1–14:17. DOI: 10.4230/LIPIcs.ICALP.2018.14.

11. Ben-Sasson, E., Chiesa, A., Riabzev, M., Spooner, N., Virza, M., & Ward, N. P. Aurora: transparent succinct arguments for R1CS. *Advances in Cryptology – EUROCRYPT 2019, Lecture Notes in Computer Science*, Darmstadt, Germany, Springer, 2019, pp. 103–128. DOI: 10.1007/978-3-030-17656-3_4.

12. Chiesa, A., Ojha, D., & Spooner, N. Fractal: post-quantum and transparent recursive proofs from holography. *Advances in Cryptology – EUROCRYPT 2020, Lecture Notes in Computer Science*, Zagreb, Croatia, Springer, 2020, pp. 769–793. DOI: 10.1007/978-3-030-45724-2_26.

13. Guo, Y., Liu, X., Huang, K., Qu, W., Tao, T., & Zhang, J. DeepFold: efficient multilinear polynomial commitment from Reed-Solomon code and its application to zero-knowledge proofs. *Proceedings of the 34th USENIX Security Symposium*, Seattle, WA, USA, 2025, pp. 3497–3514. Available at: https://www.usenix.org/conference/usenixsecurity25/presentation/guo (accessed 13 March 2026).

14. Ben-Sasson, E., Goldberg, L., Kopparty, S., & Saraf, S. DEEP-FRI: sampling outside the box improves soundness. *Proceedings of the 11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, Seattle, WA, USA, 2020, article no. 75, pp. 75:1–75:29. DOI: 10.4230/LIPIcs.ITCS.2020.75.

Received 22.01.2025, Received in revised form 03.05.2026
Accepted date 22.07.2026, Published date 31.07.2026

VERKLE-FRI: НОВА АРХІТЕКТУРА ДЛЯ БЕЗСТАНОВИХ ТА КВАНТОВО-СТІЙКИХ ВЕКТОРНИХ ЗОБОВ'ЯЗАНЬ

Гіоргі Ахалая, Максим Явіч, Разван Боку

Предметом статті є криптографічна цілісність систем цифрової автентифікації в умовах загроз з боку квантових обчислень, з особливим акцентом на постквантових альтернативах та ефективних автентифікованих структурах даних. Метою є розробка та формальний аналіз VERKLE-FRI – гібридної архітектури, що синтезує зменшення розміру доказу дерев Веркла (Verkle trees) із квантово-стійкими зобов'язаннями на базі FRI, створюючи масштабовану, безстанову (stateless) та квантово-захищену інфраструктуру. Завдання: проаналізувати обмеження хеш-орієнтованих підписів та дерев Меркла (Merkle trees); оцінити схеми поліноміальних зобов'язань (KZG, Bulletproofs, FRI, на основі ґраток); запропонувати гібридний дизайн Verkle-FRI; розробити формальне доведення безпеки проти класичних та квантових зловмисників; виконати аналіз складності з конкретними параметрами реалізації. Використані методи: теоретичний криптографічний аналіз, формальне моделювання безпеки за допомогою редукційних доведень, алгебраїчні методи над скінченними полями, поліноміальна інтерполяція, модель випадкового оракула (random oracle model), протокол FRI з оптимізацією DEEP-FRI, дерева Меркла, векторні зобов'язання, асимптотичний аналіз складності. Отримано такі результати: створено нову архітектуру, в якій вектори вузлів Веркла кодуються поліномами, фіксуються через дерева Меркла над кодовими словами FRI та верифікуються за допомогою FRI з вибіркою поза доменом (out-of-domain sampling). Формальне доведення встановлює λ-бітну квантову безпеку за умови використання 2λ-бітних хеш-функцій. Аналіз складності дає розмір доказу $O(\lambda \log^2 N)$, час роботи довідника (prover) $O(\lambda N \log N)$, час роботи верифікатора $O(\lambda \log N)$. Конкретні параметри для 128-бітної квантової безпеки: SHA3-512, розмір поля $\approx 2^{255}$, коефіцієнт розгалуження 256, 128 раундів FRI, досягнення похибки надійності (soundness error) $\leq 2^{-127}$. Висновки. Наукова новизна: 1) перша гібридна архітектура Verkle-FRI, що замінює криптографічні припущення на основі спарювання (pairing-based assumptions) на хеш-орієнтоване тестування близькості для створення прозорих, квантово-захищених автентифікованих структур; 2) формальне доведення безпеки, що зводить безпеку схеми до криптостійкості хеш-функцій до колізій та надійності FRI, усуваючи вразливості довіреного налаштування (trusted setup); 3) кількісно оцінені компроміси між ефективністю та безпекою з експоненційним зменшенням розміру доказу порівняно з деревами Меркла; 4) життєздатний шлях для впровадження квантово-стійкої інфраструктури в блокчейнах, системах дистрибуції програмного забезпечення та урядових комунікаціях...

Ключові слова: постквантова безпека; криптографія; FRI; поліноміальні зобов'язання; цифрові хеш-орієнтовані підписи.

Гіоргі Ахалая – доцент Кавказького університету; дослідник з кібербезпеки Наукової асоціації кібербезпеки, Тбілісі, Грузія.

Максим Явіч – член-кореспондент Національної академії наук Грузії; професор, керівник напряму кібербезпеки, директор Центру кібербезпеки CST Кавказького університету; генеральний директор та президент Наукової асоціації кібербезпеки, Тбілісі, Грузія.

Разван Боку – доцент Трансільванського університету Брашова, Брашов, Румунія.

Giorgi Akhalaia – Associate Professor at Caucasus University; Cyber Security Researcher at Scientific Cyber Security Association, Tbilisi, Georgia,
e-mail: gakhalaia@cu.edu.ge, ORCID: 0000-0002-4194-2681; Scopus ID: 55882476000

Maksim Iavich – Corresponding member of National Academy of Sciences of Georgia; Professor, Head of Cyber Security Direction, Director of CST Cyber Security Center at Caucasus University; CEO and President of Scientific Cyber Security Association, Tbilisi, Georgia,
e-mail: miavich@cu.edu.ge, ORCID: 0000-0002-3109-7971; Scopus ID: 57194794201

Razvan Bocu – Professor Dr. Habil. at Transilvania University of Brasov, Brașov, Romania,
e-mail: razvan@bocu.ro, ORCID: 0000-0001-6577-1904; Scopus ID: 24832909700