

Received	2026/01/02	تم استلام الورقة العلمية في
Accepted	2026/01/22	تم قبول الورقة العلمية في
Published	2026/01/23	تم نشر الورقة العلمية في

Comparative Analysis of Zero-Knowledge Proof-of-Stake Systems in Verifiable Cryptography and Their Applications in Improving the Efficiency of Decentralized Storage and Blockchains

Nuha Omran Abokhdair

Ali Alissawi Ahmed AlQudairi

abo.khdeir@zu.edu.ly

Alialgader53@gmail.com

Computer Science Department, Faculty of Science,
University of Zawia, Libya

Abstract

Decentralized storage networks increasingly rely on blockchain-based verification to ensure data integrity without centralized control; however, proof-intensive workloads introduce significant latency and on-chain cost overhead. This paper presents a lifecycle-based comparative analysis of major zero-knowledge proof (ZKP) models used in decentralized storage, focusing on zk-SNARK frameworks and transparent zk-STARK constructions. A multi-layer evaluation framework is introduced, aligning performance analysis with the core stages of the proof lifecycle: generation, aggregation, and on-chain verification. Building on this analysis, the paper proposes a hybrid architecture that combines parallel STARK-based proof generation with recursive SNARK-based compression, reducing on-chain verification complexity to near-constant. A Filecoin-inspired case study, supported by a quasi-empirical performance model, demonstrates that the proposed hybrid approach significantly reduces verification latency and data overhead while mitigating the linear growth of verification costs. The results indicate that hybrid ZKP architectures offer a scalable and economically viable solution for decentralized storage systems and large-scale blockchain networks.

Keywords: Zero-knowledge proofs, zk-SNARKs, zk-STARKs, recursive aggregation, decentralized storage, verifiable cryptography, scalability, gas cost.

تحليل مقارن لأنظمة إثبات الحصة ذات المعرفة الصفيرية في التشفير القابل للتحقق وتطبيقاتها في تحسين كفاءة التخزين اللامركزي وتقنية سلاسل الكتل

علي العيساوي احمد القديري

الدكتورة: نهى عمران أبوخدير

Alialgader53@gmail.com

abo.khdeir@zu.edu.ly

قسم الحاسوب - كلية العلوم - جامعة الزاوية - ليبيا

الملخص

تعتمد شبكات التخزين اللامركزية بشكل متزايد على التحقق القائم على تقنية البلوك تشين لضمان سلامة البيانات دون تحكم مركزي؛ إلا أن أحمال العمل التي تتطلب إثباتات مكثفة تؤدي إلى زيادة كبيرة في زمن الاستجابة وتكاليف المعالجة على البلوك تشين. تقدم هذه الورقة البحثية تحليلاً مقارناً قائماً على دورة حياة نماذج إثبات المعرفة الصفيرية (ZKP) الرئيسية المستخدمة في التخزين اللامركزي، مع التركيز على أطر عمل zk-SNARK وبنيات zk-STARK الشفافة، كما تقدم إطار عمل تقييم متعدد الطبقات، يربط تحليل الأداء بالمراحل الأساسية لدورة حياة الإثبات: التوليد، والتجميع، والتحقق على البلوك تشين، وانطلاقاً من هذا التحليل، نقترح الورقة بنية هجينة تجمع بين توليد الإثباتات المتوازي القائم على STARK والضغط المتكرر القائم على SNARK، مما يقلل من تعقيد التحقق على البلوك تشين إلى مستوى شبه ثابت. وتوضح دراسة حالة مُستوحاة من Filecoin، مدعومة بنموذج أداء شبه تجريبي، أن النهج الهجين المقترح يقلل بشكل كبير من زمن استجابة التحقق وتكاليف البيانات، مع الحد من النمو الخطي لتكاليف التحقق، تشير النتائج إلى أن بنى إثبات المعرفة الصفيرية الهجينة توفر حلاً قابلاً للتوسع ومجدياً اقتصادياً لأنظمة التخزين اللامركزية وشبكات البلوك تشين واسعة النطاق.

الكلمات المفتاحية: إثباتات المعرفة الصفيرية، zk-SNARKs، zk-STARKs، التجميع التكراري، التخزين اللامركزي، التشفير القابل للتحقق، قابلية التوسع، تكلفة الغاز.

Introduction

Privacy, scalability, and verifiability are essential elements for sustainable blockchain infrastructure. As decentralized systems evolve from simple value transfer to complex applications such as decentralized identity and storage, the tension between transparency (verifiability) and privacy (data confidentiality) becomes increasingly apparent. In decentralized storage networks, the integrity and availability of stored data must be verifiable over time. Therefore, systems rely on cryptographic proofs, such as storage integrity proofs, to demonstrate continued data possession without disclosure.

Despite these mechanisms, proof-intensive workloads present a significant challenge, requiring the generation and periodic verification of large numbers of proofs. This places pressure on the blockchain's verification layers and increases gas costs and latency. Zero-knowledge proofs (ZKPs) offer a promising approach to mitigating these challenges by enabling concise verification of claims without disclosing sensitive information. However, the main ZKP models involve inherent trade-offs. zk-SNARKs provide concise proofs and low verification costs, but common architectures rely on trusted setup and pairing-based assumptions, raising concerns about setup integrity and resilience against quantum computing [1,2]. In contrast, zk-STARKs eliminate trusted setup and offer transparency with robust properties against quantum computing, but they typically produce larger proofs and incur higher verification costs, often requiring off-chain verification [2, 3]. Current studies analyze these models separately and, to a large extent, do not provide a unified analytical framework specifically designed for the operational requirements of decentralized storage systems.

This paper's novelty compared to existing surveys lies in four practical contributions:

- (a) A ZKP lifecycle-compatible evaluation framework that correlates ZKP performance with the decentralized storage proof lifecycle (generation $\rightarrow$ aggregation $\rightarrow$ on-chain verification).
- (b) A formal hybrid architecture that combines transparent STARK proofs with recursive SNARK compression, including explicit complexity and size transformations.

- (c) A quasi-empirical analytical model that estimates verification time and on-chain data load under variable proof batch sizes.
- (d) A Filecoin-inspired case study that implements the framework and demonstrates how hybrid models reduce blockchain verification load.

Problem Definition and Cryptographic Challenges:

Traditional blockchain verification requires nodes to re-execute transactions and verify state transitions, which becomes increasingly costly as throughput increases. Zero-Knowledge Proofs (ZKPs) reduce this cost by allowing the prover to generate a cryptographic proof that the computation was performed correctly, while verifiers efficiently verify the proof. However, ZKPs face three main challenges in decentralized storage environments:

- High prover cost (generation burden) due to expensive polynomial operations and constraint systems.
- Trusted setup risks in many efficient zk-SNARKs.
- Limited scalability in recurring validations, as verifying each proof individually leads to linear growth in the verification process.

Research Question: How can the efficiency of ZKPs be improved to reduce the burden of proof generation and verification while maintaining privacy and decentralization in blockchain and decentralized storage environments?

Related Work and Selected Studies:

Recent research in zero-knowledge proofs (ZKPs) has seen rapid development aimed at addressing scalability challenges and improving the efficiency of decentralized systems. Current research focuses heavily on reducing proof size and lowering the computational burden associated with their generation and verification within blockchain environments. Given the stringent requirements of decentralized storage networks, this section presents a critical analysis of key representative studies that have addressed recursive zk-SNARKs and transparent zk-STARKs. This review aims to assess the effectiveness of these solutions in reducing proof processing time and improving verification efficiency, thus paving the way for understanding current limitations and proposed solutions for achieving sustainability in large-scale systems. In this

context, we discuss three key research papers that highlight these technical issues:

Study 1: Liu et al. (2025) – The GENES Protocol

Liu Jiayi, Guo Li, and Kang Tianyu introduced GENES, an innovative recursive zk-SNARK protocol designed to enhance blockchain scalability. The methodology utilizes an efficient framework to aggregate multiple R1CS (Rank-1 Constraint System) instances into a single, succinct verifiable proof with near-constant verification complexity. Experimental results demonstrated significant reductions in both prover and verifier time compared to traditional frameworks. Although this efficiency comes at the cost of a slightly larger proof size, the GENES protocol is highly effective for layered architectures aiming to minimize the computational load on nodes and improve on-chain confirmation speeds. [4]

Study 2: Zhang et al. (2024) – Secure Transactions in Distributed Computing

Zhang and colleagues proposed a blockchain-ZKP integrated method to secure data transactions in distributed computing environments. Their framework combines smart contracts with zero-knowledge proofs to mitigate data disclosure risks and enhance verification efficiency during exchange. By accelerating the data preparation phase and implementing effective workload decomposition, the study reported a substantial reduction in proof generation time. These findings emphasize the practical feasibility of ZKPs in decentralized storage and data-trading scenarios, where reducing response time and on-chain verification costs is critical for transaction fairness and reliability [5].

Study 3: Yuan (2025) – Decentralized Identity and Scalable Data Sharing

Yuan Hui developed a scalable, privacy-preserving framework for decentralized identity and verifiable data sharing utilizing zk-STARKs. The study highlights the qualitative advantages of STARK-based systems, specifically their transparency (eliminating the need for a trusted setup) and their inherent resistance to quantum attacks. Through benchmarking experiments, the research demonstrated improved prover efficiency and robust security properties. However, it also acknowledged a critical technical trade-off: zk-STARKs often result in larger proof sizes and higher

verification times compared to SNARK counterparts in certain scenarios, underscoring the need for optimized or hybrid implementations in resource-constrained environments [6].

Limitations of Comparison in Previous Studies:

Although the aforementioned studies are valuable, their comparisons are often not aligned with decentralized storage proof lifecycles. In particular, evaluations are frequently reported under heterogeneous assumptions (e.g., varying hardware, security parameters, and implementation settings), and they seldom integrate proof generation, aggregation frequency, and on-chain verification cost within a single, consistent framework. Moreover, storage-specific constraints such as recurring epoch-based proofs, verification multiplicity, and call data-driven gas costs are not always explicitly modeled. These limitations motivate the lifecycle-driven comparative framework and the parameterized evaluation adopted in this paper.

Methodology (Lifecycle-Based Comparative Framework):

This study employs a lifecycle-based comparative methodology to evaluate how decentralized storage proofs are generated, aggregated, compressed, and verified on the blockchain. The proposed framework divides the proof lifecycle into three operational layers to enable a consistent and performance-oriented comparison across zero-knowledge proof (ZKP) models.

These layers are as follows:

1. Proof Generation Layer:

This layer evaluates the prover-side efficiency, including generation time, memory usage, and parallelism. While zk-SNARK systems typically produce concise proofs and support efficient verification, they may require a trusted setup and incur significant proving costs depending on circuit size and constraint density [2, 4]. In contrast, zk-STARK systems offer transparent setup and high parallelism, making them well-suited for large-scale, data-intensive workloads. However, they often produce larger proof objects and may incur higher verification costs [3, 6].

2. Recurrent Aggregation/Compression Layer:

This layer assesses the possibility of aggregating and compressing multiple proofs into a smaller representation to reduce the number

of verifications. Recursive SNARK techniques can aggregate multiple proofs into a single, concise proof, transforming the verifier's workload from linear growth in the number of proofs to near-constant or logarithmic complexity [4,7]. This layer is particularly important in decentralized storage environments, where the periodic submission of proofs can lead to verification congestion.

3. On-Chain Verification (Economic Viability) Layer:

This layer assesses the cost and practical viability of on-chain verification, including gas consumption, verification time, and call data storage constraints in typical blockchain environments. Verification complexity and proof size directly impact transaction fees, throughput, and overall system sustainability [8,9].

Therefore, comparative analyses prioritize architectures that minimize on-chain verification complexity while maintaining security guarantees.

As illustrated in Figure 1, the three-layer decomposition makes the system requirements clear: proof generation must scale with data size, aggregation must control proof multiplicity, and on-chain verification must remain economically viable under network fee constraints.




Fig 1: The zk-STARK/zk-SNARK hybrid lifecycle framework for decentralized storage

Fig1. illustrates the zk-STARK/zk-SNARK hybrid lifecycle framework for decentralized storage where layer (1) creates transparent STARK proofs for multiple storage data $P_1, \dots, P_n$; layer (2) recursively combines these proofs into a single concise SNARK proof Π ; and layer (3) verifies Π on a blockchain (e.g., Ethereum/Filecoin) to achieve low gas cost and verification complexity approaching $O(1)$ or $O(\log n)$.

Comparative Results and Assumptions:

Table1 summarizes the performance comparison of the main commonly used zero-knowledge proof models for storage proof workloads in decentralized storage networks, focusing on prover and verifier response time, proof size, memory requirements, scalability, cost impacts on the blockchain, and post-quantum security characteristics.

TABLE 1. Performance Comparison of Main ZKP Models for Storage Proof Workloads

Metric	GENES / Recursive zk- SNARKs	Groth16 / Traditional zk- SNARK	zk-STARK Framework
Proof generation time	450–900 ms (aggregation overhead) [4]	100–300 ms [2]	30–80 ms (parallelizable) [3]
Verification time	2–5 ms, near-constant $O(1)$ [4]	5–10 ms (grows with volume) [2]	15–35 ms (multi-round) [3]
Proof size	1–20 KB (aggregated) [4]	5–15 KB [2]	50–500 KB [3]
Memory consumption	High (≈ 500 –1000 MB) [4]	Medium [2]	Low–Medium [3]
Gas cost suitability	Very low (single verification) [4,8]	Low–Medium [2]	High on-chain; often off-chain [3]
Post-quantum resistance	Limited (pairings) [2,4]	Limited (pairings) [2]	Strong (hash-based) [3]

Assumptions: The figures presented are based on peer-reviewed benchmarks and typical reference applications cited in the references [2, 3, 4]. All comparisons assume commercial-grade multi-core CPU environments without dedicated acceleration using FPGA/ASIC, and similar conventional security objectives (≈ 128 bits). The feasibility of on-chain implementation is interpreted under typical constraints similar to the EVM, where both verification processes and the size of proof and calldata significantly impact gas cost and throughput [8,9]. Because the results presented depend on statement size, circuit constraints, and aggregation configurations, the values are presented as ranges to reflect real-world variance.

Discussion: These results indicate that iterative aggregation of zk-SNARKs, such as GENES frameworks, is particularly beneficial for on-chain verification in proof-intensive storage environments, as it reduces multiple verification processes to a single proof check, resulting in near-constant workload for the verifier. However, this may shift the computational burden to the prover side (increasing memory usage and aggregation costs). Traditional zk-SNARK systems such as Groth16 remain effective for medium workloads with high on-chain efficiency, but they lack native scalability for large-scale aggregation, as they usually rely on a trusted setup. zk-STARK systems offer greater transparency and flexibility in the face of quantum computing, and scale well for generating proofs through parallelism. However, larger proof sizes and higher verification costs often make them more suitable for off-chain verification or hybrid designs.

Table (2) presents a structural comparison between zk-SNARK and zk-STARK in terms of setup assumptions, security foundations, proof-size characteristics, prover performance behavior, and basic cryptographic fundamentals, highlighting why each model is suitable for different decentralized storage design priorities.

TABLE 2. Structural Comparison Between zk-SNARK and zk-STARK

Criterion	zk-SNARKs	zk-STARKs
Trusted setup	Often required (e.g., Groth16, PLONK variants) [2,4]	Not required (transparent) [6]
Quantum resistance	Limited (pairing-based assumptions) [2,4]	Stronger (hash-based) [3,8]
Proof size	Very small (hundreds of bytes to few KB in some settings) [2,5]	Larger (tens to hundreds of KB) [3,10]
Prover performance	Can be heavy; higher constants [4]	Highly parallelizable; scalable for large datasets [3,10]
Main primitives	Elliptic-curve pairings [2,6]	Hash functions / FRI-style protocols [3,8]

Assumptions: This comparison reflects the common and representative characteristics of zk-SNARK and zk-STARK architectures in [2,3,4,6], namely that proof sizes and performance characteristics vary with expression complexity, circuit and path size, and implementation options. Therefore, entries describe typical ranges and qualitative behavior rather than fixed constants. The term "quantum resistance" is interpreted in the standard cryptographic sense. Pairing-based assumptions (such as those related to discrete logarithms) are less robust against future quantum adversaries, while hash-based security assumptions are more suitable for post-quantum environments [3, 8].

Discussion: These results indicate that zk-SNARK is generally preferred when on-chain constraints prevail, given its concise proofs and rapid verification, although it often requires trusted setup and relies on pairing-based cryptography [2,4]. In contrast, zk-STARK offers transparent setup and strong long-term security thanks to its use of hash functions. While characterized by high parallelism during the proof process, it can be limited by large proof sizes and high verification costs, making it difficult to deploy directly on the blockchain in high-fee environments [3,6]. For decentralized storage networks, this structural variation encourages the use of hybrid architectures that combine transparent proof generation (STARK-style) with iterative compression (SNARK-style) to achieve both transparency and efficiency on the blockchain.

Formal Hybrid Architecture (STARK-Generate → SNARK-Compress)

To mitigate the scalability and verification inefficiencies inherent in monolithic zero-knowledge proof (ZKP) systems, this work introduces a formally defined hybrid architecture that explicitly decouples proof generation from proof verification. The proposed design leverages the complementary properties of two ZKP paradigms:

- (i) The transparency and highly parallelizable proof generation of zk-STARKs.
- (ii) The succinctness and recursive aggregation capabilities of zk-SNARKs.

In this architecture, zk-STARKs are employed for large-scale, off-chain proof generation, while zk-SNARKs are used to recursively compress multiple STARK proofs into a single succinct proof suitable for efficient on-chain verification.

Stage 1: STARK Proof Generation:

For each storage epoch generating n integrity statements (e.g., per sector or shard), the prover generates a discrete STARK proof P_i for each statement $i \in \{1, \dots, n\}$. The cumulative raw proof volume before aggregation, which represents the potential on-chain data burden in a non-optimized scenario, is defined as:

$$|P_i| \sum_{i=1}^n =_{total} S \quad (1)$$

Where $|P_i|$ denotes the byte-size of the i -th proof. The primary advantage here is the avoidance of a "trusted setup" during the initial heavy-duty proof generation phase.

Stage 2: SNARK-Based Recursive Compression:

To achieve economic feasibility, a recursive SNARK aggregator circuit A is employed to compress the set of n STARK proofs into a single, succinct global proof Π defined as:

$$A(P_1, P_1, \dots, P_n) = \Pi \quad (2)$$

In this refined approach, the on-chain verifier is only required to validate the single aggregated proof Π . This architectural shift effectively reduces the verification complexity from linear growth

$O(n)$ typical of naive verification to near-constant complexity $O(1)$

Discussion:

The proposed hybrid architecture achieves a principled integration of STARK and SNARK technologies:

- Generation Phase: Preserves the transparency, parallelizability, and post-quantum security properties of zk-STARKs, while avoiding trusted setup requirements during large-scale proof generation.
- Verification Phase: Exploits the succinctness and recursive composition of zk-SNARKs to minimize on-chain verification cost and computational overhead.
- Economic Impact: Significantly reduces gas consumption and verification latency, thereby improving system throughput and enabling scalable deployment in decentralized storage and blockchain-based data availability networks.

Hybrid Architecture for Storage Proofs




Fig 2: Proposed hybrid architecture for decentralized storage proofs (STARK-Generation → SNARK-Compression → Unified On-Chain Verification).

The workflow illustrates three primary stages:

Stage 1: Parallel generation of transparent STARK proofs for individual storage integrity statements.

Stage 2: Recursive SNARK aggregation to compress the multiple proofs into a single succinct proof Π .

Stage 3: On-chain verification, where the verification complexity is reduced from linear $O(n)$ to near-constant $O(1)$, significantly minimizing gas costs and system latency.

Case Study: Optimizing Storage Proof Lifecycles in a Filecoin-Like Environment

This case study applies the proposed hybrid framework within a large-scale decentralized storage workflow, where proof generation latency and on-chain gas costs are the main scalability barriers. By aligning the system architecture with storage network requirements, we demonstrate how iterative aggregation addresses the "multi-verification" bottleneck identified in recent studies [2, 10].

Lifecycle and Implementation:

The hybrid model optimizes the storage proof lifecycle into three improved phases:

- Parallel Generation Layer (PoRep/PoSt): Instead of using homogeneous SNARKs, the system uses zk-STARKs for individual data segments. This leverages the high parallelism of STARKs to reduce proof latency [3,8], which is critical for meeting the stringent deadlines of Proof-of-Storage (PoSt) applications.
- Iterative Aggregation Layer: Using iterative frameworks such as GENES [4,7], the system compresses independent STARK proofs into a single, concise SNARK proof. This transforms the verification burden from linear growth to near-constant complexity.
- On-Chain Verification Layer: The blockchain verifier only processes the final aggregated proof, ensuring that the network is not overburdened by the size of individual storage proofs [5, 9].

Performance and resource analysis:

The following table compares the performance of the traditional approach (individual verification) with the proposed hybrid model, based on analytical criteria derived from modern performance benchmarks [1,6].

TABLE 2. Impact of Hybrid Aggregation on On-Chain Verification and Resources (Number of Proofs = 1000)

Metric	Naive Approach (Individual STARKs)	Proposed Hybrid Model (STARK + SNARK)	Improvement Factor
Verification Time (ms)	15,000 – 35,000 ms	2 – 5 ms	~7,000x Faster
On-Chain Data Size (KB)	50,000 – 200,000 KB	5 – 20 KB	~10,000x Smaller
Computational Complexity	Linear $O(n)$	Near-Constant $O(1)$	Structural Shift
Economic Feasibility	Low (High Gas Costs)	High (Sustainable)	Scalable

Discussion and Interpretation:

The data in Table 3 confirms that the key factor for scalability in decentralized storage is the reduction of multiple verifications.

Time Efficiency: By reducing on-chain verification time from approximately 35 seconds to less than 5 milliseconds for 1,000 proofs, the model eliminates block congestion and ensures that service providers can resolve claims within the required timeframe.

Economic Sustainability: The significant reduction in data size (from megabytes to kilobytes) directly lowers gas costs in EVM environments, ensuring that the cost of proof storage does not exceed the rewards received by the service provider.

Trade-off Analysis: Although the hybrid model adds a slight computational burden to off-chain aggregation, this trade-off is strategic for protecting the most expensive resource of the blockchain execution layer [4, 10].

Performance Evaluation and Discussion: A Quasi-Experimental Analysis

To provide a quantitative validation of the proposed hybrid architecture without the overhead of a full prototype deployment, this section presents a reproducible quasi-experimental evaluation. This analysis is based on established benchmarks from peer-reviewed literature, estimating how verification costs and on-chain data volume scale as the density of storage proofs increases.

Evaluation Setup and Metrics:

The evaluation contrasts two distinct design paradigms:

The Naïve Approach: Where each proof (e.g., STARK-based) is submitted and verified individually on-chain.

The Proposed Hybrid Approach: Where multiple proofs are recursively aggregated off-chain, and only a single succinct proof is submitted per cycle.

We evaluate these designs against two primary metrics: Total Verification Time (latency) and Total On-Chain Data Volume (storage overhead). These metrics are critical for assessing the economic viability and scalability of decentralized storage networks like Filecoin and IPFS, where transaction fees are tied to gas consumption [7, 8].

Semi-Experimental Results:

The values in Table 4 are derived from reference intervals in current ZKP implementations: STARK parameters for individual proof sizes and SNARK iterative aggregation criteria for the hybrid compressed proofs .

TABLE 4. Semi-Experimental Outcomes vs. Batch Size n

n (proofs/epoch)	Naïve verification time $n \cdot t_v^s$ (ms)	Hybrid verification time t_v^R (ms)	Naïve on- chain data $n \cdot \ P_s\ $ (KB)	Hybrid on- chain data $\ (\Pi)\ $ (KB)
10	150–350	2–5	500–2,000	5–20
100	1,500–3,500	2–5	5,000– 20,000	5–20
1,000	15,000–35,000	2–5	50,000– 200,000	5–20

Assumptions:

- Linearity: Individual verification time is assumed to be linearly proportional to n .
- Constant Complexity: Hybrid verification assumes recursive aggregation resulting in a proof of quasi-constant complexity.
- Data Proxy: On-chain data size serves as a proxy for the "storage burden," directly influencing gas costs in EVM-like environments [8, 9].

- Variability: The ranges reflect differences in implementation options and aggregation configurations [3,4].

Discussion and Scaling Trends:

The empirical data in Table 4 underscores that the primary bottleneck for large-scale decentralized systems is the linear accumulation of verification overhead and on-chain storage requirements.

In the Naïve Approach, both the verification time and the published data volume grow linearly with n . As shown in the scaling trends (Figures 2-3), this approach quickly becomes impractical for high-throughput storage networks due to prohibitive costs and latency. Conversely, the Hybrid Iterative Aggregation makes the verification cost almost independent of the batch size. This maintains near-stability of the on-chain data, which is particularly beneficial for networks that must repeatedly serve proofs under strict fee and throughput constraints [4,8,9]. This evidence supports the claim that the proposed architecture effectively mitigates the "linear growth" problem, offering a scalable path for next-generation verifiable cryptography.

The following figure compares the verification time with the increasing number of proofs per n epoch, and compares naïve verification per proof with hybrid iterative assembly.




Fig3. Verification time versus number of proofs per cycle n . Simple verification increases linearly with n , while hybrid iterative pooling results in a quasi-constant (or logarithmic) verification cost.

The following figure illustrates how the data footprint on the chain expands with n for a naive proof deployment versus a single pooled proof deployment in the hybrid design.




Fig 4. The size of the data on the series versus the number of proofs per cycle n . Simple diffusion is directly proportional to (n) , while the hybrid approach remains almost constant due to clustering.

Conclusion:

This paper presented a comparative analysis of Zero-Knowledge Proof (ZKP) models across their operational lifecycle within decentralized storage and blockchain environments. The findings confirm a fundamental trade-off between existing approaches: zk-SNARK-based models provide highly succinct proofs that are well suited for on-chain verification but typically rely on trusted setup procedures and stronger cryptographic assumptions, whereas zk-STARK-based models offer transparency, scalability, and post-quantum security at the cost of larger proof sizes and higher verification overhead. To address the computational and economic burden of repeated proof verification, the paper formally defined and evaluated a hybrid architecture that integrates transparent STARK-based proof generation with recursive SNARK-based aggregation. A Filecoin-inspired case study and quasi-experimental evaluation demonstrate that this hybrid approach can substantially reduce both the number of on-chain verifications and the overall data footprint, effectively rendering verification costs nearly

independent of the proof batch size. Overall, the results indicate that zero-knowledge proof systems, particularly when combined through hybrid architectures, represent a critical step toward improving the scalability and efficiency of decentralized storage and blockchain infrastructures; nevertheless, further progress requires extensive real-world experimentation, evaluation over larger datasets, and systematic comparison with alternative solutions to fully realize the potential of these technologies in future deployments.

Future Challenges and Proposed Improvements:

While this paper makes a pioneering contribution to enhancing the efficiency of Zero-Knowledge Proofs (ZKPs) in decentralized networks using techniques such as zk-SNARKs and zk-STARKs, there are several areas that can be improved or expanded to ensure the sustainability of the proposed solutions. This section identifies future challenges and proposes improvements that can support the effectiveness and scalability of the proposed solutions, including:

1. Moving from Simulation to Practical Application:

Conducting practical experiments in real-world environments (such as Ethereum testnets) is crucial for realistically evaluating gas performance and costs, rather than relying solely on theoretical or quasi-experimental models.

2. Testing Systems Using Massive Datasets:

Expanding the study to include massive datasets and workloads to simulate large decentralized storage systems like Filecoin, thereby assessing true scalability.

3. Enhancing Quantum Resilience:

Conducting an in-depth analysis of the resilience of zero-knowledge proof (ZKP) systems against advanced quantum threats, and exploring

the integration of post-quantum cryptographic techniques to ensure long-term security.

4. Integration with Layer 2 Solutions:

Exploring how zk-SNARKs and zk-STARKs interact and integrate with scalability solutions such as zk-Rollups and Optimistic Rollups to further reduce costs.

5. Expanding the Technical Comparison:

Studying other technologies such as Bulletproofs and PLONK and comparing them with existing solutions to determine the optimal economic and technical performance for large-scale decentralized storage applications.

6. Addressing On-Chain Verification Challenges:

Conducting practical tests to evaluate the actual computational load on the network in complex scenarios (such as increasing the number of nodes) to improve algorithms and reduce the impact of verification costs.

References:

- [1] M. Khaburzaniya et al., "A STARK-based Aggregator for Batch-Verify Signatures," 2021.
- [2] K. T. Vo et al., "ZCLS: A Lifecycle Strategy for Efficient ZK-Rollup Circuit Optimization in Circom," *IEEE Access*, vol. 13, pp. 101109–101123, 2025.
- [3] J. Groth, "SE-SNARK Constructions with a Single Verification," in *Proc. EUROCRYPT*, 2016.
- [4] J. Liu, L. Guo, and T. Kang, "GENES: An Efficient Recursive zk-SNARK and Its Novel Application in Blockchain," *Electronics*, vol. 14, no. 3, p. 492, Jan. 2025.
- [5] B. Zhang et al., "A Blockchain and Zero Knowledge Proof Based Data Security Transaction Method in Distributed Computing," *Electronics*, vol. 13, no. 21, p. 4260, Nov. 2024.
- [6] H. Yuan, "A Scalable Privacy-Preserving Decentralized Identity and Verifiable Data Sharing Framework based on Zero-Knowledge Proofs," *arXiv preprint arXiv:251009715*, Oct. 2025.
- [7] J. Liu et al., "Recursive IPA-based zk-SNARK Protocol and R1CS Merging," *Electronics*, vol. 14, no. 3, 2025.
- [8] V. Buterin, "Binius: Rollups with Binary Fields and Succinct Proofs," *Vitalik.ca*, 2024.
- [9] S. Gupta, "Zero-Knowledge Proofs for Privacy-Preserving Systems: A Survey Across Blockchain, Identity, and Beyond," *Engineering and Technology Journal*, vol. 10, no. 7, pp. 5755–5761, July 2025.
- [10] S. Rostamkolaei Motlagh et al., "Bottlenecks in Current ZKP Models: A Systematic Review," *IEEE Access*, vol. 13, 2025.

- [11] Q. Qamar et al., "Security Risks in Smart Contracts and ZKP Implementations," IEEE Access, vol. 13, pp. 147318–147335, 2025.