Innovations in Computing and Emerging Technologies

ISSN: 3106-6518 (Print) & 3106-6526 (Online) | Issue: 01
Peer-Reviewed Academic Journal

Research Article

A Systematic Review of Quantum-Resilient Indexing Strategies in Blockchain-Based Data Storage

Asad Ali
Department of Computer Science , Virtual University of Pakistan
Email: asadyousafzai82@gmail.com

Submitted	Accepted	Published	Article No.
7-Nov-25	24-Nov-25	19-Dec-25	ICET0125-327

Affiliations:
Asad Ali Department of Computer Science
Virtual University of Pakistan
Email: asadyousafzai82@gmail.com

Abstract

Blockchain-based databases are widely used in various applications, including health records, supply chains, and finance, due to their tamper-proof ledgers and decentralization. Data on blockchains is protected using cryptographic algorithms such as RSA and elliptic curve cryptography (ECC). As quantum computing grows, key cryptographic assumptions deserve evaluation. Public-key signature schemes are vulnerable to Shor's algorithm, whereas Grover's algorithm primarily reduces the effective security level of hash-based primitives without directly compromising them. As a result, researchers proposed several post-quantum indexing mechanisms in blockchain-based database. A systematic literature review (SLR) on the application of post-quantum cryptography (PQC) to data indexing in blockchain database has been discussed. A detailed examination and analysis have been conducted on 25 papers out of the corpus of 158 identified papers. The various classical indexing structures like Merkle trees, Patricia tries, and AVL-Merkle trees have been studied. As well as post-quantum cryptographic methods, such as hash-based (XMSS, SPHINCS+) and lattice-based (CRYSTALS-Dilithium, Kyber) schemes are identified in this literature. Extracted information tells that PQC schemes are quantum resistant but increase storage and verification costs. Various mitigation techniques like aggregation, off-chain indexing, and hybrid encryption were proposed by researchers to address these issues. This review offers insight into the importance of enhanced performance testing and a clear roadmap for transitioning to post-quantum blockchain systems.

Keywords: post-quantum cryptography, blockchain database, Merkle tree, XMSS, SPHINCS+, lattice cryptography, authenticated indexing

---

1 Introduction

Blockchain database provides decentralized, tamper-evident storage that is attractive for applications requiring high auditability; finance, supply chains, and health records, among them. The authenticated data structures such as Merkle trees, Patricia tries, and Merkleized prefix tries, and the public-key signatures used to authorize and verify state changes, are fundamental to these systems' trust model [1], [2], [3], [4]. However, development in quantum computation threaten classical assumptions. Shor?s algorithm targets integer and discrete-log-based schemes used for signatures and key exchange, and Grover?s algorithm reduces the effective security of hash functions by providing a square-root speedup to generic search attacks[5], [6], [7]. As a response, designers must consider the long-term security benefits of post-quantum cryptography (PQC) against the storage, computation, and protocol complexities that arise in real blockchain systems [8], [9]. This paper reviews existing research on quantum-resistant data indexing techniques, seeking to answer how classical and post-quantum approaches differ in terms of query efficiency, storage overhead, implementation complexity, and resilience to quantum attacks within blockchain-based database. Attention is given to practical, case-based implementations such as Electronic Health Record (EHR) ledgers. The subsequent sections, methodology, results & discussion, challenges, and conclusion observed techniques, empirical results, and open gaps to inform the case-based evaluation described in this paper.

2 RELATED WORK AND BACKGROUND

This section places the review in the context of existing research and explains the main technical concepts that form its foundation.

2.1 Authenticated Indexing in Blockchain

Verified data structures keep blockchain quick and reliable. For instance, merkle trees, Patricia Merkle tried to use in Ethereum to quickly store and access data, and AVL-Merkle trees handle range-based searches [1], [5]. These structures are used for convenient proofs of data correctness and reasonable validation, though relying on hash functions and digital signatures for the final level of verification. As a result, their overall long-term protection could be put at risk [10].

2.2 Post-Quantum Cryptographic Primitives

There is a variety of post-quantum methods, each with unique features that affect how data is indexed:

2.2.1 Hash-based schemes (XMSS, SPHINCS+)

These provide conservative security grounded purely in hash assumptions. XMSS is stateful and efficient per signature, while SPHINCS+ is stateless at the cost of larger signatures and higher computation. Hash-based schemes naturally pair with Merkleized constructions, but present signature-size and state management challenges in multi-validator systems [11], [12].

2.2.2 Lattice-based schemes (Dilithium, Falcon, Kyber as KEMs)

Lattice schemes provide relatively compact signatures and moderate verification cost, and are widely considered practical PQC candidates for ledgers [13], [14]. They become attractive when signature size and stateless operation are priorities, but their arithmetic may require optimization on constrained nodes [9], [15].

2.2.3 Others (Code-based, Multivariate)

These are less frequently used in blockchain indexing due to very large keys or verification costs in many proposed parameterizations.

2.3 Hybrid & Architectural Strategies

Because direct, naive replacement of classical primitives can inflate block sizes and verification costs, researchers propose hybrid or architectural mitigations, including:

2.3.1 Hybrid Signatures (Classical + PQC in Parallel)

Hybrid digital signature schemes combine traditional cryptographic algorithms such as ECDSA or RSA with post-quantum cryptographic (PQC) algorithms like CRYSTALS-Dilithium or Falcon. By applying both signatures to the same transaction or block, these hybrid methods ensure forward security during the migration period between classical and quantum-safe systems [13]. In essence, even if classical cryptography is later broken by quantum attacks, the PQC component still guarantees authenticity and integrity.

2.3.2 Aggregation and Batching of Signatures or Inclusion Proofs

Post-quantum signatures, such as lattice-based or hash-based schemes, produce longer keys and proofs. The computational weight of post-quantum cryptography can be handled through signature aggregation and batching. Aggregation bundles multiple signatures into a single, compact form, and batching is used for the verification of several inclusion proofs simultaneously, instead of one by one. The distribution of overhead costs across multiple transactions is achieved using these techniques. This is crucial in a blockchain context as it mitigates the scalability issues that are commonly associated with PQC adoption without jeopardizing security[3], [4].

2.3.3 Off-chain Authenticated Indexes Anchored On-chain

By offloading large cryptographic data and indexing from the main chain in this method, bottlenecks created by PQC are avoided. The efficiency of blockchain is ensured even when dealing with more difficult security procedures. Authenticated indexes like sparse Merkle trees or Merkle-based AVL trees are kept off-chain. Their state is periodically summarized using a Merkle root and recorded on-chain so that the data can be verified without storing the full index on the blockchain. The full post-quantum cryptographic data structures are stored off-chain, and their correctness is ensured by regularly anchoring cryptographic commitments on the blockchain. [8]. This design ensures that quantum-safe verification does not burden the consensus layer and allows scalability and efficiency. Moreover, transparent auditability is still maintained. Such type of information is persistent in recent literature and is the principal focus of the case-based evaluation within this paper [4], [16], [17].

3 RESEARCH METHODOLOGY

This SLR follows established guidelines adapted from evidence-based software engineering literature. The approach was implemented with reproducibility and transparency in mind.

3.1 Research Questions

The SLR addresses these research questions:

1. Which PQC primitives and indexing structures are being used in blockchain?
2. What are the measured or estimated impacts of PQC on query efficiency, storage overhead, and verification latency in blockchain indexes?
3. What architectural strategies mitigate PQC overhead for indexing (e.g., aggregation, off-chain indexing, selective application)?
4. What is missing? Where do we need more testing??

3.2 Search Strategy and Sources

Searches were carried out across multiple academic databases and augmented by backward/forward snowballing:

3.2.1 Databases

IEEE Xplore, ACM Digital Library, Scopus, Google Scholar (as primary sources), plus targeted searches of security and cryptography venues and recent preprints.

3.2.2 Search Terms

Combinations of ("blockchain" OR "distributed ledger") AND ("post-quantum" OR "quantum-resistant" OR "quantum safe") AND ("index" OR "indexing" OR "Merkle" OR "commit*" OR "searchable" OR "query") AND ("lattice" OR "hash-based" OR "XMSS" OR "SPHINCS" OR "Dilithium" OR "Kyber").

3.2.3 Date Range

2012–2025 to capture early blockchain indexing work and the full span of recent PQC development. The initial broad search and snowballing returned several hundred candidate records. Bibliographic metadata for identified records was maintained and saved in an Excel file. After removing duplicates and title/abstract screening, 169 unique records were retained for potential inclusion. Subsequently, full-text assessment produced a final set of 25 studies selected for in-depth analysis. These numbers reflect the staged screening described above and the corpus used for the Results & Discussion in Section IV.

3.3 Inclusion and Exclusion Criteria

3.3.1 Inclusion Criteria

1. Peer-reviewed conference/journal papers, and significant technical preprints that address PQC in blockchain contexts or explicitly evaluate indexing/commitment structures under PQC assumptions.

2. Works that report empirical measurements (benchmarks or simulations) or provide concrete algorithmic/design proposals for index/proof structures.

3. English language publications (2012–2025).

3.3.2 Exclusion criteria

1. Papers that discuss PQC primitives purely from a theoretical cryptography perspective without application or implication for ledger indexing or proofs.

2. Non-technical editorial pieces lacking experimental or design detail.

3. Works limited to general blockchain performance improvements, such as consensus algorithm adjustments lacking connections to indexing structures or cryptographic design, were omitted from the review.

3.4 Data Extraction and Synthesis

From each included paper, we extracted: bibliographic details (authors, year, venue); cryptographic primitives discussed; indexing/data structure(s) considered; metrics reported (signature size, proof size, storage overhead, transaction per second, verification latency); experimental platform (simulator, testbed, main net traces); and proposed architectural mitigations. The extracted fields were tabulated, enabling thematic coding and quantitative summaries of reported metrics.

A standard PRISMA-style flow (records identified ? screened ? included) was maintained in the internal review documentation and is included as Figure 1.

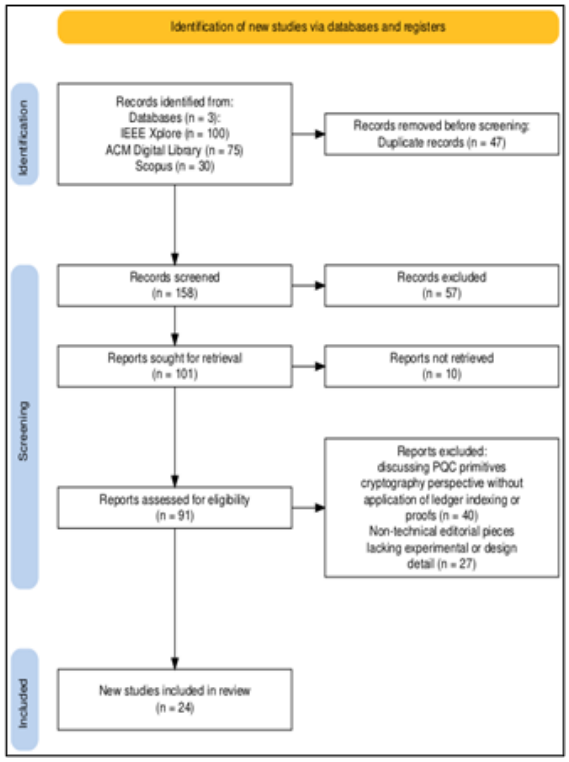

Figure 1: PRISMA flow diagram summarizing the article selection process for the systematic literature review.

The selected papers revealed that most of the work has been done in the last five years. The graph pertaining to year-wise data of selected studies is shown in Figure 2.

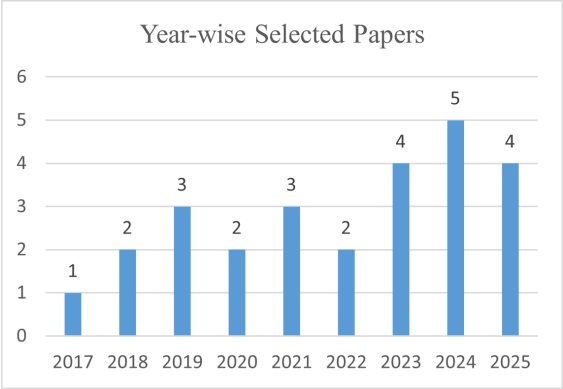

Figure 2: Year-wise data of selected papers

Out of the 25 indexing-focused studies, 20.83% examined classical indexing structures such as Merkle trees and Patricia tries. Another 20.83% explored hash-based post-quantum indexing mechanisms, including XMSS and SPHINCS+. Lattice-based proposals formed the largest category at 25%, reflecting their growing prominence in quantum-resilient blockchain design. DAG-based or hybrid architectures accounted for 16.67% of the studies. The remaining 16.67% addressed other indexing-related post-quantum techniques. A pie chart representing this distribution is shown in Figure 3.

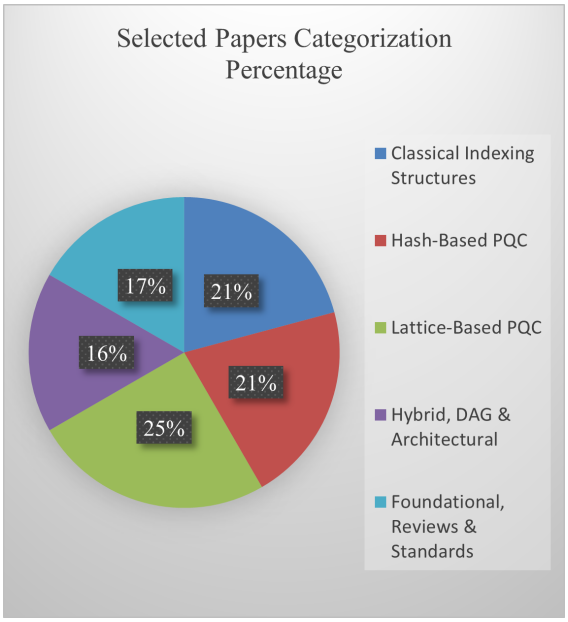

Figure 3: Distribution of Indexing types discussed in the selected paper

Selected papers are further divided into various categories based on the theme and indexing mechanism discussed in each. The summary of selected papers in fivecategories are given in Table 1.

Moreover, the summary table for the database from which these papers are extracted is given in Table 2.

4 RESULTS AND DISCUSSION

This is the heart of the review. Findings are presented by theme: (A) classical indexing structures, discussing context and limitations under quantum threats, (B) hash-based PQC and Merkle-centric adaptations, (C) lattice-based and other PQC alternatives and their indexing implications, and (D) architectural mitigations and empirical evidence. Each subsection summarizes the related studies and identifies important performance measures and design observations when possible.

4.1 Classical Indexing Structures: properties and quantum limitations

4.1.1 Merkle Trees and Merkleized Trees

Binary Merkle trees and Merkle-based prefix Trees are the elementary data structures used in blockchain ledgers. They are widespread because they use very small proof to prove that a piece of data exists in the ledger and make this whole process very efficient. These constructions also help in quick and simple updates as new data is added [1], [5], [19]. Patricia Merkle tries to use in-state storage, offering compact key-value representation and efficient path compression [18]). Balanced tree variants such as AVL-Merkle blends provide better worst-case bounds for range queries [25].

The principal quantum concern is not the tree structure itself, but the cryptographic building blocks underpinning verification (hash functions and signature schemes). Grover?s algorithm reduces the effective security of a hash function by roughly one bit of security per doubling of core operations (i.e., roughly a square-root speedup), which suggests that hash outputs and security parameters must be scaled conservatively [7]. Therefore, while Merkle trees remain conceptually valid, their parameterization (hash size, domain separation) and integration with PQC signatures must be rethought for long-term resilience [15].

Papers evaluating direct substitution of larger or quantum-hardened hashes find modest storage increases at the digest level but larger systemic impacts when signature sizes also grow: the per-node proof remains small, but the cost of storing and transmitting

Table 1: Categorization of Selected Papers

Category	Focus / Description	References
Classical Indexing Structures	Studies focusing on traditional Merkle Trees, Patricia Tries, or general blockchain querying without PQC modifications.	[1], [18], [19]
Hash-Based PQC	Schemes relying on hash functions (e.g., XMSS, SPHINCS+) and their hardware/software implementations.	[10], [11], [12], [20]
Lattice-Based PQC	Schemes using lattice structures (Dilithium, Kyber) for signatures, privacy, or infrastructure.	[9], [13], [14], [15], [17], [21]
Hybrid, DAG & Architectural	Strategies utilizing signature aggregation, off-chain indexing, DAG structures, or hybrid (Classical + PQC) designs.	[3], [4], [16], [22], [23]
Foundational, Reviews & Standards	Theoretical foundations (Shor/Grover), broad architectural surveys, NIST standardization reports, and security analyses.	[2], [5], [6], [7], [8], [21], [24]

Table 2: Summary of Selected Papers from various Databases

Scientific Database	Number of Papers	Percentage	References
IEEE Xplore	12	0.48	[1], [2], [8], [9], [10], [11], [14], [16], [17], [18], [20], [22]
ACM Digital Library	6	0.24	[4], [5], [7], [12], [13], [25]
Springer	1	0.04	[3]
Elsevier	1	0.04	[24]
World Scientific	1	0.04	[15]
Hindawi	1	0.04	[23]
Tech Science Press	1	0.04	[19]
Quantum Journal	1	0.04	[6]
NIST	1	0.04	[21]

signatures, public keys, and proofs that reference large keys can dominate block sizes in practice [21].

Patricia tries to optimize key-value retrieval and is widely used in stateful chains [18]. Their design is favorable for authenticated lookup, but their proof-size advantages can be counterbalanced if PQC increases the per-object signing metadata. Work in the literature shows that deploying PQC in systems with many small state updates (e.g., high-volume EHR writes) may disproportionately increase chain growth unless mitigations are applied [21].

4.2 Hash-Based PQC (XMSS, SPHINCS+) and Merkle-Style Integrations

4.2.1 XMSS and The Stateful Signature Model

XMSS is a Merkle-tree-based signature scheme that aggregates many one-time signatures under a single root. Its major advantage is conservative security based solely on secure hashes; its major operational challenge is statefulness, as signers must track leaf usage to prevent signature reuse [11], [20]. For single-authority append-only logs (audit trails), XMSS aligns naturally with Merkleized ledgers: the root can serve as a public verification anchor, and one-time keys correspond to per-transaction signatures.

Indexing Implications In systems where index nodes or validators are thin clients or where multiple nodes sign concurrently, XMSS state management introduces complexity: key synchronization becomes necessary, or signing must be centralized (which undermines decentralization). Storage impact is moderate (signature sizes are larger than ECDSA but vary by parameter set), and authors suggest that XMSS is best suited for append-only, controlled signing environments [20]. 4.2.2 SPHINCS+ and stateless hash-based designs SPHINCS+ eliminates the stateful requirement by building a stateless hypertree over hash-based one-time schemes. This makes it operationally attractive for distributed validator sets, but signature sizes are larger, and signing/verification computation is heavier than for most lattices [8], [12].

Indexing Implications SPHINCS+ simplifies deployment in multi-party networks but raises per-transaction bandwidth and storage needs. Studies indicate that without aggregation or careful anchoring, SPHINCS+ signatures can substantially increase block size and propagation time in high-throughput systems [8], [19], [21].

Aggregation and amortization for PQC schemes Several works propose aggregating many small proofs/signatures into batch commitments [3]. Aggregation amortizes signature overhead across many transactions, reducing per-transaction on-chain cost at the expense of latency and finality characteristics. Aggregation is particularly effective when workloads are tolerant to small batching delays [3], [4]. This aggregation strategy could improve transaction throughput by approximately 34% compared to non-aggregated approaches[26]. In the planned EHR case study, batching trade-offs must be evaluated carefully because health-

care access can be latency sensitive. 4.3 Lattice-Based & Other PQC Families

Lattice schemes (Dilithium, Falcon, Kyber) Lattice-based signatures and KEMs have found strong traction in PQC evaluations and proposals for blockchain use: they offer competitive signature sizes and verification speed relative to many hash-based stateless schemes [9], [13], [14]. Papers implementing lattice signatures in prototype ledgers report storage increases smaller than naive hash-based stateless replacements, and verification times that are acceptable when optimized or parallelized [8]. Recent frameworks have successfully integrated lattice-based primitives with Identity-Based Encryption (IBE) to enhance scalability. For instance, [26] developed a framework using lattice-based polynomials to generate keys from wallet identities. This effectively solves the key-escrow problem along with maintaining quantum resistance.

Indexing Implications Because lattice signatures are comparatively compact, they reduce the per-transaction bandwidth/registry inflation that would otherwise stress block propagation and index growth. Conversely, lattice arithmetic, such as polynomial operations, modular reductions, is heavier and may require optimized libraries or hardware support on resource-constrained nodes [13], [21].

Hybrid designs using lattices + DAG / parallel topologies There are some studies which proposed coupling lattice-based PQC with DAGs or parallel chains to merge high throughput with quantum resilience [22]. These designs distribute indexing tasks to ease the load on the main chain and lower verification costs through parallel processing. The literature here is more experimental than field-tested: prototypes show promise, moderate storage cost, and improved Transaction per second (TPS), but the complexity of implementation remains high, and the consensus/index interplay requires further study [7], [20].

4.3 Architectural Mitigations (Off-Chain Indices, Anchoring, and Selective PQC)

Off-chain authenticated indexes anchored on-chain

A common design pattern is to store large PQC artifacts or full indexes off-chain in verifiable data stores, while committing compact digests or roots to the chain periodically [16]. Off-chain indexes can provide low-latency query performance without inflating the on-chain index; the on-chain anchor preserves immutable verifiability for audits.

Trade-Offs & Risks Off-chain approaches rely on availability and trustworthy or verifiable retrieval protocols. Papers recommend combining off-chain storage with content addressing and erasure coding to manage availability and support verifiability [23]. In healthcare settings, off-chain storage can preserve patient privacy and permit richer indexing such as full-text search, but it also creates additional operational responsibilities.

Selective PQC Application and Hybrid Signatures Applying PQC selectively to long-term artifacts such as block headers, checkpoints while allowing short-term transactions to use hybrid classical+PQC signatures during migration windows has been proposed as a pragmatic transition strategy [13]. Hybrid signatures, both classical and PQC, provide immediate post-quantum auditability while preserving existing client compatibility. The downside is increased storage and verification costs while both signature types are maintained.

Compression, Deduplication, and Delta Encoding Because a nontrivial portion of on-chain growth comes from repeated key and metadata patterns such as signatures from a small set of frequent signers, deduplication and compression techniques can mitigate PQC inflation. Literature shows that effective duplication reduces on-chain index growth in practice, but its benefit depends heavily on workload patterns.

4.4 Empirical Evidence (Measured Impacts and Open Gaps)

4.4.1 Storage Overhead and Block/Transaction Size

Multiple experimental and simulation studies report that replacing small classical signatures (e.g., 64–72 bytes for ECDSA/Ed25519) with PQC signatures can multiply per-transaction metadata by 3×–30×, depending on scheme and parameterization [15], [20], [21], [24]. Aggregation and off-chain anchoring reduce effective per-transaction on-chain costs dramatically (often by 50–90%)[4], [23], but require new protocols to preserve per-transaction accountability[3].

Verification Latency and Throughput Verification cost increases are observed in many implementations [20], [24]. However, careful software optimization, parallel verification across cores, and use of optimized PQC libraries reduce end-user impact to tolerable levels for many use cases. As far as latency is concerned, lattice techniques are attractive because they typically show reduced verification overhead than stateless hash-based systems of comparable security [24]. One of the papers reported that the use of lattice-based aggregate signatures achieved a 60.5% improvement in read latency compared to standard post-quantum frameworks[26].

Benchmarking Inconsistency & Reproducibility Comparing results across existing studies is difficult. This is primarily because of researchers using very different testing setups, especially for batching techniques and workload assumptions. It is difficult to draw clear conclusions because of these variations. Standardized datasets and common evaluation benchmarks are necessary to improve consistency in future

research. Table 3 shows a gap, with reviews such as [5] and [10] highlighting the limitations of reliable decision-making due to the absence of unified data. The lack of correlation between hardware-specific experiments [12] and theoretical standards [21] validates this 'Benchmarking Inconsistency', which prevents results from being cross-verified.

4.5 Answer to research questions

4.5.1 Which PQC primitives and indexing structures are being used in blockchain?

Researchers typically follow one of two routes. First, there are traditional buildings like Patricia Tries and Merkle Trees. These remain the foundation of the system, but to withstand Grover's algorithm, they now require more stringent hash parameters. Second, there is an obvious trade-off when we examine the cryptographic tools themselves. Although hash-based schemes, such as SPHINCS+ and XMSS, are thought of as the "safe bet" for security, they have drawbacks: SPHINCS+ generates enormous signatures, and XMSS is difficult to administer. As a result, many are using lattice-based schemes (such as Kyber or CRYSTALS-Dilithium). With quick verification and smaller signatures that work better with conventional blockchains, they appear to be the "Goldilocks" choice.

4.5.2 What are the measured or estimated impacts of PQC on query efficiency, storage overhead, and verification latency in blockchain indexes?

Answer: When PQC is incorporated, you must pay for security with performance. The requirement for large storage is the biggest problem. The standard ECDSA scheme fills up the ledger quickly due to less requirement of storage. However, if we compare it with PQC, the transaction data in that case may expand from 3 to 30 times. Verification is slow, but lattice-based algorithms are generally faster than stateless hash-based schemes. Interestingly, query efficiency doesn't drop because the math gets harder (it stays logarithmic); it drops because we are forced to transmit much larger proofs and sibling hashes over the network.

4.5.3 What architectural strategies mitigate PQC overhead for indexing (e.g., aggregation, off-chain indexing, selective application)?

To stop the system from grinding to a halt, three main strategies have emerged:

1. Signature aggregation allows bundling multiple signatures into one proof. This spreads the overhead and can reduce the on-chain storage needs by as much as half, or even up to 90% in some scenarios.

2. Off-chain anchoring: The concept here is to maintain the heavy PQC indexes and data off the main chain; we only commit small, lightweight "anchors" such as Merkle roots to the ledger. This ensures that the chain can be audited without clogging it.

3. Hybrid Models: This approach is thought to be the best of the approaches mentioned earlier. Heavy PQC is only applied to high-stakes tasks, like governance or checkpoints, and we maintain faster classic schemes for regular day-to-day transactions.

4.5.4 What is missing? Where do we need more testing?

The resolution of standardization issues is necessary. The non-availability of a standard framework for testing makes it difficult to compare results across different hardware or workloads. The facts, like the length of time it takes to rebuild an index or actual database throughput, are often overlooked in current studies. Their attention is confined to the math (signing/verifying) and ignores the larger picture. Managing the state of XMSS keys in a live network and migrating the current blockchain to these new systems without breaking it are also being left out of consideration.

5 CHALLENGES AND FUTURE DIRECTIONS

It is not an easy task to adopt quantum-safe blockchain databases. Despite identifying various barriers in the literature, the main issue is how to design a system that is secure against quantum computers while still being cheap and fast enough to be utilized.

Quantum threats are effectively addressed by hash-based methods such as SPHINCS+ and XMSS. Nevertheless, their weight increases due to the increased bandwidth and storage requirements. The unfeasibility of them for high-speed databases is due to this. Thus, methods such as signature aggregation or off-chain storage [4, 12, 20] are necessary. Lattice-based algorithms (such as Dilithium) are being prioritized as they offer a better balance of speed and safety [14].

Architectural fixes are also needed. Quantum-safe crypto can only be applied when it is necessary, and we must move verified indexes off-chain. Preventing the blockchain from becoming too large while still maintaining verifiability is possible through this method.

The specific demand of the application also drives the architectural solution. For storing a permanent au-

Table 3: Table 3

Ref.	Paper Type	Focus / Topic	Relevance to Benchmarking & Reproducibility Gap
[5]	Systematic Review	Architecture Design Decisions	Designers frequently must choose system architectures that lack shared or comparable data. The uncertainty of judging a design's ability to withstand real security risks makes it unclear and largely based on experience rather than evidence.
[8]	Review Article	Post-Quantum Blockchain Overview	Comparing classical cryptography with post-quantum technology becomes difficult unless metric standards are shared. It's difficult to draw clear conclusions because different studies measure performance differently.
[10]	Systematic Review	Hash-Based Signatures (HBS)	It was brought to light that the development of optimizations for XMSS and SPHINCS+ is generally done independently. It is essential to evaluate the security performance of these individual enhancement systematically.
[12]	Experimental Study	Hardware/Software for SPHINCS+	The primary issue is that outcomes frequently depend on particular gear, such as FPGAs. Comparing them with software-only benchmarks is a challenge due to this.
[21]	Standardization Report	NIST Digital Signature Process	A formal solution is provided for inconsistent practices. The earlier problem of studies being too distinct to be compared correctly can be corrected by defining essential criteria.
[24]	Survey	Threshold Signatures & PQC	It looks at the gap between new experimental methods and final NIST standards. This gap makes it hard to compare performance reliably.

dit record, deployment of the strongest possible PQC signatures is necessary. On the other hand, off-chain indexing is a better choice if we are required to build an application which is latency-sensitive and fast. Alongside technical aspects, the operational requirements cannot be ignored. We also need to sort out critical aspects such as migration plans, management, and software upgrades. If we don?t have clear rules for the evolution of these hybrid systems or XMSS keys management, the technical reliability won't fix issues.

6 CONCLUSION

In the last decade, post-quantum indexing has shifted from theory to necessity. Reviewing 25 key studies out of 158 shows that conventional security tactics are becoming vulnerable.

Rapid developments in quantum computing are challenging our current security fundamentals. Shor?s algorithm threatens public-key systems, and Grover?s algorithm weakens hash-based protections. Switching to Post-Quantum Cryptography (PQC) is necessary to safeguard and secure data as a result.

It is clear from the findings that the current solutions are not unified and follow distinct directions. Kyber and Falcon are examples of lattice-based schemes that offer a practical balance between computational efficiency and security. Despite the increase in storage and bandwidth requirements, hash-based approaches like XMSS offer stronger long-term assurances. Database operations are under significant pressure due to increased key sizes and extended signatures, which result in slower transaction processing and verification of data, leading to a substantial increase in storage usage.

This is a significant constraint for environments that require high throughput, such as electronic health records or digital IDs. Alternative solutions are being explored by architects to solve this problem. Batching (for reducing repetitive math) and off-chain indexing (keeping the heavy data local and only placing small proofs on-chain) are some of the techniques we are seeing. We are still working on it. Testing still lacks a significant gap. The same conditions have not been used in many studies to compare traditional indexing techniques (like Merkle trees) against novel post-quantum techniques. Although real-world implementations are rare and dull, critical issues like migration planning are frequently overlooked. The focus of future research should be on broader systems in which these techniques are applied, not isolated algorithmic analysis. Applied evaluations that reflect real operational conditions are necessary, not theoretical assumptions alone. The subsequent phase of this research will be centered around the development of a practical assessment framework that can empirically evaluate the impact of post-quantum techniques on performance, cost, and scalability in a live blockchain situation.

References

[1] H. Liu, X. Luo, H. Liu, and X. Xia, Merkle Tree: A Fundamental Component of Blockchains, in 2021 International Conference on Electronic Information Engineering and Computer Sci-

ence (EIECS), Sept. 2021, pp. 556-561. doi: 10.1109/EIECS53707.2021.9588047.

[2] M. Zhaofeng, W. Lingyun, W. Xiaochang, W. Zhen, and Z. Weizhe, Blockchain-Enabled Decentralized Trust Management and Secure Usage Control of IoT Big Data, IEEE Internet Things J., vol. 7, no. 5, pp. 4000?4015, May 2020, doi: 10.1109/JIOT.2019.2960526.

[3] R. Goyal and V. Vaikuntanathan, Locally Verifiable Signature and Key Aggregation, in Advances in Cryptology CRYPTO 2022, Y. Dodis and T. Shrimpton, Eds., Cham: Springer Nature Switzerland, 2022, pp. 761-791. doi: 10.1007/978-3-031-15979-4_26.

[4] Y. Zhao, Practical Aggregate Signature from General Elliptic Curves, and Applications to Blockchain, in Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security, in Asia CCS 19. New York, NY, USA: Association for Computing Machinery, July 2019, pp. 529?538. doi: 10.1145/3321705.3329826.

[5] S. Ahmadjee, C. Mera-Gomez, R. Bahsoon, and R. Kazman, A Study on Blockchain Architecture Design Decisions and Their Security Attacks and Threats, ACM Trans. Softw. Eng. Methodol., vol. 31, no. 2, pp. 1-45, Apr. 2022, doi: 10.1145/3502740.

[6] C. Gidney and M. Eker, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits, Quantum, vol. 5, p. 433, Apr. 2021, doi: 10.22331/q-2021-04-15-433.

[7] L. K. Grover, A fast quantum mechanical algorithm for database search, in Proceedings of the twenty-eighth annual ACM symposium on Theory of Computing, in STOC 96. New York, NY, USA: Association for Computing Machinery, July 1996, pp. 212?219. doi: 10.1145/237814.237866.

[8] T. M. Fernandez-Carames and P. Fraga-Lamas, Towards post-quantum blockchain: A review on blockchain cryptography resistant to quantum computing attacks, IEEE Access, vol. 8, pp. 21091-21116, 2020.

[9] Y. Gao, X. Chen, and W. Shang, A Lattice-based Linkable Ring Signature Scheme for Blockchain Privacy Protection, in 2024 IEEE/ACIS 27th International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD), July 2024, pp. 76?80. doi: 10.1109/SNPD61259.2024.10673921.

[10] E. Fathalla and M. Azab, Beyond classical cryptography: A systematic review of post-quantum hash-based signature schemes, security, and optimizations, IEEE Access, 2024, Accessed: Oct. 19, 2025. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10731676/

[11] G. J. W. Kathrine, K. P, I. Johnraja, S. Kirubakaran, S. Salaja, and K. Arunkumar, Enhancing Smart Grid Security with XMSS-Based Blockchain Technology,in 2023 International Conference on Emerging Research in Computational Science (ICERCS), Dec. 2023, pp. 1-6. doi: 10.1109/ICERCS57948.2023.10433959.

[12] J. Lopez-Valdivieso and R. Cumplido, Design and Implementation of Hardware-Software Architecture Based on Hashes for SPHINCS+, ACM Trans. Reconfigurable Technol. Syst., vol. 17, no. 4, pp. 1-22, Oct. 2024, doi: 10.1145/3653459.

[13] P. Bagchi, B. Bera, A. K. Das, and B. Sikdar, Quantum Safe Lattice-Based Single Round Online Collaborative Multi-Signature Scheme for Blockchain-Enabled IoT Applications, ACM Trans. Sens. Netw., vol. 21, no. 2, pp. 1-33, Mar. 2025, doi: 10.1145/3715696.

[14] L. Han, G. Hu, X. Li, F. Xia, S. Wang, and L. You, A Novel Lattice-Based Blockchain Infrastructure and Its Application on Trusted Data Management, IEEE Trans. Netw. Sci. Eng., vol. 12, no. 4, pp. 2524-2536, July 2025, doi: 10.1109/TNSE.2025.3550158.

[15] A. C. H. Chen, The Security Performance Analysis of Blockchain System Based on Post-Quantum Cryptography - A Case Study of Cryptocurrency Exchanges, Jan. 23, 2024. doi: 10.1142/S2196888825500204.

[16] L. Mu, M. Lv, S. Wang, and H. Cao, A Hybrid Index-Based Block Construction and Retrieval Algorithm for Efficient Data Retrieval on Blockchain, in 2024 4th International Conference on Computer Science and Blockchain (CCSB), Sept. 2024, pp. 487-491. doi: 10.1109/CCSB63463.2024.10735653.

[17] C. Peng, H. Xu, and P. Li, Redactable Blockchain Using Lattice-based Chameleon Hash Function, in 2022 International Conference on Blockchain Technology and Information Security (ICBCTIS), July 2022, pp. 94-98. doi: 10.1109/ICBCTIS55569.2022.00032.

[18] P. Dhiman, S. K. Henge, S. Singh, A. Kaur, P. Singh, and M. Hadabou, Blockchain Merkle-Tree Ethereum Approach in Enterprise Multitenant Cloud Environment, Comput. Mater. Contin., vol. 74, pp. 3297?3313, Oct. 2022, doi: 10.32604/cmc.2023.030558.

[19] M. Liu, H. Wang, and F. Yang, An Efficient Data Query Method of Blockchain Based on Index, in

2021 7th International Conference on Computer and Communications (ICCC), Dec. 2021, pp. 1539-1544. doi: 10.1109/ICCC54389.2021.9674708.

[20] M. Ma, X. Ji, and J. Cai, A Configurable XMSS Post-Quantum Hardware Implementation with SM3 and SHA-256, in 2024 4th International Conference on Communication Technology and Information Technology (ICCTIT), Dec. 2024, pp. 220?225. doi: 10.1109/ICCTIT64404.2024.10928459.

[21] G. Alagic et al., Status report on the first round of the additional digital signature schemes for the NIST post-quantum cryptography standardization process, National Institute of Standards and Technology (U.S.), Gaithersburg, MD, NIST IR 8528, Oct. 2024. doi: 10.6028/NIST.IR.8528.

[22] X. Zhang, R. Li, W. Hou, and H. Zhao, V-Lattice: A Lightweight Blockchain Architecture Based on DAG-Lattice Structure for Vehicular Ad Hoc Networks, Secur. Commun. Netw., vol. 2021, pp. 1-17, May 2021, doi: 10.1155/2021/9942632.

[23] K. Miyachi and T. K. Mackey, hOCBS: A privacy-preserving blockchain framework for healthcare data leveraging an on-chain and off-chain system design, Inf. Process. Manag., vol. 58, no. 3, p. 102535, May 2021, doi: 10.1016/j.ipm.2021.102535.

[24] K. Sedghighadikolaei and A. A. Yavuz, A Survey of Threshold Signatures: NIST Standards, Post-Quantum Cryptography, Exotic Techniques, and Real-World Applications, ACM Comput. Surv., vol. 0, no. ja, doi: 10.1145/3772274.

[25] K.-C. Li, R.-H. Shi, W.-P. Guo, P.-B. Wang, and B.-S. Shao, Dynamic Range Query Privacy-Preserving Scheme for Blockchain-Enhanced Smart Grid Based on Lattice, IEEE Trans. Dependable Secure Comput., vol. 21, no. 4, pp. 1652?1664, July 2024, doi: 10.1109/TDSC.2023.3288228.

[26] R. Saha et al., A Blockchain Framework in Post-Quantum Decentralization, IEEE Trans. Serv. Comput., pp. 1-1, 2023, doi: 10.1109/TSC.2021.3116896.