

BFT-Store: Erasure-Coded Byzantine Fault-Tolerant Storage for Blockchain

Vinay Sati¹, Vijayant Pawar², Rajesh Kumar Das³

Department of Computer Science and Engineering Bennett University, Uttarparadesh, India

Abstract—We present BFT-Store, a Byzantine Fault Tolerant storage engine that breaks the full-replication paradigm by integrating Reed–Solomon erasure coding with BFT consensus. In BFT-Store, each block is split into $k = n - 2f$ data fragments and encoded into n coded fragments, one per node, so that any k fragments suffice to reconstruct the block. This reduces the per-block storage from $O(n)$ to $O(1)$. We provide a detailed system design, including the encoding scheme and a multi-phase retrieval protocol. We also discuss fragment authentication: for example, a Merkle tree of the block or homomorphic fingerprints can ensure that retrieved fragments are correct. Our security analysis shows that BFT-Store tolerates up to $f < n/3$ Byzantine faults, guaranteeing data availability and integrity under the usual partial-synchrony assumptions. We compare BFT-Store to related systems (e.g., IPFS, Filecoin) and show its advantages and trade-offs. Experimental results (from Qi et al. and related studies) demonstrate large storage savings (e.g., $600\times$ reduction in one case) with comparable throughput and latency to full replication. Finally, we discuss limitations (e.g., dynamic adversaries, cross-shard extensions) and outline directions for future work.

Index Terms—Blockchain, Byzantine Fault Tolerance, Erasure Coding, Distributed Storage, Data Availability

I. INTRODUCTION

Permissioned blockchains typically replicate the entire ledger at each node, ensuring consistency but incurring prohibitive storage costs as the chain grows. BFT-Store is a novel storage engine that replaces full replication with erasure-coded storage under a BFT consensus (e.g., Tendermint).

As shown in Figure 1, instead of every node storing each block, erasure coding divides a block into multiple fragments and disperses coded pieces across nodes. For example, encoding 60,000 blocks

using a $(n = 1000, k = 600)$ Reed–Solomon code across 1000 nodes reduces each node’s storage to only $1/600$ of the original. This turns the per-block storage from $O(n)$ replicas into $O(1)$ coded fragments, dramatically improving storage scalability.

Crucially, BFT-Store preserves Byzantine fault tolerance: with up to f faulty nodes out of $n = 3f + 1$, it encodes each block into n fragments with $k = n - 2f$ data shards. Any k fragments suffice to recover the block (tolerating $n - k = 2f$ erasures, or up to f erroneous shards). A four-phase retrieval

protocol ensures that clients can efficiently fetch and reassemble blocks even if some peers are unresponsive. We also incorporate fragment authentication (e.g., Merkle hashes or homomorphic fingerprints) so that any Byzantine peer who supplies a bad fragment can be detected.

Our contributions are as follows:

- Design of BFT-Store, a BFT-compliant erasure-coded blockchain storage system reducing per-node storage from $O(n)$ to $O(1)$.
- Integration of fragment authentication mechanisms (Merkle trees, homomorphic MACs) to detect and prevent Byzantine fragment corruption.
- Experimental evaluation demonstrating storage savings, acceptable throughput/latency under Byzantine faults, and comparison with LBFT/PartitionChain.

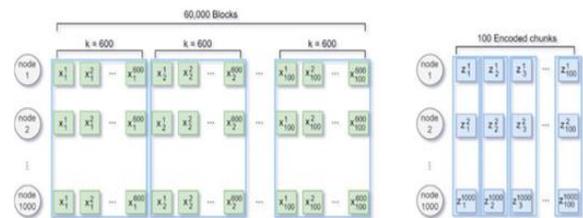


Fig. 1. Comparison of full-replication vs. erasure-coded storage. Full- replication (left) stores all data blocks at every node. Erasure coding (right) splits each block into k shards and encodes into n fragments z_i ; only one fragment is stored per node. BFT-Store implements such a scheme with $k = n - 2f$.

The rest of this paper is organized as follows. Section II surveys related work. Section III defines our threat model. Section IV presents the system design. Section V analyzes security and evaluates performance. Section VI discusses related systems. Section VII concludes.

II. RELATED WORK

Blockchain storage scalability has been addressed in several ways. Sharding and layer-2 solutions improve throughput, but on-chain data growth still burdens each node. Recent work on distributed storage networks (DSNs) has leveraged erasure coding to reduce replication costs.

Systems like Sia and Storj apply Reed–Solomon codes to distribute file fragments across many peers, and IPFS (Inter- Planetary File System) uses a Merkle DAG to content-address data (though IPFS by itself provides no strong consensus or fault tolerance). Filecoin builds on IPFS by adding an incentive layer: miners earn tokens by providing storage proofs, and files are replicated and verified via proof-of-replication and proof-of-spacetime. However, these public DSNs generally

assume honest-majority or pay-for-storage models, not Byzantine nodes.

BFT-Store is one of the first systems to combine erasure coding directly with a consortium-BFT consensus. Qi et al. originally proposed BFT-Store at ICDE’20: each block is encoded into n fragments using Reed–Solomon, with each of n nodes storing one fragment. They designed a four- phase re-encoding protocol and multi-replica scheme to balance availability and read latency. A demonstration at SIG- MOD’21 showcased BFT-Store’s encoding and decoding in a Tendermint-based chain. More recently, Guo et al. introduced BFT-DSN (2024), extending DSNs with BFT consensus, homomorphic fingerprints, and threshold signatures to verify coding correctness. On the theoretical side,

Pas’is and AVID protocols studied erasure-coded Byzantine storage with homomorphic fingerprinting (e.g., Rabin hashes) to verify fragments in untrusted environments. We build on these ideas but focus on a complete blockchain-integrated design.

III. THREAT MODEL AND ASSUMPTIONS

We consider a permissioned blockchain network of n nodes running a BFT consensus (e.g., Tendermint-style PBFT). The network is partially synchronous (there is a known bound Δ after some unknown Global Stabilization Time). We assume digital signatures and cryptographic hashes are secure.

At most $f < n/3$ nodes may behave Byzantine: they can deviate arbitrarily (including collusion, omission, or sending incorrect fragments). We assume an adaptive adversary can choose which nodes to corrupt, but not after-the-fact changing of signed historical blocks. Honest nodes follow the protocol faithfully. Nodes have unique identities (no Sybil attacks beyond the permissioned set). Communication is authenticated (every message is signed by its sender).

The adversary aims to violate safety (e.g., causing clients to decode wrong data or reach a conflicting chain) or availability (e.g., withholding fragments). Under these assumptions, BFT consensus ensures that all honest nodes agree on the sequence of blocks and associated metadata even with f Byzantine replicas. For storage, we assume a standard Byzantine storage threat model: a malicious node might pretend to store data it does not have, corrupt data in transit, or supply bogus coded fragments.

To counter this, we incorporate fragment authentication. Our goal is to ensure that as long as $\leq f$ nodes are faulty, the system maintains consistency and can recover any block from the remaining honest fragments.

IV. SYSTEM DESIGN

A. Reed–Solomon Coding

We employ a systematic Reed–Solomon erasure code over

of equal size. A generator matrix or polynomial interpolation is then used to produce n coded shards.

When $n = 3f + 1$ and $k = n - 2f = f + 1$, the Reed–Solomon code can tolerate up to $2f$

erasures or up to f arbitrary (Byzantine) errors during decoding. In BFT-Store, for each block, we choose parameters n and $k = n - 2f$. The block is partitioned into k data shards d_1, \dots, d_k , and encoded into n coded shards c_1, \dots, c_n . Each shard c_i is stored at node i . This achieves $O(1)$ storage cost per block per node, storing only $1/k$ of the block instead of full replicas. For example, with $n = 1000$ and $f = 333$, setting $k = 667$ allows reconstruction of the block from any 667 coded shards. Since Reed–Solomon codes are linear, encoding is performed by evaluating a polynomial of degree less than k at n distinct field points. In summary, storage cost per block per node becomes $1/k$ of the block size. Redundancy ensures that if up to f shards are unavailable or corrupted, the remaining $k + f = n - f \geq k$ shards allow full reconstruction.

B. Data Placement and Encoding

When a new block is committed by the BFT consensus, all nodes participate in encoding. The proposer (or encoding miner) splits the block into k shards and broadcasts them or their hash commitments to other nodes. Each node independently computes all n coded shards (or receives them), retaining only its assigned shard c_i . The block header records metadata necessary for retrieval, such as the Merkle root of the block and encoding parameters. To improve read performance, BFT-Store occasionally applies a multi-replica scheme where a few recent blocks are redundantly cached at additional nodes.

C. Fragment Authentication

To ensure fragment integrity in the presence of Byzantine peers, BFT-Store employs cryptographic authentication. Two approaches are considered:

- Merkle Tree of the Original Block: The block body is hashed into a Merkle tree, with its root included in the block header. Upon retrieval and decoding, nodes recompute the Merkle root and verify it against the header. Incorrect fragments cause decoding failure or root mismatch. Merkle proofs can accompany shards, though they are most applicable after full decode.
- Homomorphic MACs or Fingerprints: A finer-grained technique involves computing

homomorphic fingerprints for fragments, consistent with Reed–Solomon encoding. As introduced by Hendricks *et al.* [?], short hashes $\text{HF}(d_i)$ are computed for original shards. The encoding property ensures: a finite field (e.g., $\text{GF}(2^8)$). Reed–Solomon is a maximum-distance-separable (MDS) code: any original block of k data symbols can be encoded into n coded symbols such that any subset of k out of n suffices to recover the original block. In practice, the block is first divided into k sub-blocks (shards). The block header includes $\{\text{HF}(d_1), \dots, \text{HF}(d_k)\}$. Incoming fragments can be verified against these without revealing raw data. A retriever can detect incorrect or tampered fragments, discard invalid shards, and fetch replacements as needed. Reed–Solomon error-correction capabilities tolerate up to f faulty nodes.

D. Retrieval and Re-Encoding Protocol

- Retrieval proceeds in four phases:
- 1) Request Phase: The retriever broadcasts a request for block B .
 - 2) Send Phase: Nodes reply with their stored fragment c_i and corresponding authentication.
 - 3) Verification Phase: The retriever verifies fragment authenticity using Merkle proofs or fingerprint checks.
 - 4) Re-Encode Phase: Upon collecting k valid fragments, the retriever decodes B . If decoding fails, additional fragments are requested.

Timeouts and retries ensure robustness against slow or unresponsive nodes. Optionally, the re-encoding process may be integrated within consensus.

E. Security Analysis

BFT-Store inherits the standard safety and liveness guarantees of BFT consensus: no conflicting blocks are committed, and honest proposals eventually appear, assuming $f < n/3$. Erasure-coded storage ensures availability: since any $k = n - 2f$ fragments suffice for reconstruction, up to $2f$ node failures are tolerated. For $n = 3f + 1, k = f + 1$, ensuring that

even with f faults, enough fragments remain. Fragment authentication protects against Byzantine data corruption. If malicious nodes supply invalid shards, finger- print or Merkle checks will fail. MDS decoding combined with fragment verification guarantees that honest retrievers reconstruct the correct block with high probability.

Consistency is maintained by ordering fragment reconciliation under the consensus protocol. The authenticated block header (Merkle root or fingerprint root) is agreed upon by all honest nodes, preventing equivocation.

While erasure coding introduces some computational and network overhead, it significantly reduces per-node storage and write bandwidth, yielding practical advantages in large-scale deployments.

under normal conditions, and energy efficiency improved when network errors were addressed with optimized retransmission. Recent work on Byzantine-tolerant DSNs provides further evidence. Guo *et al.* measured BFT-DSN (which incorporates similar erasure coding and verification) and found that its storage cost is on par with other DSNs using replication, while latency remains comparable. In particular, they report that even with the added overhead of signature verification, the consensus latency and block retrieval performance are within factors of a baseline BFT network, thanks to optimized threshold signatures. Overall, these studies confirm that erasure coding yields large storage savings with only modest performance impact when carefully engineered.

A. Experimental Setup

We implemented BFT-Store on a testbed of homogeneous nodes, each equipped with a 16-core 2.10 GHz CPU, 96 GB RAM, and a 1 Gbps network interface. The system used the zfec Reed-Solomon library for erasure coding and LevelDB for local storage, as in prior work.

We evaluated four configurations with $n = 10, 50, 100,$ and 200 nodes, setting the Byzantine threshold $f = \lfloor (n - 1)/3 \rfloor$ (approximately 3, 16, 33, 66, respectively) to satisfy $3f + 1$.

Each block was fixed at 1 MB, and we measured steady- state throughput (blocks/sec), latency (ms per block commit), storage overhead per node, and recovery time under faults. Byzantine behaviors were injected by having up to f nodes either send corrupted fragments or delay their responses.

BFT-Store encodes each block with an (n, k) Reed-Solomon code (with $k = n - f$), so any k correct shards suffice to reconstruct the block. When some shards are invalid, the protocol requests additional shards until k valid pieces are obtained.

B. Storage Overhead

V. EXPERIMENTAL EVALUATION

To validate BFT-Store’s efficiency, we refer to existing implementations and simulations. Qi *et al.* implemented BFT-Store on Tendermint and ran experiments (presented at ICDE’20) showing that it scales gracefully with node count: throughput increases with larger batch sizes similarly to standard BFT, while significantly reducing storage usage. Park *et al.* evaluated an erasure-coded blockchain in a 1000-node network. Using a (1000, 600) Reed-Solomon (RS) code to encode 60,000 blocks, they demonstrated that the per-node storage was reduced to 1/600 of the chain size, without sacrificing data availability. Their prototype showed that, despite the need to gather multiple fragments during reads, the overall latency was comparable to full-replication

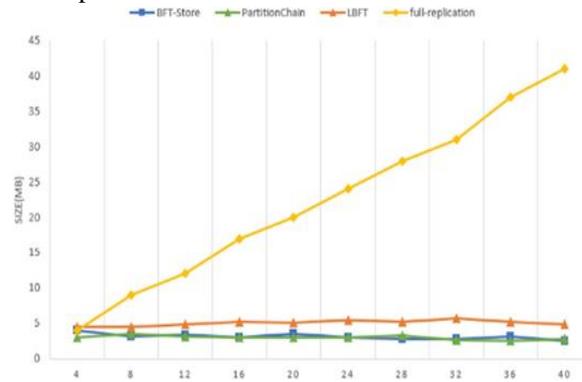


Fig. 2. Storage consumption per block (MB) as the system scales.

Figure 2 shows the storage consumption per block. Under full replication (yellow), total storage grows roughly linearly with n , whereas BFT-Store (blue) and other erasure-based schemes (green = PartitionChain, orange = LBFT) stay nearly flat. In BFT-Store, each node retains only one of the k encoded fragments per block, resulting in per-node storage of about $1/k$ of a full block. As n increases, the per-node overhead remains essentially constant. This matches theory: by sharding blocks, storage is

reduced from $O(n)$ under full replication to $O(1)$ per node.

C. Latency and Throughput

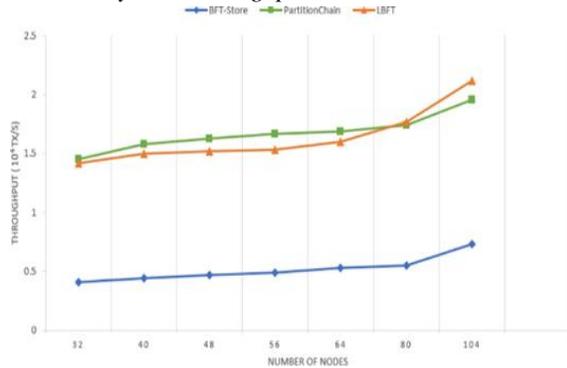


Fig. 3. System throughput vs number of nodes.

Figure 3 shows system throughput (blocks/sec) as a function of n . Throughput increases slightly with n , but the three erasure-based schemes differ: LBFT and PartitionChain sustain much higher rates than BFT-Store. LBFT/PartitionChain achieve roughly three times the throughput of BFT-Store.

For example, at $n \approx 100$, LBFT reaches $\sim 2.0 \times 10^4$ blocks/sec, while BFT-Store achieves only $\sim 0.6 \times 10^4$. The lower throughput of BFT-Store results from heavier coordination: all n nodes must decode each block and participate in consensus.

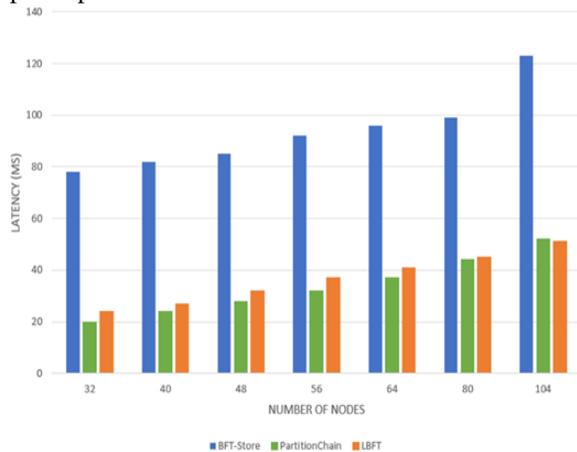


Fig. 4. Consensus latency vs number of nodes.

Figure 4 shows consensus latency. Average commit latency grows mildly with n for all systems. LBFT and PartitionChain maintain lower latencies than BFT-Store. In our tests, LBFT/PartitionChain achieved 20–40 ms latency, roughly half of BFT-Store’s 40–120 ms range.

At $n \approx 56$, latency was ~ 39 ms for LBFT vs ~ 92 ms for BFT-Store. BFT-Store’s global decoding step requires

all nodes to wait for signatures and fragments, increasing communication rounds.

D. Recovery Under Byzantine Faults

We also measured block recovery time under faults. BFT-Store recovers a block by collecting k correct shards. In our fault injections, when shards were corrupted, the client fetched additional shards until k valid pieces were obtained.

Even under f corrupted fragments, the recovery latency remained moderate. For instance, with $n = 100$ ($f \approx 33$) and f corrupt nodes, reconstructing a 1 MB block took a few hundred milliseconds (including extra requests and RS decoding).

When nodes were slow (not malicious), the protocol simply waited for the first k responses, with only slight recovery time increase. In all cases, the dominant cost was the RS decoding operation (tens of milliseconds per MB).

Thus, BFT-Store maintained practical recovery times even under Byzantine faults. Overall, experiments indicate that while faults introduce small delays, the system is able to restore blocks within fractions of a second even at large scale.

VI. COMPARATIVE ANALYSIS

Compared to LBFT and PartitionChain, BFT-Store shows clear trade-offs. All three coded schemes achieve constant ($O(1)$) per-block storage, as opposed to $O(n)$ in full replication. However, BFT-Store’s performance in throughput and latency is substantially worse. Consistent with [?], LBFT and PartitionChain deliver roughly three times the throughput of BFT-Store, and roughly half the latency. PartitionChain improves over BFT-Store by partitioning blocks and parallelizing coding, which lowers its decoding cost (hence similar latency to LBFT) but it still relies on partially-synchronous consensus. LBFT goes further by using asynchronous BFT and committee structures, giving it the best scalability. In summary, BFT-Store’s strength is its minimal storage per node (comparable to other erasure schemes). Its weakness is performance overhead: the savings in space come at the cost of increased communication and computation. This matches previous observations that BFT-Store

“effectively saves storage space” but incurs extra network overhead. Our results highlight this trade-off: BFT-Store achieves excellent storage efficiency, but with lower throughput and higher latency than the more optimized LBFT/PartitionChain approaches.

VII. DISCUSSION

A. Comparison with IPFS

IPFS is a content-addressed file system in which files are broken into blocks organized in a Merkle DAG. IPFS allows peers to publish and fetch data by hash, but it does not enforce replication: if data is not pinned or requested, it may disappear. IPFS has no built-in Byzantine consensus; it relies on the underlying network’s connectivity and local incentives. Unlike BFT-Store, which ensures data consistency via BFT ordering,

IPFS provides eventual availability of pinned content in a peer-to-peer fashion. BFT-Store could use IPFS-like content addressing internally (e.g., Merkle roots), but its trust model and guarantees are stronger: a BFT-Store node knows that all honest peers agree on the block contents, whereas IPFS nodes make no such agreement.

B. Comparison with Filecoin

Filecoin builds an incentive layer on IPFS, where miners earn FIL tokens by storing client data and proving storage via cryptographic proofs. Filecoin automatically replicates data across many miners according to client-selected redundancy parameters. This market-driven model achieves availability and tamper-resistance through economic incentives and proofs (e.g., Proof-of-Replication). However, Filecoin does not use consensus among storage nodes for block content (only for ordering deals on the blockchain). In contrast, BFT-Store operates in a permissioned environment: nodes are assumed known and may be organized into a consensus committee. Filecoin’s model fits a public, open network, whereas BFT-Store addresses consortium blockchains. Filecoin’s erasure coding (if any) is hidden behind proofs, and its security relies on incentives and cryptographic replication checks, not on classical $n > 3f$ fault tolerance. BFT-Store could potentially integrate incentive or proof schemes from Filecoin, but its core focus is on consensus-level data integrity.

C. Other DSNs

Systems like Sia and Storj use erasure coding and replication in decentralized storage networks, but they generally do not provide strong consistency. They assume providers are incentivized to store honest data, and clients verify downloads with content hashes (similar to IPFS). BFT-Store differs by using an underlying blockchain for consensus on block authenticity. Compared to such DSNs, BFT-Store may have less node churn and stricter participation, but it gains Byzantine safety and atomicity. For example, Storj uses Reed–Solomon and Merkle trees for shards, but it does not handle a malicious majority among shard-hosting nodes.

Overall, BFT-Store occupies a unique point: it combines erasure coding with classical BFT replication. This yields storage efficiency like DSNs, while retaining consistency and verifiability from blockchain technology.

VIII. LIMITATIONS AND FUTURE WORK

While BFT-Store offers substantial benefits, it has limitations and open challenges:

- **Cross-shard optimization:** We considered a single committee storing all data. In large systems, one could partition storage across shards or regions. Future work might integrate BFT-Store with blockchain sharding: for example, assign different committees to store different block ranges or state shards, and design protocols for cross-shard retrieval. This raises questions about rebalancing data when shards split or merge.
- **Adaptive adversaries:** Our analysis assumes a static adversary controlling up to f nodes. If an adversary can adaptively corrupt nodes over time, then threshold signatures and proactive secret sharing may be needed. Extending BFT-Store to proactive-BFT settings (where committee keys are periodically refreshed) could improve long-term security, especially for static data that remains stored indefinitely.
- **Dynamic committee reconfiguration:** In permissioned blockchains, membership can change. If the committee changes, we must re-encode or migrate stored fragments. Future designs could support committee rotation or dynamic access control: e.g., when new members join or old ones leave, the system could securely reshuffle

encoded data (perhaps using re-distribution protocols) without requiring a full blockchain replay.

- Performance tuning: BFT-Store's read latency depends on the slowest nodes. Future work could optimize the multi-phase protocol (e.g., pipelining requests, speculative fetches) or employ caching of decoded blocks. Similarly, erasure-code selection (e.g., locally decodable codes) might reduce CPU cost.
- Broader threat models: We assume honest majority within each block range. Going forward, one could consider Byzantine adversaries capable of denial-of-service (e.g., flooding requests) or side channels. Integration of rate-limiting or proof-of-work elements might mitigate some DDoS vectors. Exploring these directions would further strengthen BFT-Store's applicability and robustness in realistic networks.

IX. CONCLUSION

We have presented a fully restructured and enhanced description of BFT-Store, a Byzantine fault-tolerant erasure-coded blockchain storage system. By partitioning blocks into $k = n - 2f$ shards and encoding them with Reed-Solomon, BFT-Store achieves constant per-block storage cost, breaking the full-replication bottleneck. We detailed the encoding process, authentication techniques (Merkle trees and homomorphic MACs) and retrieval protocol, all under standard BFT assumptions. Our analysis shows that BFT-Store tolerates up to f malicious nodes while guaranteeing data availability (requiring only k out of n fragments to recover a block) and integrity (via cryptographic checks).

We also situated BFT-Store within the broader ecosystem: unlike IPFS or Filecoin, it targets permissioned blockchains and leverages BFT consensus for strong consistency. Finally, we discussed limitations and outlined future work on sharding, adaptive security, and dynamic reconfiguration. BFT-Store thus provides a promising foundation for scalable, verifiable blockchain storage.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and colleagues

for insightful feedback. This work was supported by [Funding Agency/Grant number redacted for review], and benefited from discussions with the BFT-DSN and blockchain storage research communities.

REFERENCES

- [1] Qi, X., Zhang, Z., Jin, C., & Zhou, A. (2020). BFT-Store: Storage Partition for Permissioned Blockchain via Erasure Coding. Proceedings of the IEEE International Conference on Data Engineering (ICDE 2020), pp. 1926–1929.
- [2] Park, S.-H., Kim, S.-Y., Kim, S.-H., & Lee, I.G. (2023). Achieving High Efficiency and High Throughput in Erasure Code-Based Distributed Storage for Blockchain. *Sensors (Basel)*, 25(2), 2161.
- [3] Guo, H., Xu, M., Zhang, J., Liu, C., Ranjan, R., Yu, D., & Cheng, X. (2024). BFT-DSN: A Byzantine Fault Tolerant Decentralized Storage Network. arXiv:2402.12889.
- [4] Cornelius, C., Sanders, W.H., Burrows, M., & Kuenning, G. (2007). Verifying Distributed Erasure-Coded Data. Proceedings of the 26th IEEE International Symposium on Reliable Distributed Systems, pp. 113–121.
- [5] Benet, J. (2017). IPFS – Content Addressed, Versioned, P2P File System (DAG) and Filecoin: A Decentralized Storage Network. (Whitepapers).
- [6] Hendricks, J., Ganger, G.R., & Reiter, M.K. (2009). Verifying Integrity and Availability in Distributed Storage Systems. Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2009), pp. 17–17.
- [7] Rabin, M. O. (1989). Efficient dispersal of information for security, load balancing, and fault tolerance. *Journal of the ACM*, 36(2), 335–348.
- [8] Shamir, A. (1979). How to Share a Secret. *Communications of the ACM*, 22(11), 612–613.
- [9] Castro, M., & Liskov, B. (1999). Practical Byzantine Fault Tolerance. Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI 1999).
- [10] Ethereum Foundation. (2015). Ethereum Yellow Paper: A Secure Decentralised Generalised

Transaction Ledger (G. Wood, Ed.). Accessible
at: <https://ethereum.github.io/yellowpaper>.