RAVEN: Retention-Aware Verifiable Extension-block Network

Ekaterina Bochvaroska

Mentor: Aleksandar Tošić

Faculty of Mathematics, Natural Sciences and Information Technologies
University of Primorska

Research Seminar (PRIN) · 2026

Outline

Topic and Research Context

Messaging systems are critical digital infrastructure.

- social communication
- economic coordination
- political interaction

Most systems today rely on centralized architectures.

End-to-end encryption does not remove the dependence on central servers.

Delivery, availability, and communication flow pass through operators.

Motivation

Centralized messaging systems create several risks:

- single points of failure
- surveillance & censorship exposure
- metadata leakage
- limited fault tolerance

Why decentralization helps?

Responsibility is distributed across many independent participants rather than one operator.

Potential benefits:

- no single infrastructure; reduced dependence on one provider
- greater fault tolerance
- stronger resistance to censorship

Blockchain

A distributed network that maintains an ordered sequence of blocks linked by hashes.

Each block typically contains:

- a reference to the previous block
- a batch of transactions
- cryptographic commitments to the block contents

Key idea

Blockchains already coordinate many distributed nodes and propagate structured data.

Related Work

- Mix Networks (Chaum), Onion Routing, GNUnet
- Early blockchain messengers (e.g., Ethereum Whisper, 2018)
- Limitations of prior P2P / decentralized systems:
  - poor adoption
  - high node maintenance
  - usability issues
  - inefficient retrieval

RAVEN builds on blockchain ordering while focusing on practical retrieval.

Core Idea of RAVEN

Hybrid approach:

- **Layer 1 – Blockchain** with cryptographic anchors (commitment stored on-chain)
- **Layer 2** – encrypted extension blocks containing the actual messages
- Clients **verify** retrieved blocks against finalized on-chain commitments
- Different retention nodes: archive and regular

Result

Lightweight ledger + verifiable off-chain storage.

RAVEN High-Level Architecture

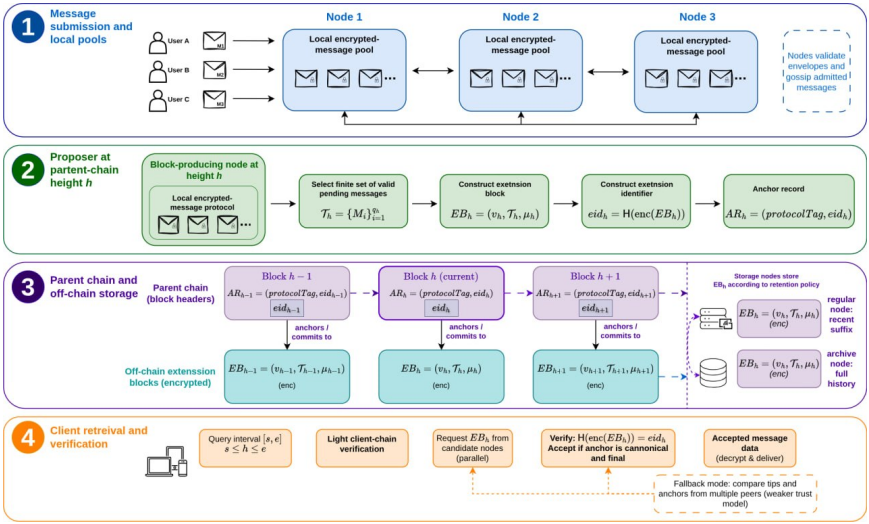

Retention-Aware Greedy Splitter

Split a request $[s, e]$ by node coverage, not into fixed-size chunks.

1. sort nodes by retention (archives last)
2. scan the range right $\rightarrow$ left
3. open a new segment only when the serving node-set changes
4. leftover old prefix $\rightarrow$ archive nodes
5. merge adjacent segments with identical candidates

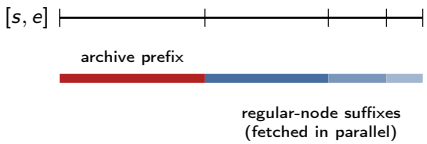

Experimental Setup

- Docker + Docker Swarm (24-node cluster)
- seed-based discovery + peer table
- failure injection (timeout, corrupt, missing, false_advertise...)
- chain height up to 20,000 blocks
- 10,000+ measured queries per strategy / concurrency
- query mix: Recent-Short 30%, Medium 30%, Long 25%, Boundary 15%
- concurrency: 1, 4, 8 , 16

Retrieval Strategies Evaluated

- Archive-only – always query full archive node(s)
- Naive whole-range – entire range, no decomposition
- RAVEN Sequential
- RAVEN Parallel
- Reorganization
- Failure Injection

Archive Load Reduction

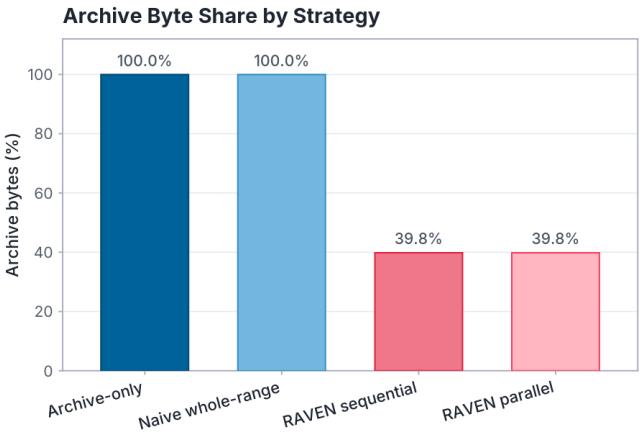

- Archive-only & Naive: 100% archive use
- RAVEN: archive usage 39.8%
- 60.2% shifted to regular nodes
- retention-aware load distribution

Completion Latency

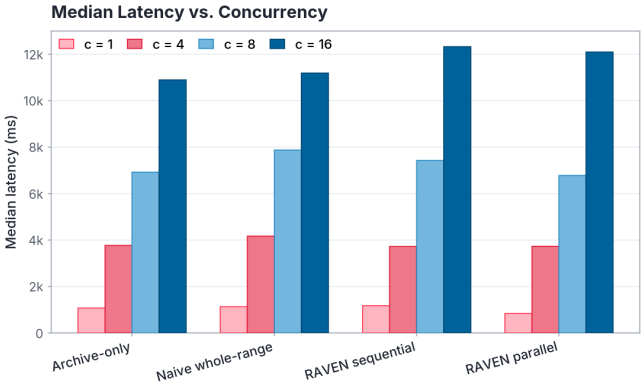

- RAVEN Parallel best at low concurrency
- parallel sub-queries reduce completion time
- higher concurrency $\Rightarrow$ fan-out & retries
- advantage decreases under heavy load

Progressive Availability & Reliability

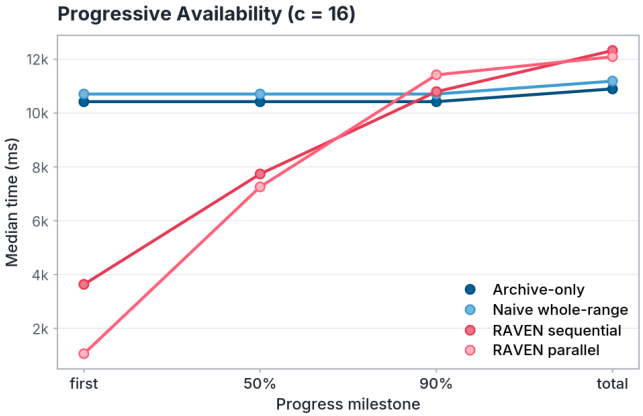

- baselines expose no data until full completion
- RAVEN Parallel: first results after 1.06 s
- enables incremental sync & earlier usability

Latency Breakdown by Phase

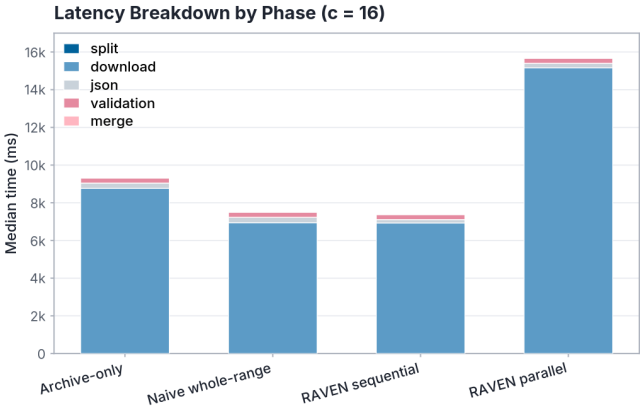

- median latency split by phase
- retrieval/response phase dominates
- split, JSON, validation, merge are small

Conclusion

RAVEN achieves its main goals:

- strong reduction in archive dependency (scalability & resilience)
- competitive performance and latency
- efficient decomposition (low subquery & retry overhead)
- balanced node utilization

Takeaway

With reorg safety and failure resilience, RAVEN is a solid foundation for decentralized messaging.

References

[1] D. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Comm. ACM*, vol. 24, no. 2, 1981.

[2] M. G. Reed, P. F. Syverson, D. M. Goldschlag, "Anonymous connections and onion routing," *IEEE JSAC*, 1998.

[3] Ethereum Foundation, "Whisper Protocol," 2018.

[4] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.

[5] "The Global Internet Shutdowns Report," Access Now, 2023.

Thank You!

Questions?