VaultX Merge: Breaking Memory Barriers in Proof-of-Space Plot Generation

Arnav Sirigere, Varvara Bondarenko, Dr. Ioan Raicu (*advisor*)
Illinois Institute of Technology, Chicago, IL

ILLINOIS TECH
College of Computing

DataSys
Data-Intensive Distributed Systems Laboratory

BACKGROUND

- MEMO is a high-throughput blockchain that uses VaultX, a lightweight Proof-of-Space (PoSp) algorithm [3] designed to enable efficient participation from both resource-constrained devices and high-performance systems [2].
- State-of-the-art Proof-of-Space systems still require substantial computational and storage resources for plot generation.
- Storing larger plots on disk is preferred, as they reduce the latency of reading from multiple smaller files during lookup.
- Large plots face memory limitations during generation, making it challenging to scale plot sizes without careful optimization.
- Our previous out-of-RAM approach offloaded intermediate data to disk, resulting in significant I/O overhead, when generating large plots.
- VaultX Merge is a new strategy that retains intermediate data in memory and merges smaller plots into larger ones, minimizing redundant I/O while maintaining high performance.
- VaultX Merge achieves up to 50% faster plot generation across diverse machines and storage configurations compared to the previous Out-of-RAM implementation.

VAULTX MERGE DESIGN

Phase 1: Subplot Generation

- VaultX Merge generates small in-memory subplots that contain Table1 and Table2 [1], reducing intermediate reads/writes.
- Subplot storage writes overlap with CPU-bound generation, enabling concurrent processing and higher throughput.

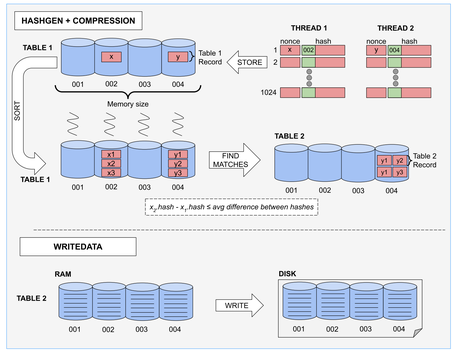

Phase 2: Merge

- A final multi-batch merge consolidates distributed bucket data into contiguous storage, improving lookup efficiency.

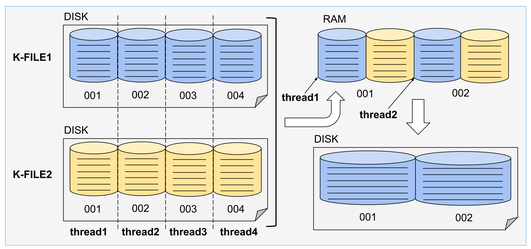

DATA COMPRESSION

- Smaller plots use 4-byte nonces with unique plot IDs and hash keys to generate hashes.
- Out-of-RAM generates larger files, thus using 5-byte nonces.
- As a result, the merged plot achieves efficient space utilization.

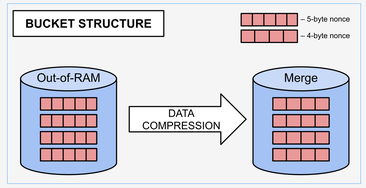

PERFORMANCE COMPARISON

- VaultX Merge performance on K-34 plots evaluated across resource-constrained and high-performance computing systems with varying storage devices, compared to previous out-of-RAM approach.
- When writing to different drives, pipelined approach is used.

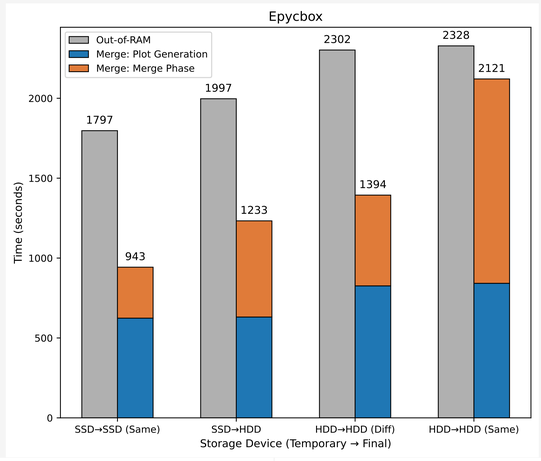

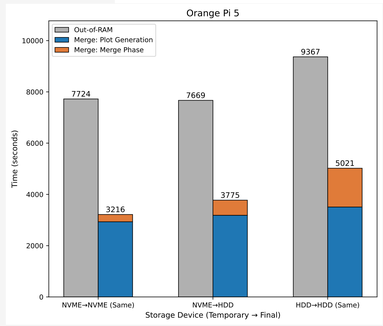

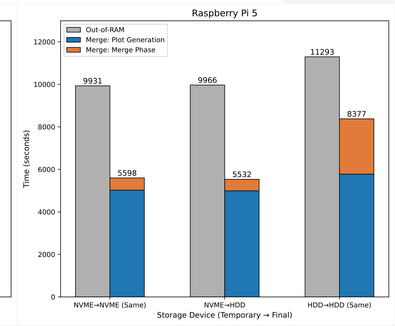

MACHINES	CPU	CORES	RAM	STORAGE	ISA
EpycBox	AMD EPYC 7501 (2.00 GHz)	64	470 GB	SATA SSD, SATA HDD	x86_64
Orange Pi 5	ARM Cortex-A76/A55 (2.3/1.8 GHz)	8	31 GB	NVMe SSD, SATA HDD	aarch64
Raspberry Pi 5	ARM Cortex-A72 (1.5 GHz)	4	8 GB	NVMe SSD, SATA HDD	aarch64

Table 1: Tech specifications of machines used for testing

VAULTX MERGE APPROACHES

Serial Merge Approach

- Causes idle CPU or storage, leading to longer runtimes on capable servers.

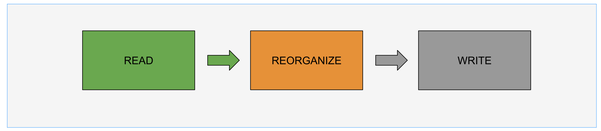

Pipelined Merge Approach

- Runs reading, reorganization, and writing concurrently using task-based parallelism.
- Uses separate temporary and final storage to avoid I/O contention.
- Maximizes CPU and storage utilization, improves throughput, and reduces total merge time.

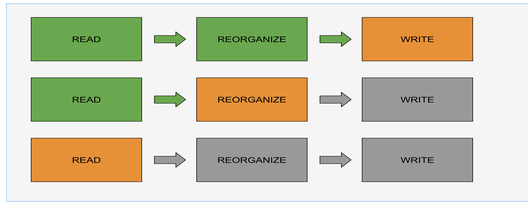

CONCLUSIONS

- VaultX Merge enables large plot generation out-of-RAM by producing multiple smaller in-memory subplots that are subsequently merged, eliminating redundant intermediate storage I/O present in prior out-of-RAM methods and significantly reducing overhead.
- The method supports scalable plot generation across a spectrum of hardware, from resource-constrained devices to high-performance computing systems.
- By leveraging unique plot IDs and smaller nonces, VaultX Merge enhances storage efficiency without compromising plot integrity.
- Overall, VaultX Merge advances the speed, scalability, and resource utilization of out-of-RAM plot generation, thereby improving the performance and practicality of the MEMO Proof-of-Space blockchain.

REFERENCES

[1] Varvara Bondarenko and Ioan Raicu. *Securing Proof-of-Space Against Hellman Attacks.*
[2] Varvara Bondarenko et al. "Improving the Performance of Proof-of-Space in Blockchain Systems". In: (2024).
[3] Stefan Dziembowski et al. "Proofs of Space". In: *Annual Cryptology Conference.* Springer, 2015, pp. 585–605.

ACKNOWLEDGEMENTS

This work is supported in part by the National Science Foundation OAC-2150500 award.