Securing Proof-of-Space against Hellman Attacks

ILLINOIS TECH
College of Computing

Varvara Bondarenko, Ioan Raicu (*advisor*)
Illinois Institute of Technology, Chicago, IL

DataSys
Data-Intensive Distributed Systems Laboratory

BACKGROUND

- State-of-the-art Proof-of-Space (PoSp) algorithms demand high-performance hardware, limiting accessibility for users with constrained computational or memory resources.
- PoSp systems can be exploited by attackers who manipulate storage proofs to gain unfair advantages, compromising system integrity [1].
- Our previous implementation focused on writing records (nonce, hash pairs) directly to disk without further processing, which was vulnerable to Hellman's time–memory trade-off attack [2].

HELLMAN ATTACK

- Stores less data to recompute it on demand.

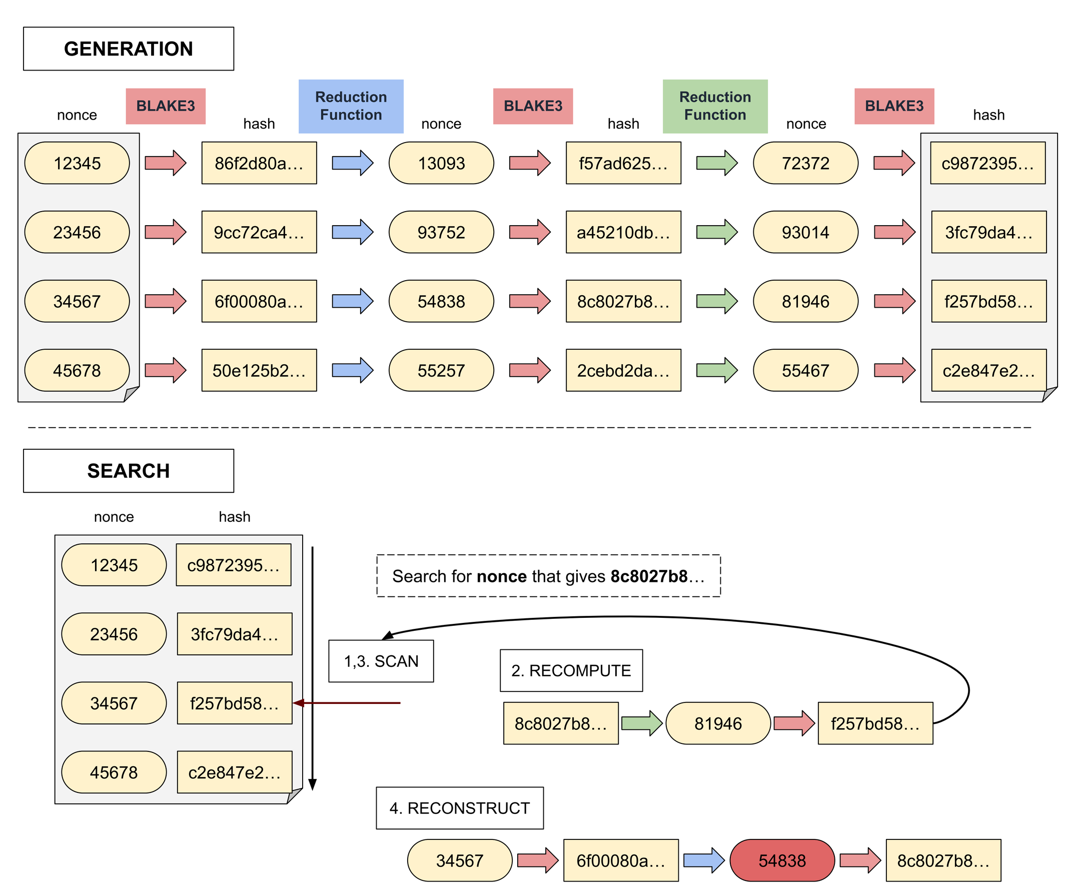

Fig. 1: Rainbow Table

VAULTX DESIGN

- VAULTX stores 2x 4- or 5-byte nonces instead of 32-byte records.
- Data compression shrinks resulting files by a factor of 4.
- VAULTX writes data in large chunks to HDD, optimizing sequential writes to improve speed.

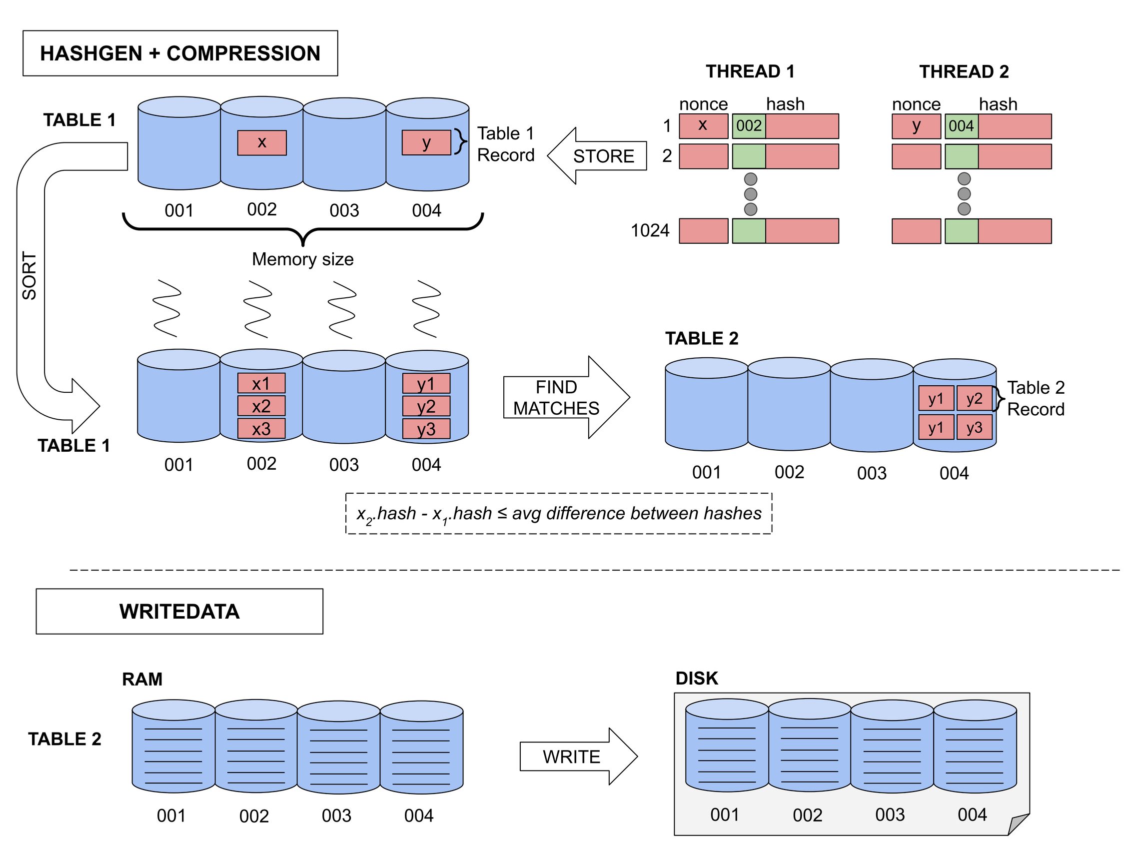

Fig. 2: VAULTX: Data Generation

PERFORMANCE COMPARISON

- Performance comparison between VAULTX and Chia in-RAM plotters on the EpycBox system with $k = 32$ plots [Table 2].

CHIA PLOTTER	RAM REQUIREMENT	K-VALUE
Bladebit Ramplot	416 GB	32
madMAx RAM disk	138 GB	32
VAULTX in-RAM	32 GB	32

Tabel 1: Chia In-RAM Plotting

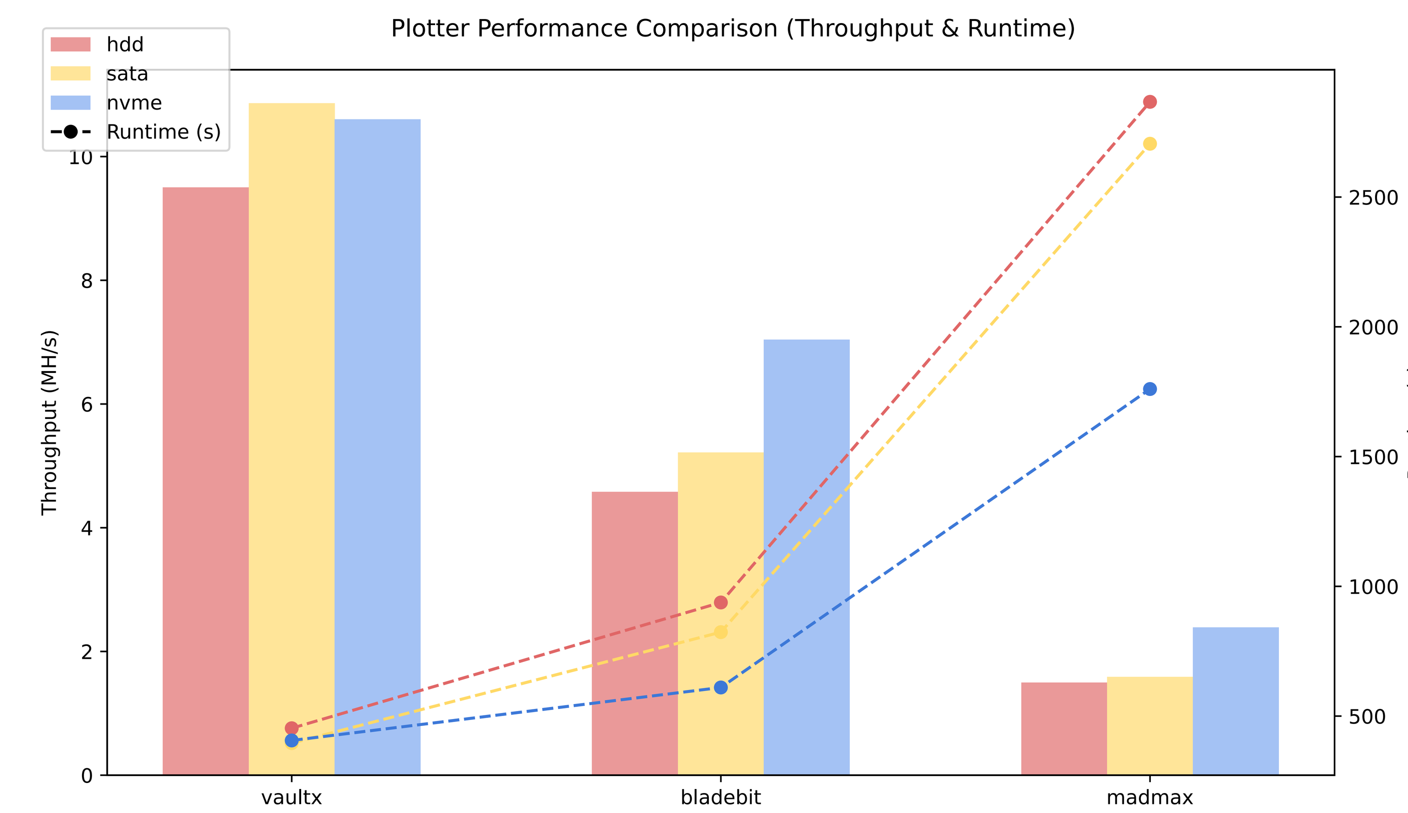

Fig. 3: Comparison between VAULTX and Chia in-RAM plotters on the EpycBox system with k = 32 plots

- Performance comparison of VAULTX across various storage mediums (HDD, SATA SSD, NVMe).
- Evaluation of VAULTX efficiency on resource-constrained and high-performance computing systems.

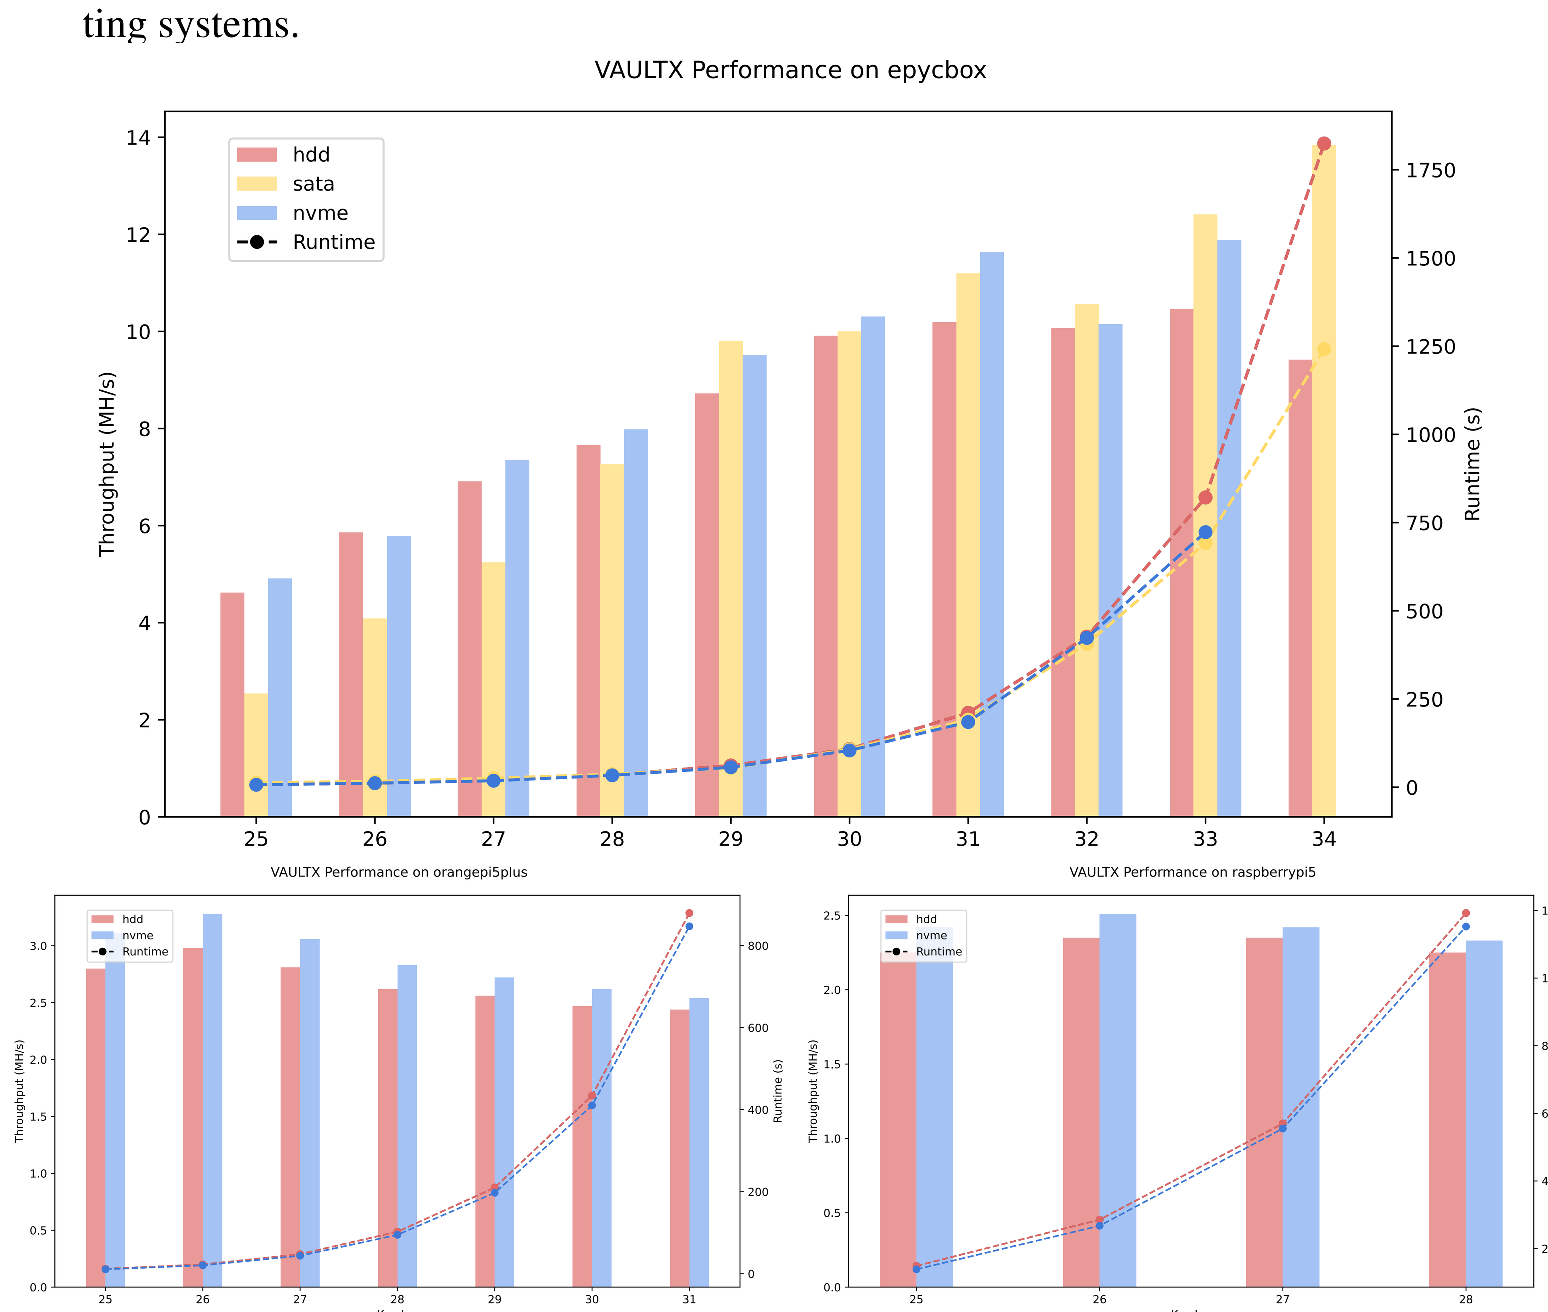

Fig. 4: Throughput and latency measurements for k=25-32 on various machines

MACHINES	CPU	CORES	RAM	STORAGE	ISA
EpycBox	AMD EPYC 7501 @2.00GHz	64	480GB	SATA, NVMe SSD, SATA HDD	x86_64
Orange Pi 5 Plus	Cortex-A55 @ 1.80GHz	8	32 GB	NVMe SSD, HDD	aarch64
Raspberry Pi 5	Cortex-A76 @ 2.50GHz	4	8 GB	NVMe SSD, HDD	aarch64

Tabel 2: Tech specifications of Mystic nodes used for testing [3]

LOOKUP COMPARISON

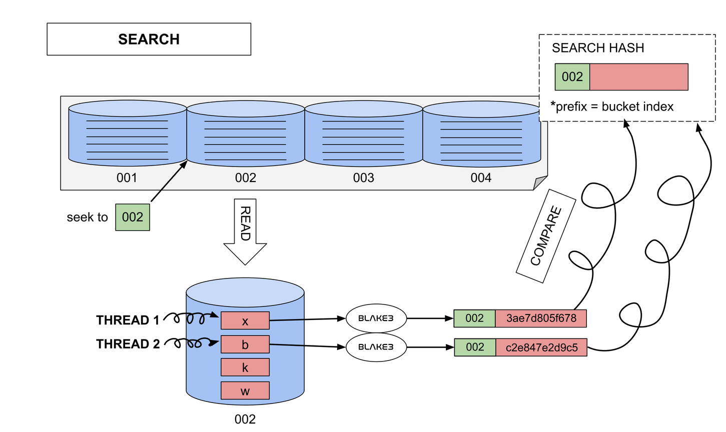

Fig. 5: VAULTX: Lookup

The lookup latency was evaluated on the EpycBox machine [Table 2] following the generation of a file with $k = 32$ records.

DRIVE	MIN (ms)	AVG (ms)	MAX (ms)
NVMe SSD	0.27	0.414	0.56
SATA SSD	0.33	0.522	0.79
SATA HDD	6.45	6.534	6.65

DRIVE	MIN (ms)	AVG (ms)	MAX (ms)
NVMe SSD	5.00	10.65	96.00
SATA SSD	4.00	9.52	94.00
SATA HDD	5.00	10.63	72.00

Tabel 3: Latency per Single Lookup: *VAULTX (left)*; *ChiaPoS (right)*;

CONCLUSIONS

- VAULTX in-RAM plotter outperforms Chia's in-RAM plotters in overall runtime and efficiency.
- VAULTX achieves comparable plot generation runtimes on HDDs relative to NVMe and SATA disks.
- VAULTX delivers high throughput and low lookup latency on resource-constrained machines, with reduced RAM requirements for in-memory plotting compared to Chia.
- VAULTX offers faster lookup performance; in contrast, Chia provides consistent lookup times across varying storage mediums.

FUTURE WORK

- Implement and benchmark an out-of-RAM version of VAULTX to evaluate performance under constrained memory conditions.
- Compare the lookup latency of VAULTX and Chia plotters as the number of generated records increases.

REFERENCES

[1] Hamza Abusalah *et al.* „Beyond Hellman's time-memory trade-offs with applications to proofs of space". Teoses: *Advances in Cryptology–ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II 23*. Springer. 2017, lk. 357–379.

[2] Varvara Bondarenko *et al.* „Improving the Performance of Proof-of-Space in Blockchain Systems" (2024).

[3] AI Orhean *et al.* „Mystic: Programmable systems research testbed to explore a stack-wide adaptive system fabric". Teoses: *8th Greater Chicago Area Systems Research Workshop (GCASR)*. 2019.

ACKNOWLEDGEMENTS

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

CONTACT: Varvara Bondarenko (vbondarenko@hawk.iit.edu)