Exploring Green Cryptographic Hashing Algorithms for Eco-Friendly Blockchains

Aahad Abubaker [1] Tanmay Anand [2] Sonal Gaikwad [2] Mahad Haider [2] Jacklyn McAninch [2]
Lan Nguyen [2] Alexandru Iulian Orhean [2] Ioan Raicu (Advisor) [2]

[1]DePaul University [2]Illinois Institute of Technology

Abstract

Cryptographic hash functions are fundamental for ensuring data security and integrity in all consensus algorithms in blockchains. While SHA256 has been widely used in many blockchain implementations, its throughput and efficiency has led the rise of a modern lightweight and speed superior implementation BLAKE3. We compared and contrasted SHA256 and BLAKE3 with a focus on blockchain workloads with small inputs and outputs. We explored different compilers and optimizations, different ways to parallelize using multi-threading and multi-processing, as well as different size systems from small Raspberry Pi 4 to a modern AMD Epyc server. We found that BLAKE3 is superior from a performance perspective. To showcase its strengths, we integrated BLAKE3 into a basic Proof-of-Space implementation that used advanced data index and search, and compared our results to the Chia blockchain plotting mechanism. Our approach offers one to two orders of magnitude higher hash generation and storage rates.

Background

- In the realm of blockchain technology, cryptographic hash functions are fundamental for ensuring data security and integrity. While SHA256 has been widely used, its throughput and efficiency have come under scrutiny. BLAKE, a more advanced hash function, has emerged as a potential alternative with superior throughput
- Blockchain's decentralized ledger system has revolutionized industries. However, cryptocurrencies like Bitcoin, which rely on energy-intensive Proof of Work (PoW) consensus, have raised environmental concerns. The rise of Proof of Space, exemplified by Chia, offers an energy-efficient alternative by using storage instead of computation for consensus.
- Our research compares the throughput of BLAKE and SHA256 and explores the potential of Proof of Space. Through benchmarking hash functions and analyzing blockchain consensus mechanisms, we can implement our findings in our implementation of a Proof of Space consensus, CryptoMemoiz.

Machine Specifications

For our experiments, we ran our hashing and subsequent plotting experiments on the following systems:

CPU Model	Sockets	Compute Power	RAM
Intel SP 8160	8	192c (384 HT) @ 2.1 GHz	768 GB
Intel HW 2620 v3	2	12c (24HT) @ 2.4 GHz	32GB
Intel Xeon Phi 7210	1	64c (256HT) @ 1.5 GHz	64GB
AMD Naples 7501	2	64c (128 HT) @ 2 GHz	128GB
Cortex-A72 (Pi)	1	4c (4 HT) @ 1.5 GHz	2 GB

Acknowledgements

This work was supported in part by the National Science Foundation (NSF) awards 2150500 and 2150501.

Hashing Benchmarking Methodology

For our testing, our hash generation process consisted of the following when comparing the throughput and number of hashes generated by the hash functions, BLAKE3 and SHA-256.

- **Compilers:** We used two different compilers, GCC and Clang, to observe differences in throughput on our hashing benchmarks, gaining insights into how compiler choice impacts the hash generation speeds.
- **Parallel Processing:** GNU Parallel and OpenMPI for parallel benchmarking were used.
- **Multithreading:** Conducted a benchmarking study on multithreading capabilities using SHA-256 and BLAKE3 as hashing functions across various computing platforms, ranging from 1 to the maximum hardware threads of the respective machine.
- **I/O Size:** To generate hashes, the input size will be 64 bytes and the output size will be 12 bytes.

Benchmarking Results

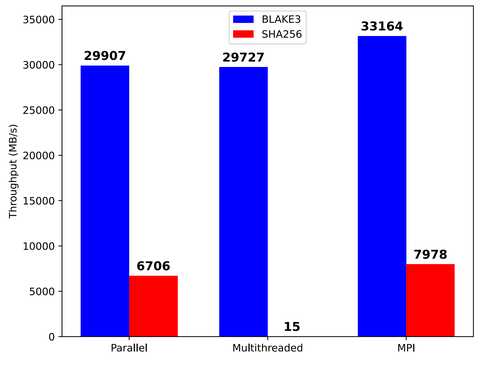

Figure 1. Hashing Performance on Intel SP 8160 Machine (Clang)

- Implementations were ran at the upper bound (384).

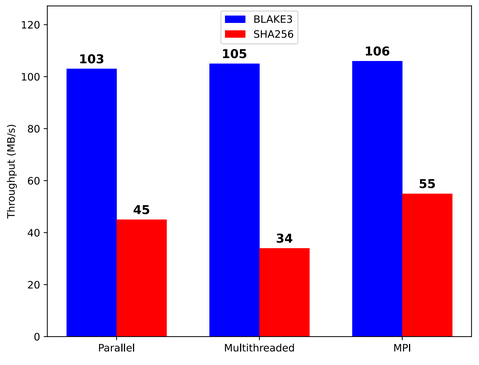

Figure 2. Hashing Performance on Cortex-A72 Raspberry Pi (Clang)

- Implementations were ran at the upper bound (4 HT).

MEMO

Memo's Proof of Space implementation, CryptoMemoiz, leverages XSearch indexing to reduce I/O, large memory, and significant compute resource requirements. Memo was designed to be lightweight enough to operate on small nodes such as Raspberry PIs, as well as making use of high-end servers with 100+ cores and hundreds of gigabytes of memory. Following are some problems with the current leading blockchains that MEMO aims to address:

- **Low Throughput:** Leading cryptocurrencies like Bitcoin can only handle around 5 transactions per second, which is significantly slower than centralized payment methods like Visa and Mastercard, which can handle around 20,000 transactions per second.
- **High Carbon Emission:** Bitcoin's energy-intensive proof-of-work consensus creates a large carbon footprint. MEMO, a proof-of-space coin, is environmentally friendly by using memoization to save hashes rather than generating new ones with power-hungry GPUs and ASICs.
- **Problems with Alternatives:** While Ethereum, the second most popular cryptocurrency in the world, has shifted from proof-of-work to proof-of-stake, the problem with the proof-of-stake validation method is that it can potentially centralize the network, as only a few people in the world may have enough money to become a validator.

Proof of Space Implementation

Using our knowledge of an optimized hash function and CryptoMemoiz' XSearch, we implemented BLAKE3 to our Proof of Space implementation, to which we generated benchmarks on the throughput of filling vaults with hashes using XSearch:

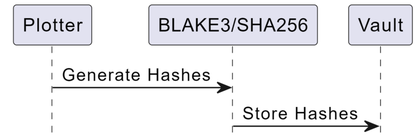

Figure 3. Implementation diagram for CryptoMemoiz

Proof of Space Results

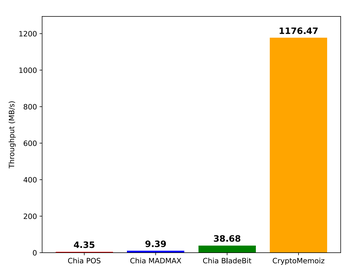

Figure 4. Plotting Performance on AMD Naples 7501 EpycBox

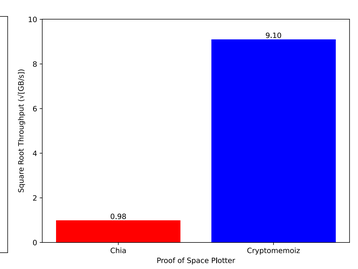

Figure 5. Plotting Performance on Raspberry Pi

Conclusions

- In summary, our comprehensive analysis conclusively establishes BLAKE3 as the superior choice for hash generation, especially when leveraged with advanced techniques like multithreading, parallel processing in GNU, and MPI.
- Furthermore, our successful integration of the optimized BLAKE3 function into CryptoMemoiz's proof of space implementation presents a significant breakthrough. The resulting plot generation performance surpasses the existing plotting mechanisms of the Chia blockchain. This achievement not only sets the stage for faster and more scalable plot creation but also aligns the aspiration for environmentally conscious and high-performance blockchain solutions.

References

[1] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. 2013. URL `https://eprint.iacr.org/2013/796`.

[2] Alexandru Iulian Orhean, Anna Giannakou, Lavanya Ramakrishnan, Kyle Chard, and Ioan Raicu. Scanns: Towards scalable and concurrent data indexing and searching in high-end computing system. In *2022 22nd IEEE International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, pages 51–60, 2022. doi:10.1109/CCGrid54584.2022.00014.