



Accelerating Merkle Patricia Trie with GPU


Yangshen Deng*, Muxi Yan*, Bo Tang
Southern University of Science and Technology
dbgroup@sustech.edu.cn



Background

Merkle Patricia Trie (MPT) is a **cryptographical index**.

- **Structure**: patricia trie (radix trie)
- **Cryptographical hashing**: merkle tree



Motivation

MPT is widely used in **immutable data systems**

- **Blockchain**: Ethereum, FISCO BCOS, Quorum, etc.
- **Verifiable database**: Alibaba LedgerDB, Spitz, GSSE, etc.

MPT update is the **bottleneck**


- Heavy **batched** write to MPT index structure
- Heavy **parallelizable** cryptographic hashing computation

GPU processes in **batch** with massive **parallelism**

Challenge


How to parallelize MPT update?

- **Node update**: concurrent modification to the index structure
- **Hash compute**: read-write dependency in concurrent hash update




Solution: Parallel Algorithms on GPU


Node Update



Hash Compute




Index Insertion Benchmark



Integrating GPU MPT Into Real-World Database Systems

Ethereum (up to **3.4X** speedup generating a block)



LedgerDB (up to **1.8X** system throughput)

