



LVMT: An Efficient Authenticated Storage for Blockchain

Chenxing Li¹, Sidi Mohamed Beillahi², Guang Yang¹, Ming Wu¹, Wei Xu³, Fan Long^{1,2}

¹Shanghai Tree-Graph Blockchain Research Institute

²University of Toronto ³Tsinghua University

Evolution of Blockchain Performance

Early blockchain systems are slow



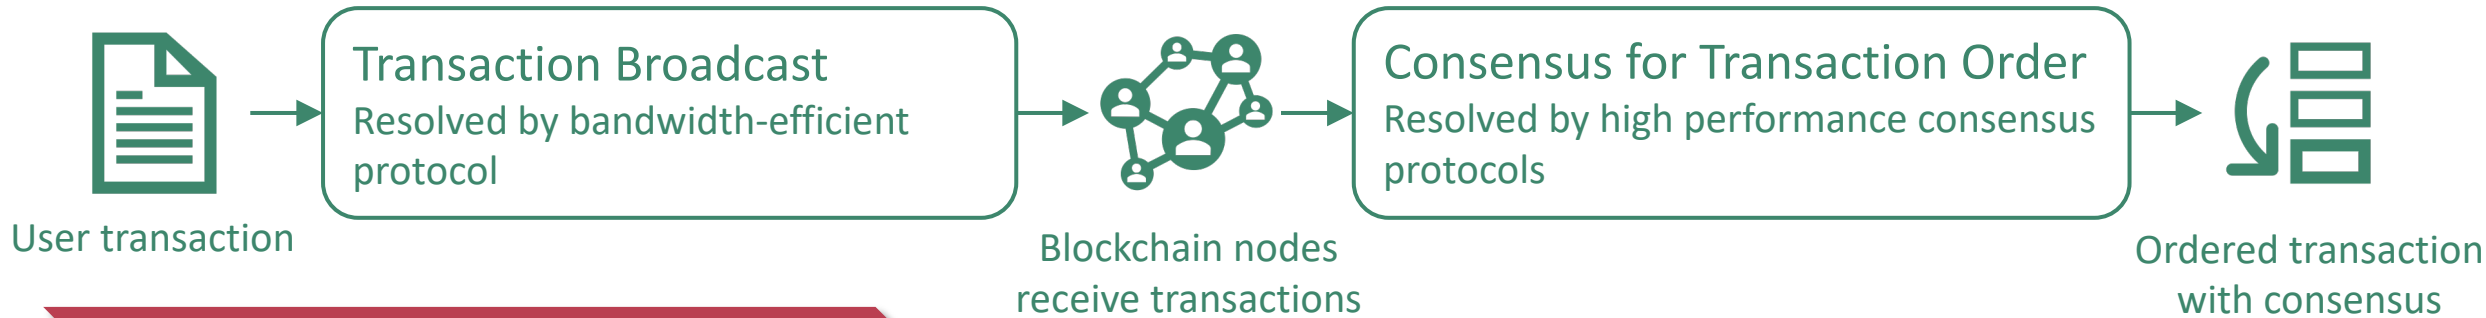
(< 30 transactions / second)



Reaches 20,000 transactions/second

Resolved bottlenecks

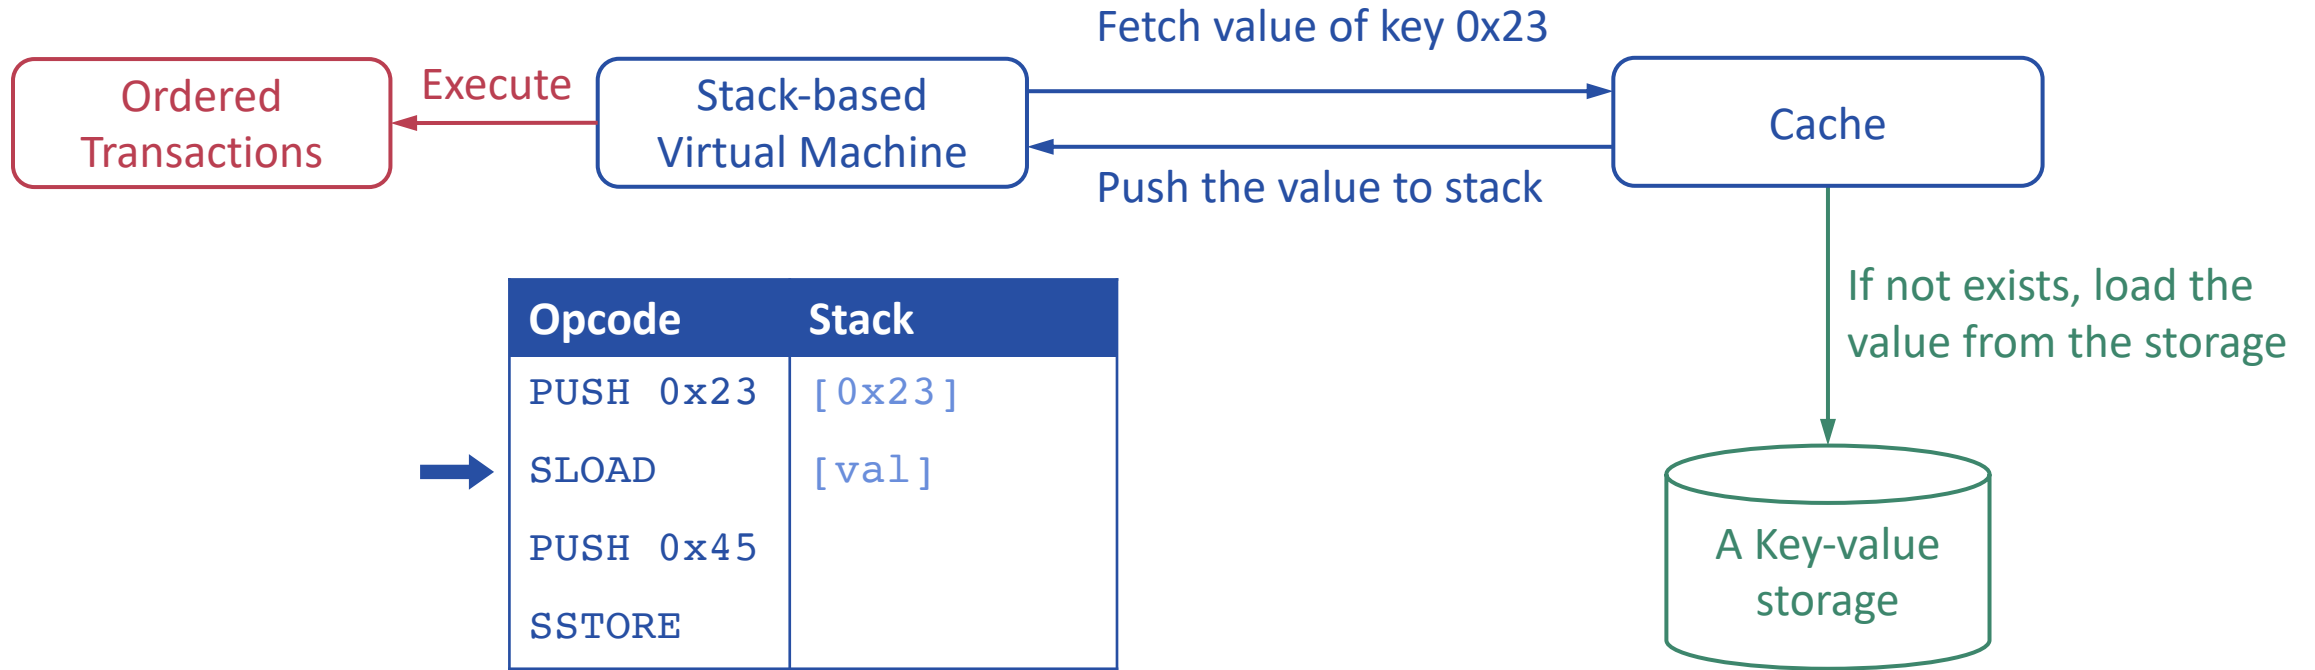
(reaches 20,000 transactions/second)



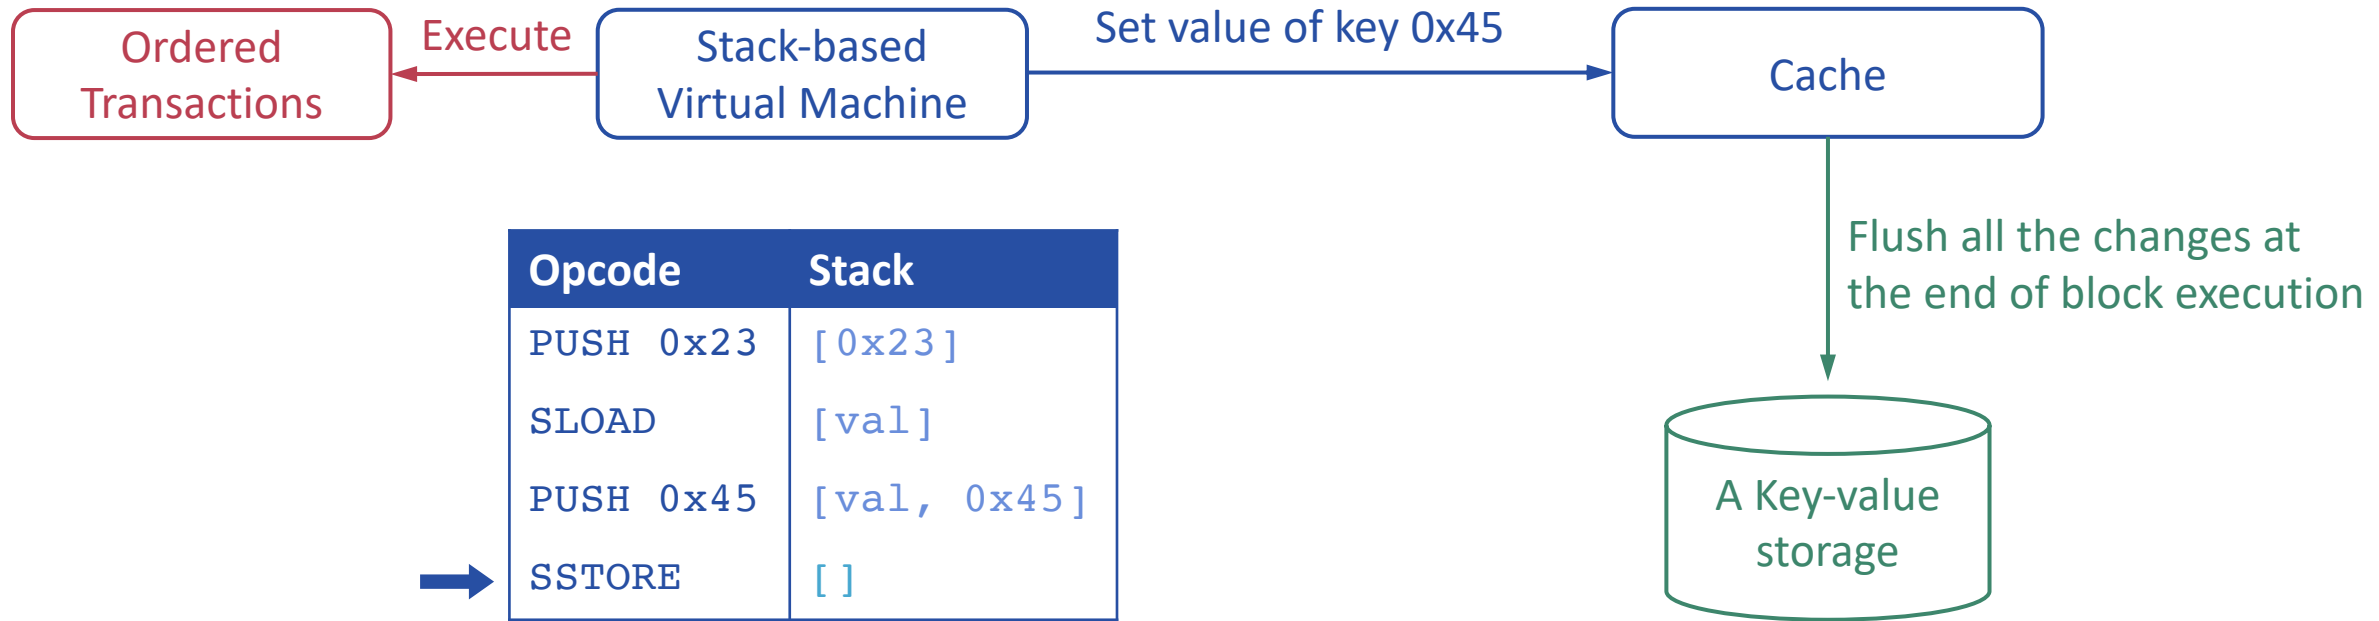
Next bottleneck



Architecture of Blockchain Execution Layer



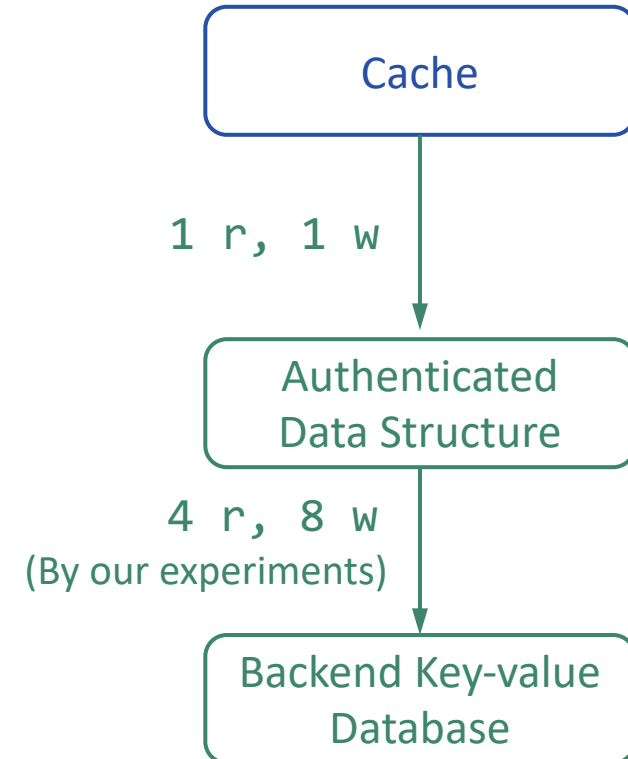
Architecture of Blockchain Execution Layer



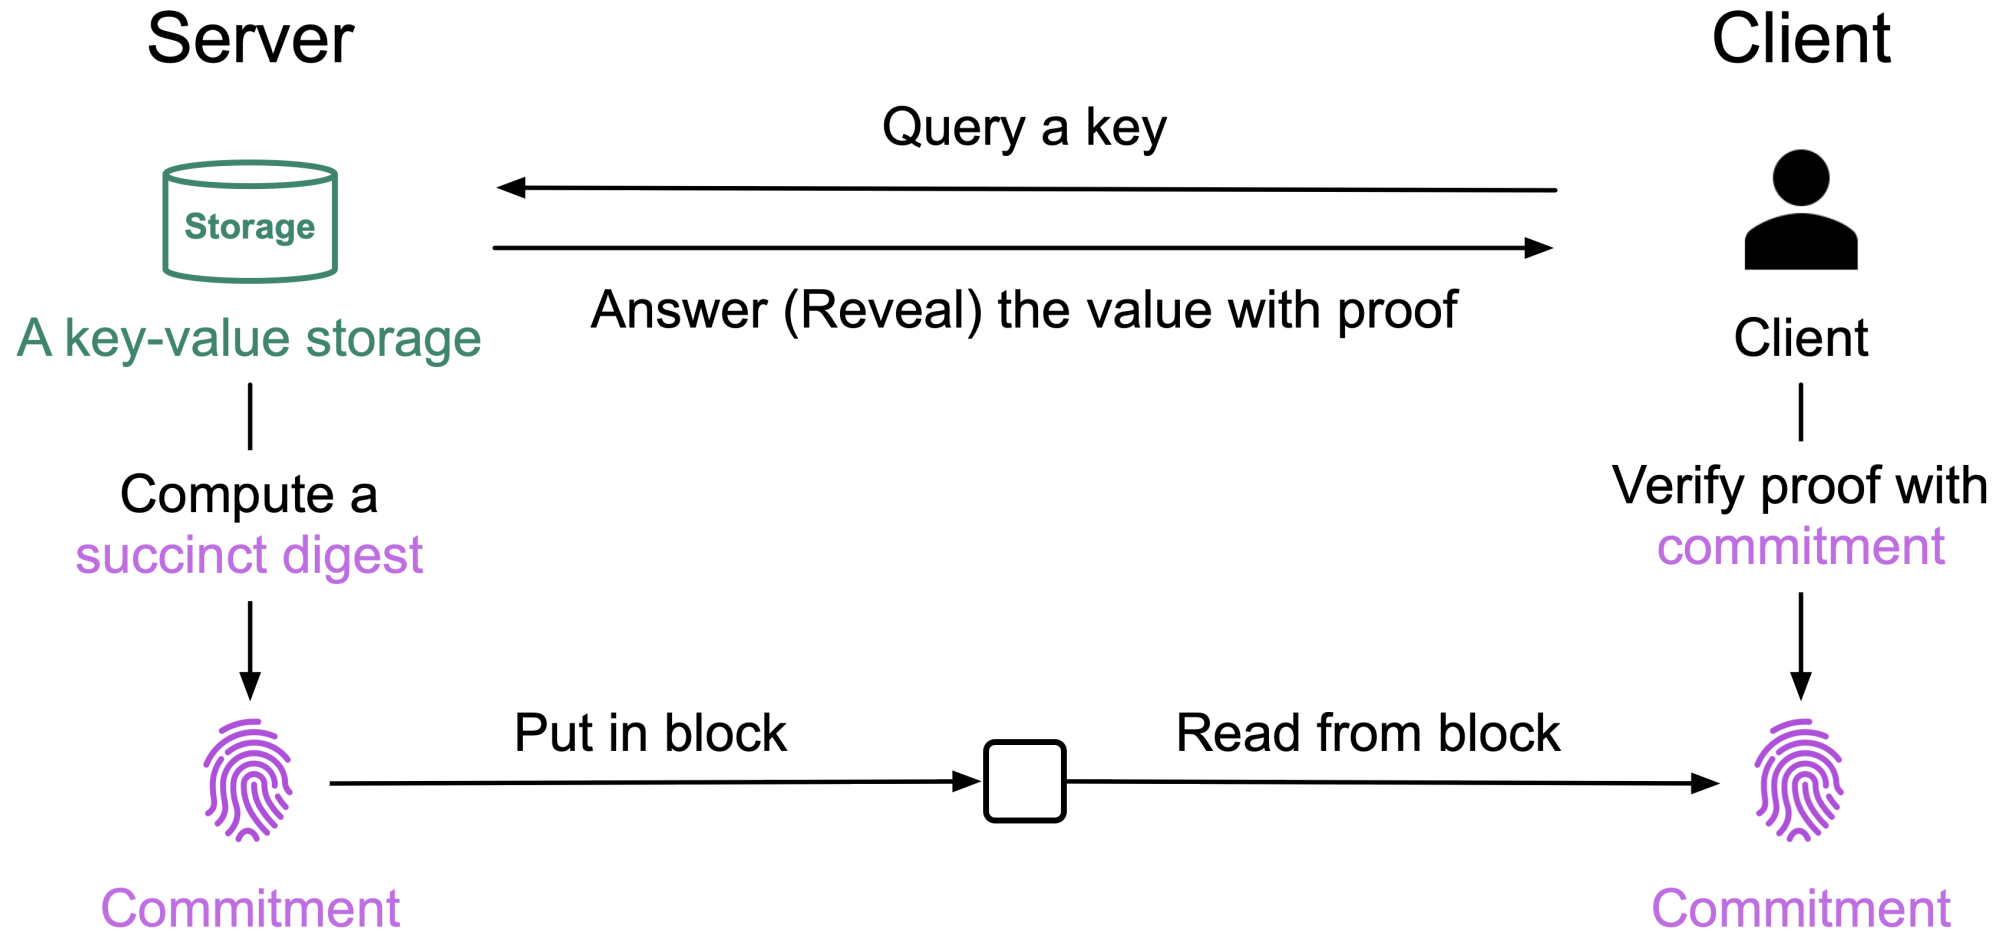
Architecture of Blockchain Execution Layer



Opcode	Stack
PUSH 0x23	[0x23]
SLOAD	[val]
PUSH 0x45	[val, 0x45]
SSTORE	[]



Authenticated Storage

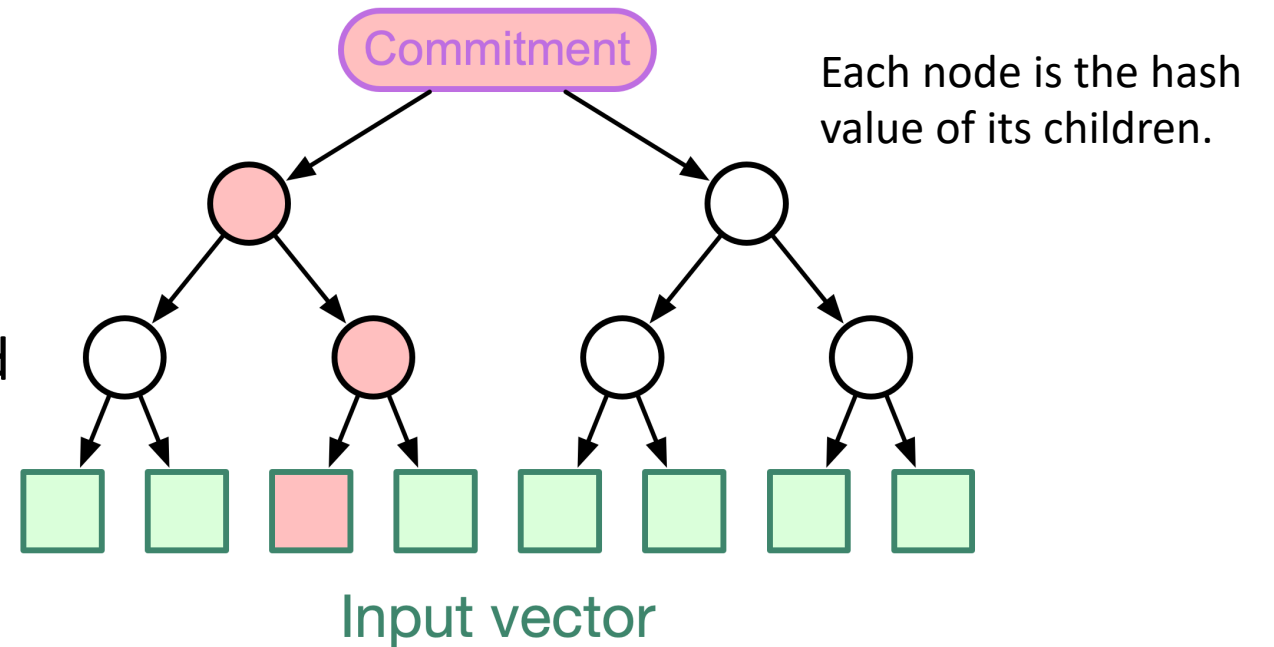


The Merkle Tree

Vector Commitment protocol: Merkle Tree

Variants for blockchain system: MPT, RainBlocks, LMPTs

- When an input element changes, the nodes along the path also changes.
- Each node is a key-value pair in backend
→ $O(\log n)$ read-write amplification



The Merkle Tree

Vector Commitment protocol: [AMT](#)

Variants for blockchain system: [LVMT \(our work\)](#)

- AMT removes the inner nodes and achieves $O(1)$ cost in maintaining commitment.


Commitment



Input vector

Challenges in using AMT

- Fast in complexity \neq fast in practice
 - AMT has slow cryptographic operations.
- AMT is not scalable.
 - Max capacity of AMT = Size of public parameters.
- Proof generation is Expensive
 - Maintaining data for generating proofs is also $O(\log n)$



Challenge 1: Costly cryptographic operations

In AMT, each time a value changes as,

$$a_i \rightarrow a_i'$$

the **commitment** adjusts accordingly

$$C \rightarrow C + (a_i' - a_i) \cdot G_i.$$

Precomputed Parameter (200 byte)

Elliptic Curve Multiplication (**92 μ s**)

Big Integer Subtraction (<0.01 μ s)

Elliptic Curve Addition (0.34 μ s)

Challenge 1: Costly cryptographic operations

In AMT, each time a value changes as,

$$a_i \rightarrow a_i' \quad (\text{Assumes } a_i' - a_i = 1)$$

the **commitment** adjusts accordingly

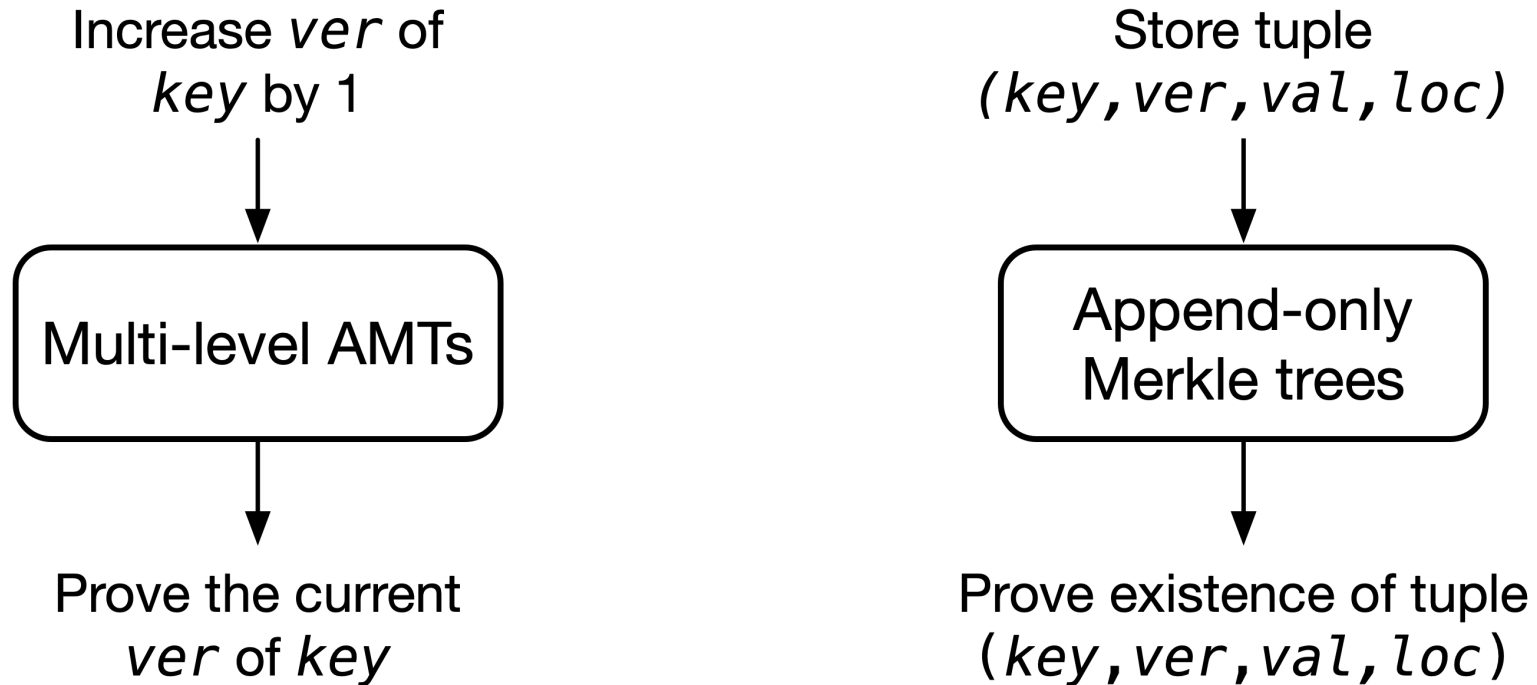
$$C \rightarrow C + 1 \cdot G_i.$$

Precomputed Parameter (200 byte)

Elliptic Curve Addition ($0.34 \mu s$)

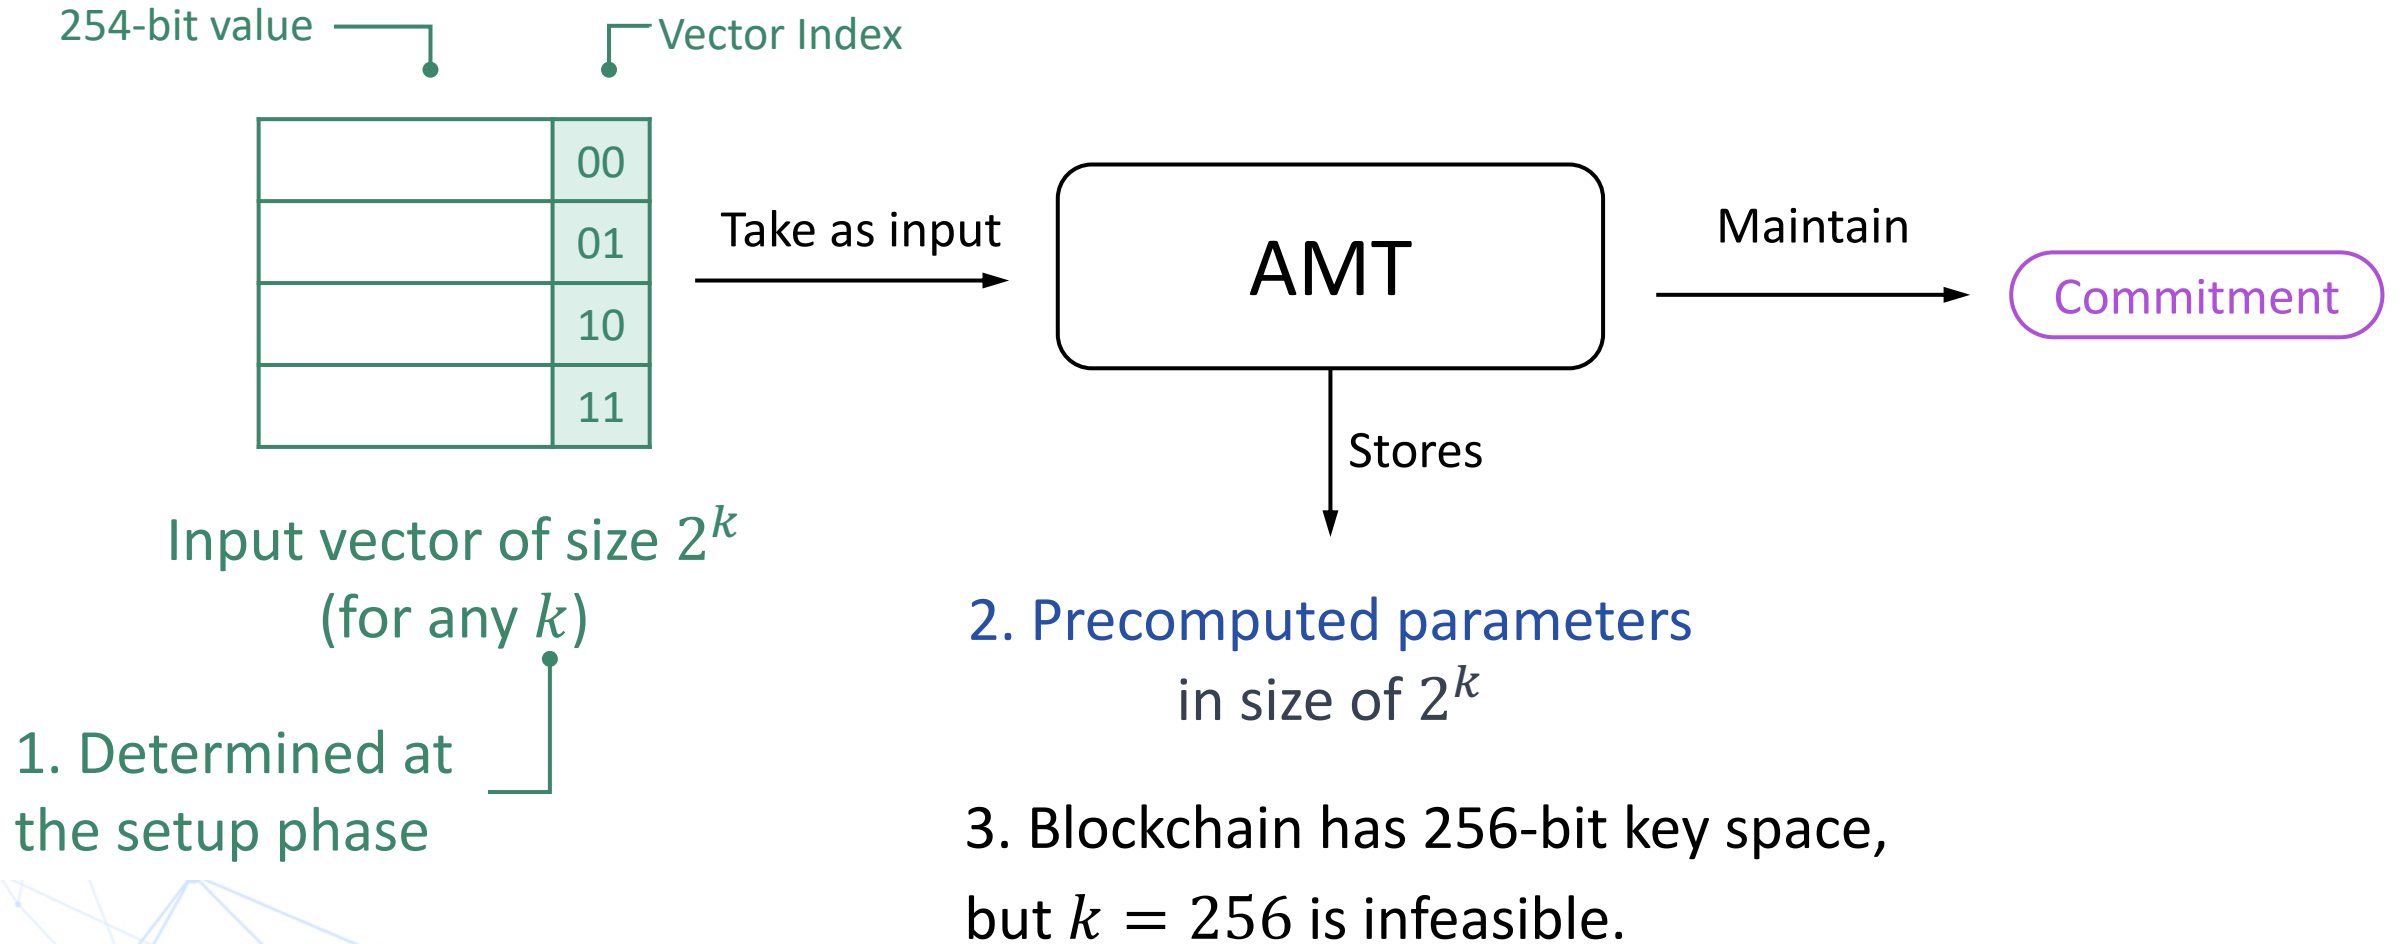
Solution 1: Version-based database

Set (*key, val*)




Prove *key*




Challenge 2: AMT is not scalable






Solution 2: Use multiple-level AMT.





Set (key, val)

- Increase version numbers in ,  and  by 1 and update commitments.
- Add the following tuples to the Merkle trees.

(key ,  , val , (3,2))

((2,1),  , )


(2 ,  , )

The slot index

The sub-AMT level allocating this key

Prove key

- Prove the version numbers with respect to the AMT commitment:



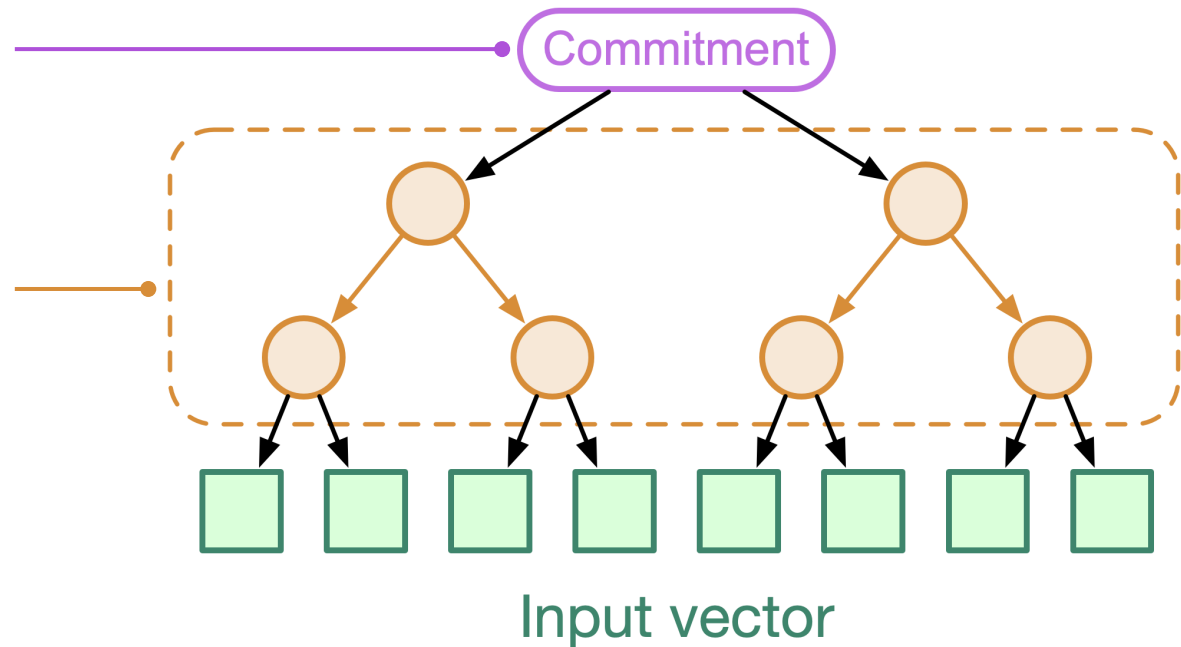
- Prove the existence of the left three tuples in Merkle trees to demonstrate the commitments at specified version numbers.



Challenge 3: Maintaining proof data incurs significant costs

Nodes *not serving clients* only need to maintain **commitment** in $O(1)$ time

Nodes *serving clients* maintain **auxiliary information** for generating proof in $O(\log n)$ time.



Solution 3: Proof Sharding

Commitment

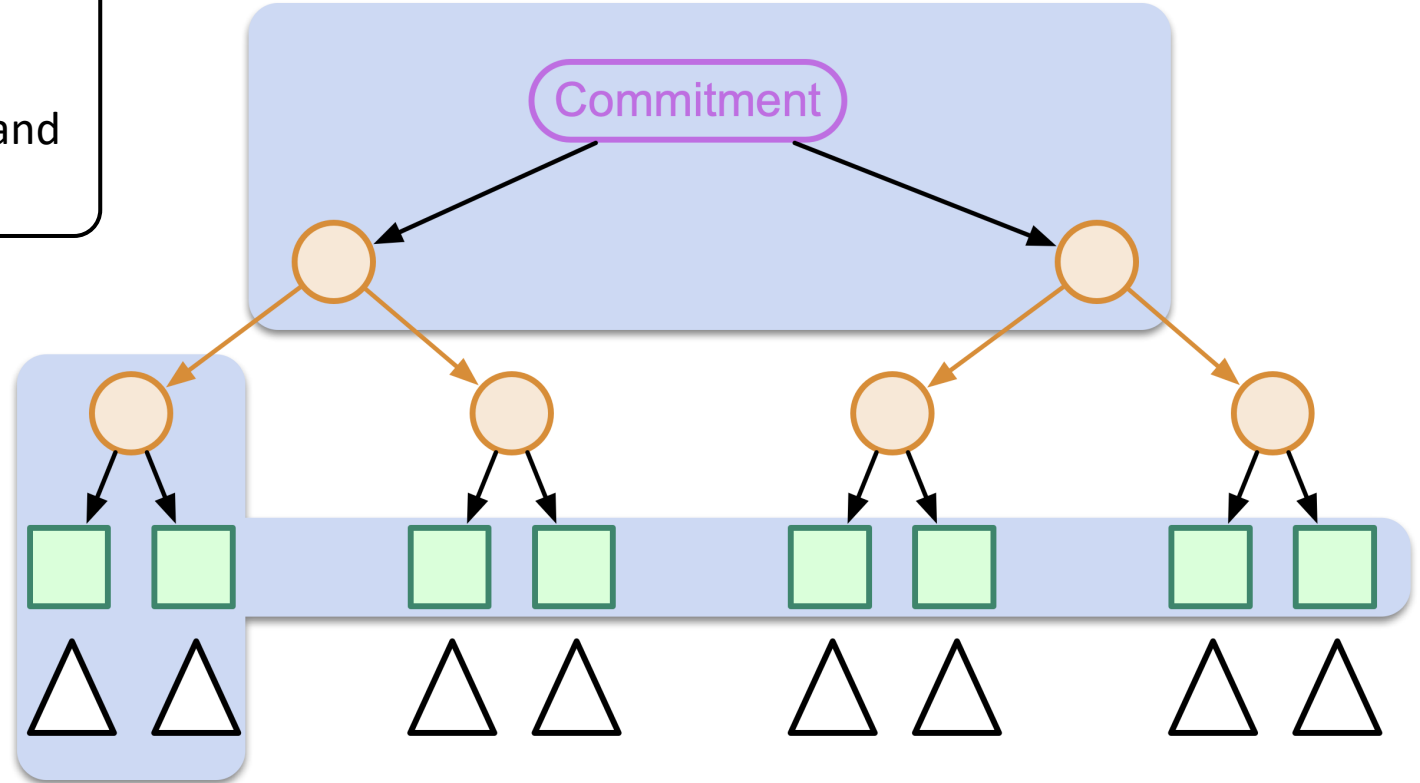


Consensus node
Don't maintain proof and
do not serve users

Auxiliary information
for generating proofs

Input vector
(Version numbers)

Sub-AMTs



RPC provider
Maintain proofs with a
cluster to serve users

Commitment

Proof shard



Commitment

Proof shard



Commitment

Proof shard





Commitment

Proof shard



Modular Authenticated Storage Benchmark Tool



1m

10m

100m

Random task on various ledger size (in million)

fresh

Randomly access new keys only

real

Real Ethereum trace

LVMT-r

Only maintains commitment

LVMT#

#: the fraction of proof shards

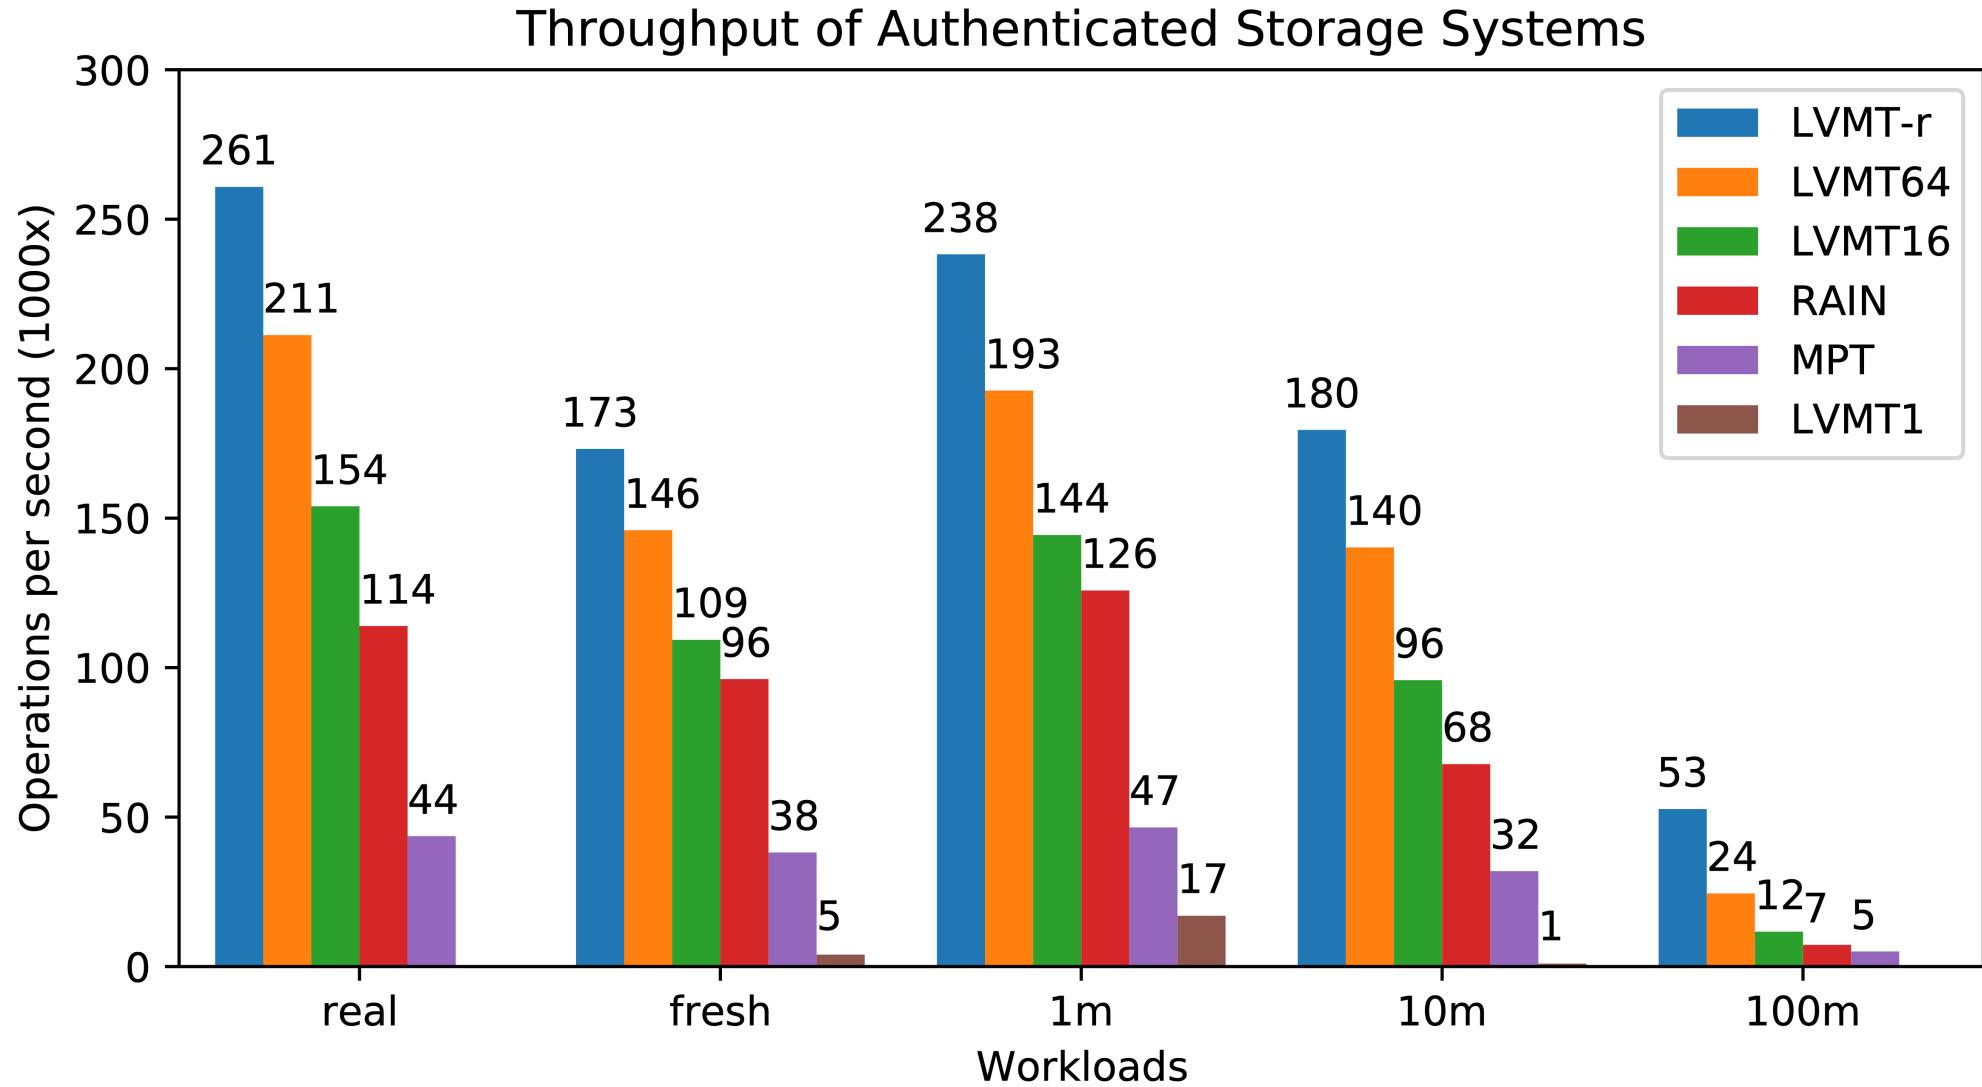
MPT

RAIN

LMPTs

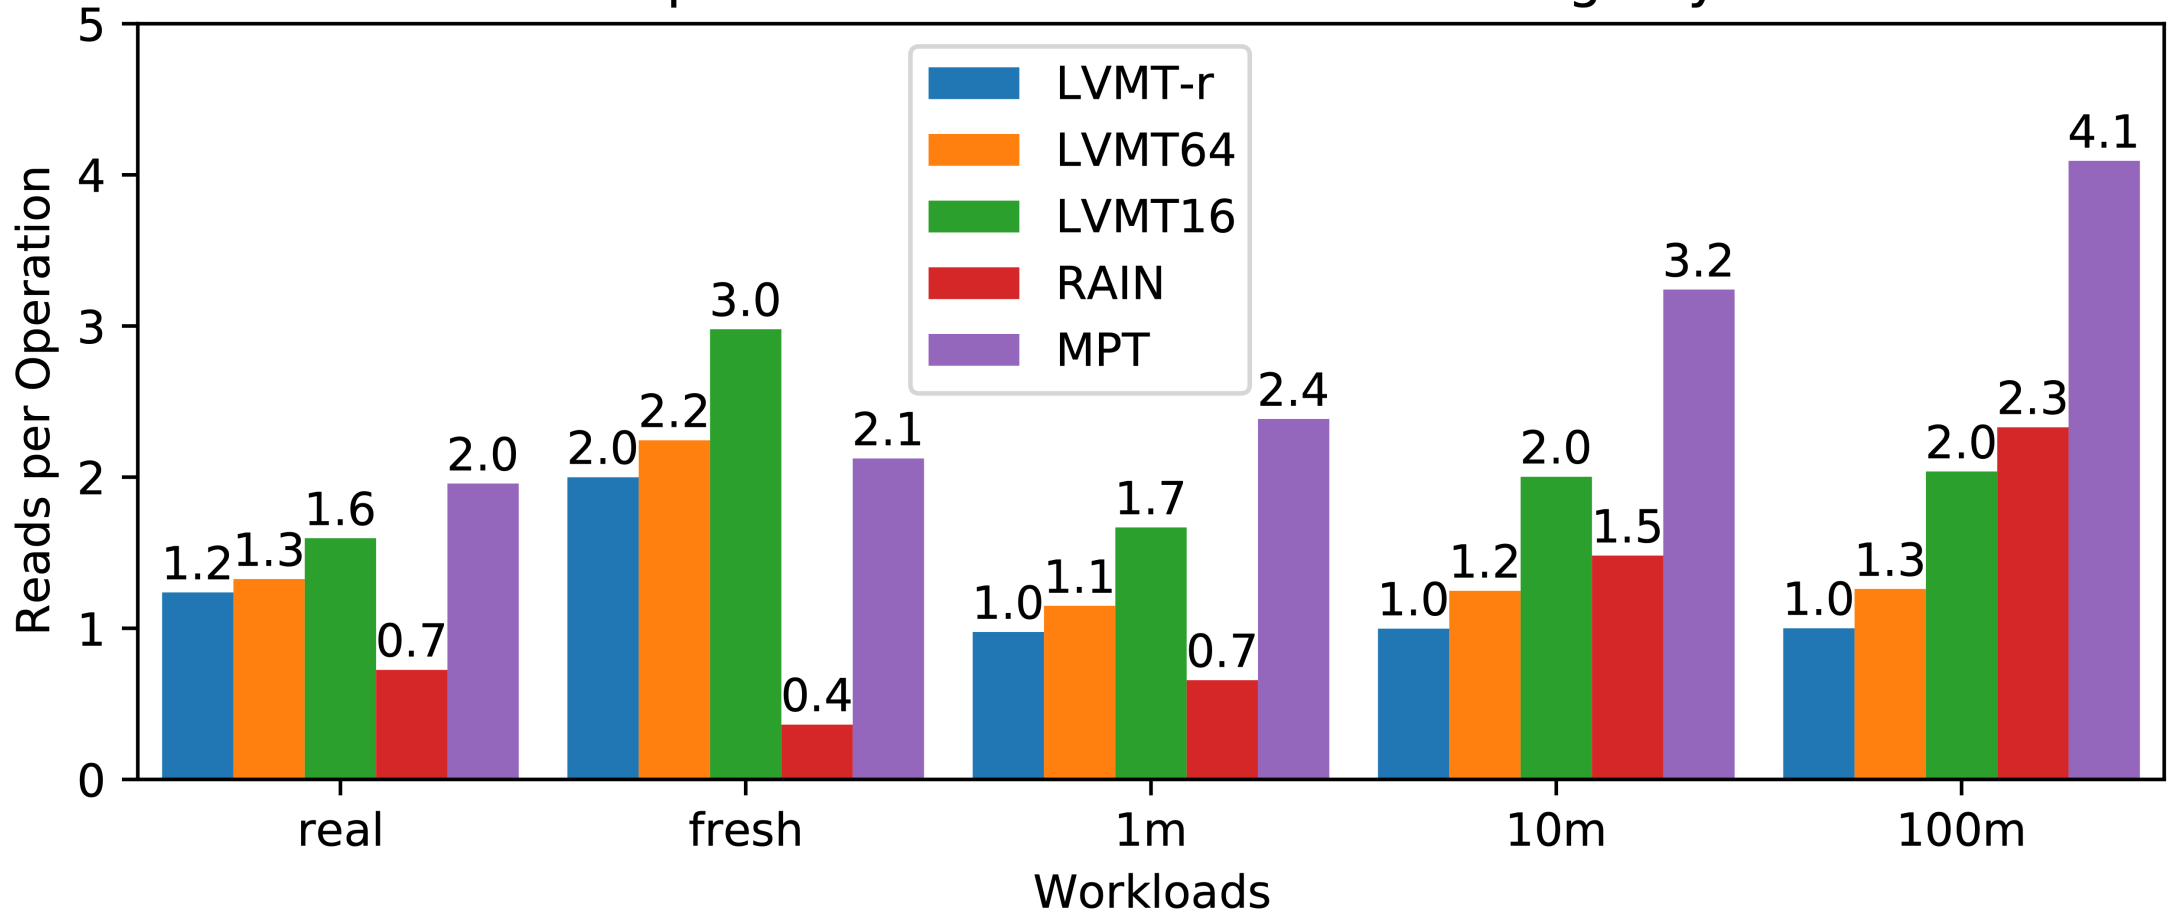
Baselines

Throughput on micro-benchmarks

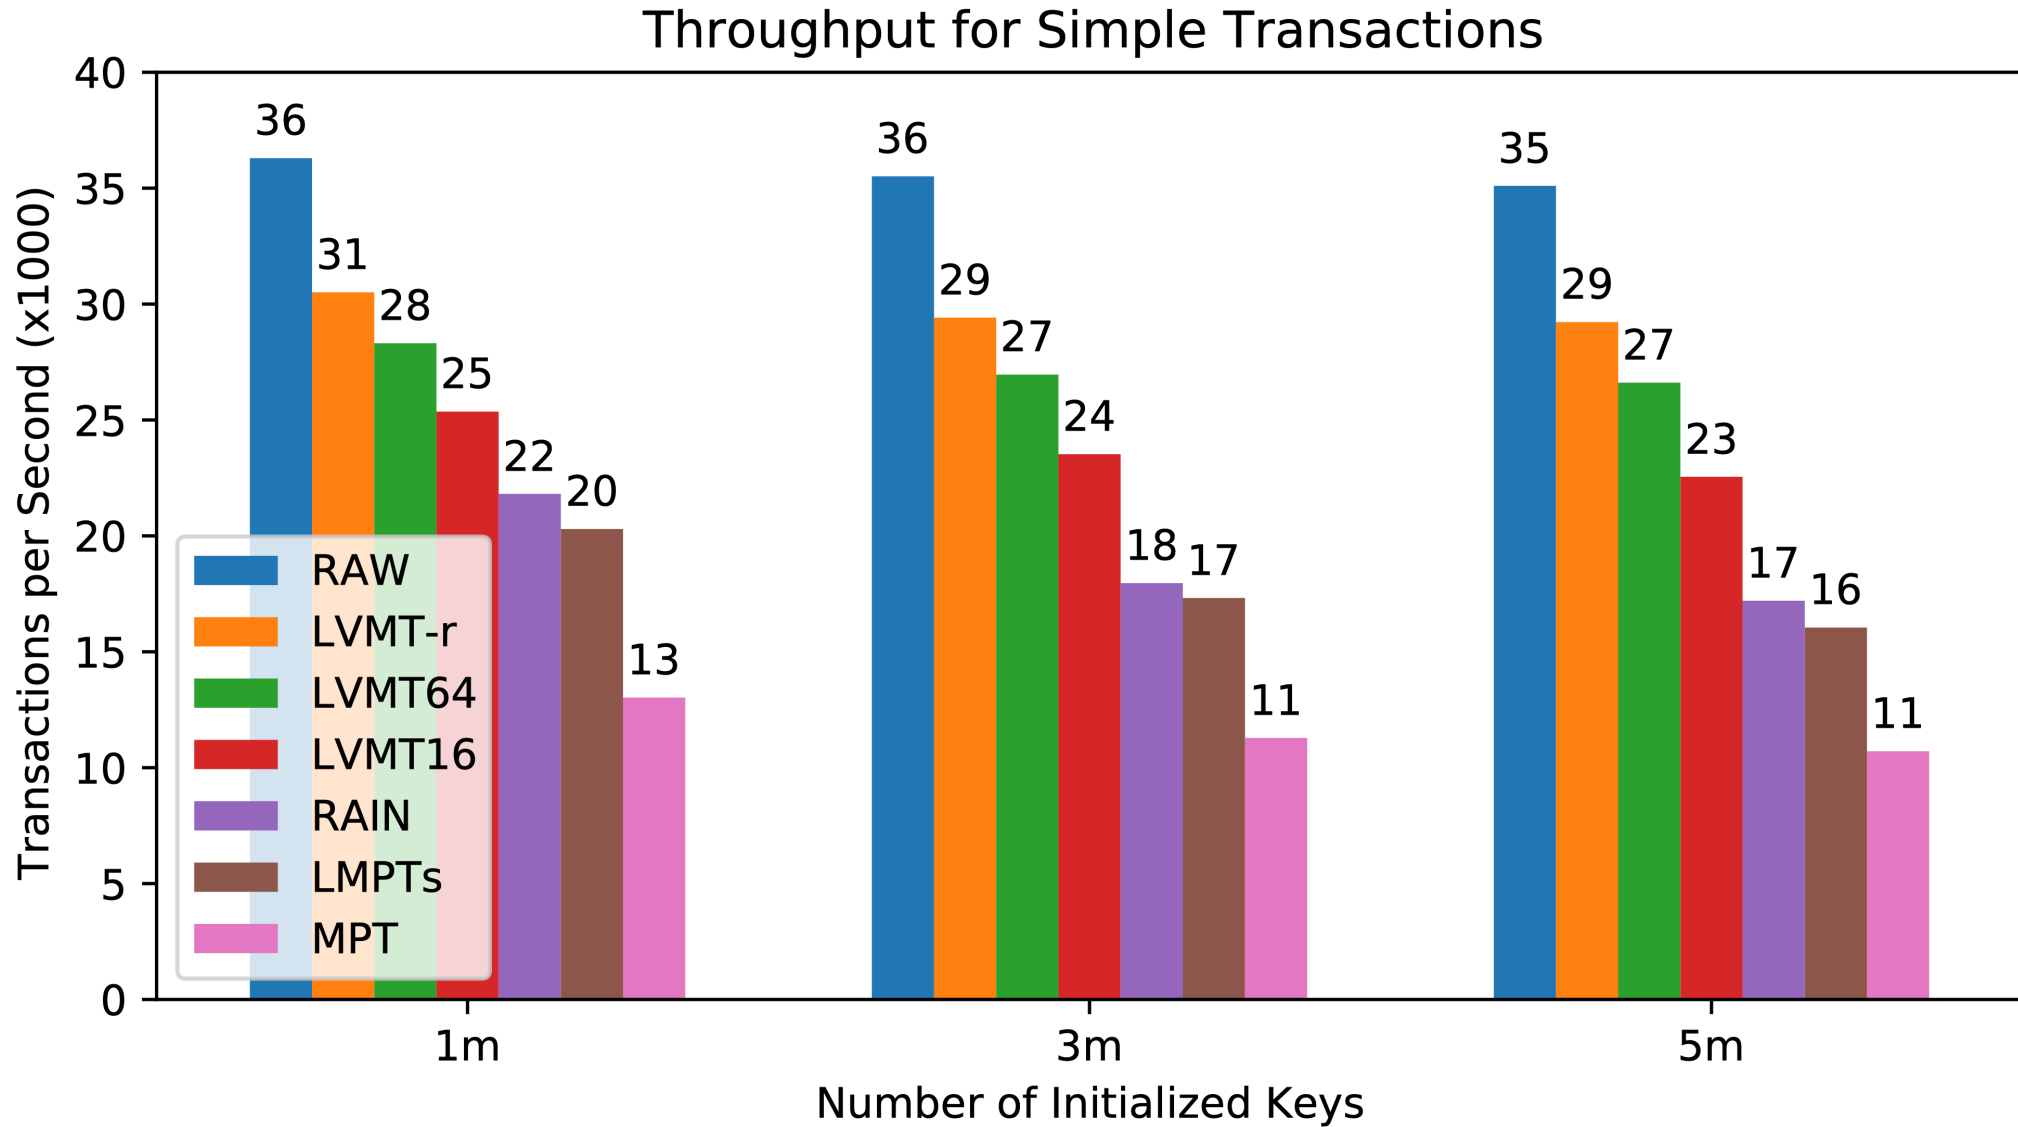


Read Amplification

Read Amplification of Authenticated Storage Systems

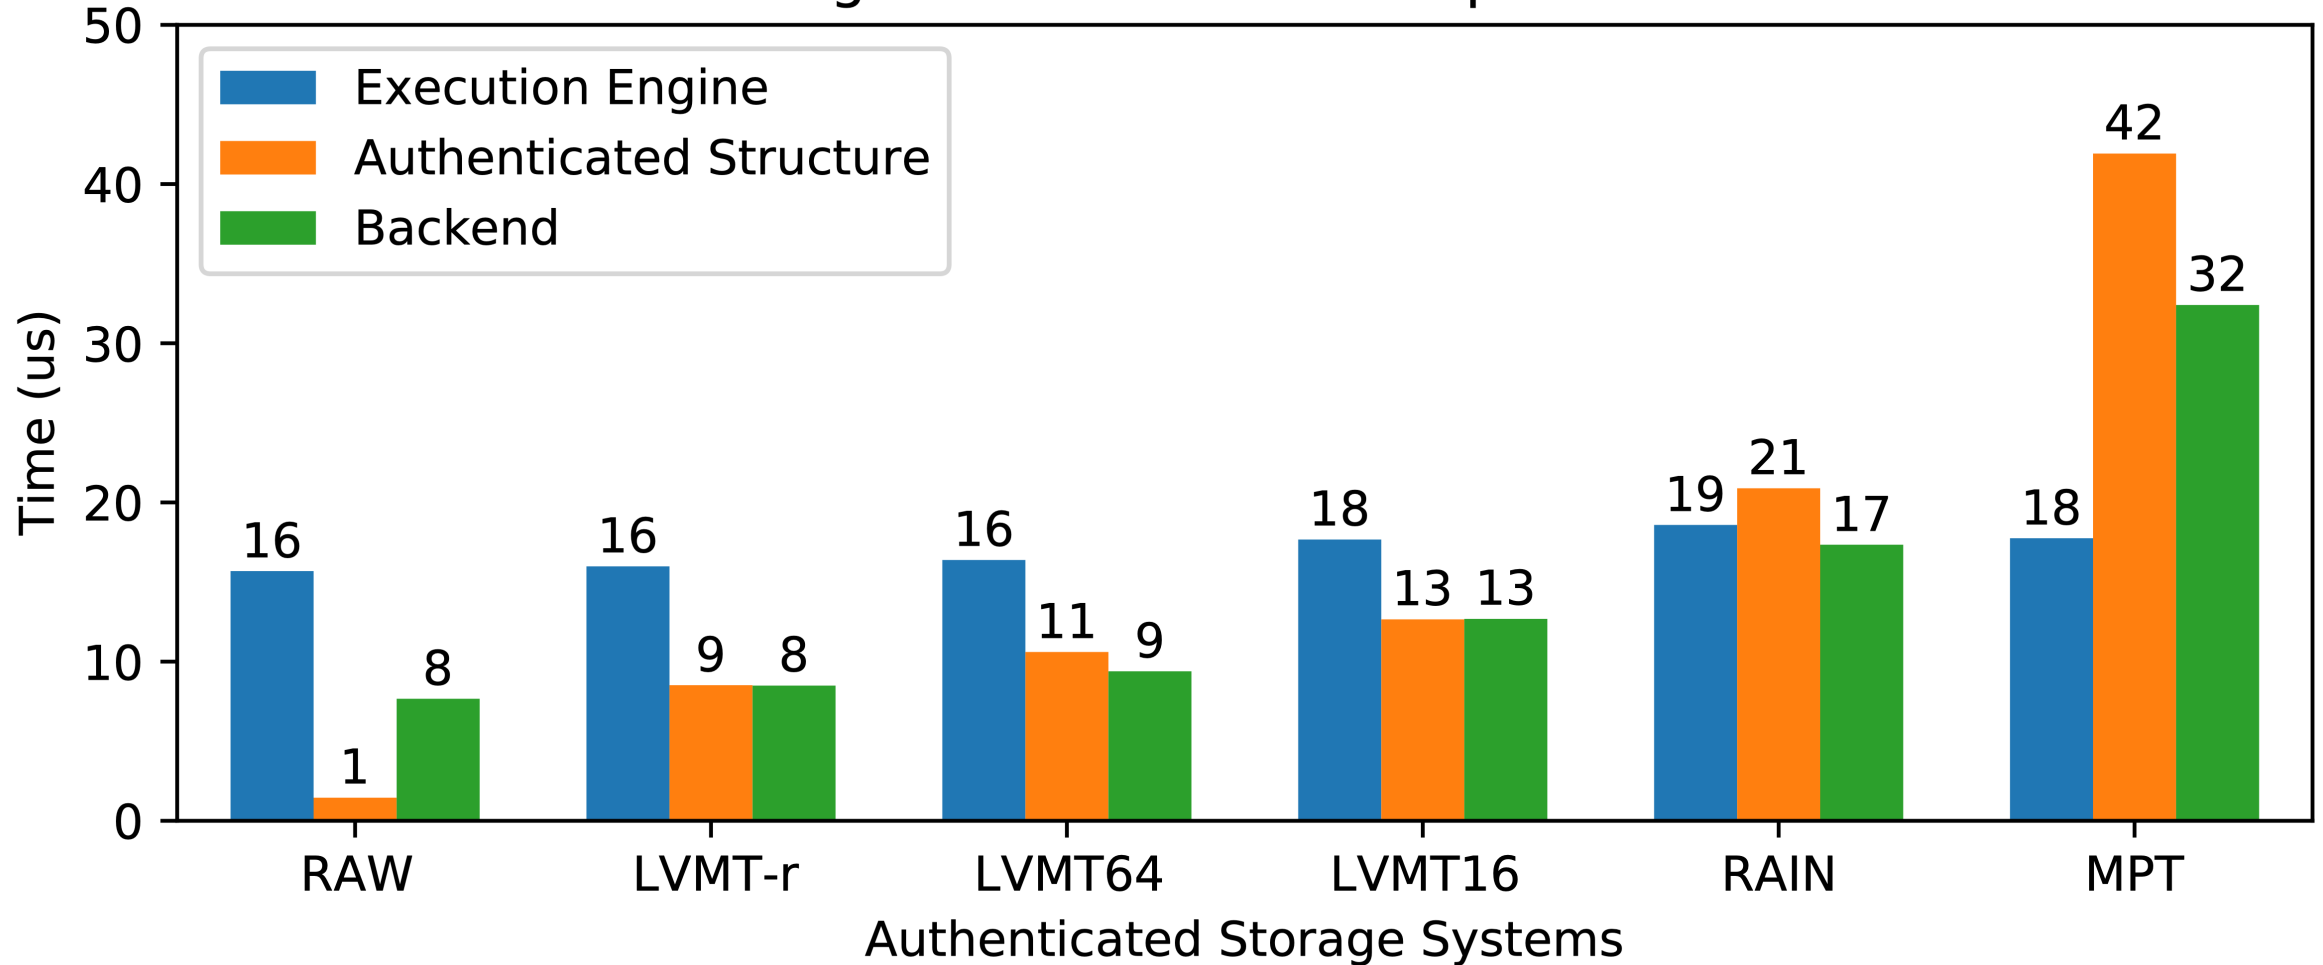


Throughput on a Blockchain Node



Time Usage Breakdown

Time Usage Breakdown for Simple Transactions



Conclusion

- LVMT utilizes the superior vector commitment protocol AMT, offering higher optimization potential.
- Through the version-based design, multi-level AMT, and proof sharding, LVMT addresses challenges effectively.
- LVMT enhances the execution throughput of a blockchain system by up to 2.7x.



Thank you and see you in Q&A

Email: lylcx2007@gmail.com

Github: <https://github.com/ChenxingLi/authenticated-storage-benchmarks>

<https://github.com/Conflux-Chain/conflux-rust/tree/asb-e2e>

A decorative network diagram in the bottom-left corner, consisting of a complex web of light blue lines connecting various nodes, some of which are highlighted with larger blue dots.