Saguaro: An Edge Computing-Enabled Hierarchical Permissioned Blockchain

Mohammad Javad Amiri[1] Ziliang Lai[2] Liana Patel[3] Boon Thau Loo[1] Eric Lo[2] Wenchao Zhou[4]

[1]University of Pennsylvania, [2]Chinese University of Hong Kong, [3]Stanford University, [4]Georgetown University

Scalable deployment of blockchain applications over wide-area networks

Edge network structure

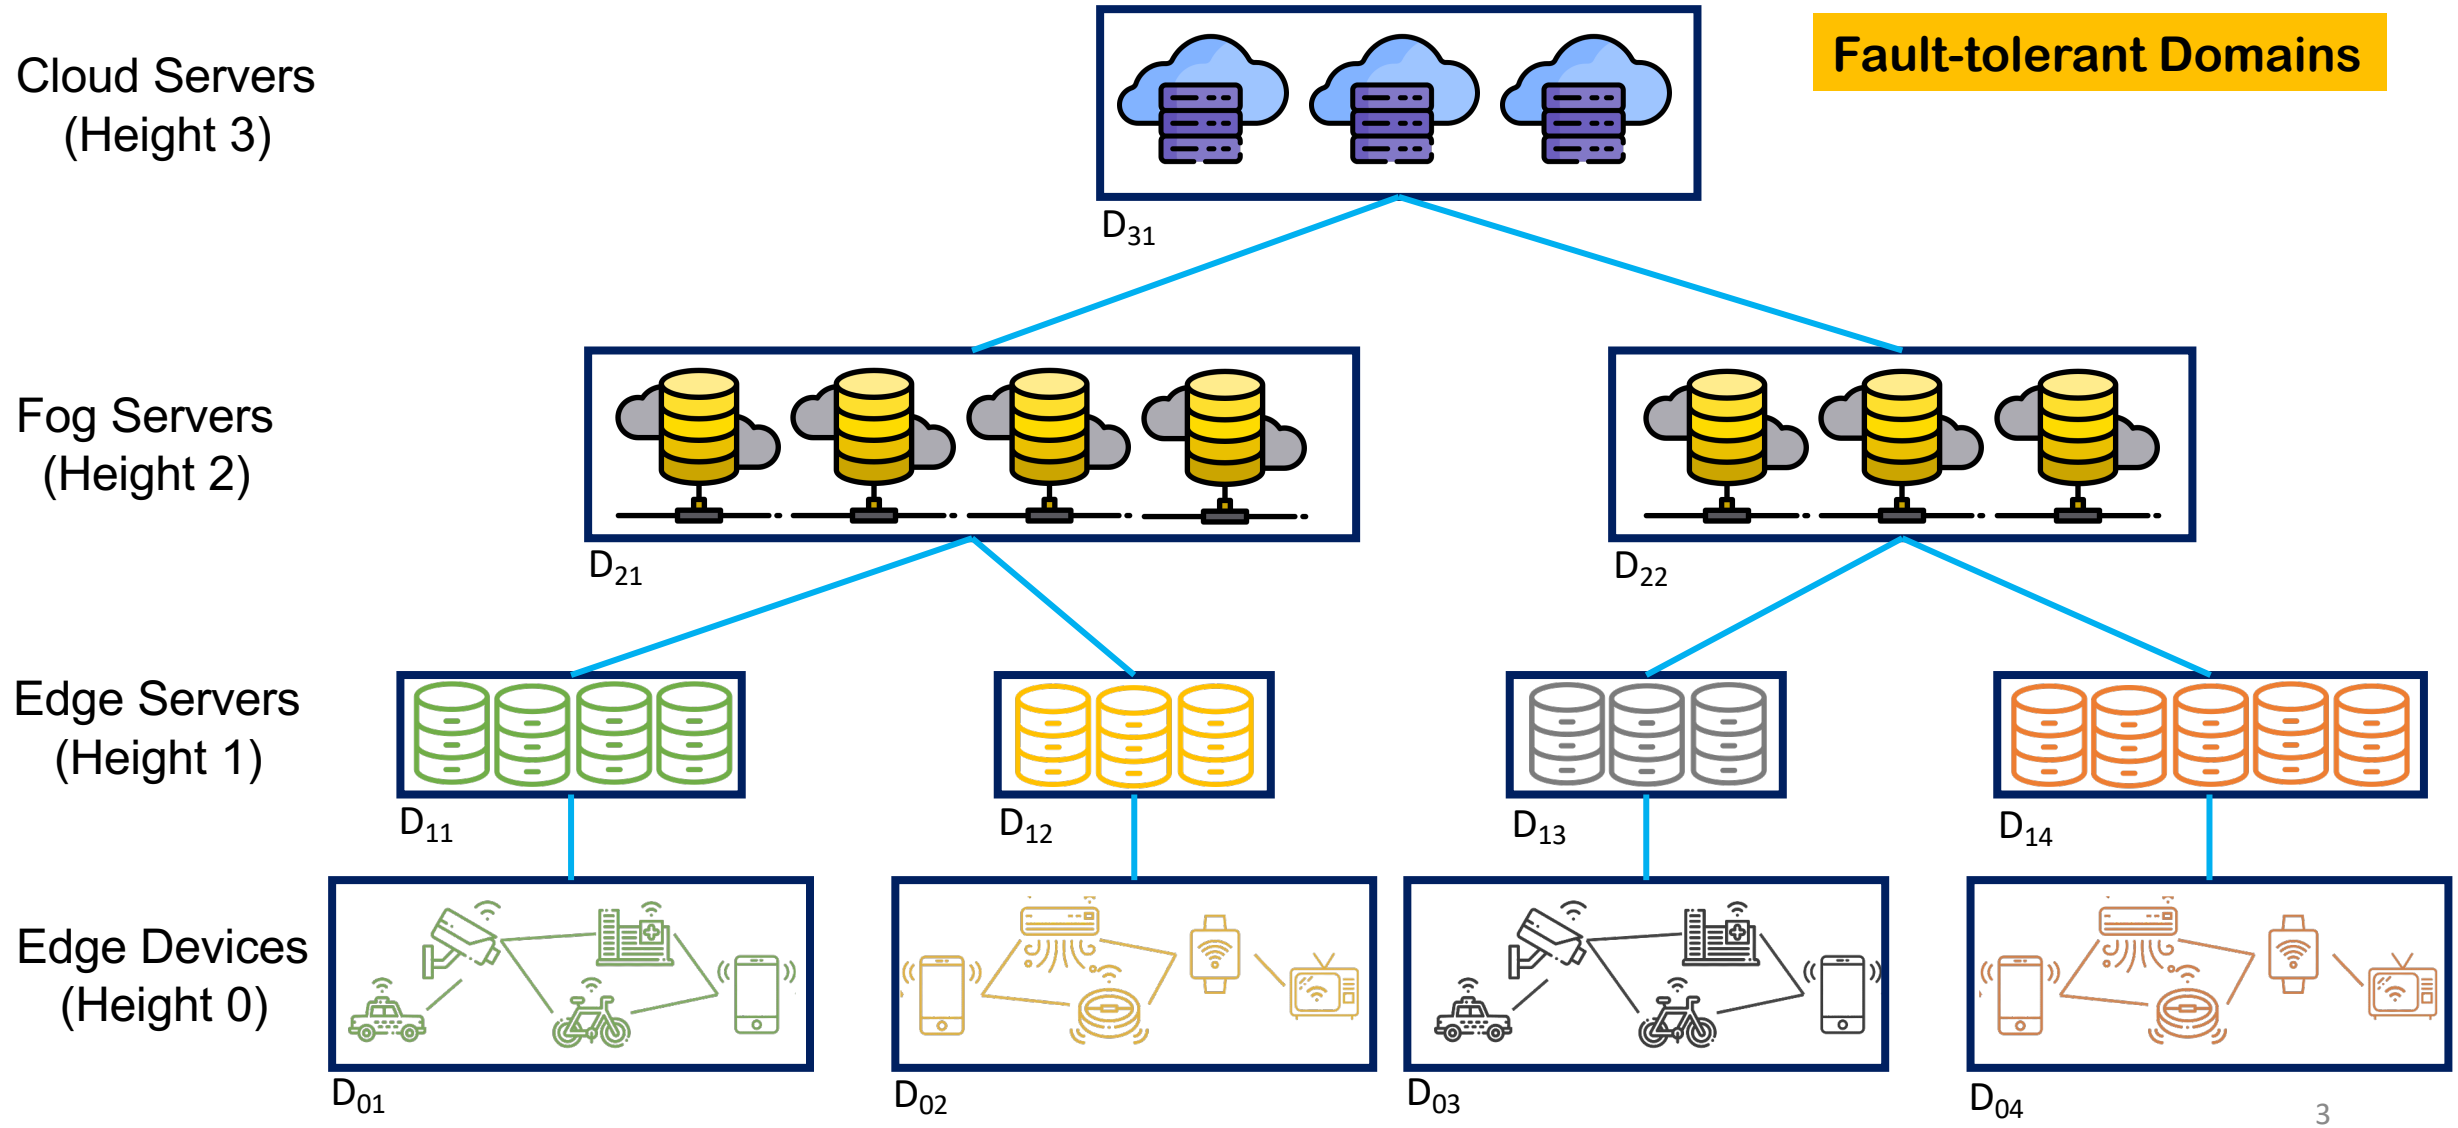

Saguaro

Processing cross-domain transactions using a coordinator-based approach by relying on the lowest common ancestor of all involved domains.

Aggregating data by propagating (a summarized version of) the ledgers up the hierarchy.

Optimistically processing cross-domain transactions and rely on higher-level nodes to detect inconsistencies.

Supports the mobility of nodes by relying on edge servers in the local and remote height-1 domains.

Scalability over wide-area networks

- Coordinator-based sharding (e.g., AHL [SIGMOD'19])
  - Runs two-phase commit on top of BFT
    - The coordinator node (cluster) is either close to clients or the data shards
    - Cannot avoid slow network links when cross-shard transactions take place.

- Flattened sharding (e.g., SharPer [SIGMOD'21])
  - Run consensus among all nodes of all involved shards
    - Requires several rounds of communication over high-latency low bandwidth Internet links.

- Full replication of the entire ledger on every cluster (e.g., GeoBFT [VLDB'20])
  - Clusters process disjoint sets of transactions and sync after each round
    - Shifts the wide-area communication from running the consensus protocol across data centers to ledger synchronization messages over a wide-area network.

Coordinator-based consensus protocol

- Transactions:
  - Initiated by edge devices (height-0)
  - Executed by edge servers in height-1 domains

- Transaction types:
  - Internal: access records within a single domain
  - Cross-domain: access records across different height-1 domains

- Consensus protocol:
  - Internal: depending on the failure model of nodes (CFT vs BFT)
  - Cross-domain: coordinator-based protocol

Internal transactions

Crash failure: fail by stopping, no malicious behavior

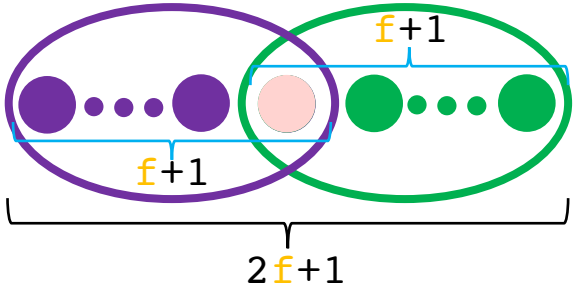

(Multi-)Paxos

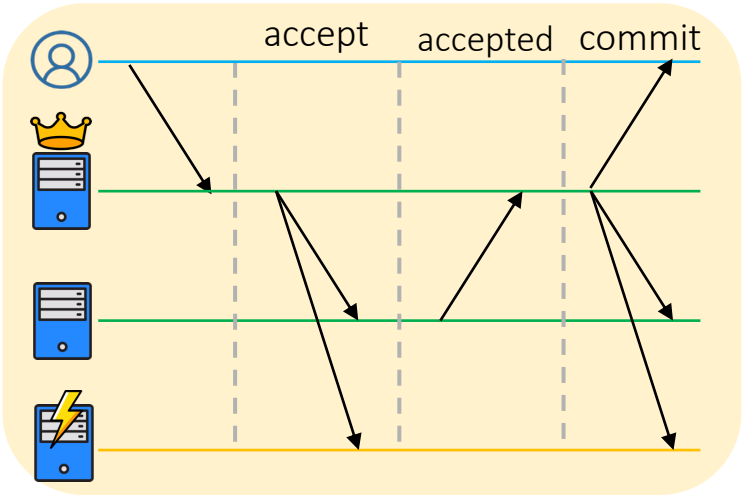

Byzantine failure: exhibit arbitrary, potentially malicious, behavior

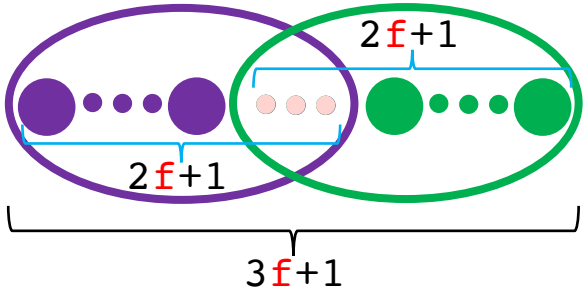

PBFT

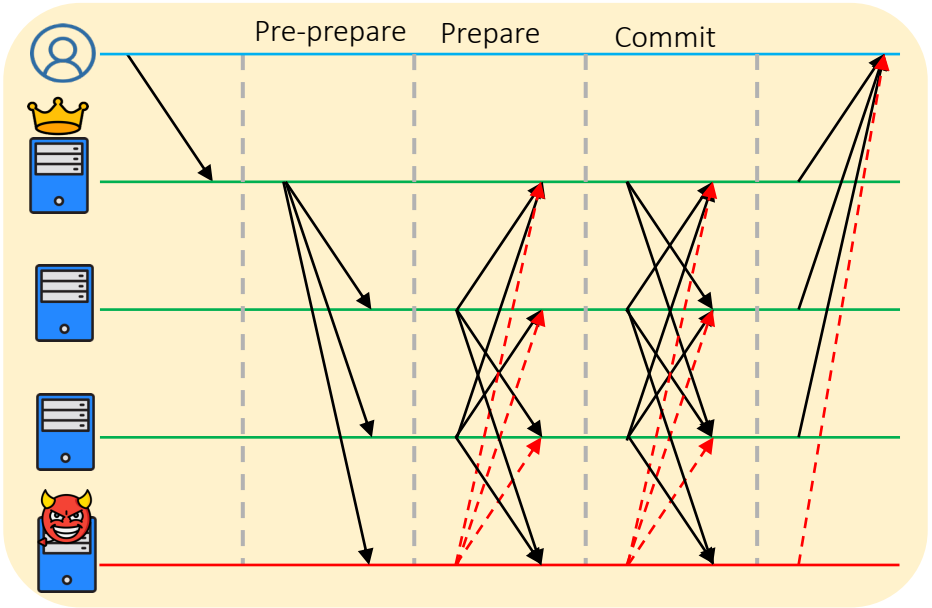

Coordinator-based cross-domain consensus

- Inspired by the traditional coordinator-based commitment protocols
- Coordinator: the Lowest Common Ancestor (LCA) of all involved height-1 domains
  - LCA domain has the optimal location to minimize the total distance

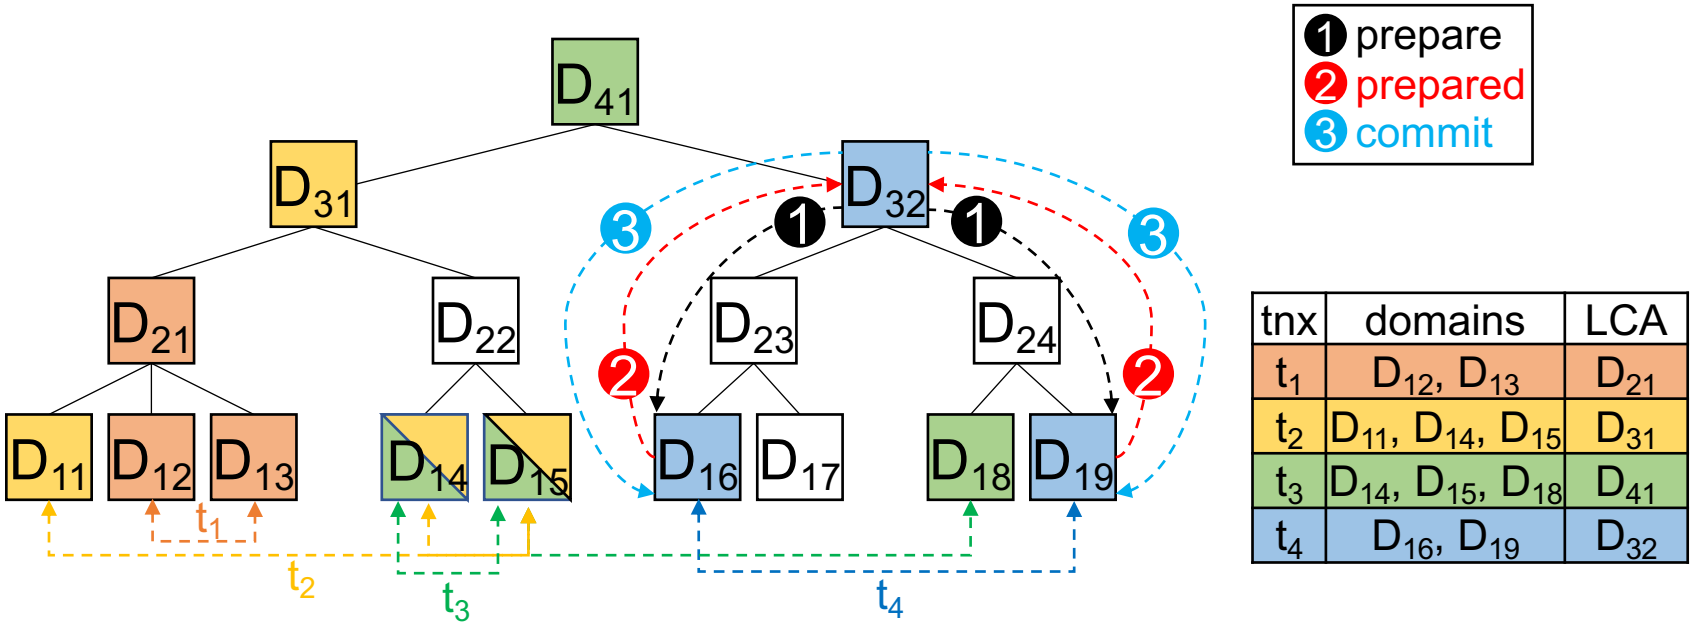

tnx	domains	LCA
t_1	D_{12}, D_{13}	D_{21}
t_2	D_{11}, D_{14}, D_{15}	D_{31}
t_3	D_{14}, D_{15}, D_{18}	D_{41}
t_4	D_{16}, D_{19}	D_{32}

An example of Saguaro blockchain ledger

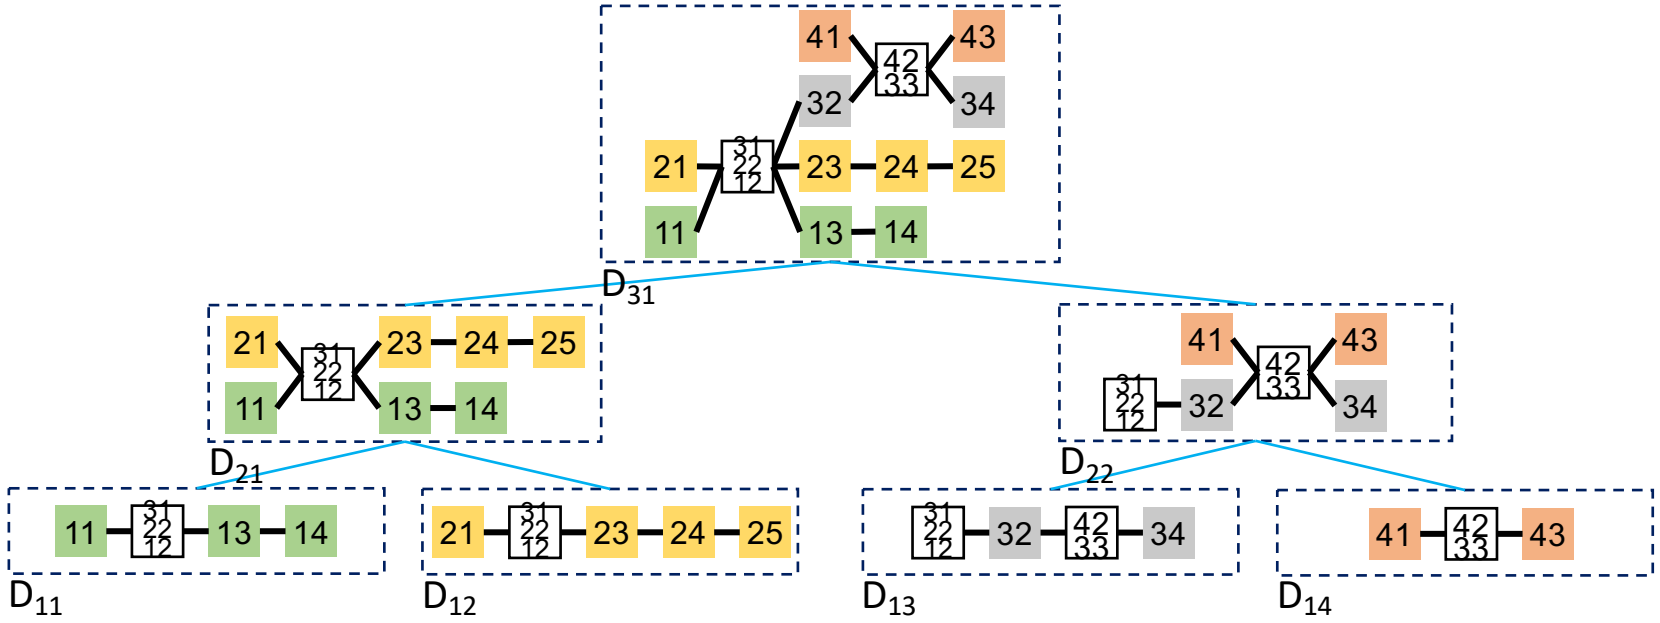

Lazy propagation of blockchain ledgers

- Perform data aggregation over transactions executed by edge servers in height-1
- Each domain maintains (a summarized version of) their child domains data.
- Block message: Transactions + an abstract version of the state updates

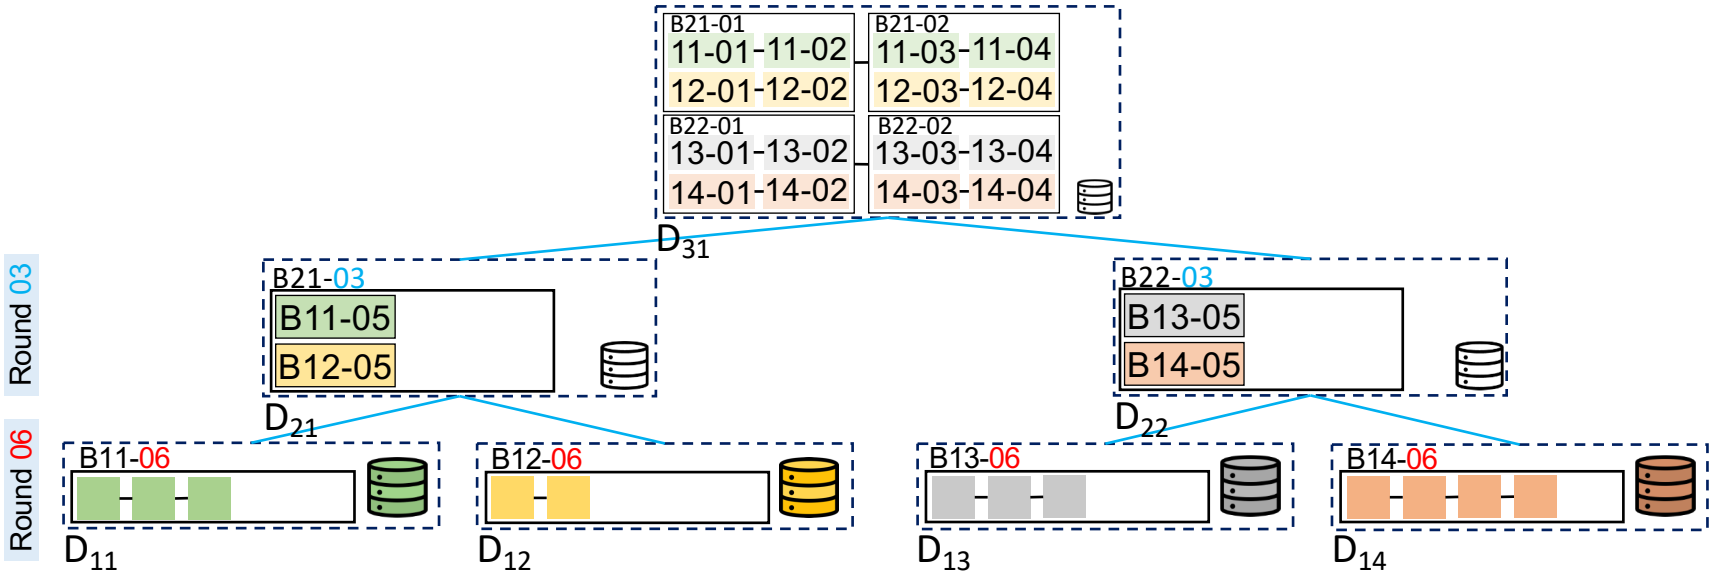

Optimistic consensus protocol

- Each involved height-1 domain optimistically commits a cross-domain transaction independent of other involved domains
- Keep a list of data-dependent transactions for each cross-domain transaction

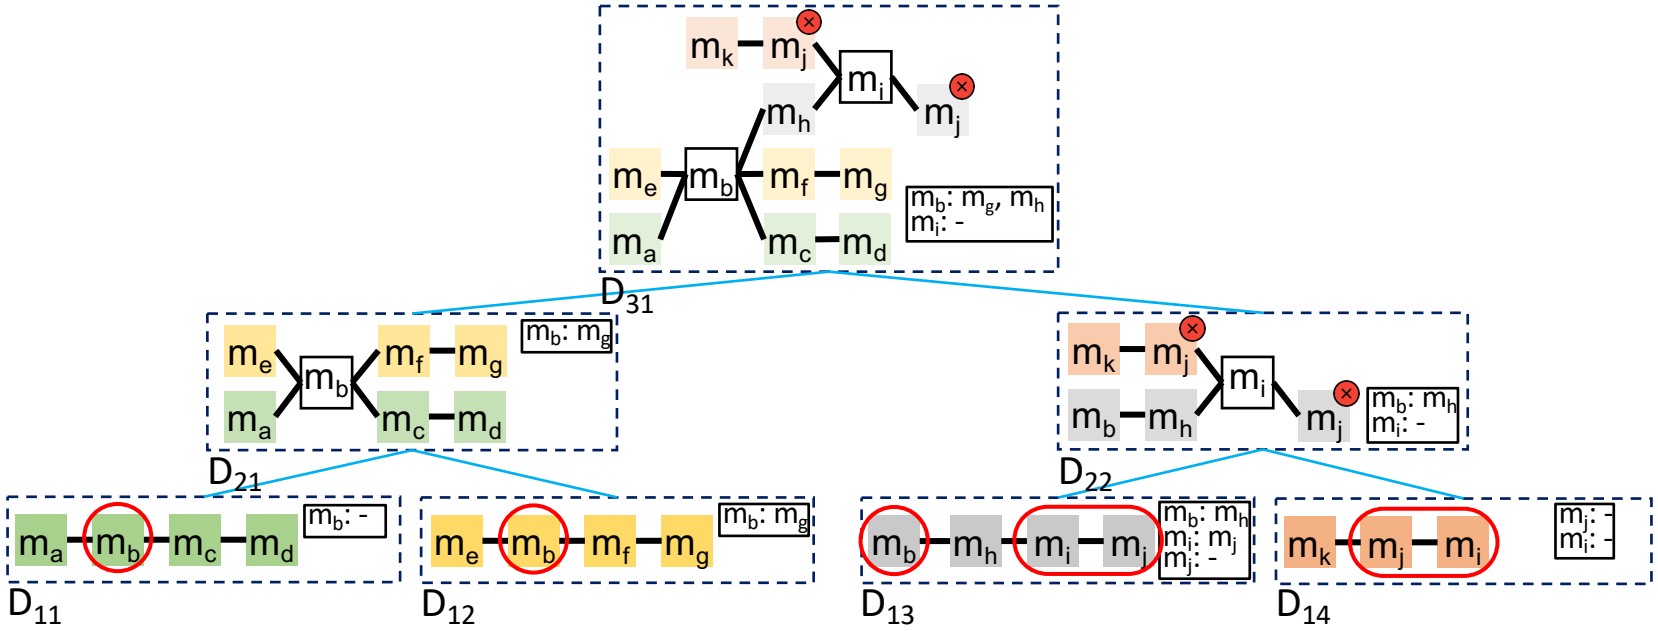

Mobile consensus

- What if a node moves from a local to a remote domain?
  - The remote domain does not have access to the state of the mobile node

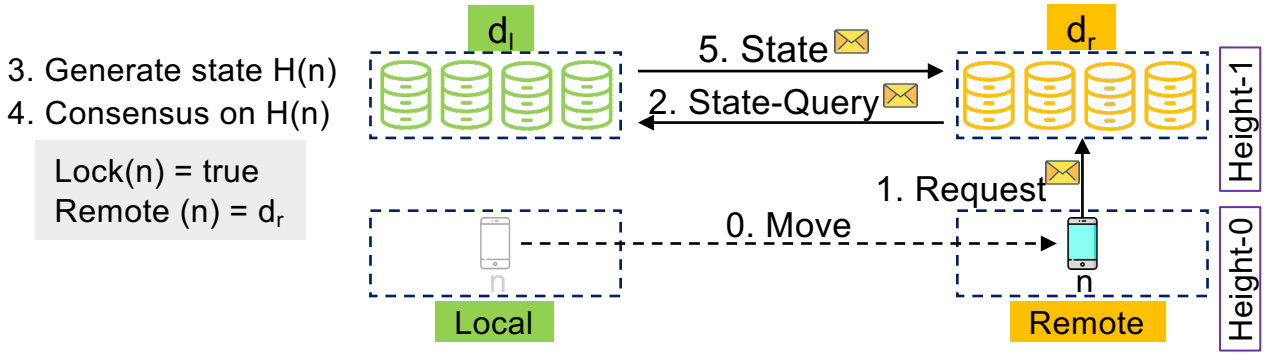

Experimental settings

- Platform: Amazon EC2
- Measuring performance
  - Throughput & Latency
- Application:
  - Micropayment
- Network:
  - A typical four-level edge network (f=1 in each cluster)
- Systems:
  - AHL [SIGMOD'19]
  - SharPer [SIGMOD'21]
  - Saguaro: Coordinator-based
  - Saguaro: Optimistic (contention: 10%, 50%, 90%)

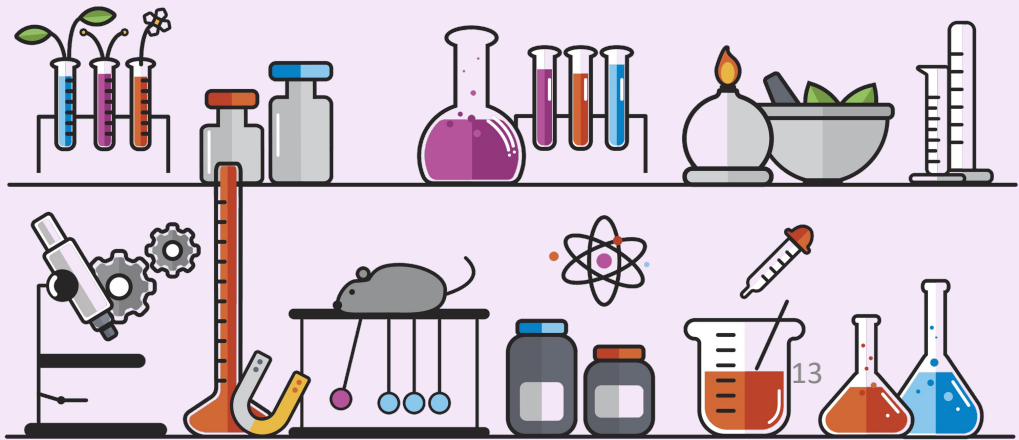

Cross-domain transactions (crash-only)

Domains: Frankfurt, Milan, London, and Paris (RTT: 9-25 ms)

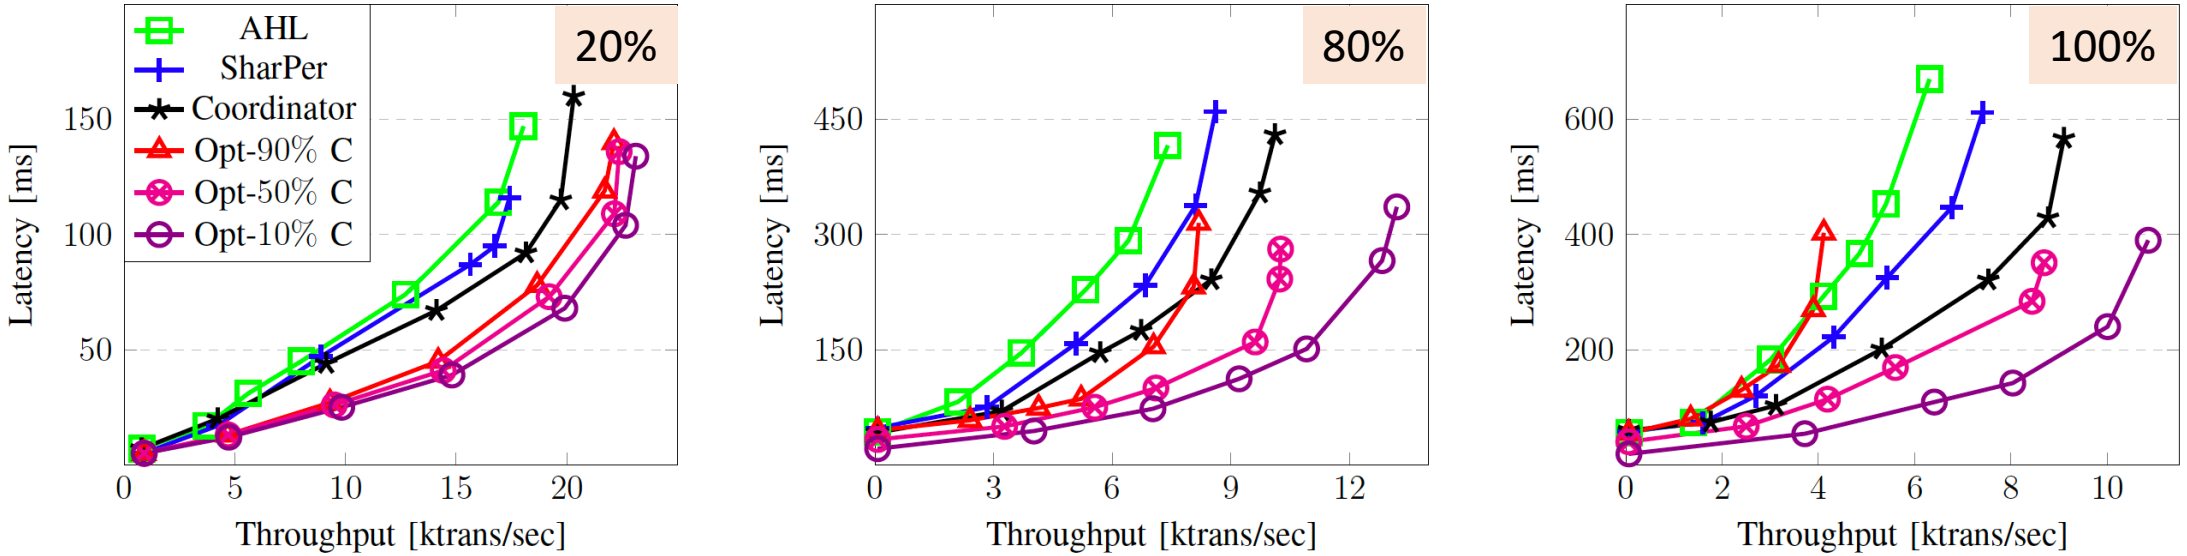

<u>20% cross-domain transactions:</u>

- Optimistic approach with 10% contention shows the best performance
  - only 0.16% of transactions appended to the ledgers in an inconsistent order
- Coordinator-based approach: 17% higher throughput compared to AHL

<u>80% & 100% cross-domain transactions:</u>

- Larger performance gap between the coordinator-based approach and existing systems

Cross-domain transactions (Byzantine)

Domains: Frankfurt, Milan, London, and Paris (RTT: 9-25 ms)

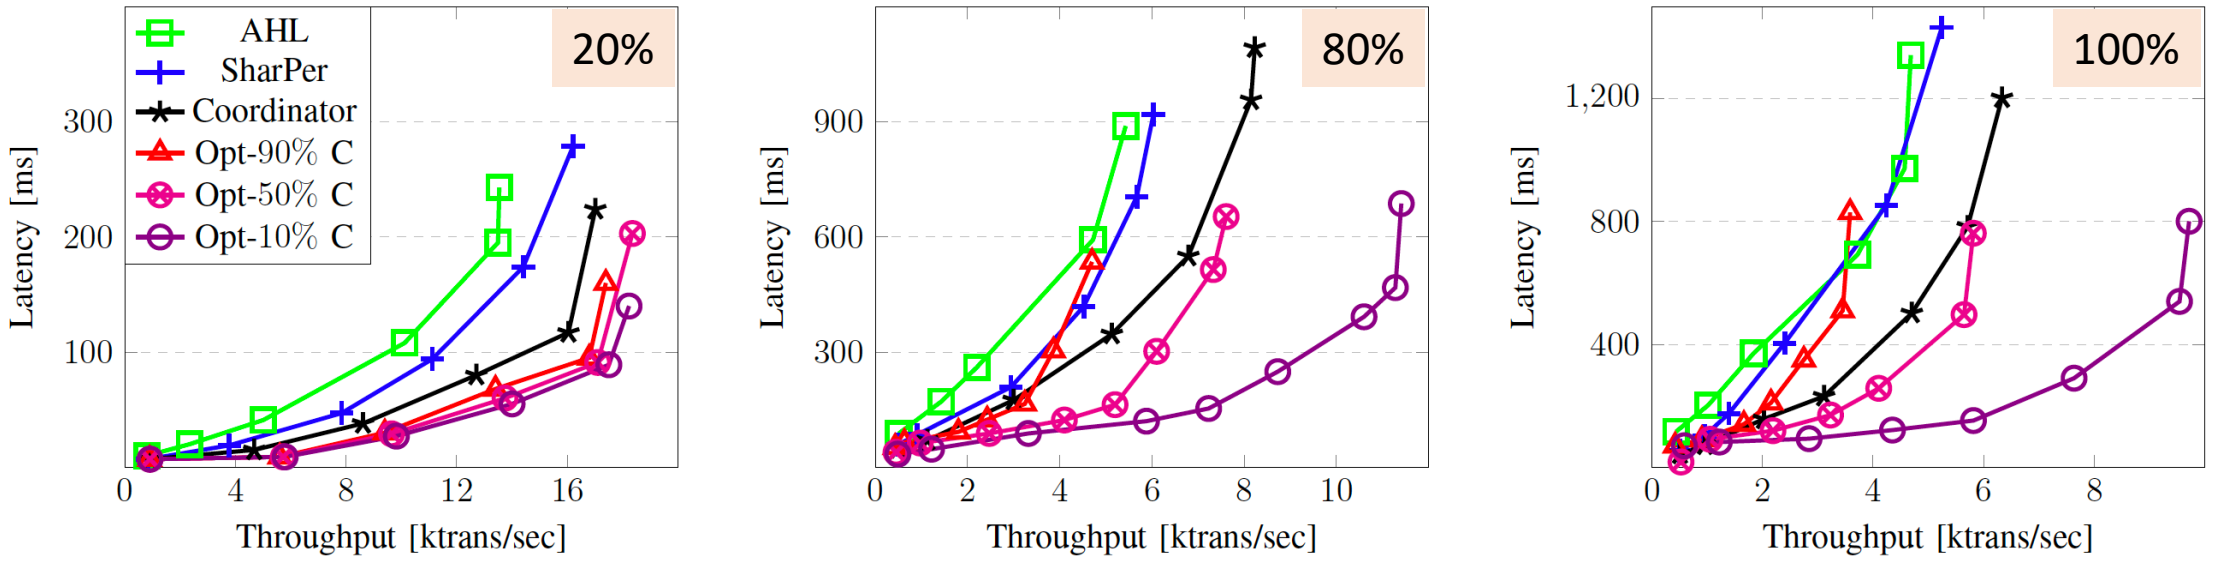

- Similar behavior, with lower throughput and higher latency

Mobile devices

A mobile node initiates 10 transactions within the remote domain before moving back to its local domain.

Crash-only

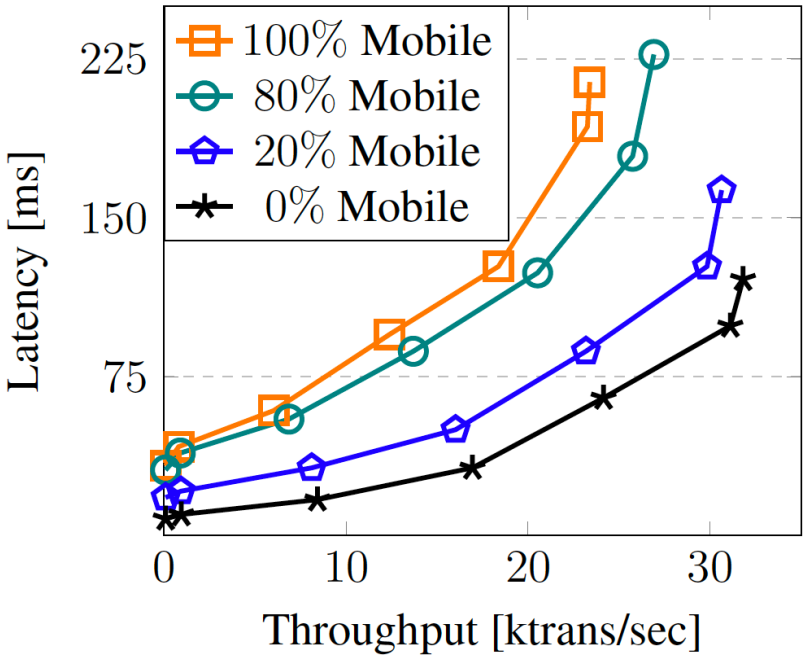

Byzantine

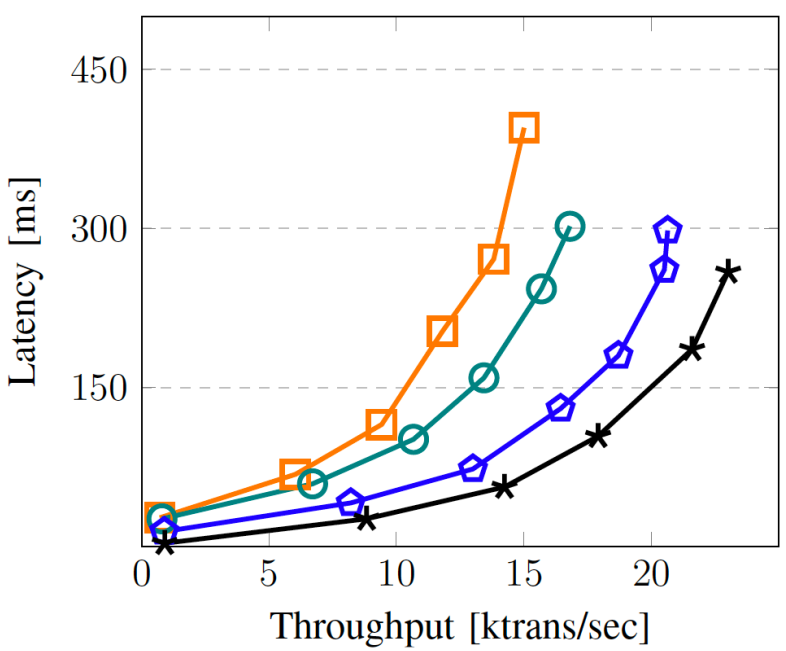

- 20% mobile transactions: 4% reduction in throughput
- increasing mobile devices from 0% to 100% (crash-only): 25% reduction in throughput
- increasing mobile devices from 0% to 100% (Byzantine): 36% reduction in throughput

Wide-area networks

Domains: California, Oregon, Virginia, Ohio, Tokyo, Seoul, and Hong Kong

Crash-only

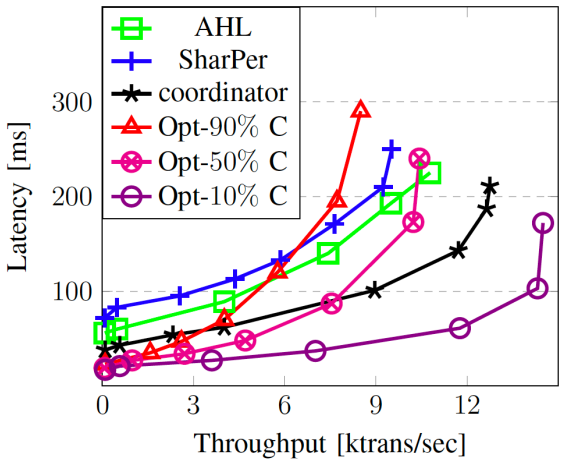

Byzantine

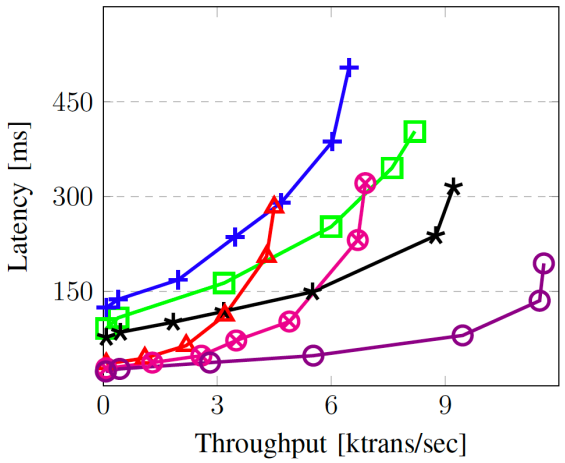

10% cross-domain transactions

Crash-only

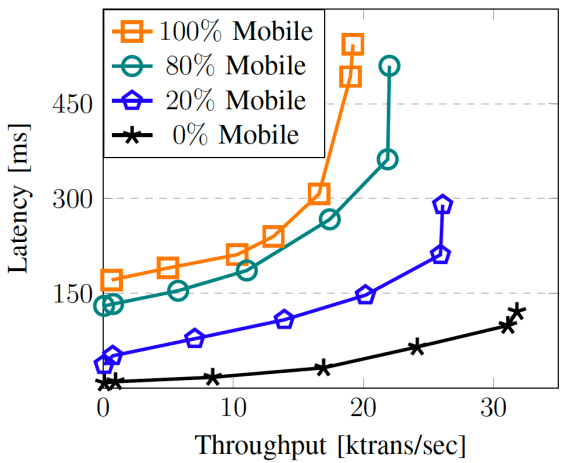

Byzantine

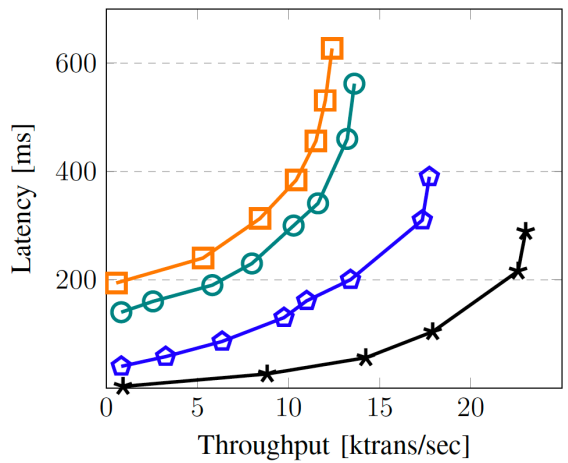

Mobile devices

- Conflicting transactions significantly reduce the performance of the optimistic protocol in high contention workloads
- Larger gap between the performance of the coordinator-based approach and AHL
- AHL demonstrates better performance compared to SharPer
- Increasing mobile devices from 0% to 100% (crash-only): 38% reduction in throughput

Evaluation Summary

- The coordinator-based protocol outperforms SharPer and AHL
  - Scalable solution that can be practically deployed over wide-area networks
- The optimistic protocol processes transactions efficiently in low-contention workloads
- The protocol performance is significantly reduced in high-contention workloads
  - due to inconsistency between the ledgers of different domains
- While SharPer outperforms AHL in nearby domains, AHL demonstrates better performance in far apart domains.
- Saguaro supports mobility over wide-area networks efficiently

Thank You!

Questions?

mjamiri@seas.upenn.edu