


When Blockchain Meets Deterministic Concurrency Control

2021 ACM SIGMOD SRC | June 20-25 | Xi'an, Shaanxi, China
Ziliang LAI (Bruce), The Chinese University of Hong Kong



Background

- Blockchains require determinism to ensure convergence.
- Concurrency may induce non-determinism because of non-deterministic OS scheduling.

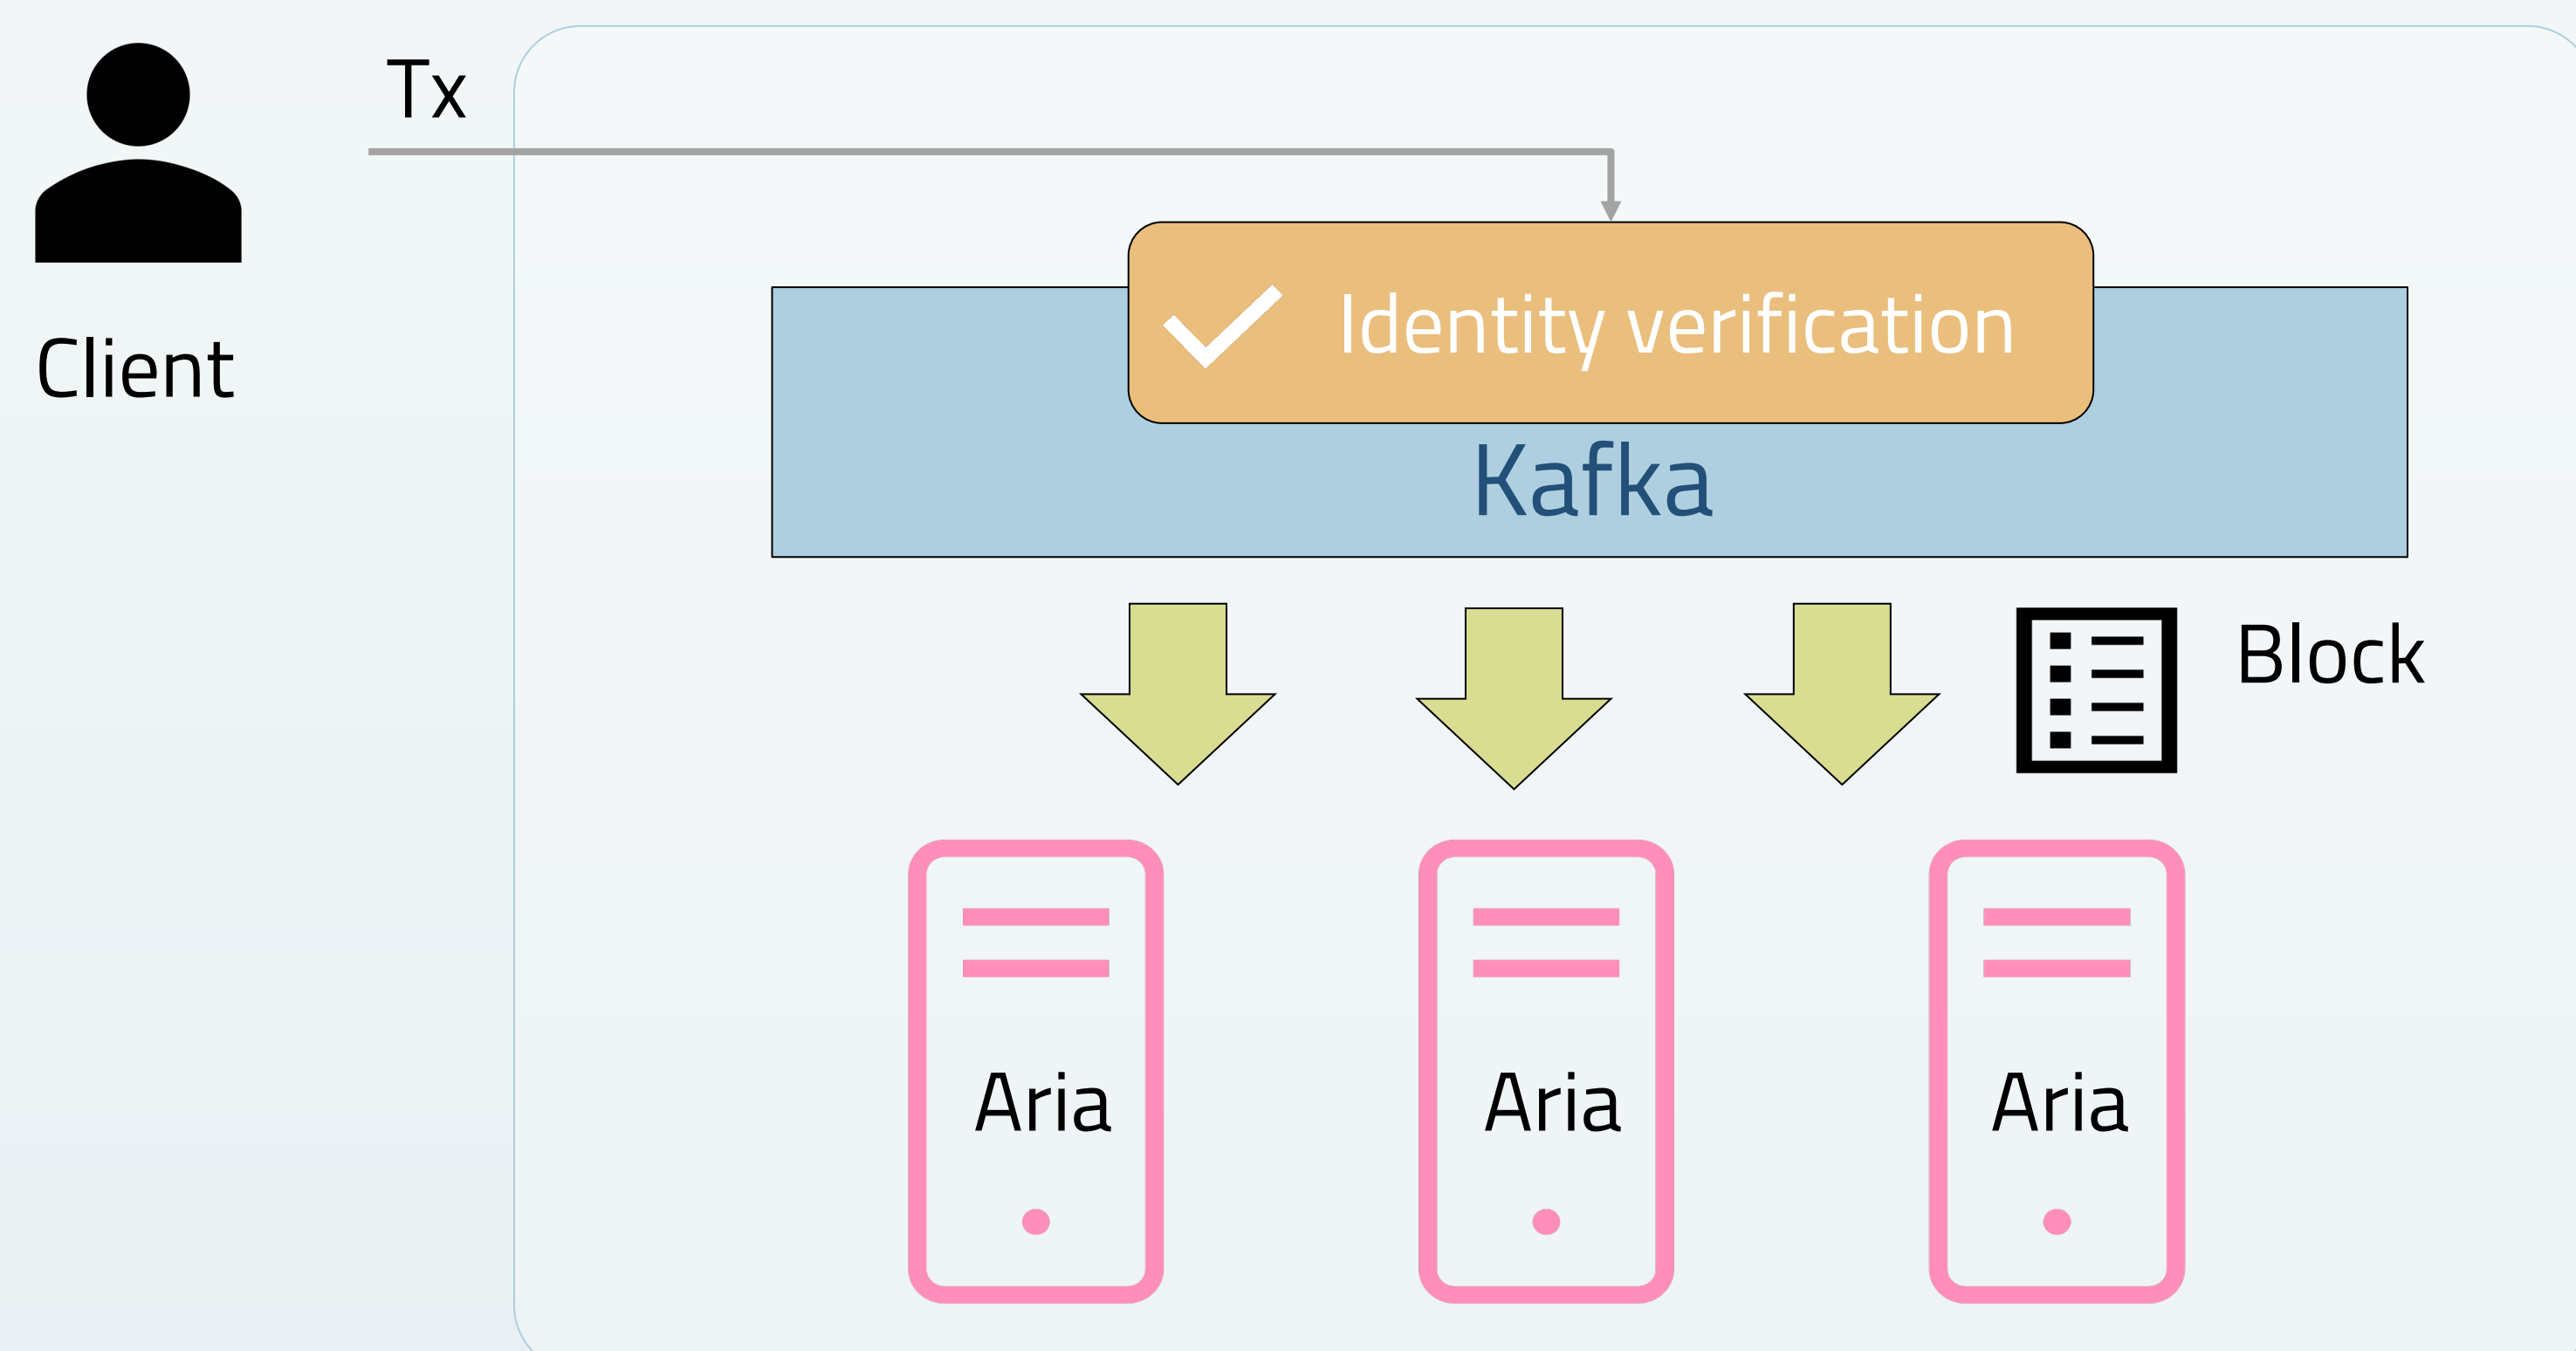
Traditional methods		
System	Method	Problem
Quorum [1]	Serial execution	No concurrency
Fabric Fabric++ Fabric# [2]	<ul style="list-style-type: none"> Parallel execution (non-deterministic) Consensus on read-write-sets (deterministic) 	Network overhead of running consensus on read-write-sets
BCR [3]	Deterministic concurrency control in disguise	Not fully concurrent

- Deterministic database : concurrency + determinism

Build a blockchain on top of a deterministic database?

Chainifying a Deterministic Database

- AriaBC: Aria [4] + Kafka (consensus) + crypto
- Reimplemented Aria on Postgresql for fair comparisons.

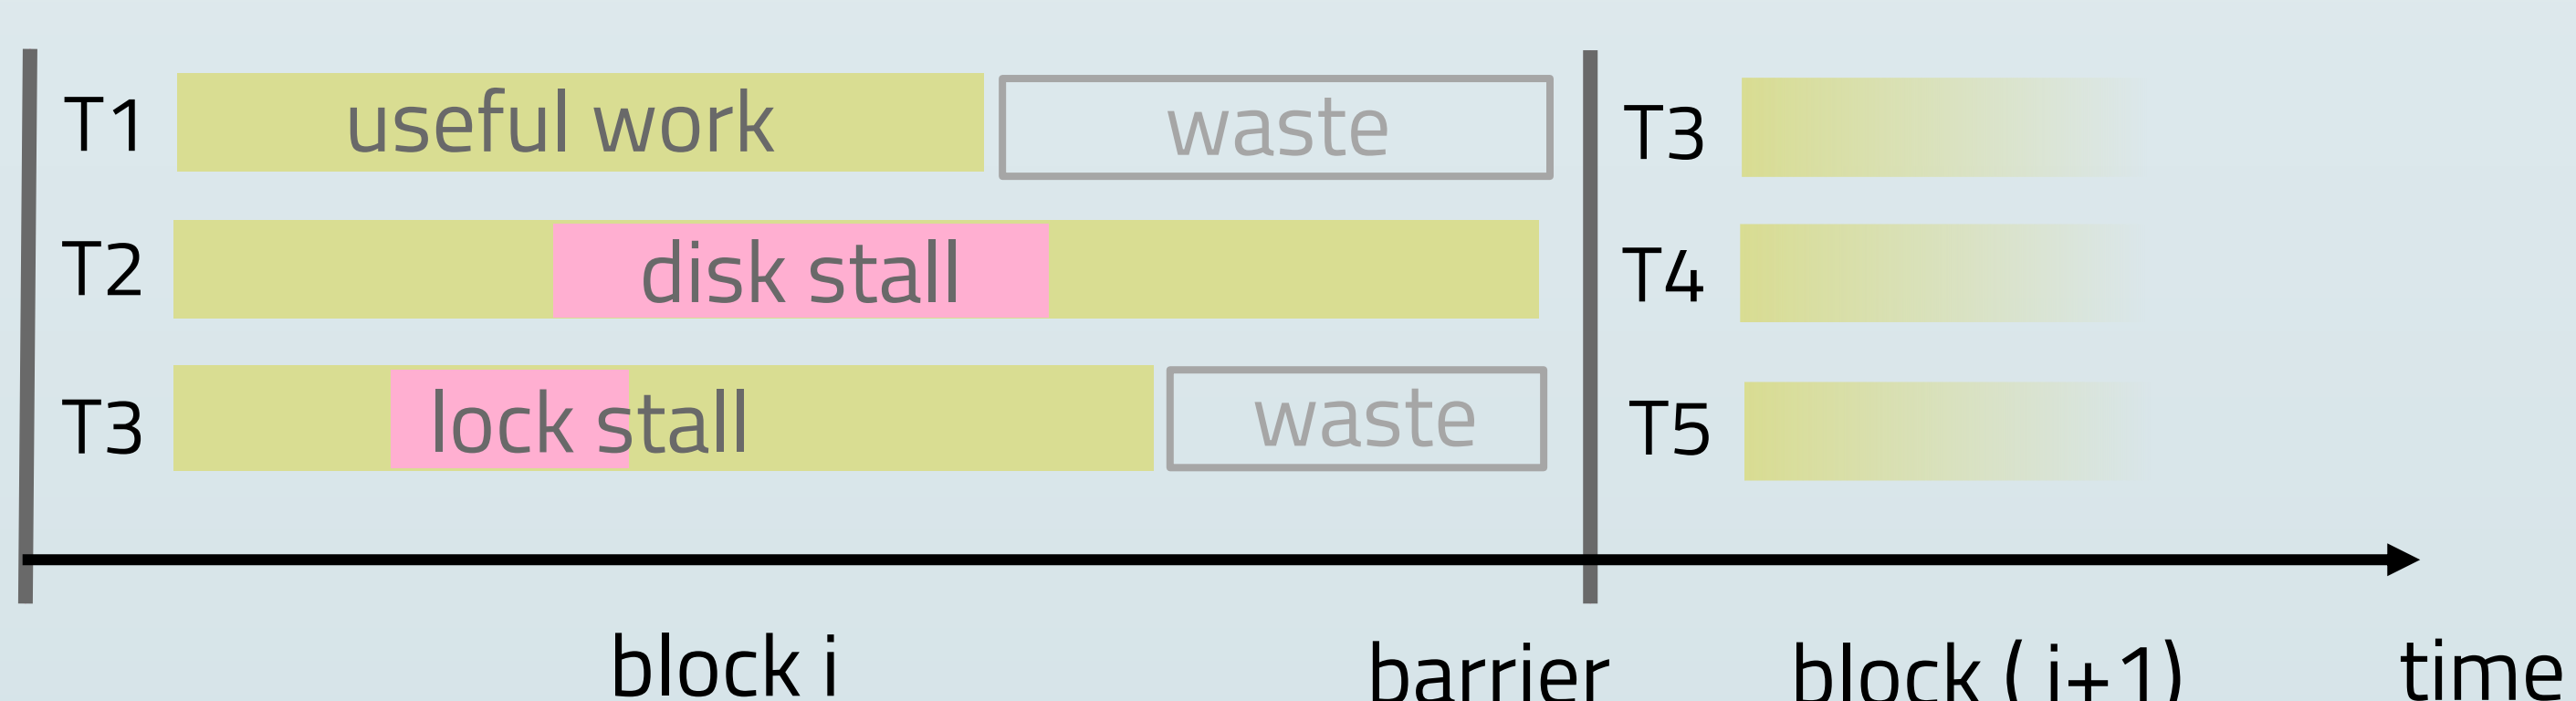


Problem 1: only one writer commits

- Aria disallows ww-conflicts.
- Example: T1, T2, T3 in the same block updates the same record, T2 and T3 must abort.
- Extensive aborts due to update hot spots


Problem 2: straggler

- Aria starts processing a block only until the previous one is finished.




DCCBC: Our Improved Deterministic CC for Blockchain

Write reordering



- No transaction is aborted due to ww-conflict.
- Maintain high throughput even with hot spots.

Breaking the block barrier

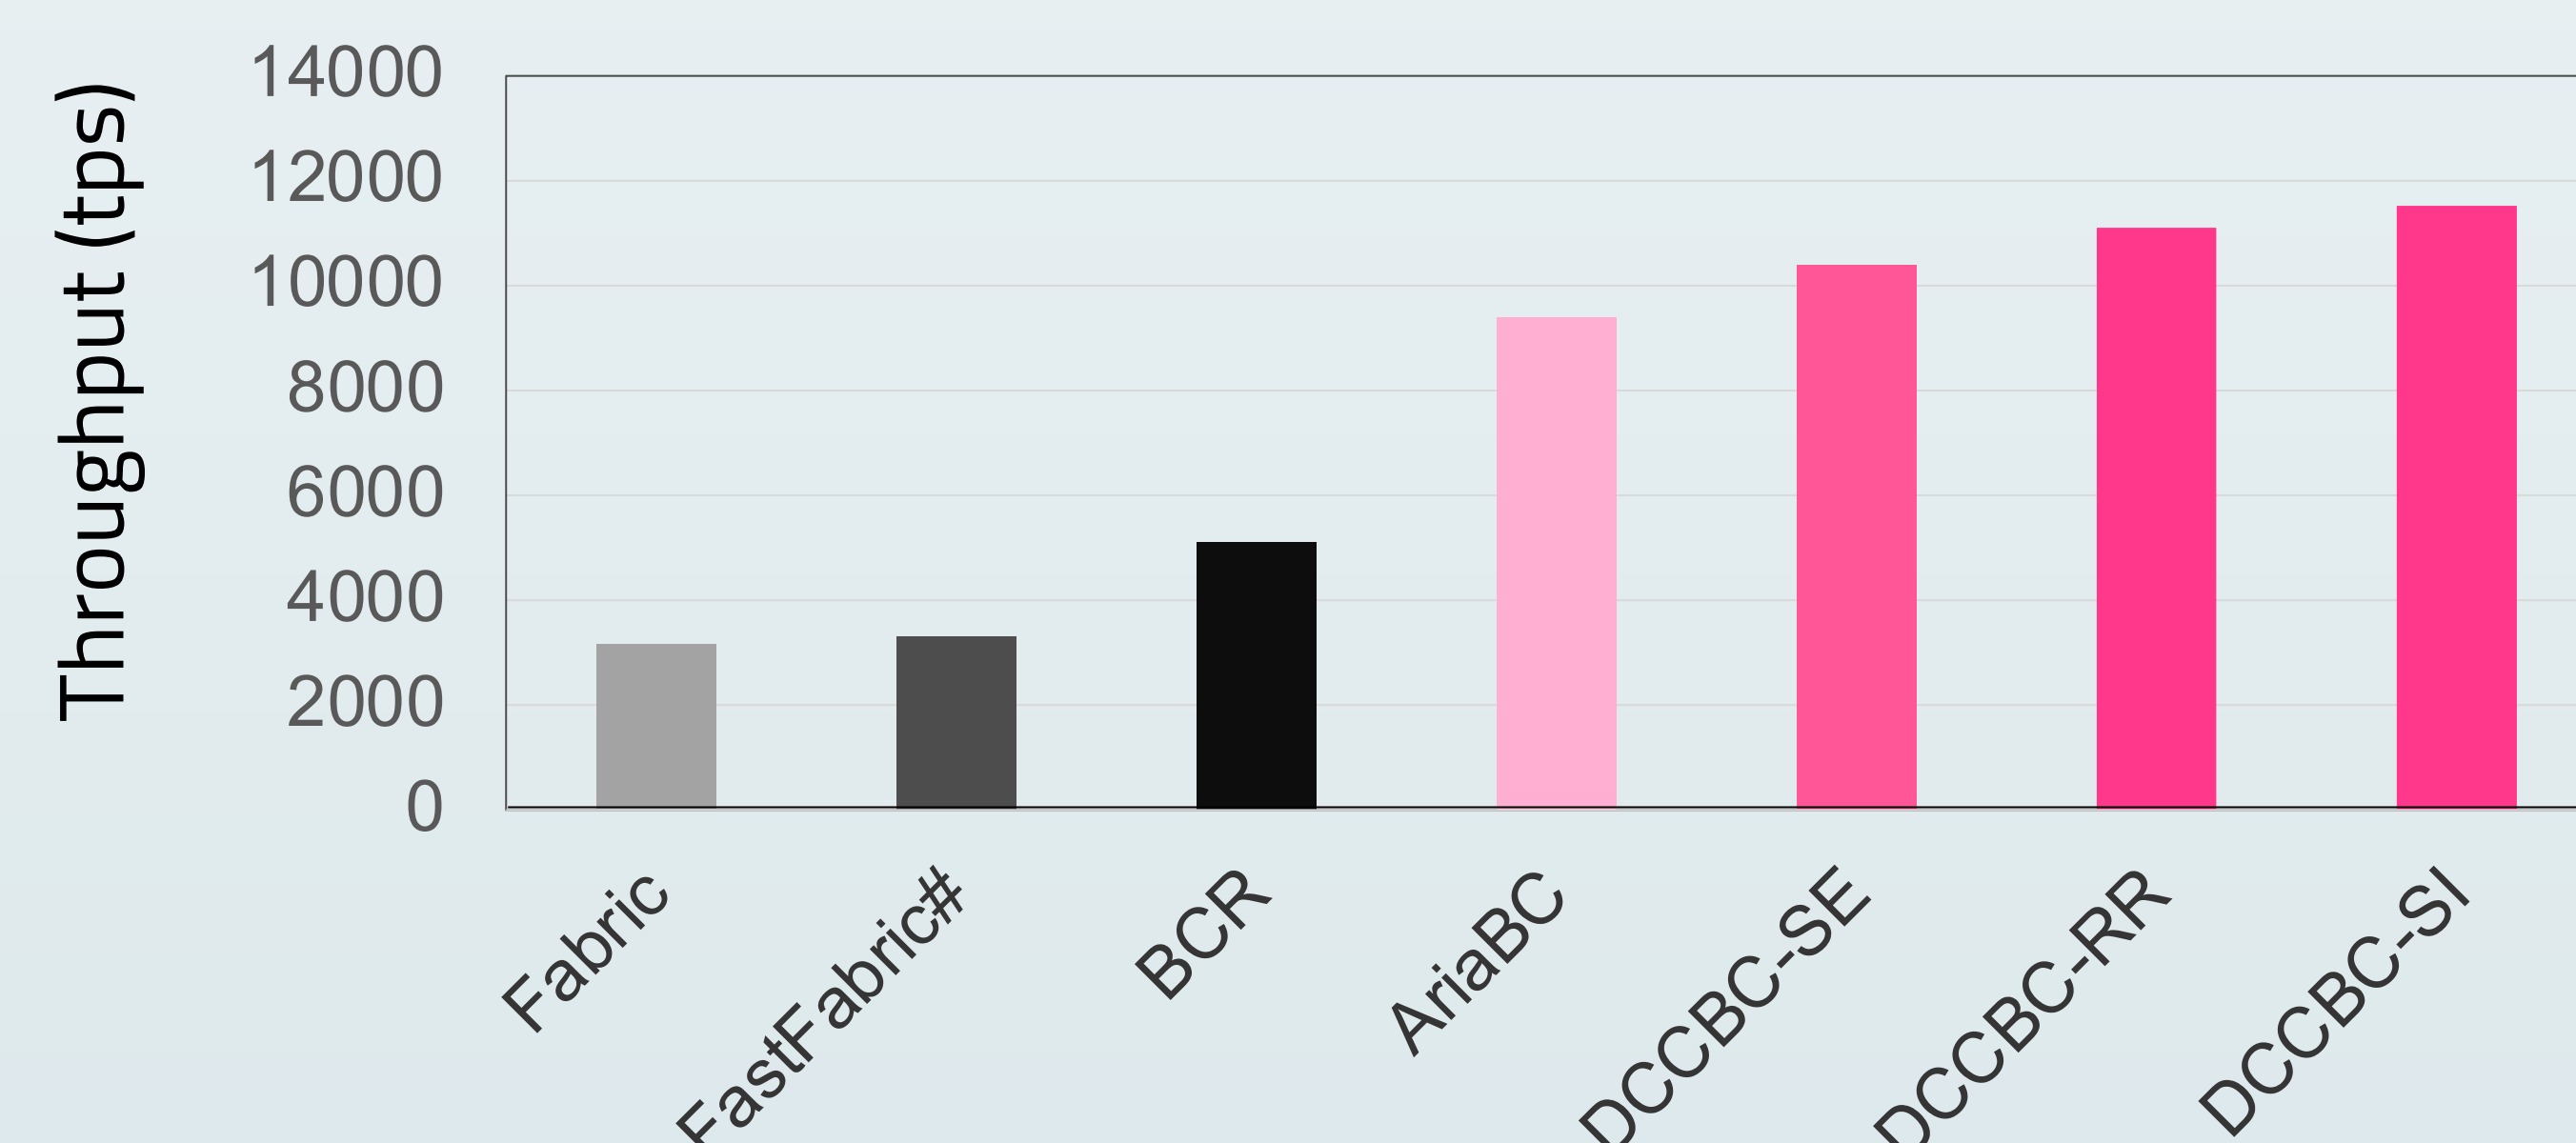


- Optimistically schedule a transaction in block (i+1) once a transaction in block i completes.
- When all transactions in block i completes, notify the optimistically executing transactions to check for stale read.
- Retry if stale read is found.

Practical Isolation Levels

- Like in relational database, many blockchain applications run safely under weaker isolation levels.
 - Smallbank benchmark: Repeatable Read
 - Money transfer (A send money to B): Snapshot Isolation
- DCCBC supports Serializable, Repeatable Read, and Snapshot Isolation.

Preliminary Results: Smallbank Benchmark



[1] Quorum. <http://www.jpmorgan.com/global/Quorum>
 [2] Ruan, Pingcheng, et al. "A transactional perspective on execute-order-validate blockchains." SIGMOD 2020
 [3] Nathan, Senthil, et al. "Blockchain Meets Database: Design and Implementation of a Blockchain Relational Database." VLDB 2019
 [4] Lu, Yi, et al. "Aria: a fast and practical deterministic OLTP database." VLDB 2020