Performance-Optimal Read-only Transactions

Haonan Lu★

Siddhartha Sen✢, Wyatt Lloyd★

★Princeton University, ✢Microsoft Research

Distributed Storage Systems

Enable Today's Web Services

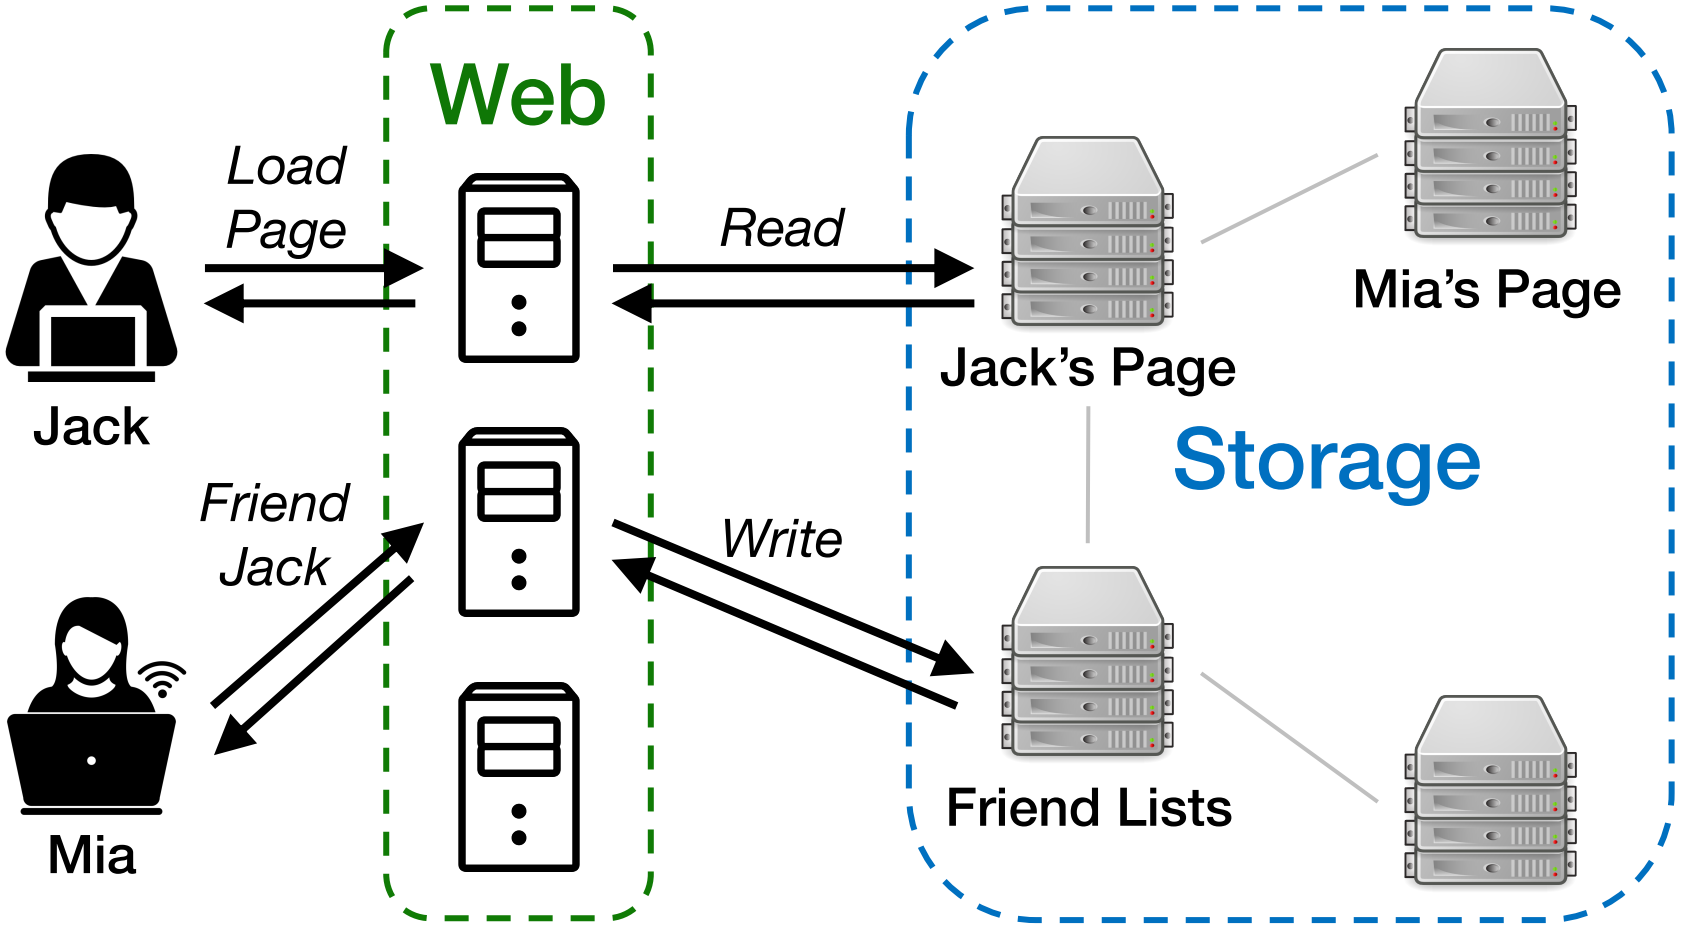

Distributed Storage Systems

Reads Dominate Workloads

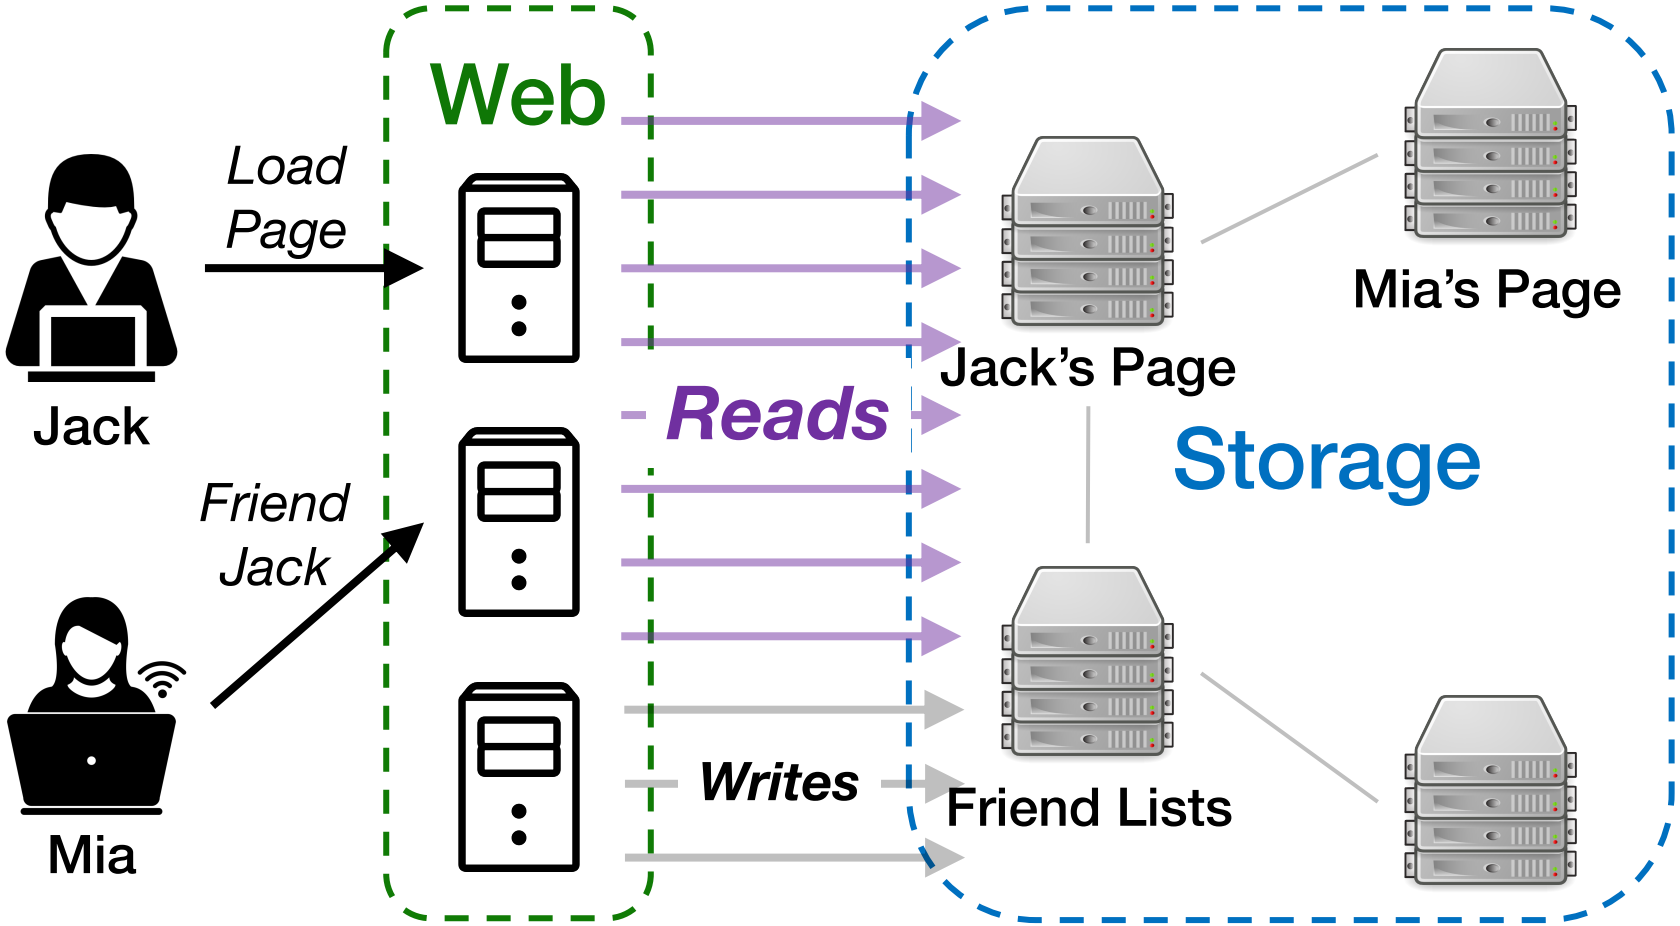

Distributed Storage Systems

Simple Reads Are Insufficient

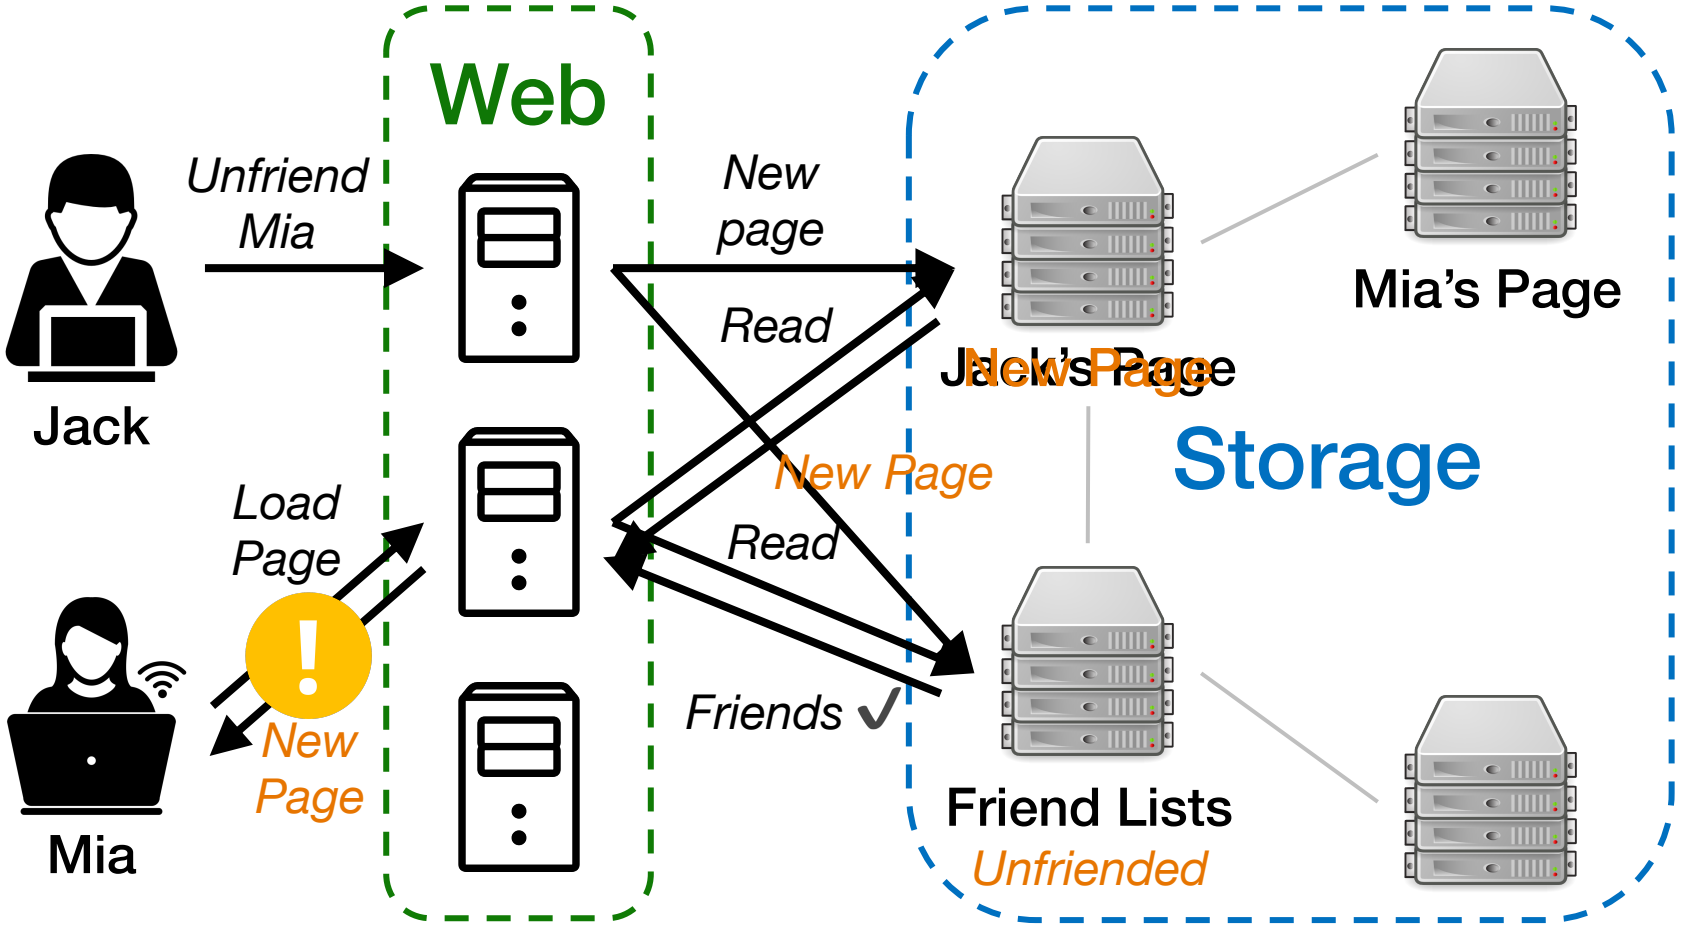

Read-Only Transactions

- A group of simple reads sent in parallel
- Do not write data
  - Writes are allowed in the system
- Coordinate a consistent view across shards

Coordination overhead causes higher latency and lower throughput

Goal:

Read-only transaction performance as close as possible to simple reads

Goal:

Read-only transaction performance as close as possible to simple reads

We answer:

- What does optimal performance mean for read-only transactions?
- When is optimal performance achievable?
- How can we design performance-optimal read-only transactions?

Performance Factors
Engineering vs. Algorithmic

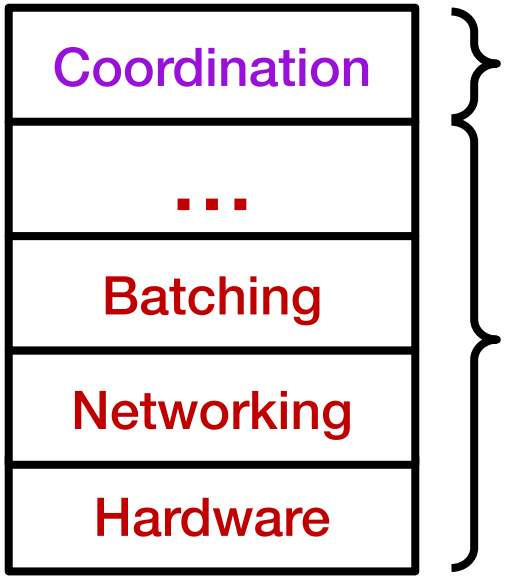

- Focus on the algorithmic properties due to coordination
- Equally impact simple reads and read-only transactions
- Abstract engineering factors by comparing to simple reads

Performance Factors

Algorithmic Properties

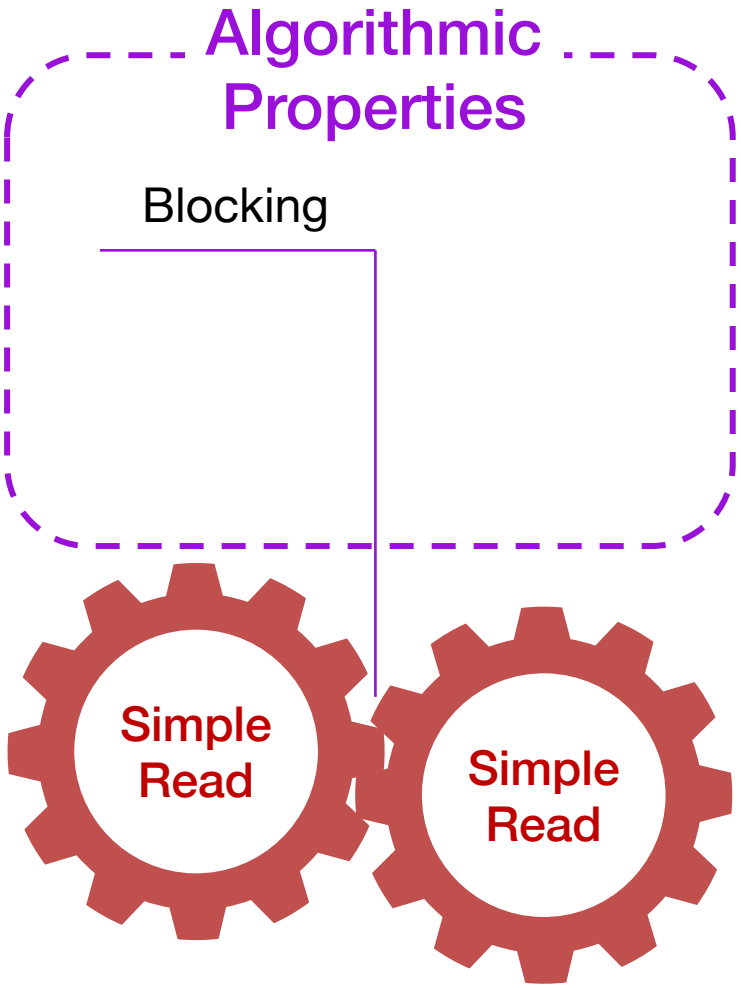

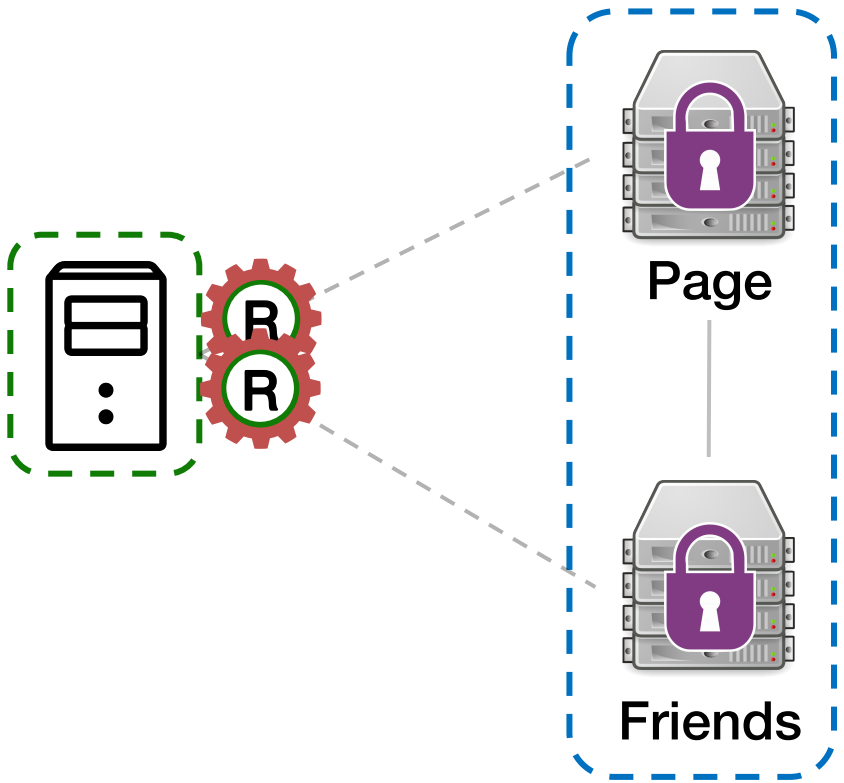

Performance Factors

Algorithmic Properties

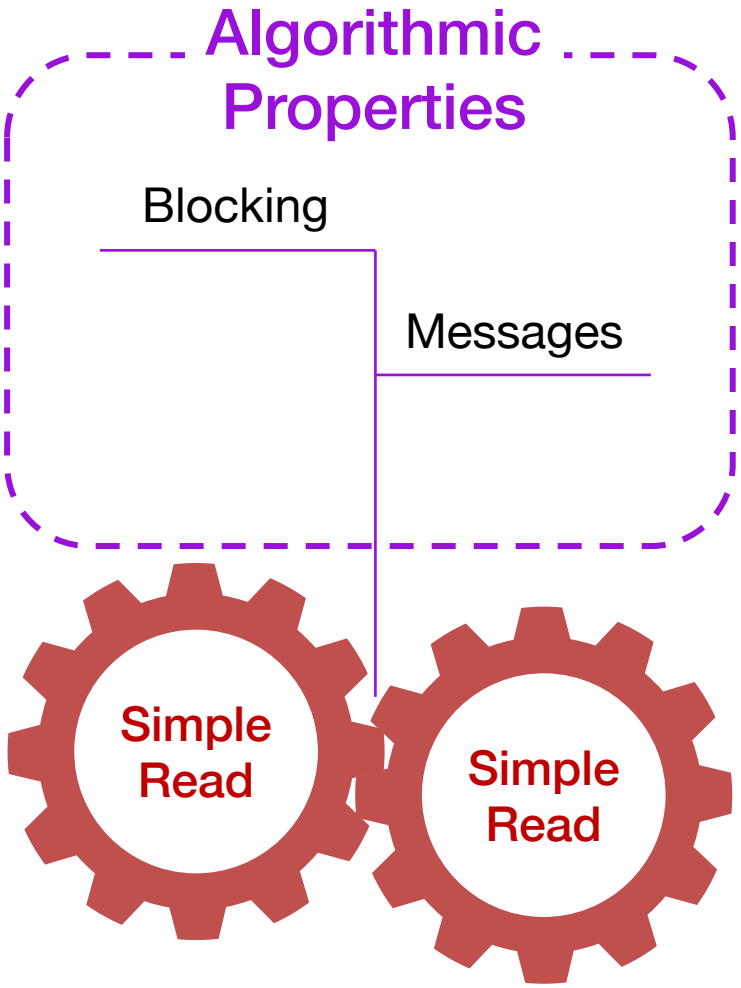

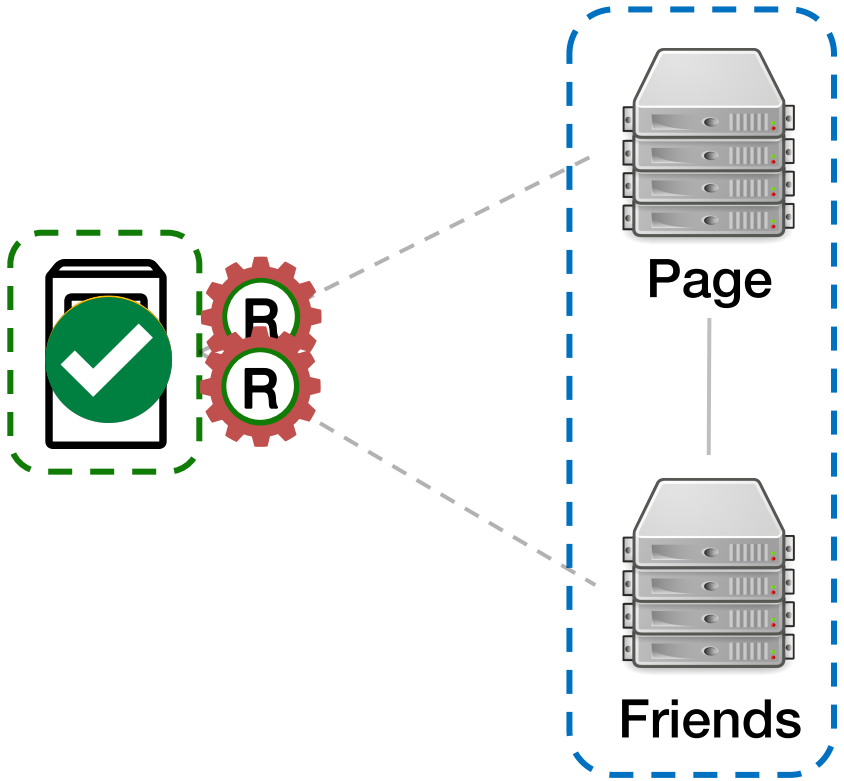

Performance Factors

Algorithmic Properties

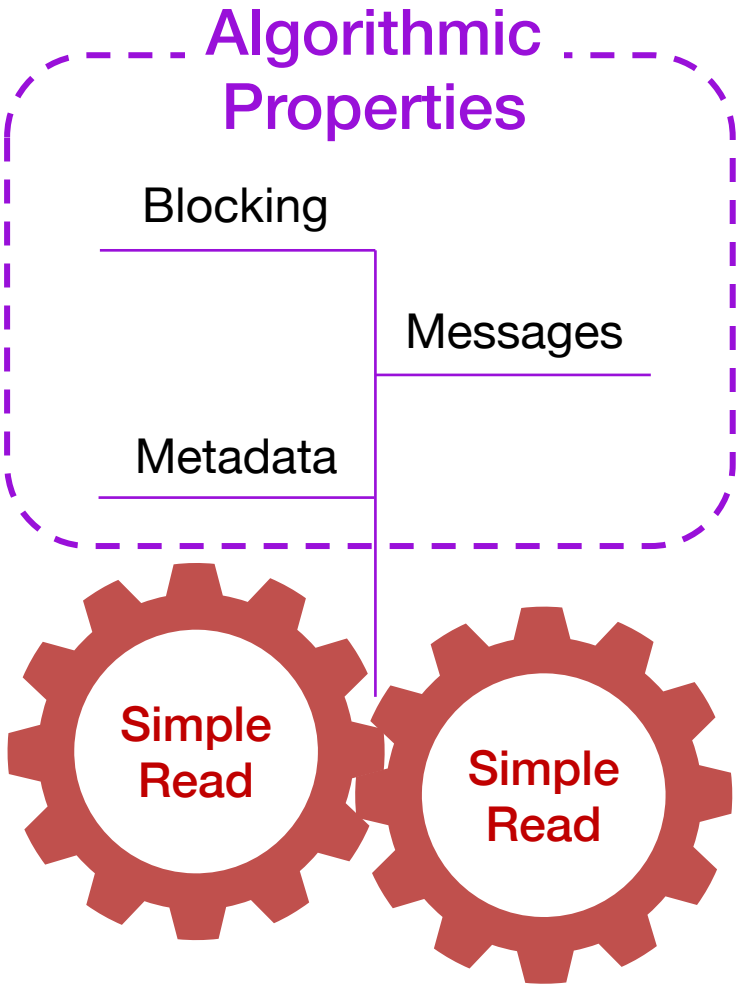

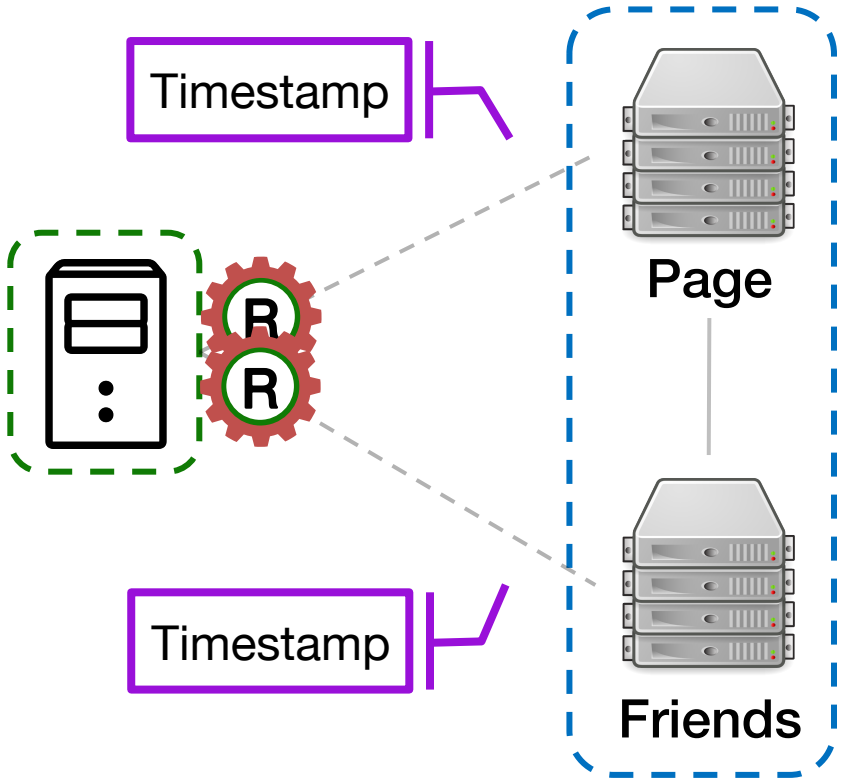

Performance Factors

Coordination Is Algorithmic

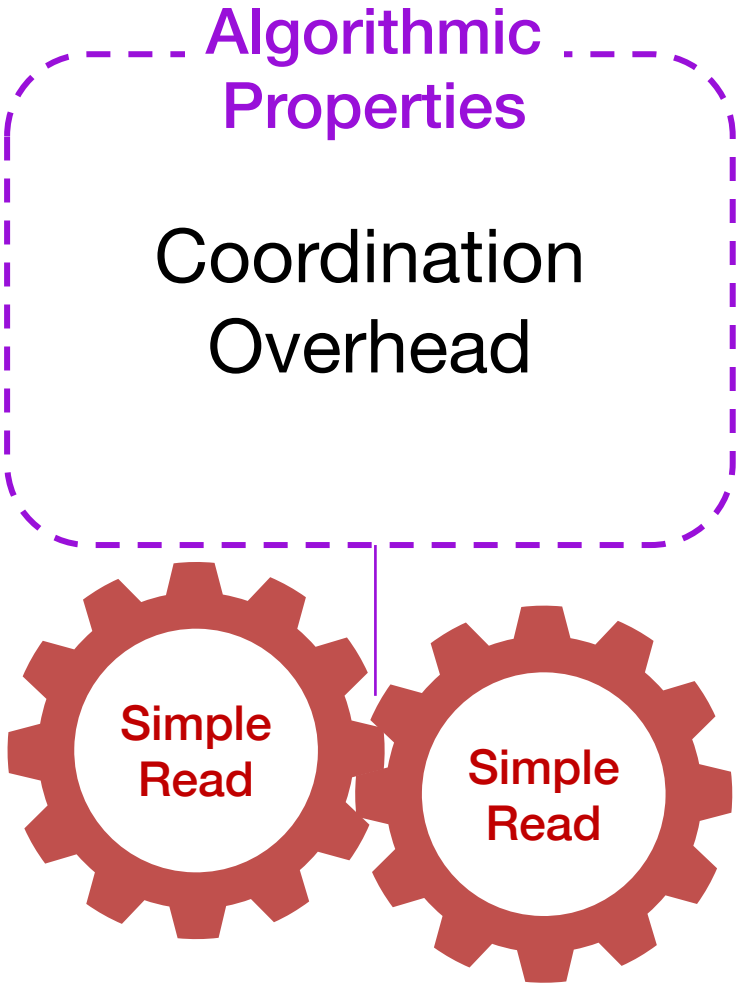

Read-Only Transactions
Optimal Performance

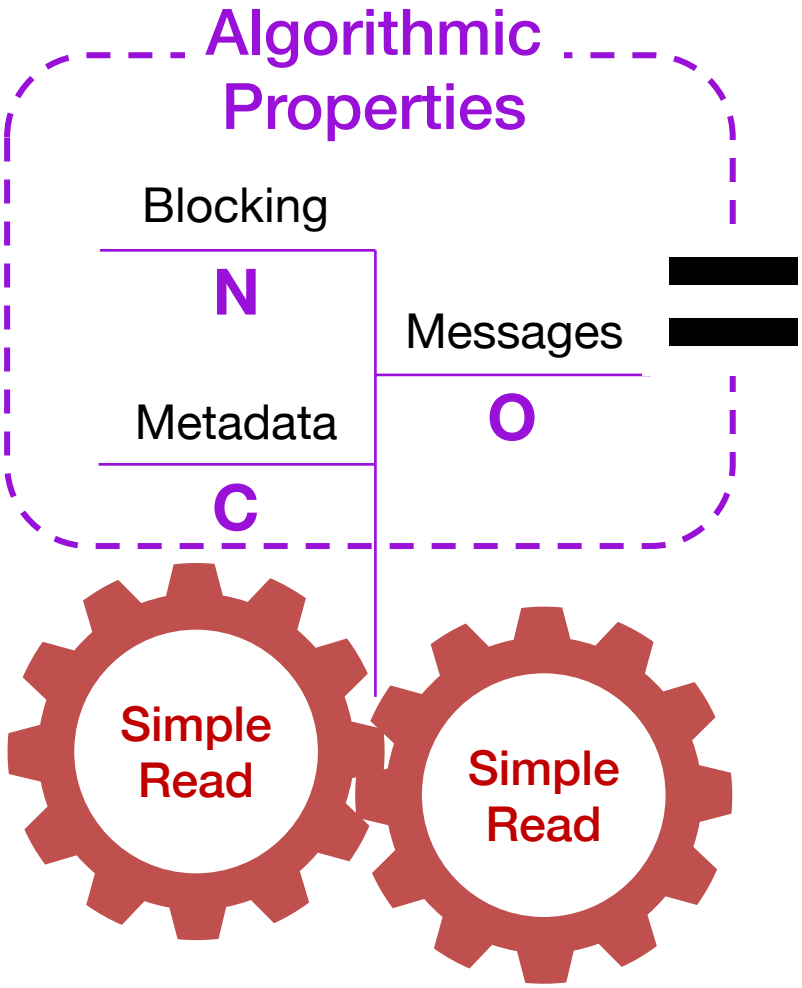

Performance-optimal Read-only Transactions (N,O,C)

Non-Blocking Reads

- Do not wait on external events
  - Distributed locks, timeouts, messages, etc.
- Lower latency
  - Avoid any time spent blocking
- Higher throughput
  - Avoid CPU cost of context switches

One-Round Communication

- One-round on-path reads
  - Succeed in one round, i.e., no retries
- No off-path messages
  - Required by reads but off the critical path
- Lower latency
  - Avoids time for extra on-path messages
- Higher throughput
  - Avoids CPU cost of processing extra messages

Constant Metadata

- Metadata
  - Information used to find a consistent view
  - Timestamps, transaction IDs, etc.
- Size of metadata remains constant regardless of contention
- Higher throughput
  - Avoids CPU cost of processing extra data

Performance-optimal read-only transactions are NOC:

Non-blocking messages

that complete in

One-round with

Constant metadata

Strict Serializability

- The strongest consistency model
  - Writing applications made easy
- Requires a total order + real-time order

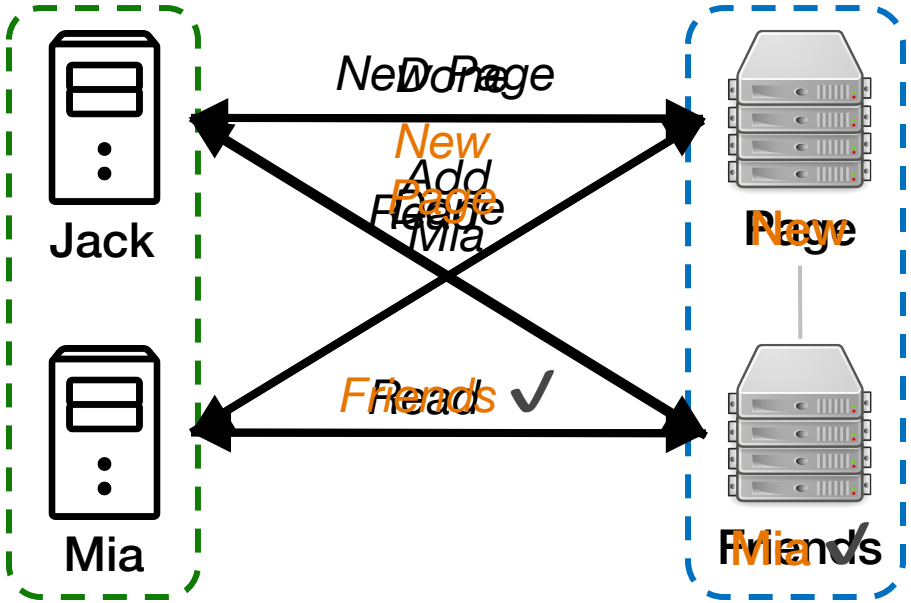

The NOCS Theorem:

Impossible for read-only transaction algorithms to achieve performance-optimality [N,O,C] and strict serializability [S]

Proof Intuition of NOCS

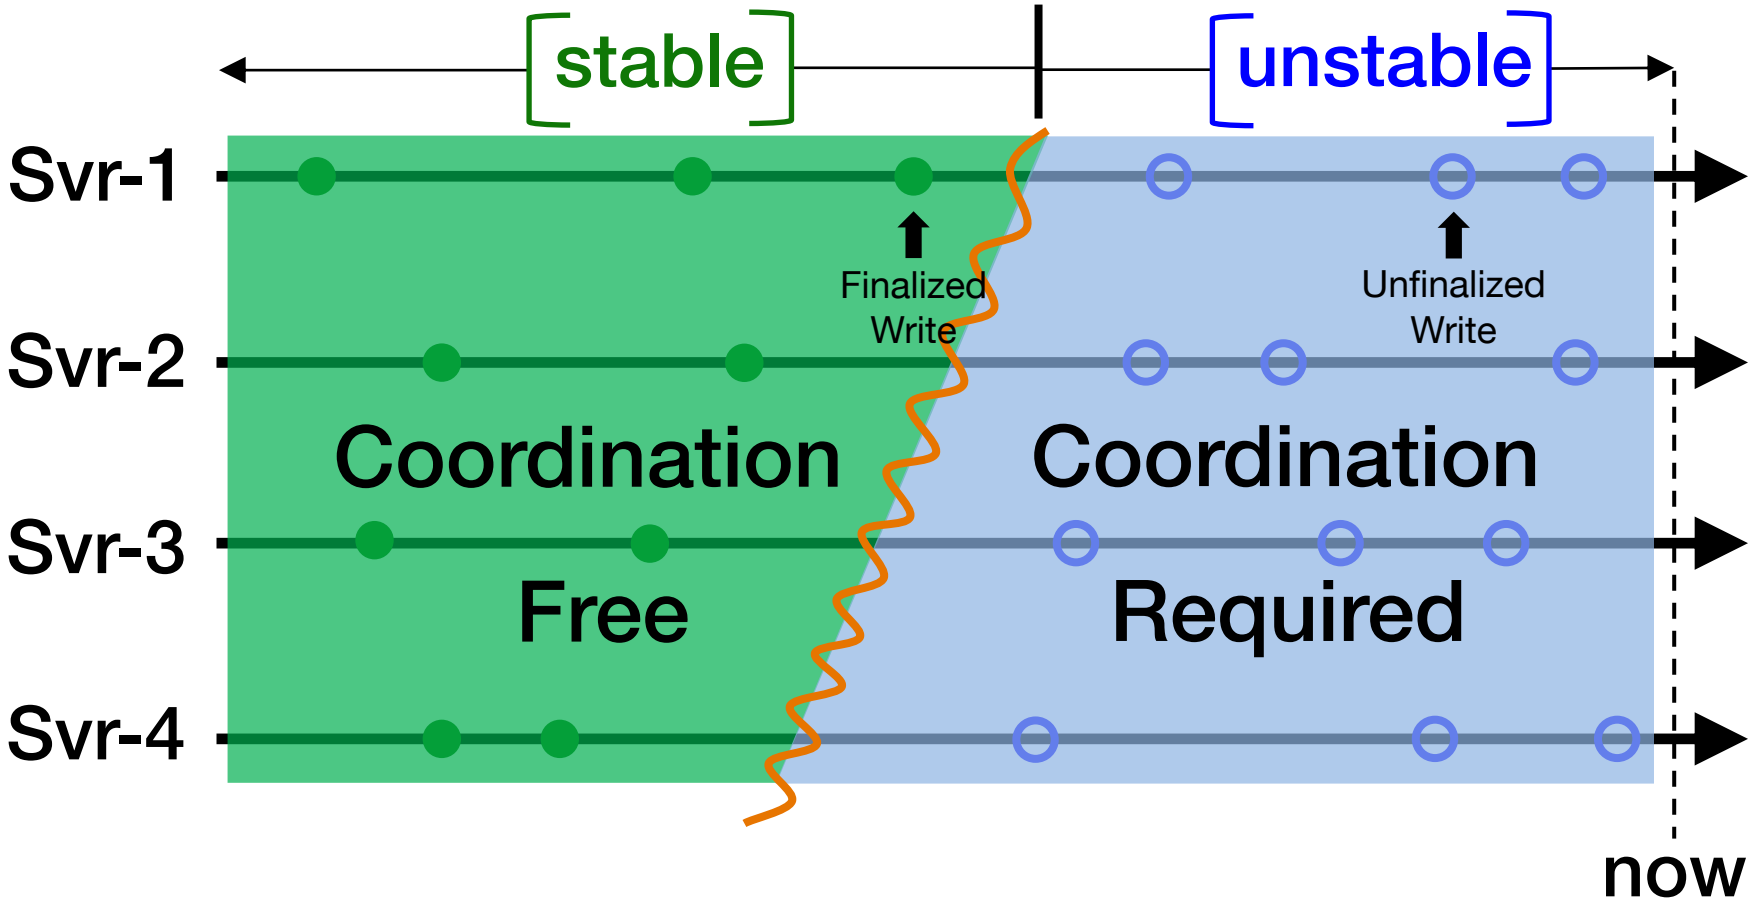

Proof Intuition of NOCS

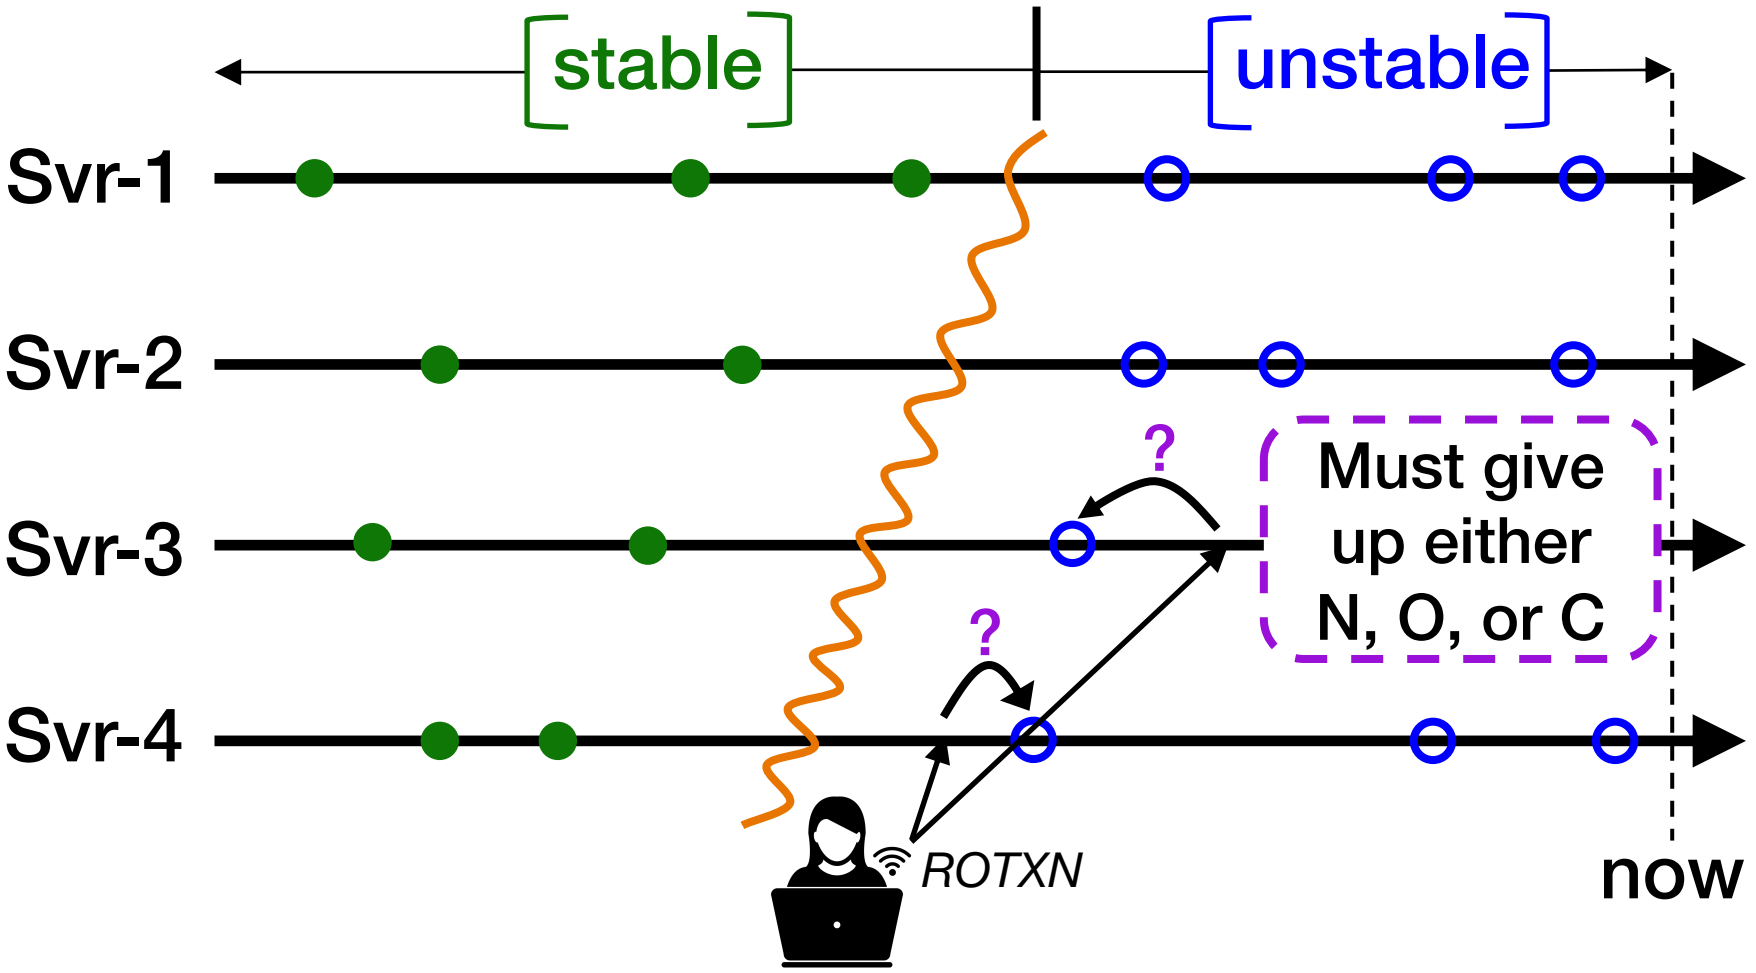

NOC Designs

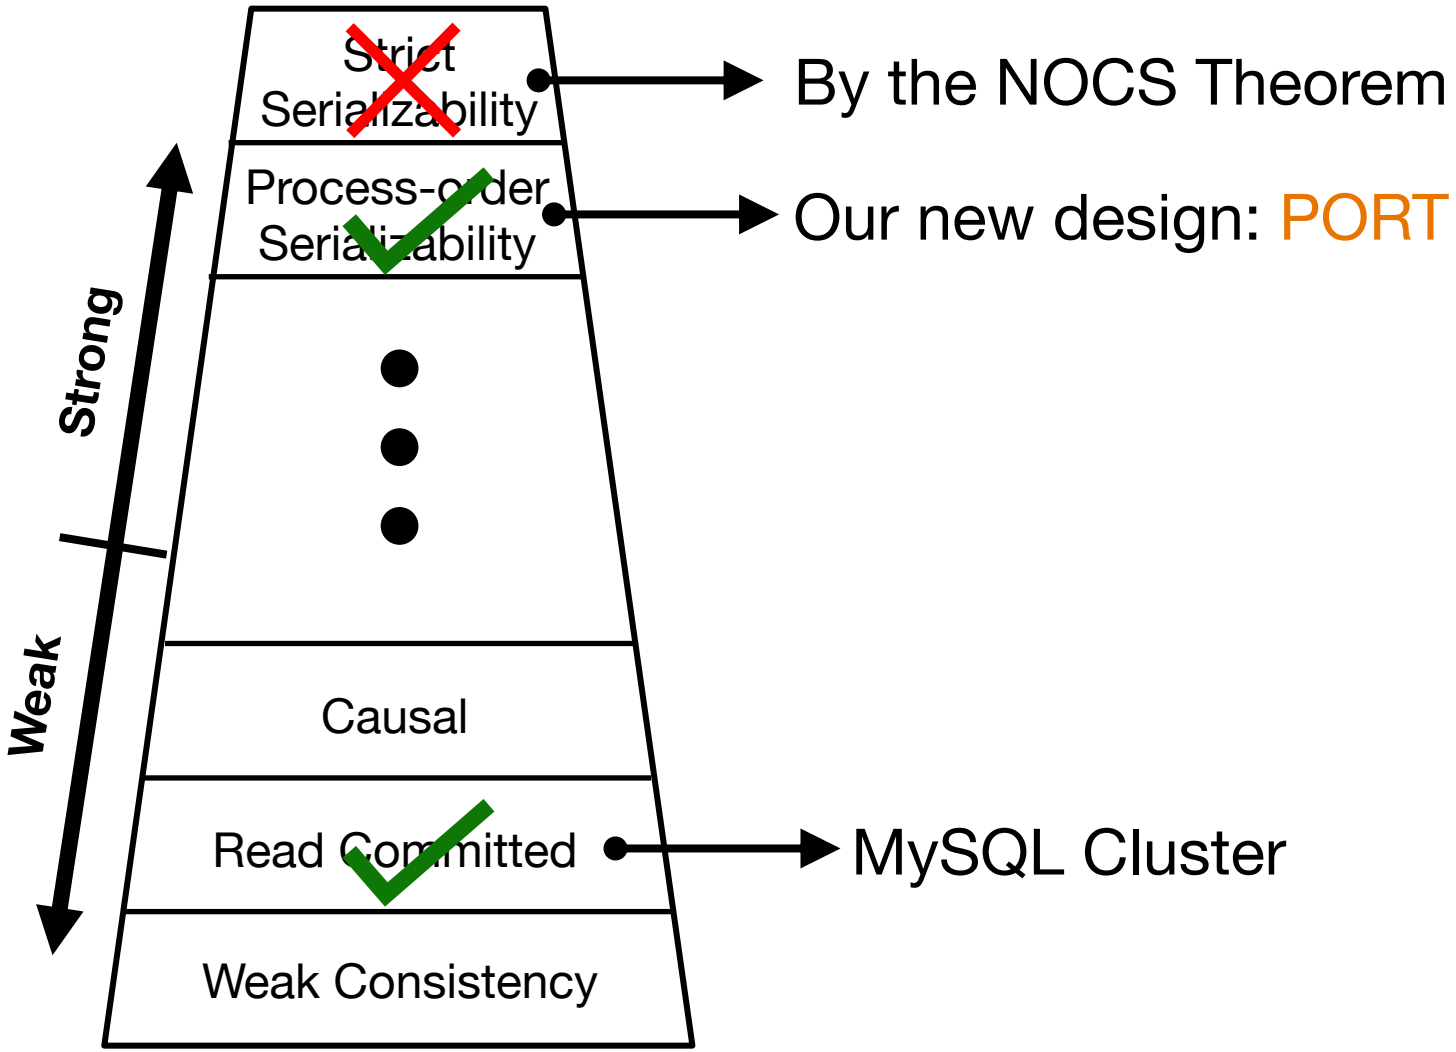
Strict Serializability
By the NOCS Theorem
Process-order Serializability
Our new design: PORT
Strong
Weak
Causal
Read Committed
MySQL Cluster
Weak Consistency

Design Insight

Capturing the Stable Frontier

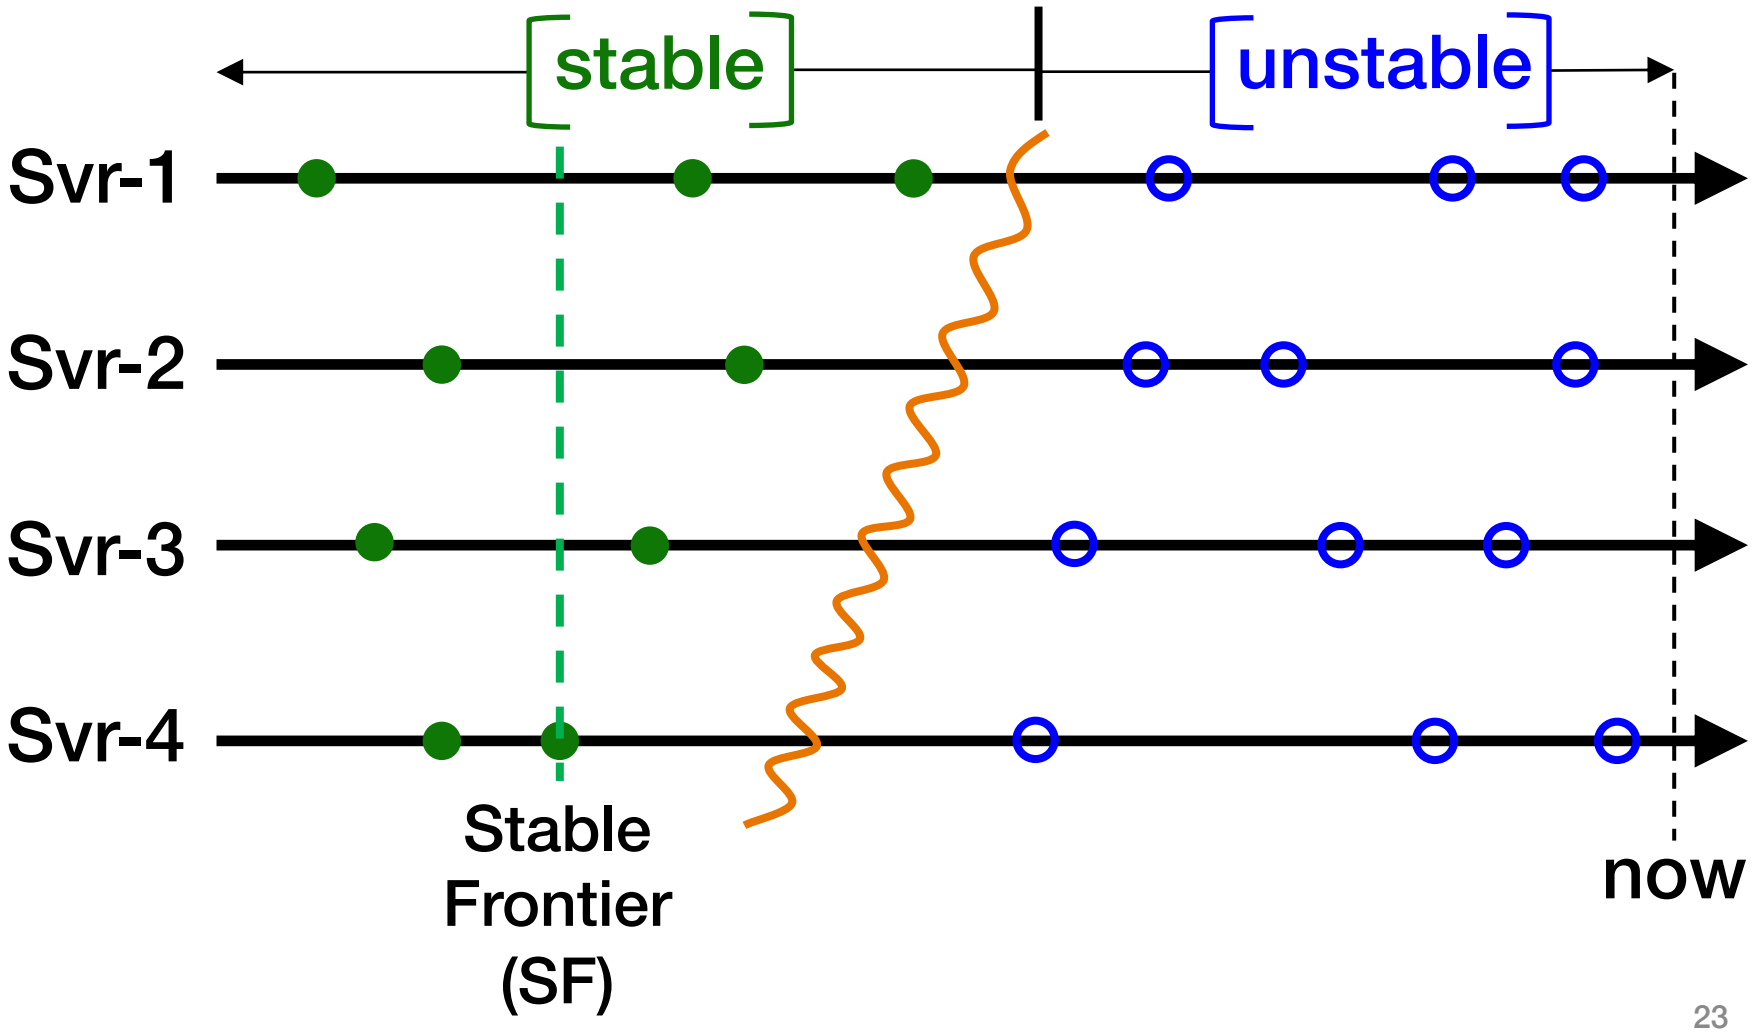

Version Clock

- A type of logical clock
  - Specialized for distributed storage systems
- Treat reads and writes differently
  - Enable optimizations for reads and writes
- Capture the stable frontier

PORT Overview

Jack

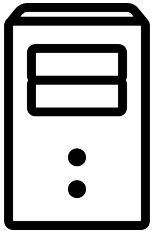

Web
Client

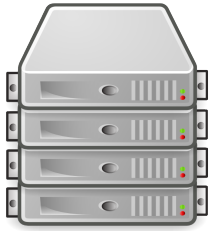

Storage
Server

PORT Overview

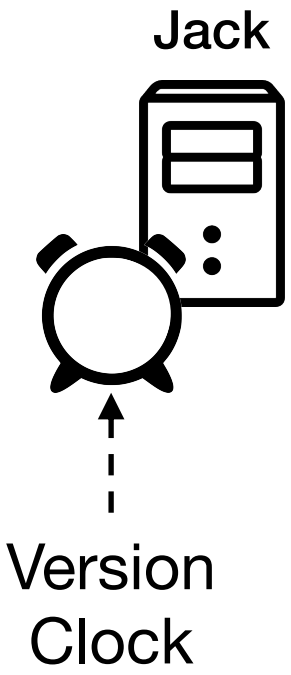

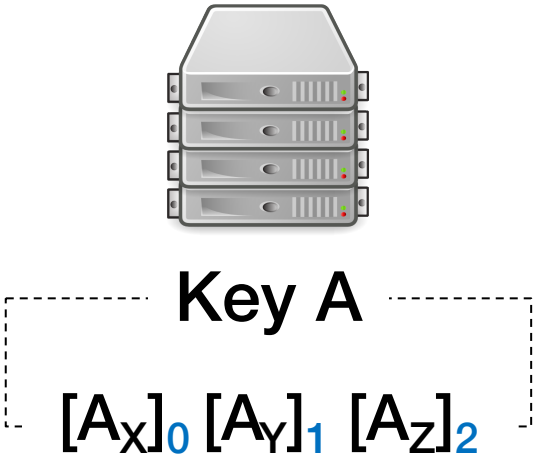

PORT Overview

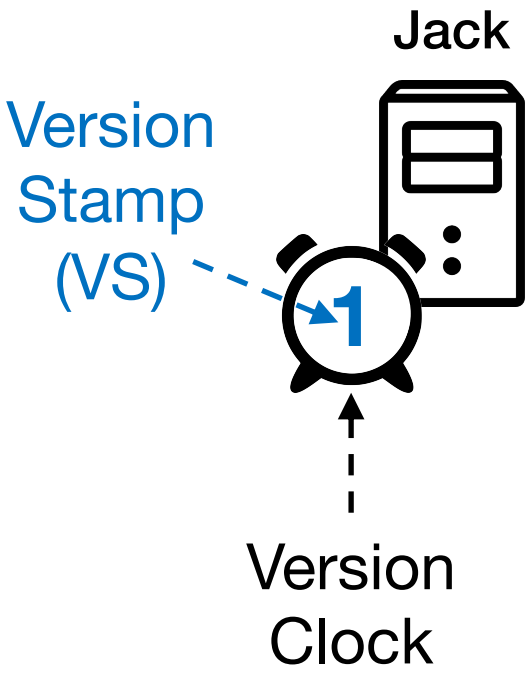

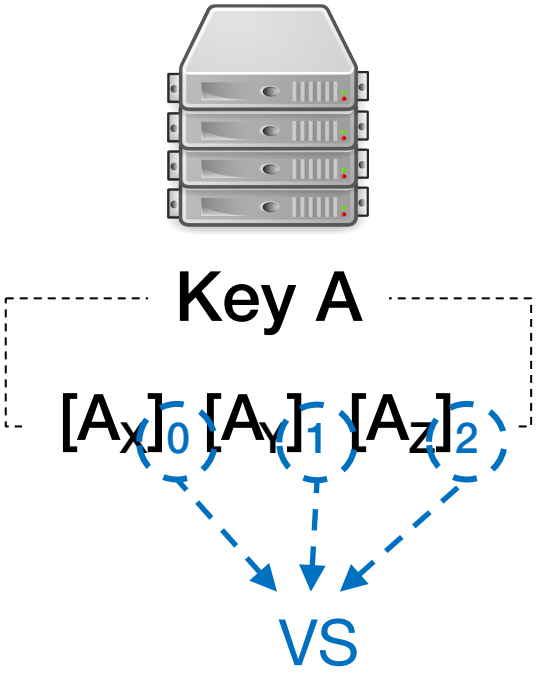

Write in PORT

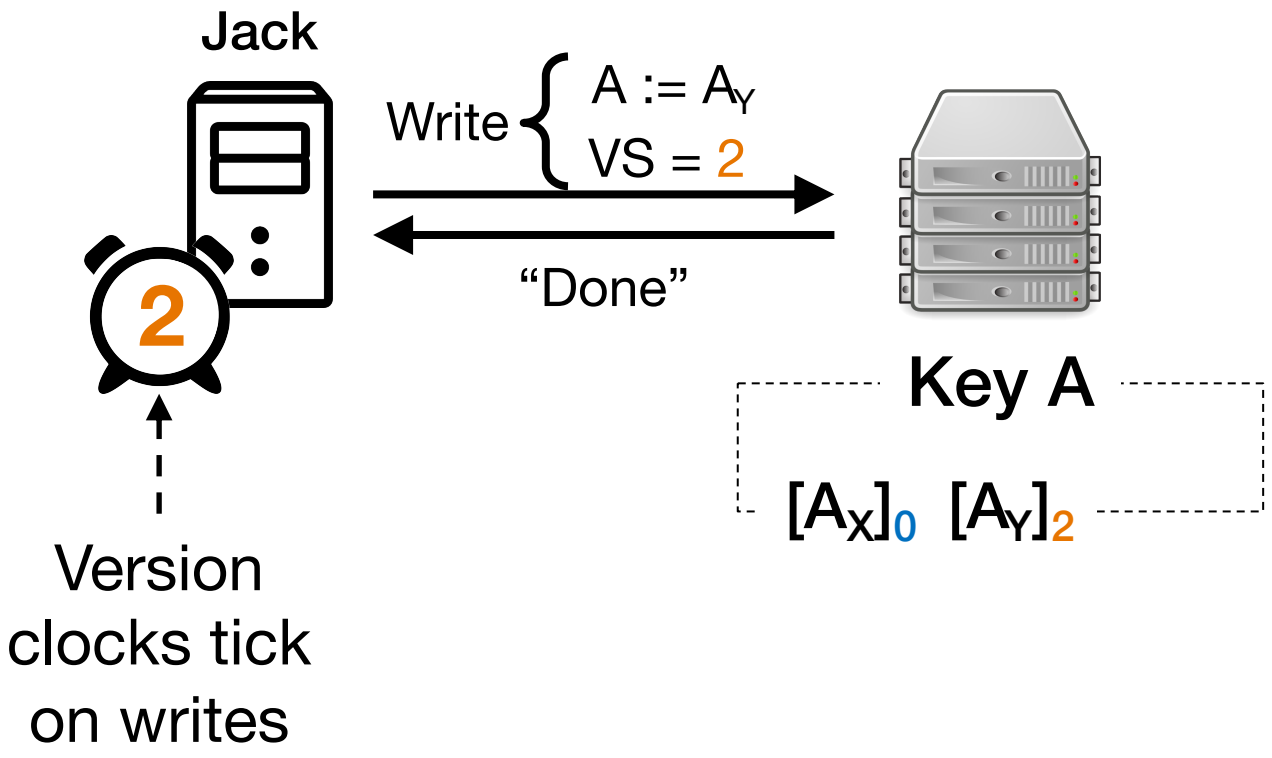

Read in Port

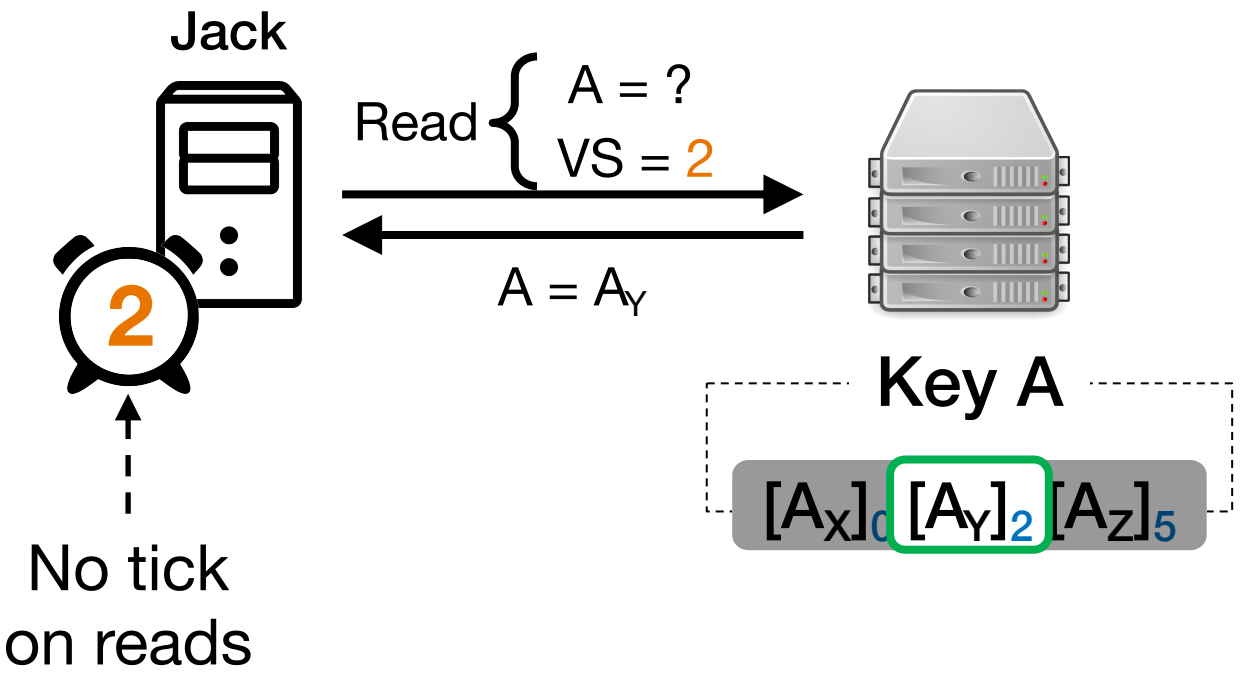

Read Promotion
Ensures a Total Order

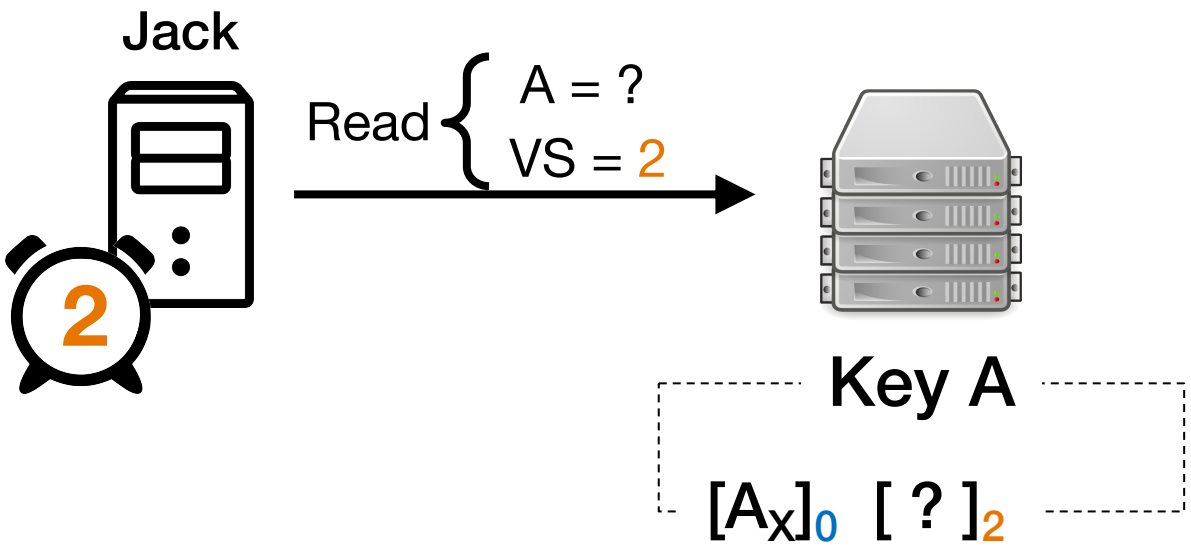

Read Promotion
Ensures a Total Order

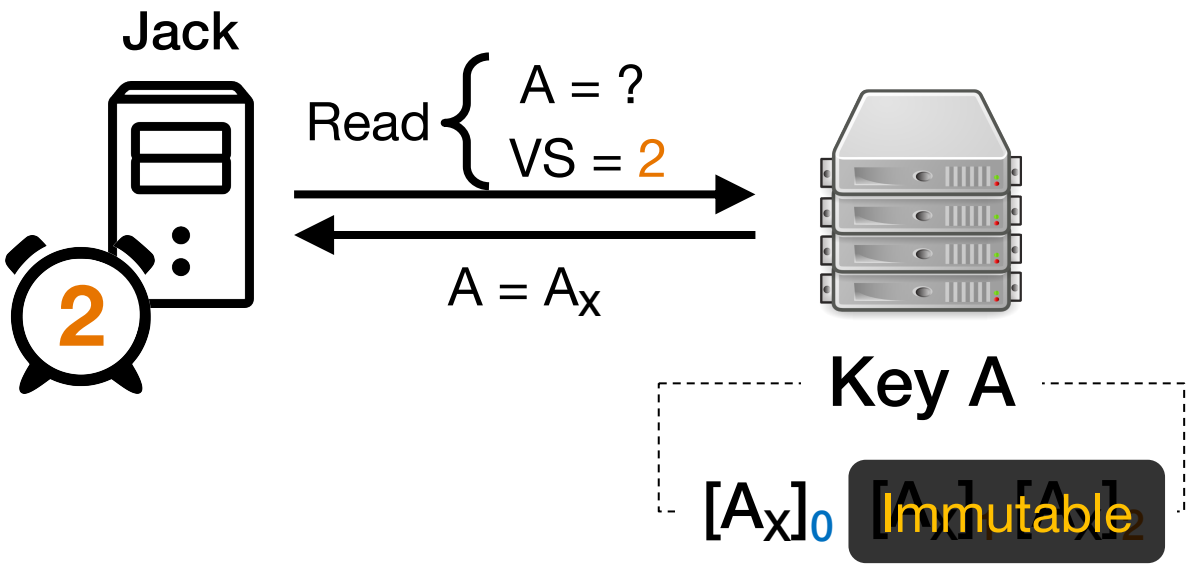

Read Promotion
Ensures a Total Order

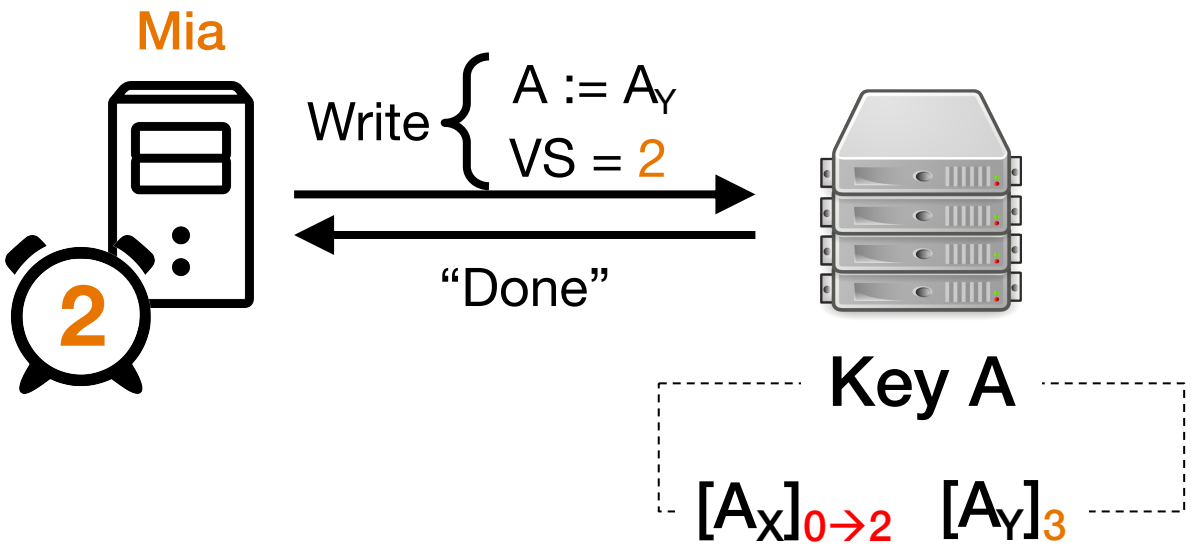

Track Stable Frontier

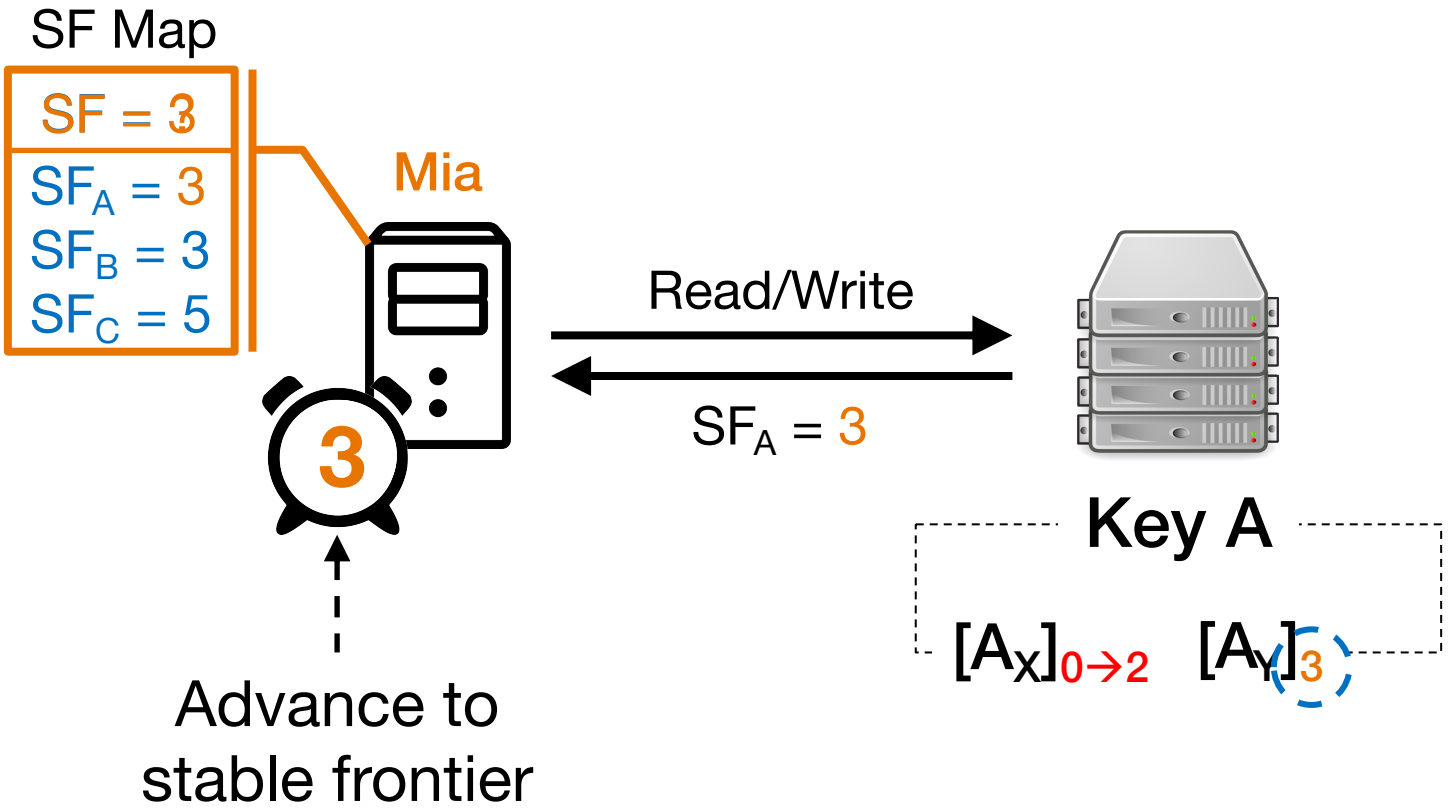

Read-Only Transaction Logic

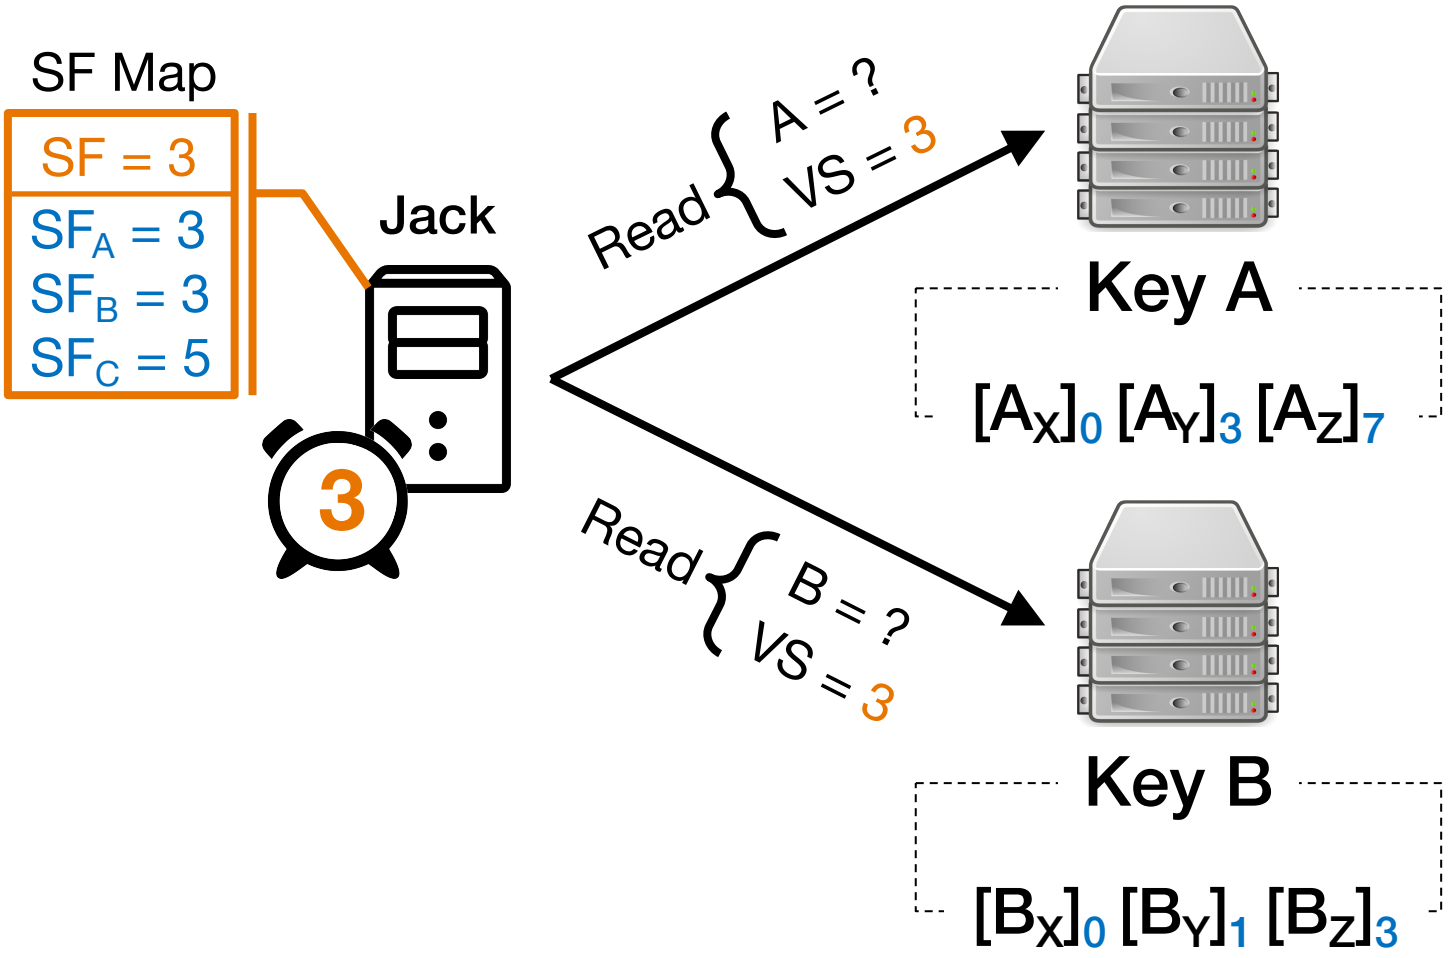

Read-Only Transaction Logic

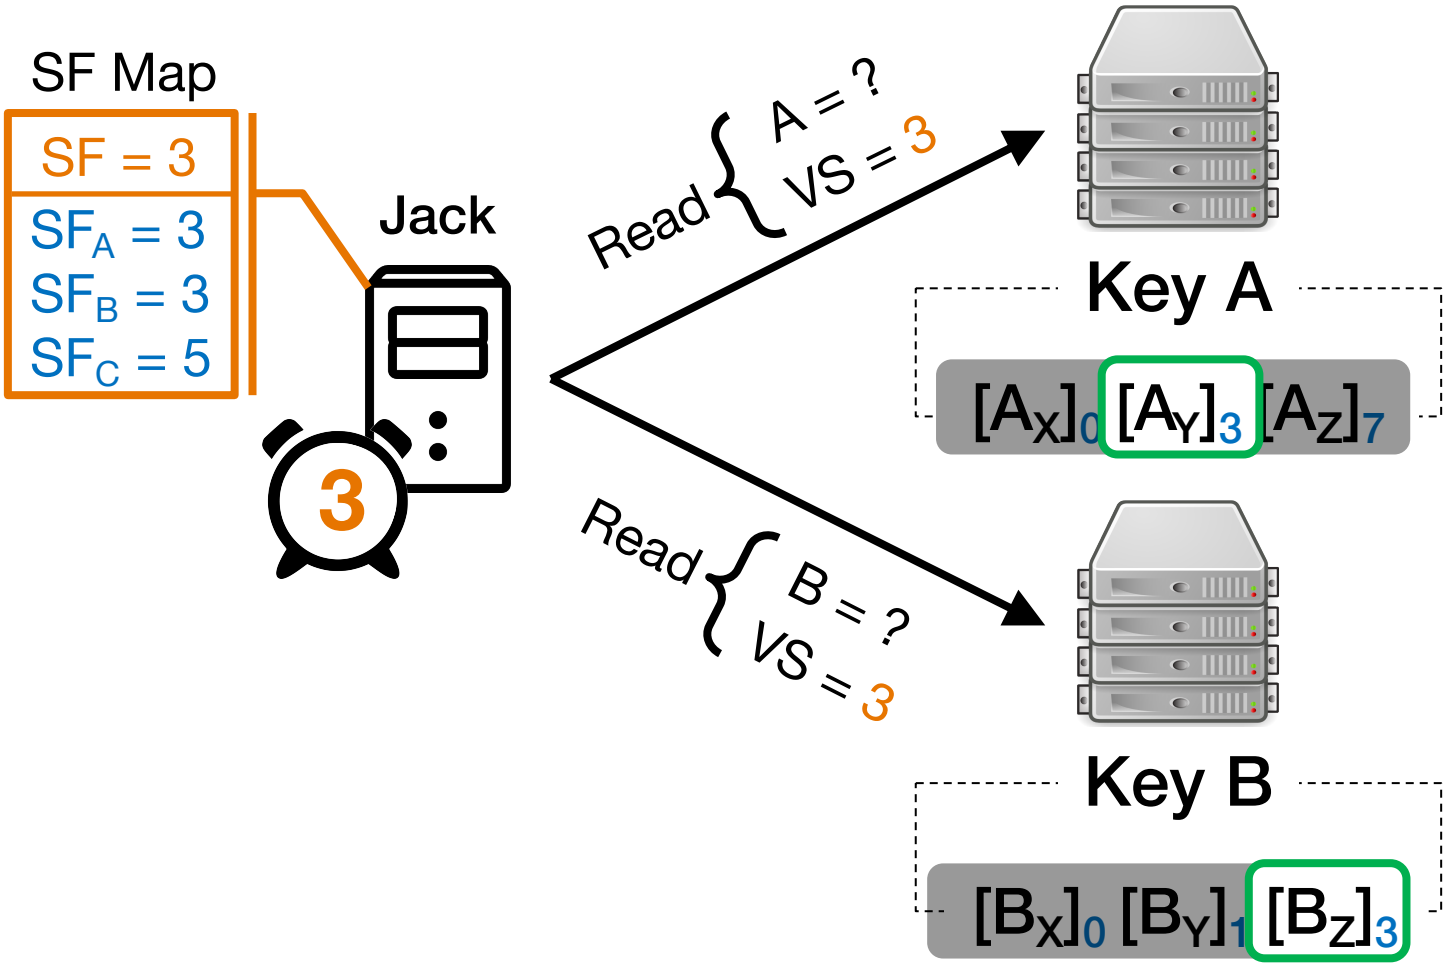

Read-Only Transaction Logic

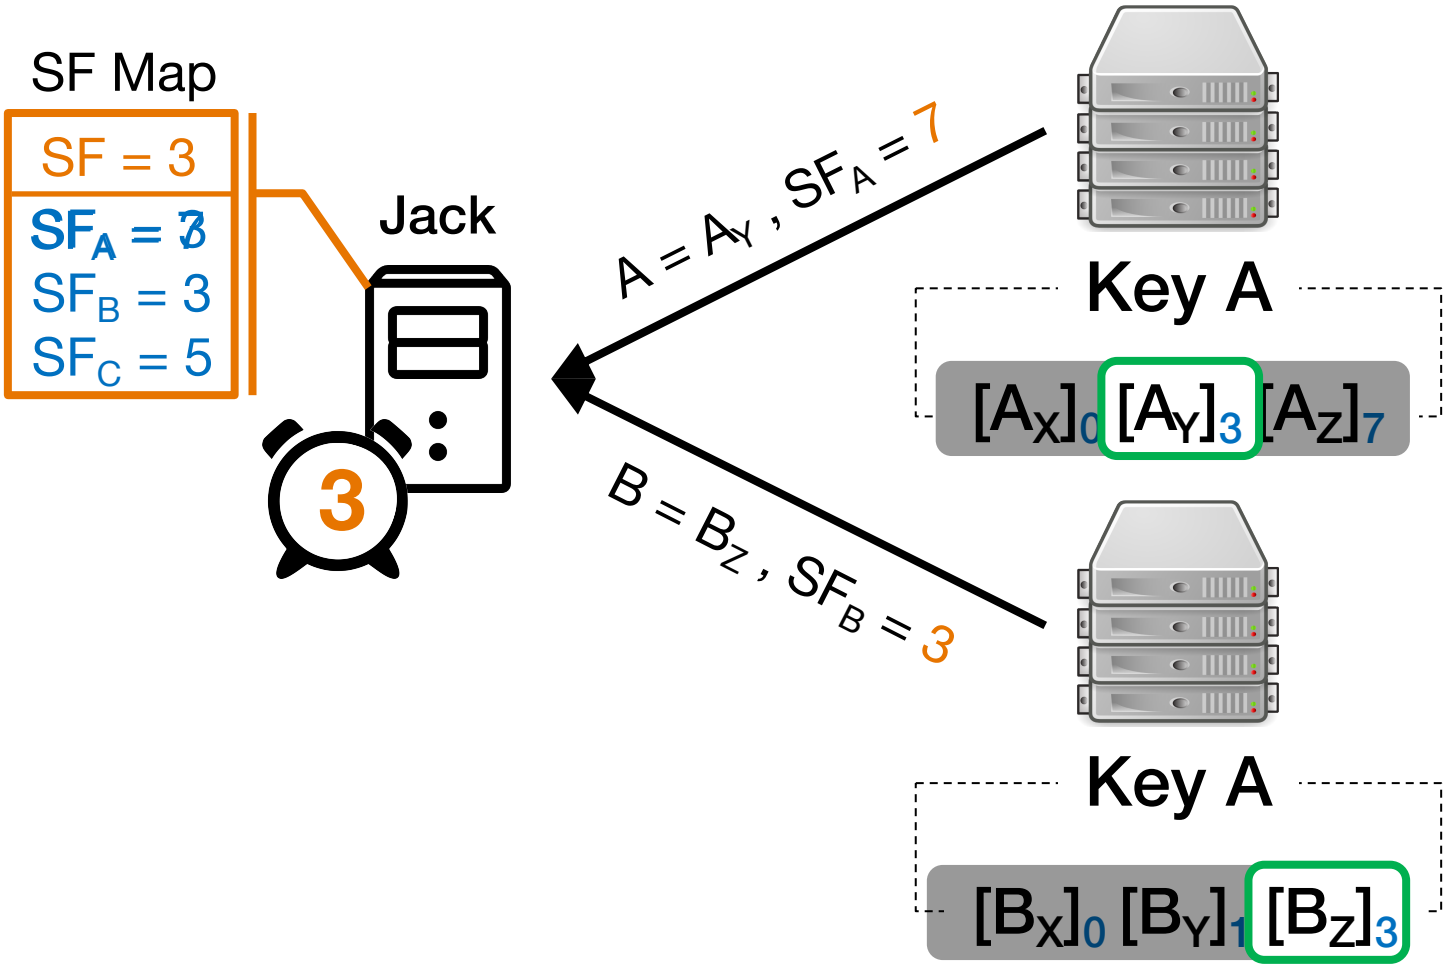

PORT Is NOC

- Reading at the stable frontier ensures reads are non-blocking (N)
- Client pre-determined snapshot with VS ensures one-round communication (O)
- One VS per read request ensure constant metadata (C)

PORT Systems

- Scylla-PORT
  - Base system: ScyllaDB (non-transactional)
    - Highly optimized → sensitive to overhead
  - NOC + Process-ordered serializability
  - Supports simple writes (not write transactions)

- Eiger-PORT
  - Base system: Eiger (N, ~~O~~, ~~C~~)
    - Existing read-only and write transactions
  - NOC + Causal consistency
  - Supports write transactions

Evaluation of Scylla-PORT

- To understand
  - Overhead in latency and throughput compared to simple reads
  - Performance advantages compared to other protocols, e.g., OCC.

- Experiment configuration
  - YCSB benchmark with customized parameters for skew and read-to-write ratios
  - Evaluated latency, throughput, scalability, freshness

Latency-Throughput
Uniform, 5% Writes

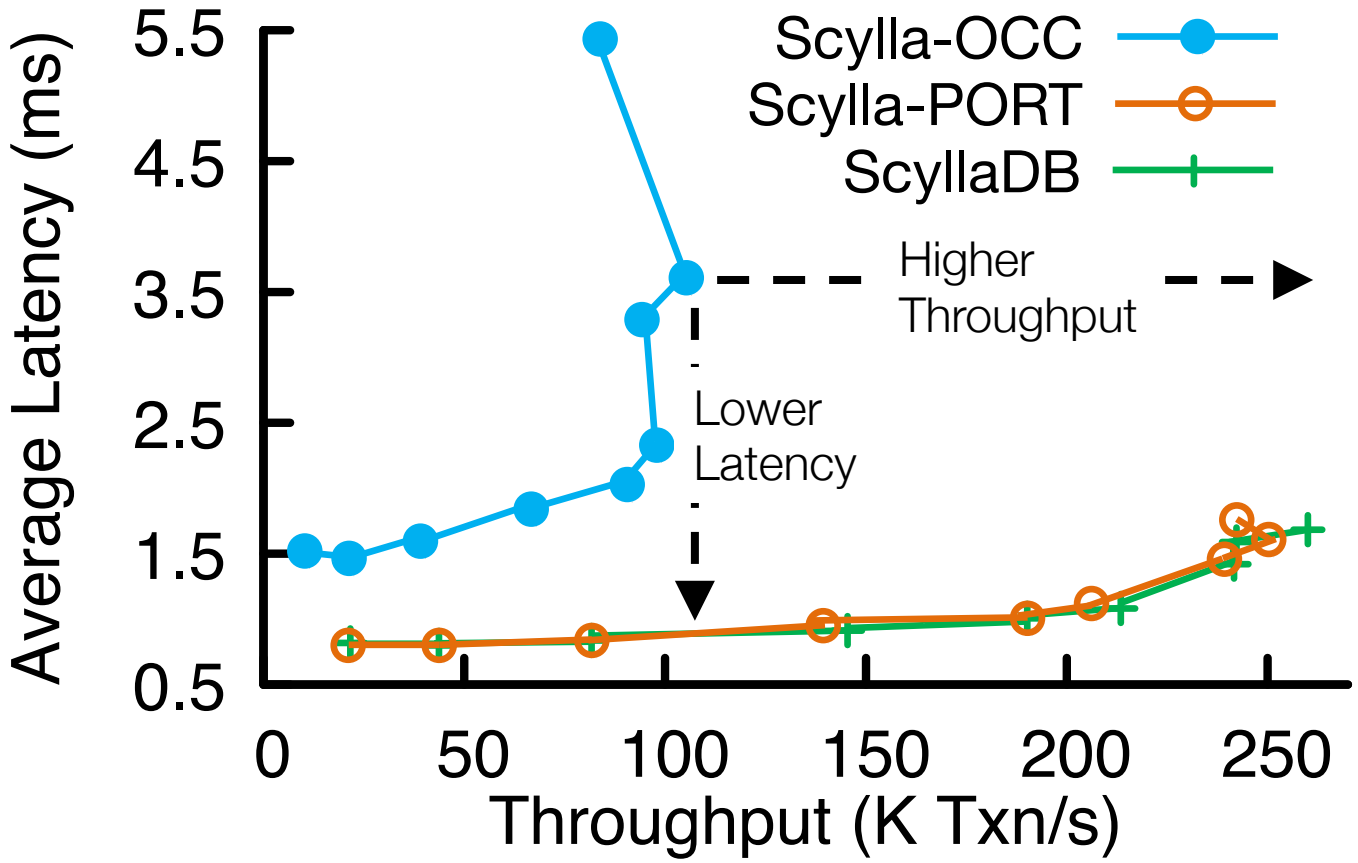

Latency-Throughput

Zipf = 0.99, 5% Writes

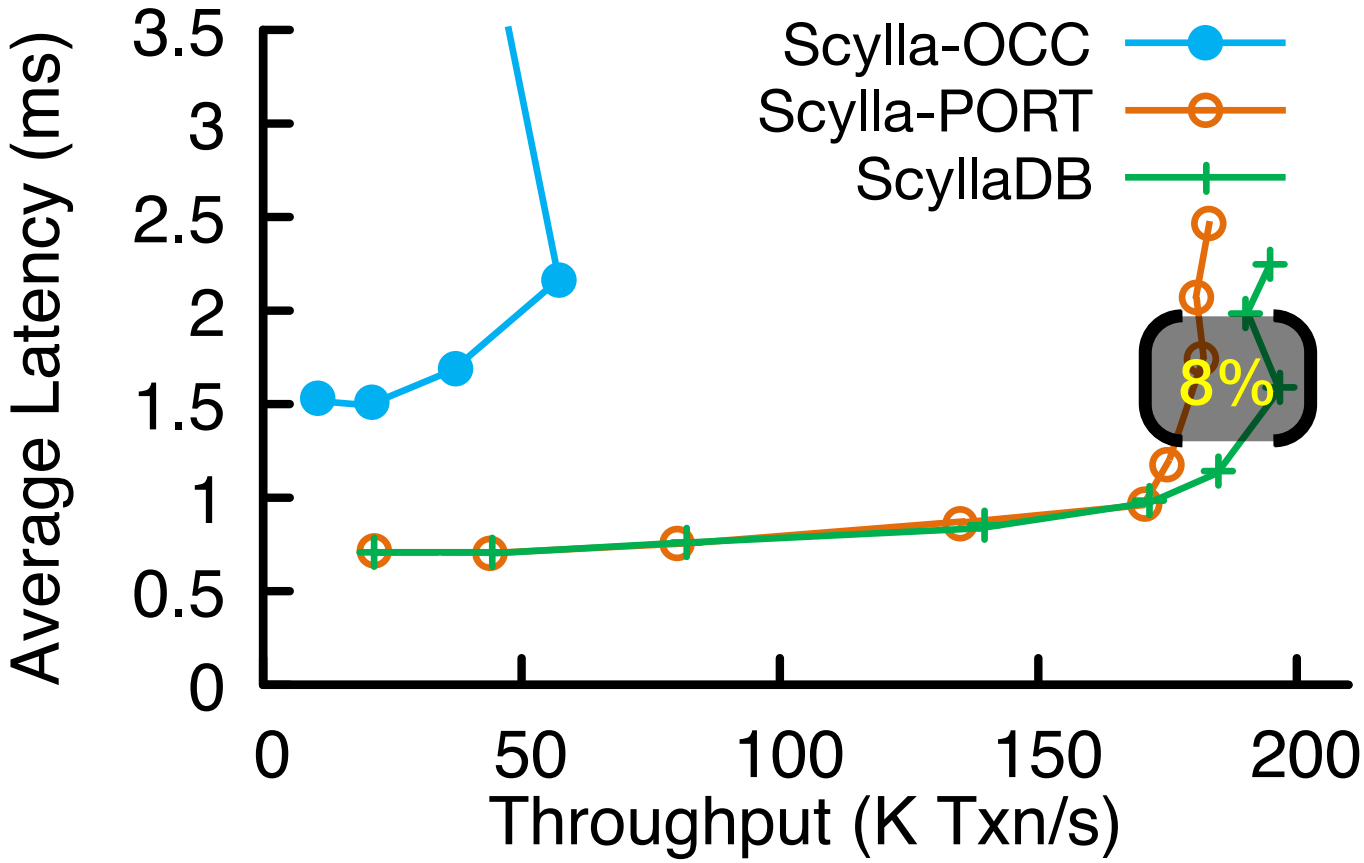

Conclusion

- Performance-optimal read-only transactions: NOC
- The NOCS Theorem for read-only transactions
  - Impossible to have all of the NOCS properties
- The design of PORT
  - NOC with the strongest consistency to date
- Scylla-PORT
  - Minimum performance overhead compared to simple reads
  - Significantly outperforms the standard OCC

Contact Information
Haonan Lu
haonanl@cs.princeton.edu