

High Performance State-Machine Replication

Parisa Jalili Marandi
jalilimp@usi.ch

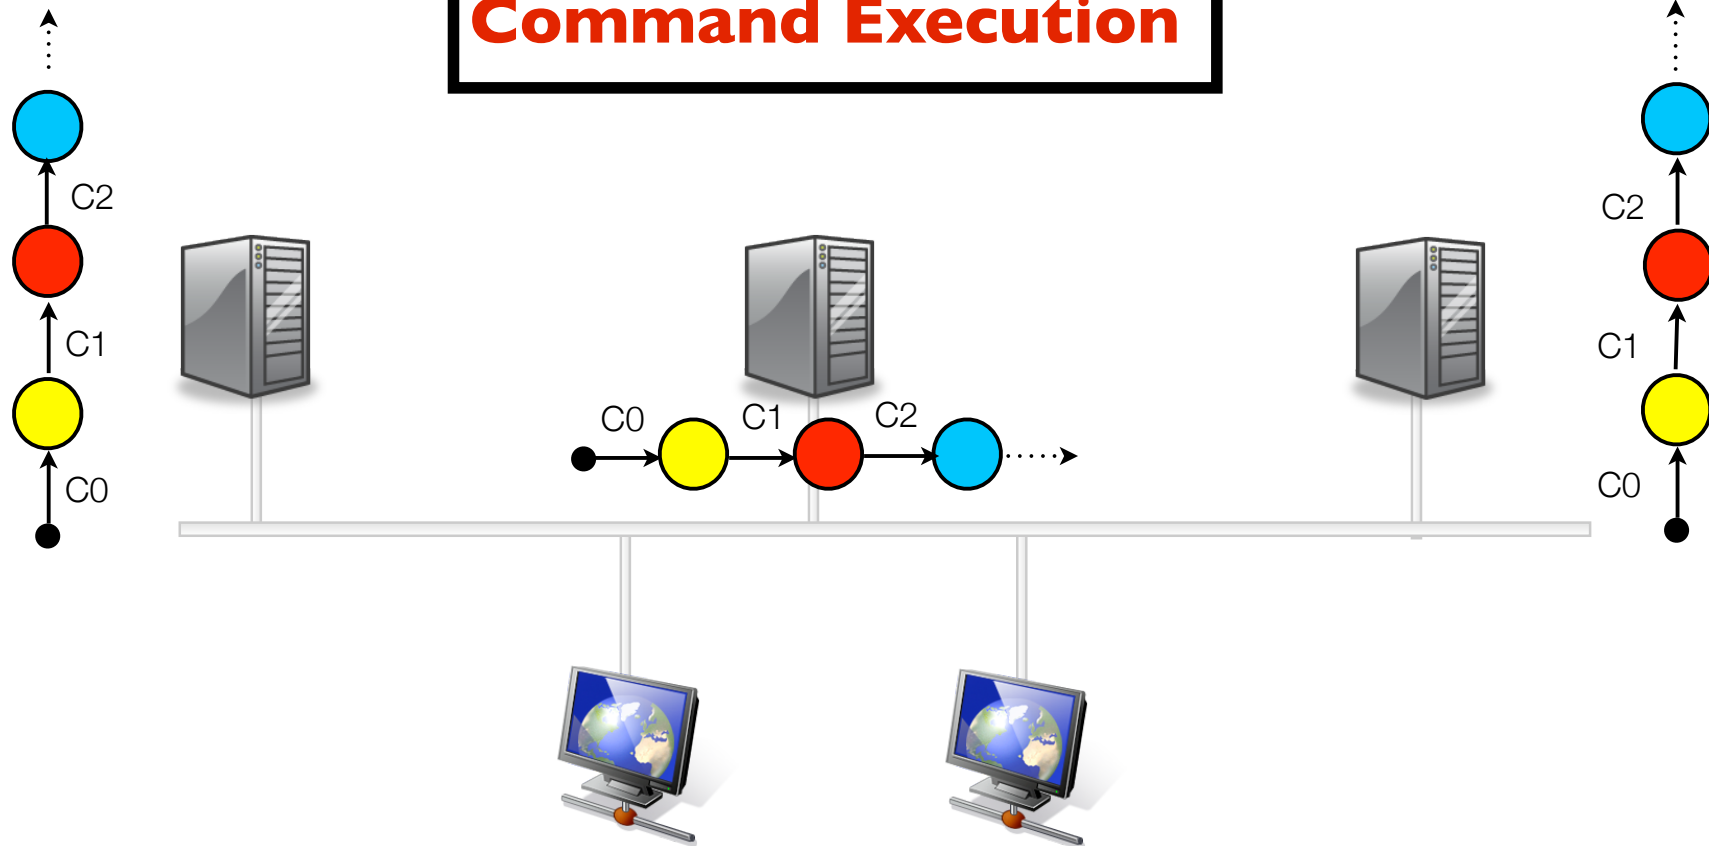
Supervisor: Prof. Fernando Pedone

High Performance State-Machine Replication

- ▶ SMR an approach to fault tolerance
- ▶ Performance: Throughput and Latency
- ▶ High Performance
 - **High Throughput**
 - **Low Latency**

State Machine Replication (SMR)

**Command Ordering
Command Execution**



Contributions

- ▶ **Efficient Ordering of Commands**
 - **Ring Paxos**

- ▶ **Inherent limitations of SMR**
 1. Ordering increases latency
 - **Speculation**
 2. SMR is not scalable
 - **Partitioning**

Ring Paxos

- ▶ A high-throughput AB-protocol based on Paxos
- ▶ Features: ring - multicasting - a majority of acceptors - indirect consensus
- ▶ **90%** of bandwidth ~ 900 Mb/s in a 1 Gb network
- ▶ Open source

Speculation & State Partitioning

▶ Speculation

- Parallelize Ordering and Execution
- Improves Latency by **16%**

▶ State Partitioning

- Replicas gain network and processing capacity
- Makes SMR scalable
- Improves Throughput by **4 times**

Future Work

- ▶ Extend Ring Paxos for x-LAN settings
- ▶ Dynamic addition and removal of learners
- ▶ Generality of State Partitioning

