Concurrency Distributed 2

COS 316 Precept 8

Outline

- Consistency Model Practice
- Replicated State Machines (Primary- Backup)

Outline

- Consistency Model Practice
- Replicated State Machines (Primary- Backup)

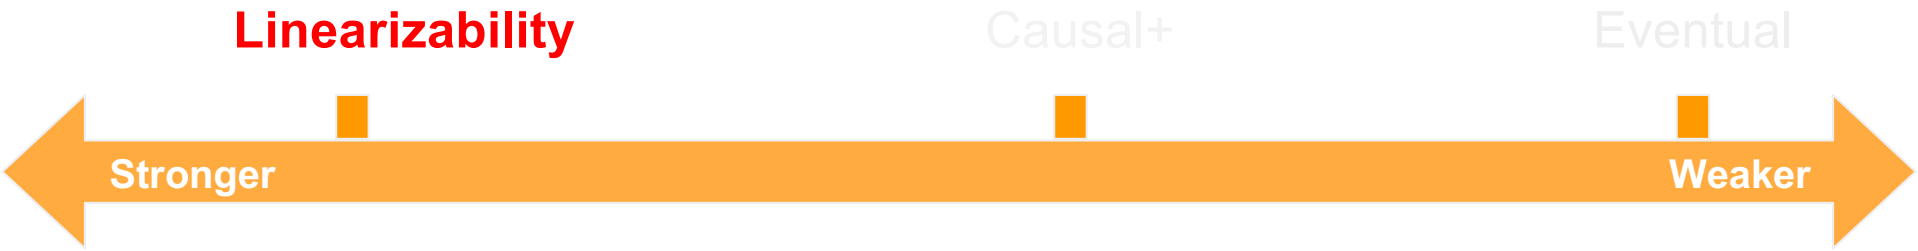
Linearizability
Causal+
Eventual
Stronger
Weaker

Linearizability

Total order: There exists some legal total order of operations.

Preserves real-time ordering: if an operation *A* completes before operation *B* begins, then op A occurs before op *B* in the total order.

A completed write op is visible to all future read ops.

Intuition: once a read “sees” a new write, all future reads must also “see” that write.

Linearizability Example

Linearizable? **No**

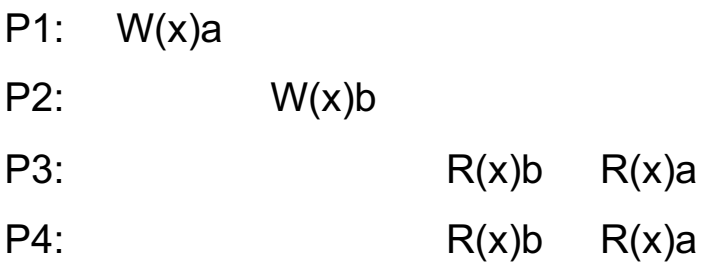

Linearizable? **Yes**

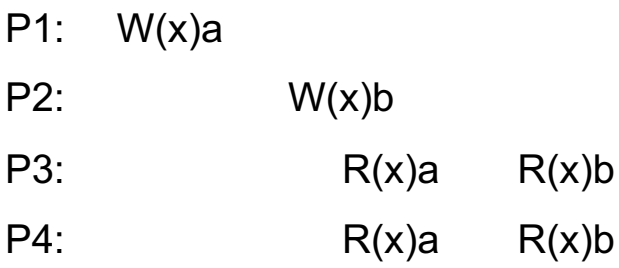

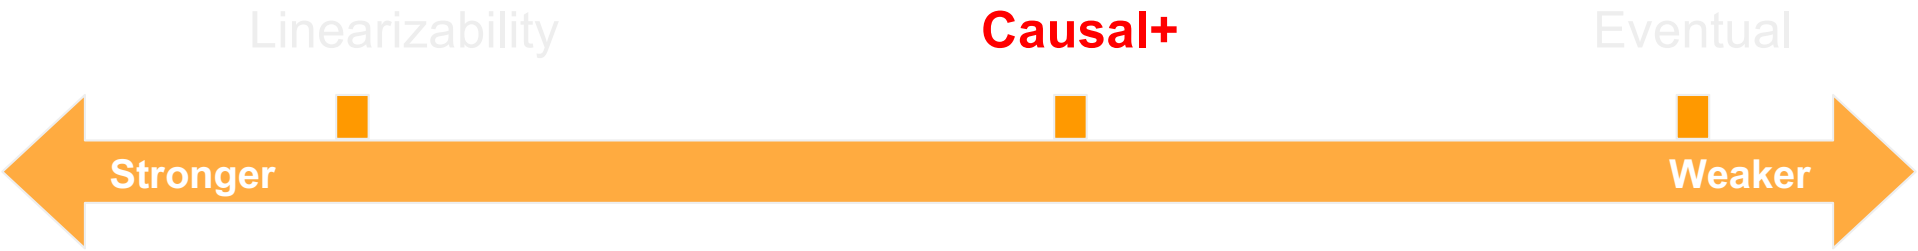
Linearizability
Causal+
Eventual
Stronger
Weaker

Causal+ Consistency

Partial order: order causally related ops the same way across all processes

+: replicas' total order eventually converges.

Difference from *Linearizability*?

Only causally related ops need to be ordered: **no guaranteed total order.**

Concurrent ops may be ordered differently across different processes.

Review Concurrency / Causality

Ops	Concurrent
a,b	No
a,e	Yes
a,g	No
c,e	Yes
c,d	No
d,g	No
d,f	No
e,g	No
a,d	No

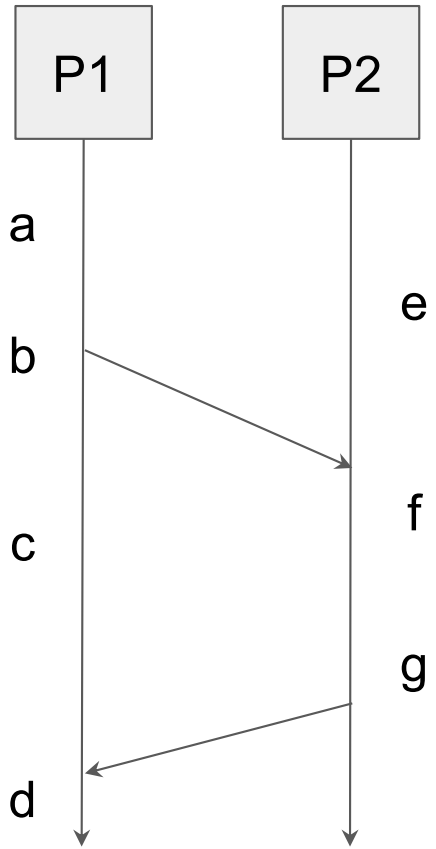

Causally+ Consistent? Yes

P1: W(x)a

P2: W(x)b

P3: R(x)b R(x)a

P4: R(x)a

Causally+ Consistent? No

P1: W(x)a

P2: R(x)a W(x)b

P3: R(x)b R(x)a

P4: R(x)a

Outline

- Consistency Model Practice
- Replicated State Machines (Primary- Backup)

Why Using Replicated State Machines?

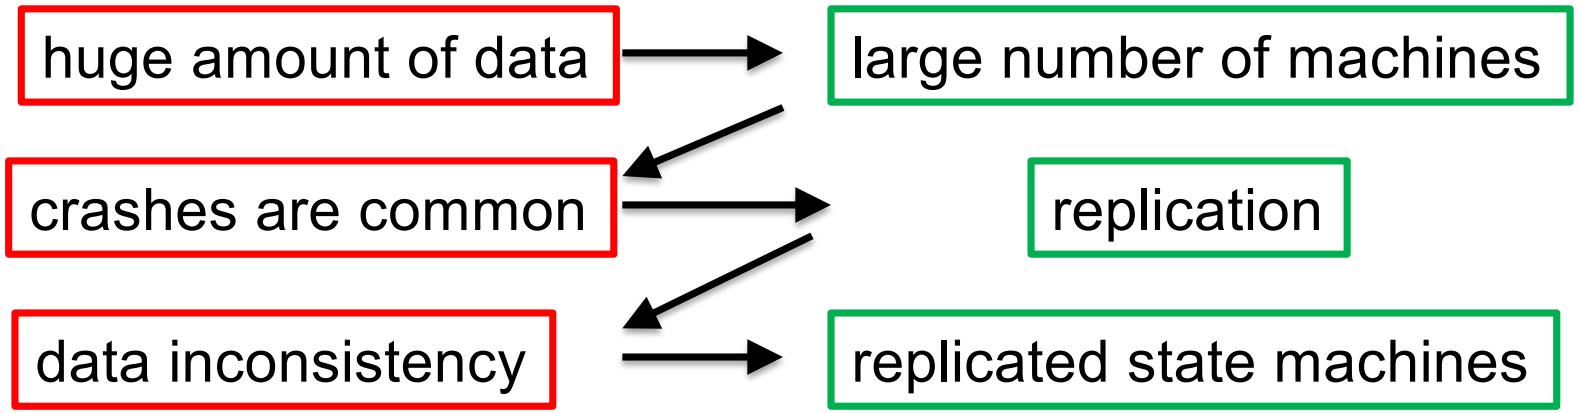

Why Using Replicated State Machines?

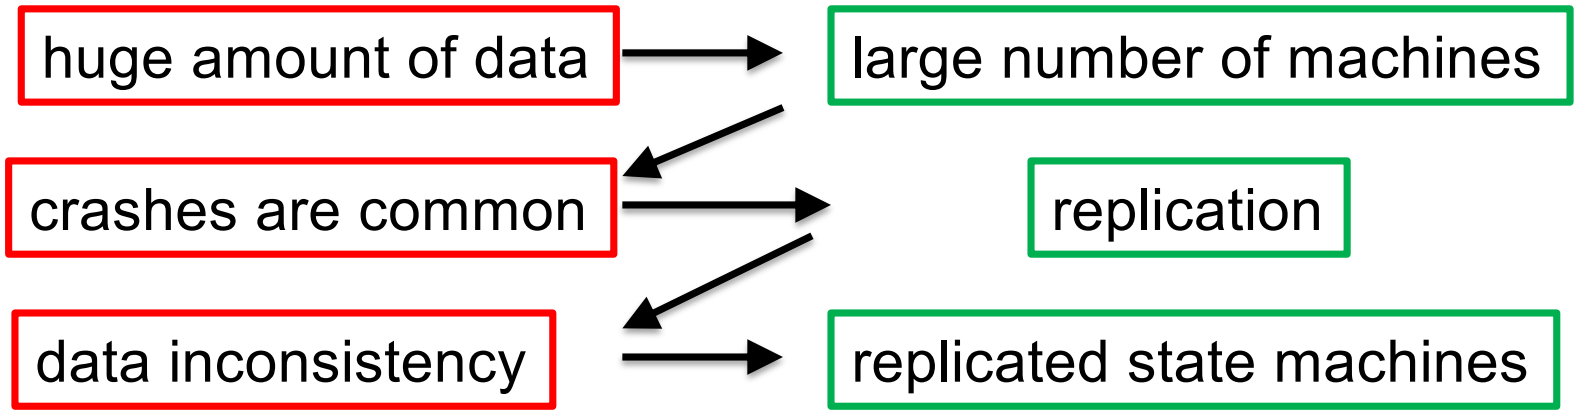

Start from the same initial state,
same sequence of state transitions

Primary-Backup

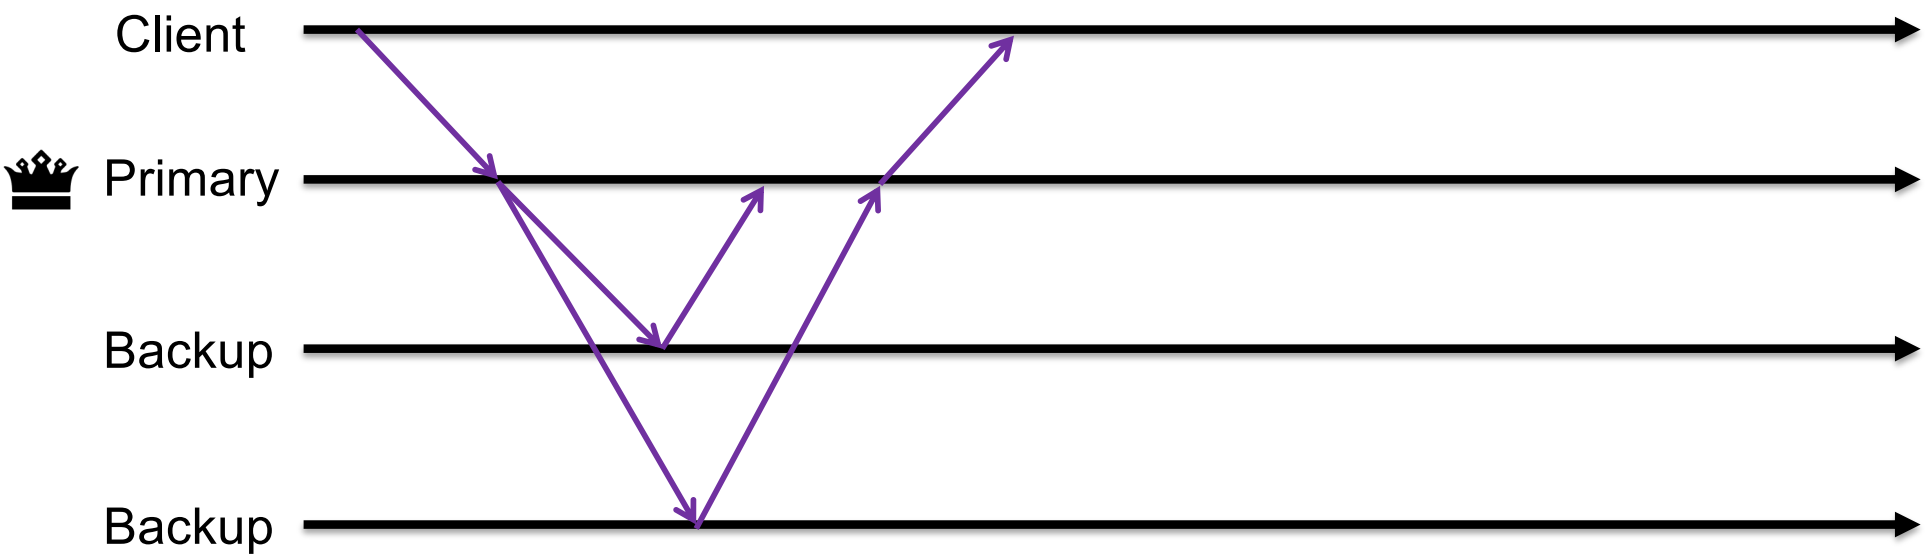

Primary-Backup

- The primary decides orders of client requests
- The primary commits a request after hearing back from all backups
- A request can only be executed if it is committed and all previous requests in the log are executed
- After executing a request, the primary replies to the client

Compared to a single machine:
primary-backup system has higher latency (more rounds of messages)
lower throughput (the primary is the bottleneck)

If adding more backups:
latency goes up (the primary needs to hear back from all backups)
throughput goes down (the primary processes more messages)