CS 425 / ECE 428
Distributed Systems
Fall 2026

Aishwarya Ganesan & Ram Kesavan

Lecture 22: Replication Control

Ack: Indranil Gupta

All slides © IG

In-class quiz (for participation only)

- Canvas “In Class Quiz3”
- Access Code: transactions

Announcements

- MP3 due Sunday 11/8, demos on 11/9
  - **Demo Signup by Friday 11/6, 11:59pm**
- HW4 released, due 12/3 (**Thursday**) at 11:59 pm
  - Deadline will creep up on you! Start a problem as soon as the topic is discussed
- Final Exam dates/times: course web-page & Piazza



Server-side Focus

- Concurrency Control = how to coordinate multiple concurrent clients executing operations (or transactions) with a single server

Next:

- Replication Control = how to handle operations (or transactions) when there are objects are stored at multiple servers, with or without replication

Replication: What and Why

- Replication = An object has identical copies, each maintained by a separate server
  - Copies are called "replicas"
- Why replication?
  - Fault-tolerance: With *k* replicas of each object, can tolerate failure of any (*k-1*) servers in the system
    - Replicas should ideally be in different failure domains
  - Load balancing: Spread read/write operations out over the *k* replicas => load lowered by a factor of *k* compared to a single replica
- Replication => Higher Availability

Availability

- If each server is down a fraction f of the time
  - Server's failure probability
- With no replication, availability of object =

  = Probability that single copy is up

  $= (1 - f)$
- With k replicas, availability of object =

  Probability that at least one replicas is up

  = 1 – Probability that all replicas are down

  $= (1 - f^k)$

Nines Availability

- With no replication, availability of object =
  $= (1 - f)$
- With k replicas, availability of object =
  $= (1 - f^k)$

Availability Table

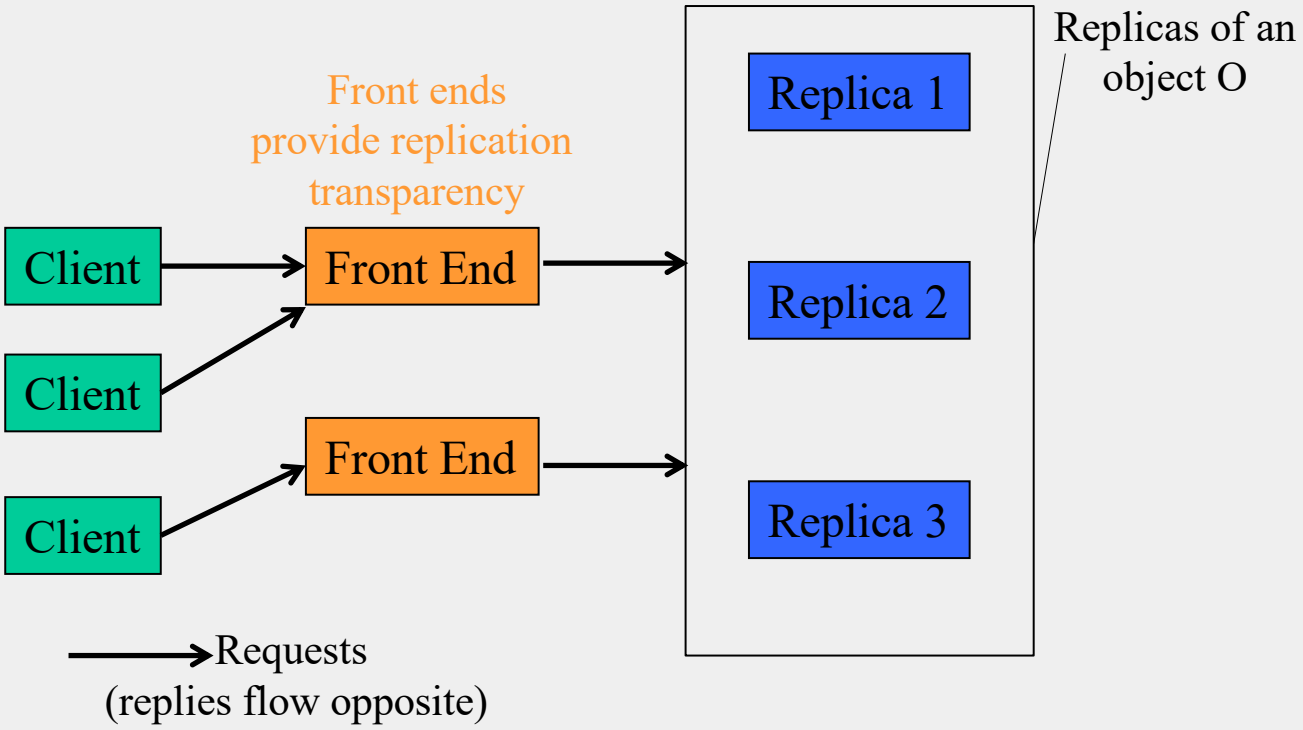
f=failure probability	No replication	k=3 replicas	k=5 replicas
0.1	90%	99.9%	99.999%
0.05	95%	99.9875%	6 Nines
0.01	99%	99.9999%	10 Nines

What's the Catch?

- Challenge is to maintain two properties

1. Replication Transparency
   - A client ought to be unaware of multiple copies of objects existing on the server side
2. Replication Consistency
   - All clients see single consistent copy of data, in spite of replication
   - For transactions, guarantee ACID

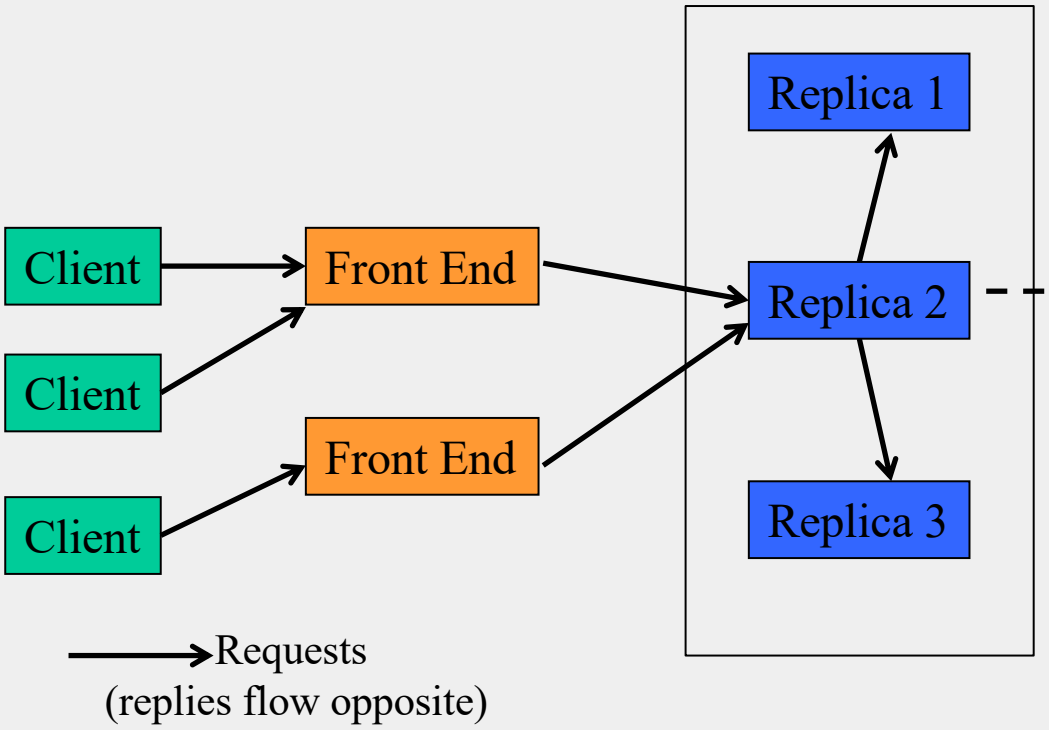
Replication Transparency



Replication Consistency

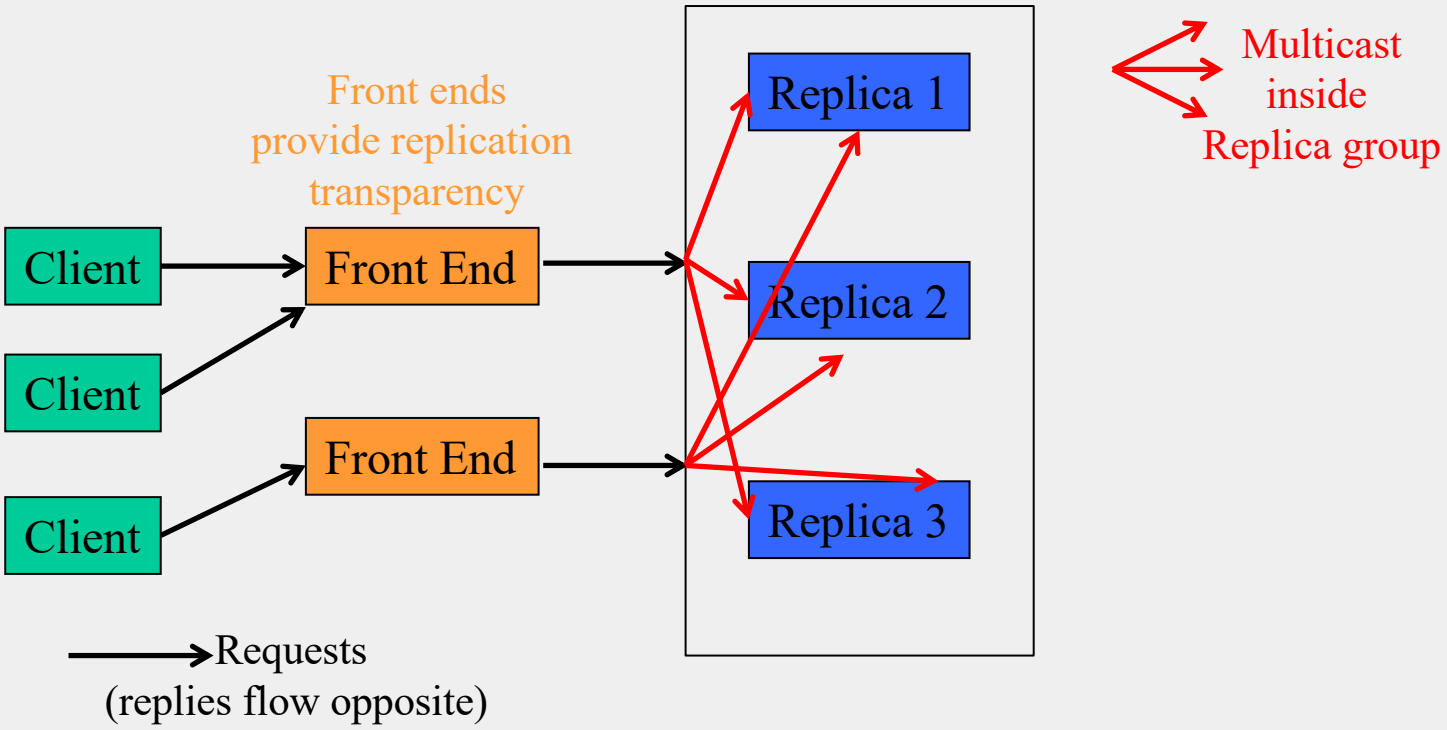
- Two ways to forward updates from front-ends (FEs) to replica group
  - Passive Replication: uses a primary replica (leader)
  - Active Replication: treats all replicas identically
- Both approaches use the concept of "Replicated State Machines"
  - Each replica's code runs the same state machine
  - *Multiple copies of the same State Machine begun in the Start state, and receiving the same Inputs in the same order will arrive at the same State having generated the same Outputs.* [Schneider 1990]

Passive Replication



- Leader => total ordering of all updates
- On leader failure, run election

Active Replication



Active Replication Using Concepts You've Learnt earlier

- Can use any flavor of multicast ordering, depending on application
  - FIFO ordering
  - Causal ordering
  - Total ordering
  - Hybrid ordering
- Total or Hybrid (*-Total) ordering + Replicated State machines approach

  => all replicas reflect the same sequence of updates to the object

Active Replication Using Concepts You've Learnt earlier (2)

- What about failures?
  - Use virtual synchrony (i.e., view synchrony)
- Virtual synchrony with total ordering for multicasts
  =>
  - All replicas see all failures/joins/leaves and all multicasts in the same order
  - Could also use causal (or even FIFO) ordering if application can tolerate it

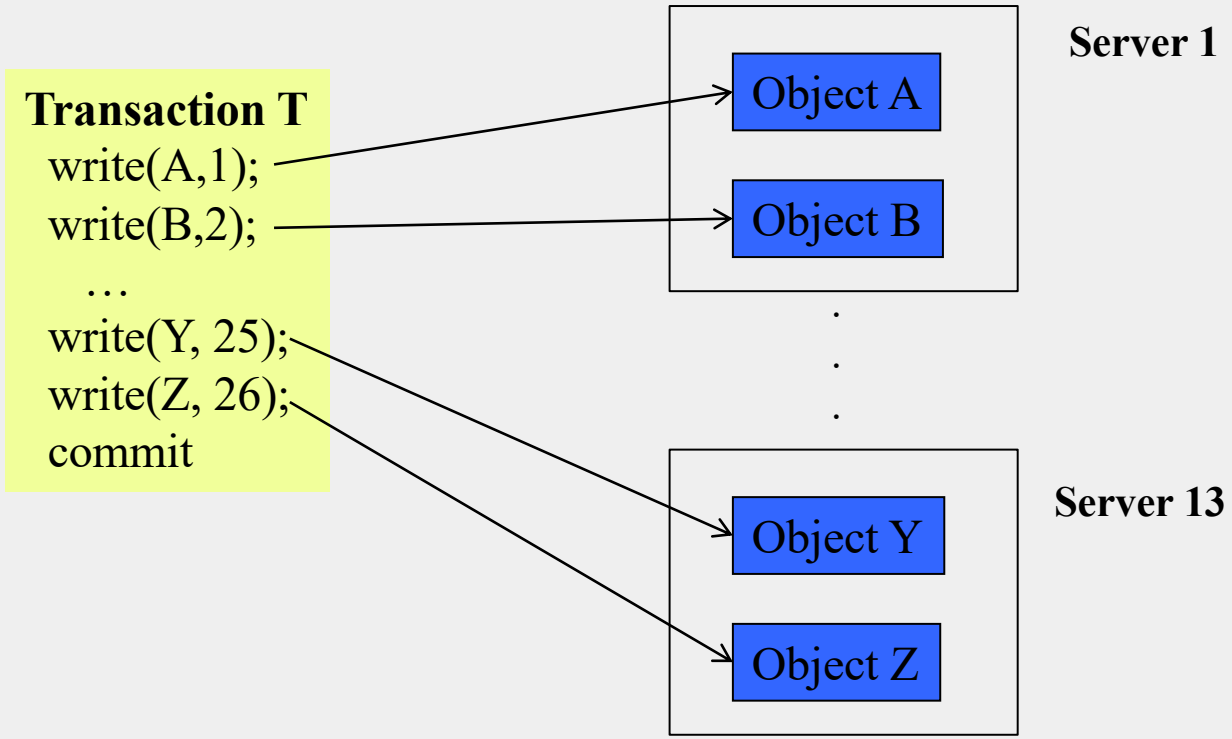
Transactions and Replication

- One-copy serializability
  - *A concurrent execution of transactions in a replicated database is one-copy-serializable if it is equivalent to a serial execution of these transactions over a single logical copy of the database.*
  - (Or) The effect of transactions performed by clients on replicated objects should be the same as if they had been performed one at a time on a single set of objects (i.e., 1 replica per object).
- In a non-replicated system, transactions appear to be performed one at a time in some order.
  - Correctness means serial equivalence of transactions
- When objects are replicated, transaction systems for correctness need one-copy serializability

Next

- Committing transactions with distributed servers

Transactions with Distributed Servers



Transactions With Distributed Servers

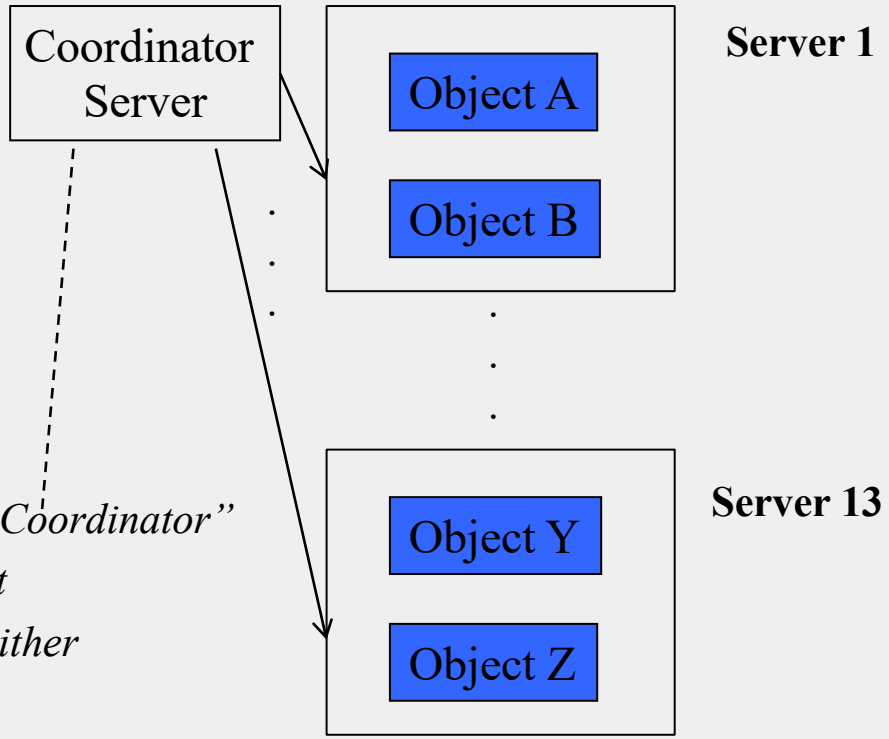
- Transaction T may touch objects that reside on different servers
- When T tries to commit
  - Need to ensure all these servers commit their updates from T => T successfully commits
  - Or none of these servers commit => T aborts
- What problem is this?

Transactions With Distributed Servers

- Transaction T may touch objects that reside on different servers
- When T tries to commit
  - Need to ensure all these servers commit their updates from T => T will commit
  - Or none of these servers commit => T will abort
- What problem is this?
  - Consensus!
  - (It's also called the "Atomic Commit problem")

One-phase Commit

Transaction T
write(A,1);
write(B,2);
…
write(Y, 25);
write(Z, 26);
commit

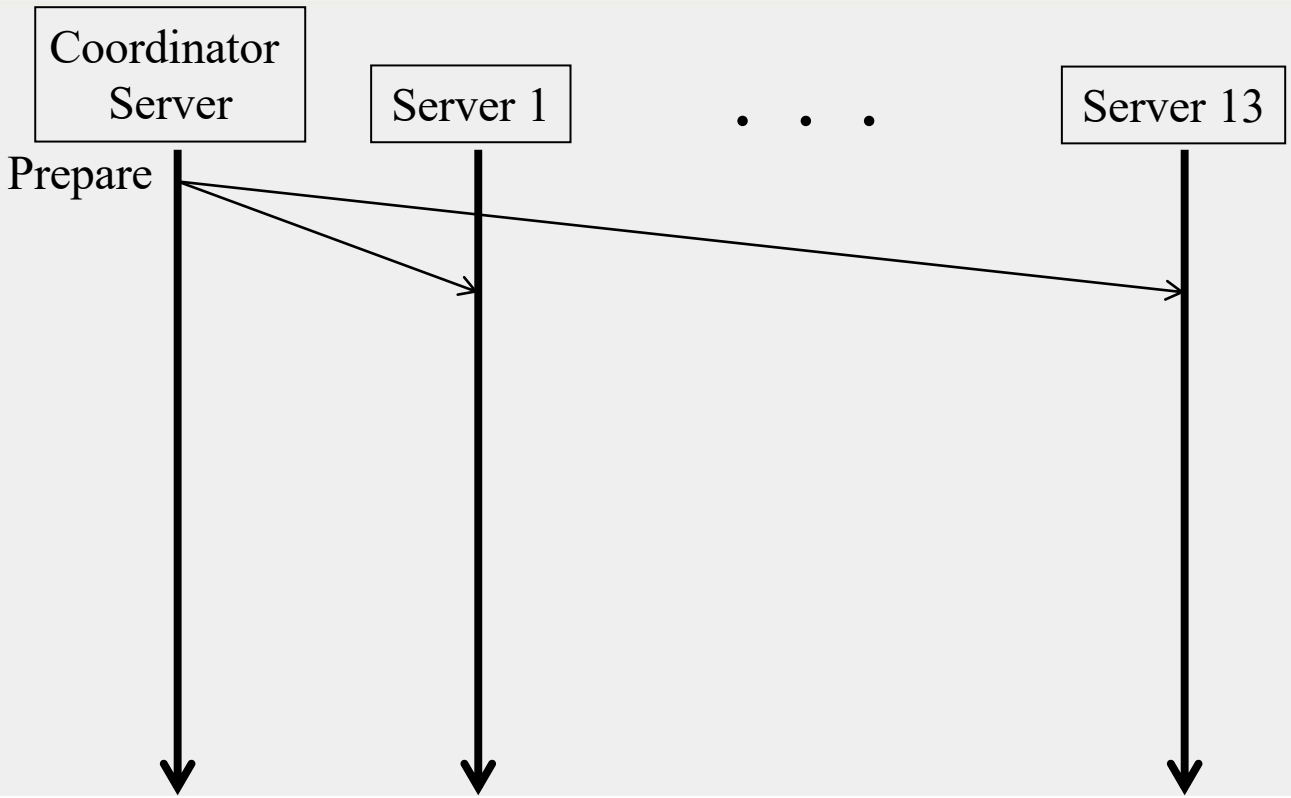


- *Special server called "Coordinator" initiates atomic commit*
- *Tells other servers to either commit or abort*

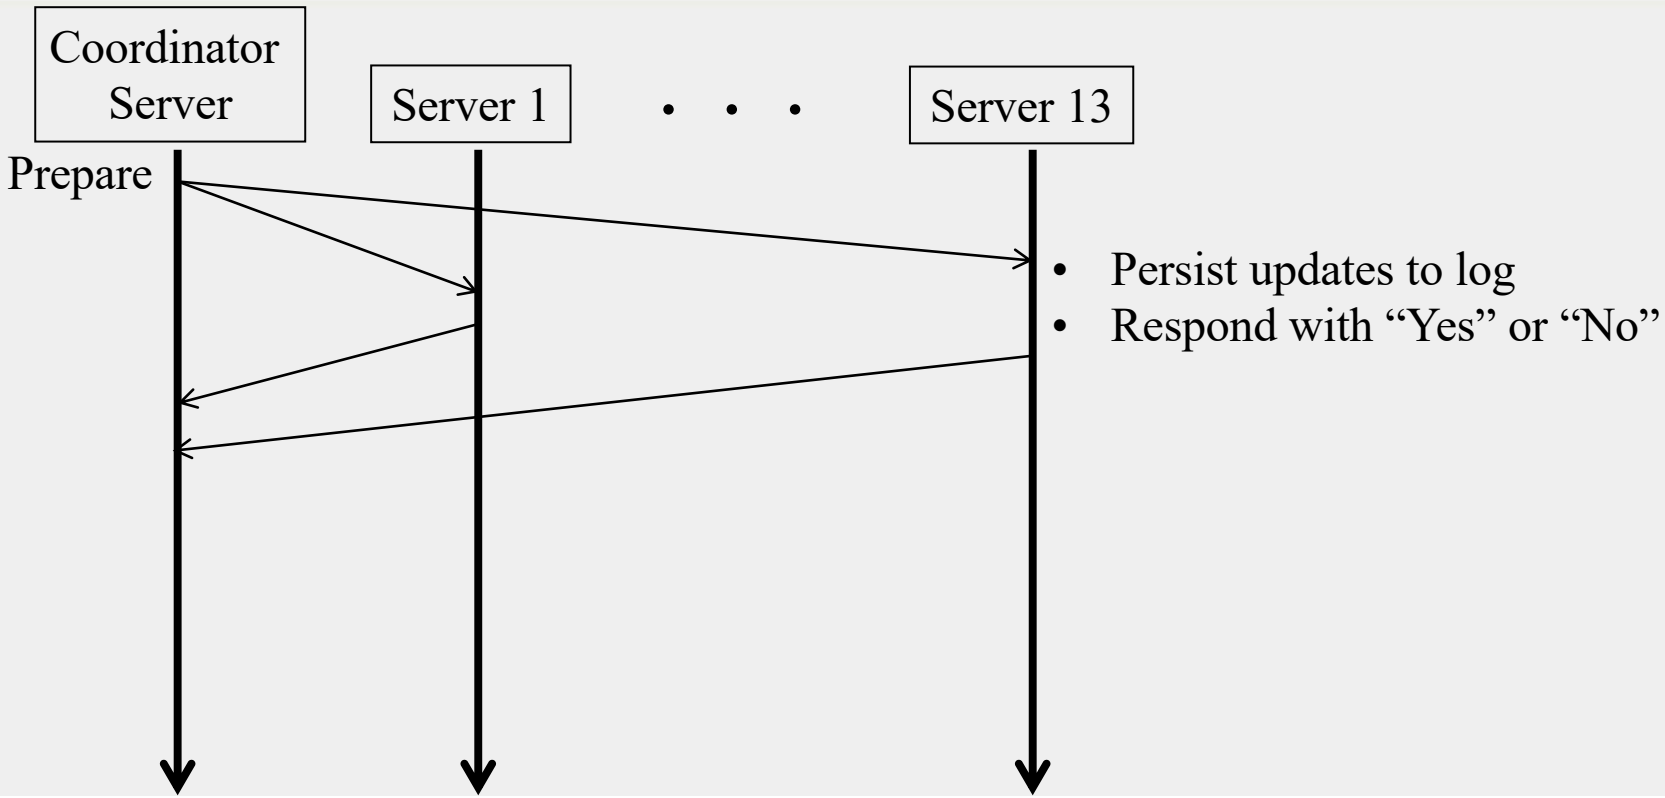
One-phase Commit: Issues

- Server with object has no say in whether transaction commits or aborts
  - If object corrupted, it just cannot commit (while other servers have committed)
- Server may crash before receiving commit message, with some updates still in memory

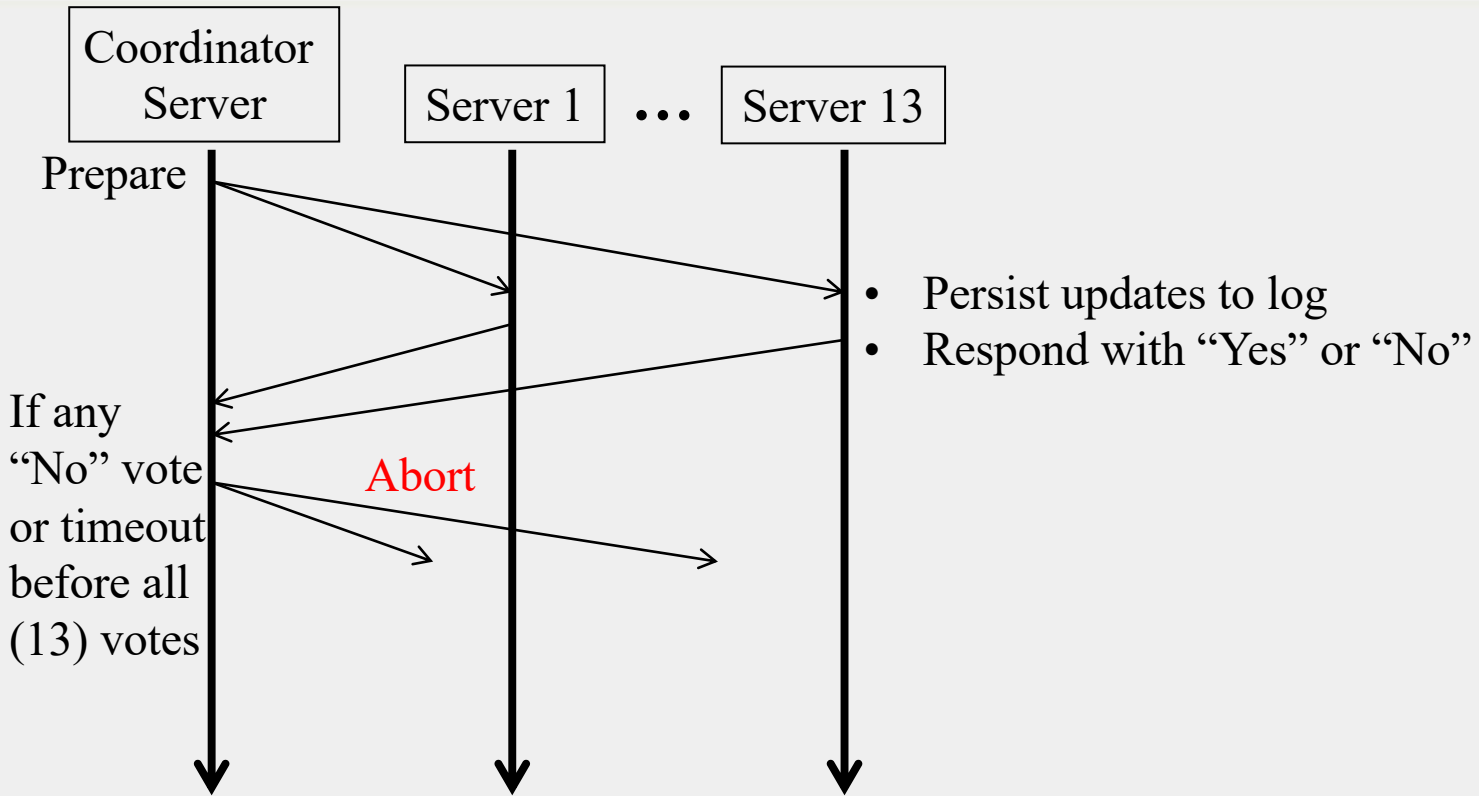
Two-phase Commit (aka 2PC)



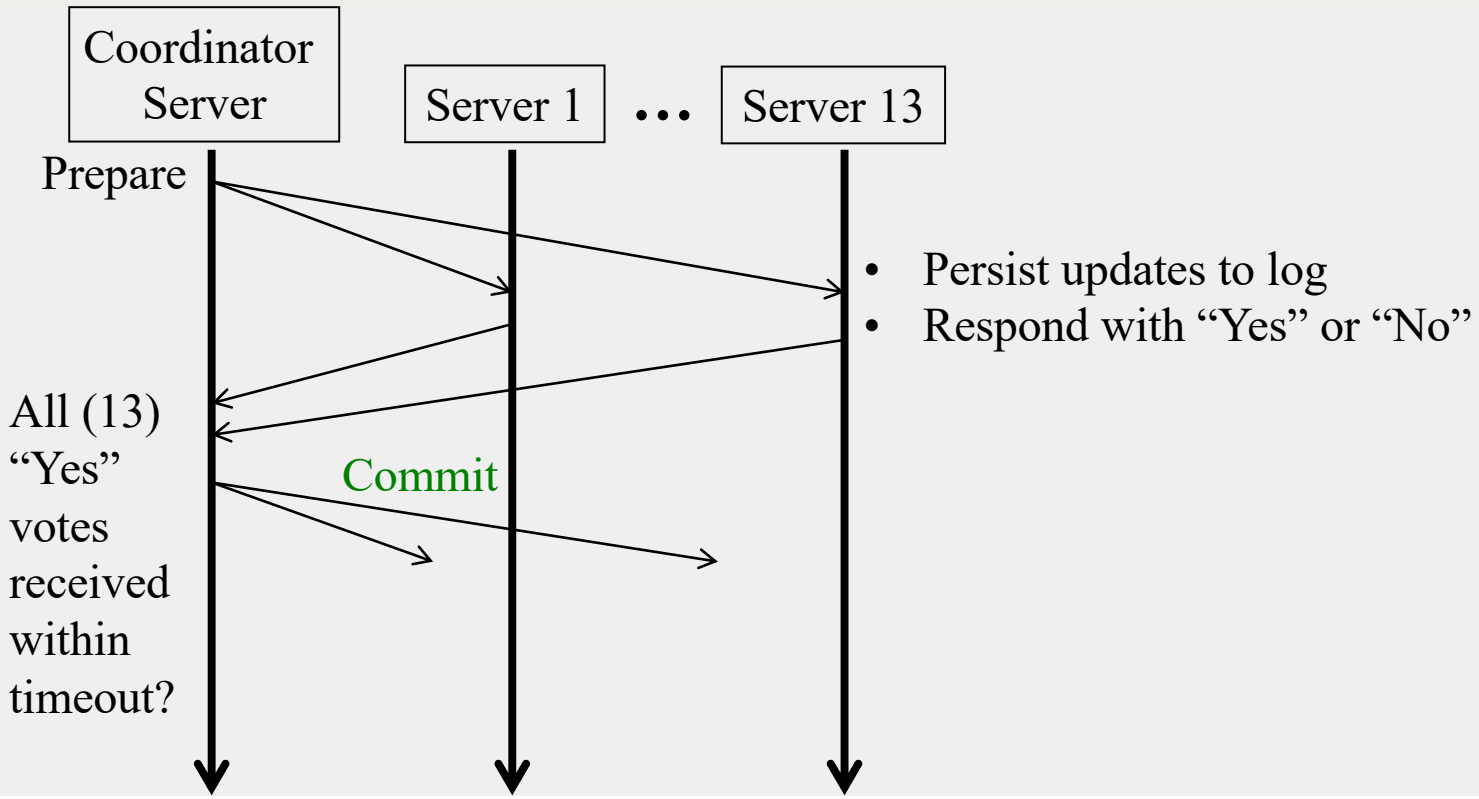
Two-phase Commit



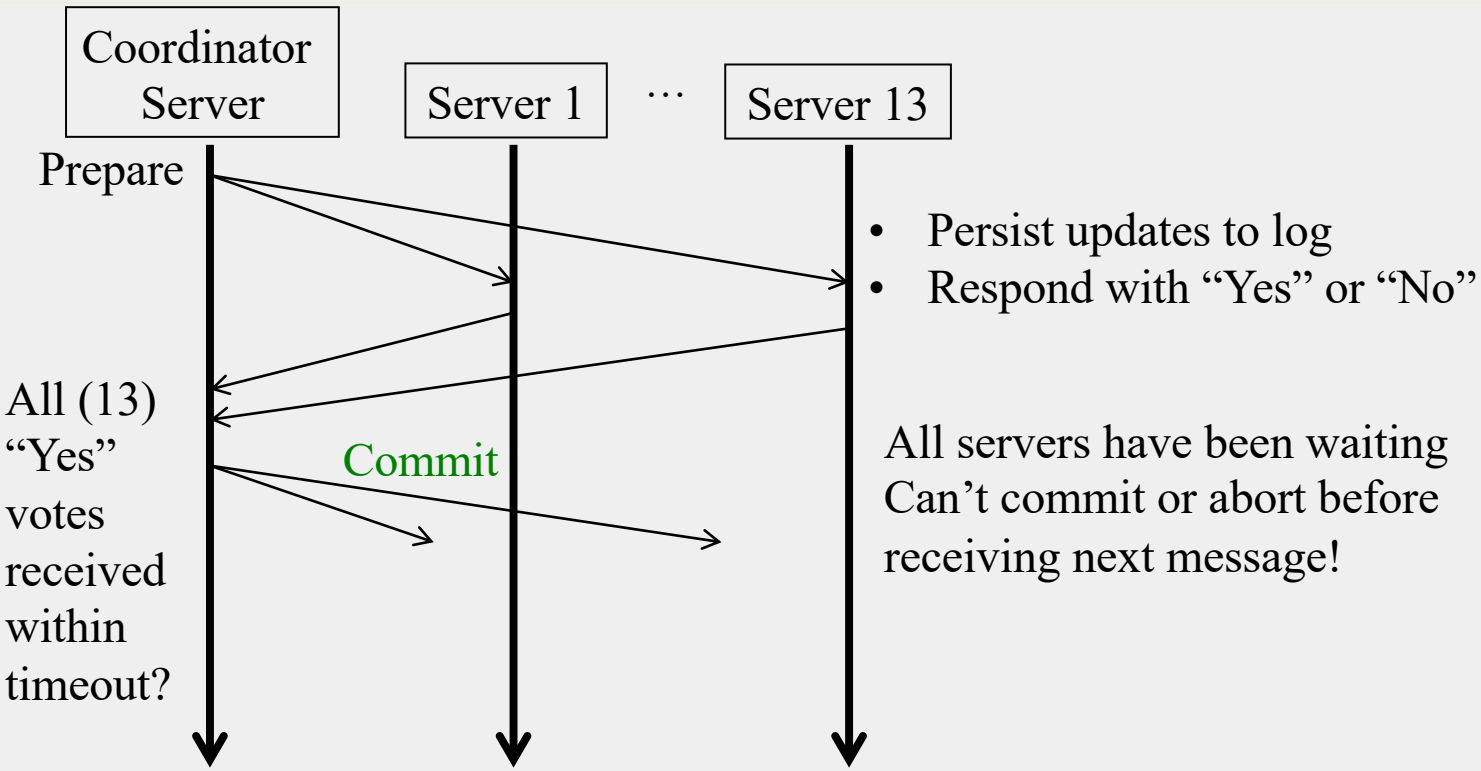
Two-phase Commit



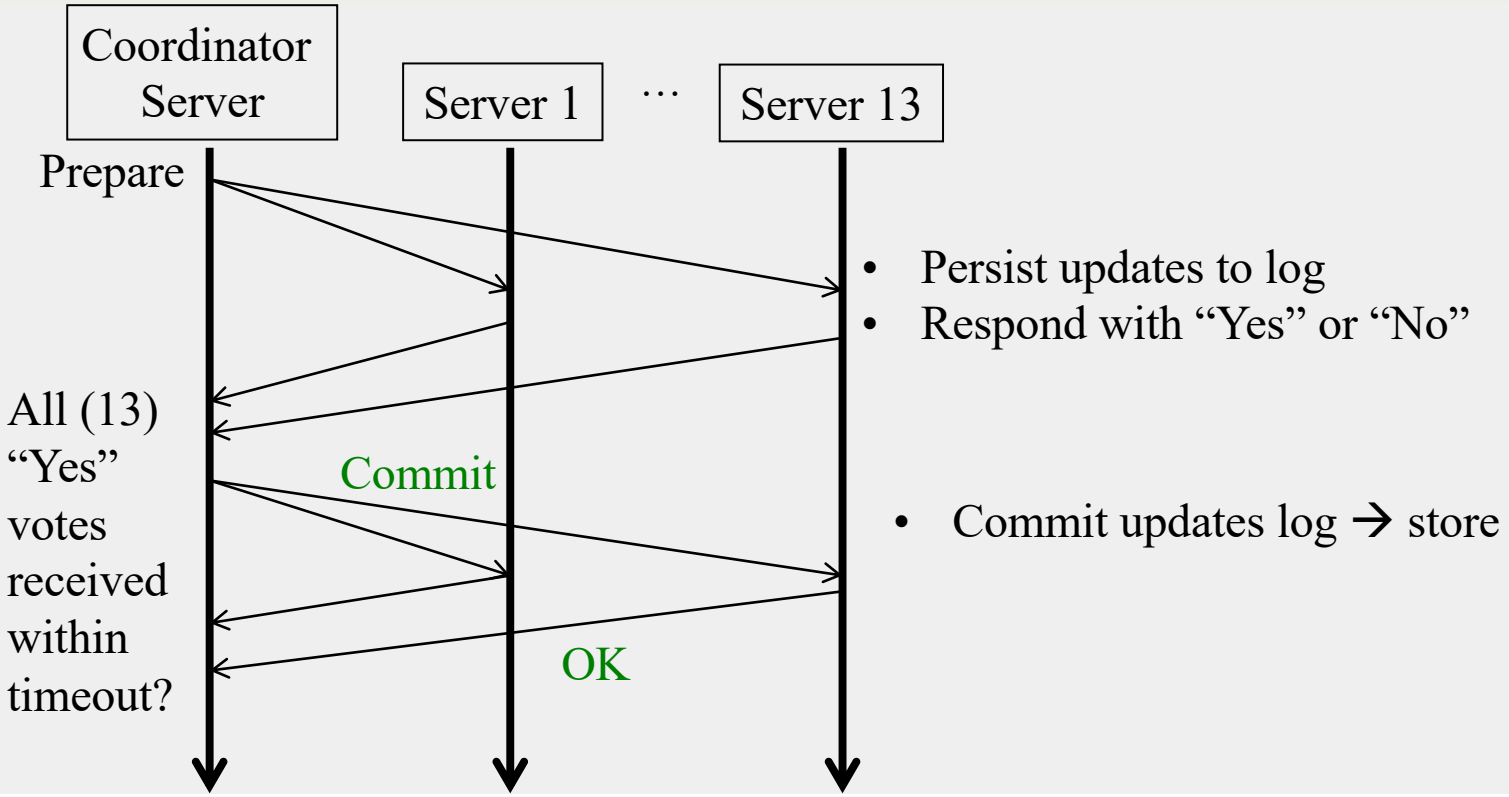
Two-phase Commit



Two-phase Commit



Two-phase Commit



Failures in Two-phase Commit

- If server voted Yes, it cannot commit unilaterally before receiving Commit message
- If server voted No, can abort right away (why?)
- What if server crashes?
  - Each server logs tentative updates into persistent storage, right before replying Yes/No in first phase. Use log to recover after crash.
- What if coordinator crashes?
  - Coordinator logs all decisions and received/sent messages to persistent storage
  - After recovery or new election => new coordinator takes over

Failures in Two-phase Commit (2)

- What if one of the Prepare messages is lost?
  - The server may decide to abort unilaterally after a timeout for first phase (server will vote No, and so coordinator will also eventually abort)
- What if a response to the Prepare is lost?
  - Coordinator pessimistically aborts the transaction after a timeout. It then sends the Abort message to all.
- What if a 2nd phase (commit or abort) message is lost?
  - Server must poll coordinator (repeatedly)
  - Coordinator cannot reclaim its log for a transaction until all servers have responded in the 2nd phase with an Ok.

Using Paxos in Distributed Servers

Atomic Commit

- Can use Paxos (instead of 2PC) for commit-v-abort decision
- But need to ensure that if any server votes No, everyone aborts

Ordering updates

- Replica group (for an object) can use Paxos to order all updates; iteratively do:
  - Server proposes message for next sequence number
  - Group reaches consensus (or not)

Use combination of 2PC and Paxos (e.g. Spanner)

  - One for commit-v-abort and the other for replica group consensus

Summary

- Multiple servers in cloud
  - Replication for Fault-tolerance
  - Load balancing across objects
- Replication flavors using concepts we learnt earlier
  - Active replication
  - Passive replication
- Transactions and distributed servers
  - Two phase commit

Announcements

- MP3 due Sunday 11/8, demos on 11/9
  - **Demo Signup by Friday 11/6, 11:59pm**
- HW4 released, due 12/3 (**Thursday**) at 11:59 pm
  - Deadline will creep up on you! Start a problem as soon as the topic is discussed
- Final Exam dates/times: course web-page & Piazza