

CSE 550: *Systems for all*

Au 2021

Ratul Mahajan

Faults

Paxos does **fail-stop** faults

- Nodes are dead or follow the protocol
- Hard part: Is the node dead or network too slow?

But faults can be **Byzantine**

- Incorrect protocol implementation
- Willful lying

Handling Byzantine faults

Replicated state machine model

- Replicated servers hold long-term state
- Clients read and modify the state

Fault model

- Up to f replicas can fail
- Same attacker can control all f replicas
 - Knows what the honest replicas do
 - Can read all network messages
- Clients aren't faulty
- Crypto cannot be broken

Strawman solution

$2f + 1$ replicas

Clients sends a command to all replicas

Waits for $f + 1$ identical replies

- With a max of f faults (and network eventually working), this will happen

What can go wrong?

Problem with the strawman

$f+1$ matching replies might be f bad nodes & 1 good

- so maybe only one good node got the operation
- next operation also waits for $f+1$
- might not include that one good node that saw op1

example: S1 S2 S3 (S1 is bad)

- everyone hears and replies to `write(value="A")`
- S1 and S2 reply to `write(value="B")`, but S3 misses it
- client can't wait for S3 since it may be the one faulty server
- S1 and S3 reply to `read(value)`, but S2 misses it; `read(value)` yields "A"
- result: client tricked into accepting out-of-date state

What happens here
with fail-stop failures?

Basic solution that works

$3f + 1$ replicas

Client waits for $2f+1$ matching replies

- f bad nodes plus a majority of the good nodes
- so all sets of $2f+1$ overlap in at least one good node

Remaining challenges

Message corruption, spoofing, ...

□ Use cryptography

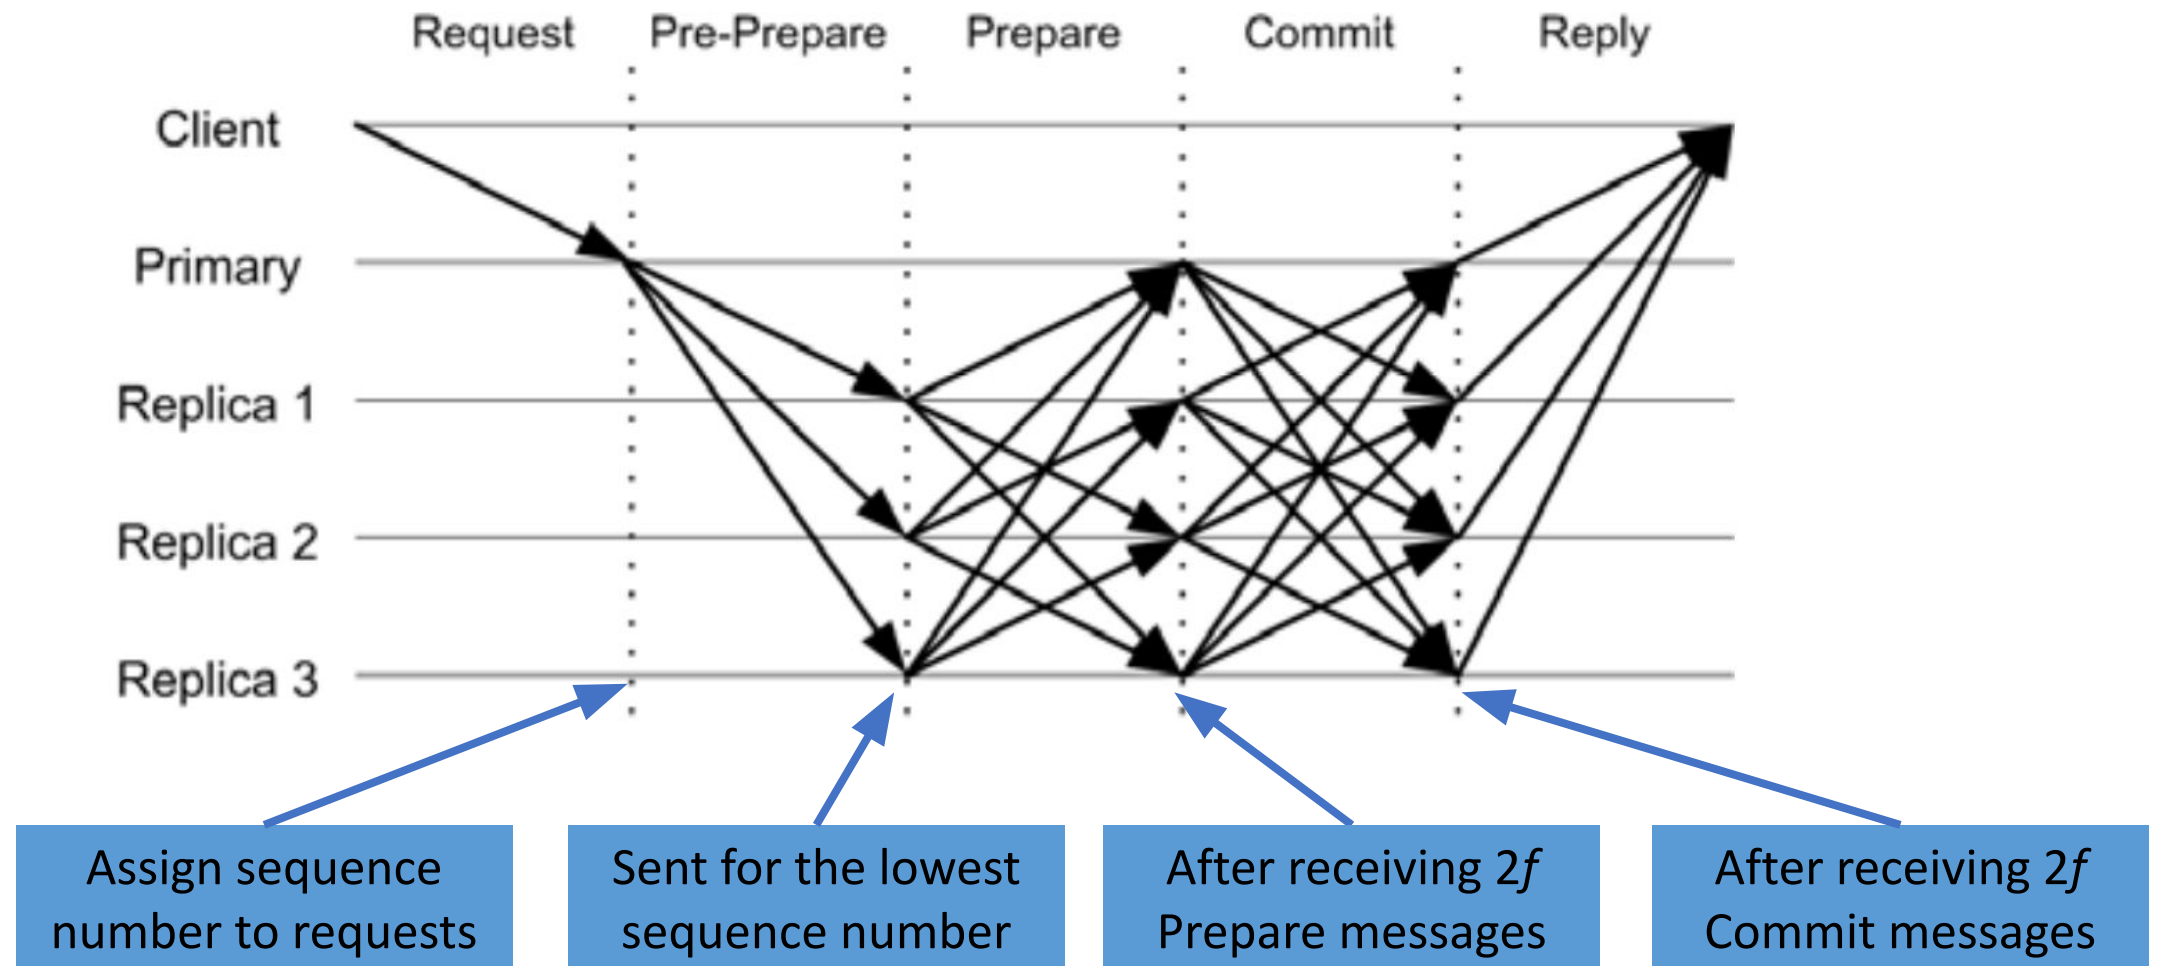
Concurrency

- Maintaining consistent state across all replicas requires that they all execute operations in the same order
- But different replicas will get messages in different orders, especially when you have multiple clients

□ Introduce a Primary

- But what if the primary is (Byzantine) faulty?

Practical Byzantine Fault Tolerance (Normal operation)



BFT principle

You cannot trust anyone but you can trust the group
(assuming that the group is big enough)

Blockchains inherit ideas from BFT solutions

Power of groups

Broadcast

Cryptography

Over to Dao and Edan