

THE UNIVERSITY *of* EDINBURGH

Distributed Systems Fall 2026

Yuvraj Patel

Disclaimer: Slides prepared using multiple sources (UW-Madison – Remzi, Andrea, Mike; Cambridge – Martin Kleepman; Distributed Systems book by MVS/ AST; University of Edinburgh – Yuvraj Patel)

Today's Agenda

Consensus

- Basics

Leader Election

Paxos

Why Consensus?

Multiple use-cases

- Replication – make sure the replicated data is same on all the nodes
- Failure Detection – a machine/leader has failed/stopped responding
- Leader Election – elect a leader to initiate a snapshot, etc.
- and many more…

All the above scenarios involve

- Multiple parties
- Presence of faults
- Coordinate amongst themselves
- Need to agree to something or arrive at a decision

Consensus Problem – Single value formulation

Consensus Protocol

Consider a distributed system with n nodes

- Each node i has an input x_i
- Faults may happen at arbitrary times

Output

- All nodes agree on a single value; Value cannot change later

Guarantee the following

- Termination (Liveness): Every non-faulty node eventually decides
- Agreement (Safety): All non-faulty nodes decide on the same value
- Validity: The decided value must be the input of at least one node

Consensus Protocol (contd…)

Not democratic; Value proposed by a small minority can be decided

Consensus possible depending on multiple parameters

Most important parameters

- System Model – Synchronous or Asynchronous
- Fault types – Crash or Byzantine

Synchronous vs. Asynchronous Systems

Synchronous systems

- Process execution speeds and message delivery times are bounded
- Can detect omission and timing failures

Asynchronous systems

- No assumptions about process execution speed or message delivery times
- Cannot reliably detect crash failures

Consensus

- Challenging in Asynchronous systems
- Solvable in Synchronous systems
- Algorithm for Asynchronous systems will work for Synchronous systems

Impossibility in Asynchronous Systems

Fischer, Lynch & Paterson show it is impossible to achieve consensus in asynchronous system with a single faulty process

They prove that no asynchronous algorithm for agreeing on a one-bit value can guarantee that it will terminate in the presence of crash faults

- With no crash too, algorithm may not terminate
- Proof constructs infinite non-terminating runs

One of the most fundamental results in distributed systems.

- Interested students can check the FLP paper -- https://dl.acm.org/doi/pdf/10.1145/3149.214121

Leader Election Problem

Need to elect leader to perform tasks and broadcast leader details

If leader fails

- Someone will detect leader failed
- Initiate a leader election to elect another leader
- Only one leader elected, and everyone agrees on who is the leader

System Model & Assumptions

System Model

- N nodes in the system; each node having unique id
- Communicate via messages; messages will eventually be delivered
- Failures/crashes may happen at arbitrary time

Assumptions

- Any node can call for an election
- Any node can call for at most one election at a time
- Multiple processes can call for an election simultaneously; still lead to a single leader
- Result independent of who calls for an election

Bully Algorithm

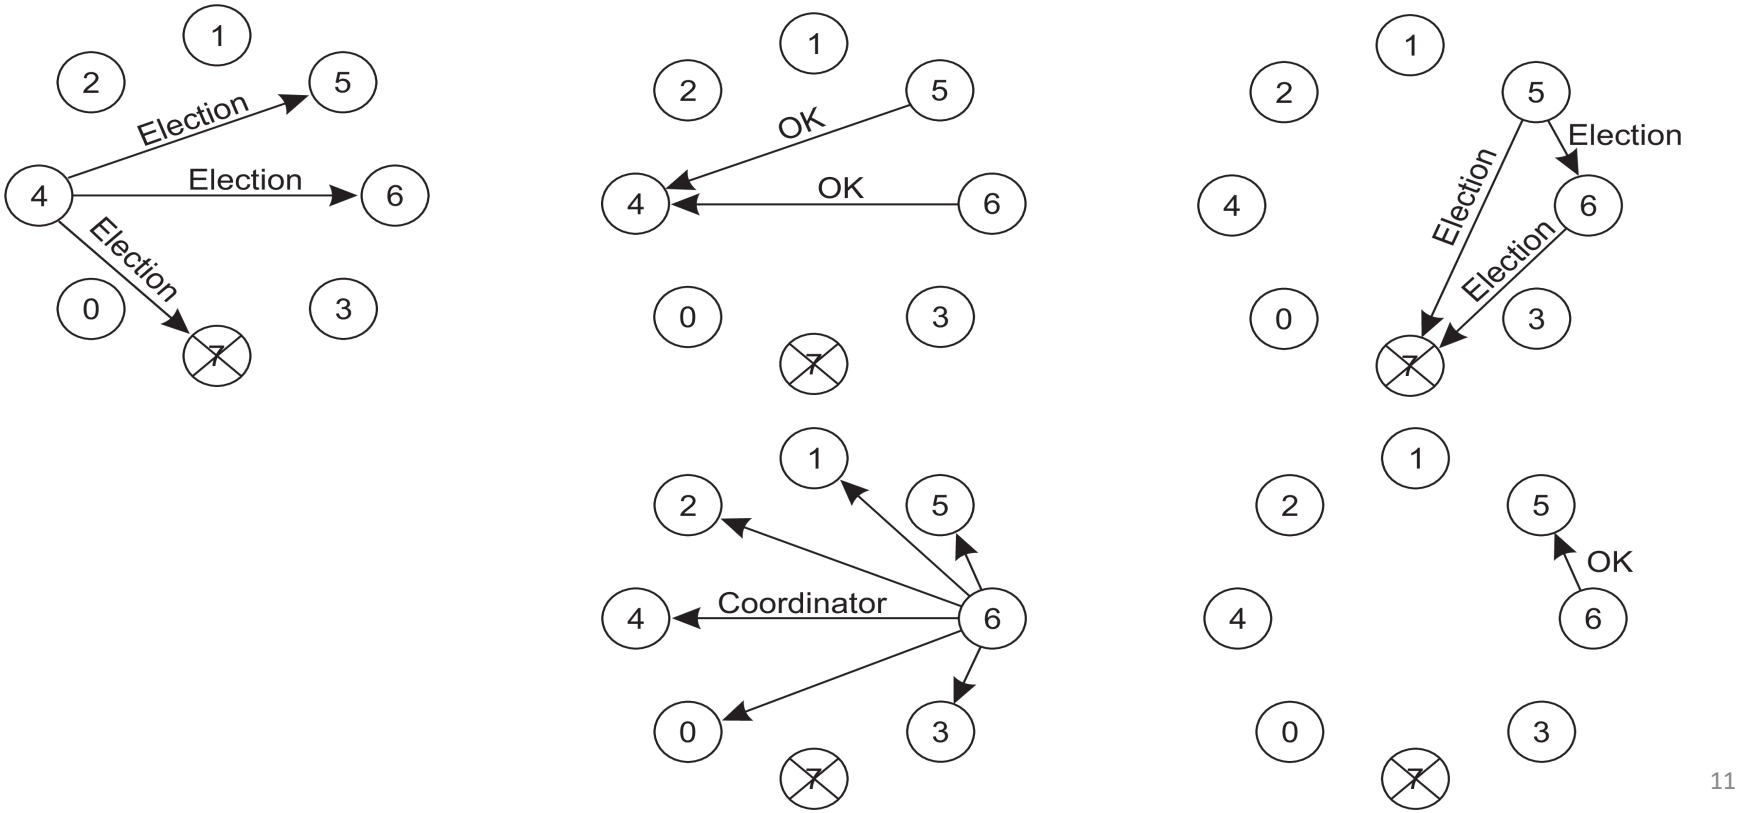
Key Idea: Node with highest ID wins

Consider N nodes {N_0, N_1, $N_{2...}$ N_n}.

Whenever a node N_k notices that the leader is unresponsive, election initiated

- N_k sends an ELECTION message to all the processes with higher IDs: N_{k+1},… N_n
- If no one responds, N_k wins
- If one of the higher-up's answers, it takes over and N_k's job is done

Example


1
2
5
Election
4
Election
6
Election
0
3
7
1
2
5
OK
4
OK
6
0
3
7
1
2
5
Election
4
Election
6
Election
0
3
7
1
2
5
Coordinator
4
6
0
3
7
1
2
5
OK
4
6
0
3
7

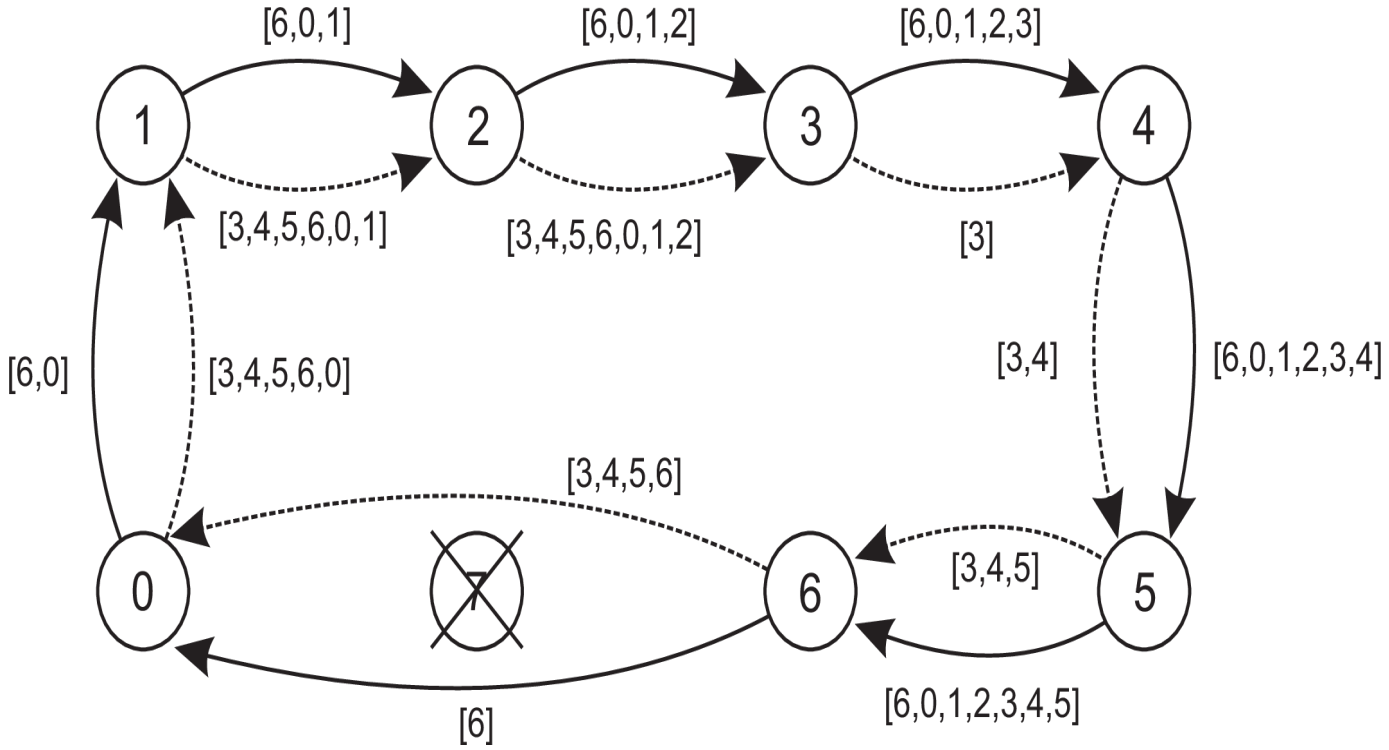
Ring Algorithm

Nodes are organized into a ring. Process with highest id is elected as coordinator

Whenever a node N_k notices that the leader is unresponsive, election initiated

- Any process can start an election by sending an election message to its successor. If a successor is down, the message is passed on the next successor
- If a message is passed on, the sender adds itself to the list.
- When the message gets back to the initiator, everyone had a chance to make its presence known.
- The initiator sends a coordinator message around the ring containing a list of all the living nodes. The one with the highest id is elected as coordinator

Example



The solid line shows the election messages initiated by N_6
The dashed one is election messages initiated by P_3

Both have the same list so it is safe to have two nodes initiating an election

How To Solve Consensus Then…

Paxos algorithm – Invented by Leslie Lamport

Most popular consensus solving algorithm

- Does not solve consensus problem (FLP still applies)

Used in many real-world systems – Yahoo, Google, etc.

Assume partially synchronous systems to avoid impossibility aspects

Paxos Algorithm

Role's node assume

- Proposers: Those who propose values
- Acceptors: Those who accept a proposed value
- Learners: Those who learn the proposed value after a consensus is reached
- One node can play two roles simultaneously

Other assumptions

- Nodes communicate with each other via messages
- Nodes operate independently and at different speed
- Nodes can crash or restart while operating
- Message receipt is asynchronous and can take longer time to be delivered, can be duplicated, and lost in the network. Messages are never corrupted

For majority, need 2m + 1 nodes to handle m failures

Proposal Numbers & Rounds

Each proposal has a unique number

- Higher numbers take priority over lower numbers (Older proposals rejected)
- Proposers always propose having a proposal number higher than it has seen/used

Simple Approach: Proposal number = Round Number + Node-ID

- Round Number – Higher than largest round number seen so far
- Need to remember largest round number so far
- Cannot reuse round number value after crash or reboots

Phases

Two phases – Prepare & Accept

Prepare Phase

- Find out any chosen values so far
- Block older and uncompleted proposals

Accept Phase

- Inform acceptors to accept a specific value

Analogous to how government passes laws

- Elect leader
- Propose a Bill
- Accept the Bill and turn in to a Law

Algorithm – Prepare Phase

Proposer

- Choose proposal number n, send <prepare, n> to acceptors

Acceptor

- Only receiving a prepare message
  - If $n > n_h$, where n_h is the highest proposal seen so far by the acceptor

    $n_h = n$. (Promise to not accept older proposals)

    If no prior proposal accepted,

    reply <promise, n, NULL>

    Else

    reply <promise, n, (n_a, v_a)>
  - Else

    Reply <prepare-failed>

Algorithm – Accept Phase

Proposer

- If receive promise from majority of the acceptors,
  Determine any earlier chosen value v_a for n_a and choose latest value or any value v selected by the proposer
  send <accept, n, v> to acceptors

Acceptors

- If n >= n_h
  $n_a = n_h = n$
  $v_a = v$
- reply <accept, n_h>

Proposer

- When responses received from the majority
  If any $n_h > n$
  Start from prepare phase again
  Else
  Value is chosen

Example – Everything works fine

Example – Acceptor failure

Accept Phase Failure

Prepare Phase Failure

Example – Proposed failure

Prepare Phase Failure

Example – Proposed failure

Accept Phase Failure

Failure Handling Summary

One proposer

- One or more acceptors fail
  - Still works as long as majority nodes are up
- Proposer fails in prepare phase
  - No-op; another proposer can make progress
- Proposer fails in accept phase
  - Another proposer overwrites or finishes the job of failed proposer

Two or more simultaneous proposers

- More complex
- Can lead to livelock (fix with leader election)

Paxos Algorithm – Safety & Liveness

Safety

- Only a single value is chosen
- Only chosen values are learned by nodes
- Only a proposed value can be chosen

Liveness

- Some proposed value eventually chosen if fewer than half of processes fail
- If value is chosen, a process eventually learns it

Paxos is safe but often live

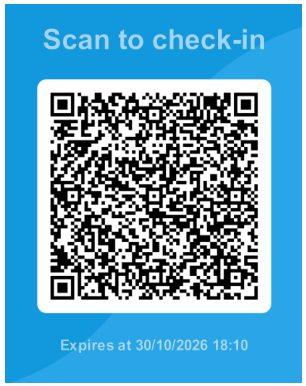
Multi Paxos

Basic Paxos comprises two rounds

For real-world systems like databases, every single operation needs to go through Basic Paxos rounds, which is costly

Multi Paxos – Creating a log of agreements

- Assume Proposer is stable
- Use Phase 1 for the Proposer election
- Use Phase 2 multiple times and work on multiple values being accepted



Coursework Survey