CS 186 - Fall 2024
Exam Prep Section 12
Consensus

1 Paxos Timing

Suppose there are five nodes P1, P2, P3, P4, and P5. We have the following information about the network:

- Nodes can broadcast messages to other nodes simultaneously.
- Each node takes 20 seconds to send a message to another node.
- Each node takes 10 seconds to flush a log record.

Suppose P1 multicasts a PREPARE(1) at t=0 seconds to P3, P4, P5. Suppose P2 multicasts a PREPARE(2) at t=15 seconds to P3, P4, P5. Answer the following questions about the Paxos execution.

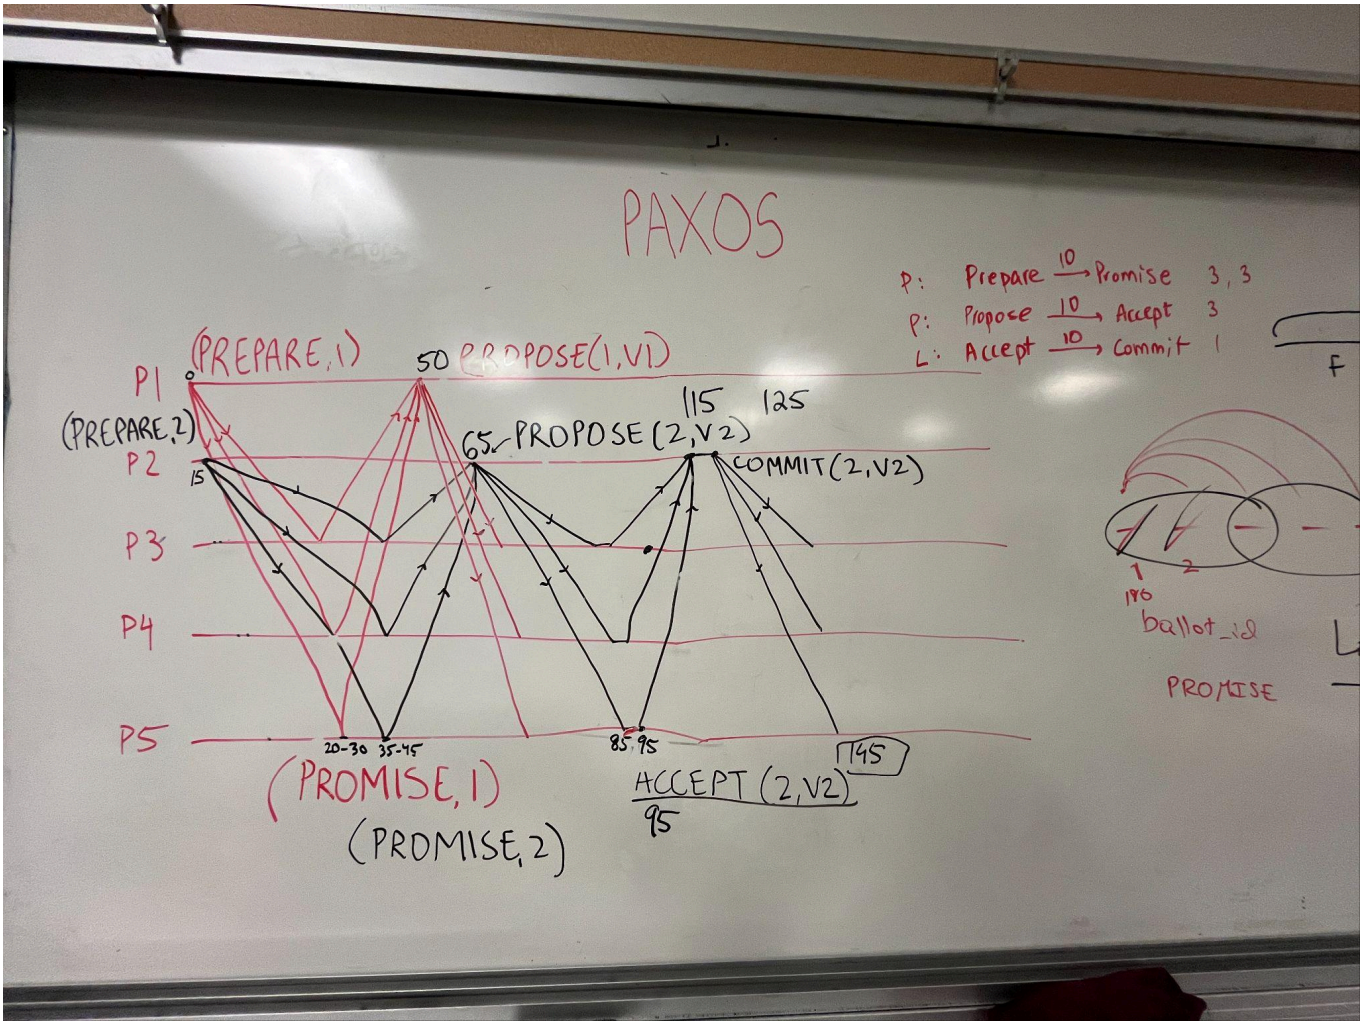

1. How many PROMISE messages are sent?

Answer:

At t=20, P3, P4, P5 receive PREPARE(1) and flush ballot 1 to disk. At t=30, they each send PROMISE(1) to P1. At t=35, P3, P4, P5 receive PREPARE(2) and flush ballot 2 to disk. At t=45, P3, P4, P5 send PROMISE(2) to P2. At t=50, P1 receives the PROMISE messages and multicasts PROPOSE(1, v1). At t=65, P2 receives the PROMISE messages and multicasts PROPOSE(2, v2). At t=70, P3, P4, P5 receive PROPOSE(1, v1) but ignore the proposal. At t=85, P3, P4, P5 receive PROPOSE(2, v2). They each flush ballot 2 to their log by t=95 and send ACCEPT(2, v2) to P2. At t=115, P2 receives the ACCEPT messages and determines that v2 was chosen. At=125 it finishes flushing ballot 2 to its disk and broadcasts a COMMIT(2, v2) message. At t=145, all nodes receive COMMIT(2, v2) and the execution finishes.

6 PROMISE messages are sent

2. How many PROPOSE messages are sent?

Answer:

2 PROPOSE messages are sent

3. How many ACCEPT messages are sent?

Answer:

3 ACCEPT messages are sent

4. How many COMMIT messages are sent?

Answer:

3 COMMIT messages are sent

5. How many log flushes occur?

Answer:

10 log flushes (1 for each PROMISE, 1 for each ACCEPT, 1 before P2 broadcasts COMMIT)

6. How long does the Paxos execution take?

Answer:

145 seconds

7. Suppose P3 multicasts a PREPARE(3) to a majority of participants at t=100 seconds. What happens?

Answer:

The nodes in the majority will first send a PROMISE(3, (2, v2)) message for ballot 3, after flushing it to disk. This PROMISE message contains ballot 2 since each node will include the highest

numbered ballot it has accepted. This forces P3 to send PROPOSE(3, v2) instead of PROPOSE(3, v3), where v3 is any arbitrarily picked value.