CSE 452

Distributed Systems

Paxos

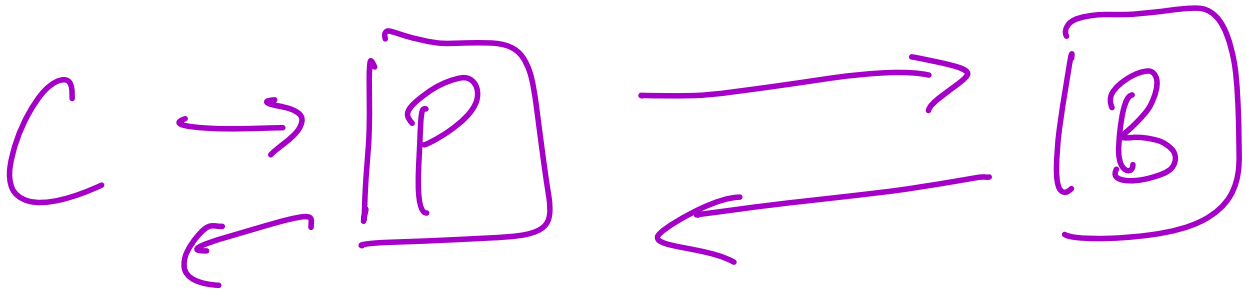

Log based SMR

slot = index into a log

MultiPaxos
- log
- decides values for each slot

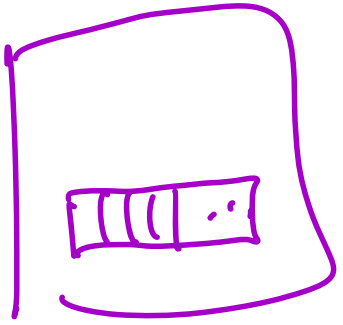

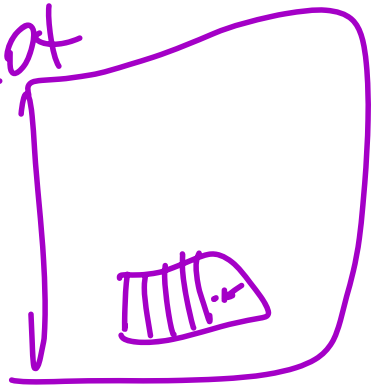

Single Decree Paxos

- Consensus problem
  - agree on a value — choose
  - nodes propose values
- requirements
  - chosen value was proposed
  - at most 1 chosen value
  - nobody thinks a value chosen unless it was

Roles (≥ nodes, but not exactly)

- proposer proposes values
- acceptor vote on proposals
- learner learns results

An execution has a fixed number of each role

Ballots (Rounds)

- ballot number (mon. increasing)
- proposed value
- allocate ballot numbers to proposer
  - even/odd numbers for two proposer

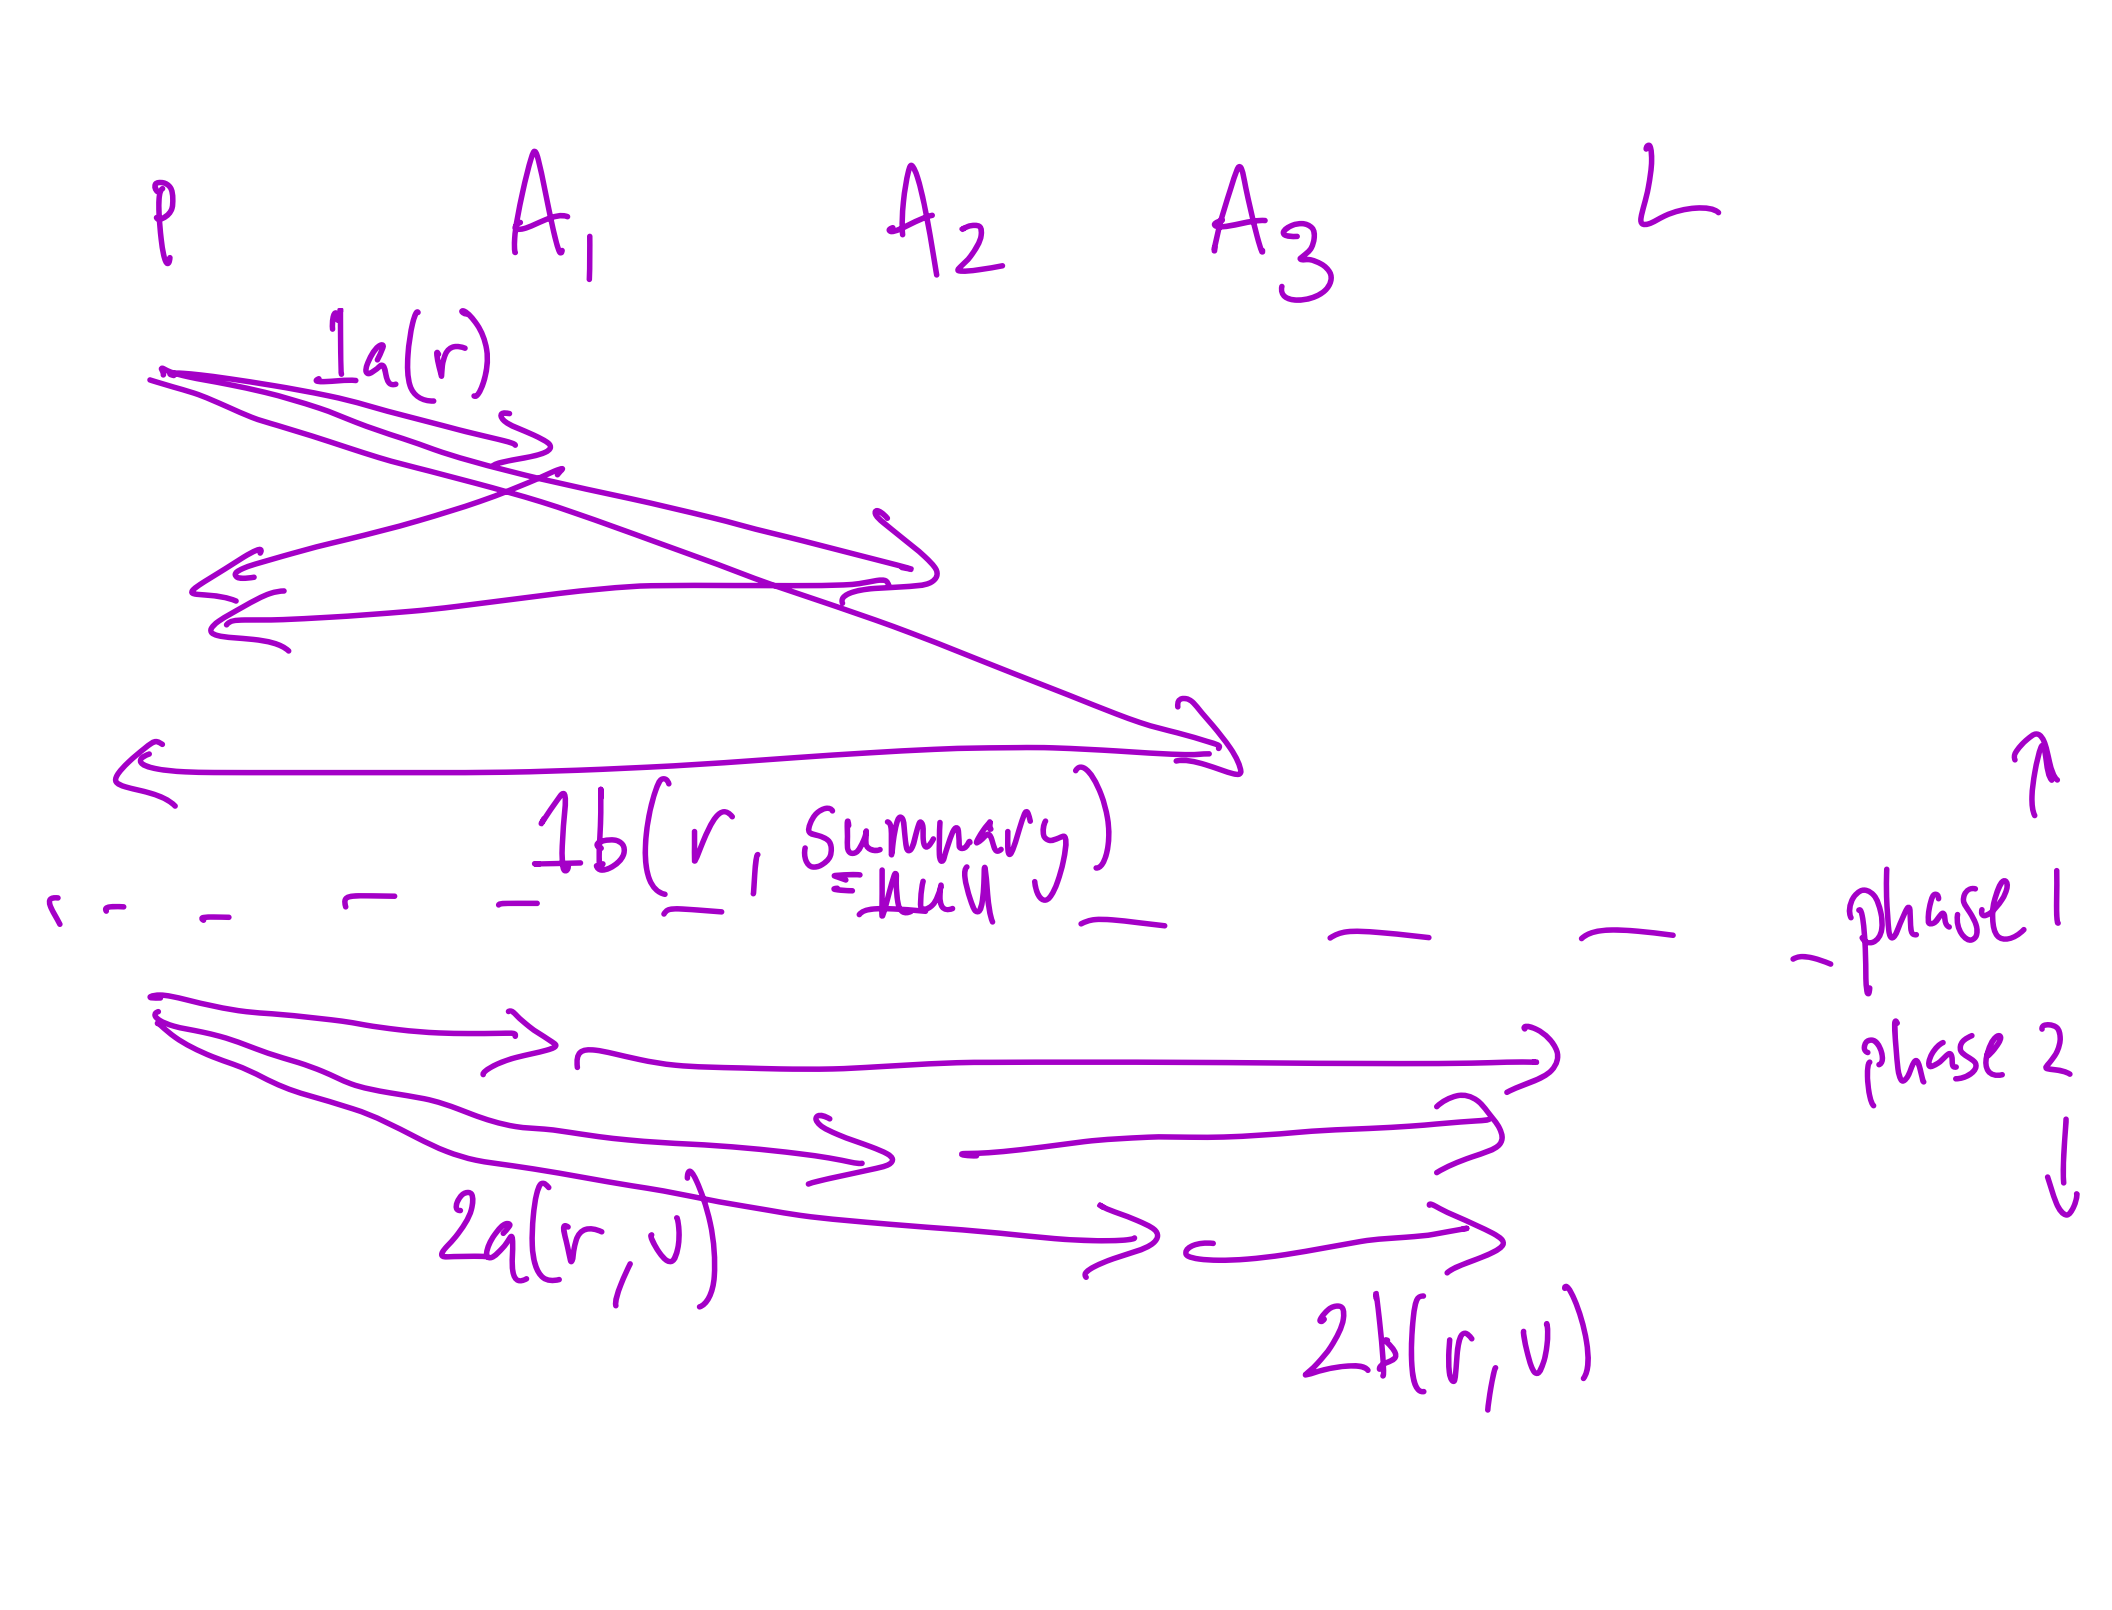
P
A1
A2
A3
L
1a(r)
1b(r, summary)
phase 1
phase 2
2a(r,v)
2b(r,v)

Phase 1

- proposer picks a round # r
- send 1a(r) to all acceptors
- acceptors respond w/ 1b(r, summary)
  - summary:
    - "I have never voted in round < r"
    - TODO: if you vote^

Phase 2

- proposer waits for a majority of 1b
- look at all 1b msgs
  - if all have Summary = null
    then propose any value
    - TODO: non null Summary
- sent 2a(r, v) to all acceptors
- acceptors send 2b(r, v) to learner if they can

Learner

- waits for a majority of 2b msgs
  (all in the same round)

Chosen(r, v) =

there is a majority of acceptors

who have sent 2b(r, v) msgs

chosen(v) = there is a round r

such that chosen(r, v)