Paxos Made Simple John Nguyen

Slides adapted from Leslie Lamport and Thomas Marshall

Problem

How to agree on a value in a distributed system that tolerate non-malicious failures?



Problem

Distributed consensus problem

- Group of processes must agree on a single value
- Value must be proposed
- After value is agreed upon, it can be learned

Requirements



Liveness

Properties: Safety

Only a proposed value is chosen
Only a single value is chosen.
Only chosen values are learned by processes

Properties: Liveness

Some proposed value is eventually chosen if fewer than half of processes fail

If a value has been chosen, a process can eventually learn the value.

Paxos's notation

Classes of agents: Proposers Acceptors Learners A process can act as more than one clients (usually 3). Assumption: asynchronous, non-byzantine model

Paxos Phase 1 (prepare)

Proposer:

Each proposal should be of form <n, v> where n is strictly increasing

> A proposer sends a prepare request with number n to majority of acceptors.

Acceptors:

- If n > n*
 - > N* = n \leftarrow promise not to accept any new proposals n' < n
 - If no prior proposal accepted
 - Reply <promise, n, Ø>

Else

- Reply <promise, n, (na , va) >
- Else
 - Reply reject

Phase 1 Example

Proposal is <5, 100> Acceptor Already accepted proposals For example, if it has accepted <1, 20>, <2,</p> 30> and <4, 40> it will respond with <4, 40> Did not accepted any proposal <5, 100>

Paxos Phase 2 (accept)

If the proposer receives a response YES to its prepare requests from a majority of acceptors, then it sends an *accept request* to each of those acceptors for a proposal numbered n with a value v which is the value of the highest-numbered proposal among the responses.

If an acceptor receives an accept request for a proposal numbered n, it accepts the proposal unless it has already responded to a prepare request having a number greater than n.

Paxos algorithm

Phase 3 (learn):

 Learners need to know which value has been chosen

Elect a set of "distinguished learners"
 Acceptors respond with to learn requests with their acceptance

These distinguished learners informs other learners

Definition of chosen

A value is chosen at proposal number n iff majority of acceptor accept that value in phase 2 of the proposal number.



Proposers can continually propose higher and higher proposal numbers without any ever being accepted

Distinguished Proposer: the only one trying to initiate proposals

Paxos's properties

▶ P1: Any proposal number is unique.

P2: Any two set of acceptors have at least one acceptor in common.

P3: the value sent out in phase 2 is the value of the highest-numbered proposal of all the responses in phase 1.









