Paxos made Practical

September 21, 2011

CS294-42 Presentation

- Gautam Kumar

Problem

- Designing a Replicated (Deterministic) State Machine.
- Requires consensus regarding the ordering of commands.
- Availability of the system
  - Paxos algorithm achieves distributed consensus
  - Can survive f out of 2f+1 failures
- The idea seems simple, the implementation does not in the scenario where machines come & go.
- This paper tries to answer some of these questions.

Models

- Original Paxos model :
  - Achieve consensus amongst the ordering of the commands as well aside from ensuring a majority of machines know about the execution of the command.

- Liskov model (conceptualized in 1988, before Lamport) :
  - No need to achieve consensus on ordering, designate a primary to do that for you, just ensure the second property.

Groups and cohorts

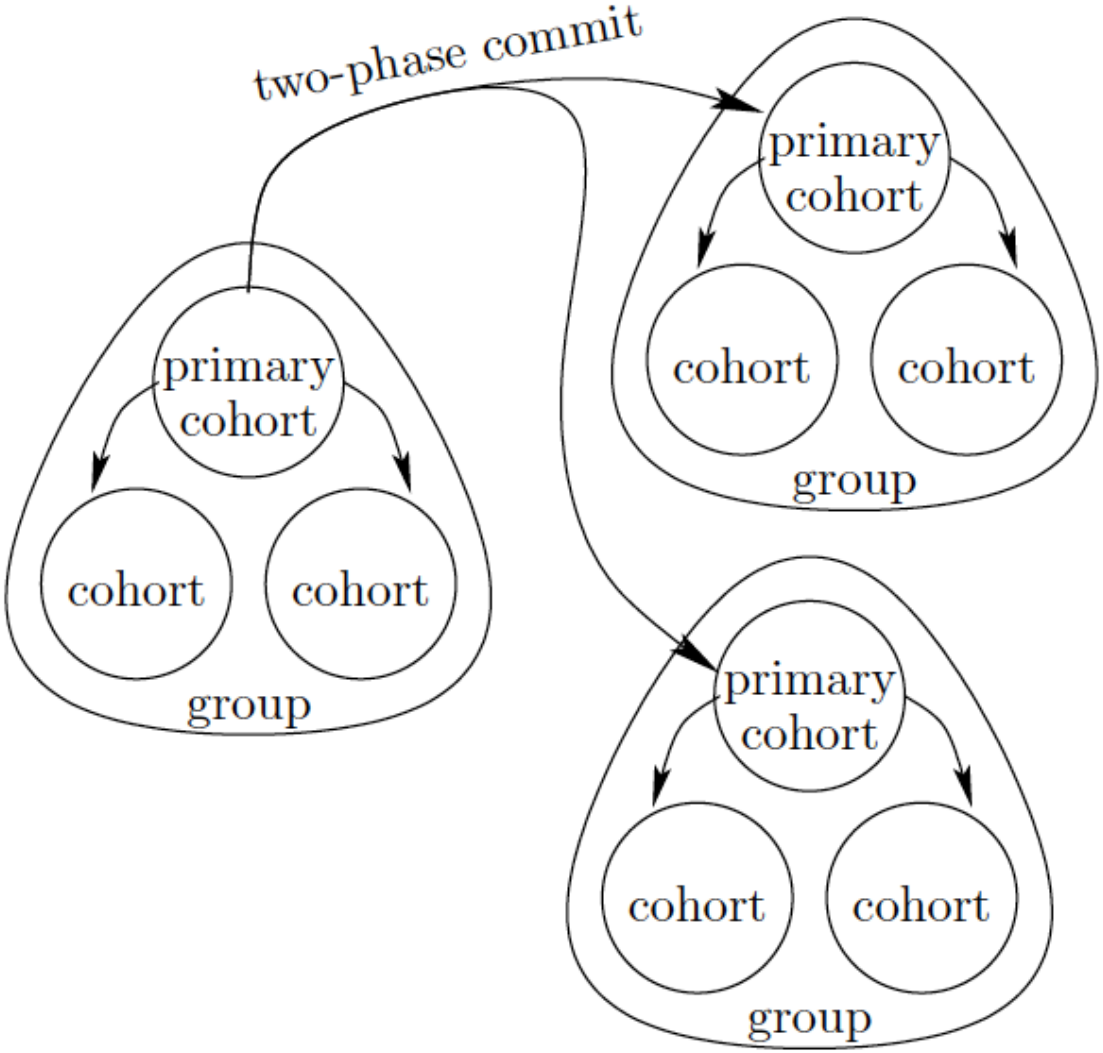

Designated Primary

- Communicates directly with the client.
- Imposes a total order on the commands.
- Removes non-determinism from the commands.
- Ensures that a majority of backup cohorts log the command.
- Executes it and sends the reply back.

Single Point of Failure?

- The primary, if it fails results in some disruption.
- But this is just temporary, since it's easy to get the system back up and running.
- Not an issue with the Lamport's model.
- This calls for a need for automatic reorganization of views (defined as a set of cohorts with a designated primary) depending on failures or cohort preferences.

Achieving a new view

- Interestingly this itself turns out to be another instance of Paxos itself.
- A view Manager first proposes a new *view_id.*
- If majority of cohorts agree, the view manager instantiates it.
- Designates a new primary.
- Good to go!

Problems with this model

- Essentially only one server executes the request.
- The others are just ‘capable’ of executing them.
- What if one is not sufficient to meet the demand?
- The previous model could get around since each cohort can be responsible for a subset of requests.
- Overheads?

Thanks!

- Gautam