State Machine Replication and Consensus Algorithms - Part II

Hui Chen [a]

[a]CUNY Brooklyn College

October 22, 2025

Raft Consensus Algorithm: Overview

- Quorum-based consensus algorithm
- Three components
  - Leader election
  - Log replication
  - Safety
- Leader must be supported by a majority of the group (a quorum)
  - Majority quorum: n servers tolerate f failures, where $n = 2f + 1$
- Availability-consistency trade-off:
  - Consistency. At most one connected subgroup can serve requests
  - Availability. Once a majority of replicas fail, the remaining replicas should not serve requests.

Terms and Elections in Raft

Raft Consensus Algorithm (Ongaro and Ousterhout 2014)

- Time divided into terms
- Each term begins with an election
- Three server states: Leader, Follower, Candidate
- Election process
  - Followers start election if no communication from leader within election timeout
  - Candidate requests votes from other servers
  - If candidate receives votes from majority, becomes leader
  - If split vote, starts new election

Log Replication in Raft

- Log structure
- Log replication process

Log Structure in Raft

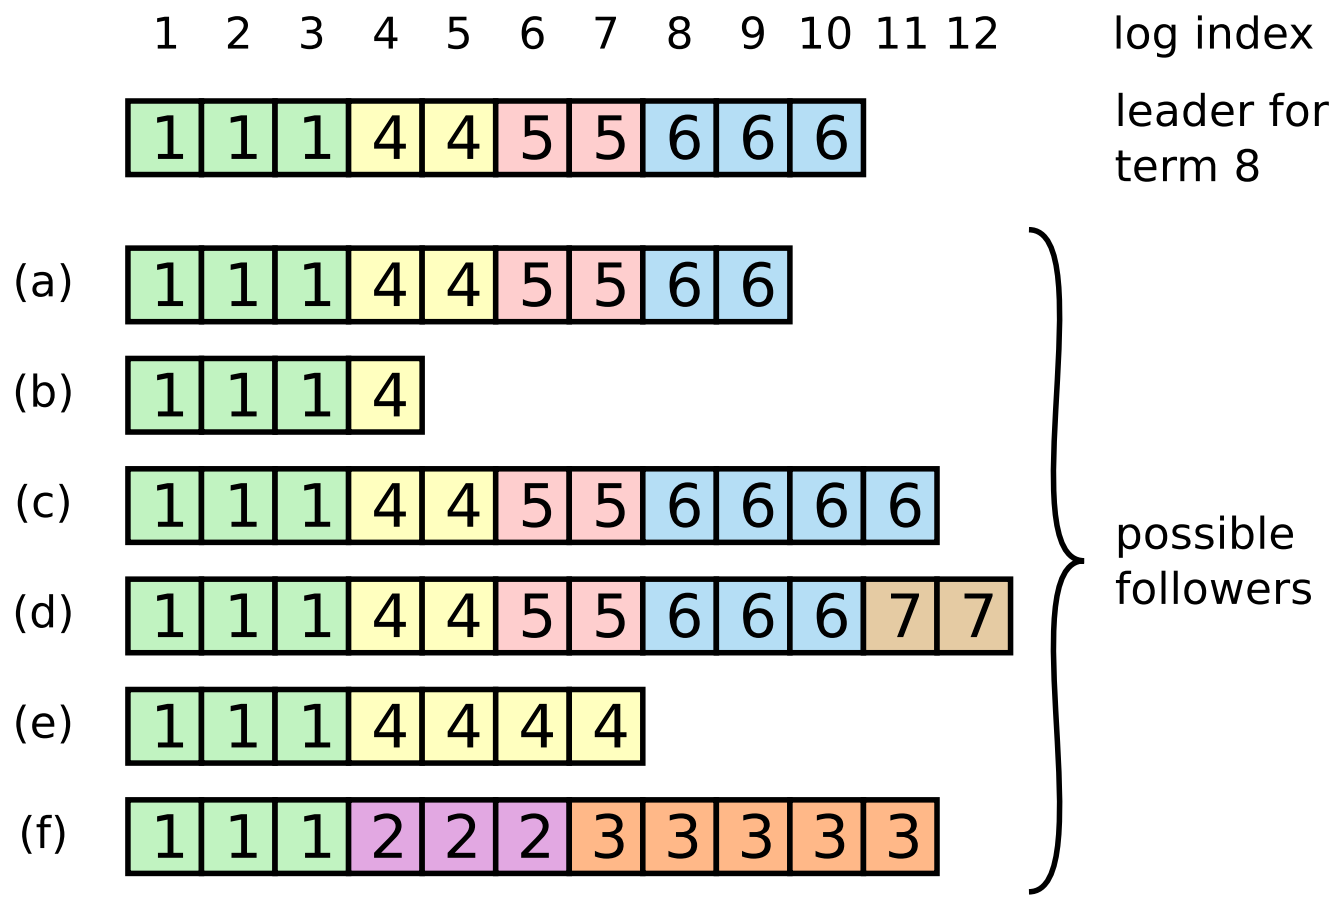

Bibliography I


Ongaro, Diego and John Ousterhout (2014). "In search of an understandable consensus algorithm". In: *2014 USENIX annual technical conference (USENIX ATC 14)*, pp. 305–319.