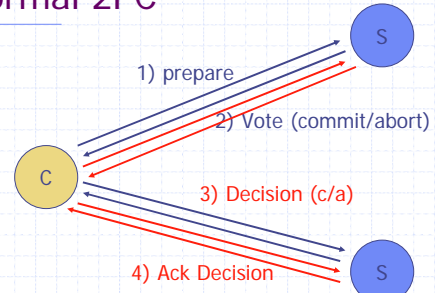


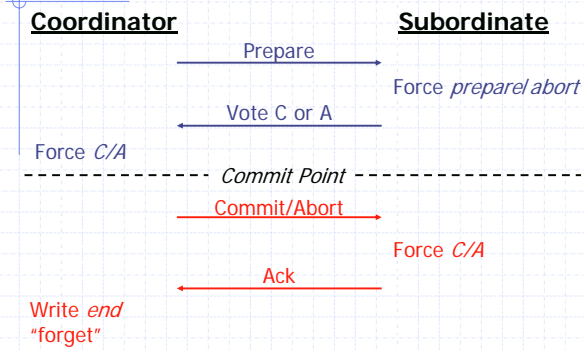
# Commit Protocols

February 26, 2004

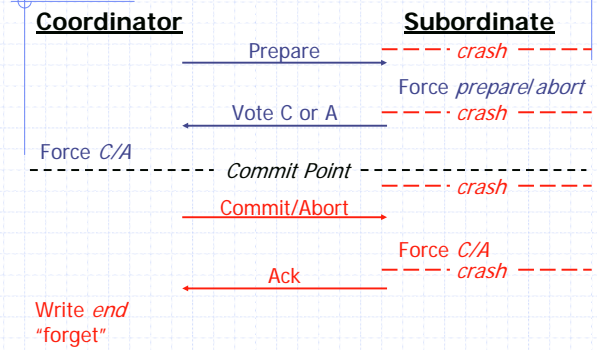
## Normal 2PC



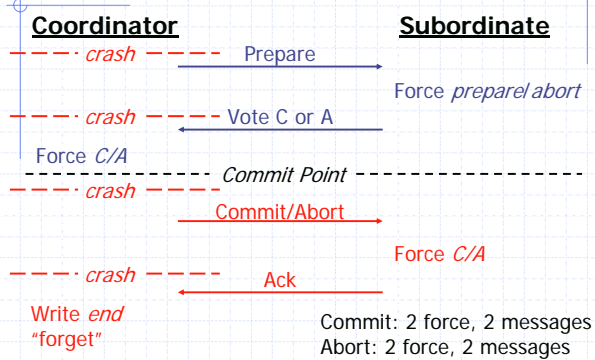
## Timeline: 2PC



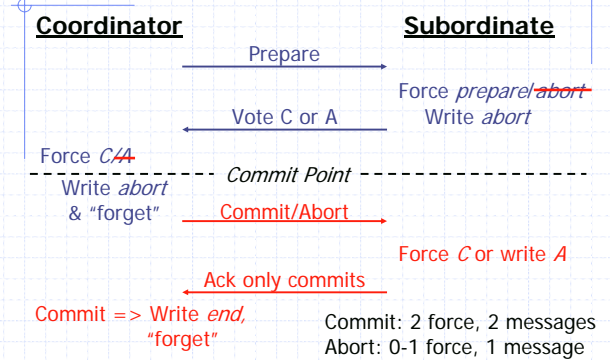
## Subordinate Crashes



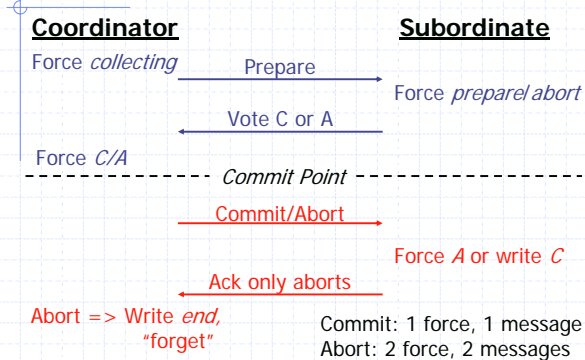
## Coordinator Crashes



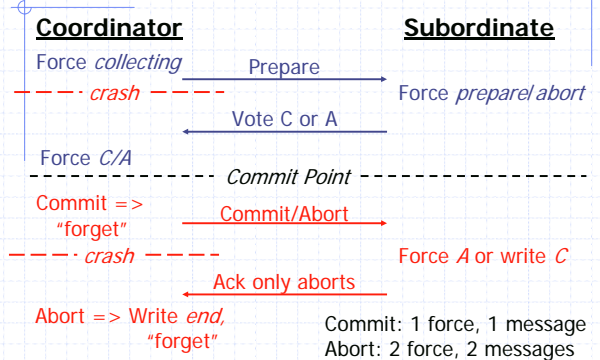
## Presumed Abort



## Presumed Commit



## Presumed Commit



## Ninja 2PC Optimizations

- ◆ Focus on availability
- ◆ Sub dies in phase 2:
  - Complete without the ack
  - Recovery will get the new version from another replica
- ◆ Coordinator dies:
  - Replicas contact each other to determine the result
  - Faster than waiting for coordinator to recover

## Ninja Apply function

- ◆ Use 2PC to apply delta to all replicas
  - Phase 1: prepare
  - Phase 2: commit
  - Ensure serializable updates
- ◆ Livelock
  - Contention may cause aborts
  - Replicas that are out of order may abort
  - No guarantee of progress
  - Real problem under high update rates

## Weak Apply

- ◆ Phase 1: Commit (!)
  - Order may vary, but fine for some uses
  - Retry commit until it succeeds on all replicas
- ◆ Phase 2: ack => stop retrying
  - Much less livelock!