
Paxos Made Simple

Lamport

Thomas Marshall

Motivation

- We need a way to maintain consistency in a distributed system in the presence of failures.
 - 2PC works, but can get “stuck”, so a consensus algorithm is better.
-

Background

- Jim Gray proposes 2PC in the 1970s, but it blocks on single node failures.
 - Dale Skeen proposes 3PC in the 1980s, but it produces incorrect results in some situations.
 - Leslie Lamport proposes Paxos in 1998; the original paper describes the ancient Greek civilization on the Paxos island.
-

Example

- Leader election (eg. Mesos and Zookeeper) – important to only have one leader at a time.
 - Some node(s) propose to be leader, other nodes can accept or reject.
-

Paxos

- We want to choose a value and have every node in the cluster agree on the value.
 - Three classes of agents: proposers, acceptors, learners.
 - Failures are possible, but non-Byzantine.
-

Safety Properties

- Only a value that has been proposed may be accepted.
 - Only a single value is chosen.
 - An agent never learns that a value has been chosen unless it actually has been.
-


Algorithm

- A proposer selects a proposal number n and sends a request to the acceptors.
 - If an acceptor has not already accepted a proposal with number greater than n , it responds that it can accept this proposal.
-


Algorithm (cont)

- If a proposer receives ready responses from a majority of acceptors, it sends an accept message.
 - An acceptor that receives an accept message accepts the proposal unless it has responded to a prepare with a number greater than n .
-

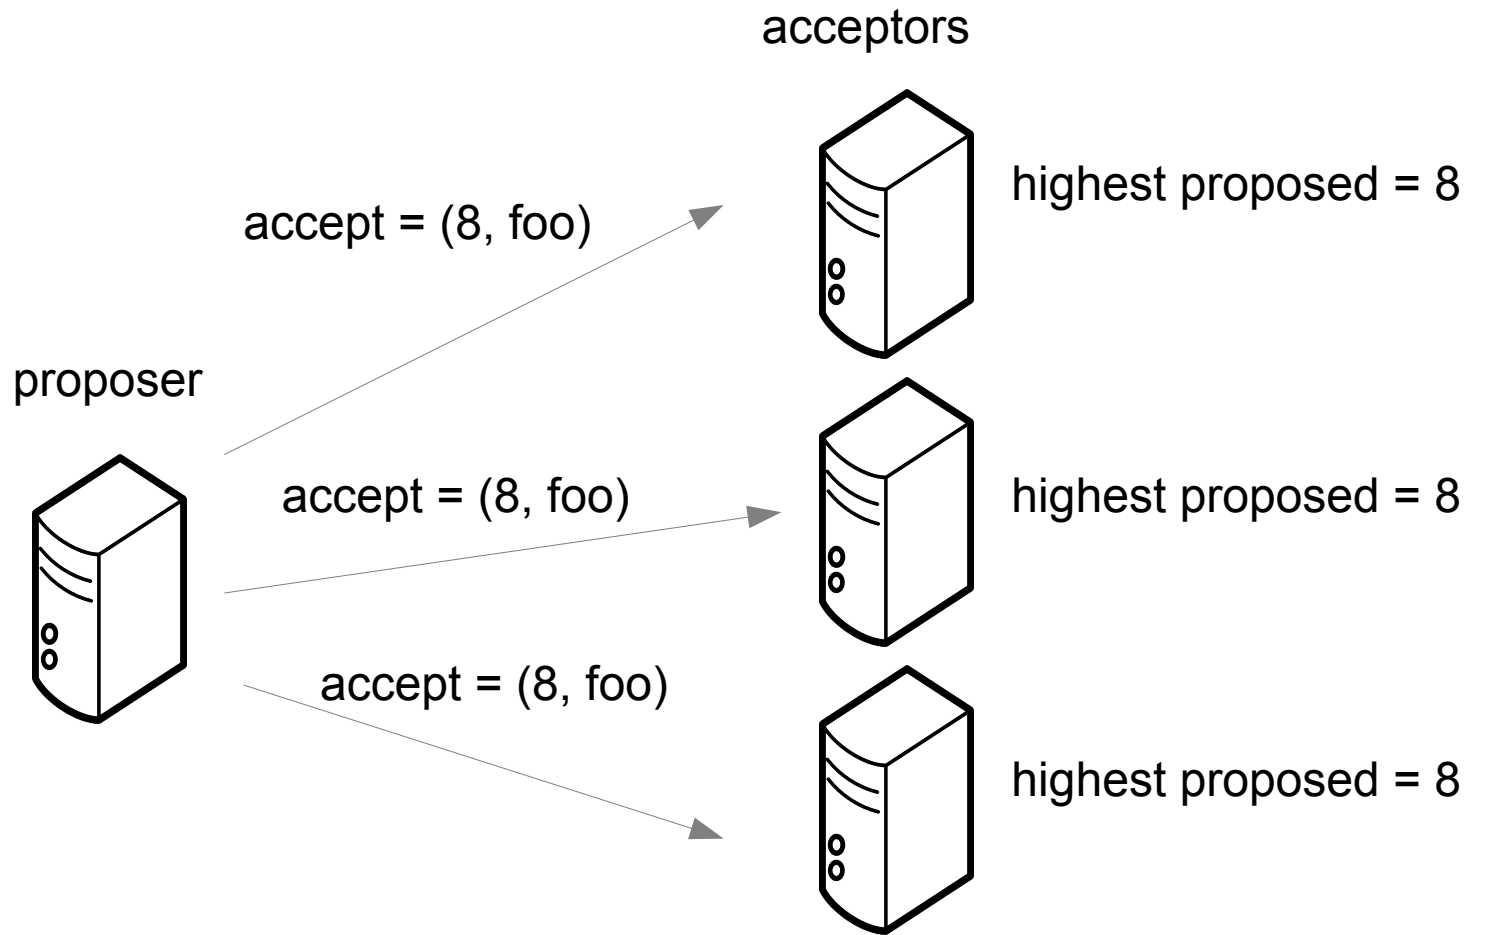
Paxos Example



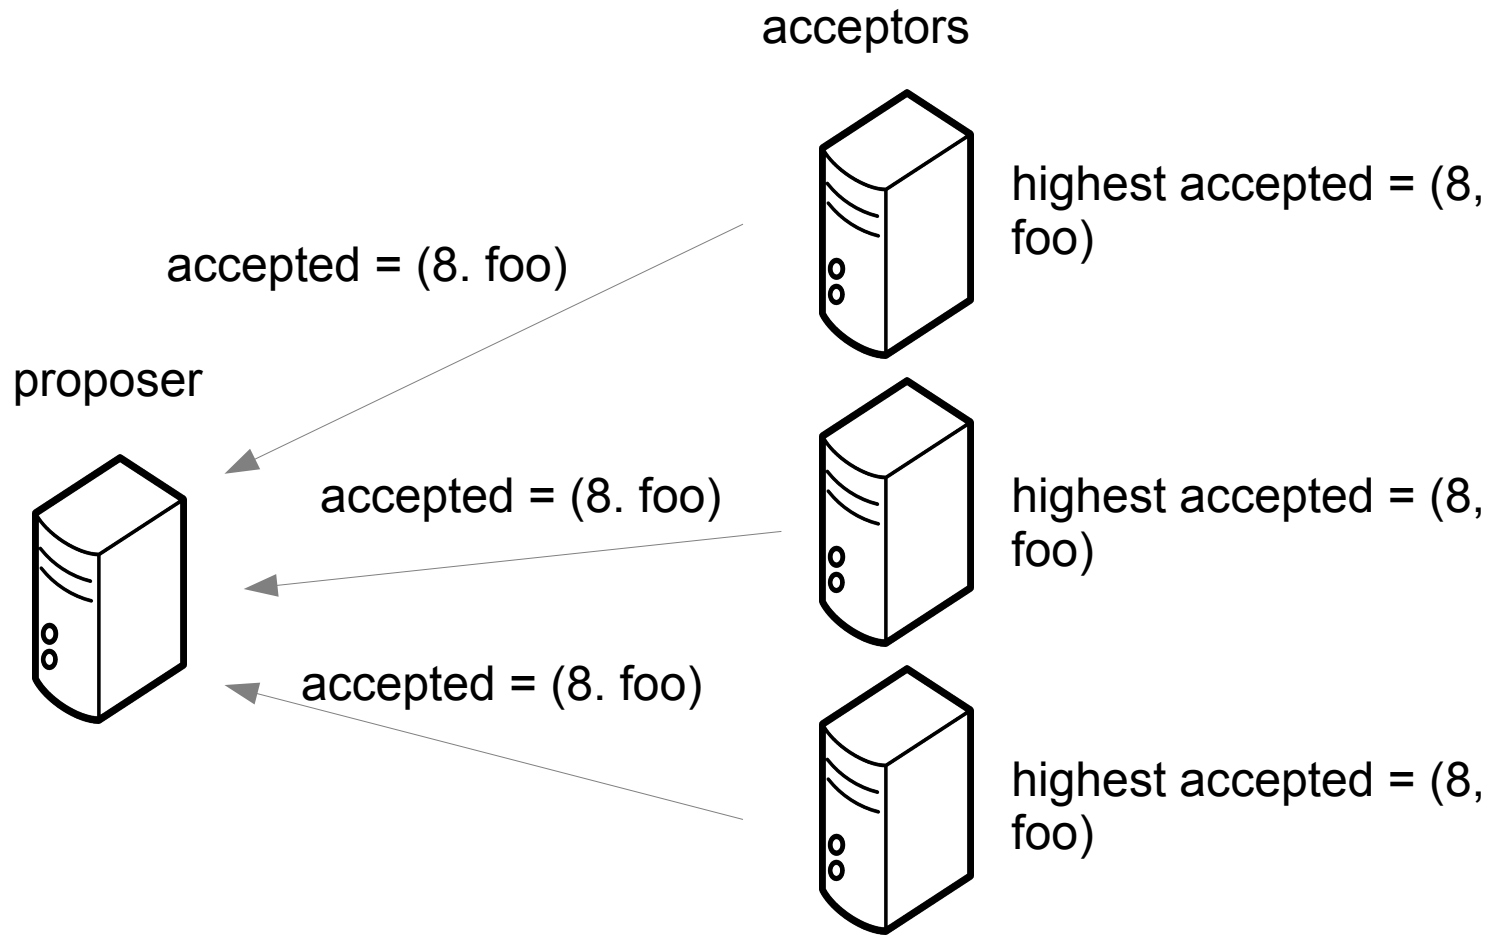
Paxos Example



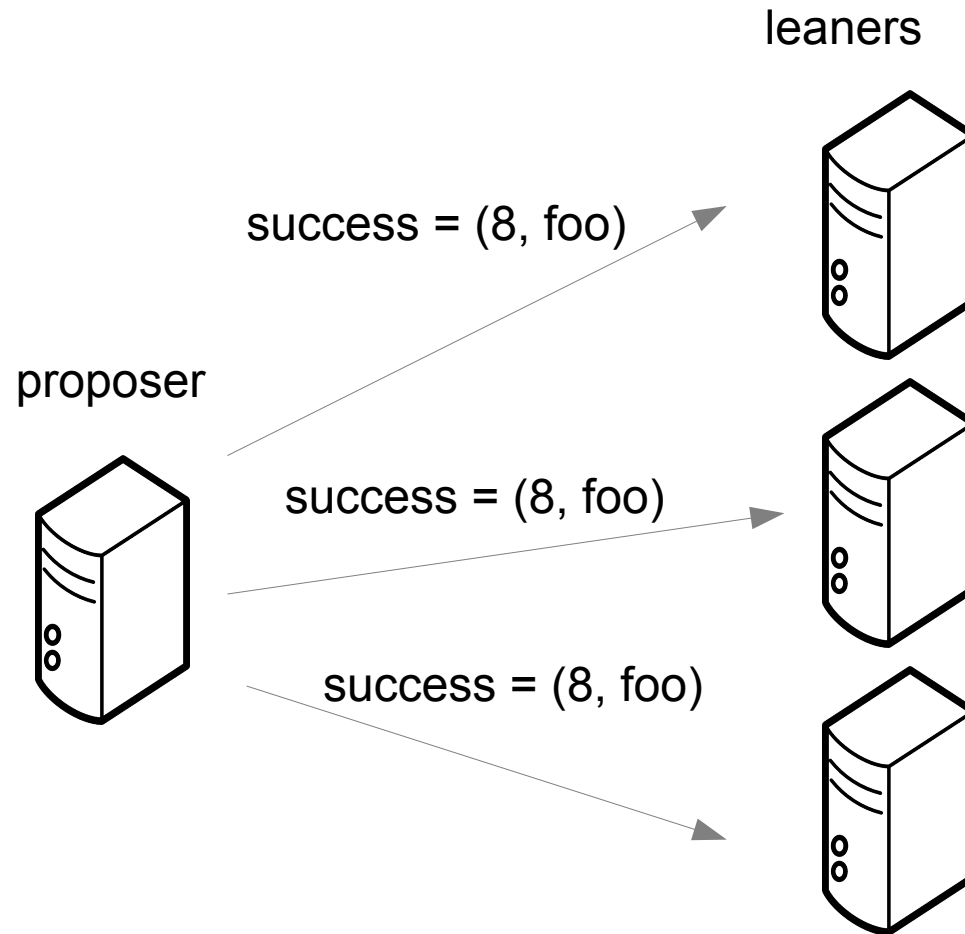
Paxos Example



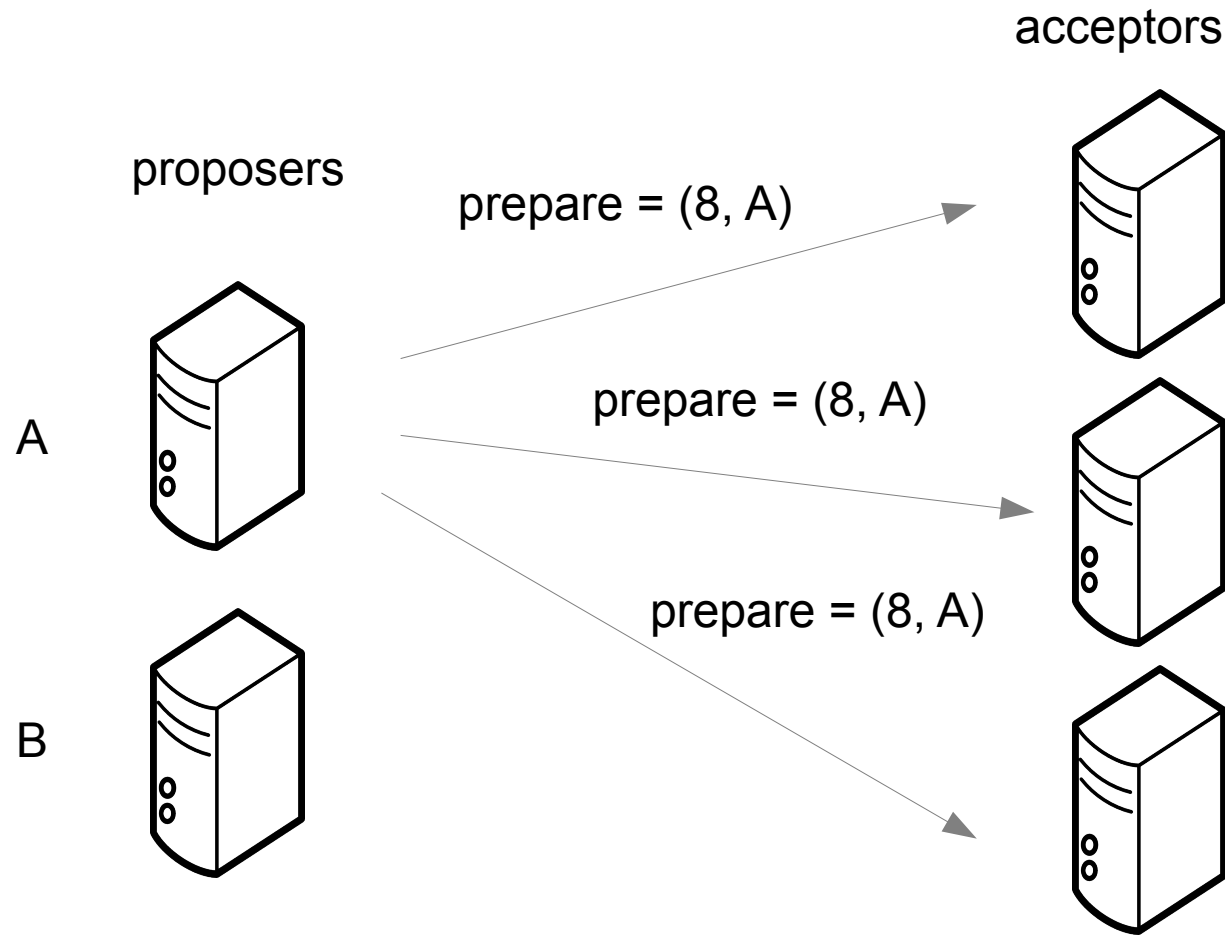
Paxos Example




Paxos Example



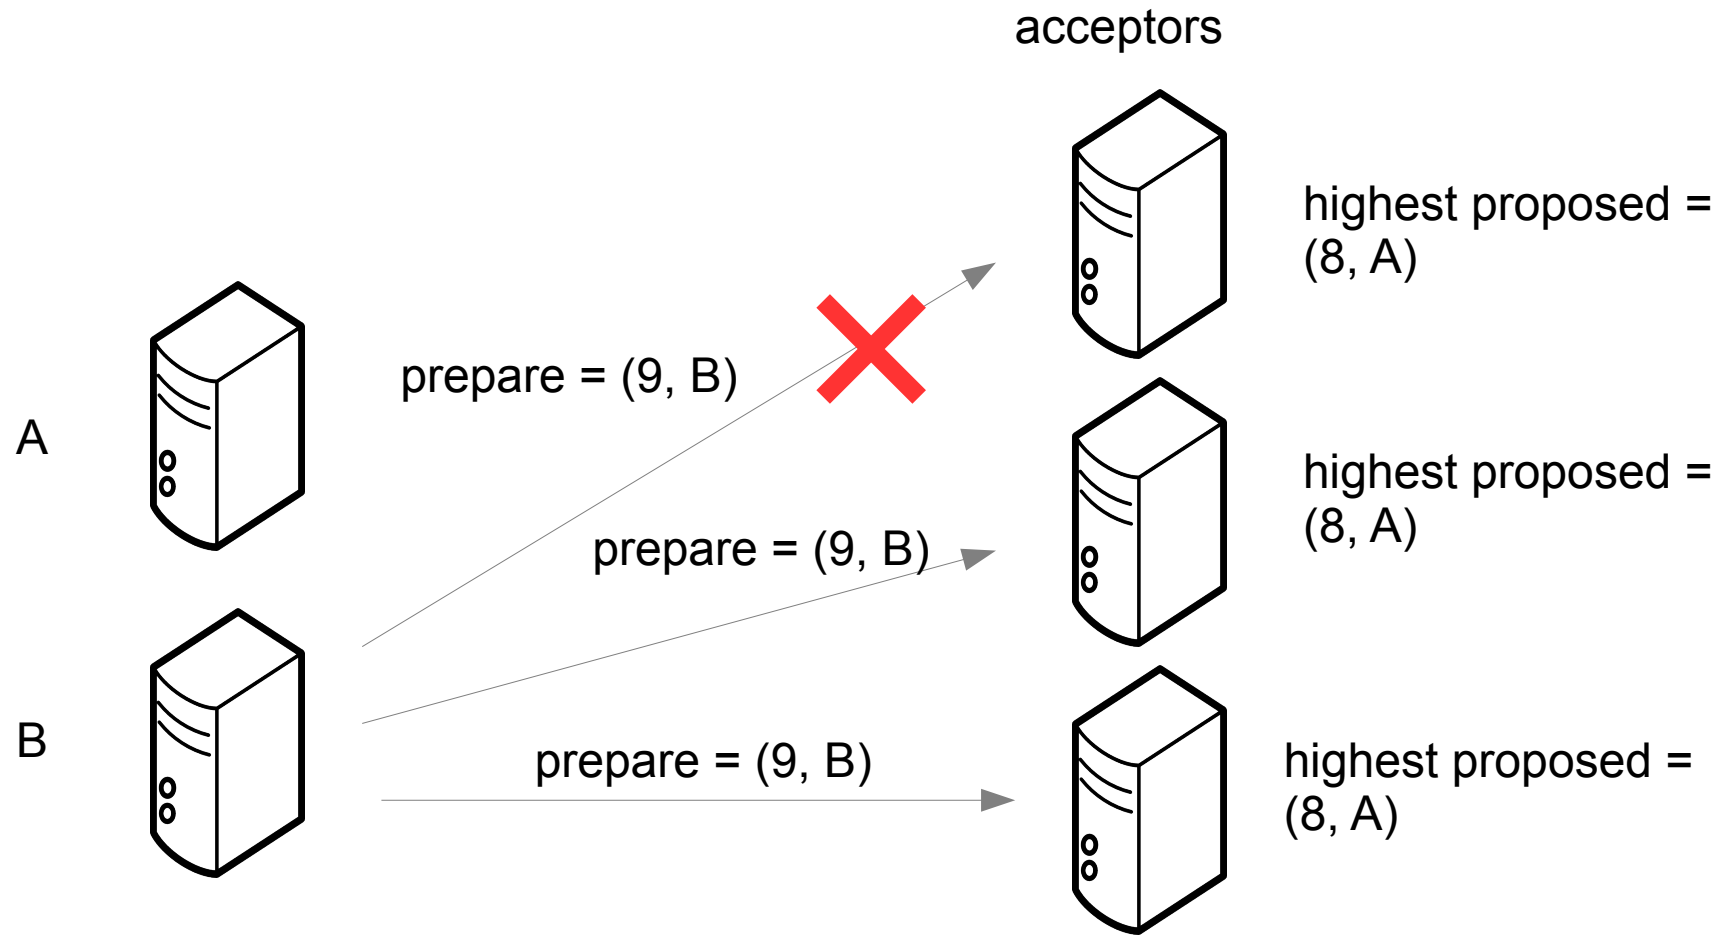
Paxos Example (2)



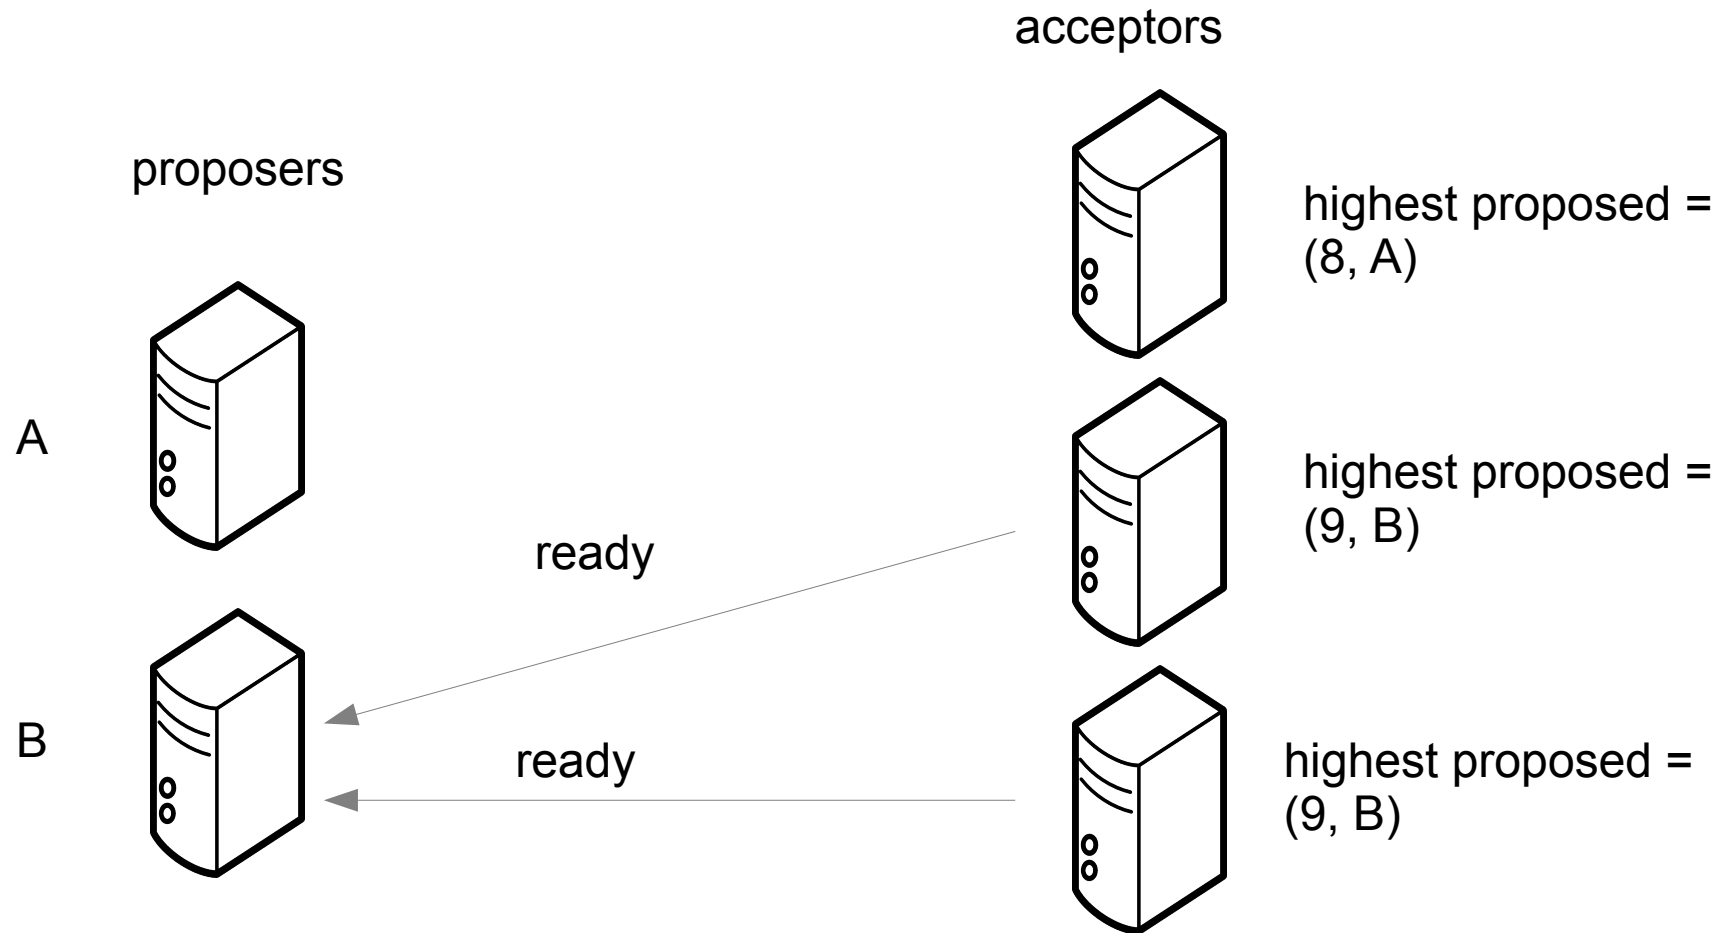
Paxos Example (2)



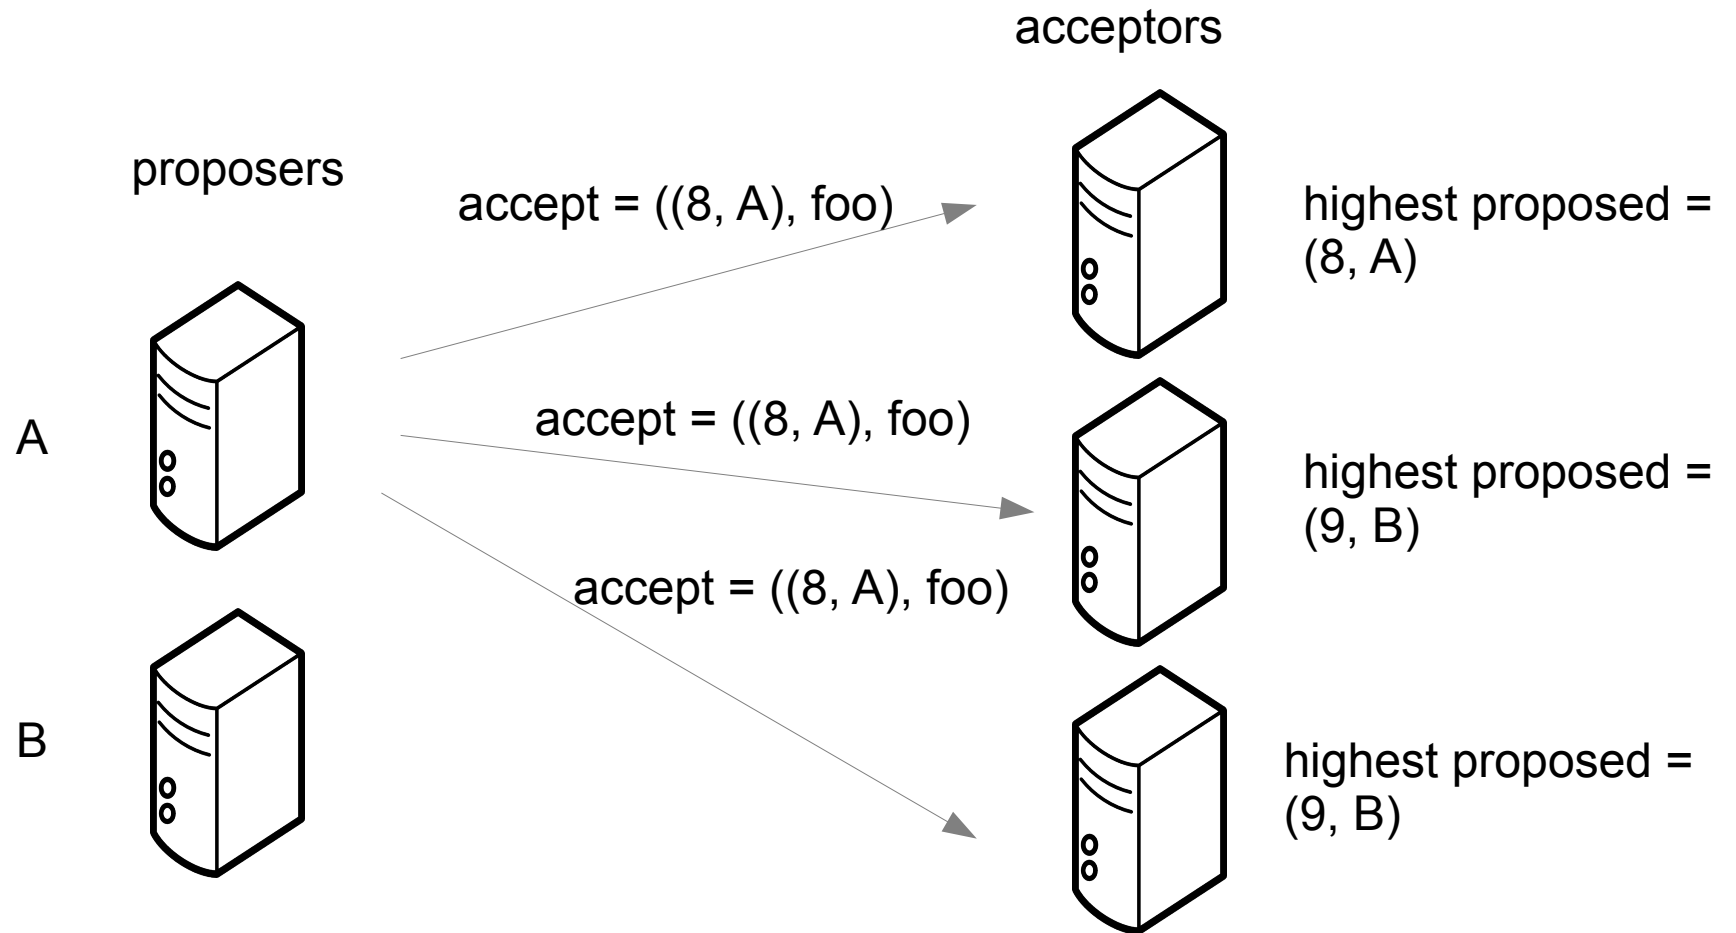
Paxos Example (2)



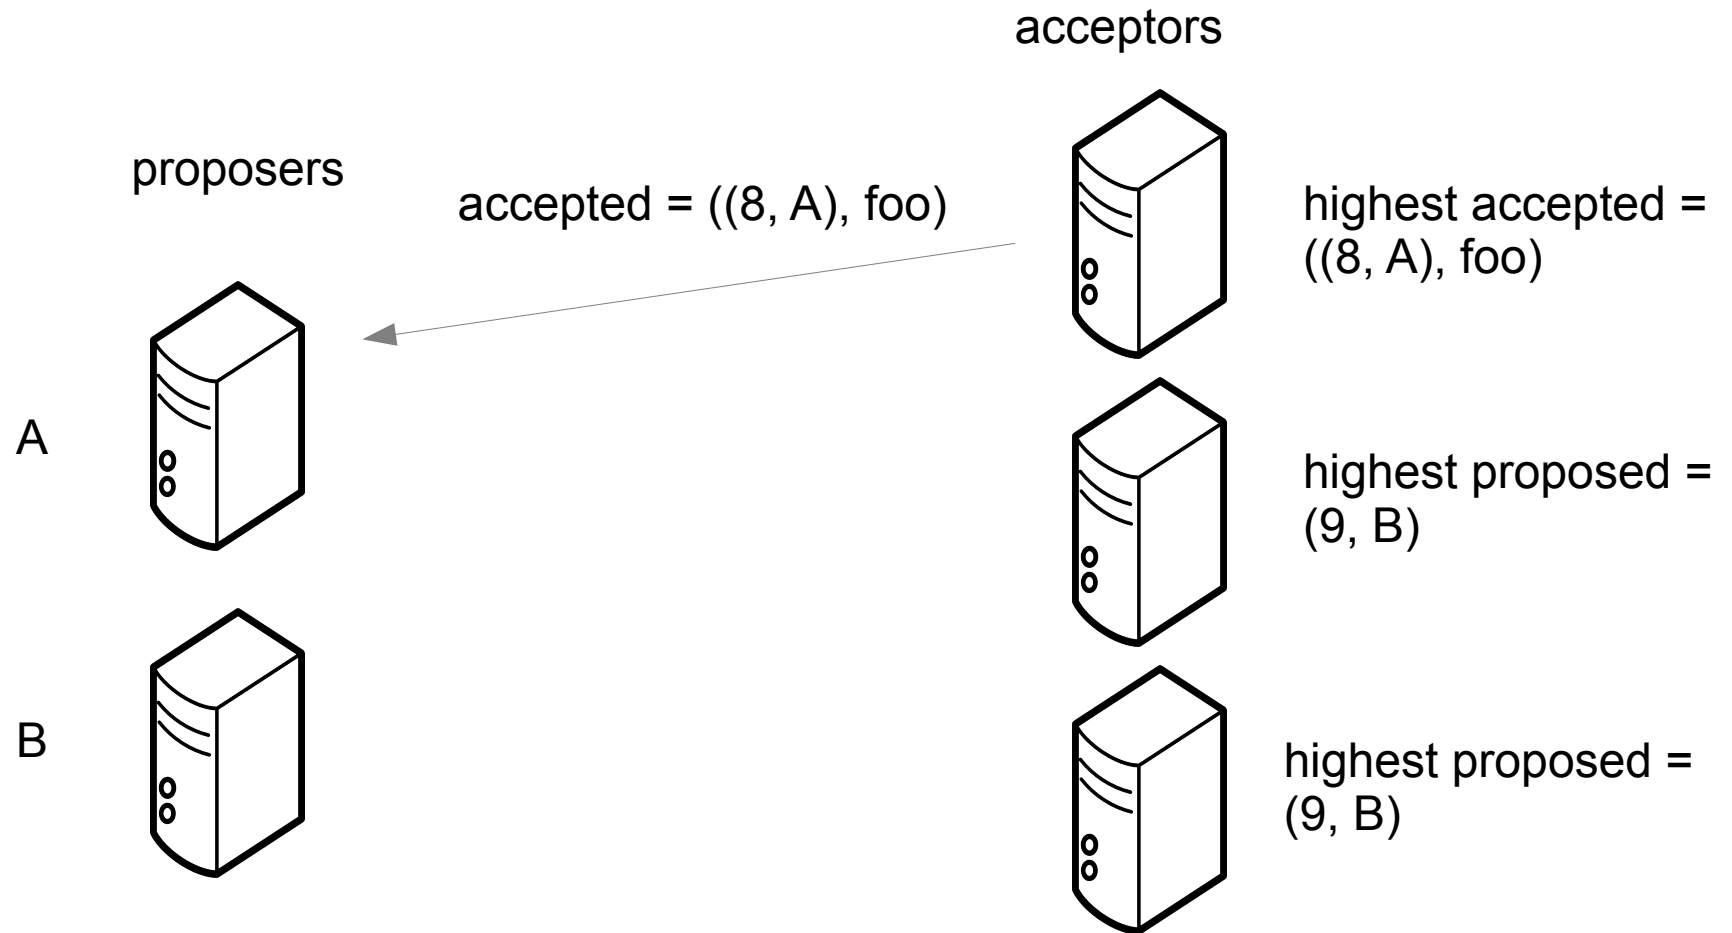
Paxos Example (2)



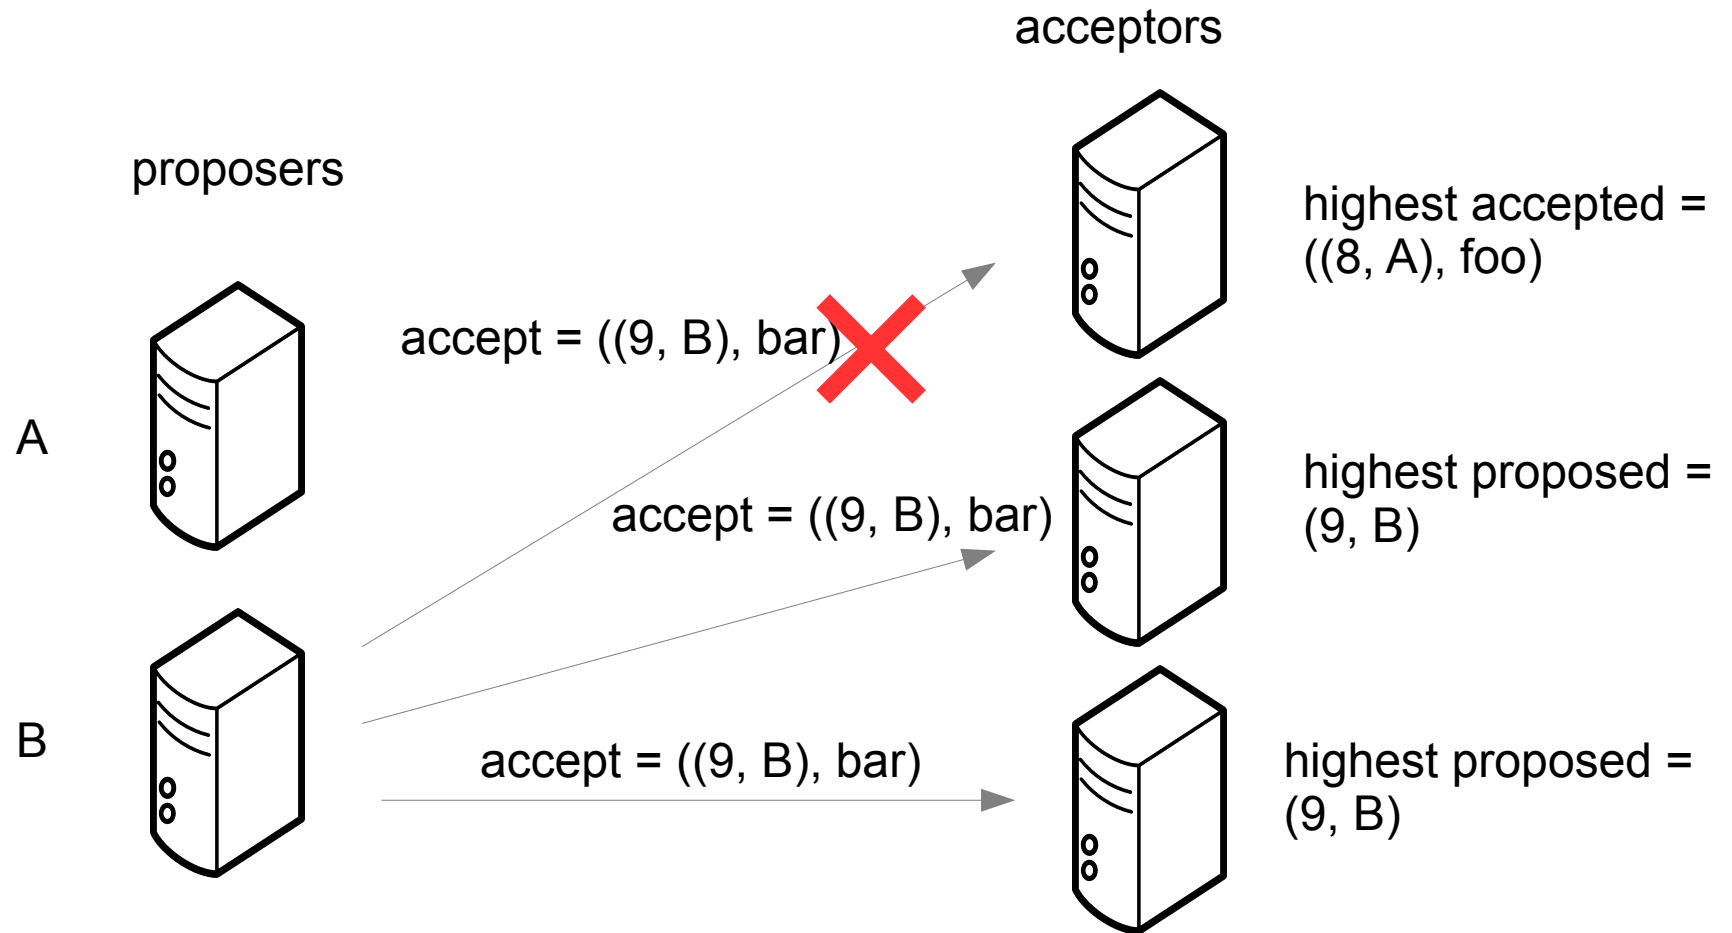
Paxos Example (2)



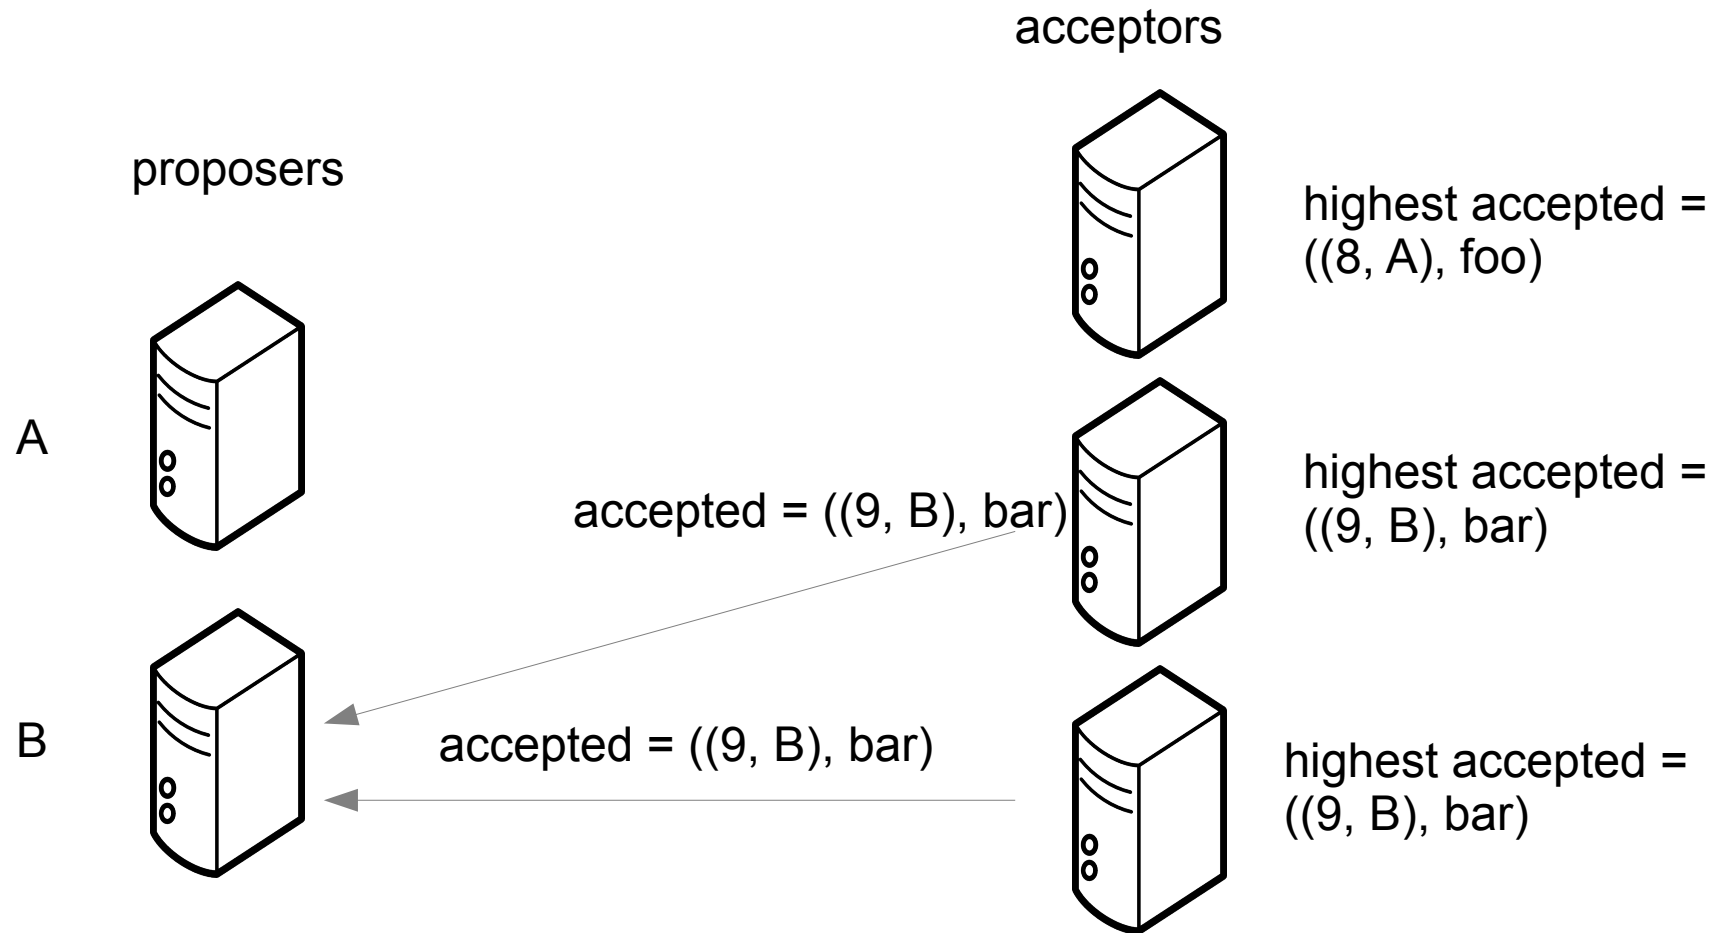
Paxos Example (2)



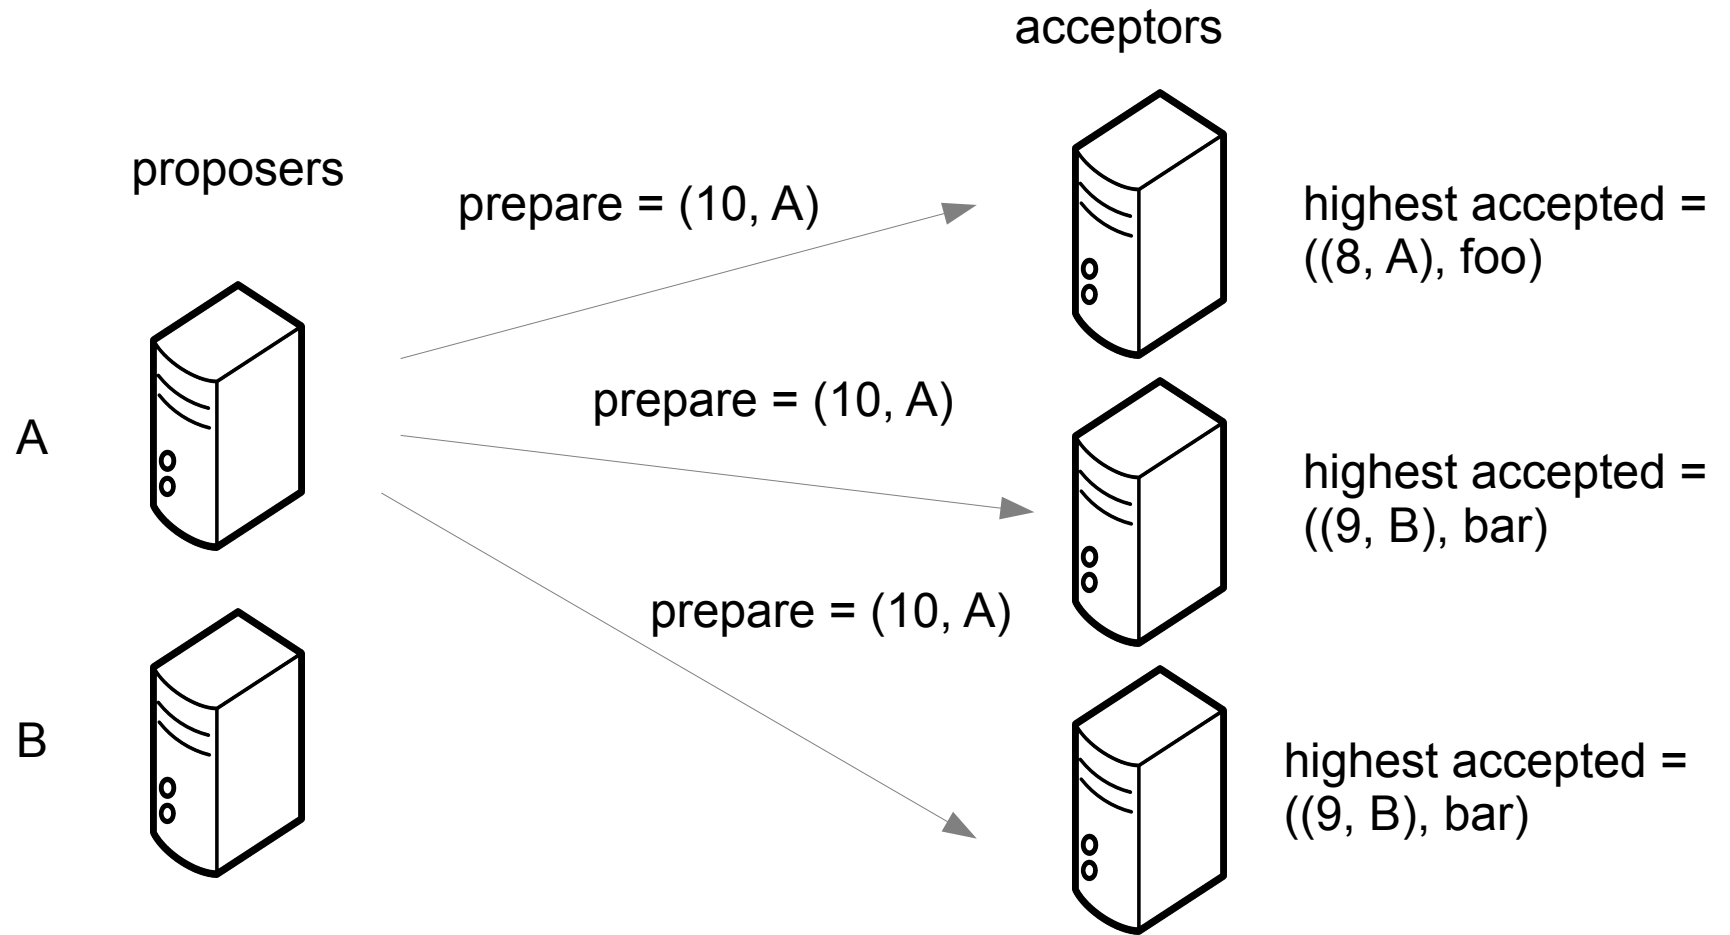
Paxos Example (2)



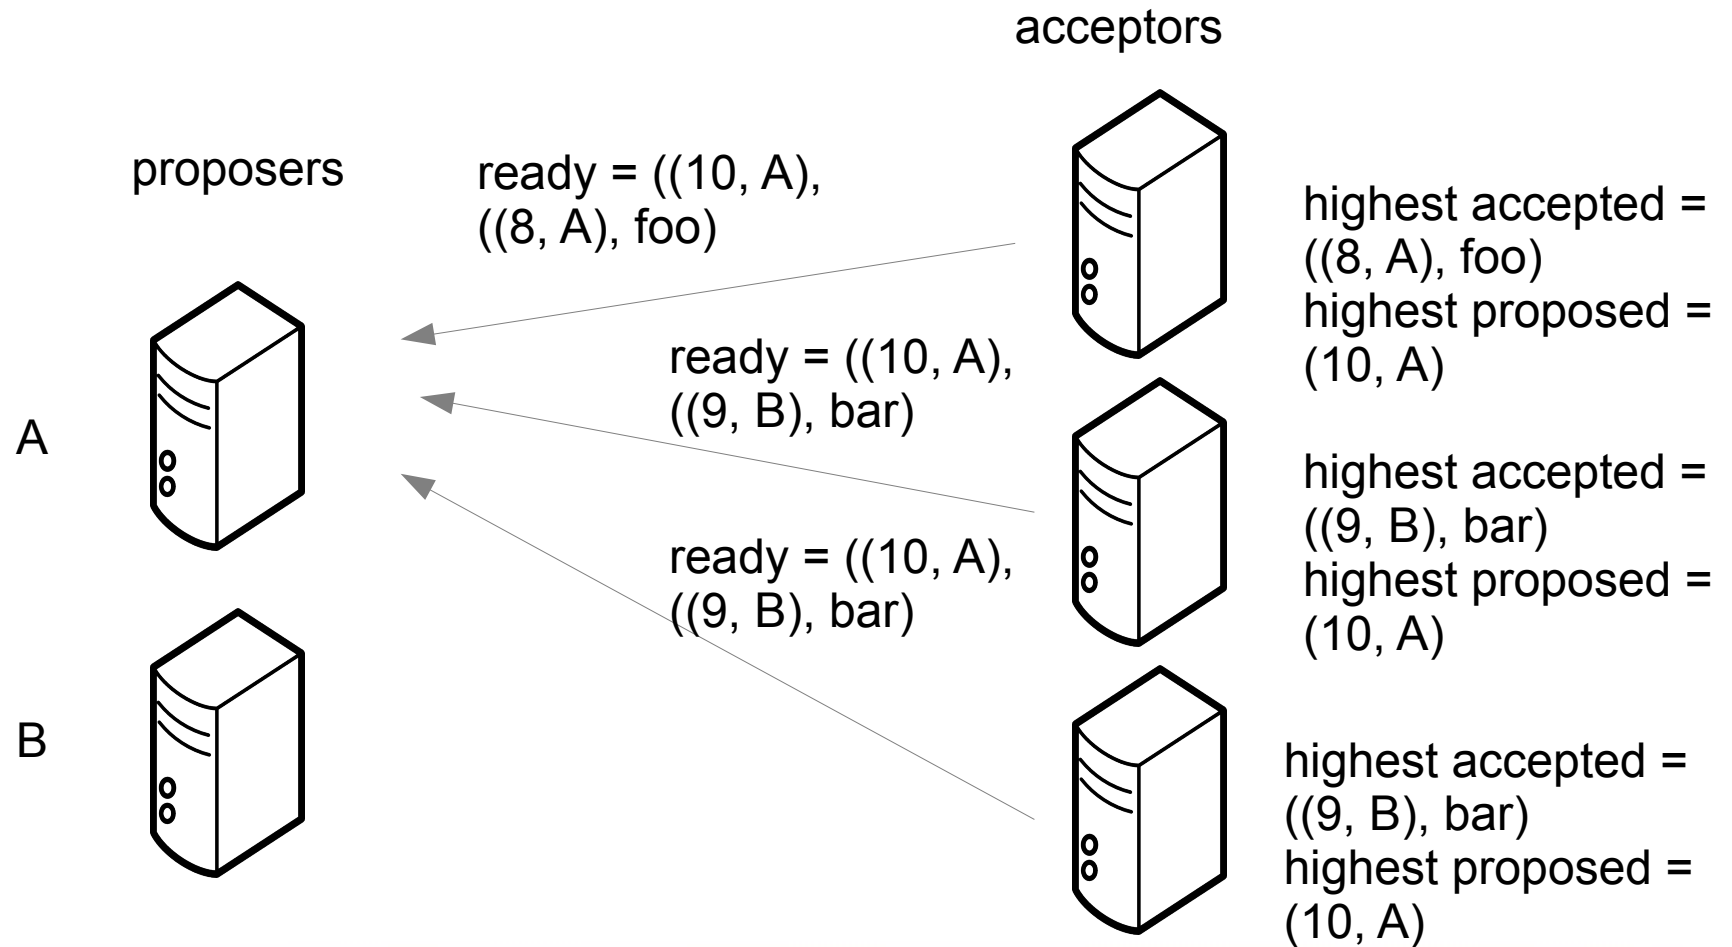
Paxos Example (2)



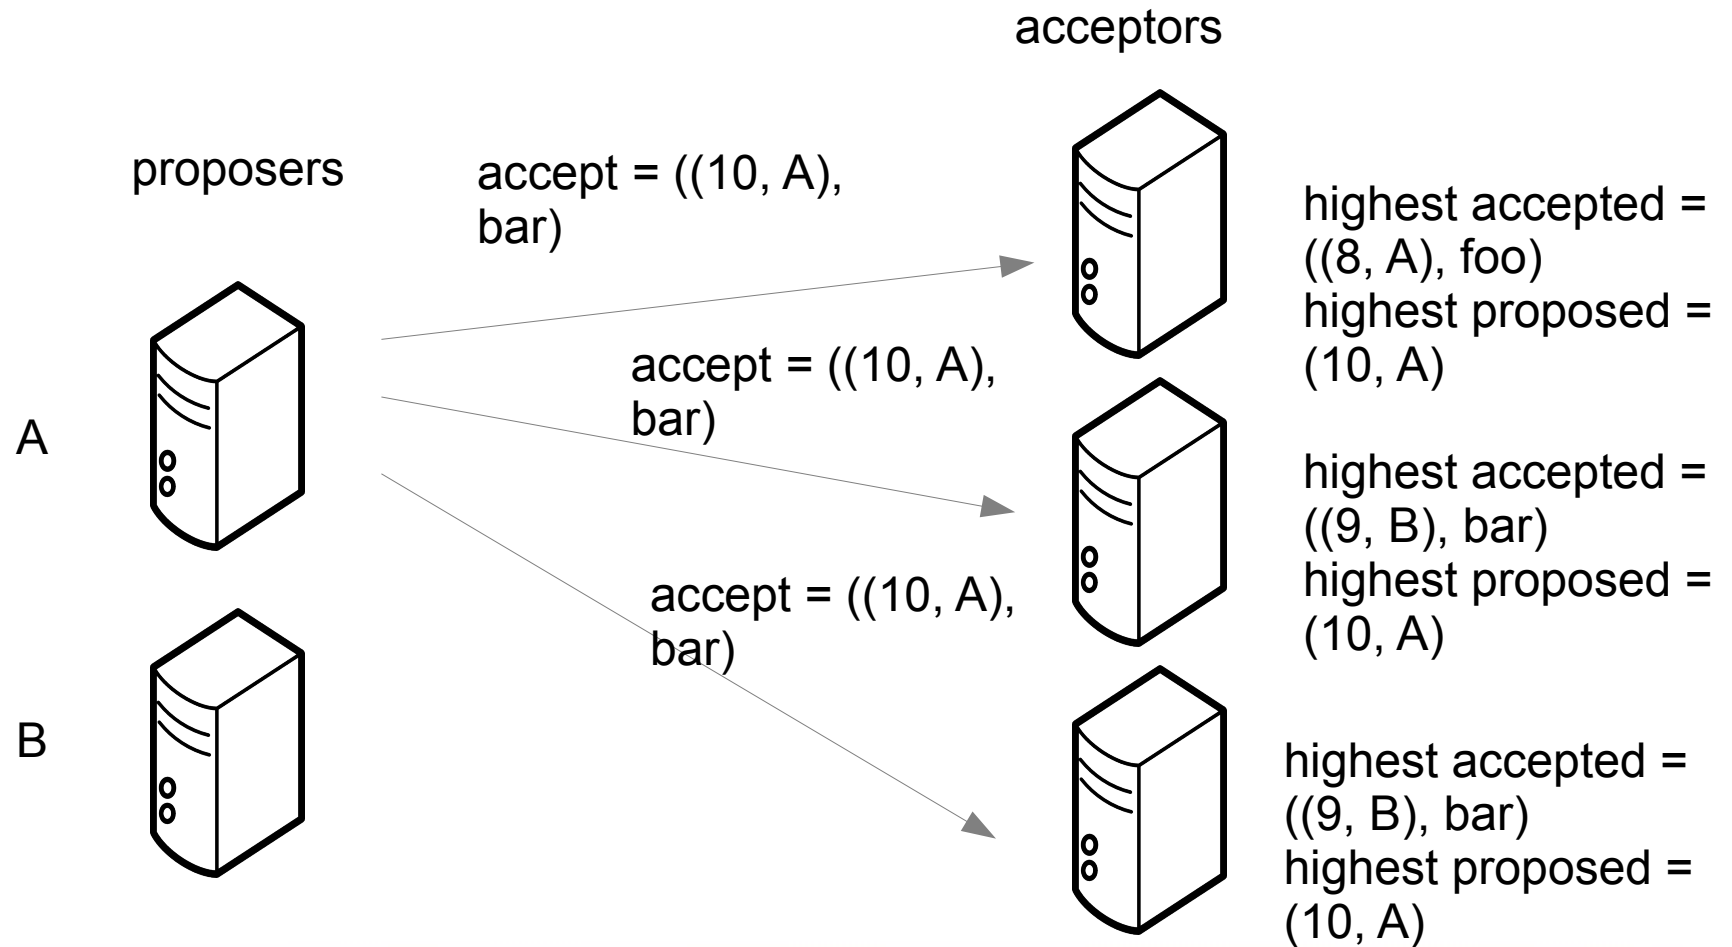
Paxos Example (2)



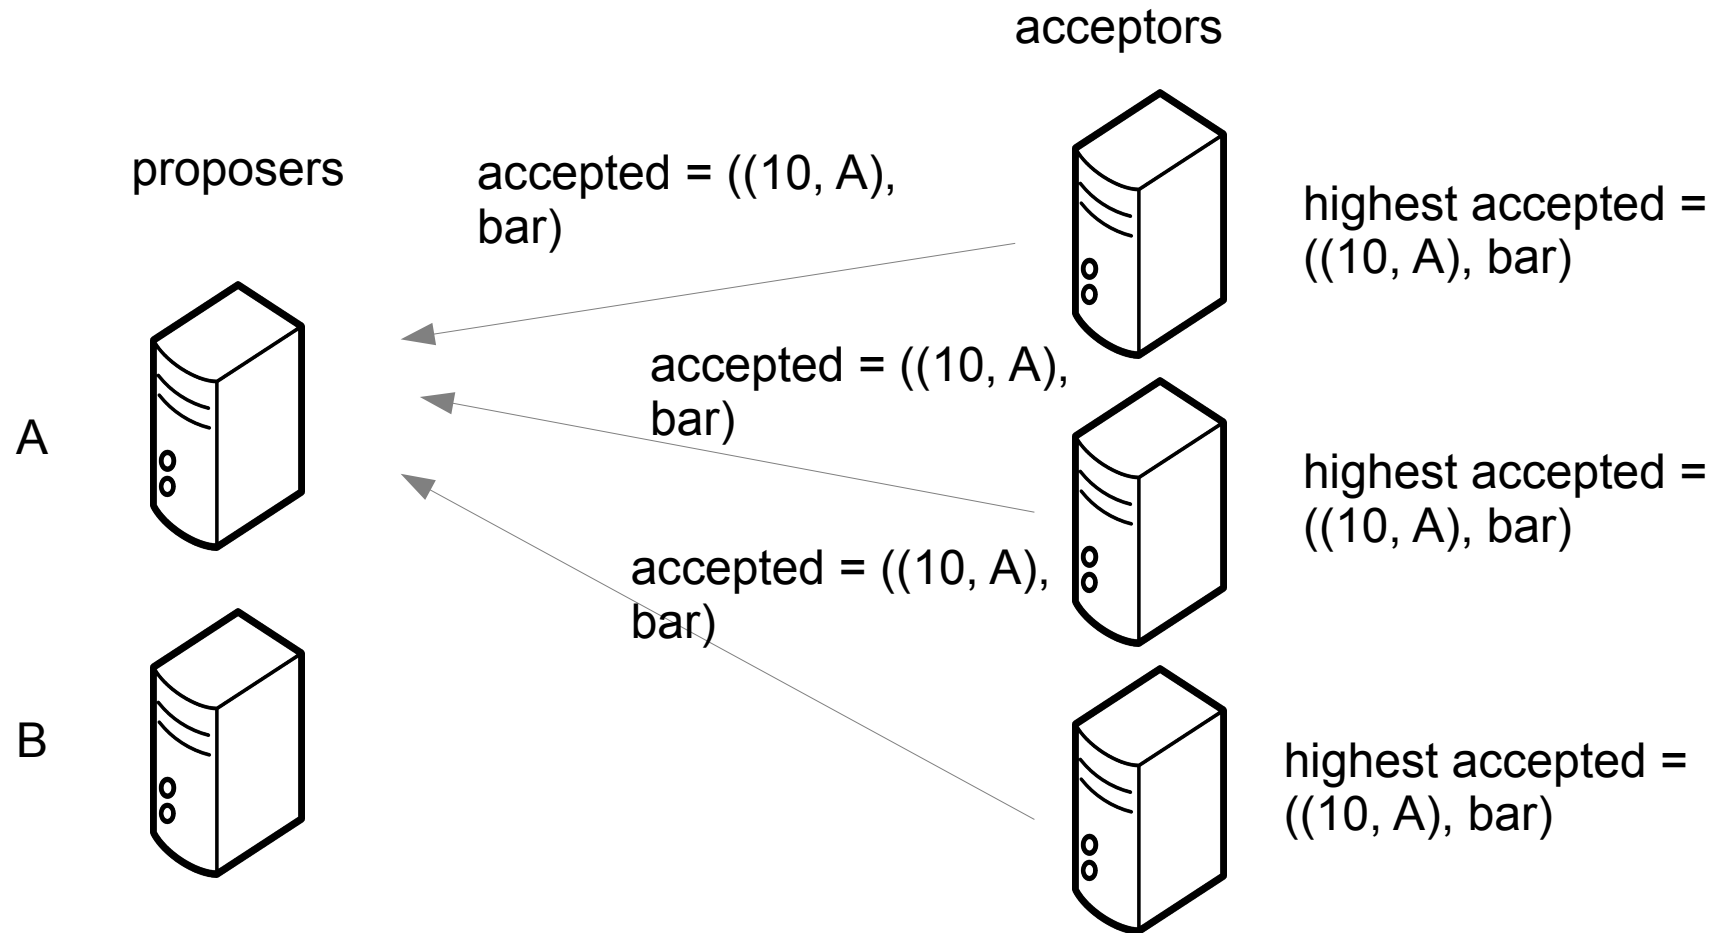
Paxos Example (2)



Paxos Example (2)



Paxos Example (2)



Progress

- You can imagine “dueling” proposers that continually propose higher and higher proposal numbers without any ever being accepted.
 - Solution: distinguished proposer.
-