

Distributed Computing


Informatyka Algorytmiczna 2021

Prof. Mirosław Kutylowski

Clients & servers, PAXOS

Client-server model

- Message passing
- Server executes commands of the client
- Single client for a server: commands with sequence numbers solve the problem



Multiple clients, multiple servers

Inconsistency problem: the same order executed on different servers may lead to different results

Client A: $x := x + 1$

Client B: $x := 2 * x$


Server 1: initially $x=0$, orders: A ($x=1$), then B ($x=2$)

Server 2: initially $x=0$, orders: B ($x=0$), then A ($x=1$)

Serializer


A straightforward but not scalable:

-
- Algorithm 15.9** State Replication with a Serializer
-
- 1: Clients send commands one at a time to the serializer
 - 2: Serializer forwards commands one at a time to all other servers
 - 3: Once the serializer received all acknowledgments, it notifies the client about the success
-



2 Lock protocol

deadlock



Algorithm 15.10 Two-Phase Protocol

Phase 1

- 1: Client asks all servers for the lock

Phase 2

- 2: if client receives lock from every server then
 - 3: Client sends command reliably to each server, and gives the lock back
 - 4: else
 - 5: Clients gives the received locks back
 - 6: Client waits, and then starts with Phase 1 again
 - 7: end if
-

2 Lock protocol

Algorithm 15.10 Two-Phase Protocol

Phase 1

1: Client asks all servers for the lock

Phase 2

2: **if** client receives lock from every server **then**

3: Client sends command reliably to each server, and gives the lock back

4: **else**

5: Client gives the received locks back

6: Client waits, and then starts with Phase 1 again

7: **end if**

2 Lock protocol

Algorithm 15.10 Two-Phase Protocol

Phase 1

1: Client asks all servers for the lock

Phase 2

2: **if** client receives lock from every server **then**

3: Client sends command reliably to each server, and gives the lock back

4: **else**

5: Client gives the received locks back

6: Client waits, and then starts with Phase 1 again

7: **end if**

What happens if some servers/clients do not respond?
Serious troubles!

Tickets concept -- PAXOS

Weaker than locks

Reissuable: new tickets can be issued even if old ones not returned

Expiration: a ticket accepted only if is it the most recent one

Ticket protocol --- 1-st trial

Algorithm 15.12 Naïve Ticket Protocol

Phase 1

1: Client asks all servers for a ticket

Phase 2

2: **if** a majority of the servers replied **then**

3: Client sends command together with ticket to each server

4: Server stores command only if ticket is still valid, and replies to client

5: **else**

6: Client waits, and then starts with Phase 1 again

7: **end if**

Phase 3

8: **if** client hears a positive answer from a majority of the servers **then**

9: Client tells servers to execute the stored command

10: **else**

11: Client waits, and then starts with Phase 1 again

12: **end if**

Ticket protocol – 1st trial

Problem:

Client A may store the commands on majority, then postpone phase 3

Client B may store some commands

Client A says to execute the stored command

PAXOS

Algorithm 15.13 Paxos

Client (Proposer)

Server (Acceptor)

Initialization

c \triangleleft *command to execute*
 $t = 0$ \triangleleft *ticket number to try*

$T_{\max} = 0$ \triangleleft *largest issued ticket*
 $C = \perp$ \triangleleft *stored command*
 $T_{\text{store}} = 0$ \triangleleft *ticket used to store C*

PAXOS

Phase 1

1: ~~$t = t + 1$~~

2: Ask all servers for ticket t

3: **if** $t > T_{\max}$ **then**

4: ~~$T_{\max} = t$~~

5: Answer with $\text{ok}(T_{\text{store}}, C)$

6: **end if**

--

Phase 2

- 7: **if** a majority answers ok **then**
- 8: Pick (T_{store}, c) with largest T_{store}
- 9: **if** $T_{\text{store}} > 0$ **then**
- 10: $c = c$
- 11: **end if**
- 12: Send propose(t, c) to same majority
- 13: **end if**

- 14: **if** $t = T_{\text{max}}$ **then**
- 15: $C = c$
- 16: $T_{\text{store}} = t$
- 17: **Answer success**
- 18: **end if**

Phase 3

```
19: if a majority answers success
    then
20:   Send execute(c) to every server
21: end if
```

Lemma 15.14. *We call a message $\text{propose}(t,c)$ sent by clients on Line 12 a proposal for (t,c) . A proposal for (t,c) is chosen, if it is stored by a majority of servers (Line 15). For every issued $\text{propose}(t',c')$ with $t' > t$ holds that $c' = c$, if there was a chosen $\text{propose}(t,c)$.*

PAXOS properties

- only one $\text{propose}(t,c)$ from a user for a given t
 - Indeed, before the next propose phase 1 must be executed with $t:=t+1$
- Assume there is $\text{propose}(t',c')$ with $t' \neq t$ and $c' \neq c$
 - let t' be the smallest one with this property
 - nonempty set S of servers that were involved in $\text{propose}(t',c')$ and $\text{propose}(t,c)$
 - a server s from S stored (t,c) , it must have occurred before accepting ticket t'
 - the client learns from s and is aware of c when issuing $\text{propose}(t',c')$, where c' is the most recent seen by the client
 - There is no more recent as c , as otherwise t' would not be minimal

PAXOS properties

COROLLARY If a command **c** is executed by some servers then eventually it will be executed by all servers.

Indeed: after the 1st **propose(t,c)** every proposal will be for **c**