

Reliable Distributed Storage

From Message Passing to Shared Memory

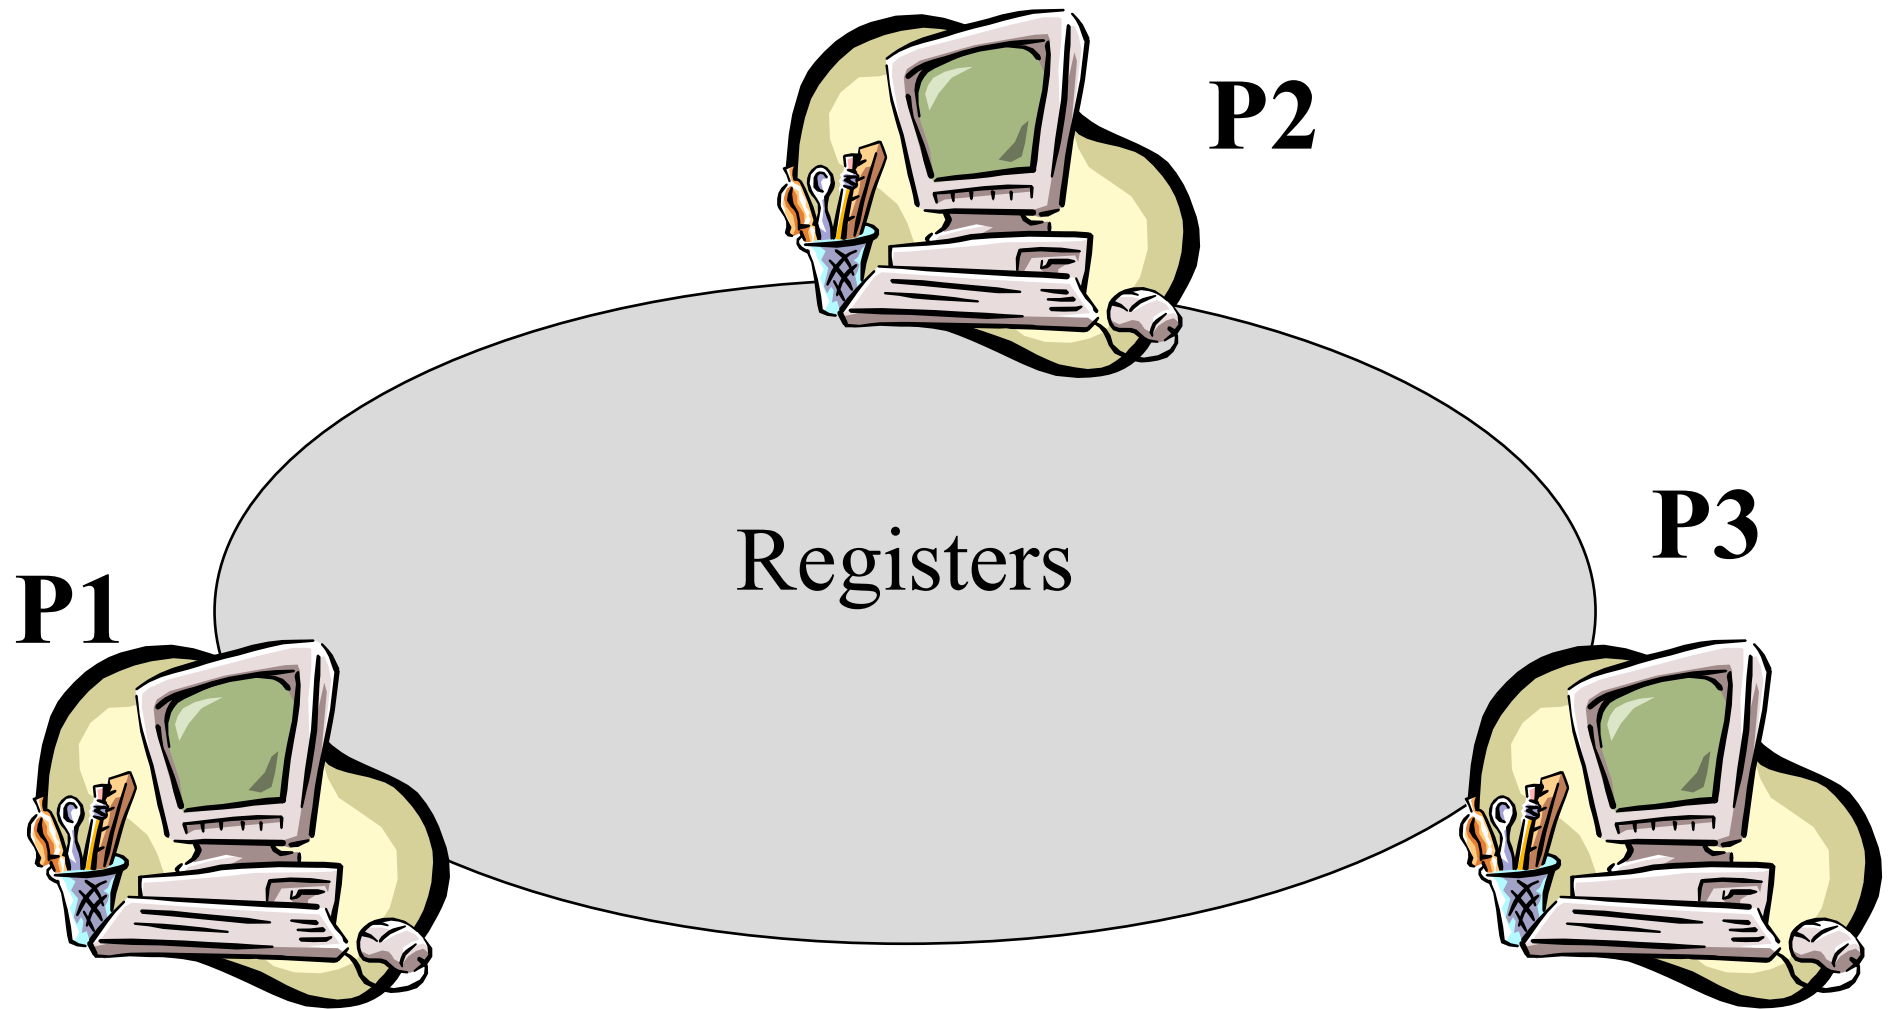
R. Guerraoui

Distributed Computing Laboratory

lcdwww.epfl.ch




The goal




Register: Specification

- A register contains ***integers*** : initial value 0
- Every value written is ***uniquely*** identified (this can be ensured by associating a process id and a timestamp with the value)
- Assume a register is local to a process, i.e., accessed only by one process: the value returned by a ***Read()*** is the last value written


Sequential execution



Concurrent execution




Execution with failures




Regular register

- Assumes only **one** writer
- Provides **strong** guarantees when there is no concurrent operations
- When some operations are concurrent, the register provides **minimal** guarantees
- **Read()** returns:
 - ✓ **the last value** written if there is no concurrent or failed operations
 - ✓ otherwise the last value written or **any** value concurrently written, i.e., the input parameter of some **Write()**


Execution




Results 1




Results 2



Results 3




Results 4



Correctness

- Results 1: non-regular register (safe)
- Results 2; 3; 4: regular register

Message passing model



Implementing a register

- Implementing ***Read()*** and ***Write()*** operations at every process
- Before returning a ***Read()*** value, or returning the ok of a ***Write()***, the process must communicate with other processes

A fail-stop algorithm

- We assume a ***fail-stop*** model:
 - Processes can fail by crashing (no recovery)
 - Channels are reliable
 - Failure detection is perfect (completeness and accuracy)

A fail-stop algorithm

- We implement a ***regular*** register
 - Every process p_i has a local copy of the register value v_i
 - Every process reads ***locally***
 - The writer writes ***globally***, i.e., at all (non-crashed) processes

A fail-stop algorithm

- Write(v) at p_i
 - send $[W, v]$ to all
 - for every p_j , wait until either:
 - receive $[ack]$ or
 - detect $[p_j]$
 - Return ok

- At p_i :
 - when receive $[W, v]$ from p_j
 - $v_i := v$
 - send $[ack]$ to p_j

- Read() at p_i
 - Return v_i

Correctness (liveness)

- ✓ A Read() is local and eventually returns
- ✓ A Write() eventually returns, by the
 - (a) the completeness property of the failure detector, and
 - (b) the reliability of the channels

Correctness (safety – 1)

- ☛ (a) In the absence of concurrent or failed operation, a Read() returns the last value written
 - ☛ Assume a Write(x) terminates and no other Write() is invoked. By the accuracy property of the failure detector, the value of the register at all processes that did not crash is x. Any subsequent Read() invocation by some process p_j returns the value of p_j , i.e., x, which is the last written value

Correctness (safety – 2)

- (b) A Read() returns the value concurrently written or the last value written
 - Let x be the value returned by a Read(): by the properties of the channels, x is the value of the register at some process. This value does necessarily come from the last or a concurrent Write().

But

- ☛ What if failure detection is not perfect?
- ☛ Can we devise an algorithm that implements a regular register and tolerates an arbitrary number of crash failures?

Lower bound

- ***Proposition:*** any wait-free asynchronous implementation of a regular register requires a majority of correct processes
- Proof (sketch): assume a $\text{Write}(v)$ is performed and $n/2$ processes crash, then a $\text{Read}()$ is performed and the other $n/2$ processes are up: the $\text{Read}()$ cannot see the value v
- The impossibility holds even with a 1-1 register (one writer and one reader)

The majority algorithm [ABD95]

- P_1 is the writer and any process can be reader
- A majority of the processes is correct (the rest can fail by crashing – no recovery)
- Channels are reliable
- Every process p_i maintains a local copy of the register v_i , as well as a sequence number sni and a read timestamp rsi
- Process p_1 maintains in addition a timestamp ts_1

Algorithm - Write()

- Write(v) at p_1
 - ✓ ts_1++
 - ✓ send $[W, ts_1, v]$ to all
 - ✓ when receive $[W, ts_1, ack]$ from majority
 - ✓ Return ok
- At p_i
 - ✓ when receive $[W, ts_1, v]$ from p_1
 - ✓ If $ts_1 > sn_i$ then
 - | $vi := v$
 - | $sn_i := ts_1$
 - | send $[W, ts_1, ack]$ to p_1


Algorithm - Read()

- Read() at p_i
 - ✓ $rsi++$
 - ✓ send $[R, rsi]$ to all
 - ✓ when receive $[R, rsi, snj, vj]$ from majority
 - ✓ $v := v_j$ with the largest snj
 - ✓ Return v
- At p_i
 - ✓ when receive $[R, rsj]$ from p_j
 - ✓ send $[R, rsj, sni, vi]$ to p_j

What if?

- Any process that receives a write message (with a timestamp and a value) updates its value and sequence number, i.e., without checking if it actually has an older sequence number

Old writes



Correctness 1

- ✓ Liveness: Any ***Read()*** or ***Write()*** eventually returns by the assumption of a majority of correct processes (if a process has a newer timestamp and does not send $[W, ts1, ack]$, then the older ***Write()*** has already returned)
- ✓ Safety 2: By the properties of the channels, any value read is the last value written or the value concurrently written

Correctness 2 (safety – 1)

- ☛ (a) In the absence of concurrent or failed operation, a ***Read()*** returns the last value written
 - ☛ Assume a Write(x) terminates and no other Write() is invoked. A majority of the processes have x in their local value, and this is associated with the highest timestamp in the system. Any subsequent Read() invocation by some process p_j returns x, which is the last written value

Atomicity

- ***An atomic register*** provides strong guarantees even when there is concurrency and failures: the execution is equivalent to a sequential and failure-free execution (***linearization***)
- Every failed (write) operation appears to be either complete or not to have been invoked at all


And

- Every complete operation appears to be executed at some instant between its invocation and reply time events


Regular vs Atomic

- For a regular register to be atomic, two successive ***Read()*** must not overlap a ***Write()***
- The regular register might in this case allow the first ***Read()*** to obtain the new value and the second ***Read()*** to obtain the old value


Sequential execution




Concurrent execution




Execution with failures




Execution 1




Execution 2




Execution 3




Execution 4




Execution 5



Execution 6



Execution 7



Fail-stop algorithms

- We first assume a fail-stop model; more precisely:
 - Any number of processes can fail by crashing (no recovery)
 - Channels are reliable
 - Failure detection is perfect: accuracy and completeness

The regular algorithm

- Consider our fail-stop ***regular*** register algorithm
 - Every process has a local copy of the register value
 - Every process reads ***locally***
 - The writer writes ***globally***, i.e., at all (non-crashed) processes


The regular algorithm

- Write(v) at p_i
 - send $[W, v]$ to all
 - for every p_j , wait until either:
 - received $[ack]$ or
 - detect $[p_j]$
 - Return ok


- At p_i :
 - when receive $[W, v]$ from p_j
 - $v_i := v$
 - send $[ack]$ to p_j

- Read() at p_i
 - Return v_i

Atomicity?




Linearization?




Fixing the pb: read-globally

- Read() at p_i
 - send $[W, v_i]$ to all
 - for every p_j , wait until either:
 - receive $[ack]$ or
 - detect $[p_j]$
 - Return v_i

Still a problem



Linearization?



A fail-stop 1-1 atomic algorithm

Write(v) at p1

- send [W,v] to p2
- Wait until either:
 - receive [ack] from p2 or
 - detect [p2]
- Return ok

At p2:

when receive [W,v]
from p1
 $v2 := v$
send [ack] to p2

Read() at p2

- Return v2

A fail-stop 1-N algorithm

- every process maintains a local value of the register as well as a sequence number
- the writer, p_1 , maintains, in addition a timestamp ts_1
- any process can read in the register

A fail-stop 1-N algorithm

Write(v) at p_1

- $ts1++$
- send $[W, ts1, v]$ to all
- for every p_i , wait until either:
 - receive $[ack]$ or
 - detect $[p_i]$
- Return ok

Read() at p_i

- send $[W, sni, v_i]$ to all
- for every p_j , wait until either:
 - receive $[ack]$ or
 - suspect $[p_j]$
- Return v_i

A 1-N algorithm (cont'd)

- At p_i

- When p_i receive $[W, ts, v]$ from p_j

- if $ts > sn_i$ then

- $v_i := v$


- $sn_i := ts$

- send [ack] to p_j

Why not N-N?


P1

$R() \rightarrow Y$




P2

$W(X)$




$W(Y)$



P3

$W(Z)$



The Write() algorithm

- Write(v) at p_i
 - ✓ send $[W]$ to all
 - ✓ for every p_j wait until
 - **receive $[W, sn_j]$ or**
 - **suspect p_j**
 - ✓ $(sn, id) := (\text{highest } sn_j + 1, i)$
 - ✓ send $[W, (sn, id), v]$ to all
 - ✓ for every p_j wait until
 - **receive $[W, (sn, id), ack]$ or**
 - **detect $[p_j]$**
 - ✓ Return ok
- At p_i
 - T1:
 - ✓ when receive $[W]$ from p_j
 - send $[W, sn]$ to p_j
 - T2:
 - ✓ when receive $[W, (sn_j, id_j), v]$ from p_j
 - ✓ If $(sn_j, id_j) > (sn, id)$ then
 - $vi := v$
 - $(sn, id) := (sn_j, id_j)$
 - ✓ send $[W, (sn_j, id_j), ack]$ to p_j

The Read() algorithm

- Read() at p_i
 - ✓ send [R] to all
 - ✓ for every p_j wait until
 - **receive [R,(snj,idj),vj] or**
 - **suspect p_j**
 - ✓ $v = v_j$ with the highest (snj,idj)
 - ✓ (sn,id) = highest (snj,idj)
 - ✓ send [W,(sn,id),v] to all
 - ✓ for every p_j wait until
 - **receive [W,(sn,id),ack] or**
 - **detect [p_j]**
 - ✓ Return v
- At p_i
 - T1:
 - ✓ when receive [R] from p_j
 - send [R,(sn,id),vi] to p_j
 - T2:
 - ✓ when receive [W,(snj,idj),v] from p_j
 - ✓ If (snj,idj) > (sn,id) then
 - $v_i := v$
 - (sn,id) := (snj,idj)
 - ✓ send [W,(snj,idj),ack] to p_j

From fail-stop to fail-silent

- We assume a majority of correct processes
- In the 1-N algorithm, the writer writes in a majority using a timestamp determined locally and the reader selects a value from a majority and then imposes this value on a majority
- In the N-N algorithm, the writers determines first the timestamp using a majority