

The Power of Registers

Prof R. Guerraoui
Distributed Programming Laboratory




© R. Guerraoui

1




Atomic execution



2

Atomic execution



3

Registers

- Question 1: what objects can we implement with registers? (this lecture)
- Question 2: what objects we cannot implement? (next lecture)

4

Wait-free implementations of atomic objects

- An **atomic** object is simply defined by its sequential specification; i.e., by how its operations should be implemented when there is no concurrency
- Implementations should be **wait-free**: every process that invokes eventually gets a reply (unless the process crashes)

5

Counter (sequential spec)

- A **counter** has two operations **inc()** and **read()** and maintains an integer x *init to 0*
- read()**:
 - return(x)
- inc()**:
 - $x := x + 1;$
 - return(ok)


6

Naive implementation

- ☞ The processes share one register Reg
- ☞ *read()*:
 - ☞ return(Reg.read());
- ☞ *inc()*:
 - ☞ temp := Reg.read()+1;
 - ☞ Reg.write(temp);
 - ☞ return(ok)

7

Atomic execution?




8

Atomic implementation

- ☞ The processes share an array of registers Reg[1,..,n]
- ☞ *inc()*:
 - ☞ temp := Reg[i].read() + 1;
 - ☞ Reg[j].write(temp);
 - ☞ return(ok)

9

Atomic execution?



10

Atomic implementation

- ☞ *read()*:
 - ☞ sum := 0;
 - ☞ for j = 1 to n do
 - ☞ sum := sum + Reg[j].read();
 - ☞ return(sum)

11

Snapshot (sequential spec)

- ☞ A *snapshot* has operations *update()* and *scan()* and maintains an array *x* of size *n*
- ☞ *scan()*:
 - ☞ return(x)
- ☞ *update(i,v)*:
 - ☞ x[i] := v;
 - ☞ return(ok)

12


Very naive implementation

Each process maintains an array of integer variables x init to $[0, \dots, 0]$

- **scan()**:
 - return(x)
- **update(i, v)**:
 - $x[i] := v$;
 - return(ok)

13

Atomic execution?



14


Less naive implementation

The processes share one array of N registers $\text{Reg}[1, \dots, N]$

- **scan()**:
 - for $j = 1$ to N do
 - $x[j] := \text{Reg}[j].\text{read}()$;
 - return(x)
- **update(i, v)**:
 - $\text{Reg}[i].\text{write}(v)$; return(ok)


15

Atomic execution?




16

Atomic execution?



17

Atomic execution?



18

Non-atomic vs atomic snapshot

What we implement here is some kind of **regular snapshot**:

A **scan** returns, for every index of the snapshot, the last written values or the value of any concurrent update

We call it **collect**

19


Key idea for atomicity

To **scan**, a process keeps reading the entire snapshot (i.e., it **collect**), until two results at the **same**

This means that the snapshot did not change, and it is safe to return without violating atomicity

20

Same value vs. Same timestamp



21

Enforcing atomicity

The processes share one array of N registers `Reg[1,..,N]`; each contains a value and a timestamp

We use the following operation for modularity

collect():

- for $j = 1$ to N do
 - `x[j] := Reg[j].read();`
- return(x)

22

Enforcing atomicity (cont'd)

scan():


- `temp1 := self.collect();`
- while(true) do
 - `temp2 := self.collect();`
 - `temp1 := temp2;`
 - if (`temp1 = temp2`) then
 - return (`temp1.val`)

update(i, v):

- `ts := ts + 1;`
- `Reg[i].write(v,ts);`
- return(ok)

23

Wait-freedom?



24

Key idea for atomicity & wait-freedom

- ☞ The processes share an array of *registers* $Reg[1,..,N]$ that contains each:
 - ☞ a value,
 - ☞ a timestamp, and
 - ☞ a copy of the entire array of values

25

Key idea for atomicity & wait-freedom (cont'd)

- ☞ To *scan*, a process keeps collecting and returns a collect if it did not change, or some collect returned by a concurrent *scan*
 - ☞ Timestamps are used to check if the collect changes or if a scan has been taken in the meantime
- To *update*, a process *scans* and writes the value, the new timestamp and the result of the scan

26

Snapshot implementation

Every process keeps a local timestamp ts

- ☞ *update(i,v)*:
 - ☞ $ts := ts + 1$;
 - ☞ $Reg[i].write(v,ts,self.scan())$;
 - ☞ return(ok)


27

Snapshot implementation

- ☞ *scan()*:
 - ☞ $t1 := self.collect()$; $t2 := t1$
 - ☞ while(true) do
 - ☞ $t3 := self.collect()$;
 - ☞ if ($t3 = t2$) then return ($t3[j,3]$);
 - ☞ for $j = 1$ to N do
 - ☞ if ($t3[j,2] \geq t1[j,2] + 2$) then
 - ☞ return ($t3[j,3]$)
 - ☞ $t2 := t3$

28

Possible execution?



29