

Computing with anonymous processes

Prof R. Guerraoui
Distributed Programming Laboratory



© R. Guerraoui

1



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)

2

Counter (atomic implementation)

The processes share an array of SWMR registers $Reg[1, \dots, n]$; the writer of register $Reg[i]$ is p_i

- *inc()*:
 - $temp := Reg[i].read() + 1$;
 - $Reg[i].write(temp)$;
 - return(ok)

3

Counter (atomic implementation)

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


4

Weak Counter

- A *weak counter* has one operation *wInc()*
- *wInc()*:
 - $x := x + 1$;
 - return(x)
- correctness: if an operation precedes another, then the second returns a value that is larger than the first one (regularity vs atomicity)

5

Weak counter execution



6

Weak Counter (lock-free implementation)


• The processes share an (infinite) array of MWMM registers $\text{Reg}[1, \dots, n, \dots]$, init to 0

• *wInc()*:

- $i := 0;$
- while ($\text{Reg}[i].\text{read}() \neq 0$) do
 - $i := i + 1;$
- $\text{Reg}[i].\text{write}(1);$
- return(i);

7

Weak counter execution



8

Weak Counter (wait-free implementation)

• The processes also use a MWMM register L

• *wInc()*:

- $i := 0;$
- while ($\text{Reg}[i].\text{read}() \neq 0$) do
 - if L has been updated n times then
 - return the largest value seen in L
 - $i := i + 1;$
- L.write(i);
- $\text{Reg}[i].\text{write}(1);$
- return(i);

9

Weak Counter (wait-free implementation)

• *wInc()*:

- $t := l := L.\text{read}(); i := 0;$
- while ($\text{Reg}[i].\text{read}() \neq 0$) do
 - if $L.\text{read}() \neq l$ then
 - $l := L.\text{read}(); t := \max(t, l); i := i + 1;$
 - if $k = n$ then return(t)
 - L.write(i);
- $\text{Reg}[i].\text{write}(1);$
- return(i);

10

Snapshot (sequential spec)

• A *snapshot* has operations *update()* and *scan()* and maintains an array x of size n

• *scan()*:

- return(x)

• NB. No component is devoted to a process

• *update(i, v)*:

- $x[i] := v;$
- return(ok)

11

Key idea for atomicity & wait-freedom

• The processes share a *Weak Counter*. Wcounter, init to 0;

• The processes share an array of *registers* $\text{Reg}[1, \dots, N]$ that contains each:

- a value,
- a timestamp, and
- a copy of the entire array of values

12

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 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

13

Snapshot implementation

Every process keeps a local timestamp ts

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

14

Snapshot implementation

- ☞ *scan()*:
 - ☞ $ts := Wcounter.wInc()$;
 - ☞ while(true) do
 - ☞ If some $Reg[j]$ contains a collect with a higher timestamp than ts , then return that collect
 - ☞ If $n+1$ sets of reads return identical results then return that one

15

Consensus (obstruction-free)

- ☞ We consider binary consensus
- ☞ The processes share two infinite arrays of registers: $Reg_0[i]$ and $Reg_1[i]$
- ☞ Every process holds an integer i init to 1
- ☞ Idea: to impose a value v , a process needs to be fast enough to fill in registers $Reg_v[i]$

16

Consensus (obstruction-free)

- ☞ *propose(v)*:
 - ☞ while(true) do
 - ☞ If $Reg_{1-v}[i] = 0$ then
 - ☞ $Reg_v[i] := 1$;
 - ☞ if $i > 1$ and $Reg_{1-v}[i-1] = 0$ then $return(v)$;
 - ☞ else $v := 1-v$;
 - ☞ $i := i+1$;
 - end

17

Consensus (solo process)

$q(1)$

Reg0(1)=0

Reg1(1)=1


Reg0(2)=0

Reg1(2)=1

Reg0(1)=0

18

Consensus (lock-step)



19

Consensus (binary)

```
propose(v):  
  while(true) do  
    If Reg1-v[i] = 0 then  
      Regv[i] := 1;  
    if i > 1 and Reg1-v[i-1] = 0 then  
      return(v);  
    else if Regv[i] = 0 then v := 1-v;  
    if v = 1 then wait(2i)  
    i := i+1;  
  end
```

20