

## FDS 2012: Implementing an atomic bit

© 2012 P. Kuznetsov

## Administrivia

- Static web page:
  - ✓ <http://www.inet.tu-berlin.de/menue/teaching0/ss2012/fds12/parameter/en/>
- Next tutorial: May 7
- Next class: May 8




© 2012 P. Kuznetsov

2

## The space of registers

- Nb of writers and readers:  
from 1W1R to NWNR
- Size of the value set: from  
binary to multi-valued
- Safety properties: safe,  
regular, atomic



All registers are (computationally) equivalent!



© 2012 P. Kuznetsov

3

## Transformations

From 1W1R binary safe to 1WNR multi-valued atomic

- I. From safe to regular (1W1R)
- II. From one-reader to multiple-reader (regular binary or multi-valued)
- III. From binary to multi-valued (1WNR regular)
- IV. From regular to atomic (1W1R)
- V. From 1W1R to 1WNR (multi-valued atomic)
- VI. From 1WNR to NWNR (multi-valued atomic)



© 2012 P. Kuznetsov

4

## This class

- Implementing an atomic bit
- Atomic snapshot: reading multiple locations atomically



5

## How to implement an atomic bit?

The problem: implement a binary 1W1R atomic register (atomic bit) from binary 1W1R safe ones (safe bits).



6

## A brute-force approach

- Binary 1W1R safe => binary 1W1R regular
- Binary 1W1R regular => multi-valued 1W1R regular
- Multi-valued 1W1R regular => multi-valued 1W1R atomic

The reader never writes (to the base objects), but...

☹ Unbounded # of base safe bits (to account for sequence numbers)



7

## An optimal solution

- No sequence numbers?
- Bounded number of safe bits,  $O(1)$ ?
- Bounded number of base actions,  $O(1)$ ?


Can we do it if the reader does not write?



8

## Safe bit to regular bit? Easy

- the writer is allowed only to *change* the value
- What about atomic?




9

## Impossible if the reader does not write for bounded # of regular bits!

**Proof sketch** (by contradiction):

- Suppose only the writer executes writes on the base (regular) bits.
- Every write operation  $W(1)$  is a sequence of writes actions  $w_1, \dots, w_k$  on base regular bits
  - ✓ Corresponds to the sequence of **shared-memory states**  $s_0, s_1, \dots, s_k$  (defined for **sequential** runs)



© 2012 P. Kuznetsov

10

## Proof (contd): digests

- There are only finitely many states! (bounded # of base registers)
- Each sequence  $s_0, s_1, \dots, s_k$  of states (though possibly unbounded) defines a **digest**  $d_0, d_1, \dots, d_m$ 
  - ✓  $d_0 = s_0, d_m = s_k$  (same global state transition)
  - ✓  $d_i = s_j \Rightarrow i=j$  (all digest elements are distinct)
  - ✓ for all  $(d_i, d_{i+1})$ , exists  $(s_j, s_{j+1})$  such that  $s_j = d_i$  and  $s_{j+1} = d_{i+1}$   
7,4,8,4,2,8,3  $\Rightarrow$  7,4,8,3
- Each write operation “looks” like its digest
- There are only finitely many digests!



© 2012 P. Kuznetsov

11

## Proof (contd.): counter-example

- Consider a run with infinitely many alternating writes:  $W_1(1), W(0), W_2(1), \dots$  (no reads)
  - ✓ Writes  $W_1, W_2, \dots$  give an infinite sequence of digests  $D_1, D_2, \dots$
- At least one digest  $D = d_0, d_1, \dots, d_m$  appears infinitely often in  $D_1, D_2, \dots$ 
  - ✓ Why?
- We can amend our run with a sequence of reads  $R_0, R_1, \dots, R_m$  (in that order), each  $R_i$  “sees” state  $d_{m-i}$ 
  - ✓ How?



© 2012 P. Kuznetsov

12

## Proof (contd.): the “switch”

- $R_0$  “sees”  $d_m$  and, thus, returns 1
  - ✓ Could have happened right after  $W(1)$
- $R_m$  “sees”  $d_0$  and, thus, returns 0
  - ✓ Could have happened right before  $W(1)$

⇒ There exists  $i$  such that  $R_i$  returns 1 and  $R_{i+1}$  returns 0 (by induction on  $i=0, \dots, m$ )




© 2012 P. Kuznetsov

13

## Proof (contd.): contradiction

- The (sequential) execution of  $R_i$  and  $R_{i+1}$  is indistinguishable (to the reader) from a concurrent one




© 2012 P. Kuznetsov

14

## The reader must write

- And the writer must read
- But how the writer would tell what it read?
  - ✓ The writer needs at least two bits!
  - ✓ Why?
- Suppose the writer writes to one bit only
  - ✓ there are exactly two digests  $0,1$  and  $1,0$
  - ✓ suppose infinitely many  $W(1)$  operations export digests  $0,1$
  - ✓ new-old inversion:



© 2012 P. Kuznetsov

15

## Optimal construction?

- Two bits for the writer
  - ✓ REG: for storing the current value
  - ✓ WR: for signaling to the reader
- One bit for the reader
  - ✓ RR: for signaling to the writer

Necessary, but is it also sufficient?



© 2012 P. Kuznetsov

16

## Evolutionary approach: Iteration 1

The reader should be able to distinguish the two cases:

- ✓ A new value was written:  $WR \neq RR$ :
- ✓ The value is unchanged:  $WR = RR$ :

### Writer:

change REG  
if  $WR = RR$  then change WR

### Reader:

if  $WR \neq RR$  then change RR  
val := REG  
return val

Does not work: the read value does not depend on RR



© 2012 P. Kuznetsov

17

## Iteration 2

Return the "old" value if nothing changed  
(local variable val initialized to the initial value of REG)

### Writer:

change REG  
if  $WR = RR$  then change WR

### Reader:

if  $WR = RR$  then return val  
change RR  
val := REG  
return val




© 2012 P. Kuznetsov

18

## Counter-example 2

Does not work: a read finds  $WR \neq RR$ , a subsequent read finds  $WR \neq RR$  and reads an old value in REG (new-old inversion)



© 2012 P. Kuznetsov

19

## Counter-example 2

1.  $w_1(1)$  completes
2.  $r_1$  reads WR, finds  $WR \neq RR$  and changes RR
3.  $w_2(0)$  begins, changes REG to 0, reads RR, finds  $WR = RR$ , changes WR, restoring the predicate  $WR \neq RR$ , and completes
4.  $w_3(1)$  begins and starts changing REG from 0 to 1
5.  $r_1$  concurrently reads REG and returns the new value 1
6.  $r_2$  begins, finds  $RR \neq WR$ , reads REG and returns the old value 0



© 2012 P. Kuznetsov

20

### Iteration 3

Read WR twice, if WR changed while the read is executed, return a conservative (old) value

#### Writer:

change REG  
if WR=RR then change WR

#### Reader:

if WR=RR then return val  
aux := REG  
change RR  
val:= REG  
if WR=RR then return val  
return aux



### Counter-example 3

Still a problem:

- *aux* gets a **new** value and *val* gets an **old** value (the two reads of REG are concurrent with a write on REG)
- succeeding read observes RR=WR (the two reads of WR are concurrent with a write on WR)

Homework: construct the counter-example



### Iteration 4

Evaluate *val* again before returning the “old” value (*aux*): to make sure *val* is at least as old as *aux*

#### Writer:

change REG  
if WR=RR then change WR

#### Reader:


if WR=RR then return val  
aux := REG  
change RR  
val:= REG  
if WR=RR then return val  
val:= REG  
return aux



### Counter-example 4

Still (!) a problem:

Unconditionally changing RR may be invalidated by a concurrent update of WR



Solution: check if  $WR \neq RR$  again before changing RR



## Final solution [Tromp, 1989]

### Writer protocol

change REG  
if WR=RR then  
change WR

### Reader protocol

- (1) if WR=RR then return val
- (2) aux := REG
- (3) if WR≠RR then change RR
- (4) val :=REG
- (5) if WR=RR then return val
- (6) val := REG
- (7) return aux



25

## Proof sketch: reading functions

A **reading function**  $\pi$ : for each complete read operation  $r$  (returning  $v$ ),  $\pi(r)$  is a write operation  $w(v)$

Show that for every run of the algorithm, there exists an **atomic** reading function  $\pi$ :

- (A0) No read  $r$  precedes  $\pi(r)$ 
  - ✓ No read returns a value not yet written
- (A1)  $w$  precedes  $r \Rightarrow w = \pi(r)$  or  $w$  precedes  $\pi(r)$ 
  - ✓ No read obtains an overwritten value
- (A2)  $r_1$  precedes  $r_2 \Rightarrow \pi(r_2)$  does not precede  $\pi(r_1)$ 
  - ✓ No new/old inversion

A run is linearizable iff an atomic reading function exists  
(Section 4.3 of the lecture notes)



© 2012 P. Kuznetsov

26

## Proof: constructing $\pi$

- Let  $r$  return a value  $v$
- Let  $\rho_r$  be the read of REG that got the value of  $r$ 
  - ✓ If  $r$  returns in line 7,  $\rho_r$  is the read action in line 2 of  $r$
  - ✓ If  $r$  returns in line 5,  $\rho_r$  is the read action in line 4
  - ✓ If  $r$  returns in line 1,  $\rho_r$  is the read in line 4 or 6 of some previous  $r'$  (depending on how  $r'$  returns)
- Let  $\phi_r$  be the last write action on REG that precedes or is concurrent to  $\rho_r$  and writes the value returned by  $r$  (and  $\rho_r$ )
- Define  $\pi(r)$  as the write operation that contains  $\phi_r$



© 2012 P. Kuznetsov

27

## Proof: show that $\pi$ is atomic

- A0 is easy: by construction of  $\pi$ ,  $\pi(r)$  precedes or is concurrent to  $r$
- A1? A2? Homework

Hint: assume the contrary and come to absurdum



© 2012 P. Kuznetsov

28