



Previously, on CS4410...

Back to Split Binary Semaphores

©




Nurse's office: critical section protecting
variables that determine when to wait


Rooms: waiting conditions

At any time, exactly one
semaphore or thread is green
(and thus, **at most one**
semaphore is green (Invariant))

Back to Split Binary Semaphores



R/W Lock



R/W Lock


Nurse's office: critical section protecting variables that determine when to wait

Rooms: waiting conditions



At any time, exactly one semaphore or thread is green
(and thus, at most one semaphore is green (Invariant))

If n readers in the
critical section,
then $n_{readers} \geq n$

WHY?



R/W Lock



R/W Lock

$n_{readers}$ incremented **inside R/W lock**
before entering the CS (i.e., the database)

Two Types of Monitors

Hoare Monitors



Tony Hoare



Mesa Monitors



Butler Lampson

Different semantics as to what happens when a thread waiting on a condition is alerted that the condition holds

Hoare Monitors

Tony Hoare, 1974

- Syntactic sugar above split binary semaphores
 - **monitor**: one thread can execute at a time
 - **wait(cond. var.)**: thread waits for given condition
 - **signal(cond. var.)**: transfer control to a thread waiting for the given condition, if any

*Similar construct
proposed by
Per Brinch Hansen*



in 1973



Hoare Monitors in Harmony

```
1  import synch
2
3  def Monitor() returns monitor:
4      monitor = synch.Lock()
5
6  def enter(mon):
7      synch.acquire(mon)
8
9  def exit(mon):
10     synch.release(mon)
11
12  def Condition() returns condition:
13     condition = { .sema: synch.BinSema(True), .count: 0 }
14
15  def wait(cond, mon):
16     cond->count += 1
17     exit(mon)
18     synch.acquire(?cond->sema)
19     cond->count -= 1
20
21  def signal(cond, mon):
22     if cond->count > 0:
23         synch.release(?cond->sema)
24         enter(mon)
```

main gate

waiting gate

passes control immediately

a no-op if no one is waiting!

What happens when a thread signals?

- **Hoare semantics:**

- signaling thread is suspended and, atomically, ownership of the lock is passed to one of the waiting threads, whose execution is immediately resumed.
- signaling thread is resumed if former waiter exits monitor, or if it waits again

Producer/Consumer with Bounded Buffer

```
1 import hoare
2
3 def BoundedBuffer(size) returns buffer:
4     buffer = {
5         .mon: hoare.Monitor(),
6         .prod: hoare.Condition(), .cons: hoare.Condition(),
7         .buf: { x:() for x in {1..size} }, circular buffer
8         .head: 1, .tail: 1,
9         .count: 0, .size: size
10    }
11
12 def put(bb, item):
13     enter monitor
14     hoare.enter(?bb→mon)
15     if bb→count == bb→size: wait if full
16         hoare.wait(?bb→prod, ?bb→mon)
17     bb→buf[bb→tail] = item
18     bb→tail = (bb→tail % bb→size) + 1
19     bb→count += 1
20     hoare.signal(?bb→cons, ?bb→mon) signal a consumer
21     hoare.exit(?bb→mon)
22     exit monitor
```


Producer/Consumer with Bounded Buffer

```
1 import hoare
2
3 def BoundedBuffer(size) returns buffer:
4     buffer = {
5         .mon: hoare.Monitor(),
6         .prod: hoare.Condition(), .cons: hoare.Condition(),
7         .buf: { x:() for x in {1..size} },
8         .head: 1, .tail: 1,
9         .count: 0, .size: size
10    }
```

circular buffer

```
11
12 def put(bb, item):
13     hoare.enter(?bb→mon)
14     if bb→count == bb→size:
15         hoare.wait(?bb→prod, ?bb→mon)
16     bb→buf[bb→tail] = item
17     bb→tail = (bb→tail % bb→size) + 1
18     bb→count += 1
19     hoare.signal(?bb→cons, ?bb→mon)
20     hoare.exit(?bb→mon)
```

enter monitor

wait if full

exit monitor

signal() passes the baton immediately if there are waiting consumers

Producer/Consumer with Bounded Buffer

```
22  def get(bb) returns next:  
    hoare.enter(?bb→mon)  
24  if bb→count == 0:  
25      hoare.wait(?bb→cons, ?bb→mon)  
26      next = bb→buf[bb→head]  
27      bb→head = (bb→head % bb→size) + 1  
28      bb→count -= 1  
29      hoare.signal(?bb→prod, ?bb→mon)  
    hoare.exit(?bb→mon)
```

enter monitor

wait if empty

signal a producer

exit monitor

*signal() passes the
baton immediately
if there are waiting
producers*

Mesa Monitors

Mesa Language, Xerox PAak 1980

- Syntactically similar to Hoare monitors

- monitors and condition variables

- Semantically closer to busy waiting

- **wait(cond. var.):** wait for condition, but may get back the CPU when condition is not satisfied (!)
- **notify(cond. var.):** move to ready queue a thread waiting for the condition, if any, **but don't transfer control** (i.e., give the CPU) **to it**
- **notifyAll(cond. var.):** move to ready queue all threads waiting for the condition, **but don't transfer control** (i.e., give the CPU) **to any of them**



What are the implications?

Hoare

- Signaling is atomic with the resumption of waiting thread
 - shared state cannot change before waiting thread is resumed
 - safety requires to signal **only** when condition holds
- Shared state can be checked using an if statement
- Makes it easier to prove liveness
- Tricky to implement

Mesa

- notify() and notifyAll() are **hints**
 - adding them affects performance, never safety
- Shared state **must be checked in a loop** (the condition could have changed since the thread was notified!)
- Simple implementation
- Resilient to **spurious wakeup**

Hoare vs Mesa Monitors

Hoare Monitors	Mesa Monitors
Baton passing approach	If at first you don't succeed... sleep & try again when the stars seem aligned!
<i>signal</i> passes baton	<i>notify(all)</i> moves waiting threads back to ready queue
used by most books	used by most real systems

*Mesa monitors won
the test of time...*

Mesa Monitors in Harmony

```
1 def Condition() returns condition:  
2   condition = bag.empty()  
3
```

*Condition: consists of a
bag of threads waiting*

```
4 def wait(c, lk):  
5   var cnt = 0  
6   let _, ctx = save():  
7     atomically:  
8       cnt = bag.multiplicity(!c, ctx)  
9       !c = bag.add(!c, ctx)  
10      !lk = False  
11      atomically when (not !lk) and (bag.multiplicity(!c, ctx) <= cnt):  
12        !lk = True  
13
```

*wait: unlock+add thread
context to bag of waiters*

```
14 def notify(c):  
15   atomically if !c != bag.empty():  
16     !c = bag.remove(!c, bag.bchoose(!c))  
17
```

*notify: remove one waiter from
the bag of suspended threads*

```
18 def notifyAll(c):  
19   !c = bag.empty()
```

*notifyAll: remove all waiters from
the bag of suspended threads*

Reader/Writer Lock Specification (again)

```
1  def RWlock() returns lock:
2      lock = { .nreaders: 0, .nwriters: 0 }
3
4  def read_acquire(rw):
5      atomically when rw→nwriters == 0:
6          rw→nreaders += 1
7
8  def read_release(rw):
9      atomically rw→nreaders -= 1
10
11 def write_acquire(rw):
12     atomically when (rw→nreaders + rw→nwriters) == 0:
13         rw→nwriters = 1
14
15 def write_release(rw):
16     atomically rw→nwriters = 0
```

Better to assert $rw \rightarrow nreaders > 0$

Reader/Writer lock with Mesa monitors

```
1  from synch import *
2
3  def RWlock() returns lock:
4      lock = {
5          .nreaders: 0, .nwriters: 0, .mutex: Lock(),
6          .r_cond: Condition(), .w_cond: Condition()
7      }
```

*It is the mutex that
protects nreaders and
nwriters, not the R/W lock!*

Invariants

- If n readers in the critical section, then $nreaders \geq n$
- If n writers in the critical section, then $nwriters \geq n$
- $(nreaders \geq 0 \wedge nwriters = 0) \vee (nreaders = 0 \wedge nwriters \leq 1)$

R/W Lock, Reader

```
9  def read_acquire(rw):
10     acquire(rw→mutex)
11     while rw→nwriters > 0:
12         wait(rw→r_cond, rw→mutex)
13     rw→nreaders += 1
14     release(rw→mutex)
15
16  def read_release(rw):
17     acquire(rw→mutex)
18     rw→nreaders -= 1
19     if rw→nreaders == 0:
20         notify(rw→w_cond)
21     release(rw→mutex)
```

*Similar to
Busy Waiting*

but needs this

R/W Lock, Writer

```
23 def write_acquire(rw):
24     acquire(rw→mutex)
25     while (rw→nreaders + rw→nwriters) > 0:
26         wait(rw→w_cond, rw→mutex)
27     rw→nwriters = 1
28     release(rw→mutex)
29
30 def write_release(rw):
31     acquire(rw→mutex)
32     rw→nwriters = 0
33     notifyAll(rw→r_cond)
34     notify(rw→w_cond)
35     release(rw→mutex)
```

*Similar to
Busy Waiting*

*don't forget
anyone!*

Conditional Critical Sections

Let me count the ways...

Busy Waiting	Split Binary Semaphores	Mesa Monitors
use a lock and a loop	use a collection of binary semaphores	use a lock, a collection of condition variables, and a loop
Easy to write the code	Just follow the recipe	Notifying is tricky
Easy to understand the code	Tricky to understand if you don't know the recipe	Easy to understand the code
Ok-ish for true multicore, but bad for virtual threads	Good for virtual threading. Thread only runs when it can make progress	Good for both multicore and virtual threading