


Conditional Waiting



Cornell CIS
COMPUTING AND INFORMATION SCIENCE

Review

- Concurrent Programming is Hard!
 - Non-Determinism
 - Non-Atomicity
- *Critical Sections* simplify things by avoiding data races
 - mutual exclusion
 - progress
- *Need both mutual exclusion and progress!*
- Critical Sections use a *lock*
 - Thread needs lock to enter the critical section
 - Only one thread can get the section's lock

How to get more concurrency?

Idea: allow multiple read-only operations to execute concurrently

- Still no data races
- In many cases, reads are much more frequent than writes

→ reader/writer lock

Either:

- multiple readers, or
- a single writer

thus not:

- *a reader and a writer, nor*
- *multiple writers*

Conditional Waiting

- Thus far we've shown how threads can wait for one another to avoid multiple threads in the critical section
- Sometimes there are other reasons:
 - Wait until queue is non-empty
 - Wait until there are no readers (or writers) in a reader/writer lock
 - ...

Reader/Writer Lock Specification

```
1  def RWlock():
2      result = { .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
```

Reader/Writer Lock Specification


```
1  def RWlock():
2      result = { .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
```

Invariants:

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

R/W Locks: test for mutual exclusion

```
1  import RW
2
3  const NOPS = 3
4
5  rw = RW.RWlock()
6
7  def thread():
8    while choose({ False, True }):
9      if choose({ "read", "write" }) == "read":
10         RW.read_acquire(?rw)
11         rcs: assert (countLabel(rcs) >= 1) and (countLabel(wcs) == 0)
12         RW.read_release(?rw)
13       else: # write
14         RW.write_acquire(?rw)
15         wcs: assert (countLabel(rcs) == 0) and (countLabel(wcs) == 1)
16         RW.write_release(?rw)
17
18  for i in {1..NOPS}:
19    spawn thread()
```



Cheating R/W lock implementation

```
1  import synch
2
3  def RWlock():
4      result = synch.Lock()
5
6  def read_acquire(rw):
7      synch.acquire(rw);
8
9  def read_release(rw):
10     synch.release(rw);
11
12 def write_acquire(rw):
13     synch.acquire(rw);
14
15 def write_release(rw):
16     synch.release(rw);
```

The *lock* protects the application's critical section

Cheating R/W lock implementation

```
1  import synch
2
3  def RWlock():
4      result = synch.Lock()
5
6  def read_acquire(rw):
7      synch.acquire(rw);
8
9  def read_release(rw):
10     synch.release(rw);
11
12 def write_acquire(rw):
13     synch.acquire(rw);
14
15 def write_release(rw):
16     synch.release(rw);
```

Allows only one reader to get the lock at a time

Does *not* have the same behavior as the specification

- it is missing behaviors
- no bad behaviors though

Busy Waiting Implementation

```
1 from synch import Lock, acquire, release
2
3 def RWlock():
4     result = { .lock: Lock(), .nreaders: 0, .nwriters: 0 }
5
6 def read_acquire(rw):
7     acquire(?rw→lock)
8     while rw→nwriters > 0:
9         release(?rw→lock)
10        acquire(?rw→lock)
11        rw→nreaders += 1
12        release(?rw→lock)
13
14 def read_release(rw):
15     acquire(?rw→lock)
16     rw→nreaders -= 1
17     release(?rw→lock)
18
19 def write_acquire(rw):
20     acquire(?rw→lock)
21     while (rw→nreaders + rw→nwriters) > 0:
22         release(?rw→lock)
23         acquire(?rw→lock)
24     rw→nwriters = 1
25     release(?rw→lock)
26
27 def write_release(rw):
28     acquire(?rw→lock)
29     rw→nwriters = 0
30     release(?rw→lock)
```

The *lock* protects *nreaders* and *nwriters*, not the critical section of the application

waiting conditions



Busy Waiting Implementation

```
1  from synch import Lock, acquire, release
2
3  def RWlock():
4      result = { .lock: Lock(), .nreaders: 0, .nwriters: 0 }
5
6  def read_acquire(rw):
7      acquire(?rw→lock)
8      while rw→nwriters > 0:
9          release(?rw→lock)
10         acquire(?rw→lock)
11         rw→nreaders += 1
12         release(?rw→lock)
13
14  def read_release(rw):
15      acquire(?rw→lock)
16      rw→nreaders -= 1
17      release(?rw→lock)
18
19  def write_acquire(rw):
20      acquire(?rw→lock)
21      while (rw→nreaders + rw→nwriters) > 0:
22          release(?rw→lock)
23          acquire(?rw→lock)
24      rw→nwriters = 1
25      release(?rw→lock)
26
27  def write_release(rw):
28      acquire(?rw→lock)
29      rw→nwriters = 0
30      release(?rw→lock)
```

Good: has the same behaviors as the implementation

Bad: process is continuously scheduled to try to get the lock even if it's not available

(Harmony complains about this as well)

Mesa Condition Variables

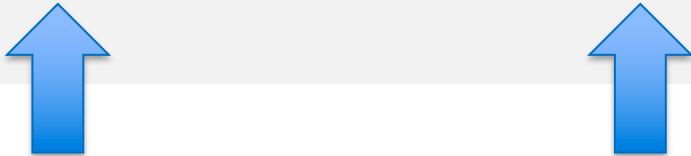
- A lock can have one or more *condition variables*
- A thread that holds the lock but wants to wait for some condition to hold can temporarily release the lock by *waiting* on some condition variable
- Associate a condition variable with each “waiting condition”
 - reader: no writer in the critical section
 - writer: no readers nor writers in the c.s.

Mesa Condition Variables, cont'd

- When a thread that holds the lock notices that some waiting condition is satisfied it should *notify* the corresponding condition variable

R/W lock with *Mesa* condition variables

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



r_cond: used by readers to wait on $nwriters == 0$

w_cond: used by writers to wait on $nreaders == 0 == nwriters$

R/W Lock, reader part

```
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)
```


R/W Lock, reader part

```
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

R/W Lock, reader part

```
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 need this

R/W Lock, reader part

```
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         rw→nreaders -= 1
18         if rw→nreaders == 0:
19             notify(?rw→w_cond)
20             release(?rw→mutex)
```

NEVER EVER USE WAIT WITHOUT AN ENCLOSING WHILE (“naked wait”)

R/W Lock, reader part

compare with busy waiting

```
def read_acquire(rw):  
    acquire(rw→lock)  
    while rw→nwriters > 0:  
        release(rw→lock)  
        acquire(rw→lock)  
    rw→nreaders += 1  
    release(rw→lock)
```

```
def read_release(rw):  
    acquire(rw→lock)  
    rw→nreaders -= 1  
    release(rw→lock)
```

```
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)
```

R/W Lock, reader part

compare with busy waiting

```
def read_acquire(rw):  
    acquire(?rw→lock)  
    while rw→nwriters > 0:  
        release(?rw→lock)  
        acquire(?rw→lock)  
        rw→nreaders += 1  
    release(?rw→lock)
```

```
def read_release(rw):  
    acquire(?rw→lock)  
    rw→nreaders -= 1  
    release(?rw→lock)
```

```
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)
```

R/W Lock, reader part

compare with busy waiting

```
def read_acquire(rw):  
    acquire(rw→lock)  
    while rw→nwriters > 0:  
        release(rw→lock)  
        acquire(rw→lock)  
    rw→nreaders += 1  
    release(rw→lock)
```

```
def read_release(rw):  
    acquire(rw→lock)  
    rw→nreaders -= 1  
    release(rw→lock)
```

```
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)
```

R/W Lock, reader part

compare with busy waiting

```
def read_acquire(rw):  
    acquire(rw→lock)  
    while rw→nwriters > 0:  
        release(rw→lock)  
        acquire(rw→lock)  
    rw→nreaders += 1  
    release(rw→lock)
```

```
def read_release(rw):  
    acquire(rw→lock)  
    rw→nreaders -= 1  
    release(rw→lock)
```

```
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)
```


R/W Lock, writer part

```
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)
```

R/W Lock, writer part

```
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)
```

R/W Lock, writer part

```
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)
```

} don't forget anybody!

R/W Lock, writer part

```
def write_acquire(rw):
    acquire(?rw→lock)
    while (rw→nreaders + rw→nwriters) > 0:
        release(?rw→lock)
        acquire(?rw→lock)
    rw→nwriters = 1
    release(?rw→lock)

def write_release(rw):
    acquire(?rw→lock)
    rw→nwriters = 0
    release(?rw→lock)
```

compare with busy waiting

```
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)
```

Why not use “if” instead of “while” around *wait()*?

- By the time waiter gets the lock back, condition may no longer hold
 - Given three threads, W1, R2, W3
 - W1 enters as a writer
 - R2 waits as a reader
 - W1 leaves, notifying R2
 - W3 enters as a writer
 - R2 wakes up
 - If R2 doesn't check condition again, R2 and W3 would both be in the critical section

Why not use “if” instead of “while” around *wait()*?

- When notifying, be safe rather than sorry
 - it’s better to notify too many threads than too few
 - in case of doubt, use `notifyAll()` instead of just `notify()`
- But this too can lead to some threads waking up when their condition is no longer satisfied

Why not use “if” instead of “while” around *wait()*?

- Because you **should** use while around wait, many condition variable implementation allow “**spurious wakeups**”
- wait() resumes even although condition variable was not notified

Naked waits: just don't do it

Busy Waiting vs Condition Variables

Busy Waiting	Condition Variables
Use a lock and a loop	Use a lock <i>and a collection of condition variables</i> and a loop
Easy to write the code	Notifying is tricky
Easy to understand the code	Easy to understand the code
Progress property is easy	Progress requires careful consideration (both for correctness and efficiency)
Ok-ish for true multi-core, but bad for virtual threads	Good for both multi-core and virtual threading

Mesa Monitors in Harmony

```
1  def Condition():
2      result = bag.empty()
3
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
14  def notify(c):
15      atomically if !c != bag.empty():
16          !c = bag.remove(!c, bag.bchoose(!c))
17
18  def notifyAll(c):
19      !c = bag.empty()
```

Condition: consists of bag of threads waiting

wait: unlock + add thread context to bag of waiters

notify: remove one waiter from the bag of suspended threads

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