

# Review

- Concurrent Programming is Hard!
  - Non-Determinism
  - Non-Atomicity
- *Critical Sections* simplify things
  - 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

# Specification in the face of Concurrency and Overlap

Is the following a possible scenario?

1. customer X orders a burger
2. customer Y orders a burger (afterwards)
3. customer Y is served a burger
4. customer X is served a burger (afterwards)

# Specification in the face of Concurrency and Overlap

Is the following a possible scenario?

1. customer X orders a burger
2. customer Y orders a burger (afterwards)
3. customer Y is served a burger
4. customer X is served a burger (afterwards)


We've all seen this happen. It's a matter of how things get scheduled!

# Specification

- One operation: order a burger
  - result: a burger (at some later time)
- Semantics: the burger manifests itself atomically *sometime during the operation*
- *Atomically: no two manifestations overlap*
- It's easier to specify something when you don't have to worry about overlap
  - i.e., you can simply give a sequential specification
- Allows many implementations


# Implementation?

- Suppose the diner has one small hot plate and two cooks
- Cooks use a lock for access to the hot plate
- Possible scenario:
  1. customer X orders burger, order ends up with cook 1
  2. customer Y orders burger, order ends up with cook 2
  3. cook 1 was busy with something else, so cook 2 grabs the lock first
  4. cook 2 cooks burger for Y
  5. cook 2 releases lock
  6. cook 1 grabs lock
  7. cook 1 cooks burger for X
  8. cook 1 releases lock
  9. customer Y receives burger
  10. customer X receives burger



# Implementation?

- Suppose the diner has one small hot plate and two cooks
- Cooks use a lock for access to the hot plate
- Possible scenario:
  1. customer X orders burger, order ends up with cook 1
  2. customer Y orders burger, order ends up with cook 2
  3. cook 1 was busy with something else, so cook 2 grabs the lock first
  4. cook 2 cooks burger for Y
  5. cook 2 releases lock
  6. cook 1 grabs lock
  7. cook 1 cooks burger for X
  8. cook 1 releases lock
  9. customer Y receives burger
  10. customer X receives burger



- *can't happen if Y orders burger after X receives burger*
- *but if operations overlap, any ordering can happen...*

# Queue test program, again

```
1  import queue
2
3  const NOPS = 4
4  q = queue.Queue()
5
6  def put_test(self):
7      print("call put", self)
8      queue.put(?q, self)
9      print("done put", self)
10
11  def get_test(self):
12      print("call get", self)
13      let v = queue.get(?q):
14          print("done get", self, v)
15
16  nputs = choose {1..NOPS-1}
17  for i in {1..nputs}:
18      spawn put_test(i)
19  for i in {1..NOPS-nputs}:
20      spawn get_test(i)
```

# How to get more concurrency?

**Idea:** allow multiple read-only operations to execute concurrently

- 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*




# 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
```

# 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);
```

# 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)
```

# 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)*



# Conditional Waiting



**Cornell CIS**  
COMPUTING AND INFORMATION SCIENCE

# 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
  - ...



# Busy Waiting: not a good way

- Wait until queue is non-empty:

*done* = **False**

**while not** *done*:

*next* = queue.get(*q*)

*done* = *next* != **None**

# Busy Waiting: not a good way

- Wait until queue is non-empty:

*done* = **False**

**while not** *done*:

*next* = queue.get(*q*)

*done* = *next* != **None**

- *wastes CPU cycles*
- *creates unnecessary contention*

# Enter *binary semaphores*



[Dijkstra 1962]

# Binary Semaphore

- Boolean variable (much like a lock)
- Three operations:
  - *binsema* = BinSema(False or True)
    - initialize *binsema*
  - *acquire(?binsema)*
    - waits until *!binsema* = False, then sets *!binsema* to True.
  - *release(?binsema)*
    - set *!binsema* to False
    - can only be called if *!binsema* = True

# Dijkstra was Dutch, like some

- He said ***P**robeer-te-verlagen* instead of acquire
- He said ***V**erhogen* instead of release
- Many people still use P/V when talking about semaphore operators
- Easier to remember:
  - ***P***rocurer (acquire)
  - ***V***acate (release)

# Difference with locks

Locks	Binary Semaphores
Initially “unlocked” (False)	Can be initialized to False or True
<i>Acquired</i> , usually <i>released</i> by same thread	Can be <i>acquired</i> and <i>released</i> by different threads
Mostly used to implement critical sections	Can be used to implement critical sections as well as waiting for special conditions

but both are much like “*batons*” that are being passed

# Binary Semaphore specification

```
def BinSema(acquired):  
    result = acquired  
  
def Lock():  
    result = BinSema(False)  
  
def acquired(binsema):  
    result = !binsema  
  
def acquire(binsema):  
    atomically when not !binsema:  
        !binsema = True  
  
def release(binsema):  
    assert !binsema  
    atomically !binsema = False
```

# Waiting with semaphores

```
import synch;
```

```
condition = BinSema(True)
```

```
def T0():
```

```
    acquire(?condition)    # wait for signal
```

```
def T1():
```

```
    release(?condition)    # send signal
```

```
spawn T0()
```

```
spawn T1()
```



# Waiting with semaphores

```
import synch;
```

```
condition = BinSema(True)
```

```
def T0():
```

```
    acquire(?condition)    # wait for signal
```

```
def T1():
```

```
    release(?condition)    # send signal
```

```
spawn T0()
```

```
spawn T1()
```

What happens if T0 runs first?  
What happens if T1 runs first?

# Semaphores can be locks too

- `lk = BinSema(False)` # False-initialized
- `acquire(?lk)` # grab lock
- `release(?lk)` # release lock

Great, what else can one do with binary semaphores??

# Conditional Critical Sections

- A critical section with a condition
- For example:
  - `queue.get()`, but wait until the queue is non-empty
    - don't want two threads to run code at the same time, but also don't want any thread to run `queue.get()` code when queue is empty
  - `print()`, but wait until the printer is idle
  - `RW.read_acquire()`, but only if there are no writers in the critical section
  - allocate 100 GPUs, when they become available
  - ...

[Hoare 1973]

# Multiple conditions

Some conditional critical sections can have multiple conditions:

- R/W lock: readers are waiting for writer to leave; writers are waiting for reader or writer to leave
- bounded queue: dequeuers are waiting for queue to be non-empty; enqueueers are waiting for queue to be non-full
- ...

# High-level idea: selective baton passing!

- When a thread wants to execute in the critical section, it needs the one baton
- Threads can be waiting for various conditions
  - such threads do not hold the baton
- When a thread with the baton leaves the critical section, it checks to see if there are threads waiting on a condition that now holds
- If so, it passes the baton to one such thread
- If not, the critical section is vacated, and the baton is free to pick up for another thread that comes along

# “Split Binary Semaphores” [Hoare 1973]

- Implement baton passing with multiple binary semaphores
- If there are  $N$  conditions, you'll need  $N+1$  binary semaphores
  - one for each condition
  - one to enter the critical section in the first place
- **Invariant: At most one of these semaphores is released (False)**
  - If all are acquired (True), baton held by some thread
  - If one semaphore is released, no thread holds the baton
    - if it's the “entry” semaphore, then no thread is waiting on a condition that holds, and any thread can enter
    - if it's one of the condition semaphores, some thread that is waiting on the condition can now enter the critical section

# “Split Binary Semaphores” [Hoare 1973]

- Implement baton passing with multiple binary semaphores
- If there are  $N$  conditions, you'll need  $N+1$  binary semaphores
  - one for each condition
  - one to enter the critical section in the first place
- **Invariant: At most one of these semaphores is released (False)**
  - If all are acquired (True), baton held by some thread
  - If one semaphore is released, no thread holds the baton
    - if it's the “entry” semaphore, then no thread is waiting on a condition that holds, and any thread can enter
    - if it's one of the condition semaphores, some thread that is waiting on the condition can now enter the critical section
      - at most one




# “Split Binary Semaphores” [Hoare 1973]

- Implement baton passing with multiple binary semaphores
- If there are  $N$  conditions, you'll need  $N+1$  binary semaphores
  - one for each condition
  - one to enter the critical section in the first place
- **Invariant: At most one of these semaphores is released (False)**
  - If all are acquired (True), baton held by some thread
  - If one semaphore is released, no thread holds the baton
    - if it's the “entry” semaphore, then no thread is waiting on a condition that holds, and any thread can enter
    - if it's one of the condition semaphores, some thread that is waiting on the condition can now enter the critical section
      - at most one
      - at least one

# Bathroom humor...

- holds baton
- does not hold baton

3 threads want to enter critical section




Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

- holds baton
- does not hold baton

3 threads want to enter critical section



semaphore = False

Bathroom: critical section


Bedrooms: waiting conditions

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

# Bathroom humor...

- holds baton
- does not hold baton

3 threads want to enter critical section




Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

- holds baton
- does not hold baton

3 threads want to enter critical section



at any time exactly one  
semaphore or thread is green

Bathroom: critical section  
Bedrooms: waiting conditions


# This is a model of:

- Reader/writer lock:
  - Bathroom: critical section
  - Bedroom 1: readers waiting for writer to leave
  - Bedroom 2: writers waiting for readers or writers to leave
- Bounded queue:
  - Bathroom: critical section
  - Bedroom 1: dequeuers waiting for queue to be non-empty
  - Bedroom 2: enqueueers waiting for queue to be non-full
- ...

# Bathroom humor...

- holds baton
- does not hold baton



3 threads want to enter critical section




Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

-  holds baton
-  does not hold baton

1 thread entered the critical section



Bathroom: critical section  
Bedrooms: waiting conditions


at any time exactly one  
semaphore or thread is green



# Bathroom humor...

- holds baton
- does not hold baton

thread needs to wait for Condition 1




Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

- holds baton
- does not hold baton

no thread waiting for condition that holds



Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...




holds baton



does not hold baton

another thread can enter the critical section




Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

- holds baton
- does not hold baton

thread entered the critical section



Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...




holds baton



does not hold baton

thread enables Condition 1 and wants to leave



Bathroom: critical section


Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

- holds baton
- does not hold baton

thread left, Condition 1 holds



Bathroom: critical section  
Bedrooms: waiting conditions


at any time exactly one  
semaphore or thread is green

# Bathroom humor...

■ holds baton

■ does not hold baton

first thread (and only first thread) can enter critical section again



Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...




holds baton



does not hold baton

first thread entered critical section again



Bathroom: critical section  
Bedrooms: waiting conditions


at any time exactly one  
semaphore or thread is green



# Bathroom humor...

- holds baton
- does not hold baton

first thread leaves




Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

- holds baton
- does not hold baton

first thread done



Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...




holds baton



does not hold baton

one thread wants to enter the critical section




Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

- holds baton
- does not hold baton

last thread entered critical section




Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

- holds baton
- does not hold baton



thread needs to wait for Condition 2




Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Bathroom humor...

-  holds baton
-  does not hold baton

thread waiting for Condition 2



Bathroom: critical section  
Bedrooms: waiting conditions

at any time exactly one  
semaphore or thread is green

# Let's build a Reader/Writer lock this way

- You may have seen other ways
- There are many ways that lead to Rome



from <https://www.r-bloggers.com/2019/07/all-roads-lead-to-rome-2/>

# Reader/Writer Lock Spec, again

```
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: implementation

```
1  from synch import BinSema, acquire, release
2
3  def RWlock():
4      result = {
5          .nreaders: 0, .nwriters: 0, .mutex: BinSema(False),
6          .r_gate: { .sema: BinSema(True), .count: 0 },
7          .w_gate: { .sema: BinSema(True), .count: 0 }
8      }
```

## Accounting:

- *nreaders*: #readers in the critical section
- *r\_gate.count*: #readers waiting to enter the critical section
- *nwriters*: #writers in the critical section
- *w\_gate.count*: #writers waiting to enter the critical section

## Invariants:

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

# Reader/writer lock: read

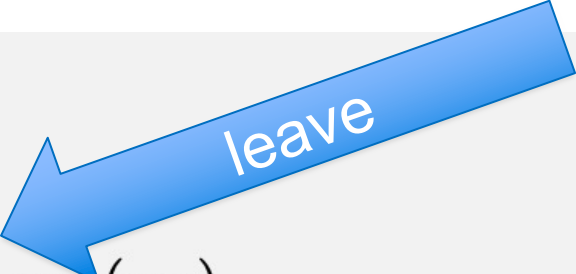
```
18     def read_acquire(rw):
19         acquire(?rw→mutex)
20         if rw→nwriters > 0:
21             rw→r_gate.count += 1; release_one(rw)
22             acquire(?rw→r_gate.sema); rw→r_gate.count -= 1
23         rw→nreaders += 1
24         release_one(rw)
25
26     def read_release(rw):
27         acquire(?rw→mutex); rw→nreaders -= 1; release_one(rw)
```

# Reader/writer lock: read

```
18  def read_acquire(rw):  
19      acquire(?rw→mutex)  ← enter main gate  
20      if rw→nwriters > 0:  
21          rw→r_gate.count += 1; release_one(rw)  
22          acquire(?rw→r_gate.sema); rw→r_gate.count -= 1  
23      rw→nreaders += 1  
24      release_one(rw)  
25  
26  def read_release(rw):  
27      acquire(?rw→mutex); rw→nreaders -= 1; release_one(rw)
```

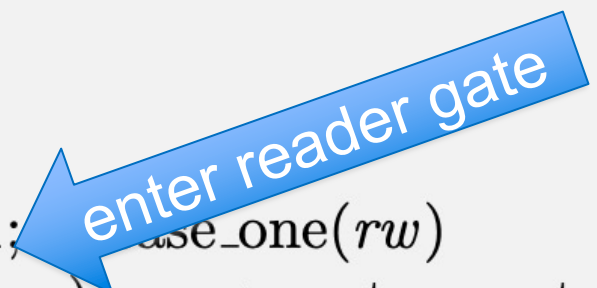
# Reader/writer lock: read

```
18  def read_acquire(rw):
19      acquire(?rw→mutex)
20      if rw→nwriters > 0:
21          rw→r_gate.count += 1; release_one(rw)
22          acquire(?rw→r_gate.sema); rw→r_gate.count -= 1
23      rw→nreaders += 1
24      release_one(rw)
25
26  def read_release(rw):
27      acquire(?rw→mutex); rw→nreaders -= 1; release_one(rw)
```



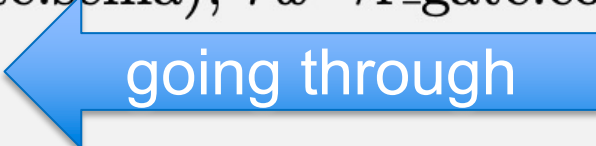
# Reader/writer lock: read

```
18  def read_acquire(rw):
19      acquire(?rw→mutex)
20      if rw→nwriters > 0:
21          rw→r_gate.count += 1; release_one(rw)
22          acquire(?rw→r_gate.sema); rw→r_gate.count -= 1
23      rw→nreaders += 1
24      release_one(rw)
25
26  def read_release(rw):
27      acquire(?rw→mutex); rw→nreaders -= 1; release_one(rw)
```




# Reader/writer lock: read

```
18  def read_acquire(rw):
19      acquire(?rw→mutex)
20      if rw→nwriters > 0:
21          rw→r_gate.count += 1; release_one(rw)
22          acquire(?rw→r_gate.sema); rw→r_gate.count -= 1
23      rw→nreaders += 1
24      release_one(rw)
25
26  def read_release(rw):
27      acquire(?rw→mutex); rw→nreaders -= 1; release_one(rw)
```



# Reader/writer lock: read

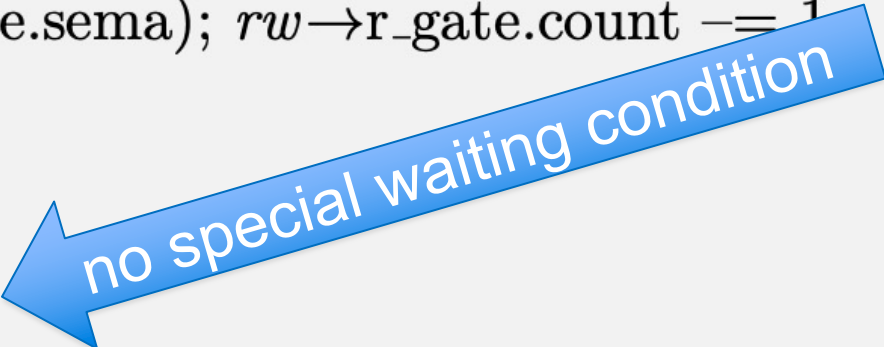
```
18  def read_acquire(rw):
19      acquire(?rw→mutex)
20      if rw→nwriters > 0:
21          rw→r_gate.count += 1; release_one(rw)
22          acquire(?rw→r_gate.sema); rw→r_gate.count -= 1
23      rw→nreaders += 1
24      release_one(rw)
25
26  def read_release(rw):
27      acquire(?rw→mutex); rw→nreaders -= 1; release_one(rw)
```



leave: let others try too

# Reader/writer lock: read

```
18  def read_acquire(rw):
19      acquire(?rw→mutex)
20      if rw→nwriters > 0:
21          rw→r_gate.count += 1; release_one(rw)
22          acquire(?rw→r_gate.sema); rw→r_gate.count -= 1
23      rw→nreaders += 1
24      release_one(rw)
25
26  def read_release(rw):
27      acquire(?rw→mutex); rw→nreaders -= 1; release_one(rw)
```



no special waiting condition



# Reader/writer lock: read

```
18  def read_acquire(rw):
19      acquire(?rw→mutex)
20      if rw→nwriters > 0:
21          rw→r_gate.count += 1; release_one(rw)
22          acquire(?rw→r_gate.sema); rw→r_gate.count -= 1
23      rw→nreaders += 1
24      release_one(rw)
25
26  def read_release(rw):
27      acquire(?rw→mutex); rw→nreaders -= 1; release_one(rw)
```

Note that acquire/release  
operations alternate

# Reader/writer lock: write

```
29  def write_acquire(rw):
30      acquire(?rw→mutex)
31      if (rw→nreaders + rw→nwriters) > 0:
32          rw→w_gate.count += 1; release_one(rw)
33          acquire(?rw→w_gate.sema); rw→w_gate.count -= 1
34      rw→nwriters += 1
35      release_one(rw)
36
37  def write_release(rw):
38      acquire(?rw→mutex); rw→nwriters -= 1; release_one(rw)
```

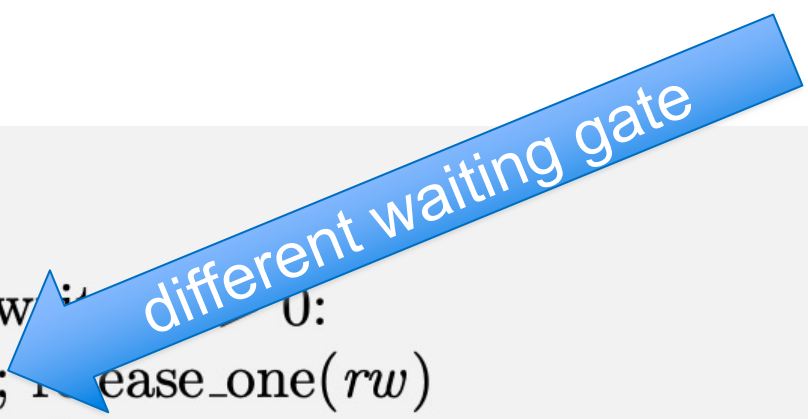
# Reader/writer lock: write

different waiting condition

```
29  def write_acquire(rw):
30      acquire(?rw→mutex)
31      if (rw→nreaders + rw→nwriters) > 0:
32          rw→w_gate.count += 1; release_one(rw)
33          acquire(?rw→w_gate.sema); rw→w_gate.count -= 1
34      rw→nwriters += 1
35      release_one(rw)
36
37  def write_release(rw):
38      acquire(?rw→mutex); rw→nwriters -= 1; release_one(rw)
```

# Reader/writer lock: write

```
29  def write_acquire(rw):
30      acquire(?rw→mutex)
31      if (rw→nreaders + rw→nwriters == 0):
32          rw→w_gate.count += 1; release_one(rw)
33          acquire(?rw→w_gate.sema); rw→w_gate.count -= 1
34          rw→nwriters += 1
35          release_one(rw)
36
37  def write_release(rw):
38      acquire(?rw→mutex); rw→nwriters -= 1; release_one(rw)
```



different waiting gate

# Reader/writer lock: leaving

```
10  def release_one(rw):
11      if (rw→nwriters == 0) and (rw→r_gate.count > 0):
12          release(?rw→r_gate.sema)
13      elif ((rw→nreaders + rw→nwriters) == 0) and (rw→w_gate.count > 0):
14          release(?rw→w_gate.sema)
15      else:
16          release(?rw→mutex)
```

# Reader/writer lock: leaving

```
10  def release_one(rw):
11      if (rw→nwriters == 0) and (rw→r_gate.count > 0):
12          release(?rw→r_gate.sema)
13      elif ((rw→nreaders + rw→nwriters) == 0) and (rw→w_gate.count > 0):
14          release(?rw→w_gate.sema)
15      else:
16          release(?rw→mutex)
```

## When leaving critical section:

- if no writers in the Critical Section and there are readers waiting  
then let a reader in
- else if no readers nor writer in the C.S. and there are writers waiting  
then let a writer in
- otherwise  
let any new thread in

# Reader/writer lock: leaving

```
10  def release_one(rw):
11      if (rw→nwriters == 0) and (rw→r_gate.count > 0):
12          release(?rw→r_gate.sema)
13      elif ((rw→nreaders + rw→nwriters) == 0) and (rw→w_gate.count > 0):
14          release(?rw→w_gate.sema)
15      else:
16          release(?rw→mutex)
```

## When leaving critical section:

- if no writers in the Critical Section and there are readers waiting  
then let a reader in
- else if no readers nor writer in the C.S. and there are writers waiting  
then let a writer in
- otherwise
  - Can the two conditions be reversed?
  - What is the effect of that?
- let any new thread in

# What happens if...

- Multiple readers are waiting and a writer leaves
- Does it let in all the readers or just one?



# Layers of Abstraction

- Note that we have two layers of abstraction:
  - The reader/writer lock object
  - The binary semaphore object
- Both can be used to implement critical sections:
  - R/W locks allow multiple readers in a critical section
  - split binary semaphores allow only one thread at a time in a critical section
- These are *not the same critical sections*
  - they occur at different levels of abstraction

# Another example: lockbox



- to enter house, you need the key
- to get the key out of the lockbox, you need the code
- the house and the lockbox are both critical sections
- **to enter the house you:**
  1. open the lockbox
  2. open the house with the key
  3. put the key back in the lockbox and close it
- **to lock the house you:**
  1. open the lockbox
  2. get the key and lock the house
  3. put the key back in the lockbox and close it

# Why is this useful?



- Because it implements an interesting rule:
  - multiple people can get into the house
  - but only if they have lockbox access
- Could design fancier rules, for example:
  - put three marbles in the lockbox
  - to enter the house, you have to remove a marble and take it with you
  - when leaving the house, you have to put the marble back in
- *What does that accomplish?*

# Same with R/W locks

- R/W lock:
  - **key to the house**
  - house allows one writer *or* multiple readers
    - but not both
- Split Binary Semaphore:
  - **lockbox**
  - + 1 marble (taken by writer)
  - + 1 (tiny) abacus (updated by readers)

# Making R/W lock starvation-free

- Last implementation suffers from starvation

# Making R/W lock starvation-free

- Last implementation suffers from starvation
  - steady stream of new readers lock out writers

# Making R/W lock starvation-free

- **change the waiting and release conditions:**
  - when a reader tries to enter the critical section, wait if there is a writer in the critical section **OR** if there are writers waiting to enter the critical section
  - exiting reader prioritizes releasing a waiting writer
  - exiting writer prioritizes releasing a waiting reader

See Harmony book

# Conditional Critical Sections

We now know of two ways to implement them:

Busy Waiting	Split Binary Semaphores
Wait for condition in loop, acquiring lock before testing condition and releasing it if the condition does not hold	Use a collection of binary semaphores and keep track of state including information about waiting threads
Easy to understand the code	State tracking is complicated
Ok-ish for true multi-core, but bad for virtual threads	Good for both multi-core and virtual threading



# Language support?

- Can't the programming language be more helpful here?
  - Helpful syntax
  - Or at least some library support

# “Hoare” Monitors

- Tony Hoare 1974
  - similar construct given by Per Brinch-Hansen 1973
- Syntactic sugar around split binary semaphores

```
single resource:monitor
begin busy: Boolean;
      nonbusy: condition;
      procedure acquire;
      begin if busy then nonbusy.wait;
            busy := true
      end;
      procedure release;
      begin busy := false;
            nonbusy.signal
      end;
      busy := false; comment initial value;
end single resource
```

“condition variable”

wait method


signal method

# Hoare Monitors in Harmony

```
1      import synch
2
3      def Monitor():
4          result = synch.Lock()
5
6      def enter(mon):
7          synch.acquire(mon)
8
9      def exit(mon):
10         synch.release(mon)
11
12     def Condition():
13         result = { .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)
```


# Hoare Monitors in Harmony

```
1      import synch
2
3      def Monitor():
4          result = synch.Lock()
5
6      def enter(mon):
7          synch.acquire(mon)
8
9      def exit(mon):
10         synch.release(mon)
11
12     def Condition():
13         result = { .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)
```



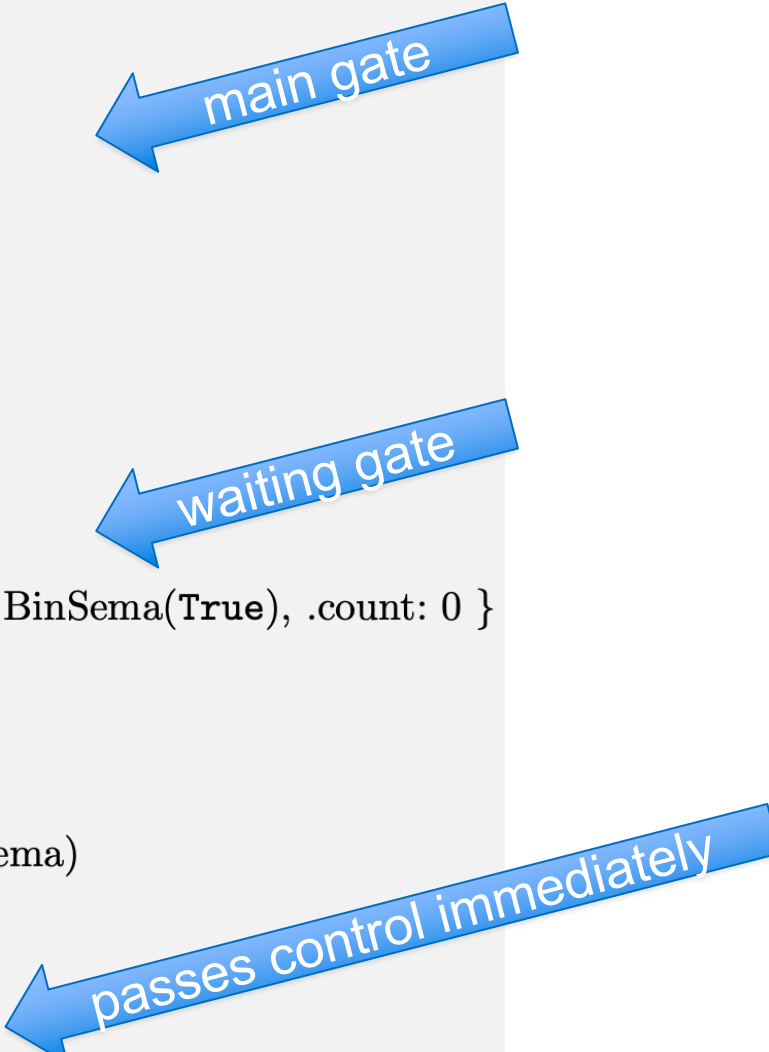
# Hoare Monitors in Harmony

```
1      import synch
2
3      def Monitor():
4          result = synch.Lock()
5
6      def enter(mon):
7          synch.acquire(mon)
8
9      def exit(mon):
10         synch.release(mon)
11
12     def Condition():
13         result = { .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)
```



# Hoare Monitors in Harmony

```
1      import synch
2
3      def Monitor():
4          result = synch.Lock()
5
6      def enter(mon):
7          synch.acquire(mon)
8
9      def exit(mon):
10         synch.release(mon)
11
12     def Condition():
13         result = { .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

# Example: bounded buffer (aka producer/consumer)

```
1      import hoare
2
3      def BB(size):
4          result = {
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         }
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)
```

# Example: bounded buffer (aka producer/consumer)

```
1  import hoare
2
3  def BB(size):
4      result = {
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     }
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)
```

← N+1 semaphores abstracted away

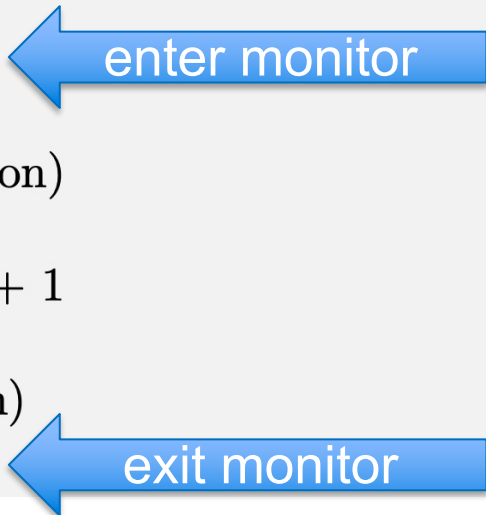


# Example: bounded buffer (aka producer/consumer)

```
1  import hoare
2
3  def BB(size):
4      result = {
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      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)
```

# Example: bounded buffer (aka producer/consumer)

```
1  import hoare
2
3  def BB(size):
4      result = {
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     }
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)
```



The diagram illustrates the execution of the `put` function within the bounded buffer. A blue arrow labeled "enter monitor" points to the `hoare.enter(?bb→mon)` line (line 13). Another blue arrow labeled "exit monitor" points to the `hoare.exit(?bb→mon)` line (line 20).

# Example: bounded buffer (aka producer/consumer)

```
1  import hoare
2
3  def BB(size):
4      result = {
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     }
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)
```



wait if full



signal a consumer

# Example: bounded buffer (aka producer/consumer)

```
1  import hoare
2
3  def BB(size):
4      result = {
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     }
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)
```

signal() passes baton  
*immediately* if there  
are threads waiting on  
the given condition  
variable

# Hoare Monitors

- Split Binary Semaphores underneath the “monitor” programming language paradigm
  - **monitor**: one thread can execute at a time
  - **wait(condition variable)**: thread waits for given condition
  - **signal(condition variable)**: transfer control to a thread waiting for the given condition, if any

# Mesa Monitors

- Introduced in the Mesa language
  - Xerox PARC, 1980
- Syntactically similar to Hoare monitors
  - monitors and condition variables
- *Semantically closer to busy waiting approach*
  - **wait(condition variable)**: wait for condition, *but may wake up before condition is not satisfied*
  - **notify(condition variable)**: wake up a thread waiting for the condition, if any, *but don't transfer control*
  - **notifyAll(condition variable)**: wake up all threads waiting for the condition, *but don't transfer control*

*This is hugely different from Hoare monitors*

# Hoare vs Mesa Monitors

Hoare monitors	Mesa monitors
Baton passing approach	Sleep + try again
signal passes baton	notify(all) wakes sleepers

Mesa monitors won the test of time...

# 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



# 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
```

# 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)
```

# R/W lock with Mesa monitors

```
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          }
```

## Invariants:

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

*mutex* protects the *nreaders/nwriters* variables, not the R/W critical section!

# 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, 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!

# Conditional Critical Sections

We now know of *three* ways to implement them:

Busy Waiting	Split Binary Semaphores	Mesa Monitors
Use a lock and a loop	Use a collection of binary semaphores	Use a lock and 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 recipe	Easy to understand the code
Ok-ish for true multi-core, but bad for virtual threads	Good for virtual threading. Thread only runs when it can make progress	Good for both multi-core and virtual threading (but not optimal)