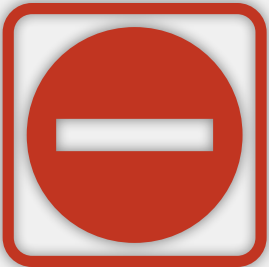


Testing a Concurrent Queue?

```
1  import queue
2
3  def sender(q, v):
4      queue.put(q, v)
5
6  def receiver(q):
7      let v = queue.get(q):
8          assert v in { None, 1, 2 }
9
10 demoq = queue.Queue()
11 spawn sender(?demoq, 1)
12 spawn sender(?demoq, 2)
13 spawn receiver(?demoq)
14 spawn receiver(?demoq)
```



Ad hoc

Unsystematic

Systematic Testing

• Sequential case:

- Try all sequences consisting of 1 operation
 - ▶ put or get
- Try all sequences consisting of 2 operations
 - ▶ put+put, put+get, get+put, get+get
- Try all sequences consisting of 3 operations
- ...

How do we know if a sequence is correct?

- We run the test program against both the sequential specification and the implementation
- We check whether running the test program against the implementation produces the behaviors (e.g., returns the same values) as running it against the sequential specification

Systematic Testing

• Concurrent case:

- Can't run same sequence of operations on both
 - ▶ even if both are correct, nondeterminism of concurrency may have the two runs produce different results
- Instead:
 - ▶ Try all interleavings of 1 operation
 - ▶ Try all interleavings in a sequence of 2 ops
 - ▶ Try all interleavings in a sequence of 3 ops
 - ▶ ...

How do we know if an interleaving is correct?

- We run the test program against both the concurrent specification and the implementation
 - this produces two DFAs, which capture all possible behaviors of the program
- We then verify whether the DFA produced running against the specification is the same as the one produced running against the implementation

Queue test program

```
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)
```

** always at least one
put and one get*

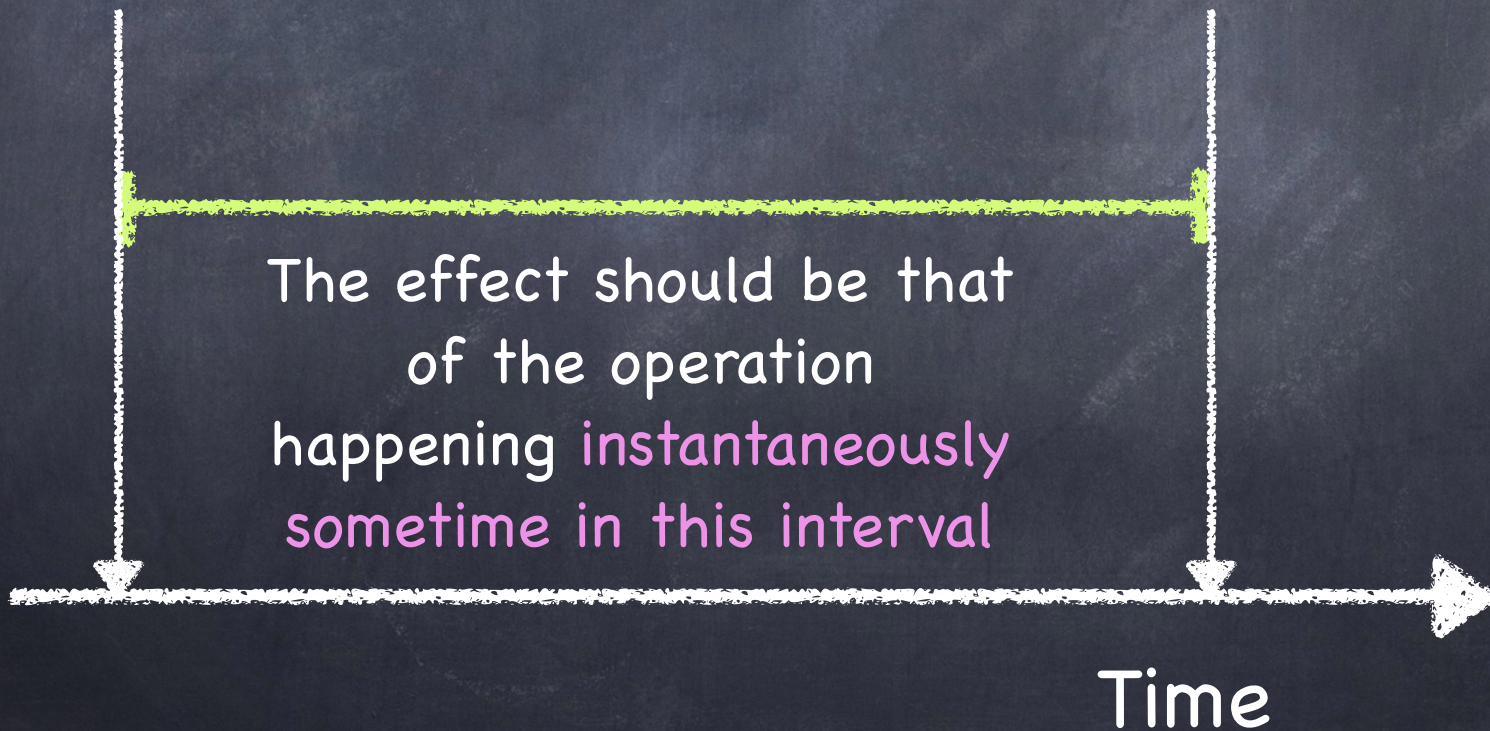
*NOPS threads,
nondeterministically
choosing* to execute
put or get*

But which behaviors
of the implementation
are **correct**?

Life of an Atomic Operation

process invokes
operation

process
continues




Life of an Atomic Operation

operation
happens
atomically

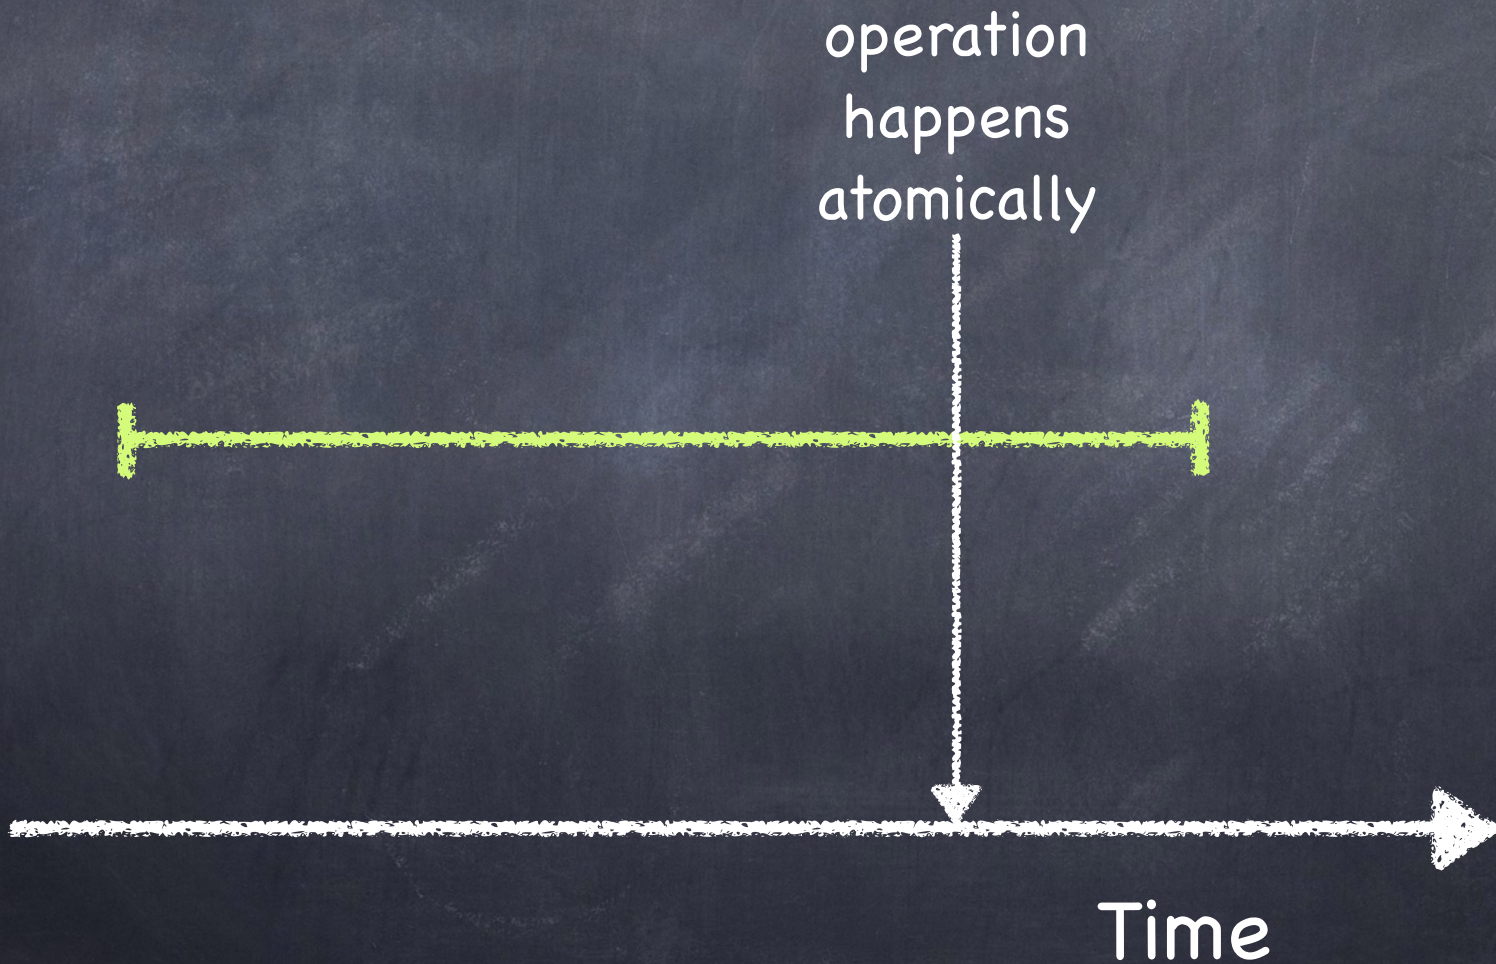


Time

Life of an Atomic Operation



Life of an Atomic Operation



Correct Behaviors

Suppose the queue is initially empty

put (3)

get () ← 3




Time

Correct Behaviors

Suppose the queue is initially empty

put (3)

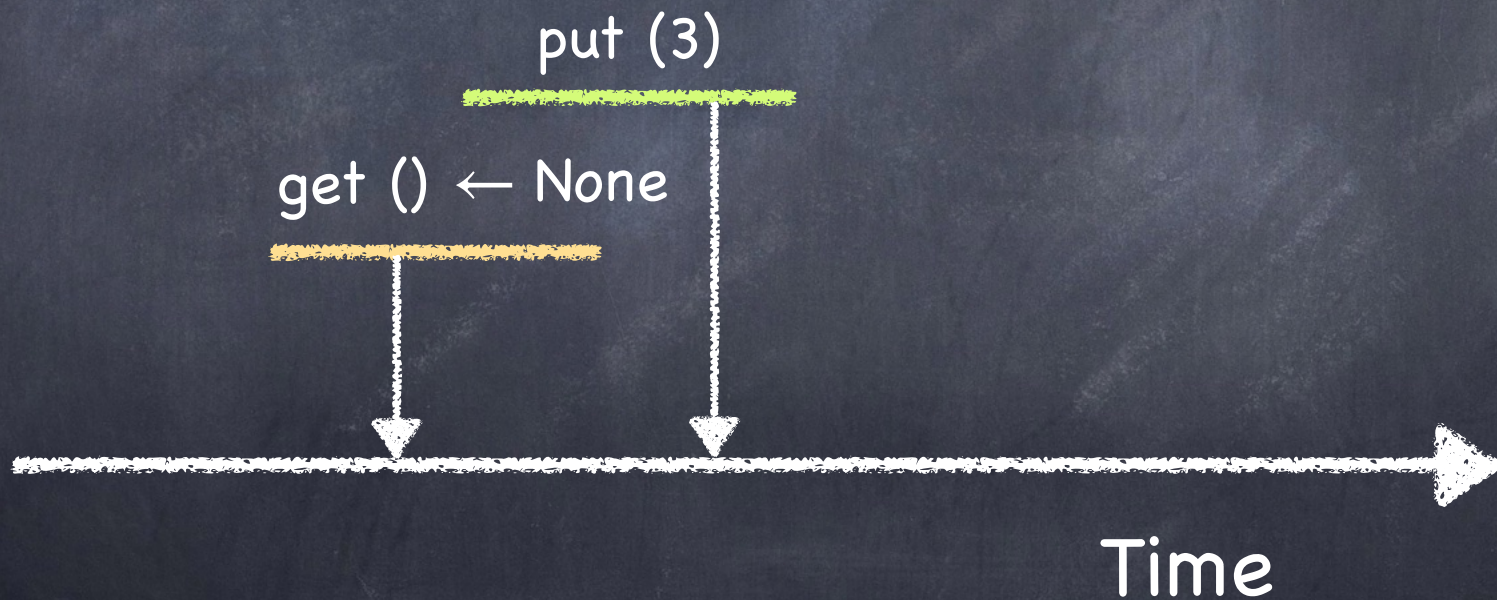
get () ← None



Time


Correct Behaviors

Suppose the queue is initially empty

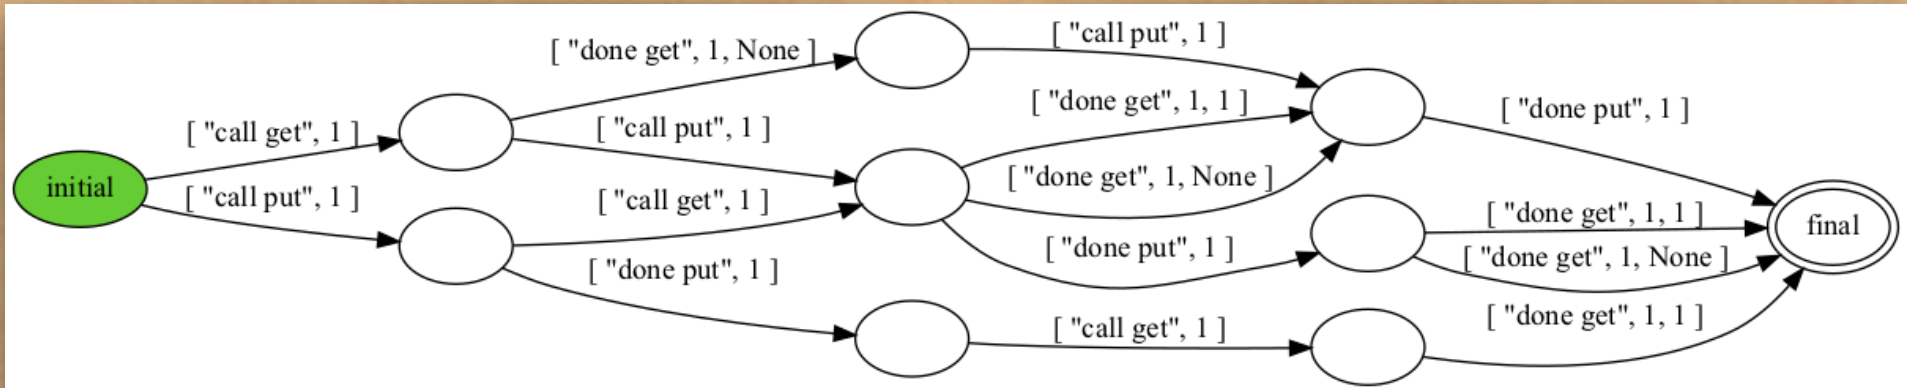


Correct Behaviors

Suppose the queue is initially empty



Queue test program



```
$ harmony -c NOPS=2 -o spec.png code/qtestpar.hny
```



Testing: comparing behaviors

```
$ harmony -o queue4.hfa code/qttestpar.hny  
$ harmony -B queue4.hfa -m queue=queueconc code/qttestpar.hny
```

- The first command outputs the behavior of the running test program against the specification in file queue4.hfa
- The second command runs the test program against the implementation and checks if its behavior matches that stored in queue4.hfa

Review

- Concurrent programming is hard!
 - Non-Determinism
 - Non-Atomicity
- Critical Sections simplify things
 - mutual exclusion
 - progress
- Critical Sections use a lock
 - Threads need lock to enter the CS
 - Only one thread can get the section's lock



Readers-Writers



• Models access to an object (e.g., a database), shared among several threads

□ some threads only read the object

□ others only write it

• Safety

$$(\#r \geq 0) \wedge (0 \leq \#w \leq 1) \wedge ((\#r > 0) \Rightarrow (\#w = 0))$$

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
 - at most one writer, and, if no writer, any number of readers

$$(\#r \geq 0) \wedge (0 \leq \#w \leq 1) \wedge ((\#r > 0) \Rightarrow (\#w = 0))$$

Reader/Writer Lock Specification

```
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
```


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 In CS rcs: assert (countLabel(rcs) >= 1) and (countLabel(wcs) == 0)
12            RW.read_release(?rw)
13         else: # write
14            RW.write_acquire(?rw)
15 In CS wcs: assert (countLabel(rcs) == 0) and (countLabel(wcs) == 1)
16            RW.write_release(?rw)
17
18 for i in {1..NOPS}:
19     spawn thread()
```

Multiple Readers

No Writer

1 Writer and No Readers

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);
```

Only 1
Reader gets
a lock at a
time!

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);
```

*Only 1
Reader gets
a lock at a
time!*

*It is
missing
behaviors
allowed by the
specification!*

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);
```

*Only 1
Reader gets
a lock at a
time!*

*It is
missing
behaviors
allowed by the
specification*

*But, at least,
no bad behavior!*

Busy-Waiting Implementation

```
1 from synch import Lock, acquire, release
2
3 def RWlock() returns lock:
4     lock = { .lock: Lock(), .nreaders: 0, .nwriters: 0 }
5
6 def read_acquire(rw):
7     acquire(?rw→lock)
8     while rw→nwriters > 0: } Busy
9         release(?rw→lock) } waiting
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)
```

Acquire the lock
Test the condition
Release the lock

Repeat

The lock
protects nreaders
and nwriters,
not the RW
critical section!

It has the same
behaviors as the
implementation!

Busy-Waiting Implementation

```
1 from synch import Lock, acquire, release
2
3 def RWlock() returns lock:
4     lock = { .lock: Lock(), .nreaders: 0, .nwriters: 0 }
5
6 def read_acquire(rw):
7     acquire(?rw→lock)
8     while rw→nwriters > 0: } Busy
9         release(?rw→lock) } waiting
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)
```

It has the same behaviors as the implementation!

Wasteful!

Process continuously scheduled to try to get the lock even if it is not available

Conditional Waiting

Conditional Waiting

- Threads wait for each other to prevent multiple threads in the CS
- But there may be other reasons:
 - Wait until queue is not empty before executing `get()`
 - Wait until there are no readers (or writers) in a reader/writer block
 - ...

Busy Waiting: not a good way

- Wait until queue is not empty:

```
done = False
while not done:
    next = get(q)
    done = next != None
```

- Wastes CPU cycles
- Creates unnecessary contention

Binary Semaphores

Dijkstra 1962



Binary Semaphore

- Boolean variable (much like a lock)
- Three operations
 - `binsema = BinSema(False or True)`
 - ▶ initializes binsema
 - `acquire(?binsema)`
 - ▶ waits until !binsema is False, then sets !binsema to True
 - `release(?binsema)`
 - ▶ sets !binsema to False
 - ▶ can only be called if !binsema = True

P & V

- ◉ Dijkstra was Dutch
 - He said **Probeer-te-verlagen** instead of acquire - and shortened it to **P**
 - He said **Verhogen** instead of release - and shortened it to **V**
 - Still very popular nomenclature
 - To remember it:
 - ▶ **P**rocurer (acquire)
 - ▶ **V**acate (release)

Binary Semaphore Specification

```
1  def BinSema(acquired):
2      result = acquired
3
4  def Lock():
5      result = BinSema(False)
6
7  def acquire(binsema):
8      atomically when not !binsema:
9          !binsema = True
10
11 def release(binsema):
12     assert !binsema
13     atomically !binsema = False
```


Semaphores v. Locks

Locks	Binary Semaphores
Initially "unlocked" (False)	Can be initialized to False or True
Usually acquired and released by the same thread	Can be acquired and released by different threads
Mostly used to implement critical sections	Can be used to implement critical sections as well as waiting for special conditions

Waiting with Semaphores

```
1  import synch
2
3  condition = BinSema(True)
4
5  √ def T0():
6      acquire(?condition)
7
8  √ def T1()
9      release(?condition)
10
11
12  spawn(T0)
13  spawn(T1)
```

Encode condition as a binary semaphore

Wait for condition to come true

Signal condition has become true

What happens if T0 runs first?

What happens if T1 runs first?

Semaphores can be locks too!

```
lk = BinSema(False)
```

Initialized to False

```
acquire(?lk)
```

grab lock

```
release(?lk)
```

release lock

What else can we do
with binary semaphores?