

# Peterson's Algorithm: Flags and Turns!

```
1 sequential flags, turn ← Prevents out-of-order execution
2
3 flags = [ False, False ]
4 turn = choose({0, 1})
5
6 def thread(self):
7     while choose({ False, True }):
8         # Enter critical section
9         flags[self] = True ← I'd like to enter...
10        turn = 1 - self ← ...but you go first!
11        await (not flags[1 - self]) or (turn == self)
12        # Wait until alone or it's my turn
13        # Critical section is here
14        cs: assert countLabel(cs) == 1
15
16        # Leave critical section
17        flags[self] = False ← Leave
18
19 spawn thread(0)
20 spawn thread(1)
```

#states = 104 diameter = 5  
#components: 37  
no issues found

# What about a proof?

- To understand **why** it works...
- We need to show that, for any execution, all states reached satisfy mutual exclusion
  - i.e., that mutual exclusion is an **invariant**
- **See the Harmony book for a proof!**
  - or come talk to me!

# Peterson's Reconsidered

- Mutual Exclusion **can** be implemented with atomic LOAD and STORE instructions
  - multiple STOREs and LOADs
- Peterson's **can** be generalized to more than 2 processes (as long as the number of processes is known) but it is a mess...
  - ...and even more STOREs and LOADs

**Too inefficient in practice!**

# Peterson's even more Reconsidered!

- It assumes LOAD and STORE instructions are **atomic**, but that is **not guaranteed** on a real processor
  - Suppose  $x$  is a 64-bit integer, and you have a 32-bit CPU
  - Then  $x = 0$  requires 2 STORES (and reading  $x$  two LOADs)
    - ▶ because it occupies 2 words!
  - Same holds if  $x$  is a 32-bit integer, but it is **not aligned on a word boundary**

# Concurrent Writing

- Say  $x$  is a 32 bit word @ 0x12340002
- Consider two threads, T1 and T2
  - T1:  $x = 0xFFFFFFFF$  (i.e.,  $x = -1$ )
  - T2:  $x = 0$
- After T1 and T2 are done,  $x$  may be any of
  - 0, 0xFFFFFFFF, 0xFFFF0000, or 0X0000FFFF
- The outcome of concurrent write operations to a variable is **undefined**

# Concurrent Reading

- Say  $x$  is a 32 bit word @ 0x12340002, initially 0
- Consider two threads, T1 and T2
  - T1:  $x = 0xFFFFFFFF$  (i.e.,  $x = -1$ )
  - T2:  $y = x$  (i.e., T2 reads  $x$ )
- After T1 and T2 are done,  $y$  may be any of
  - 0, 0xFFFFFFFF, 0xFFFF0000, or 0X0000FFFF
- The outcome of concurrent read and write operations to a variable is **undefined**



# Data Race

- When two threads access the same variable...
- ...and at least one is a STORE...
- ...then the semantics of the outcome is **undefined**

# Harmony's "sequential" statement

- *sequential* *turn, flags*
- Ensures that LOADs and STOREs are atomic
  - concurrent operations appear to be executed sequentially
  - this is called *sequential consistency*
- Say  $x$ 's current value is 3; T1 STOREs 4 into  $x$ ; T2 LOADs  $x$ 
  - with atomic LOAD/STORE, T2 reads 3 or 4
  - with modern CPUs/compiler, what T2 reads is undefined



# Sequential Consistency

- Java has a similar notion
  - volatile int x (not the same as in C/C++)
- Loading/Storing sequentially consistent variables is more expensive than loading/storing ordinary variables
  - it restricts CPU or compiler optimizations

So, what do we do?

# Interlock Instructions

- Machine instructions that do multiple shared memory accesses atomically
- TestAndSet s
  - returns the old value of s (LOAD r0,s)
  - sets s to True (STORE s, 1)
- Entire operation is atomic
  - other machine instructions cannot interleave

# Harmony Interlude: Pointers

- If  $x$  is a shared variable,  $?x$  is the **address** of  $x$
- If  $p$  is a shared variable, and  $p == ?x$ , then we say that  $p$  is a **pointer** to  $x$
- Finally,  $!p$  refers to the **value** of  $x$

# Test-and-Set in Harmony

```
1  def test_and_set(s):  
2      atomically:  
3          result = !s  
4          !s = True
```

• For example:

lock1 = False

lock2 = True

r1 = test\_and\_set(?lock1)

r2 = test\_and\_set(?lock2)

assert lock1 and lock2

assert (not r1) and r2

# Recall: bad lock implementation

```
1  lockTaken = False
2
3  def thread(self):
4      while choose({ False, True }):
5          # Enter critical section
6          await not lockTaken ← Test..
7          lockTaken = True ← ..and set
8
9          # Critical section
10         cs: assert countLabel(cs) == 1
11
12         # Leave critical section
13         lockTaken = False
14
15     spawn thread(0)
16     spawn thread(1)
```

*Test and set  
not  
atomic!!*

*Test..*

*..and set*



# A good implementation ("Spinlock")

```
1 lockTaken = False
2 |
3 def test_and_set(s):
4     ... atomically:
5     ...     result = !s
6     ...     !s = True
7     ~
8 def thread(self):
9     ... while choose ( {False, True} ):
10    ...     # enter critical section
11    ...     while test_and_set(?lockTaken):
12    ...         ... pass
13    ~
14    ...     cs: countLabel(cs) == 1
15    ~
16    ...     # exit critical section
17    ...     atomically lockTaken = False
18    ~
19 spawn thread(0)
20 spawn thread(1)
```

Same idea  
as before,  
but now  
with an  
**atomic**  
test&set!

Lock is repeatedly  
"tried", checking on a  
condition in a tight  
loop ("spinning")

# Locks

- Think of locks as “baton passing”
  - at most one thread can “hold” False



# Specifying a Lock

```
1  def Lock():
2      result = False
3
4  def acquire(lk):
5      atomically when not !lk:
6          !lk = True
7
8  def release(lk):
9      assert !lk
10     atomically !lk = False
```

An object, and the behavior of the methods that are invoked on it

- uses **atomically** to specify the behavior of these methods when executed in isolation

# Locks and Critical Sections

Two important invariants

- $T @ cs \Rightarrow T$  holds the lock
- At most one thread can hold the lock



# Implementing\* a lock

\*Just one way of doing so

```
1  def test_and_set(s):  
2      atomically:  
3          result = !s  
4          !s = True  
5  
6  def Lock():  
7      result = False  
8  
9  def acquire(lk):  
10     while test_and_set(lk):  
11         pass  
12  
13  def release(lk):  
14     atomically !lk = False
```

*Specification of the  
CPU's test-and-set  
functionality*

*Must use an atomic  
STORE instruction*

## Specification

```
1 def Lock():
2     result = False
3
4 def acquire(lk):
5     atomically when not !lk:
6         !lk = True
7
8 def release(lk):
9     assert !lk
10    atomically !lk = False
```

What an abstraction  
does

## Implementation

```
1 Def Lock()~
2 ... result = False~
3 ~
4 def test_and_set(s):~
5 ... atomically:~
6 ... | ... result = !s~
7 ... | ... !s = True~
8 ~
9 def atomic_store(var, val):~
10 ... atomically !var = val~
11 ~
12 def acquire(lk):~
13 ... while test_and_set(lk):~
14 ... | ... pass~
15 ~
16 def release(lk):~
17 ... atomic_store(lk, False)]
```

How the abstraction  
does it



# Using a lock for a critical section

```
1  import synch
2
3  const NTHREADS = 2
4
5  lock = synch.Lock()
6
7  def thread():
8      while choose({ False, True }):
9          synch.acquire(?lock)
10         cs: assert countLabel(cs) == 1
11         synch.release(?lock)
12
13     for i in {1..NTHREADS}:
14         spawn thread()
```

# Spinlocks and Time Sharing

- Spinlocks work well when threads on different cores need to synchronize
- But what if two threads are on the same core?
  - when there is no preemption?
    - ▶ all threads may get stuck while one is trying to obtain the spinlock
  - when there is preemption?
    - ▶ still delays and a waste of CPU cycles while a thread is trying to obtain a spinlock

# Beyond Spinlocks

- We would like to be able to suspend a thread that is trying to acquire a lock that is being held
  - until the lock is ready
- A context switch!

# Context switching in Harmony

- Harmony allows contexts to be saved and restored (i.e., context switch)

□ `r = stop p`

- ▶ stops the current thread and stores context in !p (p must be a pointer).

□ `go (!p) r`

- ▶ adds a thread with the given context (i.e., the one pointed by p) to the bag of threads. Threads resumes from `stop` expression, returning `r`

# Lock specification using stop and go

```
1 import list
2
3 def Lock():
4     result = { .acquired: False, .suspended: [] }
5
6 def acquire(lk):
7     atomically:
8         if lk->acquired:
9             stop ?lk->suspended[len lk->suspended]
10            assert lk->acquired
11        else:
12            lk->acquired = True
13
14 def release(lk):
15     atomically:
16         assert lk->acquired
17         if lk->suspended == []:
18             lk->acquired = False
19        else:
20            go (list.head(lk->suspended)) ()
21            lk->suspended = list.tail(lk->suspended)
```

*.acquired: boolean*

*.suspended: queue of contexts*

*add stopped context at the end  
of queue associated with lock*

*restart thread at head of queue  
and remove it from queue*



# Lock specification using stop and go

```
1 import list
2
3 def Lock():
4     result = { .acquired: False, .suspended: [] }
5
6 def acquire(lk):
7     atomically:
8         if lk->acquired:
9             stop ?lk->suspended[len lk->suspended]
10            assert lk->acquired
11        else:
12            lk->acquired = True
13
14 def release(lk):
15     atomically:
16         assert lk->acquired
17         if lk->suspended == []:
18             lk->acquired = False
19        else:
20            go (list.head(lk->suspended)) ()
21            lk->suspended = list.tail(lk->suspended)
```

*Similar to Linux  
"futex":  
with no contention  
(hopefully the common  
case) acquire() and  
release() are cheap.  
With contention, a  
context switch is  
required*



# Choosing Modules in Harmony

- “synch” is the (default) module that has the specification of a lock
- “synchS” is the module that has the **stop/go** version of the lock
- You can select which one you want”
  - **harmony -m synch=synchS x.hny**
- “synch” tends to be faster than “synchS”
  - smaller state graph

# Atomic Section ≠ Critical Section

Atomic Section	Critical Section
Only one thread can execute	Multiple threads can execute concurrently, just not within a critical section
Rare programming language paradigm	Ubiquitous: locks available in many mainstream programming languages
Good for specifying interlock instruction	Good for implementing concurrent data structures

# Using Locks

- Data structures maintain some invariant
  - Consider a linked list
    - ▶ There is a **head**, a **tail**, and a list of nodes such as the head points to the first node, tail points to the last one, and each node points to the next one, except for the tail, which points to **None**. However, if the list is empty, head and tail are both **None**
- You can assume the invariant holds right after acquiring the lock
- You must **make sure** invariant holds again right before releasing the lock

# Building a Concurrent Queue

- `q = queue.new()`: allocates a new queue
- `queue.put(q, v)`: adds `v` to the tail of queue `q`
- `v = queue.get(q)`: returns
  - `None` if `q` is empty, or
  - `v` if `v` was at the head of the queue

# Specifying a Concurrent Queue

```
1 import list
2
3 def Queue():
4     result = []
5
6 def put(q, v):
7     !q = list.append(!q, v)
8
9 def get(q):
10    if !q == []:
11        result = None
12    else:
13        result = list.head(!q)
14        !q = list.tail(!q)
15
```

*Sequential*

```
1 import list
2
3 def Queue():
4     result = []
5
6 def put(q, v):
7     atomically !q = list.append(!q, v)
8
9 def get(q):
10    atomically:
11        if !q == []:
12            result = None
13        else:
14            result = list.head(!q)
15            !q = list.tail(!q)
```

*Concurrent*



# Example of using a 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)
```

*enqueue v onto q*

*dequeue and check*

*create a queue*



# Queue implementation, v1



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def Queue():
5     result = { .head: None, .tail: None, .lock: Lock() }
6
7 def put(q, v):
8     let node = malloc({ .value: v, .next: None });
9     acquire(?q->lock)
10    if q->head == None:
11        q->head = q->tail = node
12    else:
13        q->tail->next = node
14        q->tail = node
15    release(?q->lock)
```

*dynamic memory allocation*

*create empty queue*

*allocate node*

*grab lock*

*The Hard Stuff*

*release lock*

# Queue implementation, v1



```
17 def get(q):
18     acquire(?q->lock)
19     let node = q->head:
20         if node == None:
21             result = None
22         else:
23             result = node->value
24             q->head = node->next
25             if q->head == None:
26                 q->tail = None
27             free(node)
28     release(?q->lock)
```

*grab lock*

*empty queue*

*The Hard  
Stuff*

*free dynamically allocated memory*

*release lock*

# How important are concurrent queues?

## 👁️ All important!


- ❑ any resource that needs scheduling
  - ▶ CPU ready queue
  - ▶ disk, network, printer waiting queue
  - ▶ lock waiting queue
- ❑ inter-process communication
  - ▶ Posix pipes: `cat file | sort`
- ❑ actor-based concurrency
- ❑ ...



Performance  
is  
critical!

# 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

## ⑥ 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 specification and the implementation
- We then perform the same sequence of operations using the code in both sequential specification and the implementation and check if these sequences produce the same behaviors (e.g., they return the same values)

# 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 run 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 a sequence is correct?

- We run the test program against both the 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)
```

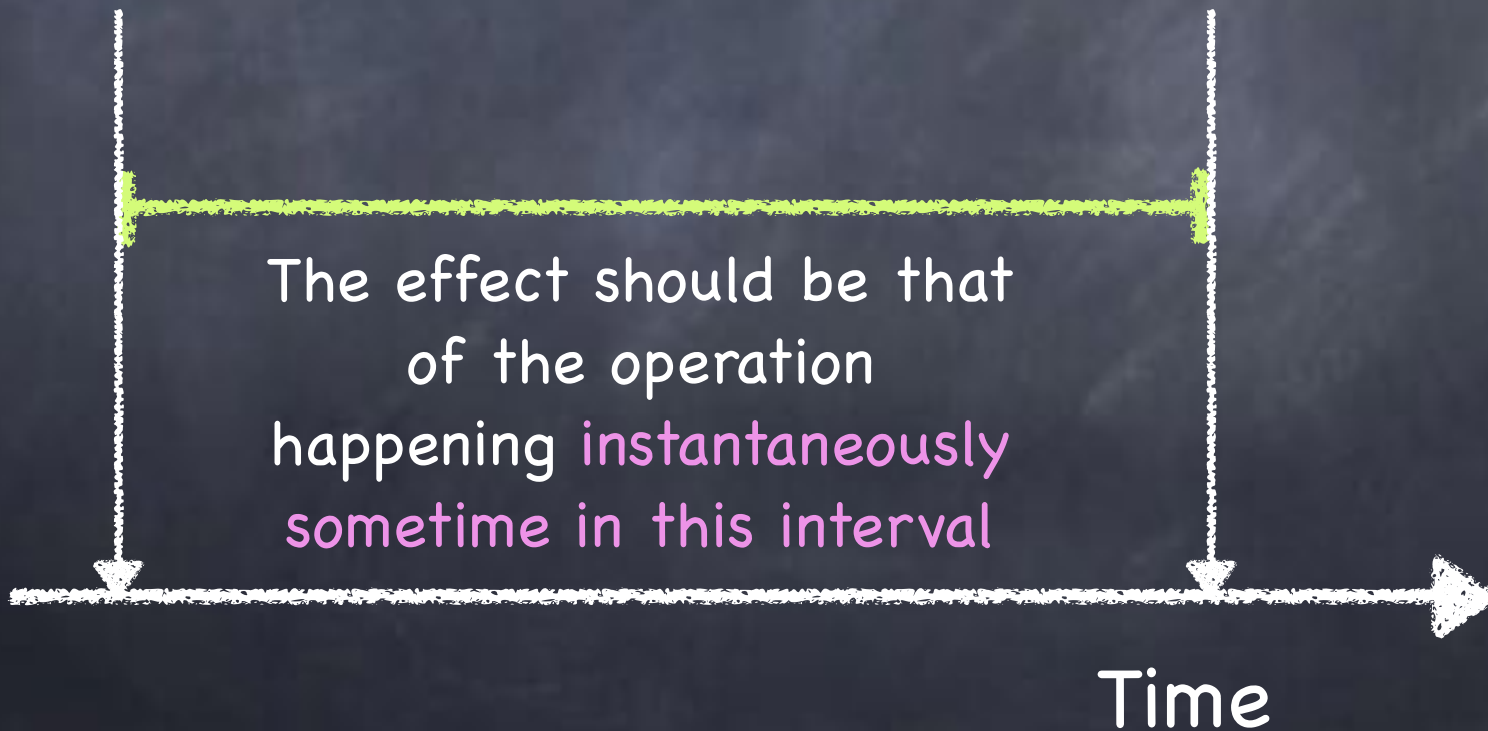
*\* always at least one  
put and one get*

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

# 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

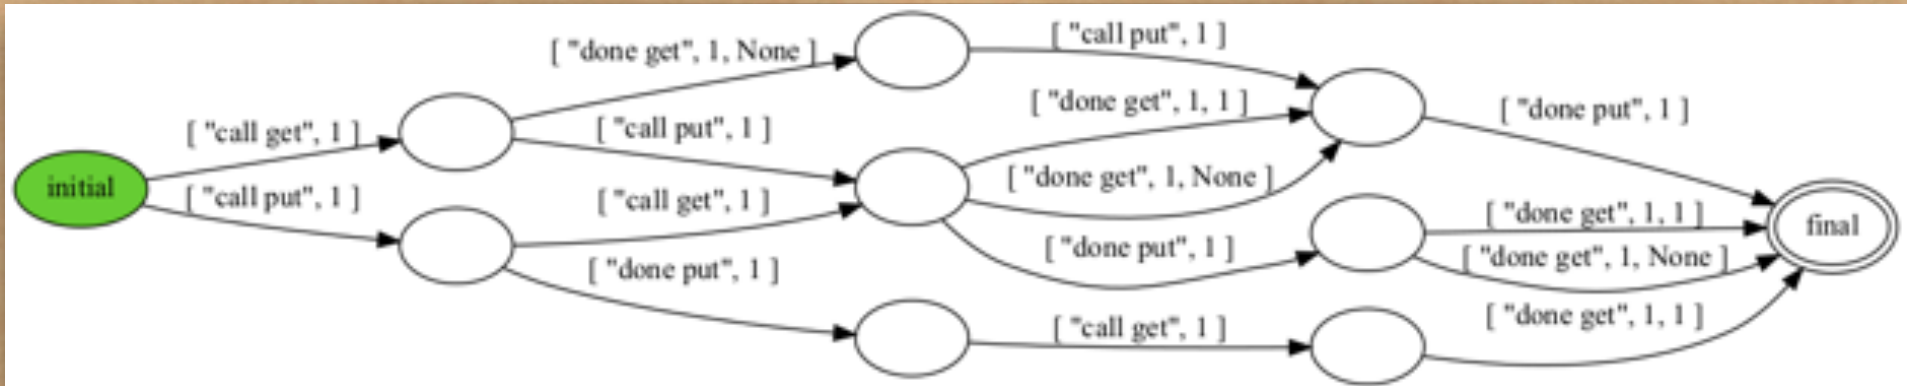


# Correct Behaviors

Suppose the queue is initially empty



# Queue test program




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

# Testing: comparing behaviors

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


# Queue implementation, v2:2 locks



- **Separate locks for head and tail**
  - put and get can proceed concurrently
- Trick: a dummy node at the head of the queue
  - last node to be dequeued (except at the beginning)
  - head and tail never **None**




# Queue implementation, v2:2 locks



```
1  from synch import Lock, acquire, release, atomic_load, atomic_store
2  from alloc import malloc, free
3
4  def Queue():
5      let dummy = malloc({ .value: (), .next: None }):
6          result = { .head: dummy, .tail: dummy, .hdlock: Lock(), .tllock: Lock() }
7
8  def put(q, v):
9      let node = malloc({ .value: v, .next: None }):
10         acquire(?q->tllock)
11         atomic_store(?q->tail->next, node)
12         q->tail = node
13         release(?q->tllock)
```



# Queue implementation, v2:2 locks



```
15 def get(q):
16     acquire(?q→hdlock)
17     let dummy = q→head
18     let node = atomic_load(?dummy→next):
19         if node == None:
20             result = None
21             release(?q→hdlock)
22         else:
23             result = node→value
24             q→head = node
25             release(?q→hdlock)
26             free(dummy)
```

*Faster!*  
*No contention for  
concurrent enqueue and  
dequeue ops ⇒ more  
concurrency*

*BUT: Data race on  
dummy → next  
when queue is empty*

# Global vs Local Locks

- The two-lock queue is an example of a data structure with **fine-grain locking**
- A global lock is easy, but limits concurrency
- Fine-grain (local) locks can improve concurrency, but tend to be tricky to get right

# Sorted lists with lock per node



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def _node(v, n): # allocate and initialize a new list node
5     result = malloc({ .lock: Lock(), .value: v, .next: n })
6
7 def _find(lst, v):
8     var before = lst
9     acquire(?before→lock)
10    var after = before→next
11    acquire(?after→lock)
12    while after→value < (0, v):
13        release(?before→lock)
14        before = after
15        after = before→next
16        acquire(?after→lock)
17    result = (before, after)
18
19 def SetObject():
20    result = _node((-1, None), _node((1, None), None))
```

*Helper routine to find and lock  
two consecutive nodes before  
and after such that:  
 $before \rightarrow value < v \leq after \rightarrow value$*

*empty list:*

```
graph LR; Node1["(-1,  
None)"] --> Node2["(1,  
None)"]; Node2 --> None;
```

# Sorted lists with lock per node



```
1 from synch import Lock, acquire, release
2 from alloc import malloc, free
3
4 def _node(v, n): # allocate and initialize a new list node
5     result = malloc({ .lock: Lock(), .value: v, .next: n })
6
7 def _find(lst, v):
8     var before = lst
9     acquire(?before→lock)
10    var after = before→next
11    acquire(?after→lock)
12    while after→value < (0, v):
13        release(?before→lock)
14        before = after
15        after = before→next
16    acquire(?after→lock)
17    result = (before, after)
18
19 def SetObject():
20    result = _node((-1, None), _node((1, None), None))
```

Hand-over-hand  
locking



empty list:





# Sorted lists with lock per node



```
22 def insert(lst, v):
23     let before, after = _find(lst, v):
24         if after->value != (0, v):
25             before->next = _node((0, v), after)
26             release(?after->lock)
27             release(?before->lock)
28
29 def remove(lst, v):
30     let before, after = _find(lst, v):
31         if after->value == (0, v):
32             before->next = after->next
33             release(?after->lock)
34             free(after)
35         else:
36             release(?after->lock)
37             release(?before->lock)
38
39 def contains(lst, v):
40     let before, after = _find(lst, v):
41         result = after->value == (0, v)
42         release(?after->lock)
43         release(?before->lock)
```

*Multiple threads can access the list simultaneously, but they can't overtake one another!*