

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

```
def Queue() returns empty:  
    empty = []  
  
def put(q, v):  
    !q += [v,]  
  
def get(q) returns next:  
    if !q == []:  
        next = None  
    else:  
        next = (!q)[0]  
        del (!q)[0]
```

Sequential

```
def Queue() returns empty:  
    empty = []  
  
def put(q, v):  
    atomically !q += [v,]  
  
def get(q) returns next:  
    atomically:  
        if !q == []:  
            next = None  
        else:  
            next = (!q)[0]  
            del (!q)[0]
```

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 dynamic memory allocation
3
4 def Queue() returns empty:
5     empty = { .head: None, .tail: None, .lock: Lock() } create empty queue
6
7 def put(q, v):
8     let node = malloc({ .value: v, .next: None }): allocate node
9     acquire(?q->lock) grab lock
10    if q->tail == None:
11        q->tail = q->head = node
12    else:
13        q->tail->next = node
14        q->tail = node
15    release(?q->lock) release lock
```

The Hard Stuff

Queue implementation, v1



```
17 def get(q) returns next:  
18   acquire(?q→lock)  
19   let node = q→head:  
20     if node == None:  
21       next = None  
22     else:  
23       next = 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!

Queue implementation, v2:2 locks




- Separate locks for head and tail
 - put and get can proceed concurrently
- Trick: put a **dummy node** at the head of the queue
 - **last node that was 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() returns empty:
5      let dummy = malloc({ .value: (), .next: None }):
6          empty = { .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) ← Why an atomic_store here?
12         q→tail = node
13         release(?q→tllock)
```

Queue implementation, v2:2 locks



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

...and here?



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
 - think of having to walk a queue...
- 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) returns node: # allocate and initialize a new list node
5     node = malloc({ .lock: Lock(), .value: v, .next: n })
6
7 def _find(lst, v) returns pair:
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    pair = (before, after)
18
19 def SetObject() returns object:
20    object = _node((-1, None), _node((1, None), None))
```

one lock per node

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["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) returns node: # allocate and initialize a new list node
5     node = malloc({ .lock: Lock(), .value: v, .next: n })
6
7 def _find(lst, v) returns pair:
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    pair = (before, after)
18
19 def SetObject() returns object:
20    object = _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             free(after)
34             release(?after→lock)
35             release(?before→lock)
36
37 def contains(lst, v) returns present:
38     let before, after = _find(lst, v):
39         present = after→value == (0, v)
40         release(?after→lock)
41         release(?before→lock)
```

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