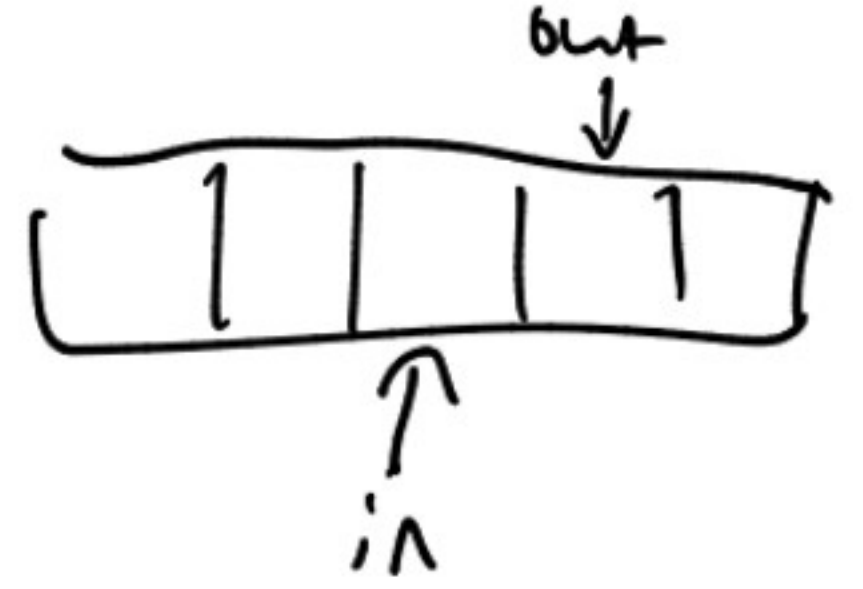


4/4/0 Lecture 9: Mesa-style monitors

- Reader/Writer problem
- MESA-style monitors
- Optimization
- (time permitting) conditions for deadlock

Bounded buffer w/ Hoare-style monitor



```
class Buffer (Monitor):
```

```
    def __init__(self, N):
```

```
        self.n = N
```

```
        self.in = 0
```

```
        self.out = 0
```

```
        self.buf = Object[N]
```

```
    def hasSpace():
```

```
        in  $\neq$  out - 1
```

mutual exclusion

```
    def put(object o):
```

- wait until hasSpace()

```
    buf[in++] = o
```

Reader-writer lock

or acquire/release

- have methods to enter/leave as a reader
enter/leave as a writer
- Either have (- any # of readers)
or (- one writer) at a time

class RWLock (Monitor):

def __init__(self):

self.lock = Lock()

self.num_readers = 0

self.writer = False

invariant:

either writer is False

or num_readers = 0

self.no_writers = Condition(lock)

predicate: writers = False

self.no_readers_writers = Condition(lock)

predicate: readers = 0 and writers = False

note: uses notify ~~ALL~~

def enter_reader(self):

with self.lock:

wait until no writers

while not (writers = False):

no_writers.wait()

num_readers ++

def leave_reader(self):

with lock:

num_readers --

self.no_readers_writers.notifyAll()

def enter_writer(self):

with self.lock:

[while not (readers = 0 and

writers = False):

no_readers_writers.wait()

self.writers = True

def leave_writer(self):

with self.lock:

self.writers = False

self.no_writers.notifyAll()

self.no_readers_writers.notifyAll()

first writer to
wake immediately
make predicate
false, so no
reason to
notify all
threads,
one is
enough.

only thing
connecting condition
variable to predicate
are (1) comment

(2) while loop where we wait

Condition needs non.lock
so it can acquire &
release it

the predicate.

Class Example:

- init -

lock

x = True

x_false = Cond(lock)

predicate: x is false

notify ~~ATK~~

def f():

while not (x = False):

x_false.wait()

x = True

def g():

while not (x = False):

x_false.wait()

~~return~~

x_false.notify()

def h():

x = False

x_false.notify ~~ATK~~ ()

↑
- might only notify
thread calling
↓

- f can make progress
after g runs
(because x = True)
but is still sleeping
because g notified.

Writing monitor

- ① write in Hoare style:
 - think about state to maintain
 - think about preconditions, postconditions, invariants (write down)
- ② add condition variables $\left\{ \begin{array}{l} \text{while not (predicate);} \\ \text{cv.wait()} \\ \text{to wait;} \end{array} \right.$ with comments indicating predicates.
- ③ compare state changes to predicates, add notifyAll when I make a predicate true
- ④ optimize by changing notifyAlls to notify where appropriate, make sure that all f's that wait for that predicate make it false before returning.

WARNING!

- use monitor design pattern:

o all CVs always have associated predicates
(boolean exprs involving member vars)

o while not(pred):
cv.wait()

Lock

with lock:

- acquiring mutual exclusion lock upon entering, releasing it when you leave

(like a waiting room)

Condition class

- ① release monitor lock
 - ② sleeping until notifyAll called
 - ③ reacquire monitor lock.
- notifyAll: wakes up all sleeping threads

shutting into waiting room.

- notify: wakes up one sleeping thread.

- instead of notifying when we might have made predicate true, check if only notify when we did make it true

```
def leave_reader():  
    with lock:
```

```
        readers--  
        optimize → if readers == 0:  
            no_readers.notify_all()
```

- use notify if only one thread will be able to make progress.

General synch. advice:

- ① use high-level primitives wherever possible:
- bounded buffer or RW lock provided by OS/language/standard library
 - monitors
 - semaphores
 - spin locks, TAS, CAS, ... ← occasionally need lock-free data structures.
- blocking primitives.

- ② don't mix primitives
- e.g. monitor code (outside of wait) shouldn't block, since it holds the lock.

- ③ document your code, use good var names, write down predicates & invariants & specifications, ...

Monitor implementation (pseudocode on website)

- very similar to semaphore impl:
 - spin lock to protect internal state
 - queue of processes trying to acquire monitor lock
 - add to queue when "with lock"
 - dequeue i when "with lock" returns.
 - queue of processes waiting on each card. var.
 - add to queue when wait()
 - remove i ~~from~~ on notify
add to monitor lock queue.

Java "synchronized" blocks are like monitors.

- Every object can be made a monitor
by using "synchronized"
- Every object is a CV.