

Lecture 7: Semaphores

- Programming with semaphores
 - Mutual exclusion
 - Counting
 - Signalling
- Implementing semaphores

Semaphore interface:

contains a non-negative counter
(counting available resources)

init (value)

P() decrement "probe"
◦ block until semaphore is at least 1 before
decrementing.

V() increment "release, signal"
(can go above initial value)

other non-standard operations to avoid:

- bounded semaphore:
{ specify an upper bound
{ block if V() would go above

- read value of semaphore: dangerous!

context switch → if $\text{sema.get_value()} > 0$:
P(sema) # know this won't block,
since sema is available

m python: acquire()

m python: release()

Semaphore design patterns:

① Mutual exclusion:

init semaphore to 1
before a crit. section: $P(\text{sema})$
after: $V(\text{sema})$
(only 1 thread in CS. at a time)

} think of semaphore
as a lock
only one lock
available.

② Counting semaphore:

init semaphore to # of available
resources
to acquire resource: P
release resource: V

③ signalling:

initialize semaphore to 0
to wait for a signal: $P(\text{sema})$
to send signal: $V(\text{sema})$

block
until
signal
arrives

call it
lock.
or better yet:
lock_resource.

call it
num_resources.
num_printers


call it
signal_description
signal_space_is_avail

- testing is very
partial because
scheduler is
non-deterministic.

- verification by
other means is
important

- need good code
hygiene

Example : Bounded buffer



interface:

put (object)

"block until space available to add o, add o to buffer, return"

put(01)
:
put(04)
put(05)
put(06)

Obj get()

"block until an object is in the buffer, remove & return object"

get() → 01
get() → 02
get() → 03
→ 04
get() → 05
get() → 06

- what code needs mutual exclusion?
- what resources need to be counted?
- what signals do you want to send?

init:

```

buffer = new Object[n]
num_empty = new Sem(N)
num_objects = new Sem(0)
in = 0 # location of next object to put in the buffer, protected by lock
lock = new Sem(1)
signal_obj_avail = new Sem(0) # used when num_objects > 0

```

put(Object o):

```

P(num_empty)
P(lock)
buffer[in++] = o
V(lock)
V(num_objects)
V(signal_obj_avail)

```

switch? →

mod n, so loop around

get:

```

P(num_objects)
P(lock)
result = buffer[in-- (mod n)]
buffer[in] = null
V(lock)
V(num_empty)

```


alternately:

```

V(num_empty)
V(lock)

```

doesn't matter



Deadlock situation!

get can't proceed until put happens, but put can't happen while get is holding lock.

Using Signalling Semaphores

threads = pool()

↙ thread represented by an object:

- identifier

- TCB.

- semaphore allocated for this thread. (input-signal)

inside thread, to wait for input:

P(self.input-signal)

inside scheduler or I/O handler:

to give input to the thread: V(thread.input-sig)

ready →

running →

waiting → dev1 → TCB → TCB.

↘ dev2

Implementing a semaphore

```
class Semaphore:
```

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

```
        self.value = value # protected by lock
```

```
        self.lock = 0 # TAS spin lock
```

```
        self.waiting = Queue() # protected by lock
```

```
    def P(self):
```

```
        while not TAS(lock):
```

```
            # wait: either
```

```
                - pass # multi-processor
```

```
                # : or put this process in ready queue.
```

```
                - yield() # single processor
```

```
        # acquired lock.
```

```
        check value:
```

```
        if > 0, decrement, release lock, return.
```

```
        if = 0, put current thread into waiting  
        state, (i.e. put it on a queue  
        of waiting procs for this  
        sema.)
```

```
        deschedule this thread  
        context switch.
```

(Full impl next lec)