



Lab Exam #2

- 25th Oct. Fri

8.30 am



9th Oct. Wed — pthreads lab
Session
meet in SL2

11th Oct. Fri — guest lecture (CC103)
measurement, analysis
& scalability reasoning
of computing systems



Lab 4a — available
4b — search +
thread pool +
pthreads.

4c — cflask

(http header parsing
+ function dispatch
+ pthreads)

Projects

— idea list available
(some ongoing edits)
— groups of 1 or 2
— original ideas
not listed
encouraged
(get 1000 bonus
points)*

* not redeemable.

② Threads & Synchronization.

Q. What is a thread?

~ an independent ordering/sequence of instructions.

~ "thread of execution"

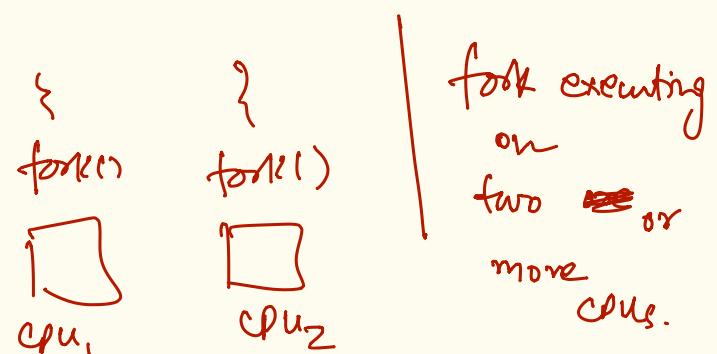
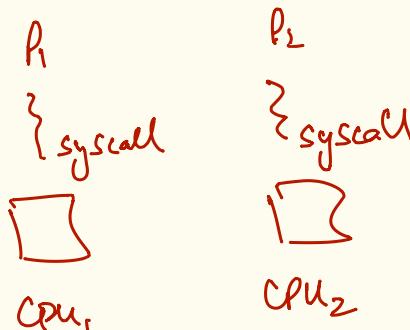
~ process is an example of thread of execution.

Why is multi-threaded execution interesting?

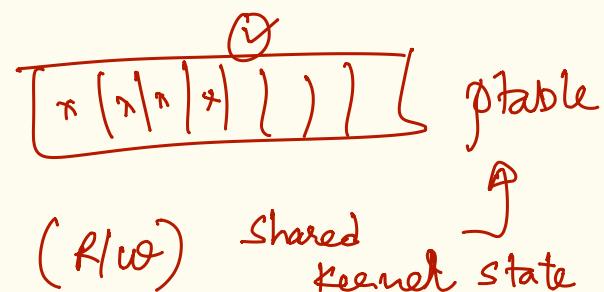
challenging?

- multiple threads of execution may simultaneously access / or have access to shared memory regions / data objects.
- all in-memory writes: read - update - write → interleaving of actions by multiple threads via multiple threads can lead to inconsistency.
 - shopping cart
 - ticket service
 - money transfer

does this happen in the kernel?



(ii) free list access
for memory allocation.

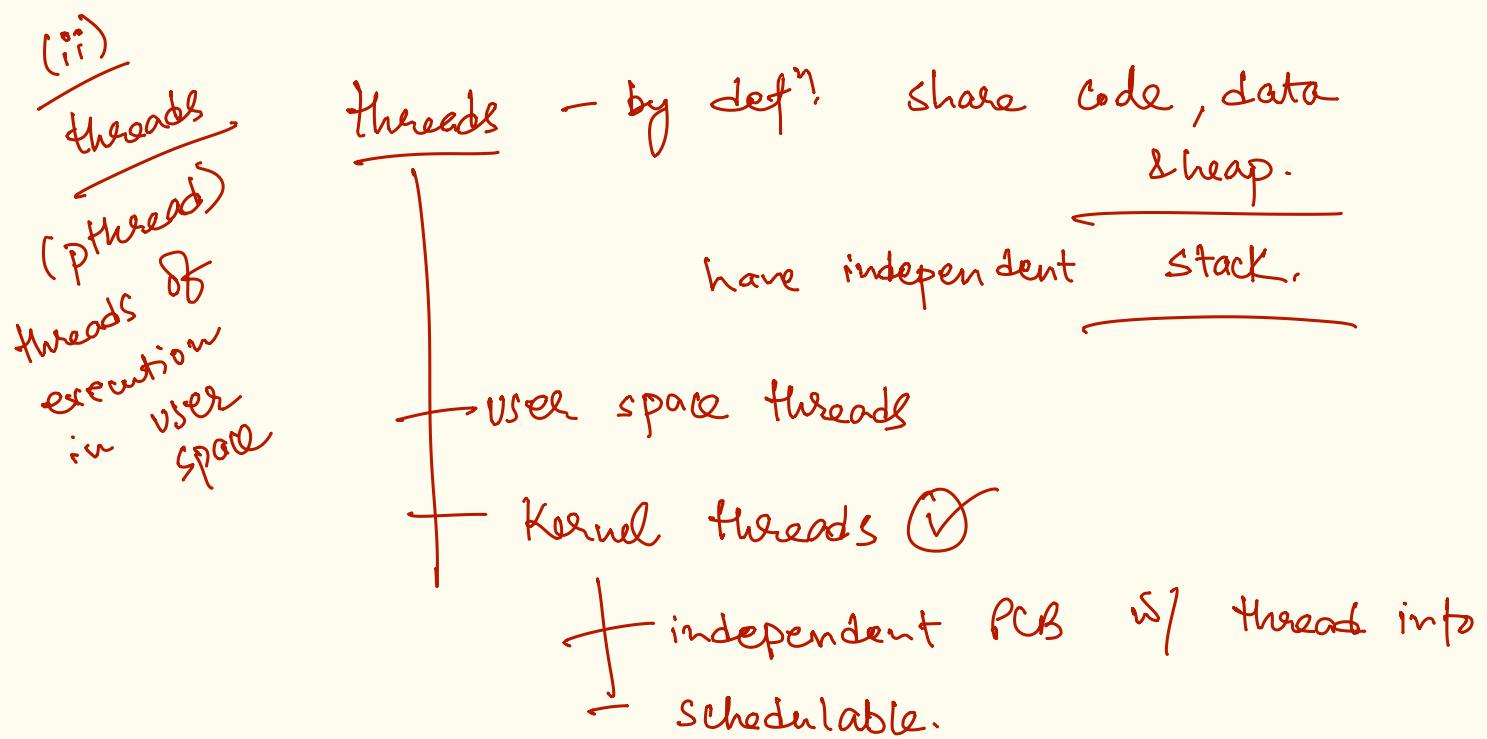
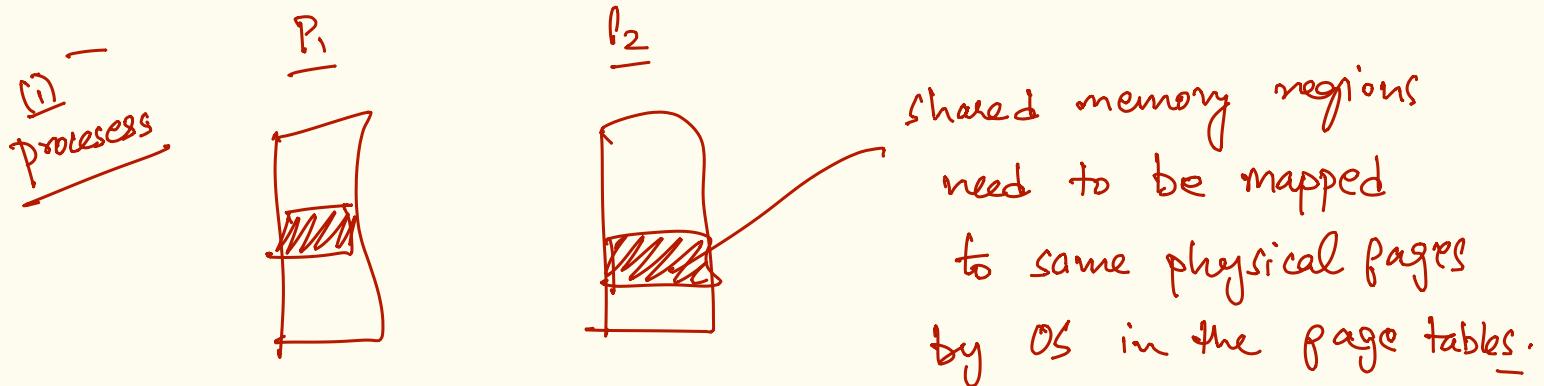


(iii) ready queue ~ list
of ready processes
for ~~scheduling~~ scheduling.

within the Kernel simultaneous access to shared state is possible —

- (i) system call + system call
- (ii) system call + interrupt
- (iii) interrupt + interrupt

Q ~ how about multiple threads of execution in user mode?



Q. how to "protect" / safely / correctly / consistently
read / update shared state?

④ critical section ~ where possibilities of
getting in inconsistent state
exist due to access of
shared state.

(i) disable preemption

[Single] CPU

depend on exit, yield, ...

(ii) disable interrupts

OK for multi-CPU

stalling all CPUs

(iii) atomic updates

ISA supported

eg. x86. lock inc R0 ;

↑ prefix locks the address bus
till instruction is complete.

(iv) locks

locking based primitives

getlock(); atomic

// critical section

release-lock();



What is a lock?

- lock is a variable (value at a memory address)

- lock $\leftarrow 0$ // available

lock $\leftarrow 1$ // not available.

Spin lock : Spins on the CPU till gets a lock.
 (i) Spin lock : Lock is the memory variable
 getlock : MOV R0, 1
 TSL LOCK, R0
 CMP R0, 0
 JNZ getlock
 ret

unlock : MOV LOCK, 0
 ret

TSL, compare & swap
 } xchg.
 atomic { oldvalue = ~~LOCK~~ LOCK
 } LOCK = R0
 return oldvalue in R0

mutex — yields CPU if lock not available
 & wakes up to test lock when lock released.

condition variables ~ any generalized condition for synchronization.

Semaphores ~ UP V — increment & wakeup
 (integers/conditions) down P — decrement & sleep

reader/writer lock —

barrier