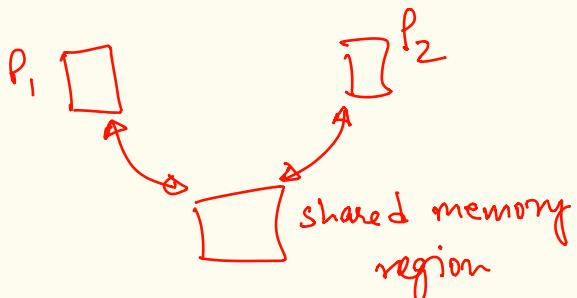


## # Concurrency &amp; Synchronization



④ fork()

↑ shares / copies

↓ lock on all  
memory regions

⑤ get address space abstraction

provides isolation

guarantee

~ Shmget } Linux-like OS  
 shmat } system calls to  
 setup & attach  
 address space region  
 to shared areas

⑥ both  $P_1$  &  $P_2$  access some common  
 memory / variable object

$P_1$  &  $P_2$   
 execute  $\rightarrow$  C-code:

$count = 3$

$P_1$        $P_2$   
 $\downarrow$        $\downarrow$   
 4      5

ISA  
 speak

$\underbrace{count} = \underbrace{count + 1};$

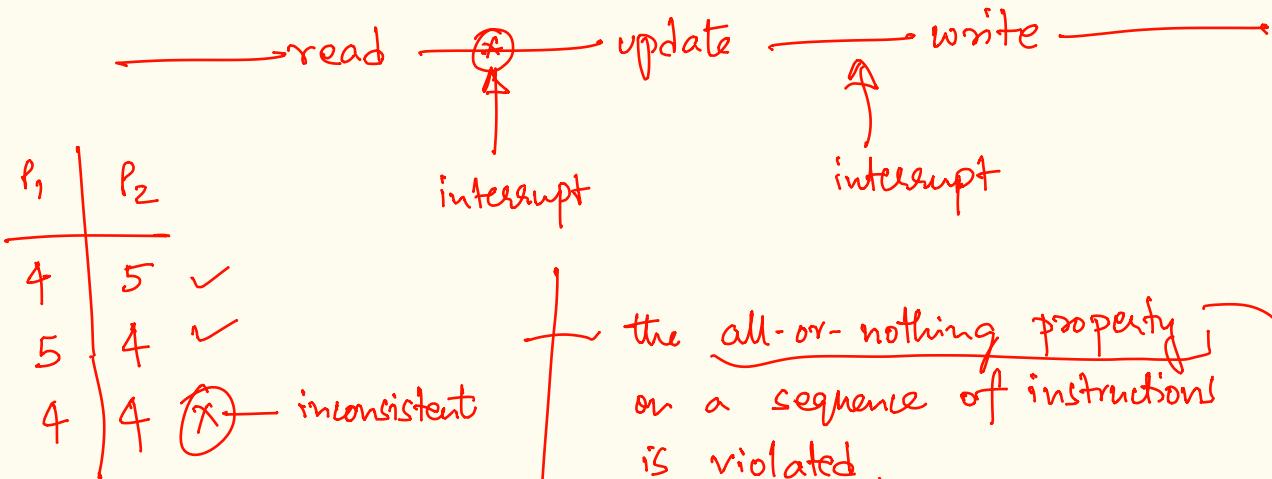
$\rightarrow$  single 'c'  
instruction

{  
 mov ebx, (%eax)  
 add ebx, 1  
 mov (%eax), ebx

move contents  
 at  $eax$  &  
 to  $ebx$  reg.

→ read → update → write →

this is the complete logical  
 action



- ② race condition happens when
- (i) shared state / objects / variables
  - (ii) multi-instance execution
  - (iii) read - update - write semantics on shared state

program instance race to execute common instructions & the order of elements instances in the race lead to inconsistent outcomes.

④ set of instructions that can lead to race conditions is the critical section.

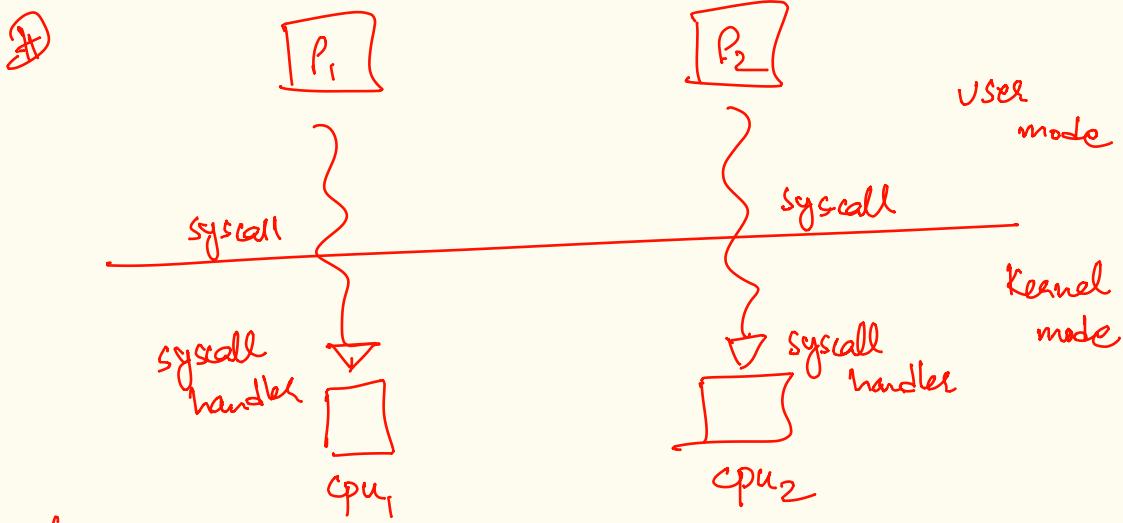
goal: is to serialize (execute one at a time) the critical section of process / program logic

to ensure atomicity



when / which execution instances ~~does~~ race conditions happen ?.

- ① multi-process + shared memory
- ② multi-threaded (processes)
- ③ multi-instances of ~~the~~ Kernel execution !



e.g.

(i) `fork()`

$\uparrow$   $\xrightarrow{\text{rv6}}$  ~ array of PCBs  
 $\downarrow$  allocproc ~ find a free PCB  
 $\quad$  if ( $\text{proc} \rightarrow \text{state} == \text{UNUSED}$ )

(ii) Scheduler as last instruction of system call

$\uparrow$   $\text{np} = \text{get next process}();$   $\leftarrow$  go through list of  
 checks for READY/RUNNABLE  
 selects a process  
 $\quad$  ( $\text{np} \rightarrow \text{state} == \text{RUNNABLE}$ )  
 $\downarrow$   $\text{np} \rightarrow \text{state} = \text{RUNNING};$   
 $\quad$  schedule ( $\text{np}$ ); //on CPU

(iii) kalloc

$\uparrow$  which fetches a free page  
 $\quad$  from the freelist of memory pages  
 $\quad$  to build  $v^2p$  mapping  
 $\quad$  in the page table.

# race conditions in the Kernel mode are possible when \_\_\_\_\_ following cases.

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

\* # Synchronization primitives/techniques to avoid race conditions

### ① disable pre-emption

~ ensure run-to-completion when in Kernel mode!

- ve's:
- okay of syscalls (okay for syscalls)
  - interrupts land in Kernel mode (multi-instance OS execution)
  - inefficient!
  - not applicable/feasible with multiple CPUs.

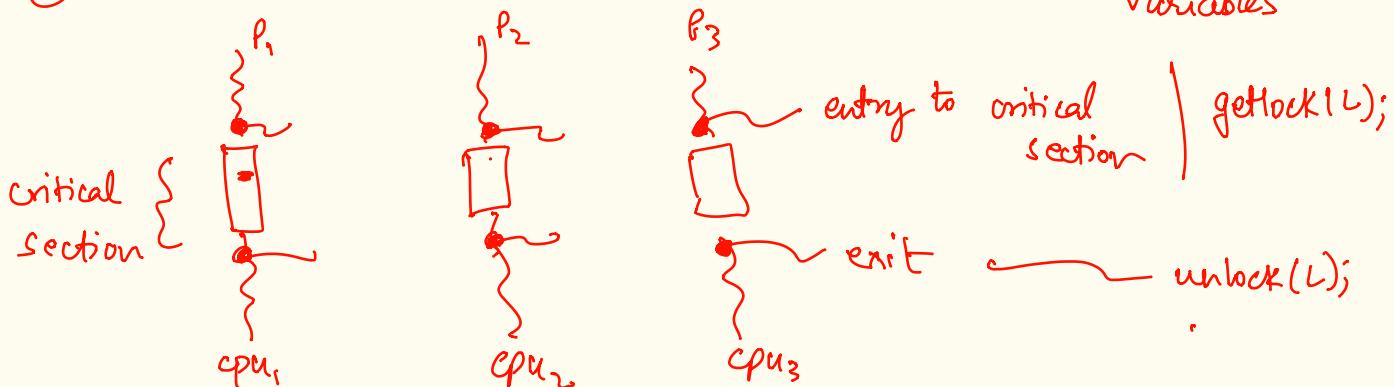
### ② + disable interrupts:

- ensures run to completion in Kernel mode

-ve's

- lost interrupts?
- not feasible multi-CPU setup.

### ③ mutual exclusion via locks / semaphores / condition variables





how to design these mutual exclusion  
primitives?

what is a lock?

variable/struct/object in memory.

int lock;

lock  $\Rightarrow$  0 ; // available

lock  $\Rightarrow$  1 ; // taken!