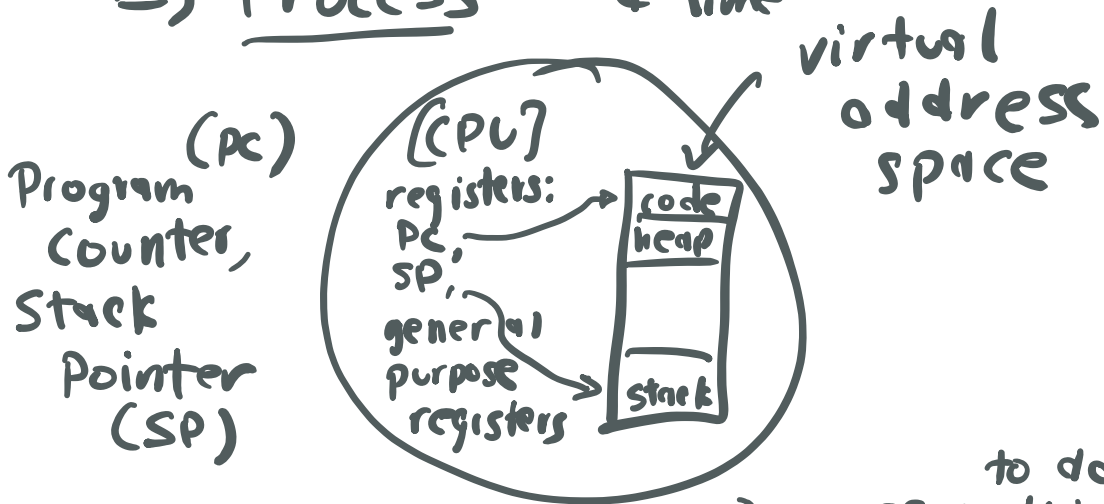


Today: OS, Part 2

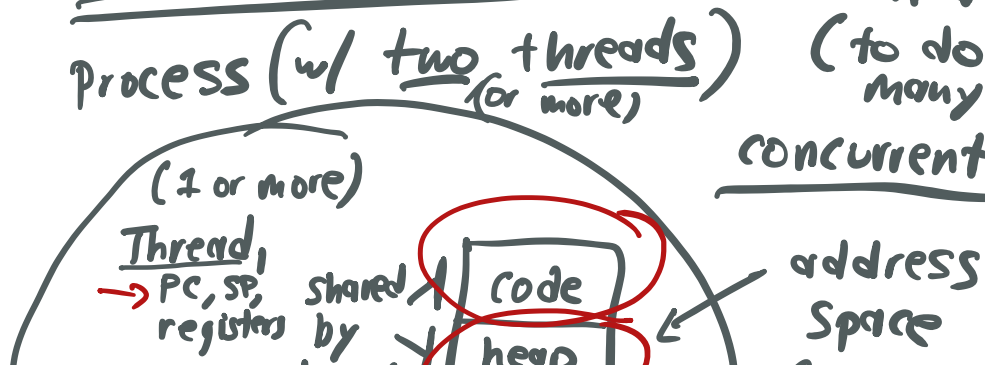
What is concurrency?
What are the major issues?
(two: atomicity and ordering)

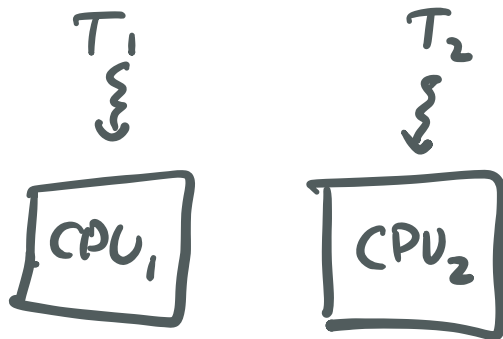
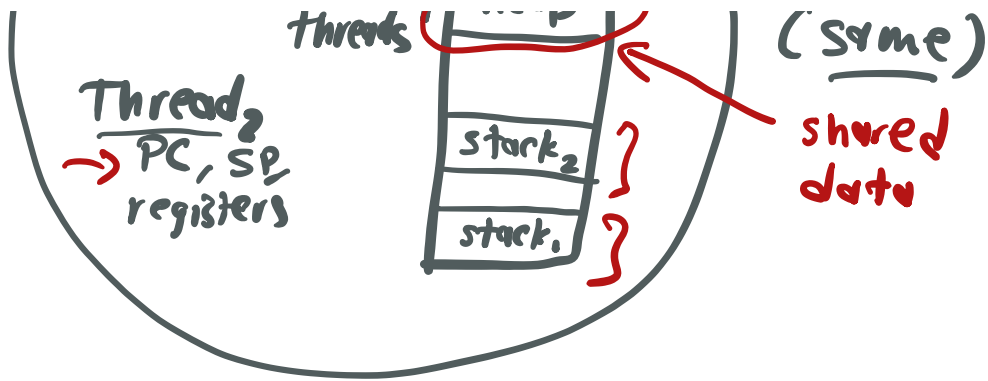
Part 1: Virtualization
Part 2: => (Concurrency)
Part 3: Persistence

To begin:
abstraction of running program
=> Process: one thing at a time



to do many things at a time
=> Multi-Threaded Process: (to do many things concurrently)





Why have multi-threaded program?

⇒ Parallelism :

to use N CPUs ($N > 1$)
to work on one problem
faster

⇒ I/O :

waiting for an I/O ⇒ thread blocked
(but others can run)

Problems: Synchronization

⇒ Atomicity :

sequence of
actions: happen
"all at once"

(not interrupted)

=> Ordering : one thing happens before the other

sequence of instructions:

inst 1 -> load addr, reg

i 2 -> add 1, reg

i 3 -> store reg, addr

(mov
in
x86)

Example run : 2 CPUs

critical section : update to shared data

-> CPU₁
=> load => reg (100)
add 1 (101)
store => 101

-> CPU₂
=> load => reg (100)
add 2 (101)
store => 101

mem
balance:

100
~~101~~
101

i 1
2
3

Problem: data race
(or race condition)

i 1
2
3

outcome:
indeterminant
(outcome is not deterministic)

can even happen on 1 CPU:

CPU₀

T₁

T₂

PC load → reg(100)

PC add → reg(101) ←

timer interrupt

OS:
thread switch

load = 100

add → 101

store 101

balance

interrupt

OS
switch

store → 101

Desire: critical section of code ⇒ atomically (all at once)

lock : lock, unlock OR acquire, release

[how to build?]
⇒ what is it?

⇒ lock ensures that only 1 thread is running w/in critical section

AT A TIME

First Try: Spin Lock

int mutex = 0; 0 → free
 1 → held

```
SpinLock () {  
    while (mutex == 1)      test mutex  
        ; // spin            branch  
    mutex = 1;  
}
```

T₁
test // 0

mutex → 1

⋮
'
⋮

both enter

T₂

test // 0
mutex → 1

←

⋮
'
⋮