

Today (3/20) "OS"

=> Bugs (Deadlock) } concurrency

=> Mid term Review } break
(cont. tomorrow)

Concurrency Bugs

Studies: bugs exist (still true)
concurrent bugs?

=> atomicity violations
(data races)

need locks

=> ordering violations


(control races)

{ need condition variables }

assume ordering: >

event 1 before event 2

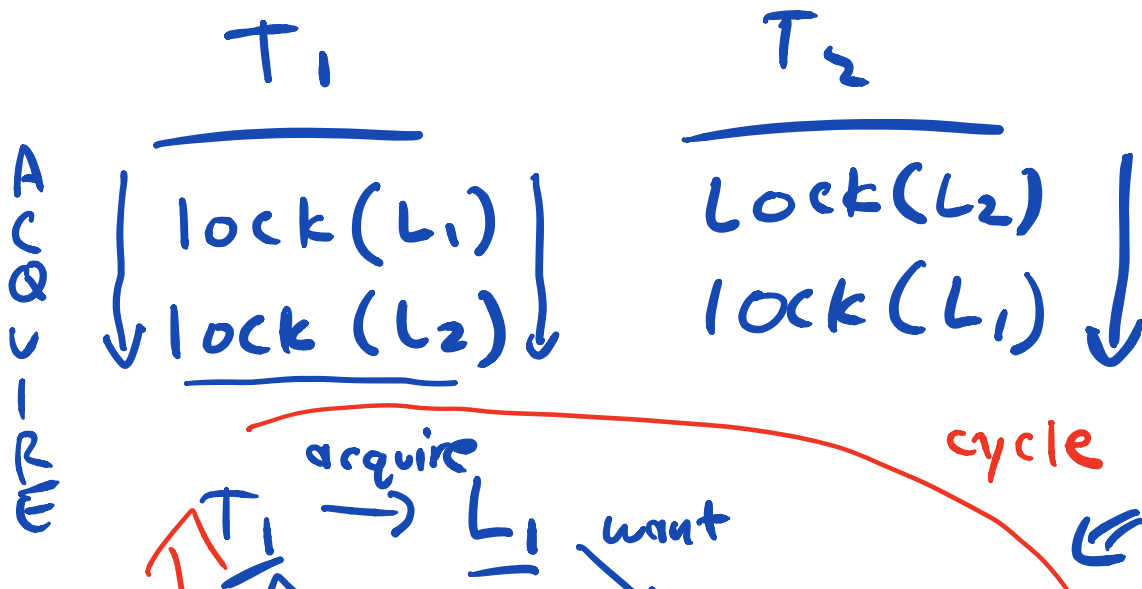
=> dead lock




helpful:
destroy car
(breaking cycle)

code: locks L_1, L_2

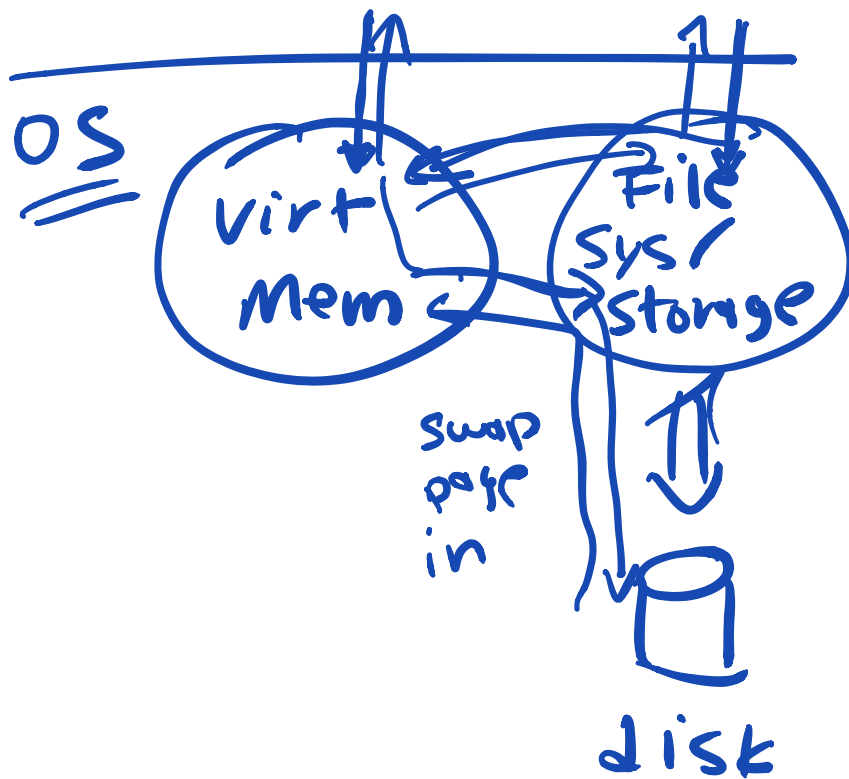
threads T_1, T_2





Examples: why deadlock occur?

\Rightarrow Software has complex dependencies



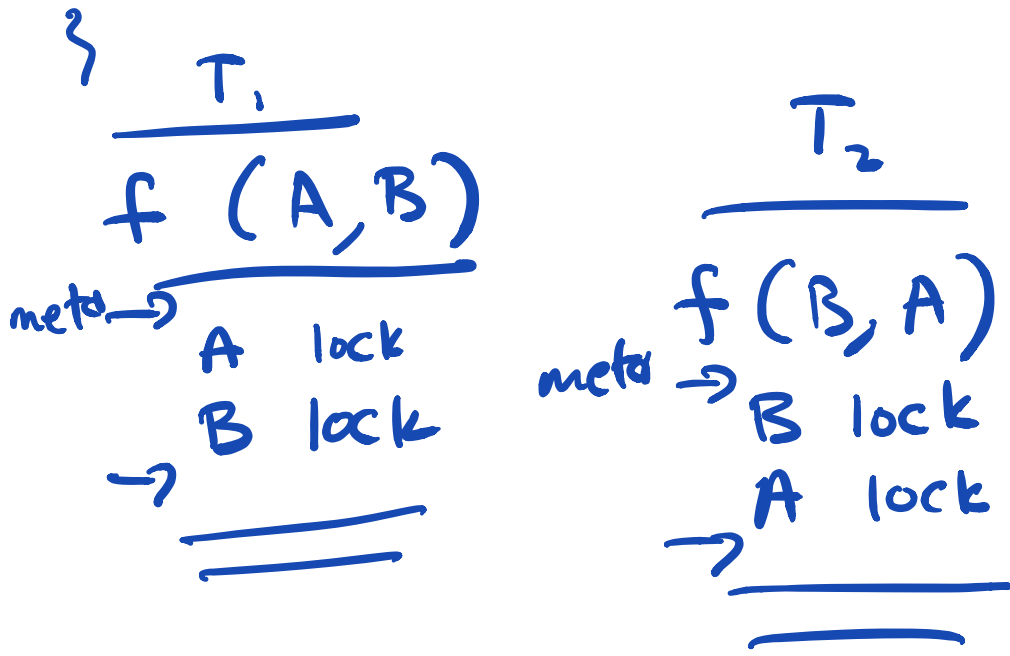
example:



```

void function ( obj1, obj2 ) {
    // perf computation
    // over both objects
    // atomically
    {
        obj1.lock();
        obj2.lock();
    } // lock by address order
    // ...
    // release locks here
}

```



⇒ Conditions: Read:

 { "Soul of a

slow Machine" }

⇒ future: ethics

Conditions: (All 4 must hold for deadlock)

- (mutual exclusion) ⇒ resource: exclusive access
- (hold-and-wait) ⇒ another req → deny
- (no preemption) ⇒ thread^(s) holding, waiting
- (circular wait) ⇒ holding resource, can't be taken away

Goal: Prevent

(by breaking

one conditions)

T_1, \dots, T_N

exist cycle

$T_1 \rightarrow T_2 \rightarrow \dots$

$\rightarrow T_N \rightarrow T_1$

Prevent: circular wait

⇒ define lock ordering

(most common solutions)

Prevent: hold + wait

=> lock acquisition lock
(meta lock)

hold ^{meta} lock

while acquiring

=> lock acquisition
& atomic

Prevent: needing locks & all
(prevent mut. excl.)

=> "lock free" concurrent
("wait free")

idea: how to build
conc. data structs
w/o locks

(use powerful

low-level (instructions)

Example:

crit section: (data race)

~~stack~~
=> count = count + 1;

~~lookup~~

powerful inst: fetch-and-add

fetch-and-add(count, 1);
(atomic h/w inst.)

Example 2: list insertion

```
void list-insert (int value) {  
  node_t *n = malloc( );  
  n->value = value;  
  n->next = head;  
  head = n;  
}
```

use CAS } needs to be atomic
compare-and-

h/w instruction: swap

```
int CAS(unsigned *addr,
         unsigned expected,
         unsigned value) {
    if (*addr == expected) {
        *addr = value;
        return 1; // success
    }
    return 0; // failure
}
```

```
list_insert(int value) {
    node_t *n = malloc(
        sizeof(node_t));
    n->value = value;
    do {
        node_t *old-head = head;
        n->next = head;
        while (CAS(&head, old-head,
            (n) == 0));
    }
}
```

potential problem: live lock
(contention)

Solution:

if fail: back off

wait (random-time)

try again }
↓

exponentially increasing, w/
each failure

Prevent: no preemption

pthread_mutex_lock()

can't preempt

other primitives:

pthread_mutex_trylock()

succeeds: grabs locks

fails: if lock held
already

worried about deadlock:

$\left[\begin{array}{l} \text{lock}(L_1) \\ \text{lock}(L_2) \end{array} \right] \}$ concern
1, 1

⇓

top:
lock(L₁)
if (trylock(L₂) == FAIL)
 unlock(L₁) ; ← may want
 goto top; to wait

Avoid : scheduling

T₁ : L₁, L₂ T₃ : L₃
T₂ : L₁, L₂

scheduler:

run T₁ to completion,
then run T₂

Detect/Recover :

Dist. Databases

Review Topics : (today)

Old Questions (tomorrow)

Midterm: Chem ← even
Location: Psych ←

student ID is odd

Format: set 7:15pm →
"easy grading" [9:15pm]

=> bring pencil(s)

A ● B ○ } mystery:
A ○ B ● } theme


low: 50% !

[Long]

Strategy:

Shortest Question First

A



↓ ?

→ OS vs H/W: (hardware support) } **someish**

Scheduling:

→ Multi-Level Feedback Queue **a lot**

→ Page-tables + Address Translation **a lot**

→ Multi-level Page Table **a lot, but...**

→ TLBs **some**

→ Conc + Semaphores **some** → **very few** H/W primitive for locks


→ Producers / Consumers → C/Us
nobody, but ... few

Sched + MLFQ:

(many)
Queues:

each queue
has priority

goals: → progression
⇒ no starvation
⇒ want to (interact)
run short jobs
quickly
(no knowledge
about jobs)



→ Highest priority job
always runs.

→ if 2 (or more) jobs have

same priority
round robin

→ Enter: at highest priority

→ if uses time quantum @ given priority level, move down one level queue

→ Periodically: boost priority of all jobs (to topmost level)

what defines time slice length?

→ min: timer interrupt every 1 ms

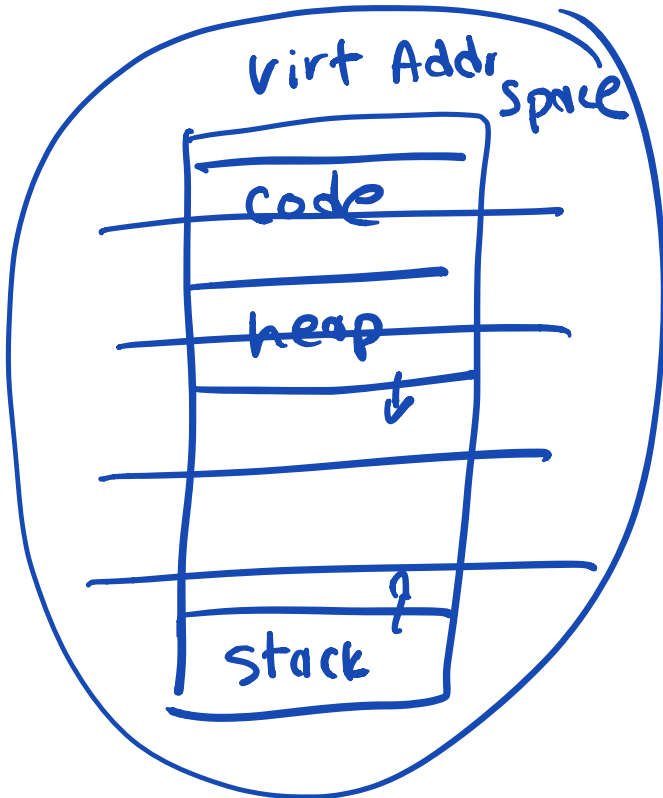
→ for each queue: define time slice is a multiple of timer

interrupt

Virtual Memory

Addr translation
TLBs,

Page tables,
↓ linear (array) ↓ multi-level



running program:
process


private

int x; => address of x
loads, stores
 x = x + 1; ↪ instructions:

hardware support : fetch

- 1) base/bounds
- 2) segmentation (3 base/bounds)

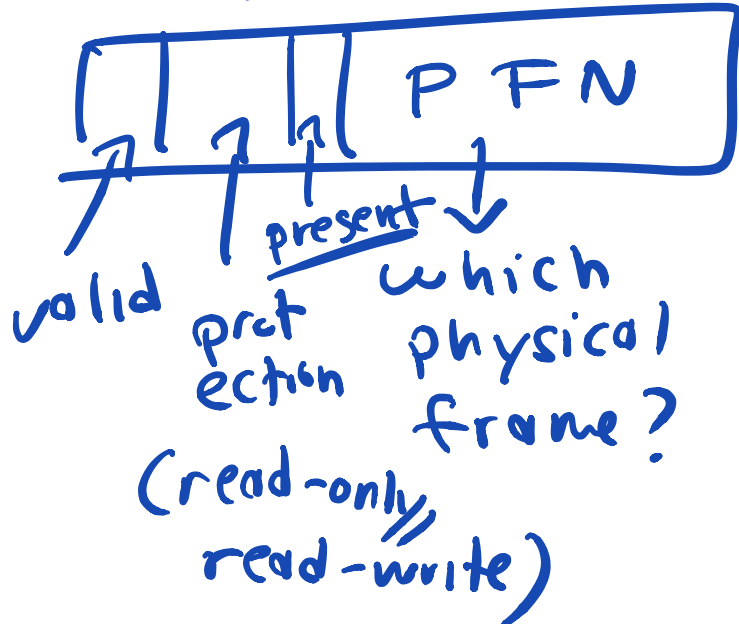
3) paging :



Page table:
 for each virt. page =>
where is this page
in phys memory?
 OR

not valid

Page Table Entry



- 1) Size of page tables => too big
- 2) Speed of paging / translation

linear (array) :
one entry per virt page

32-bit addr space

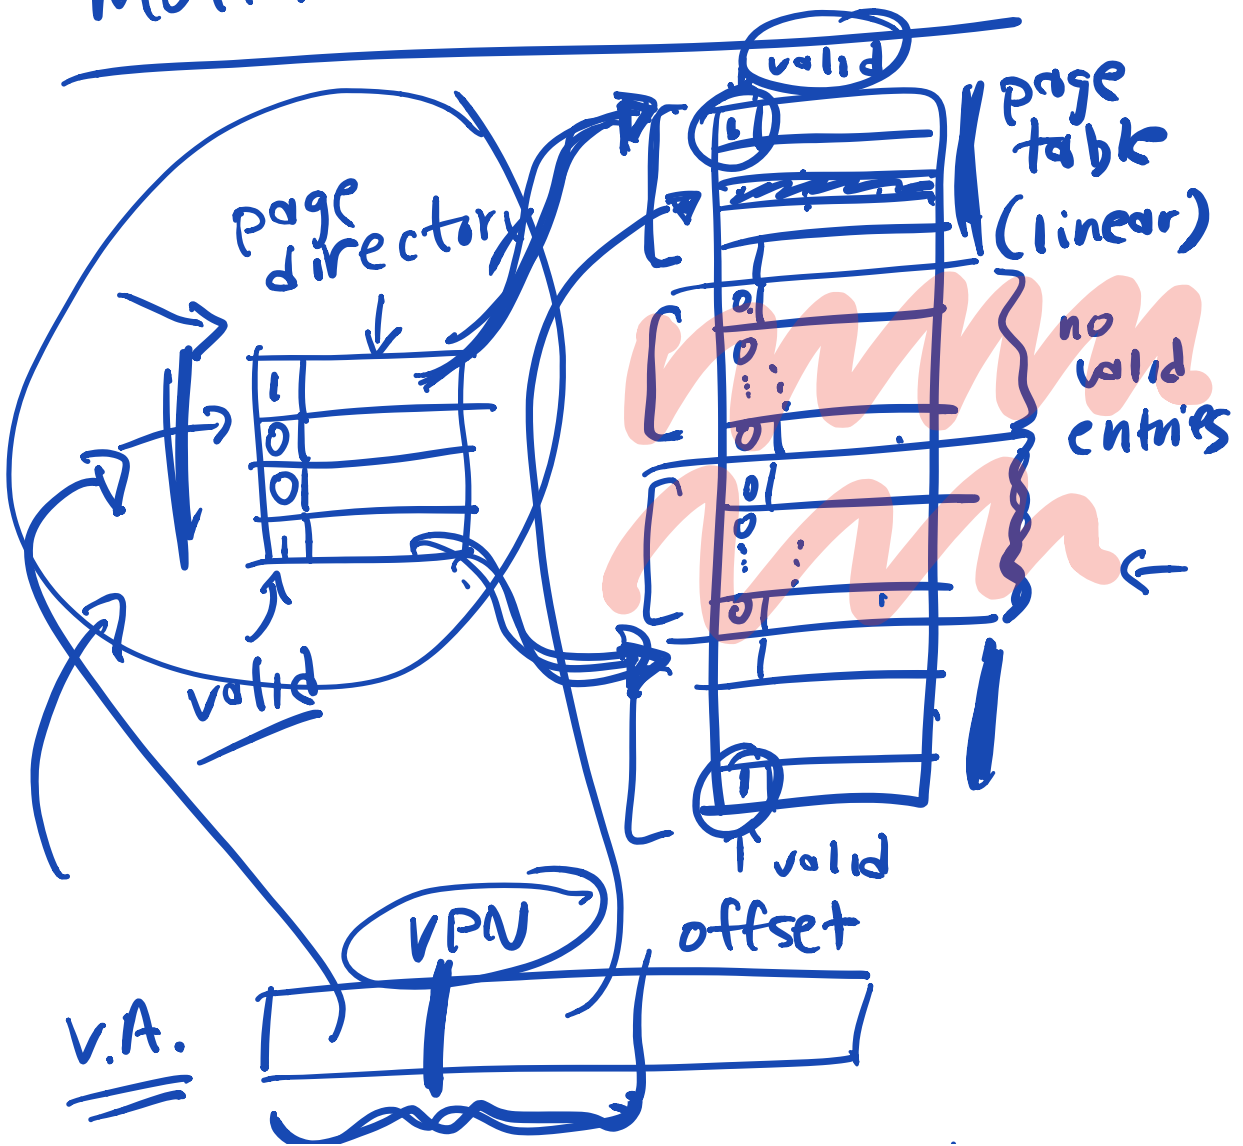
4KB page

4KB ↓


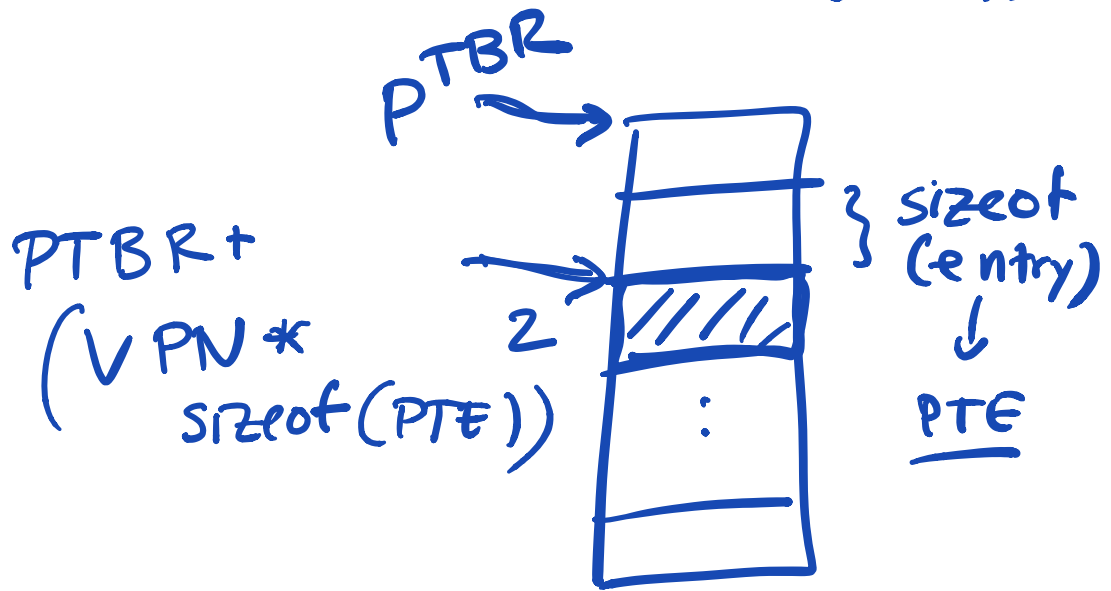



VPN 32-bit
 2^{20} virt pages
~ 1 million entries

Multi-level page table:




linear
page
table
(array)








TLBs: Speed!



w/o TLB: 2 mem refs
for every mem ref

TLB: "Address Translation
Cache"

⇒ more complexity



PA ⇒ memory

⇒ TLB miss:

H/w based:

H/w look up
entry in Page
Table (pointed to
by PTBR)



(CR3)

update TLB

(retry)