

CS 354 - Lecture 12

①

Plan:

1. Cache

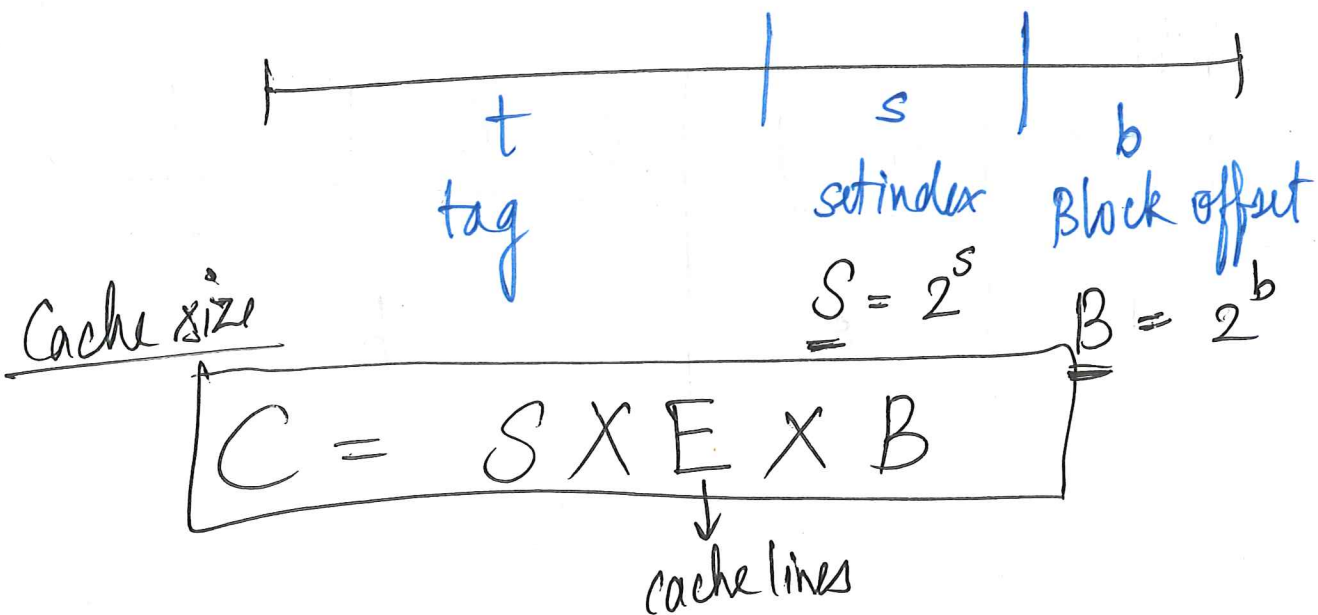
- set associative cache
- fully " "
- Writes
- Cache friendly code.

2. Intro to OS.


- CPU virtualization
- Memory "

Review

m-bit address



(2)




$E = 1$ (direct mapped cache).


$B = 16$ bytes

$$C = S \times E \times B = 4 \times 1 \times 16 = \underline{\underline{64 \text{ bytes}}}$$


addr = 8 bits



Dyn. mem. allocator: (3) Alignment = 8 bytes




malloc(n)




- char - 1
- short - 2
- int - 4
- long long - ~~4~~ or 8

```
int dot_prod(int a[4], int b[4]) {
    int sum;
    for (int i = 0; i < 4; i++)
```


sum += a[i] * b[i];




1) Read $a[0]$ ^④



2) Read $b[0]$



3) Read $a[1]$



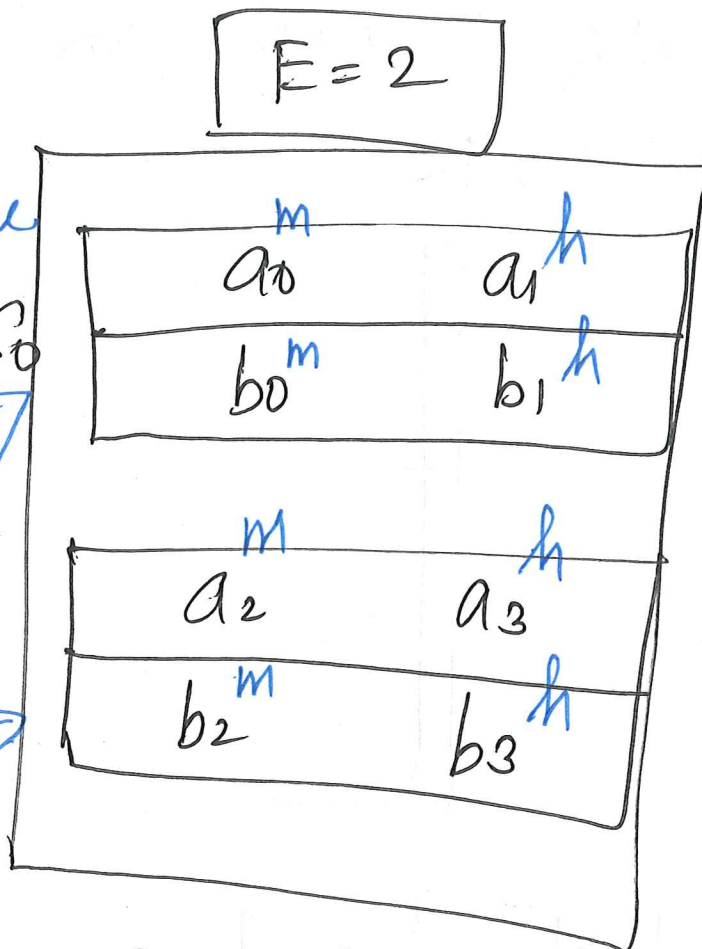
$$\text{miss rate} = \frac{\# \text{ misses}}{\# \text{ refs.}} = \frac{8}{8}$$

$$= \underline{\underline{100\%}}$$

Set associative cache (5)

2-way set associative cache
E

0	a_0	a_1
8	a_0 a_2	a_3
16	b_0	b_1
24	b_2	b_3




$$C = 2 \times 2 \times 8$$
$$= \underline{\underline{32 \text{ bytes}}}$$

$$\text{miss rate} = \frac{4}{8} = \underline{\underline{50\%}}$$

$$1 < E < \frac{C}{B}$$

Fully associative cache (8)

S=8



$$S=1 \quad E=4 \quad B=8$$



$$C = 1 \times 4 \times 8 = \underline{\underline{32 \text{ bytes}}}$$

$$E = \frac{C}{B}$$

"Main memory" is a fully associative cache.



Writes

(7)




Writes

Write through




Write back.



bo' written to K+1 only if eviction happens.

Evictions

(8)



Assume: $x_0 \dots x_7$

$y_0 \dots y_7$

Read $\boxed{x_4 \ x_5}$ \rightarrow maps to S_0 .

What are we going to evict?

$x_0 \ x_1$

$y_0 \ y_1$

Policy : Cache replacement policy.


Random

LRU
Least recently used

LFU
 \downarrow
Frequently..

(9)


LRU



References: $\underline{\underline{b_0}}$ $\underline{\underline{b_1}}$ $\underline{\underline{b_2}}$ $\underline{\underline{b_1}}$ $\underline{\underline{b_0}}$ $\underline{\underline{b_3}}$ $\underline{\underline{b_4}}$ $\underline{\underline{b_1}}$

\underline{m} \underline{m} \underline{m} \underline{h} \underline{h} \underline{m} \underline{m} \underline{m}

LFU



Ref: b_0 b_0 b_0 b_1 b_2 b_1 b_2 | $\underline{\underline{b_3}}$ $\underline{b_3}$ $\underline{b_3}$ $\underline{b_3}$ | $\underline{\underline{b_4}}$

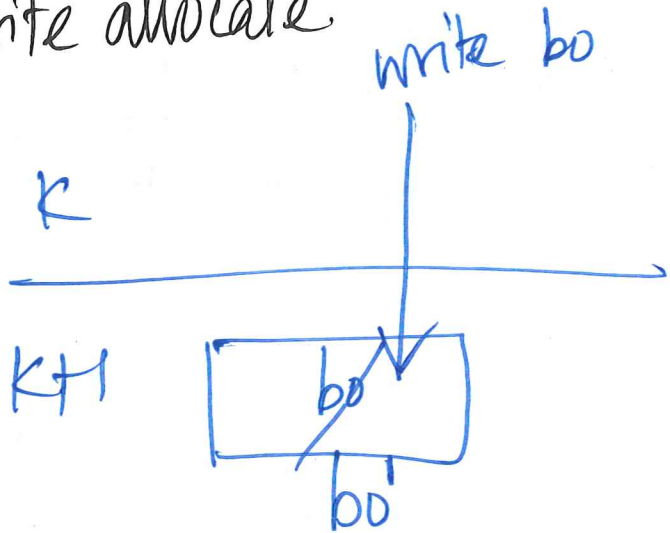
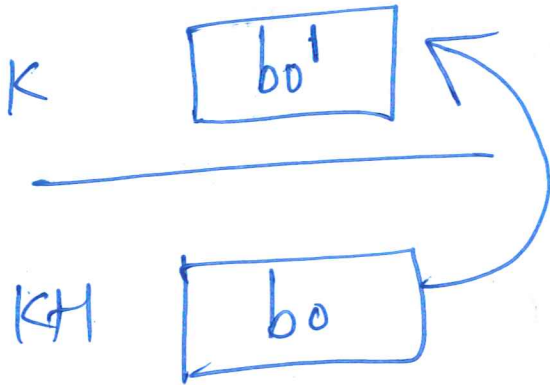
- $\underline{\underline{b_0 - 3}}$
- $\underline{\underline{b_1 - 2}}$
- $\underline{\underline{b_2 - 2}}$
- $\underline{\underline{b_3 - 4}}$

(10)

Writes

Write allocate

No-write allocate



Cache Friendly Code

1. miss rate = $\frac{\# \text{ misses}}{\# \text{ references}}$
2. hit rate = $1 - \text{miss rate}$
3. hit time = $t_{\text{set selection}} + t_{\text{line iden.}} + t_{\text{word iden.}}$
4. miss penalty.

Transpose

$$M = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$$

(src)

$$M^t = \begin{bmatrix} 1 & 3 \\ 2 & 4 \end{bmatrix}$$

(dst)

```
for ( i=0; i < 2; ++i )
    for ( j=0; j < 2; ++j )
        dst [j] [i] = src [i] [j];
```

* src starts at 0; dst starts at 16.

* Cache-direct mapped, write through,

* write allocate

* B = 8 bytes

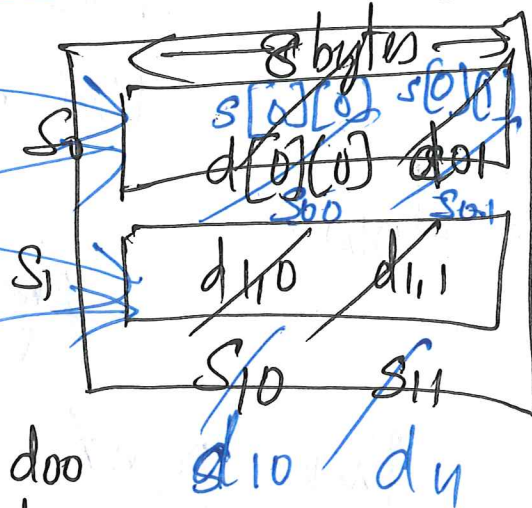
$$C = 16 \text{ bytes}$$

src

0,0 m	0,1 m
1,0 m	1,1 h

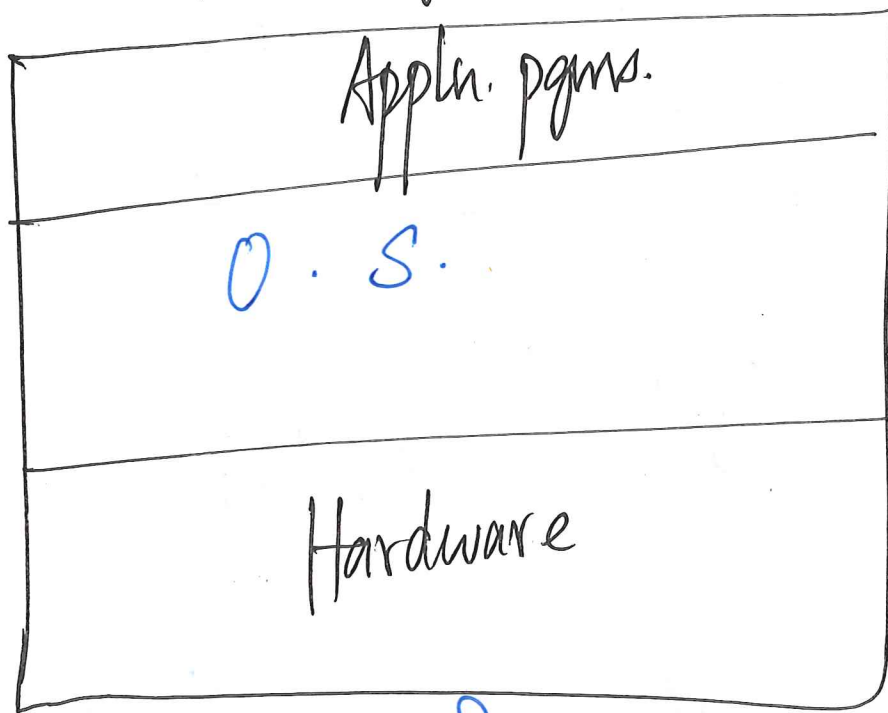
dst

0,0 m	0,1 m
1,0 m	1,1 m



- S00 → d00
- S01 → d10
- S10 → d01
- S11 → d11


Intro to Operating Systems. ⁽¹²⁾



What is an OS?

⇒ software that makes the machine much easier to use.



How?




1 CPU

Run pgm A and pgm B on CPU.

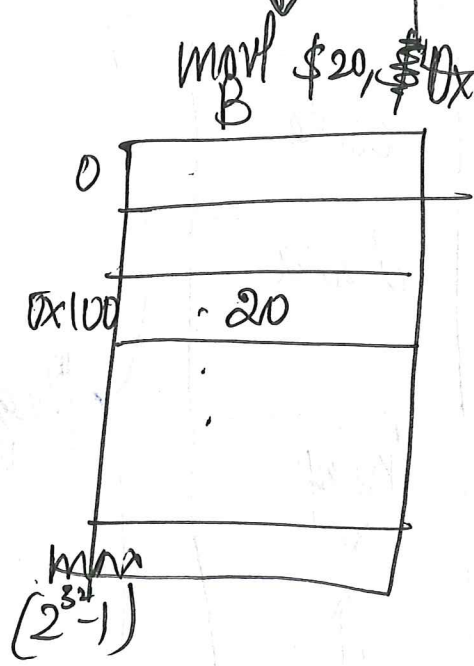
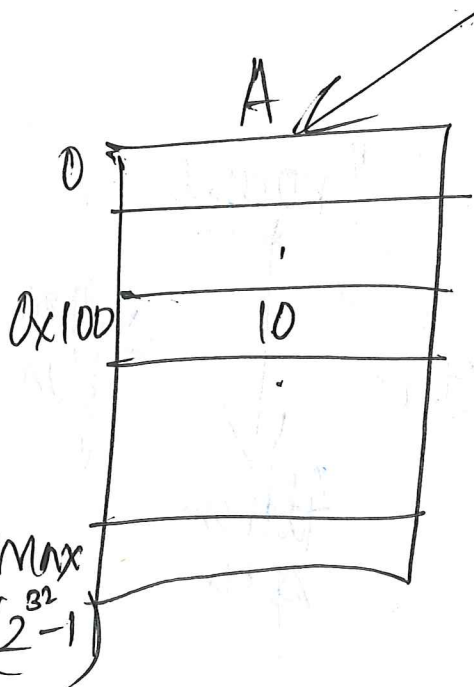


Illusion: virtual CPUs





idea: time-sharing

memory



"memory virtualization"


Virtualization

CPU


memory

Process - basic abstraction.


Memory



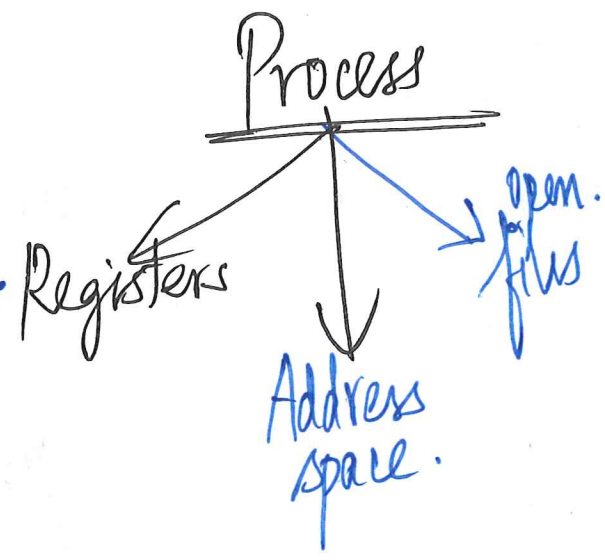

→ ./a.out →




Program



Process



(15)



How do you switch to B?

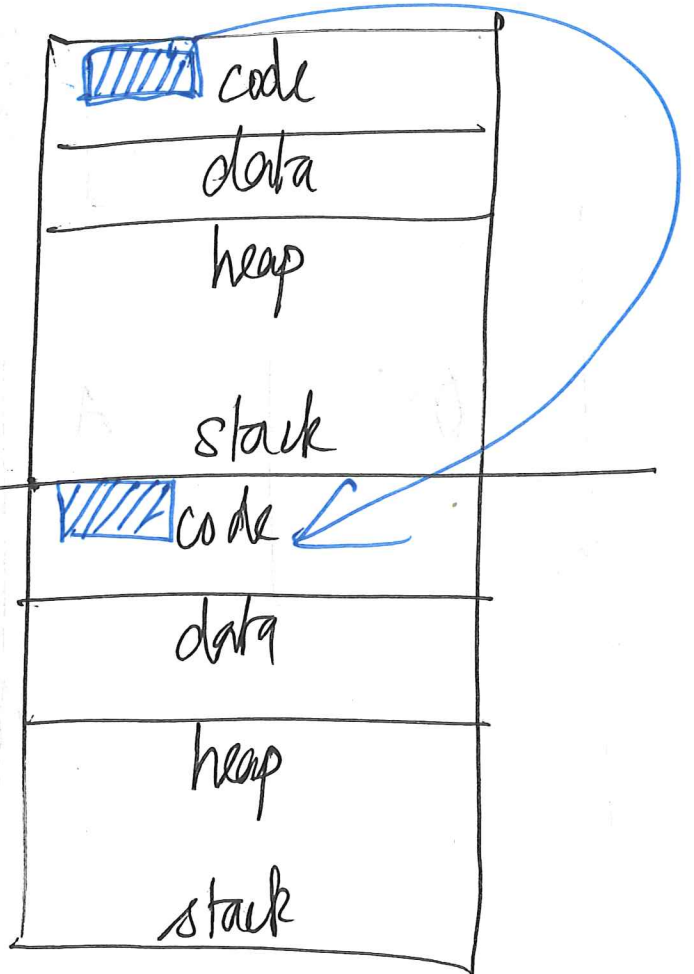


↑
choose
proc A
to run.

yield(); X

while(1)
;

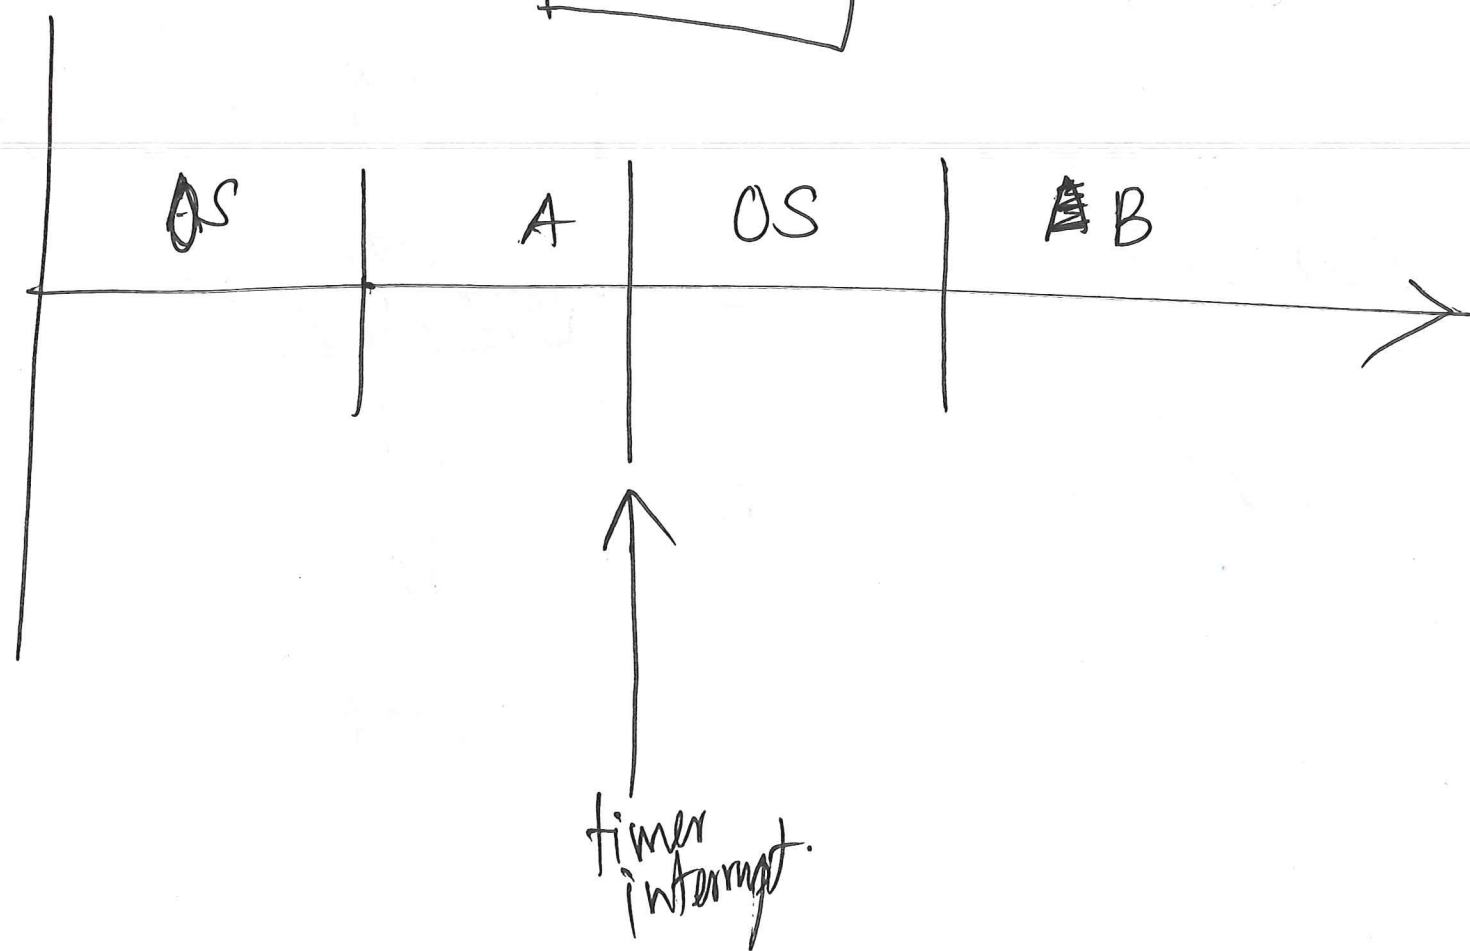
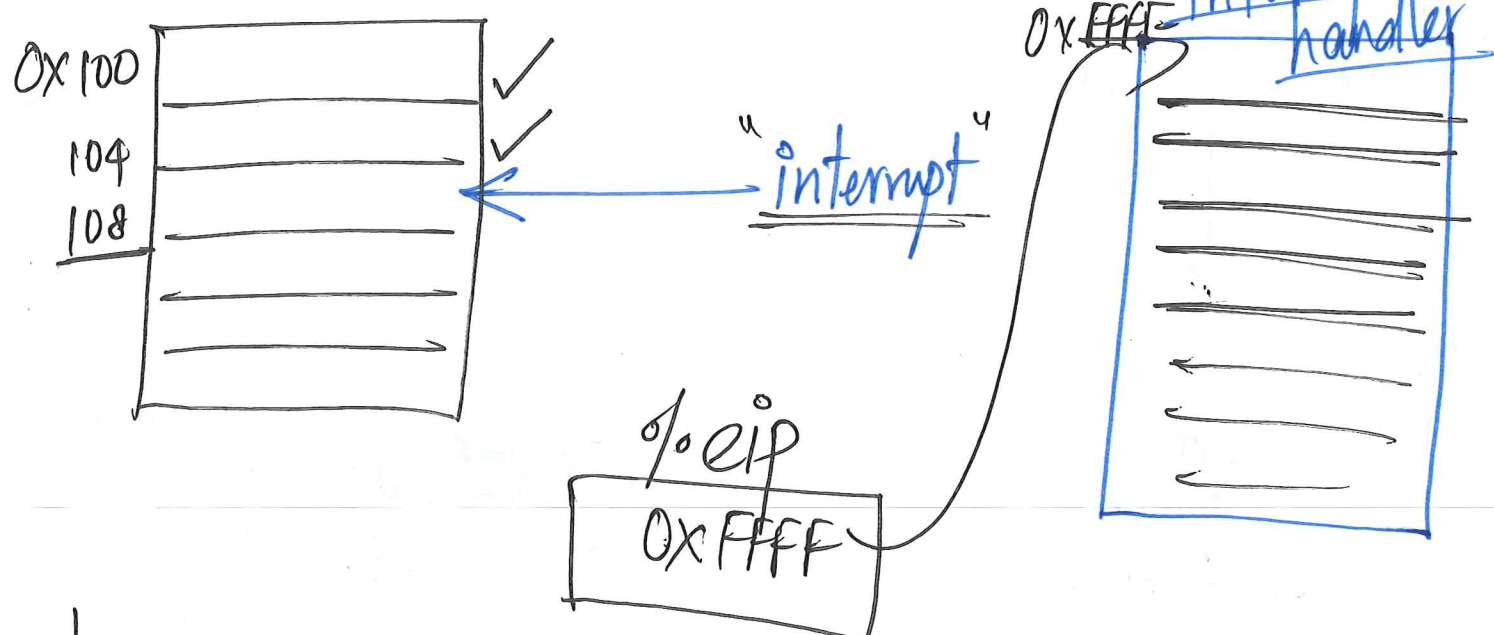
OS



(16)

Hardware support

"timer interrupt" - every 10 m.s.



CPU Virtualization (17)

1. Mechanisms
2. Policies

"timer interrupt"

mode bit

@boot: OS tells the hardware the location of the interrupt handlers.


OS ~~setup~~ starts the timer interrupt.
(10 ms)

timer interrupt: can be switched on/off.
change the addr of the timer interrupt


What if a user process switches off the timer interrupt?

Hardware support


"privilege mode" (status register).




→ user mode



→ "kernel" mode



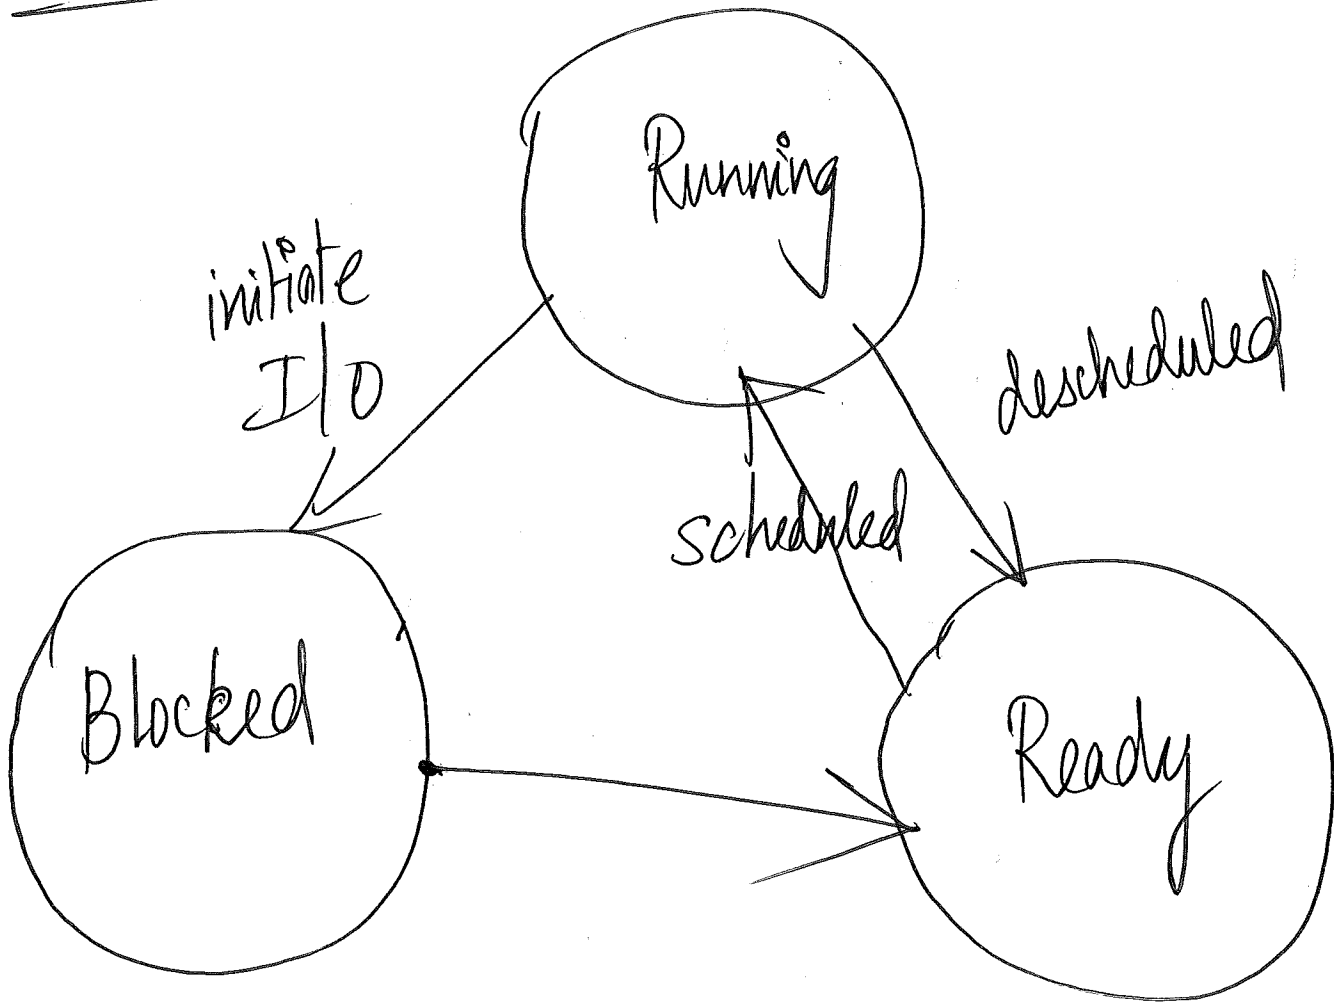
not a normal jump



jumps to user code
 +
~~decreases~~ the privilege level to ~~kernel~~ user
~~raises~~

Process states

(19)



P ₀	P ₁	P ₂	P ₃				
Ready	Running	Ready	Blocked				

A

OS

B



} "struct" proc
Process Control Block
(PCB)



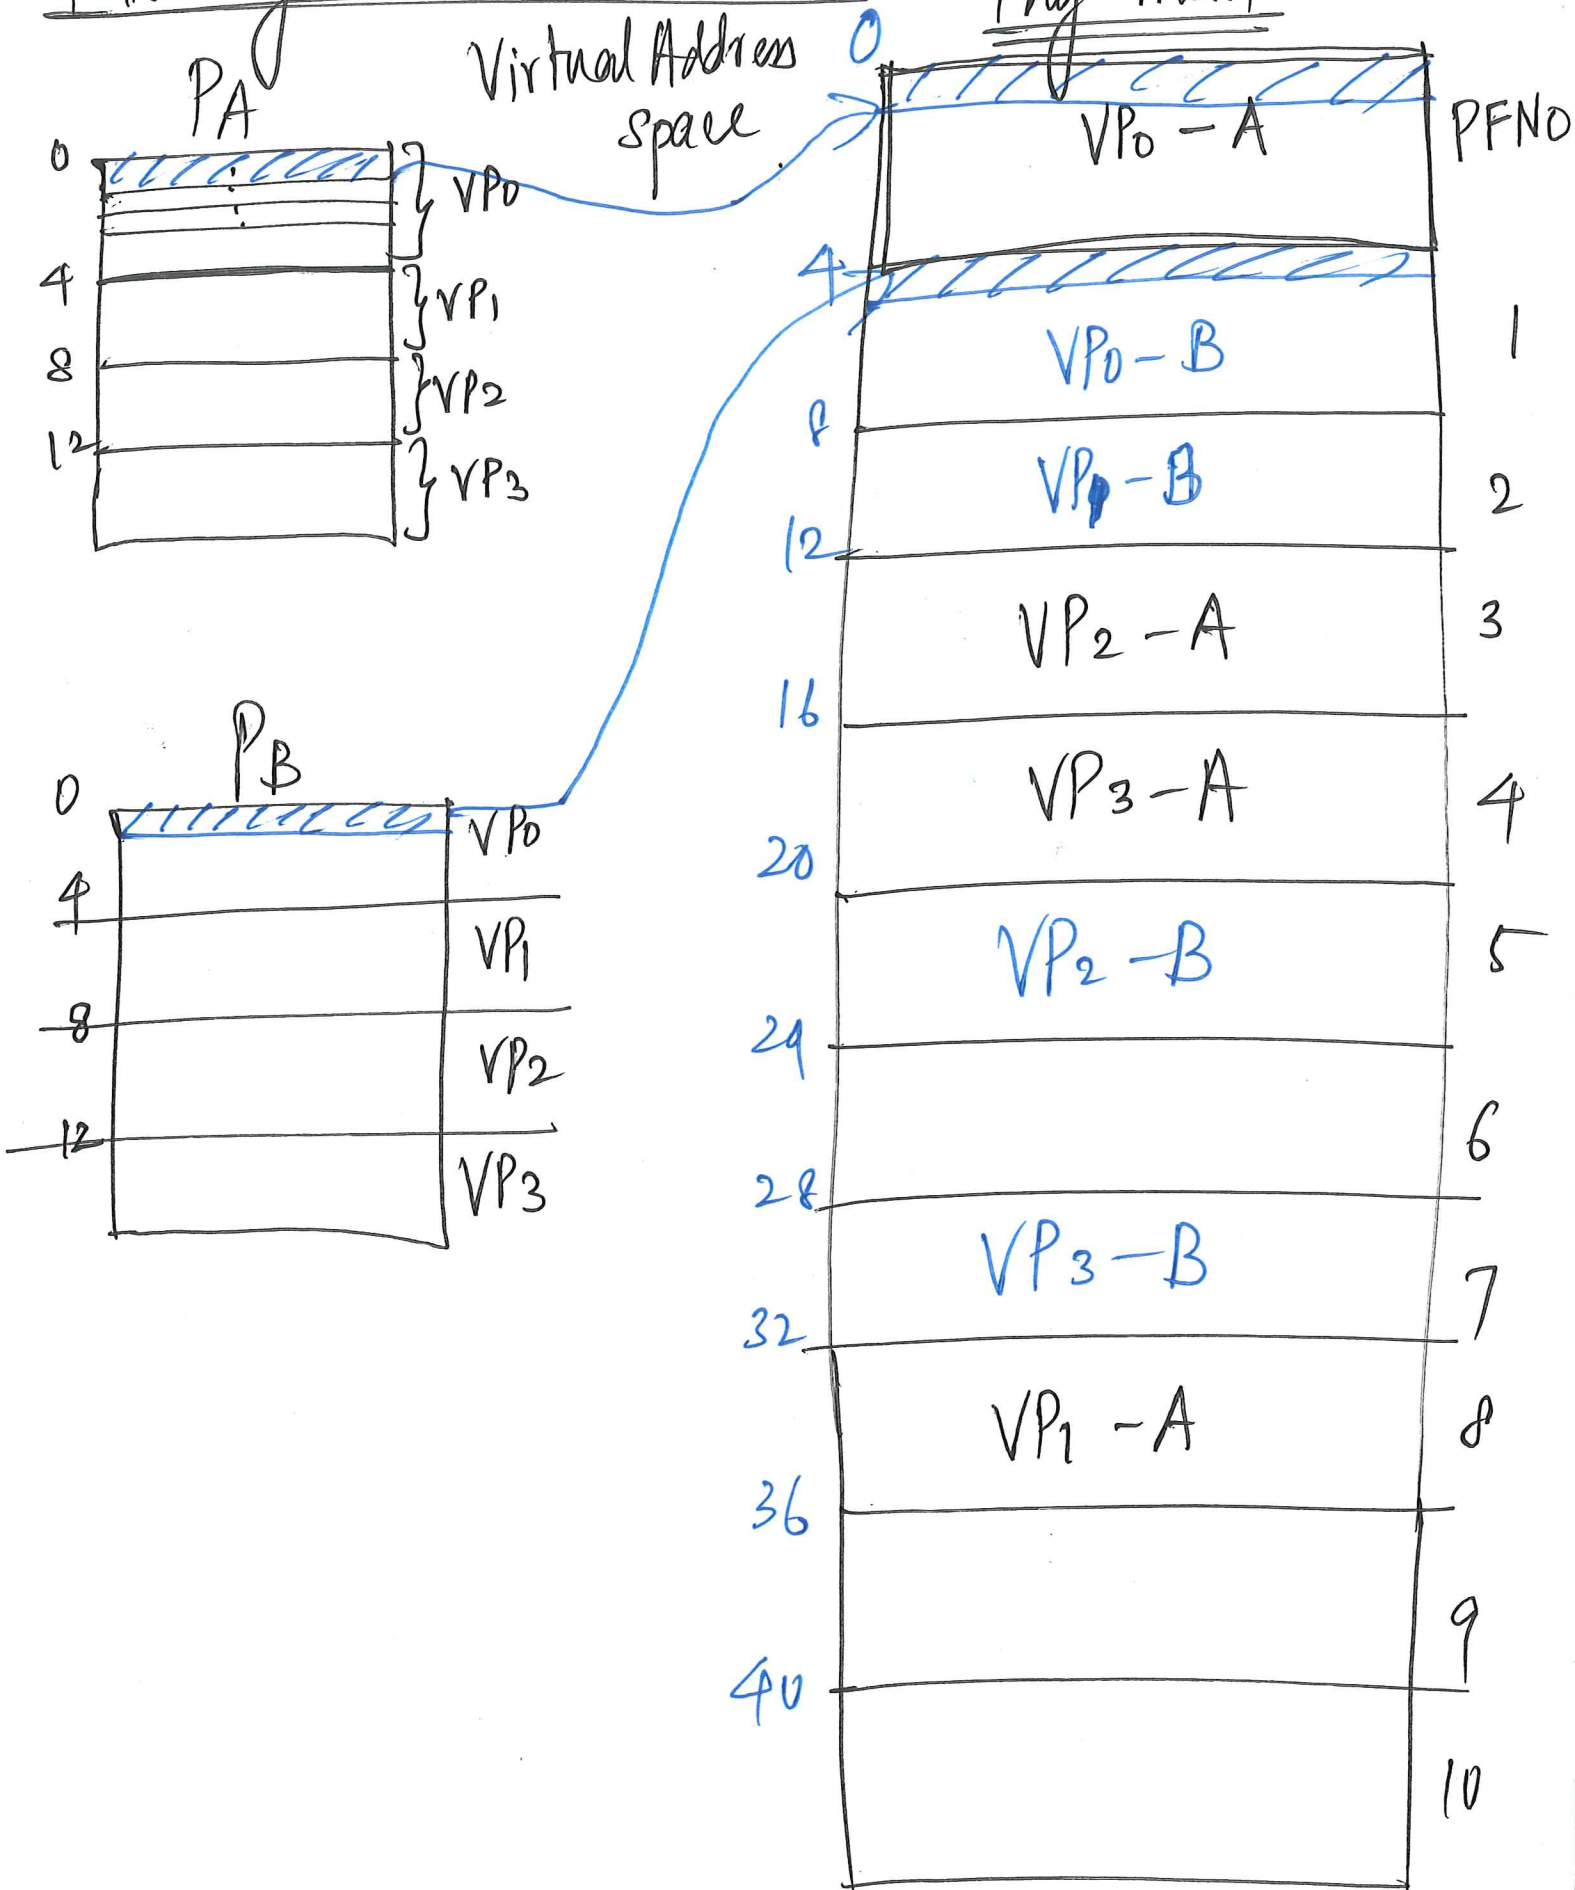
Context Switch

1. Save the state of PA.
2. Restore the state of PB
(that was previously saved).



(21)

Memory Virtualization



CS 354: Intro to Computer Systems (Spring 2018)

Lecture 13 - Set Associative Caches

The following problem concerns basic cache lookups.

- The memory is byte addressable.
- Memory accesses are to 1-byte words (not 4-byte words).
- Physical addresses are 12 bits wide.
- The cache is 4-way set associative, with a 2-byte block size and 32 total lines.

$E = 4$
 $B = 2$
 $S = 8$


In the following tables, all numbers are given in hexadecimal. The contents of the cache are as follows:

4-way Set Associative Cache																
Index	Tag	Valid	Byte 0	Byte 1	Tag	Valid	Byte 0	Byte 1	Tag	Valid	Byte 0	Byte 1	Tag	Valid	Byte 0	Byte 1
0	29	0	34	29	87	0	39	AE	7D	1	68	F2	8B	1	64	38
1	F3	1	0D	8F	3D	1	0C	3A	4A	1	A4	DB	D9	1	A5	3C
2	A7	1	E2	04	AB	1	D2	04	E3	0	3C	A4	01	0	EE	05
3	<u>3B</u>	<u>0</u>	AC	1F	<u>E0</u>	0	B5	70	<u>3B</u>	<u>1</u>	<u>66</u>	95	37	1	49	F3
4	<u>80</u>	<u>1</u>	60	35	<u>2B</u>	0	19	57	<u>49</u>	<u>1</u>	8D	0E	00	0	70	AB
5	EA	1	B4	17	CC	1	67	DB	8A	0	DE	AA	18	1	2C	D3
6	1C	0	3F	A4	01	0	3A	C1	F0	0	20	13	7F	1	DF	05
7	<u>0F</u>	<u>0</u>	00	FF	<u>AF</u>	1	B1	5F	<u>99</u>	<u>0</u>	AC	96	<u>3A</u>	<u>1</u>	22	79

Part 1

The box below shows the format of a physical address. Indicate (by labeling the diagram) the fields that would be used to determine the following:

- CO* The block offset within the cache line
- CI* The cache index
- CT* The cache tag



Part 2

For the given physical address, indicate the cache entry accessed and the cache byte value returned in hex. Indicate whether a cache miss occurs.

If there is a cache miss, enter "-" for "Cache Byte returned".

Physical address: 0x3B6

Physical address format (one bit per box)



Physical memory reference

Parameter	Value
Cache Offset (CO)	0x0
Cache Index (CI)	0x3
Cache Tag (CT)	0x3B
Cache Hit? (Y/N)	Y
Cache Byte returned	0x66

set #3

tag = 00111011
0x3B

Part 3

In the 4-way Set Associative Cache given above, list all the hex memory addresses that will hit in Set 7.

0xAF | 111_

0xAFE, 0xAFF

0x3AE, 0x3AF