

*Welcome back !!*

# MEMORY: PAGING AND TLBS

Shivaram Venkataraman

CS 537, Fall 2024

# ADMINISTRIVIA

- Project 2 done?! → *deadline today*
- Project 3 will be out! Start early? → *2 weeks*
- Project 1 grades →
- Midterm conflicts?

# AGENDA / LEARNING OUTCOMES

## Memory virtualization

What is paging and how does it work?

What are some of the challenges in implementing paging?

**RECAP**

# MEMORY VIRTUALIZATION

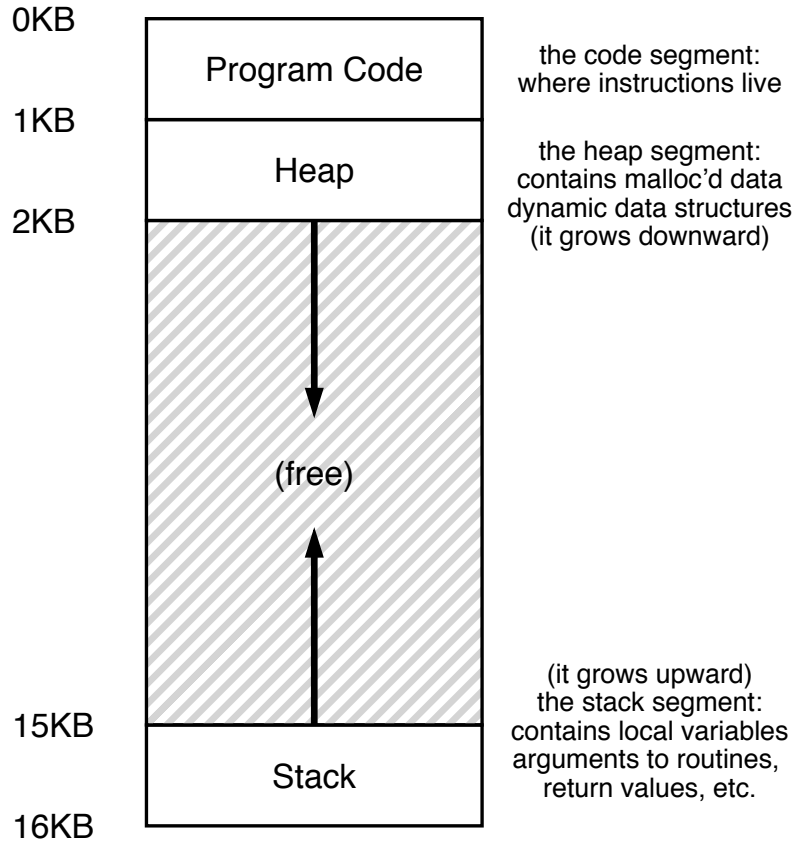
**Transparency:** Process is unaware of sharing

**Protection:** Cannot corrupt OS or other process memory

**Efficiency:** Do not waste memory or slow down processes

**Sharing:** Enable sharing between cooperating processes

# RECAP: WHAT IS IN ADDRESS SPACE?



Static: Code and some global variables

Dynamic: Stack and Heap


offset 12 bits  
 min 0  
 max  $2^{12} - 1$

# REVIEW: SEGMENTATION

max segment size = 4K

Divide address space into logical segments

14 bit virtual address



Each segment corresponds to logical entity in address space


(code, stack, heap)

2 bits to select a segment

Each segment has separate base + bounds register


How does process designate a particular segment?

- Top bits of logical address select segment
- Low bits of logical address select offset within segment



# EXAMPLE: SEGMENTATION

virtual address



%rip: 0x0010

Seg	Base	Bounds
0	<u>0x4000</u>	0xfff
1	<u>0x5800</u>	0xfff
2	0x6800	0x7ff

CPU

1. Fetch instruction at logical addr 0x0010

$$\text{Physical addr: } 0x4000 + 0x0010 \\ = 0x4010$$

2. Exec, load from logical addr 0x1100

$$\text{Physical addr: } 0x5800 + 0x0100 \\ = 0x5900$$



# QUIZ 5!

<https://tinyurl.com/cs537-fa24-q5>





VA space 1K, Physical memory 16K

Segment 0 base (grows positive) : 0x1959 (decimal 6489)  
Segment 0 limit : 316

*limit < 411*

Segment 1 base (grows negative) : 0x0b05 (decimal 2821)  
Segment 1 limit : 282

Address space with two segments:  
segment 0 (topbit=0) or segment 1 (topbit=1).



Translate

0x019b (decimal: 411):

*segmentation violation*

Segment?

0x0315 (decimal: 789):


Translate

0x0315 (decimal 789)

VA space 1K, Physical memory 16K

Segment 0 base (grows positive) : 0x1959 (decimal 6489)

Segment 0 limit : 316



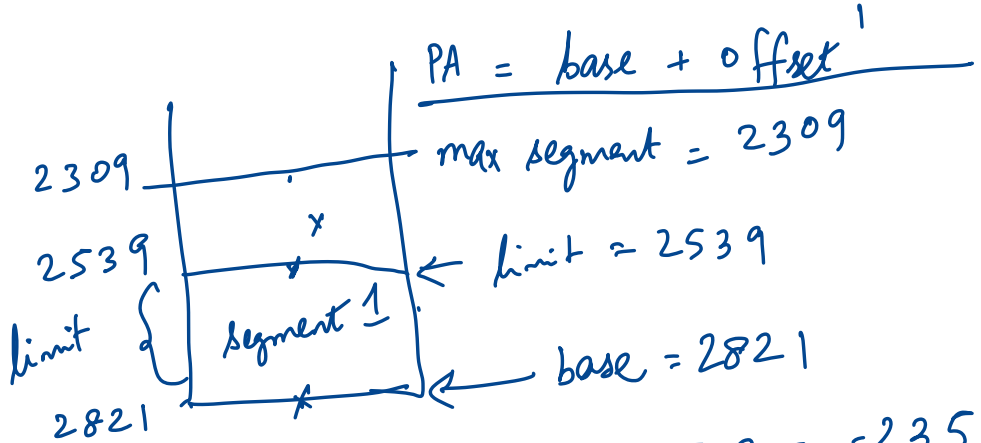
Segment 1 base (grows negative) : 0x0b05 (decimal 2821)

Segment 1 limit : 282

$$\text{offset}' = \text{offset} - \frac{\text{max seg size}}$$

Segment?

0x0315 (decimal: 789):



Translate

0x0315 (decimal 789)

$$277 - 512 = -235$$

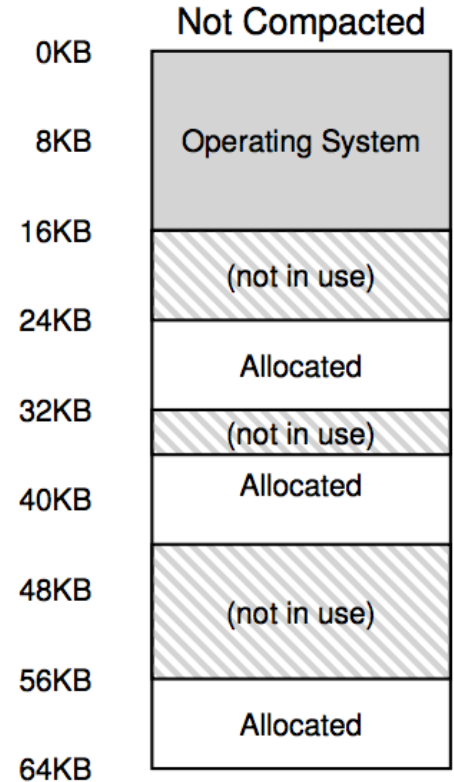
$$PA = 2821 - 235 = 2586$$

# DISADVANTAGES OF SEGMENTATION

Each segment must be allocated contiguously

May not have sufficient physical memory for large segments?

External Fragmentation



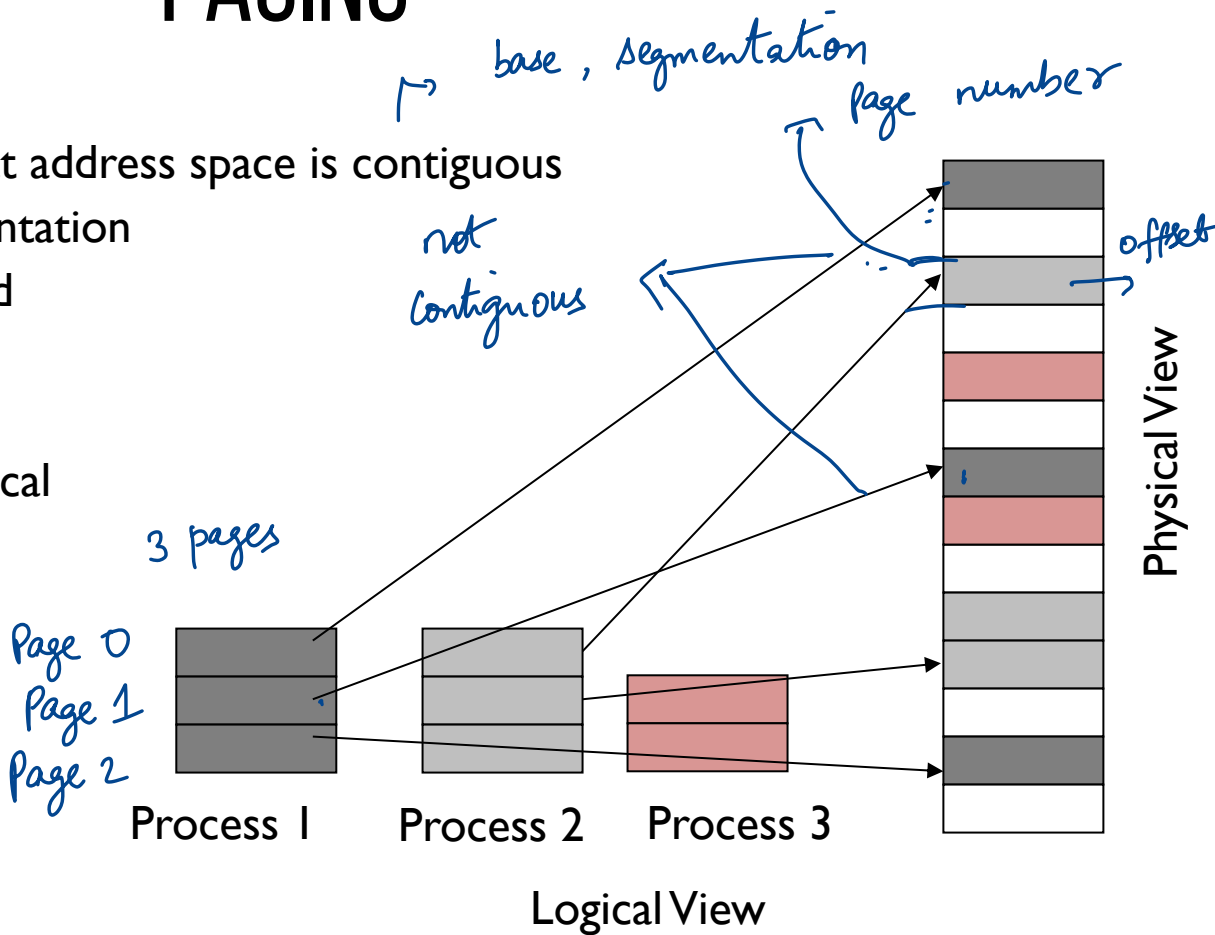
**PAGING**

# PAGING

Goal: Eliminate requirement that address space is contiguous  
Eliminate external fragmentation  
Grow segments as needed

Idea:  
Divide address spaces and physical memory into fixed-sized pages

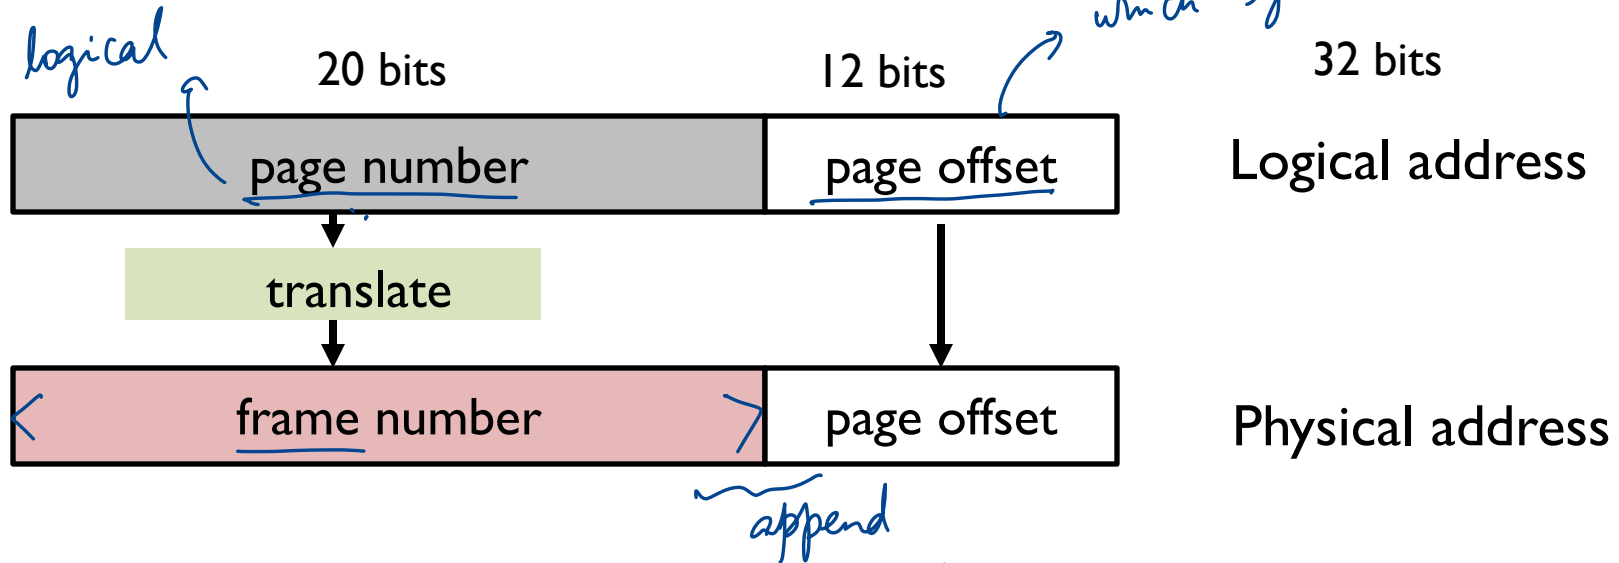
Size:  $2^n$ , Example: 4KB



# TRANSLATION OF PAGE ADDRESSES

How to translate logical address to physical address?

- High-order bits of address designate page number
- Low-order bits of address designate offset within page




No addition needed; just append bits correctly!

# ADDRESS FORMAT

Given known page size, how many bits are needed in address to specify offset in page?

Page Size	Low Bits (offset)
16 bytes	4 bits
1 KB	$\log_2(1 \text{ KB}) = 10 \text{ bits}$
1 MB	20
512 bytes	9
4 KB	12

# ADDRESS FORMAT



Given number of bits in virtual address and bits for offset, how many bits for virtual page number?


Total bits -  
low bits offset

Page Size	Low Bits(offset)	Virt Addr Total Bits	<u>High Bits(vpn)</u>
16 bytes	<u>4</u>	10	6
1 KB	10	20	10
1 MB	20	32	12
512 bytes	9	16	7
4 KB	12	32	20




# ADDRESS FORMAT

6 bits  
vpn



$2^{20}$   
2



Given number of bits for vpn, how many virtual pages can there be in an address space?

Page Size	Low Bits (offset)	Virt Addr Bits	High Bits (vpn)	<u>Virt Pages</u>
16 bytes	4	10	6	$2^6 = 64$
1 KB	10	20	10	$2^{10}$
1 MB	20	32	12	$2^{12}$
512 bytes	9	16	7	$2^7$
4 KB	12	32	20	$2^{20}$


32 bits virtual address =  $2^{20}$  pages, each page 4 KB

# VIRTUAL → PHYSICAL PAGE MAPPING

Number of bits in virtual address

need not equal

number of bits in physical address



6 bits vA

8 bits PA

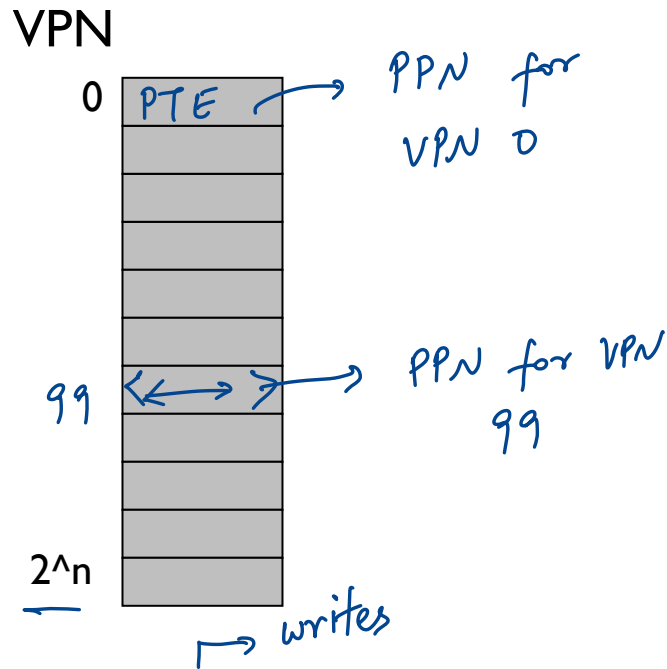
How should OS translate VPN to PPN?

# PAGETABLES


What is a good data structure ?

Simple solution: Linear page table aka *array*

with one entry for each *VPN*  
 contents = *PPN* corresponding to *VPN*




Page Table Entry






PPN

↪ read / write


# PER-PROCESS PAGETABLE



*Page table  
for each  
process*



# FILL IN PAGETABLE



Page Tables:

P1

3
1
7
10

P2

0
4
2
6

P3

8
5
9
11

# WHERE ARE PAGETABLES STORED?


Implication: Store each page table in memory

Hardware finds page table base with register (e.g., CR3 on x86)

What happens on a context-switch?

Change contents of page table base register to newly scheduled process

Save old page table base register in PCB of descheduled process



# OTHER PAGETABLE INFO

What other info is in pagetable entries besides translation?

- valid bit
- protection bits → *R/W or execute*
- present bit (needed later)
- reference bit (needed later)
- dirty bit (needed later)

Pagetable entries are just bits stored in memory

- Agreement between HW and OS about interpretation

# MEMORY ACCESSES WITH PAGING

14 bit addresses


0x0010: movl 0x1100, %edi

Assume PT is at phys addr 0x5000

Assume PTE's are 4 bytes

Assume 4KB pages

How many bits for offset? 12



Fetch instruction at logical addr 0x0010

Access page table to get ppn for vpn 0

Mem ref 1: 0x5000 → read PTE

Learn vpn 0 is at ppn 2

Fetch instruction at \_\_\_\_\_ (Mem ref 2)

0x2010  
append



# MEMORY ACCESSES WITH PAGING

14 bit addresses


0x0010: movl 0x1100, %edi

Assume PT is at phys addr 0x5000

Assume PTE's are 4 bytes

Assume 4KB pages

How many bits for offset? 12



Exec, load from logical addr 0x1100

Access page table to get ppn for vpn 1

Mem ref 3:  $0x5000 + \text{size(PTE)} \times 1$   
 $= 0x5004$

Learn vpn 1 is at ppn 0

Movl from \_\_\_\_\_ into reg (Mem ref 4)

0x0100

2 mem access for

1 read

# MEMORY ACCESSSES WITH PAGING

14 bit addresses

```
0x0010: movl 0x1100, %edi
```

Assume PT is at phys addr 0x5000

Assume PTE's are 4 bytes

Assume 4KB pages

How many bits for offset? 12

Simplified view  
of page table

2
0
80
99

Fetch instruction at logical addr 0x0010

Access page table to get ppn for vpn 0

Mem ref 1: 0x5000

Learn vpn 0 is at ppn 2

Fetch instruction at 0x2010 (Mem ref 2)

Exec, load from logical addr 0x1100

Access page table to get ppn for vpn 1

Mem ref 3: 0x5004

Learn vpn 1 is at ppn 0

Movl from 0x0100 into reg (Mem ref 4)


# PROS/CONS OF PAGING

No external fragmentation

Any page can be placed in any frame in physical memory

Fast to allocate and free

- Alloc: No searching for suitable free space
- Free: Doesn't have to coalesce with adjacent free space



*memory allocator*  
Internal fragmentation → *4 KB why not 1MB?*

- Page size may not match process needs
- Wasted memory grows with larger pages

Additional memory reference to page table →

- Page table must be stored in memory
- MMU stores only base address of page table

Storage for page tables may be substantial

- Requires PTE for all pages in address space
- Entry needed even if page not allocated?

# SUMMARY: PAGE TRANSLATION STEPS

For each mem reference:

*bit shifting* →

1. extract **VPN** (virt page num) from **VA** (virt addr)

2. calculate addr of **PTE** (page table entry)

3. read **PTE** from memory → *expensive*

4. extract **PFN** (page frame num)

5. build **PA** (phys addr)

6. read contents of **PA** from memory into register *expensive*

Which steps are expensive?

# EXAMPLE: ARRAY ITERATOR

*integer 4 bytes*

```
int sum = 0;
for (i=0; i<N; i++){
    sum += a[i];
}
```

Assume 'a' starts at 0x3000  
Ignore instruction fetches  
and access to 'i'

What virtual addresses?

*a[0]* load 0x3000

*a[1]* load 0x3004

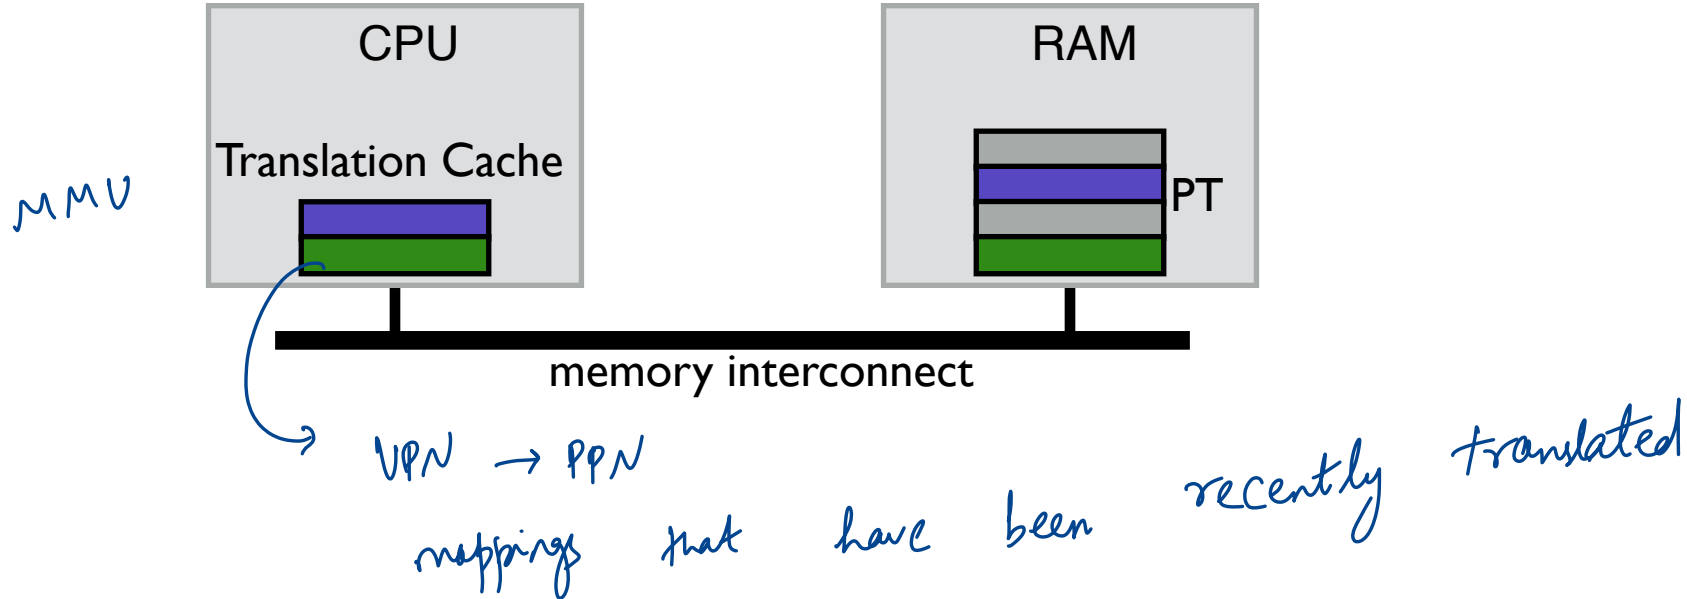
*a[2]* load 0x3008

load 0x300C

What physical addresses?

load 0x100C *↖ PTE*  
load 0x7000 *→ PA*  
load 0x100C *→ PTE*  
load 0x7004  
load 0x100C *→ PTE*  
load 0x7008  
load 0x100C *→ PTE*  
load 0x700C

# STRATEGY: CACHE PAGE TRANSLATIONS




**TLB: TRANSLATION LOOKASIDE BUFFER**

TLB

# TLB ORGANIZATION

TLB Entry



Fully associative

Any given translation can be anywhere in the TLB

Hardware will search the entire TLB in parallel



# ARRAY ITERATOR (W/ TLB)

```
int sum = 0;
for (i = 0; i < 2048; i++){
    sum += a[i];
}
```

Assume 'a' starts at 0x1000  
Ignore instruction fetches  
and access to 'i'

Assume following virtual address stream:

load 0x1000

load 0x1004

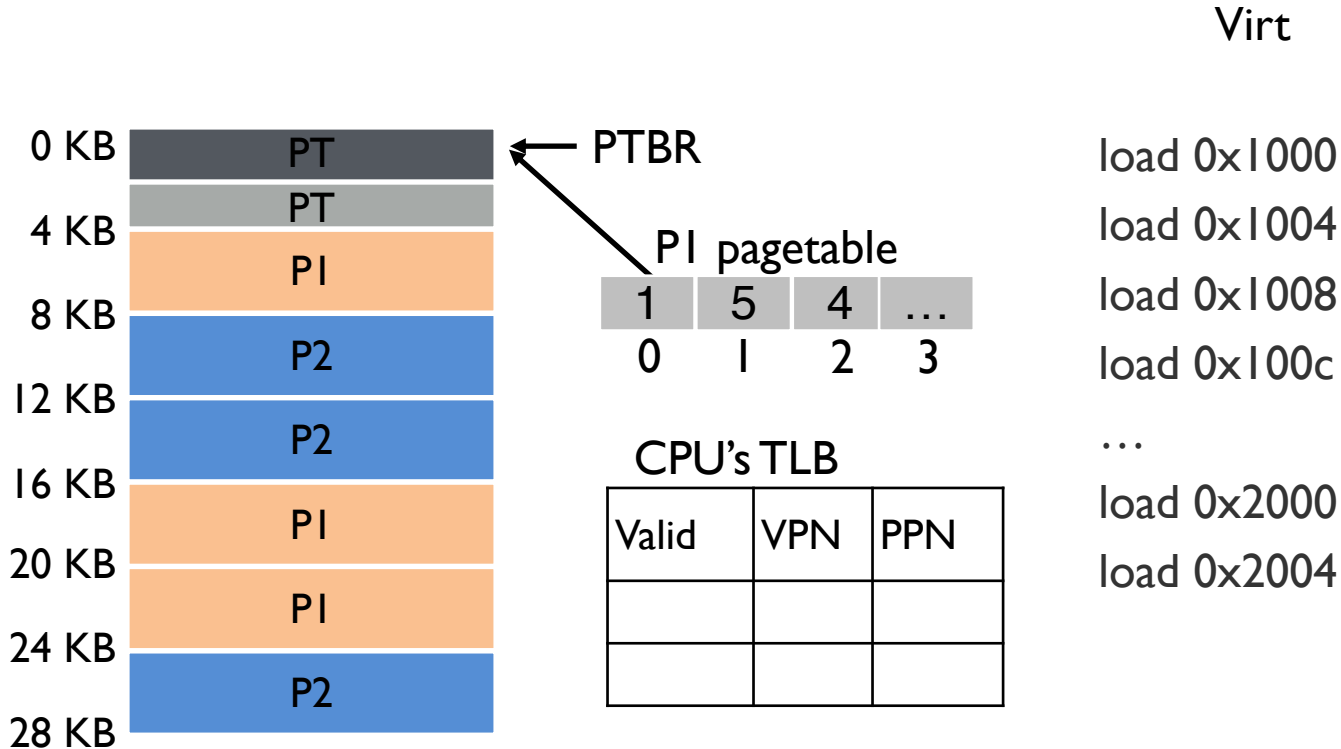
load 0x1008

load 0x100C


...

What will TLB behavior look like?

# TLB ACCESSES: SEQUENTIAL EXAMPLE



# TLB ACCESSES: SEQUENTIAL EXAMPLE



CPU's TLB

Valid	VPN	PPN
1	1	5
1	2	4

Virt	Phys
load 0x1000	load 0x0004
load 0x1004	load 0x5000
load 0x1008	load 0x5004 (TLB hit)
load 0x100c	load 0x5008 (TLB hit)
...	load 0x500c (TLB hit)
load 0x2000	load 0x500c
load 0x2004	...
	load 0x0008
	load 0x4000 (TLB hit)
	load 0x4004

# PERFORMANCE OF TLB?

Miss rate of TLB:  $\# \text{TLB misses} / \# \text{TLB lookups}$

$\# \text{TLB lookups?}$  number of accesses to a = 2048

```
int sum = 0;
for (i=0; i<2048; i++) {
    sum += a[i];
}
```

$\# \text{TLB misses?}$   
= number of unique pages accessed  
= 2048 / (elements of 'a' per 4K page)  
= 2K / (4K / sizeof(int)) = 2K / 1K  
= 2

Miss rate? =  $2/2048 = 0.1\%$

Would hit rate get better or worse  
with smaller pages?

Hit rate?  $(1 - \text{miss rate}) = 99.9\%$

# TLB PERFORMANCE

How can system improve hit rate given fixed number of TLB entries?

Increase page size:

Fewer unique page translations needed to access same amount of memory

TLB Reach: Number of TLB entries \* Page Size

# WORKLOAD ACCESS PATTERNS

## Workload A

```
int sum = 0;
for (i=0; i<2048; i++) {
    sum += a[i];
}
```

Sequential array accesses  
almost always hit in TLB!


## Workload B

```
int sum = 0;
srand(1234);
for (i=0; i<1000; i++) {
    sum += a[rand() % N];
}
srand(1234);
for (i=0; i<1000; i++) {
    sum += a[rand() % N];
}
```

# WORKLOAD ACCESS PATTERNS


Spatial Locality

Sequential Accesses



Temporal Locality

Repeated Random Accesses



# WORKLOAD LOCALITY

**Spatial Locality:** future access will be to nearby addresses

**Temporal Locality:** future access will be repeats to the same data

What TLB characteristics are best for each type?

Spatial:

- Access same page repeatedly; need same vpn  $\rightarrow$  ppn translation
- Same TLB entry re-used

Temporal:

- Access same address near in future
- Same TLB entry re-used in near future
- How near in future? How many TLB entries are there?



# OTHER TLB CHALLENGES

How to replace TLB entries ? LRU ? Random ?

TLB on context switches ? HW or OS ?

# NEXT STEPS

Project 3 out!

Next class: More TLBs and better pagetables!