

MEMORY: PAGING AND TLBS



Shivaram Venkataraman

CS 537, Spring 2020

ADMINISTRIVIA

- Project 1b done!

- Project 2a is out!  Next Friday Feb 14, 10pm

- Reminder: Midterm I makeup

- Discussion section: Process API, Project 2a

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

ABSTRACTION: ADDRESS SPACE

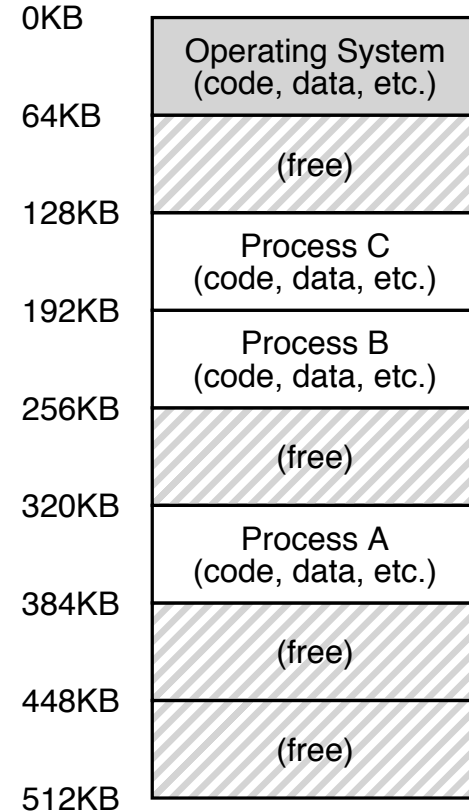
Virtual



Translated

A blue curved arrow pointing from the virtual address space diagram to the physical address space diagram.

Physical



REVIEW: SEGMENTATION

2 bit segment selector
12 bits

14 bit addressing scheme

0x0010: movl 0x1100, %edi
0x0013: addl \$0x3, %edi

per process

%rip: 0x0010

Seg	Base	Bounds
0	0x4000	0xfff
1	0x5800	0xfff
2	0x6800	0x7ff

MMU

1. Fetch instruction at logical addr 0x0010

Physical addr: $0x4000 + 0x010 = 0x4010$

2. Exec, load from logical addr 0x1100

Physical addr: $0x5800 + 0x100 = 0x5900$

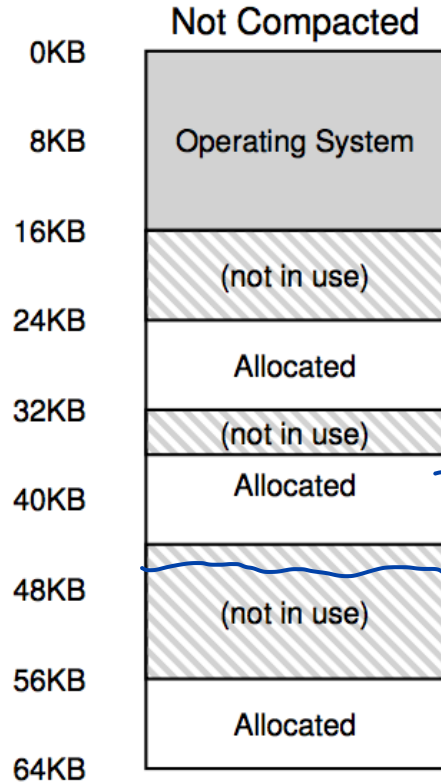
3. Fetch instruction at logical addr 0x0013

Physical addr:

4. Exec, no load

1. Extract segment bits
2. segment table
base register
3. Add base to offset
bits

FRAGMENTATION



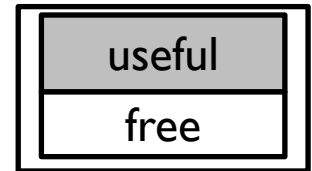
Definition: Free memory that can't be usefully allocated

Types of fragmentation

→ External: Visible to allocator (e.g., OS)

Internal: Visible to requester

Internal



PAGING

PAGING

Goal: Eliminate requirement that address space is contiguous

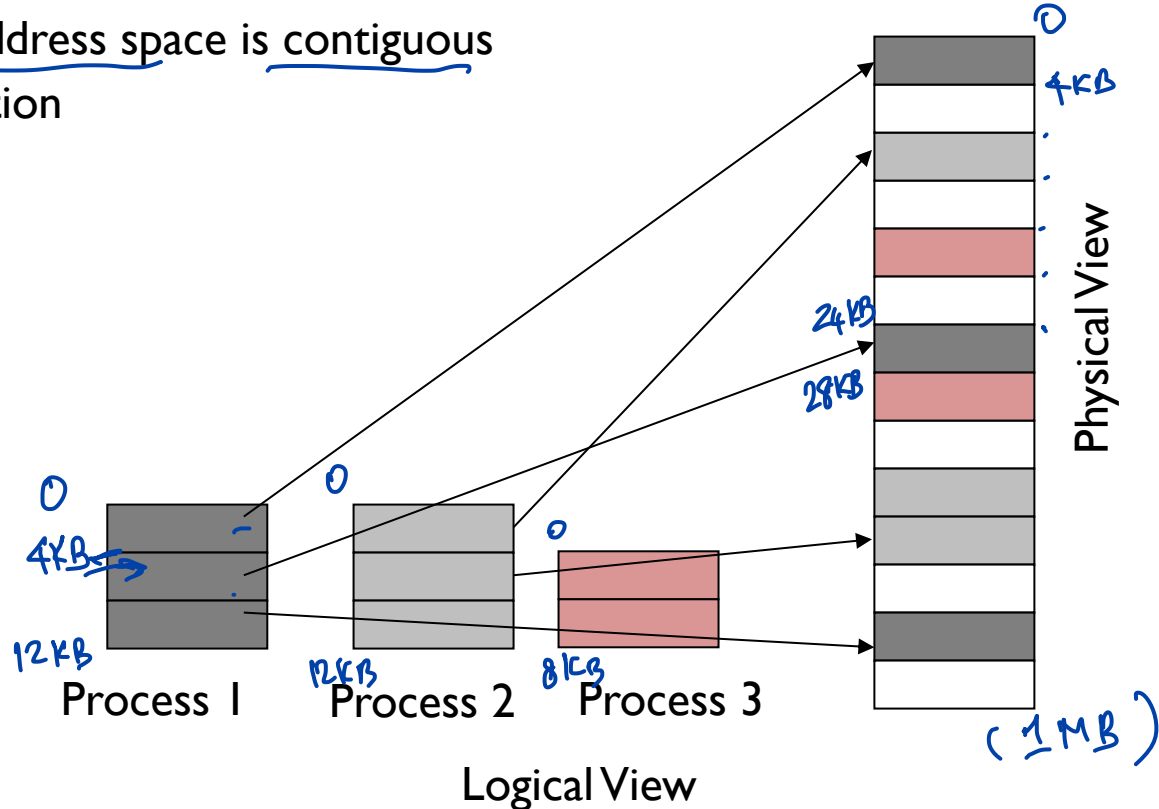
Eliminate external fragmentation

Grow segments as needed

Idea:

Divide address spaces and physical memory into fixed-sized pages

Size: 2^n , Example: 4KB

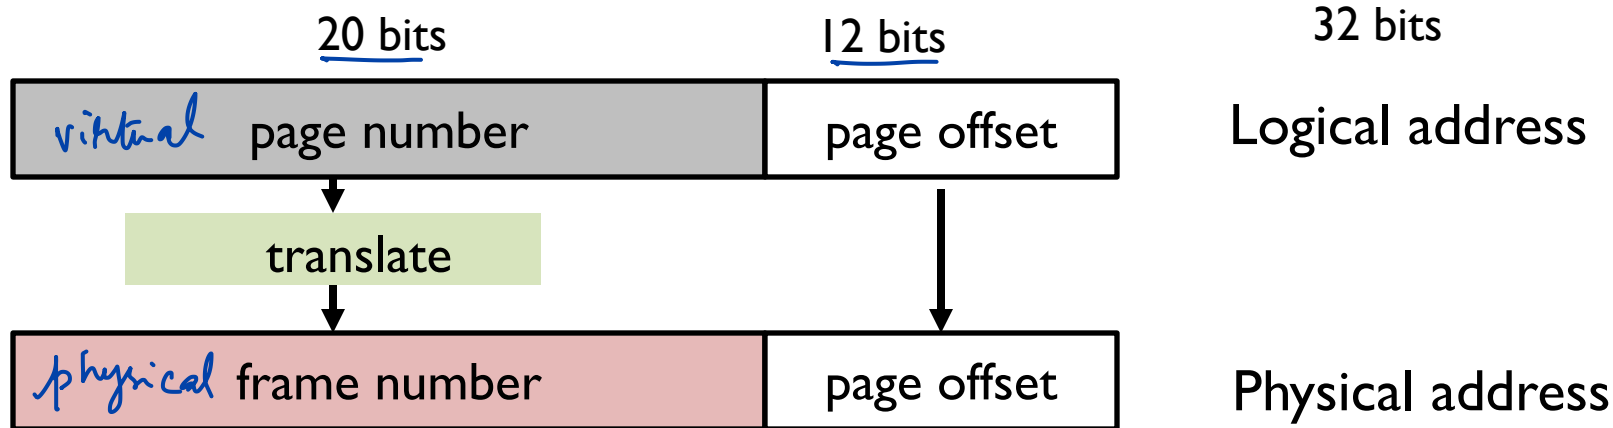


TRANSLATION OF PAGE ADDRESSES

virtual addr → phy addr

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

Page size
fixed configuration

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	10 bits
1 MB	20 bits
512 bytes	9 bits
4 KB	12 bits

1024 x 1 KB

$\log_2 (\text{Page Size})$
= number of bits in offset

ADDRESS FORMAT

virtual addr




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


Page Size

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

Virtual addr



ADDRESS FORMAT



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)	Virt Pages
16 bytes	4	10	6	$2^6 = 64$
1 KB	10	20	10	$2^{10} = 1024$
1 MB	20	32	12	$2^{12} = 4096$
512 bytes	9	16	7	$2^7 = 128$
4 KB	12	32	20	2^{20}


≈ 1 million

VIRTUAL → PHYSICAL PAGE MAPPING


Number of bits in
virtual address

need not equal

number of bits in
physical address



2 bits for
VPN



Physical
memory

How should OS translate VPN to PPN?

→ PAGETABLES

What is a good data structure ?

Simple solution: Linear page table aka *array*


Size of a page table

= Num entries \times size of entry

Page table entry

VPN


0



PPN

2^n

Present/not




PTE

4 bytes



20 bits

permission

PER-PROCESS PAGETABLE



FILL IN PAGETABLE



QUIZ 9

<https://tinyurl.com/cs537-sp20-quiz9>



Description

1. one process uses RAM at a time
2. rewrite code and addresses before running
3. add per-process starting location to virt addr to obtain phys addr
4. dynamic approach that verifies address is in valid range
5. several base+bound pairs per process

Name of approach

Time sharing
Static relocation
Base register
Base + Bounds
Segmentation

Candidates: Segmentation, Static Relocation, Base, Base+Bounds, Time Sharing

QUIZ9: HOW BIG IS A PAGETABLE?

Consider a **32-bit** address space with 4 KB pages. Assume each PTE is 4 bytes

How many bits do we need to represent the **offset** within a page?

$$12 \text{ bits} = \log_2(4\text{KB})$$

How **many virtual pages** will we have in this case?

$$20 \text{ bits VPN} = 2^{20} \approx 1 \text{ million}$$


What will be the **overall size** of the page table?

$$\begin{aligned} \text{Num Virt Page} \times \text{size(PTE)} &= 1 \text{ million} \times 4 \text{ bytes} \\ &= 4 \text{ MB} \end{aligned}$$

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
- 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

2 mem ops
→ Page Table
→ Phy addr

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


2 bits v_{PN}

Simplified view
of page table

2
0
80
99

0x5004 ←

0x5000



Fetch instruction at logical addr 0x0010

Access page table to get ppn for v_{PN} 0

Mem ref 1: read 0x5000

Learn v_{PN} 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 v_{PN} 1

Mem ref 3: 0x5004

Learn v_{PN} 1 is at ppn 0

Movl from 0x0100 into reg (Mem ref 4)

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)

ADVANTAGES 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

Simple to swap-out portions of memory to disk (later lecture)

- Page size matches disk block size
- Can run process when some pages are on disk
- Add “present” bit to PTE

↳ Swapping

DISADVANTAGES OF PAGING

Internal fragmentation: Page size may not match size needed by process

- Wasted memory grows with larger pages
- **Tension?**




Additional memory reference to page table → Very inefficient

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

Storage for page tables may be substantial

4MB

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



SUMMARY: PAGE TRANSLATION STEPS

For each mem reference:

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

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

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

(int is
4 bytes)

What virtual addresses?

load 0x3000

load 0x3004

load 0x3008

load 0x300C

What physical addresses?

load 0x100C

load 0x7000

load 0x100C

load 0x7004

load 0x100C

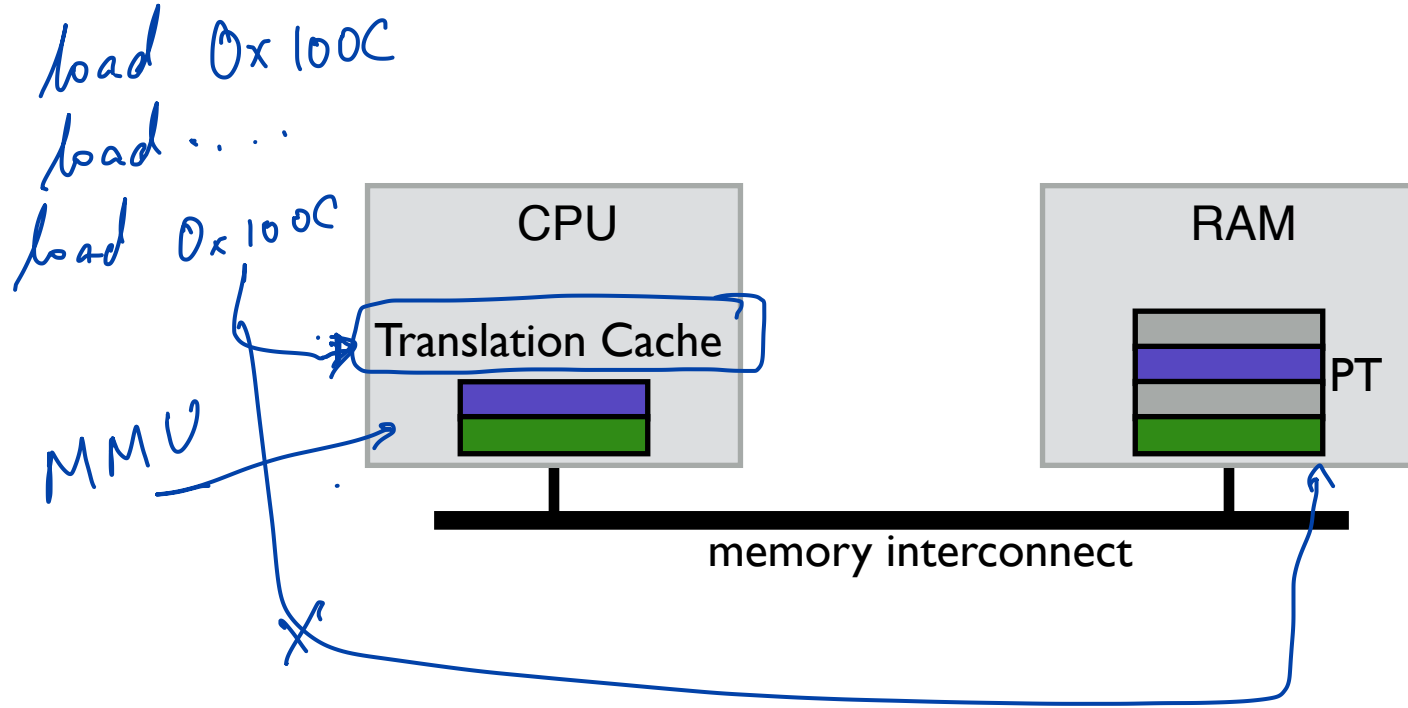
load 0x7008

load 0x100C

load 0x700C

Page table
Phys a[0]
Page table

STRATEGY: CACHE PAGE TRANSLATIONS




TLB: TRANSLATION LOOKASIDE BUFFER

TLB ORGANIZATION

TLB Entry

Tag (virtual page number)	Physical page number (page table 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

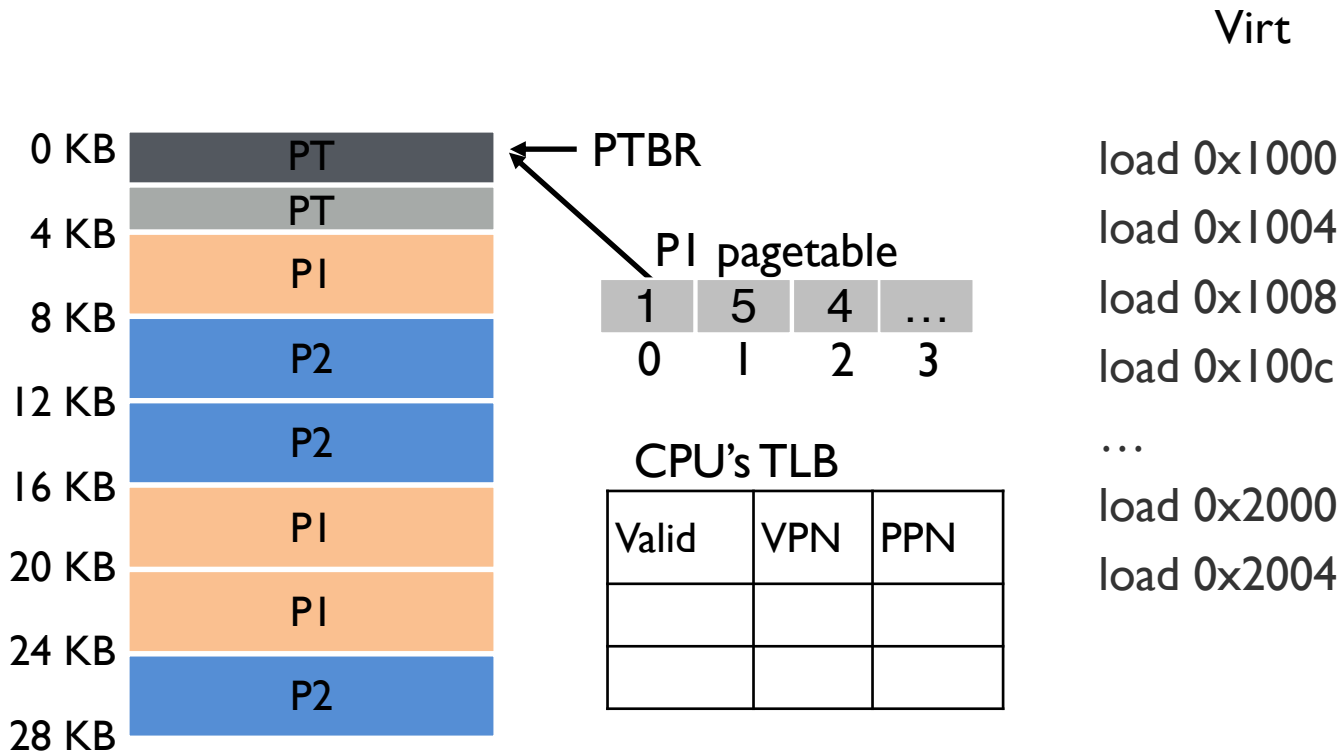
load 0x1008

load 0x100C

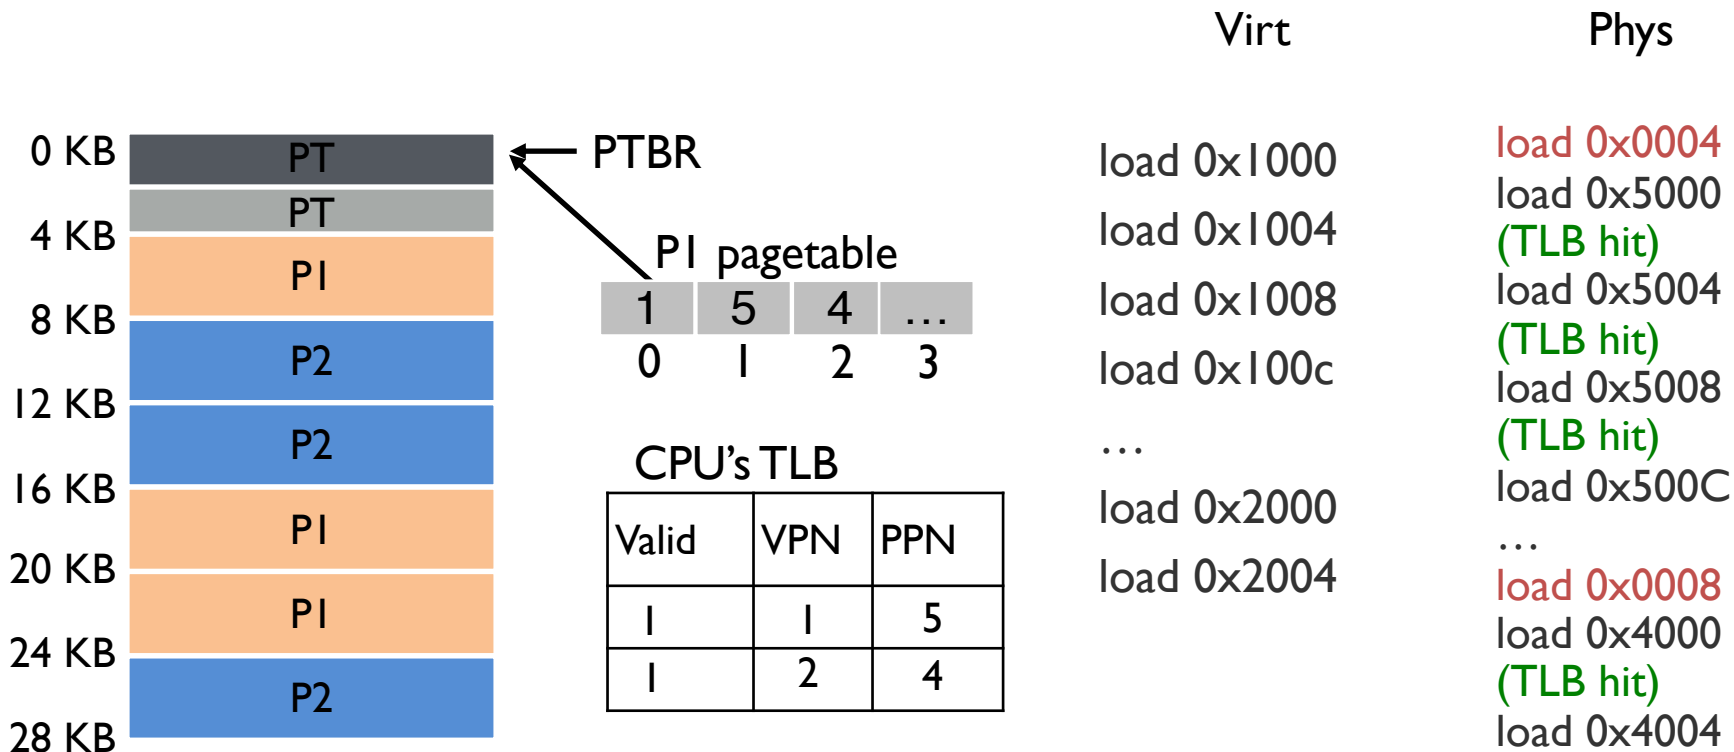
...

What will TLB behavior look like?

TLB ACCESSSES: SEQUENTIAL EXAMPLE



TLB ACCESSES: SEQUENTIAL EXAMPLE



QUIZ 10: TLBS

<https://tinyurl.com/cs537-sp20-quiz10>



Consider a processor with 16-bit address space and 4kB page size.
Assume Page Table is at 0x2000 and each PTE is of 4 bytes.

Simplified view of the PT

VPN	PPN
4	7
5	8
3	9
2	1

Virtual Addresses

0x3000: load 0x5320, %eax

0x3004: load 0x4004, %ebx

0x3008: mul %ecx, %eax, %ebx

0x300C: store %ebx, 0x5324

0x3010: load 0x5328, %ebx

Memory accesses

Total number of memory accesses

QUIZ 10: TLBS

Simplified view of the PT

VPN	PPN
4	7
5	8
3	9
2	1

Virtual Addresses

0x3000: load 0x5320, %eax

0x3004: load 0x4004, %ebx

0x3008: mul %ecx, %eax, %ebx

0x300C: store %ebx, 0x5324

0x3010: load 0x5328, %ebx

Memory accesses

Valid	VPN	PPN
0	2	6
0	7	23
0	2	5
0	3	2
0	1	89

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
= $2048 / (\text{elements of 'a' per 4K page})$
= $2\text{K} / (4\text{K} / \text{sizeof(int)}) = 2\text{K} / 1\text{K}$
= 2

Would hit rate get better or worse
with smaller pages?

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

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 2a is out!

Discussion today: Process API, Project 2a

Next class: More TLBs and better pagetables!