

# COS 140: Foundations of Computer Science

## Virtual Memory: Translation Lookaside Buffer

Fall 2018

<b>Virtual Memory</b>	<b>3</b>
Pages . . . . .	4
Paging. . . . .	5
Page faults. . . . .	6
<b>Page Tables</b>	<b>7</b>
Mapping . . . . .	7
Page tables . . . . .	8
Example . . . . .	9
Example . . . . .	10
MMU . . . . .	11
Address translation . . . . .	12
Problems . . . . .	14
Solution. . . . .	16
<b>TLB</b>	<b>17</b>
What is it? . . . . .	17
How it works . . . . .	18
Successful? . . . . .	19
<b>Multi-Level Paging</b>	<b>20</b>
Problem. . . . .	20
Solution. . . . .	21
Success? . . . . .	22
<b>Caching</b>	<b>23</b>

## Homework, announcements

- Reading: Chapter “Virtual Memory: Translation Lookaside Buffer” (Ch. 21)
- Homework: Exercises at end of chapter
- Due Wednesday, 11/14 (later than usual)
- NOTE:** Prelim II on 11/14
- Reminder: Advising!

Copyright © 2002–2018 UMaine Computer Science Department – 2 / 23

## Virtual Memory

3 / 23

### Virtual memory

- Goals:
  - Allow more processes to run simultaneously (increase *degree of multiprogramming*)
  - Allow very large processes to run
- Approach:
  - Keep most of each process on disk (in paging/swap area)
  - Only keep that part of each process in memory that is actually in use

Copyright © 2002–2018 UMaine Computer Science Department – 3 / 23

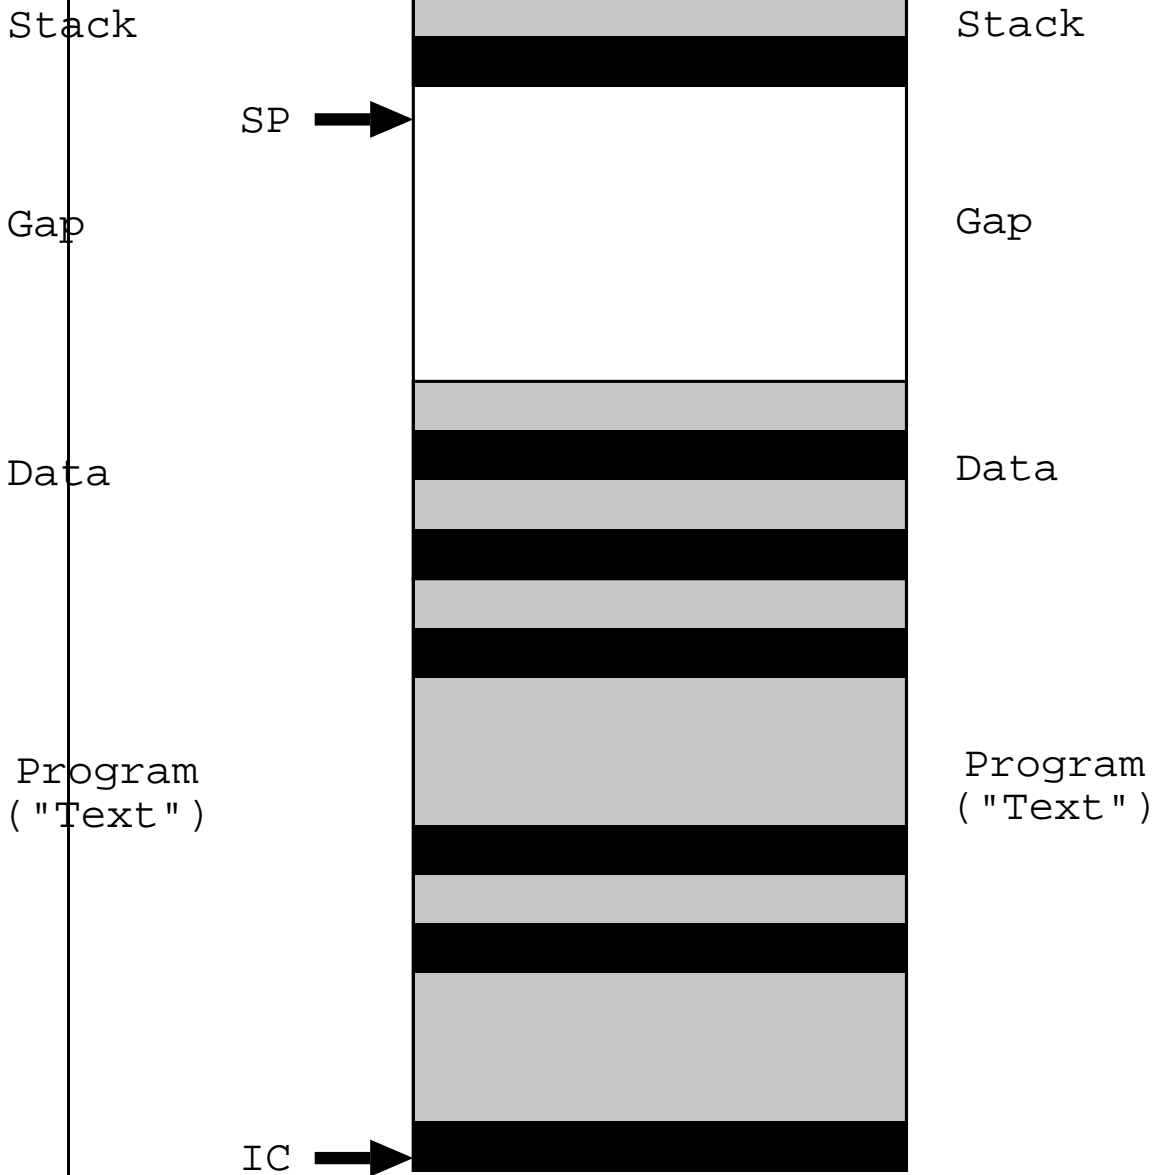
## Pages and page frames

- Divide process' address space into *pages* of some fixed size – usually 2Kb–4Kb
- Divide (physical) memory into *page frames* of same size
- Put needed pages of process into page frames: Need not be contiguous
- Move pages back and forth between disk and memory as needed: *demand paging*

Copyright © 2002–2018 UMaine Computer Science Department – 4 / 23

## When will a process need a new page?

Accessing non-resident data page IC moves to non-resident page Stack moves into non-resident-page



Copyright © 2002–2018 UMaine Computer Science Department – 5 / 23

## Page faults

- If page is *resident* – use it
- If page is not resident  $\Rightarrow$  *page fault*
- A page fault is a type of interrupt
- Operating system wakes up, tries to bring in the needed page
  - What if there are free page frames?
  - What if there are no free page frames?

Copyright © 2002–2018 UMaine Computer Science Department – 6 / 23

## Page Tables

7 / 23

### Page $\Leftrightarrow$ frame mapping


- At any given time, need to keep track of where a process' resident pages are
- Also need to keep track of which pages are not resident
- Use a *page table* for this
- Page table entries (PTEs)* map from pages to where those pages live in memory

Copyright © 2002–2018 UMaine Computer Science Department – 7 / 23

## Page tables

- One page table per process
- One page table entry per page in virtual memory (address space)
- Each entry contains:
  - Present/absent bit
  - Page frame number
  - Dirty bit (M bit)
  - Reference bit (R bit)
  - Maybe other things


**Paging Example**  
of Process (64K)



Page Table

15	
14	
13	
12	
11	
10	
9	
8	
7	
6	
5	
4	
3	
2	
1	
0	


**Virtual Memory**  
of Process (64K)



Page Table

15	3
14	
13	
12	
11	
10	
9	
8	
7	2
6	
5	
4	
3	
2	
1	
0	1

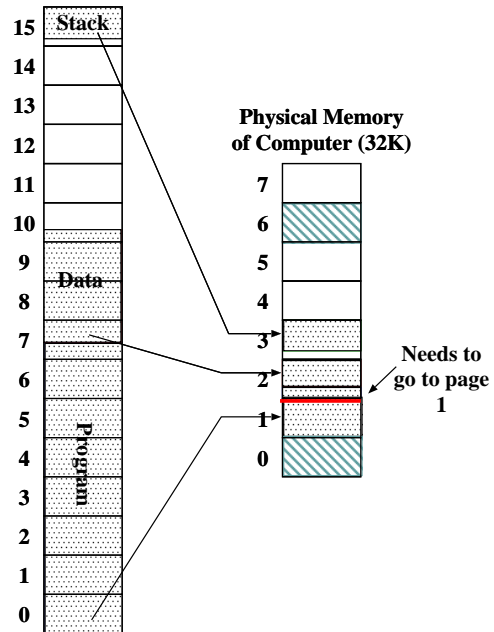
**Virtual Memory**  
of Process (64K)



Page Table

15	3
14	
13	
12	
11	
10	
9	
8	
7	2
6	
5	
4	
3	
2	
1	
0	1

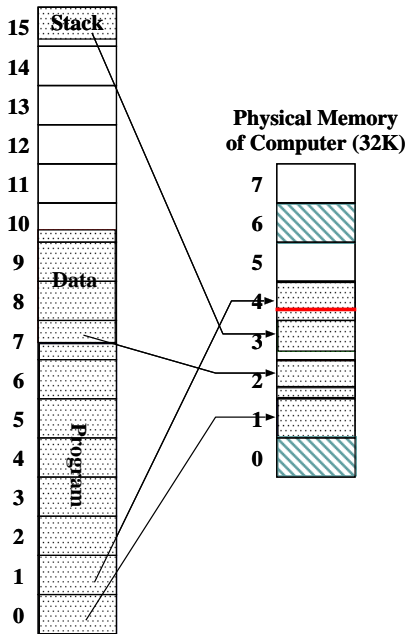
**Virtual Memory**  
of Process (64K)



Page Table

15	3
14	
13	
12	
11	
10	
9	
8	
7	2
6	
5	
4	
3	
2	
1	
0	1


**Paging Example**  
of Process (64K)



**Page Table**

15	3
14	
13	
12	
11	
10	
9	
8	
7	2
6	
5	
4	
3	
2	
1	4
0	1


**Virtual Memory**  
of Process (64K)



**Page Table**

15	3
14	
13	
12	
11	
10	
9	5
8	
7	2
6	
5	
4	
3	
2	7
1	4
0	1

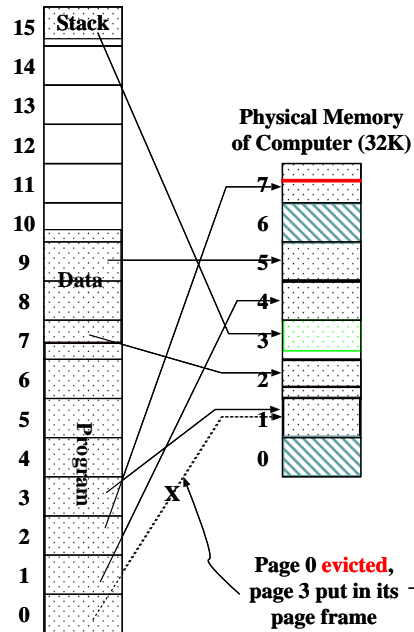
**Virtual Memory**  
of Process (64K)



**Page Table**

15	3
14	
13	
12	
11	
10	
9	5
8	
7	2
6	
5	
4	
3	
2	7
1	4
0	1

**Virtual Memory**  
of Process (64K)



**Page Table**


15	3
14	
13	
12	
11	
10	
9	5
8	
7	2
6	
5	
4	
3	1
2	7
1	4
0	<del>1</del> 3

What to do when PC gets to virtual page 3?

Page 0 evicted, page 3 put in its page frame

## Memory management unit

- How does a virtual address  $\Rightarrow$  a physical address?
- *Memory management unit* (MMU): piece of hardware that sits between the CPU and memory:



- These days: MMU is on CPU chip

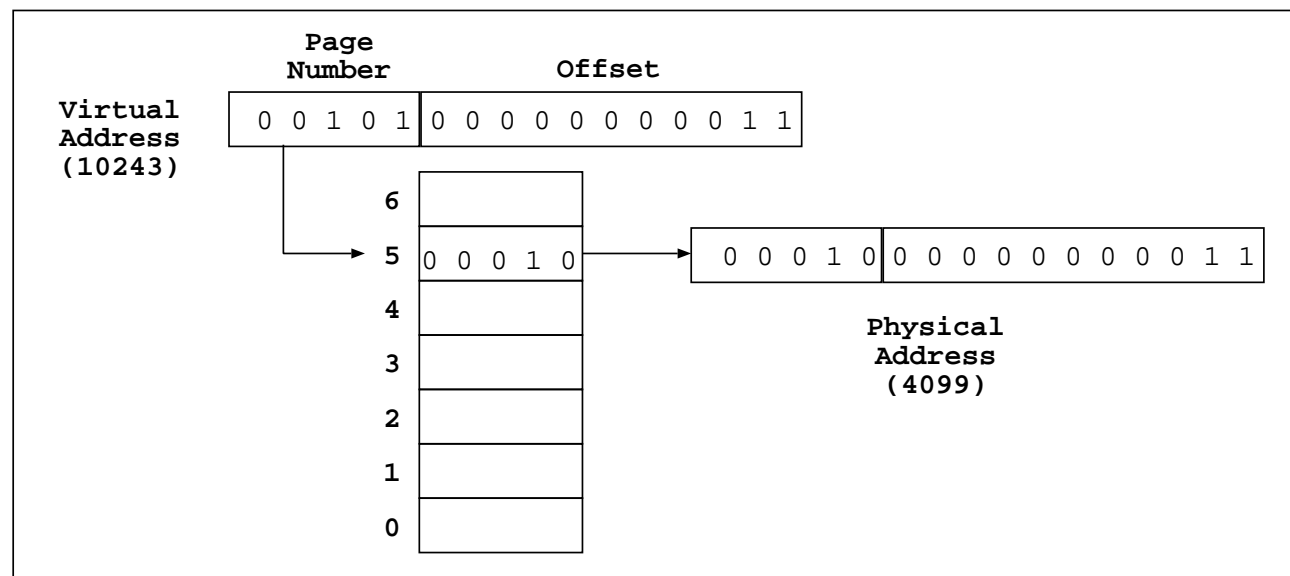


## Address translation

- How does the virtual address get translated to a physical address?
- Suppose we have 2KB pages, 16-bit machine:
  - $2\text{KB} = 2^{11}$  bytes – need 11 bits to address each byte on a page
  - Divide virtual address into 5-bit page number, 11-bit offset

Copyright © 2002–2018 UMaine Computer Science Department – 12 / 23

## Address translation



Copyright © 2002–2018 UMaine Computer Science Department – 13 / 23

## Problems with page tables

- Speed of access:
  - If table lives in memory, then for every memory access, have to look in page table to find address—a memory access itself—then do memory access
  - *Effective memory access time* is doubled
- Possible solution: Use very fast memory in MMU

Copyright © 2002–2018 UMaine Computer Science Department – 14 / 23

## Problems with page tables

- But page tables can be very large:
  - One entry per page in virtual address
  - With 2KB ( $2^{11}$  bytes) pages, 1 word/entry:
    - ▷ 16-bit machine:  $2^{16}/2^{11} = 2^5$  entries =  $2^5 \times 2 = 64$  bytes
    - ▷ 32-bit machines:  $2^{32}/2^{11} = 2^{21}$  entries =  $2^{21} \times 2^2 = 8\text{Mb}$  (!)
    - ▷ 64-bit machines:  $2^{64}/2^{11} \times 2^3 = 2^{56}$  bytes (!!)= 67,108,864 GB
- Realistically, can't afford that much fast memory!  
At \$6/GB, \$402,653,184...  
...per process!

Copyright © 2002–2018 UMaine Computer Science Department – 15 / 23

## Solution: Caching

- A *cache* is fast memory that holds *part* of what's in slower memory
- Idea: prevent accessing slower memory by keeping in cache what will be needed soon
- Cache according to (e.g.):
  - recency
  - frequency
  - predicted next use

Copyright © 2002–2018 UMaine Computer Science Department – 16 / 23

## TLB

17 / 23

### Translation Look-Aside Buffer

- TLB is a cache of page table entries
- TLB lives in MMU and is composed of very fast registers
- Special kind of registers: *associative* – MMU can look up page table entry corresponding to page number in one step

Copyright © 2002–2018 UMaine Computer Science Department – 17 / 23

## How it works

- When process starts – TLB empty, address 0 CPU  $\Rightarrow$  MMU:
  - PTE for page 0 not in TLB; read in from page table in memory
  - Page 0 not mapped  $\Rightarrow$  page fault
  - Page in page 0 into frame  $i$ , update PTE in TLB
- Next time some address on page 0 referenced:
  - Use PTE in TLB to find frame  $i$
  - No memory access for PTE
- When TLB full:
  - Have to eject some PTE from cache
  - Write it to page table first

Copyright © 2002–2018 UMaine Computer Science Department – 18 / 23

## How well does it work?

- Assume memory access time  $t_m = 100\text{ns}$ ; no TLB:  $200\text{ns}/\text{access}$
- With TLB:

$$T_{eff} = p(t_m + t_r) + (1 - p)(2t_m + t_r)$$

where  $T_{eff}$  is *effective access time*,  $p$  is probability of PTE being in the cache (*hit ratio*), and  $t_r$  is time to look something up in TLB

- With  $t_r = 20\text{ns}$ , hit ratio 0.5:  $T_{eff} = 170\text{ns}$
- Hit ratio 0.8:  $T_{eff} = 140\text{ns}$
- Hit ratio 0.98:  $T_{eff} = 122\text{ns}$


Copyright © 2002–2018 UMaine Computer Science Department – 19 / 23

## Multiple-level paging

- Problem: With large address sizes, page tables too big to keep in memory
- Would like to page out the page tables themselves!
- Can't, with monolithic tables
- Solution: multi-level page tables
- "Outer tables" act as page tables for "inner" pages – only outer needs to be resident in memory
- Price:  $\geq 2$  memory accesses/access in worst case

Copyright © 2002–2018 UMaine Computer Science Department – 20 / 23

## Multi-level page tables



Copyright © 2002–2018 UMaine Computer Science Department – 21 / 23

### Does TLB still help?

- Suppose we have 4-level paging, 100 ns memory access time, 20 ns TLB time
- Worst case: 500ns/access (4 page table accesses + desired access)
- With TLB, hit rate 0.98:  $T_{eff} = 188$  ns
- More realistic numbers:  $t_r = 1$  ns,  $t_m = 60$  ns; 4-level paging:  $T_{eff} = 68.8$  ns
- In case you're interested, the effective memory access time for  $n$ -level paging is:

$$T_{eff_n} = nt_r + t_m + n(1 - p)t_m$$

Copyright © 2002–2018 UMaine Computer Science Department – 22 / 23

## Caching

23 / 23

### Other kinds of caching

- Caching shows up many places in OS, elsewhere
- Processor caching physical memory
- Disk block caching for files
- Network file systems
- Virtual memory itself

Copyright © 2002–2018 UMaine Computer Science Department – 23 / 23