CIT 5950
Section 6

Scheduling & Virtual Memory

Logistics

- HW2 (Threads) due Next Friday, March 21st @ 11:59pm
- Check-in due before lecture next week
- Midterm is still being graded, hopefully have grades sometime next week

Exercise 1 - scheduling

Quick review of scheduling algorithms

- **FCFS (First come first served)**
  - Simple to implement
  - Low throughput, slow response, no priority
- **SJF (Shortest Job First)**
  - Minimal average turnaround time
  - Need to use estimates, no priority, possible starvation
- **Round Robin**
  - Relatively fair
  - Need to choose time quantum correctly, no priority
- **Priority Round Robin**
  - Introduce priority
  - Many design choices (# levels, time quantum, priority assignment)

Exercise 1

Name	Arrival Time	Running Time
Bert	2	11
Ernie	0	8
Oscar	12	20
Grover	7	15
Elmo	10	4

Consider the following set of tasks/processes:

Using the **Round Robin** scheduling algorithm and a time quantum/slice of 8, **what is the finishing time for each?**

What is the **average waiting time?**

Review - Virtual Memory

Memory Overload!

Suppose I want to run these games simultaneously …

8 GB RAM!!!

4 GB!!!

8 GB!!!

… on my humble laptop with 8GB random access memory (RAM)

(warning: don’t do it!)

Why not physical memory?

<u>**First-pass solution:**</u>

Can we store program contents (code + data + heap + stack) directly in physical memory (RAM)?

We can divide up RAM spaces so that each program occupies a fixed partition of RAM, preventing accidental overwrites

Problems with this approach?

Why not physical memory?

<u>**Problems with first-pass solution:**</u>

- We would run out of memory and crash!
  - We can't run any program larger than 8GB (or the partition we assign to the program)
- Potential inefficiencies in memory use
  - E.g. program 1 gets 7GB space, but only ends up using 2GB
- Compiler will need to know a program's partition beforehand
- Difficult to keep track of individual segments within a partition (e.g. enforcing read-only restriction)

Introducing Virtual Memory

- **Virtual memory is not real memory. It is an address representation system that we use to make memory appear larger than it is**
- We tell each program that it can use all the addresses in the "address space"
    - In a 64-bit machine, address space is 2^64, with addresses from 0 to 2^64 - 1
- In behind the scenes, each (used) virtual memory address is mapped to a space in the RAM or the disk
    - The Memory Management Unit (MMU) takes care of this

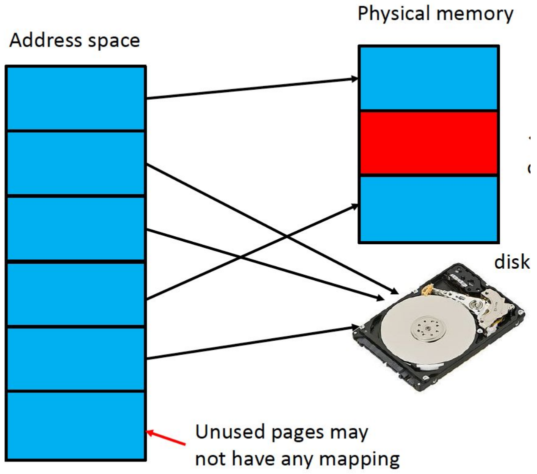

Benefits of virtual memory

- We are no longer bounded by the size of the RAM!
  - Disk space is typically much larger (and comes at a cheaper price)
- We do not need to worry about process isolation
  - The MMU will take care of that
- More flexibility and efficiency in managing memory
  - The MMU can switch the actual storage location without the user/program noticing
  - Less likely to have “gaps” / fragmentation in memory

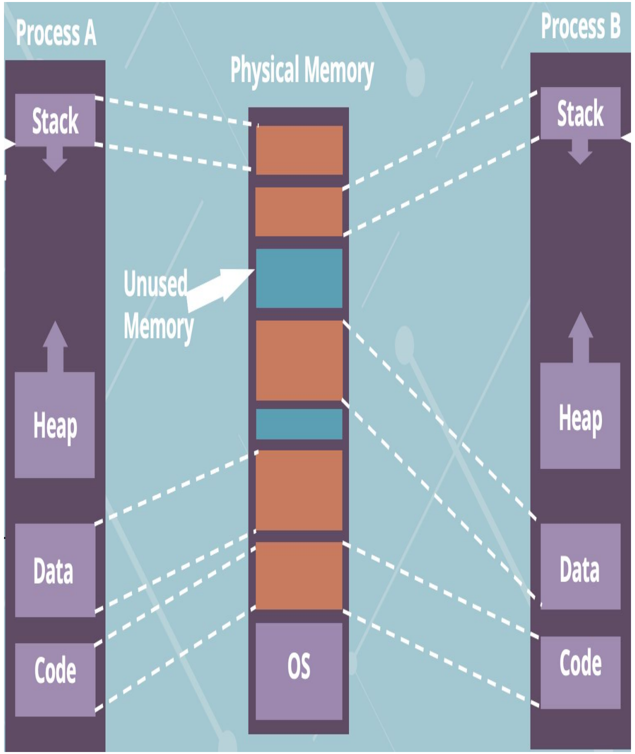

Why pages and page table?

- **OK, virtual memory is good. We get it. Why do we need this extra thing called pages and page table?**
  - Imagine a mapping of every individual virtual address to physical address - the lookup table will be as big as the address space (e.g. 2^64)!
  - Dividing virtual memory addresses into chunks make it easier to manage
- **A page is a unit of virtual memory (e.g. 4KB)**
  - Each page in virtual memory maps to a frame in physical storage (RAM + disk) of the same size
  - The mapping is recorded in the **page table**
  - Each process has one page table

This table will be very large and inefficient!

Virtual address	Valid	Physical address
x1A23	0	X2253 (RAM)
x399A	1	X5001 (Disk)
x7282	0	X3AB2 (RAM)

Keeping track of memory in page level is easier

Virtual page #	Valid	Physical Page Number
0	0	null
1	1	0
2	1	1
3	0	disk

VM Calculation Cheat Sheet

Address space = 2 ^ (# bits in the address, e.g. 64 for 64-bit machine)

Size of virtual memory = Address space * Addressability (# Bytes in each address)

Number of pages = Size of virtual memory / size of a page

Size of a page = size of a frame

Number of frames = Physical Memory (RAM) space / size of a frame

Number of bits to represent pages or frames = Log_2 (number of pages or frames)

Page number = first N bits of the virtual address, N = number of bits to represent pages or frames

1 KB = 2^10 B   1 MB = 2^20 B   1 GB = 2^30 B

Exercise 2
Pages &
Page tables

Exercise 2

Consider a system as follows:

- 32-bit address space
- 16-bit addressable
- 1GB of physical memory
- page sizes of 64kB

a) How many **pages** are there in virtual memory?

b) How many **frames** are there in physical memory?

c) How many **bits** are there in each address' page number?

d) Consider the virtual address xABCDEF01. What is its **page number** in **hexadecimal**?

Page Replacement Algorithms

Page replacement algorithms

- The RAM can only hold up to a certain number of pages
  - If the RAM is full, we need to “evict” pages or move them to the disk
  - How should we decide which pages to evict?
- Goal: optimize (minimize) number of times we need to fetch something from the disk (**page faults**)
  - **FIFO** (first in first out): evict the page that first entered physical memory
  - **LRU** (least recently used): evict the page that has not been used for the longest time

Exercise 3
Page Eviction algorithms

Exercise 3

We have a byte-addressable system that has a 16-bit address space, 32kB of physical memory, and page sizes of 8kB. Assume the page table is initially empty, and then a process generates the following sequence of virtual addresses:

x3311
x1234
x1255
x3456
xA349
x7777
xB222
x6222

a) If virtual address `x5324` is requested next, which page will be evicted if using a First In First Out (FIFO) replacement algorithm?

b) Instead of using FIFO, which page will be evicted if using a Least Recently Used (LRU) replacement algorithm?

c) Rather than using FIFO or LRU, imagine that the system could look into the future and see that the next four virtual address requests (after x5324) would be x1A23, x399A, x7282, and x4A32. Knowing this information, which page should be evicted when the request for x5324 generates a page fault?