

CMPU 334 Exam 1
Spring 2020

Name: Master Key

Instructions:

This is a closed book, closed notes exam. No electronic devices, including calculators are allowed. You have 75 minutes. There are 12 problems and 8 pages to this exam.

Good luck!

1 (8)	
2 (8)	
3 (8)	
4 (8)	
5 (8)	
6 (8)	
7 (8)	
8 (10)	
9 (8)	
10 (8)	
11 (8)	
12 (10)	
Total (100)	

1. The purpose of the operating system.

a. What are two roles the operating system plays in order to make application programmers' jobs easier?

- 1) a virtual machine
- 2) a standard library

b. What are two (often conflicting) goals/requirements for the OS while running programs?

- 1) efficiency
- 2) safety/control/isolation/security

"how"

"which"

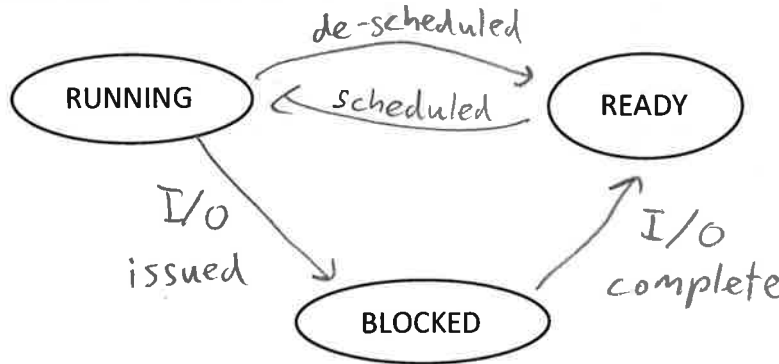
2. Mechanisms vs. Policies. Circle the items that correspond to mechanisms below:

- a. The OS uses LRU to select a page to evict upon memory pressure
- b. Hardware raises a page fault for a memory access
- c. The OS inserts a new TLB entry upon a (software-managed) TLB miss
- d. The OS schedules processes with Round Robin

Policy

policy

3. The States of a Process. Draw the arrows to complete the process state machine. Label each arrow with a word or two describing what could cause the transition. Be sure to leave out arrows for transitions that never occur.

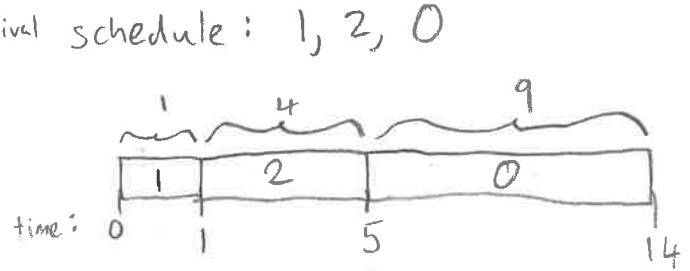


4. **Scheduling Metrics.** The following 3 jobs (with their running times listed) arrive at time 0 in the following order:
 Job 0 (length = 9)
 Job 1 (length = 1)
 Job 2 (length = 4)

Compute the response time and the turnaround time for the following scheduling algorithms. Assume the RR order is the same as the order in which the jobs are listed and a time quanta of 1.

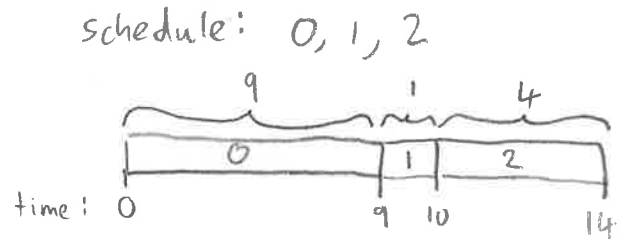
SJF (Shortest Job First)

Job	Response Time	Turnaround Time
0	$5-0=5$	$14-0=14$
1	$0-0=0$	$1-0=1$
2	$1-0=1$	$5-0=5$



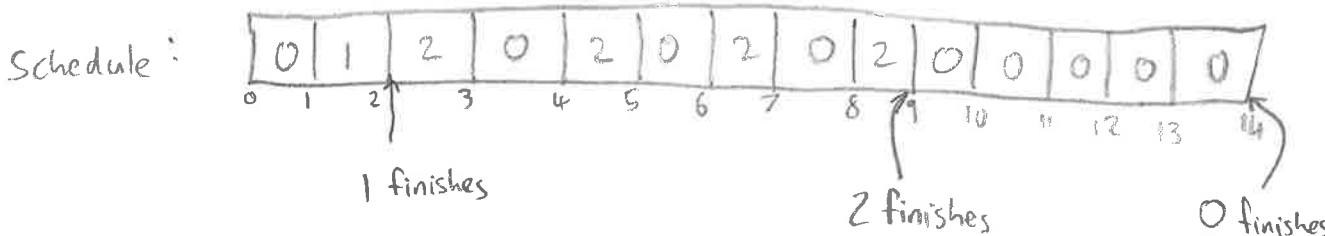
FIFO (First In First Out)

Job	Response Time	Turnaround Time
0	$0-0=0$	$9-0=9$
1	$9-0=9$	$10-0=10$
2	$10-0=10$	$14-0=14$



RR (Round Robin)

Job	Response Time	Turnaround Time
0	$0-0=0$	$14-0=14$
1	$1-0=1$	$2-0=2$
2	$2-0=2$	$9-0=9$



5. Address Translation.

a. How many virtual memory accesses must be translated for the following short program?

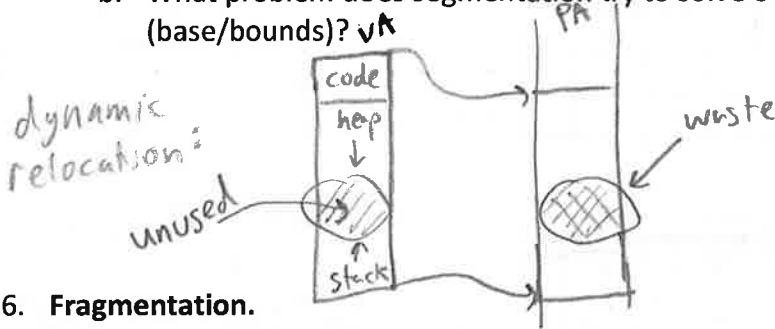
```

1, load 0x2000, %r1
2, inc  %r1
3, store %r1, 0x2000
    
```

2 load/stores
+ 3 instruction fetches

5

b. What problem does segmentation try to solve over simple dynamic relocation (base/bounds)?



Eliminate wasted space by relocating each segment independently

6. Fragmentation.

a. Describe the difference between internal vs. external fragmentation. Give an example of when the OS may encounter each.

Internal: unusable space within allocated objects → e.g. large pages

External: unusable space between allocated objects → e.g. variable-sized segments



b. Very large pages
i. What type of fragmentation could occur with large page sizes?

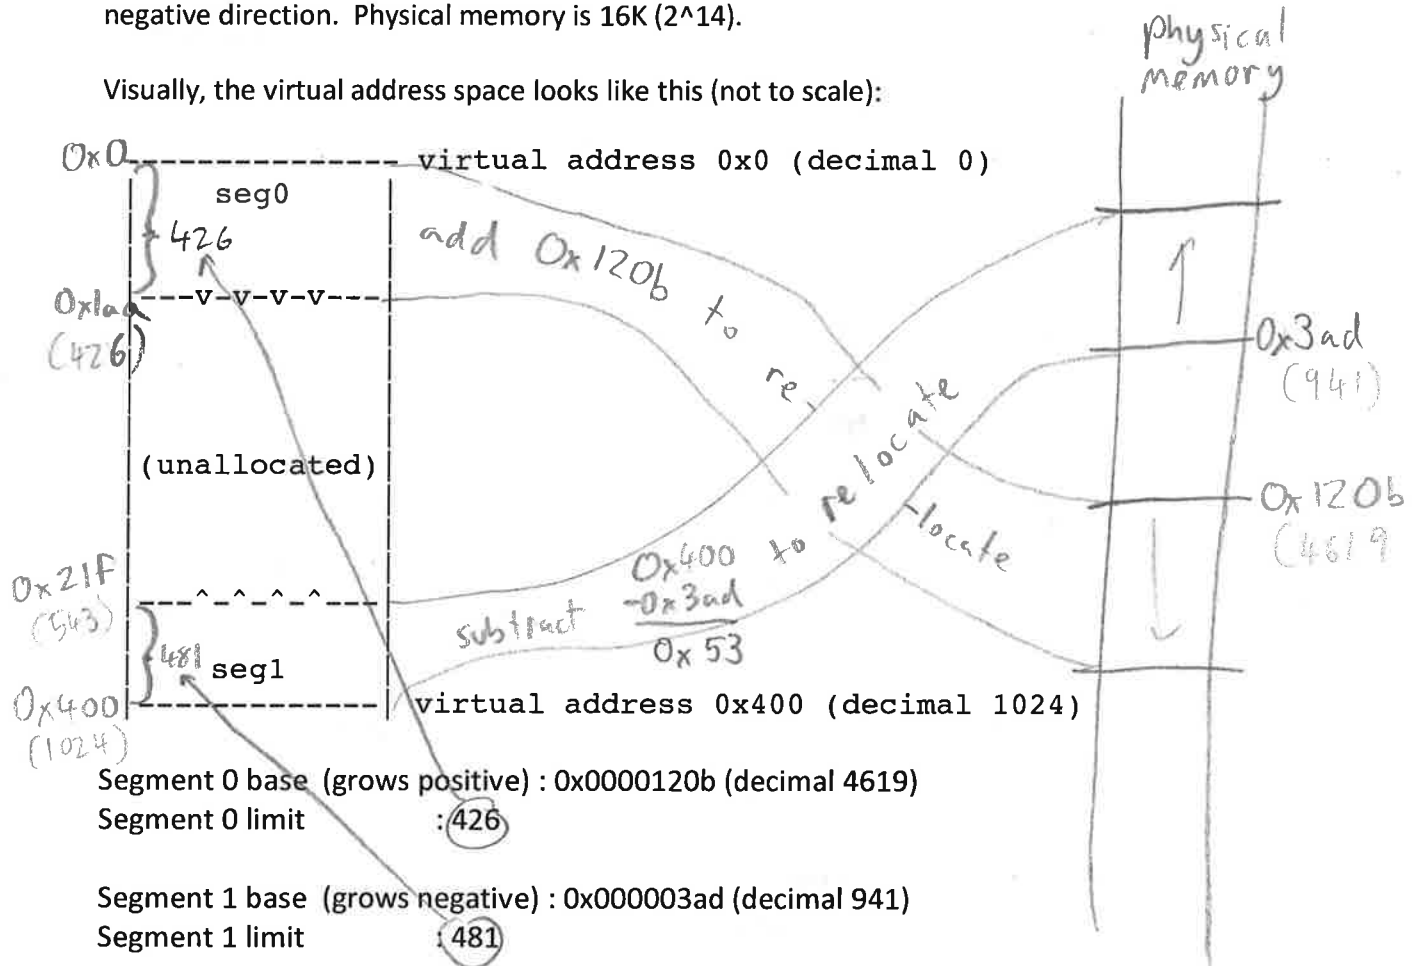
internal

ii. What are two advantages to using a large page size?

- 1) smaller page table
- 2) better TLB coverage

7. **Address translations with segmentation.** For this problem, you should assume a simple 1K (2^{10}) virtual address space with only two segments. The top bit of the 10 bit virtual address is used to check which segment the virtual address is in: 0 for segment 0 (where, say, code and the heap would reside) and 1 for segment 1 (where the stack lives). Segment 0 grows in a positive direction (towards higher addresses), whereas segment 1 grows in the negative direction. Physical memory is 16K (2^{14}).

Visually, the virtual address space looks like this (not to scale):



For each VA below, write which segment the VA is in (0 or 1) and the physical address (PA) that it translates to. If the VA is out of bounds, write SEGFAULT instead of the PA.

VA 0x0b4 (decimal: 180) --> SEG	0	PA	0x12bf	0x0b4 + 0x120b = 0x269
VA 0x269 (decimal: 617) --> SEG	1	PA	0x216	0x269 - 0x53 = 0x216
VA 0x341 (decimal: 833) --> SEG	1	PA	0x2ee	0x341 - 0x53 = 0x2ee
VA 0x1ee (decimal: 494) --> SEG	0	PA	Segfault	
VA 0x092 (decimal: 146) --> SEG	0	PA	0x129d	0x092 + 0x120b = 0x29d

Each VA is 10 bits long

hex digit hex digit

4 4

segment

which segment?

bounds check

10. Translation Lookaside Buffers.

a. Why is the principle of *locality* in programs important for TLB performance?

The TLB is a cache of address translations.

Locality means that a program's memory accesses (for code or data) will likely be near (on the same page) as other recent accesses. Thus the translation is likely to hit in

b. True or false: Circle the true statements below:

i. The primary goal of the TLB is to reduce the size of page tables

ii. The TLB is a cache of address translations

iii. The TLB resides in main memory (RAM)

iv. Depending on the architecture, TLB misses can be served by hardware or software

v. TLB entries translate segment bases to bounds

vi. TLB entries translate VPNs to PFNs

vii. The TLB translates addresses generated by instruction fetches

viii. The TLB translates addresses generated by loads

ix. The TLB translates addresses generated by stores

x. The TLB evicts pages to disk

ALL
memory
accesses

the TLB, resulting in
good performance

↳ no, speed
them up

↳ no, the point is to avoid costly RAM
access

11. Page Table Entries.

a. PTEs typically contain the PFN and some other bits that provide useful information to the OS. Give three examples of other bits one may find in a PTE.

1. valid bit
2. dirty bit (whether contents need to be written on eviction)
3. referenced bit (for e.g. clock algorithm)

b. When a page is swapped to disk, what happens to the corresponding PTE?

1. the present bit is cleared.
2. the PFN is replaced with the offset on the swap disk.

12. Page Replacement Policies.

- a. Assuming a replacement policy of OPT (optimal), and a cache of size 3 pages, figure out whether each of the following page references hit or miss in the page cache.

evict page used furthest in the future

VPN Access	Hit/Miss (Circle one)	State of memory after access (3 pages)
0	Hit/ Miss	0--
1	Hit/ Miss	01-
1	Hit /Miss	01- → used furthest in future
5	Hit/ Miss	015 → used furthest in future
6	Hit/ Miss	615
4	Hit/ Miss	415
5	Hit /Miss	415
4	Hit /Miss	415
1	Hit /Miss	415
1	Hit /Miss	415

- b. Assuming a replacement policy of LRU (least recently used), and a cache of size 3 pages, figure out whether each of the following page references hit or miss in the page cache.

VPN Access	Hit/Miss (Circle one)	State of memory after access (3 pages)
0	Hit/ Miss	0--
1	Hit/ Miss	01-
1	Hit /Miss	01- → least recently used
5	Hit/ Miss	015 → least recently used
6	Hit/ Miss	615
4	Hit/ Miss	645
5	Hit /Miss	645 → least recently used
4	Hit /Miss	645
1	Hit/ Miss	145
1	Hit /Miss	145

- c. In what circumstances is LRU a poor policy choice?

Looping over more pages than fit in memory causes every new page to be a miss!