CS 4410: Operating Systems
Homework 7

- Homework may be done in pairs, or individually. If doing in pairs, **one** of you should upload to gradescope and add your partner to the group assignment in the upper right corner of the screen. (Do **not** just upload the assignment twice or it will be graded twice, which means grading will take longer.)
- The deadline is Mon, Nov 21 at 11:59AM.
- No late submissions will be accepted.
- You must attribute every source used to complete this homework.
- For some of the problems, you will need two integers. Here is the algorithm for computing these integers:
  - If you are working with a partner, let `var` be the lexicographically smaller of the two NetIDs.
  - Let `varInt` be the integral part of `var`. That is, if `var = rst12`, then `varInt = 12`.
  - If `varInt` is a single digit integer, let `varInt = 13 × varInt`
  - Let `Int1` be the first digit of `varInt`
  - Let `Int2` be the second digit of `varInt`
- Assume the storage unit convention: $1\text{G} = 2^{10} \times \text{M} = 2^{20} \times \text{K} = 2^{30}$ (bytes).
- **For all problems that use `Int1` or `Int2`, please write down the parameters (related variables and settings calculated from your NetID) before answering each question.**

1 Hardware Upgrade

The boss of VWMare, Harold, asks you to upgrade the hardware for machines in their data center. Currently each machine has one TLB and one physically addressed cache. It uses two-level page tables, and the operating system demands paging to disk. The TLB is automatically filled by the MMU upon a TLB miss. The page tables are always kept in physical memory and not cached. Finding a page table entry takes two memory accesses: the first to access the first level, the second for the second (where the actual page table entry is found). Assume the operating system always has access to a a set of clean pages, so that pages do not need to be written back to the disk when there is a page fault.

To help you understand the current system performance, Harold gives you the following report:

Measurement	Value
$P_{\text{CacheMiss}}$ = probability of a cache miss	0.01
P_{TLBMiss} = probability of a TLB miss	0.01
P_{Fault} = probability of a page fault, given a TLB miss	0.00002
T_{Cache} = time to access cache	1ns
T_{TLB} = time to access TLB	1ns
T_{DRAM} = time to access main memory	$(100+10\cdot\texttt{Int1})$ns
T_{HDD} = time to transfer a page to/from a hard disk drive (HDD)	$(10+\texttt{Int2})$ms
$P_{\text{HDD-MISS}}$ = probability of a hard disk drive (HDD) miss	0

1. What is the memory access time on average for an application to do a single memory reference, using a machine from VWMare?

2. Given a budget of $300 per machine, which of the folowing upgrade options would you choose in order to maximize performance? What is the new average memory access time of your upgraded system?

Item	Specs	Cost (dollars)
(A) a faster HDD	transfers a page in 7ms, no miss rate	100
(B) a solid state drive (SSD) (between DRAM and HDD)	Transfers a page in $(10+\texttt{Int1})\mu s$, miss rate is 0.1	100
(C) larger DRAM	Changes the probability of a page fault to 0.00001	100
(D) a faster network adapter	$(300+10\cdot\texttt{Int2})$ms to access remote memory; probability of a remote memory miss is 0.5, given a page fault	100

2 Raid by RAID

In a RAID 5 system, suppose there are 6 hard drives, each of which has one single-sided disk platter having 40000 tracks, each of which has 6000 sectors of 512 bytes. The strip configuration of the RAID is shown as follow:

Disk 0	Disk 1	Disk 2	Disk 3	Disk 4	Disk 5
0	1	2	3	4	p_0
5	6	7	8	p_1	9
10	11	12	p_2	13	14
15	16	p_3	17	18	19
20	p_4	21	22	23	24
p_5	25	26	27	28	29
30	31	32	33	34	p_6
35	36	37	38	p_7	39

1. What is the capacity of this RAID (excluding the space reserved for redundancy)?
2. Assume disk 1 has failed and been replaced by a new disk, which disk(s) will be accessed to reconstruct the data on the newly inserted disk? Specify the disk(s) and the operations on them.
3. Suppose this RAID can service two write requests (to two different disks) in parallel. When writing to a block (a cell in the above figure), the RAID 5 system will update both the data and the corresponding parity block on the same row. Currently, block 0 is being written. There are pending write requests to the blocks listed below (order does not matter and they are independent):

   $$0, 4, 8, 21, 24, 26, 30, 38, 32$$

   The system should service a second write request in parallel with the writing of block 0. This other request must not conflict with the writing of block 0. Which of these pending requests can possibly be scheduled to execute in parallel with the write to block 0?

3 Elevator

The new scheduling algorithm for the elevator in VWMare headquarters will be provided by Harold. He asks you to analyze elevator performance by simulating a real world example for each of his scheduling algorithm candidates.

The following list shows all the requests to visit certain floors in the order that they were requested. Initially, the elevator is on the $(11+(\texttt{Int1} \bmod 3))$th floor.

$$5, 23, 9, 14, 2, 20, (\texttt{Int2}+1), 10, 12, 16, 30$$

1. **First come first served (FCFS)**. What is the total movement for the elevator (in number of floors)?
2. **Shortest seek time first (SSTF):** the elevator chooses from the list the floor closest to its current location, breaking ties by choosing the request that arrived earlier. What is the total movement for the elevator?
3. **C-LOOK:** the elevator goes as far as the furthest request in its current direction and then reverses its direction immediately. Assume that initially the elevator is going up. What is the total movement for the elevator?