CSci 402 - Operating Systems

Final Exam (TT Section)

Spring 2024

[9:00:00am-9:40:00am), Tuesday, May 7)

Instructor: Bill Cheng

Teaching Assistant: Zhuojin Li

(This exam is open book and open notes.
Remember what you have promised when you signed your
Academic Integrity Honor Code Pledge.)

Time: 40 minutes

_____-
Name (please print)

Total: 38 points

_____-
Signature

Instructions

1. This is the first page of your exam. The previous page is a title page and does not have a page number. Since this is a take-home exam, no need to sign above since you won't submit this file.
2. Read problem descriptions carefully. You may not receive any credit if you answer the wrong question. Furthermore, if a problem says *"in N words or less"*, use that as a hint that N words or less are expected in the answer (your answer can be longer if you want). Please note that points may get *deducted* if you put in wrong stuff in your answer.
3. If a question doesn't say `weenix`, please do not give `weenix`-specific answers.
4. Write answers to all problems in the **answers text file**.
5. For non-multiple-choice and non-fill-in-the blank questions, please show all work (if applicable and appropriate). If you cannot finish a problem, your written work may help us to give you partial credit. We may not give full credit for answers only (i.e., for answers that do not show any work). Grading can only be based on what you wrote and cannot be based on what's on your mind when you wrote your answers.
6. Please do *not* just draw pictures to answer questions (unless you are specifically asked to draw pictures). Pictures will not be considered for grading unless they are clearly explained with words, equations, and/or formulas. It's very difficult to draw pictures in a text file and you are not permitted to submit additional files other than the answers text file.
7. For problems that have multiple parts, please clearly *label* which part you are providing answers for.
8. Please ignore minor spelling and grammatical errors. They do not make an answer invalid or incorrect.
9. During the exam, please only ask questions to *clarify* problems. Questions such as "would it be okay if I answer it this way" will not be answered (unless it can be answered to the whole class). Also, you are suppose to know the definitions and abbreviations/acronyms of *all technical terms*. We cannot "clarify" them for you. We also will **not** answer any clarification-type question for multiple choice problems since that would often give answers away.
10. Unless otherwise specified and stated explicitly, multiple choice questions have one or more correct answers. You will get points for selecting correct ones and you will lose points for selecting wrong ones.
11. When we grade your exam, we must assume that you wrote what you meant and you meant what you wrote. So, please write your answers accordingly.

(Q1) (2 points) Which of the following statements are correct about approaches to deal with the problem caused by the **popf** instruction so that a virtual machine can be built for the **x86 processor**?

(1) VMware's solution is a compile-time solution, i.e., some sensitive instructions are replaced with hypercalls when the kernel is compiled
(2) in Intel's hardware-based solution, the `popf` instruction is disabled so that nothing happens when the `popf` instruction is executed
(3) with paravirtualization, certain sensitive instructions are replaced with hypercalls at the time the kernel is compiled
(4) in Intel's hardware-based solution, a new "processor mode" was added
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q2) (2 points) Which of the following statements are correct about the **FIFO** scheduler?

(1) the FIFO scheduler is inherently unfair to long jobs
(2) for the FIFO scheduler, average waiting time does not depend on the ordering of jobs at the run queue
(3) the FIFO scheduler has the largest variance in waiting time among all scheduling disciplines
(4) the FIFO scheduler appears to be a fair scheduling policy
(5) "starvation" at the scheduler is a common problem for the FIFO scheduling policy

Answer (just give numbers): _____

(Q3) (2 points) A **disk map** of a S5FS **inode** contains 13 disk block pointers. If a disk block is 1KB in size, which of the following are **possible file sizes** if an inode's **disk map** contains non-null block numbers in the first 12 entries and the last entry is null?

(1) 256 KB
(2) 100 KB
(3) 96 MB
(4) 64 MB
(5) 4 MB

Answer (just give numbers): _____

(Q4) (2 points) Which of the following statements are correct about the naive spin lock implementation vs. the **"better" spin lock implementation**?

(1) the naive spin lock does busy-waiting while the "better" spin lock does not do busy-waiting
(2) the naive spin lock and the "better" spin lock are functionally equivalent, i.e., they achieve the same functionality; the main difference is their performance impact on the system
(3) no matter which spin lock implementations you use, the impact on system performance is the same if the spin lock is currently unavailable
(4) using the "better" spin lock can lock the spin lock slightly faster than the naive spin lock if the spin lock is currently available
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q5) (2 points) Let's say that the address space of a user space process in **weenix** looks like the following:

```
    VADDR RANGE        PROT FLAGS    MMOBJ      OFFSET    VFN RANGE
0x08029000-0x08038000 rw-  PRIVATE 0xcfe0c034 0x00012 0x08029-0x08038
0x08038000-0x0804c000 r-x  PRIVATE 0xcfe0c004 0x00013 0x08038-0x0804c
0x0804d000-0x08060000 rw-  PRIVATE 0xcfe0c064 0x00034 0x0804d-0x08060
```

If you get a page fault with vaddr = 0x0804b668, what **pagenum** would you use to lookup a page frame when you are handling a page fault? Please just give an integer value answer (no partial credit for this problem).

(Q6) (2 points) Which of the following statements are correct about **user processes and threads in weenix**?

(1) a user process and its corresponding kernel process are represented using different kernel data structures
(2) there is no kernel data structure to represent a kernel stack
(3) a user stack is explicitly represented in the kernel as a kernel stack
(4) to "kill" a user thread, you just need to "kill" the corresponding kernel thread since they are exactly the same thread
(5) weenix uses one data structure to represent the address space of a user process and a separate data structure to represent the address space of the corresponding kernel process

Answer (just give numbers): _____

(Q7) (2 points) Which of the following statements are correct about **LFS (log-structured file system)**?

(1) LFS's append-only and never delete/modify requirements make the idea of LFS not very useful in practice for a file system on a hard drive
(2) the inode map in LFS achieves the same functionality as the disk map in S5FS
(3) LFS is designed to use close to 100% of the disk transfer capacity when writing to a hard drive
(4) the two checkpoint files in LFS achieves the same functionality as the superblock in S5FS
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q8) (2 points) Which of the following statements are correct about a **SJF (shortest job first)** scheduler?

(1) this scheduler generally has a smaller variance in waiting time than other schedulers
(2) compared with some other schedulers, this scheduler can have a large average waiting time
(3) it is possible that short jobs may "starve" if long jobs keep arriving
(4) this scheduler is inherently unfair to short jobs
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q9) (2 points) Which of the following statements are correct about using the basic **multi-level feedback queues** to schedule both interactive and non-interactive jobs?

(1) if a highest priority thread gives up the CPU voluntarily before using up a full time slice, you should decrease its priority
(2) "aging" in multi-level feedback queues means to always assign the highest priority to a job that has been running in the system the longest
(3) if a highest priority thread uses a full time slice, you should add a new higher priority queue and increase its priority
(4) if a highest priority thread uses a full time slice, you should decrease its priority
(5) when a job arrives at the run queue (implemented as multi-level feedback queues), it should get the highest priority

Answer (just give numbers): _____

(Q10) (2 points) Which of the following are bits inside a PTE (page table entry) that are **used in weenix** (which runs on an x86 CPU)?

(1) a "busy" bit
(2) a "private/shared" bit
(3) a "pinned" bit
(4) a "dirty" bit
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q11) (2 points) The following is part of the code that is used to lock a mutex in a **straight-threads** (i.e., no interrupt) mutex implementation on a **single CPU**:

```
if (!m->locked) {
  m->locked = 1;
}
```

Which of the following statements are correct about using the above code in a **multiple CPU** system to lock mutex `m` correctly and safely?

(1) it's safe and correct to use the code as-is to lock mutex `m` because it's impossible for multiple CPUs, sharing the same bus, to read from the same memory location at exactly the same time
(2) it's safe and correct to use the code as-is to lock mutex `m` because it's impossible for multiple CPUs, sharing the same bus, to write to the same memory location at exactly the same time
(3) it's safe and correct to use the code as-is to lock mutex `m` because setting the value of m->locked is an atomic operation
(4) it's safe and correct to use the code as-is to lock mutex `m` because reading the value of m->locked is an atomic operation
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q12) (2 points) Let's say that you are using a **rate-monotonic scheduler** to schedule 4 periodic tasks with $T_1 = 0.5$, $P_1 = 6$, $T_2 = 1$, $P_2 = 6.5$, $T_3 = 0.5$, $P_3 = 7$, and $T_4 = 1$, $P_4 = 7.5$ (all in seconds). If you were to simulate the **rate-monotonic scheduler** to see if it can schedule these 4 periodic tasks without any of them missing a deadline, assuming that you schedule all 4 periodic tasks to start a time = 0, how many seconds into the simulation would be the first time all 4 jobs would be scheduled to start executing at exactly the same time again? Please just give a numeric answer (no partial credit for this problem).

(Q13) (3 points) Let's say that you have four threads A, B, C, and D and you are using **stride scheduling**. You have decided to give thread A 5 ticket, thread B 9 tickets, thread C 12 tickets, and thread D 18 tickets. The initial pass values that **you must used** for the four threads are shown below along with the "winner" of the iteration 1. Please run **stride scheduling** to fill out all the entries (pass values) in the table and keep track of the "winner" in each round. For **iterations 2 through 7**, please write on your answer sheet the "winner" and the winning pass value of that iteration. (For example, you would write "D:9" for iteration 1 since D is the "winner" of iteration 1 and the winning pass value is 9.) You must use the **smallest possible integer stride values** when calculating all the pass values. If you get the stride values wrong, you will not get any partial credit for this problem.

itr	A	B	C	D
1	13	20	18	(9)
2				
3				
4				
5				
6				
7				

(Q14) (2 points) Which of the following is correct about **management of page frames**?

(1) Linux uses a clock algorithm to determine if a page frame is recently used or not
(2) if the idea of "working set" is fully implemented in the OS, thrashing can be prevented
(3) in Linux, user pages can be found in all three physical memory "zones"
(4) if the "local allocation" scheme is implemented in the OS, thrashing cannot happen
(5) in Linux, a dirty and inactive page frame must be freed/deallocated after it has been "cleaned"

Answer (just give numbers): _____

(Q15) (2 points) Which of the following statements are correct about **shadow objects** in **weenix** if we want **copy-on-write** and **fork()** to work together correctly?

(1) a shadow object holds pages that were copy-on-write but have never been modified
(2) a list of shadow objects is maintained in a doubly-linked circular list in weenix
(3) if a vmarea is privately mapped and read-writable, you must use a shadow object for its first mmobj
(4) if a vmarea is shared-mapped, you must use a shadow object for its first mmobj
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q16) (2 points) Which of the following statements are correct about **I/O virtualization**?

(1) in VMware's solution to I/O virtualization, most device drivers in the guest OS must be rewritten so that they can be supported
(2) Xen's I/O virtualization solution performs better than VMware's I/O virtualization solution
(3) I/O virtualization for desktop machines is challenging because it's virtually impossible for virtual machine vendors to support all devices
(4) in Xen's solution to I/O virtualization, only a few device drivers in the guest OS has to be modified in order for Xen to use them
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q17) (3 points) Let's say that you have four threads A, B, C, and D and you are using the basic **round-robin (RR) / time-slicing** scheduler with a very small time slice. At time zero, all four threads are in the run queue and their processing times are shown in the table below. Assuming that there are no future arrivals into the run queue, please complete the table below with the "waiting time" of all four threads and the "average waiting time" (AWT) of these four threads and write the results on your answer sheet. Please make it very clear which waiting time is for which thread and which one is the AWT. For non-integer answers, you can use fractions or decimals with two digits after the decimal point. Your answer must not contain plus or multiplication symbols. You must use the definition of "waiting time" given in lectures.

	A	B	C	D	AWT (1 pt)
T (hrs)	14	12	12	11	-
wt (hrs)					

(Q18) (2 points) Considering only **clustered hash page table** schemes and **(non-clustered) hashed page table** schemes, which of the following statements are correct?

(1) a clustered hash page table usually would take up more space than a non-clustered hash page table if address space is sparsely allocated
(2) non-clustered hash page tables typically performs better than clustered page tables
(3) the performance of non-clustered hash page tables is independent on the lengths of the hash conflict/collision resolution chains
(4) the performance of clustered hash page tables is independent of how address space is allocated
(5) none of the above is a correct answer

Answer (just give numbers): _____