CSci 402 - Operating Systems

Final Exam (DEN Section)

Fall 2025

[9:00:00am-9:40:00am), Monday, May 12)

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 virtual machine (VM) and virtual machine monitor (VMM)?

(1) VMM is a user space program that runs inside a virtual machine
(2) another name for virtual machine is "hypervisor"
(3) when virtual machine is used, the OS of the real machine is often referred to as the guest OS
(4) a "virtual machine mointor" is a terminal device that's used to interact with a virtual machine
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q2) (2 points) Which of the following statements are correct about thread implementation strategies?

(1) one main problem with the 1 x 1 model is that it's slow because locking and unlocking mutex are slow for this model
(2) one problem with the basic N x 1 model is priority inversion
(3) the scheduler activations model is a variation on the two-level model
(4) one main problem with the 1 x 1 model is that it's slow because system calls are slow
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q3) (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 depends on the ordering of jobs at the run queue
(3) "starvation" at the scheduler is not possible for the FIFO scheduling policy
(4) the FIFO scheduler has the largest variance in waiting time among all scheduling disciplines
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q4) (2 points) Which of the following statements are correct about **terminal driver** and **pseudoterminal driver**?

(1) for a pseudoterminal, the input and output (on the device end) comes from and goes to an actual device
(2) for a pseudoterminal, the input and output (on the device end) comes from and goes to a special kernel-only process
(3) for a pseudoterminal, the input (on the device end) comes from the window manager, which is a user space process
(4) a terminal driver typically runs in kernel space while a pseudoterminal driver typically runs in user space
(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
0x08039000-0x08048000 rw-  PRIVATE 0xcfe0c034 0x0002e 0x08039-0x08048
0x08048000-0x0805e000 r-x  PRIVATE 0xcfe0c004 0x0002d 0x08048-0x0805e
0x0805e000-0x0806f000 rw-  PRIVATE 0xcfe0c064 0x0002c 0x0805e-0x0806f
```

If you get a page fault with vaddr = 0x0805ae82, 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 **microkernel**?

(1) in the design of the microkernel architecture, device drivers cannot be moved into user space
(2) one main differences between a message port and a Unix pipe is that you can assign names to Unix pipes but you cannot assign names to message ports
(3) almost all microkernel implementations have good performance
(4) file access control in a microkernel system typically is based on "capabilities" just like in a traditional Unix system
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q7) (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)	10	14	15	8	-
wt (hrs)					

(Q8) (2 points) Which of the following statements are correct about the **scheduler activations model**?

(1) in the scheduler activations model, the kernel does not schedule/assign CPUs to threads; instead, the kernel schedules/assigns CPUs to processes
(2) the down side of the scheduler activations model is that it's difficult to make round-robin scheduling work well
(3) scheduler activations model is not popular because it's undesirable to let the user-space scheduler of one process to make scheduling decision for another process
(4) the down side of the scheduler activations model is that if a user thread makes a system call, another user thread in the same user process cannot make a system call until the first thread has returned from the kernel
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q9) (2 points) Which of the following statements are correct about **paravirtualization**?

(1) inside a commercial paravirtualized OS, usually there are no device drivers and there are no file systems in the guest/virtualized OS
(2) VirtualBox is well-known for its paravirtualization patent
(3) a paravirtualized OS is indistinguishable from the corresponding real OS in the sense that it can also run on the hardware the real OS was designed to run on
(4) one way to implement paravirtualization is to modify the hardware so that an OS can run inside a virtual machine without modification
(5) none of the above is a correct answer

Answer (just give numbers): _____

(Q10) (2 points) which of the following statements are correct about the **N x 1 (two-level)** thread implementation model?

(1) this model is used in the old days when a user space program didn't know that the kernel was multithreaded
(2) in this model, when a user thread makes a system call and gets blocked inside the kernel, other threads in the same process can still run as long as they don't make system calls
(3) in this model, thread creation and destruction are implemented as user space library function calls
(4) in this model, a user thread does not need to trap into the kernel to lock or unlock a mutex
(5) in this model, when one user thread wants to give up the processor to switch to another user thread in the same process, it must make a system call

Answer (just give numbers): _____

(Q11) (2 points) In **weenix** (which runs on an x86 CPU), when a **user program** makes a **memory reference**, which of the following are possible conditions that can cause a **page fault** and the user program **may not be terminated inside the page fault handler** (assuming that your kernel code is perfect)?

(1) the memory reference is a write operation, the corresponding page table entry's P bit is 1, and the corresponding memory segment's access protection is "read-only"
(2) the memory reference is a read operation (and not an "execute" operation), the corresponding page table entry's P bit is 1, and the corresponding memory segment's access protection is "execute-only"
(3) the memory reference is a write operation, the corresponding page table entry's P bit is 1, and read/write bit is set for "read-only", and the corresponding memory segment's access protection is "read+write"
(4) the memory reference is a write operation, the corresponding page table entry's P bit is 0, and the page faulted virtual address is not within the address range in any of the memory segment
(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 $T_1 = 0.5$, $P_1 = 3.5$, $T_2 = 1$, $P_2 = 4$, $T_3 = 0.5$, $P_3 = 4.5$, and $T_4 = 1$, $P_4 = 5$ (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) (2 points) Which of the following is correct about **management of page frames**?

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

Answer (just give numbers): _____

(Q14) (2 points) Which of the following statements are correct about **LFS (log-structured file system)**?

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

Answer (just give numbers): _____

(Q15) (2 points) Which of the following statements are correct about a **B+ tree of order** $m = 13$?

(1) since m is 13, the root node must have at least 7 child nodes
(2) since m is 13, it's okay for the root node to have 12 child nodes
(3) since m is 13, it's okay for an intermediate node (i.e., neither a root node nor a leaf node) to have 7 child nodes
(4) since m is 13, the height of the B+ tree must be strictly less than 13
(5) since m is 13, it's okay for an intermediate node (i.e., neither a root node nor a leaf node) to have 6 child nodes

Answer (just give numbers): _____

(Q16) (2 points) Which of the following statements are correct about **executing sensitive instructions** on an **IBM 360 machine**?

(1) when a sensitive instruction is executed in the **virtual user mode** inside a **virtual machine**, it will cause a trap and the trap will get delivered to the guest OS
(2) when a sensitive instruction is executed in the **virtual privileged mode** inside a **virtual machine**, it will cause a trap and the trap will get delivered to the guest OS
(3) when a sensitive instruction is executed in the **privileged mode** of the **real machine**, it will cause a trap into the VMM
(4) when a sensitive instruction is executed in the **VMM**, the instruction will get emulated by the VMM
(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 **stride scheduling**. You have decided to give thread A 5 ticket, thread B 6 tickets, thread C 7 tickets, and thread D 10 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 "B:4" for iteration 1 since B is the "winner" of iteration 1 and the winning pass value is 4.) 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	11	(4)	12	16
2				
3				
4				
5				
6				
7				

(Q18) (2 points) Which of the following statements are correct about the basic (two-level) virtual memory scheme where a virtual address is divided into a virtual page number (leading 20 bits) and an offset (trailing 12 bits) on a 32-bit machine?

(1) when performing a virtual to physical address translation, the least-significant 12-bits of the virtual address (i.e., the offset) must equal to the least-significant 12-bits of the translated address
(2) virtual page number can be think of as an array index into a page table which has 2^{20} entries
(3) during address translation, a physical address is obtained by adding a physical page number with the 12-bit offset
(4) during the lifetime of a user process, the mapping of virtual pages to physical pages must stay the same because this mapping is determined by the linker
(5) none of the above is a correct answer

Answer (just give numbers): _____