# EN 600.318/418 Operating Systems: Final Exam

Friday May 9th 2008

## I Frontmatter

1) This exam is take home, open text-book / notes / any literature published on the course website. You may not reference your colleagues' answers or other material.
2) This exam is not timed, but you must answer all questions in one sitting. Please note the time taken to solve the exam accurately on your answer sheet. This will have no bearing on your grade.
3) This exam is due by Monday May 12th 2008, 5pm. Late days cannot be counted for the exam.
4) You must answer all questions.
5) For Yes/No Why/Why-not style questions, you must pick exactly one side, and substantiate it. Ambivalent or unsubstantiated answers will receive no credit.
6) Please submit your solutions to cs418instructors@cs.jhu.edu. You may also submit a paper copy personally to NEB217.

## II Process Synchronization and Deadlocks

Consider the following scenario: Two roads meet at an intersection forming a 4-way stop and each road has a lane traveling in each direction. Four cars are at a complete stop outside the intersection. No car may travel in reverse and all cars are longer than a lane is wide.

1. Use the `TestAndSet` operation to model this problem.
2. Prove that this scenario can enter a deadlock using the bankers algorithm. (Identify resources and show a deadlocked matrix.)
3. For each condition for deadlock, suggest a modification to the problem that will prevent deadlock.

## III Access Control

The following questions are designed to make you think, not write. No answer should require more than one or two paragraphs.

1. We have discussed how domains in the Lampson access matrix can violate confinement using the defined access control operations. Briefly describe the steps taken, and rights required.
2. Throughout this course we have argued against global namespaces. Yet, our discussion of capabilities appears to contain a global namespace: the domains of the Access Matrix. What feature of capability invocation behaves like a local namespace? (i.e. What feature causes this namespace to be irrelevant?)
3. The operations that the textbook defines allow one domain to name another domain, for example to switch domains or perform ownership operations. Consider removing the global namespace of domains from the operations defined in the textbook. How would this impact the functionality and security of the system? Explain.

## IV Virtual Memory

Consider a machine architecture X such that:

- ➔ Register size is 64 bit
- ➔ Virtual address space is 64-bit
- ➔ Physical address space is 32-bit
- ➔ Pages are 16384 bytes
- ➔ The MMU uses a hierarchical page table system for address translation wherein each page table occupies one page.

1) Suppose that the memory map of a process P looks like the following (that is, following virtual address ranges have valid mappings; address ranges are written in hex):

   OS mappings: `[8000 0000 0000 0000 – 8000 0000 000a 0000)`

   Process mappings:

   Code: `[0000 0000 0001 0000 – 0000 0000 ab00 0000)`

   Data: `[0000 00f0 0000 0000 – 0000 00f0 0044 0000)`

   Stack: `[2000 0000 0000 0000 – 2000 0000 0004 0000)`

2) Draw the page table structure for the above process. Please show all details: number of levels, number of entries at each level of the hierarchy, number of bits translated at each level, etc.
3) Suppose that the architecture X used inverted page tables for address translation. Please draw the translation table for process P.
4) Suppose that the architecture X used Guarded page tables (GPT) for address translation. Please draw the GPT for process P.
5) Which of the address translation schemes (among 1, 2, and 3) would you recommend for architecture X? Why?

## V Virtual Machines

A Type-I Virtual Machine is one that runs directly on hardware, and faithfully emulates the hardware functionality. Execution under a VM is said to `go wrong' if the effects of the execution is different from what is guaranteed by the hardware. Suppose that the architecture X mentioned above

- ➔ Maintains two privilege levels: User/Supervisor mode
- ➔ Has TLBs that are hardware managed
  - TLB population is controlled by hardware
  - TLB is not cleared automatically when page tables are switched
  - TLB is cleared when there is a privilege level change
  - A privileged instruction `INVTLB` can be used to clear the TLB
- ➔ Does not provide hardware virtualization support

1) Suppose that we want to construct a Type I VM for architecture X using a shadow paging technique. Please explain the steps involved in the construction of such a VM. You only have to

describe:

- How many page tables are maintained by the VMM
- How and when pages are shadowed
- How the hardware user/supervisor mode switch is handled
- How the guest user/supervisor mode switch is handled
- How the INVLTLB instruction is handled

2) Can the VM constructed in (2) using shadow paging go wrong?

   If so, please construct an example (sequence of operations/events) where the VM goes wrong. How can this problem be solved?

   If not, based on the details of architecture X and your description of the VM in (2), please give an argument as to why all sequences of operations/events are correct with respect to the hardware.

3) Suppose that TLBs in X were software managed instead of hardware managed. Does the answer to question (3) change? Why or why not?

## VI File Systems

1) Consider a classic UNIX file system (FFS). Suppose that we have a file whose size is 1347 blocks. Draw a picture of the file system showing the blocks in the file and the corresponding meta-data blocks (inode and indirect blocks).
2) Suppose that we perform a series of writes that modify block 2 and block 1270 of that file. Identify all of the blocks in the file system that will be re-written (including bitmap blocks), and give the order in which these rewrites must occur to preserve consistency.
3) Briefly describe the order of updates that must be performed for the same file if we had a file system that supports
   1. Journaling
   2. Soft-updates
   3. Write anywhere file layout (WAFL).
4) Suppose that there was a system crash in the process of the above write. Upon restart, some of the above schemes require a file system integrity check (FSCK). For each file system above (FFS, journaling, soft-updates and WAFL) identify whether FSCK is necessary at startup, and explain why or why not. If FSCK is necessary, how should the check be performed (that is, what is compared to what)?
5) Suppose that you were to pick a file system for an online server in which:
   1. File-system operations are dominated by reads
   2. Writes to the file system are idempotent
   3. File system consistency/recovery is essential

   Which of the above file systems (or some hybrid combination of the same) would you prefer for such a server? Why?

** * * * HAVE A GOOD SUMMER * * * **