CS-736: Midterm Exam (Fall 2000)

	Points	Total Possible
Writing your first and last name on all **7** pages		1
Part I: Short Answers		$(17 \times 3) \rightarrow 51$
Part II: Long Answers		$(3 \times 16) \rightarrow 48$
Total		100

Name: _____

Part I: Short Questions

The following questions are multiple choice. There are two styles of questions. In the first, there is one and only one "right" answer. In the second, there are between one and four right answers - you should circle all that apply. Ones of the second style are marked as appropriate.

Each of the seventeen (17) questions is worth three (3) points for a total of fifty-one (51) points in this section.

1. In Pilot, a program can give advice to the operating system in order to tailor the OS's behavior to the program's liking. The Space.Activate() call on a region of a mapped file results in

   (a) the OS immediately fetching those pages

   (b) the OS knowing that those pages should be fetched soon

   (c) in-kernel data structures related to the space getting marked as "hot"

   (d) the pages of the space are all moved back on the global replacement list

2. Assume there are three threads in a process, labeled A, B, and C. The process wishes to allocate 50% of its time to thread A, and let B and C share the other 50%. Assuming that we are using resource containers to account for CPU time of the jobs, what is the minimal number of resource containers that are required for this?

   (a) 1

   (b) 2

   (c) 3

   (d) 1.5

3. In Elephant, three policies were used to differentiate between historical accesses. "Keep all" keeps all versions of a file ever created; "Keep Safe" just keeps enough data around such that undo is possible; "Keep Landmarks" keeps a pruned history with landmark versions intact. If "Keep Safe" and "Keep Landmarks" were removed from the system, in which of the following ways could the system be simplified? **(Circle all that apply)**

   (a) No need for inode logs

   (b) No need for copy-on-write

   (c) No need for versioned directories

   (d) No need for the cleaner

4. Which of the following are true for the VAX/VMS system? **(Circle all that apply)**

   (a) simple FIFOs are used to track process resident sets

   (b) a memory hog will not get more memory than other processes

   (c) the $system$ base register contains a virtual address

   (d) the P_0 and P_1 base registers contain virtual addresses

5. In the original Unix system, the i-number (or "index number") of a file can be used to locate the file's inode

   (a) by scanning the free list

   (b) by checking the bit map

   (c) by performing a simple "base + offset" calculation

   (d) not at all (the i-number can't be used to directly find the inode)

6. In the exokernel, an "abort" protocol is sometimes used when the exokernel wants to take back some resources from a libOS. In which of the following situations could an abort be required? **(Circle all that apply)**

   (a) The exokernel wants the libOS to give back a page in a certain amount of time, and the libOS is tardy in doing so

   (b) The exokernel wants the libOS to give back a disk block in a certain amount of time, and the libOS is tardy in doing so

   (c) The exokernel wants the libOS to give up the CPU, and the libOS is tardy in doing so

   (d) The libOS is multiplexing between two application-level threads, and one thread is not yielding the processor

7. Multics procedure segments are said to be "pure". What does that mean? **(Circle all that apply)**

   (a) They are generated by a compiler

   (b) When executed, they will run to completion

   (c) When executed, they will not write into other segments' data sections

   (d) When executed, they will not change their own contents

8. Unix FFS used which of the following techniques to reduce the ill effects of *disk seeks* that were prevalent in the old Unix file system? **(Circle all that apply)**

   (a) Bit maps

   (b) Sub-blocks

   (c) Cylinder groups

   (d) Rotational layout tables

9. When calling read() from an application to read a single data block from an open file in LFS, how many disk reads could occur? **(Circle all that apply)**

   (a) Zero

   (b) One

   (c) Two

   (d) More than two

10. In Nucleus, assume a server process implements a user-level file system. What prevents a malicious process from sending a "fake" response from the file system to a client waiting for a response?

    (a) Mutual trust among processes

    (b) The client waits for a particular buffer from the server

    (c) Only the server can send send a message to waiting clients

    (d) The Nucleus kernel monitors all communication carefully

11. Which of the following is true for RAID-5 as compared to RAID-4? **(Circle all that apply)**

    (a) RAID-5 has potentially higher read bandwidth

    (b) RAID-5 has potentially higher write bandwidth

    (c) RAID-5 has potentially lower read bandwidth under a single failure

    (d) RAID-5 potentially uses more capacity

12. In Hydra, if a process has a reference to an object, it can for certain

    (a) modify the object's entire representation

    (b) modify only the object's data part

    (c) destroy the object

    (d) perform operations in accordance with the access rights of the reference

13. In which of the following aspects would you say a deterministic proportional-share scheduler has *clear benefits* over randomized, lottery-based scheduling? **(Circle all that apply)**

    (a) predictability

    (b) amount of state per job in system

    (c) precision in allocation of proportions

    (d) amount of time to pick who runs for the next quantum

14. Mach and Nucleus are both quite similar in basic structure. Which of the following are enhancements that Mach provided that were not available in Nucleus? **(Circle all that apply)**

    (a) A communication system that is tightly integrated with virtual memory

    (b) Communication between processes that is kernel protected

    (c) Some typical "system" services can be provided by user-level processes

    (d) Multiple threads of control per address space

15. The THE system is organized in a strict hierarchy, virtualizing aspects of the machine one at a time, with higher-level abstractions built on top of the earlier layers. The lowest level of virtualization ("level 0") is of:

    (a) the processor

    (b) core memory

    (c) the drum

    (d) the keyboard console

16. A TLB miss occurs in an application running on Irix running on top of Disco. Which of the following represents a legal control flow? **(Circle all that apply)**

    (a) Irix TLB miss handler → resume application

    (b) Disco → resume application

    (c) Disco → Irix TLB miss handler → Disco → resume application

    (d) Disco → Irix TLB miss handler → Irix page fault handler → Disco → resume application

17. In Scheduler Activations, assume that we have a process running on four processors (and hence 4 scheduler activations). The kernel then decides to preempt two of those processors to give to another process. How many processor preemptions are required to remove two processors from the process?

    (a) 1

    (b) 2

    (c) 3

    (d) 4

Part II: Longer Questions

The second half of the exam consists of three longer questions. Try to sketch out the basic idea of your answers before getting bogged down in details.

Each of the three (3) questions is worth sixteen (16) points for a total of forty-eight (48) points in this section.

1. **A Question Of Structure.**

   In this question, we explore using the mechanisms in Hydra as building blocks for operating system kernels. Hydra provides a generalized protected control transfer known as a *call* and an early form of object oriented programming to enable programmers to assemble whatever type of operating environment they would like.

   a): Describe the call mechanism. How is it similar to and different from a trap into a modern-day operating system? How does one increase privilege across a call boundary?

   b 1): Assume we wish to build a communication infrastructure similar to that of the Nucleus system with the Hydra primitives. First, describe the nucleus communication system.

   b 2): Now, describe how to build the Nucleus communication system on top of the Hydra primitives. What objects would you provide? Where would you use the call mechanism?

2. **Be Very Af-RAID.**

Assume you have a RAID-5 system with D data disks and 1 parity disk. Each disk has B blocks of capacity, and all disks are of identical make, model, and performance. Unfortunately, one of your data disks fails, and you replace it with a brand new disk. The system must then carry out the task of updating the new disk so that the RAID-5 returns to its normal "working" status (*i.e.*, all the data disks have all their data and the parity disk has all its necessary parity information), a process known as *reconstruction*.

a): Describe how the process of reconstruction would work for RAID-5.

b): During this reconstruction, how many blocks must be read from the RAID-5 storage system, and how many written?

c): Now assume that we had instead created a storage system with mirroring instead of RAID-5, with a total of 2D disks for the system. Describe how reconstruction would work in the mirrored system.

d): Which is faster - reconstruction in RAID-5 or reconstruction in the mirrored system? (justify)

3. **It's Your Lucky Day.**

   In this question, we explore the application of lottery scheduling to disks. Assume we have machine with a single disk attached, and we want to provide applications with proportional share of the disks both in terms of capacity and performance.

   a): In CPU lottery scheduling, a single list was used to track each job; each node of the list held the number of tickets allocated to that job. What kind of data structures would you use to organize the disk-based lottery? (describe and justify)

   b): With CPU lotteries, a lottery was held every quantum to determine who would run for the next quantum. When applying lottery scheduling to disk performance, how often would you hold a lottery? (justify)

   c): Any time we are applying a proportional-share approach, the effect on performance comes into question. In the CPU-based system, how did lottery scheduling affect overall system performance? What kind of effect could a proportional-share system have on overall disk performance?

   d): Finally, we also wish to use lotteries to proportionally share disk capacity. When would we hold a lottery in this case? What other issues arise in using lotteries for sharing capacity instead of performance?