Name: ______________________________

CS 4/53201 **Final Exam** **Operating Systems**

Wednesday 13 December 2000

---

1. **At the beginning of the semester, I talked about how the operating system hides many of the details of the underlying hardware from the user.**

   a. **The operating system makes use of (and can inspire) features of the CPU architecture. Briefly describe <u>two</u> specific features that a CPU architecture can provide, and explain how the operating system takes advantage of those features. (10 points)**

   Bases and limit registers, memory management unit, etc. — provides support for complex memory management schemes such as segmentation and paging

   Read-modify-write instruction such as test and set — useful for implementing semaphores

   Interrupts — useful for implementing semaphores in a uniprocessor environment

   Timer — useful for implementing round-robin CPU scheduling

   User / kernel mode instructions — disallows user access to certain instructions

   b. **The operating system can also hide limitations of the underlying hardware. Give <u>one</u> example of a limitation that is built into a machine, and how the operating system makes the user unaware that the limitation exists. (5 points)**

   By providing demand paging and swapping, the OS can allow a larger address space than supported by physical memory.

   By providing time sharing, the OS gives the illusion of a dedicated CPU to each process / thread.

2. **One of the most important concepts examined in this class is the process.**

   a. **Draw the five-state process model, and clearly label all states and transitions (a detailed explanation is not needed, just simple labels). (15 points)**

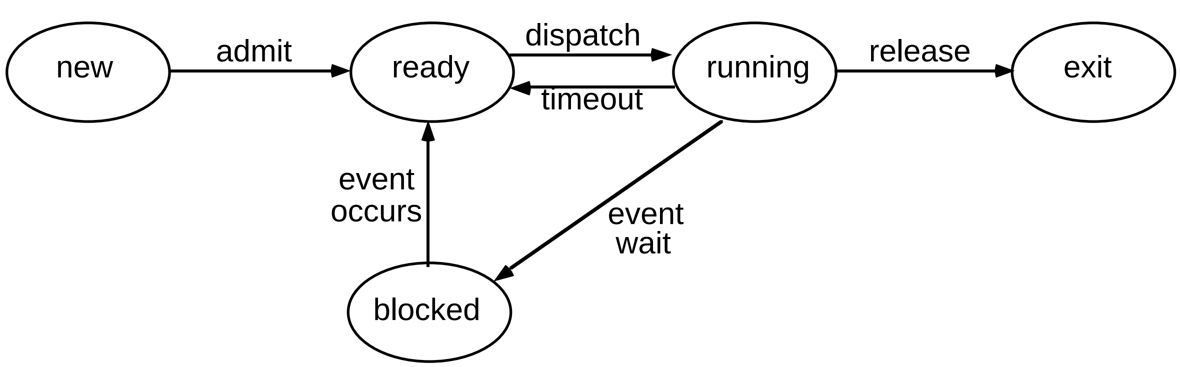

b. Explain the "blocked" state — what it means for a process to be blocked, why it would be blocked, and how it would get unblocked. (10 points)

A process goes into the blocked state while waiting on an I/O event, or a semaphore or condition variable signal. While in that state it sits on a queue, waits, and does not run. It unblocks when the event it is waiting on occurs.

3. One of the algorithms that can be used for memory management, partitioning in particular, is Best Fit.

a. Briefly explain this algorithm. (10 points)

Search all partitions, and of those sufficiently large enough, choose the smallest.

b. Explain why this algorithm may not be the best choice for memory management. (5 points)

Choosing block sizes as tightly as possible tends to leave many very small free memory partitions that are too small to be useful for another process, hence a large amount of external fragmentation. The sorting necessary for making the best choice is also more time-consuming than some other algorithms such as First Fit.

4. A central component of any file system is the file descriptor (which is called the "inode" in UNIX).

a. Circle those items in the following list that are typically stored in the file descriptor: (12 points)

file name	(access permission)	(file owner)
(file size)	read/write position in file	(pointer to blocks on disk)

b. Where are the inodes stored in UNIX? Be specific. (Hint: memory, disk) (8 points)

On the disk, the inodes are stored in the ilist, which is at the beginning of the partition.

In memory, inodes that correspond to active files are stored in the active inode table.

c. How does a UNIX inode point to the disk blocks where the corresponding file is stored? Don't worry about specific numbers, but describe enough of the concept to convince me that you understand it.. (10 points)

An inode contains a set of pointers that point to specific disk blocks. If more disk blocks are needed, another pointer points to a single indirect block, which contains pointers to yet more disk blocks. If more disk blocks are needed, another pointer points to a double indirect block, which contains pointers to yet more single indirect blocks. In earlier versions of UNIX, there was also a triple indirect block, but that is no longer used in modern versions.

Name: ________________________________

5. **Suppose a given disk has 200 tracks, numbered 0 to 199. The head has just finished accessing track 125, and is now accessing track 143. The queue of requests, in FIFO order is as follows (86 appeared first): 86, 147, 91, 177, 94, 150, 102, 175, 130.**

   a. **List *all* the head movements needed to satisfy the requests, using *FCFS* disk head scheduling. (5 points)**

      143 -> 86 -> 147 -> 91 -> 177 -> 94 -> 150 -> 102 -> 175 -> 130

   b. **List *all* the head movements needed to satisfy the requests, using *SSTF* disk head scheduling. (5 points)**

      143 -> 147 -> 150 -> 130 -> 102 -> 94 -> 91 -> 86 -> 175 -> 177

   c. **List *all* the head movements needed to satisfy the requests, using *LOOK* disk head scheduling. (5 points)**

      143 -> 147 -> 150 -> 175 -> 177 -> 130 -> 102 -> 94 -> 91 -> 86

6. **TCP is a popular protocol used on the internet at the transport layer.**

   a. **What are the responsibilities of the transport layer / TCP? (10 points)**

      To reliably transfer messages (broken into packets) between hosts, with error control for out-of-sequence and missing packets

   b. **How can TCP be said to be "reliable"? (5 points)**

      Because packets must be acknowledged, and if the receiver does not acknowledge a packet after a short period, it is retransmitted. Thus the sender can be assured that the data reached the destination, or else the sender is notified that there was an error.

7. **Define and distinguish between these three concepts: SIMD, MIMD, distributed system. (15 points)**

   SIMD — one control unit supplies instructions to all the processing elements simultaneously, so each is doing the same thing, but on different data (each processing element is operating on data in its own memory)

   MIMD — each processing element has its own control unit, so each is potentially doing something different, and on different data (each processing element is operating on data in its own memory)

   Distributed system — a MIMD system where the processing elements are loosely connected with a unifying OS of some sort

Name: ______________________________

8. Sun's NFS supports mounting of remote file systems by client machines.

a. Draw and explain a diagram illustrating this. (15 points)

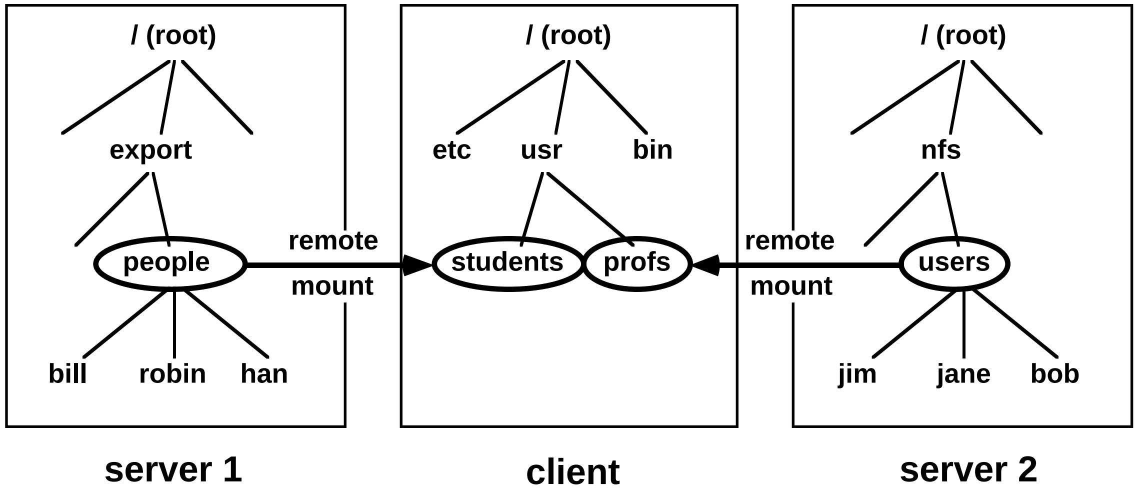

A directory on a file system (such as "users" on server 2 in the figure above) is mounted by the client onto its file system (at location "profs" in the figure above). After this mounting, references on the client to a directory such as /usr/profs/bob will actually refer to /nfs/users/bob on the remote file system.

b. What information must the client OS maintain? (5 points)

Local names and mount point, and corresponding name/path on remote machine.