

**Tuesday 10 December 2002**

---

- 1. To protect against damaging actions by users, certain CPU instructions are considered to be privileged instructions, meaning they only run when the CPU is set to kernel / monitor mode. How does the CPU get set to kernel mode? Be specific. (10 points)**

When a user program makes a system call or causes a page fault, a trap instruction is executed, which transfers control to a trap handler and sets the CPU to kernel mode.

- 2. Consider a time-sharing OS with a round-robin CPU scheduler.**

- a. Briefly summarize this scheduling algorithm. (10 points)**

The process at the head of the ready list is chosen, and run for at most one time slice. If it terminates or blocks before the end of the time slice, the next process at the head of the ready list is chosen; if it reaches the end of the time slice and is still running, it is preempted and moved to the end of the ready list.

- b. What is the effect if the time slice is chosen to be too short or too long? (5 points)**

If the time slice is too short, too much time is wasted in context switches. If it is too long, then the algorithm is essentially FCFS scheduling and the interactive rapid response is lost.

- 3. When a thread performs a semaphore Wait operation, it may have to block. How does being blocked on a semaphore Wait compare to being blocked waiting on an I/O operation to complete? (5 points)**

In either case, the thread is in the blocked state, waiting on some outside event to occur so that it can wake up. *Different details depending on how you interpreted the question...*

- 4. Deadlock can be eliminated if the “mutual exclusion” criterion can be avoided. Explain how printer spooling can avoid mutual exclusion and thus prevent deadlock from occurring. (10 points)**

Printer spooling prevents deadlock by pretending the printer is sharable — thus eliminating the need for mutual exclusion. What actually happens is that a process sends print jobs to a print sever, which can accept multiple jobs simultaneously (make it seem that the printer is being shared), but which then actually sends those jobs sequentially to the printer.

- 5. Consider the memory management technique of paging and the technique of allocating disk space in fixed-size blocks.**

Name: \_\_\_\_\_

**a. In what ways are these two techniques similar? (10 points)**

Both allocate data in fixed-size chunks, which avoids external fragmentation. However, both suffer from internal fragmentation.

**b. In what ways are these two techniques different? (10 points)**

In paging, it really doesn't matter where a page is stored in memory — accessing one frame is as fast as any other. On the disk, due to seek / rotation time, access is faster if the blocks are store contiguously, or at least close together.

**6. In demand paging, the ideal page replacement algorithm is the Optimal algorithm, where the page evicted is the page that will be referenced the farthest in the future. Since this algorithm is impossible to implement, the Least-Recently-Used (LRU) algorithm is considered instead.**

**Explain why the LRU algorithm should approximate the Optimal algorithm. (10 points)**

The Principle of Locality of Reference says that if a data item is needed now, chances are it'll be needed again in the future, so if a page hasn't been used recently, chances are it won't be needed again anytime soon.

**7. There are many techniques for organizing the disk blocks that comprise a file.**

**a. What advantages does continuous allocation have over indexed allocation? (10 points)**

The implementation is simple, storing only a pointer to the first block and the length, instead of complicated index structures. This simplicity uses less disk space for those data structures, and the small number of seeks means it can access the file quicker (i.e., there isn't a need to keep going back to index to get a pointer to the next block).

**b. What advantages does indexed allocation have over continuous allocation? (10 points)**

With indexed allocation it is easier to place or grow a file since there isn't a need for a large contiguous space, and indexed allocation also avoids the external fragmentation that is possible in continuous allocation. Indexed allocation is also more efficient for random access than continuous.

**c. How does multilevel indexed allocation improve on indexed allocation? (10 points)**

It's more efficient for small files, since the pointers can be stored in the file descriptor rather than a separate index block. It also provides better support for large files, since additional levels of indexing can be added as necessary.

8. The time required to access a particular block of data on a disk depends on the seek time, rotational delay, and data transfer rate. For each of these three factors, how can performance be increased? (15 points)

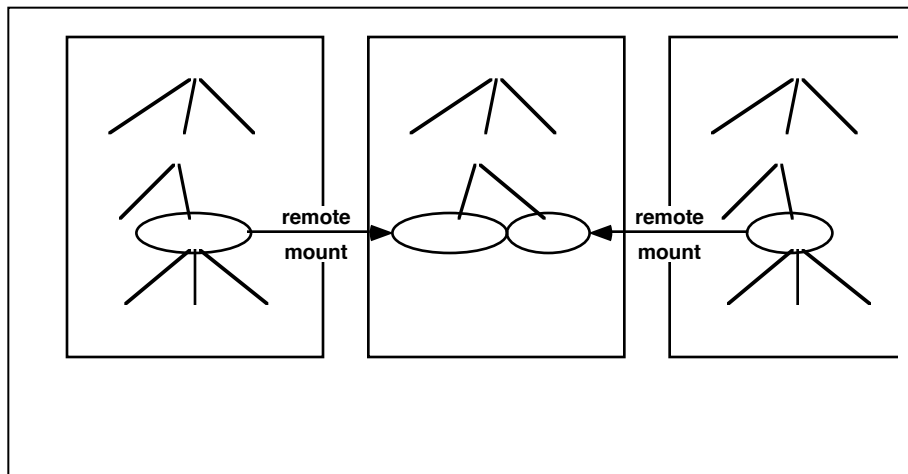
Data transfer rate — faster hardware, large blocks.

Seek time — faster hardware, disk head scheduling algorithms, contiguous allocation, allocation to same cylinder or nearby.

Rotational delay — faster hardware,

9. Explain the concept of mounting a remote file system. Draw a picture if you think that might aid in your explanation. (15 points)

A portion of a remote file system, for example the file system tree under /export/people from server 1, is mounted onto the client file system, for example at location /usr/students. If this is done, then a file access to /usr/students/bill on the client machine actually accesses Bill's files on server1.



10. Linux is distributed under the GNU General Public License (GPL). What does this mean? (10 points)

The code can be used and modified by anyone. None of the code can be turned into a commercial product, though it can be redistributed for a reasonable fee. If the code is distributed, the source code must always be distributed with compiled binaries.

11. When a new version of Windows comes out, why is it likely to run programs compiled under earlier versions of Windows (even old DOS program), yet it may require a new device driver for a printer that you bought only last year? (10 points)

Windows XP includes environmental systems to emulate Windows 95/98, DOS, OS/2 etc., so that binaries compiled for those systems should still run even though the new version of Windows has a new kernel. However, since a new Windows kernel may have new functionality, system calls, etc., device drivers (which interact directly with the OS kernel and the hardware abstraction layer) must often be rewritten.

Name: \_\_\_\_\_