EXAM ON FILE SYSTEMS AND I/O – SPRING 2004
CSC 262 – OPERATING SYSTEMS
NICHOLAS R. HOWE

1. **Layered Operating Systems** (12 points). Some operating systems are organized into layers for greater modularity of design. For example, we have looked at systems with four software layers: kernel, device drivers, service, and user-level. Specify which layer should perform each of the activities listed below.

a. Move comb over a specified cylinder group.
Driver level
b. Delete a file on the disk.
Service level
c. Suspend a process that has requested information from a disk, until it is available.
Kernel
d. Transfer data from specified sectors on a disk into main memory.
Driver level
e. Send e-mail over a network.
User level
f. Report that a requested directory cannot be found.
Service level

2. **Hierarchical File Systems** (12 points). Consider the three disk volumes shown below. Assuming that hda is mounted at **/**, hdb is mounted at **/usr**, and hdc is mounted at **/usr/local**, draw a diagram of the resulting file system.

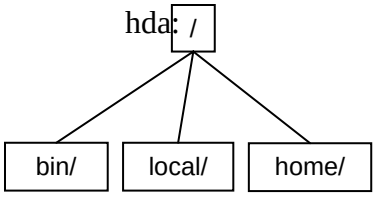

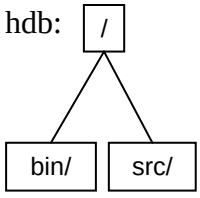

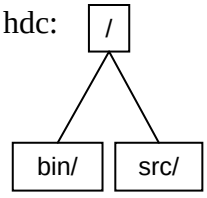

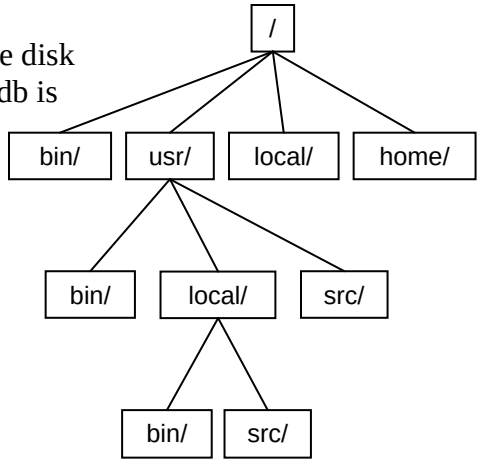

3. **Performance Measures** (8 points). A given computer has 10 jobs to do, each of which requires 100 milliseconds of CPU time. Your job is to compute utilization and throughput under the following circumstances:

a. If run sequentially, each job takes 250 milliseconds to complete.
Utilization = 1000/2500 = 40%, Throughput = 10 jobs/2.5 s = 4 jobs/s
b. If run in parallel, all the jobs finish simultaneously in 1.25 seconds.
Utilization = 1000/1250 = 80%, Throughput = 10 jobs/1.25 s = 8 jobs/s

4. **Disk Scheduling** (16 points). Consider the table below, which shows a series of requests for access to particular disk tracks, along with the time at which the request is made. Assume that the action begins at 0 ms with the first two requests already waiting and the disk head at track 0. Servicing each request takes exactly 5 ms, and that no request may be satisfied before the time at which it is been issued. Compute (i) the **order** in which the track requests would be satisfied, and (ii) the **latency** of each individual request.

Track:	12	24	15	11	19	8	18
Time:	-2 ms	-1 ms	2 ms	3 ms	6 ms	8 ms	16 ms

a. FCFS *Order: 12, 24, 15, 11, 19, 8, 18* *Latency: 7, 11, 13, 17, 19, 22, 19*

b. SSTF *Order: 12, 11, 8, 15, 18, 19, 24* *Latency: 7, 7, 7, 18, 9, 24, 36*

c. SCAN *Order: 12, 15, 19, 24, 18, 11, 8* *Latency: 7, 8, 9, 21, 9, 27, 27*

d. F-SCAN *Order: 12, 24, 19, 15, 11, 8, 18* *Latency: 7, 11, 9, 18, 22, 22, 19*

5. **File Systems** (12 points). MiniFS is a Unix-like file system built with extremely tiny disk blocks: only 4 bytes! Each MiniFS disk has a maximum of 255 data blocks, so a pointer to a block takes exactly one byte. I-nodes in MiniFS have one direct pointer, one single indirect pointer, one double indirect pointer, and one triple indirect pointer. The data blocks are numbered starting with 1, since a block number of 0 indicates a null pointer.

a. What is the maximum file size, in bytes, for MiniFS? *85*4 = 340 Bytes*

b. Given the i-node at right and data blocks shown below, what are the contents of the specified file? List all the **data block numbers** in order. Do not include any pointer blocks or blocks belonging to other files.
Block numbers: 3, 2, 7, 5.

i-node
3
1
0
0

Block 1	*Block 2*	*Block 3*	*Block 4*	*Block 5*	*Block 6*	*Block 7*	*Block 8*
2	5	8	0	3	8	1	6
7	4	3	0	1	6	5	7
5	6	1	3	0	4	3	8
0	1	0	2	0	3	7	1

6. **Protection & Security** (16 points). Some systems use a simplified security system based upon *protection rings*. In this scheme, domains are arranged in a hierarchy, and the capabilities of each domain include those of all domains lower in the hierarchy. This allows for an efficient representation in terms of capability lists, because each domain only needs to list the new

privileges gained over the next lower domain in the hierarchy. Similarly, access control lists need only to list the minimum domain level required for each operation on a resource.

Suppose that we have a system that uses protection rings for security. It has four levels, *Guest*, *User*, *Supervisor*, and *Root* (in increasing order of privilege). Consider the resource ACLs below, specified according to the convention in the preceding paragraph. Convert this to a traditional access control matrix that would be understood by a system that doesn't use protection rings. Use the * notation for the ability to confer rights.

PasswordFile: Read = Supervisor, Write = Root
DataFile: Read = Guest, Write = User, {ConferRead,ConferWrite} = Supervisor
ProgramFile: Execute = Guest, Write = Supervisor
Printer: Write = User, ConferWrite = Supervisor

	PasswordFile	*DataFile*	*ProgramFile*	*Printer*
Guest	--	R	X	--
User	--	RW	X	W
Supervisor	R	R*W*	WX	W*
Root	RW	R*W*	WX	W*

7. **Vocabulary** (16 points). Below is a schematic diagram of a disk drive. Name the parts illustrated. Each letter refers to a different item.

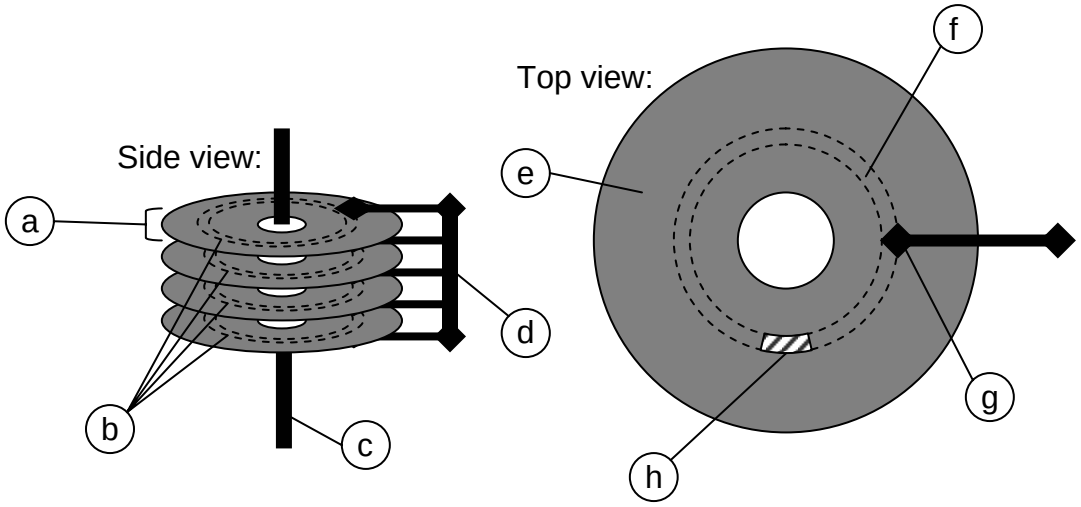

a = platter, b = cylinder, c = spindle, d = comb, e = surface, f = track, g = read/write head, h = sector.

8. **History** (8 points). Classify the following items into one of the following generations: I (1945-55), II (1955-65), III (1965-80), or IV (1980+)

_*II*_ a). Batch readers collect jobs waiting to run

_*I*_ b). Computers based on vacuum tubes

_*IV*_ c). Desktop workstations run multiple programs at once

*III* d). Mainframe computers allow timesharing between multiple users

_*II*_ e). No other processes can run while the current process uses the “line printer”

*III* f). Memory protection mechanisms first introduced

*III* g). The first personal computers are available to the public

_*IV*_ h). Personal computers more powerful than Cray-1 supercomputer are available