B3.4-R3: OPERATING SYSTEMS

NOTE:

1. **Answer question 1 and any FOUR questions from 2 to 7.**
2. **Parts of the same question should be answered together and in the same sequence.**

Time: 3 Hours **Total Marks: 100**

1. Answer the following questions.

a) Briefly describe the four major resource managers in a typical operating system.

b) Discuss why we have different processor modes and how these modes are used in typical operating systems.

c) Describe the similarities and difference between two different threads running in the same process and two independent processes.

d) What is the difference between a physical address and a virtual address?

e) How can a CPU be brought from user mode to kernel mode? How it can be brought back from the kernel mode to user mode? Justify your answer.

f) In what circumstances are interrupts more efficient than polling?

g) What is the difference between multiprocessing and multiprogramming?

(7x4)

2. As a system administrator you have noticed that usage peaks between 10:00 AM to 5:00 PM and between 7:00 PM to 10:00 PM. The company's CEO decided to call on you to design a system where during these peak hours there will be three levels of users. Users in level 1 are to enjoy better response time than users in level 2, who in turn will enjoy better response time than users in level 3. You are to design such a system so that all users will still get some progress, but with the indicated preferences in place.

a) Will a fixed priority scheme with pre-emption and 3 fixed priorities work? Why, or why not?

b) Will a UNIX-style multi-feedback queue work? Why, or why not?

c) If none of the above works, could you design a scheduling scheme that meets the requirements?

(5+5+8)

3.

a) Identify and explain three styles of switching from user mode to supervisor mode.

b) A hardware designer argues that he has enough transistors on the chip to provide 1024 integer registers and 512 floating point registers. You have been invited as the operating system guru to give opinion about the new design.

   i) What is the effect of having such a large number of registers on the operating system?

   ii) What additional hardware features you would recommend added to the design above.

(6+12)

4.

a) Most round-robin schedulers use a fixed size quantum. Give an argument in favor of and against a small quantum.

b) Consider a uniprocessor kernel that user programs can trap into using system calls. The kernel receives and handles interrupt requests from I/O devices. Would there be any need for critical sections within that kernel?

c) How one can prevent deadlock to take place in an operating system? Explain.

(6+6+6)

5.

a) Define critical section problem for process synchronization. Explain why it is difficult to implement critical section problem.

b) Consider the following program fragment:

```
P (s1);
a++;
P (s2);
v++;
V (s2);
V (s1);
```

(s1, s2 are semaphores). All variables are automatic. Now, consider two threads running this fragment of code simultaneously, can there be a deadlock? Why, or why not.

(9+9)

6.

a) A minicomputer uses the buddy system for memory management. Initially it has one block of 256K at address 0. After successive requests of 7K, 26K, 34K and 19K come in, how many blocks are left and what are their sizes and addresses?

b) Explain working of virtual memory.

c) A computer whose processes have 512 pages in their address space keeps its page tables in memory. The overhead required for reading a word from the page table is 50 nsec. To reduce this overhead, the computer has a TLB that holds 32 entries and can do lookups in 5nsec. What hit rate is needed to reduce the mean overhead to 20nsec?

(6+6+6)

7.

a) Explain salient features of I/O system of a general-purpose operating system.

b) Contiguous allocation of files leads to disk fragmentation. Is this internal or external fragmentation?

c) Can we implement symbolic links in DOS (FAT file system)? If so, show how, and if not, explain why.

(12+3+3)