NATIONAL EXAMINATIONS MAY 2018

17-COMP A-5, OPERATING SYSTEMS

3 Hour Duration

NOTES:

1. If doubts exist as to the interpretation of any question, the candidate is urged to submit with the answer paper a clear statement of any assumption made.

2. Provide justifications for your answers. Show all your work.

3. CLOSED BOOK. Candidates may use one of the two calculators, the Casio approved model or Sharp approved model. No other aids.

4. The candidate has to answer **any five questions**. Each question has multiple parts. The marks for each part is indicated within square brackets.

5. Total Marks = 100.

6. This exam has got 5 pages (including this page).

1 [20 marks]

[9] **(a)** Consider the following solution to the critical section problem involving two concurrent processes P0 and P1. Identify as many **distinct problems** as you can in the design. If similar problems occur at multiple places identify them each time but explain it only once. Your list of errors should include defects (if any) that may not necessarily give rise to incorrect results but do indicate flaws in design.

Justify your answer with the help of examples. Be as specific as you can when you describe the situations in which problems occur.

Algorithm

Process P0

```
do {
        guard0 = true;
        if guard1 {
                guard0 = false;
                while guard1 no-op;
                guard0 = true;
        };

        Code for CS

        guard 0 = false;
} while (true) ;
```

Process P1

```
        do {
         guard1 = true;
         if guard0 {
                guard1=false;
                while guard0  no-op;
                guard1 = true;
        };

                Code for CS

                guard1= false;
        } while (true);
```

Note: CS: Critical Section. no-op: no operation.

[6] **(b)** Briefly discuss whether or not each of the three requirements associated with the solution to the critical section problem is satisfied when a monitor is used to guard the access to the critical section.

[5] (c) What is a real time system? With the help of examples distinguish between a hard and a soft real time system.

2 [20 marks]

[12] **(a)** Consider a **preemptive** short term scheduling strategy in which the priority of a process may change dynamically with time. (Larger priority numbers imply higher priority). At any point in time the highest priority process is run on the system. Ties are broken in favour of the process that entered the ready to run queue first. If a running process is preempted then its time of entry into the ready to run queue is the time at which the preemption was made.

When a process enters the ready to run queue its priority is set to **P.** When the process is waiting in the ready to run queue its priority changes at a rate **a** as given by the following equation.

Priority of a process in the ready to run queue = $P + a.t^2$
(where t is the time elapsed (in seconds) after the process entered the ready to run queue).

When a process is selected to run on the CPU its priority is set to **Q**. As it starts running on the CPU its priority changes at a rate **b** as given by the following equation.

Priority of the process running on the CPU = $Q - b.t'$
(where t' is the time elapsed (in seconds) after the process started running on the CPU).

The parameters P, Q, a, and b can be set to give many different scheduling policies. Once chosen the values of these parameters become fixed and cannot change.

Determine P, Q, a, and b that will produce the Last Come First Served policy.

Under this policy whenever a process arrives on the system it preemptively captures the CPU. That is, if the CPU is free the process is allocated the CPU; if the CPU is busy then the executing process is preempted and the CPU is allocated to the process that just entered the ready to run queue. Whenever a process completes its CPU burst and the CPU is free, the process (in the ready to run queue) that was preempted most recently is allocated the CPU.

(b)

[5] (i) Briefly discuss the benefits and the shortcomings of the Round Robin CPU scheduling policy.

[3] (ii) What is the impact of the duration of time slice on the performance of the Round Robin CPU scheduling policy?

3. [20 marks]

[5] **(a)** Briefly discuss the locality of reference for programs (also known as the principles of locality) that underlie the memory referencing behavior of a program when it executes on the system.

[5] **(b)** Discuss "thrashing" in the context of memory management and explain its relationship with program locality.

[10] **(c)** Discuss with the help of examples how the *working set based memory management technique* uses the principles of locality in improving system performance and in effectively controlling thrashing.

4. [20 marks]

[6] (a) Describe the components of a resource-allocation graph used in the context of deadlocks.

Does a cycle in the resource allocation graph always imply the existence of a deadlock? Justify your answer with the help of examples.

[8] (b). Different approaches are available for the handling of deadlocks on the system. Using examples briefly discuss the "deadlock prevention" approach.

[6] (c) Consider a moving head hard disk which consists of a single platter (surface) with 250 tracks on it. The tracks are numbered 0 to 249. The disk has just served a request on track 95 and is currently serving a request at track 153. The queue of pending requests in FIFO order is:

94, 175, 80.

What is the total head movement (in number of tracks) needed to satisfy all these requests for the following disk scheduling algorithms (i) C-LOOK and (ii) Shortest Seek Time First (SSTF)?

[Assume that no further requests arrive on the system during the service of the above requests.]

5. [20 marks]

(a) Consider a demand paged virtual memory system and the following page reference string:

271, 272, 273, 274, 375, 273, 274, 271, 376, 277, 278, 277, 278, 279, 277, 278, 279

[4] (i) What is the minimum number of page faults for this reference string that can occur on the system?

[10] (ii) Determine the number of pages faults for the First In First Out (FIFO) page replacement policy when 4 frames are allocated to the program.

[6] (b) Consider the following sequence of memory addresses referenced by a program:
25, 175, 178, 177, 282, 283, 285, 25, 196, 199
Note that each address refers to a word stored in the corresponding memory location.

Determine the page reference string (reference string) that contains the corresponding sequence of page numbers when the page size is 75 words.

6. [20 marks]

[4] (a) Discuss the role of access control in the context of file systems that are used by multiple users.

[6] (b) Discuss with the help of examples any one technique used for performing access control on a multi-user file system.

[10] (c) This question focuses on two methods for storing information on the disk. Consider a file currently consisting of 150 blocks (numbered 1 - 150). Assume that the directory is available in main memory.
Compare the performance of the contiguous and linked block allocation techniques for each of the following cases. (i) For each case (A-C) compute the minimum number (M) of disk operations that are required when contiguous allocation and when linked allocation is used. (ii) Then determine which one of the techniques led to a lower average M. [Note that the average for any given technique is computed over the three cases A-C].

(A) The contents of block 131 and 95 are read.
(B) The contents of block 149 are exchanged with the contents of block 90.
(C) Block 95 is deleted.

Consider each case (A-C) separately. Note that each disk operation corresponds to the reading of a block from the disk or the writing of a block to the disk. While computing the number of disk operations, ignore the disk operations that may be required for the location and maintenance of free space. Since the directory is in main memory any operation on the directory is not counted as a disk operation.
ASSUME: The length of the file is known to the system. Assume that there is no room on the disk for the file to grow at the beginning but there is room to grow at the end.

7. [20 marks]

[6] (a) What is meant by protection in the context of operating systems? With the help of examples discuss the access matrix based protection technique used in multi-user computing systems.

[4] (b) Given that the base (relocation) register contains 1000 and the limit register contains 1400, determine the results of address translation (from logical to physical) in each of the following cases. Include the physical memory address generated.

(i) when the logical memory address is 705
(ii)when the logical memory address is 1500
[all addresses are expressed in decimal].

[5] (c) Briefly discuss the tradeoff between the minimization of average job turnaround time and starvation that is often observed in the context of CPU scheduling. Provide an example to demonstrate this tradeoff.

[5] (d) Answer any one of the following ((i) or (ii))

(i) What is a Redundant Array of Inexpensive Disks (RAID)? Discuss its impact on system reliability and performance.

OR

(ii) Discuss the advantages of using an acyclic graph directory over using a simple two-level directory in a multi-user system. Describe how file deletion is handled on such a system when an acyclic graph directory is used.