NATIONAL EXAMINATIONS December 2017

98-COMP-A5, 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 pocket 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).

5. Total Marks = 100.

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

Marking Scheme

1(a) 11 marks
1(b) 5 marks
1(c) 4 marks

2(a) 8 marks
2(b) 8 marks
2(c) 4 marks

3(a) 8 marks
3(b) 12 marks

4 (a) 9 marks
4(b) 9 marks
4(c) 2 marks

5(a) 14 marks
5(b) 3 marks
5(c) 3 marks

6(a) 10 marks
6(b) 6 marks
6(c) 4 marks

7(a) 4 marks
7(b) 4 marks
7(c) 4 marks
7(d) 4 marks
7(e) 4 marks

1 [20 marks].

(a) Consider the following arrivals on a system. Each process has a single CPU burst and does not perform any I/O.

Process	Arrival Time (seconds)	Execution Time(seconds)
Proc1	2	14
Proc2	4	7
Proc3	6	21
Proc4	8	2
Proc 5	10	1

(i) What is the minimum mean process turnaround time for any non-preemptive CPU scheduling strategy ?
(ii) What is the mean process turnaround time for the First Come First Served CPU scheduling strategy?

(b) Discuss a multiple queue based CPU scheduling technique that tries to approximate the behaviour of the Optimal CPU scheduling strategy

(c) Using examples, discuss how starvation can occur for some processes in a real time system.

2 [20 marks]

(a) Discuss whether increasing the degree of multiprogramming can improve the CPU utilization in the following scenarios. If your answer is NO for any one of these scenarios describe what change needs to be made for the improvement of CPU utilization.

(i) when the CPU utilization is 5% and the paging disk utilization is 1%
(ii) when the CPU utilization is 10% and the paging disk utilization is 98%
(iii) when the CPU utilization is 80% and paging disk utilizations are 10%
(iv) when both the CPU and paging disk utilizations are 50%

(b) Consider a demand paged virtual memory system in which a single program is currently running. The page map table is held in associative registers (associative memory). It takes 25 milliseconds to service a page fault if an empty frame is available or the replaced page is not modified, and 60 milliseconds if the replaced page is modified. Memory access time is 200 nanoseconds. Assume that for 70% of the page faults a page replacement is necessary and the page to be replaced is modified.

What is the maximum acceptable page fault rate such that the effective memory access time for the program is not greater than 300 nanoseconds?

(c) Briefly discuss the difference between internal and external fragmentation in the context of

page based memory management.

3 [20 marks].

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

171, 172, 173, 174, 275, 173, 174, 171, 276, 177, 178, 177, 178, 179, 177, 178, 179

Determine the number of page faults that would occur with the LRU page replacement algorithm when 3 frames are allocated to the program.

(b) Consider a multiprogrammed system that uses multiple partitions (of variable size) for memory management. A linked list of holes called the free list is maintained by the operating system to keep track of the available memory in the system. At a given point in time the free list consists of holes with sizes:
302K, 243K, 405K, 480K, 270K, 291K, 325K, and 350K

The free list is also ordered in the sequence given above: the first hole in the list is of size 302K words which is followed by a hole of size 243K words and so on. Jobs with different memory requirements arrive on the system in the following order:

	Arrival Time	Memory Requirement
Job 1	t1	322K
Job 2	t2	305K
Job 3	t3	403K
Job 4	t4	290K

[Given $t1 < t2 < t3 < t4$]

Explain how memory allocation would be performed in the given situation for (i) the best fit and (ii) the first fit policy. [For each policy determine which hole is allocated to each job after it arrives on the system].

4 [20 marks]

(a). Consider a multiprogrammed system consisting of 10 resources of the same type. No deadlock handling technique is employed by the system. That is, if a resource is requested by a process and one is available, a resource is allocated to the requesting process; otherwise the requesting process is blocked. P processes are run concurrently on the system. Each process can simultaneously hold up to K resources at any given point in time.

Once a resource is acquired by a process, it must be released by the process before it can be assigned to another process. Assume that each process requests and releases one resource at a time.

(i) For P = 15 and K = 3 can a deadlock occur on the system?
(ii) For K = 3, determine the maximum value of P such that a deadlock can never occur on the system

(b) Consider a moving head hard disk which consists of a single platter (surface) with 200 tracks on it. The tracks are numbered 0 to 199. The disk has just served a request on track 110 and is currently serving a request at track 153. The queue of pending requests in FIFO order is:

137, 165, 192.

(i) What is the total head movement (in number of tracks) needed to satisfy all these requests when a SCAN algorithm is used for disk scheduling?
(ii) In what order should the requests be served to minimize the total head movement?

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

(c) Briefly explain what is meant by execution time address binding in the context of memory management.

5 [20 marks]

(a) Different methods exist for storing information on the disk. Consider a file currently consisting of 120 blocks (numbered 1 - 120). Assume that the directory is available in main memory.
For each of the following cases (A-D) compute the minimum number of disk operations that are required when linked allocation (based on a singly linked list) is used.

(A) The contents of block 102, 98 and 103 are to be read.
(B) The contents of block 101 are to be exchanged with the contents of block 94.
(C) A new block is inserted after block 90. The content of the new block are the same as that of block 70.
(D) Block 65 is deleted.

Consider each case (A-D) 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.

(b) Briefly discuss the *mutual exclusion requirement* associated with the solution to the critical section problem.

(c) Is the *bounded waiting requirement* associated with the solution to the critical section problem satisfied when the entry to the critical section is controlled by a semaphore?

6 [20 marks]

(a) Briefly distinguish among a safe state, an unsafe state and a deadlock state. Using a system running five processes that use at least three different types of resources provide an example of

(i) an unsafe state and (ii) a deadlock state.

(b) Consider a demand paged virtual memory system. Can changing the system and allowing it to bring in more than one page after a page fault occurs improve system performance? Explain your answer with the help of examples.

(c) Explain what is meant by multi-threading and discuss why it is used. Discuss also the overhead associated with multithreading. [Be brief].

7 [20 marks]

(a) Briefly discuss File Allocation Table (FAT).

(b) Give examples of applications that perform (i) sequential access to files (ii) random access to files.

(c) Briefly discuss the advantages of using multiple disks for information storage.

(d) Briefly discuss the goals of *protection* in a multi-user computer system.

(e) Briefly discuss the roles of the valid/invalid bit and the dirty (modify) bit in memory management.