00039

The Open University of Sri Lanka
Faculty of Engineering Technology
Department of Electrical & Computer Engineering

Study Programme	: Bachelor of Technology Honours in Engineering
Name of the Examination	: Final Examination
Course Code and Title	**: EEX5335|EEX5535 Operating Systems**
Academic Year	: 2020/2021
Date	: 02nd of February 2022
Time	: 0930 – 1230hrs
Duration	: **3 hours**

General Instructions

1. Read all instructions carefully before answering the questions.
2. This question paper consists of **Eight (8)** questions in **Five (5)** pages.
3. Answer **any five (5)** questions given. All questions carry equal marks.
5. Answer for each question should commence from a new page.
6. This is a Closed Book Test **(CBT).**
7. Answers should be in clear handwriting.
8. Do not use red colour pen.

Question 1

i.) List five services provided by an operating system and briefly explain the benefits of these services to users. **[05 marks]**

ii.) Briefly explain the differences between *Monolithic* and *Microkernel* operating system architectures. State the advantages of these differences to provide a better service to users. **[06 marks]**

iii.) Describe the following terms using appropriate diagrams. What are they used for in an operating system? How do they work?

a. Inter Process Communication (IPC) **[03 marks]**
b. Message Passing **[03 marks]**
c. System Call **[03 marks]**

Question 2

Given five memory partitions 200kB, 500kB, 250kB, 350kB and 500kB which are fixed and in order.

i.) How would the first-fit, best-fit and worst-fit algorithms place processes of 212kB, 417kB, 112kB and 426kB in order? Illustrate the answers with suitable diagrams. State the assumptions you make. **[10 marks]**

ii.) Calculate internal or external fragmentations. Which algorithm makes the most efficient use of memory? Justify your answer in terms of internal and external fragmentation that could occur. **[10 marks]**

Question 3

i.) The following *Figure 3.1* shows a simplified layout of a process inside main memory. Briefly explain what is loaded in each section. **[04 marks]**

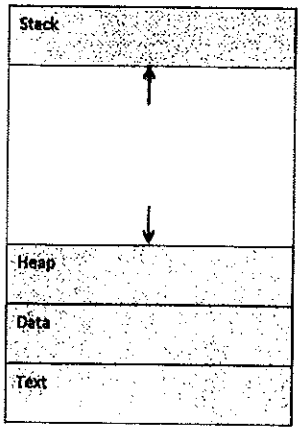

Figure 3.1 : A simplified layout of a process inside main memory

ii.) Briefly explain the states of a process referring the given "Process state diagram" as depicted in the Figure 3.2. **[06 marks]**

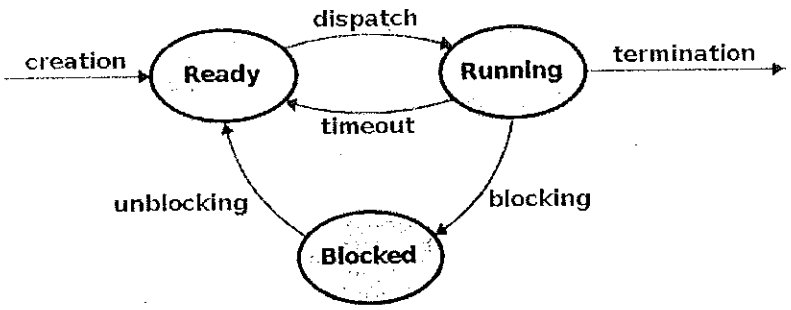

Figure 3.2 : Process state diagram

iii.) State the importance of having "Process Control Block (PCB)". **[04 marks]**

iv.) Differentiate a "Thread Control Block (TCB)" from a "Process Control Block (PCB)". **[06 marks]**

Question 4

i.) Provide two programming examples in which multithreading provides better performance than a single-threaded solution. **[04 marks]**

ii.) State two differences between user-level threads and kernel-level threads. State under what circumstances one is better than the other. **[04 marks]**

iii.) Consider a multiprocessor system and a multithreaded program written using the many-to-many threading model. Let the number of user-level threads in the program be greater than the number of processors in the system. Discuss the performance implications of the following scenarios.

a. The number of kernel threads allocated to the program is less than the number of processors. **[04 marks]**
b. The number of kernel threads allocated to the program is equal to the number of processors. **[04 marks]**
c. The number of kernel threads allocated to the program is greater than the number of processors but less than the number of user-level threads. **[04 marks]**

Question 5

The following sequence of events has occurred when allocating multiple resources to perform a function of a processor.

"**P**" indicates a process and "**R**" indicates a resource.

Event	Action
1	P1 requests and gets R1
2	P2 requests and gets R2
3	P2 requests R1
4	P1 releases R1 and P2 gets R1
5	P2 releases R2 and P3 gets R2
6	P3 requests and gets R3
7	P4 requests and gets R4
8	P3 requests R4
9	P4 requests R2
10	P4 requests R3

i.) Draw directed graphs to analyze the above scenario using Holt's modeling method. **[06 marks]**

ii.) Is this system, as a whole, deadlocked? Justify your answer. **[04 marks]**

iii.) Are there any deadlocked processes? If yes name them. If not indicate which events occurring next would cause a deadlock situation in the system. **[04 marks]**

iv.) What is the status of the system if P4's request for R2 is granted before P2's request for R2? Illustrate your answer by drawing a directed graph. **[06 marks]**

Question 6

i.) Explain the four conditions required to meet at a deadlock situation considering the given scenario in Question 5. **[08 marks]**

ii.) If your operating system supports a device allocation policy which says no event could be started unless all resources have been allocated to the process, what could be the result? Explain. **[06 marks]**

iii.) What is called "process starvation"? Explain using the given scenario in Question 5. How to overcome this problem? **[06 marks]**

Question 7

i.) Using two operating systems you are familiar with as examples, describe how the types of files in the file systems are similar or different. **[04 marks]**

ii.) Shared files and directories in a file system can be implemented in several ways. Name two common approaches adopted by operating systems for this and discuss merits and demerits of each approach. **[04 marks]**

iii.) Consider a file system that uses a modified contiguous allocation scheme with support for extents. A file is a collection of extents, with each extent corresponding to a contiguous set of blocks. A key issue in such systems is the degree of variability in the size of extents.

State the advantages and disadvantages of the following schemes.

a. All extents are the same size, and the size is predetermined. **[04 marks]**

b. Extents can be of any size and are allocated dynamically. **[04 marks]**

c. Extents can be of a few fixed sizes, and these sizes are predetermined. **[04 marks]**

Question 8

i.) State the differences between preemptive and non-preemptive scheduling algorithms and give examples for each. **[04 marks]**

ii.) What is the advantage of having different time-quantum sizes on different levels of a multilevel queuing system? **[02 marks]**

iii.) Consider a system running ten I/O bound tasks and one CPU bound task. Assume that the I/O bound tasks issue an I/O operation once for every two millisecond of CPU computing and that each I/O operation takes 20 milliseconds to complete. Also assume that the context switching overhead is 0.1 milliseconds and that all processes are long-running tasks. Calculate the CPU utilization for a round-robin scheduler when;

a) the time quantum is 1 millisecond **[04 marks]**

b) the time quantum is 10 milliseconds **[04 marks]**

iv.) Suppose that a scheduling algorithm favors the processes that have used the least processor time in the recent past. Why will this algorithm favor I/O bound programs and yet not permanently starve CPU-bound programs. State all the assumptions. **[06 marks]**