



NATIONAL OPEN UNIVERSITY OF NIGERIA
University Village, 91 Cadastral Zone, Nnamdi Azikwe Expressway, Jabi, Abuja
FACULTY OF SCIENCES
COMPUTER SCIENCE DEPARTMENT
2021_1 EXAMINATIONS

CIT 752: OPERATING SYSTEM CONCEPTS

Credit: 2 units

TIME ALLOWED: 2 Hours

INSTRUCTION: Answer Question 1 and any other FOUR (3) Questions

Question 1

- (a) State two major differences between user-level and kernel-level threads? (3 marks)
- (b) Under what circumstances is one type better than the other? (2 marks)
- (c) (i) What resources are used when a thread is created? (2 marks)
- (ii) How do the resources differ from those used when a process is created? (2 marks)
- (iii) Differentiate between parallelism and concurrency (2 marks)
- (iv) Is it possible to have concurrency but not parallelism? Explain. (2 marks)
- (d) What is the meaning of the term busy waiting? (3 marks)
- (e) What other kinds of waiting are there in an operating system? (3 marks)
- (f) (i) Can busy waiting be avoided altogether? (2 marks)
- (g) Is busy waiting always less or more efficient? (2 marks)
- (h) How do you overcome 'busy waiting'? (2 marks)

Question 2

- (a) A CPU-scheduling algorithm determines an order for the execution of its scheduled processes. Given n processes to be scheduled on one processor,
 - (i) How many different schedules are possible? (2 marks)
 - (ii) Give a formula in terms of n . to satisfy your answer in a (i) above (2 marks)
- (b) Explain the difference between preemptive and non-preemptive scheduling. (3 marks)

- (c) Suppose that the following processes arrive for execution at the times indicated in the table below and each process will run for the amount of time listed. In answering the questions, use non-preemptive scheduling, and base all decisions on the information you have at the time the decision is been made.

Process	Arrival Time	Burst Time
P1	0.0	8
P2	0.4	4
P3	1.0	1

- (i) What is the average turnaround time for these processes using the FCFS scheduling algorithm? (2 marks)
- (ii) What is the average turnaround time for these processes using the SJF scheduling algorithm? (2 marks)
- (iii) The SJF algorithm is supposed to improve performance, but notice that we chose to run process P1 at time 0 because we did not know that two shorter processes would arrive soon. Compute what the average turnaround time will be if the CPU is left idle for the first 1 unit and then SJF scheduling is used. Remember that processes P1 and P2 are waiting during this idle time, so their waiting time may increase. This algorithm could be called future-knowledge scheduling. (2 marks)
- (d) What advantage is there in having different time-quantum sizes at different levels of a multilevel queuing system? (2 marks)

Question 3

- (a) Explain why interrupts are not appropriate for implementing synchronization primitives in multiprocessor systems. (2 marks)
- (b) List three examples of deadlocks that are not related to a computer system environment. (3 marks)
- (c) Suppose that a system is in an unsafe state. Show that it is possible for the processes to complete their execution without entering a deadlocked state. (3 marks)
- (d) Explain three (3) methods for handling deadlocks. (3 marks)
- (e) Elaborate four (4) conditions that can cause a deadlock to occur. (4 marks)

Question 4

- (a) State two (2) differences between logical and physical addresses. (3 marks)
- (b) Consider a logical address space of 64 pages of 1,024 words each, mapped onto a physical memory of 32 frames.
 - (i) How many bits are there in the logical address? (3 marks)
 - (ii) How many bits are there in the physical address? (3 marks)
- (c) Why are page sizes always powers of 2? (3 marks)
- (d) Why are segmentation and paging sometimes combined into one scheme? (3 marks)

Question 5

- (a) Describe a mechanism by which one segment could belong to the address space of two different processes. (3 marks)
- (b) Explain the four (4) difference between internal and external fragmentation. (4 marks)
- (c) Explain why mobile operating systems such as iOS and Android do not support swapping. (3 marks)
- (d) Provide examples of applications that typically access files according to the following methods:
 - (i) Sequential (1 mark)
 - (ii) Random (1 mark)
- (d) If the operating system knew that a certain application was going to access file data in a sequential manner, how could it exploit this information to improve performance? (3 marks)