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

Study Programme	: Bachelor of Software Engineering Honours
	: Bachelor of Science Honours in Engineering
Name of the Examination	: Final Examination
Course Code and Title	**: EEX5563/EEX5564 Computer Architecture and Operating Systems**
Academic Year	: 2023/24
Date	: 05th February 2025
Time	: 1330 - 1630hrs
Duration	: **3 hours**

General Instructions

1. Read all the instructions carefully before answering the questions.
2. This question paper consists of **two (02)** sections, Section A and Section B on **seven (07)** pages.
3. Section A consists of twenty (20) multiple-choice questions (MCQs), while Section B comprises four (04) essay-type questions.
4. Answer all questions.
5. This is a Closed Book Test (CBT).
6. Answers should be in clear handwriting.
7. Do not use red colour pen.
8. Non-programmable calculators are allowed.
9. Use the sheet provided at the end of this exam paper to answer the MCQs (Section A) in Section A and attach it to the front of the answer sheets.

SECTION A

(This section consists of 20 MCQs. Provide the answers in the given sheet at the end)

1) In a multi-level feedback queue scheduling algorithm, a process that consumes less CPU time is likely to be:
   a) Moved to a lower priority queue
   b) Moved to a higher priority queue
   c) Terminated
   d) Assigned to a round-robin queue

2) Which of the following is true about monitors in process synchronization?
   a) A monitor is a data structure, but not a synchronization mechanism.
   b) A monitor allows multiple processes to be in the critical section at the same time.
   c) A monitor provides a higher-level abstraction than semaphores for managing process synchronization.
   d) A monitor must use busy waiting to manage access to shared resources.

3) In real-time operating systems, which of the following is NOT typically a scheduling criterion?
   a) Deadlines
   b) Priority
   c) Turnaround time
   d) Response time

4) In an I/O-bound process system, which scheduling algorithm typically performs best?
   a) First-Come, First-Served (FCFS)
   b) Shortest Job Next (SJN)
   c) Round Robin (RR)
   d) Longest Job First (LJF)

5) The thrashing phenomenon in an operating system occurs when:
   a) Processes are starved due to CPU contention.
   b) The system spends a majority of time swapping pages in and out of memory.
   c) The file system becomes fragmented due to excessive writes.
   d) Processes take too long to complete due to excessive context switching.

6) In a distributed file system, NFS (Network File System) is an example of:
   a) A file management system that uses virtual memory.
   b) A protocol for file sharing over a network.
   c) A user-level application that handles file access.
   d) A scheduling mechanism for accessing shared files.

7) In the Banker's algorithm for deadlock avoidance, the system is in a safe state if:
   a) Every process in the system can finish without causing a deadlock.
   b) The total number of resources is exhausted.
   c) Processes are allowed to execute in a cyclic manner.
   d) All processes are granted resources without any checks.

8) Livelock in a system occurs when:
   a) Processes remain in a non-terminating state, repeatedly executing the same instructions.
   b) The system fails to allocate memory resources, causing processes to terminate.
   c) Two or more processes are indefinitely swapping resources without making progress.
   d) The system is unable to allocate new resources to any process due to infinite waiting.

9) Which of the following statements is NOT true about the monolithic kernel?
   a) It provides high performance due to fewer context switches.
   b) It integrates most of the operating system services into one large kernel.
   c) It offers better security due to process isolation.
   d) It can be complex to maintain due to its large codebase.

10) Which of the following statements is true regarding thread synchronization in a multithreaded program?
   a) Mutexes and semaphores are used to allow multiple threads to enter the critical section simultaneously.
   b) A race condition can occur when threads do not properly synchronize access to shared resources.
   c) Deadlocks are impossible in multithreaded programs if threads are synchronized properly.
   d) A thread can acquire multiple mutexes simultaneously without risking deadlocks.

11) Which of the following statements about Instruction Set Architecture (ISA) is correct?
   a) ISA defines the internal structure and hardware organization of the CPU.
   b) ISA determines the execution speed of instructions by directly controlling clock cycles.
   c) ISA specifies the operations a processor can execute, including data types, registers, and addressing modes.
   d) ISA is responsible for optimizing cache hierarchies and memory bandwidth.

12) Which of the following best describes assembly language?
   a) A high-level programming language used for developing operating systems.
   b) A language that directly interacts with the hardware using symbolic instructions.
   c) A universal language that works across all computer architectures without modification.
   d) A language that eliminates the need for compilers or interpreters.

13) What is the primary advantage of interrupt-driven I/O compared to programmed I/O?
   a) Reduced hardware requirements by removing I/O controllers.
   b) Eliminates the need for the CPU to check device status repeatedly.
   c) Completely bypasses CPU involvement in all I/O operations.
   d) Transfers data faster by directly connecting devices to memory.

14) In a complex pipeline design, which of the following stages is primarily responsible for resolving data dependencies and forwarding data between instructions?
   a) Fetch
   b) Decode
   c) Execute
   d) Write-back

15) Which of the following best describes a drawback of microprogrammed realization compared to hardwired realization?
a) Increased complexity due to the use of physical logic circuits.
b) Slower execution speed due to the need for instruction decoding from memory.
c) Reduced flexibility in handling different instruction sets.
d) Higher power consumption due to the use of fixed combinational logic.

16) Which of the following is NOT a challenge associated with instruction-level parallelism (ILP)?
a) Data hazards
b) Control hazards
c) Instruction throughput
d) Resource limitations

17) Which of the following statements is false?
a) MIMD is ideal for tasks requiring high parallelism and multitasking.
b) SISD executes a single instruction stream on a single data set.
c) SIMD is used for applications requiring high throughput on data-parallel tasks.
d) MISD allows multiple processors to execute different instructions on a single data set.

18) Which of the following statements correctly describes a superscalar architecture?
a) Superscalar architectures execute multiple instructions simultaneously using multiple execution units.
b) Superscalar processors focus on reducing the number of instruction pipelines.
c) Superscalar architectures rely on a single pipeline to execute instructions sequentially.
d) Superscalar designs use very complex instructions to handle parallel execution.

19) Which of the following is a correct description of cache coherence?
a) Cache coherence ensures that all caches in a multiprocessor system have a consistent view of memory.
b) Cache coherence only applies to single-core processors.
c) Cache coherence prevents caches from communicating with each other.
d) Cache coherence is managed by the operating system and not hardware.

20) Which of the following statements is false?
a) RISC architectures use a reduced set of instructions to simplify hardware implementation.
b) CISC architectures aim to execute a wide range of instructions efficiently.
c) RISC designs tend to have longer instruction execution times due to complexity.
d) CISC processors often use microprogramming for instruction execution.

SECTION B (This section consists of four essay-type questions).

Question 1

a) Discuss the role of CPU scheduling in an operating system. Highlight why scheduling is necessary in a multiprogramming environment.

(5 marks)

b) A system has the following set of processes, each with its burst time and priority.

Process	Burst Time (ms)	Priority (Lower number = Higher priority)
P1	10	3
P2	1	1
P3	2	4
P4	1	5
P5	5	2

i. Calculate the average waiting time and turnaround time using Priority Scheduling. Repeat the calculation using the Round Robin Scheduling algorithm with a time quantum of 2 ms. Show all calculations step-by-step.

(10 marks)

ii. Evaluate the performance of the algorithms used in part (i) based on waiting time, turnaround time, and fairness. Which algorithm would you recommend for this scenario and why?

(5 marks)

Question 2

a) List and explain the four necessary conditions for deadlock to occur.

(4 marks)

b) Describe what a **page table** is and its role in virtual memory management.

(4 marks)

c) Given the following resource allocation graph, determine if a deadlock exists. State all your assumptions.

(4 marks)

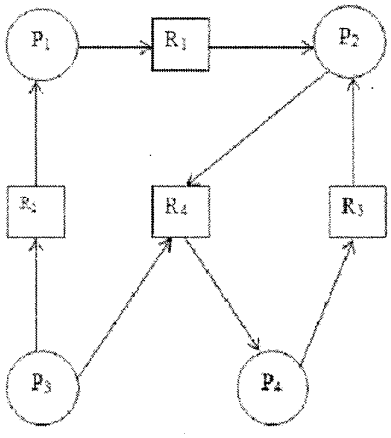

Figure 1: Resource allocation graph (Circles represent the processes and the squares represent non-sharable resources)

d) A system uses a **single-level page table**. The virtual memory size is 16GB, and the page size is 4KB.

a. How many entries are there in the page table? What is the size of the page table if each entry is 4 bytes?

(2 Marks)

b. Explain how **deadlock prevention** techniques can be applied to resource allocation in a system that uses page tables for memory management.

(2 Marks)

e) Discuss how the operating system ensures that deadlock does not occur in a system with **paging** and **virtual memory**.

(4 Marks)

Question 3

a) State Amdahl's Law and derive the equation for f, which represents the portion of the program that can be parallelized, assuming there is no scheduling overhead. Let T be the total time it takes to run the program on a single processor. Using this, calculate the speedup achieved when N parallel processors are used to fully utilize the parallelizable part of the program.

(4 marks)

b) Consider the execution of a program which results in the execution of 2 million instructions on a 2 GHz processor. The program consists of four major types of instructions. The instruction mix and the CPI for each instruction type are given below based on the results of a program analysis study. Assume a benchmark program executes 2 million instructions on a 2 GHz processor. The program includes four main types of instructions. The instruction mix and the CPI for each type are provided based on a results of a program analysis study.

* Keeping your answer to two decimal points is sufficient.

- Arithmetic and logic: CPI = 1, Instruction miss = 58%
- Load/store with cache hit: CPI = 2, Instruction miss = 16%
- Branch: CPI = 4, Instruction miss = 14%
- Memory reference with cache miss: CPI = 8, Instruction miss = 12%

i. Find the CPI and the MIPS of the above processor.

(4 marks)

Now assume that the program can be executed in 8 parallel threads (tasks) on 8 core system with roughly equal number of instructions executed in each task and each core having the same performance as above single processor. Coordination and synchronization between the parts adds an extra 20000 instruction executions to each task. Assume the instruction mix is same as above for each task but increase the CPI for memory reference with cache miss to 12 cycles due to contention for memory.

ii. Determine the new average CPI.

(2 marks)

iii. Determine the corresponding MIPS rate.

(2 marks)

iv. Calculate the speedup factor.

(4 marks)

v. If it is assumed that the parallelizable fraction of code is f = 1, then what is the actual speedup factor with the theoretical speedup factor (percentage) determined by Amdahl's law.

(4 marks)

Question 4

a) Briefly describe the main characteristics of memory hierarchy in a computer system, explaining why having a memory hierarchy is important.

(4 marks)

b) How does mapping work in Direct-mapped, Associative Mapping, and Set-associative Mapping cache techniques?

(3 marks)

c) Consider a magnetic disk drive with 4 surfaces, 250,000 tracks per surface, and 500 sectors per track. The sector size is 512 bytes. The average seek time is 8.5 ms, the track-to-track access time is 4.7 ms, and the drive rotates at 7200 rpm. Successive tracks in a cylinder can be read without head movement.

* Keeping your answer to two decimal points is sufficient.

i. What is the disk capacity?

(2 marks)

ii. What is the data transfer rate?

(3 marks)

iii. What is the average access time? Assume this file is stored in successive sectors and tracks of successive cylinders, starting at sector 0, track 0 of cylinder (i).

(4 marks)

iv. Estimate the time required to transfer a 5-MB file.

(4 marks)

Use this sheet to provide answers to the MCQ questions of SECTION A and attach to the answer script.

Index Number: .....

Question Number	Answer			
1	a	b	c	d
2	a	b	c	d
3	a	b	c	d
4	a	b	c	d
5	a	b	c	d
6	a	b	c	d
7	a	b	c	d
8	a	b	c	d
9	a	b	c	d
10	a	b	c	d
11	a	b	c	d
12	a	b	c	d
13	a	b	c	d
14	a	b	c	d
15	a	b	c	d
16	a	b	c	d
17	a	b	c	d
18	a	b	c	d
19	a	b	c	d
20	a	b	c	d