

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



Study Programme : Bachelor of Science Honours in Engineering  
Name of the Examination : Final Examination  
**Course Code and Title : EEX6236 Advanced Computer Architecture**  
Academic Year : 2024/25  
Date : 24<sup>th</sup> January 2026  
Time : 1330-1630h  
Duration : 3 hours

### General instructions

1. Read all instructions carefully before answering the questions.
2. This question paper consists of **Four (4)** questions in **Three (3)** pages.
3. Answer **all** questions.
4. Answer for each question should commence on a new page.
5. This is a Closed Book Test (CBT).

### Question 01

This question assesses understanding of multi/many-core architectural drivers and memory organization principles

- a. The drive towards multi-core and many-core architectures was largely necessitated by fundamental limitations encountered in single-core processor design. Explain the concepts of the “Power Wall” and the “Frequency Wall” and detail how these limitations hindered further performance scaling of single-core CPUs. (10 marks)
- b. Multi-core systems can be designed using either Homogeneous or Heterogeneous architectures. Compare and contrast these two architectural approaches based on three aspects: the type of cores used, the resulting complexity of programming, and their typical use cases (e.g., in server versus mobile environments). (10 marks)
- c. Non-Uniform Memory Access (NUMA) is an architecture used to improve the scalability of shared-memory systems.
  - I. Describe the challenge introduced by NUMA regarding access latency when a processor requests data from memory. (5 Marks)

- II. Explain how a mechanism like *Cache Coherent NUMA (CC-NUMA)* is necessary to maintain a consistent view of shared data across the multiple distributed memory controllers and caches in such an architecture. (5 Marks)

## Question 02

This question focuses on comparative programming models and synchronization essential for parallel execution.

- a. Granularity refers to the size of computation in a parallel task. (5 Marks)
- I. Explain the difference between *fine-grained* and *coarse-grained* parallelism.
- II. Justify why “coarse-grained parallelism” is generally preferred or more suitable for execution in a “loosely coupled Distributed System” (such as a cluster of machines) compared to a “tightly coupled multi-core system”. (5 Marks)
- b. Parallel programming relies on different communication models depending on the underlying hardware architecture. Compare the “Shared Memory Programming Model” with the “Message Passing Programming Model” based on the following criteria:
- I. The underlying memory model (e.g., global address space vs. private memory). (4 Marks)
- II. How communication and coordination are fundamentally achieved between parallel tasks. (3 Marks)
- III. The relative scalability of each model to very large numbers of processors/nodes. (3 Marks)
- c. In a multi-core shared memory system, mechanisms are required to ensure the correctness of parallel execution. Describe the primary function and a typical use case for two distinct synchronization primitives (excluding atomic operations) from the list below, and explain why they are critical for avoiding race conditions: (10 marks)
1. Locks (Mutexes)
  2. Barriers
  3. Semaphores

### Question 03

This question assesses knowledge of memory reliability mechanisms.

- a. Error detection is crucial for maintaining data integrity in memory. (10 marks)  
 Explain the primary difference in detection capability and the relative hardware complexity trade-off between "Parity Bits" and "Cyclic Redundancy Check (CRC)", and give a typical use case for each in a computer system.
- b. For high-reliability server memory systems, advanced error correction codes are deployed.
  - I. Describe the fault tolerance capability of "Single Error Correction, Double Error Detection (SECCDED)" codes. (5 Marks)
  - II. Explain how "Chipkill" technology enhances reliability beyond standard SECCDED codes, specifically addressing scenarios that SECCDED cannot handle effectively. (5 Marks)

### Question 04

This question assesses knowledge of specialized instruction sets and complex arithmetic hardware.

- a. "Short Vector Instruction Sets (SVIS)", such as SSE or AltiVec, are common extensions in modern processors.
  - I. Describe the processing model SVIS uses (Scalar or Vector) and identify one key application domain where SVIS provides significant performance enhancement. (5 Marks)
  - II. Explain how SVIS achieves "higher performance" and "reduced energy consumption" for data-intensive workloads compared to equivalent scalar instruction loops. (5 Marks)
- b. Square root computation requires specialized hardware algorithms. Compare the "Newton-Raphson method" and the "Digit Recurrence method" based on the following criteria:
  - I. The speed of convergence (e.g., linear vs. quadratic). (3 Marks)
  - II. The core arithmetic operations required (e.g., division, simple logic). (3 Marks)
  - III. Which method is generally better suited for implementation in a small, low-area embedded system? (4 Marks)

*\*All Rights Reserved\**

[illegible][illegible]