UNIVERSITY OF KWAZULU-NATAL

School of Engineering

Discipline of Electrical, Electronic and Computer Engineering

MAIN EXAMINATIONS: SEPTEMBER 2025

MODULE and CODE: Embedded Systems, ENEL4ESH2

Duration: 2 hours (120 minutes) **Total Marks: 100**

Internal Examiner: **Mr. R. Khuboni**
External Examiner: **Prof. K. Ouahada**

General Instructions:

1. Answer all questions in ink except for MCQs and diagrams, in pencil.
2. You may use a scientific calculator but clear all program memory.
3. Try to keep your sketches as clear and as readable as possible.

There are three sections:

Section A (25%)	Answer on MCQ sheet – indicate ***student number*** and ***seat number***
Section B (35%)	Answer in an **answer booklet provided**
Section C (40%)	Answer in an **answer booklet provided**

SECTION A – GENERAL MULTIPLE CHOICE

There are 25 MCQ questions based on Embedded Systems. Please answer all these questions. Each question is worth 1 mark and there is no negative marking. Please follow the instructions on the MCQ answer sheet.

[25 Marks]

A1. Which of the following best defines an embedded system?
   a. A programmable computer used for multiple applications
   b. A dedicated system integrated within a larger device
   c. A personal computer with custom peripherals
   d. A system that only uses general-purpose processors
   e. None of these

A2. Which of these is a key constraint in embedded systems such as a bike computer?
   a. High throughput
   b. Cost and power efficiency
   c. Cloud scalability
   d. Large memory usage
   e. None of these

A3. Cyclomatic complexity in software design measures:
   a. Lines of code
   b. Memory footprint
   c. Number of independent paths in control flow
   d. Average execution time
   e. None of these

A4. Programmable Logic Devices (PLDs) such as FPGAs are commonly used because:

a. They have very low unit cost for mass production
b. They allow reprogramming for prototyping and design flexibility
c. They eliminate the need for microcontrollers
d. They always consume less power than ASICs
e. None of these

A5. What does the Behavioural Domain in the Y-chart represent?

a. The high-level functionalities and operations required by the embedded system, including tasks, algorithms, and control processes.
b. The different possible implementations of the system's functionalities, such as hardware and software options.
c. The allocation decisions of functionalities to hardware or software components.
d. The physical components and their layout in the embedded system.
e. None of these

A6. What is a key feature of bank-switched memory compared to dual-port memory?

a. Bank-switched memory allows simultaneous access to different memory banks by the CPU and I/O hardware.
b. Bank-switched memory provides a single memory space that is accessed by both software and hardware.
c. Bank-switched memory has two modes of operation for exclusive access to memory banks.
d. Bank-switched memory uses a FIFO buffer for data transfer.
e. None of these

A7. What happens during a recessive to dominant transition in a CAN bus system?

a. A bit is automatically inserted by the receiver to maintain synchronization.
b. All nodes on the bus synchronize to the Start of Frame (SOF) bit.
c. The bus becomes idle, and all nodes stop transmitting data.
d. The bit-stuffing process is initiated by the receiver.
e. None of these

A8. What is the primary function of the DMARequest signal in a DMA operation?

a. To indicate that the DMA Controller has completed the data transfer.
b. To signal the peripheral device to place data onto the data bus.
c. To acknowledge that the DMA Controller has gained control of the buses.
d. To request control of the system buses from the bus arbitration unit.
e. None of these

A9. Consider the case where you have two 2K x 8 RAM chips and you are expected to create a 4K x 8 memory modules. Which of the following lines would remain the same to achieve this memory expansion?

a. Address and Data Lines
b. Control Lines
c. Address Lines
d. Control and Data Lines
e. Address and Control Lines

A10. How can cache capacity misses be mitigated?
   a. By increasing cache associativity
   b. By implementing hardware prefetching
   c. By using larger cache blocks
   d. By optimizing compiler code
   e. All of these

A11. In RTOS scheduling, the primary goal is to?
   a. Maximize average CPU utilization
   b. Meet task deadlines deterministically
   c. Minimize the number of interrupts
   d. Increase GUI responsiveness
   e. None of these

A12. A cache miss penalty includes.
   a. Only the time to transfer a block
   b. Only the time to access lower level
   c. Replacement time, access time, and transfer time
   d. Only the time to invalidate upper levels
   e. None of these

A13. Why does I2C use ACK/NACK after every byte?
   a. To confirm data integrity
   b. To allow master to terminate communication early
   c. To let slaves indicate readiness
   d. To support multi-slave addressing
   e. All of these

A14. Which miss type cannot be reduced by increasing cache size?
   a. Conflict miss
   b. Compulsory miss
   c. Capacity miss
   d. Coherence miss
   e. None of these

A15. Which of the following reduces conflict misses without enlarging the cache size?
   a. Larger block size
   b. Higher associativity
   c. Non-blocking cache
   d. Victim cache
   e. All of these

A16. A master writes data to a slave but receives a NACK. Which is not a likely reason?
   a. Incorrect slave address
   b. Slave is busy and not ready
   c. Clock stretching by the slave
   d. SDA line error
   e. None of these

A17. Why does LIN adopt a time-triggered schedule controlled by the master, instead of event-driven arbitration?

a. To simplify slave hardware and reduce cost
b. To increase bandwidth above 1 Mbps
c. To allow multiple masters on the same bus
d. To reduce frame overhead size
e. None of these

A18. Why does DMA improve throughput in sensor-heavy embedded applications?

a. Fewer pins used
b. Hardware priority arbitration
c. Frees CPU from memory transfers
d. Enables polling-free buffers
e. None of these

A19. What determines which node wins arbitration when multiple CAN nodes start transmitting simultaneously?

a. The Control field
b. The Data field
c. The Identifier field
d. The CRC field
e. None of these

A20. Each I^2C byte transfer is:

a. 7 data bits + 1 parity
b. 8 data bits + 1 ACK/NACK bit
c. 9 data bits + 1 stop bit
d. 8 data bits only
e. None of these

A21. Why does LIN not require quartz crystals on slave nodes?

a. Event-triggered scheduling
b. Master provides synchronization byte
c. Slaves always free-run at fixed baud
d. LIN uses differential signaling
e. None of these

A22. In an RTOS, a task handle is primarily used to:

a. Identify the task for later control operations
b. Allocate heap memory
c. Set interrupt priorities
d. Configure timer frequency
e. None of these

A23. What does *vTaskStartScheduler()* do in FreeRTOS?

a. Creates all tasks
b. Starts the tick timer and begins task scheduling
c. Deletes all tasks
d. Suspends the highest-priority task
e. None of these

A24. Which of the following helps reduce software complexity?

a. Modularity
b. Structured programming
c. Clear specifications
d. All of these
e. None of these

A25. How does the synchronous clock requirement help in case two nodes attempt to start transmission simultaneously on a CAN bus?

a. It ensures that one node always wins the bit arbitration.
b. It prevents any further communication on the bus.
c. It causes collision and data loss.
d. It leads to a bus error, requiring manual intervention.
e. None of these

SECTION B – GENERAL SHORT QUESTIONS

There are 3 questions based on Embedded Systems. Please answer all questions to your best ability.

[35 Marks]

B1. A processor system is divided into four components that account for 50%, 15%, 15%, and 20% of total execution time, respectively. A new design improves the performance of these components by 10%, 50%, 10%, and 20%. Using Amdahl's Law, calculate the overall system speedup achieved by this design.

[5]

B2. Assume that you are working on a project that uses an ARM72 processor with a two-level cache hierarchy (L1 cache = 64 KB, L2 cache = 2 MB), and the main memory of 4 GB is connected to the second-level cache. The processor generates an average of 1.35 memory references per instruction. The L1 cache has a hit time of 1.8 clock cycles, while the L2 cache hit time is unknown. The L2 miss penalty (accessing main memory) is 85 clock cycles. During program execution, for every 1000 memory references, there are 65 L1 cache misses and 30 L2 cache misses. The system must be designed to achieve an Average Memory Access Time (AMAT) of exactly 5 clock cycles.

a. Calculate the miss rates for L1 and L2 caches, including the local miss rate of L2 and the global miss rate with respect to total memory references.
b. Given that the AMAT target is 5 clock cycles, compute the required L2 cache hit time.
c. Determine the average number of stall cycles per instruction due to memory access, assuming 1.35 memory references per instruction.

[10]

B3. A single-purpose custom processor is required to generate the **Fibonacci output value** based on the iterative code loop shown below. The process begins when *go_i* is asserted. The seed values are presented on inputs *a_i* and *b_i,* and the iteration count is provided on input *n_i*. The result of the computation is produced on output *d_o*. [6 marks]

```
// Computes: t = nth Fibonacci number using iterative updates with seeds a and b
// t = F(n), where F(0) = a, F(1) = b, and F(k) = F(k-1) + F(k-2) for k ≥ 2
int a, b, t, i, n;
while (1) {
   while (!go_i) ;
   a = a_i;      // seed F(0)
   b = b_i;      // seed F(1)
   n = n_i;      // iteration count
   i = 0;
   while (i < n) {
      t = a + b;   // compute next term
      a = b;       // shift forward
      b = t;       // update latest term
      i = i + 1;
   }
   d_o = a;   // after n iterations, a = F(n)
}
```

Figure 1: FSM Code

a) From the Figure 1 above, sketch the logical FSMD diagram for a single purpose processor.

[6]

b) Split the FSMD into two parts, the Controller and the Data-path. Show all the units in the data-path and all the control signals between the two parts. *(Do not simplify or optimize your processor). What can be optimized within this design?*

[14]

SECTION C – GENERAL QUESTIONS

There are 3 questions based on Embedded Systems. Please answer questions to your best ability.

[40 Marks]

C1. Figure 2 shows a control system comprising of a master controller and a few slave controllers. The control system is connected through a CAN network. Slave4 has highest priority and slave1 lowest priority. With respect to this type of network, answer the following questions:

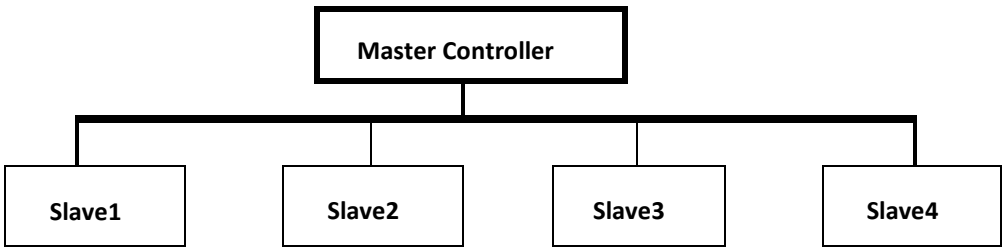


Figure 2: CAN Bus system

a. Illustrating with appropriate diagrams, explain how you can ensure that in the event of simultaneous transmission of data by two or more slaves, only the slave with the highest priority will be transmitted first.

[4]

b. Illustrating with approximate performance measure, give a brief comparative account of the following Media Access Control (MAC) protocols, ***Aloha MAC Protocol***, ***Ethernet MAC protocol*** and ***CAN Bus MAC protocol***.

[9]

C2. You are the project engineer at EmbeddedTech Solutions (Pty) Ltd, tasked with designing a simple energy monitoring module using the AVR128DA microcontroller. The system components are connected via the I²C bus, where the microcontroller operates as the master device. The connected peripherals are: DS3231 Real-Time Clock (RTC) at address 0x68, 24LC512 EEPROM (256 KB) at 0x50, ZMPT101B voltage sensor at 0x40, and ACS712 current sensor at 0x48. The system is required to log voltage and current measurements every 5 minutes, using the RTC for timestamping, and store the results in the EEPROM at a bus speed of 400 Kbps. This configuration provides continuous monitoring and storage of energy usage data for smart home automation. With reference to this design, answer the following questions:

a. Explain how the synchronization of the serial clock line in the I2C bus system works in terms of the low and high periods?

[5]

b. Consider this setup with multiple slave modules (i.e., RTC, EEPROM, voltage, and current sensors) trying to transmit data to the master device. Explain how does the I2C bus arbitration handles data transmission from these modules? Illustrate with a timing diagram.

[7]

C3. Design a minimal cost positive logic select signals for the following. Ensure to minimize the Karnaugh maps, determine their equations, and chip select signals.

RAM1 $0000 to $2FFF
RAM2 $3000 to $7FFF
Input $A000
Output $B000
ROM $D000 to $FFFF

a. What happens if the software accesses the location $6000?
b. Design a fully decoded negative logic select signal for a 4KB RAM at RAM2 memory allocation from the above memory map. Is this a regular or irregular address range?

[15]

$$2^{\text{index}} = \frac{\text{Cache size}}{\text{Block size} \times \text{Set associativity}}$$

$$\text{CPU execution time} = (\text{CPU clock cycles} + \text{Memory stall cycles}) \times \text{Clock cycle time}$$

$$\text{Memory stall cycles} = \text{Number of misses} \times \text{Miss penalty}$$

$$\text{Memory stall cycles} = \text{IC} \times \frac{\text{Misses}}{\text{Instruction}} \times \text{Miss penalty}$$

$$\frac{\text{Misses}}{\text{Instruction}} = \text{Miss rate} \times \frac{\text{Memory accesses}}{\text{Instruction}}$$

$$\text{Average memory access time} = \text{Hit time} + \text{Miss rate} \times \text{Miss penalty}$$

$$\text{CPU execution time} = \text{IC} \times \left(\text{CPI}_{\text{execution}} + \frac{\text{Memory stall clock cycles}}{\text{Instruction}}\right) \times \text{Clock cycle time}$$

$$\text{CPU execution time} = \text{IC} \times \left(\text{CPI}_{\text{execution}} + \frac{\text{Misses}}{\text{Instruction}} \times \text{Miss penalty}\right) \times \text{Clock cycle time}$$

$$\text{CPU execution time} = \text{IC} \times \left(\text{CPI}_{\text{execution}} + \text{Miss rate} \times \frac{\text{Memory accesses}}{\text{Instruction}} \times \text{Miss penalty}\right) \times \text{Clock cycle time}$$

$$\frac{\text{Memory stall cycles}}{\text{Instruction}} = \frac{\text{Misses}}{\text{Instruction}} \times (\text{Total miss latency} - \text{Overlapped miss latency})$$

$$\text{Average memory access time} = \text{Hit time}_{L1} + \text{Miss rate}_{L1} \times (\text{Hit time}_{L2} + \text{Miss rate}_{L2} \times \text{Miss penalty}_{L2})$$

$$\frac{\text{Memory stall cycles}}{\text{Instruction}} = \frac{\text{Misses}_{L1}}{\text{Instruction}} \times \text{Hit time}_{L2} + \frac{\text{Misses}_{L2}}{\text{Instruction}} \times \text{Miss penalty}_{L2}$$

Figure 5.9 Summary of performance equations in this chapter. The first equation calculates the cache index size, and the rest help evaluate performance. The final two equations deal with multilevel caches, which are explained early in the next section. They are included here to help make the figure a useful reference.