

**CST1**  
**COMPUTER SCIENCE TRIPOS Part IB**

---

Tuesday 9 June 2026 14:00 to 17:00

---

COMPUTER SCIENCE Paper 5

Answer **five** questions.

Submit the answers in five **separate** bundles, each with its own cover sheet. On each cover sheet, write the numbers of **all** attempted questions, and circle the number of the question attached.

**You may not start to read the questions  
printed on the subsequent pages of this  
question paper until instructed that you  
may do so by the Invigilator**

STATIONERY REQUIREMENTS

*Script paper*

*Blue cover sheets*

*Tags*

SPECIAL REQUIREMENTS

*Approved calculator permitted*

## 1 Computer Networking

- (a) An Ethernet packet containing TCP/IP session data typically holds pairs of src/dest addresses or identifiers for various protocol layers. Select three different pairs of addresses, where each pair belongs to a different layer in the protocol stack. Explain the function of each pair and their associated assignment and matching procedures. [6 marks]
- (b) Integrity checks include cyclic redundancy checks (CRCs) and checksums. How do these differ in complexity and performance? [2 marks]
- (c) A packet switch or router can maintain packet queues at its input, at its output, distributed through its switching fabric, or using a combination of these options. Out of these options, is one theoretically ideal? Your answer should consider router cost and FIFO and non-FIFO queueing. [6 marks]
- (d) All-optical packet switches have been designed without any packet queues. They operate by sending a packet to a random output port when there is contention for the desired port. How adaptable is TCP to such systems in terms of detecting and recovering from loss or to out-of-order reception? [6 marks]

## 2 Computer Networking

- (a) Describe the functions of *broadcast*, *multicast* and *anycast* in the context of both local-area networks (LANs) and the Internet, and compare them against each other. [4 marks]
- (b) A packet network being *reliable* implies it should always be available. List three further, typical reliable delivery guarantees or other aspects of reliability. List two guarantees provided by a *secure* network connection (e.g. as provided by TLS). [5 marks]
- (c) Techniques for providing network reliability can be broadly classified as hop-by-hop or end-to-end.
- (i) Explain what these terms mean and provide a real-world example of both. [4 marks]
- (ii) Which is possibly more important and why can it be advantageous for a long-distance connection to use both? [2 marks]
- (d) A team is implementing software to run in a set of post office checkout tills. The leader suggests it is straightforward to ensure consistent balances using reliable multicast and suggests that L2 (LAN data-link layer) broadcast will serve just as well.
- (i) Why might this look like a tempting suggestion? [1 mark]
- (ii) Can multicast ever be reliable? [2 marks]
- (iii) If not, what alternatives should be considered in this situation? [2 marks]

### 3 Computer Networking

- (a) A Fellow in a Cambridge college thinks they cannot connect to a database hosted in the Finance Office and suspects this is a matter of VLAN configuration.
- (i) What is a VLAN? [2 marks]
  - (ii) Why might a College network use VLANs? [2 marks]
  - (iii) Assuming it is a VLAN issue, should the network be reconfigured for this Fellow or is an alternative solution preferable? [2 marks]
- (b) Describe three feasible methods by which traffic per server can be balanced when multiple servers host a popular web domain such as `www.bbc.co.uk`. Each should operate in a different part of the overall system. [6 marks]
- (c) Which of the methods given in answer to Part (b) also balance load for individual objects on a web page? [2 marks]
- (d) Tunnelling is a repetition of one or more protocol layers on top of where they would normally sit, and is commonly used for home access to a work file server. Why is this used, and what are alternatives? Answer giving advantages or disadvantages. [6 marks]

#### 4 Concurrent and Distributed Systems

- (a) Describe briefly what is meant by *consensus* in distributed systems, and how it relates to *total order broadcast*. [3 marks]
- (b) Define the five properties that a total order broadcast algorithm must specify. For each property, state whether it is a safety or a liveness property. [5 marks]
- (c) Your colleague proposes implementing a total order broadcast system as follows. A server keeps an in-memory array of references to messages, initialised with nulls. When a client wants to broadcast a message, the client sends this message over the network to the server. On receipt, a thread on the server performs an atomic compare-and-swap operation that updates the first null array entry to point at that message, provided that another thread has not concurrently updated the same array entry. All clients then deliver messages from non-null array entries in the order they appear in the array.

Is this a correct total order broadcast algorithm according to your definition from Part (b)? Justify your answer. [4 marks]

- (d) How does the algorithm from Part (c) compare to Raft? [3 marks]
- (e) Imagine you have a fault-tolerant counting service, which can be queried by clients. Every query atomically increments a counter and returns the new counter value. Can this service be used to implement a total order broadcast system according to your definition from Part (b)? If yes, explain how this is done. If no, explain why it cannot be done. [3 marks]
- (f) In the scenario of Part (e), does it make a difference if you assume that there are no more than two clients? Justify your answer. [2 marks]

## 5 Concurrent and Distributed Systems

In a bank's database, let a *payment* from one account (the *source*) to another (the *destination*) be a transaction that first reads the balance of both accounts, aborting if the payment amount is greater than the source account balance. If not aborted, the transaction updates the source account balance to be the balance it previously read minus the payment amount, updates the destination account balance to be the previously read balance plus the payment amount, and then commits.

- (a) What is the purpose of strict two-phase locking (2PL) in a database, and how does it work? [3 marks]
- (b) Assume the database uses strict 2PL with a multi-reader single-writer lock for each account, which is acquired when the account is first accessed in a transaction. Which problem could occur if there is a payment from account *A* to account *B*, and concurrently another payment from account *B* to account *A*? Draw a state machine diagram showing all the possible executions of these two payments, assuming that each account balance read or update is an atomic operation. [6 marks]
- (c) Your colleague suggests changing the payment transaction as follows. Instead of accessing first the source and then the destination account, a payment transaction should first read the balance of the account with the lower account number, then read the balance of the account with the higher account number, then update the lower-numbered account, and finally update the higher-numbered account. Does this prevent the problem from Part (b)? Justify your answer. [4 marks]
- (d) As a consistency check, the bank wants to periodically read the balance of all accounts and confirm that their sum equals the bank's total holdings. Explain what problems might occur in a database that uses strict 2PL. [2 marks]
- (e) The bank is considering moving to a database that uses optimistic concurrency control instead of 2PL. Explain what impact this is likely to have on the scenarios from Parts (b), (c), and (d). What new problems might this introduce that do not occur with 2PL? [5 marks]

## 6 Introduction to Computer Architecture

- (a) For a microprocessor pipeline, what are control hazards and how can they be mitigated? [4 marks]
- (b) Consider the 32-bit RISC-V assembler code below for the three functions `foo1`, `foo2`, and `foo3`. The RISC-V reference card has been provided if you need to look up an instruction.

```

foo1:
    mv    t0, zero
    beq   a0, zero, foo1exit
foo1loop:
    andi  t1, a0, 1
    add   t0, t0, t1
    srli  a0, a0, 1
    bne   a0, zero, foo1loop
foo1exit:
    mv    a0, t0
    ret

foo2:
    mv    s3, ra
    mv    s0, a0
    andi  a0, s0, 0xff
    call  foo1
    mv    s1, a0

    srli  s0, s0, 8
    andi  a0, s0, 0xff
    call  foo1
    add   s1, s1, a0

    srli  s0, s0, 8
    andi  a0, s0, 0xff
    call  foo1
    add   s1, s1, a0

    srli  s0, s0, 8
    andi  a0, s0, 0xff
    call  foo1
    add   a0, s1, a0
    mv    ra, s3
    ret

foo3:
    li    t0, 0x55555555
    and   t1, a0, t0
    srli  t2, a0, 1
    and   t2, t2, t0
    add   a0, t1, t2
    li    t0, 0x33333333
    and   t1, a0, t0
    srli  t2, a0, 2
    and   t2, t2, t0
    add   a0, t1, t2
    li    t0, 0x0f0f0f0f
    and   t1, a0, t0
    srli  t2, a0, 4
    and   t2, t2, t0
    add   a0, t1, t2
    li    t0, 0x00ff00ff
    and   t1, a0, t0
    srli  t2, a0, 8
    and   t2, t2, t0
    add   a0, t1, t2
    andi  t1, a0, 0xff
    srli  t2, a0, 16
    add   a0, t2, t1
    ret

```

- (i) If each function is called with the argument in register `a0=0xff000000`, what answer does it return? Explain your answers. [6 marks]
- (ii) If this code was executed on a simple five-stage scalar pipeline (instruction-fetch, decode/register-fetch, execute, memory-access, write-back) with no branch prediction and branches taken in the execute state, how many clock cycles will each function take when called with `a0=0x0f`? State your assumptions. [6 marks]
- (c) A new custom instruction is proposed to speed up computation of `foo1`. Write an efficient combinational function for `foo1` in SystemVerilog. [4 marks]

## 7 Introduction to Computer Architecture

- (a) In Flynn's taxonomy of computer architectures four design spaces are identified and given the following labels: SISD, SIMD, MISD, MIMD. What do these abbreviations stand for? Give an example use case for each space in this taxonomy. [6 marks]
- (b) More recently SIMT systems have appeared. What does SIMT mean? What commercial processors implement the SIMT model? [2 marks]
- (c) Modern systems-on-chip typically contain processor cores of different size and microarchitecture. What are the advantages of this approach? [4 marks]
- (d) How does the handling of data-dependent control-flow differ between RISC-V based processors and GPUs? [4 marks]
- (e) What techniques are used to mitigate memory access latency overheads by RISC-V cores and GPUs? [4 marks]

## 8 Introduction to Computer Architecture

- (a) What is Moore's Law and does it still apply in 2026? [2 marks]
- (b) Dennard scaling stopped around 20 years ago. What is Dennard scaling and what have been its effects on clock frequency scaling? [2 marks]
- (c) How does pipelining allow clock frequency to be increased? [2 marks]
- (d) You are asked to evaluate two scalar processor designs, Fastpath and Ecopath, that both use the RISC-V ISA RV32G. The pipelines are depicted below where: IF=Instruction Fetch, DC+RF=Decode and Register Fetch, EX=Execute (including branches), MA=Memory Access, WB=Write-back. Fastpath runs at 2GHz enabled by its pipelined instruction and data caches that take two cycles to return information but can accept a new address every clock cycle. Ecopath has single cycle caches.

*Ecopath Pipeline*

IF	DC+RF	EX+MA+WB
----	-------	----------

*Fastpath Pipeline*

IF1	IF2	DC+RF	EX	MA1	MA2	WB
-----	-----	-------	----	-----	-----	----

- (i) What are feed-forward paths (sometimes called *bypass paths*) and where are they needed for the two pipelines? Provide a brief rationale for each forwarding path. [4 marks]
- (ii) Given the following memory copy code, how many pipeline stalls will each pipeline experience when executing one iteration of the loop? Assume that all memory accesses hit in the caches, that there is no branch prediction, and `a2` is initialised to 1.

```

memcpy: lw    t0, 0(a0)
        sw    t0, 0(a1)
        addi a0, a0, 4
        addi a1, a1, 4
        addi a2, a2, -1
        bge  a2, zero, memcpy
        ret

```

[4 marks]

- (iii) When executing the `memcpy` code with `a2=2`, will Fastpath or Ecopath be fastest? Justify your answer. [4 marks]
- (iv) What code optimisations could be applied to reduce the execution time? [2 marks]

**END OF PAPER**