

COMPUTER SCIENCE TRIPOS Part IB

Tuesday 2 June 2020 1.30 to 4.30

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 Design

- (a) What is Moore's law for CMOS technology scaling? [3 marks]
- (b) Why does the performance of wires not scale with transistor technology shrinks? [3 marks]
- (c) What determines the maximum clock frequency that a digital circuit can be safely clocked at? [3 marks]
- (d) In SystemVerilog, when is a variable of type `logic` implemented as a `wire` rather than a `reg`? [3 marks]
- (e) Consider the following three correct implementations of the function `fib` that compute the Fibonnaci sequence.

```

module fibA(input          clk,
            input          rst,
            input logic [5:0] n,
            input logic     start,
            output logic [20:0] r);

function automatic longint fibr(longint n);
    fibr = (n<2) ? n : fibr(n-1) + fibr(n-2);
endfunction

always @(posedge clk or posedge rst)
    if(rst)
        r <= 0;
    else
        if(start)
            r <= fibr(n);
endmodule

module fibB(input          clk,
            input          rst,
            input logic [5:0] n,
            input logic     start,
            output logic [20:0] r,
            output logic     busy);

logic [20:0] a,b;
logic [ 5:0] j;
always @(posedge clk or posedge rst)
    if(rst)
        {busy, r, a, b} <= 0;
    else
        if(start)
            {busy, r, a, b, j} <=
                {1'd1, 21'd0, 21'd1, 21'd0, n};
        else
            if(j!=0) begin
                a <= a+b;
                b <= a;
                j <= j-1;
            end
            else {busy, r} <= {1'd0, b};
endmodule

module fibC(input          clk,
            input          rst,
            input logic [5:0] n,
            input logic     start,
            output logic [20:0] r);

// initialise ROM
logic [20:0] fibrom [0:31] =
    { 0, 1, 1, 2,
      3, 5, 8, 13,
      21, 34, 55, 89,
      144, 233, 377, 610,
      987, 1597, 2584, 4181,
      6765, 10946, 17711, 28657,
      46368, 75025, 121393, 196418,
      317811, 514229, 832040, 1346269};

always @(posedge clk or posedge rst)
    if(rst)
        r <= 0;
    else if(start) r <= fibrom[n];
endmodule

```

- (i) How many clock cycles does it take for each module to compute `fib(8)` in

simulation? Justify your answer. [6 marks]

(ii) Which of these modules is synthesisable? Justify your answer. [2 marks]

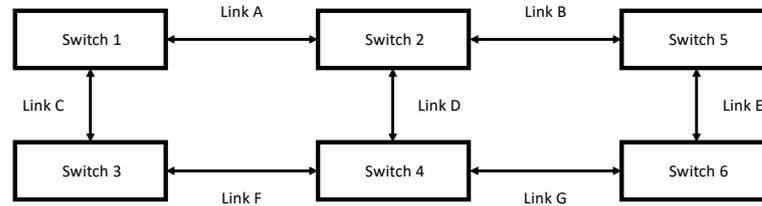
2 Computer Design

- (a) In computer architecture, why do we aim to make the common case fast? [3 marks]
- (b) If we have a new hardware divider that speeds up division by $2\times$ but results in a 2% drop in processor clock frequency, how frequently does divide need to be used to improve overall performance by 10%? [3 marks]
- (c) What properties are used to cache data efficiently? [3 marks]
- (d) What is the difference between direct-mapped, fully-associative and set-associative caches? [3 marks]
- (e) What cache-line replacement policies are typical for direct-mapped and set-associative caches? [3 marks]
- (f) Why are multi-level caches typically used? [2 marks]
- (g) What hardware is typically used to make virtual memory efficient? [3 marks]

3 Computer Design

- (a) Explain where a multicore processor with a short-vector instruction set fits within Flynn's taxonomy? [3 marks]
- (b) Describe and contrast Amdahl's and Gustafson's laws. [4 marks]
- (c) Explain how a general-purpose heterogeneous multicore, such as Arm's big.LITTLE, can out-perform a homogeneous design. [4 marks]
- (d) Computational sprinting is a technique to increase the performance of a core for a brief amount of time by temporarily raising its frequency and voltage. Describe the types of workload where computational sprinting on a general-purpose core would out-perform a specialised accelerator. [4 marks]
- (e) Comment on the advantages and disadvantages of computing within a reconfigurable fabric (e.g. an FPGA) alongside the processor within the same SoC. [5 marks]

4 Computer Networking



The network illustrated above represents an Ethernet Layer-2 network that uses spanning-tree to compute forwarding tables. Assume all links have a link-weight of one.

[Note: Tie-breaking/leader-elections use the switch identifier from this diagram.]

- (a) Compute the steady state routing/forwarding table for Switch 3.

[4 marks]

- (b) Noting which switches recompute a solution, enumerate the changed forwarding tables in switches of this network resulting from the complete failure and removal of link D.

[9 marks]

- (c) Following the removal of Link D, a new link H is added between Switch 1 and 4; however, this link fails frequently.

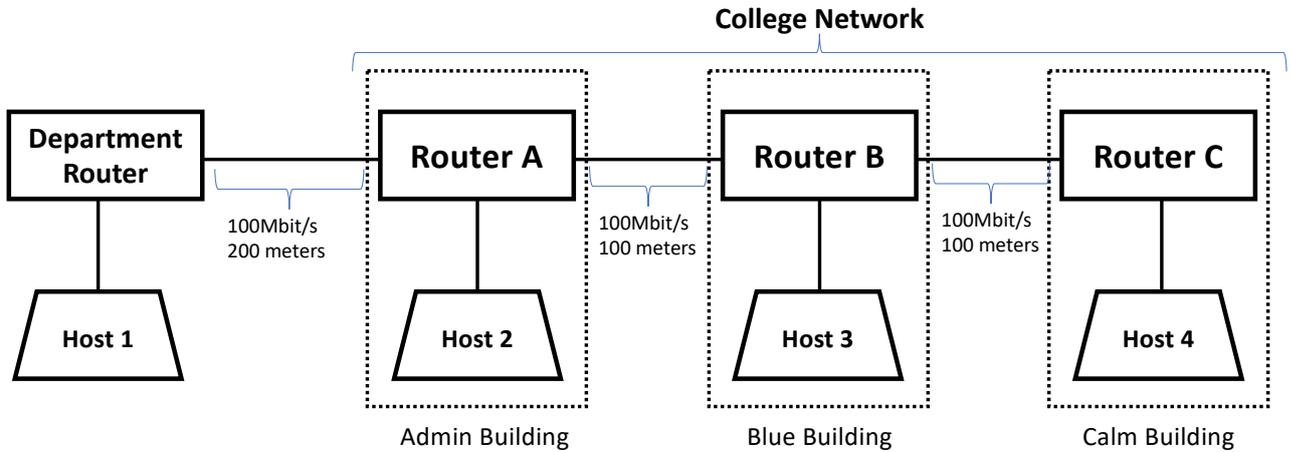
Denied access to monitor the network-traffic, outline a diagnostic strategy to identify the faulty Link H, making clear how the network-operator might use interrogation of network switch forwarding tables.

[4 marks]

- (d) Now suppose the switches do not permit interrogation of switch forwarding tables and that the link status information is untrustworthy. Outline other techniques that might be used to identify the failed link.

[3 marks]

5 Computer Networking



- (a) Using a diagram, illustrate how long a packet of length L bytes will take to travel over an idle network from Host 4 to Host 1? The routers use a store-and-forward architecture. [5 marks]
- (b) Computers in Calm building are often not getting allocated IP addresses and the performance is quite poor. The department router serves DHCP for the College network and is operating correctly. Residents in Blue report intermittent performance issues, but no one in Admin reports any problems. Network measurements reveal that the per-router packet loss for each switch under load can be as high as one packet in five thousand, but it is significantly worse for packets smaller than 1000 bytes, where as many as one packet in twenty are lost.
- With these insights, explain the cause of the problems experienced. Make clear any simplifying assumptions you have made. [5 marks]
- (c) Some students in Calm have found using IPv6 will 'work' (i.e., connecting to the wider University services is possible, but not to Internet services); although still not performing as well as when they are in the Admin building. Describe the steps by which IPv6 addresses are allocated without DHCP and consider why this service may be working more reliably than IPv4? [6 marks]
- (d) Two approaches to improve the network performance are available: one is to upgrade the performance of the physical links between the buildings to 10Gbit/s. The alternative approach is to significantly change the topology of the network by adding an additional high performance router, but leaving the performance of the physical links unchanged.

Briefly give the advantages and disadvantages of each approach. [4 marks]

6 Computer Networking

- (a) Arriving at your college room, you plug into the wired Ethernet jack for the first time. The network admin has a record of your MAC address and your machine can join the network without further action on your part.

Assume: Your laptop's Ethernet address is 0a:0b:0c:0d:0e:0f, DHCP server address is 131.111.7.3, your IPv4 address will be 131.111.7.121, the gateway's IP address is 131.111.7.1, and Ethernet address is 00:01:02:03:04:05, the network netmask is 255.255.255.0

Write the series of protocol/packet exchanges that occur on the wired Ethernet link, up until you can send a single packet to 128.232.0.20. You do not need to describe packets after this packet has left the link. Include ARP and DHCP packets, stating the IP and Ethernet addresses of the packets where possible.

[10 marks]

- (b) Consider two neighbours, Alice and Bob. Each have wireless IPv4 routers with integrated NAT. Each neighbour connects their laptop to their own wireless router, and each uses appropriate utilities to examine the IP address of each laptop. They realise the laptops have the same IP address.

(i) How is that possible? [2 marks]

(ii) Justify one reason that wide-spread deployment of IPv6 would remove the need for the NAT devices.

[2 marks]

(iii) Justify one reason that an IPv6 user might want to continue using their NAT.

[2 marks]

(iv) Further investigations show the two wireless routers are using the same wireless channel, although with different SSIDs. Detail what this situation means, why this situation is both possible and perfectly standards-compliant behaviour, and what implications there are for this situation — including how any negative effects can be made less severe.

[4 marks]

7 Concurrent and Distributed Systems

- (a) What are the distinguishing differences between a process and a thread? How are they best mapped to processors? [4 marks]
- (b) (i) What are the main features of a *monitor* in concurrency theory? [2 marks]
- (ii) How can a thread suspend itself inside a monitor? Use a state diagram or flow chart to explain your answer. [5 marks]
- (c) Define *priority inversion* and give an example of a situation where it may occur. [3 marks]
- (d) Several threads split over several processes all write lines of text to one output device. Characters from different lines must not be interleaved. The processes have different priorities. Outline the major software components (such as processes, threads, buffers, locks and system calls) involved in this task and how they interact. Identify the points where threads will block and examine whether priority inversion is likely to happen. [6 marks]

8 Concurrent and Distributed Systems

- (a) State the definition of the *happens-before* relation in a distributed system. [2 marks]
- (b) Describe how to compute Lamport timestamps. [2 marks]
- (c) Lamport timestamps are often used to provide a total ordering on events in a system with multiple communicating processes. State how this total order is defined, and explain how this order relates to the happens-before relation. [2 marks]
- (d) Explain what is meant by *FIFO broadcast* and *FIFO-total order broadcast* in the context of a process group. [2 marks]
- (e) Give pseudocode for an algorithm that implements FIFO-total order broadcast using Lamport clocks. You may assume that each process has a unique identifier, and that the set of all process IDs in the group is known. Further assume that the underlying network provides reliable FIFO broadcast; that is, you may use a function `sendFIFO(m)` that transmits message `m` to all processes in the group. The function `deliverFIFO(m)` is invoked at each process (including the sender) when message `m` is delivered at that process by the FIFO broadcast layer. Your algorithm should call `deliverTotalOrder(m)` with messages `m` in strictly monotonically increasing order of their Lamport timestamp.

Use this outline to get started:

```
// Called by the user when they want to send a message
function totalOrderBroadcast(msg) {
  let m = ...; // construct message to send by FIFO broadcast
  sendFIFO(m); // use underlying FIFO broadcast
}

// Called by the FIFO broadcast layer when a message is received
function deliverFIFO(m) {
  // Figure out the total order...
  // then call deliverTotalOrder(...) as appropriate
}
```

[10 marks]

- (f) Briefly comment on the fault-tolerance properties of your algorithm in Part (e). [2 marks]

END OF PAPER