



COMPUTER SCIENCE TRIPOS Part II

Wednesday 3 June 2015 1.30 to 4.30 pm

COMPUTER SCIENCE Paper 8

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

Rough work pad

SPECIAL REQUIREMENTS

Approved calculator permitted

1 Comparative Architectures

A multicore processor consists of eight scalar cores. Each core has private 8KB L1 instruction and data caches. The cores are supported by a 4MB L2 cache that is 8-way banked. Communication between the cores and the L2 cache banks takes place over two crossbar switches (one for communication in each direction). The L2 cache maintains a directory that shadows the L1 tags. The L1 caches are write-through, with allocate on load and no-allocate on stores.

- (a) Describe a simple cache coherence protocol suitable for this processor. [8 marks]
- (b) In the simple scheme you have described, what states may L1 cache lines be in? [2 marks]
- (c) Why might it be better to shadow the L1 tags rather than adding state to each L2 cache line? [5 marks]
- (d) What advantages does your protocol have over a simple snooping coherence protocol running over a shared bus? [5 marks]

2 Computer Systems Modelling

(a) A Poisson process of rate λ has inter-event times X_1, X_2, \dots that are independent Exponential random variables with parameter λ .

(i) Show that the random variables X_i for $i = 1, 2, \dots$ satisfy the *memoryless property*

$$\mathbb{P}(X_i > t + s | X_i > t) = \mathbb{P}(X_i > s)$$

where s and t are any positive real numbers. [2 marks]

(ii) Suppose that T_1, T_2, \dots, T_n is an observed sequence of n consecutive event times. Describe what tests you could conduct of the hypothesis that the observed events arise from a Poisson process of some given rate λ .

[6 marks]


(b) (i) Describe the M/G/1 queue model explaining the mathematical assumptions made and the required features of the queue. [4 marks]

(ii) Suppose that the arrival rate for the M/G/1 queue is λ customers per second, that the service rate is μ customers per second and assume that $\rho = \lambda/\mu < 1$. Define the utilization, U , of the server and show that $U = \rho$. [2 marks]

(iii) A *busy period* of the M/G/1 queue is a period of time which starts when the server becomes occupied and continues until the server is no longer occupied. An *idle period* is the period of time between consecutive busy periods. Explain how to find the mean length of an idle period. Using your result for the utilization of the server in part (b)(ii), derive an expression for the mean length, τ , of a busy period. [6 marks]


3 Computer Vision

- (a) A Bayesian classifier uses observations x to assign visual objects to either one of two classes, C_1 or C_2 . Their baseline prior probabilities are $p(C_1)$ and $p(C_2)$, with sum $p(C_1) + p(C_2) = 1$. Observations x have unconditional probability $p(x)$, and the class-conditional probabilities of a given observation x are $p(x|C_1)$ and $p(x|C_2)$.



- (i) Using the above quantities provide an expression for $p(C_k|x)$, the likelihood of class C_k given an observation x . [2 marks]
- (ii) Provide a decision rule using $p(C_k|x)$ and $p(C_j|x)$ for assigning classes based on observations that will minimise misclassification. [2 marks]
- (iii) Now express your decision rule instead using only the quantities $p(C_k)$, $p(C_j)$, $p(x|C_k)$, $p(x|C_j)$, and relate it to the diagram above. [2 marks]
- (iv) If the classifier decision rule assigns class C_1 if $x \in R_1$, and C_2 if $x \in R_2$ as shown in the figure, what is the total probability of error? [3 marks]
- (v) If classifier decisions are made by computing functions $y_k(x)$, $y_j(x)$ of the observations x and assigning class C_k if $y_k(x) > y_j(x) \forall j \neq k$, for example $y_k(x) = p(C_k|x)$, what are such functions $y_k(x)$ called? [1 mark]
- (b) Discuss the significance of the fact that typically in mammalian visual systems, there are almost ten times more corticofugal neural fibres sent back down from the visual cortex to the thalamus, as there are ascending neural fibres bringing visual data from the retina up to the thalamus. Does this massive neural feedback projection support the thesis of “vision as graphics” and, if so, how? [5 marks]

- (c) Discuss the theory of vision as model building, hypothesis generation and testing, and knowledge-based processing, in light of the paradoxical figure on the right. What do we learn from bistable or rivalrous percepts? Discuss how top-down context information should drive the integration of low-level data into meaningful visual wholes.



[5 marks]

4 Digital Signal Processing

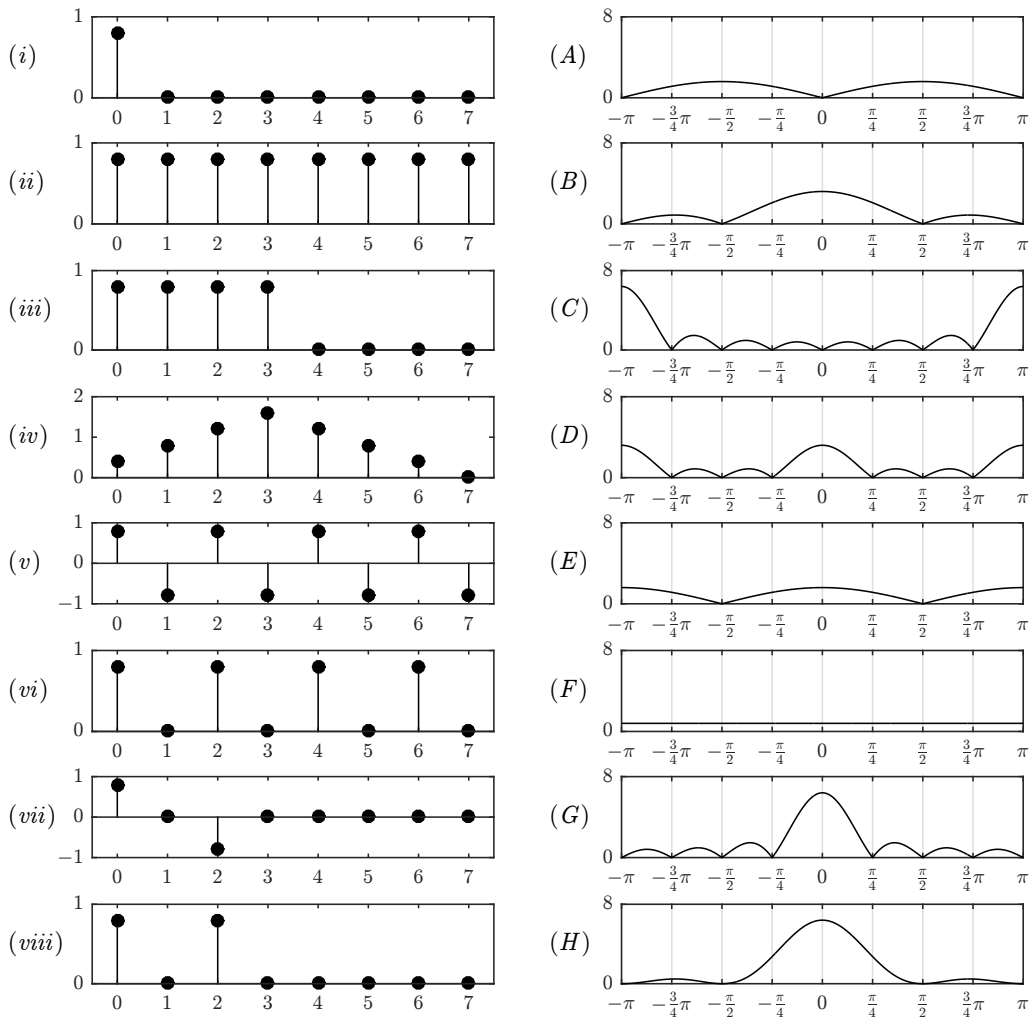
The discrete-time Fourier transform (DTFT) of a discrete sequence $\{x_n\}$ can be defined as

$$X(e^{j\omega}) = \sum_{n=-\infty}^{\infty} x_n \cdot e^{-j\omega n}$$

- (a) If $\{x_n\}$ was the result of sampling a signal at sampling rate f_s and we want to know its DTFT at frequency f , what will be the corresponding value for ω ? [2 marks]

- (b) If $\{x_n\}$ has only real values and we know the value of $X(e^{j\pi/4})$, what is the value of $X(e^{j\pi \times 3.75})$? [2 marks]

- (c) Each of the eight plots (i)–(viii) below shows real-valued samples x_0, \dots, x_7 from a discrete sequence $\{x_n\}$, with $x_n = 0$ for $n < 0$ or $n > 7$. For each of these eight sequences, identify which of the eight plots (A)–(H) shows the magnitude $|X(e^{j\omega})|$ of the corresponding discrete-time Fourier transform. [8 × 2 marks]



5 E-Commerce

- (a) What is a block chain in the context of electronic money? [4 marks]
- (b) Explain the use of a block chain in Bitcoin. Why is it required? [4 marks]
- (c) Discuss whether Bitcoin is anonymous. [6 marks]
- (d) Discuss whether Bitcoin is a currency. [6 marks]

6 Temporal Logic and Model Checking

In this question assume that p and q are atomic formulae.

- (a) Compare and contrast path formulae and state formulae in temporal logic. [4 marks]
- (b) Describe and contrast the meanings of $F(G p)$ and $AF(AG p)$. [4 marks]
- (c) Describe and contrast the meanings of $G(F p)$ and $AG(AF p)$. [4 marks]
- (d) Write down and justify a temporal logic formula that expresses the property that some state satisfying q is reachable from every state satisfying p . [4 marks]
- (e) Write down and justify a temporal logic formula that expresses the property that no path contains a consecutive sequence of 256 states satisfying p . [4 marks]

7 Information Retrieval

- (a) Consider a standard bag-of-words model for the document retrieval problem.
- (i) Give an expression for the tf-idf weighting scheme which assigns a weight to each term in a *document*. Motivate each part of your expression, using the notion of Zipf's law when appropriate. [3 marks]
 - (ii) How might you modify your tf-idf scheme for each term in a *query*? Why might you use different schemes for documents and queries? [3 marks]
- (b) Edit distance can be used for spelling correction in search queries.
- (i) Define edit distance. [1 mark]
 - (ii) As an example of how to calculate edit distance efficiently, show how dynamic programming can be used to calculate the edit distance between *able* and *belt*. [5 marks]
- (c) The PageRank algorithm uses a model of a “random surfer” to calculate the importance or validity of a page. Describe how the random surfer can be modelled as an ergodic Markov chain, and how this leads to the PageRank values being calculated as a principal left eigenvector of the transition probability matrix. (You are not required to give a formal definition of an ergodic Markov chain; an informal description will suffice.) [8 marks]

8 Quantum Computing

(a) Consider the following two-qubit quantum state, $|\phi\rangle$.

$$\frac{\sqrt{2}}{3\sqrt{3}}|00\rangle - \frac{1}{\sqrt{6}}|01\rangle + \frac{2i\sqrt{2}}{3\sqrt{3}}|10\rangle - \frac{5i}{3\sqrt{6}}|11\rangle$$

- (i) What are the probabilities of outcomes 0 and 1 if the first qubit of $|\phi\rangle$ is measured?
- (ii) What are the probabilities of outcomes 0 and 1 if the second qubit of $|\phi\rangle$ is measured?
- (iii) What is the state of the system after the first qubit of $|\phi\rangle$ is measured to be a 0?
- (iv) What is the state of the system if the second qubit of $|\phi\rangle$ is measured to be a 1?
- (v) What are the probabilities of outcomes 0 and 1 if the second qubit of the system is measured, after the first qubit of $|\phi\rangle$ has been measured to be 0?
- (vi) What are the probabilities of outcomes 0 and 1 if the first qubit of the system is measured, after the second qubit of $|\phi\rangle$ has been measured to be 1?

[2 marks each]

(b) The two qubit *quantum Fourier transform* is given by the following matrix.

$$F_2 = \frac{1}{2} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -1 & -i \\ 1 & -1 & 1 & -1 \\ 1 & -i & -1 & i \end{bmatrix}$$

Sketch a circuit for implementing the operator F_2 using any combination of 1-qubit Hadamard gates; 1-qubit Pauli gates; 2-qubit C-NOT gates; controlled phase shifts and swap gates (the *swap gate* S is defined by $S|xy\rangle = |yx\rangle$). Briefly explain your circuit. [8 marks]

9 Security II

You are working on an encryption device with your new colleague, Mallory Baish, who proposes that you use a pseudo-random generator

$$r_i = h_1(s_i), \quad s_{i+1} = h_2(s_i)$$

where $s_0 \in G$ is the random initial state and the other $s_i \in G$ are subsequent internal states, all invisible to adversaries. The $h_1, h_2 : G \rightarrow G$ are two secure one-way functions.

Adversaries may see any of the past outputs r_0, \dots, r_{n-1} . If they can predict from those, with non-negligible probability, the next value r_n , then the security of your device will be compromised.

- (a) Give a rough estimate for the probability that an adversary can predict r_n , as a function of n and $|G|$. Explain your answer. [6 marks]
- (b) Mallory also suggests a specific implementation:

$$\begin{aligned} h_1(x) &= f(u^x \bmod p) & p &= \text{a 2056-bit prime number} \\ h_2(x) &= f(v^x \bmod p) & u, v &= \text{two numbers from } \mathbb{Z}_p^* \\ f(x) &= x \bmod 2^{2048} & G &= \mathbb{Z}_{2^{2048}} \end{aligned}$$

- (i) The constants p , u and v will be known to the adversary. What conditions should they fulfill so that h_1 and h_2 can reasonably be described as one-way functions, and how would you normally generate suitable numbers u and v ? [*Hint*: quadratic residues] [4 marks]
- (ii) If f were replaced with the identity function, how could an adversary distinguish the r_i emerging from this pseudo-random generator from a sequence of elements of \mathbb{Z}_p^* picked uniformly at random? [4 marks]
- (iii) After you choose a value for p , Mallory urges you to use two particular values for u and v generated in your absence. You briefly see “ $v = u^e \bmod p$ ” scribbled on a whiteboard. You become suspicious that Mallory is trying to plant a secret backdoor into your pseudo-random generator.

Explain how Mallory could exploit such a backdoor. [6 marks]

10 System-on-Chip Design

- (a) Explain what factors limit the complexity and performance of an SoC at the heart of a portable electronic device. [4 marks]
- (b) Compare and contrast the use of hardware and software to implement a compute-intensive algorithm on an SoC, such as data encryption. Include customised processors and co-processors in your analysis. [5 marks]
- (c) (i) Define the term *fully-pipelined* with respect to a hardware component. [2 marks]
- (ii) Describe and compare three designs for a fixed or floating-point multiplier that vary in performance: one at least should be fully pipelined. [6 marks]
- (iii) Define the term *structural hazard* and explain why these can affect system performance. Which of your designs from part(c)(ii) might present such a hazard and why? [3 marks]

11 Topical Issues

- (a) Compare and contrast the radio channels used by Bluetooth Low Energy (BLE) and WiFi in the 2.4 GHz ISM radio band. Explain why BLE uses only three channels for advertisements and how the three were selected. [6 marks]
- (b) Consider a wireless mouse computer peripheral. Analyse the use of WiFi, BLE and a proprietary radio system for communication with the computer, giving advantages and disadvantages for each. [10 marks]
- (c) Bluetooth 4.0 specifies a per-packet Cyclic Redundancy Check (CRC) to protect each 31-octet BLE packet as a whole. In contrast the longer WiFi packets contain a separate CRC for each of the packet header and the packet payload. When Bluetooth 4.1 introduced longer BLE packets, it retained a single CRC for the entire packet. What are the consequences of this decision? [4 marks]

12 Topics in Concurrency

This question is on an authentication protocol using a key server and symmetric keys. $Key(X, Y)$ represents the symmetric key used to encrypt messages sent by X to Y , and symbols K and K' are used as variables over keys. **SPL** terms representing a key server S , an initiator A and responder B are:

$$\begin{aligned} S &= !(in \{X, Y\}_{Key(X,S)}. out \{Key(X, Y), Key(Y, X), Y\}_{Key(S,X)}) \\ A &= out \{A, B\}_{Key(A,S)}. in \{K, K', B\}_{Key(S,A)}. out \{m\}_K. in \{m, m\}_{K'} \\ B &= out \{B, A\}_{Key(B,S)}. in \{K', K, A\}_{Key(S,B)}. in \{\psi\}_K. out \{\psi, \psi\}_{K'} \end{aligned}$$

- (a) (i) The capabilities assumed of an attacker when public-key cryptography is used for authentication, as when studying the Needham-Schröder-Lowe protocol, are that it can pair messages, split paired messages, encrypt messages under a public key and decrypt messages under a public key if it has access to the private key.

Give four **SPL** processes Spy_1, \dots, Spy_4 representing these capabilities.

[4 marks]

- (ii) Give a further two processes Spy_5, Spy_6 representing the capability of an attacker to encrypt and decrypt messages when symmetric-key cryptography is used.

[2 marks]

- (b) Let $P_{Spy} =!(\parallel_{i \in \{1, \dots, 6\}} Spy_i)$. Draw the events of the Petri net for

$$P_{Spy} \parallel S \parallel A \parallel B.$$

For P_{Spy} , only show those from Spy_5 and Spy_6 .

[7 marks]

- (c) *Secrecy* of the message m can be viewed as m never being output directly to the network by either the participants in the protocol or the attacker.

Give a reasonable general condition on the set of messages initially assumed to have been output to the network for which secrecy of m holds. You may assume that if $Key(X, Y) = Key(X', Y')$ then $X = X'$ and $Y = Y'$.

Describe the principles underlying a proof of the secrecy of the message m .

[7 marks]

END OF PAPER