COMPUTER SCIENCE TRIPOS Part II

Wednesday 5 June 2013 1.30 to 4.30

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

SPECIAL REQUIREMENTS
Approved calculator permitted

1 Advanced Graphics

(*a*) Beginning with a sphere of stone, a sculptor slices the sphere in half and carves a curving path of semi-circular cross-section, tracing the letter S into the flattened face of the sphere as shown at left:

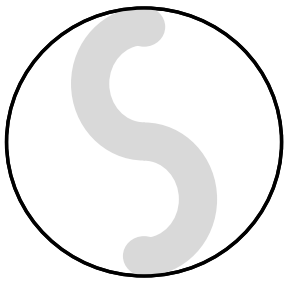

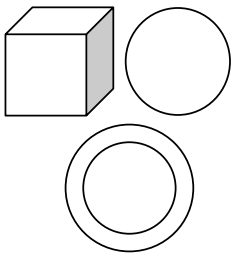

The top and bottom of the S perfectly touch the top and bottom of the hemisphere. The dimensions of the sculpture are as follows:

- Sphere radius: 100cm
- Line width: 20cm
- Radius of the ends of the S: 10cm

Describe how you would build this sculpture using the technique of *Constructive Solid Geometry*. Assume that you have only three primitives, each centred on the origin:

- A *cube*, where each edge is 10cm long.
- A *sphere*, of radius 100cm.
- A *torus*, where the radius of a cross-section of the tube is 2cm and the ring of the tube is a circle of radius 9cm. The torus lies in the xy plane.

For full marks, specify every transformation, in order, for every primitive (e.g., "Translate the cube by 1m up x" or "Rotate this object by 45 degrees around the z-axis") and every binary operation between primitives. [8 marks]

(*b*) (*i*) Given two disks of radius 1, one centered at $(1, 1, 1)$ with normal vector $(0, 0, -1)$ and the other centered at $(1, 1, -1)$ with normal vector $(0, \sqrt{2}, \sqrt{2})$, compute the exact *radiosity* view factor between them. Clearly state the equation you use. Assume there is no occlusion between the two disks. [4 marks]

(*ii*) Briefly describe an efficient mechanism for using modern hardware to compute (approximate) view factors between patches in a radiosity system, including occlusion. [4 marks]

(*iii*) Describe a hybrid method which could produce images which both solve the global illumination problem with a radiosity solution and also correctly portray lighting phenomena such as caustics. [4 marks]

2 Artificial Intelligence II

(a) Denoting the *utility* of a state s of the world by $U(s)$, and given that we have *evidence* E regarding the world and probability distribution $\Pr(s|A = a, E)$ modelling the effect of taking specific actions, define the *expected utility* associated with taking an action. [2 marks]

(b) Evil Robot is, despite his undoubted evilness, very shy where romance is concerned. He has fallen for a beautiful vacuum cleaner, called SN05718, and is wondering whether to ask her to accompany him for a change of oil. He rates the utility of going alone as -10 and the utility of being accompanied as $+100$. Being shy, he feels that if he asks her then she will accept with probability 0.1, but if he does not then there is a small chance of 0.01 that she will in fact ask him. What is the expected utility of this situation? [3 marks]

(c) If in the scenario described in part (b) we discover that it is possible to obtain further evidence E' in addition to E regarding the world, derive an expression for the *value of perfect information* associated with finding the value of E'. [5 marks]

(d) Evil Robot has a plan to acquire SN05718's diary to find out whether or not she likes him. He believes that the likelihood of her liking him is 0.3. Also, if she likes him and he asks her to accompany him he thinks she will accept with probability 0.7, whereas if he does not ask she will in any case accompany him with probability 0.4. On the other hand, if she does not like him then the corresponding probabilities are 0.05 and 0.01. It will cost Evil Robot 15 to get someone to steal the diary. Compute whether or not he should. [10 marks]

3 Comparative Architectures

(*a*) How do superblock and trace scheduling differ? [4 marks]

(*b*) How might a programmer improve the performance of a program given detailed knowledge of a processor's memory hierarchy? [6 marks]

(*c*) Larger scale networks, i.e. those involving chip-to-chip or longer distance communications, have been designed for many years. What new challenges and constraints are introduced when designing on-chip networks? [6 marks]

(*d*) As fabrication technologies scale the performance of wires improves slowly relative to that of transistors. Why is this particularly problematic when attempting to increase the performance of superscalar processors? [4 marks]

4 Computer Systems Modelling

Consider a single server queue with customer arrivals occuring at times of a Poisson process of rate λ. Customers are served with independent exponential service times which have mean $1/\hat{\mu}$ if there are less than k customers present and mean $1/\mu$ if there are k or more customers present and where k is fixed. The number present, n, can be modeled as a birth-death process with the birth rates $\lambda_n = \lambda$ for each $n = 0, 1, 2, \ldots$ and state-dependent death rates

$$\mu_n = \begin{cases} \hat{\mu} & 1 \leq n < k \\ \mu & n \geq k \end{cases}.$$

Write $\hat{\rho} = \lambda/\hat{\mu}$ and $\rho = \lambda/\mu$ and assume that $\rho < 1$.

(a) Find an expression for π_n the probability of being in state n under the equilibrium distribution which you may assume to exist. [4 marks]

(b) Show that if $\hat{\mu} = \mu$ then your result for π_n in part (a) coincides with the case of a M/M/1 queue, namely $\rho^n(1-\rho)$. [1 mark]

(c) Find an expression for L the expected number of customers in the system. [6 marks]

(d) Show that

$$L_q = L - (1 - \pi_0)$$

where L_q is the expected number of customers in the queue waiting for service. [2 marks]

(e) Find expressions for the expected time, W, that a customer spends in the system and the expected time, W_q, that a customer spends waiting for service. [2 marks]

(f) Consider an example where customers arrive according to a Poisson process with a mean inter-arrival time of 30 minutes. Suppose that the service times are exponential with mean 40 minutes if there are no customers waiting but have mean 20 minutes if there are any customers waiting. Compute π_0, L, L_q, W and W_q. [5 marks]

5 Computer Vision

(*a*) Consider the following 2D filter function $f(x, y)$ incorporating the Laplacian operator that is often used in computer vision:

$$f(x,y) = \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right) e^{-\left(x^2+y^2\right)/\sigma^2}$$

(*i*) In 2D Fourier terms, what type of filter is this? Is it a lowpass, a highpass, or a bandpass filter? Justify your answer. [2 marks]

(*ii*) Are different orientations of image structure treated differently by this filter, and if so, how? Is it isotropic, or anisotropic? [2 marks]

(*iii*) Approximately what is the spatial frequency bandwidth of this filter, in octaves? [Hint: the answer is independent of σ.] [2 marks]

(*iv*) What is meant by image operations "at a certain scale of analysis?" Explain the scale parameter σ, and define a scale-space fingerprint. [2 marks]

(*b*) Write a block of pseudo-code for convolving an image with a feature-detecting kernel. (You may ignore out-of-bounds issues at the image array boundaries.) [3 marks]

(*c*) In pattern classification with two classes, explain how an ROC curve is derived from the underlying distributions. Define a threshold-independent performance metric based on the distributions' moments. [4 marks]

(*d*) When visually inferring a 3D representation of a face, it is useful to extract separately both a shape model, and a texture model. Explain the purposes of these steps, their use in morphable models for pose-invariant face recognition, and how the shape and texture models are extracted and later re-combined. [5 marks]

6 Digital Signal Processing

Consider the discrete system

$$y_n = \sum_{i=0}^{\infty} x_{n-2i} \cdot \left(-\frac{1}{2}\right)^i$$

(*a*) Write down the first 8 samples of the impulse response of this filter. [2 marks]

(*b*) Provide the finite-difference equation of an equivalent recursive filter that can be implemented with not more than two delay elements. [4 marks]

(*c*) What is the z-transform $H(z)$ of the impulse response of this filter? [4 marks]

(*d*) Where are the zeros and poles of $H(z)$? [6 marks]

(*e*) We now operate this discrete system at sampling frequency $f_\mathrm{s} = 1$ MHz and feed it with input $x_n = \cos(2\pi f n / f_\mathrm{s})$. For which f (with $0 \leq f \leq f_\mathrm{s}/2$) will the peak amplitude of the output sequence $\{y_n\}$ be largest, and how large will it be? [4 marks]

7 E-Commerce

(*a*) Discuss how you would use social media to attract traffic to your website. [4 marks]

(*b*) Discuss the difficulties in measuring the use of your website and related social media campaign. [5 marks]

(*c*) Discuss whether the use of third party analytics creates opportunities to leak personal information. [5 marks]

(*d*) Discuss why it is sometimes good to give things away. [6 marks]

8 Hoare Logic

Use notation from logic ($\forall$, $\exists$, etc.) in your answers to the questions below.

(a) Define the semantics of the partial correctness Hoare triple, $\{P\}\,C\,\{Q\}$. Briefly explain this definition. [3 marks]

(b) Define the semantics of the total correctness Hoare triple, $[P]\,C\,[Q]$. Explain what is 'total' about total correctness. [3 marks]

(c) State an inference rule for partial correctness Hoare Logic that is not sound in total correctness Hoare Logic. Explain your choice. [3 marks]

(d) State and briefly explain the semantics of the separation logic Hoare triple. Point out at least two differences between $\{P\}\,C\,\{Q\}$ in traditional Hoare logic and separation logic. [5 marks]

(e) Carefully state an inference rule that is part of separation logic but not present in traditional Hoare logic. [3 marks]

(f) Point out at least two aspects in which the semantics of Hoare logic or separation logic do not reflect the semantics of real programming languages. [3 marks]

9 Information Retrieval

Question Answering (QA) is the task of building a system that can select answer strings from a corpus, in response to a natural language question.

(a) Mean Reciprocal Rank (MRR) has been used for several years by the TREC QA organisers as the main evaluation metric. Give the formula for MRR, and state two disadvantages of this method. [4 marks]

(b) The Microsoft system competing in the TREC-10 conference (QA track) used an unusual approach for Question Answering (QA). Give a brief overview of how the system works. [4 marks]

(c) Describe in detail, with an example, how the Microsoft system generates query strings and scores answers. [6 marks]

(d) Consider the following question answer pair from the official TREC QA corpus:

> **Question:** Who killed Abraham Lincoln?
> **Answer:** John Wilkes Booth is perhaps America's most infamous assassin. He is best known for having fired the bullet that ended Abe Lincoln's life.

Can the Microsoft system as described in part (b) be awarded a point in the TREC competition for finding this answer? If yes, describe how and under which conditions. If no, describe changes to the system in part (b) that would make it possible to do so. [6 marks]

10 Principles of Communications

The Internet makes use of distributed route computation algorithms such as link-state, distance-vector, and path-vector schemes. When these were designed and implemented in the early days, there was little information about the possible future statistical properties of the topology of the Internet. Since then, measurements show that the way the network has evolved has led to highly clustered, and possibly small-world characteristics in the graph representing the topology.

In this situation, how might you choose to design your routing system, and furthermore how might the topology impact how you choose to represent the graph for the purposes of a link-state or distance-vector computation, and for forwarding? [20 marks]

11 Quantum Computing

Let a_0a_1 be the two-bit representation of $a \in \{0, 1, 2, 3\}$. We define the 2-bit Boolean function f_a by:

$$f_a(x_0, x_1) = (a_0 \cdot x_0) \oplus (a_1 \cdot x_1).$$

where $\cdot$ denotes Boolean *and* and $\oplus$ represents *exclusive or*.

For each such function f, let U_f denote the 3-qubit unitary operator that computes f in the sense that:

$$U_f|x_0x_1y\rangle = |x_0x_1\rangle|y \oplus f(x_0, x_1)\rangle.$$

In the following, $H^{\otimes n}$ denotes the n-bit Hadamard operator.

(a) Show how U_{f_2} can be implemented with the use of C-NOT gates. [3 marks]

(b) Show that:

$$(H^{\otimes 2}\text{C-NOT}H^{\otimes 2})|xy\rangle = \text{C-NOT}|yx\rangle.$$

[3 marks]

(c) Using (b) or otherwise, show that for each of the four possible values of a, the operator $(H^{\otimes 3}U_{f_a}H^{\otimes 3})$ can be implemented as a circuit using only C-NOT gates. [6 marks]

(d) Given a black box implementing U_{f_a} for an unknown value of a, show that we can construct a quantum circuit that determines the value of a with certainty, using the black box only once.

[*Hint:* Consider the circuit from (c) applied to a suitable computational basis state.]

[8 marks]

12 Security II

The RSA public-key crypto system performs calculations in the group $\mathbb{Z}_n$, with $n = pq$ being the product of two large prime numbers p and q. The public key consists of the tuple (n, e), with $\gcd(\phi(n), e) = 1$, and the corresponding private key is (n, d). A message $m \in \mathbb{Z}_n$ is encrypted via $c = m^e \bmod n$ and decrypted as $m = c^d \bmod n$.

(*a*) Given p, q, and e, how can you apply the extended Euclidian algorithm to find a suitable d? [6 marks]

(*b*) If we modified RSA to use as the public modulus a prime number instead of a composite of two large prime numbers, that is $n = p$ instead of $n = pq$, would this affect its security, and if so how? [4 marks]

(*c*) In the *UltraSecure* virtual-private network, each router knows of each of its remote communication peers the RSA public key (n, e), which all have $e = 3$ and $2^{1023} \leq n < 2^{1024}$. If router *alice* needs to establish a shared 256-bit AES secret key k with remote router *bob*, it looks up *bob*'s (n, e) and then uses this method:

- *alice* picks $k \in \{0, 1\}^{256}$ by reading 32 bytes from `/dev/random`
- *alice* interprets k as binary integer m with $0 \leq m < 2^{256}$
- *alice* sends $c = m^e \bmod n$ to *bob*
- *bob* decrypts c into m and recovers k (by removing leading zeros)

Then *alice* and *bob* secure the rest of their communication with shared secret k.

(*i*) How could an eavesdropper obtain m from c? [4 marks]

(*ii*) Suggest a better method of using RSA to establish an AES key than the one given above. [6 marks]

13 System-on-Chip Design

(*a*) Why do System-on-Chip designs use both on/off power control over subsystems as well as adjustable supply voltages when a subsystem is switched on? [5 marks]

(*b*) How might the two techniques from part (*a*) be used in conjunction in a server that contains four similar processing elements that take jobs from a shared queue? State any assumptions you make. For instance, you might assume each processing element consists of about 50,000 gates, that job queue entries are about a kilobyte in length and that their arrival rate varies greatly. [5 marks]

(*c*) A simulation of the four processing elements that modelled each gate in detail would be slow. Briefly describe *two* alternative simulation models that respectively model less and far less detail, while still preserving accuracy in terms of the job queue length variation. [3 marks each]

(*d*) Would the supply voltage variation need to be modelled in each of your two models of part (*c*)? [2 marks]

(*e*) How can a single simulation mix a low-level model of one processing element with a high-level model of the remainder? [2 marks]

14 Topics in Concurrency

This question is on basic Petri nets (Petri nets in which the pre- and post-condition multirelations are relations and all places have capacity 1) and the modal μ-calculus.

(a) Draw basic Petri nets to illustrate each of the following:

(i) Independence (also called concurrency)

(ii) Backwards conflict

(iii) Forwards conflict

(iv) Contact [4 marks]

(b) Prove that if $\mathcal{M} \xrightarrow{e_1} \mathcal{M}_1 \xrightarrow{e_2} \mathcal{M}'$ and e_1 and e_2 are independent events in a basic Petri net then there exists a marking $\mathcal{M}_2$ such that $\mathcal{M} \xrightarrow{e_2} \mathcal{M}_2 \xrightarrow{e_1} \mathcal{M}'$. [3 marks]

(c) Draw a basic Petri net with four events $\{a, b, c, d\}$ that gives rise to the following transition system. The states of the transition system should correspond to reachable markings of the Petri net, the state corresponding to the initial marking should be s, and the transitions should be labelled by the event of the Petri net that generates them.

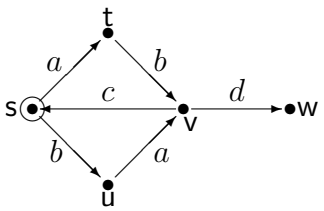

[3 marks]

(d) With respect to the transition system drawn in part (c), determine which of the states $\{\mathsf{s}, \mathsf{t}, \mathsf{u}, \mathsf{v}, \mathsf{w}\}$ satisfy the following modal μ-calculus assertions:

(i) $\mu X.(\langle d \rangle T \vee \langle . \rangle X)$

(ii) $\nu Y.([.]Y \wedge \langle . \rangle T)$

Justify your answers. [4 marks]

(e) Prove that in a finite-state transition system, $s \models \nu X.(\langle a \rangle T \wedge \langle . \rangle X)$ if and only if there exists an infinite path from s along which an a-action can occur in every visited state. [6 marks]

END OF PAPER