COMPUTER SCIENCE TRIPOS Part II

Wednesday 5 June 2002 1.30 to 4.30

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

1 Denotational Semantics

Let D be a cpo with bottom element $\bot$. Let $k : D \to D$ be a continuous function. Let $\mathbb{B} = \{true, false\}$. Define the conditional function

$$\text{if} : \mathbb{B}_\bot \times D \times D \to D$$

by $\text{if}(b, d, d') = d$ if $b = true$, d' if $b = false$, and $\bot$ otherwise. Let $h : D \to \mathbb{B}_\bot$ be a continuous function which is strict (so $h(\bot) = \bot$).

The function f^* is the least continuous function from D to D such that

$$\forall x \in D.\ f^*(x) = \text{if}(h(x),\ x,\ f^*(k(x)) \ .$$

(*a*) State the principle of fixed point induction. [3 marks]

(*b*) Show that

$$\forall x \in D.\ h(f^*(x)) = \text{if}(h(x),\ h(x),\ h(f^*(k(x)))) \ . \quad \text{[4 marks]}$$

(*c*) Prove that the property

$$Q(f) \Leftrightarrow_{def} \forall x \in D.\ h(f(x)) \sqsubseteq true$$

is admissible. [5 marks]

(*d*) Prove $Q(f^*)$ by fixed point induction. [8 marks]

2 VLSI Design

(*a*) Draw a transistor-level circuit diagram of a 2-input NAND gate in CMOS and describe its operation. [6 marks]

(*b*) Sketch a stick diagram and physical layout for the same gate, paying attention to the dimensions of the transistors to ensure balanced rise and fall times. [6 marks]

(*c*) Draw a transistor-level circuit diagram of a 2-input Muller C-element and describe its operation. Why is this sometimes described as an AND gate for events? [8 marks]

3 Digital Communication II

(*a*) Describe the design and operation of a modern high-performance IP router, such as might be found in the Internet core. [12 marks]

(*b*) Contrast the above design with that of a similar performance ATM switch. [4 marks]

(*c*) Explain why building switches and routers that are able to keep up with transmission link rates is becoming increasingly difficult. [4 marks]

4 Advanced Graphics

(*a*) Describe how an object built using constructive solid geometry (CSG) can be represented using a binary tree. Given the intersection points of a ray with each primitive in the tree, show how to calculate the first intersection point of the ray with the entire CSG object. [6 marks]

(*b*) Implicit surfaces are normally combined by adding the field functions together to create a "blobby" blended surface. Describe an alternative mechanism (or mechanisms) for combining implicit surfaces which would produce results more akin to CSG union and intersection. Explain why it produces these results. Given this mechanism, suggest a way of combining implicit surfaces to produce a result similar to CSG difference. [4 marks]

(*c*) Describe the basic radiosity algorithm. [10 marks]

5 Business Studies

(*a*) Distinguish between top-down, bottom-up and spiral (rapid prototype) development methodologies. Illustrate your answer with reference to an example of designing a building. [5 marks]

(*b*) You are in charge of commissioning the design of a new building, such as the new Computer Laboratory building. Draw up a high-level GANTT chart for this task up to the letting of the building contract. [10 marks]

(*c*) Discuss what monitoring and quality control procedures might apply to the design process. How will you get the agreement of the various stakeholders? [5 marks]

6 Security

(*a*) Explain the concept of a *Trusted Computing Base* and outline its meaning in the context of the access control provided by a typical Unix workstation. [5 marks]

(*b*) An automatic teller machine (ATM) communicates with a central bank computer for PIN verification. A 32-bit CRC code is added to each packet to detect transmission errors and then the link is encrypted using a block cipher in counter mode. Describe an attack that is possible in this setup. [5 marks]

(*c*) (*i*) Describe a cryptographic protocol for a prepaid telephone chip card that uses a secure 64-bit hash function H implemented in the card. In this scheme, the public telephone needs to verify not only that the card is one of the genuine cards issued by the phone company, but also that its value counter V has been decremented by the cost C of the phone call. Assume both the card and the phone know in advance a shared secret K. [5 marks]

(*ii*) Explain the disadvantage of using the same secret key K in all issued phone cards and suggest a way around this. [5 marks]

7 Optimising Compilers

Consider the ML-like language given by abstract syntax

$$e ::= x \mid \lambda x.e \mid e_1\ e_2 \mid \mathit{ref}\, e \mid !e \mid e_1 := e_2 \mid \mathit{if}\, e_1\ \mathit{then}\, e_2\ \mathit{else}\, e_3$$

where x ranges over variable names. A previous compiler phase has already determined a type for each variable and has checked the program is well-typed, where types have syntax

$$\tau ::= \mathit{int} \mid \mathit{intref} \mid \tau \to \tau'.$$

Note that there is no polymorphism and moreover if-then-else uses an integer rather than a boolean for a condition and the result of an assignment is the *int* value assigned rather than a unit value.

(*a*) Give an *effect system* (also known as an annotated type system) in which we can derive judgements of the form

$$\Gamma \vdash e : t, F$$

where t is an extended form of τ and Γ is a set of assumptions of the form $x : t$. Effects F are subsets of $\{A, R, W\}$ representing that the side-effects of evaluating e may respectively include a *ref* allocation (A), a dereferencing read (R) or an update (W). [12 marks]

(*b*) Give types and effects for the following programs, commenting briefly on any issues or problems your scheme encounters and how they may be resolved. (Assume g represents a global variable of type *intref*.)

- $\mathit{if}\, !g\ \mathit{then}\ \mathit{ref}\, 1\ \mathit{else}\ g$
- $\lambda x.\mathit{if}\, !g\ \mathit{then}\ \mathit{ref}\, x\ \mathit{else}\ g$
- $\mathit{if}\, !g\ \mathit{then}\ \lambda x.\mathit{ref}\, x\ \mathit{else}\ \lambda x.g$ [8 marks]

8 Advanced Algorithms

(a) Explain what is meant by the Kolmogorov Complexity $K(n)$ of a natural number n. [5 marks]

(b) Consider a graph of the function $K(n)$ plotted against n:

(i) Show that it is smooth, in the sense that for any n and fairly small value of k the value of $K(n+k)$ will be quite close to the value of $K(n)$. [3 marks]

(ii) Show that it is rough, in the sense that for any N there are two values of n_1 and n_2 between N and $2N$ such that $K(n_1)$ is about $2^{K(n_2)}$, i.e. one has a complexity exponentially bigger than the other. [3 marks]

(iii) Explain why the graph is bounded above by some straight line of the form $n + c$ and comment on what the constant represents. [3 marks]

(iv) Explain why for any constant k there will be a value N such that $n > N$ implies that $K(n) > k$. [3 marks]

(v) Demonstrate that there is no constant N such that $n > N$ implies $K(n) > \log\log\log\log n$. [3 marks]

9 Neural Computing

(*a*) Sketch a neural network that implements the following three operators for image analysis by demodulation, for purposes such as finding the important features in a face:

$$g(x,y) = \int_\alpha \int_\beta e^{-\left((x-\alpha)^2+(y-\beta)^2\right)/\sigma^2} \cos(\omega(x-\alpha))\ I(\alpha,\beta)\ d\alpha\ d\beta$$

$$h(x,y) = \int_\alpha \int_\beta e^{-\left((x-\alpha)^2+(y-\beta)^2\right)/\sigma^2} \sin(\omega(x-\alpha))\ I(\alpha,\beta)\ d\alpha\ d\beta$$

$$A^2(x,y) = g^2(x,y) + h^2(x,y)$$

where $I(x,y)$ is the image, ω is a parameter determining the approximate scale of the features being extracted, and $A^2(x,y)$ is the nonlinear output of the neural network that highlights the primary features. [10 marks]

(*b*) What is accomplished when such a network acts upon an image $I(x,y)$? In the case that the image happens to be a face, why does the output $A^2(x,y)$ detect the primary facial features? [4 marks]

(*c*) Describe the effect of your neural network in terms of the two-dimensional spatial frequency domain. To what image structure is it most sensitive, as a function of frequency, and as a function of orientation? [3 marks]

(*d*) Describe the neurobiological basis of your network in terms of the known properties of cells in the mammalian visual cortex. What is the name of the neurones that are described by the operators $g(x,y)$ and $h(x,y)$? What is the name of the class of neurones whose physiology is described by $A^2(x,y)$? [3 marks]

10 Comparative Architectures

(*a*) Many statically-scheduled super-scalar processors have long pipelines. Explain why this is so. [4 marks]

(*b*) What problems do such long pipelines introduce, and what techniques can be employed to reduce the penalty? [6 marks]

(*c*) Explain how register score boarding is used to prevent data hazards. [5 marks]

(*d*) What special difficulties does providing support for floating point over/underflow exceptions introduce? [5 marks]

11 Numerical Analysis II

(*a*) In the Chebyshev characterisation theorem, the best L_∞ approximation to $f(x)$ over a closed finite interval by a polynomial $p_{n-1}(x)$ of degree $n-1$ has the property that the maximum error $|e(x)|$ is attained at M distinct points ξ_j such that $e(\xi_j) = -e(\xi_{j-1})$. What is M? [2 marks]

(*b*) Let $x = m \times 2^k$ represent a normalised number in a floating-point implementation. When computing $\sqrt{x}$ show how the domain of the problem can be reduced to $x \in [1, 4)$. Find the coefficients a, b which minimise $||e(x)||_\infty$ over $[1, 4]$ where $e(x) = ax + b - \sqrt{x}$. [8 marks]

(*c*) Taking full account of symmetry, describe the form of the best polynomial approximation $p(x)$ to x^4 over $[-1, 1]$ by a polynomial of lower degree. Sketch x^4 and $p(x)$, showing the extreme values of $|e(x)|$ where $e(x) = x^4 - p(x)$. Hence compute the coefficients of $p(x)$. [10 marks]

12 Specification and Verification I

(*a*) Outline the steps involved in proving a specification $\{P\}C\{Q\}$ using the method of verification conditions. [6 marks]

(*b*) The familiar algorithm for generating verification conditions assumes that an annotation is added before a command C_2 in a sequence $C_1;C_2$ unless C_2 is an assignment. Extend this algorithm so that no annotation is required if C_2 is of the form `IF` B `THEN` $X_1{:=}E_1$ `ELSE` $X_2{:=}E_2$. [6 marks]

(*c*) Justify your extended algorithm by showing that if the verification conditions it generates from $\{P\}$ C; `IF` B `THEN` $X_1{:=}E_1$ `ELSE` $X_2{:=}E_2\{Q\}$ are provable, then $\vdash$ $\{P\}$ C; `IF` B `THEN` $X_1{:=}E_1$ `ELSE` $X_2{:=}E_2\{Q\}$. [8 marks]

13 Computer Vision

The following very useful operator is often applied to an image $I(x, y)$ in computer vision algorithms, to generate a related "image" $g(x, y)$:

$$g(x,y) = \int_\alpha \int_\beta \nabla^2 e^{-\left((x-\alpha)^2+(y-\beta)^2\right)/\sigma^2} I(\alpha,\beta)\ d\alpha\ d\beta$$

where

$$\nabla^2 = \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right)$$

(a) Give the general name for this type of mathematical operator, and the chief purpose that it serves in computer vision. [2 marks]

(b) What image properties should correspond to the zeroes of the equation, i.e. those points (x, y) in the image $I(x, y)$ where the above result $g(x, y) = 0$? [3 marks]

(c) What is the significance of the parameter σ? If you increased its value, would there be more or fewer points (x, y) at which $g(x, y) = 0$? [3 marks]

(d) Describe the effect of the above operator in terms of the two-dimensional Fourier domain. What is the Fourier terminology for this image-domain operator? What are its general effects as a function of frequency, and as a function of orientation? [4 marks]

(e) If the computation of $g(x, y)$ above were to be implemented entirely by Fourier methods, would the complexity of this computation be greater or less than the image-domain operation expressed above, and why? What would be the trade-offs involved? [4 marks]

(f) If the image $I(x, y)$ has 2D Fourier Transform $F(u, v)$, provide an expression for $G(u, v)$, the 2D Fourier Transform of the desired result $g(x, y)$ in terms of only the Fourier plane variables $u, v, F(u, v)$, and the above parameter σ. [4 marks]

14 Computer Systems Modelling

Consider an $M/G/1$ queue.

(a) What is the condition for the non-saturation of the server? [2 marks]

(b) State the *Pollaczek–Khintchine* formula for the steady-state average number of customers in the system. [4 marks]

(c) Suppose that Edward and Ursula are two applicants for a vacancy as a bank teller. It has been determined by empirical testing that Edward's service times are exponentially distributed with mean 0.9 minutes, while Ursula's are uniformly distributed between 0.8 and 1.2 minutes. It is known that customers arrive at the bank teller's window with inter-arrival times which are exponentially distributed and at the rate of 50 per hour.

(i) Can both applicants cope with the load?

(ii) If so, which one should be employed so as to minimise the steady-state average number of customers present?

In both cases, justify your answer. [8 marks]

(d) Given a sequence of pseudo-random numbers $U_1, U_2, \ldots$ distributed uniformly between 0 and 1 explain briefly how to construct pseudo-random sequences for the inter-arrival times of bank customers and for the service times of both Edward and Ursula. [3 marks]

(e) Briefly compare the advantages and disadvantages of the analytical queueing theory and the discrete event simulation approaches to determining performance measures by the above bank employer. [3 marks]

15 Topics in Concurrency

(a) Describe concisely a model checking algorithm for judgements of the form $p \models A$, where p is a finite-state process and A is an assertion of the modal μ-calculus. [4 marks]

(b) Show how to express a minimum fixed point assertion $\mu X.A$ in terms of a maximum fixed point assertion. [2 marks]

(c) Let $\mu X\{p_1, \cdots, p_n\}A$ mean $\mu X.(\neg\{p_1, \cdots, p_n\} \wedge A)$. From ($a$), or otherwise, show that:

(i) when $q \in \{p_1, \cdots, p_n\}$, the judgement $q \models \mu X\{p_1, \cdots, p_n\}A$ is false;

(ii) when $q \notin \{p_1, \cdots, p_n\}$,

$$q \models \mu X\{p_1, \cdots, p_n\}A \;\Leftrightarrow\; q \models A[\mu X\{q, p_1, \cdots, p_n\}A/X] \; .$$

[7 marks]

(d) From the algorithm you have described in (a), using (c) if it is helpful, decide whether or not the following judgement holds:

$$P \models \mu X.([a]F \vee \langle a\rangle X)$$

where P is the CCS process defined by

$$P \stackrel{\text{def}}{=} a.Q \qquad Q \stackrel{\text{def}}{=} a.P + a.\mathbf{nil} \; .$$ [7 marks]

END OF PAPER