COMPUTER SCIENCE TRIPOS Part II

Thursday 6 June 1996 1.30 to 4.30

Paper 9

Answer **five** *questions.*
Submit the answers in five **separate** *bundles each with its own cover sheet.*
Write on **one** *side of the paper only.*

1 VLSI

Give the transistor schematic and stick diagram for

(*a*) a static CMOS 2-input NOR gate [4 marks]

(*b*) a static CMOS 3-input NAND gate [6 marks]

(*c*) a 3-bit barrel shifter [10 marks]

2 Digital Communication II

CD quality (16 bits per sample, stereo, 22 KHz analog bandwidth) audio is to be sent over an ATM network. Assuming that cells are filled with samples what cell rate is required to carry this traffic? [4 marks]

What is the approximate packetization delay? [4 marks]

What is the effect of an undetected cell loss? [4 marks]

Design an adaptation layer to detect cell loss. What is the effect to the application of detected cell loss? [8 marks]

3 Computer System Modelling

The Erlangian distribution E_r, with parameters (r, μ) is given by

$$f_X(x) = \frac{\mu(\mu x)^{r-1}}{(r-1)!}\mathrm{e}^{-\mu x} \quad x \geqslant 0$$

(a) Given an $M/E_2/1$ queue, draw a Markov Chain describing the queueing system, and derive a formula for the variance of the service time distribution. Give an example of a queueing system in which the use of an Erlang service time distribution would be useful. [5 marks]

(b) Describe how random variables from a given distribution function $f_X(x)$ can be sampled for use in a discrete event simulator. [3 marks]

(c) Use your answer from (b) to develop pseudo-code for a function in a discrete event simulator which when called returns a sampled value from the distribution function E_r. State any assumptions made and explain any arguments to the function. Comment on the efficiency of your code. [12 marks]

4 Distributed Systems

Einstein has established that there is no universal time. For earth-based computer systems discuss how events might be assigned a time stamp which is reasonably close to conventional earth-time.

Describe the constraints on system-wide event ordering and discuss alternative approaches to meeting them. [10 marks]

For a system in which data replicas are maintained:

Either Define *total order* and *causal order* applied to the delivery of update messages to the replicas. Outline an approach to maintaining causal order. [10 marks]

or Describe algorithms and protocols which may be used to achieve consistency of replicas at all times. [10 marks]

5 Business Studies

What is meant by *SWOT analysis*? [5 marks]

A small computer company with strong and innovative hardware expertise is considering manufacturing a network interface computer (NIC). The device, which would sell for about half the current price of a PC, is based on games console technology, with a built-in modem. It would allow a user to convert his or her television to a web-browser. Apart from a small amount of parameter storage, the proposed device contains no disc or other long-term memory.

How would you determine the market for such a device? [5 marks]

Perform an analysis of this opportunity. What advice would you give the company? [5 marks]

Comment on changes to the business model that may be expected to be caused by the rapid development of the Internet. [5 marks]

6 Advanced Algorithms

Explain the steps involved in using the Miller–Rabin test to check whether a number N is composite. This will involve computing $a^{N-1} \bmod N$ for some value of a. [10 marks]

Carry out the steps for $N = 65$ and $a = 1, 2, 8$ and 12. Comment on what (if anything) each partial result tells you about N and which cases (if any) help you to decide whether N is prime or what its factors might be.

Pretend throughout the calculation that you do not know that $65 = 5 \times 13$. Proceed as though 65 were a huge number, imagining that you do not know at the outset whether it is prime or composite and that you are certainly unable to spot any factors. [10 marks]

7 Optimising Compilers

Briefly summarize the main concepts of strictness analysis including the kind of languages to which it applies, and the way in which both system-provided and user-defined functions f yield strictness properties $f^{\#}$ (relate the types of f and $f^{\#}$). [6 marks]

Give the strictness functions corresponding to the following ternary functions:

(a) `f1(x,y,z) = x*y + z`

(b) `f2(x,y,z) = if x=9 then y else z`

(c) `f3(x,y,z) = pif x=9 then y else z`

where `pif` e_1 `then` e_2 `else` e_3 is the *parallel conditional*: it behaves similarly to the standard conditional in that if e_1 evaluates to `true` or `false` then it yields e_2 or e_3 as appropriate; however, evaluation of e_2 and e_3 occurs concurrently with e_1 to allow the `pif` construct also to terminate with the value of e_2 when e_2 and e_3 both terminate with equal values (even if e_1 computes forever).

Comment briefly how your strictness property for `f1` would change if the multiplication returned zero without evaluating the other argument in the event that one argument were zero. [7 marks]

Let g, h_1 and h_2 be binary functions and recall the definition of function composition:

$$g \circ \langle h_1, h_2 \rangle = \lambda(x, y).g(h_1(x, y), h_2(x, y)).$$

Define three such functions in an ML-like syntax (whose arguments and results are integers) and which have the property that

$$(g \circ \langle h_1, h_2 \rangle)^{\#} \neq g^{\#} \circ \langle h_1^{\#}, h_2^{\#} \rangle.$$

[Hint: you might find it helpful to think of a solution where g *may* ignore one of its arguments but *always does* when composed with $\langle h_1, h_2 \rangle$.] Comment whether this inequality means that $g^{\#} \circ \langle h_1^{\#}, h_2^{\#} \rangle$ fails to be a *safe* strictness property for $g \circ \langle h_1, h_2 \rangle$. [7 marks]

8 Computational Neuroscience

It has been remarked that "neural networks are the second best way of computing just about anything." Discuss this, touching on the following issues: expressiveness; computational efficiency; generalization; sensitivity to noise; transparency (the ability to explain why a given output value is justified); the use of prior knowledge; whether neural networks fulfill our needs for a comprehensive computational theory of learning. [20 marks]

9 Security

Shamir's three-pass protocol enables Alice to send a message m to Bob in the following way:

$$\begin{array}{ll} A \rightarrow B : m^{ka} & (\text{mod } p) \\ B \rightarrow A : m^{ka\,kb} & (\text{mod } p) \\ A \rightarrow B : m^{kb} & (\text{mod } p) \end{array}$$

Explain this protocol, stating the constraint on m and the principal vulnerability. [10 marks]

It is suggested that the encryption operation $m \rightarrow m^{kx}$ be replaced with a provably secure encryption operation, namely a one-time pad. How would this affect the protocol's security? [10 marks]

10 Natural Language Processing

Describe three significant differences between programming languages and natural languages. [8 marks]

What problems do these differences pose for attempts to construct programs that "understand" a natural language? [12 marks]

11 Information Theory and Coding

Consider a noiseless analog communication channel whose bandwidth is 10,000 Hz. A signal of duration 1 second is received over such a channel. We wish to represent this continuous signal exactly, at all points in its one-second duration, using just a finite list of real numbers obtained by sampling the values of the signal at discrete, periodic points in time. What is the length of the shortest list of such discrete samples required in order to guarantee that we capture all of the information in the signal and can recover it exactly from this list of samples? [5 marks]

Name, define algebraically, and sketch a plot of the function you would need to use in order to recover completely the continuous signal transmitted, using just such a finite list of discrete periodic samples of it. [5 marks]

Consider a noisy analog communication channel of bandwidth Ω, which is perturbed by additive white Gaussian noise whose power spectral density is N_0. Continuous signals are transmitted across such a channel, with average transmitted power P (defined by their expected variance). What is the *channel capacity*, in bits per second, of such a channel? [10 marks]

12 Computer Vision

Using appropriate mathematical expressions, define the following operations commonly used in computer vision and briefly explain their function and applications:

(a) convolution [4 marks]

(b) correlation [4 marks]

(c) bandpass filtering [4 marks]

(d) edge detection by second-derivative zero-crossings [4 marks]

(e) invariant transform [4 marks]

13 Types

Briefly explain what is meant by *capture-avoiding substitution.* [3 marks]

What is a *principal typing* and why is it useful? [5 marks]

Suppose that a constant *fix* is added to the expressions of System F, with the typing rule

$$\Gamma \vdash \mathit{fix} \in \mathit{All}(X)\,(X \to X) \to X \qquad \text{(T-Fix)}$$

and principal evaluation rule:

$$\frac{f\ (\mathit{fix}\ X\ f) \Downarrow r}{\mathit{fix}\ X\ f \Downarrow r} \qquad \text{(E-Fix)}$$

Also, suppose we are given a built-in type operator *List* and the following expression constants:

$$\begin{array}{rcl} \mathit{nil} & \in & \mathit{All}(X)\ (\mathit{List}\ X) \\ \mathit{cons} & \in & \mathit{All}(X)\ X \to (\mathit{List}\ X) \to (\mathit{List}\ X) \\ \mathit{car} & \in & \mathit{All}(X)\ (\mathit{List}\ X) \to X \\ \mathit{cdr} & \in & \mathit{All}(X)\ (\mathit{List}\ X) \to (\mathit{List}\ X) \\ \mathit{null} & \in & \mathit{All}(X)\ (\mathit{List}\ X) \to \mathit{Bool} \end{array}$$

Use these primitives to write a polymorphic function *fold* of type

$$\mathit{fold} \in \mathit{All}(X)\ \mathit{All}(Y)\ (X \to Y \to Y) \to Y \to (\mathit{List}\ X) \to Y$$

that "folds a function across a list." For example, applying *fold* to $+$, 0, and a list of numbers should return the sum of the list. [8 marks]

Which of the following existential packages is most useful, and why?

$$[\mathit{Int}, \{x = 5,\ f = \mathit{fun}(i \in \mathit{Int})\ i + 1\}] \in \mathit{Some}(X)\ \{x \in X,\ f \in X \to X\}$$
$$[\mathit{Int}, \{x = 5,\ f = \mathit{fun}(i \in \mathit{Int})\ i + 1\}] \in \mathit{Some}(X)\ \{x \in \mathit{Int},\ f \in X \to \mathit{Int}\}$$
$$[\mathit{Int}, \{x = 5,\ f = \mathit{fun}(i \in \mathit{Int})\ i + 1\}] \in \mathit{Some}(X)\ \{x \in X,\ f \in X \to \mathit{Int}\}$$
$$[\mathit{Int}, \{x = 5,\ f = \mathit{fun}(i \in \mathit{Int})\ i + 1\}] \in \mathit{Some}(X)\ \{x \in \mathit{Int},\ f \in \mathit{Int} \to \mathit{Int}\}$$

[4 marks]

14 Numerical Analysis II

Let

$$p_n(x) = a_n x^n + \cdots + a_1 x + a_0$$

have n_+ positive real roots. If Descartes' rule of signs is expressed in the form

$$0 \leqslant c - n_+ = 2k,$$

what do c and k represent? [4 marks]

How many *positive* real roots do the following polynomials have?

(*a*) $x^6 - x^4 - x - 2$

(*b*) $x^4 - 2x^3 - 235x^2 - 940x + 10200$, given that $x = 5$ is a root

How many *negative* real roots does the following polynomial have?

(*c*) $2x^3 - 53x^2 + 316x + 600$ [4 marks]

Given that the polynomial

$$343x^3 - 294x^2 + 32$$

has a double root, find all of its roots. [6 marks]

Müller's method uses the formula

$$x_{i+1} = x_i - \frac{2f(x_i)}{c_i \pm \sqrt{{c_i}^2 - 4b_i f(x_i)}}.$$

What is the advantage of having a square root in the formula? How is the sign chosen in the denominator? Describe briefly the idea underlying Müller's method (omitting algebraic details) and comment on the choice of starting values.
[6 marks]

15 Pi Calculus

Define the notion of a *sorting* over a set $\mathcal{S}$ of subject sorts in the π calculus. Given a process P and a sorting *ob* over $\mathcal{S}$, explain the assertion that P *respects ob*. [6 marks]

Let $\mathcal{S} = \{A, B, C\}$ with $a : A$, $b, y : B$ and $c, z : C$.

Let $P = (\nu a)\big(a(y, z).\overline{z}\langle y\rangle \mid c(b).\overline{a}\langle b, c\rangle\big)$.

Show that P respects many different sortings over $\mathcal{S}$, and describe them.

On the other hand, let $\mathcal{S}$ contain at most two subject sorts. In this case, show that there are exactly two sortings over $\mathcal{S}$ which are respected by P. [7 marks]

Explain how recursive definition of processes in the π calculus can be represented in terms of replication. Would this be possible even in the monadic π calculus? [7 marks]