

# COMPUTER SCIENCE TRIPOS Part II

---

Thursday 8 June 2000 1.30 to 4.30

---

Paper 9

*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.*

*Write on **one** side of the paper only.*

## 1 Distributed Systems

In a distributed electronic conference application each participant has a replica of a shared whiteboard. Only one user at a time may write to the whiteboard, after which the user's update is propagated to all conference members. A dedicated process manages each whiteboard replica. Define and discuss a protocol that the replica managers could use to achieve this mutual exclusion. [10 marks]

A service database is replicated within a widely distributed system to achieve high availability and low access latency. Within a hierarchy of replicas a set of top-level primary servers maintain strong consistency of their replicas. Describe how this strong consistency could be achieved. [10 marks]

## 2 VLSI Design

Sketch a transistor-level design and give a brief description of the method of operation for each of the following memory cells:

(a) a dual-ported, static memory cell [5 marks]

(b) a dynamic memory cell using standard MOS logic [5 marks]

(c) a dynamic memory cell for dense memory [5 marks]

Explain the comparative merits of the three designs and explain where they might be used. [5 marks]

### 3 Digital Communication II

Consider the IntServ (RSVP based) and DiffServ QoS Architectures.

- (a) Compare the nature of “guarantee” that is made in each and the time scale on which these are altered. [6 marks]
- (b) Compare where state resides in each case, and how much state there is. [3 marks]
- (c) Compare the number of control interactions that a core router in each architecture would engage in. [3 marks]
- (d) Compare the resource efficiency of each scheme in terms of link bandwidth. [Hint: consider single-ended Service-Level Agreements (SLAs).] [3 marks]
- (e) Describe the architecture of a core router in a DiffServ environment, recalling that each packet will have a Type of Service field describing its designated Per Hop Behaviour. [5 marks]

### 4 Advanced Graphics and HCI

- (a) Show how you would calculate the first intersection point between an arbitrary ray and a finite length open cylinder of unit radius aligned along the  $x$ -axis. [Note: an open cylinder is one which has no end caps.]

Having calculated the intersection point, how would you calculate the normal vector? [7 marks]

- (b) A non-rational B-spline has knot vector  $[1, 2, 4, 7, 8, 10, 12]$ . Derive the first of the third order (second degree) basis functions,  $N_{1,3}(t)$ , and graph it.

If this knot vector were used to draw a third order B-spline, how many control points would be required? [7 marks]

- (c) 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]

## 5 Business Studies

What is meant by *a critical path*? [5 marks]

The village bakery has asked you to advise them about setting up a web site, including a trading function.

Draw up a project plan, illustrated by a GANTT chart, and indicate the critical path. [5 marks]

Make an estimate of the costs involved, and estimate how much working capital you would need. [5 marks]

What other advice would you give them? [5 marks]

## 6 Security

The owner of a banking system which previously used manually distributed shared keys to compute MACs on transactions decides to use public key cryptography to distribute MAC keys in future. The proposed protocol is

$$A \rightarrow B : \{ \{T_A, K_{AB}\}_{K_A^{-1}} \}_{K_B}$$

Explain the symbolism used in this description. [2 marks]

What is wrong with this protocol? [6 marks]

The protocol is changed to

$$A \rightarrow B : \{ \{A, T_A, K_{AB}\}_{K_A^{-1}} \}_{K_B}$$

What attacks might there be on the system now? [12 marks]

## 7 Computer System Modelling

A database system has a central processor and three (different) discs.

Measurements are taken for 1000 transactions on a lightly loaded system and the following observations are made.

- The CPU scheduler initiated or resumed transaction processing 10,000 times. The total CPU usage was 25 seconds.
- Disc 1 made 5000 transfers with an average transfer time of 10 ms.
- Disc 2 made 2000 transfers with an average transfer time of 50 ms.
- Disc 3 made 2000 transfers with an average transfer time of 20 ms.

Derive the visit counts, service times and transaction service demands. What is the bottleneck device? What is the maximum throughput of the system measured in transactions per second? [6 marks]

Describe *two* balanced systems which bound the throughput of the system. What is the maximum throughput of these systems? [7 marks]

Recall that the throughput of a balanced system with  $K$  devices,  $N$  customers and service demand  $D$  per device is

$$X(N) = \frac{N}{(N + K - 1)} \times \frac{1}{D}$$

How many transactions do you expect to be in the system with a throughput of 7 transactions per second? [7 marks]

## 8 Neural Computing

Give evidence supporting the view that the main computational load that has shaped the evolution of the human brain is “social computation”, with sexual success being the ultimate measure of the value of an algorithm or neural design feature. Say what implications this has for:

- The cognitive skills and perceptual faculties that have been selected for in brain evolution, as contrasted with the goals which are the traditional focus of AI.
- The design of face recognition algorithms, which aim to interpret facial expression, gesture, and intent, as well as gender and identity.
- The construction of the theory that other persons have minds, too.
- Models of action, planning, and interaction between self and others.

[8 marks]

Comment on whether this “social computation” view of human brain evolution implies that brain science is less relevant to the goals of computer science than is usually thought.

[2 marks]

Answer any *five* of the following seven short questions:

- (a) Roughly what is the total number of neurones in the human brain?
- (b) Roughly what is the total number of synapses in the human brain? How does this compare with the total number of stars in our galaxy, and with the total number of galaxies in the known universe?
- (c) Why is nerve impulse propagation described as “saltatory”, and what purposes are achieved by this method of signalling?
- (d) What is the approximate speed of nerve impulse propagation in warm-blooded animals, in metres/sec?
- (e) Why is “white matter” white, what cells are responsible for this, and what purpose do they serve?
- (f) Name the three principal ions involved in nerve membrane current flows, and identify which two of them transit through voltage-dependent conductances.
- (g) What causes the refractory deadtime of about 1 msec after each nerve impulse?

[2 marks each]

## 9 Artificial Intelligence

Consider the following story of the play Macbeth, by William Shakespeare:

The characters are Macbeth, Lady-Macbeth, Duncan and Macduff. Macbeth is an evil noble. Lady-Macbeth is a greedy ambitious woman. Duncan is a king. Macduff is a loyal noble. Macbeth is weak because Macbeth married Lady-Macbeth and because Lady-Macbeth is greedy. Lady-Macbeth persuades Macbeth to want to be king. Macbeth murders Duncan using a knife because Macbeth wants to be king and because Macbeth is evil. Lady-Macbeth kills Lady-Macbeth. Macduff is angry because Macbeth murdered Duncan and because Macduff is loyal to Duncan. Macduff kills Macbeth.

Construct a semantic network representing the above story. [8 marks]

Show the chain of reasoning leading to Macduff killing Macbeth. [5 marks]

It is possible to change the story so that Lady-Macbeth is unable to persuade Macbeth to want to be king. Augment the story to provide a reason for Lady-Macbeth's inability to persuade Macbeth to want to be king. Update the semantic network to reflect the new situation. [7 marks]

## 10 Numerical Analysis II

Explain the terms (a) *positive definite*, (b) *positive semi-definite* for a symmetric matrix  $\mathbf{A}$ . If a square matrix  $\mathbf{B}$  is non-singular, which of the properties (a) or (b) most accurately describes  $\mathbf{B}^T \mathbf{B}$ ? What if  $\mathbf{B}$  is singular? [4 marks]

State *Schwarz's inequality* for the product  $\mathbf{A}\mathbf{B}$ . In what way is this modified for the product  $\mathbf{A}\mathbf{x}$ , where  $\mathbf{x}$  is a vector? What are the *singular values* of  $\mathbf{A}$ , and how are they related to the  $l_2$  norm of  $\mathbf{A}$ ? In the *singular value decomposition*  $\mathbf{A} = \mathbf{U}\mathbf{W}\mathbf{V}^T$ , what is  $\mathbf{W}$ ? [5 marks]

Let  $\hat{\mathbf{x}}$  be an approximate solution of  $\mathbf{A}\mathbf{x} = \mathbf{b}$ , and write  $\mathbf{r} = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}$ ,  $\mathbf{e} = \mathbf{x} - \hat{\mathbf{x}}$ . Find an expression which is an upper bound for the relative error  $\|\mathbf{e}\|/\|\mathbf{x}\|$  in terms of computable quantities. Explain how this result may be interpreted if the  $l_2$  norm is used. [8 marks]

Suppose  $\mathbf{A}$  is a  $5 \times 5$  matrix and  $\mathbf{A}\mathbf{x} = \mathbf{b}$  is to be solved by singular value decomposition. If *machine epsilon*  $\simeq 10^{-15}$  and the singular values of  $\mathbf{A}$  are  $1, 10^{-6}, 10^{-10}, 10^{-17}, 0$  write down the generalised inverse  $\mathbf{W}^+$  that you would use. [3 marks]

## 11 Information Theory and Coding

- (a) Prove that the information measure is additive: that the information gained from observing the combination of  $N$  independent events, whose probabilities are  $p_i$  for  $i = 1 \dots N$ , is the *sum* of the information gained from observing each one of these events separately and in any order. [4 marks]
- (b) What is the shortest possible code length, in bits per average symbol, that could be achieved for a six-letter alphabet whose symbols have the following probability distribution?

$$\left\{ \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}, \frac{1}{32} \right\}$$

[3 marks]

- (c) Suppose that ravens are black with probability 0.6, that they are male with probability 0.5 and female with probability 0.5, but that male ravens are 3 times more likely to be black than are female ravens.

If you see a non-black raven, what is the probability that it is male? [4 marks]

How many bits worth of information are contained in a report that a non-black raven is male? [1 mark]

Rank-order for this problem, from greatest to least, the following uncertainties:

- (i) uncertainty about colour
- (ii) uncertainty about gender
- (iii) uncertainty about colour, given only that a raven is male
- (iv) uncertainty about gender, given only that a raven is non-black

[3 marks]

- (d) If a continuous signal  $f(t)$  is *modulated* by multiplying it with a complex exponential wave  $\exp(i\omega t)$  whose frequency is  $\omega$ , what happens to the Fourier spectrum of the signal?

Name a very important practical application of this principle, and explain why modulation is a useful operation.

How can the original Fourier spectrum later be recovered? [3 marks]

- (e) Which part of the 2D Fourier Transform of an image, the amplitude spectrum or the phase spectrum, is indispensable in order for the image to be intelligible?

Describe a demonstration that proves this. [2 marks]

## 12 Computer Vision

Define the *Correspondence Problem*, detailing the different forms that it takes in stereo vision and in motion vision.

- (a) In each case, explain why the computation is necessary. [5 marks]
- (b) What are the roles of space and time in the two cases, and what symmetries exist between the stereo vision and the motion vision versions of the Correspondence Problem? [5 marks]
- (c) How does the complexity of the computation depend on the number of underlying features that constitute the data? [5 marks]
- (d) Briefly describe at least one general approach to an efficient algorithm for solving the Correspondence Problem. [5 marks]



### 13 Types

Give the axiom and rules of the type system for polymorphic lambda calculus (PLC). [6 marks]

Given any function  $\rho$  mapping type variables  $\alpha$  to boolean values  $\mathbf{b} \in \{\mathbf{true}, \mathbf{false}\}$ , we extend  $\rho$  to a function on all PLC types by defining

$$\begin{aligned}\rho(\tau \rightarrow \tau') &= \rho(\tau) \Rightarrow \rho(\tau') \\ \rho(\forall\alpha(\tau)) &= \rho[\alpha \mapsto \mathbf{true}](\tau) \ \& \ \rho[\alpha \mapsto \mathbf{false}](\tau)\end{aligned}$$

where  $\Rightarrow$  and  $\&$  are the usual boolean operations of implication and conjunction, and where  $\rho[\alpha \mapsto \mathbf{b}]$  is the function mapping  $\alpha$  to  $\mathbf{b}$  and otherwise acting like  $\rho$ . For example, show that for any  $\rho$  we have  $\rho(\forall\alpha(\alpha \rightarrow \alpha)) = \mathbf{true}$  and  $\rho(\forall\alpha(\alpha)) = \mathbf{false}$ . [2 marks]

Let  $\Phi(\Gamma, M, \tau)$  be the following property of PLC typing judgements  $\Gamma \vdash M : \tau$ .

“For all  $\rho$ , if  $\rho(\tau) = \mathbf{false}$  then  $\Gamma$  contains a type assignment  $x_i : \tau_i$  with  $\rho(\tau_i) = \mathbf{false}$ .”

Show that  $\Phi(\Gamma, M, \tau)$  is closed under the axiom and rules of the PLC type system. (You may assume without proof that if  $\alpha$  is not free in  $\tau$  then  $\rho[\alpha \mapsto \mathbf{b}](\tau) = \rho(\tau)$ ; and also that type substitutions  $\tau'[\tau/\alpha]$  satisfy  $\rho(\tau'[\tau/\alpha]) = \rho[\alpha \mapsto \rho(\tau)](\tau')$ .) [10 marks]

Deduce that there is no *closed* PLC expression of type  $\forall\alpha(\alpha)$ . [2 marks]

**END OF PAPER**