

COMPUTER SCIENCE TRIPOS Part II

Thursday 4 June 1998 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

Discuss capability-based access control under the headings

- protection of capabilities
- control of transfer of capabilities between principals
- delegation of rights
- revocation of capabilities

Your answer should mention the differences between the management of capabilities in distributed as opposed to centralised systems. You should consider alternative designs of capabilities. [20 marks]

2 Digital Communication II

What is a *leaky bucket algorithm*? How can it be used to police traffic? How can it be used to shape traffic? [5 marks]

You are required to provide a flow admission controller for a router. RSVP is used to request resources for flows with leaky bucket parameters. Outline how you would design the controller, bearing in mind the soft state nature of RSVP. [15 marks]

3 Computer System Modelling

An M/M/m queue has an arrival process with mean rate λ , and processes customers at a mean rate of μ .

- (a) What are the distributions and parameters of the inter-arrival and service times of customers? [3 marks]
- (b) Sketch an outline proof showing that the distribution of the departure process from the queue is the same as that of the arrivals process. [10 marks]

Briefly contrast analytical queueing analysis and discrete event simulation with regard to their fields of applicability and other important considerations for the systems modeller. [7 marks]

4 VLSI

Write short notes on *two* of the following:

- (a) the fundamental limits which may slow down progress in semiconductor technology;
- (b) designing VLSI systems for low-power applications;
- (c) problems which can prevent an apparently properly designed chip from working to specification.

[10 marks each]

5 Business Studies

You are assigned the role of UK sales and marketing manager for a new kind of low-cost computer, primarily aimed at the educational market. Whilst not directly PC compatible, the computer includes web access, PC-compatible word processing, other PC-compatible productivity tools, and a suite of educational programs.

How would you approach this task? Draw up an outline business plan as follows.

- (a) Show what communication and distribution channels you propose. [5 marks]
- (b) Propose a selling price, and estimate the number of units you might sell at this price. [5 marks]
- (c) Estimate a 3-year budget for the sales and marketing activity. [5 marks]
- (d) State how you would refine your estimates, and what monitoring you would put in place. [5 marks]

Background information:

- There are about 32,000 schools in the UK.
- The UK government has recently published a consultative document *The National Grid for Learning* with a proposal to spend £100M on provision of internet access in schools over the next 5 years. This sum includes infrastructure provision, content development and teacher training, as well as a contribution to provision of computers in schools.
- The average school IT spend is projected to be £18,000 each year. Additional funding may be available from government and parents for specific projects.
- The unit manufacturing cost is £200, delivered.

6 Advanced Algorithms

Describe the structure of an ordinary heap, and document the costs associated with the following operations.

- (a) Create a heap from n items where the items are all available at once but are initially in a random order.
- (b) Remove the top (i.e. smallest) value stored in the heap.
- (c) Given a pointer to an arbitrary item in the heap, re-instate the heap property after the key associated with that single item is decreased in value.
- (d) Form a new heap whose elements are all those that are present in two other heaps (which may be destroyed in the combining process if that helps).

You are not expected to give detailed accounts of the algorithms involved.

[6 marks]

Now explain the structure of a Binomial Heap and compare, with some explanation of your claims, the costs incurred in the same set of operations if Binomial rather than ordinary heaps were to be used.

[14 marks]

7 Optimising Compilers

Consider the programming language with terms e having abstract syntax:

$$e ::= x \mid c \mid \lambda x.e \mid e_1 e_2 \mid \text{let } x = e_1 \text{ in } e_2$$

where x ranges over a set of identifiers and c over a set of integer constants. For the rest of the question, your answers can be illustrated by reference to the program p :

$$\lambda z.\text{let } id = \lambda x.x \text{ in } id \ id \ 7$$

State how to label terms in p uniquely so that a subterm occurring repeatedly in a term has different labels. [4 marks]

Show how such terms may be seen as a family of flowgraphs, one for each λ (you may find it useful to consider the above labelling as providing a unique function name for anonymous λ -abstractions). [4 marks]

Define the *call graph* of such a family of flowgraphs, stating clearly how indirect calls are treated. [4 marks]

Describe how to associate a flow-variable with each labelled node of a term such as p and to derive equations which can improve the above treatment of indirect calls to get a better approximation of the edges in the call graph. [8 marks]

[Hint: you may find it useful to recall the shorthand of $(\gamma \mapsto \delta) \supseteq \beta$ as representing the compound constraint that

$$\text{whenever } (\lambda x^j.e^k)^i \in \beta \text{ we have } \alpha_j \supseteq \gamma \wedge \delta \supseteq \alpha_k$$

where α_r is the flow variable associated with the node labelled r .]

8 Neural Computing

When using a feed-forward network to solve a classification problem we can interpret the network's outputs as posterior probabilities of class membership, and then subsequently use these probabilities to make classification decisions. Alternatively, we can treat the network as a discriminant function which is used to make the classification decision directly. Discuss the relative merits of these two approaches. [7 marks]

Explain the concept of a likelihood function, and the principle of maximum likelihood. [3 marks]

Consider a feed-forward network which implements a function $y(\mathbf{x}, \mathbf{w})$ in which y is the output variable, \mathbf{x} is the vector of input variables, and \mathbf{w} is the vector of weight parameters. We wish to use this network to solve a classification problem involving two classes \mathcal{A} and \mathcal{B} . The value of y , when the network is presented with an input vector \mathbf{x} , is to be interpreted as the posterior probability $P(t = 1|\mathbf{x})$ in which $t = 1$ denotes class \mathcal{A} and $t = 0$ denotes class \mathcal{B} . Write down the probability distribution of t given y . Use the principle of maximum likelihood to derive an expression for the corresponding error function defined over a set of training data comprising input vectors \mathbf{x}_n and targets t_n , where $n = 1, \dots, N$.

Write down a suitable form for the output unit activation function $y = g(a)$. Hence evaluate the derivative of $\ln P(t|y)$ with respect to a . [10 marks]

9 Security

Describe those provisions of the following Acts of Parliament that are relevant to computer security:

- The Civil Evidence Act of 1968
- The Police and Criminal Evidence Act of 1984
- The Data Protection Act of 1984
- The Computer Misuse Act of 1990

[12 marks]

A software house incorporates a time-lock which causes their product to stop working if it is not supplied with a suitable password every 6 months. What risks are they running? [4 marks]

A hospital de-identifies patient records by removing names and addresses, leaving only the patient's postcode and date of birth as an identifier. These records are then sold to researchers and drug companies. What risk is the hospital running? [4 marks]

10 Denotational Semantics

State, with justification, whether each of the following statements is true or false.

(a) The set of natural numbers, $\mathbb{N} = \{0, 1, 2, \dots\}$, equipped with the usual less-than-or-equal relation, \leq , is a domain. [3 marks]

(b) The set of all subsets of \mathbb{N} , equipped with the relation of subset inclusion, is a domain. [4 marks]

(c) For any domain D and element $d \in D$ with $d \neq \perp$

$$f_d(x) = \begin{cases} \top & \text{if } d \sqsubseteq x, \\ \perp & \text{otherwise} \end{cases}$$

defines a strict continuous function f_d from D to the flat domain $\{\top\}_\perp$. [4 marks]

(d) For any domain D and element $d \in D$ with $d \neq \perp$

$$g_d(x) = \begin{cases} \perp & \text{if } x \sqsubseteq d, \\ \top & \text{otherwise} \end{cases}$$

defines a strict continuous function g_d from D to $\{\top\}_\perp$. [4 marks]

(e) For any continuous functions $h : D \rightarrow E$ and $k : E \rightarrow D$ between domains D and E , $\text{fix}(k \circ h) \sqsubseteq k(\text{fix}(h \circ k))$. [5 marks]

11 Information Theory and Coding

Consider Shannon's third theorem, the *Channel Capacity Theorem*, for a continuous communication channel having bandwidth W Hertz, perturbed by additive white Gaussian noise of power spectral density N_0 , and average transmitted power P .

(a) Is there any limit to the capacity of such a channel if its signal-to-noise ratio $\frac{P}{N_0W}$ can be arbitrarily increased? If so, what is that limit? [2 marks]

(b) Is there any limit to the capacity of such a channel if, leaving N_0 and P unchanged, its bandwidth W in Hertz can be arbitrarily increased? If so, what is that limit? [3 marks]

[Hint: $\ln(1+x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} + \dots$ where $-1 < x < 1$.]

Explain why smoothing a signal, by low-pass filtering it *before* sampling it, can prevent aliasing. Draw a graph in the Fourier domain which illustrates the origin of aliasing, and also show in the picture how smoothing solves the problem. What would be the most effective low-pass filter to use for this purpose? Draw it in the Fourier domain. [10 marks]

Suppose that women who live beyond the age of 70 outnumber men in the same age bracket by three to one. How much information, in bits, is gained by learning that a certain person who lives beyond 70 happens to be male? [5 marks]

12 Computer Vision

It could be said that the central problem of pattern recognition is the relation between the within-class variability and the between-class variability for the patterns that one would like to recognise. Explain this problem in the case of face recognition, treating separately the problems of

- (a) *face detection* (distinguishing faces from non-faces)
- (b) *face identification*
- (c) *face interpretation* (classifying the expression and pose angle of the face)

How do the forms of variability for faces influence each of the three tasks? Is within-class variability ever helpful, and between-class variability ever harmful, to the performance of the task? What role can statistical decision-theory play in formalising and solving these problems?

[20 marks]

13 Types

What is meant by the terms *safety* and *type soundness* in connection with programming languages and their type systems? [2 marks]

Explain how a naïve combination of ML-style polymorphism and references results in an unsound type system. Your account should include the axioms and rules of the type system, but need not give a formal definition of the operational semantics. [15 marks]

State, without proof, how the rule for typing **let**-expressions can be restricted to restore type soundness. [3 marks]

14 Numerical Analysis II

Explain what is meant by *local error* and *global error* in methods for the solution of *ordinary differential equations (ODEs)*. If a typical method has local error $O(h^3)$, what would you expect the global error to be? [3 marks]

Euler's method for solution of $y' = f(x, y)$ can be expressed as $y_{n+1} = y_n + k_1$. From the Taylor series, find an expression for k_1 . [2 marks]

The Runge–Kutta method RK2 is

$$y_{n+1} = y_n + \frac{1}{2}(k_1 + k_2)$$

where k_1 is the increment used by Euler's method, and

$$k_2 = h f(x_n + h, y_n + k_1).$$

In terms of Euler's method, what does the quantity k_2 represent? [2 marks]

Assume that RK2 is carried out with step sizes h and $h/2$, and that

$$y_{(h)}(x_{n+1}) = y(x_{n+1}) + C_n h^2 + O(h^3).$$

Derive an estimate of the error $E_n = |C_n|(h/2)^2$ in $y_{(h/2)}(x_{n+1})$. [3 marks]

Let ε be the *target error per unit step*. Why, in *step-size control* for RK2, is $\varepsilon' = \varepsilon/8$ taken as the target error corresponding to half the step size? [2 marks]

A certain ODE is to be solved using RK2 with step-size control. Using computed values for y from the table below, taking $\varepsilon = 0.005$, and starting with $h = 0.1$, state at which values of x you would make the *first* and *second* changes of step size, and what new values of h you would use in each case.

		h			
		0.025	0.05	0.1	0.2
x	0.05	0.10038	0.10050		
	0.1	0.20279	0.20304	0.20400	
	0.15	0.30946	0.30981		
	0.2	0.42295	0.42341	0.42516	0.43200
	0.25	0.54649	0.54702		
	0.3	0.68434	0.68490	0.68697	
	0.35	0.84247	0.84295		
	0.4	1.02971	1.02989	1.03047	1.03373
	0.45	1.25995	1.25930		
	0.5	1.55646	1.55379	1.54484	

[8 marks]

15 Communicating Automata and Pi Calculus

Explain briefly the role played by structural congruence in defining the reaction rules of the π -calculus. Give the structural congruence rules which involve restriction, or composition, or both. [7 marks]

A π -calculus term of the form $\text{new } \bar{z}(M_1 | \cdots | M_m | !Q_1 | \cdots | !Q_n)$, where $m, n \geq 0$ and each M_i is a summation, is said to be in *standard form*. Give an argument (which need not be fully formal) to show that every term of the π -calculus is structurally congruent to a standard form. [6 marks]

Consider the term

$$P = \bar{x}\langle a \rangle \mid \text{new } a (!\bar{a}\langle c \rangle \mid (a(b).b(c).\bar{c}\langle x \rangle + x(y).y(z).\bar{z}\langle a \rangle)) .$$

Convert P to standard form by the rules of structural congruence, indicating which rules are used. Hence write down all the possible reactions $P \rightarrow P'$, and the possible reactions of each resulting P' . For each result which has no further reactions, write down the simplest term to which it is strongly equivalent (\sim), giving a brief justification. [7 marks]