COMPUTER SCIENCE TRIPOS Part II

Thursday 3 June 2004 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.*

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 Human–Computer Interaction

(*a*) If you had been one of the original inventors of the WIMP interface, and engineers on the technical team had been sceptical about the advantages that it would bring, then:

(*i*) What *two* pieces of empirical evidence could you have presented to sceptics to help convince them? [2 marks]

(*ii*) Describe the technique that you would use to collect data for *each* of these. [4 marks]

(*iii*) Describe the technique that you would use to analyse *each* of those data sets. [4 marks]

(*b*) If you were the HCI researcher on a speech interaction project intended to replace the WIMP interface, then

(*i*) Which of the techniques described above would be most relevant to your work? Why? [3 marks]

(*ii*) Explain how a modern analytic evaluation method could give design guidance to the team, including a description of how the analysis would be conducted, and *two* design improvements that might result from that analysis. [7 marks]

2 VLSI Design

(*a*) Sketch designs for an n-input NAND gate in CMOS using

(*i*) static CMOS;

(*ii*) dynamic CMOS;

(*iii*) pseudo-nMOS. [2 marks each]

(*b*) Assuming that a conducting p-channel has a resistance γ times that of a similarly sized n-channel, annotate each of your circuit diagrams with suitable widths for the transistors, and explain the reasons for their values. [2 marks each]

(*c*) For each design, calculate the logical effort and the parasitic delay. [2 marks each]

(*d*) Which is likely to be fastest for the case when $n = 4$ and $\gamma = 3$? What difference would it make if the circuit were driving a large capacitative load? [2 marks]

3 Optimising Compilers

Assume that a program consists of a sequence of declarations `Object` o`;` where o is an object name, followed by a sequence of function definitions $f(x_1, \ldots, x_k) = e$ where expressions, e, have syntax

$$e ::= n \mid o \mid x \mid f(e_1, \ldots, e_k) \mid \texttt{let } x = e_1 \texttt{ in } e_2 \mid \texttt{if } e_1 \texttt{ then } e_2 \texttt{ else } e_3.$$

where n ranges over integer constants and x over variables (which may contain integers or object *references*). Variables may not contain function values.

Alias Analysis is a technique which will determine that, during evaluation of e within

$$\texttt{let } x = o \texttt{ in let } y = o \texttt{ in } e$$

x and y alias because they are both references to the same object o.

(*a*) Show how to associate a flow variable with each variable and (sub-)expression of the program. State the values which flow variables might reasonably take in such an analysis. [4 marks]

(*b*) Show how, given a program, we can generate a set of constraint-style equations (analogously to control-flow analysis for λ-expressions) whose solution gives a superset of the values which might be returned from each (sub-)expression of the program. [Hint: suppose that each function definition has flow variables representing the value ranges of each of its arguments and of its result.] [8 marks]

(*c*) Explain what happens in, and give modifications to part (*b*) for, the generalisation whereby variables can also reference functions and be called by the syntax

$$e ::= e_0(e_1, \ldots, e_k)$$

[4 marks]

(*d*) Explain how you would respond to the criticism that your analysis may fail to terminate if your language is extended with arithmetic expressions because a single expression may give rise to an infinite set of values. [2 marks]

(*e*) Briefly describe any optimisation whereby knowing that x and y cannot alias is necessary for the optimisation to be safe. [2 marks]

4 Distributed Systems

A network-based service manages persistent objects. The service must enforce an access control policy to protect the objects.

(*a*) Discuss how this access control might best be implemented for the following example of objects and policy components:

> Objects: Files in a University Department's file service, operating behind a firewall.
>
> Policy: The owner may specify read, write and execute rights in terms of principals and groups. [4 marks]

(*b*) Discuss how this access control might best be implemented for *two* of the following examples:

(*i*) Objects: Files in a commercial, distributed, Internet-based file service.
Policy: The owner may authorise other principals to download the file.

(*ii*) Objects: Sales data relating to a company.
Policy: Those employed in the Sales Departments of all branches of the company worldwide may read the data.

(*iii*) Objects: Electronic health records (EHRs) in a nationwide service.
Policy: The owner (patient) may read from its own EHR. A qualified and employed doctor may read and write the EHR of a patient registered with him/her.

(*iv*) Object: The solution to online coursework.
Policy: The coursework setter has read and write access. A candidate has no access until after the marks have been published.

[8 marks each]

5 Advanced Systems Topics

(*a*) Describe the DHT algorithms that are used in each of CAN and Pastry to route a request for content to the node that holds that content.
[5 marks each]

(*b*) CAN and Pastry include mechanisms to reduce the latency of search. Describe how each of these works, and discuss their relative complexity and performance.
[5 marks each]

6 Advanced Graphics

(*a*) Describe an algorithm to find the intersection point between an arbitrary ray and an arbitrary plane. [5 marks]

(*b*) Explain how a plane can be used as a Computational Solid Geometry (CSG) primitive. [2 marks]

(*c*) List the three binary operations used in CSG. Explain how a CSG object can be represented as a binary tree. Describe an algorithm to find the first intersection point between a ray and an arbitrary CSG object. Assume that there are already algorithms which you can use to find the intersection points between the ray and each type of CSG primitive. Ensure that you state any assumptions you make about the information provided to you by these ray-primitive intersection algorithms. [8 marks]

(*d*) Derive the NURBS basis function $N_{4,4}$ for the knot vector $[1, 2, 3, 4, 5, 5, 5, 6, 7, 8, 9, 10]$. [5 marks]

7 Digital Communication II

(*a*) Multicast Addressing and Routing provides a set of mechanisms for senders to transmit packets that are replicated by the routers so that they can be received by multiple systems. Explain how the basic mechanisms of *IGMP*, *reverse path forwarding* based on the underlying unicast routes, *pruning* and *grafting*, fit together to create this network service. [8 marks]

(*b*) How might IP multicast be a risk for a network provider? [2 marks]

(*c*) The Resource Reservation Protocol, RSVP, is a *receiver oriented* signalling protocol to establish state in routers for the purposes of classifying packets into flows and scheduling those flows onto routers. Explain what is meant by "receiver oriented", and how this enables RSVP to be used by a multicast (many-to-many) application. [5 marks]

(*d*) Why is TCP not going to work well with multicast? [3 marks]

(*e*) What is philosophically odd about using TCP with RSVP? [2 marks]

8 Artificial Intelligence

In the following, N is a feedforward neural network architecture taking a vector

$$\mathbf{x}^T = \begin{pmatrix} x_1 & x_2 & \cdots & x_n \end{pmatrix}$$

of n inputs. The complete collection of weights for the network is denoted $\mathbf{w}$ and the output produced by the network when applied to input $\mathbf{x}$ using weights $\mathbf{w}$ is denoted $N(\mathbf{w}, \mathbf{x})$. The number of outputs is arbitrary. We have a sequence $\mathbf{s}$ of m labelled training examples

$$\mathbf{s} = ((\mathbf{x}_1, \mathbf{l}_1), (\mathbf{x}_2, \mathbf{l}_2), \ldots, (\mathbf{x}_m, \mathbf{l}_m))$$

where the $\mathbf{l}_i$ denote vectors of desired outputs. Let $E(\mathbf{w}; (\mathbf{x}_i, \mathbf{l}_i))$ denote some measure of the error that N makes when applied to the ith labelled training example. Assuming that each node in the network computes a weighted summation of its inputs, followed by an activation function, such that the node j in the network computes a function

$$g\left(w_0^{(j)} + \sum_{i=1}^{k} w_i^{(j)} \text{input}(i)\right)$$

of its k inputs, where g is some activation function, derive in full the backpropagation algorithm for calculating the gradient

$$\frac{\partial E}{\partial \mathbf{w}} = \begin{pmatrix} \frac{\partial E}{\partial w_1} & \frac{\partial E}{\partial w_2} & \cdots & \frac{\partial E}{\partial w_W} \end{pmatrix}^T$$

for the ith labelled example, where $w_1, \ldots, w_W$ denotes the complete collection of W weights in the network.

[20 marks]

9 Database Theory

Consider a database with one binary relation B.

(*a*) Write programs in stratified Datalog for the following queries:

(*i*) Give the set of elements x for which there are fewer than three elements y such that $B(x, y)$. [3 marks]

(*ii*) Give the set of elements x such that there is no path (in the graph formed by B) from x back to itself. [3 marks]

(*iii*) Give the set of elements x such that $B(x, y)$ for every y. [3 marks]

(*b*) Which of the queries defined in part (*a*) is:

(*i*) safe? [3 marks]

(*ii*) domain independent? [4 marks]

(*iii*) monotone? [4 marks]

In each case, justify your answer fully.

10 Types

Let β be a type variable and let α range over type variables distinct from β. The subsets of polymorphic lambda calculus (PLC) types that are *positive* (ranged over by τ) and *negative* (ranged over by ν) in β are defined by the following grammar:

$$\begin{aligned} \tau &::= \forall\alpha(\tau) \mid \alpha \mid \beta \mid \nu \to \tau \\ \nu &::= \forall\alpha(\nu) \mid \alpha \mid \tau \to \nu \end{aligned}$$

(*a*) Give inductive definitions, following the structure of the grammar above, of closed PLC terms P_τ for each positive type τ, and N_ν for each negative type ν, such that

$$\begin{aligned} &\emptyset \vdash P_\tau : \forall\alpha_1, \alpha_2((\alpha_1 \to \alpha_2) \to (\tau[\alpha_1/\beta] \to \tau[\alpha_2/\beta])) \\ &\emptyset \vdash N_\nu : \forall\alpha_1, \alpha_2((\alpha_1 \to \alpha_2) \to (\nu[\alpha_2/\beta] \to \nu[\alpha_1/\beta])) \end{aligned}$$

[12 marks]

(*b*) Now let τ be the type $\forall\alpha((\beta \to \alpha) \to \alpha)$, which is positive in β. Calculate the beta-normal form of P_τ. [8 marks]

11 Computer Vision

(*a*) Why is the performance of current face recognition algorithms so poor? Address both what makes the problem domain intrinsically so challenging, and the shortcomings of the strategies adopted in the design of current algorithms. Comment on possible directions for improving performance. [8 marks]

(*b*) For what size of filter kernel does it become more efficient to perform convolutions by instead computing Fourier Transforms, and why? [2 marks]

(*c*) For an aligned stereo pair of cameras separated by base distance b, each with focal length f, when a target point projects outside the central axis of the two cameras by amounts α and β:

(*i*) What is the computed target depth d? [2 marks]

(*ii*) Why is camera calibration so important for stereo vision computations? [1 mark]

(*iii*) Identify *four* relevant camera degrees-of-freedom and briefly explain their importance for stereo vision algorithms. [2 marks]

(*d*) What does the Spectral Co-Planarity Theorem assert about translational visual motion, and how the parameters of such motion can be extracted? [2 marks]

(*e*) What information about the shape and orientation of an object can be inferred, and how, from the extraction of texture descriptors; and what is the role of prior assumptions in making such inferences? [3 marks]

12 Numerical Analysis II

(*a*) A Riemann integral over $[a, b]$ is defined by

$$\int_a^b f(x)\,dx = \lim_{\substack{n\to\infty \\ \Delta\xi\to 0}} \sum_{i=1}^{n} (\xi_i - \xi_{i-1}) f(x_i) \ .$$

Explain the terms *Riemann sum* and *mesh norm.* [4 marks]

(*b*) Consider the quadrature rule

$$Qf = \frac{3h}{8}[f(a) + 3f(a+h) + 3f(a+2h) + f(a+3h)] - \frac{3f^{(4)}(\lambda)h^5}{80} \ .$$

If $[a, b] = [-1, 1]$ find $\xi_0, \xi_1, \ldots, \xi_4$ and hence show that this is a Riemann sum. [3 marks]

(*c*) Suppose R is a rule that integrates constants exactly over $[-1, 1]$, and that $f(x)$ is bounded and Riemann-integrable over $[a, b]$. Write down a formula for the composite rule $(n \times R)f$ and prove that

$$\lim_{n\to\infty} (n \times R)f = \int_a^b f(x)\,dx$$ [6 marks]

(*d*) What is the formula for $(n \times Q)f$ over $[a, b]$? [4 marks]

(*e*) Which polynomials are integrated exactly by Qf? Which monomials are integrated exactly by the product rule $(Q \times Q)F$ when applied to a function of x and y? [3 marks]

13 Specification and Verification II

(*a*) Discuss when it is appropriate to use theorem proving and when it is appropriate to use model checking for formal verification. Give an example to illustrate when theorem proving would be used, and also one in which model checking would be best. [8 marks]

(*b*) Consider the device shown below:

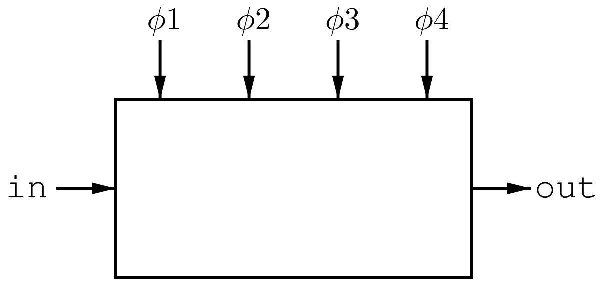

It is assumed that $\phi 1$, $\phi 2$, $\phi 3$, $\phi 4$ constitute a four-phase clock satisfying (*i*) and (*ii*), and that the device has the property (*iii*), where:

(*i*) At all times exactly one of $\phi 1$, $\phi 2$, $\phi 3$, $\phi 4$ is true.

(*ii*) If ϕi is true at time t then $\phi(\texttt{if } i < 4 \texttt{ then } i{+}1 \texttt{ else } 1)$ will be true at time $t{+}1$.

(*iii*) If $\phi 1$, $\phi 2$, $\phi 3$, $\phi 4$ satisfy (*i*) and (*ii*) above, and if $\phi 1$ is true at time t then the value at `out` at time $t{+}3$ will equal the value input at `in` at time $t{+}1$.

Express assumptions (*i*) and (*ii*) and property (*iii*) both in higher order logic and in temporal logic. [4 marks each]

14 Natural Language Processing

(*a*) Give brief definitions of the following terms:

(*i*) referring expression;

(*ii*) cataphora;

(*iii*) pleonastic pronoun.

[6 marks]

(*b*) Describe the Lappin and Leass algorithm for pronoun resolution, illustrating its operation on the text below. Exact weights for salience factors are *not* required.

> Owners love the new hybrid cars. They all say that they have much better fuel economy than conventional vehicles. And it seems that the performance of hybrid cars matches all expectations.

[14 marks]

15 Denotational Semantics

(*a*) Show that any continuous function $h : D \to D$ on a domain D has a least prefixed point $\mathit{fix}(h)$. [5 marks]

(*b*) Let $f : D \times E \to D$ and $g : D \times E \to E$ be continuous functions where D and E are domains. The continuous function $\langle f, g\rangle : D \times E \to D \times E$, acts so that $(d, e) \mapsto (f(d, e), g(d, e))$. Bekič's theorem states that the least fixed point of $\langle f, g\rangle$ is the pair $(\hat{d}, \hat{e})$ where

$$\begin{aligned} \hat{d} &= \mathit{fix}(\lambda d.\ f(d, \hat{e})) \text{ where} \\ \hat{e} &= \mathit{fix}(\lambda e.\ g(\mathit{fix}(\lambda d.\ f(d, e)), e))\ . \end{aligned}$$

You are asked to show Bekič's theorem in the following stages. Write (d_0, e_0) for the least fixed point of $\langle f, g\rangle$.

(*i*) Show that $(\hat{d}, \hat{e})$ is a fixed point of $\langle f, g\rangle$. Deduce that $(d_0, e_0) \sqsubseteq (\hat{d}, \hat{e})$. [5 marks]

(*ii*) Show the converse, that $(\hat{d}, \hat{e}) \sqsubseteq (d_0, e_0)$. [10 marks]

16 Topics in Concurrency

(*a*) Describe the semantics of the modal μ-calculus. [4 marks]

(*b*) Describe without proof the meaning of the following modal μ-calculus assertions:

(*i*) $\nu Z.\ \langle c \rangle Z$; [1 mark]

(*ii*) $\mu Z.\ \langle c \rangle Z$; [1 mark]

(*iii*) $\nu Z.\ (A \wedge ([c]F \vee \langle c \rangle Z))$ (here F means false); [2 marks]

(*iv*) $\mu Z.\ (B \vee (A \wedge \langle c \rangle Z))$; [2 marks]

(*v*) $\nu Z.\ (B \vee (A \wedge \langle c \rangle Z))$. [2 marks]

(*c*) Consider the transition system consisting of two states p, q and two transitions $p \xrightarrow{c} q$ and $q \xrightarrow{c} p$.

(*i*) Does p satisfy $\mu Z.\ ([c]F \vee (\langle c \rangle T \wedge \langle c \rangle Z))$?

(*ii*) Does p satisfy $\nu Z.\ ([c]F \vee (\langle c \rangle T \wedge \langle c \rangle Z))$?

(Again, here F means false and T means true.) In this part you should justify your answers carefully. [8 marks]

END OF PAPER