COMPUTER SCIENCE TRIPOS Part II

Thursday 3 June 1999 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

Describe *one* of the following architectures for supporting distributed programming:

Either Remote Procedure Call

or Object Request Brokerage [11 marks]

Many emerging applications require timely response to events. Discuss the following approaches to achieving this:

(*a*) polling [3 marks]

(*b*) synchronous call back [3 marks]

(*c*) asynchronous notification [3 marks]

2 VLSI

What is a Muller C-element? Show its truth table and a circuit diagram using CMOS transistors. [10 marks]

Show how it can be used

(*a*) to control the flow of data through a micro-pipeline using bundled data; [5 marks]

(*b*) to detect completion after a pipeline stage using dual-rail encoding. [5 marks]

3 Digital Communication II

The following formula is used in a TCP implementation to estimate round trip times:

$$\overline{RTT}' = \alpha \times \overline{RTT} + (1 - \alpha) \times RTT_{sample}$$

Describe what this will do. How does one find a suitable value for the initial $\overline{RTT}$? How are RTT_{sample}'s taken? Under what conditions will this formula produce bad estimates? [10 marks]

An ISP wants to ensure that its customers use conformant TCP implementations.

How can a TCP implementation which functions satisfactorily as far as its user is concerned, be non-conformant? Why should the ISP be concerned?

Outline how you would provide tools for the ISP to detect non-conformant implementations. Where would you monitor TCP streams and why? [10 marks]

4 Advanced Graphics

Describe the Model Human Processor. [5 marks]

Explain the limitations of the Model Human Processor. [2 marks]

List the three ways of combining objects using constructive solid geometry (CSG). Describe how an object built using CSG can be represented using a binary tree. Given the intersection points of a ray with each primitive in the tree, explain how these points are passed up the tree by each type of combination node to find the first intersection point of the ray with the whole CSG object. [8 marks]

Explain what a form factor is, in radiosity. Outline an implementable method of calculating form factors. [5 marks]

5 Business Studies

Describe *five* criteria a Venture Capitalist might use for assessing a business. [5 marks]

Write, in terms of these criteria, an assessment of the Apple iMac project. [3 marks for each criterion]

6 Communicating Automata and Pi Calculus

Define the notion of *sorting* in the π-calculus, based upon a set $\mathcal{S}$ of non-parametric sorts. What is meant by saying that a set of agents *respects* a sorting? Explain briefly why, if $P \xrightarrow{\tau} P'$, then P' respects any sorting respected by P. [5 marks]

Simple data can be represented as abstractions in the π-calculus. In particular, if *True* and *False* represent the truth-values, then $\mathit{True}\langle b\rangle$ and $\mathit{False}\langle b\rangle$ are processes representing each truth-value located at b. Define these abstractions. Also, for arbitrary processes P and Q, define an abstraction $\mathit{Cond}(P,Q)$ such that

$$\begin{aligned} \mathit{Cond}(P,Q)\langle b\rangle \mid \mathit{True}\langle b\rangle &\rightarrow^* P \\ \mathit{Cond}(P,Q)\langle b\rangle \mid \mathit{False}\langle b\rangle &\rightarrow^* Q \end{aligned}$$

and demonstrate these reactions. Give a sorting respected by these agents.

[6 marks]

Define π-calculus abstractions which may be used to represent lists, whose elements may in turn be represented by abstractions. By analogy with *Cond*, define an abstraction $\mathit{Listcases}(P,(v\ell)Q)$ which can analyse a list, so that if L represents the empty list then

$$\mathit{Listcases}(P,(v\ell')Q)\langle \ell\rangle \mid L\langle \ell\rangle \rightarrow^* P$$

while if L represents a list whose head and tail are represented by V and L' then

$$\mathit{Listcases}(P,(v\ell')Q)\langle \ell\rangle \mid L\langle \ell\rangle \rightarrow^* \mathsf{new}\, v\ell'\, (Q \mid V\langle v\rangle \mid L\langle \ell'\rangle)$$

Demonstrate this last reaction. [6 marks]

Write down a sorting respected by the list abstractions and *Listcases*, involving a parametric sort $\textsc{list}(\sigma)$ (where σ may be any sort for the elements of a list).

[3 marks]

7 Optimising Compilers

Explain the term *instruction scheduling* understood as part of the compilation process. Describe what effects it might have on compile-time or run-time performance when performed by a compiler. [4 marks]

For a plausible, but not necessarily implemented, machine architecture of your choice give an algorithm which performs instruction scheduling. Note carefully what information on instructions is liable to be needed and the unit of code over which scheduling is done. [10 marks]

Discuss the following statements:

(*a*) "An instruction scheduling phase cannot do as well as an assembly code programmer because

```
ld.w   r0,a
st.w   r0,b
ld.w   r0,c
st.w   r0,d
```

cannot be interleaved." [3 marks]

(*b*) "My fast new special-purpose processor has 4 Kbyte of on-chip register-speed memory and hence has no *load delays*. It is not pipelined but instead its execution unit executes simultaneously the four 32-bit instructions found starting at (PC). The provided compiler works by placing instructions one-by-one in such a group of four and starting a new group of four (by completing any current group of four with no-ops if necessary) if the current instruction reads a value written by another of the current group of four. Hence there is no need for instruction scheduling." [3 marks]

8 Artificial Intelligence

Identify *six* application areas where reasoning under conditions of uncertainty is necessary, and explain briefly why. [12 marks]

Pick *one* of these areas and design four inference rules that might be used for reasoning in that domain. [8 marks]

9 Neural Computing

Give brief explanations of the following terms:

(*a*) the curse of dimensionality

(*b*) the Perceptron

(*c*) error back-propagation

(*d*) generalisation

(*e*) loss matrix

[4 marks each]

10 Denotational Semantics

Define what it means for two closed PCF terms of the same type to be *contextually equivalent*. [5 marks]

Explain why the denotational semantics of PCF using domains and continuous functions fails to be *fully abstract* with respect to its operational semantics. You may use any standard properties of the operational and denotational semantics without proof, so long as they are clearly stated. [15 marks]

11 Information Theory and Coding

What class of continuous signals has the greatest possible entropy for a given variance (or power level)? What probability density function describes the excursions taken by such signals from their mean value? [4 marks]

What does the Fourier power spectrum of this class of signals look like? What is notable about the entropy of this distribution of spectral energy? [4 marks]

An error-correcting Hamming code uses a 7-bit block size in order to guarantee the detection, and hence the correction, of any single bit error in a 7-bit block. How many bits are used for error correction, and how many bits for useful data? If the probability of a single bit error within a block of 7 bits is $p = 0.001$, what is the probability of an error correction failure, and what event would cause this? [4 marks]

Suppose that a continuous communication channel of bandwidth W, which is perturbed by additive white Gaussian noise of constant power spectral density, has a channel capacity C. Approximately how much would C be degraded if suddenly the added noise power became 8 times greater? [4 marks]

You are comparing different image compression schemes for images of natural scenes. Such images have strong statistical correlations among neighbouring pixels because of the properties of natural objects. In an efficient compression scheme, would you expect to find strong correlations in the compressed image code? What statistical measure of the code for a compressed image determines the amount of compression it achieves, and in what way is this statistic related to the compression factor? [4 marks]

12 Computer Vision

Contrast the use of linear *versus* non-linear operators in computer vision, giving at least one example of each. What can linear operators accomplish, and what are their fundamental limitations? With non-linear operators, what heavy price must be paid and what are the potential benefits? [8 marks]

When shape descriptors such as "codons" or Fourier boundary descriptors are used to encode the closed 2D shape of an object in an image, how can invariances for size, position, and orientation be achieved? Why are these goals important for pattern recognition and classification? [6 marks]

Define the general form of "superquadrics" used as volumetric primitives for describing 3D objects. What are their strengths and their limitations? [6 marks]

13 Types

A common idiom in typed programming languages is an *option type*, adding an additional value to an existing type. For example, in ML one might use

```
datatype 'a Option = None
                   | Some of 'a
```

Give a PLC encoding of `Option`, i.e. a PLC type Opt_α with free type variable α and suitable PLC terms such that

$$\{\alpha\}, \emptyset \vdash \mathrm{None}_\alpha : \mathrm{Opt}_\alpha$$
$$\{\alpha\}, \emptyset \vdash \mathrm{Some}_\alpha : \alpha \to \mathrm{Opt}_\alpha.$$

Give a typing derivation for None_α. [8 marks]

Any function, say $f : \gamma \to \delta$, can be lifted to a function of type $\mathrm{Opt}_\gamma \to \mathrm{Opt}_\delta$ that takes None_γ to None_δ and is as f elsewhere. Give a suitable PLC term $\mathrm{Lift}_{\gamma\delta}$ such that

$$\{\gamma, \delta\}, \emptyset \vdash \mathrm{Lift}_{\gamma\delta} : (\gamma \to \delta) \to (\mathrm{Opt}_\gamma \to \mathrm{Opt}_\delta).$$

Show the beta-equivalence

$$(\mathrm{Lift}_{\gamma\delta}\, f)\, (\mathrm{Some}_\gamma\, x) \;=_\beta\; \mathrm{Some}_\delta\, (f\, x)$$

[7 marks]

Similarly, functions $f : \gamma \to \mathrm{Opt}_\delta$ and $g : \delta \to \mathrm{Opt}_\epsilon$ can be composed. Give a suitable PLC term such that

$$\{\gamma, \delta, \epsilon\}, \emptyset \vdash \mathrm{Compose}_{\gamma\delta\epsilon} : (\gamma \to \mathrm{Opt}_\delta) \to (\delta \to \mathrm{Opt}_\epsilon) \to (\gamma \to \mathrm{Opt}_\epsilon).$$

[5 marks]

14 Numerical Analysis II

With reference to solution of the differential equation $y' = f(x, y)$, explain the conventional notation x_n, $y(x_n)$, y_n, and f_n. [3 marks]

Explain the terms *local error*, *global error*, and *order* of a method. [3 marks]

Milne's method uses the *multistep* formulae

$$y_{n+1} = y_{n-3} + \frac{4h}{3}(2f_n - f_{n-1} + 2f_{n-2})$$

$$y_{n+1} = y_{n-1} + \frac{h}{3}(\tilde{f}_{n+1} + 4f_n + f_{n-1})$$

which each have local error $O(h^5)$. Outline the general technique for deriving multistep formulae. What is the meaning of the term $\tilde{f}_{n+1}$? Suggest a suitable starting procedure and explain how the Milne formulae are used. [8 marks]

Let $y(x_0) = 0, h = 0.3$ and $f(x, y) = 3y/x - 2$. Suppose the following values of f_n have been generated by the starting procedure: 1.3, 2.1, 3.4 for $n = 1, 2, 3$. Calculate the first required value of $\tilde{f}_{n+1}$. [3 marks]

Contrast Milne's method with a comparable *one step* method, commenting particularly on *stability*, *efficiency* and *step size* considerations. [3 marks]

15 Security

Give examples of the use of a block cipher to provide assurance of

(*a*) confidentiality

(*b*) integrity

(*c*) timeliness

(*d*) covertness

(*e*) resistence to jamming

In each case describe a possible application and indicate the mode of operation you would use. [4 marks each]