COMPUTER SCIENCE TRIPOS Part II

Thursday 9 June 2005 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

STATIONERY REQUIREMENTS
Script Paper
Blue Coversheets
Tags

1 Human–Computer Interaction

A company that sells both spreadsheets and word processors has received complaints from users in the banking industry. The users often copy data from spreadsheets into letters offering special finance terms to individual customers. The default behaviour of the word processor "paste" command is simply to insert the numeric value, whereas a special option (in the "paste special ..." dialog) inserts a recalculating formula. The special option is used so regularly that users have requested an extra item on the pop-up (right click) menu.

(*a*) How would you estimate the increase in operation speed that might result from this change? [6 marks]

(*b*) How would you confirm the actual speed increase after constructing a prototype? [4 marks]

(*c*) The design team suggest an alternative – that the word processor should be enhanced with sufficient calculation functions that a spreadsheet is not needed at all. What factors should be taken into account in order to assess the effect this would have on user tasks? [10 marks]

2 VLSI Design

(*a*) Sketch stick diagrams of memory cells for:

(*i*) read-only memory; [3 marks]

(*ii*) static memory; [3 marks]

(*iii*) dynamic memory using standard CMOS; [3 marks]

(*iv*) dynamic memory for dense layout. [3 marks]

(*b*) Extend your design for parts (*ii*) and (*iii*) to accommodate four read ports. What would inhibit such an extension for design (*iv*)? [6 marks]

(*c*) State *two* considerations that are becoming increasingly important in memory design as feature sizes decrease. [2 marks]

3 Digital Communication II

Outline the interaction between the end-to-end flow of packets, and the following four buffering/queueing schemes. Be sure to compare and contrast their effects when deployed in a large-scale network.

(*a*) FIFO;

(*b*) RED;

(*c*) RED with ECN;

(*d*) Fair queueing.

[5 marks each]

4 Distributed Systems

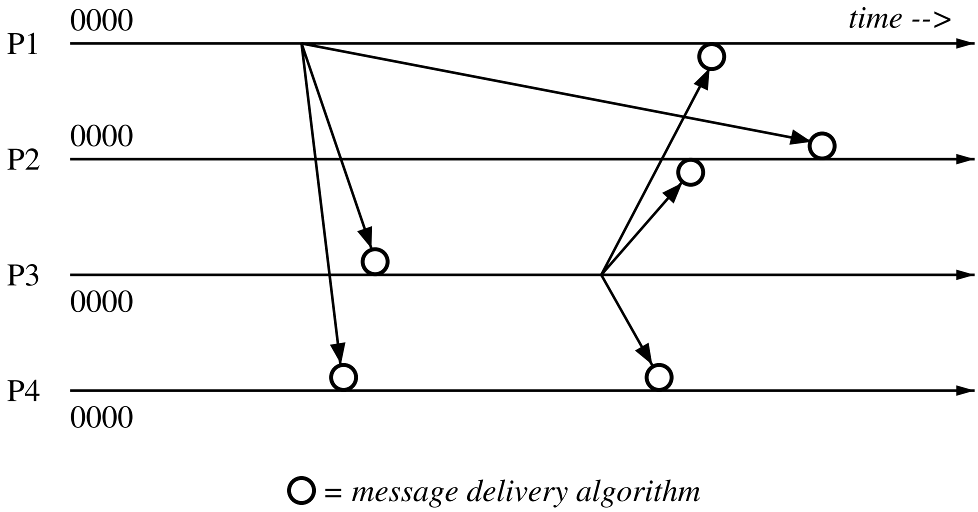

The above diagram represents a process group that communicates by means of multicast messages. At each process-hosting node, message delivery software decides whether a given incoming message should be delivered to the process or buffered for later delivery. This is achieved by the use of vector clocks.

(*a*) Describe, by means of the above example, the vector clock algorithm for delivery of messages in causal order. [15 marks]

(*b*) By means of a similar example, show that total ordering of messages is not achieved by this algorithm. [5 marks]

5 Advanced Systems Topics

(*a*) Compare and contrast the request routing mechanisms of Gnutella and Pastry. [10 marks]

(*b*) It has been said: "Unstructured peer-to-peer systems are better than structured peer-to-peer systems because they implement searching and complex queries." Describe how a structured system might be able to implement search. [10 marks]

6 Advanced Graphics

(*a*) Compare and contrast B-spline and subdivision representations of curves. [4 marks]

(*b*) Explain how B-spline basis functions are derived from the knot vector. [4 marks]

(*c*) Derive the quadratic uniform B-spline basis function (use the knot vector $[0, 1, 2, 3]$). [4 marks]

(*d*) Describe an algorithm to give the first intersection point of a ray with a closed cylinder of finite length aligned along the z-axis. [8 marks]

7 Optimising Compilers

(*a*) Summarise the basic principles behind strictness analysis including: what language paradigm it can be applied to, the representation of compile-time values expressing strictness, how these may be calculated and how the results of such calculations can be used to optimise programs. [8 marks]

(*b*) A program contains the following user function definitions. Give corresponding strictness functions assuming that `if-then-else` takes an integer as its first argument.

(*i*) `fun f(x) = 42` [1 mark]

(*ii*) `fun g(x) = g(x+1)` [1 mark]

(*iii*) `fun h(y,z) = if f(7) then y else z` [2 marks]

(*iv*) `fun k(x,y,z) = pif(x,y,z)` where `pif`(e, e', e'') is a primitive which evaluates its three arguments in parallel, returning e' if e evaluates to a non-zero integer, returning e'' if e evaluates to zero and *also* returning e' if e' and e'' evaluate to the same integer even if e is still being evaluated. [4 marks]

(*c*) "Any Boolean expression *be* containing variables $\{x_1, \ldots, x_k\}$ but not containing negation can be expressed as the strictness function for a user-defined function `fun u(`$x_1, \ldots, x_k$`) =` e." Argue that this statement is true, showing how to construct some such e from a given *be*. [4 marks]

[Hint: you may assume *be* has been written in DNF form

$$(v_{11} \wedge \cdots \wedge v_{1m_1}) \vee \cdots \vee (v_{n1} \wedge \cdots \wedge v_{nm_n})$$

where v_{ij} are members of $\{x_1, \ldots, x_k\}$.]

8 Artificial Intelligence II

We wish to model the unobservable state of an environment using a sequence $S_0 \to S_1 \to S_2 \to \cdots$ of sets of random variables (RVs) where at time i we are in state S_i and observe a set of RVs E_i. The distributions of the RVs do not change over time, and observations depend only on the current state.

(*a*) Define a *Markov process*, the *transition model* and the *sensor model* within this context. [3 marks]

(*b*) Assuming that evidence $E_{1:t} = e_{1:t} = (e_1, e_2, \ldots, e_t)$ has been observed define the tasks of *filtering*, *prediction* and *smoothing*. [3 marks]

(*c*) Derive a recursive estimation algorithm for performing filtering by combining the evidence e_t obtained at time t with the result of filtering at time $t-1$. [8 marks]

(*d*) How does a *hidden Markov model* differ from the setup described? [1 mark]

(*e*) Show how for the case of a hidden Markov model your filtering algorithm can be expressed in terms only of matrix operations. [5 marks]

9 Bioinformatics

(*a*) Present the aim of phylogeny algorithms:

(*i*) Describe the main differences between Parsimony, Distance and Likelihood-based algorithms. [5 marks]

(*ii*) Describe the input and the output of a distance-based algorithm. [5 marks]

(*iii*) Discuss the complexity of the Neighbour-Joining algorithm. [5 marks]

(*b*) Describe with one example the Needleman–Wunsch algorithm. [5 marks]

10 Types

Give a polymorphic lambda calculus (PLC) type $list_\alpha$ that contains a single free type variable α and which corresponds to the ML datatype of polymorphic lists:

```
datatype 'a list = Nil | Cons of 'a * ('a list)
```

[2 marks]

Give PLC expressions *Nil*, *Cons* and *iter* of appropriate types that encode the ML constructors `Nil` and `Cons` and the ML function `iter` given by

```
fun iter x f Nil = x
  | iter x f (Cons(h, t)) = f h (iter x f t)
```

You should prove the PLC typings you claim for these expressions. [13 marks]

Show that *iter* has β-conversion properties corresponding to the above declaration of the ML function `iter`. [5 marks]

11 Computer Vision

(*a*) Consider the following isotropic 2D filter function $f(x, y)$ incorporating the Laplacian operator $\nabla^2 = \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right)$, that is often used in computer vision:

$$f(x, y) = \nabla^2 e^{-\left(x^2+y^2\right)/\sigma^2}$$

(*i*) In 2D Fourier terms, what type of filter is this? (For example, is it a lowpass, a highpass, or a bandpass filter?) [2 marks]

(*ii*) Are different orientations of image structure treated differently by this filter, and if so, how? [2 marks]

(*iii*) Approximately what is the spatial frequency bandwidth of this filter, in octaves? [Hint: the answer is independent of σ.] [1 mark]

(*iv*) What is meant by image operations "at a certain scale of analysis"? In this context, define a scale-space fingerprint, and explain the rôle of the scale parameter. [3 marks]

(*b*) What surface properties can cause a human face to form either a Lambertian image or a specular image, or an image lying anywhere on a continuum between those two extremes? In terms of geometry and angles, what defines these two extremes of image formation? What difficulties do these factors create for efforts to extract facial structure from facial images using "shape-from-shading" inference techniques? [3 marks]

(*c*) Why can't any computer vision operations be performed directly on `.jpeg` image formats? [1 mark]

(*d*) Discuss the significance of the fact that mammalian visual systems send perhaps ten times as many corticofugal neural fibres back down from the visual cortex to the thalamus, as there are ascending neural fibres bringing visual data from the retina up to the thalamus. Does this massive neural feedback projection support the thesis of "vision as graphics", and if so how? [4 marks]

(*e*) Explain why inferring object surface properties from image properties is, in general, an ill-posed problem. In the case of inferring the colours of objects from images of the objects, how does knowledge of the properties of the illuminant affect the status of the problem and its solubility? [4 marks]

12 Numerical Analysis II

(*a*) Explain the term *positive semi-definite*. If $\mathbf{A}$ is a real square matrix show that $\mathbf{A}^T\mathbf{A}$ is symmetric and positive semi-definite. [3 marks]

(*b*) How is the l_2 *norm* of $\mathbf{A}$ defined? State Schwarz's inequality for the product $\mathbf{Ax}$. [2 marks]

(*c*) Describe briefly the properties of the matrices $\mathbf{U}$, $\mathbf{W}$, $\mathbf{V}$ in the *singular value decomposition* $\mathbf{A} = \mathbf{UWV}^T$. [3 marks]

(*d*) Let $\hat{\mathbf{x}}$ be an approximate solution of $\mathbf{Ax} = \mathbf{b}$, and write $\mathbf{r} = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}$, $\mathbf{e} = \mathbf{x} - \hat{\mathbf{x}}$. Derive a computable estimate of the relative error $\|\mathbf{e}\|/\|\mathbf{x}\|$ in the approximate solution, and show how this may be used with the l_2 norm. [8 marks]

(*e*) Suppose $\mathbf{A}$ is a 7×7 matrix whose singular values are 10^2, 10^{-4}, 10^{-10}, 10^{-16}, 10^{-22}, 10^{-29}, 10^{-56}. Construct the matrix $\mathbf{W}^+$ that you would use (*i*) if *machine epsilon* $\simeq 10^{-15}$, and (*ii*) if *machine epsilon* $\simeq 10^{-30}$. [4 marks]

13 Specification and Verification II

(*a*) Discuss the challenge of modelling transistors in a way that is both tractable for formal verification and accurate enough to be useful. [5 marks]

(*b*) Discuss the accuracy of the simple switch model of transistors and discuss how the model can be improved. [5 marks]

(*c*) Outline how transistor circuits that use precharging can be formally modelled in higher order logic. Briefly discuss potential inaccuracies of such models. [5 marks]

(*d*) Describe how to specify that a bit-level circuit with 4-bit inputs *a* and *b* and an 8-bit output *prod* performs multiplication. [5 marks]

14 Natural Language Processing

(*a*) The Figure below shows feature structures corresponding to lexical entries for *snores* and *he*.

$$
\begin{bmatrix}
\text{HEAD} & \begin{bmatrix} \text{CAT} & \textbf{verb} \end{bmatrix} \\
\text{COMP} & \textbf{filled} \\
\text{SPR} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \text{CAT} & \textbf{noun} \end{bmatrix} \\ \text{SEM} & \begin{bmatrix} \text{INDEX} & \boxed{1} \end{bmatrix} \end{bmatrix} \\
\text{SEM} & \begin{bmatrix} \text{PRED} & \textbf{snore\_v} \\ \text{ARG1} & \boxed{1} \end{bmatrix}
\end{bmatrix}
\qquad
\begin{bmatrix}
\text{HEAD} & \begin{bmatrix} \text{CAT} & \textbf{noun} \end{bmatrix} \\
\text{COMP} & \textbf{filled} \\
\text{SPR} & \textbf{filled} \\
\text{SEM} & \begin{bmatrix} \text{INDEX} & \boxed{1} \\ \text{PRED} & \textbf{pron} \\ \text{ARG1} & \boxed{1} \end{bmatrix}
\end{bmatrix}
$$

These structures can be combined using a grammar rule to give a feature structure corresponding to the phrase *he snores* with a semantic structure equivalent to $\text{pron}(x) \wedge \text{snore\_v}(x)$. Give this grammar rule as a feature structure and show the results of applying the rule to the structures in the Figure. [7 marks]

(*b*) Syntactically the verb *rains* takes the pleonastic pronoun *it* as subject but semantically it has no arguments. Give possible feature structures for *rains* and pleonastic *it*. Show how ungrammatical sentences such as *he rains* are avoided, mentioning any modifications to the lexical entries in the Figure that might be necessary. [6 marks]

(*c*) Selectional restrictions can be used to block parses of semantically anomalous sentences such as:

The pebble snores.
The pebble wrote a book.
The dog wrote a book.

Describe how selectional restrictions might be encoded in a feature structure grammar. [7 marks]

15 Denotational Semantics

(*a*) Describe how to construct the function cpo $((D \to E), \sqsubseteq)$ of two cpos $(D, \sqsubseteq_D)$ and $(E, \sqsubseteq_E)$. Prove that $((D \to E), \sqsubseteq)$ is a cpo.
(You may use general facts about least upper bounds provided you state them clearly.) [7 marks]

(*b*) The function *uncurry* is inverse to the function *curry*; it takes a continuous function in $(D_1 \to (D_2 \to E))$ as argument and yields a continuous function in $((D_1 \times D_2) \to E)$ as result. Give a definition of *uncurry* and show it is a continuous function.
(You may use general facts about continuous functions provided you state them clearly.) [6 marks]

(*c*) Exhibit two terms of PCF which are contextually equivalent and yet have distinct denotations in the domain $(\mathbb{B}_\perp \to (\mathbb{B}_\perp \to \mathbb{B}_\perp)) \to \mathbb{B}_\perp$ where $\mathbb{B} = \{true, false\}$ is the set of truth values. Explain why their denotations differ. [7 marks]

16 Topics in Concurrency

The syntax of parallel commands is given by:

$$c ::= X := a \;\mid\; c_0; c_1 \;\mid\; c_0 \,||\, c_1 \;\mid\; \textbf{if } b \textbf{ then } c \;\mid\; \textbf{while } b \textbf{ do } c$$

where X ranges over locations, a over arithmetic expressions, and b over boolean expressions.

(*a*) Give an operational semantics to parallel commands, assuming an operational semantics for arithmetic and boolean expressions. [5 marks]

(*b*) This part is concerned with a Petri net semantics for parallel commands.

There are to be two kinds of conditions: *data conditions*, pairs of locations and integers, which specify the contents of locations, and *control conditions*, which specify the local control points in parallel components of commands.

A parallel command is to be represented by a basic net (where every condition has capacity one) in which a subset of control conditions I is to be distinguished as its *initial* conditions and another subset T is to be distinguished as its *terminal* conditions; the initial conditions are precisely those control conditions which hold at the start of execution of the command; the terminal conditions are precisely those control conditions which hold if and when the command terminates.

A diagrammatic account suffices for answers to the questions below.

(*i*) Describe an (infinite) net for $X := X + 1$. [2 marks]

(*ii*) Describe a construction on nets for $c_0; c_1$. [Hint: Replace the terminal conditions T_0 of c_0 and the initial conditions I_1 of c_1 with their product $T_0 \times I_1$.] [4 marks]

(*iii*) Describe a construction on nets for $c_0 \,||\, c_1$. [2 marks]

(*iv*) Describe a construction on nets for **if** $X > 0$ **then** c. [2 marks]

(*v*) Describe a construction on nets for **while** $X > 0$ **do** c. [5 marks]

END OF PAPER