COMPUTER SCIENCE TRIPOS Part II

Wednesday 6 June 2001 1.30 to 4.30

Paper 8

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.*

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 Distributed Systems

Write brief notes on *each* of the following:

(*a*) remote procedure call [4 marks]

(*b*) object orientated middleware [4 marks]

(*c*) message orientated middleware [4 marks]

(*d*) event-based middleware [4 marks]

(*e*) publish–subscribe systems [4 marks]

Your notes should highlight the significant differences between these forms of middleware and should mention any application areas for which each is especially suited.

2 VLSI Design

(*a*) The *constant field* model of MOS scaling applies a dimensionless factor α to manufacturing dimensions (length, width and thickness), voltages and processing concentrations, so that channel thickness remains unchanged. For example, with $\alpha = 1$, the dimensions are unchanged; with $\alpha = 2$, they would be halved. Derive approximate expressions for the consequent scaling of

- gate area
- channel resistance
- current
- load capacitance
- gate delay
- static power consumption (per gate)
- power density (per unit area)
- current density (in wires)

What are the main implications for speed, size and power? [8 marks]

(*b*) *Constant voltage* is an alternative model in which the only manufacturing dimensions are scaled, leaving voltages unchanged, so the channel thickness increases by a factor α. Derive approximate expressions for the consequent scaling and summarise the main implications. [8 marks]

(*c*) What further factors make both models inappropriate as device sizes continue to decrease? [4 marks]

3 Digital Communication II

(*a*) In the context of networking, what is *congestion*? [2 marks]

(*b*) Discuss the evolution of congestion control in TCP. You may like to consider some or all of the following in your answer: *congestion window*, *slow start*, *fast retransmit*, *fast recovery*, *TCP Vegas*, *RED*, *ECN*, and *congestion pricing.* [12 marks]

(*c*) Compare and contrast congestion control in TCP with the ways in which congestion is managed in ATM networks. [6 marks]

4 Advanced Graphics and HCI

(*a*) For a given order, k, there is only one basis function for uniform B-splines. Every control point uses a shifted version of that one basis function. How many different basis functions are there for open-uniform B-splines of order k with $n+1$ control points, where $n \geqslant 2k-3$? [6 marks]

(*b*) Explain what is different in the cases where $n < 2k-3$ compared with the cases where $n \geqslant 2k-3$. [3 marks]

(*c*) Sketch the different basis functions for $k = 2$ and $k = 3$ (when $n \geqslant 2k-3$). [4 marks]

(*d*) Show that the open-uniform B-spline with $k = 3$ and knot vector [000111] is equivalent to the quadratic Bezier curve. [7 marks]

5 Business Studies

You are inspired to make your fortune by starting a distance learning enterprise to teach the world Computer Science, using multimedia lessons distributed over the Web.

Write notes for a business plan for the potential investors, under the following headings:

(*a*) The market. [5 marks]

(*b*) The team required. [5 marks]

(*c*) Outline overall project plan. [5 marks]

(*d*) Business model, with a rough estimation of capital expenditure and profitability. [5 marks]

6 Security

In the Wired Equivalent Privacy protocol used in IEEE 802.11 networks, data are protected at the link level during transmission on a wireless LAN. Each frame has a 32-bit CRC appended to it; it is then encrypted using the RC4 stream cipher, initialised with a shared key and a 24-bit initial value; and finally, the initial value is sent with the encrypted frame.

(*a*) Why is the initial value used? [4 marks]

(*b*) Is the CRC an appropriate mechanism, and, if not, what should be used instead? [4 marks]

(*c*) Describe *one* passive attack on this system. [4 marks]

(*d*) Describe *one* active attack on this system. [4 marks]

(*e*) What would be the effect of upgrading from RC4 to a stronger cipher, such as AES used in output feedback mode? [4 marks]

7 Optimising Compilers

(*a*) Explain the ideas of strictness analysis, including over what languages the ideas are applicable and what transformations are enabled by it. Describe how strictness functions for (*i*) built-in and (*ii*) user-defined functions are defined, clarifying the similarities and differences. [10 marks]

(*b*) A language has a user-defined function f which is defined in terms of built-in functions $a_1, \cdots, a_t$ and possibly recursion. Later, to aid efficiency, an additional function a_{t+1} is added to the set of system functions, but its effect (semantics) is the same as that of f. By considering examples similar to those used to show analyses are safe but imprecise, or otherwise, determine a relationship between the strictness functions $f^\sharp$ and $a^\sharp_{t+1}$. [5 marks]

(*c*) It is noted that strictness functions, e.g.

$$cond^\sharp(x, y, z) = x \wedge (y \vee z)$$

do not generally use negation in their defining boolean expressions. Show that all strictness functions can be written without negation or find a counter-example. Hint: No computable function f can have semantics such that there are x and y which satisfy

$$f(x, y) = \bot \text{ and } f(x, \bot) \neq \bot.$$

[5 marks]

8 Artificial Intelligence

Can a computer think? [20 marks]

9 Neural Computing

(*a*) (*i*) A competitive Kohonen neural network forms feature maps which can be regarded as performing dimensionality reduction. Explain this. [4 marks]

(*ii*) Is training time normally faster, or slower, in a supervised neural network compared with an unsupervised one? What is the major disadvantage inherent in the use of supervised neural networks? [2 marks]

(*iii*) What class of neural network can be used to overcome the mathematical difficulties caused by the use of non-orthogonal sensory and motor representations? [2 marks]

(*b*) (*i*) Give *three* examples of biological sensory or motor control systems that seem to rely on the use of non-orthogonal coordinates. [3 marks]

(*ii*) Explain why this creates a problem in the computational evaluation and simulation of such systems, and discuss whether or not you think this issue matters in the function of the actual neurobiological systems. [2 marks]

(*c*) (*i*) Give *four* examples of neural activity having a fundamentally quantal structure, in the sense that signals or events are quantised into discrete packages rather than being continuous. [4 marks]

(*ii*) For purposes of understanding neurobiological computation, what can be learned from studying the brain's failures, either as the consequences of specific forms of trauma or in normal function as revealed in the systematic visual illusions? [3 marks]

10 Comparative Architectures

An important application spends a large proportion of its running time executing a particular loop. The loop is responsible for summing two arrays containing unsigned eight-bit values packed in memory into a similar third array. Saturating arithmetic is used, whereby results that overflow are "clipped" to the maximum representable value. For example, for the unsigned eight-bit case, 250 plus 20 would result in 255. In this particular application, such overflows are rare in practice, and this fact may be exploited to optimise the implementation.

(*a*) Write pseudo code for a simple implementation of the inner loop for a 32-bit processor with a RISC-like instruction set. [Hint: It is possible to use the CPU's 32-bit ALU to perform four eight-bit additions with a single add instruction, and then use further code to detect if overflow occurred and correct it. You may assume the arrays start on word-aligned boundaries.] [10 marks]

(*b*) Assuming the arrays are present in the CPU's L1 data cache, estimate the number of cycles required to sum arrays of length N on a statically-scheduled two-way super-scalar processor. State any assumptions you make. [5 marks]

(*c*) Many instruction set architectures have been augmented with SIMD (Single Instruction Multiple Data) instructions to enhance processors' performance when dealing with packed arrays such as those used by this application. Demonstrate how SIMD instructions could be used to optimise the loop. Assuming that 50% of the running time of the application was spent executing your previous loop implementation, estimate the program's speedup when using the SIMD optimised loop. [5 marks]

11 Numerical Analysis II

(*a*) A cubic spline over knots $x_1, x_2, \ldots x_n$ is defined by

$$\phi(x) = \frac{(x - x_j)y_{j+1} + (x_{j+1} - x)y_j}{d_j} - \frac{(x - x_j)(x_{j+1} - x)\{(d_j + x_{j+1} - x)\mu_j + (d_j + x - x_j)\mu_{j+1}\}}{6d_j}$$

for $x \in [x_j, x_{j+1}]$ where $d_j = x_{j+1} - x_j$. The spline is continuous in its first and second derivatives.

(*i*) Find $\phi(x_j)$. [2 marks]

(*ii*) Find formulae for $\phi'(x_j)$ and $\phi'(x_{j+1})$ for $x \in [x_j, x_{j+1}]$. [4 marks]

(*iii*) What is $\phi''(x_j)$? [2 marks]

(*b*) Form a set of equations for computing the unknowns $\{\mu_j\}$, specifying suitable end conditions to simplify these equations. [10 marks]

(*c*) What are the important properties of these equations with respect to their numerical solution? [2 marks]

12 Specification and Verification I

(*a*) Describe the axioms and rules of Floyd–Hoare logic for reasoning about `FOR`-commands. Carefully explain any side conditions. [8 marks]

(*b*) Let $n!$ be the factorial of n ($0! = 1$ and $(n+1)! = (n+1) \times n!$). Give a proof of

```
{N ⩾ 0} X := 1; FOR Y := 2 UNTIL N DO X := X × Y {X = N!}
```

[12 marks]

13 Computer Vision

Understanding, classifying, and identifying human faces has been a longstanding goal in computer vision. Yet because the face is an expressive social organ, as well as an object whose image depends on identity, age, pose and viewing angle, and illumination geometry, many forms of variability are all confounded together, and the performance of algorithms on these problems remains very poor. Discuss how the different kinds and states of variability (e.g. same face, different expressions; or same identity and expression but different lighting geometry) might best be handled in a statistical framework for generating categories, making classification decisions, and recognising identity. In such a framework, what are some of the advantages and disadvantages of wavelet codes for facial structure and its variability?

[20 marks]

14 Computer Systems Modelling

Two servers operate with different performance characteristics at mean rates μ_1 and μ_2. You wish to combine them into a single system by associating each server with a separate FIFO queue and dispatching incoming work items to the first queue with probability p_1 and to the other queue with probability p_2. Incoming items arrive at a rate λ and none are discarded from the system.

You may assume that the inter-arrival-time distribution and both service-time distributions are exponential, that there is no limit on the queue lengths and that the population size is infinite.

(*a*) Using Kendall notation, describe the first server and its queue. Construct a Markov-chain model for this part of the system. [2 marks]

(*b*) Let $q_{k,i}$ denote the probability that there are exactly i items of work in server k and its queue. By using detailed flow balance equations or otherwise express $q_{k,i}$ in terms of λ, p_k and μ_k. [6 marks]

(*c*) Hence derive T_k, the mean response time of work items served at k. [6 marks]

(*d*) Suppose that the system administrator wishes to ensure that work items receive the same mean response time irrespective of which server they visit. Express p_1 in terms of λ, μ_1 and μ_2. Qualitatively, when is it reasonable to consider dispatching work to both servers to maintain an equal mean response time? How will the system behave at other times? [6 marks]

15 Topics in Concurrency

(*a*) Describe an algorithm for deciding whether or not a finite-state CCS process satisfies an assertion in the modal μ-calculus. [6 marks]

(*b*) Draw the reachable transition system of the CCS process P, where

$$P \overset{\text{def}}{=} a.P + b.(b.nil + a.nil).$$ [2 marks]

(*c*) Illustrate the use of the algorithm of part (*a*) by giving a derivation which decides whether or not the CCS process P above satisfies the modal μ-calculus assertion A, where

$$A \equiv \nu X.([b]F \vee (\langle a\rangle T \wedge [\cdot]X)).$$

(You should assume the usual properties of boolean operations.) [8 marks]

(*d*) Give, without proof, a short description of those finite-state processes which satisfy the assertion A. [4 marks]

END OF PAPER