COMPUTER SCIENCE TRIPOS Part II

Tuesday 6 June 2006 1.30 to 4.30

PAPER 7

*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 Comparative Architectures

(*a*) Give *three* fundamental differences between RISC and CISC processor design and explain why RISC might lead to improved performance. [6 marks]

(*b*) Members of a *binary compatible* family of processors may have rather different hardware designs. Define this term and explain what variations between the processors in a family might exist, why they exist and how they are made to appear the same. [7 marks]

(*c*) Processor families vary in the number of user registers in their basic programming models. How does this affect performance and what extensions are used, both above and below the programming model, to provide more register storage? [7 marks]

2 Digital Communication II

(*a*) Flow control mechanisms in networks can be loosely categorised as *open loop*, and *closed loop*.

(*i*) Describe the closed loop feedback control scheme employed by the Transmission Control Protocol (TCP) in the Internet today for flow and congestion control. [6 marks]

(*ii*) The telephone network of today and the Resource Reservation Protocol (RSVP) and associated services proposed for the future Internet provide open loop control. Briefly explain the functions of *signalling*, *admission control* and *policing* in these types of networks. [3 marks each]

(*b*) Traffic sources are described as elastic or inelastic. Outline how an inelastic traffic source can be characterised and how this relates to the type of network resource guarantees it may need. [5 marks]

3 Security

(*a*) What is meant by *Mandatory Access Control*? Give an example. [5 marks]

(*b*) How might you use mandatory access control to protect the safety-critical systems in a car (engine control unit, ABS, stability control, etc.) from user-programmable systems (telephone, entertainment, navigation, etc.)? [5 marks]

(*c*) What problems would you anticipate in keeping the implementation clean as these systems evolve? [5 marks]

(*d*) What architecture might you therefore propose a car maker adopt for its next-generation networking? [5 marks]

4 Advanced Graphics

(*a*) Describe, in outline, each of the implicit surface, NURBS surface, and constructive solid geometry methods for defining three-dimensional shapes. [4 marks each]

(*b*) Compare and contrast the three methods. [8 marks]

5 Computer Systems Modelling

(a) Describe the congruential methods for generating pseudo-random numbers from a Uniform $(0, 1)$ distribution. [3 marks]

(b) Let U be a Uniform $(0, 1)$ random variable. Show that for any continuous distribution function, $F(x)$, the random variable, X, defined by

$$X = F^{-1}(U)$$

has the probability distribution function $F(x)$. [3 marks]

(c) Apply the method of part (b) to generate random variables with the following distributions. In each case, specify the distribution function $F(x)$ that you use.

(i) Uniform distribution on the interval (a, b), for $a < b$. [2 marks]

(ii) Exponential distribution with parameter λ. [2 marks]

(d) Define the Poisson process, $N(t)$, $(t \geq 0)$ of rate λ. [2 marks]

(e) Show that for each fixed $t \geq 0$, $N(t)$ is a Poisson random variable with parameter λt. [3 marks]

(f) Show that the interarrival times of consecutive events in a Poisson process of rate λ are independent random variables each with the exponential distribution with parameter λ. Show how this leads to a method to simulate the events of a Poisson process. [5 marks]

6 Specification and Verification I

(*a*) What is total about total correctness? [2 marks]

(*b*) State the `WHILE`-Rule of Floyd–Hoare Logic. [2 marks]

(*c*) Give a proof of `{T} WHILE T DO SKIP {F}`. [2 marks]

(*d*) What does the truth of `{T} WHILE T DO SKIP {F}` show? [2 marks]

(*e*) How are expressions like `1/0` handled in Floyd–Hoare Logic? [2 marks]

(*f*) What are verification conditions? [2 marks]

(*g*) Must the verification conditions be true for correctness? Briefly justify your answer. [2 marks]

(*h*) Name *one* method used to prove verification conditions. [2 marks]

(*i*) What are the "hooked" variables in VDM used for? [2 marks]

(*j*) What are weakest preconditions and weakest liberal preconditions? [2 marks]

7 Specification and Verification II

(*a*) What are Binary Decision Diagrams and how are they used to represent state-transition functions symbolically? [4 marks]

(*b*) What is temporal abstraction? How are models at different temporal abstraction levels related? [4 marks]

(*c*) What is the difference between LTL and CTL? [4 marks]

(*d*) How do the Verilog and VHDL simulation cycles differ? [4 marks]

(*e*) What is the difference between formal verification using model checking and using theorem proving? [4 marks]

8 Information Theory and Coding

(*a*) Give *three* different expressions for mutual information $I(X;Y)$ between two discrete random variables X and Y, in terms of their two conditional entropies $H(X|Y)$ and $H(Y|X)$, their marginal entropies $H(X)$ and $H(Y)$, and their joint entropy $H(X,Y)$. Explain in ordinary language the concept signified by each of the measures $H(X|Y)$, $H(X)$, $H(X,Y)$, and $I(X;Y)$. Depict in a Venn diagram the relationships among all of the quantities mentioned here. [8 marks]

(*b*) 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? [2 marks]

(*c*) What is the shortest possible code length, in bits per average symbol, that could be achieved for a six-letter alphabet whose symbols have the following probability distribution?

$$\left\{ \tfrac{1}{2}, \tfrac{1}{4}, \tfrac{1}{8}, \tfrac{1}{16}, \tfrac{1}{32}, \tfrac{1}{32} \right\}$$

[3 marks]

(*d*) If we wish to increase the transmission capacity of a noisy communication channel, is it more effective to increase its electronic bandwidth in Hertz, or to improve its signal-to-noise ratio? Briefly say why. [2 marks]

(*e*) A continuous signal whose total bandwidth is 1 kHz and whose duration is 10 seconds may be perfectly represented (even at points in between the points at which it is sampled) by what minimal number of real numbers? [2 marks]

(*f*) Give the names of *three* functions (not necessarily their equations) which are self-Fourier. [3 marks]

9 Types

(*a*) Give the typing rules for the Mini-ML language with booleans, conditionals, function abstraction and application, local declarations, nil, cons and case expressions. Make the form of the typing judgement clear and explain any auxiliary relations used in the side-conditions of the rules. [8 marks]

(*b*) What is a *principal* type scheme for a closed Mini-ML expression? State without proof the Hindley–Damas–Milner theorem for the Mini-ML typeability problem. [5 marks]

(*c*) Which of the following Mini-ML expressions are typeable? Give the principal type scheme for any of the expressions that are typeable.

(*i*) $\lambda x(x :: \texttt{nil})$ [2 marks]

(*ii*) $\lambda f(f(f :: \texttt{nil}))$ [2 marks]

(*iii*) $\texttt{let } f = \lambda x(x :: \texttt{nil}) \texttt{ in } f(f :: \texttt{nil})$ [3 marks]

10 Topics in Concurrency

(*a*) Present the transition semantics rules for the CCS operations of prefixing, binary sum and parallel composition. [4 marks]

(*b*) Describe diagrammatically a Petri net semantics for the CCS operations of binary sum and parallel composition. [6 marks]

(*c*) (*i*) Draw the transition system associated with the CCS process P defined by $P \stackrel{\text{def}}{=} a.(P + a.b.P)$. [2 marks]

(*ii*) Does P satisfy $\mu Z.\ (\langle b\rangle T \vee (\langle a\rangle T \wedge [a]Z))$? Justify your answer. [4 marks]

(*iii*) Does P satisfy $\nu Z.\ (\langle b\rangle T \vee (\langle a\rangle T \wedge [a]Z))$? Justify your answer. [4 marks]

Here T means true. (Although a rigorous proof is not required, you should justify your answers carefully.)

11 Information Retrieval

Information retrieval systems vary in the expressivity of the query languages they employ. For instance, some systems support proximity search: if two query terms are connected by the "Next" operator, then only those documents are retrieved where the query terms appear close together (i.e., within a certain number of words of each other).

(*a*) List and briefly describe other ways in which the syntax and the interpretation of query languages may vary. [8 marks]

(*b*) Describe with an example how the "Next" operator described above is implemented efficiently in modern information retrieval systems. Your answer should include a description of the data structure(s) necessary to support it. [5 marks]

(*c*) A search engine supports error correction in the following way: If an error is suspected in a query term, the system provides a link labelled "Did you mean X?", where X is the corrected term, in addition to its normal results. The link leads to a list of retrieved documents, corresponding to a variant of the original query, with X replacing the misspelled term.

(*i*) Explain why it is non-trivial to implement this feature efficiently. [3 marks]

(*ii*) Discuss methods for implementing this feature in a realistic setting. [4 marks]

12 Natural Language Processing

The following shows a simple context free grammar (CFG) for a fragment of English.

```
S -> NP VP              Adj -> angry          Vbe -> is
VP -> Vbe Adj           Adj -> big            N -> dog
NP -> Det N             Adj -> former         N -> cat
N -> Adj N              P -> at
Adj -> Adj PP           P -> on
PP -> P NP              Det -> the
```

(*a*) Show the parse tree that this grammar would assign to (1).

(1) the dog is angry at the cat

[3 marks]

(*b*) One respect in which this grammar overgenerates is that some adjectives, including *former*, occur only before a noun (see (2)) and that PPs do not combine with adjectives occurring before a noun (see (3)).

(2) * the dog is former

(3) * the angry at the cat dog is big

Show how the grammar given above could be modified to prevent this type of overgeneration. [4 marks]

(*c*) The grammar also behaves incorrectly with examples (4), (5) and (6):

(4) * the dog is big at the cat (*big* does not take a PP)

(5) * the dog is angry on the cat (*angry* only takes PPs where the P is *at*)

(6) * the dog is angry at the cat at the cat (adjectives may not combine with multiple PPs)

Show modifications to the grammar which would prevent these types of overgeneration. [5 marks]

(*d*) Describe how the overgeneration in part (*c*) could be dealt with in a feature structure (FS) grammar, giving full lexical entries for *angry* and *big* and details of rules and other lexical entries as necessary to explain your account.

[8 marks]

13 Business Studies

(*a*) Name *five* different types of intellectual property. [5 marks]

(*b*) Distinguish between "Deep linking" and "Direct linking". Can a search engine deep link without infringing the copyright of the original site? [5 marks]

(*c*) Why is the use of thumbnails of pictures by a search engine fair use? [5 marks]

(*d*) Why are there likely to be only a few dominant search engines? [5 marks]

14 E-Commerce

(*a*) In a telecommunications business context, what is meant by *Triple* and *Quadruple Play* strategies, and what will be the effects of the adoption of such strategies? [10 marks]

(*b*) Describe possible business models for a small independent Internet television start-up company. Estimate start-up costs and profitability, and describe some of the challenges such a television station will need to overcome. [10 marks]

15 Additional Topics

(*a*) (*i*) Briefly describe the purpose of broadband fixed wireless access (FWA) and give typical situations where it could be deployed. [4 marks]

(*ii*) What advantages does radio offer over traditional wired access? [2 marks]

(*b*) Explain how intersymbol interference (ISI) arises in a wireless communication system and its effect on system performance. [2 marks]

(*c*) (*i*) Describe how

- linear equalisation
- orthogonal frequency division multiplexing (OFDM)

can be used to combat ISI. [4 marks]

(*ii*) Highlight a problem which can occur when using linear equalisation and describe an alternative equalisation approach that can overcome this problem. What problem can occur with this alternative approach? [4 marks]

(*iii*) Highlight problems that arise with the use of OFDM. In what situation is OFDM preferred over the use of equalisation? [4 marks]

16 Additional Topics

(*a*) What is an (m, M)*R-Tree*? Explain how new items can be inserted and how *range queries* can be evaluated. [6 marks]

(*b*) Explain the trade-off in varying m between small and large values. If the dataset is known in advance, would a large or small value of m be appropriate? [4 marks]

(*c*) R^+*-Trees*, R^**-Trees*, and *QSF-Trees* are special forms of R-Tree. Explain how each differs from the basic R-Tree and what advantage is presented by each modification. [2 marks each]

(*d*) As private motor cars increase in electronic sophistication, Sentient Computing becomes ever more applicable. Describe *two* context-aware behaviours that a car's electronic systems could exhibit, making use of general-purpose processing power, data storage, and wireless data communication. [2 marks each]

END OF PAPER