COMPUTER SCIENCE TRIPOS Part II

Tuesday 3 June 2003 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

1 Comparative Architectures

(*a*) Compare and contrast each of the following techniques for achieving instruction-level parallelism:

(*i*) statically-scheduled super scalar;

(*ii*) out-of-order speculative execution;

(*iii*) Very Long Instruction Word (VLIW);

(*iv*) EPIC (as used by IA-64).

[12 marks]

(*b*) Discuss hardware multi-threading, and hence the different implementation approaches that have been tried to enable a single CPU core to execute from multiple instruction streams. How can multi-threading be used to improve system performance? What are the pitfalls? [8 marks]

2 Digital Communication II

(*a*) Explain the terms *Work Conservation* and *max–min fairness*, in the context of packet switching. [8 marks]

(*b*) Outline the operation of two work-conserving queueing schemes that provide max–min fairness, and two (simpler) ones that do not. [8 marks]

(*c*) Give at least two main implementation costs associated with implementations of fairness in packet switched routers. [4 marks]

3 Security

(*a*) Describe the Bell–LaPadula security policy. [6 marks]

(*b*) Describe the Chinese Wall security policy. [6 marks]

(*c*) To what extent is the Chinese Wall policy an extension of Bell–LaPadula? [6 marks]

(*d*) Are either of these policies relevant to digital rights management? [2 marks]

4 Advanced Graphics

(*a*) We want to find the first intersection point between an arbitrary ray and a sphere of arbitrary radius at an arbitrary position in space.

(*i*) List and define all of the parameters required to specify the geometry of the ray and the sphere. [2 marks]

(*ii*) Give an algorithm which returns the desired intersection point (if it exists) and the appropriate normal vector at the intersection point. [5 marks]

(*b*) Describe a method which converts an arbitrary sphere to a triangle mesh at a desired resolution. The desired resolution is specified as a desired number of triangles, D. Your method should produce a number of triangles, N, which is within an order of magnitude of D: $D/10 < N < 10D$. [4 marks]

(*c*) The Catmull-Clark bivariate subdivision scheme is a bivariate generalisation of the univariate $\frac{1}{8}[1, 4, 6, 4, 1]$ subdivision scheme. It creates new vertices as blends of old vertices in the following ways:

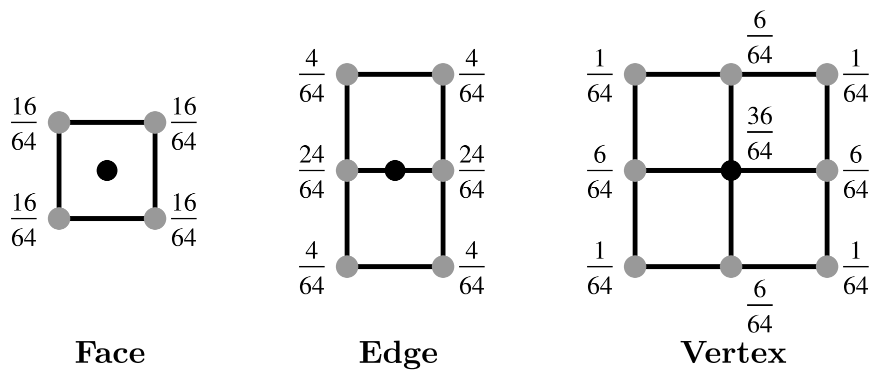

(*i*) Provide similar diagrams for the bivariate generalisation of the univariate four-point interpolating subdivision scheme $\frac{1}{16}[-1, 0, 9, 16, 9, 0, -1]$. [5 marks]

(*ii*) Explain what problems arise around extraordinary vertices (vertices of valency other than four) for this bivariate interpolating scheme and suggest a possible way of handling the creation of new edge vertices when the old vertex at one end of the edge has a valency other than four. [4 marks]

5 Computer Systems Modelling

Let $N(t)$ denote the number of events in the time interval $[0, t]$ for a (homogeneous) Poisson process of rate λ, $(\lambda > 0)$.

(*a*) State the necessary properties on $N(t)$ that define a (homogeneous) Poisson process of rate λ. [4 marks]

(*b*) By dividing the interval $[0, t]$ into equal length sub-intervals show that $N(t)$ is a Poisson random variable with mean λt. [4 marks]

(*c*) Let X_1 denote the time of the first event and for $n > 1$ let X_n denote the elapsed time between the $(n-1)$th and the nth events of the Poisson process. Determine the distribution of X_1 and the joint distribution of X_1 and X_2. [4 marks]

(*d*) Let $S_n = \sum_{i=1}^{n} X_i$ denote the time of the nth event. Derive the probability density function of the random variable $S_n(t)$. [4 marks]

(*e*) Give an algorithm to generate the first T time units of a (homogeneous) Poisson process of rate λ. [4 marks]

6 Specification and Verification I

(*a*) Explain the difference between a *variant* and an *invariant*. Briefly describe what they are used for. [4 marks]

(*b*) State and justify the verification conditions for the total correctness of `WHILE` commands. [6 marks]

(*c*) (*i*) Devise a precondition P that makes the following specification true.

```
[P]
WHILE I≤N DO SUM := SUM+(2×I); I := I+1
[SUM = N×(N+1)]
```

[2 marks]

(*ii*) Devise and justify annotations for this specification that yield provable verification conditions. [8 marks]

7 Specification and Verification II

(*a*) Describe the semantics of formulae in *linear temporal logic* (LTL) and *computation tree logic* (CTL). [2+2 marks]
Illustrate your answer by contrasting the meanings of $\mathtt{G}\,P$ in LTL with $\mathtt{AG}\,P$ in CTL (where P is a property of states). [2+2 marks]

(*b*) Give an LTL property that cannot be expressed in CTL. [2 marks]

(*c*) Give a CTL property that cannot be expressed in LTL. [2 marks]

(*d*) Describe briefly the kinds of properties that can be expressed using Sugar Extended Regular Expressions (SEREs), Foundation Language (FL) formulae and Optional Branching Extension (OBE) formulae of the *Sugar 2.0* property language. [4 marks]

(*e*) Consider the property: "whenever a, b and c occur on successive cycles, then on the cycle that c occurs, d must occur also, followed on the next cycle by e" (where a, b, c, d and e are boolean expressions). Use this property to illustrate how SEREs can sometimes help specify properties more compactly than pure LTL. [4 marks]

8 Information Theory and Coding

(a) Describe the types of data which are amenable to lossy compression with an explanation of why they are amenable to lossy compression. Describe the mechanisms that are used to perform lossy compression with an explanation of why they improve the compression rate. [8 marks]

(b) Two coding schemes are proposed for the binary coding of a four-symbol alphabet:

Symbol	Code 1	Code 2
a	00	0
b	01	10
c	10	110
d	11	111

Under what probability distributions would Code 2 be a more efficient code than Code 1? You may assume that $p(a) \geq p(b) \geq p(c) \geq p(d)$. [6 marks]

(c) For an alphabet consisting of m equiprobable symbols encoded using a binary prefix code, prove that the average length per symbol of the binary code is greater than or equal to $\log_2(m)$ bits, for any possible prefix code. [6 marks]

9 Types

(a) Give the axioms and rules for inductively generating ML typing judgements $\Gamma \vdash M : \tau$, where Γ is a finite function from type variables to type schemes and M ranges over expressions built up from variables using only function abstraction ($\lambda x(M)$), function application ($M_1\, M_2$) and local declarations (**let** $x = M_1$ **in** M_2). As part of your answer you should explain what it means for a type scheme to *generalise* a type. [7 marks]

(b) Consider the fixpoint combinator Y, which is defined to be the expression $\lambda x((\lambda y(x(y\,y)))\,\lambda y(x(y\,y)))$. State, with justification, whether there is a type τ for which $\emptyset \vdash Y : \tau$ is provable from the axioms and rules in part (a). [6 marks]

(c) Consider adding to the ML type system a 'universal' type ω together with the axiom

$$\text{(univ)} \qquad \Gamma \vdash M : \omega$$

asserting that any expression M has type ω. In this augmented type system show that $x : \omega \to \alpha \vdash \lambda y(x(y\,y)) : (\omega \to \omega) \to \alpha$ is provable, where α is any type variable. [3 marks]

Deduce that $\emptyset \vdash Y : (\omega \to \alpha) \to \alpha$ is also provable, where Y is the expression in part (b). [4 marks]

10 Topics in Concurrency

(a) Prove that if two states in a finite transition system satisfy the same assertions in the modal μ-calculus, then they are strongly bisimilar. [10 marks]

(b) Describe a Petri net semantics for the following fragment of CCS:

$$t ::= \mathit{rec}\ x\ \alpha.s \mid t_1 \parallel t_2 \mid t \setminus b$$

in which

$$s ::= x \mid \alpha.s \mid s_1 + s_2$$

where α ranges over the actions of CCS, b over non-τ actions and x over process variables.

A diagrammatic account suffices, though you should make clear the form of labelled Petri net you are using and its "token game". Although no proof is needed, your semantics should represent the independence of actions in a parallel composition and agree with the usual transition semantics of CCS. [10 marks]

11 Information Retrieval

(*a*) Why are query terms and document terms weighted in modern information retrieval systems? Describe a method for weighting terms. [6 marks]

(*b*) Describe two challenges in recognising arbitrary person names in unrestricted text. [4 marks]

(*c*) Describe how bootstrapping can be used in information retrieval. [5 marks]

(*d*) Which of the following information retrieval tasks is bootstrapping best suited for, and why?

(*i*) Recognition of person names.

(*ii*) Association of a person with his or her job title.

(*iii*) Identification of arbitrary relationship between two people, as in the following examples:

$$\text{“John did not fire Jim”} \longrightarrow \left[\begin{array}{l}\text{Relationship: Boss}\\ \text{Person 1: John}\\ \text{Person 2: Jim}\end{array}\right.$$

$$\text{“James and Jill visited their parents”} \longrightarrow \left[\begin{array}{l}\text{Relationship: Sibling}\\ \text{Person 1: James}\\ \text{Person 2: Jill}\end{array}\right.$$

[5 marks]

12 Business Studies

(*a*) Describe *three* criteria a UK patent must exhibit. [3 marks]

(*b*) What is the difference between the protection granted by a patent and that granted by copyright? Is this different in the UK from the USA? How might a computer program be protected? [5 marks]

(*c*) Explain why giving away software might be a good thing. [6 marks]

(*d*) A University Technology Transfer Office (TTO) is established to generate income from patenting and subsequently licensing intellectual property developed in a certain University. By drawing up a 5-year outline cash flow or otherwise, indicate whether this is a viable activity. How else might a University benefit from the IPR it generates? [6 marks]

13 E-Commerce

Imagine you have written a program that you want to commercialise.

(*a*) Outline a design for a web-site for the wider dissemination and/or sale of the program. Actual HTML is not required, nor details of the program. Your answer should include:

(*i*) Elements of the business model, and a description of any mechanism or legal framework needed. [5 marks]

(*ii*) A high-level site map. Outline any special features, such as registration, tracking or subscriptions that you will need to implement. [5 marks]

(*iii*) A sketch of a typical page. Explain any features you use to enhance performance. [5 marks]

(*b*) How could you market the site and drive traffic to it? [5 marks]

14 Additional Topics

(*a*) Give *three* methods of computing location of cell phones and indicate the accuracy in each case. [10 marks]

(*b*) What is the location accuracy of GPS (Global Positioning System) and how can it be improved? [5 marks]

(*c*) Give the accuracy of the ultrasonic Active Bat location system and discuss how it might compare with radio-based ultra wideband location systems. [5 marks]

15 Additional Topics

(*a*) In relation to the locational privacy problem for the Active Badge system:

(*i*) Define location privacy.

(*ii*) Define a sensible security policy for the system with respect to location privacy.

(*iii*) What elements of the system does a user need to trust?

(*iv*) What if one does not want to be tracked? [6 marks]

(*b*) In the Active Badge system, the badge emits its identifier and the building infrastructure picks it up. To protect location privacy, some have suggested to reverse this architecture: the room would transmit its identifier and the badge would pick it up. Discuss advantages and disadvantages of this arrangement. [2 marks]

(*c*) You are required to design the security architecture for a location-based system. You are the cellular phone operator, so you know the location of users; application providers selling their location-based services to users must go through you. Of course you know the position of all active phones at all times, but you want to reassure your users that application providers can't track them. State your security policy and describe your implementation that enforces it. [6 marks]

(*d*) Describe at least *two* attacks against the system you designed in part (*c*). [6 marks]

END OF PAPER