COMPUTER SCIENCE TRIPOS Part IB

Thursday 8 June 2000 1.30 to 4.30

Paper 6

Answer **five** *questions.*
No more than **two** *questions from any one section are to be answered.*
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.*

SECTION A

1 Data Structures and Algorithms

Describe an $O(n\log(n))$ algorithm based on a variation of merge sort to find the closest pair of a given set of points lying in a plane. You may assume that the set of points is given as a linked list of (x, y) coordinates. [8 marks]

Carefully prove that your algorithm can never take longer than $O(n\log(n))$. [6 marks]

Modify, with explanation, your algorithm to find the pair of points with minimum Manhattan distance. The Manhattan distance between points (x_1, y_1) and (x_2, y_2) is $|x_1 - x_2| + |y_1 - y_2|$. [6 marks]

2 Computer Design

Why are the following statements fallacies?

(*a*) MIPS is an accurate measure for comparing performance among computers. [5 marks]

(*b*) A benchmark is a typical program which accurately predicts the performance of all other applications. [5 marks]

(*c*) Complex instruction set computers minimise the semantic gap between machine code and high-level languages, thereby making applications run more quickly. [5 marks]

(*d*) Data caches always improve processor throughput. [5 marks]

3 Digital Communication I

Compare circuit switching and packet switching, paying attention to channel characteristics and resource efficiency. [7 marks]

What is *wave division multiplexing* (WDM)? Is it more like circuit switching or packet switching and why? [7 marks]

Wave length conversion is the process, either optical or optical–electronic–optical, of receiving a signal on one wavelength and transmitting on another.

How does wave length conversion ease the problem of routing optical carriers in a network? [3 marks]

"The huge capacity of WDM systems will mean that IP becomes redundant." Discuss. [3 marks]

4 Computer Graphics and Image Processing

Give an algorithm for drawing the part of a circle which lies in the first octant. Assume that the circle has integer radius and is centered at the origin. Assume that you have a function *setpixel*(x, y) which turns on pixel (x, y). [10 marks]

Derive a matrix, or a product of matrices, to perform a clockwise 2D rotation of arbitrary angle, θ, about an arbitrary point, (x_c, y_c). [4 marks]

Provide an algorithm to ascertain whether the Bezier curve defined by $P_1P_2P_3P_4$ lies within some tolerance, ϵ, of the straight line segment, $\overline{P_1P_4}$, which joins the Bezier curve's end points. Your algorithm must return *false* if the Bezier curve is outside the tolerance; it must return *true* if the curve is well inside the tolerance; it may return either *true* or *false* if the curve is inside, but not well inside, the tolerance. [6 marks]

SECTION B

5 Comparative Programming Languages

Give a brief summary of the main syntactic constructs found in the programming language Smalltalk. Other languages often have the conditional constructs if–then–else and while. Show how these two constructs can be defined in Smalltalk. [8 marks]

Illustrate the use of Smalltalk by showing how you would define a method to compute the factorial of an integer. [8 marks]

Although Smalltalk was originally designed to be an interpretive language, modern implementations are dramatically more efficient. Briefly outline what techniques might have been used to make this improvement. [4 marks]

6 Compiler Construction

Describe how a parse tree can be translated into a sequence of assembly language instructions based on a pattern matching graph derived from a set of tree rewriting rules where each rule has a cost and a corresponding fragment of code. Illustrate your answer using the following rules:

```
Ri = Kk                   LDI  Ri,Kk     Cost 2
Ri = add(Ri,Kk)           ADDI Ri,Kk     Cost 3
Ri = add(Ri,Rj)           ADD  Ri,Rj     Cost 3
Ri = add(Ri,add(Rj,Kk))   ADD  Ri,Rj,Kk  Cost 4
```

applied to the following parse tree:

```
add(K1,add(add(K2,add(K3,K4)),add(K5,K6)))
```

[15 marks]

Discuss the advantages and disadvantages of this approach to code generation. [5 marks]

7 Prolog for Artificial Intelligence

One of the regulations of the International Rugby Board (IRB) states that for a player to be eligible to play for a given country, the player's father or mother or grandfather or grandmother must have been born in that country. Assume that there is a complete genealogical database consisting of Prolog clauses of the form person(P, B, F, M), where P is a person's name, B is the country of P's birth, F is their father's name and M is their mother's name. For example, the clause

person(bruce, australia, rhodri, bronwyn).

might appear in such a database. Further assume that names in the database are constructed so as to refer uniquely to individuals. Write Prolog clauses defining the predicate eligible such that goals of the form eligible(P,C) succeed if and only if person P is eligible to play for country C according to the above regulation. [10 marks]

Given a list of players on a given country's team, define a predicate checkteam that will check each member of the team for eligibility according to the eligible predicate, and furthermore check that each player appears on the list only once. The checkteam goal will fail if any player is ineligible or if any player is listed more than once. [10 marks]

8 Databases

Describe the basic architecture of the ODMG standard for Object Data Management. [10 marks]

What support is provided for transactions? What locking modes are available, and how are they used by the database runtime systems? [4 marks]

The query language OQL is recognised as a standard by the Object Management Group (OMG). To what extent is it similar to SQL, and in what ways does it differ? [6 marks]

SECTION C

9 Semantics of Programming Languages

What does it mean to say that two configurations of a labelled transition system are *bisimilar*? [3 marks]

Describe a labelled transition system for a language of communicating processes with input prefixing ($c(x) \,.\, P$), output prefixing ($\bar{c}\langle E\rangle \,.\, P$), an inactive process ($\mathbf{0}$), choice ($P + P'$), parallel composition ($P|P'$) and channel restriction ($\nu\, c \,.\, P$). You may assume there is a relation $E \Downarrow n$ which defines when an integer expression E evaluates to an integer n. [7 marks]

For each of the following pairs of processes, say whether or not they are bisimilar. Justify your answer in each case.

(a) $\bar{c}\langle 1\rangle \,.\, ((\bar{c}\langle 2\rangle \,.\, \mathbf{0}) + (\bar{c}\langle 3\rangle \,.\, \mathbf{0}))$ and $(\bar{c}\langle 1\rangle \,.\, \bar{c}\langle 2\rangle \,.\, \mathbf{0}) + (\bar{c}\langle 1\rangle \,.\, \bar{c}\langle 3\rangle \,.\, \mathbf{0})$ [4 marks]

(b) P and $\nu\, c \,.\, ((c(x) \,.\, \mathbf{0})|(\bar{c}\langle 1\rangle \,.\, P))$, where c does not occur in P [6 marks]

10 Foundations of Functional Programming

Give as simple a set of rules as you can for transforming lambda calculus to a form where there are no bound variables mentioned, but where there are many instances of the three standard combinator constants **S**, **K** and **I**. [6 marks]

Describe tree-rewrites suitable for reducing expressions written in terms of combinators. [6 marks]

Explain how you might deal with the issue of keeping track of the values of bound variables if you were to interpret lambda calculus directly. [8 marks]

11 Logic and Proof

For each of the given pairs of terms, give a most general unifier or indicate why none exists. (Here x, y, z are variables while a, b are constant symbols.)

$$\begin{array}{lll} h(x,y,x) & \text{and} & h(y,z,u) \\ h(x,y,z) & \text{and} & h(f(y),z,x) \\ h(x,y,b) & \text{and} & h(a,x,y) \\ h(x,y,z) & \text{and} & h(g(y,y),g(z,z),g(u,u)) \end{array}$$

[4 marks]

A standard unification algorithm takes a pair of terms t_1 and t_2 and returns a substitution θ such that $t_1\theta = t_2\theta$. Show how this algorithm can be used to find the unifier of several ($n > 2$) terms t_1, t_2, ..., t_n: a substitution θ such that $t_1\theta = t_2\theta = \cdots = t_n\theta$. Indicate how the unifier is constructed from the unifiers of $n - 1$ pairs of terms. (Assume that all required unifiers exist and ignore the question of whether the unifiers are most general.) [6 marks]

Prove using resolution the formula

$$\big(\forall x\,[P(x) \leftrightarrow (Q(x) \wedge \neg Q(f(x)))]\big) \rightarrow \exists y\,\neg P(y)$$

or indicate why this formula is not a theorem. [10 marks]

12 Complexity Theory

State the hierarchy theorems for time and space. [4 marks]

A linear time reduction from a language L_1 to L_2 is a reduction that can be computed by a deterministic Turing machine in time $O(n)$.

A class of languages $\mathcal{C}$ is closed under linear time reductions if whenever $L_2 \in \mathcal{C}$ and L_1 is linear-time reducible to L_2, then $L_1 \in \mathcal{C}$.

For *each* of the following complexity classes (a) to (d), say

- whether it is closed under linear time reductions
- whether it contains problems that are complete under linear time reductions

Give full justification for your answers.

(a) $\textsc{Dspace}(n^2)$ [4 marks]

(b) L [4 marks]

(c) P [4 marks]

(d) NP [4 marks]

END OF PAPER