COMPUTER SCIENCE TRIPOS Part II (General)
DIPLOMA IN COMPUTER SCIENCE

Wednesday 2 June 2004 1.30 to 4.30

Paper 12 (Paper 3 of Diploma in Computer Science)

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 Data Structures and Algorithms

It is proposed to store a large number of records on a disk using Larsen's method so that any lookup can be done using only one disk transfer. All the records are of length 200 bytes and each contains a 20 byte key. The data is to be held on a single disk preformatted to contain 100,000,000 sectors each of size 4096 bytes. Reading multiple consecutive sectors is regarded as a single transfer.

(a) Describe Larsen's algorithm in detail and, for the records and disk specified above, state the disk block size, the signature size and the amount of main memory that you would choose to use. [10 marks]

(b) Carefully estimate the maximum number of records that could reasonably be stored on the disk assuming the sizes you gave in part (a). [6 marks]

(c) Discuss the advantages and disadvantages of different signature sizes. [4 marks]

2 Computer Design

It is possible to design a single instruction computer (SIC). For example, the instruction Subtract and Branch on Negative is sufficiently powerful. This instruction takes the form "A,B,C,D", meaning "Read A, Subtract B, Store in C, and Branch to D if negative". If a branch is not required, the address D can be set to the next instruction in the sequence so that the next instruction will be executed regardless of whether the branch is taken or not. An assembler short form for this branchless instruction is simply "A,B,C".

(*a*) Write fully commented SIC assembler which implements the following pseudo code:

```
a=1;
b=1;
for(i=1; i<n; i++) {
  a=a+b;
  b=a-b;
}
```

[9 marks]

(*b*) Reduced instruction set computers typically achieve high performance by optimising the common case. In particular, a regular and simple instruction format typically allows extensive use of pipelining. What pipeline and memory structure would you recommend in order to execute SIC code quickly? [9 marks]

(*c*) How does the density of SIC machine code compare with current commercial processors? [2 marks]

3 Digital Communication I

(*a*) Define the terms *flat* and *hierarchical* as applied to address spaces. [2 marks]

(*b*) Give *four* examples of address spaces and state whether they are flat or hierarchical, and why. [4 marks]

(*c*) Describe class-based addresses as used in the Internet. You need not worry about precise field sizes or class names. [4 marks]

(*d*) Describe classless addresses as used in the Internet. [3 marks]

(*e*) Why were they introduced? [2 marks]

(*f*) What information must be held in a routing table when classless addresses are used? [5 marks]

4 Distributed Systems

(*a*) Define strong and weak consistency. [2 marks]

(*b*) A process group manages a set of widely distributed replicas. The group is open and unstructured; that is, external processes may invoke any group member for reading or writing.

(*i*) Discuss how the replicas can be kept strongly consistent in the presence of concurrent invocations and failures. [12 marks]

(*ii*) Would it be more appropriate to use a structured group (with a single co-ordinator) to manage the replicas? Justify your answer. [6 marks]

Your solution should discuss the *selection and use* of algorithms and protocols. It is not necessary to specify them in great detail.

5 Computer Graphics and Image Processing

(*a*) We wish to produce two algorithms: one which draws the outline of a circle and one which draws a filled circle.

(*i*) Describe an efficient algorithm which will draw a one-pixel wide outline of a circle of integer radius, R, centred on the origin. [10 marks]

(*ii*) Describe the modifications required to your algorithm to make it draw a filled circle. [3 marks]

(*b*) Given a function `drawline(x1,y1,x2,y2)`, describe an algorithm for drawing a Bezier cubic curve to a specified level of accuracy using only straight lines. [7 marks]

6 Compiler Construction

You have been provided with the description of a programming language, J, intended for scripting applications. Its syntax is similar to a cut-down version of Java in that it consists of function definitions which have bodies containing if-then-else, while-do, assignments and (typed) declarations of variables. Only one statement or keyword may occur on a line so that it is sufficient to describe the start of a loop iteration with its line number. You need to explain to your boss the alternatives for implementing this so that a decision may be made as to the best implementation strategy.

The choice is between:

(a) compiling J to machine code;

(b) compiling J to "interpreted byte code", and then interpreting this;

(c) parsing J to a syntax tree representation and then interpreting this using a function which walks the tree;

(d) keeping J in a text file and then interpreting it by reading each line (and acting on it) as and when the line is required.

For *each* of (a)–(d), (i) summarise the main phases of work that are done *before* execution in each case, giving a brief explanation of the main actions of the main interpreter loop (if any) *during* execution, and (ii) for each of the following possible erroneous forms, explain whether the error would be found before or during execution: malformed syntax, undeclared variable, type error, division by zero.

[5 marks each]

You are not expected to argue for or against any of the alternatives.

7 Databases

Assume a simple movie database with the following schema. (You may assume that producers have a unique certification number, `Cert`, that is also recorded in the `Movie` relation as attribute `prodC#`; and no two movies are produced with the same title.)

```
Movie(title,year,length,prodC#)
StarsIn(movieTitle,movieYear,starName)
Producer(name,address,cert)
MovieStar(name,gender,birthdate)
```

(a) Write the following queries in SQL:

(i) Who were the male stars in the film *The Red Squirrel*? [1 mark]

(ii) Which movies are longer than *Titanic*? [2 marks]

(b) SQL has a boolean-valued operator `IN` such that the expression `s IN R` is true when `s` is contained in the relation `R` (assume for simplicity that `R` is a single attribute relation and hence `s` is a simple atomic value).

Consider the following nested SQL query that uses the `IN` operator:

```
SELECT name
FROM   Producer
WHERE  cert IN (SELECT prodC#
                FROM   Movie
                WHERE  title IN (SELECT movieTitle
                                 FROM   StarsIn
                                 WHERE  starName='Nancho Novo'));
```

(i) State concisely what this query is intended to mean. [1 mark]

(ii) Express this nested query as a single `SELECT-FROM-WHERE` query. [2 marks]

(iii) Is your query from part (b)(ii) always equivalent to the original query? If yes, then justify your answer; if not, then explain the difference and show how they could be made equivalent. [6 marks]

(c) SQL has a boolean-valued operator `EXISTS` such that `EXISTS R` is true if and only if `R` is *not* empty.

Show how `EXISTS` is, in fact, redundant by giving a simple SQL expression that is equivalent to `EXISTS R` but does not involve `EXISTS` or any cardinality operators, e.g. `COUNT`. [Hint: You may use the `IN` operator.] [8 marks]

8 Artificial Intelligence

(*a*) Describe the way in which a problem should be represented in order to allow its solution using a *heuristic search* technique. [5 marks]

(*b*) Define what it means for a search algorithm to be *complete*, and to be *optimal*. [2 marks]

(*c*) Define what it means for a heuristic function to be *admissible*, and to be *monotonic*. [2 marks]

(*d*) Describe the operation of the $A^\star$ heuristic search algorithm. [5 marks]

(*e*) Prove that the $A^\star$ heuristic search algorithm is optimal when applied in conjunction with a monotonic heuristic. State the conditions under which the algorithm is also complete, and explain why this is the case. [6 marks]

9 Numerical Analysis II

(*a*) State a recurrence formula for the sequence of Chebyshev polynomials, $\{T_k(x)\}$, and list these as far as $T_5(x)$. [4 marks]

(*b*) What is the best L_∞ polynomial approximation over $[-1, 1]$ to x^k using polynomials of lower degree, and what is its degree? Use this property to explain the method of economisation of a Taylor series. How can the error in one economisation step be estimated? [7 marks]

(*c*) It is required to approximate the function $f(x) = \lim_{k\to\infty} P_k(x)$ over $[-1, 1]$ with an absolute accuracy of 2 decimal places, where

$$P_k(x) = \sum_{n=1}^{k} \frac{x^n}{n\ n!} \ .$$

As this series converges faster than e^x, a good estimate of the error $||f(x) - P_k(x)||_\infty$ in the truncated Taylor series is given by evaluating the next term

$$\frac{x^{k+1}}{(k+1)(k+1)!}$$

at $x = 1$. Use the method of economisation to find a polynomial approximation of the required accuracy. [9 marks]

10 Introduction to Functional Programming

(a) Consider the following ML function declaration:

```
fun cat (b,f) nil     = b
  | cat (b,f) (x::xs) = f(x,(cat(b,f) xs));
```

(i) Give the type of the function `cat`. [2 marks]

(ii) Using the function `cat` define a function `filter` such that `filter p l` returns those elements in a list `l` that satisfy the predicate `p`. [3 marks]

(iii) Using the function `cat` define a function `cmap` such that `cmap f l` applies function `f` to every element in list `l`. [3 marks]

(b) Consider the following ML function declaration:

```
fun ana (p,g) b = if p(b) then
                          []
                  else    let val (a,b1)=g(b)
                          in
                            a::(ana (p,g) b1)
                          end;
```

(i) Give the type of the function `ana`. [3 marks]

(ii) Using the function `ana` define a function `zip` that converts a pair of lists into a list of pairs. [4 marks]

(iii) Using the function `ana` define a function `amap` such that `amap f l` applies function `f` to every element in list `l`. [5 marks]

11 Natural Language Processing

A context free grammar for a fragment of English is shown below:

```
S -> NP VP
NP -> Det N
N -> N N
VP -> rumbles, rusts
Det -> the, a, every
N -> bus, car, train, park, airport, station
```

(*a*) Show the parse trees for the two parses that the grammar assigns for sentence S1.

S1: the train station bus rumbles

[3 marks]

(*b*) Give an algorithm for a bottom-up passive chart parser without packing. Illustrate your answer by showing the edges constructed when parsing sentence S1. [11 marks]

(*c*) Describe how this algorithm could be modified so that edges may be *packed*, illustrating your answer by considering sentences S1 and S2. What effect does packing have on parsing efficiency?

S2: the airport car park bus rumbles

[6 marks]

12 Complexity Theory

Recall that a *simple path* in a graph is a path with no repeated nodes. Consider the following two decision problems:

- Given a graph $G = (V, E)$, a positive integer k, source $s \in V$ and a target $t \in V$, is there a simple path from s to t of length *at least* k?
- Given a graph $G = (V, E)$, a positive integer k, source $s \in V$ and a target $t \in V$, is there a simple path from s to t of length *at most* k?

One of these problems is known to be in P while the other one is known to be NP-complete.

(a) Which of the two problems is in P and which is NP-complete? [2 marks]

(b) Describe a polynomial time algorithm for the problem that is in P. [6 marks]

(c) Give a proof of NP-completeness for the problem that is NP-complete. You may assume the NP-completeness of any problem, such as *Hamiltonian Cycle*, mentioned in the lecture course. [12 marks]

END OF PAPER