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

Wednesday 7 June 1995 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.*
Write on **one** *side of the paper only.*

1 Processor Architecture

A classical RISC five-stage pipeline is depicted below:

instruction fetch	register fetch	execute	memory access	register write back

Using the above pipeline as a basis for discussion, explain the following:

(*a*) What are *data bypasses* (sometimes called feed-forward paths)? [5 marks]

(*b*) The above pipeline is likely to have two bypasses. Between which stages are the bypasses required and why? [5 marks]

(*c*) Why do load delay slots arise? [5 marks]

(*d*) Which of the following code segments will execute more quickly on the above pipeline and why (you may assume that there are no cache misses)?

Code segment 1

```
load r1,4(sp)
load r2,8(sp)
load r3,12(sp)
add r1,r2,r4 # r4=r1+r2
add r3,r4,r4 # r4=r3+r4
```

Code segment 2

```
load r1,4(sp)
load r2,8(sp)
load r3,12(sp)
add r2,r3,r4 # r4=r2+r3
add r1,r4,r4 # r4=r1+r4
```

[5 marks]

2 Designing Interactive Applications

The figure below shows the controls for a telephone answering machine, designed to meet requirements for low manufacturing cost, efficient operation and ease of learning. The controls include:

(A) a button to *start* recording (of outgoing message) or playback (of outgoing or incoming messages);

(B) a button to *stop* recording, playback or rewind;

(C) a button to *rewind* the incoming-message tape;

(D) a button to step between the four *modes*: outgoing-message-play, outgoing-message-record, incoming-messages-play and incoming-messages-record (the normal mode for awaiting calls);

(E) a display indicating the current mode;

(F) a light that flashes to show how many messages have been received.

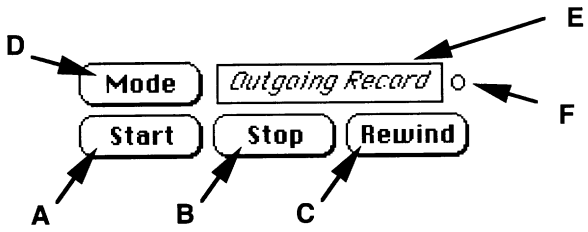

What forms of mental model are likely to be acquired by users of the answering machine, both owners (that is, recipients of calls) and callers? [5 marks]

An additional feature is proposed for the answering machine, allowing the owner to record special outgoing messages, each one to be played when calls are received from an owner-specified telephone number. To record such a message, the owner of the machine sets it in outgoing-message-record mode, keys in the telephone number in question using the attached telephone's push-buttons, and records the message in the normal way, using the *start* and *stop* buttons. Up to 16 special outgoing messages, each for a different incoming telephone number, can be recorded in this manner.

You have been asked to review this proposed new feature, in terms of the owner's ability to learn to use the machine in an error-free manner. What problems can you identify in the feature's design? Propose solutions that are consistent with the original requirements. [15 marks]

3 Digital Communication I

Compare the functions of a *MAC level bridge* with an *IP router*. In what circumstances is it more appropriate to use one than the other? [6 marks]

Discuss the tables inside both bridges and routers used to control the acceptance and forwarding of packets. Indicate both how these tables are used and how information is put in the tables. How quickly can the tables be searched in each case? [10 marks]

What features might be added to a router or bridge to improve some aspects of network security? [4 marks]

4 Graphics

Why are matrix representations used to describe point transformations in computer graphics? [6 marks]

Describe how to represent three different 2D transformations as matrices. [6 marks]

Explain how to derive a sequence of transformations to achieve the overall effect of performing a 2D rotation about an arbitrary point. [8 marks]

5 Developments in Technology

Write a short essay on the developments in electronic computers during the 1940s and 1950s. Give attention to the emergence of key ideas about machine architecture in this period, such as conditional branching and program interrupt. [20 marks]

6 Programming in C and C++

Students have been set a programming exercise where they are expected to write a subroutine to perform a certain operation, and they have been told that they should write it in the language C. One of the offerings submitted for assessment is the following. The student involved, who had read a book by Dijkstra but not understood it, proudly proclaims that of course this program has not actually been tried on a computer.

Explain what the program is (probably) supposed to do, and identify as many problems with it as you can, given that in this examination questions are expected to be completed in around 30 minutes.

```
typedef unsigned long int thing;

#define swap(p,q) v = p; p = q; q = v;

void fast(thing a[], int left, int write)
{                     /**********************
    int i, j;          * Declare variobles! *
    thing v;           **********************/
    if (write-left > 1)
    {   v = a[write];
        i = left, j = write;
        for (;;)
        {   while (a[++i] < v};
            while (a[--j] >= v);
            if (i < j) swap(a[i], a[j]);
            else break;
        }
        fast(a, left, i-1);  // re-curse here.
        fast(a, i, write)
    }
}
```

[20 marks]

7 Compiler Construction

Give a BNF syntax for unsigned floating point constants that would allow constants such as

```
1.2          3E4          0.23E-8
```

but not

```
.123         20.          21.E4
```

[6 marks]

Construct a finite state machine that could be used in a lexical analyser to read in such floating point numbers. [7 marks]

Show how this finite state machine can be extended to evaluate the floating point number it reads. [7 marks]

8 Prolog for Artificial Intelligence

Electricity consumers are supplied with electricity from an electricity generating station. Electricity is distributed from the station to the various consumers through a network of transformers as shown in this diagram:

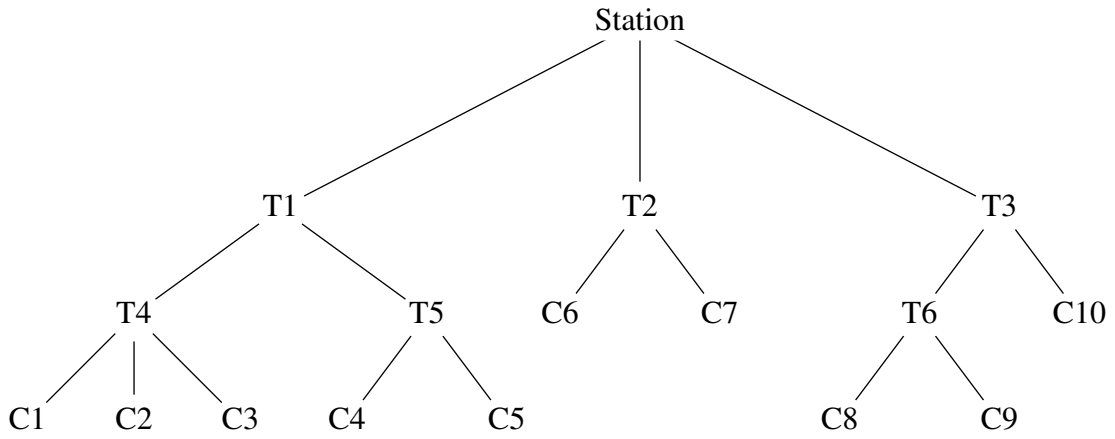

Nodes C1 to C10 are consumers; nodes T1 to T6 are transformers. Each consumer has a direct connection to only one transformer. Sometimes a transformer may malfunction or need to be taken out of service temporarily.

Devise a data structure in Prolog to represent networks like this. [5 marks]

Write a Prolog procedure to define predicate `supplies` such that goal `supplies(X,Y)` succeeds when there is an electricity supply from node `X` to node `Y` in the network. [10 marks]

Explain how to enhance your representation to permit multiple connections which can be used to make a supply around a transformer which has been taken out of service. [5 marks]

Note: The `supplies` predicate may contain arguments in addition to `X` and `Y` at your discretion.

9 Databases

What forms of error do database normal forms seek to render less probable? [3 marks]

Define Second Normal Form. [3 marks]

Define Fourth Normal Form. [4 marks]

Give an example of a database schema that is in 2NF form but *not* in 4NF, showing how it is possible for mutually contradictory data to be recorded. Explain your assumptions about the semantics underlying the database schema. [7 marks]

Are there applications where normal forms are unnecessary or even unhelpful? [3 marks]

10 Topics in Artificial Intelligence

Using appropriate mathematical expressions, define the following operations commonly used in computer vision and briefly explain their function and applications:

(a) convolution

(b) correlation

(c) bandpass filtering

(d) edge detection by derivative zero-crossings

(e) invariant transform

[20 marks]

11 Computer Systems Modelling

Given that in a balanced system with K devices and N customers, the utilisation of each device is given by

$$U = \frac{N}{N + K - 1}$$

derive a formula for the response time in terms of throughput, the number of devices and the average service demand at each device. [6 marks]

A system consists of three types of devices, A, B and C. Customers require service at each type of device but do not care at which particular device they are served. The numbers of each type of device and average service requirements per customer are

	number of devices	average service demand
A	48	48 ms
B	24	24 ms
C	18	18 ms

so that, for example, a customer requires on average 48 ms of service at a type A device.

Give bounds for the system response time at a throughput of 500 customers per second when a scheduling policy ensures that

(a) no device is more than 1.5 times as busy as the average for devices of the same type

(b) no device is more than 1.8 times as busy as the average for devices of the same type

[9 marks]

What can you say about response time if no limit on utilisation skew across devices of the same type is guaranteed? [5 marks]

12 Introduction to Functional Programming

Consider the ML definitions:

```
fun I x = x
fun curry f x y = f (x,y)
fun uncurry f (x,y) = f x y
```

What are the types of `curry` and `uncurry`? [2 marks]

Recall that f `o` g is the function that maps x to $f(g(x))$. Describe the effect of the following functions:

```
curry (fn(x,y) => x)
uncurry o curry
curry I
uncurry I
```

[4 marks]

Infinite lists can be represented in ML by functions. A function f represents the infinite list $f(0), f(1), f(2), \ldots$

(a) Give a representation for the infinite list $0, 2, 4, \ldots$ [2 marks]

(b) Code in ML a map functional for this representation; given a function f and the infinite list $x_0, x_1, x_2, \ldots$, it should yield the representation of $f(x_0), f(x_1), f(x_2), \ldots$ [2 marks]

(c) Code in ML a drop function, which given an integer $i \geqslant 0$ and an infinite list $x_0, x_1, x_2, \ldots$ returns the infinite list $x_i, x_{i+1}, x_{i+2}, \ldots$ [2 marks]

(d) Code in ML an interleave function, which combines the infinite lists $x_0, x_1, x_2, \ldots$ and $y_0, y_1, y_2, \ldots$ to yield $x_0, y_0, x_1, y_1, \ldots$ [3 marks]

(e) Code in ML a filter function, which given a predicate p and an infinite list $x_0, x_1, x_2, \ldots$ returns the infinite list obtained by deleting each x_i for which $p(x_i)$ is `false`. [5 marks]

13 Complexity Theory

Comment on each of the following statements about Computational Complexity. Indicate any places where their wording is not sufficiently precise. For each, decide whether the statement is true, partially true, true but improperly justified, false or just muddled.

(a) Given a variant on Quicksort that uses a median-of-three procedure to select pivots, it seems hard to identify exactly what ordering of input data will make the quicksort behave worst. But because there are $N!$ possible different orderings and $N!$ is a bit like 2^N the problem is an NP one.

(b) If we could solve the Boolean Satisfiability problem efficiently we could use that to simulate the behaviour of *any* Turing machine and hence solve all other problems efficiently.

(c) To keep your secret treasure safe you intend to dig out a series of caves and tunnels forming a maze. You invent a graph for which you know a Hamiltonian circuit (for example, you start by putting in edges to make that circuit and then add lots more to make the graph more complicated). The Hamiltonian circuit problem is known to be NP complete, so given just the graph nobody except you will be able to find the Hamiltonian circuit. You wire up the tunnels so that the treasure can only be reached (safely!) by traversing the Hamiltonian circuit. Being arrogant you then pin details of the graph on your door. NP completeness means that your treasure is almost certainly secure.

You should provide a brief overview of any result, construction or proof that you refer to, but you are not expected to work through the details.

[20 marks]

14 Numerical Analysis II

A cubic spline $\phi(x)$ is defined over $[a, b]$ with knots x_1, x_2, ... x_n such that $a < x_1$, $x_n < b$. The spline takes the values y_1, y_2, ... y_n at the knots. What continuity conditions are usually imposed on the cubic spline at each knot? [2 marks]

If $d_j = x_{j+1} - x_j$ and $\mu_j = \phi''(x_j)$, the spline has the following formula for $x \in [x_j, x_{j+1}]$

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

By differentiating this formula:

(a) find formulae for $\phi'(x_j)$ and $\phi'(x_{j+1})$ for $x \in [x_j, x_{j+1}]$ [4 marks]

(b) verify that $\phi''(x_j) = \mu_j$, $\phi''(x_{j+1}) = \mu_{j+1}$ [4 marks]

(c) deduce the equation which expresses the continuity condition on $\phi'(x)$ at x_j [6 marks]

If the equations derived in part (c) are solved as a simultaneous system, what are the unknowns? If the end conditions specify the spline to be linear in $[a, x_1]$ and $[x_n, b]$ how does this simplify the calculation? State the most important properties of the resulting equations. [4 marks]