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

Monday 3 June 1996 1.30 to 4.30

Paper 10 (Paper 1 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 Continuous Mathematics

We compute the representation of some continuous function $f(t)$ in a space spanned by an orthonormal family $\{\Psi_j(t)\}$ of continuous basis functions by projecting $f(t)$ onto them. We express these projections in bracket notation $<f(t), \Psi_j(t)>$ denoting $\int_{-\infty}^{\infty} f(t)\Psi_j(t)dt$, and $f(t)$ is assumed to be square-integrable.

(*a*) Give an expression for computing $f(t)$ if we know its projections $< f(t), \Psi_j(t) >$ onto this set of basis functions $\{\Psi_j(t)\}$. Explain what is happening. [5 marks]

(*b*) Now give an expression for computing $f^{(n)}(t)$, the $n$th derivative of $f(t)$ with respect to $t$, in terms of the same projections and continuous basis set. (You may assume the existence of all derivatives.) Explain your answer. [5 marks]

(*c*) Now consider a linear, time-invariant system with impulse-response function $h(t)$, having time-varying input $s(t)$ and time-varying output $r(t)$:

$$s(t) \longrightarrow \boxed{h(t)} \longrightarrow r(t)$$

In the case that the input is the complex exponential $s(t) = \exp(i\mu_j t)$ (where $i = \sqrt{-1}$ and $\mu_j$ is a constant), what can you say about the output $r(t)$ of such a system? [5 marks]

(*d*) If the input $s(t)$ has been represented in terms of a set of complex exponentials $\Psi_j(t) = \exp(i\mu_j t)$ as described at the beginning of this question, is it possible for *different* complex exponentials (not included in this set) to appear in the output $r(t)$ when it too is represented in terms of complex exponentials? Justify your answer. [5 marks]

## 2 Modula-3

Outline the syntax of a procedure *signature*. [4 marks]

Describe the different *modes* that a formal parameter may have. [4 marks]

Explain the operation of the following program and include the output in your account.

```
MODULE Chop EXPORTS Main;

IMPORT IO, Fmt;

CONST
  y = '0';      (* Character zero *)

VAR
  x := 'Z'; z := 'A';

PROCEDURE p(VAR x : CHAR) =
  BEGIN
    FOR z := y TO x DO
      IO.Put(Fmt.Char(z))
    END;
    IO.Put("\n");
    x := VAL( (ORD(x)+ORD(y)) DIV 2, CHAR)
  END p;

PROCEDURE q(z : CHAR) =
  CONST
    x = '?';
  BEGIN
    REPEAT
      p(z);
      IO.Put(Fmt.Char(x) & Fmt.Char(y) & Fmt.Char(z) & "\n")
    UNTIL y = z
  END q;

BEGIN
  q('9');
  IO.Put(Fmt.Char(x) & Fmt.Char(y) & Fmt.Char(z) & "\n")
END Chop.
```

[12 marks]

## 3 Further Modula-3

A Modula-3 library module is required to provide a buffered stream of characters between two threads, with appropriate synchronization.

Write a `Buffer` interface defining an opaque object type, `T`, with the following methods:

> `init`: taking an argument specifying the buffer size and returning a suitably initialized object
>
> `put`: inserting a single character into the buffer
>
> `get`: extracting a single character from the buffer

[4 marks]

Sketch a corresponding implementation giving

(*a*) the concrete revelation of the types [4 marks]

(*b*) the three default methods [4 marks each]

.

## 4 Compiler Construction

Describe what is meant by a *phrase structured grammar* and a *context free grammar*. [3 marks]

Describe an algorithm to calculate the set $L_T(P)$ of all terminal symbols that can start a string derived from the non-terminal $P$ using one or more productions of a given context free grammar. Illustrate your answer by calculating $L_T$ sets for the following grammar:

```
S -> U V

V -> + U V | ϵ

U -> X W

W -> * X W | ϵ

X -> ( S ) | n
```

[6 marks]

Describe an algorithm to calculate the set `FOLLOW(P)` of terminal and non-terminal symbols for a given context free grammar, where

$$\texttt{FOLLOW(P) = \{ X | S} \overset{*}{\texttt{=>}} \texttt{... P X ... \}}$$

i.e. all symbols that can follow `P` in a sentential form derived from the sentence symbol `S`. Illustrate your answer by calculating the `FOLLOW` sets for the grammar given above. [6 marks]

Outline possible ways in which the space used by the Action and Goto matrices of an SLR(1) parser can be reduced. [5 marks]

## 5 Data Structures and Algorithms

Describe and justify Dijkstra's algorithm for finding the shortest path between two vertices in a directed graph with non-negative lengths associated with its edges. [8 marks]

How can this algorithm be extended to consider graphs with some negative lengths? [6 marks]

By considering the graph on $\{A, B, C\}$ with $A \to B$ having length $-2$, $B \to A$ having length 1 and $A \to C$ having length 1, or otherwise, show that the "shortest path" is not always well defined if there are negative lengths. When is it well defined? [6 marks]

## 6 Data Structures and Algorithms

Describe and justify the Graham scan algorithm for finding the convex hull of a set of points in the plane. [8 marks]

How does its cost depend on the number of points? [4 marks]

Give a technique for heuristically eliminating a number of points before doing the scan. In what circumstances can the heuristic fail to help and what would you do about it? [8 marks]

## 7 Operating System Foundations

What is meant by *memory-mapped I/O*? How might devices be protected from uncontrolled access? [3 marks]

Describe how I/O might be programmed with and without interrupts enabled. Discuss the implications of the use of these approaches within an operating system. Discuss how interrupt handling might be supported by hardware. [5 marks]

Why does a cache alleviate the possible performance penalty of using DMA? [3 marks]

Give an example involving interrupt-driven software which illustrates the danger of race conditions. [5 marks]

Outline how semaphore operations can be made atomic. [4 marks]

**8 Operating System Functions**

A computer with a 32-bit virtual addressing scheme uses pages of size 4 Kbyte. Describe, with the aid of diagrams, two practical schemes for managing its virtual address space, comparing them with regard to speed of access, efficiency (of space), and ease of memory sharing between processes. [10 marks]

A Winchester-style disc has its head currently located at track 100, and the head is moving towards track 0. Given the reference string (27, 129, 110, 186, 147, 41, 10, 64, 120, 11, 8, 10) representing the (ordered) sequence of requests for disc tracks, give the sequence of disc addresses visited by the disc head under the SSTF, SCAN and C-SCAN disc scheduling algorithms. In each case briefly describe the algorithm, and compute the average cost of a disc access in terms of the mean number of tracks traversed per access.

In what way is each of these algorithms biased in its service of disc requests? Describe an algorithm which reduces the bias. [10 marks]

**9 Mathematics for Computation Theory**

Let $S$ be a finite alphabet, $\mathcal{E}$ be an algebra of events over $S$, and $M_{mn}$ be the algebra of $(m \times n)$ event matrices over $\mathcal{E}$. For $M$, $N \in M_{mn}$ define

$$M \leqslant N \quad \text{if and only if} \quad M + N = N$$

so that $(M, \leqslant)$ is a partially ordered set.

Let $A, B$ be $(m \times m)$, $(m \times n)$ event matrices over $\mathcal{E}$. Prove that $X = A^*B$ is the least $(m \times n)$ event matrix such that

$$X = B + AX \qquad (*)$$

stating clearly any algebraic assumptions on which your proof depends (Arden's rule). [10 marks]

Suppose now that $A = (A_{ij} \mid 1 \leqslant i, j \leqslant m)$ is such that no event $A_{ij}$ contains the null string. Prove that $X = A^*B$ is the *unique* solution of $(*)$.

[Hint: suppose if possible $X > A^*B$ is a solution of $(*)$. Let $(X - A^*B) = Y = (Y_{ij})$, and choose string $y$ of minimal length across all $Y_{ij}$.] [10 marks]

## 10 Computation Theory

Show how to code the program and initial data for an $n$-register machine into natural numbers $p$ and $d$. In what sense do the codes $p$ and $d$ determine a unique computation? [9 marks]

Using your codes establish a precise statement of the Halting Problem for Register Machines. [3 marks]

Assume that the Halting Problem is in general undecidable. Prove that it cannot be decided whether a general program $p$ terminates when the initial data is zero in every register. [8 marks]

## 11 Numerical Analysis I

Explain how the 32 bits are arranged to store the *sign*, *exponent*, and *significand* in a single precision number under the IEEE binary floating-point standard (IEEE 754). How is the exponent stored? [2 marks]

Explain the terms $e_{max}$, $e_{min}$, *normalized number*, *denormal number*, *hidden bit*, $NaN$. [5 marks]

In terms of the stored bit-pattern, how can each of the following be recognized: $\pm 0$, $\pm\infty$, *normalized number*, *denormal number*, $NaN$? [5 marks]

Suppose the principles of IEEE arithmetic are applied to a floating-point representation with only 6 stored bits. If the precision $p$ is 3 (including the hidden bit) and $e_{min}$ is $-2$, what is $e_{max}$? [2 marks]

List the 16 bit-patterns which do not represent normalized numbers, and identify what each represents. [6 marks]

## 12 Unix Case Study

For a Unix filing system:

(*a*) Describe the organization of the metadata for a filing system held on a single device or partition of a device. Outline any advantages and disadvantages of the approach. [4 marks]

(*b*) Explain how the online filing system is constructed from the separate filing systems you have described. [2 marks]

(*c*) Describe how hard links and soft (symbolic) links are implemented. Mention any restrictions on their use and any implications of their existence on other file operations. [6 marks]

(*d*) Describe how the disc blocks used to store a file or directory are recorded. Discuss whether this method is suitable for a large video file. [2 marks]

(*e*) Describe how access control policy is recorded and how access by program, as well as by user, is supported. Discuss this latter mechanism. [6 marks]

## 13 Digital Electronics

An electronically controlled mousetrap has detectors to indicate when the trap door is entirely open ($D$), entirely closed ($C$), when there is something in the trap ($S$) and when the thing in the trap is heavier than a mouse ($H$). A single control ($T$) is used for opening ($T$=1) or closing ($T$=0) the trap door. Unfortunately the $H$ signal is valid only when the door is shut.

Draw a state diagram for the system taking care not to trap the cat. [8 marks]

Provide equations for the output and controls of JK flip flops to implement a machine with the state diagram below (A and B are inputs, X is the output).

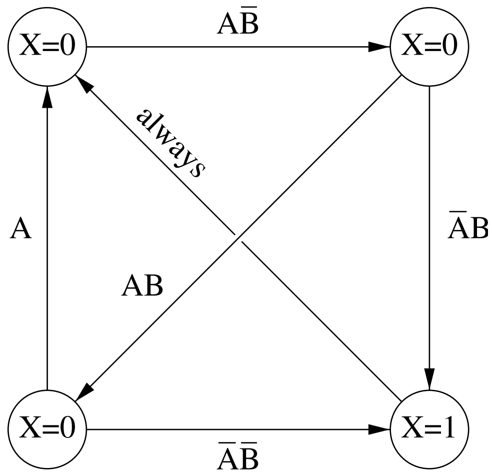


[8 marks]

Simplify the following state diagram:

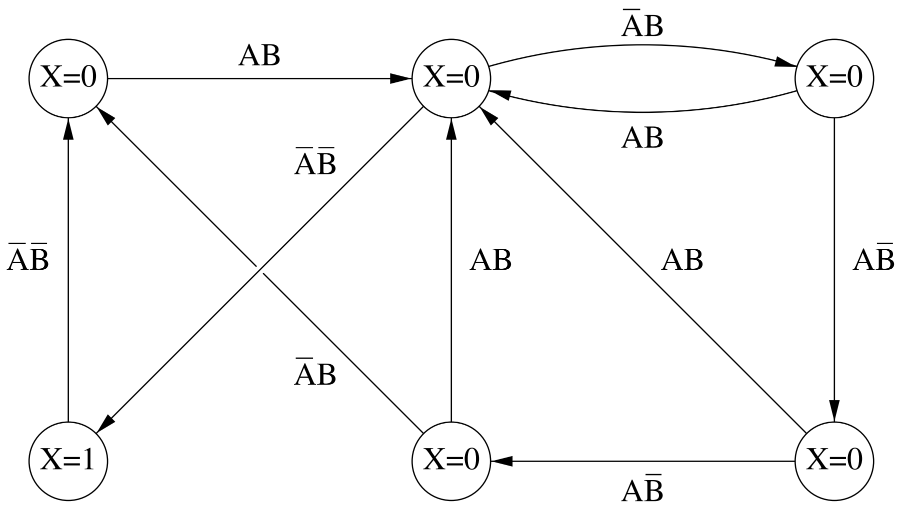


[4 marks]