

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

---

Monday 2 June 1997 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. 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.*

### 1 Continuous Mathematics

Suppose that some continuous function  $f(x)$  has Fourier Transform  $F(\mu)$ . Now consider the consequences in the Fourier domain of each of the following operations upon  $f(x)$ :

- (a) What will be the Fourier Transform of the  $n^{\text{th}}$  derivative of  $f(x)$  with respect to  $x$ :  $\frac{d^n}{dx^n} f(x)$ ? [4 marks]
- (b) What will be the Fourier Transform after shifting  $f(x)$  by a distance  $\alpha$ :  $f(x - \alpha)$ ? [4 marks]
- (c) What will be the Fourier Transform after dilating  $f(x)$  by a factor  $\alpha$ :  $f(x/\alpha)$ ? [4 marks]
- (d) Suppose that  $f(x)$  is convolved with another function  $g(x)$  whose Fourier Transform is  $G(\mu)$ . What will be the Fourier Transform  $H(\mu)$  of the convolution result  $h(x)$  where

$$h(x) = \int_{-\infty}^{+\infty} f(\alpha)g(x - \alpha)d\alpha?$$

[4 marks]

- (e) Suppose now that a two-dimensional continuous function  $f(x, y)$  has a 2D Fourier Transform  $F(\mu, \nu)$ . Define the Laplacian operator  $\nabla^2$  and express the 2D Fourier Transform of  $\nabla^2 f(x, y)$ . [4 marks]

## 2 Modula-3

The following code has been extracted from a Modula-3 program written to maintain information about the collection of bicycles and tricycles in a transport museum:

```

TYPE
  Cycle = OBJECT
    year   : CARDINAL;
    weight : REAL
  METHODS
    info () : TEXT      := NoInfo
  END;

  TriCycle = Cycle OBJECT
    dl : REAL;
    df : REAL;
    dr : REAL
  OVERRIDES
    info      := Tinfo
  END;

PROCEDURE Tinfo (c : TriCycle) : TEXT =
  BEGIN
    RETURN "Year: " & Fmt.Int(c.year) &
           "  Weight: " & Fmt.Real(c.weight) &
           "  Mean Diameter: " & Fmt.Real((c.dl+c.df+c.dr)/3.0)
  END Tinfo;

VAR
  rudge := NEW (TriCycle, year := 1927, weight := 30.6,
                dl := 0.55, df := 0.66, dr := 0.77);

```

Supply a suitable procedure `NoInfo` and then explain the code when this fragment is augmented by your new procedure. The explanation should particularly describe the object `rudge`. [8 marks]

Explain the operation and effects of each of the following `IO.Put` statements:

```

IO.Put (rudge.info () & "\n");
IO.Put (TriCycle.info (rudge) & "\n");
IO.Put (Cycle.info (rudge) & "\n");
IO.Put (NARROW (rudge, Cycle).info () & "\n");

```

[12 marks]

### 3 Further Modula-3

The thread system in Modula-3 uses mutexes and condition variables to control concurrency. An alternative scheme would be to use message passing.

Each message-passing channel would be an object with two methods:

```
send (message: REFANY)
```

and

```
receive (): REFANY
```

plus an initialisation method.

Calling `send` would append the argument to a (fixed-length) FIFO queue of data which would be collected by the thread calling `receive`. The thread would block if `receive` were called when the queue was empty or if `send` were called when the buffer was full.

Write an interface, `Message`, defining a suitable opaque object type with `init` and `send` methods, and write a further interface `MessageInternal` revealing the `receive` method. [8 marks]

Write an implementation exporting to both of these interfaces supplying the concrete revelations of the types and providing appropriate default methods. [12 marks]

## 4 Compiler Construction

Consider the following grammar for expressions ( $\langle E \rangle$ ) and commands ( $\langle C \rangle$ ).

$$\begin{aligned} \langle E \rangle & ::= i \mid n \mid \langle E \rangle - \langle E \rangle \mid \langle E \rangle ** \langle E \rangle \mid ( \langle E \rangle ) \\ \langle C \rangle & ::= i := \langle E \rangle \\ & \quad \mid \text{if } \langle E \rangle \text{ then } \langle C \rangle \mid \text{if } \langle E \rangle \text{ then } \langle C \rangle \text{ else } \langle C \rangle \\ & \quad \mid \langle C \rangle \text{ repeatwhile } \langle E \rangle \mid \langle C \rangle ; \langle C \rangle \mid \{ \langle C \rangle \} \end{aligned}$$

Show that there are syntactic ambiguities between (a) the minus (-) and exponentiation (\*\*) operators, (b) the if-command and the if-then-else-command, and (c) the if-then-else-command and the repeatwhile-command.

[4 marks]

Define, in a programming language notation of your choice, a recursive descent parser that will construct the abstract syntax tree for an input stream conforming to the above syntax for commands. You may assume the existence of a function `lex()` that will yield an integer representing the next lexical token from the input stream, and the functions `mk2(op,x)`, `mk3(op,x,y)` and `mk4(op,x,y,z)` that will construct abstract syntax tree nodes with a given operator and one, two or three operands. You should assume that exponentiation is right associative and more binding than subtraction which is left associative. The command following `then` should be the longest possible and the command before `repeatwhile` should be the shortest possible.

[12 marks]

Briefly outline how you would modify your parser if the command to the left of `repeatwhile` was changed to be the longest (rather than the shortest). [4 marks]

## 5 Data Structures and Algorithms

You are given two sets of numbers  $X_i$  and  $Y_j$ , where  $i$  and  $j$  run from 1 to  $N$ .

Devise an algorithm to find the  $M$  largest values of  $X_i - Y_j$ . This algorithm should not be quadratic in  $N$ , though it is permitted to be quadratic in  $M$ .

You should regard  $N$  as being of the order of 20,000 and  $M$  as being of the order of 1,000.

[20 marks]

## 6 Data Structures and Algorithms

Describe and justify Kruskal's algorithm for finding the minimum spanning tree of an undirected graph. [6 marks]

Suppose that all edges longer than some given  $L$  were omitted from the graph, for example as a result of not calculating them at all. Would the algorithm still give the correct result? Would you be able to tell if it had not? If it yielded a tree would it be guaranteed to be the best one? Justify your answers and consider the problem of finding a minimum spanning tree for 1,000,000 points in a plane rectangle where there is an edge between every pair of points and the cost of the edge is the Euclidean distance between the two points. [14 marks]

## 7 Operating System Foundations

In a given application, processes may be identified as "readers" or "writers". Multiple-reader, single-writer access to data is required to be implemented with priority for writers. Readers execute procedures `startread` and `endread` before and after reading. Writers execute `startwrite` and `endwrite`.

Give pseudocode for the `startread` procedure using

- (a) semaphores
- (b) conditional critical regions
- (c) monitors
- (d) guarded commands

Explain via comments in your pseudocode how concurrency control is achieved in each case and criticise each approach from the viewpoint of the programmer and the implementer.

[20 marks]

## 8 Operating System Functions

A computer is equipped with a CPU with a 32-bit virtual address space, a 32-entry TLB with access time 10ns and 32 Mbyte DRAM with access time 100ns. Its secondary storage is provided by an IDE hard disc with transfer rate 1 Mbyte/s, and which rotates at 3600 revolutions per minute and has an average seek time of 10ms. The computer uses demand paged virtual memory with 1 kbyte pages.

- (a) Explain the function of the TLB with the aid of a diagram. [4 marks]
- (b) Design a page table structure which the operating system can use to implement virtual memory on this system and describe how a virtual address is translated using it. Are there any drawbacks of your approach? [4 marks]
- (c) What is *demand paging*? Briefly describe a policy which the operating system can use to share the available physical memory between competing processes. [3 marks]
- (d) Calculate the effective memory access time for the system if virtual memory is managed using the scheme described in your answer to (b), if page tables are kept locked in memory, the probability of finding a translation in the TLB is 98%, and the probability of a page fault is  $10^{-6}$ . Exclude operating system overhead from your calculation, and briefly explain your answer. [9 marks]

## 9 Mathematics for Computation Theory

- (a) Let  $A$  be a set of basic symbols,  $\Phi = \{\phi_i\}_{i=1}^k$  be a finite set of operators, each of finite arity  $r_i = ar(\phi_i)$ . Define the language  $(A, \Phi)$  of algebraic expressions over  $A$ . [4 marks]

State the Principle of Structural Induction for a predicate  $P$  defined on  $L = (A, \Phi)$ . [4 marks]

- (b) Let  $S = \{\vdash, X\}$  be an alphabet of two symbols. The language  $L \subseteq S^*$  is defined as follows:

- (i)  $\vdash X \in L$ ;
- (ii) if  $\vdash y \in L$  for some string  $y \in S^*$ , then also  $\vdash yy \in L$ ;
- (iii) if  $\vdash y \in L$  for some string  $y \in S^*$ , then also  $\vdash yyy \in L$ ;
- (iv) if  $\vdash yXXXXX \in L$  for some string  $y \in S^*$ , then also  $\vdash y \in L$ ;
- (v) no other string  $\in L$ .

Determine whether the following strings belong to  $L$ :

- (1)  $\vdash XXXXXXXX$ ;                      (2)  $\vdash XXXXX$ .

[Hint: it may be helpful to consider a predicate of the form “the number of  $X$ s is not divisible by  $p$ ” for a suitable prime  $p$ .] [12 marks]

## 10 Computation Theory

Explain the action of a Turing machine, and show how the progress of a computation may be tracked by maintaining a record of the configuration at each time  $t$ . Prove that a computation which enters the same configuration twice will not terminate.

[8 marks]

Suppose you are given a Turing machine  $T$  having  $r$  states and  $k$  symbols. It is known that in a particular computation the head moves on the tape so that it is never more than  $l$  squares from its starting point. Calculate a bound on the number of configurations that the machine may enter during the computation.

[4 marks]

State a precise form of the unsolvability of the HALTING problem for Turing machines. Assuming this result, show that it is not possible to compute a bound on the distance of the head from its starting position during HALTING Turing machine computations.

[8 marks]

[You may assume that a Turing machine computation may be characterised by codes  $q$  and  $d$  that specify the quintuplet description and initial tape contents uniquely.]

## 11 Numerical Analysis I

Define *absolute error* and *relative error*. How are they related?

[2 marks]

Consider the quadratic expression  $Q(x) = ax^2 + bx + c$  in which  $a$ ,  $b$ ,  $c$  and  $x$  are all represented with the same relative error  $\delta$ .

In computing  $bx$ , estimate the worst-case relative error, and hence the worst-case absolute error.

[2 marks]

Now estimate the worst-case absolute error in computing  $Q(x)$ .

[4 marks]

Comment on the suitability of the formula

$$x = \frac{-b + \sqrt{b^2 - 4ac}}{2a}$$

for computing one of the roots of  $Q(x)$  in floating-point arithmetic. Derive an alternative formula and describe how it could be used in practice. Illustrate your answer by applying it to the case  $a = 4$ ,  $b = 400$ ,  $c = 7$  on a decimal machine with only 4 significant digits available. [You may assume  $\sqrt{b^2 - 4ac} \simeq b - (2ac/b)$  for the purposes of calculation.]

[12 marks]



## 12 Unix Case Study

Explain how the buffer-cache is used by the Unix file system. [10 marks]

Indicate why this approach might cause problems for applications which are concerned with the management of persistent data. [1 mark]

Give examples of when a process needs to *sleep* while accessing the buffer-cache and in each case indicate the event that causes the corresponding *wakeup* to be executed. [6 marks]

Explain the interaction of wakeup and scheduling. [3 marks]

## 13 Digital Electronics

A round robin arbiter attempts to allocate a resource in a cyclic fashion, but remembers only the last (or current) user of a resource. Outline the operation of a four-input round robin arbiter, describing the inputs, outputs and a sequence which shows its operation. [5 marks]

Draw the state diagram for a four-input round robin arbiter. [8 marks]

Derive equations for an implementation of a four-input round robin arbiter which uses J-K flip flops. *Do not minimise these equations.* [7 marks]