

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

Monday 4 June 2001 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.*

**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 Digital Electronics

- (a) Give the truth table for the exclusive or (XOR) of two boolean variables. [2 marks]
- (b) Express XOR in disjunctive normal form. [3 marks]
- (c) XOR gates are sometimes called *half-adders*. Show how two XOR gates can be used to find the low-order bit of the sum of three one-bit numbers. [3 marks]
- (d) If the XOR gates in part (c) were composed of gates corresponding to the expression in part (b), what would the delay of the circuit be? (Assume that inverters, OR gates and AND gates have a delay τ .) [4 marks]
- (e) Give the transistor level circuit diagram for a CMOS NOR gate and explain its operation. [4 marks]
- (f) Give a transistor level circuit diagram for a CMOS XOR gate. [4 marks]

2 Foundations of Programming

- (a) Briefly describe the operation of each of the following Java operators:

& | ^ << >> >>> [6 marks]

- (b) On what primitive types do these operators operate? [2 marks]

- (c) A Java program which is designed to exercise some of these operators begins thus:

```
public class Changes
{ public static void main(String[] args)
  { int n = 0x87654321;
    System.out.println(Integer.toHexString(cross(n)));
    System.out.println(Integer.toHexString(one8(n)));
  }
}
```

The method `cross()` accepts an eight-digit hexadecimal number (for example `0x87654321`) as an argument and interchanges the digits in pairs (returning, in this example, `78563412`).

The method `one8()` also accepts an eight-digit hexadecimal number (such as `0x87654321`) as an argument but interchanges only the three inner pairs of digits (thereby returning, in this example, `86745231`).

Supply the methods `cross()` and `one8()`. [8 marks]

- (d) Write a method `msb()` which accepts an `int n` and returns 0 if `n` is zero but otherwise returns an eight-digit hexadecimal number with a single bit set at the position of the most significant 1 bit in `n`. [4 marks]

3 Further Java

- (a) Describe the differences and similarities between abstract classes and interfaces in Java. How would you select which kind of definition to use? [5 marks]
- (b) An enthusiast for programming with *closures* proposes extending Java so that the following method definition would be valid:

```
Closure myCounter (int start) {
    int counter = start;
    return {
        System.out.println (counter ++);
    }
}
```

The programmer intends that no output would be produced when this method is executed, but that it would return an object of a new type, `Closure`, and that invoking `apply()` on that object would cause successive counter values to be printed.

By using an *inner class* definition, show how this example could be re-written as a valid Java program. [5 marks]

- (c) A common programming mistake is to try to define a class to have more than one superclass. For example a naïve programmer may write

```
class AutoTree extends Thread, BinaryTree {
    ...
}
```

when defining a new kind of data structure which uses an additional thread to keep the tree balanced. Describe *three* ways in which this problem can be resolved to produce (one or more) valid class definitions. State, with a brief justification, which you would use here. [10 marks]

4 Compiler Construction

A *regular grammar* is a grammar whose rules are in one of the two following forms (where A and B are non-terminal symbols and a is a terminal):

$$\begin{aligned} A &\rightarrow a \\ A &\rightarrow aB \end{aligned}$$

- (a) Give a regular grammar which generates floating point numbers of exactly the following form:

$$(0|1)^+.(0|1)^*[e(0|1)^+]$$

where “ $()$ ” indicates grouping, “[$]$ ” indicates optional item, “ ρ^+ ” indicates one or more repetitions of ρ and “ ρ^* ” indicates zero or more repetitions of ρ .

[8 marks]

- (b) Give a non-regular grammar with fewer productions than your answer to (a) but which generates the same set of strings. [4 marks]

- (c) Determine, with justification, for the following grammars

(i) whether S generates strings not generated by T ; and

(ii) whether T generates strings not generated by S .

$$\left\{ \begin{array}{l} S \rightarrow aaS \\ S \rightarrow Scc \\ S \rightarrow d \end{array} \right\} \text{ and } \left\{ \begin{array}{l} T \rightarrow aTc \\ T \rightarrow d \end{array} \right\}$$

[4 marks]

- (d) What is the significance for the compilation process of the idea of “strings which can be generated by regular grammars”? Your answer should explain where such a module recognising such strings would appear in a compiler and a possible external interface (functions, variables and/or objects) it might present to the rest of the compiler. [4 marks]

5 Introduction to Security

Which mode (or modes) of operation of the Advanced Encryption Standard (AES) block cipher would you use to protect the following? Give a brief justification for your answers.

- (a) Inter-bank funds transfers. [4 marks]
- (b) Email messages. [4 marks]
- (c) A high-frequency radio modem link. [4 marks]
- (d) Passwords stored on a local disk. [4 marks]
- (e) The pulse train from a gearbox sensor to the tachograph in a truck. [4 marks]

6 Data Structures and Algorithms

- (a) Carefully describe an implementation of quicksort to sort the elements of an integer vector, and state, without proof, its expected and worst case complexity for both time and space in terms of the size of the vector. [7 marks]
- (b) Describe a more efficient algorithm for the case where it is known that the vector has exactly 10^6 elements uniformly distributed over the range 0 to 10^6 . [7 marks]
- (c) Describe an efficient algorithm to find the median of a set of 10^6 integers where it is known that there are fewer than 100 distinct integers in the set. [6 marks]

7 Computer Design

- (a) Why does pipelining introduce data hazards? [4 marks]
- (b) What mechanism may be used to remove some data hazards? [4 marks]
- (c) Why is it not possible to remove all data hazards? [4 marks]
- (d) What hardware is required to prevent data hazards from infringing the programmer's model of instruction execution? [4 marks]
- (e) What is the difference between a data hazard and a control hazard? [4 marks]

8 Operating System Foundations

- (a) Describe how a pool of disc I/O buffers might also be used as a disc block cache. [2 marks]
- (b) What are the advantages and disadvantages of a buffer-cache? [3 marks]
- (c) Give examples of where concurrency control is needed in the management of the buffer-cache. [6 marks]
- (d) Describe how application-level read and write requests are satisfied from the buffer-cache. Include any potential delay due to concurrency control. For simplicity you may assume that a request is satisfied from a single disc block. [Hint: you may find it useful to give some detail of the data held in a disc block header in the buffer-cache.] [9 marks]

9 Operating System Functions

- (a) What are the key issues with scheduling for shared-memory multiprocessors? [3 marks]
- (b) *Processor affinity, task scheduling* and *gang scheduling* are three techniques used within multiprocessor operating systems.
- (i) Briefly describe the operation of each. [6 marks]
- (ii) Which problem does the processor affinity technique seek to overcome? [2 marks]
- (iii) What problem does the processor affinity technique suffer from, and how could this problem be overcome? [2 marks]
- (iv) In which circumstances is a gang scheduling approach most appropriate? [2 marks]
- (c) What additional issues does the virtual memory management system have to address when dealing with shared-memory multiprocessor systems? [5 marks]

10 Continuous Mathematics

- (a) The MacLaurin series for a continuous, infinitely differentiable function, $f(x)$, is:

$$f(x) = f(0) + \frac{f'(0)}{1!}x + \frac{f''(0)}{2!}x^2 + \frac{f'''(0)}{3!}x^3 + \dots + \frac{f^{(k)}(0)}{k!}x^k + \dots$$

Derive the MacLaurin series for each of $\sin(x)$, $\cos(x)$, and e^x . [6 marks]

- (b) Hence, or otherwise, prove that:

$$e^{i\phi} = \cos \phi + i \sin \phi$$

where $i = \sqrt{-1}$ [3 marks]

- (c) Prove that the box function, $b(x)$:

$$b(x) = \begin{cases} 1, & |x| \leq \frac{1}{2} \\ 0, & \text{otherwise} \end{cases}$$

has the Fourier transform, $B(\nu)$:

$$B(\nu) = \frac{\sin \pi\nu}{\pi\nu}$$

where ν is frequency measured in Hertz (cycles per second). [7 marks]

- (d) The convolution of $b(x)$ with itself is $t(x)$:

$$t(x) = b(x) * b(x) = \begin{cases} 1+x, & -1 \leq x \leq 0 \\ 1-x, & 0 \leq x \leq 1 \\ 0, & \text{otherwise} \end{cases}$$

Hence, or otherwise, find the Fourier transform, $T(\nu)$, of $t(x)$. [4 marks]

11 Mathematics for Computation Theory

(a) Define precisely what is meant by the following:

(i) \prec is a well-founded relation on the set S ;

(ii) $y \in S$ is a minimal element for \prec . [3 marks]

(b) If \prec is a well-founded relation on S , show that every non-empty subset of S contains an element that is minimal for \prec . [4 marks]

(c) Let (P, \leq) be a finite partially ordered set. A *chain* $X \subseteq P$ is a totally ordered subset of P , and an *antichain* $Y \subseteq P$ is a subset such that no two distinct elements $y, y' \in Y$ are comparable. The antichains $\{Y_i \mid 1 \leq i \leq k\}$ cover P if

$$P \subseteq \bigcup_{i=1}^k Y_i.$$

Prove that the smallest possible number of antichains in a cover of P is exactly the length of a longest chain in P . [Hint: If not, consider the set of minimal elements in a minimal counterexample.]

[13 marks]

12 Computation Theory

(a) Describe the action of a Turing machine. [4 marks]

(b) Define what is meant by a configuration of an N -state, k -symbol Turing machine. [2 marks]

(c) Explain briefly how to enumerate all possible Turing machine computations, so that a given computation can be characterised by a single natural number code c . [5 marks]

(d) Show that it is not possible to compute the maximum distance travelled by the Turing machine head from its initial position during halting computations as a function of the code c . Any results that you use should be stated clearly. [9 marks]

13 Numerical Analysis I

For IEEE Single Precision $\beta = 2$, $p = 24$, $e_{\max} = +127$, $e_{\min} = -126$. Explain these parameters. How many bits are required to store the *exponent* and the *significand*, respectively? How is the exponent stored? [6 marks]

By means of a table, or otherwise, describe how the following quantities are represented: *zero*, *denormal numbers*, *normalised numbers*, *infinities* and *Not a Number (NaN)*. [5 marks]

Let ω represent any of the operations $+$ $-$ $*$ or $/$. Let x be any normalised or denormal number or ± 0 . Writing n for any *NaN* value, what do the following evaluate to?

(a) $x \omega n$

(b) $\pm \infty \omega n$

(c) $x \omega \pm \infty$

(d) $\sqrt{\pm \infty}$

[6 marks]

Let z be the smallest representable positive normalised number. What are the values of the following?

(e) z

(f) the largest representable number smaller than z

(g) the smallest representable positive number

[3 marks]

END OF PAPER