COMPUTER SCIENCE TRIPOS Part IA

Monday 31 May 2004 1.30 to 4.30

Paper 1

Answer **two** *questions from Section A, and* **one** *question from* **each** *of Sections B, C, D and E.*
Submit the answers in six **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

SECTION A

1 Foundations of Computer Science

(*a*) What does the ML function `map` do? Give an example, first coded without `map` and then with it, to illustrate how it can lead to more compact or comprehensible code. [3 marks]

(*b*) Functions `foldl` and `foldr` might be defined as

```
fun foldl f (e, []) = e
  | foldl f (e, x::xs) = foldl f  (f(e,x), xs);

fun foldr f ([], e) = e
  | foldr f (x::xs , e) = f(x, foldr f (xs,e));
```

Explain what these two functions do and why they may be useful. [4 marks]

(*c*) Here is a typical use of `map`:

```
fun mangle n = (n-2)*(n+7);
fun manglelist x = map mangle x;
```

Show how to express `manglelist` using one of the "fold" functions rather than `map`. [3 marks]

2 Discrete Mathematics

(*a*) Given a set A, define the following terms:

(*i*) a relation on A; [1 mark]

(*ii*) an equivalence relation on A; [1 mark]

(*iii*) a partial order on A; [1 mark]

(*iv*) a well-founded relation on A. [1 mark]

(*b*) If two relations $<_A$ on A and $<_B$ on B are well founded, show that the lexicographic relation and product relation on $A \times B$ are both well founded. [6 marks]

3 Programming in Java

Describe how the following features of Java relate to the behaviour of *accessor* and *mutator* methods:

(*a*) fields;

(*b*) local variables;

(*c*) parameters;

(*d*) return values;

(*e*) visibility modifiers.

Illustrate each with brief code samples. [2 marks each]

4 Operating Systems

(*a*) Modern computers store data in a variety of "memories", each with differing size and access speeds. Briefly describe each of the following:

(*i*) cache memory; [2 marks]

(*ii*) main memory; [2 marks]

(*iii*) registers. [2 marks]

(*b*) Give an example situation in which operating systems effectively consider disk storage to be a fourth type of "memory". [2 marks]

(*c*) A researcher proposes using fast non-volatile memory for *all* data, rather than using the four separate kinds of "memory" mentioned above. Comment on the pros and cons of this approach. [2 marks]

SECTION B

5 Foundations of Computer Science

The following ML datatype can be viewed as defining a lazy or infinite sort of tree where each node in the tree holds an integer:

```
datatype tr =
    N of int * unit->tr * unit->tr;
```

(*a*) Write a function called `ndeep` such that if `n` is an integer and `z` is a tree (i.e. of type `tr`) the call `ndeep n z` will return an ordinary list of all the 2^n integers at depth exactly `n` in the tree. Note that if `n` $= 0$ it will return a list of length 1, being just the top integer in the tree. Comment on its efficiency. [8 marks]

(*b*) You are given a `tr`, and told that it contains arbitrarily large values at least somewhere in it. You want to find a value from it that is bigger than 100 (but if there are many big values it does not matter which one is returned). Because the tree is infinite you cannot use simple depth-first search: you decide to use iterative deepening. Thus you first check all integers at depth 1, then at depth 2, depth 3, ... and return when you first find a value that is greater than 100.

Use exception handling to return the large value when you find it. Present and explain code that searches the lazy tree. [12 marks]

6 Foundations of Computer Science

In ML it is possible to use functions as values: they can be passed as arguments and returned as results. Explain the notation used to write a function without having to give it a name. [2 marks]

This question looks at two different ways of implementing functional arrays.

(a) One possible functional implementation of an array is based on trees, and the path to a stored value follows the binary code for the subscript:

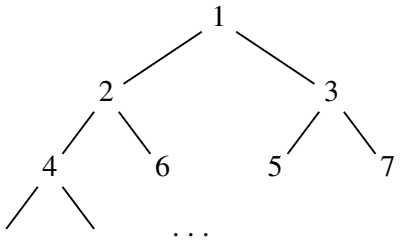

where in the above diagram the numbers show where in the tree a value with the given subscript will live.

Write code that creates, retrieves values from and updates an array that has this representation, and using big-O notation explain the associated costs. [8 marks]

(b) A different way of handling functional arrays is to represent the whole array directly by a function that maps from integers to values. To access the item at position k in such an array you just use the array as a function and give it k as its argument, and to update the array you need to create a new function reflecting the changed value.

If the array is to hold integer values, what ML type does it have? [1 mark]

Write a function `update a n v` where `a` is a functional array in this style, `n` is an integer index and `v` is a new value. The result of the call to `update` must behave as an array that stores all the values that `a` did except that it is as if an assignment of the style "`a[n] := v`" has been performed. [5 marks]

In big-O notation, what is the cost of your `update` function? After a sequence of updates what is the cost of accessing the array? [4 marks]

SECTION C

7 Discrete Mathematics

Recall the Fibonacci numbers defined by:

- $f_0 = 0$
- $f_1 = 1$
- $f_n = f_{n-1} + f_{n-2}$ for $n > 1$

Using induction on n, or otherwise, show that $f_{m+n} = f_{m-1}f_n + f_m f_{n+1}$ for $m > 0$. [4 marks]

Deduce that $\forall m, n > 0 \,.\, m|n \Rightarrow f_m|f_n$. [4 marks]

Deduce further that $\forall n > 4 \,.\, f_n$ prime $\Rightarrow n$ prime. [2 marks]

Given $n \in \mathbb{N}$, let $g_i = f_i \bmod n$, and consider the pairs (g_1, g_2), $(g_2, g_3), \ldots, (g_i, g_{i+1}), \ldots$. Show that there must be a repetition in the first $n^2 + 1$ pairs. Let $r < s$ be the least values with $(g_r, g_{r+1}) = (g_s, g_{s+1})$. Show that $g_{r-1} = g_{s-1}$, and deduce that $r = 1$. Calculate g_1 and g_2, and deduce that $g_{s-1} = 0$. Hence show that one of the first n^2 Fibonacci numbers is divisible by n. [10 marks]

8 Discrete Mathematics

Suppose that A is a finite set with a bijection: $A \to A \times A$. Calculate $|A|$. [2 marks]

Give an example of a countably infinite set B with a bijection: $B \to B \times B$, proving the result carefully. [4 marks]

Consider the following definitions:

$$M = \{n \in \mathbb{N} \mid 2|n\}, \text{the even numbers}$$
$$O = \mathbb{N} \setminus M, \text{the odd numbers}$$
$$P = \mathcal{P}(\mathbb{N}), \text{the set of subsets of } \mathbb{N}$$
$$Q = \mathcal{P}(M)$$
$$R = \mathcal{P}(O)$$

Show that P, Q and R are uncountable, and construct a bijection: $P \to Q \times R$. [12 marks]

Hence show that there is an uncountable set C with a bijection: $C \to C \times C$. [2 marks]

SECTION D

9 Programming in Java

Implement a videogame called JavaConk, a computerised version of the game of conkers. Two round conkers should be displayed on the screen. One conker is controlled by the computer, and constantly travels around the screen, bouncing off the sides if necessary. The other conker is controlled by the player, using the mouse. If the two collide, and if the sum of kinetic energies is less than a threshold set for this game, the computer conker will rebound at a new velocity determined by the collision speed. If the kinetic energy is larger than the threshold, one of the two will be smashed. The conker that will smash is determined by chance, but with probability weighted so that the one travelling more slowly is proportionally more likely to smash.

The classes `BasicGame`, `Ball`, and `OutOfBoundsException` have already been coded. Relevant method signatures for these classes (and standard library class `MouseEvent`) are as follows:

```
class Ball
    public Ball(int diameter, int xPos, int yPos, String colour)
    public int getX()
    public int getY()
    public void moveTo(int x, int y) throws OutOfBoundsException
    public boolean collidesWith(Ball other)
    public void smash()
    protected void draw()

abstract class BasicGame implements MouseMotionListener
    public void mouseMoved(MouseEvent e)

class OutOfBoundsException extends Exception
    public boolean overHorizontalBoundary()
    public int getXboundary()
    public int getYboundary()

class MouseEvent extends Event
    public int getX()
    public int getY()
```

Provide all necessary code to implement the required behaviour.
[Hint: Kinetic Energy $E_K = \frac{1}{2}mv^2$.] [20 marks]

10 Programming in Java

Consider the design of a Java application for the automated marking of exam questions on Java. The answer to each question consists of Java source code. Your application should read the text of each answer from a file, and check that actual lines of code expected by the examiner are included at any point in the answer. The examiners will want to calculate the average mark (out of 20) for each question, and also the total mark (out of 100) achieved by each candidate in a paper consisting of 5 questions. Model answers should take the form of an executable Java class, for example a class called `question1model` might contain the following method:

```
protected void defineAnswer()
  { expectedLines.add("import java.io.*;");
    expectedLines.add("public class helloWorld");
    expectedLines.add("system.out.println(\"hello world\");");
  }
```

Provide a UML diagram showing the classes in the application, and write the basic code for each class, in order to achieve the above functionality. You may ignore error processing, and indicate repetitive passages of code by an ellipsis "...". [20 marks]

SECTION E

11 Operating Systems

(*a*) What was the key innovation of the *von Neumann Architecture*? [2 marks]

(*b*) Describe (with the aid of a diagram where appropriate) the representation in main memory of:

(*i*) an unsigned integer; [2 marks]

(*ii*) a signed integer; [2 marks]

(*iii*) a floating point number; [4 marks]

(*iv*) an instruction. [4 marks]

(*c*) Does an operating system need to know whether the contents of a particular register represent a signed or unsigned integer? Justify your answer. [2 marks]

(*d*) Describe what occurs during a *context switch*. [4 marks]

12 Operating Systems

(*a*) *System calls* are part of most modern operating systems.

(*i*) What is the purpose of a system call? [2 marks]

(*ii*) What mechanism is typically used to implement system calls? [2 marks]

(*b*) Process scheduling can be *preemptive* or *non-preemptive*. Compare and contrast these approaches, commenting on issues of simplicity, fairness, performance and required hardware support. [8 marks]

(*c*) Briefly compare and contrast the notion of *process* in the Windows XP and UNIX operating systems. Describe the scheduling algorithms used in each case. [8 marks]

END OF PAPER