COMPUTER SCIENCE TRIPOS

Friday 3 June 1977. 9 to 12

PAPER 2

Candidates should answer three questions from Section A and one question from Section B.

Candidates are advised to devote equal time to each section.

SECTION A

1. 2^{20} numbers taken from a uniform distribution in the range (0,1) are to be added. Describe in detail how errors propagate when the additions are done in a straightforward manner on the IBM 370 using single precision floating point arithmetic. Design and justify some way of performing this calculation so that the result is perturbed by no more than a few rounding errors in the basic precision of the arithmetic unit.

2. Design a counter using JK flipflops whose counting sequence is 0,1,2,3,4,5,6,7,8,9,10,11,0,1,..., and explain its operation in detail.

3. Discuss possible ways in which an interactive debugging system may provide facilities for tracing the flow of control in erroneous programs. Explain how a simple break-point mechanism could be implemented for the NOVA.

A fragment of a DEBUG dialogue on the NOVA is shown below. Explain each line and comment upon any bugs or peculiarities you think the dialogue reveals. User input is underlined.

```
$;
ST+17/      LDA 3 OPN
X           JSR @GO
GO/         NEXT+3/     STA 3 OP
NEXT+4      MOVR# 3 3 SNC
NEXT+5      MOVS  3 3
GO$B
ST+17$R
```

4. Compare and contrast the facilities of an *integrated* virtual memory system, in which files are the same as segments for all purposes, and a *separated* virtual memory system in which files must be loaded into a swapping regime. Which approach would you recommend to a system designer? Why?

5. Suggest several different ways in which an Assembler language programmer may represent integers on the IBM 370, and discuss the advantages of each. Discuss which of these representations would be useful for the implementation of integer quantities in COBOL.

6. Naive use of recursion can sometimes cause extreme run-time inefficiency. Discuss, with examples, how some recursive functions may be rewritten more efficiently. Estimate the number of procedure calls required to evaluate A(100), where A is defined in BCPL as follows:

```
LET A(N)  =  N<0 → 0, B(N)+A(N-10)
AND B(N)  =  N<0 → 0, C(N)+B(N-5)
AND C(N)  =  N<0 → 0, 1
```

Suggest an alternative more efficient definition of A.

Comment on the possibility of removing recursion from the following definition of the Binomial coefficients:

```
LET  C(N,R)  =  R=0 → 1,
                R=N → 1,
                C(N-1,R-1)+C(N-1,R)
```

7. Discuss to what extent it would be possible to write a program to test whether a FORTRAN program violates the ANSI standard in any way. Show ways in which an ANSI standard FORTRAN program could produce different results on different machines.

SECTION B

The following questions involve the design of programs. This is to be taken to mean providing detailed notes, incorporating all major policy decisions, that could subsequently be used to guide a competent assistant in developing the program required. These notes may contain fragments of code or descriptions of data structures where these form an effective method of communicating design decisions to the assistant.

Any extra assumptions that you need in order to produce a satisfactory answer should be stated explicitly.

8. Machine A and Machine B are two similar computers connected by a data link. The link hardware has three registers called DATA, CTRL and STATE that are accessible to both machines. DATA is a register that holds data words transmitted from one machine to the other; CTRL is a single-bit register which specifies the interpretation of the value in DATA. When CTRL=0 the value in DATA is the count of the number of words to be transmitted to the machine that set it; thus, if machine A sets CTRL=0 and DATA=5, machine B interprets this as a request to transmit 5 data words down the link to machine A. When CTRL=1 the value in DATA is a proper data word. STATE may take the values MCA, MCB and IDLE, and is used to determine which machine currently has access to DATA and CTRL. When STATE=MCA only machine A may read or update DATA and CTRL. Machine A may change STATE from MCA to MCB by the call SEND() and from MCA to IDLE by FREELINK(). If machine A executes the function call RESERVE(), and at the critical moment the value of STATE is IDLE, then STATE is set to MCA and the result of the call is TRUE, otherwise STATE is left unchanged and the result is FALSE. The interpretation of STATE and the definition of SEND, FREELINK and RESERVE are similar for machine B. If both machines call RESERVE simultaneously when STATE=IDLE the link hardware selects which machine wins.

In each machine user access to the link is via two processes, the *reading process* and the *writing process*, which are interfaced to the link hardware by a link driver. The driver is synchronised by the function call WAITEVENT(), which will cause it to wait until a relevant event occurs; the result of the call, R say, specifies the event, which, for machine A, will be one of the following:

R=1 means that the reading process has requested to receive, from machine B, RCOUNT words of data which should be stored in the vector RBUF. When all the words have been read and properly stored, the driver returns RBUF to the reading process by the call RDONE(). The reading process is held until this call is executed.

R=2 means that the writing process has requested to transmit WCOUNT words of data from the vector WBUF to machine B. When the transmission is complete the driver should return WBUF to the writing process by the call WDONE(). The writing process is held until this call is executed.

R=3 means that the value of STATE has been changed from MCB to MCA, indicating that machine B has executed SEND() after setting suitable values in DATA and CTRL.

R=4 means that STATE has changed from MCB to IDLE.

R=5 means that STATE has changed from IDLE to MCB.

For machine B the interpretation of WAITEVENT is similar.

If events occur while the driver is active, they will be queued in order of occurrence until extracted by successive calls of WAITEVENT.

Give a detailed outline of a possible design for the driver, stating clearly any additional assumptions that you make about the problem.

9. It is proposed to provide at a certain computer installation a utility program that compares two nearly identical text files. Suggest how the differences might be presented to the user and outline an algorithm that the program could use to generate this information.

10. The map of a small island is represented by a matrix H of 100×100 elements such that H(x,y) is the height above sea level in metres of the 10 metre square at coordinate position (x,y). H(x,y) is 0 for all coordinate positions in the sea and is greater than zero for all squares on the island. Each 10 metre square can be assumed to be horizontal. There is continuous gentle rain falling over the entire area which causes lakes to form on the island. You may assume that the height of the surface of each lake is equal to the minimum height of all the squares of land which constitute its perimeter. Describe an algorithm that could be used to calculate the total volume of water that is contained in all the lakes on the island.