

THE UNIVERSITY OF AUCKLAND

EXAMINATION FOR BA BSc ETC 1998

COMPUTER SCIENCE

Operating Systems

(Time allowed : **TWO** hours)

NOTE :

Answer any **SIX** questions. The total mark for each question is 20; the total for the paper is 120 marks.

Please write your answers to questions of parts A and B in separate answer books.

PART A

QUESTION 1.

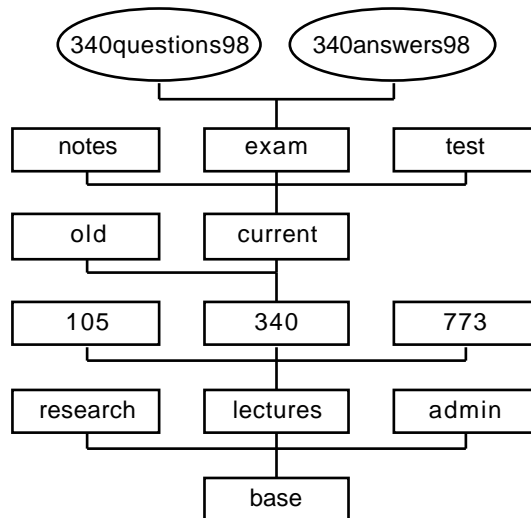
- (a) Describe the components of a disc file which must be managed by the operating system, and show how the file name is used to find the file on the disc.
- (b) How does a Macintosh file differ from a conventional file ? Explain why the additional feature of a Macintosh file might be useful. (NOTE : an answer in principle is required; details of the Macintosh system are not.)
- (c) In the design for a proposed file system, there is provision for files which can contain an arbitrary number of independent data areas; the number of data areas for a file would be defined when the file was declared, but would not be changed thereafter. Suggest how a file system offering files of this sort could be implemented, and describe its file descriptor. (The file descriptor is the entry in the file table or directory which contains the system's information about the file.) You need not go into detail about points which are the same as a conventional file system; point out differences between your suggested implementation and that of a conventional system, including differences from conventional application programmer interfaces.

CONTINUED

QUESTION 2.

Two computer systems both implement a conventional tree-structured file directory. One system is exclusively text-based (such as Unix or MS-DOS), the other is a conventional GUI system in which all file selection can only be done by using a mouse and pointer.

The file structures in both cases are identical, and are as illustrated in the diagram; rectangles denote directories, and ellipses denote files.



Starting from the base directory in the textual system, or a screen with only the window for the base directory visible in the GUI system :

- (a) If you know that the pathname of the file you are seeking is

base/lectures/340/current/exam/340questions98,

what is the easiest way to open it for editing in the two systems ? (You need not spell out every step if you can adequately describe the method more easily.) Describe your input to the system, and how it is interpreted by the system. Identify the information you use, and state what information the system must present to you. Comment in one short sentence on the major difference exhibited between graphical and textual interface styles.

- (b) If all you know is that you are looking for this year's examination questions for the lecture course 340, what is the easiest way to open it for editing in the two systems ? (You need not spell out every step if you can adequately describe the method more easily.) Identify the information you use, and state what information the system must present to you. Comment in one short sentence on the difference from the behaviour in (a).
- (c) Explain why the second operation would be harder if the directories were renamed on the pattern 1, 1.2, 1.2.2, etc. Use your explanation to comment on the function of a file name, and explain why it is useful to permit multiple levels of directories.

QUESTION 3.

- (a) What information must a computer system manage for a process's memory management in a segmented virtual memory system ? Identify the items you would expect to find in the process's segment table, and explain why they are needed.
- (b) What is meant by the "working set" of a process ?
- (c) The algorithm below transposes a matrix A. Regarding each row of the matrix (row i includes the elements { A[i] [j], 0 <= j < N }) as a segment, discuss the behaviour of the algorithm for various values of N if run in a multiprogrammed segmented virtual memory system which can conveniently accommodate a working set of size W, but not much more.

```
int A[][] = new int[ N ][ N ];  
  
int t;  
  
for (int i = 0; i < N; i++)  
{  
    for (int j = i; j < N; j++)  
    {  
        t = A[ i ][ j ];  
        A[ i ][ j ] = A[ j ][ i ];  
        A[ j ][ i ] = t;  
    }  
}
```

QUESTION 4.

- (a) Distinguish between system backup, file generations, and file archives. Explain what each of these does (details of how they work are not required), and point out its advantages from the point of view of someone using the computer system.
- (b) An automatic file archive service moves files from disc to off-line storage without explicit instructions. Describe how the system works. List the changes which you would have to make in an operating system's conventional file management software in order to use such a service. (Do NOT give details of implementation.)
- (c) As costs of storage media decrease, it becomes sensible to think about preserving backups for ever. In an incremental backup system in which a complete backup copy is taken each month, increments are saved each day, and the complete backups saved for many years. Is a file archive necessary ? Explain your answer.

PART B

QUESTION 5.

- (a) Briefly describe the differences between a thread (or lightweight process) and a traditional heavyweight process (such as used by UNIX). Remember to mention differences in context switching.
- (b) In what ways are three threads executing the same code and a UNIX process which forks itself twice producing three processes similar ? How do the threads differ from the processes with respect to open files, memory and communication ?
- (c) Explain the benefits a system level implementation of threads provides over a user level implementation (where the threads are managed by a user level library linked into the process). In particular mention the handling of system calls.

QUESTION 6.

- (a) Distributed operating systems use a variety of different methods to make the system look like one large computer. Briefly describe what the following terms mean :

Location transparency
Migration transparency
Replication transparency

- (b) Briefly describe how a Java program could get a reference to a remote object under Java RMI. What programs must be running and what must have happened before the client Java program can get the remote reference ? Do not describe the compiling of the Java source programs nor the production of the stub and skeleton classes. Do not describe how a remote method invocation is carried out.
- (c) Java RMI does not provide true location transparency because a URL is used to find the remote service. Describe how you could implement a simple system which allows a remote service to run on different locations without having to change the URL used by the client (local) class when it first gets a reference to the object.

CONTINUED

QUESTION 7.

The following program uses an event counter to synchronize activity :

```
program Alternate;  
  
var  
    turn : eventcounter;  
  
procedure OneAtATime( which : char, i : integer );  
begin  
    while true do begin  
        await( turn, i );  
        writeln( which, ': ', i );  
        advance( turn );  
        i := i + 2  
    end  
end;  
  
begin  
    cobegin  
        OneAtATime( 'A', 1 );  
        OneAtATime( 'B', 0 )  
    coend  
end.
```

- (a) Show what the output of this program would be.
- (b) What changes have to be made to the program in order for it to be able to coordinate activity between three concurrent calls to the OneAtATime procedure so that the output looks like this :

```
A : 0  
B : 1  
C : 2  
A : 3  
B : 4  
C : 5  
A : 6  
...
```

- (c) What advantages do event counters have over semaphores ? What advantages do semaphores have over event counters ?

QUESTION 8.

(a) Explain briefly what is meant by these terms :

Device table
Device descriptor
Input-output request block
Device driver
Interrupt handler

(b) With reference to the terms in part (a) describe the steps required to withdraw a device from service and install a replacement (which may be different in some details) without interrupting the running of the rest of the system.
