

Computer Science 340
(for 07.340 and T07.340)

Operating Systems

TEST, 1995

READ THIS FIRST !!!!!!!!!!!!!

Answer all questions –

- BRIEFLY. We don't think that you should need much more than one side of paper for the answer to any one question. Adjust your answer to fit.

Please don't write the answer to one of Alan's questions (1 and 2) on the same sheet as one of Robert's questions (3 and 4). (We'd like to be able to separate them for marking.)

Make sure your name is on every piece of paper which you hand in.

The same number of marks is allocated for each question.

The marks for each question are equally divided between the parts of the question, except in Question 4, where the proportions of the marks allocated to the three parts are indicated by fractions in this typeface.

There are FOUR questions; the test lasts for 90 minutes. As we imagine you can work out for yourself, that can reasonably be interpreted as about $22\frac{1}{2}$ minutes for each question. We expect you to spend about 15 minutes writing each answer; take that into account when deciding the amount of detail expected.

We expect answers in terms of the material already covered in the 340 course, augmented by such general computing knowledge as can reasonably be expected of a stage 3 Computer Science student; dissertations on topics not yet treated in the course are not required and will not receive any marks.

We try to give all the information you need to answer each question, but don't always succeed. If you consider that you have insufficient information to answer a question, don't ask a question in the test : explain your difficulty in your written answer, say what further information you would need to resolve it, and make clear any assumptions that you have made.

QUESTION 1.

- (a) The user database typical of most shared computer systems records facts about the users of the computer system. Briefly explain why such a database is necessary and discuss its security requirements. (Implementation details are *not* required.)
- (b) Name two items of information which might be stored in such a database, and explain why they must be recorded. Also name two items of information which a system might maintain for users currently active in the system (logged on, or with a job currently in progress) but which are not recorded permanently in the user database, and explain how they are used.
- (c) Many programmes designed for use on personal computers (of any sort) provide for the definition of preferred configurations (screen layouts, fonts, memory sizes, etc.), and can save these as "preference files" between sessions. These files are independent of each other, and are commonly stored either with the programmes with which they are associated or in a special preferences directory. Identify and explain similarities and differences between these files and the user database of a shared system.

NOTE : It is, of course, common to have personal preference files in shared systems, and they are generally much like those of personal systems. These are irrelevant to the question, which is about different ways of saving information between sessions.

(CONTINUED ON THE NEXT PAGE.)

QUESTION 2.

- (a) Explain briefly why computer-human interfaces should be *consistent* and *self-explanatory*. In each case, give an example of good design in an operating system interface, and state why you believe the example to be good. (If you can't think of an example, give a very brief description of how a well designed system should behave, and again state why you think it's good.)
- (b) Within any single programme, the copy-and-paste operations used for manipulating selections of text and graphical material in common GUI systems can be implemented in any way which suits the programme's author(s). Suppose you wish to broaden the notion of consistency so that material can be copied-and-pasted from any window to any other, which may or may not belong to the same process.
- (i) Define the task to be carried out by the operating system.
- (ii) Describe the facilities which the operating system should provide to programmers.
- (iii) Give a step-by-step description of a copy-and-paste operation between programmes.

Point out any assumptions about the copy-and-paste operation which you make in your answer.

NOTE that there are several possible answers to this part of the question. If in doubt, just pick one which you think is good and present that.

- (c) Suppose now that the same copy-and-paste principle is extended to include operations between a programme's window and the desktop. (Note that the desktop can reasonably be regarded as the operating system's window.)
- (i) If an icon is copied from the desktop and pasted into a window, should the result be that a picture of the icon appears in the window, or that the contents of the file represented by the icon appear in the window, or something else ? (If something else, say what.)
- (ii) If material is copied from a window to the desktop, what should happen ?
- (iii) Comment on the consistency achievable in these operations.

QUESTION 3.

- (a) When a subject requests access to an object in a capability based-system, it has to provide the appropriate capability in order to access the object. Suppose that to create an object (such as a file) the subject needs to have the capability to create new objects of the required type. Describe how a new process acquires permission to create files in such a system. Explain how your answer maintains the security of file directories.
- (b) In the Amœba operating system, capabilities are implemented cryptographically. Explain how such a system could work. In particular, explain how capabilities are created and how they are checked.
- (c) Since encryption-based capabilities do not require any special hardware assistance, it is possible to send a user-generated random sequence of bits as a capability or as the encrypted part of a capability. In theory this means that a user could attempt to access an object by repeatedly trying different sequences of bits until the correct number is found. Why is this not regarded as a disastrous security problem ? Support your argument with simple calculations based on 48-bit encrypted capabilities. (Note : $2^{10} \approx 10^3$)

(CONTINUED ON THE NEXT PAGE.)

QUESTION 4.

Information about files must be maintained in permanent storage by the file system. In the notes we called this information the file descriptor and it is found via a directory entry. To find this information for a particular file the file system usually has to traverse the directory structure.

(a) Assuming that directories are stored as ordinary files, compare the following two approaches when searching for a file's file descriptor.

(i) Directories are read into memory as required and then discarded (after writing back to disk if changed).

(ii) All directories are read into memory when the system is started and maintained both in memory and on permanent storage.

(1/3)

(b) In designing a new improved user-friendly file system we want to incorporate file names of any length. What advantages and disadvantages does this decision have for the user and the file system (and other sections of the operating system) ? Assume that all directories are read into memory when the system is started. Be sure to mention the consequences unlimited length filenames have on directory structures, and on actions such as finding and renaming files.

(1/2)

(c) Describe another way of getting the advantages of long file names without the disadvantages.

(1/6)

Robert Sheehan and Alan Creak,
June, 1995.