DISCS AS DEVICES.

We have already said quite a lot about disc files, mainly in the *FILES IN THE SYSTEM* chapter; even though we were discussing files in a general sense, the practical fact that most files are disc files inevitably coloured our treatment. Now we want to enlarge on the specific topic of files stored on discs, and also introduce some other material which is commonly stored on discs. We deal with the file systems first, as this is a direct consequence of our earlier discussion, but then we broaden our view to include other uses of the disc which are implied by previous requirements but which we have not stressed so strongly.

The properties of discs are determined by what they are; the interesting properties – from the point of view of an operating system – are determined by how they are used, which is in turn determined by the relationship between the disc and the rest of the computer hardware and operating system.

For a long time it has been sufficient to discuss discs without further qualification. From the introduction of magnetic disc stores around 1960 there has been steady development (and a most remarkable spurt when the primacy of discs as secondary storage media was threatened for a while by bubble memories[EXE8]), with storage density increasing, and variants of the original fixed discs – removable disc cartridges and floppy discs – coming into use. All these devices are sufficiently similar in operation to be grouped together under the same heading.

In the last few years, this monopoly has been threatened by the optical disc, offering greatly increased capacity, but so far mainly as an archive store because it is still not easy to rewrite the contents of an optical disc. Apart from the greater capacity, the optical discs have turned up nothing very new for operating systems, so it is still reasonable to regard discs as a single topic. We shall do so in our discussion, which will be directed mainly towards magnetic discs.

WHAT IS A DISC ?

The obvious answer – a round flat thing that spins around a central axis perpendicular to its plane – isn't a silly answer, because it determines several significant properties of the disc as it is used for storage. For example, it is hard to imagine any sensible recording scheme other than circular tracks of data which can be written and read by some sort of probe mounted near the surface. We'll give a little more detail in the chapter *DISCS – THE HARDWARE VIEW*, but not much more, as the disc machinery is incidental to our requirements except insofar as it determines what the operating system has to do.

The dependence of the disc on physical movement means that it is likely to be comparatively slow to respond; generally it is necessary to find the correct track on the disc, and to wait for the right part of the track to come round as the disc spins. The expected time lags for these two operations are called the *seek latency* and *rotational latency*, respectively. A reasonable expected overall latency for a modern disc system is of the order of a few milliseconds – which is quite fast, but not in comparison with processor speeds.

WHAT DO WE DO WITH THEM ?

We store things in them. Most human activities require storage spaces of one sort or another, if only because we can't do many things at once, and computer activities are similar. Houses usually have several sorts of storage space – cupboards for things used reasonably often, box rooms or the roof space for things which might come in useful some day, letter boxes for incoming post, specialised storage devices like refrigerators, bookshelves (and, in certain cases, the floor, chair seats, table tops, the bath .....), and so on.

Without stretching credulity too far, you can find analogues of most of these sorts of storage space in a computer system, and the disc does it all. Because of that, we must provide different ways of storing material to suit different requirements – usually because of different degrees of urgency, but also to provide privacy and security, and (in partitioning) to allow different sorts of activity – often different operating systems – to be used with the same disc system.

Detachable discs (cartridges and floppy discs) introduced additional problems when they were first used, because it was no longer possible to regard the computer's disc system as fixed. Each disc had to be self-contained, and the system had to guard against the possibility that a disc might be removed while the data on it were not self-consistent.

Once operating systems began to cater for multiprogramming, the disc systems had to be extended to provide for the simultaneous use of the disc facilities by many, essentially uncoordinated, processes. This introduces problems of contention for the use of the disc, interference between different processes' requirements, and load balancing between the discs used by the operating system – all of which must be dealt with as they arise, for, by the nature of a general-purpose operating system, it is very difficult to exercise any control over the details of what the programmes being executed might do from moment to moment.

Finally, we already know something of what we want, as we have designed the file system. Now we begin to see something of the environment in which this system must be implemented, and we shall work through this in the next few chapters.

REFERENCE.

EXE8 : R. Bernhard : "Bubbles take on disks", *IEEE Spectrum* **17#5**, 30-33 (May, 1980) (quoted in *Files and databases : an introduction* (P.D. Smith, G.M Barnes (Addison-Wesley, 1987), page 33).

QUESTIONS.

Can you think of a better way to record material on discs ? (For example, radially – record from the centre outward along radii. Perhaps you'd be able to read a whole radius at once by some means. Or perhaps a spiral track ?)
