


main DOS on the disk, and transfer it byte by byte into RAM, whereupon that DOS can take over and perform some far more sophisticated functions. This process of switching the computer on, then waiting for the DOS to take over, is called 'booting-up'. When it is finished, a greeting is printed on the screen with a prompt to indicate that the computer is ready for a command from the user.

Whichever form the DOS in a system takes, its main function is looking after the locations of the contents of the disk. You may remember that a disk (see page 114) is divided into concentric rings, called tracks, which are in turn divided into sectors; and the intersection of a track and sector is called a block. A block can typically hold 128 bytes of information, and is the smallest unit that the disk can read or write at a time. One of the main reasons for having a DOS is to enable the computer to remember the exact location of everything on the disk. This task is more awesome than it sounds. Let's suppose our disk drive has a capacity of 320 Kbytes — enough to store 20 programs of 16 Kbytes each. With each block holding 128 bytes, loading one of those programs without the benefit of a DOS would require you to specify 128 different blocks, each with its own track and sector number!

displays the BAM on the screen, and you can watch the entries being made as you save a program. When a file is erased, the DOS doesn't bother to wipe clean all the blocks used in that file; it simply changes the entries in the BAM to indicate that the contents of those blocks are now unwanted.

Room To Spare



Before the DOS can store a new file and make an entry in the directory, it must first consult the Block Availability Map (BAM) or Free Sector List. This is a section of memory in which each bit corresponds to a block on the disk. A binary 1 indicates that the block is in use, 0 that it is free (we've shown it as solid or empty squares). Notice that the innermost tracks (at the bottom of the map) have fewer sectors than the others, because they are so much shorter

Local Directory

Filename	Type	Location (Track-Sector)
Invaders	Prog	20-1,20-7,20-2...
Temperat	Prog	25-11,26-5,26-12...
Budget	Prog	23-12,24-3,24-9...
Budgetdat	Data	27-1,27-7,27-2...

The directory on a disk typically occupies the centre track. It contains a list of the filenames, file types (program, data, and perhaps other categories) and track and sector numbers where the file is stored

In order to perform this function, the DOS keeps a disk directory. This is usually located in the middle track of the disk because it has to be referenced frequently, and this minimises the distance the read/write head has to move. The speed of operation of a disk is far more dependent on the time taken to move the head from track to track than on the speed at which the disk spins.

The directory is a list of all the files (which may be programs or files of data) currently on the disk, with details of the file name, file type, and a list of the blocks (each specified by track and sector) where that file is stored. There may be some other entries, such as the date when a back-up copy of the file was last made, or a list of the users who can access a particular file.

When a new file is to be stored, the DOS must first look up something called the Free Sector List or the Block Availability Map (BAM). This has a single bit corresponding to every block on the disk, and as a block is used the value of its bit is changed from zero to one. Some home computers with disk drives feature a utility program that

Another feature of this system is that files are not stored, as would be expected, in consecutive neighbouring blocks. Suppose, for example, that a track consists of 12 sectors, numbered 1 to 12 clockwise. The first 128 bytes of a program might be found in sector 1, the second in sector 7, the third in sector 2 and so on. This is because there is a small time lapse while a block's contents are transferred to the memory buffer used to write each block. If the DOS had to write consecutive sectors, it would have to wait for one complete revolution of the disk between each write — thus slowing the system down. Furthermore, a disk



CHRIS STEVENS

I.Q. Test

Some disk drives contain their own microprocessor and RAM. These are called 'intelligent' drives, and the DOS is incorporated in the form of ROM. Where 'non-intelligent' drives are used, the DOS is stored inside the computer

that has been in use for some time, with files that change in length each day, will end up with a BAM looking like a piece of Gruyère cheese, and new files will have to be fitted into the holes.

A Disk Operating System has many other functions, including formatting new disks (marking out the tracks and sectors on a blank disk and creating an empty directory), making back-up copies, and 'tidying-up' full disks. More sophisticated versions include a variety of data handling structures (see page 204).