

Divide and Conquer

At the Ozaukee Radio Club breakfast last Saturday (8Feb97), Dave Barrow, N9UNR, and I were talking about the huge size of hard drives nowadays. He is in the market for a new computer, and knows that it is wise to think about partitioning the drive into several "logical" drives, rather than having just one huge drive C:, but couldn't remember the numbers to help him make the appropriate decision concerning how large each partition should be. I promised to look up the numbers, and decided you might like to have me share them with you. Actually, I touched on this subject in Number 29, which appeared in the March 1996 issue of BSSS.

Why partition your large drive into more than one "letter"? Overhang. Remember? Suppose you create a file with a simple text editor that contains one sentence (this one). I just did, and it occupies exactly 90 bytes, including the parentheses and period at the end. So, you have used up 90 bytes of space on the floppy, right? Wrong! The smallest "chunk" of information that DOS can write to a floppy disk is 512 bytes (the same is true for all versions of Windows, since all versions of Windows use DOS to handle reading from and writing to hard and floppy drives). That means that when I saved the file, 512 bytes of space were reserved for the file (I named it "SENTENCE.TXT", even though only 90 bytes were used. Those 512 bytes are now unusable for any other purpose whatsoever, unless I erase the file. Exactly 422 bytes of space (512 - 90) are wasted, because they hold none of my data. Those 422 bytes are called **overhang**.

Well, that is not such a big deal. Not on a floppy disk. However, the size of that "smallest chunk" is not the same on all size hard drives (by the way, that "smallest chunk" is also known as a **cluster** or **allocation unit**). Take a look at the following table.

PARTITION	BYTES	OVERHANG (with a 90 byte file)
Drives up to 16 Mb	512	422
16 to 128 Mb	2,048	1,958
128 to 256 Mb	4,096	4,006
256 to 512 Mb	8,192	8,102
512 Mb to 1Gb	16,384	16,294
1 to 2 Gb	32,768	32,678
2 to 4 Gb	65,536	65,446

You can see that there is a lot of wasted space when you write a small file, and that wasted space increases markedly as the hard drive gets larger. For example, the smallest chunk of space that can be allocated on a 512 Mb hard drive is 16 kb (16,384 bytes). Almost all of it is unused (and unusable) space with a 90-byte file; 16,294 bytes are allocated and cannot be used for any other purpose.

You might expect that this waste of space can be gotten around - just don't write any little files. Right? Wrong. If you write an 8,190-byte file to our example 512 Mb hard drive, only 2 bytes are wasted. But, if you write an 8,195-byte file, 3 bytes over the minimum sized chunk, 8,189 bytes are wasted. You completely filled up one allocation unit (8,192 bytes), and the 3 bytes left over went into another one. That second allocation unit also is reserved for exclusive use of your file, even though you only have 3 bytes to stick in it.

So what is the overhang in your machine? Of course, there is no way to know without a thorough analysis, but several writers have suggested that the average overhang equals about half the cluster size. That means if you have a new 1.2 Gb hard drive (a popular size today), you should multiply 32,768 ÷ 2 bytes (16,384) times the number of files on your drive to estimate overhang.

The drive I am writing this article on is 1.6 Gb, partitioned as one big drive C: since it is new and I have not finished organizing it yet. CHKDSK tells me it contains 4,529 files. That means I have about 72 Mb of overhang on my drive at this moment. Although that is less than 5% of the drive's size, it is larger than the total size of most hard drives in use five years ago!

Well then, what is the answer? How should one approach the partitioning problem? There are some rules of thumb, which revolve around the idea of making the drive work for you as efficiently as is possible. Don't partition a drive so that its size is equal to or just over one of the size jumps listed in the table above. Beyond that, it is best to keep your partitions to below 512 Mb. Here is the result of those suggestions:

SIZE OF DRIVE	NUMBER OF PARTITIONS	DRIVE LETTERS
Up to 127 Mb	1	C: only
128 to 255 Mb	2	C: and D:
256 to 511 Mb	1	C: only
512 Mb to 1.023Gb	2	C: and D:
1.023 to 1.535 Gb	3	C: thru E:
1.536 to 2.047 Gb	4	C: thru F:
2.048 to 2.559 Gb	5	C: thru G:
2.560 to 3.071 Gb	6	C: thru H:
3.072 to 3.583 Gb	7	C: thru I:
3.584 to 4.089 Gb	8	C: thru J:

This scheme is somewhat arbitrary and of my design, others might argue with it. On the other hand, it limits the average overhang of smaller drives to 2 kb, and the larger drives to about 4 kb. This approach also allows for dedicating logical drives to specific functions. For example, my installation of Windows 95 occupies somewhat over 100 Mb; I could nicely fit Win 95 and associated programs into one partition, Microsoft Word (about 60 Mb) and associated word processing programs into another, and so on.

Whatever you do, do it for a logical reason. However, realize that we humans are never satisfied; when you get it just the way you want it, your desires will change and it will be time to try some other scheme! Happy computing!