Understanding Disk Geometries

Because of the way that DOS and Windows store information on a hard disk, and the way that the BIOS accesses drives, some PCs may not be able to handle hard disks larger than a certain size without the addition of special software. This can cause problems for support staff.

By Bryan Betts

When IDE was invented, many thought that hard drive geometry issues were a thing of the past. IDE drives intelligently translated one set of parameters to another, so, as long as the settings chosen came to less than or equal to the capacity of the drive, all would be well.

It all changed about two years ago, when the first drives appeared which exceeded half a gigabyte in capacity. Users, support staff and those upgrading the PCs discovered that their machines couldn't see the full hard drive capacity.

This year, drives appeared with capacities in excess of 2 GB, and this time, while PCs could see the full size, the formatting programs within DOS and Windows could not.

The Physical Disk

Drives are defined in terms of cylinders, heads and sectors. A cylinder is a section, centred on the disk's spindle and cutting through all the disk platters along the line of one of the circular tracks.

The number of heads also represents the number of recording surfaces, with platters normally being double-sided. The "sectors" setting refers to the number of sectors (data blocks) written on each track. Each sector holds 512 bytes.

Thus multiplying the number of cylinders by the number of heads gives the total number of tracks, and multiplying this by the sector figure gives the total number of blocks.

Zoning

Reality isn't this simple, of course. Modern drives use zoning: because the tracks near the outer edge of a disk are longer, they can accommodate more data sectors, so drive makers vary the number of sectors per track across the platter to maximise the data capacity. Fortunately, the user doesn't see this as it is hidden behind the intelligent translation carried out by the drive itself.

With SCSI drives, which are smarter than IDE, hence the higher cost, none of this need be a problem. The computer merely asks for a specified data block from the controller, it passes the request on to the drive, and the data is retrieved.

However, the bulk of hard drives in use today are IDE, and IDE reduces costs by making use of the CHS (cylinders, heads, sectors) figures.

The IDE Limit

The IDE specification was put together nearly ten years ago as a way of extending the system bus out to the device by embedding an IBM AT-style disk controller into the hard disk itself. It was defined to support drives up to 137 GB in size, while the PC BIOS itself could support 8.4 GB.

However, the designers of the PC BIOS and of IDE had independently set different maximum limits for the number of cylinders, heads and sectors (CHS) a hard disk could have. Individually, both sets of maxima were generous enough (see Figure 1), but when they were brought together the limiting values became the lowest common denominators, fixing the maximum size of a drive at 504 MB (528 million bytes).

(Notice, by the way, that software companies tend to use 1024 bytes as 1 KB, and 1024 KB as 1 MB, whereas drive manufacturers define a megabyte as a million bytes, so their quoted drive sizes are about 5% bigger than will be reported by the software.)

This same limit was also hard-coded into Windows 3.1 by Microsoft, meaning that the standard WDCTRL 32-bit disk access driver will not support drives over 504 MB in size. If one is connected, 32-bit disk access will be disabled. You can re-enable it by obtaining a replacement driver - check the web site or CompuServe forum of your hard disk manufacturer.

In addition, WDCTRL does not support advanced features such as block mode and DMA data transfers. As a result, some systems will actually be faster with 32-bit disk access disabled. Other operating systems such as Windows 95 and OS/2 should not have this problem.

Limit Workarounds

The IDE limit was relatively painless for most PC users and manufacturers to bypass. Handling larger drives involves a software or BIOS extension to the Int13h limitations, translating the cylinder, head and sector (CHS) parameters into either a CHS array or a Logical Block Address (LBA), both of which support drives up to 8.4 GB in size.

Extended CHS alters the number of cylinders, heads and sectors to fit within the BIOS limits; typically it might halve the number of cylinders and double the number of heads, for example.

LBA translation provides a more consistent mapping of logical to physi-

cal parameters, and also allows a single access scheme to be used for all devices and drivers.

Enabling Translation

PCs with BIOS dates of 1994 or later will usually include LBA or Extended CHS translation within the BIOS, though it may need enabling. If you have a 1 GB hard disk which starts giving errors when half-way through formatting, or half full of data, this is a possible sign that LBA has inadvertently been turned off. Note that different drive access methods tend to be incompatible, so data may be lost if the translation mode is changed.

If neither LBA or Extended CHS is present then either the BIOS or motherboard must be upgraded, or a new adapter card with its own hard drive BIOS can be installed.

Alternatively, overlay software such as Ontrack Disk Manager or MicroHouse International's EZ-Drive can be used; this is often shipped with drives and it rewrites the master boot record, inserting a translating driver to enable the PC to recognise the larger drive. Note: never use FDISK/MBR (which replaces the boot sector on a PC which may have problems, or a boot sector virus) on a machine with EZ-Drive or Disk Manager installed.

A few PCs will not support drives larger than 504 MB under any circumstances, but most can use translation, either in the BIOS or via software, to support drives larger than 504 MB on standard IDE interfaces.

New 32-bit disk access drivers for Windows 3.1 can be installed to support larger drives, too, and are usually supplied along with the overlay software.

Drive speed should be enhanced if a FastDisk driver is used which supports block mode and DMA transfers. Several of these exist but they may be specific to a certain brand of hard disks - for example, Western Digital's WDCDRV.386 does not work with some non-WD drives.

Windows 95

Windows 95 does not suffer the 32-bit disk access limitation, but it is subject to the 504 MB IDE limit as it still works through the BIOS. Only operating systems which do not use the BIOS for hard disk access, such as Unix, can use large drives freely.

EIDE And ATA-2

Large drive support was addressed by Western Digital, which worked with others on the original IDE standard, when it proposed an extension to the original standard called Enhanced IDE (EIDE). This includes an expanded set of drive commands and registers, while maintaining backward compatibility.

Apart from large hard drives, EIDE also covers support for more drives including tape and CD-ROM, and faster data transfers.

The faster transfer modes within EIDE are in turn borrowed from ATA-2, a specification laid down by the Small Form Factor Committee of ANSI, the American National Standards Institute.

Following on from the original ATA (Advanced Technology Attachment) standard, ATA-2 set a range of higher speeds and commands which an IDE drive could support. Some of the transfer modes are DMA (Direct Memory Access) and some are Programmed I/O (PIO), with the new baseline being PIO Mode 3 which can yield up to 11.1 MB/sec, and Multi-word DMA Mode 1 which provides up to 13.3 MB/sec.

For comparison, standard IDE peaks at around 4 MB/sec and standard SCSI at 5 MB/sec.

When a PC boots up, the IDE adapter interrogates the drive in Mode 0 to find out the highest transfer mode it supports, and then switches into that mode. If the drive's highest mode is lower than the adapter's highest mode (for example, if a PIO Mode 3 adapter is used with a Mode 4 drive), the adapter will order the drive to switch into the lower mode.

As a system-level specification, EIDE goes beyond ATA-2 to take into account more than just hard disks. EIDE includes the option of a secondary IDE port, plus ATAPI, the ATA Packet Interface, which allows it to support CD-ROMs and tape drives using a command set based on SCSI-2. ATA-2 also gives recommendations, addresses and interrupts for tertiary and quaternary IDE ports, though most BIOSes would need proprietary drivers for such an adapter.

Indeed, ATA-2 gives manufacturers a lot of choice about which features to include or leave out, and this gave rise to concern that new features were being implemented too haphazardly. In response to this, Seagate, supported by Quantum, derived compatible subsets of ATA-2 called FastATA and FastATA-2 which actually specify the modes that a hard drive must support. FastATA is PIO Mode 3 and DMA Mode 1, while FastATA-2 includes PIO Mode 4 and DMA Mode 2.

Faster Options

Beyond PIO Mode 4, a 20 MB/s PIO Mode 5 was proposed, but the ANSI committee decided not to include it because the electrical noise issues were becoming too problematic. In any case, even PIO Mode 3 outstrips the rate at which most drives can transfer data to or from the disk platter, which typically peaks at around 5 MB/s, so it is only significant for short bursts of data where the cache buffer comes into play.

Instead, DMA support is becoming the preferred route to faster data throughput, and is included in newer chipsets such as Intel's Triton. It is not always simple to activate, though, with some motherboards apparently ignoring the BIOS option for DMA and defaulting to PIO Mode 3.

	BIOS/Int13	IDE	Lowest Common Denominator
Sectors per track	63	255	63
Number of heads	255	16	16
Number of cylinders	1024	65536	1024
Maximum capacity	7.8 GB	127.5 GB	504 MB

Figure 1 - Drive Geometry Limitations.

Disk Geometries

Software driver support may be required to fire up the DMA part of the chipset in some cases, and DMA code is more complex to implement, with the memory controller needing to be aware of scatter-gather - the possibility that it might have to pull various blocks of data from different areas of the disk and then drop them into different areas of memory.

ATA And The Future

Although it has not yet been finalised, the draft of the next specification, ATA-3, is largely complete. There are no faster transfer modes, but many new drive commands and functions are included. Among these are security commands for locking and password-ing drives, aimed at users of notebooks which have removable drives.

ATA-3 also implements Self-Monitoring, Analysis and Reporting Technology (SMART) for early diagnosis of faults, which was originally developed by Compaq for use in its own PCs.

There are more drive identification functions too, allowing drives to report their capabilities in areas such as power saving, where formerly the PC had to send the drive a command and see how it reacted. However, as with ATA-2, most of these extra features are optional for the drive manufacturer and will not be implemented across the board.

After ATA-3 comes ATA-4, which will roll in ATAPI support. A multitasking interface has also been proposed: at present IDE doesn't multitask like SCSI does - you can't issue a command to the drive, disconnect and leave it to report back later. Even with two IDE drives, it is nearly impossible to do one thing on each drive simultaneously.

Partition Limits

The IDE limit applied to the physical drive, so it could not be circumvented by partitioning the drive into sub-504 MB logical drives. However, even when you have large drive support, partitions have their own size limitation, which only really became apparent during 1996 as the first drives larger than 2 GB appeared.

The 2 GB maximum partition size is a fundamental part of the filing systems used by DOS and most DOS-derived operating systems, including Windows 3 and Windows 95. This system uses 16-bit binary numbers for the clusters on the drive; the largest 16-bit number you can have is 65,535, so at 64 sectors per cluster, and 512 KB per sector, that makes a maximum of 2 GB (2147 million bytes) per file structure.

The FAT Barrier

The affected file systems are the FAT (File Allocation Table) used by DOS and Windows 3.x, and the VFAT used by Windows 95. OS/2 and Windows NT have their own filing systems, HPFS and NTFS respectively, for their own disk volumes, but when used in DOS compatibility mode they too are FAT-based and are affected by the 2 GB limit.

Apart from switching to NTFS or HPFS and a new operating system, there has, so far, been nothing you could do about the 2 GB limit. However, in November 1996 Microsoft released a new version of Windows 95 with a 32-bit file system, which means that Win 95 users, at least, now have more flexibility in that area.

Fortunately, the 16-bit FAT limit applies only to disk partitions, not to the physical drives themselves, so by subdividing a larger drive into 2 GB partitions the full capacity can normally be accessed.

The hard drive makers have tried to make a virtue of partitioning, citing its data management virtues. In fact, it may be preferable to divide the drive into smaller partitions, because the partition size also determines the cluster size.

Cluster Considerations

Essentially, FAT and VFAT systems divide each partition into as many clusters as they can, subject to the maximum of 65,535 (see Figure 2). As a result, the cluster size increases with partition size - and because a cluster is the smallest chunk of disk drive that can be accessed, a stored file always takes at least one cluster however small it is.

On a 2 GB partition, the cluster size reaches its maximum of 32 KB. Many files will be smaller than this, so a lot of clusters will only be part full, which means there will be a significant amount of wasted space.

One way around this is to use a disk compression product that has its own secondary file system with smaller clusters. Disk compressors store the user's files in a single FAT file as a pseudo-drive, and some can be set to use a 1:1 compression ratio for minimal performance overhead.

Alternatively, the drive can be split into smaller partitions, bearing in mind the break-points on cluster size. This can increase file storage efficiency by 10 or 20%, or even more for drives holding a very high proportion of small files.

In theory, partitioning can also enhance performance. As already mentioned, most hard disks use a zoned scheme for sector placement, so the outer tracks have more sectors per track than the inner ones. This means that on the outer tracks more data passes under the drive head per disk rotation, so data transfer rates are higher here.

In practice, however, the drive's buffering and read-ahead algorithms tend to smooth out such differences.

One further problem is that of wraparound: on some PCs, drives in excess of 504 MB (if the PC does not support LBA) or 2 GB (if it does support LBA) may be incorrectly recognised, especially if automatic BIOS set-up is used. Typically only the surplus capacity over the relevant limit is seen, thus a 540 million byte drive might appear to have only a dozen or so megabytes.

When this occurs, sometimes FDISK will still recognise and partition the full capacity; otherwise disk overlay soft-

Partition Size (MB)	Cluster Size (KB)
0-127	2
128-255	4
256-511	8
512-1023	16
1024-2047	32

Figure 2 - Cluster Size Variations.

ware must be used to set the drive up correctly.

The next barrier for IDE hard disk capacities is likely to be the 7.8 GB limit imposed by the PC BIOS. Getting around this will require an operating system that does not use the BIOS for disk access, such as UNIX (although it may need patching, depending on the version), Windows NT with NTFS or OS/2 with HPFS.

Hard Drive Upgrades

Adding a new hard drive to an existing PC may be difficult or simple, depending on which of the two main advantages you wish to gain. Disks larger than 504 MB can be installed in many PCs, but the higher data transfer rates will require support from both the system hardware and software.

When installing large hard disks remember that PCs which have a translating BIOS which supports Extended CHS or LBA will accept a larger than normal entry for the number of cylinders, while others will not. Of the latter, most should be able to use overlay software to gain access to the full volume size.

To support fast data transfers, a local bus EIDE adapter will be needed. In modern motherboards such as Intel's widely used Pentium offerings, this is typically integrated onto the motherboard itself. Many VLB systems will already include a translating BIOS, but most do not have integrated EIDE and will require a separate plug-in local bus EIDE card, along with a software driver to enable the faster transfer modes.

EIDE should not be an issue on newer PCs, though new 32-bit disk access drivers for Windows 3.1 may be required. For older systems, several results are possible from installing an EIDE drive, depending on the PC's capabilities:

1. The drive does not work at all,
2. No benefit from EIDE at all, the drive functioning as an IDE drive of 528 MB or less,
3. The full drive capacity is available but there are no speed benefits, or
4. It yields both large drive support and greater transfer speed.

Even without the faster data transfer modes of EIDE, there may still be some performance benefits, as modern EIDE hard drives incorporate many technical advances over previous generations. As a result, even when EIDE drives are used with a standard IDE interface, both access times and transfer rates can be significantly better than those achieved with older IDE drives.

If the drive was delivered with disk utility software, this should handle the task of setting it up as a boot drive under DOS. A bootable DOS diskette will be needed in any case, and to help solve problems later this should also have the appropriate versions of the SYS, FDISK and FORMAT programs for your operating system copied onto it.

Most drives are now delivered preformatted; if not, then obviously it must be partitioned and formatted first. A PC with a translating BIOS will not need overlay software, such as Disk Manager or EZ-Drive, and the new disk must merely be FDISKed to make the primary partition active, and SYSed to copy over the system files.

Whether using overlay software or not, if the new drive is to be the primary boot drive, it is generally easiest to set up the system with this drive alone and boot it from floppy. Once the new drive is set up and booting correctly, the original drive can be jumpered as a slave (remember to also adjust the new drive to be master in a two-drive system, if necessary) and reinstalled for its contents to be copied over. This might need a two-way IDE cable and a two-way power splitter.

It may be necessary to obtain a listing of jumper settings, like that supplied with EZ-Drive or available on some bulletin boards, to rejumper the old drive, particularly if the original documentation is not available.

Remember also to make a note of the cylinder/head/sector parameter settings for the old drive before you remove it from the system. A truly excellent source for this information, if it is not listed on the drive, is The Hard Disk Database, available from Electrocution in the UK for around US$20.

Needless to say, data should be backed up before upgrading. It would also be a good time to reinstall DOS and Windows, and possibly the application software as well. This is particularly true if using overlay software, as there is a fair chance that the Windows drivers installed by the overlay software will conflict with something in an existing Windows setup, and a fresh Windows installation is the best way to avoid this problem. Don't forget to install a suitable FastDisk 32-bit disk access driver for Windows 3.x, too.

If possible, after copying the data onto the new drive, remove the old one (rejumpering the master for single-drive operation if necessary) and discard it - an old drive will be slower and less reliable. If you must keep it, use it as a slave, preferably on a different IDE channel as some BIOSes don't permit different drives on the same channel to use different IO modes, forcing a single mode for the whole cable instead.

Lastly, don't attach an ATAPI CD-ROM to the primary channel with a hard drive under Windows 3.x. It will prevent 32-bit disk access from working unless you have a new CD-ROM driver that supports it.

For information on The Hard Disk Database from Electrocution, email paco@cix.co.uk or telephone +44 1264 870720.

PCSA

The Author

Bryan Betts (bryan@cix.co.uk) is a freelance IT writer.

New Reviews from Tech Support Alert

Anti-Trojan Software Reviews

A detailed review of six of the best anti trojan software programs. Two products were impressive with a clear gap between these and other contenders in their ability to detect and remove dangerous modern trojans.

Inkjet Printer Cartridge Suppliers

Everyone gets inundated by hundreds of ads for inkjet printer cartridges, all claiming to be the cheapest or best. But which vendor do you believe? Our editors decided to put them to the test by anonymously buying printer cartridges and testing them in our office inkjet printers. Many suppliers disappointed but we came up with several web sites that offer good quality cheap inkjet cartridges with impressive customer service.

Windows Backup Software

In this review we looked at 18 different backup software products for home or SOHO use. In the end we could only recommend six though only two were good enough to get our “Editor’s Choice” award

The 46 Best Freeware Programs

There are many free utilities that perform as well or better than expensive commercial products. Our Editor Ian Richards picks out his selection of the very best freeware programs and he comes up with some real gems.