Understanding Drive Technology

Disk drives have come a long way since the first PCs shipped. We take an in-depth look at the development of these drives, how they work, and what you need to look out for when specifying them for your users.

The first of a two-part article.

By Ian Murphy
IT Journalist

As the computer revolution continues, we find ourselves committing more and more information to the care of electronic systems. A vast amount of that information is purely transitory, and the benefits of ownership are often questionable but, despite this, it is necessary that we store it. At present we are doubling the amount of information stored every couple of years according to the major analyst organisations. The problem of storage, retrieval and the fear of loss of data have led us into a position where much of what is stored is duplicated. Rather than creating the paperless office, the increasingly widespread use of computers has spawned the paper-hungry office, as people insist on making copies of everything in case the computer crashes.

This fear of data loss is based on the problems of early systems and the exaggerations of those who had to deal with them. Despite the reliability of modern computer systems, failures still occur, but mechanical failure represents an inordinately small proportion of lost data. A much more important area of data loss is that caused by malicious acts such as those perpetrated by virus writers, industrial sabotage, random acts of electronic vandalism, disgruntled employees and badly- written programs. Even this category pales into insignificance when compared with the amount of data that is lost through stupidity, carelessness, and either poor training or lack of it.

The introduction of high-level languages to make application development easier has led to programs that have become increasingly disk-hungry. There are several reasons for this - poor programming skills, bad programming tools, compilers that only do minimal code optimisation, and use of graphical user interfaces with their massive overhead. The average user requires almost 1 GB of disk space for their operating system and office productivity applications before allowance is even made for their data.

On top of all this comes data gathered from the Internet, email, mailing list servers and newsgroups. As organisations have allowed their users access to the Internet for research purposes, they have been rocked by the amount of information that is now being gathered. Many people still do not realise that, in order to view a Web page, the entire contents of the page are downloaded into a temporary directory before being displayed on your screen. In order to speed up access to Web pages there are numerous applications that allow you to duplicate entire Web sites onto local storage so that you can read information without incurring huge telephone bills. Even where users have access to leased lines, these programs allow overnight information gathering to take the stress away from day-to-day Internet access.

Today's Needs

To allow us to capture and store all this information, disk drives have become increasingly larger and faster, and virtually every computer sold today is equipped with one. The current range of low-end drives already exceed the capabilities of the high-end drives from five years ago, yet storing so much information on a single device requires an understanding of the capabilities of that device. Speed isn't necessarily a good thing; nor is it the answer to every problem. Instead it's important to evaluate the way your users need access to data and decide on the right disk technology. In order to make that assessment you need an understanding not only of older drive technologies, but also what is happening with current products and their future development.

Disk drives can often be in use for a considerable period of time, and sometimes additional storage is added rather than replacing the drive. The problem with replacement is that it necessitates complete reinstallation of operating systems and applications, which means you also need to back up and restore the data.

1983 saw the introduction of the hard disk as a standard component in the fledgling PC, and comparing those early 5 MB drives to the 60 GB and greater drives of today makes little sense. Although the basic idea has changed little, the way it is implemented had undergone a radical change. During that entire process the terminology has also changed, and it can almost be a full-time task keeping up to date with the latest terms and ensuring that you do understand them. SCSI (Small Computer System Interface) is a perfect example of this, and we will see later how changes here have occurred even before products have been properly agreed through the standards bodies, leading to confusion over what should be implemented in the next generation of computers and servers.

Inside The Drive

A hard disk drive is made from a number of flat platters made from aluminium disk substrate or mylar, each side of which is coated with a magnetic material composed of small metallic particles made from either iron oxide or thin-film metal media. This material is divided into a number of concentric circles called tracks, and then further sub-divided into slices, creating areas known as sectors. If you were to pass a nail through the same sector on all the platters in a disk, you would end up with a cylinder. The data is read and written by means of heads mounted on an actuator arm. One of the key speed improvements over the last fifteen years has been in the performance of the actuator arm and its accuracy in positioning over the data area on the disk.

Another element in increasing the performance, and which has also improved the storage capacity, has been the increase in the number of heads on the actuator arm. Early drives used just one head per side, and this led to a number of performance issues. As data became more fragmented over the platter, the head needed to be moved a considerable distance in order to carry out any read or write operation. The early drives were also incapable of reading the data when it was very close to the head without having to rotate the surface to gain an accurate positioning, and this led to wasted effort and a reduction in performance. To reduce this, the data was not written contiguously to the surface, but was spaced out across the sectors. This was often referred to as the “interleave factor”, with a 3:1 factor indicating the need to rotate the disk three times in order to read the required data. Today’s drives all work on a 1:1 ratio and use data buffers, therefore interleave factor is rarely mentioned any more.

Heads

With multiple heads on a single arm, the distance that needed to be traversed was much reduced with the inevitable performance gain. The heads have also been moved closer to the actual platter, and this has had two important benefits. The first is a reduction in the amount of power necessary to read and write, which means disk drives can be smaller, as the amount of heat that needs dissipating is reduced. With the reduction in power requirements we also saw more reliable notebook computers, since as storage power needs reduced, battery life was extended. The second benefit was that the amount of data stored on a platter could be significantly increased, as the heads were smaller and more accurate, and the material coating the platter was much finer than that previously used.

“Every hard disk ships with a potential storage capacity, but formatting programs determine how much data will be stored in any given sector.”

This led to another advance in the way that the data was actually stored on the disk platter. Each track on the platter is further divided into segments, rather like a cake, called sectors. Older drives were very wasteful in the way they created the sectors on the disk because each track was split into an equal number of sectors. This meant that the storage capacity of the disk was constrained by the amount of data that could be stored in a sector at the centre of each platter. This was replaced by a technique called MZR (Multiple Zone Recording), which takes into account the fact that, as you move out from the centre of a platter, tracks are bigger. By using a fixed amount of space to indicate the size of a sector, the outer tracks can be split into more sectors than the inner tracks, allowing a more effective use of the drive. There are now utilities which allow you to specify what data is stored in the sectors at the very centre of a drive, as they are likely to be accessed faster than those at the outside.

File Contiguity

The contiguity of files has always been an issue in how data is stored and retrieved. It makes sense to consider that you store the parts of any file in such a way as to keep them together, but the way we create data means that files are often broken up over time. Going back to the early days of drives, when we used an interleave factor, the last thing we actually wanted was large contiguous files, because we couldn't read the entire file in a single pass. This increased the need for wasted rotations before data could be read or written. As a result, a number of techniques were experimented with in terms of data slewing, where data is written in the most effective way for the drive to recover it. In large-scale RAID (Redundant Array of Inexpensive Disks) implementations, you can actually specify how the slewing is achieved, and the early development work in this field produced many papers showing how performance on the Zebra file system was heavily influenced by the rate and angle of slew. Today, we rely on the microelectronics of the drives allied with drive array and operating system software to ensure that, when we defragment our data, it is written away in an optimal fashion.

Rotation Speed

The speed at which the hard disk rotates is also a critical factor, with faster rotation speeds giving the highest potential for data read / write. Whilst this is theoretically true, much depends on the quality of the drive microelectronics, the quality of the heads and the ability of the drive to perform consistently at higher speeds. Most low-end drives today run at around 5,400 rpm, while performance drives are generally 7,200 rpm, and newer drives reaching 10,000 rpm. Over the next year, expect the number of 5,400 rpm drives to diminish rapidly as 7,200 rpm becomes the norm, and 10,000 rpm takes over the performance space. Some vendors have been talking about drives reaching speeds of 15,000 rpm and Seagate has already announced its first drives spinning at this speed.

At such speeds there are two primary issues that need to be dealt with - noise and heat. If you are designing drives to sit in a server that will be located in a purpose-built computer room that is air-conditioned and in which no-one spends a great deal of time, then you can get away with a reasonable variance in both these areas. Unfortunately, with the amount of high-performance computing that takes place at the desktop, and the increasing number of departmental servers, both issues must be dealt with. Eliminating noise is extremely difficult for the drive designer because they are under pressure to deliver increasingly small drives with much higher capacity, where all the available space is used for the platters and supporting microelectronics. The designer also has no way of knowing what type of drive cabinet will be used to hold the drive, and a significant amount of noise can be due to sympathetic vibration within the computer.

It is always worth checking that the hard disks are properly secured and, if the drive mounting section is removable, check that it is firmly secured to the chassis. Check how snugly the sides fit on the computer and, if there are areas where the side panels are not firmly affixed to the chassis, consider inserting small clips to form a solid mounting. This can be a real problem where you have desktop computers with the monitor sitting on the system base, because it puts a lot of strain on one part of the case and can cause distortion. Whilst you are unlikely to get vibration due to the weight of the monitor acting as a damper, any use of the system base at a later date without a monitor mounted upon it will almost certainly mean a problem with vibration noise.

"Some drive enclosure and case manufacturers now produce special cases with a fan for each drive, and these are highly recommended."

Heat Dissipation

Heat is undoubtedly the biggest problem, and both the speed and the reduced size of the drive act to exacerbate it. Fortunately, many case manufacturers still provide older 5.25 inch drive bays so, when you install drives, there is a reasonable amount of space between the drives. If you are expecting to use more than three drives in any given case, or if you expect to use other heat-generating devices such as a CD-R (CD Recorder) or CD-RW (CD Read / Write) device, you need to remember that this class of device in particular is very heat-sensitive. Get the CD-R / CD-RW drive too hot, and you will get a significant failure rate on your production of CDs.

When placing multiple drives into a computer containing an AGP card or a slot on a processor such as the Pentium II or Pentium III, you are going to have a lot of very high heat-generating devices. The standard fan on the power supply is unlikely to cool the system effectively, and although the processors should be equipped with their own fans, they only move the heat away from the physical processor itself - they are not external fans. In this case, you need to attach at least one fan to the front of the system case and another to the rear, to ensure a flow of air across the motherboard and all other heat-generating components. Some drive enclosure and case manufacturers now produce special cases with a fan for each drive, and these are highly recommended, especially in computers where high performance is necessary.

Another reason to ensure that you have a solid connection between the drive and the chassis is electrical grounding. All devices create some form of electrical interference, and in modern computers with their vast array of electronics this problem has become worse over the last few years. Many desktop users report unsatisfactory performance of microphones when using sound cards in their computers. Undoubtedly this is due to interference from other components, and the key culprits are the cooling fans on Slot 1 processors, hard disk drives and CD-ROM / DVD devices. There is no easy solution except to carefully ensure that the mains lead is properly earthed, all devices are fixed firmly to the chassis and that the chassis itself is earthed.

Terminology

When purchasing a hard disk it is important to understand the terminology, because purchasing decisions can often be based on wrong interpretation of the specification. The most important feature of a drive's performance is the internal transfer time. This averages out a number of different values and shows how long it really takes to extract data from the surface and return it to the cache buffer. This measurement is a better indicator than the more traditional access time, which is a combination of rotational latency and seek time, and which is all too commonly quoted as the ideal performance criteria. It is generally believed that the larger the cache buffer the better, although this has to be tempered by the type of drive interface, and other cache mechanisms such as those on the hard disk controller and system cache.

Each cache in the chain will build its own picture of how frequently certain areas of data are requested, and this is where cache gets its speed from. Unfortunately, because this is a chain, each cache will hold a slightly different view on the most requested piece of data. To minimise this inconsistent view, cache algorithms exchange their cache tables with each other, but if your caches are too large then the end result can be a constant process of cache rebuild that will ultimately negate the advantages and slow the system down. In addition, EIDE is not as efficient as SCSI in how it uses the cache buffer (this will be dealt with later).

Transfer Rate

"A number of techniques were experimented with in terms of data slewing, where data is written in the most effective way for the drive to recover it."

The next measurement is the external transfer or data rate of the drive interface, and this is set by the drive interface type - Ultra SCSI, SCSI 160 / m etc - and therefore is independent of the actual drive. It represents the maximum possible transfer speed between the drive and the controller card, although there are issues over command overhead and system bus performance that limit the actual speed of transfer, often meaning that drives perform below their capability. On the other hand, manufacturers of caching controller cards such as DPT (now owned by Adaptec) and Mylex (recently acquired by IBM), claim that they can consistently match the interface speed through their memory and algorithms.

Always look at the manufacturer's Web site or the drive specification to see if they provide you with an average sustained transfer rate, as this is what you are likely to experience. The only other key measurement is the rotational speed that has already been covered, and this is only truly important for a very small number of applications such as Audio Visual, CAD or maybe video streaming. In the case of video streaming it is unlikely you will be using a single drive, so the performance of the RAID controller and of the type of RAID solution implemented will also be a significant factor.

Defining Capacity

Finally, there's that old problem of understanding how much data can be stored on a disk when you format it. Every disk comes with an unformatted capacity that is a combination of the number of bits per track, the number of cylinders and the number of heads. However, when you format a disk, you impose a sector size, giving a new calculation of sector size, sectors per track and number of heads. Every hard disk ships with a potential storage capacity, but formatting programs determine how much data will be stored in any given sector. If the sector is too small fragmentation becomes a major problem, whilst a large sector size might lead to a massive waste of resources. Some operating systems do not allow you to determine the sector size; instead the vendor has already implemented an algorithm that will do this based on the size of the volume you are creating.

Try doing an examination of the data on your hard disk and see how much disk space is used to store your data. If you have a lot of Web site-sourced graphics and similar, you will almost certainly discover that you are using as much as 100% more disk space than the actual data you are storing. This is quite common, and the natural response is to resort to using a compression program to store more information in the available space. This will simply not work. No matter how large a sector, it can only contain data from a single file, so compressing the data is only effective if the majority of your files are greater than one sector in size, and if your compression algorithm reduces their storage requirements by at least 50%. However, obviously whenever you compress data you must decompress it in order to make use of it, so there will be a lot of work done by your processor, and it will have to allocate a portion of memory to hold the two versions of the data.

Drive Families

Hard disks break down into two main families, IDE (Integrated Drive Electronics) and SCSI (Small Computer Systems Interface). Each family has evolved through several different interfaces and we are going to take a look at each of them. We are also seeing the introduction of new operating systems from Microsoft and a change in underlying system architecture that will determine whether or not you can take full advantage of these newer drives.

Before the IDE interface became standardised we had a range of different drive technologies such as MFM, RLL, RLL-2 and ESDI. Each of these advanced the capabilities of the hard disk drive, but they were limited in several areas such as speed and capacity. The latter was a big problem because, even today, when you look at the system BIOS, you can see that the BIOS vendors still hold information on drives that are unlikely to be used because they are so small. This backward compatibility issue is a problem for the whole computer industry, yet there comes a time when such adherence to older standards becomes irrelevant, and for most computers manufactured today we have passed that point with regard to drive capacity.

IDE, EIDE (Enhanced IDE) and Ultra-DMA (Direct Memory Access) are all derivatives of the same architecture, in much the same way as SCSI, SCSI-2 and Ultra-SCSI are. IDE was brought about by Compaq requesting that Western Digital look at the idea of combining the hard disk drive and its electronics with the drive controller back in 1984. This was taken through to manufacture by CDC (Imprimis) in 1985, and subsequently shipped in early Compaq AT computers. The move to EIDE was not such a clear development, with two different groups competing to produce a better interface using the IDE command set. The ANSI committee finally acted to create a single standard combining elements from both the ATA (Advanced Technology Attachment) and EIDE proposals.

"Older drives were very wasteful in the way they created the sectors on the disk because each track was split into an equal number of sectors."

It should not be forgotten that the IDE interface was quite a big step forward, and the original IDE or ATA interface was designed around the original system bus still found in machines today - ISA (Industry Standard Architecture). This system bus was inherently slow and thus impacted any peripheral designed to utilise it. Early IDE drives were designed to transfer data at 4 MB / sec, but soon developed to handle 8 MB / sec. Although it has become common to refer to the newer IDE drives as EIDE drives, it is more accurate to talk about this class of drives as ATA drives, as it is the ATA numbering that is regularly updated to reflect the newer interfaces and performance characteristics. The key elements among the joint proposals were lifting the size of drive that could be recognised without special software, providing support for other devices such as tape streamers and CD-ROM drives, support for more than two devices in a single computer, and much faster data transfer.

When you read through the specifications of different hard disk drives you will see that there are two main ways of transferring information, PIO (Programmed Input/Output) and DMA. However, be wary of those who describe hard disk drives in terms of the PIO and DMA modes that they use. During the development of the IDE/ATA interface huge problems with timing and interference on both the interface and the personal computer have been overcome. The problem with relying on PIO and DMA mode number is that the changes to the interface have not been reflected through the PIO or DMA mode numbering. For example PIO Mode 4 and DMA Mode 2 are used in all ATA specifications from ATA-2 through to the newest, ATA-5. As you can see, it is irrelevant then to talk about a PIO Mode 4 or EIDE device as often occurs when talking to suppliers because this could be a hard disk whose interface is rated at between 16 MB/sec and 66 MB/sec.

PIO is the most basic mechanism and on older computers with limited clock cycles was extremely expensive on processor time. However, this reliance on the CPU was also the reason why many manufacturers favoured PIO because it meant that they didn't need any additional hardware and the different PIO modes define data transfer speeds that depend on the PIO cycle time. Arguably, this cycle time is the major limitation in performance of PIO devices. In order to take advantage of the higher PIO modes, a significant improvement in performance of the system bus and the time it takes for the command turnaround time are required. Intel's new IA-64 architecture and the vast number of available cycles on processors such as the Pentium II have enabled manufacturers at the low end to keep cranking up the PIO performance.

Advantages

DMA, however, has some significant advantages over PIO because today, we rely on multiword rather than single word DMA. The differences between single word and multiword DMA are only seen when multiple requests are made to retrieve, for example, a large file. With single word, a complete command has to be issued for each piece of the file to be transferred, whilst multiword takes the initial request and holds the signal lines high until the last piece of the file has been retrieved. As well as multiword and single word DMA, there are also first and third party DMA mechanisms. First party or bus mastering DMA requires a DMA chip on the drive allowing the drive to take control of the transfer of data to and from memory without utilising CPU cycles. This is highly efficient in resource utilisation and it is for this reason, bus mastering is found in many different devices such as network cards, as well as hard disk controllers. In addition, this is the preferred mechanism for the SCSI interface as well as the higher performing Ultra-DMA drives. However, when implementing multiple bus mastering devices inside a computer, it is important to ensure that you carefully check the architecture you are working with.

The PCI architecture requires an arbitration chip that deals with requests from bus mastering devices when they want to take control of the system bus and memory for transfer of information. As you might imagine, given the number of different motherboards using PCI chipsets, there are some good and bad implementations of this arbitration mechanism.

This article will be concluded in a future issue of PCSA.

"A hard disk drive is made from a number of flat platters made from aluminium disk substrate or mylar, each side of which is coated with a magnetic material."

PCSA

Understanding Drive Technology

We conclude our examination of hard drive technology, from MFM to Ultra SCSI, Firewire and beyond.

By Ian Murphy
IT Consultant

On many early PCI machines, there were real problems with more than two bus master devices and some slots didn't actually support bus master devices, therefore you may want to check with your hardware vendor before loading your PCI bus.

Third party DMA puts much more stress on the main CPU in that it utilises the primary DMA chip on the motherboard to control all transfers. As it is using a chip on the motherboard, it means that there is much more activity when reading and writing data, therefore transfers are typically half the speed of first party/bus mastering DMA devices. Cacheing and data buffers are both used to help reduce the physical latency involved in read and write operations. With these methods has come a range of terms for different techniques such as pre-fetch, adaptive and segmentation. In addition, high performance disk controllers use memory for cacheing and often utilise a technique known as elevator sorting or scatter/gather and this technique is now being incorporated in hard disk electronics.

Buffers

Data buffers were originally used to enable a complete sector to be read into memory before being either transferred to memory or written to disk. Whilst this provides an improvement in performance, the real gains come when the buffers are large enough to contain an entire track. At this point, a technique called pre-fetch is employed and works on the basis that when a request for a logical block X is made, it is fair to assume that the next request will be for logical block X+1. In the early days, operating system vendors built specialised commands into their operating systems to improve the performance of the cache. With the arrival of hard disk cacheing controllers, the market got a little confused with conflicts between drive vendors, operating system vendors and hard disk controller vendors, all of whom claimed that their implementation was superior.

One example of this was that using some versions of Novell NetWare you were forced to disable operating system cache in order to prevent data corruption and unexpected system crashes. It also caused incredible problems with performance monitoring tools, as the measurement programs only took account of how long it took to get data off of the physical hard disk. However, when it comes to complex hardware RAID solutions the performance improvements gained from the use of a cacheing controller are significant. In fact, the performance that can be gained from the controller card is so significant that Adaptec has been including high-end SCSI cards that can use standard memory chips for several years now.

Elevator Sorts

As you increase the size of a buffer, you need a better way to extract the data and the most important technique is called an elevator sort. Imagine getting in an elevator in a large building with 10 other people, who all want to get off at different floors. If the elevator simply went from floor to floor in the order that the buttons were originally pressed, it would mean many trips up and down the building before everyone arrived at their floor. Older drives worked in exactly this way, in that they picked up data blocks in the order that they were written to disk. This often meant many, many wasted cycles whilst the disk rotated and the heads waited for the correct sector to appear beneath them.

Elevators are much more intelligent than that, and will make a single pass up and down the building depositing and picking up new people. When we apply this technique to hard disk drives with sufficient memory, then the result is the same. A single pass is made across the disk, gathering data as it goes, and then it is reassembled in the correct sequence before passing it to the system bus. In fact, so efficient is the mechanism that some drives utilise two buffers, one for read and one for write, which has the effect of dramatically improving read / write performance.

There are two ways that vendors implement their cache, and these are known as adaptive and segmenting. Segmenting refers to the splitting of the drive cache into multiple segments, which has a benefit for multi-user systems. When program A requests a data read, it is only allowed to use one of the available segments. This means that the next program has an available segment when it wants to cache data requested by its user. Whilst the segments are small, the cache is used much more efficiently and the larger capacity drives currently ship with four segments.

Adaptive cache adds a degree of sophistication to this technique. If there are four segments in use, there can be a problem when program 5 requests that data be read from the disk. Using a simple FIFO principle, the data being cached by program 1 would be flushed from memory and replaced with that for program 5. However, if the data for program 1 were being heavily used, then this would be extremely inefficient, not to say wasteful. This is a problem experienced with operating systems that use disk space to augment their physical memory. The sophistication is in the algorithm that decides which segment to purge for the data requested by program 5 based on usage, frequency of access and size of data.

Ultra DMA

ATA-5 or Ultra DMA / 66 was designed, unsurprisingly, to double the effective data transfer rate of Ultra DMA / 33 drives. These drives have been getting a lot of good press lately but there are problems in using them in many computers. In order to increase transfer speed with ATA-2, synchronous data transfer mechanisms were introduced. In this approach the drive controls all the signals, known as a strobe, and uses the leading or rising edge to signify a signal separator. A signal may be either command or data and speed increases occur when you speed up the strobe. Earlier the problems with electromagnetic interference were mentioned and early computers had real problems with this. Today, the health and safety rules have tightened up considerably on how much radiation occurs in a computer and this has meant that speeds can be increased. However, placing high speed drives in older machines may not yield the performance gains you might expect due to the amount of retransmission needed to overcome interference.

“Placing high speed drives in older machines may not yield the performance gains you might expect due to the amount of retransmission needed to overcome interference.”

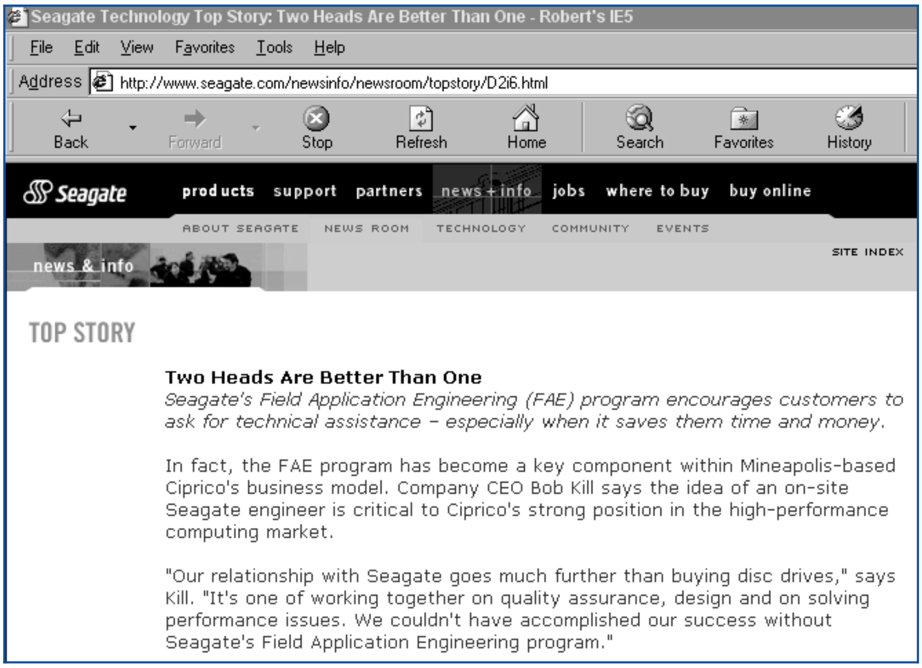

When Ultra DMA drives first appeared they used both the rising and falling edges of the strobe as signal separators and this allowed an effective doubling in speed without an increased risk of interference. This also meant that they could continue to use the existing 40-pin cable that has been in use for some time. With the newer Ultra DMA / 66 drives, the strobe speed has had to be doubled in order to increase the speed. This has the result of increased susceptibility to interference so the drive vendors have started to ship a new cable that must be used to minimise these problems. There are still only 40 data pins, but a ground pin is interleaved with each signal pin, providing what sound technicians would call a balanced interface. This results in an 80-pin cable. This approach is common to several areas of electronic signalling particularly in the data communications field.

Motherboard Upgrades

As well as requiring the newer cable, computers also need an updated motherboard to take advantage of the Ultra DMA / 66 drives. Backward compatibility means that these drives will work in older machines but only at lower speeds. You may also find that additional software is required, as in the case of Seagate drives, to toggle Ultra DMA / 66 functionality. On the positive side, there is little extra work required to manufacture these drives other than a small software change, an update to the drive electronics and the 80-pin cable. This should mean that prices will continue to fall on these drives, unlike the premium prices often demanded when newer SCSI drives appear.

SCSI

This takes us forward to SCSI devices. The original SCSI specification was well into production before the ANSI committee, X3T9.2, finally ratified it in 1986 as ANSI X3.131-1986. As with any standard that is implemented before ratification, there were a number of known problems with SCSI-1, not least because there were numerous areas that were left to vendors for customisation. That said, many different vendors were able to make their products interoperate and this provided the impetus for the SCSI-2 standard.

SCSI-2, like SCSI-1, was being implemented long before it was ever ratified and this, in the end, turned out to be what saved SCSI. Whilst the formal definition on SCSI-2 started in early 1989, the changes that ANSI was going through caused several delays before ratification of the standard could take place in 1994. Luckily the ability of the vendors to work toward a common goal was so strong that over the last three years of its development, the only changes were very minor and so X3.131-1994 was finally released.

“Low voltage differential devices are designed to combine single-ended and differential into a single electrical specification providing low-cost, low power, high signal stability and long cable lengths.”

Seagate products support partners news + info jobs where to buy buy online

NEWS RELEASE

For Immediate Release

For further information contact Joyce See (65) 488-7579

SEAGATE PROVIDES DISC DRIVE TECHNOLOGY FOR ZINWELL'S DIGITAL VIDEO RECORDER

Seagate continues its growth in the Consumer Solutions market

TAIPEI, Taiwan — 04 June 2001 — Seagate Technology, the world's leading disc drive supplier for the consumer market, announced today that its award-winning U Series disc has been chosen as the storage device in Zinwell's Digital Video Recorder (DVR). The endorsement reaffirms Seagate's leading position in the Personal Video Recorder (PVR) market, and continues to fuel its growth in Consumer Solutions.

Seagate offers a variety of testing, integration and mechanical analysis services to assist in qualifying drive-capable consumer electronics (CE) devices. These services, part of the Company's CETEC (Consumer Electronics Testing & Engineering Center) are available to key OEM customers such as Zinwell as part of Seagate's partnership qualification business model.

"Leveraging its core disc drive technology for the home CE market, Seagate has scored yet another CE win with Zinwell," said Joel Stead, Seagate senior vice president of Worldwide Sales and Marketing. "With the end-to-end relationships nurtured through its unique CETEC processes and services, Seagate maintains its advantage over other disc drive makers as the premium provider of added value services for CE customers. Today, Seagate is the most prepared, experienced disc drive company with whom to partner when bringing a CE device to market."

"We are pleased to endorse Seagate's high-quality hard disc drives for use in our DVRs," said Jack Huang, Chairman of Zinwell. "Seagate's engineering team has worked closely with our engineers to test and develop our consumer products. Leveraging Seagate's disc drive technology with the unique architecture of our products, we are delivering a new dimension to home entertainment."

SCSI-3, or Ultra SCSI, was a major upheaval for the committee, with multiple implementations of the SCSI protocol to cover emerging technologies. There was also a problem with the renaming of the standard to Ultra SCSI and this created a lot of confusion among buyers, some of which, several years on, still exists. The committee also realised that there was a need to change from a parallel SCSI interface (those that use SCSI in their names) towards a serial SCSI interface (SSA, FC-AL and P1394). This has led to SCSI evolving from an interface / protocol mechanism to a device / command set mechanism. From a performance standpoint, Ultra SCSI is rated at 20 MB / sec with Wide Ultra SCSI meeting a throughput of 40 MB / sec.

Ultra2 SCSI was a revision of Ultra SCSI and, to be accurate, is simply a subset of the Ultra SCSI standard. The key to Ultra2 SCSI was that it introduced LVD (Low Voltage Differential) drives. As well as increasing the speed of the interface to 40 MB / sec for standard Ultra2 SCSI and 80 MB / sec for Wide Ultra2 SCSI, it also increased the length of the cable that could sustain the data transfer rate. This had become an increasing problem with SCSI as cable lengths had been continually shortened to support the increase in data transfer rates with poor reliability of cables over three metres. With Ultra2 SCSI suddenly the cable length was increased to a maximum of 25 metres (12 metres with 16 devices) although the payoff was that higher quality cables must be used, not the more traditional flat ribbon cables used with Ultra SCSI. LVD cables also carry a proper terminating block so the old problem with SCSI of bus termination has been removed from the user.

Connections

To understand a little more about why Ultra2 SCSI and LVD is important, it is necessary to quickly mention the different connections used by SCSI. Single-ended / differential SCSI devices are variations on the electrical portion of the SCSI specification. Single-ended devices are the most common and are much more resilient to mistakes when plugging in cables. Differential devices, however, allow much further distances between devices, with a maximum cable length of 25 metres using HVD or 12 metres using LVD, compared to six metres for single-ended SCSI. In fact, with Ultra SCSI the standard actually dictates a maximum of three metres. Ultra-SCSI further restricts this to 1½ metres if four or more devices are used, otherwise it must be no more than three metres. In the early days of differential devices it was necessary to use adapters to convert the interface to single-ended for greater flexibility.

Low voltage differential devices are designed to combine single-ended and differential into a single electrical specification providing low-cost, low power, high signal stability and long cable lengths. In early 1998 we started to see a significant number of LVD devices other than disk drives hit the market and there was a lot of discussion

"Each LUN identifies a distinct channel to which you have devices attached. Unless you have more than one controller, or a controller with more than one channel, your LUN will be 0."

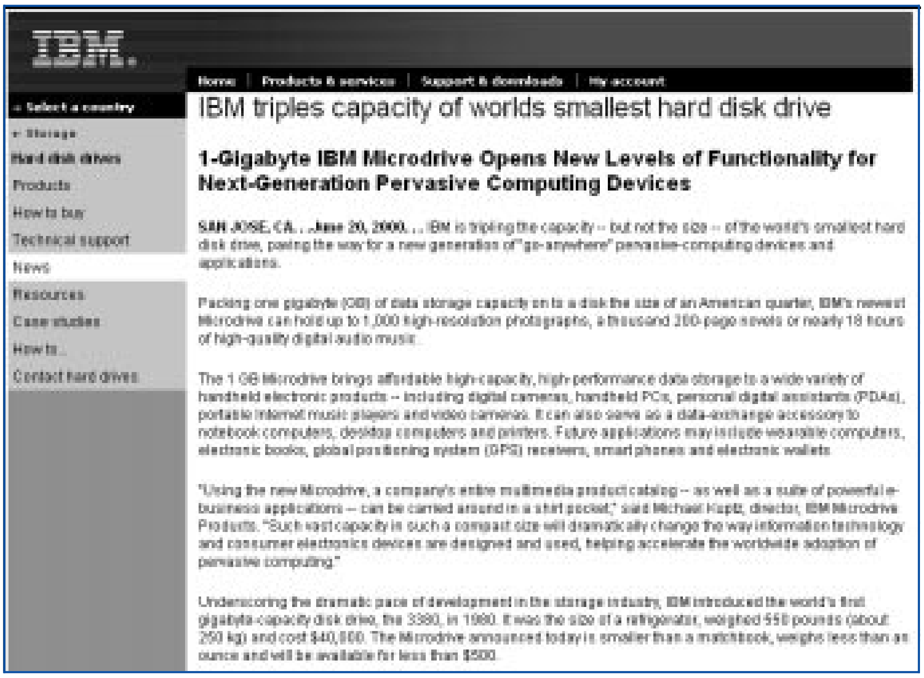

about a universal driver that would allow both single-ended and LVD mixing on the same channel. This discussion became a key component during the development of Ultra3 SCSI.

During the moves to evolve Ultra2 SCSI to the Ultra3 SCSI standard, it became clear that there were some distinct differences, again, between the vendors. The result was that of the five key proposals for Ultra3 SCSI, only three could be agreed on with two being classed as too difficult. The SCSI Trade Association defined Ultra3 SCSI as a possible set of new features that are being developed in the ANSI T10 committee, which is officially charged with the development of the SCSI technical documents. By the STA definition, an Ultra3 SCSI product may implement faster data transfers (160 MB / second) via double transition clocking (DTC), CRC, Domain Validation, Packetized Protocol or Quick Arbitrate (QA).

The result of the Ultra3 SCSI discussions was an agreement to change the SCSI numbering sequence again and to introduce a single performance speed for each future version. Starting with Ultra 160 instead of Ultra3 SCSI we already have a roadmap taking us through the next five years, with Ultra 320 and Ultra 640 drives and controllers in the labs of companies such as Seagate and Adaptec. In addition, it was felt that at this time the technology was only able to reliably support DTC, CRC and Domain Validation.

It is DTC that provides the ability for SCSI to reach 160 Mbytes / sec as it sends data on both the leading and the trailing edge of every clock pulse. With Ultra 320 and Ultra 640 the ability to maintain speeds will depend on the clock pulse being drastically reduced. This is where CRC becomes a powerful mechanism, because it goes far beyond simple parity checking and is already used inside operating systems and communication protocols without becoming a drain on resources or a bottle-neck. Reducing the clock pulse means dealing with increased electromagnetic interference and CRC is seen as being the only reliable tool.

Future SCSI

Yet it is Domain Validation that is changing the way that we will utilise SCSI in the future. At present, when you change your hard disk system, you have to change everything because it is almost impossible to mix different versions of SCSI. Using Domain Validation, however, you will be able to mix Ultra 160 and different versions of Ultra2 SCSI and Ultra SCSI on the same channel. Each device will be queried as to its operating speed and Ultra 160 controllers will then talk to each device at its optimum speed. You will need to carefully group devices together using ID numbers but the potential is vast with the ability to incrementally upgrade your disk subsys-tems. Adaptec has made the upgrade route simpler by discontinuing all disk controllers except their four new Ultra 160 controllers. As each controller supports all earlier versions of SCSI, this shouldn’t be seen as a disadvantage. You will also need to consider a motherboard upgrade if you are to reach 160 MB / second because you need a 64-bit PCI slot for the controller cards to operate at peak speed.

This doesn’t mean that you can’t use Ultra 160 cards in existing PCI slots but, if you do, you’ll only get a maximum of 80 MB / second peak throughput. Motherboards with integral Ultra 160 SCSI controllers are now widely available. The only potential fly in the ointment is that of cable length. Using just LVD devices will mean cables can be kept to 12 metres but when mixed with single ended devices, the maximum cable length drops significantly to as little as four metres.

“Using just LVD devices will mean cables can be kept to 12 metres but when mixed with single ended devices, the maximum cable length drops significantly to as little as four metres.”

Logical Units

One of the most commonly misunderstood pieces of SCSI technology is Logical Unit Numbers (LUNs). To keep it simple, each LUN identifies a distinct channel to which you have devices attached. Unless you have more than one controller, or a controller with more than one channel, your LUN will be 0.

Termination is another of the major failures of SCSI installations and attempts by manufacturers to simplify this area have often led to further confusion as they change the termination mechanism.

Each end of a SCSI chain must be terminated with the appropriate resistors and on older devices this is very simple since there are three resistor packs which can be unplugged. When you do unplug resistors, remember to bag them up and tape them to the side of the drive in case you need them again. When refitting them, there is a small dot on each resistor to tell you the correct alignment of the pack. Later SCSI devices, particularly the HP CD writers, use a simple jumper rather than resistor packs. This has been known to lead to confusion over the termination of the device as people confuse the termination jumper, spin up jumper and ID jumpers. Those that have jumpers or switches use a technique called active termination and you must check your documentation carefully when installing SCSI devices. LVD has removed all the problems of knowing how to terminate the chain by having a fixed resistor at the end of the cable.

The versions of SCSI that we have been talking about are all based on parallel technology and several years ago, during the negotiations for SCSI 2, it was generally accepted that parallel SCSI was close to its maximum speed of between 10 and 20 MB / sec. The idea that we would be taking the technology to 160 MB / sec and be planning to reach 640 MB / sec was unthinkable. Such speeds were held to be only achievable when the problems of serial SCSI were resolved and with several competing technologies, it was uncertain as to where the R&D money would be invested. One of the key reasons for this belief was the problem of reliable signalling and, in particular, the cabling. The higher the number of signals needed to accomplish a task, the more expensive the connectors become. Even more important is the risk of cross talk and with cables that are carrying data, this means corruption whilst in transit. Fibre Channel has the ability to work with either fibre optic, co-axial or 9-pin twisted pair cables, SSA uses 9-pin twisted pair and P1394 only needs six leads but must also provide power supply for the device.

We've seen how LVD provides a length of 12 metres for all speeds and a maximum of 31 devices on a single channel whilst P1394 (FireWire) allows a separation of 4½ metres between devices with 16 devices on a channel. Compare this to Fibre Channel, which allows 50 metres between devices when using twisted-pair and as much as 10km when using optical fibre with a maximum number of devices anywhere between 127 and 16 million.

Firewire

The current Big Thing in storage is P1394 (FireWire or i.Link) technology. With initial speeds ranging from 100 MB / sec to 400 MB / sec, these devices are likely to be attractive to a number of newer technologies such as set-top boxes where they can be used for video-on-demand products. The current specification allows for the doubling of speed to over 800 MB / sec.

The only realistic competitor to FireWire as a serial SCSI-based interface is Fibre Channel. Over the last six years we have listened to Fibre Channel vendors trying to promote the technology as the solution to network storage. Storage Area Networking (SAN) relies heavily on the success of Fibre Channel and whilst there are several successful products in the SAN market, there are significant problems with interoperability and management. In its most basic implementation, Fibre Channel operates at 100 MB / sec on a single channel and 200 MB / sec when running on a dual channel. The actual standard currently provides for this to be extended to speeds of 1GB / sec and 2GB / sec, providing issues of signalling and reliable management software can be overcome.

"The higher the number of signals needed to accomplish a task, the more expensive the connectors become. Even more important is the risk of cross talk."

PCSA

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.