Print Screen

Vol. XIIII April 1996 No. 04

Newsletter of Stanford / Palo Alto Users Group for PC

A Non-Profit / Educational Organization

General Meeting - May 08th @ 7:30 pm - MediaCity

ParaDisk

Bill Weber
Web-Star Computer Group
(415) 968-7351 or
billw@mediacity.com.

The ParaDisk PD350 is a Hard Disk drive kit which allows you to use up to an 8.4 gigabyte hard drive through the parallel port of any IBM compatible notebook, laptop, or desktop computer. It is SPP/EPP/ECP compatible and comes with all the necessary device drivers to get it up and running quickly.

The ParaDisk PD350 supports the most popular hard disk drives on the market. You can order your ParaDrive 525CD with a drive installed or you can order the kit and install the drive yourself.

Continued on Page 12

Iomega Zip Drive FAQs

FaxBack # 2002 12/95
(800) 456-5522
(801) 778-5763

Condensed for PrintScreen

What is the price for the drive?
The street price for the drive will be about $199.99. A 100 MB starter disk is included with this drive. Additional 100MB disk will have a street price of about $19.95 for each 100MB disk.

What Zip Drive models are available?
A Zip SCSI model and a Zip Parallel Port Model. The SCSI model can be used with Macintosh or a PC that is using a Zip compatible SCSI adapter.

Continued on Page 12

Video Via MPEG

Philip Albinus
Windows Magazine

While there is more than one way to get video on your PC, the PC industry appears to have settled on a single standard for full-screen, full-motion video. MPEG, named after the Motion Picture Experts Group, is an algorithm for compressing and decompressing digital video and audio. Since uncompressed video files are so large--a six-minute video clip with audio could occupy 35MB of hard disk space--you need to compress the files for efficient storage. You then need decompression software to recover the original digital data in order to view the video.

Continued on Page 13

A Message From The New Prez

Bob Mitchell

So there I was sitting there paying close attention to center stage at the last meeting thinking how nice it was to have somebody like Brian Christopher and all of the volunteers working with him taking on the responsibility of making the group work and all I had to do was sit back and enjoy it. That all changed with a tap on the shoulder from Bev Altman. So much for my leisure time in retirement. Welcome to a new responsibility!

First off, a very large thanks to Brian and each of the individuals on the staff this last year who did their part in putting it all together. It is a major effort to obtain good guests for our meetings each month. Brilliant and talented individuals willing to demonstrate the applications that you want to see are not necessarily standing in line waiting for invitations from user groups. The contributions of personal time put into the planning and execution of the meetings, publishing and mailing Print Screen, updating the Web Page, keeping track of membership dues, finances, publicity, picking up & distributing the mail, preparing the disk-of-the-month and all of the other behind the scene activities, are sincerely appreciated

> The real truth is, the habit of a lifetime of challenges and taking on responsibility is hard to break.

So who is Bob Mitchell? (not Robert Mitchell...he's the other club member who has the advantage of fewer miles and certainly more energy). The older fella just retired last September from 43 years in "big time show biz". I've spent the last 41 years with KGO-TV (Channel 7) in San Francisco. That's one of the ABC stations along with the ABC network and all other ABC properties that Capital Cities Broadcasting acquired a few years back. As you know it recently became part of the Disney Corporation. I left the station before having had the distinct experience of receiving a paycheck with a picture of a mouse on it, however Mickey does appear on the Disney stock certificates that were received to replace CCB certificates. During that period of 41 years I served as Film Director, Assistant Director of Programming and finally as the Program Operations Director of KGO-TV. A parent of three boys (now adults), there was past involvement in all the kid activities such as scouts, (as Scout master and Committee Chairman), and Little League (as League president). Other volunteer activities include the Chairman of a recreational group, membership on a church financial council and president of that church's Pastoral Council. This is not to say I'm that good but rather that I'm too dumb to say "No" when asked to take on a leadership role. The real truth is, the habit of a lifetime of challenges and taking on responsibility is hard to break.

There'll be more to come in the future but the first priority is to obtain from you, your interests and your likes and dislikes in the matter of SPAUG. In that connection, you will find enclosed in the SPAUG Print Screen a list of questions. Your answers and the answers of *every member* of SPAUG will serve to determine what we will strive to provide for you in future meetings. We want to offer you a compelling reason to attend every meeting. It is my firm belief that the reason for the existence of SPAUG still remains. The continuing success of other user groups attests to this. Based on this it is our intention (for openers) to restore the interest in the group and the attendance figures we once had. With your cooperation we stand a pretty fair chance of making it work.

Are you willing?

Tell the New Prez what you think!

The Official 1996 Membership Survey starts on the next page. Fax or mail the completed survey to Bob and be eligible to WIN a prize for the best ideas.

Stanford Palo Alto User's Group

Bob Mitchell - e-mail: TVBob1@aol.com
SPAUG, POB 3738, Stanford, CA 94309-3738
Fax (415) 364-3426

We need your help! We want you to tell us what you want!

In an effort to offer a program for each monthly meeting that will offer something to the most members most of the time we need your input. The attached list of questions will take a few minutes of your time but we're going to do everything we can to insure that it will have been time well spent. *Every* person's input is necessary.

The following questions will establish a general level of expertise and interest in our group. Please indicate what you consider yourself in the following areas (Don't be modest but don't feel shy about stating that you're unfamiliar with some areas):

Personal Computing in general:
☐ Very Knowledgeable ☐ Above Average ☐ Average ☐ Little knowledge

PC Hardware (Knowledge of upgrading, hard drives, memory, graphics accelerators, etc.)
☐ Very Knowledgeable ☐ Above Average ☐ Average ☐ Little knowledge

DOS and DOS commands in general
☐ Very Knowledgeable ☐ Above Average ☐ Average ☐ Little knowledge

Windows 3.x in general
☐ Very Knowledgeable ☐ Above Average ☐ Average ☐ Still a DOS diehard

Windows '95 in general
☐ Very Knowledgeable ☐ Above Average ☐ Average ☐ Little knowledge

If you don't have it now, do you plan to upgrade to "95" in the reasonably near future?
☐ Yes ☐ No ☐ Already Have it

Information Systems (American On Line, Compuserve, Prodigy, etc.)
☐ Very Knowledgeable ☐ Above Average ☐ Average ☐ Little knowledge

The Internet (Including Internet Service Providers/Internet Software)
☐ Very Knowledgeable ☐ Above Average ☐ Average ☐ Little knowledge

Do you have a modem? ☐ Yes ☐ No

What in the general program format of our meetings do you like best?

_____

What would you like to see changed or eliminated?

_____

What would you like to see added to our meetings that we currently do not have?

_____

What hardware would you like to see demonstrated?

_____

What software applications would you like to see demonstrated?

_____

How interested are you in having game software demonstrated?
☐ Very interested ☐ Somewhat interested ☐ not interested

How interested are you in having educational software demonstrated?
☐ Very interested ☐ Somewhat interested ☐ not interested

How interested are you in having various areas of the Internet explored at each meeting?
☐ Very interested ☐ Somewhat interested ☐ not interested

How interested are you having LAN hardware, software and operation demonstrated?
☐ Very interested ☐ Somewhat interested ☐ not interested

If you have not attended meetings in the past few months, why have you elected not to attend and what would it take to get you back?_____

_____

Suggestions that would be a compelling reason for you to attend future meetings: _____

_____

Your Name (optional) _____

Your E-Mail Address (if any) (optional) _____

The Digital Video Revolution

Microsoft TechNet CD

The increasing synergy among the computer, consumer electronics, entertainment and telecommunications industries all of which are moving to digital media is fueling multimedia computing that incorporates video, sound, graphics, text and animation. The result will likely be one macro digital infrastructure where any equipment with a microprocessor computers, video players, stereo sound systems, electronic games -- can access, manipulate and deliver any type of digital data. The distinctions among such hardware will become less significant, and users will be able to receive and easily use information in more natural and intuitive ways.

Nowhere is the impact of this digital revolution greater than in video. The popularity of television (particularly the growth of cable), VCRs and home video rentals would suggest there's little wrong with the analog, or non-digital, form of video. Yet traditional analog video offers slower access through linear searching, lacks interactivity and loses image quality with each generation of reproduction.The significant installed base of analog video systems (particularly in education and training) means that an ideal multimedia computing platform should continue to support analog video. But analog video, in itself, is not ideal. It will likely become less of a market factor over time as people turn to digital video, which overcomes the problems associated with analog forms. For example, with digital video:

·searching is facilitated by fast, random access
·interactivity is optimized
·copies can be easily reproduced with no loss of image clarity
·transmission is possible across local and wide area networks
·digital images can be easily and fully manipulated

Expect to find digital video behind applications as diverse as teleconferencing, entertainment broadcast/cablecast and videotelephony all using a multimedia-based desktop computer as the user's access to digital media. As digital video becomes increasingly popular, the distinctions among television broadcasts, video rentals and computer processing should blur.

A user, for example, may order and download video entertainment and application software from a "digital television network" direct to his integrated computer/VCR, add digitized representations of himself, customize the story via on-screen editing and then output a finished disk or tape. Or, consider the advantages of a palmtop unit that can download mapping data from remote sources via wireless networks, calculate the optimum route to one's destination and the time required to get there, display the route graphically and then provide video and audio narration of tourist attractions along the way.

Digital Video Compression: One Problem, Many Solutions.

Vast amounts of data are needed to represent moving video images. A single video frame could take up to a megabyte of storage. A full second of video could take about 28MB of storage. At that rate, a 650MB CD-ROM, despite its large storage capacity, would hold only about 20 seconds of video not counting any associated audio data. Meanwhile, digital video applications used over networks and phone lines also run into problems because of the heavy data transmission that must take place over limited-bandwidth lines. High speed LANs support only 10 Mbps and T1 phone lines support an even slighter 1.5 Mbps.

Vendors and industry groups have devised methods that compress data at ratios up to 200:1. That extends single-disc play-time up to 72 minutes, and makes video conferencing feasible and economical over low-bandwidth phone lines. All of these methods use mathematical formula to delete redundant information from images, store them in short-hand form, and then expand them for display

and use. But these methods vary in how, and how much, data they compress, and in their suitability for various applications. Some applications require little or no data storage, but call for equal use of real-time video compression and decompression; in video e-mail, videotelephony and teleconferencing, for example, video data is constantly being created, compressed, transmitted and decompressed for display. Other applications require only a single, non-real-time compression when the software is produced, then use frequent, real-time decompression when the data is called for viewing. These applications include electronic books, training, kiosks and point-of-purchase, and entertainment.

The various digital video solutions that provide options for this range in uses. Software-only methods, for example, make it easy to distribute video content. Full frame methods are appropriate for video editing. The following are descriptions of the major video compression formats.

Audio Video Interleaved Technology.

Microsoft's own format is implemented completely in software. It works with standard Multimedia PCs and runs without extra and potentially costly video compression hardware. In contrast to hardware-assisted methods that can require video compression boards with digital signal processors costing from $1,500 to $4,000, Microsoft's format can bring digital video within financial reach of the broadest possible market. A typical sequence of Audio Video Interleaved technology permits 15 frames-per-second of motion video to run within windows at resolutions of 160 x 120 pixels and includes support for 8-bit audio. That makes it a successful solution for mixed text/video applications -- such as interactive newsletters and books that would run video in a window much as printed books and magazines today have text pages that include photos. Audio Video Interleaved technology allows ISVs to include a video solution within the economic and technological reach of most multimedia users. It is not intended as a solution where full-screen, full-motion video is required -- for example, in interactive training manuals depicting complex images.

Motion-Joint Photographic Experts Group (JPEG) specification.
Because motion video is composed of a series of still images, one compression plan is based on JPEG, the standard for still image compression promulgated by the International Organization for Standardization (ISO) and the Consultative Committee for Telephony and Telegraphy (CCITT). Among M-JPEG's advantages: It produces a relatively high-level image that makes it suitable for all but the most image-sensitive applications. Unlike some compression plans still in development, the industry has endorsed M-JPEG and it's available today. Also, because it provides full access to randomly selected frames, it's useful for the cut-and-paste type of operations of non-linear video editing. Among its disadvantages: M-JPEG's compression ratio - about 25:1 -- makes it suitable for magnetic hard disk drives, but not for the CD-ROMs and T1 phone lines that will have to carry the bulk of video data. This format also lacks an associated audio compression specification, forcing vendors to adopt or devise their own. To achieve real-time compression rates of 30 frames per second, M-JPEG needs an assist from digital signal processing hardware. And because it's based on a single-image compression standard, M-JPEG doesn't consider the similarities between adjacent images in a motion video. While that facilitates non-linear editing, it also keeps M-JPEG from compressing data as effectively as other specifications.

Moving Pictures Experts Group (MPEG) specification
MPEG is being designed from the ground-up as a motion video standard. The ISO committee working on this has included delta framing, also called temporal compression. Delta framing allows MPEG to compare adjacent frames in a video, wringing out more of the redundant data. As a result, it achieves

compression rates of 1.2 Mbps (up to 8 Mbps) and ratios up to three-times as great as for M-JPEG, making it suitable for CD-ROMs and for T1 telephone lines. The standard also includes an 8:1 compression ratio for associated audio data. These advantages make MPEG a good format for electronic books, entertainment, and other applications that require the user to decompress pre-packaged and pre-compressed video. However, MPEG also requires additional video hardware and the standard is not yet set. Expect to see movement toward approving the standard later this year.

Digital Video Interactive. Intel acquired this technology from RCA's Sarnoff Research Institute in 1988 and has since developed it. DVI displays a high-quality image at quarter-screen size, and can display low-resolution, full-motion video in a full-screen window. DVI is implemented via two programmable Action Media chips: a pixel processor handles decompression and a display processor handles video output. Intel began by manufacturing video hardware.

Now, Intel also licenses the software, and is working with other PC makers to include its chipset on their motherboards. Originally costing more than $20,000, DVI has come down in cost by a factor of 10, and Intel reportedly expects a single-chip solution will ultimately come down in cost by another 10-fold factor. DVI has traditionally supported two modes: a high-quality production level video (PLV) and a lower-quality real-time video (RTV). Title developers send their video on magnetic media to Intel, where a 64-processor computer converts it to digital-based PLV at 30 frames-per-second and 256 x 240 resolution. Or, developers can speed the conversion process and trade down in resolution by creating their own RTV compressed video (30 fps, 128 x 120 resolution) using the Action Media board. In addition to these two modes, DVI supports JPEG and Intel is expanding it to support MPEG and others.

EDO RAM & Burst Caching

Kai Kaltenbach, Microsoft Premier Corporate Support, 9/95

Why Cache?

When a microprocessor asks for information faster than system RAM can deliver it, the processor goes into a *wait state*. Essentially the processor is sitting around doing nothing until the system RAM is ready to deliver the information it asked for. This greatly slows down system operation. When a system is running without encountering wait states, it is said to be in *zero wait-state* operation, and runs much faster.

Memory speed is measured in nanoseconds (ns). The fastest affordable DRAM (Dynamic RAM) memory chips are 60-70ns. For a processor to operate in zero wait-states at a system-board speed of 33MHz (as with a 486DX/33, 486DX2/66, 486DX4/100) the system RAM would have to have a speed of 30ns - prohibitively expensive. For zero wait-states at a system board speed of 66MHz (as with a Pentium 66, 100 or 133) the memory would have to operate at 15ns! What's more, it would have to be more expensive SRAM (Static RAM), which is faster than DRAM because it doesn't require the system to refresh its contents periodically. At the time of this writing, 15ns SRAM is *over ten times the cost* of standard 70ns DRAM.

This is where memory caching comes in, making today's systems possible at a reasonable cost. You may be familiar with using a disk cache, such as Microsoft SmartDrive, which uses a small RAM buffer to speed up access to a large hard disk. Memory caching uses a small buffer of very fast RAM to speed up a large bank of slower RAM.

Level 1 and Level 2

All Intel processors since the advent of the 486 are equipped with an integral cache of 8kb-16kb in size. When a RAM cache is

built into a CPU, it's known as a Level 1 (L1) cache.

Most systems today use a second RAM cache built onto the system board, called a Level 2 (L2) cache.

Hits and Misses

The cache is managed by an 'intelligent' circuit called the cache controller. A system with both an L1 and an L2 cache has two cache controllers; one on the CPU chip itself, and one on the motherboard. The cache controller uses various prediction algorithms to enhance cache performance. For example, it attempts to predict what memory segments the processor will ask for next, and read those segments into the cache before the processor asks for them. This is known as read-ahead caching.

When the processor asks for some data from memory, and that data can be delivered directly from the cache RAM, that's a *cache hit*. When the system has to take the performance hit of going to the main bank of memory to retrieve the data, that's a *cache miss*. The percentage of cache hits versus cache misses determines the system's performance versus other systems with the identical CPU.

Why an L2 Cache

The cache hit/miss ratio, and therefore overall system performance, is determined by several factors (see below). One of the crucial factors is the ratio between the size of the cache and the size of system RAM. As previously noted, L1 caches are generally 8kb-16kb in size. This tiny cache is not sufficient to produce a large cache hit/miss ratio with any significant amount of system RAM. Therefore, performance suffers significantly without an L2 cache. It is not uncommon, for example, for a 486 system with an efficient L2 cache to far outperform a Pentium system without a cache. In a recent industry magazine test of notebook computers, a 486 machine (with L2 cache) outperformed a Pentium 90 machine (without L2 cache) by 30%.

L2 cache sizes range from 64K-1024K, with 256K being by far the most common size. More on L2 cache sizing later.

Cache Performance Considerations

The following factors influence the performance of a cached system:

·Cache controller design efficiency

All caches are not created equal, even if they are of equal size. Given the trade press focus on cache size, most purchasers simply ask for a cache of a particular size, and don't focus specifically on performance measurements. Unfortunately, this has led some system vendors to develop very low-cost caching systems that allow them to advertise a 256K cache without regard to the performance of that cache. It's entirely possible, in fact common, for a smaller, well designed cache to outperform a larger, badly designed cache.

·Cache size to system RAM ratio

See above. All things being equal as far as cache architecture and controller design is concerned, a larger cache-to-system RAM ratio will provide better system performance, up to a point. You quickly reach a point of diminishing returns. The important thing to remember in general, is that if you wanted to maintain the same cache hit ratio when you double the amount of system RAM, you would have to double the amount of cache RAM as well (although there are other factors that do not make this a linear relationship).

·Cache RAM speed & System RAM speed

Most system boards are designed for a particular speed of system RAM and cache RAM. There are some exceptions that allow you to tune the system's cache parameters to different speeds of memory. For example, it is getting more and more common for system boards to offer a 70ns/60ns switchable memory speed option. Without such an option, adding faster system RAM than the board is designed for won't provide

any performance benefits. In Pentium systems, 20ns cache SRAM is generally used for 50-60MHz system boards (using the Pentium 75/90/100/120), and 15ns cache SRAM is normally utilized for 66MHz system boards (using the Pentium 100/133). Cache SRAM at speeds up to 8ns has recently become available, although rare and expensive.

·Software

Cache controllers are usually programmed with algorithms based on statistical analysis of memory access by popular operating systems. Many cache controllers are optimized for either 16-bit or 32-bit software systems. If your particular software accesses memory in a different pattern than the cache controller was optimized for, you can get significantly higher or lower than theoretical (benchmarked) efficiency. Upgrading an operating system from 16-bit to 32-bit can change system *hardware* performance dramatically in some cases. When evaluating systems for purchase, make sure to benchmark the systems under your operating system of choice, and if possible, the operating system you plan to implement *next*.

Software tools are available for measuring cache efficiency, such as those from Sofwin Laboratories (800-339-2579). Sofwin tools in particular have a feature to show whether a system cache is optimized for 16-bit or 32-bit operations. While such measurements can lend insight into system design, they are arguably less useful for purchasing decisions, because your real-world performance will depend on the software and operating system being used.

New RAM and Cache Technologies

Recently, several new RAM and cache technologies were introduced. These include:

·EDO DRAM

Enhanced Data Output (EDO) DRAM provides faster data throughput that partially obviates the need for an L2 cache. Systems using EDO DRAM and no L2 cache will be faster than similar systems using regular DRAM, but not as fast as systems with an L2 cache. EDO DRAM also provides a performance benefit when used with an L2 cache, but industry magazine test centers have reported that the performance difference in that case is less than 5%. Theoretically, EDO DRAM doesn't cost any more to manufacture than regular DRAM, so eventually EDO DRAM may replace regular DRAM. But at the time of this writing, EDO DRAM was significantly more than 5% more expensive than regular DRAM, and probably not worth the price/performance ratio on systems with an L2 cache.

·EDRAM

You can think of Enhanced DRAM (EDRAM) as RAM that carries its own cache on each module. In an EDRAM-based system, essentially the entire system memory bank is the cache. This can provide dramatic performance improvements. However, at this time, EDRAM is scarce, very expensive and has not been adopted by many system vendors.

·Burst Cache (Pipeline Burst or Synchronous Burst)

Burst cache technology brings a very large performance advantage to the Pentium playing field, made possible by Intel's recent introduction of the Triton chipset for Pentium systems, and also supported by other chipset vendors. Industry magazine tests show that burst cache equipped systems outperform their standard cache counterparts by 20% or more. In fact, the performance benefit is frequently more than the performance difference between Pentium chip classes, i.e. a Pentium 90 with burst cache has been shown to outperform a Pentium 100 with normal cache. Since the difference in price between normal cache and burst cache is usually less than the difference in price

between Pentium chip classes, it only makes sense to standardize on burst cache systems. There are other considerations of course, because the Intel Triton chipset does not support some features that are required by corporate standards, such as multiprocessor operation, memory parity, and over 128MB of system RAM.

General Recommendations

The following general guidelines will help you specify systems that will give you the best possible performance under Windows 95 and Windows NT. However, it's important to remember that the key measurement is how your software performs on a given system versus that system's cost, service and warranty, reliability and compatibility. And needless to say, the other components of a system; hard disk, video card, etc; can affect performance as much as anything else. The key factor is *balance*; that all the components of the system are equal in performance, and no significant bottlenecks exist. That's why a real-world benchmark of your particular operating system and applications is so important.

·L2 Caching

Industry publications clearly show the large performance advantage of an L2 cache. Since L2 caching is essentially an industry standard today, the only difficult choices you may have to make will be in the area of notebook computers, which have not yet embraced the L2 cache in significant numbers.

·L2 Cache Sizing

You will find a lot of varying opinions on the benefits of various L2 cache sizes. The concensus among industry insiders seems to be that you can get by with 128K of L2 cache up to 8MB DRAM, with 256K of L2 the standard from 16MB-32MB, and 512K optimal for 32MB and up. Again, these figures are rough estimates, and performance can vary widely due to the cache performance considerations discussed earlier.

·Burst caching

Published benchmarks definitely point towards the burst cache superiority. And since its performance boost costs less than the equivalent investment in CPU power, it's said to be a smart choice for desktop machines. You may want to forego burst caching for servers, since the system board chipset that supports burst caching doesn't provide some mission-critical features at this time (see above).

·Memory Technology

Since alternative memory technologies (EDRAM, and EDO RAM in systems with L2 cache) have not yet been shown in the media to provide a demonstrable price/performance ratio increase over standard DRAM, that remains the standard today. If you want the fastest possible system, and you're buying from a hardware vendor that doesn't give you a price hit for EDO DRAM, then by all means use it.

Conclusion

Implementing a cache system will heighten system speed and performance at a reasonable cost. By understanding how cache works, the types of cache systems available, and the factors that affect its performance one can make an informed decision configuring system cache in Windows NT and Win95.

Prizes- Fun - Excitement

Complete your Membership Survey and send it to Bob Mitchell, to be eligible for all the above. "Hints for the Winners" will include phrases like "Bob, I want to help, what can I do." Or another good one is "Bob, I heard the Group received several Win95 videos so I want to host a Win95 SIG using the videos for training." Other hot phrases are, "Bob, I just upgraded to a new Pentium system and would like to donate my 486 DX4 system to the Group." Excellent idea as SPAUG can give the donor a tax deductible receipt.

WIN95 Tips

Windows Magazine

Drivers' Ed

Win95 supports a large number of popular video cards from about 30 different vendors but that's not the same as saying it supports every feature of every card. If your card, or some special feature you need or want on your card, isn't supported, you'll need your old-style Win3.x video setup disks. Similarly, some video cards need the help of a small program (often called "monitor" and placed in the AUTOEXEC file) to properly communicate with the monitor. Win95 usually leaves this program alone which it should. But it's a wise precaution to make sure you hang onto your old Win 3.x drivers and setup disk until you're 100 percent sure your video is working properly under Win95.

Screen Gems

Logow.sys and Logos.sys are graphic screens shown during Windows shutdown. The former is the "wait" screen, and the latter is the "it is now safe to shutdown" screen. Unlike the "RLE" weirdness of past Windows versions, these files are standard bitmaps you can modify with the Paint applet or another bitmap editor. Make a backup copy first so you can get the system back the way it was, then have at it. Make your system look the way you want when it's shutting down.

Give Warm Boots the Boot

Normally, when you select Restart from the shutdown menu, the system "warm boots"—it shuts down as though you'd pressed control-alt-delete (though in a more orderly and less dangerous way). But if you just want to restart Windows without restarting your whole system, click the Restart Computer option on the Shutdown menu, hold down the Shift key and click on the Yes confirmation option. Windows 95 will now shutdown and restart without rebooting your whole system.

Measuring Up
What You See Is What You Get

Karen Strauss - Windows Magazine

You may wonder why the viewable areas of 20- and 21-inch monitors don't match the advertised sizes. Well, monitor manufacturers borrowed the idea from television manufacturers, who label products based on the size of the tube. They measure the tube diagonally, including the size of the glass fromedge to edge--before it's placed inside the case--and use this number to label the monitor. But a bezel is placed around the edges of the tube, which typically cuts down the viewable area of the screen. Some monitors lose additional viewing area because of the shadow-mask technology used to sharpen the image.

For this review, we measured the maximum area Windows occupies, ignoring all of the dead space. Large monitors are consistent (see below) with smaller monitors, generally providing a viewing area 1-2 inches smaller than the listed tube size.

Despite this confusion, there is a standard way to judge the picture size of a monitor. VESA (the Video Electronics Standards Association) has established a method of defining the usable image area for displays. Called VIAD (Video Image Area Definition), this standard defines the video-image area as the largest rectangular image, listed in horizontal length and vertical height in millimeters. Most manufacturers include this information on their product sheets, so look for it when you comparison shop.

Monitor Size (Inches) / Actual Image Size (Inches)

·14 / 12.7 to 12.8

·15 / 13.5 to 13.75

·17 / 15.25 to 16.33

·20-21 / 18.75 to 19.88

Paradisk contined from Page 1

The ParaDisk PD350 is convenient. It provides a printer pass-through port for simultaneous use of a printer.

The ParaDisk PD350 travels easily. Not only is it lightweight, but it comes with a delux carry bag.

Features

- Built-in ATAPI (Enhanced IDE) Interface
- Supports up to 8.4 gigabyte hard Drives
- Hard Disk Auto-detection
- Fast access time
- SPP - 350 KB/sec.
- EPP - 700KB/sec.
- No slot required
- Compatible with SPP/EPP/ECP parallel ports
- Printer port pass-through for simultaneous printer use.
- Ideal for Notebooks or Desktops
- Easy to install and use
- Supplied Device Drivers
- Includes a deluxe carry bag.

Zip Drive continued from Page 1

The SCSI model has two 25 pin SCSI ports (for chaining), a SCSI ID switch, and a SCSI termination switch. The Parallel Port model can only be used with a PC and connect to the computers Parallel Port.

What is included with the Zip Drive?
The Zip solution comes complete with drive, cable, power supply, operating software, and a starter 100MB disk with Zip tools software.

Which SCSI adapters can be used with the SCSI model?
Iomega provides a optional low cost SCSI adapter called the Zip Zoom Accelerator. Other ASPI compatible SCSI adapters can be used in conjunction with Iomega SCSI driver software.

What Operating Systems are supported?
MSDOS and PCDOS 4.x or greater, Win95 and Win3.1,OS/2 v.2.0 or greater, OS/2 Warp, with the purchase of the optional Iomega OS/2 driver software package (OAD for OS/2 v.2.34 or greater). Macintosh System 6.0.8 or greater (System 7 or greater is required for full functionality)

Does the Parallel Port model allow printer connection?
Yes, the Parallel Port Zip model has a second port that lets you connect to your printer.

Is there a speed difference between the SCSI and Parallel models on PC's?
The SCSI model is 2-5 times faster, depending on your system configuration, however, the parallel port model is easily connected to any PC.

Are the disks rugged?
The disk are very rugged. They can withstand a 1000Gs shock (an 8 foot drop) without any harm to your data.

How long is the disk media shelf life?
Zip disks have a long 10 year shelf life.

Can the Zip Drive be chained with other SCSI devices?
Yes, the SCSI model can be put in a SCSI chain just like any SCSI device.

How is power provided to the Zip Drive?
It uses an AC power adapter.

How do I get the accessories for the Zip Drive?
Call Iomega at 1-800-MY-STUFF.

What are the specifications for the Zip Drive?
Max Sustained transfer rate - up to 1.4MB/sec.
Average Seek time - 29 milliseconds
Buffer size - 32K
SCSI throughput - up to 60MB/minute
Parallel throughput - up to 20MB/minute
Disk storage capacity - 100MB or 25MBs
Warranty - 1 year on drive

Video continued from Page 1

With a video card, MPEG software and a fast CPU, your PC can display full-screen video files at 30 frames per second--the standard used by broadcast professionals. As more video card manufacturers choose MPEG for full-screen, full-motion video, you won't have to view video clips the size of a matchbook anymore.

The MPEG 1 algorithm has a compression ratio of 150: 1. This standard CODEC (coder/decoder) is prevalent in PCs and is optimized for playback from CD-ROMs. If you are viewing full-screen, full-motion video on a PC, it's most likely using MPEG 1. PCs can play .AVI video files with Microsoft Video for Windows, but the video quality rapidly degrades as you expand the window to near full-screen size.

Right now, there are two ways to play back MPEG: You can use a hardware-assisted or software-based method. Hardware, for example, can be an MPEG daughter card that works in tandem with an existing video accelerator card, or an all-in-one card such as the Jazz Jakarta, which includes both MPEG and graphics acceleration chips. Some video card makers, on the other hand, offer an after-the-fact MPEG software solution. The Number Nine and Diamond Multimedia cards we received both included a coupon in the box for free MPEG software. MPEG software is hardly the best solution, though. Experts doubt that an MPEG software solution is viable on a PC powered by anything less than a 90MHz Pentium CPU. There are also some rough edges. For example, getting the audio and video to synchronize--such as having an actor's lips and mouth move at same time--is often difficult.

Some argue that a hardware MPEG solution is required to free up an already overburdened microprocessor. By using an MPEG-dedicated chip on the video card or on the motherboard, as some companies are considering, the MPEG hardware solution frees up the CPU.

MPEG 2 is a follow-up to MPEG 1 and is aimed at the broadcast industry. According to John Reno, product manager for C-Cube Microsystems, MPEG 2 resolves four times as many pixels (720x480 pixels) as MPEG 1. "You really don't see that on the PC environment now," said Reno. "In two to three years, you'll see higher speed drives, cable or network connection for MPEG 2 for the PC. "Whether the final video solution is MPEG hardware or software, the pundits agree that the video output on a PC must match the performance of television and the price must ultimately drop to the $200 range to gain favor with consumers.

Spy / Thriller / Mystery Adventure on the Web

Treasure Quest - Sirius Publishing
http://www.treasurequest.com

Psychic Detective - Electronic Arts
http://www.ea.com/eastudios.html

The Dame Was Loaded - Philips
http://www.philipsmedia.com/media/games

Connections - Discovery Channel
http://www.discovery.com

The Pandora Directive - Access Software
http://www.accesssoftware.com

The Dig - LucasArts
http://www.lucasarts.com

Fox Hunt - Capcom Digital Studios
http://www.capcoment.com

Secrets of the Lost Dynasty - Smokin' Digital
http://www.smokin.com

The 11th Hour - Virgin Interactive
http://www.vie.com

Spycraft - Activision
http://www.activision.com

Future Vision

Helen Custer
Inside Windows NT

In both the CISC (complex instruction set computers) and RISC (reduced instruction set computers) arenas, promising technologies have emerged rapidly. Microsoft saw that in order to exploit these and other hardware advances, it needed to produce an operating system for the 1990s - one that was portable and able to move easily from one hardware platform to another. Although Microsoft and IBM created the OS/2 operating system in the 1980s, Microsoft recognized that the system had many shortcomings, the most obvious being that OS/2 is not portable. It was written in assembly language to run on single-processor, Intel 80286 computers. Rather than try to overhaul the OS/2 system software, Microsoft decided to build a new, portable operating system from the ground up.

Staff

Bob Mitchell
President (415) 368-9530
Walter Varner
Accounting (408) 739-3488
Mildred Kohn
Secretary (415) 949-1833
Beverly Altman
Director at Large
Membership List (415) 329-8252
Kendric Smith
WebMaster (415) 493-7210
Brian Christopher
Editor (415) 952-5632

Membership Internet Addresses

Beverly Altman	hfdj68a@prodigy.com
Audrey Borland	baha@hooked.net
Frank Campbell	rfvw93a@prodigy.com
Brian Christopher	brianc@mediacity.com
Jim Dinkey	jimdinkey@aol.com
Karen Fung	fung@hoover.stanford.edu
Dick Harding	73467.2445@compuserve.com
Nancy Helmy	nhelmy@concentric.net
Stein Hoffmoen	steinh@ix.netcom.com
Jess Kanarek	jessejk@aol.com
Floyd Kessler	75013.1017@compuserve.com
Don Kleyensteuber	donaldk@ix.netcom.com
Mildred Kohn	mildredk@aol.com
Aldora Lee	DMRAnalyst@aol.com
Bill McElhinney	mcels@aol.com
Bob Mitchell	TVBob1@aol.com
Larry Mehl	mehl@well.com
Clark Moore	vrwg62a@prodigy.com
Jim Powell	jimp1938@hooked.net
Norm Rossen	dmtt02a@prodigy.com
Ron Seltzer	ras011@aol.com
John Sessoms	adtek201@aol.com
Lamont Shadowens	shad@ix.netcom.com
Russ Smith	russfs@delphi.com
Kendric Smith	kendric@aol.com
Barry Smith	76250.2277@compuserve.com
Ben Swan	benswan@aol.com
Walt Varner	71754.135@compuserve.com
Bill Weber	billw@mediacity.com
Larry Weinberg	lewenber@aol.com
Seth Wu	seth.wu@vkn.varian.com

What Does @ Mean to You?

Kendric C. Smith
Condensed from Information Week, 3/96

The @ sign is commonly used in UNIX programming, and Internet e-mail addresses. In English we refer to it simply as the "at" sign, but in other cultures, @ has a number of different names. In Italian it is known as "chiocciolina", and in French, "petit escargot", both meaning little snail. In Germany, @ is referred to as "klammeraffe", or spider monkey. The Dutch refer to it as "api", a shortened form of "apestaart", meaning monkey's tail. In Finland it is a cat's tail or "miau". In Norway it is called the "kanel-bolle", a spiral-shaped cinnamon cake, and in Israel, it is called a "shtrudel". In Denmark, @ is a "snabel", an "a" with a trunk. The Spanish refer to it as "arroba", a unit of about 25 pounds, for which it is the sign.

Win95 Video Memory Requirements

Microsoft TechNet CD

This article describes the memory requirements for different video resolutions in Windows 95.

Colors	16	256	32K③	64K③	16.7 M③
Resolution					
640x480	256K①	256K	1 MB	1 MB	1 MB
800x600	256K	512K	1 MB	1 MB	1.5 MB
1024x768	512K	1 MB	1.5 MB	1.5 MB	2.5 MB
1280x1024	1 MB②	1.5 MB	2.5 MB	2.5 MB	4 MB

①Most VGA cards are shipped with 256K of RAM installed.

② 640K actually needed. 1 MB required due to upgrade requirements on video cards.

③ External RAMDAC (RAM digital-to-analog converter) required for these colors.

More Information

Most current video chip sets have an internal ability to display 256 simultaneous colors. This requires 3 analog-to-digital converters to convert the bit planes in video memory to analog signals for the color guns in the monitor. This function is provided by the RAMDACs.

To use more than the basic number of colors typically requires an external RAMDAC. The "High Sierra" RAMDAC is a relatively small 16-pin DIP (dual in-line package) that supports up to 32,000 colors (or 15-bit color). The newer 24-bit (16.7 million colors) RAMDACs are sometimes larger than the actual video chip set.

Speed and resolution have a different meaning for RAMDACs and the actual video chip set. The speed of the chip set determines what type of performance is delivered, and resolution is the number of pixels it can process. For RAMDACs, resolution refers to the number of colors, not the number of pixels on the screen. The RAMDAC's resolution is its ability to differentiate between subtle shades of color.

It is difficult to produce a high-speed and high-resolution RAMDAC. These chips are not only large, but produce a significant amount of heat and add considerably to the cost of the video board. Many systems simply include empty sockets for optional RAMDACs. Without these chips, however, the boards are simply 256-color boards.

The amount of RAM required for each resolution-and-color combination can be determined using the following formula:

((horizontal * vertical) / 8) * number of bits for color

For example,
((640 * 480) /8) * 24 (16.7 million colors) = 921,600 bytes or 1 MB

The following table lists the number of colors converted to bits:

Number of colors	Bits
16	4
256	8
32K	15
64K	16
16.7 M	24

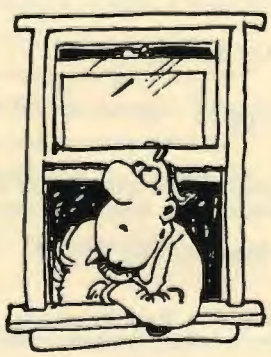

Stanford Palo Alto Users Group for PC
POB 3738
Stanford, CA. 94309-3738
http://www.mediacity.com/~spaug/

Address Correction Requested

Inside News

General Meeting

Wednesday, May 8th
7:30 PM
MediaCity
526 Bryant Street
Palo Alto
Parking under
City Hall
on Ramona St.

The SPAUG Home Page is sponsored and provided by:
MediaCity Inc. (415) 321-6800