C H A P T E R 1 5

Multimedia Services

For the past year, the home market has been the fastest-growing segment of the PC business, and multimedia titles have been one of the fastest-growing segments of the software industry. A large and increasing portion of the PCs being sold into homes are coming with the equipment that makes cool multimedia applications possible − notably CD-ROM drives, sound subsystems, horsepower, and local-bus video.

In 1993, the installed base of multimedia-capable Windows PCs grew rapidly to become the largest multimedia computing platform in the world.

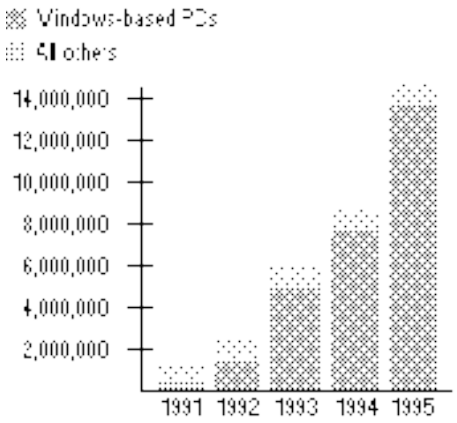

Figure 1. Estimated and forecast sales of multimedia-capable PCs.
Source: Dataquest

By Christmas of 1993, there were more multimedia titles available for Windows than there were for any other computing platform.

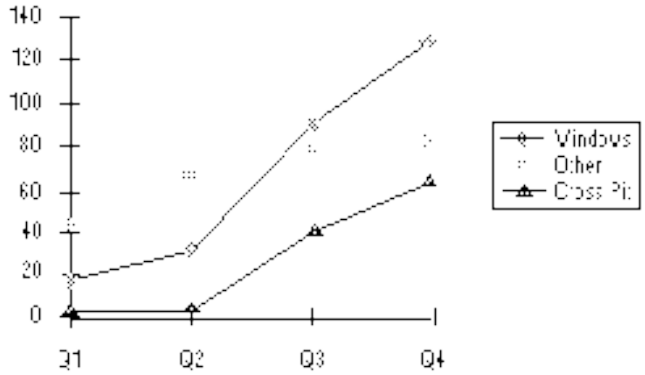

Figure 2. Number of multimedia titles sold by computer software retailers in 1993, by platform, by quarter, as reported by PC Data.

Microsoft Windows and Multimedia

Microsoft is committed to making Windows the leading force in multimedia technologies and systems for personal computers. Our commitment takes many forms, but the most important one is our ongoing investment in multimedia-related research and development. Some of the end results of the last few years of research and development are described in this chapter. This is far from the end, though. Multimedia technologies are evolving rapidly, and we will continue to press ahead in providing tools and architectural enhancements to enable developers and consumers to take advantage of new innovations.

A Little History

It is worth dwelling for a moment on how far Windows multimedia has come in the last few years. When Video for Windows 1.0 was released in 1992, sound cards and CD-ROM drives were relatively rare. Graphics subsystems were universally ISA-based, and software codec technology was in its infancy. The standard size for a digital video clip in this timeframe was 160 pixels by 120 pixels − one-sixteenth of a VGA resolution screen. Technologists (who understood how difficult this was to accomplish) cheered wildly and proclaimed the dawn of the multimedia computing era. Customers shrugged. What's so great about a video clip the size of a dancing postage stamp?

In 1993, hardware and software makers began to deliver equipment and technology that offered better-than-postage stamp performance at reasonable consumer prices. Double-speed CD-ROM drives and local bus video offered more bandwidth to support the massive data requirements of digital video and quality sound. A second generation of software codecs made more effective use of the data available. Prices on 16-bit sound cards dropped into consumer range. With Microsoft Video for Windows 1.1, the size of a digital video clip that a mainstream computer could display reliably increased to 320 x 240 − one-quarter the size of the screen. On one hand, critics have a point when they label this sort of digital video "dancing credit cards." On the other hand, digital video of this size has proven compelling enough to consumers that it has spurred a virtual tidal wave of multimedia title development. Retail software store shelves are crowded with multimedia titles and games.

Progress marches on. Installing Windows 95 will provide today's multimedia PCs with an overnight upgrade in multimedia capabilities. Based on the capabilities of high-end PCs in 1994, the mainstream PC of 1995 will be able play digital video segments that are larger, smoother, and better-looking than ever before − even up to 640 x 480 (full screen) and beyond. We are now able to look forward quite realistically to a time when the amount of data that can be stored on a CD-ROM (rather than the speed of the video subsystem) is the most relevant factor limiting the richness of a consumer's experience with a multimedia title or game.

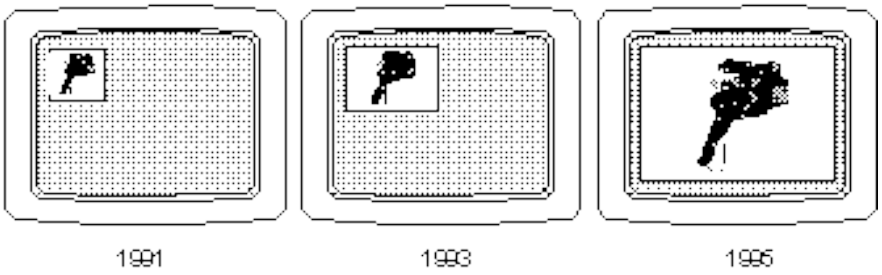

Figure 3. The dancing postage stamp era of multimedia computing is over.

Windows 95—A New High-Performance Multimedia Platform

Windows 95 delivers a new high-performance platform for PC multimedia. From a Big Picture perspective, here's the "greatest hits" of what Windows 95 contributes to the world of multimedia computing:

For consumers, Windows 95 makes multimedia easier, more fun, and more engaging.

- **Easier.** Plug and play will make it far easier for consumers to install multimedia devices successfully. All of the architectural support for digital video, audio and MIDI is built into Windows 95, so that users are liberated from setup challenges. And Windows 95 is compatible with multimedia titles and tools created for Windows 3.1.
- **More Fun.** Windows 95 is a much better platform for computer games than any version of Windows has ever been, including support for fast, intensely graphical games.
- **More Engaging.** Installing Windows 95 is an immediate multimedia upgrade that allows any PC to become a better, more exciting multimedia playback machine. Authors creating titles and games for Windows 95 will be able to make their products faster and more exciting to play.

For developers, Windows 95 offers a powerful platform for professional multimedia authoring

- **Power.** The new 32-bit architecture in Windows 95 squeezes vastly improved multimedia performance out of PCs, so developers can capture digital video and sound that is bigger and bolder than ever before. The multitasking architecture of Windows 95 makes it a much more convenient working environment for multimedia authors.
- **Professional Quality.** The streamlined architecture of digital video, digital audio, MIDI and file handling subsystems in Windows 95 enable authors and toolmakers to create very high-quality, sound, video, and animation effects. Windows 95 is a very attractive platform for professional development of multimedia effects and footage beyond the realm of the PC – TV commercials, for example.

For hardware makers, Windows 95 offers exciting new opportunities

- **Graphics.** A display driver technology called Display Control Interface (DCI) offers ways for Windows to take advantage of hardware assistance for several graphical operations such as image stretching.
- **Sound.** A new technology called Polymessage MIDI offers sound card makers a way to play very, very complex MIDI sequences with virtually no CPU use. Sound cards are improving rapidly, and there is a great deal of room for competition on a feature basis.

Making Multimedia Easier

Plug and Play Support

As multimedia applications, titles, tools, and games have become more and more compelling, consumers have begun buying add-on multimedia components (such as CD-ROM drives and sound cards). Buying these devices has been cheap and easy; installing them has been a different matter. To put it mildly, installing a CD-ROM in a PC has heretofore required... patience.

Support of Plug-and-Play in Windows 95 will make the prospect of adding a new multimedia device to a PC considerably less daunting. Just plug in a Plug and Play enabled sound card and (literally) it plays. In fact, Windows 95 even makes the prospect of installing *old* multimedia devices less daunting − Windows 95 includes tools that make it vastly easier to identify and resolve conflicts between so-called "legacy" devices that are not plug-and-play enabled. Windows 95 includes built-in drivers for the most popular sound cards to make this process as painless as we can possibly make it.

It is difficult to overstate the importance of plug and play for multimedia. Plug and play will do three things for the multimedia market:

- It will allow the base of multimedia capable PCs to grow through plug-and-play upgrade kits, rather than placing so much of the growth burden on the purchase of new CPUs. Because Windows 95 includes the basic architecture for handling sound, MIDI, and digital video, every PC running Windows 95 can easily be made into a multimedia PC − just plug in a sound card and/or CD-ROM drive.
- It will substantially diminish the cost of installing and supporting multimedia devices, which will (among other things) help speed their adoption for business use.
- As multimedia standards (such as CD-ROM speed) continue to improve, plug and play will allow consumers to upgrade multimedia components conveniently without replacing their entire PC. Plug and play support will be vital for adoption of new multimedia devices like MPEG cards.

AutoPlay: Spin and Grin

In various ways, titles and games that run off a CD-ROM feel a bit different than other applications. For one, the way to start a CD-ROM program differs from hard disk-based applications—you first have to open your drawer, extract the right disk, and place it in the CD-ROM drive. Then you can run it like any other program. Assuming, of course, that you can find the icon you created when you first installed the program—a second difference between CD-ROM products and hard disk-based applications is that CD-ROM products may be used less regularly.

In watching users run multimedia applications, we realized that the act of placing a disk in a CD-ROM drive is loaded with information. If the CD-ROM is a title, and you've never run it before, then the act of putting it in the drive means that you intend to install the program. If you already have installed the title, then the act of inserting the disk means that you intend to run it.

Simple enough. In Windows 95, we have implemented a feature called AutoPlay that allows software developers to make their products easier for customers to install and run. When you put a disk into a CD-ROM drive, Windows 95 automatically spins it and looks for a file called AUTORUN.INF. If this file exists, then Windows 95 opens it and follows the instructions.

This new feature will make the setup instructions for a Windows 95–based multimedia game or title almost absurdly easy:

> To play _____, insert the disk in your CD-ROM drive.
>
> Have a nice day!

Built-in Support for Digital Video

For the past several years, Microsoft has been developing a high-performance architecture for digital video − Microsoft Video for Windows. (For more details, see the “Multimedia Graphics Architecture” section later in this chapter.)

In the past, Video for Windows was distributed separately (principally as a Software Developers’ Kit). With the release of Windows 95, Microsoft Video for Windows is built right into the operating system. For the first time, the ability to play digital video will be built into every copy of Microsoft Windows (including Windows NT). This has several implications:

- Users and independent software vendors can use the .AVI file format to distribute digital video files with the same confidence that they today distribute files of other Windows-supported formats like .TXT, .WRI, .BMP, .PCX, and .WAV.
- The barriers to entry for would-be multimedia title and tool developers will be further lowered because the issues of licensing and installing Microsoft Video for Windows will disappear.

Built-in Support for Sound and MIDI

MIDI is the computer equivalent of sheet music. Using sheet music, you can describe how to play Beethoven’s Moonlight Sonata in a few pages − but in order to actually play the song you need to find a piano and a person who knows how to read sheet music. When you hear the music performed from the sheet music, you can expect some variation in sound depending on the circumstances − for example, if you use an expensive grand piano, the sonata will sound better than it would if you used an old upright.

Similarly, a MIDI file can contain the electronic instructions for playing Moonlight Sonata in just a few kilobytes − but in order to play the song you must have a device (such as a sound card) that knows how to "read" MIDI instructions and that can produce a piano sound. Just as the sound of pianos varies somewhat in the real world, so does the sound of a piano on a sound cards.

At the high end, MIDI is used as a development tool for musicians. Virtually all advanced music equipment today supports MIDI, and MIDI offers a convenient way to control the equipment very precisely.

At the low end, MIDI is becoming an ever more popular tool for multimedia product developers because it offers a way to add music to titles and games with a tiny investment of disk space and data rate. The majority of sound cards today have on-board MIDI support built in.

Windows 95 includes built-in support for both MIDI and waveform audio (.WAV).

CD Player: Whistle while you work

Many people like to play audio CDs in their CD-ROM drives while working. So we created a CD Player to go into Windows 95. The controls on this player look just like a regular CD player, and it supports many of the same features you find in advanced CD players (such as random play, programmable playback order and the ability to save programs) so that you don't have to re-create your playlist each time you pop in a CD.

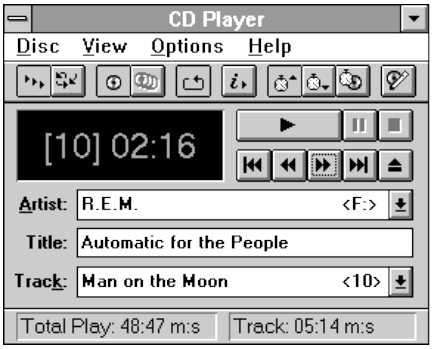

Figure 4. You already know how to use the CD Player in Windows 95. Your CD will play, uninterrupted, in the background as you work.

Making Multimedia More Engaging

CD+ Support Built Into Windows 95

In addition to making it easy for users of Windows 95 to play their favorite audio CDs from their current collection, Windows 95 is helping to define a standard for music CDs of the future. Windows 95 is the first operating system to announce support for the new Sony/Phillips "CD+" format. CD+ is a new CD format that will enable audio CD players and multimedia PCs to easily play the same compact discs. This new format allows both audio and data to be integrated on the same CD, in a manner conducive to both users of audio CDs and users of PC–based CD-ROM titles.

The CD+ format uses new technology, called stamped multisession, that solves the "track one" problem that has prevented easy use of CD-ROMs in audio CD players. Until now, CD-ROM titles have used the first track of a compact disc for data, thus producing static—and potential speaker damage—when played on audio CD players. Sony and Phillips are implementing stamp multisession under the brand name "CD+." Other music industry companies can license the CD+ brand from them, or create their own implementations of stamp multisession. Microsoft Windows 95 will accommodate all compatible implementations of the technology.

Because data and audio information can be combined on the same CD title, the new CD+ format will open up a broad, new category of CD titles that can be enjoyed fully as audio discs, and that can provide digital information in the form of music videos, song lyrics, biographies and other text, and even promote online exchanges with musicians, when inserted into a PC running Windows 95.

The new format leverages a range of new features being included exclusively in Windows 95 to help make multimedia more engaging. The AutoPlay feature in Windows 95, for example, will enable users to insert a compact disc in their CD-ROM drive and have it automatically play. Also, the 32-bit multimedia subsystems in Windows 95 enable unprecedented playback performance. The new CD file system further facilitates multimedia use, while the Plug-and-Play support in Windows 95 will make it simple for consumers to install and use CD-ROM drives and related hardware.

Bigger, Faster, Better-looking 32-bit Digital Video Playback

Displaying digital video involves moving and processing huge streams of data continuously and efficiently. The new digital video implementation in Windows 95 offers some exciting new efficiencies that will allow software developers confidently to create multimedia titles that are more compelling and good-looking than ever before.

Multimedia title and game developers are business people − when they create a product, they do so with the hope of turning a profit. To maximize the number of PCs that can run a title, most developers tend to include lowest-common-denominator digital video. To ensure that as many PCs as possible can play their title or game, developers have tended to use "postage-stamp" sized video windows with low frame rates (which make movement look "jerky") and extreme compression (which makes the video look "blocky").

Windows 95 will raise the lowest common denominator significantly.

In the past, the process of displaying digital video has relied on a series of 16-bit systems − from reading data from the disk, to decompressing the video data, to displaying it on screen. One key design goal of Windows 95 was to transition this architecture to 32 bits, and the difference is eye-popping. For multimedia users, installing Windows 95 will be the quickest and cheapest multimedia upgrade available. Without adding any hardware, Windows 95 enables customers to display bigger, smoother, more colorful digital video than ever before.

It's also important to note that multimedia in Windows 95 is fully compatible with 16-bit multimedia titles. Early testing has shown that the 32-bit improvements in file access speed and stream handling results in performance improvements even for 16-bit multimedia applications − the biggest improvements, of course, will be realized in the new generation of fully 32-bit titles that will be designed for Windows 95.

For customers who upgrade their PC to Windows 95, one easy-to-overlook source of performance improvements is the display driver. Many display drivers are updated more or less continuously, whether to fix problems, enhance performance, or to incorporate new features such as DCI. Most customers, however, don't update drivers on their system unless they are having a problem. Upgrading to Windows 95 will ensure that they have the latest and greatest.

Multitasking and Threads: "We *Don't* Interrupt This Program..."

Multimedia applications don't take well to interruption. − When you are watching a video clip or listening to a sound file, you really don't want it to stop in the middle.

The multitasking in Windows 95 is quite different from prior versions of Windows because it is preemptive. In Windows 95, multiple 32-bit processes can share the CPU at the same time, whether those processes have been initiated by different applications ("multitasking") or by one application ("threading").

This has a very important implication for how multimedia titles will feel to consumers. Threading allows multimedia titles and games to have a more smooth, finished feeling to them. For example, a game might have a thread that plays background music continuously during game play. This would help smooth out the breaks between scenes, when the game is loading new data on another thread of the program.

There is at least one other benign externality for multimedia in the move to 32-bitness. As applications, tools and codecs are gradually rewritten to 32 bits, video and other multimedia processes will become less and less likely to be interrupted by other applications. A simple example of this is that in Windows 95 you can move a video window while it is playing without interrupting it.

Built-in Support for Fast CD-ROMs

The trend toward faster CD-ROM drives (double- and triple-speed) is a very good thing for multimedia computing. To get the best possible performance from these new devices, Windows 95 includes a new 32-bit CD-ROM file system (CDFS) for reading files from CD-ROM drives as quickly and efficiently as possible. The Windows 3.1 system for reading files from CD-ROM drives (MSCDEX.DLL) will be included in Windows 95 for last-resort compatibility with products that rely on it.

Faster reading of CD-ROM data helps to make video and audio playback from CD-ROM drives look and sound better. This is an important component of the overall performance enhancements to multimedia in Windows 95.

Windows 95 also extends its support for CD-ROM to drives that read XA-encoded disks, such as Kodak PhotoCD and Video CDs.

MPEG-Hardware Support for TV-like Video from Your CD-ROM

MPEG is a very complex codec (compression/decompression system) for squeezing digital video and stereo audio into an incredibly small data stream. For example, most feature movies can fit on two CD-ROMs with MPEG compression.

Because MPEG is so complex, displaying video from an MPEG file is a very calculation-intensive process − so calculation-intensive, in fact, that the most appealing way to display MPEG video on today's PCs is by using hardware assistance.

Together with the Open PC MPEG Consortium, Microsoft has defined an industry standard for MPEG board and chip makers that want to ship MPEG devices for Windows 95. This standard will allow applications to incorporate MPEG video without worrying about precisely which vendor's MPEG device is present to decompress it.

Making Windows More Fun

Fast DIB Drawing

For the past year, the home market has been the fastest-growing segment of the PC business. More and more of our customers are telling us that they want games for Windows − and at this point, there aren't many. Games are already the largest category of multimedia application, but most of today's computer games are running on MS-DOS.

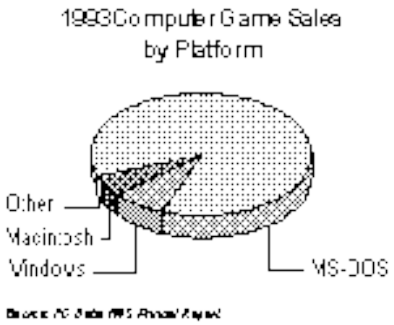

Figure 5. At the end of 1993, computer games were one of the last remaining software categories for which Windows product sales trailed MS-DOS product sales.

The speed of graphics (or, more appropriately, the lack of it) in Windows has been one of the most important obstacles keeping game developers from choosing the Windows platform for their games. We have addressed this issue head-on in Windows 95 in a way that provides substantially improved speed while preserving the device independence that makes Windows appealing in the first place.

A new 32-bit call was added to the Win32 API for Windows 95 and Windows NT, called CreateDIBSection. This new feature allows developers to get bitmaps onto the screen as quickly as possible − if there is nothing fancy (such as clipping or stretching), the CreateDIBSection call will actually allow applications to send DIBs more or less directly to the video frame buffer. (For more information, see the diagram in the "Multimedia Graphics Architecture" section later in this chapter.)

Because we recognize that this kind of graphic speed is critically important to quality games, we have moved a portion of the CreateDIBSection improvements of Windows 95 into a tool for Windows 3.1, called the WinG libraries. WinG (pronounced "Win Gee" − the "G" stands for games) libraries allow game

developers to create fast, graphical games for Windows today with the assurance that your game will be fast and compatible with Windows 95.

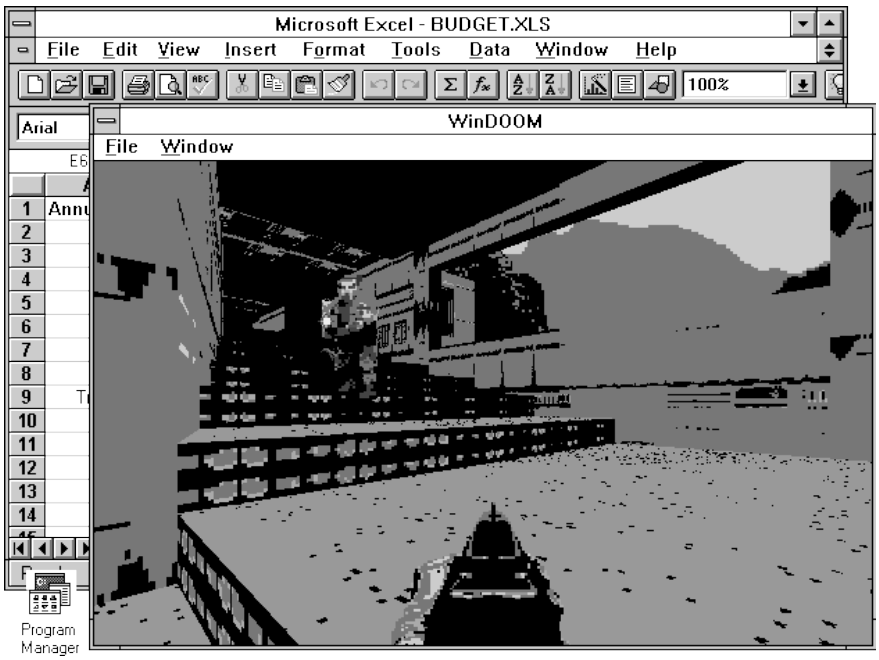

Figure 6. The graphics core of DOOM for Windows was ported from MS-DOS to the WinG library in 2 days. Id Software is scheduled to ship a full version of the product towards the end of 1994.

Built-in Joystick Support

Not much to say here, really. Windows 95 has built-in joystick support so you don't need to load external drivers. *Have fun!*

Powerful Development Environment

Sound Compression for CD-quality Sound

Sound can take up a lot of disk space. Full CD-quality, uncompressed stereo audio contains a lot of data—about 176K for every second of sound! An entire CD-ROM can contain only a little over an hour of music. It can also eat up a fair-sized chunk of the data rate that a CD-ROM drive is capable of sustaining.

To lessen the burden of storing and playing sound from an application, Windows 95 includes a family of sound compression technologies ("codecs"). These codecs can be divided into two groups:

- Music-oriented codecs (such as IMADPCM) are included that allow close to CD-quality sound to be compressed to about one-quarter size.
- Voice-oriented codecs (such as TrueSpeech™) are included to allow very, very efficient compression of voice data.

This support for compressed sound is two-way − you can play sound from a compressed sound file, or you can compress a sound file (using the built-in sound recording and editing utility). If you have a microphone, you can turn on voice compression when recording so that your file is compressed in real time.

In addition to the codecs that come with Windows 95, the audio architecture of Windows multimedia is designed to be extendible through other installable codecs. The video architecture of Windows multimedia can be extended in the same way.

Polymessage MIDI Support for Better Sound

Windows 95 comes with Microsoft's best-ever implementation of MIDI, including a new technology called polymessage MIDI support. This enhancement allows Microsoft Windows to communicate multiple MIDI instructions simultaneously within a single interrupt. The result of this change is that playing MIDI files now requires even less computing power than it did before, and allows developers to process MIDI instructions alongside graphics and other data even more successfully.

Multitasking

Multitasking makes Windows 95 a much more attractive platform for multimedia authoring. Creating multimedia content is very CPU-intensive work that can take a long time to complete. For example, compressing a digital video file can take hours, depending on the complexity of the file and the specs of the system doing the compression. Currently, authors who want to compress more than one digital video file have to do them one at a time − when one file finishes, they can start the next. The result was that video authors were virtually chained to their desks until late at night.

With Windows 95, authors still have control of their PC, even when an enormous compression operation is underway. This makes it possible for digital video authors to initiate several compression operations at once − and then head home.

Professional Quality

Capture and Compression of Bigger Digital Video

When it comes to capturing digital video, there's no avoiding the grim reality that video contains an enormous amount of data. Capturing digital video is even more data-intensive than playing it back, because raw digital video footage is uncompressed. A single frame of full-color video at 640x480 contains close to a megabyte of data. At 30 frames per second, you can fill up a 1 gigabyte hard drive with uncompressed video data in less than a minute. There are ways to compress this data in order to make your storage go further, but no matter what, the rate at which you can write data to disk is of great importance.

The 32-bit file access of Windows 95 is therefore every bit as important to digital video authors as it is to digital video users. Because you can write more data to disk more quickly in Windows 95, you can capture better-looking video—bigger, more frames per second, more colorful.

Once the raw footage is captured, the next step is to compress it − a potentially time-consuming process. Both Cinepak and Indeo™ will be available in 32-bit versions for Windows 95, and this should make the process considerably more efficient.

General MIDI: You Want a Trumpet, You Get a Trumpet!

One of the early challenges for MIDI was that it was, in a way, too flexible. Any instrument can be "connected" to any MIDI channel, so that a "sequence" (song) written for a piano might accidentally end up being played on a tuba.

The general MIDI specification is an industry standard way for MIDI authors to request particular instruments and sounds. Microsoft supports this standard.

Built-in Support for Multimedia Devices

Windows 95 includes built-in support for common multimedia authoring devices like laser disks and VCRs. This makes it easy to set up a system for "step capture," a process in which the author captures digital video data one frame at a time, usually to be compressed later. This is a slow process, but it is absolutely the best way to get the best possible quality digital video.

Frame-accurate control of the VCR is also important for recording broadcast-quality special effects to use in commercials, movies, television programs, music Multimedia PCs for 1995 videos, and the like.

Opportunities for IHVs and OEMs: Multimedia PCs for 1995

All things being equal, installing Windows 95 will upgrade any PC into a more capable multimedia tool.

Of course, all things are *not* equal. There is a great deal of variation in the quality and capability of multimedia PCs and devices. The *Microsoft PC 95 Hardware Design Guide* is being published under separate cover to help IHVs and OEMs identify opportunities to take advantage of new capabilities in Windows 95.

We are making five high-level recommendations to OEMs:

- **Balance beats horsepower**. Multimedia playback places heavy demands on many parts of the system, from the CD-ROM (reading data) to the hard disk (writing data) to the CPU (decompressing) to the video and audio subsystems (playing it). A fast CPU does not guarantee a great playback system. In fact, multimedia playback on most of today's high-end PCs is not constrained by the CPU.
- **Local bus video is indispensable**. Even OEMs creating "non-multimedia" systems should use local bus video, because doing so will enable consumers to plug-and-play their way to a multimedia system later, should they choose to do so. Without local bus video, a PC will not be able to keep up with the amount of video data that 1995's consumer multimedia titles and games will want to display continuously.
- **Double-speed CD-ROM or better**. Titles in 1995 will be written assuming double-speed data rates.
- **SVGA (800x600) or better with 16-bit color**. Why more colors than 256? Because multimedia applications use a lot of colors, and tend to compete for access to the system palette. Consider the challenge of a multimedia presentation that includes a digital video clip of an underwater scene on a slide with a smooth-shaded maroon background. There aren't enough colors in a 256-color palette to make both the slide background and the underwater scene look good.
- **16-bit audio**. The installed base of sound cards that can interpret MIDI is now large enough to be tempting to game and title developers. Not all sound systems are equal − some sound great (16-bit with sampled sounds), and some sound like Star Trek reruns. The differences are significant, and customers will be able to tell the difference.

New Opportunities for Great-Sounding Audio

There is a great deal of variation in the quality of audio cards and sound systems. Most of the time, sound cards up to now have been used principally for their ability to play waveform audio − the equivalent of recorded sound. For some uses, like voice-overs, there is no realistic alternative to recorded waveforms. However, recorded sound is very resource-intensive for both the CD-ROM and the CPU. In Windows 95, there are enhancements to the handling of MIDI that makes it an even more appealing alternative to .WAV for playing music within games and multimedia titles. There are several things that makers of audio cards and systems can do to distinguish themselves:

- **Polymessage MIDI support.** This is a very efficient new technology included in Windows 95 that makes it easier for application and game writers to use MIDI. If a sound card supports polymessage MIDI, the CPU use required to play even a very complex song is quite small.
- **16-voice or better polyphony.** Polyphony is the ability to play multiple sounds at once. Support for more concurrent sounds means fuller-sounding playback.
- **Sampled sound rather than wave-form synthesis.** Wave-form synthesis uses a mathematical approximation of a sound, such as a piano. Sampled sound is an actual recording of the piano, and sounds considerably better. Including samples of at least the most common General MIDI instruments helps ensure that music in games and titles sounds really good, instead of synthetic.

DCI: Taking Advantage of New Video Card Features

In the summer of 1994, Microsoft released the new DCI display driver development kit. This technology was developed in partnership with Intel and other makers of advanced video display cards.

DCI is a device driver level interface that allows Windows to take advantage of hardware features that are (or could be) built into advanced display adapters, specifically:

- Stretching which speeds up rendering of images that are stretched or distorted.
- Color-space conversion which assists in playback of compressed digital video by accepting YUV data instead of requiring RGB.
- Double Buffering which allows faster, smoother block transfers (BLTs) of images by providing memory space for off-screen drawing.
- Chroma key which facilitates the merging of video data streams, allowing a particular color to be treated as "transparent" in the merge operation.
- Overlay which speeds display of partly concealed objects.
- Asynchronous drawing which, along with double buffering, provides a faster method for "drawing" into off-screen memory space.

Most of the hardware features above relate to the fast, efficient decompression and playback of digital video. Applications that use the Microsoft Video for Windows architecture will benefit from these features automatically and substantially.

Multimedia Architecture

Multimedia Graphics Architecture

There are four kinds of graphics an application might want to "draw" on the screen, and four APIs that an application can use to do so:

- **"Productivity application" graphics**. Scroll bars, fonts, buttons, and the like. Applications that want the system to help them draw these things use GDI, the basics Windows graphics API.
- **Digital video**. Applications that want to play digital video use the Video for Windows API. More details on the Video for Windows architecture are provided in the following section.
- **Game graphics**. Games draw their own graphics (in memory) and want bitmaps blasted to the screen as fast as possible. That's what WinG does. It is available for Windows 3.1, and provides many of the same benefits of the CreateDIBSection function in Windows 95, as well as fast access to the frame buffer through DCI.
- **3D engineering graphics**. Applications that want the system to help them draw 3D solids use OpenGL. OpenGL is Microsoft's strategic choice of 3D application programming interface. We have a long- term commitment to deliver an implementation of OpenGL as part of the broader Win32 API, a

commitment we announced last November. Our first OpenGL implementation shipped in Windows NT 3.5.

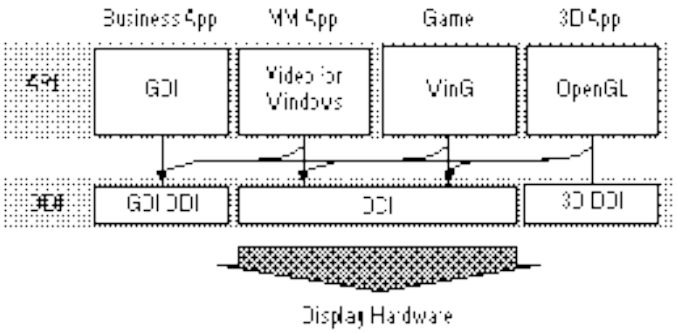

Figure 7. Windows Graphics Architecture

There are three pieces to the device driver interface in Windows, and the APIs described above are designed to take advantage of whatever DDI provides the best performance.

- GDI-DDI is the basic graphics device driver interface for Windows. It is optimized for the flexible graphics requirements described above for the GDI API.
- DCI is the new device driver interface created jointly by Microsoft and Intel. DCI drivers provide a fast, direct way for games and digital video in windows to write to the video frame buffer. It also enables digital video playback to take advantage of several specific kinds of hardware support included on advanced graphics adapters. For example, stretching hardware can allow users to scale up the size of a digital video clip with virtually no additional strain on the CPU. Color space conversion support in hardware can reduce the amount of work a codec must perform by up to 30%, allowing substantially better video playback.
- The 3D-DDI enables applications that use OpenGL to take advantage of accelerated 3D support in hardware.

How Multimedia Data is Routed in Windows

The following diagram describes (in simplified form) the path that synchronized multimedia data travels from storage to experience during playback.

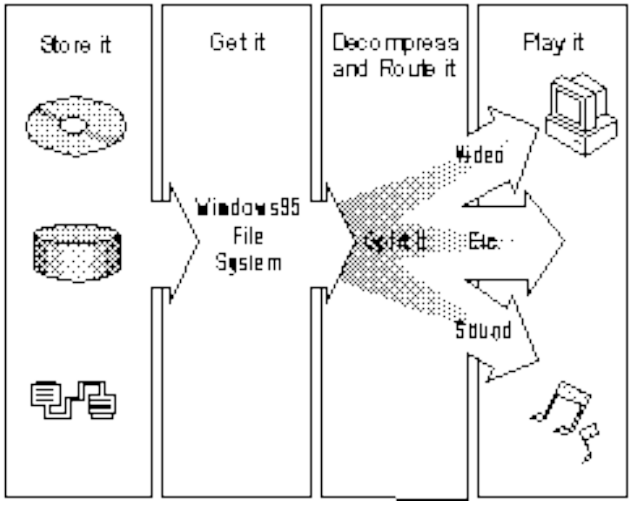

Figure 8. Windows Multimedia

First, the data (usually an .AVI file) must be stored somewhere, such as a CD-ROM, a local hard drive, a network file server, or another storage medium. The quality of the eventual playback will be constrained by the amount of data that the storage medium can supply to the file system continuously.

A command (such as Play), usually issued through the Media Control Interface (MCI), causes the relevant part of the file system in Windows 95 to retrieve the stored data. Obtaining this data swiftly and steadily is vital to the success of overall playback performance, and the 32-bit protected-mode enhancements in the new file system (and CDFS) in Windows 95 have a lot to do with the overall performance enhancements of multimedia in Windows 95.

A multimedia data stream (such as an .AVI file) generally contains multiple components, such as digital video data, audio data, text, and perhaps other data (such as hot spot information, additional audio tracks, and so forth.) As multimedia information comes off the CD-ROM, the first job of the Video for Windows architecture is to figure out what the data stream contains, and to separate and route it accordingly.

In most cases, digital video and digital audio are stored in a compressed form. Before it can be seen or heard, therefore, it must be decompressed. Frequently, this function is performed in software. If hardware support is available on the graphics adapter or sound card for all or part of the decompression work, however, Video for Windows can tap into it.

The Video for Windows architecture has been created in a way that allows *installable* codecs. Windows 95 ships with a set of useful software-only codecs for both video and audio, but you are not limited to these tools only. As new codecs become available for particular audio and digital video needs, they can be plugged into the Video for Windows architecture. For example, motion JPEG is a useful codec for multimedia authoring − capture cards that support JPEG compression and decompression are easily available, even though JPEG itself is not explicitly provided in Windows 95.