MPEG Choices for PCs Abound

Soon to Be a PC Standard—Vendors Vie with Different Solutions

by James L. Turley

MPEG-1 recently took a major leap toward becoming the de facto standard for delivering compressed video on personal computers. Compaq's decision to include MPEG-1 decompression hardware as a standard feature on future models *(see **0909MSB.PDF**)*, coupled with Microsoft's move to include MPEG-1 decompression software with Windows 95, has pushed MPEG-1 past the point of critical mass and into the mainstream of the PC business.

The accelerating adoption of MPEG-1 over competing compression algorithms will ignite a boom in the development of software that includes color video and audio. Games especially will take advantage of these features, but multimedia reference titles will follow, with increasing reliance on the format to deliver extra content to business users.

These developments are bad news for proponents of competing audio/video compression formats like Cinepak, TrueMotion, and Indeo. With no clear industry leader, software vendors are divided among these codecs. But with new impetus from Compaq and Microsoft, multimedia developers will migrate toward MPEG-1, adopting it more earnestly as the installed base grows.

Several solutions are available for MPEG-1 playback on a standard PC. While software decompression and playback is now possible, hardware decoders will be a growth industry in the coming years.

Let There Be Video

Over the past few years, several methods have emerged to provide full-motion video on standard PCs. Uncompressed, even a still 256-color image at 640 × 480 resolution requires 300K; animating such images at 30 frames/s (the accepted minimum for smooth full-motion video) is beyond the capabilities of most PCs. Even one minute's worth of this uncompressed video would fill a 600M CD-ROM. Thus, compression and decompression are mandatory before video can be stored or played back on inexpensive computers.

In the absence of any standard decompression or acceleration hardware, only the x86 instruction set and a VGA-resolution monitor are givens. In recent years, several companies strove to develop real-time decompression algorithms that balanced playback quality with reasonable storage requirements, resulting in a number of lossy, low-resolution software codecs such as Cinepak, TrueMotion, and Intel's Indeo.

These delivered small, often jerky color video with unsynchronized sound or no sound at all. The MPEG-1 standard *(see **060803.PDF**)* was touted as a better- quality alternative, compressing 352 × 240 images at 30 fps into a 176K/second data stream, including stereo audio. But its more complex algorithm was too compute-intensive to deliver a reasonable frame rate on current 386 or 486DX processors. Thus, Indeo and other codecs that delivered poorer quality but ran on available CPUs gained popularity. But as CPU performance steadily increases—from 486DX2 to DX4 and Pentium—it becomes feasible to perform MPEG-1 decompression in something approaching real time.

Another drawback of Cinepak, TrueMotion, and most forms of Indeo is that they cannot encode video in real time. This makes them less attractive to software producers than MPEG, because the source video must be stored in uncompressed form in its entirety and compressed offline at a later time. MPEG, in contrast, can be compressed on the fly in real time. Although MPEG-1 encoders are priced upwards of $1,000, this is a small price compared with the cost of massive amounts of disk storage required for uncompressed video.

Software Provides One Solution

Two companies, Xing Technology and Mediamatics, have developed software engines for decompressing and playing MPEG-1 video and audio on standard PCs. Both products run under Windows 3.x, avoiding the 64K driver limit of MS-DOS. Microsoft licensed a version of Mediamatics' MPEG Arcade Player to include in Windows 95. In an effort to improve graphics performance under Windows 95, Microsoft replaced much of the DCI (device control interface) graphics library in Windows 3.x with DirectDraw. This change requires different binaries for applications that run under both versions of Windows. It has also been a significant setback to Xing, among others, which relied heavily on DCI functions for its software MPEG player.

The Mediamatics decoder can play MPEG-1 video at approximately 24 fps on a system with a 90-MHz Pentium and a medium-priced graphics accelerator card (e.g., Diamond Stealth64). Like the Xing player, the Mediamatics software makes extensive use of the display adapter's accelerator chip to perform color-space conversion and scaling. These two functions comprise a significant portion of the MPEG decompression routine and are particularly time-consuming and inefficient on the x86 architecture. Although the software from both companies will run with an unaccelerated VGA adapter, their performance is much worse.

Mediamatics estimates that its software running on a 120-MHz Pentium will provide full 30-fps playback. Even with a 120-MHz Pentium, however, sound quality is still sub-par. An MPEG data stream includes stereo 44.1-KHz, CD-quality sound. The Mediamatics player, however, performs only a cursory demodulation of the sound stream, delivering 11-KHz stereo sound comparable to an AM radio. In taking advantage of a faster processor, Mediamatics chose to boost the frame rate rather than improve sound quality. Either way, 100% of the Pentium is devoted to processing MPEG-1 data.

This is the biggest problem with software decompression. Although acceptable MPEG-1 playback is now available essentially for free (at least on systems with a 90-MHz Pentium and Windows 95), it leaves no headroom for any additional processing. One could argue that additional processing is not needed; that video playback, by its very nature, commands the user's full attention. Playing a video in the background behind a spreadsheet, for instance, makes little sense. Apart from those people who want to keep a miniature television running in the corner of their screen, any additional CPU cycles would be wasted anyway.

While this may be true for "serious" MPEG applications in education, business, or reference works, it is certainly not the case for entertainment titles and games. Nor is it applicable to multitasking operating systems that allow file transfers, fax transmissions, or print spooling in the background. These all require significant amounts of additional processing at the same time the video and audio are playing. To avoid overwhelming the processor, a hardware assist of some kind is required. Adding a hardware MPEG-1 decoder frees the processor for other tasks and improves playback quality.

Hardware Decompression Spares CPU

In addition to its soon-to-be-popular software decoder, Mediamatics also develops hardware. With a business model similar to that of ARM or MIPS Technologies, the company has developed decompression cores for both MPEG-1 and MPEG-2, which it then licenses to silicon vendors. VLSI Technology has licensed both cores for use in its set-top-box decoder chips. Western Digital and a number of other hardware vendors have also begun development around the MPEG-1 core, with product announcements expected later in the year.

As Table 1 shows, current MPEG-1 decoder chips are similar in many ways. They all include an internal processor of some kind and several additional function units, usually a special multiply/divide unit for performing IDCT (inverse discrete cosine transform) operations. A DRAM controller for the minimal 512K buffer is also standard.

Currently, every MPEG-1 decompression chip for PCs relies on an independent graphics accelerator chip for scaling, color-space conversion (from YUV to RGB), frame-buffer management, and an interface to the video DACs. Such features are now common in mid-range and high-end graphics chips and are often found on even basic PCs. Figure 1 illustrates a typical MPEG-1 multimedia board with separate graphics, video, and digital/analog converters.

Where some of these chips differ is in the way they receive their compressed data. The S3, Winbond, and C-Cube chips, for example, all have their own ISA or PCI bus interface. This allows the chips to be added as aftermarket multimedia accelerator boards, but it necessitates irregular or nonstandard cable connections between the MPEG logic and the existing graphics accelerator. Alternatively, these chips can be integrated onto one board with the accelerator chip, sharing the latter's bus connection.

The recently announced chip set from S3 *(see* ***0909MSB.PDF****)* offers two roads to MPEG integration. The Scenic/MX1 includes a PCI interface that it uses to communicate with the company's general-purpose accelerator, thus allowing it to be used on an aftermarket PCI add-in board. The nearly identical Scenic/MX2 version

Decompression Chip	C-Cube CL450	C-Cube CL480PC	SGS-T STi3430	Winbond W9910IF	Winbond W9920IF	Winbond W9920CF	S3 Scenic/MX1	LSI L64002	Hyundai HDM8211
MPEG-1/2	MPEG-1	MPEG-1	MPEG-1	MPEG-1	MPEG-1	MPEG-1	MPEG-1	MPEG-2	MPEG-2
Bus interface	16-bit	ISA	16-bit	ISA	ISA	PCI	PCI	32-bit	16-bit
Audio?	No	Yes	Yes	Yes	No	No	Yes	Yes	Yes
Video?	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
DRAM required	512K	512K	512K	None	512K	512K	512K	512K–2M	1M–2M
DRAM data bus	16-bit	16-bit	16-bit	N/A	16-bit	16-bit	16-bit	64-bit	64-bit
Frequency	40 MHz	40 MHz	30 MHz	36 MHz	36 MHz	36 MHz	28 MHz	27 MHz	32 MHz
Package	PQFP-128	PQFP-128	PQFP-120	PQFP-80	PQFP-160	PQFP-100	PQFP-128	PQFP-160	PQFP-208
Voltage	5 V	3.3 V	3.3 V	5 V	5 V	5 V	3.3 V	3.3 V	3.3 V
Mfg process	0.8 µ	0.65 µ	0.5 µ	0.8 µ	0.6 µ	0.6 µ	0.65 µ	0.5 µ	0.65 µ
Price (10K)	$35	$45	$41	$9	$23	$20	$35	$36	$78

Table 1. MPEG-1 and MPEG-2 decoder chips are available from a number of sources, with various feature sets. Older chips separate audio and video decode functions, while newer designs integrate them. (Source: vendors)

drops the PCI interface in favor of a dedicated 8-bit bus between itself and the accelerator chip. This approach keeps video traffic off the host PCI bus. The Scenic/MX2 could also be offered as an end-user upgrade, possibly as a daughtercard on an S3-based graphics board.

Although the interchip bus is proprietary, the S3 chips can configure the interface for compatibility with existing digitizers from Philips, C-Cube, IIT, and others that output 4:2:2 YUV data.

The MPEG-1 standard specifies an output resolution of 352 × 240, but in PCs this results in a small picture, which is often scaled to a larger window. The method used to "upsample," or replicate pixels to achieve the larger size affects the perceived quality of expanded video windows. Some chips interpolate pixels, using a smoothing algorithm similar to those commonly found in laser printers. Others simply replicate the pixels, resulting in blocky, albeit larger, video images.

Audio, Video or Both?

Another feature separating MPEG chips from one another is their treatment of the audio data stream. At the low end, MPEG chips either do not decode the audio stream at all or, like Winbond's W9920IF, require a separate audio chip. For the lowest-cost implementation, omitting audio support has price advantages. But with both audio and video playback now available purely through software, the market window for these very low end devices is rapidly closing.

Many chips that include audio functions still rely on extra support, using device drivers running on the host CPU to demultiplex the audio from the video stream. This task is not extraordinarily complex, and it allows the system to handle nonstandard or encrypted MPEG data streams. S3's Scenic chips fall into this category. However, synchronizing the audio and video when they are handled by separate chips is a nontrivial task.

Standard Microprocessors Need Not Apply

Interestingly, although most vendors supplying MPEG decoder chips hold licenses for embedded microprocessor cores, only one company has chosen to use it for its MPEG products. Hyundai, which acquired a SPARC database through TI in 1994, is the only vendor to use a standard microprocessor instruction set.

LSI Logic, for example, opted for a custom logic design rather than apply a MIPS core to the problem. Winbond bypassed PA-RISC, and IBM, which has an MPEG-2 decoder, did not include PowerPC. C-Cube designed its own core, while IIT, which has x86 experience, chose to buy the MIPS-X core developed at Stanford (see MPR 10/30/91, p. 1). Cirrus Logic, while it has not announced any MPEG chips to date, is unlikely to exercise its ARM license, relying instead on the CompCore Multimedia design it licensed in 1994 *(see* ***0811MSB.PDF****)*. The C-Cube and Hyundai chips are the only ones to actually execute external code, the others relying on internal microcode ROMs.

The widespread decision to develop special MPEG cores highlights many important aspects of this market. First, price is a critical factor. Cost requirements for mass-market peripherals can't support CPU licensing fees. Second, none of the current 32-bit microprocessor architectures is well-suited to the task of performing discrete cosine transforms or other routines necessary for video decompression (witness the need for a 120-MHz Pentium to keep up with a 40-MHz CL480). Finally, software compatibility is not an issue for a chip that will execute only a carefully optimized device driver for one specific task. On all counts, mainstream microprocessors come up short.

Furthermore, video-output tasks gain very little benefit from the conventional data cache designs found on general-purpose microprocessors. Unlike computer applications, there is little or no locality of reference for video data; once it's processed, it's gone.

At this time, there are no chips that combine MPEG decoding and conventional graphics acceleration. All the MPEG-1 decoders are separate devices and require either their own connection to the system bus or a private interface into the video subsystem. For the time being, end-user MPEG upgrades are often awkward and require add-in cards or ribbon cables between boards.

Today's graphics-chip vendors argue that the dual-chip approach is not likely to change within the next year, for reasons both technical and commercial. Although the adoption of MPEG-1 will be widespread, and the demand for MPEG-1 accelerators will increase, single-chip multimedia devices are not imminent, they say. An MPEG-1 decompression engine requires approximately 400,000 transistors; graphics accelerators are not much smaller, and few functions can be shared between

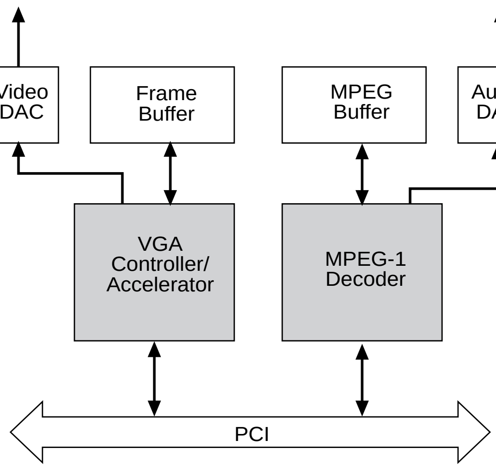

Figure 1. A typical MPEG-1 decompression solution for PCs adds a decoder to an existing GUI accelerator.

the two. The MPEG-1 decoder market is too uncertain for chip vendors to gamble on what is for them a very large die. By 1997, however, sub-half-micron technology should allow these functions to be economically merged in a single die, sharing a PCI bus interface and a single DRAM buffer.

Whither MPEG-2?

With current processor performance, MPEG-1 decompression can be performed either by software or in hardware. However, a software decoder for MPEG-2 (which delivers 4× resolution over MPEG-1) may be many years away. Even with the P6, Mediamatics claims that acceptable MPEG-2 performance is not possible without a significant hardware assist.

MPEG-1 will offer sufficient resolution and sound quality for PC users for several years to come. It has to—MPEG-2's storage requirements are better suited to broadcast media, where storage is not an issue. MPEG-1 was developed to address both the capacity and bandwidth limitations of original CD-ROM technology. A modern 4× CD-ROM drive can read data quickly enough to play back MPEG-2 streams, but the discs still do not have the capacity to make it worthwhile. MPEG-2 on a typical 4× CD-ROM is sweet and short.

Several major consumer-electronics giants are striving to change this situation. Both the Sony/Philips alliance and the Toshiba/Time Warner sodality are seeking to establish their respective high-density CD-ROM formats as industry standards. So far, both factions are accumulating licensees at a brisk pace, leading to a potential industry split reminiscent of the VHS/Beta battle of a generation ago. This would obviously be destructive to the industry as a whole, so hopefully some agreement on a common standard will be reached soon.

If one (or both) of these "quad-density" CD formats becomes popular, it will then be feasible to deliver MPEG-2 content on a static storage medium suitable for PCs. Fortunately, MPEG-2 is a superset of MPEG-1, so earlier titles should still play on an upgraded machine. Until that time, MPEG-1 will be the de facto standard for PCs, while MPEG-2 is relegated to the higher end of the market for broadcast encoding. The proposed MPC-3 (multimedia PC) standard, in fact, specifies MPEG-1 at 30 fps as a minimum requirement.

Adoption and Outlook

The MPEG-1 standard is sufficient for compression and playback on PCs, where video quality is not crucial and storage space is limited. Moreover, video playback on a PC is currently more of a novelty than a requirement, and most long-time computer users are thrilled that their PC can play small, grainy images with AM radio–quality sound at all. Except for a few games (which always seem to push the limits of PC hardware) and multimedia encyclopediae, none of today's software is near the 74-minute video limit of a standard CD-ROM. For as long as mainstream PC applications do not stretch the capabilities of the standard, MPEG-1 will continue to satisfy the majority of users and OEMs.

Including MPEG-1 decompression software (along with Cinepak and Indeo decoders) with Windows 95 will give millions of casual PC users their first taste of MPEG video, much as QuickTime did for Apple customers. Those who run Windows 95 on their new Compaq Presario or other PC with hardware MPEG support will be happier still. The widespread availability of software MPEG players will spur the purchase of MPEG-1 accelerator boards as more and more PC owners seek to satisfy an appetite whetted by Windows 95. Microsoft's move, in fact, has overjoyed the makers of 1996's "games accelerator" cards. Users who are unsatisfied with the performance of their software decoder will likely take a big step, paying perhaps $150 for a hardware accelerator that delivers noticeably better performance.

As an interesting side note, Apple will start including MPEG boards in low-end Macintoshes while its high-end systems rely on software decompression—because MPEG chips are cheaper than a faster PowerPC, which can perform MPEG-1 decoding at 30 fps with full audio.

The demand for cheap MPEG chips at the low end will come from motherboard makers and OEMs that need to offer MPEG as a checklist item. Like PCs with accelerated graphics cards today, these will become competitive necessities in 1996, but they do not have to offer competitive performance in the customer's system.

That leaves a great middle territory with few potential buyers. Many of today's chips fall into this category. Separate audio and video decoders are too expensive to include on a motherboard and not fast enough to satisfy performance-hungry upgrade customers. A battle is brewing for the high ground in PC MPEG-1 chips, and several vendors are girded for combat. ♦

For More Information

For additional information about the MPEG decoder chips and software mentioned in this article, contact the vendors listed below.

C-Cube Microsystems (Milpitas, Calif.) 408.944.6400; fax 408.944.6314. Cirrus Logic (Fremont, Calif.) 510.226.8300; fax 510.252.6020. Hyundai (San Jose, Calif.) 408.473.9200; fax 408.473.9800. LSI Logic (Milpitas, Calif.) 408.433.8000; fax 408.433.8989. Mediamatics (Santa Clara, Calif.) 408.496.6360; fax 408.496.6634. S3 (Santa Clara, Calif.) 408.980.5400; fax 408.980.5444. SGS-Thomson (San Jose, Calif.) 408.452.8585; fax 408.452.1549. Winbond (San Jose, Calif.) 408.943.6666; fax 408.943.6668. Xing Technology (Arroyo Grande, Calif.) 805.473.0145; fax 805.473.0147.